The Interface between P and NP: COL, XOR, NAE, 1-in-k, and Horn SAT

Toby Walsh
Cork Constraint Computation Center
University College Cork
Cork
Ireland
tw@4c.ucc.ie

Abstract

We study in detail the interface between P and NP by means of five new problem classes. Like the well known 2+p-SAT problem, these new problems smoothly interpolate between P and NP by mixing together a polynomial and a NP-complete problem. In many cases, the polynomial subproblem can dominate the problem's satisfiability and the search complexity. However, this is not always the case, and understanding why remains a very interesting open question. We identify phase transition behavior in each of these problem classes. Surprisingly we observe transitions with both smooth and sharp regions. Finally we show how these problem classes can help to understand algorithm behavior by considering search trajectories through the phase space.

Introduction

Where makes NP-complete problems hard to solve? Research into phase transition behavior has given much insight into this question. See, for example, (Cheeseman, Kanefsky, & Taylor 1991; Mitchell, Selman, & Levesque 1992; Gomes & Selman 1997; Walsh 1999). Propositional satisfiability (SAT) is the canonical NP-complete problem and one in which we have perhaps the most insight into phase transition behavior and problem hardness. For random SAT problems with few clauses, problems are almost surely satisfiable and it is easy to guess one of the many satisfying assignments. For random SAT problems with many clauses, problems are almost surely unsatisfiable and it is easy to prove that there can be no satisfying assignments. The hardest random SAT problems tend to be inbetween, when problems are neither obviously satisfiable or unsatisfiable. If we look more closely, especially within a search algorithm like the Davis Logemann Loveland (DLL) procedure, we see both polynomial subproblems (for example, clauses of length 2) and subproblems which are NP-complete (clauses of length 3 or more). As we explain in the next section, there may be interesting and unexpected interactions between them. To study this in more detail, Monasson et al. introduced the 2+p-SAT problem class (Monasson *et al.* 1999) which mixes together polynomial and NP-complete SAT problems and lets us explore in detail the interface between P and NP. We continue this research programme by introducing five new problems at the interface between P and NP.

2+p-SAT

A random k-SAT problem in n variables has l clauses of length k drawn uniformly at random. A sharp transition in satisfiability has been proved for random 2-SAT at $l/n = 1$ (Chvatal & Reed 1992; Goerdt 1992), and conjectured for random 3-SAT at $l/n \approx 4.3$ (Mitchell, Selman, & Levesque 1992). Associated with this transition is a rapid increase in problem difficulty. The random 2-SAT transition is continuous as the backbone (the fraction of variables taking fixed values) increases smoothly. On the other hand, the random 3-SAT transition is discontinuous as the backbone jumps in size at the phase boundary (Monasson *et al.* 1998).

To study this in more detail, Monasson et al. introduced the 2+p-SAT problem class (Monasson *et al.* 1999). This interpolates smoothly from the polynomial 2-SAT problem to the NP-complete 3-SAT problem. A random 2+p-SAT problem in n variables has l clauses, a fraction $(1-p)$ of which are 2-SAT clauses, and a fraction p of which are 3-SAT clauses. This gives pure 2-SAT problems for $p = 0$, and pure 3-SAT problems for $p = 1$. For any fixed $p > 0$, the 2+p-SAT problem class is NP-complete since the embedded 3-SAT subproblem can be made sufficiently large to encode other NP-complete problems within it.

By considering the satisfiability of the embedded 2-SAT subproblem and by assuming that the random 3-SAT transition is at $l/n \approx 4.3$, we can bound the location the random 2+p-SAT transition to:

$$1 \leq \frac{l}{n} \leq \min(\frac{1}{1-p}, 4.3)$$

Surprisingly, the upper bound is tight for $p \leq 2/5$ (Achlioptas *et al.* 2001b). That is, the 2-SAT subproblem *alone* determines satisfiability up to $p = 2/5$. Asymptotically, the 3-SAT clauses do not determine if problems are satisfiable, even though they determine the worst-case complexity. Several other phenomena occur at $p = 2/5$ reflecting this change from a 2-SAT like transition to a 3-SAT like transition. For example, the transition shifts from continuous to discontinuous as the backbone jumps in size (Monasson *et al.* 1998). In addition, the average cost to solve problems appears to increase from polynomial to exponential both for complete and local search algorithms (Monasson *et al.* 1998; Singer, Gent, & Smaill 2000). Random 2+p-SAT problem thus look like polynomial 2-SAT problems up to $p = 2/5$ and NP-complete 3-SAT problems for $p > 2/5$.

