The Designer's Perspective to Atomic Noncooperative Networks

Lavy Libman and Ariel Orda
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, Israel
{libman@tx, ariel@ee}.technion.ac.il

Abstract— **In noncooperative networks, resources are shared among selfish users, which optimize their individual performance measure. Traditional design methods tend to perform poorly in such networks, as they do not take into account the inherent noncooperative nature of the network users. Such networks require specialized design techniques in order to achieve efficient utilization of resources.**

We consider the generic and practically important class of atomic resource sharing networks, in which traffic bifurcation is not implemented, hence each user allocates its whole traffic to one of the network resources. We investigate topologies of parallel resources within a game-theoretic framework and establish the foundations of a design and management methodology that enables to operate such networks efficiently, despite both the lack of cooperation among users and the restrictions imposed by atomic resource sharing. We study various problems pertaining to capacity allocation, pricing, and admission control, and show that their solutions are substantially different from those corresponding to traditional networks.

Keywords— **Noncooperative networks, atomic (unsplittable) resource sharing, routing, capacity allocation, network management.**

I. Introduction

Traditional computer networks were designed and operated with systemwide optimization in mind. Accordingly, the actions of the network users were determined so as to optimize the overall network performance. Consequently, users would often find themselves sacrificing some of their own performance for the sake of the entire network. In recent years it has been recognized that systemwide optimization may be an impractical paradigm for the control of modern networking configurations (see, e.g., [1], [2], [3], [4], [5], and references therein). Indeed, control decisions in large-scale networks are often made by each user independently, according to its own individual performance objectives. Such networks are henceforth called *noncooperative*, and game theory [6] provides the systematic framework to study and understand their behavior. Under the noncooperative paradigm, the network is considered as a site of competition for the network resources among selfish users.

Game theoretic models have been employed in the context of flow control [1], [3], routing [2], [7], and virtual path bandwidth allocation [8]. These studies mainly investigate the structure of network operating points, i.e. the *Nash equilibria* of the respective games. Such equilibria are inherently inefficient [9] and, in general, exhibit suboptimal network performance. This deficiency can be overcome with the intervention of a *network agent*. This can be done both during the provisioning phase, i.e., when the network parameters are sized, and during the run-time phase, i.e., during the operation of the network. During the provisioning phase, the network *designer* should attempt to enforce a systemwide efficient equilibrium, by making intelligent topology design and capacity allocation decisions. During the run-time phase, the network *manager* should attempt to encourage "good behavior" among users by imposing additional rules on the game, e.g., through pricing schemes, service disciplines and admission control.

Until recently, network design and management issues were scarcely addressed under a noncooperative setting, mainly due to the complex structure — or lack thereof — of the underlying game. However, due to the emerging practical importance of the noncooperative paradigm, in recent years there is a growing effort to establish proper design and management methodologies for noncooperative networks (see, e.g., [4], [5], [10], [11], and references therein).

From the resource sharing perspective, these studies essentially assumed that *flow bifurcation* is permitted, i.e., each user may split its traffic among several network resources. However, it is often impractical to implement flow bifurcation [12], due to constraints imposed by the network or the switch architecture. For example, in ATM networks all the flow associated with a connection has to traverse through the same route. Even when not imposed by the architecture, lack of bifurcation is typical of various resource sharing scenarios. Common examples are a client connecting to only one mirror FTP site at a time, or a task that can be run in only one processor at a time in a processor sharing system. In such cases, we say that resource sharing is *atomic*, in the sense that a demand cannot be split among two or more resources.

Noncooperative networks with atomic resource sharing, henceforth referred to as *Atomic Noncooperative Networks (ANNs)*, were recently investigated in [14], where several fundamental properties of their Nash equilibria were shown. Based on these structural results, this paper establishes a framework for the proper design and management of ANNs, which enables their efficient use in spite of the inherent lack of cooperation. We demonstrate that atomic noncooperative networks call for special design and

This research was supported by the Ministry of Science in Israel.

management methods. For the typical design problem of capacity allocation, we show that traditional techniques, suitable for networks with flow bifurcation (with or without cooperation among users), are wholly inappropriate for ANNs, and then present a capacity allocation scheme that is shown to be optimal for an ANN framework. Next, we investigate management issues. Considering several system criteria, we show how to find desirable or optimal profiles of user demands. These findings lead to proper schemes for pricing and admission control in ANNs.

The text is organized as follows. Section II quotes from [14] some fundamentals required for the rest of the paper, including the definition of the model and an outline of several game-theoretic concepts. Section III demonstrates how traditional design rules can achieve poor performance in ANNs. The capacity allocation problem and management issues are discussed in sections IV and V, respectively. Conclusions and directions for future research are presented in section VI.

II. Model Formulation

We quote from [14] the definition of residual-capacity atomic noncooperative networks (ANNs). For the sake of concreteness, the analysis is presented in terms of a routing problem, as a specific case of a generic resource sharing problem. Nevertheless, as mentioned in the introduction, our study applies to a variety of atomic resource sharing problems in communication networks. Therefore, the following terms should be understood generically.

We consider a set $\mathcal{I} = \{1, \ldots, I\}$ of users, sharing a set of parallel links $\mathcal{L} = \{1, \ldots, L\}$ connecting a common source to a common destination. Each user $i \in \mathcal{I}$ has a flow rate (throughput) demand, which is an ergodic process with an average rate of $r^i > 0$. It is assumed, without loss of generality, that $r^1 \geq \ldots \geq r^i$. Each user i is allowed to select one link l^i to attend its flow demand. The selection of a user will also be referred to as its *strategy*. The set of users selecting link l is denoted by $U_l = \{i | l^i = l\}$. The *total flow* of link l is $f_l = \sum_{i \in U_l} r^i$. A *flow configuration* is the vector of strategies $\mathbf{f} = \langle l^1, \ldots, l^I \rangle$.

We assume that with each link $l \in \mathcal{L}$ is associated a parameter $C_l > 0$, called its *capacity*, and that, without loss of generality, $C_1 \geq \ldots \geq C_L$. The difference $C_l - f_l = C_l - \sum_{i \in U_l} r^i$ is termed the *residual capacity* of the link, and is denoted by RC_l. The cost that a user has to pay for routing its flow over link l is $\mathrm{J}(\mathrm{RC}_l)$, where $\mathrm{J}(\cdot)$ is a cost function that is identical in all links and decreasing in its argument. We do not make any assumptions about the cost function except that it is defined for any possible value of the residual capacity and that it is decreasing; in particular, there are no requirements about its continuity. For ease of presentation, we also define the individual link cost functions $\mathrm{J}_l(f_l) = \mathrm{J}(C_l - f_l)$; obviously, each link cost function is increasing in the link flow. Incidentally, 'cost', in this context, can mean anything that makes a user less happy when it is higher; some possible interpretations are, for example, the money paid for using the link, or the link delay.

