Webs of Trust: Choosing Who to Trust on the Internet

Matteo Dell'Amico[0000−0003−3152−4993]

NortonLifeLock Research Group
matteo.dellamico@nortonlifelock.com
https://www.nortonlifelock.com/about/corporate-profile/research-labs/matteo-dellamico

Abstract. How to decide whether to engage in transactions with strangers? Whether we're offering a ride, renting a room or apartment, buying or selling items, or even lending money, we need a degree of trust that the others will behave as they should. Systems like Airbnb, Uber, Blablacar, eBay and others handle this by creating systems where people initially start as untrusted, and they gain reputation over time by behaving well. Unfortunately, these systems are proprietary and siloed, meaning that all information about transactions becomes property of the company managing the systems, and that there are two types of barriers to entry: first, whenever new users enter a new system they will need to restart from scratch as untrusted, without the possibility of exploiting the reputation they gained elsewhere; second, new applications have a similar cold-start problem: young systems, where nobody has reputation yet, are difficult to kickstart.

We propose a solution based on a *web of trust*: a decentralized repository of data about past interactions between users, without any trusted third party. We think this approach can solve the aforementioned issue, establishing a notion of trust that can be used across applications while protecting user privacy. Several problems require consideration, such as scalability and robustness, as well as the trade-off between privacy and accountability.

In this paper, we provide an overview of issues and solutions available in the literature, and we discuss the directions to take to pursue this project.

Keywords: Reputation · Decentralization · Social networks · Trust · Privacy · Security · Scalability · Network embeddings · Sybil attack · Whitewashing · Distributed Ledgers · Smart Contracts

1 Introduction

The Internet enables decentralized point-to-point communication between billions of users, and this has unlocked an enormous potential for interactions between them. The so-called *sharing economy*, represented by companies such as Airbnb, Uber, Blablacar, etc., exploits this, by putting in contact users that would otherwise not know each other, and letting them engage in transactions (e.g., share the cost for a ride, rent a room, etc.) that often make use of resources that would otherwise be wasted. Crucially, these services provide *reputation systems* that allow us to predict whether somebody will behave in the way they should.

These services are certainly both useful and successful, but they have a couple of shortcomings that we're interested in tackling. First, they are *proprietary*: all the data about user interactions is kept and monetized by the companies handling those services, with little control by users themselves about their own data, and the company is effectively monopolistic in its market, with the possibility of requiring high transaction fees; second, they are *siloed*, meaning that a user's information—and reputation—remains confined in that particular platform. This creates two different types of barriers to entry: first, users that enter a platform will be considered as totally unknown to the world, without the possibility of leveraging the trust they may have earned in the past, for example thanks to social connections and/or past interactions in a similar platform; second, new applications have a similar cold-start problem: if nobody has reputation, fewer people will be confident enough to start interacting, creating unnecessary friction until, if ever, enough users obtain a reputation that is positive enough.

A solution where a single entity, whether a corporation or a nation state, manages the reputation of people from all points of view is obviously criticizable: the Chinese "social credit system" [37]—a reputation system that can, for example, bar people from taking planes or flights if their reputation is bad—raised the alarm of privacy advocates over the world, and has drawn comparisons to pieces of dystopic science fiction [74]. Similar comparisons were drawn for Peeple—a proposed application to leave reviews for people based on professional or personal relationships—which was harshly criticized by public opinion, and the company backtracked to the point of allowing people to veto the reviews they receive, making the usefulness of the application (not yet launched as of January 2020) questionable [57].

We consider the problem of designing a reputation system that does not suffer from the privacy shortcomings described above. We require the system to be *decentralized*, and we want users to be *in control* of whether and how they should appear in the system. The "web of trust" name was first proposed by Zimmerman [82] as a decentralized way of certifying other people's identity—as opposed to the hierarchical structure of trusted third parties like certification authorities that is used, for example, for TLS/SSL certificates. We extend this concept, and here we call "web of trust" a decentralized construction that keeps all kinds of assessments between users, with the goal of creating an efficient and privacy-conscious system without the need of any trusted third party. Garcia Lopez et al. [30] discuss the problems of incentives to cooperation, free riding and decentralized trust as key weaknesses of permissionless blockchains; we think that our effort can be useful in alleviating them.

Recent privacy regulations, such as the European Union's General Data Protection Regulation (GDPR) [26] and the recent California Consumer Privacy Act

(CCPA [1], in turn inspired by the GDPR [48]) require that users are given clear and informed opportunities to give consent about using their personal data, and that they have access to granular controls that allows them to decide which of their personal data is shared with whom. The system that we plan to build is based on these concepts: information is shared in a decentralized fashion with trusted peers, through policies which are under control of users themselves.

