

Markov Incremental Constructions

Bernard Chazelle · Wolfgang Mulzer

Received: date / Accepted: date

Abstract A classic result asserts that many geometric structures can be constructed optimally by successively inserting their constituent parts in random order. These randomized incremental constructions (RICs) still work with imperfect randomness: the dynamic operations need only be “locally” random. Much attention has been given recently to inputs generated by Markov sources. These are particularly interesting to study in the framework of RICs, because Markov chains provide highly nonlocal randomness, which incapacitates virtually all known RIC technology.

We generalize Mulmuley’s theory of Θ -series and prove that Markov incremental constructions with bounded spectral gap are optimal within polylog factors for trapezoidal maps, segment intersections, and convex hulls in any fixed dimension. The main contribution of this work is threefold: (i) extending the theory of abstract configuration spaces to the Markov setting; (ii) proving Clarkson-Shor type bounds for this new model; (iii) applying the results to classical geometric problems. We hope that this work will pioneer a new approach to randomized analysis in computational geometry.

Keywords Randomized incremental constructions · Expander graphs · Clarkson-Shor bound

1 Introduction

Randomized incremental constructions (RICs) are popular for three reasons: they are widely applicable; they are as simple as one could hope; they are often optimal under random input sequences [4–8, 11, 12, 14–16, 18, 22, 31, 32, 34–36, 40–43, 48]. But what if the sequences are not truly random? In the worst case, the running time typically goes up by a factor of n . Less obvious is the fact that perfect randomness is not actually necessary. Mulmuley [37] proved that $O(1)$ -wise independence is in fact sufficient. On the other hand, Amenta et al. [2] showed that the entropy may slowly decay during the RIC without penalty; in other words,

This work was supported in part by NSF grants CCR-0306283, CCF-0634958.

B. Chazelle · W. Mulzer
Department of Computer Science, Princeton University, 35 Olden Street, Princeton, NJ 08540, USA
E-mail: {chazelle, wmulzer}@cs.princeton.edu

the insertion sequence can afford to be less and less random as the construction progresses. Devillers and Guigue [17] introduced the *shuffling buffer* which randomly permutes contiguous subsequences of the input sequence of a certain length k , and they provided trade-offs between the length k and the running time of the RIC. What these results demonstrate is that standard RIC analysis still works as long as there is sufficient *local* randomness *early* enough. Unfortunately, these two features are precisely what is lacking in *Markov sources*.

What are those? A Markov source is a probabilistic model of input data that serializes the production of data over time by means of a random walk in a graph. It is widely used in queuing theory, speech recognition, gesture modeling, protein homology, computer graphics, robotics, web searching, etc. It captures the statistical correlations created by *time coherence*. In speech, for example, the randomness of the next utterance is heavily dependent on the previous ones; hence the use of hidden Markov models. In geometric applications, Markov sources have been used in ray tracing [25, 46], computer games [30], robotics [21], terrain generation [47], etc. In computer science, one of the main motivations has been *locality of reference*; in particular, there exists a vast body of research in online algorithms for Markov sources [10, 24, 26–29, 38, 39, 45]. The work of Amenta et al. [2] on RICs is also motivated by the desire for local access. The focus of much of modern computing has shifted over to the “data” side (partly because of the need to cope with massive data sets), and it is natural to ask what happens to a general algorithmic paradigm (RIC) when one assumes a Markov source—arguably the most widely used probabilistic model in applied science today.¹

For the purpose of this paper, our model consists of an *event graph* $G = (V, E)$ which is connected and undirected. This means that G defines a Markov chain that is irreducible and reversible but not necessarily ergodic. Each node v is associated with an item x_v in a universe \mathcal{U} . Requests are specified by following a random walk, beginning at a random start node of G and hopping from node to node, each time choosing an adjacent node v uniformly at random. Upon reaching v , item x_v is inserted into the current structure. The structure in question depends, of course, on the application. In this work we consider convex hulls, trapezoidal maps, and segment intersections. The structure is the corresponding *conflict graph*. Actually, we can use a data structure called *influence graph* [5–8, 18, 19] or *history graph* [36], which has the advantage of supporting queries and allowing for online (semi)dynamic algorithms. This means that we do not even need to know the graph G ahead of time. Our analysis, in fact, supports all known variants of RICs.

Our results and their significance The main contribution of this paper is to extend Mulmuley’s theory of Θ -series [36] to Markov chains.² In the course of doing so, we prove two results of independent interest: one is a generalization of the classic Clarkson-Shor counting technique for Markov sampling; the other is a new bound on mean first passage times for Markov chains with bounded spectral gap.

With the new tools we build, we are able to bound the expected complexity of RIC for convex hulls in d dimensions by $O(\gamma^{-d} n^{\lfloor d/2 \rfloor} (\log n)^{\lfloor d/2 \rfloor})$ for $d > 3$, and $O(n(\gamma^{-1} \log n)^d)$ for $d \leq 3$, where γ is the spectral gap of the random walk, ie, the difference between the first and second largest eigenvalues of the transition matrix—note that γ is a positive constant in the case of a random graph or an expander. For trapezoidal maps of nonintersecting segments and segment intersections, the complexity is respectively $O(n(\gamma^{-1} \log n)^4)$ and $O((n+m)(\gamma^{-1} \log n)^6)$, where m is the number of intersections.

¹ See Appendix A for technical background on Markov Chains and our terminology.

² See Appendix C for more on configuration spaces and Θ -series

The polylog factors we achieve are unlikely to be tight, and improving them remains an intriguing open problem. We suspect that further progress hinges on a better understanding of the geometry of k -sets and of the short-term behavior of random walks, a topic that seems to have been addressed only recently in the Markov chain literature [3].

A reader familiar with the role of expanders in derandomizing **BPP** might expect that Markov sources should provide more, not less, randomizing power than, say, bounded independence.³ This intuition is wrong for very interesting reasons that are important to understand. The standard analyses for RICs require *global* randomness within *local* time windows. Markov sources violate that essential feature in the worst possible way. Even post-mixing, a short walk contains virtually no global randomness. Note that all previous uses of expanders for (de)randomization rely on their randomness over *global* windows: in that regard, this paper pioneers a local approach to Markov chains that is likely to find further applications. In particular, our contribution includes new results of general interest on first passage times.

That Markov RICs come within polylog and not, say, n^ϵ of optimal is, in and of itself, a very intriguing result. In fact, we do not know any simple argument that shows that the expected RIC complexity beats that of the worst possible deterministic insertion sequence (even assuming bounded spectral gap)! Markov sources seem to shatter the foundation of RIC's analysis at its core. We show in this paper how the framework of Θ -series can be partly salvaged. It is doubtful whether backwards analysis can be similarly rescued, but this is a fascinating open question.

RICs provide essentially the simplest algorithms one can hope for. The message of this paper is that a tiny amount of local entropy (as provided by Markov sources) is sufficient to bring about almost all (though not quite all) the benefits of full randomization. What our work also shows is that there is nothing obvious at all about such a statement.

2 A Simple Example: Treaps

We begin with a toy example that avoids some of the complications of the general case: suppose that each node v of G is labeled with an element x_v from a totally ordered universe and that all the labels are distinct. The structure to be maintained is a binary search tree T . Start with an empty tree and perform a random walk on G . When the walk reaches v for the first time, insert x_v into T . Our goal is to bound the expected time for the construction of T . For convenience, we assume that G is connected and r -regular⁴, for some constant r . The complexity of the algorithm is tightly coupled to the spectral gap γ , which is the difference between the first and second largest eigenvalues of the (stochastic) transition matrix. We will prove the following result.

Theorem 1 *The expected time to construct the binary search tree is $O(\gamma^{-1}n \log n)$.*

For example, if G is the complete graph with self-loops, $\gamma = 1$ and we get the original theorem about treap construction [33,44]. More interestingly, if G is a random r -regular graph or an expander, we have $\gamma = \Theta(1)$, and the running time is still optimal. A cycle, on the other hand, has $\gamma = \Theta(1/n^2)$, and Theorem 1 predicts a running time of $O(n^3 \log n)$.

For distinct $u, v \in V$, let S_{uv} denote the set of nodes $z \in V$ such that x_z lies in the open interval bounded by x_u and x_v . Furthermore, let I_{uv} be the indicator random variable for

³ **BPP** is the class of languages that can be decided by probabilistic polynomial time Turing machines with two-sided error bounded away from $1/2$.

⁴ A graph is r -regular if all its nodes have degree r .




