

Efficiency and Equilibrium in Task Allocation Economies with Hierarchical Dependencies

William E. Walsh Michael P. Wellman

Artificial Intelligence Laboratory

University of Michigan

1101 Beal Avenue, Ann Arbor, MI 48109-2110 USA

{wew, wellman}@umich.edu

Abstract

We analyze economic efficiency and equilibrium properties in decentralized task allocation problems involving hierarchical dependencies and resource contention. We bound the inefficiency of a type of approximate equilibrium in proportion to the number of agents and the bidding parameters in a particular market protocol. This protocol converges to an approximate equilibrium with respect to all agents, except those which may acquire unneeded inputs. We introduce a decommitment phase to allow such agents to decommit from their input contracts. Experiments indicate that the augmented market protocol produces highly efficient allocations on average.

1 Introduction

We consider task allocation problems in which competing agents desire to accomplish tasks, which may require complex chains of production activity. In order to perform a particular task, an agent may need to achieve some subtasks, which may in turn be delegated to other agents, forming a supply chain through a hierarchy of task achievement. Constraints on the task assignment arise from *resource contention*, where agents would need a common resource (e.g., a subtask achievement, or something tangible like a piece of equipment) to accomplish their own tasks. We assume that agents are self-interested and have private information, and so we must allocate in a decentralized manner.

We take a market-based approach to decentralized resource allocation, utilizing the large body of solution methods and analytical techniques from economics. Auctions mediate negotiation and determine prices and allocations. Prices indicate relative values of resources to guide local agent decisions. Experience with the *market-oriented programming* approach has verified that it works predictably and effectively in convex domains [Wellman, 1993]. Discrete problems, such as the task allocation domain considered here, provide additional challenges.

In previous work [1998], we proposed a market protocol that reliably constructs supply chains in a decentralized manner. In this paper we generalize the model and protocol to account for multiple competing demands for multiple end tasks. We bound the inefficiency of a type of approximate equilibrium in proportion to the number of agents and the bidding parameters in a particular market protocol. This protocol converges to an approximate equilibrium with respect to all agents, except those which may acquire unneeded inputs. We introduce a decommitment phase to allow such agents to decommit from their input contracts. Experiments indicate that the augmented market protocol produces highly efficient allocations on average.

We describe the task allocation problem in Section 2. We discuss price systems and competitive equilibrium in Section 3, and survey some market protocols for decentralized resource allocation in Section 4. In Section 5 we analyze the equilibrium properties of the extended protocol, and in Section 6 we examine its efficiency. We describe related task allocation work in Section 7.

2 Problem Description

Tasks are performed on behalf of particular agents; if two agents need a subtask then it would have to be performed twice to satisfy them both. In this way, tasks are the same as any other discrete resource. Hence we make no distinction in our model, and use the term “good” to refer to both.

We describe the problem in terms of bipartite graphs. A **task dependency network** is a directed, acyclic graph, (V, E) . The vertices are $V = G \cup A$, where G is the set of goods, $A = C \cup \Pi \cup S$ is the set of agents, C is the set of consumers, Π is the set of producers, and S is the set of suppliers. The edges, E , connects agents with goods they can use or provide. There exists an edge $\langle g, a \rangle$ from $g \in G$ to $a \in A$ when agent a can make use of one unit of g , and an edge $\langle a, g \rangle$ when a can provide one unit of g . When an agent can acquire or provide multiple units of a good, we represent each unit as a separate edge. The goods can be traded only in integer quantities.

A **consumer** wishes to acquire one unit of one good from a set of some high-level goods. A **producer** can

produce a single unit of an **output** good conditional on acquiring a certain number each of some fixed set of **input** goods. A producer's input requirements are **complementary** in that it must acquire each of its inputs; it cannot accomplish anything with only a partial set. A **supplier** can supply a set of goods, up to some maximum quantity for each, without requiring any input goods.

An **allocation** is a subgraph $(V', E') \subseteq (V, E)$. For $g \in G$, an edge $\langle a, g \rangle \in E'$ means that agent a provides g , and $\langle g, a \rangle \in E'$ means a acquires g . An agent is in an allocation graph iff it acquires or provides a good. A good is in an allocation graph iff it is bought or sold.

