Fairness in Cellular Mobile Networks

Saswati Sarkar, *Member, IEEE,* and
Kumar N. Sivarajan, *Member, IEEE*

***Abstract*—Channel allocation algorithms for channelized cellular systems are discussed from a new perspective, *viz.*, fairness of allocation. The concepts of relative and absolute fairness are introduced and discussed. It will be shown that under certain reasonable assumptions, there exists an absolute (max-min) fair carried traffic intensity vector (a vector describing the traffic carried in the cells of the system). We also show that this vector is unique. We describe some properties of the max-min fair carried traffic intensity vector in an asymptotic limit where the traffic and the number of channels are scaled together. For each traffic pattern, we determine a fixed channel allocation which attains this max-min fair carried traffic intensity vector independent of the value of the offered traffic, in the same asymptotic limit. Finally, we discuss a tradeoff between being max-min fair and trying to maximize revenue. We conclude this correspondence by discussing some possible extensions of our work.**

***Index Terms*—Asymptotic analysis, cellular networks, channel assignment algorithms, fairness, max-min fairness, revenue maximization.**

I. Introduction and Summary

In a cellular system, the coverage area is logically divided into cells. Each cell has a cell site or a base station. The communication from the mobile user is directed to a central switching office by the base station. The central switching office directs this communication to the destination. Depending on the mode of multiple access used by the mobile customers, cellular systems can be broadly classified into *channelized* and *nonchannelized* systems. In a channelized cellular system, the multiple access is time-division multiple access (TDMA) or frequency-division multiple access (FDMA), or a combination of both. The term *channel* refers to a time slot in TDMA, a frequency slot in FDMA, and a combination of both in TDMA/FDMA systems such as the Global System for Mobile Communications (GSM). Calls arrive and depart at random times in the cells of the system and a *channel assignment algorithm* must assign a channel to each call for its duration, while obeying certain channel reuse constraints. These reuse constraints can be modeled by a hypergraph [2], as explained in [7]. This correspondence will deal entirely with channelized systems. The cellular network operator whose task it is to choose an appropriate channel assignment algorithm is usually interested in maximizing his revenue and this amounts to maximization of the total traffic carried in the system, if we assume calls in all cells are charged at the same rate (dollars per unit time). In addition, in order to keep his customers satisfied, the operator has to provide them a minimum grade of service (maximum blocking probability). However, many channel assignment algorithms provide unequal grades of service in the various cells of the system, and are thus "unfair." This correspondence attempts to study the concept of fairness in cellular networks. First, we shall motivate the study of fairness with an example.

Manuscript received February 5, 1997; revised January 18, 2000. The work of S. Sarkar was supported in part by the National Science Foundation under Grant ANI01-06984. A portion of this work was performed at the Indian Institute of Science, Bangalore, supported under a grant from Nortel Networks. The material in this correspondence was presented in part at the 34th Annual Allerton Conference on Communications, Control, and Computing, Allerton, IL, September 1996, and published in the *Proceedings* of that conference.

S. Sarkar is with the Department of Electrical Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: swati@ee.upenn.edu).

K. N. Sivarajan is with Tejas Networks, Bangalore, 560 012, India (e-mail: kumar@tejasnetworks.com).

Communicated by R. Cruz, Associate Editor for Communication Networks.
Publisher Item Identifier 10.1109/TIT.2002.800495.

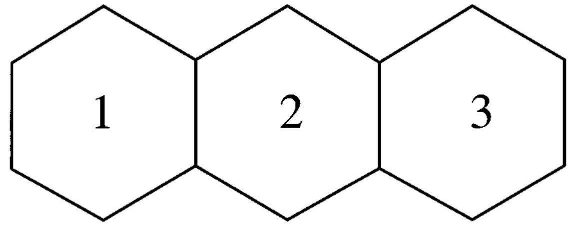

Fig. 1. Three-cell system.

Example: Consider the three-cell system shown in Fig. 1 and assume that a total of 80 channels are available. Adjacent cells are forbidden from using the same channel, but cells 1 and 3 can reuse the same channel. Consider two fixed channel algorithms, FCAA1 and FCAA2. FCAA1 allocates 40 channels to each cell and FCAA2 allocates 42 channels to cells 1 and 3 and 38 channels to cell 2. Let the offered traffic in each cell be 35 Erlangs. Let us assume that there are no handovers and no intercell calls. We also assume that neither algorithm allows any queuing of call requests. If there is a free channel in the cell in which a call is requested, the call request is honored; otherwise it is blocked. FCAA1 achieves a blocking probability of 5.4% in each cell and the total carried traffic is 99.3 Erlangs. For FCAA2, the blocking probability is 3.6% in cells 1 and 3 and 7.8% in cell 2. The total carried traffic is 99.8 Erlangs, slightly better than that of FCAA1. FCAA1 can be said to be "fairer" than FCAA2 since it treats the individual cells more "equally." It is reasonable to assume that in this case a network operator will prefer FCAA1 to FCAA2 even though his total carried traffic is slightly reduced. Otherwise, he will risk the desertion of customers in cell 2 for a competitor operator. Thus, the problem of ensuring some level of fairness in channel allocation becomes an interesting one.

We remark that fairness does not necessarily imply "equal blocking probability." Various definitions of fairness are possible but the most appropriate one appears to be the notion of "max-min fairness." A similar notion of fairness is widely used in the context of flow control [3]. Roughly speaking, max-min fairness minimizes the overall blocking probability without decreasing the blocking probability in any cell at the expense of other cells which are already worse off. More precise definitions follow in later sections.

We make the following assumptions regarding the cellular system. The system consists of N cells and the underlying offered traffic model is independent from cell to cell; in particular, we ignore the effect of call handovers and intercell calls. However, it is likely that we can extend our results to the case in which this independence assumption is dropped and handovers and intercell calls can be included. Our optimism is derived from the fact that the results of [7], to which we shall refer extensively, have been extended to include handovers in [9]. The call requests form a Poisson process and the call duration is exponentially distributed. The N cells share a common set of n channels. If A_i denotes the offered traffic in cell i (the expected number of calls that would be in progress in cell i if all call requests could be honored), then A_i/n is the *offered traffic intensity* in cell i. The offered traffic intensity in the system, r, is the sum of the offered traffic intensities in the cells, thus $r = \sum A_i/n$. The ratio $p_i = A_i/\sum_{i=1}^{N} A_i$ represents the fraction of the total offered traffic in cell i and the vector $\tilde{p} = (p_1, p_2, \ldots, p_N)$ is the *traffic pattern*. The carried traffic intensity in cell i, x_i, is the carried traffic (expected number of calls in progress) in cell i per available channel in the system. We call the vector $\tilde{x} = (x_1, x_2 \ldots, x_N)$ the *carried traffic intensity vector*. When a call request arrives in a cell, the channel assignment algorithm either assigns it a channel for its entire duration, or blocks it. (A blocked call

