Evasive Random Walks and the Clairvoyant Demon

Aaron Abrams[*], Henry Landau[†], Zeph Landau[‡], James Pommersheim[§] and Eric Zaslow[¶]

Abstract

A pair of random walks (R, S) on the vertices of a graph G is *successful* if two tokens can be scheduled (moving only one token at a time) to travel along R and S without colliding. We consider questions related to P. Winkler's *clairvoyant demon problem*, which asks whether for random walks R and S on G, $Pr[\ (R,S) \text{ is successful }] > 0$. We introduce the notion of an *evasive* walk on G: a walk S so that for a random walk R on G, $Pr[\ (R,S) \text{ is successful }] > 0$. We characterize graphs G having evasive walks, giving explicit constructions on such G. Also, we show that on a cycle, the tokens must collide quickly with high probability.

1 Introduction

In 1993, Coppersmith, Tetali, and Winkler [2] considered the following question about random walks on graphs. Suppose two tokens are taking random walks on a finite, simple, connected graph G; this means that each token sits at a vertex of G and, when instructed to move, advances to a (uniformly) randomly chosen adjacent vertex. A scheduling demon, whose goal is to keep the tokens from colliding, chooses at each time step which one of the tokens is to move. The question, then, is can the demon prevent the tokens from ever colliding? The main theorem of Coppersmith, Tetali, and Winkler is that if the demon has no knowledge other than the current locations of the tokens, then regardless of the demon's strategy, the tokens will collide in expected time at most $(4/27 + o(1))n^3$, where n is the number of vertices of G.

*Math Department, UC Berkeley, Berkeley CA 94720, abrams@math.berkeley.edu

†hjl@research.bell-labs.com

‡Sloan Center for Theoretical Neurobiology, Dept. of Physiology, Box 0444, University of California, San Francisco, 513 Parnassus Avenue, San Francisco CA, 94122, landau@math.berkeley.edu

§Department of Mathematics, Pomona College, 610 N. College Ave., Claremont CA 91711-6348, jpommersheim@POMONA.EDU

¶Department of Mathematics, Northwestern University, 2033 Sheridan Road, Evanston, IL 60208, zaslow@math.nwu.edu

This result led Winkler to ask whether the demon can do better with more information. Specifically, suppose the demon knows from the outset which (infinite) path will be travelled by each token. Can this *clairvoyant demon* succeed at preventing a collision forever?

Definition 1.1. Let G be a graph, and let R and S be two (infinite) walks on G. If two tokens T_R and T_S can be scheduled to travel along R and S without colliding, we call (R, S) a *successful pair* of walks.

For example, if G is the complete graph on the vertices A, B, C, D then any pair (R, S) where R begins $AB....$ and S begins $BA....$ is not successful, as the demon has no first move. By contrast the pair $ABABAB...$ and $CDCDCD...$ is obviously successful. Note that in general, as in the first example, if a pair is not successful then this can be detected by looking at some finite initial strings of the two walks. Thus, in particular, on a finite graph there is a positive chance that a random pair is *not* successful. Detecting success of a pair, however, is not so easy, and it is by no means obvious that the probability of a pair's being successful should be positive.

Definition 1.2. We say that the *clairvoyant demon succeeds on G* if for randomly chosen walks R and S on G, (R, S) is a successful pair with positive probability.

Thus the Clairvoyant Demon problem is to determine, for a graph G, whether the clairvoyant demon succeeds on G. It is obvious that the demon does not succeed on a path; Winkler [5] has shown that the demon doesn't succeed on a cycle either. It is not known whether there is any graph on which the demon succeeds, but Winkler has conjectured that the demon succeeds on the complete graph K_4.

The clairvoyant demon problem has a natural reformulation in the language of percolation; see [5]. From this point of view, Winkler [5] and independently Balister, Bollobás, and Stacey [1] have shown that the *fickle demon*, who is allowed to move the tokens both forward and backward along their walks, does succeed on K_4. Also Gács [3] has shown that on a graph where the clairvoyant demon succeeds, the probability of being able to advance the tokens for n steps but not indefinitely is bounded below by n^{-a}, for some $a > 0$, rather than decreasing exponentially as might have been expected.

