Node Participation in Ad Hoc and Peer-to-Peer Networks: A Game-Theoretic Formulation

Luiz A. DaSilva and Vivek Srivastava

Bradley Department of Electrical and Computer Engineering,
Virginia Polytechnic Institute and State University,
Alexandria, VA, USA – 22314.
{ldasilva, vivs}@vt.edu

Abstract

Ad hoc and peer-to-peer networks sometimes operate as voluntary resource sharing networks, relying on users' willingness to spend their own resources for the common good. As the costs of such resource sharing (what we call "node participation") outweigh the benefits perceived by the nodes, users are less likely to participate, compromising overall network goals. The contribution of this paper is to formalize some of the relevant tradeoffs as a first step toward the design of appropriate incentive structures. We formulate a game theoretic model for node participation and derive conditions that will lead to a socially desirable equilibrium. We also analyze the impact of threats posed by a rogue node in the network.

Keywords

Mobile ad hoc networks, game theory, peer-to-peer, resource sharing, incentives

1 Introduction

In this paper, we focus on networks where all nodes voluntarily perform services directly for one another, helping achieve a network-wide goal. Examples include grid computing, ad hoc, sensor, and peer-to-peer (P2P) networks[1]. We group such environments under the term *voluntary resource sharing networks*, where the resources being shared may include processing and forwarding capabilities, storage, files, and data aggregation. In many cases of interest, this cooperation happens in a distributed fashion, without a centralized controlling entity.

Participation in such environments is often voluntary, with users perceiving some benefit in contributing. For instance, grid computing may rely on users' perception of contributing towards a worthy goal by making available their idle CPU cycles for scientific research such as the SETI@Home project [12]. However, there are also costs to participating. For instance, in an ad hoc or sensor network, by forwarding packets for others a node may deplete its own limited energy resources. As the costs outweigh the perceived benefit, users are less likely to volunteer their resources, compromising the overall goal of the network.

The contribution of this paper is to analytically model node behavior in voluntary resource sharing networks and quantify the cost/benefit tradeoffs that will lead nodes to volunteer their resources. We adopt a game theoretic approach, due to its applicability to modeling conflict and cooperation among rational decision-makers [1]. Game theory is a branch of mathematics that provides a suite of analytical tools to analyze the behavior and the motivations for such behavior among rational entities [5]. A broader discussion of the applicability of game theory to the study of ad hoc networks can be found in [14].

It is intuitive that, whenever incentives do not occur naturally (e.g., through altruistic motives), artificial incentives must be offered (e.g., in the form of payments or virtual currency) to ensure node participation. We define the socially optimal outcome as the situation where all nodes are willing to make their resources available to others.

---

[1] We note that resource sharing in grid computing and ad hoc and sensor networks may be mandated, rather than voluntary, when all nodes are under the control of a single administrative entity. We do **not** include these environments under the classification of voluntary resource sharing networks.

There may be other socially desirable outcomes: for instance, enough nodes are willing to forward packets for others to maintain connectivity in an ad hoc network. A situation where no node is willing to participate is clearly undesirable.

The paper is structured as follows. We start by discussing some of the relevant literature on node participation and incentives. We then formulate a game theoretic model and derive conditions that will lead to a socially desirable equilibrium. We study such conditions under two different participation strategies adopted by the nodes. Next, we model the impact of threats imposed by a rogue node in the network. We conclude the paper by discussing practical considerations and directions for future work.

2 Related Work

Important parallels between peer-to-peer environments and ad hoc networks exist when considering the impact of selfish behavior ("free-riding") on achieving socially-desirable equilibria. In peer-to-peer, the effectiveness of the system depends on the willingness of individuals to advertise and contribute files; in ad-hoc networks, the network may become partitioned unless nodes are willing to forward packets for others. In either case, in the absence of incentives the equilibrium is for nodes not to contribute to the network [15] [13].

There has been recent research in modeling file sharing networks (such as enabled by Kazaa and Gnutella) using game theory. If the nodes in the file sharing network are assumed to be rational and homogeneous, the analysis leads to a Nash equilibrium in which nodes do not to share their files, and their best strategy is to only download files and allow zero uploads [6]. The result is not surprising, as most of the file sharing problems are modeled based on some variant of the prisoner's dilemma, which leads to socially non-optimal solutions [11] [8]. Note, however, that if this were the observed behavior of all the nodes participating in a peer-to-peer network, the network would cease to exist. [6] considers the presence of altruistic nodes (thereby some level of heterogeneity) in the network. In this heterogeneous network, not surprisingly, the Nash equilibrium is for the altruistic nodes to share their files, thereby leading to a better socially optimal state.

