On dealing with adversaries fairly

Andrei Serjantov and Ross Anderson*
Firstname.Lastname@cl.cam.ac.uk

May 6, 2004

Abstract

Peer-to-peer systems are often vulnerable to disruption by minorities. There are several strategies for dealing with this problem, but ultimately many of them come down to some kind of voting or collaborative filtering mechanism. Yet there exists a large literature on voting theory, also known as social choice theory. In this note we outline some of its key results and try to apply them to a number of recommender systems in the literature.

1 Introduction

The emergence of peer-to-peer systems has raised the question of dealing with adversaries in an open environment. Free speech in one jurisdiction may be prohibited speech in another [16]; music sharing may be seen as private by some users but as copyright infringement by others; and legitimate political protest for one group may be spam from the viewpoint of another. The openness of many peer-to-peer systems compounds the problem. If you have a system which anybody can join and anyone can leave at any time, how can you prevent attacks by small minorities? This has been a central issue in research on reputation systems, trust, and peer to peer system design in general.

One approach has been to design systems so that it is in everyone's best interest to behave in the intended way, and thus achieve the common goal. This area – mechanism design – has produced a number of interesting results [28, 19, 11]. However, mechanism design is neither always possible, nor practical.

In this paper, we investigate an alternative approach. We consider the behaviour of the system as being governed by the behaviour of its users, and let them specify preferences over the possible states of the system. For instance in a music sharing network we might decide which order the indexing servers should appear in, or simply determine music popularity.

To determine the global preference from the preferences of the users (preference aggregation), we use tools from voting theory, also known as *social choice*

*University of Cambridge Computer Laboratory, 15 JJ Thompson Avenue, Cambridge, UK

theory. This approach does mean that the system designer loses control over exactly how the system evolves as this is determined by the users. For example, if the system supports file sharing, and starts off sharing music, it may end up as a system for sharing pornography if this is what the users prefer!

Peer-to-peer systems are not the only possible application of voting theory to computer science, and a scattering of papers refer to its potential applicability – to collaborative filtering [23], service prioritisation in distributed systems [2], and voting on blacklists/whitelists for spam filtering [10]. There has also been some work on complexity; finding the winner of an election may be NP-complete [14]. No doubt there will be many more applications; for example, we can view software development models as a method of social choice, with the traditional closed-source model seen as a dictatorship while free and open source models allow users to specify their preferences with some limited decision making on whether these are accepted or not.

However, in this note we will concentrate on peer-to-peer systems. We will first give a brief introduction to voting theory, then review the preference aggregation schemes in the literature on social choice, look at their limitations and point out cases where they may be appropriate to problems in distributed systems. Then, as the principal use of voting in peer-to-peer systems is node reputation, we examine several reputation schemes from the voting-theory point of view. We point out their limitations and suggest improvements. Finally, we discuss how voting mechanisms can be applied in peer-to-peer content-sharing schemes.

2 Voting Theory – Fairness and Manipulation Resistance

There is a substantial literature on voting mechanisms. However, the computer science literature deals with different problems from the economics literature, and the experts appear mostly unaware of each others' work.

The computer science literature deals with mechanisms for casting simple yes/no votes using cryptographic mechanisms that protect voter anonymity, and may provide further properties such as receipt-freeness (the voter is unable to prove afterward which way she voted). There is also a literature on Byzantine agreement, which deals with mechanisms whereby a number of principals can resist the deception of a treacherous minority; and there has been significant work on shared-control mechanisms whereby, for example, any m out of n principals can perform some act [3]. (At this point, readers with an economics background can skip to the start of section 3.)

The economists' voting literature, also known as social choice theory, deals with the difficulties that arise when aggregating preferences, the theoretical limitations on aggregation methods, and practical methods to obtain the best trade-off in a given application. It is becoming clear that these results are also applicable in computing. In the rest of this section we will provide a brief introduction to social choice for computer scientists.

We will start off with three practical problems, which were already understood in the eighteenth century [30]. First, although simple majority-voting mechanisms can choose effectively between two outcomes, things get complicated once there are three or more outcomes, because we may have cyclic preferences. For example, Alice might prefer Labour over Liberal over Conservative, while Bob prefers Liberal over Conservative over Labour, and Charlie prefers Conservative over Labour over Liberal. Now, by a majority of two to one in each case, Labour is preferred to Liberal, Liberal is preferred to Conservative, yet Conservative is preferred to Labour.

