A Mechanism for Equitable Bandwidth Allocation under QoS and Budget Constraints

Sreenivas Gollapudi
Oracle Corporation / SUNY at Buffalo
Email: sreenivas.gollapudi@oracle.com

D. Sivakumar
IBM Almaden Research Center
Email: siva@almaden.ibm.com

***Abstract*— Equitable bandwidth allocation is essential when QoS requirements and purchasing power vary among users. To this end, we present a mechanism for bandwidth allocation based on differential pricing. In our model, the QoS vs. cost trade-off induces a minimum acceptable allocation, a maximum acceptable allocation, and a unique optimal allocation for each user.**

We analyze the fairness and truthfulness properties of our mechanism from a game-theoretic perspective. We show that it produces allocations that provably satisfy a variant of the classical notion of *max-min fairness*. It ensures that flows with higher QoS requirements need to pay at higher rates to increase their likelihood of being served. Furthermore, the *Nash equilibrium* induced by our mechanism leads to allocations that are comparable to "socially optimal" allocations; hence users gain very little by being untruthful.

I. Introduction

Fair and efficient bandwidth management is an essential component of a QoS infrastructure. Such an infrastructure is needed to provide guaranteed services within user-specified bounds on, among other things, minimum bandwidth requirement. A good method for bandwidth allocation should satisfy several, often conflicting, desiderata. For example, it should lead to alleviation of congestion problems by using a fair allocation policy while ensuring the users exhibit truthful behavior from the point of view making requests for bandwidth.

Dividing a commodity (bandwidth, in our case) among several agents (flows) is a basic problem studied in game theory, especially from the viewpoints of fairness (ensuring equitability of allocations) and truthfulness (strategyproofness—ensuring that the agents report their preferences truthfully). In these scenarios, agents first report their preferences in the form of *utility functions* that indicate how much value they derive for various allocations. In the context of bandwidth allocation, the utility is a measure of how valuable a specific allocation is to the agent, and is usually a function of several parameters internal to the agent, such as the agent's QoS requirements,budget constraints, etc. The allocation mechanism then determines, for each agent, how much of the commodity they receive. The most general notion of interest with respect to a mechanism is the *Nash equilibrium* it induces. In our scenario, this corresponds to the set of reported utility functions that no selfish and rational agent has any incentive to deviate from.

A very common and well-studied class of preference functions is that of single-peaked preferences where each agent has a unique allocation value that would maximize her utility. In particular, bandwidth requirements usually have this property. For this class of preferences, Sprumont [1] provided an elegant characterization of allocation rules that satisfy various natural desirable properties:

Feasibility: The sum of allocations to various agents equals the total availability.

Efficiency: If the sum of ideal allocations exceeds the supply, then no agent should receive more than her ideal allocation; if supply exceeds the sum of ideal allocations, then no agent should receive less than her ideal allocation.

Strategyproofness: There is no incentive for any agent to mis-report her preferences.

Envy-freeness: When the allocations have been determined, no agent would prefer another agent's allocation to hers; this can be viewed as a form of fairness in allocation.

Sprumont's characterization showed that the only allocation rule that satisfies all of the above conditions is the *uniform allocation rule*, which can be summarized as follows. Suppose that the sum of ideal allocations exceeds supply. Initially allocate an equal share to every agent; if the ideal allocation for an agent is less than this equal share, that agent will receive exactly her ideal allocation. This creates some excess, which will be equally re-apportioned among all remaining agents. This process is repeated until all of the commodity has been allocated. Subsequent to Sprumont's work, other authors [2], [3], [4], [5] have provided alternate characterizations of the uniform rule, and/or extended it to incorporate various other criteria. In all of these, however, one axiom remains unchanged, namely that of feasibility.

In this paper, we study the bandwidth allocation problem in contexts where each agent has, in addition to a single-peaked preference function, minimum and maximum values of allocation that are acceptable. Before we motivate this scenario further, we note that the presence of minimum/maximum acceptable values limits the ability to produce feasible allocations. In particular, the uniform rule does not apply anymore; this reopens the problem of producing efficient, strategyproof, and fair allocations that respect the agents' maximum/minimum requirements. A motivating factor for the introduction of min/max constraints in the preference functions is to model division problems where the agents have to pay for their allotment according to a *cost* function. In some scenarios, for every agent there is an interval of acceptable allocations outside of which cost dominates utility.

Our main result in this paper is a bandwidth allocation mechanism that has desirable properties of fairness and strategyproofness. In addition, our mechanism satisfies at least one of two additional goals—it either produces a feasible allocation, or achieves ideal allocation for all agents with positive allocation (often called welfare maximization). Most importantly, we explicitly characterize the Nash equilibrium induced by our mechanism, and show that it leads to very desirable allocations both in terms of maximizing the number of agents served, and in terms of maximizing allocations for the "socially responsible" agents.

Before we describe in detail the notions of fairness and truthfulness, we briefly sketch the structure of our mechanism. Our mechanism takes place in rounds. Each round consists of an *admission* phase, where the remaining agents are made a minimum offer of bandwidth, which they choose to accept or reject. Rejecting agents are deferred to the next round, in which the minimum offer, if any, will be strictly larger. Accepting agents submit their utility functions, and enter the *allocation* phase, during which their allocation will be completely determined. The mechanism stops when the bandwidth left over is inadequate for any remaining agent, or when all agents have received maximum welfare (utility minus cost).

Fairness. One of the most well-studied notions of fairness in the context of allocation problems is that of *max-min fairness* [6], [7], [8], [9]. Informally, a max-min fair allocation attempts to maximize the minimum allocation. In our setting, max-min fairness could be provably impossible. For example, given the resource and welfare constraints, some agents might be forced to receive zero allocation, which violates max-min fairness. Also, it might be desirable to modify the notion of max-min fairness to incorporate additional elements of equitability. Our notion of fairness incorporates into max-min fairness information about the degree to which various agents are "socially responsible." (Formal definitions appear in Section II.)

Truthfulness. Various notions of truthfulness have been studied in game theory, ranging from the dominance of truthful strategies when the participants care about worst-case scenarios (e.g., prisoner's dilemma, assuming the prisoners indeed committed the crime together) to strategyproofness and group strategyproofness (where no coalition of agents can improve their welfare by mis-reporting their utility functions). While in the case of prisoner's dilemma, the participants have no information about the strategies of the other participants, strategyproofness or group strategyproofness ensure no user has any incentive to behave untruthfully *even if* she (or a coalition) has full information about the other participants' utility functions. Group strategyproofness is one of the strongest forms of truthfulness (see, e.g., [10], [11], and references therein).

For our mechanism, we combine both types of truthfulness arguments. In particular, we show a weak form of strategyproofness for the admission phase (where the answers are Boolean); assuming that the admission phase took place truthfully, we show group strategyproofness for the allocation phase (where agents report their utility functions).

Nash equilibrium. In addition to the modular analysis in terms of fairness and truthfulness, we also provide a complete characterization of the Nash equilibrium induced by our mechanism. The classical result of Nash [12] shows that every game has a "mixed-strategy equilibrium," from which no participant has any incentive to deviate. A mixed-strategy equilibrium is one where a participant has a probability distribution over her set of "pure" strategies. (Formal definitions appear in Section II.) In many games of interest, pure-strategy equilibria are quite rare. We prove the surprising result that our mechanism leads to a pure-strategy equilibrium; moreover, we explicitly derive this equilibrium and relate it to a measure of how socially responsible various agents are.

Differential pricing. Differential pricing enables a network to support the agents' ability to pay at different price levels and receive QoS commensurate with the price level. Thus, an agent can increase her chances of receiving a certain QoS by paying more.