The 2+p-SAT problem class helps us understand the performance of the DLL algorithm for solving 3-SAT (Cocco & Monasson 2001). At each branch point in its backtracking search tree, the DLL algorithm has a mixture of 2-SAT and 3-SAT clauses. We can thus view each branch as a trajectory in the 2+p-SAT phase space. For satisfiable problems solved without backtracking (i.e. $l/n < 3$), trajectories stay within the satisfiable part of the phase space. For satisfiable problems that require backtracking (i.e. $3 < l/n < 4.3$), trajectories cross the phase boundary separating the satisfiable from the unsatisfiable phase. The length of the trajectory in the unsatisfiable phase gives a good estimate of the amount of backtracking needed to solve the problem. Finally, for unsatisfiable problems, trajectories stay within the unsatisfiable part of the phase space. The length of the trajectory again gives a good estimate of the amount of backtracking needed to solve the problem.

2+p-COL

Given the insight that 2+p-SAT has provided into computational complexity and algorithm performance, we decided to look more deeply into the interface between P and NP by means of some new problem classes. The first problem is 2+p-COL, a mix of 2-coloring and 3-coloring. A random k-COL problem consists of n vertices, each with k possible colors, and e edges drawn uniformly and at random. Like k-SAT, k-COL is NP-complete for $k \geq 3$ but polynomial for $k = 2$. To interpolate smoothly from P to NP, a random 2+p-COL problem has a fraction $(1 - p)$ of its vertices with 2 colors, and a fraction p with 3 colors. Note that the vertices with 2 colors are fixed at the start and cannot be chosen freely. We obtained similar results if the 2-colorable vertices have 2 colors chosen at random from the 3 possible or (as here) the same 2 fixed colors. Like 2+p-SAT, the 2+p-COL problem class is NP-complete for any fixed $p > 0$.

2+p-COL has one major difference to 2+p-SAT. Whilst 2-SAT, 3-SAT and 3-COL all have sharp transitions, 2-COL has a smooth transition (Achlioptas 1999). The probability that a random graph is 2-colorable is bounded away from 1 as soon as the average degree is bounded away from 0, and drops gradually as the average degree is increased, only hitting 0 with the emergence of the giant component (and an odd length cycle). Hence 2-colorability does not drop sharply around a particular average degree (as in 3-colorability) but over an interval that is approximately $0 < e/n < 1/2$.

In Figure 1, we see how the random 2+p-COL transition varies as we increase p and n. At $p \approx 0.8$, the 2+p-COL transition appears to sharpen significantly. In Figure 2, we look more closely at $p = 0.8$. For $p = 0.8$, there is a region up to around $e/n \approx 1.8$ in which the transition appears smooth and like that of 2-COL. The nature of the transition then appears to change to a sharp 3-COL like transition, with the probability of colorability dropping rapidly from around 90% to 0%. We thus appear to have both smooth and sharp regions.

We define δ_{2+p} as the largest ratio e/n at which 100% of problems are colorable:

$$\delta_{2+p} = \sup\{\frac{e}{n} \mid \lim_{n \mapsto \infty} \text{prob}(2+p\text{-colorable}) = 1\}$$

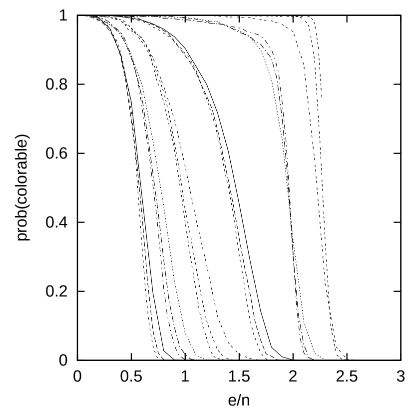

Figure 1: Probability that 2+p-COL problem is colorable (y-axis) against e/n (x-axis). Plots are for p=0 (leftmost), 0.2, 0.4, 0.6, 0.8 and 1.0 (rightmost) for 100, 200 and 300 vertex graphs (increasing sharpness).

