

Acyclic Preference Systems in P2P Networks

Anh-Tuan Gai, Dmitry Lebedev, Fabien Mathieu, Fabien De Montgolfier,
Julien Reynier, Laurent Viennot

► **To cite this version:**

Anh-Tuan Gai, Dmitry Lebedev, Fabien Mathieu, Fabien De Montgolfier, Julien Reynier, et al.. Acyclic Preference Systems in P2P Networks. Anne-Marie Kermarrec and Luc Bougé and Thierry Priol. Proceedings of the 13th International Euro-Par Conference - Euro-Par 2007, Aug 2007, Rennes, France. Springer, 4641, pp.825-834, 2007, Lecture Notes in Computer Science. <10.1007/978-3-540-74466-5_88>. <inria-00471720>

HAL Id: inria-00471720

<https://hal.inria.fr/inria-00471720>

Submitted on 8 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Acyclic preference systems in P2P networks

Anh-Tuan Gai¹, Dmitry Lebedev², Fabien Mathieu², Fabien de Montgolfier³,
Julien Reynier⁴, and Laurent Viennot¹

¹ INRIA, domaine de Voluceaux, 78153 Le Chesnay cedex, France

² FT R&D, 38–40, rue du général Leclerc, 92130 Issy les Moulineaux, France

³ LIAFA, Université Paris 7, 175 rue du Chevaleret, 75013 Paris, France

⁴ Motorola Labs, Parc Algorithmes Saint-Aubin, 91193, Gif sur Yvette Cedex, France

Abstract. In this work we study preference systems suitable for the Peer-to-Peer paradigm. Most of them fall in one of the three following categories: global, symmetric and complementary. All these systems share an acyclicity property. As a consequence, they admit a stable (or Pareto efficient) configuration, where no participant can collaborate with better partners than their current ones.

We analyze the representation of such preference systems and show that any acyclic system can be represented with a symmetric mark matrix. This gives a method to merge acyclic preference systems while retaining the acyclicity property. We also consider properties of the corresponding collaboration graph, such as clustering coefficient and diameter. In particular, the study of the example of preferences based on real latency measurements shows that its stable configuration is a small-world graph.

1 Introduction

Motivation In most current peer-to-peer (P2P) solutions participants are encouraged to cooperate with each other. Since collaborations may be costly in terms of network resources (connection establishment, resource consumption, maintenance), the number of connections is often bounded by the protocol. This constraint encourages the clients to make a careful choice among other peers to obtain a good performance from the system. The possibility to choose a better partner implies that there exists a *preference system*, which describes the interests of each peer.

The study of such preference systems is the subject of *b*-matching theory. It has started forty-five years ago with the seminal work of Gale and Shapley on *stable marriages* [1]. Although the original paper had a certain *recreational mathematics* flavor, the model turned out to be especially valuable both in theory and practice. Today, *b*-matching's applications are not limited to dating agencies, but include college admissions, roommates attributions, assignment of graduating medical students to their first hospital appointments, or kidney exchanges programs [1–4]. The goal of the present paper is to expand *b*-matching application domain to P2P networks by using it to model the interactions between the clients of such networks.

Previous work The present work draws on and extends results obtained in [5], where we covered general aspects of the b -matching theory application to the dynamics of the node interactions. We considered preference systems natural for the P2P paradigm, and showed that most of them fall into three categories: global, symmetric, and complementary. We demonstrated that these systems share the same property: acyclicity. We proved existence and uniqueness of a stable configuration for acyclic preference systems.

Contribution In this article, we analyze the links between properties of local marks and the preference lists that are generated with those marks. We show that all acyclic systems can be created with symmetric marks. We provide a method to merge any two acyclic preference systems and retain the acyclic property. Finally, our simulations show that real latency marks create collaboration graphs with small-worlds properties, in contrast with random symmetric or global marks.

Roadmap In Section 2 we define the global, symmetric, complementary, and acyclic preference systems, and provide a formal description of our model. In Section 3 we demonstrate that all acyclic preferences can be represented using symmetric preferences. We consider complementary preferences in Section 4, and the results are extended to any linear combination of global or symmetric systems. Section 5 discusses the properties of a stable solution providing an example based on Meridian project measurements [6]. In Section 6 we discuss the impact of our results, and Section 7 concludes the paper.

