contributed articles

DOI:10.1145/1953122.1953154

Decision-making procedures in online social networks should reflect participants' political influence within the network.

BY PAOLO BOLDI, FRANCESCO BONCHI, CARLOS CASTILLO, AND SEBASTIANO VIGNA

Viscous Democracy for Social Networks

ON APRIL 23, 2009, Facebook announced preliminary results of a vote in which users were asked about a change in the terms of use of the network.[a] The vote was the result of a petition by thousands of users criticizing the site for claiming too many rights over user-generated content. Attempting to justify the change, Facebook let users vote, saying that if 30% of the then roughly 200 million "active" users would vote the decision would be binding on Facebook management.[b] The outcome was 74.4% of the voters preferred the new rules. And while only 600,000 users voted (1/100th of the prefixed quorum of 60,000,000) the change in the terms of use were officially adopted.[c] The global privacy watchdog, Privacy International, called the vote a "massive confidence trick"[d]. The low voting turnout was largely foreseeable. Only a small fraction of Facebook users have the time, patience, or dedication or take the service seriously enough to actively participate in its governance.

Whether to decide on a motion (such as pick one option among a set of alternatives) or elect representatives (such as constitute a senate), voting is a collaborative decision-making process, seeking a result that reflects as much as possible the opinion of the community as a whole.

Viewed this way, the failure of the 2009 Facebook voting experiment is explained by the kind of voting system[2,14] Facebook adopted. Attempting this kind of direct democracy voting in large online communities is not necessarily the best approach; when public decisions reach a certain level of complexity, it is unrealistic to assume every participant is engaged and informed enough to contribute to the decision.[8,11]

This opinion is shared by other authors[5,16] observing that the degree of commitment of participants in online communities and collaboration systems varies greatly. Nielsen's "90-9-1 Rule for Social Design"[16] says: "In most

a http://www.facebook.com/fbsitegovernance

b An account is defined as active if it had some activity in the past 30 days.

c http://www.cio.com/article/490775/After_Vote_Facebook_Gets_New_Governing_Rules

d http://www.privacyinternational.org/article.shtml?cmd[347]=x-347-564312

» key insights

- In an online community where users typically have very different levels of engagement and activity, direct democracy may not be the best approach to decision making.
- The proposed voting system allows people to delegate their decision-making power to people they trust in their social networks; the decision weight is allocated proportionally to the people who trust you—plus the people who trust in the people who trust you.
- The insights behind the proposed method stem from a 60-year research line on ranking in social sciences that more recently found application in Web search with the PageRank index.

online communities, 90% of users are lurkers who never contribute, 9% of users contribute a little, and 1% of users account for almost all the action."

This may sound unfair but is central to the way open collaboration networks work; for example, in the case of Wikipedia, Shirky[19] wrote: "Fewer than 2% of Wikipedia users ever contribute, yet that is enough to create profound value for millions of users… among those contributors, no effort is made to even out their contribution. The spontaneous division of labor driving Wikipedia wouldn't be possible if there were concern for reducing inequality. On the contrary, most large social experiments are engines for harnessing inequality rather than limiting it."

Indeed, in well-designed collaborative systems, even those participants who show up only once can contribute positively toward achieving community goals. However, these "drive-by" participants do not (and should not) have unreasonable expectations about their weight in group decisions. For instance, in the case of the popular open source distribution of Ubuntu, Leadbeater[15] wrote: "Decision-making is very open: Anyone can see what is decided and how; anyone can make suggestions about what should be done. But the way decisions are made is rarely democratic."

In a community with only a few core members with long-term commitment and many others regularly joining and leaving, egalitarian democracy is neither expected nor appropriate. Thus, the decision-making mechanism is often meritocratic. These considerations suggest that different forms of voting systems should be considered for peculiar communities like electronically mediated social networks.

Direct, Representative, and Liquid Democracy

Direct democracy is based on the idea of maximum equality and fairness by making all constituents vote directly

PHOTOGRAPH BY FLICKR USER CALAFELLVALO

for the various motions. Direct democracy works best in practice for small cohesive groups, but when decisions are highly complex, and the community is large, it becomes impractical for every citizen to be fully informed on every issue. Furthermore, direct democracy requires deliberation to work effectively, and deliberation is relatively more difficult to achieve through electronic communications than through direct face-to-face contact. Electronic communication reduces the set of modalities by which group members are able to communicate, influencing their performance, particularly when they are new to the technology being used.[12]

Representative democracy involves a relatively small number of representatives elected by their constituents to take decisions on their behalf about many different matters over a relatively long period of time. Beyond the question of which representation structure is most appropriate for a given context, representative democracy is only weakly democratic; though citizens may participate in elections, they do not really choose their representatives in the strict sense, choosing instead among a restricted set of candidates whose views and values are often radically different from their own and who tend to make fundamentally different choices when faced with controversial social decisions. As a result, voters' apathy is more common than political interest. In 2007, James Green-Armytage, an economics Ph.D. candidate at the University of California, Santa Barbara, wrote[11]: "In traditional representation systems, voters' positions on hundreds of social issues must be reduced to choices between candidates or parties, resulting in massive information loss." Voter apathy, combined with the concentration of power in the hands of a small political elite, creates fertile ground for corruption, entrenchment, and conflicts of interest, potentially resulting in bad government.[18]

