Beyond NP: the QSAT phase transition

Ian P. Gent and **Toby Walsh** *
Department of Computer Science
University of Strathclyde
Glasgow G1 1XL United Kingdom
ipg@cs.strath.ac.uk, tw@cs.strath.ac.uk

Abstract

We show that phase transition behavior similar to that observed in NP-complete problems like random 3-SAT occurs further up the polynomial hierarchy in problems like random 2-QSAT. The differences between QSAT and SAT in phase transition behavior that Cadoli et al report are largely due to the presence of trivially unsatisfiable problems. Once they are removed, we see behavior more familiar from SAT and other NP-complete domains. There are, however, some differences. Problems with short clauses show a large gap between worst case behavior and median, and the easy-hard-easy pattern is restricted to higher percentiles of search cost. We compute the "constrainedness" of k-QSAT problems for any k, and use this to predict the location of phase transitions. We conjecture that these predictions are less accurate than in NP-complete problems because of the super-exponential size of the state space, and of the weakness of first moment methods in complexity classes above NP. Finally, we predict that similar phase transition behavior will occur in other PSPACE-complete problems like planning and game playing.

Introduction

A simple generalization of propositional satisfiability (SAT) is quantified satisfiability (QSAT). This is the prototypical PSPACE-complete problem. PSPACE is the class of problems that can be solved using polynomial space. Many search problems in AI lie within this complexity class (for example, propositional reasoning in many types of non-monotonic, modal, belief, temporal, and description logics). Do we observe phase transition behavior in this complexity class similar to that seen in P and NP? Is the definition of constrainedness proposed in (Gent *et al.* 1996) again useful?

We first introduce QSAT and the random model used in (Cadoli, Giovanardi, & Schaerf 1997). We then argue that random models should avoid unary constraints like unit clauses as they are often responsible for trivially insoluble problems. We show that Cadoli et al's model suffers from this flaw and propose instead two 'flawless' models for generating random QSAT. We concern ourselves with k-QSAT, a restricted subclass of QSAT detailed below. We define the constrainedness, κ, of k-QSAT problems for all k and predict the location of the phase transition, the first time this has been done for a complexity class above NP. For 2-QSAT, we compare this prediction with empirical results. The prediction is not always as accurate as in many NP problems, and we conjecture why.

QSAT

QSAT is the problem of deciding the satisfiability of propositional formulae in which the Boolean variables are either existentially or universally quantified. For example, $\forall x \exists y \ (x \vee \neg y) \wedge (\neg x \vee y)$ evaluates to true since whatever truth value, T or F we give to x, there is a truth value for y, namely the same value as x, which satisfies the quantified formula. We can group consecutive variables sharing the same quantifier into a set bound by a single quantifier, so we assume that the quantifiers alternate, an universal following an existential and vice versa. A k-QSAT problem is a QSAT problem in which there are k alternating quantifiers applied to disjoint sets of variables, with the innermost quantifier being existential. Our example above is in 2-QSAT, while 1-QSAT is the same as SAT. Many games like generalized versions of checkers, Go, Hex, and Othello are PSPACE-complete. Indeed, we can view QSAT as a game between the existential quantifiers, which try to pick instantiations that give a satisfiable subformula, and the universal quantifiers, which try to pick instantiations that give an unsatisfiable subformula.

Whilst SAT is NP-complete, QSAT is PSPACE-complete, and k-QSAT is Σ_kP-complete (Papadimitriou 1994). Notice that the difference between QSAT and k-QSAT is that in QSAT there is no *a priori* limit on the number k of alternations. The union of the classes Σ_kP for all k defines the 'cumulative polynomial hierarchy' PH. If for some i, $\Sigma_{i+1}\text{P} = \Sigma_i\text{P}$ then the polynomial hierarchy 'collapses' at level i and PSPACE = PH = Σ_iP.

