

MUSICFX: An Arbiter of Group Preferences

Joseph F. McCarthy

Center for Strategic Technology Research
Accenture
3773 Willow Road
Northbrook, IL 60062 USA

mccarthy@cstar.accenture.com

Abstract

Much of the research into intelligent environments has focused on how an environment might sense and respond to the presence and/or actions of a single individual in that environment. MUSICFX is a system that responds to the presence of a group of people in a shared environment by using a *Group Preference Agent* to select music that will best accommodate the musical preferences of that group.

Introduction

We spend much of our time in shared environments – spaces in which two or more people are in close proximity and are mutually affected by tactile, olfactory, visual, and aural factors such as temperature, scents, lighting, and background noise. As embedded computers permeate our environments, we are seeing new capabilities for sensing people and their activities, and for intelligently responding to these events. This new paradigm results in a shift of perspective: from viewing people as *users* of computers to a view of people as *inhabitants* of environments. One way that an intelligent environment can respond to its inhabitants is by adjusting itself to better suit the inhabitants' needs or preferences.

MUSICFX is a system that adjusts the selection of music playing in a fitness center to best accommodate the musical preferences of the people working out at any given time. The system has a database of fitness center members' preferences for a wide range of musical genres, a mechanism for identifying who is working out at any given time, and an algorithm for selecting a musical genre that will promote members' listening pleasure.

This research grew out of the juncture of two concurrent threads in the author's experience. One is a perception that most of the research in ubiquitous computing has focused on ways that an environment might respond to a single individual rather than a group of people. The other is a frustration with the "lowest common denominator" music typically played in a fitness center on a daily basis.

The group preference agent embodied in MUSICFX is not limited to selecting music in a fitness center, however. It is applicable to any shared environment in which people gather for an extended period of time. For example, a restaurant would provide another type of environment that might benefit from a group preference agent, but an elevator would not be a good candidate. Furthermore, the approach is applicable to environmental factors beyond music or other auditory input: visual, olfactory and tactile factors could also be adapted to the set of people assembled in a particular location. Finally, the set of preferences for each individual might either be explicitly specified by each person, or inferred based on knowledge of the person's behavior, e.g., a shopper's purchase history.

The following sections in this paper will discuss the general concept of a group preference agent, describe the specific social milieu in which MUSICFX operates and the technical architecture of the system, compare and contrast the group preference agent with other research into intelligent environments, and conclude with some other potential applications of this approach.

A Group Preference Agent

A group preference agent is an entity that considers the diverse preferences of a group of people in a shared environment, and affects one or more environmental factors based on that consideration. A disc jockey at a party is one example of a human performing this function; other examples include dimming the lights at the start of a presentation, turning down the heat after observing people fanning themselves, closing a window while following a fume-spewing truck, or turning the music volume way up to irritate one's parents. Note that in these examples, the preferences may either be stated explicitly ("Turn that music down!"), inferred from observation (fanning = too hot), or inferred based on an understanding of typical human preferences (noxious fumes = bad).

There are a number of factors that make the job of a group preference agent difficult:

- How do you know what people want? They might tell you explicitly, provide subtle (and not-so-subtle) cues, or just sit (or perhaps jog) there, silently and passively. In some cases, it might be possible to infer their preferences from knowledge about their past or about human behavior in general.
- How badly do they want it? Do the people who make their preferences known, loud and clear, have stronger preferences than those who do not articulate their preferences, or are they merely more willing to share their views?
- How much control do you give to inhabitants? Do you permit people to exercise direct control over environmental factors or merely seek their input?
- How often should the environment adjust itself? While we want our environments to be responsive, we don't want them to "thrash" – spending more time adjusting to inhabitants' preferences than in any stable state. Most people prefer a stable temperature and lighting level, even if these are sub-optimal, rather than frequent changes in these factors.

In MUSICFX, we know what people want because they have told us explicitly what they like, or dislike, among a wide range of musical genres. We also know the strength of their preferences by using a rating scale that differentiates genres they love from those they merely like, and genres they hate from those they merely dislike.

