

Experiments with Planning and Markets in Multiagent Systems

Naveenkumar Muguda, Peter R. Wurman, and R. Michael Young
North Carolina State University
Department of Computer Science
Raleigh, NC 27695-7535 USA
{nvmuguda,wurman,young}@csc.ncsu.edu

Mobile devices are hand-held devices used to deliver time sensitive and locale specific information to the users. Users of multi-user environments with limited resources can be enabled with agents running on mobile devices as assistants to improve their ability to plan activities in the space. This paper discusses experiments to characterize the benefits of planning in such an environment, particularly when the resources can be reserved and the reservations traded in a market. We observe that with tradable reservations the social welfare increases with increase in planning horizon. We also observe that tradable reservations and clairvoyance help the users satisfy their preferences and the constraints imposed by environment.

Categories and Subject Descriptors: K.4.4 [Computers and Society]: Electronic Commerce

Additional Key Words and Phrases: Trading Agents

1. INTRODUCTION

The impending availability of mobile, networked computing devices will enable a variety of context-aware software systems. Current efforts by the telecommunications and computing industries are focused on using the technologies for either the delivery of contextually-relevant information to mobile users, or to enhance communication between disperse parties. However, as the hardware matures and becomes more powerful, more sophisticated applications will become practical. We envision applications that enhance mobile devices with agent-based interfaces that help users plan activities in networked environments.

The problem domain consists of physical, multi-user environments in which users participate in complex activities whose sub-tasks have limited capacity. We consider a prototypical multi-user model that is particularly well suited to describe closed environments like warehouses, hospitals, amusement parks, and museums [Kurumatani 2002; Prado and Wurman 2002], but which can also easily describe more open environments with limited resources like movie theaters and restaurants.

In such environments, the quality of our experience often depends upon our implicit

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2004 ACM /2004/-0034 \$5.00

coordination with other users. If, for example, we go to the grocery store or a restaurant at the same time as all of our neighbors, the wait will be interminable. If two patients are delivered to radiology at the same time, one will have to wait while the other is served. One common solution to this problem is to make some resources reservable (e.g., restaurants, movies, appointments with doctors). Another is to make available information about the current wait-time to improve the decision making (e.g., taking a number at a deli). Still another solution is for the user to learn a function describing the expected load on a resource (e.g., grocery stores are least busy during dinner time or very late at night).

Distinguishing characteristics of these disparate environments include:

- Activities have limited capacity. The ability of a resource to serve users is measured in terms of the number of persons admitted, the frequency of admission, and the duration of the activity.
- Users have their own volition. An agent may make recommendations, but the human may exercise her autonomy and deviate from the recommended itinerary.
- User preferences over the various activities may initially be unknown. In some scenarios, a user may wish to participate in an activity multiple times, and her value for each interaction may differ.
- Congestion will affect a user's preference among activities. By implication, the user's preference among possible plans depends upon the actions of other agents.
- Some resources can be reserved ahead of time, such as a reservation for an MRI machine, tickets to a special art exhibit, an IMAX show, or an amusement ride.

Traditional approaches to managing resource contention in these environments include architecting the physical space to encourage the flow of people through it, adding distraction via televisions to entertain people in the queue, and allowing reservations for particular resource-bounded activities. Disney, for instance, permits guests to make reservations for a ride using its *fastpass* system which grants the user the right to attend the ride during a particular time window. Given the success of such systems, we expect to see more reservable activities at amusement parks in the future. However, none of these approaches enable the kind of flexibility or information needed by users to individualize their experience.

We envision two enhancements to the basic model. First, we add the ability to buy and sell reservations in an electronic marketplace. A marketplace for reservations give users the flexibility to adjust her commitments when she revises her plan in response to fatigue, changes in the environment, changes in her interests due to gained knowledge, or other factors.

Second, we extend the infrastructure to allow context-aware mobile communication. Specifically, we outfit each user with a mobile device that hosts a software assistant. These assistants help the user plan an itinerary, keep the user informed of changes in the environment that may warrant a change of plans, and dynamically manage reservations for resources that are components of the current plan.