The second (related) problem is tactical voting. This is familiar enough in recent history; in the US, many Democrats consider Nader's candidacy to be the reason George Bush won the 2000 presidential election, while in Europe the candidate for the French presidency who commanded the largest plurality of initial support did not even make the final run-off, because of a protest vote for a far-right candidate. It turns out that any voting method based on scoring (such as the *Borda rule* which simply gives three points to the first preference, two to the second and one to the third) is vulnerable to tactical voting.

The third is the cake-division paradox [26]. Suppose Alice, Bob and Charlie are voting to divide a cake. Alice and Bob can vote themselves half each, and give Charlie nothing. So while elections in which the participants are aggregating judgments (such as 'who was the greatest golfer of the nineteenth century') are hard enough, things get harder still when the participants are aggregating interests. Political elections do, of course, have a strong element of interest aggregation, as does the shared control of distributed systems.

This much was known to early researchers. During the run-up to the French Revolution, there was much discussion of voting; Borda proposed a scheme based on scoring, and his rival Condorcet came up with a circular preference example to show that tactical voting for a marginal candidate could lead to strange results. In particular, he introduced the idea of a *Condorcet winner* – a candidate who beats every other candidate in a pairwise vote – and showed that the Borda scoring system did not always pick the Condorcet winner.

Since then, many voting schemes have been proposed, and many odd features found. However, the modern study of voting theory really starts with Arrow's famous impossibility theorem, which tells us that where there are more than three voters and three choices, there is no 'best' voting system. Arrow's work also introduces a formal axiomatic basis to the field.

2.1 Arrow's Theorem

We now introduce a formal notation for preference aggregation, that is, deriving one global preference from those of many individuals. We represent preferences as reflexive, transitive and complete relations (also called orders) on a set X of states of the world. Preferences of an individual i are represented by the relation R_i, while global preferences are represented by the relation denoted R. If voter i prefers x to y, then $(x, y) \in R_i$. If i is indifferent between x and y, then

$(x, y) \in R_i \wedge (y, x) \in R_i$. If i strictly prefers x to y, then $(x, y) \in R_i \wedge (y, x) \notin R_i$. From relation R we easily derive relations I and P which express the indifference and strict preference respectively [1].

If each voter i has a reflexive, transitive and complete relation $R_i \subseteq X \times X$, expressing his preferences over the choices in X, then the function $f : \{R_i \subseteq X \times X\} \rightarrow R$ is a method of preference aggregation. If the range of such a function is the set of all orders, it is called a Social Welfare Function (SWF). What should the properties of such a function be?

- It is reasonable to expect that no individual should command global preferences. Hence, $\neg\exists i$ such that $\forall x, y \in X \ . \ xP_iy \Rightarrow xPy$. This is condition $\mathbb{D}$, known as *non-dictatorship.*

- We might require the domain of such a function to include any set of R_i's. This is condition $\mathbb{U}$, or *unrestricted domain.* Intuitively, this means that voters aren't constrained to combinations of choices that are consistent with some ideology; for any votes, we will get a meaningful outcome.

- If every individual prefers x to y, the global preference should reflect this. $(\forall i \ . \ xP_iy) \Rightarrow xPy$. This is condition $\mathbb{P}$, the *Pareto principle.* This is fairly weak as we do not insist that voting for a candidate should always improve that candidate's chances, merely that if everyone prefers a candidate then that candidate prevails.

- The global preference between x and y should depend on x and y *only* and not, say, on x vs z and y vs z. $R|_{\{x,y\}} = f_{\{x,y\}}(\{R_i|_{\{x,y\}}\})$. This condition $\mathbb{I}$ is called the *independence of irrelevant alternatives.* Intuitively, this is the requirement that adding extra candidates to an election should not affect the relative ranking of the candidates that were already there.

Unfortunately Arrow showed[2] that there is no SWF which satisfies the four axioms above where $|X| > 3$. He called this the General Possibility Theorem; it is at least as widely known as the Arrow Impossibility Theorem. It is in some sense a generalised form of the Condorcet paradox.

