Geometric Facility Location under Continuous Motion

Bounded-Velocity Approximations to the Mobile Euclidean k-Centre and k-Median Problems

by

Stéphane Durocher

B.Sc., University of Toronto, 1997
M.Sc., University of British Columbia, 1999

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Computer Science)

The University of British Columbia

April 2006

Abstract

The traditional problems of facility location are defined statically; a set (or multiset) of n points is given as input, corresponding to the positions of clients, and a solution is returned consisting of set of k points, corresponding to the positions of facilities, that optimizes some objective function of the input set. In the k-centre problem, the objective is to select k points for locating facilities such that the maximum distance from any client to its nearest facility is minimized. In the k-median problem, the objective is to select k points for locating facilities such that the average distance from each client to its nearest facility is minimized. A common setting for these problems is to model clients and facilities as points in Euclidean space and to measure distances between these by the Euclidean distance metric.

In this thesis, we examine these problems in the mobile setting. A problem instance consists of a set of mobile clients, each following a continuous trajectory through Euclidean space under bounded velocity. The positions of the mobile Euclidean k-centre and k-median are defined as functions of the instantaneous positions of the clients. Since mobile facilities located at the exact Euclidean k-centre or k-median involve either unbounded velocity or discontinuous motion, we explore approximations to these. The goal is to define a set of functions, corresponding to positions for the set of mobile facilities, that provide a good approximation to the Euclidean k-centre or k-median while maintaining motion that is continuous and whose magnitude of velocity has a low fixed upper bound. Thus, the fitness of a mobile facility is determined not only by the quality of its optimization of the objective function but also by the maximum velocity and continuity of its motion. These additional constraints lead to a trade-off between velocity and approximation factor, requiring new approximation strategies quite different from previous static approximations.

We identify existing functions and introduce new functions that provide bounded-velocity approximations of the mobile Euclidean 1-centre, 2-centre, and 1-median. We show that no bounded-velocity approximation of the Euclidean 3-centre or the Euclidean 2-median is possible. Finally, we present kinetic algorithms for maintaining these various functions using both exact and approximate solutions.

Stéphane Durocher
University of British Columbia

January 2006

Contents

List of Tables

List of Figures

Preface

Keywords
facility location, geometric facility location, mobile facility location, continuous motion, bounded velocity, mobile points, mobile clients, mobile facilities, approximation, eccentricity, stability, Euclidean distance, 1-centre, 2-centre, k-centre, 1-median, Weber point, k-median, Steiner centre, Steiner point, m-hull, projection median, reflection-based 2-centre, centre of mass, rectilinear 1-centre, rectilinear 1-median

Thesis Title
Alternative titles considered for this thesis include the brief but perhaps too general, "Mobile Facility Location", and that which was selected as the subtitle, "Bounded-Velocity Approximations to the Mobile Euclidean k-Centre and k-Median Problems", with the possibility of appending the suffix "with Applications to Mobile Facility Location". The final title was chosen for its appeal to a more general audience, as the term "geometric facility location" is somewhat less technical and potentially more widely recognized. As well, the mention of "continuous motion" signals that we are not dealing with discretized time samples, but rather with motion defined over a continuous time interval.

Previously Featured
Some of the work presented in this thesis has appeared in conference proceedings and journal publications. Some results on the Steiner centre from Ch. 4 appear in [DK03, DK04, DK06]. Some results on the projection median from Ch. 5 appear in [DK05b, DK05c]. Some results on reflection-based 2-centre approximations from Ch. 6 appear in [DK05a].

In a Rush?
If you possess only limited time with which to make your way through these pages, an overview of the motivation, main ideas, and contributions can be extracted from the obvious places such as Ch. 1, the first section of every chapter, and the last section of Chs. 4 through 6. Secs. 2.3 and 2.4 are recommended for readers unfamiliar with the Euclidean k-centre and k-median problems. Ch. 3 is recommended for readers new to problems involving continuous motion or approximation. Of particular significance to all readers is Sec. 3.6 which provides a perspective on the relevance of results developed in subsequent chapters.

Mount Columbia, May 2005

Acknowledgements

First and foremost I must acknowledge my thesis supervisor, Dr. David Kirkpatrick. Throughout my time as a graduate student David has provided me with unending guidance in research, many years of financial support, and a genuine friendship. David has dedicated hundreds of hours of his time collaborating with me on various lingering unresolved questions relating to our research, often prefacing what would turn out to be a brilliant solution with the question "What about trying this...?" David has endured numerous patient revisions of our journal and conference submissions amid my protests that no more revisions were necessary. He was correct every time and the revisions resulted in substantially improved final versions of each paper. David is truly a model faculty member who excels as a researcher and is highly regarded in his field, an outstanding lecturer who is well-respected by his students, and someone who gives generously of his time as a department member and dedicates himself to improving every aspect of the academic environment. David is humble, always acknowledging others and often understating his own contributions. David is selfless in his dedication to his graduate students, generous with his time for them, always respectful of their ideas, tactful with his suggestions, and enthusiastically committed to their success. For years to come I will undoubtedly find myself asking "What would David have done in this situation?" Finally, it is worth noting that after several years of trying, I have yet to arrive at the Sun Run finish line ahead of David. This year's race is not looking any better; as I nurse a stress fracture back to health, David appears to be in great shape as result of a dedicated daily commute by bicycle and his long weekend runs. Those who know David realize he is not surpassed easily, whether it be running 10 km, finding ingenious new solutions to difficult research problems, or by his endless supply of smiles with which he greets colleagues day after day as he walks by in the hallway.

My thesis committee members certainly deserve recognition. Dr. Will Evans, Dr. Tamara Munzner, and Dr. Alan Wagner kindly agreed to sit on my committee and give of their time and wisdom. Each of them helped improve and contribute to various areas of this thesis related to their respective specializations.

I should acknowledge the MITACS project on Facility Location Optimization headed by Dr. Binay Bhattacharya. Our semi-monthly meetings with Dr. Boaz Ben-Moshe and Qiaosheng Shi along with David, Binay, and myself provided stimulating problem-solving sessions, including a discussion of the one-dimensional geometric 2-centre which led to the algorithm described in Sec. 6.2.2.

Also, to the members of the Beta lab, the theory group, and the lunch crew who meet every day at noon in the lounge, thank you for making every day enjoyable; your company will be missed.

To my family, my parents Yves and Christiane, my sister Evelyne, and my brother Daniel, thank you for your kindness and support throughout the years. You will undoubtedly be relieved to see this thesis put to rest. Thank you especially to Daniel, who has joined me on the west coast for the past few years. Having a family member living in the same city has been truly enjoyable and our time spent together savouring a burger and a beer or scrambling along a ridge toward the summit will be remembered fondly.

Finally, my lovely wife Andrea Bunt deserves a special acknowledgement for her encouragement, her patience, her kindness, her laughter, her beautiful smile, and her love throughout the entire duration of this thesis. Andrea and I met in my first year of the Ph.D. program. As I end this final year and write the last pages of this thesis I look forward to a happy future together.

Andrea, to you I dedicate this thesis.

Chapter 1

Introduction

This chapter presents a brief introduction and motivation for the mobile problems of geometric facility location, followed by a short description of the main contributions developed in this thesis and an overview of the thesis' organization into chapters. Chapter 1 serves as an extended abstract for the following chapters.

1.1 Geometric Facility Location and Mobility

1.1.1 Continuous Motion and the Need to Approximate

The traditional problems of facility location are set in a static setting; client positions are fixed and a single location is selected for each facility. A set of n points is given as input, corresponding to the positions of clients, and a solution consisting of set of k points that optimizes some objective function of the input set is returned, corresponding to the positions of facilities. These static problems of facility location, in particular the k-centre and k-median problems, have been studied extensively (see [FMW83, HM03, HLP$^+$87, HM89, LMW88] for general overviews of static facility location). In the k-centre problem, the objective is to select k points for locating facilities such that the maximum distance from any client to its nearest facility is minimized. In the k-median problem, the objective is to select k points for locating facilities such that the average distance from each client to its nearest facility is minimized. A common setting for these problems is to model clients and facilities as points in Euclidean space and to measure distances between these by a Minkowski distance metric (typically either Euclidean distance, rectilinear distance, or Chebyshev distance). This class of facility location problems is referred to as *geometric facility location*.

Motivated by recent advances in mobile computing and telecommunications, these questions have been posed in the mobile setting (e.g., [AH01, AGHV01, AGG02, AdBG$^+$05, BBKS06, GGH$^+$03, Her05]), presenting new constraints and challenges specific to mobility. A problem instance consists of a set of mobile clients, each following a continuous trajectory through Euclidean space under bounded velocity. Since mobile facilities located at the exact Euclidean k-centre or k-median involve either unbounded velocity or discontinuous motion, we explore approximations to these. The goal is to define a set of functions, corresponding to positions for the set of facilities, that provide a good approximation to the k-centre or k-median while maintaining motion that is continuous and whose magnitude of velocity has a low fixed upper bound. Thus, the fitness of a mobile facility is determined not only by the quality of its optimization of the objective function but also by the maximum velocity and continuity of its motion. These additional constraints lead to a trade-off between velocity and approximation factor, requiring new approximation strategies quite different from previous static approximations.

Although numerous papers examine problems related to mobile, kinetic, or dynamic facility location, most are only indirectly related to this work since they involve variations such as discontinuous motion, discrete time steps, or distance metrics other than Euclidean distance. A notable exception is the work of Bereg et al. [BBKS00, BBKS02, BBKS06] which examines some early problems of mobile facility location and introduces the mobile Euclidean 1-centre and 1-median problems discussed in Chs. 4 and 5; in fact, their paper provided the catalyst that eventually led the author to select this direction for his thesis.

Why are these problems relevant? Any actual implementation involving real mobile clients and mobile facilities in the physical world requires both continuity in position and some upper bound on velocity. Scenarios involving vehicles, mobile robots, or people with wireless communication devices suggest that bounds on continuity and velocity are necessary in many applications [AOY99, KNW02, CFPS03, Sch03, CMKB04, CMB06] (See Sec. 3.7.6 for a discussion of applications of mobile facility location). Therefore, in many cases a bounded-velocity approximation of the k-centre and k-median is necessary. Finally, it should be noted that the worst-case examples resulting in unbounded velocity or discontinuity of the k-centre or k-median are easily realized by a small number of clients, for example, with as few as four clients moving at unit velocity along random linear trajectories inside the unit square on the plane.

1.1.2 Evaluating Bounded-Velocity Approximations

Given that approximation is necessary, we are motivated to ask how closely the exact positions can be approximated. The two measures by which approximation schemes are compared are *maximum velocity*, an upper bound on the maximum relative velocity of a mobile facility, and *approximation factor*, a bound on the worst-case ratio of the values of the optimization function given by the approximation function relative to the optimal value.

Maximum velocity is straightforward to define relative to the motion of the mobile clients. Since we are interested in relative velocity we may assume that the velocity of every client is at most one.

Each facility location problem has an associated optimization function g. For the k-centre, g is the maximum distance from any client to the nearest facility, whereas for the k-median, g is the sum of the distances from each client to the nearest facility. The usual measure for evaluating the approximation factor of an approximation scheme Υ relative to the optimal scheme Ξ is to bound the worst-case ratio of their optimization functions over all sets of mobile clients P. Thus, we say Υ is a λ-approximation of Ξ if

$$\forall P,\ \ f(P, \Upsilon) \leq \lambda f(P, \Xi). \tag{1.1}$$

For some problems, however, no bounded-velocity approximation is possible. That is, for any approximation scheme Υ, no fixed values of $v_{\max}$ and λ satisfy both the bounded-velocity requirement and Eq. (1.1). For these cases, a counter-example can be constructed to demonstrate that given any λ and $v_{\max}$, no approximation moving with maximum velocity at most $v_{\max}$ can guarantee an approximation factor of λ. Evaluation of potential solutions to these problems requires a different type of analysis, one which falls outside the definition of approximation considered in this thesis.

Depending on the number of facilities, k, and the dimension, d, of the problem space $\mathbb{R}^d$, either a) a mobile k-centre or k-median problem has a low upper bound on its velocity and no approximation is necessary, b) the problem has unbounded velocity or discontinuous motion but a bounded-velocity approximation is possible, or c) the problem has unbounded velocity or discontinuous motion and no bounded-velocity approximation is possible. We consider problems that corresponds to the second set of conditions, when the motion of the mobile k-centre or k-median is discontinuous or has unbounded velocity but a bounded-velocity approximation exists. Values of k and d that result in these conditions are:

1. Euclidean 1-centre in $\mathbb{R}^d$ for $d \geq 2$,

2. Euclidean 2-centre in $\mathbb{R}^d$ for $d \geq 2$, and

3. Euclidean 1-median in $\mathbb{R}^d$ for $d \geq 2$.

For each of these, although the motion of the exact centre(s) or median has unbounded velocity or is discontinuous, a bounded-velocity approximation is still possible. We develop, analyze, and compare several strategies for approximating the mobile Euclidean 1-centre, 2-centre, and 1-median problems.

1.2 Contributions and Thesis Overview

This section highlights the major contributions of my thesis research. It also provides a summary of the organization of this material into chapters.

1.2.1 Geometric Facility Location (Ch. 2)

Ch. 2 presents the definitions of two fundamental problems of geometric facility location, the Euclidean k-centre and the Euclidean k-median, as these are commonly defined in the facility location literature. Work related to static instances of these problems is discussed along with related problems including the rectilinear k-centre, the rectilinear k-median, the centre of mass, and k-means clustering. A reader familiar with these definitions may consider omitting this chapter as no new material is introduced.

1.2.2 Mobile Facility Location (Ch. 3)

This research explores new ground which initially requires establishing answers to fundamental questions such as: what properties are significant in a mobile setting, why might the exact k-centre or k-median be inadequate under motion, and how should potential alternatives be compared against each other? In answering these questions, Ch. 3 introduces concepts necessary for discussing problems of mobile facility location, including formal definitions for maximum velocity, continuity of motion, approximation factor, and the related notion of stability. In addition, work related to problems in mobile facility location is discussed. This chapter establishes the important questions regarding mobile problems in geometric facility location, questions which are addressed in the remaining chapters. Ch. 3 and, in particular, Sec. 3.6 are essential to understanding the full significance and context of the results established in this thesis and their relevance in relation to the fields of computational geometry and facility location.

1.2.3 Mobile Euclidean 1-Centre (Ch. 4)

Ch. 4 explores bounded-velocity approximations of the mobile Euclidean 1-centre. Motivated to define a weighted mean of the extreme points of the client set whose weights change continuously under motion, our search led (somewhat indirectly) to the definition of the Steiner centre. Interestingly, the Steiner centre can also be defined by projection of the client positions onto a line through the origin, and integrating over all such projections. Previous applications of the Steiner centre have been mostly limited to its definition on convex polytopes as applied to topological problems in differential geometry involving surface curvature. To the author's knowledge, the Steiner centre had never been examined as approximation to the Euclidean 1-centre nor had it been applied to the setting of mobile facility location. The proof of the approximation factor (eccentricity) of the Steiner centre established in Ch. 4 suggests the Steiner centre as a bounded-velocity approximation to the mobile Euclidean 1-centre, providing a good compromise between maximum velocity and approximation factor, which compares well against other natural approximation functions of the Euclidean 1-centre including the rectilinear 1-centre and the centre of mass.

1.2.4 Mobile Euclidean 1-Median (Ch. 5)

Ch. 5 explores bounded-velocity approximations of the mobile Euclidean 1-median. The Euclidean 1-median problem is perhaps more complicated to approximate than the Euclidean 1-centre, since not only does the Euclidean 1-median move discontinuously, but, in general, its exact position cannot be calculated for five or more clients [Baj88]. Not knowing the position of the Euclidean 1-median increases the difficulty of measuring approximation factors since the exact value of the optimization function cannot be known for the majority of client sets. Nevertheless, we introduce the projection median, a new median function which we analyze as a bounded-velocity approximation of the Euclidean 1-median. Indeed, good upper and lower bounds on its approximation factor as well as a tight bound on its maximum relative velocity compare well against other natural approximation functions of the Euclidean 1-median, including the rectilinear 1-median and the centre of mass.

1.2.5 Mobile Euclidean 2-Centre (Ch. 6)

Ch. 6 explores bounded-velocity approximations of the mobile Euclidean 2-centre. Multiple-facility problems differ in several aspects from single-facility problems. Specifically, a static k-centre problem usually involves two steps: clients are partitioned into k clusters and a 1-centre is identified for each cluster. These two steps are not independent; the optimal positioning of k facilities induces an optimal partition of the client set and vice-versa. Identifying a Euclidean k-centre requires coordination between the positions of the facilities. In the mobile setting, explicit clustering is infeasible because client movement between partitions inevitably leads to discontinuities in the motion of the mobile k-centre. Again, our goal remains to maintain an approximation function that returns a set of k mobile facilities, each moving continuously and under bounded velocity.

To overcome problems of discontinuous motion, we introduce reflection-based 2-centre functions, a new set of approximations to the mobile Euclidean 2-centre that involves coordinating the positions of the two facilities without explicit partitioning of the clients. Our analysis capitalizes on our results from Ch. 4. We derive bounds on the approximation factor and on the maximum velocity of these approximations of the Euclidean 2-centres. To the author's knowledge, no previous bounded-velocity approximations to the mobile Euclidean 2-centre have been defined.

1.2.6 Mobile Geometric k-Centres and k-Medians (Ch. 7)

In Ch. 7 we show that even in one dimension, no bounded-velocity approximation is possible for any geometric k-centre when $k \geq 3$ nor for the geometric k-median when $k \geq 2$. For any fixed values of λ and $v_{\max}$, we construct a counter-example to show that either λ or $v_{\max}$ is insufficient. Finally, we ask whether there exist bounded-velocity approximations to these problems when the approximation function has greater than k facilities.

1.2.7 Implementation (Ch. 8)

The development of kinetic data structures (KDS) by Basch et al. [BGH99] within the last decade precipitated research on a series of algorithmic problems in computational geometry involving sets of mobile points (clients) moving continuously through Euclidean space. Implementation issues for our work on mobile facility location fall very naturally within this framework, allowing for several existing algorithms and data structures to be adapted to the various approximation functions, including the mobile Steiner centre, the mobile projection median, and reflection-based mobile 2-centres. In the case of the mobile Steiner centre, achieving an efficient implementation using a KDS required defining a new approximation of the convex hull, a concept of independent interest.

In addition to the theoretical work, a Java applet was coded to provide a visual demonstration of numerous approximation functions of the mobile Euclidean 1-centre, 1-median, and 2-centre. This implementation was used to collect statistics for an empirical average-case evaluation of these approximation functions.

1.2.8 Conclusion and Open Problems (Ch. 9)

Ch. 9 provides a brief summary and concludes with some direction for future research.

Chapter 2

Geometric Facility Location

2.1 Introduction

2.1.1 Chapter Objectives

Chapter 2 motivates and defines the traditional static problems of facility location. Of particular relevance to this thesis are centre and median problems when the domain of input points is Euclidean space and the distance metric is Euclidean distance. The chapter closes with a discussion of related work. The mobile versions of these problems are addressed in Chapter 3.

2.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Defining Geometric Facility Location (Sec. 2.2)
Sec. 2.2 begins by identifying the four attributes of a facility location problem instance: the universe, the distance metric, the number of facilities to be located, and the optimization criterion. Of relevance to this thesis are the problems of geometric facility location, where clients and facilities are modelled as points in Euclidean space, distance between these is measured by a Minkowski distance metric, and facilities are not restricted to coincide with client positions.

Centre Problems (Sec. 2.3)
Sec. 2.3 defines the centre problem, one of two elementary problems of facility location. The first problem introduced is the 1-centre on any metric δ and in any universe. The problem is generalized to the k-centre, where multiple facilities serve the client set. We focus specifically on the Euclidean 1-centre and the Euclidean k-centre, where the metric δ corresponds to Euclidean (ℓ_2) distance. We provide an overview of previous work.

Median Problems (Sec. 2.4)
Sec. 2.4 introduces the median problem, the second of two elementary problems of facility location. The 1-median and k-median are introduced first, followed by a discussion of the Euclidean 1-median and the Euclidean k-median and an overview of previous work.

Related Work (Sec. 2.5)
Sec. 2.5 provides an overview of related static problems in facility location. These include the rectilinear k-centre and k-median, the centre of mass, k-means clustering, and continuous facility location. Related problems in mobile facility location are discussed in Sec. 3.7.

2.2 Defining Geometric Facility Location

2.2.1 Facility Location: Clients and Facilities

Facility location describes an extensive range of problems that have been examined within numerous fields including computational geometry (e.g., [RT90, Wel91, AS94, BE97, AS98, Est01, BMM03]), operations research (e.g., [Hak64, HPRT85, DTW86, MS02, HM03]), data analysis (e.g., [JD88, ME98, Est99, GMMO00, EH01]), computational complexity (e.g., [KH79a, FPT81, MIH81, MS84]), and graph theory (e.g., [HM72, WH73b, DF74, Min77, Sla81, HLN91]). This research area has a long history and continued activity accompanied by a rich literature. The traditional problem of locating a facility to optimize some function of the input set of client positions was first formally defined by Weber [Web22] early in the last century. Problems that would eventually define the fundamental questions of facility location, however, were examined much earlier. These include finding the centre of the circumcircle of a triangle (the circle on which all three triangle vertices lie) examined as early as the 3rd century B.C. by Archimedes and finding the Fermat-Torricelli point of a triangle (the point that minimizes the sum of the distances to the triangle vertices) which was first posed by Fermat [Fer91] and solved by Torricelli early in the 17th century [KV97].

Although countless variations of problems are classified as facility location, every instance of a problem in facility location is characterized by the following input parameters:

- a universe, U, from which a set (or multiset) P of input client positions is selected,
- a distance metric, $\delta : U \times U \to \mathbb{R}^+$, defined over the universe U,
- an integer, $k \geq 1$, denoting the number of facilities to be located, and
- an optimization function, $g : \mathscr{P}(U) \times U^k \to \mathbb{R}$, that takes as input a set (or multiset) of client positions and a set of k facility positions and returns a function of their distances as measured by metric δ, where $\mathscr{P}(A)$ denotes the power set of set A.

The corresponding problem statement consists of selecting locations for a set of facilities to optimize the value of function g;

- select a set (or multiset) F of k facility positions in universe U that minimizes $g(P, F)$.

In a typical scenario, the client set[1] represents demand points. The cost for providing a service to these clients is represented as a function, g, of the distances between their respective positions and the positions of a set of facilities. Regardless of whether the metric and optimization function are measuring distance travelled, bits transferred, length of cable required, amplitude of signal transmitted, dollars spent, or time required to complete service, the underlying problems share similar instance attributes and problem statements.

Given a universe U and an integer k, we refer to any function $\Upsilon : \mathscr{P}(U) \to \mathscr{P}(U)$ that takes a multiset of elements from U as input (clients) and returns a multiset of k elements of U (facilities) as a **facility function**. When $k = 1$, we say Υ is a **single-facility function**. When $k > 1$, we say Υ is a **multiple-facility function**. Given specific values for U, δ, k, and g describing a particular facility location problem, an exact or approximate solution to the problem is given by defining a facility function. Common facility functions have specific names such as the Euclidean k-centre, the Euclidean k-median, and the centre of mass described in Secs. 2.3, 2.4, and 2.5.2.

Many facility location problems are NP-hard when greater than one facility is being located. These include the k-centre and k-median problems in $\mathbb{R}^2$ under Euclidean or rectilinear distance [MS84], both of which remain NP-hard to approximate within some fixed factor, and the k-centre and k-median problems on graphs [KH79a], where in all cases k is assumed to be an arbitrary input parameter (discussed in greater detail in Secs. 2.3 through 2.5). When only one or two facilities are to be located, the majority of these same problems can be solved either exactly or within close approximation of the exact solution in polynomial time.

Reviews of the relevant problems and histories of significant discoveries within the field of facility location are provided in [FMW83, HLP+87, HM89, LMW88] and more recently in [HM03].

2.2.2 Geometric Facility Location: Euclidean Space and Minkowski Distance

Problems are classified by differences in the domain universe and in the choice of distance metric. In addition, the fundamental problems of facility location have spawned countless variations by augmentations incorporating additional constraints such as weights, capacities, costs, or obstacles; some of these related problems are examined in Sec. 2.5.

The universe of allowable positions for clients and facilities for a particular problem is modelled by one of three spaces: continuous space, discrete space, or network space [HLP+87, HN98, Pao99, HM03]. **Continuous space** refers to a universe defined as a region, typically within $\mathbb{R}^d$, such that clients and facilities may be positioned anywhere within the continuum, and the number

[1] Unless otherwise specified, both the set of input client positions and the set of facility positions may include multiplicities. By convention, we use the same set theory terminology, notation, and operations for multisets as we do for sets.

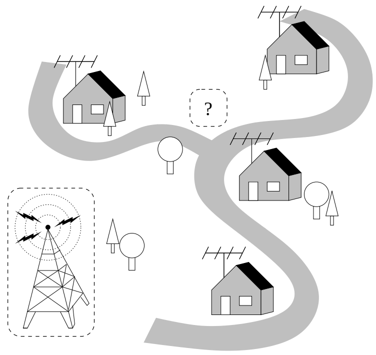

Figure 2.1: Where should the radio tower be located to minimize the maximum distance from any house to the tower?

of possible locations is uncountably infinite. **Discrete space** refers to a universe defined by a finite set of predefined positions. **Network space** refers to a universe defined by an undirected weighted graph, whose edge weights respect the triangle inequality. Possible client positions are given by vertices. Depending on the problem, facilities may be positioned anywhere along an edge, or restricted to graph vertices. Common definitions for the universe include the real numbers in one or more dimensions, $\mathbb{R}^d$; a d-dimensional grid, $\mathbb{Z}^d$; some d-dimensional region, $[a_1, b_1] \times \ldots \times [a_d, b_d]$; a finite set of allowable client positions, $\{u_1, \ldots, u_k\}$; and an undirected graph whose edge weights respect the triangle inequality. A natural domain for describing problems involving continuous motion and the domain used within this work is unbounded d-dimensional Euclidean space, namely, $\mathbb{R}^d$.

The following example illustrates the differences between the three possible domains for the universe. Say the client set corresponds to the positions of houses within some neighbourhood for which a single radio tower (the facility) must be positioned with the objective to optimize the worst-case signal reception (minimize the distance from the tower to the furthest house). See Fig. 2.1. If the tower may be positioned anywhere, then the universe is a continuous space. If the radio tower may only be constructed on one of five available tower-sized lots, then the universe is a discrete space. Finally, if instead the problem is to position a mailbox along the roadside somewhere within the neighbourhood so as to minimize the worst-case bicycling distance from the mailbox to the furthest house, then the universe is a network space (assuming bicycles ride on the road).

For a given universe, the distance metric further differentiates between specific problems. Within continuous-space universes, the most prevalent convention for measuring distance between two points u and v in $\mathbb{R}^d$ is to use the Minkowski distance given by the ℓ_p norm of the vector $u - v$ for some $p \geq 1$. In general, the class of continuous-space facility location problems whose input universe is $\mathbb{R}^d$ and whose distance metric is a Minkowski distance is referred to as **geometric facility location** [HRS04, FMW05] (and sometimes as *planar facility location* when $d = 2$ [FMW83]). As expressed by Fekete et al.,

> With many practical motivations, geometric instances of facility location problems have attracted a major portion of the research to date. In these instances, the sets D of demand locations and F of feasible placements are modelled as points in some geometric space, typically $\mathbb{R}^2$, with distances measured according to the Euclidean (ℓ_2) or rectilinear (ℓ_1) metric. [FMW05, p. 61]

The Minkowski distance between two points x and y in $\mathbb{R}^d$ corresponds to the Minkowski norm $||x - y||_p$. For a given $p \geq 1$, the ℓ_p norm[2] of a point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$ is given by

$$||x||_p = \left(\sum_{i=1}^{d} |x_i|^p \right)^{1/p}. \tag{2.1}$$

Almost exclusively, the three values of p used are $p = 1$, i.e., the rectilinear distance $||u - v||_1$; $p = 2$, i.e., the Euclidean distance $||u - v||_2$; and $p = \infty$, i.e., the Chebyshev distance $||u - v||_\infty$.

Distance and velocity in Euclidean space are commonly measured using Euclidean distance. As discussed in Chs. 3 through 6, many of the interesting challenges of mobile facility location arise under Euclidean distance. Thus, we adopt Euclidean distance as our distance metric, although rectilinear and Chebyshev distance arise occasionally throughout the thesis. By Eq. (2.1), the Euclidean distance between two points $x = (x_1, \ldots, x_d)$ and $y = (y_1, \ldots, y_d)$ in $\mathbb{R}^d$ is given by

$$||x - y|| = \sqrt{\sum_{i=1}^{d} (x_i - y_i)^2}. \tag{2.2}$$

Since it is the norm most frequently referred to in this work, we omit the subscript and write simply $||x||$ to denote the ℓ_2 norm of $x \in \mathbb{R}^d$.

2.2.3 Optimization Function: Sum or Maximum

Perhaps the most significant characteristic that distinguishes problems of facility location is the objective function. A typical instance of a facility location problem can be classified as either a *centre* problem or a *median* problem, corresponding to the two major underlying classes of objective functions involving either a maximum or a sum, respectively. As stated by Hale and Moberg,

> There exists two predominant objective functions in location science: minisum and minimax. These are also known as the median and center problems, respectively. [HM03, p. 22]

[2] The Minkowski norm is defined for any $p > 0$. When $p \in (0, 1)$, however, the resulting norm is not a metric since it does not respect the triangle inequality.

We now introduce definitions for various centre problems (Sec. 2.3) and median problems (Sec. 2.4) specific to geometric facility location.

2.3 Centre Problems

2.3.1 1-Centre and k-Centre

Given a set of points P, a fundamental problem of geometry and data analysis concerns the characterization and computation of points that are central to P.

Definition 2.1. *Given a universe U, a finite set of points $P \in \mathscr{P}(U)$, and a metric $\delta : U \times U \to \mathbb{R}^+$, a* **1-centre** *of P is a point $c \in U$ that minimizes*

$$\max_{p \in P} \delta(p, c). \tag{2.3}$$

Set P must be nonempty for the value of Expr. (2.3) to be defined. Furthermore, multiplicities of points in P do not alter the value of Expr. (2.3). Consequently, it is irrelevant whether P is a set or a multiset. As we will see in Sec. 2.4, the same observation is not true of the 1-median.

A 1-centre problem is also known as a *minimax* problem [Han73, EH72, DF74, HPRT85, MC86b, DTW86, HM03].

Of course, the term *1-centre* derives directly from the more general *k-centre*, where k denotes the number of facilities. We now generalize the single-facility definition of the 1-centre to multiple facilities.

Definition 2.2. *Given a universe U, a finite set of points $P \in \mathscr{P}(U)$, a metric $\delta : U \times U \to \mathbb{R}^+$, and a positive integer k, a* **k-centre** *of P is a set of k points $F \subseteq U$, that minimizes*

$$\max_{p \in P} \min_{c \in F} \delta(p, c). \tag{2.4}$$

That is, C is a set of k points such that the maximum distance, $\delta(p, c)$, from any point $p \in P$ to the nearest facility $c \in F$ is minimized.

When the domain is Euclidean space and δ is a Minkowski distance, we refer to a k-centre as a **geometric k-centre**. If δ is the Euclidean distance metric, then we call it a *Euclidean k-centre*.

2.3.2 Euclidean 1-Centre

A natural, and for many applications the default, metric for measuring distance between two points is the Euclidean distance. The corresponding Euclidean 1-centre of P is the (unique) centre of the smallest enclosing hypersphere of P. See Fig. 2.2.

Definition 2.3. *Given a finite set $P \in \mathscr{P}(\mathbb{R}^d)$, the* **Euclidean 1-centre** *of P is the function whose value, $\Xi_d(P)$, is the point in $\mathbb{R}^d$ that minimizes*

$$\max_{p \in P} ||p - \Xi_d(P)||. \tag{2.5}$$

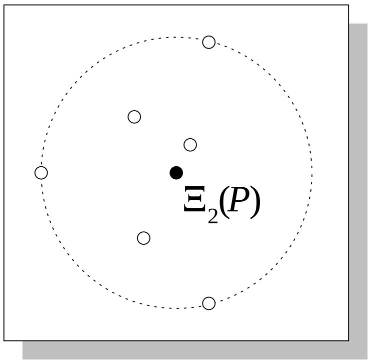

Figure 2.2: The Euclidean 1-centre, $\Xi_2(P)$, of a set of points P in $\mathbb{R}^2$ corresponds to the centre of the smallest enclosing circle of P.

The value $\max_{p \in P} ||p - \Xi_d(P)||$ is referred to as the **Euclidean radius** of P, also known as the *circumradius* of P.

It is straightforward to show that the Euclidean 1-centre is invariant under similarity transformations. Note, this property does not hold for all geometric 1-centres. For example, the rectilinear 1-centre is not invariant under rotation and reflection (see Sec. 2.5.1).

Synonyms for the Euclidean 1-Centre

The Euclidean 1-centre is also known as Euclidean centre [DK06], 1-centre, circumcentre [Hon95, Kim98], centre of the circumcircle, ℓ_2 centre, unweighted Euclidean minimax problem [HLP+87], midpoint (in $\mathbb{R}$), spatial 1-centre [HM03], centre of the smallest enclosing circle/sphere/hypersphere or disc/ball/d-ball (in $\mathbb{R}^2/\mathbb{R}^3/\mathbb{R}^d$) [Wel91], and Kimberling triangle centre $X(3)$ [Kim98] (when $|P| = 3$). In contrast, the *centrepoint* of P sometimes refers to a point c such that for every line l through c, at least $k|P|$ points of P lie on either side of l, where $0 < k < 1$ is some fixed fraction [JM93] (such a point c is also sometimes called a median of P [AdBG+05]).

Algorithms for Finding the Euclidean 1-Centre

Although the question of finding the minimum enclosing circle of a triangle is thought to have been first posed by Archimedes, the general question of finding the minimum enclosing circle for n points in the plane was first posed by Sylvester in 1857 [Syl57]. An early algorithm was provided by Chrystal in 1885 [Chr85]. Since then, the minimum enclosing circle (for points in $\mathbb{R}^2$) and minimum enclosing sphere (for points in $\mathbb{R}^3$) problems have been well studied with both deterministic and randomized linear-time algorithmic solutions. Megiddo [Meg83] gives a deterministic $\Theta(n)$-time linear programming solution in $\mathbb{R}^2$, where $n = |P|$. This result has been extended to $\mathbb{R}^d$ for any fixed d in $O(d^{O(d)}n)$ time by Agarwal et al. [AST93] and by Chazelle and Matoušek

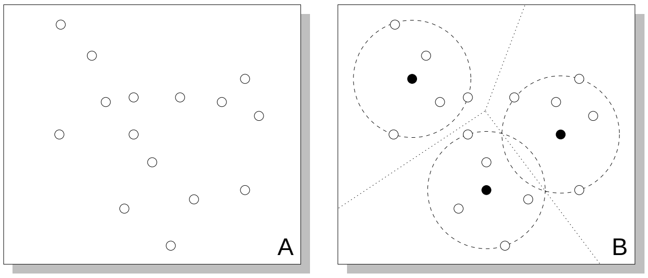

Figure 2.3: a set of points P in $\mathbb{R}^2$, a Euclidean 3-centre of P, the associated three minimum enclosing circles of P, the corresponding Voronoi diagram, and the induced 3-partition

[CM96]. Since every point must be examined, these results are asymptotically optimal when d is fixed. Welzl [Wel91] gives a simpler randomized algorithm with $\Theta(n)$ expected time in $\mathbb{R}^d$ for any fixed d. Xu et al. [XFS03] review solutions to the minimum enclosing circle problem while Nielsen and Nock [NN04] review solutions to the minimum enclosing sphere problem.

Applications of the Euclidean 1-Centre
As stated by Nielsen and Nock, applications for the Euclidean 1-centre problem span a wide array of fields:

> The smallest enclosing ball, as a fundamental primitive, finds many applications in computer graphics (collision detection, visibility culling, ...), machine learning (support vector clustering, similarity search, ...), metrology (roundness measurements, ...), facility locations (base station locations, ...), and so on. [NN04, p. 147]

2.3.3 Euclidean k-Centre

The Euclidean k-centre is the natural generalization of the Euclidean 1-centre to multiple facilities.

Definition 2.4. *Given a finite set $P \in \mathscr{P}(\mathbb{R}^d)$ and a positive integer k, a* **Euclidean k-centre** *of P is a set of k points in $\mathbb{R}^d$, $\{\Xi_d^1(P), \ldots, \Xi_d^k(P)\}$, that minimizes*

$$\max_{p \in P} \min_{1 \leq i \leq k} ||p - \Xi_d^i(P)||. \tag{2.6}$$

The value $\max_{p \in P} \min_{1 \leq i \leq k} ||p - \Xi_d^i(P)||$ is referred to as the **Euclidean k-radius** of P.

Given a set P, the points of a Euclidean k-centre of P correspond to the centres of k hyperspheres whose union encloses the points of P such that the radius of the largest hypersphere is minimized. See Fig. 2.3. Unlike the Euclidean 1-centre, a Euclidean k-centre is not unique, even for $k = 2$. For example, let

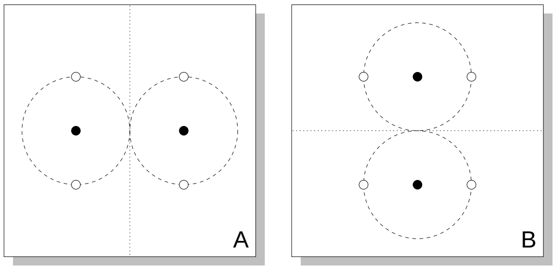

Figure 2.4: non-uniqueness of the Euclidean 2-centre

four points be located at the vertices of the unit square in the plane. Two distinct 2-centres are possible, corresponding to the midpoints of opposite pairs of edges of the square. See Fig. 2.4.

Synonyms for the Euclidean k-Centre

The Euclidean k-centre is also known as Euclidean p-centre [Dre84a, HLC93b], Euclidean m-centre [KLC90], minimax location-allocation problem [DC97], minmax multicentre problem [GJ79, CPP02], minimax radius clustering [BE97], and planar k-centre problem (in $\mathbb{R}^2$) [Cha99, Dre84b].

Algorithms for Finding the Euclidean 2-Centre

A Euclidean 2-centre is straightforward to find in linear time in $\mathbb{R}$ (see Sec. 6.2). Drezner [Dre84b] provides an $O(n^3)$-time algorithm for finding a Euclidean 2-centre in $\mathbb{R}^2$. Eppstein [Epp92] gives algorithms requiring $O(n^2 \log^2 n \log^2 \log n)$ expected time and $O(n^2 \log^4 n)$ worst-case time. Hershberger and Suri [HS91] provide an $O(n^2 \log n)$-time solution to the corresponding decision problem; this result is improved to $O(n^2)$ time by Hershberger [Her93]. Agarwal and Sharir [AS91, AS94] and Katz and Sharir [KS93] give $O(n^2 \log^3 n)$-time solutions for finding the Euclidean 2-centre in $\mathbb{R}^2$. Jaromczyk and Kowaluk [JK94] give an $O(n^2 \log n)$-time algorithm. Sharir [Sha97] reduces the time complexity to $O(n \log^9 n)$. Eppstein [Epp97] gives a simpler randomized algorithm in $O(n \log^2 n)$ expected time. Finally, Chan [Cha99] gives a deterministic algorithm in $O(n \log^2 n \log^2 \log n)$, still in $\mathbb{R}^2$. Agarwal and Sharir [AS98] mention a generalization of Drezner's algorithm from $\mathbb{R}^2$ to $\mathbb{R}^d$ to give an algorithm requiring $O(n^{d+1})$ time.

Algorithms for Finding the Euclidean k-Centre

No efficient algorithm is known for the Euclidean 3-centre in $\mathbb{R}^2$ [Sha97] (with running time comparable to those of algorithms for the Euclidean 2-centre described above). The problem is solved in linear time in $\mathbb{R}$ using the algorithms of Drezner [Dre87] and Hoffmann [Hof05]; these same algorithms also solve the rectilinear 3-centre in linear time in $\mathbb{R}^2$.

	$k = 1$	$k = 2$	k fixed	k arbitrary
$\mathbb{R}$	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$
$\mathbb{R}^2$	$\Theta(n)$	$O(n \log^2 n \log^2 \log n)$	$n^{O(\sqrt{k})}$	NP-hard
$\mathbb{R}^d$	$O(d^{O(d)} n)$	$O(n^{d+1})$	$n^{O(k^{1-1/d})}$	

Table 2.1: time complexities of algorithmic solutions to the Euclidean k-centre

The Euclidean k-centre can be solved in linear time in $\mathbb{R}$ using the algorithm of Frederickson [Fre91] for finding the k-centre in a tree. Frederickson's algorithm does not restrict centres to be located at vertices but also allows them to be located along the interiors of edges [BBK+02, CPP02].

When k is an arbitrary input parameter, Megiddo and Supowit [MS84] show the Euclidean k-centre problem is NP-hard in $\mathbb{R}^2$; they also show that finding an ϵ-approximation remains NP-hard for any $\epsilon < 2/\sqrt{3} \approx 1.1547$. This bound was increased by Feder and Greene [FG88] who show the problem remains NP-hard for any $\epsilon < (1 + \sqrt{7})/2 \approx 1.8229$. Exponential-time solutions exist. The current best algorithm for the Euclidean k-centre in $\mathbb{R}^2$ requires $n^{O(\sqrt{k})}$ time [HLC93b]. Drezner [Dre84a] gives an $O(n^{2k+1})$-time algorithm for solving the Euclidean k-centre problem in $\mathbb{R}^d$. Agarwal and Procopiuc [AP98] provide a algorithm that improves the time to $n^{O(k^{1-1/d})}$. These results are summarized in Tab. 2.1.

Many approximation algorithms exist for the Euclidean k-centre problem. Gonzalez [Gon85] and Hochbaum and Shmoys [HS86] provide 2-approximation algorithms for the Euclidean k-centre in $\mathbb{R}^2$ that requires $O(nk)$ time. This time was reduced to $\Theta(n \log k)$ and generalized to $\mathbb{R}^d$ by Feder and Greene [FG88]. Agarwal and Procopiuc [AP98] give a $(1 + \epsilon)$-approximation algorithm for the Euclidean k-centre problem in $\mathbb{R}^d$ running in time $O(n \log k + k/\epsilon^{dk})$.

2.4 Median Problems

2.4.1 1-Median and k-Median

Whereas the optimization function in a centre problem has as its goal to minimize the *maximum* distance to any point, the goal of the median's optimization function is to minimize the *sum* (or average) of the distances to the points.

Definition 2.5. *Given a universe U, a finite multiset of points $P \in \mathscr{P}(\mathbb{R}^d)$, and a metric $\delta : U \times U \to \mathbb{R}^+$, a* **1-median** *of P is a point $m \in U$ that minimizes*

$$\sum_{p \in P} \delta(p, m). \tag{2.7}$$

A 1-median problem is also known as a *minisum* problem (e.g., [HPRT85, DTW86]). Set P must be nonempty for the value of Expr. (2.7) to be defined. Observe that minimizing the sum of the distances to the points of P is equivalent to minimizing the average distance to the points of P.

Unlike the 1-centre, the position of a 1-median is affected by multiplicities of points in P. Consequently, we consider finite multisets $P \in \mathscr{P}(\mathbb{R}^d)$.

As with the 1-centre, the term *1-median* derives directly from the more general *k-median*, where k denotes the number of facilities. We now generalize the single-facility definition of the 1-median to multiple facilities.

Definition 2.6. *Given a universe U, a finite multiset of points $P \in \mathscr{P}(U)$, a metric $\delta : U \times U \to \mathbb{R}^+$, and a positive integer k, a* **k-median** *of P is a set of k points $F \subseteq U$, that minimizes*

$$\sum_{p \in P} \min_{c \in F} \delta(p, c). \tag{2.8}$$

When the domain is Euclidean space and δ is a Minkowski distance, we refer to a k-median as a **geometric k-median**. If δ is the Euclidean distance metric, then we call it a *Euclidean k-median*.

2.4.2 Euclidean 1-Median

The corresponding Euclidean 1-median of P is a point in $\mathbb{R}^d$ that minimizes the sum of the Euclidean distances to points of P.

Definition 2.7. *Given a finite multiset $P \in \mathscr{P}(\mathbb{R}^d)$, a* **Euclidean 1-median** *of P is a function whose value, $M_d(P)$, is a point in $\mathbb{R}^d$ that minimizes*

$$\sum_{p \in P} ||p - M_d(P)||. \tag{2.9}$$

The value $\sum_{p \in P} ||p - M_d(P)||$ is referred to as the **Euclidean median sum** of P.

Unlike the Euclidean 1-centre, the Euclidean 1-median not always uniquely defined. If the points of P are not collinear, then the Euclidean 1-median is unique [KM97]. Similarly, if $|P|$ is odd, then the Euclidean 1-median is also unique. When the points of P are collinear and $|P| = 2n$, the points of P can be ordered and any point that lies on the line segment between the nth and $n+1$st largest elements is a 1-median of P. See Fig. 2.5. Since M_d is a function and must return a single point in $\mathbb{R}^d$, the common convention when defining the Euclidean 1-median, is to let M_d be the midpoint between these two elements [Wei]. Finally, the Euclidean 1-median of some (but not all) unbounded sets is uniquely defined. See Sec. 2.5.3.

Synonyms for the Euclidean 1-Median

The Euclidean 1-median has been rediscovered in a variety of contexts resulting in numerous names being assigned to it. The most common of these is *Weber point* [Baj88, BMM03, FMW05, Wes93]. Other names include Torricelli point [Kim98, Wei], Fermat point [Kim98], first Fermat point [Wei], generalized Fermat point [Wes93], first isogonic centre [Kim98, Wei], isogonal point

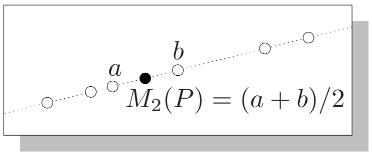

Figure 2.5: When the points of P are collinear and $|P|$ is even, by convention we define the Euclidean 1-median to be the midpoint of the two middle elements of P.

[KM97], ℓ_2 median, Euclidean median [DK05c, DK05b], median centre [Wes93], minisum problem [HLP+87, Wes93], spatial median [Wes93], Steiner problem [KM97, Wes93], bivariate median [Wes93], minimum aggregate travel point [Wes93], the point of equilibrium in a Varignon frame [Wes93], Kimberling triangle centre $X(13)$ [Kim98], or any combination of Fermat-Steiner-Torricelli-Weber point [BMM03, CT90, KM97, Pla95, Wes93]. Note, the centre of mass (centre of gravity) is sometimes incorrectly identified as being equivalent to the Euclidean 1-median; the centre of gravity is the point that minimizes the sum of the squares of distances [Sch73]. In addition, the term "median" sometimes refers to alternate generalizations of the median to higher dimensions. For example, Agarwal et al. [AdBG+05], use the term in reference to a point m such that for every line l through m, at least $k|P|$ points of P lie on either side of l, where $k \in [0, \frac{1}{2}]$ is fixed. Finally, the Euclidean 1-median is sometimes defined with a non-negative weight assigned to each point [CT90, Wes93]; when the weights are rational this reduces to Def. 2.7 since we allow multiplicities of points. An overview of the history and solutions to the Euclidean 1-median problem can be found in [DKSW02, KM97, Wes93].

Algorithms for Finding the Euclidean 1-Median

The Euclidean 1-median problem on three points in the plane was first posed by Fermat [Fer91] and solved geometrically by Torricelli early in the 17th century [KV97]. Alternate geometric solution techniques were subsequently found by Cavalieri and Simpson [DKSW02]. In $\mathbb{R}$, a Euclidean 1-median is easily found in $\Theta(n)$ time, where $n = |P|$, by a linear-time selection algorithm. In general, solving for the exact location of the Euclidean 1-median in two or more dimensions is difficult. Bajaj states, "there exists no exact algorithm under models of computation where the root of an algebraic equation is obtained using arithmetic operations and the extraction of kth roots" [Baj88, p. 177]. Indeed, no polynomial-time algorithm is known, nor has the problem been shown to be NP-hard [Hak00]. The most common approximation algorithm is Weiszfeld's algorithm [Wei37], an iterative procedure that converges to the Euclidean 1-median. Chandrasekaran and Tamir [CT90] give a polynomial-time algorithm for an ϵ-approximation of the Euclidean 1-median. More recently, Indyk [Ind99] and Bose et al. [BMM03] both give randomized ϵ-approximations algorithms with running times linear in n and polynomial in $1/\epsilon$. Bose et al. also give an

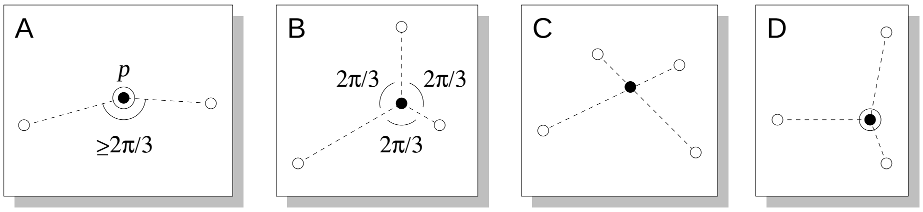

Figure 2.6: the Euclidean 1-median when $|P| = 3$ and $|P| = 4$

$O(n \log n)$-time deterministic ϵ-approximation algorithm.

Most approximation algorithms exploit the convexity of the objective function [LMW88, KM97]. That is, since the function $f(x) = ||p - x||$ is convex and the sum of convex functions remains convex, Expr. (2.9) must also be convex.

Finding the Euclidean 1-Median of Small Point Sets

The position of the Euclidean 1-median of P is well defined when $|P| \leq 4$. As shown by Heinen [Hei34], when $|P| = 3$, either a) the interior angle formed at some client $p \in P$ on the convex hull of P is at least $2\pi/3$ and the Euclidean 1-median of P coincides with p (see Fig. 2.6A) or b) all three interior angles are less than $2\pi/3$ and the Euclidean 1-median of P is located at the unique point whose position induces three angles of $2\pi/3$ with the clients of P (see Fig. 2.6B).

When $|P| = 4$ in $\mathbb{R}^2$, three cases are possible [KM97]: either a) the points of P are convex and the Euclidean 1-median of P is defined by the intersection of the two lines induced by opposite points (see Fig. 2.6C) b) the points P are collinear and any point between two middle points of P defines a Euclidean 1-median of P (see Fig. 2.5) or c) the points of P are neither convex nor collinear and the Euclidean 1-median of P coincides with the unique client of P located inside the convex hull of P (see Fig. 2.6D).

When $|P| = 4$ in $\mathbb{R}^3$, the cases are analogous to those for $|P| = 3$ in $\mathbb{R}^2$ [KM97]: either a) the points of P are collinear, b) the interior solid angle formed at some client $p \in P$ on the convex hull of P has measure at least π and the Euclidean 1-median of P coincides with p, or c) all interior solid angles have measure less than π and the Euclidean 1-median of P is the unique point inside the convex hull of P that forms four solid angles of measure π with the clients of P.

The Euclidean 1-median is invariant under similarity transformations. When $|P| \geq 5$, this invariance property allows the exact position of the Euclidean 1-median to be calculated for some configurations of points P by exploiting a rotational or reflectional symmetry to reduce the set of allowable locations for $M_d(P)$ (for example, see the proof of Thm. 5.21).

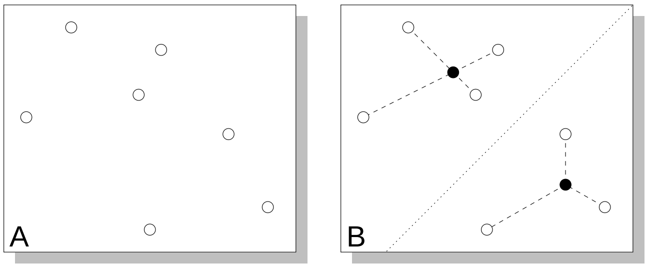

Figure 2.7: a set of points P in $\mathbb{R}^2$, a Euclidean 2-median of P, the corresponding Voronoi diagram, and the induced 2-partition

2.4.3 Euclidean k-Median

Definition 2.8. *Given a finite multiset $P \in \mathscr{P}(\mathbb{R}^d)$ and a positive integer k, a* **Euclidean k-median** *of P is a set of k points in $\mathbb{R}^d$, $\{M_d^1(P), \ldots, M_d^k(P)\}$, that minimizes*

$$\sum_{p \in P} \min_{1 \leq i \leq k} ||p - M_d^i(P)||. \tag{2.10}$$

The value $\sum_{p \in P} \min_{1 \leq i \leq k} ||p - M_d^i(P)||$ is referred to as the **Euclidean k-median sum** of P.

Synonyms for the Euclidean k-Median

The Euclidean k-median is also known as Euclidean p-median [TFL83a, Est01], spatial k-median problem [HM03], minisum location-allocation problem [DC97], multisource Weber problem [HMT98, BHMT00], generalized multi-Weber problem [DC97], k-hub location problem [SP97], multi-switch location problem [VP03], multi-depot location problem [DC97], two centre location-allocation problem (when $k = 2$) [Ost75], and planar k-median (in $\mathbb{R}^2$) [Dre84b].

Algorithms for Finding the Euclidean 2-Median

The solution space for the Euclidean k-median is neither convex nor concave [Coo67]. Naturally, finding a Euclidean k-median is at least as difficult as finding a Euclidean 1-median, meaning that no algorithms are known for finding an exact Euclidean k-median, even in $\mathbb{R}^2$ when $k = 2$.

Since a client is always served by the nearest facility, the Voronoi diagram of a Euclidean 2-median of a set of points P consists of a single straight line (plane, hyperplane) that partitions P into two sets. See Fig. 2.7. The number of possible dividing lines (planes, hyperplanes) is proportional to the number of pairs of clients [CRW91]. Algorithms for finding an approximate Euclidean 2-median in $\mathbb{R}^2$ provided by Ostresh [Ost75] and Drezner [Dre84b] exploit this property by exhaustively examining all such possible partitions of the client set, identifying an approximate Euclidean 1-median for every subpartition, and taking the minimum over all possible solutions. Consequently, both algorithms

	$k = 1$	$k = 2$	k fixed	k arbitrary
$\mathbb{R}$	$\Theta(n)$	$\Theta(n)$	$\Theta(kn)$	$\Theta(kn)$
$\mathbb{R}^2$	complexity unknown			NP-hard
$\mathbb{R}^d$				

Table 2.2: time complexities of algorithmic solutions to the Euclidean k-median

have time complexity $O(n^2 f(n))$, where $f(n)$ is the time complexity of finding an approximate Euclidean 1-median for a set of n points. Rosing [Ros92] extends this technique to find an approximate Euclidean 3-median in $\mathbb{R}^2$.

Algorithms for Finding the Euclidean k-Median

In $\mathbb{R}$, Hassin and Tamir [HT91] show that the Euclidean k-median can be solved exactly in $O(kn)$ time. Like the Euclidean k-centre, when k is an arbitrary input parameter, Megiddo and Supowit [MS84] show the Euclidean k-median problem is NP-hard in $\mathbb{R}^2$; they also show that finding an ϵ-approximation remains NP-hard for any $\epsilon < 3/2$. These results are summarized in Tab. 2.2.

As for approximate solutions, Jain and Vazirani [JV99] give a 6-approximation algorithm in $O(n^2)$ time. Charikar and Guha [CG99] give a 4-approximation algorithm in $O(n^3)$ time. Finally, Arora et al. [ARR98] give an $O(n^{O(1+1/\epsilon)})$-time ϵ-approximation. Kolliopoulos and Rao [KR99] provide a randomized approximation scheme that returns a solution expected to be within a factor of $1+\epsilon$ of the optimum, requiring $O(2^{1/\epsilon^d} n \log n \log k)$ time. Given fixed values for ϵ, k, and d, Har-Peled and Mazumdar [HM04] provide a linear-time algorithm for finding a $(1+\epsilon)$-approximation of the Euclidean k-median using coresets.

2.5 Related Work in Geometric Facility Location

This section provides a brief overview of other key areas within the wide range of problems of facility location. The rectilinear k-centre, rectilinear k-median, centre of mass, k-means clustering, and continuous facility location are concepts that will be revisited in Chs. 4 through 7. This section provides an overview of results related to these problems.

The section concludes with a brief description of other key areas within facility location, both those outside geometric facility location and those involving restrictions, augmentations, or interesting variations of problems in geometric facility location. These include two other major areas of facility location: facility location on networks and discrete facility location. These topics define tangential areas only indirectly related to the work of this thesis. Familiarity with the contents of this section is not essential to understanding the results and contributions of this thesis.

2.5.1 Rectilinear k-Centre and k-Median

Problems that are difficult under Euclidean distance are sometimes solved more easily under a distance metric not defined in terms of radicals. Chebyshev distance and rectilinear distance are the only two such Minkowski distance metrics. Although ambiguous, the facility location literature defines the *rectilinear k-centre* as a geometric k-centre for which the distance metric δ is Chebyshev (ℓ_∞) distance whereas the *rectilinear k-median* is defined as a geometric k-median for which the distance metric δ is rectilinear (ℓ_1) distance. The definitions of the rectilinear k-centre and rectilinear k-median which we establish formally in this section will be used in Chs. 4 through 6.

Rectilinear k-Centre

The definition of the rectilinear k-centre is analogous to that of the Euclidean k-centre, but with respect to the Chebyshev (ℓ_∞) distance metric.

The **Chebyshev distance** between two points $x = (x_1, \ldots, x_d)$ and $y = (y_1, \ldots, y_d)$ in $\mathbb{R}^d$ is given by

$$||x - y||_\infty = \lim_{p \to \infty} ||x - y||_p = \max_{1 \le i \le d} |x_i - y_i|. \tag{2.11}$$

Chebyshev distance is also known as Tchebychev, chessboard, maximum, minimax, or ℓ_∞ distance. Although ambiguous, Chebyshev distance is sometimes also referred to as *rectilinear distance*, the name more commonly used to refer to ℓ_1 distance, the Minkowski distance metric when $p = 1$.

Definition 2.9. *Given a finite set $P \in \mathscr{P}(\mathbb{R}^d)$ and a positive integer k, a* **rectilinear k-centre** *of P is a set of k points in $\mathbb{R}^d$, $\{R_d^1(P), \ldots, R_d^k(P)\}$, that minimizes*

$$\max_{p \in P} \min_{1 \le i \le k} ||p - R_d^i(P)||_\infty. \tag{2.12}$$

The value $\max_{p \in P} \min_{1 \le i \le k} ||R_d(P) - p||_\infty$ is referred to as the **rectilinear k-radius** of P. The rectilinear k-centre is also known as rectangular p-centre [Dre87] and rectilinear minimax.

Just as a Euclidean k-centre of a set P in $\mathbb{R}^d$ corresponds to the centres of k hyperspheres of minimum radius that enclose the points of P, a rectilinear k-centre of P corresponds to the centres of k axis-parallel hypercubes, such that the width of the largest hypercube is minimized while also enclosing the points of P. See Fig. 2.9B. Since the enclosing hypercube of a set of points is not unique, it is common to select the centre of the corresponding bounding box. See Fig. 2.8A.

Since the Chebyshev norm is not invariant under rotation or reflection, the rectilinear k-centre is not invariant under these transformations. See Fig. 2.8B. It is straightforward to show that the rectilinear k-centre is invariant under translation and scaling.

Just as the Euclidean 1-centre is the centre of the smallest enclosing hypersphere and the rectilinear 1-centre is the centre of the smallest enclosing

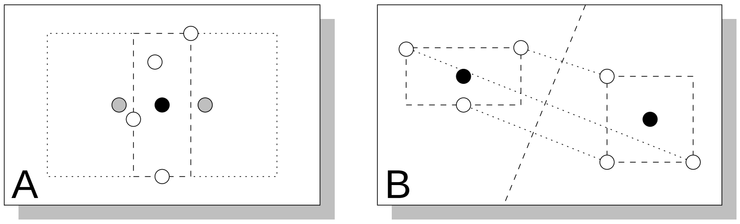

Figure 2.8: **A.** A rectilinear 1-centre is not unique whereas the centre of the bounding box is unique. **B.** The rectilinear 1-centre is not invariant under reflection or rotation.

axis-parallel box, other generalizations are possible to the centre of the smallest enclosing diamond (the ℓ_1 1-centre), ellipsoid, cylinder, tetrahedron, rectangular box, and parallelepiped in $\mathbb{R}^3$ [VW04].

Algorithms for Finding the Rectilinear k-Centre

All Minkowski norms coincide in $\mathbb{R}$. Therefore, finding a rectilinear k-centre in $\mathbb{R}$ corresponds to finding a Euclidean k-centre in $\mathbb{R}$. Finding a rectilinear 1-centre of P in $\mathbb{R}^d$ reduces to solving d independent one-dimensional geometric 1-centre problems. That is, the rectilinear 1-centre is found by identifying the extreme points of P along each dimension and returning the midpoint of each. A single scan of the clients of P suffices, requiring $\Theta(nd)$ time.

Drezner [Dre84b] provides a linear-time solution to the rectilinear 2-centre in $\mathbb{R}^2$. Still in linear time, this result is extended to higher dimensions and to the weighted case by Ko and Ching [KC92]. Similarly, the rectilinear 3-centre in $\mathbb{R}^2$ is solved in linear time by Hoffmann [Hof05]. The analogous problem in $\mathbb{R}^d$ can be solved in time $O(n \log n)$ [AK99], matching the lower bound of $\Omega(n \log n)$ [Hof01]. The rectilinear 4-centre is solved in $O(n \log n)$ by Sharir and Welzl [SW96]. Chan [Cha98] gives a randomized algorithm with $O(n \log n)$ expected time for the rectilinear 5-centre. Sharir [SW96] provides the fastest deterministic algorithm for the 5-centre in $\mathbb{R}^2$, requiring $O(n \log^4 n)$ time.

Nussbaum [Nus97] gives an $O(n^{k-4} \log n)$-time algorithm for the rectilinear k-centre in $\mathbb{R}^2$ for a fixed k. As stated by Agarwal and Procopiuc [AP98], the techniques of Hwang et al. [HLC93b] for solving the Euclidean k-centre in $\mathbb{R}^2$ (k fixed) in $n^{O(\sqrt{k})}$ can be generalized to the rectilinear k-centre with the same running time. Finally, Agarwal and Procopiuc [AP98] provide an algorithm for the rectilinear k-centre in $\mathbb{R}^d$ (k fixed) in $n^{O(k^{1-1/d})}$.

In $\mathbb{R}^2$, Feder and Greene [FG88] and Ko et al. [KLC90] show it is NP-hard to approximate the rectilinear k-centre problem with an approximation factor less than 2 when k is an arbitrary input parameter. These results are summarized in Tab. 2.3.

Gonzalez [Gon85] gives a 2-approximation for the rectilinear k-centre in $\mathbb{R}^2$ in $O(nk)$ time. This time was reduced to $O(n \log k)$ by Feder and Greene [FG88]. In $\mathbb{R}^d$, Hochbaum and Shmoys [HS86] give a 2-approximation algorithm in $O(nk)$

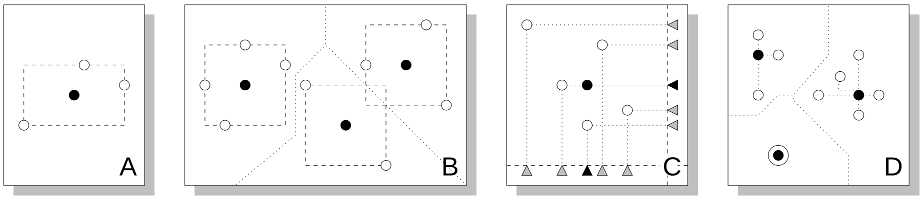

Figure 2.9: examples of the rectilinear 1-centre, 3-centre, 1-median, and 3-median in $\mathbb{R}^2$ including the corresponding ℓ_∞ and ℓ_1 Voronoi diagrams and the projection of client positions onto the axes

time. Still in $\mathbb{R}^d$, Agarwal and Procopiuc [AP98] provide a $(1+\epsilon)$-approximation algorithm for the rectilinear k-centre problem in $O(n \log k + k/\epsilon^{dk})$ time.

Rectilinear k-Median

The definition of the rectilinear k-median is analogous to that of the Euclidean k-median, but with respect to the rectilinear (ℓ_1) distance metric.

The **rectilinear distance** between two points $x = (x_1, \ldots, x_d)$ and $y = (y_1, \ldots, y_d)$ in $\mathbb{R}^d$ is given by

$$||x - y||_1 = \sum_{i=1}^{d} |x_i - y_i|. \tag{2.13}$$

Rectilinear distance is also known as Manhattan, city block, taxicab, rectangular, metropolitan, or ℓ_1 distance [SS01].

Definition 2.10. *Given a finite multiset $P \in \mathscr{P}(\mathbb{R}^d)$ and a positive integer k, a* **rectilinear k-median** *of P is a set of k points in $\mathbb{R}^d$, $\{S_d^1(P), \ldots, S_d^k(P)\}$, that minimizes*

$$\sum_{p \in P} \min_{1 \le i \le k} ||p - S_d^i(P)||_1. \tag{2.14}$$

The value $\sum_{p \in P} \min_{1 \le i \le k} ||p - S_d^i(P)||_1$ is referred to as the **rectilinear k-median sum** of P. The rectilinear k-median is also known as rectangular p-median [Dre87], rectilinear minisum, and coordinate median [Wes93] (when $k = 1$).

Since the rectilinear distance metric is not invariant under rotation or reflection, the rectilinear k-median is not invariant under these transformations. It is straightforward to show that the rectilinear k-median is invariant under translation and scaling.

Algorithms for Finding the Rectilinear k-Median

As is the case for the rectilinear k-centre, finding a rectilinear k-median in $\mathbb{R}$ corresponds to finding a Euclidean k-median in $\mathbb{R}$ since all Minkowski norms coincide in one dimension. As shown by Wendell and Hurter [WH73a], to find

rectilinear k-centre

<table>
<tr><td></td><td>$k=1$</td><td>$k=2$</td><td>k fixed</td><td>k arbitrary</td></tr>
<tr><td>$\mathbb{R}$</td><td colspan="4">same as Euclidean k-centre</td></tr>
<tr><td>$\mathbb{R}^2$</td><td>$\Theta(n)$</td><td>$\Theta(n)$</td><td>$n^{O(\sqrt{k})}$</td><td rowspan="2">NP-hard</td></tr>
<tr><td>$\mathbb{R}^d$</td><td>$\Theta(dn)$</td><td>$\Theta(n \log n)$</td><td>$n^{O(k^{1-1/d})}$</td></tr>
</table>

rectilinear k-median

<table>
<tr><td></td><td>$k=1$</td><td>$k=2$</td><td>k fixed</td><td>k arbitrary</td></tr>
<tr><td>$\mathbb{R}$</td><td colspan="4">same as Euclidean k-median</td></tr>
<tr><td>$\mathbb{R}^2$</td><td>$\Theta(n)$</td><td>$O(n^5)$</td><td>$O(n^{2k+1})$</td><td rowspan="2">NP-hard</td></tr>
<tr><td>$\mathbb{R}^d$</td><td>$\Theta(dn)$</td><td>$O(n^{2d+1})$</td><td>$O(n^{dk+1})$</td></tr>
</table>

Table 2.3: time complexities of algorithmic solutions to the rectilinear k-centre and k-median

possible locations for a rectilinear k-median of a set P in $\mathbb{R}^d$, one need only consider intersection points within the convex hull of P. That is, the projection of P onto each axis induces a grid of $O(n^d)$ points. Those that lie within the convex hull of P are candidates for defining a rectilinear k-median of P.

The rectilinear 1-median in $\mathbb{R}^d$ is solved in linear time by solving d independent one-dimensional 1-median problems along each dimension [Baj84]. See Fig. 2.9C.

The hardness results of Megiddo and Supowit [MS84] for the Euclidean k-median also apply to the rectilinear version of the problem. That is, the rectilinear median problem is NP-hard in $\mathbb{R}^2$ when k is an arbitrary input parameter. Furthermore, the problem remains hard when approximating to within a factor of $3/2$.

Little is known on the complexity of the rectilinear k-median when k is fixed in two or more dimensions. The property restricting possible solutions to intersection points leads to a brute-force algorithm requiring $O(n^{dk+1})$ time. These results are summarized in Tab. 2.3.

Several approximation algorithms for the Euclidean k-median can be used to find approximate solutions to the rectilinear k-median in $\mathbb{R}^d$. Jain and Vazirani [JV99] give a 6-approximation algorithm in $O(n^2)$ time. Charikar and Guha [CG99] give a 4-approximation algorithm in $O(n^3)$ time. Arora et al. [ARR98] give an $O(n^{O(1+1/\epsilon)})$-time ϵ-approximation. In addition, heuristic solutions are common, including heuristics using Tabu search [Ohl97] and Kohonen self-organizing feature maps [HT04].

2.5.2 Centre of Mass and k-Means Clustering

The objective of the Euclidean k-median optimization function involves minimizing the sum of the Euclidean distances from the clients to their respective nearest facilities. The k-means clustering problem is a close relative of the Euclidean k-median, for which the optimization function involves minimizing the sum of the squared Euclidean distances. For a single facility ($k = 1$) this function has a single optimum, more commonly known as the *centre of mass*.

Centre of Mass
We define the centre of mass.

Definition 2.11. *Given a finite multiset $P \in \mathscr{P}(\mathbb{R}^d)$, the* **centre of mass** *of P is the function whose value, $C_d(P)$, is the point in $\mathbb{R}^d$ given by*

$$C_d(P) = \frac{1}{|P|} \sum_{p \in P} p. \tag{2.15}$$

The centre of mass is the point that minimizes the sum of the square distances [Sch73, Wes93]. This is easily seen by the following derivation. Given a set P in $\mathbb{R}^d$, let c be a point in $\mathbb{R}^d$ that minimizes the sum of the squares of the distances to the points of P. That is, c minimizes

$$\sum_{p \in P} ||c - p||^2. \tag{2.16}$$

The partial derivatives of Expr. (2.16) with respect to c must all be zero. Thus, for all $1 \leq i \leq d$,

$$\begin{aligned} & \frac{\partial}{\partial c_i} \sum_{p \in P} ||c_i - p||^2 = 0, \\ \Rightarrow \quad & \sum_{p \in P} (c_i - p_i) = 0, \\ \Rightarrow \quad & c_i = \frac{1}{|P|} \sum_{p \in P} p_i, \end{aligned} \tag{2.17}$$

where $c = (c_1, \ldots, c_d)$ denote the components of c (respectively, p) in dimensions $1, \ldots, d$. Eq. (2.17) matches the definition of the centre of mass in Eq. (2.15).

The centre of mass is also known as geometric centroid [Wei], least squares point, centroid, mean, 1-mean, centre of gravity [Sch73], and Kimberling triangle centre $X(2)$ [Kim98].

Function C_d is invariant under affine transformations. The position of the centre of mass is is easily constructed in $\Theta(n)$ time.

k-Means Clustering
The k-means clustering problem is the generalization of the centre of mass to multiple facilities.

Definition 2.12. *Given a finite multiset $P \in \mathscr{P}(\mathbb{R}^d)$ and a positive integer k, a* **k-means clustering** *of P is a set of k points in $\mathbb{R}^d$, $\{C_d^1(P), \ldots, C_d^k(P)\}$, that minimizes*

$$\sum_{p \in P} \min_{1 \leq i \leq k} ||p - C_d^i(P)||^2. \tag{2.18}$$

Observe that the squared Euclidean norm $||\cdot||^2$ is not a distance metric since the triangle inequality does not hold. The optimization function of a k-means cluster in Expr. (2.18) corresponds to the sum of the variance of each cluster [HE02].

Synonyms for k-Means Clustering

The k-means clustering problem is also known as least squares clustering [EE04], least squares quantization [Llo82], minimum variance clustering [DM00], variance-based k-clustering [IKI94], k-cluster centroid [DM00], k-cluster mean [DM00], generalized Lloyd Max problem [GJW82], ℓ_2^2 k-median clustering [VKKR03], and minimum sum-of-squares clustering [HM01]. The measure (value of the optimization function, Expr. (2.18)) is also known as squared error distortion [KMN+02b] and root-mean-square distance [AM04].

Complexity of k-Means Clustering

The exact complexity of k-means clustering is unclear. Some go so far as to claim that k-means clustering is NP-complete even for $k = 2$ [VKKR03, KSS04, SS05]. This is surprising given that polynomial-time algorithms exist for solving 2-means clustering in $\mathbb{R}^d$ when d is fixed. The source of this confusion may be a citation by Sabharwal and Sen [SS05] of an article by Drineas et al. [DFK+99]; Drineas et al. state that *discrete* 2-means clustering is NP-hard in $\mathbb{R}^d$, which is cited by Sabharwal and Sen [SS05] and interpreted to mean that (non-discrete) 2-means clustering is NP-hard. In general, it seems widely believed that k-means clustering is NP-complete when k is an arbitrary input parameter in two or more dimensions [DM00, KMN+02a, Mer03] but again, the source of this result is unclear. The literature makes frequent reference to the paper of Brucker [Bru78] as evidence that k-means clustering is NP-complete. Although Brucker's paper proves NP-completeness for several related clustering problems, the hardness of k-means clustering does not appear to be an immediate consequence of these results. Brucker's results are often referenced indirectly via Garey and Johnson [GJ79]. A second source cited is the work of Garey et al. [GJW82, DM00] showing NP-hardness for the generalized Lloyd-Max problem. Again, the hardness of k-means clustering does not appear to be an immediate consequence of these results. According to Mount [Mou05], an expert on k-means clustering, no proof of NP-hardness nor any polynomial-time algorithm has been presented to date.

Algorithms for Finding a k-Means Clustering

In $\mathbb{R}$, k-means clustering can be solved using dynamic programming in $O(kn^3)$ time. The clients are first sorted. Partial solutions for an optimal k'-means clustering on the first n' clients of P are stored in a $k \times n$ array. The optimal solution for a k-means clustering of all n clients is given by examining all possible $(k-1)$-means clusterings for 1 through $n-1$ clients, calculating the sum of the squared distances for the last cluster in each case, and selecting the minimum over all cases.

	$k=1$	$k=2$	k fixed	k arbitrary
$\mathbb{R}$	$\Theta(n)$	$O(n^2)$	$O(kn^3)$	$O(kn^3)$
$\mathbb{R}^2$	$\Theta(n)$	$O(n^3)$	$O(n^{2k(k-1)/2+1})$	complexity unknown
$\mathbb{R}^d$	$\Theta(dn)$	$O(n^{d+1})$	$O(n^{dk(k-1)/2+1})$	

Table 2.4: time complexities of algorithmic solutions to the k-means clustering problem

The facility nearest to a client p under Euclidean distance remains nearest to p under the squared Euclidean distance. Consequently, given a 2-means clustering of a set of clients P, the corresponding 2-partition of P must be separable by a hyperplane, as is the case for the Euclidean 2-centre and 2-median [Ost75, Dre84b]. In $\mathbb{R}^2$ there are $O(n^2)$ possible choices for a dividing line to partition the clients. The centre of mass for each partition is found in $O(n)$ time, giving a total time of $O(n^3)$. In $\mathbb{R}^d$ a hyperplane is uniquely defined by d linearly independent points, of which there are $O(n^d)$ possible choices [Ost75]. Thus, this algorithm solves the 2-means clustering problem in $\mathbb{R}^d$ in $O(n^{d+1})$ time [HII$^+$93, IKI94]. The 3-means clustering problem is solved in $\mathbb{R}^2$ in $O(n^5 \log n)$ time by Hasegawa et al. [HII$^+$93]. For a fixed k, the k-means clustering problem can be solved in $O(n^{dk(k-1)/2+1})$ time in $\mathbb{R}^d$ [HII$^+$93]. These results are summarized in Tab. 2.4.