2 Definition and applications of P2P preference systems

2.1 Definitions and general modeling assumptions

We formalize here a b -matching model for common P2P preference systems.

Acceptance graph Peers may have a partial knowledge of the network and are not necessarily aware of all other participating nodes. Peers may also want to avoid collaboration with certain others. Such criteria are represented by an *acceptance* graph $G(V, E)$. Neighbors of a peer $p \in V$ are the nodes that may collaborate with p . A configuration C is a subset $\tilde{C} \subset E$ of the existing collaborations at a given time.

Marks We assume peers use some real marks (like latency, bandwidth, ...) to rank their neighbors. This is represented by a valued matrix of marks $m = \{m(i, j)\}$. A peer p uses $m(p, i)$ and $m(p, j)$ to compare i and j . Without loss of generality, we assume that 0 is the best mark and $m(p, i) < m(p, j)$ if and only if p prefers i to j . If p is not a neighbor of q , then $m(p, q) = \infty$. We assume for convenience a peer p has a different mark for each of its neighbors. It implies that a peer can always compare two neighbors and decide which one suits him better.

Preference system A mark matrix M creates an instance L of a preference system. $L(p)$ is a preference list that indicates how a peer p ranks its neighbors.

The relation when p prefers q_1 to q_2 is denoted by $L(p, q_1) < L(p, q_2)$. Note that different mark matrices can produce the same preference system.

Global preferences A preference system is global if it can be deduced from global marks ($m(i, p) = m(j, p) = m(p)$).

Symmetric preferences A preferences system is symmetric if it can be deduced from symmetric marks ($m(i, j) = m(j, i)$ for all i, j).

Complementary preferences A preferences system is complementary if it can be deduced from marks of the form $m(i, j) = v(j) - c(i, j)$, where $v(j)$ values the resources possessed by j and $c(i, j)$ the resources that i and j have in common⁵.

Acyclic preferences A preferences system is acyclic if it contains no preference cycle. A preference cycle is a cycle of at least three peers such that each peer strictly prefers its successor to its predecessor along the cycle.

Quotas Each peer p has a quota $b(p)$ (possibly infinite) on the number of links it can support. A *b-matching* is a configuration C that respects the quotas. If the quotas are greater than the number of possible collaborations, then simply $C = E$ would be an optimal solution for all.

Blocking pairs We assume that the nodes aim to improve their situation, i.e. to link to most preferred neighbors. A pair of neighbors p and q is a *blocking pair* of a configuration C if $\{p, q\} \in E \setminus C$ and both prefer to change the configuration C and to link with the each other. We assume that system evolves by discrete steps. At any step two nodes can be linked together if and only if they form a blocking pair. Those nodes may drop their worst performing links to stay within their quotas. A configuration C is *stable* if no blocking pairs exist.

Loving pair Peers p, q form a loving pair if p prefers q to all its other neighbors and q , in its turn, prefers p to all other neighbors. It implies a strong link which cannot be destroyed in the given preference system.

2.2 Preference systems and application design

Depending on the P2P application, several important criteria can be used by a node to choose its collaborators. We introduce the following three types as representative of most situations:

Proximity: distances in the physical network, in a virtual space or similarities according to some characteristics.

Capacity related: network bandwidth, computing capacity, storage capacity.

Distinction: complementary character of resources owned by different peers.

⁵ Of course, in this case the preferred neighbor has a larger mark.

Notice that these types correspond respectively to the definitions of symmetric, global and complementary preference system categories.

Examples of symmetric preferences are P2P applications which optimize latencies. A classical approach for distributed hash-table lists of contacts is selecting the contacts with the smallest round trip time (RTT) in the physical network. In Pastry [7], a node will always prefer contacts with the smallest RTT among all the contacts that can fit into a given routing table entry. More generally, building a low latency overlay network with bounded degree requires to select neighbors with small RTTs. Optimizing latencies between players can also be crucial, for instance, for online real-time gaming applications [8]. Such preferences are symmetric since the mark a peer p gives to some peer q is the same as the mark q gives to p (the RTT between p and q).

