A General-Equilibrium Approach to Distributed Transportation Planning

Michael P. Wellman
USAF Wright Laboratory
WL/AAA-1
Wright-Patterson AFB, OH 45433
wellman@wl.wpafb.af.mil

Abstract

Market price mechanisms from economics constitute a well-understood framework for coordinating decentralized decision processes with minimal communication. WALRAS is a general "market-oriented programming" environment for the construction and analysis of distributed planning systems, based on general-equilibrium theory. The environment provides basic constructs for defining computational market structures, and a procedure for deriving their corresponding competitive equilibria. In a particular realization of this approach for a simplified form of distributed transportation planning, we see that careful construction of the decision process according to economic principles can lead to effective decentralization, and that the behavior of the system can be meaningfully analyzed in economic terms.

Distributed Planning

In a *distributed* or multiagent planning system, the plan for the system as a whole is a composite of plans produced by its constituent agents. Planning might be distributed because agents are separated geographically, have different information, possess distinct capabilities or authority, or were designed and implemented separately. In any case, because each agent has limited competence and awareness of the decisions produced by others, some sort of coordination is required to maximize the performance of the overall system. However, central control or extensive communication is deemed infeasible, as it violates whatever constraints dictated distribution of the planning task in the first place.

The task facing the designer of a distributed planning system is to define a computationally efficient coordination mechanism and its realization for a given or constructed configuration of agents. By the term *agent*, I refer to a module that acts within the mechanism according to its own knowledge and interests. The capabilities of the agents and their organization in an overall decision-making structure determine the behavior of the overall system. Because it concerns the collective behavior of self-interested decision makers, the design of this decentralized structure is fundamentally an exercise in economics. The problem of developing architectures for distributed planning is largely one of ***mechanism design*** [Hurwicz, 1977; Reiter, 1986], and many ideas and results from economics are directly applicable. In particular, the class of mechanisms based on price systems and competition has been deeply investigated by economists, who have characterized the conditions for its efficiency and compatibility with other features of the economy. When applicable, the competitive mechanism achieves coordination with minimal communication requirements (in a precise sense related to the dimensionality of messages transmitted among agents [Reiter, 1986]).

The theory of ***general equilibrium*** [Hildenbrand and Kirman, 1976] provides the foundation for a general approach to the construction of distributed planning systems based on price mechanisms. In this approach, we regard the constituent planning agents as consumers and producers in an artificial economy, and define their individual activities in terms of production and consumption of commodities. Interactions among agents are cast as exchanges, the terms of which are mediated by the underlying economic mechanism, or protocol. By specifying the universe of commodities, the configuration of agents, and the interaction protocol, we can achieve a variety of interesting and often effective decentralized behaviors. Furthermore, we can apply economic theory to the analysis of alternative architectures, and thus exploit a wealth of existing knowledge in the design of distributed planners.

In the following, I describe this general approach and a programming environment based on it. An example problem in distributed transportation planning demonstrates the feasibility of decentralizing a problem with nontrivial interactions, and the applicability of economic principles to collective problem solving.

Transportation Example

In a simplified version of the transportation planning problem, the task is to allocate a given set of cargo movements over a given transportation network. The transportation network is a collection of locations, with links (directed edges) identifying feasible transportation operations. Associated with each link is a specification of the cost of moving cargo along it. Suppose further that the cargo is homogeneous, and amounts of cargo are arbitrarily divisible. A movement requirement associates an amount of cargo with an origin-destination pair. The planning problem is to determine the amount to transport on each link in order to move all the cargo at the minimum cost.

A distributed version of the problem would decentralize the responsibility for transporting separate cargo elements. For example, planning modules corresponding to geographically or organizationally disparate units might arrange the transportation for cargo within their respective spheres of authority. Or decision-making activity might be decomposed along hierarchical levels of abstraction, gross functional characteristics, or according to any other relevant distinction. This decentralization might result from real distribution of authority within a human organization, from inherent informational asymmetries and communication barriers, or from modularity imposed to facilitate software engineering.

Consider, for example, the abstract transportation network of Figure 1, taken from Harker [1988].[1] There are four locations, with directed links as shown. Consider two movement requirements. The first is to transport cargo from location 1 to location 4, and the second in the reverse direction. Suppose we wish to decentralize authority so that separate agents (called shippers) decide how to allocate the cargo for each movement. The first shipper decides how to split its cargo units between the paths $1 \to 2 \to 4$ and $1 \to 2 \to 3 \to 4$, while the second figures the split between paths $4 \to 2 \to 1$ and $4 \to 2 \to 3 \to 1$. Note that the latter paths for each shipper share a common resource: the link $2 \to 3$.

