

# Heavy-Tailed Distributions in Combinatorial Search

Carla P. Gomes<sup>1</sup>, Bart Selman<sup>2</sup>, and Nuno Crato<sup>3</sup>

<sup>1</sup> Rome Laboratory, Rome Lab, NY 13441-4505, gomes@ai.rl.af.mil

<sup>2</sup> Computer Science Department, Cornell University, Ithaca, NY 14853,  
selman@cs.cornell.edu

<sup>3</sup> Dept. of Mathematics, New Jersey Institute of Technology, Newark, NJ 07102,  
USA, ncrato@m.njit.edu

**Abstract.** Combinatorial search methods often exhibit a large variability in performance. We study the cost profiles of combinatorial search procedures. Our study reveals some intriguing properties of such cost profiles. The distributions are often characterized by very long tails or “heavy tails”. We will show that these distributions are best characterized by a general class of distributions that have no moments (i.e., an infinite mean, variance, etc.). Such non-standard distributions have recently been observed in areas as diverse as economics, statistical physics, and geophysics. They are closely related to fractal phenomena, whose study was introduced by Mandelbrot. We believe this is the first finding of these distributions in a purely computational setting. We also show how random restarts can effectively eliminate heavy-tailed behavior, thereby dramatically improving the overall performance of a search procedure.

## 1 Introduction

Combinatorial search methods exhibit a remarkable variability in the time required to solve any particular problem instance. For example, we see significant differences on runs of different heuristics, runs on different problem instances, and, for stochastic methods, runs with different random seeds. The inherent exponential nature of the search process appears to magnify the unpredictability of search procedures. It is not uncommon to observe a combinatorial method “hang” on a given instance, whereas a different heuristic, or even just another stochastic run, solves the instance quickly.

We explore the cost distribution profiles of search methods on a variety of problem instances. Our study reveals some intriguing properties of such cost profiles. The distributions are often characterized by very long tails or “heavy tails”. We will show that these distributions are best captured by a general class of distributions that have no moments, *i.e.*, they have infinite mean, variance, etc.

---

Carla P. Gomes works for Rome Laboratory as a Research Associate. This work was performed while the second author was at AT&T Laboratories, Florham Park, NJ 07932-0971.

Heavy-tailed distributions were first introduced by the Italian-born Swiss economist Vilfredo Pareto in the context of income distribution. They were extensively studied mathematically by Paul Lévy in the period between the world wars. Lévy worked on a class of random variables with heavy tails of this type, which he called *stable* random variables. However, at the time, these distributions were largely considered probabilistic curiosities or pathological cases mainly used in counter-examples. This situation changed dramatically with Mandelbrot’s work on fractals. In particular, two seminal papers of Mandelbrot (1960, 1963) were instrumental in establishing the use of stable distributions for modeling real-world phenomena.

Recently, heavy-tailed distributions have been used to model phenomena in areas as diverse as economics, statistical physics, and geophysics. More concretely, they have been applied in stock market analysis, Brownian motion, weather forecasts, earthquake prediction, and recently, for modeling time delays on the World Wide Web (e.g., Mandelbrot 1983; Samorodnitsky and Taquq 1994). We believe our work provides the first demonstration of the suitability of heavy-tailed distributions in modeling the computational cost of combinatorial search methods.




Figure 1a: Erratic behavior of mean cost value.

Various researchers studying the computational nature of search problems have informally observed the erratic behavior of the mean and the variance of the search cost. This phenomenon has led them to use the median cost to characterize search difficulty. The heavy-tailed distributions provide a formal framework explaining the erratic mean and variance behavior. See Figure 1, for a preview of this phenomenon. Figure 1a shows the mean cost calculated over an increasing number of runs, on the same instance, of a backtrack style search procedure (details below). Contrast this behavior with that of the mean




Figure 1b: Mean for a standard distribution (gamma).

of a standard probability distribution (a gamma distribution; no heavy tails) as given in Figure 1b. In Figure 1b, we see that the sample mean converges rapidly to a constant value with increasing sample size. On the other hand, the heavy-tailed distribution in Figure 1a shows a highly erratic behavior of the mean that does not stabilize with increasing sample size.<sup>4</sup>

As a direct practical consequence of the heavy-tailed behavior of cost distributions, we show how randomized *restarts* of search procedures can dramatically reduce the variance in the search behavior. In fact, we will demonstrate that a search strategy with restarts can eliminate heavy-tailed distributions. This may explain the common informal use of restarts on combinatorial search problems.

