

A Multi-Agent System for Meting Out Influence in an Intelligent Environment

M. V. Nagendra Prasad & Joseph F. McCarthy

Center for Strategic Technology Research

Accenture

3773 Willow Road

Northbrook, IL 60062 USA

+1 847 714 {2062,2260}

nagendrap@yahoo.com, mccarthy@cstar.accenture.com

ABSTRACT

Intelligent environments are physical spaces that can sense and respond to the people and events taking place within them, providing opportunities for people to influence environmental factors that affect them, such as the lighting, temperature, décor or background music in the common areas of an office building. The designer of an environment that can be influenced by a group of collocated people rather than a single individual must decide how to accord influence among the individuals in the group. We have designed two multi-agent group preference arbitration schemes and tested them out in an intelligent environment, MUSICFX, which controls the selection of music played in a fitness center. One scheme seeks to maximize the average satisfaction of the inhabitants, the other seeks to maximize the equitable distribution of satisfaction among the inhabitants. We present the results of a series of experiments using real data collected from the deployed system, and discuss the ramifications of these two potentially conflicting goals.

Keywords

Multi-agent systems, intelligent environments, ubiquitous computing, agents¹.

INTRODUCTION

An intelligent environment can detect the people within its space and can then adapt itself to those people. Most of the research into intelligent environments and other applications of ubiquitous computing has focused on how an environment can sense and respond to a single individual [1, 2, 4, 5, 6, 7, 8, 12]. Our research, in contrast, explores the issue of how an environment can effectively adapt to a *group* of people, even when these people have a diverse set of preferences. We are interested in exploring how *background* environmental factors might be better adapted to the preferences, rather than direct commands, of

inhabitants. In this regard, our research is more akin to the ideas of calm technology [13] and ambient media [3, 14].

In this paper, we will describe the design of a multi-agent system that adapts to changes in the environment and the needs of its inhabitants. The system is set up as an artificial economy of agents serving as proxies for actual inhabitants. Within this dynamic market economy, different control strategies – implemented as simple market rules – can produce different effects on the inhabitants. The specific issue we address in this paper is the potential conflict between a strategy to maximize average inhabitant satisfaction and a strategy to maximize the equitable distribution of satisfaction among inhabitants. The former goal may lead to a “tyranny of the majority” wherein a small number of inhabitants, with preferences that vary significantly from the norm, never achieves satisfaction. The latter goal may lead to instances where the preferences of a small minority override the preferences of a large majority. One might characterize these goals as trying to please some of the people all of the time versus trying to please all of the people some of the time.

As an example, suppose an intelligent meeting room adjusts its temperature in accordance to the thermal preferences of its inhabitants. Most people prefer something close to “room temperature” (68°F/20°C), but invariably, some like it hot and some like it cold. A room that wants to maximize average satisfaction would set its temperature to the mean of the thermal preferences of its inhabitants; this would mean the outliers who like it hotter or colder would never be maximally comfortable. A room that seeks to achieve the most equitable distribution of satisfaction might often set its temperature to the average preference, but it would also occasionally set its temperature higher and occasionally set its temperature lower, so that the people on the fringes would have higher overall satisfaction over time.

Another example involves background music playing in a fitness center. An intelligent fitness center environment that seeks to maximize average satisfaction might play only the music that is most popular among the majority of its

¹ Copyright ? 1999, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

members. An environment that is more concerned with an equitable distribution of satisfaction might occasionally play music that is not the most popular with the most people, but is popular among some of members with more eclectic musical tastes.

MUSICFX [9] is a realization of such a system, having been installed in a corporate fitness center – the Fitness Xchange (FX) at Accenture Technology Park (ATP) – where it acts as an automatic disc jockey, deciding what kind of music to play for the fitness center members working out at any given time. The system has been very popular, with over 70% of members reporting they like the MUSICFX-controlled music selection better than the previous human-controlled music selection. Although much of the discussion in this paper focuses on the MUSICFX system, we believe that the issues we are addressing are of a much more general nature, and need to be considered by the designers of any intelligent environment. MUSICFX, as a deployed system, can be regarded as an interim testbed, which provides us considerable data collected from an instantiated intelligent environment.

