Kinetic Maintenance of Mobile k -Centres on Trees

Stephane Durocher*

School of Computer Science, McGill University
Montréal, Canada

Christophe Paul

CNRS & School of Computer Science, McGill University
Montréal, Canada

& Laboratoire d'Informatique de Robotique et de Microélectronique de Montpellier
Montpellier, France

LIRMM Research Report 06057

December 4, 2006

Abstract

Given a set P of points (clients) on a weighted tree T , the k -centre of P corresponds to a set of k points (facilities) on T such that the maximum graph distance between any client and its nearest facility is minimized. We consider the *mobile k -centre* problem on trees. Let C denote a set of n mobile clients, each of which follows a continuous trajectory on T . We establish tight bounds on the maximum relative velocity of the 1-centre and 2-centre of C . When each client in C moves with linear motion along a path on T , the motions of the corresponding 1-centre and 2-centre are piecewise linear; we derive a tight combinatorial bound of $\Theta(n)$ on the complexity of the motion of the 1-centre and corresponding bounds of $O(n^2\alpha(n))$ and $\Omega(n^2)$ for the 2-centre, where $\alpha(n)$ denotes the inverse Ackermann function. We describe efficient algorithms for calculating the trajectories of the 1-centre and 2-centre of C : the 1-centre can be found in optimal time $O(n \log n)$ when the distance function between mobile clients is known or $O(n^2)$ when the function must be calculated, and a 2-centre can be found in time $O(n^2 \log n)$. These algorithms lend themselves to implementation within the framework of kinetic data structures, resulting in structures that are compact, efficient, responsive, and local.

Keywords. facility location, tree, mobile, centre, 2-centre, continuous, velocity, kinetic data structure

*Research for this work was supported by NSERC.

1 Introduction

Motivation. Finding a set of k points that are central to a collection of data points drawn from a metric space is a fundamental problem of geometry and data analysis. Within the context of facility location, this problem is commonly known as the k -centre problem; given a set of points (clients) in a metric space S , the k -centre of P is a set of k points (facilities) such that the maximum distance from any client to its nearest facility is minimized. Two common choices for S are a Minkowski distance (typically ℓ_1 , ℓ_2 , or ℓ_∞) in Euclidean space and graph distance on a weighted graph.

Recently, the k -centre problem has been explored under mobility. In one dimension, the mobile 1-centre problem reduces to maintaining the extrema of a set of mobile clients as these move along the real line [3, 4, 8, 20]. Natural generalizations of the mobile 1-centre to higher dimensions in $\mathbb{R}^d$ lead to the mobile Euclidean 1-centre [4, 11, 16], the mobile rectilinear 1-centre [4, 12], and the kinetic convex hull [8, 9, 20]. In this paper, we consider a different generalization of the one-dimensional mobile problem to the metric space of graph distance on a weighted graph and, in particular, on a weighted tree.

Although the static k -centre problem on graphs is well understood (see Sec. 3), the corresponding mobile problem remained unexplored. Any path in a weighted graph is isometric to a line segment; we generalize the motion of a single client on the line to motion on a path in a graph. The union of the trajectories of a set of clients forms a graph. That is, given a weighted graph G , each mobile client follows a continuous trajectory along the edges and vertices of G . Continuity and bounded velocity are natural constraints on any physical moving object. It is straightforward to show that for any graph G that contains a cycle, there exist sets of mobile clients on G whose 1-centre is discontinuous. We restrict our attention to metric spaces for which the k -centre is continuous. In particular, graph distance on a tree maintains many properties of Minkowski distance on the real line, such as a unique shortest path between two points and a unique, continuous 1-centre (see Sec. 4.1), while introducing interesting algorithmic challenges to the problem of maintaining a mobile k -centre. Many of the properties we examine are familiar to geometry: continuity of motion, bounded velocity, extremal points, reflection, etc. As such, these questions share common aspects with both kinetic/mobile problems within computational geometry and problems in facility location on graphs.

Main Results. The 1-centre on a tree is unique [21]. We show its motion is continuous and has relative velocity at most one when the motion of clients is continuous. Since a 2-centre of a tree is not unique, we identify a particular 2-centre which we call the *equidistant 2-centre* and show that its motion is continuous and has relative velocity at most two when the motion of clients is continuous. The 3-centre is discontinuous even on a line segment; furthermore, no bounded-velocity approximation is possible for the mobile 3-centre [15]. We consider values of k for which the mobile k -centre is continuous: $k \leq 2$.

When each client in C moves with linear motion along a path on T , the motions of the corresponding 1-centre and equidistant 2-centre are piecewise linear. We derive a tight combinatorial bound of $\Theta(n)$ on the complexity of the motion of the 1-centre, an upper bound of $O(n^2\alpha(n))$ on the complexity of the motion of the equidistant 2-centre, and a worst-case lower bound of $\Omega(n^2)$ on the complexity of the motion of any 2-centre, where $\alpha(n)$ denotes the inverse Ackermann function. We describe efficient algorithms for calculating the trajectories of the 1-centre and 2-centre of C . When the all-pairs distance function between mobile clients is known at all times, the 1-centre can be found in optimal time $O(n \log n)$. The distance function can be calculated in time $O(n^2)$. The 2-centre can be found in time $O(n^2 \log n)$. Our algorithms have natural implementations as kinetic data structures (KDS), introduced by Basch et al. [8]. Applications of KDSs include collision detection (e.g., [1, 7, 30]), proximity problems (e.g., [10]), extent and extremal elements (e.g., [3, 4, 8, 20]), the convex hull (e.g., [8, 9, 20]), and the k -centre and k -median (e.g., [2, 4, 11, 15, 16, 19, 27]).

2 Definitions

Since a *point* refers to a fixed position in a metric space, we refer to a *client* in the context of motion. Let $C = \{c_1, \dots, c_n\}$ denote a set of mobile clients, where $I = [0, t_f]$ denotes a time interval, U_T denotes the continuum of points defined by a weighted tree $T = (V, E)$, and each c_i is a continuous function $c_i : I \rightarrow U_T$.

For every $t \in I$, let $C(t) = \{c(t) \mid c \in C\}$ denote the set of points in U_T that corresponds to the positions of clients in C at time t . The position of a mobile *facility* f is a function of the positions of a set of clients, $f : \mathcal{P}(U_T) \rightarrow U_T$.

A common assumption in problems involving motion in Euclidean space is that the position of a mobile client is a linear function over time (e.g., [3, 4, 8]). We make a similar assumption and consider clients with linear motion on trees to establish combinatorial bounds. As we discuss in Sec. 4.3, our algorithms generalize to allow algebraic motion (although the derived bounds no longer apply). A mobile client or facility a has *linear motion* if for all $t \in I$, $d(a(0), a(t)) = t \cdot v_a$, where v_a is a non-negative constant and $d(b, c)$ denotes the graph distance between points b and c in U_T . We refer to v_a as the *velocity* of a . That is, a follows a continuous trajectory along the path on T between $a(0)$ and $a(t_f)$ with velocity v_a . The union of the trajectories of a set of n mobile clients that move with linear motion is a subgraph of U_T that has at most $2n$ vertices of degree one. Therefore, we assume that T has at most $2n$ leaves and at most $4n - 1$ vertices, and that $c(0)$ and $c(t_f)$ are vertices of T , for each $c \in C$.

