

Non-cooperative Planning in Multi-Agent, Resource-Constrained Environments with Markets for Reservations

Jorge E. Prado and Peter R. Wurman

North Carolina State University
Suite 110, Venture 1; Centennial Campus
Raleigh, NC 27695-7535 USA
jprado@acm.org , wurman@csc.ncsu.edu

Abstract

Queuing and accepting reservations are common social methods for allocating scarce resources in our daily lives. However, neither method is flexible when they are components of large, complex plans. In this paper we investigate the use of mobile devices that provide timely information, facilitate planning, and enable the trading of reservations. We investigate the behavior of a closed society of simple agents competing for scarce resources. The results of the experiments demonstrate that a simple reservation mechanism can actually reduce the social welfare under certain conditions, but tradable reservations and clairvoyance each improve it. While in many situations queues are unavoidable, better information and more flexibility in reservation handling can facilitate planning.

Introduction

In modern society, we spend a great deal of time performing activities in physical, multi-user environments that have limited capacity to serve their users. The quality of our experience often depends upon our implicit coordination with other users. If, for example, we go to the grocery store or to a restaurant at the same time as all of our neighbors, the wait will be interminable. One common solution to this problem is to make some resources reservable (i.e., restaurants, movies). Another is to make available information about the current wait-time to improve the decision making (i.e., traffic reports). Still another solution is for the user to learn a function describing the expected load on a resource (i.e., grocery stores are less busy very late at night).

Software agents operating on mobile devices have the potential to greatly improve our ability to operate in these complex environments. These agents can help the user plan an itinerary, can keep the user informed of changes in the environment that may warrant a change of plans, and can dynamically manage reservations for resources that are required for the current plan.

In this paper we consider a prototypical multi-user environment that has reservable, capacity-constrained resources. The model we examine is particularly well suited to describe closed, densely populated attractions like amusement parks, museums, and aquariums, but can also easily describe more

general environments. We analyze the system performance, measured in terms of social welfare, as we adjust the following characteristics:

1. **Information Quality:** Often a person arrives at a restaurant and finds it crowded. If she had known in advance the number of people waiting at the restaurant, she may have chosen another time, or another restaurant. A *clairvoyant* person would know the expected wait time of every place before deciding where to go.
2. **Tradable Reservations:** Imagine a person who bought the last ticket for the evening show at the theatre. His plan changes suddenly and now he has to catch a plane this afternoon, so he has to get rid of the theatre reservations, and get an airline ticket. He may not be able to buy a ticket from the airline, but there is another traveler who would be willing to cancel her trip if she could get theatre tickets. Markets for the reservations would enable these two people to adjust their plans in a satisfactory manner.
3. **Planning:** Consider a person in an amusement park who really wants to ride a roller coaster. Acquiring a reservation for the roller coaster right now is impossible, and the queue is very long. If he plans ahead, he may find that the reservations for the roller coaster in a few hours are easy to get, and he can construct a plan around the roller coaster reservation that involves other attractions.

Model

We consider a discrete time, finite horizon model of the environment. This would be consistent, for instance, with a single day visit to an amusement park. A discrete time interval, t , is an element of the ordered set $\{0, \dots, T\}$. The rest of the model consists of two parts: the environment and the agents.

The Environment

We model the environment as a connected graph. The nodes of the graph represent the activities (rides, exhibitions, restaurants, etc.). The edges of the graph represent the walkways connecting the activities. The set of all nodes is denoted N and the set of all edges L . Individual nodes and edges are designated n and l , respectively.

Nodes have attributes that govern how agents interact with them. The *maximum capacity* of node n , denoted c_n , is the

maximum number of agents that can use activity n at the same time. Naturally, $c_n > 0$. The *duration* is the amount of time an agent spends inside node n , and is represented by $s_n > 0$. The *admittance frequency* is the amount of time between the admittance of one group of agents to the activity and the admittance of the next group of agents. Admittance is periodic starting at $t = 0$. The admittance frequency of node n is denoted f_n , where $f_n > 0$.