Fig. 1 All elements $x_z, z \in \bar{S}$, are stored in a subtree rooted at x_w . (a) If $w = u$, then $I_{uv} = 0$; (b) if $w = v$, then $I_{uv} = 1$; and (c) if $w \in S_{uv}$, we have $I_{uv} = 0$.

the event that x_u is compared to x_v when x_u is inserted into T , ie, the event that x_v is an ancestor of x_u in the binary search tree. Clearly, the time to insert x_u into T is proportional to $\sum_{v \in V \setminus u} I_{uv}$. We claim that $I_{uv} = 1$ precisely if the random walk encounters v before any other node in $\bar{S} = S_{uv} \cup \{u, v\}$ [44, Lemma 4.3]: let w be the first node in \bar{S} that is encountered during the random walk, and let T^* be the tree just before the insertion of x_w . Since the labels of the nodes in \bar{S} constitute an interval, for every element x in T^* the comparison of x with any $x_z, z \in \bar{S}$, yields the same result, irrespective of which $z \in \bar{S}$ is chosen. Therefore, the search paths in T^* for all $x_z, z \in \bar{S}$, are identical. It follows that all those elements will be stored in a subtree rooted at x_w . We can now cover all the cases (see Figure 1): if $w = u$, then $I_{uv} = 0$, and if $w = v$, then $I_{uv} = 1$. Finally, if $w \in S_{uv}$, then x_u and x_v will be stored in different subtrees of x_w , and hence they will never be compared to each other, ie, $I_{uv} = 0$. Thus, the expected time to build the binary search tree is proportional to

$$\Theta = \sum_{\substack{u, v \in V \\ u \neq v}} \mathbf{E}[I_{uv}] = \sum_{\substack{u, v \in V \\ u \neq v}} \Pr[\text{the random walk meets } v \text{ before any node in } S_{uv} \cup \{u\}]. \quad (1)$$

To get a handle on this sum, we need some random walk theory. Recall that the *transition matrix* of a Markov process with n states is the $n \times n$ matrix P in which entry P_{ij} is the probability of a transition from state i to state j . The transition matrix of a random walk on a graph G is its adjacency matrix, normalized so that each row sums to one. Furthermore, for any initial probability distribution $\pi_0 \in \mathbb{R}^n$, the distribution after t steps equals $\pi_0^T P^t$.

For technical reasons, we assume a lazy walk with $P = \frac{1}{2}(I + M/r)$, where M is the adjacency matrix of G . This is only for analytical convenience, and an actual implementation could assume a random walk in the original graph G . For the cost of a constant-factor slowdown, the lazy walk brings with it well-known analytical benefits. For example, P is positive semidefinite and the walk is ergodic. Fix a node $u_0 \in V$. Given any nonempty set $S \subset V$ and $u \in V \setminus S$, let $\Pr[u_0 \xrightarrow{u} S]$ be the probability that an infinite walk from u_0 reaches u before any node in S , and let $t_0 = \lceil c(1 - \lambda)^{-1} \log n \rceil$ be an upper bound on the mixing time, where λ is the second largest eigenvalue of P and c is a large enough constant [13]. Note that $\lambda = 1 - \gamma/2$ (the factor $1/2$ comes from the lazy walk) and that

$$\lambda^{t_0} \leq 1/n, \quad (2)$$

for appropriate c , since $\lambda = 1 - (1 - \lambda) \leq \exp(\lambda - 1)$. We begin with a technical result of independent interest.⁵

⁵ For convenience, we use the Vinogradov notation \ll and \gg for $O(\cdot)$ and $\Omega(\cdot)$, respectively.

Lemma 1 For any given $u_0 \in V$, nonempty $S \subset V$, $u \in V \setminus S$,

$$\Pr[u_0 \xrightarrow{u} S] \ll \sum_{0 \leq t < 3t_0} (P^t)_{u_0 u} + \frac{1}{(1-\lambda)|S|}.$$

Proof We may assume that $u_0 \neq u$, since otherwise the sum on the right-hand side is at least 1 and the lemma holds trivially. Similarly, we assume $u_0 \notin S$, since otherwise $\Pr[u_0 \xrightarrow{u} S] = 0$. Let Q be the matrix derived from P by zeroing out any entry P_{vw} with either v or w (or both) in $S \cup \{u\}$. (We index matrix elements and vector coordinates by their corresponding nodes in G .) Being positive semidefinite, Q has a (real) spectral decomposition $\sum_i \mu_i \mathbf{z}_i \mathbf{z}_i^T$ such that $\mu_1 \geq \dots \geq \mu_n = 0$ and the \mathbf{z}_i constitute an orthonormal basis of eigenvectors. By the Perron-Frobenius theorem (Theorem 5 in Appendix B), $\lambda < 1$. We also have $\mu_1 < 1$. To see why, note that the components of $G \setminus (S \cup \{u\})$ induce a decomposition of Q into block matrices. If $\mu_1 = 1$, one of these block matrices Q' would have principal eigenvalue 1 and Perron-Frobenius would yield a corresponding eigenvector with all positive entries. But this is impossible, since Q' has a row whose entries sum to less than 1 (and all other rows sum to at most 1).⁶ By the eigenvalue interlacing lemma (Theorem 4), $\mu_2 \leq \lambda$, so for any $v, w \in V \setminus (S \cup \{u\})$

$$\begin{aligned} (Q^t)_{vw} &= z_{1v} z_{1w} \mu_1^t + \sum_{i>1} z_{iv} z_{iw} \mu_i^t && \text{(spectral decomposition)} \\ &\leq z_{1v} z_{1w} \mu_1^t + \mu_2^t \sqrt{\sum_{i>1} z_{iv}^2 \sum_{i>1} z_{iw}^2} && \text{(Cauchy-Schwarz and } \mu_i \leq \mu_2) \\ &\leq z_{1v} z_{1w} \mu_1^t + \lambda^t. && \text{(orthonormality of the } \mathbf{z}_i \text{ and } \mu_2 \leq \lambda) \end{aligned} \quad (3)$$

Since $\mathbf{1}/\sqrt{n}$ is the principal unit eigenvector of P for the eigenvalue 1, an analogous calculation for P yields for any $v, w \in V$:

$$(P^t)_{vw} \leq \frac{1}{n} + \lambda^t. \quad (4)$$

To bound $\Pr[u_0 \xrightarrow{u} S]$, we proceed as follows: first, we distinguish between short paths (with less than $3t_0$ steps) and long paths (with at least $3t_0$ steps). The contribution of the short paths constitutes the first summand in the bound of Lemma 1. To analyze the contribution of the long paths, we break down every long path from u_0 to u into a pre-mixing part, a mixed portion, and the premixed part of the reverse path. We then assess the contribution of each piece. Let N_u denote the set of nodes in $V \setminus S$ adjacent to u via a nonloop edge. Since G is r -regular, $|N_u| \leq r$. Note that $(Q^t)_{u_0 v}$ is the probability that a t -step random walk from u_0 ends in v while avoiding $S \cup \{u\}$. Therefore,

$$\Pr[u_0 \xrightarrow{u} S] = \frac{1}{r} \sum_{t=0}^{\infty} \sum_{v \in N_u} (Q^t)_{u_0 v} \leq \sum_{t < 3t_0} (P^t)_{u_0 u} + \frac{1}{r} \sum_{t \geq 3t_0} \sum_{v \in N_u} (Q^t)_{u_0 v}. \quad (5)$$

⁶ Let \mathbf{v} be this eigenvector, and suppose that the i th row of Q' sums to less than 1. Since $\mu_1 = 1$, there must be an index j such that $Q'_{ij} > 0$ and $v_j > v_i$, where v_i, v_j denote the corresponding components of \mathbf{v} . The j th row sums to at most 1, and by symmetry $Q'_{ji} > 0$. Hence there must be a j' with $Q'_{jj'} > 0$ and $v_{j'} > v_j$. Repeating this argument yields an arbitrarily long strictly increasing sequence of components of \mathbf{v} , which is impossible since \mathbf{v} has finite dimension.

We now break down the long paths. The last summand in (5) is bounded by

$$\begin{aligned}
& \frac{1}{r} \sum_{v \in N_u} \sum_{t \geq t_0} \sum_{a, b \in V} (P^{t_0})_{u_0 a} (Q^t)_{ab} (P^{t_0})_{bv} && \text{(break-up and } (Q^{t_0})_{uv} \leq (P^{t_0})_{uv}\text{)} \\
& \leq \left(\frac{1}{n} + \lambda^{t_0} \right)^2 \sum_{t \geq t_0} \sum_{a, b \in V} (Q^t)_{ab} && \text{(by (4) and } |N_u| \leq r\text{)} \\
& \leq \frac{4}{n^2} \sum_{t \geq t_0} \sum_{a, b \in V} (Q^t)_{ab}. && \text{(by (2))}
\end{aligned} \tag{6}$$

Since $\|\mathbf{z}_1\|_2 = 1$ and since at least $|S| + 1$ of its coordinates are zero (an easy consequence of being an eigenvector for Q), Cauchy-Schwarz yields $\|\mathbf{z}_1\|_1^2 \leq n - |S| - 1$. By applying Perron-Frobenius to the parts of the block decomposition of Q induced by the components of $G \setminus (S \cup \{u\})$, we can assume that \mathbf{z}_1 is nonnegative, and so

$$\sum_{a, b \in V} z_{1a} z_{1b} = \|\mathbf{z}_1\|_1^2 \leq n - |S| - 1. \tag{7}$$

We have

$$\begin{aligned}
\Pr[u_0 \xrightarrow{u} S] - \sum_{t < 3t_0} (P^t)_{u_0 u} &\leq \frac{1}{r} \sum_{t \geq 3t_0} \sum_{v \in N_u} (Q^t)_{u_0 v} && \text{(by (5))} \\
&\leq \frac{4}{n^2} \sum_{a, b \in V} \sum_{t \geq t_0} (z_{1a} z_{1b} \mu_1^t + \lambda^t) && \text{(by (3,6))} \\
&\leq \frac{4}{n(1 - \mu_1)} + \frac{4\lambda^{t_0}}{1 - \lambda}. && \text{(geometric sum and (7))}
\end{aligned} \tag{8}$$

