Coordination Failure and Congestion in Information Networks

A. M. Bell*, W. A. Sethares†, and J. A. Bucklew‡

Abstract

Coordination failure, or agents' uncertainty about the action of other agents, may be an important source of congestion in large decentralized systems. The *El Farol* or Santa Fe bar problem provides a simple paradigm for congestion and coordination problems that may arise with over utilization of the Internet. This paper recasts the problem in a stochastic framework and derives a simple adaptive strategy that has intriguing optimization properties; a large collection of decentralized decision makers, each acting in their own best interests and with limited knowledge, converge to a solution that (optimally) solves a complex congestion and social coordination problem. A variation in which agents are allowed access to full information is not nearly as successful.

Introduction

This paper focuses on imperfect information and coordination failure across agents as a source of congestion in large decentralized systems. We utilize the scenario posed by Arthur (1994) as a simplified model of a large class of congestion and coordination problems that arise in modern engineering and economic systems. *El Farol* is a bar in Santa Fe. The bar is popular, but becomes overcrowded when more than sixty people attend on any given evening. Everyone enjoys themselves when fewer than sixty people go, but no one has a good time when the bar is overcrowded.

The *El Farol* or Santa Fe bar problem emphasizes the difficulty of coordinating the actions of independent agents without a centralized mechanism. The analogy between the bar problem and decentralized resource allocation is noted by Greenwald, Mishra and Parikh (1998), as well as in our previous work (Sethares and Bell 1998). Glance and Huberman (1994) and Huberman and Lukose (1997) consider the dynamics of congestion on the Internet when externalities similar to those found with public goods prevail. Unlike the standard public good framework, in this scenario fully informed optimizing agents will not increase consumption of a publicly available resource until it experiences an inefficient level of congestion. if agents could predict the behavior of other agents perfectly the bar would never be crowded and all patrons would have a good time. The only source of congestion, at least in a deterministic framework, is the inability of agents to coordinate their actions.

In a previous treatment (Sethares and Bell 1998) we proposed a deterministic adaptive algorithm based on habit formation which enabled agents to coordinate in a decentralized environment while avoiding the seemingly random fluctuations in aggregate attendance that Arthur's simulations demonstrated.

Here we consider the bar problem in a stochastic setting where agents' strategies are characterized by a probability of attending that evolves over time. There are several advantages to considering the stochastic version of the adaptive learning rule: a clearer problem statement, a simpler algorithm

*NASA Ames Research Center, Mail Stop 269-3, Moffett Field, CA 94035-1000, abell@mail.arc nasa.gov.

†Department of Electrical and Computer Engineering, University of Wisconsin-Madison.

‡Department of Electrical and Computer Engineering, University of Wisconsin-Madison.

that is amenable to detailed analysis, and more general results We analyze the dynamic and equilibrium characteristics of the system in relation to the mixed and pure strategy equilibria of the corresponding game.

The type and characteristics of the equilibria actually observed depend crucially on the nature of the information available to agents. In particular, we show that limiting the information available to agents leads them to successfully coordinate on a Pareto efficient equilibrium while providing more information leads to an inefficient outcome. Our results emphasize the critical role that increased information exchange may play in creating and alleviating congestion that arises from coordination failure. For example, in a complex environment such as the Internet supplying individual routers with more information about congestion may not result in better system-wide performance.

Algorithm Statement

Let agents have identical payoffs: b is the payoff an agent receives for attending a crowded bar and g is the payoff an agent receives for attending an uncrowded bar. Without loss of generality let h, the payoff received for staying home, be zero. Let M be the total number of agents and $\mathcal{N}$ be the maximum capacity of an uncrowded bar. The game is then $G = [M, \{S_i\}, u_i(s_i, s_{-i})]$ where S_i consists of two strategies, go to the bar (indicated by 1) and stay home (indicated by 0) and $u_i(0, s_{-i}) = 0$ for all s_{-i}, $u_i(1, s_{-i}) = g$ when $\sum_{s_{-i}} \leq \mathcal{N} - 1$, and $u_i(1, s_{-i}) = b$ when $\sum_{s_{-i}} > \mathcal{N} - 1$.

In a deterministic setting where agents utilize only pure strategies a Nash equilibrium occurs when exactly sixty agents choose to attend. There are $\binom{100}{60}$ such equilibria. There are no symmetric pure strategy Nash equilibria. Pure strategy Nash equilibria are Pareto efficient: in equilibrium no agent can be made better off without making another agent worse off. There is one symmetric mixed strategy equilibrium where every agent has the same probability of attending each period. Because of the variance in attendance that results from agents' mixed strategies the equilibrium is not Pareto efficient.

