Efficient Signal Proportional Allocation (ESPA) Mechanisms: Decentralized Social Welfare Maximization for Divisible Resources

Rajiv Maheswaran and Tamer Başar, *Fellow, IEEE*

***Abstract*—We address the problem of devising efficient decentralized allocation mechanisms for a divisible resource, which is critical to many technological domains such as traffic management on the Internet and bandwidth allocation to agents in ad hoc wireless networks. We introduce a class of efficient signal proportional allocation (ESPA) mechanisms that yields an allocation which maximizes social welfare with minimal signaling and computational requirements for the resource. Revenue limits for this class are obtained and a sequence of schemes that approach these limits arbitrarily closely are given. We also present a locally stable negotiation scheme applicable to the entire class and illustrate efficiency and revenue properties through simulation.**

***Index Terms*—Communication system economics, game theory, mechanism design.**

I. Introduction

THE FOUNDATIONS and future of information technology, and communication networks in particular, are large syntheses of independent components which require decentralized control, as the scale of these systems and heterogeneous ownership prevent timely or imposable centralization, respectively. Consequently, clients desiring access to these resources must use local information to negotiate for service, often via electronic proxies. The quality-of-service (QoS) received by end users depends critically on the actions of the population of clients (or their proxies) simultaneously requesting a portion of a common scarce resource. While current protocols [such as transmission control protocol (TCP) [1]] are not tailored to a user's specific need, the emergence of voice, video, and peer-to-peer applications portends an evolution toward more intelligent proxies (or *agents*) that capture user or application preferences more explicitly and adapt their actions to fulfill their needs more closely. This interplay between "selfish" agents competing for a common good has invited game-theoretic and economic principles as prevalent tools for analysis and design of network resource allocation.

The key components of most communication networks (bandwidth on an Internet or satellite link, buffer space) along with other vital components of information technology (processor share, memory space) can be characterized as an arbitrarily divisible resource due to the magnitude of their capacity and the flexibility with which they can be partitioned. Classical economic literature with respect to mechanisms for allocating divisible resources is sparse and often inapplicable to technological domains where signaling size and computation to calculate an allocation are significant design metrics. Wilson's *share auctions* [2], and the Generalized Vickrey Auction (GVA) [3] which is an instantiation of a Vickrey–Clarke–Groves (VCG) mechanism [4]–[6], require agents to submit a signal characterizing their entire valuation structure which in this case is infinite dimensional. Even partitioning the resource into reasonable bundle sizes would lead to a vast combinatorial space for which the solution is NP-complete [7]. More recently, Mackie–Mason and Varian introduced the notion of "smart markets" for Internet pricing, where packets would carry bids which determine their level of service [8]. Even though their model could lead to the resource being utilized at less than full capacity and requires a computation of a sort to order bids, the notion of applying market-based control to communication networks persisted and proliferated. Semret and Lazar proposed the progressive second price (PSP) auction [9], an extension of the Vickrey auction, which required two-dimensional signaling and $O(N^2)$ computation at the resource to calculate the allocation, where N is the number of competing agents. While PSP has near-optimal efficiency and desirable revelation properties, in implementation, the auction is indeterminate for five minutes before closing, and that is not on a time-scale which is functional for many highly dynamic environments.

In most current work, network resource allocation is based on a *proportionally fair* (PF) pricing mechanism by Kelly *et al.* [10]. However, when agents incorporate the relationship of the price to the bids into their strategies (turning the mechanism into an auction), the efficiency of the pricing mechanism (with respect to social welfare maximization) is undermined. This loss of efficiency has been studied under cooperative [11] and competitive [12] formulations. Johari and Tsitsiklis [12] show that the worst case performance of the PF auction is no worse than 75% of optimal. Sanghavi and Hajek [13] have suggested a mechanism that improves this to 87.5% for a two-user case and

Manuscript received February 15, 2005; revised January 15, 2006. The work of T. Başar was supported in part by the National Science Foundation (NSF) under Grant ANI-031976.

R. Maheswaran is with the Information Sciences Institute, University of Southern California, Marina Del Rey, CA 90292 USA (e-mail: maheswar@isi.edu).

T. Başar is with the Coordinated Sciences Laboratory, University of Illinois at Urbana–Champaign, Urbana, IL 61801 USA (e-mail: tbasar@control.csl.uiuc.edu).

Digital Object Identifier 10.1109/JSAC.2006.872881

conjecture that the efficiency degrades slightly as the number of users increase.

In this paper, we investigate mechanisms with single-dimensional signaling and $O(N)$ cost of computation for finding an allocation for N signals, as it is the minimal levels of each for arbitrary partitioning a divisible resource. We show that the PF auction, while optimal in an initial expansion where costs are known *a priori* and the allocation scheme is signal proportional, is in general never efficient. Our main contribution is obtaining an infinite subclass of efficient signal proportional allocation (ESPA) mechanisms that always maximize social welfare for an arbitrary collection of agents with quasi-linear utilities. Thus, ESPA mechanisms are the optimal tools for allocating divisible network resources when efficiency (social welfare for strategic agents), computational cost, and signaling space are the metrics of interest. Given this infinite set of mechanisms, we can optimize over a secondary metric, such as revenue generation, while maintaining efficiency. We obtain revenue limits for the ESPA class and discuss how one can approach these revenue limits to arbitrary precision. A locally stable decentralized negotiation protocol and simulations that illustrate the effectiveness of ESPA and various revenue generation properties are also presented. This paper is based and developed on work presented in [14]. The proofs of selected propositions in the exposition are omitted here and the reader is referred to [14].

II. Models and Equilibrium

Transparent mechanisms (or auctions) are characterized by an allocation rule $x(s)$ and a cost rule $c(s)$, where $s = [s_1 \cdots s_N]$ represents the signals from a population of N agents and $x_i(s)$ and $c_i(s)$ are, respectively, the allocation and cost to the ith agent. We work in the one-dimensional signaling space, where $s_i \in I\!\!R$ and $s_i \geq 0$. One subset of this space of auctions is the collection of those that can be characterized by the following allocation rule:

$$x_i(s) = \frac{w_i(s)}{\sum_{j=1}^{N} w_j(s) + \epsilon}$$

where ϵ is a parameter controlled by the resource (e.g., the resource's signal). Signals are translated to weights, denoted by the functions $\{w_i(\cdot)\}_{i=1}^{N}$, which determine the proportions of the allocation. Allocation rules of this form fit nicely with Generalized Processor Sharing models for flow control in networks [15]. We begin our analysis with the case for which $w_i(s) = w(s_i)$and $c_i(s) = c(s_i)$. These restrictions incorporate the notion of fairness where each agent is given the same weight and pays the same cost as any other agent who makes the same signal. It also removes the coupling of the signals away from the weight and cost functions and isolates the interaction in the allocation rule. We assume that the weights and costs are strictly increasing functions of their arguments. We also assume that a signal of zero will yield a weight and cost of zero as well. We consider this class of rules to be a reasonable and tractable initial expansion around the proportionally fair (PF) auction which is the "point" in mechanism space characterized by $w(s_i) = s_i$ and $c(s_i) = s_i$. It can be shown that we do not need to express both $w(s_i)$ and $c(s_i)$. By making the substitutions $t_i = c(s_i)$ and $\tilde{w}(t_i) := w(c^{-1}(t_i))$ (c is invertible if it is monotonically increasing), we can express this class of mechanisms with the rules

$$x_i(t) = \frac{\tilde{w}(t_i)}{\sum_{j=1}^{N} \tilde{w}(t_j) + \epsilon}, \quad c_i(t) = t_i.$$

With similar substitutions, we can equivalently express this class with the rules

$$x_i(s) = \frac{s_i}{\sum_{j=1}^{N} s_j + \epsilon}, \quad c_i(s) = c(s_i). \tag{1}$$

We choose to work with the characterization described in (1), where $c(s_i) \in C^2$ is a twice differentiable increasing function of s_i. We denote a mechanism where $x_i(s)$ is as in (1) and $c_i(s)$ is such that the costs can be computed in $O(N)$ cycles as a signal proportional allocation (SPA) mechanism. SPA mechanisms have the minimal signaling and computational costs for allocation determination that we desire in many communication network contexts.