We already noted $\mu_1 < 1$. However, to bound (8), we need a better estimate on $1 - \mu_1$. This can be done using an argument similar to one given by Broder and Karlin [9]: since \mathbf{z}_1 is nonnegative, $n\mathbf{z}_1 - \|\mathbf{z}_1\|_1 \mathbf{1}$ is normal to the principal eigenvector $\mathbf{1}$ of P , and since P is symmetric, by Courant-Fischer (Theorem 3),

$$\begin{aligned}
\lambda &\geq \frac{(n\mathbf{z}_1 - \|\mathbf{z}_1\|_1 \mathbf{1})^T P (n\mathbf{z}_1 - \|\mathbf{z}_1\|_1 \mathbf{1})}{\|n\mathbf{z}_1 - \|\mathbf{z}_1\|_1 \mathbf{1}\|_2^2} \\
&= \frac{n^2 \mathbf{z}_1^T P \mathbf{z}_1 - n \|\mathbf{z}_1\|_1 (\mathbf{z}_1^T P \mathbf{1} + \mathbf{1}^T P \mathbf{z}_1) + \|\mathbf{z}_1\|_1^2 \mathbf{1}^T P \mathbf{1}}{n^2 \|\mathbf{z}_1\|_2^2 - n \|\mathbf{z}_1\|_1 (\mathbf{z}_1^T \mathbf{1} + \mathbf{1}^T \mathbf{z}_1) + \|\mathbf{z}_1\|_1^2 \mathbf{1}^T \mathbf{1}}.
\end{aligned}$$

Now, since $\mathbf{1}$ is a left and right eigenvector of P , $\mathbf{1}^T P = \mathbf{1}^T$ and $P \mathbf{1} = \mathbf{1}$. Furthermore, $\mathbf{1}^T \mathbf{1} = n$, $\mathbf{z}_1^T \mathbf{1} = \mathbf{1}^T \mathbf{z}_1 = \|\mathbf{z}_1\|_1$, and $\|\mathbf{z}_1\|_2 = 1$. Hence,

$$\lambda \geq \frac{n\mathbf{z}_1^T P \mathbf{z}_1 - \|\mathbf{z}_1\|_1^2}{n - \|\mathbf{z}_1\|_1^2} \geq \frac{n\mu_1 - \|\mathbf{z}_1\|_1^2}{n - \|\mathbf{z}_1\|_1^2},$$

because $\mathbf{z}_1^T P \mathbf{z}_1 \geq \mathbf{z}_1^T Q \mathbf{z}_1 = \mu_1$. It follows that $n\mu_1 \leq n\lambda + (1 - \lambda)\|\mathbf{z}_1\|_1^2$, and using (7) we get $\mu_1 \leq 1 - (1 - \lambda)(|S| + 1)/n$. Plugging this bound into (8) completes the proof, as $\lambda^{t_0} \leq 1/n \leq 1/|S|$ by (2). \square

Proof (of Theorem 1) By (1), the expected running time is

$$\Theta = \sum_{\substack{u,v \in V \\ u \neq v}} \frac{1}{n} \sum_{u_0 \in V} \Pr[u_0 \xrightarrow{v} S_{uv} \cup \{u\}],$$

where u_0 is the random start node of the walk. By Lemma 1,

$$\Theta \ll \sum_{\substack{u,v \in V \\ u \neq v}} \frac{1}{n} \sum_{u_0 \in V} \left(\sum_{t=0}^{3t_0-1} (P^t)_{u_0v} + \frac{1}{(1-\lambda)(|S_{uv}|+1)} \right) = \sum_{\substack{u,v \in V \\ u \neq v}} \left(\frac{3t_0}{n} + \frac{1}{(1-\lambda)(|S_{uv}|+1)} \right),$$

since $\sum_{u_0 \in V} (P^t)_{u_0v} = 1$, as $(P^t)_{u_0v}$ is the probability that a t -step random walk ending in v started out at u_0 . Hence,

$$\Theta \ll nt_0 + (1-\lambda)^{-1} \sum_{\substack{i,j=1 \\ i \neq j}}^n \frac{1}{|i-j|+1} \ll \gamma^{-1} n \log n.$$

□

3 Θ -Series for Markov Sources

We use the classical notion of *configuration spaces* (see [36] or Appendix C for a primer) and adapt it to the Markov model. This is done as follows: fix a natural number d , the *degree* of the configuration space. Each node v of G is assigned an object x_v chosen from a geometric universe (eg, points, hyperplanes, segments), and to each d -tuple $\mathbf{u} = (u_1, \dots, u_d)$ of distinct $u_i \in V$ we assign a (possibly empty) $S_{\mathbf{u}} \subseteq V$ disjoint from \mathbf{u} . We denote by f_k the number of \mathbf{u} 's such that $|S_{\mathbf{u}}| = k$ and by $f_{\leq k}$ the prefix sum $f_0 + \dots + f_k$. We write $f_k(n)$ and $f_{\leq k}(n)$ to refer to the maximum such values over all subsets of the universe of size n . The coordinates of a d -tuple \mathbf{u} play the role of the *triggers* and the sets $S_{\mathbf{u}}$ that of the *stoppers*. Naturally, f_k counts the k -sets of the underlying range space.

The apparent simplifications of our model do not, in fact, restrict the generality of the results in any way. Indeed, our framework can just as easily handle cases where \mathbf{u} is not a sequence but a multiset, where it maps to several stopper sets, or where the degree d is not unique. Given a random ordered $\mathbf{u} = (u_1, \dots, u_d)$ with distinct elements, perform an infinite random walk from a random node in G . If the walk first reaches u_1, \dots, u_d in that order before hitting any node in $S_{\mathbf{u}}$, then set $\Phi = n^d |S_{\mathbf{u}}|$; else set $\Phi = 0$. Standard Θ -series theory shows that the expectation of Φ determines the expected amortized complexity of RIC [36]. As before, we assume that the graph is connected and r -regular, and we let γ denote the spectral gap. We postpone the proof of this result:

Theorem 2 (Master Theorem) *If there exists a constant $\alpha > 0$ such that $f_0(n) = O(n^\alpha)$, then $\mathbf{E}[\Phi] \ll \gamma^{-d} n^\alpha (\log n)^{d-\alpha}$ for $\alpha > 1$ and $\mathbf{E}[\Phi] \ll \gamma^{-d} n (\log n)^d$ for $\alpha \leq 1$.*

We apply the theorem to three problems: convex hulls (and hence Voronoi diagrams); trapezoidal maps of disjoint segments; and line segment intersections. For simplicity, we assume that the input is in general position. The algorithms themselves operate in standard incremental fashion by inserting objects online with the help of the history graph. The algorithms do not require knowledge of the Markov chain (which is why we do not use conflict graphs).

- CONVEX HULLS IN \mathbb{R}^d : The convex hull of n points in \mathbb{R}^d has $O(n^{\lfloor d/2 \rfloor})$ faces, which implies that $\alpha = \lfloor d/2 \rfloor$. The algorithm runs in time $O(\gamma^{-d} n^{\lfloor d/2 \rfloor} (\log n)^{\lfloor d/2 \rfloor})$ for $d > 3$, and $O(n(\gamma^{-1} \log n)^d)$ for $d \leq 3$.
- TRAPEZOIDAL MAPS: At each node, the trapezoidal map formed by a set of (nonintersecting) segments is maintained. The relevant configuration space is made of three subconfiguration spaces of respective degrees 2, 3, and 4. Hence, the time required by the algorithm is $O(n(\gamma^{-1} \log n)^4)$.
- SEGMENT INTERSECTIONS: The m intersections among n segments are computed in $O((n+m)(\gamma^{-1} \log n)^6)$ steps. The proof depends on an extension of the Master Theorem discussed in Section 4.

To bound the expectation of Φ , we need to understand a certain stochastic process, which we proceed to describe. A random *thread* refers either to a single node w_1 chosen uniformly at random (thread size of 1) or to a sequence w_1, \dots, w_l (thread size of $l > 1$), where w_1 is random and, for each $i > 0$, w_{i+1} is the end node of a random walk from w_i of length $t_i > 0$. The time sequence $\theta = (t_1, \dots, t_{l-1})$ parametrizes the thread. Given $1 \leq \mu \leq d$, a random μ -thread is a sequence of μ threads whose sizes add up to d : each thread is drawn independently and has its own size and time sequence. Its time sequence θ refers now to the collection of its constituent threads' time sequences. A μ -thread forms a d -tuple \mathbf{u} and is therefore associated with a stopper set⁷ $S_{\mathbf{u}}$. Let $g_k^{(\mu)}$ be the probability that a random μ -thread (with a given time sequence) produces \mathbf{u} such that $|S_{\mathbf{u}}| = k$.

$$g_k^{(\mu)} = \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u} : |S_{\mathbf{u}}| = k], \quad (9)$$

and let $g_{\leq k}^{(\mu)} = \sum_{0 \leq i \leq k} g_i^{(\mu)}$.