These attributes allow us to simulate a wide variety of activities. For instance, roller coasters are modeled as having small capacity, short duration, and frequent admittance. Theatre shows have long duration, relatively large capacity, and infrequent admittance. A sit-down restaurant has moderately frequent admittance and long duration, while a cafeteria may admit one person every time step.

Many of the agents' decisions will require the agent to determine the *next admittance time*—the number of time steps from time t until node n admits another group of agents. Next admittance time is denoted $e_{n,t}$, and can be computed as:

$$e_{n,t} = \begin{cases} 0 & \text{if } (t \bmod f_n) = 0, \\ f_n - (t \bmod f_n) & \text{otherwise.} \end{cases}$$

Each activity has an ordered set of agents waiting in line to enter the attraction, called the *queue*. The queue at time t , is denoted by $Q_{n,t}$. The agents are represented by α_i , and the position of agent α_i is $p_{\alpha,i}$. The *queue length*, $q_{n,t}$, is the number of agents waiting in line to enter node n at time t , and is equal to the cardinality of $Q_{n,t}$.

An agent can compute the amount of time it expects to wait in line for node n with queue length $q_{n,t}$. The expected wait time, denoted $w_{n,t}$, is:

$$w_{n,t} = f_n * \left\lceil \frac{q_{n,t}}{c_n} \right\rceil + e_{n,t}.$$

The queues are FIFO (unless the agent has a reservation, described later). Agents can abandon the line at any time.

The edges of the graph represent walkways (or other means of transportation), that we call links. Each link connects two nodes, and has an associated traversal time, $d_{m,n}$, where m is the origin node, n is the destination node, and $d_{m,n} > 0$. If nodes m and n are not directly connected, then $d_{m,n}$ is undefined.¹ Note that $d_{m,n} = d_{n,m}$.

The graph generation process used in the simulation ensures that the graph will always be connected, that is, there is always a path from one node to another, using one or more links. Sometimes, the direct link from m to n is not the shortest path (i.e., it may be the scenic route). Let $D_{m,n}$ be the shortest path from m to n , using one or more links. $D_{m,n}$ may be less than $d_{m,n}$.

The Agents

The agents represent people who use the facilities in the environment. Each agent has attributes that describe its history,

¹Links could be modeled as activities with infinite capacity, admittance every time step, and duration $d_{m,n}$.




Figure 1: Utility Functions.

its current state, its utility for various actions, and its plan. An individual agent is denoted α , and the set of all agents A .

Agent α 's *position*, $p_{\alpha,t}$, is the node or link at which it is located at time t , where $p_{\alpha,t} \in N \cup L$. The agent's *time to finish*, $z_{\alpha,t}$, is the amount of time until agent α completes its current action.

$$z_{\alpha,t} = \begin{cases} d_{n,m} & \text{if agent } \alpha \text{ enters link } n \leftrightarrow m, \\ s_n & \text{if agent } \alpha \text{ enters activity } n, \\ z_{\alpha,t-1} - 1 & \text{if agent } \alpha \text{ is occupied} \\ & \text{and } z_{\alpha,t-1} > 0, \\ 1 & \text{if agent } \alpha \text{ decides to do nothing} \\ & \text{for 1 time step,} \\ 0 & \text{otherwise (including standing} \\ & \text{in line).} \end{cases}$$

The agent keeps track of its *history* in the form of the number of visits it has made to each activity. An agent is considered to have visited a node only if it actually partakes in the activity. The number of times α has entered node n is denoted $h_{\alpha,n,t}$.

$$h_{\alpha,n,t} = \begin{cases} 0 & \text{if } t = 0, \\ h_{\alpha,n,t-1} + 1 & \text{if agent } \alpha \text{ enters node } n \text{ at} \\ & \text{time } t, \\ h_{\alpha,n,t-1} & \text{otherwise.} \end{cases}$$

The agent derives value by participating in the activities of the environment. The amount of value is defined by the agent's *utility function*. In general, the utility that an agent gets from node n is a function of the number of times that the agent has already visited n . For simplicity, we consider only three types of utility functions, each representing a simplified, but common, scenario in these environments. Figure 1 shows the three types of utility functions: (a) utility that is constant, (b) utility that decreases linearly with the number of visits, and (c) utility that is constant to some threshold number of visits, and zero thereafter.