Suppose that the agent initially attends p percent of the time. Consistent with the desire to maximize pleasure and minimize painful experiences, the agent goes more often (increases p slightly) if the bar is uncrowded, but prefers to go less often (to decrease p) if the bar is crowded Over time, the agent gathers information about the state of the bar, and 'remembers' this in the form of the parameter p. This learning rule can be interpreted as a kind of habit formation or stimulus-response.

Summarizing the previous notation there are M agents competing for the $\mathcal{N}$ spaces at the bar. The probability that the ith agent attends is p_i. Let k be a time (iteration) counter and $N(k)$ be the number of agents attending at time k. Let μ be a characteristic parameter that defines how much each agent changes p_i in response to new information and let $p_i(k)$ designate the instantaneous value of p_i at the time k. Let

$$N(k) = \sum_{i=1}^{M} x_i(k) \tag{1}$$

where the $x_i(k)$ are independent Bernoulli random variables that are 1 with probability $p_i(k)$ and zero otherwise. The evolution of the $p_i(k)$ is then defined by $p_i(k+1) =$

$$\begin{array}{ll} 0 \quad \text{if} \quad p_i(k) - \mu(N(k) - \mathcal{N})\, x_i(k) < 0 & \\ 1 \quad \text{if} \quad p_i(k) - \mu(N(k) - \mathcal{N})\, x_i(k) > 1 & \\ p_i(k) - \mu(N(k) - \mathcal{N})\, x_i(k) \quad \text{otherwise} & (2) \end{array}$$

The operation of the algorithm is uncomplicated. At each time k the agent flips a biased coin, attending with probability $p_i(k)$. When the agent attends, then the parameter $p_i(k)$ is adjusted, increasing it proportionally to $N(k) - \mathcal{N}$ if the bar is uncrowded and decreasing it proportionally to $N(k) - \mathcal{N}$ if the bar is crowded. Since the $p_i(k)$ represent probabilities, they must be constrained to lie within 0 and 1. When the agent does not attend $x_i(k)$ is zero and $p_i(k+1) = p_i(k)$. Note that the stepsize does not decrease over time. The simplicity of the scheme makes it feasible to analyze the resulting behavior, and as demonstrated in section .

In Arthur's formulation of the problem, agents have access to information about attendance at the bar even on evenings when they do not themselves attend. This can be incorporated into the algorithm (2), giving the update $p_i(k+1) =$

$$0 \quad \text{if}\ p_i(k) - \mu(N(k) - \mathcal{N}) < 0$$

$$1 \quad \text{if } p_i(k) - \mu(N(k) - \mathcal{N}) > 1$$

$$p_i(k) - \mu(N(k) - \mathcal{N}) \quad \text{otherwise} \qquad (3)$$

which mimics the information structure used by Arthur's agents. As will become clear, this information structure is a key element in the behavior of the algorithm. When agents base their updates on only their own experiences as in (2) they utilize "partial information". In contrast, (3) utilizes "full information" because agents base their decisions on the full record of attendance.

A related version of the stochastic algorithm updates according to whether the bar is crowded or not: $p_i(k+1) =$

$$0 \quad \text{if } p_i(k) - \mu \ \mathrm{sgn}(N(k) - \mathcal{N}) \ x_i(k) < 0$$

$$1 \quad \text{if } p_i(k) - \mu \ \mathrm{sgn}(N(k) - \mathcal{N}) \ x_i(k) > 1$$

$$p_i(k) - \mu \ \mathrm{sgn}(N(k) - \mathcal{N}) \ x_i(k) \quad \text{otherwise} \qquad (4)$$

Generic Behavior of the Algorithms

This section explores the generic behavior of the system when each of the $M = 100$ agents follows the strategy defined by (2) above. Though details of the various simulations differ, a typical case is illustrated in Figure 1. The probabilities $p_i(0)$ were initialized randomly.

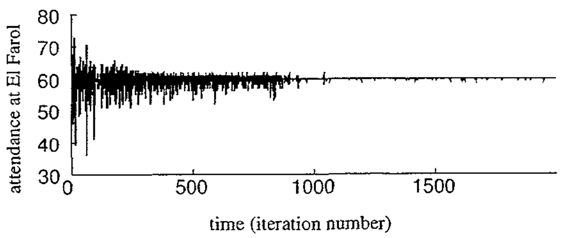

Figure 1: Attendance with Partial Information

Perhaps the most striking aspect of these simulations is the rapid convergence to near the optimal value of $\mathcal{N} = 60$ and the associated decline in the variance of attendance. The outcome approaches that which would be chosen with centralized control, despite the fact that each agent is autonomous, and makes the decision to go (or not to go) based on local information, that is, on its own experiences. In comparison, in Arthur's setup, there are far greater excursions about the optimal value and the bar is overcrowded about half the time.