The typical approach in multiagent systems research views agents as taskable entities, capable of being delegated duties and performing them with little or no user supervision [Greenwald and Stone 2001; Maes 1994; Rich and Sidner 1997]; the agent performs tasks on behalf of a user, and retains responsibility for executing the plan that achieves the goals. In contrast, the agents we envision rely on an alternate form of agent-human interaction where the primary activities in the domain are carried out by human users rather than

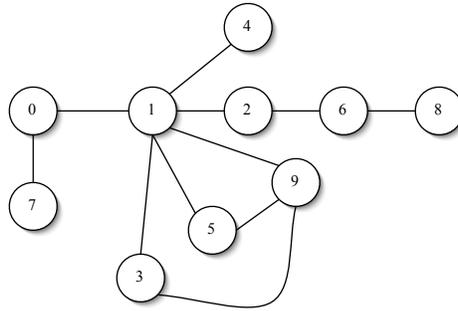


Fig. 1. An environment with ten nodes.

by computational agents.

An agent that can support this kind of interactivity must be able to communicate with the human user—we assume plan-based dialog—and be able to analyze and make marketplace decisions. In this paper we report on our initial investigation of the interaction between planning and trading, and the effect these capabilities have on the system as a whole. In Section 2 we describe our abstract model of agents and the environment in which they operate. Section 3 describes the marketplace that enables trading, and the planning and bidding processes of the agents. Section 4 presents the results of our experiments with the environment. Section 5 describes the related and future work.

2. 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, or a shift in a hospital. A discrete time interval, t , is an element of the ordered set $\{0, \dots, T\}$. The model consists of agents, α , and an undirected graph representing the environment.

2.1 The Environment

We model the environment as a connected graph in which the nodes represent the activities (rides, exhibitions, restaurants, etc.) and the edges 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. Figure 1 illustrates the graph used later in our experiments.

Nodes have attributes that govern how agents interact with them. The *admittance frequency* of node n , denoted f_n , 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 capacity*, denoted c_n , is the size of the admitted group. The *duration* is the amount of time an agent spends engaged in the activity at node n , and is represented by s_n . The parameters c_n , s_n , and f_n are strictly positive.

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. In this version of the model, however, we do assume

each activity has a deterministic duration.

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 potentially has an ordered set of agents waiting in line to engage in the activity, called the *queue*. Node n 's queue at time t is denoted by $Q_{n,t}$. 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 $Q_{n,t}$. The length of the queue in front of agent α is denoted $q_{n,t}^\alpha$ which is used when an agent needs to compute the amount of time it expects to wait in line for node n . The expected wait time, denoted $w_{n,t}$, is:

$$w_{n,t} = f_n * \left\lfloor \frac{q_{n,t}^\alpha}{c_n} \right\rfloor + e_{n,t}.$$

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), called 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. In the model, $d_{m,n} = d_{n,m}$.

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. Let $D_{m,n}$ be the shortest path from m to n , 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).

2.2 The Agents

In the simulation, agents represent people equipped with mobile devices while engaged in the environment. Each agent has attributes that describe its history, its current state, its utility for various actions, and its plan. 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, and is determined by the transition function:

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

The agent's *history* is simply the number of times it has engaged in each activity. The number of times α has entered node n is denoted $h_{\alpha,n,t}$, and is initially zero. Thereafter, it is incremented each time the agent enters the activity at node n .

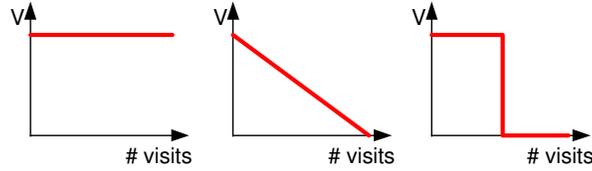


Fig. 2. Utility Functions.

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 the three prototype utility functions shown in Figure 2: (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.

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.

Let $v_{\alpha,n}(h_{\alpha,n,t})$ represent agent α 's marginal value for entering n as a function of the number of previous visits. The utility achieved by agent α between time t_0 and time t is

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

where $\Delta\mu_{\alpha}$ is the change in the monetary state of the agent between t_0 and t .