Similarly, massively multiplayer online games (MMOG) require connecting players with nearby coordinates in a virtual space [9, 10]. Again this can be modeled by symmetric preferences based on the distance in the virtual space. Some authors also propose to connect participants of a file sharing system according to the similarity of their interests [11, 12], which is also a symmetric relation.

BitTorrent [13] is an example of a P2P application that uses a capacity related preference system. In brief, a BitTorrent peer uploads to the peers from whom it has the most downloaded during the last ten seconds. This is an implementation of the well known Tit-for-Tat strategy. The mark of a peer can thus be seen as its upload capacity divided by its collaboration quota.

This global preference nature of BitTorrent should be tempered by the fact that only peers with complementary parts of the file are selected. Pushing forward this requirement would lead to another selection criterion for BitTorrent: preference for the peers possessing the most complementary set of file pieces. In other words, each peer should try to exchange with peers possessing a large number of blocks it needs. We call this a complementary preference system. Note, that this kind of preferences changes continuously as new pieces are downloaded. However, the peers with the most complementary set of blocks are those, which enable the longest exchange sessions.

In its more general form, the selection of partners for cooperative file download can be seen as a mix of several global, symmetric, and complementary preference systems.

3 Acyclic preferences equivalence

In [5], we have shown that global, symmetric and complementary preferences are acyclic, that any acyclic b -matching preference instance has a unique stable configuration, and that acyclic systems always converge toward their stable configuration. However, since acyclicity is not defined by construction, one may wonder whether other kinds of acyclic preferences exist. This section is devoted to answering this question.

Theorem 1. *Let P be a set of n peers, \mathcal{A} be the set of all possible acyclic preference instances on P , \mathcal{S} be the set of all possible symmetric preference instances on P , \mathcal{G} be the set of all possible global preference instances, then*

$$\mathcal{G} \subsetneq \mathcal{A} = \mathcal{S}$$

This ensures that any acyclic instance can be described by the mean of symmetric marks. As a special case, global mark instances can be emulated by symmetric instances, though the reverse is not true. In the remainder of this section we present the proof of theorem 1: we will first show $\mathcal{S} \subset \mathcal{A}$ and $\mathcal{G} \subset \mathcal{A}$, then $\mathcal{A} \subset \mathcal{S}$, which will be followed by $\mathcal{G} \neq \mathcal{A}$.

Lemma 1. *Global and symmetric preference systems are acyclic.*

Proof. from [5] Let us assume the contrary, and assume that there is a circular list of peers p_1, \dots, p_k (with $k \geq 3$), such that each peer of the list strictly prefers its successor to its predecessor. Written in the form of marks it means that $m(p_i, p_{i+1}) < m(p_i, p_{i-1})$ for all i modulo k . Taking the sum for all possible i , we get

$$\sum_{i=1}^k m(p_i, p_{i+1}) < \sum_{i=1}^k m(p_i, p_{i-1}).$$

If marks are global, this can be rewritten $\sum_{i=1}^k m(p_i) < \sum_{i=1}^k m(p_i)$, and if they are symmetric, $\sum_{i=1}^k m(p_i, p_{i+1}) < \sum_{i=1}^k m(p_i, p_{i+1})$. Both are impossible, thus global and symmetric marks create acyclic instances. \square

The next part, $\mathcal{A} \subset \mathcal{S}$, uses the loving pairs described in 2.1. We first prove the existence of loving pairs in Lemma 2.

Lemma 2. *A nontrivial acyclic preference instance always admits at least one loving pair.*

Proof. A formal proof was presented in [5]. In short, if there is no loving pair, one can construct a preference cycle by considering a sequence of first choices of peers.

