Resource Allocation and Emergent Coordination in Wireless Sensor Networks

Aram Galstyan
Information Sciences Institute
University of Southern California
Marina del Rey, California
galstyan@isi.edu

Bhaskar Krishnamachari
Department of Electrical Engineering
University of Southern California
Los Angeles, California
bkrishna@usc.edu

Kristina Lerman
Information Sciences Institute
University of Southern California
Marina del Rey, California
lerman@isi.edu

Abstract

Coordination in wireless sensor networks (WSN) is required for many tasks that are best achieved collectively, such as coverage and medium access. One of the major challenges in the design of WSN are the strong limitations imposed by finite onboard power capacity. Because communication requires considerable energy, it is imperative to have a coordination mechanism that requires little or no communication. Moreover, since WSN are likely to operate in unstructured and dynamic environments, the coordination mechanism has to be adaptive and robust with respect to environmental changes. Lack of centralized control in WSN requires alternative means for coordinating actions and resources of individual nodes to achieve long network lifetime, while not severely compromising network task performance. In this paper we explore the paradigm of *emergent coordination* as a mechanism for adaptive, distributed coordination in WSN. Specifically, we study a WSN composed of self–interested nodes that utilize a simple reinforcement learning scheme and achieve coordination by playing repeated resource allocation (load balancing) games with changing resource (load) capacities. Our results indicate that for a certain range of parameters the network is very adaptive to these changes. Although we formulate the problem in rather abstract settings of repeated games, the methods can be applied to a range of specific sensor coordination problems such as network coverage and medium access.

Introduction

Thanks to recent technological advances in wireless communications and embedded systems there has been a growing interest in wireless sensor networks (WSN). WSN involve a large number of spatially distributed sensor nodes with heterogenous sensing and computational capabilities that communicate over wireless channels for collaborative information processing. WSN have variety of potential application such as environmental and habitat monitoring, military surveillance, maintenance and control of complex systems, etc.

Future WSN are envisioned to consist of hundreds to thousands of unattended sensor nodes, making centralized control of information processing in such networks infeasible. In order for wireless sensor nodes to operate autonomously, distributed coordination mechanism for an efficient system–wide performance (i.e., in target tracking, data routing, etc) are needed. One of the major challenges in WSN design that makes it crucially different from conventional networks are that individual nodes are severely energy-constrained. This fact puts strict limits on the communications and computational overhead on any control algorithms. Hence, coordination mechanisms that allow for little or no communication are of interest. A successful coordination mechanism should meet the following stringent requirements: it should be scalable, robust to individual agents/communication channels failures, and computationally efficient. Moreover, since WSN are likely to operate in unstructured and dynamic environments, the coordination mechanism has to be adaptive and robust with respect to environmental changes. Since communication is very energy-intensive, it is also desirable that a coordination mechanism require minimal or no communication.

In this paper, we explore the paradigm of emergent coordination as an efficient distributed control mechanism for WSN. Instead of concentrating on a specific sensor coordination problem, we present our results for rather general settings of repeated games. Specifically, we treat nodes as autonomous self-interested agents that utilize a simple reinforcement learning scheme and achieve coordination by playing repeated resource allocation (load balancing) games with changing resource (load) capacities. Our results indicate that for a range of parameters the system as a whole adapts efficiently to these changes. More importantly, the range of parameters for which coordination arises is independent of the number of nodes in the system. This property is very important for the WSN where the number of nodes might change in time (e.g., some nodes will run out of the power, while other nodes might be introduced to an already existing system).

Problem Formulation

Many problems in WSN can be stated in terms of distributed resource allocation. For example, in the network coverage problem one is interested in a distributed coordination mechanism that allows for only a certain fraction of nodes to be actively sensing at any given time (this could be phrased more specifically like all regions should be covered by k active sensors) — with the others in sleep mode to conserve energy. This can be termed as a resource allocation problem where the resource is the node's energy. Another example concerns the use of sensor nodes with CDMA spread-spectrum radios where, since multi-user interference reduces communication quality, it is desirable to ensure that no more than L nodes are transmitting over the medium at any given time. Here the problem is to design a coordination mechanism that will allow the nodes to use the common resource (channel bandwidth) efficiently.