We assume an upper bound of one on the velocity of clients since we are interested in *relative velocity*. Unlike mobile clients, a mobile facility is not required to travel along a path in T nor is its velocity required to remain constant. A mobile facility f has *maximum velocity* v_f if

$$\forall t_1, t_2 \in I, d(f(C(t_1)), f(C(t_2))) \leq v_f |t_1 - t_2|, \quad (1)$$

for all sets of mobile clients C defined on any tree T and any time interval I . Observe that continuity is a necessary condition for any fixed upper bound on velocity. Similarly, we say the *rate of change* of the k -radius is bounded by r_f if

$$\forall t_1, t_2 \in I, |r(C(t_1)) - r(C(t_2))| \leq r_f |t_1 - t_2|, \quad (2)$$

for all sets of mobile clients C defined on any tree T and any time interval I .

We say that two clients a and b *cross* at time t_0 if

$$a(t_0) = b(t_0) \text{ and } \exists \epsilon > 0 \text{ s.t. } \forall t \in (t_0 - \epsilon, t_0), a(t) \neq b(t).$$

In most cases, clients a and b coincide only at the instant t_0 . However, if a and b have the same velocity, then their trajectories may merge such that the positions of a and b coincide until their trajectories diverge again. Since clients a and b have constant velocity and their trajectories intersect in a path, a and b may cross at most once. We define the crossing event as the instant t_0 when their two positions first coincide.

We say client $c \in C$ is *extreme* at time t if c does not lie in the interior of any path through T between two clients in $C(t)$. The *convex hull* of $C(t)$ corresponds to the union of all paths between two clients in $C(t)$. Whereas some definitions of the convex hull on a graph refers to a subset of the vertices [13], we refer to the continuous subset of U_T .

We recall the (static) definition of a k -centre of a client set on a tree.

Definition 1. *Given a weighted tree T and a set of points C in U_T , a k -centre of C is a set of k points in U_T , denoted $\Xi_1(C), \dots, \Xi_k(C)$, that minimizes*

$$\max_{c \in C} \min_{1 \leq i \leq k} d(c, \Xi_i(C)). \quad (3)$$

When $k = 1$, we omit the subscript and write $\Xi(C)$. The definition of a mobile k -centre of a set of mobile clients C follows directly from this static definition. That is, the instantaneous positions of a mobile k -centre of C at time t is given by Definition 1 in terms of $C(t)$.

We refer to the value of (3) as the k -radius of C or simply as its *radius* when $k = 1$. The diameter of C is twice the radius of C [22] (for graphs, the diameter is at most twice the radius). A *diametric path* of C is a path between two clients c_1 and c_2 in C such that the distance between them is the diameter of C . We refer to $\{c_1, c_2\}$ as a *diametric pair* and to c_1 and c_2 as *diametric clients*. The 1-centre of C is the unique midpoint of all diametric paths of C [21].

The 1-centre problem on graphs is also known as the absolute centre [21, 22, 23], single centre [22], and minimax location problem [14, 21]. A common variation of the k -centre problem on graphs is known as the vertex k -centre or discrete k -centre problem, for which the choice of locations for the facility is restricted to vertices (similarly, clients) of the graph G . Observe that maintaining continuity in the motion of a mobile facility is impossible in the vertex centre model, as a facility could be required to jump discontinuously from vertex to vertex (similarly, client to client).

3 Related Work

Handler [21] gives linear-time algorithms for identifying the 1-centre and 2-centre of a tree. Frederickson gives a linear-time algorithm for finding a k -centre of a tree [18]. Kariv and Hakimi [28] provide an $O(mn + n \log^2 n)$ -time algorithm for the 1-centre problem on graphs, where $n = |V|$ and $m = |E|$. A review of single-facility location problems on graphs is given in [24]. As for multiple-facility location problems, [28, 29], and [31] provide reviews of the k -centre and k -median problems on graphs, while [32] reviews these problems on trees.

Kinetic data structures (KDS), introduced by Basch et al. [8], allow the maintenance of an attribute (called the configuration function) of a set of mobile objects moving continuously in some metric space. To do so, a KDS maintains a dynamic set of certificates that guarantee the correctness of the configuration function at any time during the motion. Each certificate c is associated with a small set of mobile objects for which some property is verified. The failure time of certificate c (called an event) is calculated as a function of the motion of these objects. The failure time is added to a priority queue. Restoring the configuration function following a certificate failure requires updating the set of certificates (and the corresponding events in the queue). The *compactness* of a KDS measures the maximum number of certificates active at any given time. The *responsivity* of a KDS measures the maximum number of certificates associated with any one mobile object. The *locality* of a KDS measures the number of certificates that require updating as a result of a certificate failure. The *efficiency* of a KDS compares the total number of certificate failures relative to the number of external events (changes to the configuration function). See [6, 8, 9, 20] for a more complete description of the KDS framework.

In relation to our work on the mobile k -centre, KDSs have been constructed to maintain various attributes of a set of mobile clients; these include extremal elements in $\mathbb{R}$ [3, 4, 8, 20], the extent and approximate extent (e.g., diameter and width) in $\mathbb{R}^2$ [3, 4], approximations to the mobile 1-centre in $\mathbb{R}^2$ [4, 11, 15, 16], approximations to the mobile 2-centre in $\mathbb{R}^2$ [15], the mobile rectilinear 1-centre in $\mathbb{R}^2$ [4, 12], the kinetic convex hull [8, 9, 20], and the approximate discrete rectilinear k -centre [19, 27].

In any metric space, identifying a pair of furthest clients in a set of mobile clients corresponds to finding the upper envelope (the maximum function) of a set of distance functions. This problem is related to Davenport-Schinzel sequences [5, 17, 26, 25, 33]. In particular, the upper (lower) envelope of a set of n line segments is a piecewise-linear function that consists of $\Theta(n\alpha(n))$ linear segments [25] in the worst case. Hershberger [26] provides an algorithm for computing the upper envelope in optimal $O(n \log n)$ time.

4 The Mobile 1-Centre on Trees

4.1 Properties of the Mobile 1-Centre

The mobile 1-centre is continuous in $\mathbb{R}$ and $\mathbb{R}^2$ [15]. Although the mobile 1-centre has at most unit relative velocity in $\mathbb{R}$, its relative velocity is unbounded in $\mathbb{R}^2$ [12]. As we show, the mobile 1-centre remains continuous on trees (but not on graphs). Furthermore, the mobile 1-centre has at most unit relative velocity on trees.

Theorem 1. *The mobile 1-centre has relative velocity at most one on trees. This bound is tight.*

Proof. Choose any $t_1, t_2 \in I$ and let $\delta = |t_1 - t_2|$. If $\Xi(C(t_1)) = \Xi(C(t_2))$, then the velocity bound holds trivially. Therefore, assume $\Xi(C(t_1)) \neq \Xi(C(t_2))$. Let P denote the path in T between $\Xi(C(t_1))$ and

$\Xi(C(t_2))$. Let r_1 and r_2 denote the respective radii of $C(t_1)$ and $C(t_2)$. Let L_1 denote the subtree of T that includes all branches of $\Xi(C(t_1))$ except P . Note, L_1 includes $\Xi(C(t_1))$. Similarly, let L_2 denote the subtree of T that includes all branches of $\Xi(C(t_2))$ except P .