Because of their overlapping resource demands, the shippers' decisions appear to be necessarily intertwined. In a congested network, for example, the cost for transporting a unit of cargo over a link is increasing in the overall usage of the link. A shipper planning its cargo movements as if it were the only user on a network would thus underestimate its costs and potentially misallocate transportation resources.

For the analysis of networks such as this, transportation researchers have developed equilibrium concepts describing the collective behavior of the shippers. In a

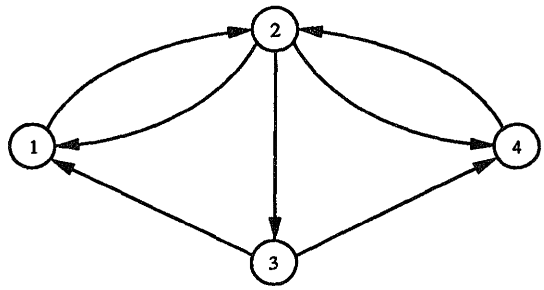

Figure 1: A simple network (from Harker [1988]).

system equilibrium, the overall transportation of cargo proceeds as if there were an omniscient central planner directing the movement of each shipment so as to minimize the total aggregate cost of meeting the requirements. In a *user equilibrium*, the overall allocation of cargo movements minimizes each shipper's total cost, with shippers sharing proportionately the cost of shared resources. There are also some intermediate possibilities, corresponding to game-theoretic equilibrium concepts such as the Nash equilibrium, where each shipper behaves optimally given the transportation policies of the remaining shippers [Harker, 1986].

From our perspective as designer of the distributed planner, we seek a decentralization mechanism that will reach the system equilibrium, or come as close as possible given the distributed decision-making structure. In general, however, we cannot expect to derive a system equilibrium or globally optimal solution without central control. Limits on coordination and communication may prevent the distributed resource allocation from exploiting all opportunities and inhibiting agents from acting at cross purposes. But under certain conditions decision making can indeed be decentralized effectively via market mechanisms. General-equilibrium analysis can help us to recognize and take advantage of these opportunities.

The WALRAS Environment

To explore the use of market mechanisms for distributed planning, I have developed a prototype environment for specifying and simulating computational markets. The system, called WALRAS,[2] provides basic mechanisms implementing various sorts of agents, auctions, and bidding protocols. To specify a computational economy, one defines a set of goods and instantiates a collection of agents that produce or consume those goods. The simulation engine of WALRAS then "runs" these agents to determine an equilibrium allocation of goods and activities in the economy.

[1] Models of this sort are employed in transportation analysis to predict cargo movements and hence characterize the effect of variations in transportation infrastructure or policy. Their intent is descriptive, as the agents are private individuals or firms outside the policymaker's control. Although the overall role of a planning system is to *prescribe* behavior, the designer of a distributed architecture also requires a descriptive model of the modules' behavior to characterize the effect of alternative configurations and coordination mechanisms.

[2] Named for the 19th-century French economist Léon Walras, who was the first to envision a system of interconnected markets in price equilibrium.

Market Configuration and Equilibrium

Agents fall in two general classes. *Consumers* can buy, sell, and consume goods, and their preferences for consuming various combinations of goods are specified by their *utility function*. *Producers* can transform some sorts of goods into some others, according to their technology or *production function*. Each type of agent may start with an initial allocation of some goods, termed their *endowment*. The objective of a consumer is to maximize its utility, subject to the constraint (the *budget constraint*) that the cost of its consumption bundle does not exceed the value of its endowment at the going prices. The objective of a producer is to maximize profits given the going price of its output good and the inputs required for its production (the *factor goods*).

WALRAS associates an *auction* with each distinct good. Agents act in the market by submitting *bids* to auctions. The form of a bid is determined by the auction protocol. In a price mechanism, bids specify a correspondence between prices and quantities of the good that the agent offers to demand or supply. The auction derives a market-clearing price, at which the quantity demanded balances that supplied, within some prespecified tolerance. When the current price is clearing with respect to the current bids, we say the market for that commodity is in equilibrium.

