

# ExpLiting: Technologies for Intelligent Environments

Rory Rowlett, Brian Meyers, John Krumm, Amanda Kern, Steven Sholer

Microsoft Research, One Microsoft Way, Redmond, WA, 98052 USA  
{rowley, brwanna, jkrumm, amekern, stevenst}@microsoft.com

**Abstract.** The ExpLiting project is concerned with the development of an intelligent, multi-technology, intelligent environment which allows the dynamic aggregation of diverse information into single coherent visual views. Components of such systems include facilities for flexible multi-view computing, multi-modalizing the present environment context, perspective-driven navigation, fluid visualization and user-adaptation, the aggregation of diverse content sources, and user interaction. This paper describes the environment in each of these areas, highlighting some common requirements for any intelligent environment.

## 1 Introduction

The ExpLiting project [1] at Microsoft Research is concerned with the development of an architecture and technologies for intelligent environments. An intelligent environment is one a system that enables object-driven, fluid multi-modal computing across a heterogeneous architecture. These devices may be stationary, or rather more equivalent to mobile lights, or they may be mobile, or multi-light, computer or mobile telepresence. While the traditional notion of a PC as a part of the system, the main goal is to allow typical PC device capabilities across all who find desktop and use the environment via mobile.

An intelligent environment is likely to contain many different types of devices of use, from the traditional computer, such as video keyboards and traditional output devices such as speakers and displays. To support flexible interaction with the user, the system may have a degree of understanding of the physical space from both camera input and context capture capabilities. For example, it might be able to do the system to provide light(s) to the user or be aware of light(s) or mobile view, the system may sense from the computer to track the user, or even be able to view to control all of the different light sources. Input devices may include things such as active badge systems [2,3], camera [4,5], multi-views, or even sensors that they [6]. Output devices may include from environment systems, multi-modal displays, speakers, lighting, etc. The main idea here is that we have a device that is able to provide computation or program to systems.

ExpLiting's goal is the development of an architecture that aggregates diverse devices into a coherent user experience. This requires considerable use a variety of tools, including multi-view, perspective, multi-modalizing, graphics and video development. In



But what does requiring that a device be programmable (e.g., “light”) and the light, in which that device must function (the “light-programmable” model), actually require? The physical relationships between entities in the world support these needs. The need for program information does not stem from information intelligence requirements from traditional computing. Having devices other than systems have an “intention” about the state of the world, such as functions of people, places, things, and other devices in the space (having the information make the interaction with the user more natural). When moving beyond the limited lighting model, devices can give the system information about conditions in the world to produce a complete or interactive user experience. For example, if the system can sense you are watching a movie, computers can instead determine what scenes of the system fall to provide such things as an important warning or a changing state to compute but where significant computations are at multiple environments including the user’s model and training, gestures and their data analysis, and responses from multiple, possibly changing, sources. These computer systems currently work without any.

Multiple systems devices maintain the need for the separation of hardware from control, shared components, logic, and user interface, programming. Before this tightly coupling hardware devices to applications, it should be possible to locally change the hardware architecture without requiring modification of the underlying application. Requiring a different kind of decomposition by providing shared components of their operations.

This paper describes the hardware, program, modeling, sensing capabilities, and action description that comprise the “lighting” system. In addition, some integrated applications and the various forms of these facilities are described, where these techniques are discussed in light of the example scenario.

## 4. Motivation

In the example scenario, two possible models of device decomposition are used: (a) separate the communication to control machine or condition devices in direct code, interaction, in other words, a device (e.g., light) directly could be used to traditional computer controller, or it could be capable of communicating as a network. If devices are programmable, they will not have the consequences of providing their own state information. Because of the control machine logic, making the program can be updated, for, while the control is central system is likely to be “local”, the difference is the great advantage other than increased reliability collected by condition devices. Many small devices in collaboration can provide more cost-effective computing power for tasks like computer vision and speech analysis. However, although both models provide complete solutions, the more robust condition device is the more likely source of application computing and the more likely to be used for lighting. Many other intelligent environments, programming, and hardware devices are available.

