Dependent Percolation and Colliding Random Walks

Peter Winkler*
Bell Labs 2C–365
700 Mountain Ave.
Murray Hill, NJ 07974 U.S.A.
pw@lucent.com

October 12, 1999

Abstract

Let G be a connected, undirected graph and $\mathbf{X} = X_0X_1X_2\dots$ and $\mathbf{Y} = Y_0Y_1Y_2\dots$ two simple random walks on G. Let $\mathbb{N}\Box\mathbb{N}$ be the non-negative quadrant of the plane grid, and H the subgraph of $\mathbb{N}\Box\mathbb{N}$ induced by the sites (i,j) for which $X_i \neq Y_j$. We say that G is "navigable" if with probability greater than 0, the origin belongs to an infinite component of H.

We determine which finite graphs are navigable, in particular that K_4, the complete graph on four nodes, is navigable but K_3 is not. Navigability of G is equivalent to the statement that with positive probability, two tokens taking random walks on G can be moved forward and backward along their paths, and ultimately advanced arbitrarily far, without colliding.

The problem is generalized to finite-state Markov chains, and a complete characterization of navigable chains is given. Similar results have been obtained simultaneously and independently by Balister, Bollobás and Stacey, using different methods; our classification theorem relies on a surprising diamond lemma which may be of independent interest.

1 Introduction

Random walks on graphs are useful, well-behaved and widely-studied processes (see, e.g., [4] for an introduction or [1] for greater depth). Much less is known about the behavior of multiple, simultaneous random walks.

In this work we study a problem concerning two random walks on a graph, which is equivalent to a dependent percolation problem on the plane grid. For background on percolation, we refer the reader to [6]. Relatively few theorems in the literature, however, pertain to *dependent* percolation.

The problem's ancestry can be traced to a question addressed in [3] which arose in asynchronous distributed computing. In a certain token management protocol two tokens take simple random walks on an n-vertex graph G, but a "schedule demon" gets to decide, at each point in time, which token moves next.

The demon sees where each token moves, but only after the move is made; the question was whether the demon can prevent the tokens from colliding in (expected) time polynomial in n. The answer (from [3]) is no; the expected number of steps to collision is at most $\frac{4}{27}n^3$ regardless of the

*Work supported in part by the Institute for Elementary Studies.

demon's strategy, and the consequence is that a certain token management protocol for networks of processors self-stabilizes in polynomial time.

In 1990, while working on the schedule demon, this author was led to the following conjecture: that on any sufficiently complex graph, if the demon is *clairvoyant* (that is, he knows each token's entire future walk) then with positive probability, he can keep the tokens apart *forever*. The problem was disseminated orally at Oberwolfach and elsewhere, and eventually appeared in print in [3]; it seems to have attracted quite a lot of attention. However, at the time of this writing, no finite connected graph is known to have this property, even though it is widely believed to be true even on K_4, the complete graph on four vertices.

Note that we cannot ever expect the clairvoyant demon to win with *certainty*, since e.g. the tokens may be headed straight for each other along the same path. On the other hand, it makes no difference where the tokens start (as long as they are on distinct vertices) because with positive probability the demon can move them anywhere he wants and, in effect, start over.

A weaker, and in some sense more realistic, version of the problem awards the demon a power greater than clairvoyance: the ability to retract errors. The "fickle demon" does not see the future, and indeed may as well move randomly, but she is permitted to move a token *backwards* along its walk. The fickle demon has at most four choices at each tick of the clock: move the first token forward, move it backward, move the second token forward or move it backward. She may not take a move which causes a collision or pushes a token back before time 0; the question is, can she—with positive probability—advance both tokens arbitrarily far in their walks? Like the clairvoyant demon, the fickle demon cannot be expected to win with probability 1, and again the starting positions are unimportant.

When the clairvoyant demon problem was brought to the attention of Noga Alon in 1990, he immediately observed the connection to percolation. In fact both demons can be put in a percolation setting. Since we will be making use of more than one type of graph product, we adopt Jaroslav Nešetřil's idea of using a symbol which pictures the product of two edges; in this notation $\mathbb{N}\Box\mathbb{N}$ is the non-negative quadrant of the plane grid, with two points adjacent when their Euclidean distance is 1. We label the columns of the grid by the vertices $X_0X_1X_2\ldots$ visited by the first token, which we will call the "X-token". Label the rows by the walk $Y_0Y_1Y_2\ldots$ of the Y-token. Now strike out all vertices ("sites") of the grid which received the same label from row and column; these points are collision points for the two tokens.

A particle begins at the origin, which represents the initial state of the system with the X-token at X_0 and the Y-token at Y_0. Each decision by the clairvoyant demon pushes the particle one step east (if the X-token is the one which is moved) or one step north (if it is the Y-token which is moved). The clairvoyant demon wins if he can push the particle to infinity, avoiding the deleted sites; this is precisely the issue of "oriented percolation". It is easy to see that the demon has no chance to win if he can see only finitely far ahead.

In contrast, the fickle demon can move the particle west or south as well as east or north. She can never get stuck (once she gets started); and even a random strategy will suffice to push the token arbitrarily far, as long as there are infinitely many sites connected to the origin. This is the issue of ordinary "undirected" percolation, typically easier to deal with than the oriented case. Figure 1 illustrates a piece of $\mathbb{N}\Box\mathbb{N}$, with lines labeled and sites stricken, and three possible escape routes for the fickle demon.

What takes both demons out of the realm of the ordinary is the process for deleting sites. In ordinary percolation each point is deleted independently with some fixed probability q, but here the following dependence appears: if three corners of a rectangle are deleted then so is the fourth. Since both oriented and undirected *independent* percolation do take place when q is small enough, we are in effect asking whether this rectangular dependence has a detrimental effect.

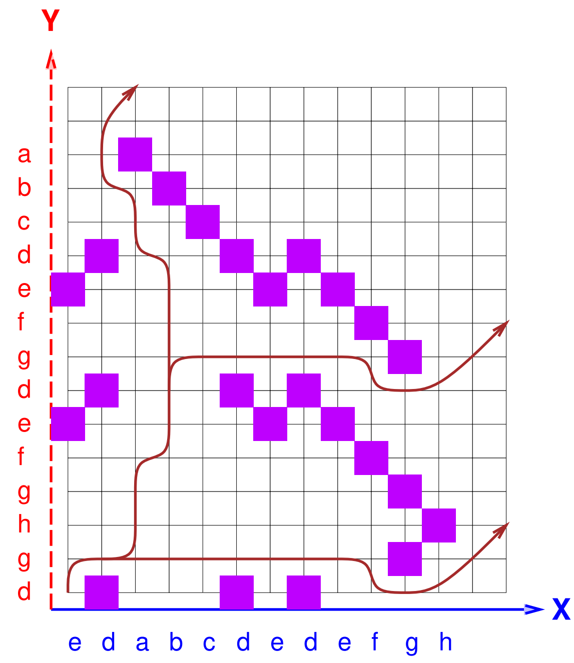

Figure 1: The fickle demon wins (so far) but the clairvoyant demon is blocked

On the one hand, the rectangular dependence is third-order and only affects some triples of sites; on the other hand, our model clearly has far less independence than ordinary percolation because in our model the fates of the sites in (say) an $n \times n$ square region are decided by only $2n$ events.

In fact we will show that undirected percolation does take place in our model, when G is sufficiently rich. Included is the case where G is a looped K_4, meaning that the rows and columns of $\mathbb{N}\square\mathbb{N}$ are independently labeled from a uniform set of four labels; this corresponds to $q = 1/4$ where ordinary undirected percolation also takes place. On the other hand ordinary percolation takes place even at $q = 1/3$ (the threshold is at $q \sim 0.41$—see e.g. [6]) and we will see that the fickle demon does *not* succeed on any 3-vertex graph.

Our approach involves determining precisely what circumstance prevents escape from the rectangle $[0, m] \times [0, n]$ on the plane grid, then showing that unless G is a path, cycle or the specific graph $K_{1,3}$, the probability that some rectangle is blocked is less than one.

For the sake of readability, we begin with the case $G = K_n$ and give a straightforward inductive proof, introducing the concepts of projections and prejections from one walk to another. We then prove a diamond lemma for projections, which enables us to refine the probabilistic part of the argument so as to settle the percolation question for arbitrary G. Finally, we generalize the problem to finite-state Markov chains, determining precisely which cases permit percolation with positive probability. (We have learned that similar results for the fickle demon have been obtained independently by Balister, Bollobás and Stacey [2], using a different approach.)

At the conclusion we make some remarks about the difficulties encountered when our methods are applied to the case of the clairvoyant demon.

2 Blockers

Let $G = \langle V, E\rangle$ be a (finite, undirected, simple, loopless, connected) graph. We denote by boldface $\mathbf{X}$ and $\mathbf{Y}$ the (infinite) walks taken by our two tokens, with $\mathbf{X} = X_0X_1\dots$ and $\mathbf{Y} = Y_0Y_1\dots$; finite initial segments of $\mathbf{X}$ and $\mathbf{Y}$ will be denoted by lightface X and Y.

Let $\mathbb{N}\Box\mathbb{N}$ be the graph whose vertices ("sites") are the integer points in the non-negative quadrant of the Euclidean plane, with $P \sim Q$ when $\rho(P,Q) = 1$, where ρ is ordinary Euclidean distance.

Let H be the subgraph of $\mathbb{N}\Box\mathbb{N}$ induced by the sites (u, v) for which $X_u \neq Y_v$, and let C be the connected component of H to which the origin belongs. If C is finite we say that the pair $\langle\mathbf{X},\mathbf{Y}\rangle$ is *blocked*; otherwise the fickle demon will eventually push the two tokens arbitrarily far along their walks. (It is easy to see that the event that the demon can push *one* token infinitely far while the other remains bounded has probability zero.)

Our object is to show that, for suitable G, the probability that $\langle\mathbf{X},\mathbf{Y}\rangle$ is blocked is less than one. Note that the presence of loops in G would not affect the issue; consecutive rows or columns with the same label can be collapsed to a single row (resp. column) with no effect.

Suppose that $\langle\mathbf{X},\mathbf{Y}\rangle$ is blocked, and let ∂C be the exterior boundary of C, that is,

$$\partial C = \{P = (u, v) :\ X_u = Y_v,\ \text{but } P \sim Q \text{ for some } Q \in C\}\ .$$

Two deleted sites P and Q cannot be adjacent in $\mathbb{N}\Box\mathbb{N}$ since G has no loops. In our notation, $\mathbb{N}\times\mathbb{N}$ is the graph on the same vertices (sites) as $\mathbb{N}\Box\mathbb{N}$ but with P adjacent to Q *diagonally*, i.e. when $\rho(P,Q) = \sqrt{2}$. Then as a subgraph of $\mathbb{N}\times\mathbb{N}$, ∂C has a connected component which contains a point from each of the axes.

This component in turn contains a *path* $\mathcal{P} = P_0P_1\dots P_t = (u_0, v_0)\ (u_1, v_1)\dots(u_t, v_t)$ in $\mathbb{N}\times\mathbb{N}$ with $u_0 = 0 = v_t$ and $X_{u_i} = Y_{v_i}$ for every i, $0 \le i \le t$. Any such path $\mathcal{P}$ will be called a *blocker* for the pair $\langle\mathbf{X},\mathbf{Y}\rangle$, and we will then say $\langle\mathbf{X},\mathbf{Y}\rangle$ is blocked by $\mathcal{P}$.

For any path $\mathcal{P}$ from the y- to x-axes, let $B_\mathcal{P}$ be the probability that $\langle\mathbf{X},\mathbf{Y}\rangle$ is blocked by $\mathcal{P}$. For independent percolation, the next step would be to show that $\sum_\mathcal{P} B_\mathcal{P}$ is finite. Later, that sum is beheaded by assuming no blockage within some fixed finite range of the origin, so that $\sum_\mathcal{P} B_\mathcal{P} < 1$. The conclusion is that percolation takes place with positive probability. This is often called a "Peierls argument" (see e.g. [6]).