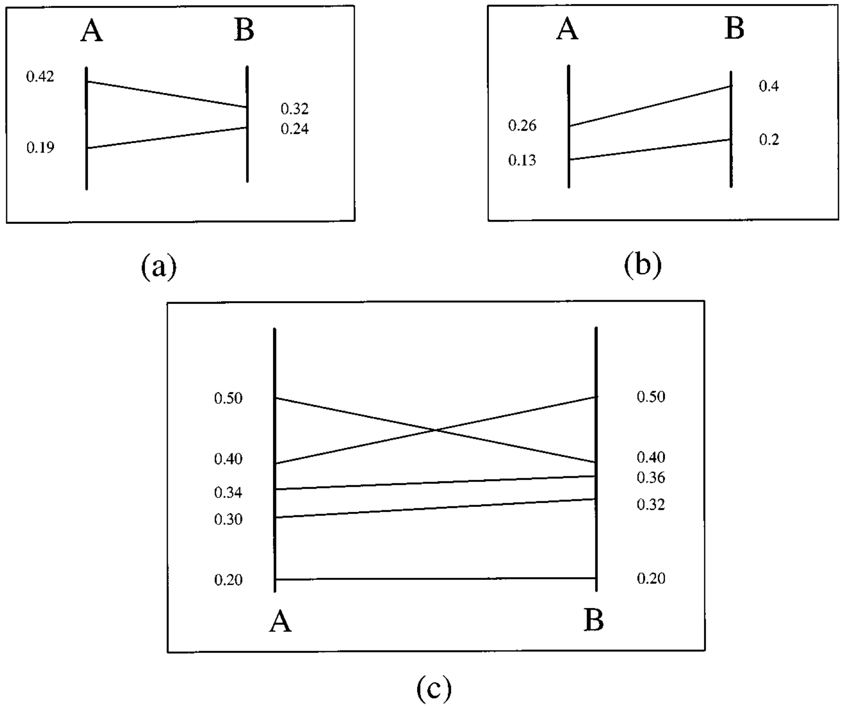

Fig. 2. Scale diagram. In all three cases illustrated B is fairer than A. In (a) and (b) $N = 2$. In (c) $N = 5$. For both A, B of (a) and (b) and A of (c) components are represented in order, with component 1 the lowermost. For B of (c) the order is 1, 2, 3, 5, 4, component 1 being the lowermost on the vertical line. Note that corresponding components of B are joined with those of A. The numbers are the values of the corresponding call acceptances.

disappears from the system.) The decision regarding blocking or honoring a call request is taken by the channel assignment algorithm based entirely on the current value of a suitably defined *state vector*. Further, we assume that the state vector assumes a finite, or countably infinite, set of values. In a nutshell, we assume that channel assignment algorithms can be modeled by continuous-time Markov chains (CTMCs). Further, we assume that there exists a probability distribution on the states of the CTMC, which satisfies the balance equations and all state probabilities are positive.[1] An example of a state vector is a list that specifies the calls in progress in each cell of the system and the channels assigned to them. The assignment of channels to calls at any stage must satisfy certain reuse constraints which essentially specify the sets of cells that may simultaneously use the same channel. Each possible channel assignment algorithm corresponds to one and only one carried traffic intensity vector under fixed r, $\tilde{p}$, and n. We denote the set of achievable carried traffic intensity vectors under fixed r, $\tilde{p}$, and n, by $F(r, \tilde{p}, n)$ or simply F. The blocking probability in cell i is $b_i = 1 - x_i/p_i r$. We denote c_i as the call acceptance in cell i, where $c_i = x_i/p_i$. $\tilde{c} = (c_1, c_2, \ldots, c_N)$ is denoted as the *call acceptance vector*. c_i/r is the call acceptance probability in cell i. We consider only systems in which $p_i > 0$ for all i.[2] The carried traffic, blocking probability, and call acceptance probability depend on the channel assignment algorithm used and we shall denote this by an appropriate superscript. For example, $\tilde{x}^A$ is the carried traffic intensity vector under channel assignment algorithm A.

We will proceed as follows. In Section II, we formally define the concept of relative fairness and present a simple necessary and sufficient condition for relative fairness. In Section III, we introduce the notion of absolute fairness and prove its existence and uniqueness. In Section IV, we deal with fairness in the asymptotic limit, i.e., when the number of channels is arbitrarily large. We also discuss the computation of the absolutely fair carried traffic intensity vector in the asymptotic limit and give the corresponding channel allocation strategy to attain it. In Section V, we discuss the loss of revenue brought about by fairness and discuss a tradeoff between being fair and maximizing the revenue. We conclude this correspondence by discussing how our work can be extended.

II. Relative Fairness

Informally, a channel assignment algorithm A is fairer than another channel assignment algorithm B, for given r, $\tilde{p}$, and n, if for every cell i whose blocking probability is decreased by B compared to A, there is some other cell j whose blocking probability was already no less than that of i under A and has been increased further by B. A more formal definition of relative fairness is as follows.

A channel assignment algorithm A is *fairer* than another channel assignment algorithm B, under the same r, $\tilde{p}$ and n if

- $\tilde{x}^A \neq \tilde{x}^B$ and
- if there exists an i such that $b_i^A > b_i^B$, then there exists a j such that $b_j^A \geq b_i^A$ and $b_j^B > b_j^A$. Equivalently, if there exists an i such that $c_i^A < c_i^B$, then there exists a j such that $c_j^A \leq c_i^A$ and $c_j^B < c_j^A$.

Since a channel assignment algorithm corresponds to a unique carried traffic intensity vector for fixed r, $\tilde{p}$, and n, we will usually speak of the fairness of carried traffic intensity vectors rather than the fairness of channel assignment algorithms. Note that given any one of the vectors, $\tilde{x}$, $\tilde{b} = (b_1, \ldots, b_N)$, and $\tilde{c} = (c_1, \ldots, c_N)$, we can determine the other two uniquely (since r and $\tilde{p}$ are fixed). Thus, we may also equivalently consider the fairness of the vectors $\tilde{b}$ or $\tilde{c}$ instead of $\tilde{x}$.

Fig. 2[3] shows pairs of carried traffic intensity vectors. In each case, one is fairer than the other.

[1]This is equivalent to assuming the CTMCs to be positive, recurrent, and regular.

[2]This means that we eliminate from the system cells with no offered traffic.

[3]Scale diagram is a schematic representation of vectors. The components (call acceptances for the cellular examples) are represented as points on a vertical line. Corresponding components of vectors are joined.

Theorem 1: A channel assignment algorithm A is fairer than another channel assignment algorithm B if and only if there exists an $A_{\min} \in S \subseteq U = \{1, 2, \ldots N\}$ such that $c^A_{A_{\min}} = \min_{i \in S} c^A_i$ and $c^A_{A_{\min}} > c^B_{A_{\min}}$, where $S = \{i: c^A_i \neq c^B_i\}$.

Remarks: In other words, this theorem states that the necessary and sufficient condition for fairness of a channel assignment algorithm A over another B is that the least call acceptance probability (over all cells, ignoring those with equal call acceptance probabilities under both A and B) under A is strictly greater than the corresponding call acceptance probability in B. This theorem will be useful in discussing absolute fairness. Note that relative fairness is different from lexicographic ordering.[4] If a vector is fairer than another, it is lexicographically greater as well. However, a vector may be lexicographically greater than another, but neither of the two may be fairer than the other. Thus, the necessary and sufficient condition in this theorem is sufficient but not necessary for lexicographic comparison.

Proof of Theorem 1: The proof will be given in two steps. First, the sufficiency will be proved and then the necessity

$$\text{Let } c^A_{A_{\min}} > c^B_{A_{\min}}, \qquad \text{where } c^A_{A_{\min}} = \min_{t \in S} c^A_t, \; A_{\min} \in S. \tag{1}$$

Then for every $j \in U$, for which $c^A_j < c^B_j$, there exists $A_{\min} \in U$ such that $c^B_{A_{\min}} < c^A_{A_{\min}}$ and

$$c^A_{A_{\min}} = \min_{t \in S} c^A_t \leq c^A_j.$$

(Since $c^A_j \neq c^B_j$, $j \in S$.) Since $A_{\min} \in S$, S is nonempty and thus $\tilde{x}^A \neq \tilde{x}^B$. Thus, A is fairer than B.

Let A be fairer than B. Thus, $\tilde{x}^A \neq \tilde{x}^B$ and hence S is nonempty. If possible, let

$$c^A_{A_{\min}} \leq c^B_{A_{\min}}$$

for all $A_{\min}$ satisfying (1). Since $A_{\min} \in S$, $c^A_{A_{\min}} \neq c^B_{A_{\min}}$. Hence,

$$c^A_{A_{\min}} < c^B_{A_{\min}} \tag{2}$$

fo all $A_{\min}$ satisfying (1).