Nevertheless, there are arguments that his result is not quite as universally pessimistic as first appears. There are a number of ways in which we can try to deal with it. We can change some of the assumptions; we can introduce randomisation; we can iterate, and hope to arrive at an equilibrium after repeated elections [18]; we can redefine fairness; or we might settle for a choice mechanism that only functions most of the time. But no-one has found a way of relaxing the Arrow conditions that works for every application, and so the best election design will depend on our application.

[1] The readers who are unfamiliar with our notation are invited to refer to Appendix A.

[2] The original formulation is presented in [4], for a modern version, see e.g. [25]

2.2 Using majority rule anyway

Majority rule is the best election algorithm where there are only two candidates, but with three or more – as Condorcet pointed out – it can lead to circular preferences. Sometimes this won't happen, or won't matter.

We may avoid cyclic preferences where voters have relatively similar interests (see Sen, [25], Chapter 10.), or the domain is somewhat restricted. One possibility is where we have single-peaked preferences; another is where we have an ideological constraint. For example, if we are picking candidates on a left-right axis, we might assume that voters preferring a socialist candidate over a conservative one would rank a far-right candidate below either[3]. However, where not too many voters violate an ideological ordering, majority rule yields transitive preferences, and in this case it is the best election algorithm [7].

2.3 Voting with multiple winners or randomised winner choice

The result above is negative, but preference aggregation methods do exist. One of the early results is due to Sen [25], who shows that if the global preferences do not need to be complete, reflexive and transitive, but merely able to choose a best alternative from any subset (the Condorcet winner if there is one), then preference aggregation schemes satisfying all the conditions above exist[4]. Such a preference aggregation scheme is given by the simple rule that if at least one voter prefers A to B, and no voter strictly prefers B to A, then A is preferred globally. As Sen formalises it:

$$xRy \leftrightarrow \neg((\forall i.yR_ix) \wedge (\exists i.yP_ix))$$

In practice, such a scheme is likely to lead to a number of alternatives which are equally as good as each other. This method may be viable where it is sufficient to elect any one of a number of acceptable candidates, with the actual choice perhaps made randomly (one might think of electing a superpeer in a file-sharing system).

2.4 Maximum likelihood mechanisms

A very interesting class of preference aggregation schemes is possible if we can slightly weaken the assumption of independence of irrelevant alternatives. Recall that the Borda Rule, and scoring-based schemes in general, are vulnerable to manipulation by the introduction of new candidates; an electorate that prefers A to B may change its mind when offered also a no-hope candidate C. However, we

[3] This may not hold in real life: Jospin failed to make the final run-off in the 2002 French presidential election because a number of his supporters made a protest vote for Le Pen, assuming incorrectly that Le Pen would be knocked out in the first round so they could vote for Jospin in the run-off

[4] In the technical literature, a scheme that merely needs to pick one winner rather than the winner is called a voting scheme

can exclude this type of pathology by the weaker criterion of *local independence of irrelevant alternatives*. Here we only insist that independence hold within every interval of the proposed ordering: so that the ranking of A and B is unaffected by the introduction of a weaker candidate C, or even a stronger candidate D.

There is then a (unique reasonable) ranking rule that will work, namely maximum likelihood voting. Here we assume that there is a correct answer, and that the voters' rankings are noisy approximations to it. In other words, given a choice between two alternatives, voters can pick the correct one with some probability > 0.5. We choose the ranking R which has the greatest likelihood given the observed voting outcome. This can be computed fairly efficiently by constructing a *vote graph* whose nodes are candidates and whose edge ij has weight equal to the number of candidates preferring i to j. It suffices to find the maximum-weight set of edges that does not contain a cycle [30].

Maximum likelihood is another strong candidate for preference aggregation in distributed systems. For a concrete example, imagine users ranking the speed of indexing servers on a file-sharing network. Some are objectively faster than others, but the view of each user may be obscured by local congestion. In this case, the method of maximum likelihood could be applied.

2.5 The Kemeny approach

Another probabilistic approach was pioneered by John Kemeny [15]. Assuming again that voters each submit a noisy approximation to a true ordering, we can create a metric between these submissions based on edit distance. We can then find a compromise ordering between these data points. A mean ranking minimises the sum of the squares of the distances between the compromise and the data points, while a median ranking minimises the sum of distances. The median ranking turns out to be less vulnerable to manipulation, as a malicious voter can get less leverage by exaggerating his preference [30].