In this paper, we are interested in a distributed coordination mechanism for resource allocation/load balancing scenarios similar to ones described above. Instead of explicitly coordinating their actions through algorithms such as leader election which might require significant communication overhead, we describe a scheme based on game dynamics. Specifically, we consider a WSN composed of N nodes, where each node has a finite set of available actions (e.g., what resource to use). There is a target number $L_i(t)$ (e.g., resource capacity), which can change in time, associated with each action. At each time step t, a node decides which action to take, and receives a positive payoff or reward whenever the number of nodes $n_i(t)$ that chose the i^{th} action is less than or equal to the target value for that action, $n_i(t) \leq L_i(t)$. The node will receive a negative payoff when $n_i(t) > L_i(t)$. Note that nodes that choose the same action receive the same payoff. There are two basic assumptions that we make: $i)$ nodes know exactly what their payoff is and $ii)$ nodes know the actions their neighbors took in the previous time step. While the second assumption is plausible in WSN where the nodes can communicate locally with their neighbors, the first one is more problematic. One solution is to have a base station that broadcasts to all nodes the output of the game at each time step (note that this scenario was used in (Iyer & Kleinrock 2003)). The other alternative is for the nodes to get some local feedback from their environment to estimate their payoff.

El Farol Bar Problem and Minority Games

Since its introduction in 1994, Arthur's El Farol Bar problem (Arthur 1994) has been one of the most widely studied examples of emergent coordination in complex adaptive systems. The model consists of N individuals/agents who have to decide independently whether to attend the El Farol bar in Santa Fe on a given night. The bar has a limited capacity, and people try to avoid attending it when it is overcrowded. There is no explicit communication between individuals, and the only information available to them is the time series of past attendance numbers. Since no deductively rational solution is possible, Arthur suggested to use inductive reasoning instead: Each agent has a set of "predictors" (strategies) that predict next week's attendance given the history of past attendance. Agents keep track of the performance of their predictors, and reinforce them according to their reliability. Numerical simulations of this simple model showed that the system self-organizes so that attendance fluctuates around the bar capacity.

The Minority Game (Challet & Zhang 1997) (MG) was introduced by Challet and Zhang as a simplified version of the El Farol Bar problem, the main difference being that instead of the actual history the agents are provided only with a binary string indicating whether the bar was overcrowded or not. More precisely, let us consider N nodes/agents in a WSN that repeatedly choose between two alternatives labelled 1 and 0 (*e.g.*, transmitting a packet through a common channel or not). If at a given time step the number of transmitting agents is less than or equal to the channel capacity, then the winning choice is 1; otherwise, it is 0.[1] As in the bar problem, each node uses a set of S strategies to decide its next move and reinforces strategies that would have predicted the winning group. The main advantage of the MG model is that strategies can be easily parameterized: a strategy is simply a lookup table that prescribes a binary output for all possible inputs, where the input is a binary string containing the last m outcomes of the game. Thus, for each choice of m, there are $P = 2^m$ possible histories (inputs), and $\Omega = 2^P$ strategies. Note that in this model the agents interact by sharing the same global signal. Despite its simplicity MG has been demonstrated to have a very rich and complex dynamics. The most interesting phenomenon of the minority model is the emergence of a coordinated phase, where the standard deviation of attendance, the volatility, becomes smaller than in the random choice game, where each agents makes either choice with probability $1/2$.[2] In particular, it has been established that coordination is achieved for memory sizes for which the dimension of the reduced strategy space is comparable to the number of agents in the system, $2^m \sim N$ (Challet & Zhang 1998; Savit, Manuca, & Riolo 1999).

