COORDINATING LOCALLY CONSTRAINED AGENTS USING AUGMENTED PRICING

RINALDO A. JOSE, PATRICK T. HARKER, AND LYLE H. UNGAR

Abstract. The growing interest in distributed and agent-based computing has drawn researchers to auctions and market-based approaches for optimally coordinating self-interested agents—that is, pricing system resources so that the amounts the agents buy correspond to what a central planner would optimally select. Typically, researchers use linear prices in theoretical and applied auction work. However, linear prices fail to coordinate certain classes of agents—for example, agents which are locally constrained. Jennergren noticed this failure in linear programs (LPs) and suggested a form of augmented pricing as a solution. Although his proof only applies to agents described by LPs, augmented pricing can still be effective when agents are described by nonlinear programs (NLPs). In this paper, we extend Jennergren's result by showing that augmented prices can optimally coordinate agents if they have concave objectives and compact, convex feasible sets; we illustrate our points using numerical examples.

The authors express their gratitude to the NSF (SBR 9708965) for partial support of this project.

Rinaldo A. Jose, Dept. of Chemical Engineering, University of Pennsylvania, Philadelphia, PA 19104-6393, joser@eniac.seas.upenn.edu.

Patrick T. Harker, Dept. of Operations & Information Management, University of Pennsylvania, Philadelphia, PA 19104-6366, harker@wharton.upenn.edu.

Lyle H. Ungar, Dept. of Computer and Information Sciences, University of Pennsylvania, Philadelphia, PA 19104-6389, ungar@cis.upenn.edu.

Introduction

Many of the decomposition techniques in the 1960s and 1970s were motivated by an inability to solve large-scale, centralized problems with the computing power of that time [1]. The dramatic improvement in computing technology since then has allowed operations researchers to solve very large problems with relative ease. Consequently, interest in decomposition techniques has waned.

There is, however, an increasingly important class of problems for which decomposition techniques are once again becoming more relevant. Along with the increase in computing power, there has also been a move towards highly distributed computing—the Internet, perhaps, being the most visible example. Artificial Intelligence (AI) researchers often view distributed computer networks as systems of agents which draw on shared system resources [2, 3, 4]. In principle, one can optimize such a system of agents by constructing a large-scale mathematical program and solving it centrally using the currently available computing power and solution techniques. In practice, however, this is often impossible. In order to solve a problem centrally, one needs complete agent information such as local objective functions and constraints. When agents are separated geographically, this information may be unobtainable or prohibitively expensive to retrieve. More importantly, agents may be unwilling to share or report their private information as it is not incentive compatible to do so; i.e., agents may have an incentive to misrepresent their true preferences.

In order to optimize these distributed systems, one must turn to the coordination aspects of decomposition research. Specifically, with limited information one must coordinate agents to an optimal solution. Recently, market-based or price-directed strategies have received much attention [5, 6, 7]. Here, one identifies the resources over which agents compete and then charges each agent based on its resource usage. Such systems can be made incentive-compatible [8]. The goal is to coordinate agents by finding equilibrium resource prices.

Researchers have tended to focus on linear pricing where the cost per unit of resource is constant and the same for everyone in the system. Although this type of pricing has the benefits of fairness and simplicity, it often fails when agents have local constraints. For instance, there is no set of prices which can coordinate agents described by linear programs (LPs). Jennergren solved the particular case of coordinating LP agents by proposing a form of nonlinear or augmented pricing and proving that equilibrium prices exist [9]. Jennergren's proof applies only to LP agents, but augmented pricing should still be useful for coordinating agents of a more general type.

This paper presents a proof that augmented pricing can optimally coordinate systems of agents with concave objectives and compact, convex feasible sets (systems of LP agents being a special case). We describe why augmented pricing pricing works, and we illustrate our points with numerical examples. In closing, we discuss the relationship between the Dantzig-Wolfe decomposition technique and price-directed methods in general.

The Need for Augmented Pricing

One of the most common failures of linear pricing occurs, curiously enough, when agents are described by linear programs. Consider the following definition of an LP agent.

Definition 1 (LP Agent)**.** Suppose there are m system-wide resources. When LP Agent j engages in an activity level $x^j \in \Re^{n_j}$ (n_j being the number of Agent j's activities), it uses $R^j(x^j) = A^j x^j$ of the system-wide resources ($R^j \in \Re^m$ being the resource usage vector and A^j being an $m \times n_j$ matrix which maps activity level into resource usage). Agent j chooses its activity level at a given price vector $p \in \Re^m_+$ by solving the following linear program:

$$\begin{array}{ll} \text{Maximize} & {c^j}^T x^j - p^T R^j(x^j) \\ \text{subject to:} & \\ & x^j \in X^j \end{array}$$

where
c^j is Agent j's objective coefficient vector,
$X^j = \{x^j : B^j x^j \leq b^j\}$ is Agent j's feasible set,
B^j is an $m_j \times n_j$ matrix (m_j being the number of local constraints), and
$b^j \in \Re^{m_j}$.

