Games and Economic Behavior 82 (2013) 702–716

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Information-sharing in social networks

Jon Kleinberg [a,1], Katrina Ligett [a,b,*,2]

[a] *Department of Computer Science, Upson Hall, Cornell University, Ithaca, NY 14853, USA*
[b] *Department of Computing and Mathematical Sciences and Division of the Humanities and Social Sciences, California Institute of Technology, MC 305-16, Pasadena, CA 91125, USA*

ARTICLE INFO

Article history:
Received 25 May 2012
Available online 23 October 2013

JEL classification:
C70

Keywords:
Social networks
Information propagation
Networks
Graphs

ABSTRACT

We present a new model for reasoning about the way information is shared among friends in a social network and the resulting ways in which the social network fragments. Our model formalizes the intuition that revealing personal information in social settings involves a trade-off between the benefits of sharing information with friends, and the risks that additional gossiping will propagate it to someone with whom one is not on friendly terms but who is within one's community. We study the behavior of rational agents in such a situation, and we characterize the existence and computability of stable information-sharing configurations, in which agents do not have an incentive to change the set of partners with whom they share information. We analyze the implications of these stable configurations for social welfare and the resulting fragmentation of the social network.

1. Introduction

A growing line of work on privacy has investigated ways for people to engage in *transactions* — purchases, queries, participation in activities, and related types of behavior — while revealing very little private information about themselves. This research has implicitly construed the problem of privacy as one of a trade-off between the concrete tasks that a person wants or needs to accomplish, and the "leakage" of personal information that might result from the interactions required to perform the task.

If one takes this view of privacy, however, it becomes very hard to reason about the kinds of simple, privacy-revealing activities that people engage in when they share personal information with friends. In fact, a large body of research (Adams and Blieszner, 1994; Aries and Johnson, 1983; Bazerman et al., 1998; Reiman, 1976) indicates that people derive utility from sharing such information as part of the formation and maintenance of friendships. Thus, we argue that it is natural to try modeling the trade-off between the benefits of sharing information with friends and the dangers that this information will reach people from whom the originator wishes to hide it.

In defining the model, we abstract away the decisions a person makes about any one piece of information. We imagine a relatively small community of people V, and an undirected *conflict graph* H over the nodes V in which an edge between individuals i and j indicates that each of i and j wishes to keep her personal information from reaching the other. For

* Corresponding author at: Department of Computing and Mathematical Sciences and Division of the Humanities and Social Sciences, California Institute of Technology, MC 305-16, Pasadena, CA 91125, USA. Fax: +1 626 792 4257.

E-mail addresses: kleinber@cs.cornell.edu (J. Kleinberg), katrina@caltech.edu (K. Ligett).

[1] This work has been supported in part by NSF grants CCF-0325453, BCS-0537606, IIS-0705774, IIS-0910664, CCF-0910940, IIS-1016099, a Simons Investigator Award, an ARO MURI grant, a Google Research Grant, a Yahoo! Research Alliance Grant, and the John D. and Catherine T. MacArthur Foundation.

[2] Work completed in part while the author was a postdoctoral associate at Cornell University. This work has been supported in part by an NSF Computing Innovation Fellowship (NSF Award CCF-0937060), an NSF Mathematical Sciences Postdoctoral Fellowship (NSF Award DMS-1004416), NSF grants CCF-0910940 and CNS-1254169, the Charles Lee Powell Foundation, and a Microsoft Faculty Fellowship.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.10.002

simplicity, we will refer to such pairs i and j as *enemies*. To represent the sharing of information, we think of the community as partitioned into disjoint "information-sharing groups," corresponding to the groups who are privy to each other's personal information. We model individuals as unwilling to be a member of such a group if it contains one of her enemies.[3] Any allowable partition of the community accordingly corresponds to the assignment of a label to each node (specifying the node's set in the partition), in such a way that if two nodes are connected by an edge in H, then they must receive different labels. In the terminology of graph theory, this is known as the problem of *graph coloring* (or, more specifically, *vertex coloring*). We discuss this graph-theoretic background further in Section 2.1.