A Kemeny-optimal compromise ordering will select the Condorcet winner where these is one, but finding it is NP-hard. A recent result is *local Kemeny optimality*, which has been proposed for aggregating spam whitelists and blacklists, and is efficient [10].

2.6 Market-based mechanisms

Market mechanisms provide another means of collective choice. For example, one might issue each voter with 1000 tokens of notional currency and let them particupate in a policy auction. This is the idea behind a proposal of Tideman and Tullock [27], in which a Vickrey-Clark-Groves (VCG) auction is then used to ensure that voters have no incentive to lie about their preferences. The effect that that if people offer more money for an alternative than is necessary to secure it, people who offer more than the winning margin have to pay a tax equal to their contribution to the victory.

There has been recent interest in using VCG auctions in computing applications such as finding optimal routes. Advances in algorithmic mechanism design by Nisan and Ronen opened up the possibility of distributing algorithms among participants who cannot be assumed to follow the algorithm but rather their own self-interest. So the algorithm designer must ensure that agents' interests are best served by behaving correctly. Feigenbaum, Papadimitriou, Sami and Shenker applied this to a VCG auction, used to find the bext routes in a network[11]. There appears to be scope for extending these mechanisms to the problems of interest here (though the current generation of algorithms is still rather slow).

2.7 Manipulability

The above schemes give us a number of options for distributed system applications. In the context of peer-to-peer systems, we are particularly concerned with attacks where a coordinated minority of nodes manipulates the choice mechanism so as to frustrate the will of the majority. So what can the existing social choice literature tell us about manipulation?

A fundamental theorem about voting schemes was established independently by Gibbard and Satterthwaite [13, 24]. In its axiomatic form, it states that preference aggregation schemes with three or more candidates, which give a single answer and which do not use randomness, are either manipulable or dictatorial. Gardenfors extended the result to show that the ordering can be manipulated, as well as the winner [12]. However, one can escape Gibbard-Satterthwaite through the use of randomness; while this is frowned upon in the more traditional applications of voting (e.g. choosing a president), it may be quite appropriate in the case of distributed systems. We might merely need to choose one superpeer out of a number of acceptable candidates.

There are further interesting results on manipulability. In some cases the outcome is still the same [22], as each side can manipulate and end up with the same outcome. In others the problem of manipulating the vote is NP-complete [5]. A survey of manipulability by Nurmi classified election schemes into four levels of manipulation-resistance, depending on the information a manipulator needed in order to interfere purposefully: voters' first-place choices, their approvals, their pairwise comparisons or their orderings [20]. Results such as these enable us to compare the manipulability of different systems.

However, in peer-to-peer applications we are dealing with different kinds of interference than in political elections. In an election, all voters may try to vote strategically, given their beliefs about the intentions of other voters. In a peer-to-peer system, it will typically be a minority – often a colluding minority – that attempts to disrupt the service enjoyed by the majority. This minority might be dishonest merchants trying to swindle eBay customers, or RIAA bots trying to disrupt a music-sharing system.

That said, we will now consider some of the reputation systems that are used in real life, or that have been proposed in the literature.

3 Unfairness in Reputation Systems

We will now consider the preference aggregation performed by reputation systems in online envornments. Not all reputaiton systems involve social choice; we are only interested in those which compute some kind of global ordering out of the individual orderings. First of all we note that in this case the exercise is that of *interest aggregation*; every participant wants to have a higher reputation. As noted above, this is a harder problem than simple judgement aggregation. It is worth examining some of the implemented and proposed solutions from the social choice point of view.

3.1 Immunizing reputation systems

In [8], Dellarocas proposed using cluster filtering to identify and eliminate preferences which look too far away from everyone else's value, and adding the remaining preferences up. His reputation scheme involves adding up ratings (which range from 0 to 100) from different users, and the purpose of filtering was to exclude small minorities of disruptive voters who deliberately gave excessively large or small scores.

This would have been less innovative from the viewpoint of social choice. Aggregating preferences while giving less weight to outlying views has been well studied in this context; see for example the algorithms inspired by Kemeny in section 2.5 above.

