

surements. Thus, OCT is a promising technique for both basic research and clinical applications.

presently with the Department of Ophthalmology, Tufts University School of Medicine, Boston, MA 02111. We thank T. F. Deutsch and R. Birngruber for helpful discussions, and we thank W. J. Woods for technical assistance. This work was supported by

National Institutes of Health contract 1-R01-GM35459-06 and Office of Naval Research contract N-00014-91-J-1956.

1 August 1991; accepted 23 September 1991

REFERENCES AND NOTES

1. G. N. Hounsfield, *Br. J. Radiol.* **46**, 1016 (1973).
2. R. Damadian, M. Goldsmith, L. Minkoff, *Physiol. Chem. Phys.* **9**, 97 (1977).
3. J. J. Wild and J. M. Reid, *Science* **115**, 226 (1952).
4. B. Chance et al., *Proc. Natl. Acad. Sci. U.S.A.* **85**, 4971 (1988).
5. J. G. Fujimoto et al., *Opt. Lett.* **11**, 150 (1986); D. Stern et al., *Inv. Ophthalmol. Vis. Sci.* **30**, 99 (1989).
6. K. G. Spears et al., *IEEE Trans. Biomed. Eng.* **36**, 1210 (1989); H. Chen et al., *Opt. Lett.* **16**, 487 (1991).
7. S. Andersson-Engels et al., *Opt. Lett.* **15**, 1179 (1990); L. Wang, P. P. Ho, C. Liu, G. Zhang, R. R. Alfano, *Science* **253**, 769 (1991).
8. R. C. Youngquist, S. Carr, D. E. M. Davies, *Opt. Lett.* **12**, 158 (1987); K. Takada et al., *Appl. Opt.* **26**, 1603 (1987).
9. G. S. Kino and S. S. C. Chim, *Appl. Opt.* **29**, 3775 (1990); B. S. Lee and T. C. Strand, *ibid.*, p. 3784.
10. A. F. Fercher, K. Mengedocht, W. Werner, *Opt. Lett.* **13**, 186 (1988); C. K. Hitzinger, *Inv. Ophthalmol. Vis. Sci.* **32**, 616 (1991).
11. D. Huang et al., *Lasers Surg. Med.* **11**, 5 (1991).
12. H. H. Gilgen et al., *J. Lightwave Technol.* **7**, 1255 (1989).
13. P. Davidovits and M. D. Egger, *Appl. Opt.* **10**, 1615 (1971).
14. R. B. Webb, G. W. Hughes, F. C. Delori, *ibid.* **26**, 1492 (1987).
15. T. Wilson and C. Sheppard, *Theory and Practice of Scanning Optical Microscopy* (Academic Press, London, 1984).
16. The refractive index is used because it closely approximates group index. The retinal refractive index of 1.36 is found in D. R. Williams, *Inv. Ophthalmol. Vis. Sci.* **19**, 653 (1980). The same value is used in other tissues.
17. A. Sommer, *Am. J. Ophthalmol.* **107**, 186 (1989).
18. H. A. Quigley and E. M. Addicks, *Arch. Ophthalmol.* **100**, 807 (1982); H. A. Quigley, *Trans. Am. Ophthalmol. Soc.* **84**, 920 (1986); H. A. Quigley, G. R. Dunkelberger, W. R. Green, *Am. J. Ophthalmol.* **107**, 453 (1989).
19. A. Sommer et al., *Arch. Ophthalmol.* **109**, 77 (1991); R. N. Weinreb et al., *ibid.* **108**, 557 (1990); L. F. Jindra et al., *Ophthalmic Surg.* **20**, 370 (1989).
20. R. C. Zeimer et al., *Arch. Ophthalmol.* **107**, 1006 (1989); T. Takenori and B. Schwartz, *Ophthalmology* **96**, 1315 (1989).
21. F. E. Kruse et al., *Ophthalmology* **96**, 1320 (1989); R. N. Weinreb, A. W. Dreher, J. F. Bille, *Int. Ophthalmol.* **13**, 25 (1989).
22. R. W. Knighton, S. G. Jacobson, C. M. Kemp, *Inv. Ophthalmol. Vis. Sci.* **30**, 2393 (1989).
23. The resolution of ultrasound is limited by the relatively long wavelengths needed for tissue penetration. Clinical ocular ultrasound typically uses 10-MHz ultrasound, which has a wavelength of 150 μm in tissue. The depth resolution of confocal laser scanning systems is limited by the aberration in ocular optics to roughly 200 μm (FWHM) according to J. F. Bille, A. W. Dreher, and G. Zinser, *Noninvasive Diagnostic Techniques in Ophthalmology*, B. R. Masters, Ed. (Springer-Verlag, Berlin, 1990), p. 528.
24. S. L. Jacques, *IEEE Trans. Biomed. Eng.* **36**, 1155 (1989).
25. The ocular safe exposure standard for permanent intrabeam viewing of 830-nm light with full 7-mm pupil aperture. ANSI Z136 (American National Standards Institute, 1986).
26. H. A. Haus, *Waves and Fields in Optoelectronics* (Prentice-Hall, Englewood Cliffs, NJ, 1984), p. 379.
27. K. M. Yoo, F. Liu, R. R. Alfano, *Opt. Lett.* **16**, 1068 (1991).
28. C. P. Lin, J. S. Schuman, and C. A. Puliafito are