Not all colorings of H are "safe," however, for the following reason: An individual i might still worry that someone in her group might start gossiping with someone in a different group, and thus might spread i's information beyond her group, and potentially to enemies of i's. Thus, we would like the boundaries of the information-sharing groups to be self-enforcing, in the sense that no small group of people could engage in cross-group gossip without opening their own information up to their enemies — a risk that deters such gossip purely through the self-interest of the parties involved.

We thus arrive at a model that is an interesting variant of graph coloring, in which the coloring problem (that of partitioning the nodes of the graph into disjoint groups such that no group contains nodes with an edge between them) is augmented with an additional constraint requiring that the partition into color classes be stable, or robust, to a certain kind of defection. In this framework, we study the properties of stable partitions, with particular attention to how the stability requirement limits the kinds of groups that can form and leads to greater fragmentation of the community. Our goal is to begin with essentially the cleanest model that captures the issues in the discussion thus far; however, most aspects of the model can be generalized, and we explore some of the consequences of these generalizations in Section 4.

2. Theory

2.1. Preliminaries

The following is a summary of concepts and terminology from graph theory and computer science that are key to the exposition of this work, but may be unfamiliar to some readers.[4]

An (undirected) *graph* consists of a set of *nodes* and a set of *edges* between pairs of nodes. A *path* between two nodes v and w in a graph is a sequence of nodes beginning at v and ending at w, with the property that each consecutive pair of nodes in the sequence is linked by an edge. A graph is *connected* if there exists a path between every two of its nodes. A *connected component* is a subset of a graph's nodes with the property that every two nodes in the subset are connected by a path.

The *graph coloring* (or, more specifically, *vertex coloring*) problem is that of assigning labels ("colors") to the nodes of a graph in such a way that any two nodes linked by an edge are assigned different labels. Such an assignment using at most c colors is called a (proper) c-coloring. The smallest number of colors needed to color a given graph H is called its *chromatic number*, denoted $\chi(H)$. Computing the chromatic number of a graph is known to belong to the well-studied class of computationally difficult problems referred to as *NP-hard* problems; see, e.g., Jensen and Toft (1994). A *bipartite graph* is one with chromatic number less than or equal to two; more generally, a c-partite graph is one with chromatic number at most c.

In graph theory, an *independent set* is a subset of the nodes such that no two are linked by an edge. Thus, a graph coloring is a partition of the nodes into independent sets. A *maximal* independent set is an independent set such that no additional node could be added without forcing the set to contain an edge. A *maximum* independent set is an independent set of maximum possible size for a given graph H.

The *set-cover problem* is as follows: given a universe of elements $\{1, 2, \ldots, n\}$ and a set system on it of size m (that is, any collection of m subsets of $\{1, 2, \ldots, n\}$ whose union comprises the universe), find the smallest number of those sets whose union contains all elements in the universe. This is a well-studied computationally difficult (NP-hard) problem. The *greedy algorithm* for the set-cover problem simply repeatedly chooses the set containing the largest number of uncovered elements; this algorithm is known to choose in the worst case $\ln s$ times as many sets as the optimum, where $s \leqslant n$ is the size of the largest set.

2.2. Formulating the information-sharing problem

We assume as input an undirected *conflict graph* H over the nodes (individuals) V with an edge between nodes i and j indicating that they are *enemies*. The "information-sharing groups" in the discussion above correspond to sets in a partition

[3] This group-based model is best suited for reasoning about the sharing of a single type of information in a small community — one can easily imagine different partitions of the community corresponding to the sharing of personal information about friends, personal information about family members, work-related gossip, and so forth. Indeed, such a structure of overlapping partitions is arguably a plausible way to think about the full spectrum of an individual's non-transactional information-sharing activities. For our purposes, however, we will focus on the partition induced by the sharing of a single kind of information.

[4] See any introductory texts on graph theory (e.g., Chartrand and Zhang, 2012) and theoretical computer science (e.g., Kleinberg and Tardos, 2005), for more discussion of these definitions.