Driven by the appeal of meritocracy, some online communities have implemented decision systems midway between the universal voting of direct democracy and representative democracy. In many online communities, the right to vote is given to a subset of the community, not selected by citizens through another vote but on the basis of commitment. For example, while the online community of editors at Wikipedia explicitly discourages voting to resolve editorial disputes[e], voting may take place for changes affecting the entire Wikipedia; an example was the April 2009 vote[f] on whether to adopt a Creative Commons license when not all Wikipedia editors were allowed to vote, only those who had contributed 25 edits or more by a certain date.

While such a solution implemented a meritocratic decision system, it did not consider the specific nature of large online social networks or their structure, properties[g], or high conductance for viral phenomena. Rather than select which people have the right to vote by a rigid threshold on commitment and activity (not necessarily implying someone is trustworthy), it would be more appropriate to adopt a fluid system based on people's trust.

The idea of fluidity leads us to consider another form of voting system—delegative democracy[h] or liquid democracy[i] based on transitive proxy voting.[8,11,23] Under it, constituents either express their opinion directly on an issue or delegate their vote on the issue (or multiple related issues) to a proxy who is another citizen they trust. If the proxy votes directly on the issue, the weight of all delegated votes the proxy received are added to the proxy's vote. Proxy delegation may be transitive, with one's vote further delegated to the proxy's proxy.

Proxy-voting systems encourage participants to cooperate to build direct, permanent political and social relationships with one another and with individual supporters, forming a web of trust. Participants in such a system can achieve political influence proportional to their level of public support in social networks usually related to their connectedness.

From Liquid to Viscous

Liquid democracy is based on local, personal acquaintance, trust, social relationships, and the principle of transitive delegation. The collaborative decision-making process in liquid democracy is a social cascade well suited to online social networks. We therefore take it as a starting point for developing our proposal for how to vote in social networks. Here, we present how it can

Figure 1. Example delegation graph.

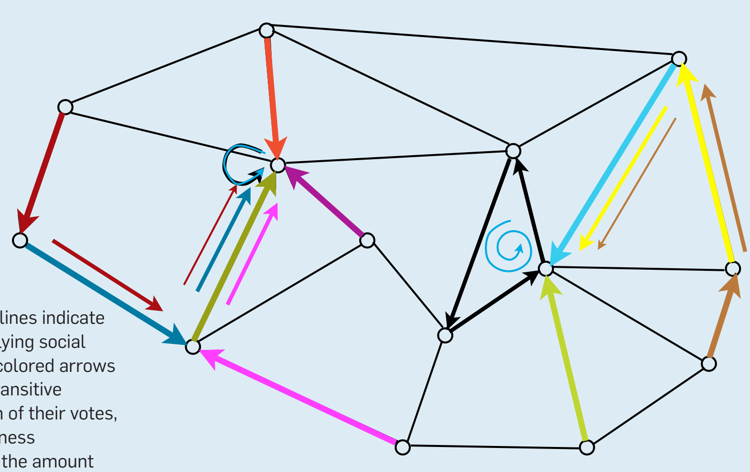

Thin gray lines indicate the underlying social network; colored arrows indicate transitive delegation of their votes, with thickness indicating the amount of "voting weight" actually transferred.

e http://en.wikipedia.org/wiki/Wikipedia:VOTE

f http://meta.wikimedia.org/wiki/Licensing_update/Result

g Social networks are characterized by several properties: a heavy tail in the distribution of the degrees of their nodes, or number of edges incident in a node; a small diameter, or maximum-possible distance between two nodes measured as length of the shortest path connecting them[1]; a small-world structure[21]; and a distinctive community organization[9]; see also Willinger et al.[22]

h The idea of delegate voting is not recent; for example, an early proposal, dated 1884, was due to Lewis Carroll.[6]

i http://democracialiquida.es/

be used for both deciding on concrete matters (such as between two alternative motions by computing the weighted count of votes received by each motion) and selecting a committee of representatives to deal with a particular group of issues during a particular period of time relatively shorter than that of representative democracy. For the time being, assume we are dealing with how to decide an issue; we later discuss how to elect a committee.

As in liquid democracy, the key aspect of the proposed voting system is that votes can be delegated transitively along the existing links of the social network. That is, members of the network can choose a proxy among their contacts. Alternatively, citizens can also choose not to delegate their votes and instead express their opinion on the matter of direct voting.

Besides the obvious organizational advantages, the constraint that votes can be assigned only to a direct connection has a twofold rationale: Voters can base their decision on direct personal knowledge of the person they might vote for, making direct propaganda essentially useless and thus decoupling popularity from credibility; and attributing mandates through a chain of direct connections should ensure a stronger sense of responsibility.

