

An Optimal Medium Access Control with Partial Observations for Sensor Networks

Răzvan Cristescu

*Center for the Mathematics of Information, California Institute of Technology, Caltech 13693, Pasadena, CA 91125, USA
Email: razvanc@caltech.edu*

Sergio D. Servetto

*School of Electrical and Computer Engineering, College of Engineering, Cornell University, 224 Philips Hall, Ithaca, NY 14853, USA
Email: servetto@ece.cornell.edu*

Received 10 December 2004; Revised 13 April 2005

We consider medium access control (MAC) in multihop sensor networks, where only partial information about the shared medium is available to the transmitter. We model our setting as a queuing problem in which the service rate of a queue is a function of a partially observed Markov chain representing the available bandwidth, and in which the arrivals are controlled based on the partial observations so as to keep the system in a desirable mildly unstable regime. The optimal controller for this problem satisfies a separation property: we first compute a probability measure on the state space of the chain, namely the information state, then use this measure as the new state on which the control decisions are based. We give a formal description of the system considered and of its dynamics, we formalize and solve an optimal control problem, and we show numerical simulations to illustrate with concrete examples properties of the optimal control law. We show how the ergodic behavior of our queuing model is characterized by an invariant measure over all possible information states, and we construct that measure. Our results can be specifically applied for designing efficient and stable algorithms for medium access control in multiple-accessed systems, in particular for sensor networks.

Keywords and phrases: MAC, feedback control, controlled Markov chains, Markov decision processes, dynamic programming, stochastic stability.

1. INTRODUCTION

1.1. Multiple access in dynamic networks

Communication in large networks has to be done over an inherently challenging multiple-access channel. An important constraint is associated with the nodes that relay transmission from the source to the destination (relay nodes, or routers). Namely, the relay nodes have an associated maximum bandwidth, determined for instance by the limited size of their buffers and the finite rate of processing. Thus, the nodes using the relay need usually to contend for the access.

A typical example of such a system is a sensor network, where deployed nodes measure some property of the environment like temperature or seismic data. Data from these nodes is transmitted over the network, using other nodes as relays, to one or more base stations, for storage or control purposes. The additional constraints in such networks result

from the fact that the resources available at nodes, namely battery power and processing capabilities, are limited. Nodes have to decide on the rate with which to inject packets into a commonly shared relay, but the multiple-access strategy cannot be controlled in a centralized manner by the node that is acting as a relay, since communication with the children is very costly. Moreover, since nodes need to preserve their energy resources, they only switch on when there is relevant/new data to transmit, otherwise they turn idle. As a result, the number of active sources is variable, thus the amount of bandwidth the nodes get is variable as well. A poorly chosen algorithm for rate control may result in a large number of losses and retransmissions. In the case of sensor networks, this is equivalent with a waste of critical resources, like battery power. It is thus needed to design simple decentralized algorithms that adaptively regulate the access to the shared medium, by maintaining the system stable but still providing reasonable throughput. A realistic assumption is that nodes have only limited information available about the state of the system. Thus, the algorithms for rate control, implemented by the data sources, should rely only on limited feedback from the routing node.




FIGURE 1: Multiple access in a simple network.

We illustrate these issues with a simple network example shown in Figure 1. Nodes 1 and 2 need to control their rate of sending further their measured and/or relayed data, while relying only on feedback from the router. Node 3 serves one single packet at a time. If the relay is aware of the numbers of nodes that access it at a certain time moment (in this case, zero, one or two), it can just allocate some fair proportion of its bandwidth to each of them, avoiding thus collisions. However, such an information is not available in general neither at the relay, nor at the nodes accessing it.

Suppose each of the two nodes 1 and 2 employs a simple random medium access protocol, defined by two Bernoulli random variables u_1 , u_2 that determine the injection probabilities. Due to the above mentioned power and communication limitations, the nodes are not able to communicate between them. For the same reasons of minimizing the overhead, they need to control the rate of transmission by using only limited information (feedback) from the relay node. This feedback is usually restricted only to acknowledgments of whether the packet sent was accepted or not. Most current protocols for data transmission, including Aloha and TCP, use this kind of information for the rate control. Current proposals for medium access protocol in sensor networks make use of randomized controllers. The study of performance and stability of such protocols is thus of obvious importance.

As an example, suppose node 1 uses a probability of injection $u_1 = 0.5$, that is, it will try to inject on average a packet every two time slots. If it sends a packet and this is accepted (there is free place in the buffer of node 3), an adequate policy will consequently increase its rate u_1 , since it is probable that node 2 is not active at that particular time. As a result, node 1 accesses the buffer more often. If on the contrary the packet is rejected, then it is probable that node 2 is accessing the channel in the same time, too. Then, node 1 will decrease its rate. Note that care must be taken so that neither of the nodes alone take full use of the buffer. This fairness can be achieved, for instance, by drastically reducing the injection probability when losses are experienced. The design and analysis of such control policies is the goal of this work.

For such a setting, due to frequent failures on links and frequent need of rerouting, protocols like TCP are not suitable (e.g., the IEEE 802.11 protocol is based on a *random* access algorithm). On the other hand, stability of random access systems (like e.g., Aloha [1]), but with private feedback,

is hard to analyze. Our goal is to provide an analysis of systems under variable conditions, where there are only partial observations available, and the rate control actions are based on those partial observations.

In this paper, we set up a “toy” problem which is analytically tractable, and which captures in a clean manner some of these issues. We propose a hybrid model, in which nodes get only private feedback from the router, like in TCP. However, TCP behavior (including fairness) is not explicitly imposed, but as we will see further, the resulting system has the slow increase/fast decrease type of behavior specific to TCP. Note that an Aloha type of contention resolution, where if there is collision no packet goes through, does not take full advantage of the buffering available at relaying nodes. Thus, unlike in Aloha, in our model one packet goes always out of the queue (since the relay has a finite buffer, and filling of the buffer is prevented by the rate control at nodes).

The key property of our model is that the control decisions, on what rate to be used by a node, are based on *all the history* that is locally available at that node. For a network with partial observations, intuitively this is the best that can be done.

1.2. Related work

The problem of how different sources gain access to a shared queue is an abstraction of the thoroughly studied *flow control* problem in networks. Many practical and well-debugged algorithms have been developed over the years [2, 3], and more recently, formulations of this problem have taken more analytical approaches, based on game theoretic, optimization, and flows-as-fluids concepts [4, 5, 6, 7]. More recently, the flow control problem has been addressed in sensor networks [8, 9].

Several important issues appear in studying the MAC problem in the sensor network context, including limited power and communication constraints, as well as interference. Contention-based algorithms include the classical examples of Aloha and carrier-sense multiple access (CSMA) [1]. Recently proposed algorithms adapted to the specific requirements of sensor networks are presented in [10, 11, 12, 13]. Scheduling-based algorithms include TDMA, FDMA, and CDMA (time/frequency/code-division multiple access) [14, 15, 16, 17, 18].

The need of a unified theory of control and information in the case of dynamic systems is underlined in the overview of [19], where the author discusses topics related to the control of systems with limited information. These issues are discussed in the context of several examples (stabilizing a single-input LTI unstable system, quantization in a distributed control two-stage setting, and LQG), where improvements in the considered cost functions can be obtained by considering information and control together, namely by “measuring information upon its effect on performance.” Extensive work along these lines is presented in [20], where the author derived techniques which consider the use of partial information, for capacity optimization of Markov sources and channels, formulated as dynamic programming problems.




FIGURE 2: The problem of N sources sharing a single finite buffer. When each source gets to observe the state of the entire network, this problem degenerates to the single-source case. The interesting case however occurs when sources only have partial information about the state of the system, and they must base decisions about when to access the channel only on that partial data.

The main tool we use in this work is the control theory with partial information. An important quantity in this context is the information state, which is a probability vector that weighs the most that can be inferred about the state of the system at a certain time instance, given the system behavior at previous time instances. There are some important results in the literature dealing with related results on convergence in distribution of the information state, in which the state of a system can only be inferred from partial observations. Kaijser proved convergence in distribution of the information state for finite-state ergodic Markov chains, for the case when the chain transition matrix and the function which links the partial observation with the original Markov chain (the observation function) satisfy some mild conditions [21]. Kaijser’s results were used by Goldsmith and Varaiya, in the context of finite-state Markov channels [22]. This convergence result is obtained as a step in computing the Shannon capacity of finite-state Markov channels, and it holds under the crucial assumption of i.i.d. inputs: a key step of that proof is shown to break down for an example of Markov inputs. This assumption is removed in a recent work of Sharma and Singh [23], where it is shown that for convergence in distribution, the inputs need not be i.i.d., but in turn the pair (channel input, channel state) should be drawn from an irreducible, aperiodic, and ergodic Markov chain. Their convergence result is proved using the more general theory of regenerative processes. However, using directly these results in our setting does not yield the sought result of weak convergence and thus stability, as we will show that the optimal control policy is a function of the information state, whereas in previous work, inputs are independent of the state of the system. This dependence due to feedback control is the main difference between our setup and previous work.

1.3. Main contributions and organization of the paper

We formulate, analyze, and simulate a MAC system where only partial information about the channel state is available. The optimal controller for this problem satisfies a separation property: we first compute a probability measure on the state space of the chain, namely the information state, then use this measure as the new state based on which to make control decisions. Then, we show numerical simulations to illustrate

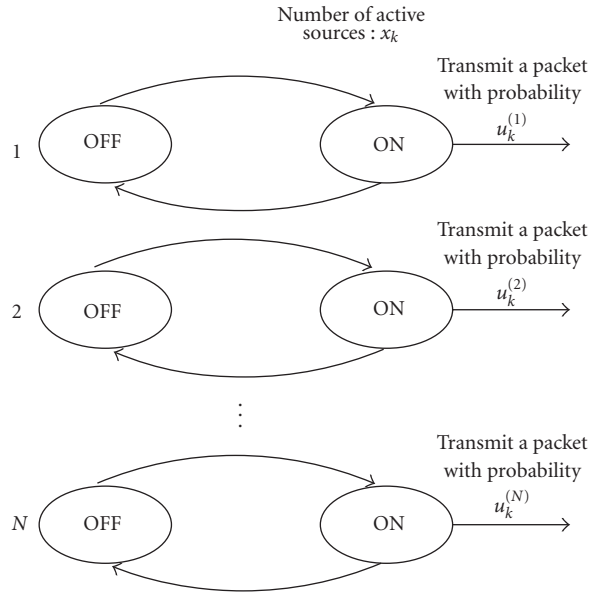


FIGURE 3: To illustrate the proposed model. N sources switch between ON/OFF states. When a source is in the ON state, it generates symbols with a (controllable) probability $u_k^{(i)}$. When it is in the OFF state, it is silent.

with concrete examples properties of the optimal control law. Finally, we show how the ergodic behavior of our queuing model is characterized by an invariant measure over all possible information states, and we construct that measure.

This paper is organized as follows. In Section 2, we set up a model of a queuing system in which multiple sources compete for access to a shared buffer, we describe its dynamics, we formulate and solve an appropriate stochastic control problem. We also present results obtained in numerical simulations to illustrate with concrete examples properties of these control boxes. Then, in Section 3, we study ergodic properties of the queuing model that result from operating the system of Section 2 under closed-loop control. There, we show how long-term averages are described succinctly in terms of a suitable invariant measure, whose existence is first proved, and then effectively constructed. The paper concludes with Section 4.

2. THE CONTROL PROBLEM

2.1. System model and dynamics

Consider the following discrete-time model (see Figure 2).

- (i) N sources feed data into the network, switching between ON/OFF states in time. While ON, source $S^{(i)}$ generates a symbol at time k with probability $u_k^{(i)}$, and remains silent with probability $1 - u_k^{(i)}$; while OFF, the source remains silent with probability 1. Given the intensity value $u_k^{(i)}$, this coin toss is independent of everything else (see Figure 3).