Since A is fairer than B, there exists $j \in U$, for which

$$c^A_j \leq c^A_{A_{\min}} \tag{3}$$

and

$$c^A_j > c^B_j. \tag{4}$$

For $j \in U \backslash S$, $(c^A_j = c^B_j)$ and, hence, inequality (4) is never satisfied for any $j \in U \setminus S$. For $j \in S$, from inequality (3), and the definition of $A_{\min}$ in (1)

$$c^A_j = c^A_{A_{\min}} = \min_{t \in S} c^A_t$$

i.e., j satisfies (1). Hence, from (2), (4) cannot be satisfied for any $j \in S$. Therefore, there is no $j \in U$ for which both inequalities (3) and (4) are satisfied and that is a contradiction. $\diamond$

We now prove another interesting result which we shall use in obtaining some results on absolute fairness.

Lemma 1: If the channel assignment algorithm A is fairer than the channel assignment algorithm B, then B cannot be fairer than A. (In other words, fairness is an antisymmetric relation.)

[4]Two vectors can be lexicographically compared as follows. If the minimum components are unequal, then the vector with a larger minimum component is lexicographically greater. If the minimum components are equal, then the second minimum component must be considered, and so on.

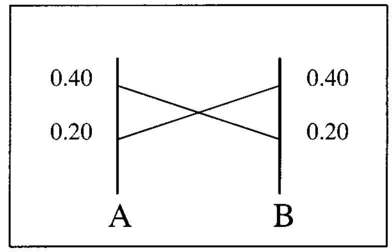

Fig. 3. Scale diagram. This figure compares the fairness of two vectors A and B with two components each. Neither A nor B is fairer than the other. For A, components are represented in order, with component 1 the lowermost. For B, the order is 2,1, component 2 being the lowermost on the vertical line. Note that the corresponding components of B are joined with those of A. The numbers are the values of the corresponding call acceptances.

Proof of Lemma 1: Assume there exist channel assignment algorithms A and B, each fairer than the other. Since A is fairer than B, by Theorem 1

$$\begin{aligned} &\exists\, A_{\min} \in S \text{ such that } c^A_{A_{\min}} = \min_{i \in S} c^A_i > c^B_{A_{\min}} \\ &c^B_{A_{\min}} \geq \min_{i \in S} c^B_i \text{ since } A_{\min} \in S. \end{aligned} \tag{5}$$

Since B is fairer than A, by Theorem 1

$$\exists\, B_{\min} \in S \text{ such that } \min_{i \in S} c^B_i = c^B_{B_{\min}} > c^A_{B_{\min}}.$$

Again $c^A_{B_{\min}} \geq \min_{i \in S} c^A_i = c^A_{A_{\min}}$, since $B_{\min} \in S$ and by inequality (5).

Combining the above inequalities we get

$$c^A_{A_{\min}} > c^B_{A_{\min}} \geq c^B_{B_{\min}} > c^A_{B_{\min}} \geq c^A_{A_{\min}}.$$

This is a contradiction. Hence the existence of two channel assignment algorithms mutually fairer than each other is not possible. $\square$

III. Absolute Fairness

A channel assignment algorithm is *absolutely fair* or *max-min fair* if it is fairer than any other channel assignment algorithm which achieves a different carried traffic intensity vector under the same r, $\tilde{p}$, and n. We shall refer to the carried traffic intensity vector corresponding to the max-min fair channel assignment algorithm as the MMF carried traffic intensity vector, or simply the MMF. Informally, a carried traffic intensity vector is the MMF if the blocking (probability) in any cell under it cannot be decreased without increasing the blocking in any cell already experiencing greater or equal blocking.

It is not obvious that every set of carried traffic intensity vectors F has an MMF. Consider the following examples.

Example III.1: Let F consist of only the carried traffic intensity vectors corresponding to the vectors $\tilde{c}^A$ and $\tilde{c}^B$ shown in Fig. 3. Neither A nor B is fairer than the other. Hence, no MMF exists in F in either case.

Example III.2: Let $F = R^N$. (R^N is the set of N-dimensional vectors with real components.) For any vector $\tilde{x}^V \in F$, there exists another vector, $\tilde{x}^C \in F$, each of whose components is greater than the corresponding one in $\tilde{x}^V$. Thus, $\tilde{x}^V$ is not fairer than $\tilde{x}^C$ and hence is not the MMF. The same observation applies for

$$F = \{\tilde{x} = (x_1, x_2, \ldots, x_N): 0 \leq x_i < 1\}.$$

In view of the above examples, it becomes necessary to prove the existence of an MMF. In this section, we will first show that any set of vectors with certain properties has an MMF. The definition of fairness in any set of N-dimensional vectors remains the same as that for carried traffic intensity vectors. $c_i^V = x_i^V / p_i$, where x_i^V, is the ith component of the vector $\tilde{V}$. $\tilde{p} = (p_1, p_2, \ldots, p_N)$ is the same for the entire set. $p_i > 0$, for all $i \in U$. The necessary and sufficient condition given in Theorem 1 for relative fairness of channel assignment algorithms holds for that of vectors in any N-dimensional vector set as well. A vector is the MMF if it is fairer than all other vectors in the set. Next, we discuss whether the set of achievable carried traffic intensity vectors of a cellular system has these properties. We end this section with the result that the MMF is unique.

Theorem 2: If P is a nonempty, coordinate convex,[5] convex, closed, bounded subset of R^N, P has an MMF.

Proof of Theorem 2: For brevity, we only give a construction for a vector that can be proved to be max-min fair using Theorem 1. For details, refer to [11]. Unless otherwise stated, $U = \{1, 2, \ldots, N\}$.

Construction: Start from the all-zero vector $c_i = 0, \forall i \in U$.

Always find the largest subset $I_l \subseteq U$ such that c_i's, $i \in I_l$, can be increased equally without decreasing any c_i. Increase the c_i's of this subset equally without decreasing any c_i, till it is no longer possible to do so.

The procedure terminates when the largest subset which can be increased equally without decreasing any c_i is empty, i.e., no c_i can be increased without decreasing some other c_i.

Informally, the procedure goes as follows: Start from the vector $(0, 0 \ldots, 0)$. Suppose the largest subset which can be increased is $\{1, 2, \ldots, k\}$, where $1 \le k \le N$. Increase $c_1, \ldots, c_k$ equally till a vector $(p_1\alpha, \ldots, p_k\alpha, 0, \ldots, 0)$ is reached and $\alpha \ge 0$ is such that

$$(p_1(\alpha+\epsilon), \ldots, p_k(\alpha+\epsilon), 0, \ldots, 0) \notin P$$

for any $\epsilon > 0$. Again increase the largest possible subset equally without decreasing any component. The largest possible subset will be a proper subset of $\{1, \ldots, k\}$. Let it be $\{1, \ldots, m\}$, $m < k$. The increase of this subset continues till

$$(p_1\beta, \ldots, p_m\beta, p_{m+1}\alpha, \ldots, p_k\alpha, 0, \ldots, 0)$$

is reached and

$$(p_1(\beta+\epsilon), \ldots, p_m(\beta+\epsilon), p_{m+1}\alpha, \ldots, p_k\alpha, 0, \ldots, 0) \notin P$$

for any $\epsilon > 0$. The procedure terminates when no element can be increased without decreasing some other(s).

Note that this is a generalization of the construction of max-min fair rate allocation presented in [3]. The construction in [3] applies to polytope feasible sets defined by linear constraints of the type, $AX \le B$. The construction we present here applies to more general feasible sets.

We discuss the validity of the assumptions that the set of carried traffic intensity vectors is nonempty, closed, bounded, convex and coordinate convex.