An agent acts competitively when it takes prices as given, neglecting any impact of its own behavior on the market-clearing price. Perfect competition reflects individual rationality when there are numerous agents, each small with respect to the entire economy. However, when an individual agent is large enough to affect prices significantly, it forfeits utility or profits by failing to take this into account.

Under the assumption of perfect competition, each agent's constrained optimization problem is parameterized by the prices of goods. We say that an agent is in equilibrium if its set of outstanding bids corresponds to the solution of its optimization problem at the going prices. If all the agents and commodity markets are in equilibrium, the allocation of goods dictated by the auction results is a *competitive equilibrium*. From the perspective of mechanism design, competitive equilibria possess several desirable properties, in particular, the two fundamental welfare theorems of general equilibrium theory: (1) all competitive equilibria are *Pareto optimal* (no agent can do better without some other doing worse), and (2) *any* Pareto optimum is a competitive equilibrium for some initial endowment. These properties seem to offer exactly what we need: a bound on the quality of the solution, plus the prospect that we can achieve the most desired behavior by carefully engineering the configuration of the computational market. Moreover, in equilibrium, the prices reflect exactly the information required for distributed agents to optimally evaluate perturbations in their behavior without resorting to communication or reconsideration of their full set of possibilities.

Computing Competitive Equilibria

Under certain "classical" assumptions (essentially continuity, monotonicity, and convexity of preferences and technologies), competitive equilibria exist, and are unique given some further restrictions. They are also computable, and algorithms based on fixed-point methods [Scarf, 1984] and optimization (variational inequality) techniques have been developed. Both sorts of algorithms in effect solve the simultaneous equilibrium equations by convergent iteration. However, by employing the equilibrium equations, these techniques violate the decentralization considerations underlying our distributed planning application.[3] For example, the constraint that profits be zero is a consequence of competitive behavior and constant-returns technology. Since information about the form of the technology and bidding policy is considered private to producer agents, it would not be permissible to embed the zero-profit condition into the equilibrium derivation procedure, as is sometimes done in computable general-equilibrium models. Similarly, explicitly examining the joint commodity space in the search for equilibrium undercuts our original motive for decomposing complex activities into consumption and production of separate goods.

WALRAS's procedure is a decentralized relaxation method, akin to the mechanism of *tatonnement* originally sketched by Léon Walras to explain how prices might be derived. A tatonnement method iteratively adjusts prices up or down as there is an excess of demand or supply, respectively (e.g., in proportion to the excess). The method employed by WALRAS successively computes an equilibrium price in each separate market, in a manner detailed below.[4] Like tatonnement, it involves an iterative adjustment of prices based on reactions of the agents in the market. However, it differs from traditional tatonnement procedures in that (1) agents submit supply and demand *curves* rather than single quantities for a particular price, and (2) the auction adjusts individual prices to clear, rather than the entire price vector by some function of summary statistics such as excess demand.

Figure 2 presents a schematic view of the WALRAS bidding process. There is an auction for each distinct good, and for each agent, a link to all auctions in which it has an interest. There is also a "tote board" of current prices, kept up-to-date by the various auctions.

[3]These methods are typically applied to the analysis of existing decentralized structures, such as transportation industries or even entire economies [Shoven and Whalley, 1984]. Because our purpose is to *implement* a distributed system, we must obey *computational* distributivity constraints not relevant to the usual purposes of applied general-equilibrium analysis.

[4]This general approach is called "progressive equilibration" by Dafermos and Nagurney [1989]. WALRAS performs progressive equilibration in their sense, but adopts a different model of market structure in general, and transportation networks in particular.

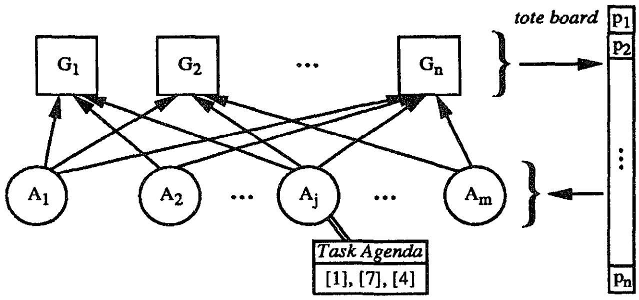

Figure 2: WALRAS's bidding process. G_i denotes the auction for the ith good, and A_i the ith trading agent.