Kanungo et al. [KMN$^+$02b] give a $(9+\epsilon)$-approximation algorithm to k-means clustering in $\mathbb{R}^d$ (no time complexity is given). Still in $\mathbb{R}^d$, Matoušek [Mat00] gives a $(1+\epsilon)$-approximation algorithm to the k-means clustering problem that runs in time $O(n \log^k n \epsilon^{-2k^2 d})$ for any fixed k and d. De la Vega et al. [VKKR03] describe a randomized algorithm that returns a $(1+\epsilon)$-approximation with constant probability in $\mathbb{R}^d$ requiring $O(g(k,\epsilon) n \log^k n)$ time, where $g(k,\epsilon) = \exp(k^3 \ln k [\ln(1/\epsilon) + \ln k]/\epsilon^8)$. Hasegawa et al. [HII$^+$93] give an $O(n^{k+1})$ time 2-approximation to k-means clustering in $\mathbb{R}^d$ for a fixed d. Inaba et al. [IKI94] provide a $(1+\epsilon)$-approximate randomized algorithm to 2-means clustering in $\mathbb{R}^d$ that runs in $O(n(1/\epsilon)^d)$ time. Sabharwal and Sen [SS05] provide a $(1+\epsilon)$-approximate randomized algorithm to 2-means clustering in $\mathbb{R}^d$ that runs in time $O(1/\epsilon^{O(1/\epsilon)}(d/\epsilon)^d n)$ with constant probability. Given fixed values for ϵ, k, and d, Har-Peled and Mazumdar [HM04] provide a linear-time algorithm for finding a $(1+\epsilon)$-approximation of k-means clustering using coresets. Agarwal et al. [AHV05] provide a survey of approximation algorithms for k-means clustering that make use of coresets.

The term *k-means algorithm* sometimes refers to Lloyd's method, an iterative heuristic used for approximating a solution to the k-means clustering problem in $\mathbb{R}^d$ [Llo82, KMN$^+$00, EE04]. Additional popular heuristics solutions to k-means clustering were introduced by MacQueen [Mac67] and Ball and Hall [BH67].

2.5.3 Continuous Facility Location

In facility location problems, clients positions are typically defined by a finite set of points. However, one may model a large set of clients, P, by a continuous[3] region of points over which we define a **client density function**, $\rho : P \to [0,1]$, such that $\int_{p \in P} \rho(p)\ dp = 1$. Equivalently, we can define $\rho : U \to [0,1]$ and add the requirement that $\rho(p) = 0$ for all $p \notin P$.

As argued by Drezner [Dre95], continuous facility location is useful for modelling large clients sets. In these cases, numerical error from discretization is reduced by employing a client density function.

The definition of the continuous Euclidean 1-centre is a straightforward generalization of Def. 2.3. The client set (now a region) must be bounded for the 1-centre to be defined.

Definition 2.13. *Given a bounded set $P \in \mathscr{P}(\mathbb{R}^d)$, the* **continuous Euclidean 1-centre** *of P is the function whose value, $\Xi_d(P)$, is the point in $\mathbb{R}^d$ that minimizes*

$$\max_{p \in \overline{P}} ||p - \Xi_d(P)||, \tag{2.19}$$

where $\overline{P}$ denotes the closure of set P.

The continuous Euclidean k-centre is defined similarly by the corresponding generalization of Def. 2.4.

The objective function for the continuous 1-median is described in terms of an integration over the client set as opposed to a sum. Thus, when $|P|$ is infinite, Def. 2.7 is generalized by integrating.

Definition 2.14. *Given an arbitrary set $P \in \mathscr{P}(\mathbb{R}^d)$, let $\rho : P \to [0,1]$ denote the client density function of the points of P within $\mathbb{R}^d$ such that $\int_{p \in P} \rho(p)\ dp = 1$. A* **continuous Euclidean 1-median** *of P is a function whose value, $M_d(P)$, is a point in $\mathbb{R}^d$ that minimizes*

$$\int_{p \in P} \rho(p) ||p - M_d(P)||\ dp. \tag{2.20}$$

Note that Def. 2.14 reduces to Def. 2.7 when P is finite. Under Euclidean distance in $\mathbb{R}^d$ for $d \geq 2$, the continuous Euclidean 1-median problem is at least as hard as the Euclidean 1-median on a finite universe [Baj88, FMW05]. Fekete et al. [FMW00, FMW05] give an $O(n)$-time algorithm for solving the continuous 1-median in $\mathbb{R}^2$ under rectilinear distance, where the region of clients is polygonal and n denotes the number of vertices on its boundary. Carmi et

[3] The term *continuous facility location* is sometimes used to refer to any problem in facility location for which the universe is a *continuous space* (see Sec. 2.2.2), even when the set of clients is itself finite (for example, [Pla95, Pla02]). We differentiate between these by referring to a *continuous space universe* versus a *discrete space universe* and apply the more common definition of *continuous facility location* to describe those problems that include a continuous region of clients representing an infinite distribution.

al. [CHPK05] provide a linear-time algorithm for finding an approximation of the Euclidean 1-median of a convex region.

The continuous Euclidean k-median is defined similarly by the corresponding generalization of Def. 2.8:

Definition 2.15. *Given a set $P \in \mathscr{P}(\mathbb{R}^d)$, and a positive integer k, let $\rho : P \to [0,1]$ denote the client density function of the points of P within $\mathbb{R}^d$ such that $\int_{p \in P} \rho(p)\ dp = 1$. A* **continuous Euclidean k-median** *of P is a set of k points in $\mathbb{R}^d$, $\{M_d^1(P), \ldots, M_d^k(P)\}$ that minimizes*

$$\int_{p \in P} \min_{1 \leq i \leq k} \rho(p)||p - M_d^i(P)||\ dp. \tag{2.21}$$

Fekete et al. [FMW00, FMW05] show the continuous rectilinear k-median is NP-hard in $\mathbb{R}^2$ when k is an input to the problem.

When $|P|$ is infinite, the centre of mass (Def. 2.11) is also generalized by integrating.

Definition 2.16. *Given a set $P \in \mathscr{P}(\mathbb{R}^d)$, let $\rho : P \to [0,1]$ denote the client density function of the points of P within $\mathbb{R}^d$ such that $\int_{p \in P} \rho(p)\ dp = 1$. The* **continuous centre of mass** *of P is the function whose value, $C_d(P)$, is the point in $\mathbb{R}^d$ given by*

$$C_d(P) = \int_{p \in P} \rho(p) p\ dp. \tag{2.22}$$

Note that Def. 2.16 reduces to Def. 2.11 when P is finite.

In this thesis we implicitly apply definitions of continuous facility location. That is, whenever P is not finite, a facility function F is understood to refer to the definition of F under continuous facility location.

2.5.4 Additional Constraints and Related Problems

This section identifies and briefly describes additional major areas within facility location that are only indirectly related to the focus of this thesis.

Facility Location on Networks

Unlike geometric facility location, where clients and facilities reside at any point in $\mathbb{R}^d$ and distances between points are defined by a Minkowski distance metric, facility location on graphs allows more general distance metrics while restricting the set of possible locations for facilities. Facility location on graphs has a well-explored set of problems accompanied by an extensive literature. A problem instance consists of a weighted graph, $G = (V, E, d, w)$, where the vertex set V corresponds to the set of clients, $d : E \to \mathbb{R}^+$ assigns positive weights to edges, and $w : V \to \mathbb{R}^+$ assigns non-negative weights to vertices. The weight, $d(e)$, of an edge, $e = (u, v) \in E$, is a positive real representing the distance between u and v. The weight w_v of vertex v is a non-negative real representing the demand of client v. Weights may be normalized. Typically it is required that the triangle inequality hold for edge weights.

A facility p must be located either exclusively on a vertex or anywhere along an edge of the graph. If p lies on edge $e = (u, v)$, its position is defined by u and v and a real value $a \in [0, 1]$ such that the distance from p to u is $a \cdot d(u, v)$ and the distance from p to v is $(1 - a)d(u, v)$. Hakimi gives this definition for distance on a graph:

> If x and y are any two points on G, the *distance* $d(x, y)$ is the length of the shortest path between x and y in G, where the length of a path is the sum of the lengths of the edges (or partial edges) in the path. [Hak00, p. 987]

Given any graph G, a k-median on G exists such that all k facilities are vertices of G [Hak00]. This theorem does not hold for the k-centre on graphs. In the *vertex k-centre* problem, facilities are required to lie on a vertex while the *absolute k-centre* problem allows facilities to lie on a vertex or anywhere along an edge.

Let $n = |V|$ and let $m = |E|$. Hakimi and Kariv [KH79a] give $O(mn + n^2 \log n)$-time algorithm for the unweighted centre problem on graphs and time $O(mn \log n)$ in the weighted case. The latter problem has its runtime reduced to $O(n^2)$ if an all-pairs shortest path distance matrix is given [EL95]. Hakimi [Hak64] gives an $O(n^3)$-time algorithm for the median problem on graphs; this is also reduced to $O(n^2)$ if an all-pairs shortest path distance matrix is given [EL95]. When k is an input parameter, weighted or unweighted k-centre and k-median are NP-hard [KH79a]. When k is fixed, k-centre can be solved in $O(m^k n^k \log^2 n)$ and k-median can be solved in $O(n^{k+1})$ [Tam88].

Frederickson [Fre91] shows that the k-centre of a tree can be found in linear time. Tamir [Tam96] gives an algorithm for finding the k-median of a tree that requires $O(n^2 k)$. Hakimi gives linear algorithms for the median problem and the unweighted centre problem on trees [Hak00]. Hakimi and Kariv [KH79a] give an $O(n \log n)$-time algorithm for the weighted centre problem on a tree.

A review of single-facility location problems on networks is given in [HLPT87]. As for multiple-facility location problems, [KH79a, KH79b], and [TFL83a] provide reviews of the k-centre and k-median problems on graphs while [TFL83b] reviews these problems on trees.

Discrete Facility Location

A problem in discrete facility location is any facility location problem for which the domain of allowable facility positions is restricted to a finite set. In particular, a common restriction within discrete facility location is that the set of allowable facility locations be restricted to the positions of the input client set. That is, given a set of clients P contained in some universe U and an integer k, select a set $F \subseteq P$, $|F| = k$, that minimizes the optimization function. Discrete facility location is sometime referred to as *metric facility location* [CGTS99].

The discrete k-median problem on any distance metric is NP-hard when k is an arbitrary input parameter [CGTS99]. In $\mathbb{R}^d$, Bereg et al. [BKST99] give a

$\Theta(nd)$-time algorithm for the discrete rectilinear 1-median problem. Under ℓ_∞ they give an $O(n \log^2 n)$-time algorithm. Charikar et al. [CGTS99] provide a $6\frac{2}{3}$-approximation algorithm to the discrete k-median problem under any distance metric.

In $\mathbb{R}^2$, the Euclidean discrete k-centre is NP-hard if k is an input variable [AS98]. Still in $\mathbb{R}^2$, Hwang et al. [HLC93a] give a $n^{O(\sqrt{k})}$-time algorithm for the discrete Euclidean k-centre. Agarwal et al. [ASW98] give an algorithm for the discrete Euclidean 2-centre in $\mathbb{R}^2$ that runs in time $O(n^{4/3} \log^5 n)$. As noted in [ASW98], the discrete Euclidean 1-centre in $\mathbb{R}^2$ is solved in $O(n \log n)$ time by finding the furthest-neighbour Voronoi diagram of the set of client positions and selecting the client with the nearest furthest neighbour. Finally, Agarwal and Procopiuc [AP98] provide an algorithm for the discrete Euclidean k-centre in $\mathbb{R}^d$ (k fixed) in $n^{O(k^{1-1/d})}$ time.

Bereg and Kirkpatrick [BK99] give an algorithm for the discrete rectilinear 2-centre in $\mathbb{R}^d$ in $O(n \log^{d-2} n \log \log n + n \log n)$ running time. In $\mathbb{R}^2$, Bereg and Segal [BS99, BS01] give an optimal algorithm for the discrete rectilinear 2-centre problem in $O(n+m)\log(n+m)$ time, where m denotes the cardinality of the set of points from which the facilities are drawn.

See Mirchandani and Francis [MF90] for an overview of the problems of discrete facility location.

Discrete facility location is not conducive to continuous motion of facilities. As is discussed in Ch. 3, one of the objectives of this thesis is to model situations in which clients and facilities are free to move continuously, providing yet further motivation for examining points in continuous space, as opposed to discrete space.

The term "facility location" sometimes refers to discrete facility location with costs associated with each potential facility. Given a set of potential facilities, each with some associated cost, a set of clients, and a distance metric (proportional to cost), the problem is to select a subset of the facilities (any number of them) to minimized the total cost of shipping a product to every client from the facility nearest to that client while including the cost of opening each new facility.

Capacitated Facility Location

To model facilities more realistically, a fixed upper bound $\alpha > 0$ is introduced such that the sum of the weights of all clients served by a single facility may not exceed α. Let F, $|F| = k$, denote the set of facilities to be positioned, let P denote the set of clients, and let $w(p)$ denote the weight of client p. For a solution to exist, a necessary condition is that

$$k \cdot \alpha \geq \sum_{p \in P} w(p). \tag{2.23}$$

Hakimi [Hak00] describes two formulations for capacitated facility location. The first requires a discrete allocation of clients to facilities such that each client is served by the nearest facility. If a client lies an equal distance from two or

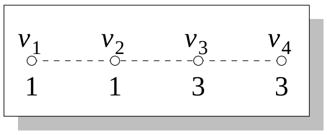

Figure 2.10: When $\alpha = 4$, no capacitated 2-centre or 2-median exists.

more facilities, then it may be served by any of these. Alternatively, clients may distribute their demand among several facilities, without requiring to be served solely by the nearest facility. These formulations differ in two ways. The first formulation requires a client to be served by a single facility while the second formulation allows a client to have its demand split between several facilities, so long as these sum to the client's demand. The first formulation requires a facility to be served by the closest facility to it while the second formulation allows a client to be served by a distant facility.

For some problem instances, the first formulation of the problem may not have any solution. For example, let four points lie on a line such that the first two have weight 1 and the last two have weight 3. When $\alpha = 4$, no capacitated 2-centre or 2-median exists under Euclidean distance. See Fig. 2.10.

In general, the additional constraints of capacity correspond to combinatorial problems as opposed to the geometric constraints of position and velocity under Euclidean distance which we consider.

Obnoxious Facility Location

The goal of every facility location problem described until now has been to minimize the sum or the maximum of the distances between clients and facilities. This goal can be reversed to maximize these distances, modelling the selection of positions for a set of undesirable facilities. Of course, the domain must be finite, otherwise a facility could be positioned at a distance approaching infinity. For models in $\mathbb{R}^d$, the domain is typically restricted to a d-dimensional region $[a_1, b_1] \times \ldots \times [a_d, b_d]$.

Unlike a k-median on a graph, an obnoxious k-median on a graph G does not always consist of vertices of G.

Eiselt and Laporte [EL95] give a good overview of obnoxious facility location. Moon and Chaudry [MC84] also include a brief overview of these problems. Ben-Moshe et al. [BMKS99] show that in $\mathbb{R}^2$ under ℓ_∞, the obnoxious 1-median is solvable in $O(n \log^2 n)$ time and the corresponding decision problem is solvable in $O(n \log n)$ time. They also show that in $\mathbb{R}^2$ under ℓ_2, the obnoxious 1-median problem is solvable in $O(n \operatorname{polylog} n)$ time and the corresponding decision problem is solvable in $O(n \log n)$ time. For additional results in obnoxious facility location within polygonal regions in $\mathbb{R}^2$ under ℓ_2 and ℓ_∞ distance metrics see [MC86a] and [MC86b].

Bereg et al. [BKST99] examine the obnoxious k-centre problem in $\mathbb{R}^2$ under multiple weights (one for each dimension). They examine both the discrete and continuous cases.

Regions, Generalized Distance Metrics, Obstacles, and New Facilities
Either clients, facilities, or both may be modelled as regions in space instead of single points. For example, the facility may be a line and the clients points in $\mathbb{R}^d$. The problem is then a linear regression style problem of positioning a line to minimize the distance from the points to the line [LMW88]. Applications to this problem involve selecting positions for roads, railroads, power lines, sewage pipes, etc. The problem has a natural generalization to locating a hyperplane that minimizes the maximum Euclidean distances to a set of points in $\mathbb{R}^d$ [SS97].

Gao et al. [GLS06] examine the inverse problem for which clients consist of lines in the plane and the facility consists of a point. Given a set of lines in the plane, Gao et al. give an algorithm that identifies the smallest circle (Euclidean 1-centre) that intersects every line.

Defining points or facilities as regions results in a variety of distance metrics being employed, ranging from object-to-object distance (for example, Hausdorff distance) to non-linear distance metrics (involving an additive factor). In general, multitudes of distance metrics are considered under various models for the universe including spherical distance [Pla95] and combinations of ℓ_p distance metrics [HLP$^+$87, LMW88, BL95].

All problems described thus far assumed unobstructed paths between facilities and clients. The introduction of obstacles into a continuous space universe, around which the path from a client to a facility (and its corresponding length) must wind, alters the distance metric and, correspondingly, the optimal positioning of facilities [CSK98].

Another common variation involves selecting positions for a set of new facilities, given the positions of existing facilities [LMW88].

Probabilistic Facility Location
Yet another model of facility location incorporates problems for which the exact position of clients is unknown, but some probability distribution is given on these positions [LMW88, HM89, Sny04]. See [Sny05] and [BJSL95] for reviews of the problems of probabilistic facility location.

Weighted Clients
A common variation of the centre and median involves assigning weights to clients. Since multiplicities of clients are permitted in the k-median problem, when the client weights are rational, any instance of a weighted client set can be reduced to an equivalent unweighted client set by addition the corresponding number of new clients coinciding with the positions of the weighted clients. The definition of the k-centre problem, however, is altered by the addition of weights. The optimization function of Expr. (2.1) becomes

$$\max_{p \in P} w(p)\delta(p, c).$$

For example, if $P = \{0, 1, 2\}$ with corresponding weights $w(0) = 1$, $w(1) = 8$, and $w(2) = 4$, the weighted 1-centre of P lies at $4/3$. In general, weighted

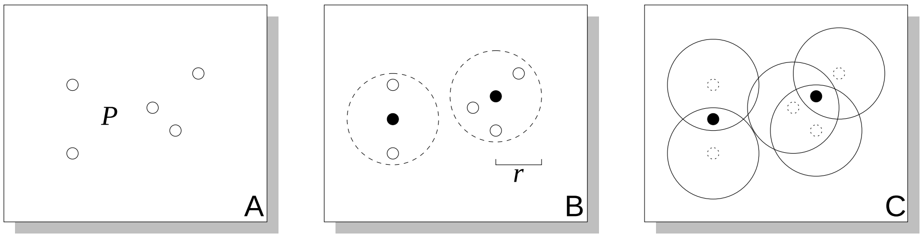

Figure 2.11: **A.** a set of points P in $\mathbb{R}^2$, **B.** a Euclidean 2-centre of P of radius r, **C.** the corresponding 2-piercing for a set of disc of radius r centred at the points of P

problems are more difficult to solve than the corresponding unweighted problem. In this thesis we consider only unweighted problem instances.

Piercing Sets and Covering

Set covering and piercing problems are closely related to the problems of facility location. For a discussion of mobile piercing and set covering problems, see [HRS04, KNS00, Seg99].

Sharir and Welzl [SW96] describe a reduction from the decision problem for a geometric k-centre problem to the k-piercing problem. Given a set P of client positions in $\mathbb{R}^d$ and a positive integer k, the decision problem associated with the Euclidean k-centre involves fixing the Euclidean radius r, and asking whether there exists a set F of k points in $\mathbb{R}^d$ such that the points of P are contained within at least one of k hyperspheres of radius r centred at the points of F. The reduction to the piercing problem is achieved by instead positioning $|P|$ hyperspheres of radius r centred at the points of P, and asking whether there exists a set F of k points in $\mathbb{R}^d$ such that each hypersphere contains at least one point of F. This reduction generalizes to any Minkowski distance metric. See Fig. 2.11.

Chapter 3

Mobile Facility Location

3.1 Introduction

3.1.1 Chapter Objectives

Chapter 2 introduced the Euclidean k-centre and Euclidean k-median problems, two fundamental problems of geometric facility location, as they are traditionally presented in a *static* setting; a problem instance consists of a set of fixed points in Euclidean space, corresponding to client positions, and a problem solution consists of a second set of fixed points in Euclidean space, corresponding to locations for facilities. In this thesis, the Euclidean k-centre and Euclidean k-median are examined in a *mobile* setting with the objective of maintaining bounded-velocity approximations to these. Chapter 3 formalizes concepts pertinent to discussing mobile problems, including maximum velocity and approximation factor. The chapter closes with a discussion of related work.

Chapter 3 establishes the important questions regarding mobile problems in geometric facility location. We identify the relevant open problems which are subsequently addressed in the remainder of the text. Of particular significance is Sec. 3.6 which provides a contextual perspective of the relevance of results developed in subsequent chapters.

3.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Continuous Motion (Sec. 3.2)
Sec. 3.2 introduces continuous motion and definitions for mobile clients and mobile facilities. We consider mobile clients whose positions are defined over a time interval by a continuous function in $\mathbb{R}^d$. The position of a mobile facility is defined as a function of the instantaneous client positions.

Velocity and Continuity (Sec. 3.3)
Sec. 3.3 formalizes the notions of velocity and continuity, two natural properties of mobile problems. By adding these constraints we alter our criteria for defining a good approximation of a facility function. The fitness of a mobile facility is measured in terms of two parameters: approximation factor and maximum velocity (which requires continuity). In this section we discuss maximum velocity, denoted $v_{\max}$.

Approximation Factor (Sec. 3.4)
Given the unbounded velocity of the mobile Euclidean 1-centre and the discontinuity of the mobile Euclidean 1-median, our search for bounded-velocity facility functions leads us to consider approximations to these. Along with maximum velocity, approximation factor, denoted λ, defines the second of two measures by which we compare mobile facility functions. Thus, the fitness of an approximation is measured both by the quality of its optimization of the objective function and also by its maximum velocity and continuity of its motion.

Stability (Sec. 3.5)
We note an inverse relationship between the stability of a facility function, as defined statically, and the maximum velocity of a mobile facility function, providing additional motivation for considering bounded-velocity approximations with implications to problems of static facility location.

Taking Perspective (Sec. 3.6)
Having motivated the importance of identifying approximations to the mobile Euclidean k-centre and k-median and having defined the concepts of maximum velocity, continuity, approximation factor, and stability necessary to evaluating and comparing approximation functions, Sec. 3.6 classifies the problems into those which do not require approximation, those that can be approximated with bounded velocity, and those for which no bounded-velocity approximation is possible. The resulting subdivision of problems corresponds to the organization of Chs. 4 through 7.

Related Work (Sec. 3.7)
Sec. 3.7 provides an overview of recent related work in mobile facility location, including work in discrete mobile facility location and dynamic facility location. Details of some results directly related to the Euclidean 1-centre, Euclidean 1-median, rectilinear 1-centre, rectilinear 1-median, and centre of mass are mentioned briefly here but are described more completely alongside detailed analyses of these concepts in Chs. 4 and 5.

3.2 Continuous Motion

The traditional problems of facility location are set in a *static* setting; client positions are fixed and a single location is selected for each facility. The problems of static facility location have been studied extensively. Within the last few years, partly motivated by the applicability of mobile computing to the wireless telecommunication industries involving cellular and radio ethernet, these questions have been posed in the *mobile* setting [AH01, AGHV01, AGG02, AdBG+05, BBKS00, BBKS06, DK03, DK04, DK05a, DK05b, DK05c, DK06, GGH+03, Her05].

We consider *continuous motion* in $\mathbb{R}^d$. That is, each client's position traces a continuous trajectory through Euclidean space, defined as a function over a continuous temporal dimension. Furthermore, we assume no prior knowledge of future client positions.

Definition 3.1. *Let $T = [0, t_f]$ denote a time interval. Let $P = \{p_1, \ldots, p_n\}$ be a set of* **mobile clients** *such that for every i, $p_i : T \to \mathbb{R}^d$ is a bounded continuous function that defines the position of client i in $\mathbb{R}^d$ at every instant $t \in T$.*

For every $t \in T$, let $P(t) = \{p_i(t) \mid p_i \in P\}$ denote the set of points corresponding to the positions of clients in P at time t.

For every client $p \in P$, the point $p(t)$ is defined at every instant t over the interval T. As such, Def. 3.1 differs significantly from another common notion of mobility, often called *dynamic motion* [Wes73, WT75, APP96, BGKS98, HP98, AHT00]. Although both continuous motion and dynamic motion refer to change in position over a temporal axis, the distinction between these two models of mobility is noted because their often disjoint objectives result in fundamentally different solution techniques. Dynamic motion involves discretized time steps for which the position of a mobile client or facility a is described by a sequence of discrete points in Euclidean space, corresponding to sampling the position of a at regular intervals in time. Hence, constraints of velocity and continuity are typically inconsequential to a solution. Work related to this alternate notion of mobility is discussed in Sec. 3.7.5. We restrict our attention to continuous motion as described by Def. 3.1.

Having defined mobile clients, we augment the definition of a facility function to the mobile setting. Given a set of mobile clients, the location of a mobile facility is specified by a given facility function Υ_d of the client positions.

Definition 3.2. *Let $T = [0, t_f]$ denote a time interval. Given a facility function $\Upsilon : \mathscr{P}(\mathbb{R}^d) \to \mathscr{P}(\mathbb{R}^d)$, the corresponding* **mobile facility function***, $\Upsilon^* : \mathscr{P}(\mathbb{R}^d) \times T \to \mathscr{P}(\mathbb{R}^d)$, is given by*

$$\Upsilon^*(P, t) = \Upsilon(P(t)). \tag{3.1}$$

Thus, the position of a mobile facility function Υ_d at time $t \in T$ corresponds its static definition applied to the set of points $P(t)$ induced by the positions of a set of mobile clients P at time t.

When $k > 1$, the facility function returns a set of single-facility functions. That is, $\Upsilon_d(P(t)) = \{\Upsilon_d^1(P(t)), \ldots, \Upsilon_d^k(P(t))\}$, where $\Upsilon_d^i : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$. For simplicity, we present definitions in terms of single-facility functions throughout this chapter; all definitions given are easily generalizable to multiple-facility functions.

Following Bereg et al. [BBKS00], we define the mobile Euclidean 1-centre, the mobile Euclidean 1-median, the mobile rectilinear 1-centre, the mobile rectilinear 1-median, and the mobile centre of mass as a direct extensions of their respective static definition: $\Xi_d(P(t))$, $M_d(P(t))$, $R_d(P(t))$, $S_d(P(t))$, and $C_d(P(t))$. See Fig. 3.1.

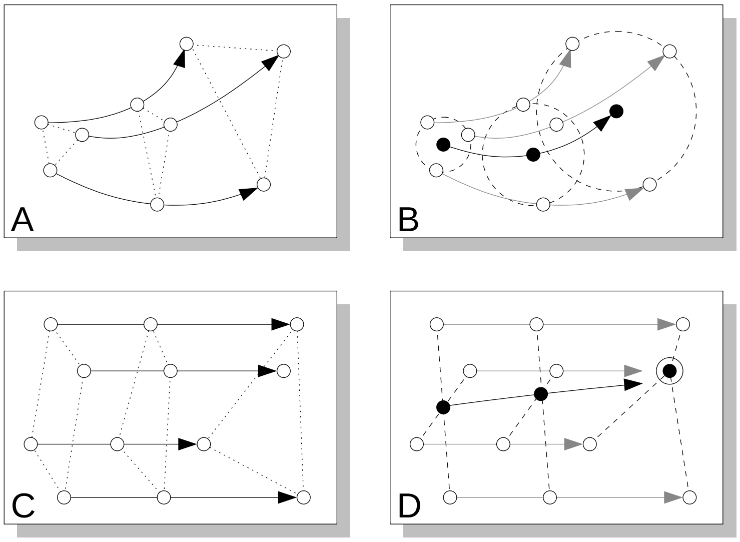

Figure 3.1: **A.** a set of mobile clients in $\mathbb{R}^2$, **B.** the corresponding mobile Euclidean 1-centre, **C.** a second set of mobile clients in $\mathbb{R}^2$, and **D.** the corresponding mobile Euclidean 1-median

3.3 Velocity and Continuity

Continuity of motion and a finite upper bound on velocity[1] impose natural constraints on any physical moving object. Scenarios involving vehicles, mobile robots, or people with wireless communication devices suggest that bounds on continuity and velocity are necessary in many applications [AOY99, KNW02, CFPS03, Sch03, CMKB04, CMB06]. Thus, the fitness of a mobile facility is determined not only by the quality of its optimization of the objective function but also by the maximum velocity and continuity of its motion.

We consider clients whose motion is continuous[2] and we assume that their velocity is bounded by a constant $\sigma > 0$. That is,

$$\forall p_i \in P,\ \forall t_1, t_2 \in T,\ ||p_i(t_1) - p_i(t_2)|| \leq \sigma \cdot |t_1 - t_2|. \tag{3.2}$$

When p_i is differentiable, then $\forall t \in T,\ ||p_i'(t)|| \leq \sigma$. Throughout this thesis we assume a constant upper bound of $\sigma = 1$ on the velocity of clients since we are interested in *relative velocity*.

Continuity is a necessary condition for bounded velocity.

[1] We use the term *bounded velocity* to mean *bounded magnitude of velocity*.

[2] We restrict our attention to motion that is *temporally continuous*; that is, given a mobile client or facility $a : T \to \mathbb{R}^d$, we require that $a(t)$ be defined for every $t \in T$, where $T = [0, t_f]$ for some $t_f > 0$. We use the term *continuous* in reference to a mobile client or facility that is *spatially continuous* (see Def. 3.3).

Definition 3.3. *Mobile facility function* $\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ *is* **continuous** *if for any time interval* $T = [0, t_f]$ *and any set of mobile clients* P *defined over* T,

$$\forall t_0 \in (0, t_f),\ \forall \epsilon > 0,\ \exists \delta > 0,\ \forall t \in (t_0 - \delta, t_0 + \delta),\ ||\Upsilon_d(P(t_0)) - \Upsilon_d(P(t))|| < \epsilon. \tag{3.3}$$

That is, a mobile facility whose motion is continuous follows a continuous trajectory through Euclidean space.

As will be shown in Chs. 4 through 6, although mobile clients are limited to at most unit velocity, the maximum velocity of a mobile facility must sometimes exceed unit velocity to guarantee a good approximation of the objective function. Maintaining a low upper bound on the relative velocity of a mobile facility function remains a primary objective, which we now define:

Definition 3.4. *Let* P *be a set of mobile clients. Let* $\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ *be a mobile facility function. The* **maximum velocity** *of a mobile facility whose location is determined by* Υ_d *is bounded by* $v_{\max}$ *if*

$$\forall t_1, t_2 \in T,\ ||\Upsilon_d(P(t_1)) - \Upsilon_d(P(t_2))|| \leq v_{\max}|t_1 - t_2|. \tag{3.4}$$

We say $v_{\max}$ is **tight** if $v_{\max}$ is the infimum over all $v'_{\max}$, where the maximum velocity of Υ_d is bounded by $v'_{\max}$ as defined in Eq. (3.4). Equivalently, velocity $v_{\max}$ is realizable; that is, there exists a set of mobile clients P in $\mathbb{R}^d$ defined over a time interval T such that $\Upsilon_d(P(t))$ moves with (instantaneous) velocity $v_{\max}$ at some instant $t \in T$.

For some mobile facility functions, even when clients are limited to unit velocity, no finite upper bound on velocity exists; that is, $v_{\max} = \infty$. For example, it is not possible to bound the velocity of the Euclidean 1-centre by any fixed constant $v_{\max}$. Specifically, for any $v_{\max} \geq 0$, Bereg et al. [BBKS00] construct an example of a set of mobile clients in $\mathbb{R}^2$, each moving in a linear trajectory with unit velocity, such that the Euclidean 1-centre moves with average velocity at least $v_{\max}$ over some time interval T, $|T| = \delta > 0$, where δ depends on $v_{\max}$ (see Sec. 4.2). Consequently, given any $v_{\max}$, Eq. (3.4) does not hold for the Euclidean 1-centre. Similarly, the velocity of the Euclidean 1-median is also unbounded, as is straightforward to demonstrate by a set of four mobile clients in $\mathbb{R}^2$ (see Sec. 5.2).

Although the velocity of the mobile Euclidean 1-centre is unbounded, its motion is continuous. The motion of the mobile Euclidean 1-median, however, is discontinuous. Again, this property is straightforward to demonstrate by a set of four mobile clients in $\mathbb{R}^2$ (see Sec. 5.2).

Empirical evidence suggests that these examples resulting in unbounded velocity or discontinuity of the Euclidean k-centre or Euclidean k-median are easily realized by a small number of mobile clients, for example, with as few as four clients moving at unit velocity along random linear trajectories inside the unit square on the plane. See Sec. 8.4.2.

3.4 Approximation Factor

Given their unbounded velocities, the Euclidean 1-centre and the Euclidean 1-median may be unfit for certain applications and impossible to maintain exactly within specific mobile contexts. A function that approximates a facility function while maintaining some fixed upper bound on its maximum velocity may be better suited. We refer to such a function as an **approximation function**.

As discussed in Ch. 2, approximation algorithms are commonly used to solve static Euclidean k-centre and k-median problems. Unlike the static setting, where the fitness of an approximation strategy is determined solely by the quality of its optimization of the objective function, in the mobile setting the fitness of an approximation strategy is also determined by the maximum velocity and continuity of its motion. As discussed in Ch. 4 and 5, these dual objectives cannot both be simultaneously satisfied optimally, leading to the development of new approximation strategies quite different from previous static approximations.

Let $\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ denote an arbitrary approximation function. Within the context of the Euclidean 1-centre, we refer to Υ_d as a **centre function** and measure the quality of Υ_d's approximation of Ξ_d in terms of their relative values of the optimization function. Within the context of the Euclidean 1-median, we refer to Υ_d as a **median function** and measure the quality of Υ_d's approximation of M_d, also in terms of their relative values of the optimization function.

If Υ_d is a centre function with maximum relative velocity bounded by $v_{\max}$, then $\max_{p\in P} ||p - \Upsilon_d(P)||$ must exceed the Euclidean radius of P for some $P \in \mathscr{P}(\mathbb{R}^d)$. That is, for some $P \in \mathscr{P}(\mathbb{R}^d)$, the ratio of the values of the optimization function for Υ_d and Ξ_d must exceed one. Similarly, if Υ_d is a median function with maximum relative velocity bounded by $v_{\max}$, then $\sum_{p\in P} ||p - \Upsilon_d(P)||$ must exceed the Euclidean median sum of P for some $P \in \mathscr{P}(\mathbb{R}^d)$. We formalize the notion of the relative value of the optimization function in terms of the *approximation factor* of Υ_d.

Definition 3.5. *Given an optimization function $g : \mathscr{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$, a facility function $F_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ that optimizes g, and an approximation function $\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$, Υ_d is a* **$\boldsymbol{\lambda}$-approximation** *of F_d if*

$$\forall P \in \mathscr{P}(\mathbb{R}^d),\ g(P, \Upsilon_d(P)) \le \lambda g(P, F_d(P)). \tag{3.5}$$

Specifically, when F_d is the Euclidean 1-centre, Ξ_d, Eq. (3.5) becomes

$$\forall P \in \mathscr{P}(\mathbb{R}^d),\ \max_{p\in P} ||p - \Upsilon_d(P)|| \le \lambda \max_{q\in P} ||q - \Xi_d(P)||. \tag{3.6a}$$

Within the context of centre functions, we refer to the approximation factor as **eccentricity**. Similarly, when F_d is the Euclidean 1-median, M_d, Eq. (3.5) gives

$$\forall P \in \mathscr{P}(\mathbb{R}^d),\ \sum_{p\in P} ||p - \Upsilon_d(P)|| \le \lambda \sum_{q\in P} ||q - M_d(P)||. \tag{3.6b}$$

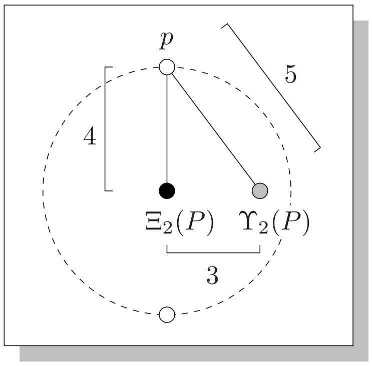

Figure 3.2: In this example, the Euclidean radius is four and the distance between $\Upsilon_2(P)$ and client p is five. Consequently, the approximation factor of Υ_2 (over all possible sets of clients) is at least $5/4$.

We say λ is **tight** if λ is the infimum over all λ', where the approximation factor of Υ_d is bounded by λ' as defined in Eq. (3.5). Equivalently, the approximation factor λ is realizable; that is, there exists a set of clients P in $\mathbb{R}^d$ such that $g(P, \Upsilon_d(P)) = \lambda g(P, F_d(P))$.

Since Ξ_d and M_d optimize their respective objective functions, the approximation factor λ ranges from 1 to ∞, with a 1-approximation function being the best approximation. Any such approximation function necessarily has maximum velocity at least one; otherwise, any parallel translation of the clients at unit velocity leads to an unbounded approximation factor because the mobile facility is unable to keep up with the client set. Thus, we consider approximation functions whose maximum velocity is in the range $[1, \infty)$ and whose approximation factor is also in the range $[1, \infty)$, and strive to attain values close to one for both properties.

Although it may also seem natural instead to define approximation as a function of relative proximity to the exact position of the Euclidean k-centre or k-median, our measure of approximation defined in terms of the objective function allows for a bounded-velocity approximation function whose position can lie relatively far away from the exact position of the Euclidean k-centre or k-median while still providing a good approximation of the objective function. This consideration is particularly relevant when the mobile facility being approximated moves arbitrarily quickly (e.g., the Euclidean 1-centre), when its position changes discontinuously (e.g., the Euclidean 1-median), when its position is not uniquely defined (e.g., the Euclidean 2-centre), or when its position in unknown (e.g., the Euclidean 1-median). Approximations to the Euclidean k-centre and k-median are typically defined in accordance with our definition of approximation factor (e.g., [AP98, ARR98, BMM03, Ind99]).

Since approximation factor is defined in terms of a worst-case configuration, it is independent of motion of points. Thus, the approximation factor of a mobile facility whose position is defined by approximation function Υ_d is simply the approximation factor of Υ_d on a static set of clients.

In addition to being defined by lower maximum velocity and a lower approximation factor, natural properties of a "better" approximation function may also include invariance under similarity transformations and consistency of definition across dimensions. Secs. 4.3 and 5.3 discuss these properties in details specific to the contexts of centre functions and median functions, respectively.

In summary, the maximum velocity and approximation factors, $v_{\max}$ and λ, allow us to compare the utility of different approximation functions. In general, functions with lower approximation factors have a higher maximum velocity, and vice-versa. Subject to this trade-off, we seek approximation functions with low maximum velocity and low approximation factor.

3.5 Stability

The notion of stability is directly related to maximum velocity.

Point coordinates are commonly represented by discretization of real positions to nearby grid coordinates. That is, each point is approximated by the nearest grid point. Given a finite set of points $P \in \mathscr{P}(\mathbb{R}^d)$ and its Euclidean 1-centre $\Xi_d(P)$, small perturbations at only a few points of P can result in a relatively large change (error) in the corresponding position of $\Xi_d(P)$ [Dre95, BBKS00]. The same is true of the Euclidean 1-median $M_d(P)$. In this sense, both the Euclidean 1-centre and the Euclidean 1-median are *unstable*.

We formalize the notion of stability by defining κ-stability for an approximation function Υ_d as a measure of its maximum volatility. This requires preliminary definitions for an ϵ-perturbation.

Definition 3.6. *Given $\epsilon > 0$, function $f : P \to \mathbb{R}^d$ is an* **ϵ-perturbation** *on $P \in \mathscr{P}(\mathbb{R}^d)$ if for all $p \in P$, $||p - f(p)|| \leq \epsilon$.*

Let F_ϵ^P denote the set of all ϵ-perturbations on P.

Definition 3.7. *A function $\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ is* **κ-stable** *if*

$$\forall \epsilon > 0,\ \forall f \in F_\epsilon^P,\ \kappa||\Upsilon_d(P) - \Upsilon_d(f(P))|| \leq \epsilon, \tag{3.7}$$

for all $P \in \mathscr{P}(\mathbb{R}^d)$.

The similarity of the definitions of stability (Def. 3.7) and maximum velocity (Def. 3.4) is perhaps not surprising; the maximum velocity, $v_{\max}$, and the stability, κ, of a mobile facility function are inversely related. Maximum velocity and approximation factor describe the fitness of a mobile approximation function's approximation of a mobile facility function just as stability and approximation factor describe the fitness of an approximation function's approximation to a static facility function.

Observation 3.1. *$\Upsilon_d : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ is a κ-stable centre function if and only if a mobile facility whose position is determined by Υ_d has maximum velocity bounded by*

$$v_{\max} = \frac{1}{\kappa}. \tag{3.8}$$

Proof. Let $\epsilon = |t_1 - t_2|$, $Q = P(t_1)$, and $f(Q) = P(t_2)$. It follows that

$$||\Upsilon_d(P(t_1)) - \Upsilon_d(P(t_2))|| \leq v_{\max}|t_1 - t_2| \Leftrightarrow \kappa||\Upsilon_d(Q) - \Upsilon_d(f(Q))|| \leq \epsilon, \quad (3.9)$$

where $\kappa = 1/v_{\max}$. □

An immediate consequence of Obs. 3.1 is that neither the Euclidean 1-centre nor the Euclidean 1-median is κ-stable for any $\kappa > 0$; that is, Ξ_d and M_d are 0-stable. This observation provides further motivation for the identification of approximation functions that achieves both low maximum velocity and a low approximation factor.

3.6 Taking Perspective: When is Bounded-Velocity Approximation Possible, Necessary, or Impossible?

Presumably, by now our motivation has convinced the reader of the value in seeking mobile facility functions to approximate the mobile Euclidean k-centre and the mobile Euclidean k-median in d-dimensional Euclidean space. We now examine which specific values of k and d define mobile problems for which bounded-velocity approximation is both necessary and possible.

A mobile problem is well-behaved if it always has at least one solution that moves continuously with a fixed upper bound on its velocity. Depending on the number of facilities, k, and the dimension, d, of the problem space $\mathbb{R}^d$, either a) a mobile Euclidean k-centre or k-median problem is well-behaved and no approximation is necessary, b) the problem is ill-behaved but can be approximated with bounded velocity, or c) the problem is ill-behaved and no bounded-velocity approximation is possible. We consider problems that corresponds to the second set of conditions, for which the velocity of the Euclidean k-centre or k-median is unbounded, but for which a bounded-velocity approximation is possible. Tab. 3.1 summarizes these conditions for specific values of d and k, with the corresponding cases annotated by a star ($\star$). Observe that if a Euclidean k-centre problem does not have bounded-velocity approximation in $\mathbb{R}^d$ for some d, then no bounded-velocity approximation exists in any higher dimension (the same is true for the Euclidean k-median). Similarly, if a specific Euclidean k-centre problem cannot be approximated, then the analogous j-centre problem cannot be approximated for any $j \geq k$ (again, the same is true for the Euclidean k-median). These claims and the classification implied by Tab. 3.1 are formalized in Chs. 4 through 7.

Tab. 3.1 displays an overview of the mobile Euclidean k-centre and k-median problems in $\mathbb{R}^d$. Three sets of parameters result in problems whose velocity is unbounded and for which bounded-velocity approximation is possible ($\star$):

1. Euclidean 1-centre in $\mathbb{R}^d$ for $d \geq 2$,
2. Euclidean 2-centre in $\mathbb{R}^d$ for $d \geq 2$, and

k-centre in $\mathbb{R}^d$	$d = 1$	$d \geq 2$
$k = 1$	continuous, $v_{\max} = 1$	continuous, $v_{\max} = \infty$, bounded-velocity approximable ($\star$)
$k = 2$	continuous, $v_{\max} = 2$	discontinuous, bounded-velocity approximable ($\star$)
$k \geq 3$	discontinuous, no bounded-velocity approximation	

k-median in $\mathbb{R}^d$	$d = 1$	$d \geq 2$
$k = 1$	continuous, $v_{\max} = 1$	discontinuous, bounded-velocity approximable ($\star$)
$k \geq 2$	discontinuous, no bounded-velocity approximation	

Table 3.1: taking perspective: feasibility of bounded-velocity approximation for the mobile Euclidean k-centre and k-median problems

3. Euclidean 1-median in $\mathbb{R}^d$ for $d \geq 2$.

In each of these cases, although the motion of the exact k-centres or k-medians has unbounded velocity or is discontinuous, a bounded-velocity approximation is still possible. For each case we develop, analyze, and compare possible approximation strategies in Chs. 4 through 6.

The sets of parameters for which no bounded-velocity approximation exists can be reduced to two cases:

1. Euclidean k-centre in $\mathbb{R}^d$ for any $k \geq 3$ and any d, and
2. Euclidean k-median in $\mathbb{R}^d$ for any $k \geq 2$ and any d.

These cases are examined in Ch. 7, in which we prove the infeasibility of approximation by k mobile facilities and examine whether these problems can be approximated by greater than k mobile facilities.

In summary, a primary objective of this thesis is to identify and analyze bounded-velocity strategies for approximating the mobile 1-centre, 2-centre, and 1-median problems in Euclidean space.

3.7 Related Work in Mobile Facility Location

Similarly to Sec. 2.5, which addresses related work in static facility location, Sec. 3.7 provides an overview of related work in mobile facility location.

3.7.1 Data Structures for Mobile Data

We examine four classes of data structures developed to maintain one or more properties of a set of mobile clients. In particular, we make use of kinetic data structures in our algorithm for maintaining the mobile Steiner centre described in Ch. 8.

Early Work in Dynamic Computational Geometry

Atallah [Ata85] considers problems for which a set of mobile clients moves continuously over time. He examines the combinatorial complexity of maintaining the position of the minimum client in $\mathbb{R}$ and maintaining the two-dimensional convex hull, when the motion of the clients is linear or bounded-degree algebraic. Specifically, the number of times a client can join or leave the boundary is calculated as a function of the cardinality of the client set and the algebraic degree of the motion. Also examined is the steady-state relative configuration of client positions after all combinatorial change events have terminated as time approaches infinity.

Kinetic Data Structures

Kinetic data structures (KDS) introduced by Basch et al. [BGSZ97, Gui98, BGH99, Bas99b] allow for efficient implementation and maintenance of various attributes of a set of mobile clients under piecewise-linear (or bounded-degree algebraic) motion. In brief, a KDS algorithm maintains a set of certificates that validate a specific property of a set of mobile clients. Each certificate corresponds to a simple geometric assertion (for example, "client p_1 lies below the line induced by clients p_2 and p_3"). The constraint on the degree of the motion of the client set allows for the occurrence of change events related to the trajectories of mobile clients to be calculated exactly. Whenever a client decides to change its trajectory, a flight update event is submitted. The list of certificates is updated accordingly whenever either a flight update or certificate failure event occurs. Basch et al. describe criteria by which a KDS is evaluated:

> For a KDS to be of good quality, the following criteria should be met:
>
> - the certificate list does not change too much when an event occurs (*responsiveness*);
> - the overhead of internal events with respect to external events is reasonable (*efficiency*);
> - the KDS itself is of small size, typically linear or slightly superlinear (*compactness*); and
> - each client is involved in only a small number of certificates (*locality*). [BGSZ97, p. 388]

Several KDS algorithms are related to our work on mobile facilities, some of which are employed in our implementations described in Ch. 8. These include the bounding box [AH01], the two-dimensional convex hull [BGSZ97, BGH99, Gui98], a $(1+\epsilon)$-approximate Euclidean 1-centre [AH01], and the extent of a set of mobile clients in $\mathbb{R}$ [AH01, Gui98, BGH99].

Real-Time Kinetic Algorithms

Recent work by Uthaisombut [Uth05a, Uth05b] suggests a compromise between dynamic facility location, in which client position are reported and facility locations are calculated at discrete time steps (see Sec. 3.7.5), and mobile facility location, in which the positions of clients and facilities are defined at all points over a continuous time interval.

Uthaisombut proposes that the time required by an algorithm for online computation should be factored into the event framework. In his model, clients move continuously and under bounded velocity, but without additional constraints on the complexity of the motion. The positions of clients are sampled at regular intervals, where the frequency of sampling is an input parameter whose value affects the maximum error between samples. As an example, Uthaisombut describes maintaining the order of a set of mobile clients in $\mathbb{R}$, such that a user may query the data structure to determine the kth largest element at any time. A naïve algorithm simply sorts the points every time step, requiring that the time interval be at least $\Theta(n \log n)$ in duration. Naturally, between time steps, a query will result in some degree of error, as a function of the proximity and maximum velocity of clients. Uthaisombut proposes an algorithm that allows sampling at regular time intervals of duration $O(n)$, resulting in a reduction in the magnitude of possible error and allowing for the necessary maintenance of the data structure to occur within the alloted time.

The primary difference between real-time kinetic algorithms and algorithms involving discretized time incrementation, is that the time interval T remains continuous. Although the algorithm samples client positions at fixed points in T, error is calculated not only for client configurations at sample times, but rather it is defined as the maximum error occurring at any point in T, including the intervals between sample times.

Incremental Motion

Motivated to generalize kinetic data structure to allow unconstrained motion, recent work of Mount et al. [MNP+04] provides data structures for incremental motion. Unlike a KDS, the motion of clients is incremental. Like a KDS, this algorithm relies on a set of certificates to validate some property of a set of mobile clients. A lower level algorithm provides estimates of future client trajectories and projected certificate failures. Whenever a client's trajectory deviates from its predicted course, the data structure is updated as necessary to validate certificates.

3.7.2 Mobile Euclidean k-Centre and k-Median

Perhaps the most significant work related to the mobile Euclidean k-centre and the mobile Euclidean k-median is the work of Bereg et al. [BBKS00, BBKS02, BBKS06] who first raised some of the fundamental questions relating to bounding the velocity of a mobile facility function. They show the velocity of the mobile Euclidean 1-centre is unbounded in $\mathbb{R}^2$. Bereg et al. consider the the mobile centre of mass, the mobile rectilinear 1-centre, and the mobile rectilinear 1-median as approximations of the Euclidean 1-centre and 1-median and examine bounds on their respective approximation factors and maximum velocities. The properties of each of these mobile facility functions are examined in greater detail in Chs. 4 and 5, in which we refer to [BBKS00, BBKS02, BBKS06].

Agarwal and Har-Peled [AH01] maintain an approximation to the mobile Euclidean 1-centre in $\mathbb{R}^2$. Their approximation does not require continuity or bounded velocity in the motion of the centre function; their objective, rather, is to minimize the number of events processed and the update cost per event using a KDS to maintain a $(1+\epsilon)$-approximation on the extent of the point set. Agarwal et al. [AGG02] use a KDS to maintain a kd-tree of the points and an ϵ-approximate mobile median in $\mathbb{R}$. In [AdBG+05], a KDS maintains the exact (expensive) and ϵ-approximate (less expensive) mobile Euclidean 1-median in $\mathbb{R}$ and $\mathbb{R}^2$.

As discussed in Secs. 4.3.2 and 5.3.2, no bounded-velocity approximation function can guarantee an approximation factor of λ to either the Euclidean 1-centre or the Euclidean 1-median for an arbitrary λ and a fixed maximum velocity $v_{\max}$ that is independent of λ. Thus, although the algorithms of Agarwal and Har-Peled [AH01] and Agarwal et al. [AdBG+05] provide excellent approximations of the mobile Euclidean 1-centre and 1-median, their maximum velocities cannot be bounded independently of the approximation factor.

3.7.3 Mobile Rectilinear k-Centre and k-Median

A natural question might be to ask why the rectilinear k-centre and k-median are not included in the set of geometric facility location problems listed in Tab. 3.1. Firstly, the distance metric selected should be applied uniformly within the optimization function (to measure distances between clients and facilities) and to describe velocity (defined as a rate of change in distance over time). Although velocity and distance are most naturally described using Euclidean distance, one may still wish to examine the mobile problems of geometric facility location under rectilinear or Chebyshev distance metrics.

As mentioned in Sec. 3.7.2, Bereg et al. [BBKS00] show the mobile rectilinear 1-centre and 1-median move with bounded velocity (under any Minkowski distance metric). Thus, no approximation is necessary. Detailed discussions of the rectilinear 1-centre and 1-median can be found in Chs. 4 and 5.

When $k \geq 2$, no bounded-velocity approximation exists for any geometric k-median in $\mathbb{R}^d$ for any d. Similarly, when $k \geq 3$, no bounded-velocity approximation exists for any geometric k-centre in $\mathbb{R}^d$ for any d. See Ch. 7 for a discussion of bounded-velocity approximations of the rectilinear 2-centre. This result implies that no bounded-velocity approximation exists for the rectilinear 2-median or for the rectilinear 3-centre in any dimension.

Finally, just like the Euclidean 2-centre, the rectilinear 2-centre exhibits discontinuous motion in $\mathbb{R}^2$. The same strategies used to approximate the mobile Euclidean 2-centre apply in this case. See Ch. 6.

Agarwal et al. [AGHV01] maintain the exact mobile rectilinear 1-centre in $\mathbb{R}^2$ under Chebyshev distance using a KDS. Agarwal and Har-Peled [AH01] maintain an approximation to the mobile rectilinear 1-centre in $\mathbb{R}^2$. Given a set P of mobile clients in $\mathbb{R}^d$, Hershberger [Her05] introduces a new KDS for maintaining a set of unit hypercubes that cover P. The number of boxes is within 3^d of the optimal value. These boxes, however, are not constrained to move continuously. Furthermore, the number of boxes is not constant; boxes are added or removed as the clients move.

3.7.4 Mobile Discrete Facility Location

By their nature, problems of discrete facility location serve different objectives than do problems in a continuous space (see Sec. 2.5.4). Since each facility is restricted to having a position that coincides with that of a client, while a facility follows a particular client, its relative velocity is at most one. However, the facility is obligated to change positions instantaneously from one client to another, resulting in discontinuities in its motion.

Given a fixed w and a set of mobile clients P in $\mathbb{R}^d$, Gao et al. [GGH+03] provide a KDS for maintaining a set of discrete k-centre of P such that every client is contained within a d-dimensional hypercube of width w whose centre is a client in P. The number of hypercubes k is within a constant factor of the minimum value for k.

3.7.5 Dynamic Facility Location and Discretized Time

Until recently, only discrete changes to the location of clients have been considered. These problems, termed **dynamic facility location** [Wes73, WT75], attempt to optimize the objective function summed over a finite set of discrete time slots, $T = \{t_1, \ldots, t_f\}$. This model does not incorporate continuity or bounded-velocity constraints in the motion of the facility. Thus, the techniques employed to solve dynamic facility location problems do not necessarily extend to their counterparts involving continuous motion.

Bhatia et al. [BGKS98] examine dynamic facility location on graphs, where edges are assigned two weights corresponding to two times t_1 and t_2. They problem consists of selecting k vertices as facilities such that the maximum graph distance from any vertex to the nearest facility is minimized over t_1 and t_2. Thus, the facilities are not mobile, but distances between clients may change. The motivating example for this model is to identify locations for facilities to serve a road network at both rush-hour and lower traffic times. Bhatia et al. give a constant factor approximation for the problem and show that no constant factor-approximation is possible for greater than two time slots.

Related to the work of Bhatia et al., Hochbaum and Pathria [HP98] examine the dynamic 2-centre on graphs. They provide a polynomial-time 3-approximation algorithms for the k-centre problem on a graph for which two time slots are given. They go on to show that the problem is NP-hard when three or more time slots are given. See Alstrup et al. [AHT00] and Auletta et al. [APP96] for a discussion of dynamic facility location on trees. See Johansson and Carr-Motyčková [JCM03] for a discussion of the mobile discrete k-clustering on ad-hoc networks.

Suzuki and Okabe [SO95] consider problems of dynamic facility location for which facilities are only available for specific time slots. For example, for some set of time slots, k facilities are to be located while for the remaining time slots, an additional j facilities are available.

Also considered by Suzuki and Okabe is a problem which they refer to as mobile facility location, where clients are static but the facility is mobile. The mobile facility has some maximum distance d it may move over the time interval. Thus, the problem is similar to a static problem of locating k facilities with the additional constraint that there must exist a path of length at most d connecting the k facilities.

Finally, we mention the work of Har-Peled [HP04], who examines the problem of finding positions for a set of static facilities for a given set of trajectories for mobile clients. Under this model, Har-Peled provides a 2-approximation algorithm to the k-centre in $\mathbb{R}^d$ in linear time.

3.7.6 Applications

In addition to the theoretical interest of generalizing the static problems of facility location to the mobile realm, problems of mobile facility location are motivated by a broad set of applications across a variety of fields, ranging everywhere from statistics to economics to robotics to telecommunications. We briefly list some of these applications in this section, first generally for problems of geometric facility location, and then specifically for the problems addressed in this thesis.

According to Nielsen and Nock [NN04], the Euclidean 1-centre finds applications in computer graphics, machine learning, and metrology. Classical applications for both 1-centre and 1-median problems in facility location include identifying sites for mobile emergency services, bus stops, or hospital sites [MS02]. Closely related is the common problem of selecting a location for a hub, be it for an airline, a rapid transit provider, a postal network, or a freight carrier [CEK02].

Telecommunications and network configuration are traditional applications for facility location on graphs [GLY02]. The advent of wireless telecommunications and wireless ethernet defines a new set of mobile facility location problems, introduced in this chapter. As stated by Karch et al. [KNW02], these mobile problems have applications in robotics including maintaining oil platforms, exploring Mars, disarming bombs, cleaning, and moving hazardous substances. To this list, Agarwal et al. [AGG02] add air-traffic control, mobile communication, navigation systems, and geographic information systems. To motivate the mobile Euclidean 1-centre and 1-median, Bereg et al. [BBKS00] suggest the problem of locating a mobile utility within a factory such as a welding robot in a manufacturing plant. Cortés et al. [CMKB04, CMB06] cite applications involving control and coordination for groups of autonomous vehicles. The framework discussed by Cortés et al. is developed for d-dimensional Euclidean space for an arbitrary d. Ando et al. [AOY99] and Cortés et al. [CMB06] suggest algorithms for identifying the Euclidean 1-centre of a set of mobile autonomous agents; their objective is to define a point of convergence on which the agents uniformally agree. These ideas are related to the work of Cieliebak et al. [CFPS03] and Schlude [Sch03] who suggest the Euclidean 1-median as a point of convergence for mobile robots that remains constant as the clients converge toward it.

Recent developments in mobile computing, and more specifically ad-hoc networks, presents further applications for the techniques in mobile facility location [GT95, Sha96, Bas99a, CWLG97, HRS04]. Huang et al. [HRS04] examine ad-hoc networks where each client's range of communication is modelled by a unit disc; as a solution to this problem, they give approximation algorithms for the mobile piercing set problem. Gao et al. [GGH+03] mention the applicability of mobile centre solutions to the fields of mobile computing, specifically within ad-hoc networks. Additional related problems from the networks community are described in the work of Gerla and Tsai [GT95], Sharony [Sha96], Basagni [Bas99a], and Chiang et al. [CWLG97].

Additional potential applications include positioning tow trucks to serve a fleet of taxis or buses, positioning police cruisers to assist patrolling officers on foot, positioning a helicopter to oversee a rescue operation, and positioning coast guard ships within proximity of a fleet of freight ships. Cortés et al. [CMKB04] mention an oceanographic sampling network, in which a series of underwater robots communicate via a local acoustic network. Closely-related are applications of mobile problems in discrete facility location, including the work of Wang and Olariu [WO04] on cluster maintenance in mobile ad-hoc wireless networks.

3.7.7 Other Related Questions

Given a set P of mobile clients in $\mathbb{R}$ or $\mathbb{R}^2$, Agarwal et al. [AdBG+05] examine exact and approximate KDS algorithms for maintaining a mobile facility function Υ such that any line passing through the point $\Upsilon(P(t))$ has at least $2|P|/3$ clients on either side of it at all times. Agarwal et al. observe the same trade-off between quality of approximation and stability of a mobile facility function.

Huang et al. [HRS04] provide approximation algorithms for the mobile piercing problem. Given a set of unit disks whose centres move continuously, Huang et al. provide a 7-approximate solution in $\mathbb{R}^2$ and a 21-approximate solution in $\mathbb{R}^3$. Given a set of unit hypercubes whose centres move continuously in $\mathbb{R}^d$, Huang et al. provide a 2^d-approximate solution. See Sec. 2.5.4 for a description of piercing problems and their relevance to the k-centre problem.

Chapter 4

Mobile Euclidean 1-Centre

4.1 Introduction

4.1.1 Chapter Objectives

The previous chapter provides us with tools for evaluating bounded-velocity approximations the mobile Euclidean k-centre and k-median problems. In Chapter 4, we address the first of these: the Euclidean 1-centre. Our exploration of approximation functions (referred to as *centre functions* in the context of the Euclidean 1-centre) leads us to consider the centre of mass, the rectilinear 1-centre, the Steiner centre, and convex combinations of these, for which we examine the maximum velocity and approximation factor (referred to as *eccentricity* in the context of the Euclidean 1-centre). Kinetic algorithms for maintaining these various mobile centre functions are discussed in Ch. 8; for now we focus on their respective qualities as approximation functions.

Although previously defined, the notion of a Steiner centre had not been analyzed in terms of its approximation of the Euclidean 1-centre nor had it been considered with respect to a set of mobile clients. Exploiting the equivalence of the two definitions of the Steiner centre established by Shephard [She66], we show the Steiner centre successfully balances the conflicting goals of closeness of approximation and low maximum velocity. Summaries of the chapter's significant results and their implications are found in Secs. 4.1.2 and 4.9.

4.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Properties of the Mobile Euclidean 1-Centre (Sec. 4.2)
Sec. 4.2 briefly examines additional properties of the mobile Euclidean 1-centre, Ξ_d. Specifically, we show that the motion of the mobile Euclidean 1-centre is continuous and we quote a theorem of Bereg et al. [BBKS00, BBKS06] proving that the velocity of the mobile Euclidean 1-centre is unbounded.

Comparison Measures (Sec. 4.3)
Building on work of Bereg et al. [BBKS00, BBKS06], Sec. 4.3 expands on the measures of eccentricity and maximum velocity and explores bounds on their relationship in terms specific to the approximation of the Euclidean 1-centre. Additional natural properties of centrality are also considered.

Rectilinear 1-Centre (Sec. 4.4)
Sec. 4.4 analyzes the properties of the mobile rectilinear 1-centre, R_d, in terms of its approximation of the Euclidean 1-centre. The rectilinear 1-centre minimizes the maximum Chebyshev (ℓ_∞) distance between itself and any client in P, suggesting it as a candidate for approximating the Euclidean 1-centre. In particular, we show that in $\mathbb{R}^d$, the rectilinear 1-centre has eccentricity $\frac{1}{2}(1+\sqrt{d})$ and we refer to a result of Bereg et al. [BBKS06] showing that its maximum velocity is $\sqrt{d}$.

Centre of Mass (Sec. 4.5)
Sec. 4.5 analyzes the properties of the mobile centre of mass, C_d, in terms of its approximation of the Euclidean 1-centre. In particular, we refer to results of Bereg et al. [BBKS06] showing that in $\mathbb{R}^d$, the centre of mass has eccentricity 2 and maximum velocity 1.

Steiner Centre (Sec. 4.6)
Sec. 4.6 presents two definitions of Steiner centre, Γ_d, first by Gaussian weights and then by projection. The core of Ch. 4 consists of the derivations of the eccentricity and maximum velocity of the Steiner centre in two and three dimensions contained in this section. In particular, we show that in $\mathbb{R}^2$, the Steiner centre has eccentricity approximately 1.1153 and maximum velocity $4/\pi$. In $\mathbb{R}^3$, we provide a lower bound of approximately 1.2017 on the eccentricity of the Steiner centre and show that its maximum velocity is $3/2$.

Triangle Centres (Sec. 4.7)
Sec. 4.7 briefly explores additional common functions that might initially suggest themselves as candidate centre functions but upon examination exhibit either discontinuity, high eccentricity, or inability to generalize to greater than three clients, making them poor centre functions.

Convex Combinations (Sec. 4.8)
Sec. 4.8 examines convex combinations of centre functions. In particular, a convex combination of a set of centre functions defines a new centre function whose maximum velocity and approximation factor can be bounded in terms of the maximum velocities and approximation factors of the component centre functions.

Evaluation (Sec. 4.9)
Sec. 4.9 summarizes the results derived in Ch. 4 by comparison of the Steiner centre, the rectilinear 1-centre, the centre of mass, and convex combinations of these in terms of their approximation of the Euclidean 1-centre. The primary measures for evaluating the quality of each centre function are eccentricity and maximum velocity (inversely related to stability) but also include consideration of whether each centre function generalizes to higher dimensions and whether it preserves various properties of invariance and consistency.

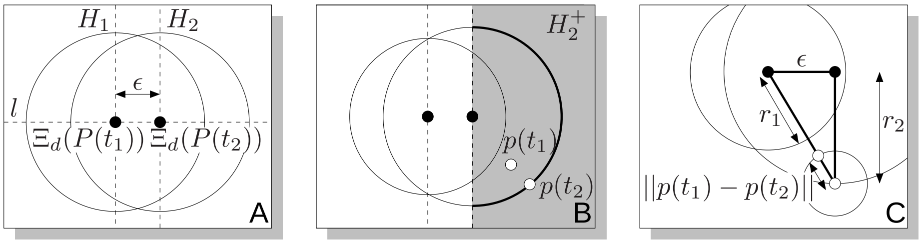

Figure 4.1: illustration in support of Obs. 4.1

4.2 Properties of the Mobile Euclidean 1-Centre

This section briefly explores the continuity and velocity of the mobile Euclidean 1-centre. Refer to Sec. 2.3.2 for the static definition of the Euclidean 1-centre.

We begin by verifying that the motion of the mobile Euclidean 1-centre is continuous. Although it seems unlikely that this result is new, the author was unable to find its proof in the literature. For completeness, the result is proved formally here.

Observation 4.1. *The mobile Euclidean 1-centre, Ξ_d, is continuous.*

Proof. Assume the mobile Euclidean 1-centre, Ξ_d, is discontinuous. Therefore, there exists a time interval $T = [0, t_f]$, an instant $t_1 \in (0, t_f)$, a fixed positive integer d, and a set of mobile clients P in $\mathbb{R}^d$ defined over T such that the motion of Ξ_d is discontinuous at t_1. By Def. 3.3, this implies

$$\exists \epsilon > 0, \ \forall \delta > 0, \ \exists t_2 \in (t_1 - \delta, t_1 + \delta), \ ||\Xi_d(P(t_1)) - \Xi_d(P(t_2))|| \geq \epsilon. \tag{4.1}$$

Let r_1 denote the Euclidean radius of $P(t_1)$. Let $\epsilon > 0$ be fixed such that Eq. (4.1) holds. Choose any $\delta \in (0, \sqrt{r_1^2 + \epsilon^2} - r_1)$. Choose $t_2 \in (t_1 - \delta, t_1 + \delta)$ such that Eq. (4.1) holds. Let r_2 denote the Euclidean radius of $P(t_2)$. Recall that the Euclidean 1-centre of a set of clients is unique. Therefore $\Xi_d(P(t_1))$ and $\Xi_d(P(t_2))$, the respective centres of the minimum enclosing hyperspheres of $P(t_1)$ and $P(t_2)$, are two distinct points that lie at least ϵ apart from each other. See Fig. 4.1A.

Let l be the line that passes through $\Xi_d(P(t_1))$ and $\Xi_d(P(t_2))$. Let H_1 and H_2 be the hyperplanes perpendicular to line l that pass through $\Xi_d(P(t_1))$ and $\Xi_d(P(t_2))$, respectively. Let H_2^+ denote the half-space induced by H_2 that lies opposite $\Xi_d(P(t_1))$. See Fig. 4.1B.

Case 1. Assume $r_2 \geq r_1$. There must be some client p in P whose position at time t_2 lies both in H_2^+ and on the minimum enclosing hypersphere of $P(t_2)$. We bound the displacement of p by

$$||p(t_1) - p(t_2)|| \leq |t_1 - t_2| < \delta < \sqrt{r_1^2 + \epsilon^2} - r_1 \leq \sqrt{r_2^2 + \epsilon^2} - r_1. \tag{4.2}$$

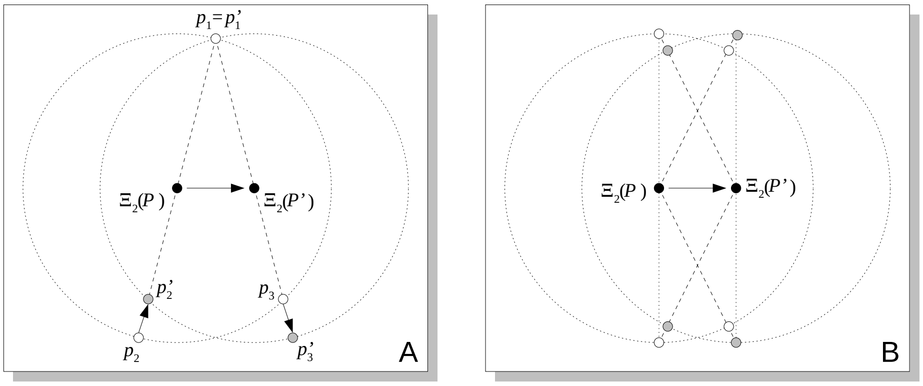

Figure 4.2: illustrations supporting Thm. 4.2 (reproduced from [BBKS00])

Since p lies in $\overline{H_2^+}$ and $r_2 \geq r_1$,

$$||\Xi_d(P(t_1)) - p(t_2)|| \geq \sqrt{r_2^2 + \epsilon^2}. \tag{4.3}$$

See Fig. 4.1C. By the triangle inequality,

$$\begin{aligned} ||\Xi_d(P(t_1)) - p(t_2)|| &\leq ||\Xi_d(P(t_1)) - p(t_1)|| + ||p(t_1) - p(t_2)|| \\ &\leq r_1 + ||p(t_1) - p(t_2)|| \\ &< \sqrt{r_2^2 + \epsilon^2}, \qquad \text{by Eq. (4.2).} \end{aligned} \tag{4.4}$$

Eqs. (4.3) and (4.4) derive a contradiction; our assumption must be false and the mobile Euclidean 1-centre, Ξ_d, must be continuous.

Case 2. Assume $r_2 < r_1$. The argument is analogous to Case 1, except we reverse t_1 and t_2. □

Although the motion of the mobile Euclidean 1-centre is continuous, Bereg et al. demonstrate that its velocity is unbounded in two or more dimensions. Specifically,

Theorem 4.2 (Bereg et al. 2006 [BBKS06]). *For any velocity $v \geq 0$ there is a set of three sites s_1, s_2, s_3 in $\mathbb{R}^d$, $d \geq 2$ such that a unit velocity motion of two of the sites induces an instantaneous velocity greater than v of the Euclidean 1-center.*

An example of a set of three mobile clients that realizes Thm. 4.2 is displayed in Fig. 4.2A. Bereg et al. also give a similar example using four mobile clients displayed in Fig. 4.2B.

As immediate consequence of Thm. 4.2, no bounded-velocity facility function can follow a trajectory that remains within an arbitrarily-small ϵ-neighbourhood around $\Xi_d(P(t))$ for $d \geq 2$.

When $d = 1$, the mobile Euclidean 1-centre moves with at most unit velocity relative to the velocity of clients.

Observation 4.3. *The one-dimensional mobile Euclidean 1-centre, Ξ_1, has maximum velocity 1. Furthermore, this velocity is realizable.*

Proof. Choose any time interval T and any finite set of mobile clients P in $\mathbb{R}$ defined over T.

$$\begin{aligned}\forall t_1, t_2 \in T,\ &||\Xi_1(P(t_1)) - \Xi_1(P(t_2))|| \\ &= \left|\left| \frac{1}{2}\left(\min_{p \in P(t_1)} p + \max_{q \in P(t_1)} q\right) - \frac{1}{2}\left(\min_{r \in P(t_2)} r + \max_{s \in P(t_2)} s\right)\right|\right| \\ &\leq \frac{1}{2}\left|\left|\min_{p \in P(t_1)} p - \min_{q \in P(t_2)} q\right|\right| + \frac{1}{2}\left|\left|\max_{p \in P(t_1)} p - \max_{q \in P(t_2)} q\right|\right| \\ &\leq \max_{p \in P} ||p(t_1) - p(t_2)|| \\ &\leq |t_1 - t_2|.\end{aligned}$$

The bound is realized when the endpoints of P move with unit velocity in a common direction. $\square$

4.3 Comparison Measures

This section expands on the comparison measures defined in Ch. 3 in terms specific to centre functions. We examine bounds on the relationship between eccentricity and maximum velocity and enumerate additional properties naturally associated with notions of centrality. Due to their relevance to the topics of this section, we refer to several results from the work of Bereg et al. [BBKS00, BBKS06].

4.3.1 Bounds on Eccentricity and Maximum Velocity

We are motivated to define centre functions that approximate the Euclidean 1-centre in the sense that they come close to minimizing Expr. (2.5) and yet have bounded maximum velocity. Thus, we examine centre functions with the twofold objective of minimizing both eccentricity and maximum velocity.

Let P denote a finite set of clients. The simplest definition of a centre function $\Upsilon_d(P)$ (that is not independent of P) simply assigns $\Upsilon_2(P) = p$, for some client $p \in P$. Of course, the velocity of Υ_d cannot exceed that of client p, and thus its maximum velocity is one. Since all clients of P must be contained within the minimum enclosing hypersphere of P, the distance from p to any client $q \in P$ is at most the diameter of the hypersphere, namely, twice the Euclidean radius. This bound is tight; the worst case is realized when $\Upsilon_d(P)$ lies opposite $\Xi_d(P)$ from some client $q \in P$ and the line segment $\overline{pq}$ forms a diameter of the minimum enclosing hypersphere of P. See the example for $d = 2$ in Fig. 4.3. Therefore, Υ_d is 2-eccentric.

In fact, as demonstrated by Bereg et al., a similar property holds for any centre function on P contained within the convex hull of P:

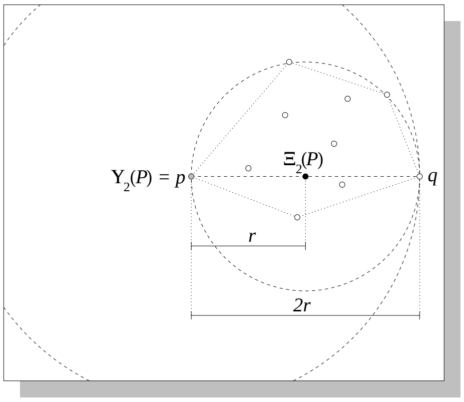

Figure 4.3: Centre function $\Upsilon_2(P) = p$ has eccentricity 2 and maximum velocity 1, where p is a client in P.

Lemma 4.4 (Bereg et al. 2006 [BBKS06]). *Let f be the initial position of a facility in $\mathbb{R}^d$.*

1. *If f is contained in the convex hull of P then there is an efficiently maintained unit velocity-bounded motion for f that guarantees a 2-approximation of the Euclidean 1-center [of P].*

2. *If f lies outside of the convex hull of P then no constant approximation factor can be guaranteed for any unit velocity-bounded motion for f.*

This sets an upper bound for the eccentricity factor λ; any reasonable bounded-velocity λ-eccentric centre function should have an eccentricity factor $\lambda \leq 2$. Furthermore, since the Euclidean 1-centre is defined as the point that minimizes Expr. (2.5), the eccentricity of any centre function must be at least 1.

As for maximum velocity, Bereg et al. show the following:

Theorem 4.5 (Bereg et al. 2006 [BBKS06]). *There exist arbitrarily large sets P of mobile sites in $\mathbb{R}^d$, $d \geq 2$, with velocities bounded by 1, such that no mobile facility that moves with velocity at most 1 can maintain a λ-approximation of the Euclidean 1-center of P, for $\lambda < 2$.*

Thm. 4.5 implies that any centre function with maximum velocity at most one has eccentricity at least two. It is straightforward to show that if a centre function has maximum velocity less than one, then its eccentricity must be infinite.

Since our goal is the identification of bounded-velocity centre functions, this sets an upper bound for the range of maximum velocities for centre functions we consider; any reasonable bounded-velocity λ-eccentric centre function will have maximum velocity $v_{\max} \geq 1$.