FIGURE 4: The only information a source has about the network is a sequence of 3-valued observations: acknowledgments, if the symbol was accepted by the buffer; losses if it is rejected due to overflow, and nothing if the decision was not to transmit at the current moment (denoted by 1, -1, 0, resp.).

- (ii) The queue has a finite buffer. When a source generates a symbol to put in this buffer, if the buffer is full, then the symbol is dropped and the source is notified of this event; if there is room left in the buffer, the symbol is accepted, and the source is notified of this event as well. Note that feedback is sent *only* to the source that generates a symbol, and not to all of them.
- (iii) The control task consists of choosing values for all $u^{(i)}$'s, at all times. A basic assumption we make is that sources are *not* allowed to coordinate their efforts in order to choose an appropriate set of control actions $u^{(i)}$ ($i = 1, \dots, N$): instead, the only cooperation we allow is in the form of having all sources implementing the same control technique, based on feedback they receive from the queue.
- (iv) The service rate of the queue is deterministic.

An illustration of this proposed model is shown in Figure 4. The dynamics of this system are modeled as follows.

- (i) $x_k \in \mathcal{S} = \{1, \dots, N\}$ is the number of ON-sources at time k , modeled as a finite-state Markov chain¹ with known matrix P of transition probabilities P_{ij} given by $p(x_k = j \mid x_{k-1} = i)$ (independent of the source intensities $u_k^{(i)}$ and of the time index k), and known $p(x_0)$ (the initial distribution over states).
- (ii) $r_k^{(i)} \in \mathcal{O} = \{-1, 0, 1\}$ is the ternary feedback from the queue to the source. The convention we use is that -1 denotes losses, 0 denotes idle periods, and 1 denotes positive acknowledgments.
- (iii) $u_k^{(i)} \in \mathcal{U}$, where $\mathcal{U} = (0, 1]$ source intensities, controllable (as defined above).
- (iv) $q_{k+1} = \min(\max(q_k + a_k - c, 0), B)$ is the queue size at moment k , with a_k the number of accepted packets, c the number of departing packets (c has a constant value), and B the maximum buffer size. If a new packet

¹For example, it is straightforward to prove that if the ON/OFF process of each source is modeled as a two-state Markov process, then also the total number of active sources is a finite-state Markov chain.




FIGURE 5: Consider a fixed (observed) state i , and assume a large *finite* shared buffer (for simplicity—if not, these curves would have to be replaced by curves derived from large deviations estimates such as given by the Chernoff bound). The probability of a packet loss is zero until the injection rate hits the fairness point $1/i$, beyond which it increases linearly, and the probability of a packet finding available space in the shared buffer increases linearly up until the fairness point $1/i$, beyond which it remains constant. Note that $u^* > 1/i$ is the largest $u \in (0, 1]$ such that $p(-1 \mid i, u) \leq T$ —the gap between $1/i$ and u^* is the “margin of freedom;” we will have to risk the loss of packets, in the case when i cannot be observed.

is accepted, the queue generates an $r_k = 1$ private acknowledgment to the source from which the packet is originated, and if the packet is not accepted, the queue generates an $r_k = -1$ acknowledgment.

- (v) $p(r \mid x, u)$ is the probability of occurrence of an observation $r \in \mathcal{O}$, when x sources are active, and when symbols are generated by all active sources at an average rate u . These probabilities can be computed a priori: for a finite but large enough buffer, a good approximation for $p(r \mid x, u)$ is illustrated in Figure 5. Note that in this approximation, the values of $p(r \mid x, u)$ do not depend on the maximum size of the buffer B , nor on the instantaneous queue size q_k .

These dynamics are illustrated in Figure 6.

There are two important observations to make about how we have chosen to set up our model. Describing the probabilities of observations $p(r \mid x, u)$ only in terms of the number of active sources x and the *average* injection rate u of all the active sources does require some justification: how can we assume that all sources inject the same amount of data, when the data on which these decisions are based (feedback from the queue), is not shared, and each source gets its own private feedback? Although this might seem unjustified, that is not the case. Once we study in some detail the control problem we are setting up here, we will find that the optimal control action u_k at time k is given by a memoryless function $u_k = g(\pi_k)$ of a random vector π that has the same distribution for *all* sources, and with well-defined ergodic properties—a precise study of these ergodic properties is the subject of Section 3. Therefore, even though at any point in time there will likely be some sources getting more and some other sources getting less than their fair share, on average all




FIGURE 6: An illustration of the model from the point of view of a single source, based on a simple birth-and-death chain for the evolution of the number of active sources.

get the same. This issue is further discussed below, both analytically (in Section 3) and in terms of numerical results (in Figure 10).

Another important thing to note is that there are strong similarities between our model and the formalization of multiaccess communication that led to the development of the Aloha protocol. However, the fact that feedback is not broadcast to all active sources in our model is a major difference between our formulation and that one. In fact, we conceived our model as an analytically tractable “hybrid” between Aloha and TCP. Like in slotted Aloha, time is discrete, feedback is instantaneous, and the state follows a Markovian evolution; but like in TCP, feedback is private only to the source that generated a transmitted packet.

Hajek [24] reviews a series of results for the two usual models for Aloha (finite number of users and one packet at a time, and infinite number of users). Decentralized policies for the injection probabilities, that maintain stability in the case of private acknowledgment feedback, are hard to be derived for the infinite-nodes case with Poisson arrivals. There is however important work [24] about stability in the finite-nodes study of Aloha. The theory in [24] is applied, as an example, to finding conditions of stability for multiplicative policies for sources that are supplied with Poisson arrivals. We expect that the theory we develop in this paper will provide a useful background for an Aloha model with random arrivals (not necessarily Poisson), with a finite number of backlogged packets, and its extension to the infinite-user model.

2.2. Formal problem statement

Intuitively, what we would like to do is maximizing the rate at which information flows across this queue, subject to the constraint of not losing too many packets. Since each time we attempt to put a packet into the shared buffer there is a chance that this packet may be lost, it seems intuitively clear that without accepting the possibility of losing a few packets, the throughput that can be achieved will be low; at the same time, we do not want a high packet loss rate, as this would correspond to a highly unstable mode of operation for our system.

This intuition is formalized as follows. Our goal is to find a policy $g = (u_1, \dots, u_K)$ that solves

$$\begin{aligned} \max_g \limsup_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K p(r_k = 1 | x_k, u_k), \\ \text{subject to } p(r_k = -1 | x_k, u_k) \leq T, \quad \forall k, \end{aligned} \tag{1}$$

where $T \in (0, 1]$ is a parameter that specifies the maximum acceptable rate of packet losses.² Note that we use a lim sup in the definition of our utility function (instead of a regular limit) because we do not know yet that the limit actually exists—although it certainly does, as will be shown later.

2.3. Warming up: finite horizon and observed state

We start with the solution to an “easier” version of our control problem: one in which the state of the chain (i.e., the number of active sources at any time) is known to all the sources. Although this would certainly *not* be a reasonable assumption to make (it does trivialize the problem), we find that looking at the solution to the general problem in this specific case is actually quite instructive, and so we start here as a step towards the solution of the case of true interest (hidden state).

The problem formulated above is a textbook example of a problem of optimal control for controlled Markov chains, and its solution is given by an appropriate set of dynamic programming equations [25]. Define $\mathbf{c}(u) = [p(1 | i, u) \cdots p(1 | N, u)]^\top$, and then

$$V_K(i) = 0, \tag{2}$$

$$\begin{aligned} V_k(i) &= \sup_{u: p(-1|i, u) \leq T} \{ \mathbf{c}(u) + PV_{k+1} \} \\ &= \sup_{u: p(-1|i, u) \leq T} \{ \mathbf{c}(u) + C \} \quad (C \text{ independent of } u). \end{aligned} \tag{3}$$

Equation (2) is set to 0 because this is only a finite-horizon approximation, but we are interested in the infinite-horizon

²In Figure 9, on numerical simulations, we illustrate how this parameter affects the behavior of the controller.




FIGURE 7: Illustrates the separation of estimation and control. Suppose we have a controlled system, which produces certain observable quantities related to its unobserved state. Based on these observations, we compute an *information state*, a quantity that somehow must capture all we can infer about the state of the system given all the information we have seen so far (this concept will be made rigorous later). This information state is fed into a control law that uses it to make a decision of what control action to choose, and this action is fed back into the system.

case, and in this case, the boundary condition given by $V_K = 0$ has a vanishing effect as we let $K \rightarrow \infty$. What is more interesting is that from (3), it follows that a *greedy* controller is optimal: this is not at all unexpected, since in our model the transition probabilities P are not affected by control, only observations are. The interplay among control and the different probabilities of observations are illustrated in Figure 5.

2.4. One step closer to reality: partial information

Definition 1. Denote the simplex of N -dimensional probability vectors by $\Pi = \{(p_1, \dots, p_N) \in \mathcal{R}^N : p_i \geq 0, \sum_{i=1}^N p_i = 1\}$.

The case of partial information (i.e., when the underlying Markov chain cannot be observed directly) poses new challenges. The problem in this case is that Markovian control policies based on state estimates are not necessarily optimal. Instead, optimal policies satisfy a “separation” property, illustrated in Figure 7 and extensively discussed in [25, pages 84–87].

Formally, an information state π_k is a function of the entire history of observations and controls $r_0 \cdots r_{k-1} u_0 \cdots u_{k-1}$, with the extra requirement that π_{k+1} can be computed from π_k, r_k, u_k .³ A typical choice is to let π_k be $p(x_k | r^{k-1}, u^{k-1})$, the conditional probability of x_k given all the past observations and applied controls. Then, an optimal controller for partially observed Markov chains also satisfies a set of dynamic programming equations, but instead of being over the states of the chain (a finite number), these equations are defined over information states [25] (i.e., over all points in the simplex of Π probabilities over N points):

$$V_K(\pi) = 0, \quad (4)$$

$$V_k(\pi) = \sup_{u: E_\pi p(-1|i, u) \leq T} E_\pi \{c(i, u) + V_{k+1}(F[\pi, u, r])\},$$

where F denotes the recursive updates of π , and where the notation E_π denotes expectation relative to the measure π .

³Note that this is a very reasonable requirement to make of something that we would like to think of as capturing some notion of state for our system.

A straightforward derivation gives the information-state transition function F :

$$\begin{aligned} \pi_{k+1} &= F[\pi_k, u_k, r_k] \\ &= C_{\pi_k} \cdot \pi_k \cdot D(u_k, r_k) \cdot P, \end{aligned} \quad (5)$$

with C_{π_k} a normalizing constant, P the transition-probability matrix of the underlying chain, and $D(u, r) = \text{diag}[p(r | 1, u) \cdots p(r | N, u)]$ a diagonal matrix. This is essentially the same set of DP equations as before, but where dependence on states is removed by averaging with respect to the current information state π_k . And as before, the optimal controller is chosen by recording for each π the value of u that achieves the supremum in the left-hand side of (4). The optimal control will thus be a function of only the information state, $u = g(\pi)$.

