ON A GENERAL RULE FOR ACCESS CONTROL*
OR, SILENCE IS GOLDEN . . .

Yechiam Yemini
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90291
USA

Leonard Kleinrock
University of California, Los Angeles
Computer Science Department
Los Angeles, CA 90024
USA

ABSTRACT

We consider the problem of multi-access control of a shared communication channel in a multi-hop Packet Radio Network. A decentralized notion of optimal channel-sharing is introduced and necessary conditions for optimal transmission policies are derived. These conditions possess an intuitively simple yet powerful interpretation of a balance principle. The optimality rules are very general and imply, as specific instances, a diverse class of known rules such as Abramson's optimal Slotted-Aloha and the optimal Urn scheme, previously introduced by the authors. Moreover, the rules lend themselves to a simple implementation of a class of decentralized hierarchical decision algorithms, which includes some Slotted-Aloha control algorithms that have previously been explored as well as new, though still unexplored, algorithms. The proposed algorithms decompose the problem into a short-term distributed decision problem of adapting transmission rights to local loads and the determination of long-range centralized policies to obtain global goals (e.g., priorities).

*This research has been supported by the Advanced Research Projects Agency of the Department of Defense under contract number MDA903-77-C-0272.

INTRODUCTION

The environment and the model

In what follows we consider a multi-hop Packet Radio Network (PRNET) [1,2,3,4,5] whose member Packet Radio Units (PRUs), $PR_1, PR_2, \ldots PR_N$, all use a common radio communication channel to communicate with each other. Data packets are passed around the network in a hop-by-hop fashion until they reach their destinations. At each hop, a data packet is stored and then forwarded using a broadcast transmission over the common channel. Not all PRUs can hear each other due to range and topography limitations. The topology of channel sharing is specified in terms of a *hearing graph* H, whose i-th vertex represents PR_i and whose edges represent the existence of a communication link (i.e., a hop) between the respective PRUs. Note that we implicitly assume the hearing relation between PRUs to be symmetric. This assumption is not necessary for the analysis that we carry out, but it simplifies the computational details.

We assume that the communication channel is time-slotted, i.e., transmissions are synchronized to slots. The channel time slots may be slightly longer than a packet transmission time and possibly of variable length (i.e., to account for propagation delays and/or access control mechanisms) [5,6]. This assumption, while essential for the analysis carried out in this paper, does not present a severe limitation upon our approach and results, which can be generalized to non-slotted channels.

Packets arrive at the different PRUs from external, *bursty* [7,8], sources. They are stored and forwarded to their destinations, following some routing algorithm. For simplicity we assume that all packets are destined to a single sink and that the routing is fixed (at least over the time period of interest). Thus the routing is completely described by specifying for each PRU, say PR_i, its immediate next hop destination $PR_{d(i)}$. Again, the assumption of fixed routing may be easily replaced with a model of dynamic routing (e.g., by replacing the single destination function d(i) with a probability distribution over the immediate destinations [9]; this method may also be used to account for multiple destinations); however, it simplifies the computational model.

PR_i is said to *interfere* with PR_j if $PR_{d(j)}$ hears PR_i. Interference is a nonsymmetric binary relation upon the nodes of the PRNET. We shall use the notation $I(j) = \{i \mid PR_{d(j)} \text{ hears } PR_i\}$ to denote the set of all PRUs that interfere with PR_j. We shall assume that when PR_j transmits a packet, it succeeds iff all PRUs in I(j) are silent. Again it is possible to include capture effects in the model at the price of some extra computational details.

With the above model, the PRNET is a large distributed service facility that shares the communication resource, i.e., channel slots. An *access scheme* is an algorithm to decide which *busy* PRUs (i.e., those having transmission-pending packets in their buffer) should have the right to transmit during the coming slot.