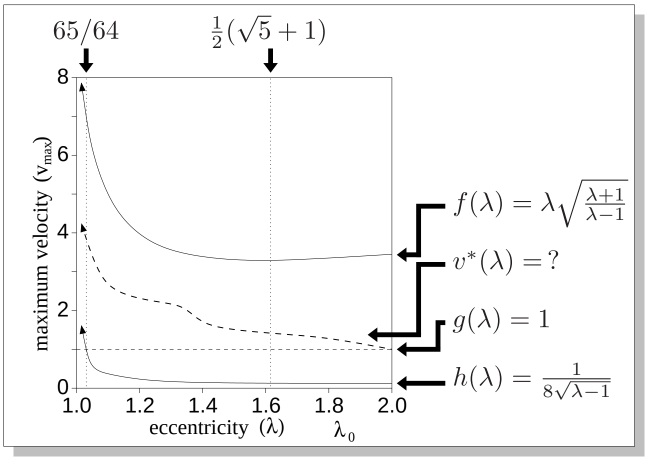

Figure 4.4: $v^*(\lambda)$ defines a theoretical lower bound on the maximum velocity of a centre function Υ_d with eccentricity λ. Although its behaviour is understood, the precise value of $v^*(\lambda)$ is unknown.

In summary, centre functions that define candidates for good bounded-velocity approximations of the mobile Euclidean 1-centre have eccentricity in the range $\lambda \in [1, 2]$ and maximum velocity in the range $v_{\max} \in [1, \infty)$; a good centre function will have eccentricity and maximum velocity both close to 1.

4.3.2 Maximum Velocity as a Function of Eccentricity

Reducing eccentricity increases maximum velocity and vice-versa. The challenge lies in understanding the trade-off between the degree of eccentricity (in the range $[1, 2]$) and the maximum velocity (in the range $[1, \infty)$). To express the actual correlation between $v_{\max}$ and λ, we define a function v^* over all centre functions Υ_d and all sets of clients P. For any fixed $\lambda \in [1, 2]$, let $v^*(\lambda)$ denote the lowest maximum velocity over all centre functions with eccentricity λ. This defines a function $v^* : [1, 2] \to [1, \infty)$, where $v^*(2) = 1$ and $\lim_{\lambda \to 1} v^*(\lambda) = \infty$. Thus, the maximum velocity of any centre function Υ_d with eccentricity λ is at least $v^*(\lambda)$. While the precise value of function $v^*(\lambda)$ for any fixed $\lambda \in (1, 2)$ remains unknown, the asymptotic behaviour of $v^*(\lambda)$ is understood, as shown by Bereg et al.:

Theorem 4.6 (Bereg et al. 2006 [BBKS06])**.** *For every $\epsilon > 0$, any $(1 + \epsilon)$-approximate mobile Euclidean 1-center has velocity at least $1/(8\sqrt{\epsilon})$ in the worst case.*

In terms of function v^*, Thm. 4.6 implies

$$v^*(\lambda) \geq \frac{1}{8\sqrt{\lambda - 1}}, \tag{4.5}$$

since $\lambda = 1 + \epsilon$. See Fig. 4.4. Since any centre function must have at least unit velocity, the lower bound implied by Thm. 4.6 is only valid in the range $\lambda \in [1, 65/64]$. This follows directly from Eq. (4.5):

$$\frac{1}{8\sqrt{\lambda - 1}} \geq 1 \Leftrightarrow \lambda \leq \frac{65}{64}. \tag{4.6}$$

Although this lower bound may not be realizable, Thm. 4.6 implies that no bounded-velocity λ-approximate centre function is possible for an arbitrary $\lambda \geq 1$ and a fixed v_{max} that is independent of λ. We refer to this lower bound again in Sec. 4.9, upon comparing the eccentricity and maximum velocity of various centre functions.

Bereg et al. consider the strategy of always moving the centre function toward the current position of the Euclidean 1-centre. The corresponding relationship between maximum velocity and eccentricity for this strategy leads to the following upper bound on v^* in $\mathbb{R}^2$:

Theorem 4.7 (Bereg et al. 2006 [BBKS06])**.** *For any $\epsilon > 0$ there is a strategy for moving a facility such that* (i) *the location of the facility provides an approximation of the Euclidean 1-center of a set P of points in $\mathbb{R}^2$ that is never worse than $1 + \epsilon$, and* (ii) *the velocity of the facility never exceeds*

$$\frac{(2 + \epsilon)(1 + \epsilon)}{\sqrt{2\epsilon + \epsilon^2}}. \tag{4.7}$$

In terms of λ, Thm. 4.7 implies

$$v^*(\lambda) \leq \lambda \sqrt{\frac{\lambda + 1}{\lambda - 1}}, \tag{4.8}$$

when the set of clients, P, lies in $\mathbb{R}^2$. See Fig. 4.4. Eq. (4.8) achieves a local minimum at $\lambda = \frac{1}{2}(1 + \sqrt{5}) \approx 1.61803$. Thus, the range over which Eq. (4.8) is decreasing corresponds to values $\lambda \in (1, \frac{1}{2}(\sqrt{5} + 1))$.

The bounds on function v^* given in Eqs. (4.5) and (4.8) are by no means tight; in particular, the centre functions we examine provide data points much closer to the true value of $v^*(\lambda)$ than the upper bound of Eq. (4.8).

Observation 4.8. *Function v^* is non-increasing.*

Proof. Suppose v^* is not non-increasing. That is, $v^*(\lambda') > v^*(\lambda)$ for some $1 < \lambda < \lambda'$. Let Υ_d be a centre function with eccentricity λ and maximum velocity $v^*(\lambda)$ (such a function must exist by the definition of v^*). Similarly, let Υ'_d be a centre function with eccentricity λ' and maximum velocity $v^*(\lambda')$. Since Υ_d is λ-eccentric and $\lambda' > \lambda$, Υ_d is also λ'-eccentric. Therefore, $v^*(\lambda') \leq v^*(\lambda)$, contradicting our original assumption. $\square$

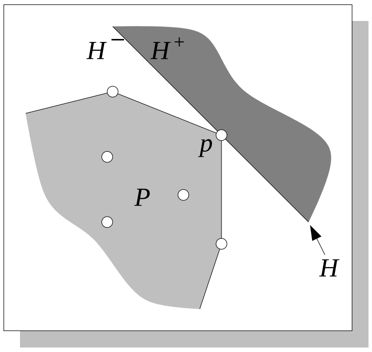

Figure 4.5: Point p is an extreme point of set P.

4.3.3 Additional Notions of Centrality

Although eccentricity and maximum velocity define the two principal measures by which we evaluate centre functions, the following define additional natural properties for a centre function Υ_d, all of which are properties exhibited by the Euclidean 1-centre:

1. $\Upsilon_d(P)$ should depend only on the extreme points of P (see Def. 4.1).
2. $\Upsilon_d(P)$ should be invariant under rotation, uniform scaling, reflection, and translation.
3. If P resides in a $(d-i)$-flat in $\mathbb{R}^d$, then the d-dimensional definition, $\Upsilon_d(P)$, should coincide with the $(d-i)$-dimensional definition, $\Upsilon_{d-i}(P)$.

For each centre function Υ_d examined, we evaluate the fitness of Υ_d and compare it against other centre functions primarily in terms of its eccentricity and maximum velocity. In addition, to further understand the behaviour of Υ_d, we also determine whether each of the properties listed above also holds for Υ_d.

Since its definition recurs frequently, we formalize the notion of an extreme point of a set of clients. A $(d-1)$-dimensional hyperplane H partitions $\mathbb{R}^d$ into three regions: H itself and the two open connected components of $\mathbb{R}^d - H$, which we denote by H^+ and H^-.

Definition 4.1. *A point p is an* **extreme point** *of the set P in $\mathbb{R}^d$ if and only if for some $(d-1)$-dimensional hyperplane H and associated half-space H^+, p satisfies $\overline{P} \cap \overline{H^+} = \{p\}$.*

Note that the extreme points of P are just the vertices of $CH(P)$, where $CH(A)$ denotes the convex hull of a set A in $\mathbb{R}^d$. See Fig. 4.5.

4.4 Rectilinear 1-Centre

This section examines properties of the mobile rectilinear 1-centre as an approximation to the mobile Euclidean 1-centre. Refer to Sec. 2.5.1 for a definition of

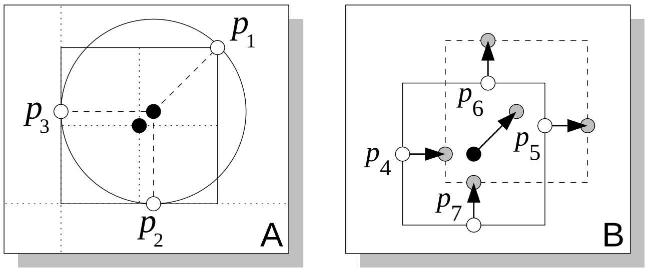

Figure 4.6: example realizing the eccentricity and maximum velocity of the rectilinear 1-centre (reproduced from [BBKS00])

the rectilinear 1-centre.

Recall that the rectilinear 1-centre of P, denoted $R_d(P)$, is a point that minimizes the maximum Chebyshev (ℓ_∞) distances from any client in P to $R_d(P)$. Given that its maximum velocity is bounded, this property suggests the rectilinear 1-centre as a natural candidate for providing an approximation of the Euclidean 1-centre.

As mentioned in Sec. 2.5.1, it is straightforward to demonstrate that the rectilinear 1-centre is invariant under translation and scaling, but not under rotation or reflection. Also, the rectilinear 1-centre depends only on the extreme points of P and its definition is consistent across dimensions.

4.4.1 Rectilinear 1-Centre: Eccentricity

As shown by Bereg et al. [BBKS00], the rectilinear 1-centre in $\mathbb{R}^2$ has eccentricity $(1+\sqrt{2})/2 \approx 1.2071$. The worst-case eccentricity is achieved by the following example. Let $p_1 = (1+1/\sqrt{2}, 1+1/\sqrt{2})$, $p_2 = (1,0)$, and $p_3 = (0,1)$ [BBKS00]. See Fig. 4.6A. We generalize this result to $\mathbb{R}^d$:

Theorem 4.9. *The d-dimensional rectilinear 1-centre, R_d, is $\frac{1}{2}(1+\sqrt{d})$-eccentric.*

Proof. Choose any $d \geq 1$. Assume P is a set of clients in $\mathbb{R}^d$ that maximizes the eccentricity of R_d. Since R_d is invariant under translation, assume the bounding box of P is $BB(P) = [0, x_1] \times \ldots \times [0, x_d]$, for some $x = (x_1, \ldots, x_d) \in [0, \infty)^d$. Every face of $BB(P)$ must be supported by the position of some client $p \in P$. The rectilinear 1-centre of P is located at $R_d(P) = x/2$.

Let $a \in P$ be a furthest client from $R_d(P)$. Since R_d is invariant under reflection, assume $a_i > R_d(P)_i$ for all $1 \leq i \leq d$. That is, a lies in the region $[R_d(P)_1, x_1] \times \ldots \times [R_d(P)_d, x_d]$. Let D denote the region $[0, a_1] \times \ldots \times [0, a_d]$. Let $P' = D \cap P$. Observe that a must lie in a corner of $BB(P)$ (that is, $D = BB(P)$) otherwise, the maximum distance from $R_d(P')$ to any client of P' would exceed the maximum distance from $R_d(P)$ to any client of P while the Euclidean radius of P' would be less than or equal to that of P. Consequently,

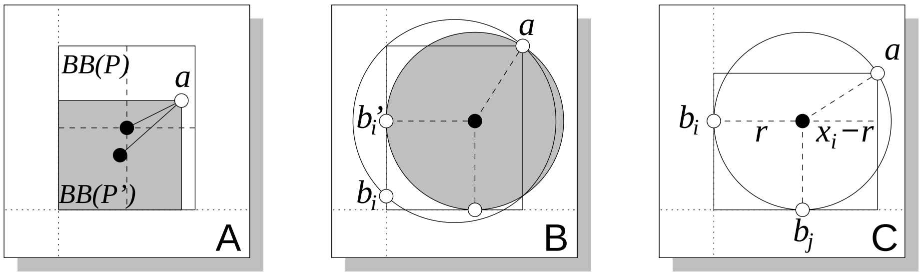

Figure 4.7: illustrations supporting Thm. 4.9

the eccentricity of P' would be greater than the eccentricity of P, contradicting our assumption. See Fig. 4.7A.

Client a supports the faces of $BB(P)$ that are not adjacent to the origin. The remaining faces of $BB(P)$ must be supported by other clients of P. Observe that a supporting client b_i must be perpendicular to $\Xi_d(P)$ with respect to the face, otherwise, moving b_i to the perpendicular position b_i' would reduce the Euclidean radius while the maximum distance from $R_d(P)$ to any client of P remained constant. Again, the eccentricity of P' would be greater than the eccentricity of P, contradicting our assumption. See Fig. 4.7B. Without loss of generality, we may scale P such that the coordinates of each supporting point, b_i, are

$$b_i = (\underbrace{1, \ldots, 1}_{i-1}, 0, \underbrace{1, \ldots, 1}_{d-i}).$$

The maximum distance from $R_d(P)$ to any client of P is $\frac{1}{2}||x||$. Let r denote the Euclidean radius of P. For every dimension i, $\Xi_d(P)$ lies a distance r from point b_i. Therefore, $\Xi_d(P)$ lies a distance $x_i - r$ from the opposite face of the bounding box. See Fig. 4.7C. The distance from $\Xi_d(P)$ to a also corresponds to the Euclidean radius, r. Consequently,

$$\begin{aligned} r^2 &= ||\Xi_d(P) - a||^2 \\ &= \sum_{i=1}^{d} (x_i - r)^2, \\ \Leftrightarrow 0 &= (d-1)r^2 - 2r \sum_{i=1}^{d} x_i + \sum_{i=1}^{d} x_i^2 \\ &= (d-1)r^2 - 2r||x||_1 + ||x||^2. \end{aligned}$$

Solving for r gives

$$r = \frac{||x||^2}{||x||_1 \pm \sqrt{||x||_1^2 - (d-1)||x||^2}}. \tag{4.9}$$

The plus/minus in Eq. (4.9) must be plus since by assuming it is minus we derive the following contradiction,

$$r = \frac{||x||^2}{||x||_1 - \sqrt{||x||_1^2 - (d-1)||x||^2}} \geq \frac{||x||^2}{||x||_1} = \left(\frac{||x||}{||x||_1}\right)||x|| > \frac{1}{2}||x||.$$

Since the Euclidean radius cannot be greater than the maximum distance from $R_d(P)$ to any client of P, the sign must be plus. The eccentricity, λ, of $R_d(P)$ is given by the ratio of the the Euclidean radius and the maximum distance from $R_d(P)$ to any client of P. That is,

$$\begin{aligned}
\lambda &= \frac{||x||}{2r} \\
&= \frac{||x||\left(||x||_1 + \sqrt{||x||_1^2 - (d-1)||x||^2}\right)}{2||x||^2} \\
&= \frac{||x||_1 + \sqrt{||x||_1^2 - (d-1)||x||^2}}{2||x||} \\
&= \frac{1}{2}\left(\frac{||x||_1}{||x||} + \sqrt{\left[\frac{||x||_1}{||x||}\right]^2 - d + 1}\right). \qquad (4.10)
\end{aligned}$$

Since all values are non-negative, λ is maximized if and only if $||x||_1/||x||$ is maximized. To locate this maximum, we examine the partial derivatives of $||x||_1/||x||$ with respect to x_i for all $1 \leq i \leq d$.

$$\begin{aligned}
& \frac{\partial}{\partial x_i}\frac{||x||_1}{||x||} = 0, \\
\Rightarrow\quad & \frac{\left(\frac{\partial}{\partial x_i}||x||_1\right)||x|| - ||x||_1\left(\frac{\partial}{\partial x_i}||x||\right)}{||x||^2} = 0, \\
\Rightarrow\quad & \frac{1}{||x||} - \frac{x_i||x||_1}{||x||^3} = 0, \\
\Rightarrow\quad & x_i = \frac{||x||^2}{||x||_1}. \qquad (4.11)
\end{aligned}$$

Since Eq. (4.11) must hold for all $1 \leq i \leq d$, it must be the case that $x_1 = \ldots = x_d$. Consequently, $||x||_1/||x|| = \sqrt{d}$ and Eq. (4.10) becomes

$$\lambda = \frac{1}{2}(\sqrt{d} + 1). \qquad \square$$

The bound on the eccentricity of the rectilinear 1-centre is tight. Thus,

Corollary 4.10. . *The d-dimensional rectilinear 1-centre, R_d, cannot guarantee λ-eccentricity for any $\lambda < \frac{1}{2}(\sqrt{d}+1)$.*

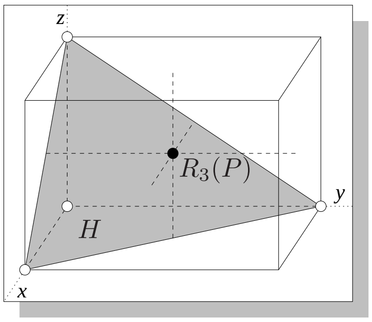

Figure 4.8: In $\mathbb{R}^3$, the rectilinear 1-centre may lie outside the convex hull.

Proof. The result follows from the example derived in the proof of Thm. 4.9 which achieves eccentricity $\frac{1}{2}(\sqrt{d}+1)$. □

As a consequence of Thm. 4.9, observe that in higher dimensions, the rectilinear 1-centre is quite eccentric, more so than any fixed client of P or any client contained within the convex hull of P (these both have eccentricity at most 2 by Lem. 4.4). Specifically, Thm. 4.9 implies that when $d = 9$, R_9 is 2-eccentric and when $d > 9$, R_d has eccentricity greater than 2. This is because although $R_d \in CH(P)$ in $\mathbb{R}$ and $\mathbb{R}^2$, the rectilinear 1-centre may lie outside the convex hull in $\mathbb{R}^d$ for $d \geq 3$. We formalize this observation with an example in $\mathbb{R}^3$.

Observation 4.11. *For some sets of clients P in $\mathbb{R}^d$, the rectilinear 1-centre of P lies outside the convex hull of P.*

Proof. Let $d = 3$ and let $P = \{(0,0,0),(0,0,1),(0,1,0),(1,0,0)\}$. Let H^- denote the half-space induced by the clients $\{(0,0,1),(0,1,0),(1,0,0)\}$ with corresponding equation $x + y + z - 1 \leq 0$. See Fig. 4.8. The convex hull of P is simply $CH(P) = H^- \cap [0,1]^3$. The rectilinear 1-centre of P lies at $R_3(P) = \frac{1}{2}(1,1,1)$. Since $\frac{1}{2}+\frac{1}{2}+\frac{1}{2}-1 = \frac{1}{2} > 0$, $R_3(P) \notin CH(P)$. □

4.4.2 Rectilinear 1-Centre: Maximum Velocity

Bereg et al. give the following tight bound on the velocity of the d-dimensional rectilinear 1-centre:

Observation 4.12 (Bereg et al. 2000 [BBKS00])**.** *For any instance of the mobile 1-center problem in $\mathbb{R}^d$, $d \geq 1$ there is a rectilinear 1-center whose velocity is bounded by $\sqrt{d}$. Furthermore, there is an instance of the problem with a unique solution moving with velocity $\sqrt{d}$.*

An example in $\mathbb{R}^2$ that achieves velocity $\sqrt{2}$ is given by four clients, $p_4, \ldots,$ p_7, such that each edge of the bounding box contains one client in its interior. Pairs of clients opposite each other move in a common direction perpendicular to the adjacent edge [BBKS00]. See Fig. 4.6B. This example is easily generalized to higher dimensions.

4.5 Centre of Mass

This section discusses properties of the centre of mass as an approximation to the mobile Euclidean 1-centre. Refer to Sec. 2.5.2 for the definition of the centre of mass.

The centre of mass of a set of clients P, denoted $C_d(P)$, is commonly used to define a point that is "central" to P. The centre of mass performs reasonably well in the mobile setting as a bounded-velocity approximation of the Euclidean 1-centre. By Lem. 4.4, assigning $\Upsilon_d(P)$ to be any client $p \in P$ defines a 2-eccentric centre function with maximum velocity one. As we show in Sec. 4.5, defining Υ_d to be the average position of clients in P improves neither eccentricity nor maximum velocity.

Unlike centre functions whose definitions depend only on the extreme points of P, the centre of mass assigns equal importance (hence its name) to every client of P. The invariance of the centre of mass under similarity transformations is straightforward to demonstrate. Also, the definition of the centre of mass is consistent across dimensions.

4.5.1 Centre of Mass: Eccentricity

Since $C_d(P)$ is a convex combination of the positions of clients of P, $C_d(P)$ must lie in the convex hull of P. Consequently, C_d has eccentricity at most 2 by Lem. 4.4. More precisely, Bereg et al. give the following bound on the eccentricity of the d-dimensional centre of mass:

Lemma 4.13 (Bereg et al. 2006 [BBKS06])**.** *The centre of mass of a set of n sites P in $\mathbb{R}^d$, $d \geq 1$, provides a $(2-\frac{2}{n})$-approximation of the Euclidean 1-center of P.*

This bound is tight:

Corollary 4.14. *The d-dimensional centre of mass, C_d, cannot guarantee λ-eccentricity for any $\lambda < 2$.*

Proof. The result follows from Thm. 4.5 and Cor. 4.15. □

The worst case is realized even in one dimension by $n-1$ clients located at the origin and a single client located at any fixed distance away from the origin.

4.5.2 Centre of Mass: Maximum Velocity

The maximum velocity of the centre of mass is an immediate consequence of a result of Bereg et al. [BBKS06], which we mention in our discussion of convex combinations of centre functions in Sec. 4.8.

Corollary 4.15. *The d-dimensional mobile centre of mass, C_d, has maximum velocity 1.*

Proof. The result follows from Obs. 4.32. □

The bound on maximum velocity is tight.

Observation 4.16. *The d-dimensional mobile centre of mass cannot guarantee relative velocity less than* 1.

Proof. When all clients of P move with unit velocity in a common direction, the velocity of the centre of mass is exactly one. □

4.6 Steiner Centre

The Steiner centre is named after Jakob Steiner who first introduced this point in the late nineteenth century [Ste81]. The original definition of the Steiner centre was phrased in terms of projection and integration, leading to the definition in Sec. 4.6.2. A second, fundamentally different definition, phrased in terms of Gaussian weights given by turn angles at the extreme points leads to the definition in Sec. 4.6.1. The equivalence of these two definitions was shown by Shephard [She66]. The dual definitions allow for numerous properties to be established; the definition by Gaussian weights lends itself to implementation within a kinetic data structure while the definition by projection allows the demonstration of bounds on eccentricity and maximum velocity.

The properties of the mobile Steiner centre compare very well against those of other mobile centre functions and suggest the Steiner centre as a natural choice for a bounded-velocity approximation to the mobile Euclidean 1-centre.

Synonyms for the Steiner centre include the Steiner curvature centroid [Buc80, Hon95], Steiner point [Grü67, Sal66, She64, She66, She68], Kimberling triangle centre $X(1115)$ [Kim], Gaussian centre [DK03, DK04], and projection centre [DK04].

The Steiner centre of a static set of clients can be found in $O(n \log n)$ time in $\mathbb{R}^2$ and $\mathbb{R}^3$ by direct implementation of the definition by Gaussian weights described below. This time complexity derives from finding the convex hull of the client set, requiring $\Theta(n \log h)$ time in $\mathbb{R}^2$ [KS86] and $\mathbb{R}^3$, [Cha96], where h denotes the number of clients on the convex hull boundary. Developing efficient algorithms for maintaining the mobile Steiner centre require us to examine the m-hull of a set of clients, an approximation of the two-dimensional convex hull which is detailed Sec. 8.2.2. In this section we focus on examining the quality of the Steiner centre as an approximation of the mobile Euclidean 1-centre and we postpone algorithmic considerations until Ch. 8.

4.6.1 Definition by Gaussian Weights

Euclidean space, specifically, $\mathbb{R}$, $\mathbb{R}^2$, and $\mathbb{R}^3$, define the most common settings for a variety of geometric problems including centre functions. Whereas the centre problem is simpler in one dimension (even the Euclidean 1-centre has maximum velocity one in $\mathbb{R}$), the more interesting characteristics of a centre function are of exhibited in $\mathbb{R}^2$ and $\mathbb{R}^3$.

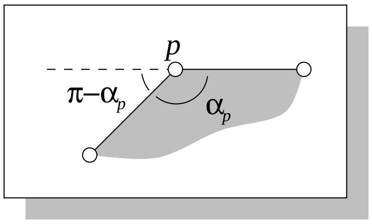

Figure 4.9: The turn angle at vertex p on the convex hull of P defines the two-dimensional Gaussian weight of p, $w_2(p) = \pi - \alpha_p$.

In this section we motivate defining a centre function using Gaussian weights, namely, the Steiner centre, first for a finite set P in $\mathbb{R}^2$, and then, more generally, for any bounded set P in $\mathbb{R}^2$. We then generalize the definitions to $\mathbb{R}$ and $\mathbb{R}^3$. As will be discussed in Sec. 8.2, the simple and intuitive definition of the Steiner centre by Gaussian weights will prove effective in efficiently calculating the position of a mobile facility that balances low maximum velocity and low eccentricity.

Steiner Centre Definition by Gaussian Weights in Two Dimensions

The simplest setting in which to examine the problem of finding the Euclidean 1-centre is $\mathbb{R}$. In this domain, the Euclidean 1-centre of a finite set of clients P in $\mathbb{R}$ satisfies

$$\Xi_1(P) = \frac{1}{2}\left(\min_{p\in P} p + \max_{q\in P} q\right). \tag{4.12}$$

Given a finite set of clients P in $\mathbb{R}$, $\Xi_1(P)$ is the average of the two extreme points of P. The same is true in any dimension; $\Xi_d(P)$ is determined by the extreme points of P. In two or more dimensions the mean of the extreme points of P is discontinuous and provides a poor approximation (eccentricity 2) to the Euclidean 1-centre. The discontinuity of such a centre function Υ_2 becomes evident whenever the motion of clients in P alters the composition of the set of extreme points. For the same reason, any centre function defined as a fixed weighted average of the extreme points of P is also discontinuous. Nevertheless, while the mean of the extrema does not provide a robust centre function, Eq. (4.12) suggests other possible generalizations to higher dimensions. By choosing weights that depend on the degree of extremity of individual clients it is possible to ensure not only continuity but also a low upper bound on maximum velocity.

For clarity, Defs. 4.2, 4.3, 4.5, and 4.6 assume $|P| \geq 2$. In the case when $|P| = 1$ (that is, $P = \{p\}$, for some p) we simply define the Steiner centre of P to be $\Gamma_d(P) = p$.

Definition 4.2. *Let P in $\mathbb{R}^2$ be a finite set of clients with $|P| \geq 2$. Let V_P denote the set of extreme points of P. For every $p \in V_P$, let α_p denote the interior angle formed on the convex hull boundary at p. The* **two-dimensional**

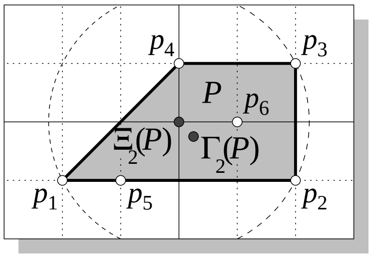

Figure 4.10: the Steiner centre $\Gamma_2(P)$ of the set $P = \{p_1, \ldots, p_6\}$

Gaussian weight *of* p *is*

$$w_2(p) = \begin{cases} \pi - \alpha_p & \text{if } p \in V_P \\ 0 & \text{if } p \in P - V_P. \end{cases} \tag{4.13}$$

For $p \in V_P$, $w_2(p)$ corresponds to the turn angle at p on $CH(P)$. Consequently, $\sum_{p \in P} w_2(p) = 2\pi$. See Fig. 4.9. Note, $w_2(p) > 0$ if and only if p is an extreme point of P. Expressed in terms of Gaussian weight, the Steiner centre is defined as the normalized weighted centre of mass of P, with weights specified by the Gaussian weights of P.

Definition 4.3. *Let P in $\mathbb{R}^2$ be a finite set of clients with $|P| \geq 2$. The* **two-dimensional Steiner centre** *of P is the normalized weighted mean of P:*

$$\Gamma_2(P) = \frac{1}{2\pi} \sum_{p \in P} w_2(p)p, \tag{4.14}$$

where $w_2(p)$ is the two-dimensional Gaussian weight of client $p \in P$.

For example, let $P = \{p_1, \ldots, p_6\} = \{(-2, -1), (2, -1), (2, 1), (0, 1), (-1, -1), (1, 0)\}$, respectively. See Fig. 4.10. Since $w_2(p) = \pi - \alpha_p$, clients $p_1, \ldots, p_6$ have weights $3\pi/4$, $\pi/2$, $\pi/2$, $\pi/4$, 0, and 0, respectively. The Steiner centre of P, $\Gamma_2(P)$, lies in position $(1/4, -1/4)$. The Euclidean 1-centre of P, $\Xi_2(P)$, lies at the origin.

Since the Steiner centre of P depends only on the extreme points of P, $\Gamma_2(P)$ remains well defined for infinite bounded sets P provided the set of extreme points is finite (for example, a polygonal region P). When the set of extreme points is infinite, the interior angle of clients on the convex hull is not well defined since distances between neighbouring clients may be infinitesimal. Thus, we provide a generalized definition of the Steiner centre equivalent to Def. 4.3.

Let $\text{ext}(P, \theta)$ denote an extreme point of set P in direction $(\cos\theta, \sin\theta)$. That is, $p = \text{ext}(P, \theta)$ if and only if there exists a half-plane H^+ with outer normal $(\cos\theta, \sin\theta)$ such that $\overline{P} \cap \overline{H^+} = \{p\}$. Note, the extreme point in a given direction θ may not exist (if it does exist, then it is unique by the above definition). We select a unique extreme point for every θ by defining $\text{Ext}(P, \theta) = \lim_{\phi \to \theta^+} \text{ext}(P, \phi)$.

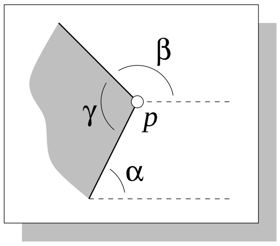

Figure 4.11: alternative definition for the Gaussian weight of p

Thus, just as we defined the continuous Euclidean 1-centre and the continuous centre of mass in Sec. 2.5.3, we now define the continuous Steiner centre of any bounded set of clients P:

Definition 4.4. *Let P in $\mathbb{R}^2$ be a bounded set of clients. The* **two-dimensional continuous Steiner centre** *of P is*

$$\Gamma_2(P) = \frac{1}{2\pi} \int_0^{2\pi} \mathrm{Ext}(P, \theta)\, d\theta. \tag{4.15}$$

Eq. (4.15) is easily shown to be equivalent to Def. 4.3 for any finite set P.

Observation 4.17. *Given a finite set P in $\mathbb{R}^2$, the Steiner centre of P coincides with the continuous Steiner centre of P.*

Proof. Let p be an extreme point of P such that α and β define the angles of the edges adjacent to p relative to the x-axis. See Fig. 4.11. The Gaussian weight of p multiplied by its position is

$$w_2(p)p = (\pi - \gamma)p = (\beta - \alpha)p = \int_\alpha^\beta p\, d\theta = \int_\alpha^\beta \mathrm{Ext}(P, \theta)\, d\theta. \qquad \square$$

The Gaussian weight formulation of Steiner centre exhibits several desirable properties of a bounded-velocity centre function. The Steiner centre is defined solely in terms of the geometry of the boundary of the convex hull of P. Small changes in the convex hull result in small changes in the weights of clients. Specifically, if a client p moves continuously, then the weight of p changes continuously, even when p moves along, joins, or leaves the convex hull boundary. This continuous change in weights results in continuity in the motion of the Steiner centre by smoothly blending the contribution of each client. Furthermore, as proved in Sec. 4.6.3, the relative position of the Steiner centre is invariant under similarity transformations. Specifically, it is straightforward to show that $\Gamma_2(g(P)) = g(\Gamma_2(P))$, where $g : \mathbb{R}^2 \to \mathbb{R}^2$ is any translation, uniform scaling, rotation, or reflection for any bounded set of clients P in $\mathbb{R}^2$. The quality of Γ_2 as a bounded-velocity centre function is evaluated in terms of eccentricity and maximum velocity in Secs. 4.6.4 and 4.6.5 and compared against other centre functions in Sec. 4.9.

See Secs. 8.2.1 through 8.2.3 for a description of algorithms for maintaining the mobile Steiner centre in two dimensions.

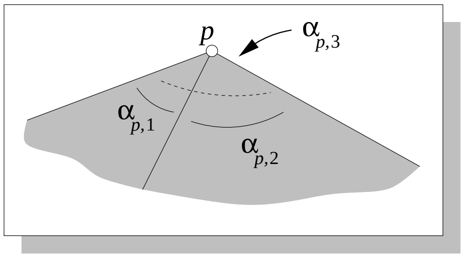

Figure 4.12: The plane angles at vertex p on the convex hull of P define the three-dimensional Gaussian weight of p, $w_3(p)$.

Steiner Centre Definition by Gaussian Weights in Three Dimensions In three dimensions, the Steiner centre of a set of clients P in $\mathbb{R}^3$ is again defined as a weighted mean. This time, however, a client $p \in V_P$ is adjacent to a set of faces on the boundary of the convex hull of P; the three-dimensional Gaussian weight of p is defined in terms of the angles formed at the faces that meet at p. See Fig. 4.12. Again, the turn angle at p (corresponding to the measure of the space of all supporting half-spaces of p) is directly proportional to the contribution of p to $\Gamma_3(P)$.

Definition 4.5. *Let P in $\mathbb{R}^3$ be a finite set of clients with $|P| \geq 2$. Let V_P denote the set of extreme points of P. For every $p \in V_P$, let F_p denote the set of faces that meet at p. For every face $f_j \in F_p$, let $\alpha_{p,j}$ denote the interior plane angle on f_j at p. The* **three-dimensional Gaussian weight** *of p is*

$$w_3(p) = \begin{cases} 2\pi - \displaystyle\sum_{f_j \in F_p} \alpha_{p,j} & \text{if } p \in V_P \\ 0 & \text{if } p \in P - V_P. \end{cases} \tag{4.16}$$

The sum of the plane angles at a client $p \in V_P$ ranges from 2π (when p is coplanar with its neighbours) and approaches a limit of 0 (when the neighbours of p approach collinearity). By Euler's theorem, the three-dimensional Gaussian weights of any arrangement of clients sum to 4π. Thus, in three dimensions we normalize by $1/4\pi$.

Definition 4.6. *Let P in $\mathbb{R}^3$ be a finite set of clients with $|P| \geq 2$. The* **three-dimensional Steiner centre** *of P is the normalized weighted mean of P:*

$$\Gamma_3(P) = \frac{1}{4\pi} \sum_{p \in P} w_3(p)p, \tag{4.17}$$

where $w_3(p)$ is the three-dimensional Gaussian weight of client $p \in P$.

As we did in $\mathbb{R}^2$, the definition of the Steiner centre generalizes to any bounded set of clients P in $\mathbb{R}^3$. Let $\text{ext}(P, \theta, \phi)$ denote an extreme point of a nonempty and bounded set $P \in \mathscr{P}(\mathbb{R}^3)$ in direction $(\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi)$.

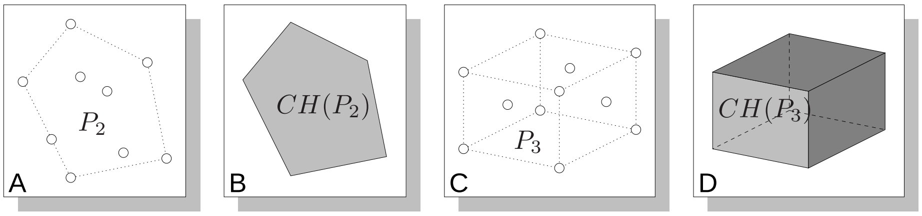

Figure 4.13: sets P_2 in $\mathbb{R}^2$ and P_3 in $\mathbb{R}^3$ and their corresponding 2- and 3-polytopes, $CH(P_2)$ and $CH(P_3)$.

That is, $p = \mathrm{ext}(P, \theta, \phi)$ if and only if there exists a half-plane $H^+ \subseteq \mathbb{R}^3$ with outer normal $(\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi)$ such that $\overline{P} \cap \overline{H^+} = \{p\}$. Again, the extreme point in a given direction may not exist (if it does exist, then it is unique by the above definition). We select a unique extreme point for every pair $\theta \in [0, 2\pi)$ and $\phi \in [0, \pi)$ by defining $\mathrm{Ext}(P, \theta, \phi) = \lim_{\alpha\to\theta^+, \beta\to\phi^+} \mathrm{ext}(P, \alpha, \beta)$.

Definition 4.7. *Let P in $\mathbb{R}^3$ be a bounded set of clients. The* **three-dimensional continuous Steiner centre** *of P is*

$$\Gamma_3(P) = \frac{1}{4\pi} \int_0^{2\pi} \int_0^{\pi} \sin\phi \, \mathrm{Ext}(P, \theta, \phi) \; d\theta. \tag{4.18}$$

The factor $\sin\phi$ is required for uniform integration over points on a sphere.

Steiner Centre Definition in One Dimension

A set of collinear client positions P can be viewed as a degenerate set of points in $\mathbb{R}^2$ whose convex hull (a line segment) has two extreme points, each with interior angle 0. This leads to a one-dimensional definition of Γ_1 that assigns equal weight to the extreme points of P. Thus,

Definition 4.8. *Let P in $\mathbb{R}$ be a bounded set of clients. The* **one-dimensional Steiner centre** *of P is*

$$\Gamma_1(P) = \frac{1}{2}\left(\min_{p\in\overline{P}} p + \max_{p\in\overline{P}} p\right). \tag{4.19}$$

Observe that $\Gamma_1(P) = \Xi_1(P)$.

Correspondence between Gaussian Weights and the Gauss Map

The Gauss map (normal map) provides an alternative interpretation for Gaussian weights. Since a Gauss map is typically defined on a polytope, we begin our discussion of Gauss maps by first introducing the polytope as an alternative to a set of points as an input parameter for some centre functions.

When Υ_d is one of the Euclidean 1-centre, the rectilinear 1-centre, or the Steiner centre, $\Upsilon_d(P) = \Upsilon_d(CH(P))$ for any finite set of clients P in $\mathbb{R}^d$. This property is not true of the centre of mass since its definition includes points in the interior of $CH(P)$.

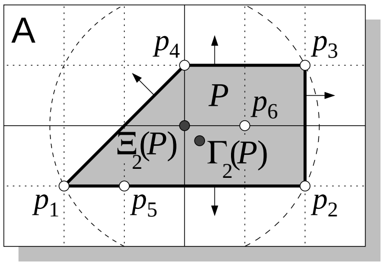

Figure 4.14: correspondence between Gaussian weights and the Gauss map in $\mathbb{R}^2$

Definition 4.9. *A* ***d*-polytope** *is the convex hull of a finite set of points in* $\mathbb{R}^d$*. We refer simply to a* **polytope** *if d is understood to be arbitrary.*

Def. 4.9 implies that a polytope P is nonempty, bounded, closed, convex, and has a finite number of extreme points. See Fig. 4.13.

The two-dimensional Gauss map of a 2-polytope P is the set of normals to edges of P projected from the origin as vertices on the unit circle. See [Car76] and [Got96] for discussions of the Gauss map. Given a finite set of clients P in $\mathbb{R}^2$, the Gauss map G_P of $CH(P)$ divides the unit circle into sectors such that the Gaussian weight of each extreme point of P is given by the length of its corresponding arc in G_P or, equivalently, the corresponding sector angle. The Gaussian weight of client $p \in V_P$ corresponds to the angular difference between the normals of the edges incident on p. The example in Fig. 4.14B displays the Gauss map of the set of clients P from Fig. 4.14A.

Similarly, the Gauss map in $\mathbb{R}^3$ can be used to interpret three-dimensional Gaussian weights. The three-dimensional Gauss map G_P of a 3-polytope P is the set of normals to faces of P projected from the origin onto the unit sphere. G_P divides the unit sphere into spherical sectors such that the Gaussian weight of each client $p \in V_P$ is given by the surface area of its corresponding spherical polygon in G_P.

For example, let $P = \{a, b, c, d\}$ be a regular tetrahedron. At each client in P, three faces meet, each forming a plane angle of $\pi/3$. Every client $p \in P$ has Gaussian weight $w_3(P) = 2\pi - (3 \cdot \pi/3) = \pi$. Observe that the spherical polygon

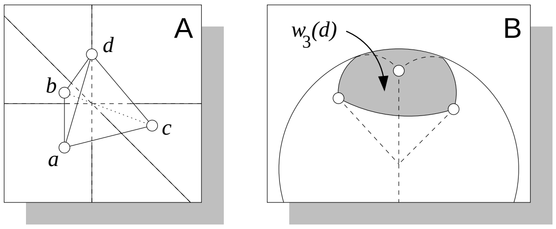

Figure 4.15: the Gauss map in three dimensions

corresponds to one quarter of the surface area of the unit sphere, an area equal to π. In this example, the Steiner centre coincides with the Euclidean 1-centre. See Fig. 4.15.

This equivalence is immediate in two dimensions. In three dimensions, the Gaussian weight of a vertex $p \in V_P$ of degree k is equal to the area of the corresponding k-sided spherical polygon as given by Girard's formula for the area of a spherical polygon. Plane angle $\theta_{p,i}$ at vertex p on the 3-polytope corresponds to interior angle $\pi - \theta_{p,i}$ on the spherical polygon. Thus, the Gaussian weight of client p, $w_3(p)$, is equivalent to the corresponding surface area, $A(p)$, on G_P:

$$w_3(p) = 2\pi - \sum_{j=1}^{k} \theta_{p,j} = \left[\sum_{j=1}^{k} (\pi - \theta_{p,j}) \right] - (k-2)\pi = A(p), \tag{4.20}$$

where $\theta_{p,j}$ are the plane angles at client p on the original 3-polytope.

Steiner Centre Definition by Gaussian Weights in Higher Dimensions
Just as the Steiner centre's definition was generalized from $\mathbb{R}^2$ to $\mathbb{R}^3$, a similar extension to $\mathbb{R}^4$ or $\mathbb{R}^d$ is possible by defining the appropriate higher-dimensional generalization of Gaussian weights. Where two-dimensional Gaussian weight corresponds to the turn angle and the three-dimensional Gaussian weight corresponds to the solid turn angle, a four-dimensional Gaussian weight could be defined in terms of turn angles at the extreme points of a 4-polytope, also called a *polychoron*. Just as the Gauss map provided an alternative interpretation for Gaussian weights in $\mathbb{R}^2$ and $\mathbb{R}^3$, the four-dimensional Gauss map provides a simple analogue for turn angles in $\mathbb{R}^4$, where the Gaussian weight of a client corresponds to a volume on the boundary of the unit hypersphere $\mathbb{S}^3$. For the remainder of the chapter, we restrict our attention to the Steiner centre in $\mathbb{R}$, $\mathbb{R}^2$, and $\mathbb{R}^3$. See Sec. 4.6.2 for a similar discussion on the generalization to higher dimensions of the Steiner centre's definition by projection.

4.6.2 Definition by Projection

Sec. 4.6.1 introduced the definition of the Steiner centre by Gaussian weights. The second definition which we now provide is by projection and integration over the positions of clients, first for a finite set in $\mathbb{R}^2$, and then in $\mathbb{R}$ and $\mathbb{R}^3$. As will become evident in Secs. 4.6.4, and 4.6.5, this second definition by projection lends itself to proving bounds on the eccentricity and maximum velocity of the Steiner centre.

Steiner Centre Definition by Projection in Two Dimensions
In one dimension, the Euclidean 1-centre of a finite set of clients P is simply

$$\Xi_1(P) = \frac{1}{2} \left(\min_{p \in P} p + \max_{q \in P} q \right). \tag{4.21}$$

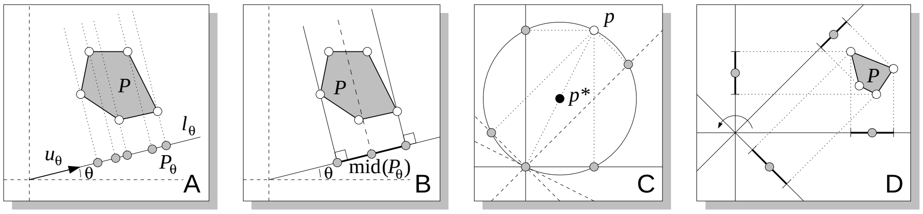

Figure 4.16: defining the Steiner centre Γ_2 by projection

That is, $\Xi_1(P)$ is the average of the two extreme points of P. As discussed in Sec. 4.6.1, while the mean of the extrema does not provide a robust centre function, Eq. (4.21) suggests other possible generalizations to higher dimensions.

One possibility is to project client positions onto a line through the origin, to find the one-dimensional Euclidean 1-centre of the projection, and to average these one-dimensional Euclidean 1-centres for all lines through the origin.

Let line l_θ be the line through the origin parallel to the unit vector $u_\theta = (\cos\theta, \sin\theta)$. Expressed in slope-intercept form, l_θ is the line $y = \tan\theta\ x$. Given a finite set of clients P in $\mathbb{R}^2$ and an angle $\theta \in [0, \pi)$, let P_θ denote the projection of P onto the line l_θ. See Fig. 4.16A. That is,

$$P_\theta = \{u_\theta \langle p, u_\theta \rangle \mid p \in P\}. \tag{4.22}$$

The midpoint of P_θ is just the Euclidean 1-centre of P_θ,

$$\mathrm{mid}(P_\theta) = \frac{u_\theta}{2}\left(\min_{p\in P}\langle p, u_\theta\rangle + \max_{q\in P}\langle q, u_\theta\rangle\right) = \Xi_2(P_\theta). \tag{4.23}$$

See Fig. 4.16B. Let $p \in \mathbb{R}^2$ be any fixed client. The average over all projections of p onto lines l_θ is

$$\frac{1}{\pi}\int_0^\pi u_\theta\langle p, u_\theta\rangle\ d\theta = \frac{p}{2}.$$

See Fig. 4.16C. Equivalently, if $P = \{p\}$,

$$p = \frac{2}{\pi}\int_0^\pi u_\theta\langle p, u_\theta\rangle\ d\theta = \frac{2}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\ d\theta.$$

This suggests the following definition of a centre function, shown to be equivalent to Def. 4.3 by Shephard [She66]:

Definition 4.10. *Let P in $\mathbb{R}^2$ be a finite set of clients. The* **two-dimensional Steiner centre** *of P is*

$$\Gamma_2(P) = \frac{2}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\ d\theta, \tag{4.24}$$

where $\mathrm{mid}(P_\theta)$ is the midpoint of the projection of P onto line $y = \tan\theta\ x$.

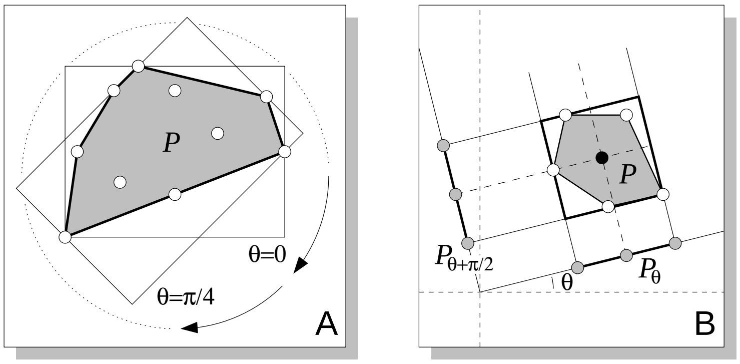

Figure 4.17: illustrations supporting Lem. 4.18

This second definition of the Steiner centre of P can be interpreted in terms of bounding boxes of P. The bounding box of P with orientation θ is simply $CH(P_\theta) + CH(P_{\theta+\pi/2})$, where addition denotes the Minkowski sum. Its centre is the point $\operatorname{mid}(P_\theta) + \operatorname{mid}(P_{\theta+\pi/2})$. See Fig. 4.17. Hence,

Lemma 4.18. *The Steiner centre of a set of clients P in $\mathbb{R}^2$, $\Gamma_2(P)$, is equivalent to the average of the centres of all bounding boxes of P.*

Proof.

$$\begin{aligned}\Gamma_2(P) =& \frac{2}{\pi} \int_0^\pi \operatorname{mid}(P_\theta)\ d\theta \\ =& \frac{2}{\pi} \left[\int_0^{\pi/2} \operatorname{mid}(P_\theta)\ d\theta + \int_{\pi/2}^{\pi} \operatorname{mid}(P_\theta)\ d\theta \right] \\ =& \frac{2}{\pi} \int_0^{\pi/2} \left[\operatorname{mid}(P_\theta) + \operatorname{mid}(P_{\theta+\pi/2}) \right]\ d\theta. \qquad \square\end{aligned}$$

Observe that the minimum of P_θ corresponds to the maximum of $P_{\theta+\pi}$. Specifically, we can rewrite Eq. (4.24) as

$$\begin{aligned}\Gamma_2(P) =& \frac{2}{\pi} \int_0^\pi \operatorname{mid}(P_\theta)\ d\theta \\ =& \frac{2}{\pi} \int_0^\pi \frac{u_\theta}{2} \left(\min_{p\in P} \langle p, u_\theta \rangle + \max_{q \in P} \langle q, u_\theta \rangle \right)\ d\theta \\ =& \frac{1}{\pi} \int_0^{2\pi} u_\theta \cdot \max_{q\in P} \langle q, u_\theta \rangle\ d\theta. \end{aligned} \tag{4.25}$$

The latter, Eq. (4.25), is used in the proof of Thm. 4.20.

Steiner Centre Definition by Projection in Three Dimensions

In three dimensions, we express the Steiner centre by projection in terms of spherical coordinates. Let $l_{\theta,\phi}$ be the line through the origin parallel to the unit

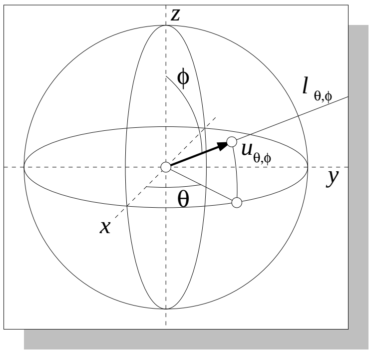

Figure 4.18: In $\mathbb{R}^3$, the Steiner centre is defined in terms of spherical coordinates, parameterized by θ and ϕ.

vector $u_{\theta,\phi} = (\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi)$. See Fig. 4.18. Let $P_{\theta,\phi}$ and $\mathrm{mid}(P_{\theta,\phi})$ be the natural generalizations of P_θ and $\mathrm{mid}(P_\theta)$ to spherical coordinates in $\mathbb{R}^3$, respectively. Thus,

$$P_{\theta,\phi} = \{u_{\theta,\phi}\langle p, u_{\theta,\phi}\rangle \mid p \in P\} \tag{4.26}$$

and

$$\mathrm{mid}(P_{\theta,\phi}) = \frac{u_{\theta,\phi}}{2}\left(\min_{p\in P}\langle p, u_{\theta,\phi}\rangle + \max_{q\in P}\langle q, u_{\theta,\phi}\rangle\right) = \Xi_3(P_{\theta,\phi}). \tag{4.27}$$

Let $p \in \mathbb{R}^3$ be any fixed point. The average over all projections of p onto all lines $l_{\theta,\phi}$ is

$$\frac{1}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \cdot p \; d\phi \; d\theta = \frac{p}{3}. \tag{4.28}$$

The factor $\sin\phi$ is required for uniform integration over points on a sphere. The factor $1/2\pi$ normalizes over the range of the integration as shown by $\int_0^\pi \int_0^\pi \sin\phi \; d\phi \; d\theta = 2\pi$. Adding a factor of three returns p instead of $p/3$, suggesting the following definition for a centre function:

Definition 4.11. *Let P in $\mathbb{R}^3$ be a finite set of clients. The* **three-dimensional Steiner centre** *of P is*

$$\Gamma_3(P) = \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \, \mathrm{mid}(P_{\theta,\phi}) \; d\phi \; d\theta, \tag{4.29}$$

where $\mathrm{mid}(P_{\theta,\phi})$ is the midpoint of the projection of P onto the line through the origin parallel to $u_{\theta,\phi} = (\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi)$.

This definition is the natural analogue of the two-dimensional projection centre, expressing $\Gamma_3(P)$ as the average midpoint over all projections of P onto lines $l_{\theta,\phi}$.

Steiner Centre Definition by Projection in Higher Dimensions
The definition of the Steiner centre by projection has a natural generalization to $\mathbb{R}^d$. We simply integrate the midpoint of the projection of P onto all lines through the origin and normalize by the volume of the unit hypersphere.

Definition 4.12. *Let P in $\mathbb{R}^d$ be a finite set of clients. Given a fixed $d \in \mathbb{N}$, the d-***dimensional Steiner centre** *of P is*

$$\Gamma_d(P) = \frac{d \displaystyle\int_{u \in \mathbb{S}^{d-1}} \mathrm{mid}(P_u)\ du}{\displaystyle\int_{u \in \mathbb{S}^{d-1}} 1\ du}, \tag{4.30}$$

where $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d \mid ||x|| = 1\}$ is the unit hypersphere and $\mathrm{mid}(P_u)$ is the midpoint of the projection of P onto the line through the origin parallel to vector u.

Again, we focus mainly on the definition of the Steiner centre in $\mathbb{R}^2$ and $\mathbb{R}^3$.

4.6.3 Properties of the Steiner Centre

To our knowledge, previous to our work, neither had the Steiner centre been evaluated as a stable approximation to the Euclidean 1-centre nor had its quality in defining the position of a mobile facility been examined. However, several useful related properties of the Steiner centre have been established which we mention here.

The Steiner centre is local [AF90]. That is, $\Gamma_d(P) \in CH(P)$. As is necessary for any bounded-velocity centre function, Γ_d is continuous [She64, She68].

When Γ_d is defined over polytopes P and Q, their respective Steiner centres are invariant under addition [She68]. That is, $\Gamma_d(P) + \Gamma_d(Q) = \Gamma_d(P + Q)$, where $P + Q$ denotes the Minkowski sum of sets P and Q. Furthermore, Γ_d is invariant under similarity transformations [She68]. Thus, for any similarity transformation $f : \mathbb{R}^d \to \mathbb{R}^d$, $\Gamma_d(f(P)) = f(\Gamma_d(P))$.

Sallee [Sal66] was the first to establish a relationship on the convex decomposition of a polytope. Given polytopes $P_1, \ldots, P_n$ such that $P = P_1 \cup \ldots \cup P_n$ is also a polytope, then

$$\Gamma_d(P) = \sum \Gamma_d(P_i) - \sum_{i<j} \Gamma_d(P_i \cap P_j) + \ldots + (-1)^{n-1}\Gamma_d(P_1 \cap \ldots \cap P_n). \tag{4.31}$$

Related to this idea, Shephard [She66] shows a relationship by a decomposition of a polytope P into its j-faces. That is, $\Gamma_d(P)$ can be defined in terms of the Steiner centres of the faces, edges, and vertices of P (in an arbitrary dimension d):

$$(1 + (-1)^{d-1})\Gamma_d(P) = \sum \Gamma_d(F_i^0) - \sum \Gamma_d(F_i^1) + \ldots + (-1)^{d-1} \sum \Gamma_d(F_i^{d-1}), \tag{4.32}$$

where F_i^j are the j-faces of P.

It is straightforward to show that the definition of Γ_d is consistent across dimensions. Finally, observe that $\Gamma_d(P)$ remains well defined when P is any nonempty and bounded region in $\mathbb{R}^d$.

4.6.4 Steiner Centre: Eccentricity

In this section we prove that the eccentricity of the Steiner centre is at most $\lambda \approx 1.1153$ in $\mathbb{R}^2$. We show that this maximum is achieved when the extreme points form an arc opposite an isolated point on the circle as displayed in Fig. 4.20B, where $\alpha = 0$ and $\beta = \gamma \approx 0.8105$. We generalize this worst-case example to $\mathbb{R}^3$ to provide a lower bound of $\lambda \approx 1.2017$ on the eccentricity of the Steiner centre in three dimensions.

Recall that the worst-case eccentricity of a centre function Υ_d is defined solely in terms of the position of Υ_d relative to the positions of the clients and of the corresponding Euclidean radius. That is, bounds on the eccentricity of Υ_d are independent of motion and are realized by the instantaneous position of a set of mobile clients. As such, we examine the set of positions $P(t_0)$ of a set of clients at some instant $t_0 \in T$. The value of t_0 is unimportant; for simplicity, we write simply P to denote $P(t_0)$ throughout Sec. 4.6.4

Eccentricity of the Steiner Centre in Two Dimensions
We first derive the eccentricity of the Steiner centre in $\mathbb{R}^2$.

Lemma 4.19. *Among all closed sets of clients P in $\mathbb{R}^2$ with Euclidean radius $r > 0$, the worst-case eccentricity of Γ_2 is realized when the extreme points of P consist of an arc A and an isolated point m on the circle C with radius r and centre $\Xi_2(P)$.*

Proof. Since $\Gamma_2(P) = \Gamma_2(CH(P))$ and $\max_{p\in P} ||\Gamma_2(P) - p||$ is realized at an extreme point of P, we can assume that P is a convex set. Let $m \in P$ be a furthest client from $\Gamma_2(P)$. Let a_x (respectively, a_y) denote the x-coordinate (respectively, y-coordinate) of a point $a \in \mathbb{R}^2$. Since Γ_2 is invariant under rotation and translation, without loss of generality, we can further assume that $m_y = \Gamma_2(P)_y$ and $m_x \geq \Gamma_2(P)_x$. Since $\max_{p\in P} ||\Gamma_2(P) - p|| \geq r > 0$, the line induced by m and $\Gamma_2(P)$ is well defined.

For $p \in P$, let

$$p' = \begin{cases} \text{left translation of } p \text{ to } C & \text{if } p \neq m \\ \text{right translation of } p \text{ to } C & \text{if } p = m \end{cases}. \tag{4.33}$$

Let set $P' = \{p' \mid p \in P\}$. Observe that every point in P' corresponds to a horizontal translation of some point in P. See Fig. 4.19. The x-coordinate of the Steiner centre of P' is given by

$$\Gamma_2(P')_x = \frac{2}{\pi} \int_0^{\pi} \mathrm{mid}(P'_\theta)_x \, d\theta. \tag{4.34}$$

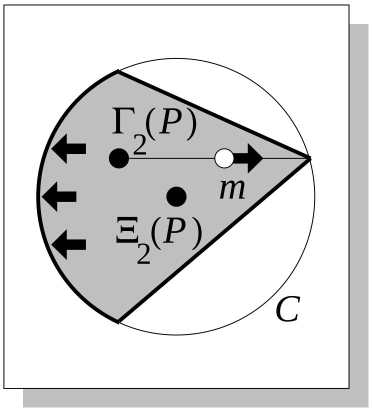

Figure 4.19: illustrations supporting Lem. 4.19

Since all clients of $P' - \{m'\}$ are left translations of clients in P,

$$\mathrm{mid}(P'_\theta)_x \leq \mathrm{mid}(P_\theta)_x + \frac{m'_x - m_x}{2}, \tag{4.35}$$

for any $\theta \in [0, \pi]$. Therefore,

$$\Gamma_2(P')_x \leq \Gamma_2(P)_x + (m'_x - m_x), \tag{4.36a}$$

and hence

$$m'_x - \Gamma_2(P')_x \geq m_x - \Gamma_2(P)_x. \tag{4.36b}$$

Since $m_x \geq \Gamma_2(P)_x$ and $m'_x \geq \Gamma_2(P')_x$,

$$|m'_x - \Gamma_2(P')_x| \geq |m_x - \Gamma_2(P)_x|. \tag{4.37}$$

Therefore,

$$||m' - \Gamma_2(P')|| \geq |m'_x - \Gamma_2(P')_x| \geq |m_x - \Gamma_2(P)_x| = ||m - \Gamma_2(P)||. \tag{4.38}$$

Since all points of P' lie within the minimum enclosing circle of P, the Euclidean radius of P' is at most the Euclidean radius of P. Therefore, Eq. (4.38) implies that the eccentricity of P' is at least as great as the eccentricity of P. The extreme points of set P' consist of an arc of C opposite the isolated point m'. □

Theorem 4.20. *The two-dimensional Steiner centre Γ_2 has eccentricity $\lambda \approx 1.1153$.*

Proof. It follows from Lem. 4.19 that to understand the eccentricity of Γ_2 it suffices to study point sets P formed by an arc A of a circle C and an isolated point m on C. Since Γ_2 is preserved by translation, reflection, rotation, and uniform scaling, we can assume C is the unit circle centred at the origin such

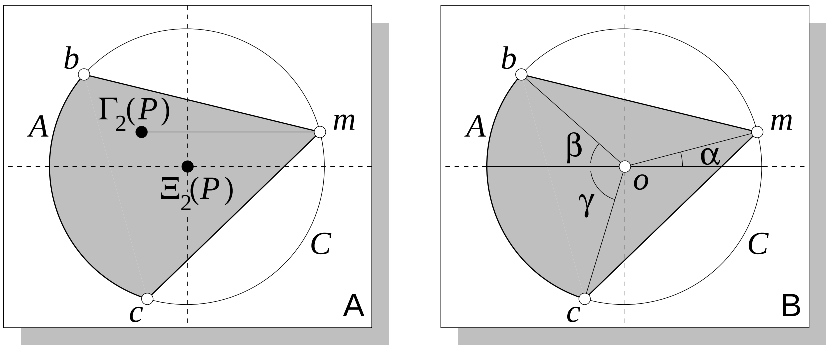

Figure 4.20: Thm. 4.20: maximizing the eccentricity of the Steiner centre

that m lies in the first quadrant and the line induced by m and $\Gamma_2(P)$ lies parallel to the x-axis. See Fig. 4.20A. Thus, point sets of interest are completely characterized by three parameters which specify the angles α, β, and γ formed, respectively, by the position of m relative to the positive x-axis and the endpoints of A relative to the negative x-axis. See Fig. 4.20B. Let $P_{\alpha,\beta,\gamma}$ denote such a set of points. To find a point set that realizes the worst-case eccentricity of Γ_2 we need only maximize $||\Gamma_2(P_{\alpha,\beta,\gamma}) - m||$. Since $\Gamma_2(P_{\alpha,\beta,\gamma})_y = m_y$ and $\Gamma_2(P_{\alpha,\beta,\gamma})_x \leq m_x$, this corresponds to maximizing $m_x - \Gamma_2(P_{\alpha,\beta,\gamma})_x$.

The Steiner centre of $P_{\alpha,\beta,\gamma}$ is straightforward to calculate by examination of the various cases for which specific extreme points of $P_{\alpha,\beta,\gamma}$ remain extreme in P_θ. The coordinates of the extreme points of P are $m = (\cos\alpha, \sin\alpha)$, $b = (-\cos\beta, \sin\beta)$, $c = (-\cos\gamma, -\sin\gamma)$, and $u_\theta = (\cos\theta, \sin\theta)$, for $\theta \in [\pi-\beta, \pi+\gamma]$.

Table 4.1 divides the range of integration, $\theta \in [0, 2\pi]$, into intervals for which each of the points m, b, c, and u_θ induce a maximum of P_θ.

interval of θ		$\arg\max_{p \in P} \langle p, u_\theta \rangle$
$[0,$	$(\pi+\alpha-\beta)/2]$	m
$[(\pi+\alpha-\beta)/2,$	$\pi-\beta]$	b
$[\pi-\beta,$	$\pi+\gamma]$	u_θ
$[\pi+\gamma,$	$(3\pi+\alpha+\gamma)/2]$	c
$[(3\pi+\alpha+\gamma)/2,$	$2\pi]$	m

Table 4.1: case analysis of extreme points in $\Gamma_2(P_{\alpha,\beta,\gamma})$ in Thm. 4.20

The x-coordinate of the Steiner centre of $P_{\alpha,\beta,\gamma}$ is given by

$$\begin{aligned}
\Gamma_2(P_{\alpha,\beta,\gamma})_x =& \frac{1}{\pi}\int_0^{2\pi} \cos\theta \cdot \max_{p\in P}\langle u_\theta, p\rangle \; d\theta \\
=& \frac{1}{\pi}\left[\int_0^{(\pi+\alpha-\beta)/2} \cos\theta\langle u_\theta, m\rangle \; d\theta + \int_{(\pi+\alpha-\beta)/2}^{\pi-\beta} \cos\theta\langle u_\theta, b\rangle \; d\theta \right.\\
&+ \int_{\pi-\beta}^{\pi+\gamma} \cos\theta\langle u_\theta, u_\theta\rangle \; d\theta + \int_{\pi+\gamma}^{(3\pi+\alpha+\gamma)/2} \cos\theta\langle u_\theta, c\rangle \; d\theta \\
&\left. + \int_{(3\pi+\alpha+\gamma)/2}^{2\pi} \cos\theta\langle u_\theta, m\rangle \; d\theta \right] \\
=& \frac{1}{4\pi}\left[-2\sin\beta - 2\sin\gamma - (\pi-\alpha-\beta)\cos\beta\right. \\
&\left. -(\pi+\alpha-\gamma)\cos\gamma + (2\pi-\gamma-\beta)\cos\alpha\right]. \qquad (4.39)
\end{aligned}$$

Let f denote the function $f(\alpha,\beta,\gamma) = m_x - \Gamma_2(P_{\alpha,\beta,\gamma})_x$. Values of α, β, and γ that define a local maximum of f must satisfy the following conditions:

$$\frac{\partial}{\partial\alpha}f = \frac{\partial}{\partial\beta}f = \frac{\partial}{\partial\gamma}f = 0.$$

Specifically,

$$\begin{aligned}
\frac{\partial}{\partial\beta}f =& \frac{1}{4\pi}[\cos\beta - (\pi-\alpha-\beta)\sin\beta + \cos\alpha] &= 0, \qquad (4.40a)\\
\frac{\partial}{\partial\gamma}f =& \frac{1}{4\pi}[\cos\gamma - (\pi+\alpha-\gamma)\sin\gamma + \cos\alpha] &= 0, \qquad (4.40b)\\
\text{and } \frac{\partial}{\partial\alpha}f =& \frac{1}{4\pi}[\cos\gamma - \cos\beta - (2\pi+\beta+\gamma)\sin\alpha] &= 0. \qquad (4.40c)
\end{aligned}$$

We now show that the constraints imposed by Eqs. (4.40a) through (4.40c) imply that for $(\alpha,\beta,\gamma) \in [0,\pi/2]^3$, f has only one local (and hence global) maximum occurring at $\alpha = 0$ and $\beta = \gamma \approx 0.81047$.

Since α, β, and γ lie in the interval $[0,\pi/2]$, the term $-(2\pi+\beta+\gamma)\sin\alpha$ in Eq. (4.40c) is nonpositive, meaning that $\cos\gamma - \cos\beta \geq 0$ and, consequently, $\gamma \leq \beta$. Furthermore, in order for the unit circle to define the minimum enclosing circle of $P_{\alpha,\beta,\gamma}$, line segment $\overline{cm}$ must pass below the origin, implying that $\gamma \geq \alpha$. See Fig. 4.20B. These constraints impose an ordering on the angles: $0 \leq \alpha \leq \gamma \leq \beta \leq \pi/2$.

We bound the value of α. Solving for $\sin\alpha$ in Eq. (4.40c) gives

$$\begin{aligned}
\sin\alpha =& \frac{\cos\gamma - \cos\beta}{2\pi+\beta+\gamma} \\
\leq& \frac{1}{2\pi}.
\end{aligned}$$

Therefore,

$$0 \le \alpha \le \arcsin\left(\frac{1}{2\pi}\right) \approx 0.159835 < \frac{3\pi}{50}. \tag{4.41}$$

We derive an upper bound on β using this bound on α. By Eq. (4.40a),

$$\begin{aligned} 0 &= \cos\beta - (\pi - \alpha - \beta)\sin\beta + \cos\alpha, \\ &\le \cos\beta - \left(\frac{47\pi}{50} - \beta\right)\sin\beta + 1, \text{ since } \alpha \in \left[0, \frac{3\pi}{50}\right]. \end{aligned} \tag{4.42}$$

Let $g(\beta) = \cos\beta - (47\pi/50 - \beta)\sin\beta + 1$. Observe that $g'(\beta) \le 0$ for $\beta \in [0, \pi/2]$. Furthermore, $g(1) < 0$. Consequently, $g(\beta) < 0$ for all $\beta \in [1, \pi/2]$. Since $g(\beta)$ must be nonnegative by Eq. (4.42), it follows that $\gamma \le \beta < 1$.

We now take a linear combination of Eqs. (4.40a), (4.40b), and (4.40c).

$$\begin{aligned} &4\pi\left(\frac{\partial}{\partial\gamma}f - \frac{\partial}{\partial\beta}f - \frac{\partial}{\partial\alpha}f\right) &&= 0, \\ &\Rightarrow (2\pi + \beta + \gamma)\sin\alpha - (\pi + \alpha - \gamma)\sin\gamma + (\pi - \alpha - \beta)\sin\beta &&= 0, \\ &\Rightarrow \underbrace{\beta\sin\alpha - \alpha\sin\beta}_{t_1} + \underbrace{\gamma\sin\alpha - \alpha\sin\gamma}_{t_2} \\ &\quad + \underbrace{(\pi - \beta)\sin\beta - (\pi - \gamma)\sin\gamma}_{t_3} + \underbrace{2\pi\sin\alpha}_{t_4} &&= 0. \end{aligned} \tag{4.43}$$

We examine terms t_1 through t_4 from Eq. (4.43). Let $h(x) = x/\sin x$. Observe that $\lim_{x\to 0} h'(x) = 0$ and $h''(x) \ge 0$ for $x \in [0, \pi/2]$. Thus, $h(x)$ is nondecreasing on the interval $[0, \pi/2]$, meaning that for any $0 \le \alpha \le \gamma \le \beta \le \pi/2$,

$$\frac{\beta}{\sin\beta} \ge \frac{\alpha}{\sin\alpha} \quad \text{and} \quad \frac{\gamma}{\sin\gamma} \ge \frac{\alpha}{\sin\alpha}. \tag{4.44}$$

Therefore, terms t_1 and t_2 in Eq. (4.43) are nonnegative.

Let $i(x) = (\pi - x)\sin x$. Observe that $i''(x) \le 0$ for $x \in [0, \pi/2]$ and $i'(1) > 0$. Therefore, $i(x)$ is nondecreasing on the interval $[0, 1]$. Consequently, since $0 \le \gamma \le \beta < 1$, we get

$$(\pi - \beta)\sin\beta - (\pi - \gamma)\sin\gamma \ge 0. \tag{4.45}$$

Therefore, term t_3 in Eq. (4.43) is nonnegative. Since terms t_1, t_2, and t_3 are nonnegative and Eq. (4.43) is equal to zero, term t_4 must be nonpositive. Thus,

$$2\pi\sin\alpha \le 0 \Rightarrow \alpha = 0. \tag{4.46}$$

Furthermore, by Eq. (4.40c),

$$\cos\gamma - \cos\beta = 0 \Rightarrow \gamma = \beta, \tag{4.47a}$$

and by Eq. (4.40a),

$$\cos\beta - (\pi - \beta)\sin\beta + 1 = 0. \tag{4.47b}$$

Eq. (4.47b) has a single root on $\beta \in [0, \pi/2]$. This can be seen by the fact that its derivative is nonpositive and its second derivative is strictly positive on this interval. This root occurs near $\beta = 0.81047$. These values are substituted into $f(\alpha, \beta, \gamma)$ to give

$$\sup_{(\alpha,\beta,\gamma)\in[0,\pi/2]^3} ||\Gamma(P_{\alpha,\beta,\gamma}) - m|| \approx 1.1153. \tag{4.48}$$

Since the Euclidean radius of P is one, this implies the eccentricity of the Steiner centre is also approximately 1.1153. □

It follows that the bound on the eccentricity of Γ_2 is tight.

Corollary 4.21. *The two-dimensional Steiner centre, Γ_2, cannot guarantee λ-eccentricity for any $\lambda < 1.1153$.*

Proof. The result follows from the worst-case example derived in the proof of Thm. 4.20 which achieves eccentricity $\lambda \approx 1.1153$. □

Eccentricity of the Steiner Centre in Three Dimensions

The Steiner centre of a set of clients in $\mathbb{R}^3$ whose positions are coplanar coincides with the corresponding two-dimensional definition of the Steiner centre in the plane. Consequently, the eccentricity of Γ_3 is at least that of Γ_2, namely, 1.1153. In fact, in this section we show it is at least 1.2017. This lower bound on eccentricity is achieved by generalizing the worst-case example from $\mathbb{R}^2$ to $\mathbb{R}^3$. Although the bound is conjectured to be tight in $\mathbb{R}^3$, the techniques used to prove the upper bound in $\mathbb{R}^2$ do not immediately generalize to $\mathbb{R}^3$.

Theorem 4.22. *The three-dimensional Steiner centre, Γ_3, cannot guarantee λ-eccentricity for any λ less than*

$$\lambda < \frac{405 + 51\sqrt{17}}{512} \approx 1.2017. \tag{4.49}$$

Proof. Let S denote the unit sphere centred at the origin. Let $a = (0, 0, -1)$. Let T denote conic region with lower apex at point a and central axis that coincides with the z-axis. Let P denote a set of clients whose convex hull is the intersection of sphere S with conic region T. See Figs. 4.21A and 4.21B. That is, P is the convex hull of a spherical cap and its opposite pole. Let C denote the circle at which the boundaries of S and T intersect. Let b_1 denote a client on C and let d denote the pole opposite a, $d = (0, 0, 1)$. Let α denote angle $\angle dob_1$, where $o = (0, 0, 0)$. Assume $\alpha \in [0, \pi/3]$. In terms of α, $b_1 = (\sin\alpha\cos\theta, \sin\alpha\sin\theta, \cos\alpha)$, for some θ. Let b_2 denote the reflection of b_1 across the z-axis, $b_2 = (-\sin\alpha\cos\theta, -\sin\alpha\sin\theta, \cos\alpha)$. Finally, let $c_1 = (\sin\phi\cos\theta, \sin\phi\sin\theta, \cos\phi)$ and $c_2 = (-\sin\phi\cos\theta, -\sin\phi\sin\theta, \cos\phi)$ denote an arbitrary client on the the boundary of $CH(P) \cap S$ and its reflection across the z-axis (where $\phi \in [0, \alpha]$ is variable). See Fig. 4.21B.

For a given θ, Tab. 4.2 divides the range of integration, $\phi \in [0, \pi)$, into subranges for which specific points of P induce the extrema of $P_{\theta,\phi}$. See Fig. 4.21B.