Permitting people to directly control the selection of music in the fitness center would likely result in a free-for-all – a Darwinian environment in which the strongest, or most persistent, get to listen to their music (although we haven't tested this hypothesis). Instead of direct control, the group preference agent seeks input from inhabitants – their preferences for different genres – and then arbitrates among these preferences in an equitable way. However, we do permit the fitness center staff to exert direct control over the system and also to exert indirect control by adjusting certain parameters.

One thing we want to avoid is excessive "surfing" – changing the music selection so often that no one ever gets to hear a complete song. Balanced against this, however, is a desire to avoid playing music that reflects the preferences of a group of people who have already left the fitness center – music which may or may not reflect the preferences of the current group of people working out – i.e., we don't want to wait too long to change the music selection. Since we don't know how long is "just right," we have allowed the fitness center staff to adjust this timeout period based on their experience with the system.

In order to provide a better context within which to understand how MUSICFX works, the following section will describe the environment in which the system operates.

The Fitness Center Environment

The music played in a fitness center provides an ideal laboratory for experimenting with a group preference agent. People are gathered together in a common location, each performing separate (though related) activities that do not typically require much attention. Since the foreground activities are not engrossing, background environmental factors become more important. Most people want to hear *some* music while they are working out, the problem for the staff of a fitness center is to decide which music to play.

The fitness center at the Accenture Technology Park (ATP), called the Fitness Xchange (FX), is located in the lower level of our Northbrook facility, and is open to all ATP residents from 6:00 a.m. through 8:00 p.m. each weekday. The FX has 24 cardiovascular machines (treadmills, stationary bikes, stair climbing machines, and so on), 14 strength-training machines and an assortment of free weight equipment. The number of people working out ranges from a high of around 25 at peak times (before work, lunchtime, and after work) to one or two people during mid-morning and mid-afternoon hours.

ATP subscribes to the PrimeStar™ direct broadcast satellite service, which includes the Digital Music eXpress (DMX ©) music service, providing 91 stations of commercial-free music, each representing a different musical genre. The variety of genres includes Album Rock, Classic Jazz and Symphonic, as well as Flamenco Music, German Oldies and Beach Party. Under the current arrangement, the FX staff may manually select any DMX station, but must walk 150 feet to the room containing the satellite receiver in order to do so.

The FX is the most popular service at ATP, with over 600 members among the Park's 1500 residents. However, despite the general popularity of the FX, the music played in the FX is a source of dissatisfaction, accounting for 25% of the written "suggestions" submitted anonymously by members. There are a small number of members who are quite willing to express their musical preferences (or, more often, dispreferences), and in the absence of input from the less vocal members, there tends to be a "squeaky wheel syndrome" where the complainants get their way. This factor, combined with the inconvenience of physically changing stations, results in three or four stations – the ones that generate the fewest complaints – being played, out of a the set of 91 possible stations.

The primary research goal of MUSICFX is to explore the social ramifications of a group preference agent in this shared environment. Related to this research goal, our more pragmatic goals for the MUSICFX system are to:

- *Democratize the music selection process.* Accommodate the silent majority rather than the vocal minority, resulting in greater listening enjoyment for FX members. The more vocal members may not be pleased as often, but the less vocal members will now have their voices “heard”.
- *Increase the variety of music played.* Provide airtime for the stations outside of the lowest common denominator set. Members who work out at off-peak hours may now be able to listen to music that would never be played when larger numbers are present.
- *Offload the music selection task from the FX staff responsibilities.* Enable the staff to spend more time attending to fitness related matters rather than acting as disc jockeys.

Having described the environment in which the MUSICFX system operates, let us now turn to a description of the system itself.

The MUSICFX System

The MUSICFX system runs on two Windows 95 computers that are linked via the ATP local area network. One computer is located in the room that houses the satellite receiver: it has an infrared (IR) remote control device attached to its serial port. The other computer is in the FX: all the software modules described below reside on this computer. When the system selects a new station, it sends a command from the FX computer to the computer near the satellite receiver, which translates it into a channel change signal sent to the remote control device.