In an *atomic noncooperative network (ANN)*, each user selfishly attempts to optimize its performance by selecting a single link in which its incurred cost is smallest, or, equivalently, where the residual capacity is largest. This setting gives rise to a *noncooperative game* [6]. A flow configuration $\mathbf{f}$ is said to be a *Nash Equilibrium Point (NEP)* if each user considers its chosen link to be the best under the given choices of other users. More formally, user i that currently utilises link l^i may decide to switch to link m^i if and only if

$$\mathrm{J}_{l^i}(f_{l^i}) > \mathrm{J}_{m^i}\left(f_{m^i} + r^i\right), \quad (1)$$

or, in terms of the link capacities,

$$C_{l^i} - f_{l^i} + r^i < C_{m^i} - f_{m^i},$$

i.e.

$$(C_{m^i} - f_{m^i}) - (C_{l^i} - f_{l^i}) > r^i. \quad (2)$$

When a user switches from its current link to the link with the highest residual capacity at the moment, we say that it makes a *best-reply move*. A flow configuration is an NEP if and only if, for no user i and the corresponding link l^i, there exists another link m^i which satisfies (1) or (2). Note that the exact nature of the cost function $\mathrm{J}(\cdot)$ is irrelevant.

A detailed analysis of the structural properties of ANNs, as well as a discussion of extensions of the model for general link cost functions (i.e. not necessarily defined by residual capacity) and for general topologies, can be found in [14]. For the rest of this paper, however, we need to quote only one fundamental theorem, essential for ANN design.

Theorem 1: *Any sequence of best-reply moves from any initial flow configuration always converges to an NEP.*
Corollary: *An ANN possesses at least one NEP.*

The proof of the theorem can be found in [14]. Also, note that the NEP, in general, is not unique. In section V-B we present necessary and sufficient conditions for an ANN to possess a unique NEP.

III. The need for a new design approach

As mentioned in the introduction, traditional design methods in computer networks assumed that users cooperate in order to reach network optimality. Problems such as capacity allocation were solved under the assumption that the routing of the users' traffic would be optimal [13]. In noncooperative networks, however, routing is performed at a Nash equilibrium point rather than an optimal point. Nash equilibria are inherently inefficient due to the noncooperative nature of the users, and, in our setting, also due to the further constraint imposed by the atomicity of resource sharing. Any network design methodology must take these facts into consideration. The importance of this proposition will become apparent by an example presented later in this section.

The formulation of design problems, and that of capacity allocation in particular, includes the definition of a criterion that assesses the quality of the solution, in terms of network performance at the operating point. The following are typical criteria for our setting:

- the average user cost, $\sum_i r^i J_{l^i}$;
- the worst user cost, $\max_i J_{l^i}$;
- the worst weighted user cost, $\max_i \left\{ r^i J_{l^i} \right\}$.

As we shall see, the worst user cost criterion is especially manageable in ANNs, where it can be handled in terms of maximizing the least residual capacity, regardless of the exact structure of the cost function. The other criteria do depend on the cost function.

Consider the following capacity allocation problem: given an ANN and an additional amount of capacity ΔC, the question is how to distribute ΔC among the links so that the resulting ANN exhibits the best possible performance. The capacity allocation problem has been analyzed in the past for bifurcated (i.e. non-atomic) routing. In [13] it is shown that, if the users' routing is optimal for the average cost, the optimal solution is to put the entire amount ΔC in the link with the largest capacity. The problem was treated in a noncooperative non-atomic setting in [10], where it was shown that adding any amount of capacity to the link with the originally largest capacity improves the performance of each individual user in the resulting (unique) NEP. The following simple example demonstrates how traditional design methods lead to network performance substantially worse than what is achieved when the noncooperative and atomic behavior of the users is accounted for. Moreover, careless addition of capacity (and, in particular, according to the above rule) may degrade the performance of the network.

Example: Let an ANN have M/M/1 links with $C_1 = 12$ (i.e. $J_1(f_1) = \frac{1}{12-f_1}$), $C_2 = 5$ (i.e. $J_2(f_2) = \frac{1}{5-f_2}$), $r^1 = 8$, $r^2 = 3$. It can be easily verified that there exists a unique NEP, namely $\langle 1, 2 \rangle$. The average cost is $\frac{8}{4} + \frac{3}{2} = 3.5$. Let us observe how the average cost is affected by adding more capacity to the network.

Let $\Delta C = 2$ be added to link 1 (the link with the originally largest capacity). In the resulting network, $C_1 = 14$ and $C_2 = 5$, and the only NEP is $\langle 1, 1 \rangle$, with an average cost of $\frac{8}{3} + \frac{3}{3} \approx 3.667$ — worse than the original average cost. On the other hand, investing ΔC in link 2 would result in a network with $C_1 = 12$, $C_2 = 7$, and the cost in the (again, unique) NEP $\langle 1, 2 \rangle$ would be $\frac{8}{4} + \frac{3}{4} = 2.75$.

However, this is still not the best solution. It turns out that the best average cost is achieved if $\Delta C_1 = \frac{44-16\sqrt{6}}{5} \approx 0.962$ is added to link 1, and $\Delta C_2 = \frac{16\sqrt{6}-34}{5} \approx 1.038$ is added to link 2. The only NEP in the resulting network is $\langle 1, 2 \rangle$, and the average cost then becomes $\frac{8}{4.962} + \frac{3}{3.038} \approx 2.6$.

Consider a similar network with the same original capacities, but with user flow demands of $r^1 = 6$, $r^2 = 3$. The only NEP is then $\langle 1, 1 \rangle$ with an average cost of $\frac{6}{3} + \frac{3}{3} = 3.0$. It is obvious that if $\Delta C = 1.001$, it is better to put it in the empty link 2, so that the NEP becomes $\langle 1, 2 \rangle$ with a cost of $\frac{6}{6} + \frac{3}{3.001} \approx 2.0$, rather than to link 1, for which the cost is $\frac{6}{4.001} + \frac{3}{4.001} \approx 2.25$. On the other hand, if $\Delta C = 0.999$, then $\langle 1, 1 \rangle$ continues to be the unique NEP no matter how ΔC is split, so it is best to add it entirely to link 1. ■

IV. Capacity allocation

The reason that atomic noncooperative networks present a particular design challenge is that they may inherently have multiple equilibria. Indeed, even the formulation of the classic capacity allocation problem — namely, how to allocate link capacities in order to achieve the 'best' operating point — is ill-defined in our case, because the operating point that the users achieve after the capacity has been allocated may not be unique. One possible way of tackling this difficulty is by refining further the meaning of the term *operating point* in a network that possesses multiple NEPs. Our approach in the sequel is to analyze the following capacity allocation problem.

The Online Capacity Allocation Problem

Input:

- existing link capacities, $C_1, \ldots, C_L$,
- user flow demands, $r^1, \ldots, r^I$,
- a flow configuration, $\mathbf{f} = \langle l^1, \ldots, l^I \rangle$, which is an NEP in the existing network,
- additional available capacity, ΔC.