For our dependent percolation, this route is doomed to failure. The number of suitable paths $\mathcal{P}$ in the $n\times n$ box containing the origin is of order 3^{cn^2} for some constant c, and every one is a blocker for at least one pair of length-n walks. If G has k vertices then every pair of such walks has probability at least $((k-1)^{-n})^2$, hence the aforementioned sum is bounded below by

$$\sum_{n=2}^{\infty} \frac{3^{cn^2}}{(k-1)^{2n}}$$

which is infinite.

To overcome this difficulty we must get a handle on what it is that stops the fickle demon from completing her escape. This requires some new notions.

3 Projections and Prejections

If $U = U_0U_1\dots U_m$ and $V = V_0V_1\dots V_n$ are finite walks on G, we write $V \to U$ if there is a *projection* from V onto U, i.e. a map π from $\{0, 1, 2, \dots, n\}$ *onto* $\{0, 1, 2, \dots, m\}$ such that

(i) $\pi(0) = 0$ and $\pi(n) = m$;

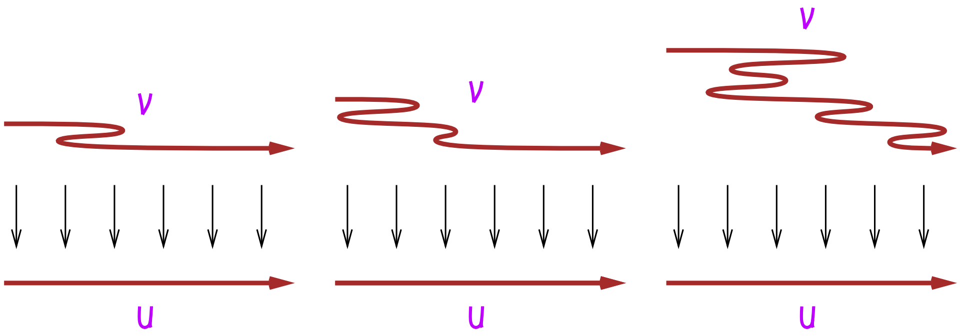

Figure 2: Some projections, in schematic form

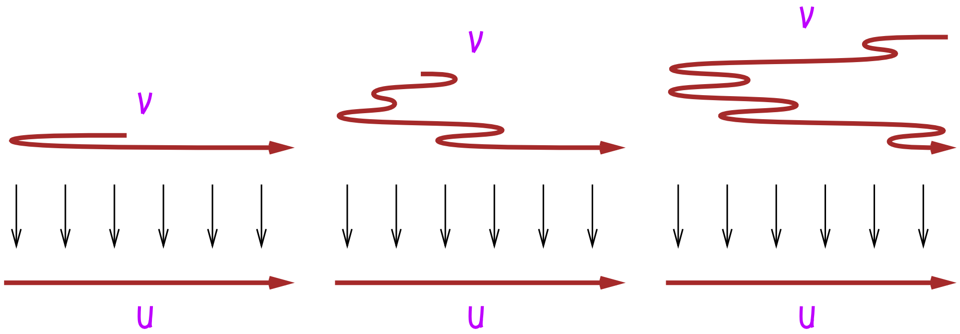

Figure 3: Some prejections, in schematic form

(ii) $|\pi(i+1) - \pi(i)| = 1$ for every i with $0 \leq i < n$;
(iii) $U_{\pi(i)} = V_i$ for every i, $1 \leq i \leq n$.

When $V \to U$ we imagine that V is represented by a piece of string, which is then folded upon itself to produce the shorter string U, as in Figure 2.

We shall also require a slightly weaker form of map between strings, which we call a *prejection*; here condition (i) above is replaced by

(i') $\pi(n) = m$.

Thus a *pre*jection from V to U may map the beginning of V to some node in the middle of U. We write (suggestively) $V \hookrightarrow U$ if there is a prejection from V to U, and imagine folded strings as in Figure 3. The definition of prejection is deliberately asymmetric, so that e.g. $V \hookrightarrow U$ does not imply $V^{-1} \hookrightarrow U^{-1}$, where U^{-1} is our notation for the reverse of the walk U. A projection is, of course, a special case of a prejection, and here $V \to U$ does imply $V^{-1} \to U^{-1}$.

We define the *length* $|V|$ of $V = v_0 v_1 \dots v_n$ by $|V| = n$; since a prejection is required to be a surjection, we have that $V \hookrightarrow U$ implies $|V| \geq |U|$. However, the definition of prejection extends naturally to (right-)infinite walks, with condition (i') omitted, where there is no issue of length.

A prejection π is defined by its *bends*; a bend is an index i for which $\pi(i+1) = \pi(i-1)$. If i is the only bend of a prejection τ then $\tau(i) = 0$ and we call τ a *tuck*. Any *pro*jection π must have an even number of bends; if it has precisely two we call it a *fold*.

It is easily checked that we can have $V \hookrightarrow U$ via different prejections. When $V \to U$ there may be many different projections from V onto U, some of which are folds and some not.

Condition (*ii*) of the definition of a projection or prejection can be replaced by "$|\pi(j) - \pi(i)| \leq j - i$ for every $0 \leq i < j \leq n$", since $\pi(i+1) = \pi(i)$ is impossible, there being no loops in G. This makes transitivity obvious:

Lemma 3.1. *The composition of prejections is a prejection, and the composition of projections is a projection. In particular $(W \to V) \wedge (V \to U)$ implies $W \to U$, and $(W \hookrightarrow V) \wedge (V \hookrightarrow U)$ implies $W \hookrightarrow U$.*

The fickle demon is in obvious trouble if the two tokens are slated to head directly toward each other's initial positions along the same path. A more general situation in which the fickle demon is blocked is illustrated in Figure 4, where we see a blocked pair of walks on a graph G, its appearance in the percolation setting, and prejections of the walks onto a third walk Z and its reverse. The relevance of the prejections is established by Theorems 3.2 and 4.1 below.

Theorem 3.2. *Let X and Y be finite initial segments of $\mathbf{X}$ and $\mathbf{Y}$, and suppose that there is a walk Z such that $X \hookrightarrow Z$ and $Y \hookrightarrow Z^{-1}$. Then the pair $\langle X, Y \rangle$ is blocked, that is, tokens taking walks which begin with X and Y cannot be scheduled so as to avoid a collision.*

Proof. We show that, roughly speaking, whichever token finishes first must have passed the other on Z, an impossibility. Let $X = X_0X_1 \ldots X_m$, $Y = Y_0Y_1 \ldots Y_n$ and $Z = Z_0Z_1 \ldots Z_k$. Let α be a prejection from X to Z, and β a prejection from Y to Z^{-1}. We assume that the values of β are interpreted as indices of the forward walk Z, so that in particular $\beta(n) = 0$. Suppose that the fickle demon has a winning strategy, the result of which is that at time t the X-token is found at step i_t of its walk, and the Y-token at step j_t. Let s be the first time at which either $i_s = m$ or $j_s = n$. Then $\alpha(i_t)$ and $\beta(j_t)$ are well defined for $0 \leq t \leq s$. Let A_t be the event that $\alpha(i_t) > \beta(j_t)$.

Clearly A_t holds at $t = s$. But $\alpha(i_t) = \beta(j_t)$ implies collision, $\alpha(i_t)$ and $\beta(j_t)$ each change only in increments or decrements of 1, and only one can change at a time. Thus A_t must hold for *all* times $t \leq s$. But if $i_s = m$ then $\alpha(i_t) = 0$ for some $t < s$, a contradiction; and similarly if $j_s = n$ then there is a $t < s$ for which $\beta(j_s) = k$. ☐

If such a Z exists we say that $\langle X, Y \rangle$ is *foiled* by Z. To show that this is the only way to stop the fickle demon is a bit trickier, and is the objective of the next section. First, however, we make use of Theorem 3.2 to obtain some negative results.

It is obvious that the fickle demon cannot win on a finite path, since the X-token will eventually migrate to the left end and then to the right end, and the y-token to the right end and then the left, thus the path itself serves as its own Z. Less obvious is what happens on an *infinite* path.

Theorem 3.3. *Let G be the two-way infinite path $\mathbb{Z}$, on which two (simple symmetric) random walks $\mathbf{X}$ and $\mathbf{Y}$ are taken from distinct starting points. Then $\mathbf{X}$ and $\mathbf{Y}$ are blocked with probability 1.*

Proof. We may suppose that the X-token begins at -1 and the Y-token at 1.

Let $Z(k)$ be the path from -2^k to 2^k, $k = 1, 2, \ldots$. Let A_k be the event that the X-token reaches 2^k before it gets to -2^k, and B_k the event that the Y-token reaches -2^k before it gets to 2^k.

A_1 and B_1 are independent and each occurs with probability 1/4, so with probability 1/16 both occur and in that case $Z(0)$ foils the tokens. If neither A_1 nor B_1 occurs then the probability that both A_2 and B_2 occur is again 1/16, because the condition then puts the X-token at -2 and the Y-token at 2. If just one of A_1 and B_1 fails, then $\Pr(A_2 \wedge B_2) = 3/16 > 1/16$.

In like manner we see that for any k, the probability of $A_k \wedge B_k$, given that $A_j \wedge B_j$ has failed to occur for any $j < k$, is at least 1/16. Thus the probability that some $Z(k)$ foils $\langle \mathbf{X}, \mathbf{Y} \rangle$ is 1. ☐

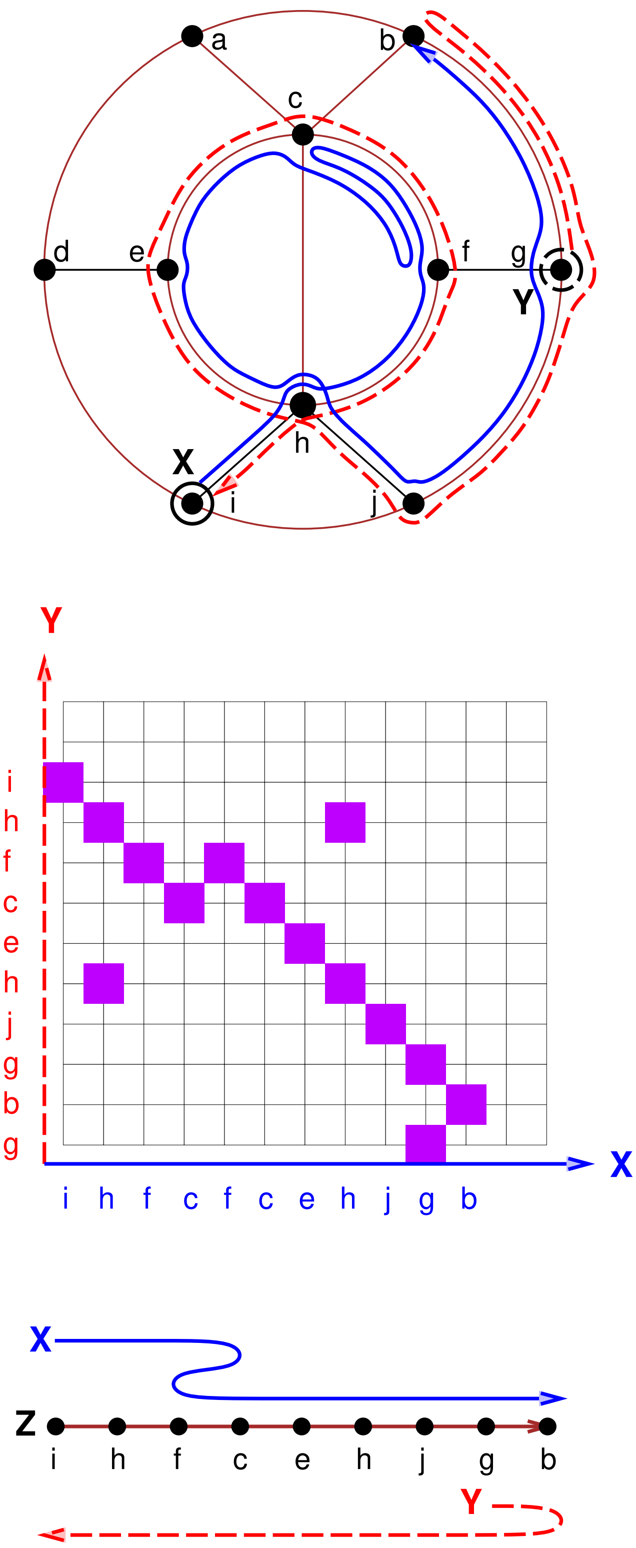

Figure 4: The fickle demon is foiled: on the graph, on the grid, and by prejections

Note that the expected maximum length to which $\mathbf{X}$ and $\mathbf{Y}$ can be advanced on $\mathbb{Z}$ is infinite, since even if the demon pushes each token alternately and only forward, they may move apart with infinite expected return time.

