Information Processing Letters ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

The Fermat star of binary trees

Fabrizio Luccio *, Linda Pagli

Dipartimento di Informatica, Università di Pisa, Italy

ARTICLE INFO

Article history:
Received 1 July 2008
Available online xxxx
Communicated by S.E. Hambrusch

Keywords:
Fermat point
Fermat star
Steiner tree
Binary tree
Rotation distance
Combinatorial problems
Design of algorithms

ABSTRACT

A Fermat point P is one that minimizes the sum δ of the distances between P and the points of a given set. The resulting arrangement, called here a *Fermat star*, is a particular Steiner tree with only one intermediate point. We extend these concepts to rooted binary trees under the known rotation distance that measures the difference in shape of such trees. Minimizing δ is hard, due to the intrinsic difficulty of computing the rotation distance. Then we limit our study to establishing significant upper bounds for δ. In particular, for m binary trees of n vertices, we show how to construct efficiently a Fermat star with $\delta \leqslant mn - 3m$, with a technique inherited from the studies on rotation distance.

1. Introduction

Pierre de Fermat is credited for having proposed the following problem: given three points A, B, C in a plane, find a point P such that the sum of the three distances of A, B, C from P is minimal. The problem was solved by Torricelli around 1640 with a beautiful geometric construction that applies if the three angles of the triangle (ABC) are $< 120°$ [2]. Point P deriving from this construction was henceforth called the *Torricelli point*. Incidentally P is such that the three angles APB, APC, BPC are of $120°$. If the triangle (ABC) has an angle $> 120°$ in vertex (say) A, this vertex is the minimizing point.

In 1750 the problem was extended by Simpson to an arbitrary number $m \geqslant 3$ of points, in an arbitrary d-dimensional space, seeking a point P with a minimal sum of distances from the m points. Let us call *Fermat star* the set of the given points plus P, together with the connecting segments of minimal total length. The center P of the star will be called *Fermat point*. As it may not be easy determining a point P that strictly minimizes the sum δ of the lengths of the star segments, we will refer to a Fermat star even if δ stays within a certain upper bound.

From the problem of Fermat mathematicians passed to the now famous problem of finding a connection of minimal length among m points, giving rise to an arrangement of points and segments that, recently and improperly, has been called *Steiner tree*, although Steiner seems not to have been substantially involved in studying the problem. The obvious difference with the Fermat star is that, in the Steiner tree, more than one new intermediate point may be present, although for $m = 3$ the two arrangements coincide. Due to its importance in real life connection problems, Steiner trees have been extensively studied in many forms [2,7]. In most cases minimizing the total length of connections as required is an NP-hard problem.

We suggest considering the above problems in a new environment of theoretical nature, where the given points are binary trees and the distance between them is measured in number of rotations. So, also the new intermediate points will be binary trees. In this paper we limit our investigation to constructing a Fermat star whose new intermediate point will be called *Fermat tree*. For this purpose we recall some definitions and properties of tree transformation.

* Corresponding author.
E-mail addresses: luccio@di.unipi.it (F. Luccio), pagli@di.unipi.it (L. Pagli).

0020-0190/$ – see front matter
doi:10.1016/j.ipl.2009.02.001

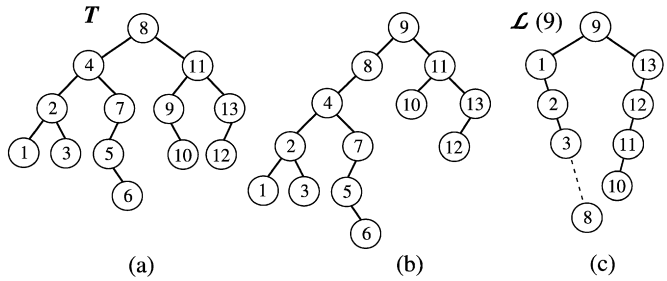

Fig. 1. (a) Arms and legs of a vertex, e.g., $L(4) = 4, 2, 1$; $R(4) = 4, 7$; $F(4) = 2, 3$; and $G(4) = 7, 5$. (b) Lifting vertex 9 to the root with $k = 2$ rotations: the value of $\phi(9) + \gamma(9)$ passes from 1 to 3. (c) The leg-tree $\mathcal{L}(9)$ obtained with $13 - (3 + 1) = 9$ additional rotations.

