Content-Driven Reputation for Collaborative Systems

Luca de Alfaro[1] and Bo Adler[2]

[1] Department of Computer Science, University of California, Santa Cruz, USA
[2] Facebook Inc., USA

Abstract. We consider collaborative editing systems in which users contribute to a set of documents, so that each document evolves as a sequence of versions. We describe a general technique for endowing such collaborative systems with a notion of content-driven reputation, in which users gain or lose reputation according to the quality of their contributions, rather than according to explicit feedback they give on one another. We show that content-driven reputation systems can be obtained by embedding the document versions in a metric space with a pseudometric that is both effort preserving (simple changes lead to close versions) and outcome preserving (versions that users perceive as similar are close). The quality of each user contribution can be measured on the basis of the pseudometric distances between appropriately chosen versions. This leads to content-driven reputation systems where users who provide contributions of positive quality gain reputation, while those who provide contributions of negative quality lose reputation. In the presence of notification schemes that prevent the formation of "dark corners" where closed groups of users can collaborate without outside interference, these content-driven reputation systems can be made resistent to a wide range of attacks, including attacks based on fake identities or specially-crafted edit schemes.

1 Introduction

In many collaborative systems, users can edit or modify the documents in the system, giving rise to a sequence of evolving versions for each document. We call such systems *collaborative editing systems.* The most prominent example of collaborative editing systems is wikis, but other systems can be similarly described. For instance, in Google Maps users can edit business listings, giving rise to a series of versions for the listings' content. A non-textual example consists in the process of uploading and revising 3D models to the Trimble 3D Warehouse [3]. Open software repositories and collaboration on shared documents are other examples.

We describe a general technique for developing *content-driven* reputation systems for collaborative editing systems. The idea behind content-driven reputation is simple: judge users by their actions, rather than by the word of other users. In a content-driven reputation system, users gain or lose reputation according to how their contributions fare: users who contribute content that is preserved, or built-upon, by later users gain reputation; users whose work is undone lose reputation. Thus, content-driven reputation systems do not require users to express judgements on one another.

Reputation systems for collaboration provide an incentive for users to contribute constructively to the system. The power of reputation in motivating users is evident in many sites, such as Stack Overflow [4]. Another use of reputation is to help predict the future behavior of users; the predictive power of reputation has been demonstrated in the Wikipedia in [7,11]. Indeed, each time we use reputation to grant privileges to users, such as the ability to perform specific system actions, we trust in part the predictive power of reputation: if we did not believe that users who contributed greatly in the past are likely to continue to provide useful contributions, there would be little reason to grant such users additional privileges. A third use of reputation is to estimate content quality and identify vandalism [6,9,8].

Content-driven reputation systems have several advantages over systems that rely primarily on user feedback [7]. User-generated rating information can be quite sparse, especially in young editing systems. Gathering the feedback and ratings requires the implementation of user interfaces that are secondary to the goal of collaboration, and can be distracting or ineffective. Content-driven reputation comes “for free”: it can be computed from content evolution information that is always present, without need for additional feedback or rating mechanisms. In content-driven reputation systems every user is turned into an active evaluator of other users’ work, by the simple act of contributing to the system. By deriving the reputation signals from content evolution, rather than separate ratings, content-driven reputation prevents schemes such as *badmouthing:* a user cannot keep a contribution, while giving poor feedback on its author. Indeed, content-driven reputation systems can be made resistant to broad categories of attacks [11].

To endow a collaborative editing system with a notion of content-driven reputation, it suffices to provide a *pseudometric* on the space of document versions. A pseudometric is a function that satisfies the same axioms as a distance (positivity, symmetry, triangular inequality), except that distinct elements of the metric space (distinct versions, in our case) can have distance 0. The pseudometric between versions should satisfy two natural requirements:

- *Outcome preserving.* If two versions look similar to users, the pseudometric should consider them close. In particular, the pseudometric should assign distance 0 to versions that look identical or that are functionally identical.
- *Effort preserving.* If a user can transform one version of a document into another via a simple transformation, the pseudometric should consider the two versions close.

These two requirements are stated in an approximate way, and meeting them perfectly in a concrete collaborative editing system may not be possible. However, the closer we get to satisfying these requirements, the higher-quality and harder-to-game the resulting reputation system will be.