## 2 Structured Search Problems

The study of the complexity and performance of search procedures when applied to realistic problems is greatly hampered by the difficulty in gathering realistic data. As an alternative, researchers heavily resort to randomly generated instances or highly structured problems from, *e.g.*, finite algebra. The random instances clearly lack sufficient structure, whereas the finite algebra problems are, in some sense, too regular. In order to bridge this gap, we introduced a new benchmark domain, the *Quasigroup Completion Problem* (Gomes and Selman 1997a).

A quasigroup is an ordered pair  $(Q, \cdot)$ , where  $Q$  is a set and  $(\cdot)$  is a binary operation on  $Q$  such that the equations  $a \cdot x = b$  and  $y \cdot a = b$  are uniquely solvable for every pair of elements  $a, b$  in  $Q$ . The *order*  $N$  of the quasigroup is the cardinality of the set  $Q$ . The best way to understand the structure of a

<sup>4</sup> The median, not shown here, stabilizes rather quickly at the value 1.

quasigroup is to consider the  $N$  by  $N$  multiplication table as defined by its binary operation. The constraints on a quasigroup are such that its multiplication table defines a *Latin square*. This means that in each row of the table, each element of the set  $Q$  occurs exactly once; similarly, in each column, each element occurs exactly once (Denes and Keedwell 1974).

An *incomplete* or *partial latin square*  $P$  is a partially filled  $N$  by  $N$  table such that no symbol occurs twice in a row or a column. The Quasigroup Completion Problem is the problem of determining whether the remaining entries of the table can be filled in such a way that we obtain a complete latin square, that is, a full multiplication table of a quasigroup. We view the pre-assigned values of the latin square as a *perturbation* to the original problem of finding an arbitrary latin square. Another way to look at these pre-assigned values is as a set of additional problem constraints on the basic structure of the quasigroup.

There is a natural formulation of the problem as a Constraint Satisfaction Problem. We have a variable for each of the  $N^2$  entries in the multiplication table of the quasigroup, and we use constraints to capture the requirement of having no repeated values in any row or column. All variables have the same domain, namely the set of elements  $Q$  of the quasigroup. Pre-assigned values are captured by fixing the value of some of the variables.

Colbourn (1983) showed the quasigroup completion problem to be NP-complete. In previous work, we identified a clear phase transition phenomenon for the quasigroup completion problem (Gomes and Selman 1997a). See Figures 2 and 3. From the figures, we observe that the costs peak roughly around the same ratio (approximately 42% pre-assignment) for different values of  $N$ . (Each data point is generated using 1,000 problem instances. The pre-assigned values were randomly generated.) This phase transition with the corresponding cost profile allows us to tune the difficulty of our problem class by varying the percentage of pre-assigned values.

An interesting application area of latin squares is the design of statistical experiments. The purpose of latin squares is to eliminate the effect of certain systematic dependency among the data (Denes and Keedwell 1974). Another interesting application is in scheduling and timetabling. For example, latin squares are useful in determining intricate schedules involving pairwise meetings among the members of a group (Anderson 1985). The natural perturbation of this problem is the problem of completing a schedule given a set pre-assigned meetings.

The quasigroup domain has also been extensively used in the area of automated theorem proving. In this community, the main interest in this domain has been driven by questions regarding the existence and nonexistence of quasigroups with additional mathematical properties (Fujita et al. 1993; Lam et al. 1989).

### 3 Computational Cost Profiles

In this section, we consider the variability in search cost due to different search heuristics. As our basic search procedure, we use a complete backtrack-style




Figure 2: The Complexity of Quasigroup Completion (Log Scale).




Figure 3: Phase Transition for the Completion Problem.

search method. The performance of such procedures can vary dramatically depending on the way one selects the next variable to branch on (the “variable selection strategy”) and in what order the possible values are assigned to a variable (the “value selection strategy”).