Each agent maintains an agenda of bid tasks, specifying the markets in which it must update its bid or compute a new one. In Figure 2, agent A_j has pending tasks to submit bids to auctions G_1, G_7, and G_4. The bid for a particular good corresponds to one dimension of the agent's solution to its constrained optimization problem, which is parameterized by the prices for all relevant goods. Acting as a perfect competitor, a WALRAS agent bids for a good under the assumption that prices for the remaining goods are fixed at their current values. The bid itself is a schedule of quantities and prices (encoded in any of a variety of formats) specifying the amount of the good demanded or supplied as its own price varies.

As new bids are received at auction, the previously computed clearing price becomes obsolete. Periodically, each auction computes a new clearing price (if any new or updated bids have been received) and posts it on the tote board. When a price is updated, this may invalidate some of an agent's outstanding bids, since these were computed under the assumption that prices for remaining goods were fixed. On finding out about a price change, an agent augments its task agenda to include the potentially affected bids.

At all times, WALRAS maintains a vector of going prices and quantities that would be exchanged at those prices. While the agents have nonempty bid agendas or the auctions new bids, some or all goods may be in disequilibrium. When all auctions clear and all agendas are exhausted, however, the economy is in competitive equilibrium (up to some numeric tolerance). This process is highly distributed, in that each agent need communicate directly only with the auctions for the goods of interest (those in the domain of its utility or production function, or for which it has nonzero endowments). Each of these interactions concerns only a single good; auctions never coordinate with each other. Agents need not negotiate directly with other agents, nor even know of each other's existence.

It is well known that tatonnement processes may not converge to equilibrium (but convergent results are indeed competitive equilibria) [Scarf, 1984]. The class of economies in which tatonnement works are those with stable equilibria: those without complementarities in preferences or technologies [Arrow and Hurwicz, 1977]. I have been unable thus far to characterize WALRAS's adjustment process simply enough to analyze its dynamics mathematically. Similar progressive equilibration algorithms are known to converge for certain special cases [Eydeland and Nagurney, 1989].

Market-Oriented Programming

As described above, WALRAS provides facilities for specifying market configurations and computing their competitive equilibrium. We can also view WALRAS as a programming environment for decentralized resource allocation procedures. The environment provides constructs for specifying various sorts of agents and defining their interactions via their relations to common commodities. After setting up the initial configuration, the market can be run to determine the equilibrium level of activities and distribution of resources throughout the economy.

To cast a distributed planning problem as a market, one needs to identify (1) the goods traded, (2) the agents trading, and (3) the agents' bidding behavior. Finally, it might be advantageous to adjust some general parameters of the bidding protocol. These design steps are serially dependent, as the definition of what constitutes an exchangeable or producible commodity severely restricts the type of agents that it makes sense to include. Below, I illustrate the design task with a WALRAS formulation of the transportation example.

Implementation

WALRAS is implemented in Common Lisp and the Common Lisp Object System (CLOS). The current version provides basic infrastructure for running computational economies, including the underlying bidding protocol and a library of CLOS classes implementing a variety of agent types. The object-oriented implementation supports incremental development of market configurations. In particular, new types of agents can often be defined as slight variations on existing types, for example by modifying isolated features of the demand policy or bid format.

Although it models a distributed system, WALRAS runs serially on a single processor. Distribution constraints on information and communication are enforced by programming and specification conventions rather than by fundamental mechanisms of the software environment. Asynchrony is simulated by randomizing the bidding sequences so that agents are called on unpredictably. Indeed, artificial synchronization can lead to an undesirable oscillation in the clearing prices, as agents collectively overcompensate for imbalances in the preceding iteration.[5]

[5] In some formal dynamic models [Huberman, 1988; Kephart *et al.*, 1989], homogeneous agents choose instantaneously optimal policies without accounting for others

The current experimental system runs transportation models of the sort described in the next section, as well as some abstract exchange and production economies with parameterized utility and production functions (including the expository examples of Scarf [1984] and Shoven and Whalley [1984]). Customized tuning of the basic bidding protocol has not been necessary. In the process of getting WALRAS to run on these examples, I have produced some substantial intermediate object structure, but much more is required to fill out a comprehensive taxonomy of agents, bidding strategies, and auction policies.

WALRAS Transportation Market

Market Structure

The primary commodity of interest in this problem is movement of cargo. Because the value and cost of a cargo movement depends on location, we designate as a distinct good capacity on each origin-destination pair in the network (see Figure 1). To capture the cost or input required to move cargo, we define another good denoting generic transportation resources. In a more concrete model, these might consist of vehicles, fuel, labor, or other factors contributing to transportation.