Let $v_{\alpha,n}(h_{\alpha,n,t})$ represent agent α 's value for entering n as a function of the number of previous visits. We assume that agents have quasilinear utility; they are not budget constrained, but they do have some other use for money outside the environment. In practice, we expect that it will be more acceptable to be granted a budget in a currency local to the environment.

Agents are utility maximizers and will follow the plan that gives them the highest utility in their planning horizon.

The accumulated utility earned by agent α from time t_0 until time t is represented by $U_{\alpha,t}$.

$$U_{\alpha,t} = \sum_n \sum_{i=1}^{h_{\alpha,n,t}} v_{\alpha,n}(i).$$

We assume that agents don't share information between them. One agent doesn't know the valuations or the plans of other agents, and it does not try to predict what the other agents will do next.

The agent's current state is the agent's position and the agent's current action. The possible actions are:

Walk: An agent can move from one place to another when it is not inside an attraction. Once it decides to go to another place, it cannot change actions until it reaches another node. For example, if an agent is going from A to D, using the links $A \leftrightarrow B$, $B \leftrightarrow C$, and $C \leftrightarrow D$ on its way, it can only change its destination when reaching B, C or D. The result of agent α walking from node m to node n at time t assuming the link $m \leftrightarrow n$ exists, is:

$$p_{\alpha,t+k} = n, \text{ where } k = D_{m,n}.$$

All agents walk at the same pace; there are no faster or slower agents. Although this assumption is not realistic, relaxing it would not change the behavior of the model.

Enter Queue: When an agent arrives at a node and wants to enter, it has to go to the end of the line and wait its turn. When two or more agents arrive to a queue at the same time, the tie is broken randomly.

Enter Activity: When the agent is within c_n of the front of the queue and the node is admitting users, the agent enters the activity. The amount of time the agent will be inside node n is s_n , and the agent is removed from the queue. After entering a node, agents cannot leave it before the activity is over.

Wait: An agent may prefer to wait, doing nothing.

The perception of the agent determines how far the agent can see. Specifically, it represents the list of places where the agent can see the queue length, and can compute the expected wait time. A *myopic* agent can see only its current location, and the neighbors that surround it (nodes directly connected to the node where the agent is located). If the agent is located in node n , then it can see the queue at node n and the queues at all nodes m where $d_{n,m}$ exists. *Clairvoyant* agents can see the queue length of all the activities.

In Figure 2, an agent in node 1 with myopic perception can see the queue lengths of its current position (node 1) and the immediate neighbors (nodes: 2, 3, 4 and 7). With clairvoyance, it would see all the nodes.

Enhanced Environments

We enhance the environments described in the previous section by adding tradable reservations and clairvoyance to improve planning. The reservations can be exchanged in markets, and we assume that the agents (on mobile devices) are endowed with communication technology that enables them to participate in the market while being carried around the environment.




Figure 2: Myopic and Clairvoyant Perception

A *reservation* for node n at time t is denoted $r_{n,t}$. If an agent holds $r_{n,t}$, it will be admitted to node n at time t without waiting in line. The number of reservations distributed in node n for each admittance time is represented by ρ_n , $0 \leq \rho_n \leq c_n$. Note that $r_{n,t}$ cannot be used after time t .

If an agent doesn't hold a reservation, it must account for the effect that other agent's reservations will have on its wait time. The expected time spent in line for an agent that doesn't have a reservation is

$$w_{n,t} = f_n * \left[\frac{q_{n,t}}{c_n - \rho_n} \right] + e_{n,t}.$$

We refer to the reservations that an agent currently holds as its *endowment*, denoted $E_{\alpha,t}$. During the trading stage of iteration t , the agent is free to sell part or all of its endowment. Reservations that the agent purchases during the trading stage are either used immediately or added to the agent's endowment for the next iteration.