A producer is **active** iff it provides its output. A **producer is feasible** iff it is inactive or acquires all its inputs. Consumers and suppliers are always feasible. An **allocation is feasible** iff all producers are feasible and all goods are in **material balance**, that is the number of edges into a good equals the number of edges out.

A **solution** is a feasible allocation such that one or more consumers acquire a desired good. If $c \in C \cap V'$ for solution (V', E') , then (V', E') is a **solution for c** .

Each supplier s has some **opportunity cost** $oc_s(g)$ for supplying one unit $\langle s, g \rangle$ of good g . The total opportunity cost to s for allocation E' is $oc_s(E') \equiv \sum_{\langle s, g \rangle \in E'} oc_s(g)$.¹ The cost might represent the value s could obtain from putting the goods to some other use, or some actual, direct cost incurred in supplying the goods.

We assume that a consumer has preferences over different possible goods, but wishes to obtain only a single unit of one good.² Thus, a consumer c obtains **value** $v_c(g)$ for obtaining a single unit of good g , and, for allocation E' , obtains value $v_c(E') \equiv \max_{\langle g, c \rangle \in E'} v_c(g)$.

Definition 1 (value of an allocation) *The value of allocation (V', E') is:*

$$value((V', E')) \equiv \sum_{c \in C} v_c(E') - \sum_{s \in S} oc_s(E').$$

Definition 2 (efficient allocations) *The set of efficient allocations contains all feasible allocations (V^*, E^*) such that $value((V^*, E^*)) = \max_{(V', E') \text{ feasible}} (value((V', E')) \mid (V', E') \text{ is feasible})$.*

¹We overload the notation for opportunity cost and agent value, using as an argument either a single good or a set of edges.

²More complicated consumer value functions can be expressed through combinations of multiple consumers and producers. For instance, a group of consumers, each desiring a single separate good, effectively represents an additive value function over the set of goods. If a consumer has preferences over **bundles** of goods, the bundles can be represented as output goods of producers, with the inputs of a producer corresponding to the components of the respective output bundle. Note that transformations such as this may not be strategically equivalent, in that the group of consumers may not behave as would the single agent they represent.

One could find efficient allocations using centralized search techniques, but we assume that we are constrained to solve the problem in a decentralized fashion. In the following, we examine an abstract framework for how a price system can guide decentralized task allocation. We also examine market protocols for decentralized task allocation.

3 Price Systems and Competitive Equilibrium

In the general-equilibrium approach to economic resource allocation, we posit a **price system** p , which assigns to each good g a nonnegative number $p(g)$ as its **price**. Intuitively, prices indicate the relative global value of the goods. Therefore, agents may use the prices as a guide to their local decision making.

We assume each agent has a **quasilinear** utility function. Its utility is the sum of the “money” it holds and the value (or negative cost) obtained from its allocation of goods. Agents wish to maximize their **surplus**.

Definition 3 (surplus) *The surplus, $\sigma(a, E', p)$, of agent a with allocation E' at prices p , is the utility gain from E' , defined as follows:*

- $v_a(E') - \sum_{\langle g, a \rangle \in E'} p(g)$, if $a \in C$,
- $\sum_{\langle a, g \rangle \in E'} p(g) - \sum_{\langle g, a \rangle \in E'} p(g)$, if $a \in \Pi$,
- $\sum_{\langle a, g \rangle \in E'} p(g) - oc_a(E')$, if $a \in S$.

Informally, an allocation (V', E') is a **competitive equilibrium** at prices p if (V', E') is feasible and assigns to each agent an allocation that optimizes the agent's surplus at p . A competitive equilibrium allocation is stable in the sense that no agent would want a different allocation at the equilibrium prices.

We should generally expect that iterative auction protocols with discrete bid adjustments would overshoot exact equilibria by at least a small amount. However, approximate equilibrium is a useful concept for analyzing such protocols [Demange *et al.*, 1986; Walsh *et al.*, 1998]. We define and discuss properties and existence of a **λ - δ -competitive equilibrium**—a particular type of approximation equilibrium relevant to task dependency networks. The λ and δ parameterize an agent's maximum error in surplus optimization.