To achieve a socially optimal equilibrium for a network with homogeneous nodes, different incentive mechanisms have been proposed in the literature. These incentives include establishing and maintaining a reputation index for every node in the network [11] [7], incorporating a tit-for-tat behavior based on past history of the other peers' behavior [8], or providing virtual or real monetary incentives [3] [4].

It is interesting to note the significant overlap in the type of incentive mechanisms that have been suggested to achieve social optimality in peer-to-peer and wireless ad hoc networks. [9] and [2] suggest a reputation based scheme to invoke cooperation among nodes in an ad hoc network for relaying packets for one another in multi-hop communication. Also, tit-for-tat behavior based mechanisms [13] have been shown to be effective in solving the problem of misbehaving nodes in routing and forwarding. We also note that, since these incentive mechanisms require repeated interaction, it might be difficult to implement them effectively if the network exhibits high node mobility. Node mobility is a crucial consideration in repeated games, since it affects the chances of the nodes to play again with one another. It can improve the efficiency of the incentive mechanisms [15] or lead to better decision making by the nodes, as we will show in our analysis.

The contribution of our work is to quantify the tradeoffs between the costs incurred and benefits accrued from participation, as perceived by nodes in the network. This is a necessary step in the design of appropriate incentive structures.

3 Game-Theoretic Model

We model node participation in an ad-hoc network as a strategic-form game G, where N is the finite, non-empty set of players, $(S_j)_{j\in N}$ are the sets of actions available to each player, and $\{U_j(\mathbf{s}):\times_{i\in N} S_i \to \Re\}_{j\in N}$ are the utilities derived by each player when joint action **s** is taken by all players. (Note that we use bold notation for vectors.) In short, we can write:

$$G=(N,(S_j)_{j\in N},(U_j)_{j\in N}) \ . \quad \textbf{(1)}$$

We consider homogeneous actions to be available to all users: to share their resources ($s_j = 1$) or to refrain from sharing ($s_j = 0$). The joint action set is, therefore, $S = \times_{j \in N} S_j = \{0,1\}^n$, where $n = |N|$ (the cardinality of set N).

In such a game-theoretic formulation, the utility function is often the "weakest link," due to the difficulty in assessing tradeoffs as perceived by individual users. In this work, we adopt general, intuitive assumptions about the utility function, without attempting to completely characterize such functions. In particular, we consider a user's utility function to be the sum of two components:

$$U_j(\mathbf{s}) = \alpha_j(\mathbf{s}) + \beta_j(\mathbf{s}) \ . \tag{2}$$

- $\alpha_j(\mathbf{s}) = \alpha_j(\sum_{i \in N, i \neq j} s_i)$ is the benefit accrued by a user from others' sharing of their resources. We assume $\alpha_j(0) = 0$ and $\alpha_j(\mathbf{s}) > 0$ if $\exists k \neq j$ such that $s_k \neq 0$, as it is intuitive that a user will accrue non-negative benefit from others' willingness to perform services for it.

- $\beta_j(\mathbf{s}) = \beta_j(s_j)$ is the benefit (or cost) accrued by sharing one's own resources with others. This may be negative, since there may be a cost to participating in the network (such as faster depletion of a node's energy resources); it may also be positive, if there exist financial incentives for participation or if the user derives satisfaction in doing so. In either case, we assume this part of the utility functions to be dependent only on the node's own chosen strategy. (Note that in peer-to-peer networks the cost of sharing ones' resources may depend on how many other nodes also share, as this affects the number of file requests received by each node. Similarly, in ad hoc networks, the number of routing requests may increase if few nodes in the network are willing to forward packets. These effects are not captured in our current model.) Also, either way, $\beta_j(0) = 0$.

The Nash equilibrium is considered a consistent prediction of the outcome of a game. A joint strategy **s** is a Nash equilibrium (NE) if no user can benefit from unilaterally deviating. If $s_j > 0 \Rightarrow \beta_j(s_j) < 0 \ \forall j$, then the only Nash equilibrium is for no nodes to participate. In other words, in the absence of incentives there is no voluntary resource sharing. We illustrate with a concrete example.