□

At any set of linear prices, the resource demands, $r = R^j(x^j)$, of each LP Agent j are defined by the vertices of a polyhedral set—the optimal resource demands will never be defined at points in the interior of an agent's feasible space. Associated with a system of LP agents is a composite LP whose feasible space is defined by the system-wide resource constraints along with all of the LP agents' local constraints. The system-wide solution occurs at one of the extreme points of the system's feasible region. From an agent's point of view, the system solution sits, in general, in the interior of that agent's feasible simplex since it will have been formed from a constraint set more restricted than the agent's. Therefore, there is no set of linear prices which can drive the agents to select the optimal solution. Price adjustment will only move the agents' resource demands from vertex to vertex (Fig. 1).

To illustrate, consider the following example from Bradley et al.[10],

Example 1. Two divisions in a large firm represented by LP agents compete for system-wide resources: seven units of working capital (Resource 1) and sixteen units of raw materials (Resource 2).

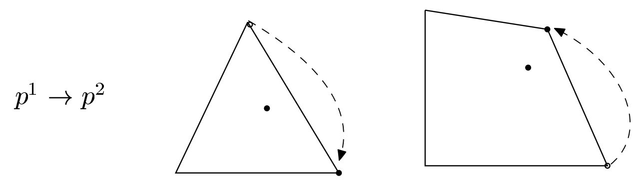

FIGURE 1. Linear prices cannot support interior points

TABLE 1. Agents in Example 1

Agent 1	Agent 2
$\max 2.5x_1 + 1.75x_2 + .75x_3$	$\max 3y_1 + 2y_2$
$4x_1 + 4x_2 + 5x_3 \leq 20$	$2y_1 + y_2 \leq 6$
$4x_1 + 2x_2 \leq 8$	$y_1 + y_2 \leq 4$
$x \geq 0$	$y \geq 0$
$r_1(x) = x_1 + x_2 + x_3$	$r_1(y) = y_1 + y_2$
$r_2(x) = 3x_1 + 2x_2 + 4x_3$	$r_2(y) = 5y_1 + 2y_2$

A centralized authority must allocate these resources optimally between the LP agents. The relevant information for this example is given in Table 1. (Note that the agents' "revenue functions" are shown—objective functions in the absence of resource costs), and the global solution is given in Table 3. Not surprisingly, an auctioneer using standard Walrasian tâtonnement to adjust resource prices is unable to find any set of equilibrium prices (Fig 2).

□

Trying to find coordinating prices is more difficult when agents have nonlinear features:

Example 2. Consider the system described in Example 1, but modify Agent 1's feasible space slightly by making one of its constraints nonlinear (see Table 2). Linear pricing is still unable to coordinate this system. Attempts to find prices fail to converge (see Fig 3).

□

To coordinate LP agents, Jennergren proposed augmented resource pricing of the form

$$p^+(R^j) = p + qR^j, \tag{1}$$

where p is the original linear price vector, q is some small, fixed constant, and R^j is the resource demand vector for Agent j. The price an agent pays for resources depends on how much of the resource it uses. This form of pricing can coordinate LP agents; we illustrate this in the following example.

Example 3. Consider the problem described in Example 1 and apply augmented pricing to it with $q = 0.01$. Adjusting prices in proportion to the excess demand (using a proportionality factor

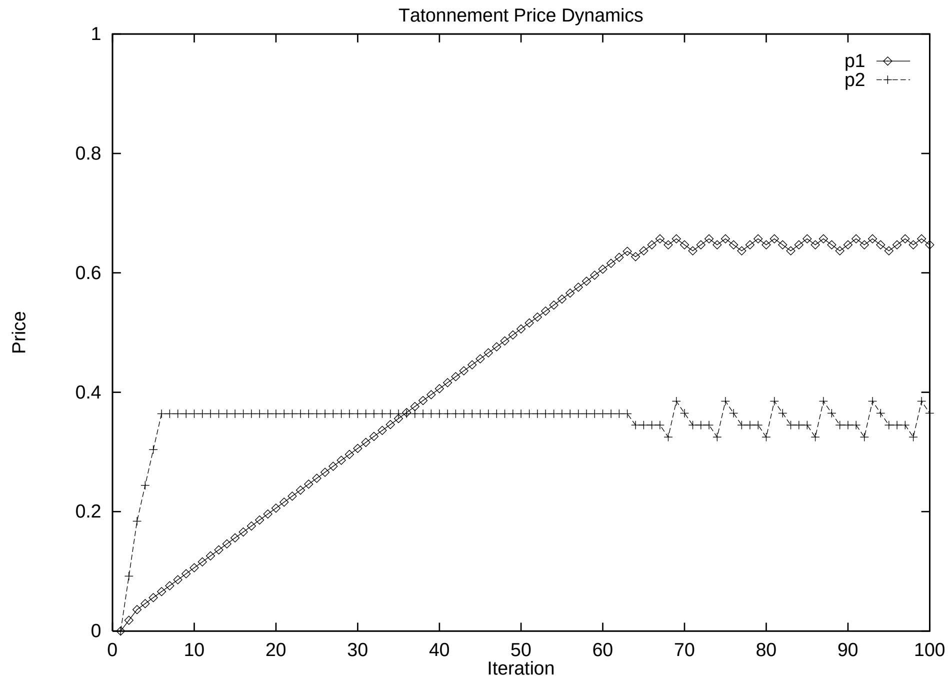

FIGURE 2. Linear pricing fails on Example 1

TABLE 2. Agents in Example 2

AGENT	OBJECTIVE (MAX)	CONSTRAINTS	RESOURCE USE
Agent 1	$2.5x_1 + 1.75x_2 + .75x_3$	$4x_1 + 4x_2 + 5x_3 \leq 20 + (4x_1 + 4x_2 + 5x_3)^{1/4}$ $4x_1+2x_2 \leq 8+(4x_1+2x_2)^{1/2}$ $x \geq 0$	$r_1(x) = x_1 + x_2 + x_3$ $r_2(x) = 3x_1 + 2x_2 + 4x_3$
Agent 2	$3y_1 + 2y_2$	$2y_1 + y_2 \leq 6$ $y_1 + y_2 \leq 4$ $y \geq 0$	$r_1(y) = y_1 + y_2$ $r_2(y) = 5y_1 + 2y_2$