Let us define the space of all possible randomized actions that the PRNET may choose during each slot as $A = [0,1]^N$ where an *action* (also *transmission policy*) $\underline{p} \in A$, $\underline{p} = (p_1, p_2, \ldots p_N)$ is a choice of a probability $0 \leq p_i \leq 1$ for each PR_i. The probability p_i is that of PR_i having a *transmission right* during the coming slot. A busy PRU tosses a coin with a transmission-bias p_i and transmits according to the result.

The objective of the decision mechanism is to maximize the rate of successful transmissions (i.e., those not interfered with) by selecting a proper transmission policy. Note that our model is general enough to describe any multi-access scheme to a shared, time-slotted channel. One scheme and another differ only in the method used for selecting a transmission policy.

Decentralization of the objective

The key constraint to the decision problem that the PRNET faces is the absence of a centralized mechanism that is aware of the network state and that may decide and enforce the decision upon the network members.

There are a few possible approaches to the problem of decentralization. One possible approach is to retain the objective function, derive a centralized solution, and relax the result to enable decentralized implementation. In what follows we assume a dual approach. Namely, we retain the decentralization but relax the objective to enable a solution.

The idea is simple: if the community of PRUs is incapable of optimizing a centralized performance index, let us replace the objective with local, decentralized, objective functions. A natural choice for local performance measures is to let PR_i choose p_i so as to maximize his individual throughput. Formally, let $\underline{\tilde{b}}^t = (\tilde{b}_1^t, \tilde{b}_2^t, \ldots \tilde{b}_N^t)$* designate the occupancy process, that is,

$$\tilde{b}_i^t \triangleq \begin{cases} 1 & \text{if } PR_i \text{ has a packet ready for transmission, at slot } t. \\ 0 & \text{otherwise.} \end{cases}$$

Let $\pi^t(\underline{b})$ designate the distribution of $\underline{\tilde{b}}^t$.

We define the throughput process

$$\tilde{s}_i^t \triangleq \begin{cases} 1 & \text{if } PR_i \text{ successfully delivers a packet to his destination at slot } t. \\ 0 & \text{otherwise.} \end{cases}$$

---

*We use $\tilde{a}$ and $\bar{a}$ to denote a random variable and its mean, respectively. Vectors are underlined and sets are capitalized.

We can express the mean throughput of PR_i, conditioned upon $\tilde{\underline{b}}^t = \underline{b}$, as follows:

$$(1)\quad s_i^t(\underline{b},\underline{p}) = E[\tilde{s}_i^t \mid \tilde{\underline{b}}^t = \underline{b}] = \begin{cases} p_i \prod_{j \in I_i(\underline{b})} (1-p_j) & \text{if } b_i = 1 \\ 0 & \text{if } b_i = 0 \end{cases}$$

where

$$(2)\quad I_i(\underline{b}) \triangleq \{ j \mid j \neq i,\ b_j = 1 \text{ and } PR_j \text{ interferes with } PR_i \}$$

The expected throughput of PR_i at slot t, is given by

$$(3)\quad \bar{s}_i^t(\underline{p}) = \sum_{\underline{b} \in \{0,1\}^N} \pi(\underline{b})\, s_i^t(\underline{b},\underline{p})$$

where $\pi^t(\underline{b})$ is the distribution of $\tilde{\underline{b}}^t$.

To simplify the notation we shall eliminate the time indexing t from all expressions whenever there is no danger of confusion. Equation (3) defines a continuous map $\underline{S} = \underline{S}(p)$ of policies onto attainable throughputs at slot t (here $\underline{S} = (\bar{s}_1, \bar{s}_2, \ldots \bar{s}_N)$). The hypercube of policies $A = [0,1]^N$ is mapped onto a compact domain $S = \underline{S}(A)$ of all attainable throughputs. We shall call the map $\underline{S} = \underline{S}(\underline{p})$: the *Abramson throughput operator*.