We consider rooted binary trees of n vertices, simply called *trees* in the following. The vertices are identified with the integers from 1 to n in *infix order*. Then, for any vertex v, the integers of the left (respectively, right) subtree are smaller (respectively, larger) than v. See the trees in Fig. 1, with $n = 13$. The following concepts and definitions are standard and well known.

Vertex rotations are local changes in a tree T preserving the infix integer assignment. If a vertex x is the left child of y, a *right rotation* at x raises x to the place of y while this vertex becomes the right child of x, and the original right subtree of x becomes the left subtree of y. In Fig. 1(a), a right rotation at vertex 4 would bring this vertex to the root of T, lower 8 as a right child of 4, and attach the subtree rooted at 7 as left subtree of 8. A *left rotation* is symmetrical. If such a rotation is applied at a vertex y lowered by the right rotation described above, it reconstructs the original tree.

Given two trees T_1, T_2, the *rotation distance* $d(T_1, T_2)$ is the minimum number of left and right rotations by which T_1 can be transformed into T_2. The transformation of T_1 into T_2 is denoted by $T_1 \to T_2$. As right and left rotations are one the inverse of the other, we have $d(T_1, T_2) = d(T_2, T_1)$. In a seminal paper based on hyperbolic geometry [6], Sleator, Tarjan, and Thurston proved that $d(T_1, T_2) \leqslant 2n - 6$ for any pair T_1, T_2, and that this bound is sharp for any large value of n. Luccio and Pagli [3] then gave a combinatorial proof of the upper bound. No efficient algorithm is known for computing the rotation distance of two arbitrary trees, but estimates of this value were given, e.g., by Rogers [5] and Pallo [4]. The technique of [4] was then adopted by Baril and Pallo for evaluating lower and upper bounds on this distance efficiently [1]. Many variations and extensions of the problem have been proposed, that have no particular connection with the present work. In proving our results we use a technique similar to the one of [3].

As an efficient procedure for determining the rotation distance of two trees is unknown, we will study upper bounds on the total number δ of rotations needed for transforming the given trees to a Fermat tree. δ is now called *global distance*, and we do not insist on finding a minimal value for it. Since the transformation process can be reversed by exchanging right and left rotations, a transformation $T_i \to T_j$ can be studied through the two transformations $T_i \to U$, $T_j \to U$ where U is a proper intermediate tree, and then reversing the second transformation. A Fermat star for m trees then provides a path for transforming anyone of them into another, passing through the Fermat tree.

2. Building a Fermat star of trees

As we shall see, important features for our problem are left and right chains of vertices. For any vertex v of a tree T, let the *left arm* $L(v)$, and the *right arm* $R(v)$, be the longest chains of vertices traversed from v when descending along T via left edges, and via right edges, respectively (v is included in both $L(v)$ and $R(v)$). And let the *left leg* of v be $F(v) = R(\mathit{leftchild}(v))$, and the *right leg* of v be $G(v) = L(\mathit{rightchild}(v))$ (see Fig. 1(a) and its caption). Note that only F and G may be void. Denote by $\lambda(v) \geqslant 1$, $\rho(v) \geqslant 1$, $\phi(v) \geqslant 0$, $\gamma(v) \geqslant 0$ the number of vertices in $L(v)$, $R(v)$, $F(v)$, $G(v)$, respectively. Furthermore we set:

Definition 1. A *leftchain-tree* $\mathcal{C}_L$ is the left chain rooted at vertex n. A *rightchain-tree* $\mathcal{C}_R$ is the right chain rooted at vertex 1. A *leg-tree* $\mathcal{L}(v)$ is a tree with vertex v at the root, such that the left and the right subtrees of v (if any) are a rightchain-tree for vertices $\{1, \ldots, v-1\}$, and a leftchain-tree for vertices $\{v+1, \ldots, n\}$, respectively (see Fig. 1(c)).