2.5. Infinite horizon

In the previous sections, we derived the solution for the optimal control in the case of partial observations when the time horizon is finite. We can get back now to the infinite-horizon problem stated in (1). The dynamic programming algorithm becomes a fixed-point system of equations with the unknowns spanning the simplex Π . Indeed, we start from the finite horizon case:

$$V_K(\pi) = \sup_{u: E_\pi p(-1|i, u) \leq T} E_\pi [c(i, u) + V_{K-1}(F[\pi, u, r_k])]. \quad (6)$$

We rewrite (6) as [25]

$$\begin{aligned} \frac{V_K(\pi)}{K} &= \sup_{u: E_\pi p(-1|i, u) \leq T} E_\pi \left[c(i, u) + V_{K-1}(F[\pi, u, r_k]) \right. \\ &\quad \left. - V_K(\pi) + \frac{V_K(\pi)}{K} \right]. \end{aligned} \quad (7)$$

Assume that the following limits exist for all $\pi \in \Pi$ and some J^* :

$$\lim_{K \rightarrow \infty} (V_K(\pi) - KJ^*) = V_\infty(\pi). \quad (8)$$

Then by taking the limit $K \rightarrow \infty$ in (7), we finally get

$$J^* + V_\infty(\pi) = \sup_{u: E_\pi p(-1|i, u) \leq T} E_\pi [c(i, u) + V_\infty(F[\pi, u, r_k])]. \quad (9)$$

The DP equation in (9) holds actually under more general conditions easy to verify for our model [25]. The transition-probability matrix P does not depend in our model on the control policy. Further, the Markov chain given by the number of active sources is irreducible in normal circumstances. Then it is shown in [25] that if these conditions are fulfilled, then the DP equation system for the average cost criterion is as in (9) and there exist $V(\pi)$, $\pi \in \Pi$ and J^* that solve it. Also, J^* is the minimum average cost and a policy g is optimal if $g(\pi)$ attains the minimum in (9).

One might attempt to solve the fixed-point system in (9) with an iteration algorithm on a discretized version of the equations system. However, there are practical difficulties to implement and simulate the optimal controller in the partial information case as defined above, having to do with the fact that our state space is the whole simplex of probability distributions Π . Our approach to find an approximate solution for the optimization problem (1) is to solve the dynamic programming system for the finite-horizon case (finite K), and study the properties of the obtained control policy by numerical simulations.

2.6. Numerical simulations

To help develop some intuition for what kind of properties result from the optimal control laws developed in previous sections, in this section we present results obtained in numerical simulations. Our approximation consists in choosing the maximum control at time k that still obeys the loss constraint, since this will also maximize the throughput. In Figure 8, we present a typical evolution over time of the information state, in Figure 9, we illustrate how different values of the threshold T influence the behavior of the controller, and in Figure 10, we address the fairness issue raised at the end of Section 2.1.

In all our simulations, we compare our controller with partial observation, with the optimal genie-aided controller that would be used if the number of active sources were known. Note that the difference between the optimal genie-aided controller and the controller derived by our algorithm is dependent on the two defining parameters of the system: the loss threshold T , and the transition-probability matrix P . Namely, our controller adapts faster to the network conditions if the transition matrix P corresponds to a slowly changing Markov chain; on the other hand, the larger the threshold T implies better adaptation, at the expense of an increased level of losses.

3. PERFORMANCE ANALYSIS

3.1. Overview

3.1.1. Problem formulation

In Section 2, we gave a model for the system of interest, we described its dynamics, we formulated an optimal control problem, we showed how this problem can be solved using standard techniques developed in the context of controlled Markov chains [25], and we developed numerical simulations to illustrate with concrete examples properties of the queues operating under feedback control. Now, once we have that optimal control algorithm, each source gets to operate the queue based on its local controller, thus resulting in a “decoupling” of the problem, as illustrated in Figure 11.

Perhaps the first question that comes to mind once we formulate the picture shown in Figure 11 is about ergodic properties of the resulting controlled queues. Specifically, we will be interested in two quantities.

(i) Average throughput:

$$J(g) = \lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K p(1 | x_k, g(\pi_k)) \quad (10)$$

$$\stackrel{?}{=} \int_{\{x, \pi\}} p(1 | x, g(\pi)) d\nu(x, \pi).$$

(ii) Average loss rate:

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K p(-1 | x_k, g(\pi_k)) \quad (11)$$

$$\stackrel{?}{=} \int_{\{x, \pi\}} p(-1 | x, g(\pi)) d\nu(x, \pi).$$

Therefore we see that, in both cases, the questions of interest are formulated in terms of a suitable invariant measure. Since we have assumed the underlying finite-state Markov chain to be irreducible and aperiodic, this chain does admit a stationary distribution. Therefore, a sufficient condition for the existence of the sought measure ν is the weak convergence of the sequence of information states π_k to some limit distribution over the simplex Π_N of probability distributions on N points. And to start developing some intuition on what to expect in terms of the sought convergence result, it is quite instructive to look at typical trajectories of the information state, as shown in Figure 12.

We state now the main theorem of this paper.

Theorem 1. *The sequence π_k converges weakly to a limit distribution ν over the simplex Π .*

The proof will follow after we briefly review some previous related work.

3.1.2. Some related work

Note that the stability of the control policy cannot in general be proven using a Lyapunov function, since the dependence of the optimal control on the information state is not a closed-form function.

In view of the previous results [21, 22, 23], a seemingly feasible approach to establish the sought convergence for our system would have been considering the control action $u \in \mathcal{U}$ to play the role of a channel input in the setup of [22], while the observations $r \in \mathcal{O}$ could have played the role of a channel output (thus making the control u and the observation r the available partial observations). However, this approach does not yield the sought result. In our system, the control u is a function of the information state, that is, it depends on the state of the system, but in those previous papers, inputs are independent of the state of the system.

3.1.3. Weak convergence of the information state—steps of the proof

The proof of weak convergence of π involves five steps.

(1) First, we show that the sequence of information states

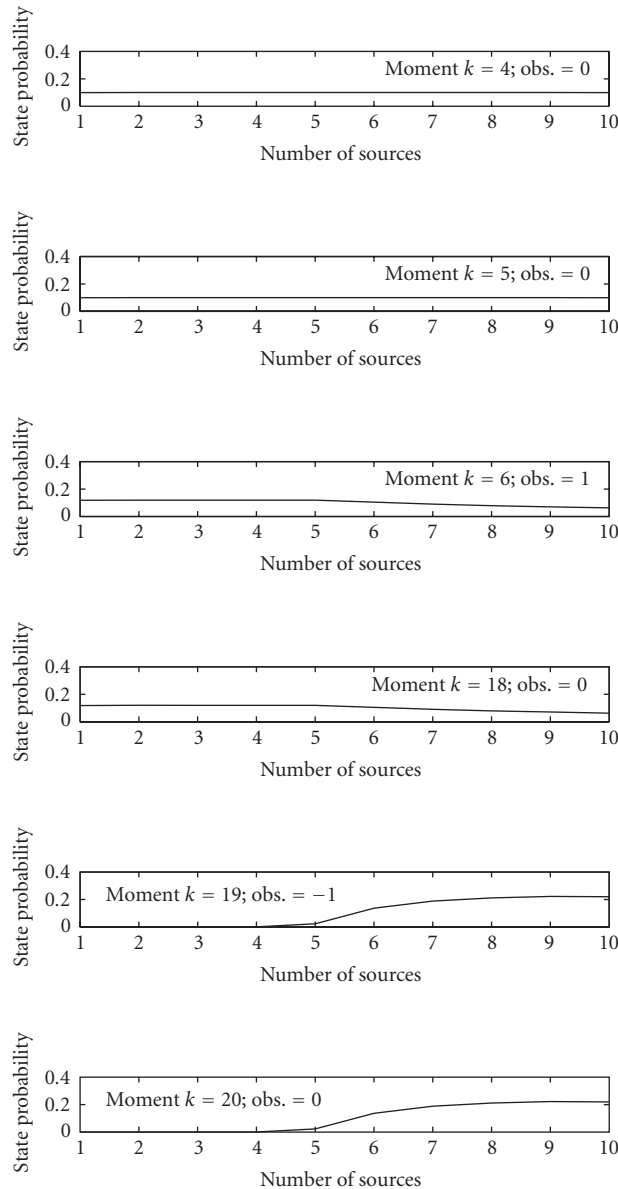


FIGURE 8: Illustrates typical dynamics of π . This plot corresponds to a symmetric birth-and-death chain as shown in Figure 6, with probability of switching to a different state $p = 0.001$, $N = 10$ sources, and loss threshold $T = 0.04$. At time 0, the initial π_0 is taken to be $\pi^s(i) = 1/N$, the stationary distribution of the underlying birth-and-death chain. While there are no communication attempts (up until time $k = 6$), π_k remains at π^s . Then at time 6, a packet is injected into the network and it is accepted, and as a result, there is a shift in the probability mass towards the region in which there is a small number of active sources. Then at time 19, another communication attempt takes place but this time the packet is rejected, and as a result, now the probability mass shifts to the region of a large number of active sources. This type of oscillations we have observed repeatedly, and gives a very pleasing intuitive interpretation of what the optimal controller does: keep pushing the probability mass to the left (because that is the region where more frequent communication attempts occur, and therefore leads to maximization of throughput), but dealing with the fact that losses push the mass back to the right. Similar oscillations are also typical of linear-increase multiplicative-decrease flow control algorithms such as the one used in TCP.

π_k has the Markov property itself. This is a Markov chain taking values in an uncountable space, though (the simplex Π).

- (2) Then we discretize the simplex Π . And we show that for all “small enough” discretizations, there is at least one observation taking π_k out of any cell with positive

probability. With this, we make sure that there are no absorbing cells, in the sense that once the chain hits that cell, it gets stuck there forever.

- (3) Then we show that the stationary distribution π^s of the underlying (finite-state) Markov chain is a point reachable from anywhere in the simplex. With this,