Nonempty: The set of carried traffic intensity vectors is clearly nonempty.

Closed: In an asymptotic limit where the number of channels and the offered traffic are made arbitrarily large while keeping the ratio finite, and which we consider in the next section, for the model we have assumed (underlying model of offered traffic independent from cell to cell, etc.), the set of carried traffic intensity vectors is given by the following inequalities (refer to [7]):

$$x_i \le \sum_{j=1}^{M} X_j a_{ij}, \qquad i = 1, 2, \ldots N \tag{6}$$

$$X_j \ge 0 \tag{7}$$

$$\sum_{j=1}^{M} X_j = 1 \tag{8}$$

$$x_i \le p_i r, \qquad i = 1, 2, \ldots N \tag{9}$$

$$x_i \ge 0, \qquad i = 1, 2, \ldots N \tag{10}$$

where $a_{ij} = 0$ or 1 depending on the system configuration and channel reuse constraint. The a_{ij} can be found by modeling the system by a hypergraph [7]. This set is closed. For the finite channel case, we could not prove that the set of carried traffic intensity vectors F is closed. This remains an assumption. However, even if the assumption turns out to be invalid our result is not seriously weakened. It can be shown that the closure of any convex, coordinate convex, and bounded set is also convex, coordinate convex, and bounded and obviously closed. Also, clearly closure of a nonempty set is nonempty. We are going to show that the set of carried traffic intensity vectors is convex, coordinate convex, and bounded. Thus, the closure of the set of carried traffic intensity vectors is nonempty, convex, coordinate convex, bounded, and a closed subset of R^N and hence has an MMF $\tilde{x}^M$ (by Theorem 2). Even if $\tilde{x}^M$ does not belong to F (closedness is not necessary for the existence of MMF, $\tilde{x}^M$ may belong to F), we have carried traffic intensity vectors in F arbitrarily close to $\tilde{x}^M$. For practical purposes, it suffices to assume that the set of carried traffic intensity vectors is closed.

Bounded: The set of carried traffic intensity vectors F is bounded. This is because the carried traffic intensity can neither exceed the offered traffic intensity in any cell nor can it exceed 1 (since the total number of channels available in the system is n, the traffic carried in any cell cannot exceed n) and it is also nonnegative. Thus, $0 \le x_i \le \min(p_i r, 1)$ for each i and hence $0 \le c_i \le \min(r, 1/p_i)$ for each i, where $0 < p_i < 1$.

Theorem 3: The set of achievable carried traffic intensity vectors for a cellular system F is convex.

For brevity, we give an outline of the proof of this important result. Refer to [11] for details. Our assumptions about channel assignment algorithms, the call arrival process, and call duration distribution allow the operation of the channel assignment algorithms to be modeled by a CTMC. Furthermore, we assume the CTMCs to be positive recurrent and irreducible. Let there be two carried traffic intensity vectors $\tilde{x}^A$ and $\tilde{x}^B$ realized by channel assignment algorithms A and B, respectively. Let $S_A = \{A_1, A_2, \ldots, \}$ (resp., $S_B = \{B_1, B_2, \ldots, \}$) denote the state space of A (resp., B). Denote one null state (state in which the number of calls in progress in the system is 0) of A by T_A and one of B by T_B. A third channel assignment algorithm C, whose state space is the union of S_A and S_B, operates as follows. When in T_A it switches over to T_B at rate γ_{AB} and when in T_B it switches over to T_A at rate γ_{BA}. Otherwise, C behaves in the same manner as A or B depending on whether it is in a state that belongs to S_A or S_B, respectively (all other transition rates remain the same). It can be shown that the CTMC corresponding to C is positive recurrent and irreducible. $\tilde{x}^C = \alpha\tilde{x}^A + (1-\alpha)\tilde{x}^B$, where $\alpha = (1+\gamma\pi^A_{T_A}/\pi^B_{T_B})^{-1}$, $\gamma = \gamma_{AB}/\gamma_{BA}$, and $\{\pi^A\}$ and $\{\pi^B\}$ are the steady-state probability distributions of algorithms A and B, respectively. By adjusting γ we can get any $\alpha \in (0, 1)$.

Theorem 4: The set of achievable carried traffic intensity vectors for a cellular system F is coordinate convex.

We need to prove that if there exists a carried traffic intensity vector $\tilde{x}^A \in F$, then any vector of the form

$$\tilde{x}^B = (x_1^A, x_2^A, \ldots, \alpha_i x_i, \ldots, x_N^A) \in F$$

[5]A subset of R_+^N (R_+^N is the set of all N-dimensional vectors with real nonnegative components), F, is said to be coordinate convex if for each vector $\tilde{x} = (x_1, x_2, \ldots, x_N)$ in F, $\tilde{x}' = (x_1, x_2, \ldots, \alpha x_i, \ldots, x_N)$, belongs to F for all α satisfying $0 \le \alpha < 1$ and for all $i \in \{1, 2, \ldots, N\}$.

where $0 \le \alpha_i < 1$, for any $i \in U$. A slight variation of the technique used to prove Theorem 3 in conjunction with Theorem 3 can be used to prove this theorem. Refer to [11] for details.

Thus, the set of achievable carried traffic intensity vectors for a cellular system is closed, bounded, convex, and coordinate convex. It is also a subset of R^N. Hence, it has an MMF by Theorem 2.

Theorem 5: MMF is unique.

Proof of Theorem 5: Let the max-min fair carried traffic intensity vector (MMF) not be unique. Hence, there exist at least two different MMFs, $\tilde{x}^{M_1}$ and $\tilde{x}^{M_2}$. Thus, the corresponding channel assignment algorithms, M_1, M_2 are both fairer than each other ($\tilde{x}^{M_1} \neq \tilde{x}^{M_2}$). This violates Lemma 1. Hence this is not possible. ◊

IV. Properties and Computation of MMF in the Asymptotic Limit

In this section, we shall consider an asymptotic limit where $A_i \to \infty$ and $n \to \infty$ but

$$\lim_{A_i \to \infty,\, n \to \infty} A_i/n = p_i r$$

is finite. Thus, both the offered traffic and the number of channels are made arbitrarily large while keeping the ratio—the offered traffic intensity—finite. The properties of channel assignment algorithms in this asymptotic limit were studied in [7]. The feasible set of carried traffic vectors can be described by some linear inequalities in this case. Incidentally, [5] studied a different notion for fairness, proportional fairness, for resource allocation in wireline case, and presented computational strategies for obtaining the same. Also, [5] showed that max-min fairness is a limiting case of a generalization of proportional fairness, and as such a max-min fair allocation can be approximated arbitrary closely using the techniques for computing proportionally fair allocation. The similarity between the wireline case considered in [5] and the asymptotic case for cellular networks considered in this section is that the feasible sets can be described by linear inequalities in both cases (the nature of the inequalities differ in the two cases though). The computational approaches differ in the two cases. The important advantage of our approach is that we exploit specific properties of max-min fair allocation and the feasible set in the cellular mobile case to develop parametrized closed-form expressions for MMFs and the corresponding channel allocations. The parameters in the closed-form expressions can be obtained by solving linear programs. The approach in [5] is to use a nonlinear optimization based iterative update technique for obtaining the proportionally fair allocation, and the iterative procedure is not guaranteed to converge in finite number of iterations. We present our results in what follows.

Let $\tilde{x}^{M\infty}(r)$ denote the MMF carried traffic intensity vector at load r in this asymptotic limit. In this limit, the set of achievable carried traffic intensity vectors $F_\infty(r)$ at r is described by inequalities (6)–(10) for our model [7]. The proofs in this section use results from [7].