We model the agents with quasi-linear utility functions

$$U_i(s) = v_i\left(x_i(s)\right) - c_i(s)$$

where v_i is a twice differentiable concave increasing function ($v_i'(\cdot) > 0$, $v_i''(\cdot) \leq 0$). We have the derivatives

$$U_i'(s) = v_i'\left(x_i(s)\right) x_i'(s) - c'(s_i)$$

and

$$U_i''(s) = v_i''\left(x_i(s)\right)\left[x_i'(s)\right]^2 + v_i'\left(x_i(s)\right) x_i''(s) - c''(s_i)$$

where

$$x_i'(s) = \frac{s_{-i} + \epsilon}{(s_i + s_{-i} + \epsilon)^2} > 0$$

$$x_i''(s) = \frac{-2(s_{-i} + \epsilon)}{(s_i + s_{-i} + \epsilon)^3} < 0.$$

If $c_i''(s_i) \geq 0$, then we have $U_i''(s) < 0$. The strict concavity of the ith agent's utility implies that it will have a unique optimal response to each opponent state $s_{-i}+\epsilon$. For the optimal response to be nonzero, we need the marginal utility when bidding zero to be positive. This occurs when

$$\frac{v_i'(0)}{s_{-i} + \epsilon} - c_i'(0) > 0 \;\Rightarrow\; \frac{v_i'(0)}{c_i'(0)} > s_{-i} + \epsilon.$$

The ith agent's response can then be determined from

$$v_i'\left(\frac{s_i}{s_i + s_{-i} + \epsilon}\right)\frac{s_{-i} + \epsilon}{(s_i + s_{-i} + \epsilon)^2} - c'(s_i) = 0 \tag{2}$$

which yields the unique optimal s_i when facing $s_{-i} + \epsilon$.

Let us define $p := \sum_j s_j + \epsilon$. Then, p serves as a measure of demand for the resource and allows us to characterize agents' optimal responses with respect to a parameter which is identical for all agents at equilibrium. Let us call this characterization a demand function $d(p)$, which captures an agent's allocation as a function of p when it uses the strategy obtained from (2). Thus, the demand function captures that $s_i = d(p)p$ is the optimal response to $s_{-i} + \epsilon = d(p)(1-p)$. We now investigate the effects if $c''(s_i) < 0$. The following result, whose proof can be found in [14] captures this.

Proposition 1: If the cost function is concave, then there exists a valuation function for which the optimal response is not unique.

Thus, we restrict our analysis to allocation mechanisms described by (1), where the cost function $c(s_i)$ is convex. We denote this class of mechanisms by $\mathcal{C}$. The intuition behind convex cost functions is that agents who receive larger allocations (due to greater signals) pay a higher cost per unit resource obtained. These occur for strictly convex cost functions and are classified as *discriminatory price* auctions. Mechanisms in $\mathcal{C}$ with linear cost functions such as the proportionally fair auction are *uniform price* auctions. For games played by agents attempting to gain access using a resource allocated through a mechanism from $\mathcal{C}$, it is important to know whether we can obtain a unique operating point, i.e., a unique Nash equilibrium.

Proposition 2: For every mechanism in $\mathcal{C}$, there is a unique Nash equilibrium.

Proof: Making the substitutions $p = \sum_j s_j + \epsilon$ and $x_i = s_i/(\sum_j s_j + \epsilon)$ into (2), we can express the first-order necessary condition for the optimal response as

$$v_i'(x_i)(1-x_i) = pc_i'(px_i). \tag{3}$$

Every pair (p, x_i) that satisfies the previous equation represents an optimal state for the ith agent. We can interpret these states as demand functions (where x_i is a function of p). By treating the previous equation as an identity, we obtain

$$[v_i''(x_i)(1-x_i) - v_i'(x_i)]\frac{\partial x_i}{\partial p} = c_i'(px_i) + pc_i''(px_i)\left[x_i + p\frac{\partial x_i}{\partial p}\right]$$

which implies

$$\frac{\partial x_i}{\partial p} = \frac{c_i'(px_i) + px_i c_i''(px_i)}{v_i''(x_i)(1-x_i) - v_i'(x_i) - p^2 c_i''(px_i)}.$$

Because $c''(s_i) \geq 0$ for all mechanisms in $\mathcal{C}$, and the valuations are increasing concave functions, we have that $\partial x_i/\partial p < 0$. This implies that the demand functions $\{d_i(p)\}_{i=1}^N$ which characterize the optimal responses of agents are decreasing, where $d_i(p) := x_i(p)$ is obtained from the unique value of x_i which solves (3) for a particular p. We note that $d_i(0) = 1\ \forall i$. Following the reasoning in [16], since all agents are characterized by decreasing demand functions, the total demand will be a decreasing function. The Nash equilibrium point is defined by total demand being one which occurs at only one value of p^*. Thus, there is a unique Nash equilibrium with signals $s_i = d_i(p^*)p^*$. ■

Given that we have a class of auctions that yield the desirable property of a unique Nash equilibrium, a natural question is how we go about choosing a mechanism within $\mathcal{C}$. In the next sections, we consider this question with social welfare maximization as a metric.

III. Efficiency in $\mathcal{C}$

A common performance measure of a mechanism (especially in distributed settings) is the efficiency it induces. Whether allocating bandwidth or buffer space, it is desirable to have the resource partitioned in a way that yields the greatest benefit to those accessing it. In economic terms, efficiency is also referred to as the social welfare of those participating in the allocation process. Social welfare is the sum of the valuations of allocations to all agents receiving service, i.e., $\sum_{i=1}^N v_i(x_i)$, where $\{v_i(x_i)\}_{i=1}^N$ are increasing concave functions. Social welfare can also be thought of as the sum of the utilities of all agents, where the resource is a player with utility $U_0(s) = v_0(x_0(s)) + \sum_{j\neq 0} c_j(s)$. Given a scarce resource, the optimal allocations are obtained from the solution of the optimization problem

$$x^* = \arg\max_{x\in X} \sum_{i=1}^N v_i(x_i) \text{ where } X = \left\{x : x_i \geq 0, \sum_{i=1}^N x_i = 1\right\}$$

which can be found by maximizing the Lagrangian

$$\mathcal{L} = \sum_{i=1}^N v_i(x_i) + \lambda\left(1 - \sum_{i=1}^N x_i\right) + \sum_{i=1}^N \mu_i x_i.$$

This is a classical optimization problem whose solution is discussed in [17]. Essentially, an optimal solution is an allocation of the entire available resource where the marginal valuations for agents with positive allocations are all equal. The value of the identical marginal valuations is the Lagrange multiplier. Agents that do not receive positive allocations are those whose highest marginal valuations are less than the value of the Lagrange multiplier. Mathematically, the optimal allocations are characterized as follows:

$$x_i > 0 \Leftrightarrow v_i'(x_i) = \lambda, \quad x_i = 0 \Leftrightarrow v_i'(0) \leq \lambda$$

where λ is chosen such that

$$\sum_{i:x_i>0} x_i = \sum_{i:x_i>0} v_i'^{-1}(\lambda) = 1$$

where $v_i'^{-1}(\lambda)$ is the inverse of the ith agent's valuation function. The intuition behind the solution is as follows. Given an allocation, if one agent (say the ith agent) has a higher marginal valuation than another (say the jth agent), the social welfare can be improved by marginally increasing x_i and marginally decreasing x_j. Thus, in an optimal allocation, all active agents should have identical marginal valuations. We refer to a mechanism that yields an allocation that maximizes social welfare for all collections of agents characterized by valuation functions $\{v_i(x_i)\}_{i=1}^N$ (with $v_i(x_i)$ as defined earlier) as *efficient*.