For wikis, the outcome-preserving requirement means that the version pseudometric should be insensitive to differences in markup language that do not alter the way a wiki page is rendered. The effort-preserving requirement means that text that is moved from one place to the other in a document should yield a smaller pseudometric distance than separate, unrelated deletions and insertions of text. Pseudometrics suited to wikis have been analyzed in depth in [5].

Devising a suitable pseudometric is not necessarily trivial. Once a suitable pseudometric is available, however, we can use it to measure the quality of edits, by measuring how much the edits are preserved in future versions of the documents. We attribute positive quality to edits that bring the document closer to how it will be in the future, and negative quality to edits that make the document more different from how it will be in the future (these edits are thus reverted). This yields the foundation of the content-driven reputation system: users whose edits have positive quality gain reputation, while users whose edits have negative quality lose reputation.

We present and justify in detail the connection between version pseudometric distance and edit quality, and we describe how the resulting reputation system can be made resistant to broad types of attacks. The results we present are a synthesis of results from [7,11,5]. In those papers, the results were presented in the special context of text documents such as wikis. Here, we put the results in a general context, removing side-issues and complications that are particular to wikis, and showing how content-driven reputation systems can be adapted to broad classes of collaborative editing systems.

2 Collaborative Editing Systems

A *collaborative editing system* (CES) consists of a set $\mathcal{D} = \{D_1, D_2, D_3, \ldots\}$ of documents, where each document $D_i \in \mathcal{D}$ is composed of a series $v_0^i, v_1^i, v_2^i, \ldots, v_{N_i}^i$ of *versions*. The version v_0^i is a null version, indicating that the document has not been created yet. Each subsequent version v_k^i, for $0 < k \leq N_i$, is obtained from r_{k-1}^i via an *edit* $e_k^i : v_{k-1}^i \to v_k^i$. We denote by $a(v)$ the author of version v, and for brevity, we denote by $a_0^i, a_1^i, a_2^i, \ldots$ the authors $a(v_0^i), a(v_1^i), a(v_2^i), \ldots$. In the following, we will often omit the superscript i denoting the document when clear from the context, or when not relevant.

We assume that the versions of the documents of the CES belong to a metric space $\mathcal{M} = (V, d)$, where V is the set of all possible versions, and $d : V \times V \mapsto \mathbb{R}_{\geq 0}$ is a pseudometric that is symmetrical and satisfies the inequality properties: for all $u, v, w \in V$,

$$
\begin{aligned}
&d(u,u) = 0 \\
&d(u,v) = d(v,u) \\
&d(u,v) + d(v,w) \leq d(u,w) \ .
\end{aligned}
$$

We ask that d be a pseudometric, rather than a distance, because we do not require that $d(u,v) > 0$ for all distinct $u, v \in V, u \neq v$. Indeed, we will see that one of the desirable properties of the pseudometric d is that it assigns distance 0 to versions that are indistinguishable to users of the system.

The model of collaborative editing systems was inspired by wikis [7], but it can be widely applied to collaborative systems. For instance, the editing of business listings on Google Maps [2] and the editing of SketchUp models in the Trimble 3D Warehouse [3] can also be modeled as collaborative editing systems.

Wiki pages and their versions directly correspond to the documents and versions in a CES. As a pseudometric, we can use one of several notions of edit distance that satisfy the triangular inequality; see [5,18] for an in-depth discussion.

In the case of Google Maps, a business listing is comprised of various fields (title, categories, location, phone, and url, among others). Users can create new listings, and

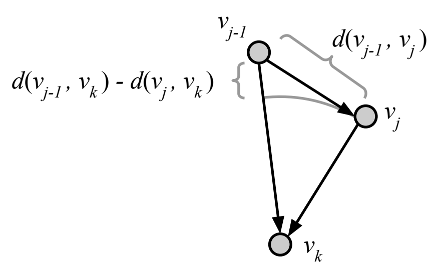

Fig. 1. The triangle of versions used to compute edit quality.

they can edit the values of the fields. The set of documents consists in the set of all business listings, and the user edits give rise to the sequence of versions. As pseudometric between fields, we can use the sum of the pseudometrics distances of the individual fields, perhaps using scaling factors that weigh the relative importance of each field. The physical distance between places on the Earth surface can be used as metric for locations; suitable distances for phone numbers and URLs consists in defining $d(u, v) = 2^{-\alpha m}$, where $\alpha > 0$ and m is the length of the longest common prefix of u and v. Distances for sets of categories are not difficult to define. These distances for the individual fields can then be combined in an overall distance for entire listings.

