Random Walks and Other Homomorphisms

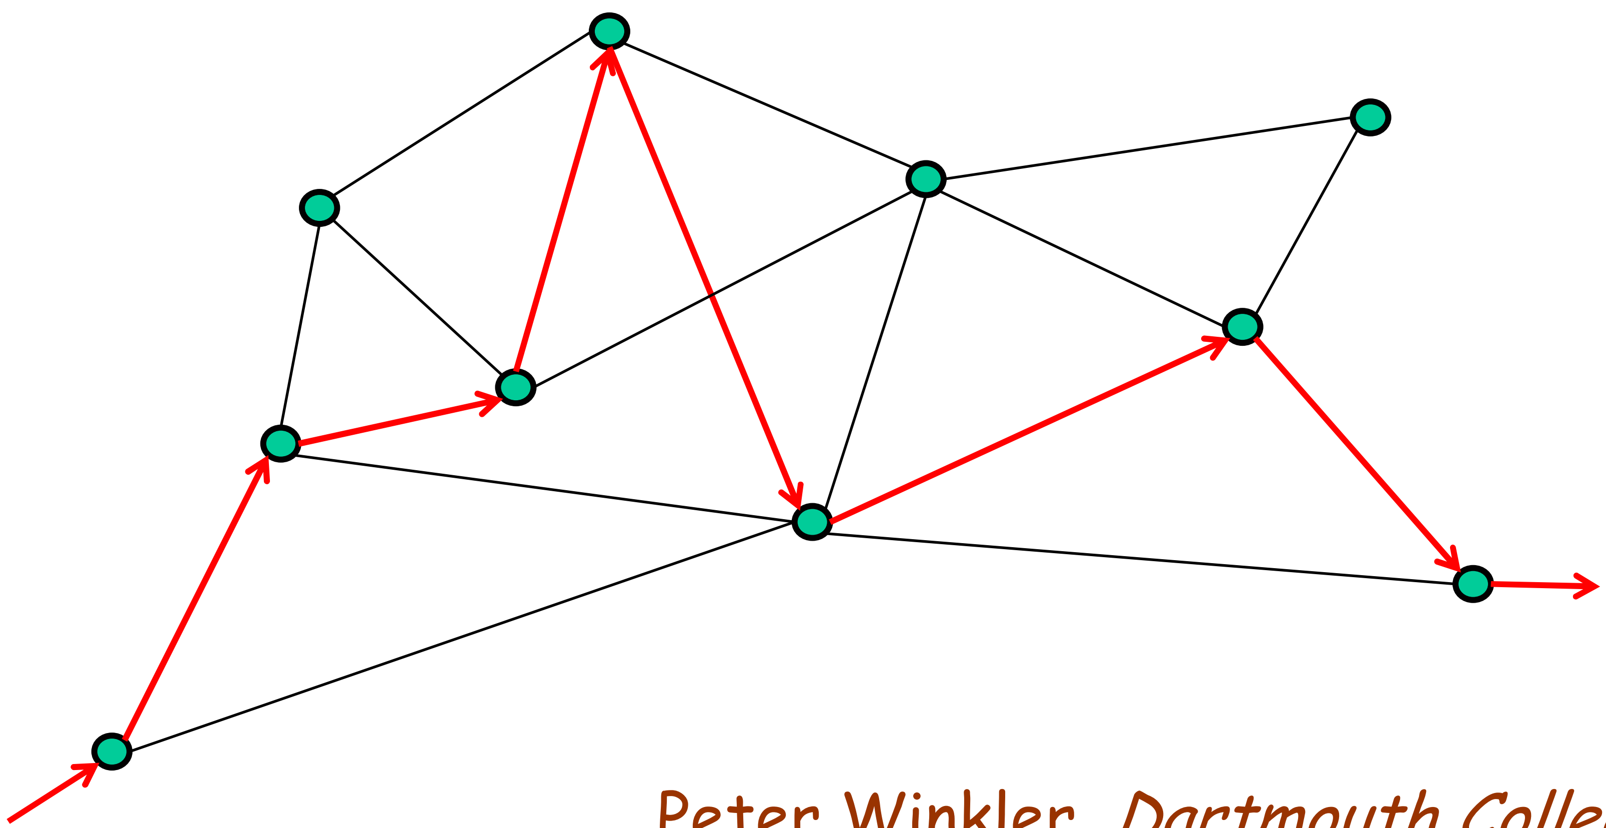

Peter Winkler, *Dartmouth College*

Random walk niceties:

Markov chain (usually ergodic)

Stationary distribution: proportional to degree

Reversible and torqueless: CTW identity

Hitting probabilities and expected hitting times given by effective resistance

Lots of applications: graph search, randomness amplification, universal cover, MCMC, spanning tree generation, protein-matching algorithms, polymers . . .

Random walk questions:

Cover time: How long does it take to hit every vertex of H ?

Mixing time: How long does it take to "get lost" in H ?

Cover time results:

Upper bound for covering all vertices: n^3 [Aleliunas et al. '79, Chandra et al. '89] reduced to $4n^3/27$ [Feige '95].

Lower bound for covering all vertices: $(1-o(1))\, n \ln n$ [Feige '95].

Upper bound for covering all edges, in terms of #edges (or total edge length): $2m^2$ to hit each edge one way, $3m^2$ both ways [Georgakopoulos & W. '11].

Two difficult new results:

PTAS for the expected cover time of a tree

Proof of the "blanket time conjecture" [Zuckerman & W. '96]: there is a universal constant *C* such that the expected time to hit all vertices a representative number of times is at most *C* times the cover time.