TABLE 3. Solution to Example 1

AGENT	RESOURCE 1	RESOURCE 2	GROSS PROFIT
Agent 1	3	7	6
Agent 2	4	9	8.33
Total	7	16	14.33

K_p=0.01) results in the price dynamics shown in Fig. 4. The final equilibrium prices coordinates the LP agents to the optimal solution (see Table 3).

□

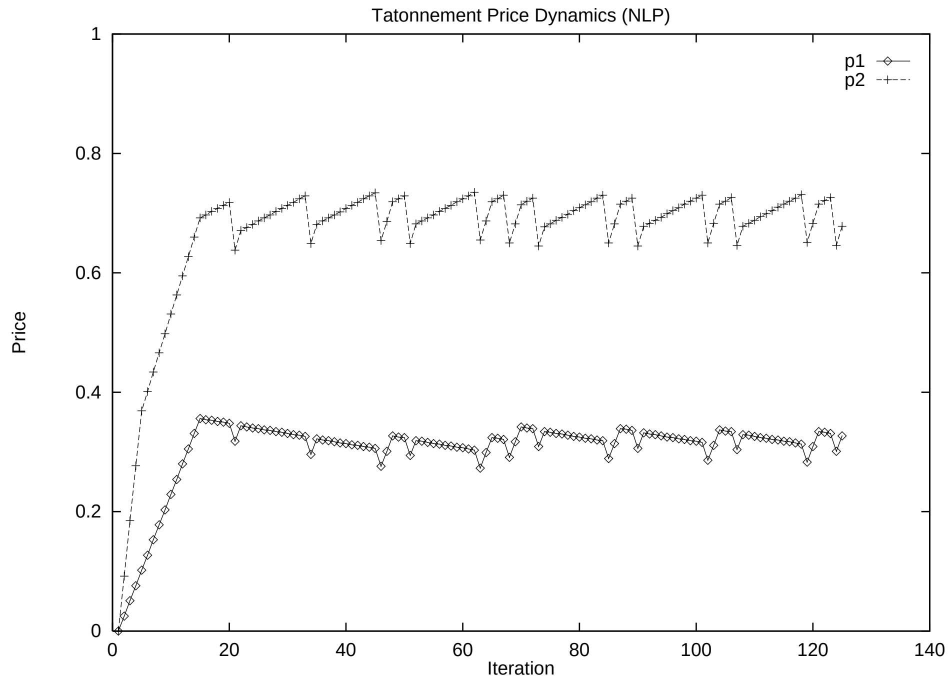

FIGURE 3. Linear pricing fails on Example 2

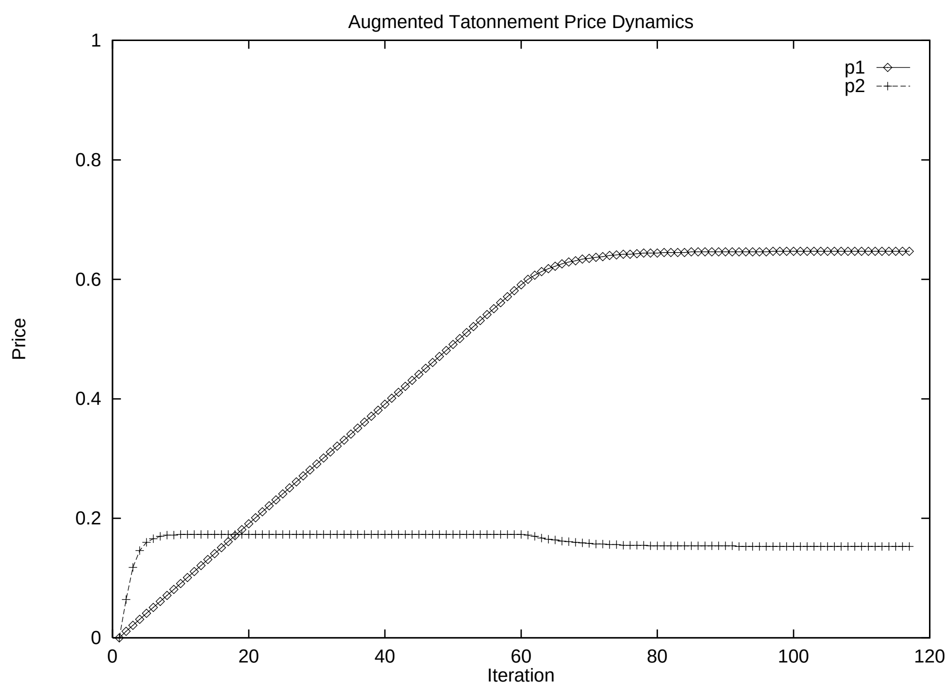

FIGURE 4. Augmented pricing solves Example 1

TABLE 4. Solution to Example 2

AGENT	RESOURCE 1	RESOURCE 2	GROSS PROFIT
Agent 1	3	8	6.75
Agent 2	4	8	8
Total	7	16	14.75