Agents are boundedly-rational utility maximizers and will follow the plan that gives them the highest utility in their planning horizon. We assume that agents do not share information between them. One agent does not 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 engaged in an activity. Once it decides to go to another place, it cannot change actions until it reaches another node. All agents walk at the same pace; there are no faster or slower agents. Although this assumption is not realistic, relaxing it would not materially 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 in the same time step, they are sequenced 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 and is removed from the queue. The amount of time the agent will be inside node n is s_n .

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

The perception of the agent determines what information about the environment the agent can see. Specifically, it represents the list of places the agent can see, and the subset of those whose queue length is visible. A *clairvoyant* agent can see the geometry and queue length of all the nodes. A *myopic* agent can see only its current location and directly

connected neighbors, and their respective queue lengths. If the agent is located in node n , it will consider plans that involve only node n and nodes m where $d_{n,m}$ exists. *Myopic agents with maps* know the geometry of the environment, but can see the queue lengths for only those nodes directly connected to their current node. For nodes not directly connected, the myopic agent with a map uses a default queue length.

3. ENHANCEMENTS TO THE ENVIRONMENT

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.

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 before or after time t .

When computing its expected wait time for a node for which it does not have a reservation, the agent must account for the fact that the node's capacity to admit users from the queue is reduced by the portion set aside for reservations. Thus,

$$w_{n,t} = f_n * \left\lfloor \frac{q_{n,t}^\alpha}{c_n - \rho_n} \right\rfloor + e_{n,t}.$$

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.

3.1 The Marketplace

Reservations can be traded in an electronic marketplace with one auction for each possible node and time where reservations are possible.¹ The auctions follow the k -double auction [Satterthwaite and Williams 1989] rules, 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 et al. 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.

¹As a practical matter, the number of reservable times may be regulated to keep the number of auctions 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 et al. 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.

If the market fails to reach equilibrium, we introduce bias in the utility computed by the agents by manipulating the price quotes. When a convergence failure occurs, the market modifies 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 are 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 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 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 particular auction are identical.

3.2 Planning

At each decision point, the agent evaluates its possible actions and selects a plan that generates the greatest utility at the current market prices. We assume the agent's planner generates only feasible plans, but that candidate plans may have different lengths. Thus, the agent compares plans based on utility per unit time. Let ψ denote a feasible plan specifying the sequence of nodes and whether a visit involves the use of a reservation. Let $\tau(\psi)$ be the duration of the plan, $|n \in \psi|$ be the number of occurrences of n in ψ , and $R(\psi)$ be the set of reservations required to execute ψ . The value the agent expects to get by participating in the plan ψ is

$$V_{\alpha,t}(\psi) = \sum_n \sum_{i=1}^{|n \in \psi|} v_{\alpha,n}(h_{\alpha,n,t} + i).$$

The utility per unit time of plan ψ is the plan's value plus any changes in monetary state divided by the duration of the plan. For brevity, denote the set of reservations in its endowment that the agent will not need as $S = E_{\alpha,t} \setminus R(\psi)$, and the set that it will need to buy as $B = R(\psi) \setminus E_{\alpha,t}$. The utility of ψ is,

$$U_{\alpha,t}(\psi) = \frac{V_{\alpha,t}(\psi) + \sum_{r \in S} \pi_r^{\text{bid}} - \sum_{r \in B} \pi_r^{\text{ask}}}{\tau(\psi)}. \quad (1)$$

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 simple planning algorithm with a finite horizon. Unlike our previous work [Prado and Wurman 2002] in which the horizon was restricted to four time steps, the plans described in this study are limited by the number of actions, and we look ahead as far as eight nodes. The algorithm uses beam search to enumerate a subset of the feasible plans which are then evaluated based upon their utility. Our larger research agenda includes incorporating a traditional AI planner as part of the plan-based dialog interface. The subset of plans computed by our beam search can be viewed as linearizations passed down from the user interface.