In this paper we rephrase the clairvoyant demon problem as a problem about single random walks, rather than pairs. Specifically, let S be a fixed walk. If, when R is chosen randomly, the pair (R, S) is successful with positive probability, then we call S *evasive*. If a random choice of S is evasive with positive probability, then the clairvoyant demon succeeds on G. We therefore focus on evasive walks.

The main result of this paper is

Theorem 1.3. *Let G be a finite, simple, connected graph. Then the following are equivalent:*

(i) there exists an evasive walk on G;

(ii) there exist uncountably many evasive walks on G;

(iii) G is not a path, a cycle, or the graph $K_{1,3}$ consisting of one vertex of degree 3 and three pendant vertices.

Section 2 is devoted to constructions (Theorems 2.5, 2.9, 2.11) of evasive walks with various properties. In particular Theorem 2.11 shows that on the complete graph K_n, $n \geq 5$, there are evasive walks which are random to first order. In Section 3 we show that on a cycle, the chance that the demon can advance the tokens N steps without collision shrinks exponentially in $N^{1/3}$. This can be viewed as a partial converse to Gács's result; it also implies that there are no evasive walks on a cycle. The proof of Theorem 1.3 is completed by combining Theorem 2.5 and Corollary 3.5.

2 Evasive walks

All graphs in this paper are finite and simple (so they have no loops or multiple edges). By a walk R on G we mean a sequence $R(1), R(2), \ldots$ of vertices of G such that $R(i)$ is adjacent to $R(i+1)$ for each i. We also sometimes think of R as a function from $\mathbb{N}$ to the vertex set of G, using notation such as $R^{-1}(v)$ to denote the set of "times" when R visits the vertex v.

Definition 2.1. A walk R on G is *evasive* if for a randomly chosen walk S on G the pair (R, S) is successful with positive probability.

Remark 2.2. By Fubini's theorem, the clairvoyant demon succeeds on G if and only if the set of evasive walks on G has positive measure in the set of all walks on G.

Thus the clairvoyant demon problem is equivalent to asking if a random walk on G is evasive with positive probability. In this section we give several constructions of evasive walks on graphs, though we remain unable to construct a set of such walks of positive measure on any graph.

We will also have occasion to think about finite walks. A *finite walk* is a finite initial segment $R(1), \ldots, R(n)$ of an infinite walk; its *length* is the number n of vertices it visits (which may be zero in the case of the empty walk). Note that our length is one more than the usual definition of length, which is the number of edges traversed.

Definition 2.3. Let R and S be finite walks on G, with $length(S) > 0$. If tokens T_R and T_S can be advanced to the end of R and S without colliding, then we call S an *allowed walk* for R.

If G is a path, the walk R which traverses G from end to end has no allowed walks. It is not hard to see that this is the only such case:

Remark 2.4. If H is a connected subgraph of G that is not a path, then every finite walk R on G has an allowed walk on H.

If G is a graph and v a vertex of G, then $G - v$ denotes the graph obtained from G by deleting v and all edges incident with v. (This is the subgraph induced by the vertices of G other than v.)

We now give the constructions. Let K_n denote a complete graph on n vertices; we note that our first result applies already to K_4.

Theorem 2.5. *Let G be a connected graph which is not a path, a cycle, or the graph $K_{1,3}$. Then there exist uncountably many evasive walks on G.*

Remark 2.6. Winkler [5] has shown that the fickle demon succeeds on G if and only if G is not a path, a cycle, or $K_{1,3}$.

For the proof we record the following easy characterization.

Remark 2.7. Let G be a connected graph. Then G is not a path, a cycle, or $K_{1,3}$ if and only if there exists a vertex v of G such that some component of $G - v$ is not a path.

Proof of Theorem 2.5. Choose v according to Remark 2.7, and let G' be a component of $G - v$ which is not a path. Choose S to be any walk on G' with the property that each finite walk on G' occurs as a subwalk of S infinitely many times. (A random walk on G' has this property with probability one.) We will show that S is evasive.

Let n be the number of vertices of G; then with probability $1/n$ a random walk R on G has $R(1) = v$ and $R^{-1}(v)$ infinite. Choosing R with these properties, set $R^{-1}(v) = \{1, i_2, \ldots\}$ with $i_j < i_{j+1}$. Consider the finite walk $R(2), R(3), \ldots, R(i_2 - 1)$ on G. By Remark 2.4, this walk has an allowed walk on G', which by choice of S occurs as a subwalk of S. While T_R sits on $R(1) = v$, advance T_S to the beginning of this allowed sequence. Then advance T_R and T_S until T_R can return to v, and repeat the process.