However, we cannot ignore the fact that personal ties in online networks are not as strong or direct as those in real-world communities. Social connections are a mixture of strong ties (family, close friends) and weak ties (distant friends, acquaintances),[10] and electronically mediated networks allow people to maintain many more weak connections.[7] Having multiple weak connections means the number they have is larger than one could consider an actual "friendship" network and members' trust in their connections is weaker on average. For this reason, it seems appropriate to introduce some reluctance in the delegation process to reduce the amount of transitivity, as we now explore in detail.

The ballot and the tally. To describe our "Viscous Democracy" voting system, we first specify how voters express their preference (sometimes called the ballot) and the algorithm that determines the final outcome (the tally). Note that related technical issues concern how the voting is carried out (such as establishing participant identity[j]), an important topic solved through various means but that we do not address here.

In a community with only a few core members with long-term commitment to the project and many others regularly joining and leaving, egalitarian democracy is neither expected nor appropriate.

The ballot. The ballot can be defined in various ways. For example, "one-vote" voting systems involve a voter picking exactly one candidate, or, in our case, one contact. In a "ranked" voting system, individual voters rank their contacts in order of preference. In a "rated" voting system, voters give a score to each contact.

For the rest of the article, we consider the simpler one-vote kind of ballot, where participants choose to delegate their decisions to exactly one of their contacts or vote for themselves, corresponding to not delegating the vote further. This ballot can be interpreted as a delegation graph or a directed graph built over the undirected underlying friendship social graph. It can contain cycles and self-loops representing the choice of some electors not to delegate their vote, instead directly expressing their opinions on the matter of voting; Figure 1 is an example of a delegation graph induced by a hypothetical vote over a social network.

The tally. Our system considers that each person in the network receives a certain amount of score (weight) used to decide among alternative motions or elect a committee, though the way scores are used is not part of the voting system[k] discussed here.

Scores from a delegation graph are computed in many ways, a trivial one being the sum of all votes received. Here, we propose a more complex tally—transitive proxy voting with exponential damping—similar to standard proxy voting of liquid democracy but with a damping factor that introduces some reluctance in the way delegated

j A popular anecdote concerns "Hank the Angry Drunken Dwarf" winning the 1998 *People* magazine online poll to identify the most beautiful people in the world; http://en.wikipedia.org/wiki/Hank_the_Angry_Drunken_Dwarf

k When making a decision, one may count only the votes from people who decided not to delegate their votes, making them choose among possible alternatives, and weighting their choices proportionally to the scores obtained. If all voters delegate their votes, then one may count only voters belonging to "delegation cycles," as they retain part of their voting power, albeit in a weaker sense.

A Brief History of Spectral Ranking: From the 1940s to PageRank

The renowned PageRank index, the basis of the initial success of Google's ranking algorithm, has been rediscovered over and over for the past 60 years.

Spectral techniques for computing "best" entities when some relationship between them is known date to at least the late 1940s when John R. Seeley published his study on ranking children ("The Net of Reciprocal Influence: A problem in Treating Sociometric Data" in *Canadian Journal of Psychology*, 1949), given a matrix with information about whether a particular child likes another child. Seeley said a child's rank should be given according the sum of the ranks of the children who like him or her. Seeley imposed this requirement through a linear system; in contemporary terms, he computed the left dominant eigenvector of the normalized matrix to establish which children were most popular.

In 1952, T.H. Wei in his unpublished Ph.D. thesis (*The Algebraic Foundations of Ranking Theory*, University of Cambridge, 1952) discussed how to rank football teams, given a matrix representing how much each team is better than another team (such as 1 for a win, 1/2 for a tie, and 0 for a loss). He said that given an equal starting score for each team, we get a more precise score by adding for each team the scores of the team it defeated, plus half the score of the team it tied, showing that iterating this process would produce the right dominant eigenvector and using it to identify the best teams.

At the time, spectral ranking—using eigenvectors to compute ranks—was an established idea. In the 1950s, Leo Katz introduced its index ("A New Status Index Derived from Sociometric Analysis" in *Psychometrika*, 1953). Katz began with the notion that given a zero/one matrix expressing whether each person in a group "likes" or "votes for" another person (implying endorsement), we should estimate the importance of a person not only from the number of his or her voters but also from the number of his voters' voters, and so on ad infinitum. To obtain a finite value, Katz suggested using an attenuation factor α to reduce the weight of longer and longer voting paths. When the attenuation factor of Katz's index approaches the reciprocal of the dominant eigenvalue of the matrix, the result (as a limit) is the standard spectral ranking of Seeley and Wei, suggesting the name "damped spectral ranking" for the former.[4]

PageRank combined the dominant eigenvector ideas of Seeley and Wei with Katz's approach (damping). Again, the motivation was different and related to the random surfer model, with PageRank modeling the behavior of a surfer moving randomly through links and, with probability α, jumping to a random node. Seeley's ranking is PageRank without jumps to random nodes, whereas Katz's index is PageRank without the normalization (divide by the number of linked nodes) usually applied to each matrix row.