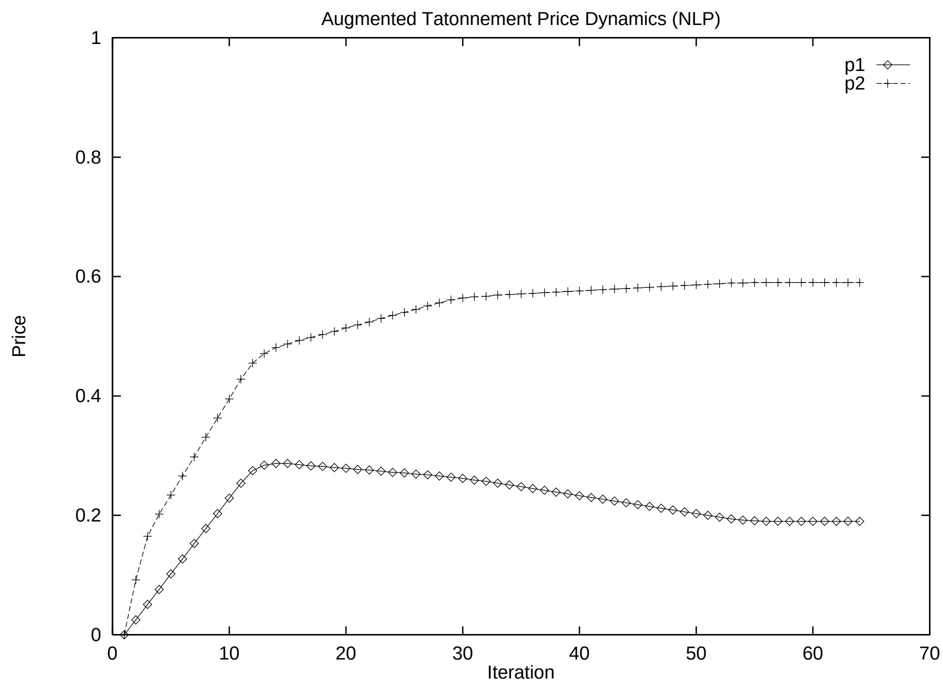

FIGURE 5. Augmented pricing solves Example 2

Although Jennergren's proof does not apply to non-LP agents, augmented pricing can still be effective:

Example 4. Consider the problem described in Example 2. Apply augmented pricing with $q = 0.01$ and adjust prices in proportion to the excess demand using a proportionality factor K_p=0.01. The resulting price dynamics are shown in Fig 5. As in Example 3, the final equilibrium prices coordinate the agents to the system-wide solution (see Table 4).

□

EXISTENCE CONDITIONS FOR AUGMENTED PRICES

Thus far, we have discussed augmented pricing somewhat informally and only with respect to specific examples. Here we develop the conditions under which coordinating augmented prices exist for systems of agents described by nonlinear programs (NLPs). Consider the definition of the NLP agent which we will use throughout this section:

Definition 2 (NLP Agent)**.** Suppose there are m system-wide resources. When NLP Agent j engages in an activity level $x^j \in \Re^{n_j}$ (n_j being the number of Agent j's activities) it uses $R^j(x^j)$ of the system-wide resources ($R^j(x^j) \in \Re^m_+$ being the resource usage vector). Agent j chooses its activity level at a given price vector $p \in \Re^m_+$ and small scalar $q > 0$ by solving the following nonlinear program:

$$\max_{x^j \in X^j} \qquad f_j(x^j) - \left(p + qR^j(x^j)\right)^T R^j(x^j) \tag{2}$$

where
f_j is a concave continuous "revenue function",
$X^j = \{x : g(x) \geq 0\}$ is Agent j's convex, compact feasible set; $0 \in X^j$, and
$R^j_i(x^j)$ is convex and continuous in x^j $\forall i = 1, \ldots, m$; $R^j_i(0) = 0$.

□

The main result of this paper is the following proposition.

Proposition 1. *Let X^{j*} be the set of optimal solutions to Eq. 2 for NLP agent j. Then, $\forall \overline{R} \in \Re^m_+$ and $q \in \Re_+$, $\exists$ prices $p^* \in \Re^m_+$, such that $X^{j*} = x^{j*}$ (a singleton) and $\sum_j R^j(x^{j*}) - \overline{R} \leq 0$.*

□

In other words, if agents have concave, continuous revenue functions and compact, convex feasible sets, there will exist equilibrium augmented prices of the form $p + qR$ which optimally coordinate the agents given the system resources $\overline{R}$.

Our first goal is to show that the resources demanded by an NLP agent at price $p + qR$ are unique and continuous in p for a given q. This allows us to cast the search for augmented prices as a nonlinear complementarity problem (NCP) [11] for which there is rich theoretical support.