3.3 Bidding

Agents interact with the market by placing bids to buy and sell reservations. We assume that the agents are competitive and believe that π^{bid} and π^{ask} represent the true costs of selling or buying. An agent bids by computing its marginal value for any reservation and offering to buy it at that price (if it does not own the reservation), or sell it at that price (if it does own the reservation).

In the following, it is convenient to assume that the agent sells its endowment into the market and then buys back the reservations that are required to complete a plan. Let ψ^* denote the plan with the highest utility that uses r^* , and ψ be the plan with the highest utility without using r . To compute the marginal value of a reservation, the agent computes the price for r^* that would make it indifferent between ψ^* and ψ . In other words, let

$$U_{\alpha,t}(\psi^*) = U_{\alpha,t}(\psi)$$

and solve for the price of r^* . Substituting equation (1) with R^* denoting the reservations required for plan ψ^* and R the reservations for ψ , and rearranging the terms, we get

$$\pi_{r^*} = V_{\alpha,t}(\psi^*) - \sum_{r \in R^* \setminus r^*} \pi_r - \frac{\tau(\psi^*)}{\tau(\psi)} (V_{\alpha,t}(\psi) - \sum_{r \in R} \pi_r).$$

If r^* is part of the endowment of the agent, then the agent would place a sell offer at this price. If r^* is not part of the agent's endowment, it would place a buy offer. If an agent holds a reservation that is not part of any feasible plan (because, say, it cannot get to the node by the reservation time), it will offer to sell the reservation for 0.

4. RESULTS

The model used for the simulations has 10 nodes, 100 agents and 100 time slots. The total capacity of the model is 25.8 agents per time slot. Thus, only a quarter of all agents can be engaged at one time. At the beginning of the simulation, each agent is randomly positioned next to an activity. If reservations are permitted, they are randomly distributed

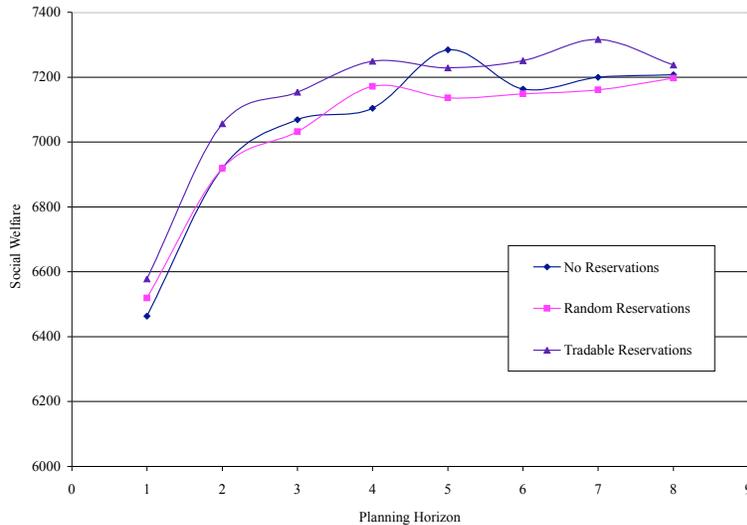


Fig. 3. Social welfare as a function of planning horizon for clairvoyant agents.

at the beginning of the simulation. The experiments were designed to capture the effects of providing better information (in the form of clairvoyance), enabling tradable reservations, and varying the level of reservable capacity. In addition, we studied the effects of having ordering constraints, geographic islands, and aversions. In most cases, the results represent a single run of the experiment, which explains some of the variation. However, we believe the trends are quite clear.

The primary measure of the performance of the system is the social welfare; an increase in the social welfare implies the system is doing a better job of allocating the resources to the agents who care the most for them. Figures 3 and 4 show the social welfare for clairvoyant and variations of myopic agents, respectively, when 40% of the capacity is reservable. In the case of clairvoyant and myopic agents with maps, increasing the planning horizon to at least three or four actions is clearly beneficial. Also, in the clairvoyant case, there is a clear benefit to allowing tradable reservations. When adjusted for the differing scales, the curves for clairvoyance and myopic with maps are quite similar. This suggests that the extra information a clairvoyant agent has—the queue length at nodes not directly connected to the current location—is not particularly valuable because it is not predictive of the future queue length.