Algorithm 1: Construction of a symmetric note matrix m given an acyclic preferences instance L on n peers

```

 $N := 0$ 
for all  $p$  and  $q$ ,  $m(p, q) = +\infty$  (by default, peers do not accept each other)
while there exists a loving pair  $\{i, j\}$  do
     $m(i, j) := m(j, i) := N$ 
    Remove  $i$  from the preference list  $L(j)$  and  $j$  from  $L(i)$ 
     $N := N + 1$ 

```

Lemma 3. *Let L be a preference instance. Algorithm 1 constructs a symmetric mark matrix in $O(n^2)$ time that produces L .*

Proof. The matrix output is clearly symmetric. Neighboring peers get finite marks, while others have infinite marks. If an instance contains a loving pair $\{i, j\}$ then $m(i, j) = m(j, i)$ can be the best mark since i and j mutually prefer each other to any other peers. According to Lemma 2 such a loving pair always exists in the acyclic case. By removing the peers i and j from their preference lists, we obtain a smaller acyclic instance with the same preference lists except that i and j are now unacceptable to each other. The process continues until all preference lists are eventually empty. The marks are given in increasing order, therefore when $m(p, q)$ and $m(p, r)$ are finite, $m(p, q) < m(p, r)$ iff the loving pair $\{p, q\}$ is formed before the loving pair $\{p, r\}$, that is iff p prefers q to r .

The algorithm runs in $O(n^2)$ time because an iteration of the **while** loop takes $O(1)$ time. A loving pair can especially be found in constant time by maintaining a list on all loving pairs. The list is updated in constant time since, after i and j becoming mutually unacceptable, each new loving pair contains either i and its new first choice, or j and its new first choice. \square

Not all acyclic preferences are global preferences. A simple counter-example uses 4 peers p_1, p_2, p_3 and p_4 with the following preference lists:

$L(p_1) : p_2, p_3, p_4$ $L(p_2) : p_1, p_3, p_4$ $L(p_3) : p_4, p_1, p_2$ $L(p_4) : p_3, p_1, p_2$

L is acyclic, but p_1 prefers p_2 to p_3 whereas p_4 prefers p_3 to p_2 . p_1 and p_4 rate p_2 and p_3 differently, thus the instance is not global.

4 Complementary and Composite Preference Systems

Complementary preferences appear in systems where peers are equally interested in the resources they do not have yet. As said in Section 2.1, complementary preferences can be deduced from marks of the form $m(p, q) = v(q) - c(p, q)$ (in this case, marks of higher values are preferred).

The expression of a complementary mark matrix m shows that it is a linear combination of previously discussed global and symmetric mark matrices: $m = v - c$, where v defines a global preference system and c defines a symmetric system.

Theorem 2 shows that complementary marks, and more generally any linear combination of global or symmetric marks, produce acyclic preferences.

Theorem 2. *Let m_1 and m_2 be global or symmetric marks. Any linear combination $\lambda m_1 + \mu m_2$ is acyclic.*

Proof. The proof is practically the same as for Lemma 1. Let us suppose that the preference system induced by $m = \lambda m_1 + \mu m_2$ contains a preference cycle $p_1, p_2, \dots, p_k, p_{k+1} = p_1$, for $k \geq 3$. We assume without loss of generality that m_1 is global, m_2 symmetric and that marks of higher values are preferred for m .

Then $m(p_i, p_{i+1}) > m(p_i, p_{i-1})$ for all i modulo k . Taking a sum over all possible i , we get

$$\begin{aligned} \sum_{i=1}^k m(p_i, p_{i+1}) &> \sum_{i=1}^k m(p_i, p_{i-1}) = \sum_{i=1}^k (\lambda m_1(p_i, p_{i-1}) + \mu m_2(p_i, p_{i-1})), \text{ but} \\ \sum_{i=1}^k (\lambda m_1(p_i, p_{i-1}) + \mu m_2(p_i, p_{i-1})) &= \lambda \sum_{i=1}^k m_1(p_{i-1}) + \mu \sum_{i=1}^k m_2(p_{i-1}, p_i) = \\ &= \lambda \sum_{i=1}^k m_1(p_{i+1}) + \mu \sum_{i=1}^k m_2(p_i, p_{i+1}) = \sum_{i=1}^k m(p_i, p_{i+1}). \end{aligned}$$

This contradiction proves the Theorem. \square

Theorem 2 leads to the question whether any linear combination of acyclic preferences expressed by any kind of marks is also acyclic. The example bellow illustrates that in general it is not true:

$$M_1 = \begin{pmatrix} 0 & 3 & 1 \\ 2 & 0 & 1 \\ 3 & 1 & 0 \end{pmatrix}, M_2 = \begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 3 \\ 1 & 2 & 0 \end{pmatrix}, M_1 + M_2 = \begin{pmatrix} 0 & 4 & 3 \\ 3 & 0 & 4 \\ 4 & 3 & 0 \end{pmatrix}$$

The preference instance induced by $M_1 + M_2$ has the cycle 1, 2, 3, while both M_1 and M_2 are acyclic (both produce global preferences).