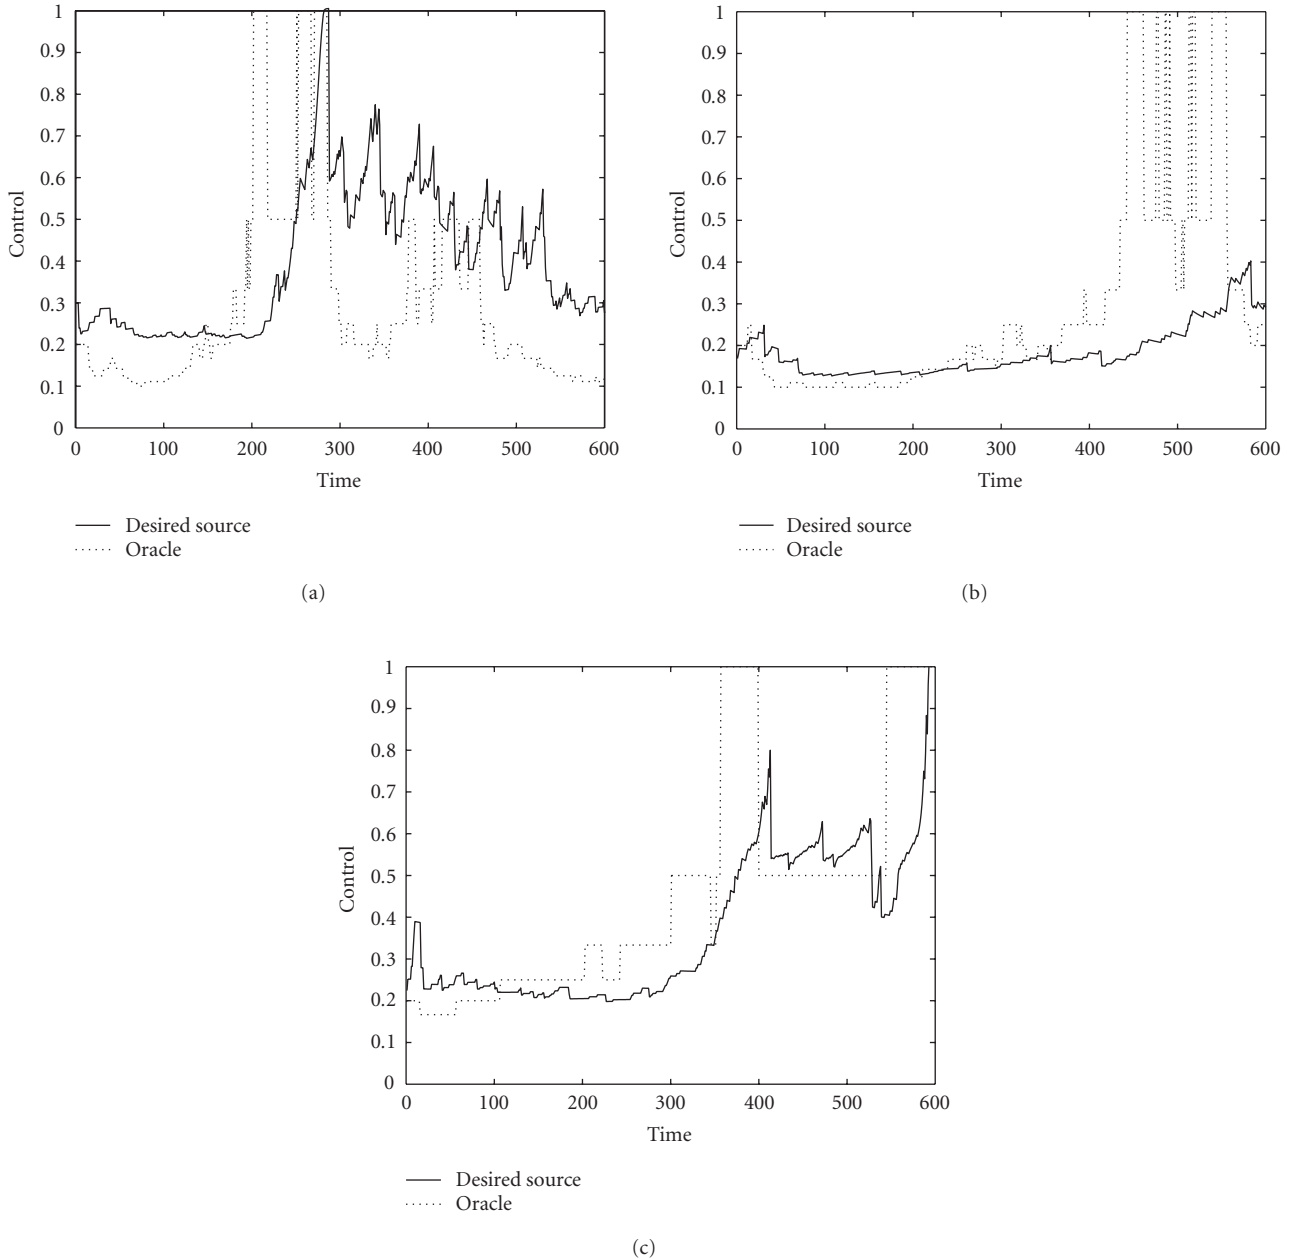


FIGURE 9: Illustrates how the value of the loss threshold T affects the optimal control law. In this case, we consider the same birth-and-death model considered in Figure 8, with three different values for T : top-left, $T = 0.1$; top-right, $T = 0.02$; bottom, $T = 0.05$. In all plots, the horizontal axis is time, the vertical axis is control intensity, and two controllers are shown: the thick black line corresponds to our optimal control law, the thin dotted line corresponds to a genie-aided controller that can observe the hidden state. And we observe a number of interesting things: (i) when T is large (a), our optimal control stays most of the time above the fair share point determined by the actions of the genie-aided controller; (ii) also when T is large, we see that sudden increases in bandwidth are quickly discovered by our optimal law; (iii) when T is small (b), the gap between the control actions of our optimal law and the genie-aided law is smaller, but our law has a hard time tracking a sudden increase in available bandwidth; (iv) for intermediate values of T (c), both the size of the gap and the speed with which changes in available bandwidth can be tracked are in between the previous two cases. These plots also suggest another intuitively very pleasing interpretation: T is a measure of how “aggressive” our optimal control law is.

we make sure that there is at least one cell which can be reached from any initial point in Π , and hence that the set of recurrent cells is not empty.

(4) Consider next any “small enough” discretization of the

space, and define a new process whose values are the cells of this discretization, based on whether π_k hits a particular cell. Then, this new process is (finite-state) Markov, and positive recurrent on a nonempty subset

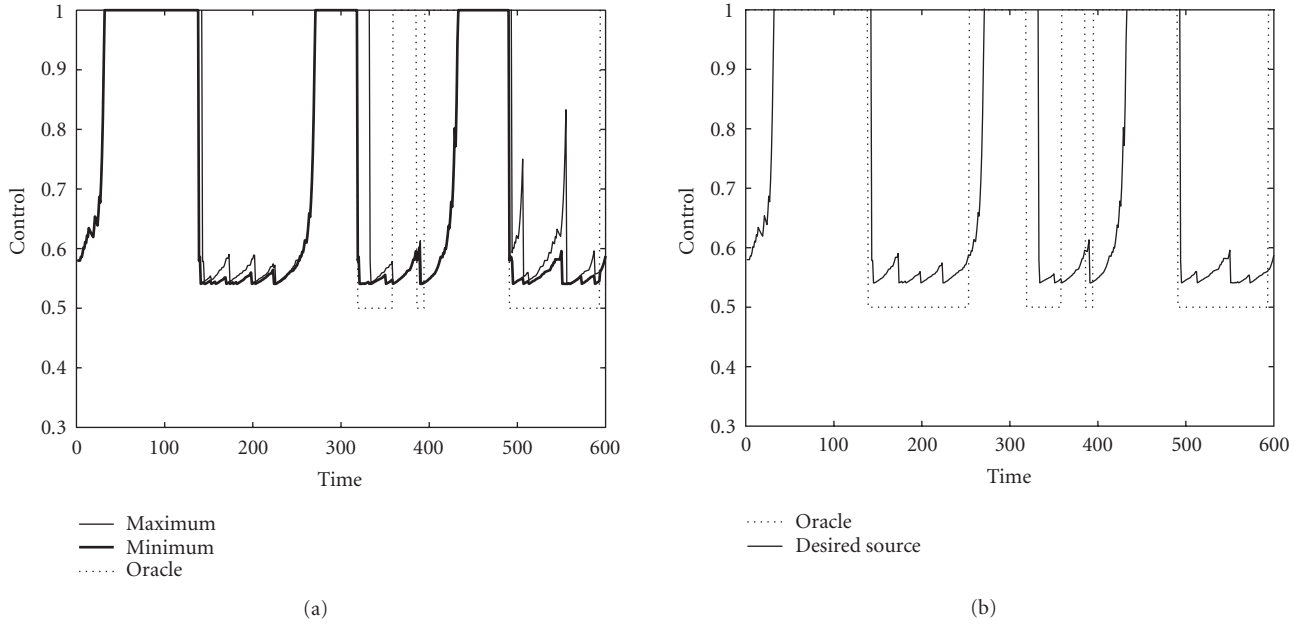


FIGURE 10: Illustrates the fairness issue raised at the end of Section 2.1. In this case, we also consider a birth-and-death chain model as in previous examples, but now with only two sources ($N = 2$). In (a), we show the maximum and the minimum control values chosen by either one of the sources over time: thick black line shows the minimum, thin solid line shows the maximum (for reference, the genie-aided controller is also shown); in (b), the thick line corresponds to the control actions of only one of the sources, all the time. Observe how, around time steps 150–250, the source shown at the bottom is the one that achieves the *maximum* at the top; but around time steps 500–600, the same source achieves the *minimum* of those injection rates. This is yet another intuitively very pleasing pattern that we have observed repeatedly in many simulations: the control law is essentially fair in the sense that, although we do not have enough information to make sure that at any time instant all controllers will use the same injection rate, at least over time the different controllers “take turns” to go above and below each other.




FIGURE 11: Illustrates how the original problem is broken into N independent identical subproblems. Since all the nodes execute exactly the same control algorithm, the distribution of π is the same for all nodes. But other than through this statistical constraint, all decisions are taken locally by each node, based on private data that is not available to any other node, and therefore completely independent.

of the cells, and therefore it admits an invariant measure itself.

- (5) Finally, we construct a measure as the limit of the “simple” measures from step 4 (as we let the size of the discretization vanish), and we show that this limit is invariant over Π . This requires some further steps, largely based on the elegant framework of [26] as follows.

- (5.1) We show that the limit exists and is well defined (it is independent of the particular sequence of discretizations considered).
- (5.2) We construct a simple φ -irreducibility measure on Π , and from there, we conclude the existence of a

unique maximal ψ -irreducibility measure.

- (5.3) We construct a family of accessible atoms in Π , and show that π_k is positive recurrent. From this and from 5.2, using a theorem from [26], we conclude that there exists a unique invariant measure on Π .
- (5.4) We show that the limit measure of (5.1) is indeed invariant, and therefore conclude that it must be the unique measure of (5.3).

Although steps 2–4 can be dealt with using classical finite-state Markov chain theory, steps 1 and 5 cannot. This is because π_k is a Markov chain defined on an (uncountable) metric space, and therefore to analyze its properties, we need




FIGURE 12: Looking at the evolution in time of the information state. For a system with three sources, π is a point in the 2D simplex as shown in this figure. And after letting the system run for some time, we find that there are regions of space visited fairly often (bottom right), regions visited less often (bottom left), and regions never visited (top right). Yet each point on this simplex determines a choice of an injection rate, and therefore the frequency with which each point is visited is clearly a fundamental performance analysis tool.

to resort to a more general theory of Markov processes. Meyn and Tweedie provide an excellent coverage of the problem of Markov chains on general spaces [26], which we found to be an invaluable tool in our work.

We continue now with the formal proofs for the steps of the proof above.

3.2. Weak convergence—steps 1–4

Step 1: π is Markov

Although involving a chain defined over a metric space, this proof is elementary, since all we need to invoke is the standard definition of the Markov property and the total probability law:

$$\begin{aligned}
& p(\pi_{k+1} \mid \pi_k, \dots, \pi_0) \\
& \stackrel{(a)}{=} \sum_{r_k} p(\pi_{k+1} \mid \pi_k, \dots, \pi_0, r_k) p(r_k \mid \pi_k, \dots, \pi_0) \\
& \stackrel{(b)}{=} \sum_{r_k} p(\pi_{k+1} \mid \pi_k, r_k, u_k) \sum_{x_k} p(r_k \mid x_k, \pi_k, \dots, \pi_0) \\
& \quad \times p(x_k \mid \pi_k, \dots, \pi_0) \\
& \stackrel{(c)}{=} \sum_{r_k} p(\pi_{k+1} \mid \pi_k, g(\pi_k), r_k) \\
& \quad \times \sum_{x_k} p(r_k \mid x_k, g(\pi_k)) \pi_k(x_k) \\
& = p(\pi_{k+1} \mid \pi_k),
\end{aligned} \tag{12}$$

where (a) results from the total probability law, (b) is because when conditioned on π_k , we can add in the conditioning u_k , since $u_k = g(\pi_k)$, and the total probability law, and (c) because conditioned on anything else, r_k depends only on x_k , u_k , and π_k contains all information about x_k given the

past. So we see that when conditioning on the past values, π_{k+1} depends only on π_k , and hence π is Markov.

An interesting observation to make here—which gives some insight into structural properties of our model that will allow us to prove the sought weak convergence result—is that the intensity of the arrivals process is a *memoryless* function of π . Although we have not attempted to prove this, it seems at least intuitively clear to us that if instead of the optimal controller we used a suboptimal one (typically based on the formation of an estimate of the current state), then the optimal decision would not be a memoryless function of the state estimate, but would actually require past state estimates as well.

Step 2: nonabsorbing small discretization cells

The next step is to show that there is a constant $C > 0$ such that, for any information state $\pi \in \Pi$, there exists an observation $r \in \mathcal{O}$ for which the distance between π_k and the next-step information state π_{k+1} corresponding to r is larger than C . This allows us to *quantize* the simplex Π and make sure that, provided that the size of a quantization cell is small enough, at least one observation will take the current information state to a different cell.

Lemma 2. *There exists a constant C such that for any $\pi \in \Pi$, there is an observation r for which $\|\pi - F[\pi, g(\pi), r]\| \geq \epsilon$, for all $0 < \epsilon \leq C$, and for any norm $\|\cdot\|$.*

Proof. This basically means that for any state, there is at least one observation that moves the chain at a finite nonzero distance away from that given state. We prove this by contradiction. We show that if all jumps are infinitesimally small, then the only information state that can satisfy this condition is the stationary distribution π^s of the original chain. But for this particular information state, any observation different from $r = 0$ does allow jumps of finite size away from it.