Both due to [Ding, Lee and Peres '10] using results of Talagrand and others.

Generalizing toward Statistical Physics?

Random walks on H are random maps from a path to H. What about random homomorphisms from other graphs G (sometimes called G-indexed random walks)?

These are rampant in statistical physics, but the usual questions asked about them are quite different from the random walk case.

Statistical physics	Graph Homomorphisms
hard-core model	random independent sets
monomer-dimer	random matchings
Potts model	random colorings
percolation	random subgraphs
linear polymers	self-avoiding random walks
branched polymers	random lattice trees
ice model	random Eulerian orientations

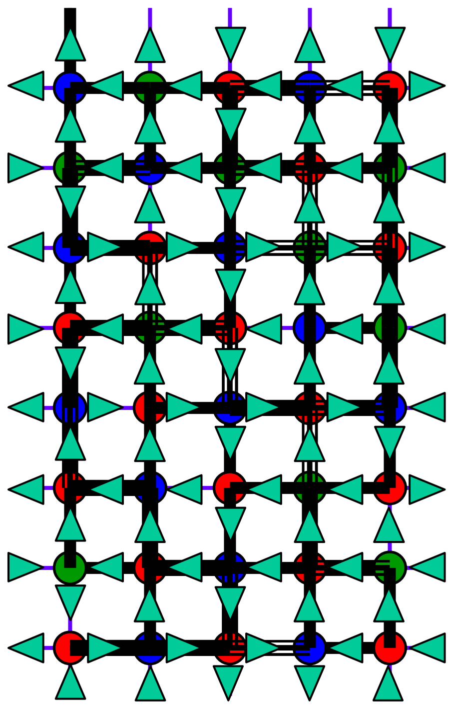

Hard constraints for nearest neighbors

"Nearest neighbor" hard constraint models abound and are particularly suitable for combinatorial analysis.

A typical situation: A small "constraint graph" H is fixed. You are given a large (possibly infinite) graph G, often a lattice, and wish to study random homomorphisms from G to H.

You can associate an activity to each vertex of H. The relative probability of a homomorphism (when G is finite) is determined by the product of the activities of the images.