Problem:

Find a vector of additional capacities $\Delta C_1, \ldots, \Delta C_{\hat{L}}$, subject to $\hat{L} \geq L$, $\Delta C_l \geq 0$ for every $1 \leq l \leq \hat{L}$, and $\sum_{l=1}^{\hat{L}} \Delta C_l \leq \Delta C$, for which the performance (according to a predefined criterion) of $\mathbf{f}_{opt}\left(C_1 + \Delta C_1, \ldots, C_L + \Delta C_L, \Delta C_{L+1}, \ldots, \Delta C_{\hat{L}}\right)$ is at least as good as the performance of $\mathbf{f}_{opt}\left(C_1 + \Delta C'_1, \ldots, C_L + \Delta C'_L, \Delta C'_{L+1}, \ldots, \Delta C'_{\hat{L}}\right)$ for any other vector of additional capacities satisfying the same conditions; here $\mathbf{f}_{opt}(\cdot)$ denotes the NEP with the highest performance among the NEPs reachable in the new ANN (with link capacities as specified by the vector $\langle\cdot\rangle$) from the initial flow configuration $\mathbf{f}$ by a sequence of best-reply moves.

In this formulation, the operating point is defined as the best among all the equilibria reachable from an initial flow configuration in the original network. The new capacity is installed during network operation (hence the term *online allocation*), and it is implicitly assumed that the network can steer the users to the best NEP by controlling the order of their moves (this assumption will be concretized and considerably alleviated in Theorem 2). An alternative (*offline*) approach could define the operating point as a certain distiguished NEP, e.g. the optimal NEP (i.e. the one that achieves the best value for the chosen criterion). Conceptually, if a solution to the offline problem is to be implemented for an existing, operational ANN, then it has to be 'shut down' before the new capacity is allocated to the links, and only then it can be 'turned on' once again to let the users reach the new NEP. Apparently, such a solution is infeasible, especially for large-scale networks that cannot be shut down completely during maintenance. Because of this observation, we focus on the online version of the problem. It will be seen that the online problem renders itself to an efficient solution, even though the problem of detecting the optimal NEP is generally NP-complete [14]. It should be emphasized that the distinction between the approaches is only relevant to networks with (potentially) multiple operating points; thus, it did not have to be introduced in the past for "socially"-controlled networks and for noncooperative, non-atomic networks, as they always possess a unique operating point.

A. Capacity allocation for the worst-cost criterion

In this section we present the solution of the online version of the capacity allocation problem for the worst-cost criterion. Because the link costs decrease in the residual capacities, the worst cost is attained in the link with the smallest residual capacity. Hence, an alternative formulation for this criterion is $\min_{l:U_l\neq\emptyset}\{C_l - f_l\} = \min_{l:U_l\neq\emptyset}\{\mathrm{RC}_l\}$, and it does not depend on the exact nature of the cost functions.

First, a few definitions to simplify the subsequent presentation. Denote by $R_1,\ldots,R_L$ the values of residual capacities in the initial flow configuration $\mathbf{f}$ in decreasing order, e.g. $R_1 = \max_l\{\mathrm{RC}_l\},\ldots,R_L = \min_l\{\mathrm{RC}_l\}$.* For the time being, we assume that in the initial NEP $\mathbf{f}$ all the links accomodate strictly positive flow (hence, $R_1 \leq C_L$); the reason that this assumption can be made without loss of generality will become evident after presenting the structure of the optimal capacity allocation. Next, we define the following sequence:

$$M_k = \begin{cases} \infty & k=0 \\ \sum_i r^i - \sum_{l=1}^{k}(C_l - C_{k+1}) & 1\leq k\leq L-1 \\ \sum_l (R_1 - R_l) = \sum_i r^i - \sum_{l=1}^{L}(C_l - R_1) & k=L \end{cases} \tag{3}$$

The sequence M_k is non-increasing: for $1 < k \leq L-1$, $M_{k-1} - M_k = k(C_k - C_{k+1}) \geq 0$, whereas $M_{L-1} - M_L = L(C_L - R_1) \geq 0$ follows from the assumption that $R_1 \leq C_L$.

It should be observed that, since the total residual capacity in the new ANN is constant, irrespective of how ΔC is split among the links, the optimal solution cannot be better than if that amount is distributed equally among the links in which there are users. In any other solution, the residual capacity will be smaller in at least one link, and the worst cost will be larger. More formally, for any subset of links $\overline{\mathcal{L}} \subseteq \mathcal{L}$, the smallest residual capacity in the new ANN provided that the users occupy only links from $\overline{\mathcal{L}}$ cannot be higher than $\frac{1}{|\overline{\mathcal{L}}|}\left(\sum_{l\in\overline{\mathcal{L}}} C_l + \Delta C - \sum_i r^i\right)$. Since, among different subsets with the same number of links, it is clearly best to take the ones with the largest capacities, we can conclude that the best capacity allocation cannot reach a worst cost that is better than the one corresponding to a residual capacity of

$$\max_{1\leq\overline{L}\leq L}\left\{\frac{1}{\overline{L}}\left(\sum_{l=1}^{\overline{L}} C_l + \Delta C - \sum_i r^i\right)\right\}, \tag{4}$$

*Note that the link ordering implied by $R_1,\ldots,R_L$ is unrelated to the original link indexes, as it is not necessarily the case that the larger residual capacities are found in the links with larger total capacities. See also the example at the end of this section.

where $\overline{L}$ is used to denote the size of the set $\overline{\mathcal{L}}$. Hence, if a capacity allocation scheme that reaches the value in (4) can be found, it is optimal. The next lemma matches the ranges of ΔC with the values of $\overline{L}$ for which the maximum in (4) is reached.

Lemma 1: *If $M_{k^*-1} > \Delta C \geq M_{k^*}$, where M_k are as defined in (3), then the maximum in (4) corresponds to $\overline{L} = k^*$, and the residual capacity in the first k^* links in the new ANN is, accordingly,*

$$\frac{1}{k^*}\left(\sum_{l=1}^{k^*} C_l + \Delta C - \sum_i r^i\right) =$$

$$\begin{cases} C_{k^*+1} + \frac{1}{k^*}(\Delta C - M_{k^*}) & k^* < L \\ R_1 + \frac{1}{L}(\Delta C - M_0) & k^* = L \end{cases}.$$

Proof: In this proof, so as not to drag the special definition of M_L, we define $C_{L+1} = R_1$; the definition of M_L is then identical to M_k for $1 \leq k \leq L-1$. In terms of M_k, the expression (4) can be written as

$$\max_{1\leq\overline{L}\leq L}\left\{\frac{1}{\overline{L}}\left(\sum_{l=1}^{\overline{L}} C_l + \Delta C - \sum_i r^i\right)\right\} =$$

$$\max_{1\leq\overline{L}\leq L}\left[C_{\overline{L}+1} + \frac{1}{\overline{L}}\left(\Delta C - M_{\overline{L}}\right)\right]. \tag{5}$$