Besides fairness and truthfulness, we are able to integrate a differential pricing policy (e.g., [14], [15], [16], [17], [18]) into our allocation mechanism. The form of differential pricing that we incorporate is similar to the Paris Metro Pricing (PMP) policy which has been proposed by Odlyzko [13] in the context of congestion control. We accomplish this by the introduction of various *price levels* at which agents will request bandwidth. This feature is also quite germane to the fairness/sovereignty issues (as well as feasibility): when the agents have varying abilities to pay, richer agents should, at the very least, be allowed to pay more to obtain the same service. In addition, multiple service levels with differential pricing also addresses the possible criticism that our basic mechanism favors agents who make smaller bandwidth requests.

Organization. In Section II, we describe the basic setup, outline the model used for the agents' utility functions, and formally define the notions of fairness, truthfulness, etc. In Section III, we describe a basic mechanism $\mathcal{M}$ for sharing bandwidth on a single link at different price levels, and summarize the fairness and truthfulness properties of our mechanism. We then extend the mechanism to the case of a full network. Section IV presents some simulation results that illustrate various characteristics of our mechanism, including an analysis of the resulting Nash equilibrium and comparison with a "socially optimal" allocation mechanism.

II. Models and Preliminaries

We use the term "vendor" to denote the seller of bandwidth and "agents" to denote the consumers of bandwidth, usually a single point-to-point flow. The kind of mechanisms we develop in this paper are suitable for deployment at various routing points in the Internet traffic. For simplicity, one may think of a router that controls allocation of bandwidth on a single link; this abstraction is often enough to consider more general point-to-point multiple-link connections in the Internet. We also define a *network coordinator* that handles the functions of admission control of a flow along path P based on capacity and utilization information of all links along P. The network coordinator is used to deploy the bandwidth allocation mechanism in a general network.

Let us consider the case of a single link along path P. Consider n agents requesting bandwidth on a link with bandwidth capacity is B. Let C denote the revenue goal for the vendor of bandwidth.

The *utility*, as a function of bandwidth, for agent i is denoted by $u_i(b) : [0, B] \longrightarrow \mathbf{R}$. We assume that the u_i's are continuously differentiable, monotonically non-decreasing, and satisfy the following property—for each i, there is a value $\beta_i \in [0, B]$ such that $u_i(b)$ is convex in the interval $b \in [0, \beta_i)$ and concave in the interval $b \in (\beta_i, B]$. We find that multimedia traffic tend to follow a QoS profile similar to the utility function illustrated in Figure 1 [19].

Let us assume the vendor's *cost* function, $c(b)$, is given by the linear function $c(b) = \lambda b$ where $\lambda B \geq C$). This definition of the cost function is easily extended to support differential pricing corresponding to different levels of QoS an agent can expect. Let ℓ_i denote the price level at which agent i decides to request bandwidth from the network. To reiterate, an agent can choose to pay more for a QoS requirement to increase the likelihood of receiving the service. Then, the cost function for agent i can be defined as $c_i(l, b) = \lambda \ell_i b$, where $\lambda \sum_i \ell_i B \geq C$. Note that the idea of using pricing as a method to support QoS in a fair and effective manner also appears in earlier work [20]. In [20], a network broker determines a price based on the competition for bandwidth, and subject to this price, agents make their requests satisfying their budget and QoS constraints. Our approach is dual to this, where we assume a fixed pricing method and provide a mechanism that determines the allocation for each agent.

The *welfare* of an agent i, as a function of bandwidth, is defined as $w_i(b) = u_i(b) - c_i(\ell_i, b)$. (In some literature, "utility" is called "value," and "welfare" is denoted by "utility.") A linear cost function typically intersects the utility function at two points β_i^0 and β_i^1, where the welfare for user i equals zero. In the interval $[0, \beta_i^0)$, the welfare is negative; this corresponds to the situation where the bandwidth is too low to be of use to meet the QoS requirements for the flow. In the interval $(\beta_i^1, B]$, the welfare is once again negative; this corresponds to the situation where the bandwidth is useful but not affordable. In the interval (β_i^0, β_i^1), the welfare is positive and reaches a maximum at some value β_i^*.

At the outset, we let the vendor publish his cost function. During the course of the mechanism, each agent will report her utility and the price level to the vendor. Based on the utility functions, price level and the cost function, the vendor apportions the bandwidth among the agents.

An *allocation mechanism* is a mapping from a vector of n utility functions $u_1, \ldots, u_n$, to a vector of n (allocation, cost) pairs $(a_1, c_1), \ldots, (a_n, c_n)$ subject to the constraints $\sum_i a_i \leq B$ (bandwidth availability), and for $1 \leq i \leq n$, $u(a_i) - c_i \geq 0$ (non-negative welfare for users).

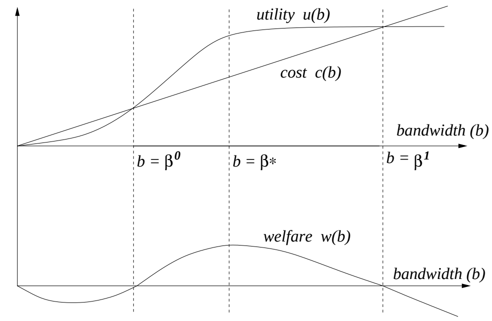

Fig. 1. A typical utility function

We say that an allocation mechanism satisfies *feasibility* if the total allocated bandwidth is B; note that, since $\lambda B \geq C$, this implies that the total cost the vendor charges the agents is at least C.

A. Fairness

In the context of sharing bandwidth, the question of what constitutes a fair allocation is not obvious. One allocation could be the vector $(a_1, a_2, \cdots, a_n)$ such that $\forall i, 0 < i \leq n$, $w_i(a_i) \geq 0$ and $\sum_{0<i\leq n} a_i \leq B$. Clearly, if $a_i = \beta_i^*, 0 < i \leq n$, we have an allocation that maximizes welfare of all agents. In case $\sum_{0<i\leq n} a_i < B$ where the vendor will be left with some excess bandwidth, it is important to find if any notion of fairness should also specify a fair way to do a feasible re-apportioning, leading to the main question: which, if any, of the agents are "more deserving" of the excess bandwidth than the others? Our notion of fairness is based on the classical max-min fairness.

Definition 1 (Max-min fairness): Let Z denote the set of agents, and for $i \in Z$, let u_i denote the utility function of agent i. Let a denote a vector of allocations and let c denote a vector of costs, both indexed by $i \in Z$, satisfying $\sum_{i \in Z} a_i \leq B$. Allocation a with cost c is said to be *max-min fair* if, for every allocation a' with cost c', the following condition holds. If, for some $i \in Z$, $a'_i > a_i$, then one of the following must be true:

(1) (capacity violation) $\sum_i a'_i > B$;

(2) (welfare reduction for agent i) $u_i(a'_i) - c'_i < u_i(a_i) - c_i$;

(3) (max-min violation) for some $j \in Z, j \neq i$, such that $a_j \leq a_i$, we have $a'_j < a_j$.

The more natural way to think about this is to assume that conditions (1) and (2) are not violated, that is, (a', c') is a feasible allocation that doesn't hurt agent i; then condition (3) says that some other agent who received no more than agent i under allocation (a, c) now receives less under allocation (a', c').

Remark 1: A more general way to state condition (3) in the definition of max-min fairness is:

(3′) for some $j \in Z, j \neq i$, such that $a_j \leq a_i$, we have $u_j(a'_j) - c'_j < u_j(a_j) - c_j$.

However, under a fairly natural assumption about the allocations, the two are equivalent. Namely, we restrict ourselves to allocations that are *stable*; an allocation is stable if there do not exist two agents $i \neq j$ such that simultaneously reducing a_i and increasing a_j by the same amount will improve the welfare of both agents i and j. Stability, together with our model of agent utility (which, in its positive region, is first strictly increasing, then strictly decreasing), implies that conditions (3) and $(3')$ are equivalent.