To begin, we show that the addition of augmented resource pricing, $p + qR$, to NLP agent j's nominal objective function, $f_j(x^j)$, yields a strictly concave function. Agent j's overall objective (the maximand of Eq. 2) can be written as:

$$\begin{gathered} f_j(x^j) - (p + qR^j(x^j))^T R^j(x^j) = \\ f_j(x^j) - p^T R^j(x^j) + qR^j(x^j)^T R^j(x^j) \end{gathered} \tag{3}$$

or

$$f_j(x^j) - p^T R^j(x^j) - q||R^j(x^j)||^2 \tag{4}$$

The first term of the right-hand side of Eq. 4 is concave by assumption; the second term is concave because $p \geq 0$ and each component of $R^j(x^j)$ is convex. To show that the NLP agent's objective function is strictly concave, it is then sufficient to show that the third term of Eq. 4 is strictly concave.

Lemma 1. *Suppose $g(x) : \Re^m \mapsto \Re_+$ is convex in x, and $\Psi : \Re_+ \mapsto \Re$ is nondecreasing, proper, and strictly convex. Then $\Psi(g(x))$ is strictly convex in x.*

Proof. Since $g(x)$ is convex,

$$g(\theta x + (1-\theta)y) \leq \theta g(x) + (1-\theta)g(y), \text{ for } \theta \in (0,1). \tag{5}$$

Since Ψ is nondecreasing, the above equation implies

$$\Psi\big(g(\theta x + (1-\theta)y)\big) \leq \Psi\big(\theta g(x) + (1-\theta)g(y)\big). \tag{6}$$

But since Ψ is strictly convex in g,

$$\Psi\big(\theta g(x) + (1-\theta)g(y)\big) < \theta\Psi\big(g(x)\big) + (1-\theta)\Psi\big(g(y)\big). \tag{7}$$

And so, from Eqs. 6 and 7 we have

$$\Psi\big(g(\theta x + (1-\theta)y)\big) < \theta\Psi\big(g(x)\big) + (1-\theta)\Psi\big(g(y)\big). \tag{8}$$

Hence, $\Psi(g(x))$ is strictly convex in x as well. □

Now we can easily show the following:

Lemma 2. *Let $X^j \subseteq \Re^{n_j}$ be nonempty and $R^j(x^j) : X^j \mapsto \Re^m_+$. Suppose for all resources i and for all agents j, $R^j_i(x^j)$ is convex in x^j. Then $||R^j(x^j)||^2$ is strictly convex in x^j for all agents j.*

Proof. The squared norm of $R^j(x^j)$ can be written as

$$||R^j(x^j)||^2 = \sum_i R^{2j}_{\ i}(x^j). \tag{9}$$

Since the sum of strictly convex functions is strictly convex, it is sufficient to show that the square of each component of $R^j(x^j)$ is strictly convex in x^j. Notice that the square operation over the positive orthant is strictly convex and nondecreasing. Since $R^j_i(x^j)$ is a convex function whose range lies in the positive orthant, we can invoke Lemma 1 with $\Psi(x) = x^2$ for $x \geq 0$ and with $g(x) = R^j_i(x)$ to conclude that $R^{2j}_{\ i}(x^j)$ is strictly convex in x^j. The result of the lemma follows immediately. □

Thus, an NLP Agent that faces augmented resource pricing has a strictly concave objective function. Next, we use the following result to show that each agent's resource demand exists and is unique for a given p.

Lemma 3 ([12]). *Consider the following parameterized maximization problem:*

$$\max_{x \in G(p)} f(x, p).$$

1. *If the constraint set $G(p)$ is nonempty and compact, and the function f is continuous, then there exists a solution x^* to the above problem.*
2. *If the function f is strictly concave, and the constraint set is convex, then a solution, should it exist, is unique.*

□

Since Agent j's feasible set X^j is nonempty, compact and convex, and the maximand of Eq. 2 is both continuous and strictly concave in x^j, Lemma 3 implies that the optimal activity level, $x^*(p;q)$, exists and is unique. Therefore, it is possible to identify $R^j(x^{j*}(p;q))$ as Agent j's demand function, $R^{j*}(p;q)$. Consequently, the aggregate resource demand for a set of n NLP agents at a given augmented resource price may be defined as

$$R^*(p;q) = \sum_{j=1}^{n} R^{j*}(p;q) \tag{10}$$

We can show that this aggregate demand function is continuous in p using the following result.

Lemma 4 ([12]). *Let $f(x;p)$ be a continuous function with a compact range and suppose that the constraint set $G(p)$ is a nonempty, compact-valued, continuous correspondence of p. Then $x(p)$ is a continuous correspondence.*

□

Invoking this lemma for an NLP agent, implies that $x^{j*}(p;q)$ is continuous in p, and consequently that $R^{j*}(p;q) = R(x^{j*}(p;q))$ is continuous in p due to the continuity of $R^j(x)$. Since the sum of continuous functions is continuous, the aggregate demand function, $R^*(p;q)$ is continuous as well.

We are now in the position to formulate the augmented pricing procedure as a nonlinear complementarity problem and to show the existence of an equilibrium price. The goal of augmented tâtonnement is to find a price vector p for some small q such that the excess resource demand, $\Delta(p;q) = R^*(p;q) - \bar{R}$, is less than or equal to zero, where $\bar{R}$ is the resource supply vector. The associated NCP is

NCP 1 (Augmented tâtonnement). Find $p^* \in \Re^m_+$ such that

$$-\Delta(p^*;q) \geq 0$$
$$p^{*T}(-\Delta(p^*;q)) = 0$$

□

The next lemma states that if no activity is an option for all the agents, there will always be a set of augmented prices high enough so that doing nothing is optimal.