For more mathematical detail and historical context, see Vigna.[20]

votes are transferred. This reluctance, controlled by a parameter α, corresponds intuitively to the idea that, in an electronically mediated social network, participants typically cannot fully trust their connections and want to refrain from giving them all their delegation. More important, we do not know how far our liquid vote might go hop-by-hop on the network. Even if we fully trust our proxy, can we transitively fully trust our proxy's proxies?

Reluctance makes the vote less liquid, reducing its strength with each delegation step, thus limiting the distance it travels. Reluctance makes the vote viscous. We might call this form of proxy voting "viscous democracy" due to the way trust (and consequently a vote's weight) decays with distance.

Viscous Democracy and Spectral Ranking

The computation we suggest is known to sociologists as Katz's index[13]; every vote transfers by transitivity to distances larger than one, but with an attenuation factor. The delegation graph has out-degree one (because a one-vote ballot is used), making our case much easier to analyze. The score of node i is simply proportional to:

$$\sum_{p \in \mathrm{Path}(-,i)} \alpha^{|p|},$$

where Path(–,i) is the set of all delegation paths ending at node i and $|p|$ is the length of one such path.[l] Computing Katz's index on the delegation graph is completely different from computing Katz's index on the social network (its standard application).

Techniques like Katz's index, or the so-called spectral ranking methods, have been used by psychologists, sociologists, and management theorists for at least the past 60 years to estimate authority, power, influence, and centrality. The most popular incarnation of this idea today is probably Google's PageRank[17]; see the sidebar "A Brief History of Spectral Ranking: From the 1940s to PageRank." Our proposal adds a new flavor to an old ingredient, showing how decades-old techniques can be applied to voting in online social networks[m].

The Parameter α

A single parameter $\alpha \in]0, 1[$ controlling the voting process can be understood as the delegation factor, or the amount of one's own power a person can delegate to another person; that is, $1 - \alpha$ is the amount of viscosity in the system.

If the delegation factor is small (close to 0), mandates become undelegable, meaning that if a person receives enough delegations (votes) from other nodes, that person alone cannot make a third party "more powerful" than herself, even if she votes for him/her. In the limit, when the delegation factor tends to zero, only direct votes count, and the resulting process is essentially a simple majority vote.

If the delegation factor is large (close to 1), most nodes delegating their mandate to someone else will not have high scores. In the limit, when the delegation factor tends to one, the system becomes liquid democracy, and the winners are chosen simply by the size of the sub-tree to which they

l Given the possible presence of cycles in the delegation graph, some of these paths may be arbitrarily long; however, the sum converges for $|\alpha| < 1$.

m Yamakawa et al.[23] used a weighted form of Katz's index but on a matrix representing voters and motions. Each voter was able to cast a positively weighted vote to one or more voters and motions; the sum of the outgoing weights was constrained by design as one. No underlying social network was considered in their voting system.

belong, or the number of people who voted for them, directly or indirectly.

Figure 2 outlines an example of viscous democracy showing how the delegation factor determines the degree of viscosity, hence when $\alpha = 0.2$ (left), the system is more viscous, and node 5, with many more direct supporters than the others, has the highest score. Meanwhile, when $\alpha = 0.9$ (right), the higher degree of transitivity makes 7 the node with the highest score. Also, in this case, node 8 is slightly stronger than node 9 (contrary to what happens for smaller values of α), because, though 8 has fewer direct supporters, she receives part of the influence of node 7.

Finally, two more properties of the system should be mentioned: Given the right delegation graph, everyone could be a winner (provided the social network is connected), so there are no a-priori losers in the social network, while some nodes have a better chance than others. Moreover, there is continuity in the decision with respect to α, so if one node obtains a higher or lower score than another with appropriate choices of α, there is always another choice that makes the two nodes have the same score. For a more complete list of features, as well as a detailed comparison with PageRank, see Boldi et al.[3]

The Social Network and Abstentionism

The system raises two interesting questions: To what extent does the structure of the social network determine the outcome of the election? and How should it deal with missing votes? The answers are closely related; indeed, when an individual abstains from voting, the system can use the local structure of the network around the individual to learn how her vote might have influenced the final outcome.

The results of an election in viscous democracy depend on the delegation factor α and on the delegation graph. The delegation graph is in turn constrained by the underlying social network, as each voter votes for only one neighbor. Earlier, we said that anybody can in theory win an election, given the right delegation graph. However, this result is optimistic, because in practice the likelihood that a node on the fringe of the network will win an election is small.

Our system deals with the problem of missing votes naturally, considering all possible what-if scenarios; specifically, if a member with k contacts does not vote, her vote may have been cast in k different ways, with each producing a certain score of the nodes. Lacking further information, the system considers these scores equally probable and takes their average as the result of the vote. That is, the system treats non-voters as if they are equally likely to trust their contacts, making the final outcome dependent on the votes that have been cast and on the local neighborhood structure around non-voters for the votes that have not been cast.