* We thank Marco Cadoli for binaries for the EVALUATE algorithm. We are members of the cross-site APES Research Group, http://www.cs.strath.ac.uk/~apes, and we thank our fellow members from both Leeds and Strathclyde Universities. We thank reviewers of AAAI for helpful comments. Toby Walsh is supported by EPSRC grant GR/L24014.

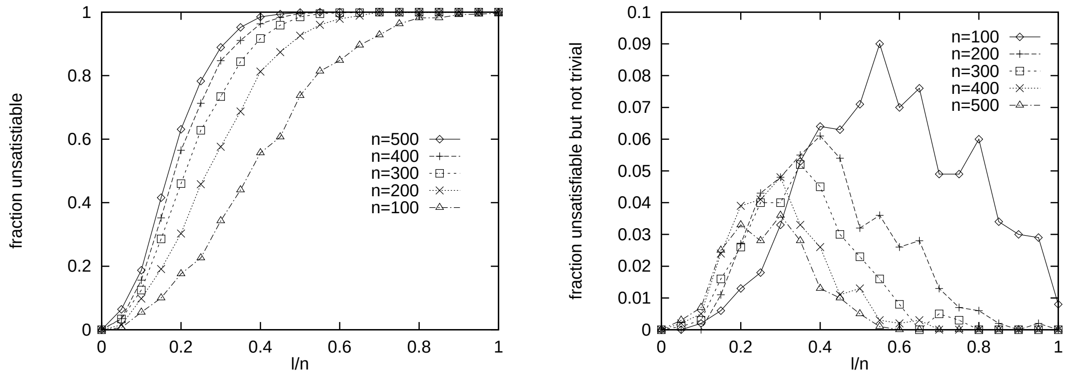

Figure 1: Random 2-QSAT problems from Cadoli *et al*'s flawed model, generated with $k = 3$ and varying n and l/n. **(left)** fraction of unsatisfiable problems; **(right)** fraction of unsatisfiable problems that are not trivial. The scale on the y-axis shows that at least 90% of all unsatisfiable problems were trivially insoluble in every case.

It is conjectured that no such collapse occurs.

In this paper, we analyse k-QSAT for all k and perform experiments on 2-QSAT. As in SAT, we can restrict the quantified formulae in a QSAT problem to conjunctive normal form (CNF). QSAT with formulae in CNF remains PSPACE-complete, and k-QSAT with formulae in 3-CNF remains Σ_kP-complete. (Cadoli, Giovanardi, & Schaerf 1998) propose an algorithm for solving QSAT problems in CNF which we use throughout this paper.

(Cadoli, Giovanardi, & Schaerf 1997) generalize the well known fixed clause model from SAT to QSAT. In this model, we fix the number of alternating quantifiers k, the cardinality of the set of variables to which each quantifier applies (typically an uniform size, n), the number of clauses l, and the size of the clauses h. Each clause is generated by choosing h distinct variables, negating each with probability 1/2. Repeated clauses or clauses just containing universals (which are trivially unsatisfiable) are discarded.

Flawed and Flawless Problems

In SAT, empty and unit clauses are normally omitted in random generation methods where the number of literals in each clause varies. For example, in the 'constant probability' model proposed in (Mitchell, Selman, & Levesque 1992), each variable is included in a clause with some constant probability, but if only zero or one variable is included, the clause is discarded. An empty clause immediately makes a problem insoluble, but the reason for omitting unit clauses is more subtle. Suppose the model did not exclude unit clauses. Each clause generated from the n variables would have a certain probability of being unit. If the average clause size is constant, then for all n this probability is above some non-zero value q. As the l clauses are generated independently, about ql will be units. As there are only $2n$ different unit clauses, we expect to generate complementary unit clauses when $ql \approx \sqrt{2n}$, just as we expect to find two people with the same birthday in a group of about $\sqrt{365}$ people. If an instance contains complementary unit clauses it is trivially unsatisfiable. So we expect problems to be trivially unsatisfiable when $l = O(\sqrt{n})$, but non-trivial unsatisfiability occurs at $l = O(n)$. Phase transition behavior is therefore eventually dominated by trivial insolubility. In short, a naive version of the constant probability model would be *flawed*, but Mitchell et al's version is *flawless*.