A graph G on which the fickle demon has positive probability of success is said to be *navigable*, thus Theorem 3.3 says that $\mathbb{Z}$ is not navigable. In fact it has implications for finite graphs as well.

Corollary 3.4. *No cycle is navigable.*

Proof. We employ $\mathbb{Z}$ as a covering graph for the cycle C_n; since collision on $\mathbb{Z}$ implies collision on C_n, the demon is doomed on the cycle as well. □

Since a connected graph on 3 vertices can only be a path or cycle, the fickle demon will need more vertices to work with if she is to succeed.

4 Foiling the Demon

Theorem 4.1. *Suppose that $\langle \mathbf{X}, \mathbf{Y} \rangle$ is a blocked pair of (infinite) walks. Then $\langle \mathbf{X}, \mathbf{Y} \rangle$ is foiled by some Z, in other words, there are finite initial segments X of $\mathbf{X}$ and Y of $\mathbf{Y}$ and a walk Z such that $X \hookrightarrow Z$ and $Y \hookrightarrow Z^{-1}$.*

Proof. A blocker $\mathcal{P}$ marks a walk W in G via $W_i := X_{u_i} = Y_{v_i}$. The projection $\pi_x :\ i \mapsto u_i$ of $\mathcal{P}$ down to the x-axis is a *pre*jection π_x from W^{-1} to X^{-1}, where X is a finite initial segment of $\mathbf{X}$; similarly, the projection $\pi_y :\ i \mapsto v_i$ of $\mathcal{P}$ over to the y-axis is also a prejection, this time from W itself onto Y^{-1} where Y is a finite initial segment of $\mathbf{Y}$. We will not make any direct use of these prejections but note that with them, Theorem 4.1 can be viewed as a kind of geometric diamond lemma. Later we will prove a genuine diamond lemma for projections.

In general a blocker looks like a zig-zag of diagonal segments running across the quadrant from the y-axis to the x-axis, as in Figure 5. To be precise we define a *turn* in the blocker $\mathcal{P}$ to be an index $s \in [1, t-1]$ for which $\rho(P_{s+1}, P_{s-1}) = 2$; in particular an *$x$-turn* if $u_{s+1} = u_{s-1}$ and a *y-turn* if $v_{s+1} = v_{s-1}$. An *interior segment* of P is an interval between consecutive turns; the *starting segment* of $\mathcal{P}$ is $[0, s]$ where s is the first turn, and the *ending segment* is defined similarly. Identifying turns and segments with the sites and sub-paths which they induce in $\mathcal{P}$, we have that $\mathcal{P}$ is the union of its segments, overlapping at their endpoints, which are precisely the turns of $\mathcal{P}$.

We prove the theorem by induction on the length $|\mathcal{P}| = t$ of the blocker $\mathcal{P}$. Observe first that if $\mathcal{P}$ consists of a single segment, necessarily running from $(0, n)$ to $(n, 0)$ for some n, then we can take $|X| = |Y| = n$ and $Z := X = Y^{-1} = W$ with identity prejections, which trivially foils $\langle X, Y \rangle$. Thus $\mathcal{P}$ has at least two segments and we let σ be a segment of minimum length.

Suppose that $\sigma = [0, s]$ is the starting segment of $\mathcal{P}$. If σ ends with an x-turn, then $u_{2s} = 0$ so $\mathcal{P}' := P_{2s}P_{2s+1} \dots P_t$ is a shorter blocker $\mathcal{P}'$ for the same pair $\langle \mathbf{X}, \mathbf{Y} \rangle$, and we may use for $\mathcal{P}$ the same X, Y and Z guaranteed for $\mathcal{P}'$ by the induction hypothesis.

If σ ends with a y-turn, then $v_{s+i} = v_{s-i}$ for $0 \le i \le s$, and $u_j = j$ for all $j \le 2s$. Hence $X_{s+i} = Y_{v_{s+i}} = Y_{v_{s-i}} = X_{s-i}$, again for $0 \le i \le s$. Thus $\mathbf{X}$ admits a tuck $\tau :\ \mathbf{X} \hookrightarrow \mathbf{X}'$ onto $\mathbf{X}' = X'_0 X'_1 \cdots = X_s X_{s+1} \dots$, and we can in effect apply the same tuck to $\mathcal{P}$, obtaining a set Q of sites defined by

$$Q := \{(u_{\tau(i)}, v_i) :\ 0 \le i \le t\} \ .$$

Since τ preserves the diagonal edges of $\mathbb{N} \times \mathbb{N}$, Q still connects the y- and x-axes within $\mathbb{N} \times \mathbb{N}$ and therefore contains a blocker $\mathcal{P}'$ for $\langle \mathbf{X}', \mathbf{Y} \rangle$ with $|\mathcal{P}'| < |\mathcal{P}|$. Let X', Y and Z be the walks guaranteed by the induction hypothesis for $\mathcal{P}'$; then $\mathbf{X} \hookrightarrow \mathbf{X}'$ via τ, so if X is the inverse image

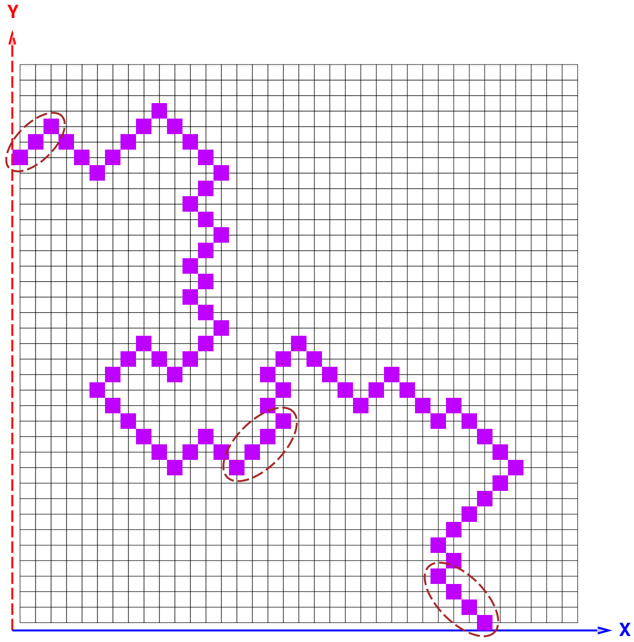

Figure 5: Blocker with starting, ending and internal segment marked

of X' under τ then $X \hookrightarrow Z$ by Lemma 3.1. Note that X' has length at least that of the second segment of $\mathcal{P}$, in particular at least s, thus X is an initial segment of $\mathbf{X}$.

The argument when σ is the final segment of $\mathcal{P}$ is symmetrical.

Now we suppose that σ is the internal segment $[s, s+j]$, and we consider first the case where σ begins with a y-turn and ends with an x-turn. Then the points P_{s-j}, P_s, P_{s+j} and P_{s+2j} are the corners of a (tilted) square, three sides of which are contained in $\mathcal{P}$. But then also the $j-1$ points lying on the shortest path in $\mathbb{N}\times\mathbb{N}$ between P_{s-j} and P_{s+2j} qualified for deletion. It follows that there is a shorter blocker $\mathcal{P}'$ which uses these points in place of $P_{s-j+1} \dots P_{s+2j-1}$, and we again apply the induction hypothesis to $\mathcal{P}'$. The argument is symmetrical when σ begins with an x-turn and ends with a y-turn.

Finally we address the case where $\sigma = [s, s+j]$ begins and ends with turns of the same type, say y-turns; and suppose u_i is increasing in i over σ (therefore in the whole range $[s-j, s+2j]$ as well). Then $v_{s-i} = v_{s+i} = v_{s+2j-i}$ for $0 \leq i \leq j$ hence $X_{v_s-i} = X_{v_s+i} = X_{v_s+2j-i}$, also for $0 \leq i \leq j$. It follows that $\mathbf{X}$ admits the fold φ which bends at v_s and $v_s + j$.

We apply the fold to $\mathcal{P}$ as we did with the tuck τ before, letting

$$Q := \{(u_{\varphi(i)}, v_i) : \ 0 \leq i \leq t\} \ .$$

and choosing a blocker $\mathcal{P}'$ from Q with $|\mathcal{P}'| < |\mathcal{P}|$. Let X', Y and Z be the walks guaranteed by the induction hypothesis for $\mathcal{P}'$; $\mathbf{X} \hookrightarrow \mathbf{X}'$ via φ so if X is the inverse image of X' under φ then $X \hookrightarrow Z$ by Lemma 3.1. The argument is the similar if u_i is decreasing in i over σ, and again if σ terminates in x-turns instead of y-turns.

This concludes the induction and thus also the proof of Theorem 4.1. $\square$

5 Percolation

Now that we know what causes blockage, we are in a position to compute the probability that $\langle \mathbf{X}, \mathbf{Y} \rangle$ will be blocked. For the sake of simplicity of presentation, we deal initially with the special case where G is the complete graph K_4 on four vertices.

Let Z be a walk of length n, and let $b(Z)$ be the probability that $\langle \mathbf{X}, \mathbf{Y} \rangle$ is foiled specifically by Z, in other words that there are initial segments X and Y of $\mathbf{X}$ and $\mathbf{Y}$ and prejections $\alpha : X \hookrightarrow Z$ and $\beta : Y \hookrightarrow Z^{-1}$.

Suppose that α does preject X onto $Z = Z_0 \dots Z_n$, and suppose $\alpha(0) = m$. Then α can be "untucked" into a *pro*jection α' from X onto

$$Z' = Z'_0 \dots Z'_{n+m} := Z_m Z_{m-1} \dots Z_1 Z_0 Z_1 \dots Z_n \ .$$

We are interested in bounding the probability $P'_{Z'}$ that some initial segment X of $\mathbf{X}$ projects onto Z'. Since a projection determines the X on which it operates, that probability is the highest when each qualifying X has only *one* such projection. Thus for this purpose, we may imagine that $Z'(i+1) \neq Z'(i-1)$ for any i, $0 < i < n+m$.

We may also assume that the qualifying X is a minimal initial segment of $\mathbf{X}$, in which case it is simply a walk on Z' which begins at Z'_0 and ends when it first reaches Z'_{m+n}. Let us compute the probability p_{j+1} that $\mathbf{X}$, given that it has successfully reached Z'_j, gets to Z'_{j+1} without falling off the path.

To do this the next node of $\mathbf{X}$ must be either Z'_{j+1} or Z'_{j-1}, and in the latter case $\mathbf{X}$ must struggle back to Z'_j and then to Z'_{j+1}. We thus obtain the recursion

$$p_{j+1} = \frac{1}{3} + \frac{1}{3} p_j p_{j+1}$$

beginning with $p_1 = 1/3$. Solving, we have $p_j < p := (3 - \sqrt{5})/2 \sim 0.382$ for all j. The key is that $p^2 < 1/3$.