Constraint graphs for particular models

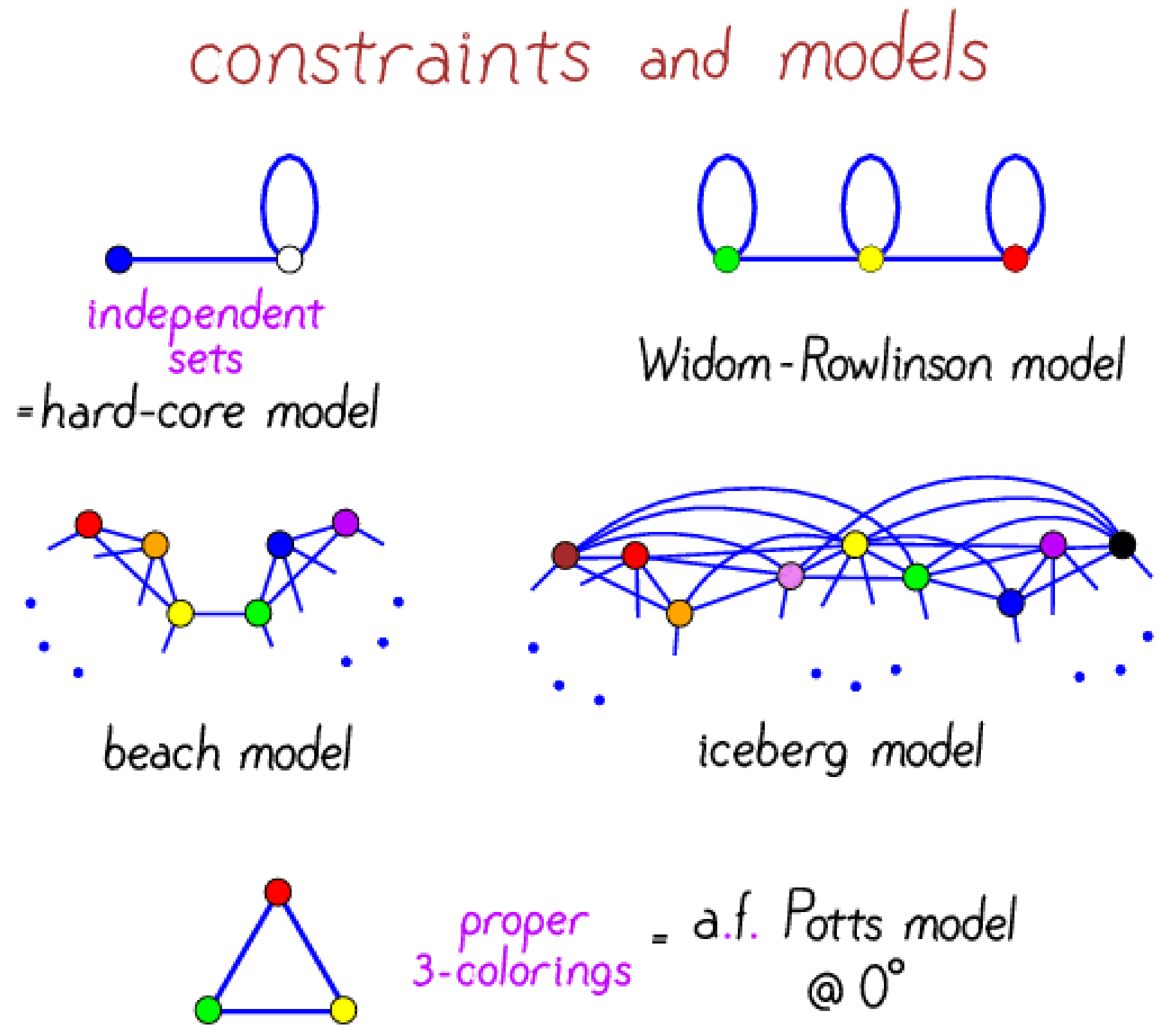

Which graphs cause a phase transition?