Theorem 6: There exist finite nonnegative $\{\beta_i\}$ such that for each $i \in U$ and $r \ge 0$

$$c_i^{M\infty}(r) = \min(r,\, \beta_i).$$

Proof of Theorem 6: First we shall prove that $c_i^{M\infty}(r)$ is a nondecreasing function of r. Consider $0 \le r_1 < r_2$. We show that $c_i^{M\infty}(r_1) \le c_i^{M\infty}(r_2)$ for all $i \in U$. Carry out the construction given in proof of Theorem 2 for both, in the asymptotic limit. The outputs will be the respective unique MMF's, $\tilde{x}^{M\infty}(r_1)$ and $\tilde{x}^{M\infty}(r_2)$. Let the construction procedures for both remain identical upto a certain point, when the carried traffic intensity vector reached is $\tilde{x}^V$ (since construction for both start from the same vector, i.e., the null vector, the construction remains identical for some time. $\tilde{x}^V$ may be the all-zero vector) and then let the constructions differ (if the constructions never differ $\tilde{x}^{M\infty}(r_1) = \tilde{x}^{M\infty}(r_2)$, $c_i^{M\infty}(r_1) = c_i^{M\infty}(r_2)$ for all $i \in U$ and there is nothing to prove). Let the largest subset of components of $\tilde{x}^V$ whose c_is can be increased equally without decreasing others be I_1 and I_2 at r_1 and r_2, respectively. Since the constructions differ, henceforth $I_1 \neq I_2$. Let, if possible, I_1 be nonempty. At any r, if I is the largest subset of components of any vector $\tilde{x}^V \in F_\infty(r)$, whose c_i's can be increased equally without decreasing any other component, then no c_j^V of $\tilde{x}^V$ with j in $U \setminus I$ can be increased without decreasing any other component, else from convexity and coordinate convexity of $F_\infty(r)$, all c_i's of $I \cup \{j\} \supset I$ can be increased equally without decreasing any other component. Thus, no component of $\tilde{x}^V$ in $U \setminus I_1$ ($U \setminus I_2$) can be increased at r_1 (r_2). Observe $F_\infty(r_1) \subseteq F_\infty(r_2)$ for $r_1 < r_2$. Thus, the c_i's of $\tilde{x}^V$ in I_1 can also be increased equally without decreasing others at r_2. Hence, $I_1 \subset I_2$. Since c_i's of $\tilde{x}^V$ in $I_2 \setminus I_1$ can be increased at r_2, but not at r_1, without decreasing others, from inequalities (6)–(10), $c_i^V = r_1$, for $i \in I_2 \setminus I_1$, but from construction principle at r_2, $c_i^V = c_j^V$ for all $i,\, j \in I_2$ (till this point c_i's of sets $\supseteq I_2$ have been increased equally and that is how $\tilde{x}^V$ has been reached). Thus, $c_i^V = r_1$ for all $i \in I_2 \supset I_1$ and no component of $\tilde{x}^V$ in I_1 can be increased at r_1 (inequality (9)). Thus, $I_1 = \phi$, i.e., construction at r_1 has terminated and $\tilde{x}^V = \tilde{x}^{M\infty}(r_1)$. Since $\tilde{x}^V$ is an intermediate vector in the construction procedure for MMF at r_2, from construction principle

$$c_i^{M\infty}(r_2) \ge c_i^V = c_i^{M\infty}(r_1)$$

for all $i \in U$ and this part of the theorem is proved.

$c_i^{M\infty}(r) = 0$ for each $i \in U$ at $r = 0$. Thus, for each $i \in U$, there exists some $\alpha_i \ge 0$ (possibly 0) such that for all $r \le \alpha_i$, $c_i^{M\infty}(r) = r$. Clearly, for each $i \in U$, there exists some finite $r \ge 0$, such that $c_i^{M\infty}(r) < r$ (e.g., for $r > M$). Let

$$\beta_i = \inf_{r:\, c_i^{M\infty}(r) < r} r < \infty.$$

$c_i^{M\infty}(r) = r$ for $r < \beta_i$. Since $c_i^{M\infty}(r)$ is a nondecreasing function of r, and $c_i^{M\infty}(\beta_i) \le \beta_i$ (inequality (9))

$$c_i^{M\infty}(\beta_i) \in [\beta_i - \epsilon,\, \beta_i], \qquad \text{for all } \epsilon > 0.$$

Hence, $c_i^{M\infty}(\beta_i) = \beta_i$. Consider any $r_1 > \beta_i$, for which $c_i^{M\infty}(r_1) < r_1$ (clearly, there exists at least one such finite r_1 since β_i is finite). Consider any $r_2 > r_1$. Let $\tilde{x}^V$ be a vector with $c_j^V = \min(r_1,\, c_j^{M\infty}(r_2))$, for each $j \in U$. Clearly, $\tilde{x}^V \in F_\infty(r_1)$. Since $c_j^{M\infty}(r)$ is a nondecreasing function of r, and $c_j^{M\infty}(r) \le r$ (from inequality (9))

$$c_j^{M\infty}(r_1) \le \min(r_1,\, c_j^{M\infty}(r_2)) = c_j^V, \qquad \text{for each } j \in U.$$

Since $\tilde{x}^V$ is not fairer than $\tilde{x}^{M\infty}(r_1)$

$$c_j^{M\infty}(r_1) = \min(r_1,\, c_j^{M\infty}(r_2))$$

for each $j \in U$ and $c_i^{M\infty}(r_1) < r_1$. Thus, $c_i^{M\infty}(r_1) = c_i^{M\infty}(r_2)$ and from the definition of β_i

$$c_i^{M\infty}(r_1) = c_i^{M\infty}(r_2), \qquad \text{for all } r_1,\, r_2 > \beta_i.$$

Again, from the nondecreasing property of $c_i^{M\infty}(r)$, and since it is upper-bounded by r

$$c_i^{M\infty}(r) \in [\beta_i,\, \beta_i + \epsilon]$$

for each $\epsilon > 0$ and for any $r > \beta_i$. Thus, $c_i^{M\infty}(r) = \beta_i$, for all $r > \beta_i$. ◊

The construction procedure given in proof for Theorem 2 can be implemented using the inequalities (6)–(10) to yield $\tilde{x}^{M\infty}(r)$, at any

$r \geq 0$, as follows. Start from the null vector. Find the largest α (possibly 0), such that $(p_1\alpha, p_2\alpha, \ldots, p_N\alpha) \in F_\infty(r)$

$$\alpha = \max\left\{ \begin{array}{l} s\colon s \leq \left(\sum_{j=1}^{M} X_j a_{ij}\right)\Big/ p_i,\ i=1, 2, \ldots, N; \\ \sum_{j=1}^{M} X_j = 1;\ X_j \geq 0,\ j=1, 2, \ldots, M;\ s \leq r \end{array} \right\}.$$

Next find $I = \{i\colon \tilde{x}^{V_i} \in F_\infty(r)\}$, where $\tilde{x}^{V_i}$ is any carried traffic intensity vector of the form

$$c_j^{V_i} \begin{cases} > \alpha, & j = i \\ \geq \alpha, & \text{otherwise} \end{cases}$$

(inequalities (6)–(10) may be used to test whether $\tilde{x}^{V_i} \in F_\infty(r)$). I is thus the largest subset whose call acceptances can be increased equally without decreasing any other component of $(p_1\alpha, p_2\alpha, \ldots, p_N\alpha)$. Without loss of generality, let $I = \{1, 2, \ldots, k\}$, $k < N$. All call acceptances of I cannot be increased equally without decreasing any other component when a carried traffic intensity vector $(p_1\beta, p_2\beta, \ldots, p_k\beta, p_{k+1}\alpha, \ldots, p_{N+1}\alpha)$ is obtained such that

$$(p_1(\beta+\epsilon), p_2(\beta+\epsilon), \ldots, p_k(\beta+\epsilon), p_{k+1}\alpha, \ldots, p_{N+1}\alpha) \notin F_\infty(r)$$

for any $\epsilon > 0$

$$\beta = \max\left\{ \begin{array}{c} s\colon s \leq \left(\sum_{j=1}^{M} X_j a_{ij}\right)\Big/ p_i,\ i \in I; \\ \sum_{j=1}^{M} X_j a_{ij} \geq p_i\alpha,\ i \in U \setminus I; \\ \sum_{j=1}^{M} X_j = 1;\ X_j \geq 0,\ j=1, 2, \ldots, M;\ s \leq r \end{array} \right\}.$$