Denote as $H_a(p)$, the maximum surplus that agent a can obtain in (V, E) , at some prices p , subject to feasibility. That is,

$$H_a(p) \equiv \max_{E' \subseteq E} \sigma(a, E', p)$$

such that a is feasible at E' .

Definition 4 (λ - δ -competitive equilibrium) *Given the parameters:*

- $\delta \in \mathbb{R}^+ \cup \{0\}$,
- $\lambda_\pi \in \mathbb{R}^+ \cup \{0\}$ for all $\pi \in \Pi$,
- N_π , the number of input units required by $\pi \in \Pi$




Figure 1: A λ - δ -equilibrium for $\lambda_{B-C-D-to-E} = \delta = 1$.

an allocation (V', E') is in λ - δ -competitive equilibrium at prices p iff:

1. For all $a \in A$, $\sigma(a, E', p) \geq 0$.
2. For all $c \in C$, $\sigma(c, E', p) \geq H_c(p) - \delta$.
3. For all $\pi \in \Pi$, $\sigma(\pi, E', p) \geq H_\pi(p) - (N_\pi \lambda_\pi + \delta)$, and π is feasible at E' .
4. For all $s \in S$, $\sigma(s, E', p) = H_s(p)$.
5. All goods are in material balance at p .

Figure 1 shows a λ - δ -equilibrium for $\lambda_{B-C-D-to-E} = \delta = 1$. 1. Goods are indicated by circles, consumers and suppliers are represented as boxes, and producers are indicated by curved boxes. A solid arrow from a good to an agent indicates that the agent buys the good, and a solid arrow from an agent to a good indicates that the agent sells the good. Dashed arrows indicate input/output capabilities not part of the allocation. Shaded agents receive non-null allocations. Agent values and good prices are shown under their respective nodes. Note that for $\pi = B-C-D-to-E$, $\sigma(\pi, E', p) = 0 \geq H_\pi(p) - (N_\pi \lambda_\pi + \delta) = 1 - 4$. Thus π obeys the producer conditions for λ - δ -competitive equilibrium.

A λ - δ -competitive equilibrium corresponds to the standard notion of a competitive equilibrium when $\lambda_\pi = \delta = 0$ for all producers π . Bikhchandani and Mamer [1997] and Gul and Stacchetti [1997] show that, in an exchange economy, any competitive equilibrium set of prices supports an optimal allocation. We extend this result below to the class of production economies represented by task allocation economies. We show this by proving the more general result that a λ - δ -competitive equilibrium is suboptimally efficient by a fixed bound, $\sum_{\pi \in \Pi} [N_\pi \lambda_\pi + \delta] + |C|\delta$.

Lemma 1 *The value of a feasible allocation (V', E') , at any prices p , can be expressed as:*

$$\text{value}((V', E')) = \sum_{a \in A} \sigma(a, E', p). \quad (1)$$

*Proof sketch.*³ Since supply equals demand in a feasible allocation, all the price terms cancel out and we are

³Complete proofs of all results appear at <http://www-personal.engin.umich.edu/~wew/Papers/ijcai99-extended.ps.Z>.

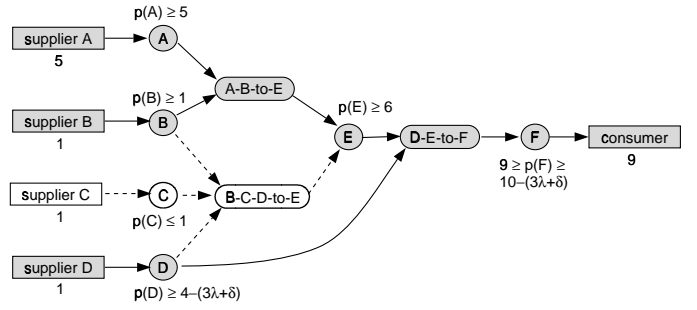


Figure 2: An economy with a λ - δ -competitive equilibrium solution when $3\lambda_{B-C-D-to-E} + \delta > 1$.