An analogous flaw was identified by (Achlioptas *et al.* 1997) in much-used random models of binary constraint satisfaction problems, although most experiments reported in the literature use parameters that are too small to be affected by flaws. (Gent *et al.* 1998) proposes a *flawless* model which eliminates the unit constraints that lead to trivially insolubility.

Cadoli et al's random QSAT model also contains a *flaw*. A QSAT instance is trivially unsatisfiable if it contains one clause with a single existential and the rest universal, and a second clause with the negation of this existential and the rest universals distinct from the first set. Such a pair of clauses is unsatisfiable since, when all the universals are F, the two units that remain after simplification are contradictory. In Cadoli et al's model with equal numbers of existential and universal variables and a fixed clause size h, the probability of each clause being unit-existential is $h/(2^h - 1)$. This is independent of n and so bounded above 0 as $n \to \infty$. As before, we expect to see two clauses with the single existential literals complementary when $l = O(\sqrt{n})$. With h fixed and $n \to \infty$, these two clauses will almost certainly have disjoint sets of universals. Unlike the constraint satisfaction models, the flaw occurs at sufficiently small problem sizes to have had a significant impact on previous experimental studies. Table 1 in (Cadoli, Giovanardi, & Schaerf 1997) appear to confirm this argument. The phase transition in solubility occurs when l is approximately proportional to $\sqrt{n}$.

To remove such trivial problems, we propose two new generation methods. We propose the name 'unit-flawless' for these methods since instances are immune from the flaws we have identified caused by unit clauses. Since it is possible that other flaws might exist the name 'flawless' is not justified, but we use it below as shorthand for 'unit-flawless' in the context of this paper. In

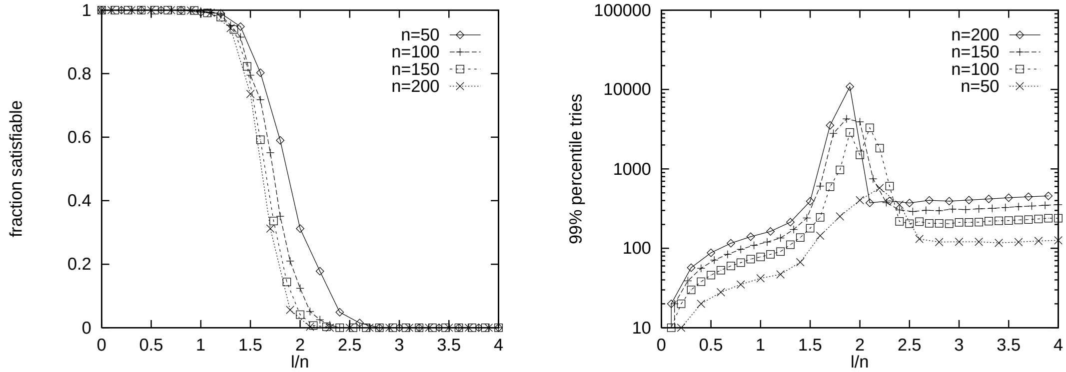

Figure 2: Random 2-QSAT problems from model A, $k = 3$, varying n and l/n. **(left)** fraction of satisfiable problems; **(right)** 99% percentile in search cost.

model A, we simply discard a clause that contains one or fewer existentials, and replace with a newly generated clause. In model B, we fix the number of existentials $e > 1$ that occur in each clause. We cannot generate a problem that is trivially unsatisfiable in either model.

Experimental verification

To show that trivially unsatisfiable problems dominate behavior in Cadoli et al's model, we ran an experiment with similar parameters to (Cadoli, Giovanardi, & Schaerf 1997). We use random 2-QSAT problems with $h = 3$. In this and subsequent experiments, we generate 1000 problems at each data point. Figure 1 shows the fraction of unsatisfiable problems and the fraction of these that are trivially unsatisfiable. We see that the phase transition is almost entirely due to trivially unsatisfiable problems. Only a few problems are unsatisfiable and not trivial, and the fraction of non-trivial ones goes down as n increases.

