

“Token” Equilibria in Sensor Networks with Multiple Sponsors*

David A. Miller
University of California, San Diego
Economics Department
9500 Gilman Dr. #0508, La Jolla, CA 92093
d9miller@ucsd.edu

Sameer Tilak
Binghamton University and
San Diego Supercomputing Center
9500 Gilman Dr. #0505, La Jolla, CA 92093
tilak@sdsc.edu

Tony Fountain
San Diego Supercomputing Center
9500 Gilman Dr. #0505, La Jolla, CA 92093
fountain@sdsc.edu

October 28, 2005

Abstract

When two sponsoring organizations, working towards separate goals, can employ wireless sensor networks for a finite period of time, it can be efficiency-enhancing for the sponsors to program their sensors to cooperate. But if each sensor privately knows whether it can provide a favor in any particular period, and the sponsors cannot contract on ex post payments, then no favors are performed in any Nash equilibrium. Allowing the sponsors to contract on ex post payments, we construct equilibria based on the exchange of “tokens” that yield significant cooperation and increase expected sponsor payoffs. Increasing the sponsors’ liability is beneficial because it enables them to use more tokens.

1 Introduction

First generation Wireless Sensor Networks (WSNs) are primarily application-specific in terms of their infrastructure, protocol/architecture, and user set, and typically all sensors within a given network are assumed to be deployed and controlled by a single administrative entity, or (equivalently) by a mutually altruistic consortium. The first generation networks are self configuring—able to configure themselves to collaborate with other sensors on their own network—but they are not generally able to interoperate

with other sensor networks, or to make their resources available to different clients.

The success and increasing deployment of WSNs is motivating a second generation that allows a diverse set of clients to learn about and access resources spanning multiple networks. Clients compose end-to-end services using the services provided by individual networks. If sensors that belong to different organizations can cooperate with each other to provide various services like sensing, routing, data processing, and storage, overall utility to a diverse set of end users can be significantly improved. Since different networks may be deployed and managed by independent organizations, they may have different protocols, architectures, security policies, and pricing mechanisms. When the WSNs of multiple sponsoring organizations interact, strategic considerations become salient.

Resource management in such large scale, multi-sponsor distributed systems is a daunting task. Game theory brings two important benefits to bear on such problems. First, it provides a set of tools for constructing and evaluating equilibria in WSNs, where each sensor uses locally available information to inform its behavior. Second, it provides a method for analyzing the human and institutional motives that underly the use of WSNs.

This second benefit is the main focus of this study, where we consider an network made up of sensors belonging to two different sponsoring organizations. Each sponsor programs its own sensors to further its own objectives, but both sponsors can gain if their sensors share information and resources in the field. In a simple model of this type of situation, we construct an equilibrium in the WSN that yields a

*Forthcoming in the Proceedings of the Workshop on Stochasticity in Distributed Systems (StoDiS’05), San Jose, CA, December 19, 2005. Copyright IEEE.

significantly higher payoff to the sponsors than they could attain separately. Further, this equilibrium in the WSN corresponds to a Nash equilibrium in the game between the sponsors; i.e., neither sponsor can gain by programming its sensors differently.

Once the sensors are in the field, they can request favors from one another. (We use the words “favors” and “services” interchangeably.) Canonical examples of services are routing, data storage, and data aggregation. The benefit of receiving a favor is greater than the cost of providing one, so for efficiency favors should always be provided. However, whether or not each sensor has the ability to provide a favor varies stochastically over time, in a way that is undetectable to other sensors. Hence a sensor from one sponsor could “claim” to be unable to provide a favor to a sensor from another sponsor, and by doing so save on the costs of providing the favor. Thus a sponsor could potentially gain by deviantly programming its sensors to request favors, but not provide them. This is similar to the classic prisoners’ dilemma problem, and in a finite horizon setting without contracts, the unique Nash equilibrium is that neither sponsor ever allows its sensors to provide favors to the other sponsor.

The problem, then, is to construct an equilibrium in which both sponsors are willing to program their sensors to cooperate. In the equilibrium we construct, each sponsor starts with an agreed-upon number of tokens, which it distributes among its sensors. When a sensor from one sponsor requests a favor from another sponsor’s sensor, it offers a token in exchange for the favor. If the requestee is able to provide the favor, it does so in return for the token.