Since all the nodes of Z' must be reached, $P_{Z'} < p^{m+n}$ and therefore the probability P_Z that some initial segment of $\mathbf{X}$ *pre*jects onto Z is less than

$$\sum_{m=0}^{n} p^{m+n} \leq \frac{p^n}{1-p} \ .$$

Back in G, the number of walks Z of length n is $4 \cdot 3^n$. Since $\mathbf{X}$ and $\mathbf{Y}$ are independent, the probability that an initial segment of $\mathbf{X}$ prejects onto a fixed Z while an initial segment of $\mathbf{Y}$ prejects onto Z^{-1} is P_Z^2. The probability that there is *some* Z of length n onto which they both preject is thus bounded by

$$4 \cdot 3^n \left(\frac{p^n}{1-p} \right)^2 = \frac{4}{(1-p)^2} (3p^2)^n \ .$$

This means that the total probability of $\langle \mathbf{X}, \mathbf{Y} \rangle$ being blocked is bounded by

$$\sum_{n=1}^{\infty} \frac{4}{(1-p)^2} (3p^2)^n = \frac{4 \cdot 3p^2}{(1-p)^2(1-3p^2)} \sim 8.151421 \ ,$$

a less than devastating conclusion. However, since we only wish to prove that percolation occurs with positive probability, we may assume that $\mathbf{X}$ and $\mathbf{Y}$ begin with particular finite initial segments; let us fix some number s, to be chosen later, and let $\mathbf{X}$ begin with $abcdabcd \dots abcd = (abcd)^s$, and $\mathbf{Y}$ with $(cadb)^s$, where a, b, c and d are the vertices of G.

Since these strings admit no tucks or folds, and do not overlap, any Z which witnesses the blockage of $\langle \mathbf{X}, \mathbf{Y} \rangle$ must contain each of $(abcd)^s$ and $((cadb)^s)^{-1} = (bdac)^s$ intact, although either may conceivably be reversed. The main possibility is that Z begins with $(abcd)^s$ and ends with $(bdac)^s$, with a middle section Z' of length (say) $n-1$. (The empty word, which is a possibility for Z', has length -1 by our convention.)

We have at most 3^n choices for Z'. The X-token will travel "for free" along the first $4s$ nodes of Z, by construction, but from there will make it to the end of Z only with probability bounded by p^{n+4s}. Since a similar calculation holds for the Y-token, the probability that initial segments of $\mathbf{X}$ and $\mathbf{Y}$ preject onto Z and Z^{-1} respectively for some Z of this sort is bounded by

$$\sum_{n=0}^{\infty} 3^n \cdot (p^{n+4s})^2 = \frac{p^{8s}}{1-3p^2}$$

which we can make as small as we like.

The probability that $\langle \mathbf{X}, \mathbf{Y} \rangle$ is blocked via some Z of the form $Z''(dcba)^s Z'(bdac)^s$, or one of the several other possibilities, is minuscule and can be dealt with in the same manner as the first case. Thus we have at last proved that with positive probability, two random walks on K_4 can be "scheduled" so that both are advanced arbitrarily far without a collision—in other words,

Theorem 5.1. *K_4 is navigable.*

It is interesting to compare the proof of Theorem 5.1 in [2], in which a different and simpler "compression" rule replaces our projections. Their approach does not yield an explicit characterization of when percolation occurs (as does Theorem 4.1), but it gives a much better lower bound on the probability of percolation.

We do feel that our notion of projection is natural and therefore its diamond lemma is of independent interest—the more so, as it appears to be a more subtle result than one might expect.

6 The Diamond Lemma

In what follows we will fix an arbitrary finite, undirected, simple, loopless, connected graph $G = \langle V, E \rangle$, upon which all walks take place. In order to determine precisely which graphs (and later, which pairs of Markov chains) are navigable, it is useful to know that every walk can be reduced by projections to a unique irreducible walk.

In fact nothing is lost if the reader assumes G is complete and looped, so that every string on the alphabet of vertices of G is a walk; then the set of finite walks on G has the structure of a semigroup, and a fold can be thought of as an application of the reduction rule

$$UxVyV^{-1}xVyW \to UxVyW$$

where x and y are vertices (generators of the semigroup), and U, V and W arbitrary (possibly null) words.

If we imagine that the walks in G are represented by strings of *arcs* (directed edges) instead of strings of vertices, then the reduction rule takes a more familiar form:

$$UVV^{-1}VW \to UVW$$

where here $(e_1 e_2 \dots e_k)^{-1} := e_k^{-1} e_{k-1}^{-1} \dots e_1^{-1}$, e_i^{-1} being the reverse of the arc e_i. Thus U^{-1} and U are inverses in the semigroup sense and the result is a sort of "free regular semigroup". The well-studied free *inverse* semigroup (see e.g. [7]) has *unique* inverses, but this is not the case here; for example the word $ee^{-1}ff^{-1}$ has both itself and the distinct word $ff^{-1}ee^{-1}$ as inverses. As far as we have been able to determine, our semigroup has not previously been studied, and the normal form theorem below is new.

Theorem 6.1. *For every finite walk W on G there is a (necessarily unique) walk W^* on G such that $U \to W^*$ for every U with $W \to U$.*

Let us remark here that Theorem 6.1 does *not* claim the following amalgamation property for projections: "if μ and ν are projections of W onto U and V, respectively, then there are projections φ and ψ from U and V onto some word X such that $\varphi\mu = \psi\nu$." This property would indeed imply the statement of Theorem 6.1 but it is false; a simple counterexample is provided by the walk ABCBABCBA and the two projections

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 1 & 2 & 3 & 2 & 1 & 2 & 3 & 4 & 5 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ 1 & 2 & 3 & 4 & 5 & 4 & 3 & 4 & 5 \end{pmatrix}.$$

In this case the projection images U and V are both irreducible, and would constitute a counterexample to Theorem 6.1 itself were it not for the fact that they are equal! We would like to blame the failure of amalgamation for the complexity of the following proof, but it is of course possible that our own lack of insight is at fault.

Proof. An index k for a walk $W = W_0W_1 \dots W_n$ will be called a *point* (of W). The walk W maps each point i to a vertex W_i of G.

A projection π on W is characterized by its *bends*; recall that a bend is a point k for which $\pi(k+1) = \pi(k-1)$, and that a projection with just two bends is called a *fold*. We write $U \prec V$ or $V \succ U$ if there is a fold mapping V onto U. If a fold φ bends at k and $k+i$, then the points $k-i$ and $k+2i$ which are mapped by φ to the same points as the bends are called "borders"; the interval $[k-i, k+2i]$ which covers all four of these points is called the *span* of φ.

We have already observed that the space of walks on G (or, more specifically, walks from one fixed vertex to another) is a partially ordered set under the relation "$\to$". The next lemma implies that if V "covers" U in this partial order then $U \prec V$ (the converse fails).

Lemma 6.2. *If $W \to U$ then W can be reduced to U by a series of folds.*

Proof. Let $W \to U$ via the projection π and let i and j be a pair of bends minimizing $j-i$ subject to $i < j$. Let φ be the fold on W which bends at i and j, reducing W to some V. Then the map ψ defined on the points of V by

$$\psi(k) = \begin{cases} \pi(k) & \text{for } k \leq i \text{ and} \\ \pi(k + 2(j-i)) & \text{for } k > i \end{cases}$$

is a projection onto U. (See Figure 6.)

We thus have $W \succ V \to U$ and since $|V| \leq |W| - 2$, where $|V|$ denotes the length of the walk V, this process terminates after at most $(|V| - |U|)/2$ folds. □

A familiar argument (given later, and sometimes known as the Church-Rosser Theorem) asserts that to show the existence of W^* it is sufficient to prove a "diamond lemma", namely:

Lemma 6.3. *If $W \succ U$ and $W \succ V$ then there is a walk Z with $U \to Z$ and $V \to Z$.*

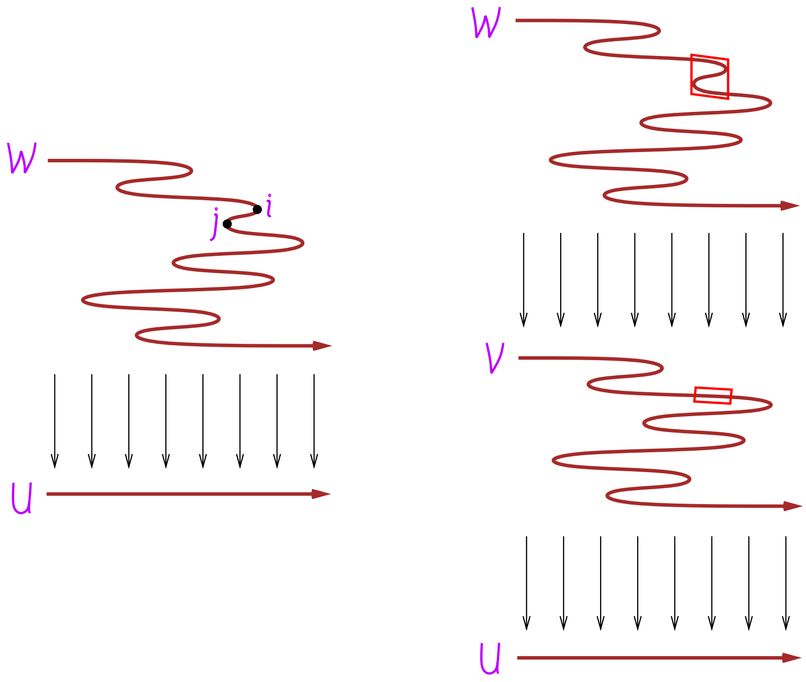

Figure 6: Decomposing a projection into a fold followed by a projection

Proof. The diamond lemma does not say that the two folds commute (whatever that means) nor that $Z \prec U$ or $Z \prec V$; in fact, we have seen that sometimes $U = V = Z = W^*$. In other cases many folds may be required to reduce U and V to Z.

Let μ be a fold which projects W onto U, and ν a fold projecting W onto V. We may assume (e.g. by induction on the length of W) that the points 0 and $n = |W|$ are among the borders of μ and ν. The general plan of the proof is first to create maps φ from W, ψ from U and χ from V, all onto the same walk X, which do produce a commuting diagram in the sense that $\varphi = \psi \circ \mu = \chi \circ \nu$. These maps are not always projections. When they are not we must tailor them to produce genuine projections which, though they no longer commute, agree on a common image Z. To effect the plan we introduce concomitant new terminology.

A *step* of W is a point together with a direction (forward or backward); the forward step from point i is denoted by i^+, the backward step by i^-. (The steps 0^- and n^+ are not considered to lie in W.) The steps of W are mapped to steps of U and V by the maps μ and ν, in the obvious way; e.g. $\mu(i^-) = \mu(i^+) = i^-$ if i is the left bend of μ.

Steps map naturally onto *tails* of G; a tail is a pair consisting of a vertex and an edge incident to that vertex. If a is a step we denote its image under W by W_a, so that e.g. $W_{i^+} := \langle W_i, \{W_i, W_{i+1}\}\rangle$.

Legitimacy of the projections μ and ν forces certain points of W to map to the same vertex, and similarly forces certain steps to map to the same tail. If a and b are steps, let us say that $a \smile b$ if $\mu(a) = \mu(b)$ and $a \frown b$ if $\nu(a) = \nu(b)$.

We say that steps a and b are *equivalent* if $a \equiv b$, where "$\equiv$" is the smallest equivalence relation containing both "$\smile$" and "$\frown$". Then of course $a \equiv b$ implies $W_a = W_b$. Equivalence classes of steps under $\equiv$ will be called *step classes*.