Figure 2 shows values of the probabilities $p_i(k)$ over the course of a typical simulation run. By the final iteration, the agents have divided themselves into two groups. The probability parameter for 60 of the agents has risen very near 1, indicating that they go to *El Farol* nearly every time. The remaining 40 agents attend less and less frequently, with their probability parameter very near zero. This division of the population appears nowhere in the algorithm statement; rather, it is an emergent property of the adaptive solution to the *El Farol* problem. When agents follow this adaptive strategy, *El Farol* looks more like *Cheers*. Despite the stochastic nature of agents of the adaptive learning rule it converges to a pure strategy Nash equilibrium.

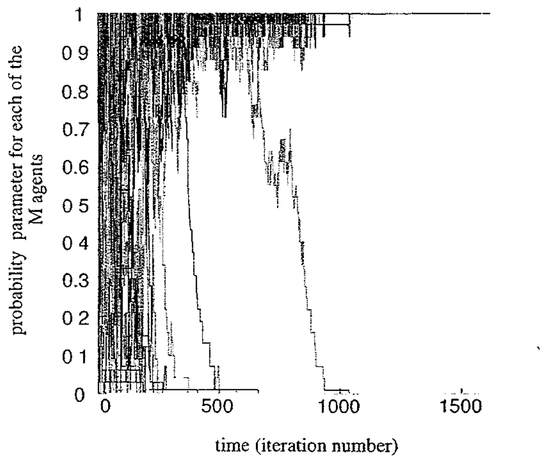

Figure 2: Probabilities with Partial Information

Finally Figures 3 and 4 use the "full information" algorithm (3) to investigate the effect of allowing the agents to update their probabilities at every iteration, whether they have personally attended the bar or not. This reflects the information structure in Arthur's simulations. Mean attendance is approximately 60, but the variance does not decline over time, indicating that seats in the bar often remain unfilled, and often the bar is overcrowded. Note that the transient behavior in the initial periods is masked by the long time scale. Figure 4 should be compared to Figure 2; the probability parameters for these agents continue to bounce randomly about some fixed value as their probabilities all increase or decrease simultaneously in response to the same signals.

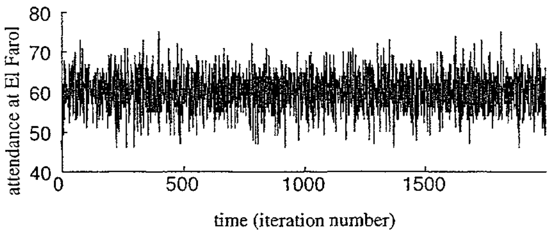

Figure 3: Attendance with Full Information

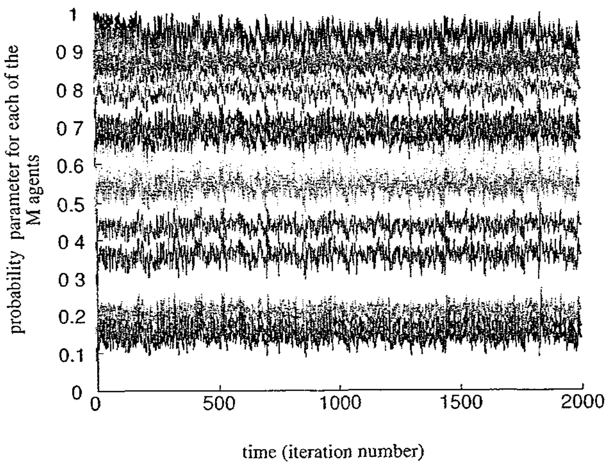

Figure 4: Probabilities with Full Information

Somewhat paradoxically, agents successfully coordinate their behavior and the system achieves a Pareto efficient outcome only when agents have access to less information. Several authors have noted a similar phenomena in transportation routing. Mahmassani and Jayakrishnan (1991) use simulations to demonstrate that when individuals pursue a strict best response strategy, changing their route no matter how small the improvement over their current choice, the performance of the system as a whole degrades if more than 25% of drivers have access to real time information about congestion. Arnott, De Palma and Lindsey (1996, 1991) show that congestion can arise because of "concentration," or similar responses to common information, and that consequently, more information can lead to increased congestion.

Analysis of the Adaptive Solutions

The first step in the analysis of the dynamic behavior of the algorithms is to determine the conditions under which the means of the $p(k)$ remain fixed; that is, to determine the steady states of the averaged system. We consider the partial information case first.