We begin by providing an overview of the MUSICFX system and the environment in which it operates. We then present a hypothetical scenario that highlights the tradeoffs between popularity and fairness. The next section defines our general framework for multi-agent group arbitration systems and describes two specific group preference arbitration schemes, MAX-SAT and EQUITABLE, that seek to maximize average satisfaction and maximize equitable distribution of satisfaction, respectively. We report on a series of experiments using these two schemes, and we conclude with a discussion of the ramifications this research has for the design of intelligent environments that take inhabitant preferences into account.

SYSTEM OVERVIEW

Any intelligent environment that adapts to the preferences of its inhabitants needs three main components: a mechanism for detecting inhabitants and their activities, a representation of inhabitant preferences, and an algorithm for deciding how to adapt based on those preferences.

MUSICFX detects inhabitants by requiring members to login, using a proximity badge reader and standard-issue ATP badges, as they enter the fitness center. Rather than requiring a member to explicitly logout – for which there exists no significant incentive – we set an expiration timeout of 90 minutes after each login, after which time the system presumes the member has left.² MUSICFX assumes its inhabitants’ activities can be broadly classified as exercising, and doesn’t require finer distinctions.

² For convenience, we will refer to the virtual “logout” events that occur 90 minutes after entering a fitness center as members “leaving” or “exiting” the center.

The MUSICFX preference database represents members’ ratings of each of 91 genres of music, each available on a separate station from a satellite music service. Each genre is rated on a 5-point scale, from +2 through -2, interpreted as “I {love, like, don’t care, dislike or hate} this kind of music.” The initial set of preferences is submitted to the system remotely via an electronic enrollment form; members can update these preferences in the fitness center at any time.

When a member logs in to the system, that person’s preferences are retrieved and added to the current pool of preferences. The MUSICFX Group Preference Arbitrator sorts the list of genres from most popular to least popular, and then uses a weighted random selection algorithm³ to select one of the most popular genres to play. The Arbitrator is invoked each time a person enters or leaves the fitness center, each time a person updates his or her preferences, each time a fitness center staff member adjusts a system parameter, or after a maximum play time for a single genre has been exceeded.

A more comprehensive description of the system can be found in McCarthy & Anagnost [9].

GROUP PREFERENCES

A sample set of music preferences, for five people (Al, Barb, Carl, Deb and Ed) and ten stations, is shown in Figure 1.

<i>i</i>	<i>Genre</i>	<i>Person</i>	A	B	C	D	E	GP_i	Pr_i
1	Alternative Rock		2	2	2	-1	-1	50	0.42
2	Hottest Hits		2	1	1	0	-2	38	0.32
3	New Music		1	1	1	-2	0	31	0.26
4	Dance		0	0	-1	2	-1	26	0.00
5	Hot Country		0	0	-2	-1	2	25	0.00
6	World Beat		0	1	-1	-1	-2	15	0.00
7	Traditional Country		-1	0	-1	1	-2	15	0.00
8	50's Oldies		0	0	-1	-1	-1	11	0.00
9	Heavy Metal		-1	-1	-1	-1	-2	4	0.00
10	Polka		-1	-1	-2	-2	-2	2	0.00

If we further simplify the scenario by supposing that these five people work out together all the time, then we can see that an algorithm seeking to maximize average satisfaction will always choose the top-rated station (based on simply summing the individual preferences), “Alternative Rock,” even though two inhabitants (Deb and Ed) dislike this station.

³ Rather than always selecting the most popular station, which could result in a tyranny of the majority, an element of randomness – with probabilities distributed according to popularity – was introduced in the original algorithm in order to inject a degree of equitability.

Choosing the second most popular station would be an improvement for Deb, who at least doesn't mind "Hottest Hits", but Ed hates that station. Selecting the third most popular station, "New Music," is less distasteful to Ed, but anathema to Deb. In order to please Deb or Ed, an algorithm would need to play the fourth or fifth most popular stations, though these selections would not be popular among the other inhabitants.

This scenario highlights the tension between trying to achieve maximum average satisfaction and trying to achieve equitability for all inhabitants. We will return to this scenario to illustrate the behavior of the two schemes we have created in order to investigate these potentially conflicting goals.

A MULTI-AGENT GROUP ARBITRATION SYSTEM

A *multi-agent group arbitration system* consists of a set of agents (A), where each *agent* (a_i) represents a single person's preferences.⁴ A market-based economy governs the arbitration process for selecting among several available *options* regarding environmental factors. The arbitration process consists of a *bid-select-redistribute* cycle, as shown in Figure 2.