One of the most effective strategies is the so-called First-Fail heuristic.<sup>5</sup> In the First-Fail heuristic, the next variable to branch on is the one with the smallest remaining domain (*i.e.*, in choosing a value for the variable during the backtrack search, the search procedure has the fewest possible options left to explore —

<sup>5</sup> This is really a prerequisite for any reasonable backtrack-style search method. In theorem proving and Boolean satisfiability, the rule is related to the powerful unit-propagation heuristic.




Figure 4: Finding quasigroups of order 20 with 10% pre-assigned values.

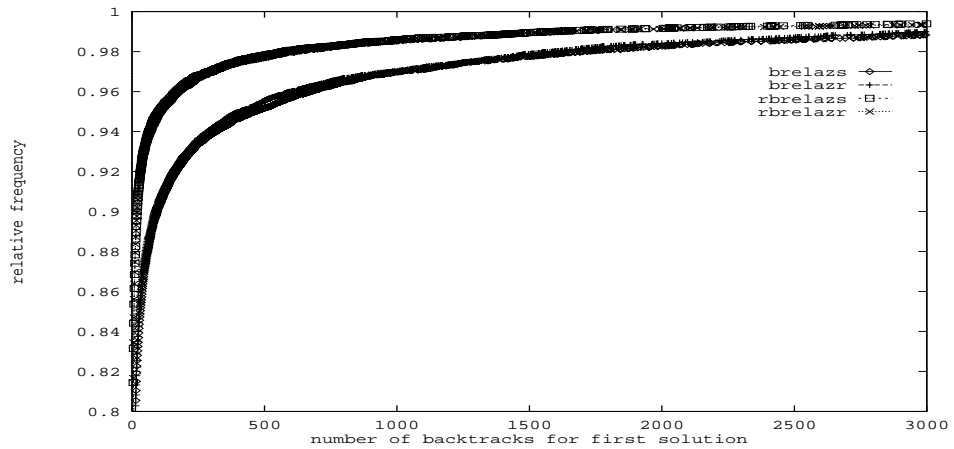


Figure 5: Finding quasigroups of order 10 at the phase transition.

leading to the smallest branching factor). We consider a popular extension of the First-Fail heuristic, called the Brelaz heuristics (Brelaz 1979), which was originally introduced for graph coloring procedures.

The Brelaz heuristic specifies a way for breaking ties in the First-fail rule: If two variables have equally small remaining domains, the Brelaz heuristic proposes to select the variable that shares constraints with the largest number of the remaining unassigned variables. A natural variation on this tie-breaking rule is what we call the “reverse Brelaz” heuristic, in which preference is given to the variable that shares constraints with the *smallest* number of unassigned variables. Any remaining ties after the (reverse) Brelaz rule are resolved randomly. (Note that such tie breaking introduces a stochastic element in our complete

search method.) One final issue left to specify in our search procedure is the order in which the values are assigned to a variable. In the standard Brelaz, value assignment is done in lexicographical order (*i.e.*, systematic). In our experiments, we consider four strategies:

- *Brelaz-S* — Brelaz with systematic value selection,
- *Brelaz-R* — Brelaz with random value selection,
- *R-Brelaz-S* — Reverse Brelaz with systematic value selection, and
- *R-brelaz-R* — Reverse Brelaz with random value selection.

We encoded this problem in C++ using ILOG SOLVER, a powerful C++ constraint programming library (Puget 1994). ILOG provides a backtracking mechanism that allows us to keep track of variables and their domains, while maintaining arc-consistency (van Hentenryck *et al.* 1992).

Figure 4 shows the performance profile of our four strategies for an instance of the quasigroup completion problem of order 20 with 10% pre-assigned values, *i.e.*, in the underconstrained area. Each curve gives the cumulative distribution obtained for each strategy by solving the problem 10,000 times. The cost (horizontal axis) is measured in number of backtracks, which is directly proportional to the total runtime of our strategies. For example, the figure shows that R-Brelaz-R, finished roughly 80% of the 10,000 runs in 15 backtracks or less.