Let a be a client in C such that $a(t_1) \in L_1$ and $d(a(t_1), \Xi(C(t_1))) = r_1$. Similarly, let b be a client in C such that $b(t_2) \in L_2$ and $d(b(t_2), \Xi(C(t_2))) = r_2$. Such clients must exist since $\Xi(C(t))$ is the midpoint of a diametric path of $C(t)$ for all t . See Fig. 1 below.

Figure 1: Relative positions of $a(t_1)$, $b(t_2)$, $\Xi(C(t_1))$ and $\Xi(C(t_2))$

Therefore,

$$d(a(t_1), b(t_2)) \leq d(a(t_1), \Xi(C(t_1))) + d(\Xi(C(t_1)), b(t_2)) + d(b(t_1), b(t_2)) \leq 2r_1 + \delta \quad (4a)$$

$$\text{and } d(a(t_1), b(t_2)) \leq d(a(t_1), a(t_2)) + d(a(t_2), \Xi(C(t_2))) + d(\Xi(C(t_2)), b(t_2)) \leq 2r_2 + \delta. \quad (4b)$$

Consequently,

$$\begin{aligned} d(a(t_1), b(t_2)) &= d(a(t_1), \Xi(C(t_1))) + d(\Xi(C(t_1)), \Xi(C(t_2))) + d(\Xi(C(t_2)), b(t_2)), \\ \Rightarrow d(\Xi(C(t_1)), \Xi(C(t_2))) &= d(a(t_1), b(t_2)) - d(a(t_1), \Xi(C(t_1))) - d(\Xi(C(t_2)), b(t_2)) \end{aligned}$$

$$\begin{aligned} d(\Xi(C(t_1)), \Xi(C(t_2))) &= d(a(t_1), b(t_2)) - r_1 - r_2 \\ &\leq \delta, \end{aligned}$$

by (4a) and (4b). Therefore, (1) holds for Ξ when $v_f = 1$. The bound is realized when the two diametric clients move in a parallel direction. $\square$

Corollary 2. *The mobile 1-centre is continuous on trees.*

Corollary 3. *The relative rate of change of the radius is at most one on trees.*

We refer to the following lemma by Handler:

Lemma 4 (Handler 1973 [21]). *Given a set of clients C on a tree T , clients $a, b \in C$ are a diametric pair of C if and only if $d(a, b) \geq \max\{d(a, c), d(b, c)\}$ for all $c \in C$.*

4.2 Complexity of the Motion of the 1-Centre

When n clients move along the real line, each with some constant velocity, the identity of the client that realizes either extremum changes $\Theta(n)$ times in the worst case [8]. In particular, any given client realizes each extremum at most once in the sequence of changes. When n clients move in $\mathbb{R}^2$ along linear trajectories with constant velocity, the diametric pair of clients changes $\Omega(n^2)$ times in the worst case [3]. As we show in Theorem 8, for a set C of n clients with linear motion on a tree T , the identity of the diametric pair of C changes $\Theta(n)$ times in the worst case. We begin with a definition.

Given a client c moving with velocity v_c , the *outward velocity* of c at time t , denoted $\vec{v}(c(t))$, is given by

$$\vec{v}(c(t)) = \begin{cases} -\infty & \text{if } c(t) \text{ is not extreme in } C(t), \\ -v_c & \text{if } c(t) \text{ moves towards the interior of the convex hull of } C(t), \\ v_c & \text{otherwise.} \end{cases} \quad (5)$$

Lemmas 5 through 7 assume linear motion of a set of clients C on a tree T . In addition, we assume that the diameter of C is non-zero at all times; a zero diameter implies that all clients in C coincide in a point and any two clients define a diametric pair. Furthermore, the interior of the convex hull is empty and, consequently, outward velocity is ill defined. We consider a zero diameter in the proof of Theorem 8.

Lemma 5. *The outward velocity of client $c \in C$ is non-decreasing while c remains in a diametric pair of C .*

Proof. Two cases are possible while c remains in a diametric pair of C .

Case 1. Assume c moves away from the interior of the convex hull of C initially. Client c has linear motion along a path $P \subseteq T$. The subpath of P that remains to be travelled by c lies outside the convex hull of C . Therefore the outward velocity of c remains constant.

Case 2. Assume c moves towards the interior of the convex hull of C initially. The outward velocity of c remains constant until c branches and turns away from the interior of the convex hull. The remainder of the motion corresponds to Case 1. $\square$

As we show in Lemma 6, any change in the outward velocity at either endpoint of a diametric path must be increasing.

Lemma 6. *Choose any $t_1 \in I$ and let $\{a_1, b_1\}$ be a diametric pair of $C(t_1)$. If $\{a_2, b_2\}$ is a diametric pair of $C(t_2)$ and a_1 is not in any diametric pair of $C(t_2)$ for some $\epsilon > 0$ and all $t_2 \in (t_1, t_1 + \epsilon)$, then $\vec{v}(a_1(t_1)) < \min\{\vec{v}(a_2(t_2)), \vec{v}(b_2(t_2))\}$.*

Proof. Since a_1 is in a diametric pair of $C(t_1)$,

$$\forall c \in C, d(a_1(t_1), b_1(t_1)) \geq d(a_1(t_1), c(t_1)). \quad (6)$$

Since a_2 and b_2 are a diametric pair of $C(t_2)$ but a_1 is not in any diametric pair of $C(t_2)$,

$$\forall c \in C, d(a_1(t_2), c(t_2)) < d(a_2(t_2), b_2(t_2)). \quad (7)$$

Since client motion is continuous and by (6) and (7),

$$d(a_1(t_1), b_1(t_1)) = d(a_2(t_1), b_2(t_1)). \quad (8)$$

The result follows from (7) and (8). $\square$

Lemma 7. *A client $c \in C$ becomes an endpoint of a diametric path of C at most four times.*

Proof. By (5), the outward velocity of a client c in a diametric pair (c is extreme) is one of two values: $\pm v_c$. By Lemma 6, a change in a diametric pair corresponds to an increase in outward velocity. Therefore, a client c realizes either endpoint of a diametric path at most twice, for a total of at most four times. $\square$

Theorem 8. *When each client in C moves with linear motion along a path on T , the motion of the 1-centre of C is piecewise linear and is composed of $\Theta(n)$ linear segments in the worst case, where $n = |C|$.*

Proof. Case 1. Assume the diameter of C is non-zero throughout the motion. The upper bound $O(n)$ follows from Lemmas 5 and 7 and the fact that the 1-centre of C is the midpoint of a diametric pair.

Case 2. Assume the diameter of C is zero at some time during the motion. A zero diameter implies that all clients in C coincide at a point; that is, all clients cross simultaneously. This degeneracy occurs at most once since any two clients cross at most once. Since clients in C have linear motion, the motion of the 1-centre of C has linear motion while all clients coincide. Before and after the degeneracy, the motion of clients in C corresponds to Case 1. Therefore, the sum of the number of linear segments of the motion of the 1-centre remains $O(n)$.

The worst-case lower bound of $\Omega(n)$ follows from the corresponding result in one dimension [8]. $\square$

4.3 Kinetic Maintenance of the Mobile 1-Centre

Given a set C of n mobile clients, each moving with linear motion in $\mathbb{R}$, the 1-centre of C is the midpoint of the extrema of C . The position of each extremum is given by the upper (respectively, lower) envelope of the set of n linear functions that correspond to the positions of clients in C relative to a fixed point in $\mathbb{R}$. Hershberger [26] gives an $O(n \log n)$ time algorithm which finds the upper envelope by dividing the set of linear functions in two, recursively finding the upper envelope of each set, and recombining the two envelopes to give the global upper envelope.