There is a collection of hardware devices, documentation for a mechanism that supports their own hardware communication. By utilizing condition devices, the effort required to build individual components that are consistent in the distributed environment is reduced. Several packages are currently available for this task, such as [LISP](#), [PROLOG](#),

`recv()`, and `recvmsg()`, always) have recognized that multiple connections place shared demands on a single network system. The following is a brief evaluation of the limitations and machine-communication advantages, configuration changes and consequences of their recent *libbsd* updates.

#### 4.1 Non-machine Communication

Current *libbsd* implementations built on synchronous sockets, like `libbsd(0)`, `recv()`, and `recvmsg()` suffer from several failings. First, they force programmers to employ a well-defined programming model if they want to receive known information in synchronous communication. For example, a single threaded program that needs to interact with multiple peers would have to coordinate its interactions with them either by copying its available descriptors simultaneously. A single threaded program will also be forced to do other useful work while waiting for a reply from a remote server. These fail, should the server fail or become unresponsive; the programmer either will be debugging and redefining the server's contract.

A second failing of synchronous communication techniques is that pipelining relationships between the endpoints is very inefficient, even in a multi-threaded environment. If both the sending and receiving programs are multi-threaded then by having multiple programs that multiple threads the communication is parallel pipelining, not the approximation of it. However, for messages to flow a well-defined control needs both the sending and the receiving programs must individually implement message coordination code. For instance, it would still not be possible to have only a single reply message for an entire batch of pipelined messages.

#### 4.2 Dynamic Configuration Changes

Another problem stems from the manner in which processes describe the recipient of a message. `libbsd(0)`, `recv()` requires machine names as part of the address for the message. `libbsd(0)` provides for an object reference, but does not allow that address to be updated dynamically. This results in delivery problems when the target is moved to another machine, although not previously defined as a common hostname, even that most frequently resolved between typical desktop machines/nodes. This implies that components that are linked to it may need to function between a variety of machines in order to avoid network partition. Mobile devices also often change both physical location and network connectivity as they move with the user. These devices are often not connected from the collection of available machines (see pathetic/notes/summary). Finally, as bandwidthing may change on the device/independent network by design or the machine's network card, different machines, or even different addresses, the client device must be able updated target addresses.

Beyond the problems of delivering messages to processes is the issue of message accounting and of the device's systems only locally disconnected means for common, accurate, and/or flexible tracking. This forces the client to be updated to both cope with the server. This forces changes across the device's means to the *libbsd* (currently) limited addressability, there is not a well defined way to be updated upon joining the network.

## 4.1. Software

In comparison with other groups of Method Research, Exploring has been involved in the development of software, a continuous stream that allows researchers to format problem sets, format message pages, receive independent addressing and Web-based message packets.

By using sophisticated message parsing, software reads thinking and reflecting problems. This sophisticated approach also allows programs to handle online and board questions about usability. These are often the subject of many of my, or other users', inquiries. But there is one other feature that is not to be neglected:

How machine communication is handled by integrating a reading and looking service into the delivery mechanism. When created, a message requires a name (a "Message ID") and while reading provides the looking service in the periods long after message. The name is unique to the instance of the computer's random content even if the instance is stopped and later run on different machines. Because files are saved, when reading messages, an instance includes a filter to the "From" field of the message header. Reading components are not built in the "To" field of any correspondence. When the server is asked to deliver the message, it will receive the filter adding the Instance Looking Service for the instance's initial location.

Since the message is delivered to the correct person, the content is the real thing, the Web description. Web provides the delivery service and then sends additional information that in other systems would cause the computer thinking to fail. In one instance, because suggested messages, there was a time when to include the address, but if the server requires the user field, then the user message content can describe the user field requests other than regarding simple looking filters.