Lemma 2 *Let f_0 be monotonically increasing. For any μ -thread and any corresponding time sequence $\theta_1, \dots, \theta_{\mu}$, we have $g_{\leq k}^{(\mu)} \ll (k/n)^{\mu} f_0(n/k)$, for $k > 0$.*

Proof We use a Clarkson-Shor type counting argument [15] tailored for Markov chains. As usual, the idea is to use sampling in order to bound $g_{\leq k}^{(\mu)}$ in terms of f_0 . More precisely, we sample a set $R^v \subseteq V$ of size about n/k . Then, for a configuration $\mathbf{u} \subseteq R^v$ with $|S_{\mathbf{u}}| \leq k$, we argue that with constant probability \mathbf{u} is active in R^v , ie, $S_{\mathbf{u}} \cap R^v = \emptyset$. Together with a bound on the probability that a given configuration \mathbf{u} appears in R^v , this yields the desired result. We may assume that $k \leq n/2d$, since for larger k the bound becomes constant and $g_{\leq k}^{(\mu)} \leq g_{\leq n}^{(\mu)} \leq 1$.

All μ -threads in this proof share the given time sequence $\theta_1, \dots, \theta_{\mu}$. Let $s \leq n$ be an integer to be determined later. For each $i = 1, \dots, \mu$, pick s random threads of type θ_i , and define R as the set of \mathbf{u} 's formed by taking all possible s^{μ} combinations of the resulting threads, one of each type. Given a fixed (nonrandom) $\mathbf{u} \in V^d$, let $p_{\mathbf{u}}$ denote the probability that \mathbf{u} is chosen by a random μ -thread. Since each starting node is chosen independently, $p_{\mathbf{u}}$ is of the form $\prod_{1 \leq i \leq \mu} \frac{p_{\mathbf{u},i}}{n}$, where $p_{\mathbf{u},i}$ is the probability that the i -th thread visits the relevant nodes of \mathbf{u} in the correct order, given that the first node of the i -th thread equals the corresponding node in \mathbf{u} . Therefore, \mathbf{u} ends up in R with probability at least $\prod_{1 \leq i \leq \mu} (1 - (1 - p_{\mathbf{u},i}/n)^s)$. Now, since $p_{\mathbf{u},i} s/n \leq 1$, we have

$$\left(1 - \frac{p_{\mathbf{u},i}}{n}\right)^s \leq 1 - \binom{s}{1} \frac{p_{\mathbf{u},i}}{n} + \binom{s}{2} \left(\frac{p_{\mathbf{u},i}}{n}\right)^2 \leq 1 - \frac{p_{\mathbf{u},i} s}{n} + \frac{1}{2} \left(\frac{p_{\mathbf{u},i} s}{n}\right)^2 \leq 1 - \frac{p_{\mathbf{u},i} s}{2n}; \quad (10)$$

⁷ This is not true if \mathbf{u} contains less than d distinct nodes. Since Lemma 2 deals only with finite stopper sets, we can invalidate this case by setting $S_{\mathbf{u}} = \mathbb{R}$, or any other infinite set.

hence,

$$\Pr[\mathbf{u} \in R] \geq \prod_{i=1}^{\mu} \frac{p_{\mathbf{u},i^s}}{2n} \gg p_{\mathbf{u}} s^{\mu}. \quad (11)$$

Let R^v be the collection of nodes appearing among the d -tuples of R . Given a fixed \mathbf{u} with $|\mathcal{S}_{\mathbf{u}}| \leq n/2d$, conditioned upon $\mathbf{u} \in R$, what is the probability that $R^v \cap \mathcal{S}_{\mathbf{u}} = \emptyset$, ie, that configuration \mathbf{u} is active in R^v ? Being in R , \mathbf{u} itself is a μ -thread formed by picking exactly one thread per type out the s available ones in R . The d nodes of \mathbf{u} lie outside $\mathcal{S}_{\mathbf{u}}$, so the only possibility for R^v to intersect $\mathcal{S}_{\mathbf{u}}$ is for any of the $(s-1)\mu$ other threads to pass through $\mathcal{S}_{\mathbf{u}}$. Take one of them: it is a random walk $w_1 \dots w_l$. The starting node w_1 is random, so its distribution forms an eigenvector for the thread's transition matrix with eigenvalue 1 (also true if $l = 1$). This means that each w_i lies in $\mathcal{S}_{\mathbf{u}}$ with probability $|\mathcal{S}_{\mathbf{u}}|/n$. These events are not independent, so we use a union bound to argue that the thread $w_1 \dots w_l$ remains outside $\mathcal{S}_{\mathbf{u}}$ with probability at least $1 - l|\mathcal{S}_{\mathbf{u}}|/n \geq 1 - d|\mathcal{S}_{\mathbf{u}}|/n$. The $(s-1)\mu$ threads that are candidates for passing through $\mathcal{S}_{\mathbf{u}}$ are independent, however, and thus refrain from doing so with probability at least $(1 - d|\mathcal{S}_{\mathbf{u}}|/n)^{(s-1)\mu}$. For any $x \in [0, 1/2]$ we have $(1-x)^{-1} \leq (1+x)^2 \leq \exp(2x)$. Thus, $(1 - d|\mathcal{S}_{\mathbf{u}}|/n) \geq \exp(-2d|\mathcal{S}_{\mathbf{u}}|/n)$, since $|\mathcal{S}_{\mathbf{u}}| \leq n/2d$. It follows that

$$\Pr[R^v \cap \mathcal{S}_{\mathbf{u}} = \emptyset | \mathbf{u} \in R] \geq \left(1 - \frac{d|\mathcal{S}_{\mathbf{u}}|}{n}\right)^{(s-1)\mu} \geq e^{-2d(s-1)\mu|\mathcal{S}_{\mathbf{u}}|/n}.$$

If $r_{\mathbf{u}}$ denotes the probability that both $\mathbf{u} \in R$ and $\mathcal{S}_{\mathbf{u}} \cap R^v = \emptyset$ then, by (11), setting $s = \frac{n}{dk}$ yields

$$r_{\mathbf{u}} = \Pr[\mathbf{u} \in R] \times \Pr[R^v \cap \mathcal{S}_{\mathbf{u}} = \emptyset | \mathbf{u} \in R] \gg \left(\frac{n}{dk}\right)^{\mu} p_{\mathbf{u}} e^{-2\mu|\mathcal{S}_{\mathbf{u}}|/k};$$

therefore, since $\mu \leq d$ and $k \leq n/2d$,

$$\sum_{\mathbf{u}: |\mathcal{S}_{\mathbf{u}}| \leq n/2d} r_{\mathbf{u}} \gg \sum_{\mathbf{u}: |\mathcal{S}_{\mathbf{u}}| \leq n/2d} \left(\frac{n}{dk}\right)^{\mu} p_{\mathbf{u}} e^{-2\mu|\mathcal{S}_{\mathbf{u}}|/k} \gg \sum_{\mathbf{u}: |\mathcal{S}_{\mathbf{u}}| \leq k} \left(\frac{n}{dk}\right)^{\mu} p_{\mathbf{u}} \gg \left(\frac{n}{k}\right)^{\mu} g_{\leq k}^{(\mu)},$$

as $g_{\leq k}^{(\mu)} = \sum_{\mathbf{u}: |\mathcal{S}_{\mathbf{u}}| \leq k} p_{\mathbf{u}}$. Since $|R^v| \leq ds$, by definition and by the monotonicity of f_0 , $|\{\mathbf{u} \in R : |\mathcal{S}_{\mathbf{u}} \cap R^v| = 0\}| \leq f_0(ds)$; therefore,

$$\sum_{\mathbf{u}: |\mathcal{S}_{\mathbf{u}}| \leq n/2d} r_{\mathbf{u}} \leq f_0(n/k).$$

Note that this holds uniformly over all time sequences for the μ -thread. \square

Proof (of Theorem 2) Recall that our goal is to bound the expectation of a random variable Φ defined as follows: pick a random d -tuple $\mathbf{u} = (u_1, u_2, \dots, u_d)$ of distinct $u_i \in V$ and perform an infinite random walk in G starting at a random node u_0 . If the walk encounters all the nodes in \mathbf{u} in that order before any node in $\mathcal{S}_{\mathbf{u}}$, let $\Phi = n^d |\mathcal{S}_{\mathbf{u}}|$, otherwise, let $\Phi = 0$. The expectation of Φ is given by

$$\mathbf{E}[\Phi] = \frac{n^d d!}{\binom{n}{d}} \sum_{\mathbf{u}} \frac{1}{n} \sum_{u_0 \in V} |\mathcal{S}_{\mathbf{u}}| \prod_{i=0}^{d-1} \Pr \left[u_i \xrightarrow{u_{i+1}} \mathcal{S}_{\mathbf{u}} \cup \{u_{i+2}, \dots, u_d\} \right],$$

where $\sum_{\mathbf{u}}$ ranges over all ordered subsets of d distinct nodes: obviously, we may restrict the sum to $\{\mathbf{u} : |\mathcal{S}_{\mathbf{u}}| > 0\}$. The sum $d! \binom{n}{d}^{-1} \sum_{\mathbf{u}}$ represents the random choice of \mathbf{u} , the sum $n^{-1} \sum_{u_0 \in V}$ accounts for the random starting vertex. The product denotes the probability that

a random walk from u_0 visits the nodes u_1, \dots, u_d in that order before it encounters any node in $S_{\mathbf{u}}$. Note that removing elements from S cannot decrease $\Pr[u_0 \xrightarrow{u} S]$; therefore,

$$\mathbf{E}[\Phi] \ll \sum_{\mathbf{u}} \frac{A_{\mathbf{u}}}{n} |S_{\mathbf{u}}| \prod_{i=1}^{d-1} \Pr[u_i \xrightarrow{u_{i+1}} S_{\mathbf{u}}],$$

where

$$\begin{aligned} A_{\mathbf{u}} &\stackrel{\text{def}}{=} \sum_{u_0 \in V} \Pr[u_0 \xrightarrow{u_1} S_{\mathbf{u}}] \\ &\ll \sum_{u_0 \in V} \left(\sum_{0 \leq t < 3t_0} (P^t)_{u_0 u_1} + \frac{1}{(1-\lambda)|S_{\mathbf{u}}|} \right) \quad (\text{Lemma 1}) \\ &= 3t_0 + \frac{n}{(1-\lambda)|S_{\mathbf{u}}|} \quad (\text{by } \sum_{u_0} (P^t)_{u_0 u_1} = 1) \\ &\ll \frac{n}{1-\lambda} \left(\frac{1}{|S_{\mathbf{u}}|} + \frac{\log n}{n} \right). \quad (\text{by } t_0 \ll (1-\lambda)^{-1} \log n) \end{aligned}$$

