The Gaussian Centre of a Set of Mobile Points

Extended Abstract

Steph Durocher* and David Kirkpatrick†
Department of Computer Science, University of British Columbia
Vancouver BC, Canada

Abstract

Given a set of client positions as input, facility location attempts to find positions for a set of facilities to optimize some objective function. In mobile facility location, clients undergo continuous motion with bounded velocity and the problem becomes maintaining the position of a mobile facility. This new area within the classical field of facility location is attractive for the many new challenges it presents, problems which did not exist in static facility location. The velocity of the mobile Euclidean centre in two or more dimensions (centre of the smallest enclosing sphere) is unbounded [BBKS00]. It is natural to impose some upper bound on the velocity of a facility. Thus, one must approximate the position of the centre. The goal is to balance a good approximation factor while maintaining low velocity. To solve these two opposing goals, we present the *Gaussian centre* as an approximation to the mobile Euclidean centre.

1 Introduction

Facility location has a considerable history as a highly-active area of research. The literature within the communities of operations research, computational geometry, networks, graph theory, and complexity theory presents extensive examinations and solutions to the problems of *static* facility location. The traditional problem of locating a facility to optimize some function of the input set of client positions was first formally defined by Alfred Weber [Web22] early in the last century. The field has been studied extensively since the 1960's, but only within the last few years have these questions been posed in the *mobile* setting [BBKS00, AH01]. Given a set of clients moving continuously over a temporal dimension, completely new problems arise. These include bounding velocity, maintaining continuity, and approximating the location of a mobile facility. Furthermore, the techniques employed to solve a particular facility location problem do not necessarily extend to a solution to its mobile counterpart. The challenges presented by mobile facility location find themselves particularly relevant given the applicability of mobile computing to the wireless telecommunication industries involving cellular and radio ethernet.

*durocher@cs.ubc.ca
†kirk@cs.ubc.ca

2 Static Facility Location

2.1 Position and Distance

Problems in facility location involve optimizing some function of the positions of a set of clients. Given a set of n clients, $S = \{v_1, \ldots, v_n\}$, select locations for a set of k facilities, $F = \{f_1, \ldots, f_k\}$, to optimize the objective function $g(S, F)$.

Client positions are often represented by points in Euclidean space, $\mathbb{R}^d$. Let $p_i \in \mathbb{R}^d$ be the position of client $v_i \in S$. If $|S|$ is infinite, we require clients to be located within some bounded region of $\mathbb{R}^d$.

We examine problems under Euclidean distance given by the ℓ_2 norm:

$$\forall u, v \in \mathbb{R}^d,\ \ell_2(u, v) = ||u - v||_2 = \sqrt{\sum_{i=1}^{d} (u_i - v_i)^2}\ . \quad (1)$$

Section 6 discusses the mobile centre problem under ℓ_1 and ℓ_∞.

2.2 The Euclidean Centre

Locating the centre of a set of points in the plane is a fundamental problem of facility location. Under ℓ_2 we refer to this centre as the Euclidean centre.

Definition 1 *Let S be a set of clients within some bounded region of $\mathbb{R}^2$. The* Euclidean centre *of S is the point $\eta \in \mathbb{R}^2$ that minimizes the maximum ℓ_2 distance to any client in S.*

The Euclidean centre η is the unique point x that minimizes

$$g(S, x) = \max_{v_i \in S} \ell_2(p_i, x)\ . \quad (2)$$

Interpreted geometrically, η is the centre of the smallest circle that contains S. Under $\mathbb{R}^d$, η is the centre of the smallest d-sphere that contains S. Megiddo [Meg83] gives a $\Theta(n)$ linear programming solution to the static Euclidean centre problem in $\mathbb{R}^2$. Agarwal et al. [AST93] extend this result to $\mathbb{R}^d$ for any fixed d in $O(d^{O(d)}n)$ time. Every client must be examined, giving a lower bound of $\Omega(n \cdot d)$, or $\Theta(n)$.