The relations "$\smile$" and "$\frown$" extend naturally from steps to points, e.g. if i and j are points then $i \smile j$ if $\mu(i) = \mu(j)$. Equivalence among points is defined as for steps, as the smallest equivalence relation containing both "$\smile$" and "$\frown$"; again, $i \equiv j$ evidently implies $W_i = W_j$. However, the somewhat finer equivalence on steps will be more useful to us than the equivalence on points. If the step classes of i^- and i^+ are η and ζ, in either order, then we say the point i is of "type" $\eta|\zeta$. In general, for points i and j, being of the same type is strictly stronger than saying $i^+ \equiv j^+$ or $i^+ \equiv j^-$, either of which already implies $i \equiv j$.

We are particularly interested in the step classes of the borders of μ and ν. Let 0 and b_{left} be the left borders of μ and ν (in either order) and b_{right} and n the right borders. Denoting by $[a]$ the equivalence class of the step a, we define six "special" step classes as follows:

$\alpha := [0^+]$;
$\beta := [n^-]$;
$\gamma := [b_{\text{right}}^-]$;
$\gamma' := [b_{\text{right}}^+]$, if $b_{\text{right}} < n$;
$\delta := [b_{\text{left}}^+]$;
$\delta' := [b_{\text{left}}^-]$, if $b_{\text{left}} > 0$.

The classes γ' and δ' may not exist, and in any case the special step classes may not all be distinct. Any point of W either of whose steps lies in a special step class will be called a *joint*; a joint k for which $[k^-] = [k^+]$ is called a *hinge*, otherwise it is a *weld*.

Figures 7and 8 below illustrate a few of the infinitely many[1] possibilities for labeling of W by special step classes.

In the figure, the spans of μ and ν are indicated by horizontal bars (above for μ, below for ν), each with four vertical prongs indicating the locations of borders and bends. The joints are marked in approximate positions, with their special step classes labeled. The actual points of the walk W, of which there may be arbitrarily many in any of the cases, are suppressed. (The folded arrows beneath each case will come into play later.)

An interval $[h, i]$ of points (and steps) between two joints, but not containing any other joint in its interior, is called a *rod*; two rods are *equivalent* if they are step-by-step equivalent, possibly with one of them reversed.

Lemma 6.4.

1. *If k is a joint in the interior of W then k is of type $\alpha|\alpha$, $\beta|\beta$, $\gamma|\gamma$, $\delta|\delta$, $\gamma|\gamma'$ or $\delta|\delta'$.*

2. *If $[h, i]$ and $[j, k]$ are two rods with $[h^+] = [j^+]$ or $[i^-] = [k^-]$ then they are (forward) equivalent.*

3. *If $[h, i]$ and $[j, k]$ are two rods with $[h^+] = [k^-]$ or $[i^-] = [j^+]$ then they are (reverse) equivalent.*

Proof. Let $\mathcal{J}_0$ be the set of borders and for each $k \geq 0$, put

$$\mathcal{J}_{k+1} := \{j \notin \bigcup_{k' \leq k} \mathcal{J}_{k'} : j \smile i \text{ or } j \frown i \text{ for some } i \in \mathcal{J}_k\}.$$

Then $\mathcal{J} := \bigcup_k J_k$ is the union of the (point-) equivalence classes of the borders, and contains all the joints. In fact $\mathcal{J}$ contains only joints. To see this, and with it statement (1) of the lemma, let $j \in \mathcal{J}_{k+1}$ with $i \smile j \in \mathcal{J}_k$. Let a and b be the steps (in either order) incident to i and c and d the steps incident to j. Then $\mu(a) = \mu(c)$ implies $\mu(b) = \mu(d)$ unless i is a border of μ and j is a bend; this happens only when $k = 0$ and results in $[c] = [d] = \alpha$, β, γ or δ. The case $i \frown j$ is similar.

[1] There are only 151 ways the borders and bends of μ and ν can interleave, but most of these produce infinitely many labelings. For example, let $j,\ k > 1$ and suppose μ bends at jk and $2jk$ while ν bends at $(j+1)k$ and $(2j+1)k$. Then the result is $3j$ hinges in all.

Note that the bends of μ and ν are in $\mathcal{J}_1 \cup \mathcal{J}_0$ and are hinges; we call them the "original hinges".

To show statements (2) and (3), suppose that $[h,i]$ is a rod and $h^+ \smile j^+$. Then $[j, j+(i-h)]$ must also be a rod, because $\mu((j+i-h)^-) = \mu(i^-)$ and if there were a joint between j and $j+(i-h)$ it would be $\smile$-equivalent to a joint in (h,i). Moreover, corresponding steps a and b of $[h,i]$ satisfy $a \smile b$. The argument is mirrored if $h^+ \smile j^-$. The same is true for $\frown$ and hence $\equiv$. □

It follows from Lemma 6.4 that there is an involution ι of the set of special step classes such that if $[h,i]$ is a rod then $\iota[h^+] = [i^-]$ and $\iota[i^-] = [h^+]$. In fact we will see that ι has no fixed point, that is, we cannot have $[h^+] = [i^-]$ when $[h,i]$ is a rod.

We now define a map φ from the indices of W to the non-negative integers $\mathbb{N}$ as follows:

1. $\varphi(0) = 0$ and $\varphi(1) = 1$;
2. $\varphi(k+1) - \varphi(k) \in \{-1, +1\}$ for all $k \in [0,n]$;
3. $\varphi(k+1) = \varphi(k-1)$ iff k is a hinge.

Thus φ bends at all the hinges of W, but (so far) is several properties short of being a projection.

Lemma 6.5.

1. *If $i \equiv j$ then $\varphi(i) = \varphi(j)$.*
2. *If $[h,i]$ is a rod then $h \not\equiv i$.*
3. *There are at most three equivalence classes of rods in W.*
4. *If i and j are equivalent welds then they are of the same type, i.e. either $[i^-] = [j^-]$ and $[i^+] = [j^+]$, or $[i^-] = [j^+]$ and $[i^+] = [j^-]$.*
5. *If $\varphi(i) = \varphi(j)$ then $i \equiv j$.*

Proof. If h, i, j and k are the borders and bends of μ from left to right, then there are hinges at i and j, and the pattern of hinges is the same on the intervals $[h,i]$ and $[j,k]$ and reversed on $[i,j]$; thus any two points which are $\smile$-equivalent are mapped to the same point by φ. The same applies for $\frown$ and hence for $\equiv$, establishing statement (1).

Statement (2) follows immediately, since the two ends of a rod cannot be mapped to the same point by φ. Statement (3) is now a consequence of Lemma 6.4 and the fact that there are at most six special step classes.

Statement (4) is *not* an immediate consequence of Lemma 6.4 because it appears to be possible that there are welds of types $\gamma|\gamma'$ and $\delta|\delta'$, with $\gamma = \delta$ but γ, γ' and δ' all distinct. However, no two of the step classes γ, γ' and δ' can be found at opposite ends of a rod, since the two types of welds and the (γ,γ) hinges are all equivalent points. Thus the set of distinct classes $\{\iota(\gamma), \iota(\gamma'), \iota(\delta')\}$ is disjoint from $\{\gamma, \gamma', \delta'\}$, but this is impossible because the coincidence of γ and δ leaves us with at most 5 distinct special step classes.

Finally, we need to show that φ does not collapse inequivalent points; in fact we show that any two inequivalent *steps* of W are mapped to different steps.

Suppose the contrary and let b be the earliest step for which there is some step a to the left of b, with $a \not\equiv b$ but $\varphi(a) = \varphi(b)$. Since φ maps equivalent rods onto the same interval of $\mathbb{N}$, b must be the initial step of some rod, say $b = i^+$. Then i cannot be a hinge; furthermore, some other weld $h < i$ must be equivalent to i, since the image of $[0, i-1]$ had to cross the point $\varphi(i)$ of $\mathbb{N}$.

But from statement (4) we know that the steps of h match the steps of i; let c be the step of h which is equivalent to b. Since $\varphi(c) = \varphi(b) = \varphi(a)$ but $c \not\equiv a$, either a or c contradicts the minimality of b. $\square$

We now know that φ is a consistent map on W, in the sense that $\varphi(i) = \varphi(j)$ implies $W_i = W_j$. Thus we can think of it as a map from W to the walk X defined by $X_{\varphi(k)} = W_k$. The walk X consists of precisely one rod from each equivalence class (thus at most three rods), joined at the images under φ of the welds of W.

Since all three of the maps μ, ν and φ have the property of collapsing steps only when they are equivalent, the special step classes are preserved and we may speak of joints, hinges and welds in U, V and X as well as in W. One might visualize W as a sequence of straight segments connected by hinges, like an old-fashioned hinged carpenter's yardstick. Each segment consists of one, two or three welded rods; the rods are oriented and are painted with one of (at most) three colors. Whether the yardstick is twice bent to form U or V, or fully folded to form X, overlapping rods match perfectly in color and orientation.

Let us now consider the walk $U = U_0, \ldots, U_s$. Since inequivalent steps of W cannot be mapped by μ to the same step of U, the step classes of W are projected faithfully onto U and we may speak of joints, welds and hinges of U as well as of W. Let ψ be the map defined on U the same way φ is defined on W, namely by bending at all hinges.

The hinges of W which appear between the borders of μ, other than μ's own bends, are arrayed symmetrically forward, reversed and then forward again, between the borders and bends of μ. They are thus folded three deep on top of one another when μ is applied. Hinges outside the borders of μ are mapped injectively. Thus for each hinge of U, none of the hinges in its μ-preimage were bent by μ itself. It follows that the result of composing μ with ψ is to bend at each hinge of W exactly once, duplicating φ.

Obviously we can apply a similar argument to $V = V_1, \ldots, V_t$ and the projection ν, obtaining a map χ such that $\chi \circ \nu = \varphi$.

Since μ and ν are projections, we have $\mu(0) = \nu(0) = 0$, $\mu(n) = s = |U|$, and $\nu(n) = t = |V|$. It follows that the maps ψ and χ begin and end in the same direction and at the same points of X as does φ; in a sense, they are simplified versions of φ.

Summarizing:

Lemma 6.6. *The maps $\varphi : W \to X$, $\psi : U \to X$, and $\chi : V \to X$ satisfy $\chi \circ \nu = \varphi = \psi \circ \mu$; $\psi(0) = \chi(0) = \varphi(0)$; $\psi(1) = \chi(1) = \varphi(1)$; $\psi(s) = \chi(t) = \varphi(n)$; and $\psi(s-1) = \chi(t-1) = \varphi(n-1)$.*

Of course, it is an immediate corollary that if φ happens to be a projection than so are ψ and χ, hence we can set $Z = X$ and we are done. This happens quite often. Figure 7 illustrates five such cases; in each case the three projections φ, ψ and χ are depicted (respectively left, center and right) as folded arrows.

The next lemma will help us to sort out and deal with the various cases where φ is not a projection.