One way to understand this way of dealing with abstentionism is by running an election in which nobody expresses a preference, or an election with 100% abstention; this is a computation of the expected outcome of each node that considers only the social network.

Though this computation yields a measure of centrality of the nodes (called "voting centrality" in Boldi et al.[3]), it does not seem much correlated to other simple centrality measures (such as number of contacts). This score is also not the standard PageRank computed on the social network graph, as there is a subtle, but important, difference between a node being equally likely to trust any of its contacts and a node spreading its trust equally among its contacts, as it would in the standard PageRank computation. A deep understanding of what

Figure 2. Scores computed on the same delegation graph with a low delegation factor ($\alpha = 0.2$) and a high delegation factor ($\alpha = 0.9$); the size of the nodes is proportional to their score.

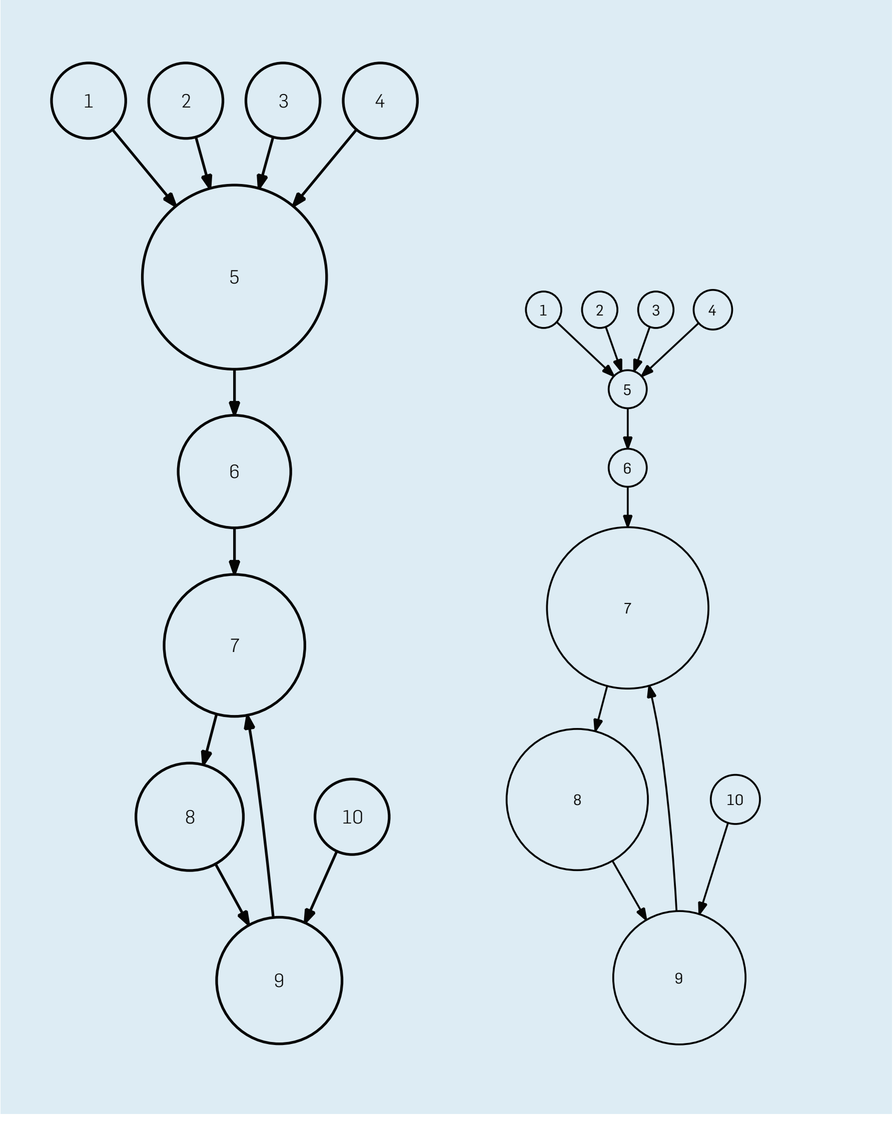

the measure actually means is still missing, but, ultimately, in viscous democracy, the social network structure influences but does not determine the outcome of all possible elections; for example, as discussed earlier, anyone can, in theory, be a winner, provided the network is connected.

Electing a Committee: Proportionality and Nonmonotonicity

Here, we discuss how viscous democracy helps select a committee of representatives to deal with some set of issues, with self-loops in the delegation graph indicating the citizens who accept the possibility of being elected to the committee; that is, nodes with a self-loop indicate their willingness to be considered candidates.