3 Mobile Facility Location

3.1 Mobile Clients

Until recently, only discrete motion of clients had been considered. Such problems, termed *dynamic facility location*, attempt to optimize the objective function summed over a finite set of discrete time slots, $T = \{t_1, \ldots, t_f\}$ [Wes73, BGKS98].

Bespamyatnikh et al. [BBKS00] first examined clients moving under *continuous motion*. Let $T = [t_0, t_f]$ be a time interval. Let $S = \{v_1, \ldots, v_n\}$ be a set of mobile clients. Let $p_i : T \to \mathbb{R}^2$ be a continuous function that assigns a position in $\mathbb{R}^2$ to client $v_i \in S$ at every instant $t \in T$. We represent the position of point v_i by $p_i(t)$ or simply by p_i if v_i is static. We impose a constant upper bound (usually 1) on the magnitude of the velocity of points[1]. If $p_i(t)$ is differentiable, then $||p_i'(t)||_2 \leq 1$. Otherwise, this condition is generalized as

$$\forall t_1, t_2 \in T,\ ||p_i(t_1) - p_i(t_2)||_2 \leq |t_1 - t_2| \ . \tag{3}$$

3.2 The Mobile Euclidean Centre

The definition of the mobile Euclidean centre is a direct extension of its static definition.

Definition 2 *Let $T = [t_0, t_f]$ be a time interval. Let S be a set of mobile clients within some bounded region of $\mathbb{R}^2$. The* mobile Euclidean centre *of S is a function, $\eta : T \to \mathbb{R}^2$, such that for every $t \in T$, $\eta(t)$ is the Euclidean centre of $S(t)$.*

3.3 Velocity of the Euclidean Centre

Bespamyatnikh et al. [BBKS00] show that in $\mathbb{R}^2$ under ℓ_2, for any $\sigma \geq 0$, there exists a configuration of mobile clients such that the Euclidean centre moves with velocity at least σ. Thus, even when clients are limited to unit velocity, no bound exists on the velocity of the Euclidean centre.

It is natural to impose some upper bound on the velocity of facilities. Therefore, we approximate the mobile centre.

3.4 Approximation Metrics

Let $\eta : T \to \mathbb{R}^2$ be the position of the Euclidean centre. Let $a : T \to \mathbb{R}^2$ be the position of the approximated centre. A good approximation $a(t)$ optimizes two criteria:

1. The maximum distance from $a(t)$ to any client remains close to the maximum distance from $\eta(t)$ to any client.
2. The velocity of $a(t)$ has a low upper bound.

The relative difference between the maximum distances to any client is

$$\lambda = \max_t \frac{\max_i ||p_i - a(t)||_2}{\max_j ||p_j - \eta(t)||_2} \ . \tag{4}$$

The velocity of the approximate centre is bound by b if

$$\forall t_1, t_2,\ ||a(t_1) - a(t_2)||_2 \leq b \cdot |t_1 - t_2| \ . \tag{5}$$

For a given approximation $a(t)$, let $v_{\max}$ be the infimum over all such b.

3.5 Related Work

Agarwal and Har-Peled [AH01] maintain the approximate mobile centre in $\mathbb{R}^2$ under ℓ_∞ and ℓ_2. Their approximations do not require continuity or bounded velocity in the motion of the approximated centre; rather, their aim is to minimize the number of events processed and the update cost per event using kinetic data structures. Agarwal et al. [AGG02] maintain the approximate mobile median in $\mathbb{R}$. Bespamyatnikh et al. [BBKS00] maintain approximations to the mobile centre and mobile median in $\mathbb{R}^2$ under ℓ_∞ and ℓ_2. These include a point on the convex hull, a bounding box, the centre of mass, and linear combinations of these. We discuss these further in Sections 4.1 and 6.