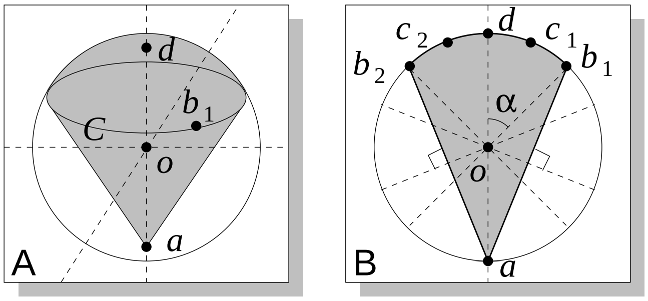

Figure 4.21: illustrations supporting Thm. 4.22

The Steiner centre of P lies at

$$\begin{aligned}
\Gamma_3(P) &= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \,\mathrm{mid}(P_{\theta,\phi})\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\left[\int_0^{\alpha} \sin\phi\cdot\frac{1}{2}\left(u_{\theta,\phi}\langle u_{\theta,\phi},c_1\rangle + u_{\theta,\phi}\langle u_{\theta,\phi},a\rangle\right)\ d\phi\right. \\
&+ \int_{\alpha}^{(\pi-\alpha)/2} \sin\phi\cdot\frac{1}{2}\left(u_{\theta,\phi}\langle u_{\theta,\phi},b_1\rangle + u_{\theta,\phi}\langle u_{\theta,\phi},a\rangle\right)\ d\phi \\
&+ \int_{(\pi-\alpha)/2}^{(\pi+\alpha)/2} \sin\phi\cdot\frac{1}{2}\left(u_{\theta,\phi}\langle u_{\theta,\phi},b_1\rangle + u_{\theta,\phi}\langle u_{\theta,\phi},b_2\rangle\right)\ d\phi \\
&+ \int_{(\pi+\alpha)/2}^{\pi-\alpha} \sin\phi\cdot\frac{1}{2}\left(u_{\theta,\phi}\langle u_{\theta,\phi},a\rangle + u_{\theta,\phi}\langle u_{\theta,\phi},b_2\rangle\right)\ d\phi \\
&+ \left.\int_{\pi-\alpha}^{\pi} \sin\phi\cdot\frac{1}{2}\left(u_{\theta,\phi}\langle u_{\theta,\phi},a\rangle + u_{\theta,\phi}\langle u_{\theta,\phi},c_2\rangle\right)\ d\phi\right]\ d\theta \\
&= \left(0,0,\cos^2\left(\frac{\alpha}{2}\right)\left[\cos^2\left(\frac{\alpha}{2}\right) - 1 + \sin\left(\frac{\alpha}{2}\right)\right]\right).
\end{aligned}$$

Observe that a is the furthest client from $\Gamma_3(P)$.

$$||\Gamma_3(P) - a|| = \Gamma_3(P)_z - a_z = 1 + \cos^2\left(\frac{\alpha}{2}\right)\left[\cos^2\left(\frac{\alpha}{2}\right) - 1 + \sin\left(\frac{\alpha}{2}\right)\right]. \quad (4.50)$$

interval of ϕ	$\arg\min_{p\in P}\langle p, u_{\theta,\phi}\rangle$	$\arg\max_{p\in P}\langle p, u_{\theta,\phi}\rangle$
$[0, \alpha]$	c_1	a
$[\alpha, (\pi-\alpha)/2]$	b_1	a
$[(\pi-\alpha)/2, (\pi+\alpha)/2]$	b_1	b_2
$[(\pi+\alpha)/2, \pi-\alpha]$	a	b_2
$[\pi-\alpha, \pi]$	a	c_2

Table 4.2: case analysis of extreme points in $\Gamma_3(P)$ in Thm. 4.22

The angle α that maximizes Eq. (4.50) is found by differentiating.

$$
\begin{aligned}
& & \frac{\partial}{\partial\alpha}||\Gamma_3(P) - a|| = 0, \\
\Rightarrow \quad & \frac{1}{2}\cos\left(\frac{\alpha}{2}\right)\left[2\sin\left(\frac{\alpha}{2}\right) + 3\cos^2\left(\frac{\alpha}{2}\right) - 4\sin\left(\frac{\alpha}{2}\right)\cos^2\left(\frac{\alpha}{2}\right) - 2\right] = 0, \\
\Rightarrow \quad & & \sin\left(\frac{\alpha}{2}\right) \in \left\{1, \frac{-1 \pm \sqrt{17}}{8}, \pi\right\}.
\end{aligned}
$$

Since $\alpha \in [0, \pi/3]$, therefore,

$$
\begin{aligned}
\sin\left(\frac{\alpha}{2}\right) &= \frac{-1+\sqrt{17}}{8} \\
\Rightarrow \quad \alpha &= 2\arcsin\left(\frac{-1+\sqrt{17}}{8}\right) \qquad (4.51) \\
\Rightarrow \quad &\approx 0.8021.
\end{aligned}
$$

Since P is contained within the unit sphere and both $(0, 0, 1)$ and $(0, 0, -1)$ are clients in P, the Euclidean radius of P is 1. Consequently, the eccentricity of Γ_3 is bounded from below by

$$
\begin{aligned}
\frac{\max_{p\in P}||\Gamma_3(P) - p||}{\max_{s\in P}||\Xi_3(P) - s||} &= \max_{p\in P}||\Gamma_3(P) - p|| \\
&= ||\Gamma_3(P) - a|| \\
&= 1 + \cos^2\left(\frac{\alpha}{2}\right)\left[\cos^2\left(\frac{\alpha}{2}\right) - 1 + \sin\left(\frac{\alpha}{2}\right)\right], \quad \text{by Eq. (4.50)}, \\
&= \frac{405 + 51\sqrt{17}}{512}, \qquad \text{by Eq. (4.51)}, \\
&\approx 1.2017. \qquad \square
\end{aligned}
$$

Since $\Gamma_3(P) \in CH(P)$, Lem. 4.4 implies an upper bound of 2 on the eccentricity of Γ_3. The lower bound from Thm. 4.22 is conjectured to be tight.

Conjecture 4.23. *The three-dimensional Steiner centre, Γ_3, has eccentricity*

$$\frac{405 + 51\sqrt{17}}{512} \approx 1.2017.$$

4.6.5 Steiner Centre: Maximum Velocity

In this section we derive the maximum velocity of the Steiner centre and show it is at most $4/\pi$ in $\mathbb{R}^2$ and $3/2$ in $\mathbb{R}^3$. In addition, we provide worst-case examples that realize each of these bounds.

Closely related to our definition of stability (see Sec. 3.5), Alt et al. [AAR97] define the *quality* of a reference point using Hausdorff distance and show that the quality of the Steiner centre is $4/\pi$ in $\mathbb{R}^2$ and $3/2$ in $\mathbb{R}^3$, matching our respective bounds on the maximum velocity of the Steiner centre. Our definition of stability lends itself better to the notion of a perturbation of a set of clients and allows us to exploit the inverse relationship between stability and maximum velocity (see Obs. 3.1). Our more inclusive definition allows for a broader set of centre functions to be considered (for example, the centre of mass). For completeness, we include our proofs of the maximum velocity of the Steiner centre.

Maximum Velocity of the Steiner Centre in Two Dimensions

We first bound the maximum velocity of the Steiner centre in two dimensions and provide an example that realizes this velocity, showing the bound is tight.

Theorem 4.24. *The two-dimensional mobile Steiner centre, Γ_2, has maximum velocity $4/\pi$.*

Proof. Choose any time interval T and any set of mobile clients P in $\mathbb{R}^2$ defined over T. Choose any $t_1, t_2 \in T$. Since Γ_2 is invariant under rotation, without loss of generality assume $\Gamma_2(P(t_1))_y = \Gamma_2(P(t_2))_y$. We bound the maximum velocity of Γ_2 from above:

$$\begin{aligned}
||\Gamma_2(P(t_1)) - \Gamma_2(P(t_2))|| &= |\Gamma_2(P(t_1))_x - \Gamma_2(P(t_2))_x| \\
&= \left| \frac{2}{\pi} \int_0^\pi \mathrm{mid}(P(t_1)_\theta)_x \; d\theta - \frac{2}{\pi} \int_0^\pi \mathrm{mid}(P(t_2)_\theta)_x \; d\theta \right| \\
&= \frac{2}{\pi} \left| \int_0^\pi \mathrm{mid}(P(t_1)_\theta)_x - \mathrm{mid}(P(t_2)_\theta)_x \; d\theta \right| \\
&\leq \frac{2}{\pi} \int_0^\pi |\,\mathrm{mid}(P(t_1)_\theta)_x - \mathrm{mid}(P(t_2)_\theta)_x| \; d\theta \\
&= \frac{2}{\pi} \int_0^\pi |\cos\theta| \cdot ||\,\mathrm{mid}(P(t_1)_\theta) - \mathrm{mid}(P(t_2)_\theta)|| \; d\theta \\
&\leq \frac{2}{\pi} \int_0^\pi |\cos\theta| \cdot \max_{p \in P} ||p(t_1) - p(t_2)|| \; d\theta, \qquad (4.52a)
\end{aligned}$$

since the velocity of $\mathrm{mid}(P_\theta)$ is at most the velocity of the endpoints of P_θ,

$$\leq \frac{2}{\pi} \int_0^\pi |\cos\theta| \cdot |t_1 - t_2| \; d\theta, \qquad (4.52b)$$

since client p has at most unit velocity,

$$\begin{aligned}
&= \frac{2}{\pi} |t_1 - t_2| \int_0^\pi |\cos\theta| \; d\theta \\
&= \frac{4}{\pi} |t_1 - t_2|. \qquad (4.52c)
\end{aligned}$$

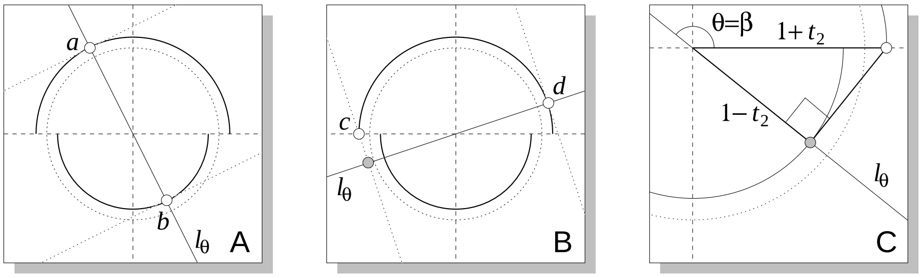

Figure 4.22: illustrations supporting Thm. 4.25

Therefore,

$$\forall t_1, t_2 \in T,\ ||\Gamma_2(P(t_1)) - \Gamma_2(P(t_2))|| \leq \frac{4}{\pi}|t_1 - t_2|, \tag{4.53}$$

for any set of mobile clients P in $\mathbb{R}^2$. □

The following example shows that the bound on maximum velocity is tight.

Theorem 4.25. *The two-dimensional mobile Steiner centre cannot guarantee relative velocity less than* $4/\pi$*.*

Proof. Let $t_1 = 0$. Let P be a set of clients with initial positions given by $P(t_1) = \{(\cos\theta, \sin\theta) \mid 0 \leq \theta < 2\pi\}$. That is, $P(t_1)$ is the set of points on the unit circle centred at the origin. Let $t_2 > 0$ be fixed and let function $f : \mathbb{R}^2 \to \mathbb{R}^2$ be defined by

$$f(p) = \begin{cases} (1+t_2)p & \text{if } p_y \geq 0 \\ (1-t_2)p & \text{if } p_y < 0 \end{cases}. \tag{4.54}$$

Let set $P(t_2) = f(P(t_1))$. Thus, $P(t_2)$ corresponds a perturbation of $P(t_1)$ such that points on or above the x-axis are scaled outward by t_2 and points below are scaled inward by t_2. See Fig. 4.22A.

For every $\theta \in [0, \pi]$,

$$\text{mid}(P(t_1)_\theta) = \frac{1}{2}[u_\theta + (-u_\theta)] = (0, 0). \tag{4.55}$$

Consequently, $\Gamma_2(P(t_1)) = (0, 0)$. The midpoint of $P(t_2)_\theta$ can be described by three cases. The simplest case occurs when one extremum of $P(t_2)_\theta$ lies on the outer semicircle and the second extremum lies on inner semicircle. For example, see points a and b in Fig. 4.22A. The second case occurs for angles θ near zero; in this case, one extremum of $P(t_2)_\theta$ is defined by the projection of one endpoint of the outer semicircle onto line l_θ whereas the other extremum remains on the outer semicircle. For example, see points c and d in Fig. 4.22B. The final case is analogous to the second case and occurs for angles θ near π. The angles θ for which a transition occurs from one case to the next are given by $\alpha = \arccos\left(\frac{1-t_2}{1+t_2}\right)$ and $\beta = \pi - \alpha = \arccos\left(\frac{t_2-1}{1+t_2}\right)$. See Fig. 4.22C.

The Steiner centre $\Gamma_2(P(t_2))$ is defined in terms of $\mathrm{mid}(P(t_2)_\theta)$. We examine the value $\mathrm{mid}(P(t_2)_\theta)$ over the three intervals, $[0,\alpha]$, $[\alpha,\beta]$, and $[\beta,\pi]$. For $\theta \in [0,\alpha]$,

$$\mathrm{mid}(P(t_2)_\theta) = \frac{1}{2}[u_\theta\langle -(1+t_2,0), u_\theta\rangle + u_\theta(1+t_2)] = \frac{u_\theta(1+t_2)(1-\cos\theta)}{2}. \tag{4.56a}$$

For $\theta \in [\alpha,\beta]$,

$$\mathrm{mid}(P(t_2)_\theta) = \frac{1}{2}[(1+t_2)u_\theta + (1-t_2)(-u_\theta)] = t_2 \cdot u_\theta. \tag{4.56b}$$

Finally, for $\theta \in [\beta,\pi]$,

$$\mathrm{mid}(P(t_2)_\theta) = \frac{1}{2}[u_\theta(1+t_2) + u_\theta\langle(1+t_2,0), u_\theta\rangle] = \frac{u_\theta(1+t_2)(1+\cos\theta)}{2}. \tag{4.56c}$$

The Steiner centre of set $P(t_2)$ is

$$\begin{aligned}
\Gamma_2(P(t_2)) =& \frac{2}{\pi}\int_0^\pi \mathrm{mid}(P(t_2)_\theta)\ d\theta \\
=& \frac{2}{\pi}\left[\int_0^\alpha \mathrm{mid}(P(t_2)_\theta)\ d\theta + \int_\alpha^\beta \mathrm{mid}(P(t_2)_\theta)\ d\theta + \int_\beta^\pi \mathrm{mid}(P(t_2)_\theta)\ d\theta\right] \\
=& \frac{2}{\pi}\left[\frac{1+t_2}{2}\int_0^\alpha u_\theta(1-\cos\theta)\ d\theta + \epsilon\int_\alpha^\beta u_\theta\ d\theta\right. \\
&\left. + \frac{1+t_2}{2}\int_\beta^\pi u_\theta(1+\cos\theta)\ d\theta\right],
\end{aligned}$$

by Eqs. (4.56a) through (4.56c),

$$\begin{aligned}
=& \frac{2}{\pi}\left[\left(\frac{\sqrt{t_2}(1+3t_2)}{2(1+t_2)} - \frac{1+t_2}{4}\arccos\alpha, \frac{t_2^2}{1+t_2}\right) + \left(0, \frac{2t_2(1-t_2)}{1+t_2}\right)\right. \\
&\left. + \left(-\frac{\sqrt{t_2}(1+3t_2)}{2(1+t_2)} + \frac{1+t_2}{4}\arccos\alpha, \frac{t_2^2}{1+t_2}\right)\right] \\
=& \left(0, \frac{4t_2}{\pi(1+t_2)}\right).
\end{aligned} \tag{4.57}$$

The maximum velocity of Γ_2 must be at least as great as the average velocity of P for any value of t_2. Therefore,

$$\begin{aligned}
v_{\max} &\geq \lim_{t_2 \to 0} \frac{||\Gamma_2(P(t_1)) - \Gamma_2(P(t_2))||}{|t_1 - t_2|} \\
&= \lim_{t_2 \to 0} \frac{1}{t_2} ||\Gamma_2(P(t_2))|| \\
&= \lim_{t_2 \to 0} \frac{1}{t_2} \left|\left| \left(0, \frac{4t_2}{\pi(1+t_2)}\right) \right|\right| \\
&= \lim_{t_2 \to 0} \frac{4}{\pi(1+t_2)} \\
&= \frac{4}{\pi},
\end{aligned} \tag{4.58}$$

where $P(t_1)$ and $P(t_2)$ are as described above. □

It follows from Thms. 4.24 and 4.25 that the maximum velocity of Γ_2 is $4/\pi$ and that this velocity is realizable.

Maximum Velocity of the Steiner Centre in Three Dimensions

Using a technique similar to that used in two dimensions, we bound the maximum velocity of the Steiner centre in three dimensions and provide an example that realizes this velocity, showing the bound is tight.

Theorem 4.26. *The three-dimensional mobile Steiner centre, Γ_3, has maximum velocity* $3/2$.

Proof. Choose any time interval T and any set of mobile clients P in $\mathbb{R}^3$ defined over T. Choose any $t_1, t_2 \in T$. Since Γ_3 is invariant under rotation, without loss of generality assume $\Gamma_3(P(t_1))_y = \Gamma_3(P(t_2))_y$ and $\Gamma_3(P(t_1))_z = \Gamma_3(P(t_2))_z$. We bound the maximum velocity of Γ_3 from above:

$$\begin{aligned}
&||\Gamma_3(P(t_1)) - \Gamma_3(P(t_2))|| \\
&\quad = |\Gamma_3(P(t_1))_x - \Gamma_3(P(t_2))_x| \\
&\quad = \left| \frac{3}{2\pi} \int_0^{\pi} \int_0^{\pi} \sin\phi \operatorname{mid}(P(t_1)_{\theta,\phi})_x \, d\phi \, d\theta \right. \\
&\qquad \left. - \frac{3}{2\pi} \int_0^{\pi} \int_0^{\pi} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})_x \, d\phi \, d\theta \right| \\
&\quad = \frac{3}{2\pi} \left| \int_0^{\pi} \int_0^{\pi} \sin\phi [\operatorname{mid}(P(t_1)_{\theta,\phi})_x - \operatorname{mid}(P(t_2)_{\theta,\phi})_x] \, d\phi \, d\theta \right| \\
&\quad \leq \frac{3}{2\pi} \int_0^{\pi} \int_0^{\pi} |\sin\phi [\operatorname{mid}(P(t_1)_{\theta,\phi})_x - \operatorname{mid}(P(t_2)_{\theta,\phi})_x]| \, d\phi \, d\theta \\
&\quad = \frac{3}{2\pi} \int_0^{\pi} \int_0^{\pi} \sin\phi |\operatorname{mid}(P(t_1)_{\theta,\phi})_x - \operatorname{mid}(P(t_2)_{\theta,\phi})_x| \, d\phi \, d\theta,
\end{aligned} \tag{4.59a}$$

since $\phi \in [0, \pi]$ implies $\sin\phi \geq 0$,

$$=\frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin^2\phi |\cos\theta| \cdot || \operatorname{mid}(P(t_1)_{\theta,\phi}) - \operatorname{mid}(P(t_2)_{\theta,\phi})|| \, d\phi \, d\theta$$

$$\leq\frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin^2\phi |\cos\theta| \cdot \max_{p\in P} ||p(t_1) - p(t_2)|| \, d\phi \, d\theta, \tag{4.59b}$$

since the velocity of $\operatorname{mid}(P_{\theta,\phi})$ is at most the velocity of the endpoints of $P_{\theta,\phi}$,

$$\leq\frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin^2\phi |\cos\theta| \cdot |t_1 - t_2| \, d\phi \, d\theta, \tag{4.59c}$$

since client p has at most unit velocity,

$$=\frac{3}{2\pi}|t_1 - t_2| \int_0^\pi \int_0^\pi \sin^2\phi |\cos\theta| \, d\phi \, d\theta$$

$$=\frac{3}{2}|t_1 - t_2|. \tag{4.59d}$$

Therefore,

$$\forall t_1, t_2 \in T, \; ||\Gamma_3(P(t_1)) - \Gamma_3(P(t_2))|| \leq \frac{3}{2}|t_1 - t_2|, \tag{4.60}$$

for any set of mobile clients P in $\mathbb{R}^3$. □

The following example shows that our bound on maximum velocity is tight.

Theorem 4.27. *The three-dimensional mobile Steiner centre cannot guarantee relative velocity less than* 3/2.

Proof. Let $P = \{u_{\theta,\phi} \mid 0 \leq \theta < 2\pi, \; 0 \leq \phi < \pi\}$ denote the set of clients on the unit sphere centred at the origin. Let $t_2 > 0$ be fixed and let function $f : \mathbb{R}^3 \to \mathbb{R}^3$ be defined by

$$f(a, b, c) = \begin{cases} (1 + t_2)(a, b, c) & \text{if } c \geq 0 \\ (1 - t_2)(a, b, c) & \text{if } c < 0 \end{cases}. \tag{4.61}$$

Let set $P(t_2) = f(P(t_1))$. $P(t_2)$ corresponds to a perturbation of $P(t_1)$ such that points on or above the xy-plane are scaled outward by t_2 and points below the xy-plane are scaled inward by t_2. See Fig. 4.23A. For every $\theta \in [0, \pi]$ and every $\phi \in [0, \pi]$,

$$\operatorname{mid}(P(t_1)_{\theta,\phi}) = \frac{1}{2}[u_{\theta,\phi} + (-u_{\theta,\phi})] = (0, 0, 0). \tag{4.62}$$

Consequently, $\Gamma_3(P(t_1)) = (0, 0, 0)$. The midpoint of $P(t_2)_{\theta,\phi}$ can be described by four cases. The first two cases occurs when one extremum of $P(t_2)_{\theta,\phi}$ lies on the outer hemisphere and the second extremum lies on the inner hemisphere

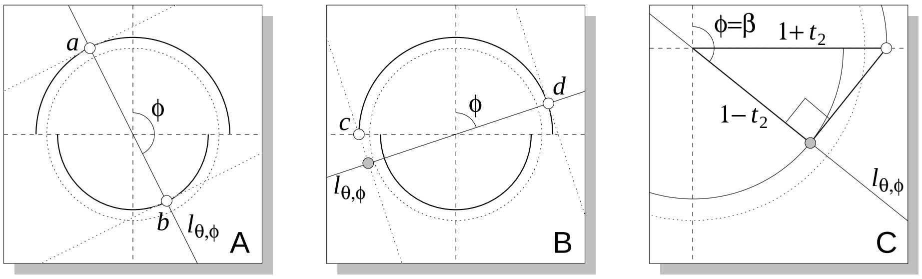

Figure 4.23: illustrations supporting Thm. 4.27

(ϕ near 0 or ϕ near π). For example, see points a and b in Fig. 4.23A. Observe that $a = -(1+t_2)u_{\theta,\phi}$ and $b = (1-t_2)u_{\theta,\phi}$. The third case occurs for angles ϕ near $\pi/2$ when $\phi < \pi/2$; in this case, one extremum of $P(t_2)_{\theta,\phi}$ is defined by the projection of one endpoint of the outer hemisphere onto line $l_{\theta,\phi}$ whereas the other extremum remains on the outer outer hemisphere. For example, see points c and d in Fig. 4.23B. Observe that $c = u_{\theta,\phi}\langle -(1+t_2)u_{\theta,\pi/2}, u_{\theta,\phi}\rangle = -(1+t_2)\sin\phi \cdot u_{\theta,\phi}$ and $d = (1+t_2)u_{\theta,\phi}$. The final case is analogous to the third case and occurs for angles ϕ near $\pi/2$ when $\phi \geq \pi/2$. The angles ϕ for which a transition occurs from one case to the next are given by $\alpha = \frac{\pi}{2} - \arccos\left(\frac{1-t_2}{1+t_2}\right)$, $\pi/2$, and $\beta = \pi - \alpha = \frac{\pi}{2} + \arccos\left(\frac{1-t_2}{1+t_2}\right)$. See Fig. 4.23C.

The Steiner centre $\Gamma_3(P(t_2))$ is defined in terms of $\mathrm{mid}(P(t_2)_{\theta,\phi})$. We examine the value $\mathrm{mid}(P(t_2)_{\theta,\phi})$ over the four intervals: $\phi \in [0,\alpha)$, $\phi \in [\alpha, \pi/2)$, $\phi \in [\pi/2, \beta)$, and $\phi \in [\beta, \pi)$. For $\phi \in [0,\alpha)$,

$$\begin{aligned}\mathrm{mid}(P(t_2)_{\theta,\phi}) &= \frac{1}{2}\Big[(1+t_2)u_{\theta,\phi} + (1-t_2)(-u_{\theta,\phi})\Big] \\ &= t_2 \cdot u_{\theta,\phi}. \end{aligned} \tag{4.63a}$$

For $\phi \in [\alpha, \pi/2)$,

$$\begin{aligned}\mathrm{mid}(P(t_2)_{\theta,\phi}) &= \frac{1}{2}\Big[u_{\theta,\phi}\langle -(1+t_2)u_{\theta,\pi/2}, u_{\theta,\phi}\rangle + u_{\theta,\phi}(1+t_2)\Big] \\ &= \frac{u_{\theta,\phi}(1+t_2)(1-\sin\phi)}{2}. \end{aligned} \tag{4.63b}$$

For $\phi \in [\pi/2, \beta]$,

$$\begin{aligned}\mathrm{mid}(P(t_2)_{\theta,\phi}) &= \frac{1}{2}\Big[-u_{\theta,\phi}(1+t_2) + u_{\theta,\phi}\langle (1+t_2)u_{\theta,\pi/2}, u_{\theta,\phi}\rangle\Big] \\ &= \frac{u_{\theta,\phi}(1+t_2)(\sin\phi - 1)}{2}. \end{aligned} \tag{4.63c}$$

Finally, for $\phi \in [\beta, \pi)$,

$$\begin{aligned}\mathrm{mid}(P(t_2)_{\theta,\phi}) &= \frac{1}{2}\Big[(1-t_2)u_{\theta,\phi} + (1+t_2)(-u_{\theta,\phi})\Big] \\ &= -t_2 \cdot u_{\theta,\phi}. \end{aligned} \tag{4.63d}$$

The Steiner centre of $P(t_2)$ is

$$\begin{aligned}
\Gamma_3(P(t_2)) &= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})\; d\phi\; d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\left[\int_0^{\alpha} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})\; d\phi + \int_{\alpha}^{\pi/2} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})\; d\phi\right. \\
&\quad \left. + \int_{\pi/2}^{\beta} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})\; d\phi + \int_{\beta}^{\pi} \sin\phi \operatorname{mid}(P(t_2)_{\theta,\phi})\; d\phi\right]\; d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\left[t_2\left(\int_0^{\alpha} \sin\phi\cdot u_{\theta,\phi}\; d\phi - \int_{\beta}^{\pi/2} \sin\phi\cdot u_{\theta,\phi}\; d\phi\right)\right. \\
&\quad \left. + \frac{1+t_2}{2}\left(\int_{\alpha}^{\pi/2} \sin\phi(1-\sin\phi)u_{\theta,\phi}\; d\phi - \int_{\pi/2}^{\beta} \sin\phi(1-\sin\phi)u_{\theta,\phi}\; d\phi\right)\right]\; d\theta,
\end{aligned}$$

by Eqs. (4.63a) through (4.63d),

$$= \left(0, 0, \frac{t_2(t_2^2+3)}{2(1+t_2)^2}\right). \tag{4.64}$$

The maximum velocity of Γ_3 must be at least as great as the average velocity of P for any value of t_2. Therefore,

$$\begin{aligned}
v_{\max} &\geq \lim_{t_2\to 0} \frac{||\Gamma_3(P(t_1)) - \Gamma_3(P(t_2))||}{|t_1 - t_2|} \\
&= \lim_{t_2\to 0} \frac{1}{t_2}||\Gamma_2(P(t_2))|| \\
&= \lim_{t_2\to 0} \frac{1}{t_2}\left|\left|\left(0, \frac{t_2(t_2^2+3)}{2(1+t_2)^2}\right)\right|\right| \\
&= \lim_{t_2\to 0} \frac{t_2^2+3}{2(1+t_2)^2} \\
&= \frac{3}{2},
\end{aligned} \tag{4.65}$$

where $P(t_1)$ and $P(t_2)$ are as described above. □

It follows from Thms. 4.26 and 4.27 that the maximum velocity of Γ_3 is $3/2$ and that this velocity is realizable.

4.7 Triangle Centres

An extensive set of triangle centres exists. These are not necessarily centre functions in the sense of defining a point that is somehow central to a triangle.

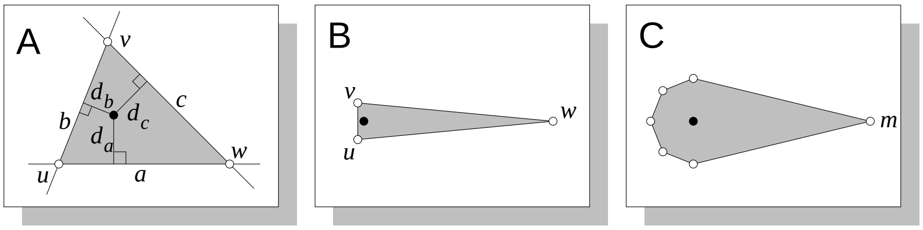

Figure 4.24: the least squares point of a triangle and its generalization to a polygon

Rather, the term *triangle centre* refers to a function on three clients in the plane that returns a fourth point, also in the plane. In this section we discuss three triangle centres that initially suggest themselves as potentially good centre functions: the least squares point, the incentre, and the orthocentre.

Kimberling provides a very extensive catalogue of triangle centres [Kim98, Kim]. At last check, the collection included 3053 triangle centres. Common triangle centres include the Euclidean 1-centre of a triangle (Kimberling triangle centre $X(3)$), the Euclidean 1-median of the vertices of triangle (Kimberling triangle centre $X(13)$) and the centre of mass of the vertices of a triangle (Kimberling triangle centre $X(2)$). Yiu [Yiu04] and Weisstein [Wei] provide overviews of common triangle centres.

We briefly examine some of the more common triangle centres, those whose definitions permit generalization to the mobile setting such that the position of the triangle centre potentially identifies a centre function with low eccentricity. For each we evaluate the maximum velocity and whether the definition generalizes to greater than three clients. Although the triangle centres we examine possess interesting geometric properties, it is straightforward to demonstrate that each results in either high eccentricity or discontinuous motion.

4.7.1 Least Squares Point

According to Honsberger, the least squares point, also known as Lemoine point, Grebe point, symmedian, and Kimberling triangle centre $X(6)$ [Kim98], is "one of the crown jewels of modern geometry" [Hon95, p. 53]. Winkler [Win79] mentions its significance in reducing numerical error in navigation. The least squares point of a triangle T in $\mathbb{R}^2$ is the unique point that minimizes the sum of the squares of the distances to the three edges of T. See Fig. 4.24A.

When triangle T becomes elongated (see Fig. 4.24B) its least squares point approaches the boundary of T and the eccentricity approaches 2. Although the least squares point has a natural generalization to a sets containing greater than three clients [Tha03], the point is easily seen to be discontinuous (and, consequently, has unbounded velocity) when clients of P join or leave the convex hull boundary of P. See Fig. 4.24C.

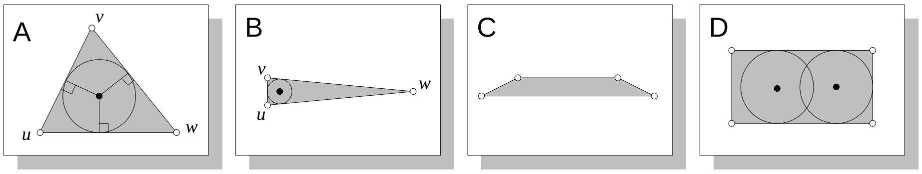

Figure 4.25: the incentre of a triangle and its generalization to a polygon

4.7.2 Incentre

The incentre of a triangle T in $\mathbb{R}^2$ is the unique point that lies an equal distance from each of the edges of T. Conversely to the Euclidean 1-centre which defines the centre of the smallest circle that encloses the convex hull of P, the incentre is the centre of largest circle contained within the convex hull of P (when $|P| = 3$). See Fig. 4.25A. The incentre is Kimberling triangle centre $X(1)$ [Kim98].

Like the least squares point, the eccentricity of the incentre is easily seen to be 2. See Fig. 4.25B. When $|P| > 3$, two generalizations are possible. The first is to locate a point that is equidistant from the edges of the convex hull of P. Such a point may not exist (for example, see the trapezoid in Fig. 4.25C). The second is to locate the centre of the largest circle contained within the convex hull of P. This second point is not unique (see Fig. 4.25D).

Finally, observe that with either of these definitions, as was the case with the least squares point, any change in the composition of the convex hull may cause a discontinuity in the position of the incentre of P.

4.7.3 Orthocentre

The orthocentre of a triangle T in $\mathbb{R}^2$ is the intersection of the altitudes of T, where the altitude of an edge is the shortest line segment from that edge to the opposite vertex. The notion of altitudes does not suggest any natural generalization to greater than three clients. The orthocentre is Kimberling triangle centre $X(4)$ [Kim98].

4.8 Convex Combinations

A new centre function can be defined by a convex combination of existing centre functions. This section examines the eccentricity and maximum velocity of the resulting convex combination in terms of the eccentricities and maximum velocities of the component centre functions.

Bereg et al. [BBKS06] refer to a *mixing strategy* in $\mathbb{R}^2$ in which a centre function Υ_2 is defined as a convex combination of the rectilinear 1-centre and the centre of mass. That is, $\Upsilon_2(P) = kR_2(P) + (1-k)C_2(P)$ for some $k \in [0, 1]$. They show the corresponding bound on the eccentricity of Υ_2, $\lambda \leq k(1+\sqrt{2})/2 + (1-k)2$ [BBKS06]. In Thms. 4.30 and 4.33 and Cors. 4.31 and 4.34, we generalize this notion to any convex combination of centre functions in any dimension and derive bounds on the eccentricity and maximum velocity of the resulting convex combination.

4.8.1 Euclidean Norm of a Convex Combination

We first examine the Euclidean norm of the convex combination of two points in $\mathbb{R}^d$.

Lemma 4.28. *Let a and b be arbitrary points in $\mathbb{R}^d$. Let $k \in [0,1]$. The Euclidean norms of a, b, and $ka + (1-k)b$ are related by*

$$||ka + (1-k)b|| \leq k||a|| + (1-k)||b||. \tag{4.66}$$

Proof.

$$\begin{aligned}
||ka + (1-k)b||^2 &= \langle ka + (1-k)b, ka + (1-k)b \rangle \\
&= k^2\langle a, a\rangle + (1-k)^2\langle b, b\rangle + 2k(1-k)\langle a, b\rangle \\
&= k^2\langle a, a\rangle + (1-k)^2\langle b, b\rangle + 2k(1-k)||a|| \cdot ||b|| \cos\theta \\
&\leq k^2\langle a, a\rangle + (1-k)^2\langle b, b\rangle + 2k(1-k)||a|| \cdot ||b|| \\
&= k^2||a||^2 + (1-k)^2||b||^2 + 2k(1-k)||a|| \cdot ||b|| \\
&= \Big[k||a|| + (1-k)||b||\Big]^2,
\end{aligned}$$

where θ is the angle between vectors a and b. Since the Euclidean norm is non-negative and $k \in [0,1]$, this gives

$$||ka + (1-k)b|| \leq k||a|| + (1-k)||b||.$$

□

Corollary 4.29. *Let c, d, and p be arbitrary points in $\mathbb{R}^d$. Let $k \in [0,1]$. The Euclidean norms of c, d, and $kc + (1-k)d$ are related by*

$$||p - [kc + (1-k)d]|| \leq k||p - c|| + (1-k)||p - d||. \tag{4.67}$$

Proof. Since Lem. 4.28 holds for all a and b in $\mathbb{R}^d$, it must hold for $a = p - c$ and $b = p - d$. Eq. (4.67) follows upon substitution of these values into Eq. (4.66), giving a generalization that corresponds to invariance under translation. See Fig. 4.26A. □

4.8.2 Convex Combinations: Eccentricity

We show that the eccentricity of a convex combination can be bounded by the corresponding convex combination of the eccentricities of its component centre functions.

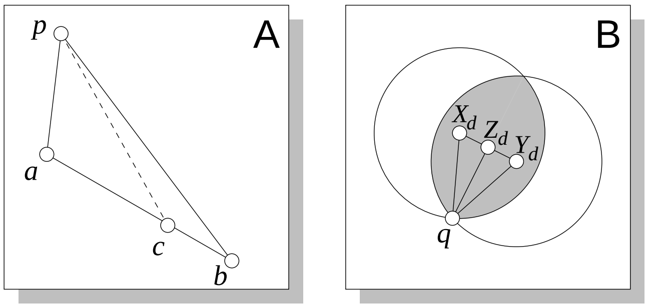

Figure 4.26: illustrations supporting Cor. 4.29 and Thm. 4.30.

Theorem 4.30. *Let X_d and Y_d be centre functions in $\mathbb{R}^d$ that are invariant under translation. Let $k \in [0,1]$. Let $Z_d(P) = kX_d(P) + (1-k)Y_d(P)$ define a third centre function. Let λ_X and λ_Y denote the respective eccentricities of X_d and Y_d. Centre function Z_d is λ_Z-eccentric, where*

$$\lambda_Z = k\lambda_X + (1-k)\lambda_Y.$$

Proof. Choose any set of client positions P in $\mathbb{R}^d$.

$$\begin{aligned}\frac{\max\limits_{p\in P}||Z_d(P)-p||}{\max\limits_{r\in P}||\Xi_d(P)-r||} =& \frac{\max\limits_{p\in P}||[kX_d(P)+(1-k)Y_d(P)]-p||}{\max\limits_{r\in P}||\Xi_d(P)-r||}\\ \leq& \frac{\max\limits_{p\in P}\Big[k||X_d(P)-p||+(1-k)||Y_d(P)-p||\Big]}{\max\limits_{r\in P}||\Xi_d(P)-r||},\end{aligned}$$

by Cor. (4.29),

$$\begin{aligned}\leq& \frac{k\left[\max\limits_{s\in P}||X_d(P)-s||\right]+(1-k)\left[\max\limits_{t\in P}||Y_d(P)-t||\right]}{\max\limits_{r\in P}||\Xi_d(P)-r||}\\ =& k\frac{\max\limits_{s\in P}||X_d(P)-s||}{\max\limits_{r\in P}||\Xi_d(P)-r||}+(1-k)\frac{\max\limits_{t\in P}||Y_d(P)-t||}{\max\limits_{r\in P}||\Xi_d(P)-r||}\\ \leq& k\lambda_X+(1-k)\lambda_Y,\end{aligned}$$

by Def. 3.5. Therefore,

$$\forall P,\ \max_{p\in P}||Z_d(P)-p|| \leq [k\lambda_X+(1-k)\lambda_Y]\max_{r\in P}||\Xi_d(P)-r||. \qquad \square$$

Corollary 4.31. *Given $n \in \mathbb{Z}^+$, for every $1 \leq i \leq n$, let $X_d^i : \mathscr{P}(\mathbb{R}^d) \to \mathbb{R}^d$ denote a centre function that is invariant under translation and let $k_i \in [0,1]$*

such that $\sum_{i=1}^{n} k_i = 1$*. Let* Z_d *denote the centre function defined by*

$$Z_d(P) = \sum_{i=1}^{n} k_i X_d^i(P).$$

For each i*, let* λ_i *denote the eccentricity of* X_d^i*.* Z_d *is* λ_Z*-eccentric where*

$$\lambda_Z = \sum_{i=1}^{n} k_i \lambda_i.$$

Proof. The result follows by induction on n using Thm. 4.30. $\square$

4.8.3 Convex Combinations: Maximum Velocity

The analogous relationship for maximum velocity follows from the definition of velocity. That is, the maximum velocity of a convex combination can be bounded by the corresponding convex combination of the maximum velocities of its component centre functions. Since the derivation is independent of the optimization function and depends only on velocity, the result applies to any approximation function (as opposed to only centre functions).

We first reproduce a related result by Bereg et al.:

Observation 4.32 (Bereg et al. 2006 [BBKS06]). *Let* $\alpha_1, \alpha_2, \ldots, \alpha_n$ *be fixed numbers such that* $\alpha_i \geq 0$ *for all* i *and* $\sum_{i=1}^{n} \alpha_i = 1$*. If all of the sites,* $s_1, \ldots, s_n$*, move with velocity at most* 1*, then the point* p *defined as the convex combination* $\sum_{i=1}^{n} \alpha_i s_i$ *of the sites moves with velocity at most* 1*.*

Obs. 4.32 implies that if approximation functions X_d^1 through X_d^n have velocity at most one, then any convex combination of X_d^1 through X_d^n also has maximum velocity at most 1. In Thm. 4.33 and Cor. 4.34 we demonstrate that a more general relationship holds for convex combinations.

Theorem 4.33. *Let* X_d *and* Y_d *be centre functions in* $\mathbb{R}^d$ *that are invariant under translation. Let* $k \in [0, 1]$*. Let* $Z_d(P) = kX_d(P) + (1-k)Y_d(P)$ *define a third approximation function. Let* v_X*,* v_Y*, and* v_Z *denote the respective maximum velocities of* X_d*,* Y_d*, and* Z_d*.* Z_d *has maximum velocity at most* v_Z*, where*

$$v_Z = kv_X + (1-k)v_Y. \tag{4.68}$$

Proof. Choose any time interval T and any set of mobile clients P in $\mathbb{R}^d$ defined over T. Choose any $t_1, t_2 \in T$.

$$\begin{aligned}
&\frac{||Z[P(t_1)] - Z[P(t_2)]||}{|t_1 - t_2|} \\
&\quad = \frac{||kX[P(t_1)] + (1-k)Y[P(t_1)] - kX[P(t_2)] - (1-k)Y[P(t_2)]||}{|t_1 - t_2|} \\
&\quad \leq k\frac{||X[P(t_1)] - X[P(t_2)]||}{|t_1 - t_2|} + (1-k)\frac{||Y[P(t_1)] - Y[P(t_2)]||}{|t_1 - t_2|},
\end{aligned}$$

by Lem. 4.28,

$$\leq kv_X + (1-k)v_Y,$$

by Def. 3.4. Therefore,

$$\forall t_1, t_2 \in T,\ ||Z[P(t_1)] - Z[P(t_2)]|| \leq [kv_X + (1-k)v_Y]|t_1 - t_2|. \qquad \square$$

Cor. 4.34 is the generalization of Thm. 4.33 analogous to Cor. 4.31.

Corollary 4.34. *Given $n \in \mathbb{Z}^+$, for every $1 \leq i \leq n$, let X_d^i denote an approximation function in $\mathbb{R}^d$ that is invariant under translation and let $k_i \in [0,1]$ such that $\sum_{i=1}^{n} k_i = 1$. Let Z_d denote an approximation function defined by*

$$Z_d(P) = \sum_{i=1}^{n} k_i X_d^i(P).$$

Let v_i denote the maximum velocity of X_d^i and let v_Z denote the maximum velocity of Z_d. Z_d has maximum velocity at most v_Z, where

$$v_Z = \sum_{i=1}^{n} k_i v_i.$$

Proof. The result follows by induction on n using Thm. 4.33. $\square$

4.8.4 Using Convex Combinations to Compare Centre Functions

Thms. 4.30 and 4.33 allow us to evaluate the significance of a centre function's eccentricity and maximum velocity. That is, if we have three centre functions Υ_d^1, Υ_d^2, and Υ_d^3 such that their respective maximum velocities are sorted in increasing order, we can define a fourth centre function Υ_d^4 by a convex combination of Υ_d^1 and Υ_d^3 such that the maximum velocity of Υ_d^4 matches that of Υ_d^2. Comparing the eccentricities of Υ_d^2 and Υ_d^4 helps determine whether Υ_d^2 is beneficial as a centre function. In Sec. 4.9, we use this technique to compare the Steiner centre against a convex combination of the centre of mass and the rectilinear 1-centre.

4.9 Evaluation

In Secs. 4.4 through 4.8 we explored candidate functions for defining bounded-velocity approximations of the mobile Euclidean 1-centre. In this section we compare these various centre functions against each other, in terms of eccentricity, maximum velocity, independence of non-extreme client positions, invariance under similarity transformations, and consistency of definition across dimensions.

Those centre functions which we identified as good bounded-velocity approximations of the mobile Euclidean 1-centre are the rectilinear 1-centre, the centre of mass, and the Steiner centre. To these we add the centre function defined by a client p in P (see Sec. 4.3.1), three triangle centres (the least squares point, the incentre, and the orthocentre), as well as a convex combination of the rectilinear 1-centre and the centre of mass which we discuss below.

Rectilinear 1-Centre

In Sec. 4.4 we examined the rectilinear 1-centre, R_d. In $\mathbb{R}^d$, we showed a tight bound of $\frac{1}{2}(1 + \sqrt{d})$ on the eccentricity of R_d and we referred to a result of Bereg et al. [BBKS06] showing a tight bound of $\sqrt{d}$ on its maximum velocity. As mentioned in Sec. 2.5.1, R_d is not invariant under rotation or reflection. It is, however, invariant under translation and scaling. The definition of R_d is consistent across dimensions. Finally, $R_d(P)$ is induced by the extreme points of P.

Centre of Mass

In Sec. 4.5 we examined the centre of mass, C_d. In $\mathbb{R}^d$, we referred to results of Bereg et al. [BBKS06] showing tight bounds of 2 on the eccentricity of C_d and 1 on its maximum velocity. We observed that $C_d(P)$ is not defined exclusively in terms of extreme points of P; rather, $C_d(P)$ assigns uniform weight to all clients in P. As mentioned in Sec. 2.5.2. C_d is invariant under similarity transformations and its definition is consistent across dimensions.

Steiner Centre

In Sec. 4.6 we examined two definitions of Steiner centre, Γ_d, first by Gaussian weights and then by projection. In $\mathbb{R}^2$, we showed tight bounds of approximately 1.1153 on the eccentricity of Γ_2 and $4/\pi$ on its maximum velocity. In $\mathbb{R}^3$, we showed a lower bound of approximately 1.2017 and an upper bound of 2 on the eccentricity of Γ_3 as well as a tight bound of $3/2$ on its maximum velocity. The definition of Γ_d by Gaussian weights assigns a weight of zero to all clients in the interior of the convex hull; thus, $\Gamma_d(P)$ is defined exclusively by the extreme points of P. We observed that the definition of Γ_d is consistent across dimensions. Finally, we referred to a result of Shephard [She68] who showed the invariance of Γ_d under similarity transformations.

Triangle Centres

In Sec. 4.7 we evaluated three triangle centres (the least squares point, the incentre, and the orthocentre) each of which fails to define a suitable candidate for a bounded-velocity centre function due either to discontinuity in its motions or lack of a natural definition for greater than three clients.

Convex Combinations

In Sec. 4.8 we presented a discussion of convex combinations of centre functions, including results on bounding the eccentricity and maximum velocity of a convex combination in terms of the eccentricities and maximum velocities of its component centre functions. We now examine specific convex combinations involving the centre functions described above.

Any convex combination Υ_d that includes the Euclidean 1-centre as a component of non-zero weight, regardless of the combination of centre functions that completes the definition of Υ_d, results in unbounded velocity for Υ_d. Consequently, we consider only convex combination whose composition does not include Ξ_d.

The maximum velocity of R_d is greater than that of the Steiner centre while the maximum velocity of C_d less than that of the Steiner centre. Thus, we consider the convex combination of R_d and C_d given by $kR_d(P)+(1-k)C_d(P)$ for some $k \in [0,1]$. We select values of k such that the maximum velocity of $kR_d(P)+(1-k)C_d(P)$ is equal to $4/\pi$ in $\mathbb{R}^2$ and equal to $3/2$ in $\mathbb{R}^3$, allowing us to compare the convex combination directly against the Steiner centre for a fixed maximum velocity. The specific values of k are given by solving for k_2 and k_3 in

$$\begin{aligned} & k_2\sqrt{2}+(1-k_2) = \frac{4}{\pi}, && \text{in } \mathbb{R}^2, \\ \text{and} \quad & k_3\sqrt{3}+(1-k_3) = \frac{3}{2}, && \text{in } \mathbb{R}^3. \end{aligned}$$

Solving for these values gives $k_2 = (4-\pi)/[\pi(\sqrt{2}-1)] \approx 0.6597$ and $k_3 = 1/[2(\sqrt{3}-1)] \approx 0.6831$. The corresponding bounds on the approximation factors are

$$\begin{aligned} & k_2\frac{1+\sqrt{2}}{2}+2(1-k_2) = \frac{(4+3\pi)\sqrt{2}-12-\pi}{2\pi(\sqrt{2}-1)} && \approx 1.4770, \text{ in } \mathbb{R}^2, \\ \text{and} \quad & k_3\frac{1+\sqrt{3}}{2}+2(1-k_3) = \frac{9\sqrt{3}-11}{4(\sqrt{3}-1)} && \approx 1.5670, \text{ in } \mathbb{R}^3. \end{aligned}$$

Finally, since R_d is neither invariant under rotation nor reflection, it follows that these properties do not hold for any convex combination whose composition includes R_d. Similarly, since the definition of $C_d(P)$ does not depend exclusively on the positions of the extreme points of P, it follows that this property does not hold for any convex combination whose composition includes C_d.

Comparison of Centre Functions

The values for the eccentricity and maximum velocity of these various centre functions are displayed in Tab. 4.3 for $\mathbb{R}^2$ and in Tab. 4.4 for $\mathbb{R}^3$.

First, observe that the convex combination $k_2R_2+(1-k_2)C_2$ is more eccentric than Γ_2 for the same maximum velocity. That is, the Steiner centre provides a better approximation of the Euclidean 1-centre than does the corresponding convex combination of C_2 and R_2, even though both have the same

centre function		eccentricity	maximum velocity
Euclidean 1-centre	Ξ_2	$\lambda = 1$	$v_{\max} = \infty$
single client $p \in P$	p	$\lambda = 2$	$v_{\max} = 1$
centre of mass	C_2	$\lambda = 2$	$v_{\max} = 1$
rectilinear 1-centre	R_2	$\lambda = (1+\sqrt{2})/2 \approx 1.2071$	$v_{\max} = \sqrt{2} \approx 1.4142$
Steiner centre	Γ_2	$\lambda \approx 1.1153$	$v_{\max} = 4/\pi \approx 1.2732$
convex combination of R_2 and C_2		$\lambda \leq 1.4770$	$v_{\max} \leq 4/\pi \approx 1.2732$

Table 4.3: comparing centre functions in $\mathbb{R}^2$

centre function		eccentricity	maximum velocity
Euclidean 1-centre	Ξ_3	$\lambda = 1$	$v_{\max} = \infty$
single client $p \in P$	p	$\lambda = 2$	$v_{\max} = 1$
centre of mass	C_3	$\lambda = 2$	$v_{\max} = 1$
rectilinear 1-centre	R_3	$\lambda = (1+\sqrt{3})/2 \approx 1.3660$	$v_{\max} = \sqrt{3} \approx 1.7320$
Steiner centre	Γ_3	$1.2017 \leq \lambda \leq 2$	$v_{\max} = 1.5$
convex combination of R_3 and C_3		$\lambda \leq 1.5670$	$v_{\max} \leq 1.5$

Table 4.4: comparing centre functions in $\mathbb{R}^3$

maximum velocity. As for the rectilinear 1-centre, the Steiner centre has both lower eccentricity and lower maximum velocity in $\mathbb{R}^2$. Since the lowest possible eccentricity is one, the difference in the eccentricities of R_2 and Γ_2 in $\mathbb{R}^2$ from 1.2071 to 1.1153 corresponds to a relative improvement of 44.3%. Similarly, since any bounded-velocity approximation must have velocity at least one, the difference in the maximum velocities of R_2 and Γ_2 in $\mathbb{R}^2$ from 1.4142 to 1.2732 corresponds to a relative improvement of 34.0%.

Bereg et al. [BBKS06] suggested a centre function that always moves toward the current position of the Euclidean 1-centre (see Sec. 4.3.2). If such a centre function maintains eccentricity at most 1.1153 then the corresponding upper bound on its maximum velocity is approximately 4.7771 (the value of Expr. 4.7 when $\epsilon = 0.1153$). This value is far greater than the maximum velocity of the Steiner centre.

Experimentation suggests that the Steiner centre performs well not only in the worst case but also in the average case. Empirical evidence is provided in Sec. 8.4.2 in the form of test results from simulations of sets of 6 clients and 16 clients for which the eccentricities and velocities of the Euclidean 1-centre, centre of mass, rectilinear 1-centre, and Steiner centre of a set of mobile clients are measured over 10000 time units. See Figs. 8.9 and 8.10.

All centre functions mentioned in this section are defined consistently across dimensions; that is, the position of $\Upsilon_d(P)$ coincides with $\Upsilon_{d-1}(P)$ when the positions of clients in P lie in a $(d-1)$-dimensional flat. All centre functions mentioned are invariant under similarity transformations except for the rectilinear 1-centre and its convex combinations which are not invariant under rotation or reflection. Finally, the centre of mass and its convex combinations are the only centre functions whose definitions depend on non-extreme points.

Chapter 5

Mobile Euclidean 1-Median

5.1 Introduction

5.1.1 Chapter Objectives

Chapter 5 examines various bounded-velocity approximations to the Euclidean 1-median. Our exploration of approximation functions (referred to as *median functions* in the context of the Euclidean 1-median) leads us to consider the centre of mass, the rectilinear 1-median, and the projection median (a new median function which we now introduce), along with convex and linear combinations of these, for which we examine the maximum velocity and approximation factor. Kinetic algorithms for maintaining these various mobile median functions are discussed in Ch. 8; for now we focus on their respective qualities as approximation functions.

The main contribution of this chapter is the definition of the projection median, which we show reasonably balances the conflicting goals of approximating the Euclidean median sum while maintaining a low maximum velocity. Summaries of the chapter's significant results and their implications are found in Secs. 5.1.2 and 5.9.

5.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Properties of the Mobile Euclidean 1-Median (Sec. 5.2)
Sec. 5.2 briefly examines additional properties of the mobile Euclidean 1-median. Specifically, we show that the motion of the mobile Euclidean 1-median is discontinuous, which implies that its velocity is unbounded.

Comparison Measures (Sec. 5.3)
Sec. 5.3 expands on the measures of approximation factor and maximum velocity and explores bounds on their relationship in terms specific to the approximation of the Euclidean 1-median. Additional natural properties of medians are also considered.

Rectilinear 1-Median (Sec. 5.4)
Sec. 5.4 analyzes the properties of the mobile rectilinear 1-median, S_d, in terms of its approximation of the Euclidean 1-median. The rectilinear 1-median of

P minimizes the sum of the rectilinear (ℓ_1) distances between itself and clients in P, suggesting it as a candidate for approximating the Euclidean 1-median. In particular, we generalize a result of Bereg et al. [BBKS00] to show that in $\mathbb{R}^d$, the rectilinear 1-median has approximation factor of $\sqrt{d}$. We show that the rectilinear 1-median cannot guarantee an approximation factor lower than $(1+\sqrt{d-1})/\sqrt{d}$. In $\mathbb{R}^2$, the upper and lower bounds coincide to give a tight bound of $\sqrt{2}$. The bounds diverge in $\mathbb{R}^d$ for $d \geq 3$. Unlike the Euclidean 1-median whose motion is discontinuous, the motion of the rectilinear 1-median is both continuous and has bounded velocity. We show a tight bound of $\sqrt{d}$ on the maximum velocity of the rectilinear 1-median in $\mathbb{R}^d$.

Centre of Mass (Sec. 5.5)
Sec. 5.5 analyzes the properties of the mobile centre of mass, C_d, in terms of its approximation of the Euclidean 1-median. The centre of mass is a point that minimizes the sum of the squared Euclidean distances between itself and clients in P, suggesting it as a candidate for approximating the Euclidean 1-median. We show that the centre of mass provides a $(2-\frac{2}{n})$-approximation in $\mathbb{R}^d$, where $n = |P|$. We refer to results of Bereg et al. [BBKS06] mentioned in Sec. 4.5.2 showing that the centre of mass has maximum velocity 1 in $\mathbb{R}^d$.

Projection Median (Sec. 5.6)
Sec. 5.6 introduces the projection median, Π_d, as a new median function defined in terms of projection of client positions onto a line through the origin and integration of the one-dimensional median of the projected point set over all such lines. A significant portion of Ch. 5 consists of the derivations of the approximation factor and maximum velocity of the projection median in two and three dimensions. In particular, we show an upper bound of $4/\pi$ and a lower bound of $\sqrt{4/\pi^2+1}$ on the approximation factor of the projection median in $\mathbb{R}^2$. It follows that the lower bound also applies in higher dimensions. We show tight bounds on the maximum velocity of the projection median of $4/\pi$ in $\mathbb{R}^2$ and $3/2$ in $\mathbb{R}^3$.

Convex Combinations (Sec. 5.7)
Sec. 5.7 examines convex combinations of median functions. In particular, a convex combination of a set of median functions defines a new median function whose maximum velocity and approximation factor can be bounded in terms of the maximum velocities and approximation factors of the component median functions.

Gaussian Median (Sec. 5.8)
Exploiting the success of the Steiner centre at defining a good approximation of the Euclidean 1-centre, Sec. 5.8 introduces the Gaussian median, G_d, a normalized weighted mean of the client positions, using a weighting function inversely related to the Gaussian weight defined in Sec. 4.6. We show the Gaussian median can be defined by a linear combination of the Steiner centre and the centre

of mass. Also, we show that for $|P| \leq 4$, the Gaussian median coincides with the projection median in $\mathbb{R}^2$ and $\mathbb{R}^3$. We derive a lower bound of 3/2 on the approximation factor of G_2. Finally, we establish upper bounds on the maximum velocity of the Gaussian median of $3 + 8/\pi$ in $\mathbb{R}^2$ and 6 in $\mathbb{R}^3$.

Evaluation (Sec. 5.9)

Sec. 5.9 summarizes the results derived in Ch. 5 by comparison of the projection median, the rectilinear 1-median, the centre of mass, and convex combinations of these in terms of their approximation of the Euclidean 1-median. The primary measures for evaluating the quality of each median function are approximation factor and maximum velocity (inversely related to stability) but also include consideration of whether each median function generalizes to higher dimensions and whether it preserves various properties of invariance and consistency.

5.2 Properties of the Mobile Euclidean 1-Median

This section briefly explores the discontinuity of the mobile Euclidean 1-median. Refer to Sec. 2.4.2 for the static definition of the Euclidean 1-median.

Theorem 5.1. *The mobile Euclidean 1-median is discontinuous in two or more dimensions.*

Proof. Choose any $\epsilon > 0$. Let $P(0) = \{(0,0),(0,0),(1,0),(1,\epsilon)\}$ and let $P(\epsilon) = \{(0,0),(0,\epsilon),(1,0),(1,0)\}$. Since $|P| = 4$ and two clients of $P(0)$ coincide at $(0,0)$, $M_2(P(0)) = (0,0)$ [KM97]. Similarly, $M_2(P(\epsilon)) = (1,0)$. Since

$$\forall \delta > 0, \exists t \in (0,\delta),\ ||\Upsilon_2(P(0)) - \Upsilon_2(P(t))|| = 1, \tag{5.1}$$

it follows that the Euclidean 1-median is not continuous by Def. 3.3. □

As an immediate consequence of Thm. 5.1, no fixed upper bound exists on the maximum velocity of the Euclidean 1-median. Furthermore, no bounded-velocity facility function can follow a trajectory that remains within an arbitrarily-small ϵ-neighbourhood around $M_d(P(t))$ for $d \geq 2$.

In one dimension, the median is both continuous and moves with at most unit velocity:

Observation 5.2. *The one-dimensional mobile Euclidean 1-median, M_d, has maximum velocity 1. Furthermore, this velocity is realizable.*

Proof. The median is defined as the $\lceil |P|/2 \rceil$nd largest client in P when $|P|$ is odd and as the the midpoint of the $(|P|/2)$nd and $(|P|/2+1)$st largest clients in P when $|P|$ is even. Each of these moves with at most unit velocity. Furthermore, this property is maintained when two or more clients coincide or cross. The upper bound follows from Obs. 4.32. The bound is realized when all clients move with unit velocity in a common direction. □

5.3 Comparison Measures

This section expands on the comparison measures defined in Ch. 3 in terms specific to median functions. We examine bounds on the relationship between approximation factor and maximum velocity and enumerate additional properties naturally associated with notions of medians.

5.3.1 Bounds on Approximation Factor and Maximum Velocity

We are motivated to define alternative median functions that approximate the Euclidean 1-median in the sense that they come close to minimizing Expr. (2.9) and yet have lower maximum velocity. Thus, we examine median functions with the twofold objective of minimizing both approximation factor and maximum velocity.

As discussed in Ch. 4, Bereg et al. [BBKS06] show that any point located within the convex hull of the set of client positions defines a centre function with approximation factor (eccentricity) at most two. This property holds for centre functions, since the Euclidean 1-centre of a set of clients P is induced by a subset of the extreme points of P. The corresponding property is not true of the Euclidean 1-median. Consequently, a median function whose position is defined to coincide with the position of a particular client cannot guarantee any fixed approximation factor:

Observation 5.3. *Let P denote a finite multiset of clients in $\mathbb{R}^d$. Let median function $\Upsilon_d(P(t)) = p(t)$, where p is a fixed client in P. Median function Υ_d cannot guarantee any fixed approximation factor of the Euclidean 1-median.*

Proof. Let $n-1$ clients be located at the origin and let client p be located at a distance d away from the origin, for some $d > 0$. The Euclidean 1-median of P lies at the origin and the Euclidean median sum is d. The sum of the distances from $\Upsilon_d(P)$ to the clients of P is $(n-1)d$. The corresponding approximation factor is $n-1$. As such, Υ_d cannot guarantee any fixed approximation factor (that is independent of $|P|$). □

Consequently, we consider median functions with approximation factors in the range $[1, \infty)$ and maximum velocities in the range $[1, \infty)$.

5.3.2 Approximation Factor as a Function of Maximum Velocity

Of course, as was the case for centre functions, the maximum velocity and approximation factor of a median function are correlated. As we show formally in Thm. 5.4, no median function can ensure any fixed maximum velocity while also guaranteeing an arbitrarily-close approximation of the Euclidean median sum. That is, no λ-approximation of the Euclidean 1-median is possible for an arbitrary $\lambda > 1$ and a fixed $v_{\max} > 1$ that is independent of λ.

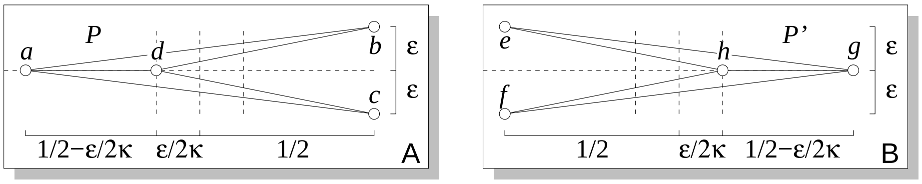

Figure 5.1: illustration in support of Thm. 5.4

Theorem 5.4. *For every $v_{\max} > 0$, if Υ_d is a median function with maximum velocity $v_{\max}$, then there exists some $\lambda_0 > 1$ such that Υ_d cannot guarantee an approximation factor less than λ_0.*

Proof. Choose any $v_{\max} > 0$, any $\epsilon \in (0, 1/v_{\max})$, and any median function Υ_d with maximum velocity $v_{\max}$. Let $P(0) = \{(0,0), (0,0), (1,\epsilon), (1,-\epsilon)\}$ and let $P(\epsilon) = \{(0,\epsilon), (0,-\epsilon), (1,0), (1,0)\}$. Let $d = (1/2 - \epsilon v_{\max}/2, 0)$ and let $h = (1/2 + \epsilon v_{\max}/2, 0)$. See Figs. 5.1A and 5.1B.

Since $|P| = 4$ and two points of P coincide at $(0,0)$, $M_2(P(0)) = (0,0)$ [KM97]. Similarly, $M_2(P(\epsilon)) = (1,0)$. The Euclidean median sum of $P(0)$ (and, by symmetry, $P(\epsilon)$) is $2\sqrt{1+\epsilon^2}$.

By Def. 3.4,

$$||\Upsilon_d(P(0)) - \Upsilon_d(P(\epsilon))|| \leq \epsilon v_{\max}. \tag{5.2}$$

Let p_x denote the x-coordinate of p, for any point p in $\mathbb{R}^d$. Consequently, either $\Upsilon_d(P(0))_x \geq d_x$ or $\Upsilon_d(P(\epsilon))_x \leq h_x$. Without loss of generality, assume $\Upsilon_d(P(0))_x \geq d_x$.

It is straightforward to show that for any point d', where $d'_x \geq d_x$, $\sum_{p \in P(0)} ||d'_x - p|| \geq \sum_{p \in P(0)} ||d_x - p||$. Therefore,

$$\begin{aligned}\sum_{p \in P(0)} ||\Upsilon_d(P(0)) - p|| \geq &\sum_{p \in P(0)} ||d - p|| \\ =&2\left(\frac{1}{2} - \frac{\epsilon v_{\max}}{2}\right) + 2\sqrt{\epsilon^2 + \left(\frac{1}{2} + \frac{\epsilon v_{\max}}{2}\right)^2}. \end{aligned} \tag{5.3}$$

By Def. 3.5, if Υ_d is a λ-approximation, then

$$\begin{aligned}\lambda \geq &\frac{\sum_{p \in P} ||\Upsilon_d(P) - p||}{\sum_{q \in P} ||M_2(P) - q||} \\ \geq &\frac{\frac{1}{2} - \frac{\epsilon v_{\max}}{2} + \sqrt{\epsilon^2 + \left(\frac{1}{2} + \frac{\epsilon v_{\max}}{2}\right)^2}}{\sqrt{1+\epsilon^2}}, \qquad \text{by Eq. (5.3).}\end{aligned} \tag{5.4}$$

Let λ_1 denote the righthand value in Eq. (5.4). It is straightforward to show that $\lambda_1 > 1$ for any $\epsilon \in (0, 1/v_{\max})$. Therefore, for any $\lambda_0 \in (1, \lambda_1)$, Υ_d is not a λ_0-approximation of the Euclidean 1-median. □

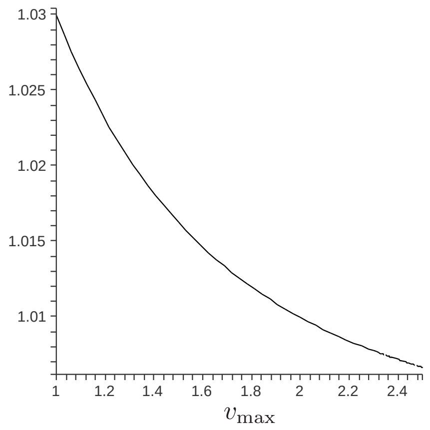

Figure 5.2: Eq. (5.4): a lower bound on λ as a function of $v_{\max}$

Although it is possible to solve for a positive value of ϵ that maximizes the final expression in Eq. (5.4) in terms of an arbitrary $v_{\max}$, the resulting ϵ has a complex representation involving numerous cubic roots and terms up to $O(v_{\max}^9)$ for which no simple representation was found. Not surprisingly, upon substituting this maximizing value for ϵ into Eq. (5.4), the resulting expression is even more complex. The resulting function is displayed in Fig. 5.2. As such, these expressions are not reproduced here; Eq. (5.4) suffices to prove that the approximation factor increases as the velocity decreases.

Reducing the approximation factor increases the maximum velocity and vice-versa. The challenge lies in understanding the trade-off between the degree of approximation factor (in the range $[1, \infty)$) and the maximum velocity (also in the range $[1, \infty)$). Thm. 5.4 implies that no bounded-velocity λ-approximate median function is possible for an arbitrary $\lambda \geq 1$ and a fixed $v_{\max}$ that is independent of λ.

5.3.3 Additional Properties of Median Functions

Although approximation factor and maximum velocity define the principal measures by which we evaluate median functions, the following define additional natural properties for a median function Υ_d, both of which are properties exhibited by the Euclidean 1-median:

1. $\Upsilon_d(P)$ should be invariant under rotation, uniform scaling, reflection, and translation.

2. If P resides in a $(d-i)$-flat in $\mathbb{R}^d$, then the d-dimensional definition, $\Upsilon_d(P)$, should coincide with the $(d-i)$-dimensional definition, $\Upsilon_{d-i}(P)$.

For each median function Υ_d examined, we evaluate the fitness of Υ_d and compare it against other median functions primarily in terms its approximation factor and maximum velocity. In addition, to further understand the behaviour of Υ_d, we also determine whether each of the properties listed above also holds for Υ_d.

5.4 Rectilinear 1-Median

This section discusses properties of the rectilinear 1-median as an approximation to the mobile Euclidean 1-median. Refer to Sec. 2.5.1 for the definition of the rectilinear k-median.

Recall that the rectilinear 1-median of P, denoted $S_d(P)$, is a point that minimizes the sum of the rectilinear (ℓ_1) distances to the positions of clients in P. Given that its maximum velocity is bounded, this property suggests the rectilinear 1-median as a natural candidate for providing an approximation of the Euclidean 1-median.

The invariance of the rectilinear 1-median under similarity transformations is straightforward to demonstrate. Also, the definition of the rectilinear 1-median is defined consistently across dimensions.

5.4.1 Rectilinear 1-Median: Approximation Factor

Bereg et al. [BBKS00] show that the rectilinear 1-median provides a $\sqrt{2}$-approximation of the Euclidean 1-median in $\mathbb{R}^2$. We show this bound is tight in the following example.

Observation 5.5. *The two-dimensional rectilinear 1-median, S_2, cannot guarantee a λ-approximation of the Euclidean 1-median for any $\lambda < \sqrt{2}$.*

Proof. Let $2k$ clients lie at $(1, 0)$, let $k+1$ points lie at $(0, 1)$, and let $k+1$ clients lie at $(0, -1)$. See Fig. 5.3. The unique rectilinear 1-median of P lies at $(0, 0)$. Since the clients of P are not collinear, the position of the Euclidean 1-median of P is also unique. Consequently, by the symmetry of P and the invariance of $M_2(P)$ under reflection, $M_2(P)$ must lie on the x-axis. The Euclidean median sum of P is

$$f(x) = 2k|1 - x| + 2(k + 1)\sqrt{x^2 + 1}, \tag{5.5}$$

where $x = M_2(P)_x$. To find the value of x that minimizes Eq. (5.5), we set its derivative to zero. Since $x \in [0, 1]$, we replace $|1 - x|$ by $(1 - x)$:

$$\begin{aligned}
& & \frac{\partial f(x)}{\partial x} &= 0 \\
\Leftrightarrow & & -2k + \frac{2x(k+1)}{\sqrt{x^2+1}} &= 0 \\
\Leftrightarrow & & x &= \frac{k}{\sqrt{2k+1}} \qquad (5.6)
\end{aligned}$$

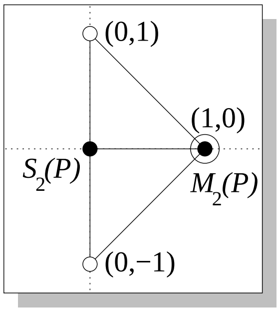

Figure 5.3: example realizing the approximation factor of the rectilinear 1-median in $\mathbb{R}^2$

It is straightforward to confirm that Eq. (5.6) is increasing for $k > 0$ by examining its derivative with respect to k. Consequently, Eq. (5.6) implies that $x = 1$ if and only if $k = 1 + \sqrt{2}$. Therefore, for any $k \geq 3$, Eq. (5.5) is minimized by some $x \geq 1$. Since $M_2(P)_x \in [0, 1]$, for any $k \geq 3$, Eq. (5.5) is minimized at $M_2(P)_x = 1$. Consequently, the Euclidean 1-median of P lies at $(1, 0)$.

We obtain the following lower bound on the approximation factor of S_2:

$$\begin{aligned}\lambda \geq & \lim_{k\to\infty} \frac{\sum_{p\in P} ||p - S_2(P)||}{\sum_{q\in P} ||q - M_2(P)||} \\ = & \lim_{k\to\infty} \frac{2(k+1) + 2k}{2(k+1)\sqrt{2}} \\ = & \lim_{k\to\infty} \frac{2k+1}{(k+1)\sqrt{2}} \\ = & \sqrt{2}. \end{aligned}$$

□

We generalize the result of Bereg et al. to $\mathbb{R}^d$ using an analogous proof.

Theorem 5.6. *The d-dimensional rectilinear 1-median, S_d, provides a $\sqrt{d}$-approximation of the Euclidean 1-median.*

Proof. By Def. 2.10, the rectilinear 1-median of P is a point $S_d(P)$ that minimizes

$$\sum_{p\in P} ||p - S_d(P)||_1. \tag{5.7}$$

Similarly, by Def. 2.7, the Euclidean 1-median of P is a point $M_d(P)$ that minimizes

$$\sum_{p\in P} ||p - M_d(P)||. \tag{5.8}$$

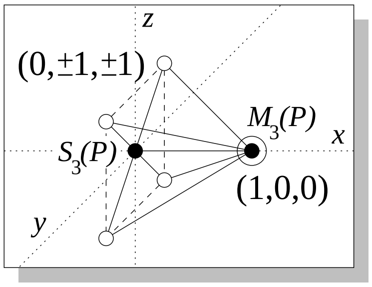

Figure 5.4: example realizing the lower bound of the approximation factor on the rectilinear 1-median in $\mathbb{R}^3$

For any $x \in \mathbb{R}^d$, $||x|| \leq ||x||_1 \leq \sqrt{d}||x||$. Consequently,

$$\begin{aligned}\sum_{p\in P} ||p - M_d(P)|| &\leq \sum_{p\in P} ||p - M_d(P)||_1 \\ &\leq \sum_{p\in P} ||p - S_d(P)||_1 \\ &\leq \sqrt{d}\sum_{p\in P} ||p - S_d(P)||. \qquad \square\end{aligned}$$

By generalizing the example from the proof of Obs. 5.5 to $\mathbb{R}^d$, we show a lower bound of $(1+\sqrt{d-1})/\sqrt{d}$ on the approximation factor of the rectilinear 1-median in $\mathbb{R}^d$. When $d = 2$, the lower bound matches the upper bound of $\sqrt{2}$ shown in Thm. 5.6.

Theorem 5.7. *The d-dimensional rectilinear 1-median, S_d, cannot guarantee a λ-approximation of the Euclidean 1-median for any*

$$\lambda < \frac{1+\sqrt{d-1}}{\sqrt{d}}. \tag{5.9}$$

Proof. Case 1. Assume $d = 1$. In $\mathbb{R}$, $S_1(P)$ and $M_1(P)$ coincide. In this case, Eq. (5.9) holds since the rectilinear median sum and the Euclidean median sum are equal.

Case 2. Assume $d \geq 2$. Let $k+1$ clients lie at $(0, \pm 1, \dots, \pm 1)$ for all 2^{d-1} combinations of ± 1. Let $2^{d-1}k$ clients lie at $(1, 0, \dots, 0)$. See Fig. 5.4. The unique rectilinear 1-median of P lies at $(0, \dots, 0)$. Since the clients of P are not collinear, the position of the Euclidean 1-median of P is also unique. Consequently, by the symmetry of P and the invariance of $M_d(P)$ under reflection, $M_d(P)$ must lie on the x-axis; that is $M_d(P) = (M_d(P)_x, 0, \dots, 0)$. The Euclidean median sum of P is

$$2^{d-1}k|1 - M_d(P)_x| + 2^{d-1}(k+1)\sqrt{M_d(P)_x^2 + d - 1}. \tag{5.10}$$

Using an argument analogous to that used in the proof of Obs. 5.5, it is straightforward to show that Expr. (5.10) is minimized at $M_d(P)_x = 1$ for any $k \geq (1+\sqrt{d})/(d-1)$. Therefore, the Euclidean 1-median of P lies at $(1, 0, \ldots, 0)$ if $k \geq (1+\sqrt{d})/(d-1)$. We obtain the following lower bound on the approximation factor of S_d:

$$\begin{aligned}
\lambda &\geq \lim_{k\to\infty} \frac{\sum_{p\in P} ||p - S_d(P)||}{\sum_{q\in P} ||q - M_d(P)||} \\
&= \lim_{k\to\infty} \frac{2^{d-1}(k+1)\sqrt{d-1} + 2^{d-1}k}{2^{d-1}(k+1)\sqrt{d}} \\
&= \lim_{k\to\infty} \frac{k(\sqrt{d-1}+1) + \sqrt{d-1}}{(k+1)\sqrt{d}} \\
&= \frac{1+\sqrt{d-1}}{\sqrt{d}}.
\end{aligned}$$

□

Although equal for $d = 2$, the lower bound of Thm. 5.7 and the upper bound of Thm. 5.6 diverge as k increases. No tight bound on the approximation factor of S_d is known for $d \geq 3$.

5.4.2 Rectilinear 1-Median: Maximum Velocity

Bereg et al. [BBKS00], observe that the relative velocity of the rectilinear 1-median of a set of mobile points in $\mathbb{R}^2$ is at most $\sqrt{2}$. Furthermore, this bound is tight. This observation is straightforward to generalize to $\mathbb{R}^d$.