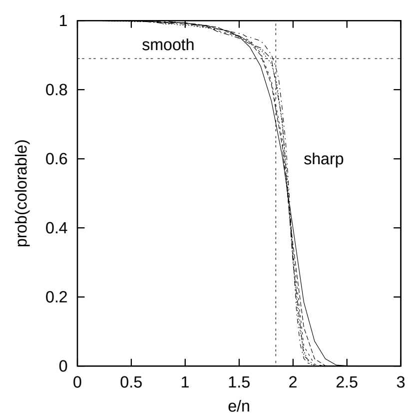

Figure 2: Probability that 2+p-COL problem is colorable (y-axis) against e/n (x-axis) for $p = 0.8$. Plots are for n=50, 100, 150, 200 and 250 vertex graphs.

From (Achlioptas 1999), $\delta_2 = 0$ and $1.923 < \delta_3 < 2.522$. For any fixed $p < 1$, a random 2+p-COL problem contains a 2-COL subproblem that grows in size with n and has average degree bounded away from 0. This subproblem has a probability of being 2-colorable that asymptotically is less than 1. Hence, the random 2+p-COL problem has a probability of being 2+p-colorable that asymptotically is also less than 1, and $\delta_{2+p} = 0$ for all $p < 1$.

We can define a dual parameter γ_{2+p}, which is the smallest ratio e/n at which 0% of problems are colorable:

$$\gamma_{2+p} = \inf\{\frac{e}{n} \mid \lim_{n \mapsto \infty} \text{prob}(2+p\text{-colorable}) = 0\}$$

Since colorability is a monotonic property (adding edges can only ever turn a colorable graph into an uncolorable graph), $\gamma_{2+p} \geq \delta_{2+p}$. Note that δ_{2+p} marks the start of the phase transition whilst γ_{2+p} marks its end. The start stays fixed at $\delta_{2+p} = 0$ for all $p < 1$ and jumps discontinuously to γ_3 at $p = 1$. The end appears to behave more smoothly, increasing smoothly as we increase p. From (Achlioptas 1999), $\gamma_2 \approx \frac{1}{2}$, and $1.923 < \gamma_3 < 2.522$. For a sharp transition like 3-coloring, $\gamma_3 = \delta_3$. As with 2+p-SAT:

$$\gamma_2 \leq \gamma_{2+p} \leq \min(\gamma_3, \frac{\gamma_2}{1-p})$$

In Figure 3, we have estimated experimentally the location of γ_{2+p} by analysing data for graphs with up to 300 vertices and compared the observed location of the (end of the) phase transition with the upper and lower bounds. As with 2+p-SAT, the upper bound (which looks just at the 2-COL subproblem) appears to be tight for $p < 0.8$.

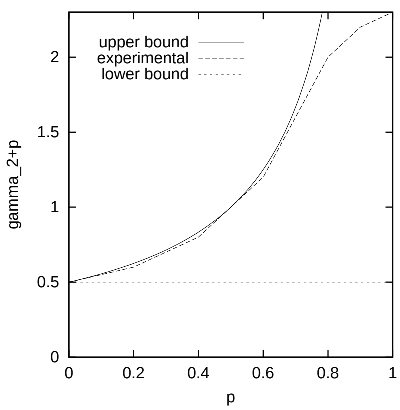

Figure 3: The location of the end of the 2+p-COL phase transition, γ_{2+p} (y-axis) against p (x-axis) for $p = 0$ to 1 in steps of 1/10.

The cost of 2+p-coloring also increases around $p \approx 0.8$. To solve 2+p-COL problems, we encode them into SAT and use zchaff, which is currently the fastest DLL procedure. Our results are thus algorithm dependent and should be repeated with other solvers. Note that the encodings of 2+p-COL problems into SAT give 2+p-SAT problems (but does not sample uniformly). In Figure 4, we see that there is a change in the search cost around $p \approx 0.8$ where we appear to move from polynomial to exponential search cost. This is despite 2+p-COL being NP-complete for all fixed and non-zero p. However, this is perhaps not so surprisng as up to $p \approx 0.8$, the polynomial 2-COL subproblem *alone* appears to determine asymptotically if the problem is colorable. Beyond this point, the NP-complete 3-COL subproblem contributes to whether the problem is colorable or not.