Again, $I_1 = \{i\colon \tilde{x}^{V_i} \in F_\infty(r)\}$ is the largest subset of components whose call acceptances can be increased without decreasing any other component of $(p_1\beta, p_2\beta, \ldots, p_k\beta, p_{k+1}\alpha, \ldots, p_{N+1}\alpha)$, where $\tilde{x}^{V_i}$ is any carried traffic intensity vector of the of the form:

$$c_j^{V_i} \begin{cases} > \beta, & j = i \\ \geq \beta, & j \in I \setminus \{i\} \\ \geq \alpha, & \text{otherwise}. \end{cases}$$

The procedure terminates when no call acceptance can be increased without decreasing others. The output will be $\tilde{x}^{M_\infty}(r)$.

If $\tilde{x}^{M_\infty}(r)$ is found at some sufficiently large r, say r' for which $c_i^{M_\infty}(r') < r'$, for all $i \in U$, then $c_i^{M_\infty}(r') = \beta_i$, for all $i \in U$, from Theorem 6. Thus, the β_i's can be determined from $\tilde{x}^{M_\infty}(r')$. Using these β_i's, $\tilde{x}^{M_\infty}(r)$ can be determined for all $r \geq 0$ from Theorem 6. Thus, $\tilde{x}^{M_\infty}(r)$ need not be found separately at all values of r. (Using inequalities (6), (8), and (9), for any $r > 1/\min_{i\in U} p_i$, $c_i^V < r$, for all $i \in U$, and any $\tilde{x}^V \in F_\infty(r)$. Hence, $c_i^{M_\infty}(r) < r$ for each $i \in U$, if $r > 1/\min_{i\in U} p_i$.]

Theorem 7: The carried traffic intensity vector corresponding to the fixed channel assignment algorithm which allocates $\lfloor np_i\beta_i \rfloor$ channels to cell i, for all $i \in U$, is the MMF for all $r \geq 0$ in the asymptotic limit.

Remark: This fixed channel allocation scheme is similar to the one used to attain the highest carried traffic intensity in [7].

First, we show that it is possible to allocate $\lfloor np_i\beta_i \rfloor$ channels to each cell i at all r. For some $r_l \geq \max_{i\in U} \beta_i$, $c_i^{M_\infty}(r_l) = \beta_i$, by Theorem 6 ($r_l$ is finite). Thus, $x_i = p_i\beta_i$ satisfies inequality (6) at all r where $\{X_j\}$ of inequality (6) satisfies inequalities (7) and (8) at all r. From [7] $\lfloor n \sum_{j=1}^{M} X_j a_{ij} \rfloor$ channels can be allocated to each cell i where $\{X_j\}$ satisfies inequalities (7) and (8). Thus, from inequality (6) $\lfloor np_i\beta_i \rfloor$ channels can be allocated to each cell i at all r.

Again from [7], in the asymptotic limit any allocation which gives $\lfloor nf_i \rfloor$ channels to a cell i attains a carried traffic intensity of $\min(f_i, r_i)$ in cell i, if offered traffic intensity in cell i is r_i for the model we have assumed. Thus, the algorithm allocating $\lfloor np_i\beta_i \rfloor$ channels to each cell i attains a carried traffic intensity of $p_i \min(\beta_i, r)$ in each call i in the asymptotic limit, for all r, i.e., $c_i(r) = \min(\beta_i, r) = c_i^{M_\infty}(r)$ at all r (Theorem 6). $\diamondsuit$

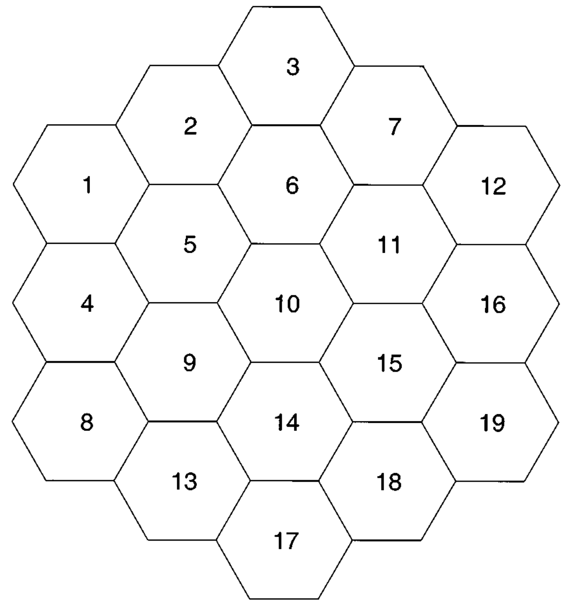

Fig. 4. Nineteen-cell system.

Example IV.1: Consider the system of 19 cells (see Fig. 4) with the following reuse constraint. A set of cells can use the same channel simultaneously if the interference in each cell of the set is $\leq 3/8$. The interference in any cell u is given by $\sum_{v\in S,\, v\neq u} d(u, v)^{-4}$, where S is the set of cells using the same channel simultaneously and $d(u, v)$ is the distance between the centers of cells u and v. We assume the cell radius to be $1/\sqrt{3}$ or equivalently, the distance between the centers of adjacent cells to be unity. (This system and reuse constraints are the same as those in [7, Example 1.2].)

1) We shall consider uniform traffic, i.e., $p_i = 1/19$ for each cell $i \in U$. The carried traffic intensity in the system under an MMF is given by

$$T_{M_\infty}(r) = \sum_{i=1}^{N} x_i^{M_\infty}(r) = \min(r, 247/49)$$

(see Fig. 6 (a)). $x_i^{M_\infty}(r) = \min(r/19, 13/49)$, $\forall i \in U$. Thus, $c_i^{M_\infty}(r) = x_i^{M_\infty}(r)/p_i = \min(r, 247/49)$, $\forall i \in U$. In Fig. 5(a), these call acceptances have been plotted. A fixed channel allocation algorithm allocating $\lfloor 13n/47 \rfloor$ channels to each cell achieves the MMF carried traffic intensity vector at all r in the asymptotic limit.

2) Let us consider nonuniform traffic

$$p_i = \begin{cases} 1/24, & i \in \{1, 2, 3, 4, 7, 8, 12 \\ & \quad 13, 16, 17, 18, 19\} \\ 1/16, & i \in \{5, 6, 9, 11, 14, 15\} \quad \text{and} \\ 1/8, & i = 10. \end{cases}$$

This pattern of nonuniform traffic resembles that in cities in which traffic is maximum in the central portion and decreases as we move toward the outskirts. The max-min fair carried traffic intensity in the system $T_{M_\infty}(r) = \min(r, 2 + \frac{r}{2}, \frac{13}{3})$ (refer to Fig. 6 (b)].

$$x_i^{M\infty}(r) = p_i \min(r, \beta_i), \forall i \in U.$$

$$\beta_i = \begin{cases} 14/3, & i \in \{1, 2, 3, 4, 7, 8, 12, \\ & \quad 13, 16, 17, 18, 19\} \\ 4, & \text{otherwise.} \end{cases}$$