If a sensor runs out of tokens, it can no longer request favors. So the more tokens are available to the sensors, the more favors will be performed and the more efficient will be the equilibrium. We consider a game with a fixed, finite horizon, but we allow the sponsors to write a contract at the outset of the game that can obligate them to make payments to each other that depend on the distribution of tokens in the final period. Each sponsor has limited liability, so its *ex post* payment is uniformly bounded across all realizations of the final distribution of tokens. Since the equilibrium value of a token is determined endogenously, this exogenous bound yields an endogenous limit on the number of tokens available.

Related literature There is a substantial literature applying game theoretic models and tools to network problems, particularly relating to the internet, for example applying incentive compatible mechanism design to distributed computing [4] and community resource sharing [7, 10], or considering more abstract issues [11]. Specifically for the setting of WSNs, recent work has applied game theoretic tools to models involving a single sponsor, such as load-balanced

tree construction for data gathering [12] and sensor role assignment [3]. Other work has investigated network architecture when each node is controlled by a different strategic agent [13]. We depart from the existing literature by considering a higher level problem: that of coordinating groups of sensors sponsored by multiple organizations. Rather than evaluating tradeoffs between specific favors (e.g., routing vs. sensing), we consider abstract favors.

We make use of two classic results from the game theory literature. First, as a benchmark we employ a theorem from [2], which implies that, when the sponsors cannot write a contract before programming their sensors, the unique Nash equilibrium outcome is that the sponsors do not program their sensors to cooperate.

Second, our result is in the spirit of the classic “folk theorem,” which states that any level of cooperation can be supported in an equilibrium of an infinitely repeated game if the players are sufficiently patient. Recent work has extended the folk theorem to games with private information (in particular, [5] and [8]). Still, no folk theorem applies to our setting, since costly communication is required to implement cooperation. But the equilibria we construct are similar to the types of equilibria that are used to prove the folk theorem.

Several authors have recently considered games in which strategic agents trade favors when their ability to perform favors is private information. Our token equilibria are based on [9], which constructs a token equilibrium in a two-player favors game; [9] also constructs simpler equilibria that induce cooperation among a network of players. [6] constructs optimal equilibria in a continuous-time two-player favors game, and [1] constructs cooperative equilibria in a more complex game with both favors and stochastic investments. Although in contrast to these authors we consider a finite horizon environment, we allow the sponsors to write a contract compelling them to pay monetary transfers based on what occurs during the game. This means that the strategic considerations in our setting are quite similar to those of infinite horizon environments.

2 The Model

Two sponsoring organizations, $i \in \{A, B\}$, each employ K sensors, s_{i1}, \dots, s_{iK} , on a rectangular grid with $2K$ nodes. Each sensor can communicate only with its immediate neighbors (vertically and horizontally, not diagonally). There are $t = 1, \dots, T$ periods. At the start of a period, each sensor learns the time during the period when it will need a favor. (The distribution from which the timing of these needs should be represented by a bounded probability density function. For simplicity, we use the uniform distribution over the length of a time period.) Also at the beginning of the period, each sensor independently learns its stochastic ability to provide favors to other sensors: with

probability λ it can provide favors in period t , and with probability $1 - \lambda$ it cannot. Favors can be provided only to immediate neighbors.

The period then plays out in real time. When its need arises, each sensor in turn can request a favor from one of its neighbors; the process of requesting and receiving a favor will be described in the subsequent section. The ‘‘payoff’’ for sensor s_{ik} in period t is

$$u_{ik}(t) = \alpha R_t - \beta P_t - \gamma C_t,$$

where α , β , and γ are positive parameters; R_t is the number of favors received by s_{ik} ; P_t is the number of favors provided by s_{ik} ; and C_t is the number of communication signals sent by s_{ik} . The utility for a sponsor in the entire game is

$$U_i = \sum_{t=1}^T \sum_{k=1}^K u_{ik}(t) + \tau_i,$$

where τ_i is the monetary transfer received by sponsor i after the end of period T . The monetary transfers must be zero sum ($\tau_A = -\tau_B$) and are bounded above by $\bar{\tau} < \infty$. Throughout, we assume $\gamma/\lambda < \alpha - \beta$, so the sponsors can gain by cooperating.

