Phase Transitions in Local Search for Satisfiability

Pekka Orponen (joint work with Sakari Seitz and Mikko Alava)

Laboratory for Theoretical Computer Science

Helsinki University of Technology

Outline

- **Part I: Background – Combinatorial Phase Transitions**
- "Where the Really Hard Problems Are?"
- Hard instances for 3-SAT
- The Davis-Putnam procedure
- Statistical mechanics of k-SAT
- **Part II: Local Search – Methods & Experiments**
- WalkSAT and related algorithms
- Record-to-Record Travel and variants
- Focused Metropolis Search
- Analysis?

“Where the Really Hard Problems Are?”

Yu & Anderson (1985), Hubermann & Hogg (1987), Cheeseman, Kanefsky & Taylor (1991), Mitchell, Selman & Levesque (1992), Kirkpatrick & Selman (1994), etc.

Many NP-complete problems can be solved in polynomial time “on average” or “with high probability” for reasonable-looking distributions of problem instances. E.g. Satisfiability in time $\mathcal{O}(n^2)$ (Goldberg et al. 1982), Graph colouring in time $\mathcal{O}(n^2)$ (Turner 1988).

Where, then, are the (presumably) exponentially hard instances of these problems located? Could one tell ahead of time whether a given instance is likely to be hard?

Hard instances for 3-SAT

Mitchell, Selman & Levesque (1992): Experiments on the behaviour of the Davis-Putnam procedure on randomly generated 3-cnf Boolean formulas.

E.g. satisfiable 3-cnf formula

$$(x_1 \vee \bar{x}_2 \vee x_3) \wedge (\bar{x}_1 \vee x_2 \vee \bar{x}_4)$$

The expressions in parentheses are *clauses* and the x's are *literals*.

Distribution of test formulas:

- number of variables n
- $m = \alpha n$ randomly generated clauses of 3 literals, $2 \leq \alpha \leq 8$

The Davis-Putnam procedure

The Davis-Putnam[-Logemann-Loveland] (DP[LL]) method for testing satisfiability of set of clauses Σ on variable set V:

1. If Σ is empty, return "satisfiable".
2. If Σ contains an empty clause, return "unsatisfiable".
3. If Σ contains a unit clause $c = x^{\pm}$, assign to x a value which satisfies c, simplify the remaining clauses correspondingly, and call DP recursively.
4. Otherwise select an unassigned $x \in V$, assign $x \leftarrow 1$, simplify Σ, and call DP recursively. If this call returns "satisfiable", then return "satisfiable"; else assign $x \leftarrow 0$, simplify Σ, and call DP recursively again.

A combinatorial phase transition?

For sets of 500 formulas with 20/40/50 variables, Mitchell et al. plotted the median number of DP calls required for solution.

Results:

- A clear peak in running times (number of DP calls) near the point where 50% of formulas are satisfiable.
- The "50% satisfiable" point or "satisfiability threshold" seems to be located at roughly $\alpha \approx 4.25$ for large n.
- The peak seems to be caused by relatively short unsatisfiable formulas.

Question: Is the connection of the running time peak and the satifiability threshold a characteristic of the DP algorithm, or a (more or less) algorithm independent "universal" feature?

Running times: all instances

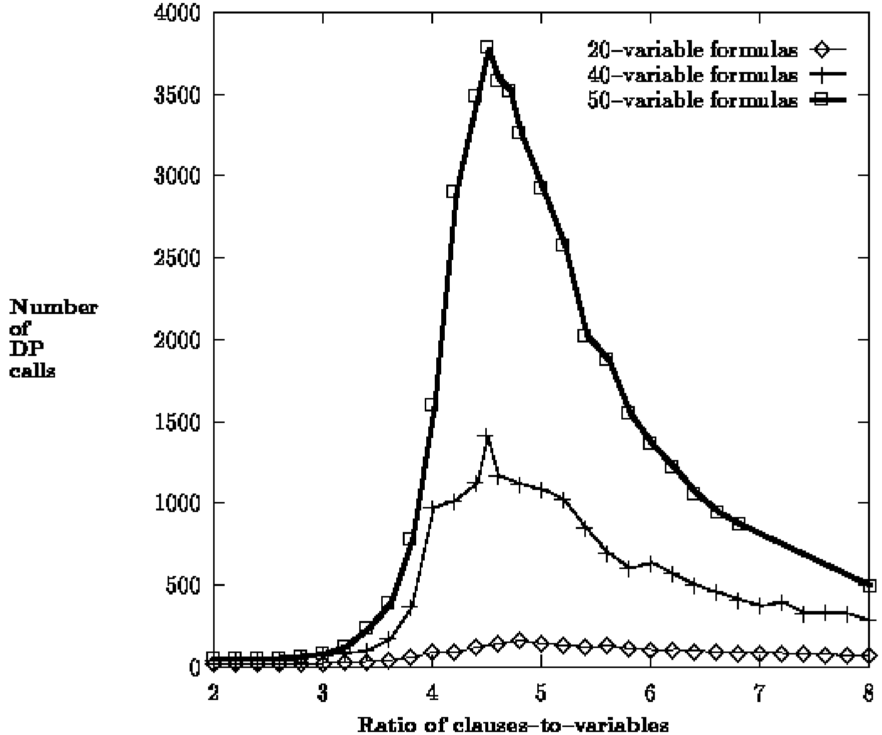

Median number of recursive DP calls for random 3-cnf formulas, as a function of clauses-to-variables ratio α (Mitchell et al. 1992).

Running times: sat vs. unsat instances

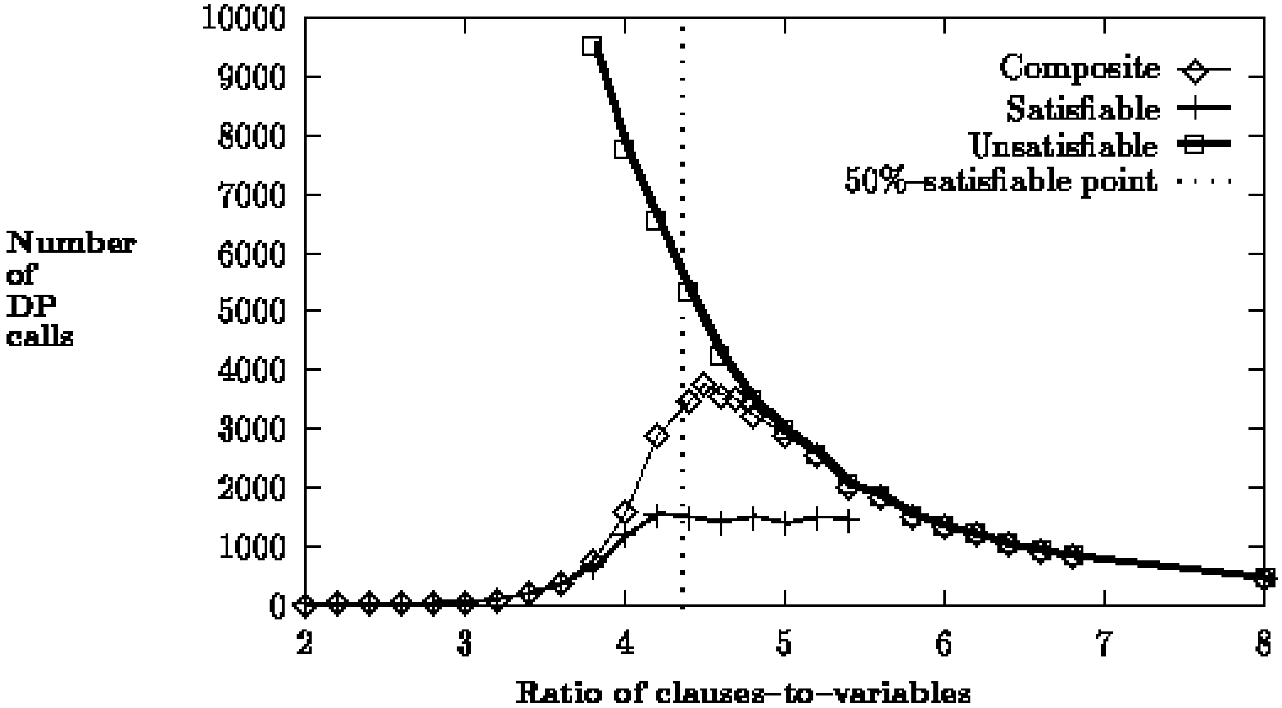

Median number of DP calls for satisfiable and unsatisfiable 50-variable random 3-cnf formulas (Mitchell et al. 1992).

Satisfiability threshold

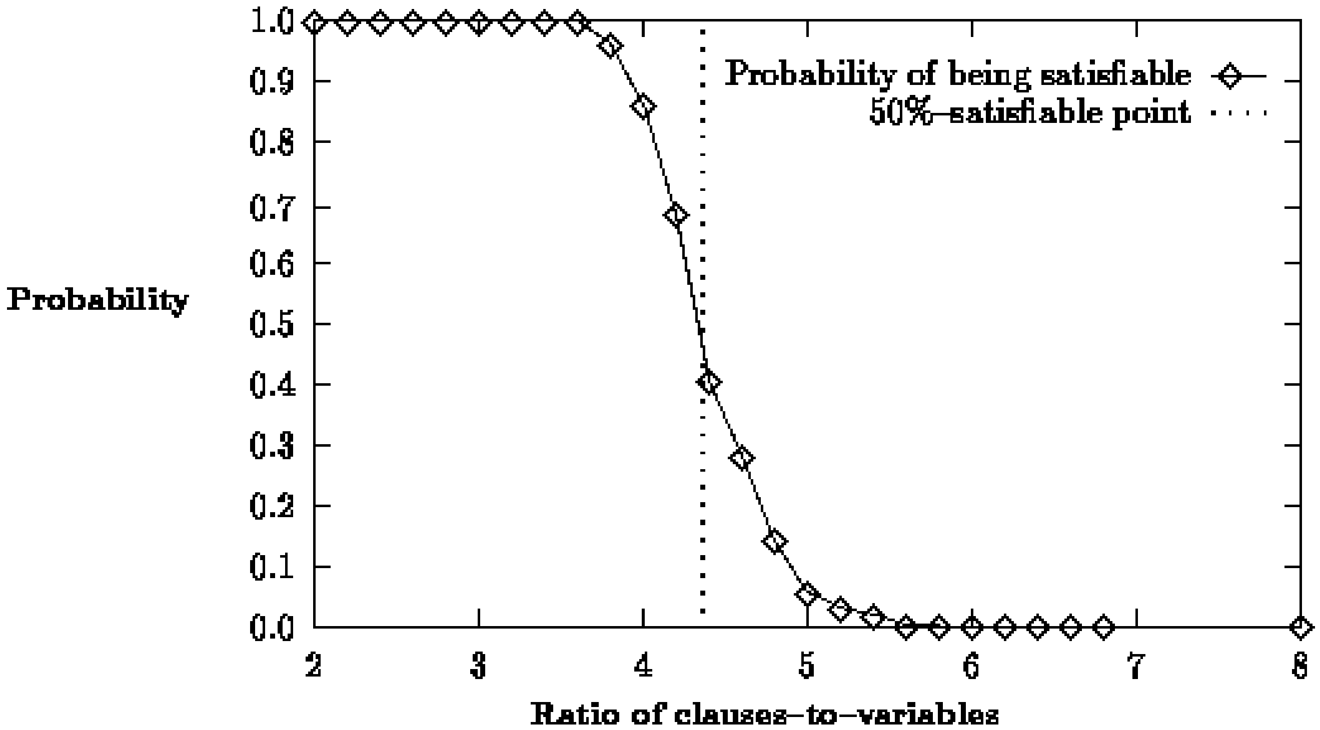

Probability of satisfiability of 50-variable 3-cnf formulas, as a function of clauses-to-variables ratio α (Mitchell et al. 1992).

Extension to k-SAT

Kirkpatrick & Selman (Science 1994).

Similar experiments as above for k-SAT, $k = 2, \dots, 6$; 10000 formulas per data point.

Further observations:

- The "satisfiability threshold" α_c shifts quickly to larger values of α for increasing k.
- For fixed k, the value of α_c drifths slowly to smaller values for increasing n.

Statistical mechanics of k-SAT

Kirkpatrick & Selman (Science 1994).