left with the original formula for the value of a solution (Definition 1). \square

Theorem 2 *If (V', E') is a λ - δ -competitive equilibrium for (V, E) at some prices p , then (V', E') is a feasible allocation with a nonnegative value that differs from the value of an efficient allocation by at most $\sum_{\pi \in \Pi} [N_\pi \lambda_\pi + \delta] + |C|\delta$.*

Proof sketch. We can compare the value of (V', E') to another feasible allocation (V^*, E^*) , agent-wise, using Equation (1). The λ - δ -competitive equilibrium conditions imposed on the agents imply a global value difference that obeys the stated bound. \square

Not all task allocation economies have competitive equilibria (technically, this is due to complementarity of inputs for producers). However, we can always specify δ and λ_π values such that a λ - δ -competitive equilibrium exists, even when a competitive equilibrium does not exist. For example, we can do this by setting prices of the goods desired by the consumers higher than their values, and the prices of all other goods to zero (if there are any suppliers that could sell directly to consumers with positive surplus on both sides, we can set the prices so that only these suppliers trade with those consumers). We then set δ and the λ_π values sufficiently high such that $H_\pi(p) - (N_\pi \lambda_\pi + \delta) \leq 0$ for all $\pi \in \Pi$. Hence producers could obey the λ - δ -competitive equilibrium conditions by being inactive. Unfortunately, we cannot generally specify a λ - δ -competitive equilibrium for some fixed, problem-independent values of δ and λ_π . Indeed, the necessary values to obtain a *non-solution* λ - δ -equilibrium are proportional to the highest consumer value and the depth of the network.

Figure 2 shows the constraints on $\lambda_{B-C-D-to-E}$ and δ in order to have a *solution* in λ - δ -competitive equilibrium for a particular economy. The only solution involves all agents except B-C-D-to-E. The constraints on the prices of goods A and B ensure that suppliers A and B sell their goods. The constraint on the price of good C ensures that supplier C does not sell its good. The constraint on the price of good E ensures that A-B-to-E is active. The constraint on the price of good D ensures that B-C-D-to-E does not trade any goods. Fi-

nally, the constraint on the price of good F ensures that D-E-to-F is active and also that the consumer buys good F. The constraint on the price of good F requires $3\lambda_{B-C-D-to-E} + \delta > 1$ to have a λ - δ -competitive equilibrium solution. Since this solution is the only efficient allocation, it is the only candidate allocation for a pure competitive equilibrium. But since the parameters must be nonzero, there is no pure competitive equilibrium. Note that to sustain the solution in λ - δ -competitive equilibrium the constraints on the prices for goods A and D, and hence the δ and $\lambda_{B-C-D-to-E}$ values, rise linearly with the cost of supplier A.

4 Market Protocols

The previous section sidesteps the issue of how we might compute an approximate competitive equilibrium. The market-based approach, which we adopt, is to design an *auction mechanism* that mediates negotiation and determines prices and allocations. The auctions along with the agents constitute an *economy*, or *market*, for the task allocation problem.

The agents' *bidding policies* govern their interaction with the auctions. The key distinction between the auction mechanism and bidding policies is that the former is under the control of the system designers, whereas the latter are determined by individual agents. Together, these specifications for behavior constitute a *market protocol*. The market protocol as a whole is our subject of analysis.

In economies characterized by infinite divisibility of goods, nonincreasing returns to scale, and gross substitutes, equilibria always exist and the *tatonnement* protocol converges to a competitive equilibrium. The WALRAS protocol is a variant of tatonnement, that can be implemented in a distributed manner across goods [Cheng and Wellman, 1998]. Highly efficient allocations can be achieved in discrete exchange economies for which each agent wants only a single good [Friedman and Rust, 1993; McAfee and McMillan, 1987; Vickrey, 1961] or have non-complementary preferences for multiple goods [Demange *et al.*, 1986]. However, economies with discrete goods and complementary preferences of agents can lack equilibria, and thus provide a much greater challenge to obtaining efficient allocations.