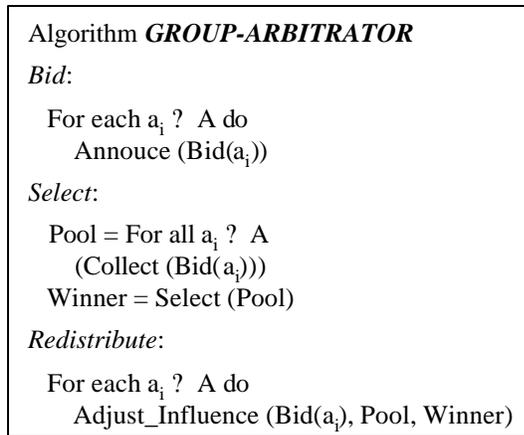


Figure 2: Bid-Select-Redistribute Cycle

In the *bid* stage, each of the agents announces its bid for the different options available. The *select* stage involves the arbitration process where the bids by all agents involved are pooled and a "winning" option is chosen. The last stage is the *redistribute* stage where the future potential for an agent's capability to influence a choice may be readjusted based on the present option selected. In the case of MUSICFX, an agent represents a person's preferences for different genres of music. A bid by an agent involves an announcement of an agent's strength of desire for (or aversion to) each genre. The group arbitration

algorithm chooses a particular genre to be played based on the combined pool of bids by all the agents presently working out in the fitness center.

We can devise various schemes by instantiating the generic *Group Arbitrator* algorithm in different ways. Below we provide the details of two schemes we devised and studied using MUSICFX as a test-bed: MAX-SAT and EQUITABLE. MAX-SAT seeks to maximize the average satisfaction of all inhabitants. The EQUITABLE algorithm, in contrast, seeks to maximize equitability of satisfaction among all inhabitants.

Scheme MAX-SAT

The MAX-SAT scheme is based on the original algorithm used in MUSICFX (without the weighted random selection operator); it is designed to select the most popular station during each cycle. The popularity of a station is defined as a function of the individual preferences of all present inhabitants, but no history of past selections is maintained.

Within the framework of *Group Arbitrator*, we can specify the *bid* and *select* functions for MAX-SAT as follows:

Bid

The *Bid* function looks at the integer-valued individual preferences ($IP_{i,j}$) of agent a_i ranging from -2 to $+2$ for each of the M options that are rated, normalizes those preferences to non-negative integers, and squares the result to broaden the gap between stations at different levels of preference, e.g., those that are loved and those that are merely liked. This results in an M -component vector :

$$Bid(a_i) = \{ b_{ij} = (IP_{ij} + 2)^2 \mid j = 1 \dots M \}$$

Select

The *Select* function takes as input a *Pool*, which is an $N \times M$ matrix where M is the number of categories being rated (musical genres) and N is the number of inhabitants (FX members who are currently working out). For each category j , and each agent i , that agent's individual preference for that category ($IP_{i,j}$) is used by the algorithm to compute the overall group preference for that category (GP_j) using the following summation formula:

$$GP_j = \sum_{i=1}^N b_{ij}$$

Select then chooses the winning option (w) that maximizes group preference (GP_j):

$$w = \arg \max_j \{ GP_j \mid j = 1 \dots M \}$$

Redistribute

MAX-SAT has no redistribution function, since it is explicitly not concerned with achieving equitability.

Scheme EQUITABLE

The EQUITABLE scheme takes an egalitarian approach to the choice of an option, based on *state* information stored

⁴ We will use the terms "agent" and "person" (or "inhabitant") interchangeably in our presentation of these algorithms.

in the agents. The state information is stored in the form of *cash*, which represents a coarse representation of the history of an agent. Each agent starts life with the same amount of cash. During each cycle, an agent's bid for a particular option is proportional to its preference for that option and the cash it has. Once an option is selected, every agent pays or receives an amount proportional to its preference for or against that option – agents that have an unfavorable option imposed on them receive payment for the inconvenience they suffer from those agents who prefer the option selected. Thus any agent that is subjected to low preference options for a long time will accumulate enough cash to dominate the bidding process at some point and thereby have one of its preferred options selected.

