A Note on the Flip Distance between Non-crossing Spanning Trees

Nicolas Bousquet
Univ. Lyon, Université Lyon 1, CNRS, LIRIS UMR CNRS 5205, F-69621, Lyon, France

Valentin Gledel
Department of Mathematics and Mathematical Statistics, Umeå University, Sweden

Jonathan Narboni
Theoretical Computer Science Department, Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland

Théo Pierron
Univ. Lyon, Université Lyon 1, CNRS, LIRIS UMR CNRS 5205, F-69621, Lyon, France

Abstract

We consider spanning trees of n points in convex position whose edges are pairwise non-crossing. Applying a flip to such a tree consists in adding an edge and removing another so that the result is still a non-crossing spanning tree. Given two trees, we investigate the minimum number of flips required to transform one into the other. The naive $2n - \Omega(1)$ upper bound stood for 25 years until a recent breakthrough from Aichholzer et al. yielding a $2n - \Omega(\log n)$ bound in the worst case. We improve this result with a $2n - \Omega(\sqrt{n})$ upper bound, and we strengthen and shorten the proofs of several of their results.

Keywords and phrases spanning tree, flip distance, reconfiguration

Digital Object Identifier 10.57717/cgt.v2i1.36

Funding ANR project GrR (ANR-18-CE40-0032)
Jonathan Narboni: National Science Center of Poland grant 2019/34/E/ST6/00443

1 Introduction

We fix a set $P = \{v_1, \dots, v_n\}$ of n points in the plane in convex position (and we assume that $v_1, \dots, v_n$ appear in this order on the convex hull of P). We consider spanning trees of the complete graph whose set of vertices is P and whose edges are straight line segments. If no two edges of such a spanning tree intersect (except maybe on their endpoints), we say that it is *non-crossing*. A *flip* removes an edge of a non-crossing spanning tree and adds another one so that the result is again a non-crossing spanning tree of P. A *flip sequence* is a sequence of non-crossing spanning trees such that consecutive spanning trees in the sequence differ by exactly one flip. We study the problem of transforming a non-crossing spanning tree into another via a sequence of flips.

Given two non-crossing spanning trees T_1 and T_2, observe that the size $|T_1 \Delta T_2|$ of the symmetric difference between their sets of edges may decrease by at most 2 when applying a flip, hence $|T_1 \Delta T_2|/2$ flips are required (note that this quantity can be as large as n when T_1 and T_2 have no common edge). Hernando et al. [3] proved that there exist two trees T_1 and T_2 such that any flip sequence needs at least $\frac{3}{2}n - 5$ flips. Regarding upper bounds, Avis and Fukuda [2] proved that there always exists a flip sequence between T_1 and T_2 using at most $2n - 4$ flips. This simple $2n - \Omega(1)$ upper bound was not improved in the last 25 years until a recent work of Aichholzer et al. [1] who improved the upper bound into $2d - \Omega(\log d)$, where d is the number of edges of T_1 not appearing in T_2. In the worst case $d = n$, and this yields a $2n - \Omega(\log n)$ bound. They also proved that there exists a flip sequence of length $\frac{3}{2}n - h$ if the two spanning trees share h edges and one of them is a path.

In this paper, we improve the worst case upper bound of the recent result [1] by proving that there exists a transformation of length at most $2n - \Omega(\sqrt{n})$ (Corollary 4). We also provide a short proof of the existence of a transformation of length $\frac{3}{2}n$ when one tree is a path (Corollary 8), which reproves in a shorter way a result from [1] in the worst case. We relax this result by showing that if one of the trees contains an induced path of length t then there exists a transformation of length $2n - \frac{t}{3}$ (Corollary 7).

Finally, we prove that if one of the trees is a nice caterpillar (whose precise definition will be given in Section 2), the shortest transformation has length at most $\frac{3}{2}n$ (Corollary 5). This is remarkable since, as far as we know, in all the examples where $\frac{3}{2}n - \Omega(1)$ flips are needed, at least one of the two non-crossing spanning trees is a nice caterpillar [3, 1]. So our statement essentially ensures that, if $\frac{3}{2}n$ is not the tight upper bound, then the spanning trees between which a larger transformation is needed should be constructed quite differently.