Many researchers have proposed combinatorial auctions to address the problem resulting from discrete goods and complementarities [Rassenti *et al.*, 1982; Sandholm, to appear]. Combinatorial auctions allocate and price *bundles* of goods. Wurman and Wellman have shown that there always exists a competitive equilibrium set of prices on bundles that supports an efficient allocation [1999]. Approximation techniques can be used to make the computation tractable [Fujishima *et al.*, 1999; Sandholm, to appear]. However, except in some restricted domains [Rothkopf *et al.*, 1998; Sandholm, to appear] a combinatorial auction must generally perform a search over a combinatorial number of bundles.

In this work, we focus on computationally feasible auc-

tions that price individual goods. The ideal such auction would induce a protocol resulting in a (near) competitive equilibrium when one exists, and would produce a highly efficient allocations otherwise. Unfortunately, no protocol has yet been proposed that always converges to near competitive equilibrium, even when one exists, in economies with discrete goods and complementarities.

Lacking an ideal mechanism, many researchers have considered simultaneous ascending auctions [Milgrom, 1997]. When agents bid according to a straightforward rule, the auctions guarantee near-optimal performance in economies with no complementarities [Demange *et al.*, 1986]. Observations of their performance in the United States FCC radio spectrum sale suggests that they can produce high-quality allocations even when complementarities exist [McAfee and McMillan, 1996]. In the following sections we analyze a variant simultaneous ascending auction for the task allocation problem.

5 SAMP-SB Protocol

In the "Simultaneous Ascending (M+1)st-Price with Simple Bidding" (*SAMP-SB*) protocol (extended from our previous work [1998] to allow multiple consumers with preferences for multiple goods), agents negotiate for the goods through auction mediators, one for each good. An auction in turn determines the price and allocation of its respective good. We assume reliable, asynchronous message passing.

5.1 Auction Mechanism

The task allocation market includes a separate auction for each good of potential value. Auctions operate simultaneously and asynchronously. Agents submit bids for goods they wish to buy or sell. A *bid* specifies the price below/above which the agent is willing to buy/sell. Auctions respond with price quotes. A *price quote* specifies the current going price and the number of units the recipient would trade at what price, given the current bid state. Agents may in turn respond with further bids. Each auction requires that an agent's successive bids increase by no less than some (generally small) positive increment δ .

When the market reaches *quiescence*—a state in which no new bids or price quotes are issued—the auctions *clear*. Each bidder is notified of the final prices and how many units it transacted in each good. The ascending rule serves a key role in establishing relationships between market quiescence and solution convergence of the economy [Walsh and Wellman, 1998]. In particular, the market always reaches quiescence with the bidding policies described in Section 5.2.

According to the the (M+1)st-price rules [Satterthwaite and Williams, 1989; Wurman *et al.*, 1998], an auction balances reported supply and demand at a uniform *clearing price*. Winners include all buyers/sellers strictly above/below the price, and, to maximize the benefits from trade, some agents at the clearing price.




Figure 3: Prices and allocations determined by a run of SAMP-SB when $\lambda_\pi = \delta = 1$.

5.2 Bidding Policies

Although multiagent system designers do not generally have control over the agents' behaviors, any conclusions about the outcome of a protocol must be based on some assumptions about these behaviors. In this work, we investigate the equilibrium and efficiency properties of a set of policies that obey the ascending auction bid constraints and require only local, private information.

Let the current going prices specified by the last price quotes be p . Supplier s places a one-time bid of $oc_s(g)$, for each unit $\langle s, g \rangle \in E$ it can supply. When consumer c is not winning a good, it bids $p(g^*) + \delta$ for good $g^* = \arg\max_{g \in G} [v_c(g) - p(g)]$ if $v_c(g^*) - p(g^*) \geq \delta$, otherwise it stops bidding. A producer π initially bids zero for its output g_π . When the prices of its inputs change, π bids $\max(\sum_{\langle g, \pi \rangle \in E} p(g) + k\lambda_\pi, \beta + \delta)$ for g_π , where k is the number of input bids π is currently losing, and β is its previous bid for g_π . A producer initially bids zero for each input good, and raises its bid on input g by δ when it is winning its output but losing its bid for g . The λ_π and δ parameters correspond to parameters in a λ - δ -competitive equilibrium, as described in the following section.