Observation 5.8. *The d-dimensional mobile rectilinear 1-median, S_d, has maximum velocity $\sqrt{d}$.*

Proof. As shown by Bajaj [Baj84] and as observed by Bereg et al. [BBKS00], the coordinates of $S_d(P)$ correspond to the one-dimensional 1-median of the clients of P with respect to each dimension. In the worst case, distinct clients induce the 1-median of P in each dimension, where the corresponding client moves with unit velocity in a direction parallel to the dimensional axis, resulting in a velocity of $\sqrt{d}$ of the rectilinear 1-median. See the example in Fig. 5.5 for $d = 2$. □

Corollary 5.9. *The d-dimensional rectilinear 1-median, S_d, cannot guarantee relative velocity less than $\sqrt{d}$.*

Proof. The example described in the proof of Obs. 5.8 is realizable. □

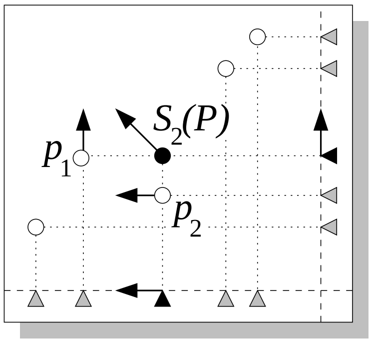

Figure 5.5: The rectilinear 1-median has maximum velocity $\sqrt{2}$ in $\mathbb{R}^2$ when clients p_1 and p_2 move with unit velocity in directions parallel to the y and x-axes, respectively.

5.5 Centre of Mass

This section discusses properties of the centre of mass as an approximation to the mobile Euclidean 1-median. Refer to Sec. 2.5.2 for the definition of the centre of mass.

Recall that the centre of mass of P, denoted $C_d(P)$, is a point that minimizes the sum of the square distances to the positions of clients in P [Sch73, Wes93]. See Sec. 2.5.2 for a derivation of this result. Given the low upper bound on its velocity, this property suggests the centre of mass as a natural candidate for providing an approximation of the Euclidean 1-median.

The invariance of the centre of mass under similarity transformations is straightforward to demonstrate. Also, the definition of the centre of mass is defined consistently across dimensions.

5.5.1 Centre of Mass: Approximation Factor

Bereg et al. [BBKS00] show that the centre of mass has approximation factor $\sqrt{2}(2-2/n)$. Using different techniques, we now establish a tight bound of $2-2/n$ on the approximation factor of the centre of mass in Thm. 5.12. Necessary to the proof of Thm. 5.12 is Lem. 5.11 which shows that for any finite multiset of clients P, if some client $a \neq M_d(P)$ is moved to coincide with $M_d(P)$, then the Euclidean 1-median of the new multiset P' remains unchanged. Lems. 5.10 and 5.11 and Thm. 5.12 refer to the following definitions for P, a, x, and n. Let P denote a finite multiset of clients in $\mathbb{R}^d$ such that client $a \neq M_d(P)$ for some $a \in P$. Let $a' = M_d(P)$, let $P' = (P - \{a\}) \cup \{a'\}$, let $x = ||a - a'||$, and let $n = |P|$. See Fig. 5.6.

Lemma 5.10. *Point $M_d(P)$ is a Euclidean 1-median of P'.*

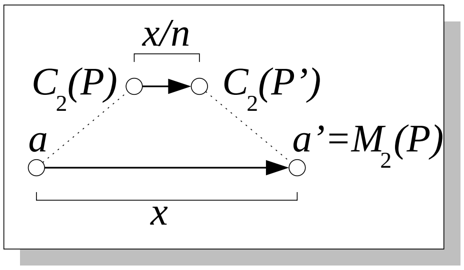

Figure 5.6: illustration in support of Lem. 5.11

Proof. Assume $M_d(P)$ is not a Euclidean 1-median of P'. Thus,

$$\sum_{p \in P'} ||p - M_d(P')|| < \sum_{p \in P'} ||p - M_d(P)||. \tag{5.11}$$

Therefore,

$$\begin{aligned}
\sum_{p \in P} ||p - M_d(P')|| &= ||a - M_d(P')|| + \sum_{p \in P - \{a\}} ||p - M_d(P')|| \\
&\leq x + ||a' - M_d(P')|| + \sum_{p \in P - \{a\}} ||p - M_d(P')|| \\
&= x + ||a' - M_d(P')|| + \sum_{p \in P' - \{a\}} ||p - M_d(P')|| \\
&= x + \sum_{p \in P'} ||p - M_d(P')|| \\
&< x + \sum_{p \in P'} ||p - M_d(P)||, \text{ by our assumption,} \\
&= \sum_{p \in P} ||p - M_d(P)||.
\end{aligned} \tag{5.12}$$

Thus, $M_d(P)$ did not minimize $\sum_{p \in P} ||p - M_d(P)||$. Consequently, $M_d(P)$ cannot be a median of P, deriving a contradiction. Therefore $M_d(P') = M_d(P)$. □

Also necessary to the proof of Thm. 5.12 is Lem. 5.11 which relates the sum of the distances between $C_d(P)$ and the clients of P to the corresponding value for P'.

Lemma 5.11.

$$\sum_{p \in P} ||p - C_d(P)|| - \sum_{p \in P'} ||p - C_d(P')|| \leq 2x\left(1 - \frac{1}{n}\right). \tag{5.13}$$

Proof. Since all clients remain static except for client a, $C_d(P) - C_d(P') = \frac{1}{n}(a - a')$. See Fig. 5.6. Consequently, the distance from a to the centre of mass changes by at most $\pm(x - x/n)$. For each of the $n-1$ points in $P - \{a\}$, the corresponding distance changes by at most $\pm x/n$. The result follows. □

Theorem 5.12. *The d-dimensional centre of mass, C_d, provides a $(2-2/n)$-approximation of the Euclidean 1-median.*

Proof. Let a, a', x, and P' be as defined in Lem. 5.10. Let $m = \sum_{p \in P} ||p - M_d(P)||$ and let $c = \sum_{p \in P} ||p - C_d(P)||$. Let m' and c' denote the corresponding values for P'. Assume P is a multiset of clients in $\mathbb{R}^d$ that maximizes the approximation factor of C_d such that $c > m(2 - 2/n)$. Observe that a client $a \neq M_d(P)$ must exist under this assumption, otherwise all clients of P would be collocated with $M_d(P)$ and $C_d(P)$. Thus,

$$c > m\left(2 - \frac{2}{n}\right),$$

$$\Rightarrow \qquad cx - cm > 2mx\left(1 - \frac{1}{n}\right) - cm,$$

since $a \neq a'$ and, consequently, $x = ||a - a'|| > 0$,

$$\Rightarrow \qquad c(x-m) > m\left[2x\left(1 - \frac{1}{n}\right) - c\right],$$

$$\Rightarrow \qquad c(m-x) < m\left[c - 2x\left(1 - \frac{1}{n}\right)\right],$$

$$\Rightarrow \qquad c(m-x) < mc',$$

by Lem. 5.11,

$$\Rightarrow \qquad cm' < mc',$$

since $M_d(P) = M_d(P')$ by Lem. 5.10 and, consequently, $m = m' + x$,

$$\Rightarrow \qquad \frac{c}{m} < \frac{c'}{m'}, \tag{5.14}$$

since m and m' are sums of non-negative terms.

This contradicts our assumption that P maximizes the approximation factor of C_d. Therefore, $c \leq m(2 - 2/n)$. That is, for all finite multisets P,

$$\sum_{p \in P} ||p - C_d(P)|| \leq \left(2 - \frac{2}{n}\right) \sum_{p \in P} ||p - M_d(P)||,$$

where $n = |P|$. □

This bound is tight:

Corollary 5.13. *The d-dimensional centre of mass, C_d, cannot guarantee a λ-approximation of the Euclidean 1-median for any $\lambda < 2$.*

Proof. The approximation bound $2-2/n$ is realized in any dimension d by $n-1$ clients located at the origin and a single client located away from the origin. The bound $2-2/n$ approaches 2 as n increases. □

5.5.2 Centre of Mass: Maximum Velocity

The velocity of the centre of mass is independent of which facility function is being approximated; refer to Sec. 4.5.2 where the velocity of the centre of mass is examined in the context of centre functions.

In brief, Bereg et al. [BBKS00] show that any function defined by a convex combination of a set of mobile points moves with maximum relative velocity at most one (Obs. 4.32). Consequently, the centre of mass has maximum velocity one (Cor. 4.15). The bound is tight, as demonstrated by any translation of the positions of clients in P at unit velocity (Obs. 4.16).

5.6 Projection Median

The definition of the Euclidean 1-median is a natural generalization of the one-dimensional median to higher dimensions. Expr. (2.9), however, suggests other possible generalizations. One possibility is to project clients onto a line through the origin, to find the one-dimensional median of the projection, and to integrate these one-dimensional medians for all lines through the origin. Using this idea, which derives from the definition of the Steiner centre by projection, we define a median function of a set of mobile clients P in $\mathbb{R}^d$, which we call the projection median of P and denote $\Pi_d(P)$.

In Sec. 5.6 we show that the projection median has a low upper bound on its maximum velocity and that it guarantees a low approximation factor of the Euclidean 1-median. In addition, we establish the invariance of the projection median under similarity transformations and demonstrate that its definition is consistent across dimensions.

5.6.1 Definition

Projection Median Definition in Two Dimensions

Let l_θ denote the line through the origin parallel to the unit vector $u_\theta = (\cos\theta, \sin\theta)$. Expressed in slope-intercept form, l_θ is the line $y = \tan\theta\, x$. Given a multiset of clients P in $\mathbb{R}^2$ and an angle $\theta \in [0, \pi)$, let P_θ denote the multiset defined by the projection of P onto line l_θ. See Fig. 5.7A. That is,

$$P_\theta = \{u_\theta \langle p, u_\theta\rangle \mid p \in P\}. \tag{5.15}$$

The median of P_θ is simply the Euclidean 1-median of P_θ,

$$\mathrm{med}(P_\theta) = M_2(P_\theta). \tag{5.16}$$

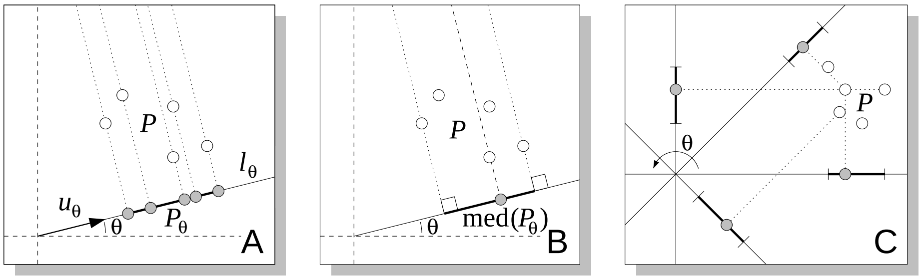

Figure 5.7: defining the projection median

Let $p \in \mathbb{R}^2$ be any fixed client. The average over all projections of p onto lines l_θ is

$$\frac{1}{\pi}\int_0^\pi u_\theta \langle p, u_\theta\rangle \; d\theta = \frac{p}{2},$$

suggesting an additional factor of 2 is necessary in the following definition for a median function:

Definition 5.1. *The* **two-dimensional projection median** *of a finite multiset P in $\mathbb{R}^2$ is*

$$\Pi_2(P) = \frac{2}{\pi}\int_0^\pi \mathrm{med}(P_\theta) \; d\theta, \tag{5.17}$$

where $\mathrm{med}(P_\theta)$ is the median of the projection of P onto the line $y = \tan\theta \; x$.

If $|P|$ is even, then P_θ may not have a unique median. In this case, let $\mathrm{med}(P_\theta)$ denote the midpoint of the region of points on l_θ that define medians of P_θ.

The formulation of the projection median displays some resemblance to the Steiner centre, which can be expressed similarly to Eq. (5.17) in $\mathbb{R}^2$ by replacing $\mathrm{med}(P_\theta)$ with $\frac{u_\theta}{2}(\min_{p\in P}\langle p, u_\theta\rangle + \max_{q \in P}\langle q, u_\theta\rangle)$, the centre of P_θ. See Sec. 4.6.2.

Although we focus on median functions defined over finite multisets, the definition of the projection median is easily generalized to a continuous definition, where the set of clients P is a bounded region in $\mathbb{R}^d$ with an associated density function. In this case, $\mathrm{med}(P_\theta)$ corresponds to the one-dimensional continuous median.

The definition of the projection median can be interpreted in terms of the rectilinear 1-median. Let d_ϕ denote the ℓ_1 norm relative to a rotation by ϕ of the reference axis. That is, $d_\phi(x) = ||f_\phi(x)||_1$, where f_ϕ is a clockwise rotation about the origin by ϕ. Let $S_\phi(P) = f_\phi^{-1}(S_2(f_\phi(P)))$ denote the rectilinear 1-median with respect to norm d_ϕ. We show the following relationship between the projection median of P and the rectilinear 1-medians of P relative to rotation:

Lemma 5.14.

$$\Pi_2(P) = \frac{2}{\pi}\int_0^{\pi/2} S_\theta(P) \; d\theta, \tag{5.18}$$

where $S_\phi(P)$ denotes the rectilinear 1-median relative to a rotation by ϕ of the reference axis.

Proof. The rectilinear 1-median can be defined in terms of the Euclidean 1-medians of its respective x- and y-coordinates. The corresponding property also holds for S_ϕ:

$$S_\phi(P) = \mathrm{med}(P_\phi) + \mathrm{med}(P_{\phi+\pi/2}).$$

Consequently,

$$\begin{aligned}
\Pi_2(P) =& \frac{2}{\pi}\int_0^\pi \mathrm{med}(P_\theta)\ d\theta \\
=& \frac{2}{\pi}\left[\int_0^{\pi/2} \mathrm{med}(P_\theta)\ d\theta + \int_{\pi/2}^{\pi} \mathrm{med}(P_\theta)\ d\theta\right] \\
=& \frac{2}{\pi}\int_0^{\pi/2} \mathrm{med}(P_\theta) + \mathrm{med}(P_{\theta+\pi/2})\ d\theta \\
=& \frac{2}{\pi}\int_0^{\pi/2} S_\theta(P)\ d\theta.
\end{aligned}$$

□

See Sec. 8.3.1 for a description of algorithms for maintaining the mobile Gaussian median in two dimensions. See Sec. 8.3.2 for a description of an algorithm for finding the static projection median in two dimensions and Sec. 8.3.3 for a description of algorithms for maintaining the mobile projection median in two dimensions.

Projection Median Definition in Three Dimensions

The definition of the projection median has a natural generalization to three dimensions, analogous to the generalization of the Steiner centre's definition by projection from two to three dimensions (see Sec. 4.6.2).

In three dimensions, we express the projection median in terms of spherical coordinates. Let $l_{\theta,\phi}$ be the line through the origin parallel to the unit vector $u_{\theta,\phi} = (\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi)$. Let $P_{\theta,\phi}$ and $\mathrm{med}(P_{\theta,\phi})$ be the natural generalizations of P_θ and $\mathrm{med}(P_\theta)$ to spherical coordinates in $\mathbb{R}^3$, respectively. Thus,

$$P_{\theta,\phi} = \{u_{\theta,\phi}\langle p, u_{\theta,\phi}\rangle \mid p \in P\} \tag{5.19}$$

and

$$\mathrm{med}(P_{\theta,\phi}) = \Xi_3(P_{\theta,\phi}). \tag{5.20}$$

Let $p \in \mathbb{R}^3$ be any fixed point. The average over all projections of p onto all lines $l_{\theta,\phi}$ is

$$\frac{1}{2\pi}\int_0^\pi\int_0^\pi \sin\phi \cdot p\ d\phi\ d\theta = \frac{p}{3}. \tag{5.21}$$

The factor $\sin\phi$ is required for uniform integration over points on a sphere. The factor $1/2\pi$ normalizes over the range of the integration as shown by $\int_0^\pi\int_0^\pi \sin\phi\ d\phi\ d\theta = 2\pi$. Adding a factor of three returns p instead of $p/3$, suggesting the following definition for a median function:

Definition 5.2. *Let P in $\mathbb{R}^3$ be a bounded and finite set of clients. The* **three-dimensional projection median** *of P is*

$$\Pi_3(P) = \frac{3}{2\pi} \int_0^{\pi} \int_0^{\pi} \sin\phi \, \mathrm{med}(P_{\theta,\phi}) \; d\phi \; d\theta, \tag{5.22}$$

where $\mathrm{med}(P_{\theta,\phi})$ *is the median of the projection of P onto the line through the origin parallel to* $u_{\theta,\phi} = (\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi)$.

This definition is the natural analogue of the two-dimensional projection median, expressing $\Pi_3(P)$ as the average median over all projections of P onto lines $l_{\theta,\phi}$.

Projection Median Definition in Higher Dimensions

The structures of the definition of the projection median in both $\mathbb{R}^2$ and $\mathbb{R}^3$ is motivated by the one-dimensional definition of the Euclidean 1-median. As such, it makes sense to define the one-dimensional projection median such that

$$\Pi_1(P) = M_1(P).$$

The definition of the projection median has a natural generalization to $\mathbb{R}^d$. We simply integrate the median of the projection of P onto all lines through the origin and normalize by the volume of the unit hypersphere.

Definition 5.3. *Given a fixed $d \in \mathbb{N}$ and a finite multiset of clients P in $\mathbb{R}^d$, the* **d-dimensional projection median** *of P is*

$$\Pi_d(P) = \frac{d \displaystyle\int_{u\in\mathbb{S}^{d-1}} \mathrm{med}(P_u) \; du}{\displaystyle\int_{u\in\mathbb{S}^{d-1}} 1 \; du}, \tag{5.23}$$

where $\mathbb{S}^{d-1} = \{x \in \mathbb{R}^d \mid ||x|| = 1\}$ *is the unit hypersphere and* $\mathrm{med}(P_u)$ *is the median of the projection of P onto the line through the origin parallel to vector u.*

We focus exclusively on the definition of the projection median in $\mathbb{R}^2$ and $\mathbb{R}^3$.

5.6.2 Properties of the Projection Median

In this section we establish properties of the projection median relating to invariance under similarity transformations and consistency of definition across dimensions.

Projection Median Invariance

We show that the projection median is invariant under similarity transformations for any multiset of clients P in $\mathbb{R}^3$. By Lem. 5.18, it follows that the projection median is also invariant under similarity transformations in $\mathbb{R}^2$.

Lemma 5.15. *The projection median is invariant under translation and uniform scaling transformations in* $\mathbb{R}^3$.

Proof. Let $f : \mathbb{R}^3 \to \mathbb{R}^3$ denote the composition of any translation transformation $g(p) = p + q$ and any uniform scaling transformation $h(p) = ap$. Thus, function f has the form $f(p) = ap + q$ for some $a \in \mathbb{R}^{3\times3}$ and $q \in \mathbb{R}^3$. Function f is an affine transformation and therefore it preserves relative ordering. Let P be any finite multiset of clients in $\mathbb{R}^3$.

$$
\begin{aligned}
\Pi_3(f(P)) &= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \operatorname{med}(f(P_{\theta,\phi}))\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \cdot f(\operatorname{med}(P_{\theta,\phi}))\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi[a \operatorname{med}(P_{\theta,\phi}) + q]\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} a\sin\phi \operatorname{med}(P_{\theta,\phi})\ d\phi\ d\theta + \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \cdot q\ d\phi\ d\theta \\
&= a\left[\frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \operatorname{med}(P_{\theta,\phi})\ d\phi\ d\theta\right] + q \\
&= f\left(\frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \operatorname{med}(P_{\theta,\phi})\ d\phi\ d\theta\right) \\
&= f(\Pi_3(P)).
\end{aligned}
$$

□

Lemma 5.16. *The projection median is invariant under rotation transformations in* $\mathbb{R}^3$.

Proof. Choose any $\alpha \in [0, 2\pi)$. Let $R_\alpha : \mathbb{R}^3 \to \mathbb{R}^3$ denote a rotation about the x-axis by α. That is,

$$
R_\alpha(p) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{bmatrix} p,
$$

for any fixed point $p \in \mathbb{R}^3$. It follows that

$$\begin{aligned}
\Pi_3[R_\alpha(P_{\theta,\phi})] &= \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \,\mathrm{med}[R_\alpha(P_{\theta,\phi})]\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \, R_\alpha[\mathrm{med}(P_{\theta,\phi})]\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{bmatrix} \mathrm{med}(P_{\theta,\phi})\ d\phi\ d\theta \\
&= \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \begin{bmatrix} \mathrm{med}(P_{\theta,\phi})_x \\ \cos\alpha\,\mathrm{med}(P_{\theta,\phi})_y + \sin\alpha\,\mathrm{med}(P_{\theta,\phi})_z \\ -\sin\alpha\,\mathrm{med}(P_{\theta,\phi})_y + \cos\alpha\,\mathrm{med}(P_{\theta,\phi})_z \end{bmatrix} d\phi\ d\theta \\
&= \frac{3}{2\pi}\begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{bmatrix} \int_0^\pi \int_0^\pi \sin\phi \,\mathrm{med}(P_{\theta,\phi})\ d\phi\ d\theta \\
&= R_\alpha[\Pi_3(P)]. \qquad (5.24)
\end{aligned}$$

Analogously, Eq. (5.24) holds when the rotation is about the y-axis, S_ϕ, or the z-axis, T_ψ:

$$S_\beta(p) = \begin{bmatrix} \cos\beta & 0 & -\sin\beta \\ 0 & 1 & 0 \\ -\sin\beta & 0 & \cos\beta \end{bmatrix} p \quad \text{and} \quad T_\gamma(p) = \begin{bmatrix} \cos\gamma & \sin\gamma & 0 \\ -\sin\gamma & \cos\gamma & 0 \\ 1 & 0 & 1 \end{bmatrix} p.$$

Since any rotation about the origin in $\mathbb{R}^3$ is defined by composition of R_α, S_β, and T_γ, Eq. (5.24) holds for any rotation $f : \mathbb{R}^3 \to \mathbb{R}^3$ about the origin in $\mathbb{R}^3$. Furthermore, any arbitrary rotation in $\mathbb{R}^3$ is defined by composition of rotation about the origin and translation. Thus, by Lem. 5.15 the projection median is invariant under rotation in $\mathbb{R}^3$. □

Lemma 5.17. *The projection median is invariant under reflection transformations in $\mathbb{R}^3$.*

Proof. For any point $p = (p_x, p_y, p_z)$ in $\mathbb{R}^3$, let $f : \mathbb{R}^3 \to \mathbb{R}^3$ denote the function $f(p) = (-p_x, p_y, p_z)$. Point $f(p)$ corresponds to the reflection of p across the yz-plane. For any multiset of clients P in $\mathbb{R}^3$,

$$\begin{aligned}
\Pi_3(f(P))_x &= \int_0^\pi \int_0^\pi \sin\phi \,\mathrm{med}(f(P_\theta))_x\ d\phi\ d\theta \\
&= -\int_0^\pi \int_0^\pi \sin\phi \,\mathrm{med}(P_\theta)_x\ d\phi\ d\theta \\
&= f(\Pi_3(P))_x. \qquad (5.25)
\end{aligned}$$

Since $\Pi_3(f(P))_y = f(\Pi_3(P))_y$ and $\Pi_3(f(P))_z = f(\Pi_3(P))_z$, we get that $\Pi_3(f(P)) = f(\Pi_3(P))$.

Any reflection $g : \mathbb{R}^3 \to \mathbb{R}^3$ can be described by some composition of f with rotation and translation transformations. It follows that the projection median is invariant under any reflection transformations in $\mathbb{R}^3$. □

Consistency of Definition Across Dimensions
As one should expect, when the clients of multiset P in $\mathbb{R}^3$ are coplanar the two- and three-dimensional definitions of the projection median coincide. Similarly, when the clients of a multiset P in $\mathbb{R}^2$ are collinear the one-, two-, and three-dimensional definitions of the projection median coincide.

Lemma 5.18. *Let P be a finite multiset of clients in $\mathbb{R}^3$ such that the positions of clients in P are coplanar. The three-dimensional projection median of P coincides with the two-dimensional projection median of P.*

Proof. Let $\Pi_2(P)$ and $\Pi_3(P)$ denote the respective two- and three-dimensional projection medians of P. Since Π_3 is invariant under rotation and translation, assume P is coplanar with the plane $z = 0$. For any $p \in P$, $p = (x, y, 0)$ and

$$\begin{aligned}\int_0^\pi \sin\phi \cdot u_{\theta,\phi}\langle (x,y,0), u_{\theta,\phi}\rangle \, d\phi &= \int_0^\pi \sin^2\phi \cdot u_{\theta,\phi}\langle (x,y), u_\theta\rangle \, d\phi \\ &= \langle (x,y), u_\theta\rangle \int_0^\pi \sin^2\phi \cdot u_{\theta,\phi} \, d\phi \\ &= \frac{4}{3} u_{\theta,\pi/2}\langle (x,y), u_\theta\rangle.\end{aligned}$$

Since the median $\mathrm{med}(P_{\theta,\phi})$ is defined in terms of $u_{\theta,\phi}\langle p, u_{\theta,\phi}\rangle$, it follows that,

$$\int_0^\pi \sin\phi \, \mathrm{med}(P_{\theta,\phi}) \, d\phi = \frac{4}{3}\mathrm{med}(P_{\theta,\pi/2}).$$

Furthermore, since $\{l_{\theta,\pi/2} \;:\; \theta \in [0,\pi]\}$ is the set of lines through the origin in plane $z = 0$, the projection $P_{\theta,\pi/2}$ is equivalent to P_θ. Therefore,

$$\begin{aligned}\Pi_3(P) &= \frac{3}{2\pi}\int_0^\pi \int_0^\pi \sin\phi \, \mathrm{med}(P_{\theta,\phi}) \, d\phi \, d\theta \\ &= \frac{3}{2\pi}\int_0^\pi \frac{4}{3}\mathrm{med}(P_{\theta,\pi/2}) \, d\theta \\ &= \frac{2}{\pi}\int_0^\pi \mathrm{med}(P_{\theta,\pi/2}) \, d\theta \\ &= \frac{2}{\pi}\int_0^\pi \mathrm{med}(P_\theta) \, d\theta \\ &= \Pi_2(P).\end{aligned}$$

□

Similarly, when the clients of P in $\mathbb{R}^2$ are collinear, the one-, two-, and three-dimensional definitions of the projection median coincide with the one-, two-, and three-dimensional definition of the Euclidean 1-median.

Lemma 5.19. *Let P be a finite multiset of clients in $\mathbb{R}^2$ such that the positions of clients in P are collinear.*

$$\Pi_i(P) = M_j(P), \;\forall i, j \in \{1, 2, 3\}.$$

Proof. Since Π_3 and M_d are invariant under translation and rotation, assume the clients of P lie on the x-axis such that $M_3(P)$ lies at the origin. Observe that $M_1(P) = M_2(P) = M_3(P)$. For any θ, $\text{med}(P_{\theta,\phi}) = (0,0,0)$. Therefore,

$$\begin{aligned}\Pi_3(P) &= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} \sin\phi \,\text{med}(P_{\theta,\phi})\; d\phi\; d\theta \\ &= \frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi} (0,0,0)\; d\phi\; d\theta \\ &= (0,0,0) \\ &= M_3(P). \end{aligned} \tag{5.26}$$

By Lem. 5.18, $\Pi_3(P) = \Pi_2(P)$. Recall that $\Pi_1(P) = M_1(P)$. □

5.6.3 Projection Median: Approximation Factor

In this section we prove that the approximation factor of the projection median is at most $4/\pi \approx 1.2732$ in $\mathbb{R}^2$. By giving an example, we provide a lower bound of $\sqrt{4/\pi^2+1} \approx 1.1854$ on the approximation factor in $\mathbb{R}^2$.

Approximation Factor of the Projection Median in Two Dimensions
In this section we bound the approximation factor of the projection median in two dimensions from above and below.

Theorem 5.20. *The two-dimensional projection median provides a $(4/\pi)$-approximation of the Euclidean 1-median.*

Proof. Let P denote any finite multiset of clients in $\mathbb{R}^2$. We bound the approximation factor of $\Pi_2(P)$:

$$\frac{\sum_{p\in P}||\Pi_2(P)-p||}{\sum_{q\in P}||M_2(P)-q||} = \frac{\sum_{p\in P}\left|\left|\frac{2}{\pi}\int_0^{\pi/2} S_\theta(P)\; d\theta - p\right|\right|}{\sum_{q\in P}||M_2(P)-q||},$$

by Lem. 5.14,

$$\begin{aligned} &= \frac{\sum_{p\in P}\left|\left|\frac{2}{\pi}\int_0^{\pi/2} S_\theta(P)\; d\theta - \frac{2}{\pi}\int_0^{\pi/2} p\; d\theta\right|\right|}{\sum_{q\in P}||M_2(P)-q||} \\ &= \frac{2}{\pi}\frac{\sum_{p\in P}\left|\left|\int_0^{\pi/2} S_\theta(P) - p\; d\theta\right|\right|}{\sum_{q\in P}||M_2(P)-q||} \\ &\le \frac{2}{\pi}\frac{\sum_{p\in P}\int_0^{\pi/2}||S_\theta(P)-p||\; d\theta}{\sum_{q\in P}||M_2(P)-q||}, \end{aligned}$$

by the $\triangle$ inequality,

$$\leq \frac{2}{\pi} \frac{\sum_{p \in P} \int_0^{\pi/2} d_\theta(S_\theta(P) - p)\ d\theta}{\sum_{q \in P} ||M_2(P) - q||}, \tag{5.27a}$$

since $\forall x\ ||x||_1 \geq ||x||$ and, similarly, $\forall x \forall \phi\ d_\phi(x) \geq ||x||$,

$$= \frac{2}{\pi} \frac{\int_0^{\pi/2} \sum_{p \in P} d_\theta(S_\theta(P) - p)\ d\theta}{\sum_{q \in P} ||M_2(P) - q||},$$

$$\leq \frac{2}{\pi} \frac{\int_0^{\pi/2} \sum_{p \in P} d_\theta(M_2(P) - p)\ d\theta}{\sum_{q \in P} ||M_2(P) - q||}, \tag{5.27b}$$

since $S_\phi(P)$ minimizes the sum of the d_ϕ distances to points of P,

$$= \frac{2}{\pi} \frac{\sum_{p \in P} \int_0^{\pi/2} d_\theta(M_2(P) - p)\ d\theta}{\sum_{q \in P} ||M_2(P) - q||},$$

$$= \frac{2}{\pi} \frac{\sum_{p \in P} \int_0^{\pi/2} \left[|\sin(\theta - \alpha_p)| + |\cos(\theta - \alpha_p)|\right] \cdot ||M_2(P) - p||\ d\theta}{\sum_{q \in P} ||M_2(P) - q||}, \tag{5.27c}$$

where $\alpha_p = \arctan[(M_2(P)_y - p_y)/(M_2(P)_x - p_x)] \bmod \frac{\pi}{2}$ (see Fig. 5.8),

$$\begin{aligned}
&= \frac{2}{\pi} \frac{\sum_{p \in P} \int_0^{\pi} |\sin(\theta - \alpha_p)| \cdot ||M_2(P) - p||\ d\theta}{\sum_{q \in P} ||M_2(P) - q||} \\
&= \frac{2}{\pi} \frac{\sum_{p \in P} \int_0^{\pi} |\sin\theta| \cdot ||M_2(P) - p||\ d\theta}{\sum_{q \in P} ||M_2(P) - q||} \\
&= \frac{2}{\pi} \frac{\sum_{p \in P} ||M_2(P) - p||}{\sum_{q \in P} ||M_2(P) - q||} \int_0^{\pi} |\sin\theta|\ d\theta \\
&= \frac{2}{\pi} \int_0^{\pi} |\sin\theta|\ d\theta \\
&= \frac{4}{\pi} \\
&\approx 1.2732.
\end{aligned} \tag{5.27d}$$

Therefore, for any finite multiset of points P in $\mathbb{R}^2$,

$$\sum_{p \in P} ||\Pi_2(P) - p|| \leq \frac{4}{\pi} \sum_{q \in P} ||M_2(P) - q||. \tag{5.28}$$

$\square$

Although we have not shown that the bound in Eq. (5.28) is tight, we provide the following lower bound:

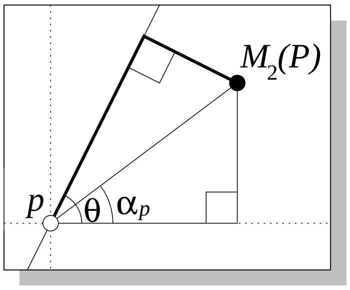

Figure 5.8: illustration in support of Thm. 5.20: $d_\theta(M_2(P) - p) = \left[|\sin(\theta - \alpha_p)| + |\cos(\theta - \alpha_p)|\right] \cdot ||M_2(P) - p||$

Theorem 5.21. *The two-dimensional projection median cannot guarantee an approximation factor less than* $\sqrt{4/\pi^2 + 1}$ *in the worst case.*

Proof. Let multiset P be defined by k clients located at $b = (0, 1)$, k clients located at $c = (0, -1)$, and a single client located at $d = (x, 0)$, for some $k \in \mathbb{N}$ and $x \in \mathbb{R}^+$. Let $\alpha = \pi/2 - \arctan(1/x) = \arctan x$. See Fig. 5.9.

We first derive the position of $M_2(P)$. Since the points of P are not collinear, the position of the Euclidean 1-median of P is unique. Consequently, by the symmetry of P and the invariance of $M_2(P)$ under reflection, $M_2(P)$ must lie on the x-axis. The Euclidean median sum of P is

$$2k\sqrt{1 + M_2(P)_x^2} + |x - M_2(P)_x|. \tag{5.29}$$

It is straightforward to show that Expr. (5.29) is minimized at $M_2(P)_x = 1/\sqrt{4k^2 - 1}$. Consequently, $M_2(P) = (1/\sqrt{4k^2 - 1}, 0)$.

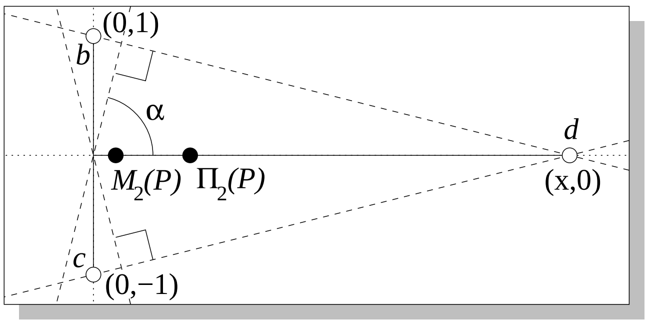

Figure 5.9: example realizing the lower bound in Thm. 5.21

By Eq. (5.17), the projection median of P is located at

$$\begin{aligned}
\Pi_2(P) &= \frac{2}{\pi}\left[\int_0^{\alpha} u_\theta \langle b, u_\theta\rangle \; d\theta + \int_{\alpha}^{\pi-\alpha} u_\theta \langle d, u_\theta\rangle \; d\theta + \int_{\pi-\alpha}^{\pi} u_\theta \langle c, u_\theta\rangle \; d\theta\right] \\
&= \frac{2}{\pi}\left[\int_0^{\alpha} u_\theta \sin\theta \; d\theta + \int_{\alpha}^{\pi-\alpha} x u_\theta \cos\theta \; d\theta - \int_{\pi-\alpha}^{\pi} u_\theta \sin\theta \; d\theta\right] \\
&= \left(\frac{2x}{\pi}\arctan\left(\frac{1}{x}\right), 0\right). \qquad (5.30)
\end{aligned}$$

The approximation factor λ is at least

$$\begin{aligned}
\lambda &\geq \lim_{\substack{x\to\infty \\ k\to\infty}} \frac{\sum_{p\in P} ||\Pi_2(P) - p||}{\sum_{q\in P} ||M_2(P) - q||} \\
&= \lim_{\substack{x\to\infty \\ k\to\infty}} \frac{2k\sqrt{\frac{4x^2}{\pi^2}\arctan^2\left(\frac{1}{x}\right) + 1} + x - \frac{2x}{\pi}\arctan\left(\frac{1}{x}\right)}{2k\sqrt{\frac{1}{4k^2-1} + 1} + x - \frac{1}{\sqrt{4k^2-1}}} \\
&= \lim_{x\to\infty} \sqrt{\frac{4x^2}{\pi^2}\arctan^2\left(\frac{1}{x}\right) + 1} \\
&= \sqrt{\frac{4}{\pi^2} + 1} \\
&> 1.1854. \qquad \square
\end{aligned}$$

Approximation Factor of the Projection Median in Three Dimensions It seems probable that the definition of the three-dimensional projection median can be interpreted in terms of the rectilinear 1-median as was done in two dimensions. If true, this equivalence may lead to a generalization of the two-dimensional upper bound on the approximation factor of the projection median. Should Thm. 5.20 generalize, the value corresponding to Eq. (5.27d) in $\mathbb{R}^3$ gives an upper bound of $3/2$. At the very least, this equivalence implies an upper bound of $\sqrt{3}$ on the approximation factor of the projection median, since $\Pi_3(P)$ can be defined as a convex combination of the corresponding median functions $S_{\theta,\phi}(P)$, each of which has approximation factor $\sqrt{3}$.

As for lower bounds, the consistency of the definition of the projection median from two to three dimensions shown in Thm. 5.6.2 implies that the two-dimensional lower bound of $\sqrt{4/\pi^2 + 1}$ established in Thm. 5.21 also holds in three dimensions.

5.6.4 Projection Median: Maximum Velocity

In this section we derive the maximum velocity of the projection median and show it is at most $4/\pi$ in $\mathbb{R}^2$ and $3/2$ in $\mathbb{R}^3$. In addition, we provide worst-case examples that realize each of these bounds.

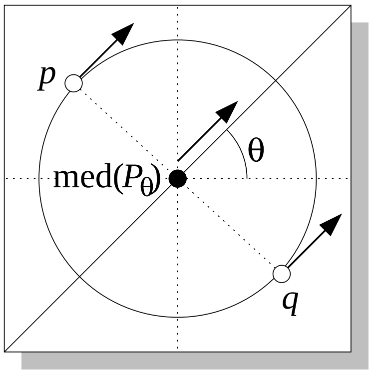

Figure 5.10: example realizing the bound in Thm. 5.22

Maximum Velocity of the Projection Median in Two Dimensions

We first bound the maximum velocity of the projection median in two dimensions and provide an example that realizes this velocity, showing the bound is tight.

Theorem 5.22. *The two-dimensional projection median, Π_2, has maximum velocity $4/\pi$.*

Proof. The proof is analogous to the proof of Thm. 4.24 except that Π_2 replaces Γ_2, $\mathrm{med}(P_\theta)$ replaces $\mathrm{mid}(P_\theta)$, and Eq. (4.52a) follows since the velocity of $\mathrm{med}(P_\theta)$ is at most the velocity of the fastest client in P_θ. □

The following example shows that the bound on maximum velocity is tight.

Theorem 5.23. *The two-dimensional projection median cannot guarantee relative velocity less than $4/\pi$.*

Proof. Let $P(0)$ be an infinite number of clients uniformly distributed on the unit circle centred at the origin. We assign instantaneous velocity to the clients of P at time $t = 0$ such that clients on or above the x-axis move right (clockwise) in a direction tangent to the circle while clients below the x-axis move right (counter-clockwise) in the opposite direction. See Fig. 5.10. Every client p in P has a corresponding client in P, $q = -p$, opposite the origin from p. By the symmetry of $P(0)$, for any line through the origin, an equal density of clients of $P(0)$ lie on either side of the line. Thus, the midpoint of each such pair of clients p and q defines $\mathrm{med}(P_\theta)$ for some P_θ (corresponding to the projection onto the line perpendicular to $p - q$). That is, for any θ, $\mathrm{med}(P(0)_\theta) = (0, 0)$. Furthermore, the resulting change in the position of $\mathrm{med}(P_\theta)$ is identical to the change at p and q. That is, $\frac{\partial}{\partial t}\,\mathrm{med}(P(0)_\theta) = u_\theta$. The velocity of $\Pi_2(P(t))$ at time $t = 0$ is given by

$$\begin{aligned}
\left|\left|\frac{\partial}{\partial t}\Pi_2(P(t))\right|\right| &= \left|\left|\frac{\partial}{\partial t}\frac{2}{\pi}\int_0^\pi \operatorname{med}(P_\theta)\ d\theta\right|\right| \\
&= \left|\left|\frac{2}{\pi}\int_0^\pi \frac{\partial}{\partial t}\operatorname{med}(P_\theta)\ d\theta\right|\right|, \qquad (5.31a)
\end{aligned}$$

by the Leibniz integral rule,

$$\begin{aligned}
&= \left|\left|\frac{2}{\pi}\int_0^\pi u_\theta\ d\theta\right|\right| \\
&= \sqrt{\left[\frac{2}{\pi}\int_0^\pi \cos\theta\ d\theta\right]^2 + \left[\frac{2}{\pi}\int_0^\pi \sin\theta\ d\theta\right]^2} \\
&= \left|\frac{2}{\pi}\int_0^\pi \sin\theta\ d\theta\right| \\
&= \frac{4}{\pi}. \qquad \square
\end{aligned}$$

The resulting velocity matches the upper bound derived in Thm. 5.22.

Maximum Velocity of the Projection Median in Three Dimensions
Using a technique similar to that used in two dimensions, we bound the maximum velocity of the projection median in three dimensions and provide an example that realizes this velocity, showing the bound is tight.

Theorem 5.24. *The three-dimensional mobile projection median, Π_3, has maximum velocity $3/2$.*

Proof. The proof is analogous to the proof of Thm. 4.26 except that Π_3 replaces Γ_3, $\operatorname{med}(P_{\theta,\phi})$ replaces $\operatorname{mid}(P_{\theta,\phi})$, and Eq. (4.59a) follows since the velocity of $\operatorname{med}(P_{\theta,\phi})$ is at most the velocity of the fastest client in $P_{\theta,\phi}$. $\square$

We generalize the two-dimensional worst-case example described in Thm. 5.23 to three dimensions to show that the bound on maximum velocity in three dimensions is tight.

Theorem 5.25. *The three-dimensional projection median cannot guarantee relative velocity less than $3/2$.*

Proof. Let $P(0)$ be an infinite number of clients uniformly distributed on the unit sphere centred at the origin. For simplicity, let P denote $P(0)$. We assign instantaneous velocity to the clients of P at time $t = 0$ such that clients move toward the positive x-axis in a direction tangent to the surface of the sphere. Every client p in P has a corresponding client in P, $q = -p$, opposite the origin from p. By the symmetry of P, for any plane through the origin, an equal density of clients of P lie on either side of the plane. Thus, for any θ and

ϕ, $\mathrm{med}(P_{\theta,\phi}) = (0,0,0)$. Furthermore, the resulting change in the position of $\mathrm{med}(P_{\theta,\phi})$ is identical to the change at the corresponding p and q. That is, $\frac{\partial}{\partial t}\mathrm{med}(P_{\theta,\phi}) = u_{\theta,\phi}$. The velocity of $\Pi_3(P(t))$ at time $t = 0$ is given by

$$\begin{aligned}\left\|\frac{\partial}{\partial t}\Pi_3(P(t))\right\| &= \left\|\frac{\partial}{\partial t}\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\phi\,\mathrm{med}(P_{\theta,\phi})\;d\phi\;d\theta\right\| \\ &= \left\|\frac{3}{2\pi}\int_0^\pi\int_0^\pi \frac{\partial}{\partial t}\sin\phi\,\mathrm{med}(P_{\theta,\phi})\;d\phi\;d\theta\right\|, \qquad (5.32a)\end{aligned}$$

by the Leibniz integral rule,

$$\begin{aligned}&= \left\|\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\phi\cdot u_{\theta,\phi}\;d\phi\;d\theta\right\| \\ &= \left(\left[\frac{3}{2\pi}\int_0^\pi\int_0^\pi \cos\theta\sin^2\phi\;d\phi\;d\theta\right]^2\right. \\ &\quad+\left[\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\theta\sin^2\phi\;d\phi\;d\theta\right]^2 \\ &\quad\left.+\left[\frac{3}{2\pi}\int_0^\pi\int_0^\pi \cos\phi\sin\phi\;d\phi\;d\theta\right]^2\right)^{\frac{1}{2}} \\ &= \left|\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\theta\sin^2\phi\;d\phi\;d\theta\right| \\ &= \frac{3}{2}. \qquad \square\end{aligned}$$

The resulting velocity matches the upper bound derived in Thm. 5.24.

5.6.5 Generalized Definition of the Projection Median

The structure common to both the projection median and the Steiner centre's definition by projection involves projecting the clients of P onto a line through

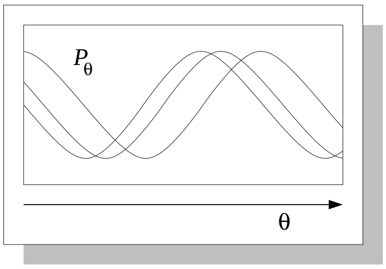

Figure 5.11: correspondence between the k-level of P_θ as a function of θ and the generalized definition of the projection median

the origin, finding the kth largest element along this projected set of points, and integrating this element over all such lines through the origin. In two dimensions, the following function captures the precise definition of the generalization of Γ_2 and Π_2:

$$\Upsilon_d^k(P) = \frac{1}{\pi} \int_0^{2\pi} m(P, \theta, k)\ d\theta, \tag{5.33}$$

where $m(P, \theta, k)$ returns $u_\theta \langle p_k, u_\theta \rangle$, such that $\langle p_1, u_\theta \rangle \leq \ldots \leq \langle p_k, u_\theta \rangle \leq \ldots \leq \langle p_n, u_\theta \rangle$, for $p_1, \ldots, p_n \in P$, and $|P| = n$. That is, $m(P, \theta, k)$ returns the kth largest element of P_θ relative to the orientation of θ. Note that θ ranges from 0 to 2π; thus the kth largest element relative to θ corresponds to the $(|P|-k+1)$st largest element relative to $\theta + \pi$. This definition generalizes to three or more dimensions in the same manner as we generalized the projection median and Steiner centre by projection to higher dimensions.

When $k = 1$ or $k = |P|$, Eq. (5.33) simplifies to Eq. (4.14): the definition of the Steiner centre by projection. Similarly, when $k = \lfloor |P|/2 \rfloor$ or $k = \lceil |P|/2 \rceil$, Eq. (5.33) simplifies to Eq. (5.17): the definition of the projection median. Thus,

$$\Upsilon_d^1(P) = \Gamma_d(P) = \Upsilon_d^{|P|}(P) \quad \text{and} \quad \Upsilon_d^{\lfloor |P|/2 \rfloor}(P) = \Pi_d(P) = \Upsilon_d^{\lceil |P|/2 \rceil}(P).$$

The significance of the value of Eq. (5.33) remains to be understood for value of $2 \leq k \leq \lfloor |P|/2 \rfloor - 1$ and $\lceil |P|/2 \rceil + 1 \leq k \leq |P| - 1$. The problem can be understood in terms of a k-level of the set of $|P|$ functions $f_i : [0, 2\pi) \to \mathbb{R}$, given by $f_i = \langle u_\theta, p_i \rangle$, for $1 \leq i \leq |P|$.

5.7 Convex Combinations

A new median function can be defined by a convex combination of existing median functions. Closely related to our discussion of convex combinations of centre functions in Sec. 4.8, this section examines the approximation factor and maximum velocity of the resulting convex combination in terms of the approximation factors and maximum velocities of the component median functions. The results established in Sec. 5.7 are used to bound the approximation factors and maximum velocities of the Gaussian median in Sec. 5.8 and of convex combinations of the centre of mass and the rectilinear 1-median in Sec. 5.9.

5.7.1 Convex Combinations: Approximation Factor

We show that the approximation factor of a convex combination can be bounded by the corresponding convex combination of the approximation factors of its component median functions.

Theorem 5.26. *Let X_d and Y_d denote median functions in $\mathbb{R}^d$ that are invariant under translation. Let $k \in [0, 1]$. Let $Z_d(P) = kX_d(P) + (1-k)Y_d(P)$ define a third median function. Let λ_X and λ_Y denote the respective approximation* *factors of X_d and Y_d. Median function Z_d provides a λ_Z-approximation of the Euclidean 1-median, where*

$$\lambda_Z = k\lambda_X + (1-k)\lambda_Y. \tag{5.34}$$

Proof. Choose any set of client positions P in $\mathbb{R}^d$.

$$\begin{aligned}\frac{\sum_{p\in P}||p - Z_d(P)||}{\sum_{p\in P}||p - M_d(P)||} &= \frac{\sum_{p\in P}||p - [kX_d(P) + (1-k)Y_d(P)]||}{\sum_{p\in P}||p - M_d(P)||} \\ &\le \frac{\sum_{p\in P} k||p - X_d(P)|| + (1-k)||p - Y_d(P)||}{\sum_{p\in P}||p - M_d(P)||},\end{aligned}$$

by Cor. 4.29,

$$\begin{aligned} &= k\frac{\sum_{p\in P}||p - X_d(P)||}{\sum_{p\in P}||p - M_d(P)||} + (1-k)\frac{\sum_{p\in P}||p - Y_d(P)||}{\sum_{p\in P}||p - M_d(P)||} \\ &\le k\lambda_X + (1-k)\lambda_Y,\end{aligned}$$

by Def. 3.5. Therefore,

$$\forall P,\ \sum_{p\in P}||p - Z_d(P)|| \le [k\lambda_X + (1-k)\lambda_Y]\sum_{p\in P}||p - M_d(P)||. \qquad \square$$

Corollary 5.27. *Given $n \in \mathbb{N}$, for every $1 \le i \le n$, let X_d^i denote a median function in $\mathbb{R}^d$ that is invariant under translation and let $k_i \in [0,1]$ such that $\sum_{i=1}^n k_i = 1$. Let Z_d denote the median function defined by $Z_d(P) = \sum_{i=1}^n k_i X_d^i(P)$. For each i, let λ_i denote the approximation factor of X_d^i. Median function Z_d is a λ_Z-approximation of the Euclidean 1-median, where*

$$\lambda_Z = \sum_{i=1}^n k_i\lambda_i. \tag{5.35}$$

Proof. The result follows by induction on n using Thm. 5.26. $\square$

5.7.2 Convex Combinations: Maximum Velocity

The analogous results for maximum velocity were shown in Thm. 4.33 and Cor. 4.34 in our discussion of convex combinations of centre functions in Sec. 4.8.

5.7.3 Using Convex Combinations to Compare Median Functions

Thms. 5.26 and 4.33 allow us to evaluate the significance of a median function's approximation factor and maximum velocity. That is, if we have three median functions Υ_d^1, Υ_d^2, and Υ_d^3 such that their respective maximum velocities are sorted in increasing order, we can define a fourth median function Υ_d^4 by a convex combination of Υ_d^1 and Υ_d^3 such that the maximum velocity of Υ_d^4 matches that of Υ_d^2. Comparing the approximation factors of Υ_d^2 and Υ_d^4 helps determine whether Υ_d^2 is beneficial as a median function. In Sec. 5.9, we use this technique to compare the projection median against a convex combination of the centre of mass and of the rectilinear 1-median.

5.8 Gaussian Median

The definition of the projection median exhibits similarities in structure to the definition of the Steiner centre by projection. As discussed in Sec. 4.6, the Steiner centre can also be defined as a normalized weighted mean of the client positions. Can such a generalization be applied to the projection median? Indeed, a simple transformation of the Gaussian weight of each client (see Defs. 4.2 and 4.5) provides a useful definition for a median function. We call this new median function the *Gaussian median* and briefly examine its properties in $\mathbb{R}^2$ and $\mathbb{R}^3$ in this section.

In particular, we show that the location of the Gaussian median of a multiset of clients P coincides with the projection median of P when $|P| \leq 4$ and that the position of the Gaussian median can be defined as a linear combination of the centre of mass and the Steiner centre. The Gaussian median is included in this chapter not with the intent to define a competitive median function, but rather to provide insight into properties related to the projection median, the Steiner centre, and the centre of mass. Indeed there remains room for improvement in the bounds on the maximum velocity and approximation factor of the Gaussian median that follow.

5.8.1 Definition

Gaussian Median Definition in Two Dimensions

The Steiner centre of P is defined as a normalized weighted mean of clients in P, where the weight, called Gaussian weight, of each client whose position is an extreme point of P is proportional to the turn angle at that point on the convex hull of P. Clients in the interior of P have weight 0.

In a sense, the notions of median and centre are opposites in client sets of small cardinality; a centre is determined by the extreme points whereas a median is determined by interior clients. See Fig. 5.12. To exploit this property, we take the inverse of the Gaussian weight of each client and examine the resulting median function:

Definition 5.4. *Let P in $\mathbb{R}^2$ be a finite multiset of clients with $|P| \geq 3$. The* **two-dimensional Gaussian median** *of P is the normalized weighted mean of P:*

$$G_2(P) = \frac{1}{(|P|-2)\pi} \sum_{p \in P} g_2(p)p, \tag{5.36}$$

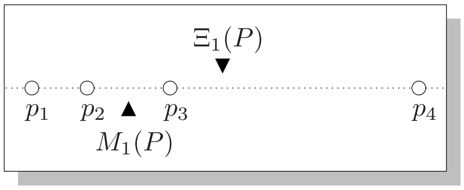

Figure 5.12: The Euclidean 1-median is defined in terms of interior points, $M_1(P) = (p_2 + p_3)/2$, whereas the Euclidean 1-centre is defined in terms of extreme points, $\Xi_1(P) = (p_1 + p_4)/2$.

where $g_2(p)$ is the **two-dimensional Gaussian median weight** *assigned to client $p \in P$ by*

$$g_2(p) = \begin{cases} \alpha_p & \text{if } p \in V_P \\ \pi & \text{if } p \in P - V_P. \end{cases},$$

where V_P is the set of extreme points of P and α_p is the interior angle formed at p on the convex hull of P.

Observe that for all p, $g_2(p) = \pi - w_2(p)$, where $g_2(p)$ is the Gaussian median weight and w_2 is the Gaussian weight (see Def. 4.2). The sum of the turn angles of $CH(P)$ is independent of $|P|$ whereas the sum of the interior angles of $CH(P)$ is $(|CH(P)| - 2)\pi$, hence the normalizing factor in Eq. (5.36).

See Sec. 8.3.1 for a description of algorithms for maintaining the mobile Gaussian median in two dimensions.

Gaussian Median Definition in Three Dimensions

Similarly, we define the three-dimensional Gaussian median:

Definition 5.5. *Let P in $\mathbb{R}^3$ be a finite multiset of clients with $|P| \geq 3$. The* **three-dimensional Gaussian median** *of P is the normalized weighted mean of P:*

$$G_3(P) = \frac{1}{(|P| - 2)2\pi} \sum_{p \in P} g_3(p)p, \tag{5.37}$$

where $g_3(p)$ is the **three-dimensional Gaussian median weight** *assigned to client $p \in P$ by*

$$g_3(p) = \begin{cases} \sum_{f_j \in F_p} \alpha_{p,j} & \text{if } p \in V_P \\ 2\pi & \text{if } p \in P - V_P. \end{cases},$$

where V_P is the set of extreme points of P, F_p denotes the set of faces that meet at p for every $p \in V_P$, and $\alpha_{p,j}$ denotes the interior plane angle on f_j at p for every face $f_j \in F_p$.

A natural definition for G_d when $|P| \leq 2$ assigns $G_d(P) = (p_1 + p_2)/2$ when $P = \{p_1, p_2\}$ and $G_d(P) = p$ when $P = \{p\}$. That is, $G_d(P) = \Xi_d(P) = M_d(P) = \Pi_d(P) = \Gamma_d(P)$ for $|P| \leq 2$.

It is straightforward to show that the Gaussian median is invariant under similarity transformations and that the three-dimensional definition of the Gaussian median of a set of coplanar points coincides with the two-dimensional definition.

5.8.2 Properties of the Gaussian Median

In this section we establish properties of the Gaussian median that relate its definition to the definitions of the Steiner centre, the centre of mass, and the projection median.

Gaussian Median Definition by a Linear Combination

The similarity in the definition of Gaussian weights and Gaussian median weights allows the Gaussian median to be expressed in terms of the Steiner centre and the centre of mass.

Observation 5.28. *For all finite multisets P in $\mathbb{R}^2$ where $|P| \geq 3$,*

$$G_2(P) = \frac{1}{|P| - 2}[|P|C_2(P) - 2\Gamma_2(P)]. \tag{5.38}$$

Proof. Choose any P in $\mathbb{R}^2$ such that $|P| \geq 3$.

$$\begin{aligned}
G_2(P) &= \frac{1}{(|P| - 2)\pi} \sum_{p \in P} g_2(p)p \\
&= \frac{1}{(|P| - 2)\pi} \sum_{p \in P} [\pi - w_2(p)]p, && \text{by Defs. 4.2 and 5.4,} \\
&= \frac{1}{|P| - 2} \left[\sum_{p \in P} p - \frac{1}{\pi} \sum_{p \in P} w_2(p)p \right] \\
&= \frac{1}{|P| - 2} \big[|P|C_2(P) - 2\Gamma_2(P)\big], && \text{by Defs. 4.3 and 5.4,}
\end{aligned}$$

where $g_2(p)$ denotes the Gaussian median weight of client p and $w_2(p)$ denotes its Gaussian weight. $\square$

The analogous property holds in three dimensions:

Observation 5.29. *For all finite multisets P in $\mathbb{R}^3$ where $|P| \geq 3$,*

$$G_3(P) = \frac{1}{|P| - 2}[|P|C_3(P) - 2\Gamma_3(P)]. \tag{5.39}$$

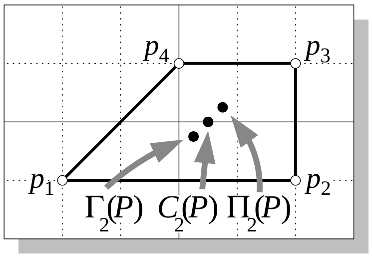

Figure 5.13: As a consequence of Cor. 5.30 and Thm. 5.31, $C_2(P) = \frac{1}{2}[\Gamma_2(P) + \Pi_2(P)]$ when $|P| = 4$.

Proof. The proof is identical to the proof of Obs. 5.28 except that the value π is replaced by 2π and references to Defs. 4.2, 4.3, and 5.4 refer instead to Defs. 4.5, 4.6, and 5.5. □

These results imply the following relationship between the projection median and the Steiner centre:

Corollary 5.30. *For any multiset P in $\mathbb{R}^d$ where $|P| = 4$ and $d \leq 3$,*

$$C_d(P) = \frac{1}{2}\big[\Pi_d(P) + \Gamma_d(P)\big], \tag{5.40}$$

where C_d denotes the centre of mass, Π_d denotes the projection median, and Γ_d denotes the Steiner centre.

Proof. The result follows from Obs. 5.28 and 5.29 and Thms. 5.31 and 5.32. □

We revisit the example presented in Sec. 4.6.1. Let $P = \{p_1, \ldots, p_4\} = \{(-2,-1),(2,-1),(2,1),(0,1)\}$, respectively. See Fig. 5.13. The Steiner centre of P, $\Gamma_2(P)$, lies in position $(1/4, -1/4)$, the Gaussian median of P, $G_2(P)$, lies in position $(3/4, 1/4)$, and the centre of mass of P, $C_2(P)$, lies in position $(1/2, 0)$. As shown in the next section, $\Pi_2(P) = G_2(P)$ since $|P| \leq 4$.

Gaussian Median Equivalence with Projection Median

We now prove that the projection median and the Gaussian median coincide in $\mathbb{R}^2$ when $|P| \leq 4$.

Theorem 5.31. *For any multiset P in $\mathbb{R}^2$ where $|P| \leq 4$,*

$$G_2(P) = \Pi_2(P).$$

Proof. When $|P| \leq 2$, $G_2(P) = \Pi_2(P)$ by the definitions of $G_2(P)$ and $\Pi_2(P)$. Therefore, choose any multiset P in $\mathbb{R}^2$ such that $|P| \in \{3, 4\}$. Observe that for

$|P| \in \{3,4\}$, $\mathrm{med}(P_\theta) = [\sum_{p\in P_\theta} p - 2\,\mathrm{mid}(P_\theta)]/(|P|-2)$.

$$\begin{aligned}
\Pi_2(P) &= \frac{2}{\pi}\int_0^\pi \mathrm{med}(P_\theta)\; d\theta \\
&= \frac{2}{\pi}\int_0^\pi \frac{1}{|P|-2}\left[\sum_{p\in P_\theta} p - 2\,\mathrm{mid}(P_\theta)\right]\; d\theta \\
&= \frac{1}{|P|-2}\left(\frac{2}{\pi}\int_0^\pi \sum_{p\in P_\theta} p\; d\theta - \frac{4}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\; d\theta\right) \\
&= \frac{1}{|P|-2}\left(\frac{2}{\pi}\int_0^\pi \sum_{p\in P} u_\theta\langle p, u_\theta\rangle\; d\theta - \frac{4}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\; d\theta\right) \\
&= \frac{1}{|P|-2}\left(\sum_{p\in P}\left(\frac{2}{\pi}\int_0^\pi u_\theta\langle p, u_\theta\rangle\; d\theta\right) - \frac{4}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\; d\theta\right) \\
&= \frac{1}{|P|-2}\left(\sum_{p\in P} p - \frac{4}{\pi}\int_0^\pi \mathrm{mid}(P_\theta)\; d\theta\right) \\
&= \frac{1}{|P|-2}\big[|P|C_2(P) - 2\Gamma_2(P)\big], && \text{by Defs. 2.11 and 4.3,} \\
&= G_2(P), && \text{by Obs. 5.28.}
\end{aligned}$$

□

The analogous result holds in three dimensions:

Theorem 5.32. *For any multiset P in $\mathbb{R}^3$ where $|P| \le 4$,*

$$G_3(P) = \Pi_3(P).$$

Proof. When $|P| \le 2$, $G_3(P) = \Pi_3(P)$ by the definition of $G_3(P)$. Therefore, choose any multiset P in $\mathbb{R}^3$ such that $|P| \in \{3,4\}$.

$$\begin{aligned}
\Pi_3(P) =& \frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\phi\,\mathrm{med}(P_{\theta,\phi})\; d\phi\; d\theta \\
=& \frac{3}{2\pi}\int_0^\pi\int_0^\pi \frac{\sin\phi}{|P|-2}\left(\sum_{p\in P_{\theta,\phi}} p - 2\,\mathrm{mid}(P_{\theta,\phi})\right)\; d\phi\; d\theta \\
=& \frac{1}{|P|-2}\left(\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\phi \sum_{p\in P_{\theta,\phi}} p\; d\phi\; d\theta\right. \\
&\left. - \frac{3}{\pi}\int_0^\pi\int_0^\pi \sin\phi\,\mathrm{mid}(P_{\theta,\phi})\; d\phi\; d\theta\right) \\
=& \frac{1}{|P|-2}\left(\frac{3}{2\pi}\int_0^\pi\int_0^\pi \sin\phi \sum_{p\in P} u_{\theta,\phi}\langle p, u_{\theta,\phi}\rangle\; d\phi\; d\theta - 2\Gamma_3(P)\right),
\end{aligned}$$

by Def. 4.6,

$$
\begin{aligned}
&= \frac{1}{|P|-2}\left(\sum_{p\in P}\left(\frac{3}{2\pi}\int_0^{\pi}\int_0^{\pi}\sin\phi\cdot u_{\theta,\phi}\langle p, u_{\theta,\phi}\rangle\ d\phi\ d\theta\right) - 2\Gamma_3(P)\right)\\
&= \frac{1}{|P|-2}\left(\sum_{p\in P} p - 2\Gamma_3(P)\right)\\
&= \frac{1}{|P|-2}\left[|P|C_3(P) - 2\Gamma_3(P)\right], && \text{by Def. 2.11,}\\
&= G_3(P), && \text{by Obs. 5.29.} \qquad \square
\end{aligned}
$$

We give a counter-example to show that Thms. 5.31 and 5.32 do not generalize to $|P| \geq 5$.

Observation 5.33. *In general, $G_d(P) \neq \Pi_d(P)$ when $|P| \geq 5$.*

Proof. Let $P = \{a, b, c, d, e\}$ such that $a = (0,0)$, $b = (1,0)$, $c = d = (4,0)$, and $e = (5,0)$. The median $M_2(P)$ is located at $(4,0)$. Furthermore, the median of any projection P_θ will be the projection of $c = d$. Therefore, $\Pi_2(P) = (4,0)$. As for the Gaussian median, $g_2(a) = g_2(e) = 0$ and $g_2(b) = g_2(c) = g_2(d) = \pi$. The Gaussian median lies at

$$G_2(P) = \frac{1}{3\pi}\sum_{p\in P} g_2(p)p = \frac{b+c+d}{3} = (3,0) \neq \Pi_2(P). \qquad \square$$

Note on Equivalence with Projection Median
An alternative to Gaussian median weight as defined in Defs. 5.4 and 5.5 would be to weight each client p in P by the fraction of turn angles $\theta \in [0,\pi)$ (respectively, $(\theta,\phi) \in [0,\pi)^2$) for which p induces a median of P_θ (respectively, $P_{\theta,\phi}$). Such a definition, although significantly more difficult to analyze, may lead to a generalization of the Steiner centre's definition by Gaussian weights that coincides with the projection median for $|P| \geq 5$.

5.8.3 Gaussian Median: Approximation Factor

We briefly examine the approximation factor of the Gaussian median.

Theorem 5.34. *For $d \in \{2,3\}$, the d-dimensional Gaussian median, G_d, cannot guarantee a λ-approximation of the Euclidean 1-median for any $\lambda < 3/2$.*

Proof. Although the example is described in $\mathbb{R}^2$, it implies the same result in $\mathbb{R}^3$. Let $n \geq 3$ be an integer and let $\epsilon > 0$. Let $n-2$ clients be located at $(0,0)$. and let one client be located at $(1,0)$. Let client a be located at $(-\epsilon,0)$ and let client b be located at $(1+\epsilon,0)$. The Gaussian median of P assigns equal weight to all clients except a and b which have weight 0. The Gaussian median of P lies at $(1/(n-2),0)$. The Euclidean 1-median of P lies at $(0,0)$. The Euclidean median sum is $2 + 2\epsilon$ and the sum of the distance from the Gaussian median of P to the clients of P is $1 + 2\epsilon + 2(n-3)/(n-2)$.

We get the following lower bound on the approximation factor of the Gaussian median, λ:

$$\begin{aligned}
& \sup_{\substack{\epsilon>0 \\ n\geq 3}} 1 + 2\epsilon + \frac{2(n-3)}{n-2} \leq \lambda(2+2\epsilon) \\
\Rightarrow \quad & \lim_{\substack{\epsilon\to 0 \\ n\to\infty}} 1 + 2\epsilon + \frac{2(n-3)}{n-2} \leq \lambda(2+2\epsilon) \\
\Rightarrow \quad & \lambda \geq \frac{3}{2}.
\end{aligned}$$ ☐

5.8.4 Gaussian Median: Maximum Velocity

We briefly examine the maximum velocity of the Gaussian median.

Theorem 5.35. *The two-dimensional mobile Gaussian median, G_2, has maximum velocity $3 + 8/\pi$.*

Proof. When $|P| = 1$, the velocity of $G_2(P(t))$ matches the velocity of the single client in P. When $|P| = 2$, the velocity of $G_2(P(t))$ is at most the velocity of the midpoint of the two clients of P. Thus, when $|P| \leq 2$, the maximum velocity of G_2 is one. Assume $|P| \geq 3$.

It is straightforward to generalize the proof of Lem. 4.28 to any linear combination of median functions. That is, we do not require k to be in the interval $[0, 1]$:

$$||ka + (1-k)b|| \leq |k| \cdot ||a|| + |1-k| \cdot ||b||. \tag{5.41}$$

Recall that

$$G_d(P(t)) = \frac{1}{|P|-2}[|P|C_d(P(t)) - 2\Gamma_d(P(t))],$$

by Obs. 5.28. By Cor. 4.15, C_2 has maximum velocity one and by Thm. 4.24, Γ_2 has maximum velocity $4/\pi$. The generalization in Eq. (5.41) allows for the corresponding generalization of Thm. 4.33. It follows that

$$\begin{aligned}
\forall t_1, t_2 \in T,\ ||G_2(P(t_1)) - G_2(P(t_2))|| &\leq \left(\frac{|P|}{|P|-2} + \frac{2}{|P|-2}\frac{4}{\pi}\right)|t_1 - t_2| \\
&\leq \left(3 + \frac{8}{\pi}\right)|t_1 - t_2|, \qquad (5.42)
\end{aligned}$$

since $|P| \geq 3$. ☐

Theorem 5.36. *The three-dimensional mobile Gaussian median, G_3, has maximum velocity 6.*

Proof. The proof is analogous to the proof of Thm. 5.35, except that the reference to Obs. 5.28. refers instead to Obs. 5.29 and the reference to Thm. 4.24 refers instead to Thm. 4.26, in which the maximum velocity of Γ_3 is shown to be 3/2. It follows that

$$\forall t_1, t_2 \in T,\ ||G_3(P(t_1)) - G_3(P(t_2))|| \leq \left(\frac{|P|}{|P|-2} + \frac{3}{|P|-2} \right) |t_1 - t_2| \\ \leq 6|t_1 - t_2|, \tag{5.43}$$

since $|P| \geq 3$. □

The bounds shown in Thms. 5.35 and 5.36 are unlikely to be tight; they are included to demonstrate that these values have fixed upper bounds.

5.9 Evaluation

In Secs. 5.4 through 5.8 we analyzed candidate functions whose properties are most applicable for defining good bounded-velocity approximations of the mobile Euclidean 1-median. In this section we compare these various median functions against each other, in terms of approximation factor, maximum velocity, invariance under similarity transformations, and consistency of definition across dimensions.

Those median functions which we identified are the rectilinear 1-median, the centre of mass, the projection median, and the Gaussian median. To these we add the median function defined by a client p in P (see Sec. 5.3.1), and a convex combination of the rectilinear 1-median and the centre of mass which we discuss below.

Rectilinear 1-Median

In Sec. 5.4 we examined the rectilinear 1-median, S_d. In $\mathbb{R}^d$, we showed an upper bound of $\sqrt{d}$ and a lower bound of $(1 + \sqrt{d-1})/\sqrt{d}$ on the approximation factor of S_d. When $d = 2$, the upper and lower bounds coincide at $\sqrt{2}$. For $d \geq 3$, the bounds diverge. Still in $\mathbb{R}^d$, we referred to a result of Bereg et al. [BBKS06] showing a tight bound of $\sqrt{d}$ on its maximum velocity. As mentioned in Sec. 2.5.1, S_d is not invariant under under rotation or reflection. It is, however, invariant under translation and scaling. The definition of S_d is consistent across dimensions.

Centre of Mass

In Sec. 5.5 we examined the centre of mass, C_d. In $\mathbb{R}^d$, we showed a tight bound of 2 on the approximation factor of C_d and we referred to a result of Bereg et al. [BBKS06] showing a tight bound of 1 on its maximum velocity. As mentioned in Sec. 2.5.2. C_d is invariant under similarity transformations and its definition is consistent across dimensions.

Projection Median

In Sec. 5.6 we introduced the projection median, Π_d. In $\mathbb{R}^2$, we showed a lower bound of $\sqrt{4/\pi^2+1}$ and an upper bound of $4/\pi$ on the approximation factor of Π_2 as well as a tight bound of $4/\pi$ on its maximum velocity. In $\mathbb{R}^3$, we observed that the same lower bound applies on the approximation factor of Π_3 and we showed a tight bound of $3/2$ on its maximum velocity. We showed that the definition of Π_d is consistent across dimensions and we demonstrated the invariance of Π_d under similarity transformations.

Convex Combinations

In Sec. 5.7 we presented a discussion of convex combinations of median functions, including results on bounding the approximation factor and maximum velocity of a convex combination in terms of the approximation factors and maximum velocities of its component median functions. We now examine specific convex combinations involving the median functions described above.

Any convex combination Υ_d that includes the Euclidean 1-median as a component of non-zero weight, regardless of the combination of median functions that completes the definition of Υ_d, results in discontinuous motion for Υ_d. Consequently, we consider only convex combinations whose composition does not include M_d.

The maximum velocity of S_d is greater than that of the projection median while the maximum velocity of C_d less than that of the projection median. Thus, we consider the convex combination of S_d and C_d given by $kS_d(P)+(1-k)C_d(P)$ for some $k \in [0,1]$. We select values of k such that the maximum velocity of $kS_d(P)+(1-k)C_d(P)$ is equal to $4/\pi$ in $\mathbb{R}^2$ and equal to $3/2$ in $\mathbb{R}^3$, allowing us to compare the convex combination directly against the projection median for a fixed maximum velocity. The specific values of k is given by solving for k_2 and k_3 in

$$\begin{aligned} k_2\sqrt{2}+(1-k_2) &= \frac{4}{\pi}, && \text{in } \mathbb{R}^2, \\ \text{and} \quad k_3\sqrt{3}+(1-k_3) &= \frac{3}{2}, && \text{in } \mathbb{R}^3. \end{aligned}$$

Solving for these values gives $k_2 = (4-\pi)/[\pi(\sqrt{2}-1)] \approx 0.6597$ and $k_3 = 1/[2(\sqrt{3}-1)] \approx 0.6831$. The corresponding bounds on the approximation factors are

$$\begin{aligned} k_2\sqrt{2}+2(1-k_2) &= \frac{\sqrt{2}(4+\pi)-8}{\pi(\sqrt{2}-1)} && \approx 1.6136, \text{ in } \mathbb{R}^2, \\ \text{and} \quad k_3\sqrt{3}+2(1-k_3) &= \frac{5\sqrt{3}-6}{2(\sqrt{3}-1)} && \approx 1.8170, \text{ in } \mathbb{R}^3. \end{aligned}$$

Finally, since S_d is neither invariant under rotation nor reflection, it follows that these properties do not hold for any convex combination whose composition includes S_d.

median function		approximation factor	maximum velocity
Euclidean 1-median	M_2	$\lambda = 1$	$v_{\max} = \infty$
single client $p \in P$	p	$\lambda = \infty$	$v_{\max} = 1$
centre of mass	C_2	$\lambda = 2$	$v_{\max} = 1$
rectilinear 1-median	S_2	$\lambda = \sqrt{2} \approx 1.4142$	$v_{\max} = \sqrt{2} \approx 1.4142$
projection median	Π_2	$\sqrt{4/\pi^2 + 1} \leq \lambda \leq 4/\pi$ $\Rightarrow 1.1854 \leq \lambda \leq 1.2732$	$v_{\max} = 4/\pi \approx 1.2732$
convex combination of S_2 and C_2		$\lambda \leq 1.6136$	$v_{\max} \leq 4/\pi \approx 1.2732$
Gaussian median	G_2	$1.5 \leq \lambda$	$v_{\max} \leq 3 + 8/\pi$ $\Rightarrow v_{\max} \leq 5.5465$

Table 5.1: comparing median functions in $\mathbb{R}^2$

median function		approximation factor	maximum velocity
Euclidean 1-median	M_3	$\lambda = 1$	$v_{\max} = \infty$
single client $p \in P$	p	$\lambda = \infty$	$v_{\max} = 1$
centre of mass	C_3	$\lambda = 2$	$v_{\max} = 1$
rectilinear 1-median	S_3	$(1 + \sqrt{2})/\sqrt{3} \leq \lambda \leq \sqrt{3}$ $\Rightarrow 1.3938 \leq \lambda \leq 1.7321$	$v_{\max} = \sqrt{3} \approx 1.7321$
projection median	Π_3	$1.1854 \approx \sqrt{4/\pi^2 + 1} \leq \lambda$	$v_{\max} = 1.5$
convex combination of S_3 and C_3		$\lambda \leq 1.8170$	$v_{\max} \leq 1.5$
Gaussian median	G_3	$1.5 \leq \lambda$	$v_{\max} \leq 6$

Table 5.2: comparing median functions in $\mathbb{R}^3$

Gaussian Median

In Sec. 5.8 we introduced the Gaussian median, G_d, as a linear combination of the Steiner centre, Γ_d, and the centre of mass, C_d. We showed that for $|P| \leq 4$, $G_d(P)$ and $\Pi_d(P)$ coincide for $d \in \{2, 3\}$. In $\mathbb{R}^2$ and $\mathbb{R}^3$, we showed a lower bound of $3/2$ on the approximation factor of G_2. In $\mathbb{R}^2$ we showed an upper bound of $3+8/\pi$ on the maximum velocity of G_2 and in $\mathbb{R}^3$ we showed an upper bound of 6 on the maximum velocity of G_3.

Comparison of Median Functions

The values for the approximation factor and maximum velocity of these various median functions are displayed in Tab. 5.1 for $\mathbb{R}^2$ and in Tab. 5.2 for $\mathbb{R}^3$.