Resource Allocation Games with Changing Resource Capacity

In a previous study (Galstyan & Lerman 2002) we demonstrated that if one introduces time-dependent capacities to the MG model defined in the previous section, the system does not adapt well. We also showed that one can achieve adaptation if instead of the interaction via a global signal, agents are allowed to interact locally. Note, that this is more realistic for the WSN where nodes communicate only locally.

Our model consist of a large number of simple autonomous nodes that interact locally: each node gets an input from K neighbors (describing what actions these nodes took in the previous time step) and maps the input to an action:

$$s_i(t+1) = F_i^j(s_{k_1}(t), s_{k_2}(t), \ldots, s_{k_K}(t)) \qquad (1)$$

where $s_{k_i}(t)$, $i = 1, \ldots, K$ are the choices made by its neighbors during the previous time step, and $F_i^j, j = 1, \ldots, S$ are randomly chosen boolean functions (called

[1] In the original MG model, the number of agents was taken to be odd, and the capacity was fixed to $(N-1)/2$, so that the agents who made the minority decision won, hence the name Minority Game.

[2] In the random choice game the average number of agents choosing "1" is $(N-1)/2$ with standard deviation $\sigma = \sqrt{N}/2$ in the limit of large N.

strategies hereafter) used by the i^{th} node. The strategies are chosen randomly and quenched throughout the game. If $S = 1$ such a network is the well known NK or Kauffman network (Kauffman 1993). As in the traditional MG, the node keeps a score for each strategy F_i^j that monitors the performance of that strategy, adding (subtracting) a point if the strategy predicted the winning (loosing) choice. We let $A(t)$ be the cumulative resource utilization at time t, $A(t) = \sum_{i=1}^{N} s_i(t)$. Then the winning choice is "1" if $A(t) \leq N\eta(t)$, and "0" otherwise. Those in the winning group are awarded a point while the others loose one. Nodes play the strategies that have predicted the winning side most often, with ties are broken randomly.

As a measure of efficiency we introduce $\delta(t) = A(t) - N\eta(t)$, that describes the deviation from the most optimal resource utilization. We are primarily interested in the cumulative "waste" over a certain time window:

$$\sigma^2 = \frac{1}{T_0} \sum_{t=t_0}^{t_0+T_0} \delta(t)^2 \tag{2}$$

We can compare the performance of our system to a default random choice game, defined as follows: assume that the nodes can query the capacity $\eta(t)$ at a given time step, and they choose action "1" with probability proportional to $\eta(t)$. In this case the average number of nodes choosing "1" is close to $\eta(t)N$ at each time step, and the fluctuations around the mean are given by the standard deviation

$$\sigma_0^2 = N\frac{1}{T} \int_{T_0}^{T_0+T} dt' \eta(t')[1 - \eta(t')] \tag{3}$$

The results of our simulation indicate that networks with $K = 2$ achieve very effective utilization of resources, even when the changes in the capacity level are relatively large. In Fig. 1 we plot resource utilization $A(t)$ for a sinusoidal change in capacity. One can see that the system follows the changes in capacity level very effectively. The inset of Fig. 1(a) shows the time series of the deviation $\delta(t)$ for $K = 2$. Initially there are strong fluctuations, hence poor utilization of the resource, but after some transient time the system as a whole adapts and the strength of fluctuations decreases. In fact, the standard deviation of the fluctuations is considerably smaller than in the random choice game as defined by Eq. (3).

In Fig. 1 (b) we plot the variance (strength of fluctuations) per node versus the parameter K, for system sizes $N = 100, 500, 1000$. For each K we performed 32 runs and averaged results. Our simulations suggest that the details of this dependence are not very sensitive to the particular form of the perturbation $\eta_1(t)$, and the general picture is the same for a wide range of functions, provided that they are smooth enough. The variance attains its minimum for $K = 2$ independent of the number of agents in the system.