Thus, using Lemma 1 once more,

$$\mathbf{E}[\Phi] \ll \frac{1}{1-\lambda} \sum_{\mathbf{u}} \left(1 + \frac{|S_{\mathbf{u}}| \log n}{n} \right) \prod_{i=1}^{d-1} \left\{ \sum_{0 \leq t < 3t_0} (P^t)_{u_i u_{i+1}} + \frac{1}{(1-\lambda)|S_{\mathbf{u}}|} \right\}. \quad (12)$$

Writing (12) as $\mathbf{E}[\Phi] \ll (1-\lambda)^{-1} \sum_{\mathbf{u}} (1 + |S_{\mathbf{u}}|(\log n)/n) B_{\mathbf{u}}$, we begin with the sum $\sum_{\mathbf{u}} B_{\mathbf{u}}$. Expanding the $(d-1)$ -fold product $B_{\mathbf{u}}$ produces 2^{d-1} terms of the form

$$\left(\frac{1}{1-\lambda} \right)^j \frac{1}{|S_{\mathbf{u}}|^j} \prod_{i \in L} \sum_{0 \leq t < 3t_0} (P^t)_{u_i u_{i+1}} = \sum_{\substack{\theta=(t_i)_{i \in L} \\ 0 \leq t_i < 3t_0}} \left(\frac{1}{1-\lambda} \right)^j \frac{1}{|S_{\mathbf{u}}|^j} \prod_{i \in L} (P^{t_i})_{u_i u_{i+1}} \quad (13)$$

where $L \subseteq \{1, \dots, d-1\}$ and $j + |L| = d-1$. Let

$$C_{\mathbf{u}}^{L, \theta} = \left(\frac{1}{1-\lambda} \right)^j \frac{1}{|S_{\mathbf{u}}|^j} \prod_{i \in L} (P^{t_i})_{u_i u_{i+1}},$$

so $\sum_{\mathbf{u}} B_{\mathbf{u}} = \sum_{\mathbf{u}} \sum_L \sum_{\theta} C_{\mathbf{u}}^{L, \theta}$. The index set L specifies the parameters of a μ -thread (except for its time sequence). Indeed, break $1, \dots, d$ into $\mu = j+1$ intervals by applying the rule that i and $i+1$ are in the same interval precisely if $i \in L$. In Figure 2, $d = 11$, $\mu = 5$, $j = 4$, $L = \{1, 2, 4, 7, 9, 10\}$, and the threads are $[1, 2, 3]$, $[4, 5]$, $[6]$, $[7, 8]$, $[9, 10, 11]$. All we can say about the time sequences is that the total number of elements t_1, t_2, \dots in all of them is exactly $|L| = d - \mu$. We use the superscripts L, θ in the sums to indicate a fixed L or a fixed time sequence θ (or both).

$$\sum_{\mathbf{u}} B_{\mathbf{u}} = \sum_L \sum_{\theta} \sum_{\mathbf{u}} C_{\mathbf{u}}^{L, \theta} = \sum_L \sum_{\theta} \left(\frac{1}{1-\lambda} \right)^j \sum_{\mathbf{u}} \frac{n^{\mu}}{|S_{\mathbf{u}}|^j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}]. \quad (14)$$

Note the presence of the factor n^{μ} to make up for the fact that in a μ -thread each thread starts from a random vertex, whereas in $C_{\mathbf{u}}^{L, \theta}$ each thread starts from the corresponding u_i .




Fig. 2 The index set $L = \{1, 2, 4, 7, 9, 10\}$ defines a 5-thread.

Assume that $j > 0$. Now, the sum $\sum_{\mathbf{u}}^{L, \theta} |S_{\mathbf{u}}|^{-j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}]$ can be upper-bounded using summation by parts:

$$\begin{aligned}
\sum_{\mathbf{u}}^{L, \theta} \frac{\Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}]}{|S_{\mathbf{u}}|^j} &= \sum_{k=1}^n \frac{g_k^{(\mu)}}{k^j} && \text{(group by } |S_{\mathbf{u}}|) \\
&= \frac{g_{\leq n}^{(\mu)}}{n^j} - g_0^{(\mu)} + \sum_{k=1}^{n-1} g_{\leq k}^{(\mu)} \left(\frac{1}{k^j} - \frac{1}{(k+1)^j} \right) && \text{(sum by parts)} \\
&\ll \frac{1}{n^j} + \sum_{k=1}^{n-1} \frac{g_{\leq k}^{(\mu)}}{k^{j+1}} && (g_{\leq n}^{(\mu)} \leq 1) \quad (15) \\
&\ll \frac{1}{n^j} + \frac{1}{n^\mu} \sum_{k=1}^{n-1} \frac{f_0(n/k)}{k^{j+1-\mu}}. && \text{(Lemma 2)}
\end{aligned}$$

Let $\sum_{\mathbf{u}}^L B_{\mathbf{u}}$ denote the sum obtained by collecting all summands in $\sum_{\mathbf{u}} B_{\mathbf{u}}$ with a fixed L . By (14) and using the identity $\mu = j + 1$,

$$\begin{aligned}
\sum_{\mathbf{u}}^L B_{\mathbf{u}} &\ll \sum_{\theta}^L \left(\frac{1}{1-\lambda} \right)^j n^\mu \sum_{\mathbf{u}}^{L, \theta} |S_{\mathbf{u}}|^{-j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}] \quad \text{(by (14))} \\
&\ll (3t_0)^{d-\mu} \left(\frac{1}{1-\lambda} \right)^j \left(n + \sum_{k=1}^{n-1} f_0(n/k) \right) && ((15), \mu = j + 1, \text{ and } |\theta| = d - \mu) \\
&\ll (\log n)^{d-\mu} \left(\frac{1}{1-\lambda} \right)^{d-1} \left(n + \sum_{k=1}^{n-1} \left(\frac{n}{k} \right)^\alpha \right). && \text{(since } f_0(n) = n^\alpha) \quad (16)
\end{aligned}$$

We can now easily cover all cases:

(I) $j > 0$ and $\alpha \leq 1$: $\sum_k (n/k)^\alpha \ll n \log n$; hence, using $\mu = j + 1$,

$$\sum_{\mathbf{u}}^L B_{\mathbf{u}} \ll (1-\lambda)^{1-d} n (\log n)^{d-j} \ll (1-\lambda)^{1-d} n (\log n)^{d-1}.$$

(II) $j > 0$ and $\mu \geq \alpha > 1$: $\sum_k (n/k)^\alpha \ll n^\alpha$; hence

$$\sum_{\mathbf{u}}^L B_{\mathbf{u}} \ll (1-\lambda)^{1-d} n^\alpha (\log n)^{d-\mu} \ll (1-\lambda)^{1-d} n^\alpha (\log n)^{d-\alpha}.$$

(III) $\alpha > \mu$ or $j = 0$: by (14),

$$\sum_{\mathbf{u}}^L B_{\mathbf{u}} \ll (1-\lambda)^{1-d} n^\mu (\log n)^{d-1-j}.$$

If $\alpha > \mu$, then $\sum_{\mathbf{u}}^L B_{\mathbf{u}} = o((1-\lambda)^{1-d} n^\alpha)$. If $j = 0$, then $\sum_{\mathbf{u}}^L B_{\mathbf{u}} \ll (1-\lambda)^{1-d} n (\log n)^{d-1}$.

Since the bounds are independent of L for which there are only constantly many choices, we conclude that

$$\sum_{\mathbf{u}} B_{\mathbf{u}} \ll (1-\lambda)^{1-d} \max \left\{ n (\log n)^{d-1}, n^\alpha (\log n)^{d-\alpha} \right\}.$$

Going back to (12), recall that

$$\mathbf{E}[\Phi] \ll (1-\lambda)^{-1} \sum_{\mathbf{u}} \left(1 + \frac{|S_{\mathbf{u}}| \log n}{n} \right) B_{\mathbf{u}}.$$

To handle $D = \frac{\log n}{n} \sum_{\mathbf{u}} B_{\mathbf{u}} |S_{\mathbf{u}}|$, first note that since $|S_{\mathbf{u}}| \leq n$, we never lose more than a factor of $\log n$. However, in Case (II) we can do better. Assume $j > 0$ and $\mu \geq \alpha > 1$. We revisit the above calculation. With the additional factor of $|S_{\mathbf{u}}|(\log n)/n$, (14) becomes

$$D = \sum_L \sum_{\theta} \log n \left(\frac{1}{1-\lambda} \right)^j \sum_{\mathbf{u}}^{\theta} \frac{n^{\mu-1}}{|S_{\mathbf{u}}|^{j-1}} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}]. \quad (17)$$

First, if $j = 1$, then $\mu = 2$ and $\sum_{\mathbf{u}}^{\theta} n^{\mu-1} |S_{\mathbf{u}}|^{1-j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}] \leq n$. Therefore in this case $D^L \ll (1-\lambda)^{1-d} n (\log n)^{d-1}$, where D^L denotes the sum obtained by collecting all the terms of D with a fixed L .