5.3 Equilibrium Convergence of SAMP-SB

λ - δ -competitive equilibria do not always exist for fixed λ_π and δ , and even when they do, the SAMP-SB protocol is not guaranteed to converge to one. Figure 3 shows the results of a run⁴ of SAMP-SB when $\lambda_\pi = \delta = 1$ for all producers π . Note that B-C-D-to-E has a negative surplus, which is a violation of defining Condition (1) of a λ - δ -competitive equilibrium. However, if no inactive producer buys a positive-price input in quiescence, then the economy is in λ - δ -competitive equilibrium.

Theorem 3 *The prices and allocation determined in quiescence by the SAMP-SB protocol is a λ - δ -competitive equilibrium iff no inactive producer buys any positive-price input.*

Proof sketch. *Case only if:* Condition (3) of λ - δ -competitive equilibrium fails if an inactive producer buys any positive-price input. *Case if:* The bidding policies,

⁴The nondeterminism of an asynchronous system implies that different runs may produce different results.

auctions, and theorem conditions ensure $\sigma(a, E', p) \geq 0$ for all $a \in A$. Consumer $c \in C$ maintains at most a single bid such that $\sigma(c, E', p) \geq H_c(p) - \delta$. If $\pi \in \Pi$ is active in quiescence, $\sigma(\pi, E', p) = H_\pi(p)$. The theorem guarantees $\sigma(\pi, E', p) = 0$ when π is inactive, which occurs only if $H_\pi(p) \leq N_\pi \lambda_\pi + \delta$. The bidding policies ensure π is feasible. Supplier $s \in S$ sells g iff $p(g) \geq oc_s(g)$, hence $\sigma(s, E', p) = H_s(p)$. The auctions ensure material balance. \square

5.4 Contract Decommitment

Value is lost when inactive producers purchase some inputs at positive prices. Figure 3 shows a run of SAMP-SB in which B-C-D-to-E buys good C, even though it does not sell its output. A straightforward protocol extension for correcting the inefficiency is to remove the wasteful dead ends by selective contract decommitment.

We propose a **contract decommitment protocol** that is applied after SAMP-SB reaches quiescence. Each inactive producer that wins some inputs at a positive price can decommit from its contracts for its inputs. The protocol is applied recursively to the producers that lose their outputs due to decommitment (we refer to SAMP-SB with decommitment as **SAMP-SB-D**). When the decommitment process terminates, agents exchange goods as specified by the remaining contracts.

In Figure 3, both B-C-D-to-E would decommit from its contract with supplier C. Clearly, Theorem 3 implies that no agent decommits iff SAMP-SB produced a λ - δ -competitive equilibrium. Moreover, if we remove all producers that decommit, and all supplier resources corresponding to decommitted contracts, the remaining agents are in λ - δ -competitive equilibrium. Note that whereas some producers can lose money in the SAMP-SB protocol, no agent receives a negative surplus from participating in SAMP-SB-D.

6 Empirical Efficiency and Equilibrium Analysis of SAMP-SB

We conducted a series of 3300 randomly generated simulations in order to empirically evaluate how often SAMP-SB converges to equilibrium as well as its average quality. We ran 100 trials for each of 5–15 goods and each of 1–3 consumers.

For a trial with N goods, we imposed a total order on the goods. Goods N and $N - 1$ had a number—chosen from $[1, 4]$ —of suppliers. All other goods i had a number—chosen from $[1, 4]$ —of sellers with output i , each of which had a $2/N$ chance of being a supplier and a $(N - 2)/N$ chance of being a producer. To ensure the network was acyclic, a producer with output i had two randomly chosen input goods j and k , such that $j, k > i$. Each consumer desired a single good. One consumer desired good 1, and the other consumers desired randomly chosen goods. We drew consumer valuations uniformly from $[0, 10N]$ and supplier costs uniformly from $[0, 5N]$. We fixed $\delta = \lambda_\pi = 1$.