In the case of myopic agents without maps, social welfare actually decreases as the planning horizon increases, and the benefits of trading are less pronounced. We suggest that this result is due to the fact that the agents make plans that do not include nodes

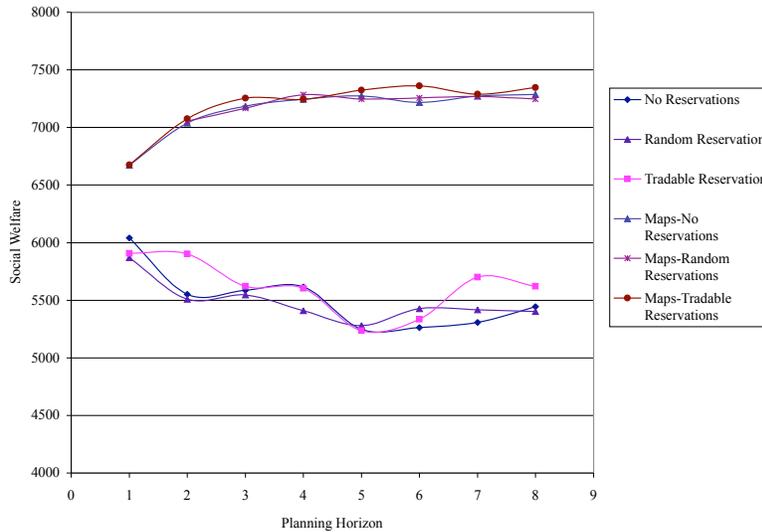


Fig. 4. Social welfare as a function of planning horizon for myopic agents.

that they cannot see, and as they move around the graph frequently abandon plans as they become exposed to other nodes in the graph.

Figure 5 shows the effects of changing the reservable capacity and planning horizon. Notice that permitting more of the capacity to be reserved generally resulted in an increase in social welfare; this is to be expected because more reservations enable the marketplace to facilitate the allocation of resources. Further, increasing the planning horizon is more beneficial as the percentage of reservable capacity increases. As more reservations enter the marketplace, agents have the ability to construct more reliable plans. The increased benefits of planning and trading is illustrated by increase in the number of trades as the planning horizon increases, as shown in Figure 6.

Another measure of the impact of reservable capacity is the effect it has on the behavior of the queues. We hypothesized that reservations would reduce the variance. This is somewhat borne out the the experiments. Table I shows the mean queue size and variance for each node when no reservations are allowed and when 40% of the capacity is reservable and agents are clairvoyant. Queue sizes are normalized to be the number of admittances that an agent would have to wait to engage in the activity (i.e., $q_{n,t}/(c_n - \rho_n)$). To avoid the startup and end effects, only data from time steps between 20 and 80 are used. Comparing the means between the two cases is inconclusive, but in seven of the ten nodes, the variance is reduced when reservations are allowed. This suggests that the reservations have a “smoothing” effect on the queue behavior.

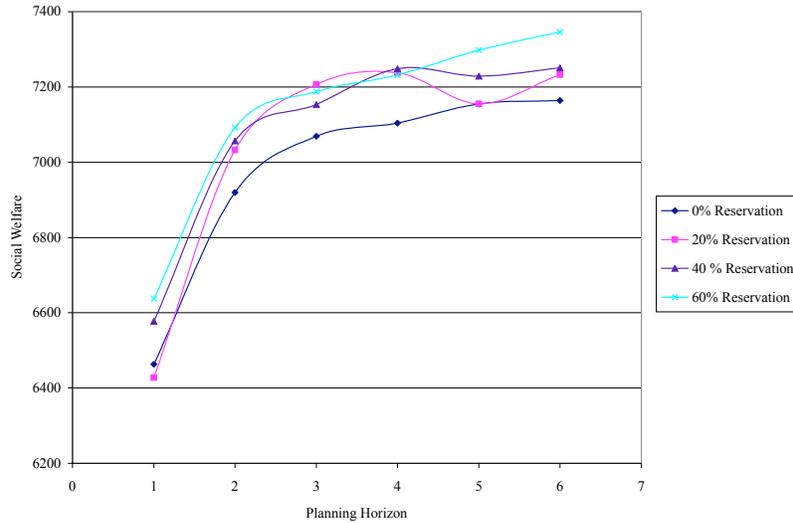


Fig. 5. Social welfare as a function of planning horizon for clairvoyant agents and varying reservable capacity.