To decentralize the decision making, we identify two groups of agents. The consumers, or shippers, have an interest in moving various units of cargo between specified locations. We identify each movement requirement with a single shipper agent. The producers, or carriers, have the capability to transport cargo units over specified links, given varying amounts of transportation resources. In the model described here, we associate one carrier with each available link. To achieve a global movement of cargo, shippers obtain transportation services from carriers in exchange for the necessary transportation resources.

The interconnectedness of agents and goods defines the market configuration. Figure 3 depicts the WALRAS configuration for the example network of Figure 1. Let $C_{i,j}$ denote the carrier that transports cargo from location i to location j, and $G_{i,j}$ the good representing an amount of cargo moved over that link. Each carrier $C_{i,j}$ is connected to the auction for $G_{i,j}$, its output good, along with G_0—its input in the production process. Shipper agents are also connected to G_0, as they are endowed with transportation resources to exchange for transportation services. In this model there are two shippers, $S_{1,4}$ and $S_{4,1}$, where $S_{i,j}$ denotes a shipper with a requirement to move goods from origin i to destination j. On the demand side, shippers connect to goods that might serve their objectives: in this case, movement along links that belong to some simple path from the shipper's origin to its destination.

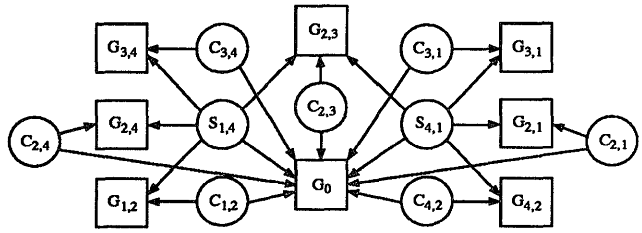

Figure 3: WALRAS market configuration for the example transportation network.

The model we employ for transportation costs is based on a network with congestion, thus exhibiting diseconomies of scale. In other words, the marginal and average costs (in terms of transportation resources required) are both increasing in the level of service on a link. Using Harker's data, costs are quadratic.[6] Let $c_{i,j}(x)$ denote the cost in transportation resources (good G_0) required to transport x units of cargo on the link from i to j. The complete cost functions are:

$$c_{1,2}(x) = c_{2,1}(x) = c_{2,4}(x) = c_{4,2}(x) = x^2 + 20x,$$

$$c_{3,1}(x) = c_{2,3}(x) = c_{3,4}(x) = 2x^2 + 5x.$$

Finally, each shipper's objective is to transport 10 units of cargo from its origin to its destination.

Agent Behavior

In the case of a decreasing returns technology, the producer's (carrier's) optimization problem has a unique solution. The optimal level of activity maximizes revenues minus costs, which occurs at the point where the output price equals marginal cost. Using this result, carriers submit supply bids specifying transportation services as a function of link prices (with resource price fixed), and demand bids specifying required resources as a function of input prices (for activity level computed with output price fixed).

For example, consider carrier $C_{1,2}$. At output price $p_{1,2}$ and input price p_0, the carrier's profit is

$$p_{1,2}y - p_0 c_{1,2}(y),$$

where y is the level of service it chooses to supply. Given the cost function above, this expression is maximized at $y = (p_{1,2} - 20p_0)/2p_0$. Taking p_0 as fixed, the carrier submits a supply bid with y a function of $p_{1,2}$.

that are simultaneously making the same choice, detracting from its value. Oscillation occurs as the result of delayed feedback about the others' decisions. I have also observed this phenomenon empirically in a synchronized version of WALRAS. By eliminating the synchronization, agents tend to work on different markets at any one time, and hence do not suffer as much from this time delay.

[6] The quadratic cost model is posed simply for concreteness, and does not represent any substantive claim about transportation networks. The important qualitative feature of this model is that it exhibits decreasing returns, a defining characteristic of congested networks. Note also that Harker's model is in terms of monetary costs, whereas we introduce an abstract input good.

On the demand side, the carrier takes $p_{1,2}$ as fixed and submits a demand bid for enough good G_0 to produce y, where y is treated as a function of p_0.