Suppose that for any $C > 0$, there exists a point $\pi \in \Pi$ such that for any observation $r \in \mathcal{O}$, $\|\pi - F[\pi, g(\pi), r]\| < C$. Denote by Q_C the set of points π verifying the above condition for a given C . Then, if $C_1 > C_2$, then $Q_{C_2} \subseteq Q_{C_1}$. Denote by Q_0 the intersection of all Q_C sets. Then the supposition that we want to contradict is equivalent to $Q_0 \neq \emptyset$.

Consider now any $\pi \in Q_0$. Then for any C arbitrarily close to 0, and for any observation r , $\|\pi - F[\pi, g(\pi), r]\| < C$ (all jumps are arbitrarily close to zero).

In what follows, $k_{(\cdot)}$ are normalizing constants. If $r = 0$, it results that π is arbitrarily close to πP . This means that π is arbitrarily close to π^s (the stationary distribution of P). Also, for $r = -1$, or $r = 1$, consider the respective $D(g(\pi), r)$ diagonal matrices. It results that π is arbitrarily close to $1/k_{\pi, r} \pi D(g(\pi), r) P$. But π is arbitrarily close to π^s as well, so it results that π^s is arbitrarily close to $1/k_{\pi^s, r} \pi^s D(g(\pi^s), r) P$. In the limit, $\pi^s = 1/k_{\pi^s, r} \pi^s D(g(\pi^s), r) P$. But this cannot be true because D is not the identity matrix. Actually, D is a diagonal matrix with increasing or decreasing diagonal elements (d_1, \dots, d_N) , for $r = 1$, respectively, $r = -1$. If, for example, $r = 1$, then $\pi^s = 1/k_{\pi^s, 1} \pi^s D(g(\pi^s), 1) P$.




FIGURE 13: A sequence of $r = 0$ observations leads the chain arbitrarily close to π^s .

This would mean that there exists $\pi_1 = 1/k_{\pi^s,1} \pi^s D(g(\pi^s), 1)$ with $\pi^s = \pi_1 P$. We know that $\pi^s = \pi^s P$, so if the chain admits only one stationary distribution, it results that $\pi_1 = \pi^s$. However this is not possible, since $\pi^s D(g(\pi^s), 1)$ moves towards $(1, 0, \dots, 0)$ the mass function of the new probability vector with respect to π^s . \square

Step 3: π^s is reachable from anywhere

Lemma 3. For any $\pi \in \Pi$, there is a nonzero probability that in the limit, the chain reaches arbitrarily close to the state π^s , when starting in state π .

Proof. We illustrate in Figure 13 the intuition on which we base our proof. The proof relies on the observation that finite-length sequences of $r = 0$ observations move the state arbitrarily closer to π^s . If the observation at time k is $r_k = 0$,

$$\begin{aligned} p(q_{k+1}^\epsilon | q_k^\epsilon, q_{k-1}^\epsilon, \dots) &= \int_{\pi_{k+1} \in q_{k+1}^\epsilon, \pi_k \in q_k^\epsilon, \pi_{k-1} \in q_{k-1}^\epsilon, \dots} p(\pi_{k+1} | \pi_k, \pi_{k-1}, \dots) d\pi_{k+1} d\pi_k d\pi_{k-1} \dots \\ &= \int_{\pi_{k+1} \in q_{k+1}^\epsilon, \pi_k \in q_k^\epsilon} p(\pi_{k+1} | \pi_k) d\pi_{k+1} d\pi_k \\ &= p(q_{k+1}^\epsilon | q_k^\epsilon) \end{aligned} \quad (13)$$

since the process π_k is Markov. The measure with respect to which we are integrating is Lebesgue measure over the Π space. We just count how often the continuous chain falls in a given cell. Thus the process $q^\epsilon(\pi_k)$ forms a *finite-state* chain, also having the Markov property (inherited from the continuous chain).

Lemma 4. For any $\epsilon \leq C$, there is a subset $\mathcal{P}^\epsilon \subseteq \Pi$, which contains the stationary distribution π^s of the x_k original chain, and on which the discretized chain $q^\epsilon(\pi_k)$ is positive recurrent.

Proof. We show in Figure 14 a typical behavior of the chain,

then the matrix $D(u_k, r_k)$ becomes diagonal with elements $d_{ii} = 1 - u_k$, so it equals the identity matrix multiplied with a constant; then the recursion for the information state, whenever the source decides not to transmit, can be expressed as $\pi_{k+1} = \pi_k P$. This vector equation has as solution the stationary distribution π^s . It follows that for any $\pi \in \Pi$ as initial state of the chain, there is a path by which the chain reaches in the limit the stationary distribution state π^s , via for example a sequence of successive $r_k = 0$ observations. But any arbitrary finite-length sequence of $r_k = 0$ observations may happen with nonzero probability, so for any $\epsilon^s > 0$, there is a finite time K with $r_k = 0$, $k \leq K$, in which the chain can reach with nonzero probability a state π^{ϵ^s} such that $\|\pi^{\epsilon^s} - \pi^s\| \leq \epsilon^s$. \square

Step 4: positive recurrent discretization on a nonempty subset

We consider now quantizations of the Markov chain formed by the sequence of information states, with quantization cells of size $\epsilon \leq C$. If the cell size is small enough, then from Lemma 2, it follows that for any π inside a discretization cell, there is at least one observation happening with nonzero probability for which the chain jumps outside the cell. This ensures that there is no state of the chain in which the system stays forever, so the recurrent irreducible subset of discretized cells has more than one element. With this procedure, we define a family of quantizations of Π , with members of the family of the form $q^\epsilon = \{q_1^\epsilon, \dots, q_{N_\epsilon}^\epsilon\}$, where q_i^ϵ are the N_ϵ compact sets contained in q^ϵ and $\bigcup_i q_i^\epsilon = \Pi$, $\bigcap_i q_i^\epsilon = \emptyset$. For simplicity, we will denote by $q^\epsilon(\pi)$ the cell to which the instantaneous information state π belongs. We note that

that shows the existence of a recursive subset $\mathcal{P}^\epsilon \subseteq \Pi$. We base our proof on the fact that π^s is recurrent so its properties will be induced on a recurrent closure of the discretized version of the simplex Π . As we showed in Lemma 3, the information state π^s can be reached in the limit with nonzero probability from any π_0 initial state of the Markov chain. For any initial π_0 , there is a sequence $\pi_0, \dots, \pi_k, \dots$ such that $\pi_k \rightarrow \pi^s$ when $k \rightarrow \infty$, and the size of the quantization cell is strictly positive. Then the time in which the discretized chain $q^\epsilon(\pi_k)$ reaches the cell containing state π^s is finite, so without loss of generality, we may consider our limit results with π^s as initial value for the information state.




FIGURE 14: After passing through a sequence of transient states, the chain reaches a recursive subset of the discretized simplex Π .

Denote by \mathcal{P}_ϵ the set of reachable quantization cells $q^\epsilon(\pi)$ if the chain starts in π^s . We already proved that the cell containing π^s is accessible in a finite number of steps from any other information state π , so implicitly from any cell $q_i^\epsilon \in \mathcal{P}_\epsilon$ as well. Moreover, by construction, any cell $q_i^\epsilon \in \mathcal{P}_\epsilon$ is reachable from $q^\epsilon(\pi^s)$. Since $\epsilon \leq C$, then for any $q^\epsilon(\pi_k)$, there is at least one observation r for which the transition from π_k to π_{k+1} leads to $q^\epsilon(\pi_{k+1}) \neq q^\epsilon(\pi_k)$. It follows that the chain with states in \mathcal{P}_ϵ is irreducible (and aperiodic as well, since the cell containing π^s is one-step reachable from itself, via an $r = 0$ observation). The state space is finite, so the chain is positive recurrent, and thus it has a stationary distribution. Denote by p_ϵ this probability distribution over the \mathcal{P}_ϵ state space. If $q_i^\epsilon \notin \mathcal{P}_\epsilon$, then $p_\epsilon(q_i^\epsilon) = 0$. \square

We will prove now that there is a limit probability measure on Π to which p_ϵ converges in the limit $\epsilon \rightarrow 0$, and study the properties of that measure.

3.3. Weak convergence—step 5

There exists a unique limit invariant measure over Π , $\nu_\epsilon \rightarrow \nu$, when $\epsilon \rightarrow 0$.

Step 5.1: existence of the limit measure

We will show that the limit measure exists by considering, for any subset A of Π , sequences of measures on subsets of the discretized simplex that cover, and respectively, intersect A . We show that they converge to the same limit.

Definition 2. Define the *inner* and *outer* sequences of measures over the simplex Π , corresponding to the set of ϵ -discretizations:

$$\begin{aligned} \nu_\epsilon^I(A) &= \sum_{S \in \mathcal{P}_\epsilon: S \subseteq A} p_\epsilon(S), \\ \nu_\epsilon^O(A) &= \sum_{S \in \mathcal{P}_\epsilon: S \cap A \neq \emptyset} p_\epsilon(S), \end{aligned} \quad (14)$$

where A is any subset in the σ -algebra of Π .

We want to prove that, for any given A , both $\nu_\epsilon^I(A)$ and $\nu_\epsilon^O(A)$ converge to the same limit, as $\epsilon \rightarrow 0$. That limit will be our limit invariant measure $\nu(A)$.

We will prove first convergence of each of the limits. Consider ν_ϵ^I ; we will prove that the sequence is Cauchy for any set A , and it trivially has a convergent subsequence, which will mean that the whole sequence is convergent.

For a given set A , denote by $A_n = \{\cup S \in \mathcal{P}_{\epsilon_n} : S \subseteq A\}$ the inner cover of set A corresponding to discretization step ϵ_n . We will prove first that the normalized volume of the difference set between two inner covers of the set A tends to the empty set, and consequently the probability measure over that difference set tends to zero. Define the metric $d(X, Y) = \mu^{\text{Leb}}((X - Y) \cup (Y - X)) / \mu^{\text{Leb}}(\Pi)$, on the σ -algebra \mathcal{B} of Π , $X, Y \in \mathcal{B}$ (this represents the normalized volume of the set where the two subsets X, Y differ from each other— μ^{Leb} is Lebesgue measure). It is easy to verify that $d(\cdot, \cdot)$ is indeed a valid metric.

Let ϵ_n be a decreasing sequence of discretization steps, with $\lim_{n \rightarrow \infty} \epsilon_n = 0$. Then due to the fact that \mathcal{P}_{ϵ_n} is a sequence of subsequent discretizations of the space Π when $n \rightarrow \infty$, it follows that $\lim_{n \rightarrow \infty} A_n = A$. Since A_n is convergent, it is also Cauchy in the metric space (\mathcal{B}, d) . This means that for any $\delta > 0$, there exists n_δ such that $d(A_n, A_m) < \delta$, for any $m > n \geq n_\delta$. So the normalized volume of the set difference between two set elements of the sequence becomes arbitrarily small. That also means that if $\epsilon_n, \epsilon_m \rightarrow 0$, then $\nu_{\epsilon_n \epsilon_m}^I(d(A_n, A_m)) \rightarrow 0$, as $\nu_{\epsilon_n \epsilon_m}$ is a stationary distribution over finite spaces with decreasing cell size. Then for any $\delta_\nu > 0$, there is n_{δ_ν} such that $\nu_{\epsilon_n \epsilon_m}^I((A_n - A_m) \cup (A_m - A_n)) < \delta_\nu$, for any $m > n \geq n_{\delta_\nu}$. Note that $\nu_{\epsilon_n \epsilon_m}^I((A_n - A_m) \cup (A_m - A_n)) \geq |\nu_{\epsilon_n \epsilon_m}^I(A_n) - \nu_{\epsilon_n \epsilon_m}^I(A_m)|$.