This is an ambitious problem, and luckily an important corpus of research can be exploited to solve these problems: the goal of this paper is to organize the related work—in sometimes disparate communities—that can be harnessed to reach our goal, and to highlight the most important open questions.

We tackle the problem of representing user reputation, as discussed in Section 2, along with ways to formalize it in such a way that there exist sound incentives to cooperation even in a completely decentralized setting. We then discuss in Section 3 the security and privacy issues connected to this, in particular as connected to the question of pseudonimity and the opportunity for *whitewashing*—i.e., discarding a user's past bad reputation—and Sybil attacks—i.e., creating large numbers of fake users to subvert the system. We then move on to architectural concerns dealing with decentralization and scalability, discussed in Section 4, including decentralized approaches to represent social networks in a privacy-aware way. In Section 5, we discuss how our design can be helped by distributed ledgers such as blockchains, and smart contracts on top of it, in order to provide consistency in a completely decentralized architecture. The conclusions of Section 6 summarize the state of the art, with potential problem, solutions, and open issues.

2 Formalizing Reputation

Here we provide an overview of the concepts and formalisms used to represent trust and reputation in a computational fashion. For more in-depth discussions, we refer the interested reader to more comprehensive works [20, 39].

It is known that reciprocative behavior can make cooperation evolve between selfish actors. For example, the game-theoretic work of Axelrod and Hamilton [3] has shown that simple "tit-for-tat" strategies—where players rewards peer that cooperate and punish those that defect—are successful in various settings of the iterated prisoner's dilemma. Cohen [14] applied successfully this strategy when designing the BitTorrent P2P file-sharing protocol.

Simple reciprocation is effective when two users interact frequently with each other, such that opportunities for reciprocation happen often. Unfortunately, this is not always the case: for one-off interactions, tit-for-tat strategies are not enough. *Reputation*, in this case, can be a means to enable *indirect reciprocation*, based on the idea that my cooperation with others will boost my reputation, and when I have a high reputation others will be more cooperative with me.

In the absence of central authorities that attribute a reputation score to everybody, the concept of reputation is still viable. Consider the example by Jösang et al. [39] in Fig. 1: Alice trusts Bob, and Bob trusts David. Bob, there-

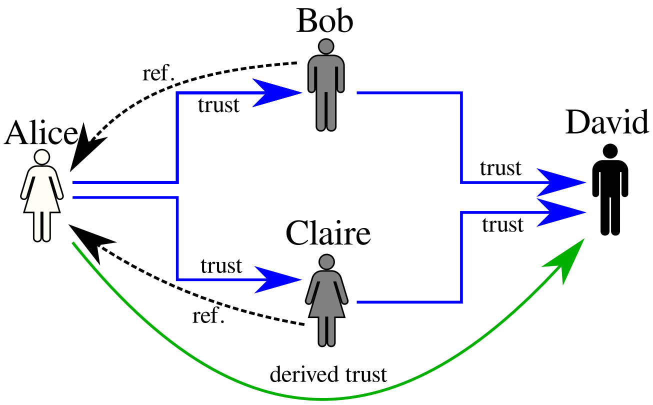

Fig. 1. Transitive trust propagation.

fore, recommends David to Alice, who obtains derived trust in David. Also Claire recommends David to Alice, giving her another reason to increase her trust in David. Since now Alice trusts David to some extent, she might also want to put trust in further people recommended by him. Note that, in this framework, reputation is *subjective*, depending on the subject that evaluates it—to put it another way, there are *no pre-trusted entities*, and everybody can *choose their trust anchors freely*; in the example, Alice's trust anchors are Bob and Claire.

In the following, we will consider a *web of trust* as a graph $G = (V, E)$ where the nodes V are our users and the edges E are *feedback* given by a user for others. A *reputation scoring function* $r : V \times V \rightarrow \mathbb{R}$, such that $R(a, b)$ is the reputation for user b as seen by a, will be the way that we use to compute a reputation score. If the transitive trust propagation pattern is used, r will be such that it will depend on paths in G from a to b. There are other possible trust propagation patterns: for example, when judgments are a matter of taste (say, we consider evaluating how somebody cooks), we may want to trust more people that have tastes, and hence judgments, similar to ours. The HITS [45] and SALSA [49] algorithms use a zig-zag propagation pattern that can be used to reflect this.

In our view, users should be free to adjust the parameters of the reputation score function—or choose different reputation score functions altogether—depending on their own preferences and on the domain at hand: for example, which types of endorsements to consider, how to aggregate them, etc. Choosing a suitable reputation scoring function is not trivial, in particular because they should be defined in such a way that they cannot be manipulated by attackers: results on security of reputation metrics are discussed in Section 3.

3 Privacy, Robustness and Accountability

Reputation systems should of course not be gameable by adversaries; here we give an overview of the issues we consider most closely related to our problems; for a more in-depth look we defer to Hoffman et al. [36].