The equilibrium concept we employ is Nash equilibrium between the sponsors, even though the events of the game play out dynamically. In dynamic games, theorists usually employ subgame perfect equilibrium or one of its refinements, to rule out the possibility that a strategic agent could threaten to take actions that are not sequentially optimal. In our setting, however, all strategic activity takes place simultaneously at the outset of the game, when the sponsors write a contract (see below) and program their sensors. Since the sensors themselves are automata rather than strategic agents, it is reasonable in this context for them to be able to commit to threats that may not be sequentially optimal. This actually makes our analysis more difficult, because we cannot employ the kind of backward induction arguments that are typically used to analyze dynamic games.

Token equilibria At the outset of a token equilibrium, the sponsors jointly agree on a number of tokens, n_i for each sponsor i to distribute among its sensors. They can also write a legally binding contract that obliges them to make payments at the end of period T that depend on the data stored in the sensors. For simplicity, we assume that this data is complete and verifiable. (This assumption could be weakened by requiring tokens to be encrypted and requiring sensors to confirm each transaction, at the cost of additional communication.) When the sponsors program their sensors according to the equilibrium, the payments need depend only on the number of tokens held by each sponsor at the end of period T .

When (for example) sensor s_{A1} has a token and needs a favor, it chooses randomly among its neighbors; say it chooses sensor s_{B1} . Then s_{A1} sends a request to s_{B1} , and tentatively forwards a token along with the request. If s_{B1} provides the requested favor, then s_{B1} obtains the token; otherwise s_{A1} retains the token. The number of communication signals sent by a particular sensor is the total number of requests it made. (If s_{B1} is unwilling or unable to provide a favor, it does not incur a communication cost when s_{A1} sends its request. If it does provide a favor, the communication cost of doing so is incorporated into β .)

3 Results

Proposition 1. *Suppose that $T = 1$ and the sponsors cannot write a contract. Then the unique Nash equilibrium outcome is that no sensor requests a favor. Equilibrium utilities are $U_A = U_B = 0$.*

Proof. If all sensors are programmed never to provide favors, then it must also be that all sensors are programmed never to request favors, since requesting favors is costly. So suppose that some sensors are programmed to provide favors. At a point in time during the period, consider a sensor whose own need for a favor has already passed, which has not yet granted any favor, and which receives a request for a favor. (In a Nash equilibrium there is no restriction on the behavior of a sensor in states of the world that, its sponsor, by means of its own programming, prevents from being reached; i.e., if the sensor is programmed not to grant a first favor, then it does not matter whether it is programmed to grant a second favor.) It should not grant the favor if the cost of doing so outweighs the benefit, and the only possible benefit is the prospect that one of its siblings might subsequently receive a favor in turn. (Two sensors are called *siblings* if they belong to the same sponsor.) Given the time remaining in the period, the probability that any of its siblings will subsequently have a need for a favor is less than $(K - 1)q$. Thus the benefit of performing a favor is less than the cost if $q \leq q^* \equiv \beta/(K - 1)\alpha$, and so along any Nash equilibrium path each sensor must be programmed to ignore any favor request that arrives after both the sensor’s need for a favor and a fraction $1 - q^*$ of the period have passed.

Now consider whether a sensor should request a favor after at least a fraction $1 - q^*$ of the period has passed, and it has not yet granted any favor. Along the equilibrium path, the probability that any of its neighbors can be programmed to grant favors at such a time is less than $4q$, and so the benefit of requesting a favor is less than the cost if $q \leq \hat{q} \equiv \min\{\gamma/4\alpha, q^*\}$. Hence along any Nash equilibrium path no sensor will request a favor after a fraction $1 - \hat{q}$ of the period has passed. Now apply the same reasoning as if the period

were only $1 - \hat{q}$ as long, and continue inductively to obtain the result. \square

Proposition 2. *Suppose that $T < \infty$ and the sponsors cannot write a contract. Then the unique Nash equilibrium outcome is that no sensor requests a favor. Equilibrium utilities are $U_A = U_B = 0$.*