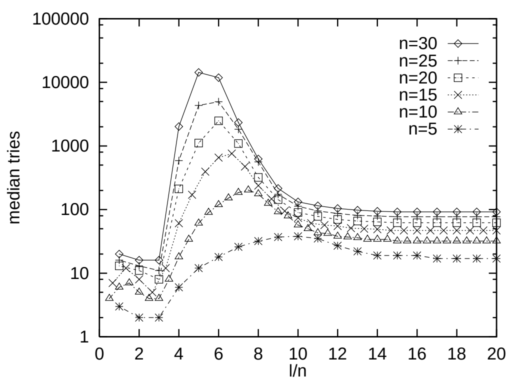

Figure 3: Median search cost of model A Random 2-QSAT problems, $h = 5$, varying n and l/n.

We next tested our proposed flawless model A, in which we discard clauses containing one or no existentials. We now observe a phase transition at an approximately fixed value of l/n. In Figure 2, we plot the fraction of satisfiable problems for random 2-QSAT problems generated by model A, with $h = 3$. The phase transition occurs around $l/n \approx 2$. There is an easy-hard-easy pattern in search cost but only in the higher percentiles. Notice that there is an increase in search cost after the phase transition, probably associated with the overheads of dealing with more clauses. Median search cost is rather uniform across the phase transition. Model A problems are significantly harder to solve than problems from the flawed model. With the flawed model, we easily ran a phase transition experiment at $n = 500$. With model A, we were unable to run a complete phase transition experiment for $n > 200$. For problems with larger clauses, the easy-hard-easy pattern is not restricted to the higher percentiles. For example, in Figure 3 we plot the median search cost for random 2-QSAT with $h = 5$. The phase transition now occurs around $l/n \approx 6$. and we observe an easy-hard-easy pattern in median cost.

We also tested our second flawless model, model B, in which number of existentials is fixed. With two or more existentials in every clause, we again see a phase transition at an approximately fixed value of l/n. For example, in Figure 4 we plot results for random 2-QSAT problems generated by model B in which 2 out of 3 literals in each clause are existentials. The phase transition occurs around $l/n \approx 1.4$. As in Figure 2, search cost increases with l/n after the phase transition. Although these are the highest costs, we expect that the peaks in the phase transition region will dominate as n increases since they will grow faster than the overheads. The easy-hard-easy pattern in search cost is again restricted to the higher percentiles, as in model A with $h = 3$. Problems from model B are more uniform, and they tend to have less variation in problem difficulty than problems from model A. As has been seen in NP problems, more uniform models tend to lead to fewer exceptionally hard problems (Gent & Walsh 1994; Smith & Grant 1995).

State space

We can use the theory of constrainedness of search problems proposed in (Gent *et al.* 1996) to predict the location of phase transitions like this. Whilst this the-