The Cost of Cheap Pseudonyms A key trade-off between privacy and accountability is the possibility of creating cheap pseudonyms: from a user's privacy point of view it is of course desirable to have several different, unlinked and possibly disposable identities. This, however, means that identities with bad reputation can simply be forgotten (*whitewashing*) and will not be linked with another user's profiles, giving them the opportunity to misbehave without paying consequences. Friedman and Resnick [29] found that, in this case, *"a large degree of cooperation can still emerge, through a convention in which newcomers 'pay their dues' by accepting poor treatment from players who have established positive reputations."* Cheap pseudonyms, hence, do not make reputation systems useless but they limit their positive impact by introducing a kind of "cold-start" problem, as also corroborated by Feldman et al. [28] in the context of P2P systems.

Reputation systems can still be useful when whitewashing is present, but this essentially rules out large transactions with important losses in case somebody misbehaves (think, e.g., of a large loan). In our scenario, we consider we should handle *both* persistent and disposable identities, handling them distinctly in the reputation system.

An interesting possibility with respect to privacy is the field of zero-knowledge proofs, with cryptographic constructs such as zk-SNARKs [7] and zk-STARKs [4]: we will investigate to which extent they can be used to prove a user's reputation score without disclosing too much information about their other past interactions. Another interesting approach to provide anonymity in reputation systems is the mix-net strategy adopted by AnonRep [79]; unfortunately, besides not allowing subjective reputation evaluation, AnonRep is susceptible to Sybil attacks. Lifting this restriction would be an interesting research direction.

Sybil attack A problem related to cheap pseudonyms is the Sybil attack [23], where a system, or parts of it, is subverted by creating a large—and possibly unlimited—number of fake identities. Cheng and Friedman [11] show that some reputation mechanisms—e.g., those based on the MaxFlow measure—are immune to Sybil attacks, in the sense that attackers cannot gain reputation score by creating fake identities; Dell'Amico and Capra [21] show that metrics such as Personalized PageRank (PPR) also give some guarantees against Sybil attacks, and propose new metrics that are Sybil-resistant while also employing the trust propagation pattern observed in HITS, SALSA and the large majority of recommender systems. A related line of research is the one by Yu et al. [77, 78], who use social networks to limit the number of Sybil users that are accepted into a system. This approach is based on the assumption that the benign part of a social network will be *fast-mixing*, i.e., random walks will quickly become uncorrelated with the place they started from; measurement studies have shown that this property is not always verified, in particular when the creation of a link in a social network requires co-location, resulting in geographically clustered networks [22].

Since we consider a system that allows for disposable identities, we must take into account the Sybil attack. Rather than building on the possibly non-verified hypothesis of a fast-mixing network, we find preferable the more solid guarantees of Sybilproof or Sybil-resistant mechanisms [11, 21].

Negative Feedback Reports of bad experiences should be taken into account as well [2, 31], but we need to make sure that the feature shouldn't be abused, for example through blackmail or retaliation [44]. This can be done by designing asymmetric systems (e.g., only one partner in an interaction can give negative reviews [12]), by associating feedback with only verified interactions, and/or by hiding user reviews until all those involved in a transaction are also committed.

4 Decentralization and Scalability

Computing reputation poses scalability issues: with the approach described in Section 2, reputation is a function of the paths on the Web of trust between two nodes; if the edges of this social network (assessments) are not public or they are simply too many, then finding them becomes tricky.

This problem—finding paths in social networks in a decentralized fashion—goes back to 1967, when it was popularized as the "small-world problem" [73]: in fact, social networks connect seemingly remote people through rather short chains of acquaintances and, perhaps surprisingly, people are able to efficiently find those chains (i.e., short paths in the social network) even without knowing the full network; Kleinberg [46] provided a mathematical model that possesses these properties, whereby nodes in a graph are placed as points in a circle, and each node has short-range links to their neighbors and long-range links to faraway nodes. If a node knows the circle position of the destination and each of its neighbor, a simple *greedy* routing strategy routing towards the neighbor that is closest to the destination is sufficient to quickly reach the destination. We can see the circle positions as a *network embedding* in a space that is in this case one-dimensional; while the circular embedding of Kleinberg is a good one for the particular kind of synthetic small-world networks created in that work, however, the same kind of embedding is not ideal for real-world social networks that can be better represented in more complex spaces.

Besides computing reputation, our problem of routing in opaque networks arises for routing in friend-to-friend networks [5, 13, 63], and to discover suitable paths for off-chain payment channels [60, 62]. In most cases, the problem is solved through an embedding: a set of coordinates associated to each node such that close nodes in the embedding are likely to be also close in the original graph. The routing algorithm can be the simple greedy approach described before, or some generalization of it (for example, keeping a queue of discovered nodes to limit the likelihood of getting stuck in a local minimum, like the solution adopted by Malkov et al. [52]).