Finally, we note that $|\nu_{\epsilon_n \epsilon_m}^I(A_n) - \nu_{\epsilon_n \epsilon_m}^I(A_m)| = |\nu_{\epsilon_n}^I(A_n) - \nu_{\epsilon_m}^I(A_m)|$, due to the property of inclusion for the sequence of measures (sum of probabilities of ϵ' -discretization cells that cover exactly an ϵ -discretization cell is equal to the probability of that ϵ -discretization cell), and the way A_n, A_m are constructed (cells corresponding to multiples of ϵ are all included in the cells corresponding to ϵ).

We conclude that for any $\delta_\nu > 0$, there is n_{δ_ν} such that $|\nu_{\epsilon_n}^I(A_n) - \nu_{\epsilon_m}^I(A_m)| < \delta_\nu$, for any $m > n \geq n_{\delta_\nu}$. This means that the sequence $\nu_{\epsilon_n}^I(A_n)$ is Cauchy as well. It is trivial to show that there is a convergent subsequence of $\nu_{\epsilon_n}^I(A_n)$. Pick, for example, $\epsilon_n = \epsilon_0 / 2^n$, then the corresponding subsequence is bounded from above by 1, and monotonically increasing; it follows that the subsequence is convergent. But a Cauchy sequence with a convergent subsequence is convergent, which proves that $\nu_\epsilon^I(A)$ is convergent for any set A .

The proof for convergence of ν_ϵ^O is similar and we will omit it. Both limits exist, and it is obvious that they fulfill the inequality

$$\lim_{\epsilon \rightarrow 0} \nu_\epsilon^I(A) \leq \lim_{\epsilon \rightarrow 0} \nu_\epsilon^O(A), \quad \text{for any } A \subset \Pi. \quad (15)$$

We want to prove that the inequality above holds in fact with equality. Assume that the inequality is strict; then let $\delta = \nu_0^O(A) - \nu_0^I(A) > 0$. But this would mean that there exists at least a cell in any partition \mathcal{P}_ϵ of size $\delta > 0$, for all $\epsilon_n \rightarrow 0$. However, in the limit, the two sets of summation become equal (with union A), so a contradiction results.

Definition 3. Define the measure ν over the simplex Π as the common limit of the two sequences of measures:

$$\nu(A) = \lim_{\epsilon \rightarrow 0} \nu_{\epsilon}^I(A) = \lim_{\epsilon \rightarrow 0} \nu_{\epsilon}^O(A) \quad (16)$$

for a given $A \subseteq \Pi$.

For the proofs in the next two sections, we will use definitions and notations also found in [26].

Step 5.2: existence of a unique maximal ψ -irreducibility measure

Definition 4. Denote by $B(\pi_0, \delta) = \{\pi \in \Pi : \|\pi - \pi_0\| < \delta\}$ the open ball, with $\delta > 0$.

Definition 5. Denote by $\mathcal{B}(\Pi)$ the σ -field generated by the open balls in Π .

Definition 6. Denote, for any state $\pi \in \Pi$ and subset $A \in \mathcal{B}(\Pi)$, the probability that, when starting in state π , the chain reaches subset A :

$$L(\pi, A) = P_{\pi}(\tau_A < \infty). \quad (17)$$

Lemma 5. Let $\pi_n \in \Pi, n = 0, 1, \dots$ be a sequence of information states. Then π_n is ϕ -irreducible on $\mathcal{B}(\Pi)$.

Proof. Let π^s be the stationary distribution of the underlying chain. Define the measure ϕ on $\mathcal{B}(\Pi)$ as

$$\begin{aligned} \phi(B(\pi^s, \delta)) &= \mu^{\text{Leb}}(B(\pi^s, \delta)), \\ \phi(A) &= 0 \quad \text{otherwise.} \end{aligned} \quad (18)$$

In step 3 of the proof, we proved that π^s is reachable from anywhere. Hence, for all $\pi \in \Pi$, we have $L(\pi, B(\pi^s, \delta)) > 0$, and ϕ is an irreducibility measure. \square

Note. If a ϕ -irreducibility measure exists, then there is a part of the space reachable from anywhere, so one might expect independence of the chain from the initial conditions, by analogy with finite chains.

Proposition 1. If π_n is ϕ -irreducible, then there exists a unique “maximal” measure ψ on $\mathcal{B}(\Pi)$ such that π_n is ψ -irreducible and $\phi \leq \psi$. Denote by $\mathcal{B}^+(\Pi)$ the σ -algebra of Π with sets on which ψ is positive.

Proof. The proof involves concepts outside the scope of this paper, but is standard for chains fulfilling the previous conditions, and it can be found in [26]. \square

Step 5.3: uniqueness of the invariant measure on Π

Definition 7. $\alpha \in \mathcal{B}(\Pi)$ is called an *atom* for a sequence π_n if there exists a measure μ on $\mathcal{B}(\Pi)$, such that, for any $\pi \in \alpha$, $P(\pi, A) = \mu(A)$ (for any $A \in \mathcal{B}(\Pi)$).

Definition 8. α is called an *accessible atom* for a sequence π_n , if π_n is ψ -irreducible and $\psi(\alpha) > 0$.

Note. Atoms behave like states in finite chains. From the development in [26], it turns out that the reason why so many results about finite chains carry over to more general settings is precisely the fact that it is always possible to construct atoms.

Proposition 2. All balls $B(\pi^s, \delta)$, with $\delta > 0$, are accessible atoms for any sequence π_n .

Proof. Let α be a set in $\mathcal{B}(\Pi)$, and let $\pi \in \alpha$. Then, depending on the current observation r , there are three possible transitions from π , via the recursion function $F[\pi, g(\pi), r]$. Then for any $A \in \Pi$, we can consider the measure

$$\begin{aligned} \mu(A) &= p(r) \quad \text{if } F[\pi, g(\pi), r] \in A, \\ \mu(A) &= 0 \quad \text{otherwise.} \end{aligned} \quad (19)$$

Then any $\alpha = B(\pi^s, \delta)$ is an accessible atom. \square

Definition 9. Denote by $E_{\pi}[\eta_A]$ the expected number of returns of the chain to subset $A \in \Pi$ when starting in state π .

Definition 10. A set $A \in \mathcal{B}(\Pi)$ is called *recurrent* if $E_{\pi}[\eta_A] = \infty$ for all $\pi \in A$ (when starting in A , the expected number of returns to A is infinite).

Lemma 6. If π_n is ψ -irreducible and admits a recurrent atom α , then every set in $\mathcal{B}^+(\Pi)$ is recurrent.

Proof. If $A \in \mathcal{B}^+(\Pi)$, then for any π , there exist r, s such that $P^r(\pi, \alpha) > 0$, $P^s(\alpha, A) > 0$ and we can write, by considering the paths of the chain that go from π to A via the atom α ,

$$\sum_n P^{r+s+n}(\pi, A) \geq P^r(\pi, \alpha) \left[\sum_n P^n(\alpha, \alpha) \right] P^s(\alpha, A) = \infty. \quad (20)$$

Since α , being an atom, implies that $\sum_n P^n(\alpha, \alpha)$ diverges. \square

Note. Observe again the analogy between atoms and states of a finite chain.

Definition 11. A sequence π_n is called *recurrent* if and only if it is ψ -irreducible, and $E_{\pi}[\eta_A] = \infty$ for any $\pi \in \Pi$ and $A \in \mathcal{B}^+(\Pi)$.

Lemma 7. Any sequence of information states drawn from the Markov chain π_n is recurrent.

Proof. From Lemma 5 and Proposition 1, it results that π_n is ψ -irreducible. Furthermore, from step 3 of the proof, it results that all the balls $B(\pi^s, \delta)$, with $\delta > 0$, are recurrent atoms. Then every $A \in \mathcal{B}^+(\Pi)$ is recurrent, and from Definition 10, it results that $E_{\pi}[\eta_A] = \infty$ for all $\pi \in \Pi$. We still need to prove that even if $\pi \notin A$, we still have $L(\pi, A) = 1$. By definition, $L(\pi, \cup \mathcal{B}^+(\Pi)) = 1$. Suppose that the sequence π_n hits at some time a set $B \in \mathcal{B}^+(\Pi)$.

If $A = B$, then the sought result follows. Otherwise, we will have $L(y, A) > 0$ for all $y \in B$, because of the ψ -irreducibility over \mathcal{B}^+ . But $B \in \mathcal{B}^+(\Pi)$ and $E_y[\eta_B] = \infty$, so it results that $L(y, A) = 1$. So finally, $L(\pi, A) = 1$ and this case is reduced to the previous one (where $\pi \in A$). Hence, π_n is recurrent. \square

Definition 12. A sequence π_n is called *positive* if and only if it is ψ -irreducible and admits an invariant measure γ .

Lemma 8 (Kac's theorem). *If a sequence π_n is recurrent and admits an atom $\alpha \in \mathcal{B}^+(\Pi)$, then π_n is positive if and only if $E_\alpha[\tau_\alpha] < \infty$.*

Proof. If $E_\alpha[\tau_\alpha] < \infty$, then obviously $L(\alpha, \alpha) = 1$, so it results π_n is recurrent. It also results from the structure of γ (see [26]) that γ is finite, so is positive as well. The converse results from the structure of γ as well. \square

Lemma 9. *The sequence of information states π_n is positive.*

Proof. From Lemma 7, it results that π_n is recurrent. Also, from step 3 of the proof, it results that every ball $\alpha = B(\pi^s) \in \mathcal{B}^+(\Pi)$ is an atom, and $E_\alpha(\tau_\alpha) < \infty$. Then it results that π_n is positive. \square

Theorem 10. *There exists a unique invariant probability measure of π_n .*

Proof. The proof for this theorem is valid for chains having the properties we have analyzed until now, and can be found in [26]. \square

Step 5.4: invariance of ν

We prove now Theorem 1. The measure ν (as constructed in step (5.1)) is the unique invariant probability measure on Π .

Proof. For invariance of ν , we need to prove that

$$\nu(A) = \int_{\Pi} \nu(dy)P(y, A). \quad (21)$$

From the definition of ν , we have that for any $\epsilon > 0$,

$$\nu_\epsilon^I(A) \leq \nu(A) \leq \nu_\epsilon^O(A). \quad (22)$$

If we denote by $P_\epsilon(\cdot, \cdot)$ the transition-probability kernel for the ϵ -discretization, then we can rewrite the rightmost term of the inequality (22) as

$$\begin{aligned} \nu_\epsilon^O(A) &= \sum_{S \in \mathcal{P}_\epsilon: S \subseteq A} p_\epsilon(S) \\ &= \sum_{S \in \mathcal{P}_\epsilon: S \subseteq A} \sum_{T \in \mathcal{P}_\epsilon} p_\epsilon(T)P_\epsilon(T, S) \\ &= \sum_{T \in \mathcal{P}_\epsilon} \sum_{S \in \mathcal{P}_\epsilon: S \subseteq A} p_\epsilon(T)P_\epsilon(T, S) \\ &= \sum_{T \in \mathcal{P}_\epsilon} p_\epsilon(T)P_\epsilon(T, \cup\{S \in \mathcal{P}_\epsilon : S \subseteq A\}). \end{aligned} \quad (23)$$

In a similar manner, we can rewrite the expression for $\nu_\epsilon^O(A)$:

$$\nu_\epsilon^O(A) = \sum_{T \in \mathcal{P}_\epsilon} p_\epsilon(T)P_\epsilon(T, \cup\{S \in \mathcal{P}_\epsilon : S \cap A \neq \emptyset\}). \quad (24)$$

By taking now the limit in expression (22), we know that both left and right limits exist and are equal, so it results that

$$\begin{aligned} \nu(A) &= \lim_{\epsilon \rightarrow 0} \sum_{T \in \mathcal{P}_\epsilon} p_\epsilon(T)P_\epsilon(T, \cup\{S \in \mathcal{P}_\epsilon : S \subseteq A\}) \\ &= \lim_{\epsilon \rightarrow 0} \sum_{T \in \mathcal{P}_\epsilon} p_\epsilon(T)P_\epsilon(T, \cup\{S \in \mathcal{P}_\epsilon : S \cap A \neq \emptyset\}) \\ &\stackrel{(a)}{=} \int_{\Pi} \nu(dy)P(y, A), \end{aligned} \quad (25)$$

where equality (a) holds because, under some continuity conditions, in the limit $\epsilon \rightarrow 0$, the sum becomes integral; the probability limit ν exists; the quantization cell $T \in \mathcal{P}_\epsilon$ becomes the infinitesimal integration variable $T \rightarrow dy$; the transition-probability kernel $P_\epsilon(\cdot, \cdot) \rightarrow P(\cdot, \cdot)$; and both the reunions of cells included in A and, respectively, intersecting A cover whole set A .