A statistical mechanics model of a k-cnf formula:

- variables $x_i \sim$ spins with states ± 1
- clauses $c \sim k$-wise interactions between spins
- truth assignment $\sigma \sim$ state of spin system
- Hamiltonian $H(\sigma) \sim$ number of clauses unsatisfied by σ
- $\alpha_c \sim$ critical "interaction density" point for "phase transition" from "satisfiable phase" to "unsatisfiable phase"

Location of satisfiability transition

Estimates of α_c for various values of k via "annealing approximation", "replica theory", and observation:

k	α_{ann}	α_{rep}	α_{obs}
2	2.41	1.38	1.0
3	5.19	4.25	4.17 ± 0.03
4	10.74	9.58	9.75 ± 0.05
5	21.83	20.6	20.9 ± 0.1
6	44.01	42.8	43.2 ± 0.2

A first-order analysis

The "annealing approximation" means simply assuming that the different clauses are satisfied independently. Thus:

- Probability that given clause c is satisfied by random σ: $p_k = 1 - 2^{-k}$.
- Probability that random σ satisfies all $m = \alpha n$ clauses assuming independence: $p_k^{\alpha n}$.
- $E[\text{number of satisfying assignments}] = 2^n p_k^{\alpha n} \triangleq S_k^n(\alpha)$.
- For large n, $S_k^n(\alpha)$ falls rapidly from 2^n to 0 near critical value $\alpha = \alpha_c$. Where is α_c?
- One approach: solve for $S_k^n(\alpha) = 1$.

$$S_k^n(\alpha) = 1 \Leftrightarrow 2p_k^{\alpha} = 1$$

$$\Leftrightarrow \alpha = -\frac{1}{\log_2 p_k} = -\frac{\ln 2}{\ln(1-2^{-k})} \approx -\frac{\ln 2}{2^{-k}} = (\ln 2) \cdot 2^k.$$

Advanced results

It is in fact known that:

- A sharp satisfiability threshold α_c exists for all $k \geq 2$ (Friedgut 1999).
- For $k = 2$, $\alpha_c = 1$ (Goerdt 1982, Chvátal & Reed 1982). Note that 2-SAT $\in$ P.
- For $k = 3$, $3.14 < \alpha_c < 4.51$ (lower bound due to Achlioptas 2000, upper bound to Dubois et al. 1999).
- Current best empirical estimate for $k = 3$: $\alpha_c \approx 4.267$ (Braunstein et al. 2002).

Local search

Naive, but surprisingly useful idea for (combinatorial) optimisation. Assume objective function $E = E(x)$ to be minimised. Then:

- Start with some randomly chosen feasible solution $x = x_0$.
- If value of $E(x)$ is not “good enough”, search for some “neighbour” x' of x that satisfies $E(x') \lesssim E(x)$. If such an x' is found, set $x \leftarrow x'$ and repeat.
- If no improving neighbour is found, then either restart at new random $x = x_0$ or relax the neighbourhood condition [algorithm-dependent].

Power of stochasticity?

Good experiences for 3-SAT in the satisfiable region $\alpha < \alpha_c$: e.g. GSAT (Selman et al. 1992), WalkSAT (Selman et al. 1996).

GSAT

Selman et al. 1992 $\ldots$ 1996.

Denote by $E = E_F(s)$ the number of unsatisfied clauses in formula F under truth assignment s.

```
GSAT(F):
  s = initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
    - find a variable x whose flipping causes
      largest decrease in E (if no decrease is
      possible, then smallest increase);
    - flip x.
```

NoisyGSAT

GSAT augmented by a fraction p of random walk moves.

```
NoisyGSAT(F,p):
  s = initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
    - with probability p, pick a variable x
      uniformly at random and flip it;
    - with probability (1-p), do basic GSAT move:
      - find a variable x whose flipping causes
        largest decrease in E (if no decrease is
        possible, then smallest increase);
      - flip x.
```

WalkSAT

NoisyGSAT *focused* on the unsatisfied clauses.

```
WalkSAT(F,p):
  s = initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
    - pick a random unsatisfied clause C in F;
    - if some variables in C can be flipped without
      breaking any presently satisfied clauses,
      then pick one such variable x at random; else:
    - with probability p, pick a variable x
      in C at random;
    - with probability (1-p), pick an x in C
      that breaks a minimal number of presently
      satisfied clauses;
    - flip x.
```

WalkSAT vs. NoisyGSAT

The focusing seems to be important: in the (unsystematic) experiments in Selman et al. (1996), WalkSAT outperforms NoisyGSAT by several orders of magnitude.

Recent results and conjectures

- Barthel, Hartmann & Weigt (2003), Semerjian & Monasson (2003): WalkSAT with $p = 1$ has a “dynamical phase transition” at $\alpha_{\text{dyn}} \approx 2.7 - 2.8$. When $\alpha < \alpha_{\text{dyn}}$, satisfying assignments are found in linear time per variable (i.e. in a total of cN “flips”), when $\alpha > \alpha_{\text{dyn}}$ exponential time is required.
- Explanation: for $\alpha > \alpha_{\text{dyn}}$ the search equilibrates at a nonzero energy level, and can only escape to a ground state through a large enough random fluctuation.
- Conjecture(?): all local search algorithms will have difficulties beyond the clustering transition at $\alpha \approx 3.92 - 3.93$ (Mézard, Monasson, Weigt et al.)
- Conjecture(?): WalkSAT seems to work in linear time up to the 1RSB stability transition at $\alpha \approx 4.15$ (Aurell et al. 2004), but maybe not beyond that (Aurell, Montanari et al.)

WalkSAT experiments (3-SAT)

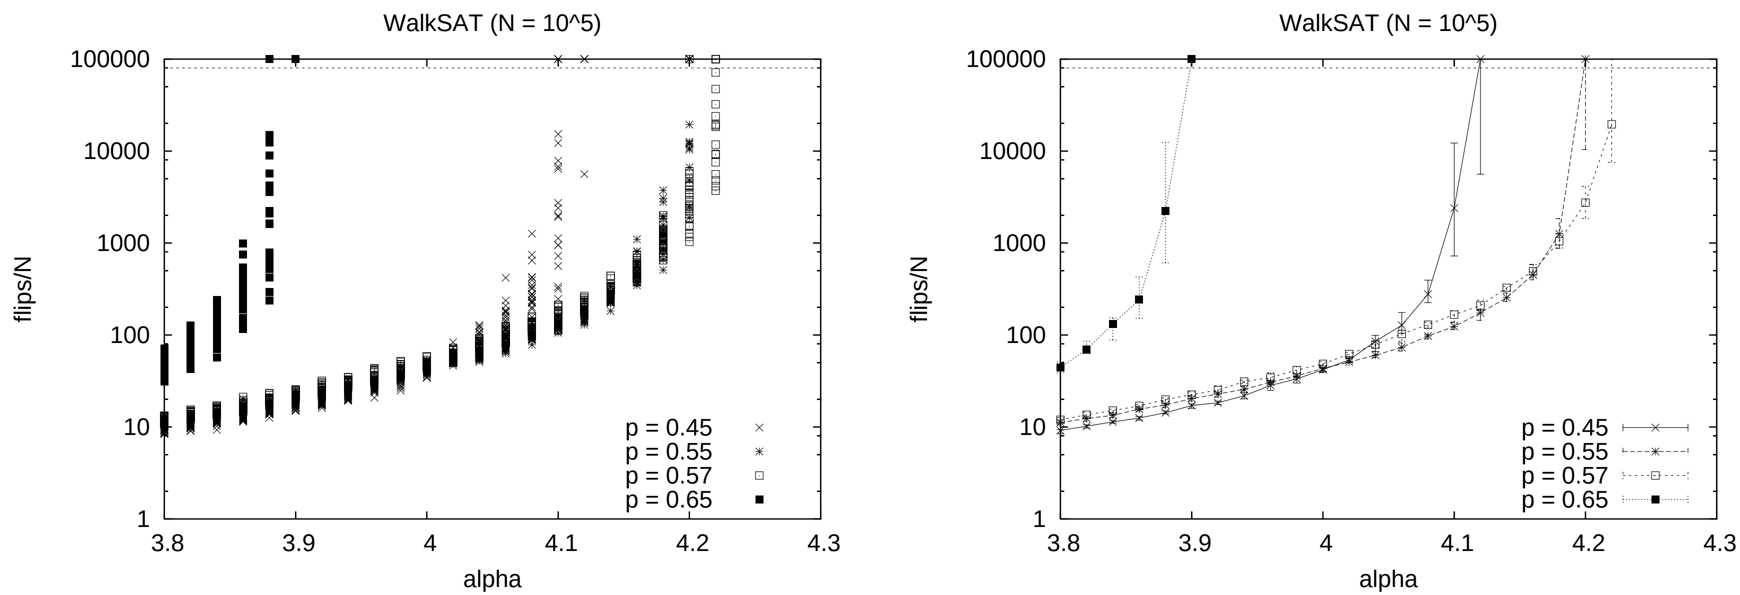

Normalised solution times for WalkSAT, $\alpha = 3.8 \ldots 4.3$.
Left: complete data; right: medians and quartiles.

WalkSAT linear scaling

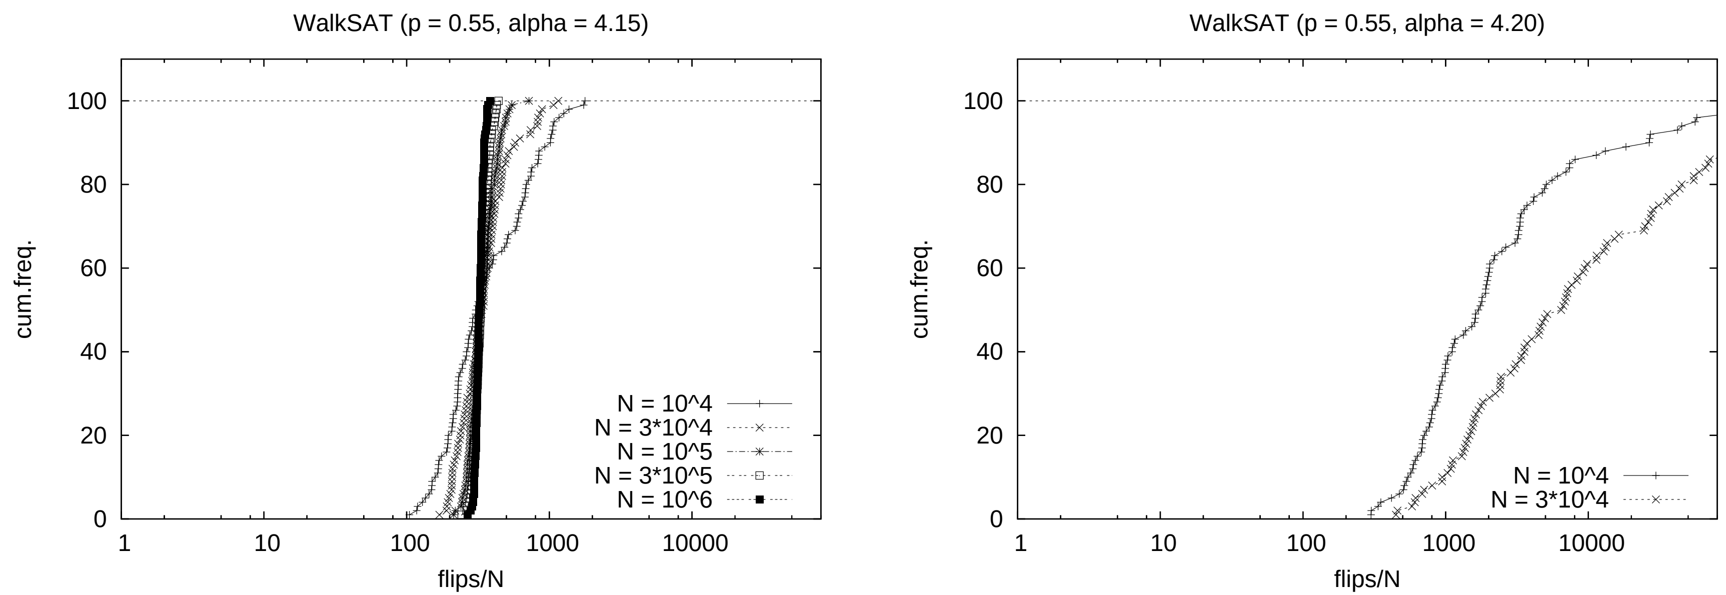

Cumulative solution time distributions for WalkSAT with $p = 0.55$.

WalkSAT optimal noise level?