The bidding behavior of shippers is more complicated. Rather than explicitly consider utility maximization, we take the shipper's objective to be to ship as much as possible (up to its movement requirement) in the least costly manner. Given a network with prices on each link, the cheapest cargo movement corresponds to the shortest path in the graph, where distances are equated with prices. Thus, for a given link, a shipper would prefer to ship its entire quota on the link if it is on the shortest path, and zero otherwise. In the case of ties, it is indifferent among the possible allocations. To bid on link i, j, the shipper can derive the threshold price that determines whether it is on a shortest path by taking the difference in shortest-path distance between the networks where link i, j's distance is set to zero and infinity, respectively.

In incrementally changing its bids, the shipper should also consider its outstanding bids and the current prices. The value of reserving capacity on a particular link is zero if it cannot get service on the other links on the path. Similarly, if it is already committed to shipping cargo on a parallel path, it does not gain by obtaining more capacity (even at a lower price) until it withdraws these other bids.[7] Therefore, the actual demand policy of a shipper is to spend its uncommitted income on the potential flow increase (derived from maximum-flow calculations) it could obtain by purchasing capacity on the given link. It is willing to spend up to the threshold value of the link, as described above. This determines one point on its demand curve. If it has some unsatisfied requirement and uncommitted income it also indicates a willingness to pay a lower price for a smaller amount of capacity. Boundary points such as this serve to bootstrap the economy; from the initial conditions it is typically the case that no individual link contributes to overall flow between the shipper's origin and destination. Finally, the demand curve is completed by an arbitrary smoothing operation on these points.

Results

With the configuration and agent behaviors described, WALRAS derives the system equilibrium (SE), that is, the cargo allocation minimizing overall transportation costs. The derived cargo movements are correct to within 10% in 36 bidding cycles, and to 1% in 72, where in each cycle every agent submits an average of one bid to one auction. The total cost (in units of G_0), its division between shippers' expenditures and carriers' profits, and the equilibrium prices are presented in Table 1. That the decentralized process produces a global optimum is perfectly consistent with competitive behavior—the carriers price their outputs at marginal cost, and the technologies are convex.

Why, then, is the perfectly competitive model typically associated with user equilibrium (UE)? The answer is that in most models (including Harker's), carriers are not expressly modeled as agents, and it is assumed that shippers share proportionately the exact cost of transportation. Under such a regime, the economic cost facing shippers is the average cost of shipment along a link. We can realize this policy in WALRAS by modifying the carriers' supply policy so that they just cover their average cost. The resulting solution is indeed UE.[8]

The lesson from this exercise is that we can achieve qualitatively distinct results by simple variations in the market configuration or agent policies. From our designers' perspective, we prefer the configuration that leads to the more transportation-efficient SE. Examination of Table 1 reveals that we can achieve this result by allowing the carriers to earn nonzero profits (economically speaking, these are really rents on the fixed factor represented by the congested channel) and redistributing these profits to the shippers to cover their increased expenditures.[9]

(Some of the) Limitations

One serious limitation of WALRAS is the assumption that agents act competitively. There are two approaches toward alleviating this restriction in a computational economy. First, we could simply adopt models of imperfect competition, perhaps based on specific forms of imperfection (e.g., spatial monopolistic competition) or on general game-theoretic models. Second, as architects we can configure the markets to promote competitive behavior. For example, decreasing the agent's grain size and enabling free entry of agents should enhance the degree of competition. Perhaps most interestingly, by controlling the

[7]Even if a shipper could simultaneously update its bids in all markets, it would not be a good idea to do so here. A competitive shipper would send all its cargo on the least costly path, neglecting the possibility that this demand may increase the prices so that it is no longer cheapest. This policy can cause perpetual oscillation. In the example network, the unique equilibrium has each shipper routing portions of its cargo on all available paths. Policies allocating all cargo to one path can never lead to this result, and hence convergence to competitive equilibrium depends on the incrementality of bidding behavior.

[8]Average-cost pricing is perhaps the most common mechanism for allocating costs of a shared resource. Shenker [1991] points out problems with this scheme—with respect to both efficiency and strategic behavior—in the context of allocating access to congested computer networks, a problem analogous to our transportation task.