The problem of finding a good embedding for a graph is a recurrent and important one in computer science [32, 35], with a variety of applications such as visualization, link prediction, community detection in addition to finding shortest paths in a graph [6, 18, 34, 50, 80, 81]. Several recent approaches, such as DeepWalk [58], LINE [72], PTE [71] and node2vec [33] have been unified as related to the factorization of network's Laplacian matrices [61].

A related field is the one of Internet coordinate systems [15, 55, 56, 59, 69], which assign coordinates to Internet nodes, with the goal of making the distance between any pair of nodes an estimation of the Internet network latency between them; a few fully decentralized approaches exist [9, 16, 17, 53]. Approaches taken in this space can be of inspiration for our problem, even though assumptions are different. In particular, (i) Internet coordinate systems are based on a real-world infrastructure with geographical constraints, hence the final layout will be influenced by those geographical characteristics; (ii) nodes can freely ping each other and so-called *beacons* that serve as references, in order to obtain better precision in network distance estimation; (iii) in Internet coordinate systems nodes can only lie in one direction: while they can artificially make their latencies appear higher, they cannot answer pings faster than what the network infrastructure allows. In short, while techniques used for Internet coordinate systems are certainly a good source of inspiration, we cannot directly use these approaches for our goal.

Fortunately, there exist decentralized approaches to network embeddings that do not leverage on the assumptions above. In the following, we outline the ones that we are aware of:

- Sandberg's approach based on Kleinberg's model for routing in the Freenet friend-to-friend network ("darknet"), embedding nodes in a circle [67]. This clean and simple approach, however, appears not very well suited to some more complex social networks [19];
- Approaches based on spectral analysis: as discussed before, Qiu et al. [61] showed that many network embedding approaches can be unified as ones based on spectral properties of the graph adjacency matrix or some related ones, like its Laplacian matrix. Dell'Amico [19] proposes a decentralized implementation of an embedding algorithm by Koren [47] (initially conceived for graph drawing) and evaluates the approach in the context of finding short paths in social networks; Kempe and McSherry [42] describe a generic approach for distributed network factorization. Ling et al. [51] propose an alternative approach which however appears less suited to our case, because the number of nodes in the network (and hance, users in the social network) is required to be fixed and known a priori.
- A couple of approaches based on spanning trees. Roos et al. [64] show how one can build multiple spanning trees for the same graph and use them for routing in friend-to-friend networks; a subsequent work [65] adapts the same

approach to payment networks. These approaches appear effective, but they have a possible limitation on the side of centralization, because the number of spanning trees that can be built is small (the papers experiment with around 10 spanning trees per network); the roots of those spanning trees introduce an element of centralization.

- A piece of work by Kermarrec et al. [43] based on a force-based layout: nodes repel each other while edges bind them together with a force proportional to their weight. Here, it is critical to find efficient ways of discovering other nodes that are close in the embedding to compute the repulsive forces; Kermarrec et al. propose a gossip-based approach inspired by Voulgaris and Van Steen [75], which may be problematic for our privacy requirements as it would require to share information with strangers. Alternative solutions (e.g., one based on a solution where all communication is tunneled through paths in the social network) may be possible, but scalability trade-offs should be evaluated.

It is interesting to see that while these approaches attempt to solve the same of similar problems, work that compares and contrasts them is lacking–these pieces of work actually rarely even reference each other. We are working on bridging this gap by implementing and comparing these approach, in order to gain further insight on the weaknesses and merits of each.

Of course, in an adversarial setting like ours, security is a key requirement: malicious users shouldn't be able to subvert the routing algorithm such that paths that would discover trusted users are not found. Kaafar et al. [41] show that decentralized network coordinate systems can be compromised by malicious nodes, and propose a system based on pre-trusted supernodes to mitigate this problem [40], while Sherr et al. [70] propose a fully decentralized countermeasure based on voting. Chen et al. [10] discuss attacks to matrix factorization-based network coordinates approaches, and interestingly propose a reputation-based countermeasure to counter them: this style of defense may be effective in a network whose very purpose is to compute reputation and whose edges do represent trust relationships. Evans et al. [27] and Schiller et al. [68] discuss attacks to Sandberg [67]'s Freenet routing algorithm, and propose countermeasures to secure it. Finally, in a recent piece of work, Bojchevski and Günnemann [8] discuss attacks on centralized node embeddings through the lenses of adversarial machine learning.

In summary, we see that an impressive amount of relevant work exists in terms of network embeddings. In our view, what is needed is a comprehensive and systematic evaluation in light of our goal of computing reputation in a privacy-aware way. Four properties are fundamental to our goals, and should be evaluated: level of decentralization, scalability, security and privacy. Once this comparison is made, we will be able to evaluate the best architectural choice for this problem, and outline if there are any major problems that still need to be solved.