In the case of the 3D Warehouse of SketchUp models, the documents correspond to the designs that have been contributed by users. Users can upload updated versions of the designs, giving rise to the sequence of versions for each design. We can measure the distance between models by considering the edit distance between text descriptions of the vertices, planes, surfaces, textures, etc, comprising the designs.

In the next section, we describe some requirements of the psedumetrics that lead to useful measures of edit quality.

3 Measuring the Quality of Contributions

As a first step towards a reputation system for contributors to collaborative editing systems, we consider the problem of measuring the quality of each individual edit. We follow the idea that the quality of an edit can be measured by how long the edit survives in the subsequent history of the document [7]. To make this precise, we measure the quality of an edit $e_j : v_{j-1} \to v_j$ with the help of two versions: the previous version v_{j-1}, and a *judge* version v_k, where $j < k$. We define the *quality* $q(v_j \mid v_{j-1}, v_k)$ *of* v_j, *with respect to judge* v_k *and reference* v_{j-1}, as follows:

$$q(v_j \mid v_{j-1}, v_k) = \frac{d(v_{j-1}, v_k) - d(v_j, v_k)}{d(v_{j-1}, v_j)} \ . \tag{1}$$

To understand this definition, it might help to refer to Figure 1, and consider the situation from the point of view of the author a_k of v_k. Clearly, the author a_k prefers version

(a) $q(v_j \mid v_{j-1}, v_k) = 1$ (b) $q(v_j \mid v_{j-1}, v_k) = -1$

Fig. 2. Edits having good and bad quality.

v_k to any previous version of the document, since a_k contributed v_k. Thus, it is natural to assume that a_k will regard positively changes that bring the current version closer to v_k, and negatively changes that make the document more different from v_k. The quantity (1) captures this idea. The numerator $d(v_{j-1}, v_k) - d(v_j, v_k)$ measures how much closer the version has become to v_k due to edit e_j. The denominator $d(v_{j-1}, v_j)$ measures the total change caused by e_j. Their ratio $q(v_j \mid v_{j-1}, v_k)$ measures thus how much of the change introduced by e_j contributes to bringing the document closer to v_k.

From the triangular inequality, we have $q(v_j \mid v_{j-1}, v_k) \in [-1, 1]$ for all versions v_j, v_k.

- The maximum quality 1 is achieved when $d(v_{j-1}, v_k) = d(v_{j-1}, v_j) + d(v_j, v_k)$, which corresponds to Figure 2(a). In this case, all the change done going from v_{j-1} to v_j is preserved in going to v_k.
- The minimum quality -1 is achieved when $d(v_j, v_k) = d(v_j, v_{j-1}) + d(v_{j-1}, v_k)$, which corresponds to Figure 2(b). In this case, all the change from v_{j-1} to v_j is undone in the subsequent change from v_j to v_k: this corresponds to a reversion.

Choice of pseudometric. The definition of edit quality relies on a choice of pseudometric on the versions of the documents. To obtain a useful measure of edit quality, the pseudometric must be *effort-preserving* and *outcome-preserving*.

A pseudometric d is *effort preserving* if the distance between versions that can be easily obtained one from the other is small. An example of pseudometric that is not effort preserving is the text edit distance, measured according to the text diff tools commonly included in text revision systems, such as cvs or git [1]. The text differences computed by such tools do not model text transposition: when a block of text is moved, the resulting difference is large, even though the act of moving the text does not require much effort.

A pseudometric d is *outcome preserving* if versions that are similar to users are close in distance. In wikis, many changes to the whitespace (spaces, newlines, and so forth) do not result in visible changes of the corresponding document. If a user make changes to the whitespace of a document, these changes, having no effect, are unlikely to be reverted, even though they might not serve any purpose. If such whitespace changes resulted in non-negligible distance, they would provide users with an artificial opportunity for doing positive quality edits, while not contributing in any meaningful way to the wiki.

For wikis, the question of appropriate pseudometrics has been studied in depth in [5], where the quality of pseudometrics is measured according to the ability of the resulting reputation system to predict the quality of the future work of users. The pseudometrics that perform best are all insensitive to whitespace changes that do not affect the way in which the markup is rendered into HTML, in accordance with the outcome-presering requirement. Furthermore, unlike the Unix `diff` command, the pseudometrics that perform well track the movement of blocks of text across versions, and distinguish between text that is inserted and deleted, from text that is moved to another location in the document. This is in compliance with the effort preserving requirement: since block moves are easy to perform via cut-and-paste, they should give rise to small distances. Indeed, the best pseudometrics experimentally are those that explain the change from one version to the other via an edit list that contains a minimal amount of text insertion, deletion, and displacement: these functions measure thus the minimum amount of edit work requierd to go from one version to the other [5].