MUSICFX provides two interfaces accessible on the FX computer. The *FX member interface* allows people to login to the system, and to update their preferences for any station. The *FX staff interface* allows the staff to monitor the system, manually select new stations and to adjust certain parameters that will be described below.

Underlying these interfaces, the FX computer hosts three components that enable the operation of the system: a database of FX members’ musical preferences, a mechanism for identifying who is working out in the FX at any given time, and an algorithm for selecting one among the 91 genres music that will promote members’ listening pleasure. The algorithm is invoked whenever an event occurs, e.g., a member enters the FX to begin a workout. Each of these components will be described in greater detail below.

Preference Database

In order for a group preference agent to make an informed decision about how to affect environmental factors, it must know something about what the current inhabitants prefer.

In MUSICFX, each FX member specifies his or her preference for each musical genre. The preference rating for a genre is represented by a number ranging from +2 to -2, where the numbers represent the following levels of preference:

Rating	Interpretation
+2	I love this music
+1	I like this music
0	I don't know / don't care about this music
-1	I dislike this music
-2	I hate this music

A member fills out and submits an electronic enrollment form upon first joining the fitness center; the FX member interface also contains an update screen that permits a member to enter or update his or her musical preferences in the fitness center at any time.

Presence Detection

A group preference agent must know the composition of the group – who are the current inhabitants? – in order to make decisions about how to adjust environmental factors. The first version of MUSICFX detects who is present in the FX by requiring members to login when they enter the fitness center. Members already use a computer in the FX to log their fitness activities at the end of a workout session – this enables FX members to track their progress and FX staff to track equipment usage. The same login ID that is used for the fitness activity tracking system is used by the MUSICFX, and since members are accustomed to using a computer at the end of each session, we do not perceive the extra login process as unduly burdensome.

Group Preference Arbitration Algorithm

We considered a number of potential algorithms for selecting a musical genre based on awareness of the preferences of FX members working out at any given time. For our initial implementation, we decided it would be best to keep the algorithm simple, so that we could implement it quickly and install the system as early as possible. Future versions may incorporate refinements based on user studies.

The group preference arbitration algorithm takes as input an $M \times N$ table of integer-valued preferences ranging from -2 to +2, 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 i , and each member j , that member’s individual preference for that category (IP_{ij}), the algorithm computes the group preference for that category (GP_i) using the following summation formula:

$$GP_i = \sum_{j=1}^N (IP_{i,j} + 2)^2$$

The formula first converts all individual preference ratings to non-negative numbers, so that we can later apply a weighted random selection operator (described below). These values are then squared in order to widen the gap between more popular categories and less popular categories.

Once this group preference value is computed for all categories, the list of values is sorted in descending order, so that the most popular category is first and the least popular is last.

We considered the policy of always selecting the top category, but since most people typically workout at the same time each day, they will tend to hear the same music (unless they update their individual preferences). The second most popular station might be the favorite among a few of those people, but it may never be played. Therefore, the system uses a weighted random selection policy for selecting one of the top m stations (we call this the *candidate set*), where m is a parameter whose value is set by the FX staff.

i	<i>Genre</i> \ <i>Person</i>	A	B	C	D	E	GP_i	Pr_i
1	Alternative Rock	1	2	0	2	2	289	0.48
2	New Music	1	1	1	0	0	169	0.28
3	Hottest Hits	1	1	2	0	-2	144	0.24
4	Hot Country	2	0	0	0	-2	100	0.00
5	Dance	2	-1	1	-1	-2	81	0.00
6	World Beat	0	1	-1	1	-2	81	0.00
7	50's Oldies	0	0	0	-1	-1	64	0.00
8	Traditional Country	1	0	0	-2	-2	49	0.00
9	Heavy Metal	-1	-1	-1	-1	-2	16	0.00
10	Polka	-1	-1	-2	-2	-2	4	0.00

Figure 1: Sample Preferences

Figure 1 depicts a set of sample preferences for five people (A through E) and ten musical genres. The group preference value calculated for each genre (GP_i) is shown in the second column from the right. Assuming that $m=3$, i.e., the candidate set is limited to the three most popular genres, the probability of selecting each genre (Pr_i) is shown in the rightmost column. The sum of the GP values for the three most popular genres is 602, so the probability of selecting the most popular genre, Alternative Rock, is 0.48, twice the probability of selecting the third most popular genre, Hottest Hits.