Taking the expectation of both sides of (2) gives

$$E\{p_i(k+1)\} = E\{p_i(k)\} - \mu E\{(N(k) - \mathcal{N})\ x_i(k)\},$$

assuming that the $p_i(k)$ are not at the boundary points 0 or 1. This expectation remains unchanged exactly when the update portion is zero, that is, when

$$E\{(N(k) - \mathcal{N})\ x_i(k)\} = 0.$$

Using (1) this can be rewritten

$$E\{(\sum_{j=1}^{M} x_j(k) - \mathcal{N})\ x_i(k)\} = E\{(\sum_{\substack{j=1 \\ j\neq i}}^{M} x_j(k) + 1 - \mathcal{N})\ p_i(k)\}$$

since the $x_i(k)$ is 1 with probability $p_i(k)$ and zero otherwise. Because the term in parenthesis is independent of $p_i(k)$ (recall that the $x_j(k)$ are independent Bernoulli random variables) this becomes

$$= (1 - \mathcal{N} + \sum_{\substack{j=1 \\ j\neq i}}^{M} E\{x_j(k)\})\ p_i(k)$$

$$= (1 - \mathcal{N} + \sum_{\substack{j=1 \\ j\neq i}}^{M} p_j(k))\ p_i(k).$$

In full vector form this is

$$\begin{pmatrix} (1 - \mathcal{N} + \sum_{j\neq 1}^{M} p_j(k))\ p_1(k) \\ (1 - \mathcal{N} + \sum_{j\neq 2}^{M} p_j(k))\ p_2(k) \\ \vdots \\ (1 - \mathcal{N} + \sum_{j\neq M}^{M} p_j(k))\ p_M(k) \end{pmatrix}. \qquad (5)$$

Consider any candidate steady state p^* with $\mathcal{N}$ ones and $\mathcal{N} - M$ zeroes. Let I_1 be the indices of the ones and I_0 be the indices of the zeroes. Then there are two kinds of terms in (5). When $i \in I_1$, $\sum_{\substack{j=1 \\ j\neq i}}^{M} p_j^* = \mathcal{N} - 1$ and so

$$(1 - \mathcal{N} + \sum_{\substack{j=1 \\ j\neq i}}^{M} p_j^*)\ p_i^* = (1 - \mathcal{N} + \mathcal{N} - 1)\ p_i^* = 0, \quad (6)$$

When $i \in I_0$, $\sum_{\substack{j=1\\j\neq i}}^{M} p_j^* = \mathcal{N}$, $p_i^* = 0$, and hence

$$(1 - \mathcal{N} + \sum_{\substack{j=1\\j\neq i}}^{M} p_j^*)\ p_i^* = (1 - \mathcal{N} + \mathcal{N})\ 0 = 0.$$

Hence p^* is a steady state.

Now consider any p^* for which $\sum_{j=1}^{M} p_j^* = \mathcal{N}$ that is not of the form of $\mathcal{N}$ ones and $\mathcal{N} - M$ zeroes. Thus $0 < p_n^* < 1$ for at least one n. In this case, the relevant term in (5) is

$$(1 - \mathcal{N} + \sum_{\substack{j=1\\j\neq n}}^{M} p_j^*)\ p_n^* = (1 - \mathcal{N} + \sum_{j=1}^{M} p_j^* - p_n^*)\ p_n^*$$

$$= (1 - N + \mathcal{N} - p_n^*)\ p_n^* = (1 - p_n^*)\ p_n^*.$$

This cannot be zero and hence p^* is not a steady state. Hence the only steady states of algorithm (2) are at p^* consisting of $\mathcal{N}$ ones and $M - \mathcal{N}$ zeroes. In particular, the symmetric mixed strategy Nash equilibrium at $p_j^* = .6$ for all j (for $g = 1$, $b \approx -0.98$) is not a steady state of this algorithm.

In contrast, consider a similar analysis carried out for the "full information" algorithm. Taking the expectation of both sides of (3) gives

$$E\{p_i(k+1)\} = E\{p_i(k)\} - \mu E\{\sum_{j=1}^{M} x_j(k) - \mathcal{N}\}.$$

Steady states occur when $E\{\sum_{j=1}^{M} x_j(k) - \mathcal{N}\} = 0$, i e., whenever

$$E\{\sum_{j=1}^{M} x_j(k)\} = \sum_{j=1}^{M} E\{x_j(k)\} = \sum_{j=1}^{M} p_j(k) = \mathcal{N}.$$

Hence any p^* with $\sum_{j=1}^{M} p_j^* = \mathcal{N}$ is a steady state of this algorithm. Note that these are not mixed strategy equilibria of the *El Farol* game unless $p_i = .6$ for every agent. The expected return in the steady state is lower for agents whose individual probabilities are lower than average as they face a higher probability that the bar will be crowded, and vice versa for those whose individual probabilities are higher than average. Consequently, a sensible learning rule might converge to a pure strategy equilibrium assuming agents adjust their parameters slowly overtime.