We also studied how the system behaves when capacity changes suddenly. Remarkably, the system is able to adapt in this case too. This is illustrated in Fig.2, where we plot the resource utilization for random step-like changes in the capacity.

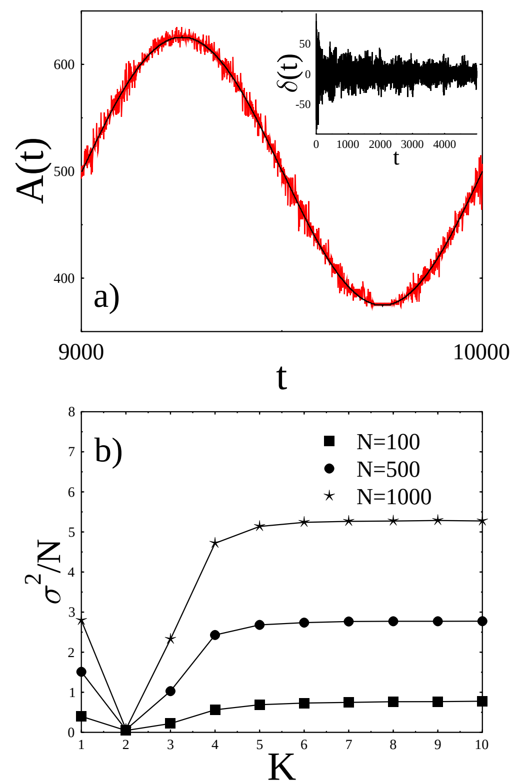

Figure 1: (a)A segment of the resource usage time series for $K = 2, \eta(t) = 0.5 + 0.12 sin(2\pi t/T), T = 1000$, (b) Waste per node vs the parameter K for different system sizes

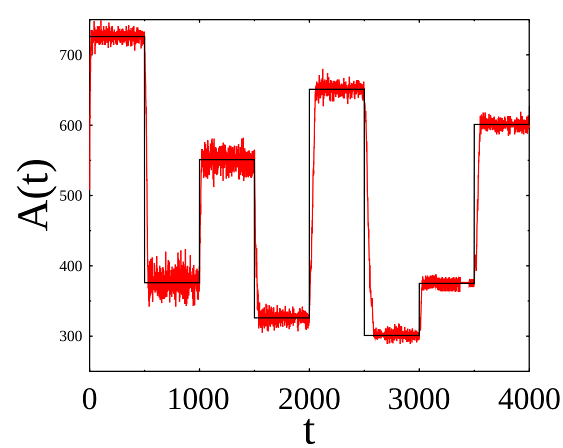

Figure 2: Resource utilization for step-like changes in the capacities

$S_{k_1}(t)$	$S_{k_2}(t)$	S_i(t+1)
0	0	0
0	1	2
0	2	1
1	0	1
1	1	0
1	2	0
2	0	2
2	1	0
2	2	2

Figure 3: Example of a strategy for $Q = 3$ and $K = 2$.

Multi-choice Games

In the previous section we considered a situation where the nodes have to make a binary decision (e.g., whether to transmit a packet or not, turn themselves on or not). In many practical situations, however, the number of choices might be greater. For example, in the presence of different sensing modalities, one might be interested to have a certain number of a specific sensors to be on at any given time. In this section we extend our model to account for the multi-choice scenario. Namely, nodes have to choose one of Q available actions, characterized by their target values $L_i(t), i = 1, \dots, Q$. To study this scenario in the context of previous sections, we use multi-state Kauffman networks (Sole, Luque, & Kauffman 2000) to model inter-node interactions. The actions of nodes are given by $s_i = \{0, 1, \dots, Q-1\}$, $i = 1, \dots, N$. The dynamics is specified analogously to the binary game, ***i.e.***, each node receives its inputs from K other nodes and maps it to one of the Q available actions.

$$s_i(t+1) = \Lambda_i^j(s_{k_1}(t), s_{k_2}(t), ..., s_{k_K}(t)) \qquad (4)$$

As before, each node has S functions (strategies), keeps a virtual score for its strategies and plays the strategy with the highest number of accumulated points. Let $A_i(t)$ be the number of nodes that chose i^{th} resource at time t. Then these nodes will be rewarded if $A_i(t) \le L_i(t)$ and punished otherwise.