First, we note that the (cumulative) distributions have surprising long tails after a steep initial climb. We will return to this issue below. We also see that R-Brelaz-R dominates the other strategies over almost the full range of the distribution. (Brelaz-S dominates very early on but the difference is not statistically significant.) Figure 5 shows the performance profile on an instance of the quasigroup completion problem in the critically constrained area. The initial climb followed by a long tail is even more dramatic. In this case, R-Brelaz-R and R-Brelaz-S give virtually the same performance, and both dominate the other two strategies.

These profiles suggest that it is difficult to take advantage of combining different heuristics in order to reduce variability. It was our initial intention to build so-called algorithm portfolios to reduce variability (Huberman *et al.* 1997 and Gomes and Selman 1997b). However, with one strategy dominating over the full profile there is no clear payoff in combining different heuristics, at least in this domain. In fact, it may well be the case that on a given problem domain, one can often find a single dominating heuristic. Our study here is not meant to be exhaustive regarding the full spectrum of search heuristics. In particular, we restricted ourselves to variations on the well-known Brelaz search heuristic.

In the next section, we concentrate on a perhaps more striking feature of the cost distributions: the *long tails*. As we will see in our section on “restarts”, the heavy tail behavior can be exploited effectively to reduce variability in the search cost.




Figure 6a: Probability distribution exhibiting heavy-tailed behavior.




Figure 6b: Log-log plot of heavy-tailed behavior.

## 4 Heavy-Tailed Distributions

Figure 6a shows the heavy-tailed nature of our cost distributions in a more direct manner. The probability distribution was obtained using R-Brelaz-R on an instance of the quasigroup completion problem of order 20 with 5% preplacement.<sup>6</sup>

<sup>6</sup> Work on exceptionally hard problems provides further support for the heavy tailed nature of the distributions (Gent and Walsh 1993; Smith and Grant 1995). However, the heavy tails we observed appear more ubiquitous: We observed heavy-tails in the majority of solvable instances in the under-constrained area and also in the majority of solvable instances in the critically constrained area. For other recent related work




Figure 6c: Log-log plot of standard distributions (no heavy tails).

We considered the distributions of over two dozen randomly picked instances from both the under-constrained and the critically constrained area, as well as some aggregate distributions. We found heavy-tailed distributions for almost all of our solvable instances and aggregate distributions. Some very easy solvable instances did not exhibit heavy tails. Interestingly, the unsolvable instances do not exhibit heavy-tails. The gamma and normal distributions were the best fit for the majority of our unsolvable instances (see also Frost et al. (1997)).

In order to model the long tail behavior of our distributions, we will consider distributions which asymptotically have tails of the Pareto-Lévy form, *viz.*

$$\Pr \{X > x\} \sim C \cdot x^{-\alpha}, \quad x > 0 \quad (1)$$

where  $\alpha > 0$  is a constant. These are distributions whose tails have a *hyperbolic decay*. For the case which concerns us it suffices to consider this tail behavior for the positive values of the random variable  $X$ . So, in what follows we will assume that the distribution has support on the positive half line, i.e.,  $\Pr \{0 \leq X < \infty\} = 1$ .

Mandelbrot (1983) provides an excellent introduction to these distributions with a discussion of their inherently self-similar or fractal nature. For a complete treatment of stable distributions see either Zolotarev (1986), or the more modern approach of Samorodnitsky and Taqqu (1994). See also de Lima (1997). In what follows, we simply outline the main results we will need to use.

A random variable  $X$  is said to have a *stable distribution* if for any  $n > 1$  there is a positive number  $C_n$  and a real number  $D_n$  such that

$$X_1 + X_2 + \cdots + X_n \stackrel{\mathcal{D}}{=} C_n X + D_n, \quad (2)$$

---

on cost distributions, see Frost et al. (1997) and Kwan (1995).

where  $X_1, X_2, \dots, X_n$  are independent copies of  $X$  and  $\stackrel{\mathcal{D}}{=}$  stands for equality in distribution. From this definition, it can be shown that the following is implied

$$C_n = n^{1/\alpha} \quad (3)$$

for some  $0 < \alpha \leq 2$  (Samorodnitsky and Taqqu 1994). The constant  $\alpha$  is called the *index of stability* of the distribution. Stable distributions with  $\alpha < 2$  have heavy tails of the Pareto-Lévy type. The index of stability is the same  $\alpha$  which appears in equation (1).