Next, we consider the case $j \geq 2$. Similarly to (15),

$$\sum_{\mathbf{u}}^{\theta} |S_{\mathbf{u}}|^{1-j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}] \ll \frac{1}{n^{j-1}} + \frac{1}{n^\mu} \sum_{k=1}^n \frac{f_0(n/k)}{k^{j-\mu}}. \quad (18)$$

Hence, (16) becomes

$$\frac{\log n}{n} \sum_{\mathbf{u}}^L B_{\mathbf{u}} |S_{\mathbf{u}}| \ll \sum_{\theta}^L \log n \left(\frac{1}{1-\lambda} \right)^j n^{\mu-1} \sum_{\mathbf{u}}^{\theta} |S_{\mathbf{u}}|^{1-j} \Pr[\mu\text{-thread} \leftrightarrow \mathbf{u}] \quad (\text{by (17)})$$

$$\ll (\log n)^{d+1-\mu} \left(\frac{1}{1-\lambda} \right)^{d-1} \left(n + \frac{k}{n} \sum_{k=1}^n \left(\frac{n}{k} \right)^\alpha \right) \quad (\text{by (18)})$$

$$= (\log n)^{d+1-\mu} \left(\frac{1}{1-\lambda} \right)^{d-1} \left(n + \sum_{k=1}^n \left(\frac{n}{k} \right)^{\alpha-1} \right),$$

and since $\alpha > 1$, $D^L = o((1-\lambda)^{1-d} n^\alpha)$.

Thus, accounting for the extra $\log n$ -factor in Cases (I) and (III),

$$\mathbf{E}[\Phi] \ll (1-\lambda)^{-d} \max \left\{ n (\log n)^d, n^\alpha (\log n)^{d-\alpha} \right\}.$$

This completes the proof of the Master Theorem. \square

4 Extensions

Segment Intersections. The Master Theorem cannot be used for the trapezoidal map of intersecting segments. The reason is that the complexity of an arrangement of n segments depends on both n and the number m of intersections. We show how to extend the Master Theorem to handle this case. The problem can be described by a configuration space that is made of subconfiguration spaces of degrees 3,4,5,6 with $f_0(n,m) = O(n+m)$. We need to strengthen Lemma 2:

Lemma 3 *Let $g_{\leq k}^{(\mu)}$ be as defined in Lemma 2; for any μ -thread, any corresponding time sequence and any $k > 0$,*

$$g_{\leq k}^{(\mu)} \ll (k/n)^\mu O((n+m)/k).$$

Proof We use the same notation as in Lemma 2. We only need a better upper bound on $\sum_{\mathbf{u}} r_{\mathbf{u}}$, the expected complexity of the trapezoidal map for the sample R^v . To do this, we bound the expected number of intersections among the line segments $x_z, z \in R^v$. Let I be an intersection and let x_u be one of its defining segments: I can only be present in the trapezoidal map for R^v if $u \in R^v$. This happens with probability at most $1 - (1 - d/n)^s$: we have s independent samples of d nodes, each of which could be u with probability $1/n$. Since $d/n \leq 1$, we have $1 - (1 - d/n)^s \leq ds/n = 1/k$. By linearity of expectation, it follows that the expected number of intersections is $O(m/k)$, which gives the desired upper bound of $O((n+m)/k)$ on the expected complexity of the trapezoidal map for R^v . Together with the lower bound from the proof of Lemma 2, this completes the proof. \square

The desired result follows now by repeating the proof of the Master Theorem with the bound $g_{\leq k}^{(\mu)} \ll (n+m)/k$ in (15). Then, (16) becomes

$$\sum_{\mathbf{u}}^L B_{\mathbf{u}} \ll (\log n)^{d-\mu} \left(\frac{1}{1-\lambda} \right)^{d-1} \left(n + \sum_{k=1}^{n-1} \frac{n+m}{k} \right),$$

and as in Cases **(I)** and **(III)** we find $\sum_{\mathbf{u}} B_{\mathbf{u}} \ll (1-\lambda)^{1-d} (n+m) (\log n)^{d-1}$. Accounting for the additional $\log n$ -factor (and the factor $(1-\lambda)^{-1}$), this yields $\mathbf{E}[\Phi] \ll (1-\lambda)^{-d} (n+m) (\log n)^d$. To summarize, the m intersections among n segments are computed in time $O((n+m)(\gamma^{-1} \log n)^6)$, as we claimed earlier.

Revisiting the Clarkson-Shor bound. While proving the Master Theorem, we obtained a variant of the Clarkson-Shor bound suited for our Markov model (Lemma 2). We believe that this lemma is of independent interest and could lead to new bounds on the number of k -sets when certain restrictions on the defining elements are imposed. Here is a toy example: let $P \subseteq \mathbb{R}^3$ be a set of n points in general position. Let H be the set of planes in \mathbb{R}^3 spanned by triplets of the form $(x, y, \mathbf{n}(x))$ for $x, y \in P$, where $\mathbf{n}(x)$ denotes a neighbor of x in the Euclidean minimum spanning tree (EMST) of P . A plane $h \in H$ *conflicts* with a point $p \in P$ if p lies below h . Let $f_{\leq k}$ denote the number of planes in H that conflict with at most k points.

Corollary 1 $f_{\leq k} = O(nk)$.

Let H' denote the planes spanned by triplets of the form $(x, y, \mathbf{nn}(x))$, $x, y \in P$, where $\mathbf{nn}(x)$ denotes the nearest neighbor of x in P . Let $f'_{\leq k}$ count the planes in H' with at most k conflicts. Since the EMST contains the nearest neighbor graph [20], we also have

Corollary 2 $f'_{\leq k} = O(nk)$.

Compare this with the well-known Clarkson-Shor bound of $O(nk^2)$ for the unrestricted case.

Proof (of Corollary 1) As our event graph G we take the EMST of P . It is connected and has bounded degree [1, Lemma 4]. Let m be the number of edges in G . We choose $d = 3$ and $\mu = 2$. The first thread has size two with time sequence (1), the second thread has size one. For each thread the probability of picking $v \in V$ as the initial vertex is $\deg(v)/2m$. In other words, the sampling is defined as follows: pick $v \in V$ with probability $\deg(v)/2m$ and take one random step in G . Then pick another random node v according to the same distribution. This yields a triplet of points spanning a plane in H , and each triplet appears with probability $\Theta(1/n^2)$. We have $f_0(n) = O(n)$, since every plane that is spanned by a triplet in P^3 and has no conflicts supports a facet of the lower convex hull of P , and since the number of such facets is $O(n)$ and each facet is supported by exactly one plane. Thus, by Lemma 2, the probability of sampling a plane in conflict with at most k points is $O((n/k)(k/n)^2) = O(k/n)$. Since every plane is sampled with probability $\Omega(1/n^2)$, the claim follows.

Technically, Lemma 2 applies only to regular graphs, while G has bounded, but possibly varying, degree. However, our discussion easily generalizes to the nonregular case—at a loss of only a constant factor. We will show this in Lemma 4 below. \square

Lemma 4 *Let G be a connected graph with n nodes, m edges, and degree bounded by r , and let f_0 be monotonically increasing. Define μ -threads as in Lemma 2, the only difference being that the initial node of each thread is sampled according to the distribution π with $\pi_v = \deg(v)/2m$. For $k > 0$, any μ -thread, and any corresponding time sequence, we have $g_{\leq k}^{(\mu)} \ll (k/n)^\mu f_0(n/k)$.*

Proof Consider the proof of Lemma 2. We may assume that $k \leq n/2dr$. Set $s = n/dkr$. Then Equation (11) still holds, since for a given $\mathbf{u} \in V^d$, the probability $p_{\mathbf{u}}$ that \mathbf{u} is chosen by a random μ -thread is now of the form $\prod_{1 \leq i \leq \mu} \frac{d_i p_{\mathbf{u},i}}{2m}$, where d_i is the degree of the node in \mathbf{u} corresponding to the initial vertex of the i -th thread, and since by our choice of s we have $d_i p_{\mathbf{u},i} s / 2m \leq r p_{\mathbf{u},i} s / 2(n-1) \leq 1$.

Next, we need to bound the probability that a configuration \mathbf{u} with $|S_{\mathbf{u}}| \leq n/2dr$ is active, given that $\mathbf{u} \in R$. Since we sample according to the stationary distribution of G , each node of a μ -thread lies in $S_{\mathbf{u}}$ with probability at most $r|S_{\mathbf{u}}|/2m \leq r|S_{\mathbf{u}}|/n$. Proceeding as before, we now get

$$\Pr[R^v \cap S_{\mathbf{u}} = \emptyset \mid \mathbf{u} \in R] \geq e^{-2d(s-1)\mu r|S_{\mathbf{u}}|/n},$$

and

$$r_{\mathbf{u}} \gg \left(\frac{n}{dkr}\right)^\mu p_{\mathbf{u}} e^{-2d|S_{\mathbf{u}}|/k}.$$

Thus, as before,

$$\sum_{\mathbf{u}: |S_{\mathbf{u}}| \leq n/2dr} r_{\mathbf{u}} \gg \left(\frac{n}{k}\right)^\mu g_{\leq k}^{(\mu)}$$

and

$$\sum_{\mathbf{u}: |S_{\mathbf{u}}| \leq n/2dr} r_{\mathbf{u}} \leq f_0(ds) = f_0(n/rk) \leq f_0(n/k),$$

since f_0 is monotone. This finishes the extension of Lemma 2 to the bounded degree case. \square

Acknowledgements We wish to thank Alistair Sinclair for helpful discussions. We would also like to thank the anonymous referees for pointing us to additional references and providing numerous suggestions for improving the exposition of this paper.