4 The Gaussian Centre

4.1 Motivation

To find the Euclidean centre of a set of clients S, one needs only to consider the extreme points: those that define the convex hull of S. Bespamyatnikh et al. [BBKS00] note that selecting a point on the convex hull to approximate the Euclidean centre guarantees $\lambda \leq 2$ and $v_{\max} \leq 1$. In terms of defining a central point, this approximation is poor; however, the convex-hull approximation provides a reference to which we can compare other approximations. Reducing λ increases $v_{\max}$ and vice-versa. The challenge lies in understanding the trade-off between the degree of the approximation, $1 \leq \lambda \leq 2$, and the maximum velocity with which the approximated centre may move, $1 \leq v_{\max} < \infty$.

4.2 Gaussian Centre Definition

We define the Gaussian centre as a weighted centre of mass of the clients. The weights are defined to change continuously as clients move along, join, or leave the convex hull.

[1] Any constant upper bound on velocity may be used. Since there is no unit of reference, without loss of generality, we select 1 for simplicity.

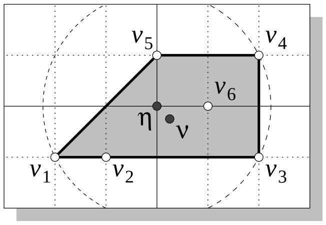

Figure 1: the Gaussian and Euclidean centres, ν and η

Let $C_S(t) \subseteq S$ be the convex hull of S at time t. For each $v_i \in C_S(t)$, let $w_i(t) = \pi - \theta_i(t)$, where $\theta_i(t)$ is the interior angle formed on the convex hull at v_i at time t. For each $v_i \in S - C_S(t)$, let $w_i(t) = 0$.

Definition 3 *The* Gaussian centre *of S, $\nu(t)$, is the normalized weighted mean of the clients in S:*

$$\nu(t) = \frac{1}{2\pi} \sum_i w_i(t) p_i(t) \ . \tag{6}$$

For example, let static clients $v_1, \ldots, v_6$ have positions $(-2,-1), (-1,-1), (2,-1), (2,1), (0,1), (1,0)$, respectively. See Figure 1. Since $w_i = \pi - \theta_i$, clients have weights $3\pi/4,\ 0,\ \pi/2,\ \pi/2,\ \pi/4, 0$, respectively. The Gaussian centre of S, ν, lies in position $(1/4, -1/4)$. The Euclidean centre of S, η, is at the origin.

The Gaussian diagram provides an equivalent definition for client weights. In two dimensions, the Gaussian diagram G_S of the convex hull C_S divides the unit circle into sectors such that the weight of each client in C_S is given by the length of its corresponding arc in G_S. The weight of a client v_i on the convex hull corresponds to the angular difference between the normals of the edges adjacent to v_i.

The Gaussian centre has several desirable properties. First, only those clients on the convex hull can determine the location of the centre. Intuitively, the more significant a client's presence on the convex hull is, the greater the weight assigned to it. When clients move, their weights change continuously, including clients that move from the interior onto the convex hull or vice-versa. This continuous change in weights ensures the Gaussian centre always moves continuously.

4.3 Approximation Factor

The Gaussian centre approximates the Euclidean centre to a factor of $\lambda \approx 1.1153$. This maximum is achieved by an arc opposite a lone vertex as displayed in Figure 2 when $\theta_1 = \theta_2 \approx 0.81047$. The exact values are given by the solution to

$$\cos\theta_1 + (\theta_1 - \pi)\sin\theta_1 + 1 = 0 \ . \tag{7}$$

The complete proof will be available in the full paper. Below is a sketch of the proof.

Let m be the furthest client from ν.

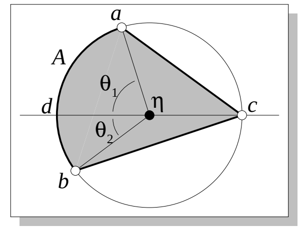

Figure 2: maximizing λ

Definition 4 *The motion of client a is* lengthening *if it causes an increase in $\ell_2(m, \nu)$.*

When $a \neq m$, the projection of a lengthening velocity of a onto $\overline{m\nu}$ points away from m.