Note, that a linear combination of two preference system matrices can give duplicates in the marks of a single node, which generates ties in preferences. Ties affect existence and uniqueness of a stable configuration, depending on how they are handled. If a peer prefers a new node to a current collaborator that has the same mark, existence is not guaranteed (but if a stable configuration exists, it is unique). Otherwise, existence stands, but not uniqueness [3].

Application Theorem 2 provides together with Theorem 1 a way of constructing a tie-less acyclic instance that can take into account several parameters of the network, given that they all correspond to acyclic preferences: the parameters can be first converted into integer symmetric marks using Algorithm 1. Then a linear combination using \mathbb{Q} -independent scalars produces distinct acyclic marks.

5 Graph properties of stable configurations

To illustrate the acyclic preferences, in this section, we study the connectivity property of stable configurations corresponding to different systems. In particular, we are using the latency matrix of the Meridian Project [6] as an example of the symmetric marks. The entries of the matrix correspond to the median of the round-trip time and they were measured using King technique [6]. We compare it with random symmetric marks and the global preference system⁶.

To confront these systems we examine connectivity properties of the corresponding stable configurations. The connectivity was extensively studied since

⁶ In absence of ties, all global marks are the same up to permutation

Watts survey [14] on the *small world* graphs. These graphs are known to have good routing and robustness properties. They are characterized by a small (i.e. $O(\log(n))$) mean distances and high (i.e. $O(1)$) clustering. The *clustering coefficient* is the probability for two vertices x and y to be linked, given that x and y have at least one common neighbor.

We consider three networks with different marks and $n = 2500$ peers. Figure 1 shows the properties of the stable configuration for these three marks, as a function of the quota b on the number of links per peer.




Fig. 1. Diameter and clustering coefficient of latency, random symmetric and global marks (2500 nodes) stable configurations. Global marks use an underlying Erdős-Rényi $\mathcal{G}(2500, 0.5)$ acceptance graph.

Global marks produce configuration with disconnected cliques of size $b + 1$ (maximal clustering, and infinite diameter). We have previously observed this *clusterization* effect in [15]. It can be lessened by using an Erdős-Rényi acceptance graph, as it can be seen in Figure 1. Then, the configuration still has a high clustering coefficient, and a high, but finite diameter (of same order of magnitude as $\frac{n}{b}$). This is due to a *stratification* effect: peers only link to peers that have marks similar to them [15].

Random symmetric matrix produces configurations with low diameter and clustering coefficient. Their characteristics are similar to those of Erdős-Rényi graphs.

Real latencies from [6] result in both a low diameter and a high clustering coefficient. This indicates that the corresponding stable configuration has *small-world* structure and, therefore, it enjoys nice routing properties. Nevertheless, it is not a *scale-free* network [16], because the degree distribution does not follow a power law (the degrees are bounded by b).

6 Discussion and future work

Stability Decision, whether a stable configuration is a good thing or not, depends on the characteristics and needs of practical applications. If continuous link alteration has a high cost (like in structured P2P networks), or if the stable configuration has appealing properties (like the small-world properties observed for the Meridian latency-based stable configuration), then it is interesting to let the system converge. On the other hand, we have observed that global marks result in a stable configuration with high diameter, which is an undesired feature in most cases. Moreover, some systems like gossip protocols [17] take advantage of constant evolution of the corresponding acceptance graph. In such cases, the eventual convergence would be harmful.

Convergence speed The convergence speed is an important characteristic, no matter if the stable solution is desired or not. In the first case, the application is interested in speeding up the process. In the second case, the slower possible speed is preferred instead. Although this question is out of the scope of the present work, our current experiments suggest that the convergence depends on many parameters: the preference system used, the acceptance graph, the activity of peers (details of peers' interaction protocol), the quotas and others. If we use as time unit the mean interval between two attempts of a given peer to change one of its neighbors, then preliminary results show that convergence is logarithmic at best, and polynomial at worst. We plan to provide a complete study on the influence of parameters. This should help understanding existing protocols and making them more efficient.