Reservations can be traded in an electronic marketplace. The market has one auction for each possible node and time where reservations are possible. As a practical matter, the number of reservable times may be regulated to keep the number of markets down. For instance, the markets may be opened for only the next two hours worth of activities, or only for the entry to activities on the quarter hour mark.

General Equilibrium Theory provides a set of sufficient conditions that ensure the existence of equilibrium prices and the Pareto optimality of the supported allocation (Mas-Colell, Whinston, & Green 1995). Unfortunately, the exchange economy defined by the agents in our model does not satisfy the conditions of the First Welfare Theorem. In particular, the goods are discrete and therefore violate the condition that preferences be convex. In addition, as the planning horizon increases, agents begin to construct plans that include complementary reservations, which violates the gross substitutes condition. The presence of these two violations will sometimes prevent the markets from converging. In order to make progress in the face of these failures, we have manipulated the auction to improve convergence at the expense of optimality.

Auction Rules

The reservations are traded in k -double auctions (Satterthwaite & Williams 1989) where $k = 1/2$. For the purposes

of this initial study, we assume that agents bid truthfully and state their actual willingness to pay for (or minimum willingness to sell) a particular reservation. Moreover, we assume the agents act competitively and take the prices announced by the auctioneer at face value.

The price quote generated by the auction is computed according to the standard M^{th} and $(M + 1)^{\text{st}}$ price rules. The buy and sell bids are sorted, and the M^{th} and $(M + 1)^{\text{st}}$ highest bids are identified, where M is the number of sell offers. These prices delineate the range of prices that will balance supply and demand. The M^{th} price corresponds to the ask quote, π_{ask} , and the $(M + 1)^{\text{st}}$ price is the bid quote, π_{bid} (Wurman, Walsh, & Wellman 1998).

After all the bids have been received, new quotes are computed and communicated to the agents. Agents can then update their bids in response to the new prices. The auction will reach equilibrium when no agent wants to change its bids, given the current prices of the reservations.

If the market fails to reach equilibrium, we introduce error in the utility computed by the agents by manipulating the price quotes. When a convergence failure occurs, the market will modify the prices announced by adding (subtracting) ϵ to the ask (bid) quote. This has the effect of making reservations seem more expensive to buyers, and less valuable to sellers. When buying r , agents will be told the ask price of r is not π_{ask} but $(\pi_{\text{ask}} + \epsilon)$, thus decreasing their expected utility of buying it. When selling r , agents will use the bid quote $(\pi_{\text{bid}} - \epsilon)$, decreasing their expected utility of selling it. If the market still fails to converge, ϵ is increased. Eventually the announced prices will reach values where no agent wishes to place a new bid, but in so doing the market sacrifices social efficiency.

After reaching quiescence, the markets clear. The exchange price is determined by using $k = 1/2$, that is, the middle point between the bid and ask quotes. All the sellers with bids below the exchange price will sell, and all the buyers with bids above the exchange price will buy. The sellers will transfer the reservation to the buyers, and buyers will reciprocate with money in the amount equal to the trading price. The detail of who exchanges the goods with whom is not important, because all the goods in a specific market are identical.

Planning

At each decision point, the agent evaluates its possible actions and selects a plan that generates the greatest utility. We assume the agent's planner generates only feasible plans.