Lemma 6.7. *Exactly two of the special step classes fail to appear in welds; one of these (say, η) is mapped by φ to the first step of X, the other (ζ) to the last.*

Proof. We have already seen that φ sends every step in a class to the same step of X. Since φ does not bend at welds, no class which appears in a weld can be sent to an end-step of X. Conversely, if the special class appears in no weld, the walk X cannot cross its image so it must be mapped to an end-step. $\square$

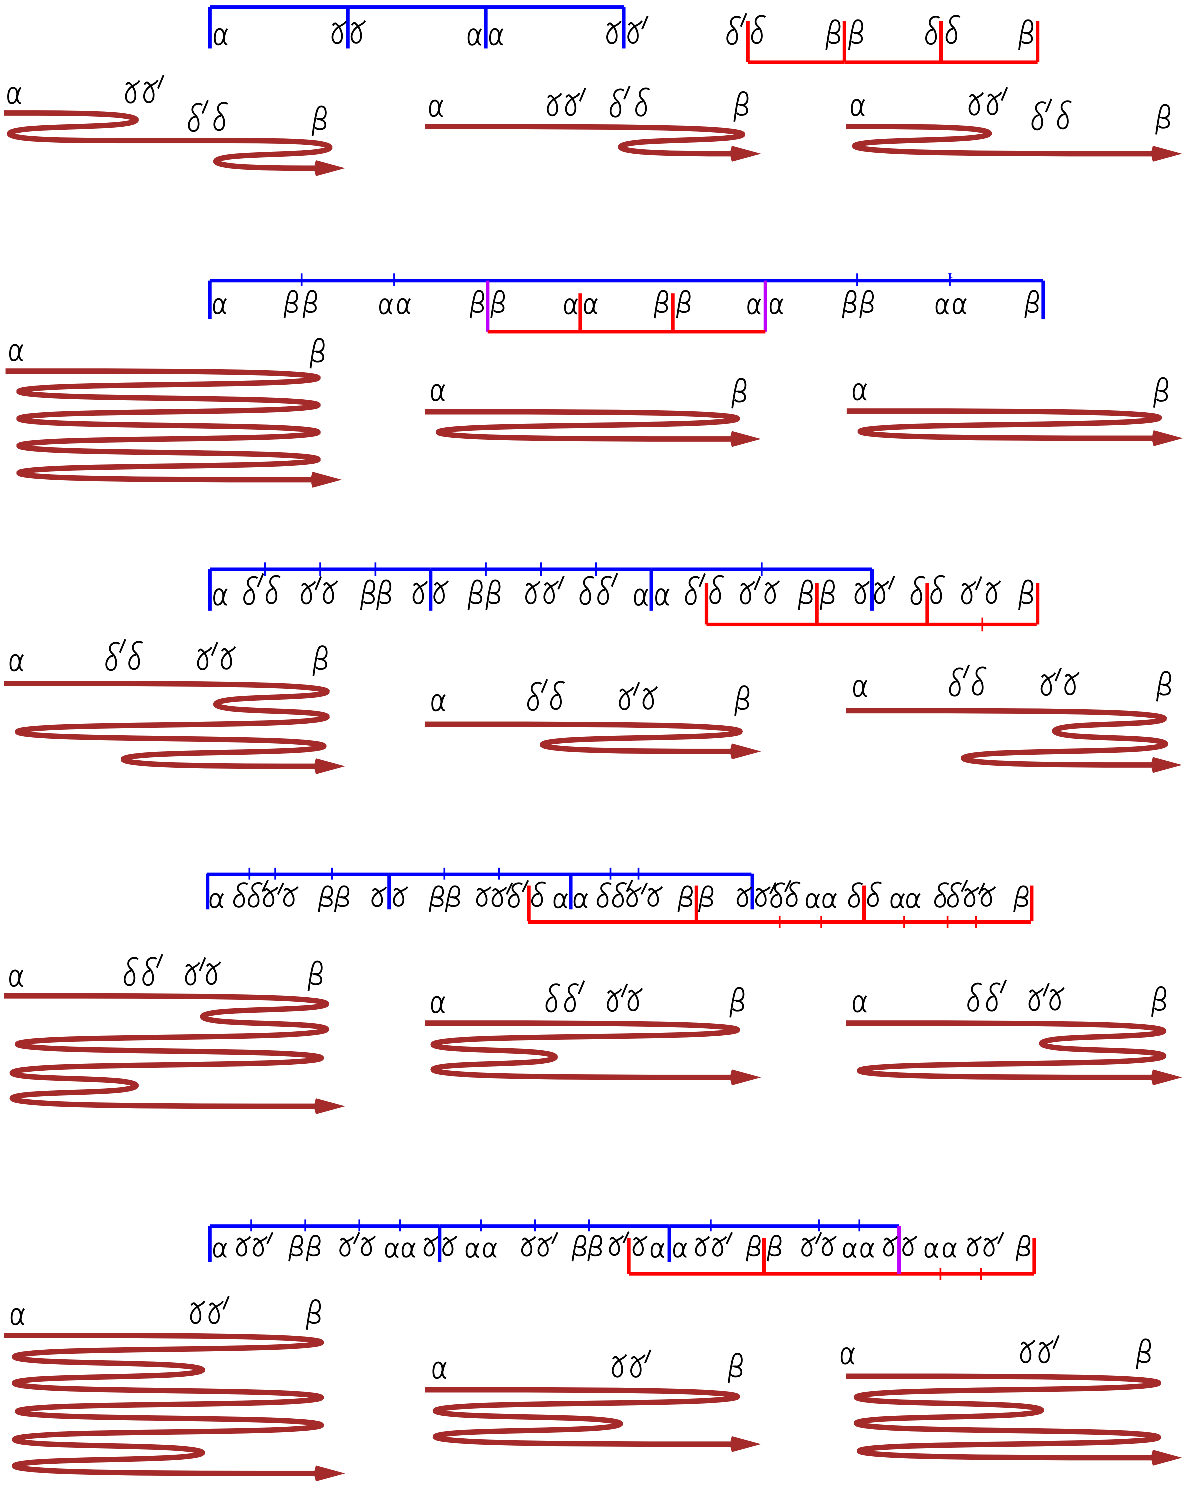

Figure 7: Five ways to label W, in which φ, ψ and χ are all projections

In W the leftmost step class is α, and the rightmost β; if these are distinct and neither appears in a weld, then we see from Lemma 6.7 that $\varphi(0) = 0$ and $\varphi(n)$ is the rightmost point of X, that is, φ is a projection. This would apply if all six step-classes are distinct, and in various other cases as well, some of which appear in Figure 7.

Lemma 6.8. *Suppose $\alpha = \beta$. Then W, U and V all project onto the walk Z consisting of X followed by its reverse X^{-1}, with the last point of X and the first of X^{-1} identified.*

Proof. If $\alpha = \beta$ and α appears in a weld, necessarily with γ' or δ', then $\alpha = \gamma$ or $\alpha = \delta$ and there is at most one special step class left out of the welds, contradicting Lemma 6.7. Thus, we must have $\alpha = \beta = \eta$, in the notation of Lemma 6.7, and ζ equal to δ or γ or both; let k be any point of W of type $\zeta|\zeta$, that is, any point mapped by φ to the last point $m = |X|$ of X. Define the map φ' from W to Z by bending at every hinge *except* k.

It is easy to see that φ' is a genuine projection. Certainly φ' maps 0 to 0 and n to the rightmost point of Z. Suppose φ' is inconsistent, that is, $\varphi'(i) \neq \varphi'(j)$ for some i, j with $W_i = W_j$. Then $i \leq k \leq j$ since φ is consistent, but then $\varphi'(i) \leq \varphi(k) \leq \varphi'(j)$ so all three are equal to $\varphi(i)$ and $\varphi(j)$.

Now we need to map U and V onto Z as well. We can do this by an identical procedure, simply changing ψ to ψ' by unbending any point in $\psi^{-1}(m)$—that is, any $\zeta|\zeta$ hinge of U. We replace χ by χ' similarly. Lemma 6.6 assures us that the same argument used for φ' works as well to show that ψ' and χ' are projections.

If the hinge k (the one unbent by φ') is neither a bend of μ nor a bend of ν, then one can in effect unbend the same hinge for each of the three maps, preserving commutativity. Since there may not be such a hinge, it is fortunate that we do not require a commuting diagram, but only a common image Z. ☐

The first two of the three cases illustrated in Figure 8 below are of the type covered by Lemma 6.8. In each there are six maps indicted by folded arrows, namely φ, ψ and χ left to right, then φ', ψ' and χ' left to right. The points at which the first three were unfolded to produce the others are indicated by black dots.

We are left with the case that $\alpha \neq \beta$ and one of them is part of a weld. Suppose for example that $\alpha|\gamma'$ is a weld so that $\alpha = \gamma$ and it is thus β and δ which are left to be the two classes not appearing in welds, as in the last example of Figure 8. The other cases are similar.

As above we let ζ stand for whichever of β and δ is mapped by φ to the right-hand end of X. Let j be the leftmost point of W other than 0 satisfying $\varphi(j) = \varphi(0)$; then $\varphi(i) > \varphi(0)$ for all i with $0 < i < j$, and there is at least one non-$\alpha|\alpha$ hinge k between 0 and j. By elimination this hinge must be of type $\zeta|\zeta$, and it follows that $\varphi(n) = 0$.

Now we let Z be the walk $X_{\varphi(0)}, X_{\varphi(0)+1}, \ldots, X_{\varphi(k)-1}, X_m, X_{m-1}, \ldots, X_1, X_0$. The map φ' from W to Z which bends at every hinge except k is seen, by an argument similar to that in the proof of Lemma 6.8, to be a projection.

As before we perform the same operation on ψ and χ to get projections ψ' and χ' from U and V to the same Z. We must be slightly more careful here to make sure that there is a $\zeta|\zeta$ hinge in U (say) between U_0 and the first α-weld of U, at which to unbend. (Certainly there is a $\zeta|\zeta$ hinge in U since some point of U must be mapped to m, but unbending at an arbitrary such hinge may produce too large an image.)

However, Lemma 6.6 assures us that ψ "starts" just as φ does, from $X_{\varphi(0)}$ to the right, and there is no place other than a $\zeta|\zeta$ hinge for it to turn back to the left. A similar argument applies to V, and the proof of the diamond lemma is (finally) complete. ☐