We shall consider $\bar{s}_i(\underline{p})$ as the *payoff function* that PR_i attaches to the policy $\underline{p}$. The problem of centralized optimization is replaced with a decentralized reconciliation of the individual objective functions of community members. This is a typical problem of mathematical economics. We shall assume the standard tools thereof to derive a solution [10,11].

A throughput vector $\underline{S}$ is called *Pareto optimal* iff

a) it is *attainable*, i.e., $\underline{S} = \underline{S}(\underline{p})$ for some transmission policy $\underline{p}$.

b) it is not dominated by any other throughput, i.e., there exists no attainable $\underline{S}'$, such that $\underline{S}' > \underline{S}$. (Here $\underline{S}' > \underline{S}$ means that $\forall i\ \bar{s}_i \geq \bar{s}'_i$, with at least one strict inequality.)

A policy $\underline{p}$ that obtains a Pareto-optimal throughput is called *Pareto-optimal policy*. Pareto-optimal policies are precisely those policies for which we cannot improve the payoff of one PRU without decreasing the payoff of some fellow PRUs.

With the model of the previous section and the decentralized notion of optimality introduced in this section, we are now ready to solve the problem of optimal decentralized policies. In the following section we characterize Pareto-optimal policies using some standard analysis. We then interpret the results in terms of a surprisingly simple rule of optimality.

Let us reconsider the Abramson map of the previous section:

$$(1) \quad \underline{S} = \sum_{\underline{b} \in \{0,1\}^N} \pi(\underline{b})\, \underline{S}(\underline{b},\underline{p})$$

Consider a Pareto optimal policy $\underline{p}^o$ obtaining a throughput $\underline{S}^o = \underline{S}(\underline{p}^o)$. Let $\underline{p}$ be an admissible transmission policy differing from $\underline{p}^o$ by a small perturbation $\Delta\underline{p} \triangleq \underline{p}-\underline{p}^o$. If $\underline{S}$ is the throughput obtained by $\underline{p}$, then $\underline{S}$ differs from $\underline{S}^o$ by a small perturbation $\Delta\underline{S} \triangleq \underline{S}-\underline{S}^o$. The conditional throughput $\underline{S}(\underline{b},\underline{p})$, as given by equation (1) of the previous section, is a smooth function of $\underline{p}$. Therefore $\Delta\underline{S}$ is related to $\Delta\underline{p}$ through a linear transformation

$$(2) \quad \Delta\underline{S} = \partial\underline{S}\, \Delta\underline{p}$$

described by the matrix

$$(3) \quad \partial\underline{S} \triangleq \sum_{\underline{b}} \pi(\underline{b}) \frac{\partial \underline{S}(\underline{b},\underline{p}^o)}{\partial \underline{p}} = \left[\sum_{\underline{b}} \pi(\underline{b}) \frac{\partial s_i(\underline{b},\underline{p}^o)}{\partial p_j} \right]_{ij}$$

which we call the *Jacobian matrix of the network*. The transformation described by (2) is a linear approximation of the nonlinear Abramson map near the Pareto optimal policy $\underline{p}^o$.

Let $D(\underline{p}^o)$ denote the set of admissible perturbations of $\underline{p}^o$, i.e., those perturbations corresponding to feasible policies. Clearly, $\underline{0} \in D(\underline{p}^o)$. Moreover, if $\underline{p}^o$ is an internal point of the set of admissible policies, then $D(\underline{p}^o)$ contains a neighborhood of zero. If the Jacobian matrix (3) is non-singular, then the image of $D(\underline{p}^o)$ under the map defined by the Jacobian matrix, i.e., the set of admissible perturbations of $\underline{S}^o$, contains a neighborhood of zero. This contradicts the extremality of $\underline{S}^o$. Therefore the Jacobian matrix of the network must be singular at $\underline{p}^o$.