For the allocation problem of the type we consider, max-min fairness is provably impossible. For example, given the constraints on resources and the utility functions of the agents, it might be inevitable that some agent receives zero allocation. For example, there are 3 agents, and the only acceptable allocation for all three agents is $1/2$ of available bandwidth. When this happens, clearly no mechanism can be max-min fair. To accommodate the possibility of different agents receiving different amounts of bandwidth, we make an extension to max-min fairness that is biased in favor of the agents who are more inclined to share (that is, agents who are socially responsible and not greedy).

Definition 2 (Max-min fairness with respect to g): Let Z denote the set of agents, and for $i \in Z$, let u_i denote the utility function of agent i. Let a denote a vector of allocations and let c denote a vector of costs, both indexed by $i \in Z$, satisfying $\sum_{i \in Z} a_i \leq B$. Let g denote an auxiliary vector, also indexed by $i \in Z$ (think of g as the "greed" vector). Allocation a with cost c is said to be *max-min fair with respect to* g if, for every allocation a' with cost c', the following condition holds. If, for some $i \in Z$, $a'_i > a_i$, then one of the following must be true:

(1) (capacity violation) $\sum_i a'_i > B$;

(2) (welfare reduction for agent i) $u_i(a'_i) - c'_i < u_i(a_i) - c_i$;

(3) (max-min violation) for some $j \in Z, j \neq i$, such that $g_j \leq g_i$, we have $a'_j < a_j$ (more generally, under the assumption of stability, $u_j(a'_j) - c'_j < u_j(a_j) - c_j$).

Note that conditions (1) and (2) are just to ensure non-triviality of the allocation a'. Also, we may obtain the classical definition of max-min fairness by taking $g = a$. Informally, our definition says that if (a', c') is a feasible allocation that doesn't hurt agent i, then it allocates strictly less to some agent j who is no greedier than agent i.

B. Truthfulness

A typical proof of truthfulness of participants in a cooperative game depends on modeling the participants as selfish and rational. Namely, each agent cares most about her own welfare; even if an agent decides to participate in a coalition that deviates from truthful reporting of utilities, the ultimate aim of each agent is to improve (or at least preserve) their own welfare. The following is a formal abstraction of this behavior.

For a vector of utility functions u, let $a_M(u)$ denote the allocation vector and $c_M(u)$ denote the cost vector determined by a mechanism M when the utility functions reported are u.

Definition 3 (Manipulation property): This property states that no agent, either individually or as part of a coalition, will not mis-report her utility if there is risk of reducing her welfare by doing so.

Fix a mechanism M, and a set of agents Z. Let u denote the vector of true utility functions of the agents, let $a = a_M(u)$ and $c = c_M(u)$. Then, for every vector u' of reported utilities, if $a' = a_M(u')$ and $c' = c_M(u')$, and for all potential coalition of agents, $T \subseteq Z$, we have:

$$\big((\forall j \notin T)[u'_j = u_j]\big) \implies (\forall i \in T)[u'_i \neq u_i \implies u_i(a'_i) - c'_i \geq u_i(a_i) - c_i]$$

The hypothesis says that the agents not in T behave truthfully, and the consequence says that every agent in T that mis-reports her utility only does so if her welfare is no worse than under truthful reporting of all utilities.

Definition 4 (Group strategyproofness): A mechanism is said to be *group strategyproof* if, under the assumption of the manipulation property, every agent reports her utility truthfully.

Definition 5 (Weak strategyproofness): A mechanism M is said to be *weakly strategyproof* if, for all i, the true utility u_i of agent i maximizes the minimum of $u_i(a_i) - c_i$ over all possible reported utilities u_j for $j \neq i$, where $a = a_M(u)$ and $c = c_M(u)$.

C. Nash equilibrium

We define the notion of Nash equilibria in the context of our bandwidth allocation problem. In a general game with n agents, each agent has a set of *strategies* from which she chooses her "move." A *payoff function* takes the vector of moves of all agents and maps it into a vector of benefits for the agents. An agent has the option of employing a *pure strategy*, that is, a fixed move from her set of strategies, or a *mixed strategy*, which is a probability distribution over her set of strategies. In the latter case, the moves of the agents, and consequently the payoffs, are random variables. For our scenario of bandwidth allocation, strategies correspond to the possible utility functions that the agents report to the mechanism.

Definition 6 (Nash equilibrium): Fix a mechanism M for allocation, and a cost function c. Let Z denote the set of agents, and for $i \in Z$, let u_i denote the true utility function of agent i. Let U_i denote a random variable corresponding to the reported utility of agent i, and let D_i denote the distribution of U_i. The distributions $U_1, \ldots, U_n$ are said to be in equilibrium if for every i and every distribution F_i for U_i,

$$\mathop{\mathrm{E}}_{\substack{U_i \sim D_i \\ U_j \sim D_j, j \neq i}} [A_i(U_i) - c(A_i)] \;\geq\; \mathop{\mathrm{E}}_{\substack{U_i \sim F_i \\ U_j \sim D_j, j \neq i}} [A_i(U_i) - c(A_i)], \tag{1}$$

where $A = a_M(U)$ denotes the allocation by mechanism M when the utility reported is U.

The left hand side of Equation (1) is the expected welfare obtained by agent i when all the agents, including i, use the mixed strategy corresponding to the distributions D; the right hand side of Equation (1) is the expected welfare obtained by agent i when agent i uses the mixed strategy F_i for U_i and the other agents j, $j \neq i$, use the mixed strategy given by the distributions D_j. Thus if all agents use the strategies prescribed by the distributions $D_j, j \in Z$, then none of them has any incentive to adopt a different strategy. The classical theorem of John Nash [12] implies, in particular, that a Nash equilibrium always exists, and in general, might consist of mixed strategies for the agents.

III. A MECHANISM $\mathcal{M}$ FOR BANDWIDTH ALLOCATION

In this section, we present our bandwidth allocation mechanism (shown in Figure 2). Note that for each agent i, the welfare function $w_i(b)$ is monotonic non-decreasing for $b \in (\beta_i^0, \beta_i^*)$ and monotonic non-increasing for $b \in (\beta_i^*, \beta_i^1)$. This implies that mechanism $\mathcal{M}$ may be implemented so that each agent first responds with a series of yes/no answers, followed by reporting four numbers—β_i^0, β_i^*, β_i^1 and price level ℓ_i. Our iterative formulation can be viewed as an extension of Sönmez's [3] characterization of the uniform rule.

The mechanism $\mathcal{M}$ proceeds in rounds with each round comprising two phases — an *admission phase* and an *allocation phase*. In the admission phase of a round, say j, the network computes an initial allocation for each agent who still needs bandwidth. Let T_j denote the set of all agents needing bandwidth in round j. An agent i whose initial allotment is outside the interval $[\beta_i^0, \beta_i^1]$, turns down the allocation (line 4). The resulting "excess" bandwidth is then partitioned among the remaining agents in T_j in the *allocation phase* of round j. Let us denote this set of agents by S_j.

The initial allocation μ_j in the admission phase of round j is computed by finding the largest integer t such that, $\mu_j = B_j/t > \mu_{j-1}$ (so that in each round, the minimum allocation, if any, is strictly more than the previous round). The mechanism treats this case as if there were t "virtual agents" in round j and agent $i (\in T_j)$ at level ℓ_i is offered $\ell_i \mu_j$.

In the allocation phase, the mechanism uniformly reapportions the remaining bandwidth among all agents in S_j. The increment in bandwidth to each agent $i \in S_j$ is computed as $b_i = min(\beta_i^*, \ell_i B_j / L'_j)$, where $L'_j = \sum_{i \in S_j} \ell_i$ is the total number of "virtual" agents participating in the admission phase of round j (lines 9–12). If the number of virtual agents are too many such that $L'_j \mu_j > B_j$, then we randomly pick an agent in S_j and drop her from round j and repeat this procedure (line 7) till we satisfy $L'_j \mu_j \leq B_j$. Choosing the excess agents randomly ensures that no (coalition of) agents can falsely report their utilities to get bandwidth in round j. Any remaining bandwidth from round j, namely $B_j - \sum_{i \in S_j} b_i$, is carried over to the round $j+1$ (line 13) and the agent list for the round $j+1$ is initialized with all agents who chose not to participate in round j (line 14).