We will make use of the following inequality, the correctness of which can be verified in a straightforward manner: for every $1 \leq k_1 \leq k_2 \leq L$,

$$k_1\left(C_{k_1+1} - C_{k_2+1}\right) \leq M_{k_1} - M_{k_2} \leq$$

$$k_2\left(C_{k_1+1} - C_{k_2+1}\right). \tag{6}$$

Assume that $M_{k^*-1} > \Delta C \geq M_{k^*}$. For any $\overline{L} \neq k^*$, define the difference

$$D = \left[C_{k^*+1} + \frac{1}{k^*}\left(\Delta C - M_{k^*}\right)\right] - \left[C_{\overline{L}+1} + \frac{1}{\overline{L}}\left(\Delta C - M_{\overline{L}}\right)\right] =$$

$$C_{k^*+1} - C_{\overline{L}+1} + \frac{1}{k^*\overline{L}}\left[\left(\overline{L} - k^*\right)\Delta C + k^* M_{\overline{L}} - \overline{L} M_{k^*}\right]$$

and consider two separate cases. If $\overline{L} < k^*$, then $\Delta C < M_{\overline{L}}$, and it follows due to the left-hand side of (6) that

$$D \geq C_{k^*+1} - C_{\overline{L}+1} + \frac{1}{k^*\overline{L}}\left(\overline{L}M_{\overline{L}} - \overline{L}M_{k^*}\right) =$$

$$\frac{1}{k^*}\left(M_{\overline{L}} - M_{k^*}\right) - \left(C_{\overline{L}+1} - C_{k^*+1}\right) \geq 0,$$

while if $\overline{L} > k^*$, then, since $\Delta C \geq M_{k^*}$,

$$D \geq C_{k^*+1} - C_{\overline{L}+1} + \frac{1}{k^*\overline{L}}\left(k^* M_{\overline{L}} - k^* M_{k^*}\right) =$$

$$C_{k^*+1} - C_{\overline{L}+1} - \frac{1}{\overline{L}}\left(M_{k^*} - M_{\overline{L}}\right) \geq 0,$$

this time due to the right-hand side of (6). In both cases, $D \geq 0$, which means that the maximum in (5) (or equivalently in (4)) is indeed reached for $\overline{L} = k^*$. ■

Lemma 1 provides an important clue for deriving the optimal allocation when $\Delta C \geq M_L$, as it distinguishes the links in which some capacity needs to be added from those that should be left intact. The actual values of the optimal allocation are given by the following theorem.

Theorem 2: *For $M_{k^*-1} > \Delta C \geq M_{k^*}$, the optimal distribution of $\{\Delta C_l\}_{1 \leq l \leq L}$ is as follows:*

- *If $k^* = L$,*

$$\Delta C_l = (R_1 - \mathrm{RC}_l) + \frac{1}{L}(\Delta C - M_L), \tag{7}$$

with the NEP remaining the same as in the original configuration;

- *If $k^* = 1$,*

$$\begin{array}{ll} \Delta C_l = 0, & 2 \leq l \leq L \\ \Delta C_1 = \Delta C & \end{array}, \tag{8}$$

with all the users moving to link 1 in the new NEP;

- *If $1 < k^* < L$,*

$$\begin{array}{cc} \Delta C_l = 0, & l > k^* \\ \Delta C_l = (C_{k^*+1} - \mathrm{RC}_l) + \frac{1}{k^*}(\Delta C - M_{k^*}), & 2 \leq l \leq k^* \\ \Delta C_1 = \Delta C - \sum\limits_{l=2}^{k^*} \Delta C_l & \end{array}, \tag{9}$$

*with all the users occupying only links $1, \ldots, k^*$ in the new NEP, provided that the users from links $k^*+1, \ldots, L$ are allowed to move first after the allocation.**

Proof: It is sufficient to show that after the capacity is allocated according to (7)-(9) and the users reach a new NEP in links $\{1, \ldots, k^*\}$, the residual capacity in those links is given by (4) for $\overline{L} = k^*$. Let us consider the cases separately.

1. $M_{L-1} > \Delta C \geq M_L$. Then the amount of capacity to be added to every link is given by (7), and the residual capacity in link l after the addition is

$$\mathrm{RC}_l + (R_1 - \mathrm{RC}_l) + \frac{1}{L}(\Delta C - M_L) =$$

$$\frac{1}{L}\Delta C + \frac{1}{L}\left(\sum_{l=1}^{L} C_l - \sum_i r^i\right),$$

i.e. as per (4) for $\overline{L} = L$. It is also a simple matter to verify that $\sum\limits_{l=1}^{L} \Delta C_l = \Delta C$. In addition, the flow configuration $\mathbf{f}$ remains an NEP in the new ANN, because there is no residual capacity advantage to any link over any other one, hence no user wants to make a move.

2. $\Delta C \geq M_1$. This is the simplest case; in fact, it is included in the next case, but does not need the proviso about the ability to control the users' moves order. Lemma 1 states that the largest residual capacity is reached for $\overline{L} = 1$, which means that the entire amount of ΔC should be allocated to link 1. It can be observed that in this case, all the users route to link 1 in a unique NEP, independently of the original flow configuration $\mathbf{f}$. Accordingly, allocating ΔC to link 1 entirely also solves the offline capacity allocation problem in this case. We note in passing that this solution coincides with that of the capacity allocation problem for non-atomic networks (see the references to [13] and [10] made in Section III).

3. $M_{k^*-1} > \Delta C \geq M_{k^*}$, where $1 < k^* < L$. After the new capacity is allocated as per (9), link 1 initially has the largest residual capacity, and all users at links $\{k^*+1, \ldots, L\}$ will request to move there, as will do some small users from the other links. Provided that the network can guarantee that all users from links $\{k^*+1, \ldots, L\}$ move first, a configuration is reached in which all users occupy links $1, \ldots, k^*$ only, and the residual capacity is balanced in them and given by (4) for $\overline{L} = k^*$. Indeed, the residual capacity in links $2, \ldots, k^*$ is

$$\mathrm{RC}_l + (C_{k^*+1} - \mathrm{RC}_l) + \frac{1}{k^*}(\Delta C - M_{k^*}) =$$

$$\frac{1}{k^*}\Delta C + \frac{1}{k^*}\left(\sum_{l=1}^{k^*} C_l - \sum_i r^i\right),$$

while in link 1, it is therefore

$$\mathrm{RC}_1 + \Delta C -$$

$$\left\{(k^*-1)\left[\frac{1}{k^*}\Delta C + \frac{1}{k^*}\left(\sum_{l=1}^{k^*} C_l - \sum_i r^i\right)\right] - \sum_{l=2}^{k^*} \mathrm{RC}_l\right\}$$

$$- \sum_{i \in \bigcup\limits_{l > k^*} U_l} r^i = \frac{1}{k^*}\Delta C + \frac{1}{k^*}\left(\sum_{l=1}^{k^*} C_l - \sum_i r^i\right).$$

■

The case $\Delta C < M_L$ is not handled by Lemma 1 and Theorem 2. With the additional capacity so small that it may not be enough to equalize the residual capacity, there does not appear to be an efficient way to calculate the optimal allocation; moreover, as we shall see in a later example, an optimal allocation may not even exist. A reasonable approach can be to try to improve the residual capacity in the worst links. To formalize this approach, define λ as the index that satisfies $\sum\limits_{l=\lambda}^{L}(R_\lambda - R_l) \leq \Delta C < \sum\limits_{l=\lambda-1}^{L}(R_{\lambda-1} - R_l)$; apparently, $\lambda \geq 2$. Then, if $\mathrm{ind}(l)$ denotes the link index in which the residual capacity is R_l, the capacity distribution among the links is

$$\Delta C_{\mathrm{ind}(l)} =$$

$$\begin{cases} (R_\lambda - R_l) + \frac{1}{L-\lambda+1}\left(\Delta C - \sum\limits_{j=\lambda}^{L}(R_\lambda - R_j)\right) & l \geq \lambda \\ 0 & l < \lambda \end{cases}.$$

*Note that this is the only case where the assumption about the ability of the network to control the order of the users' moves is required.