Lemma 5. *If* $0 \in X^j$ *for all agents* j *as defined in this section, then there exists some price vector* $\bar{p} < \infty$ *such that* $\Delta(p;q) \leq 0$, $\forall p \geq \bar{p}$.

Proof. Since X^j is compact for all agents $j = 1, \ldots, n$, there exists $\underline{f}, \overline{f} \in \Re$ such that

$$-\infty < \underline{f} \leq f(x^j) \leq \overline{f} < \infty \qquad \forall x^j \in X^j, \tag{11}$$

where $f(x^j)$ is Agent j's nominal objective or revenue function. Suppose $R_i^j(x^j) = \epsilon > 0$. Let $f^* = \max(|\overline{f}|, |\underline{f}|)$ and let $\epsilon^* = \min(\bar{R}_i/n, \epsilon)$. Choose $\bar{p}^j$ such that $\bar{p}_i^j = f^*/\epsilon^*$. Then for Agent j,

$$f_j(x^j) - p^T R^j(x^j) - q||R^j(x^j)||^2 < \underline{f}. \tag{12}$$

But for $x^j = 0$, $R^j = 0$, and the left hand side of Eq. 11 is at least $\underline{f}$. Therefore, at price $\bar{p}^j$, $x^j = 0$ is preferred to any element of the set $A = \{x : R_i^j(x) = \epsilon\}$. Furthermore, because $\epsilon^* \leq \bar{R}^i/n$, any element of the set $B = \{x : R_i^j(x) < \bar{R}_i/n\}$ is preferred to any element of the set A (since $f_j(x^j)$ is concave and $R_i^j(x^j)$ is convex).

There are now two possibilities:

- None of the elements of the set $C = \{x : 0 < R_i^j(x) < \bar{R}_i/n\}$ is feasible, and so the agent must choose $x^j = 0$ at price $\bar{p}^j$.
- At least one element of the set C is feasible, and so the agent can choose x^j from the feasible points of C, or it can choose $x^j = 0$.

In either case, the agent's demand for Resource i is less than or equal to $\bar{R}_i/n$ for any price p such that $p_i > \bar{p}_i^j$. Now take all the agents together and choose $\bar{p}$ such that $\bar{p}_i = \max(\bar{p}_i^1, \ldots, \bar{p}_i^n)$. Then,

$$\sum_{i=i}^{n} R_i^j \leq \bar{R}_i \qquad \Longrightarrow \qquad \Delta_i(\bar{p};q) \leq 0 \tag{13}$$

But since this argument was independent of the other resources, we can choose the rest of the components of $\bar{p}$ in the same way. Thus, one can construct $\bar{p}$ such that for any $p > \bar{p}$, $\Delta(p;q) \leq 0$. □

Now we can invoke the following result from Harker:

Lemma 6 ([13]). *If* $\Delta(p;q)$ *is a continuous function over* $\Re_+^m$ *and there exists a price* $\bar{p}$ *strictly less than infinity such that* $\Delta(p;q) \leq 0$, $\forall p > \bar{p}$, *then* $\exists p^*$ *that solves NCP1* $\forall q > 0$.

□

Therefore, for sufficiently small q, there exists a set of augmented prices which can optimally coordinate a system of NLP agents as defined in this section. Thus, our main proposition is proved.

Note that the resource cost function,

$$(p + qR(x))^T R(x) = p^T R(x) + q||R(x)||^2, \tag{14}$$

is only one member of a family of augmented prices for which the results in this section apply. There is nothing particularly special about the squared norm of the resource use; what *is* important is that it is strictly convex in x. Since each component of $R(x)$ is convex, any function Ψ that is strictly convex and non-decreasing over the positive orthant can be used to generate another member of the augmented pricing family, $p^T R(x) + \Psi\big(R(x)\big)$ (cf. Lemma 1).

Discussion

Why Augmented Pricing Works. Consider the objective function of an LP Agent j (ignoring resource costs for now):

$$u_j(x^j) = c^{j\,T} x^j \tag{15}$$

where x^j is the decision vector and c^j the objective coefficient vector. Since the resource used by Agent j is

$$R^j(x^j) = A^j x^j, \tag{16}$$

the objective in terms of resource usage is

$$u_j(R^j) = c^{j\,T} \left((A^{j\,T} A^j)^{-1} A^{j\,T} R^j \right). \tag{17}$$

Subtracting the resource cost yields

$$\begin{aligned} u_j(R^j) &= c^{j\,T}(A^{j\,T}A^j)^{-1}A^{j\,T}R^j - (p + qR^j)^T R^j \\ &= \left(c^{j\,T}(A^{j\,T}A^j)^{-1}A^{j\,T} - p^T \right) R^j - qR^{j^T} R^j. \end{aligned} \tag{18}$$

If we make the substitution

$$\bar{c}^j = c^{j\,T}(A^{j\,T}A^j)^{-1}A^{j\,T}, \tag{19}$$

it is clear that the objective can be viewed as the sum of linear and quadratic pieces (without cross terms):

$$u_j(R^j) = \underbrace{(\bar{c}^{j\,T} - p^T)R^j}_{\text{linear}} - \underbrace{qR^{j^T} R^j}_{\text{quadratic}} \tag{20}$$

During price adjustment, the quadratic surface remains fixed since q is fixed, but the hyperplane defined by $(\bar{c}^{j\,T} - p^T)R^j$ changes its tilt as p changes. If we restrict our attention to the positive orthant of the resource plane, an increase in prices tilts the hyperplane downwards, while a decrease tilts it upwards.