We now present some properties of arms and legs. For a vertex v, let $RS(v)$ and $LS(v)$ be the sets of all the vertices in the subtrees rooted at the right children of $L(v)$, or rooted at the left children of $R(v)$, respectively. E.g., $RS(4) = \{7, 5, 6, 3\}$ and $LS(4) = \{2, 1, 3, 5, 6\}$ in Fig. 1(a). We have:

Fact 1. *The vertices in $RS(v)$ can be brought into $L(v)$ with $|RS(v)|$ left rotations; similarly the vertices in $LS(v)$ can be brought into $R(v)$ with $|LS(v)|$ left rotations; and these bounds are sharp.*

To prove Fact 1 note that, for bringing the vertices of $RS(v)$ into $L(v)$ each left rotation must be performed at a vertex adjacent to $L(v)$ in the current step. Symmetrical operations are done for $LS(v)$. For example, in the tree of Fig. 1(a) the vertices of $RS(4)$ can be moved into $L(4)$ with four left rotations. From Fact 1 we immediately have:

Fact 2. *A tree of root r can be transformed into the leftchain-tree $\mathcal{C}_L$ with $|RS(r)| = n - \lambda(r)$ left rotations, or into the rightchain-tree $\mathcal{C}_R$ with $|LS(r)| = n - \rho(r)$ left rotations, and these bounds are sharp.*

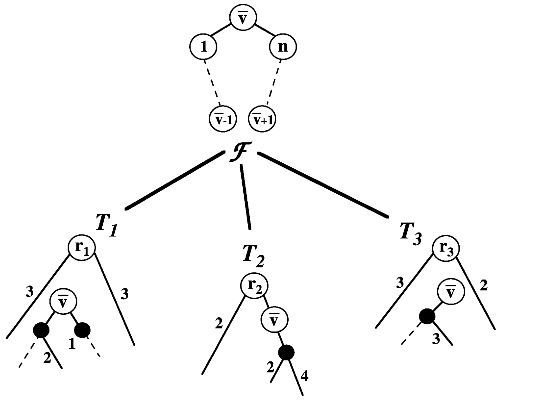

Fig. 2. A Fermat star for T_1, T_2, T_3, assuming $n = 20$, $\lambda = 3 + 2 + 3 = 8$, $\rho = 3 + 4 + 2 = 9$, $\sigma(\bar{v}) = 2 + 0 + 3 + 1 + 2 + 0 = 8$. By Theorem 1, the Fermat tree $\mathcal{F} = \mathcal{L}(\bar{v})$ is at global distance $\delta = 3 \cdot 20 - 8 - 3 = 49$ from T_1, T_2, T_3.

For the legs we state the following facts whose proofs are also immediate:

Fact 3. *A right rotation at vertex v increases by one the length of its right leg, while its left leg is unchanged, i.e. $\gamma(v) = \gamma(v) + 1$. A left rotation at vertex v increases by one the length of its left leg, while its right leg is unchanged, i.e. $\phi(v) = \phi(v) + 1$.*

Fact 4. *A tree T can be transformed into a leg-tree $\mathcal{L}(v)$ by lifting vertex v to the root with k consecutive rotations at v, where k is the distance of v from the root; then transforming the left subtree of the root into a right chain, and the right subtree of the root into a left chain, as indicated in Fact 2. The procedure requires $n - (\phi(v) + \gamma(v) + 1)$ rotations in total. We do not know whether this bound is sharp.*

The two main stages of the transformation $T \to \mathcal{L}(v)$ can be studied in Fig. 1 (a) to (c). The declared number of rotations is due to Facts 2 and 3, the latter implying that the initial k steps for bringing v to the root also increase by k the value $\phi(v) + \gamma(v)$. The transformations in Facts 2 and 4 are basic tools for building a Fermat tree $\mathcal{F}$ for a set of trees, establishing a constructive upper bound on the global distance δ of a Fermat star.

Consider $m \geqslant 3$ trees $T_1, T_2, \ldots, T_m$ of n vertices, with roots $r_1, r_2, \ldots, r_m$, and let $\lambda = \sum_{i=1,\ldots,m} \lambda(r_i)$, $\rho = \sum_{i=1,\ldots,m} \rho(r_i)$. For each vertex v, let $\phi_i(v)$ and $\gamma_i(v)$ denote the lengths of the two legs of v in T_i. Let $\sigma(v) = \sum_{i=1,\ldots,m} (\phi_i(v) + \gamma_i(v))$, and let $\bar{v}$ be a vertex that maximizes the value of σ. From Facts 2 and 4 we immediately have:

Theorem 1. *For the trees $T_1, T_2, \ldots, T_m$, a Fermat tree $\mathcal{F}$ can be built as*: (i) $\mathcal{F} = \mathcal{C}_L$, *with* $\delta = mn - \lambda$; (ii) $\mathcal{F} = \mathcal{C}_R$, *with* $\delta = mn - \rho$; (iii) $\mathcal{F} = \mathcal{L}(\bar{v})$, *with* $\delta = mn - \sigma(\bar{v}) - m$.