What if $\underline{p}^o$ is not an internal point of A? If $\underline{p}^o$ is internal to any face of the hypercube A, then it is possible to show that by properly restricting the Abramson operator to a subnetwork, the resulting operator must have a singular Jacobian matrix. The demonstration involves some lengthy combinatorial arguments that are of no interest for us. Therefore we shall proceed to derive necessary conditions for Pareto optimality assuming that $\underline{p}^o$ is an internal point of A. The conditions that we derive may be easily verified to hold when $\underline{p}^o$ is not internal to A and even when $\underline{p}^o$ is an extreme point of A.

We conclude: If $\underline{p}^o$ is a Pareto-optimal policy obtaining a throughput $\underline{S}^o$, then the Jacobian determinant of the network (or a proper subnetwork) must be zero. That is,

$$(4)\quad 0 = \left| \sum_{\underline{b}} \pi(\underline{b}) \frac{\partial s_i(\underline{b},\underline{p}^o)}{\partial p_j} \right|$$

The generic elements of the Jacobian are easily computed to be

$$(5)\quad \underline{S}_{ij} \triangleq \begin{cases} \dfrac{E_i}{1-p_i} & i=j \\[2ex] \dfrac{S_{i/j}}{1-p_j} & i \neq j \end{cases}$$

where

$$E_i \triangleq \sum_{\{\underline{b} \mid b_i=1\}} \pi(\underline{b})(1-p_i) \prod_{j \in I_i(\underline{b})} (1-p_j)$$

is the *expected number of slots that are empty* at the destination of PR_i given that it is busy, and

$$S_{i/j} = \begin{cases} \displaystyle\sum_{\{\underline{b} \mid b_i=b_j=1\}} \pi(\underline{b})\, p_i \prod_{k \in I_i(\underline{b})} (1-p_k) & \text{if } j \text{ interferes with } i \\[4ex] 0 & \text{if } j \text{ does not interfere with } i \end{cases}$$

is the *expected number of successful packets* that PR_i delivers given that PR_j is busy and that PR_j interferes with PR_i; it is 0 otherwise.

Let us consider a typical one-hop network (in which all PRUs hear each other). In this case the optimality condition (4) may be expressed as

$$(6)\quad 0 = |\partial \underline{S}| = \begin{vmatrix} E_1 & -S_{1/2} & -S_{1/3} & \cdots & -S_{1/N} \\ -S_{2/1} & E_2 & -S_{2/3} & \cdots & -S_{2/N} \\ -S_{3/1} & -S_{3/2} & E_3 & \cdots & -S_{3/N} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ -S_{i/1} & -S_{i/2} & \cdots E_i \cdots & & -S_{i/N} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ -S_{N/1} & -S_{N/2} & \cdots & \cdots & E_N \end{vmatrix}$$

APPLICATIONS

The ecology of channel sharing

The optimality conditions of the previous section represent the rational behavior of a resource-sharing network community. If the users are too rude, the common resource may be polluted with collisions. If the users are too polite, they may leave much of the resource under-utilized. Pareto-optimal policies make the best expected use of the channel in the sense that every slot wasted in silence by one PRU is, on the average, utilized by another. The pollution of the channel by collisions or its under-utilization during some slots is not the result of imprudent network behavior but rather the consequence of the statistical fluctuations of the service demands. Pareto-optimality thus represents maximal prudence in the face of lady luck.

The optimality conditions are expressed in terms of a set of Lagrangian multipliers that we interpreted in the previous section as throughput prices. For each Pareto-optimal policy there exists a price vector $\underline{c}$ for which "silence" equals "throughput" at each network node. Looked at in another way, it is possible to select a price vector and consider Pareto-optimal policies corresponding to that pricing policy. By properly choosing a pricing policy for the network, we may approximate global objectives through the decentralized optimization mechanism.

For instance, let us consider an extreme case of dynamic price adjustment whereby a set of PRUs is selected during each slot to have infinite cost, while all other PRUs are assigned a zero cost. By properly selecting the PRUs whose cost is infinite, a perfect slot-by-slot scheduling of transmissions is obtained. This merely hypothetical scheme is meant only to illustrate the potential power of dynamic pricing.