On the Bethe lattice (the k-regular tree, with $k > 2$):

Theorem: phase transition iff constraint graph H contains one of

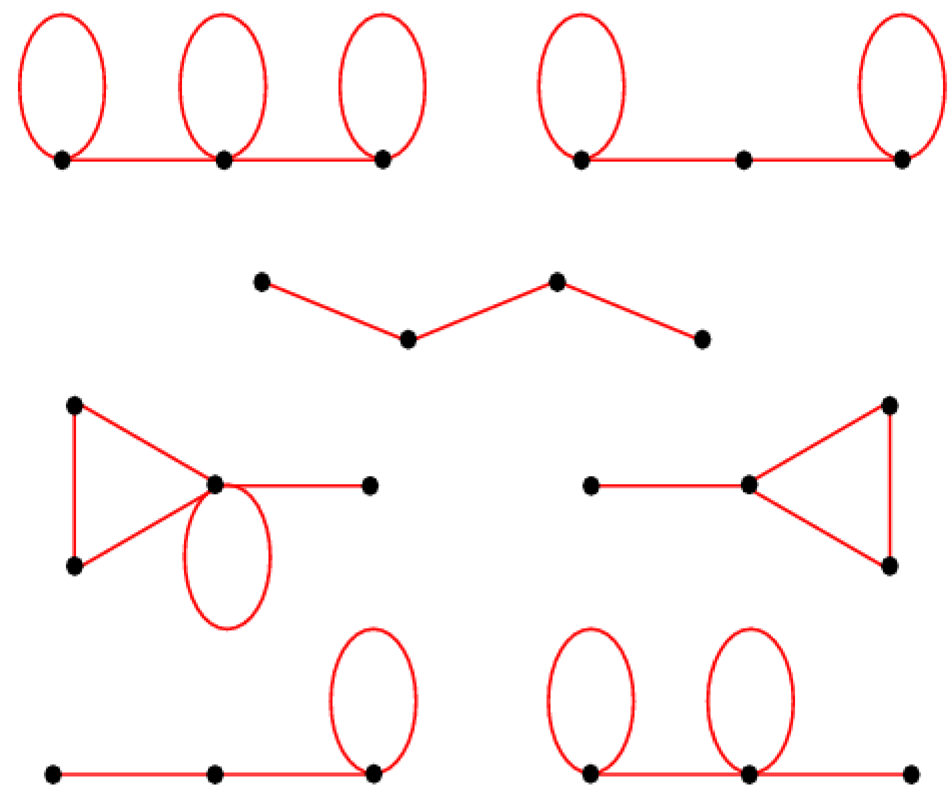

[*Brightwell & W.* '99]

A surprise consequence

Theorem: [*Brightwell, Haggstrom & W.* '99] In the discrete hard core model (and the Widom-Rowlinson model) it is possible to have multiple critical points---that is, one can go in and out of the ordered regime as activity increases.

Two shortcomings of the B.-W. theory: (1) It considers only "nice" phase transitions (involving simple, invariant Gibbs measures) and (2) It's on the Bethe lattice, which does not resemble low-dimensional space.

Branching random walks: mixing vs recovery

For ordinary random walks: Question is: how long does it take to get lost?

For branching random walks: Question is: do you ever get lost? In other words, can you tell from where you are, where you started? (Spitzer recovery, random recovery: statistical, pattern, strong)

But, for branching random walks: Maybe we should ask: if you do get lost, how long does it take? How long before your "tokens" cover H uniformly?

What if we replace the tree by a graph?

Theorem: TFAE for any fixed restraint graph H:

(1) For any infinite bounded-degree G, there are activities for which G->H has no phase transition;

(2) For any *finite* G, the graph of homomorphisms from G to H is connected;

(3) H is *dismantlable* (cop-win for the "active game").