To summarize, agents with similar utilities at different levels of pricing receive allocations whose ratio is roughly the ratio of the levels.

A. Fairness, feasibility, and welfare maximization

We will now summarize the properties of mechanism $\mathcal{M}$ from the point of view of achieving feasibility, welfare maximization, and fairness. The proof of Theorem 1 is straightforward, and follows from the details of mechanism $\mathcal{M}$.

Mechanism $\mathcal{M}$:

(0) Initially, the vendor announces the cost function parameter λ.
Let $T_1 = Z$, the set of agents who wish to obtain bandwidth.
Let $B_1 = B$, the total bandwidth available.
Let $j = 1$.

(1) **while** $T_j \neq \emptyset$ AND $B_j \neq 0$ **do**: (while agents left and bandwidth available)

(2) Default allocation: For $i \in T_j$, let $b_i = 0$.

Admission phase:

(3) The vendor computes the value μ_j as follows, and announces μ_j to the agents:

$$\mu_j = \begin{cases} B_1/L & \text{if } j = 1, \text{ where } L = \sum_{i \in T_j} \ell_i \\ B_j/t & \text{if } j > 1, \text{ where } t \text{ is the largest integer such that } B_j/t > \mu_{j-1} \end{cases}$$

(Note: for $j > 1$, we ensure that $\mu_j > \mu_{j-1}$.)

(4) Each agent then sends a bit to the vendor, indicating their decision to participate in round j of the allocation.

Allocation phase:

(5) Let S_j be the set of agents who decide to participate in round j, and let $L'_j = \sum_{i \in S_j} \ell_i$.

(6) Users in S_j then report their utility functions $u_i : [0, B] \longrightarrow \mathbf{R}^+$, $i \in S_j$, subject to the constraint $\beta_i^0 \leq \mu_j$ (β_i^0, β_i^1 and β_i^* are as defined above).

(7) **if** $j > 1$ and $L'_j > t$ (thus $L'_j \mu_j > B_j$) then:
choose randomly agents from S_j and delete them from S_j till $L'_j \mu_j \leq B_j$;
for each agent i remaining in S_j, allocate (at most) $\mu_j \ell_i$ units of bandwidth and **stop**.
end if

(8) Initial allocation: For $i \in S_j$, let $b_i = \mu_j \ell_i$.

(9) If, for some i, $b_i \geq \beta_i^*$, then set $b_i = \beta_i^*$, and delete i from S_j.

(10) **while** $S_j \neq \emptyset$ **do**:

(11) Allocate more bandwidth:
Increase b_i uniformly for all $i \in S_j$ until one of the following happens:
(case 1) $\sum_{i \in T_j} b_i = B_j$, or
(case 2) for some $i \in S_j$, we have $b_i = \beta_i^*$.

(12) Case 1 (exhausted bandwidth):
If $\sum_{i \in T_j} b_i = B_j$, then **stop**.
Case 2 (agent i has reached maximum welfare):
If $b_i = \beta_i^*$, then delete i from S_j, and continue.

end while

(13) $B_{j+1} = B_j - \sum_{i \in T_j} b_i$.

(14) $T_{j+1} = \{i \in T_j \mid b_i = 0\}$ (agents in T_j who did not participate in round j).

(15) $j = j + 1$.

end while

Fig. 2. Mechanism for bandwidth allocation

Let Z denote the set of all agents, and say $\mathcal{M}$ proceeded for $r \geq 1$ rounds. For $1 \leq j \leq r$, let Z_j denote the set of agents that actively participate in round j of the mechanism $\mathcal{M}$, that is, these are the agents that entered the set S_j in step (5) of the algorithm. For convenience, we allow for the possibility that in the r-th round, step (3) sets $t = 0$, and assume that when this happens, every agent joins S_j; note that in this case, step (7) will ensure that none of these agents are allocated any bandwidth, and the mechanism stops. Thus, Z is partitioned into $Z_1, \ldots, Z_r$.

Users in Z_r can be further partitioned into three sets: Z_r^*, consisting of agents whose welfare is maximized; Z_r^+, consisting of agents who receive positive allocation (of at least μ_r) but whose welfare is not maximized; and Z_r^0, consisting of agents who receive zero bandwidth. (Note that some of these may be empty; also note that $Z_r^0 = \emptyset$ iff either $Z_r^* \neq \emptyset$ or $Z_r^+ \neq \emptyset$.)

Theorem 1 (Fairness of $\mathcal{M}$): Let Z^* denote the union of Z_j, for $j < r$, and Z_r^*. Let $Z^+ = Z_r^+$, and let $Z^0 = Z_r^0$. Define the vector g, indexed by $i \in Z$, by $g_i = \mu_j$, if $i \in Z_j$ (this is the smallest offer μ_j of bandwidth such that $\mu_j \ell_i \geq \beta_i^0$, so i is admitted in round j). The mechanism $\mathcal{M}$ satisfies the following properties.

(Case 1 — bandwidth exhausted) $\mathcal{M}$ terminates with $B_{r+1} = 0$. In this case, Z^* and Z^+ are possibly non-empty, Z_r^0 is empty. Furthermore, (1a) the mechanism achieves feasibility; (1b) agents in Z^* have maximum welfare; (1c) all agents in $Z^+ = Z_r^+$ have the same allocation, and this allocation is at least as much as any agent in Z_r^*; (1d) the allocation is max-min fair with respect to g.

(Case 2 — bandwidth not exhausted) $\mathcal{M}$ terminates with $B_r > 0$ in step (7) with $t = 0$. In this case, Z^+ is empty, Z^* and $Z^0 = Z_r^0$ are possibly non-empty. Furthermore, (2a) the mechanism does not achieve feasibility, but any bandwidth left is not sufficient for any of the remaining agents in Z^0; (2b) agents in Z^* have maximum welfare; (2c) Z^+ is empty; (2d) the allocation is max-min fair with respect to g.

(3) In addition, for every j, $1 \leq j \leq r$, the allocation produced by mechanism $\mathcal{M}$, restricted to the agents in Z_j, is max-min fair.

(4) Allocations produced by mechanism $\mathcal{M}$ are efficient.

B. Truthfulness of $\mathcal{M}$

Recall that in the course of the execution of $\mathcal{M}$ on a set Z of agents, the only steps in which the agents send any information to the mechanism are steps (4) and (6). In step (4) of the admission phase, they indicate their willingness to participate in some round j of the protocol, and in step (6), the beginning of the allocation phase, they submit their utility functions.

Theorem 2: Consider any execution of $\mathcal{M}$ on a set Z of agents.

(1) *(Truthfulness in admission phase)* The admission phase of mechanism $\mathcal{M}$ is weakly strategyproof.

(2) *(Truthfulness in allocation phase)* Assuming that all the admission phase queries were answered truthfully, the allocation phase of mechanism $\mathcal{M}$ is group strategyproof.

The proof of this theorem for the case $\ell_i = 1$, $i \in S_j$, is presented in the Appendix.

C. Nash equilibrium of $\mathcal{M}$

In this Section, we explicitly derive a pure-strategy Nash equilibrium of the allocation game induced by mechanism $\mathcal{M}$. Recall that in mechanism $\mathcal{M}$, each agent reports three values representing her utility—β^0, β^*, and β^1—and her price level ℓ. Since β^1 is not used by mechanism $\mathcal{M}$, and since the level parameter ℓ is simply a scale factor, the set of possible strategies available to agent i involve only β^0 and β^*.