[9]In the model of general equilibrium with production, consumers own shares in the producers' profits. This closes the loop so that all value is ultimately realized in consumption. We can specify these shares as part of the initial configuration, just like the endowment. In this example, we distribute the shares evenly between the two shippers.

pricing	total cost	shipper	carrier	$p_{1,2}$	$p_{2,1}$	$p_{2,3}$	$p_{2,4}$	$p_{3,1}$	$p_{3,4}$	$p_{4,2}$
marginal cost (SE)	1136	1514	378	40.0	35.7	22.1	35.7	13.6	13.6	40.0
average cost (UE)	1143	1143	0	30.0	27.1	16.3	27.1	10.7	10.7	30.0

Table 1: Equilibria derived by WALRAS for the transportation example. Total cost = shipper expense−carrier profit.

agents' knowledge of the market structure (via standard information-encapsulation techniques), we can degrade their ability to exploit whatever market power they possess. Uncertainty has been shown to increase competitiveness among risk-averse agents in some formal bidding models [McAfee and McMillan, 1987], and in a computational environment we have substantial control over this uncertainty.

The existence of competitive equilibria and efficient market allocations also depends critically on the assumption of nonincreasing returns to scale. Although congestion is a real factor in transportation networks, for many modes of transport there are often other economies of scale and density that may lead to returns that are increasing overall [Harker, 1987].

Having cast WALRAS as a general environment for distributed planning, it is natural to ask how universal "market-oriented programming" is as a computational paradigm. We can characterize the computational power of this model easily enough, by correspondence to the class of convex programming problems represented by economies satisfying the classical conditions. However, the more interesting issue is how well the conceptual framework of market equilibrium corresponds to the salient features of distributed planning problems. Although it is too early to make a definitive assertion about this, it seems clear that many planning tasks are fundamentally problems in resource allocation, and that the units of distribution often correspond well with units of agency. Economics has been the most prominent (and arguably the most successful) approach to modeling resource allocation with decentralized decision making, and it is reasonable to suppose that the concepts economists find useful in the social context will prove similarly useful in our analogous computational context. Of course, just as economics is not ideal for analyzing all aspects of social interaction, we should expect that many issues in the organization of distributed planning will not be well accounted-for in this framework.

Finally, the transportation network model presented here is a highly simplified version of the actual planning problem for this domain. A more realistic treatment would cover multiple commodity types, discrete movements, temporal extent, hierarchical network structure, and other critical features of the problem. Some of these may be captured by incremental extensions to the simple model, perhaps applying elaborations developed by the transportation science community.[10]

Related Work

The techniques and models described here obviously build on much work in economics and transportation science. The intended research contribution here is to neither of these fields, but rather in their application to the construction of a computational framework for decentralized decision making in general and distributed transportation planning in particular.

The basic idea of applying economic mechanisms to coordinate distributed problem solving is not new to the AI community. Starting with the contract net [Davis and Smith, 1983], many have found the metaphor of markets appealing, and have built systems organized around markets or market-like mechanisms [Malone *et al.*, 1988]. Miller and Drexler [1988] have explored this approach in depth, presenting some underlying rationale and addressing specific issues salient in a computational environment. Recently, Waldspurger et al. [1992] have developed computational market mechanisms to allocate computational resources in a distributed operating system. For further remarks on this line of work, see [Wellman, 1991].

WALRAS is distinct from these prior efforts in two primary respects. First, it is constructed expressly in terms of concepts from general equilibrium theory, to promote mathematical analysis of the system and facilitate the application of economic principles to architectural design. Second, WALRAS is designed to serve as a general programming environment for implementing computational economies. Although not developed specifically to allocate computational resources, there is no reason these could not be included in market structures configured for particular application domains. Indeed, the idea of grounding measures of the value of computation in real-world values (e.g., cargo movements) follows naturally from the general-equilibrium view of interconnected markets, and is one of the more exciting prospects for future applications of WALRAS to distributed problem-solving.

Finally, market-oriented programming shares with Shoham's [1990] *agent-oriented programming* the view that distributed problem-solving modules are best designed and understood as rational agents. The two approaches support different agent operations (trans-

[10]For example, many transportation models (including Harker's more elaborate formulation [Harker, 1987]) allow for variable supply and demand of the commodities and more complex shipper-carrier relationships. Concepts of spatial price equilibrium, based on markets for commodities in each location, seem to offer the most direct approach toward extending the transportation model within WALRAS.

actions versus speech acts), adopt different rationality criteria, and emphasize different agent descriptors, but are ultimately aimed at achieving the same goal of specifying complex behavior in terms of agent concepts (e.g., belief, desire, capability) and social organizations. Combining individual rationality with laws of social interaction provides perhaps the most natural approach to generalizing Newell's [1982] "knowledge level analysis" idea to distributed computation.