Proof. The unique Nash equilibrium outcome in Proposition 1 yields the minimax utility for each sponsor. Apply Proposition 0 of [2] to obtain the result. \square

Proposition 3. *Now suppose that $T < \infty$ and the sponsors can write a contract as described in Section 2. If $\bar{\tau} \geq \beta/2$, then there exists a token equilibrium in which the expected sum of equilibrium utilities is positive.*

Proof. Suppose sponsor A begins with 1 token, sponsor B begins with none, and the contract binds the sponsor without the token at the end of period T to pay an amount x to the other sponsor. Suppose A programs each of its sensors to wait until period T , and then request a favor if it has the token; and B programs its sensors to wait until period T , and then provide a favor if able when a token is offered. Then B receives expected utility of $\lambda(x - \beta) - (1 - \lambda)x$, compared to $-x$ if it programmed its sensors to ignore such requests. Hence B 's strategy is a best response to A 's strategy if $x \geq \beta/2$. In turn, A receives expected utility of $\lambda(\alpha - x) + (1 - \lambda)x - \gamma$, compared to x if it programmed its sensors not to make such requests. Hence the sponsors' strategies are mutual best responses if $\alpha - \frac{\gamma}{\lambda} \geq 2x \geq \beta$; e.g., $x = \beta/2 \leq \bar{\tau}$. Finally, neither sponsor, acting unilaterally, can gain by programming its sensors to make or respond to requests in earlier periods. The expected sum of equilibrium utilities is $\lambda(\alpha - \beta) - \gamma > 0$. \square

Proposition 4. *Suppose that $T < \infty$ and the sponsors can write a contract as described in Section 2. Then there exists a token equilibrium in which, along the equilibrium path, each sensor requests a favor if and only if it holds a token when its need for a favor arises, and each sensor grants a favor if it is able to do so when it receives a request. No more than $2\bar{\tau}/\beta$ tokens that can be accommodated in such an equilibrium.*

Proof. Suppose the contract binds sponsor A to pay $(p_B - p_A)x$ to sponsor B , where p_i is the number of tokens held by i 's sensors at the end of period T , and $x = \beta/2$. Suppose that A programs its sensors as described. Consider sensor s_{Bk} in period t , and suppose it receives a request for a favor. As in Proposition 3, the cost of providing the favor is outweighed by the benefit, which is at least x (since the token can at least be held until the end of period T), B 's best response is to program its sensors to provide favors upon request, if able.

Now suppose that sensor s_{Bk} holds one or more tokens in period t . Let z_t be the value of s_{Bk} 's marginal token in period t , when s_{Bk} employs the token optimally by either requesting a favor or holding it until period $t + 1$. Then

$$z_t = \max\{\lambda(\alpha - y_t) + (1 - \lambda)z_{t+1} - \gamma, z_{t+1}\} \geq z_{t+1},$$

where y_t is expected cost of providing a future favor or payment as a consequence of passing the token to one of A 's sensors. Hence if s_{Bk} ever finds it optimal to request a favor, then it should request a favor whenever it holds a token. Finally, as computed in Proposition 3, sponsors will optimally program their sensors to request favors in period T if $\alpha - \gamma/\lambda \geq 2x$. The number of tokens is maximized subject to these constraints by setting $x = \beta/2$. \square

4 Discussion

Our model features one-shot contractual interaction between the sponsors with exogenously limited liability. However, it is well known that results in such environments are qualitatively similar, and can be made quantitatively similar by appropriate choice of the liability limit, to settings of infinitely repeated interaction without contracts. As in the classic full information folk theorem (so called due to the lack of a definitive originator), in an infinitely repeated setting the non-contractual monetary transfers can be enforced by the threat of cutting off cooperation permanently. The effective liability limit can then be endogenously derived from the rate of time preference.

Furthermore, in the infinitely repeated setting, there are additional complicated, but well-understood modifications by which monetary transfers can be dispensed with entirely, and their incentive effects replicated by changing the number of tokens allocated to each sponsor in each successive project as a function of the outcomes in previous projects (see [5, 8]).