Cooperative Solution of Constraint Satisfaction Problems

SCOTT H. CLEARWATER, BERNARDO A. HUBERMAN, TAD HOGG

It is widely believed that a group of cooperating agents engaged in problem solving can solve a task faster than either a single agent or the same group of agents working in isolation from each other. Nevertheless, little is known about the quantitative improvements that result from cooperation. A number of experimental results are presented on constraint satisfaction that both test the predictions of a theory of cooperative problem solving and assess the value of cooperation for this class of problems. These experiments suggest an alternative methodology to existing techniques for solving constraint satisfaction problems in computer science and distributed artificial intelligence.

THAT COOPERATION LEADS TO IMPROVEMENTS in the performance of a group of individuals underlies the founding of a firm, the existence of scientific and professional communities, and the establishing of committees charged with solving particular problems. In computation, the emergence of massively parallel machines underscores the assumed power of concurrency for solving very complex tasks that can be decomposed into smaller pieces, and a large effort is being devoted to the design of parallel algorithms for the solution of computationally hard problems.

Many of these tasks can be viewed as searches in large problem spaces. For realistic problems, where no algorithmic solution is known, heuristic methods are used to prune the search. Moreover, even with massively parallel machines, the huge size of the search space means that there will still be a large amount of search per processor. Recently, a theory that elucidates the performance of cooperative processes searching through a large problem space was developed (1). It showed that cooperative searches, when sufficiently large, can display universal characteristics, independent of the detailed nature of either the individual processes or the particular problem being tackled. This universality manifests itself in two separate ways: first, the existence of a sharp transition from exponential to polynomial time required to find the solution as heuristic effectiveness is improved (2); second, the appearance of a lognormal distribution in the effectiveness of an individual agent's

problem solving. The enhanced tail of this distribution guarantees the existence of some agents with superior performance. This can bring about a combinatorial implosion (3) with superlinear speedup in the time to find the answer with respect to the number of processes.

We present a number of experimental results on cooperative problem solving that both test the predictions of the theory and provide a quantitative assessment of the value of cooperation in solving a class of problems. These experiments were carried out by having a number of computational agents solve a set of cryptarithmic problems and measuring their individual and global performance. These results provide a striking example of the improvements in performance that result from cooperation and suggest an alternative methodology to existing techniques for solving constraint satisfaction problems in computer science and distributed artificial intelligence (4).

Cryptarithmic codes are typical of constraint satisfaction problems that lie at the heart of studies of human and computer problem solving (5, 6). The task is to find unique digit assignments to each of the letters so that the numbers represented by the words add up correctly. The constraint of a unique digit for a unique letter in a decimal representation reduces the total number of possible states, N_{states} , from 10^n to $10!/(10-n)!$ where n is the number of unique letters in the problem. A familiar example of such problems is provided by the sum: DONALD + GERALD = ROBERT. This particular problem has $n = 10$ and one solution, which is given by $A = 4$, $B = 3$, $D = 5$, $E = 9$, $G = 1$, $L = 8$, $N =$

Dynamics of Computation Group, Xerox Palo Alto Research Center, Palo Alto, CA 94304.