PEV range	SAMP-SB	SAMP-SB-D
< 0	79	0
0	7	46
1–9	6	0
10–19	10	1
20–29	22	1
30–39	24	4
40–49	32	1
50–59	55	7
60–69	80	15
70–79	158	23
80–89	329	58
90–99	1203	331
100	1305	2813

Mean PEV for SAMP-SB	83
Mean PEV for SAMP-SB-D	97
Percentage of trials in λ - δ -equilibrium	37
Mean non-equilibrium PEV of SAMP-SB	73
Mean non-equilibrium PEV of SAMP-SB-D	95

Table 1: Distributions over percentage of the efficient value (PEV) ranges, mean PEV, and equilibrium attainment for 3300 trials.

As a benchmark, for each economy we computed the value of an efficient (optimal) allocation using A* search with an admissible heuristic. We ran SAMP-SB and SAMP-SB-D, and computed the resulting *percentage of the efficient value (PEV)*. We discarded all trials for which the efficient value was zero, and substituted a new trial using a different configuration but the same number of consumers and goods.

Table 1 shows the distributions over ranges of PEV and mean PEV for SAMP-SB and SAMP-SB-D. The decommitment protocol greatly increases the PEV of SAMP-SB, raising the overall mean from 83% to 97%. The distribution of the PEV of SAMP-SB-D is quite high: 85% of the trials had 100% efficiency and 95% of the trials had at least 90% efficiency.

The SAMP-SB protocol reached λ - δ -equilibrium in over 37% of the trials. When we remove all λ - δ -equilibrium trials from the trial set, the mean PEV of SAMP-SB decreases from 83% to 73%. This is not surprising because the trials that reach λ - δ -equilibrium have high PEV values (Theorem 2). However, the PEV of SAMP-SB-D decreases only a small amount, from 97% to 95%. We might conclude from these simulations that the majority of efficiency loss from the SAMP-SB protocol can be attributed to inactive producers that acquire some inputs, with the resulting wasted costs of suppliers that feed into them.

We note that the PEV numbers we report are somewhat arbitrary in that they depend on the particular method for generating trials. Moreover, if a solution forms for a consumer in a particular economy, then we can get an arbitrarily high PEV by setting the consumer’s value arbitrarily high, while holding constant

the rest of the configuration. We attempted to mitigate this effect by setting consumer values at reasonable levels.

7 Related Work in Task Allocation

The CONTRACT NET protocol forms supply chains top-down in a greedy fashion [Davis and Smith, 1983]. This approach produces satisficing allocations when there are no resource limitations. However, the firm resource constraints we impose in our model would generally require CONTRACT NET to backtrack in order to guarantee a feasible allocation. Additionally, since the allocation policy is loosely specified in the original work, we cannot draw any conclusions about the efficiency of the approach.

Sandholm [1993] examines a specialization of CONTRACT NET for a generalization of Task Oriented Domains (TODs) [Rosenschein and Zlotkin, 1994]. Agents begin with an initial allocation of tasks and negotiate task exchanges until there are no more mutually beneficial trades. The trades can greatly increase the allocation quality, but the system may get stuck in local minima.

Sandholm’s model includes local constraints on task achievement, but does not impose a hierarchical dependency structure. Thus each locally feasible bilateral trade can be executed immediately and independently of other trades. We cannot generally apply an incremental trading protocol to our task allocation model. A local exchange may require reallocation throughout the entire network to maintain production feasibility.

Andersson and Sandholm [1998] find that decommitment protocols increase the quality of the resulting allocations in variants of TODs. With incremental trading, decommitment gives agents the opportunity to engage in other more cost-effective task allocation contracts.

8 Conclusions and Future Work

We have shown that we can bound the inefficiency of a λ - δ -competitive equilibrium in a task dependency network in proportion to the number of agents and the bidding parameters in the SAMP-SB protocol. SAMP-SB can sometimes converge to a λ - δ -competitive equilibrium, when one exists. When combined with contract decommitment, SAMP-SB produces highly efficient allocations on average.