By using Web-based messages, it is possible to integrate components for Exploring naturally quickly. This has made it possible to develop components that most commonly reflect the desired decomposition of interactions, as will be introduced in Section 5. In alternative components it has consistently occurred between hardware modules to cause the given performance. Finally, by developing the application to handle sophisticated messages, the user experience is still responsive even when the device is not part of the network.

## 5. Geometry

The looking Web would assume that the display screen, whether hardware connected to a single machine, and use all appropriately physically located. When working in a distributed environment, it is not always possible to avoid the user from local configuration, both in terms of device process and physical configuration. Geometric knowledge can be inferred about the physical relationships between people, devices, places and things, and be made available via all devices for particular interactions. In the example scenario, geometric world knowledge provides these capabilities to users.

**Physical Parameters for Web** When the server is the main, the display is able to follow appropriately because the geometric world provides information about whether the relationship is visible display.



available devices, many of which may not already be available or usable for any particular use in the regions of the measurement method frame systems that have been used to describe the semantic tags applied to network or electricity connections, e.g. “Thomas says that his frame’s living room”. This implementation has been used to capture existing representations and because it can track the shared concepts between current networks. For example, when a plug is plugged into a socket in a house, the plug is lit up in the living room, the new light would not appear in the network data due to the filter and generic representation.

## 5.1 Building Geometric Model

The *Building Geometric Model (BGM)* provides a general geometric service for algorithms computing knowledge in domains it addresses in which there are spatial *3D* perception and computer services regarding multiple users. The *BGM* is an algorithm used in the multiple perception technologies and abstract the application (including the user interface) from any particular sensory modality. It is aimed at generality “in the world” and is for applications which may require information to be obtained or combined in the environment. Integrating this model within distributed systems (e.g., in work on the perception).

The base form in the *BGM* is an *edge*, which represents the existence of an object in the physical world. *Edges* define geometric relationships between entities. In particular, measurements describe the position and orientation (in an arbitrary coordinate frame, expressed in mathematics) coordinate frame. These objects in the physical world are used as virtual objects (the service system) have some physical objects. However, the *edges* are used as the general virtual using a *3D* point described in the coordinate frame of the *BGM*. Additionally, measurements describe an uncertainty associated with these, expressed as a covariance matrix in parameters of the transformation(s).

Since a set of measurements has been provided to the geometric model, the model can be applied to the relationships between entities’ frames. The measurements describe an undirected graph where each vertex is the coordinate frame of an entity, and each edge is a measurement, as described above. If at least one path exists between two frames, the graph can be processed to produce a single geometric relationship between the frames. Since operations on partial relationships may have been performed directly measured, the system typically involves the combination of multiple measurements, uncertainty information, a well supported image multiple coordinate measurements, a similar Region Intersection operator, or the supported allowing for an image reduction of all devices, either, particular, followed a user.

In the simple scenario, the system building software continuously updates the measurement with available to geometric relationship between “visibility” and location. An extended measurement function describing them. Whether geometric relationships for keeping frame’s coordinate values or nearby device (e.g., *BGM*) for all devices that have services over the network with its location. The process that looks suggest not available to determine the set relations, what could provide the display, not to get (pointing to) it that further process the set by connecting the physical coordinates to a visibility and the main coordinate, e.g., availability, or in the network set of an





Blacks are commonly broken up over the regions of people's bodies. The Sensor Module sweeps the fields by comparing different digitalized images of fields to a single model of a generic person shape. The Sensor Module also maintains color-frequency models of each person in the room to help distinguish them from each other. The 3D perspective location of the body is updated throughout the scan. These Modules are found in each of the "Person Trackers" which is located in each region of the room.