Initially, consider plans that involve choosing only the next activity. Let ψ denote such a feasible plan. Suppose agent α is at node m and plan ψ involves going to node n and participating without a reservation. The value of going to node n at time t is $v_{\alpha,n}(h_{\alpha,n,t})$. The total amount of time that α will spend to complete the activity of node n is the time to finish its current activity $z_{\alpha,t}$ at node m , plus the travel time $D_{m,n}$, plus the time in line $w_{n,t}$, plus the duration s_n of n . Let

$$\tau(\psi) = z_{\alpha,t} + D_{m,n} + w_{n,t} + s_n.$$

The utility that α receives from ψ , will also include the

market value of the unused portion of the agent's endowment (i.e., the sum of bid prices of everything that it owns). Thus, the utility of the plan to go to node n is

$$u_{\alpha,t}(\psi) = \frac{v_{\alpha,n}(h_{\alpha,n,t})}{\tau(\psi)} + \sum_{r \in E_{\alpha,t}} \pi_r.$$

If the agent decides to use a reservation that it owns, the time required to enter node n , $w_{n,t}$, will be the time until the reservation can be used, and the income from the agent's endowment will not include the reservation \hat{r} that it intends to use. The utility will be:

$$u_{\alpha,t}(\psi) = \frac{v_{\alpha,n}(h_{\alpha,n,t})}{\tau(\psi)} + \sum_{r \in E_{\alpha,t}} \pi_r - \pi_{\hat{r}}.$$

If the agent needs to buy the reservation, r , that completes the plan, the utility expression will have the same form as the above with the exception that the agent will need to use the ask price of r , and r is not an element of $E_{\alpha,t}$. Notice that in addition to being truthful, the agent is pessimistically using the bid and ask quotes rather than computing the true $k = 1/2$ transaction prices.

When planning farther ahead, it is possible that a plan will require more than one reservation. To compute the utility of a complex plan, ψ , we compute the average value of all the activities considered in the plan horizon. Let V denote the average value of a complex plan. The utility will be:

$$u_{\alpha,t} = V + \sum_{r \in E_{\alpha,t}} \pi_r - \sum_{\hat{r} \text{ used in } \psi} \pi_{\hat{r}}.$$

When two plans provide the same utility to an agent, the agent selects one according to the following rules:

- A plan that uses reservations has priority over plans that involve waiting in line (because of the uncertainty of the future queue length).
- A plan that uses a reservation that the agent owns has priority over a plan that involves buying reservations.
- Doing nothing has the lowest priority.

These rules basically encode an agent's preference of more certain actions over those that involve more uncertainty.

For the simulations described in this paper, we have used a very simple planning algorithm with a limited horizon. As part of the larger research agenda, we plan to integrate state of the art planning systems into the simulation.

Bidding

Agents interact with the market by placing bids to buy and sell reservations. The first step in determining bids is to compute the utility of each plan. We will also assume that the agents are truthful and bid their exact willingness to pay (sell).

Each agent will find the plan that gives him the highest utility. We will denote the highest utility $u_{\alpha,t}^*$ and the next highest utility $\hat{u}_{\alpha,t}$. Let V^* be the average value of the plan and π^* the price of the reservation used in the highest utility

Plan	$D_{m,n}$	$w_{n,t}$	s_n	$V_{n,t}$	$V = \frac{V_{n,t}}{\pi}$	$V - \pi$	offer	$b_{\alpha,n,t}$
Ride 1 with res.	0	1	1	10	$10/2 = 5$	$5 - \pi_1$	Sell	$5 - 3.75 + \pi_2$
Ride 2 with res.	2	0	2	15	$15/4 = 3.75$	$3.75 - \pi_2$	Buy	$3.75 - 5 + \pi_1$
Ride 2	2	3	2	15	$15/7 = 2.14$	2.14		
Ride 1	0	5	1	10	$10/6 = 1.66$	1.66		

Table 1: Example bid calculation when the agent owns a reservation for Ride 1.

plan, as calculated before. Similarly, \hat{V} and $\hat{\pi}$ represent the valuation and the price of the reservation used in the second highest utility plan.