It results that ν is invariant, and thus it is the unique invariant measure on Π . \square

3.4. Numerical simulations

In this section, we show results of numerically evaluating the integrals above. We simulated a system with $N = 2$, $N = 4$, and $N = 8$ sources, and with different values for the loss threshold $T = 0.02$, $T = 0.05$, and $T = 0.1$. The chain is birth-and-death with probability P_{switch} . We let the system run for $t = 100\,000$ time steps. We plot the average throughput and loss as a function of the transition probability P_{switch} . The resulting plots are shown in Figure 15. We see that the plots do not depend significantly on P_{switch} . The dependence is essentially on the stationary probability π_s of the original chain P , which is the same for any symmetric birth-and-death chain. As expected, large values for T imply larger throughput, as the controller is allowed to probe more often the environment; this is on the expense of increased losses.

4. CONCLUSIONS

We proposed a new mechanism for rate control in sensor networks, based on partial observation about the state of the system. We show that, when the system state is Markov, the optimal controller essentially depends on the information state, a quantity that takes into account all previous controls and observations about the system. Then, we show results on the convergence of the information state, which imply stability of our control policy. Namely, (a) we formulated a queueing problem in which the process of arrivals is *not* independent of the (partially observed) state of the queue, (b) we solved the corresponding optimal control problem, and (c) we proved a theorem regarding its ergodic behavior—the existence of a suitable invariant measure. Our main insight to

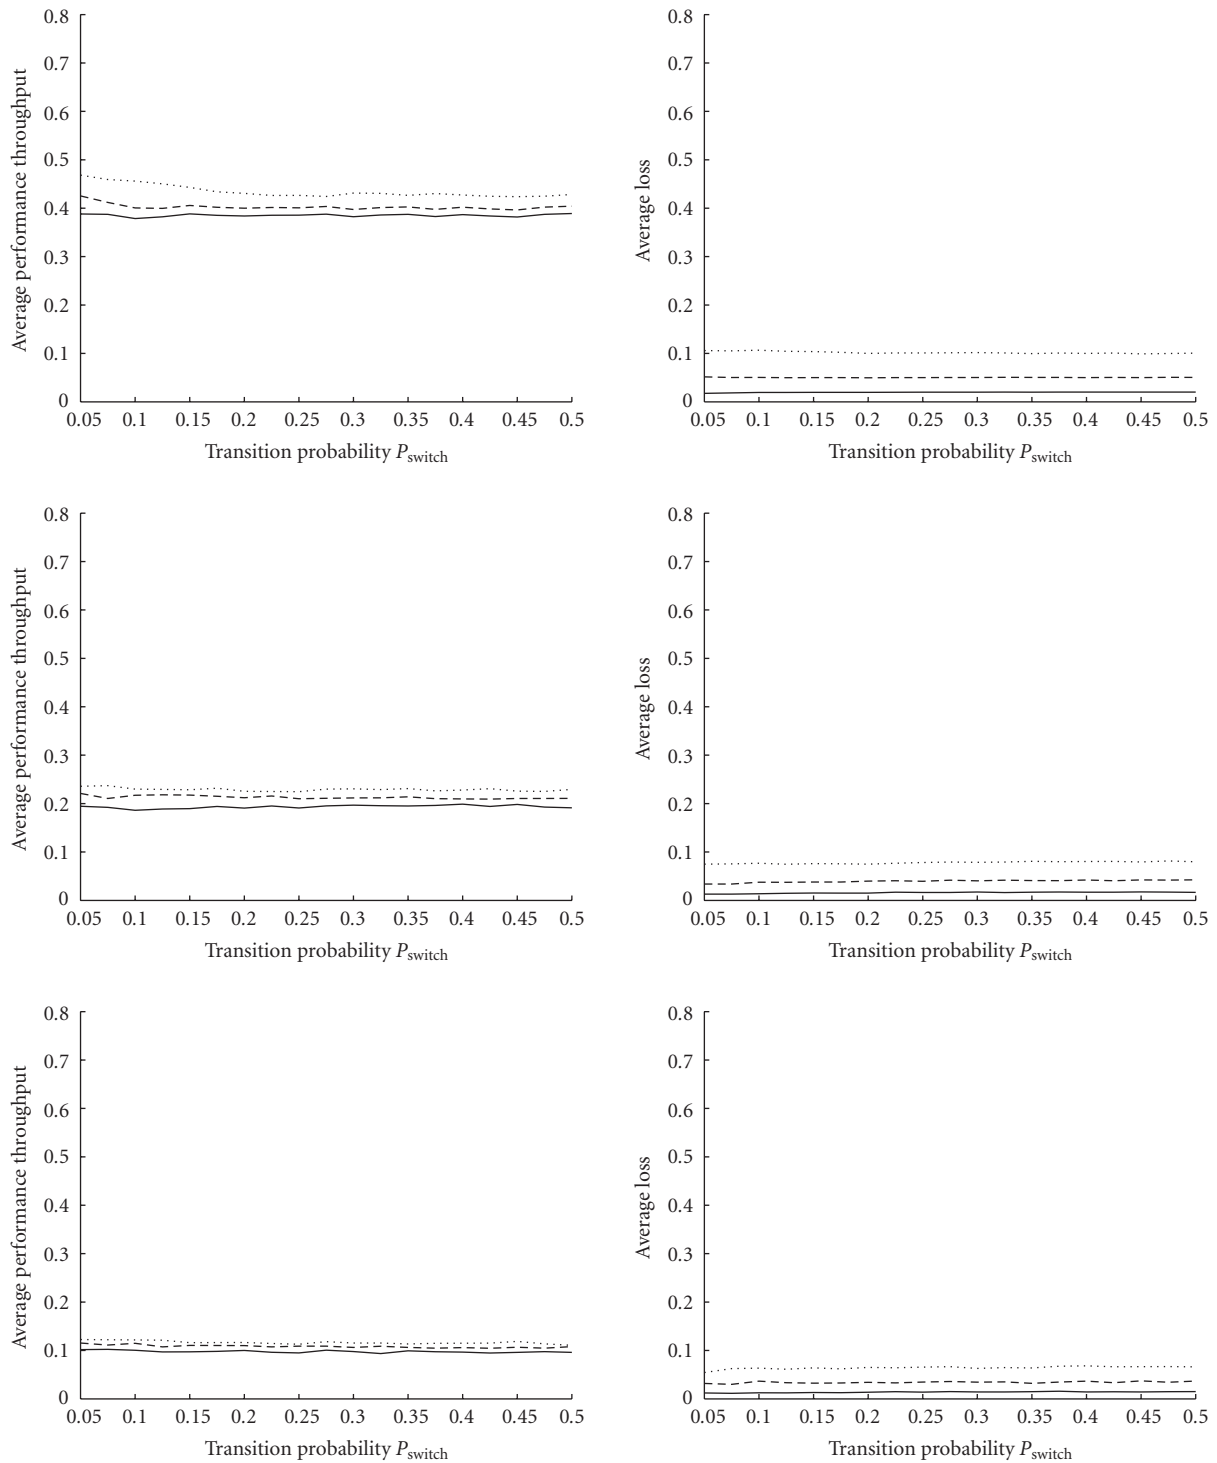


FIGURE 15: The plots from up to bottom: $N = 2, 4, 8$ sources. Left: average throughput; right: average loss. Legend: dotted plot, $T = 0.1$; dashed plot, $T = 0.05$; solid plot, $T = 0.01$.

tackle these problems was that by conditioning on information states, the arrivals process does become a process with independent increments—but since this conditioning term is itself Markov, this is how the model is rendered analytically tractable.

An interesting conclusion that we draw from our results relates to the observation that to make a system efficient in an information-theoretic sense requires driving the system very close to the instability point. Consistent with this observation, for our model to be efficient, the queue needs to be

driven *past* the stability point and into the instability region. This is because, without being able to observe the state of the system but being able to observe when the system becomes unstable, this instability is the only indication that we are operating at peak efficiency. Note that this is a key idea behind the implementation of TCP (increase window size while packets get acknowledged, decrease when packets get lost). In our model, without forcing “TCP compatibility,” we observe the exact same type of behavior. This provides a strong argument for the (additive probing probability)/(multiplicative back-off) paradigm in the design of MAC access protocols in sensor networks, where only limited partial observations of the state system are available to the medium access controllers.

ACKNOWLEDGMENTS

This work was largely completed while both authors were with l'École Polytechnique Fédérale de Lausanne (EPFL), and supported primarily by the Swiss National Science Foundation under Grant 21-61831.00 (with additional support from the National Science Foundation under Grant CCR-0227676). Parts of this work have been presented at the 35th Annual Conference on Information Sciences and Systems (CISS 2001), and at the IEEE International Symposium on Information Theory (ISIT 2002).