6, $O = 2$, $R = 7$, $T = 0$. Solving this problem involves performing a search. The speed at which an agent can solve the problem depends on the initial conditions and the particular sequence of actions it chooses as it moves through a search space. This sequence relies on the knowledge, or hints, that an agent has about which step should be taken next.

A simple noncooperative search strategy is to generate and test; that is, at each step an assignment is made to each letter and tested to see if the problem is solved. This search can be described by the probability of finding a solution at step t , $P_{\text{success}}(t)$. The probability distribution that an agent finds a solution at its r th step, or trial, is given by

$$P(t) = \prod_{r=0}^{t-1} [1 - P_{\text{success}}(r)] P_{\text{success}}(t) \quad (1)$$

where t runs from 0 (finding the correct answer immediately) to infinity (never finding a solution). For instance, if an agent chooses assignments randomly, then $P_{\text{success}}(t) = N_{\text{sol}}/N_{\text{states}}$, where N_{sol} is the number of solutions to the problem and N_{states} is the number of states in the search space that gives rise to a geometric distribution of $P(t)$.

A second noncooperative strategy is to partition the search space and restrict each agent to perform the same generate and test search only in a particular partition. Unlike the previous case, the partitioning prevents agents from duplicating each other's work and can be expected to give somewhat better performance, as we will show below.

We now turn to the case of cooperation, which takes the form of reading and writing hints to a central blackboard (7) that can be accessed by all agents. The use of a blackboard as the method of communication is not critical to our model or results. Other forms of cooperation, such as genetic algorithms (8), can also be used.

In the cryptarithmic case, hints are lists of letter-digit assignments that add up correctly modulo 10 for at least one column (9). For the example considered above $N = 7$, $R = 2$, and $B = 9$. Each agent asynchronously chooses hints randomly from the blackboard. If there are no hints, or if the agent has already used the chosen one, a random letter-digit assignment is chosen as in the first noncooperating strategy. Once the agent obtains the new state, it generates all possible hints from its state and posts them to the blackboard unless they already exist. Assignments that work for more than one column are posted as several different hints.

From the derivation in (1), which assumed the independence of hints, we expect

the probability of cooperating agents solving the problem in t steps is given by the lognormal distribution

$$P(t) = \frac{N(\mu, \sigma, \ln t)}{t} = \Lambda(\mu, \sigma, t) = \frac{1}{\sigma t \sqrt{2\pi}} \exp \frac{(\ln t - \mu)^2}{2\sigma^2} \quad (2)$$

which is unimodal and has a very long tail. Here μ and σ denote the mean and standard deviation of the logarithm of the fraction of the search space pruned by a hint.

These theoretical discussions are confirmed by our experimental results. Specifically, the speed distribution for the cases of (i) noncooperating agents, (ii) noncooperating agents with a partitioned search space, and (iii) cooperating agents was compiled for agents trying to solve the problem $WOW + HOT = TEA$. Figure 1A shows the resulting speed distributions for a typical run for each of the three cases, where speed




Fig. 1. (A) Probability density distribution of $\ln(\text{speed})$ obtained by fitting typical runs for agents trying to solve $WOW + HOT = TEA$. The leftmost curve is for 100 noncooperating agents, the center curve is for partitioned search (using 720 partitions and agents), and the rightmost curve is for 100 cooperating agents. The left curve corresponds to a geometric time distribution and the other two to lognormals. The conversion to a log scale allows all three cases to be shown on the same plot and also makes them appear similar to a normal distribution. (B) Speed probability density on a linear scale for the rightmost curve of (A). The speed of the fastest agent is 100 times greater than for the noninteracting case. The error bars are proportional to the square root of the number of agents. In (A) and (B) the results were not statistically different from that predicted by the fitted distributions using the $P = 0.01$ level of significance with the Kolmogorov-Smirnov test.

is defined as the inverse time to solve the problem for an agent. The striking overall improvement in performance becomes apparent once we note that the speedup is about 100 times larger for the performance distribution of the cooperating agents than that obtained for the noncooperating case and about 10 times larger than for the partitioned case. The distribution for the cooperative case is shown with a linear scale in Fig. 1B. The actual speedup for the cooperative case varies greatly from run to run because it is very sensitive to the quality of the first few hints posted to the blackboard. Similar results were obtained on other cryptarithmic problems that used between five and ten different letters and had from 1 to over 100 solutions.