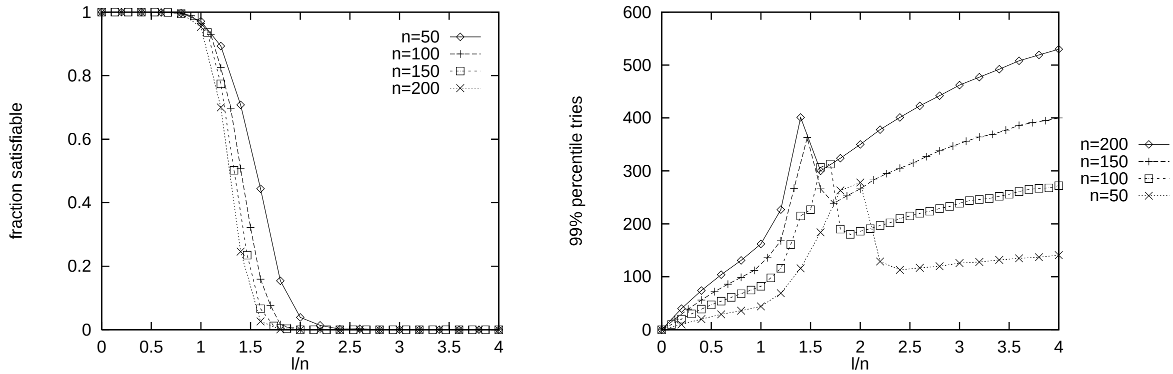

Figure 4: Random 2-QSAT problems from model B generated with 2 out of 3 literals in each clause being existentials, and varying n and l/n. **(left)** fraction of satisfiable problems; **(right)** 99% percentile in search cost.

ory was developed for NP-complete problems such as SAT, it has also been used in the complexity class P (Gent *et al.* 1997). To determine the constrainedness κ of random QSAT problems, we first identify the state space. A state is described by a set of substitutions for the existential variables. We need a set as there is a (possibly different) substitution for each set of values of the universal variables. The size of problems is equal to the number of bits needed to describe a state. For random 1-QSAT problems, we need n bits to describe a state as there are n existentials needing 1 bit each. For 2-QSAT problems, we need $n2^n$ bits to describe a state since there are 2^n different values for the universals, each of which requires n bits to specify the values for the existentials. In general, if we need s_k bits to specify a k-QSAT state for k even, we need $s_k + n$ bits for a k+1-QSAT state. And if we need s_k bits to specify a k-QSAT state for k odd, we need $2^n s_k$ bits to specify one k-QSAT state for each of the 2^n values of the new universal variables and hence a k+1-QSAT state. Thus, $2i$-QSAT problems have size $n\sum_{j=1}^{i} 2^{jn}$ and $(2i$+$1)$-QSAT problems have size $n\sum_{j=0}^{i} 2^{jn}$.

Constrainedness of QSAT

The informal notion of the constrainedness of a search problem has been formalised by the introduction of a constrainedness parameter, κ (Gent *et al.* 1996). Search problems with small κ values are underconstrained and almost surely soluble. Problems with large values of κ are over-constrained and almost surely insoluble. Inbetween, near $\kappa = 1$, problems are "critically constrained" and on the knife-edge between solubility and insolubility: in this region of κ we expect to see phase transition behavior and the hardest search problems. By definition (Gent *et al.* 1996) we have,

$$\kappa \quad =_{\text{def}} \quad -\frac{\log_2(\Pr\{\text{a state is a solution}\})}{\text{size of problem}}$$

We first derive a general formula for the constrainedness, κ, of k-QSAT problems. We will then specialize this formula for models A and B. For the general case, we assume that each of the l clauses has a probability p_j of containing exactly j existentials but that $p_0 = 0$ to exclude clauses without any existentials. To simplify computation, we assume that clauses are generated independently of each other.

$$\begin{aligned}
& \Pr\{\text{a state is a solution}\} \\
= \; & \Pr\{\text{a state satisfies a set of } l \text{ clauses}\} \\
= \; & (\Pr\{\text{a state satisfies a clause}\})^l \\
= \; & (\sum_j p_j \Pr\{\text{a state satisfies a clause of } j \text{ existentials}\})^l
\end{aligned}$$

Given a state, and a clause with j existentials, at least one of the $h - j$ universals is true in all but 1 out of 2^{h-j} of the substitutions in the state. Hence, for random $2i$- or $2i + 1$-QSAT problems, we need not consider 2^{in} different substitutions for the existentials, but just $2^{in}/2^{h-j} = 2^{in+j-h}$, each of which is assumed to be independent[1]. One of the j existentials in a clause is true in all but 1 out of 2^j cases. Hence, $\Pr\{\text{a state satisfies a clause of } j \text{ existentials}\} = (1 - 1/2^j)^{2^{in+j-h}}$. For random $2i$-QSAT, this gives us

$$\kappa \quad = \quad -\frac{l}{n}\frac{1}{\sum_{j=1}^{i} 2^{jn}} \log_2(\sum_{j=1}^{h} p_j (1 - 1/2^j)^{2^{in+j-h}})$$