We begin with some observations that lead to the derivation of the equilibrium.

(Observation 1) There is no incentive for any agent to mis-report her β^* value. If she reports a smaller value β', then mechanism $\mathcal{M}$ will limit her allocation to β', which has less welfare; similarly, if she reports a value larger than β^*, she incurs negative welfare. (Note that β^* is reported only in the allocation phase, where $\mathcal{M}$ has been shown to be group strategyproof (cf. Theorem 2).)

For agent j, let $\tilde{\beta}_j$ denote the reported value of the minimum acceptable allocation (recall that β_j^0 is the true minimum acceptable allocation).

(Observation 2) There is no incentive to have $\tilde{\beta}_j > \beta_j^0$, for doing so would only delay the entry of agent j into the allocation phase of $\mathcal{M}$, and any possible allocation derived from reporting $\tilde{\beta}_j$ could also have been derived by reporting β_j^0. (Once again, the proof of Theorem 2 presents a more rigorous version of this observation.)

Thus, the only possible "strategizing" that agent j can possibly do is to employ a mixed strategy, that is, a probability distribution for $\tilde{\beta}_j$ in the interval $[0, \beta_j^0]$. We show next that there is a Nash equilibrium where agent j picks either $\tilde{\beta}_j = 0$ or $\tilde{\beta}_j = \beta_j^0$ with probability 1.

Definition 7 (Fairness rank): For an agent $j \in Z$, define the *fairness value* of j to be $\varphi(j) = \beta_j^0/\ell_j$, and the *fairness rank* of j to be $rank(j) = |\{i \in Z \mid \varphi(i) \leq \varphi(j)\}|$, that is, the fairness rank of j is the rank of j in the list of agents sorted by their fairness value.

The next definition is crucial in the pure strategy Nash equilibrium that we define.

Definition 8 (Domination index): For an agent $j \in Z$, define the *domination index* of j to be

$$\Delta(j) = \min \left\{ d \,\middle|\, \left(\sum_{i:rank(i) \leq d} \min(\beta_i^*, \frac{\ell_i}{\ell_j}\beta_j^0) \right) > B \right\}.$$

Informally, the domination index identifies the set of agents with lowest fairness values that dominate agent j, that is, the set S of agents such that if all agents in $i \in S$ report $\tilde{\beta}_i = 0$, then agent j has no possibility of obtaining an allocation of at least β_j^0 from mechanism $\mathcal{M}$, *regardless* of what $\tilde{\beta}_j$ is.

Lemma 1 (Crossover lemma): Let $j \in Z$ denote an agent, and let $r = \Delta(j) \leq rank(j) = s$. Every $i \in Z$ such that $rank(i) < r$ must satisfy $\Delta(i) \geq r$. (By symmetry, if $rank(i) > r = \Delta(j) \geq rank(j) = s$, then $\Delta(i) \leq r$.)

Proof: Assume to the contrary that $d = \Delta(i) < r$. Since $rank(i) < r \leq s = rank(j)$, we have $\varphi(i) \leq \varphi(j)$. Then

$$\begin{aligned} B &< \sum_{k:rank(k) \leq d} \min(\beta_k^*, \frac{\ell_k}{\ell_i}\beta_i^0) && \text{(By defn. of } \Delta(i)) \\ &\leq \sum_{k:rank(k) \leq d} \min(\beta_k^*, \frac{\ell_k}{\ell_j}\beta_j^0) && \text{since } \varphi(i) = \frac{\beta_i^0}{\ell_i} \leq \frac{\beta_j^0}{\ell_j}, \end{aligned}$$

which implies that $\Delta(j) \leq d < r$, contradicting the minimality of r as the choice for $\Delta(j)$. ■

Theorem 3 (Nash equilibrium of $\mathcal{M}$): Define, for $j \in Z$, $\tilde{\beta}_j = 0$ if $\Delta(j) > rank(j)$, and $\tilde{\beta}_j = \beta_j^0$ if $\Delta(j) \leq rank(j)$. The collection of reported utility functions where the minimum acceptable allocations are reported as $\{\tilde{\beta}_j\}_{j \in Z}$ is a pure-strategy Nash equilibrium of $\mathcal{M}$.

Proof:

(Case 1) Suppose $j \in Z$ is such that $r = \Delta(j) \leq rank(j) = s$. Thus $\tilde{\beta}_j = \beta_j^0$. Suppose agent j reports a minimum acceptable allocation of $b < \beta_j^0$. By the crossover lemma, we know that for every i such that $rank(i) < r$, we have $\Delta(i) \geq r$, which implies that $rank(i) < \Delta(i)$, and hence $\tilde{\beta}_i = 0$. Additionally, since $\Delta(j) = r$, we know that $B < \sum_{i:rank(i) \leq r} \min(\beta_i^*, (\ell_i/\ell_j)\beta_j^0$. This implies that when all agents (except j) report their utilities according to $\tilde{\beta}$ and j reports b, mechanism $\mathcal{M}$ would have admitted all i's such that $rank(i) < r$ in the first round. The earliest that agent j could enter the allocation phase is in the first round, by setting $b = 0$. In any event, if $\mathcal{M}$ gets to a round R such that $\mu_R \geq \beta_j^0/\ell_j$ (which would truly satisfy agent j's min. requirement), it would have allocated at least $\min(\beta_i^*, \ell_i \mu_R) \geq \min(\beta_i^*, (\ell_i/\ell_j)\beta_j^0)$ to every agent i such that $rank(i) < r$ (because i enters in round 1). On the other hand, the sum, over all i such that $rank(i) < r$ of this quantity, plus β_j^0 exceeds B, so it would be impossible to allocate at least β_j^0 to agent j. Reporting $\tilde{\beta}_j = \beta_j^0$ as the min. acceptable allocation has the same effect, that is, no allocation for agent j. Therefore, there is no incentive for agent j to report any minimum acceptable allocation other than β_j^0.

(Case 2) Now suppose $j \in Z$ is such that $r = \Delta(j) > rank(j) = s$. Thus $\tilde{\beta}_j = 0$. Suppose agent j reports a minimum acceptable allocation of b such that $0 < b \leq \beta_j^0$. By the crossover lemma, we know that for every i such that $rank(i) \geq r$, we have $\Delta(i) \leq r$, which implies that $rank(i) \leq \Delta(i)$, and hence $\tilde{\beta}_i = \beta_i^0$. Additionally, since $\Delta(j) = r > s = rank(j)$, we know that $B \geq \sum_{i:rank(i) \leq s} \min(\beta_i^*, (\ell_i/\ell_j)\beta_j^0$. In particular, if agent j reports her min. acceptable allocation is 0, then she is assured of receiving at least β_j^0, when all the other agents i report $\tilde{\beta}_i$. On the other hand, if she reports the value b s.t. $b > 0$, her allocation can be no more than what she obtains by reporting 0 (since reporting $b > 0$ might admit her in a later round), and could be possibly less than β_j^0 (in which case she has negative welfare). Thus $\tilde{\beta}_j = 0$ dominates any other possible value for min. acceptable allocation. ■

D. The general network case

We described mechanism $\mathcal{M}$ to allocate bandwidth on a single link. Now we consider the more general case where we have a graph $G = (V, E)$, a capacity function $B : E \longrightarrow \mathbf{R}^+$, and where agent i holds a path P_i in G and a utility function $u_i : \mathbf{R}^+ \longrightarrow \mathbf{R}^+$. The utility function u_i denotes, for each b, the value derived by agent i if she is allocated bandwidth b along the entire path P_i. We will sketch an extension of $\mathcal{M}$, denoted by $\mathcal{M}_g$, to allocate bandwidth in this case. Mechanism $\mathcal{M}_g$ runs in the *network coordinator*.