Events

There are five events that trigger the execution of the group preference arbitration algorithm:

- Member Entrance
- Member Exit
- Individual Preference Update
- System Parameter Adjustment
- Adverse Reaction

As mentioned earlier, in the first version of MUSICFX, members are required to manually login at the computer; this login process triggers the Member Entrance event. We decided not to require explicit logouts to trigger Member Exit events, for fear that people might forget; also, the music plays both in the fitness center and the adjacent locker rooms, and we wanted members' preferences to still be considered while they are in the locker rooms. Therefore, we set a default time – 90 minutes – for the duration of a workout session (plus changing time). Whenever a member updates his or her preferences, we trigger an Individual Preference Update event, since the previous set of preferences is no longer valid. Likewise, when the FX staff changes a system parameter, e.g., narrowing the range of top-rated stations to be considered, the currently playing station may or may not still be a member of the candidate set, so a Parameter Adjustment event is triggered. Lastly, we provide a button on the Member Interface to permit a member to tell the system “I Hate This Song”; when this button is pressed, an Adverse Reaction event is triggered.

System Parameters

The FX staff can constrain the operation of MUSICFX system by adjusting any of the following four parameters:

- Individual Preference Filter
- Group Preference Filter
- Maximum Station Play Time
- Adverse Reaction Threshold

Each of these will be discussed in more detail below.

In addition to controls for the parameters listed above, the interface for FX staff also provides the capability to manually select new stations to play (pressing a button rather than having to walk down the hallway to another room). The FX staff interface also includes a button to turn off the algorithm, requiring manual intervention by the FX staff for all station changes.

Individual Preference Filter

A primary motivation behind MUSICFX is to increase the listening pleasure of the FX members working out at any given time. This can be accomplished by playing more of the music that members want to listen to or by playing less of the music that members do not want to listen to (ideally, both). One way to ensure that the system plays less music that people do not want to listen to is to prohibit the system from playing any station for which anyone present has

specified a low rating. We provide the FX staff a way to specify an individual preference filter threshold between 2 and -2; any station for which any present member has specified a preference value less than that threshold is eliminated from consideration by the algorithm. For example, a threshold setting of -1 would prohibit any station for which anyone currently working out had specified a rating of -2 (i.e., at least one member hates this kind of music).¹

Group Preference Filter

The group preference filter threshold permits the FX staff to determine how many of the top-rated stations – according to the GP formula listed above – will be candidates for the weighted random selection procedure. The possible values range from 1 to 91, where 1 would force the system to select the top choice each time and 91 would permit the system to select any station (though selection of popular stations would be more probable than selection less popular stations). This parameter was included to give the staff some control over the variability of the music being played. Most FX members follow fairly regular workout schedules, e.g., Monday, Wednesday and Friday from 7:00 to 8:00 in the morning; thus, many of the same people work out at the same time each day. If the system always chose the top-rated station, then members would be likely to listen to the same station every time they work out. While we still permit this policy – with the group preference filter threshold set to 1 – we also allow the music selection to vary among the set of most popular stations.

Since the group preference filter threshold may arbitrarily partition a set of stations that share the same group preference rating (equivalent GP values), our sorting routine ensures that stations within an equivalence class are randomly distributed within each segment of the sorted list. Without this, the regularity of member workout times combined with a constant threshold value would result in the same stations being prohibited each day.

Maximum Play Time

Due to the diverse, and sometime conflicting, musical tastes of the FX members working out at any given time, we recognize our inability to please all the people all the time. One situation that we endeavor to avoid is “starvation” – one or more members never hearing music they enjoy due to the differences between their music preferences and those of the majority of people with whom they regularly workout. The weighted random selection operator is one strategy we use to avoid starvation. Another strategy we use is to limit the period of time that any one genre will

play – regardless of how popular it is – before the selection algorithm is invoked.