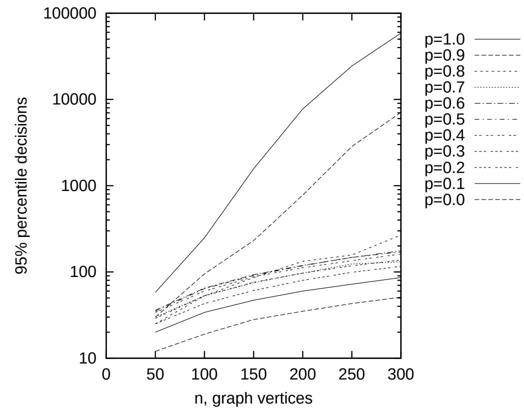

Figure 4: 95% percentile in the search cost to solve 2+p-COL problems at the phase boundary (y-axis logscale) against number of vertices (x-axis). Plots are for p=0 to 1 in steps of 1/10 (increasing hardness)

2+p-COL can be used in a similar way to 2+p-SAT to study 3-coloring algorithms. At each branching decision in a coloring algorithm, some vertices have three colors available, whilst others only have two. If any vertex has only a single color available, we are not at a branching point as we can commit to this color and simplify the problem. The algorithm thus has a sequence of 2+p-COL problems to solve, and we can view each branch in the algorithm's search tree as a trajectory in the 2+p-COL phase space. In Figure 5, we plot a number of trajectories for 3-coloring graphs with Brelaz's algorithm (Brelaz 1979). The trajectories are qualitatively very similar to those of the DLL algorithm in the 2+p-SAT phase space (Cocco & Monasson 2001). For $e/n \leq 1.5$, problems are solved without backtracking. Trajectories trace an arc that stays within the 'colorable' part of the phase space. On the other hand, the algorithm backtracks for problems with $e/n \geq 2$. For graphs with $e/n = 2$, the trajectory starts in the 'colorable' part of the phase space and crosses over into the 'uncolorable' part of the phase space. The algorithm then backtracks till we return to the 'colorable' part of the phase space. This sort of knowledge might be used both to model algorithms and to improve them. For example, we could develop a randomization and restart strategy which restarts when we estimate to have branched into an insoluble part of the phase space.

XOR SAT

We now turn to some other tractable satisfiability problems. Schaeffer's famous dichotomy result (Schaeffer 1978)

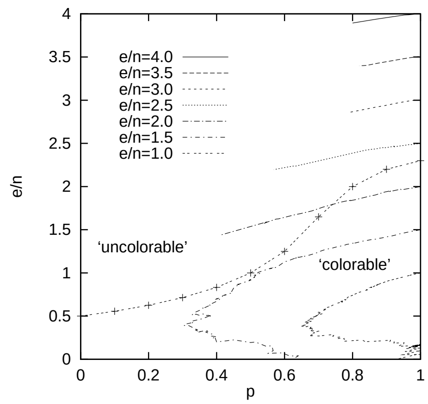

Figure 5: Trajectories in the 2+p-COL phase space for Brelaz's graph coloring algorithm on 3-COL problems with $n = 300$ and e/n from 1 to 4. The crossed line gives the experimental observed values of γ_{2+p}. The region marked 'uncolorable' is where graphs are asymptotically not 2+p-colorable, whilst the region marked 'colorable' is where graphs are asymptotically sometimes 2+p-colorable.

identifies the four non-trivial tractable and maximal restrictions of propositional satisfiability: 2-SAT, HORN SAT, dual HORN SAT, and XOR SAT. The later class is where clauses have the usual "or" replaced with an "exclusive or" (or its negation). This reduces the complexity from NP to P.