When a committee having s seats must be selected, the system can choose the s top-scoring candidates. However, ensuring proportionality provides an opportunity to select a committee that represents the diversity of users. The criterion of proportionality requires that, ideally, each political alliance has a share of the seats proportional to its share of the votes.[14]

The concept of "party" or "alliance" can be mapped onto the system. Absent specific alliances declared beforehand, a voting system for social networks may interpret the connected components of the delegation graph as alliances, as they represent communities of like-minded people delegating to other members of the community but not to non-members. This way of selecting communities is much more fine-grain than simply choosing the connected components of the underlying social network; casting a vote implies an expression of will that singles out one special relationship among a set of contacts that in practice contains many weak ties.[7]

The connected components of the delegation graph allow for proportionality enforced by picking, for each connected component, the top-k scoring nodes, in which k is proportional to the size of the connected component. For example, suppose the system must assign s seats and have c communities with $n_1, \ldots, n_c$ members, respectively; then it can assign to the i-th community $k_i = n_i \cdot c/(n_1 + \ldots + n_c)$ seats, choosing the k_i top-scoring nodes within the community. However, as in all voting systems based on proportionality, attention must be paid to how fractional seats are assigned; in the formula, k_i may not even be an integer. There are many known (and incompatible) solutions to this problem, the most widely used (including for electing the European Parliament) being the so-called D'Hondt rule,[14] which assigns the available seats one at a time, giving it to the community with largest ratio $n_i/(s_i + 1)$, where s_i is the number of seats assigned to the i-th community so far.

If viscous democracy is used for proportional voting, as in many other multiple-winner voting procedures, there is no guarantee of monotonicity. This means, counterintuitively, that not voting for node y may under certain circumstances increase y's chances of being elected.

As an example of non-monotonicity, consider the situation in Figure 3 and assume that the voting system will assign three seats using the D'Hondt rule.

If C votes for herself, the delegation graph has three communities with 12, 8, and 5 members each, respectively, with C not elected, because the leftmost community receives two seats, and the one in the middle receives the remaining one[n]. Instead, if C votes for B, we have two communities of sizes 12 and 13, with C elected, as the merged community receives two seats[o] assigned to its top-two-scoring members: B and C. In this case C is better off voting for B instead of for herself, meaning there is no monotonicity property.

Non-monotonicity makes the system vulnerable to tactical voting, though, in practice, no user has enough information about the structure of the network and about other voters' decisions to implement it. In particular, if all votes are cast simultaneously and kept secret ahead of the poll, tactical voting is even more difficult, though not, in principle, impossible. Worth noting is that the whole voting system is essentially as reliable as the underlying social network; forms of collusion (such as accepting friendship from strangers if they promise to vote for me) are possible, unless the social network manages to completely stop voters from accepting strangers as contacts.

Figure 3. If *C* wants to be elected as one of the three representatives of this network, should she, under D'Hondt proportionality rules, vote for herself or for *B*?

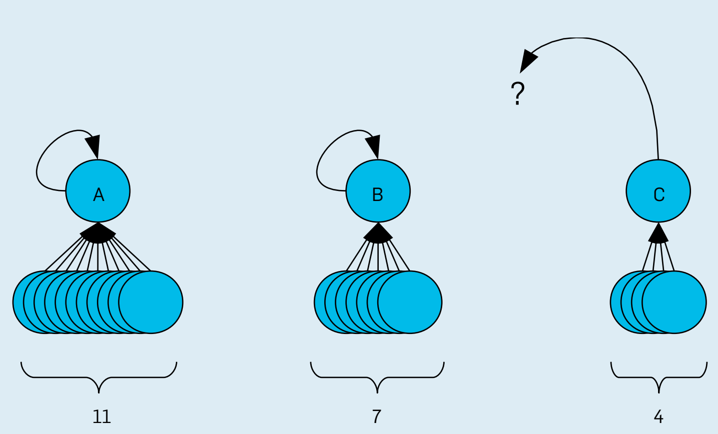

n Applying D'Hondt's rule (basically dividing the number of votes of each party first by 1, then by 2, then by 3, and so on, as seats are allocated), the initial multipliers are respectively 12, 8, and 5 units. The first community wins one seat, as it had the largest multiplier. Now multipliers are 6 (12 divided by 2), 8, and 5. The second community wins one seat, and multipliers are now 6, 4 (8 divided by 2), and 5. The first community wins one more seat, and multipliers are now 3 (6 divided by 2), 4, and 5. It would then be the turn of a third community, but no more seats are left to fill.

o Following the same procedure as before, D'Hondt multipliers are 12 and 13. The second community wins one seat, and multipliers are now 12 and 6.5 (13 divided by 2). The first community wins one seat, and multipliers are now 6 (12 divided by 2) and 6.5. The second community gets the last seat.

Simulated Voting in DBLP and Y! Meme

The effectiveness of a voting system in practice depends on qualitative factors like whether the voters are able to understand the voting system and accept it, whether the decisions reached by the community are in some sense correct, and whether the members of the community agree to go along with such decisions. Though running a real voting experiment in a sufficiently large social network would be extremely difficult, we offer several simulations suggesting viscous democracy produces results consistent with common sense and reflective of the structural properties of the social networks on which they are run.