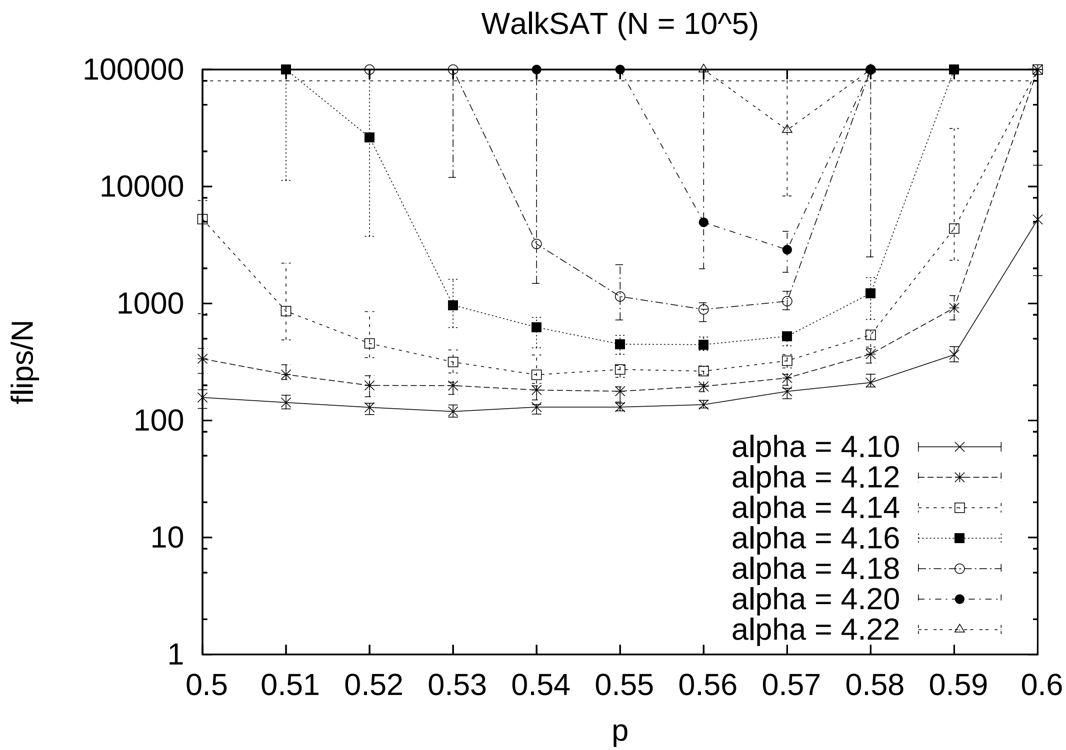

Normalised solution times for WalkSAT with $p = 0.50 \ldots 0.60$, $\alpha = 4.10 \ldots 4.22$.

WalkSAT sensitivity to noise

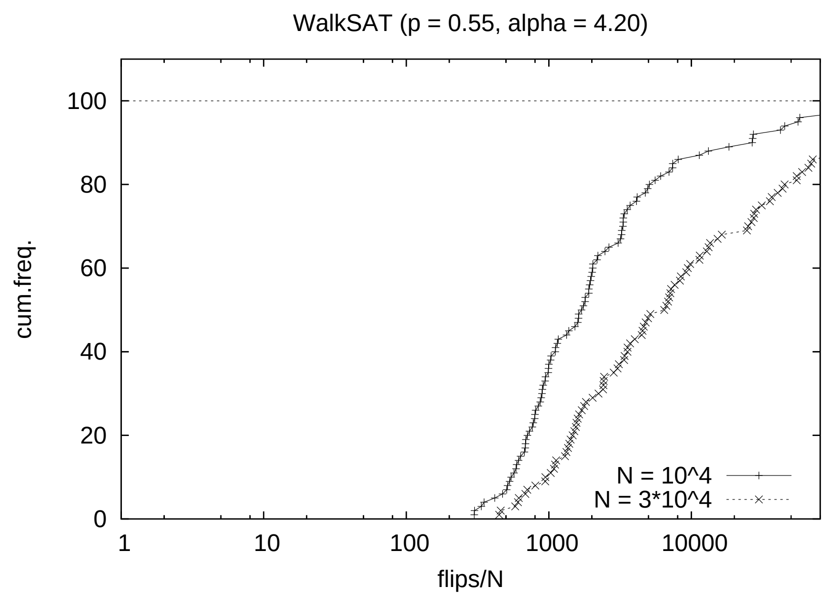

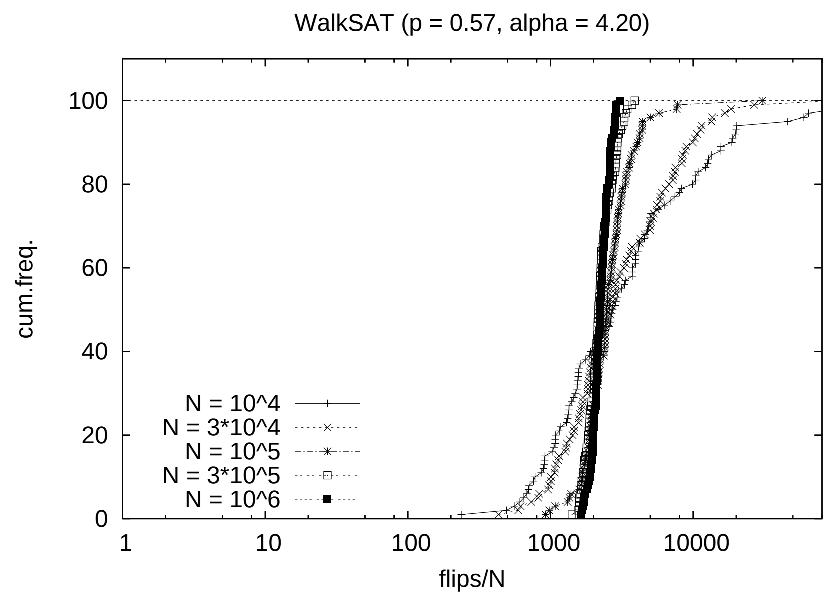

Cumulative solution time distributions for WalkSAT at $\alpha = 4.20$ with $p = 0.55$ and $p = 0.57$.

WalkSAT small-scale effects

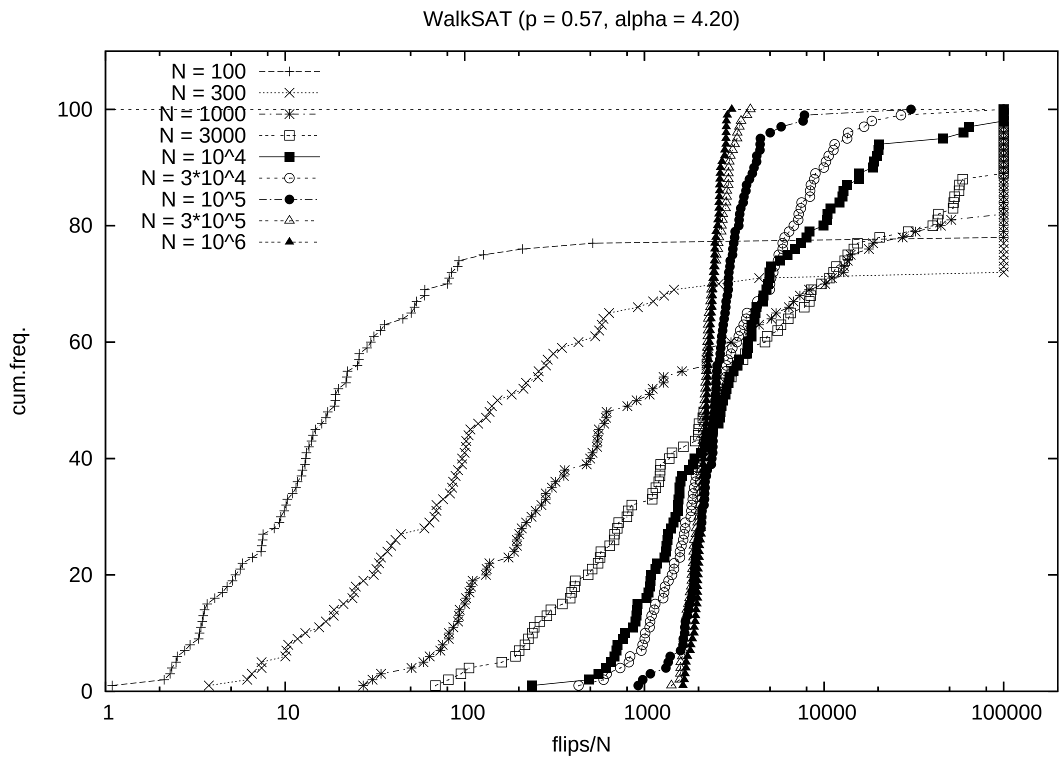

Cumulative solution time distributions for WalkSAT at $\alpha = 4.20$ with $p = 0.57$.

Record-to-Record Travel (RRT)

Very simple stochastic local optimisation algorithm introduced by Dueck (1993). Dueck claimed good results on solving 442-city and 532-city TSP's; after that little used.

```
RRT(E,d):
  s = initial feasible solution;
  s* = s; E* = E(s);
  while moves < max_moves do
    if s is a global min. of E then output s & halt,
    else:
      pick a random neighbour s' of s;
      if E(s') <= E* + d then let s = s';
      if E(s') < E* then:
        s* = s'; E* = E(s').
```

RRT in action (d = 2)

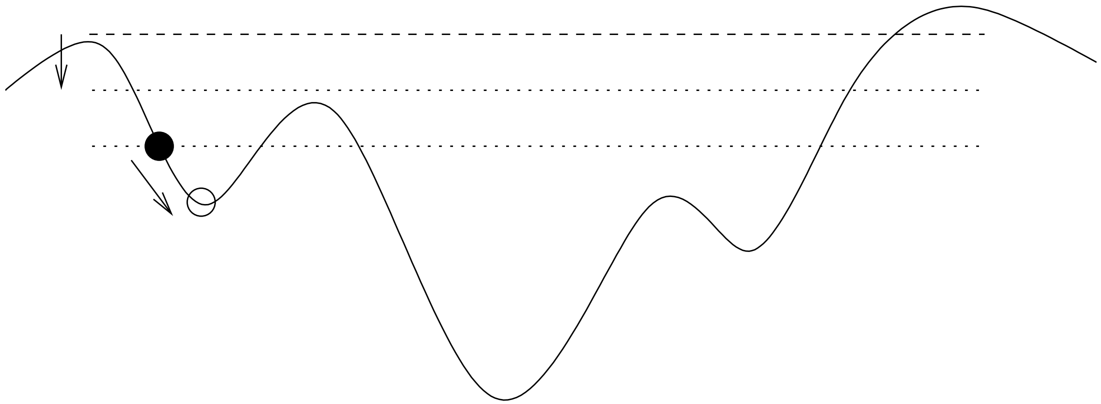

RRT in action (d = 2)

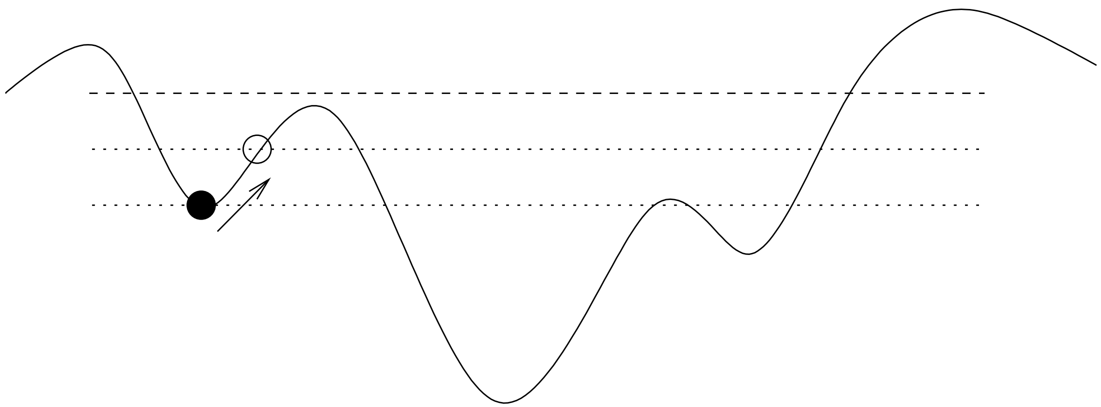

RRT in action (d = 2)