If this approach is used, the users do not deviate from the original flow configuration $\mathbf{f}$ in the resulting ANN, because $\mathbf{f}$ is an NEP in the original ANN and the residual capacity advantage in any pair of links does not grow. Indeed, it can be proved that the above method is optimal among all allocations that do not cause the users to make moves. However, it cannot be guaranteed that a different allocation does not 'strike gold' by causing the ANN to settle to a flow configuration in which the residual capacities in all links are balanced; moreover, it seems that there is no efficient way to detect such a possibility.

The following example illustrates the various cases discussed above.

Example: Consider an ANN of three links, $C_1 = 40$, $C_2 = 30$, $C_3 = 21$, and four users, $r^1 = 28$, $r^2 = 12$, $r^3 = 10$, $r^4 = 5$. The initial flow configuration is $\mathbf{f} = \langle 1, 2, 2, 3\rangle$, which can be verified to be an NEP. The residual capacities are $\mathrm{RC}_1 = C_1 - r^1 = 12$, $\mathrm{RC}_2 = C_2 - \left(r^2 + r^3\right) = 8$, $\mathrm{RC}_3 = C_3 - r^4 = 16$. In decreasing order, the values become $R_1 = 16$, $R_2 = 12$, $R_3 = 8$; the indexes of the residual capacities in decreasing order are seen to be unrelated to the original link indexes.

The M-sequence, calculated by (3), is: $M_1 = 45$, $M_2 = 27$, $M_3 = 12$. Accordingly, if $\Delta C \geq 45$, it is best to allocate all of it to link 1. In that case, the capacity of link 1 is so large that all the users move there eventually, to reach a new NEP of $\langle 1, 1, 1, 1\rangle$. If $27 > \Delta C \geq 12$, then the optimal allocation is $\Delta C_1 = 4 + \frac{\Delta C - 12}{3} = \frac{\Delta C}{3}$, $\Delta C_2 = 8 + \frac{\Delta C - 12}{3} = \frac{\Delta C}{3} + 4$, $\Delta C_3 = \frac{\Delta C - 12}{3} = \frac{\Delta C}{3} - 4$. This allocation causes the residual capacities in all links to become equal, while the users stay in the original configuration $\mathbf{f}$.

In the range $45 > \Delta C \geq 27$, the optimal allocation is $\Delta C_1 = 1 + \frac{\Delta C - 1}{2}$, $\Delta C_2 = \frac{\Delta C - 1}{2}$, $\Delta C_3 = 0$. After the allocation, user 4 no longer stays in link 3, because the residual capacity in link 1 is $13 + \frac{\Delta C - 1}{2} \geq 26 > C_3 = 21$; it therefore moves to link 1, bringing the ANN to its new NEP $\langle 1, 2, 2, 1\rangle$, in which the residual capacities of links 1 and 2 are equal.

Finally, for $\Delta C < 12$, the general method provides no answer about the optimal allocation. Indeed, consider the following two cases. For $\Delta C = 1$, there is no way that the additional capacity can be allocated to make the users make any moves; accordingly, the best improvement in worst cost is reached if the entire amount is simply allocated to the link with the smallest residual capacity in $\mathbf{f}$, which is link 2. With $\Delta C = 3$, however, one can do better. Suppose the allocated capacity is distributed as follows: $\Delta C_1 = 1 - \epsilon$, $\Delta C_2 = 0$, $\Delta C_3 = 2 + \epsilon$, where $\epsilon \ll 1$ is an infinitesimal quantity, just enough to make sure that user 3 makes a switch after the allocation. Then, after the allocation, the residual capacity advantage of link 3 over link 2 is $16 + (2 + \epsilon) - 8 = 10 + \epsilon$, barely enough to cause $r^3 = 10$ to make the switch to link 3. In turn, this causes user 4 to move to link 2, eventually reaching the new NEP $\langle 1, 2, 3, 2\rangle$, in which the residual capacities are $\mathrm{RC}_1 = 13 - \epsilon$, $\mathrm{RC}_2 = 30 - (12 + 5) = 13$, $\mathrm{RC}_3 = 21 + (2 + \epsilon) - 10 = 13 + \epsilon$, i.e. practically balanced. Note that, in this case, there is no optimal capacity allocation; as $\epsilon > 0$ becomes smaller, the allocation achieves a better performance. The above phenomenon happens for all $12 > \Delta C \geq 3$; just add an additional $\frac{\Delta C - 3}{3}$ to every link, keeping the eventual residual capacities 'almost' balanced. ■

Recapping: whenever $\Delta C \geq M_L$, the optimal capacity allocation can be stated explicitly. It should be noted that the optimal allocation does not involve creating new links, but merely enhancing existing links. Also, since the optimal allocation depends only on the differences of the link capacities rather than the capacities themselves in the original ANN, it is invariant to an addition or subtraction of a common amount of capacity in all the links. For that reason, if there are links that are unoccupied in the initial flow configuration $\mathbf{f}$, the optimal capacity allocation is the same as in an equivalent ANN, in which all the link capacities are demoted accordingly to make the unused links irrelevant (i.e. with zero capacity). This explains why our assumption that there is at least one user in every link in the initial configuration $\mathbf{f}$ does not cause a loss in generality.

As a final comment, note that the optimal solution of the capacity allocation problem is not additive; in other words, the optimal allocation for $(\Delta C_I + \Delta C_{II})$ is not obtained by implementing the optimal allocation for ΔC_I and subsequently for ΔC_{II}. In practical terms, this means that the best policy for the allocation of an amount of capacity may be to hold it up completely and wait for the next time that yet another quota of capacity is available. Then, implementing the optimal allocation for the total amount at once results in a network with a better performance than if the capacities had been dealt with separately.

This section presented a solution to the problem of optimal capacity allocation under the criterion of worst user cost. The analysis was manageable mostly due to the fact that maximizing the worst user cost is equivalent to maximizing the minimum residual capacity, and hence does not depend on the exact nature of the cost function. Tackling the capacity allocation problem for other criteria, such as average user cost, is substantially more difficult because the specific properties of the cost function must then be considered.

V. Management issues

Network management policies attempt to improve the network performance during the run-time phase of the network. In the context of noncooperative networks, a major challenge of network management is to encourage 'good behavior' among users, in order to achieve a reasonably good performance in spite of the inherent lack of cooperation.

Once the resources are designed and provisioned, the user cost functions, and consequently their performance, depend solely on the user demands. In other words, network performance is a function of the profile of demands. Accordingly, network management should encourage demand profiles that result in a desirable performance. One way to end up with the desired profile is to employ a pricing scheme that provides incentives to users whose demands fit well in the desired profile. Clearly, a prerequisite for providing such incentives is to identify the desirable profiles of user demands, which we investigate in section V-A.