Consider $|N| = 3$, $\alpha_j(\mathbf{s}) = \sum_{i \in N, i \neq j} s_i$ and $\beta(\mathbf{s}) = -1.5 s_j$. The choice of utility function here is arbitrary, as the same conclusions hold for any functions obeying the assumptions outlined above. Similarly, the number of participants in the game is chosen for ease of visualization and does not affect the results for any n > 1. For this game, the utilities accrued by each player for every strategy profile are tabulated in Figure 1. The reader can verify that the only NE is (0,0,0); clearly, this is an inefficient outcome, as (1,1,1) would be a Pareto optimal strategy.

$s_3 = 0$

	$s_2 = 0$	$s_2 = 1$
$s_1 = 0$	0,0,0	1,-1.5,1
$s_1 = 1$	-1.5,1,1	-0.5,-0.5,2

$s_3 = 1$

	$s_2 = 0$	$s_2 = 1$
$s_1 = 0$	1,1,-1.5	2,-0.5,-0.5
$s_1 = 1$	-0.5,2,-0.5	0.5,0.5,0.5

Fig.1. Concrete example. Boxes are labeled with utility values for users 1, 2 and 3, respectively, for each possible joint strategy in $\{0,1\}^3$.

4 Repeated Games

The prisoner's dilemma is probably the most well known example in game theory. Similarly to the example above, it achieves a non-optimum equilibrium when played once. However, other equilibria are achievable when the game is repeated, provided that players do not know a priori how many repetitions of the game there will be. This provides the inspiration for the development below.

Consider a repeated game, played K times, where K is a geometrically distributed discrete random variable with parameter $0 < p < 1$. Therefore, $P[K = k] = p(1-p)^k, k = 0,1,2,...$ and $E[K] = \frac{1-p}{p}$. Note that, as $p \to 1$, the probability that the game will be repeated approaches 0. The geometric distribution is chosen for its memoryless property.

4.1 Grim Trigger Strategy

Consider a grim-trigger strategy [5] adopted by all nodes: share as long as all other nodes share; do not share if any of the others have deviated in the previous round. The trigger is activated when any one node decides to switch from the desired behavior (sharing resources) and is grim as the nodes do not switch their action back once the punishment (not sharing) is initiated.

If no node deviates, at any time a user's expected payoff from that point forward is:

$$[\alpha_i(N-1) + \beta_i(1)] \cdot [1 + \sum_{k=0}^{\infty} pk(1-p)^k] = \frac{\alpha_i(N-1) + \beta_i(1)}{p} \quad . \tag{3}$$

If, on the other hand, a node unilaterally deviates, its expected payoff from that point forward is simply $\alpha_i(N-1)$. So, it is a Nash equilibrium for all nodes to participate as long as, $\forall i$,

$$\alpha_i(N-1) > \frac{-\beta_i(1)}{1-p} \quad . \tag{4}$$

We offer an interpretation of this result. If $\beta_i(1) > 0$, i.e., if the user derives some benefit or satisfaction from sharing her own resources with others, then, not surprisingly, it is always an equilibrium to participate. More interestingly, when $\beta_i(1) < 0$ (i.e., there is a cost in sharing one's resources), then a socially optimal equilibrium is still sustainable. The precise cost/benefit tradeoff is given by the inequality above.

We also note that, in an ad hoc network, the time horizon for the repetitions of a game (characterized by parameter p) can be interpreted as having a direct relationship with a user's mobility. In this sense, the more mobile users are, the less incentive there is to share one's resources (the closer p is to 1).

The results above assume an all-or-nothing policy: if any node deviates (refuse to share) in one round, all others will deviate in the next round. In the next section, we explore the robustness of a "softer" policy that does not require *all* nodes to share their resources.

4.2 An Alternative Strategy

Let us denote by $s_i^{(k)}$ the strategy adopted by node i in the k^{th} round of the game. Suppose a node adopts the following strategy: in every round *k* of the repeated game, a node decides to share its resources as long as the following condition is satisfied:

$$\alpha_j\left(\sum_{i \in N, i \neq j} s_i^{(k-1)}\right) > \frac{-\beta_j(1)}{1-p} \quad . \tag{5}$$

We study the stability of this desirable equilibrium in relation to the probability of the game being repeated. We perform a simulation to determine whether a deviation in a single node's strategy (due to variation in nodes' perceptions of whether the game is likely to be repeated) will result in a cascading effect on the other nodes and lead to a shift to an equilibrium where all nodes decide not to share their resources.