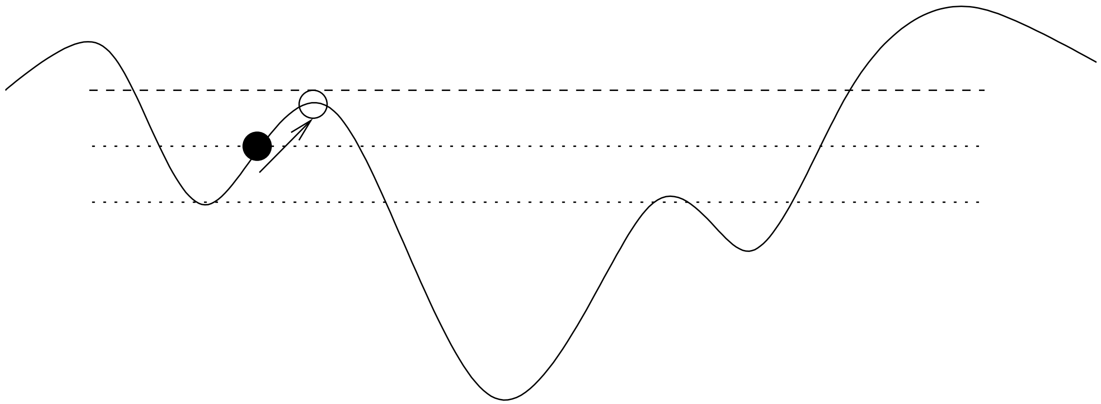

RRT in action (d = 2)

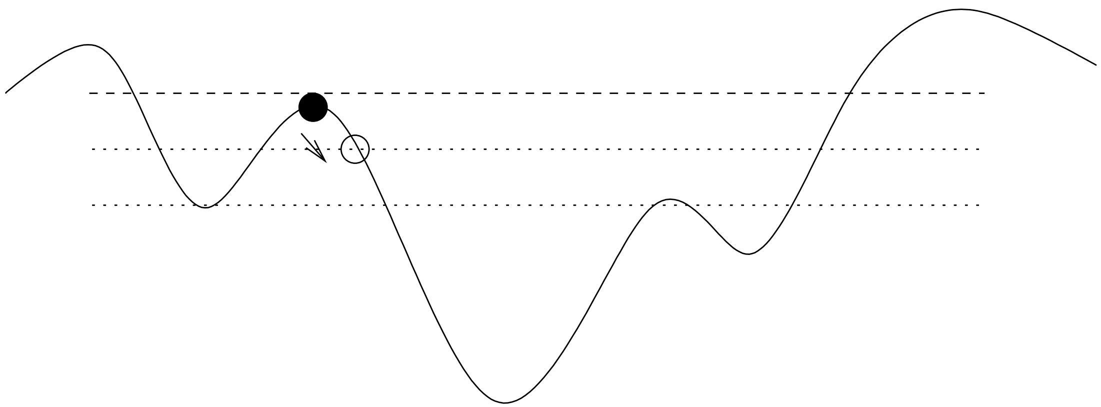

RRT in action (d = 2)

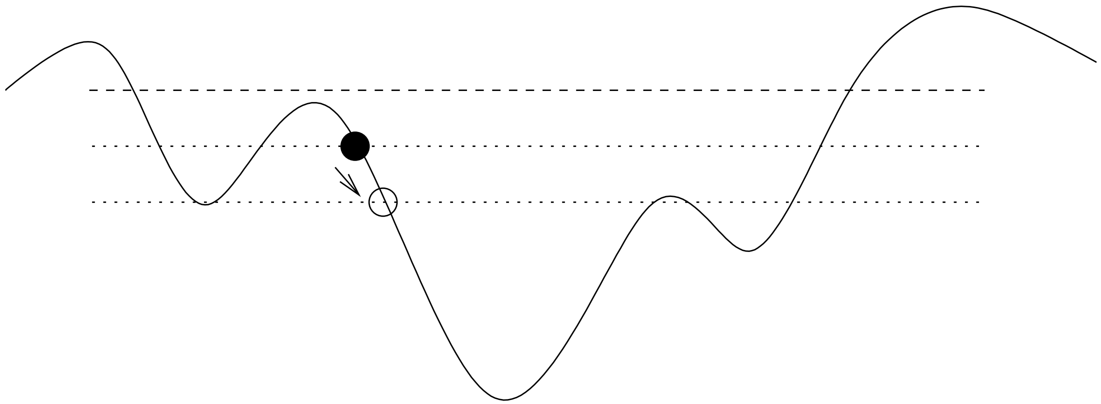

RRT in action (d = 2)

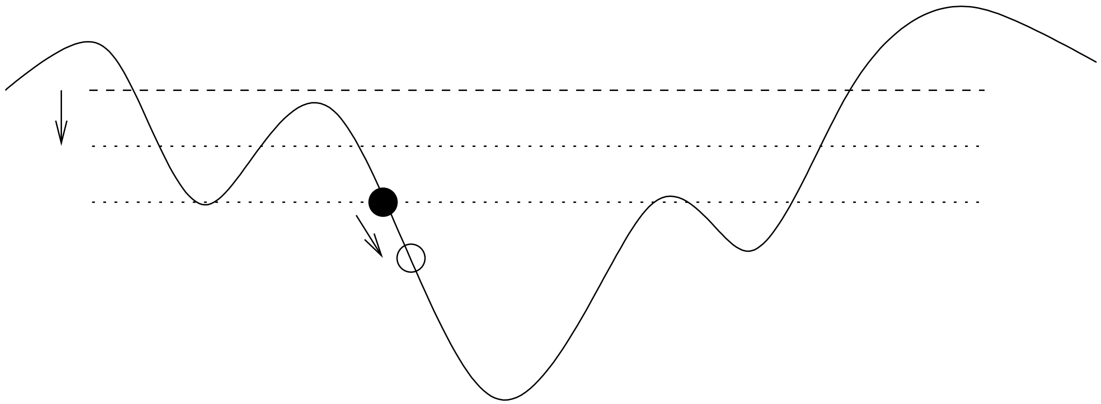

RRT in action (d = 2)

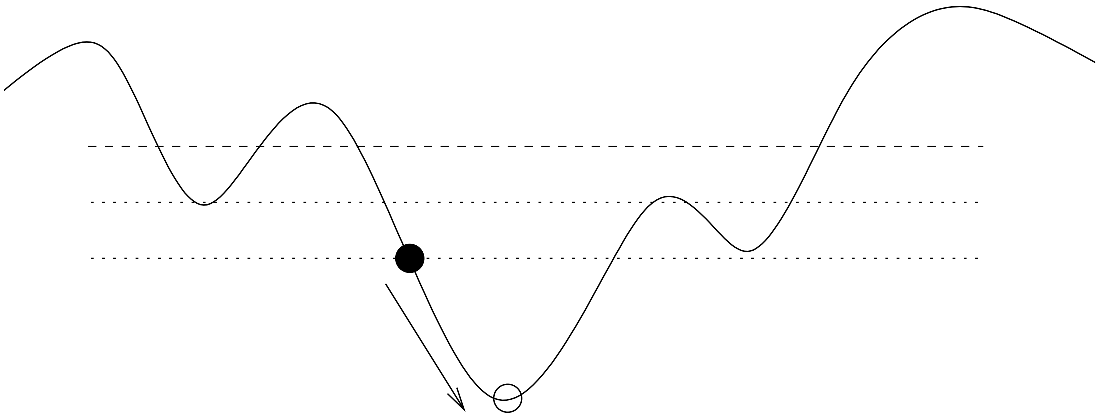

RRT in action (d = 2)

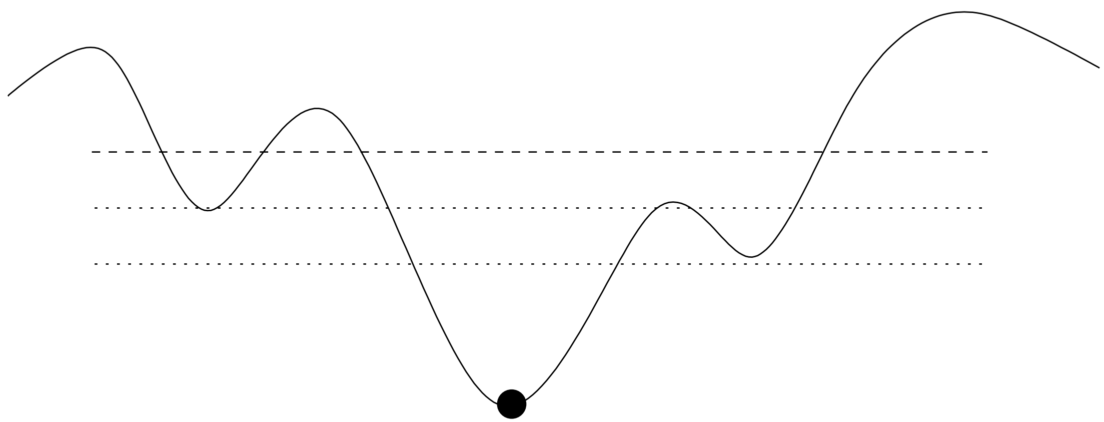

Focused RRT

In applying RRT to SAT, $E(s) =$ number of clauses unsatisfied by truth assignment s. Single-variable flip neighbourhoods.

Focusing: flipped variables chosen from unsatisfied clauses. (Precisely: one unsatisfied clause is chosen at random, and from there a variable at random.)

$\Rightarrow$ FRRT = focused RRT.

FRRT experiments (3-SAT)

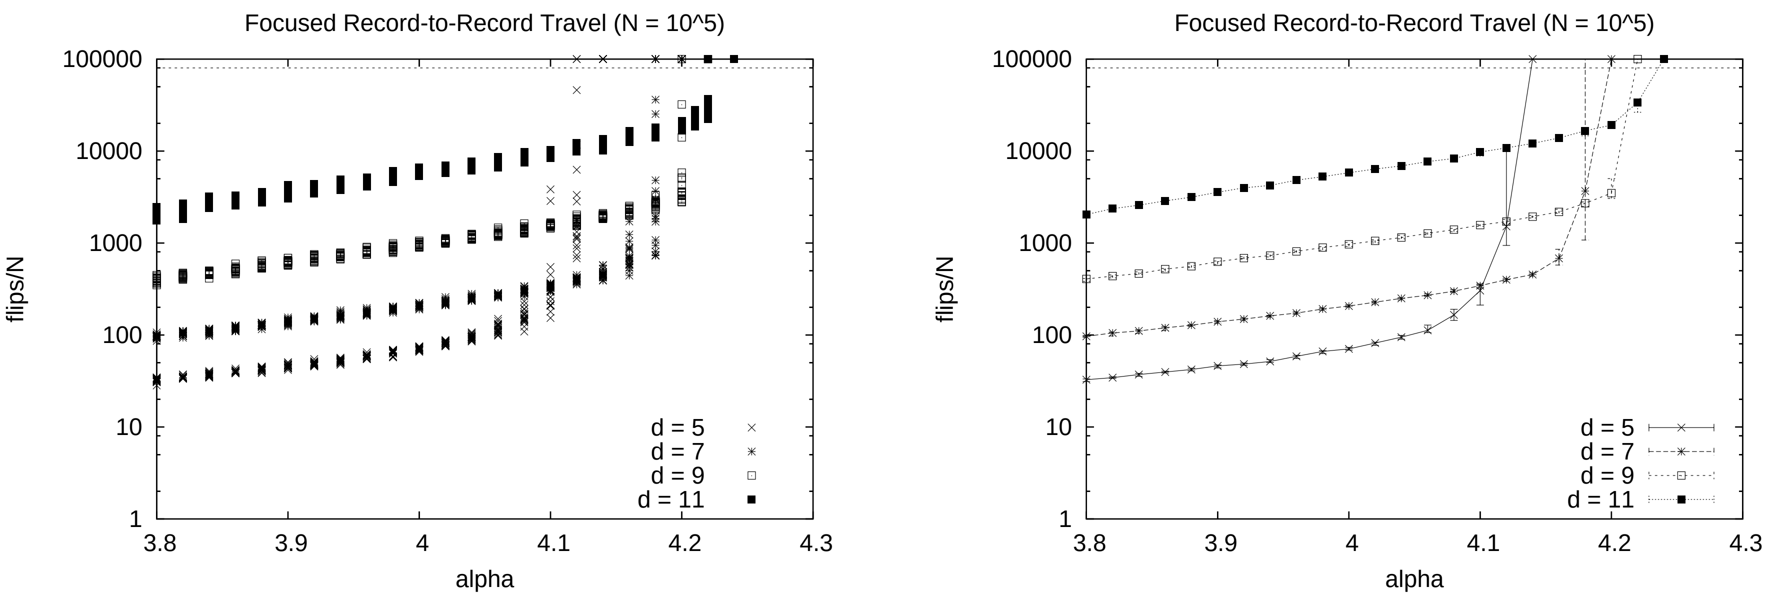

Normalised solution times for FRRT, $\alpha = 3.8 \ldots 4.3$.
Left: complete data; right: medians and quartiles.