Since the existence or nonexistence of moments is completely determined by the tail behavior, it is simple to check that the index of stability  $\alpha$  is the *maximal moment exponent* of the distribution. For  $\alpha < 2$ , moments of  $X$  of order less than  $\alpha$  are finite while all higher order moments are infinite, i.e.,  $\alpha = \sup\{a > 0 : E|X|^a < \infty\}$ . For example, when  $\alpha = 1.5$ , the distribution only has a finite mean but no finite variance. When  $\alpha = 0.6$ , the distribution does not have a finite mean nor a finite variance.

While it is relatively easy to define a stable distribution, only in a few particular cases the density of the stable distributions is known in its *closed* form.

It should be noted, however, that distributions with tails of the form (1) are in the *domain of attraction* of stable distributions, i.e., properly normalized sums of variables with tails of the Pareto-Lévy type converge in distribution to an  $\alpha$ -stable random variable. This additive character of stable distributions matches the additive nature of the number of nodes searched in subtrees of the backtrack tree. This provides some intuition behind the suitability of the stable distributions for modeling search cost distributions.

In order to check for the existence of heavy tails in our distributions, we proceed in two steps. First, we graphically analyze the tail behavior of the sample distributions. Second, we formally estimate the index of stability.

If a Pareto-Lévy tail is observed, then the rate of decrease of the estimated density is hyperbolic — i.e., slower than the exponential rate. The complement to one of the cumulative distribution also displays a hyperbolic decay

$$1 - F(x) = \Pr\{X > x\} \sim C.x^{-\alpha}. \quad (4)$$

Then, for an heavy-tailed random variable, a log-log plot of the frequency of observed backtracks after  $x$  should show an approximate linear decrease at the tail. Moreover, the slope of the observed linear decrease provides an estimate of the index  $\alpha$ . In contrast, for a distribution with an exponentially decreasing tail, the log-log plot should show a faster-than-linear decrease of the tail.

Since the described behavior is a property of the tail we should mainly be concerned with the last observations, say the 10% observations that display a higher number of backtracks.

In Figure 6b, we have plotted three empirical cumulative distributions. One based on the probability distribution from Figure 6a (under-constrained), another for a medium constrained (solvable) instance, and a third for a critically constrained (solvable) instance. The linear nature of the tails in this log-log plot directly reveals tails of the Pareto-Lévy type.

For contrast we show in Figure 6c the log-log plots of two standard probability distributions. We see sharp rounded drop-off of both curves — indicating the absence of heavy tails. The distributions are given by the cost profiles on two unsolvable instances of our quasigroup completion problem. One is a rare unsolvable problem in the underconstrained area (best fit: a gamma distribution), the other is an unsolvable instance in the critically constrained region (best fit: normal distribution).

To complement our visual check of Figure 6b, and obtain an estimate of the index of stability (the value of  $\alpha$ ), we use the method of Hall (1982), which performs a regression on the extreme tails. Let  $X_{n1} \leq X_{n2} \leq \dots \leq X_{nn}$  be the order statistics, i.e., the ordered values of the sample  $X_1, X_2, \dots, X_n$  of the obtained number of backtracks. Set  $r < n$  as a truncation value which allows us to consider only the extreme observations. We obtain the estimator

$$\hat{\alpha}_r = \left( r^{-1} \sum_{j=1}^r \log X_{n,n-j+1} - \log X_{n,n-r} \right)^{-1} \quad (5)$$

This is a maximum likelihood estimator and Hall (1982) has established its asymptotic normality. Hall has also determined the optimal choice of the truncation parameter  $r$ . However, since this parameter is a function of the *unknown* parameters of the distribution, we adhere here to the common practice of using a set of values in the range  $\{n/10, n/25\}$ . This corresponds to severe truncations, which allow us to be more confident in our results.

We examined over two dozen distributions, and found values for  $\alpha$  that are consistent with the infinite variance hypothesis ( $\alpha < 2$ ) and, in many cases, they point to the nonexistence of the mean ( $\alpha < 1$ ). The estimates of  $\alpha$  for the distributions in Figure 6b were consistent with the hypothesis of infinite variance and infinite mean. The standard deviation in the estimates of the  $\alpha$  values were consistently an order of magnitude smaller than the estimates themselves, pointing to highly significant coefficients.