It is clear that if nodes do not believe the game will be repeated (*p = 1)* the game reduces to a single stage game and all nodes shift their strategy. However, we wish to determine, as a function of the number of nodes in the network, the value of *p* at which the cascading behavior is observed and the equilibrium shifts for every node. In the simulation, for a fixed value of cost incurred by a node in sharing (i.e., $\beta_i(1) = \beta$), we vary the number of nodes in the system and establish what values of *p* will still support a socially desirable equilibrium. In our simulation, each node accrues a different benefit that it derives from other node's willingness to share resources. The function $\alpha_j(\mathbf{s}) = A_j \sum_{i \in N, i \neq j} s_i$, with A_j being taken from a uniform distribution between [0,1]. However, the cost incurred by each node in sharing is same. The simulation is repeated for different values of β. Each point plotted is the average of 200 repetitions of the simulation.

The plot in Figure 2 can be interpreted as follows. For each value of β, we plot the maximum value of p that will still lead to a desirable equilibrium. When more nodes are present, the desirable equilibrium is more robust to players' exogenous beliefs about the repeatability of the game (and, in a practical interpretation for ad hoc networks, that equilibrium is more robust to node mobility). Also, as the cost of participation (β) increases, the desirable equilibrium requires players to believe that the game has a high probability of being repeated (corresponding to a low value of p). As discussed before, $p = 1$ corresponds to a single stage game; the node knows that the game will not be repeated (at least, not with the same neighbors), always leading to non-cooperation. It is interesting to note, however, that if the number of players is high enough (i.e., if the network is dense enough), a socially desirable equilibrium is achievable even for values of p arbitrarily close to 1.

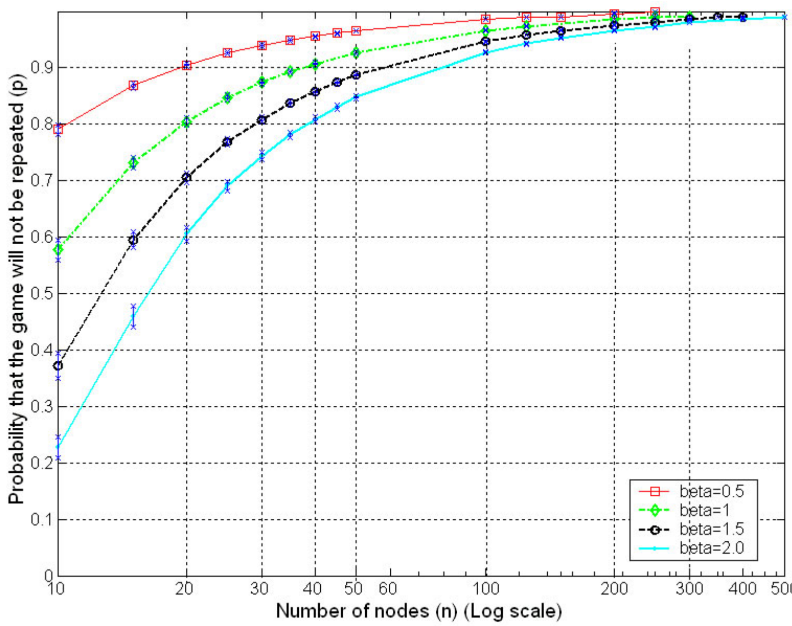

Fig. 2. Depending on node density, desirable equilibria are achieved even if nodes think there is a low likelihood of the game being repeated (corresponding to $p \rightarrow 1$)

5 Rogue Nodes

We now build upon this game to consider the risk posed by rogue nodes in the network. Some examples of the penalty that regular nodes in a resource sharing network may suffer from the presence of rogue nodes are enumerated in Table 1. Since we consider the presence of different kinds of nodes in the network, this can be modeled as a game of incomplete information, where the identity of rogue nodes in the network is not known a priori by other nodes. We introduce parameter T_j to characterize the identity of different types of nodes in the network. N is the set of players, $(S_j)_{j \in N}$ the action set, $(T_j)_{j \in N}$ describes the possible types of player j, and $\{p_j(\cdot \mid t_j): \times_{i\in N-j} T_i \rightarrow [0,1]^{n-1}\}_{j\in N}$ are conditional probabilities ascribed by player j to the types of other players, given that player j is of type t_j. The utility functions are now a function of both the collective actions of players and their types, as denoted by $\{U_j(\mathbf{s},t)\}_{j\in N}$. The Bayesian game Γ [10] can be expressed as:

$$\Gamma = (N,(S_j)_{j\in N},(T_j)_{j\in N},(p_j)_{j\in N},(U_j)_{j\in N}) \ . \tag{6}$$