Instantiating this strategy within a *Group Arbitrator* framework, we have:

Bid

The *Bid* function looks at the integer-valued preferences of agent a_i ranging from -2 to $+2$, normalizes each preference rating to a non-negative integer, and multiplies this value by the amount of cash possessed by the agent. A bid factor (*bf*), a scaling constant between 0 and 1, is used to modulate the amount of cash tendered during any given cycle. This results in an M -component vector:

$$Bid(a_i) = \{ b_{ij} = bf \cdot (IP_{ij} + 2) \cdot Cash(a_i) \mid j = 1 \dots M \}$$

Select

The *Select* function for EQUITABLE is the same as for MAX-SAT.

Redistribute

The *Adjust_Influence* function determines how to reallocate wealth among the agents representing the current inhabitants. The amount of compensation (C_i) paid to agents who endure unfavorable options is proportional to the difference between the maximum preference – in this case, 2 – and the agent's individual preference for the winning option (w), multiplied by a compensation deceleration function (*cdf*) of the agent's cash. This amount is adjusted by a compensation factor (*cf*), another scaling constant between 0 and 1:

$$C_i = cf \cdot (Max-Pref - IP_{i,w}) \cdot cdf(Cash(a_i))$$

The compensation deceleration function is defined as

$$cdf(\text{amount}) = 1 / (1 + e^{2 \cdot (\text{amount} - \text{initial-cash}) / \text{initial-cash}})$$

where *initial-cash* is a constant for all agents.

Designing general market-based multi-agent schemes involves parameterizing the system along a number of dimensions. Setting the parameters' forms and values involves a good understanding of the dynamics of the environment being modeled. The specific definition of the *cdf* function is a good example of such a parameter. Inhabitants who dislike or hate the vast majority of options are very likely to garner a huge amount of the net wealth in the system through repeated compensations. This may lead

to a net drain of wealth from the rest of the system, which, in turn, may cause imbalances in the market when it starts functioning in regions of the parameter space characterized by "extreme" cash values. The *cdf* function has been designed to reduce the compensation received by agents with very high amounts of cash.

Compensatory updating of the agent cash is done as follows:

$$Cash(a_i) = Cash(a_i) + C_i$$

In order to determine how much each agent pays, we first compute the total compensation, C , as the sum of all individual agent compensation needs (C_i). For each agent a_i , its payment P_i is a fraction of the total compensation, based on a payment factor (*pf*), a scaling constant between 0 and 1, a payment deceleration function (*pdf*) of the agent's cash, and the agent's share of the overall group preference for the winning option selected:

$$P_i = C \cdot pf \cdot pdf(Cash(a_i)) \cdot (IP_{i,w}) / Total_Pref$$

Total_Pref, the normalized total preference for N agents, can be defined as follows:

$$Total_Pref = \sum_{i=1}^N (IP_{i,w})$$

The payment deceleration function (*pdf*) is defined as

$$pdf(\text{amount}) = 1 / (1 + e^{2 \cdot (\text{initial-cash} - \text{amount}) / \text{initial-cash}})$$

Just as the *cdf* function was designed to reduce compensation received by extremely wealthy agents, the *pdf* function was designed to reduce the compensation paid out by very poor agents.

EXPERIMENTS

Our goal is to create a multi-agent framework in which we can explore different schemes for allocating influence in an intelligent environment. In particular, we wanted to explore the tradeoff between popularity and fairness in such a setting. To this end, we ran simulations of the two algorithms described in the previous section.

The deployed MUSICFX system has extensive logs that track events that take place in the environment. In particular, we know when each person has entered and left the fitness center, what each person's preferences were at any given time, and which station has been selected at any given time. This event data can be used in our simulator, however, we wanted to first run the algorithms on smaller sets of data so that we could better understand their behavior under more tractable conditions.

We will present the behavior of our schemes using three sets of data: one corresponding to the simplified data shown in Figure 1, another using randomly generated data with a random selection of inhabitants, and a final experiment using real data from the event log from MUSICFX over a one month period.

In each experiment, we measure each inhabitant's satisfaction with the option selected. For example, using the data in Figure 1, if Deb spent four time units listening to Alternative Rock, two time units listening to Dance, and one time unit listening to Hot Country, her individual satisfaction would be $(4 \cdot -1) + (2 \cdot 2) + (1 \cdot -1) = -1$.

Two metrics serve as a basis for comparing the performance of the algorithms over entire populations. The first measures the total of individual satisfaction levels obtained by a scheme; the second measures the equitability of individual satisfaction levels obtained by a scheme.