□

The set of evasive walks arising from the construction of Theorem 2.5 has measure zero, since each such walk misses some vertex of G. We give a more elaborate construction to improve this situation, after first introducing some more notation.

For G' a connected subgraph of a graph G, let

$$\langle G' \rangle_l$$

denote any finite walk S on G' with the property that for every finite walk R on G of length $\leq l$, S contains a finite subwalk allowed for R. Such walks exist by Remark 2.4, provided that G' is not a path. For vertices $v_1, \ldots, v_n$ we write

$$\langle v_1 \cdots v_n \rangle_l$$

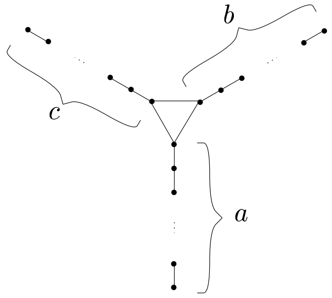

Figure 1: The graph Y_{abc}.

to denote $\langle G' \rangle_l$ where G' is the subgraph of G induced by the v_i.

For example, on K_4 (with vertices A, B, C, D), the symbol $\langle ABC \rangle_3$ could be used to denote the walk $ABCACBA$, as this contains subwalks allowed for each of the 53 possible walks of length at most 3 on K_4.

Definition 2.8. Denote by Y_{abc} the graph consisting of a triangle with paths of lengths a, b, c attached to the vertices. See Figure 1.

The next construction is slightly complicated; it may help the reader to first pretend that $G = K_4$.

Theorem 2.9. *Let G be a connected graph which is not a tree, a cycle, or one of the graphs Y_{abc}. Then there exist uncountably many evasive walks on G which visit each vertex of G infinitely many times.*

For the proof we again use a characterization.

Remark 2.10. Let G be a connected graph. Then G is not a tree, a cycle, or one of the graphs Y_{abc} if and only if there exists a vertex v of G such that v is connected by at least two edges to some non-path component G' of $G - v$.

Proof of Theorem 2.9. Choose v and G' according to Remark 2.10, and let $w \in G'$ be a neighbor of v. Suppose first that $G - v$ is connected. Set

$$S = w\langle G' \rangle_{l_{11}} wvw \langle G' \rangle_{l_{21}} \langle G' \rangle_{l_{22}} wvw \cdots,$$

with the integers l_{ij} to be determined later. (In order to ensure that S is a walk, we extend the ends of each $\langle G' \rangle_{l_{ij}}$ if necessary so that they start and end at a neighbor of w in G'.) Clearly S visits each vertex of G infinitely many times; we show that for suitable choices of the l_{ij}, S is an evasive walk.

Again let n be the number of vertices of G; then with probability $1/n$ a random walk R starts at v and visits it infinitely often. For such an R, let $R^{-1}(v) = \{i_1, i_2, \ldots\}$ where $i_j < i_{j+1}$ (thus $i_1 = 1$), and let $d_j = i_{j+1} - i_j - 1$. Likewise, let $S^{-1}(v) = \{n_1, n_2, \ldots\}$ denote the successive visits of S to v.

Let E_j be the event that $R(i_j+1)=w$ or $R(i_{j+1}-1)=w$. Let $p_l = Pr\{d_j \le l \mid E_j\}$ and $q = 1 - Pr\{E_j\}$; note that these are independent of j and that $p_l \to 1$ as $l \to \infty$.

To begin, T_R is on $R(1)=v$ and T_S is on $S(1)=w$. We say that the tokens are *reset* initially, and again whenever T_R is on v and T_S is on $S(i+1)=w$, where $S(i)=v$. We compute a lower bound on the probability of being able to reset the tokens.

Case 1: If $R(2) \ne w$ and $R(i_2-1) \ne w$, then with T_R held on v we can advance T_S through $\langle G' \rangle_{l_{11}}$ to $S(n_1-1)=w$, move T_R to $R(2) \ne w$, T_S to $S(n_1)=v$, T_R to $R(i_2-1) \ne w$, T_S to $S(n_1+1)=w$, and T_R to $R(i_2)=v$, resetting the tokens.

