

Cooperative Coloring

By William Kolata

One goal of computer science is to develop algorithms for solving classes of problems that have practical import. A useful question researchers might ask in pursuing this goal is how humans solve such problems. A further interesting question is how the structure of the problem affects the ability to solve it.


In a recent article in *Science* (Volume 313, August 11, 2006, pages 824–827), Michael Kearns and colleagues in the Department of Computer and Information Science at the University of Pennsylvania report on an experiment in distributed problem solving: a coloring problem on a series of graphs. The researchers gave a set of students at the university an interface to a distributed computer system that provided a common set of colors and a local view of the current state of a graph coloring; each student, given control of a single vertex, could change the color of that vertex at any time. The number of colors in each case was equal to the chromatic number of the graph (the minimum number of colors needed to color the vertices of a graph so that no two adjacent vertices have the same color). If the students collectively colored the graph within the allotted time (300 seconds), each received a reward of \$5.

Graph coloring has practical import. It is also a difficult problem: As the authors point out, graph coloring is known to be intractable even when the number of colors allowed is greater than the chromatic num-

ber. The problem “generalizes many traditional problems in logistics and operations research,” they write; it pops up, for example, in channel selection for wireless communication.

Six graphs were used in the experiment. The first three—one simple cycle and two cycles with additional randomly chosen chords—simulated the combination of local and long-distance connectivity found in social networks. The fourth graph was a “leader cycle” (a cycle with two distinguished vertices, the “leaders,” each of which is highly connected). The final two graphs were more complicated; these “preferential graphs” were formed incrementally, with new links gradually and stochastically added to already highly connected vertices. Preferential graphs are known to have certain structures, such as hubs. The chromatic numbers of the graphs used in the experiment are 2 for the first four and 3 and 4, respectively, for the last two.

So how successful were the students, and what features can be discerned from their solution strategies? In almost all cases (a total of 38 trials, 6 or 7 for each graph), the students were able to color the first four graphs, the cycles, in the allotted time; the simple cycle proved to be the most difficult. Solution times were proportional to the average distance of the graph (the average number of links between any two vertices). The authors conjecture that, even though the cycles with chords have some vertices with more complicated local connectivity, they have shorter solution times because they



A three-colorable preferential-attachment graph. Courtesy of Michael Kearns.

reduce the length of the paths required to disseminate changes to the current state.

The preferential graphs proved to be more difficult, accounting for six of the seven failures to reach a solution. In contrast to the cycle results, solution times for the preferential graphs were inversely proportional to the average distance.

The authors also compared the human solutions with solutions produced by a simple heuristic algorithm. In the case of a conflict, if there are one or more unused colors in the local neighborhood of a vertex, the algorithm chooses among them at random, thereby eliminating the conflict; otherwise, it chooses at random among all colors for the vertex. The results are the opposite of those for the human subjects: Cycles were more difficult than preferential graphs, and the relation between average distance and solution difficulty was reversed for both cyclic and preferential graphs.

To investigate the impact of allowing the students to have more information on the status of a coloring, the authors extended

the local view with two options: a local view with summary information on the connectivity of the locally connected vertices and a global view of the entire graph. Increasing the information reduced solution times for the cycles but had the opposite effect for the preferentially connected graphs.

The authors view these results as “suggestive.” For the future, they envision larger-scale experiments on the Web, a wider range of problem scenarios, with subjects participating in the formation of the network structure rather than reacting to an imposed structure, and a wider range of problems for which collective solutions can be implemented.

William Kolata is SIAM’s technical director.