In a XOR k-SAT problem, we have l clauses defined over n variables, and each clause is an "exclusive or" of k literals (which ensures that an odd number of the literals are true) or its negation. Random XOR 3-SAT has a sharp threshold in the interval $0.8894 \leq l/n \leq 0.9278$ (Creognou, Daude, & Dubois 2001). Experiments put the transition at $l/n \approx 0.92$, whilst statistical mechanical calculations place it at $l/n = 0.918$ (Franz *et al.* 2001). To interpolate smoothly from P to NP, we introduce the random XOR2SAT problem, with a fraction $(1-p)$ of XOR 3-SAT clauses and a fraction p of 3-SAT clauses. Like 2+p-SAT, XOR2SAT is NP-complete for any fixed $p > 0$. Like 2+p-SAT but unlike 2+p-COL, the XOR2SAT threshold is always sharp.

In Figure 6, we have estimated experimentally the location of the phase boundary and compared it with the bounds:

$$0.92 \leq \frac{l}{n} \leq \min(\frac{0.92}{1-p}, 4.3)$$

The upper bound (which looks just at the polynomial XOR 3-SAT subproblem) appears to be loose for all $p > 0$. The NP-complete 3-SAT subproblem thus contributes to satisfiability for all $p > 0$. This contrasts with 2+p-SAT and 2+p-COL where the polynomial subproblem alone determines satisfiability for $0 < p \leq p_c$. It remains a very interesting open question why XOR2SAT is so different to 2+p-SAT in this respect. We expected the polynomial XOR 3-SAT subproblem in XOR2SAT to be even more dominate than the polynomial 2-SAT subproblem in 2+p-SAT. Each XOR 3-SAT clause rules out twice as many assignments as a 2-SAT clause, and the XOR 3-SAT phase transition occurs at a smaller clause to variable ratio than the 2-SAT transition. It is therefore very surprising that the XOR 3-SAT subproblem does not dominate the XOR2SAT phase transition like the 2-SAT subproblem dominates the 2+p-SAT phase transition.

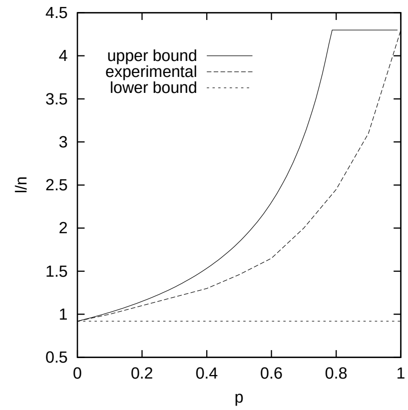

Figure 6: The location of the phase boundary for random XOR2SAT problems (y-axis) against p (x-axis) for $p = 0$ to 1 in steps of 1/10.

We encoded XOR2SAT problems into 3-SAT and solved them with zchaff. The search cost at the phase boundary appears to grow exponentially with n for all p. Given that the NP-complete 3-SAT subproblem contributes to satisfiability for all $p > 0$, it should perhaps not be surprising that we observe exponential growth in search cost for all $p > 0$.

Horn SAT

To complete the coverage of Schaeffer's dichotomy result, we turn to HORN SAT. Horn clauses of a fixed size $k > 1$ are trivially always satisfiable; every clause contains at least one negative literal and is satisfied by assigning all variables to false. We therefore consider a less trivial problem class in which we sample Horn clauses uniformly up to some fixed size. Such problems can contain (positive) unit clauses and so are not always satisfiable. In a k-HORN SAT problem, each clause is up to length k and Horn (i.e. contains at most one positive literal). Random k-HORN SAT has a smooth threshold whose shape is known analytically for $k = 2$ (Istrate 2000). Unlike 2-COL (whose smooth threshold starts at $e/n > 0$), the 2-HORN SAT threshold only starts at $l/n = 3/2$.