In the absence of a centralized scheduling mechanism, it is still possible to adapt the pricing to the long-range statistics of the network state and/or global priorities. The decomposition of the decision mechanism to local individual decisions (i.e., adjustment of individual transmission policies to balance "silence" with "throughput") and to a global pricing mechanism generates a useful control hierarchy. Short-term decisions requiring rapid adaptation are arrived at by local PRUs, while long-range, global adaptation is decided by higher (more coordinated) mechanisms (e.g., through pricing).

Optimal Slotted-ALOHA policies

Let us consider the case of a one-hop PRNET where the transmission policy is fixed. If the queuing system is to be stable, then it should be able to properly handle a *heavy traffic condition*. We shall employ our general optimality conditions to derive the heavy-traffic results of Abramson [3]. The heavy-traffic condition may be expressed in terms of the occupancy distribution. Namely, $\pi(\underline{b}) = \delta(\underline{b},\underline{1})$, where $\underline{1}$ is the vector all of whose coordinates are 1 and $\delta(\underline{b},\underline{1})$ is Dirac's delta distribution, concentrated on the occupancy-state $\underline{1}$.

If all PRUs are identical, then we may apply the optimality condition expressed by equation (7) of the previous section. That is, "silence" is equal to $(1-p)^N$ and "throughput" equals $(N-1)p(1-p)^{N-1}$. The two expressions are equal to each other iff $p=1/N$, which is an instance of Abramson's optimality conditions. If the PRUs are not identical, the more general condition $p_1 + p_2 + \ldots p_N = 1$ may be obtained if we note that the Jacobian determinant of the network reduces to the Jacobian determinant obtained by Abramson [3] (Abramson's G_i correspond under heavy traffic to p_i).

In particular, consider the symmetric case when all PRUs are identical; the determinant becomes (after some easy algebra)

$$|\partial S| = [\hat{E}-(N-1)\hat{S}]\,[\hat{E}+\hat{S}]^{N-1}$$

where $\hat{E} \triangleq E_1 = E_2 = \ldots = E_N$ and $\hat{S} \triangleq S_{i/j}$ for all i and j ($i \neq j$).

This expression is zero iff

$$(7) \quad \hat{E} = (N-1)\,\hat{S}$$

The left-hand side represents the expected number of slots that a busy PRU leaves empty at his destination; we shall call this expression *silence*. The right-hand side represents the amount of throughput, produced by the rest of the network, that a busy PRU sees. We shall call this last expression *throughput*. Thus, *a necessary condition for an optimal selection of transmission policies is that each PRU equate "silence" with "throughput"*.

We conclude for the symmetric one-hop PRNET that if a policy is Pareto-optimal, then each busy PRU trades the slots that he wastes in silence for an equal number of slots successfully used by all those PRUs that may be harmed by his transmission.

Let us return to the determinant of the general network (4). If it is to be zero, there should be a linear combination of its rows that yields a zero vector. Let us denote the coefficients of such a linear combination $\underline{c} = (c_1, c_2, \ldots c_N)$. The optimality condition

$$c_1 E_1 = \sum_{\{j \mid 1 \in I(j)\}} c_j S_{j/1}$$

$$c_2 E_2 = \sum_{\{j \mid 2 \in I(j)\}} c_j S_{j/2}$$

$$\vdots$$

$$c_N E_N = \sum_{\{j \mid N \in I(j)\}} c_j S_{j/N}$$

may be interpreted as follows. Each PRU PR_i is endowed with a slot dollar cost c_i reflecting his relative significance in the network. The left-hand side of each expression represents the value of silence of the respective PR_i. The right-hand side represents the value of success for the busy PRUs whose transmission PR_i may harm, given that PR_i is busy. *The optimality condition requires that each PRU, once busy, should consider the needs of his fellow PRUs. He should be ready to trade his silence for an equivalent dollar-worth amount of their throughput, by selecting his transmission policy to balance the two quantities.*

Equating the two quantities we find that k should satisfy

$$\frac{1}{n-1} = \frac{k}{N-k-n+1}$$

from which follows

$$k = \frac{N-n+1}{n}$$

This is precisely the optimal choice of k for the urn scheme as given by equation (14) of section 2.2.2 in [12]. Therefore, a busy PRU should select k so as to equate his expected silence with the throughput of others.