We now investigate the efficiency of SPA mechanisms. We begin with the class of mechanisms in $\mathcal{C}$ and argue that while linear cost functions outperform strictly convex cost functions with respect to social welfare, they are not efficient. We then expand our rule space beyond $\mathcal{C}$ to devise efficient SPA mechanisms. We also assume that $\epsilon = 0$ (i.e., the entire resource is allocated) for mathematical and notational simplicity. Extensions to the case where $\epsilon > 0$ are straightforward.

Proposition 3: Within the class of mechanisms in $\mathcal{C}$, social welfare is maximized when the cost functions are linear.

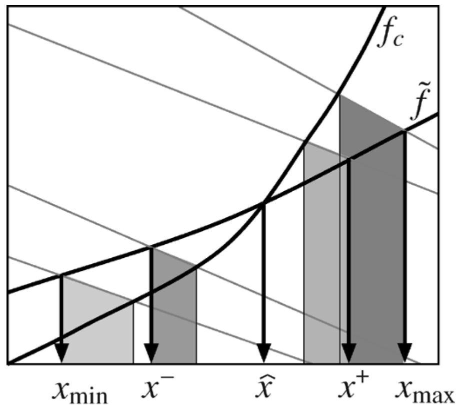

Fig. 1. Illustration of $\tilde{f}$ and f_c along with marginal valuations for four agents.

Proof: The allocation for the ith agent is determined by

$$v_i'(x_i) = \frac{pc'(px_i)}{1-x_i} =: f(p, x_i).$$

We have $c''(\cdot) \geq 0$ for all cost functions in $\mathcal{C}$, which implies that $f(p, x_1) > f(p, x_2)$ if $x_1 > x_2$ for all $p > 0$. If $v_i'(0) \leq f(p, 0)$, then the ith agent will not receive a positive allocation at p. Otherwise, the ith agent will receive a unique positive allocation $x_i(p)$ which is the solution of $v_i'(x_i) = f(p, x_i)$. The uniqueness is because v_i' is a decreasing function of x_i and f is a strictly increasing function of x_i. We also have $f(p_1, x) > f(p_2, x)$ if $p_1 > p_2$ for all $x \in (0, 1)$. This implies that $x_i(p_1) < x_i(p_2)$ if $p_1 > p_2$. Given any cost function $c(s)$, we can find the equilibrium p by solving $\sum_i x_i(p) = 1$, from which we can obtain the equilibrium allocations $\{x_i(p)\}$.

For a linear cost function $c(s) = ks$, we have $f(p, x_i) = pk/(1-x_i)$. We note that every linear cost function yields the same allocations to participating agents, though the equilibrium p might differ. Let us assume that for $c(s) = k_1 s$, the equilibrium is at p_1. Then, if $c(s) = k_2 s$, $p_2 = p_1 k_1/k_2$ satisfies all the conditions for equilibrium with the same allocations as the case with $c(s) = k_1 s$. Alternatively, we can let $\tilde{p} = pk$, and obtain $x_i(\tilde{p})$ from the equations $v_i'(x_i) = \tilde{p}/(1-x_i) =: \tilde{f}(\tilde{p}, x_i)$. The equilibrium allocations for all linear cost functions $\{x_i(\tilde{p}^*)\}$ can be expressed in terms of $\tilde{p}^*$, which is the solution to $\sum_i x_i(\tilde{p}) = 1$.

Let us now consider a strictly convex cost function $c(s)$. Let us assume that for some $\hat{x}$ and $\hat{p}$, we have

$$\frac{\hat{p}c'(\hat{p}\hat{x})}{1-\hat{x}} = \frac{\tilde{p}^*}{1-\hat{x}} \Rightarrow \hat{p}c'(\hat{p}\hat{x}) = \tilde{p}^*.$$

Because $c''(\cdot) > 0$, we have $\forall x > \hat{x}$

$$\frac{\hat{p}c'(\hat{p}x)}{1-x} > \frac{\hat{p}c'(\hat{p}\hat{x})}{1-x} = \frac{\tilde{p}^*}{1-x}$$

and $\forall x < \hat{x}$, we have

$$\frac{\hat{p}c'(\hat{p}x)}{1-x} < \frac{\hat{p}c'(\hat{p}\hat{x})}{1-x} = \frac{\tilde{p}^*}{1-x}.$$

What we have shown is that if $f_c(p, x) := pc'(px)/(1-x)$ (obtained from a strictly convex function) intersects $\tilde{f}(\tilde{p}, x)$ at some point $\hat{x}$, then f_c will be larger than $\tilde{f}$ for $x > \hat{x}$ and f_c will be less than $\tilde{f}$ for $x < \hat{x}$.

Let us now assume that we have an arbitrary heterogeneous agent population. Let $x_i^c(p)$ be obtained from the solution of $v_i'(x_i) = f_c(p, x_i)$ and $\tilde{x}_i$ be obtained from the solution of $v_i'(x_i) = \tilde{f}(\tilde{p}^*, x_i)$. Let $x_{\min} := \min\{\tilde{x}_i : \tilde{x}_i > 0\}$ be the smallest positive allocation and $x_{\max} := \max\{\tilde{x}_i\}$ be the largest allocation under a linear cost function. Since we are comparing the performance over all agent populations, we consider one for which $x_{\min} < x_{\max}$. If $p_{\min}$ is where

$$\begin{aligned} f_c(p_{\min}, x_{\min}) &= \frac{p_{\min}c'(p_{\min}x_{\min})}{1-x_{\min}} \\ &= \frac{\tilde{p}^*}{1-x_{\min}} = \tilde{f}(\tilde{p}^*, x_{\min}) \end{aligned}$$

then $f_c(p_{\min}, x) > \tilde{f}(\tilde{p}^*, x)$ for all $x > x_{\min}$, and $\sum_i x_i^c(p_{\min}) < \sum_i \tilde{x}_i = 1$. Similarly, if $p_{\max}$ is where

$$\begin{aligned} f_c(p_{\max}, x_{\max}) &= \frac{p_{\max}c'(p_{\max}x_{\max})}{1-x_{\max}} \\ &= \frac{\tilde{p}^*}{1-x_{\max}} = \tilde{f}(\tilde{p}^*, x_{\max}) \end{aligned}$$

then $f_c(p_{\max}, x) < \tilde{f}(\tilde{p}^*, x)$ for all $x < x_{\max}$, and $\sum_i x_i^c(p_{\max}) > \sum_i \tilde{x}_i = 1$.