Indeed, given the social choice toolkit, we can improve on Dellarocas' approach. His decision to simply ignore outliers is suboptimal on several counts. First, it violates the basic idea that "everyone's preferences count"; second, interpersonal comparisons of the kind he proposes have been long frowned upon; and third, users are unlikely to be able to express the relative strength of their preferences consistently on such a wide scale. We suggest that Kemeny, local-Kemeny, maximum-likelihood, or simple median would be an improvement.

3.2 Reputation system metrics

Next, let us look at the voting reputation system proposed by Marti and Garcia-Molina in [17]. In this reputation scheme, there is no global reputation value for a node; instead if a node q wants to compute the reputation of node r, it asks a set of nodes (assume for simplicity that this is all the nodes in the network) for their opinions of r. Their opinions are then weighted by what q thinks of each of the respective nodes. Formally if V is the set of nodes in the system,

$$p'(r) = \frac{\sum_{v \in V} R(q,v)R(v,r)}{\sum_{v \in V} R(q,v)}$$

This is the "preference aggregation" part of the reputation which is of interest to us here, rather than the whole defintion.

In effect, a database of votes is available, and each user then chooses his own weighting to construct a welfare function of his own. The rationale is the following: a user doing transactions online with experience of the success or failure of previous transactions with particular counterparties may want to discriminate against them based on this prior experience. The multiplicity of possible choice functions does not, of course, make the problems go away; inexperienced users (at least) will presumably use equal weights, which turns the scheme into a simple scoring system. Such users can be deceived in the usual ways.

Even once a user has experience of the dependability of other users, social choice theory does have something useful to say. In effect, one is using a maximum likelihood rule, and in the case that we have n voters each of whom is likely to speak the truth about the state of the world with probability p, the weight that should be given to their view is $\log(p_i/(1-p_i))$ [30].

3.3 Trust in a Peer-to-Peer Information System

Our third example is a proposal by Despotovic [1]. Here, P is the set of all nodes in the network and $c(p,q)$ is defined as the number of complaints filed by p against q. The reputation of p is then defined as:

$$T(p) = \sum_{q \in P} c(p,q) \times \sum_{q \in P} c(q,p)$$

High values of $T(p)$ indicate that p is not trustworthy.

Multiplicative utilities are not unknown in the choice literature (see for example [6]). Used like this, though, they give rise to an immediate problem: if I do not complain, then no complaints against me (however justified) are registered. So there is a Nash equilibrium at "no complain, no complain", we do not go into the details here. This can be fixed in various ways, such as replacing multiplication by addition, or by adding 1 before multiplying.

Assuming such a bug-fix, let us evaluate this aggregation scheme against the conditions we considered above. We define the relation R_i as follows:

$$(1+c(i,x)) \times (1+c(x,i)) \leq (1+c(i,y)) \times (1+c(y,i)) \Rightarrow xR_iy$$

Then, from before,

$$\sum_i (1+c(i,x)) \times \sum_i (1+c(x,i)) \leq \sum_i (1+c(i,y)) \times \sum_i (1+c(y,i)) \Rightarrow xRy$$

- The above definition does not satisfy Pareto. Consider the following exmaple: There are four nodes in the network: $\{p,q,r,s\}$. $c(p,s) = c(s,p) = c(p,r) = c(r,p) = 1$, $c(s,q) = 0$, $c(q,s) = 2$, $c(r,q) = 2$, $c(q,r) = 0$. Hence, qP_sp, qP_rp, and yet pIq which violates $(\forall ixP_iy) \Rightarrow xPy$[5].

[5] Note that we have excluded p and q from the calculations which determine their relative reputations.

- This preference aggregation scheme satisfies $\mathbb{I}$, independence of irrelevant alternatives.

The scheme does satisfy nondictatorship and unrestricted domain; but can we do better? Suppose we go instead for the obvious ordering on tuples:

$$T(x) = \left(\sum_i c(i,x), \sum_i c(x,i) \right)$$

and define: $(a,b)R(c,d) \Leftrightarrow (a \leq c) \wedge (b \leq d)$. We will find that this relation is not a total order (and hence the funciton is no longer an SWF). However, in all the cases where we start off with an ordering, this is majority voting (on two different parameters) which satisfies Pareto ($\mathbb{P}$), Independence of Irrelevant Alternatives ($\mathbb{I}$) and Non-dictatorship ($\mathbb{D}$), but not Unrestricted Domain ($\mathbb{U}$). In effect, by trying out all the possible peer-to-peer rating systems, we just end up reinventing social choice theory. This is clearly suboptimal.