In Fig. 5(b), call acceptances have been plotted for these cells. The following fixed channel allocation algorithm gives the MMF at all r in the asymptotic limit (n_i is the number of channels allocated to the ith cell):

$$n_i = \begin{cases} \lfloor 7n/36 \rfloor, & i \in \{1, 2, 3, 4, 7, 8, 12, \\ & \quad 13, 16, 17, 18, 19\} \\ \lfloor n/4 \rfloor, & i \in \{5, 6, 9, 11, 14, 15\} \quad \text{and} \\ \lfloor n/2 \rfloor, & i = 10. \end{cases}$$

V. Tradeoff Between Fairness and Revenue Maximization

More often than not, MMF is not the carried traffic intensity vector which yields the maximum revenue. At present, we consider the rate-per-call-per-unit time to be uniform throughout the system and hence the total carried traffic intensity in the system gives the rate at which revenue is earned, except for a multiplicative constant which we take as unity. We define the *marginal revenue* (see [8]) of a channel assignment algorithm A (or the corresponding carried traffic intensity vector $\tilde{x}^A(r)$) at load r and number of channels n, as

$$T_A(r, n) = \sum_{i=1}^{N} x_i^A(r).$$

So the marginal revenue of a channel assignment algorithm is the rate at which revenue is earned per channel using that algorithm. Let $T_M(r, n)$ be the marginal revenue of the MMF. Let

$$T_R(r, n) = \sum_{i=1}^{N} x_i^R(r)$$

be the maximum marginal revenue [8] and $\tilde{R}(r)$ the channel assignment algorithm ($\tilde{x}^R(r)$ the corresponding carried traffic intensity vector at load r) yielding this marginal revenue at load r. Let $T_{M_\infty}(r)$ and $T_{R_\infty}(r)$ be the corresponding asymptotic marginal revenues. Consider the following examples.

Example V.1:

1) Consider the system of Example IV.1 1). $T_{M_\infty}(r)$ and $T_{R_\infty}(r)$[6] have been plotted versus r in Fig. 6(a). For $r > 247/49$, $T_{M_\infty}(r) < T_{R_\infty}(r)$.

2) Consider the system of Example IV.1 2). For $r > 4.0$, $T_{M_\infty}(r) < T_{R_\infty}(r)$ (refer to Fig. 6(b)).

Clearly, $T_M(r, n) \leq T_R(r, n)$ for all systems. Generally this inequality is strict. Whenever this inequality is strict, the channel assignment algorithm attaining $T_R(r, n)$ is unfair compared to that attaining $T_M(r, n)$. In other words, revenue is maximized at the expense of fairness. If the operator maximizes revenue at the cost of fairness, it may cause customer dissatisfaction (customers will be dissatisfied throughout the system as they will experience poor-quality service when they move to certain cells) and possibly customer desertion to a competitor operator, which will result in reduced r and hence reduced revenue. Instead, if the network operator had opted for a carried traffic intensity vector which fetches less revenue but is fairer, customer desertion may have been less and it is possible that the ultimate revenue earned would be greater. While we are still in the process of modeling the effects of customer desertion on fairness and revenue, it seems plausible that the ideal channel assignment algorithm in many situations would earn less marginal revenue than $T_R(r, n)$ (which we get by not caring about fairness) and more than $T_M(r, n)$ (which we get by being max-min fair). It would also, in some sense (not necessarily that of Section II), be "fairer" than $\tilde{R}(r)$ but not "as fair" as the MMF channel assignment algorithm $\tilde{M}(r)$ at load r. A possible approach is to fix the

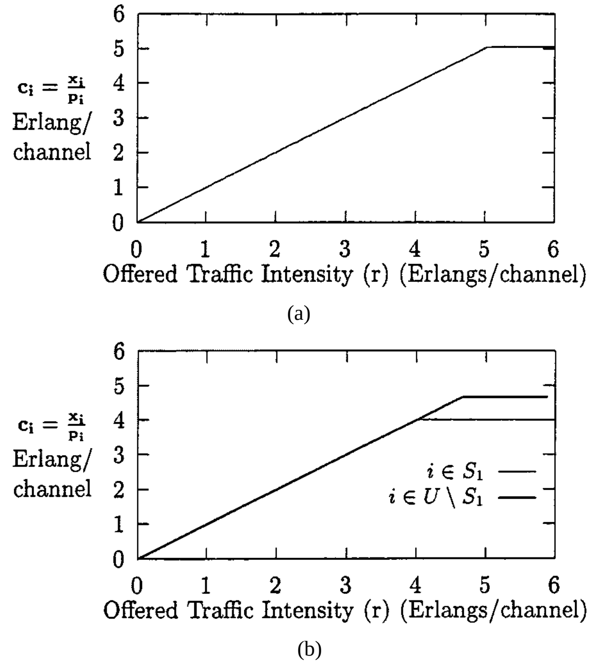

Fig. 5. Maxmin fair call acceptances (c_i) in a 19-cell system. $S_1 = \{5, 6, 9, 10, 11, 14, 15\}$ (central cell and its neighbors). (a) Uniform traffic. (b) Nonuniform traffic.

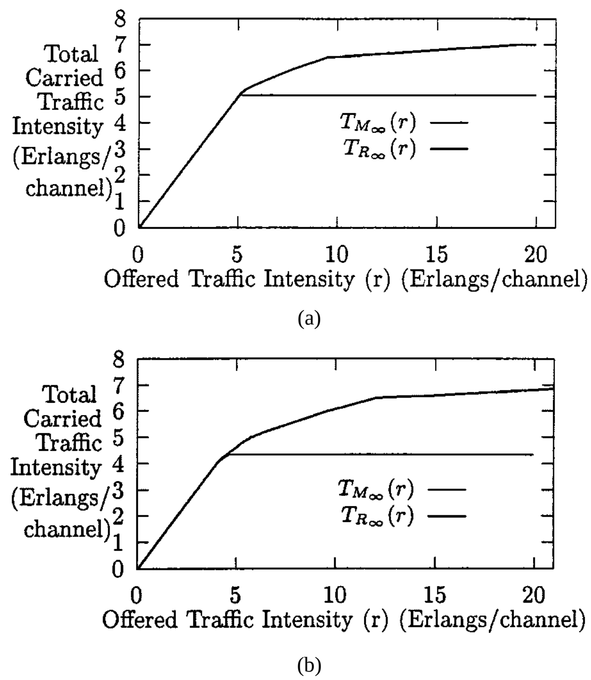

Fig. 6. Loss in revenue brought about by max-min fairness for a 19-cell system. (a) Uniform traffic. (b) Nonuniform traffic.

[6]Computation of $T_{R_\infty}(r)$ has been discussed in [7].

revenue at some value between $T_M(r, n)$ and $T_R(r, n)$, say m, and choose the channel assignment algorithm corresponding to the MMF among those carried traffic intensity vectors that *earn a marginal revenue of at least* m. We denote this carried traffic intensity vector by $\tilde{x}^M(m, r, n)$ ($\tilde{x}^{M\infty}(m, r)$ in the asymptotic limit). Consider the following example.

Example V.2:

1) Consider the system of 19 cells with uniform traffic (refer to Example IV.1 1)). At $r = 38/7, T_{R_\infty}(r) = 37/7 = 5.2856$ and $T_{M_\infty}(r) = 247/49 = 5.0408$. A carried traffic intensity vector $\tilde{x}^{R\infty}(r)$ which maximizes the marginal revenue in the asymptotic limit (computed as per [7]) has a blocking probability of 50% in the central cell (cell 10) and 0% in all other cells. On the other hand, $\tilde{x}^{M\infty}(m, r)$ at $m = 5.2$ gives following blocking probabilities:

$$b_i = \begin{cases} 11.175\%, & i \in \{5, 6, 9, 10, 11, 14, 15\} \\ 0.287\%, & i \in \{2, 4, 7, 13, 16, 18\} \\ 0\%, & \text{otherwise.} \end{cases}$$

Clearly, $\tilde{x}^{M\infty}(m, r)$ is much fairer than $\tilde{x}^{R\infty}(r)$, while it fetches only slightly less revenue (5.2 as compared to 5.286) at $r = 38/7$. This suggests that a network operator may be better off in the long run sacrificing the additional revenue achieved by $\tilde{x}^{R\infty}(r)$ and using $\tilde{x}^{M\infty}(5.2, r)$ instead, at $r = 38/7$.