DECENTRALIZED ALGORITHMS TO OPTIMIZE TRANSMISSION POLICIES

The characterization of optimal transmission policies, obtained in the previous section, may serve as a basis for a set of distributed access control algorithms. The algorithms are quasi-static in the sense of [14]. The decision making is completely decentralized. The information required for a decision is available to each PRU, at no extra cost, through the acknowledgment mechanism.

The algorithms consist of a decentralized iterative process that tries to balance "throughput" and "silence". The problem is essentially that of solving a large stochastic system of balance equations through gradient iterations. The details of The algorithms consist of a decentralized iterative process that tries to balance "throughput" and "silence". The problem is essentially that of solving a large stochastic system of balance equations through gradient iterations. The details of the algorithms, as well as the problems of convergence, are beyond the scope of this paper and will be presented in the future.

Let us recall the optimality rule: For each PR_i

(1) $\hat{E}_i = \hat{S}_i$

where $\hat{E}_i$ is the expected dollar-value of empty slots at the destination of PR_i, when PR_i is busy. $\hat{S}_i$ is the expected dollar-value of the throughput of those PRUs with which PR_i may interfere when it is busy.

The algorithm to implement the rule (1) consists of

1. *Estimation* Each PRU gathers acknowlegments statistics during his busy periods. We assume that both successful packet deliveries and collisions are acknowledged. PR_i can monitor the acknowledgments sent by all PRUs that he may hear. Therefore PR_i observes $\tilde{S}^i_j$ for all PR_j such that $i \in I(j)$, from which he may compute $\tilde{S}_i = \sum c^i_j \tilde{S}^i_j$, where the summation extends over all PR_j with which PR_i may interfere ($i \in I(j)$). The conditional expectation of S^i_j given that PR_i is busy, is precisely the "throughput"

The optimal URN-Scheme

Let us consider the optimal Urn-scheme developed in [12,13]. We consider a one-hop PRNET in which each member is aware, at each slot, of the total number n of busy PRUs (or an estimate of n). Of all the PRUs, k are selected to have transmission rights, while the rest keep quiet. Each PRU selects the same k lucky PRUs by drawing k numbers from a pseudo-random number generator. All PRUs use the same generator and the same seed (which may be a simple time function), thus deciding on identical k "lucky" PRUs. Therefore, the selection of the k PRUs is decentralized and randomized, yet coordinated. The urn-scheme provides for a smooth adaptivity to the load by adjusting the value of k as a function of n.

The name "urn-scheme" derives from the analogy to the problem of sampling k balls from an urn containing n black (for busy) balls and N-n white (for idle) balls. The objective is to choose k to maximize the probability that the sample contains exactly one black ball. The optimal Urn-scheme selects the value of k to be $\lfloor N/n \rfloor$.* For instance, when n=1, all k=N PRUs are given the right of transmission (but only one will actually transmit); when n approaches N (i.e., heavy-traffic), k approaches 1, i.e., only one PRU is given the right of transmission (due to the heavy load, it is very likely to be a busy PRU). It was shown in [12,13] that the urn-scheme converges to optimal Slotted-ALOHA for light traffic and to Time-Division-Multiple-Access for heavy traffic, while outperforming both for a medium load.

In [13] we have described one possible implementation of the urn-scheme. We shall now see that the optimal value of k, selected by the urn-scheme, is the very value of k that satisfies the optimality conditions of the previous section.