First, observe that the convex combination $k_2S_2+(1-k_2)C_2$ provides a worse approximation factor than does Π_2 for the same maximum velocity. That is, the projection median provides a better approximation of the Euclidean 1-median than does the corresponding convex combination of C_d and S_d, even though both have the same maximum velocity. As for the rectilinear 1-median, the projection median has both a lower approximation factor and lower maximum velocity in $\mathbb{R}^2$. The upper bound on the approximation factor of the projection median needs to be improved before a similar claim can be made in $\mathbb{R}^3$. Since the lowest possible approximation factor is one, the difference in the approximation factors of S_2 and Π_2 in $\mathbb{R}^2$ from 1.4142 to 1.2732 corresponds to a relative improvement of 34.0%. Similarly, since any bounded-velocity approximation must have velocity at least one, the difference in the maximum velocities of S_2 and Π_2 in $\mathbb{R}^2$ corresponds to the same relative improvement of 34.0%.

Experimentation suggests that the projection median performs well not only in the worst case but also in the average case. Empirical evidence is provided in Sec. 8.4.2 in the form of test results from simulations of sets of 6 clients and 16 clients for which the approximation factors and velocities of the Euclidean 1-median, centre of mass, rectilinear 1-median, projection median, and Gaussian median of a set of mobile clients are measured over 10000 time units. See Figs. 8.11 and 8.12.

All median functions mentioned in this section are defined consistently across dimensions; that is, the position of $\Upsilon_d(P)$ coincides with $\Upsilon_{d-1}(P)$ when the positions of clients in P lie in a $(d-1)$-dimensional flat. All median functions mentioned are invariant under similarity transformations except for the rectilinear 1-median and its convex combinations which are not invariant under rotation or reflection.

Chapter 6

Mobile Euclidean 2-Centre

6.1 Introduction

6.1.1 Chapter Objectives

Chapter 6 examines bounded-velocity approximations to the Euclidean 2-centre.

The static Euclidean 2-centre problem reduces to identifying an optimal partition of the client set and finding a Euclidean 1-centre within each partition. The two subproblems are not independent; rather, given a set of clients P, the coordination of the positions of the two facilities of the Euclidean 2-centre of P corresponds to a global minimum of the Euclidean 2-radius of P. The motion of clients causes discontinuous changes in the optimal partition, resulting in discontinuity in the motion of the mobile Euclidean 2-centre. This discontinuity introduces new challenges in addition to the unbounded velocity inherited from the mobile Euclidean 1-centre.

As a natural progression from our discussion of bounded-velocity approximations of the Euclidean 1-centre, our exploration of approximation functions of the Euclidean 2-centre (referred to as *2-centre functions*) initially leads us to consider the rectilinear 2-centre and 2-means clustering, the respective generalizations of the rectilinear 1-centre and the centre of mass to two facilities. We show that neither of these is continuous. Thus, we are motivated to explore alternative methods for defining 2-centre functions, using techniques different from those presented in Chapter 4.

Although the one-dimensional mobile Euclidean 1-centre is not unique, reflection can be used to define a unique bounded-velocity Euclidean 2-centre. We employ this strategy in two or more dimensions to define *reflection-based 2-centre functions*, a new set of 2-centre functions which we now introduce. The choice of a mobile function for the point of reflection is a critical factor in the eccentricity and maximum velocity of the resulting 2-centre function. We capitalize on our results from Chapter 4 and consider the rectilinear 1-centre, the Steiner centre, the Euclidean 1-centre, and the centre of mass as candidates for the point of reflection.

The main contribution of this chapter is the definition of these reflection-based 2-centre functions, in particular, the Steiner reflection 2-centre and the rectilinear reflection 2-centre, which we show successfully balance the conflicting goals of approximating the Euclidean 2-radius while maintaining a low maximum velocity.

Kinetic algorithms for maintaining these various mobile 2-centre functions are discussed in Ch. 8; for now we focus on their respective qualities as approximation functions. Summaries of the chapter's significant results and their implications are found in Secs. 6.1.2 and 6.7.

6.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Properties of the Mobile Euclidean 2-Centre (Sec. 6.2)
Sec. 6.2 examines additional properties of the mobile Euclidean 2-centre. Specifically, we show that even in one dimension the Euclidean 2-centre (and, furthermore, any geometric 2-centre) is not unique. That is, for some sets of clients P in $\mathbb{R}^d$, two or more pairs of points in $\mathbb{R}^d$ realize the Euclidean 2-radius of P, even when $d = 1$. We show that no one-dimensional Euclidean 2-centre can guarantee relative velocity less than two (unlike the one-dimensional Euclidean 1-centre and 1-median that both have maximum velocity one). Furthermore, we show that the Euclidean 2-centre is discontinuous in two or more dimensions.

Comparison Measures (Sec. 6.3)
Sec. 6.3 expands on the measures of eccentricity and maximum velocity and explores bounds on their relationship in terms specific to the approximation of the Euclidean 2-centre. We show that no bounded-velocity approximation of the Euclidean 2-centre can guarantee eccentricity less than $\sqrt{2}$ or maximum velocity less than $1 + \sqrt{3}/2$ in two or more dimensions (unlike the Euclidean 1-centre for which bounded-velocity approximations exist for any eccentricity $\lambda > 1$ and any maximum velocity $v_{\max} \geq 1$, although these are not necessarily simultaneously achievable by any single centre function).

Single-Facility Approximation Functions (Sec. 6.4)
Sec. 6.4 briefly addresses 2-centre functions whose two facilities always coincide. We show that such a 2-centre function cannot guarantee any bound on eccentricity.

Rectilinear 2-Centre and 2-Means Clustering (Sec. 6.5)
Sec. 6.5 analyzes the properties of the mobile rectilinear 2-centre and the mobile 2-means clustering in terms of their approximation of the Euclidean 2-centre. As seen in Ch. 4, both the rectilinear 1-centre and the centre of mass provide good bounded-velocity approximations of the Euclidean 1-centre, suggesting the corresponding two-facility functions as candidates for approximating the Euclidean 2-centre. We show that like the Euclidean 2-centre, both the rectilinear 2-centre and 2-means clustering are discontinuous.

Reflection-Based 2-Centre Functions (Sec. 6.6)
Sec. 6.6 introduces the idea of defining an approximation to the Euclidean 2-centre of a set of clients P by selecting an arbitrary client $p_0 \in P$ and its

reflection q across a mobile function F_d, where the position of $F_d(P(t))$ is central to $P(t)$. We examine reflection-based 2-centre functions when F_d corresponds to the Euclidean 1-centre, the rectilinear 1-centre, the Steiner centre, and the centre of mass.

We show general lower bounds of 2 on the eccentricity and 3 on the maximum velocity of any reflection-based 2-centre function. When the point of reflection is the Euclidean 1-centre ($F_d = \Xi_d$) we show unbounded velocity and a tight bound of 4 on eccentricity in $\mathbb{R}^d$. When the point of reflection is the rectilinear 1-centre ($F_d = R_d$) we show a tight bound of $2\sqrt{d}+1$ on maximum velocity, an upper bound of $2\sqrt{d}$ on eccentricity, and a lower bound of $2\sqrt{2}$ on eccentricity in $\mathbb{R}^d$. When the point of reflection is the Steiner centre ($F_2 = \Gamma_2$) we show a tight bound of $8/\pi + 1$ on maximum velocity, an upper bound of $8/\pi$ on eccentricity, and a lower bound of $2\sqrt{1+1/\pi^2}$ on eccentricity in $\mathbb{R}^2$. Finally, when the point of reflection is the centre of mass ($F_d = C_d$) we show a tight bound of 3 on maximum velocity and unbounded eccentricity in $\mathbb{R}^d$.

Evaluation (Sec. 6.7)

Sec. 6.7 summarizes the results derived in Ch. 6 by comparison of the various 2-centre functions discussed in terms of their approximation of the Euclidean 2-centre. The primary measures for evaluating the quality of each 2-centre function are eccentricity and maximum velocity (inversely related to stability).

Bounded-Velocity Approximations of the Rectilinear 2-Centre (Sec. 6.8)

Sec. 6.8 briefly addresses the problem of identifying bounded-velocity approximations of the rectilinear 2-centre, where maximum velocity and approximation factor are defined with respect to Chebyshev distance. We show that the rectilinear reflection 2-centre provides a 2-approximation of the rectilinear 2-centre when distance and velocity are measured using the Chebyshev norm.

6.2 Properties of the Mobile Euclidean 2-Centre

This section explores the existence of multiple solutions (non-uniqueness) of the Euclidean 2-centre, establishes a tight bound of two on the maximum velocity (and continuity) of the Euclidean 2-centre in one dimension, and demonstrates the discontinuity (and unbounded velocity) of the mobile Euclidean 2-centre in two or more dimensions. Refer to Sec. 2.3.3 for the static definition of the Euclidean k-centre.

Given a set of mobile clients P in $\mathbb{R}^d$, recall that a Euclidean 2-centre of P consists of a set of two mobile facility functions which we denote $\Xi_d(P(t)) = \{\Xi_d^1(P(t)), \Xi_d^2(P(t))\}$. Similarly, a 2-centre function of P is denoted $\Upsilon_d(P(t)) = \{\Upsilon_d^1(P(t)), \Upsilon_d^2(P(t))\}$.

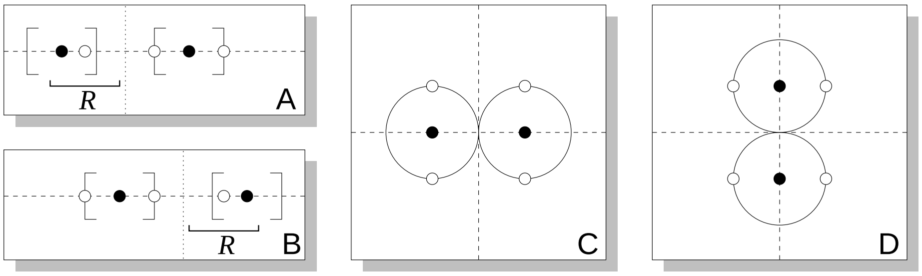

Figure 6.1: non-uniqueness of the geometric 2-centre in $\mathbb{R}$ and the Euclidean 2-centre in $\mathbb{R}^2$

6.2.1 Non-Uniqueness of the Geometric 2-Centre

The Euclidean 1-centre of a finite set of clients P in $\mathbb{R}^d$ is unique for any $d \geq 1$. Similarly, the Euclidean 1-median of a finite multiset of clients P in $\mathbb{R}^d$ is unique for any $d \geq 1$ if either $|P|$ is odd or the clients of P are not all collinear. In general, the Euclidean 2-centre of a finite set P, however, is not unique.

Given a set of clients, P, and a Euclidean 2-centre of P, $\Xi_d^1(P)$ and $\Xi_d^2(P)$, set P can be partitioned into two subsets, P_1 and P_2, such that clients in partition P_i are nearest to facility $\Xi_d^i(P)$, for $i \in \{1, 2\}$. If the Euclidean 2-radius is achieved in one partition only, say P_1, (that is, the local Euclidean radius of P_1 is greater than that of P_2) then there is some connected region R that intersects P_2 such that any point in R may be selected as the position of the facility to serve clients in P_2. For example, let $P = \{0, 1, 2\}$. See Figs. 6.1A and 6.1B. The clients of P can be partitioned either as sets $\{\{1\}, \{2, 3\}\}$ or $\{\{1, 2\}, \{3\}\}$. In either case, the interval R corresponds to a set of values whose choice for the position of the second facility does not increase the Euclidean 2-radius.

In $\mathbb{R}^2$, the choice for partitions might not be unique, even if both partitions achieve the Euclidean 2-radius simultaneously. For example, let $P = \{(1, 1), (1, -1), (-1, 1), (-1, -1)\}$. See Figs. 6.1C and 6.1D. Jaromczyk and Kowaluk [JK95] show a tight bound on the worst-case possible number of ambiguous solutions to the Euclidean 2-centre problem; specifically, for a finite set of clients P in $\mathbb{R}^2$, $|P|$ possible Euclidean 2-centre solutions may exist corresponding to $|P|$ mutually distinct partitions of P (for example, when n clients are positioned at the vertices of a regular n-gon and n is odd).

In the mobile setting, we ask whether there exists some mobile facility whose motion is continuous and whose velocity is bounded while maintaining a bounded approximation factor. Since we approximate the Euclidean 2-radius and not the exact position of the Euclidean 2-centre, the non-uniqueness of the Euclidean 2-centre has no effect on the definition of the approximation factor of a 2-centre function.

In one dimension, there exists a mobile Euclidean 2-centre whose motion is continuous with maximum velocity 2. In two or more dimensions, no mobile Euclidean 2-centre is continuous. We establish these properties formally in the

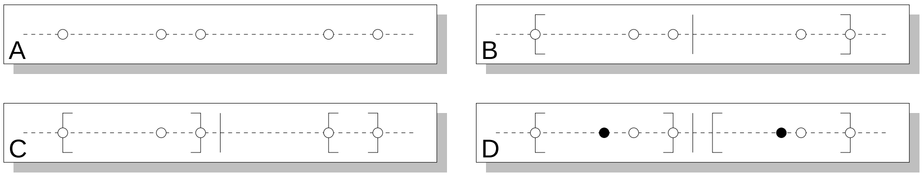

Figure 6.2: one-dimensional algorithm for the mobile geometric 2-centre

next two sections.

6.2.2 Bounded Velocity in One Dimension

We begin by describing an algorithm for finding a Euclidean 2-centre of a finite set of clients P in $\mathbb{R}$. The algorithm is described for a static set of clients but is easily generalized to the mobile setting. Although previous chapters have focused on properties of mobility and approximation and have postponed detailed descriptions of algorithm until Ch. 8, in this instance the location for two facilities returned by the algorithm are used in our proofs of the maximum velocity of the Euclidean 2-centre.

Recall that the Euclidean 2-centre and the rectilinear 2-centre are equivalent in $\mathbb{R}$ (both are geometric 2-centres) since all Minkowski distance metrics are equal in $\mathbb{R}$. Therefore, the following algorithm can be used to find any one-dimensional geometric 2-centre.

The algorithm begins by identifying the Euclidean 1-centre of P, $\Xi_1(P)$, as the midpoint of the extreme points of P. See Fig. 6.2B. Clients are partitioned about $\Xi_1(P)$. Any client whose position coincides with $\Xi_1(P)$ may be included in either partition arbitrarily. The extreme points of each partition are then identified. See Fig. 6.2C. The Euclidean 2-radius of P, r, is determined by the partition of greater diameter. Without loss of generality, assume this is the left partition. Therefore, the location of the first facility , $\Xi_1^1(P)$, must coincide with the Euclidean 1-centre of the left partition. It follows that any point within distance r from the extreme points of the right partition can be selected to define the location of the second facility. A natural choice for selecting the position of the second facility is to employ symmetry and define $\Xi_1^2(P)$ by reflecting $\Xi_1^1(P)$ across the Euclidean 1-centre of P, $\Xi_1(P)$. See Fig. 6.2D. This algorithm is straightforward to implement in $\Theta(n)$ time, where $n = |P|$. See Ch. 8 for mobile implementation details using a KDS.

The algorithm described above is not an approximation algorithm but, rather, it returns the exact positions of a Euclidean 2-centre of P. As we show in Thm. 6.3, although the Euclidean 2-centre is not unique in $\mathbb{R}$, the position returned by this algorithm moves continuously and with maximum velocity two when the clients of P move continuously with maximum velocity one. Furthermore, this velocity is optimal; as we show in Thm. 6.4, in the worst-case, any Euclidean 2-centre moves with relative velocity at least two in $\mathbb{R}$.

To prove an upper bound on the maximum velocity of the Euclidean 2-centre in $\mathbb{R}$ we first prove two lemmas. In Lem. 6.1 we bound the maximum velocity of the facility whose position is defined by reflection across the Euclidean 1-centre. In Lem. 6.2 we bound the maximum velocity of the Euclidean 2-centre when no clients change partitions. These lemmas are used in the proof of Thm. 6.3, showing that the one-dimensional Euclidean 2-centre has maximum velocity two.

Lems. 6.1 and 6.2 and Thm. 6.3 refer to the following definition for function $R : \mathbb{R} \to \mathbb{R}$. Let $R[x(t)]$ denote the reflection of $x(t)$ across $\Xi_1[P(t)]$. That is, $R[x(t)] = 2\Xi_1[P(t)] - x(t)$.

Lemma 6.1. *Let P denote a set of mobile clients in $\mathbb{R}$ defined over a time interval T. Let A denote a subset of P such that for every $p \in A$ and every $t \in T$, $p(t) \leq \Xi_1(P(t))$. The reflection of $\Xi_1(A(t))$ across $\Xi_1(P(t))$ has maximum velocity two.*

Proof. We bound the velocity of $R(\Xi_1[A(t)])$:

$$
\begin{aligned}
\forall t_1, t_2 \in T,\ &||R(\Xi_1[A(t_1)]) - R(\Xi_1[A(t_2)])|| \\
&= ||2\Xi_1[P(t_1)] - \Xi_1[A(t_1)] - 2\Xi_1[P(t_2)] + \Xi_1[A(t_2)]|| \\
&= \left|\left| \left(\min_{p \in P(t_1)} p + \max_{q \in P(t_1)} q \right) - \frac{1}{2} \left(\min_{p \in A(t_1)} p + \max_{q \in A(t_1)} q \right) \right.\right. \\
&\qquad \left.\left. - \left(\min_{p \in P(t_2)} p + \max_{q \in P(t_2)} q \right) + \frac{1}{2} \left(\min_{p \in A(t_2)} p + \max_{q \in A(t_2)} q \right) \right|\right| \\
&= \left|\left| \left(\min_{p \in P(t_1)} p + \max_{q \in P(t_1)} q \right) - \frac{1}{2} \left(\min_{p \in P(t_1)} p + \max_{q \in A(t_1)} q \right) \right.\right. \\
&\qquad \left.\left. - \left(\min_{p \in P(t_2)} p + \max_{q \in P(t_2)} q \right) + \frac{1}{2} \left(\min_{p \in P(t_2)} p + \max_{q \in A(t_2)} q \right) \right|\right| \\
&= \left|\left| \frac{1}{2} \min_{p \in P(t_1)} p - \frac{1}{2} \min_{q \in P(t_2)} q + \max_{r \in P(t_1)} r \right.\right. \\
&\qquad \left.\left. - \max_{s \in P(t_2)} s - \frac{1}{2} \max_{t \in A(t_1)} t + \frac{1}{2} \max_{u \in A(t_2)} u \right|\right| \\
&\leq \frac{1}{2} \left|\left| \min_{p \in P(t_1)} p - \min_{q \in P(t_2)} q \right|\right| + \left|\left| \max_{r \in P(t_1)} r - \max_{s \in P(t_2)} s \right|\right| \\
&\qquad + \frac{1}{2} \left|\left| \max_{t \in A(t_1)} t - \max_{u \in A(t_2)} u \right|\right| \\
&\leq \frac{1}{2} \max_{p \in P} ||p(t_1) - p(t_2)|| + \max_{p \in P} ||p(t_1) - p(t_2)|| \\
&\qquad + \frac{1}{2} \max_{p \in P} ||p(t_1) - p(t_2)|| \\
&= 2 \max_{p \in P} ||p(t_1) - p(t_2)|| \\
&\leq 2|t_1 - t_2|.
\end{aligned}
$$

□

Lemma 6.2. *Let P denote a finite set of mobile clients in $\mathbb{R}$ defined over a time interval T such that at any $t \in T$, no client $p \in P(t)$ has a position that coincides with $\Xi_1(P(t))$. There exists a one-dimensional mobile geometric 2-centre of P with maximum velocity 2.*

Proof. Choose any $t_1, t_2 \in T$. Let the geometric 2-centres of $P(t_1)$ and $P(t_2)$ be defined by the algorithm described above. Let $A(t)$ and $B(t)$ denote the respective left and right partitions of $P(t)$ across $\Xi_1[P(t)]$. Without loss of generality, assume the diameter of $A(t_1)$ is greater than or equal to the diameter of $B(t_1)$.

Since the motion of clients is continuous, the constraint imposed on client positions, $p(t) \neq \Xi_1[P(t)]$, implies that all clients in P remain in their respective partitions for all $t \in T$.

Case 1. Assume the diameter of $A(t_2)$ is greater than or equal to the diameter of $B(t_2)$. Thus, the position of the facility serving $A(t)$ is given by $\Xi_1[A(t)]$ and the position of the facility serving $B(t)$ is given by $R(\Xi_1[A(t)])$ for both $t = t_1$ and $t = t_2$. By Obs. 4.3, $||\Xi_1[A(t_1)] - \Xi_1[A(t_2)]|| \leq |t_1 - t_2|$ and by Lem. 6.1, $||R(\Xi_1[A(t_1)]) - R(\Xi_1[A(t_2)])|| \leq 2|t_1 - t_2|$. Therefore the Euclidean 2-centre of P has maximum velocity two.

Case 2. Assume the diameter of $A(t_2)$ is less than the diameter of $B(t_2)$. Therefore, there exists some $t_3 \in [t_1, t_2]$ such that the diameter of $A(t_3)$ is equal to that of $B(t_3)$. Observe that $\Xi_1[A(t_3)] = R(\Xi_1[B(t_3)])$ and $\Xi_1[B(t_3)] = R(\Xi_1[A(t_3)])$. We bound the velocity of the facility serving partition A:

$$\begin{aligned}
\forall t_1, t_2 \in T,\ & ||\Xi_1[A(t_1)] - R(\Xi_1[B(t_2)])|| \\
&= ||\Xi_1[A(t_1)] - \Xi_1[A(t_3)] + R(\Xi_1[B(t_3)]) - R(\Xi_1[B(t_2)])|| \\
&\leq ||\Xi_1[A(t_1)] - \Xi_1[A(t_3)]|| + ||R(\Xi_1[B(t_3)]) - R(\Xi_1[B(t_2)])|| \\
&\leq |t_1 - t_3| + 2|t_3 - t_2|,
\end{aligned}$$

by Obs. 4.3 and Lem. 6.1,

$$\leq 2|t_1 - t_2|. \tag{6.1}$$

An argument analogous to Eq. (6.1) provides the corresponding bound on the velocity of the facility serving partition B:

$$\forall t_1, t_2 \in T,\ ||R(\Xi_1[B(t_1)]) - \Xi_1[A(t_2)]|| \leq 2|t_1 - t_2|.$$

Therefore, the Euclidean 2-centre of P has maximum velocity two. □

We now remove the restriction on client positions and show this maximum velocity holds for any set of mobile clients in $\mathbb{R}$.

Theorem 6.3. *There exists a one-dimensional mobile geometric 2-centre with maximum velocity 2.*

Proof. Choose any finite set of clients in $\mathbb{R}$ defined over a time interval $T = [0, t_f]$. Let the geometric 2-centres of $P(t)$ be defined by the algorithm described above. For every $t \in T$, let $A(t)$ and $B(t)$ denote the respective left and right partitions of $P(t)$ across $\Xi_1[P(t)]$.

Let n denote the number of clients in P that change partitions at least once over interval T. We say client p changes partitions if there exist $t_1, t_2 \in T$ such that $p(t_1) \in A(t_1)$ and $p(t_2) \in B(t_2)$. We use induction on n.

If no client in P changes partitions, then the result follows from Lem. 6.2. Choose any $k \geq 1$. Assume the one-dimensional mobile geometric 2-centre has maximum velocity 2 when fewer than k clients change partitions. We now prove that the one-dimensional mobile geometric 2-centre has maximum velocity 2 when k clients change partitions.

Let $p \in P$ denote a client that changes partitions at least once. Without loss of generality, assume $p(t_1) \in A(t_1)$ and $p(t_2) \in B(t_2)$ for some $0 \leq t_1 < t_2 \leq t_f$. Let $t_0 \in [t_1, t_2)$ denote a point at which p moves from A to B. That is, assume $p(t_0) \in A(t_0)$ and for all $\delta > 0$ there exists an $\epsilon \in (0, \delta)$ such that $p(t_0 + \epsilon) \in B(t_0 + \epsilon)$. Observe that $p(t_0) = \Xi_1[P(t_0)]$.

Let $f_A(t)$ and $f_B(t)$ denote the positions of the two facilities that serve partitions A and B, respectively, as defined by our one-dimensional 2-centre algorithm. Let $T_1 = [0, t_0]$ and let $T_2 = (t_0, t_f]$. By our inductive hypothesis, f_A and f_B each have maximum velocity 2 over intervals T_1 and T_2. That is,

$$\forall t_1, t_2 \in T_1,\ ||f_A(t_1) - f_A(t_2)|| \leq 2|t_1 - t_2|, \tag{6.2a}$$

and

$$\forall t_3, t_4 \in T_2,\ ||f_A(t_3) - f_A(t_4)|| \leq 2|t_3 - t_4|. \tag{6.2b}$$

The corresponding bounds hold for f_B. In addition,

$$\forall t_1 \in T_1,\ ||f_A(t_1) - f_A(t_0)|| \leq 2|t_1 - t_0|, \tag{6.3a}$$

and

$$\forall t_4 \in T_2,\ \lim_{\epsilon \to 0^+} ||f_A(t_0 + \epsilon) - f_A(t_4)|| \leq 2|t_0 - t_4|. \tag{6.3b}$$

Again, the corresponding bounds hold for f_B.

Since $p(t_0) \in A(t_0)$ and $p(t_0) = \Xi_1[P(t_0)]$, the diameter of $A(t_0)$ must be greater than the diameter of $B(t_0)$. Therefore, our algorithm locates the two facilities at $f_A(t_0) = \Xi_1[A(t_0)]$ and $f_B(t_0) = R(\Xi_1[A(t_0)])$. Similarly, for all $\delta > 0$ there exists an $\epsilon \in (0, \delta)$ such that the diameter of $A(t_0 + \epsilon)$ must be less than the diameter of $B(t_0 + \epsilon)$. Therefore, our algorithm locates the two facilities at $f_A(t_0 + \epsilon) = R(\Xi_1[B(t_0 + \epsilon)])$ and $f_B(t_0 + \epsilon) = \Xi_1[B(t_0 + \epsilon)]$. Since $p(t_0) = \Xi_1[P(t_0)]$, $f_A(t_0) = \lim_{\epsilon \to 0^+} f_A(t_0 + \epsilon)$ and $f_B(t_0) = \lim_{\epsilon \to 0^+} f_B(t_0 + \epsilon)$.

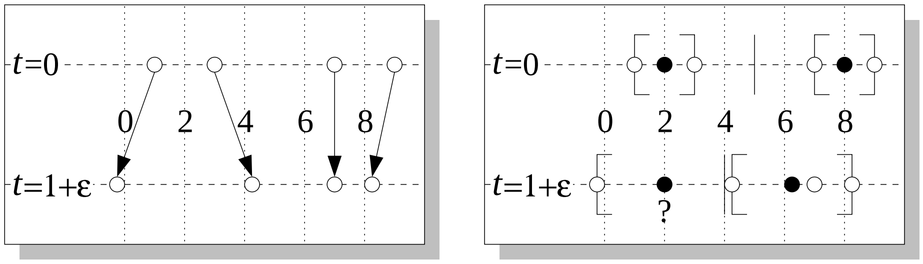

Figure 6.3: illustration in support of Thm. 6.4

Therefore,

$$\begin{aligned}
\forall t_1 \in T_1,\ \forall t_2 \in T_2,\ &||f_A(t_1) - f_A(t_2)|| \\
&= ||f_A(t_1) - f_A(t_0) + f_A(t_0) - f_A(t_2)|| \\
&\leq ||f_A(t_1) - f_A(t_0)|| + ||f_A(t_0) - f_A(t_2)|| \\
&= \lim_{\epsilon \to 0^+} ||f_A(t_1) - f_A(t_0)|| + ||f_A(t_0 + \epsilon) - f_A(t_2)|| \\
&\leq 2|t_1 - t_0| + 2|t_0 - t_2|, && \text{by Eq. (6.3)}, \\
&= 2|t_1 - t_2|.
\end{aligned}$$

The analogous bound holds for f_B. Thus, the one-dimensional mobile geometric 2-centre of P has maximum velocity 2. □

We show this bound is tight by providing an example.

Theorem 6.4. *No one-dimensional mobile geometric 2-centre can guarantee relative velocity less than 2.*

Proof. Let $P(0) = \{1, 3, 7, 9\}$ and let $P(1+\epsilon) = \{-\epsilon, 4+\epsilon, 7, 8+\epsilon\}$ for some $0 < \epsilon \leq 1$. See Fig. 6.3. Observe that there exists a set of four mobile clients with velocity at most one whose positions realize $P(0)$ and $P(1+\epsilon)$. The Euclidean 2-radius of $P(0)$ is easily seen to be 1 by the symmetry of $P(0)$. The unique set of facilities to realize the Euclidean 2-radius of $P(0)$ is $\{2, 8\}$. The Euclidean 1-centre of $P(1+\epsilon)$, $\Xi_1(P(1+\epsilon))$, lies at 4, partitioning the clients of $P(1+\epsilon)$ into two sets. The Euclidean 2-radius of $P(1+\epsilon)$ is 2, realized by the rightmost partition of $P(1+\epsilon)$. Although multiple positions are possible for the facility that serves the left partition, the facility serving the right partition has a unique position at $6+\epsilon$ that realizes the Euclidean 2-radius of 2. Therefore, the respective geometric 2-centres of P at times $t = 0$ and $t = 1+\epsilon$ are given by $\{2, 8\}$ and $\{x, 6+\epsilon\}$, where $x \leq 4$. It follows that some facility must have moved at least $2 - \epsilon$. In the limit as $\epsilon \to 0$, we get a lower bound of 2 on the velocity of the geometric 2-centre in $\mathbb{R}$. □

Both the Euclidean 1-centre and 1-median have maximum velocity one in one dimension and unbounded velocity in two or more dimensions. This property

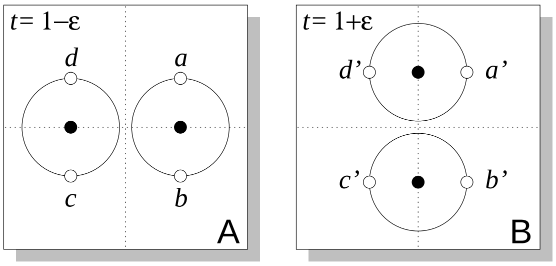

Figure 6.4: illustration in support of Thm. 6.5

may have suggested the more general property that for a fixed k, either a) some Euclidean k-centre (k-median) has maximum velocity one, or b) no Euclidean k-centre (k-median) can guarantee any fixed upper bound on velocity. It is interesting that this is indeed not the case since the tight bound of 2 on the velocity of the one-dimensional Euclidean 2-centre shown in Thms. 6.3 and 6.4 disproves such a hypothesis.

6.2.3 Discontinuity in Two Dimensions

In two or higher dimensions, the Euclidean 2-centre is no longer continuous. We give an example of a set of mobile clients in $\mathbb{R}^2$ for which no continuous Euclidean 2-centre exists.

Theorem 6.5. *The mobile d-dimensional Euclidean 2-centre is discontinuous for $d \geq 2$.*

Proof. Let $P = \{a, b, c, d\}$ denote a set of four mobile clients such that

$$a(t) = \begin{cases} (2-t, 1) & t \leq 1 \\ (1, t) & t > 1 \end{cases}, \qquad b(t) = \begin{cases} (2-t, -1) & t \leq 1 \\ (1, -t) & t > 1 \end{cases},$$

$$c(t) = \begin{cases} (t-2, -1) & t \leq 1 \\ (-1, -t) & t > 1 \end{cases}, \qquad \text{and } d(t) = \begin{cases} (t-2, 1) & t \leq 1 \\ (-1, t) & t > 1 \end{cases}.$$

Observe that each client moves with unit velocity. When $t < 1$, the unique Euclidean 2-centre of $P(t)$ is $\{\Xi_2^1(P(t)), \Xi_2^2(P(t))\} = \{(2-t, 0), (t-2, 0)\}$. See Fig. 6.4A. Similarly, when $t > 1$, the unique Euclidean 2-centre of $P(t)$ is $\{\Xi_2^1(P(t)), \Xi_2^2(P(t))\} = \{(0, t), (0, -t)\}$. See Fig. 6.4B. The corresponding Euclidean 2-radius is one in both instances. It follows that

$$\forall t_1 < 1,\ \forall t_2 > 1,\ ||\Xi_2^i(P(t_1)) - \Xi_2^j(P(t_2))|| \geq \sqrt{2}, \tag{6.4}$$

for any combination of i and j in $\{1, 2\}$. Consequently, the Euclidean 2-centre is discontinuous at $t = 1$. by Def. 3.3. $\square$

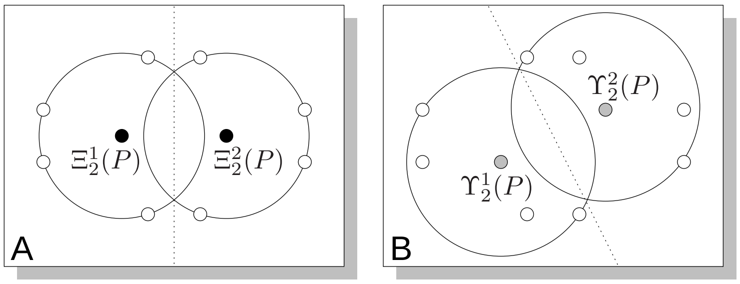

Figure 6.5: different partitions induced by Ξ_2 and Υ_2

6.3 Comparison Measures

This section expands on the comparison measures defined in Ch. 3 in terms specific to bounded-velocity approximations of the Euclidean 2-centre. We examine bounds on the relationship between eccentricity and maximum velocity. In particular, we show that if Υ_d is any λ-eccentric 2-centre function with maximum velocity $v_{\max}$, then $\lambda \geq \sqrt{2}$ and $v_{\max} \geq 1 + \sqrt{3}/2$.

6.3.1 Bounds on Eccentricity and Maximum Velocity

With two facilities instead of one, the definition of eccentricity (approximation factor) includes a minimization over the set of facilities. A bounded velocity approximation now corresponds to a pair of approximation functions Υ_d^1 and Υ_d^2. We say that $\Upsilon_d = \{\Upsilon_d^1, \Upsilon_d^2\}$ is λ-eccentric (equivalently, Υ_d is a λ-approximation of the Euclidean 2-centre) if

$$\forall P \in \mathscr{P}(\mathbb{R}^d),\ \max_{p \in P} \min_{i \in \{1,2\}} ||p - \Upsilon_d^i(P)|| \leq \lambda \max_{q \in P} \min_{j \in \{1,2\}} ||q - \Xi_d^j(P)||. \tag{6.5}$$

For a fixed $t \in T$, the positions of the Euclidean 2-centre induce a partition of the clients in P, such that each client p in P is served by the facility nearest to p. Similarly, Υ_d^1 and Υ_d^2 induce a partition of the clients in P. These two partitions of P are not necessarily identical. See Fig. 6.5.

Thm. 6.4 shows that even in one dimension the Euclidean 2-centre has velocity two in the worst case. In our discussion of bounded-velocity approximations of the Euclidean 1-centre and the Euclidean 1-median, we examine approximation functions that guarantee both a fixed approximation factor and maximum velocity as low as one (for example, the centre of mass). We now show that velocity $1 + \sqrt{3}/2$ is sometimes necessary in order to guarantee *any fixed approximation factor* in $\mathbb{R}^d$ for any $d \geq 2$.

Theorem 6.6. *No mobile 2-centre function in $\mathbb{R}^d$ with maximum velocity less than $1 + \sqrt{3}/2$ can guarantee λ-eccentricity for any fixed $\lambda > 0$ and any $d \geq 2$.*

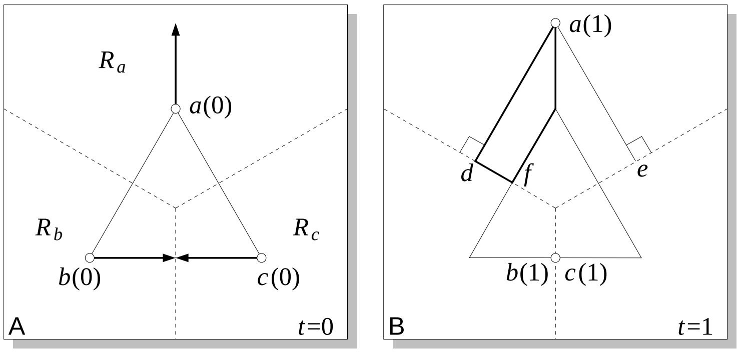

Figure 6.6: illustration in support of Thm. 6.6

Proof. Let $P = \{a, b, c\}$ denote a set of three mobile clients with initial positions (at time $t = 0$) at the vertices of an equilateral triangle in $\mathbb{R}^2$ such that any two clients in P lie a distance two from each other. Let R_a, R_b, and R_c denote the Voronoi regions induced by $a(0)$, $b(0)$, and $c(0)$, respectively. See Fig. 6.6A. Choose any $\Upsilon_2^1(P(t))$ and $\Upsilon_2^2(P(t))$ in $\mathbb{R}^2$ for the positions of the 2-centre function. The interior of at least one of R_a, R_b, or R_3 must be empty of $\Upsilon_2^1(P(0))$ and $\Upsilon_2^2(P(0))$. Without loss of generality assume R_a is empty. Let b and c move toward each other at unit velocity until they meet at their midpoint after one time unit. Let a move away from their midpoint with unit velocity. See Fig. 6.6B. Thus, the Euclidean 2-radius of $P(1)$ is zero. If Υ_2 has any fixed approximation factor, then $\Upsilon_2^1(P(1))$ and $\Upsilon_2^2(P(1))$ must coincide with $a(1)$ and $b(1) = c(1)$. Two points lie nearest to $a(1)$ along the boundary of R_a, which we denote d and e. Since these two cases are symmetric, we examine the left point, d. Let $f = [a(0) + b(0)]/2$. Either $\Upsilon^1(P(1))$ or $\Upsilon^2(P(1))$ must travel from the boundary of R_a to $a(1)$ during the time interval $T = [0, 1]$. This distance is at least as great as the length of the longer edge of the right trapezoid induced by f, $a(0)$, $a(1)$, and d. Angle $\angle da(1)b(1) = \angle fa(0)b(1) = \pi/6$. Since $||f - a(0)|| = ||a(1) - a(0)|| = 1$, it follows that $||d - a(1)|| = 1 + \sqrt{3}/2 \approx 1.8660$. □

Thus, no mobile 2-centre function in $\mathbb{R}^d$ with maximum velocity less than $1 + \sqrt{3}/2$ can guarantee λ-eccentricity for any fixed $\lambda > 0$ and any $d \geq 2$. This property highlights a significant difference between approximations of the Euclidean 1-centre and approximations of the Euclidean 2-centre; in particular, we examined bounded-velocity approximations of the Euclidean 1-centre that guarantee eccentricity 2 while only requiring unit velocity in $\mathbb{R}^d$ for any $d \geq 1$.

With respect to eccentricity λ, Thm. 4.7 by Bereg et al. [BBKS06] show that for every $\lambda > 1$ there is a fixed $v_{\max} \geq 1$ such that there exists an approximation of the Euclidean 1-centre that guarantees eccentricity λ and maximum velocity $v_{\max}$ in $\mathbb{R}^d$ for any $d \geq 1$. Again, the situation differs when approximating the Euclidean 2-centre. As we now prove, if a 2-centre function Υ_d is continuous

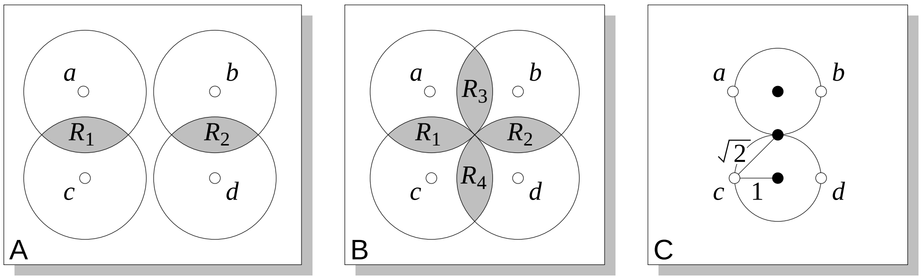

Figure 6.7: illustration in support of Thm. 6.7

(a necessary condition for Υ_d to guarantee any fixed upper bound on velocity) then Υ_d cannot be λ-eccentric for any $\lambda < \sqrt{2}$ in $\mathbb{R}^d$ for any $d \geq 2$ (where λ is independent of $v_{\max}$).

Theorem 6.7. *No continuous mobile 2-centre function in $\mathbb{R}^d$ can guarantee λ-eccentricity for any $\lambda < \sqrt{2}$ and any $d \geq 2$.*

Proof. The result follows from the example described in the proof of Thm. 6.5. If Υ_2 guarantees eccentricity $\lambda = \sqrt{2}$, then for any t there exists a partition of $P(t)$ into two sets P_1 and P_2 such that $\Upsilon_2^1(P(t))$ is contained within the intersection of circles of radius $\sqrt{2}$ centred at each of the clients in P_1 and the same holds for $\Upsilon_2^2(P(t))$ and P_2. These circles have a fixed radius of $\sqrt{2}$ because the Euclidean 2-radius remains one throughout the motion of the clients. When $t < 1 - \sqrt{2}$, a unique partition of $P(t)$ exists such that this intersection is nonempty. We denote the corresponding regions R_1 and R_2. See Fig. 6.7A. The same holds for $P(t)$ when $t > 1 + \sqrt{2}$, for which we denote the corresponding regions R_3 and R_4. At some point t_0, $\Upsilon_2^1(P(t_0))$ and $\Upsilon_2^2((t_0))$ must make a transition from regions R_1 and R_2 over to R_3 and R_4. Since the motion of Υ_2 must be continuous, the transition must occur when the regions overlap. The regions have a unique point of intersection occurring at $t_0 = 1$ at the origin. See Fig. 6.7B. Therefore, $\Upsilon_2^1(P(1)) = \Upsilon_2^2((1)) = (0, 0)$. Thus, the lower bound on eccentricity is realized at time $t = 1$. If the radius of the circles is decreased to less than $\sqrt{2}$, then no such intersection exists. $\square$

6.3.2 Maximum Velocity as a Function of Eccentricity

In Sec. 4.3.2 we examined bounds relating the maximum velocity and eccentricity of a bounded-velocity approximation of the mobile Euclidean 1-centre. Within each partition of the client set, the Euclidean 2-centre behaves locally like a Euclidean 1-centre problem. Although no upper bounds can be concluded from those mentioned in Sec. 4.3.2, the lower bounds examined certainly extend to the Euclidean 2-centre.

In particular, it follows from Thm. 4.6 by Bereg et al. [BBKS06] that any λ-eccentric approximation of the Euclidean 1-centre has velocity at least $1/8\sqrt{\lambda - 1}$.

This result implies a similar bound on the maximum velocity of any λ-eccentric approximation of the Euclidean 2-centre. However, Thm. 6.7 shows that $\lambda \geq \sqrt{2}$, for which function $1/8\sqrt{\lambda - 1}$ has value at most $1/8\sqrt{\sqrt{2} - 1} \approx 0.1942$. Since any 2-centre function with bounded eccentricity must have $v_{\max} \geq 1$, Thm. 4.6 does not further constrain the range of allowable maximum velocities for 2-centre functions.

6.4 Single-Facility Approximation Functions

This section briefly discusses 2-centre functions for which both facilities coincide.

In Chs. 4 and 5 we examined several bounded-velocity approximation functions for a single facility. A single facility can be used to define the positions of two facilities whose positions coincide. We show that no single-facility function, regardless of constraints on its velocity or continuity, can guarantee any fixed approximation factor of the Euclidean 2-centre.

Observation 6.8. *Let Υ_d denote any mobile facility function in $\mathbb{R}^d$. Function Υ_d cannot guarantee eccentricity λ for any fixed λ, for any $d \geq 1$.*

Proof. Let $P = \{0, 1\}$ denote a set of clients in $\mathbb{R}$. The unique Euclidean 2-centre of P coincides with the two client positions of P. The corresponding Euclidean 2-radius is zero. Let $\Upsilon_d(P)$ lie at any point in $\mathbb{R}$. The distance from $\Upsilon_d(P)$ to some client in P must be at least $1/2$. Consequently, no λ exists that satisfies Eq. (6.5). □

Obs. 6.8 implies that no single-facility function can guarantee any approximation of a geometric 2-centre. In Sec. 7.3 we address the related question of whether a $(k+1)$-facility function can provide a bounded-velocity approximation of the geometric k-centre.

6.5 Rectilinear 2-Centre and 2-Means Clustering

As we saw in Ch. 4, the rectilinear 1-centre and the centre of mass both provide bounded-velocity approximations of the mobile Euclidean 1-centre. As we now show, neither of these approximation functions generalizes to define positions of two mobile facilities whose motion is continuous. Recall that 2-means clustering corresponds to the two-facility generalization of the centre of mass. Refer to Sec. 2.5.1 for a definition of the rectilinear 2-centre and to Sec. 2.5.2 for a definition of 2-means clustering.

Corollary 6.9. *The mobile 2-means clustering is discontinuous for $d \geq 2$.*

Proof. The result follows from the example described in the proof of Thm. 6.5. It is straightforward to show that the 2-means clustering for this example is unique and coincides with the Euclidean 2-centre. Consequently, the 2-means clustering is discontinuous in $\mathbb{R}^d$ for $d \geq 2$. □

Corollary 6.10. *The mobile rectilinear 2-centre is discontinuous for $d \geq 2$.*

Proof. The result follows from the example described in the proof of Thm. 6.5 upon rotating the clients of P by $\pi/4$ about the origin. It is straightforward to show that the rectilinear 2-centre for this example is unique and coincides with the Euclidean 2-centre. Consequently, the rectilinear 2-centre is discontinuous in $\mathbb{R}^d$ for $d \geq 2$. □

Velocity and continuity aside, the eccentricity of both the rectilinear 2-centre and 2-means clustering coincide with their corresponding single-facility approximation factors.

Observation 6.11. *The d-dimensional rectilinear 2-centre has eccentricity $\lambda = (1+\sqrt{d})/2$. Furthermore, the d-dimensional rectilinear 2-centre cannot guarantee λ-eccentricity for any λ less than $(1+\sqrt{d})/2$.*

Proof. Let P denote any finite set of clients in $\mathbb{R}^d$. Let $\Xi_d^1(P)$ and $\Xi_d^2(P)$ denote a Euclidean 2-centre of P. Let $R_d^1(P)$ and $R_d^2(P)$ denote a rectilinear 2-centre of P. Let P_1 and P_2 denote the partition of P induced by $\Xi_d^1(P)$ and $\Xi_d^2(P)$ such that $\Xi_d^1(P)$ is the facility closest to any client in P_1 and $\Xi_d^2(P)$ is the facility closest to any client in P_2. If any client p in P is equidistant from $\Xi_d^1(P)$ and $\Xi_d^2(P)$, then assume p is assigned to either partition arbitrarily. The upper bound follows from Thm. 4.9, since the rectilinear 2-radius cannot exceed the maximum of the rectilinear radii of either partition relative to its respective rectilinear 1-centre. Similarly, the lower bound follows from Cor. 4.10 by taking two instances of a client set that realizes the eccentricity of the rectilinear 1-centre and positioning these two sets sufficiently far apart. □

Observation 6.12. *The d-dimensional 2-means clustering has eccentricity 2. Furthermore, d-dimensional 2-means clustering cannot guarantee λ-eccentricity for any λ less than 2.*

Proof. The proof is analogous to the proof of Obs. 6.11 except that references to Thm. 4.9 and Cor. 4.10 refer instead to Lem. 4.13 and Cor. 4.14. □

See Sec. 6.8 for a discussion of bounded-velocity approximations of the rectilinear 2-centre, where maximum velocity and approximation factor are defined with respect to Chebyshev distance.

6.6 Reflection-Based 2-Centre Functions

This section introduces reflection-based 2-centre functions. Four specific functions are defined and analyzed: the Euclidean reflection 2-centre, the rectilinear reflection 2-centre, the Steiner reflection 2-centre, and the mean reflection 2-centre. We derive bounds on the eccentricity and maximum velocity of each or show that none is possible.

6.6.1 Motivation

As shown in Sec. 6.5, neither the rectilinear 2-centre nor 2-means clustering defines a bounded-velocity 2-centre function. That is, solutions employed to provide bounded-velocity approximations of the Euclidean 1-centre cannot be generalized to solve the corresponding problem on two facilities.

As discussed in Ch. 4, any centre function whose position is contained within the convex hull of P guarantees a 2-approximation of the Euclidean 1-centre (see Lem. 4.4). In particular, when the clients of P coincide at a point a, a centre function located within the convex hull of P will also coincide with a. A natural strategy for finding a static approximation to the Euclidean 2-centre problem involves partitioning the clients into two sets and subsequently identifying an approximation to the Euclidean 1-centre of each partition. As mentioned earlier, such strategies generalize poorly to the mobile setting because discontinuities in the position of a mobile 2-centre function can result from changes in the partition of the client set. Thus, a continuous mobile 2-centre function cannot be defined in terms of partitions of the client set.

Nevertheless, if the clients of P form two obvious clusters, then a 2-centre function Υ_d should position one facility close to each cluster. In particular, when the clients of P coincide at two points a and b, the Euclidean 2-radius of P is zero, and $\Upsilon_d^1(P)$ and $\Upsilon_d^2(P)$ must coincide with a and b in order to guarantee any fixed upper bound on eccentricity. When this occurs, observe that any client p_0 in P and its reflection across the midpoint of P coincide with $\{a, b\}$.

As described Sec. 6.2.2, a natural definition for the one-dimensional Euclidean 2-centre of P is provided by viewing the position of the second facility as the reflection of the first facility across the Euclidean 1-centre of P. In one dimension, the first facility can be specified by the position of the Euclidean 1-centre of the cluster with greater diameter. This strategy does not generalize to higher dimensions because of the discontinuity of the Euclidean 2-centre. Furthermore, the unbounded velocity of the Euclidean 1-centre in two or more dimensions precludes it from being used to define a bounded-velocity facility. Instead, we identify a mobile facility function, denoted F_d, that remains central to P while moving under bounded velocity. A client of P, say p_0, is selected arbitrarily and the position of the first facility is set to coincide with that of p_0. The position of the second facility, q, is found by reflecting p_0 across F_d. See Fig. 6.8.

Definition 6.1. *Given a finite set of mobile clients P in $\mathbb{R}^d$, an arbitrarily-selected client p_0 in P, and a centre function F_d, a* **reflection-based 2-centre function** *consists of two facility functions, Υ_d^1 and Υ_d^2, whose positions are given by the position of client $p_0(t)$ and its reflection across $F_d(P(t))$.*

We refer to F as the **reflection function**. We select bounded-velocity approximations of the mobile Euclidean 1-centre as natural candidates for F_d. These include the mobile rectilinear 1-centre, the mobile Steiner centre, and the centre of mass. For comparison, we also examine the case when F_d is the

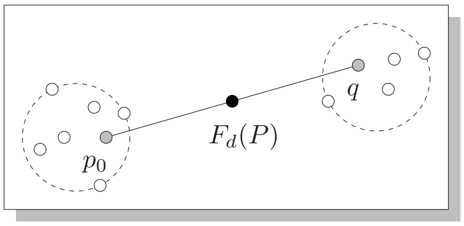

Figure 6.8: When the clients of P form two clusters, client p_0 and its reflection q across $F_d(P)$ define an approximation to the Euclidean 2-centre.

mobile Euclidean 1-centre. We refer to each of these approximation functions respectively as the rectilinear reflection 2-centre, the Steiner reflection 2-centre, the mean reflection 2-centre, and the Euclidean reflection 2-centre.

The definition of a particular reflection-based 2-centre function is consistent across dimensions if the corresponding property holds for the reflection function, as is the case for the rectilinear 1-centre, the Steiner centre, the Euclidean 1-centre, and the centre of mass. In addition, invariance under similarity transformations also follows if the corresponding property holds for the reflection function, given a fixed choice for client p_0. Therefore, the rectilinear reflection 2-centre is invariant under translation and uniform scaling, but not under reflection and rotation whereas the Steiner reflection 2-centre, the Euclidean reflection 2-centre, and the mean reflection 2-centre are invariant under all similarity transformations.

6.6.2 Reflection-Based Approximations: Maximum Velocity

As we now show, tight bounds on relative velocity are straightforward to establish for all four reflection-based approximations we examine. The worst case is achieved when the reflection function F_d and the client p_0 being reflected move toward or away from each other at their respective maximum velocities.

Theorem 6.13. *Let a and b denote mobile clients or mobile facility functions with respective maximum velocities v_a and v_b. The maximum velocity of the reflection of a across b is $2v_b + v_a$. Furthermore, this bound is realizable if the maximum velocities of a and b are simultaneously realizable in opposite directions.*

Proof. Choose any time interval T and any $v_a, v_b > 0$. Choose any functions $a : T \to \mathbb{R}^d$ and $b : T \to \mathbb{R}^d$ such that

$$\forall t_1, t_2 \in T,\ ||a(t_1) - a(t_2)|| \le v_a|t_1 - t_2|,$$

and

$$\forall t_1, t_2 \in T,\ ||b(t_1) - b(t_2)|| \le v_b|t_1 - t_2|.$$

The reflection of $a(t)$ across $b(t)$ corresponds to the function $c(t) = 2b(t) - a(t)$. We bound the velocity of c:

$$\begin{aligned}\forall t_1, t_2 \in T,\ ||c(t_1) - c(t_2)|| &= ||2[b(t_1) - b(t_2)] - [a(t_1) - a(t_2)]|| \\ &\leq 2||b(t_1) - b(t_2)|| + ||a(t_1) - a(t_2)|| \\ &\leq (2v_b + v_a)|t_1 - t_2|.\end{aligned}$$

Therefore, the velocity of c is at most $2v_b + v_a$. As shown by the following example, this bound is realizable in any dimension $d \geq 1$.

Let $a(0) = 1$, $a(1) = 0$, $b(0) = 2$, and $b(1) = 4$. Observe that a and b move with respective average velocities $v_a = 1$ and $v_b = 2$ over the time interval $T = [0, 1]$. The reflection of a across b lies at $c(0) = 3$ and $c(1) = 8$ corresponding to a displacement of $||c(0) - c(1)|| = 5 = 2v_b + v_a$. Therefore, c has average velocity $2v_b + v_a$ over interval T. □

Corollary 6.14. *The d-dimensional Euclidean reflection 2-centre cannot guarantee relative velocity $v_{\max}$ for any fixed $v_{\max}$, for any $d \geq 2$.*

Proof. The result follows from Thms. 4.2 and 6.13. □

Corollary 6.15. *The rectilinear reflection 2-centre has maximum velocity $2\sqrt{d}+1$ in $\mathbb{R}^d$. Furthermore, the rectilinear reflection 2-centre cannot guarantee relative velocity less than $2\sqrt{d}+1$ in $\mathbb{R}^d$, for any $d \geq 1$.*

Proof. The result follows from Obs. 4.12 and Thm. 6.13 because the velocity of R_d is independent of the velocity of p_0 whenever p_0 not an extreme point of P. □

Corollary 6.16. *The Steiner reflection 2-centre has maximum velocity $8/\pi+1$ in $\mathbb{R}^2$ and 6 in $\mathbb{R}^3$. Furthermore, the Steiner reflection 2-centre cannot guarantee relative velocity less than $8/\pi+1$ in $\mathbb{R}^2$ and 6 in $\mathbb{R}^3$.*

Proof. The result follows from Thms. 4.24, 4.25, 4.26, 4.27, and 6.13 because the velocity of Γ_d is independent of the velocity of p_0 whenever p_0 is not an extreme point of P. □

Corollary 6.17. *The mean reflection 2-centre has maximum velocity 3. Furthermore, the mean reflection 2-centre cannot guarantee relative velocity less than 3 in $\mathbb{R}^d$ for any $d \geq 1$.*

Proof. The result follows from Obs. 4.16, Cor. 4.15, and Thm. 6.13. Although the velocity of C_d is not independent of the velocity of p_0, the contribution of p_0 to the velocity of C_d approaches zero as $|P|$ increases. □

6.6.3 Reflection-Based Approximations: Lower Bounds on Eccentricity and Maximum Velocity

We derive a lower bound on the eccentricity of any reflection-based 2-centre function:

Observation 6.18. *Any approximation function for which the position of one facility is set to coincide with the position of a mobile client cannot guarantee eccentricity less than two.*

Proof. The result follows from the the eccentricity of a single client's approximation of the Euclidean 1-centre described in Sec. 4.3.1. See the example for $d = 2$ in Fig. 4.3. □

As discussed in Sec. 4.3.1, the bound of 2 is tight. The bound on maximum velocity follows from By Thm. 6.13:

Corollary 6.19. *No λ-eccentric reflection-based 2-centre function can guarantee relative velocity less than three.*

Proof. A reflection-based 2-centre function Υ_d is defined in terms of a reflection function F_d. If the eccentricity of Υ_d is bounded, then F_d must have maximum velocity at least as great as the maximum velocity of clients. By Thm. 6.13 it follows that no reflection-based approximation function can guarantee relative velocity less than three. □

Consequently, all reflection-based approximation functions have maximum velocity at least three and an approximation factor of at least two.

6.6.4 Reflection Across the Euclidean 1-Centre

Although Cor. 6.14 shows that the Euclidean reflection 2-centre has unbounded velocity in two or more dimensions, we examine its approximation factor. Perhaps surprisingly, reflection across the Euclidean 1-centre results in a 2-centre function with greater eccentricity and higher maximum velocity than both the Steiner reflection 2-centre and the rectilinear reflection 2-centre.

Theorem 6.20. *The two-dimensional Euclidean reflection 2-centre is 4-eccentric.*

Proof. Let P denote any finite set of clients in $\mathbb{R}^2$. Let p_0 denote a client of P whose position corresponds to the first facility, $\Upsilon_2^1(P)$. Let q denote the reflection of p_0 across $\Xi_2(P)$. The position of the second facility, $\Upsilon_2^2(P)$ is given by q. Let C denote the minimum enclosing circle of P and let s denote the radius of C.

Let $\Xi_2^1(P)$ and $\Xi_2^2(P)$ denote a Euclidean 2-centre of P. Let r denote the Euclidean 2-radius of P. Let P_1 and P_2 denote the partition of P induced by $\Xi_2^1(P)$ and $\Xi_2^2(P)$ such that $\Xi_2^1(P)$ is the facility closest to any client in P_1 and $\Xi_2^2(P)$ is the facility closest to any client in P_2. If any client p in P is equidistant

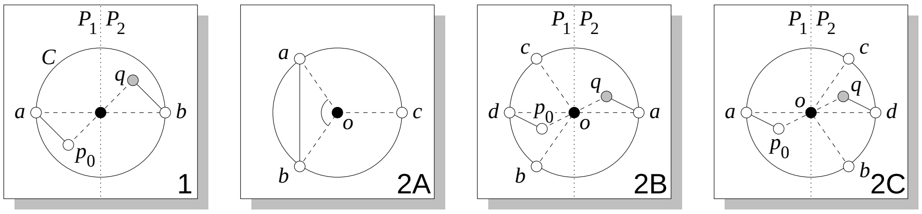

Figure 6.9: illustration in support of Thm. 6.20

from $\Xi_2^1(P)$ and $\Xi_2^2(P)$, then assume p is assigned to either partition arbitrarily. Without loss of generality assume $p_0 \in P_1$. Since $p_0 \in CH(P_1)$,

$$\forall p \in P_1, \; ||p_0 - p|| \leq 2r, \tag{6.6}$$

by Lem. 4.4. Therefore, we need only to verify that $||q - p|| \leq 4r$ for all clients $p \in P_2$.

Case 1. Assume C is supported by two clients $a, b \in P$ that lie opposite each other on C. Clients a and b must lie in opposite partitions, otherwise the Euclidean 2-radius equals the Euclidean radius. Without loss of generality assume $a \in P_1$. See Fig. 6.9(1). Since $b \in CH(P_2)$, for all clients $p \in P_2$, $||p - b|| \leq 2r$ by Lem. 4.4. Observe that $||b - q|| = ||a - p_0||$. Therefore, by Eq. (6.6),

$$\forall p \in P_2, \; ||p-q|| \leq ||p-b||+||b-q|| \leq 2r+||b-q|| = 2r+||a-p_0|| \leq 4r. \tag{6.7}$$

Case 2. Assume no two clients in P lie opposite each other on circle C. Let o denote the centre of C. At least three clients $a, b, c \in P$ must support C such that the angles $\angle aob$, $\angle aoc$, and $\angle boc$ are all less than π. Without loss of generality, assume $\angle boc$ corresponds to the minimum of the three angles. Since the angles sum to 2π, we get $\angle boc \leq 2\pi/3$. Since all angles are less than π, we get $\angle aob \geq \pi/3$ and $\angle aoc \geq \pi/3$. Furthermore, at least one of a, b, or c must lie in each partition.

Case 2a. Assume a and b lie in the same partition. Therefore, c must lie in the opposite partition. See Fig. 6.9(2A). Since $\pi/3 \leq \angle aob < \pi$, it follows that $||a - b|| \geq s$. Consequently, $2r \geq s$. The distance between any two points contained within C is at most $2s$. Therefore,

$$\forall p \in P_2, \; ||p - q|| \leq 2s \leq 4r. \tag{6.8}$$

Since every point inside circle C lies at most $2s$ from any client in P, it follows that q lies at most $2s \leq 4r$ from any client in P.

Case 2b. Assume $c, b \in P_1$ and $a \in P_2$. See Fig. 6.9(2B). Let d denote the point opposite a on circle C. Since $\angle aob$, $\angle aoc$, and $\angle boc$ are all less than π, one of c and b must lie above d and the other must lie below d on circle C. Consequently, $d \in MEC(P_1)$. Therefore, for all clients $p \in P_1$, $||p - d|| \leq 2r$ by Lem. 4.4. Similarly, since $a \in CH(P_2)$, for all clients $p \in P_2$, $||p - a|| \leq 2r$ by

Lem. 4.4. Since $p_0 \in CH(P_1)$, for all points $p \in MEC(P_1)$, $||p - p_0|| \leq 2r$. In particular, $||d - p_0|| \leq 2r$. Observe that $||d - p_0|| = ||a - q||$. Thus,

$$\forall p \in P_2,\ ||p - q|| \leq ||p - a|| + ||a - q|| = ||p - a|| + ||d - p_0|| \leq 4r. \tag{6.9}$$

Case 2c. Assume $c, b \in P_2$ and $a \in P_1$. See Fig. 6.9(2C). Let d denote the point opposite a on circle C. Since $\angle aob$, $\angle aoc$, and $\angle boc$ are all less than π, one of c and b must lie above d and the other must lie below d on circle C. Consequently, $d \in MEC(P_2)$. Therefore, for all clients $p \in P_2$, $||p - d|| \leq 2r$ by Lem. 4.4. Similarly, since $a \in CH(P_1)$ for all clients $p \in P_1$, $||p - a|| \leq 2r$ by Lem. 4.4. By Eq. (6.6), $||a - p_0|| \leq 2r$. Observe that $||a - p_0|| = ||d - q||$. Thus,

$$\forall p \in P_2,\ ||p - q|| \leq ||p - d|| + ||d - q|| = ||p - d|| + ||a - p_0|| \leq 4r. \tag{6.10}$$

Case 2d. Assume $c, a \in P_1$ and $b \in P_2$. This case is analogous to Case 2b since we have not made any assumptions to differentiate a from b.

Case 2e. Assume $c, a \in P_2$ and $b \in P_1$. This case is analogous to Case 2c since we have not made any assumptions to differentiate a from b.

The result follows from Eq. (6.6) through Eq. (6.10). □

Theorem 6.21. *The d-dimensional Euclidean reflection 2-centre cannot guarantee λ-eccentricity for any λ less than 4 when $d \geq 2$.*

Proof. Let $\theta \in (0, \pi/4)$. Let $P = \{p_0, p_1, p_2, p_3\}$ where $p_0 = (-\cos\theta, -\sin\theta)$, $p_1 = (-1, 0)$, $p_2 = (1, 0)$, and $p_3 = (\cos\theta, -\sin\theta)$. The Euclidean 1-centre of P lies at the origin. The unique Euclidean 2-centre of P lies at $(p_0 + p_1)/2$ and $(p_2 + p_3)/2$. Let the first facility, $\Upsilon_2^1(P)$, coincide with p_0. Let q denote the reflection of p_0 across $\Xi_2(P)$. The position of the second facility, $\Upsilon_2^2(P)$ is given by q. See Fig. 6.10.

The Euclidean 2-radius is $\frac{1}{2}\sqrt{(1 - \cos\theta)^2 + \sin^2\theta} = \frac{1}{2}\sqrt{2(1 - \cos\theta)}$. The furthest client from q is p_2, separated by a distance of $2\sin\theta$. It follows that

$$\lambda \geq \lim_{\theta \to 0^+} \frac{2\sin\theta}{\frac{1}{2}\sqrt{2(1 - \cos\theta)}} = 4\sqrt{\lim_{\theta \to 0^+} \frac{\sin^2\theta}{2(1 - \cos\theta)}} = 4.$$ □

6.6.5 Reflection Across the Rectilinear 1-Centre

This section examines properties of the mobile rectilinear reflection 2-centre as an approximation to the mobile Euclidean 2-centre. Refer to Sec. 2.5.1 for a definition of the rectilinear 1-centre.

Given the unbounded velocity of the Euclidean reflection 2-centre, bounded-velocity centre functions provide natural candidates for defining the reflection function F_d. We first consider the the rectilinear 1-centre of P, denoted $R_d(P)$, which we examined as a bounded-velocity approximation of the mobile Euclidean 1-centre in Sec. 4.4. The properties of low maximum velocity and low

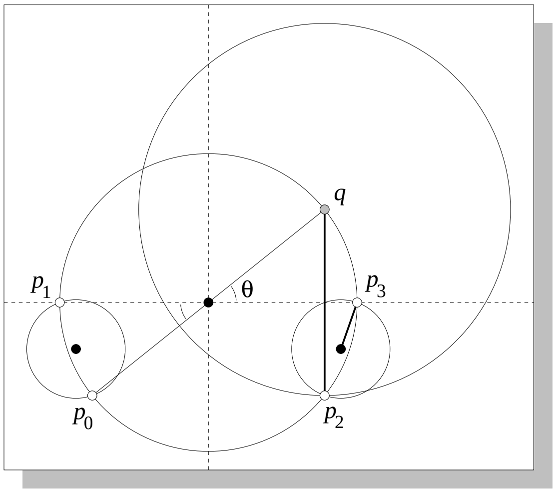

Figure 6.10: illustration in support of Thm. 6.21

eccentricity exhibited by the rectilinear 1-centre suggest it as a natural candidate for defining the point of reflection in a reflection-based 2-centre function.

In Sec. 6.6.6 we consider reflection across the Steiner centre and discover that both the rectilinear reflection 2-centre and the Steiner reflection 2-centre have lower eccentricity than the Euclidean reflection 2-centre.

As shown in Cor. 6.15, the d-dimensional rectilinear reflection 2-centre has velocity $2\sqrt{d} + 1$ in the worst case.

We first bound the eccentricity of the one-dimensional rectilinear reflection 2-centre in Lem. 6.22. This result allows us to derive a bound for a general d in $\mathbb{R}^d$ in Thm. 6.23.

Lemma 6.22. *The one-dimensional rectilinear reflection 2-centre is 2-eccentric.*

Proof. Let P denote any finite set of clients in $\mathbb{R}$. Let p_0 denote a client of P whose position corresponds to the first facility, $\Upsilon_1^1(P)$. Let q denote the reflection of p_0 across $R_1(P)$. The position of the second facility, $\Upsilon_1^2(P)$ is given by q.

Let P_1 and P_2 denote the partition of P induced by clients positioned respectively to the left and right of $\Xi_1(P)$. If any client p in P coincides with $\Xi_1(P)$, then assume p is assigned to partition P_1. There exists a Euclidean 2-centre of P, $\Xi_1^1(P)$ and $\Xi_1^2(P)$, such that $\Xi_1^1(P)$ is the facility closest to any client in P_1 and $\Xi_1^2(P)$ is the facility closest to any client in P_2. Let d denote the maximum of the diameters P_1 and P_2. It follows that $d = 2r$, where r denotes the Euclidean 2-radius of P. Without loss of generality, assume $p_0 \in P_1$. Therefore,

$$\max_{p \in P_1} ||p - p_0|| \leq d \quad \text{and} \quad \max_{p \in P_2} ||p - q|| \leq d,$$
$$\Rightarrow \max_{p \in P} \min_{i \in \{1,2\}} ||p - \Upsilon_1^i(P)|| \leq 2r. \qquad \square$$

Theorem 6.23. *The d-dimensional rectilinear reflection 2-centre is $2\sqrt{d}$-eccentric.*

Proof. Let P denote any finite set of clients in $\mathbb{R}^d$. Let p_0 denote a client of P whose position corresponds to first facility, $\Upsilon_d^1(P)$. Let q denote the reflection of p_0 across $R_d(P)$. The position of the second facility, $\Upsilon_d^2(P)$ is given by q.

Recall that the d-dimension rectilinear 1-centre of P is defined by finding the one-dimensional rectilinear 1-centre of P in each dimension. That is, $R_d(P)_i = R_1(P_i)$, where $P_i = \{p_i \mid p \in P\}$. Since $q = 2R_d(P) - p_0$, therefore $q_i = 2R_1(P_i) - [p_0]_i$. Consequently,

$$\begin{aligned}
\max_{p\in P} \min_{j\in\{1,2\}} ||p - \Upsilon_d^j(P)|| &= \max_{p\in P} \min_{j\in\{1,2\}} \sqrt{\sum_{i=1}^{d} |p_i - \Upsilon_d^j(P)_i|^2} \\
&\leq \sqrt{\sum_{i=1}^{d} \left[\max_{p\in P} \min_{j\in\{1,2\}} |p_i - \Upsilon_d^j(P)_i|\right]^2} \\
&\leq \sqrt{\sum_{i=1}^{d} (2r_i)^2},
\end{aligned}$$

by Lem. 6.22, where r_i denotes the Euclidean 2-radius of P_i,

$$\begin{aligned}
&\leq \max_{1\leq i\leq d} \sqrt{d(2r_i)^2} \\
&= \max_{1\leq i\leq d} 2r_i\sqrt{d} \\
&\leq 2r\sqrt{d}, \qquad (6.11)
\end{aligned}$$

where r denotes the Euclidean 2-radius of P. □

Theorem 6.24. *The two-dimensional rectilinear reflection 2-centre cannot guarantee λ-eccentricity for any λ less than $2\sqrt{2}$.*

Proof. Let $P = \{(-2,0), (-2,-2), (0,2), (2,2)\}$. Let $p_0 = (-2,0)$. The unique Euclidean 2-centre of P has positions $(-2,1)$ and $(1,2)$. The Euclidean 2-radius of P is 1. The rectilinear 1-centre of P, $R_2(P)$, is located at the origin. Let $p_0 = (-2,0)$. The reflection of p_0 across $R_2(P)$, denoted q, is located at $(2,0)$. Client $a = (0,2)$ lies a distance $2\sqrt{2}$ from both q and p_0. See Fig. 6.11A. □

Although the lower bound of $2\sqrt{2}$ applies in $\mathbb{R}^d$ for any $d \geq 2$, the worst-case example described in the proof of Thm. 6.24 does not provide an upper bound of $2\sqrt{d}$ when generalized to $\mathbb{R}^d$ for $d \geq 3$. See Fig. 6.11B.

6.6.6 Reflection Across the Steiner Centre

This section examines properties of the mobile Steiner reflection 2-centre as an approximation to the mobile Euclidean 2-centre. Refer to Sec. 4.6 for a definition of the Steiner centre.

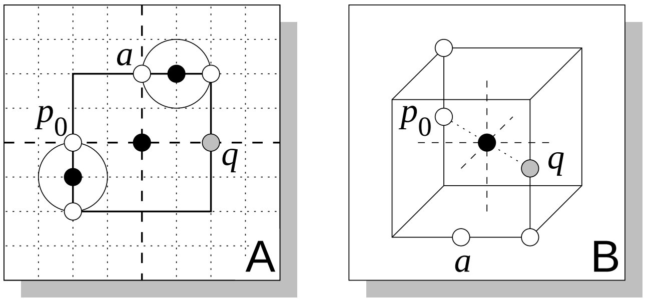

Figure 6.11: illustration in support of Thm. 6.24

We examined the Steiner centre of P, denoted $\Gamma_d(P)$, as a bounded-velocity approximation of the mobile Euclidean 1-centre in Sec. 4.6. The properties of low maximum velocity and low eccentricity exhibited by the Steiner centre suggest it as a natural candidate for defining the point of reflection in a reflection-based 2-centre function.

As shown in Cor. 6.16, the two-dimensional Steiner reflection 2-centre has velocity $8\pi+1$ and the three-dimensional Steiner reflection 2-centre has velocity 6 in the worst case.

Let d_ϕ denote the ℓ_∞ norm relative to a rotation by ϕ of the reference axis. That is, $d_\phi(x) = ||f_\phi(x)||_\infty$, where f_ϕ is a clockwise rotation about the origin by ϕ. Let $R_\phi(P) = f_\phi^{-1}(R_2(f_\phi(P)))$ denote the rectilinear 1-centre with respect to norm d_ϕ. As shown in Lem. 4.18, the Steiner centre of a set of clients P in $\mathbb{R}^2$ can be defined as the limit of the convex combinations of the rotated rectilinear 1-centres of P. That is, $\Gamma_2(P) = \frac{2}{\pi}\int_0^{\pi/2} R_\theta(P)\, d\theta$.

Theorem 6.25. *The two-dimensional Steiner reflection 2-centre is $8/\pi$-eccentric.*

Proof. Let P denote any finite set of clients in $\mathbb{R}^2$. Let p_0 denote a client of P whose position corresponds to the first facility, $\Upsilon_2^1(P)$. Let q denote the reflection of p_0 across $\Gamma_2(P)$. The position of the second facility, $\Upsilon_2^2(P)$ is given by q.

Let $\Xi_2^1(P)$ and $\Xi_2^2(P)$ denote a Euclidean 2-centre of P. Let r denote the Euclidean 2-radius of P. Let P_1 and P_2 denote the partition of P induced by $\Xi_2^1(P)$ and $\Xi_2^2(P)$ such that $\Xi_2^1(P)$ is the facility closest to any client in P_1 and $\Xi_2^2(P)$ is the facility closest to any client in P_2. If any client p in P is equidistant from $\Xi_2^1(P)$ and $\Xi_2^2(P)$, then assume p is assigned to either partition arbitrarily. Without loss of generality assume $p_0 \in P_1$. Since $p_0 \in CH(P_1)$,

$$\forall p \in P_1,\ ||p_0 - p|| \leq 2r, \tag{6.12}$$

by Lem. 4.4. Therefore, we need only to verify that $||q-p|| \leq (8/\pi)r$ for all clients $p \in P_2$.

