Behavioral Graph Coloring

Michael Kearns, Nick Montfort, Siddharth Suri
Department of Computer and Information Science
University of Pennsylvania
3330 Walnut Street
Phildadelphia, PA 19104
(mkearns|nickm|ssuri)@cis.upenn.edu

Introduction

Beginning with the pioneering experiments of Travers and Milgram in 1969, there is a long and fascinating literature examining the structural and navigational properties of natural social networks. Findings range from the now familiar folklore of "six degrees of separation" to more recent theoretical explanations of the heuristics people might employ to exploit such structure. This line of investigation can be summarized in computer science terminology: Using relatively local information, distributed human organizations can collectively compute good approximations to the all-pairs shortest paths problem.

Given the volume and visibility of this research, it is perhaps suprising that there is little work on its natural generalization — namely, what *other* types of distributed optimization problems can humans networks solve? In this paper we describe the preliminary findings of a series of behavioral experiments we have been conducting at the University of Pennsylvania. Human subjects attempt to perform distributed graph coloring in a setting in which each subject "plays" a single vertex in a large and potentially complex graph. Players asynchronously update their choice of color for their vertex in an attempt to avoid or remove coloring conflicts. The experimental system allows us to vary the graph topology, the locality of information given to subjects, and the incentive scheme[1].

System and Experimental Methodology

Both our experiments and system were designed to permit the investigation of three main variables: the structure or topology of the graph being colored; the amount and locality of information given to each subject; and the incentive or payment scheme. Before providing details on the system, our procedures, and the values for these experimental variables, we first comment on our choice of graph coloring as the optimization problem to examine, for which the main reasons were threefold.

First, we were interested in choosing a problem which had a notably different status, from the computer science perspective, than the shortest paths problem, which has been the focus of the long literature that partially inspired this research. Thus, whereas shortest paths is known to be a computationally easy problem for centralized computation, graph coloring is notoriously hard, since it is NP-hard to even weakly approximate the chromatic number. Second, we were interested in a problem that, despite being possibly challenging to solve, was easy for human subjects with no special background to quickly understand. Third, we sought a problem which requires global coordination for its solution, but in which each subject could locally verify their contribution or hinderance to this solution.

[1] As we shall describe, subjects were paid according to performance under two different schemes.

System Description

The system we built for our experiments provides subjects with a simple browser-based visual interface allowing asynchronous updating of color choice and a view of the current experiment's state. The system permits us to execute a pre-planned series of graph coloring problems with specified graph topologies and information views (discussed below). During a coloring exercise, each subject sees an interface divided into two panels. The left-hand or *action panel* provides colored buttons that can be used to change the color of the subject's vertex. The right-hand or *information panel* provides varying amounts of information (see below), continually refreshed, on the color choices of other players, but always includes at least the color choices of the subject's immediate neighbors. The information panel always indicates how many color conflicts there are in the subject's neighborhood, if any, and graph edges with color conflicts are highlighted in bold lines. In addition, the right-hand panel always includes a "progress bar" at the bottom indicating how many conflicts remain globally.

The system logs fine-grained temporal data on the exact sequence of events in each coloring exercise. This log contains every color-change event, indexed by vertex or subject number, the color selected, and a timestamp with 1-second resolution. The system also administers and logs entry and exit questionnaires to each subject.

Experimental Procedures

We now briefly describe or experimental protocol, which was approved by Penn's Institutional Review Board process. Sessions were held in a laboratory containing 38 workstations, which determined the size of our subject pool for each session. Subjects were drawn from a Penn undergraduate computer science class on a related topic with no prerequisites, and were required to attend a preliminary lecture in which they were instructed about the graph coloring problem, the workings of the system, and the specifics of what they would see and how they would be paid.

Each experimental session consisted of a series of 19 consecutive graph coloring problems, for which a maximum of 5 minutes each was allocated. The experimental sessions thus lasted between one and two hours. During the experiments, physical partitions were erected to prevent subjects from glancing at other subjects or their screens. A timekeeper called out how much of the allowed 5 minutes was left at various points during each coloring problem. Subjects were carefully observed throughout the session to make sure they were violating any protocols, which including not speaking or communicating with any other subjects, attempting to look at the workstations of other subjects, etc. Each problem ended either after 5 minutes or when a proper coloring was found by the subjects, whichever occurred first, and the session proceeded to the next coloring problem. It is important to note that in each coloring problem, the number of colors provided to subjects was exactly equal to the chromatic number of the graph (which was computed in advance off-line). Thus we deliberately held subjects to the highest standard of optimal coloring, rather than exploring approximations.