Definition 5 *A* maximizing *configuration is a set of clients S whose positions maximize λ.*

Since λ is a ratio, λ is independent of scaling. Without loss of generality, we assume $\max_j \ell_2(p_j, \eta) = 1$. Therefore, all clients lie within or on the unit circle centred at η.

Lemma 1 *In any maximizing configuration, all clients in S lie on the unit circle centred at η.*

Proof sketch. By contradiction. Assume S is a maximizing configuration such that some client a on the convex hull lies inside the unit circle. Let b and c be a's neighbours on the hull. Let $\beta/2\pi$ be the weight of a. See Figure 3a. When a moves with unit velocity toward b, ν moves with velocity $\frac{1}{2\pi}(-\sin\beta\cos\beta + \beta, -\sin^2\beta)$ relative to $\overline{ab}$. When the velocity of a is perpendicular to $\overline{ab}$, away from the interior of S, ν moves with velocity $\frac{1}{2\pi}(-\sin^2\beta, \sin\beta\cos\beta + \beta)$. These are linearly independent for $\beta \neq 0$. If the motion of a is reversed, the resulting motion of ν is reversed. When $a \neq m$, at least one of these four motions must be lengthening. If $a = m$, then moving m away from ν is a lengthening motion. Thus, S was not a maximizing configuration. Contradiction. Therefore, in any maximizing configuration, all clients must lie on the unit circle. □

Let l be the line through ν perpendicular to $\overline{\nu m}$. Lemmas 2 and 3 address points on either side of l, first those opposite m, then those on the same side as m.

Lemma 2 *In any maximizing configuration, all chords on the convex hull of S must have normals whose projections onto $\overline{\nu m}$ point toward m.*

Proof sketch. By contradiction. Assume S is a maximizing configuration of points such that two neighbouring points, b and c, lie on the convex hull and the projection of the normal of $\overline{bc}$ onto $\overline{\nu m}$ point away from m. By Lemma 1, b and c must lie on the unit circle. Add a point a at the midpoint of $\overline{bc}$. See Figure 3b. When a

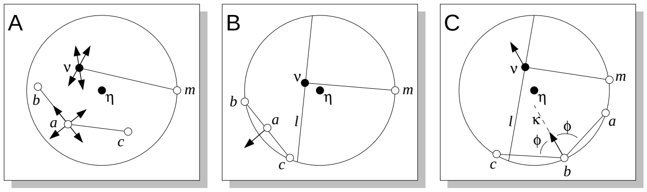

Figure 3: Lemmas 1, 2, and 3

moves along the normal to $\overline{bc}$ away from η, the weight of b and c decreases by some amount ϵ and the weight of a increases from 0 to ϵ. The normal of $\overline{bc}$ projected onto $\overline{\nu m}$ points away from m. Thus, the motion of a is lengthening. Therefore, S was not a maximizing configuration. Contradiction. Therefore, all chords on the convex hull of S must have normals whose projections onto $\overline{\nu m}$ point toward m. □

Corollary 1 *In any maximizing configuration, any points opposite l from m lie on a single continuous arc on the unit circle.*

Proof sketch. By Lemma 2, otherwise any gap would form a chord whose normal points away from m. □

Lemma 3 *In any maximizing configuration, only a single point may lie on the same side of l as m.*

Proof sketch. By contradiction. Assume S is a maximizing configuration such that two or more points lie on the convex hull on the same side of l as m. Let a and b be any two such neighbouring points such that $b \neq m$. By Lemma 1, a and b lie on the unit circle. See Figure 3c. Let c be b's neighbour opposite a. Let κ be the bisector of $\angle abc$. When b moves along κ, ν moves in the same direction. By Lemma 2, chords $\overline{ab}$ and $\overline{bc}$ must have normals whose projections onto $\overline{m\nu}$ point toward m. The orientation of κ lies between the reflections of these two normals. The projection of κ onto $\overline{m\nu}$ must point away from m. Thus, the motion of a is lengthening, meaning S was not maximized. We derive a contradiction. Therefore, in any maximizing configuration, the only point that may lie on the same side of l as m is m itself. □

By Lemmas 1, 2, and 3, a maximizing configuration consists of a continuous arc opposite at most one point. See Figure 2.