We consider the decision making process from the point of view of a given busy PRU. "Silence" occurs if all k owners of transmission rights happen to be other nonbusy PRUs. The probability that this occurs, given that our PRU is busy, is

$$E_i = \frac{\binom{N-k-1}{n-1}}{\binom{N-1}{n-1}}$$

The probability of a successful use of a slot by another PRU, given that our designated PRU is busy, is given by

$$S_i = \frac{\binom{k}{1}\binom{N-k-1}{n-2}}{\binom{N-1}{n-1}}$$

*Here $\lfloor x \rfloor$ is the integer part of x.

that we wish to estimate. The estimation of the last parameter can follow standard methods for estimating point processes [15].

Similarly, "silence" can be estimated by monitoring the acknowledgements sent by $PR_{d(i)}$ (in fact, by monitoring slots in which d(i) does not acknowledge anything). Again the problem is that of estimating the conditional expected value of an observed point process.

If only successful packets are to be acknowledged, then E_i may be estimated from the unconditional expected silence at $PR_{d(i)}$ to be monitored by the latter. If the acknowledgement mechanism is not collision free, further sophistication must be introduced into the estimation mechanism.

2. *Adaptation* Here the rule is simple: if $\hat{E}_i > \hat{S}_i$, then PR_i knows that he wastes too many slots, which nobody else uses anyway, in silence. PR_i will increase p_i. If $\hat{E}_i < \hat{S}_i$ then, by the same token, PR_i knows that he is talking too much, preventing fellow PRUs from getting a fair portion of the channel. PR_i should decrease p_i.

There is only one problem: the rule (1) is necessary but insufficient. For example, the rude policy $\underline{p} = (1, 1, \ldots 1)$ satisfies the rule independently of the input structure giving $E_i = S_i = 0$. Clearly, our algorithm may lead the network to choose this policy even when the results are disastrous. We have to design some precautionary measures to prevent our algorithm from converging to the rude policy when it should not.

The required modifications are simple: Each PRU monitors his own throughput. If, as a result of an increase in p_i, PR_i watches his throughput dropping, he knows that he may have increased it beyond the optimal value. The natural response is to decrease the value of p_i.

The above control mechanism may be implemented in a similar fashion to the control mechanisms presented by L. Kleinrock and M. Gerla in [16,17]. Moreover, when the traffic becomes heavy the algorithm will be an implementation of Abramson's optimality criteria. Thus, it will become identical to the mechanisms of [16,17].

In a similar fashion we may use the optimality rule to implement a version of the Urn scheme. Indeed, as demonstrated in the previous section, the optimal selection of k (the number of PRUs possessing a transmission right) satisfies the optimality rule. Therefore, k may be adjusted using the information acquired from the acknowledgment traffic only. The algorithm to implement the scheme consists of two parts similar to those above

1. *Estimation* same as before.

2. *Adaptation* If $\hat{E}_i < \hat{S}_i$, then using the expressions for $\hat{E}_i$ and $\hat{S}_i$, given in the previous section, we find that

$$\frac{N-n+1}{n} < k$$

That is, k is too large; PR_i should lower his estimate of k. Similiarly, if $\hat{E}_i > \hat{S}_i$, PR_i should increase his estimate of k.

By combining the above decentralized decision mechanism with a higher-order pricing mechanism, it is possible to establish priority mechanisms over the network. This possibility of establishing a family of hierarchical resource-sharing mechanisms is an appealing subject for further research and experiments. Further work is also required to develop the details of the estimation and adaptation mechanisms, prove the convergence of the algorithms, compare their performance in the context of one-hop systems with that of known control schemes [4, 12, 13, 16, 17, 18, 19,], and test them in a multi-hop environment.