5 Consistency: Distributed Ledgers and Smart Contracts

The problem of finding a consistent state in a decentralized network was an unsolved problem, until 2008, when Nakamoto [54] introduced the disruptive concept of blockchain. Blockchains allow creating a *distributed ledger* (DL), that is, an append-only, unmodifiable data structure that is readable and writeable by everybody. While the concept of distributed ledger was famously created to enable cryptocurrencies, the data structure itself lends itself to several other important uses: a rather trivial one—which is indeed the cornerstone for many others—is providing a place where transactions are sorted and recorded forever univocally.

Probably the most generic way application of a distributed ledger is using it as input to a Turing-complete virtual machine: this is the approach taken by Ethereum [76], which is centered on the concept of *smart contracts*, which are programs for a "world computer", as the Ethereum creators informally call the Ethereum Virtual Machine (EVM) [25]. In this way, anything that can be written in software can be represented, in a consistent and completely decentralized way, in the state of the EVM. Among the software running on Ethereum, Decentralized Autonomous Organizations (DAOs) [24, 38]—completely self-organizing, decentralized organizations—are particularly fascinating experiments; we think that efforts that try to build sustainable economies on DAOs [66] may benefit from a distributed reputation metric.

Distributed ledgers and smart contracts are certainly powerful and flexible constructs, but they also have shortcomings: all data and programs for them need to be written on a DL, and writing on the DL is expensive and has high latency (improving latency and throughput of distributed ledgers is a very active area of research). The consequence is that we think that our architecture can and should benefit from DLs in order to achieve features like non-repudiable ratings (for example, this avoids that a user associates different ratings to the same transaction when reporting it to different recipients); similarly, smart contracts can be used to obtain agreement on computations when otherwise it would be problematic. Another use case where DLs are beneficial is associating ratings between users in the web of trust to transactions that already happened on DLs: since creating those transactions is expensive, this gives us a higher confidence that those transactions actually happened, and they are not part of a Sybil attack aimed to subvert the reputation system. Ironically, in a case like this, a shortcoming of a DL (the cost of writing on it) makes them more apt to our use case.

In general, rather than centering the design on a DL-based solution, we envision a system that should avoid the costs of DLs as much as possible, and use them only when and where they are essential to create a completely decentralized architecture.

6 Conclusion

We have explored the challenges that arise when designing a global-scale web of trust for reputation, together with a discussion of related pieces of work and directions towards its implementation. While the objective is ambitious, it is encouraging to see that many difficult sub-problems have been tackled by an impressive array of research, in many cases from different communities.

Our immediate plan now is evaluating and comparing the approaches to decentralized routing discussed in Section 4, in order to find to what extent these approaches are usable for our tasks, and which ones would be preferable: we think that this could be the core of a system that can be used effectively.

This document has been written with the goal of inspiring discussion, critique, and collaboration; while much has to be done, we think that this document can provide useful information on key design issues and related work with the goal of developing an open, decentralized and privacy-conscious reputation system for the Internet.

References

[1] Assembly, C.S.: The California Consumer Privacy Act of 2018 (2018), URL https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201720180AB375

[2] Avesani, P., Massa, P., Tiella, R.: Moleskiing. it: a trust-aware recommender system for ski mountaineering. International Journal for Infonomics **20**(35), 1–10 (2005)

[3] Axelrod, R., Hamilton, W.D.: The evolution of cooperation. science **211**(4489), 1390–1396 (1981)

[4] Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity. Cryptol. ePrint Arch., Tech. Rep **46**, 2018 (2018)

[5] Bennett, K., Grothoff, C., Horozov, T., Patrascu, I., Stef, T.: Gnunet-a truly anonymous networking infrastructure. In: In: Proc. Privacy Enhancing Technologies Workshop (PET, Citeseer (2002)

[6] Berchenko, Y., Teicher, M.: Graph Embedding through Random Walk for Shortest Paths Problems. In: Watanabe, O., Zeugmann, T. (eds.) Stochastic Algorithms: Foundations and Applications, pp. 127–140, Lecture Notes in Computer Science, Springer Berlin Heidelberg (2009), ISBN 978-3-642-04944-6

[7] Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, pp. 326–349, ITCS '12, ACM, New York, NY, USA (2012), ISBN 978-1-4503-1115-1, https://doi.org/10.1145/2090236.2090263, URL http://doi.acm.org/10.1145/2090236.2090263

[8] Bojchevski, A., Günnemann, S.: Adversarial Attacks on Node Embeddings via Graph Poisoning. In: International Conference on Machine Learning, pp. 695–704 (2019)