Using a related idea, we describe an algorithm for identifying a sequence of diametric pairs of a set of mobile clients, each moving with linear motion on a tree. We then describe how to implement the algorithm as a KDS. The algorithm makes use of the distance function d , where $d(a(t), b(t))$ returns the graph distance on tree T between mobile clients a and b at time t . We begin with the following lemma upon which our algorithm relies.

Lemma 9. *Let C_1 and C_2 be sets of points on U_T for some tree T . Let a_i and b_i denote a diametric pair of C_i , for $i = 1, 2$. Points e and f are a diametric pair of $C_1 \cup C_2$, where*

$$\{e, f\} = \underset{\{e', f'\} \subseteq \{a_1, b_1, a_2, b_2\}}{\operatorname{argmax}} d(e', f'). \quad (9)$$

Proof. By Lemma 4 we know

$$\forall c \in C_1, \max\{d(c, a_1), d(c, b_1)\} \leq d(a_1, b_1) \text{ and } \forall c \in C_2, \max\{d(c, a_2), d(c, b_2)\} \leq d(a_2, b_2). \quad (10)$$

Similarly, by Lemma 4 it suffices to show

$$\forall c \in C_1 \cup C_2, \max\{d(c, e), d(c, f)\} \leq d(e, f).$$

Without loss of generality, choose any c in C_1 . Let T' denote the spanning tree of a_1, b_1, a_2 , and b_2 in T . Let q denote the vertex of T' that is closest to c . By (10) we know

$$\begin{aligned} & d(a_1, c) \leq d(a_1, b_1), \\ \Rightarrow & d(a_1, q) + d(q, c) \leq d(a_1, q) + d(q, b_1), \\ \Rightarrow & d(q, c) \leq d(q, b_1). \end{aligned}$$

Similarly, $d(q, c) \leq d(q, a_1)$. It follows that

$$\begin{aligned} \max\{d(c, e), d(c, f)\} &= d(c, q) + \max\{d(q, e), d(q, f)\} \\ &\leq \min\{d(q, a_1), d(q, b_1)\} + \max\{d(q, e), d(q, f)\}, \\ &\leq d(e, f), \end{aligned}$$

because q lies on a path in T between e and a or b , and on a path in T between f and b or a . $\square$

Algorithm Description. The set of mobile clients C is partitioned arbitrarily into sets C_1 and C_2 of size $\lfloor n/2 \rfloor$ and $\lceil n/2 \rceil$. For each $i = 1, 2$, the algorithm is called recursively to find a sequence of diametric pairs of C_i , denoted $\{a_{i,1}, b_{i,1}\}, \dots, \{a_{i,m_i}, b_{i,m_i}\}$, and a corresponding partition of the time interval I , denoted $I_{i,1}, \dots, I_{i,m_i}$, such that for each j , $a_{i,j}(t)$ and $b_{i,j}(t)$ are a diametric pair of $C_i(t)$ for all $t \in I_{i,j}$. The recursion terminates when $n \leq 2$, in which case each client in C is in a diametric pair. We now describe how to compute a corresponding sequence for C .

Consider a third partition of the time interval I , denoted $I_1, \dots, I_m$, such that for each i , $I_i = I_{1,j} \cap I_{2,k}$, for some j, k . For all $t \in I_i$, diametric pairs of $C_1(t)$ and $C_2(t)$ consist of four clients in C , say a_1, b_1, a_2 , and b_2 . Let e and f be defined as in (9). By Lemma 9, e and f are a diametric pair of $C(t)$. The sequence of pairs of clients in $\{a_1, b_1, a_2, b_2\}$ that realize e and f corresponds to the sequence of pairs whose relative distance is maximized. That is, there are six combinations of pairs in $\{a_1, b_1, a_2, b_2\}$, each of which

corresponds to an inter-client distance function. The upper envelope of these six functions determines the sequence of identities of e and f during I_i . Thus, solutions to the recursive subproblems are combined to find the sequence of diametric pairs of C .

Time Complexity. By Theorem 8, the complexity of the motion of the 1-centres of C_1 and C_2 is $O(n)$. That is, the time interval I can be partitioned into $O(n)$ subintervals such that the motion of each 1-centre is linear within every subinterval (i.e., $m \in O(n)$). Within each subinterval, we find the maximum of six piecewise-linear functions, each composed of at most four linear segments. Therefore, the maximum function is also piecewise linear, consists of at most 24 linear segments, and can be found in constant time. Thus, the solutions to the two subproblems are combined in $O(n)$ time. The recursion tree has depth $\lfloor \log_2 n \rfloor$, resulting in a total time of $f(n) = 2f(n/2) + O(n)$ which is $O(n \log n)$. The worst-case lower bound of $\Omega(n \log n)$ follows from the corresponding one-dimensional problem [26].

Distance Function. Depending on the formulation of the problem, the input may not include the distance function. In this case, the input is given simply as the set of clients, each of which specifies origin and destination vertices in T . In particular, the path of a client's trajectory is not given.

We assume only a basic weighted edge adjacency list or matrix for the tree T . Build a table $A[i, j]$ that stores the following information for each vertex u_i and each client c_j : $d(u_i, c_j(0))$, the velocity of $c_j(0)$ relative to u_i , and the instant in I (if any) at which the velocity of c_j relative to u_i becomes negated (that is, c_j takes a branch such that its motion changes from towards u_i to away from u_i). This information encodes the two-segment piecewise-linear function $d(u_i, c_j(t))$. Table $A[i, j]$ has size $O(n^2)$ and can be calculated in time $O(n^2)$ by considering each client c_j and tracing its trajectory through T . Values for vertices in branches not followed by the trajectory of c_j are easily calculated recursively.

For any clients c_1 and c_2 in C , the client-to-client distance function $d(c_1(t), c_2(t))$ can be calculated in constant time from table A . While c_1 and c_2 move towards each other, $d(c_1(t), c_2(t)) = |d(c_1(0), c_1(t)) - d(c_1(0), c_2(t))|$. After one client, say c_1 , turns away from the other, $d(c_1(t), c_2(t)) = |d(c_1(t_f), c_1(t)) - d(c_1(t_f), c_2(t))|$. Recall that $c_1(0)$ and $c_1(t_f)$ are vertices of T .

KDS Implementation. We describe a KDS that maintains a diametric pair over time along with a set of certificates that validates the identity of the pair at any time during the motion.

The set of certificates corresponds to the recursive hierarchy described in our algorithm. At any time t , for each set C in the hierarchy, the certificate for $C(t)$ consists of five inequalities that confirm the maximum of six functions. That is, the certificate verifies the identity of a diametric pair of $C(t)$ in terms of the diametric pairs of the subsets $C_1(t)$ and $C_2(t)$ by Lemma 9. The corresponding properties are certified recursively for $C_1(t)$ and $C_2(t)$. Each set maintains a single certificate defined in terms of four clients and the total number of certificates is $O(n)$; therefore, the KDS is compact. Each client is contained in at most $O(\log n)$ sets and, consequently, is associated with at most $O(\log n)$ certificates. As a result, a motion plan update for a client results in changes to the failure times of $O(\log n)$ certificates; therefore, the KDS is local.