Similarly, for random $2i + 1$-QSAT,

$$\kappa \quad = \quad -\frac{l}{n}\frac{1}{\sum_{j=0}^{i} 2^{jn}} \log_2(\sum_{j=1}^{h} p_j (1 - 1/2^j)^{2^{in+j-h}})$$

Model A

As clauses containing zero or one existentials are discarded, $p_0 = p_1 = 0$ and $p_j = \binom{h}{j} / \sum_{i=2}^{h} \binom{h}{i} = \binom{h}{j} / (2^h - h - 1)$ for $j > 1$. For random 2-QSAT,

$$\kappa \quad = \quad -\frac{l}{n}\frac{1}{2^n} \log_2(\sum_{j=2}^{h} p_j (1 - \frac{1}{2^j})^{2^{n+j-h}})$$

[1] For 2-QSAT, this assumption is correct. For k-QSAT for $k > 2$, the innermost existentials can vary more than the outermost, so the assumption will start to break.

For large h, we make a mean-field approximation that each clause has $h/2$ existentials and universals. Hence,

$$\begin{aligned}\kappa &\approx -\frac{l}{n}\frac{1}{2^n}\log_2((1-\frac{1}{2^{h/2}})^{2^{n-h/2}}) \\ &= -\frac{l}{n}\frac{1}{2^{h/2}}\log_2(1-\frac{1}{2^{h/2}}) \\ &\approx -\frac{l}{n}\log_2(1-\frac{1}{2^h})\end{aligned}$$

Model B

Each of the l clauses contains exactly e existentials. That is, $p_e = 1$ and $p_j = 0$ for $j \neq e$. For random 2-QSAT , this gives,

$$\kappa = -\frac{l}{n}2^{e-h}\log_2(1-\frac{1}{2^e})$$

As $\log_2(1+x) \approx x/\ln(2)$ for small x, if h and e are large then,

$$\kappa \approx \frac{l}{n}2^{e-h}\frac{1}{2^e\ln(2)} = \frac{l}{n}\frac{1}{2^h\ln(2)} \approx -\frac{l}{n}\log_2(1-\frac{1}{2^h})$$

This is the same approximation as we derived for model A. Note that the constrainedness is independent of e, the number of existentials provided this and the clause size are large. Where there are n universal and n existential variables, the constrainedness of random 2-QSAT problems from either model is approximately *double* the constrainedness of a random SAT problem with the same number of variables, clauses, and clause size. This is perhaps not too surprising. Universally quantifying half the variables in a SAT problem is likely to give a much more constrained problem.

Location of phase transition

Constrainedness can be used to predict the location of phase transitions. In many NP-complete problems, phase transition behavior is seen at $\kappa \approx 1$ (Gent *et al.* 1996). For model B problems with 2 out of 3 literals in each clause being existential, the phase transition occurs at $\kappa \approx 0.30$, ($l/n \approx 1.4$). For model A problems with $h = 5$, the transition is at $\kappa \approx 0.28$, ($l/n \approx 6$).

To investigate why phase transitions occur earlier than predicted, we ran experiments with model B problems with $n = 25$ and $h = 5$, i.e. 5 literals in each clause, and varying numbers of existentials e in each clause. The phase transition occurs at larger values of l/n as we increase the number of existentials. With fewer, more constraining universals, we need more constraints on each of the less constraining existentials. As we move from $e = 4$ to $e = 5$, we move down a complexity class to NP. Not too surprisingly, there is a significant drop in problem hardness as we move from $e = 4$ to $e = 5$. More surprisingly, problem hardness increases as we increase the number of existentials (and reduce the number of universals) in each clause from $e = 2$ (and 3 universals in each clause) to $e = 4$ (and just 1 universal). In Figure 5, we plot the constrainedness, κ, against the 99% percentile of search cost: as expected the peaks in search cost line up closely with the satisfiability transition. The location of the phase transition approaches $\kappa \approx 1$ as we increase the number of existentials in each clause. The rather constant shift of the phase transition as e increases may indicate a systematic error in our estimate for κ.