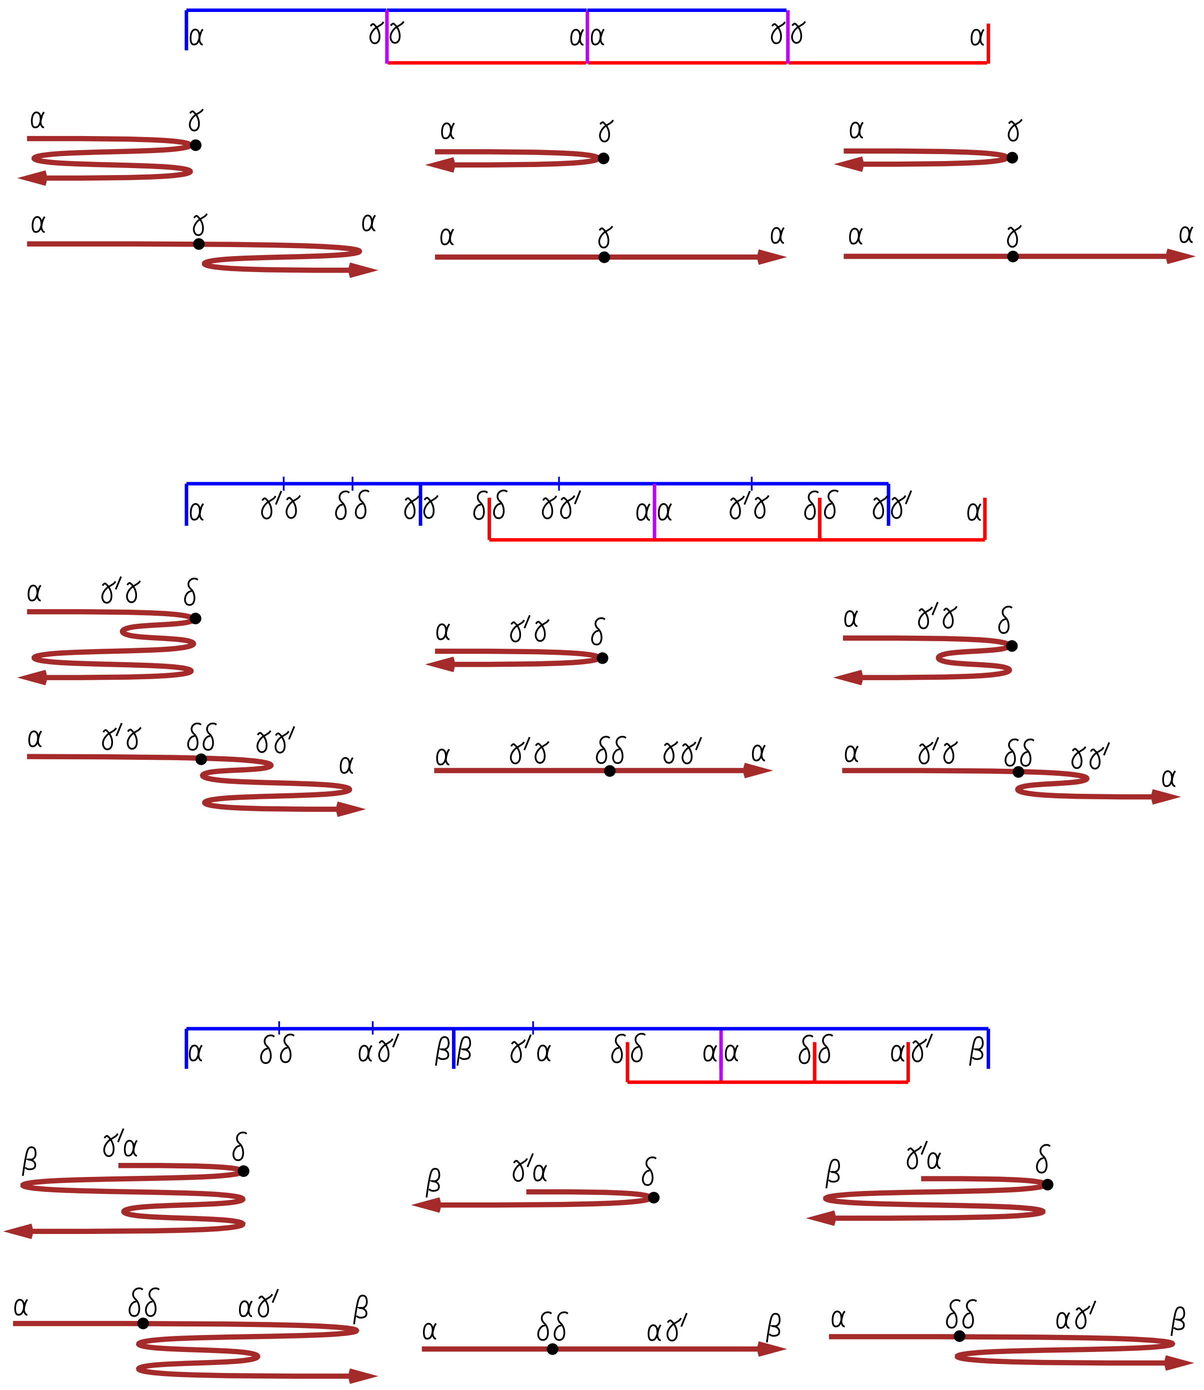

Figure 8: Three more ways to label W, in which φ, ψ and χ are altered to create projections with a common image

It remains only to wind up the proof of Theorem 6.1 from the diamond lemma. If the theorem is false then there are walks X, Y and W such that $W \to X$ and $W \to Y$, with X and Y distinct and irreducible (permitting no folds). Assume further that W has minimum length subject to this property.

Let $W \succ U \to X$ and $W \succ V \to Y$. Then the diamond lemma provides a walk Z with $U \to Z$ and $V \to Z$. Since $|U| < |W|$ and $|V| < |W|$, the minimal reductions U^* and V^* exist and must be equal to X and Y respectively. But then $Z \to X$ and $Z \to Y$, again contradicting the minimality of $|W|$. □

7 Navigable Graphs and Reversible Markov Chains

With the Diamond Lemma in hand we are in a position to classify the navigable graphs G. It will turn out that the addition of positive real weights to the edges of G (so that the probability distribution for a token stepping from vertex u is proportional to the weights of the neighbors of u) has no effect on the classification or its proof, so we can allow $\mathbf{X}$ and $\mathbf{Y}$ to be sample walks from a general finite-state, reversible Markov chain.

We first consider the case where G is *almost* as simple as K_3, in the sense that it has just one additional vertex which is connected via a pendant edge. Specifically, Let D (for "diode") be the graph consisting of vertices a, b, c d with $a \sim b \sim c \sim a$ and $c \sim d$; fixed, positive real weights are assigned to the four edges of D.

Let X be a random walk on D of length n on D. Recall that X^* is our notation for the unique walk, guaranteed by Theorem 6.1, such that $V \to X^*$ for any V with $X \to V$. Let Z be any fixed finite walk on D.

Lemma 7.1. *There exists a real $\alpha < 1$ (depending on the edge-weights of the diode) such that for any Z and any n, $\Pr(X^* = Z) < \alpha^n$.*

Proof. We begin by dividing X into segments $\sigma_0, \ldots, \sigma_t$, each of which is a maximal subwalk which begins and ends with a c and contains either c's and d's only, or c's, a's and b's only. (We can assume there is a c at each end of X so that the segments cover X completely.) Note that segments are defined so that they are forced to alternate in type between those containing d's and those containing a's and/or b's.

Now each segment σ_i has a uniquely determined minimal projection image σ'_i. If σ_i contains d's then $\sigma'_i = cdc$; otherwise both σ_i and σ'_i are walks on the cycle $C_3 := \{a, b, c\}$. Therefore they lift uniquely to walks beginning at 0 and ending at some point $3w_i$ on the doubly infinite path $\mathbb{Z}$, labeled $\ldots cabcabcabc \ldots$ with a c at 0. The integer w_i is the common "winding number" of σ_i and σ'_i.

If $U = U_0, \ldots, U_i$ and $V = V_0, \ldots, V_j$ are any two walks with $U_i = V_0$, it is convenient to denote by $U|V$ the concatenation of U and V with endpoints overlapping; thus $U|V := U_0, \ldots, U_{i-1}, U_i = V_0, V_1, \ldots, V_j$. Let $X' := \sigma'_0|\sigma'_1|\cdots|\sigma'_t$; then $X \to X' \to X^*$.

To each reduced segment σ'_i we associate an even simpler segment which we call σ_i^-, as follows: if $\sigma'_i = cdc$ then $\sigma_i^- := \sigma'_i$; if $\sigma'_i \neq cdc$ and σ'_i has winding number $w_i > 0$, then $\sigma_i^- := c(abc)^{w_i}$; if $w_i < 0$ then $\sigma_i^- := (c(abc)^{w_i})^{-1} = c(bac)^{w_i}$. If $w_i = 0$, σ_i^- becomes the null segment. We define $X^- := \sigma_0^-|\sigma_1^-|\cdots|\sigma_t^-$ except that the elimination of ab-type segments with winding number 0 may bring some d-type segments into juxtaposition, and those will be collapsed into a single "cdc". Thus the segments of X^- will still alternate in type.

Claim 7.2. $X^- \to (X^*)^-$.

Proof. Let f be the minimum number of folds necessary to reduce X' to X^*; we proceed by induction on f.

Let $\varphi : X' \to Y$ be a fold such that only $f-1$ folds are necessary to reduce Y to $Y^* = X^*$. Then our induction hypothesis implies that $Y^- \to (X^*)^-$, so we need only demonstrate that $X^- \to Y^-$.

The projection φ cannot bend at a segment border since each such $c \in X'$ has a d on one side and not the other. Suppose that φ bends inside segments σ'_k and σ'_{k+j}; note that $j \neq 0$ since the segments of X' are irreducible.

It follows that the borders of φ are in segments σ'_{k-j} and σ'_{k+2j}, and:

- each of σ'_k and σ'_{k+j} is either of type d or has winding number 0;
- $\sigma'_{k-j+i} = (\sigma'_{k+j-i})^{-1} = \sigma'_{k+j+i}$ for each i, $1 \leq i \leq j-1$;
- $Y = \sigma'_0|\sigma'_1|\cdots|\sigma'_{k-2}|\sigma'_{k-1}|\sigma'_{k+2j}|\sigma'_{k+2j+1}|\cdots|\sigma'_t$.

If both bends of φ are at d's, then folding X^- by bending at the same d's produces Y^-. If it happens that those two d's are collapsed in X^- on account of intervening segments having winding number 0, then already $Y^- = X^-$.

If σ'_k is not of type d then, since it has winding number 0, it collapses in X^-; we bend X^- inside the d-type segment which results when σ'_k's two neighboring segments are identified. The same applies when σ'_{k+j} is not of type d, or when both bends of φ occur in non-d segments, reducing all three cases to the previous one. □

We have succeeded in showing that if $X^* = Z$ then $X^- \to Z^-$; we now show that the latter has exponentially small probability when Z is fixed and X random of length n. Suppose that the non-type-d segments of X^- are $\mu_1, \ldots, \mu_s$ with winding numbers $u_1, \ldots, u_s$, and the non-d segments of Z^- are $\nu_1, \ldots, \nu_t$ with winding numbers $v_1, \ldots, v_t$.

Observe that a projection π from X^- can bend only at d's; thus it maps each non-d segment μ_i of X^- either forward (which we record by setting $\pi[i] = +1$) or backward (-1) onto the segment ν_j where $j = 1 + \sum_{h=1}^{i-1} \pi[h]$. Let us fix any of the 2^s maps $\pi[\cdot] : \{1, \ldots, s\} \to \{+1, -1\}$.

In order for this $\pi[\cdot]$ to correspond to a legitimate projection, it must match segments with the right winding numbers, that is, if μ_i maps to ν_j then $v_j = \pi[i] \cdot u_i$. The winding numbers w_i of the non-d segments of X are i.i.d. *symmetric* random variables, since X is a reversible Markov chain, and the same therefore applies to X^- except that u_i is never zero.

With $\pi[\cdot]$ still fixed and X random, we have $\Pr(u_i = u) \in (0, p]$ for each $u \neq 0$ and some $p < \frac{1}{2}$; thus the probability that $\pi[\cdot]$ projects onto $\mathbb{Z}^-$ is at most p^s. It follows that $\Pr(X^- \to Z^-) < 2^s \cdot p^s = \beta^s$, where $\beta = 2p < 1$.

It remains only to link s with n, but this is easy; if q is the stationary probability of state d and r the probability that between two d's the a's, b's and c's wind nontrivially, then $\Pr(s \leq Cqrn)$ is exponentially small for any $C < 1$. This concludes the proof of Lemma 7.1. □

Before going on to show that D is navigable, we take another look at the relation between projections and prejections. If W is any walk, we let $W[i,j]$ be the subwalk of W running between indices i and j, inclusive.

Lemma 7.3. *Let X be a walk of length m on any graph. If $X \hookrightarrow Z$ then there are indices $i \leq j$ such that $X[i,m]^* = Z^* = (X[0,i]^{-1}|X[j,m])^*$.*

Proof. Let μ be a prejection from X onto Z; choose any i for which $\mu(i) = 0$. Then $\mu \restriction [i,m]$ projects $X[i,m]$ onto Z, hence $X[i,m]^* = Z^*$ as claimed. By the "discrete intermediate value theorem" applied to $\mu \restriction [i,m]$, there is an index $j \geq i$ for which $\mu(j) = \mu(0)$; then $\mu \restriction ([0,i] \cup [j,m])$ projects $X[0,i]^{-1}|X[j,m]$ onto Z and the lemma follows. □

Theorem 7.4. *The diode is navigable.*

Proof. Fix weights on the diode D, letting $\alpha < 1$ be as in the statement of Lemma 7.1. Let $\mathbf{X}$ and $\mathbf{Y}$ be infinite random walks sampled from the reversible Markov chain determined by D. Let $B(m,n)$ stand for the event that the fickle demon is blocked by $\mathbf{X} \restriction [0,m]$ and $\mathbf{Y} \restriction [0,n]$, in other words, that there is a walk Z such that $\mathbf{X}[0,m] \hookrightarrow Z$ and $\mathbf{Y}[0,n] \hookrightarrow Z^{-1}$. We begin by showing that $\Pr(B(m,n))$ is exponentially small in $m+n$.