Mechanism $\mathcal{M}_g$ builds on the ideas of mechanism $\mathcal{M}$, and can be thought of as the simultaneous execution of $|E|$ mechanisms, one for each edge $e \in E$, subject to the constraints imposed by paths. (For example, an edge cannot increase the allocation of an agent that would violate the capacity constraint on another edge along that agent's path.)

In the admission phase, each edge $e \in E$, independently computes a quantity μ_j^e as $\mu_j^e = B_j^e/L^e$, where B_j^e is the bandwidth available along e in round j and $L^e = \sum_i \ell_i,$ *s.t.* $e \in P_i$ The network coordinator collects the values of minimum allocation from all edges in the network and offers, to an agent i, a minimum allocation defined by $\mu_j = \min_{e \in P_i} \mu_j^e$. Users who accept the minimum offer turn in their utility functions, and enter the allocation phase. The mechanism then increases every agent's allocation uniformly until one of the following conditions occurs: (1) some agent has reached maximum welfare, in which case she is deleted from the set S_j (where j denotes the round); or (2) some edge e is tight, that is, exhausts its bandwidth. In the latter case, all agents whose paths include e are deleted from the set S_j (and their allocations frozen at the current values). We state the following theorem without proof; the proof is similar to the proofs for $\mathcal{M}$.

Theorem 4: Mechanism $\mathcal{M}_g$ satisfies fairness, truthfulness, achieves feasibility and has welfare maximization properties similar to $\mathcal{M}$.

IV. Experimental Results

In this section, we present simulation results to show the effectiveness of the Mechanism $\mathcal{M}$. We study the allocations resulting from deploying $\mathcal{M}$, as well as the allocations that result from the induced Nash equilibrium (cf. Section III-C). We compare these allocations to a "socially optimal" allocation policy. In all cases, there are two parameters of primary interest:

(1) *Quality of Service*: For an agent with optimal allocation β^*, an allocation b of bandwidth is defined to achieve a QoS value of $Q(\beta^*, b) = b/\beta^*$. Note that $0 \leq Q(\cdot, \cdot) \leq 1$.

(2) *Probability of allocation*: Under some distribution of the agents' utility functions, we define the probability of allocation to be the fraction of agents who receive at least β^0, their minimum acceptable allocation.

Recall that the fairness rank of an agent (denoted by $rank(\cdot)$) is defined as the fraction of agents whose β^0/ℓ values are no larger than that of this agent. In this section, we will work with the following version of $rank$, $F(\cdot) = 1 - rank(\cdot)/n$, thus $0 \leq F(\cdot) \leq 1$. This index measures how closely selfish an agent is relative to the other agents. Agent i s.t. $F(i)$ is close to 1 can be considered to be socially responsible, while a value close to 0 indicates a very selfish agent, *i.e.*, one who requests a lot of bandwidth but is willing to pay very little. A mechanism that is fair with respect to such an index will allocate more bandwidth to an agent who is less selfish.

We will also be interested in studying the correlation between the quality of service received by an agent and the fairness rank of the agent. Since both Q and F map the set of agents into the interval $[0, 1]$, we will use the standard definition of correlation defined as follows. For an allocation $b(\cdot)$, we define $\chi(F, Q) = (1/n) \sum_i F_i Q_i$, where n denotes the number of agents, $F_i = F(i)$, and $Q_i = Q(\beta_i^*, b(i))$. Note that $\chi(F, Q) \in [0, 1/2]$.

As a comparison point for $\mathcal{M}$ and the induced Nash equilibrium, we will study the allocation policy that maximizes the number of agents who receive at least their minimum acceptable bandwidth. We will call this the *socially optimal* allocation.

In evaluating mechanism $\mathcal{M}$ and the socially optimal mechanism, we will work with the true utility functions of the agents.

Our experiments are conducted for the case of a single link of capacity B. Agent i discloses her utility by a tuple $\langle\beta_i^0, \beta_i^*, \beta_i^1\rangle$. The optimum allocation for agent i, β_i^*, is assigned according to a power-law distribution where the probability of a flow requesting $\beta_i^* = b$, $b \in [0, B]$, is proportional to $1/b^2$. This is in accordance with the discovery of [21], [22] that several parameters of interest, especially the packet rate, in Internet traffic measurement obey a power law. For an agent i, given β_i^*, the lower and upper bounds on acceptable bandwidth, β_i^0 and β_i^1, respectively, are computed using a normal distribution around β_i^* with standard deviation randomly chosen in the interval $[0, \beta_i^*/2)$. The price level ℓ_i for agent i is computed using a normal distribution whose mean is proportional to her optimum allocation value β_i^*.

In Figures 3 and 4, we study the effect of an agent's fairness rank on her QoS and her probability of receiving (at least) minimum allocation. These experiments were carried out using 5000 agents with a total bandwidth of 10^6 units and an average bandwidth request of 10^3 units. The number of price levels was set to 5. We now summarize the salient observations from these figures.

(1) $\mathcal{M}$ and the induced Nash equilibrium have rather similar behavior. $\mathcal{M}$ results in higher QoS values for agents with high values of F, that is, the socially responsible agents. The Nash allocation, on the other hand, leads to a higher probability of allocation for more agents than $\mathcal{M}$ does.

(2) The socially optimal mechanism does significantly better than $\mathcal{M}$ and the Nash allocations in terms of serving more agents. However, as can be seen from the droop in Figure 3, it does rather poorly in terms of the QoS received by the agents who are most socially responsible. (Indeed, it is able to serve more agents precisely at the cost of the QoS of the socially responsible agents.)

(3) With the parameter settings chosen, $\mathcal{M}$ and Nash allocations lead to nearly perfect QoS for about 50% of the agents. As the congestion increases, this percentage naturally shifts to a lower value, while the shape of the plots (not shown) remains essentially the same.

In Figures 5 and 6, we study the correlation $\chi(F, Q)$ and the probability of allocation, both as functions of the congestion. In this set of experiments, we fixed the total bandwidth at 10^6, and varied the number of agents from 100 to 10000. The number of price levels was set to 5. Once again, we summarize the main observations below.

(1) In Figure 5, $\mathcal{M}$ and Nash allocations result in perfect QoS when the congestion is low, therefore χ achieves its best possible value of $1/2$. As the congestion increases, the correlation between F and Q degrades gracefully, with a minimum of $1/4$ even at high levels of congestion. This shows that $\mathcal{M}$ and Nash allocations continue to favor the socially responsible agents, even when congestion increases.

(2) The allocation resulting from the Nash equilibrium achieves the best of both worlds (QoS and Prob. of allocation) when congestion increases significantly. Namely, it achieves the best correlation between QoS and fairness rank, and also serves almost as many agents as the socially optimal allocation. (This feature is only more accentuated with higher congestion, plot omitted.)