3.4 Real Reputation Systems

Where computer science can bring some novel experience to the social choice debate may lie in the experience of online rating systems developed by sites such as Amazon and eBay.

Ebay lets buyers and sellers at its auction site rate each other according to whether they paid promptly, whether the goods were shipped on time, whether they were as described, and so on. There is now a body of research on the effect of this ratings system (and of other similar ones); Resnick finds, for example, that despite obvious incentives to free-ride, a rating is provided more than half the time, that it is almost always positive, that rating scores are predictive of future performance, that buyers and sellers reciprocate and retaliate, and that good ratings did not let sellers boost their prices [21].

Amazon displys the average rating of a book based on scores given by customers who reviewed it. Social choice theory tells us that scoring systems are open to manipulation and indeed it is so: after a faculty colleague of ours stopped supplying course notes to students and pushed them to buy her book instead, the students retaliated by flooding Amazon with poor reviews and comments such as recommending the book as a cure for insomnia. Its rating plummeted.

4 Preference Aggregation in P2P Systems

This brief look at the social choice literature and its application to distributed system design leads us to make several observations.

First, we notice that unlike in elections, committee meetings, or other traditional exercises of social choice, in distributed systems the election is likely to run many times, rather than just for a single round. This can make the social choice problem much easier. There is a large literature on multi-stage elections, and a much smaller one on iterative elections that converge to an equlibrium. In a peer-to-peer system, a wrong choice of a supernode would degrade the performance noticeably, leading to a re-run of the election.

Another phenomenon of peer to peer systems is that the attacker is likely to have many nodes colluding with each other [9], while the honest users are unlikely to employ any sort of strategic voting (they just run the software). Moreover, the attacker is likely to know what the preferences of the users are going to be, and hence will have almost complete information unless we deliberately randomise honest-user behaviour.

On a positive note, if the problem in hand is one of estimating a true fact about the network (such as the speed of a particular server), then the maximum likelihood method can be used to give optimal results, or we can use the median ranking to give less weight to the outliers. Both have their own merits, but one has to think about the problem carefully – is there indeed a true fact about the world out there to be discovered?

Yet another question which can be simply resolved is "what is the best of two options?". For instance, should we admit new users today or should we try to optimise network topology? Majority rule will give the best answer.

Finally, there are a host of issues relating to computational effort and to locality. This may swing our choice in favour of methods such as approval voting [29] for situation where we want to minimize the effort made by the users, and possibly distribute it in space and/or in time.

5 Future Work

On the practical side, social choice rules give us valuable insights into design choices arising in peer-to-peer systems, such as ways to choose the supernodes in a network. Next-generation peer-to-peer systems could be a useful testbed to explore how the theory works in a new environment. From the theoretical point of view, building a model of a peer-to-peer system, and then working out which voting schemes (or reputation systems) yield acceptable properties is an interesting, and potentially worthwhile exercise.

So far, most reputation schemes have been evaluated merely through simulations (e.g. [17, 8]). Evaluating them against the existing literature on voting and social choice should provide interesting insights in this field too. Ad-hoc networks should also provide interesting problems – such as how to deal with highly mobile nodes, nodes that are intermittently off-line, and nodes subject to occasional, transient subversion. Perhaps this will provide stimulating new problems for the existing social choice community.

6 Conclusions

Until now, computer scientists and economists have developed separate literatures on elections and voting. In this paper we have argued the case for bridging the gap. We have introduced the social choice tools which economists have developed over the last 50 years, and shown that they are immediately applicable to practical problems in the design of reputation systems and peer-to-peer networks. We have been able to identify problems, and suggest improvements, with existing designs.

The broader conclusions of this paper are twofold. Firstly, the domain of distributed system design (and peer to peer systems in particular) has many potential applications of social choice tools. Secondly, the engineering problems here are quite different to the traditional social choice subject matter. This provides opportunities for new work, such as in designing schemes which provide statistical properties over multiple rounds. Perhaps the medium-term goal is a new way of designing distributed systems with evolving behaviour.