SAMP-SB relies on the *competitive assumption* that agents take price quotes as the actual prices, and do not attempt to manipulate them. Although this assumption is reasonable when there are many agents trading in each good—and hence no agent has significant market power—this assumption is less realistic in thinner markets. An agent may have opportunities to gain higher surplus by using other bidding policies, particularly if it has knowledge of other agents’ preferences or behavior. An enforcement mechanism could ensure that agents correctly follow the decommitment protocol, but the potential to decommit may still affect agents’ negotiation policies. We seek to establish conditions under which it

is rational for agents to play SAMP-SB-D, and to analyze strategic behavior when these conditions do not hold.

We wish to examine the possibility of including combinatorial auctions with SAMP-SB in the task allocation market. If we can identify portions of task networks with structures amenable to feasible bundle pricing, we may be able to increase the efficiency of allocations.

Acknowledgments

This work was supported by a NASA/Jet Propulsion Laboratory Graduate Student Researcher fellowship, a DARPA Information Survivability grant, and a NSF/DARPA/NASA Digital Library Initiative grant. A portion of this work was performed at the 1998 Santa Fe Institute Workshop on Computational Economics. We also thank the anonymous reviewers for helpful suggestions.

References

- [Andersson and Sandholm, 1998] Martin R. Andersson and Tuomas W. Sandholm. Leveled commitment contracting among myopic individually rational agents. In *Third International Conference on Multi-Agent Systems*, pages 26–33, 1998.
- [Bikhchandani and Mamer, 1997] Sushil Bikhchandani and John W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74:385–413, 1997.
- [Cheng and Wellman, 1998] John Q. Cheng and Michael P. Wellman. The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes. *Computational Economics*, 12:1–24, 1998.
- [Davis and Smith, 1983] Randall Davis and Reid G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.
- [Demange *et al.*, 1986] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.
- [Friedman and Rust, 1993] Daniel Friedman and John Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1993.
- [Fujishima *et al.*, 1999] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Sixteenth International Joint Conference on Artificial Intelligence*, 1999.
- [Gul and Stacchetti, 1997] Faruk Gul and Ennio Stacchetti. Walrasian equilibrium without complementarities. Technical report, Princeton University and University of Michigan, 1997.
- [McAfee and McMillan, 1987] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.
- [McAfee and McMillan, 1996] R. Preston McAfee and John McMillan. Analyzing the airwaves auction. *Journal of Economic Perspectives*, 10(1):159–175, Winter 1996.
- [Milgrom, 1997] Paul Milgrom. Auction theory in practice: The simultaneous ascending auction. Technical Report TR 98-0002, Stanford University, December 1997. Department of Economics.
- [Rassenti *et al.*, 1982] S. J. Rassenti, V. L. Smith, and R. L. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13:402–417, 1982.
- [Rosenschein and Zlotkin, 1994] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [Rothkopf *et al.*, 1998] Michael H. Rothkopf, Aleksandar Pekeć, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.
- [Sandholm, 1993] Tuomas W. Sandholm. An implementation of the CONTRACT NET protocol based on marginal cost calculations. In *Eleventh National Conference on Artificial Intelligence*, pages 256–262, 1993.
- [Sandholm, to appear] Tuomas Sandholm. Approaches to winner determination in combinatorial auctions. *Decision Support Systems*, to appear.
- [Satterthwaite and Williams, 1989] Mark A. Satterthwaite and Steven R. Williams. Bilateral trade with the sealed bid k-double auction: Existence and efficiency. *Journal of Economic Theory*, 48:107–133, 1989.
- [Vickrey, 1961] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [Walsh and Wellman, 1998] William E. Walsh and Michael P. Wellman. A market protocol for decentralized task allocation. In *Third International Conference on Multi-Agent Systems*, pages 325–332, 1998.
- [Walsh *et al.*, 1998] William E. Walsh, Michael P. Wellman, Peter R. Wurman, and Jeffrey K. MacKie-Mason. Some economics of market-based distributed scheduling. In *Eighteenth International Conference on Distributed Computing Systems*, pages 612–621, 1998.
- [Wellman, 1993] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.
- [Wurman and Wellman, 1999] Peter R. Wurman and Michael P. Wellman. Equilibrium prices in bundle auctions. Santa Fe Institute working paper, 1999.
- [Wurman *et al.*, 1998] Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24:17–27, 1998.