As in the binary game, the strategies can be represented as lookup tables, that for each of the Q^k possible inputs assign an output (action) from the set $\{0, 1, \dots, Q-1\}$ (see Fig.). Strategy bias P is the parameter that determines the relative homogeneity of the output column of the strategy table. The entries in the output column are chosen as follows: first, with probability $1/Q$ one of the resources is chosen, let us say Q_k. Then, in the output column we choose $i)$ Q_k with probability P, $ii)$ Q_i, $i \neq k$ with probability $(1-P)/(Q-1)$. Thus, for a binary choice game ($Q = 2$), for $P = 0.5$ there are (in average) equal numbers of '0's and '1's in the output column; while for $P = 1.0$, the entries are all '0's or '1's, whichever symbol has been picked randomly.

The baseline solution we compare against is the random choice game where agents choose a particular resource with probability $\eta_i(t) = L_i(t)/N$. If the capacities are constant, then the standard deviation of each A_i that characterize the waste of resource is $\sigma_i^0 = N\eta_i(1-\eta_i)$ (the upper index 0 stays for the random choice game). In the case of changing capacities, one has to take an integral over the time window for which waste is calculated as in Eq. 3. In the results presented below we will primarily consider the waste averaged over the resources, i.e.,

$$\sigma_{tot}^2 = \frac{1}{Q}\sum_{i=1}^{Q}\sigma_i^2 \qquad (5)$$

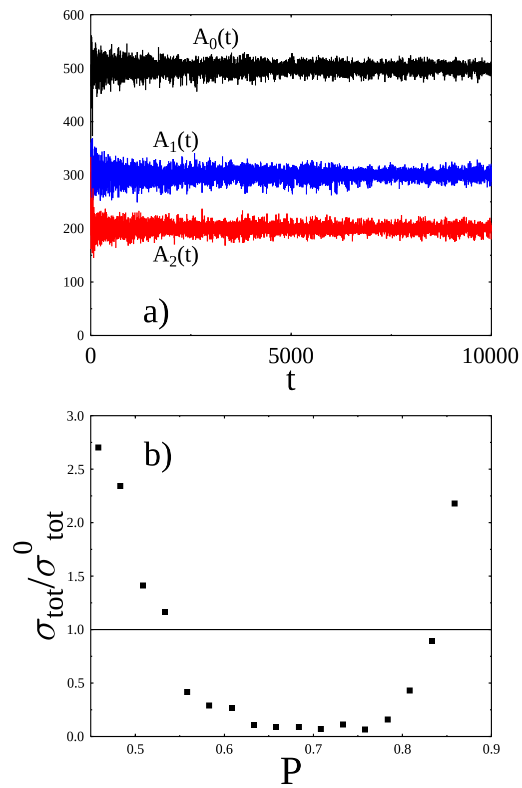

Figure 4: a) Time series for the resource usage for $Q = 3$ resources with capacities $\eta_0 = 0.5$, $\eta_1 = 0.3$, $\eta_2 = 0.2$. The simulation parameters are $N = 1000$, $S = 2$, $K = 2$, and $P = 0.6$. b)Normalized variance versus the strategy bias P (the horizontal line is the variance for the random choice game).