References

1. Aldous, D., Steele, J.M.: Asymptotics for Euclidean minimal spanning trees on random points. *Probab. Theory Related Fields* **92**(2), 247–258 (1992)
2. Amenta, N., Choi, S., Rote, G.: Incremental constructions con BRIO. In: SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry, pp. 211–219. ACM, New York, NY, USA (2003)
3. Barnes, G., Feige, U.: Short random walks on graphs. *SIAM J. Discrete Math.* **9**(1), 19–28 (1996)
4. de Berg, M., van Kreveld, M., Overmars, M., Schwarzkopf, O.: *Computational Geometry: Algorithms and Applications*. Springer Verlag (2000)
5. Boissonnat, J.D., Devillers, O., Schott, R., Teillaud, M., Yvinec, M.: Applications of random sampling to on-line algorithms in computational geometry. *Discrete Comput. Geom.* **8**(1), 51–71 (1992)
6. Boissonnat, J.D., Teillaud, M.: The hierarchical representation of objects: the Delaunay tree. In: SCG '86: Proceedings of the second annual symposium on Computational geometry, pp. 260–268. ACM, New York, NY, USA (1986)
7. Boissonnat, J.D., Teillaud, M.: On the randomized construction of the Delaunay tree. *Theoret. Comput. Sci.* **112**(2), 339–354 (1993)
8. Boissonnat, J.D., Yvinec, M.: *Algorithmic geometry*. Cambridge University Press, New York, NY, USA (1998)
9. Broder, A.Z., Karlin, A.R.: Bounds on the cover time. *J. Theoret. Probab.* **2**(1), 101–120 (1989)
10. Chassaing, P.: Optimality of move-to-front for self-organizing data structures with locality of references. *Ann. Appl. Probab.* **3**(4), 1219–1240 (1993)
11. Cheong, O., Mulmuley, K., Ramos, E.A.: Randomization and derandomization. In: J.E. Goodman, J. O'Rourke (eds.) *Handbook of discrete and computational geometry*, 2nd edn., chap. 40, pp. 895–926. CRC Press, Inc., Boca Raton, FL, USA (2004)
12. Chew, L.P.: Building Voronoi diagrams for convex polygons in linear expected time. Tech. Rep. PCS-TR90-147, Dartmouth College, Hanover, NH, USA (1990)
13. Chung, F.R.K.: *Spectral Graph Theory* (CBMS Regional Conference Series in Mathematics, No. 92). American Mathematical Society, Providence, RI, USA (1997)
14. Clarkson, K.L., Mehlhorn, K., Seidel, R.: Four results on randomized incremental constructions. *Comput. Geom.* **3**(4), 185–212 (1993)
15. Clarkson, K.L., Shor, P.W.: Applications of random sampling in computational geometry. II. *Discrete Comput. Geom.* **4**(5), 387–421 (1989)
16. Devillers, O.: The Delaunay hierarchy. *Internat. J. Found. Comput. Sci.* **13**, 163–180 (2002)
17. Devillers, O., Guigue, P.: The shuffling buffer. *Int. J. Comput. Geometry Appl.* **11**(5), 555–572 (2001)
18. Devillers, O., Meiser, S., Teillaud, M.: Fully dynamic Delaunay triangulation in logarithmic expected time per operation. *Comput. Geom.* **2**(2), 55–80 (1992)
19. Dobrindt, K., Yvinec, M.: Remembering conflicts in history yields dynamic algorithms. In: *Algorithms and computation (Hong Kong, 1993)*, *Lecture Notes in Comput. Sci.*, vol. 762, pp. 21–30. Springer, Berlin (1993)
20. Eppstein, D., Paterson, M.S., Yao, F.F.: On nearest-neighbor graphs. *Discrete Comput. Geom.* **17**(3), 263–282 (1997)
21. Fox, D., Burgard, W., Thrun, S.: Markov localization for reliable robot navigation and people detection. In: *Selected Papers from the International Workshop on Sensor Based Intelligent Robots*, pp. 1–20. Springer-Verlag, London, UK (1999)
22. Guibas, L.J., Knuth, D.E., Sharir, M.: Randomized incremental construction of Delaunay and Voronoi diagrams. *Algorithmica* **7**(4), 381–413 (1992)
23. Horn, R.A., Johnson, C.R.: *Matrix analysis*. Cambridge University Press, Cambridge (1990)
24. Hotz, G.: Search trees and search graphs for Markov sources. *Elektronische Informationsverarbeitung und Kybernetik* **29**(5), 283–292 (1993)
25. Jensen, H.W.: *Realistic image synthesis using photon mapping*. A K Peters Ltd., Natick, MA (2001)
26. Kapoor, S., Reingold, E.M.: Stochastic rearrangement rules for self-organizing data structures. *Algorithmica* **6**(2), 278–291 (1991)
27. Karlin, A.R., Phillips, S.J., Raghavan, P.: Markov paging. *SIAM J. Comput.* **30**(3), 906–922 (2000)

28. Konneker, L.K., Varol, Y.L.: A note on heuristics for dynamic organization of data structures. *Inf. Process. Lett.* **12**(5), 213–216 (1981)
29. Lam, K., Leung, M.Y., Siu, M.K.: Self-organizing files with dependent accesses. *J. Appl. Probab.* **21**(2), 343–359 (1984)
30. Lu, P., Zeng, X., Huang, X., Wang, Y.: Navigation in 3D game by Markov model based head pose estimating. In: *ICIG '04: Proceedings of the Third International Conference on Image and Graphics*, pp. 493–496. IEEE Computer Society, Washington, DC, USA (2004)
31. Matoušek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. *Algorithmica* **16**(4-5), 498–516 (1996)
32. Mehlhorn, K., Sharir, M., Welzl, E.: Tail estimates for the efficiency of randomized incremental algorithms for line segment intersection. *Comput. Geom.* **3**, 235–246 (1993)
33. Motwani, R., Raghavan, P.: *Randomized algorithms*. Cambridge University Press, Cambridge (1995)
34. Mulmuley, K.: A fast planar partition algorithm. I. *J. Symbolic Comput.* **10**(3-4), 253–280 (1990)
35. Mulmuley, K.: A fast planar partition algorithm. II. *J. ACM* **38**(1), 74–103 (1991)
36. Mulmuley, K.: *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice-Hall, Englewood Cliffs (1994)
37. Mulmuley, K.: Randomized geometric algorithms and pseudorandom generators. *Algorithmica* **16**(4-5), 450–463 (1996)
38. Phatarfod, R.M., Pryde, A.J., Dyte, D.: On the move-to-front scheme with Markov dependent requests. *J. Appl. Probab.* **34**(3), 790–794 (1997)
39. Schulz, F., Schömer, E.: Self-organizing data structures with dependent accesses. In: *ICALP*, pp. 526–537 (1996)
40. Schwarzkopf, O.: Dynamic maintenance of geometric structures made easy. In: *Proceedings of the 32nd annual symposium on Foundations of computer science*, pp. 197–206. IEEE Computer Society Press, Los Alamitos, CA, USA (1991)
41. Seidel, R.: A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Comput. Geom.* **1**, 51–64 (1991)
42. Seidel, R.: Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.* **6**(5), 423–434 (1991)
43. Seidel, R.: Backwards analysis of randomized geometric algorithms. In: *New trends in discrete and computational geometry, Algorithms Combin.*, vol. 10, pp. 37–67. Springer, Berlin (1993)
44. Seidel, R., Aragon, C.R.: Randomized search trees. *Algorithmica* **16**(4-5), 464–497 (1996)
45. Shedler, G.S., Tung, C.: Locality in page reference strings. *SIAM Journal on Computing* **1**(3), 218–241 (1972)
46. Veach, E., Guibas, L.J.: Metropolis light transport. In: *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pp. 65–76. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA (1997)
47. Wellington, C., Courville, A., Stentz, A.T.: A generative model of terrain for autonomous navigation in vegetation. *Int. J. Rob. Res.* **25**(12), 1287–1304 (2006)
48. Welzl, E.: Smallest enclosing disks (balls and ellipsoids). In: *New results and new trends in computer science (Graz, 1991), Lecture Notes in Comput. Sci.*, vol. 555, pp. 359–370. Springer, Berlin (1991)