The key difference between cooperating and noncooperating agents is that hints effectively reduce the size of the search space by focusing the agents on much more plausible courses of action. Moreover, an agent searching one part of the space may find a hint useful to another agent in a different part of the space. From our results, we have seen that a collection of cooperating agents can improve their overall performance with this simple mechanism. In fact, this improvement takes place throughout the search; the cooperating agents increase their number of correct columns at a faster rate than noncooperating ones.

We have also examined a number of extensions to our basic results, all of which showed lognormal behavior for the cooperating case. These were motivated by situations found in other constraint satisfaction problems.

In many situations accessing the blackboard will incur some cost. In our basic analysis the hint blackboard access was assumed to be cost-free, which is reasonably valid for cases where the agents select hints randomly off the blackboard.

In reality, agents often have specific expertise. We modeled this by specializing the agents to deal with only certain kinds of hints. We observed that this restriction reduced the average speedup. However, this resulted in a wider range of speed. Hence, with specialization the faster performers benefit while the slower performers suffer.

Also, agents often start a problem with some prior knowledge. We modeled this by using an initial blackboard that had hints on it. The effect of the quality of the agents' initial knowledge of the search space was modeled by the inclusion of hints on the initial blackboard. We found that a non-empty initial blackboard significantly increases the number of unique solutions found. This is easily understood because the initial collection of hints will likely point to

different solutions, whereas with an empty blackboard the direction of the agents' searches is highly focused by the first few arriving hints. Also, we observed that a nonempty initial blackboard leads to a smaller range of speeds because, with many hints already available, the importance of hint selection strategy becomes less important.

In conclusion, we have shown how cooperating toward the solution of a constraint satisfaction problem can increase the speed with which it is solved as compared to either the noncooperating case or a partitioned problem space, even with very simple agents and hints. The resulting distribution of performance agrees with the theoretical predictions and provides a quantitative assessment of the value of cooperating in problem solving. These agents were very sensitive to the first few hints. There remains the question of how much individual expertise is required to give more reproducible results. This would allow prediction of μ and σ , and hence scaling with the number of agents.

The way agents interpret their hints has a strong effect on the rate at which they solve the problem. This is especially so for the fastest and slowest agents. For a sufficiently large number of agents, the group with the highest diversity in interpretation was able to solve the problem first. Interestingly, high diversity not only leads to very fast performers but to very slow ones as well. These slow performers are necessary because they provide some hints used by the fastest agents.

This work suggests an alternative to the current mode of constructing task-specific computer programs that deal with constraint satisfaction problems. Rather than spending all the effort in developing a monolithic program or perfect heuristic, it may be better to have a set of relatively simple cooperating processes work concurrently on the problem while communicating their partial results. This would imply the use of "hint engineers" for coupling previously disjoint programs into interacting systems that can use each other's (imperfect) knowledge.

Because our results confirm a theory that provides a quantitative relation between performance, number of agents, and the ability of agents to use diverse hints, this new methodology may be particularly useful in areas of artificial intelligence such as design, qualitative reasoning, truth maintenance systems, and machine learning. Researchers in these areas are just starting to consider the benefits brought about by massive parallelism and concurrency.

2. ——— and T. Hogg, *Artif. Intell.* **33**, 155 (1987).
3. W. A. Kornfeld, *Commun. Assoc. Comput. Mach.* **25**, 734 (October 1982).
4. L. Gasser and M. Huhns, *Distributed Artificial Intelligence* (Morgan Kaufmann, San Mateo, CA, 1989), vol. 2.
5. A. Newell and H. A. Simon, *Human Problem Solving* (Prentice-Hall, Englewood Cliffs, NJ, 1972).
6. W. A. Kornfeld and C. E. Hewitt, *IEEE Trans. Syst. Man Cybern.* **SMC-11**, 24 (1981).
7. R. Englemore and T. Morgan, *Blackboard Systems* (Addison Wesley, Reading, MA, 1988).