Total Satisfaction. This is the sum of individual satisfaction levels for all of the inhabitants in the population.

Gini Coefficient. The Gini Coefficient is a measure of how much a given distribution of wealth (satisfaction, in this case) departs from the ideal egalitarian distribution [10]. This statistic can be explained with reference to the so-called Lorenz curve. In plotting a Lorenz curve, measures of individual wealth, or in this case, satisfaction, are sorted in an increasing order and then cumulative measures are derived. The X-axis represents the percentage of the population; the Y-axis represents the percentage of cumulative wealth. If wealth is distributed completely equitably, then the top 10% of the population owns 10% of the wealth, the top 20% owns 20% of the wealth, and so on. The ideal curve thus has a 45° slope; any deviation from this ideal curve represents a measure of the inequity across the population. The Gini Coefficient measures the difference between the Lorenz curve for a given distribution of wealth and the ideal curve representing an egalitarian distribution; thus lower values represent more egalitarian distributions than higher values. The Gini Coefficient can be defined by the following formula:

$$1 - \frac{1}{n} \sum_{i=1}^n \frac{z_i}{Z} = \frac{2}{n^2 Z} \sum_{i=1}^n z_i^2$$

where z_1, \dots, z_n represent individual levels of wealth in decreasing order of size, Z is the total income, and n is the number of individuals. For our experiments, wealth is defined as satisfaction with selected options.

Experiment Set I

In the first set of experiments, we used the hypothetical data shown in Figure 1. We wanted to see whether there was any measurable difference between the behavior of the schemes when they were run with data that was constructed specifically to highlight the potential conflict between the goals of popularity and fairness. We ran each algorithm for 500 time units, tracking individual satisfaction levels.

Figure 3 shows the Lorenz curve for the performance of the two schemes on this data set, and Table 1 gives the statistics. As predicted, EQUITABLE trades off lower total satisfaction (by 17.5%) for greater equitability (by 94.2%).

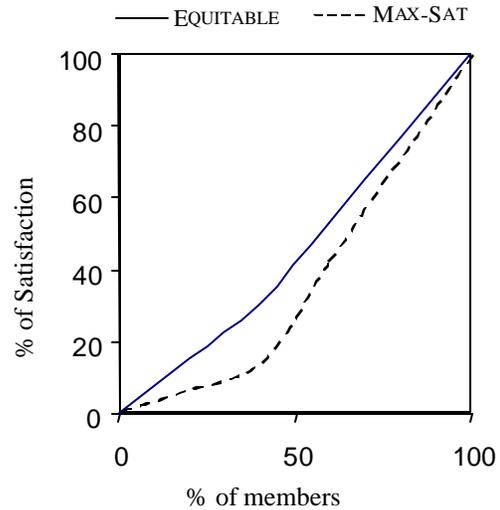


Figure 3: Lorenz curve for Experiment I

Total Satisfaction			Gini Coefficient		
Max-Sat	Equitable	% change	Max-Sat	Equitable	% change
7000	5776	-17.50%	0.26	0.01	94.20%

Table 1: Results of Experiment I

In fact, EQUITABLE nearly achieves a perfectly egalitarian distribution of satisfaction (a Gini Coefficient of 0.0 represents the ideal value).

Experiment Set II

Having convinced ourselves that the EQUITABLE scheme achieves a much more equitable distribution of satisfaction than MAX-SAT when presented with inhabitant preference data that was specially contrived to achieve this result, we next set about experimenting with randomly generated preference data. For the second set of experiments, we created a population of 15 people, each of whom were assigned randomly generated preference ratings for each of five stations. For each epoch of the experiment, we randomly selected 10 people from this population, provided them each with the same initial allocation of cash, and ran MAX-SAT for 500 time units; we then reinitialized the cash allocation for each person and ran EQUITABLE for 500 time units. We ran 10 epochs of the experiment, with the results shown in Table 2.