5 Rotated Bounding-Box Centre

5.1 Bounding Box Centres

The Gaussian centre assigns weights to extreme points. The simplest measure of extreme points is the bounding box, defined relative to some axis of orientation. The centre of the bounding box approximates the Euclidean

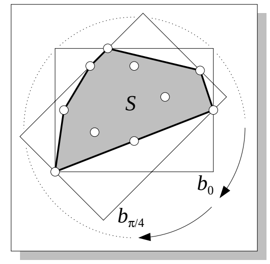

Figure 4: a continuously-rotated bounding box

centre with $\lambda = (1+\sqrt{2})/2$ and $v_{\max} = \sqrt{2}$ [BBKS00]. To improve the approximation, one could average the centres of several bounding boxes, each rotated by some angle $\phi \in [0, \pi/2)$.

Let S be a set of mobile clients within some bounded region of $\mathbb{R}^2$. For a given angle $\phi \in [0, \pi/2)$, let $b_\phi(t)$ be the bounding box of S at time t such that for each edge $\{e_1, \ldots, e_4\}$ of $b_\phi(t)$, the angle formed between e_i and the x-axis is $\phi \mod \frac{\pi}{2}$. Let $c(b_\phi(t))$ be the centre of bounding box $b_\phi(t)$. When $\phi = 0$, $b_0(t)$ is equivalent to the mobile ℓ_∞ centre. Similarly, when $\phi = \pi/4$, $b_{\pi/4}(t)$ is equivalent to the mobile ℓ_1 centre. As the number of bounding boxes approaches infinity, over all $\phi \in [0, \pi/2)$, we get the *rotated bounding-box centre*:

Definition 6 *Let S be a set of mobile clients within some bounded region of $\mathbb{R}^2$. The* rotated bounding-box centre *of S is*

$$\frac{2}{\pi}\int_0^{\pi/2} c(b_\phi(t))\ d\phi\ . \qquad (8)$$

Note, this definition could also be defined over the domains $[0, \pi]$ or $[0, 2\pi]$. The domain $[0, \pi/2]$ is sufficient since any two bounding boxes $b_{\phi_1}(t)$ and $b_{\phi_2}(t)$ are equal if $\phi_1 = \phi_2 \mod \pi/2$.

5.2 Velocity of Bounding-Box Centre

For any given bounding box, $b_\phi(t)$, the magnitude of the velocity of its centre is bounded by $\sqrt{2}$. This maximum is achieved along the diagonal. See Figure 5a. For a fixed direction of motion, say the x-axis, bounding box $b_{\phi-\pi/4}(t)$ contributes $\sqrt{2}\cos\phi$. See Figure 5b. Thus, the maximum velocity of the rotated bounding-box centre is

$$\frac{2}{\pi}\int_{-\pi/4}^{\pi/4} \sqrt{2}\cos\phi\ d\phi = \frac{4}{\pi} \approx 1.2732\ . \qquad (9)$$

Interestingly, for any set of mobile clients, the Gaussian centre is equal to the rotated bounding-box centre. The proof will be available in the full paper. This equivalence allows us to bound the velocity of the Gaussian centre.