Case 2: If $R(2)=w$ or $R(i_2-1)=w$ but $d_1 \le l_{11}$, $R(i_2+1) \ne w$, and $R(i_3-1) \ne w$ then with T_R on $R(1)=v$ we move T_S to the beginning of an allowed walk in $\langle G' \rangle_{l_{11}}$ for $R(2)R(3)\cdots R(i_2-1)$. We then advance T_R to $R(i_2)=v$ and T_S through the remainder of $\langle G' \rangle_{l_{11}}$ to $S(n_1-1)=w$. Then as before, move T_R to $R(i_2+1) \ne w$, T_S to $S(n_1)=v$, T_R to $R(i_3-1) \ne w$, T_S to $S(n_1+1)=w$, and T_R to $R(i_3)=v$, resetting the tokens.

The probability of Case 1 or 2 occurring is at least $q+(1-q)qp_{l_{11}}$.

We now try to reset the tokens again. Note that T_R is now on $R(i_2)$ or $R(i_3)$; as it makes no difference for the argument, we suppose it is on $R(i_2)$. Again, the tokens can be reset if $R(i_2+1) \ne w$ and $R(i_3-1) \ne w$ (as in Case 1); or if $R(i_2+1)=w$ or $R(i_3-1)=w$, but $d_2 \le l_{21}$, $R(i_3+1) \ne w$, and $R(i_4-1) \ne w$ (as in Case 2, using an allowed walk from $\langle G' \rangle_{l_{21}}$).

The new case is that $R(i_2+1)=w$ or $R(i_3-1)=w$, and $R(i_3+1)=w$ or $R(i_4-1)=w$, but $d_2 \le l_{21}$, $d_3 \le l_{22}$, $R(i_4+1) \ne w$, and $R(i_5-1) \ne w$. Here we proceed similarly, but we must use allowed walks from both $\langle G' \rangle_{l_{21}}$ and $\langle G' \rangle_{l_{22}}$.

These three cases occur with probability at least $q+(1-q)qp_{l_{21}}+(1-q)^2qp_{l_{21}}p_{l_{22}}$.

Continuing, we see that we can reset the tokens infinitely many times with probability

$$p \ge \left(q+(1-q)qp_{l_{11}}\right)\left(q+(1-q)qp_{l_{21}}+(1-q)^2qp_{l_{21}}p_{l_{22}}\right)\cdots,$$

thereby making (R,S) a successful pair if this product converges. The reader can verify that if each $p_{l_{ij}}$ were equal to 1, then the product would converge (recall that $q>0$ is constant). Thus, since $p_l \to 1$ as $l \to \infty$, we can choose the l_{ij} so that S is evasive, and the conclusion follows.

It remains to consider the case in which $G-v$ is not connected. Here in order to make S visit each vertex of G infinitely many times we replace each occurrence of v in S with a finite walk on $G-G'$ starting and ending at v and visiting each vertex of $G-G'$. The proof that S is evasive in this case is similar to the proof when $G-G'$ is connected; one has to consider at each stage the additional possibility that from time i_j+1 to $i_{j+1}-1$, the walk R may not be on G'. We omit the details.

□

The set of evasive walks arising from the construction of Theorem 2.9 still has measure zero, since the set of times that each such walk visits the vertex v has upper density zero. (Recall that the *upper density* of a set of integers X is $\limsup_{n\to\infty} \frac{|X\cap\{1,\ldots,n\}|}{n}$.) This is nevertheless the best we can do on K_4. However, on K_n for $n \geq 5$ we have the following improvement.

Theorem 2.11. *For $n \geq 5$, there exist (uncountably many) evasive walks S on K_n for which $S^{-1}(v)$ has upper density at least $1/n$ for each vertex v of G.*

Proof. Denote the vertices of K_n by $v_1, \ldots, v_n$. Let

$$S = \langle v_1v_2v_3\rangle_{l_1}\langle v_2v_3v_4\rangle_{l_2}\cdots\langle v_kv_{k+1}v_{k+2}\rangle_{l_k}\cdots,$$

where the subscripts on the v's are taken modulo n and the l_i are yet to be determined. We will show that the l_i can be chosen so that S is evasive.

Choose R at random, subject to $R(1) = v_{n-1}$. Let $i_1 = 1$ and $i_j = \min\{i > i_{j-1} : R(i) = v_{j-2}\}$ for $j = 2, 3, \ldots$, taking the indices mod n; thus $\{R(i_j)\}$ is the earliest subsequence of R which looks like $v_{n-1}v_nv_1v_2v_3v_4\ldots$. Let $d_j = i_{j+1} - i_j - 1$ be the length of the subwalk between $R(i_j)$ and $R(i_{j+1})$.