Apart from efficient performance at equilibrium, network management may attempt to achieve other desirable properties. For example, possessing a unique NEP is usually considered a desirable property. In fact, noncooperative networks which do not possess unique operating points are generally more difficult to design and manage, since a unique and predictable operating point is a prerequisite for many design methods. Moreover, a unique NEP is independent of the initial configuration and hence stable even if the flow rates of the users experience short-term changes. Accordingly, in section V-B we demonstrate how the NEP uniqueness can be guaranteed by the employment of an admission control policy.

A. Optimal profiles of user demands

Given a network that is designed to support some (expected) total flow demand R, our aim is to identify the best way to partition R among various users so that the resulting NEP will be the best possible. The problem is formally stated as follows.

Optimal NEP flow distribution

Input:
- a set of links with capacities $C_1, \ldots, C_L$,
- a total flow demand of R.

Problem:

Find I and a vector $\langle r^1, \ldots, r^I \rangle$, such that $\sum_{i=1}^{I} r^i = R$ and the performance (according to a predefined criterion) of $\mathbf{f}_{opt}\left(\langle r^1, \ldots, r^I \rangle\right)$ is at least as good as the performance of $\mathbf{f}_{opt}\left(\left\langle r'^1, \ldots, r'^{I'} \right\rangle\right)$ for any I' and $\left\langle r'^1, \ldots, r'^{I'} \right\rangle$ satisfying $\sum_{i=1}^{I'} r'^i = R$, where $\mathbf{f}_{opt}\left(\langle \cdot \rangle\right)$ denotes the optimal NEP, i.e. the one with the highest performance among all the NEPs that exist in the ANN with user demands as specified by the vector $\langle \cdot \rangle$.

In the context of a bifurcating (i.e. non-atomic) framework, this question has a trivial answer, namely have just one user with a flow demand of R; indeed, that user could achieve the ideal network optimum by splitting its flow among the various links according to the optimal routing strategy. In the context of atomic routing, however, this would typically be a bad idea, since a single user must assign the whole of R to a single link.

The performance criterion plays an important part in this problem. If the performance is measured by the worst cost that a user attains, then the solution is rather trivial. Indeed, for a given amount of total residual capacity, the least worst cost is attained if the residual capacities are balanced in all the links in which there are users. But a configuration with balanced residual capacity is automatically reached by an infinite number of infinitesimal users, which sum up to a total flow of R. Indeed, when such users reach an NEP, the residual capacity advantage of any link may not be higher than a single user's flow demand (otherwise that user would want to switch), which means that it must be zero. Therefore, a management policy that attempts to minimize the worst cost will simply prefer small users over large ones.

The problem becomes complex when average cost is used to measure the network performance. Of course, the average cost depends on the precise nature of the link cost functions. In this section, following [4], the problem is tackled for M/M/1 link cost functions, defined as $\mathrm{J}\left(\mathrm{RC}_l\right) = \frac{1}{\mathrm{RC}_l}$; the average cost is then defined to be $\frac{1}{R}\sum_i \frac{r^i}{\mathrm{RC}_{l^i}} = \frac{1}{R}\sum_l \frac{f_l}{C_l - f_l}$.

The analysis performed in the sequel will show that the best NEP is characterized by the total link flows only, and not by the actual user flow demands. The network manager may decide how to translate the optimal NEP link flows to an actual flow demand profile. One possibility is to define user flows equal to the optimal link flows; if this definition is employed, the optimal NEP is also the *natural NEP* of the users (see definition in section V-B.1 below), and the network should be able to force the users reach it through an appropriate admission policy.

We proceed to find the (atomic) NEP with the least average cost among all NEPs having a total flow of R. First, we quote the structure of the 'ideal' (cooperative and non-atomic) optimum [10]:

$$f_l^* = \begin{cases} C_l - \left(\sum_{n=1}^{L^*} C_n - R\right) \frac{\sqrt{C_l}}{\sum_{n=1}^{L^*} \sqrt{C_n}}, & 1 \le l \le L^* \\ 0, & l > L^* \end{cases} \quad (10)$$

where L^* is defined through a monotonous non-decreasing sequence,

$$G_l = \sum_{n=1}^{l-1} C_n - \sqrt{C_l} \sum_{n=1}^{l-1} \sqrt{C_n}, \quad 2 \le l \le L$$
$$G_1 = 0, \quad G_{L+1} = \sum_{n=1}^{L} C_n = C,$$

as the index satisfying

$$G_{L^*} < R \le G_{L^*+1}.$$

Denote the flows of the lowest-cost NEP by $\hat{f}_l$, and denote by $\hat{L}$ the index of the last link with strictly positive flow in that NEP. Since a strictly positive flow $\hat{f}_{\hat{L}}$ stays in link $\hat{L}$, the residual capacity in no other link is greater than $C_{\hat{L}}$, so $\hat{f}_l \ge C_l - C_{\hat{L}}$ for any $l < \hat{L}$. It can be noted that $\hat{L} \le L^*$; this is because in the ideal optimum, the residual capacity in all the links with positive flow is strictly higher than C_{L^*+1}, because $\frac{C_l}{\left(C_l - f_l^*\right)^2} \le \frac{1}{C_{L^*+1}}$ for every $l \le L^*$, while $C_l > C_{L^*+1}$. Therefore, it is impossible that $\hat{L} > L^*$ without the total flow exceeding R.

Consequently, the NEP with the lowest average cost is, for some $\hat{L}$, an optimal distribution of the flow subject to the constraints that the residual capacity in any link is not higher than $C_{\hat{L}}$ and there is no flow beyond link $\hat{L}$. The next proposition specifies this distribution.

Proposition 1: *Given a link index $\hat{L}$ and the constraints*

$$\begin{array}{ll} f_l \ge C_l - C_{\hat{L}}, & 1 \le l \le \hat{L} \\ f_l = 0, & l > \hat{L} \end{array} \quad (11)$$

if the constraints are feasible (that is, they can be satisfied by a total flow of R), then the optimal distribution of the flow that satisfies the constraints is

$$\hat{f}_l = \begin{cases} C_l - C_{\hat{L}}, & 1 \le l \le L' \\ C_l - \left(\sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R\right) \frac{\sqrt{C_l}}{\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}}, & L' < l \le \hat{L} \\ 0, & l > \hat{L} \end{cases} \quad (12)$$

where L' is defined through the monotonously decreasing sequence

$$\Delta_i = C_{\hat{L}} \sum_{n=i+1}^{\hat{L}} \left(1 - \frac{\sqrt{C_n}}{\sqrt{C_i}}\right), \quad 1 \leq i \leq \hat{L} - 1$$
$$\Delta_0 = \hat{L} C_{\hat{L}}, \quad \Delta_{\hat{L}} = 0$$

as the index which satisfies

$$\Delta_{L'} \geq R - \sum_{n=1}^{\hat{L}} \left(C_n - C_{\hat{L}}\right) > \Delta_{L'+1}. \tag{13}$$

<u>Proof</u>: It is easily seen that Δ_i is indeed a decreasing sequence; as i becomes larger, less elements are summed and each one is smaller. It is also easily verified that $\sum_l \hat{f}_l = R$.