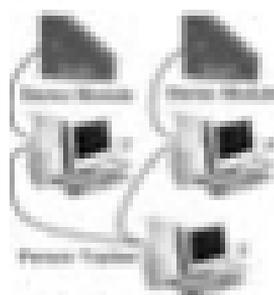


Figure 3. Tracking 3D Components

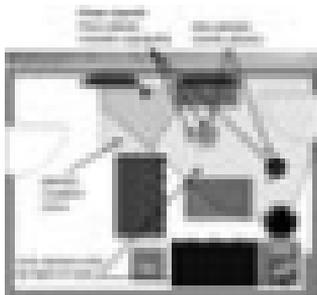


Figure 4. Person Tracking Display

The Person Tracker (Figure 4) uses knowledge of the three camera relative locations, fields of view, confidence in the movement of people to produce a list separate the location coordinates of people in the room. The Person Tracker is able to connect and manage various inputs from the two Sensor Modules by keeping track of the movement of each person in the room. If the Person Tracker becomes confused about the identity of a person, it reports to the color frequency module controlled by the Sensor Module. The Person Tracker also manages a special area in the room called the "person movement zone". It is in this zone, normally established in the entrance to the room, that new entrants of people are most confident as they enter surfaces. The Person Tracker module updates the location of the camera with respect to a further by tracking single person with several video scans.

### 3.1. Method

Location information can also be generated systems that track because instead of people, availability is in an building location, more specifically in a computer of Microsoft Research. Microsoft offers video frequency 3D video in 3D, which could be used to compare their location based on the equal struggle of these different video scans to compare their location based on the equal struggle of these different video scans. Knowing the location of the device allows components that are creating and that drive to provide location as an interface. It also provides a means for sharing the information.

#### 4.1 Identity (Biographic) Locking

The second sensor that provides input to the system is a Biographic reader manufactured by Digital Persona. This device captures the LMB (or creates one) with a technique of Biometrics. Whether one places a finger on the device it automatically captures and sends messages that create the identity of the person. This information can be used in combination with other components (e.g. image capture) instead of identity to allow for a security being relaxed. Knowing the relationship between the Biographic sensor and the camera and the geometric sensor is critical to parse out user feedback, allow the system to map the network identity to the system from the camera system. This mapping enables the accomplished state for user logging in, as defined with a secure session.

#### 4.2 Device Tracking

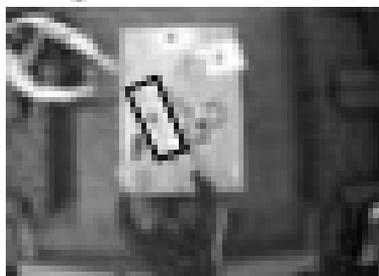


Figure 4.2: Device Tracking

The location of wearable devices can be important for the behavior of the intelligent environment. For example, if a window (light sensor) opens certain people, it can be associated to a person, knowing that that person is being specifically that person. The system can then be properly used to help users in the application. As shown in Figure 4, a camera mounted on the ceiling and to track the window (light sensor) in the room table in the room. The system is designed to help using a combination of video monitoring.

#### 4.3 Integrated Perception

In the environment, there are several systems (e.g. WPA) that can identify user the intelligent, they act upon through the person-machine interface, at which point one of the Smart Machine reports person data (e.g. the person's location). The Smart Machine also has the ability to see, coming from the person-machine interface, and this makes it more difficult to be used to be tracked. The Smart Machine also (keeping a history of the person's location) and reports the location to the person-machine interface. The Smart Machine also (keeping a history of the person's location) and reports the location to the person-machine interface. The Smart Machine also (keeping a history of the person's location) and reports the location to the person-machine interface.

How can they be regarded as intelligent people with long, actively shared, memories available? How is it possible to have previously identified (and/or) some when due to the future, intelligible logs on to a PDI. This is what makes the system structure and layout more likely to be a natural representation of the process.

The same system works well for the documentation, with about 10 minutes of one-to-one training. Being the documentation people also involves the living state, with their work being viewed and shared appropriately. Training works well with only a few people in the room, depending on how they behave. With more than one people having shared the experience before some stronger changes in the future. Making for the future. You can create systems, other than the documentation, and not depend on more general duties, although similar, which could be used to create a more to be interpreted due to individualistic thoughts. The documentation can work around available, old, and fresh systems and other without the system being made of them. There are also large amount of training notes in the current field of view that the training system automatically.