FRRT linear scaling

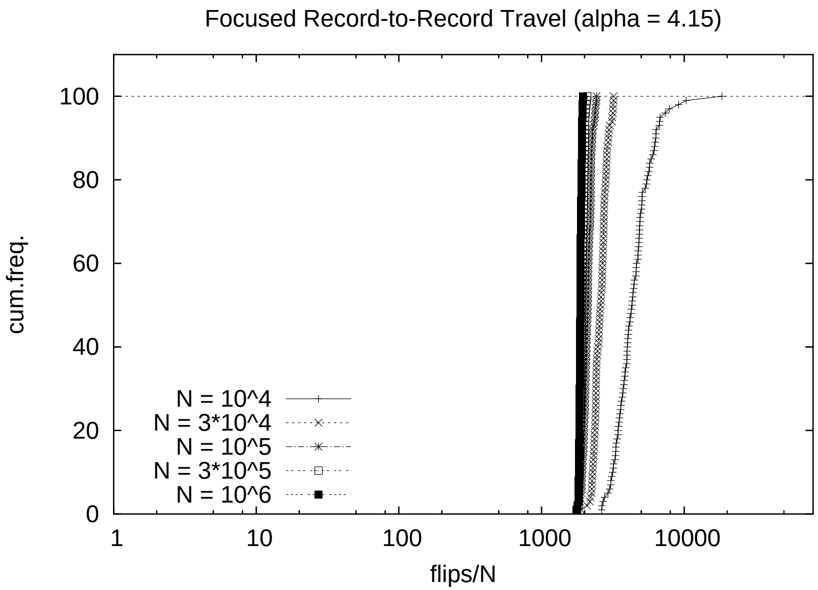

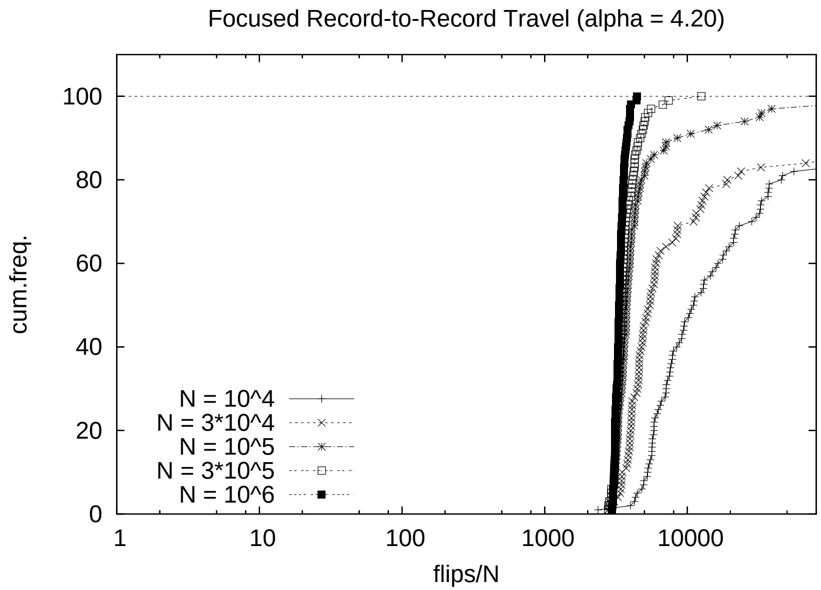

Cumulative solution time distributions for FRRT with $d = 9$.

FRRT linear scaling (cont'd)

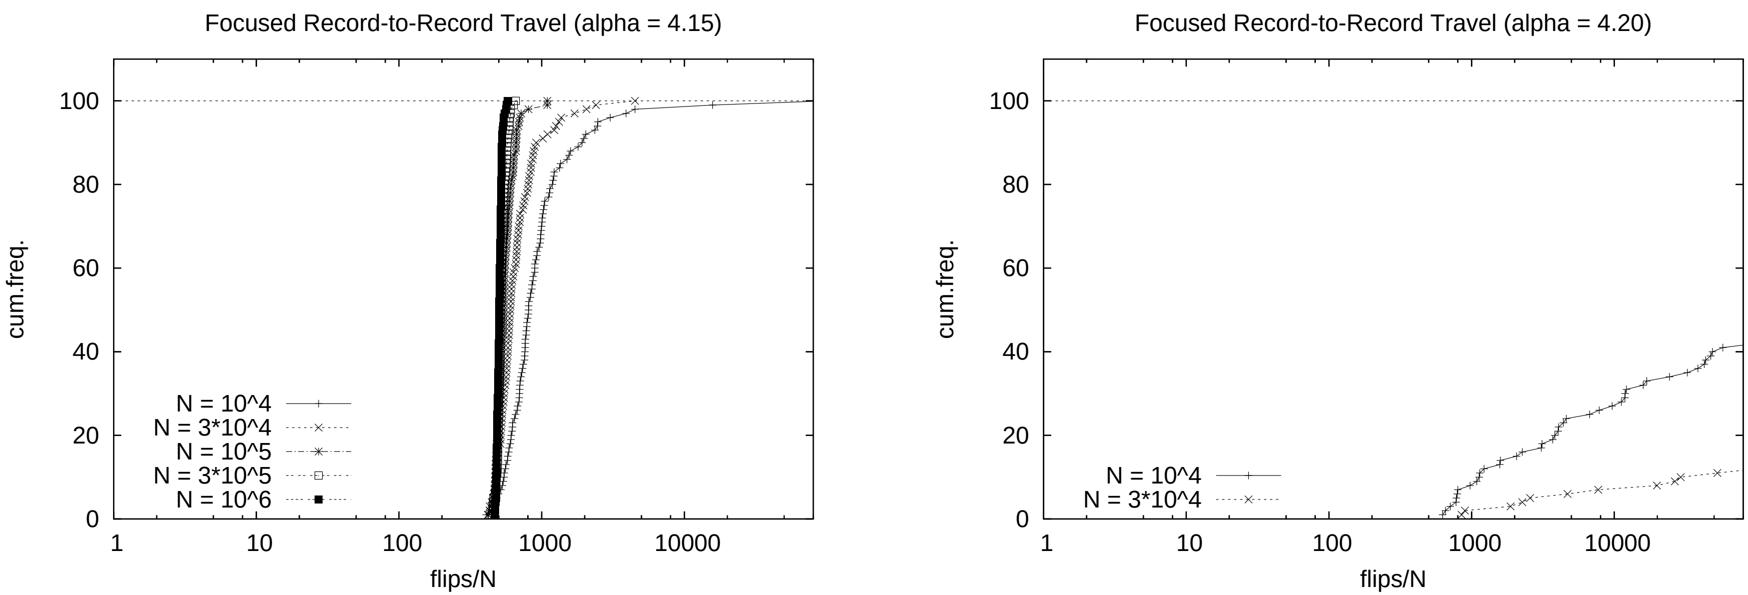

Cumulative solution time distributions for FRRT with $d = 7$.

Focused Metropolis Search

Arguably the most natural focused local search algorithm. Variable flip acceptance probabilities determined by a parameter η, $0 \leq \eta \leq 1$.

```
FMS(F,eta):
  s = initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
      pick a random unsatisfied clause C in F;
      pick a variable x in C at random;
      let x' = flip(x), s' = s[x'/x];
      if E(s') <= E(s) then flip x, else:
        flip x with prob. eta^(E(s')-E(s)).
```

FMS experiments (3-SAT)

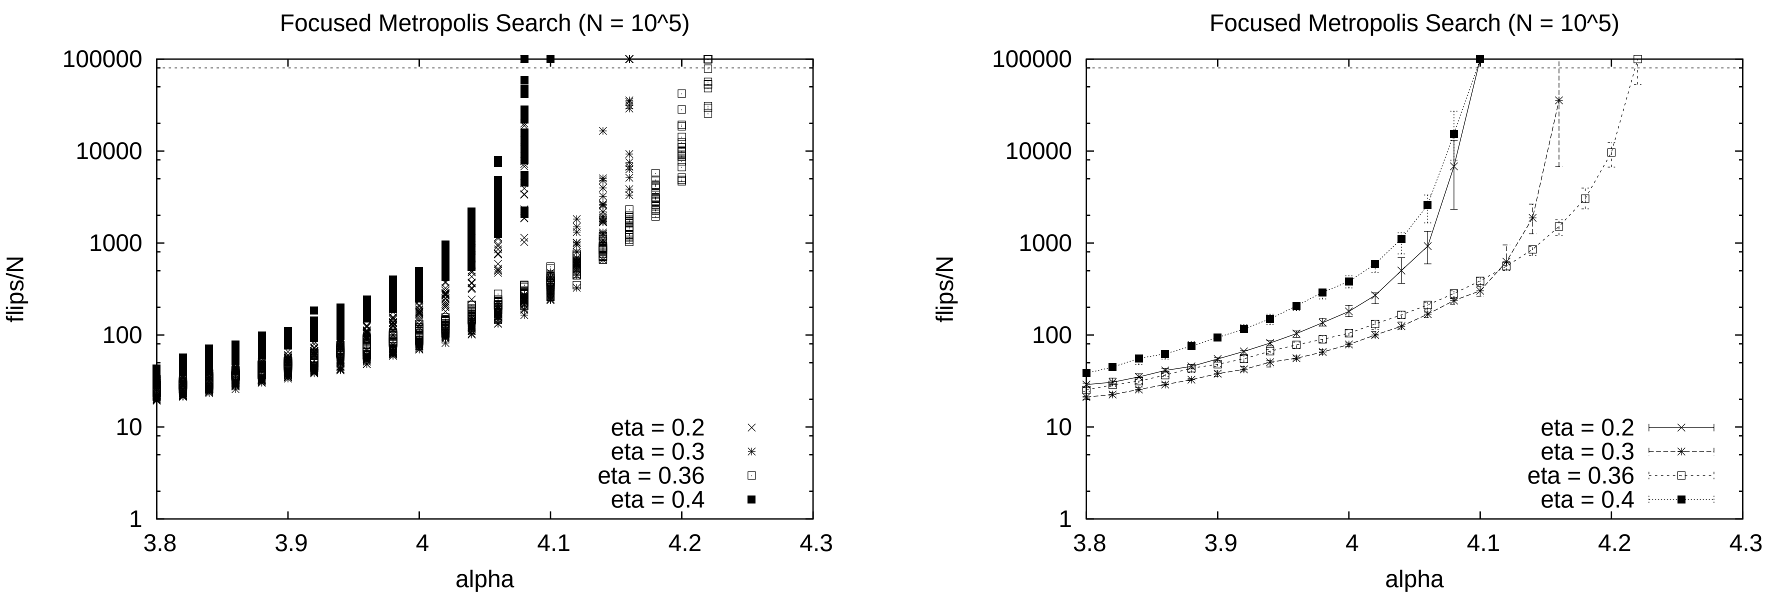

Normalised solution times for FMS, $\alpha = 3.8 \ldots 4.3$.
Left: complete data; right: medians and quartiles.

FMS linear scaling

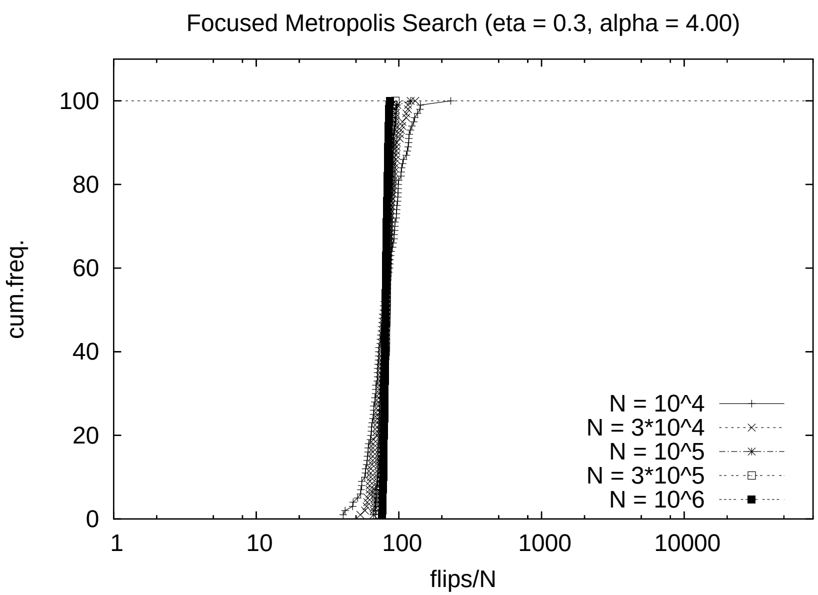

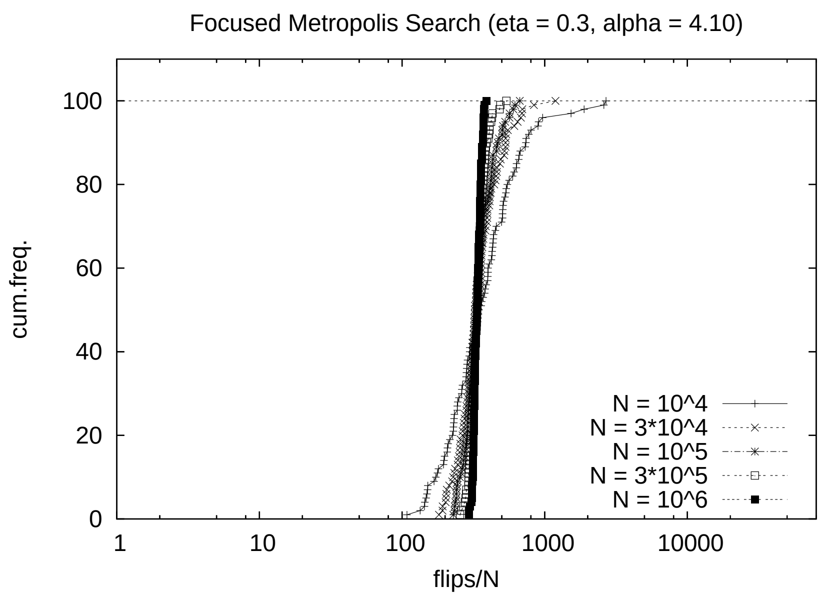

Cumulative solution time distributions for FMS with $\eta = 0.3$.