The first step is to establish that $\hat{f}_l$ satisfy the constraints. This is easily seen for $1 \leq l \leq L'$ and for $l > \hat{L}$. As for $L' < l \leq \hat{L}$, note that the rightmost inequality in (13) can be weakened to

$$R - \sum_{n=1}^{\hat{L}} \left(C_n - C_{\hat{L}}\right) > \Delta_{L'+1} =$$
$$C_{\hat{L}} \sum_{n=L'+1}^{\hat{L}} \left(1 - \frac{\sqrt{C_n}}{\sqrt{C_{L'+1}}}\right) \geq$$
$$C_{\hat{L}} \sum_{n=L'+1}^{\hat{L}} \left(1 - \frac{\sqrt{C_n}}{\sqrt{C_l}}\right),$$

which implies, by a different grouping of some terms,

$$C_{\hat{L}} \frac{\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}}{\sqrt{C_l}} > \sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R,$$

therefore

$$C_{\hat{L}} > \left(\sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R\right) \frac{\sqrt{C_l}}{\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}} = C_l - \hat{f}_l,$$

which means that the constraint is satisfied.

To prove that $\hat{f}_l$ are optimum flows, it suffices to show that they satisfy the Kuhn-Tucker conditions, which, for this case, state that there exists a positive number λ such that

$$\begin{array}{ll} \frac{C_l}{\left(C_l - \hat{f}_l\right)^2} = \lambda, & \hat{f}_l > C_l - C_{\hat{L}} \text{ (i.e. } L' < l \leq \hat{L}) \\ \frac{C_l}{\left(C_l - \hat{f}_l\right)^2} = \frac{C_l}{C_{\hat{L}}^2} \geq \lambda, & \hat{f}_l = C_l - C_{\hat{L}} \text{ (i.e. } 1 \leq l \leq L') \end{array} \tag{14}$$

The first condition of (14) is easily seen to be satisfied, because

$$\frac{C_l}{\left(C_l - \hat{f}_l\right)^2} = \frac{\left(\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}\right)^2}{\left(\sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R\right)^2} = \lambda$$

is independent of l. To see that the second condition of (14) is also satisfied, note that the leftmost inequality in (13) can be weakened to

$$C_{\hat{L}} \sum_{n=L'+1}^{\hat{L}} \left(1 - \frac{\sqrt{C_n}}{\sqrt{C_l}}\right) \geq$$
$$\Delta_{L'} = C_{\hat{L}} \sum_{n=L'+1}^{\hat{L}} \left(1 - \frac{\sqrt{C_n}}{\sqrt{C_{L'}}}\right) \geq$$
$$R - \sum_{n=1}^{\hat{L}} \left(C_n - C_{\hat{L}}\right).$$

Repositioning a few terms, this implies

$$C_{\hat{L}} \frac{\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}}{\sqrt{C_l}} \leq \sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R,$$

and now taking the squares of both sides and inverting yields

$$\frac{C_l}{C_{\hat{L}}^2} \geq \frac{\left(\sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}\right)^2}{\left(\sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R\right)^2} = \lambda,$$

which means that the second condition is satisfied. ■

The average cost of the above flow distribution, for a given $\hat{L}$, is

$$\mathrm{J} = \frac{1}{R} \sum_l \frac{f_l}{C_l - f_l} = \frac{1}{R} \sum_{l=1}^{L'} \left(\frac{C_l}{C_{\hat{L}}} - 1\right) +$$
$$\frac{1}{R} \sum_{l=L'+1}^{\hat{L}} \left(\frac{\sqrt{C_l} \sum_{n=L'+1}^{\hat{L}} \sqrt{C_n}}{\sum_{n=1}^{\hat{L}} C_n - L' C_{\hat{L}} - R} - 1\right).$$

This expression does not seem to exhibit any sort of monotonous dependency on $\hat{L}$. Hence, in order to find the NEP with the least average cost, one has to enumerate on all possible values of $\hat{L}$, verify for each one that the constraints are feasible, calculate the optimal flow distribution and its average cost, and finally select the value of $\hat{L}$ that achieves the lowest average cost. It can be observed (informally) that, as a rule, the higher the differences among link capacities, the greater is the distance between the lowest-cost NEP and the ideal optimum, because the constraints (11) affect a larger portion of the flow.

B. Maintaining a unique NEP through admission control

This section first quotes from [14] the conditions for the uniqueness of the NEP in an ANN (for completeness), and then discusses an admission control policy that can be employed for maintaining a unique NEP.

B.1 NEP uniqueness conditions

We begin by observing that condition (2) (in section II) implies that if a user considers its link to be the best at any moment, so do all users with a larger flow demand in the same link. Consequently, an NEP can always be found in an ANN in the following manner: order the users by decreasing flow demands, and put each one in turn in the link with the largest residual capacity so far. Applying the above observation after each user in turn shows that the resulting configuration is an NEP. This NEP is henceforth termed the *natural NEP* of the ANN.

For a network to have a unique NEP, then, it is required that no other NEP but the natural one exists. This requirement is tested by the following algorithm, whose validity is proved in [14].

Algorithm NEP UNIQUENESS (NU):

(1) If $\sum_{i=2}^{I} r^i \geq C_1 - C_2$, go to step 6.
(2) Assign $C_1 \leftarrow C_1 - r^1$, and reorder the indexes to maintain the capacities in decreasing order.
(3) Assign $r^1 \leftarrow r^2, \ldots, r^{I-1} \leftarrow r^I$ and $I \leftarrow I - 1$.
(4) If $I > 1$, return to step 1.
(5) Stop: there is a unique NEP in the network.
(6) Stop: the NEP in the network is not unique.

B.2 The admission control policy

The target of guaranteeing the uniqueness of the operating point at all times can be achieved by the employment of the following admission control policy.

- When a new user arrives, run algorithm **NU** and admit the user only if the network continues to have a unique NEP with it.
- When a user wishes to leave, it is allowed to do so immediately.

Note that this (rather straightforward) policy does not provide incentives to users with desired flow demands; it simply rejects those that do not conform to a strict rule. For the method to work, it must be established that the NEP uniqueness property is not harmed when a user leaves. The next proposition states that this is always the case.

Proposition 2: *If there is a unique NEP in an ANN $(\mathcal{L},\mathcal{I})$, then there is a unique NEP in all ANNs with the same set of links $\mathcal{L}$ and a set of users that is a subset of $\mathcal{I}$.*

Proof: Evidently, it is sufficient to prove that the NEP uniqueness property remains if a single user is removed from the ANN. The proposition then follows by induction.