We think that all our partial results give additional credit to the following conjecture:

▶ **Conjecture 1.** *There is a flip sequence between any pair of non-crossing spanning trees of length at most $\frac{3}{2}n$.*

All our proofs are simple, self-contained, and mainly follow from simple applications of a lemma stated at the beginning of the next section.

2 Results

Recall that all along the paper we consider a set of n points $V = \{v_1, \dots, v_n\}$ in convex position appearing in that order. A *leaf* of a tree T is a vertex of degree one. An *internal node* of T is a vertex that is not a leaf. A *border edge* is an edge of the convex hull, *i.e.* $v_i v_{i+1}$ for some i. Given two points v_i, v_j, $V \setminus \{v_i, v_j\}$ has two *parts* (possibly empty), namely $v_{i+1}, \dots, v_{j-1}$ and $v_{j+1}, \dots, v_n, v_1 \dots, v_{i-1}$. Finally, we say that t edges $a_1b_1, \dots, a_tb_t$ (that may share endpoints) are *parallel* if $a_1, a_2, \dots, a_t, b_t, b_{t-1}, \dots, b_1$ appear in that order in the cyclic ordering of V. If moreover all their endpoints are pairwise distinct, then we say the edges are *strictly parallel*.

Let us first prove the following claim:

▷ Claim 2. Let T be a non-crossing spanning tree and e be a border edge. Then we can add e in T with an edge-flip without removing any border edge (except if T only contains border edges).

Proof. Adding e to T does not create any crossing, since e belongs to the convex hull of P. Moreover, the unique cycle in $T \cup \{e\}$ must contain at least an edge e' that does not belong to the set of border edges, since otherwise $T \cup \{e\}$ is precisely the convex hull of P. ◀

All our results follow from the following simple but very useful lemma:

▶ **Lemma 3.** *Let $i \leqslant n$. Let T_1, T_2 be two non-crossing spanning trees of P such that T_1 contains all the edges $v_j v_{j+1}$ for $j < i$ and T_2 has no edge $v_k v_\ell$ with $k > i$ and $\ell > i$. Then there exists a flip sequence between T_1 and T_2 of length at most $|T_1 \Delta T_2|/2$.*

Proof. Let us denote by X the subset of points $\{v_1, \dots, v_i\}$ and E_X the set of edges $v_j v_{j+1}$ for every $j < i$.

Let us first prove that there is a transformation from T_2 into a tree T_2' containing all the edges of E_X where the size of the symmetric difference decreases at each step. Indeed, assume that T_2 does not contain an edge e of E_X. Since T_1 contains all edges in E_X, the cycle obtained by adding e to T_2 contains an edge f in $T_2 \setminus T_1$ (and thus in $T_2 \setminus E_X$). Flipping

e with f in T_2 decreases the size of the symmetric difference between T_1 and the resulting tree T_2^* by 2. We repeat this operation between T_1 and T_2^* until all the edges of E_X have been added. Denote by T_2' the resulting tree. Observe that T_2' consists in a path (containing the set of edges E_X) with leaves attached to it. In particular, all the vertices of $V \setminus X$ are leaves.

We now prove that, as long as $T_1 \neq T_2'$, one can find a *good* flip in T_1, *i.e.* such that after applying it to T_1, the resulting tree T_1' still contains all the edges of E_X, and $|T_1' \Delta T_2'| = |T_1 \Delta T_2'| - 2$. The conclusion immediately follows by iterating this argument on T_1' until we reach T_2'.

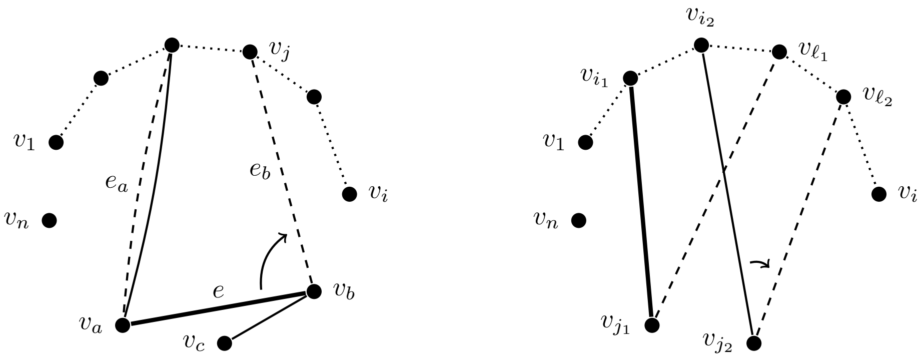

(a) The flip $e \to e_b$ is always good. Moreover, $v_b v_c$ is not external, while $v_a v_b$ is.