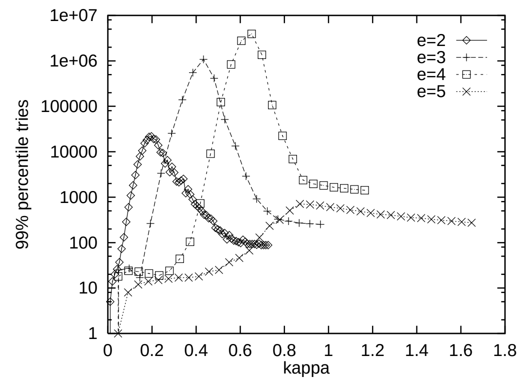

Figure 5: Model B problems with $n = 25$ and $h = 5$, varying number of existentials e, search cost *vs.* κ.

There is an alternative explanation for these errors. The prediction that the phase transition occurs around $\kappa \approx 1$ is based in part by the first moment Markov bound (that is, $prob(sol) \leq \langle Sol \rangle$). The location of phase transitions in NP problems can be predicted better by a second moment method using the variance in the number of solutions (Smith & Dyer 1996). In fact, at the SAT phase transition, an exponentially small number of problems have an exponential number of solutions (Kamath *et al.* 1995). In a PSPACE problem like QSAT, the variance in the number of solutions, and the super-exponential size of the state space, may result in the Markov bound being a less good predictor for the location of the phase transition.

Related and Further Work

This work is entirely novel in showing that the theory of constrainedness developed for NP problems (Gent *et al.* 1996) can be applied to a PSPACE problem. Since the theory has also been applied to a problem in P, that of establishing arc-consistency (Gent *et al.* 1997), constrainedness can be used both up and down the complexity hierarchy.

Phase transitions in PSPACE problems have been studied outside a general framework like the theory of constrainedness. Cadoli et al introduced an algorithm for QSAT and performed experimental evaluations (Cadoli, Giovanardi, & Schaerf 1997; 1998). As we discussed above, their randomly generated instances suffer from a flaw that can make them trivially insoluble. This may have given misleading impressions about the efficiency of their algorithm. We have introduced two new methods for generating 'flawless' random QSAT problems that are typically much harder.

The modal propositional logic K, which is PSPACE-complete, displays a phase transition and an easy-hard-easy pattern in search cost (Giunchiglia *et al.* 1998). The problem generator used in these experiments improves upon an earlier one that was criticised for giving instances that are 'trivial' as they are propositionally unsatisfiable (Hustadt & Schmidt 1997). Our use of the term 'trivial' is a complexity class lower: the unsatisfiable instances we identify in Cadoli et al's model can be found in almost linear time.

Bylander performed an average-case analysis on certain classes of problems within his model of random propositional STRIPS problems (Bylander 1996). He concluded by "suggesting that PSPACE-complete problems exhibit threshold phenomena similar to NP-complete problems." We can extend this conjecture by suggesting that the prediction of threshold phenomena using constrainedness can be extended from NP-complete problems through the polynomial hierarchy and to PSPACE-complete problems.

We have only tested our predictions experimentally for 2-QSAT. It would be interesting to investigate k-QSAT for $k > 2$, to see whether κ still makes reasonable predictions of the location of phase transitions. There are technical issues extending our definition of κ beyond k-QSAT to full QSAT (assuming the polynomial hierarchy does not collapse). Since any instance of a QSAT problem has some maximum number of alternating quantifiers, it would seem that our definition of κ for the relevant $k-$QSAT would apply. It is likely that this would work, provided that the values of p_j used in the derivation were correct. In particular, p_j would have to be the conditional probability of j existentials existing in a clause *given that* the instance contained exactly k alternating quantifiers. Even then, it is possible that technical difficulties would arise in defining κ, and we leave this question open for further investigation.