[9] Chen, Y., Wang, X., Shi, C., Lua, E.K., Fu, X., Deng, B., Li, X.: Phoenix: A weight-based network coordinate system using matrix factorization. IEEE Transactions on Network and Service Management **8**(4), 334–347 (2011)

[10] Chen, Y., Wu, S., Li, J., Fu, X.: NCShield: Protecting Decentralized, Matrix Factorization-Based Network Coordinate Systems. IEEE Transactions on Services Computing **10**(2), 244–257 (Mar 2017), ISSN 1939-1374, https://doi.org/10.1109/TSC.2015.2437383

[11] Cheng, A., Friedman, E.: Sybilproof reputation mechanisms. In: Proceedings of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems, pp. 128–132, ACM (2005)

[12] Chwelos, P., Dhar, T.: Caveat emptor: Differences in online reputation mechanisms. Tech. rep., Working Paper, Sauder School of Business, University of British Columbia (2006)

[13] Clarke, I., Sandberg, O., Wiley, B., Hong, T.W.: Freenet: A distributed anonymous information storage and retrieval system. In: Designing privacy enhancing technologies, pp. 46–66, Springer (2001)

[14] Cohen, B.: Incentives build robustness in bittorrent. In: Workshop on Economics of Peer-to-Peer systems, vol. 6, pp. 68–72 (2003)

[15] Costa, M., Castro, M., Rowstron, R., Key, P.: PIC: Practical Internet coordinates for distance estimation. In: 24th International Conference on Distributed Computing Systems, 2004. Proceedings., pp. 178–187, IEEE (2004)

[16] Cox, R., Dabek, F., Kaashoek, F., Li, J., Morris, R.: Practical, distributed network coordinates. ACM SIGCOMM Computer Communication Review **34**(1), 113–118 (2004)

[17] Dabek, F., Cox, R., Kaashoek, F., Morris, R.: Vivaldi: A decentralized network coordinate system. In: ACM SIGCOMM Computer Communication Review, vol. 34, pp. 15–26, ACM (2004)

[18] Das Sarma, A., Gollapudi, S., Najork, M., Panigrahy, R.: A sketch-based distance oracle for web-scale graphs. In: Proceedings of the third ACM international conference on Web search and data mining, pp. 401–410, ACM (2010)

[19] Dell'Amico, M.: Mapping small worlds. In: Peer-to-Peer Computing, 2007. P2P 2007. Seventh IEEE International Conference on, pp. 219–228, IEEE (2007)

[20] Dell'Amico, M.: Exploiting Social Networks in Robust P2P Applications. Ph.D. thesis, Università degli Studi di Genova (2008), URL https://www.disi.unige.it/person/DellamicoM/research/phd-thesis.pdf

[21] Dell'Amico, M., Capra, L.: Dependable filtering: Philosophy and realizations. ACM Transactions on Information Systems (TOIS) **29**(1), 5 (2010)

[22] Dell'Amico, M., Roudier, Y.: A measurement of mixing time in social networks. In: Proceedings of the 5th International Workshop on Security and Trust Management, Saint Malo, France, p. 72 (2009)

[23] Douceur, J.R.: The Sybil attack. In: International workshop on peer-to-peer systems, pp. 251–260, Springer (2002)

[24] DuPont, Q.: Experiments in algorithmic governance: A history and ethnography of “the dao,” a failed decentralized autonomous organization. In: Bitcoin and Beyond, pp. 157–177, Routledge (2017)

[25] Ethereum: Ethereum: the world computer (2015), URL https://www.youtube.com/watch?v=j23HnORQXvs

[26] European Union: Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Official Journal of the European Union (2016)

[27] Evans, N.S., GauthierDickey, C., Grothoff, C.: Routing in the dark: Pitch black. In: Twenty-Third Annual Computer Security Applications Conference (ACSAC 2007), pp. 305–314, IEEE (2007)

[28] Feldman, M., Papadimitriou, C., Chuang, J., Stoica, I.: Free-riding and whitewashing in peer-to-peer systems. IEEE Journal on Selected Areas in Communications **24**(5), 1010–1019 (2006)

[29] Friedman, E.J., Resnick, P.: The social cost of cheap pseudonyms. Journal of Economics & Management Strategy **10**(2), 173–199 (2001)

[30] Garcia Lopez, P., Montresor, A., Datta, A.: Please, do not Decentralize the Internet with (Permissionless) Blockchains! In: 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), pp. 1901–1911 (Jul 2019), https://doi.org/10.1109/ICDCS.2019.00188, iSSN: 1063-6927

[31] Golbeck, J.A.: Computing and applying trust in web-based social networks. Ph.D. thesis, University of Maryland (2005)

[32] Goyal, P., Ferrara, E.: Graph embedding techniques, applications, and performance: A survey. Knowledge-Based Systems **151**, 78–94 (2018)