A certificate failure occurs whenever the diametric pair of a set C changes. Locally, the certificate for C is restored in constant time; however, a change in the diametric pair of C may percolate upwards in the tree, resulting in $O(\log n)$ additional certificate updates; therefore, the KDS is responsive. By Theorem 8, each set C contributes at most $O(|C|)$ certificate failures, resulting in a total of $O(n \log n)$ certificate failures over the entire motion. Although this value is asymptotically greater than $\Theta(n)$ (the worst-case number of external events for a set of n clients), any offline algorithm for finding the trajectory of the 1-centre requires $\Omega(n \log n)$ time in the worst case, even in one dimension. Therefore, the KDS is efficient.

Note, linear motion is not required by this KDS. In particular, the KDS applies to any algebraic motion for which the client-to-client distance function permits calculating the failure time of a certificate. In general, the combinatorial bounds and running times mentioned earlier do not apply to non-linear motion.

5 The Mobile 2-Centre on Trees

5.1 Properties of the Mobile 2-Centre

Although a 2-centre of a set of clients C on a tree is not unique (this is the case even in one dimension [15]), any 2-centre of C , $\Xi_1(C)$ and $\Xi_2(C)$, defines a natural bipartition of C , denoted $\{C_1, C_2\}$, such that

$$\forall c \in C_1, d(c, \Xi_1(C)) \leq d(c, \Xi_2(C)) \quad \text{and} \quad \forall c \in C_2, d(c, \Xi_1(C)) \geq d(c, \Xi_2(C)). \quad (11)$$

We refer to $\{C_1, C_2\}$ as a *diametric partition* of C and to C_1 and C_2 as *diametric subsets* of C . A diametric partition induced by a given 2-centre is not unique since any client that is equidistant from $\Xi_1(C)$ and $\Xi_2(C)$ may belong to either C_1 or C_2 . Since the 2-radius of C is at most the radius, it follows that there exists a diametric pair $\{a, b\}$ such that $a \in C_1$ and $b \in C_2$. As shown by Handler [22], following property is equivalent to (11):

$$\forall c \in C_1, d(c, a) \leq d(c, b) \quad \text{and} \quad \forall c \in C_2, d(c, a) \geq d(c, b). \quad (12)$$

We say $\{C_1, C_2\}$ is a diametric partition of C *induced* by $\{a, b\}$. We refer to the *local 1-centre*, *local radius*, and *local diametric pair/path*, respectively, in reference to the 1-centre, radius, and diametric pair/path of C_1 or C_2 . The local 1-centres of C_1 and C_2 are a 2-centre of C [22]. In the mobile setting, we say a diametric partition $\{C_1, C_2\}$ is *unchanged* during time interval I if and only if

$$\forall t_1, t_2 \in I, \forall c \in C, c(t_1) \in C_1(t_1) \Leftrightarrow c(t_2) \in C_1(t_2). \quad (13)$$

Recall that $C_1(t_1)$ denotes a set of points in U_T ; therefore, (13) does not imply $C_1(t_1) = C_1(t_2)$ for $t_1, t_2 \in I$.

We refer to the following lemma by Handler:

Lemma 10 (Handler 1978 [22]). *Any local diametric pair includes one diametric client in C .*

5.2 Equidistant 2-Centre

Even in one dimension the motion of a 2-centre defined by two local 1-centres is not continuous. This is easily demonstrated by an example: position a client at each endpoint of a line segment and let a third client move from one endpoint to the other. Not all 2-centres are discontinuous; by generalizing the one-dimensional 2-centre strategy described in [15], we obtain a strategy for defining the positions of a 2-centre on a tree whose motion is continuous and whose relative velocity is at most two. We refer to this particular 2-centre as the *equidistant 2-centre*:

Definition 2. *Let $\{a, b\}$ be a diametric pair of C . An equidistant 2-centre of C , denoted $\{\dot{\Xi}_1(C), \dot{\Xi}_2(C)\}$, is a pair of points that lie on the path between a and b at a distance ρ from a and b , respectively, where ρ denotes the 2-radius of C .*

See Fig. 2 for an example. As we now demonstrate, the equidistant 2-centre is independent of the choice of the diametric pair $\{a, b\}$.

Lemma 11. *The equidistant 2-centre is unique.*

Proof. If C has a unique diametric pair, then the equidistant 2-centre is also unique by Definition 2. Therefore, assume C has two or more diametric pairs. Choose any two diametric pairs, $\{a_1, b_1\}$ and $\{a_2, b_2\}$. Without loss of generality, assume

$$d(a_1, a_2) \leq d(a_1, b_2). \quad (14)$$

Let $\{\dot{\Xi}_1(C), \dot{\Xi}_2(C)\}$ denote the equidistant 2-centre defined in terms of $\{a_1, b_1\}$ and let $\{\ddot{\Xi}_1(C), \ddot{\Xi}_2(C)\}$ denote the equidistant 2-centre defined in terms of $\{a_2, b_2\}$. Without loss of generality, assume $d(a_1, \dot{\Xi}_1(C)) \leq d(a_1, \dot{\Xi}_2(C))$ and $d(a_2, \ddot{\Xi}_1(C)) \leq d(a_2, \ddot{\Xi}_2(C))$. Let ρ denote the 2-radius of C . By symmetry, it suffices to show $\dot{\Xi}_1(C) = \ddot{\Xi}_1(C)$.

Figure 2: **Equidistant 2-centre example.** In tree T_1 , $\{a, b\}$ is the unique diametric pair. In tree T_2 , any two clients form a diametric pair; the diametric path between e and f is displayed. The corresponding diametric subsets are A_1 and B_1 in T_1 (unique) and A_2 and B_2 in T_2 (not unique). In tree T_1 , the 2-radius is realized by the local diametric pair $\{b, d\}$. Consequently, the local 1-centre of B_1 coincides with the equidistant 2-centre in B_1 . In tree T_2 , the 2-radius is equal to the radius. Consequently, the local 1-centre of A_2 coincides with the 1-centre and with both equidistant 2-centres. The local 1-centre of B_2 coincides with client f .

If $a_1 = a_2$, then $\dot{\Xi}_1(C) = \ddot{\Xi}_1(C)$ by Definition 2. Therefore, assume $a_1 \neq a_2$. By (3) and the triangle inequality, $\min\{d(a_1, a_2), d(a_1, b_2)\} \leq 2\rho$. By (14), $d(a_1, a_2) \leq 2\rho$. Let v denote the vertex of T that joins the branches containing a_1 , a_2 , and $\Xi(C)$, respectively. Since a_1 and a_2 are diametric clients, $d(a_1, \Xi(C)) = d(a_2, \Xi(C))$ and, therefore, $d(a_1, v) = d(a_2, v)$. Consequently, $d(a_1, v) = d(a_2, v) \leq \rho$. The point that lies at a distance ρ from a_1 on the path between a_1 and $\Xi(C)$ coincides with the point that lies at a distance ρ from a_2 on the path between a_2 and $\Xi(C)$. That is, $\dot{\Xi}_1(C) = \ddot{\Xi}_1(C)$. $\square$

Corollary 12. $\dot{\Xi}_1(C)$ and $\dot{\Xi}_2(C)$ lie in the intersection of all diametric paths of C .

Lemma 13. The equidistant 2-centre of C is a 2-centre of C .