References

[1] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM)*. 2001.

[2] R. Anderson. *Invted Talk*, PODC 2003.

[3] R. Anderson. *Security Engineering: A guide to building dependable distributed systems*. Wiley, 2001. ISBN 0-471-38922-6.

[4] K. Arrow. *Social Choice and Individual Values*. Wiley, 1951.

[5] J. J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[6] L. F. Cranor. *Declared-strategy voting: an instrument for group decision-making*. Ph.D. thesis, Washington University, 1996.

[7] P. Dasgupta and E. Maskin. Is majority rule the best voting method? Available at `www.econ.cam.ac.uk/faculty/dasgupta/MajRuVot.pdf`.

[8] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behaviour. In *ACM Conference on Electronic Commerce*. ACM, 2003. ISBN 1-58113-679-X.

[9] J. R. Douceur. The sybil attack. In *First International Workshop on Peer-to-Peer Systems (IPTPS '02)*. Cambridge, MA, March 2002.

[10] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation, spam resistance and social choice. `theory.stanford.edu/muri/reports/1999-2000/cynthia2.ps`.

[11] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A bgp-based mechanism for lowest-cost routing. In *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, pages 173–182. ACM Press, 2002. ISBN 1-58113-485-1.

[12] P. Gardenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13:217–228, 1976.

[13] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601, 1973.

[14] E. Hemaspaandra and L. Hemaspaandra. Computational politics: Electoral systems. In *MFCS 2000*, pages 64–83. Springer LNCS 1893.

[15] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:571–591, 1959.

[16] L. Lessig and P. Resnick. Zoning speech on the internet: A legal and social model. *Michigan Law Review*, 98.2:396–431.

[17] S. Marti and H. Garcia-Molina. Examining metrics for peer-to-peer reputation systems. Technical Report 2003-39, Stanford University, July 2003.

[18] R. B. Myerson and R. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.

[19] N. Nisan and A. Ronen. Algorithmic mechanism design. In *ACM Symposium on Theory of Computing (STOC)*. 1999.

[20] H. Nurmi. *Comparing Voting Systems*. D Reidel Publishing Company, Theory and Decision Library, Dordrecht, 1987.

[21] R. Z. P. Resnick. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In *NBER workshop*. 2001. Draft.

[22] B. Peleg. Consistent voting systems. *Econometrica*, 46(1):153–161, 1978.

[23] D. Pennock, E. Horvitz, and C. Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *17th National Conference on Artificial Intelligence*, pages 729–734. 2000.

[24] M. Satterthwaite. Strategy-proofness and arrow's condition: existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[25] A. Sen. *Collective Choice and Social Welfare*. Holden-Day and Oliver and Boyd, 1970.

[26] A. Sen. Social choice theory: A re-examination. *Econometrica*, 45, 1977.

[27] T. Tideman and G. Tullock. A new and superior process for making social choices. *Journal of Political Economy*, 84:1145–1159, December 1976.

[28] H. Varian. Economic mechanism design for computerized agents. In *Usenix Workshop on Electronic Commerce*. 1995.

[29] R. J. Weber. Approval voting. *The Journal of Economic Perspectives*, 9(1):39–49, 1995.

[30] P. Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995. Available from `http://uk.jstor.org/view/08953309/di980582/98p02187/0`.

A Notation

- xPy is defined as $(xRy) \wedge \neg(yRx)$.
- xIy is defined as $(xRy) \wedge (yRx)$.
- A relation $R \subseteq X \times X$ is said to be *reflexive* if $\forall x \in X.(x, x) \in R$.
- A relation $R \subseteq X \times X$ is said to be *transitive* if $(x, y) \in R \wedge (y, z) \in R \Rightarrow (x, z) \in R$.
- A relation $R \subseteq X \times X$ is said to be *complete* if $\forall x, y \in X.(x, y) \in R \vee (y, x) \in R$.
- The relation $R|_{\{x,y\}}$ denotes the part of relation R which involves x and y. Formally, $R|_X = \{(x, y)|(x, y) \in R \wedge x \in X \wedge y \in X\}$.