A Markov Chains

A Markov chain M over a finite *state space* Q is an infinite sequence of random variables X_0, X_1, \dots with the following properties: (i) $X_t \in Q$ for $t \geq 0$; (ii) X_0 is drawn from a given initial distribution π_0 over Q ; and (iii) there are $p_{qr} \in [0, 1]$, $q, r \in Q$, such that $\Pr[X_{t+1} = q_{t+1} | X_0 = q_0, X_1 = q_1, \dots, X_t = q_t] = p_{q_t q_{t+1}}$ for $t \geq 0$, ie, the distribution of X_{t+1} depends only on X_t . The variable X_t is called the *state* at time t . The $|Q| \times |Q|$ matrix P formed by the p_{qr} is called the *transition matrix* of M . The distribution of X_t can be computed as $\pi_0^t P^t$. We say that M is *irreducible* if for any two states $q, r \in Q$ there exists $t \geq 0$ such that $\Pr[X_t = r | X_0 = q] > 0$, ie, every state can reach all other states with positive probability after a finite number of steps. A state $q \in Q$ is *periodic* if there exists an integer $\Delta > 0$ such that $\Pr[X_t = q | X_0 = q] = 0$, unless t is a multiple of Δ . Furthermore, q is *non-null persistent* if

$$\Pr[\exists t > 0 : X_t = q | X_0 = q] = 1$$

and

$$\sum_{t>0} t \Pr[X_t = q | X_0 = q, X_1, \dots, X_{t-1} \neq q] < \infty,$$

ie, if every state is revisited with probability 1 after a finite number of steps. The chain is *aperiodic* if none of its states is periodic. It is *ergodic* if it is aperiodic and if all its states are non-null persistent. Any finite, irreducible, aperiodic Markov chain is ergodic. This implies that it has a unique *stationary distribution* π , ie, there exists a unique distribution π with $\pi^T = \pi^T P$. Finally, M is *reversible* if there is a distribution π such that for any $q, r \in Q$ we have $\pi_q p_{qr} = \pi_r p_{rq}$.

Given an undirected graph $G = (V, E)$ and an initial distribution π_0 on V , a *random walk* on G is a sequence of vertices v_0, v_1, \dots , where v_0 is chosen according to π_0 and v_{t+1} is found by following a random edge out of v_t . A random walk induces a Markov chain with state space V . This chain is always reversible. It is irreducible if and only if G is connected, and aperiodic if and only if G has no bipartite components. Thus, any connected, non-bipartite graph induces an ergodic Markov chain. In this case, any random walk converges to the stationary distribution given by $\pi_v = \deg(v)/2|E|$ for $v \in V$. In particular, π is uniform if all vertices have the same degree.

B Matrix Theory

We recall some basic facts from matrix theory [23]. Let $A \in \mathbb{R}^{n \times n}$. We say that A is *symmetric* if $A^T = A$. A symmetric matrix is *positive semidefinite* if $\mathbf{v}^T A \mathbf{v} \geq 0$ for every $\mathbf{v} \in \mathbb{R}^n$. We call $\mathbf{v} \in \mathbb{R}^n \setminus \mathbf{0}$ an *eigenvector* of A if there exists an *eigenvalue* $\lambda \in \mathbb{R}$ with $A\mathbf{v} = \lambda\mathbf{v}$. Every symmetric matrix has n real eigenvalues $\lambda_1, \dots, \lambda_n$ and a corresponding orthonormal basis $\mathbf{v}_1, \dots, \mathbf{v}_n$ of eigenvectors. They can be characterized as follows [23, Theorem 4.2.11]:

Theorem 3 (Courant-Fisher) *Let $A \in \mathbb{R}^{n \times n}$ be symmetric with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$ and corresponding eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$. Then, for $k = 1, \dots, n$,*

$$\lambda_k = \max_{\substack{\mathbf{v} \in \mathbb{R}^n \setminus \mathbf{0} \\ \mathbf{v} \perp \mathbf{v}_1, \dots, \mathbf{v}_{k-1}}} \frac{\mathbf{v}^T A \mathbf{v}}{\mathbf{v}^T \mathbf{v}}.$$

The Courant-Fisher theorem allows us to relate the eigenvalues of any principal submatrix of A to those of A [23, Theorem 4.3.15].

Theorem 4 (Interlacing Theorem) *Let $A \in \mathbb{R}^{n \times n}$ be symmetric with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$, and let A_r be obtained from A by deleting $n - r$ rows and the corresponding columns from A . Let $\mu_1 \geq \dots \geq \mu_r$ be A_r 's eigenvalues. Then, for $1 \leq k \leq r$,*

$$\lambda_k \geq \mu_k \geq \lambda_{n-r+k}.$$

For $A \in \mathbb{R}^{n \times n}$, let G_A be the directed graph on $\{1, \dots, n\}$ which contains an edge from i to j precisely if $A_{ij} \neq 0$. We call A *irreducible* if G_A is strongly connected. If all entries of A are nonnegative, we can say more about its principal eigenvalue and eigenvector [23, Theorem 8.4.4].

Theorem 5 (Perron-Frobenius) *Let $A \in \mathbb{R}^{n \times n}$ be nonnegative, irreducible, and symmetric with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$. Then $\lambda_1 > \lambda_2$, and λ_1 has an eigenvector with all positive entries.*

C Configuration Spaces

A *configuration space* of degree d over a universe \mathcal{U} is a set C of *configurations*. A configuration $\sigma \in C$ is a pair (D_σ, S_σ) , where $D_\sigma, S_\sigma \subseteq \mathcal{U}$ are disjoint with $|D_\sigma| \leq d$. D_σ are the *triggers* and S_σ the *stoppers* of σ . Given a subset $U \subseteq \mathcal{U}$, we say σ is *active* in U if $D_\sigma \subseteq U$ and $S_\sigma \cap U = \emptyset$. The configuration space framework is powerful enough to capture many geometric construction problems, as we will explain below.

The generic construction problem can be phrased as follows: given $U \subseteq \mathcal{U}$, find all active configurations in U . The randomized incremental construction (RIC) paradigm solves this problem by picking a random permutation of U and inserting the elements one by one, creating and destroying configurations according to which trigger and stopper sets contain the newly inserted element. In order to locate the conflicting configurations for the new element quickly, the RIC maintains a *conflict graph* \mathcal{C} , ie, a bipartite graph representing the conflicts between the currently active configurations and the elements in U that still need to be processed. The graph \mathcal{C} is updated after each insertion, and in our examples this takes time linear in the number of edges in \mathcal{C} that are modified. Algorithms that rely on \mathcal{C} are *static*, since all objects in U need to be known in advance. The influence or history graph [8, 36] keeps track of all the configurations that have been active




Fig. 3 (a) A trapezoid is defined by 2,3 or 4 line segments. (b) The handle e of a racquet (τ, e) is defined by 1 or 2 segments. (c) In an opaque representation, the trapezoid σ is incident to 4 vertices, in a planar graph representation, it is incident to 10 vertices.

in the construction so far and stores information about their adjacencies that makes it possible to postpone the conflict updates for an element until it is inserted, with the same asymptotic cost. Thus, algorithms that use this structure are *online*, ie, they do not need to know the input beforehand. Using any of the above data structures, the expected running time of RIC for our examples is

$$\Theta = \sum_{\sigma \in C} |S_\sigma| \Pr[\sigma \text{ becomes active during the construction}].$$

We call this sum the Θ -series of the RIC. In this paper, we use the following configuration spaces:

- CONVEX HULLS IN \mathbb{R}^d AND VORONOI DIAGRAMS IN \mathbb{R}^{d-1} [36, Example 3.4.2]: Let $P \subseteq \mathbb{R}^d$ be in general position. The set C consists of all open half-spaces σ whose bounding hyperplane is spanned by a d -tuple D_σ of distinct points in P . The stopper set S_σ contains all points in $P \cap \sigma$. Clearly, the active configurations in a subset $U \subseteq P$ correspond to the facets of the convex hull of U . Note that a d -tuple D_σ defines two half-spaces, but this can be disambiguated using the ordering of D_σ . By a standard reduction this configuration space also handles Voronoi diagrams in \mathbb{R}^{d-1} .
- TRAPEZOIDAL MAPS [36, Example 3.4.1]: Let L be a set of nonintersecting planar line segments in general position. To avoid unbounded trapezoids, we assume a large bounding box that contains L . The set C consists of all trapezoids σ that can be defined by a set D_σ of line segments in L (and parts of the bounding box). It is easily seen that $|D_\sigma| \in \{2, 3, 4\}$ (see Figure 3(a)). The stopper set S_σ contains the line segments in L that cross the interior of σ . Again, a tuple D_σ may define more than one trapezoid, which we disambiguate with the ordering information of D_σ . Note that this configuration space gives an *opaque representation* of the map: each trapezoid is incident to at most 6 vertices, even though its bounding segments may be subdivided by trapezoids on the other side (see Figure 3(c)). Since L is nonintersecting, this is sufficient to capture the running time of the RIC, because newly inserted segments cross only vertical boundaries of trapezoids that are destroyed.
- SEGMENT INTERSECTIONS [36, Example 3.4.4]: Let L be a set of planar line segments in general position. Again, we assume a large bounding box for L . Since now a newly inserted segment can cross other line segments, we need a planar graph representation of the trapezoidal map. This is achieved using *racquets*, ie, pairs $\sigma = (\tau, e)$, where τ is a trapezoid and e a vertical attachment. The endpoint of e is defined by 1 or 2 segments, while τ is defined by 2, 3, 4 segments (see Figure 3(a,b)). Thus, $|D_\sigma| \in \{3, 4, 5, 6\}$. The stopper set S_σ of a racquet $\sigma = (\tau, e)$ contains all segments in L that intersect τ or e .