It is not difficult to devise appropriate pseudometrics for business listings, as previously mentioned. On the other hand, devising appropriate pseudometrics for complex domains, such as the 3D solids generated in SketchUp, is not an easy problem. The main difficulty lies in meeting the outcome preserving criterion, which requires the metric to consider close the designs that are visually similar.

4 Content-Driven Reputation

To construct our content-driven reputation system, we associate a reputation $r(a) \in \mathbb{R}_{\geq 0}$ with every author a. The initial value of user reputations corresponds to the amount of reputation we can accord to users whom we have never seen in the system before, and it depends on how difficult it is to create new user accounts. In Wikipedia, where there are no restrictions to the creations of user accounts, WikiTrust gives new users reputation equal to 0: if we gave new users any larger amount $r > 0$, users whose reputation fell below r could simply open another account to get back to reputation r. In systems where users cannot easily create many accounts, we can afford giving new users some amount of reputation. This is akin to social interaction: when we deal with a perfect stranger hiding behind a nickname on the internet, we usually accord very little trust to the stranger, since obtaining such fake identities is essentially free. When we deal with a real person, whose name we know, we usually accord to that person some trust, since we know that the person cannot easily change identity if the person breaks our trust.

We update user reputation as follows. For each edit $e_j : v_{j-1} \rightarrow v_j$ done by a_j, we measure the quality of e_j with respect to set $F_j \subseteq \{v_{j+1}, \ldots, v_N\}$ of future versions; the precise rule for choosing F_j will be discussed later. For each version $v \in F_j$, we update the reputation of a_j via:

$$r(a_j) := r(a_j) + q(v_j \mid v_{j-1}, v)\, d(v_{j-1}, v_j)\, f(r(a(v)))\ , \tag{2}$$

where $f : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$ is a monotonic function. Thus, the reputation of the author of e_j is incremented in proportion to the amount $d(v_{j-1}, v_j)$ of work done, multiplied by its quality $q(v_j \mid v_{j-1}, v)$, and multiplied by the reputation of the author of the reference revision v, rescaled according to a function $f(\cdot)$.

In (2), the version v has the role of "judge" in measuring the quality of the edit: the factor $f(r(a(v)))$ ensures that the higher the reputation of the author of v, the higher the weight we give to the quality judgement that uses v as reference. We rescale the reputation $r(a(v))$ using a monotonic function f to limit the influence of high-reputation users over the overall system. In most collaborative systems, including the Wikipedia, there is a group of long-term users who are responsible for a large fraction of the work; these users tend to accumulate large amounts of reputation. If in (2) we used $r(a(v))$ directly, this would give these top users an outsized influence over the reputation system. In the Wikipedia, we rescale reputations via $f(x) = log(1 + \max\{0, \varepsilon + x\})$, where $\varepsilon \geq 0$ allows us to tune the amount of influence of new users on the system. Such a logarithmic rescaling function is a natural choice when the user contribution amounts and reputations follow a power-law distribution [10,15,13], and worked well in practice for Wikipedia editions in different languages [7,6].

In order to choose the set F_j of reference versions, we first remove from $v_{j+1}, v_{j+2}, v_{j+2}, \ldots$ all the versions by the same author as v_j: we do not want a user to be a judge of his or her own work. Let $\sigma_j = v'_{j+1}, v'_{j+2}, v'_{j+3}, \ldots$ be the resulting sequence. One choice for F_j consists in taking the first K revisions of σ_j for some fixed $K > 0$; this is the choice followed in WikiTrust [7]. Another choice consists in taking F_j to be the whole σ_j, using geometrically-decaying weights for reference revisions farther in the future, to ensure that each edit causes a bounded change in the user's reputation. Under this choice, (2) becomes:

$$r(a_j) := r(a_j) + \sum_{k \geq j+1} (1-\alpha)\, \alpha^{j-k+1} q(v_j \mid v_{j-1}, v_k)\, d(v_{j-1}, v_j)\, f(r(a(v))) \tag{3}$$

for a geometric decay factor $0 < \alpha < 1$.