(b) The flip $v_{i_1} v_{j_1} \to v_{j_1} v_{\ell_1}$ is not good but $v_{i_2} v_{j_2} \to v_{j_2} v_{j'}$ is.

Figure 1 Dotted edges denote the path E_X, edges of T_2' are dashed and edges of T_1 are full.

Let us distinguish two cases:

Case 1. All the edges of $T_1 \setminus T_2'$ have both endpoints in $V \setminus X$ (see Figure 1 (a)).

An edge e of $T_1 \setminus T_2'$ is *external* if there is no other edge uv in $T_1 \setminus T_2'$ such that the endpoints of e not in $\{u, v\}$ do not lie in the same part as X in $V \setminus \{u, v\}$.

Since T_1 contains all the edges of E_X, if $T_1 \setminus T_2'$ is non empty, then it should contain an external edge $e = v_a v_b$. Since v_a and v_b are leaves of T_2' and all the edges of T_2' have an endpoint in X, there are two edges e_a, e_b linking X to respectively v_a and v_b. By symmetry, we may assume that the path from v_b to X in T_1 goes through v_a, and in particular T_1 does not contain e_b (since it contains all the edges of E_X). Therefore flipping $v_a v_b$ with e_b in T_1 yields a tree. Moreover it is non-crossing since if two edges of the resulting tree cross, one of them should be e_b. Since e_b is in T_2' and, by assumption, all the edges of T_1 between X and $V \setminus X$ are in T_2', the other edge f should have both endpoints in $V \setminus X$. This is impossible since e is external. So, $e \to e_b$ is a good flip, which concludes this case.

Case 2. $T_1 \setminus T_2'$ contains an edge between X and $V \setminus X$ (see Figure 1 (b)).
Let $v_{i_1} v_{j_1}$ with $i_1 < j_1$ be the edge of $T_1 \setminus T_2'$ between X and $V \setminus X$ such that i_1 is minimum and, with respect to that condition, $j_1 > i$ is maximum. Recall that v_{j_1} is a leaf of T_2', and let v_{ℓ_1} be its neighbor in X in T_2'.

Since T_1 contains E_X, exchanging $v_{i_1} v_{j_1}$ with $v_{\ell_1} v_{j_1}$ in T_1 still yields a tree. If it is a good flip, then we are done. Otherwise, there must be an edge $v_{i_2} v_{j_2}$ of T_1 crossing $v_{\ell_1} v_{j_1}$ and this edge is not in T_2' (since it crosses $v_{\ell_1} v_{j_1} \in T_2'$). Moreover, since both $v_{i_1} v_{j_1}$ and $v_{i_2} v_{j_2}$ are in T_1, they are parallel.

So either $i_2 \leqslant i_1$ and $j_2 \geqslant j_1$ or $i_2 \geqslant i_1$ and $j_2 \leqslant j_1$ (and at least one of the inequalities is strict). Our choices of i_1, j_1 ensures that the first case is impossible. So $i_2 \geqslant i_1$ and $j_2 \leqslant j_1$.

We now conclude with the following iterative argument. Since v_{j_2} is a leaf of T_2' connected to X, let v_{ℓ_2} be its neighbor in X. Note that $\ell_2 \geqslant \ell_1$ since $v_{j_1}v_{\ell_1}$ and $v_{j_2}v_{\ell_2}$ are parallel in T_2'. If exchanging $v_{i_2}v_{j_2}$ and $v_{j_2}v_{\ell_2}$ is a good flip, the conclusion follows. Otherwise an edge $v_{i_3}v_{j_3}$ of $T_1 \setminus T_2'$ is crossing $v_{j_2}v_{\ell_2}$. And $v_{i_1}v_{j_1}$, $v_{i_2}v_{j_2}$ and $v_{i_3}v_{j_3}$ are parallel since $\ell_2 > \ell_1$ and $v_{i_3}v_{j_3}$ cannot cross $v_{i_2}v_{j_2}$.