[33] Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 855–864, ACM (2016)

[34] Gubichev, A., Bedathur, S., Seufert, S., Weikum, G.: Fast and accurate estimation of shortest paths in large graphs. In: Proceedings of the 19th ACM international conference on Information and knowledge management, pp. 499–508, ACM (2010)

[35] Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584 (2017)

[36] Hoffman, K., Zage, D., Nita-Rotaru, C.: A survey of attack and defense techniques for reputation systems. ACM Computing Surveys (CSUR) **42**(1), 1 (2009)

[37] Hvistendahl, M.: Inside China's vast new experiment in social ranking. Wired (2017), URL https://www.wired.com/story/age-of-social-credit/

[38] Jentzsch, C.: Decentralized autonomous organization to automate governance. White paper, November (2016)

[39] Jösang, A., Ismail, R., Boyd, C.: A survey of trust and reputation systems for online service provision. Decision support systems **43**(2), 618–644 (2007)

[40] Kaafar, M.A., Mathy, L., Barakat, C., Salamatian, K., Turletti, T., Dabbous, W.: Securing internet coordinate embedding systems. In: ACM SIGCOMM Computer Communication Review, vol. 37, pp. 61–72, ACM (2007)

[41] Kaafar, M.A., Mathy, L., Turletti, T., Dabbous, W.: Real attacks on virtual networks: Vivaldi out of tune. In: Proceedings of the 2006 SIGCOMM workshop on Large-scale attack defense, pp. 139–146, ACM (2006)

[42] Kempe, D., McSherry, F.: A decentralized algorithm for spectral analysis. Journal of Computer and System Sciences **74**(1), 70–83 (2008)

[43] Kermarrec, A.M., Leroy, V., Trédan, G.: Distributed social graph embedding. In: Proceedings of the 20th ACM international conference on Information and knowledge management, pp. 1209–1214, ACM (2011)

[44] Klein, T.J., Lambertz, C., Spagnolo, G., Stahl, K.O.: Last minute feedback. Tech. rep., SFB/TR 15 Discussion Paper (2006)

[45] Kleinberg, J.M.: Authoritative sources in a hyperlinked environment. Journal of the ACM (JACM) **46**(5), 604–632 (1999)

[46] Kleinberg, J.M.: Navigation in a small world. Nature **406**(6798), 845 (2000)

[47] Koren, Y.: On spectral graph drawing. In: International Computing and Combinatorics Conference, pp. 496–508, Springer (2003)

[48] Lapowsky, I.: California Unanimously Passes Historic Privacy Bill. Wired (Jun 2018), URL https://www.wired.com/story/california-unanimously-passes-historic-privacy-bill/

[49] Lempel, R., Moran, S.: Salsa: the stochastic approach for link-structure analysis. ACM Transactions on Information Systems (TOIS) **19**(2), 131–160 (2001)

[50] Liao, Y., Du, W., Geurts, P., Leduc, G.: DMFSGD: A decentralized matrix factorization algorithm for network distance prediction. IEEE/ACM Transactions on Networking (TON) **21**(5), 1511–1524 (2013)

[51] Ling, Q., Xu, Y., Yin, W., Wen, Z.: Decentralized low-rank matrix completion. In: 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2925–2928, IEEE (2012)

[52] Malkov, Y., Ponomarenko, A., Logvinov, A., Krylov, V.: Approximate nearest neighbor algorithm based on navigable small world graphs. Information Systems **45**, 61–68 (Sep 2014), ISSN 0306-4379, https://doi.org/10.1016/j.is.2013.10.006, URL http://www.sciencedirect.com/science/article/pii/S0306437913001300

[53] Mao, Y., Saul, L.K., Smith, J.M.: Ides: An internet distance estimation service for large networks. IEEE Journal on Selected Areas in Communications **24**(12), 2273–2284 (2006)

[54] Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008), URL https://bitcoin.org/bitcoin.pdf

[55] Ng, T.E., Zhang, H.: Predicting Internet network distance with coordinates-based approaches. In: Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies, vol. 1, pp. 170–179, IEEE (2002)

[56] Ng, T.E., Zhang, H.: A Network Positioning System for the Internet. In: USENIX Annual Technical Conference, General Track, pp. 141–154 (2004)

[57] Pearson, J.: Peeple has backtracked to the point of pointlessness. Motherboard (2015), URL https://motherboard.vice.com/en_us/article/vv74z3/peeple-has-backtracked-to-the-point-of-pointlessness

[58] Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 701–710, ACM (2014)

[59] Pias, M., Crowcroft, J., Wilbur, S., Harris, T., Bhatti, S.: Lighthouses for scalable distributed location. In: International Workshop on Peer-To-Peer Systems, pp. 278–291, Springer (2003)