It can be easily verified that the three forms for $\mathcal{F}$ indicated in Theorem 1 can be built in time $O(mn)$. In particular, $\bar{v}$ can be found within the same time bound. Clearly the form for $\mathcal{F}$ chosen among the above three will be the one that minimizes δ. See Fig. 2 where the best choice for the Fermat tree is $\mathcal{L}(\bar{v})$ with $\delta = 49$. This bound, however, can be re-formulated in a weaker but simpler form. We have:

Theorem 2. *For $n > \lambda + \rho$ we have $\sigma(\bar{v}) \geqslant 2m$.*

Proof. First note that in a tree T of root r we have: $\sum_{v \in T} (\phi(v) + \gamma(v)) = 2n - \lambda(r) - \rho(r)$. This is an immediate consequence of Lemma 1 of [3], and can be proved here by noting that each vertex in $RS(r)$ is in the right leg of exactly one other vertex of T, and each vertex in $LS(r)$ is in the left leg of exactly one other vertex of T. We then have: $\sum_{i=1,\ldots,m} \sum_{v \in T_i} (\phi(v) + \gamma(v)) = 2mn - \lambda - \rho$, and this value is $\geqslant 2mn - n$ for $n > \lambda + \rho$. That is, the average value of $\sum_{i=1,\ldots,m} (\phi_i(v) + \gamma_i(v))$ for all v is greater than $2m - 1$, and the thesis follows. $\square$

We can now state the following corollary of Theorems 1 and 2, that applies in the vast majority of cases for increasing n.

Corollary 1. *For $n > \lambda + \rho$ a Fermat tree can be built as $\mathcal{L}(\bar{v})$ with $\delta \leqslant mn - 3m$.*

In Fig. 2, for example, we have $n > \lambda + \rho$ and Corollary 1 applies, indicating that $\mathcal{L}(\bar{v})$ can be built with $\leqslant 3 \cdot 20 - 3 \cdot 3 = 51$ rotations.

3. A concluding remark

We have seen that the minimizing Fermat point in a plane may coincide with one of the given points. Similarly, the Fermat tree that minimizes the value of δ for m trees $T_1, \ldots, T_m$ may coincide with one of them, say T_i. E.g., this happens if the other trees can be transformed into T_i with one rotation each. Although such particular conditions can sometimes be checked in polynomial time, it is difficult in general to find a Fermat tree that minimizes δ, due to the difficulty of computing the rotation distance.

In fact a Fermat star for $T_1, \ldots, T_m$ can be built using the known transformation between two arbitrary trees within the upper bound of $2n - 6$ rotations (e.g., see [3]). One of the trees T_i can then be taken as the Fermat tree, and the other trees are transformed into T_i with $\delta \leqslant 2n(m-1) - 6(m-1)$. Clearly the upper bound of Theorem 1, or of Corollary 1, are asymptotically much lower.

Extending this work to the construction of a Steiner tree of trees would raise a similar discussion. As we do not pretend attaining a minimum value of δ, we simply note that the Fermat star for $T_1, \ldots, T_m$ is also a Steiner tree and take the former as a solution. We are unable at the moment of suggesting a construction with a better upper bound.

References

[1] J.L. Baril, J. Pallo, Efficient lower and upper bounds of the diagonal-flip distance between triangulations, Information Processing Letters 100 (2006) 131–136.

[2] F.K. Hwang, D.S. Richards, P. Winter, The Steiner Tree Problem, Annals of Discrete Mathematics, vol. 53, 1992.
[3] F. Luccio, L. Pagli, On the upper bound on the rotation distance of binary trees, Information Processing Letters 31 (1989) 57–60.
[4] J. Pallo, An efficient upper bound of the rotation distance of binary trees, Information Processing Letters 73 (3–4) (2000) 87–92.
[5] R.O. Rogers, On finding shortest paths in the rotation graph of binary trees, Congressus Numerantium 137 (1999) 77–95.
[6] D.D. Sleator, R.E. Tarjan, W.R. Thurston, Rotation distance, triangulations, and hyperbolic geometry, J. Amer. Math. Soc. 1 (3) (1988) 647–681.
[7] W.D. Smith, P. Shor, Steiner tree problems, Algorithmica 7 (1992) 329–332.