The period of time that one genre can play without interruption is limited by the maximum play time parameter, which varies between 1 and 60 minutes. We have used an initial setting of 30 minutes, which seems about right. Coupled with our estimate of workout sessions lasting 90 minutes, on average, this results in at least three genre changes – for a total of at least four genres – that will be played for any one member’s workout. A few people have complained that the genres change too frequently, and a few have complained that they don’t change often enough, so we think we’ve arrived at a reasonable setting for this parameter, although the FX staff is free to adjust it at any time.

Adverse Reaction Threshold

There are situations in which the FX members are enjoying – or at least not complaining about – the station currently playing, and then one song is played that is strongly disliked by one or more people. Prior to MUSICFX, these people had to communicate this dislike to the FX staff, and if enough people complained, the FX staff would then walk down the hallway to the satellite receiver room to change the station. The FX member interface, through which people login to the system and/or update their preferences, provides an “I Hate This Song!” button so that people can complain to the system rather than the FX staff. The FX staff interface provides a corresponding control through which the staff can specify a threshold for how many members have to press their “I Hate This Song!” button before the system will respond to this adverse reaction by selecting another station in the candidate set.

Related Research

Other researchers have explored how environments might sense and respond to inhabitants, but most of them have focused on single inhabitants.

The Olivetti Active Badge™ system [Hopper, et al., 1993; Harter and Hopper, 1994] provides a mechanism for locating and tracking individuals throughout a building using infrared badges and a network of transceivers. The system, as originally designed, did not include a representation of preferences, and was primarily focused on how artifacts (computers, doors, or telephones) might respond to an individual rather than a group. This technology could be quite useful in the MUSICFX system, since it would eliminate the need for a manual login or a preset timeout to trigger entrance and exit events – the system could simply poll periodically for the presence of active badges in the fitness center and locker rooms.

¹ In a preliminary survey of 18 people, we found that all but two stations would be prohibited with an individual preference filter threshold of -1, i.e., only two stations were not hated by at least one person.

The Xerox PARCTAB system [Want, et al., 1995] also provides a mechanism for locating and tracking people who are carrying a handheld device. One application allows individual inhabitants of a room to vote on the quality or pace of a presentation using their PARCTAB; the presenter can then respond to this feedback, but the environment itself does not respond to voting. A PARCTAB could be used to control the lighting or temperature of a room – much as we can change channels with a television remote control device – but it does not appear that the designers were concerned with using this functionality in rooms with multiple inhabitants.

The Responsive Environment Project at Xerox [Elrod, et al., 1993] explored how an environment might conserve energy by adjusting the lighting and temperature, based on an awareness of who was present (or scheduled to be present) in offices and common areas within a building. This work differs from MUSICFX in that it was basically a two-state system – an office or common area was either empty or non-empty – and the preferences of the inhabitants was not considered in its control strategy.

The Intelligent Room at MIT [Coen, 1997] is able to track multiple inhabitants in the room, and supports a number of methods for inhabitants to give commands to the room. While it has some capability for noting individual preferences (e.g., not playing Mozart as someone is dozing), it has no explicit mechanism for arbitrating among preferences of a group of people.

The Reactive Room [Cooperstock, 1997] is a shared telepresence environment which responds to its inhabitants, and has a mechanism for storing preferences for videoconference equipment usage. This work focuses on a very difficult problem – how a shared *virtual* environment might better adapt to its inhabitants, when the inhabitants are distributed across multiple physical sites. One way that the room reacts is to adjust a remote camera based on someone leaning left or right; however, it's not clear how the room would react to different people leaning in different directions simultaneously.

Cooperstock, et al. [1997], posit four important factors that affect any intelligent environment: the invisibility of the technology, the capability of manually overriding the system, a mechanism for providing feedback to users, and an ability to adapt to the preferences of users. MUSICFX meets all four of these criteria: other than requiring manual login, the system does not require conscious interaction on the part of the FX members, e.g., requiring people to input their preferences each time (or voting on each song); the FX staff can manually override the station selected by the system at any time; the primary feedback that the system provides is in the selection of music played, but the system also provides some feedback to each member with respect to how well the member's preferences align with those of the current group; finally, the system's primary purpose is

to continually adapt to the preferences of a changing group of people working out in the FX.