First, we examine the resource allocation problem with fixed capacities. The time series of the resource usage for $Q = 3$ choices, and capacities $\eta_0 = 0.5$, $\eta_1 = 0.3$, $\eta_2 = 0.2$ is shown in Fig. 4(a). After a short transient time, the number of nodes $A_i(t)$ using the resource i start to fluctuate around the fixed capacity level. The strength of these fluctuations determines the global waste. In Fig. 4(b) we show the dependence of waste (averaged over the resources and normalized to σ_{tot}^0) on the strategy bias. The horizontal line shows the value of σ_0 for the random choice game. Remarkably, for some values of P the efficiency of resource allocation is an order of magnitude better than in the random choice game.

Next we consider the case with time-dependent capacities. In Fig. we plot standard deviation *vs* the strategy bias

P for $Q = 10$ and number of strategies $S = 2, 3, 5$, for sinusoidal changes in the capacity. Again, an average over choices has been taken. For each set of parameters we did 8 trials, and averaged the results and normalized the cumulative waste by its value for the random choice game (horizontal line). One can see that for the certain range of the bias the efficiency of the system is better than in the random choice game, and there is a well pronounced minimum in σ_{tot}. Note also that the minimum is the deepest for $S = 2$, *i.e.,* two strategies per agent. One of the factors contributing to this is that it takes longer for the system to achieve the coordinated phase as one increases S, and since we have used the same duration of simulation for each S, this would result in a slightly greater value of σ_{tot}. However, we have verified that even if the simulation are carried out for long enough times, the system with $S = 2$ performs slightly better.

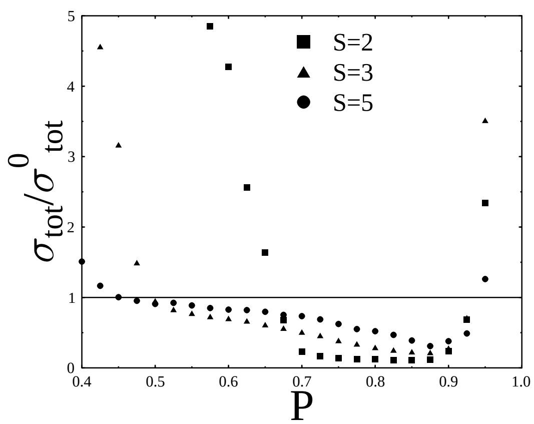

Figure 5: Normalized variance versus the strategy bias P for $Q = 10$, $N = 1000$, and time dependent (sinusoidal) capacities. The horizontal lines ($\sigma_{tot}/\sigma^0_{tot} = 1$) show the variance for the random choice game.

Related Work

Traditional alternatives to the techniques that we have described in this paper include (a) centralized tasking/planning and (b) distributed leader election. In the former approach, the central base station needs to maintain a global view of the system, perform the resource allocation and distribute the allocation plan to all nodes in the network. In the context of sensor networks, such centralized approaches have been rejected as being inherently non-scalable (Deborah Estrin & Kumar 1999). In the latter approach, nodes contend through a distributed election process using classical leader-election techniques such as the wave algorithm (Lynch 1996); these techniques have been used previously for cluster-head election, multi-target counting and feature extraction in sensor networks (Heinzelman *et al.* 2000; Fang, Zhao, & Guibas 2002; Krishnamachari & Iyengar 2003). Both approaches suffer from drawbacks compared to our proposed techniques – both can involve excessive communication between the nodes, an expensive operation in energy-starved sensor networks; and neither is particularly adaptive to changing loads/resource requirements.

Also somewhat related is the ASCENT protocol, which provides for distributed adaptive topology control for routing in sensor networks (Cerpa & Estrin 2002). ASCENT provides for resource-adaptation by waking up or putting to sleep sensor nodes depending on the level of activity in the network. It differs from our work in that ASCENT is not based on game dynamical techniques, and relies on explicit coordination between through the communication of “help” messages between nodes.