4.1 Truthfulness

A reputation system based on (2) or (3) is a *truthful* mechanism in the game-theoretic meaning of the term: if a user wants to modify a document, a dominating strategy (an optimal strategy for the user) consists in performing the modification as a single edit [11,17]. Users have no incentive to play complicated strategies in which the modification is broken up into a sequence of edits having the same cumulative effect. This property is fundamental in a reputation system. If users derived more reputation by breaking up edits into many small steps, or by performing every edit by first deleting the entire document, then replacing it with the new version, the evolution of the content in the collaborative system could be severely disrupted by users trying to maximize their reputation.

To prove the truthfulness of the reputation systems based on (2) or (3), we consider the case of an edit $e_j : v_{j-1} \to v_j$ being split into two edits having the same cumulative effect: $e'_j : v_{j-1} \to v'$ and $e''_j : v' \to v_j$; the general case is analogous. We analyze the case for (2); the same argument works also for (3). Consider a fixed version $v \in F_j$ used to judge e_j, and let $c = f(r(a(v)))$. For the edit e_j, the total amount of reputation gained by the author of e_j from judge v is:

$$\begin{aligned} c\, q(v_j \mid v_{j-1}, v)\, d(v_{j-1}, v_j) &= c \frac{d(v_{j-1}, v) - d(v_j, v)}{d(v_{j-1}, v_j)} d(v_{j-1}, v_j) \\ &= c\big[d(v_{j-1}, v) - d(v_j, v)\big] \ . \end{aligned} \tag{4}$$

When the edit e_j is split into e'_j, e''_j, the total amount of reputation gained due to judge v is:

$$\begin{aligned} & c\left[q(v' \mid v_{j-1}, v)\, d(v_{j-1}, v') + q(v_j \mid v', v)\, d(v', v_j)\right] \\ &= c\Big[\big[d(v_{j-1}, v) - d(v', v)\big] + \big[d(v', v) - d(v_j, v)\big]\Big] \\ &= c\big[d(v_{j-1}, v) - d(v_j, v)\big] \ . \end{aligned} \tag{5}$$

The result follows by comparing (4) and (5).

4.2 Resistance to attacks and dark corners in collaboration

The content-driven reputation defined by (2) or (3) is susceptible to attacks in which a user controls several user accounts, and coordinates the actions of these accounts in order to increase the reputation of a subset of these accounts; these attacks are broadly known as Sybil attacks or, less formally, sock-puppet attacks [14,12,16,11]. The accounts that are controlled by a user in order to enhance the reputation of the user's main account are known as *sock-puppet* accounts.

A detailed description of defense mechanisms that can be used in content-driven reputation systems against Sybil attacks appeared in [11]. We survey here the main idea, which consists in limiting the amount of reputation that can be gained from an interaction with other users, unless the contribution itself has stood the test of time.

The technique is applicable to the Wikipedia, and to other collaborative systems that, like the Wikipedia, have no "dark corners": all edits are viewed in timely fashion by honest users. More precisely, we say that a collaborative system has no dark corners within time constant T if there is a set U of *good users* such that, for every version v, v has been viewed by a user in U with probability at least $1 - e^{-t/T}$, where t is the time since the version was created. This set of good users must consists of users who are both well-intentioned, and willing to repair vandalism or damage to documents via edits. The Wikipedia, with its *recent-changes patrol* (or RC patrol), feeds of recent edits and page creations, and editors who subscribe to notifications to changes in pages, has no dark corners within a time constant of less than a day. When a collaborative system has no dark corners, a group of users cannot work at length in secrecy, hidden from view: every edit is eventually subjected to the judgement of users that do not belong to the select group.

In collaborative systems with no dark corners, the technique advocated in [11] calls for the author of a version v_j gaining reputation from a future reference version v, via (2), only in two cases:

- the reputation of the author of v is greater than the reputation of the author of v_j;
- the amount of time elapsed between v_j and v is longer than a pre-determined amount T, and for all versions v_i, v_k separated from v by less than time T, and with $i < j < k$, we have $q(v_j \mid v_i, v_k) > 0$.

These conditions ensure that a user can gain reputation only from users of higher reputation, or if no other users objected to the edits performed, for a pre-determined length of time T. Under these two conditions, [11] showed that if a user controls a set V of accounts, the user cannot raise the reputation of any account in U above the maximum max$\{r(u) \mid u \in V\}$ already held, without performing work that is recognized as useful also by the broader community of users.