n	c_n	f_n	No reservations		40% reservable	
			μ	Variance	μ	Variance
0	7	4	0.84	0.41	0.93	0.64
1	10	1	0	0	0	0
2	6	3	2.62	2.49	1.68	1.33
3	8	5	0.71	0.44	0.64	0.35
4	9	3	1.80	1.62	0.84	0.60
5	6	2	0.56	0.35	0.38	0.15
6	5	2	0.87	1.02	0.35	0.16
7	2	3	1.72	4.16	1.85	1.66
8	4	5	1.88	4.16	2.54	1.89
9	1	3	1.97	4.26	2.40	6.73

Table I. Mean and variances of queue sizes for each node in the graph.

One experiment we ran involved partitioning the environment into fully connected sub-graphs, or islands, with a single link between islands. This graphical structure captures designs like that of Disney's Animal Kingdomtm, in which travel between the regions representing different continents literally requires walking to the island in the middle of the park. The interesting conclusion from these experiments with myopic agents is that planning is beneficial up to a horizon about the cardinality of the island. Planning further ahead provides little added benefit.

We also constructed an experiment in which agents belong to one of three types. Two of

SIGecom Exchanges, Vol. 5, No. 1, 6 2004.

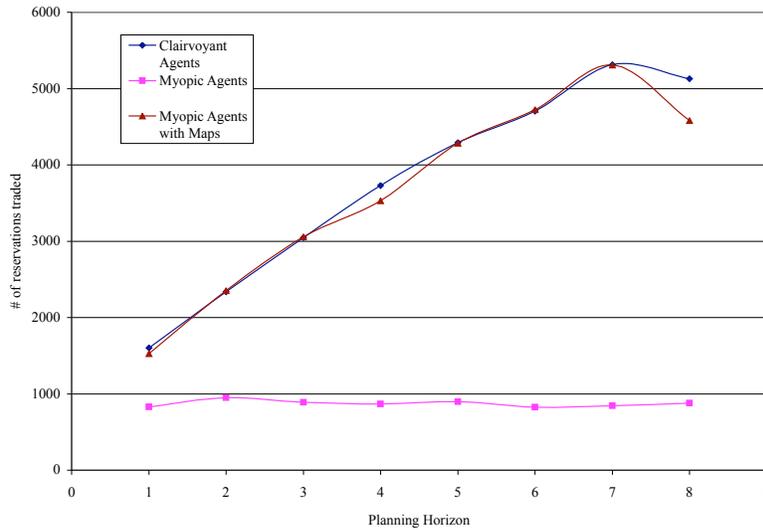


Fig. 6. The total number of trades with clairvoyant agents and 40% reservable capacity.

the types had an aversion to (mutually exclusive) subset of the nodes. This model captures the scenario where there are activities targeting different age groups: toddlers can't go on roller coasters, and teenagers abhor kiddy rides. Naturally, agents with an aversion to a subset of nodes spent all of their time in the other nodes. More interestingly, agents with no aversion spent considerable time in the subsets of nodes that were disliked by one of the other groups because the reduced competition for those activities led to shorter queues.

Another experiment involved adding precedence constraints which require that an agent visit one node before another. The constraints are loose requirements in that they do not require that the successor node be visited immediately after the predecessor. In other words, if the constraints requires that you visit A before B, visiting A essentially earns you a coupon to visit B at a later time. Figure 7 shows the transitions from one node in the graph to another when the constraint requires that node 5 come before 4, which must come before 1. Interestingly, although the constraint is loose, the most frequent transitions are from 5 to 4 (twelve times) and 4 to 1 (sixteen times). This result is due to the fact that the predecessor node serves as a funnel that limits the number of agents competing for the successor nodes, thus keeping the queue sizes down. Thus, when an agent completes node 5, the most attractive node to visit is 4 because it has no queue.