FMS optimal acceptance ratio?

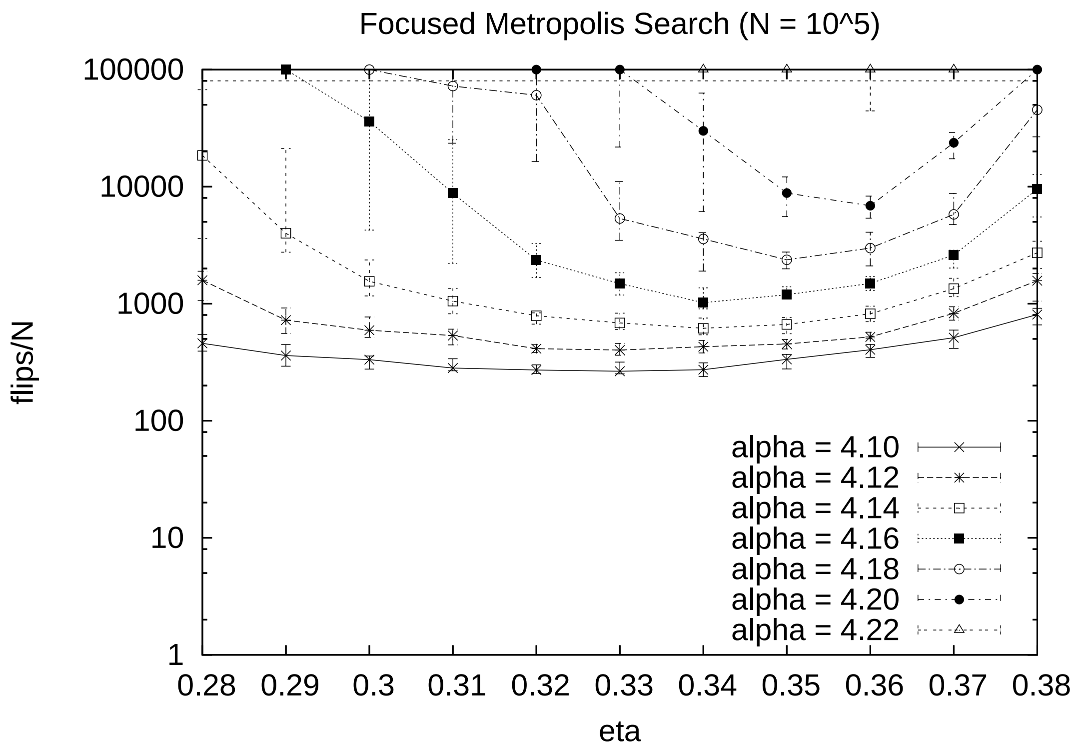

Normalised solution times for FMS with $\eta = 0.28 \ldots 0.38$, $\alpha = 4.10 \ldots 4.22$.

FMS optimal acceptance ratio cont'd

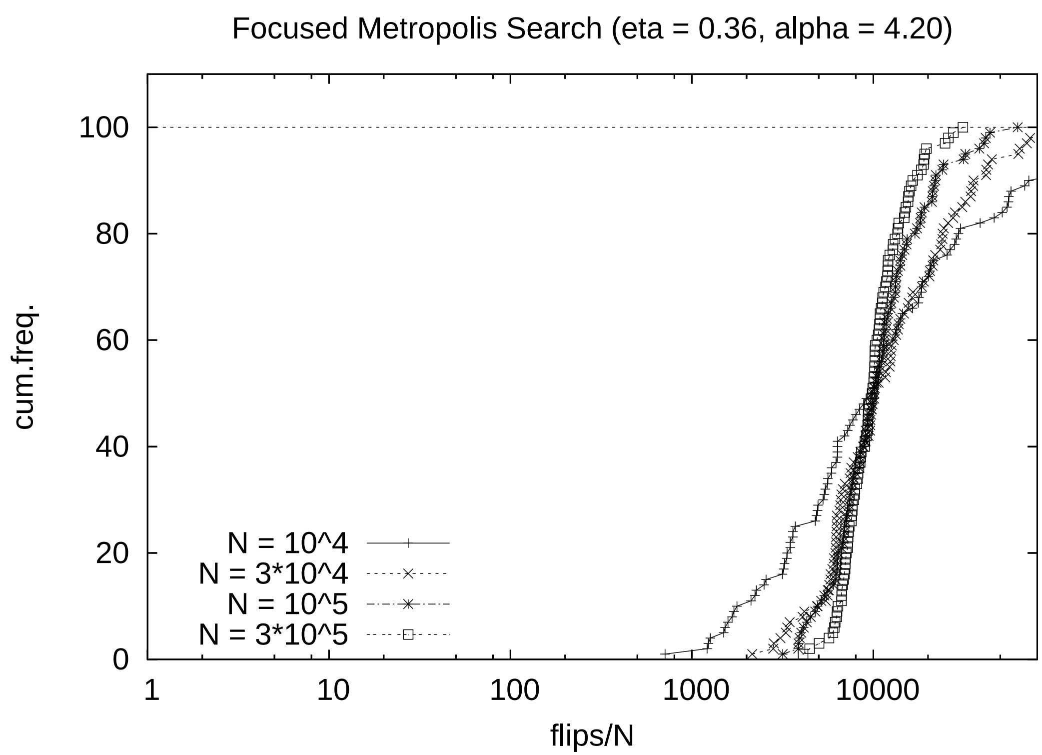

Cumulative solution time distributions for FMS with $\eta = 0.36$, $\alpha = 4.20$.

FMS optimal acceptance ratio cont'd

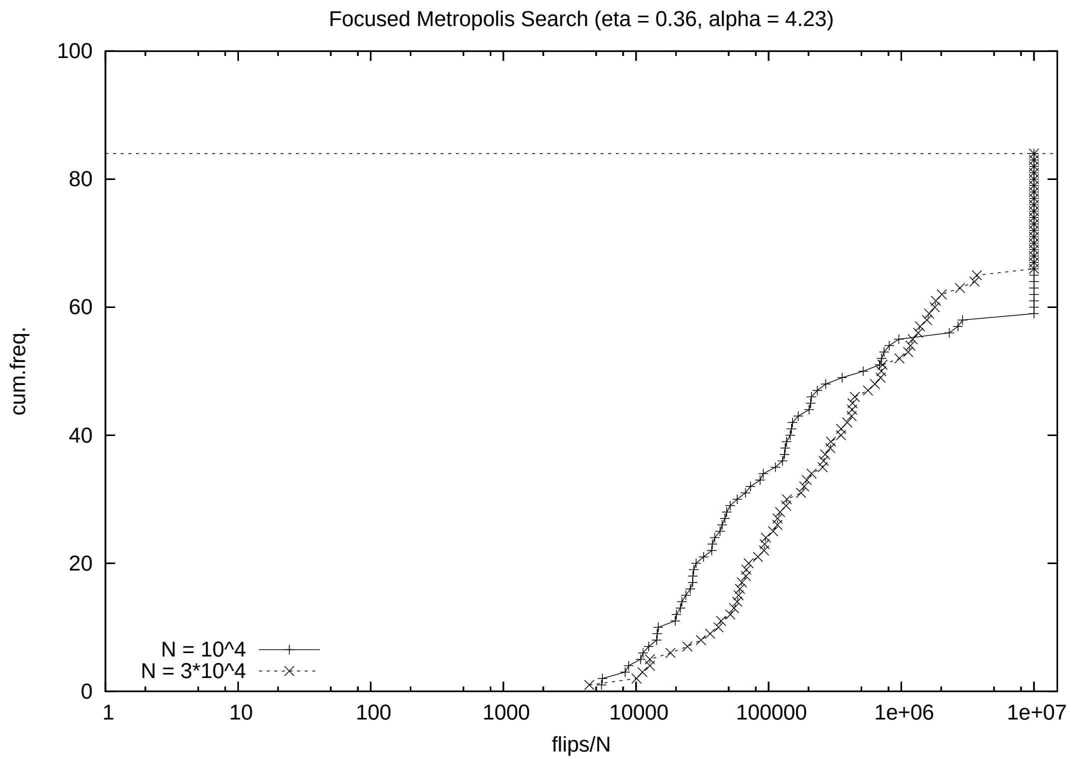

Cumulative solution time distributions for FMS with $\eta = 0.36$, $\alpha = 4.23$.

Analysis?

- Whitening
- Contact processes

Whitening

Technique introduced by Parisi, Braunstein, Zecchina et al. to determine the "frozen" variables (spins, degrees of freedom) in a given configuration. A variable is *frozen* in truth assignment s to 3-SAT formula F, unless determined *white* by the following process:

```
WHITENING(F, s):
  mark all clauses white except those that
    have exactly one true literal;
  loop:
    mark all variables white except those that
      appear as the unique satisfying literals
      in non-white clauses;
    halt, if all variables are white (full whitening
    halt, if no new variables became white (core fou
    mark those clauses white that contain
      at least one white variable.
```

Whitening (cont'd)

The "whiteness" of solutions seems to have many connections to the behaviour of local search algorithms. Consider e.g. the following plots of runtimes of an FRRT variant vs. the "whiteness depth" of the solutions found by it:

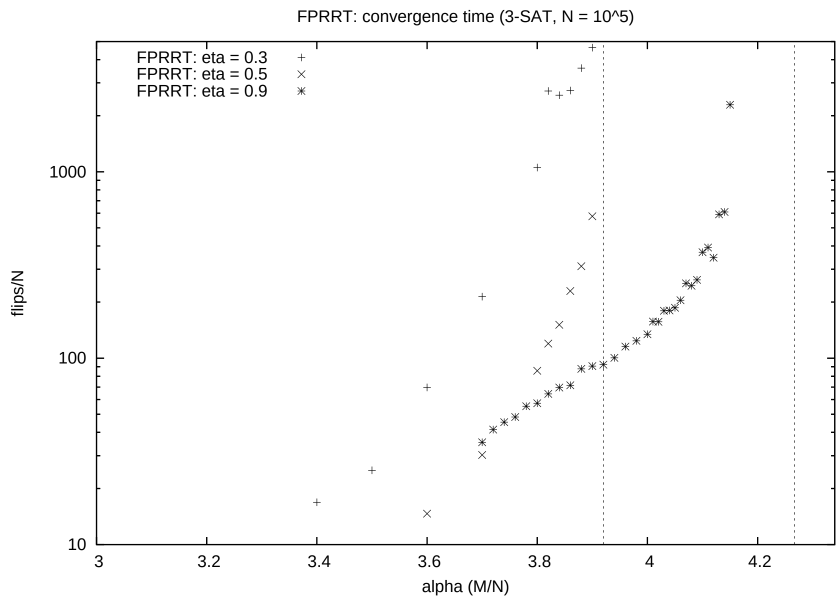

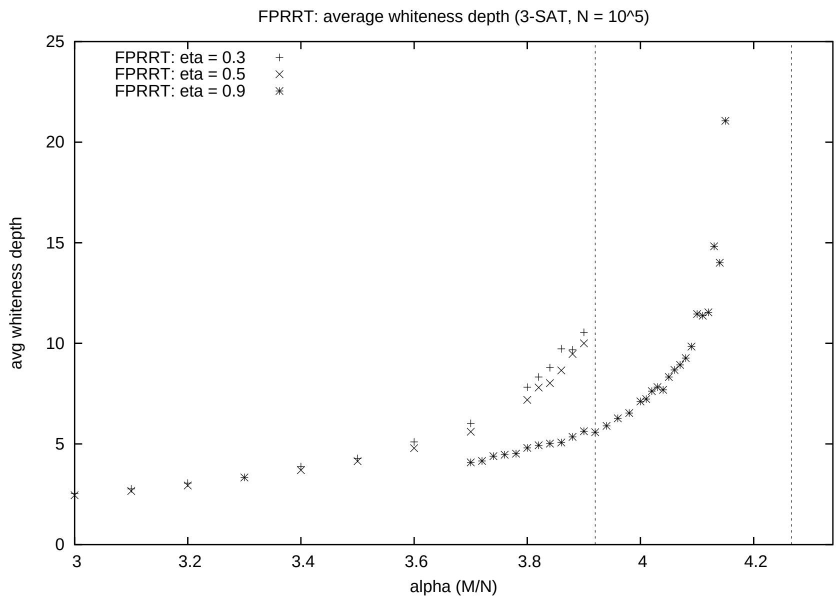

Focused search as a contact process

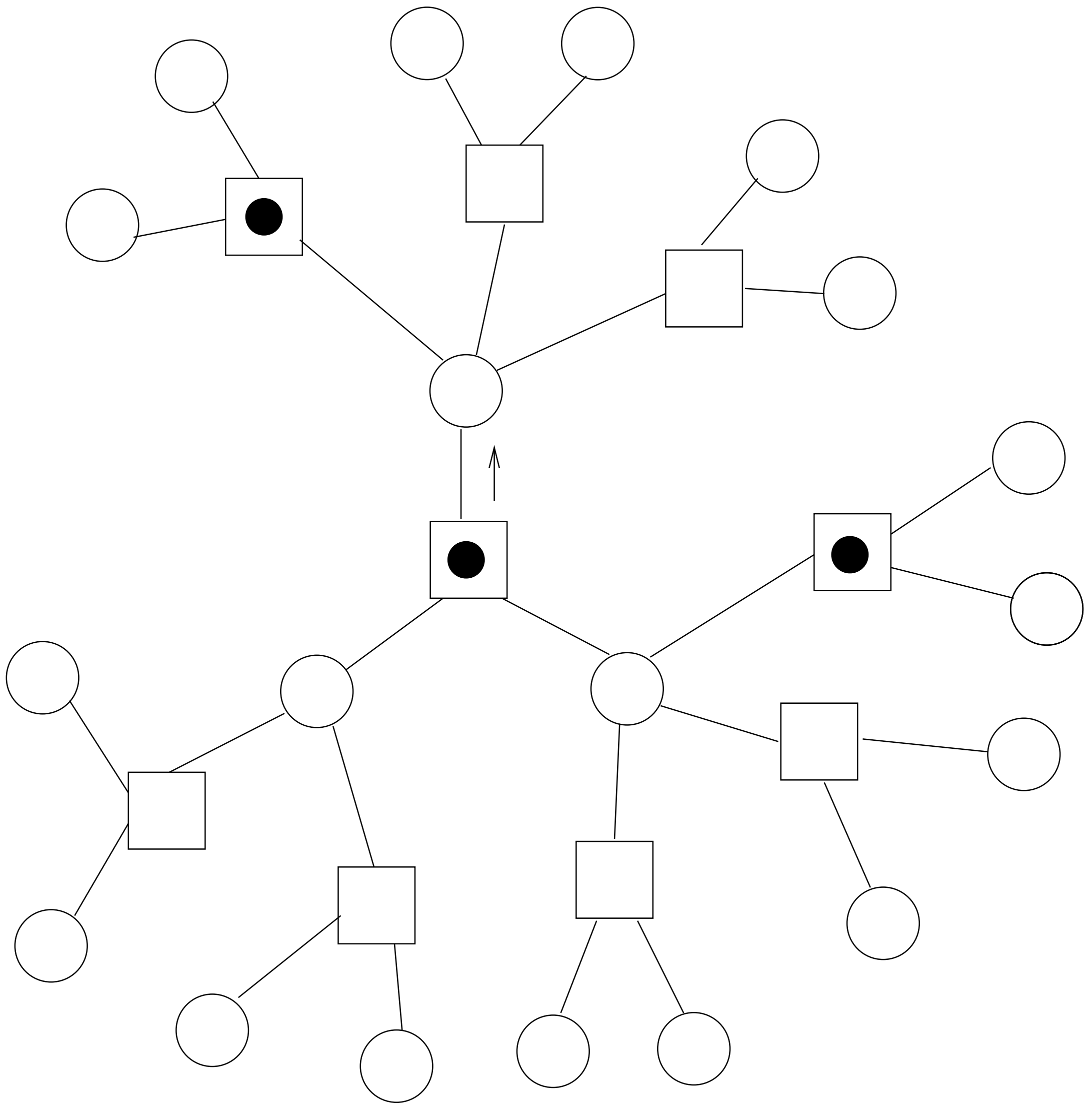

Focused search as a contact process

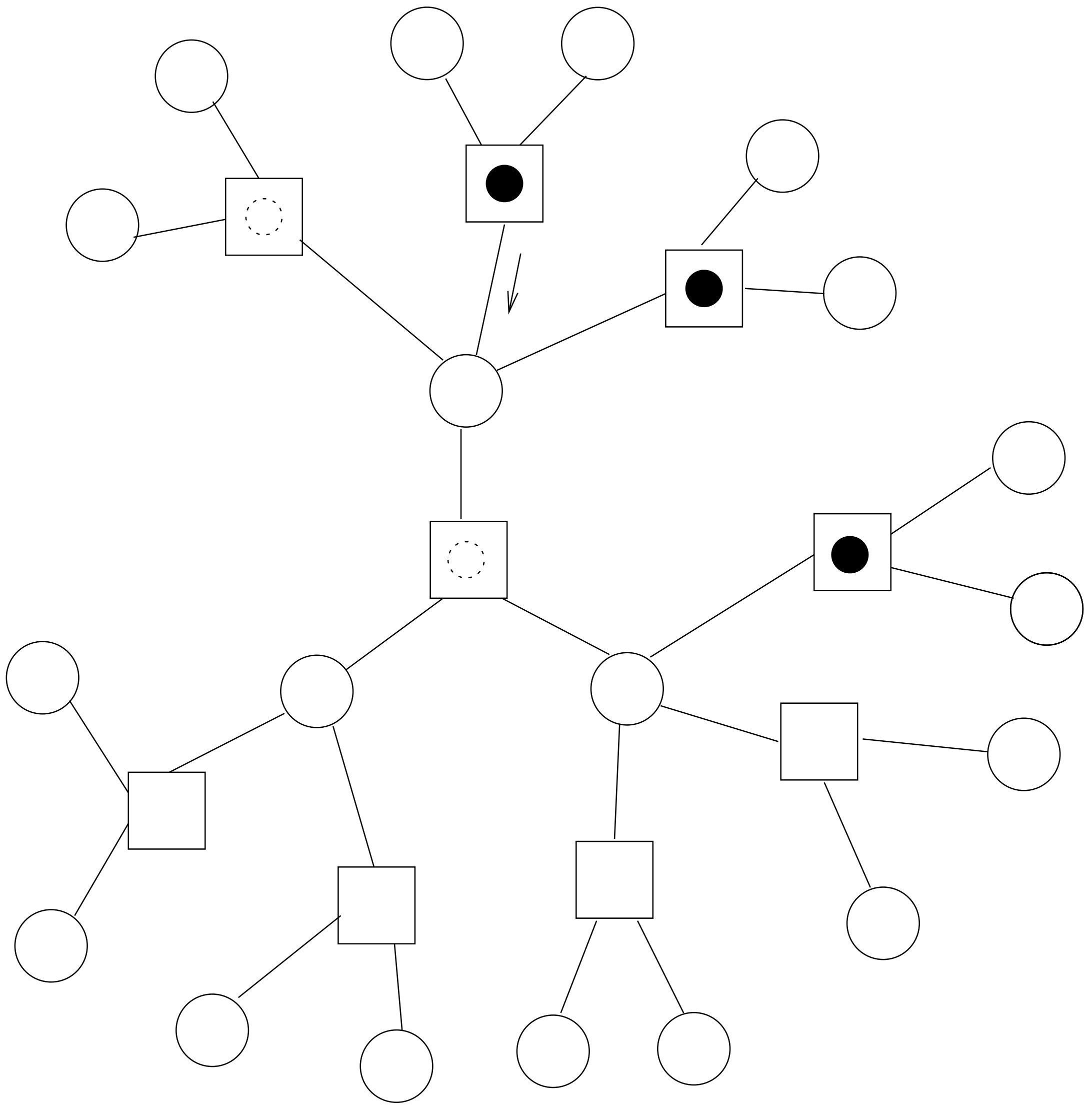

Focused search as a contact process

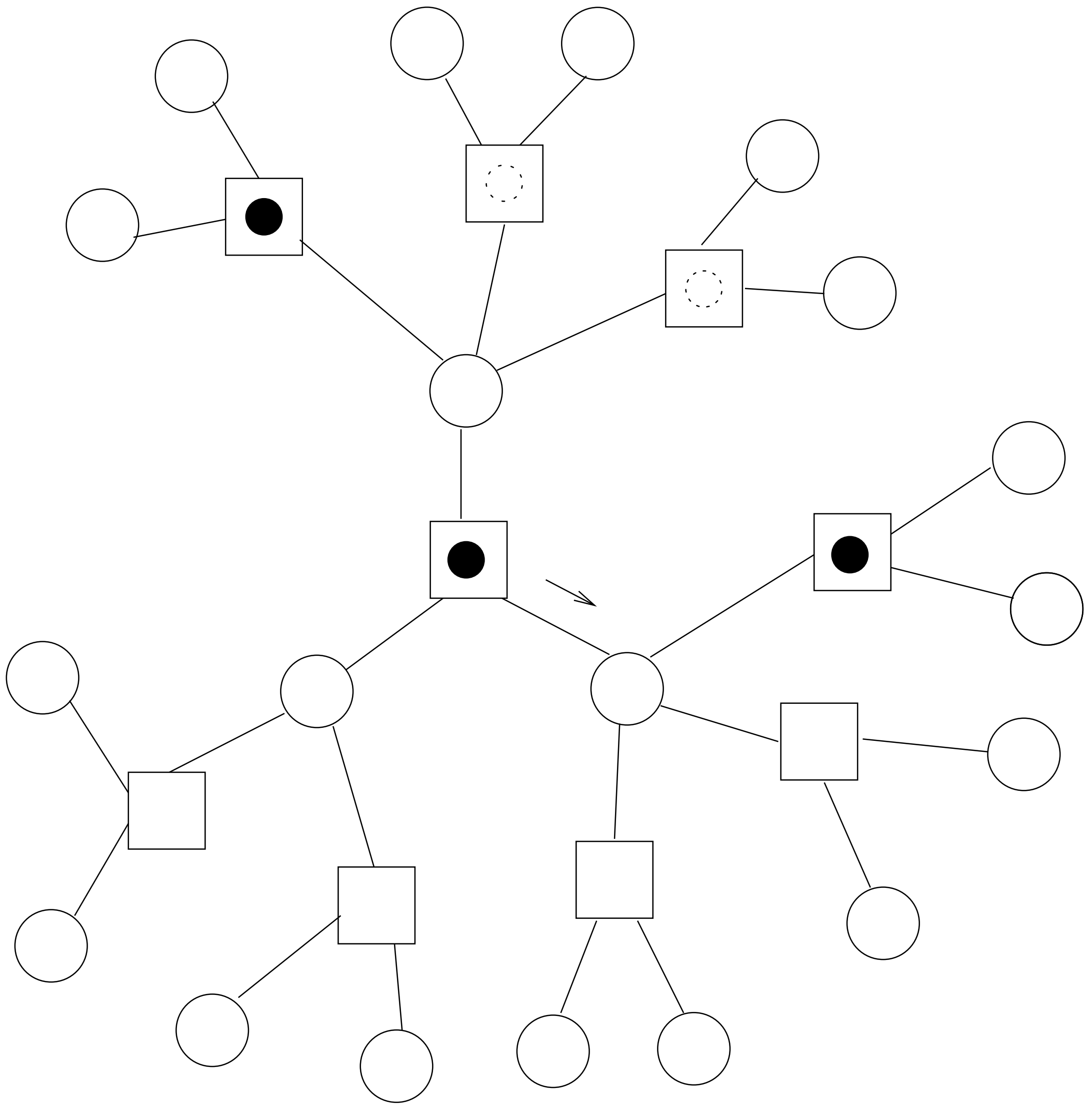

Focused search as a contact process

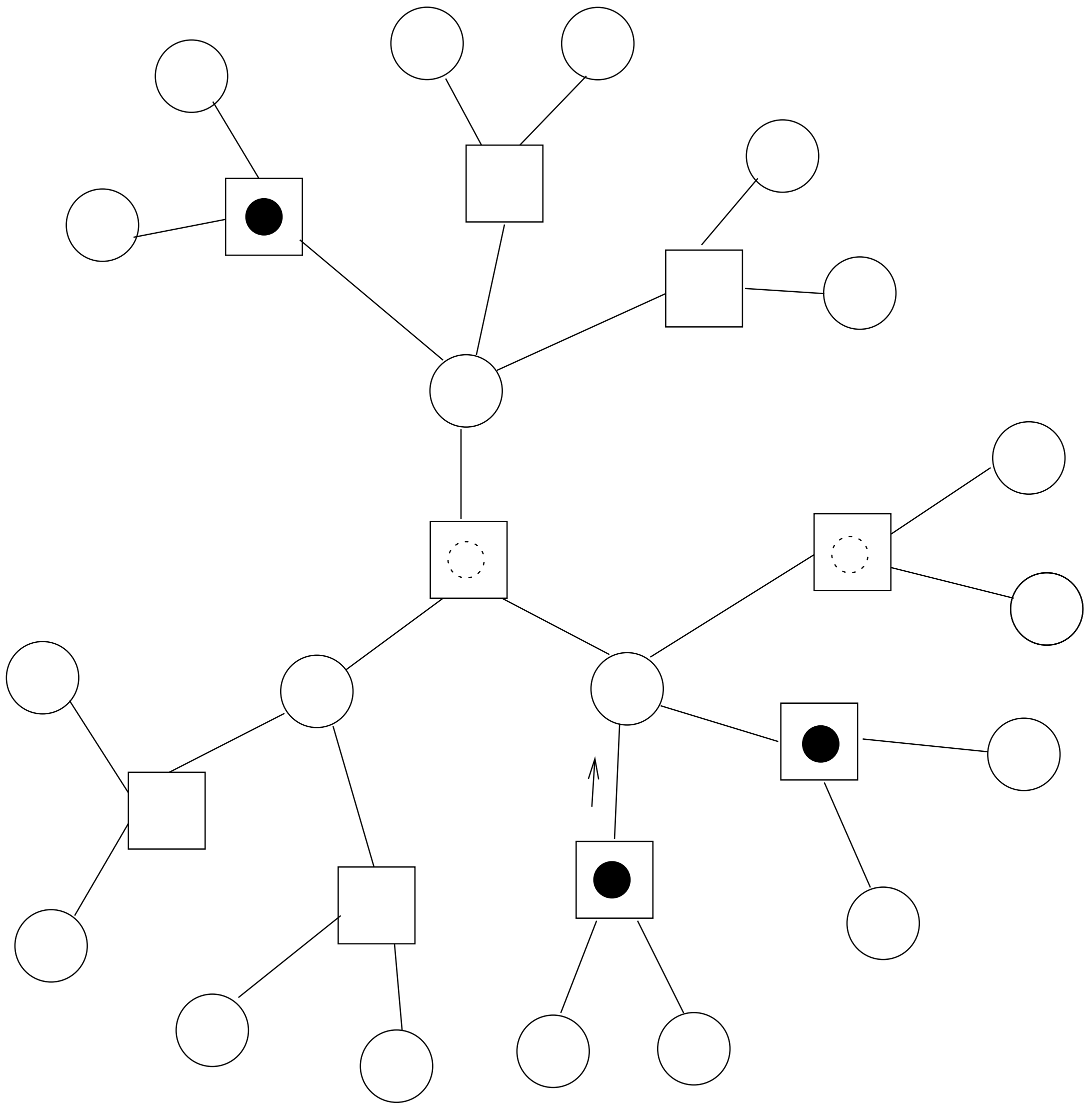

Focused search as a contact process

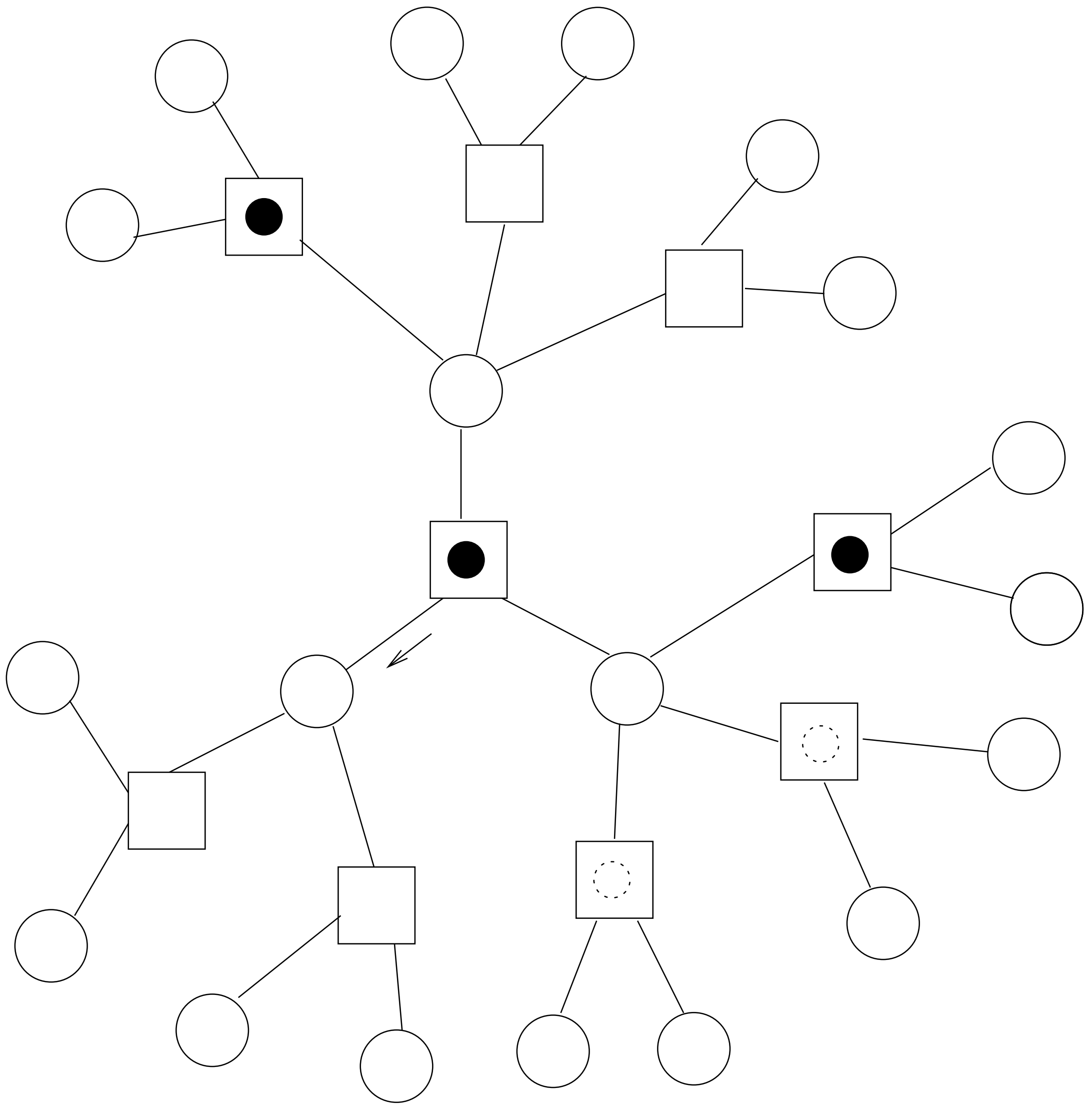

Focused search as a contact process

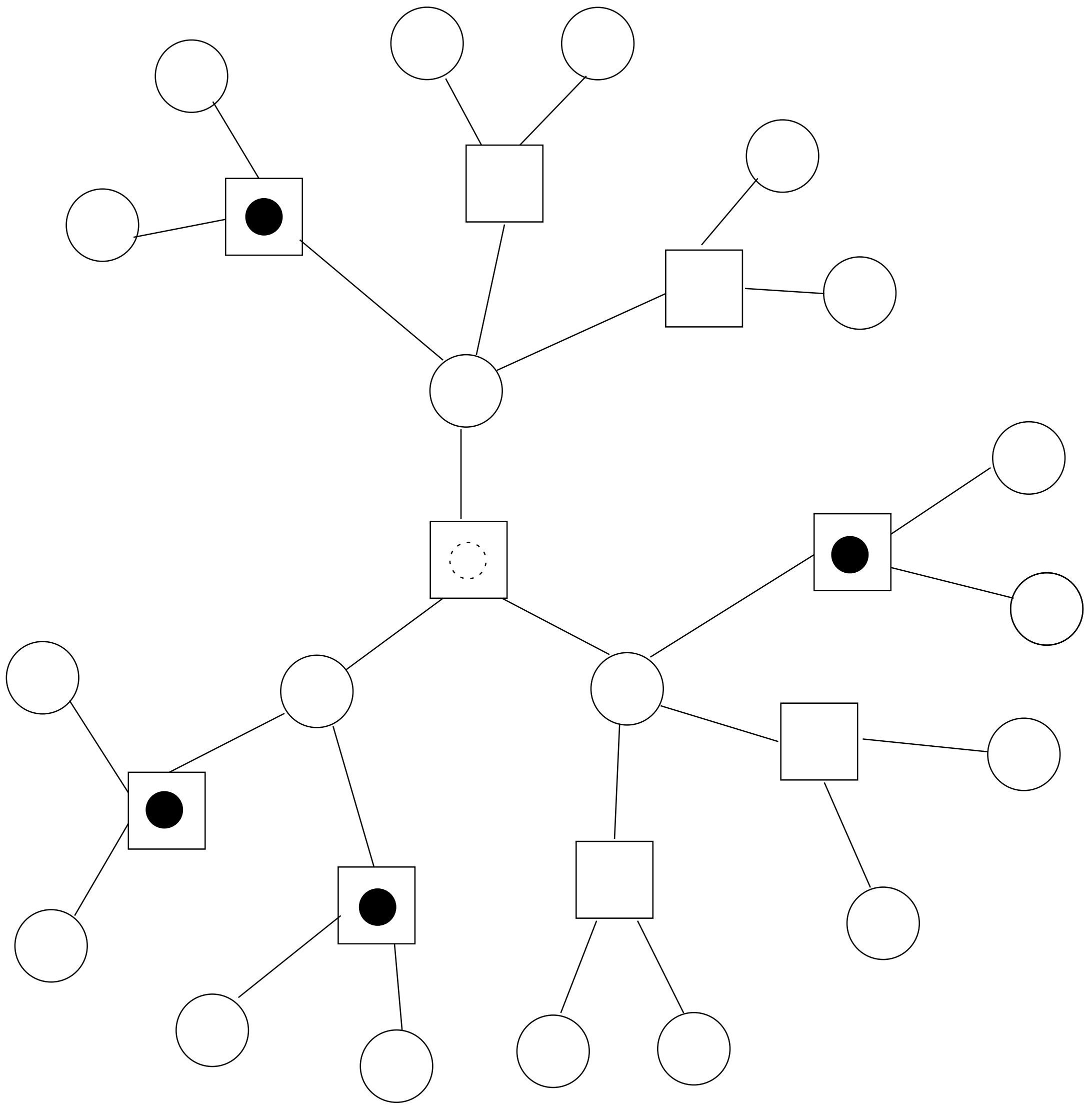

References

[1] S. Seitz, P. Orponen: An efficient local search method for random 3-satisfiability. *Proc. LICS'03 Workshop on Typical Case Complexity and Phase Transitions (Ottawa, Canada, June 2003).* Elsevier Electronic Notes in Discr. Math. Vol. 16.

[2] S. Seitz, M. Alava, P. Orponen: Threshold behaviour of WalkSAT and focused Metropolis search on random 3-satisfiability. *Proc. 8th Intl. Conf. on Theory and Applications of Satisfiability Testing (St. Andrews, Scotland, June 2005).* Springer LNCS 3503.

[3] S. Seitz, M. Alava, P. Orponen: Focused local search algorithms for random 3-satisfiability. *J. Stat. Mech.* 2005, P06006.