This implies that the $\hat{p}$ for which $\sum_i x_i^c(\hat{p}) = 1$ must lie strictly between $p_{\min}$ and $p_{\max}$. Furthermore, $f_c(\hat{p}, x)$ intersects $\tilde{f}(\tilde{p}^*, x)$ at a point $\hat{x} \in (x_{\min}, x_{\max})$. For $x > \hat{x}$

$$f_c(\hat{p}, x) > \tilde{f}(\tilde{p}^*, x) \Rightarrow \tilde{x}_i > x_i^c(\hat{p}) > \hat{x}.$$

For $x < \hat{x}$

$$f_c(\hat{p}, x) < \tilde{f}(\tilde{p}^*, x) \Rightarrow \tilde{x}_i < x_i^c(\hat{p}) < \hat{x}.$$

If ΔSW is the difference in social welfare between the allocations for a linear cost function $(\tilde{x})$ and the allocations for a strictly convex cost function $(x^c(\hat{p}))$, we have

$$\begin{aligned} \Delta SW &= \sum_i v_i(\tilde{x}_i) - \sum_i v_i\left(x_i^c(\hat{p})\right) \\ &= \sum_i \int_0^{\tilde{x}_i} v_i'(x)dx - \sum_i \int_0^{x_i^c(\hat{p})} v_i'(x)dx \\ &= \sum_{i:\tilde{x}_i > \hat{x}} \int_{x_i^c(\hat{p})}^{\tilde{x}_i} v_i'(x)dx - \sum_{i:\tilde{x}_i < \hat{x}} \int_{\tilde{x}_i}^{x_i^c(\hat{p})} v_i'(x)dx. \end{aligned}$$

If $x^+ := \min\{\tilde{x}_i : \tilde{x}_i > \hat{x}\}$ and $x^- := \max\{\tilde{x}_i : \tilde{x}_i < \hat{x}\}$, then $\Delta SW >$

$$\begin{aligned} &\sum_{i:\tilde{x}_i > \hat{x}} \int_{x_i^c(\hat{p})}^{\tilde{x}_i} \frac{\tilde{p}^*}{1-x^+}dx - \sum_{i:\tilde{x}_i < \hat{x}} \int_{\tilde{x}_i}^{x_i^c(\hat{p})} \frac{\tilde{p}^*}{1-x^-}dx \\ &= \frac{\tilde{p}^*}{1-x^+} \sum_{i:\tilde{x}_i > \hat{x}} \left[\tilde{x}_i - x_i^c(\hat{p})\right] - \frac{\tilde{p}^*}{1-x^-} \sum_{i:\tilde{x}_i < \hat{x}} \left[x_i^c(\hat{p}) - \tilde{x}_i\right]. \end{aligned}$$

Since the sums of allocations in both cases are one, we have

$$\begin{aligned}\sum_{i:\tilde{x}_i>\hat{x}}[\tilde{x}_i - x_i^c(\hat{p})] &= \left[\sum_{i:\tilde{x}_i>\hat{x}}\tilde{x}_i\right] - 1 + 1 - \left[\sum_{i:\tilde{x}_i>\hat{x}} x_i^c(\hat{p})\right] \\ &= \left[\sum_{i:\tilde{x}_i>\hat{x}}\tilde{x}_i\right] - \left[\sum_i \tilde{x}_i\right] \\ &\quad + \left[\sum_i x_i^c(\hat{p})\right] - \left[\sum_{i:\tilde{x}_i>\hat{x}} x_i^c(\hat{p})\right] \\ &= \sum_{i:\tilde{x}_i<\hat{x}}[x_i^c(\hat{p}) - \tilde{x}_i] =: \alpha > 0.\end{aligned}$$

Incorporating this into the bound on social welfare difference, we have

$$\Delta SW > \alpha\tilde{p}^*\left[\frac{1}{1-x^+} - \frac{1}{1-x^-}\right] > 0.$$

∎

The proof (shown first in [14]) exploits the properties of (3) under the different cost structures. The optimality of linear cost functions (including the PF auction) among mechanisms in $\mathcal{C}$ in addition to their practical benefits (ease of implementation, etc.) gives credence to their use. However, we still are unable to induce efficiency by limiting ourselves to this class of allocation and cost rules.

Corollary 1: There is no mechanism in $\mathcal{C}$ which maximizes social welfare for all agent populations.

Proof: We consider the case of two agents with valuation functions such that $v_1'(x) > v_2'(x)$ for all $x \in [0,1]$. We consider a mechanism with a linear cost function as it yields the optimal efficiency among all cost functions in $\mathcal{C}$. The allocation for Agent 1 is obtained from the solution to $v_1'(x_1) = p/(1-x_1)$ for some $p > 0$. This implies that $x_1 < 1$, which in turns implies that $x_2 > 0$. Since both agents are active, their marginal valuation functions must intersect $p/(1-x)$, which is an increasing function of x. This implies that the agents' marginal valuations are not equal at equilibrium unless $x_1 = x_2 = 1/2$, and that cannot occur because $v_1'(1/2) > v_2'(1/2)$. Since both agents are active and their marginal valuations are not equal, the social welfare can be improved and thus the mechanism is not efficient. ∎

IV. Designing Efficient Auctions

The intuition behind why mechanisms in $\mathcal{C}$ cannot be efficient lies in the equilibrium condition $v_i'(x_i) = pc'(px)/(1-x) =: f(p,x)$. Here, $f(p,x)$ plays the role of the Lagrange multiplier. For cost functions that are convex, $f(p,x)$ is an increasing function of x, which yields unequal marginal valuations at equilibrium. For optimality we need a cost function that would yield a function $f(p,x)$ which induces identical marginal valuations at equilibrium, i.e., $f(p,x)$ which is independent of x. This would make $f(p,x) = g(p)$ "flat" across $x \in [0,1]$ and the marginal valuation functions of all active agents would intersect f at the same value. We refer to $g(p)$ as the *generator function*. The generator function serves the purpose of the Lagrange multiplier in the social welfare maximization problem. Thus, if we are to use this generator function to obtain a maximum efficiency cost function for all agent populations, it must be able to span all viable values that a Lagrange multiplier might take, i.e., all nonnegative real numbers. Furthermore, we need the generator function to be one-to-one. Otherwise, an agent population whose optimal allocations occur at a particular Lagrange multiplier value could be reached at two different values of p, which indicates multiple equilibria. We now show that by using an appropriate generator function, we can construct cost functions that induce an equilibrium at which active agents have the same marginal valuation for their allocations.

Proposition 4: Let $g(p)$ be a one-to-one function whose range space is the set of all nonnegative real numbers. Consider the mechanism

$$x_i(s) = \frac{s_i}{s_i + s_{-i}} \quad c_i(s) = s_{-i}\int_0^{s_i}\frac{g(t+s_{-i})}{(t+s_{-i})^2}dt$$

where $s_{-i} = \sum_{j\neq i} s_j + \epsilon$. For an arbitrary agent population with quasi-linear utilities and concave increasing valuations, any equilibrium under this mechanism will yield a solution where all active agents have the same marginal valuation.

Proof: First, we explain the construction of the cost function. If we set the marginal valuation to be equivalent to the generator function, we have

$$v_i'(x_i) = \frac{pc'(px_i)}{1-x_i} = g(p) \quad \Rightarrow$$
$$c'(px_i) = \frac{g(p)}{p}(1-x_i) = \frac{g(p)}{p^2}p(1-x_i).$$

At equilibrium, we have $px = s_i$, $p(1-x) = s_{-i}$, and $p = s_i + s_{-i}$. We realize that we cannot express the marginal cost solely as a function of s_i. However, we can express it as a function of s_i and s_{-i} as follows:

$$c_i'(s_i; s_{-i}) = \frac{g(s_i + s_{-i})}{(s_i + s_{-i})^2}s_{-i}. \tag{4}$$

Given that this expression for marginal cost holds for all equilibrium solutions, we can integrate over s_i to obtain

$$c_i(s) = c(s_i; s_{-i}) = s_{-i}\int_0^{s_i}\frac{g(t+s_{-i})}{(t+s_{-i})^2}dt. \tag{5}$$