We consider that nodes in a network can be of two types ($T_j = \{0,1\}\ \forall\, j \in N$): *regular* nodes ($t_j = 0$) accrue benefits from participation as described in the previous game we considered; *rogue* nodes ($t_j = 1$) accrue benefits from others' active participation in sharing their own resources. The presence of rogue nodes in the network decreases the utility of regular nodes. Examples include any node that is responsible for the threats summarized in Table 1. For concreteness, we assume regular nodes accrue a fixed reward R > 0 from sharing their resources when no rogue nodes are present, and a fixed penalty P < 0 when rogue nodes are present. We summarize the components of the utility function of regular nodes below:

For $\{j : t_j = 0\}$:

$$\alpha_j(\mathbf{s},t) = \alpha_j(\mathbf{s}) = \alpha_j(\sum_{i\in N, i\neq j} s_i) \ . \tag{7}$$

$$\beta_j(\mathbf{s},t) = \begin{cases} P & s_j = 1 \ and\ \exists k \in N-j \ such\, that\ t_k = 1 \\ 0 & s_j = 0 \\ R & otherwise \end{cases} \ . \tag{8}$$

The interesting question becomes: what values of reward and penalty and what sets of beliefs by player j will lead her to decide to share her resources with the network?

We assume that rogue nodes know about the presence of regular nodes in the network, but not vice-versa, and let θ_j denote the probability ascribed by player j that there is a rogue node in the network. It is reasonable to assume that it is a dominant strategy for rogue nodes to always share their own resources, for instance so that they can "blend in" better with regular nodes[2]. Let us then consider the expected utilities of a regular node j:

$$\begin{aligned} & E[U_j(\mathbf{s},t) \mid s_j = 0] = \alpha_j(\mathbf{s}) \\ & E[U_j(\mathbf{s},t) \mid s_j = 1] = \theta_j(\alpha_j(\mathbf{s}) + P) + (1-\theta_j)(\alpha_j(\mathbf{s}) + R) \ . \\ & = \alpha_j(\mathbf{s}) + P\theta_j + (1-\theta_j)R \end{aligned} \tag{9}$$

A benefit for sharing one's own resources exists if and only if:

$$P\theta_j + (1-\theta_j)R > 0 \Rightarrow \left|\frac{R}{P}\right| > \frac{\theta_j}{1-\theta_j} \ . \tag{10}$$

Therefore, the higher the probability ascribed by a player to the presence of a rogue node in the network, the larger the reward needs to be with respect to the penalty, for the player to share her resources.

Social welfare in voluntary resource sharing networks is maximized when all nodes volunteer their resources. The inequality above implies that welfare is maximized when

[2] This assumption is not required for the results we derive to hold; however, they simplify the derivation that follows.

$$\left|\frac{R}{P}\right| > \max_{\{j \in N | t_j = 0\}} \left(\frac{\theta_j}{1-\theta_j}\right) \quad . \tag{11}$$

We are currently exploring mechanisms for detecting the presence of rogue nodes, such as distributed trust management schemes. Robust mechanisms to this end will in effect reduce θ_j in our model and decrease the incentives (reward R) required to lead nodes to participate.

6 Conclusions and Future Work

A non-trivial conclusion from the results presented here is that nodes may agree to share their resources even if they perceive a cost in doing so. This happens as the nodes recognize that refusing to participate will result in similar behavior by others, which ultimately would compromise the viability of the network as a whole.

Clearly, a game theoretic model does not completely capture all aspects of node participation in a real ad hoc network. However, it provides useful insight into incentive mechanisms that are needed to induce node participation. It also opens up some important questions, which are the subject of our current research:

- If one's decision to participate is dependent on other nodes' behavior, how does a given node reliably assess other node's decision to make their resources available?
- Mobility may influence the incentives necessary for participation. In our model, this is captured by the parameter p, which expresses a node's belief that the game will be repeated with the same neighbor. A better understanding of the effects of mobility is needed.
- In formulating the impact of the threat of rogue nodes on other nodes' behavior, we assumed exogenous beliefs. In other words, regular nodes have pre-established beliefs about whether rogue nodes are present. We are exploring distributed trust management mechanisms that would result in nodes' ability to isolate uncooperative or malicious participants in the network.