By choosing 6 different graph topologies, 3 different information views, and 2 different incentive schemes, we generated a total of $6 \times 3 \times 2 = 36$ unique experimental conditions. All 18 corresponding to one of the incentive schemes were conducted on the evening of January 24, 2005, and the 18 corresponding to the other incentive scheme were given the following evening. The order of problems within each session were chosen randomly. In addition, to examine potential "learning" effects, each evening ended with the identical problem it began with, for a total of 19 coloring problems per session.

We now proceed to describe our choices of values for our three main experimental variables, beginning with the graph topologies used.

Graph Topologies

The space of possible graph topologies is obviously immense. Given our interest in contrasting behavioral graph coloring with the aforementioned literature on social network theory and navigation, our choices were closely guided by the generative models proposed in that line of research. However, since it was an open question whether human subjects could solve these kinds of problem efficiently under *any* conditions, we also desired a certain breadth of approach. For these reasons, we drew topologies from two recent but rather different stochastic models for network formation.

The first of these was the so-called "small worlds" family, in which a simple cycle is augmented with a variable number of randomly chosen chords. Larger numbers of chords are known to dramatically decrease the (average or worst-case) diameter, and are meant to model long-distance relationships in social networks arising from chance encounters and the like. We examined three topologies from this family: a simple 38-cycle with no chords, a graph consisting of a cycle with 5 chords, and a cycle with 20 chords. Rather than choosing the chords uniformly at random, we selected them at random from among all chords that would not increase the chromatic number beyond the 2 colors required for the simple cycle[2]. This has the advantage of allowing us to model long-distance connections (and thus reduce diameter) while in a mathematical sense make the problem strictly harder (since we have simply added constraints without requiring more colors).

The second model we examined is known as *preferential attachment*. In this model, a graph is built incrementally by adding one new vertex at a time. A new vertex is given a fixed number ν of edges to the existing graph; but rather than these edges being chosen uniformly at random, they are directed to an existing vertex with a probability proportional to its current degree. Among other properties, this stochastic model is known to generate heavy-tailed distributions of degrees (modeling the social phenomenon of "connectors") as well as small diameter.

Finally, we created one topology in the cycle family intended to experiment with more "engineered" or hierarchical structures, such as one might find in corporations or the military. In this graph, a 36-cycle is augmented by two "leader" vertices, one of which is connected to all even vertices on the cycle, the other to all the odd vertices. The leaders are also connected to each other. The resulting graph remains two colorable, but now has very low diameter and two vertices with very high degree.

Information Views

Our system can be configured to provide three different information views in the right-hand panel of the user interface; as has been mentioned, each of the 6 graph topologies was presented to the population under all three information views, as well as under both incentive schemes discussed below. We emphasize that while the information view varied from problem to problem, in any given exercise *all 38 participants* were given the same view. We have not experimented with different subjects having different views.

In the *low* information view, subjects could see only the chosen color of their own vertex, and the colors of their immediate neighbors in the graph. The *medium* information view is identical to the low, except now each neighbor is annotated with the (static) value of its degree. This view was motivated by the desire to provide subjects with some minimal additional information on the local structure that suggested which of their neighbors might have a more difficult coloring task. In the *high* information view, each subject could see the entire graph of 38 vertices as well as the dynamic color choices. In all three information views, the display was continually refreshed to provide subjects with the latest color choices.

[2] This amounts to restricting to chords between vertices whose indices have opposite parity.

Incentive Schemes

In line with the standards of behavioral economics and related fields, we paid subjects according to their performance, but examined two different schemes for doing so. In the *collective* incentive scheme (which was used on the first of two evenings of experiments), for each of the 19 coloring problems, each subject was paid $5 for each graph that was properly colored (no coloring conflicts anywhere in the graph within 5 minutes). If even a single conflict remained after 5 minutes, none of the 38 subjects received any payment for that problem. In the *individual* incentive scheme (used the second evening), each subject was paid $5 if at the conclusion of a problem (either due to proper coloring or the end of the 5 minutes) if *they* participated in no color conflicts, regardless of the global outcome.

These two schemes were introduced to allow the study of possible behavioral differences between a “team” and “greedy” incentive mechanism. A natural question to ask is whether such differences can arise in a problem such as coloring, where a subject’s contribution to the global solution is already locally determined.