We now verify that this yields equal marginal valuations at equilibrium for active agents by analyzing agents' optimal responses. If an agent with utility $U_i(s) = v_i(x_i(s)) - c_i(s)$ has a positive allocation at equilibrium, its optimal signal is an extremal point obtained as a solution of

$$U_i'(s) = v_i'\left(\frac{s_i}{s_i + s_{-i}}\right)\frac{s_{-i}}{(s_i + s_{-i})^2} = c'(s_i; s_{-i}).$$

Substituting the marginal cost

$$v_i'\left(\frac{s_i}{s_i + s_{-i}}\right)\frac{s_{-i}}{(s_i + s_{-i})^2} = \frac{g(s_i + s_{-i})}{(s_i + s_{-i})^2}s_{-i}$$
$$\Rightarrow \quad v_i'\left(\frac{s_i}{s_i + s_{-i}}\right) = g(s_i + s_{-i}) = g(p).$$

Thus, the marginal valuations of active agents for any set of bids that form an equilibrium solution are identical. In fact, the value of the marginal valuations is the output of the generator function at the equilibrium value of p. Furthermore, for any inactive agent at equilibrium, we have

$$v_i'(0) \le s_{-i}c'(0;s_{-i}) = s_{-i}\frac{g(0+s_{-i})}{(0+s_{-i})^2}s_{-i} = g(s_{-i}) = g(p)$$

which meets our conditions for a solution to the social welfare maximization problem. ■

We note that the key to these mechanisms is the s_{-i} factor in the cost functions as it cancels the s_{-i} that appears when we take the partial derivative of $v_i(x_i(s))$ with respect to s_i. In essence, by making agents account for increased demand in their costs as well as the allocation, we are able to achieve maximum efficiency. Table I displays cost functions associated with various generator functions.

We see that we can generate a diverse set of cost functions that yield equal marginal valuations at equilibrium. All cost functions yield a cost of zero if the agent bids zero. The cost function with the simplest and most intuitive form may be that generated by $g(p) = p^2$, which yields $c(s) = s_i s_{-i}$. This states that an agent's cost depends linearly on both its own signal and the sum of signals of all other agents.

To this point, we have neglected to analyze the effect of the cost function on equilibrium. We know that if an equilibrium exists, it will maximize social welfare. The question that follows is what generator functions yield cost functions that lead to the existence of a unique equilibrium.

Proposition 5: Let $g(p)$ be a one-to-one function whose range space is the set of all nonnegative real numbers. Furthermore, let $g'(p)$ exist and be positive for all $p \ge 0$. Then, the mechanism using the cost function generated by $g(p)$ yields a unique equilibrium.

Proof: We show this by obtaining demand functions and showing that they are decreasing functions of p. We already know that $v'(x) = g(p)$ at equilibrium. Taking this as an identity, we have

$$v''(x)\frac{\partial x}{\partial p} = g'(p) \quad \Rightarrow \quad \frac{\partial x}{\partial p} = \frac{g'(p)}{v''(x)} < 0$$

for all concave valuation functions. Since the demand functions are decreasing, we can apply similar reasoning from Proposition 2 to state that we have a unique Nash equilibrium. ■

We refer to the class of auctions created from the generator functions as described in Proposition 4 as ESPA mechanisms. To obtain a unique Nash equilibrium, we must limit ourselves to strictly increasing generator functions. The functions listed in Table I all satisfy this requirement. If we generated a cost function using $g(p) = p^k$ with $k \le 0$, the resulting demand functions would not decrease and no equilibrium solution would exist. The set of strictly increasing functions is an infinite set from which we can obtain generator functions. Thus, we can obtain an infinite number of efficient mechanisms.

V. Revenue Generation

Given that we have an infinite number of mechanisms that maximize social welfare, we can optimize a secondary metric over this class. A natural choice would be the revenue generated from the allocation. Given a collection of agents characterized by their valuation functions, $\mathcal{V} := \{v_1, \ldots, v_N\}$, we know λ^* and $\{x_i^*\}_{i=1}^N$ from the solution of the problem (described earlier), where

$$x^* = \arg\max_{x\in X}\sum_{i=1}^N v_i(x_i) \text{ where } X = \left\{x : x_i \ge 0, \sum_{i=1}^N x_i = 1\right\}.$$

Here, λ^* is the optimal value of the Lagrange multiplier λ when

$$\mathcal{L} = \sum_{i=1}^N v_i(x_i) + \lambda\left(1 - \sum_{i=1}^N x_i\right) + \sum_{i=1}^N \mu_i x_i.$$

When using an efficient allocation mechanism, an agent characterized by the valuation function $v_i(x)$ will receive an allocation x_i^* at equilibrium. Also, since the marginal valuations of active agents at equilibrium will be equal to the Lagrange multiplier, we have $g(p^*) = \alpha h(p^*) = \lambda^*$, where α is a constant scale factor and $h(\cdot)$ has the properties of an ESPA generator function.

Corollary 2: The cost paid by an agent receiving a proportion x of a divisible resource in a collection of agents, where the Lagrange multiplier for an efficient allocation is λ, is

$$c(\lambda, x) = \alpha h^{-1}\left(\frac{\lambda}{\alpha}\right)(1-x)\int_{h^{-1}\left(\frac{\lambda}{\alpha}\right)(1-x)}^{h^{-1}\left(\frac{\lambda}{\alpha}\right)} \frac{h(u)}{u^2}du. \quad (6)$$

Proof: Dropping the * superscript and making the substitutions $s_i = x_i p$, $s_{-i} = (1-x_i)p$, and $p = h^{-1}(\lambda/\alpha)$ into (5), we obtain (6). ■

Table II displays the costs paid by an agent under cost rules from various generator functions expressed as a function of share received and marginal valuation at equilibrium. We note that the costs are independent of the scale factor α. We now investigate the upper limit on the revenue that can be generated by an ESPA mechanism.

Proposition 6: Given a collection of agents $\mathcal{V}$, an ESPA mechanism cannot generate revenue at equilibrium which is greater than $\lambda^*(\mathcal{V})$.

TABLE I
Cost Rules for Given Generators

$g(p)$	$c(s_i; s_{-i})$
p	$s_{-i}\log(1 + s_i/s_{-i})$
$p^k, k > 0, k \ne 1$	$(s_{-i}/k-1)[(s_i+s_{-i})^{k-1} - s_{-i}^{k-1}]$
$p^{1/k}, k > 1$	$\frac{s_{-i}}{1/k-1}[(s_i+s_{-i})^{(1/k-1)} - s_{-i}^{(1/k-1)}]$

TABLE II
Share Costs for Given Generators

$g(p)$	$c(\lambda, x)$
p	$-\lambda(1-x)\log(1-x)$
$p^k, k > 0, k \ne 1$	$\frac{\lambda(1-x)}{k-1}[1-(1-x)^{(k-1)}]$
$p^{1/k}, k > 1$	$\frac{\lambda(1-x)}{1/k-1}\left[1-(1-x)^{(1/k-1)}\right]$