To interpolate smoothly from P to NP, we introduce the HORN2SAT problem, which has a fraction $(1-p)$ of 2-HORN SAT clauses and a fraction p of 3-SAT clauses. Like 2+p-SAT, HORN2SAT is NP-complete for any fixed $p > 0$. Although the HORN2SAT transition appears to sharpen immediately $p > 0$, there again appears to be a change around $p \approx 0.6$. The location of the transition starts to increase rapidly for $p > 0.6$. At the same time, search cost (especially in the higher percentiles) appears to go from polynomial to exponential

1-in-k-SAT

To finish our study of the interface between P and NP, we look at two more satisfiability problems in which phase transition behavior has been observed: 1-in-k-SAT and NAE SAT. In a 1-in-k-SAT problem, each clause specifies that exactly one out of k literals is true. For $k \geq 3$, the phase transition is sharp, occurs at $l/n = 2/k(k-1)$ and is continuous (Achlioptas *et al.* 2001a). Like k-SAT, 1-in-k-SAT is NP-complete for $k \geq 3$ and polynomial for $k = 2$. A 1-in-2-SAT problem can be readily mapped into a 2-COL problem and vice versa. Each 1-in-2-SAT clause fixes two variables either to opposite truth values (which is equivalent to an edge fixing two vertices to different colors) or to the same truth value (which is equivalent to merging two vertices so that they have the same color). We therefore expect the 1-in-2-SAT phase boundary to be smooth like that for 2-COL. To interpolate smoothly from P to NP, we introduce the 1-in-2+p-SAT problem, with a fraction $(1-p)$ of 1-in-2-SAT clauses and a fraction p of 1-in-3-SAT clauses. Like 2+p-SAT, this problem class is NP-complete for all fixed $p > 0$.

Although 1-in-3-SAT is NP-complete, random 1-in-3-SAT problems are much easier to solve than random 3-SAT problems of a similar size. This may be related to the fact that we can prove the exact location of the 1-in-3-SAT phase transition. The proof bounds the location of the phase transition from the satisfiable side by showing that a simple unit clause heuristic will almost surely decide all satisfiable problems. On the unsatisfiable side, the proof uses the fact that the backbone is of size $O(n)$ to show that adding a single clause is likely to cause unsatisfiability with constant probability. As satisfying assignments are easy to guess and proofs of unsatisfiability are likely to be short, random 1-in-3-SAT problems are unlikely to be hard to solve. Since there is no rapid jump in problem hardness, it is hard to be sure where the 1-in-2+p-SAT transition shifts from a smooth 1-in-2-SAT like transition to a sharp 1-in-3-SAT like transition. Our results suggest that this occurs in the interval $0.3 < p < 0.6$. As with 2+p-COL, the transition appears to have both smooth and sharp regions when p is small.

NAE SAT

Our final problem class is NAE SAT. In a NAE k-SAT problem, each clause specifies that k literals cannot all take the same truth value (i.e. they are Not All Equal). Like k-SAT, this problem class is NP-complete for $k \geq 3$ but is polynomial for k=2. Random NAE 3-SAT has a sharp threshold in the interval $1.514 \leq l/n \leq 2.215$ (Achlioptas *et al.* 2001a). Experimental results put the transition at $l/n \approx 2.1$.

To interpolate smoothly from P to NP, we introduce the random NAE 2+p-SAT problem, with a fraction $(1-p)$ of NAE 2-SAT clauses and a fraction p of NAE 3-SAT clauses. For any fixed $p > 0$, NAE 2+p-SAT is NP-complete. As with 2+p-COL, the transition appears to go from smooth (at $p = 0$) to sharp (at $p = 1$) and has both smooth and sharp regions for intermediate values of p. We again define a parameter γ_{2+p} identifying the end of the phase transition; this is the smallest clause to variable ratio at which 0% of problems are satisfiable. Like 2+p-COL,

$$\gamma_2 \leq \gamma_{2+p} \leq \min(\gamma_3, \frac{\gamma_2}{1-p})$$

In Figure 7, we have estimated experimentally the location of γ_{2+p}. The upper bound (which looks just at the NAE 2-SAT subproblem) appears to be tight for a large range of p. The polynomial NAE 2-SAT subproblem appears to determine satisfiability for p up to around 0.8. Not surprisingly, the search cost only appears to go from polynomial to exponential for $p \geq 0.8$.