2) Consider the system of 19 cells with nonuniform traffic (refer to Example IV.1 2)). At $r = 4.5, T_{R_\infty}(r) = 4.348$ and $T_{M_\infty}(r) = 4.25$. A carried traffic intensity vector $\tilde{x}^{R\infty}(r)$ which maximizes the marginal revenue in the asymptotic limit (computed as per [7]) has a blocking probability of 26.98% in the central cell (cell 10) and 0% in all other cells. On the other hand, $\tilde{x}^{M\infty}(m, r)$ at $m = 4.3$ gives following blocking probabilities:

$$b_i = \begin{cases} 15.553\%, & i = 10 \\ 6.668\%, & i \in \{5, 6, 9, 11, 14, 15\} \\ 0\%, & \text{otherwise.} \end{cases}$$

Clearly, $\tilde{x}^{M\infty}(4.3, r)$ is much fairer than $\tilde{x}^{R\infty}(r)$, while it fetches only slightly less revenue (4.3 as compared to 4.348) at $r = 4.5$. Again this suggests that a network operator may be better off in the long run sacrificing the additional revenue achieved by $\tilde{x}^{R\infty}(r)$ and using $\tilde{x}^{M\infty}(4.3, r)$ instead, at $r = 4.5$. MMF gives a blocking probability of 11.11% in cells 5, 6, 9, 10, 11, 14, 15, and 0% blocking in other cells.

The value of m could possibly be the result of modeling customer desertion in an appropriate manner. Further investigation in this direction is an interesting topic for future research.

We have so far assumed that the set of carried traffic intensity vectors $F(m, r, \tilde{p}, n)$ which fetch a marginal revenue of at least m at r has an MMF. This can be proved using the following theorem.

Theorem 8: If D is any nonempty, convex, closed, bounded subset of R_+^N, D has an MMF.

Note that the assumption of coordinate convexity is not required. Consider a set

$$D_s = \bigcup_{\tilde{A} \in D} S(\tilde{A})$$

where

$$S(\tilde{A}) = \{\tilde{V}: (0, 0 \ldots, 0) \le \tilde{V} \le \tilde{A}\}$$

where we say $\tilde{V_1} \le \tilde{V_2}$ if $x_i^{V_1} \le x_i^{V_2}$, for each $i \in U$. We call D_s the coordinate convex extension of D. Thus, for every vector $\tilde{V} \in D_s$, there exists some vector $\tilde{A} \in D$ such that $0 \le c_i^V \le c_i^A$ for all $i \in U$. Clearly, $D_s \supseteq D$, hence D_s is coordinate convex.

D_s is closed, bounded, convex, nonempty if D has these properties. (It can be shown that coordinate convex extension of a closed, bounded, convex, nonempty set in R_+^N is also closed, bounded, convex, and nonempty.) Thus, by Theorem 2 D_s has an MMF. Let it be $\tilde{M}$. Clearly, $\tilde{M}$ is fairer than all vectors in $D \subseteq D_s$. Hence, $\tilde{M}$ is the MMF in D if it belongs to D. If $\tilde{M} \notin D$, $\tilde{M} \in D_s \setminus D$. Thus, there exists $\tilde{V} \in D$ such that

$$c_i^M \begin{cases} < c_i^V, & \text{for some } i = j \\ \le c_i^V, & i \ne j. \end{cases}$$

$\tilde{M}$ is not fairer than $\tilde{V} \in D \subseteq D_s$. This contradicts the fact that $\tilde{M}$ is the MMF in D_s. □

$F(m, r, \tilde{p}, n) = F \cap R(m)$, where F is the set of carried traffic intensity vectors at r, n (and traffic pattern $\tilde{p}$) and

$$R(m) \subset R^N = \left\{\tilde{V}: \sum_{i=1}^N x_i^V \ge m \text{ where } x_i^V \text{ is the } i\text{th component of } \tilde{V}\right\}$$

Clearly, $R(m)$ is closed and convex. F is also closed and convex. Thus, the intersection of these two is also closed and convex. Since F is bounded, $F(m, r, \tilde{p}, n) \subseteq F$ is also bounded. $m \le T_R(r, n)$ means $x^R(r)$ which fetches revenue equal to $T_R(r, n) \ge m$ is in set $R(m)$ and it is also there in F. Thus, the intersection is nonempty. $F(m, r, \tilde{p}, n) \subseteq F \subset R_+^N$. $F(m, r, \tilde{p}, n)$ is a nonempty closed, bounded, convex subset of R_+^N. Hence, $F(m, r, \tilde{p}, n)$ has an MMF by Theorem 8.

VI. Conclusion and Future Work

We have developed the notions of relative and absolute fairness. We have obtained a simple necessary and sufficient condition for relative fairness which is useful in determining if a vector is fairer than another. We have shown that any subset of R^N satisfying certain properties has an MMF and the set of carried traffic intensity vectors satisfies those properties. We have shown that the MMF is unique. We have been able to specify a fixed channel assignment algorithm which yields the MMF at all loads, in the asymptotic limit. Future research may be directed toward determining optimal channel assignment algorithms that would maximize revenue taking into account the effects of customer desertion.

We would like to point out that many of our results are very general in nature and apply to a large class of other networks as well. The problem of max-min fair bandwidth allocation is very relevant in context of other networks as well, e.g., automated teller machine (ATM) networks. The set of feasible allocations often satisfies the conditions of Theorem 8. Thus, we know from Theorem 8 that the max-min fair allocation exists uniquely (uniqueness follows from Theorem 5) and the construction of Theorem 2 actually yields this allocation if the set of feasible allocations is known. Our results may find application in the problem of max-min fair allocation of available bandwidth to competing available bit rate (ABR) traffic in ATM networks with minimum cell rate requirements (MCR) [1].

REFERENCES

[1] S. P. Abraham and A. Kumar, "Algorithms for max–min fair rate control of ABR connections with nonzero MCR in ATM networks," in *Proc. Conf. Communication Technologies*, 1996.

[2] C. Berge, *Hypergraphs*. Amsterdam, The Netherlands: North-Holland, 1989.

[3] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1987.

[4] E. Kreyszig, *Introductory Functional Analysis With Applications*. New York: Wiley, 1978.

[5] F. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, no. 3, pp. 237–252, Mar. 1998.

[6] V. H. Macdonald, "The cellular concept," *Bell Syst. Tech. J.*, vol. 58, no. 1, Jan. 1979.

[7] R. J. McEliece and K. N. Sivarajan, "Performance limits for channelized cellular telephone systems," *IEEE Trans. Inform. Theory*, vol. 40, pp. 21–34, Jan. 1994.

[8] ——, "Maximizing marginal revenue in generalized blocking service networks," in *Proc. 30th Annu. Allerton Conf. Communications, Control, and Computing*, vol. 18, 1992, pp. 455–464.

[9] K. N. Sivarajan, "Performance limits for cellular systems with handoffs," in *Proc. 31st Annu. Allerton Conf. Communications, Control, and Computing*, vol. 19, 1993, pp. 887–897.

[10] R. W. Wolff, *Stochastic Modeling and the Theory of Queues*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[11] S. Sarkar and K. N. Sivarajan, "Fairness in cellular mobile networks," Lab. Resource Allocation in Networks, Elec. Comput. Eng. Dept., Indian Institute of Science, Bangalore, Tech. Rep.LRAN-1, 1997.