In a CS community. First, we simulated an election in the overall computer science community using the DBLP co-authorship network[p] in which each node represents a computer scientist and the system interprets co-authorship relationships as social ties, indicating two scientists are connected if they have co-authored an article. We simulated a vote to elect the most representative author in each area, using the following criteria: Each author x considers her co-authors in decreasing order of number of papers co-authored and votes for the first one more "productive" than x, or that has written more papers than x; if no such co-author exists, x votes for herself.

The protocol we used for this simulated vote included admittedly more information than the underlying social network of co-authorship. This fact would also be true of a real-world voting experiment; a voter would choose a proxy among her contacts on the basis of information not available or deducible from the network alone, but be aware that the results of the voting depend largely on the tallying rules adopted.

Concluding the simulation, we considered the top-scoring authors within each connected component of the delegation graph; such components correspond roughly to the different research communities, as explained earlier. The table here lists the top-10 results of two communities, broadly identified as "Theory and Algorithms" and "Databases and Data Management"[q]. Due to the way the simulation was run, the most prolific authors tended to be favored, but the relationship between number of publications and score is not trivial, as reflected even in the few examples in the table[r]. In any case, the results are in line with what common sense would suggest.

In a social network. Yahoo! Meme[19] is a microblogging platform allowing for the viral spread of information memes, most posted by users as interesting or funny photos or cartoons. The social network here is directed, and a link represents a follower-followed relationship; all the viscous democracy machinery naturally extends to this

p http://www.informatik.uni-trier.de/~ley/db/

q The data and code for this experiment are available at http://law.dsi.unimi.it/dblp/.

r Kendall's τ correlation coefficient between number of publications and voting scores is 0.27, where 0 means complete lack of correlation and ± 1 means perfect positive/negative correlation.

Computer scientists with top-10 scores in two communities broadly identified as "theory and algorithms" (left) and "databases and data management" (right). The score is computed using our viscous-democracy-voting algorithm, with $\alpha = .85$; the number of publications is in parentheses.

Micha Sharir	.01594 (411)
Noga Alon	.00178 (402)
Erik D. Demaine	.00050 (305)
Avi Wigderson	.00042 (243)
Oded Goldreich	.00042 (269)
Leonidas J. Guibas	.00038 (299)
David Eppstein	.00037 (304)
Michael T. Goodrich	.00037 (228)
David Peleg	.00031 (279)
Mikhail J. Atallah	.00031 (190)

Hector Garcia-Molina	.00683 (374)
Jeffrey D. Ullman	.00097 (244)
Michael Stonebraker	.00034 (230)
Randy H. Katz	.00028 (179)
David Maier	.00027 (207)
David A. Patterson	.00021 (144)
David J. DeWitt	.00020 (179)
Rajeev Motwani	.00016 (182)
Raghu Ramakrishnan	.00015 (208)
David E. Culler	.00014 (140)

Figure 4. Representative images from two example clusters after simulated voting in Yahoo! Meme.

Example Cluster (a)

Example Cluster (b)

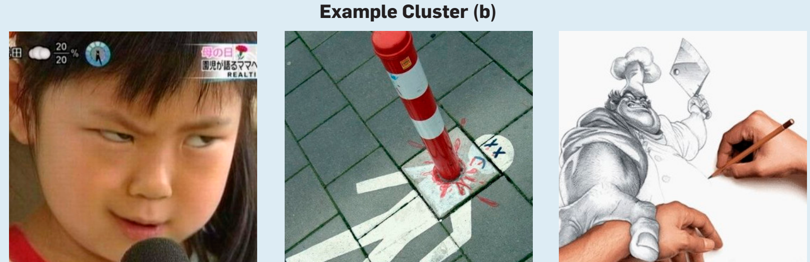

case, in which followers can vote for only one of the users they follow.

We performed a simulated election on an early snapshot of Yahoo! Meme with tens of thousands of users. First, we computed the influence of a user u considering all the memes u had posted and summing the number of copies of the memes that had been re-posted (or "retweeted" in microblogging jargon) by followers u and, recursively, by followers-of-followers. Next, we assumed that users would vote for the person they follow most often—the one from whom they have re-posted the most memes, breaking ties arbitrarily. Finally, users voted for themselves if they were not following anyone more influential than themselves.

The result of this simulated election was a series of influential users "elected" in different communities. As with the DBLP simulated election, this type of voting helps motivate the emergence of homogeneous communities. Communities tend to be homogeneous in terms of the countries where their users are located, and inspecting the memes with the largest number of re-posts in these communities, we see they correspond to coherent topics; for example, Figure 4 includes thumbnails of memes posted in two communities that tend to post "cute" and "funny" images, respectively.

Conclusion

As in any voting system, viscous democracy is not exempt from trade-offs (such as vulnerability to tactical voting and dependence on the choice of the parameter α), along with possible extensions (such as giving different weights to each delegation arc and allowing users to vote for motions and delegate fractionally[23]). But all these variants could make the system overly complicated and thus more difficult to understand and use.