As shown in the proof of Thm. 6.24 with respect to the rectilinear reflection

2-centre, if $q_R = 2R_2(P) - p_0$, then

$$\max_{p \in P_2} |p_x - [q_R]_x| \leq 2r \quad \text{and} \quad \max_{p \in P_2} |p_y - [q_R]_y| \leq 2r.$$

That is, every client in P_2 is contained within a box of width and height $4r$ centred at q_R. Let $q_\theta = 2R_\theta(P) - p_0$. It follows that every client in P_2 is contained within a box of width and height $4r$, whose axis is rotated by θ relative to the x-axis, and whose centre is q_θ. Consequently,

$$\max_{p \in P_2} |p_x - [q_\theta]_x| \leq 2\sqrt{2}\cos(\pi/4 - \theta)r. \tag{6.13}$$

We now bound the maximum distance in the x-coordinates from any client in P_2 to q.

$$\begin{aligned}
\max_{p \in P_2} |p_x - q_x| &= \max_{p \in P_2} |p_x - (2\Gamma_2(P)_x - [p_0]_x)| \\
&= \max_{p \in P_2} \left| p_x - \left(2\left[\frac{2}{\pi} \int_0^{\pi/2} R_\theta(P)_x \; d\theta \right] - [p_0]_x \right) \right|, \quad \text{by Lem. 4.18,} \\
&= \max_{p \in P_2} \left| \frac{2}{\pi} \int_0^{\pi/2} p_x - (2R_\theta(P)_x - [p_0]_x) \; d\theta \right| \\
&= \max_{p \in P_2} \left| \frac{2}{\pi} \int_0^{\pi/2} p_x - [q_\theta]_x \; d\theta \right| \\
&\leq \max_{p \in P_2} \frac{2}{\pi} \int_0^{\pi/2} |p_x - [q_\theta]_x| \; d\theta \\
&\leq \frac{2}{\pi} \int_0^{\pi/2} \max_{p \in P_2} |p_x - [q_\theta]_x| \; d\theta \\
&\leq \frac{2}{\pi} \int_0^{\pi/2} 2\sqrt{2}\cos(\pi/4 - \theta)r \; d\theta, \quad \text{by Eq. (6.13),} \\
&= \frac{4r\sqrt{2}}{\pi} \int_0^{\pi/2} \cos(\pi/4 - \theta) \; d\theta \\
&= \frac{8r}{\pi}.
\end{aligned}$$

The Steiner reflection 2-centre is invariant under rotation. Consequently,

$$\max_{p \in P_2} |p_x - q_x| \leq \frac{8r}{\pi} \Rightarrow \max_{p \in P_2} ||p - q|| \leq \frac{8r}{\pi}. \qquad \square$$

Theorem 6.26. *The two-dimensional Steiner reflection 2-centre cannot guarantee λ-eccentricity for any λ less than $2\sqrt{1 + 1/\pi^2}$.*

Proof. Let a continuous arc of clients lie on a unit semicircle centred at the origin on the positive x-axis. Let two clients lie opposite the arc at $a = (-1, 1)$

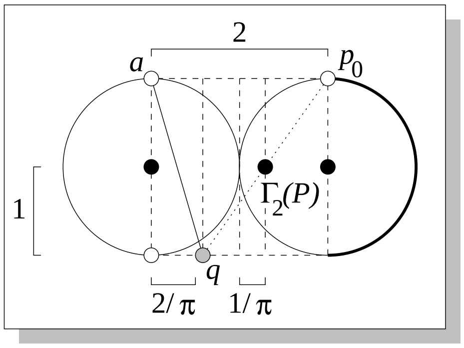

Figure 6.12: illustration in support of Thm. 6.26

and $b = (1,1)$. The unique Euclidean 2-centre of P lies at $(0,0)$ and $(-1,0)$. The corresponding Euclidean 2-radius, r, is one. The Steiner centre of P lies at $(1/\pi - 1/2, 0)$. Let $p_0 = (0,1)$ defines the position of the first facility, $\Upsilon_d^1(P)$. Let q denote the reflection of p_0 across $\Gamma_2(P)$. Observe that $q = (2/\pi - 1, -1)$. The position of the second facility, $\Upsilon_d^2(P)$ is given by q. See Fig. 6.12.

The reflection of q across $\Gamma_2(P)$ lies a distance $\sqrt{2^2 + (2/\pi)^2}$ from client a. □

6.6.7 Reflection Across the Centre of Mass

Given that we examined the centre of mass as a bounded-velocity approximation of both the Euclidean 1-centre in Sec. 4.5 and the Euclidean 1-median in Sec. 5.5, one might naturally consider the centre of mass to define the position of the reflection function F_d. As we show in this section, the resulting reflection-based approximation function cannot guarantee any fixed degree of approximation. Refer to Sec. 2.5.2 for a definition of the centre of mass.

As shown in Cor. 6.17, the d-dimensional mean reflection 2-centre has velocity 3 in the worst case.

Theorem 6.27. *The d-dimensional mean reflection 2-centre cannot guarantee eccentricity λ for any fixed λ, for any $d \geq 1$.*

Proof. We define a set P of three clients in $\mathbb{R}$. Let two clients be located at the origin and let a single client a be located at 1. The centre of mass of P lies at $C_d(P) = 1/3$. The reflected facility has position either $1/3$ or $2/3$, depending on the position of p_0. The Euclidean 2-radius of P is zero. The distance from client a to the nearest facility is at least $1/3$. Consequently, no λ satisfies Eq. (6.5). □

6.7 Evaluation

Following our exploration of 2-centre functions in Secs. 6.5 through 6.4, we have identified and analyzed candidate functions most applicable to defining good bounded-velocity approximations of the Euclidean 2-centre. These are the rectilinear reflection 2-centre and the Steiner reflection 2-centre. We compare these against each other and against the Euclidean reflection 2-centre, the mean reflection 2-centre, the rectilinear 2-centre, 2-means clustering, any two clients in P, and single-facility centre functions.

Lower Bounds on Eccentricity and Maximum Velocity
In Sec. 6.3 we showed that if Υ_d is any 2-centre function with guaranteed fixed bounds on eccentricity, λ, and maximum velocity, $v_{\max}$, then $\lambda \geq \sqrt{2}$ and $v_{\max} \geq 1 + \sqrt{3}/2$. This result may seem surprising since no such lower bounds exist for approximations of the Euclidean 1-centre or the Euclidean 1-median. In fact, for every $\lambda > 1$ there exists a bounded-velocity approximation of the Euclidean 1-centre with approximation factor at most λ.

In Sec. 6.6.3 we showed that if Υ_d is any reflection-based 2-centre function with guaranteed fixed bounds on eccentricity, λ, and maximum velocity, $v_{\max}$, then $\lambda \geq 2$ and $v_{\max}3$.

Rectilinear 2-Centre and 2-Means Clustering
In Sec. 6.5 we showed that although the rectilinear 2-centre and 2-means clustering have eccentricity $(1 + \sqrt{2})/2$ and 2, respectively, both are discontinuous in two or more dimensions, and, consequently, cannot provide bounded-velocity approximations of the Euclidean 2-centre in $\mathbb{R}^d$ for any $d \geq 2$. As mentioned in Sec. 2.5.1, the rectilinear 2-centre is not invariant under under rotation or reflection. It is, however, invariant under translation and scaling. The 2-means clustering is invariant under all similarity transformations. The definition of both 2-centre functions is consistent across dimensions.

Rectilinear Reflection 2-Centre
In Sec. 6.6.5 we examined the rectilinear reflection 2-centre. In $\mathbb{R}^d$, we showed an upper bound of $2\sqrt{d}$ and a lower bound of $2\sqrt{2}$ on the eccentricity of the rectilinear reflection 2-centre for $d \geq 2$. When $d = 2$, the upper and lower bounds coincide at $2\sqrt{2}$. For $d \geq 3$, the bounds diverge. Still in $\mathbb{R}^d$, we showed a tight bound of $1+2\sqrt{d}$ on the maximum velocity of the the rectilinear reflection 2-centre for any $d \geq 1$. The rectilinear reflection 2-centre is not invariant under under rotation or reflection. It is, however, invariant under translation and scaling. Finally, its definition is consistent across dimensions.

Steiner Reflection 2-Centre
In Sec. 6.6.6 we examined the Steiner reflection 2-centre. In $\mathbb{R}^2$, we showed an upper bound of $8/\pi$ and a lower bound of $2\sqrt{1 + 1/\pi^2}$ on the eccentricity of the Steiner reflection 2-centre. In $\mathbb{R}^2$ and $\mathbb{R}^3$, we showed tight bounds of $8/\pi+1$ and 6, respectively, on the maximum velocity of the the Steiner reflection 2-centre. The Steiner reflection 2-centre is invariant under similarity transformations. Finally, its definition is consistent across dimensions.

2-centre function	eccentricity	maximum velocity
Euclidean 2-centre	$\lambda = 1$	$v_{\max} = \infty$
rectilinear 2-centre	$\lambda = (1+\sqrt{2})/2 \approx 1.2071$	$v_{\max} = \infty$
2-means clustering	$\lambda = 2$	$v_{\max} = \infty$
clients $p_1, p_2 \in P$	$\lambda = \infty$	$v_{\max} = 1$
single-facility function	$\lambda = \infty$	
bounded eccentricity		$1+\sqrt{3}/2 \leq v_{\max}$ $\Rightarrow 1.8660 \leq v_{\max}$
continuous	$1.4142 \approx \sqrt{2} \leq \lambda$	
reflection-based 2-centre functions		
point of reflection	**eccentricity**	**maximum velocity**
Euclidean 1-centre	$\lambda = 4$	$v_{\max} = \infty$
centre of mass	$\lambda = \infty$	$v_{\max} = 3$
rectilinear 1-centre	$\lambda = 2\sqrt{2} \approx 2.8284$	$v_{\max} = 2\sqrt{2}+1 \approx 3.8284$
Steiner centre	$2\sqrt{1+1/\pi^2} \leq \lambda \leq 8/\pi$ $\Rightarrow 2.0989 \leq \lambda \leq 2.5465$	$v_{\max} = 8/\pi + 1 \approx 3.5465$
bounded eccentricity		$3 \leq v_{\max}$
any point of reflection	$2 \leq \lambda$	

Table 6.1: comparing 2-centre functions in $\mathbb{R}^2$

Euclidean Reflection 2-Centre

In Sec. 6.6.4 we examined the Euclidean reflection 2-centre. In $\mathbb{R}^2$, we showed a tight bound of 4 on the eccentricity of the Euclidean reflection 2-centre. We showed that the Euclidean reflection 2-centre cannot guarantee any bound on maximum velocity in $\mathbb{R}^d$ for any $d \geq 2$. The Euclidean reflection 2-centre is invariant under similarity transformations. Finally, its definition is consistent across dimensions.

Mean Reflection 2-Centre

In Sec. 6.6.7 we examined the mean reflection 2-centre. In $\mathbb{R}^d$, we showed a tight bound of 3 on the maximum velocity of the mean reflection 2-centre for any $d \geq 1$. We showed that mean reflection 2-centre cannot guarantee any bound on eccentricity in $\mathbb{R}^d$ for any $d \geq 1$. The mean reflection 2-centre is invariant under similarity transformations. Finally, its definition is consistent across dimensions.

Comparison of Approximation Functions

The values for the eccentricity and maximum velocity of these various median functions are displayed in Tab. 6.1 for $\mathbb{R}^2$ and in Tab. 6.2 for $\mathbb{R}^3$.

A scan of Tabs. 6.1 and 6.2 reveals that only two 2-centre functions were identified that have fixed upper bounds on both eccentricity and maximum velocity: the rectilinear reflection 2-centre and the Steiner reflection 2-centre. In $\mathbb{R}^2$, the upper bound on the eccentricity of the Steiner reflection is less than

2-centre function	eccentricity	maximum velocity
Euclidean 2-centre	$\lambda = 1$	$v_{\max} = \infty$
rectilinear 2-centre	$\lambda = (1+\sqrt{3})/2 \approx 1.3660$	$v_{\max} = \infty$
2-means clustering	$\lambda = 2$	$v_{\max} = \infty$
clients $p_1, p_2 \in P$	$\lambda = \infty$	$v_{\max} = 1$
single-facility function	$\lambda = \infty$	
bounded eccentricity		$1 + \sqrt{3}/2 \leq v_{\max}$ $\Rightarrow 1.8660 \leq v_{\max}$
continuous	$1.4142 \approx \sqrt{2} \leq \lambda$	
reflection-based 2-centre functions		
point of reflection	**eccentricity**	**maximum velocity**
Euclidean 1-centre	$\lambda = 4$	$v_{\max} = \infty$
centre of mass	$\lambda = \infty$	$v_{\max} = 3$
rectilinear 1-centre	$2\sqrt{2} \leq \lambda \leq 2\sqrt{3}$ $\Rightarrow 2.8284 \leq \lambda \leq 3.4641$	$v_{\max} = 2\sqrt{3}+1 \approx 4.4641$
Steiner centre	$2.0989 \approx 2\sqrt{1+1/\pi^2} \leq \lambda$	$v_{\max} = 6$
bounded eccentricity		$3 \leq v_{\max}$
any point of reflection	$2 \leq \lambda$	

Table 6.2: comparing 2-centre functions in $\mathbb{R}^3$

the tight bound on the eccentricity of the rectilinear reflection 2-centre. Furthermore, the maximum velocity of the Steiner reflection 2-centre is lower than the maximum velocity of the rectilinear reflection 2-centre. Since the lowest possible eccentricity is $\sqrt{2}$, the difference in the eccentricities of the rectilinear and Steiner reflection 2-centres in $\mathbb{R}^2$ from 2.8284 to 2.5465 corresponds to a relative improvement of 19.9%. Similarly, since any bounded-velocity approximation must have velocity at least $1 + \sqrt{3}/2$, the difference in the maximum velocities of the rectilinear and Steiner reflection 2-centres in $\mathbb{R}^2$ from 3.8284 to 3.5456 corresponds to a relative improvement of 14.4%.

We were unable to show an upper bound on the eccentricity of the Steiner reflection 2-centre in $\mathbb{R}^3$. The rectilinear reflection 2-centre has eccentricity in the range $[2\sqrt{2}, 2\sqrt{d}]$ and maximum velocity $2\sqrt{d}+1$ in $\mathbb{R}^d$ for $d \geq 2$.

Experimentation suggests that the Steiner reflection 2-centre and the rectilinear reflection 2-centre perform well not only in the worst case but also in the average case. Empirical evidence is provided in Sec. 8.4.2 in the form of test results from simulations of sets of 6 clients and 16 clients for which the eccentricities and velocities of the Euclidean 2-centre, rectilinear reflection 2-centre, and Steiner reflection 2-centre of a set of mobile clients are measured over 10000 time units. See Figs. 8.13 and 8.14.

Both the rectilinear reflection 2-centre and the Steiner reflection 2-centre are defined consistently across dimensions; that is, $\{\Upsilon_d^1(P), \Upsilon_d^2(P)\}$ coincides with $\{\Upsilon_{d-1}^1(P), \Upsilon_{d-1}^2(P)\}$ when the positions of clients in P lie in a $(d-1)$-dimensional flat. The reflection 2-centre is invariant under translation and uniform scaling but not under rotation or reflection. The Steiner reflection 2-centre is invariant under all similarity transformations.

6.8 Bounded-Velocity Approximations of the Rectilinear 2-Centre

In Ch. 3 we motivated examining bounded-velocity approximations of the Euclidean 1-centre, the Euclidean 1-median, and the Euclidean 2-centre. As we show in Ch. 7, even in one dimension, no bounded-velocity approximation is possible for any geometric k-centre for $k \geq 3$. Similarly, we show that even in one dimension, no bounded-velocity approximation is possible for any geometric k-median for $k \geq 2$. Both the rectilinear 1-centre and the rectilinear 1-median have bounded velocity. Consequently, our examination of bounded velocity approximations in Chs. 4 through 6 has focused on objective functions defined in terms of the Euclidean distance metric. This leaves open the problem of defining a bounded-velocity approximation of the rectilinear 2-centre, which we show is discontinuous in Cor. 6.10 (the discontinuity applies whether distance is measured using Euclidean or Chebyshev distance). This question is straightforward to address by use of ideas developed in this chapter. As we now show, the rectilinear reflection 2-centre provides a 2-approximation of the rectilinear 2-centre when distance and velocity are measured using the Chebyshev norm.

Recall that the Chebyshev distance between points $x = (x_1, \ldots, x_d)$ and $y = (y_1, \ldots, y_d)$ in $\mathbb{R}^d$ is given by

$$||x - y||_\infty = \max_{1 \leq i \leq d} |x_i - y_i|.$$

It follows that a mobile facility Υ_d has **maximum Chebyshev velocity** $v_{\max}$ if

$$\forall t_1, t_2 \in T, \; ||\Upsilon_d(P(t_1)) - \Upsilon_d(P(t_2))||_\infty \leq v_{\max}|t_1 - t_2|. \tag{6.14}$$

We say 2-centre function Υ_d provides a **Chebyshev $\boldsymbol{\lambda}$-approximation** of the rectilinear 2-centre if

$$\forall P \in \mathscr{P}(\mathbb{R}^d), \; \max_{p \in P} \min_{i \in \{1,2\}} ||p - \Upsilon_d^i(P)||_\infty \leq \lambda \max_{p \in P} \min_{i \in \{1,2\}} ||p - R_d^i(P)||_\infty. \tag{6.15}$$

Corollary 6.28. *The rectilinear reflection 2-centre provides a Chebyshev 2-approximation of the rectilinear 2-centre. Furthermore, the rectilinear reflection 2-centre cannot guarantee a Chebyshev λ-approximation of the rectilinear 2-centre for any $\lambda < 2$.*

Proof. A sufficient condition for Eq. (6.15) is provided if the corresponding bound holds in every dimension. That is,

$$\forall P \in \mathscr{P}(\mathbb{R}^d), \; \forall 1 \leq j \leq d, \; \max_{p \in P} \min_{i \in \{1,2\}} |p_j - \Upsilon_d^i(P)_j| \leq \lambda \max_{p \in P} \min_{i \in \{1,2\}} |p_j - R_d^i(P)_j|. \tag{6.16}$$

Lem. 6.22 shows that the one-dimensional rectilinear reflection 2-centre provides a 2-approximation of a one-dimensional geometric 2-centre. The definition of a Chebyshev λ-approximation coincides with that of a λ-approximation in one dimension. Therefore, Lem. 6.22 implies Eq. (6.16) and the rectilinear reflection 2-centre provides a Chebyshev 2-approximation of the rectilinear 2-centre as defined in Eq. (6.15). The worst case is realized when client p_0 lies at a corner of the bounding box of P. □

Corollary 6.29. *The rectilinear reflection 2-centre has maximum Chebyshev velocity at most 3. Furthermore, the rectilinear reflection 2-centre cannot guarantee relative Chebyshev velocity less than 3.*

Proof. The Chebyshev velocity of a mobile client or mobile facility is realized in one dimension. Therefore, the Chebyshev velocity of the rectilinear 1-centre is at most one. The result follows from Thm. 6.13. □

Chapter 7

Mobile Geometric k-Centre and k-Median

7.1 Introduction

7.1.1 Chapter Objectives

In Ch. 3 we motivated identifying bounded-velocity approximations of the Euclidean k-centre and k-median. In Chs. 4 through 6 we addressed the mobile Euclidean 1-centre, 1-median, and 2-centre problems. For each problem we analyzed existing approximation functions and/or introduced new ones, each of which was evaluated in terms of its maximum velocity and approximation factor. In Chapter 7, we briefly address the mobile Euclidean k-centre for $k \geq 3$ and the mobile Euclidean k-median for $k \geq 2$.

As we show, even in one dimension, no three-facility function can guarantee both continuity and a bounded approximation of the Euclidean 3-centre. Similarly, even in one dimension, no two-facility function can guarantee both bounded velocity and a bounded approximation of the Euclidean 2-median. Compared to results of the previous chapters, the findings of this chapter are modest contributions. Most of the results are straightforward observations included for completeness. The results suggest that evaluation of potential strategies for approximating the Euclidean k-centre for $k \geq 3$ and the Euclidean k-median for $k \geq 2$ requires a different type of analysis, one which falls outside the definitions of approximation considered in this thesis.

Finally, the chapter closes with a brief examination of whether these Euclidean k-centre and k-median problems can be approximated by introducing additional facilities. For example, we ask whether there exists a set of $k + 1$ mobile facilities that provides a bounded-velocity approximation of the mobile Euclidean k-centre.

7.1.2 Chapter Overview

Geometric 3-Centre and Geometric 2-Median (Sec. 7.2)
Sec. 7.2 examines properties of the one-dimensional geometric 3-centre (equivalent to the Euclidean 3-centre in $\mathbb{R}$) and the one-dimensional geometric 2-median (equivalent to the Euclidean 2-median in $\mathbb{R}$). We demonstrate that even in one dimension the geometric 3-centre and 2-median are discontinuous. We show that no bounded-velocity λ-approximation is possible for either the geometric

3-centre or the geometric 2-median in $\mathbb{R}$. In each cases, given any fixed values for λ and $v_{\max}$, a counter-example can be constructed to demonstrate that no set of three (respectively, two) mobile facilities moving with maximum velocity at most $v_{\max}$ can guarantee an approximation factor of λ of the mobile geometric 3-centre (respectively, 2-median).

Additional Facilities (Sec. 7.3)
In Sec. 7.3 we relax the requirement that the number of facilities of an approximation match the number of k-centres or k-medians. We show that a bounded-velocity approximation of the geometric 3-centre is possible with four mobile facilities. However, this result does not generalize; we show that no bounded-velocity approximation of the geometric k-centre is possible with $k+1$ mobile facilities when $k \geq 4$. Similarly, we show that no bounded-velocity approximation of the geometric k-median is possible with $k+1$ mobile facilities when $k \geq 3$.

7.2 Geometric 3-Centre and Geometric 2-Median

7.2.1 Properties of the Mobile 3-Centre and Mobile 2-Median

This section explores the existence of multiple solutions (non-uniqueness) of the Euclidean 3-centre and the Euclidean 2-median and shows that both are discontinuous, even in one dimension. Refer to Sec. 2.3.3 for the static definition of the Euclidean k-centre and to Sec. 2.4.3 for the static definition of the Euclidean k-median.

As shown in Sec. 6.2, the Euclidean 2-centre is not unique in $\mathbb{R}^d$ for any $d \geq 1$. This property extends to the k-centre for any $k \geq 2$. An analogous proof can be used to show that the Euclidean k-median is not unique in $\mathbb{R}^d$ for any $k \geq 2$ and any $d \geq 1$.

With respect to continuity, we now show that both the geometric 3-centre and 2-median are discontinuous in one dimension. Although both the Euclidean 1-median and 2-centre were shown to be discontinuous in Chs. 5 and 6 in two or more dimensions, both problems are continuous (and have bounded velocity) in one dimension.

Discontinuity of the Geometric 3-Centre
We first demonstrate the discontinuity of the geometric 3-centre.

Theorem 7.1. *The mobile d-dimensional geometric 3-centre is discontinuous for any $d \geq 1$.*

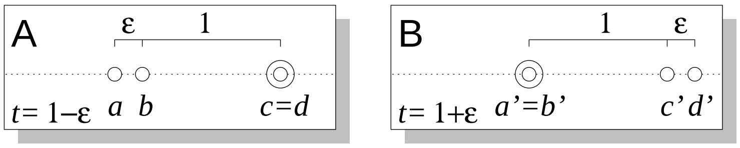

Figure 7.1: illustration in support of Thm. 7.1

Proof. Let $P = \{a, b, c, d\}$ denote a set of four mobile clients such that

$$a(t) = \begin{cases} t-1 & \text{if } t \leq 1 \\ 0 & \text{if } t > 1 \end{cases}, \; b(t) = 0, \quad c(t) = 1, \text{ and } d(t) = \begin{cases} 1 & \text{if } t \leq 1 \\ t & \text{if } t > 1 \end{cases}.$$

Observe that each client moves with at most unit velocity. Fig. 7.1 displays $P(1-\epsilon)$ and $P(1+\epsilon)$ for some $\epsilon > 0$. The geometric 3-radius of $P(t)$ is 0 for all t. When $t \neq 1$, the 3-radius is realized by a unique geometric 3-centre whose positions coincide with $P(t)$. That is, $\{\Xi_1^1(P(t)), \Xi_1^2(P(t)), \Xi_1^3(P(t))\} = \{t-1, 0, 1\}$ for all $t < 1$ and $\{\Xi_1^1(P(t)), \Xi_1^2(P(t)), \Xi_1^3(P(t))\} = \{0, 1, t\}$ for all $t > 1$. It follows that

$$\forall t_1 < 1, \; \forall t_2 > 1, \; ||\Xi_1^i(P(t_1)) - \Xi_1^i(P(t_2))|| \geq 1, \tag{7.1}$$

for some i in $\{1, 2, 3\}$. Consequently, the geometric 3-centre is not continuous by Def. 3.3. □

Discontinuity of the Geometric 2-Median
Similarly, we demonstrate the discontinuity of the geometric 2-median.

Theorem 7.2. *The mobile d-dimensional geometric 2-median is discontinuous for any $d \geq 1$.*

Proof. Let $P = \{a, b, c, d\}$ denote a set of four mobile clients such that

$$a(t) = \begin{cases} t-2 & \text{if } t \leq 1 \\ -1 & \text{if } t > 1 \end{cases}, \quad b(t) = c(t) = 0, \quad \text{and } d(t) = \begin{cases} 1 & \text{if } t \leq 1 \\ t & \text{if } t > 1 \end{cases}.$$

Observe that each client moves with at most unit velocity. Fig. 7.2 displays $P(1-\epsilon)$ and $P(1+\epsilon)$ for some $\epsilon > 0$. The geometric 2-median sum of $P(t)$ is

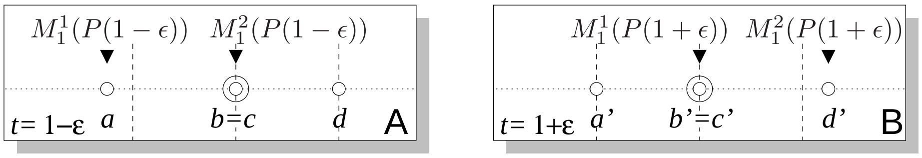

Figure 7.2: illustration in support of Thm. 7.2

1 for all t. When $t \neq 1$, the 2-median sum is realized by a unique geometric 2-median given by $\{M_1^1(P(t)), M_1^2(P(t))\} = \{t-2, 0\}$ for all $t < 1$ and $\{M_1^1(P(1+t)), M_1^2(P(1+t))\} = \{0, t\}$ for all $t > 1$. It follows that

$$\forall t_1 < 1,\ \forall t_2 > 1,\ ||M_1^i(P(t_1)) - M_1^i(P(t_2))|| \geq 1, \tag{7.2}$$

for some i in $\{1, 2\}$. Consequently, the geometric 2-median is not continuous by Def. 3.3. □

7.2.2 Inapproximability of the Mobile 3-Centre and Mobile 2-Median

In this section we show that no continuous λ-approximation of the geometric 3-centre is possible for any fixed $\lambda > 0$ and no bounded-velocity λ-approximation of the geometric 2-median is possible for any fixed $\lambda > 0$. We prove these results in one dimension, implying that the corresponding k-centre and j-median problems are inapproximable in $\mathbb{R}^d$ for any $d \geq 1$, any $k \geq 3$, any $j \geq 2$, and any Minkowski distance metric.

Inapproximability of the Geometric 3-Centre
We begin by demonstrating the inapproximability of the geometric 3-centre. The proof generalizes the example developed in the proof of Thm. 7.1.

Corollary 7.3. *No continuous λ-approximation to the geometric 3-centre exists in $\mathbb{R}$, for any fixed λ.*

Proof. Let P be defined as in the proof of Thm. 7.1. Choose any $\lambda > 0$. Let $\Upsilon_1^1(P(t))$, $\Upsilon_1^2(P(t))$, and $\Upsilon_1^3(P(t))$ denote the positions of three facilities with approximation factor λ. The Euclidean 3-radius of $P(t)$ is zero for all t. It follows that $\{\Upsilon_1^1(P(t)), \Upsilon_1^2(P(t)), \Upsilon_1^3(P(t))\} = \{t-1, 0, 1\}$ for all $t < 1$ and $\{\Upsilon_1^1(P(t)), \Upsilon_1^2(P(t)), \Upsilon_1^3(P(t))\} = \{0, 1, t\}$ for all $t > 1$. Consequently,

$$\forall t_1 < 1,\ \forall t_2 > 1,\ ||\Upsilon_1^i(P(t_1)) - \Upsilon_1^i(P(t_2))|| \geq 1, \tag{7.3}$$

for some i in $\{1, 2, 3\}$. Therefore, Υ_1 is not continuous by Def. 3.3. □

Inapproximability of the Geometric 2-Median
We now demonstrate the inapproximability of the geometric 2-median.

Theorem 7.4. *No bounded-velocity λ-approximation to the geometric 2-median exists in $\mathbb{R}$, for any fixed λ.*

Proof. Choose any $\lambda \geq 1$ and $v_{\max} \geq 1$. Choose any $\epsilon \in (0, 1/v_{\max})$. Choose an integer k such that $k > \max\{\lambda/\epsilon, \lambda/(1 - v_{\max}\epsilon)\}$. Let

$$a(t) = 0, \quad b(t) = -t, \quad \text{and } c(t) = 1.$$

Let P denote a multiset of $2k + 1$ clients such that k clients have positions determined by $a(t)$, k clients have positions determined by $b(t)$, and a single

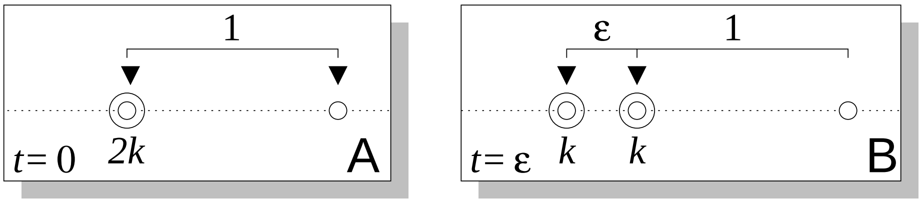

Figure 7.3: illustration in support of Thm. 7.4

client has position determined by $c(t)$. See Fig. 7.3. Observe that each client moves with at most unit velocity.

The geometric 2-median sum of $P(0)$ is zero, realized by the unique geometric 2-median, $\{M_1^1(P(0)), M_1^2(P(0))\} = \{0, 1\}$. Observe that $k > \lambda/\epsilon \geq 1/\epsilon$. Therefore, $k\epsilon > 1$ and the two clusters of k clients must each be served by a different median of $P(\epsilon)$. Thus the geometric 2-median sum of $P(\epsilon)$ is one, realized by the unique geometric 2-median, $\{M_1^1(P(\epsilon)), M_1^2(P(\epsilon))\} = \{-\epsilon, 0\}$.

Let $\Upsilon_1^1(P(t))$ and $\Upsilon_1^2(P(t))$ denote the positions of two facilities with approximation factor λ. Since the geometric 2-median sum is zero at $t = 0$, these two facilities must coincide with the geometric 2-median at $t = 0$. Since the geometric 2-median sum is one at $t = \epsilon$, no facility can lie to the right of λ/k at time $t = \epsilon$ (otherwise the corresponding sum of distances would exceed λ). That is, $\{\Upsilon_1^1(P(0)), \Upsilon_1^2(P(0))\} = \{0, 1\}$ and $\{\Upsilon_1^1(P(\epsilon)), \Upsilon_1^2(P(\epsilon))\} = \{x_1, x_2\}$, where $x_1 \leq \lambda/k$ and $x_2 \leq \lambda/k$. Therefore,

$$|\Upsilon_1^i(P(0)) - \Upsilon_1^i(P(\epsilon)| \geq 1 - \frac{\lambda}{k} > \epsilon v_{\max},$$

for some $i \in \{1, 2\}$. That is, the average velocity of some facility exceeds $v_{\max}$ over the time interval $[0, \epsilon]$. $\square$

7.2.3 Alternate Notions of Approximation

The common definition of approximation which is used in this thesis (Def. 3.5) does not allow for the comparison of two bounded-velocity facility functions in terms of the quality of their respective approximations of the mobile Euclidean k-centre when $k \geq 3$ and k-median when $k \geq 2$.

Client sets that exemplify the inapproximability tend to occur when clients are collocated at k points and any approximation of the geometric k-centre or k-median whose facilities fail to coincide with these k points results in an unbounded approximation factor. Measuring a relative difference instead of a ratio may allow for informed conclusions to be drawn about the quality of facility functions relative to each other.

Natural possibilities for measuring the quality of approximation of the Euclidean k-centre include taking a difference of the optimization functions and normalizing by the Euclidean radius. In the case of the Euclidean k-median, one may consider an additional normalizing factor corresponding to the cardinality of the client set.

Evaluating mobile approximation strategies in terms of these alternatives requires a change of framework that does not fit with the results presented in this thesis; within this work we restrict our analysis to approximation as defined in Def. 3.5. In the following section we briefly consider the question of whether bounded-velocity approximation is possible given additional mobile facilities.

7.3 Additional Facilities

Given that the geometric 2-median and the geometric 3-centre cannot be approximated by a bounded-velocity facility function, we relax the requirement that the number of approximate facilities must correspond to the original number of facilities. That is, we consider approximations of a k-centre or k-median problem by use of $k+1$ or greater mobile facilities. As we show, one additional facility permits approximation for small values of k, but the problems remain inapproximable for most values of k.

7.3.1 Geometric k-Centre with Additional Facilities

We begin by showing that the geometric 3-centre can be approximated with four mobile facilities. Lem. 7.5 is used in the proof of Obs. 7.6.

Lemma 7.5. *Given a set of clients P in $\mathbb{R}$, let p_l and p_r denote clients at the extrema of P and let p_m denote a client closest to $\Xi_1(P)$. There exists a geometric 3-centre of P such that p_l, p_r, and p_m are each served by different facilities, $\Xi_1^1(P)$, $\Xi_1^2(P)$, and $\Xi_1^3(P)$.*

Proof. Let r denote the Euclidean 3-radius of P. Observe that $r \leq |p_r - p_l|/6$. Assume all clients in the interval $[p_l, p_l + 2r]$ are served by $\Xi_1^1(P)$. Similarly, assume all clients in the interval $[p_r - 2r, p_r]$ are served by $\Xi_1^2(P)$. Consequently, all clients in the interval $[p_l + 2r, p_r - 2r]$ are served by $\Xi_1^3(P)$.

Case 1. Assume $p_m \in (p_l + 2r, p_r - 2r)$. Client p_m must be served by $\Xi_1^3(P)$.

Case 2. Assume $p_m \notin (p_l + 2r, p_r - 2r)$. It follows that $P \cap (p_l + 2r, p_r - 2r) = \varnothing$. Consequently, $P \subseteq [p_l, p_l + 2r] \cup [p_r - 2r, p_r]$ and all clients in P are served by $\Xi_1^1(P)$ or $\Xi_1^2(P)$. Therefore $\Xi_1^3(P)$ can be assigned to p_m. □

Observation 7.6. *Given a finite set of mobile clients P in $\mathbb{R}$, there exists a set of four mobile facilities with maximum velocity three that provides a 2-approximation of the geometric 3-centre of P.*

Proof. Since the worst-case approximation factor is realized independently of motion, we consider a static set of clients P in $\mathbb{R}$ to simplify notation. Let the first two facilities, $\Upsilon_1^1(P)$ and $\Upsilon_1^2(P)$, have positions that coincide with the extrema of P. Let the position of the third facility, $\Upsilon_1^3(P)$, coincide with the position of a client of P closest to $\Xi_1(P)$ and let the position of the fourth facility, $\Upsilon_1^4(P)$, be given by the reflection of $\Upsilon_1^3(P)$ across $\Xi_1(P)$. Observe that if two clients are nearest to $\Xi_1(P)$ on opposite sides, then they are each other's reflections across $\Xi_1(P)$.

If $|P| \leq 2$, then $\Upsilon_1^1(P)$ and $\Upsilon_1^2(P)$ coincide with P. Therefore, assume $|P| \geq 3$. Let P_1, P_2, and P_3 denote the partition of P induced by $\Xi_1^1(P)$, $\Xi_1^2(P)$, and $\Xi_1^3(P)$ such that $\Xi_d^i(P)$ is the facility closest to any client in P_i for $i \in \{1,2,3\}$. If any client p in P is equidistant from $\Xi_d^i(P)$ and $\Xi_d^j(P)$, then assume p is assigned to partitions P_i or P_j arbitrarily.

By Lem. 7.5, the leftmost client, the rightmost client, and a client nearest to the midpoint must all lie in different partitions. Within each partition, the Euclidean 3-centre behaves locally like a Euclidean 1-centre problem. It follows from Lem. 4.4 that any client within P_i provides a 2-approximation of Ξ_1^i, for $i \in \{1,2,3\}$.

With respect to a set of mobile clients P, the positions of Υ_1^3 and Υ_1^4 are interchanged periodically. Whenever such an interchange occurs, $\Upsilon_1^3(P(t))$ and $\Upsilon_1^4(P(t))$ are equidistant from $\Xi_1(P(t))$; that is, $\Upsilon_1^3(P(t))$ and $\Upsilon_1^4(P(t))$ correspond to each other's reflection across $\Xi_1(P(t))$. To avoid discontinuity, assume $\Upsilon_1^{3'}(P(t))$ refers to the leftmost point of $\{\Upsilon_1^3(P(t)), \Upsilon_1^4(P(t))\}$ and $\Upsilon_1^{4'}(P(t))$ refers to the rightmost point.

Facilities Υ_1^1 and Υ_1^2 each have velocity at most one since their positions coincide with those of clients in P. Similarly, either $\Upsilon_1^{3'}$ or $\Upsilon_1^{4'}$ has instantaneous velocity at most one since its position coincides with that of some client in P. By Thm. 6.13 and Obs. 4.3, the last facility, $\Upsilon_1^{4'}$ or $\Upsilon_1^{3'}$, has maximum velocity three. □

Although four facilities are sufficient (and necessary) to provide a bounded-velocity approximation of the geometric 3-centre, the corresponding property does not hold for the geometric 4-centre.

Observation 7.7. *No five mobile facilities can guarantee a continuous λ-approximation of the geometric 4-centre in $\mathbb{R}$ for any fixed λ.*

Proof. Choose any $\lambda > 0$ and any $\epsilon \in (0, 1/2)$. Let $P = \{a, b, c, d, e, f\}$ denote a set of six mobile clients such that

$$a(t) = 0, \qquad b(t) = \begin{cases} t - \epsilon & t \leq \epsilon \\ 0 & t > \epsilon \end{cases},$$

$$c(t) = 1, \qquad d(t) = \begin{cases} 1 & t \leq \epsilon \\ 1 + t - \epsilon & \epsilon < t \leq 2\epsilon \\ 1 + 3\epsilon - t & 2\epsilon < t \leq 3\epsilon \\ 1 & t > 3\epsilon \end{cases},$$

$$e(t) = 2, \qquad \text{and } f(t) = \begin{cases} 2 & t \leq 3\epsilon \\ 2 + t - 3\epsilon & t > 3\epsilon \end{cases}.$$

Observe that each client moves with at most unit velocity. See Fig. 7.4. The geometric 4-radius of $P(t)$ is 0 for all t, realized by a unique geometric 4-centre whose positions coincide with $P(t)$ for all $t \notin \{\epsilon, 3\epsilon\}$.

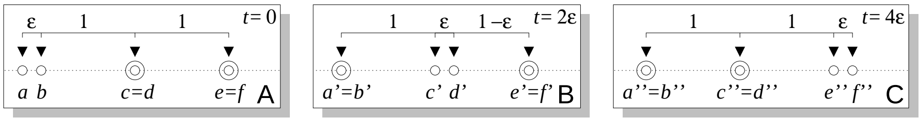

Figure 7.4: illustration in support of Obs. 7.7

By the pigeonhole principle and an argument analogous to the proof of Cor. 7.3, it follows that no set of five bounded-velocity mobile facilities always coincides with the positions of $P(t)$.

If the motions of clients d and f at time $t = \epsilon$ are selected by an adversary as a function of the positions of the five facilities, the claim can be strengthened to state that no set of five continuous mobile facilities always coincides with the positions of $P(t)$. That is, at time $t = \epsilon$, two facilities must be located at $a(\epsilon)$ and $b(\epsilon)$, a third facility must be located at $c(\epsilon) = d(\epsilon)$ and a fourth facility must be located at $e(\epsilon) = f(\epsilon)$. The fifth facility cannot coincide with both $c(\epsilon) = d(\epsilon)$ and $e(\epsilon) = f(\epsilon)$. The adversary selects to move either d or f at time ϵ, whichever lies further away from the fifth facility. □

Corollary 7.8. *No $k+1$ mobile facilities can guarantee a continuous λ-approximation of the geometric k-centre in $\mathbb{R}$ for any $k \geq 4$ and any fixed λ.*

Proof. The argument follows from Obs. 7.7 upon augmenting the client set P with $k - 4$ clients located at $3, 4, \ldots, k - 2$ since each additional client must coincide with a facility. □

7.3.2 Geometric k-Median with Additional Facilities

We show similar bounds on the geometric k-median.

Observation 7.9. *No four mobile facilities can guarantee a bounded-velocity λ-approximation of the geometric 3-median in $\mathbb{R}$ for any fixed λ.*

Proof. Choose any $\lambda \geq 1$ and $v_{\max} \geq 1$. Choose any $\epsilon \in (0, 1/v_{\max})$. Choose an integer k such that $k > \max\{\lambda/\epsilon, \lambda/(1 - v_{\max}\epsilon)\}$. Let

$$a(t) = -1, \qquad b(t) = \begin{cases} t - 1 - \epsilon & t \leq \epsilon \\ -1 & t > \epsilon \end{cases},$$

$$c(t) = 0, \qquad d(t) = \begin{cases} 1 & t \leq \epsilon \\ t + 1 - \epsilon & t > \epsilon \end{cases},$$

and $e(t) = 1$.

Let P denote a multiset of $4k + 1$ clients such that k clients have positions determined by $a(t)$, k clients have positions determined by $b(t)$, a single client has position determined by $c(t)$, k clients have positions determined by $d(t)$, and k clients have positions determined by $e(t)$. See Fig. 7.5. Observe each client moves with at most unit velocity.

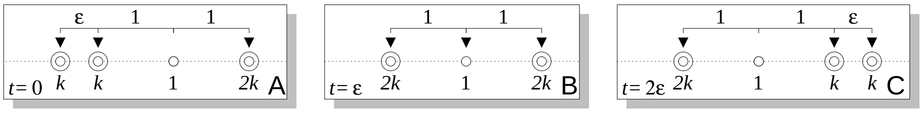

Figure 7.5: illustration in support of Obs. 7.9

By the pigeonhole principle and an argument analogous to to the proof of Thm. 7.4, it follows that the average velocity of some facility exceeds $v_{\max}$ over the time interval $[0, 2\epsilon]$. That is, no four mobile facilities with maximum velocity $v_{\max}$ can coincide with $\{-1-\epsilon, -1, 1\}$ at time $t = 0$, $\{-1, 0, 1\}$ at time $t = \epsilon$, and $\{-1, 1, 1+\epsilon\}$ at time $t = 2\epsilon$. □

Corollary 7.10. *No $k+1$ mobile facilities can guarantee a bounded-velocity λ-approximation of the geometric k-median in $\mathbb{R}$ for any $k \geq 3$ and any fixed λ.*

Proof. The argument follows from Obs. 7.9 upon augmenting the client set P with $k-3$ clients located at $2, 3, \ldots, k-2$ since each additional client must coincide with a facility. □

The question of whether there exists a set of three mobile facilities that provides a bounded-velocity approximation of the geometric 2-median remains open.

The negative results in Obs. 7.7 and 7.9 and Cor. 7.8 and 7.10 imply the corresponding results in $\mathbb{R}^d$ for any $d \geq 1$. The positive result in Obs. 7.6, however, does not imply the corresponding results in higher dimensions.

Finally, it is straightforward to show that the negative results imply similar results for various combinations of facilities k and $k+j$. For example, it follows from Obs. 7.7 and Cor. 7.8 that no eleven mobile facilities can guarantee a bounded-velocity λ-approximation of the geometric 9-centre in $\mathbb{R}$ for any $k \geq 4$ and any fixed λ ($4+5 = 9$ facilities and $5+6 = 11$ facilities). Again, combinations of positive results do not follow from Obs. 7.6.

Chapter 8

Implementation

8.1 Introduction

8.1.1 Chapter Objectives

This chapter provides descriptions of algorithms that implement the key approximation functions discussed in Chs. 4 through 6. Specifically, we describe kinetic data structures for maintaining the two-dimensional Steiner centre, the projection median, and reflection-based 2-centre functions.

A graphical demonstration of these approximation functions was implemented in Java with the dual objective of providing visual intuition of the behaviour of these functions in a mobile setting as well as gathering informal empirical data to confirm that the relative ordering of the respective worst-case theoretical bounds is realized in practice. We provide a brief description of the implementation followed by a short discussion of the data collected. The reader should not draw overly-strong conclusions about bounds on average-case performance from the statistics presented; rather, these are included to provide informal evidence that the worst-case bounds on velocity and approximation factor are not overly pessimistic or unrealistic. In particular, the data suggest that the mobile Steiner centre and the mobile projection median provide better approximations of the mobile Euclidean 1-centre and 1-median, respectively, more often than do the mobile centre of mass, the mobile rectilinear 1-centre, the mobile rectilinear 1-median, and the mobile Gaussian median.

8.1.2 Chapter Overview

Below is a summary of the sections presented in this chapter.

Maintaining Mobile Centre and 2-Centre Functions (Sec. 8.2)

Sec. 8.2 briefly addresses maintaining the mobile centre of mass and refers to related work for KDS algorithms that maintain the mobile rectilinear 1-centre. We describe two KDS algorithms for maintaining the mobile Steiner centre. The first maintains the exact position of the mobile Steiner centre. For the second algorithm we introduce the m-hull, a discretization of the convex hull, allowing us to develop a more efficient algorithm for maintaining a close approximation of the mobile Steiner centre. The section closes with a discussion of algorithms for maintaining reflection-based 2-centre functions.

Maintaining Mobile Median Functions (Sec. 8.3)
Sec. 8.3 begins by referring to related work for KDS algorithms that maintain the mobile rectilinear 1-median. We employ the Steiner centre algorithm described in Sec. 8.2.3, in an algorithm for maintaining a close approximation of the mobile Gaussian median. We then discuss algorithms for finding the static projection median and for maintaining its exact and approximate position in the mobile setting.

Java Visualization (Sec. 8.4)
Sec. 8.4 begins with a brief description of a graphical demonstration of these approximation functions implemented in Java. Next we provide a summary of data collected by the implementation presented in the form of cumulative distribution plots and percentile tables.

8.2 Maintaining Mobile Centre and 2-Centre Functions

In Ch. 4 we identified four bounded-velocity approximations of the mobile Euclidean 1-centre: the Steiner centre, the rectilinear 1-centre, the centre of mass, and a linear combination of the latter two. In Ch. 6 we identified two bounded-velocity approximations of the mobile Euclidean 2-centre: the Steiner reflection 2-centre and the rectilinear reflection 2-centre. We describe algorithms to maintain these that make use of kinetic data structures (KDS). See Sec. 3.7.1 for a discussion of KDS.

The rectilinear 1-centre of a set of mobile clients is maintained by the mobile minima and maxima of each dimension. Agarwal and Har-Peled [AH01] describe an efficient implementation using a KDS. Closely-related is the maintenance of the extent of a set of mobile clients in one dimension [AH01, AGHV01, Gui98, BGH99].

The centre of mass of a set of mobile clients is straightforward to maintain as the average of the positions of all mobile clients. No actual KDS is required since no certificate validation is necessary. If the clients of P have motion that is degree j polynomial, then the motion of $C_d(P(t))$ is also degree j polynomial. The motion of $C_d(P(t))$ requires only a constant-time update whenever a client's flight plan is updated.

Maintaining any convex combination of the mobile centre of mass and the mobile rectilinear 1-centre is achieved by maintaining the motions of both centre functions as described above and returning the convex combination of their trajectories.

We now examine implementation issues involving KDS for the maintenance of both exact and approximate mobile Steiner centres of a set of mobile clients. We describe a simple algorithm to maintain an arbitrarily-close approximation of the Steiner centre of a set of mobile clients by using a KDS to maintain the m-hull of the clients (see Def. 8.1). We show the motion of the Steiner centre of

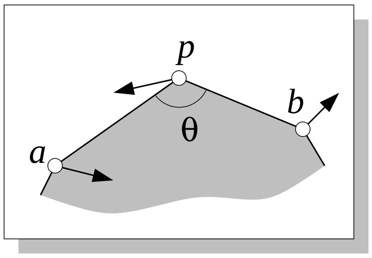

Figure 8.1: The Gaussian weight of client p is defined in terms of the positions of clients p, a, and b.

the m-hull follows a piecewise-linear trajectory if the motion of clients in P is also piecewise-linear. Although the Steiner centre has two equivalent definitions, in this context maintaining the Steiner centre of the m-hull is much simpler by its formulation by Gaussian weights.

8.2.1 Maintaining the Mobile Steiner Centre

The Gaussian weight $w_2(p(t))$ of an extreme point p in P at time t is defined in terms of the position of client p and its two neighbouring clients, a and b, on the convex hull boundary of P. Using the inner product of the vectors $a - p$ and $b - p$, it follows that

$$w_2(p(t)) = \pi - \arccos\left(\frac{\langle a(t) - p(t), b(t) - p(t)\rangle}{||a(t) - p(t)|| \cdot ||b(t) - p(t)||}\right). \tag{8.1}$$

See Fig. 8.1. Even if the motion of clients is linear, function $w_2(p(t))$ remains trigonometric. As a consequence, the position of the Steiner centre is not expressible as a polynomial and its description requires a number of terms proportional to the size of the convex hull, $\Theta(|P|)$. Similarly, even under linear motion of clients, the trajectory of the Euclidean 1-centre Ξ_2 cannot be expressed algebraically. At any given time, the position of $\Xi_2(P(t))$ is defined by at most three clients and, unlike Γ_2, the trajectory of Ξ_2 is expressible by a constant number of terms (while the same three clients define Ξ_2).

Given this constraint on the complexity of a description of the Steiner centre's trajectory, the position of Γ_2 may be maintained by any KDS that maintains the convex hull of a set of mobile clients. For any mobile client p, the description of its Gaussian weight $w_2(p(t))$ changes only when the neighbours of p change along the convex hull boundary or when p joins or leaves the convex hull boundary. Each such update requires only constant time. Therefore, the number of KDS events processed remains unchanged and the complexity of the new KDS is not increased. Thus, a KDS may be used to maintain the Steiner centre with responsiveness, efficiency, locality, and compactness identical to that for maintaining the convex hull. However, the expression for the position of the Steiner centre requires $\Theta(n)$ terms, where $n = |P|$.

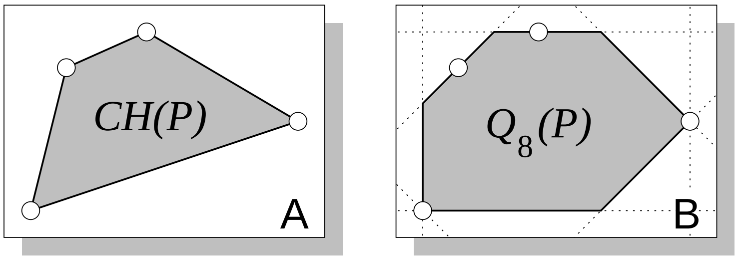

Figure 8.2: the convex hull and the 8-hull of P

8.2.2 The Steiner Centre of the m-Hull

The definition of many centre functions (like the Euclidean 1-centre and the Steiner centre) depends only on extreme points of the set P. Of course, the convex hull of any (possibly infinite) bounded set of clients P can be closely approximated by some finite set of points P'. We formalize this notion by defining the m-hull of a set of clients. We then show that when any set of clients $P \in \mathscr{P}(\mathbb{R}^d)$ is approximated by its m-hull, $Q_m(P)$, the relative distance between $\Gamma_2(P)$ and $\Gamma_2(Q_m(P))$ is $O(\frac{1}{m})$.

Definition 8.1. *Let $P \in \mathscr{P}(\mathbb{R}^2)$ and let $m \in \mathbb{Z}, m \geq 4$, be fixed. The* **m-hull** *of P, denoted $Q_m(P)$, is defined by the intersection of all half-planes H^+ such that $P \subseteq H^+$ and the outer normal to the boundary line of H^+ is $u_\phi = (\cos\phi, \sin\phi)$ for some $\phi = 0 \bmod \frac{2\pi}{m}$.*

See the example in Fig. 8.2 for $m = 8$. The boundary of $Q_m(P)$ is a polygon with at most m sides whose edges have normals parallel to $(\cos(\frac{2\pi j}{m}), \sin(\frac{2\pi j}{m}))$ for some $j \in \mathbb{Z}$. As m increases, the m-hull of P approaches the convex hull of P.

We show that when a client set P is approximated by its m-hull, $Q_m(P)$, the relative distance between $\Gamma_2(P)$ and $\Gamma_2(Q_m(P))$ is $O(\frac{1}{m})$.

Lemma 8.1. *Let $P \in \mathscr{P}(\mathbb{R}^2)$ and let $m \in \mathbb{Z}^+$ be fixed, $m \geq 4$. Let $Q_m(P)$ denote the m-hull of P and let r be the Euclidean radius of P. The distance between $\Gamma_2(P)$ and $\Gamma_2(Q_m(P))$ satisfies*

$$||\Gamma_2(Q_m(P)) - \Gamma_2(P)|| \leq \frac{16r}{\pi m}. \tag{8.2}$$

Proof. Since $\Gamma_2(P) = \Gamma_2(CH(P))$ and the m-hull of P is equal to the m-hull of $CH(P)$, assume without loss of generality that $P = CH(P)$. Choose any $m \in \mathbb{Z}$, $m \geq 4$. Let $Q_m(P)$ denote the m-hull of P. Let r be the Euclidean radius of P. Let f be an ϵ-perturbation of $Q_m(P)$ such that for every $q \in Q_m(P)$, $f(q)$ is a nearest client in P to q (the value of ϵ is chosen below). For every edge l of the boundary of $Q_m(P)$, there is a client $p \in P$ tangent to l. Let a and b be extreme points in P defining adjacent boundary edges l_1 and l_2 on the boundary of $Q_m(P)$. Let point $c \in Q_m(P)$ denote the intersection of l_1 and l_2.

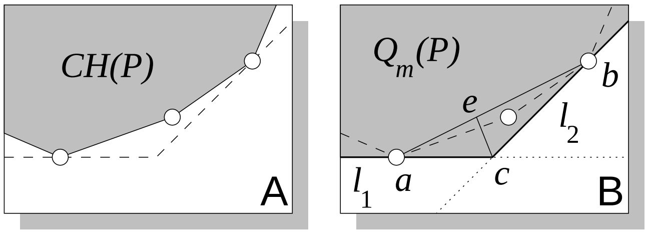

Figure 8.3: illustration in support of Lem. 8.1

See Fig. 8.3B. If $c \in P$ then locally $||f(c) - c|| = 0$. Therefore assume $c \notin P$. The distance from c to line $\overline{ab}$ is maximized when $||a-c|| = ||b-c||$. Since $c \notin P$, angle $\angle acb = \pi - 2\pi/m$. Consequently, $\angle cab = \angle cba = \pi/m$. Since $a,b \in P$, $||a - b|| \leq 2r$. Let e denote the midpoint of $\overline{ab}$. Therefore, $||a - e|| \leq r$ and $||e-c|| \leq r\tan(\pi/m)$. Thus, no point in $Q_m(P)$ may lie farther than $r\tan(\pi/m)$ from the convex hull of P. Therefore, the maximum distance between $Q_m(P)$ and the convex hull of P is at most $\max_{q \in \overline{Q_m(P)}} ||q - f(q)|| \leq r\tan(\pi/m)$. Thus, f is an $r\tan(\pi/m)$-perturbation of $Q_m(P)$. Thus, let $\epsilon = r\tan(\pi/m)$. By Def. 3.7 and Thm. 4.24,

$$||\Gamma_2(Q_m(P)) - \Gamma_2(f(Q_m(P)))|| \leq \frac{4r}{\pi}\tan\left(\frac{\pi}{m}\right). \tag{8.3}$$

Since $P \subseteq Q_m(P)$ and by the definition of f, observe that $\Gamma_2(f(Q_m(P))) = \Gamma_2(P)$. Also note that if $\theta < \pi/4$, then $\tan\theta < 4\theta/\pi$. Therefore,

$$\begin{aligned}||\Gamma_2(P) - \Gamma_2(Q_m(P))|| &= ||\Gamma_2(Q_m(P)) - \Gamma_2(f(Q_m(P)))|| \\ &= \frac{4}{\pi}r\tan\left(\frac{\pi}{m}\right) \\ &\leq \frac{4}{\pi}r\left[\frac{4}{\pi}\frac{\pi}{m}\right] \\ &= \frac{16r}{\pi m}. \qquad \square\end{aligned}$$

The idea of the m-hull is related to that of the *strong convex hull* introduced by Fink and Wood [FW03] and the *convex F-hull* introduced by Rawlins and Wood [RW87]. Both the strong convex hull and the convex F-hull correspond to the smallest convex region that contains a set of points P and whose edges are parallel to a given set of lines. Note, the term *k-hull* has alternative definitions. In particular, Cole et al. [CSY87] define the k-hull of a set of points P such that any hyperplane H passing through a vertex of the k-hull of P has at least k points of P contained within each of the two half-spaces induced by H. Thus, the convex hull is a 1-hull.

8.2.3 Mobile Implementation Using the m-Hull

For implementation it may be desirable to define a mobile centre function that carries the benefits of the Steiner centre but has a simple polynomial description. Under linear motion of clients, we describe a simple discretization using the m-hull that allows the motion of the Steiner centre to be closely approximated by a piecewise-linear function.

Let $Q_m(P)$ denote the m-hull of a set $P \in \mathscr{P}(\mathbb{R}^2)$. See the example in Fig. 8.2 for $m = 8$. The boundary of $Q_m(P)$ is a polygon with at most m sides with turn angles that are multiples of $\frac{2\pi}{m}$ between 0 and π. These correspond to Gaussian weights. Therefore, the Gaussian weight of $q \in Q_m(P)$ is $w_2(q(t)) = j\frac{2\pi}{m}$ for some $j \in \{0, \ldots, \lfloor m/2 \rfloor\}$. Furthermore, while the adjacencies between edges of $Q_m(P)$ to points on the convex hull boundary of P remains unchanged, the Gaussian weight of any extreme point $q \in Q_m(P)$ remains constant. Since the weights are constant, the Steiner centre $\Gamma_2(Q_m(P))$ is simply a linear combination of the vertices of $Q_m(P)$. Therefore, under linear motion of clients of P, between events along the convex hull boundary of $Q_m(P)$, the motion of $\Gamma_2(Q_m(P))$ is also linear (and continuous). In general, the motion of $\Gamma_2(Q_m(P))$ is piecewise-linear.

Maintaining the mobile m-hull of P in a KDS is simple. Associated with the m-hull are m normal vectors, $u_\phi = (\cos\phi, \sin\phi)$, where ϕ is drawn from the set of m angles $\Phi = \{j\frac{2\pi}{m} \mid 0 \le j \le m-1\}$. For each $\phi \in \Phi$, let $P_\phi = \{u_\phi \langle p, u_\phi \rangle \mid p \in P\}$ be the projection of P onto the line through the origin that lies parallel to u_ϕ. We maintain the maximum client in each of the m sets P_ϕ. As described by Guibas [Gui98], a KDS that maintains the maximum of a set of clients in $\mathbb{R}$, each moving with linear motion, is responsive, efficient, compact, and local. Under linear motion the maximum client of each set P_ϕ changes at most $n = |P|$ times. We require maintaining m instances of this KDS. Therefore, the total number of times a maximum client changes is at most $m \cdot n$.

The set of m maximum clients defines the m-hull, $Q_m(P)$, and ultimately the Steiner centre of $Q_m(P)$, $\Gamma_2(Q_m(P))$. Associated with each maximum client is a tangent line with normal u_ϕ. These lines are ordered and we maintain the m intersection points that define the boundary of the m-hull (intersection points may be collocated resulting in fewer than m points). Since the clients of P move linearly, the motion of the intersection points is also linear. Furthermore, an intersection point only requires updating whenever the maximum client of one of its defining lines is updated. For each such event, the Gaussian weight of a client on the boundary of $Q_m(P)$ requires a constant-time update. Between events, weights of clients in $Q_m(P)$ remain constant.

Although Richardson [Ric97] provides an approximation of the convex hull of P to within $O(1/m^2)$ (measured by the Hausdorff distance between the two hulls) while requiring at most m vertices, the m-hull has the advantage that interior angles at the vertices of $Q_m(P)$ are multiples of $2\pi/m$. Consequently, maintaining the kinetic m-hull is straightforward and only requires maintaining the m supporting planes with outer normals $j \cdot 2\pi/m$, for $j = 0 \ldots m-1$.

In summary, given a set of mobile clients P each moving in linear trajectories, the Steiner centre of P does not move with algebraic motion. However, the m-hull allows the maintenance of an approximation to the mobile Steiner centre of P that moves with piecewise-linear motion. Furthermore, m can be selected independently of $|P|$ to ensure the approximate Steiner centre is made arbitrarily close to the Steiner centre of P. Finally, maintaining the m-hull and approximate Steiner centre of a set of mobile clients P using a KDS is responsive, efficient, compact, and local.

The size of the KDS can be reduced from $\Theta(mn)$ to $\Theta(n \log m)$ by using a natural generalization of a kinetic tournament [BGH99, Gui98]. The size bound exploits the fact that the m-hull of a set P can be efficiently represented in $\Theta(\min\{m, |P|\})$ space. The total number of change events remains $O(mn)$ [DK06].

8.2.4 Maintaining Mobile 2-Centre Functions

In Ch. 6 we identified two bounded-velocity approximations of the mobile Euclidean 2-centre: the Steiner reflection 2-centre and the rectilinear reflection 2-centre.

Each of these is defined by reflecting the position of some client p_0 in P across the corresponding centre function. It follows that the kinetic maintenance of any reflection-based 2-centre function is achieved by maintaining the corresponding mobile centre function as described earlier in this section.

8.3 Maintaining Mobile Median Functions

In Ch. 5 we identified five bounded-velocity approximations of the mobile Euclidean 1-median: the projection median, the Gaussian median, the rectilinear 1-median, the centre of mass, and a linear combination of the latter two.

The rectilinear 1-median of a set of mobile clients is maintained by the respective one-dimensional mobile median in each dimension. Agarwal et al. [AGG02] describe an efficient implementation using a KDS.

As described in Sec. 8.2, the centre of mass of a set of mobile clients is straightforward to maintain as the average of the positions of all mobile clients.

Maintaining any convex combination of the mobile centre of mass and the mobile rectilinear 1-median is achieved by maintaining the motions of both centre functions as described above and returning the convex combination of their trajectories.

We now describe algorithms for maintaining the mobile Gaussian median and the mobile projection median.

8.3.1 Maintaining the Mobile Gaussian Median

The definition of the Gaussian median of P is a normalized weighted average of the positions of clients in P. The weight of client p at time t is $\pi - w_2(p(t))$, where $w_2(p(t))$ corresponds to Eq. (8.1) (see Def. 5.4). Consequently, our observations about maintaining the exact position of the mobile Steiner centre with a KDS also apply to maintenance of the mobile Gaussian median. See Sec. 8.2.1. Thus, the position of the Gaussian median is not expressible as a polynomial and its description requires a number of terms proportional to the size of the convex hull, $\Theta(|P|)$. Given this constraint on the complexity of a description of the Gaussian median's trajectory, the position of G_2 may be maintained by any KDS that maintains the convex hull of a set of mobile clients.

As shown in Obs. 5.28, the Gaussian median of a set of mobile clients P can be described as a function of the Steiner centre of P and the centre of mass of P. Just as we employed the m-hull to approximate the mobile Steiner centre, the correspondence between the Gaussian median and the Steiner centre suggests we examine the quality of the approximation of the Gaussian median of the m-hull of a set of mobile clients.

Corollary 8.2. *Let $P \in \mathscr{P}(\mathbb{R}^2)$ and let $m \in \mathbb{Z}^+$ be fixed, $m \geq 4$. Let $Q_m(P)$ denote the m-hull of P and let r be the Euclidean radius of P. The distance between $G_2(P)$ and $G_2(Q_m(P))$ satisfies*

$$||G_2(Q_m(P)) - G_2(P)|| \leq \frac{(12\pi + 32)r}{\pi m}. \tag{8.4}$$

Proof.

$$\begin{aligned}&||G_2(P) - G_2(Q_m(P))|| \\ &\qquad = \frac{1}{|P|-2}\left|\left||P|[C_2(P) - C_2(Q_m(P))] - 2[\Gamma_2(P) - \Gamma_2(Q_m(P))]\right|\right|,\end{aligned}$$

by Obs. 5.28,

$$\begin{aligned}&\leq \frac{1}{|P|-2}\big(|P| \cdot ||C_2(P) - C_2(Q_m(P))|| + 2||\Gamma_2(P) - \Gamma_2(Q_m(P))||\big) \\ &\leq \frac{1}{|P|-2}\left(|P|\frac{4r}{m} + 2\frac{16r}{\pi m}\right),\end{aligned}$$

by Thm. 8.1,

$$\leq \frac{(12\pi + 32)r}{\pi m}.$$ ☐

Again, our observations about maintaining the mobile Steiner centre of the m-hull with a KDS also apply to maintenance of the mobile Gaussian median of the m-hull. See Sec. 8.2.3.

8.3.2 Algorithm for the Static Projection Median

The projection median can be found using techniques similar to those developed by Bereg et al. [BKS00]. In brief, as θ varies from 0 to π, the client(s) in P that

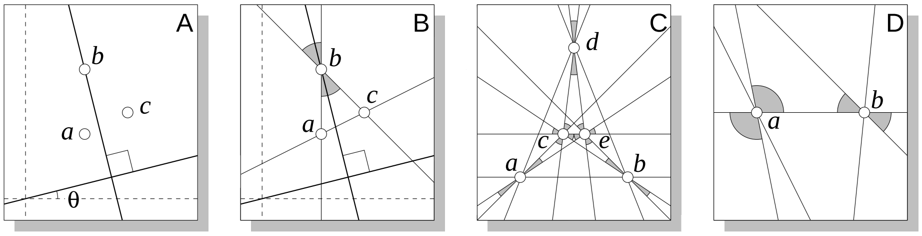

Figure 8.4: For every client p in P, the static projection median algorithm identifies the range of angles for which p induces induces a median of P_θ.

induce $\text{med}(P_\theta)$ are identified by maintaining a line (perpendicular to l_θ) that partitions P into two sets of at most $\lfloor n/2 \rfloor$ clients each (see Def. 5.1).

The range of integration $\theta \in [0, \pi)$ can be partitioned into subintervals $[\alpha_i, \alpha_{i+1})$ such that for every $\theta \in [\alpha_i, \alpha_{i+1})$, the median of P_θ is given by the projection of client $p_{a_i} \in P$ onto line l_θ. That is,

$$\frac{1}{2\pi} \int_0^\pi \text{med}(P_\theta)\ d\theta = \frac{1}{2\pi} \sum_{i=0}^{m-1} \int_{\alpha_i}^{\alpha_{i+1}} u_\theta \langle p_{a_i}, u_\theta \rangle\ d\theta, \tag{8.5}$$

for some $0 = \alpha_0 < \ldots < \alpha_m = \pi$ and $p_{a_0}, \ldots, p_{a_{m-1}} \in P$. Eq. (8.5) has a closed form integration consisting of $\Theta(m)$ terms. Each endpoint of interval $[\alpha_i, \alpha_{i+1})$ at client a coincides with an interval $[\alpha_{i+1}, \alpha_{i+2})$ at an adjacent client, say b, along edge $\overline{ab}$. See Fig. 8.4D.

These subintervals are identified by maintaining the convex hull of each partition as the line rotates, requiring $O(\log^2 n)$ time per update [OL81]. See Fig. 8.5. Since the dual problem to maintaining these partitions corresponds to an $n/2$-level, we require at most $O(n^{4/3})$ such updates [Dey98]. That is $m = O(n^{4/3})$. Between updates, the contribution to $\Pi_2(P)$ of the client(s) that induce $\text{med}(P_\theta)$ is calculated in $O(1)$ time. Together, these give an $O(n^{4/3} \log^2 n)$-time algorithm. This can be improved to $O(n^{4/3} \log^{1+\epsilon} n)$ amortized time using the dynamic convex hull data structure of Chan [Cha01].

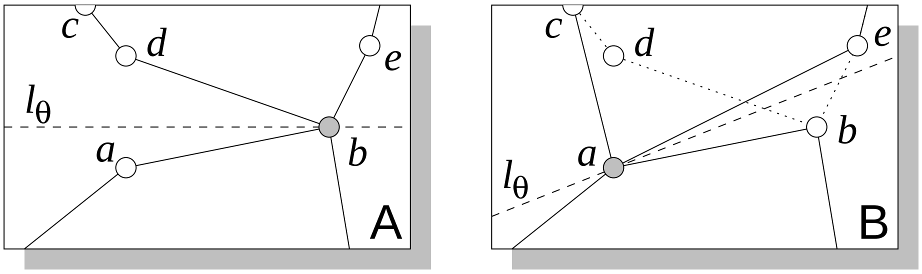

Figure 8.5: maintaining the convex hulls of both partitions as l_θ rotates

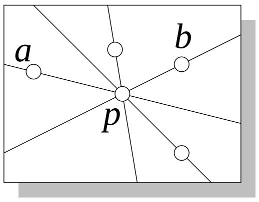

Figure 8.6: Edge $\overline{pa}$ crosses edge $\overline{pb}$ when a and b lie in the same half-plane relative to p and $\arctan[(a_y - p_y)/(a_x - p_x)] = \arctan[(b_y - p_y)/(b_x - p_x)]$.

8.3.3 Maintaining the Mobile Projection Median

The static projection median algorithm generalizes to a mobile algorithm using a KDS. We first employ the static algorithm on the initial set of client positions such that for every client p in P, we identify the subintervals of angles for which p induces a median of P_θ. These subintervals are maintained as the clients move. Let a, b, and p denote mobile clients in P. Edges $\overline{ap}$ and $\overline{bp}$ cross if

$$\arctan\left(\frac{a_y(t) - p_y(t)}{a_x(t) - p_x(t)}\right) = \arctan\left(\frac{b_y(t) - p_y(t)}{b_x(t) - p_x(t)}\right)$$
$$\Rightarrow \quad [a_y(t) - p_y(t)][b_x(t) - p_x(t)] = [a_x(t) - p_x(t)][b_y(t) - p_y(t)], \tag{8.6}$$

as long as a and b lie in the same half-plane relative to p. See Fig. 8.6. Eq. (8.6) is a piecewise-quadratic polynomial given piecewise-linear client motion. Furthermore, between client flight updates, every two edges cross at most once. There are $\Theta(n^2)$ edges and n clients, resulting in $O(n^3)$ edge crossings. For each edge crossing, the corresponding intervals must be updated at three clients involved. Four cases are possible as displayed in Fig. 8.7.

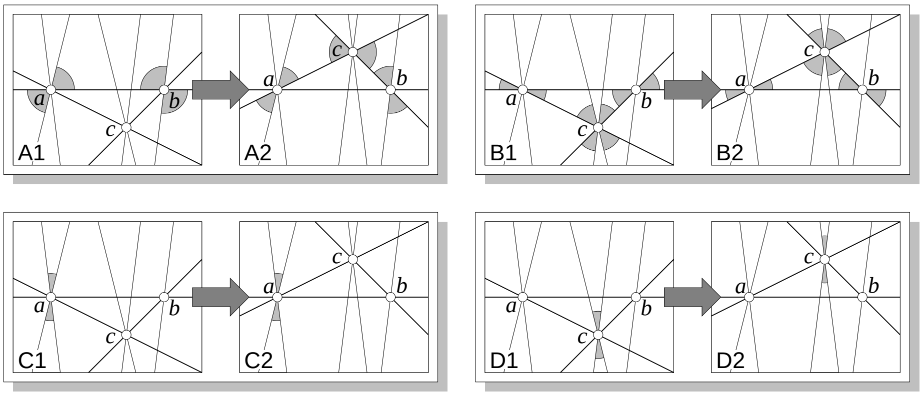

Figure 8.7: Four cases are possible when c crosses the edge between a and b.

The resulting KDS has similar properties to those of a KDS that maintains the exact Steiner centre. That is, the position of the projection median is not expressible as a polynomial. Furthermore, its description requires $O(n^{4/3})$ terms, corresponding to the number of subintervals of $[0, \pi)$ described in the static algorithm. Given this constraint on the complexity of a description of the projection median's trajectory, the position of Π_2 may be maintained by any KDS that maintains the set of subintervals for each client. For any mobile client p, the description of its subintervals change only when some client crosses and edge adjacent to p. Each such update requires only constant time. For comparison, in general the position of the Euclidean 1-median cannot be expressed exactly (see Sec. 2.4.1).

8.3.4 Discretized Approximation of the Projection Median

Just as the m-hull enabled for a simpler KDS to maintain a close approximation of the Steiner centre, we propose maintaining a discretization of the projection median. In Lem. 5.14 we showed that the projection median can be expressed in terms of the rectilinear 1-median:

$$\Pi_2(P) = \frac{2}{\pi}\int_0^{\pi/2} S_\theta(P)\, d\theta = \lim_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1} S_{i\pi/2n}(P). \tag{8.7}$$

Consequently, we propose approximating the projection median by

$$T_m(P) = \frac{1}{m}\sum_{i=0}^{m-1} S_{i\pi/2m}(P), \tag{8.8}$$

for a fixed $m \geq 1$. Every $T_m(P)$ is maintained kinetically by an instance of the rectilinear 1-median KDS of Agarwal et al. [AGG02] mentioned earlier.

By Thms. 5.6 and 5.26, T_m provides a $\sqrt{2}$-approximation of the Euclidean 1-median for any $m \geq 1$. We provide a better bound on the approximation factor of T_m using techniques similar to those used in the proof of Thm. 5.20. This bound is tight as $m \to \infty$.

Theorem 8.3. *The two-dimensional discretized projection median provides a λ-approximation of the Euclidean 1-median, where*

$$\lambda = \frac{1}{m}\csc\left(\frac{\pi}{4m}\right). \tag{8.9}$$

Proof. Choose any fixed $m \geq 1$. We bound the median sum of T_m.

$$\frac{\sum_{p\in P}||T_m(P)-p||}{\sum_{q\in P}||M_2(P)-q||}$$
$$= \frac{\sum_{p\in P}\left|\left|\frac{1}{m}\sum_{i=0}^{m-1} S_{i\pi/2m}(P)-p\right|\right|}{\sum_{q\in P}||M_2(P)-q||}$$
$$= \frac{\sum_{p\in P}\left|\left|\frac{1}{m}\sum_{i=0}^{m-1} S_{i\pi/2m}(P)-\frac{1}{m}\sum_{i=0}^{m-1} p\right|\right|}{\sum_{q\in P}||M_2(P)-q||}$$
$$\leq \frac{\sum_{p\in P}\frac{1}{m}\sum_{i=0}^{m-1}||S_{i\pi/2m}(P)-p||}{\sum_{q\in P}||M_2(P)-q||},$$

by the $\triangle$ inequality,

$$\leq \frac{\sum_{p\in P}\frac{1}{m}\sum_{i=0}^{m-1} d_{i\pi/2m}(S_{i\pi/2m}(P),p)}{\sum_{q\in P}||M_2(P)-q||},$$

since $\forall x\ ||x||_1 \geq ||x||$ and, similarly, $\forall x \forall \phi\ d_\phi(x) \geq ||x||$,

$$\leq \frac{\sum_{p\in P}\frac{1}{m}\sum_{i=0}^{m-1} d_{i\pi/2m}(M_2(P),p)}{\sum_{q\in P}||M_2(P)-q||}, \tag{8.10a}$$

since $S_\phi(P)$ minimizes the sum of the d_ϕ distances to points of P,

$$= \frac{\sum_{p\in P}\frac{1}{m}\sum_{i=0}^{m-1}\left[\cos\left(\frac{i\pi}{2m}-\alpha_p\right)+\sin\left(\frac{i\pi}{2m}-\alpha_p\right)\right]\cdot||M_2(P)-p||}{\sum_{q\in P}||M_2(P)-q||}, \tag{8.10b}$$

where $\alpha_p = \arctan[(M_2(P)_y - p_y)/(M_2(P)_x - p_x)] \bmod \frac{\pi}{2}$ (see Fig. 5.8),

$$= \frac{\sum_{p\in P}\left[\cos\left(\frac{\pi}{2m}\right)\cos\alpha_p - \sin\left(\frac{\pi}{2m}\right)\sin\alpha_p + \cos\alpha_p\right]\cdot||M_2(P)-p||}{m\sin\left(\frac{\pi}{2m}\right)\sum_{q\in P}||M_2(P)-q||} \tag{8.10c}$$
$$\leq \frac{\sum_{p\in P}\left[\cos\left(\frac{\pi}{2m}\right)\cos\left(\frac{\pi}{4m}\right) + \sin\left(\frac{\pi}{2m}\right)\sin\left(\frac{\pi}{4m}\right) + \cos\left(\frac{\pi}{4m}\right)\right]\cdot||M_2(P)-p||}{m\sin\left(\frac{\pi}{2m}\right)\sum_{q\in P}||M_2(P)-q||}, \tag{8.10d}$$

by Eq. (8.11),

$$= \frac{\csc\left(\frac{\pi}{4m}\right)\sum_{p\in P}||M_2(P)-p||}{m\sum_{q\in P}||M_2(P)-q||}$$
$$= \frac{1}{m}\csc\left(\frac{\pi}{4m}\right). \tag{8.10e}$$

Eq. (8.10d) follows from Eq. (8.10c) by taking the derivative and setting it to zero:

$$\begin{aligned}
&\frac{\partial}{\partial \alpha_p}\left[\cos\left(\frac{\pi}{2m}\right)\cos\alpha_p - \sin\left(\frac{\pi}{2m}\right)\sin\alpha_p + \cos\alpha_p\right] = 0,\\
\Rightarrow\quad & -\cos\left(\frac{\pi}{2m}\right)\sin\alpha_p - \sin\left(\frac{\pi}{2m}\right)\cos\alpha_p - \sin\alpha_p = 0,\\
\Rightarrow\quad & \alpha_p = -\arctan\left(\frac{\sin\left(\frac{\pi}{2m}\right)}{\cos\left(\frac{\pi}{2m}\right)+1}\right)\\
& \quad = -\frac{\pi}{4m}. \qquad (8.11)
\end{aligned}$$

Therefore, for any finite multiset of points P in $\mathbb{R}^2$ and for any fixed $m \geq 1$,

$$\sum_{p\in P} ||T_m(P) - p|| \leq \frac{1}{m}\csc\left(\frac{\pi}{4m}\right)\sum_{q\in P}||M_2(P) - q||. \qquad \square$$

Observe that

$$\lim_{m\to\infty}\frac{1}{m}\csc\left(\frac{\pi}{4m}\right) = \frac{4}{\pi}. \qquad (8.12)$$

Eq. (8.12) converges rapidly to $4/\pi$, as shown by the Laurent series expansion of cosecant which gives

$$\frac{1}{m}\csc\left(\frac{\pi}{4m}\right) = \frac{4}{\pi} + \frac{\pi}{24m^2} + O\left(\frac{1}{m^4}\right).$$

8.4 Java Visualization

The worst-case approximation factor of an approximation function is realized by a static set of client positions. As such, intuition about the approximation factor of the two-dimensional Euclidean 1-centre, 2-centre, and 1-median is effectively communicated though the use of figures that exemplify specific characteristics of various sets of client positions. Intuition about velocity, however, is more difficult to convey in a figure. Even more difficult to represent graphically is the correlation between velocity and approximation factor.