Proof: The notation $\lambda^*(\mathcal{V})$ denotes that the Lagrange multiplier of the social welfare maximization problem depends on the collection of agents. Let us assume that for some collection of agents characterized by the valuation set $\mathcal{V} := \{v_1, \ldots, v_N\}$, the revenue generated by an ESPA mechanism is greater than λ^*. This implies that $\exists i \in \{1, \ldots, N\}$ such that $c(\lambda^*, x_i^*) = \lambda^* x_i^* + \delta$, for some $\delta > 0$. Let $\tilde{\mathcal{V}} := \{v_1, \ldots, \tilde{v}_i, \ldots, v_N\}$ be another collection of agents identical to $\mathcal{V}$ except for the ith agent who is now characterized by a valuation function, where $\tilde{v}_i'(x) = \lambda^* + (x_i^* - x)\epsilon$ with $\epsilon > 0$ and $\lambda^* = \lambda^*(\mathcal{V})$. Because $v_i'(x_i^*) = \tilde{v}_i(x_i^*) = \lambda^*(\mathcal{V})$, all ESPA mechanisms would yield the same allocations, $\{x_j^*\}_{j=1}^N$, equilibrium signals and costs for both collections. However, in $\tilde{\mathcal{V}}$, we have

$$\begin{aligned} U_i\left(s_i^*; s_{-i}^*\right) &= U_i\left(\lambda^*, x_i^*\right) = v\left(x_i^*\right) - c\left(\lambda^*, x_i^*\right) \\ &= \left[\lambda^* x_i^* + \frac{\epsilon\left(x_i^*\right)^2}{2}\right] - \left[\lambda^* x_i^* + \delta\right] \\ &= \frac{\epsilon\left(x_i^*\right)^2}{2} - \delta < 0 \end{aligned}$$

for ϵ sufficiently small. This implies that the ith player has a negative utility at the efficient equilibrium point, and thus will not to participate at that allocation. This further implies that if an ESPA mechanism generates revenue greater that λ^*, there is an agent collection for which it does not induce an efficient allocation, and hence is not an ESPA mechanism. Thus, ESPA mechanisms cannot generate revenue greater than λ^*. ∎

We now investigate how close we can reach this limit. We introduce the notion of an *extremal* optimal allocation as one where a single agent obtains access to the entire resource. This would occur when $v_i'(1) > v_j'(0)\ \forall j \neq i$.

Proposition 7: For agent collections with nonextremal optimal allocations, we can generate revenue arbitrarily close to λ^* by using the ESPA mechanism associated with the generator function $g(p) = p^{1/k}$ with k sufficiently large.

Proof: From the last row of Table II, we have that each agent with a positive allocation pays a cost

$$c\left(\lambda^*, x_i^*\right) = \frac{\lambda^*\left(1 - x_i^*\right)}{\frac{1}{k} - 1}\left[1 - \left(1 - x_i^*\right)^{\left(\frac{1}{k} - 1\right)}\right].$$

Taking the limit of $c(\lambda^*, x_i^*)$ as $k \to \infty$, we have

$$c\left(\lambda^*, x_i^*\right) \to -\lambda^*\left(1 - x_i^*\right)\left[1 - \left(1 - x_i^*\right)^{-1}\right] = \lambda^* x_i^*.$$

Thus, for any $\epsilon > 0$, we can find a k^* such that $c(\lambda^*, x_i^*) > \lambda^* x_i^* - \epsilon/N$. Using the ESPA mechanism associated with $g(p) = p^{1/k^*}$, the revenue generated will be

$$\sum_i c\left(\lambda^*, x_i^*\right) > \sum_i \lambda^* x_i^* - \frac{\epsilon}{N} = \lambda^* - \epsilon.$$

∎

We note that calculating k^* *a priori* to guarantee revenue with a certain proximity to λ^* would require some knowledge of the minimum positive allocation that might result for expected populations of agents. As we will see later, there can be undesirable effects of using an ESPA mechanism generated from $g(p) = p^{1/k}$ with k very large, which will temper the temptation to extract revenue very tightly. We also note that revenue generation for agent collections with extremal optimal allocations can at first inspection cause some difficulties, namely because the winning agent's signal cost becomes zero. This can be countered with the resource sending a fixed signal ϵ which is also useful for countering collusion [17]. As extremal allocations are rare in most communication network contexts, we leave this discussion for another forum.

VI. Negotiation

In settings with distributed control, it is unlikely for agents to reach an equilibrium state after one round of signaling. Hence, we require a protocol for negotiating or discovering a stable operating point, typically based on some network feedback. For ESPA mechanisms, returning the value of p is a single-dimensional parameter that naturally serves this purpose as it represents a measure of aggregate demand. We propose the following relaxed update scheme for negotiation where the superscript indicates the round and α_i is the relaxation parameter:

$$s_i^{n+1} = \alpha_i\left[x_i(s^n) p_i\left(x_i(s^n)\right)\right] + (1 - \alpha_i) s_i^n.$$

Intuitively, an agent approaches the signal for which it would accept the current allocation. Here, $p_i(x_i)$ is the inverse of the demand function $d_i(p)$ described earlier and it represents an agent's optimal response as a mapping that gives the value of p at which the ith agent would demand a share x_i of the resource. Obtaining this function is not trivial for an arbitrary mechanism, but for an ESPA mechanism, we know that the marginal valuation for any active agent at an equilibrium allocation is captured by the generator function, i.e., $v_i'(x_i) = g(p)$. This relationship was determined from a relation representing an agent's optimal response, thus treating this as an identity, we have $p_i(x_i) = g^{-1}(v_i'(x_i))$. The local evolution of the system of agents employing negotiation strategies described in Section VI can be expressed as

$$\left(s^{n+1} - s^*\right) = \left[A(1 - X)Q + (1 - A)\right]\left(s^n - s^*\right)$$

where

$$J = A(1 - X)Q + (1 - A)$$

is the Jacobian for which A is a diagonal matrix of the relaxation parameters α_i, $X = x^* 1_N^T$, where x^* is a column vector of optimal resource shares, 1_N is a length-N column vector of ones, and Q is a diagonal matrix with the ith diagonal element

$$q_i = \frac{\left(p_i\left(x_i^*\right) + x_i^* p_i'\left(x_i^*\right)\right)}{\left(\sum_j s_j^*\right)}.$$

Proposition 8: If $\alpha_i < 2/\hat{q}_i$ where $\hat{q}_i = \max_{x_i \in [0,1]} - x_i p_i'(x_i)/p_i(x_i)$, the relaxed update scheme described in Section VI is locally stable.

Proof: The proof is analogous to that presented in [16] for a stable negotiation protocol for a PF auction, where the functions $\{p_i(x_i)\}$ are obtained by applying the inverse of the generator function to the marginal valuation function as described above. ■

Proposition 8 leads to an interesting insight regarding the ESPA mechanisms that approach the revenue limits. We can show that for the class of revenue maximizing generator functions where $g(p) = p^{1/k}$ for $k > 1$, we have $\hat{q}_i = k\hat{x}_i v_i''(\hat{x}_i)/v_i'(\hat{x}_i)$, where $\hat{x}_i = \arg\max_{x_i} x_i v_i''(x_i)/v_i'(x_i)$. As we increase $k \to \infty$, presumably to approach the revenue limits, we have $\hat{q}_i \to \infty$ which implies the value of the relaxation parameter necessary to guarantee local stability approaches zero, which indicates a slower convergence to equilibrium. Thus, there seems to exist a tradeoff between revenue and rate of convergence.