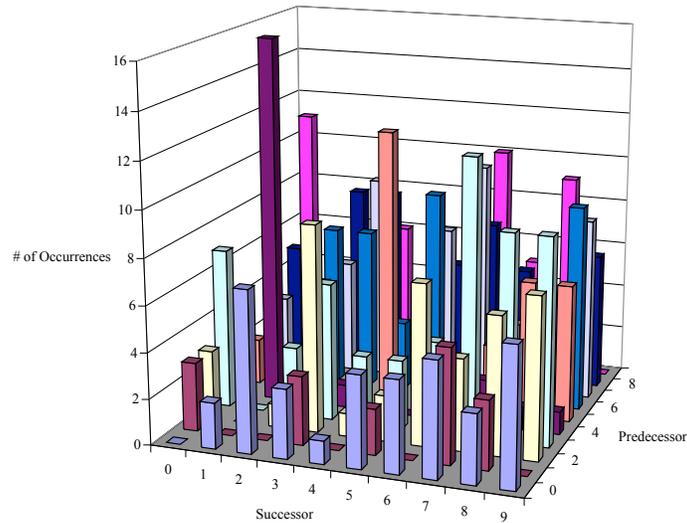


Fig. 7. The number of transitions between nodes when constraint $5 \rightarrow 4 \rightarrow 1$ is enforced.

5. RELATED WORK AND FUTURE DIRECTIONS

For over a decade, researchers have studied software agents in electronic markets [Chavez and 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 workgroups coordinate their activities, but to our knowledge the project does not involve market interactions. The Supply Chain Management version of the Trading Agent Competition [Sadeh et al. 2003] is a framework for studying trading strategies in a complex marketplace with temporal scheduling. However, the scheduling aspects of TAC SCM, to date, have been less of an issue than the strategic interactions between agents. Some work has been done on market-based scheduling [Clearwater 1995; Wellman et al. 2001] but the constraints on the models differ in significant ways from the model presented here.

We plan to continue to extend the model. In particular, we are interested in exploring better market mechanisms (eg., combinatorial exchanges) and 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.

6. ACKNOWLEDGEMENTS

This paper is based on the first author's Masters Thesis. We wish to thank Munindar Singh, and the members of the Intelligent Commerce Research Group at NCSU 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*. Seattle, WA, 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*. London.
- CLEARWATER, S. 1995. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific.
- GREENWALD, A. AND STONE, P. 2001. Autonomous bidding agents in the Trading Agent Competition. *IEEE Internet Computing*, 52–60.
- KURUMATANI, K. 2002. User intention market for multi-agent navigation—an artificial intelligent problem in engineering and economic context. In *AAAI Workshop on Multi-Agent Modeling and Simulation of Economic Systems*. Edmonton, CA, 1–4.
- MAES, P. 1994. Agents that reduce work and information overload. *Communications of the ACM* 37, 31–40.
- MAS-COLELL, A., WHINSTON, M. D., AND GREEN, J. R. 1995. *Microeconomic Theory*. Oxford University Press, Oxford.
- PRADO, J. E. AND WURMAN, P. R. 2002. Non-cooperative planning in multi-agent, resource-constrained environments with markets for reservations. In *AAAI 2002 Workshop on Planning with and for Multiagent Systems*. 60–66.
- RICH, C. AND SIDNER, C. 1997. Collagen: When agents collaborate with people. In *Proceedings of the First International Conference on Autonomous Agents*. 284–291.
- SADEH, N., ARUNACHALAM, R., ERIKSSON, J., FINNE, N., AND JANSON, S. 2003. TAC-03: A supply-chain trading competition. *AI Magazine* 24, 1, 92–94.
- 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. 1993. A market-oriented programming environment and its application to distributed multi-commodity flow problems. *Journal of Artificial Intelligence Research* 1, 1–23.
- WELLMAN, M. P., WALSH, W. E., WURMAN, P. R., AND MACKIE-MASON, J. K. 2001. Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35(1-2), 271–303.
- 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.