We also study the effect of differential pricing on the probability of receiving allocation. Figure 7 illustrates this behavior using mechanism $\mathcal{M}$ for various link bandwidths. We set the number of agents to 10000, while the available bandwidth is varied from 0.5×10^6 to 2×10^6; β^* has a power law distribution with mean 10^3. We report the probability of allocation as a function of the ratio of an agent's minimum acceptable allocation to the average minimum acceptable allocation, $\beta^0/\overline{\beta^0}$. It is clear that increasing the level parameter results in higher probability of allocation. It is also interesting to note the following natural trade-off: to achieve a probability of allocation of, say $3/4$, an agent can either enter at level 4

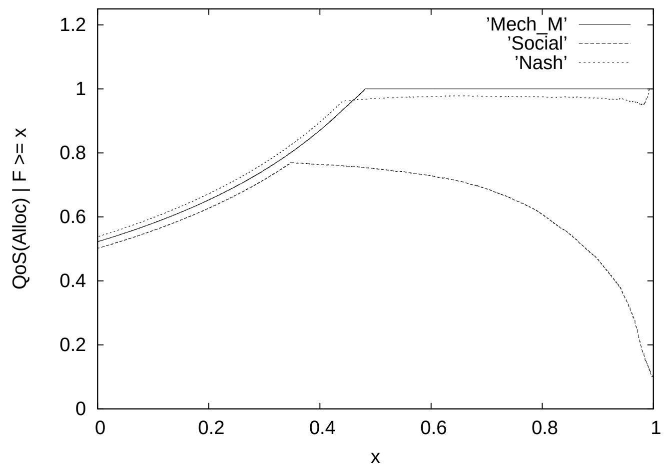

Fig. 3. Fairness rank vs. QoS

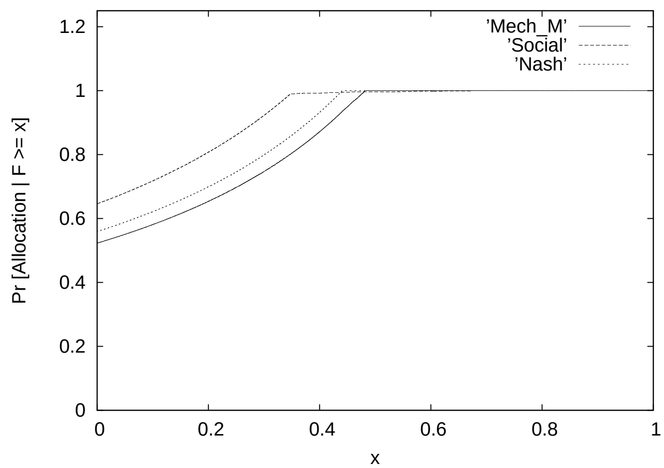

Fig. 4. Fairness rank vs. Prob. allocation

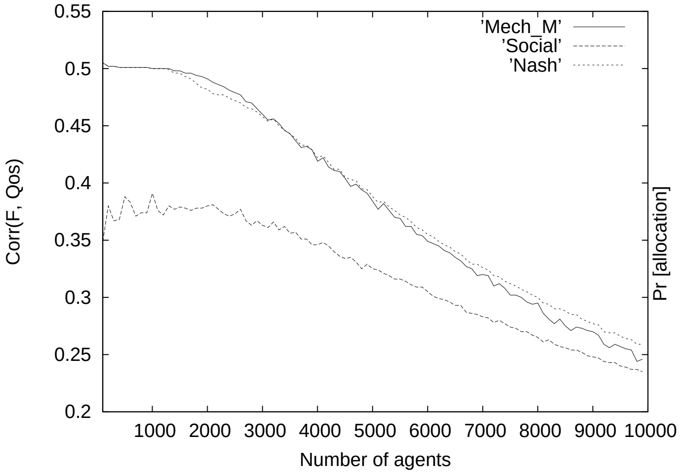

Fig. 5. Congestion vs. $\chi(F, Q)$

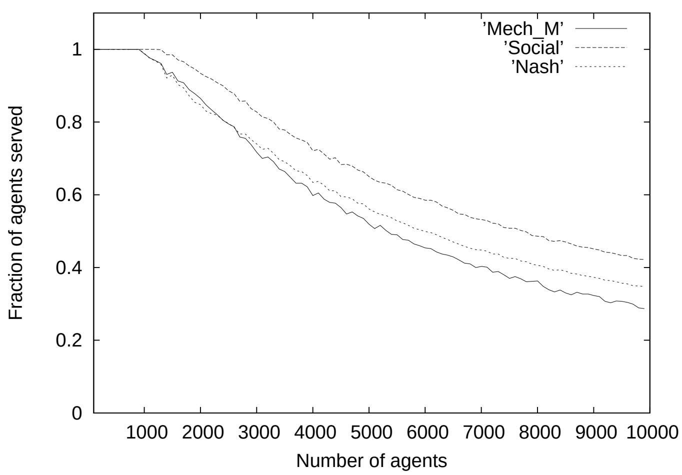

Fig. 6. Congestion vs. Prob. allocation

with $\beta^0/\overline{\beta^0}$ close to 1, or she can enter at level 1 with a $\beta^0/\overline{\beta^0}$ close to $1/4$.

V. Conclusions

One of the fundamental goals of mechanism design is to induce Nash equilibria that have desirable properties (e.g., truthfulness, fairness, approximate optimality, etc.). Our mechanism $\mathcal{M}$ leads to a Nash equilibrium where the allocations are more balanced, more agents receive at least their minimum acceptable allocation, and the less selfish agents tend to receive higher levels of QoS. An interesting open problem is to show that the Nash equilibrium of Theorem 3 yields allocations that are provably good approximations to allocations that optimize some natural criterion (e.g., the number of agents served, average QoS, etc.).

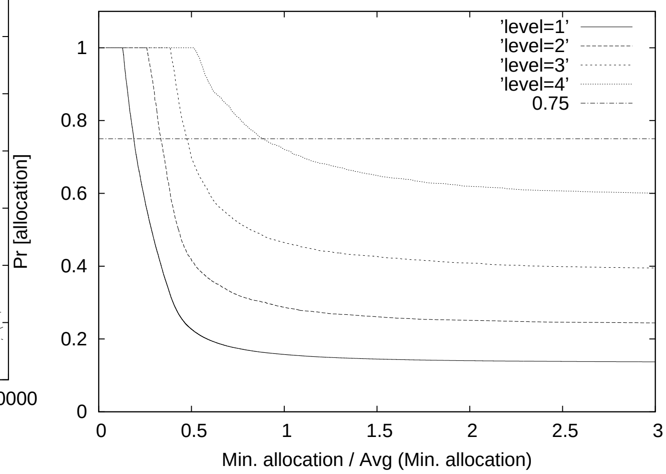

Fig. 7. $\beta^0/\overline{\beta^0}$ vs. Prob. allocation

References

[1] Y. Sprumont, "The division problem with single-peaked preferences: a characterization of the uniform allocation rules," *Econometrica*, vol. 59, pp. 509–519, 1991.

[2] S. Ching, "A simple characterization of the uniform rule," *Economics Letters*, vol. 40, pp. 57–60, 1992.

[3] T. Sönmez, "Consistency, monotonicity, and the uniform rule," *Economic Letters*, vol. 46, pp. 229–235, 1994.

[4] S. Barberà, M. O. Jackson, and A. Neme, "Strategy-proof allotment rules," *Games and Economic Behavior*, vol. 18, pp. 1–21, 1997.

[5] W. Thomson, "Consistent solutions to the problem of fair division where preferences are single peaked," *Journal of Economic Theory*, vol. 63, pp. 219–245, 1994.

[6] J. Rawls, *A Theory of Justice*. Harvard University Press, 1971.

[7] J. Jaffe, "Bottleneck flow control," *IEEE Trans. Communications*, vol. 26, pp. 954–962, 1981.

[8] A. Kumar and J. Kleinberg, "Fairness measures for resource allocation," in *Proceedings of the 41stAnnual IEEE Symposium on Foundations of Computer Science*, 2000, pp. 75–85.

[9] P. Marbach, "Priority service and max-min fairness," in *Proceedings of INFOCOM*, 2002.

[10] H. Moulin, "Incremental cost sharing: characterization by strategyproofness," *Social Choice and Welfare*, vol. 16, pp. 279–320, 1999.

[11] H. Moulin and S. Shenker, "Strategyproof sharing of submodular costs: budget balance vs. efficiency," *Economic Theory*, vol. 18, pp. 511–533, 2001.

[12] J. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, pp. 289–295, 1951.

[13] A. Odlyzko, "Paris metro pricing for the Internet," in *proceedings on ACM Conference on Electronic Commerce*, Denver, Colorado, November 1999, pp. 140–147.

[14] R. Gibbens and P. Key, "The use of games to assess user strategies for differential quality of service in the internet," in *Proc. of the Workshop on Internet Service Quality Economics (MIT)*, 1999.