To further understand the global behavior of the system we relate the algorithms utilized by individuals to a global cost function. The algorithm can be derived as an approximation to an instantaneous gradient descent for minimization of the cost function

$$J(k) = (E\{N(k)\} - \mathcal{N})^2 \tag{7}$$

where

$$E\{N(k)\} = E\{\sum_{i=1}^{M} x_i(k)\} = \sum_{i=1}^{M} E\{x_i(k)\} = \sum_{i=1}^{M} p_i(k) \tag{8}$$

is the expected number of attendees at time k. The typical gradient strategy is to update the state using

$$p_i(k+1) = p_i(k) - \mu(k)\frac{dJ(k)}{dp_i(k)}. \tag{9}$$

With $J(k)$ as in (7),

$$\frac{dJ(k)}{dp_i(k)} = (E\{N(k)\} - \mathcal{N})\ \frac{dE\{N(k)\}}{dp_i(k)}.$$

From (8), the derivative is $\frac{dE\{N(k)\}}{dp_i(k)} = 1$, and hence

$$\frac{dJ(k)}{dp_i(k)} = E\{N(k)\} - \mathcal{N}.$$

Replacing $E\{N(k)\}$ by its instantaneous value gives

$$\frac{dJ(k)}{dp_i(k)} \approx N(k) - \mathcal{N}$$

which is an instantaneous approximation to the gradient of $J(k)$. Substituting this into (9) gives

$$p_i(k+1) = p_i(k) - \ \mu\ (N(k) - \mathcal{N}). \tag{10}$$

In the limited information case this update occurs only when $x_i(k) = 1$, in the full information case this update occurs every iteration regardless of the agent's attendance. Adding the *a priori* limits on $p_i(k)$ then gives the algorithms (2) and (3). For both algorithms $E\{N(k)\} = \mathcal{N}$ in a steady state. However, because the limited information algorithm converges to a pure strategy equilibria the actually observed costs will be 0, whereas with full information the expected costs will be $\frac{1}{2}Var[N(k)]$.

Similarly, the algorithm based on the sign of $(N(k) - \mathcal{N})$, 4, can be derived from the absolute value cost function $J(k) = |E\{N(k)\} - \mathcal{N}|$. By analogy, these algorithms are variants of the Least Mean Square (LMS) algorithms which are common in the context of linear system identification and adaptive filtering; (4) is an analog of the signed LMS algorithm.

The adaptive solution thus provides a simple mechanism whereby a large collection of decentralized decision makers, each acting in their own best interests and with only limited knowledge, can solve a complex congestion and social coordination problem. Moreover, convergence to the solution is relatively rapid (depending on the initial conditions) and robust.

References

Arnott, R.; De Palma A.; and Lindsey, R. 1996. Information and Usage of Free-access Congestible Facilities with Stochastic Capacity and Demand. *International Economic Review* 37(1):181-203.

Arnott, R.; De Palma A.; and Lindsey, R. 1991. Does Providing Information to Drivers Reduce Traffic Congestion? *Transportation Research A* 25A(5):309-318.

Arthur, W. B. 1994. Inductive Reasoning and Bounded Rationality: The *El Farol* Problem. *American Economic Review: Papers and Proceedings 1994* 84(May):406-411.

Glance, N S.; and Huberman, B. A. 1994. The Dynamics of Social Dilemmas. *Scientific American* (March):76-83.

Greenwald, A.; Mishra, B.; and Parikh, R. 1998. The Santa Fe Bar Problem Revisited: Theoretical and Practical Implications. Technical Report, New York University.

Huberman, B. A.; and Lukose, R. M. 1997. Social Dilemmas and Internet Congestion. *Science* 277(July 25):535-537.

Mahmassani, H. S.; and Jayakrishnan, R. 1991. System Performance and User Response Under Real-time Information in a Congested Traffic Corridor," *Transportation Research A* 25A(5):293-307.

Sethares, W. A.; and Bell, A. M. 1998. An Adaptive Solution to the *El Farol* Problem. *Proceedings of the 36th Annual Allerton Conference on Communication, Control, and Computing, Allerton IL, Sept. 1998.*