The repetition of this argument either provides a good flip or extracts a sequence $v_{i_1}v_{j_1}, \ldots, v_{i_r}v_{j_r}$ of parallel edges in T_1. Since T_1 contains $n-1$ edges, the process must stop after at most $n-1$ steps and yields a good flip, which concludes the proof. ◀

In the rest of the paper, we derive corollaries from that lemma. First, we provide a generic upper bound, which improves the one from [1] in the worst case:

▶ **Corollary 4.** *There exists a flip sequence of length at most $2n - \Omega(\sqrt{n})$ between any pair of non-crossing spanning trees.*

Proof. Let T_1, T_2 be two non-crossing spanning trees. Partition arbitrarily the set P into $\sqrt{n}$ sections of size $\sqrt{n}$. If one section does not contain any edge of T_2 with both endpoints in it, then we use Claim 2 to add to T_1 all the border edges outside of this section (in $n - \sqrt{n}$ flips), and then apply Lemma 3 to transform the resulting tree into T_2 with n additional flips.

Therefore, we can assume that each section contains both endpoints of an edge in T_2. In particular, the shortest edge contained in each section is a border edge (since T_2 is a non-crossing spanning tree). Applying again $n - \sqrt{n}$ times Claim 2 to all the non-border edges of T_2 and n times to those of T_1, we can transform both trees into a tree only containing border edges. This yields a flip sequence between T_1 and T_2 in at most $2n - \sqrt{n} + 1$ steps (since any two trees only containing border edges can be obtained from each other via a single edge-flip). ◀

A *caterpillar* is a tree such that the set of internal nodes induces a path. It is moreover *nice* if:

- it is a star, or
- it has exactly two internal nodes vw such that the two parts of $V \setminus \{v, w\}$ are the open neighborhood of v (except w) and the open neighborhood of w (except v), or
- it has at least three internal nodes such that, for every set of three consecutive internal nodes u, v, w, the neighbors of v are exactly u, w and all the points in the part of $V \setminus \{u, w\}$ which does not contain v (see Figure 2 for an illustration).

As far as we know, in all the examples where $\frac{3}{2}n - \Omega(1)$ flips are needed, at least one of the two non-crossing spanning trees is a nice caterpillar [3, 1].

▶ **Corollary 5.** *Let T_1, T_2 be non-crossing spanning trees such that T_2 is a nice caterpillar. There exists a flip sequence between T_1 and T_2 of length at most $\frac{3}{2}n$.*

Proof. Let us denote by $w_1, \ldots, w_k$ the set of internal nodes of T_2. Up to renaming, we can assume that $w_1 = v_1$ and denote by i the index such that $v_i = w_k$. Up to reversing the ordering of the vertices, we may also assume that $i \leqslant n/2$. Applying Claim 2 at most i times, we can add all edges v_jv_{j+1} for $j < i$ to T_1 and obtain a tree T_1'.

Since T_2 is a nice caterpillar, $\{v_1, \ldots, v_i\}$ contains every other w_j and all the leaves of T_2 attached to the w_j's in $\{v_i, \ldots, v_n\}$. Therefore, every edge of T_2 has an endpoint in

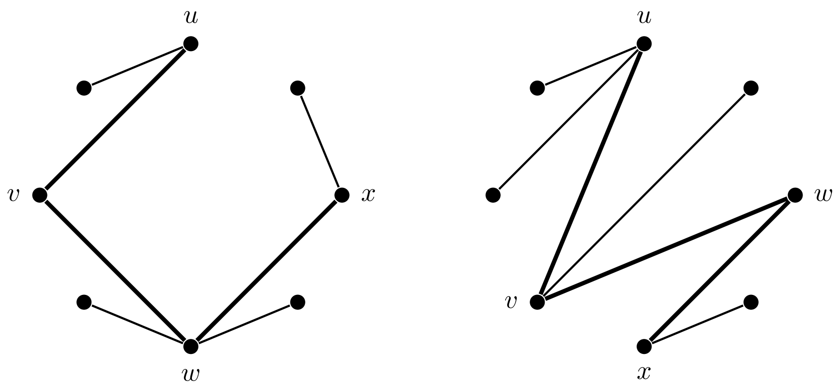

Figure 2 Two caterpillars, whose internal path $uvwx$ is highlighted in bold. The left one is not nice (v does not satisfy the desired property), while the right one is.