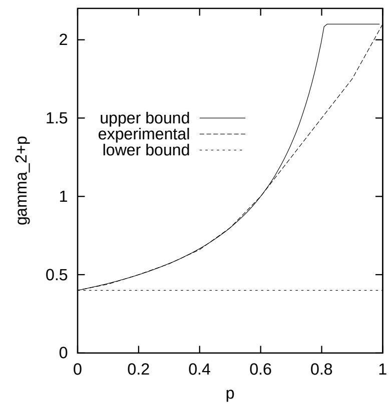

Figure 7: The location of the end of the NAE 2+p-SAT phase transition, γ_{2+p} (y-axis) against p (x-axis) for $p = 0$ to 1 in steps of 1/10.

A NAE 3-SAT clause on the literals a, b and c can be represented by the 3-SAT clauses $a \vee b \vee c$ and $\neg a \vee \neg b \vee \neg c$. We can therefore encode a NAE 3-SAT problem into 3-SAT by doubling the number of clauses (but keeping the number of variables constant). Achlioptas et al. observe that it would be "*truly remarkable*" if the random NAE 3-SAT phase transition occured at a clause to variable ratio half that of random 3-SAT since this encoding introduces significant correlations between the clauses (Achlioptas *et al.* 2001a). However, their experimental results put the random NAE 3-SAT phase transition at $l/n \approx 2.1$, which is tantalisingly close to half the 4.3 value believed to hold for random 3-SAT. To look at this issue in more detail, we introduce the NAE2SAT problem, with a fraction $(1-p)$ of NAE 3-SAT clauses and a fraction p of 3-SAT clauses.

If we ignore correlations between clauses, each NAE 3-SAT clause is twice as constraining as a 3-SAT clause. Hence $(1-p)l$ NAE 3-SAT clauses and pl 3-SAT clauses should behave like $2(1-p)l + pl$ 3-SAT clauses. That is, $(2-p)l$ 3-SAT clauses. The "effective" clause to variable ratio (in terms of 3-SAT clauses) is thus $(2-p)l/n$. To test this idea, Figure 8 plots the probability of satisfiability against $(2-p)l/n$. Very surprisingly, the phase transition occurs around an "effective" 3-SAT clause to variable ratio of 4.3. It appears that correlations between the NAE 3-SAT clauses can be almost completely ignored. To steal Achlioptas et al.'s words, this is "truly remarkable". NAE 3-SAT behaves like 3-SAT at twice the clause to variable ratio.

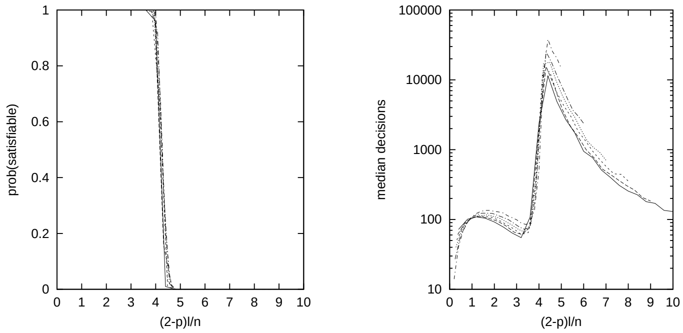

Figure 8: Probability that a NAE2SAT problem is satisfiable (left graph y-axis) and median search cost to solve a NAE2SAT problem (right graph, y-axis) against $(2-p)l/n$ (x-axes) for $n = 200$ and $p = 0$, 0.2, 0.4, 0.6, 0.8 and 1.

Conclusions

We have studied in detail the interface between P and NP by means of five new problem classes: 2+p-COL, XOR2SAT, HORN2SAT, NAE 2+p-SAT and 1-in-2+p-SAT. These problems smoothly interpolate between P and NP. In many cases, the polynomial subproblem dominates the problem's satisfiability and the search complexity up to some $p_c > 0$. For example, 2+p-COL behaves like the embedded polynomial 2-COL subproblem up to $p \approx 0.8$. Similarly, NAE 2+p-SAT behaves like the embedded polynomial NAE 2-SAT subproblem also up to $p \approx 0.8$. However, this is not always the case. In particular, in the XOR2SAT problem, both the 3-SAT clauses and the polynomial XOR 3-SAT clauses appear to contribute to the problem's satisfiability for all $p > 0$.