VII. Simulation

We now illustrate the efficiency, revenue, and convergence properties of the ESPA mechanisms through simulation using the decentralized relaxed update scheme described in Section VI. We consider a collection of N agents which are characterized by valuation functions of the following form: $\{v_i(x_i) = b_i x_i - (1/2)a_i x_i^2\}_{i=1}^N$. Given any such collection, we can find the Lagrange multiplier to the social welfare maximization problem as follows:

$$\lambda^* = \lambda_{\hat{n}+1}, \quad \text{where} \quad \hat{n} = \arg\max_{n>1} \{n : \lambda_n \le b_n\},$$

$$\text{and} \quad \lambda_n = \frac{\left[\left(\sum_{i=1}^{n-1} \frac{b_i}{a_i}\right) - 1\right]}{\sum_{i=1}^{n-1} \frac{1}{a_i}}.$$

From this we can obtain the optimal allocations ($x_i^* = (b_i - \lambda^*)/a_i$ if $b_i > \lambda^*$ and $x_i^* = 0$, otherwise), which then yields the value of social welfare for an efficient allocation mechanism: $\sum_{i=1}^N b_i x_i^* - (1/2)a_i(x_i^*)^2$. For any given agent collection, we can apply the relaxed update scheme (with some random initial signals) which we know will converge with a sufficiently small relaxation parameter (we never encountered a case over hundreds of tests where the local condition was met but the global system did not converge, though proof of global stability is an open question). By comparing the resulting allocations at equilibrium and the sum of their valuations to the optimal values derived theoretically, we could verify if the ESPA mechanisms indeed were efficient. Agent populations varied from 10 to 100, and the valuation functions were chosen such that $b_i^k \in [0\ 1]$ and $(a_i/b_i)^k \in [0\ 1]$ with uniform probability (k is a factor which can be used to determine what percentage of agents had positive allocations at equilibrium). We tested the ESPA mechanisms for which $g(p) \in \{p^2, p^{1/2}, p^{1/4}, p^{1/6}, p^{1/8}, p^{1/10}\}$ in addition to the PF mechanism. As expected, all ESPA mechanisms yielded an efficient allocation for every test run, while the PF mechanism never did (though it did typically yield allocations with greater than 90% efficiency). Sample evolutions for a collection of 10 and 100 agents are displayed in Fig. 2(a) and (b).

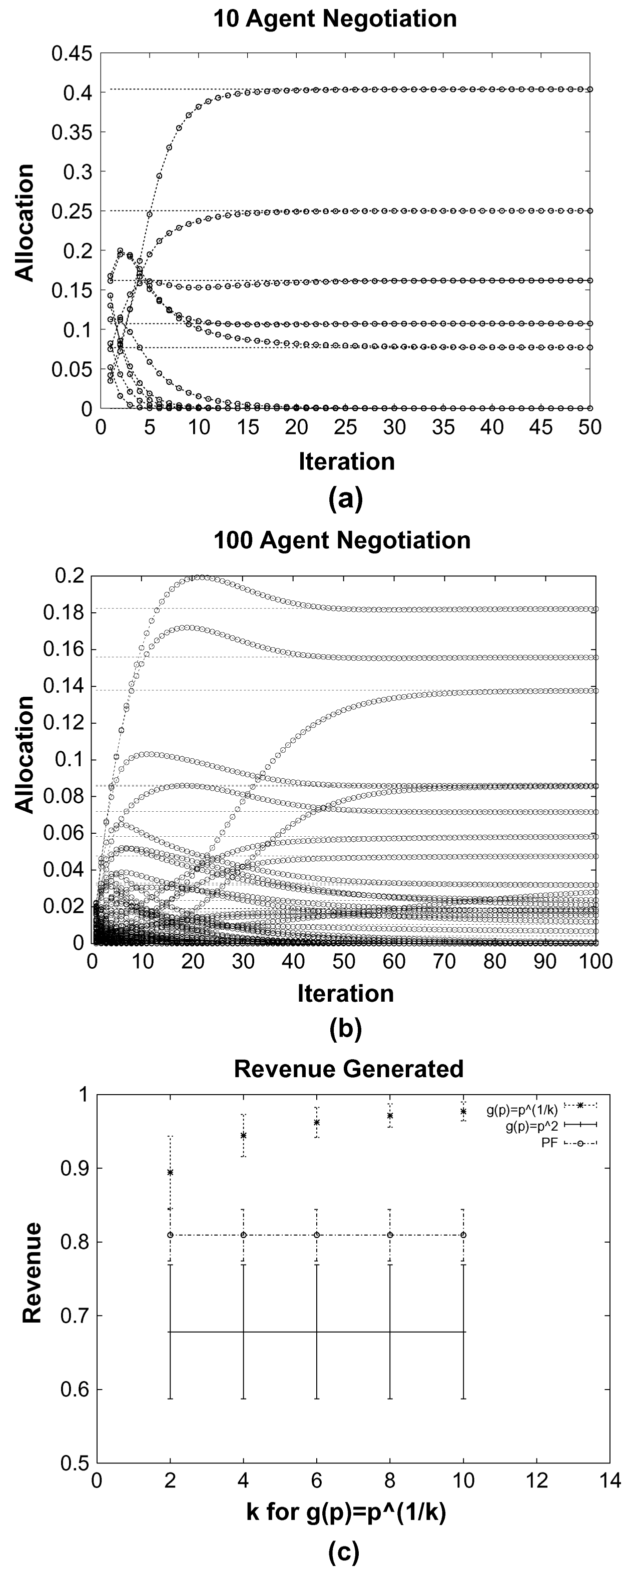

Fig. 2. (a) Ten agent negotiation. (b) 100 agent negotiation. (c) Mean and standard deviation of revenue generated for 100 runs of various mechanisms.

To illustrate the revenue limits, we obtained the equilibrium signal values, computed the corresponding costs which were then divided by λ^* for that particular collection so that we could aggregate data over various collections. For each mechanism described above, we ran 100 scenarios with 10 agents. The means

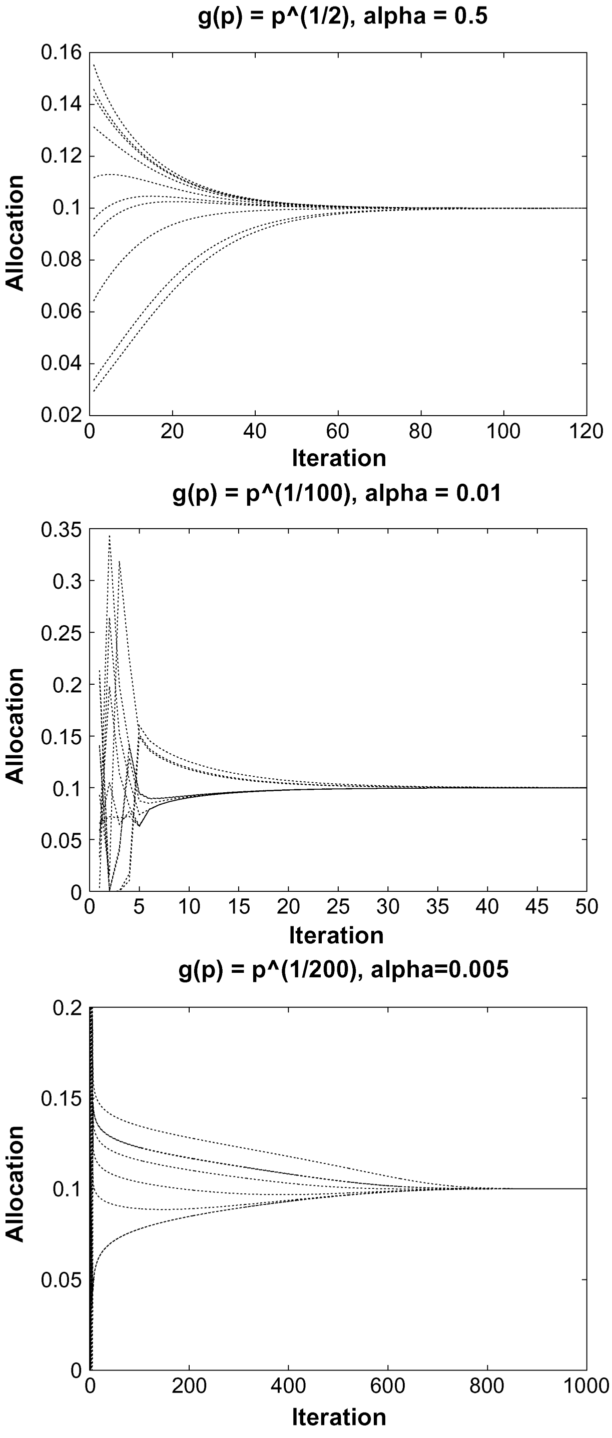

Fig. 3. Sample evolution of ESPA-$p^{1/k}$ for $k = 2, 100, 200$ with α that guarantees local stability.

and standard deviations of the revenue generated at equilibrium is displayed in Fig. 2(c). There exist ESPA mechanisms $(g(p) = p^2)$ that generate less revenue than the PF auction, but the mechanisms corresponding to the generator set $g(p) = p^{1/k}$ do indeed approach the limit λ^*, which was never exceeded at equilibrium. We note that this limit is exceeded in transience. In fact the ESPA-p^2 mechanism which performed worst with respect to revenue generation in the set considered here, actually generated significantly more revenue in transience than the others. Understanding transience is a key area for future research.