We claim that if $d_j \leq l_j$ for all j, then (R, S) is a successful pair. To see this, assume $d_1 \leq l_1$. Then $\langle v_1v_2v_3\rangle_{l_1}$ contains an allowed subwalk for the (possibly empty) subwalk $R(2), R(3), \cdots, R(i_2 - 1)$. Holding T_R at $R(1) = v_{n-1}$, we advance T_S to the start of this allowed walk. We next advance T_R and T_S until T_R arrives at $R(i_2) = v_n$, and finally, keeping T_R fixed, we advance T_S to the end of $\langle v_1, v_2, v_3\rangle_{l_1}$.

This process can be repeated indefinitely provided that each $d_i \leq l_i$; this establishes the claim.

Now, for fixed j, we have $Pr(d_j \leq l_j) = 1 - (\frac{n-2}{n-1})^{l_j}$. Thus

$$Pr((R,S) \text{ is a successful pair}) \geq \prod_j \left(1 - \left(\frac{n-2}{n-1}\right)^{l_j}\right),$$

a product which converges (to a positive number) whenever $\sum_j (\frac{n-2}{n-1})^{l_j}$ does. Thus S is evasive for any choice of $\{l_j\}$ which makes the above sum converge.

To verify the claim about upper densities, we choose $l_j = \lceil j/n \rceil$. The infinite sum clearly converges, and by symmetry we may choose S so that each block

$$\langle v_1v_2v_3\rangle_k\langle v_2v_3v_4\rangle_k\cdots\langle v_nv_1v_2\rangle_k$$

visits each vertex the same number of times. $\square$

An understanding of how the length of $\langle v_1v_2v_3\rangle_k$ grows as a function of k might provide information about the lower densities of the sets $S^{-1}(v)$.

3 Collisions on Cycles

In this section we show that on a cycle, a collision will occur quickly with high probability. This implies that there are no evasive walks on a cycle. The proof of the main theorem is completed as follows.

Proof of Theorem 1.3. That (iii) implies (ii) is the content of Theorem 2.5. Obviously (ii) implies (i). It remains to show (i) implies (iii). There are clearly no evasive walks on a path, and Corollary 3.5 of this section shows that there are no evasive walks on a cycle. We leave it as an exercise to show that the existence of an evasive walk on $K_{1,3}$ would imply the existence of an evasive walk on the cycle K_3. (In general the clairvoyant demon succeeds on K_n iff it succeeds on $K_{1,n}$.) This completes the proof. $\square$

The essential ingredient in obtaining the bound of Theorem 3.4 is a winding number, defined next.

Definition 3.1. Let G be a cycle on n vertices. With (R, S) a pair of infinite walks on G, define the *fractional winding sequence* of R by $f_R(1) = 0$ and

$$f_R(i+1) = \begin{cases} f_R(i) + 1/n \text{ if } R(i+1) \text{ is a clockwise step from } R(i), \text{ or} \\ f_R(i) - 1/n \text{ if } R(i+1) \text{ is a counter-clockwise step from } R(i). \end{cases}$$

Define $f_S(i)$ similarly, except that $f_S(1) = k/n$ where $0 \leq k < n$ is the number of clockwise steps from the start of R to the start of S. The *winding number of R at i* is $w_R(i) = \lfloor f_R(i) \rfloor$, and likewise for S. Define the *winding sequence* of R, W_R, to be the subsequence of $w_R(1), w_R(2), \ldots$ that ignores repetition; i.e. a maximal subsequence of $w_R(1), w_R(2), \ldots$ such that consecutive elements are distinct. Define W_S likewise.

The winding numbers for a successful pair are related:

Remark 3.2. For (R, S) a pair of walks on a cycle, if the clairvoyant demon can advance T_R to $R(i)$ and T_S to $S(j)$ then

$$w_R(i) \leq w_S(j) \leq w_R(i) + 1.$$

The following lemma is used in the proof of Theorem 3.4.