What important lessons can be learnt from this study? First, we can have transitions with both smooth and sharp regions. Problems like 2+p-COL and NAE 2+p-SAT let us study in detail how transitions sharpen and the large impact this has on search cost. Second, whilst the polynomial 2-SAT subproblem dominates 2+p-SAT up to $p = 2/5$, there are problems like XOR2SAT in which the polynomial subproblem does not dominate even though it is more tightly constraining than 2-SAT. Understanding this phenomenon is likely to bring fresh insight into problem hardness. Third, problem classes like these can help us understand algorithm behavior. For instance, we can view Brelaz's 3-coloring algorithm as searching trajectories in the 2+p-COL phase space. And finally, given the insights gained from studying the interface between P and NP, it is may be worth looking at the interface between other complexities classes. For example, we might study the interface between NP and PSpace.

Acknowledgements

This research was supported in full by an EPSRC advanced research fellowship. The author wishes to thank the members of the APES research group, as well as Carla Gomes, Bart Selman, and Christian Bessiere.

References

Achlioptas, D.; Chtcherba, A.; Istrate, G.; and Moore, C. 2001a. The phase transition in 1-in-k SAT and NAE SAT. In *Proc. of 12th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA'01)*, 719–720.

Achlioptas, D.; Kirousis, L.; Kranakis, E.; and Krizanc, D. 2001b. Rigorous results for (2+p)-SAT. *Theoretical Computer Science* 265(1-2):109–129.

Achlioptas, D. 1999. *Threshold phenomena in random graph colouring and satisfiability*. Ph.D. Dissertation, Dept. of CS, University of Toronto.

Brelaz, D. 1979. New methods to color the vertices of a graph. *Communications of ACM* 22:251–256.

Cheeseman, P.; Kanefsky, B.; and Taylor, W. 1991. Where the really hard problems are. In *Proc. of 12th IJCAI*, 331–337.

Chvatal, V., and Reed, B. 1992. Mick gets some (the odds are on his side). In *Proc. of 33rd Annual Symp. on Foundations of Computer Science*, 620–627. IEEE.

Cocco, S., and Monasson, R. 2001. Trajectories in phase diagrams, growth processes and computational complexity: how search algorithms solve the 3-satisfiability problem. *Phys. Rev. Letters* 86(8):1654–1657.

Creognou, N.; Daude, H.; and Dubois, O. 2001. Approximating the satisfiability threshold for random k-XOR-formulas. Tech. Report LATP/UMR6632 01-17, Lab. d'Informatique Fondamentale de Marseille.

Franz, S.; Leone, M.; Ricci-Tersenghi, F.; and Zecchina, R. 2001. Exact solutions for diluted spin glasses and optimization problems. *Phys. Rev. Letters* 87(12):127209.

Goerdt, A. 1992. A theshold for unsatisfiability. In Havel, I., and Koubek, V., eds., *Mathematical Foundations of Computer Science*, LNCS. 264–274.

Gomes, C., and Selman, B. 1997. Problem structure in the presence of perturbations. In *Proc. of 14th National Conference on AI*, 221–226.

Istrate, G. 2000. Dimension-dependent behavior in the satisfiability of random k-Horn formulae. Tech. report, Center for Nonlinear Studies, Los Alamos National Lab. Available from http://xxx.lanl.gov/ as cs.DS/0007029.

Mitchell, D.; Selman, B.; and Levesque, H. 1992. Hard and Easy Distributions of SAT Problems. In *Proc. of 10th National Conference on AI*, 459–465.

Monasson, R.; Zecchina, R.; Kirkpatrick, S.; Selman, B.; and Troyansky, L. 1998. Determining computational complexity for characteristic 'phase transitions'. *Nature* 400:133–137.

Monasson, R.; Zecchina, R.; Kirkpatrick, S.; Selman, B.; and Troyansky, L. 1999. 2+p SAT: Relation of typical-case complexity to the nature of the phase transition. *Random Structures and Algorithms* 15(3-4):414–435.

Schaeffer, T. 1978. The complexity of satisfiability problems. In *Proc. of 10th ACM Symp. on Theory of Computation*, 216–226.

Singer, J.; Gent, I.; and Smaill, A. 2000. Local search on random 2+p-SAT. In *Proc. of 14th ECAI*.

Walsh, T. 1999. Search in a small world. In *Proc. of 16th IJCAI*.