Future Work

We are considering extending this work along four different dimensions: new ways of identifying who is present in the FX at any given time, different representations of member preferences, variations on the current group preference arbitration algorithm and new applications of this approach to other environments.

The current method of identifying member entrance and exit events is simple but inaccurate. We have set the expected workout duration time high enough so that nearly all members will have their preferences considered while they are working out. Unfortunately, this means that in most cases, members' preferences will continue to be considered by the system after they have left the fitness center. Ideally, the system could periodically poll to see who is present in the FX or locker rooms, and not require any manual login or default timeout. One way to eliminate the manual login would be to allow members to swipe their ATP badges through a proximity badge reader when they enter the FX – there are already similar badge readers controlling access to the locker rooms. A more comprehensive solution to both the login and logout problem would be to use an active badge system, as described earlier.

Member preferences are represented by numerical ratings of +2 to -2 for each musical genre. This captures two types of information – whether the genre is favored or disfavored by a member, and the strength of the member's preference for that genre – in a single number. Unfortunately, there is no disincentive for a member to be “open-minded” about their preferences – a member may rate a small number of favorite genres as +2 and all the others as -2, giving such a member potentially greater influence over the music selected by the group preference agent. Other, more equitable, schemes may include allocating a maximum number of most loved and most hated stations (to force more ratings in the middle ground), or adding an additional factor to weight each member's preferences according to their distribution of preferences (e.g., members with an overwhelming proportion of -2 ratings would have their preferences discounted). We are also considering market-based schemes for allocating preference resources.

The group preference arbitration algorithm makes its decision based on current information only. We believe that incorporating historical information could improve its decisions. For example, a member with musical preferences that are very different from most of the other members who typically work out at the same time may experience “starvation”, i.e., that member may never listen to the type of music that he or she enjoys. If the system can detect such situations, it would be able to distribute its

selection of music more equitably. Another way that the history mechanism might be used is to break ties among stations in the same equivalence class by favoring infrequently played stations over frequently played stations that are at the same level of group preference. The first version of the system tracks the information that could be used for these algorithm enhancements, but the information has not yet been used.

Fitness centers are not the only environments where adaptation to musical preferences might be beneficial. Any environment in which groups of people are gathered for significant periods of time – say, more than 15 minutes – and in which it would be preferable to listen to or watch something rather than nothing is a candidate for this approach. For example, restaurants might more effectively cater to their customers by playing music that the customers really want to listen to, rather than music that the staff thinks the customers want to listen to (or music that the restaurant staff wants to listen to). Perhaps the issuance of frequent diners cards could include musical preferences as well as culinary preferences and other factors that might help make the restaurants atmosphere and service more personalized.

Music is not the only environmental factor that could adapt to a group of inhabitants. Visual displays could adapt to a group of shoppers in a store, or perhaps in a region of a store, promoting items that are likely to be of interest to the current group. User preferences in this case might be inferred from the purchase history of the shoppers rather than explicitly requested, since shopping goals vary more frequently than music preferences. In fact, shoppers might be encouraged to identify themselves to such a system – using smart cards or some special courtesy card – by the prospect of a MUSICFX system that will play the music that they want to hear while shopping.

Acknowledgements

The author wishes to thank Theodore Anagnost and Chirag Patel for contributing ideas to MUSICFX and for writing most of the code to implement the system., Kelly Dempski and Eric Meidel for their assistance with the networking and hardware for the system, Madeleine Milan, Sharon Kordik and Fran Scott for their help in pushing this project out of the research lab and into the real world, and Andy Fano, Naghi Prasad and Kishore Swaminathan for helpful comments on an earlier draft of this paper.

References

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.

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.

Scott Elrod, Gene Hall, Rick Costanza, Michael Dixon, and Jim Des Rivieres. 1993. Responsive Office Environments. *Communications of the ACM* 36(7):84-85.

Andy Harter and Andy Hopper. 1994. A Distributed Location System for the Active Office. *IEEE Network* 8(1):62-70.

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).

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.