REFERENCES

- [1] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, Englewood Cliffs, NJ, USA, 2nd edition, 1992.
- [2] S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *IEEE/ACM Trans. Networking*, vol. 1, no. 4, pp. 397–413, 1993.
- [3] V. Jacobson, “Congestion avoidance and control,” in *Proc. ACM Symposium on Communications Architectures and Protocols (ACM SIGCOMM '88)*, pp. 314–329, Stanford, Calif, USA, August 1988.
- [4] F. Kelly, “Mathematical modelling of the Internet,” in *Mathematics Unlimited - 2001 and Beyond*, B. Engquist and W. Schmid, Eds., pp. 685–702, Springer, Berlin, Germany, 2001, available from <http://www.statslab.cam.ac.uk/frank/PAPERS/>.
- [5] R. J. La and V. Anantharam, “Utility-based rate control in the Internet for elastic traffic,” *IEEE/ACM Trans. Networking*, vol. 10, no. 2, pp. 272–286, 2002.
- [6] S. H. Low and D. E. Lapsely, “Optimization flow control. I. Basic algorithm and convergence,” *IEEE/ACM Trans. Networking*, vol. 7, no. 6, pp. 861–874, 1999.
- [7] S. H. Low, F. Paganini, and J. C. Doyle, “Internet congestion control,” *IEEE Control Syst. Mag.*, vol. 22, no. 1, pp. 28–43, 2002.
- [8] Y. Sankarasubramaniam, I. F. Akyildiz, and S. W. McLaughlin, “Energy efficiency based packet size optimization in wireless sensor networks,” in *Proc. 1st IEEE International Workshop on Sensor Network Protocols and Applications (SNPA '03)*, pp. 1–8, Anchorage, Alaska, USA, May 2003.
- [9] C.-Y. Wan, S. Eisenman, and A. Campbell, “CODA: COngestion Detection and Avoidance in sensor networks,” in *Proc. ACM SenSys '03*, Los Angeles, Calif, USA, November 2003.
- [10] M. Carvalho and J. J. Garcia-Luna-Aceves, “A Scalable model for channel access protocols in multihop Ad-hoc networks,” in *Proc. 10th Annual ACM International Conference on Mobile Computing and Networking (MobiCom '04)*, Philadelphia, Pa, USA, September–October 2004.
- [11] V. Naware and L. Tong, “Smart antennas, dumb scheduling for medium access control,” in *Proc. 37th Annual Conference on Information Sciences and Systems (CISS '03)*, Baltimore, Md, USA, March 2003.
- [12] S. Singh and C. S. Raghavendra, “PAMAS: power aware multi-access protocol with signalling for Ad-hoc networks,” *ACM Computer Communication Review*, vol. 28, no. 3, pp. 5–26, 1998.
- [13] A. Woo and D. Culler, “A transmission control scheme for media access in sensor networks,” in *Proc. 7th Annual ACM International Conference on Mobile Computing and Networking (Mobicom '01)*, pp. 221–235, Rome, Italy, July 2001.
- [14] L. Bao and J. J. Garcia-Luna-Aceves, “Transmission scheduling in Ad-hoc networks with directional antennas,” in *Proc. 8th Annual ACM International Conference on Mobile Computing and Networking (MobiCom '02)*, pp. 48–58, Atlanta, Ga, USA, September 2002.
- [15] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, “Energy-efficient communication protocol for wireless microsensor networks,” in *Proc. 33rd IEEE Annual Hawaii International Conference on System Sciences (HICSS '00)*, vol. 2, Maui, Hawaii, USA, January 2000.
- [16] E.-S. Jung and N. H. Vaidya, “A power control MAC protocol for Ad-hoc networks,” in *Proc. 8th Annual ACM International Conference on Mobile Computing and Networking (MobiCom '02)*, pp. 36–47, Atlanta, Ga, USA, September 2002.
- [17] K. Sohrabi and G. J. Pottie, “Performance of a novel self-organization protocol for wireless Ad-hoc sensor networks,” in *Proc. 50th IEEE Vehicular Technology Conference (VTC '99)*, vol. 2, pp. 1222–1226, Amsterdam, Netherlands, September 1999.
- [18] W. Ye, J. Heidemann, and D. Estrin, “Medium access control with coordinated adaptive sleeping for wireless sensor networks,” *IEEE/ACM Trans. Networking*, vol. 12, no. 3, pp. 493–506, 2004.
- [19] S. Mitter, “Control with limited information,” *European Journal of Control*, vol. 7, no. 2-3, pp. 122–131, 2001.
- [20] S. Tatikonda, *Control Under Communication Constraints*, Ph.D. thesis, M. I. T. Press, Cambridge, Mass, USA, 2000.
- [21] T. Kaijser, “A limit theorem for partially observed Markov chains,” *The Annals of Probability*, vol. 3, no. 4, pp. 667–696, 1975.
- [22] A. J. Goldsmith and P. P. Varaiya, “Capacity, mutual information, and coding for finite-state Markov channels,” *IEEE Trans. Inform. Theory*, vol. 42, no. 3, pp. 868–886, 1996.
- [23] V. Sharma and S. K. Singh, “Entropy and channel capacity in the regenerative setup with applications to Markov channels,” in *Proc. IEEE International Symposium on Information Theory (ISIT '01)*, Washington, DC, USA, June 2001.
- [24] B. Hajek, “Stochastic approximation methods for decentralized control of multiaccess communications,” *IEEE Trans. Inform. Theory*, vol. 31, no. 2, pp. 176–184, 1985.
- [25] P. R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1986.
- [26] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*, Springer, London, UK, 1993.

Răzvan Cristescu was born in Ploiesti, Romania, in 1974. He received his Ph.D. degree from l'École Polytechnique Fédérale de Lausanne (EPF, Lausanne) in 2004, his Lic. Tech. in 2000 from Helsinki University of Technology, and his M.S. degree in 1998 from Polytechnic University of Bucharest. Since October 2004, he has been a Postdoctoral Fellow in the Center for the Mathematics of Information at California Institute of Technology.

His research interests include information theory, network theory, computational complexity, signal processing for communications, and applications in sensor networks.



Sergio D. Servetto was born in Argentina, on January 18, 1968. He received a Licenciatura en Informatica from Universidad Nacional de La Plata (UNLP, Argentina) in 1992, and the M.S. degree in electrical engineering, and the Ph.D. degree in computer science from the University of Illinois at Urbana-Champaign (UIUC), in 1996 and 1999. Between 1999 and 2001, he worked at the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

Since fall 2001, he has been an Assistant Professor in the School of Electrical and Computer Engineering at Cornell University, and a member of the field of applied mathematics. He was the recipient of the 1998 Ray Ozzie Fellowship, given to "outstanding graduate students in computer science," and of the 1999 David J. Kuck Outstanding Thesis Award for the best doctoral dissertation of the year, both from the Department of Computer Science at UIUC. He is also the recipient of a 2003 NSF Career Award. His research interests are centered around information-theoretic aspects of networked systems, with a current emphasis on problems that arise in the context of large-scale sensor networks.



Special Issue on

Wireless Access in Vehicular Environments

Call for Papers

Wireless access in vehicular environments (WAVE) technology comes into sight as a state-of-the-art solution to contemporary vehicular communications, which is anticipated to be widely applied in the near future to radically improve the transportation environment in aspects of safety, intelligent management, and data exchange services. WAVE systems will fundamentally smooth the progress of intelligent transportation systems (ITSs) by providing them with high-performance physical platforms. WAVE systems will build upon the IEEE 802.11p standard, which is still active and expected to be ratified in April, 2009. Meanwhile, the VHF/UHF (700MHz) band vehicular communication systems are attracting increasingly attention recently.

The fast varying and harsh vehicular environment brings about several fresh research topics on the study of WAVE systems and future vehicular communication systems, which include physical layer challenges associated with mobile channels, capacity evaluation, novel network configuration, effective media access control (MAC) protocols, and robust routing and congestion control schemes.

The objective of this special section is to gather and circulate recent progresses in this fast developing area of WAVE and future vehicular communication systems spanning from theoretical analysis to testbed setup, and from physical/MAC layers' enabling technology to network protocol. These research and implementation activities will be considerably helpful to the design of WAVE and future vehicular communications systems by removing major technical barriers and presenting theoretical guidance. This special issue will cover, but not limited to, the following main topics:

- 700MHz/5.8GHz vehicle-to-vehicle (V2V)/vehicle-to-infrastructure (V2I) single-input single-output (SISO)/MIMO channel measurement and modeling, channel spatial and temporal characteristics exploration
- Doppler shift study, evaluation and estimate, time and frequency synchronizations, channel estimate and prediction
- Utilization of MIMO, space-time coding, smart antenna, adaptive modulation and coding
- Performance study and capacity analysis of V2V and V2I communications operating over both 5.8GHz and 700Mz
- Software radio, cognitive radio, and dynamic spectrum access technologies applied to WAVE and future vehicular communication systems
- Mesh network and other novel network configurations for vehicular networks
- Efficient MAC protocols development

- Routing algorithms and congestion control schemes for both real-time traffic warning message broadcasting and high-speed data exchange
- Cross-layer design and optimization
- Testbed or prototype activities

Authors should follow the EURASIP Journal on Wireless Communications and Networking manuscript format described at the journal site <http://www.hindawi.com/journals/wcn/>. Prospective authors should submit an electronic copy of their complete manuscript through the EURASIP Journal on Wireless Communications and Networking's Manuscript Tracking System at <http://mts.hindawi.com/>, according to the following timetable.

| | |
|------------------------|-----------------|
| Manuscript Due | April 1, 2008 |
| First Round of Reviews | July 1, 2008 |
| Publication Date | October 1, 2008 |

Guest Editors

Weidong Xiang, University of Michigan, Dearborn, USA; xwd@umich.edu

Javier Gozalvez, University Miguel Hernández, Spain; j.gozalvez@umh.es

Zhisheng Niu, Tsinghua University, China; niuzhs@tsinghua.edu.cn

Onur Altintas, Toyota InfoTechnology Center, Co. Ltd, Tokyo, Japan; onur@jp.toyota-itc.com

Eylem Ekici, Ohio State University, USA; ekici@ece.osu.edu



CALL FOR PAPERS

The 2008 European Signal Processing Conference (EUSIPCO-2008) is the sixteenth in a series of conferences promoted by EURASIP, the European Association for Signal, Speech, and Image Processing (www.eurasip.org). Formerly biannual, this conference is now a yearly event. This edition will take place in Lausanne, Switzerland, organized by the Swiss Federal Institute of Technology, Lausanne (EPFL).

EUSIPCO-2008 will focus on the key aspects of signal processing theory and applications as listed below. Exploration of new avenues and methodologies of signal processing will also be encouraged. Accepted papers will be published in the Proceedings of EUSIPCO-2008. Acceptance will be based on quality, relevance and originality. Proposals for special sessions and tutorials are also invited.

Areas of Interest

- *Audio and Electroacoustics*
- *Design and Implementation of Signal Processing Systems*
- *Image and Multidimensional Signal Processing*
- *Multimedia Signal Processing*
- *Signal Detection and Estimation*
- *Sensor Array and Multichannel Processing*
- *Signal Processing for Communications*
- *Speech Processing*
- *Education in Signal Processing*
- *Nonlinear Signal Processing*
- *Medical Imaging and Image Analysis*
- *Signal Processing Applications (Biology, Geophysics, Seismic, Radar, Sonar, Remote Sensing, Astronomy, Bio-informatics, Positioning etc.)*
- *Emerging Technologies*

Best Paper Awards

There will be a best paper and a student paper contest. Student authors who appear as first authors in a paper may enter the best student paper contest.

Submissions

Procedures to submit a paper, proposals for special sessions/tutorials, will be detailed at www.eusipco2008.org. Submitted papers must be camera-ready, final, no more than five pages long all inclusive and conforming to the format that will be specified on the EUSIPCO web site.

Important Dates

- Proposals for Special Sessions and Tutorials: **December 7, 2007**
- Electronic submission of Full papers (5 pages A4): **February 8, 2008**
- Notification of Acceptance: **April 30, 2008**
- Submission of Camera-Ready Papers: **May 31, 2008**
- Conference: **August 25-29, 2008**

www.eusipco2008.org



European Association
for Signal, Speech and
Image Processing

General Chairman

J.-Ph. Thiran, EPFL, Switzerland
JP.Thiran@epfl.ch

General Co-Chairman

P. Vanderghyest, EPFL, Switzerland
Pierre.Vanderghyest@epfl.ch

Technical Program Co-Chairmen

P. Frossard, EPFL, Switzerland
Pascal.Frossard@epfl.ch

A. Cavallaro, Queen Mary, University of
London, UK
andrea.cavallaro@elec.qmul.ac.uk

Plenary Talks

M. Unser, EPFL, Switzerland
Michael.Unser@epfl.ch

S. Godsill, Cambridge Univ, UK
sjq@eng.cam.ac.uk

Special Sessions

C. De Vleeschouwer, UCL, Belgium
devlees@tele.ucl.ac.be

Tutorials

N. Paragios, Ecole Centrale, France
nikos.paragios@ecp.fr

F. Bimbot, IRISA, Rennes, France
Frederic.Bimbot@inria.fr

Publications

R. Reilly, UC Dublin, Ireland
richard.reilly@ucd.ie

Local Arrangements

M. Marion, EPFL, Switzerland
marianne.marion@epfl.ch

Finances

M. Bach Cuadra, EPFL, Switzerland
meritxell.bach@epfl.ch

Publicity

G. Olmo, Politecnico di Torino, Italy
gabriella.olmo@polito.it

M. Gabbouj, Tampere UT, Finland
Moncef.Gabbouj@tut.fi

International Liaisons

B. Macq, UCL, Belgium
macq@tele.ucl.ac.be

U.B. Desai, Indian Inst. of Tech., India
ubdesai@ee.iitb.ac.in

D. Erdogmus, Orgon H&S University, USA
derdogmus@ieee.org

Exhibits

J. Reichel, Spinetix S.A., Switzerland,
reichel@ieee.org



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

l a u s a n n e