Perhaps, our approach is the closest to that of Iyer and Kleinrock (Iyer & Kleinrock 2003), who use the Gur game formalism to achieve coordination among the sensor nodes. In the Gur game the nodes are learning automata that interact with a base station. At each time step, the base station probabilistically distributes payoff to the nodes based on the number of nodes that chose to turn on. However, they are mainly concerned with the static situation, and it is not clear how Gur game will adapt to the changing target functions. In addition, our approach seems to be more general in that it allows us to treat coordination problems with more than two available actions.

Discussion

We presented an efficient mechanism for emergent coordination between autonomous nodes based on game dynamics. While we formulated the problem in rather abstract settings, the mechanism for emergent coordination through game–dynamics can be applied to various WSN problems such as coverage and medium access. We showed that this mechanism achieves robust and efficient coordination in resource allocation/load balancing tasks, such as network coverage and medium access, even when resource capacities are changing in time.

We studied the system numerically for a wide range of parameters and system sizes. Our results indicate that some parameters lead to a very adaptive and efficient global behavior, even though the nodes make decisions based mainly on local information. In some cases the resource utilization is by an order of magnitude better than in the baseline random choice game. Remarkably, the nodes are able to adapt without explicitly knowing the target values or the total number of nodes in the system. These features suggest that our model may serve as a highly distributed, robust and scalable mechanism for resource allocation in large scale WSN.

There are still many interesting research questions that were not addressed. For instance, we assumed that nodes rely on the binary signal to carry information about the game outcome. It will be interesting to come up with a scenario when instead of the global payoff signal (i.e., broadcasted by a base station), nodes can estimate their payoff outcome by local measurement of environmental state. For instance, in the channel allocation problem nodes can learn whether the packet they sent was received that would indicate whether the channel was over-utilized or not. Another way to determine the winning state of the game is to estimate it from the actions took by the neighbors. Preliminary simulations results suggest that this is a feasible way of achieving distributed coordination without relying on a global signal.

Acknowledgements

The research reported here was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contracts number F30602-00-2-0573. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of any of the above organizations or any person connected with them.

References

Arthur, W. B. 1994. Inductive reasoning and bounded rationality. *American Economic Review* 84:406–411.

Cerpa, A., and Estrin, D. 2002. Ascent: Adaptive self-configuring sensor networks topologie. In *INFOCOM*.

Challet, D., and Zhang, Y.-C. 1997. Emergence of cooperation and organization in an evolutionary game. *Physica A* 407.

Challet, D., and Zhang, Y.-C. 1998. On the minority game: Analytical and numerical studies. *Physica A* 256:514.

Deborah Estrin, Ramesh Govindan, J. H., and Kumar, S. 1999. Next century challenges: Scalable coordination in sensor networks. In *ACM Mobicom*.

Fang, Q.; Zhao, F.; and Guibas, L. 2002. Counting targets: Building and managing aggregates in wireless sensor networks. Technical Report P2002-10298, Palo Alto Research Center.

Galstyan, A., and Lerman, K. 2002. Adaptive boolean networks and minority games with time-dependent capacities. *Physical Review* E66:015103.

Heinzelman, W.; Chandrakasan, A.; ; and Balakrishnan, H. 2000. Energy-efficient communication protocol for wireless microsensor networks. In *33rd International Conference on System Sciences (HICSS'00)*.

Iyer, R., and Kleinrock, L. 2003. Qos control for sensor networks. In *Proceedings of ICC '03*.

Kauffman, S. A. 1993. *The Origins of Order*. New York: Oxford University Press.

Krishnamachari, B., and Iyengar, S. 2003. Efficient and fault-tolerant feature extraction in sensor networks.

Lynch, N. A. 1996. *Distributed Algorithms*. Morgan Kaufmann.

Savit, R.; Manuca, R.; and Riolo, R. 1999. Adaptive competition, market efficiency, phase transition. *Phys. Rev. Lett.* 82(10):2203.

Sole, R. V.; Luque, B.; and Kauffman, S. A. 2000. Phase Transitions in Random Networks with Multiple States. *SFI Working Papers* 00-02-011(5):4.