To help convey various notions described in this thesis, a visual demonstration was implemented to display a set of mobile clients P in $\mathbb{R}^2$ and the significant approximation functions of P described in Chs. 4 through 6. We briefly describe the implementation in this section.

In addition, the implementation provides a source of data on the approximation factor and average velocity of each approximation function at every time step. Data collected from extended runs is included in this section to reinforce conclusions drawn in Chs. 4 and 5 and suggest that the worst-case bounds derived are are not overly pessimistic or unrealistic. It should be emphasized, however, that the reader is not expected to draw precise numerical conclusions on average-case bounds on approximation factors or velocities; rather these statistics are included only to provide informal evidence to reinforce the corresponding formal results established in Chs. 4, 5, and 6.

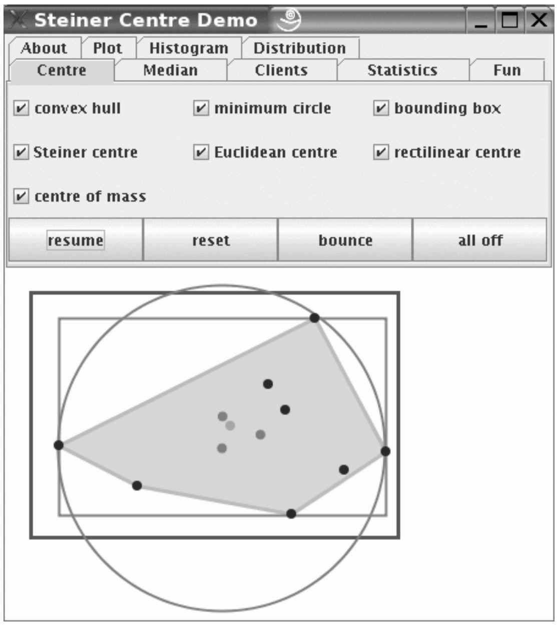

Figure 8.8: a screen capture of the applet (the different colours are not easily distinguishable in this grey-scale image)

8.4.1 Java Applet Overview

The implementation as described was written in Java. The interface consists of a tabbed control panel above a display window. See Fig. 8.8. The window displays a set of mobile clients moving inside a red rectangle. Each client follows a linear trajectory until it bounces upon coming into contact with a rectangle edge. The bounce angle is randomized so that client trajectories do not remain constant.

A set of check boxes and a tabbed menu allows the user to select which approximation functions of the client set are to be displayed. These include the Euclidean 1-centre, the rectilinear 1-centre, the centre of mass, the Steiner centre, the Euclidean 1-median, the rectilinear 1-median, the projection median, the Gaussian median, the Euclidean 2-centre, the Steiner reflection 2-centre, and the rectilinear reflection 2-centre. Exact positions (to within standard floating point error) are calculated for all facilities except the Euclidean 1-median, for which Weiszfeld's approximation algorithm is implemented, and the projection median, for which the discretization described in Eq. (8.8) is calculated. Additional features of P that can be displayed include the minimum enclosing circle(s), the bounding box, the convex hull, client projections onto the x- and y-axes, and the Voronoi diagram of P corresponding to a partition of clients in P based on proximity to either facility of the Euclidean 2-centre. Buttons and slide bars for controlling the velocity, position, and num-

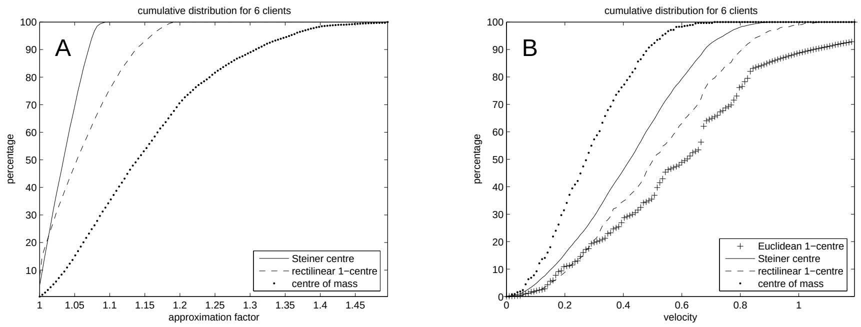

Figure 8.9: centre function cumulative plots for approximation factor and average velocity for six mobile clients

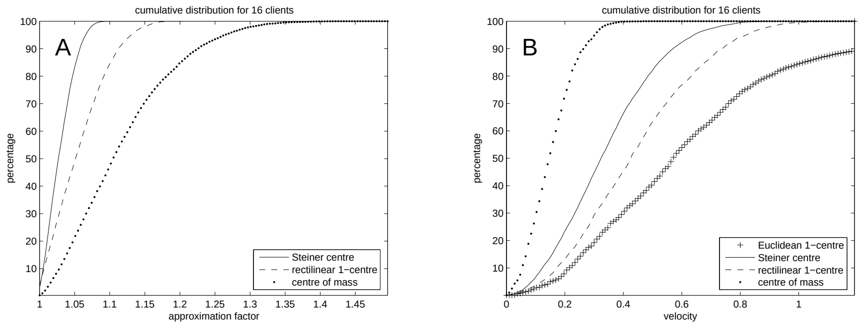

Figure 8.10: centre function cumulative plots for approximation factor and average velocity for sixteen mobile clients

ber of clients are also included. The applet is available for viewing online at `http://www.cs.ubc.ca/~durocher/gaussianDemo.html`.

In addition to displaying the various mobile objects described above, the applet includes three visualizations that display the approximation factor and velocity of the approximation functions. These are intended to help a viewer get a sense of the relative approximation factors and velocities of the approximation functions for various configurations of client positions. These include scrolling overlayed plots of the approximation factors (respectively, velocity) of various centre (respectively, median, 2-centre) functions over time, a histogram of the instantaneous approximation factors (velocity) of various centre (median, 2-centre) functions, and a cumulative distribution plot displaying the distribution of the approximation factors (velocity) of various centre (median, 2-centre) functions over the entire duration of a simulation.

centre function		50th percentile	75th percentile	90th percentile
Euclidean	6	1.0000 | 0.6188	1.0000 | 0.7969	1.0000 | 1.0594
1-centre	16	1.0000 | 0.5719	1.0000 | 0.8156	1.0000 | 1.1906
centre	6	1.1406 | 0.2812	1.2188 | 0.3937	1.3068 | 0.4875
of mass	16	1.1055 | 0.1500	1.1641 | 0.2156	1.2266 | 0.2719
rectilinear	6	1.0547 | 0.5062	1.1016 | 0.6750	1.1406 | 0.8063
1-centre	16	1.0547 | 0.4312	1.0859 | 0.5906	1.1133 | 0.7406
Steiner	6	1.0353 | 0.4312	1.0547 | 0.5719	1.0703 | 0.6844
centre	16	1.0273 | 0.3281	1.0430 | 0.4594	1.0586 | 0.5719

Table 8.1: summary of centre function statistics displayed as percentiles: $\lambda | v_{\max}$

8.4.2 Empirical Evidence

In addition to visual display, the implementation was used to collect empirical data on the performance of the approximation functions implemented. Each run consisted of 10,000 time steps. For each instance, a set of mobile client positions was randomly generated over a uniform distribution of x- and y-coordinates inside a bounded rectangular area. Clients were assigned velocities consisting of a magnitude in the range $[0, v]$ for a fixed v and a direction angle in the range $[0, 2\pi]$, both of which were generated uniformly and at random. The data collected at every time step consisted of instantaneous approximation factor and average velocity between the current time and previous time stamp. This information was recorded for the Euclidean 1-centre, the rectilinear 1-centre, the Steiner centre, the centre of mass (as a centre function), the Euclidean 1-median, the rectilinear 1-median, the projection median, the Gaussian median, the centre of mass (as a median function), the Euclidean 2-centre, the rectilinear reflection 2-centre, and the Steiner reflection 2-centre. Since these statistics depend on the cardinality of the client set, the simulation was run for 6 clients and again for 16 clients, resulting in twelve sets of data: 6 vs. 16 clients, centre function vs. median function vs. 2-centre function, and velocity vs. approximation factor.

The results of these tests are displayed in Figs. 8.9 through 8.14. The plots display the cumulative distribution as a percentage. Recall that the approximation factor is always at least one. Consequently, the domain displayed for approximation factor is $[1, 1.5]$ for centre functions, $[1, 1.1]$ for median functions, and $[1, 3]$ for 2-centre functions, whereas the domain displayed for velocity is $[0, 1.2]$ for centre and median functions and $[0, 4]$ for 2-centre functions. The test data is also summarized in Tabs. 8.1, 8.2. and 8.3.

In each of the plots, the topmost line best satisfies the property being plotted. For instance, in Fig. 8.9A, the topmost line corresponds to the Steiner centre, suggesting that on average, the approximation factor of the Steiner centre was less than that of the rectilinear 1-centre, which was in turn lower than that of the centre of mass. Similarly, in Fig. 8.9B, the topmost line corresponds to the centre of mass, suggesting that on average, the velocity of the centre of mass was less than that of the Steiner centre, followed by the rectilinear 1-centre, and

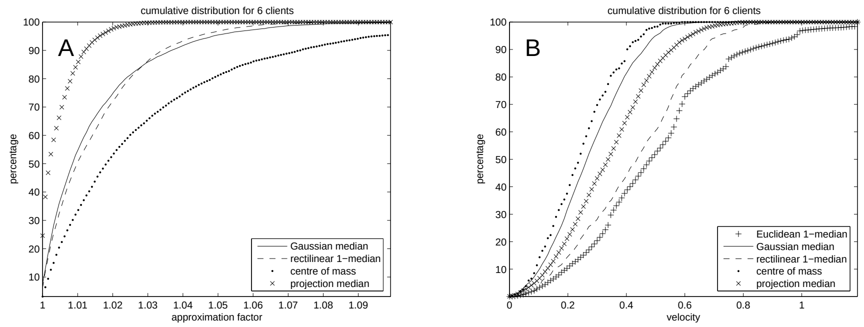

Figure 8.11: median function cumulative plots for approximation factor and average velocity for six mobile clients

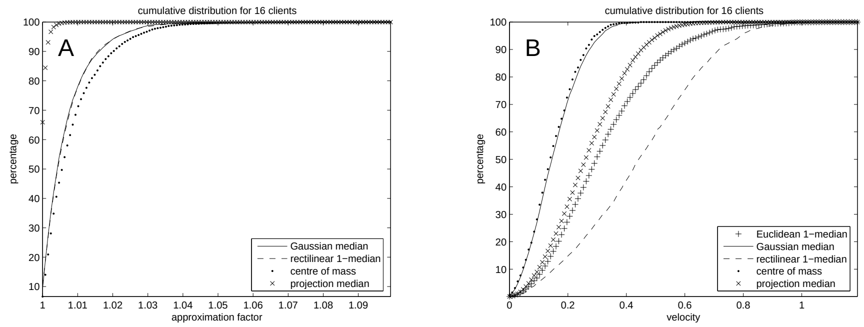

Figure 8.12: median function cumulative plots for approximation factor and average velocity for sixteen mobile clients

finally the Euclidean 1-centre. The results of Fig. 8.10 are similar.

Similarly, we present data collected for mobile median functions. In Fig. 8.11A, the topmost line corresponds to the projection median, suggesting that on average, the approximation factor of the projection median was less than that of the rectilinear 1-median and the Gaussian median, followed by the centre of mass. In Fig. 8.11B, the topmost line corresponds to the centre of mass, suggesting that on average, the velocity of the centre of mass was less than that of the Gaussian median, followed by the projection median, followed by the rectilinear 1-median, followed by the Euclidean 1-median. The results of Fig. 8.12 are similar.

Finally, we present data collected for mobile 2-centre functions. In Fig. 8.13A, the topmost line corresponds to the Steiner reflection 2-centre, suggesting that on average, the approximation factor of the Steiner reflection 2-centre was slightly less than that of the rectilinear reflection 2-centre. Observe that the plot lines in Fig. 8.13A appear smooth with the exception of a spike occurring at approximation factor $\lambda = 2$. This cluster is explained by the fact that

median function		50th percentile	75th percentile	90th percentile
Euclidean	6	1.0000 | 0.4875	1.0000 | 0.6281	1.0000 | 0.8344
1-median	16	1.0000 | 0.3000	1.0000 | 0.4312	1.0000 | 0.5719
centre	6	1.0188 | 0.2437	1.0406 | 0.3375	1.0742 | 0.4031
of mass	16	1.0063 | 0.1406	1.0117 | 0.2062	1.0203 | 0.2625
rectilinear	6	1.0102 | 0.4406	1.0219 | 0.5719	1.0344 | 0.6750
1-median	16	1.0047 | 0.4500	1.0094 | 0.6000	1.0164 | 0.7219
Gaussian	6	1.0086 | 0.2719	1.0203 | 0.3750	1.0375 | 0.4688
median	16	1.0047 | 0.1500	1.0094 | 0.2156	1.0164 | 0.2812
projection	6	1.0023 | 0.3469	1.0070 | 0.4500	1.0125 | 0.5625
median	16	1.0000 | 0.2625	1.0008 | 0.3656	1.0016 | 0.4594

Table 8.2: summary of median function statistics displayed as percentiles: $\lambda|v_{\max}$

whenever client p_0 lies on the minimum enclosing circle of the partition of P with larger Euclidean radius, the corresponding approximation factor is exactly 2. No such spike is evident in Fig. 8.14A; this property is expected since the probability that p_0 lies on the minimum enclosing circle of the partition of P with larger radius decreases inversely with $|P|$. In Fig. 8.13B, both the rectilinear reflection 2-centre and the Steiner reflection 2-centre show a sharp spike near velocity $v_{\max} = 1$. Whenever the reflected centre, q, moves slower than p_0, the maximum velocity of a reflection-based 2-centre function is realized by p_0. Again, this property is less pronounced when $|P| = 16$ in Fig. 8.14B. Note that since velocity is measured as average over a time interval, discontinuities in the position of the Euclidean 2-centre are not recorded as arbitrarily large velocities.

In Ch. 4 we showed that in $\mathbb{R}^2$ the worst-case approximation factors of centre functions in order from lowest (best) to highest (worst) correspond to the Steiner centre, the rectilinear 1-centre, and the centre of mass. Similarly, we showed that in $\mathbb{R}^2$ the worst-case maximum velocities of centre functions in order from slowest to fastest correspond to the centre of mass, the Steiner centre, the rectilinear 1-centre, and the Euclidean 1-centre. Although average-case values tend to be significantly lower that the corresponding worst-case bounds, Figs. 8.9 and 8.10 show that the relative ordering remains unchanged in practice.

In Ch. 5 we showed that in $\mathbb{R}^2$ the worst-case approximation factors of median functions in order from lowest (best) to highest (worst) correspond to the projection median, the rectilinear 1-median, and the centre of mass. Similarly, we showed that in $\mathbb{R}^2$ the worst-case maximum velocities of median functions in order from slowest to fastest correspond to the centre of mass, the projection median, the rectilinear 1-median, and the Euclidean 1-median. Although average-case values tend to be significantly lower that the corresponding worst-case bounds, Figs. 8.11 and 8.12 show that the relative ordering remains unchanged in practice. Although we did not derive tight bounds on the maximum

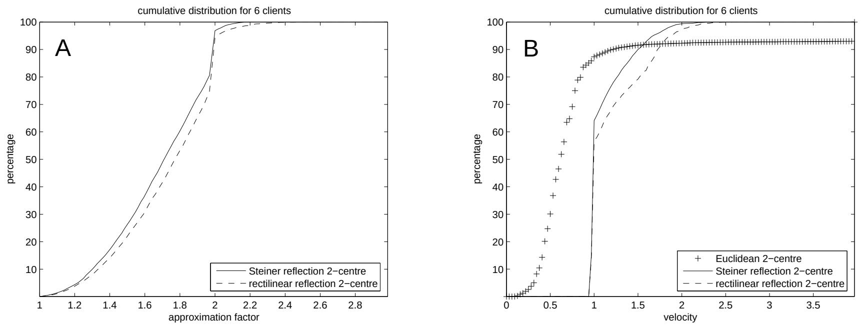

Figure 8.13: 2-centre function cumulative plots for approximation factor and average velocity for six mobile clients

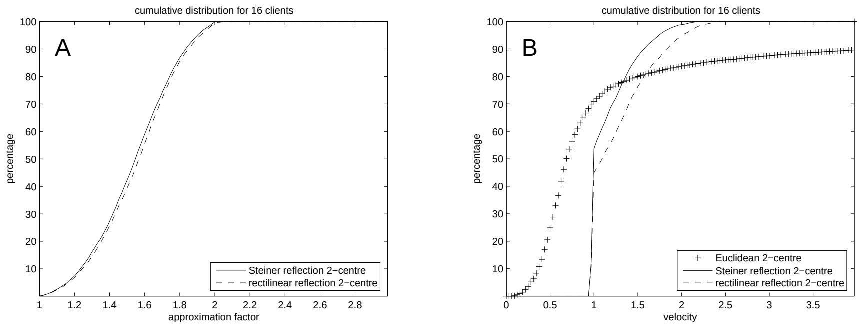

Figure 8.14: 2-centre function cumulative plots for approximation factor and average velocity for sixteen mobile clients

velocity and approximation factor of the Gaussian median, Figs. 8.11 and 8.12 suggest that the upper bounds may indeed be lower that those we have established in Sec. 5.8.

In Ch. 6 we showed that in $\mathbb{R}^2$ the worst-case approximation factors of 2-centre functions in order from lowest (best) to highest (worst) correspond to the Steiner reflection 2-centre and the rectilinear reflection 2-centre. Similarly, we showed that in $\mathbb{R}^2$ the worst-case maximum velocities of 2-centre functions in order from slowest to fastest correspond to the Steiner reflection 2-centre, the rectilinear reflection 2-centre, and the Euclidean 2-centre. Although average-case values tend to be significantly lower that the corresponding worst-case bounds, Figs. 8.13A and 8.14A show that the relative ordering of approximation factors remains unchanged in practice. The analogous conclusion on velocity is not immediate from this analysis, partly due to the fact that this data does not account for discontinuities in the motion of the Euclidean 2-centre.

2-centre function		50th percentile	75th percentile	90th percentile
Euclidean	6	1.0000 | 0.6250	1.0000 | 0.7812	1.0000 | 1.2188
2-centre	16	1.0000 | 0.6875	1.0000 | 1.1562	1.0000 | 3.9688
Steiner	6	1.7188 | 1.0000	1.9375 | 1.1875	2.0000 | 1.5000
reflection 2-centre	16	1.5469 | 1.0000	1.7188 | 1.3125	1.8438 | 1.5938
rectilinear	6	1.7812 | 1.0000	1.9844 | 1.3750	2.0000 | 1.7500
reflection 2-centre	16	1.5781 | 1.0938	1.7344 | 1.5000	1.8594 | 1.8438

Table 8.3: summary of 2-centre function statistics displayed as percentiles: $\lambda | v_{\max}$

Chapter 9

Conclusions and Directions for Future Research

This chapter provides a brief conclusion and lists possible future research directions that result from this work.

9.1 Bounded-Velocity Approximations of the Mobile Euclidean k-Centre and k-Median

This thesis seeks to identify bounded-velocity approximations to the mobile Euclidean k-centre and the mobile Euclidean k-median, two sets of mobile facilities whose positions move with unbounded velocity and/or discontinuous motion, even when $k = 1$.

We identified the rectilinear 1-centre, the centre of mass, and the Steiner centre as bounded-velocity approximations of the mobile Euclidean 1-centre. In particular, we established that the Steiner centre, which had not previously been evaluated as an approximation of the Euclidean 1-centre, successfully balances low maximum velocity and a low approximation factor.

We identified the rectilinear 1-median and the centre of mass and we introduced the projection median and the Gaussian median as bounded-velocity approximations of the mobile Euclidean 1-median. The definition of the projection median provides a new generalization of the one-dimensional median to higher dimensions that successfully balances low maximum velocity and a low approximation factor.

We introduced the rectilinear reflection 2-centre and the Steiner reflection 2-centre as bounded-velocity approximations of the Euclidean 2-centre. These two mobile approximation functions overcome the challenges of discontinuity and implicit partitioning imposed by multiple facilities.

We addressed the Euclidean k-centre for $k \geq 3$ and the Euclidean k-median for $k \geq 2$ and showed that no bounded-velocity approximation can be guaranteed for either of these.

Finally, we presented kinetic algorithms for maintaining these various mobile approximation functions on a set of mobile clients using both exact and approximate solutions.

9.2 Directions for Future Research

In this section we briefly mention some of the open problems that arise as a result of this research. The first few questions involve tightening bounds on approximation factor or generalizing two-dimensional results to three dimensions.

Steiner Centre

In Thm. 4.20 and Cor. 4.21 we showed a tight bound of approximately 1.1153 on the approximation factor (eccentricity) of the Steiner centre's approximation of the Euclidean 1-centre in two dimensions. In Thm. 4.22 we showed a corresponding lower bound of approximately 1.2017 in three dimensions. Lem. 4.4 implies an upper bound of 2 on the approximation factor. The question of finding a tight bound $\lambda \in [1.2017, 2]$ on the approximation factor of the Steiner centre in three dimensions remains open.

Rectilinear 1-Median

In Thm. 5.6 we showed an upper bound of $\sqrt{d}$ on the approximation factor of the rectilinear 1-median's approximation of the Euclidean 1-median in d dimensions. In Thm. 5.7 we showed a corresponding lower bound of $(1+\sqrt{d-1})/\sqrt{d}$. The question of finding a tight bound $\lambda \in [(1+\sqrt{d-1})/\sqrt{d}, \sqrt{d}]$ on the approximation factor of the rectilinear 1-median in d dimensions remains open.

Projection Median

In Thm. 5.20 we showed an upper bound of $4/\pi$ on the approximation factor of the projection median's approximation of the Euclidean 1-median in two dimensions. In Thm. 5.21 we showed a corresponding lower bound of $\sqrt{4/\pi^2+1}$. The question of finding a tight bound $\lambda \in [\sqrt{4/\pi^2+1}, 4/\pi]$ on the approximation factor of the projection median in two dimensions remains open.

It seems probable that the definition of the three-dimensional projection median can be interpreted in terms of the rectilinear 1-median as was done in two dimensions in Lem. 5.14. If true, this equivalence may lead to a generalization of the two-dimensional upper bound on the approximation factor of the projection median. Should Thm. 5.20 generalize, the three-dimensional upper bound corresponding to Eq. (5.27d) simplifies to $3/2$.

Reflection-Based 2-Centre Functions

In Thm. 6.25 we showed an upper bound of $8/\pi$ on the approximation factor (eccentricity) of the Steiner reflection 2-centre's approximation of the Euclidean 2-centre in two dimensions. In Thm. 6.26 we showed a corresponding lower bound of $2\sqrt{1+1/\pi^2}$. The question of finding a tight bound $\lambda \in [2\sqrt{1+1/\pi^2}, 8/\pi]$ on the approximation factor of the Steiner reflection 2-centre in two dimensions remains open. No greater lower bound is known in three dimensions nor is any upper bound currently known. The generalization of Lem. 4.18 to three dimensions would allow for the proof of Thm. 6.25 to be generalized to three dimensions, resulting in an upper bound of 5 on the approximation factor of the Steiner reflection 2-centre in three dimensions.

In Thm. 6.23 we showed an upper bound of $2\sqrt{d}$ on the approximation factor of the rectilinear reflection 2-centre's approximation of the Euclidean 2-centre in d dimensions. In Thm. 6.24 we showed a corresponding lower bound of $2\sqrt{2}$. The question of finding a tight bound $\lambda \in [2\sqrt{2}, 2\sqrt{d}]$ on the approximation factor of the rectilinear reflection 2-centre in d dimensions remains open.

Approximating the Euclidean 2-Centre without Reflection

The only bounded-velocity approximations of the Euclidean 2-centre currently known are the Steiner reflection 2-centre and the rectilinear reflection 2-centre. Straightforward variations of these (e.g., linear combinations) are likely to produce additional related bounded-velocity approximations. The problem of defining a bounded-velocity approximation whose position is independent of reflection of a client position remains open.

Additional Facilities

The question of whether there exists a set of three mobile facilities that provides a bounded-velocity approximation of the geometric 2-median in one dimension remains open. The analogous questions for the geometric 3-centre and for k mobile facilities and the $(k+1)$-median were addressed in Sec. 7.3.

Implementation and Average-Case Analysis

As mentioned in Sec. 8.4.2, the data collected by the Java applet are not intended as a formal statistical analysis. The following changes are required if one were interested in collecting more robust statistics:

- the positions and velocities of clients should be randomized after each time step,
- each test should be run for a greater number of time steps,
- results from several runs should be averaged,
- greater care needs to be taken with respect to numerical error in small values of λ, and
- potential error with respect to approximation of the positions of the Euclidean 1-median and the projection median must be addressed.

For a more formal average-case analysis, the precise definition of the model needs to be addressed. Specifically, what defines average motion or randomized motion of a mobile client? Does this motion occur within a bounded area (if so, what is its shape)? How many mobile clients should be included? Large sets of mobile clients tend to induce slow-moving centre functions clustered near the middle whereas smaller sets allow for larger variation and rapid changes in the positions of centre functions. What random distribution best describes an "average" distribution of velocities of mobile clients (simpler if all clients have unit velocity)?

Bibliography

[AAR97] Helmut Alt, Oswin Aichholzer, and Günter Rote. Matching shapes with a reference point. *International Journal of Computational Geometry and Applications*, 7:349–363, 1997.

[AdBG$^+$05] Pankaj K. Agarwal, Mark de Berg, Jie Gao, Leonidas J. Guibas, and Sariel Har-Peled. Staying in the middle: Exact and approximate medians in $\mathbb{R}^1$ and $\mathbb{R}^2$ for moving points. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 17, pages 42–45, 2005.

[AF90] Jean-Pierre Aubin and Hélène Frankowska. *Set-valued Analysis*, volume 2 of *Systems & Control: Foundations & Applications*. Birkhäuser, Boston, 1990.

[AGG02] Pankaj K. Agarwal, Jie Gao, and Leonidas J. Guibas. Kinetic medians and *kd*-trees. In *Proceedings of the Tenth European Symposium on Algorithms*, volume 2461 of *Lecture Notes in Computer Science*, pages 5–16, 2002.

[AGHV01] Pankaj K. Agarwal, Leonidas J. Guibas, John Hershberger, and Eric Veach. Maintaining the extent of a moving point set. *Discrete and Computational Geometry*, 26:353–374, 2001.

[AH01] Pankaj K. Agarwal and Sariel Har-Peled. Maintaining approximate extent measures of moving points. In *Proceedings of the Symposium on Discrete Algorithms*, pages 148–157. ACM Press, 2001.

[AHT00] Stephen Alstrup, Jacob Holm, and Mikkel Thorup. Maintaining center and median in dynamic trees. In *Proceedings of the Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 46–56. Springer-Verlag, 2000.

[AHV05] Pankaj K. Agarwal, Sariel Har-Peled, and Kasturi R. Varadarajan. Geometric approximation via coresets. In Emo Welzl, editor, *Current Trends in Combinatorial and Computational Geometry*. Cambridge University Press, New York, 2005. To appear.

[AK99] Eyal Assa and Matthew J. Katz. 3-piercing of *d*-dimensional boxes and homothetic triangles. *International Journal of Computational Geometry and Applications*, 8:249–260, 1999.

[AM04] Pankaj K. Agarwal and Nabil H. Mustafa. k-means projective clustering. In *Proceedings of the Symposium on Principles of Database Systems*, volume 23, pages 155–165, 2004.

[AOY99] Hideki Ando, Yoshinobu Oasa, and Ichiro Suzuki Masafumi Yamashita. Distributed memoryless point convergence algorithm for mobile robots with limited visibility. *IEEE Transactions on Robotics and Automation*, 15(5):818–828, 1999.

[AP98] Pankaj K. Agarwal and Cecilia M. Procopiuc. Exact and approximation algorithms for clustering. In *Proceedings of the Symposium on Discrete Algorithms*. ACM Press, 1998.

[APP96] Vincenzo Auletta, Domenico Parente, and Giuseppe Persiano. Dynamic and static algorithms for optimal placement of resources in a tree. *Theoretical Computer Science*, pages 441–461, 1996.

[ARR98] Sanjeev Arora, Prabhakar Raghavan, and Satish Rao. Approximation schemes for Euclidean k-medians and related problems. In *Proceedings of the Symposium on Theory of Computing*, pages 106–113. ACM, 1998.

[AS91] Pankaj K. Agarwal and Micha Sharir. Planar geometric location problems and maintaining the width of a planar set. In *Proceedings of the Symposium on Discrete Algorithms*, pages 449–458. ACM Press, 1991.

[AS94] Pankaj K. Agarwal and Micha Sharir. Planar geometric location problems. *Algorithmica*, 11:185–195, 1994.

[AS98] Pankaj K. Agarwal and Micha Sharir. Efficient algorithms for geometric optimization. *ACM Computing Surveys*, 30:412–458, 1998.

[AST93] Pankaj K. Agarwal, Micha Sharir, and Sivan Toledo. An efficient multi-dimensional searching technique and its applications. Technical Report CS-1993-20, Department of Computer Science, Duke University, 1993.

[ASW98] Pankaj K. Agarwal, Micha Sharir, and Emo Welzl. The discrete 2-center problem. *Discrete and Computational Geometry*, 20:287–305, 1998.

[Ata85] Mikhail J. Atallah. Some dynamic computational geometry problems. *Computers and Mathematics with Applications*, 11(12):1171–1181, 1985.

[Baj84] Chanderjit Bajaj. *Geometric Optimization and Computational Complexity*. PhD thesis, Cornell University, 1984.

[Baj88] Chanderjit Bajaj. The algebraic degree of geometric optimization problems. *Discrete and Computational Geometry*, 3:177–191, 1988.

[Bas99a] Stefano Basagni. Distributed clustering for ad hoc networks. In *Proceedings International Symposium on Parallel Architectures, Algorithms, and Networks (I-SPAN)*, pages 310–315, 1999.

[Bas99b] Julien Basch. *Kinetic Data Structures*. PhD thesis, Stanford University, 1999.

[BBK$^+$02] Sergei Bespamyatnikh, Binay Bhattacharya, Mark Keil, David Kirkpatrick, and Michael Segal. Efficient algorithms for centers and medians in interval and circular-arc graphs. *Networks*, 39(3):144–152, 2002.

[BBKS00] Sergei Bespamyatnikh, Binay Bhattacharya, David Kirkpatrick, and Michael Segal. Mobile facility location. In *Proceedings of the International ACM Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, volume 4, pages 46–53, 2000.

[BBKS02] Sergei Bespamyatnikh, Binay Bhattacharya, David Kirkpatrick, and Michael Segal. Lower and upper bounds for tracking mobile servers. In R. Baeza-Yates, U. Montanari, and N. Santoro, editors, *Foundations of Information Technology in the Era of Network and Mobile Computing*, volume 2, pages 47–58. Kluwer Academic Publisher, 2002. Proceedings of the IFIP International Conference on Theoretical Computer Science.

[BBKS06] Sergey Bereg, Binay Bhattacharya, David Kirkpatrick, and Michael Segal. Competitive algorithms for mobile centers. *Mobile Networks and Applications*, 11(2):177–186, 2006.

[BE97] Marshall Bern and David Eppstein. Approximation algorithms for geometric problems. In Dorit Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*, pages 296–345. PWS Publishing Co., 1997.

[BGH99] Julien Basch, Leonidas Guibas, and John Hershberger. Data structures for mobile data. *Journal of Algorithms*, 31(1):1–28, 1999.

[BGKS98] Randeep Singh Bhatia, Sudipto Guha, Samir Khuller, and Yoram Sussmann. Facility location with dynamic distance functions. *Journal of Combinatorial Optimization*, 2:199–217, 1998.

[BGSZ97] Julien Basch, Leonidas J. Guibas, Craig Silverstein, and Li Zhang. A practical evaluation of kinetic data structures. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 388–390, 1997.

[BH67] G. Ball and D. Hall. A clustering technique for summarizing multivariate data. *Behavioral Sciences*, 12(2):153–155, 1967.

[BHMT00] Jack Brimberg, Pierre Hansen, Nenad Mladenović, and Éric D. Taillard. Improvements and comparison of heuristics for solving the multisource Weber problem. *Operations Research*, 48(1):129–135, 2000. Technical report IDSIA-33-97.

[BJSL95] Oded Berman, Patrick Jaillet, and David Simchi-Levi. Location-routing problems with uncertainty. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 427–452. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[BK99] Sergei Bespamyatnikh and David Kirkpatrick. Rectilinear 2-center problems. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 11, pages 68–71, 1999.

[BKS00] Sergei Bespamyatnikh, David Kirkpatrick, and Jack Snoeyink. Generalizing ham sandwich cuts to equitable subdivisions. *Discrete and Computational Geometry*, 24(4):605–622, 2000.

[BKST99] Sergei Bespamyatnikh, Klara Kedem, Michael Segal, and Arie Tamir. Optimal facility location under various distance functions. In *Proceedings of the Workshop on Algorithms and Data Structures*, volume 1663 of *Lecture Notes in Computer Science*, pages 318–329. Springer-Verlag, 1999.

[BL95] Jack Brimberg and Robert F. Love. Estimating distances. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 9–32. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[BMKS99] Boaz Ben-Moshe, Matthew J. Katz, and Michael Segal. Obnoxious facility location: Complete service with minimal harm. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 11, 1999.

[BMM03] Prosenjit Bose, Anil Maheshwari, and Pat Morin. Fast approximations for sums of distances, clustering and the Fermat-Weber problem. *Computational Geometry: Theory and Applications*, 24(3):135–146, 2003.

[Bru78] Peter Brucker. On the complexity of clustering problems. In R. Henn, B. Korte, and W. Oettli, editors, *Proceedings of the Workshop on Optimization and Operations Research*, volume 157 of *Lecture Notes in Economics and Mathematical Systems*, pages 45–54. Springer-Verlag, 1978.

[BS99] Sergei Bespamyatnikh and Michael Segal. Rectilinear static and dynamic discrete 2-center problems. In *Proceedings of the Workshop on Algorithms and Data Structures*, volume 1663 of *Lecture Notes in Computer Science*, pages 276–293. Springer-Verlag, 1999.

[BS01] Sergei Bespamyatnikh and Michael Segal. Rectilinear static and dynamic discrete 2-center problems. *International Journal of Mathematical Algorithms*, 2:149–162, 2001.

[Buc80] Su Buchin. *Lectures on Differential Geometry.* World Scientific, Singapore, 1980.

[Car76] Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice-Hall, Englewood Cliffs, New Jersey, 1976.

[CEK02] James F. Campbell, Andreas T. Ernst, and Mohan Krishnamoorthy. Hub location problems. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 373–407. Springer, New York, 2002.

[CFPS03] Mark Cieliebak, Paola Flocchini, Giuseppe Prencipe, and Nicoal Santoro. Solving the robots gathering problem. In *Proceedings of the Thirtieth International Colloquium on Automata, Languages and Programming*, volume 2719 of *Lecture Notes in Computer Science*, pages 1181–1196, 2003.

[CG99] Moses Charikar and Sudipto Guha. Improved combinatorial algorithms for the facility location and k-median problems. In *Proceedings of the Symposium on Foundations of Computer Science*, pages 378–388. IEEE Computer Society Press, 1999.

[CGTS99] Moses Charikar, Sudipto Guha, Eva Tardos, and David B. Shmoys. A constant-factor approximation algorithm for the k-median problem. In *Proceedings of the Symposium on Theory of Computing*, pages 1–10. ACM, 1999.

[Cha96] Timothy M. Chan. Optimal output-sensitive convex hull algorithms in two and three dimensions. *Discrete and Computational Geometry*, 16:361–368, 1996.

[Cha98] Timothy M. Chan. Geometric applications of a randomized optimization technique. In *Proceedings of the ACM Symposium on Computational Geometry*, volume 14, pages 269–278. ACM Press, 1998.

[Cha99] Timothy M. Chan. More planar two-center algorithms. *Computational Geometry: Theory and Applications*, 13(3):189–198, 1999.

[Cha01] Timothy M. Chan. Dynamic planar convex hull operations in near-logarithmic amortized time. *Journal of the ACM*, 48(1):1–12, 2001.

[CHPK05] Paz Carmi, Sariel Har-Peled, and Matthew J. Katz. On the Fermat-Weber center of a convex object. *Computational Geometry: Theory and Applications*, 32(3):188–195, 2005.

[Chr85] Chrystal. On the problem to construct the minimum circle enclosing n given points in a plane. In *Proceedings of the Edinburgh Mathematical Society*, volume 3, pages 30–33, 1885.

[CM96] Bernard Chazelle and Jiří Matoušek. On linear-time deterministic algorithms for optimization problems in fixed dimensions. *Journal of Algorithms*, 21:579–597, 1996.

[CMB06] Jorge Cortés, Sonia Martínez, and Francesco Bullo. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Transactions on Automatic Control*, 51(8):1298–1298, 2006.

[CMKB04] Jorge Cortés, Sonia Martínez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.

[Coo67] L. Cooper. Solutions of generalized locational equilibrium models. *Journal of Regional Science*, 7:1–18, 1967.

[CPP02] Stephen A. Cook, Jan Pachl, and Irwin S. Pressman. The optimal location of replicas in a network using READ-ONE-WRITE-ALL policy. *Distributed Computing*, 15(1):57–66, 2002.

[CRW91] Vasilis Capoyleas, Günter Rote, and Gerhard Woeginger. Geometric clusterings. *Journal of Algorithms*, 12:341–356, 1991.

[CSK98] Joonsoo Choi, Chan-Su Shin, and Sung Kwon Kim. Computing weighted rectilinear median and center set in the presence of obstacles. In *Proceedings of the Ninth International Symposium on Algorithms and Computation*, volume 1533 of *Lecture Notes in Computer Science*, pages 29–40. Springer-Verlag, 1998.

[CSY87] Richard Cole, Micha Sharir, and Chee K. Yap. On k-hulls and related problems. *SIAM Journal on Computing*, 16(1):61–77, 1987.

[CT90] Ramaswamy Chandrasekaran and Arie Tamir. Algebraic optimization: The Fermat-Weber location problem. *Mathematical Programming*, 46:219–224, 1990.

[CWLG97] Ching-Chuan Chiang, Hsiao-Kuang Wu, Winston W. Liu, and Mario Gerla. Routing in clustered multihop, mobile wireless networks with fading channel. In *Proceedings of IEEE SICON*, pages 197–211, 1997.

[DC97] Zuo Dai and To-yat Cheung. A new heuristic approach for the p-median problem. *Journal of the Operational Research Society*, 48(9):950–960, 1997.

[Dey98] Tamal K. Dey. Improved bounds for planar k-sets and related problems. *Discrete and Computational Geometry*, 19:373–382, 1998.

[DF74] P. M. Dearing and Richard L. Francis. A minimax location problem on a network. *Transportation Science*, 8:333–343, 1974.

[DFK+99] P. Drineas, Alan Frieze, Ravi Kannan, Santosh Vempala, and V. Vinay. Clustering in large graphs and matrices. In *Proceedings of the Symposium on Discrete Algorithms*, volume 10. ACM Press, 1999.

[DK03] Stephane Durocher and David Kirkpatrick. The Gaussian centre of a set of mobile points. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 15, pages 123–127, 2003.

[DK04] Stephane Durocher and David Kirkpatrick. The Gaussian centre and the projection centre of a set of points in $\mathbb{R}^3$. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 16, pages 140–144, 2004.

[DK05a] Stephane Durocher and David Kirkpatrick. Bounded-velocity approximations of the mobile Euclidean 2-centre. In *Proceedings of the Annual Fall Workshop on Computational Geometry and Visualization*, volume 15, pages 48–50, 2005.

[DK05b] Stephane Durocher and David Kirkpatrick. The projection median of a set of points. *Computational Geometry: Theory and Applications*, 2005. Submitted.

[DK05c] Stephane Durocher and David Kirkpatrick. The projection median of a set of points in $\mathbb{R}^2$. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 17, pages 46–50, 2005.

[DK06] Stephane Durocher and David Kirkpatrick. The Steiner centre: Stability, eccentricity, and applications to mobile facility location. *International Journal of Computational Geometry and Applications*, 16(4):345–371, 2006.

[DKSW02] Zvi Drezner, Kathrin Klamroth, Anita Schöbel, and George O. Wesolosky. The Weber problem. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 1–36. Springer, New York, 2002.

[DM00] Inderjit Dhillon and Dharmendra Modha. A data-clustering algorithm on distributed memory multiprocessors. In Mohammed J. Zaki and Ching-Tien Ho, editors, *Proceedings of the Workshop on*

Large-Scale Parallel Data Mining, volume 1759 of *Lecture Notes in Artificial Intelligence*, pages 245–260, 2000.

[Dre84a] Zvi Drezner. The p-center problem: heuristic and optimal algorithms. *Journal of the Operational Research Society*, 35:741–748, 1984.

[Dre84b] Zvi Drezner. The planar two-center and two-median problems. *Transportation Science*, 18(4):351–361, 1984.

[Dre87] Zvi Drezner. On the rectangular p-center problem. *Naval Research Logistics*, 34:229–234, 1987.

[Dre95] Zvi Drezner. Replacing discrete demand with continuous demand. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 33–42. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[DTW86] Zvi Drezner, Jacques-François Thisse, and George O. Wesolowsky. The minimaxmin location problem. *Journal of Regional Science*, 26:87–101, 1986.

[EE04] H. M. M. ten Eikelder and A. A. van Erk. Unification of some least squares clustering methods. *Journal of Mathematical Modelling and Algorithms*, 3:105–122, 2004.

[EH72] D. Jack Elzinga and Donald W. Hearn. Geometrical solutions for some minimax location problems. *Transportation Science*, 6:379–394, 1972.

[EH01] Vladimir Estivill-Castro and Michael E. Houle. Robust distance-based clustering with applications to spatial data mining. *Algorithmica*, 30:216–242, 2001.

[EL95] Horst A. Eiselt and Gilbert Laporte. Objectives in location problems. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 151–180. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[Epp92] David Eppstein. Dynamic three-dimensional linear programming. *ORSA Journal on Computing*, 4(4):360–368, 1992.

[Epp97] David Eppstein. Faster construction of planar two-centers. In *Proceedings of the Symposium on Discrete Algorithms*, pages 131–138, 1997.

[Est99] Vladimir Estivill-Castro. Convex group clustering of large georeferenced data sets. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 11, pages 31–34, 1999.

[Est01] Vladimir Estivill-Castro. Computational geometry provides techniques for approximately solving the p-median problem. Technical Report 2001-01, Department of Computer Science and Software Engineering, University of Newcastle, 2001.

[Fer91] Pierre de Fermat. Tome I. In M. M. P. Tannery and C. Henry, editors, *Oeuvres*. Gauthier-Villars et Fils, Paris, 1891.

[FG88] Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *Proceedings of the Symposium on Theory of Computing*. ACM, 1988.

[FMW83] Richard L. Francis, Leon F. McGinnis, and John A. White. Location analysis. *European Journal of Operational Research*, 12:220–252, 1983.

[FMW00] Sándor P. Fekete, Joseph S. B. Mitchell, and Karin Weinbrecht. On the continuous Weber and k-median problems. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 70–79, 2000.

[FMW05] Sándor P. Fekete, Joseph S. B. Mitchell, and Karin Weinbrecht. On the continuous Fermat-Weber problem. *Operations Research*, 53:61–76, 2005.

[FPT81] Robert J. Fowler, Michael S. Paterson, and Steven L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Information Processing Letters*, 12:133–137, 1981.

[Fre91] Greg N. Frederickson. Parametric search and locating supply centers in trees. In *Proceedings of the Workshop on Algorithms and Data Structures*, volume 519 of *Lecture Notes in Computer Science*, pages 299–319. Springer-Verlag, 1991.

[FW03] Eugene Fink and Derick Wood. Planar strong visibility. *International Journal of Computational Geometry and Applications*, 13(2):173–187, 2003.

[GGH$^+$03] Jie Gao, Leonidas J. Guibas, John Hershberger, Li Zhang, and An Zhu. Discrete mobile centers. *Discrete and Computational Geometry*, 30(1):45–65, 2003.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability*. W. H. Freeman and Company, New York, 1979.

[GJW82] Michael R. Garey, David S. Johnson, and Hans S. Witsenhausen. The complexity of the generalized Lloyd-Max problem. *IEEE Transactions on Information Theory*, 28:255–256, 1982.

[GLS06] Jie Gao, Michael Langberg, and Leonard J. Schulman. Analysis of incomplete data and an intrinsic-dimension Helly theorem. In *Proceedings of the Symposium on Discrete Algorithms*. ACM Press, 2006. To appear.

[GLY02] Eric Gourdin, Martine Labbé, and Hande Yaman. Telecommunication and location. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 275–305. Springer, New York, 2002.

[GMMO00] Sudipto Guha, Nina Mishra, Rajeev Motwani, and Liadan O'Callaghan. Clustering data streams. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 359–366, 2000.

[Gon85] Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.

[Got96] Daniel Gottlieb. All the way with Gauss-Bonnet and the sociology of mathematics. *American Mathematical Monthly*, 103:457–469, 1996.

[Grü67] Branko Grünbaum. *Convex Polytopes*. Wiley and Sons, New York, second edition, 1967.

[GT95] Mario Gerla and Jack Tzu-Chieh Tsai. Multicluster, mobile, multimedia radio network. *Wireless Networks*, 1(3):255–265, 1995.

[Gui98] Leonidas J. Guibas. Kinetic data structures: a state of the art report. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, pages 191–209. A. K. Peters, Ltd., 1998.

[Hak64] S. Louis Hakimi. Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research*, 12(3):450–459, 1964.

[Hak00] S. Louis Hakimi. Location theory. In Rosen, Michaels, Gross, Grossman, and Shier, editors, *Handbook of Discrete and Combinatorial Mathematics*. CRC Press, 2000.

[Han73] Gabriel Y. Handler. Minimax location of a facility in an undirected tree graph. *Transportation Science*, 7:287–293, 1973.

[HE02] Greg Hamerly and Charles Elkan. Alternatives to the k-means algorithm that find better clusterings. In *Proceedings of the Conference on Information and Knowledge Management*, volume 11, pages 600–607, 2002.

[Hei34] F. Heinen. *Über Systeme von Kräften*. Gymnasium zu Cleve (gedrickt bei G.D. Bädeker, Essen), 1834.

[Her93] John Hershberger. A faster algorithm for the two-center decision problem. *Information Processing Letters*, 47:23–29, 1993.

[Her05] John Hershberger. Smooth kinetic maintenance of clusters. *Computational Geometry: Theory and Applications*, 31:3–30, 2005.

[HII+93] S. Hasegawa, H. Imai, M. Inaba, N. Katoh, and J. Nakano. Efficient algorithms for variance-based k-clustering. In *Proceedings of the Pacific Conference on Computer Graphics and Applications*, volume 1, pages 75–89. World Scientific, 1993.

[HLC93a] R. Z. Hwang, Richard Chia-Tung Lee, and R. C. Chang. The generalized searching over separators strategy to solve some np-hard problems in subexponential time. *Algorithmica*, 9:398–423, 1993.

[HLC93b] R. Z. Hwang, Richard Chia-Tung Lee, and R. C. Chang. The slab dividing approach to solve the Euclidean p-center problem. *Algorithmica*, 9:1–22, 1993.

[HLN91] Pierre Hansen, Martine Labbé, and Brigitte Nicolas. The continuous center set of a network. *Discrete Applied Mathematics*, 30:181–195, 1991.

[HLP+87] Pierre Hansen, Martine Labbé, Dominique Peeters, Jacques-François Thisse, and Vernon J. Henderson. *Systems of Cities and Facility Location*. Harwood Academic Publishers, New York, 1987.

[HLPT87] Pierre Hansen, Martine Labbé, Dominique Peeters, and Jacques-François Thisse. Single facility location on networks. *Annals of Discrete Mathematics*, 31:113–146, 1987.

[HM72] S. Louis Hakimi and Shachindra N. Maheshwari. Optimum locations of centers in networks. *Operations Research*, 20(5):967–973, 1972.

[HM89] Arthur P. Hurter and Joseph S. Martinich. *Facility Location and the Theory of Production*. Kluwer Academic Press, Boston, Massachusetts, 1989.

[HM01] Pierre Hansen and Nenad Mladenović. j-means: a new local search heuristic for minimum sum of squares clustering. *Pattern Recognition*, 34(2):405–413, 2001.

[HM03] Trevor S. Hale and Christopher R. Moberg. Location science research: A review. *Annals of Operations Research*, 123:21–35, 2003.

[HM04] Sariel Har-Peled and Soham Mazumdar. Coresets for k-means and k-median clustering and their applications. In *Proceedings of the Symposium on Theory of Computing*, pages 291–300. ACM, 2004.

[HMT98] Pierre Hansen, Nenad Mladenović, and Éric D. Taillard. Heuristic solution of the multisource Weber problem as a p-median problem. *Operations Research Letters*, 22:55–62, 1998.

[HN98] Horst W. Hamacher and Stefan Nickel. Classification of location models. *Location Science*, 6:229–242, 1998.

[Hof01] Michael Hoffmann. Covering polygons with few rectangles. In *Abstracts of the European Workshop on Computational Geometry*, pages 39–42. Universität Berlin, 2001.

[Hof05] Michael Hoffmann. A simple linear algorithms for computing rectilinear 3-centers. *Computational Geometry*, 31:150–165, 2005.

[Hon95] Ross Honsberger. *Episodes in Nineteenth and Twentieth Century Euclidean Geometry.* New Mathematical Library. Mathematical Association of America, New York, 1995.

[HP98] Dorit S. Hochbaum and Anu Pathria. Locating centers in a dynamically changing network, and related problems. *Location Science*, 6:243–256, 1998.

[HP04] Sariel Har-Peled. Clustering motion. *Discrete and Computational Geometry*, 31(4):545–565, 2004.

[HPRT85] Pierre Hansen, Dominique Peeters, Denis Richard, and Jacques-François Thisse. The minisum and minimax location problems revisited. *Operations Research*, 33:1251–1265, 1985.

[HRS04] Hai Huang, Andréa W. Richa, and Michael Segal. Approximation algorithms for the mobile piercing set problem with applications to clustering in ad-hoc networks. *Mobile Networks and Applications*, 9(2):151–161, 2004.

[HS86] Dorit S. Hochbaum and David B. Shmoys. A unified approach to approximate algorithms for bottleneck problems. *Journal of the ACM*, 33:533–550, 1986.

[HS91] John Hershberger and Subhash Suri. Finding tailored partitions. *Journal of Algorithms*, 12:431–463, 1991.

[HT91] Refael Hassin and Arie Tamir. Improved complexity bounds for location problems on the real line. *Operations Research Letters*, 10:395–402, 1991.

[HT04] Kuang-Han Hsieh and Fan-Chih Tien. Self-organizing feature maps for solving location-allocation problems with rectilinear distances. *Computers and Operations Research*, 31:1017–1031, 2004.

[IKI94] Mary Inaba, Naoki Katoh, and Hiroshi Imai. Applications of weighted voronoi diagrams and randomization to variance-based k-clustering. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 332–339, 1994.

[Ind99] Piotr Indyk. Sublinear time algorithms for metric space problems. In *Proceedings of the Symposium on the Theory of Computing*, volume 31, pages 428–434, 1999.

[JCM03] Tomas Johansson and Lenka Carr-Motyčková. On clustering in ad hoc networks. In *First Swedish National Computer Networking*, page 14T, 2003.

[JD88] Anil K. Jain and Richard C. Dubes. *Algorithms for Clustering Data*. Prentice Hall, Englewood Cliffs, New Jersey, 1988.

[JK94] Jerzy W. Jaromczyk and Miroslaw Kowaluk. An efficient algorithm for the Euclidean two-center problem. In *Proceedings of the Symposium on Computational Geometry*, volume 10, pages 303–311. ACM Press, 1994.

[JK95] Jerzy W. Jaromczyk and Miroslaw Kowaluk. A geometric proof of the combinatorial bounds for the number of optimal solutions for the 2-center Euclidean problem. In *Proceedings of the Canadian Conference on Computational Geometry*, volume 7, pages 19–24, 1995.

[JM93] Shreesh Jadhav and Asish Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. In *Proceedings of the ACM Symposium on Computational Geometry*, volume 9, pages 83–90, 1993.

[JV99] Kamal Jain and Vijay V. Vazirani. Primal-dual approximation algorithms for metric facility location and k-median problems. In IEEE, editor, *Proceedings of the Symposium on Foundations of Computer Science*, pages 2–13. IEEE Computer Society Press, 1999.

[KC92] Ming-Tat Ko and Y. T. Ching. Linear time algorithms for the weighted tailored 2-partition problem and the weighted rectilinear 2-center problem under L_∞-distance. *Discrete and Applied Mathematics*, 40:397–410, 1992.

[KH79a] Oded Kariv and S. Louis Hakimi. An algorithmic approach to network location problems. I: The p-centers. *SIAM Journal of Applied Mathematics*, 37(3):513–538, 1979.

[KH79b] Oded Kariv and S. Louis Hakimi. An algorithmic approach to network location problems II: The p-medians. *SIAM Journal of Applied Mathematics*, 37(3):539–560, 1979.

[Kim] Clark Kimberling. Clark Kimberling's encyclopedia of triangle centers. http://faculty.evansville.edu/ck6/encyclopedia/ETC.html.

[Kim98] Clark Kimberling. Triangle centers and central triangles. *Congressus Numerantium*, 129:1–295, 1998.

[KLC90] Ming-Tat Ko, Richard Chia-Tung Lee, and J. S. Chang. An optimal approximation algorithm for the rectilinear m-center problem. *Algorithmica*, 5:341–352, 1990.

[KM97] Yaakov S. Kupitz and Horst Martini. Geometric aspects of the generalized Fermat-Torricelli problem. In *Intuitive Geometry*, volume 6, pages 55–127. Bolyai Society Mathematical Studies, Budapest, 1997.

[KMN+00] Tapas Kanungo, David M. Mount, Nathan Netanyahu, Christine Piatko, Ruth Silverman, and Angela Y. Wu. The analysis of a simple k-means clustering algorithm. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 100–109, 2000.

[KMN+02a] Tapas Kanungo, David M. Mount, Nathan Netanyahu, Christine Piatko, Ruth Silverman, and Angela Y. Wu. An efficient k-means clustering algorithm: Analysis and implementation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):881–892, 2002.

[KMN+02b] Tapas Kanungo, David M. Mount, Nathan Netanyahu, Christine Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for k-means clustering. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 10–18, 2002.

[KNS00] Matthew J. Katz, Frank Nielsen, and Michael Segal. Maintenance of a piercing set for intervals with applications. In D. T. Lee and Shang-Hua Teng, editors, *Proceedings of the Eleventh International Conference on Algorithms and Computation (ISAAC)*, volume 1969 of *Lecture Notes in Computer Science*, pages 552–563. Springer, 2000.

[KNW02] Oliver Karch, Hartmut Noltemeier, and Thomas Wahl. Location and robotics. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 409–438. Springer, New York, 2002.

[KR99] Stavros G. Kolliopoulos and Satish Rao. A nearly linear-time approximation scheme for the Euclidean k-median problem. In *Proceedings of the Seventh Annual European Symposium on Algorithms*, volume 1643 of *Lecture Notes in Computer Science*, pages 387–389, 1999.

[KS86] David G. Kirkpatrick and Raimund Seidel. The ultimate planar convex hull algorithm? *SIAM Journal on Computing*, 15:287–299, 1986.

[KS93] Matthew J. Katz and Micha Sharir. An expander-based approach to geometric optimization. In *Proceedings of the ACM Symposium on Computational Geometry*, volume 9, pages 198–207. ACM Press, 1993.

[KSS04] Amit Kumar, Yogish Sabharwal, and Sandeep Sen. A simple linear time $(1+\epsilon)$-approximation algorithm for k-means clustering in any dimensions. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*, pages 454–462, 2004.

[KV97] Jakob Krarup and Steven Vajda. On Torricelli's geometrical solution to a problem of Fermat. *IMA Journal of Management Mathematics*, 8(3):215–224, 1997.

[Llo82] Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[LMW88] Robert F. Love, James G. Morris, and George O. Wesolowsky. *Facilities Location*. North-Holland, New York, 1988.

[Mac67] James B. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. University of California Press, 1967.

[Mat00] Jiří Matoušek. On approximate geometric k-clustering. *Discrete and Computational Geometry*, 24:61–84, 2000.

[MC84] I. Douglas Moon and Sohail S. Chaudhry. An analysis of network location problems with distance calculations. *Management Science*, 30(3):290–307, 1984.

[MC86a] Emanuel Melachrinoudis and Thomas P. Cullinane. Locating an obnoxious facility within a polygonal region. *Annals of Operations Research*, 6:137–145, 1986.

[MC86b] Emanuel Melachrinoudis and Thomas P. Cullinane. Locating an undesirable facility with a minimax criterion. *European Journal of Operational Research*, 24:239–246, 1986.

[ME98] Alan T. Murray and Vladimir Estivill-Castro. Cluster discovery techniques for exploratory spatial data analysis. *International Journal of Geographical Information Science*, 12(5):431–443, 1998.

[Meg83] Nimrod Megiddo. Linear-time algorithms for linear programming in R^3 and related problems. *SIAM Journal on Computing*, 4:759–776, 1983.

[Mer03] Peter Merz. An iterated local search approach for minimum sum-of-squares clustering. In M. R. Berthold et al., editor, *Advances in Intelligent Data Analysis V*, volume 2810 of *Lecture Notes in Computer Science*, pages 286–296. Springer-Verlag, 2003.

[MF90] Pitu B. Mirchandani and Richard L. Francis, editors. *Discrete Location Theory*. John Wiley and Sons, Inc., New York, 1990.

[MIH81] Shigeru Masuyama, Toshihide Ibaraki, and Toshiharu Hasegawa. The computational complexity of the m-center problems on the plane. *Transactions of the Institute of Electronics and Communication Engineers of Japan*, 64E:57–64, 1981.

[Min77] Edward Minieka. The centers and medians of a graph. *Operations Research*, 25:641–650, 1977.

[MNP+04] David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A computational framework for incremental motion. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 200–209, 2004.

[Mou05] David M. Mount, 2005. Personal communication.

[MS84] Nimrod Megiddo and Kenneth J. Supowit. On the complexity of some common geometric location problems. *SIAM Journal on Computing*, 13(1):182–196, 1984.

[MS02] Vladimir Marianov and Daniel Serra. Location problems in the public sector. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 119–150. Springer, New York, 2002.

[NN04] Frank Nielsen and Richard Nock. Approximating smallest enclosing balls. In *Proceedings of the Workshop on Computational Geometry and Applications*, volume 3045 of *Lecture Notes in Computer Science*, pages 147–157. Springer-Verlag, 2004. Part III.

[Nus97] Doron Nussbaum. Rectilinear p-piercing problems. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 316–323. ACM Press, 1997.

[Ohl97] M. Ohlemüller. Tabu search for large location-allocation problems. *Journal of the Operational Research Society*, 48:745–750, 1997.

[OL81] Mark H. Overmars Overmars and Jan van Leeuwen. Maintenance of configurations in the plane. *Journal of Computer and System Sciences*, 23(2):166–204, 1981.

[Ost75] Lawrence M. Ostresh, Jr. An efficient algorithm for solving the two center location-allocation problem. *Journal of Regional Science*, 15(2):209–216, 1975.

[Pao99] Cappanera Paola. A survey on obnoxious facility location. Technical Report TR-99-11, Dipartimento di Informatica, Università di Pisa, 1999.

[Pla95] Frank Plastria. Continuous location problems. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 225–262. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[Pla02] Frank Plastria. Continuous covering location problems. In Zvi Drezner and Horst W. Hamacher, editors, *Facility Location: Applications and Theory*, pages 37–80. Springer, New York, 2002.

[Ric97] Thomas J. Richardson. Approximation of planar convex sets from hyperplane probes. *Discrete and Computational Geometry*, 18:151–177, 1997.

[Ros92] Kenneth E. Rosing. An optimal method for solving the (generalized) multi-Weber problem. *European Journal of Operational Research*, 58(3):414–426, 1992.

[RT90] Jean-Marc Robert and Godfried T. Toussaint. Computational geometry and facility location. In *Proceedings of the International Conference on Operations Research and Management Science*, pages B–1–B–10, 1990.

[RW87] Gregory J. E. Rawlins and Derick Wood. Optimal computation of finitely oriented convex hulls. *Information and Computation*, 72:150–166, 1987.

[Sal66] George Thomas Sallee. A valuation property of Steiner points. *Mathematika*, 13:76–82, 1966.

[Sch73] Alain Schärlig. About the confusion between the center of gravity and Weber's optimum. *Regional and Urban Economics*, 3(4):371–382, 1973.

[Sch03] Konrad Schlude. From robotics to facility location: Contraction functions, Weber point, convex core. Technical Report 403, Eidgenössische Technische Hochschule Zürich, 2003.

[Seg99] Michael Segal. *Covering Point Sets and Accompanying Problems*. PhD thesis, Ben-Gurion University, 1999.

[Sha96] Jacob Sharony. An architecture for mobile radio networks with dynamically changing topology using virtual subnets. *Mobile Networks and Applications*, 1(1):75–86, 1996.

[Sha97] Micha Sharir. A near-linear algorithm for the planar 2-center problem. *Discrete and Computational Geometry*, 18(2):125–134, 1997.

[She64] Geoffrey C. Shephard. Approximation problems for convex polyhedra. *Mathematika*, 11:9–18, 1964.

[She66] Geoffrey C. Shephard. Steiner points of convex polytopes. *Canadian Journal of Mathematics*, 18:1294–1300, 1966.

[She68] Geoffrey C. Shephard. A uniqueness theorem for the Steiner point of a convex region. *Journal of the London Mathematical Society*, 43:439–444, 1968.

[Sla81] Peter J. Slater. On locating a facility to service areas within a network. *Operations Research*, 29(3):523–531, 1981.

[Sny04] Lawrence V. Snyder. Facility location under uncertainty: A review. Technical Report 04T-015, Department of Industrial and Systems Engineering, Lehigh University, 2004.

[Sny05] Lawrence V. Snyder. Facility location under uncertainty: A review. Technical Report 04T-015, Department of Industrial and Systems Engineering, Lehigh University, 2005.

[SO95] Atsuo Suzuki and Atsuyuki Okabe. Using Voronoi diagrams. In Zvi Drezner, editor, *Facility Location: A Survey of Applications and Methods*, volume 106, pages 103–118. Springer Series in Operations Research, Springer Verlag, New York, 1995.

[SP97] Jinhyeon Sohn and Sungsoo Park. A linear program for the two-hub location problem. *European Journal of Operational Research*, 100:617–622, 1997.

[SS97] Georg Still and Martin Streng. The Chebyshev hyperplane optimization problem. *Journal of Global Optimization*, 11:361–376, 1997.

[SS01] Maria Paola Scaparra and Maria Grazia Scutellà. Facilities, locations, customers: Building blocks of location models - a survey. Technical Report TR-01-18, Dipartimento di Informatica, Università di Pisa, 2001.

[SS05] Yogish Sabharwal and Sandeep Sen. A linear time algorithm for approximate 2-means clustering. *Computational Geometry: Theory and Applications*, page 14, 2005. To appear.

[Ste81] Jakob Steiner. *Gesammelte Werke*, volume 1. Reimer, Berlin, 1881.

[SW96] Micha Sharir and Emo Welzl. Rectilinear and polygonal p-piercing and p-center problems. In *Proceedings of the ACM Symposium on Computational Geometry*, volume 12, pages 122–132. ACM Press, 1996.

[Syl57] James Joseph Sylvester. A question in the geometry of situation. *Quarterly Journal of Mathematics*, 1:79, 1857.

[Tam88] Arie Tamir. Improved complexity bounds for center location problems on networks by using dynamic data structures. *SIAM Journal on Discrete Mathematics*, 1(3):377–396, 1988.

[Tam96] Arie Tamir. An $O(pn^2)$ algorithm for the p-median and related problems on tree graphs. *Operations Research Letters*, 19(2):59–64, 1996.

[TFL83a] Barbaros C. Tansel, Richard L. Francis, and Timothy J. Lowe. Location on networks: A survey. part I: The p-center and p-median problems. *Management Science*, 29(4):482–497, 1983.

[TFL83b] Barbaros C. Tansel, Richard L. Francis, and Timothy J. Lowe. Location on networks: A survey. part II: Exploiting tree network structure. *Management Science*, 29(4):498–511, 1983.

[Tha03] Charles Thas. A generalization of the Lemoine point. *Forum Geometricum*, 3:161–167, 2003.

[Uth05a] Patchrawat Uthaisombut. Real-time kinetic algorithms. Technical report, Department of Computer Science, University of Pittsburgh, 2005.

[Uth05b] Patchrawat Uthaisombut. Real-time kinetic algorithms. In *Proceedings of the Annual Fall Workshop on Computational Geometry and Visualization*, volume 15, pages 3–5, 2005.

[VKKR03] W. Fernandez de la Vega, Marek Karpinski, Claire Kenyon, and Yuval Rabani. Approximation schemes for clustering problems. In *Proceedings of the Symposium on Theory of Computing*, volume 35, pages 50–58, 2003.

[VP03] Boris S. Verkhovsky and Yuriv S. Polyakov. Highly efficient algorithm for two-switch location problem. In K. J. Engemann and G. E. Lasker, editors, *Advances in Decision Technology and Intelligent Information Systems*, volume 4, pages 51–55. The International Institute for Advanced Studies in Systems Research and Cybernetics, Windsor, Canada, 2003.

[VW04] Frédéric Vivien and Nicolas Wicker. Minimal enclosing parallelepiped in 3D. *Computational Geometry*, 29:177–190, 2004.

[Web22] Alfred Weber. *Über den Standort der Industrie.* J. C. B. Mohr (Paul Siebeck), Tübingen, 1922.

[Wei] Eric W. Weisstein. Mathworld – a Wolfram web resource. http://mathworld.wolfram.com.

[Wei37] Endre Weiszfeld. Sur le point pour lequel la somme des distances de n points donnés est minimum. *Tôhoku Mathematical Journal*, 43:355–386, 1937.

[Wel91] Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In H. Maurer, editor, *New Results and New Trends in Computer Science*, volume 555 of *Lecture Notes in Computer Science*, pages 359–370. Springer-Verlag, 1991.

[Wes73] George O. Wesolowsky. Dynamic facility location. *Management Science*, 19:1241–1248, 1973.

[Wes93] George O. Wesolowsky. The Weber problem: History and perspectives. *Location Science*, 1(1):5–23, 1993.

[WH73a] Richard E. Wendell and Arthur P. Jr. Hurter. Locations theory, dominance, and convexity. *Operations Research*, 21:314–320, 1973.

[WH73b] Richard E. Wendell and Arthur P. Jr. Hurter. Optimal locations on a network. *Transportation Science*, 7:18–33, 1973.

[Win79] Peter M. Winkler. Most probable position from three L.O.P.'s. Manuscript, 1979.

[WO04] Lan Wang and Stephan Olariu. A unifying look at clustering in mobile ad hoc networks. *Wireless Communications and Mobile Computing*, 4(6):623–637, 2004.

[WT75] George O. Wesolowsky and William G. Truscott. The multiperiod location-allocation of facilities. *Management Science*, 22:57–65, 1975.

[XFS03] Sheng Xu, Robert M. Freund, and Jie Sun. Solution methodologies for the smallest enclosing circle problem. *Computational Optimization and Applications*, 25:283–292, 2003.

[Yiu04] Paul Yiu. A tour of triangle geometry. In *Proceedings of the Annual Meeting of the Florida Section of the Mathematical Association of America*, volume 37, 2004.

Appendix A

List of Symbols

$\mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$ is the set of integers.

$\mathbb{Z}^+ = \{x \mid x \in \mathbb{Z},\ x \geq 0\}$ is the set of non-negative integers.

$\mathbb{N} = \{x \mid x \in \mathbb{Z},\ x > 0\}$ is the set of natural numbers. By convention, we assume $0 \notin \mathbb{N}$.

$\mathbb{R}$ is the set of real numbers.

$\mathbb{R}^+ = \{x \mid x \in \mathbb{R},\ x \geq 0\}$ is the set of non-negative real numbers.

$\mathbb{Q} = \{\frac{p}{q} \mid p \in \mathbb{Z},\ q \in \mathbb{N}\}$ is the set of rational numbers.

$\mathbb{Q}^+ = \{x \mid x \in \mathbb{Q},\ x \geq 0\}$ is the set of non-negative rational numbers.

(p_x, p_y) denote the x- and y-coordinates of a point $p \in \mathbb{R}^2$.

(p_x, p_y, p_z) denote the x-, y-, and z-coordinates of a point $p \in \mathbb{R}^3$.

p_i denotes the ith coordinate of a point $p = (p_1, \ldots, p_i, \ldots, p_d) \in \mathbb{R}^d$.

$\mathscr{P}(A)$ denotes the power set of set A.

$\widetilde{\mathscr{P}}(A)$ denotes the set of all nonempty subsets of set A. $\widetilde{\mathscr{P}}(A) = \mathscr{P}(A) - \{\varnothing\}$.

$\widehat{\mathscr{P}}(A)$ denotes the set of all nonempty bounded subsets of set A.

$\mathring{\mathscr{P}}(A)$ denotes the set of all nonempty finite subsets of set A. $\mathring{\mathscr{P}}(A) \subseteq \widehat{\mathscr{P}}(A) \subseteq \widetilde{\mathscr{P}}(A) \subseteq \mathscr{P}(A)$.

$\overline{A}$ denotes the closure of set A. That is, $\overline{A}$ is the intersection of all closed sets that contain A.

$BB(A)$ denotes the bounding box of set A (including its interior).

$CH(A)$ denotes the convex hull of set A. That is, $CH(A)$ is the intersection of all convex sets that contain A.

$MEC(A)$ denotes the minimum enclosing circle of set A in $\mathbb{R}$ (including the interior of the circle).

$\mathrm{circ}(A)$ denotes the circumference of a minimum enclosing circle of set A in $\mathbb{R}^2$.

V_A denotes the set of extreme points of set A.

$\partial(A)$ denotes the boundary of set A. That is, $\partial(A) = \overline{A^c} \cap \overline{A}$.

$\deg(p)$ denotes the degree of vertex p. When p is an extreme point of P, $\deg(p)$ denotes the number of neighbours of p on the graph induced by $CH(P)$.

id_U denotes the identity function $\mathrm{id} : U \to U$ on universe U.

$|x|$ denotes the absolute value of $x \in \mathbb{R}$.

$||x||_p$ denotes the ℓ_p (Minkowski) norm of a point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, where $||x||_p = \left(\sum_{i=1}^d |x_i|^p \right)^{1/p}$.

$||x||_1$ denotes the ℓ_1 (rectilinear) norm of a point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, where $||x||_1 = \sum_{i=1}^d |x_i|$.

$||x||$ denotes the ℓ_2 (Euclidean) norm of a point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, where $||x|| = \sqrt{\sum_{i=1}^d x_i^2}$.

$||x||_\infty$ denotes the ℓ_∞ (Chebyshev) norm of a point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, where $||x||_\infty = \lim_{p\to\infty} ||x||_p = \max_{i=1}^d |x_i|$.

Γ_d denotes the d-dimensional Steiner centre.

G_d denotes the d-dimensional Gaussian median.

C_d denotes the d-dimensional centre of mass.

R_d denotes the d-dimensional rectlinear 1-centre.

S_d denotes the d-dimensional rectlinear 1-median.

M_d denotes the d-dimensional Euclidean 1-median.

Π_d denotes the d-dimensional projection median.

Ξ_d denotes the d-dimensional Euclidean 1-centre.

Υ_d denotes a d-dimensional approximation function.

F_d denotes a d-dimensional function used to define the point of reflection in a reflection-based 2-centre function.

w_d denotes the d-dimensional Gaussian weight.

g_d denotes the d-dimensional Gaussian median weight.

u_θ denotes a unit vector in $\mathbb{R}^2$, $u_\theta = (\cos\theta, \sin\theta)$.

$u_{\theta,\phi}$ denotes a unit vector in $\mathbb{R}^3$, $u_{\theta,\phi} = (\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi)$.

l_θ denotes the line in $\mathbb{R}^2$ through the origin parallel to vector u_θ.

$l_{\theta,\phi}$ denotes the line in $\mathbb{R}^3$ through the origin parallel to vector $u_{\theta,\phi}$.

P_θ denotes the projection of set $P \in \mathscr{P}(\mathbb{R}^2)$ onto line l_θ.

$P_{\theta,\phi}$ denotes the projection of set $P \in \mathscr{P}(\mathbb{R}^3)$ onto line $l_{\theta,\phi}$.

$\langle u, v\rangle$ denotes the inner product (dot product) of two points $u, v \in \mathbb{R}^d$.

$\max(A)$ denotes the maximum extreme point of a the closure of a set of collinear points, A in $\mathbb{R}$.

$\operatorname{med}(A)$ denotes the median of a set of collinear points, A in $\mathbb{R}$.

$\operatorname{mid}(A)$ denotes the midpoint of the closure of a set of collinear points, A in $\mathbb{R}$.

$\min(A)$ denotes the minimum extreme point of the closure of a set of collinear points, A in $\mathbb{R}$.

$v_{\max}$ denotes the maximum velocity of a mobile approximation function.

κ refers to the stability of an approximation function.

λ refers to the approximation factor of an approximation function.

$\operatorname{ext}(P, \theta)$ denotes an extreme point of set P in $\mathbb{R}^2$ in direction $(\cos\theta, \sin\theta)$. That is, $p = \operatorname{ext}(P, \theta)$ if and only if there exists a half-plane $H^+ \subseteq \mathbb{R}^2$ with outer normal $(\cos\theta, \sin\theta)$ such that $\overline{P} \cap \overline{H^+} = \{p\}$. Note, the extreme point in a given direction θ may not exist (if it does exist, then it is unique by the above definition).

$\operatorname{Ext}(P, \theta) = \lim_{\phi \to \theta^+} \operatorname{ext}(P, \phi)$. *See* $\operatorname{ext}(P, \theta)$.

Index