Are heavy-tailed distributions able to explain the strange sample mean discussed in the introduction? In other words, are stable distributions with index of stability of the order of magnitude of those estimated, able to generate data which reproduces the pattern shown in Figure 1a?

By using the method of Chambers, Mallows, and Stuck (1976), we generated random samples from a stable distribution, and calculated the mean as function of the number of samples. The resulting sequence of partial means is portrayed on Figure 7. The comparison between Figures 1a and 7 is striking, as the general wild oscillations are very similar and characteristic of heavy-tailed distributions.

## 5 Exploiting Heavy-Tailed Behavior

For our heavy-tailed distributions, we see that our procedures are in some sense most effective early on in the search. This suggests that a sequence of short runs instead of a single long run may be a more effective use of our computational




Figure 7: Behavior of mean for example stable distribution.

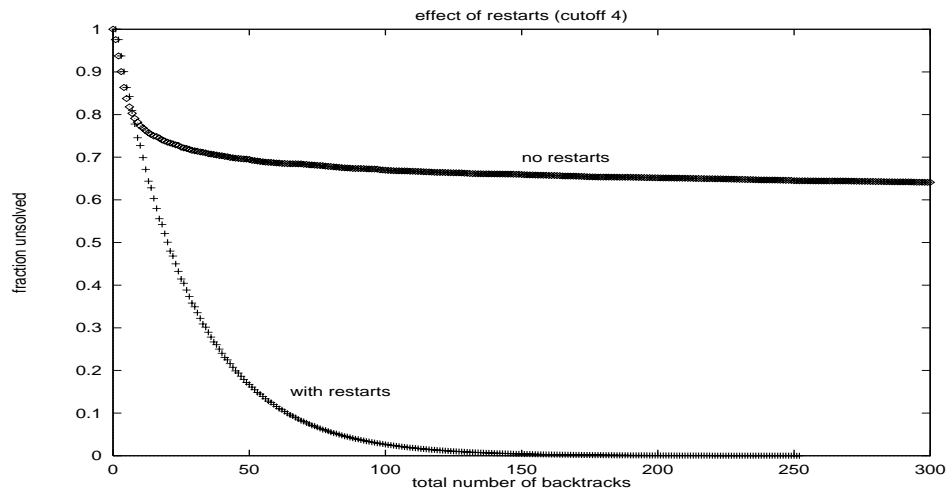


Figure 8a: Using restarts to exploit heavy-tailed behavior.

resources. We explore this idea by considering a fixed limit  $L$  on our overall cost (“run time”). From the cumulative cost distribution and  $L$ , we can determine what our expected probability of *not* solving the instance is because the search procedure runs out of time. We can also compute this “probability of failure” for a procedure that quickly restarts. Figures 8a and 8b give the results of such an analysis. (For more detailed results on the derivations of the probability distributions for restarts, see Gomes and Selman, Rome Lab Technical Report, 1997. For related work, see Gomes and Selman 1997b.)

The analysis was done for the completion problem of an instance of order 20 with 5% pre-placed. See distribution in Figure 6a. From Figure 8a, we see

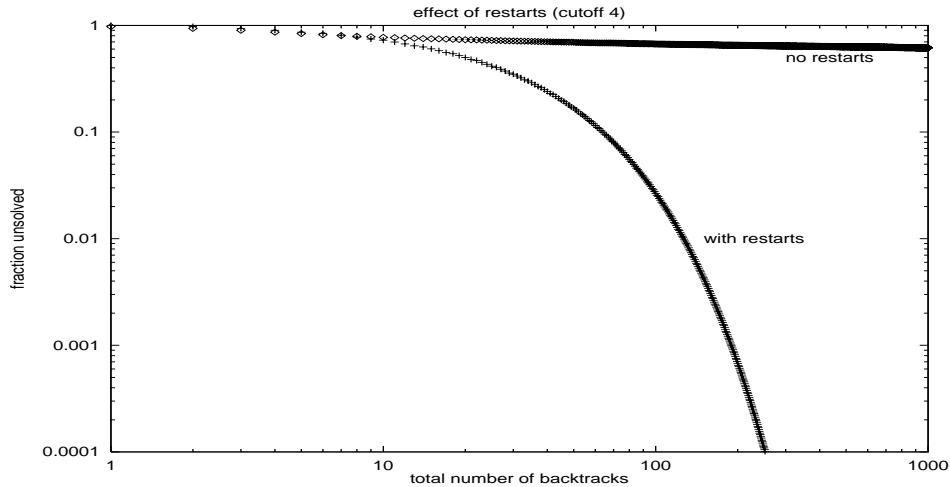


Figure 8b: Log-log plot for restarts.