CONCLUSIONS

We have seen that Pareto-optimality provides an excellent norm for rational decentralized resource-sharing. Using this criterion, we have derived an intuitively simple, yet powerful, rule-of-thumb for optimal, decentralized, multi-access control. Not only does the optimality rule encompass previous results (e.g., optimal Slotted-ALOHA and the optimal Urn Scheme), but it also enhances our understanding of proper hierarchical, decentralized, channel-allocation policies. Moreover, it yields a class of simple distributed decision mechanisms to implement the optimal policies. Further research is required to explore the stability of the proposed access-control mechanisms and to apply our methods to solve other problems of decentralized resource-sharing in computer communication networks.

REFERENCES

[1] R. E. Kahn, "The Organization of Computer Resources into a Packet Radio Network," *IEEE Transactions on Communications*, Vol. Com-25, No. 1, January 1977.

[2] L. Kleinrock, *Queueing Systems*, Vols. I and II, John Wiley & Sons, 1975 and 1976.

[3] N. Abramson, "Packet Switching Via Satellite," *AFIPS Conference Proceedings*, NCC (1973), pp. 695-702.

[4] S. S. Lam, *Packet Switching in Multi-access Broadcast Channel with Applications to Satellite Communications in Computer Network*, Computer Science Department, School of Engineering and Applied Science, Report UCLA-ENG-7429, March 1974.

[5] F. A. Tobagi, *Random Access Techniques for Data Transmission Over Packet Switched Radio Networks*, Computer Science Department, School of Engineering and Applied Science, Report UCLA-ENG-7499, December 1974.

[6] M. Scholl, Multiplexing Techniques for Data Transmission Over Packet Switched Radio Systems, Ph.D Dissertation, Computer Science Department, School of Engineering and Applied Science, Report UCLA-ENG-76123, December 1976.

[7] G. Akavia, Hierarchical Organization of Computer Communication Networks, Ph.D Dissertation, Computer Science Department, School of Engineering and Applied Science, UCLA, 1978.

[8] S. S. Lam, "A New Measure to Characterize Data Traffic," *IEEE Transactions on Communications*, Vol. Com-26, pp. 137-140, January 1978.

[9] L. Kleinrock, *Communication Nets: Stochastic Message Flow and Delay*, McGraw-Hill, 1964 (reprinted by Dover, 1972).

[10] R. D. Luce and H. Raiffa, *Games and Decisions*, John Wiley & Sons, 1967.

[11] G. Debreu, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, Cowles Foundation Monograph 17, Yale University Press, 1959.

[12] Y. Yemini, On Channel Sharing In Discrete-Time, Packet-Switched, Multi-access Broadcast Communication, Ph.D Dissertation, Computer Science Department, School of Engineering and Applied Science, UCLA, 1978.

[13] L. Kleinrock and Y. Yemini, "An Optimal Adaptive Scheme for Multiple-Access Broadcast Communication," *Proceedings of ICC Conference*, pp. 7.2.1-7.2.5, Toronto, Canada, June 1978.

[14] R. Gallager, "A Minimum Delay Routing Algorithm Using Distributed Computation," *IEEE Transactions on Communications*, Vol. Com-25, January 1977.

[15] A. Segall, "Recursive Estimation from Discrete-Time Point Processes," *IEEE Transactions on Information Theory*, Vol. IT-22, July 1976.

[16] M. Gerla and L. Kleinrock, "Closed Loop Stability Controls for S-ALOHA Satellite Communications," *Fifth Data Communications Symposium*, Snowbird, Utah, September 1977.

[17] L. Kleinrock and M. Gerla, "On the Measured Performance of Packet Satellite Access Schemes," *Proceedings ICCC Conference*, Tokyo, Japan, September 1978.

[18] G. Fayolle, E. Gelenbe, and J. Labetoulle, "The Stability and Optimal Control of the Packet Switching Broadcast Channel," *Journal of the ACM*, Vol. 24, No. 3, July 1977, pp. 375-386.

[19] T. A. Banh and E. Gelenbe, "Near Optimal Behavior of the Packet Switching Broadcast Channel," *Rapport de Recherche*, No. 299, Iria Laboria, May 1978.