Finally, we illustrate the tradeoff incurred when attempting to approach the revenue limit. While using a relaxation parameter greater than the one specified in Proposition 8 often leads to a stable evolution, to guarantee local stability, we require α's smaller than the prescribed values. For the valuations we considered, this requirement translates to $\alpha < 2(b_i/a_i - 1)/k$. We show the evolutions of systems with ten agents, where $b_i = 3$ and $a_i = 2$ for all agents under the ESPA-$p^{1/k}$ mechanisms, where $k \in \{2, 100, 200\}$ in Fig. 3. While our supposition that convergence is delayed for large k is supported by the $k = 200$ case, we see that $k = 100$ actually converges faster that $k = 2$. Understanding the relationship between relaxation (or decentralized algorithms) and convergence is another key area for further study.

VIII. Conclusion

The main contribution of this paper is the development of a method to construct an infinite class of mechanisms that maximize social welfare under the lowest possible signaling and computational costs for auctioning a divisible resource. The key factor was coupling signals in the cost rule as well as the allocation rule. Open problems for future work include transforming these schemes to domains with different signal-cost-allocation restrictions (e.g., signals must be proportional to cost) and extending to divisible resources with alternate properties (e.g., excess demand spills over into finite buffers). System designers can now address secondary performance measures, while maintaining efficiency by optimizing over this class. We illustrate this by deriving a revenue limit for the ESPA class and identifying mechanisms that approach the limit arbitrarily closely. We also provide a locally stable negotiation protocol to reach an operating point, which also identifies a possible tradeoff between revenue and rate of convergence.

Because of transparency, efficiency, and minimal overhead, one might surmise that the auctions presented here are "ideal" allocation mechanisms. However, these mechanisms are not strictly incentive compatible, though calculating how to exploit this is an open area for research. Furthermore, as for all mechanisms, the effects of cooperative (or collusive) phenomena and other higher-degree responses must be investigated to fully understand the consequences of implementing these systems.

References

[1] L. Peterson and B. S. Davie, *Computer Networks*, 1st ed. San Francisco, CA: Morgan Kauffman, 1996.

[2] R. Wilson, "Auctions of shares," *Quart. J. Economics*, vol. 93, pp. 675–689, 1979.

[3] H. R. Varian, "Economic mechanism design for computerized agents," in *Proc. 1st Usenix Conf. Electronic Commerce*, New York, Jul. 1995, pp. 13–21.

[4] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *J. Finance*, vol. 16, pp. 8–37, 1961.

[5] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, vol. 11, pp. 17–33, 1971.
[6] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, pp. 617–631, 1973.
[7] N. Nisan and A. Ronen, "Computationally feasible VCG mechanisms," in *Proc. ACM Conf. Electronic Commerce*, Minneapolis, MN, Oct. 2000, pp. 242–252.
[8] J. K. MacKie-Mason and H. R. Varian, "Pricing the internet," in *Public Access to the Internet*, B. Kahin and J. Keller, Eds. Cambridge, MA: MIT Press, 1995, pp. 269–314.
[9] N. Semret, "Market mechanisms for network resource sharing," Ph.D. dissertation, Columbia Univ., New York, 1999.
[10] 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.
[11] H. Yaiche, R. R. Mazumdar, and C. Rosenberg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 667–678, Oct. 2000.
[12] R. Johari and J. Tsitsiklis, "Efficiency loss in a network resource allocation game," *Math. Oper. Res.*, vol. 29, no. 3, pp. 407–435, 2004.
[13] S. Sanghavi and B. Hajek, "Optimal allocation of a divisible good to strategic buyers," in *Proc. 43rd IEEE Conf. Decision and Control*, Paradise Island, Bahamas, Dec. 14–17, 2004.
[14] R. T. Maheswaran and T. Başar, "Social welfare for selfish agents: Motivating efficiency for divisible resources," in *Proc. 43rd IEEE Conf. Decision and Control*, Paradise Island, Bahamas, Dec. 14–17, 2004.
[15] A. K. Parekh and R. G. Gallagher, "A generalized processor sharing approach to flow control in integrated services networks: The single-node case," *IEEE/ACM Trans. Netw.*, vol. 1, no. 3, pp. 344–357, Jun. 2004.
[16] R. T. Maheswaran and T. Başar, "Nash equilibrium and decentralized negotiation in auctioning divisible resources," *Group Decision and Negotiation*, vol. 12, no. 5, pp. 361–395, Sep. 2003.
[17] ——, "Coalition formation in proportionally fair divisible auctions," in *Proc. 2nd Int. Joint Conf. Autonomous Agents and Multi-Agent Syst.*, Melbourne, Australia, Jul. 2003, pp. 25–32.

Rajiv Maheswaran received the B.S. degree in applied mathematics/physics/engineering from the University of Wisconsin-Madison, in 1995, and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois at Urbana–Champaign, Urbana, in 1998 and 2003, respectively.

He is a Research Scientist at the Information Sciences Institute, University of Southern California, Marina Del Rey. His research interests span various aspects of multi-agents systems including resource allocation, communication networks and distributed artificial intelligence, focusing on game-theoretic and decision-theoretic frameworks and solutions.

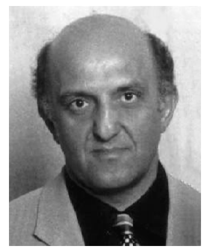

Tamer Başar (S'71–M'73–SM'79–F'83) received the B.S.E.E. degree from Robert College, Istanbul, and the M.S., M.Phil., and Ph.D. degrees in engineering and applied science from Yale University, New Haven, CT.

After working at Harvard University, Cambridge, MA, and Marmara Research Institute, Gebze, Turkey, he joined the University of Illinois at Urbana–Champaign in 1981, where he is currently the Fredric G. and Elizabeth H. Nearing Professor of Electrical and Computer Engineering, and Center for Advanced Study Professor. He has published extensively in systems, control, communications, and dynamic games, and has current interests in modeling and control of communication networks, control over heterogeneous networks, resource management and pricing in networks, and robust identification and control.

Dr. Başar is a member of the National Academy of Engineering, a member of the European Academy of Sciences, a Fellow of the International Federation of Automatic Control (IFAC), a Past President of the IEEE Control Systems Society (CSS), and the founding President of Annals of the International Society of Dynamic Games (ISDG). He has received several awards and recognitions over the years, among which are the Medal of Science of Turkey (1993), the Distinguished Member Award (1993), Axelby Outstanding Paper Award (1995) and Bode Lecture Prize (2004) of the IEEE CSS, and the Triennial Quazza Medal (2005) of IFAC. He is the Editor-in-Chief of *Automatica*, Editor of the *Birkhäuser Series on Systems & Control*, Managing Editor of the *Annals of the International Society of Dynamic Games* (ISDG), and member of editorial and advisory boards of several international journals.