Conclusions

What general lessons can be learnt from this study? First, we can define the constrainedness of problems in PSPACE in a similar way to problems in NP. A phase transition in satisfiability and an easy-hard-easy pattern again occur at a critical value of constrainedness. However, predictions made by our theory are less accurate than in NP: we conjecture that this may be due to the huge state spaces of PSPACE-complete problems.

Second, we must take care to avoid trivially insoluble problems when generating random problems in new domains. Trivially insoluble problems have caused difficulties in propositional satisfiability, binary constraint satisfaction problems, and as we have shown here, QSAT. Since unit constraints are often the cause of trivially insolubility, we can usually generate *unflawed* problems by simply disallowing unit constraints.

Third, QSAT plays a similar role in PSPACE to the role played by SAT in NP. Because of this, we conjecture that constrainedness will be a useful theory in the study of many PSPACE problems of great interest in AI. Outstanding examples include games playing and planning problems. While it will be interesting to see phase transitions in these problems, it will be fascinating if constrainedness can be used, as it has in NP, to suggest new search methods.

References

Achlioptas, D.; Kirousis, L.; Kranakis, E.; Krizanc, D.; Molloy, M.; and Stamatiou, Y. 1997. Random constraint satisfaction: A more accurate picture. In *Proc. CP97*, 107–120. Springer.

Bylander, T. 1996. A probabilistic analysis of propositional STRIPS planning. *Artificial Intelligence* 81:241–271.

Cadoli, M.; Giovanardi, A.; and Schaerf, M. 1997. Experimental analysis of the computational cost of evaluation quantified Boolean formulae. In *Proceedings of the AI*IA-97*, 207–218. Springer-Verlag. LNAI-1321.

Cadoli, M.; Giovanardi, A.; and Schaerf, M. 1998. An algorithm to evaluate quantified Boolean formulae. In *Proc. AAAI-98*, 262–267.

Gent, I. P., and Walsh, T. 1994. Easy problems are sometimes hard. *Artificial Intelligence* 335–345.

Gent, I.; MacIntyre, E.; Prosser, P.; and Walsh, T. 1996. The constrainedness of search. In *Proc. AAAI-96*, 246–252.

Gent, I.; MacIntyre, E.; Prosser, P.; Shaw, P.; and Walsh, T. 1997. The constrainedness of arc consistency. In *Proc. CP-97*, 327–340. Springer.

Gent, I.; MacIntyre, E.; Prosser, P.; Smith, B.; and Walsh, T. 1998. Random constraint satisfaction: Flaws and structure. Technical Report APES-08-1998, APES research group. http://www.cs.strath.ac.uk/~apes/apereports.html.

Giunchiglia, E.; Giunchiglia, F.; Sebastiani, R.; and Tacchella, A. 1998. More evaluation of decision procedures for modal logics. In *Proc. KR 98.* Morgan Kauffmann.

Hustadt, U., and Schmidt, R. 1997. On evaluating decision procedures for modal logic. In *Proc. IJCAI-97*, 202–207.

Kamath, A.; Motwani, R.; Palem, K.; and Spirakis, P. 1995. Tail bounds for occupancy and the satisfiability threshold conjecture. *Randomized Structure and Algorithms* 7:59–80.

Mitchell, D.; Selman, B.; and Levesque, H. 1992. Hard and easy distributions of SAT problems. In *Proc. AAAI-92*, 459–465.

Papadimitriou, C. 1994. *Computational Complexity.* Addison-Wesley.

Smith, B., and Dyer, M. 1996. Locating the phase transition in binary constraint satisfaction problems. *Artificial Intelligence* 81:155–181.

Smith, B., and Grant, S. 1995. Sparse constraint graphs and exceptionally hard problems. In *Proc. IJCAI-95*, 646–651.