We are also in the process of extending the model presented here to consider individual strategy sets $S_j = [0,1]$, i.e., a continuum that allows for partial participation by a node. As mentioned earlier in the paper, we also investigate utility functions where the cost of participation $\beta_j(\mathbf{s})$ depends on the joint actions by all players, rather than only on node j's action.

Acknowledgements

This work is partially supported by a grant from the Office of Naval Research.

References

[1] Robert Axelrod, "The Evolution of Cooperation," *Basic Books,* Reprint edition, New York, 1984.
[2] S. Buchegger and J.Y. Le Boudec, "Performance analysis of the CONFIDANT protocol: cooperation of nodes – fairness in dynamic ad-hoc networks," *Proc. ACM MobiHoc*, June 2002.
[3] L. Buttyan and J. P. Hubaux, "Nuglets: A virtual currency to stimulate cooperation in self organized mobile ad-hoc networks," Swiss Federal Institute of Technology, Lausanne, Switzerland, Report no. DSC /2001/001, January 2001.
[4] R. Buyya and S. Vahzkudai, "Compute Power Market: Towards a Market-Oriented Grid," *Proc. 1st IEEE/ACM Int'l Symp. Cluster Computing and the Grid* (CCGRID 01), May 2001, pp. 347-354.
[5] Prajit K. Dutta, "Strategies and Games: Theory and Practice," *MIT Press*, 1999.
[6] P. Golle, K. Leyton-Brown, I. Mironov and M. Lillibridge, "Incentives for sharing in peer-to-peer networks," *Proc. 2nd International Workshop on Electronic Commerce*, November 2001.
[7] H.T. Kung and C-H. Wu, "Differentiated admission for peer-to-peer systems: incentivizing peers to contribute their resources," *Proc. Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 2003.
[8] K. Lai, M. Feldman, I. Stoica and J. Chuang, "Incentives for cooperation in peer-to-peer networks," *Proc. Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 2003.

[9] P. Michiardi and R. Molva, "Game theoretic analysis of security in mobile ad hoc networks," *Institute Eurecom*, Sophia-Antipolis, France, Report no.RR-02-070, April 2002.
[10] R. B. Myerson, *Game Theory: Analysis of Conflict*, Harvard University Press, 1991.
[11] K. Ranganathan, M. Ripeanu, A. Sarin and I. Foster, "To Share or not to Share: an Analysis of Incentives to Contribute in Collaborative File Sharing Environments," *Proc. Workshop on Economics of Peer to Peer Systems*, Berkeley, CA, June 2003.
[12] "Search for Extraterrestrial Intelligence at Home," 2003; http://setiathome.ssl.berkeley.edu/
[13] V. Srinivasan, P. Nuggehalli, C.F. Chiasserini and R.R. Rao, "Cooperation in wireless ad hoc networks," *Proc. IEEE INFOCOM,* vol.2, April 2003, pp. 808-817.
[14] V. Srivastava, J. Neel, A. MacKenzie, J. Hicks, L.A. DaSilva, J. Reed and R.P. Gilles, " Employing game theory to analyze software radio enabled wireless ad hoc networks," submitted for publication.
[15] A. Urpi, M. Bonuccelli and S. Giordano, "Modeling cooperation in mobile ad hoc networks: a formal description of selfishness," *Proc. Workshop on Modeling and Optimization in Mobile and Wireless Ad Hoc Networks,* March 2003.

Table 1 - Benefits and threats from voluntary resource sharing.

Environment	Benefit perceived by participating nodes	Cost to participating nodes	Threats from the presence of rogue nodes
Grid computing	Societal benefits (e.g., research advances); financial incentives (e.g., [3])	Usage of CPU cycles affecting the performance of the node's own applications	Compromising of the integrity and secrecy of local data; protection of local data and processing resources from unauthorized access
Ad hoc networks	Enabling of multi-hop communications	Increased energy consumption leading to a reduction in node lifetime	Increased likelihood of detection in hostile environments
Sensor networks	Increased confidence in sensed information; aggregation of data	Increased energy and bandwidth consumption	Interception of critical or confidential information
Peer to peer	Distribution of information; trade of music/video files	Sharing of bandwidth and disk space	Collection of personal information for marketing purposes