The combination of these linear and quadratic pieces allows the auctioneer to support points in the interior of Agent j's feasible space. Consider the case of a single system resource and select a price such that Agent j's resource constraint is non-binding. The optimal resource demand $\bar{R}^j$ occurs at the associated stationary point of Agent j's objective. Raising the resource price lowers hyperplane component of Agent j's objective causing its resource demand to move smoothly in towards the origin from the positive side. Lowering the resource price, on the other hand, moves demands smoothly away from the origin. Since the agent has local constraints restricting its activity level (and hence its resource use), there will be an upper bound on its resource demand. Thus, assuming the auctioneer controls all of the system resources at the start of the auction, the agent's resource demands will be non-negative. Therefore, by choosing the appropriate price, the auctioneer can make the agent respond with any point in its feasible space.

If, however, the price vector p is constrained to be nonnegative, the auctioneer will only be able to elicit a subset of feasible resource points from the agent—those sufficiently close to the origin. Consider the stationary point of $u_j(R^j)$ at prices p:

$$\left(\frac{\bar{c}_1^j - p_1}{2q}, \cdots, \frac{\bar{c}_m^j - p_m}{2q} \right) \tag{21}$$

Since lowering resource prices moves this stationary point away from the origin (in the positive orthant of the resource plane), the farthest demand point, R_0^j, that the agent can ever respond with occurs when the resource price vector is identically zero:

$$R_0^j(q) = \left(\frac{\bar{c}_1^j}{2q}, \cdots, \frac{\bar{c}_m^j}{2q} \right). \tag{22}$$

By increasing p, the only demand points the auctioneer can support are members of the box $B(q)$ defined by

$$B(q) = \{ r \in R | r \leq R_0^j(q) \} \tag{23}$$

where R is the resource plane (see Fig. 6). Therefore, if any of the agent's feasible points lie outside of this box, there may not be a price vector p, given this q, which can support the system solution. However, since the size of this box is inversely related to q, the auctioneer can always find some box $B(q^*)$ which contains the agent's feasible region by choosing q^* sufficiently small (see Fig. 7). Hence, there is some $q^{\max}$ for which augmented pricing is guaranteed to support all of the agent's feasible points.

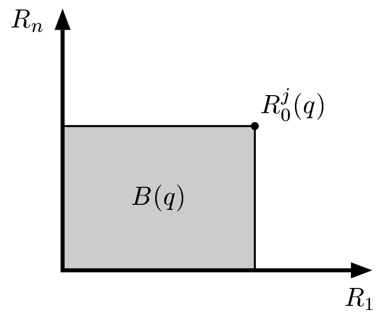

FIGURE 6. Auctioneer can only support points in $B(q)$

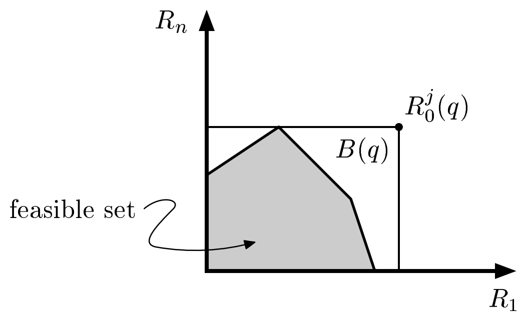

FIGURE 7. Smallest box $B(q)$ for which augmented pricing can support all feasible points

Although we have focused on an individual LP agent, from the auctioneer's perspective, a collection of LP agents is no different from a single one—the collective's objective still has linear and quadratic parts, and the feasible points still form a simplex. Therefore, for a set of LP agents, there will also be some $q^{\max}$ such that for any $q < q^{\max}$, there will exist augmented pricing capable of optimally coordinating the agents.

Better Ways to Adjust Prices. Though we have presented results which establish sufficient existence conditions for equilibrium resource prices, we have guaranteed nothing about reaching these prices. In all our examples, we have simply adjusted prices in proportion to the excess demand. In a sense, we have used proportional controllers [14] which manipulate resource prices in order to maintain an excess demand setpoint of zero. Although this works well enough for simple examples, this type of price adjustment may have problems with more complicated systems. In particular, one must be concerned about how the excess demand controllers are "tuned"—too rigorous a tuning will result in price oscillations and instability. If the center has enough information about the agents, one alternative to proportional adjustment might be homotopy-based price adjustment—this may be more intelligent and more reliable although probably not more efficient if proportional adjustment is satisfactory. Likewise, some form of randomized search might also prove to be a practical alternative if the auctioneer has sufficiently tight bounds on the resource prices.

Dantzig-Wolfe as a Price Directed Coordination Technique. At this point, one might object to the need for augmented pricing citing the Dantzig-Wolfe decomposition technique as an apparent counterexample. To be sure, Dantzig-Wolfe works by adjusting linear prices, but that part of the algorithm is only an information gathering stage. The final prices in Dantzig-Wolfe are not equilibrium prices. Dantzig-Wolfe identifies the system solution but must *impose* it on the agents.