$\{v_1, \dots, v_i\}$ hence T_1' and T_2 satisfy the hypothesis of Lemma 3 and we can transform T_1' into T_2 in at most n steps for a total of $n + i \leqslant 3n/2$ steps. ◀

Using Lemma 3, one can also prove the following that will lead to another interesting corollary:

▶ **Lemma 6.** *Let T_1, T_2 be two non-crossing spanning trees such that T_2 has t parallel edges (resp. strictly parallel edges). There exists a flip sequence between T_1 and T_2 of length at most $2n - \frac{t+1}{2}$ (resp. $2n - t$).*

Proof. Let t be the maximum number of parallel edges in T_2. Let $e_1 = a_1b_1, \dots, e_t = a_tb_t$ be t parallel edges of T_2. We say that vertices between b_t and b_1 (resp. a_1 and a_t) in the cyclic ordering in the section that does not contain a_1 (resp. b_1) are *bottom vertices* (resp. *top vertices*). Let b be the number of bottom vertices.

Since T_2 is non-crossing, for every $i \leqslant t - 1$, there is a shortest path Q_i from an endpoint of e_i to an endpoint of e_{i+1} in T_2. Note that this path might be reduced to a single vertex if the two edges share an endpoint (we say that Q_i is *trivial*). Observe that the same trivial path may appear several times if several edges share the same endpoint. By maximality of t, there cannot be an edge in Q_i between a top vertex and a bottom vertex. Therefore we can classify the $t - 1$ paths Q_i in two types: Q_i is a *top path* if it only contains edges between top vertices and a *bottom path* otherwise. By symmetry, we can assume that at least $w := (t - 1)/2$ paths are top paths.

We claim that we can transform T_2 into a tree T_2' such that all the edges of the tree T_2' have at least one endpoint between a_1 and a_t in at most $b - w$ steps.

Denote by A_i the set of top vertices between a_i and a_{i+1} and by B_i the set of bottom vertices between b_{i+1} and b_i. Recall that by maximality of t there is no edge between A_i and B_i, except a_ib_i and $a_{i+1}b_{i+1}$. If Q_i is a non-trivial bottom path, we can remove one edge of the bottom part and add one edge in the top part to get a top path. We then aim at removing all edges of Q_i and connect all their endpoints in B_i to a_i. To this end, we say that an edge v_pv_q with $p < q$ with both endpoints in B_i is *exterior* if no edge v_rv_s distinct from v_pv_q with both endpoints in B_i satisfies $r \leqslant p < q \leqslant s$. One can easily remark that we can iteratively replace an exterior edge v_pv_q by an arc connecting v_p or v_q to a_i until no edge with both endpoints in B_i remains. So we can ensure that no edge with both endpoints in B_i remains in at most $|B_i| - 1$ steps (-2 if Q_i was initially a top path). If we sum over all the sections, since $\sum_i(|B_i|-1) = b-1$ and we remove 1 additional flip for each of the w top paths, this process yields a tree T_2' where all the edges have at least one endpoint between a_1 and a_t in $b-w-1$ flips.

Now we can transform T_1 into a tree T_1' that contains all border edges except maybe between bottom vertices in at most $n-b$ steps by Claim 2. Finally, we may apply Lemma 3 to transform T_1' into T_2' in at most n steps, which in total gives a flip sequence of length at most $(b-w-1)+(n-b)+n = 2n-w-1$, as claimed.

In the strictly parallel case, let t' be the maximum number of strictly parallel edges. Observe that each non-trivial top path must contain a border edge in T_2 (and then in T_2'). Note that there are at most $t-t'$ trivial top paths, hence T_2' and T_1' share at least $w-t+t'$ border edges, and by Lemma 3, the flip sequence from T_1' to T_2' costs at most $n-w+t-t'$. The total length of the flip sequence between T_1 and T_2 is thus at most $(b-w-1)+(n-b)+(n-w+t-t') = 2n-2w-1+t-t' = 2n-t'$. ◀

Lemma 6 immediately implies:

▶ **Corollary 7.** *Let T_1, T_2 be two non-crossing spanning trees such that T_1 contains a subpath of length t. There exists a flip sequence between T_1 and T_2 of length at most $2n - \frac{t}{3}$.*