Lemma 3.3. *There exists an integer N and a constant $c > 0$ such that for all $n > N$, a simple random walk of length n^3 on a line remains within n units of its starting point with probability less than or equal to e^{-cn}.*

Proof. For a walk to stay within n units of its starting point it is necessary (though not sufficient) that no subwalk of the walk moves more than $2n$ units in either direction. By the central limit theorem, there is an integer N_0 and a constant $c_0 > 0$ so that for $k > N_0$, a walk of length $c_0 k^2$ remains within k units of its starting point with probability less than e^{-1}. Given a walk of length n^3, we break it into pieces of length $c_o(4n^2)$ and apply the previous fact with $k = 2n$ to each of the $n/4c_0$ pieces. It follows that the probability that the walk remains within n units of its starting point is at most $e^{-n/4c_0}$. $\square$

Theorem 3.4. *Let G be a cycle with n vertices. Then there exists a constant $c > 0$ such that for sufficiently large N, the probability that a random pair (R, S) can be advanced N steps without collision is less than $e^{-cN^{\frac{1}{3}}}$.*

Proof. Let R and S be walks on G, and let N be as in Lemma 3.3. Since $f_R(i)$ is a random walk on the integer multiples of $1/n$, Lemma 3.3 provides a constant $c_0 > 0$ such that

$$Pr\left(|w_R(i)| \leq \frac{N^{1/3}}{n} \quad \forall i \leq N\right) \leq e^{-c_0 N^{1/3}}.$$

Ignoring these cases, we find that the winding number exceeds $\frac{1}{n}N^{1/3}$ (for some $i \leq N$) with probability at least $1 - e^{-c_0 N^{1/3}}$. (If R achieves $-\frac{1}{n}N^{1/3}$ as a winding number instead, we replace all winding numbers by their negatives in the ensuing argument.)

We will say that the winding sequence of a walk *decreases at k* if subsequent to the first occurrence of $k+1$, the next occurrence of $k-1$ is before the first occurence of $k+2$, and *increases strongly at k* if from the first occurrence of k it proceeds $k, k+1, k+2$.

By Remark 3.2, any scheduling of the tokens has the property that $W_S \leq k$ until the first time W_R reaches k (at which time $W_S = k$). The key observation is that, again by Remark 3.2, the pair (R, S) is not successful if there is a k at which W_R increases strongly while W_S decreases. In such a case, collision of the tokens is ensured by the time W_R reaches $k+2$. Let E_k be the event that W_R increases strongly at k and W_S decreases at k. Now, there is some $p > 0$ (depending on n) such that $Pr(E_k) > p$ for all k. For the walk to be free of collision for more than N steps, none of the events E_k may occur for $0 \leq k \leq \frac{1}{n}N^{1/3}$. The events E_{2k} are independent (though the E_k aren't); the probability that none of them occurs is at most $(1-p)^{\frac{1}{2n}N^{1/3}} = e^{-cN^{1/3}}$. This completes the proof. $\square$

Corollary 3.5. *There are no evasive walks on a cycle.*

Proof. Let S be a walk on a cycle. If W_S decreases at only a finite number of integers then W_S is bounded below. The pair (R, S) can then be successful only if W_R is bounded below as well, an event of probability zero.

In the remaining case W_S decreases at an infinite set of integers T. By the argument of Theorem 3.4, for (R, S) to be a successful pair, W_R cannot increase strongly at any of the integers in T. As T is infinite, the probability of this event is zero. Consequently S is not an evasive walk. $\square$

4 Acknowledgements

The authors would like to thank Peter Winkler, who introduced us to the clairvoyant demon problem. We also thank Julie Landau, without whose support and cross-court forehand this paper would not have been possible.

References

[1] P. N. Balister, B. Bollobás, and A. N. Stacey, *Dependent percolation in two dimensions*, Probab. Theory Related Fields **117** (2000), no. 4, 495–513.

[2] D. Coppersmith, P. Tetali, and P. Winkler, *Collisions among random walks on a graph*, SIAM J. Discrete Math **6** (1993), no. 3, 363–374.

[3] P. Gács, *The Clairvoyant Demon Has a Hard Task*, to appear in Combinatorics, Probability, and Computing.

[4] R. K. Miller, *Automated Guided Vehicles and Automated Manufacturing*, first edition, Society of Mechanical Engineers, Dearborn, Michigan, 1987.

[5] P. Winkler, *Dependent percolation and colliding random walks*, Random Structures and Algorithms **16** (2000), no. 1, 58–84.