that without restarts and given a total of 50 backtracks, we have a failure rate of around 70%. Using restarts (every 4 backtracks), this failure rate drops to around 10%. With an overall limit of only 150 backtracks, the restart strategy solves the instance almost always, whereas the original procedure still has a failure rate of around 70%. Such a dramatic improvement due to restarts is typical for heavy tailed distributions — in particular, we get similar results on critically constrained instances. Finally, Figure 8b shows a clear downward curve for the restart strategy. This suggests that the heavy-tailed nature of the cost distribution has disappeared. And, thus, we see that random restarts provide an effective mechanism for dealing with heavy-tailed cost distributions. These results explain the informal popularity of restart strategies in combinatorial search methods.

## 6 Conclusions and Future Work

We have revealed the special heavy-tailed nature of the cost distribution of combinatorial search procedures. We showed how such distributions can be modeled as stable distributions with heavy Pareto-Lévy type tails. Our analysis explains the empirically observed erratic behavior of the mean and variance of the cost of combinatorial search. And, more generally, the high variability observed between runs of such procedures.

Stable distributions have recently been used to capture a variety of real-world phenomena, such as stock market and weather patterns. We believe our results are the first indication of the occurrence of such distributions in purely computational processes. We hope that our results will further stimulate research along these lines by employing the special statistical tools available in this area.

We also showed how a “restart” strategy is an effective remedy against the heavy-tailed phenomena. Restarts drastically reduce the probability of failure

under limited time resources and reduce the overall variability of the method. Of course, when heavy tails are absent, restarts are much less effective. In our study, we did not encounter heavy tails for unsolvable instances.

**Acknowledgments** We would like to thank Karen Alguire for developing an exciting tool for experimenting with the quasigroup completion problem. We also would like to thank Nort Fowler for many useful suggestions and discussions, and Neal Glassman for suggesting the domain of combinatorial design as a potential benchmark domain. The first author is a research associate with Rome Laboratory and is funded by the Air Force Office of Scientific Research, under the New World Vistas Initiative (F30602-97-C-0037 and AFOSR NWV project 2304, LIRL 97RL005N25).

## References

- Anderson, L. (1985). Completing Partial Latin Squares. *Mathematisk Fysiske Meddelelser*, 41, 1985, 23–69.
- Brelaz, D. (1979). New methods to color the vertices of a graph. *Comm. of the ACM* (1979) 251–256.
- Chambers, John M., Mallows, C.L., and Stuck, B.W. (1976) A method for simulating stable random variables. *Journal of the American Statistical Association* **71**, 340–344.
- Cheeseman, Peter and Kanefsky, Bob and Taylor, William M. (1991). Where the Really Hard Problems Are. *Proceedings IJCAI-91*, 1991, 163–169.
- Colbourn, C. (1983). Embedding Partial Steiner Triple Systems is NP-Complete. *J. Combin. Theory (A)* 35 (1983), 100-105.
- Dechter, R. (1991) Constraint networks. *Encyclopedia of Artificial Intelligence* John Wiley, New York (1991) 276-285.
- de Lima, Pedro J.F. (1997). On the robustness of nonlinearity tests to moment condition failure. *Journal of Econometrics* **76**, 251–280.
- Denes, J. and Keedwell, A. (1974) Latin Squares and their Applications. *Akademiai Kiado, Budapest, and English Universities Press, London*, 1974.
- Frost, Daniel , Rish, Irina, and Vila, Lluís (1997) Summarizing CSP hardness with continuous probability distributions. *Proc. AAAI-97*.
- Fujita, M., Slaney, J., and Bennett, F. (1993). Automatic Generation of Some Results in Finite Algebra *Proc. IJCAI*, 1993.
- Freuder, E. and Mackworth, A. (Eds.). *Constraint-based reasoning*. MIT Press, Cambridge, MA, USA, 1994.
- Gent, I. P. and Walsh, T.. (1993). Easy Problems are Sometimes Hard, the DIMACS Challenge on Satisfiability Testing. Piscataway, NJ, Oct. 1993. Full version AIJ (Hogg et al. 1996).
- Gent, I. and Walsh, T. (1996) The Satisfiability Constraint Gap. *Artificial Intelligence*, 81, 1996.
- Gomes, C.P. and Selman, B. (1997a) Problem Structure in the Presence of Perturbations *Proc. AAAI-97*, Providence, RI, 1997.
- Gomes, C.P. and Selman, B. (1997b) Algorithm Portfolio Design: Theory vs. Practice, *Proc. UAI-97*, Providence, RI, 1997.
- Hall, Peter (1982) On some simple estimates of an exponent of regular variation. *Journal of the Royal Statistical Society*, **B 44**, 37–42.
- Huberman, B.A., Lukose, R.M., and Hogg, T. (1997). An economics approach to hard computational problems. *Science*, 265, 51–54.
- Hogg, T., Huberman, B.A., and Williams, C.P. (Eds.) (1996). Phase Transitions and Complexity. *Artificial Intelligence*, 81 (Spec. Issue; 1996)