8. D. E. Glover, in *Genetic Algorithms and Simulated Annealing*, L. Davis, Ed. (Morgan Kaufmann, Los Altos, CA, 1987), pp. 12–31.
9. Although a particular column may add up correctly, those letter-digit assignments may not lead to a solution, that is, the column hints are only locally correct.
10. This work was supported in part by Air Force Office of Scientific Research contract F49620-90-C-0086.

4 June 1991; accepted 11 September 1991

Radical Reactions of C₆₀



P. J. KRUSIC, E. WASSERMAN, P. N. KEIZER, J. R. MORTON, K. F. PRESTON

Photochemically generated benzyl radicals react with C₆₀ producing radical and nonradical adducts R_nC₆₀ (R = C₆H₅CH₂) with $n = 1$ to at least 15. The radical adducts with $n = 3$ and 5 are stable above 50°C and have been identified by electron spin resonance (ESR) spectroscopy as the allylic R₃C₆₀• (3) and cyclopentadienyl R₅C₆₀• (5) radicals. The unpaired electrons are highly localized on the C₆₀ surface. The extraordinary stability of these radicals can be attributed to the steric protection of the surface radical sites by the surrounding benzyl substituents. Photochemically generated methyl radicals also add readily to C₆₀. Mass spectrometric analyses show the formation of (CH₃)_nC₆₀ with $n = 1$ to at least 34.

THE RECENT DISCOVERY THAT C₆₀ can be produced in macroscopic quantities (1) has sparked much interest in the chemistry of this unusual molecule. An important aspect of the chemistry of C₆₀ is its reactivity towards free radicals. The molecule has, in effect, 30 carbon-carbon double bonds to which free radicals can add. Our previous results (2) indicated that multiple additions of a variety of radicals can indeed take place very readily, warranting characterization of this molecule as a radical sponge. No structural information was available, however. Here we identify by electron spin resonance (ESR) two extraordinarily stable, prototypical radical types formed in the addition of benzyl radicals to C₆₀. Crucial to the analysis of the complex mixture was information obtained by: (i) ¹³C labeling of the entering benzyl radicals in the α position; and (ii) the different ESR power saturation behavior of the two radical structures. The latter allowed selection of either spectrum simply by varying the microwave power level.

C₆₀ dissolves in a limited number of organic solvents, notably toluene. Since this solvent is susceptible to attack by photolytically generated *tert*-butoxy radicals with the formation of benzyl radicals (3), we used

this route to study the addition of benzyl radicals to C₆₀.



In a typical experiment, 50 μl of di-*tert*-butyl peroxide were added to 350 μl of a saturated (~ 3 mM) solution of C₆₀ in sodium-dried, oxygen-free toluene in a quartz ESR tube. The solution was then irradiated at various temperatures in the cavity of an ESR spectrometer. Focused ultraviolet light (UV) of a high-pressure mercury discharge lamp was used that was filtered to remove the visible and much of the infrared radiation by an aqueous NiSO₄·CoSO₄ filter (4).

The UV irradiation of such solutions at room temperature produced a single ESR absorption that grew steadily to a maximum and that did not decay when the light was extinguished. This absorption had a most unusual microwave power saturation behavior. Unlike the ESR spectra of most carbon-centered radicals, the spectrum did not power saturate even with the full output of the microwave source (~ 200 mW). Also, the g factors measured at high (200 mW) and at low (200 μW) incident powers were significantly different (2.00221 and 2.00250, respectively) so that the absorption shifted noticeably (~ 0.5 G) along the magnetic

REFERENCES AND NOTES

1. B. A. Huberman, *Physica D* **42**, 38 (1990).

P. J. Krusic and E. Wasserman, Central Research and Development, E. I. du Pont de Nemours & Company, Wilmington, DE 19880-0328.

P. N. Keizer, J. R. Morton, K. F. Preston, Steacie Institute for Molecular Sciences, National Research Council of Canada, Ottawa, Canada K1A 0R9.