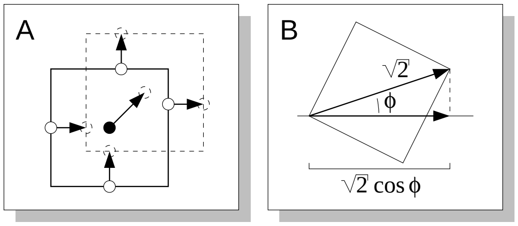

Figure 5: maximum velocity of a bounding box centre

Table 1: Approximations to the Euclidean Centre

Approximation	λ	$v_{\max}$
Euclidean centre	1	∞
point on convex hull	2	1
centre of mass	2	1
bounding box	$\frac{1+\sqrt{2}}{2} \approx 1.2071$	$\sqrt{2} \approx 1.4142$
Gaussian centre	≈ 1.1153	$\frac{4}{\pi} \approx 1.2732$

6 Other Approximations

How does the Gaussian centre compare to other approximations for the Euclidean centre? Bespamyatnikh et al. [BBKS00] examine a point on the convex hull, the centre of mass, and the bounding box. Furthermore, they explore a mixed strategy consisting of a linear combination of the bounding box and the centre of mass. See Table 1 for λ and $v_{\max}$ values.

Our discussion has focused on the ℓ_2 distance metric. Under ℓ_1 the velocity of the centre is bounded by $\sqrt{2}$ and under ℓ_∞ it is bounded by 1. Unlike ℓ_2, where the velocity of the centre is unbounded, approximation is not required since the exact centre can be followed. Furthermore, under ℓ_2, the position of the centre is unique whereas under ℓ_1 and ℓ_∞, even simple examples give rise to multiple centre points.

7 Future Work: 3 Dimensions

The Gaussian centre definition has a natural extension into three dimensions.

Let $S = \{v_1, \ldots, v_n\}$ be a set of clients. Let $p_i \in \mathbb{R}^3$ be the position of client v_i. Let $C_S \subseteq S$ be the three-dimensional convex hull of S. Let F_i be the set of faces that meet at vertex $v_i \in C_S$. Let $\theta_{i,j}$ be the interior plane angle on face f_j at vertex v_i. Let $w_i = 2\pi - \sum_{f_j \in F_i} \theta_{i,j}$ be the weight of client p_i. For each $v_i \in S - C_S$, let $w_i = 0$.

The sum of the plane angles at a vertex v ranges from 0 to 2π. By Euler's theorem, the sum of all weights for any arrangement of points is 4π.

Definition 7 *In three dimensions, the* Gaussian centre *of S, ν, is the normalized weighted mean of the clients of S:*

$$\nu(t) = \frac{1}{4\pi} \sum_i p_i(t) w_i(t) \ . \qquad (10)$$

When vertices are coplanar, a client set S in three dimensions reduces to the two-dimensional case. This is reflected by the fact that such an instance induces two symmetric faces. The Gaussian centre is given by

$$\nu = \frac{1}{4\pi} \sum_i (2\pi - 2\theta_i) = \frac{1}{2\pi} \sum_i (\pi - \theta_i) \ . \qquad (11)$$

As was the case in two dimensions, the Gaussian diagram provides an interesting alternative interpretation of vertex weights. In three dimensions, the Gaussian diagram G_S of a polyhedron C_S divides the unit sphere into spherical sectors such that the weight of each client in C_S is given by the surface area of its corresponding spherical polygon in G_S.

References

[AGG02] Pankaj K. Agarwal, Jie Gao, and Leonidas J. Guibas. Kinetic medians and *kd*-trees. In *10th Annual European Symposium on Algorithms*, 2002.

[AH01] Pankaj K. Agarwal and Sariel Har-Peled. Maintaining approximate extent of moving points. In *Symposium on Discrete Algorithms*, pages 148–157. ACM Press, 2001.

[AST93] Pankaj 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.

[BBKS00] Sergei Bespamyatnikh, Binay Bhattacharya, David Kirkpatrick, and Michael Segal. Mobile facility location. In *Fourth International ACM Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, pages 46–53, 2000.

[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.

[Meg83] Nimrod Megiddo. Linear-time algorithms for linear programming in r^3 and related problems. *SIAM J. Comput.*, 4:759–776, 1983.

[Web22] Alfred Weber. *Über den Standort der Industrie.* J. C. B. Mohr (Paul Siebeck), Tübingen, 1922.

[Wes73] George O. Wesolowsky. Dynamic facility location. *Management Science*, 19:1241–1248, 1973.