Let us first fix $\mathbf{Y}$ and assume that $B(m,n)$ holds, with $\nu : \mathbf{Y}[0,n] \hookrightarrow Z^{-1}$ and $\nu(k) = 0$. Then $X[0,m] \hookrightarrow Z^*$, thus $X[i,m]^* = Z^* = (X[0,i]^{-1}|X[j,m])^*$ for some $i,j \le m$. Since at least one of $X[i,m]$ and $X[0,i]^{-1}|X[j,m]$ has length $\ge m/2$, we have, regardless of $\mathbf{Y}$, n and k, that

$$\Pr(B(m,n)) < m^2\alpha^{m/2} .$$

Applying the same argument with $\mathbf{X}$ fixed instead, we have

$$\Pr(B(m,n)) < \min\{m^2\alpha^{m/2}, n^2\alpha^{n/2}\} < (m^2+n^2)\alpha^{(m+n)/4} .$$

It follows that $\sum_{m,n}^{\infty} \Pr(B(m,n)) < \infty$; to reduce this sum to less than 1, we employ the technique used previously, assuming that $\mathbf{X}$ and $\mathbf{Y}$ each begin with particular uncontractible walks U and V of sufficient length.

For the diode, we can take $U = (c(abc)^9d)^k$ and $V = (c(abc)^8d)^k$, choosing k so that

$$\sum_{m=29k}^{\infty} \sum_{n=26k}^{\infty} (m^2+n^2)\alpha^{(m+n)/4} < 1 .$$

We omit the routine verification that $B(m,n)$ is not increased when $\mathbf{X}$ is assumed to begin with U and $\mathbf{Y}$ with V. $\square$

If an edge (say, between b and d) is added to D we obtain another graph D^+. For walks on D^+, segments are defined as maximal subwalks beginning and ending with c's and containing no d's, together with the ("type d") c-to-c subwalks between them. Elimination of every a and b appearing in a type-d segment cannot free a pair of finite walks which was blocked, nor does it reduce the length of the walks by more than a constant factor, so we are reduced to the diode case. (One can also apply the argument of the diode case directly, with $\sigma^* = cdc$ when σ is a type-d segment, regardless of its complexity.) The result:

Theorem 7.5. *The graph D^+ is navigable.*

We now have the tools with which to classify all the navigable graphs. Let H be a subset of the nodes of a graph G. We define a graph structure on H not in the usual way, but as follows: $u \bowtie v$ if there is a path P from u to v in G which contains no other vertex of H. The point of defining adjacency in H in this manner is that if W is a (finite) walk on G then $W' := W \restriction H$ is a walk on $\langle H, \bowtie \rangle$, which we denote simply by H, and conversely every walk on H is the restriction to H of some walk on G.

We say that G is "substantial" if there is a four-point subset $H = \{a,b,c,d\}$ such that the diode D is contained in $\langle H, \bowtie \rangle$ as a (not necessarily induced) subgraph.

Let C_n denote the cycle on n vertices, P_n the path on $n+1$ vertices, and $K_{1,3}$ the star consisting of one vertex of degree 3 and three pendant vertices of degree 1.

Lemma 7.6. *Let G be a connected graph which is not substantial; then either $G = C_n$ for some $n \ge 3$, or $G = P_n$ for some $n \ge 0$, or $G = K_{1,3}$.*

Proof. It is easy to check that C_n, P_n and $K_{1,3}$ are not substantial; suppose that G is not isomorphic to any of these. If G is a tree (and not a path) it contains a vertex of u of degree at least 3; let a, b and c be any three of its neighbors. Since $G \ncong K_{1,3}$ there is another vertex $d \notin \{a,b,c,u\}$, but adjacent to either u or one of a, b, and c; say c. Either way, we have our H.

If G contains a cycle, fix some $C_n \subset G$, and let d be any vertex of G such that $d \notin C_n$ but $d \sim c \in C_n$. Now we let a and b be any two other vertices on C_n. $\square$

We have already seen from Theorem 3.3 and Corollary 3.4 that paths and cycles are not navigable, and these results are easily extended to the case where edges are weighted. The graph $K_{1,3}$ is for our purposes the same as K_3, since neither token can be permitted to stay on the central vertex while the other moves. Thus it is immediate that all navigable graphs are substantial.

To show the converse, assume G is substantial and reduce any finite walk W on G to a walk W^D on $\langle H, \bowtie \rangle$ by eliminating all nodes from W other than a, b, c and d. Note that W^D is also a sample walk from a reversible Markov chain, since $(W^D)^{-1} = (W^{-1})^D$. We now observe that if $W \hookrightarrow Z$ then $W^D \hookrightarrow Z^D$, and it follows that navigability of G is implied by navigability of $\langle H, \bowtie \rangle$. However, the only possibilities for $\langle H, \bowtie \rangle$ are D, D^+ and K_4 and all three are navigable. Thus:

Theorem 7.7. *G is navigable if and only if it is not isomorphic to a path, a cycle or $K_{1,3}$.*

8 General Markov Chains

If $\mathbf{X}$ and $\mathbf{Y}$ are obtained from a non-reversible Markov chain, slightly greater care must be taken with starting states; the argument that the initial positions of the X and Y tokens are unimportant no longer applies when there are transitions which are permitted in one direction and not the other. Thus, we assume that the starting positions are chosen randomly (e.g. via the stationary distribution of the chain); the probability distribution does not matter as long as all states are possible. This allows us to eliminate any transient states, thus we can assume that the underlying digraph $\vec{G}$ of the Markov chain is strongly connected.

Since non-reversible Markov chains constitute a richer class than reversible ones, it is perhaps a bit surprising that:

Theorem 8.1. *Every non-reversible Markov chain is navigable.*

Proof. The "cycle-reversing identity" of [3] states that in a reversible chain started at state c, the expected time to hit states a, then b, then c again is the same as for the order b, a, c. Tetali [8] has observed that in fact the cycle-reversing identity *characterizes* reversible chains; in any non-reversible chain, there are states a, b, c which are more often traversed in the order c, a, b, c than in the order c, b, a, c.

Let $\mathbf{W}$ be an infinite sample walk from a non-reversible chain and $\mathbf{W}^C$ the result of deleting every node of $\mathbf{W}$ other than a, b and c; again, since $W \hookrightarrow Z$ implies $W^C \hookrightarrow Z^C$, we are reduced to showing navigability for a three-state Markov chain with positive torque. Let $\mathbf{X}$ be an infinite random walk from this reduced chain.

The winding number of $\mathbf{X}[0,m]$ is a random walk on $\mathbb{Z}$ with positive drift, hence with positive probability $\mathbf{X}$ begins $c, a, b, c \ldots$ and $\mathbf{X}[0,m]$ has strictly positive winding number for every $m \geq 4$. It follows that for every m, if $\mathbf{X}[0,m] \hookrightarrow Z$ then Z^* is an initial segment of the string $(cab)^\infty$.

We apply a similar argument to $\mathbf{Y}$, concluding that with positive probability, $\mathbf{Y}$ begins $a, b, c, a \ldots$ and $\mathbf{Y}[0,n]$ has positive winding number for every $n \geq 4$. Then for every n, if $\mathbf{Y}[0,n] \hookrightarrow Z^{-1}$, $(Z^*)^{-1}$ must be an initial segment of $(abc)^\infty$. Since no initial segment of $(cab)^\infty$ is equal to the reverse of an initial segment of $(abc)^\infty$, percolation is assured. $\square$

9 The Clairvoyant Demon

We conclude with some remarks about the failure (so far) of our techniques to prove *oriented* percolation, even for $G = K_{1,000}$. Even a quick look at the obstacles facing the clairvoyant demon suggests the following analog to projections and prejections:

Let $U = U_0U_1 \dots U_m$ and $V = V_0V_1 \dots V_n$ be finite walks on G. A map π from $\{0, 1, 2, \dots, n\}$ *onto* $\{0, 1, 2, \dots, m\}$ will be called a *production* if

(*i*) $\pi(0) = 0$ and $\pi(n) = m$;
(*ii*) $\pi(i+1) \leq \pi(i) + 1$ for every i with $0 \leq i < n$;
(*iii*) $U_{\pi(i)} = V_i$ for every i, $1 \leq i \leq n$.

The simplest nontrivial kind of production π is one which satisfies $\pi(i+1) \leq \pi(i)$ for precisely one value of i; this we will call a *shift*.

If, in the definition of production, condition (*i*) is weakened to just $\pi(n) = m$, we have a *pre*duction.

A blocker for oriented percolation, when projected downward onto the X-axis, induces a preduction onto (the reverse of) an initial segment of $\mathbf{X}$ and similarly for $\mathbf{Y}$. It is straightforward to show that if for some Z and initial segments X of $\mathbf{X}$ and Y of $\mathbf{Y}$ there are preductions from X to Z and from Y to Z^{-1}, then the clairvoyant demon is blocked. However:

1. the clairvoyant demon can be blocked in other situations as well;
2. the composition of two productions may not be a production, and similarly for preductions;
3. not every production can be decomposed as the composition of shifts; and
4. the diamond lemma fails both for productions and for compositions of productions.

There is a piece of good news; the semigroup reduction rule

$$UVVW \to UVW$$

suggested by productions, when expanded to a congruence, creates a well-known semigroup called the *free band* (see e.g. [7]). However, it turns out that in the free band two words of forms UVW and $UV'W$ are equivalent under very weak conditions, for example whenever each of U and W hits every vertex of G at least once. For any G, it will happen with probability 1 that for some initial segment X of $\mathbf{X}$ and Y of $\mathbf{Y}$, $X = UVW$ and $Y^{-1} = UV'W$ with these conditions obtaining. Thus the full congruence generated by productions is too big to allow the clairvoyant demon through; perhaps something tractable between it and the (X, Y, Z)-production condition can be found.

On top of these difficulties, it turns out that the clairvoyant demon cannot achieve the same exponentially-diminishing blocking probability enjoyed by the fickle demon, and commonly observed in percolation problems. The following result was recently communicated to us by Peter Gács:

Theorem 9.1. (Gács [5]) *Suppose for a given fixed G that the clairvoyant demon percolates with positive probability. Then there is a constant $\alpha = \alpha(G)$ such that the probability that he can reach distance n but is ultimately blocked is at least $n^{-\alpha}$.*

Despite all this, we still believe the clairvoyant demon wins even on K_4.

Acknowledgments: The author is indebted to a great many of his colleagues for helpful insights concerning both the fickle and clairvoyant demons. Included are Noga Alon, Graham Brightwell, Charles Grinstead, Sasha Kostochka, Prasad Tetali, and all of the members of the Institute for Elementary Studies. Thanks are due also to the referees for the detection and resulting elimination of many errors and ambiguities.

References

[1] D.J. Aldous and J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, to appear (1998). Draft on web at http://www.stat.berkeley.edu/users/aldous/book.html.

[2] P.N. Balister, B. Bollobás and A.N. Stacey, Dependent Percolation in Two Dimensions, draft (1999).

[3] D. Coppersmith, P. Tetali and P. Winkler, Collisions among random walks on a graph, *SIAM J. Disc. Math.* **6** #3 (Aug. 1993), 363–374.

[4] P.G. Doyle and J.L. Snell, *Random Walks and Electric Networks*, Mathematical Assoc. of America, Washington, DC 1984.

[5] P. Gács, Prescient scheduling (work in progress), preprint (1999).

[6] G. Grimmett, *Percolation*, Springer-Verlag, Heidelberg (1989).

[7] J.M. Howie, *Fundamentals of Semigroup Theory*, London Math. Soc. Monograph (New Series) #12, Clarendon Press, Oxford (1995)

[8] P. Tetali, An extension of Foster's network theorem, *Comb., Prob. & Comput.* **3** (1994), 421–427.