Proof. Choose any client $c \in C$. Let $\{a, b\}$ be a diametric pair of C . Let v denote the point in U_T that joins the branch containing c to the path between a and b (c may coincide with v). Let $\{C_1, C_2\}$ be a diametric partition of C induced by a and b such that $a \in C_1$. Without loss of generality, assume $c \in C_1$ and $d(a, \dot{\Xi}_1(C)) \leq d(a, \dot{\Xi}_2(C))$. Let ρ denote the 2-radius of C . By Corollary 12, $\dot{\Xi}_1(C)$ and v lie on the path between a and b . By Definition 2, $d(a, \dot{\Xi}_1(C)) = \rho$. If v lies between $\dot{\Xi}_1(C)$ and a , then $d(\dot{\Xi}_1(C), c) \leq d(\dot{\Xi}_1(C), a) = \rho$, otherwise a is not a diametric client. Therefore, assume $\dot{\Xi}_1(C)$ lies between a and v . Since $a, c \in C_1$, $d(a, c) = d(a, \dot{\Xi}_1(C)) + d(\dot{\Xi}_1(C), c) \leq 2\rho$ and, consequently, $d(\dot{\Xi}_1(C), c) \leq \rho$. $\square$

Lemma 14. The relative rate of change of the 2-radius is at most one on trees.

Proof. We show $|\rho(C(t_1)) - \rho(C(t_2))| \leq |t_1 - t_2|$, for any $t_1, t_2 \in I$, where $\rho(C(t_i))$ denotes the 2-radius of $C(t_i)$. Choose any $t_1, t_2 \in I$. Let $\delta = |t_1 - t_2|$. Let a_i and b_i be clients in C such that $\{a_i(t_i), b_i(t_i)\}$ is a diametric pair of $C(t_i)$, for $i = 1, 2$. Since a local 1-centre is the midpoint of a local diameter, the 2-radius of $C(t_i)$ can be expressed as,

$$\rho(C(t_i)) = \frac{1}{2} \max_{c \in C} \min\{d(c(t_i), a_i(t_i)), d(c(t_i), b_i(t_i))\}. \quad (15)$$

Since clients move with at most unit velocity,

$$\forall \{c, e\} \subseteq C, |d(c(t_1), e(t_1)) - d(c(t_2), e(t_2))| \leq 2\delta. \quad (16)$$

Let $\{A_2(t_2), B_2(t_2)\}$ denote the diametric partition of $C(t_2)$ induced by $\{a_2(t_2), b_2(t_2)\}$ such that $a_2(t_2) \in A_2(t_2)$ and $b_2(t_2) \in B_2(t_2)$.

Case 1. Without loss of generality, assume $a_1(t_2) \in A_2(t_2)$ and $b_1(t_2) \in B_2(t_2)$.

$$\begin{aligned}
\rho(C(t_1)) &= \frac{1}{2} \min \left\{ d(\tilde{c}(t_1), a_1(t_1)), d(\tilde{c}(t_1), b_1(t_1)) \right\}, & \text{where } \tilde{c} \in C \text{ maximizes (15) when } i = 1, \\
&\leq \frac{1}{2} \min \left\{ d(\tilde{c}(t_2), a_1(t_2)), d(\tilde{c}(t_2), b_1(t_2)) \right\} + \delta, & \text{by (16),} \\
&\leq \frac{1}{2} \max_{c \in C} \min \left\{ d(c(t_2), a_1(t_2)), d(c(t_2), b_1(t_2)) \right\} + \delta \\
&\leq \frac{1}{2} \max_{c \in C} \min \left\{ d(c(t_2), a_2(t_2)), d(c(t_2), b_2(t_2)) \right\} + \delta,
\end{aligned}$$

since $\{a_2(t_2), b_2(t_2)\}$ is a diametric pair of $C(t_2)$, $\{a_1(t_2), a_2(t_2)\} \subseteq A_2(t_2)$, and $\{b_1(t_2), b_2(t_2)\} \subseteq B_2(t_2)$,

$$= \rho(C(t_2)) + \delta, \quad \text{by (15).}$$

Case 2. Assume $a_1(t_2)$ and $b_1(t_2)$ are in the same diametric subset of $\{A_2(t_2), B_2(t_2)\}$.

$$\begin{aligned}
\rho(C(t_1)) &\leq r(t_1), & \text{where } r(t_1) \text{ denotes the radius of } C(t_1), \\
&= \frac{1}{2} d(a_1(t_1), b_1(t_1)), & \text{since } \{a_1(t_1), b_1(t_1)\} \text{ is a diametric pair of } C(t_1), \\
&\leq \frac{1}{2} d(a_1(t_2), b_1(t_2)) + \delta, & \text{by (16),} \\
&\leq \rho(C(t_2)) + \delta,
\end{aligned}$$

since $a_1(t_2)$ and $b_1(t_2)$ are in the same diametric subset of $C(t_2)$. □

Theorem 15. *Each facility in the mobile equidistant 2-centre has relative velocity at most two.*

Proof. Choose any $t_1, t_2 \in I$. By symmetry of $\dot{\Xi}_1(C)$ and $\dot{\Xi}_2(C)$, it suffices to show that

$$\min\{d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_1(C(t_2))), d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_2(C(t_2)))\} \leq 2\delta, \quad (17)$$

where $\delta = |t_1 - t_2|$. Let a_i and b_i be clients in C such that $\{a_i(t_i), b_i(t_i)\}$ is a diametric pair of $C(t_i)$, for $i = 1, 2$.

Case 1. Assume $\dot{\Xi}_1(C(t_1))$ lies between $\dot{\Xi}_1(C(t_2))$ and $\dot{\Xi}_2(C(t_2))$. By Corollary 12, $\dot{\Xi}_1(C(t_2))$ and $\dot{\Xi}_2(C(t_2))$ lie between $a_2(t_2)$ and $b_2(t_2)$. Without loss of generality, assume the points are ordered $a_2(t_2)$, $\dot{\Xi}_1(C(t_2))$, $\dot{\Xi}_1(C(t_1))$, $\dot{\Xi}_2(C(t_2))$, $b_2(t_2)$. Furthermore, assume $\dot{\Xi}_1(C(t_2))$ is closer to $\dot{\Xi}_1(C(t_1))$ than to $\dot{\Xi}_2(C(t_1))$. (See Fig. 3a.)

Figure 3: The two different cases for the relative positions of $\Xi_1(C(t_2))$, $\Xi_1(C(t_1))$ and $\Xi_2(C(t_2))$.

$$\begin{aligned}
d(\dot{\Xi}_1(C(t_1)), a_2(t_2)) &\leq d(\dot{\Xi}_1(C(t_1)), a_2(t_1)) + d(a_2(t_1), a_2(t_2)) \\
&\leq \rho(C(t_1)) + d(a_2(t_1), a_2(t_2)) \\
&\leq \rho(C(t_1)) + \delta, && \text{since } d(a_2(t_1), a_2(t_2)) \leq \delta, \\
&\leq \rho(C(t_2)) + 2\delta, && \text{by Lemma 14.}
\end{aligned} \tag{18}$$

Observe that

$$\begin{aligned}
d(a_2(t_2), \dot{\Xi}_1(C(t_1))) &= d(a_2(t_2), \dot{\Xi}_1(C(t_2))) + d(\dot{\Xi}_1(C(t_2)), \dot{\Xi}_1(C(t_1))) \\
&= \rho(C(t_2)) + d(\dot{\Xi}_1(C(t_2)), \dot{\Xi}_1(C(t_1))).
\end{aligned} \tag{19}$$

Therefore, by (18) and (19),

$$\min\{d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_1(C(t_2))), d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_2(C(t_2)))\} \leq 2\delta.$$

Case 2. Assume $\dot{\Xi}_1(C(t_1))$ does not lie on the path between $\dot{\Xi}_1(C(t_2))$, and $\dot{\Xi}_2(C(t_2))$. By Corollary 12, $\dot{\Xi}_1(C(t_1))$ lies between $a_1(t_1)$ and $b_1(t_1)$. Without loss of generality, assume $a_1(t_1)$ lies on the branch rooted at $\dot{\Xi}_1(C(t_1))$ that does not contain $\dot{\Xi}_1(C(t_2))$ and $\dot{\Xi}_2(C(t_2))$. (See Fig. 3b.)