[*Brightwell & W.* '00]

Switching domains and questions

What if we ask random-walk-type questions for the more general random homomorphisms, and stat-phys-type questions for random walks?

Example: [Brightwell & W. '00] define H to be *4-mobile* (!) if Hom(G,H) is connected whenever G has max degree 2.

This is basically a statement about walks (open and closed) on H . But there's an embarrassing open problem: how can you tell if H is 4-mobile??

Switching domains and questions, cont.

Here's a reverse problem: "cover time" for random homomorphisms from a class G of graphs other than the class of paths.

We ask the following way: how big a graph *G* from G do you need so that a random homomorphism from *G* to your fixed target H is onto?

Example: suppose G is the class of discrete hypercubes, and H is a path of length 6.

Switching domains and questions, cont.

Pictorially:

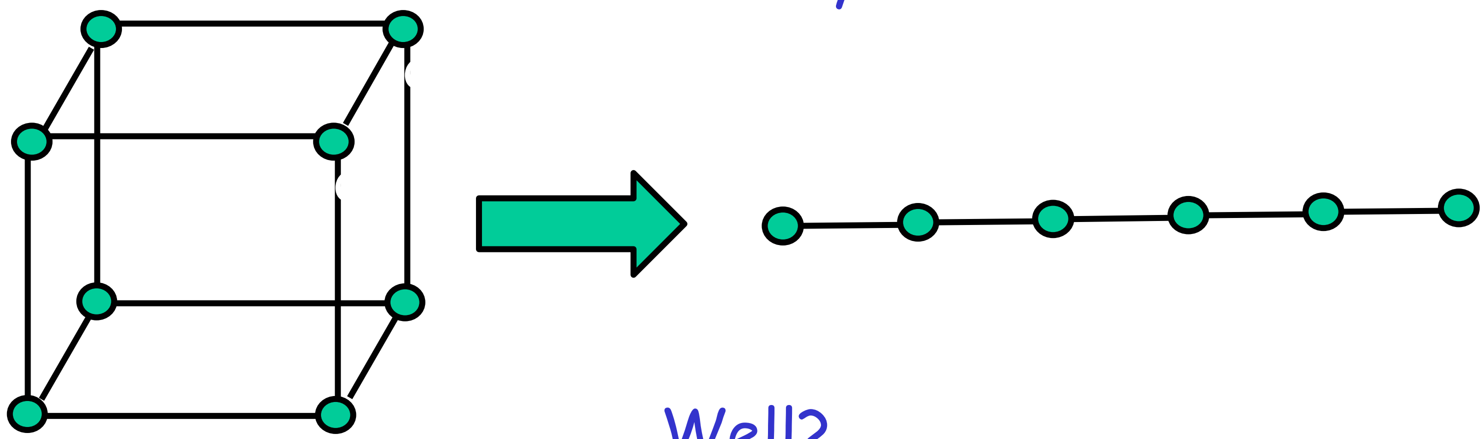

Well?

Answer: No matter how big the hypercube is, its image on the path will be (w.h.p.) only of size 5. [Kahn '01, Galvin '03]

[Benjamini et al. '07]: does not happen if graphs in G have sub-logarithmic degree.

On the other hand ...

Suppose G of is the class of b-ary trees of depth n (e.g., partial Bethe lattices), for some $b > 1$.

Then you're taking a branching random walk on H, and your cover time ought to be much smaller than with paths---often, not much more than the diameter of H. Is it?

More generally (but also more vaguely): What can you say about the probability that a random homomorphism from G to H is onto? Which domain graphs are more likely to cover?

Does mixing make sense?

Again suppose G is a class of graphs, one of which is G.

Suppose we have a random homomorphism from G to H with u v. Let w be some other vertex of G. What is the probability distribution of its image? When is it the case that as w gets farther from u, these distributions converge?

If they do converge, how fast?