Epoch	Total Satisfaction			Gini Coefficient		
	Max-Sat	Equitable	% change	Max-Sat	Equitable	% change
1	14000	12521	-10.6%	0.19	0.10	50.5%
2	15500	13624	-12.1%	0.19	0.11	39.9%
3	14500	13162	-9.2%	0.22	0.13	41.8%
4	14000	12713	-9.2%	0.23	0.15	34.7%
5	12000	10825	-9.8%	0.30	0.10	68.3%
6	14500	11243	-22.5%	0.19	0.08	56.0%
7	15000	12830	-14.5%	0.16	0.15	5.4%
8	15000	12564	-16.2%	0.19	0.15	20.5%
9	14500	11939	-17.7%	0.19	0.10	49.6%
10	13000	11851	-8.8%	0.28	0.12	59.6%
Average	14200	12327	-13.1%	0.21	0.12	42.6%

Table 2: Results of Experiment II

Once again, the EQUITABLE scheme results in a lower total satisfaction than MAX-SAT (13.1% less), but achieves greater equitability of satisfaction (42.6% more). It is interesting to note that the differences in this data set are less dramatic than they are for the specially contrived data used in the first experiment.

Experiment Set III

In the final experiment, we used real event log data from the deployed MUSICFX system to test the two schemes. For each arrival event in a segment of the event log covering a period of one month, a {person id, arrival time} pair was extracted, yielding a data set that includes various groupings of 166 fitness center members. The results of this simulation are shown in Table 3.

Total Satisfaction			Gini Coefficient		
Max-Sat	Equitable	% change	Max-Sat	Equitable	% change
6767	5953	-12.0%	0.14	0.09	31.9%

Table 3: Results of Experiment III

As in the previous experiments, EQUITABLE sacrifices total satisfaction (12.0% less than MAX-SAT) for increased equitability of satisfaction (31.9% more).

Discussion

In all three experiments, EQUITABLE led to a more egalitarian distribution of satisfaction, at the cost of a lower total satisfaction among inhabitants. Although the results were most dramatic in the first experiment, using data specially constructed to highlight the conflict between popularity and fairness, there were considerable differences seen in results from the real data set.

RELATED WORK

The research described in this paper deals with the design of a multi-agent system for group preference arbitration schemes. Other researchers [8, 11] have explored applications of market-based multi-agent systems.

Huberman and Clearwater [8] created a market-based system in which a set of agents, each representing the temperature controller of an individual office within a building, bid to buy or sell thermal units. While the goal of their system –

maximizing comfort – is similar to the goal of the MAX-SAT algorithm, their agents do not retain money between auctions (which are held every minute), and thus the system does not have the capability to maximize the equitability of comfort distribution over time. Another difference is that Huberman and Clearwater’s agents are bidding over a variable amount (units) of a fixed resource (hot or cold air), whereas our agents are bidding over a variable amount (time) of a variable resource (91 options simultaneously available). It is interesting to note that Huberman and Clearwater also contend with the issue of people exhibiting extreme preferences, which in their case corresponds to extreme thermostat settings in individuals’ offices.

Walsh and Wellman [11] present a decentralized protocol for allocating tasks among agents that contend for scarce resources. The framework described in their work is potentially applicable to a large class of multi-agent problems, including the search for information in a digital library. However, their framework does not appear to be well suited to the problem of allocating a resource (or good) that is inherently shared, such as the music played in a fitness center.

CONCLUSION

In this paper, we present a Multi-Agent Group Preference Arbitration system. At the core of this framework is a multi-agent system based on market mechanisms for resolving multiple, conflicting preferences among a group of people inhabiting a shared environment. Based on our experiences with MUSICFX, a deployed system for group preference arbitration used in the selection of music in a fitness center, we designed two distinct schemes: MAX-SAT and EQUITABLE. These schemes were tested both on artificial data and real data derived from the event logs maintained by the MUSICFX system.

The results of our experiments have provided strong empirical evidence demonstrating that we can affect the tradeoff between popularity and fairness. However, determining which of these goals should be emphasized within a given environment is a difficult policy question, the answer to which is beyond the scope of the work reported here.

One of the shortcomings of the current EQUITABLE scheme is that it does not take into account a person’s overall preference distribution. Intuitively, someone who hates nearly everything should not be paid as much for his or her inconvenience as someone who has mostly positive preferences, since such a person is difficult to please anyhow. Likewise, someone who hates nearly everything should have to pay more for his or her preferred option(s), since there are so few alternatives available to that person. Future versions of the algorithm will investigate ways to take these factors into consideration.

Another future direction involves exploring how group preference arbitration can affect other factors in a shared environment. For example, we plan to develop an application that affects visual aspects of an environment in response to the presence of different inhabitants.