The value that α has for a reservation, r , is the amount that would make the agent indifferent between the plan with r and the best alternative. To compute a bid, $b_{\alpha,n,t}$, for a reservation, r^* , that is part of the best plan, the agent compares the plan to the second best plan. The agent is indifferent when $u_{\alpha,t}^* = \hat{u}_{\alpha,t}$, that is,

$$\begin{aligned} V^* - \pi^* &= \hat{V} - \hat{\pi}, \\ \pi^* &= V^* - (\hat{V} - \hat{\pi}). \end{aligned}$$

Thus, the truthful agent will bid $b_{\alpha,n,t} = V^* - (\hat{V} - \hat{\pi})$. If $r^* \in E_{\alpha,t}$, α is willing to sell r^* as long as it receives at least $b_{\alpha,n,t}$. If r^* is not in α 's endowment, α is willing to buy r^* for up to $b_{\alpha,n,t}$.

Bids for reservations that are part of the other plans are assessed in the same manner. Let $u_{\alpha,t}$ be the utility of a arbitrary other plan. Again, the agent finds the indifference point where $u_{\alpha,t}^* = u_{\alpha,t}$. It follows that

$$\begin{aligned} V^* - \pi^* &= V - \pi, \\ \pi &= V - (V^* - \pi^*). \end{aligned}$$

And the truthful bid will be $b_{\alpha,n,t} = V - (V^* - \pi^*)$.

Table 1 shows an example of the bids computed by an agent, assuming the agent is in node 1 at time 0 and there are only 2 nodes. The price of the reservation for ride 1 at time 1 (that he owns) is π_1 and for ride 2 at time 2 is π_2 . Also $z_{\alpha,t} = 0$, $\pi_2 < 1.61$ and $\pi_1 < 1.25 + \pi_2$.

The process is more difficult when plans involve more than one activity because of the difficulty of ascribing the value of a plan to the component reservations. One solution is to permit combinatorial bidding (de Vries & Vohra to appear). However, due to the size of real problem instances, that may prove intractable. Investigating combinatorial bidding in this environment is part of the future work.

In the meantime, we ascribe value to component reservations with the following heuristic. Let the superscript $*$ denote the highest utility plan that uses r . Let π_{others}^* denote the sum of the prices of all other reservations required in plan $*$. Using the same nomenclature presented before,

the amount to bid for each reservation in the plan will be:

$$\begin{aligned} u_{\alpha,t}^* &= \hat{u}_{\alpha,t}, \\ V^* - \pi_{n,t}^* - \pi_{others}^* &= \hat{V} - \hat{\pi}, \\ \pi_{n,t}^* &= V^* - \pi_{others}^* - (\hat{V} - \hat{\pi}) \\ b_{\alpha,n,t} &= V^* - \pi_{others}^* - (\hat{V} - \hat{\pi}) \end{aligned}$$

All the agents will place their bids: buy bids for reservations they don't have, and sell bids for reservations that they have. If an agent holds a reservation that it cannot use in any of its plans (because, say, it can't get to the node by the reservation time), it will offer to sell the reservation for 0.

Results

The model used for the simulations has 10 nodes, 100 agents and 100 time slots. The total capacity of the model is 36.24 agents per time slot. Thus only slightly more than a third of all agents can be active at one time. At the beginning of the simulation, each agent is randomly positioned next to an activity.

We included two methods of allocating the reservations. One is to randomly distribute reservations at the beginning of the simulation. The second is to assign reservations on a first-come, first-served basis (e.g. non-random).

We measured the social welfare of the system under different conditions: no reservations, randomly distributed reservations without trading, non-randomly distributed reservations without trading, and trading. The four types of simulations were run with both myopic and clairvoyant perception and with varying levels of reservable capacity. Figures 3, 4 and 5 show the results when 20%, 40% and 70% (resp.) of the capacity is reservable. As a baseline, we observe that the lines corresponding to myopic and clairvoyant agents using no reservations are constant in the three graphs.