[60] Poon, J., Dryja, T.: The Bitcoin Lightning network: Scalable off-chain instant payments. See https://lightning. network/lightning-network-paper. pdf (2016)

[61] Qiu, J., Dong, Y., Ma, H., Li, J., Wang, K., Tang, J.: Network Embedding As Matrix Factorization: Unifying DeepWalk, LINE, PTE, and Node2vec. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 459–467, WSDM '18, ACM, New York, NY, USA (2018), ISBN 978-1-4503-5581-0, https://doi.org/10.1145/3159652.3159706, URL http://doi.acm.org/10.1145/3159652.3159706, event-place: Marina Del Rey, CA, USA

[62] Raiden: What is the Raiden network? (2019), URL https://raiden.network/101.html

[63] Rogers, M., Bhatti, S.: How to disappear completely: A survey of private peer-to-peer networks. RN **7**(13), 1 (2007)

[64] Roos, S., Beck, M., Strufe, T.: Anonymous addresses for efficient and resilient routing in f2f overlays. In: Computer Communications, IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on, pp. 1–9, IEEE (2016)

[65] Roos, S., Moreno-Sanchez, P., Kate, A., Goldberg, I.: Settling payments fast and private: Efficient decentralized routing for path-based transactions. arXiv preprint arXiv:1709.05748 (2017)

[66] Rozas, D., Tenorio-Fornés, A., Díaz-Molina, S., Hassan, S.: When ostrom meets blockchain: Exploring the potentials of blockchain for commons governance. Available at SSRN 3272329 (2018)

[67] Sandberg, O.: Distributed routing in small-world networks. In: 2006 Proceedings of the Eighth Workshop on Algorithm Engineering and Experiments (ALENEX), pp. 144–155, SIAM (2006)

[68] Schiller, B., Roos, S., Hofer, A., Strufe, T.: Attack resistant network embeddings for darknets. In: 2011 IEEE 30th Symposium on Reliable Distributed Systems Workshops, pp. 90–95, IEEE (2011)

[69] Shavitt, Y., Tankel, T.: Big-bang simulation for embedding network distances in euclidean space. IEEE/ACM Transactions on Networking (TON) **12**(6), 993–1006 (2004)

[70] Sherr, M., Blaze, M., Loo, B.T.: Veracity: practical secure network coordinates via vote-based agreements. In: Proceedings of the 2009 conference on USENIX Annual technical conference, pp. 13–13, USENIX Association (2009)

[71] Tang, J., Qu, M., Mei, Q.: Pte: Predictive text embedding through large-scale heterogeneous text networks. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1165–1174, ACM (2015)

[72] Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: Line: Large-scale information network embedding. In: Proceedings of the 24th international conference on world wide web, pp. 1067–1077, International World Wide Web Conferences Steering Committee (2015)

[73] Travers, J., Milgram, S.: The small world problem. Phychology Today **1**(1), 61–67 (1967)

[74] Vincent, A.: Black Mirror is coming true in China, where your 'rating' affects your home, transport and social circle. The Telegraph (2017), URL https://www.telegraph.co.uk/on-demand/2017/12/15/black-mirror-coming-true-china-rating-affects-home-transport/

[75] Voulgaris, S., Van Steen, M.: Epidemic-style management of semantic overlays for content-based searching. In: European Conference on Parallel Processing, pp. 1143–1152, Springer (2005)

[76] Wood, G.: Ethereum: A secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper **151**, 1–32 (2014)

[77] Yu, H., Gibbons, P.B., Kaminsky, M., Xiao, F.: Sybillimit: A near-optimal social network defense against sybil attacks. In: 2008 IEEE Symposium on Security and Privacy (S&P 2008), pp. 3–17, IEEE (2008)

[78] Yu, H., Kaminsky, M., Gibbons, P.B., Flaxman, A.: Sybilguard: defending against sybil attacks via social networks. In: ACM SIGCOMM Computer Communication Review, vol. 36, pp. 267–278, ACM (2006)

[79] Zhai, E., Wolinsky, D.I., Chen, R., Syta, E., Teng, C., Ford, B.: AnonRep: Towards Tracking-Resistant Anonymous Reputation. In: 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16), pp. 583–596, USENIX Association, Santa Clara, CA (Mar 2016), ISBN 978-1-931971-29-4, URL https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/zhai

[80] Zhao, X., Sala, A., Wilson, C., Zheng, H., Zhao, B.Y.: Orion: shortest path estimation for large social graphs. networks **1**, 5 (2010)

[81] Zhao, X., Sala, A., Zheng, H., Zhao, B.Y.: Efficient shortest paths on massive social graphs. In: 7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), pp. 77–86, IEEE (2011)

[82] Zimmerman, P.: PGP user's guide (1994)