Conclusion

In summary, WALRAS represents a general approach to the construction and analysis of distributed planning systems, based on general equilibrium theory and competitive mechanisms. The approach works by deriving the competitive equilibrium corresponding to a particular configuration of agents and commodities, specified using WALRAS's basic constructs for defining computational market structures. In a particular realization of this approach for a simplified form of distributed transportation planning, we see that qualitative differences in economic structure (e.g., cost-sharing among shippers versus ownership of shared resources by profit-maximizing carriers) correspond to qualitatively distinct behaviors (user versus system equilibrium). This exercise demonstrates that careful design of the distributed decision structure according to economic principles can sometimes lead to effective decentralization, and that the behaviors of alternative systems can be meaningfully analyzed in economic terms.

Acknowledgment

I have benefited from discussions of computational economies with several colleagues, and would like to thank in particular Jon Doyle, Anna Nagurney, Scott Shenker, Carl Waldspurger, and Martin Weitzman for helpful comments and suggestions.

References

Arrow, Kenneth J. and Hurwicz, Leonid, editors 1977. *Studies in Resource Allocation Processes.* Cambridge University Press, Cambridge.

Dafermos, Stella and Nagurney, Anna 1989. Supply and demand equilibration algorithms for a class of market equilibrium problems. *Transportation Science* 23:118–124.

Davis, Randall and Smith, Reid G. 1983. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence* 20:63–109.

Eydeland, Alexander and Nagurney, Anna 1989. Progressive equilibration algorithms: The case of linear transaction costs. *Computer Science in Economics and Management* 2:197–219.

Harker, Patrick T. 1986. Alternative models of spatial competition. *Operations Research* 34:410–425.

Harker, Patrick T. 1987. *Predicting Intercity Freight Flows.* VNU Science Press, Utrecht, The Netherlands.

Harker, Patrick T. 1988. Multiple equilibrium behaviors on networks. *Transportation Science* 22:39–46.

Hildenbrand, W. and Kirman, A. P. 1976. *Introduction to Equilibrium Analysis: Variations on Themes by Edgeworth and Walras.* North-Holland Publishing Company, Amsterdam.

Huberman, B. A., editor 1988. *The Ecology of Computation.* North-Holland.

Hurwicz, Leonid 1977. The design of resource allocation mechanisms. In Arrow and Hurwicz [1977]. 3–37. Reprinted from *American Economic Review Papers and Proceedings*, 1973.

Kephart, J. O.; Hogg, T.; and Huberman, B. A. 1989. Dynamics of computational ecosystems. *Physical Review A* 40:404–421.

Malone, T. W.; Fikes, R. E.; Grant, K. R.; and Howard, M. T. 1988. Enterprise: A market-like task scheduler for distributed computing environments. In Huberman [1988]. 177–205.

McAfee, R. Preston and McMillan, John 1987. Auctions and bidding. *Journal of Economic Literature* 25:699–738.

Miller, Mark S. and Drexler, K. Eric 1988. Markets and computation: Agoric open systems. In Huberman [1988]. 133–176.

Newell, Allen 1982. The knowledge level. *Artificial Intelligence* 18:87–127.

Reiter, Stanley 1986. Information incentive and performance in the (new)^2 welfare economics. In Reiter, Stanley, editor 1986, *Studies in Mathematical Economics.* MAA Studies in Mathematics.

Scarf, Herbert E. 1984. The computation of equilibrium prices. In Scarf, Herbert E. and Shoven, John B., editors 1984, *Applied General Equilibrium Analysis.* Cambridge University Press, Cambridge. 1–49.

Shenker, Scott 1991. Congestion control in computer networks: An exercise in cost-sharing. Prepared for delivery at Annual Meeting of the American Political Science Association.

Shoham, Yoav 1990. Agent-oriented programming. Technical Report STAN-CS-90-1335, Stanford University Computer Science Department.

Shoven, John B. and Whalley, John 1984. Applied general-equilibrium models of taxation and international trade: An introduction and survey. *Journal of Economic Literature* 22:1007–1051.

Waldspurger, Carl A.; Hogg, Tad; Huberman, Bernardo A.; et al. 1992. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering* 18:103–117.

Wellman, Michael P. 1991. Review of Huberman [1988]. *Artificial Intelligence* 52:205–218.