- Kirkpatrick, S. and Selman, B. (1994) Critical Behavior in the Satisfiability of Random Boolean Expressions. *Science*, 264 (May 1994) 1297–1301.
- Kwan, Alvin C. M. (1995) Validity of normality assumption in CSP research, *PRICAI'96: Topics in Artificial Intelligence. Proceedings of the 4th Pacific Rim International Conference on Artificial Intelligence*, 459–465.
- Lam, C., Thiel, L., and Swiercz, S. (1989) The Non-existence of Finite Projective Planes of Order 10. *Can. J. Math.*, Vol. XLI, 6, 1989, 1117–1123.
- Mandelbrot, Benoit B. (1960) The Pareto-Lévy law and the distribution of income. *International Economic Review* 1, 79–106.
- Mandelbrot, Benoit B. (1963) The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Mandelbrot, B. (1983) *The fractal geometry of nature*. Freeman: New York. 1983.
- Mitchell, D., Selman, B., and Levesque, H.J. (1989) Hard and easy distributions of SAT problems. *Proc. AAAI-92*, San Jose, CA (1992) 459–465.
- Puget, J.-F. (1994) A C++ Implementation of CLP. *Technical Report 94-01* ILOG S.A., Gentilly, France, (1994).
- Russell, S and Norvig P. (1995) *Artificial Intelligence a Modern Approach*. Prentice Hall, Englewood Cliffs, NJ. (1995).
- Samorodnitsky, Gennady and Taqqu, Murad S. (1994) *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance*, Chapman and Hall, New York.
- Selman, B. and Kirkpatrick, S. (1996) Finite-Size Scaling of the Computational Cost of Systematic Search. *Artificial Intelligence*, Vol. 81, 1996, 273–295.
- Smith, B. and Dyer, M. Locating the Phase Transition in Binary Constraint Satisfaction Problems. *Artificial Intelligence*, 81, 1996.
- Smith, B. and Grant S.A., Sparse Constraint Graphs and Exceptionally Hard Problems. *IJCAI-95*, 646–651, 1995. Full version in AIJ (Hogg et al. 1996).
- van Hentenryck, P., Deville, Y., and Teng Choh-Man (1992) A generic arc consistency algorithm and its specializations. *Artificial Intelligence*, 57, 1992.
- Williams, C.P. and Hogg, T. (1992) Using deep structure to locate hard problems. *Proc. AAAI-92*, San Jose, CA, July 1992, 472–277.
- Zhang, W. and Korf, R. A Study of Complexity Transitions on the Asymmetric Travelling Salesman Problem. *Artificial Intelligence*, 81, 1996.
- Zolotarev, V.M. (1986) One-dimensional Stable Distributions. Vol. 65 of “Translations of mathematical monographs”, American Mathematical Society. Translation from the original 1983 Russian Ed.