Proof. Let $Q := x_1, \dots, x_{t+1}$ be a subpath of T_1 of length t. For every $2 \leqslant i \leqslant t-1$, we say that the edge x_ix_{i+1} of Q is *separating* if x_{i-1} and x_{i+2} are separated by x_i, x_{i+1} (*i.e.* exactly one of x_i, x_{i+1} appear between x_{i-1} and x_{i+2} in the cyclic ordering of the vertices). We say that the edge is a *series edge* otherwise. By convention, the first and last edges of Q are both series and separating. Denote by s (resp. p) the number of series edges (resp. separating edges), so that $s+p=t+2$.

Observe that the set of separating edges of Q are parallel, hence Lemma 6 ensures that there exists a flip sequence of length at most $a = 2n - \frac{p-1}{2}$. Moreover, if x_ix_{i+1} is a series edge then there is a border edge in T_1 between x_i and x_{i+1} (in the part that does not contain the vertices x_{i-1} and x_{i+2}). So there exists also a flip sequence of length at most $b = 2n-s+1$ from T_1 to T_2 (passing through a border tree) by Claim 2. Now observe that $2a+b = 6n-t$, hence either a or b must be at most $\frac{6n-t}{3}$, which concludes. ◀

In the case of paths, we can actually improve Corollary 7 by finding a flip sequence of length at most $\frac{3}{2}n$. This reproves in a shorter way a result of [1] the worst case scenario, namely when the two trees do not share any edge[1].

▶ **Corollary 8.** *Let T_1, T_2 be two non-crossing spanning trees such that T_2 is a path. There exists a flip sequence between T_1 and T_2 of length at most $\frac{3}{2}n$.*

Proof. Let $x_1, \dots, x_n$ be the vertices of the path T_2 (in order). Deleting x_1 and x_n in the cyclic ordering separates the set of vertices into two parts called the *top* and the *bottom* parts. We consider that x_1 and x_n appear in both parts. Observe that all the edges of T_2 are either border edges (between two consecutive vertices of the top or the bottom part) or *traversing edges* with one endpoint in each part.

Let us denote by n_t (resp. n_b) the number of vertices of the top part (resp. bottom part) including x_1 and x_n. Note that $n_t + n_b = n+2$. Let us denote by b_t, b_b the number of border edges in T_2 respectively in the top and bottom parts.

---

[1] The result of [1] ensures that there exists a transformation whose length is at most $\frac{3}{2} \cdot |T_2 \setminus T_1|$ when T_2 is a path.

We add all the $n_t - 1$ border edges of the top part to T_1 by Claim 2. Then, we transform in T_2 all the b_b border edges of the bottom part into traversing edges as follows: flip each bottom border edge $x_i x_{i+1}$ with $x_j x_{i+1}$ where j is the largest index of a top vertex less than i. Observe that the two resulting trees share b_t common border edges in the top part, and satisfy the hypothesis of Lemma 3. Therefore there is a flip sequence of length at most $n - b_t$ between them, and thus we can transform T_1 in T_2 with at most $(n_t - 1) + b_b + (n - b_t)$ flips.

Exchanging the top and bottom parts in the previous argument yields another flip sequence of length $(n_b - 1) + b_t + (n - b_b)$. The sum of these lengths is at most $2n + n_b + n_t - 2 = 3n$ which ensures one of the two sequences has length at most $\frac{3}{2}n$, which completes the proof. ◀

References

1 Oswin Aichholzer, Brad Ballinger, Therese Biedl, Mirela Damian, Erik D. Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Josef Tkadlec, and Yushi Uno. Reconfiguration of non-crossing spanning trees. *CoRR*, abs/2206.03879, 2022. `arXiv:2206.03879`, `doi:10.48550/arXiv.2206.03879`.

2 David Avis and Komei Fukuda. Reverse search for enumeration. *Discret. Appl. Math.*, 65(1-3):21–46, 1996. `doi:10.1016/0166-218X(95)00026-N`.

3 M. Carmen Hernando, Ferran Hurtado, Alberto Márquez, Mercè Mora, and Marc Noy. Geometric tree graphs of points in convex position. *Discret. Appl. Math.*, 93(1):51–66, 1999. URL: `https://doi.org/10.1016/S0166-218X(99)00006-2`.