One issue that arises in the deployed MUSICFX system is the requirement that a member fill out a questionnaire with 91 questions (corresponding to the different stations). This might be a disincentive for registering with the system. In addition, in a number of environments, explicit questionnaires may not be a feasible way of deriving user preferences. We are beginning to look at techniques from machine learning and collaborative filtering to induce user preferences from observation or sparse data. Another interesting direction of future work involves dealing with more than one shared resource and group arbitration in such situations. The problem here is complicated by the fact that in addition to arbitrating the users sharing a particular resource, we also have to deal with optimally partitioning users across resources.

We believe that as intelligent environments gain increasingly sophisticated ways of sensing and responding to their inhabitants, there are many important issues to explore. MUSICFX provides one environment in which to investigate some of these issues, but we look forward to a future filled with many intelligent environments in which rich and complex intra- and inter-environmental interactions can evolve.

ACKNOWLEDGMENTS

The authors wish to thank Theodore Anagnost, who has been responsible for the implementation and maintenance of the real MUSICFX system, and who has contributed many ideas to the design of both the actual algorithm used in the system and the variations described in this paper. We also wish to thank Andy Fano for helpful comments on an earlier draft of this paper.

REFERENCES

1. Michael H. Coen. 1997. Building Brains for Rooms: Designing Distributed Software Agents. In *Proceedings of the Ninth Conference on Innovative Applications of Artificial Intelligence (IAAI 97)*. 971-977. Menlo Park, CA: AAAI Press..
2. Jeremy R. Cooperstock, Sidney S. Fels, William Buxton, and Kenneth K. Smith. 1997. Reactive Environments: Throwing Away Your Keyboard and Mouse. *Communications of the ACM* 40(9): 65-73.
3. Andrew Dahley, Craig Wisneski and Hiroshi Ishii. Water Lamp and Pinwheels: Ambient Projection of Digital Information into Architectural In *Proceedings of the 1998 ACM Conference on Human Factors in Computing Systems (CHI '98)*, (Los Angeles, April 1998), ACM Press, pp. 269-270.
4. Scott Elrod, Gene Hall, Rick Costanza, Michael Dixon, and Jim Des Rivieres. 1993. Responsive Office Environments. *Communications of the ACM* 36(7): 84-85.
5. Jorg Geibler. Shuffle, Throw or Take It!: Working Efficiently with an Interactive Wall. In *Proceedings of the 1998 ACM Conference on Human Factors in Computing Systems (CHI 98 Summary)*, (Los Angeles, April 1998), ACM Press, pp. 265-266.
6. Andy Harter and Andy Hopper. 1994. A Distributed Location System for the Active Office. *IEEE Network* 8(1): 62-70.
7. Andy Hopper, Andy Harter, and Tom Blackie. 1993. The Active Badge System. In *Proceedings of the Conference on Human Factors in Computing Systems (InterCHI 93)*.
8. Bernardo A. Huberman and Scott H. Clearwater 1995. A Multi-agent System for Controlling Building Environments. In the *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS 95)*.
9. Joseph F. McCarthy and Theodore D. Anagnost. 1998. MUSICFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. In *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work (CSCW '98)*, pp. 363-372.
10. David W. Pearce (ed.), *The MIT Dictionary of Modern Economics*, 1992, MIT Press, Cambridge, MA.
11. William E. Walsh and Michael P. Wellman 1998. A Market Protocol for Decentralized Task Allocation. In the *Proceedings of the Third International Conference on Multi-agent Systems (ICMAS 98)*.
12. Roy Want, Bill N. Schilit, Norman I. Adams, Rich Gold, Karin Petersen, David Goldberg, John R. Ellis, and Mark Weiser. 1995. An Overview of the PARCTAB Ubiquitous Computing Experiment. *IEEE Personal Communications* 2(6): 28-43.
13. Mark Weiser and John Seely Brown. The Coming Age of Calm Technology. In *Beyond Calculation: The Next Fifty Years of Computing*. Springer-Verlag, 1997, pp. 75-85.
14. Craig Wisneski, Hiroshi Ishii, Andrew Dahley, Matt Gorbet, Scott Brave, Brygg Ullmer and Paul Yarin. Ambient Displays: Turning Architectural Space into an Interface between People and Digital Information. In *Proceedings of the International Workshop on Cooperative Buildings (CoBuild '98)*. Darmstadt, Germany, February 1998. Springer Press, pp. 22-32.