Using an argument similar to Case 1, we get

$$\begin{aligned}
&\min\{d(a_1(t_2), \dot{\Xi}_1(C(t_2))), d(a_1(t_2), \dot{\Xi}_2(C(t_2)))\} \leq \rho(C(t_2)), \\
\Rightarrow \min\{d(a_1(t_1), \dot{\Xi}_1(C(t_2))), d(a_1(t_1), \dot{\Xi}_2(C(t_2)))\} &\leq \rho(C(t_2)) + \delta, && \text{since } d(a_1(t_1), a_1(t_2)) \leq \delta, \\
&\leq \rho(C(t_1)) + 2\delta, && \text{by Lemma 14.}
\end{aligned} \tag{20}$$

Observe that

$$\begin{aligned}
&\min\{d(a_1(t_1), \dot{\Xi}_1(C(t_2))), d(a_1(t_1), \dot{\Xi}_2(C(t_2)))\} \\
&= d(a_1(t_1), \dot{\Xi}_1(C(t_1))) + \min\{d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_1(C(t_2))), d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_2(C(t_2)))\} \\
&= \rho(C(t_1)) + \min\{d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_1(C(t_2))), d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_2(C(t_2)))\}.
\end{aligned} \tag{21}$$

Therefore, by (20) and (21),

$$\min\{d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_1(C(t_2))), d(\dot{\Xi}_1(C(t_1)), \dot{\Xi}_2(C(t_2)))\} \leq 2\delta.$$

Therefore, (17) holds in all cases. $\square$

Since no mobile 2-centre can guarantee relative velocity less than two in one dimension [15], it follows that the maximum velocity of the equidistant 2-centre is optimal.

Corollary 16. *Each facility in the mobile equidistant 2-centre is continuous.*

5.3 Complexity of the Motion of the 2-Centre

When clients move with linear motion, we derive combinatorial bounds of $O(n^2\alpha(n))$ on the complexity of the motion of the equidistant 2-centre and $\Omega(n^2)$ on the worst-case complexity of the motion of any 2-centre.

Theorem 17. *When each client in C moves with linear motion along a path on T , the motion of each facility in the equidistant 2-centre of C is piecewise linear and is composed of $O(n^2\alpha(n))$ linear segments, where $n = |C|$.*

Proof. By Theorem 8, there exists a sequence of diametric pairs of C , denoted $\{a_1, b_1\}, \dots, \{a_m, b_m\}$, and a corresponding partition of the time interval I , denoted $I_1, \dots, I_m$, such that $m \in O(n)$. It suffices to show that for every i , the motion of each facility in the equidistant 2-centre of C is piecewise linear and is composed of $O(n\alpha(n))$ linear segments during I_i .

Choose any $i \in \{1, \dots, m\}$ and consider the motion of clients in C during I_i . For every t , let $C_1(t)$ and $C_2(t)$ be a diametric partition induced by $a_i(t)$ and $b_i(t)$. By Lemma 10, $a_i(t)$ is in a local diametric pair of $C_1(t)$ for all t . The second client opposite $a_i(t)$ in the local diametric pair corresponds to a furthest client from $a_i(t)$ in $C_1(t)$. For any client $c \in C$, $d(c(t), a_i(t))$ and $d(c(t), b_i(t))$ are piecewise-linear functions composed of at most four linear segments; consequently, c changes partitions $O(1)$ times. Within C_1 , therefore, the function $d(c(t), a_i(t))$, may be partially defined, with $O(1)$ intervals over which it is undefined. Finding the furthest client from a_i in C_1 corresponds to finding the upper envelope of the $n - 2$ distance functions for all clients in $C \setminus \{a_i, b_i\}$. Since the functions are partially defined, the upper envelope consists of $O(n\alpha(n))$ linear segments [5]. This function corresponds to the local diameter of $C_1(t)$. The maximum of the two local diameters determines the 2-radius; therefore, the 2-radius of C also consists of $O(n\alpha(n))$ linear segments during I_i . Since a_i and b_i have linear motion, the result follows by Definition 2. $\square$

Theorem 18. *There exists a set of mobile clients C , each moving with linear motion on the real line, such that the motion of some facility in any 2-centre of C whose motion is piecewise linear is composed of $\Omega(n^2)$ linear segments, where $n = |C|$.*

Proof. We define a set of n clients for any even n . For $0 \leq i \leq n/2 - 1$, let client c_i have position $c_i(t) = \text{sign}(i)2(n^2 - i^2 + it)$ and velocity (relative to $-\infty$) $v_i = \text{sign}(i)i$. Let the remaining $n/2$ clients have velocity zero and be positioned at distinct points in $(-1, 1)$. See Fig. 4.

Figure 4: The initial configuration of clients when $n = 12$

For $0 \leq i \leq n/2 - 1$, client c_i passes client c_{i-2} at time $t = 2(i - 1)$. Observe that $\Xi(C(t)) = 1$ when $t \bmod 4 = 0$ and $\Xi(C(t)) = -1$ when $t \bmod 4 = 2$ for all $t \in [0, n]$. Therefore, the 1-centre traverses the interval $[-1, 1]$ $n/2$ times, crossing each of the $n/2$ static clients on each traversal.

The partition of C defined by (12) is unique whenever $\Xi(C(t))$ does not coincide with any client in C . The 2-radius is uniquely determined by the partition of larger local radius. Furthermore, any 2-centre of C must include one facility whose position is uniquely determined by the midpoint of the local diametric path of the partition with larger local radius. The motion of this facility changes $\Omega(n)$ times between each change to the motion of $\Xi(C(t))$, resulting in $\Omega(n^2)$ changes in total. $\square$

5.4 Kinetic Maintenance of the Mobile 2-Centre

Capitalizing on our 1-centre results, we describe an algorithm for identifying local 1-centres and the equidistant 2-centre of a set of mobile clients, each moving with linear motion on a tree.

By Theorem 18, even under linear motion of clients in C the motion of any 2-centre of C has complexity $\Omega(n^2)$ in the worst case. It follows that any algorithm for calculating the trajectories of a mobile 2-centre of C requires $\Omega(n^2)$ time in the worst case. Since it can be calculated in $O(n^2)$ time, we assume a fully-defined client-to-client distance function.

Algorithm Description. We first run our 1-centre algorithm to find a sequence of diametric pairs of C , denoted $\{a_1, b_1\}, \dots, \{a_m, b_m\}$, and a corresponding partition of the time interval I , denoted $I_1, \dots, I_m$, such that $m \in O(n)$. For each time interval I_i , determine when each client c is closer to a_i and when it is closer to b_i . This determines the sets $C_1(t)$ and $C_2(t)$ for all $t \in I_i$. Consider C_1 (an analogous algorithm applies to C_2). A diametric pair of $C_1(t)$ is given by $a_i(t)$ and a furthest client from $a_i(t)$ in $C_1(t)$. Each local

diametric pair determines the motion of the corresponding local 1-centre and the local radius, from which the motion of the equidistant 2-centre is straightforward to calculate.