Several trends are visible in the graphs. First, the overall quality of the solutions increases as the proportion of the capacity increases. Second, the difference between the myopic results and clairvoyant results is quite large. This result primarily stresses the importance of making good information available to the agents. Third, with or without clairvoyance, trading improves social welfare. However, adding reservations without trading sometimes decreases the social welfare because a large percentage of the reservations go unused, increasing the uncertainty in predicting the queue's progress. Finally, increasing the planning horizon does not clearly help or hurt. This inconclusive result is due, in part, to the fact that the horizons are all quite small and the planning algorithm is not sophisticated. This is a primary area for future study.

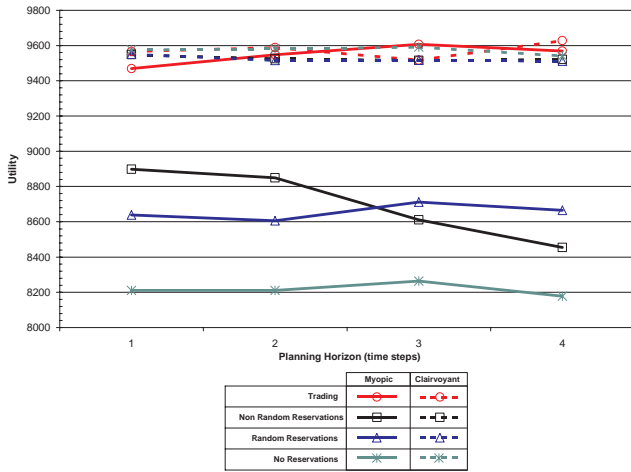


Figure 3: Utility vs Planning Horizon with 20% reservable capacity




Figure 5: Utility vs Planning Horizon with 70% reservable capacity




Figure 4: Utility vs Planning Horizon with 40% reservable capacity

Non-randomly distributed reservations give better results than randomly distributed reservations when the planning horizon is small. When planning ahead, more reservations requested by the agents are not used and wasted, because they change their plans and the unused reservations cannot be reallocated.

Figure 6 analyzes the evolution of prices over time, using Clairvoyance, trading and planning 1 time step ahead. The dotted line shows the queue length of the node analyzed. Each continuous line represents the price of a reservation over time. A circle over a continuous line indicate that the corresponding reservation was traded at that price and time. It is clear from the graph that the longer queue length makes the reservations more valuable, and that price goes up as reservation time nears. In most cases the reservations are traded early and the agents keep their plans of using a reservation once they buy it.

There are two sources of “noise” in the agent’s decision process that may help explain the results.

1. **Uncertain queue length:** When an agent decides to go to a node without a reservation, the queue length will change by the time it arrives. If many agents converge on a node because they heard the line was short, a good portion of them will be disappointed.
2. **Convergence problems:** Each time the markets fail to converge, a bias is introduced in the quotes to force convergence. This bias can produce a sub-optimal allocation, although the result is still guaranteed to be better than not trading. We expect the convergence problems to be more frequent when the planning horizon is increased.

Even with those problems, it is clear that using tradable reservations improves the social welfare.

Related Work

A great deal of recent work has studied software agents in electronic markets (Chavez & Maes 1996; Wellman 1993). However, to our knowledge, no one has modeled the types of environments that we have addressed in this paper. The Electric Elves project (Chalupsky *et al.* 2001) is one project that studies the impact of mobile assistants that help work-groups coordinate their activities, but currently the project does not involve market interactions. The Trading Agent Competition (Wellman *et al.* 2001b) is a framework for studying trading strategies in a marketplace for travel resources. The agents in TAC, however, represent groups of people with travel preferences, but do not individually have an overly complex scheduling problem. Some work has been done on market-based scheduling (Clearwater 1995; Wellman *et al.* 2001a) but the constraints on the models differ in significant ways from the model presented here.