Also by standard arguments, our qualitative results extend further to infinitely repeated, non-contractual environments with many sponsors, in which pairs of sponsors match randomly each period, and in which each sponsor's history of cooperation is publicly observed. Such an environment may come closest to describing the potential use of multi-sponsor WSNs within a relatively small scientific community. But as the use of WSNs diffuses to wider audiences of users, such communal reputation mechanisms may begin to break down, making explicit contracts more useful.

5 Future Work

We are now simulating the operation of small scale multiple-sponsor networks using the ns-2 network simulator to estimate the fraction of efficient payoffs attained in

token equilibria as a function of the sponsors' liability. We are developing methods to evaluate the performance of token equilibria for a large-scale sensor networks under various realistic scenarios, including a range of probability distributions for requesting and providing favors.

Analytically, we have drastically simplified the analysis by assuming that if a sensor can provide favors during a period, it can provide as many favors as are requested of it. Under this assumption, the probability that a sensor can grant a favor does not vary during the period. If each sensor could grant at most one favor in each period, due to communication costs it might be optimal for sponsors to program each of their sensors not to request favors after a fraction of each period has passed—moreover, this fraction might differ for sensors located at different nodes on the grid. Adding this complication does not appear to make the analysis intractable, but it may be better addressed by simulation than by abstraction.

It has been shown by [6] in a related model that there can be efficiency gains to gradually forgiving debts of favors. In our context this suggests that it may be helpful to depreciate disparities in token holdings over time. Ultimately, under any given liability limit the optimal value of each sensor's marginal token ought to vary with the number of tokens held by it and its siblings. The fact that sensors do not know how many tokens their siblings hold may make computing an optimal equilibrium an intractable problem, but token equilibria are certainly amenable to tractable improvements, which we intend to pursue.

Existing middlewares for sensor networks typically focus on rapid development and integration, but within a network for a single application and administrative domains. To that end, we would like to propose an *agoric* approach to middleware design. We have started the development of its initial prototype. We then plan to deploy this middleware on the real sensors for performance evaluation. For scalability, we plan to leverage this middleware substrate towards development of *autonomous agents* to automate a broad range of activities including resource discovery and price negotiation.

References

- [1] A. Abdulkadiroğlu and K. Bagwell. Trust, reciprocity and favors in cooperative relationships. Working paper, May 2005.
- [2] J.-P. Benoit and V. Krishna. Nash equilibria of finitely repeated games. *International Journal of Game Theory*, 16:197–204, 1987.
- [3] J. Byers and G. Nasser. Utility-based decision-making in wireless sensor networks. In *Proceedings of the ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'00)*, Boston, MA, August 11, 2000.
- [4] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *Proceedings of the 6th Annual International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (Dial-M'02)*, Atlanta, GA, September 28, 2002.
- [5] D. Fudenberg, D. K. Levine, and E. S. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, September 1994.
- [6] C. Hauser and H. A. Hopenhayn. Trading favors: Optimal exchange and forgiveness. Working Paper, June 2005.
- [7] D. Irwin, J. Chase, L. Grit, and A. Yumerefendi. Self-recharging virtual currency. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'05)*, Philadelphia, PA, August 22–26, 2005.
- [8] D. A. Miller. Attainable payoffs in repeated games with interdependent private information. Working Paper, August 2005.
- [9] M. Möbius. Trading favors. Working paper, May 2001.
- [10] C. Ng, P. Buonadonna, B. N. Chun, A. C. Snoeren, and A. Vahdat. Addressing strategic behavior in a deployed microeconomic resource allocator. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'05)*, Philadelphia, PA, August 22–26, 2005.
- [11] C. H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*, Hersonissos, Crete, Greece, July 6–8, 2001.
- [12] N. Sadagopan and B. Krishnamachari. Decentralized utility-based design of sensor networks. In *Proceedings of the 2nd IEEE Workshop on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt'04)*, Cambridge, UK, March 24–26, 2004.
- [13] S. Zhong, L. E. Li, Y. G. Liu, and Y. R. Yang. On designing incentive-compatible routing and forwarding protocols in wireless ad-hoc networks—an integrated approach using game theoretical and cryptographic techniques. In *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom'05)*, Cologne, Germany, August 28–September 2, 2005.