Sorted by increasing complexity, people in social networks share information, cooperate, work collaboratively, and take collective action.[19] Though few online communities would be expected to collectively take binding decisions today, they will, increasingly, in the future. Novel interactive environments call for development of novel collective decision systems over the coming years.

Driven by the appeal of meritocracy, some online communities have implemented decision systems midway between the universal voting of direct democracy and representative democracy.

References

1 Albert, R., Jeong, H., and Barabasi, A.L. Diameter of the World Wide Web. *Nature 401* (Sept. 9, 1999), 130–131.
2. Arrow, K.J. *Social Choice and Individual Values.* John Wiley & Sons, New York, 1951; Second Edition, 1963.
3. Boldi, P., Bonchi, F., Castillo, C., and Vigna, S. Voting in social networks. In *Proceedings of the 18th Conference on Information and Knowledge Management* (Hong Kong, Nov. 2–6). ACM Press. New York, 2009, 777–786.
4. Boldi, P., Santini, M., and Vigna, S. PageRank: Functional dependencies. *ACM Transactions on Information Systems 27*, 4 (Nov. 2009), 1–23.
5. Butler, B., Sproull, L., Kiesler, S., and Kraut, R. Community effort in online groups: Who does the work and why? In *Leadership At a Distance*, S. Weisband, Ed. Lawrence Erlbaum Associates, Hillsdale, NJ, 2007, 171–194.
6. Carroll, L. *The Principles of Parliamentary Representation.* Harrison and Sons, London, 1884.
7. Donath, J. and Boyd, D. Public displays of connection. *BT Technology Journal 22*, 4 (Oct. 2004), 71–82.
8. Ford, B. *Delegative Democracy.* May 2002; http://www.bford.info/deleg/deleg.pdf
9. Girvan, M. and Newman, M.E.J. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences 99*, 12 (2002), 7821–7826.
10. Granovetter, M.S. The strength of weak ties. *The American Journal of Sociology 78*, 6 (May 1973), 1360–1380.
11. Green-Armytage, J. *An Introduction to Proxy-Based Democratic Systems.* May 2007; http://www.econ.ucsb.edu/~armytage/voting/proxy2007.htm
12. Hollingshead, A.B., McGrath, J.E., and O'Connor, K.M. Group task performance and communication technology: A longitudinal study of computer-mediated versus face-to-face work groups. *Small Group Research 24*, 3 (Aug. 1993), 307–333.
13. Katz, L. A new status index derived from sociometric analysis. *Psychometrika 18*, 1 (Mar. 1953), 39–43.
14. Katz, R.S. *Democracy and Elections.* Oxford University Press, 1997.
15. Leadbeater, C. *We-Think: Mass Innovation, Not Mass Production.* Profile Books Ltd., London, July 2009.
16. Nielsen, J. *Participation Inequality: Encouraging More Users to Contribute.* Jakob Nielsen's Alertbox, Oct. 9, 2006; http://www.useit.com/alertbox/participation_inequality.html
17. Page, L., Brin, S., Motwani, R., and Winograd, T. *The PageRank Citation Ranking: Bringing Order to the Web. Technical Report 66.* Stanford University, Palo Alto, CA, 1999.
18. Pani, M. *Hold Your Nose and Vote: Why Do Some Democracies Tolerate Corruption? Working Paper WP/09/83.* International Monetary Fund, Middle East and Central Asia Department, Apr. 2009.
19. Shirky, C. *Here Comes Everybody: The Power of Organizing Without Organizations.* The Penguin Press, New York, Feb. 2008.
20. Vigna, S. *Spectral Ranking* (2009); arXiv:0912.0238v2
21. Watts, D.J. and Strogatz, S.H. Collective dynamics of 'small world' networks. *Nature 393* (June 4, 1998), 440–442.
22. Willinger, W., Alderson, D., and Doyle, J.C. Mathematics and the Internet: A source of enormous confusion and great potential. *Notices of the American Mathematical Society 56*, 5 (May 2009), 586–599.
23. Yamakawa, H., Yoshida, M., and Tsuchiya, M. Toward delegated democracy: Vote by yourself, or trust your network. *International Journal of Human and Social Sciences 1*, 2 (Spring 2007), 146–150.

Paolo Boldi (boldi@dsi.unimi.it) is an associate professor in the Dipartimento di Scienze dell'Informazione of the Università degli Studi di Milano, Milano, Italy.

Francesco Bonchi (bonchi@yahoo-inc.com) is a researcher in Yahoo! Research, Barcelona, Spain.

Carlos Castillo (chato@yahoo-inc.com) is a researcher in Yahoo! Research, Barcelona, Spain.

Sebastiano Vigna (vigna@dsi.unimi.it) is an associate professor in the Dipartimento di Scienze dell'Informazione of the Università degli Studi di Milano, Milano, Italy.