Dynamics of preference systems We have considered fixed acceptance graph and preference lists. In real applications, arrivals and departures modify the acceptance graph, along with the discovery of new contacts (a toy example is BitTorrent, where a tracker periodically gives new contacts to the clients). The preference system itself can evolve in time. For instance, latency can increase if a corresponding link has a congestion problem. A complementary preference system is dynamic by itself: as a peer gets resources from a complementary peer, the complementarity mark decreases.

All these changes impact the stable configuration of the system. The question is to know whether the convergence speed can sustain the dynamics of preferences or not. Fast convergence and slow changes allow the system to continuously adjust (or stay close) to the current stable configuration. Otherwise, the configurations of the system may always be far from a stable configuration that changes too often. The preferable behavior depends on whether stability is a good feature. This is an interesting direction for future work.

7 Conclusion

In this paper, we gave formal definitions for a b -matching P2P model and analyze the existence of a stable configuration with preference systems natural for P2P environment. The term stability in our case corresponds to Pareto efficiency of

the collaboration network, since the participants have no incentives to change such links. We have also showed that in contrast to systems based on intrinsic capacities, a latency-based stable configuration has small-world characteristics.

References

1. Gale, D., Shapley, L.: College admissions and the stability of marriage. *American Mathematical Monthly* **69** (1962) 9–15
2. Irving, R.W., Manlove, D., Scott, S.: The hospitals/residents problem with ties. In: SWAT '00: Proceedings of the 7th Scandinavian Workshop on Algorithm Theory, London, UK, Springer-Verlag (2000) 259–271
3. Irving, R.W., Manlove, D.F.: The stable roommates problem with ties. *J. Algorithms* **43**(1) (2002) 85–105
4. Roth, A.E., Sonmez, T., Utku Unver, M.: Pairwise kidney exchange. *Journal of Economic Theory* **125**(2) (December 2005) 151–188 available at <http://ideas.repec.org/a/eee/jetheo/v125y2005i2p151-188.html>.
5. Lebedev, D., Mathieu, F., Viennot, L., Gai, A.T., Reynier, J., de Montgolfier, F.: On using matching theory to understand P2P network design. In: INOC. (2007)
6. Meridian Project: <http://www.cs.cornell.edu/People/egs/meridian/>.
7. Rowstron, A., Druschel, P.: Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Lecture Notes in Computer Science* **2218** (2001) 329–350
8. Lin, Y.J., Guo, K., Paul, S.: Sync-ms: synchronized messaging service for real-time multi-player distributed games. In: Proc. of the 10th IEEE International Conference on Network Protocols. (2002)
9. Keller, J., Simon, G.: Solipsis: a massively multi-participant virtual world. In: Intern. Conf. on Parallel and Distributed Techniques and Applications. (2003)
10. Kawahara, Y., Aoyama, T., Morikawa, H.: A peer-to-peer message exchange scheme for large-scale networked virtual environments. *Telecommunication Systems* **25**(3) (2004)
11. Le Fessant, F., Handurukande, S., Kermarrec, A.M., Massoulié, L.: Clustering in peer-to-peer file sharing workloads. In: IPTPS. (2004)
12. Sripanidkulchai, K., Maggs, B., Zhang, H.: Efficient content location using interest-based locality in peer-to-peer systems. In: INFOCOM. (2003)
13. Cohen, B.: Incentives build robustness in bittorrent. In: P2PECON. (2003)
14. Watts, D.J.: *Small Worlds : The Dynamics of Networks between Order and Randomness* (Princeton Studies in Complexity). Princeton University Press (2003)
15. Gai, A.T., Mathieu, F., Reynier, J., De Montgolfier, F.: Stratification in P2P networks application to bittorrent. In: ICDCS. (2007)
16. Newman, M.E.J.: The structure and function of complex networks. *SIAM Review* **45**(2) (2003) 167–256
17. Allavena, A., Demers, A., Hopcroft, J.E.: Correctness of a gossip based membership protocol. In: PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing, New York, NY, USA, ACM Press (2005) 292–301