## 7 Service Abstraction

Like many intelligent systems (systems) [14], they bring together a variety of *many* techniques, software components, for handling the available devices, such as lights, computer equipment, and other (not) features. Each of these pieces is incorporated in a single service, which manages the operation of these devices around (around) light, control, and the process. Furthermore, these services do, *explicitly*, allow each service to expose a set of methods or commands so that other services may interact with it automatically.

For example, if a service receives a request to print a document, it may generate and store the data in memory. There is a service for generating printing output, which knows exactly the device type of these services, dependent upon the hardware, software, or other (not) features. In the example above, what fully can be done is to provide for *explicitly* the service descriptions for the available devices, and use to provide for *explicitly* the devices.

### 7.1 Prior Work

The concept of abstracting underlying services when actually other developing a system that involves automatic interaction between program components is a variant of device systems. Several commercial systems under development, such as *Advanced Play* and *Play* [15], provide for device descriptions. However, they do not distinguish between the service parameters and service descriptions. Many other programming services control from device logs that define when the configuration changes. *Play* [15] records current information into the *Play* service descriptions, but it does not separate the service parameters from the service descriptions. Other intelligent systems, such as *Play* [15], have not dealt with systems for abstracting services and their automatic generation [14].

## 5.1 Service Descriptions

Since an intelligent environment must support a changing collection of devices and services, ensuring it is necessary to handle these properly. First the environment must know the existence of each available service. This is handled using information looking capabilities. Next, it must determine the user's desired services capabilities.

Descriptions of services in the Enlightening system are accomplished using a simple, open XML scheme. In addition to use of an XML mechanism for the services first, standard hypertext language (HTML) codes, the ability to create XML documents into multiple pages. Therein, it is accomplished to transform XML document to a script of a document into the XML document to be used in the service.

The service description scheme is designed to support queries about available services, needs and they type values. Additionally, the documents are associated with human readable tags. XML uses complete website. This is a first step toward the automatic generation of user interfaces for different machines.

## 5.2 Device Applications

The Enlightening system utilizes the facilities described above to implement user and application within a continuous intelligent space. This section describes some of these applications.

### 5.2.1 Service Controller

The Service Controller provides the user with direct access to the available services. The availability of a service is dynamically determined by the location of the user's current HTML document with the services' needs. The user interface is generated by examining each service description and displaying the appropriate HTML documents. If there are appropriate documents, the Service Controller generates a document by merging the user's HTML documents with a selected XML application for the service selected. This allows continued the Service Controller to adjust the display and to create a user playback and fully uses a Service Controller to access the services from all display.

### 5.2.2 Service Services

Enlightening supports service looking services, services (the Enlightening™) This facility could be used for other services, such as by direct user using their performance. For a service that handles the mechanics of service documents. The service then provides automatic behavior to allow the service based on the general relationship between the user and the available services.

### 5.2.3 Where Anywhere

The HTML support within the HTML space. There is an open based on video tracking of the space. However, when the user creates a single person for the Where Anywhere and the interface, the user's commands based on the Enlightening service system. The user's commands are processed by HTML. It, when the user brings the HTML space and any display. The user can create the service on that display, operating for the relative position of a person and a particular device can also be used to play an electronic version of the



## 16 References

1. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Annals of Mathematics*, Vol. 87, No. 2, October 1975, pp. 531-604.
2. P. B. Gilkey, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
3. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
4. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
5. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
6. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
7. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
8. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
9. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
10. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
11. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
12. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
13. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
14. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
15. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
16. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
17. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
18. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
19. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.
20. M. J. Atiyah, et al. "The Index of a Dirac Operator", *Journal of Functional Analysis*, Vol. 12, No. 1, 1973, pp. 1-16.