This result indicates how patroling mechanisms such as notification feeds and the RC patrol contribute to the quality of a collaborative system, and how content-driven reputation can leverage such mechanisms and achieve resistance to Sybil attacks.

5 Conclusions

Content-driven notions of edit quality and reputation are well suited to a large class of collaborative editing systems, in which content evolves as a sequence of versions, each version produced by a user edit. These collaborative editing systems are common: examples include wikis, but also contributing to on-line shared documents, contributing to sofware repositories, collaboratively designing 3D objects, and editing business listings in Google Maps. Content-driven reputation systems provide a notion of user reputation that can be computed objectively, from the evolution of the content itself, without need for asking users for feedback on other user's work.

Two main requirements are needed for obtaining robust content-driven reputation systems. The first requirement is the ability to embed document versions in a metric space, so that the distance between versions is both *effort-preserving* (easy to do changes lead to close versions) and *outcome-preserving* (similar versions are close). Suitable metrics are available for text, and we believe can be developed in a great number of collaborative systems. The second requirement is the presence of patrolling mechanisms that ensure that the system does not have "dark corners" where users can work for a long time in secret, using various schemes to unduly raise their reputation. Under these two conditions, content-driven reputation systems can reward contributors whose work is preserved in the system, and are robust with respect to large categories of attacks, including Sybil attacks.

There is much research that needs to be done in furthering the use of content-driven reputation. One direction of work consists in identifying suitable notions of distance for more general collaborative editing domains. Another direction of work consists in studying the social consensus dynamics that the systems induce. For instance, the reputation-rescaling function f in (2) is used to prevent a class of users from deriving such high values of reputation, that their opinion trumps that of everyone else — creating a "reputation oligarchy". It would be of high interest to study under what conditions systems develop dominating sets of users, who cannot be replaced in spite of the constant influx of new users. A third direction of work consists in studying how to best integrate content-driven reputation with information derived from user-provided feedback and ratings.

References

1. `http://git-scm.com`.
2. `http://maps.google.com`.
3. `http://sketchup.google.com/3dwarehouse/`.
4. `http://stackoverflow.com`.
5. B.T. Adler. *WikiTrust: Content-Driven Reputation for the Wikipedia*. PhD thesis, University of California, Santa Cruz, 2012.
6. B.T. Adler, K. Chatterjee, L. de Alfaro, M. Faella, and I. Pye. Assigning trust to Wikipedia content. In *WikiSym 08: Proceedings of the International Symposium on Wikis*. ACM Press, 2008.
7. B.T. Adler and L. de Alfaro. A content-driven reputation system for the Wikipedia. In *Proc. of the 16th Intl. World Wide Web Conf. (WWW 2007)*. ACM Press, 2007.
8. B.T. Adler, L. de Alfaro, S.M. Mola-Velasco, P. Rosso, and A.G. West. Wikipedia vandalism detection: Combining natural language, metadata, and reputation features. In *Computational Linguistics and Intelligent Text Processing*, pages 277–288. Springer, 2011.
9. B.T. Adler, L. de Alfaro, and I. Pye. Detecting Wikipedia vandalism using WikiTrust. In *PAN lab report, CLEF (Conference on Multilingual and Multimodal Information Access Evaluation)*. 2010.
10. A.-L. Barabási and R Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
11. K. Chatterjee, L. de Alfaro, and I Pye. Robust content-driven reputation. In *First ACM Workshop on AISec*. ACM Press, 2008.
12. A. Cheng and E. Friedman. Sybilproof reputation mechanisms. In *Proc. of the ACM SIGCOMM workshop on Economics of peer-to-peer systems*. ACM Press, 2005.
13. A. Clauset, C.R. Shalizi, and M. E.J. Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
14. J.R. Douceur. The sybil attack. In *Peer-to-Peer Systems: First Intl. Workshop*, volume 2429 of *Lect. Notes in Comp. Sci.*, pages 251–260, 2002.
15. M. Girvan and M.EJ Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.
16. B.N. Levine, C. Shields, and N.B. Margolin. A survey of solutions to the sybil attack. Technical Report Technical Report 2006-052, Univ. of Massachussets Amherst, 2006.
17. M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
18. D. Sankoff and J.B. Kruskal, editors. *Time Warps, String Edits, and Macromolecules: the Theory and Practice of Sequence Comparison*. CSLI Publications, 1999.