[15] R. J. Gibbens and F. Kelly, "Resource pricing and the evolution of congestion control," *Automatica*, vol. 35, pp. 1969–1985, 1999.

[16] P. Key and D. McAuley, "Differential qos and pricing in networks: where flow-control meets game theory," *Proc. IEE Proceedings Software*, vol. 146, no. 2, 1999.

[17] Z. Wang, *Internet QoS - Architectures and Mechanisms for Quality of Service*. Morgan Kaufmann Publishers, 2001.

[18] M. Carlson, W. Weiss, S. Blake, Z. Wang, D. Black, and E. Davies, "An architecture for differentiated services," RFC 2475, December 1998.

[19] D. Reininger and R. Izamailov, "Soft quality-of-service for vbr+ video," in *proceedings of the International Workshop on Audio-Visual Services over Packet Networks, AVSPN '97*, Denver, Colorado, September 1997.

[20] E. Fulp, M. Ott, D. Reininger, and D. Reeves, "Paying for QoS: An optimal distributed algorithm for pricing network resources," in *Proceedings of the 6th IWQoS*, 1998, pp. 75–84.

[21] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of ethernet traffic," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 1–15, 1994.

[22] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1995.

Appendix

Proof: We first consider the admission phase; assuming truthfulness in this phase, we then show group strategyproofness for the allocation phase.

(Admission phase) Here we have two cases to consider.

Case 1. In some round j, the utility function for some agent i satisfies $u_i(\mu_j) - \lambda\mu_j < 0$, that is, agent i has negative welfare at the minimum offer μ_j in round j. Here we employ weak strategyproofness to argue that agent i has no incentive to say "yes" and enter S_j in this round. Specifically, agent i, with offer μ_j, must consider the possibility that all the other agents accept the offer μ_j, and the vendor allocates exactly μ_j to some t agents (computed by $\mathcal{M}$ in step (3)). In particular, agent i might be allocated μ_j units of bandwidth, resulting in negative welfare. Notice that, in this case, the vendor has no bandwidth left, so there is no possible increase to b_i in step (11) of $\mathcal{M}$. Therefore, the response that minimizes worst-case risk (where risk is the negative of welfare) is "no." There are simple examples to show that even weak versions of group strategyproofness do not hold.

Case 2. In some round j, the utility function for some agent i satisfies $u_i(\mu_j) - \lambda\mu_j \geq 0$, that is, agent i has non-negative welfare at the minimum offer μ_j in round j. We wish to argue that the agent has no incentive to say "no" and not enter S_j in this round. Let us consider what happens when the agent says "no." There are two possibilities—the agent is either deleted from $S_{j'}$ in some future round $j' > j$ and receives zero bandwidth, or the agent is allocated some bandwidth b in some future round $j' > j$. In the first case, by reporting "yes" in round j, the agent assures herself of non-negative welfare, so there is no reason to say "no." In the second case, the initial offer $\mu_{j'}$ in round j is strictly higher than μ_j; *however*, the final bandwidth allocation b for agent i *cannot* exceed what is left at the end of round j. If agent i truthfully reported "yes" in round j, she would have been eligible for this bandwidth in round j, so is guaranteed at least as much welfare. Therefore, there is no incentive to report "no" when the truthful answer is "yes." Indeed, this analysis does not depend on whether agent i is part of an untruthful coalition, and so we actually have group strategyproofness for this case.

(Allocation phase) In the allocation phase, we show group-strategyproofness, assuming that all the queries in the admission phase were answered truthfully.

Suppose that C is a coalition of agents that mis-report their utility functions. Let u_i, $i = 1, \ldots, n$, denote the true utility functions of the n agents; let $u'_i, i = 1, \ldots, n$, denote their reported utility functions. Note that for $i \notin C$, $u'_i = u_i$. Let b and b' denote the vectors of allocations, respectively, when the utilities reported are u and u'. Similarly, let w and w' denote the vectors of welfare functions. The manipulation property (see Section II-B) implies that for every $i \in C$, $u_i(b'_i) - \lambda b'_i \geq u_i(b_i) - \lambda b_i$.

To prove group strategyproofness, it is sufficient to derive a contradiction to the manipulation property, assuming that when agents in C mis-report their utilities, one of the following happens: (1) some agent not in the coalition has a strictly worse welfare, or (2) some agent (whether or not in the coalition) has a strictly better welfare. (If an agent in C has strictly less welfare, it contradicts the manipulation property.)

Let us consider the allocations for each agent at the end of each round of mechanism $\mathcal{M}$. Let d denote the earliest round at the end which the allocation under u' is different from the allocation under u. At the beginning of round d, the bandwidth available under u and u' are identical, but at the end of round d, one of conditions (1) and (2) above occurs.

Case 1. First we will suppose that the welfare of some agent $j \notin C$ is worse under u' than under u, that is, $w(b'_j) < w(b_j)$. Since $\beta^0_j \leq b'_j, b_j \leq \beta^*_j$, and since w is non-decreasing in the interval (β^0_j, β^*_j), we have $b'_j < b_j$. Furthermore, $b_j \leq \beta^*_j$ (since $j \notin C$), so $b'_j < b_j \leq \beta^*_j$. This implies that when $\mathcal{M}$ is executed with u', agent j remains in the set S_d till the end of round d. Thus at the termination of round d of $\mathcal{M}$ with u', all the bandwidth is exhausted, so $\sum_{k \in Z_d} b_k \leq \sum_{k \in Z_d} b'_k$; on the other hand, we have $b'_j < b_j$, so there must be some agent i for which $b'_i > b_i$.

We now show that i must be in C. That is, for every $i \notin C$, we have $b'_i \leq b_i$. Suppose not, and we have $i \notin C$ such that $b'_i > b_i$ at the end of round d. We noted that $j \in S_d$ at the end of round d with utilities u', so $b'_j \geq b'_i$ (since every agent in S_d at the end of round d has at least as much allocation as any other agent—see Theorem 1, part (1c)). Also, since $b'_i \leq \beta^*_i$, we have $\beta^*_i \geq b'_i > b_i$, which implies that i was in S_d at the termination of round d under utilities u. This implies, again by Theorem 1, part (1c), that $b_i \geq b_j$. Combining these inequalities, we have $b'_j \geq b'_i > b_i \geq b_j$, a contradiction to $b'_j < b_j$.

Thus we have shown that there exists an agent $i \in C$ for whom we have $b'_i > b_i$, that is, the bandwidth lost by agent j could only have been allocated to some agent $i \in C$. This reduces to Case 2.

Case 2. Some agent i has a strictly better welfare. Again, since $b_i \leq \beta^*_i$, it can be seen that $b'_i > b_i$.

(Case 2.1) $b_i = \beta^*_i$, in which case we have $b'_i > b_i = \beta^*_i$; since welfare is non-increasing in the interval (β^*_i, β^1_i), the welfare at b'_i is strictly less than the welfare at b_i, contradicting the manipulation property.

(Case 2.2) Here $b_i < \beta^*_i$ and $b'_i > b_i$. The assumption $b_i < \beta^*_i$ implies that i stays in S_d at the termination of round d under u. Here we crucially use the assumption that the admission step in round d was carried out truthfully, so the same set of agents that enter S_d under u also enter S_d under u'. The condition $b'_i > b_i$ implies that i stays in S_d even longer under the utility profile u'. Where could the excess bandwidth $b'_i - b_i$ have come from? Note that no truthful agent j can be deleted from S_d prior to reaching allocation β^*_j. Thus, the only way the mechanism didn't exhaust its bandwidth under u' is if some agent $k \in C$ has been deleted earlier under u' than under u. Since k is deleted from S_d earlier under u' than under u, the welfare for k is strictly less under u' than under u, a contradiction to the manipulation property. ■