We expect that there is a lot of work in Queuing Theory that is relevant to the model presented here. However, we are unaware of any that ties queuing directly to market mechanisms. We do expect Queuing Theory to have relevant

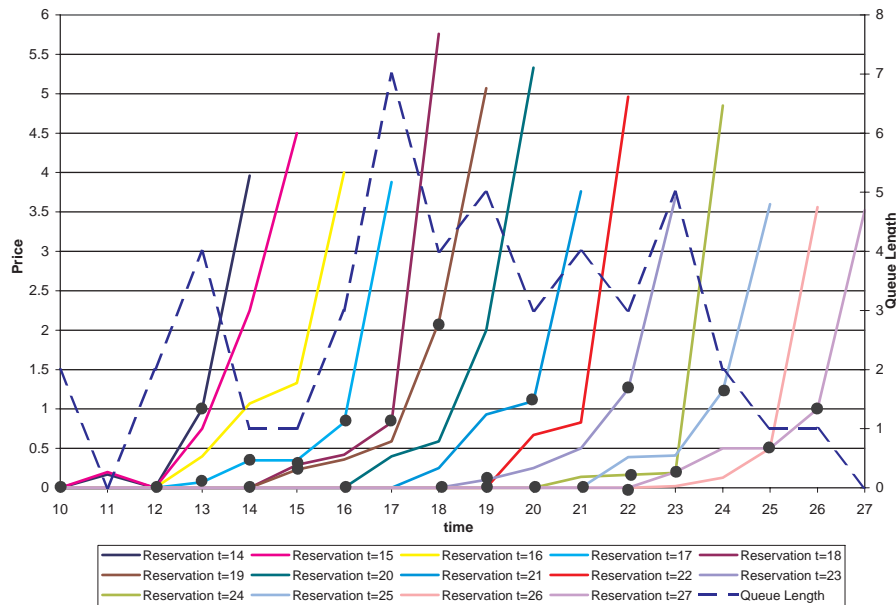


Figure 6: Evolution of Queue Length and Prices over Time (Circles Indicate Trades)

connections to the planning aspects of the agent's decision problems.

Conclusions

We present a formal model of a common, multi-agent coordination problem in which agents are non-cooperative and resources are limited. Through simulation, we have experimented with the performance of the social system when the resources are reservable, and when agents can trade the reservation in a marketplace. We found that better information and trading reservations both improve the social welfare, but reservations alone are not always beneficial.

We plan to continue to extend the model. In particular, we are interested in exploring better market mechanisms (eg. combinatorial auctions), better planning methods (eg. partial-order planning). We also plan to study the ability of the enhanced system to accommodate plan deviations, as when a human user suddenly becomes interested in an activity that was not previously in the plan.

Acknowledgements

This paper is based on the first author's Masters Thesis. We wish to thank R. Michael Young, Robert St. Amant, the members of the Intelligent Commerce Research Group at NCSU, and the anonymous reviewers for their valuable comments.

References

Chalupsky, H.; Gil, Y.; Knoblock, C. A.; Lerman, K.; Oh, J.; Pynadath, D.; Russ, T.; and Tambe, M. 2001. Electric elves: Applying agent technology to support human

organizations. In *Thirteenth Innovative Applications of Artificial Intelligence Conference*, 51–58.

Chavez, A., and Maes, P. 1996. Kasbah: An agent marketplace for buying and selling goods. In *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*.

Clearwater, S. 1995. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific.

de Vries, S., and Vohra, R. to appear. Combinatorial auctions: A survey. *INFORMS Journal on Computing*.

Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.

Satterthwaite, M. A., and Williams, S. R. 1989. Bilateral trade with the sealed bid k -double auction: Existence and efficiency. *Journal of Economic Theory* 48:107–133.

Wellman, M. P.; Walsh, W. E.; Wurman, P. R.; and MacKie-Mason, J. K. 2001a. Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35(1-2):271–303.

Wellman, M. P.; Wurman, P. R.; O'Malley, K. A.; Banger, R.; Lin, S.-D.; Reeves, D.; and Walsh, W. E. 2001b. Designing the market game for a trading agent competition. *IEEE Internet Computing* 5(2):43–51.

Wellman, M. P. 1993. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research* 1:1–23.

Wurman, P. R.; Walsh, W. E.; and Wellman, M. P. 1998. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems* 24:17–27.