Time Complexity. For a client $c \in C$, the functions $d(c(t), a_i(t))$ and $d(c(t), b_i(t))$ are piecewise linear, each composed of at most four linear segments. Therefore, c changes partitions $O(1)$ times during interval I_i and calculating the interval for which c resides in either partition is achieved in constant time. Finding a furthest client from $a_i(t)$ for all $t \in I_i$ corresponds to finding the upper envelope of $n - 2$ partially-defined, piecewise-linear functions, which can be done in $O(n \log n)$ time using Hershberger’s [26] algorithm. Since there are $O(n)$ time intervals, the total runtime is $O(n^2 \log n)$.

KDS Implementation. We describe a KDS that maintains the equidistant 2-centre over time along with a set of certificates that validates its identity.

We augment the 1-centre KDS described in Sec. 4.3. We require one additional certificate per client c to verify whether c is in C_1 or C_2 . We require a maximum KDS for C_1 (and a second one for C_2) that maintains the furthest client from a_i (respectively, b_i). The kinetic tournament KDS described by Basch et al. [8] allows for clients to be inserted and deleted from the set (recall that each client changes sets $O(1)$ times between changes to the diametric pair). This latter KDS has efficiency, compactness, locality, and responsiveness that is comparable to our 1-centre KDS.

In terms of performance, the worst case occurs whenever the diametric pair changes and $O(n)$ certificates must be updated. Therefore, this KDS has locality $O(n)$. The total number of certificate failures is $O(n \log n)$ between changes to the diametric pair, or $O(n^2 \log n)$ in total. By Theorem 18, the number of external events is $\Omega(n^2)$ in the worst case; therefore, the KDS has good efficiency (but it may not be optimal). The total number of certificates remains $O(n)$; therefore, the KDS is compact. Finally, $O(n)$ certificates are associated with each diametric client; therefore, the KDS has responsiveness $O(n)$.

6 Future Work

Discrete k -Centre. The mobile discrete k -centre (when each facility is a client in C) is discontinuous. Maintaining a sequence of clients that realize a discrete k -centre of a set of mobile clients on a tree is an open problem. Maintaining a discrete 1-centre of C corresponds to maintaining the identity of a client in C that is closest to $\Xi(C(t))$. For the discrete 2-centre, however, the problem is complicated by the fact that a diametric partition does not necessarily correspond to a discrete diametric partition; that is, (11) and (12) are not necessarily equivalent in the discrete case.

k -Centre on Graphs. We have omitted a detailed discussion of the k -centre on graphs due to space considerations. As mentioned earlier, it is not difficult to show that for any graph G that contains a cycle, there exist sets of mobile clients on G whose 1-centre is discontinuous. This motivates the search for bounded-velocity approximations of the k -centre. For the 1-centre on graphs, we have preliminary results showing that no continuous $(2 - \epsilon)$ -approximation is possible for any $\epsilon > 0$. A unit-velocity 2-approximation is given by selecting an arbitrary client $c_0 \in C$ and setting the position of the facility to coincide with $c_0(t)$. It is unknown whether any bounded-velocity approximation exists for the k -centre on graphs when $k \geq 2$.

References

- [1] Pankaj K. Agarwal, Julien Basch, Leonidas J. Guibas, John Hershberger, and Li Zhang. Deformable free-space tilings for kinetic collision detection. *International Journal of Robotics Research*, 21(3):179–197, 2002.
- [2] 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.

- [3] 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.
- [4] Pankaj K. Agarwal and Sarel Har-Peled. Maintaining approximate extent measures of moving points. In *Proceedings of the Symposium on Discrete Algorithms*, pages 148–157. ACM Press, 2001.
- [5] Pankaj K. Agarwal and Micha Sharir. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, 1995.
- [6] Julien Basch. *Kinetic Data Structures*. PhD thesis, Stanford University, 1999.
- [7] Julien Basch, Jeff Erickson, Leonidas J. Guibas, John Hershberger, and Li Zhang. Kinetic collision detection between two simple polygons. *Computational Geometry*, 27(3):211–235, 2004.
- [8] Julien Basch, Leonidas Guibas, and John Hershberger. Data structures for mobile data. *Journal of Algorithms*, 31(1):1–28, 1999.
- [9] 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.
- [10] Julien Basch, Leonidas J. Guibas, and Li Zhang. Proximity problems on moving points. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 344–351. ACM-SIAM, 1997.
- [11] Sergey Bereg, Binay Bhattacharya, David Kirkpatrick, and Michael Segal. Competitive algorithms for mobile centers. *Mobile Networks and Applications*, 11(2):177–186, 2006.
- [12] 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.
- [13] Gary Chartrand and Ping Zhang. h -convex graphs. *Mathematica Bohemica*, 126(1):209–220, 2001.
- [14] P. M. Dearing and Richard L. Francis. A minimax location problem on a network. *Transportation Science*, 8:333–343, 1974.
- [15] Stephane Durocher. *Geometric Facility Location under Continuous Motion*. PhD thesis, University of British Columbia, 2006.
- [16] 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.
- [17] Herbert Edelsbrunner, Leonidas J. Guibas, and Micha Sharir. The upper envelope of piecewise linear functions: Algorithms and applications. *Discrete and Computational Geometry*, 4:311–336, 1989.
- [18] 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.
- [19] Jie Gao, Leonidas J. Guibas, John Hershberger, Li Zhang, and An Zhu. Discrete mobile centers. *Discrete and Computational Geometry*, 30(1):45–65, 2003.
- [20] 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.
- [21] Gabriel Y. Handler. Minimax location of a facility in an undirected tree graph. *Transportation Science*, 7:287–293, 1973.

- [22] Gabriel Y. Handler. Finding two-centers of a tree: The continuous case. *Transportation Science*, 12:93–106, 1978.
- [23] Gabriel Y. Handler and Pitu B. Mirchandani. *Location on Networks: Theory and Algorithms*. MIT Press, Cambridge, Massachusetts, 1979.
- [24] Pierre Hansen, Martine Labbé, Dominique Peeters, and Jacques-François Thisse. Single facility location on networks. *Annals of Discrete Mathematics*, 31:113–146, 1987.
- [25] Sergiu Hart and Micha Sharir. Nonlinearity of Davenport-Schinzel sequences and of a generalized path compression scheme. In *Proceedings of Foundations of Computer Science*, volume 25, pages 313–319, 1984.
- [26] John Hershberger. Finding the upper envelope of n line segments in $O(n \log n)$ time. *Information Processing Letters*, 33(4):169–174, 1989.
- [27] John Hershberger. Smooth kinetic maintenance of clusters. In *Proceedings of the ACM Symposium on Computational Geometry*, pages 48–57. ACM Press, 2003.
- [28] 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.
- [29] 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.
- [30] David Kirkpatrick, Jack Snoeyink, and Bettina Speckmann. Kinetic collision detection for simple polygons. *International Journal of Computational Geometry and Applications*, 12(1&2):3–27, 2002.
- [31] 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.
- [32] 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.
- [33] Ady Wiernik and Micha Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete and Computational Geometry*, 3(1):15–47, 1988.