The Dantzig-Wolfe decomposition locates the system-wide solution of a collection of LP agents by iteratively collecting agent vertex points and solving for the convex combination of them which has the greatest value. First, the center announces linear resource prices associated with each of the linking constraints. Based on these prices, agents maximize their net profit and transmit their gross profit and resource demands to the center. The center solves a master problem which maximizes the sum of the convex combinations of all bids reported so far and chooses the next resource prices to be the dual prices of the corresponding resource constraints. When these prices converge, the center has collected enough vertices to determine the optimal convex combination. The appropriate weighting coefficients come from the solution of the final master problem.

Although the Dantzig-Wolfe center can find the global solution, the center cannot support it. The final solution chosen by the Dantzig-Wolfe center is unnatural in the sense that it does not correspond to any resource price vector—it is simply some convex combination of the bids reported so far. The price "convergence" in Dantzig-Wolfe signals that no new bids have been generated by the agents, *not* that the agents' demands are in equilibrium in any sense. Moreover, since prices are non-zero, the global resource constraints *must* be binding, implying that the aggregate demand exceeds total supply. The Dantzig-Wolfe center is therefore compelled to impose a solution, but it cannot expect to enforce it. After-markets and agent-agent swaps would erode the global solution and redistribute resources sub-optimally. For instance, suppose there is a unique optimal allocation for the agents, and suppose the center makes this allocation under the final, non-zero, reported prices. If agents are not all identical, then there exists a pair of agents such that the marginal gain of a unit of resource for one agent exceeds the marginal loss for the other agent. Between these agents, a resource trade will be made and the optimal allocation will vanish.

Note that this failure of Dantzig-Wolfe to coordinate LP agents occurs because we assume that agents can communicate. In a standard Dantzig-Wolfe application, communication between agents is implicitly forbidden. The Dantzig-Wolfe center is able to impose its own solutions because the agents have no way of redistributing their allocations after the procedure finishes. When agents can exchange information, a price-directed resource allocation is stable only when equilibrium prices exist. Because there are no linear equilibrium prices for LP agents, Dantzig-Wolfe fails.

Conclusion

Although linear prices can often coordinate agents which are described by objective or utility functions only, equilibrium prices may not even exist when agents have local constraints. The simplest example of this problem occurs with LP agents. Although Jennergren treated this problem by defining a form of augmented pricing and proving the existence of coordinating prices, his proof relied heavily on the LP nature of the problem. In this paper, we have developed a more general proof which not only applies to systems of LP agents but also to systems of agents which have concave objectives and compact, convex feasible sets.

Coordination of these types of nonlinear agents often arises in practice: optimizing units in chemical plants, scheduling track time for trains, allocating computer resources for users, and so on. Because agents are typically self-interested and not willing to reveal private information, one cannot solve these kinds of coordination problems by simply maximizing the sum of the agents' utility functions with respect to their constraints. One must use methods which can operate with limited information. Auction-based techniques are ideally suited here because they only require resource demands—auctions can also be made incentive compatible to motivate the truthful revelation of those demands. Auctions, however, can only work successfully when coordinating prices exist. Showing when such prices exist has been the subject of this paper.

References

[1] A. M. Geoffrion, "Elements of Large-scale Mathematical Programming," *Management Science*, vol. 16, no. 11, pp. 652–691, 1970.

[2] T. Sandholm and V. Lesser, "Issues in automated negotiation and electronic commerce: Extending the contract net framework," in *Readings in Agents* (M. Huhns and M. Singh, eds.), Morgan Kaufmann, 1997.

[3] M. P. Wellman, "A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems," *Journal of Artificial Intelligence*, pp. 1–23, 1993.

[4] M. P. Wellman, "The economic approach to artificial intelligence," *ACM Computing Surveys*, vol. 27, pp. 360–362, Sept. 1995.

[5] R. L. Graves, L. Schrage, and J. Sankaran, "An auction method for course registration," *Interfaces*, vol. 23, no. 5, pp. 81–92, 1993.

[6] B. A. Huberman and T. Hogg, "Distributed computation as an economic system," *Journal of Economic Perspectives*, vol. 9, no. 1, pp. 141–152, 1995.

[7] G. Y.-J. Lin and J. J. Solberg, "Integrated shop floor control using autonomous agents," *IIE Transactions*, vol. 24, pp. 57–71, July 1992.

[8] S. J. Grossman and O. D. Hart, "An analysis of the principal-agent problem," *Econometrica*, vol. 51, pp. 7–45, Jan. 1983.

[9] L. P. Jennergren, "A price schedules decomposition algorithm for linear programming problems," *Econometrica*, vol. 41, pp. 965–980, Sep 1973.

[10] S. P. Bradley, A. C. Hax, and T. L. Magnanti, *Applied Mathematical Programming*. Addison-Wesley, 1977.

[11] P. T. Harker and J. S. Pang, "Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications," *Mathematical Programming*, vol. 48B, pp. 161–220, 1990.

[12] H. R. Varian, *Microeconomic Analysis*. W. W. Norton & Company, 1992.

[13] P. T. Harker, "Existence of competitive equilibria via smith's non-linear complementarity result," *Economics Letters*, vol. 19, pp. 1–4, 1985.

[14] C. A. Smith and A. B. Corripio, *Principles and Practice of Automatic Process Control*. John Wiley & Sons, Inc., 1985.