Consider an ANN for which algorithm **NU** reports a unique NEP. The algorithm iterates $I-1$ times on steps (1)-(4), and at all iterations the inequality tested in step (1), namely

$$\sum_{i=2}^{I} r^i \geq C_1 - C_2, \tag{15}$$

is false. Suppose that the algorithm is run on the same ANN, except that user j has been removed. Denote the new flow demand values by ρ^i; thus, $\rho^i = \begin{cases} r^i & i < j \\ r^{i+1} & i \geq j \end{cases}$. In the first $j-1$ loops of the new run, the inequality (15) is false again, because at each iteration the value of the sum is smaller (by precisely r^j) than the sum at the respective step in the original ANN. Now focus on iteration number j: in the original run, it is known that at this point $\sum_{i=2}^{I} r^i < C_1 - C_2$, but in terms of the new ANN, because user j (now r^1) is not there any more, $\sum_{i=1}^{I} \rho^i < C_1 - C_2$. This inequality continues to be true in the subsequent loops as well, because each time the value of ρ^1 is subtracted from both sides. Hence, $\sum_{i=2}^{I} \rho^i < C_1 - C_2$ holds from iteration j and on, and the algorithm reports a unique NEP for the new ANN as well. ■

VI. Conclusion

In this paper we provided a game-theoretic framework for the design and management of atomic noncooperative networks (ANNs). Other models of resource sharing in noncooperative networks have been studied in the past, but they dealt exclusively with *bifurcated sharing*. The ANN model is important in that it deals with resource sharing networks where bifurcation of the demand is impossible or impractical. This is indeed the case in many practical scenarios, such as utilisation of virtual circuits in packet-switched networks, mirror database sites and World-Wide Web servers, and also beyond the context of communication networking, e.g. in multiprocessor systems.

We have shown that the distinct properties of ANNs require special design and management techniques, different from the ones for networks with systemwide optimization or with bifurcated sharing. In particular, the optimal solutions of such problems as capacity allocation and admission control have different structures than the traditional ones, and several results in that respect have been established in this paper.

For the provisioning phase, we have made the distinction between offline and online capacity allocation. While noting that, generally, optimal offline capacity allocation is an NP-complete problem, we presented a solution of the online capacity allocation problem for the criterion of worst cost.

For the run-time phase, we have shown how to find the optimal profile of user demands for the purpose of devising a pricing scheme, and, in particular, presented a complete analytical solution for the case of M/M/1 link cost functions. An argument has been made in favor of networks with a unique equilibrium. Accordingly, we presented an efficient algorithm for detecting the uniqueness conditions, which can be employed in an admission control policy that puts as its goal to maintain a unique equlibrium point at all times.

The issue of design and management of ANNs provides a fertile ground for further research. One important direction is to establish a proper design methodology for more general cost functions and network topologies. A first step has been taken in [14], where the structural properties of the equilibrium have been established also for such generalizations of the basic model. Another interesting direction is to explore the various design and management issues under additional system criteria. Finally, concretizing these findings into actual management schemes, and in particular pricing schemes, is a challenging issue for future research.

The establishment of a design and management methodology, which fully addresses the various problems and challenges posed by noncooperative networks with atomic resource sharing, remains a hard problem, with many facets still to be addressed. Yet, we believe that our study provides some important guidelines and tools for reaching the solution.

References

[1] Z. Zhang and C. Douligeris. Convergence of Synchronous and Asynchronous Greedy Algorithms in a Multiclass Telecommunications Environment. *IEEE Transactions on Computers*, 40(8): 1277-1281, August 1992.

[2] A. Orda, R. Rom, and N. Shimkin. Competitive Routing in Multiuser Communication Networks. *IEEE/ACM Transactions on Networking*, 1(5): 510-521, October 1993.

[3] Y. A. Korilis and A. A. Lazar. On the Existence of Equilibria in Noncooperative Optimal Flow Control. *Journal of the ACM*, 42(3): 584-613, May 1995.

[4] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving Network Optima Using Stackelberg Routing Strategies. *IEEE/ACM Transactions on Networking*, 5(1): 161-173, February 1997.

[5] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang. Pricing in Computer Networks: Motivation, Formulation and Example. *IEEE/ACM Transactions on Networking*, 1(6): 614-627, December 1993.

[6] R. B. Myerson. *Game Theory: Analysis of Conflict.* Harvard University Press, Cambridge, Massachusetts, 1991.

[7] A. A. Economides and J. A. Silvester. Multi-Objective Routing in Integrated Services Networks: A Game Theory Approach. *Proceedings of INFOCOM'91*, pp. 1220-1225, Bal Harbor, FL, April 1991.

[8] A. A. Lazar, A. Orda, and D. E. Pendarakis. Virtual Path Bandwidth Allocation in Multi-User Networks. *IEEE/ACM Transactions on Networking*, 5(6): 861-871, December 1997.

[9] P. Dubey. Inefficiency of Nash Equilibria. *Mathematics of Operations Research*, 11(1): 1-8, February 1986.

[10] Y. A. Korilis, A. A. Lazar and A. Orda. Capacity Allocation under Noncooperative Routing. *IEEE Transactions on Automatic Control*, 42(3): 309-325, March 1997.

[11] J. K. MacKie-Mason and H. R. Varian. Pricing Congestible Network Resources. *IEEE Journal on Selected Areas in Communications*, 13(7): 1141-1149, September 1995.

[12] A. Kershenbaum. *Telecommunications Network Design Algorithms.* McGraw-Hill, 1993.

[13] D. Bertsekas and R. Gallager. *Data Networks.* Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1992.

[14] L. Libman and A. Orda. Atomic Resource Sharing in Noncooperative Networks. To appear in *Telecommunication Systems (Special Issue on Network Economics).* Conference version: *Proceedings of INFOCOM'97*, Kobe, Japan, April 1997.

Lavy Libman received his B.Sc. degrees (summa cum laude) in Electrical Engineering and in Computer Engineering and his M.Sc. degree in Electrical Engineering from the Technion - Israel Institute of Technology, Haifa, Israel, in 1992 and 1997, respectively. He is currently a Ph.D. candidate in the Department of Electrical Engineering, Technion - Israel Institute of Technology. In 1992–1999 he served as an engineer in the Israel Defense Forces (with a research rank since 1996).

His current research interests are in the application of game theory, pricing, and economic tools to computer networks, in particular to routing, flow control, and provision of QoS.

Mr. Libman received the Merlin Award for distinguished students in Electrical Engineering and the Knesset (Israeli Parliament) Award for distinguished students.

His e-mail address is: libman@tx.technion.ac.il

Ariel Orda (S'94-M'92-SM'97) received the B.Sc. (summa cum laude), M.Sc., and D.Sc. degrees in Electrical Engineering from the Technion - Israel Institute of Technology, Haifa, Israel, in 1983, 1985, and 1991, respectively. Since 1994, he has been with the Department of Electrical Engineering at the Technion, where he is currently a a Senior Lecturer and the Academic Head of the Computer Networking Laboratory. In 1993–1994, he was with the Center for Telecommunication Research, Columbia University, New York, NY, as a Visiting Scientist, and with AT&T Bell Laboratories, Murray Hill, NJ, as a Scientific Consultant. During the summers of 1992 and 1995–1999, he held visiting positions at AT&T Bell Laboratories, Murray Hill, NJ, IBM Watson, Hawthorne, NY, and Bell Laboratories — Lucent Technologies, Holmdel, NJ. Since 1991 he has held several consulting positions with the Israeli industry.

His current research interests include QoS routing, the application of game theory to computer networking, network pricing, and distributed network algorithms.

Dr. Orda received the Award of the Chief Scientist in the Ministry of Communication in Israel, a Gutwirth Award for outstanding distinction, and the Research Award of the Association of Computer and Electronic Industries in Israel.

His e-mail address is: ariel@ee.technion.ac.il