

# THE VIRTUAL CHOCOLATE FACTORY; BUILDING A REAL WORLD MIXED-REALITY SYSTEM FOR INDUSTRIAL COLLABORATION AND CONTROL

Maribeth Back, Anthony Dunnigan, Jonathan Foote, Sagar Gattepally, Don Kimber,  
Bee Liew, Eleanor Rieffel, Jun Shingu, Jim Vaughan  
FX Palo Alto Laboratory

## ABSTRACT

Virtual, mobile, and mixed reality systems have diverse uses for data visualization and remote collaboration in industrial settings, especially factories. We report our experiences in designing complex mixed-reality collaboration, control, and display systems for a real-world factory, for delivering real-time factory information to multiple types of users.

In collaboration with TCHO, a chocolate maker in San Francisco, our research group is building a virtual “mirror” world of a real-world chocolate factory and its processes. Sensor data is imported into the multi-user 3D environment from hundreds of sensors on the factory floor. The resulting “virtual factory” is designed for simulation, visualization, and collaboration, using a set of interlinked, real-time layers of information about the factory and its processes.

We are also looking at appropriate industrial uses for mobile devices such as cell phones and tablet computers, and how they intersect with virtual worlds and mixed realities. For example, an experimental iPhone web app provides mobile laboratory monitoring and control. The mobile system is integrated with the database underlying the virtual factory world.

These systems were deployed at the real-world factory and lab in 2009, and are now in beta development. Through this mashup of mobile, social, mixed and virtual technologies, we hope to create industrial systems for enhanced collaboration between physically remote people and places – for example, factories in China with managers in Japan or the US.

**Keywords:** mixed reality, 3D applications, advanced visualization, remote collaboration, virtual worlds, remote monitoring, remote control, data visualization, pervasive computing, simulation, mobile computing, collaborative tools.

## 1. INTRODUCTION

Virtual and mixed-reality environments create billions of dollars in revenue every year, in applications like simulation, games and education. What are the appropriate uses of these technologies for the enterprise? We are exploring a variety of uses of virtual, augmented, and mixed-reality environments for industrial systems, such as real machines and processes in a factory. We are interested in understanding appropriate methods for communicating and collaborating in these environments, and, to aid this

kind of collaboration, in importing real-world sensor data (such as temperature controls, machine state, and environmental status) and multi-camera imagery from the real factory floor.

This research is aided by an ongoing collaboration between our research lab in Silicon Valley and a local company that runs a real-world chocolate factory. TCHO, the chocolate company, is a start-up focused on creating “high-tech chocolate” by applying new technologies with careful attention to detail at every step from cacao bean sourcing through chocolate production, shipping, and marketing [12]. As such, they are open to the shared experiment of inventing and installing mixed-reality systems for factory applications.

### 1.1 Mixed-reality systems: related work

Mixed-reality (or cross-reality) systems are those where ubiquitous sensor/actuator network infrastructures in the real world and augmented or virtual world systems interact, creating a “seamless electronic nervous system that extends across people, things, and places” [9]. Physical/digital integration systems such as Phidgets and Arduino, developed for ubiquitous computing projects, are sometimes used for the sensor/actuator aspect of mixed reality [3, 1]. Often, as in the present case, off-the-shelf network and control components combined with custom hardware and code may be used. Several types of interaction and display may co-exist in such a system, creating a mixed-reality continuum ranging from small mobile devices to large displays, and from simple data display to full-featured visualization and control. Applications for these systems have included data center management, home monitoring, inventory tracking, and social network awareness augmentation [4, 5, 6, 8].

For the enterprise, ease of use, security, and integration with existing systems are key to adopting new technologies. The Virtual Factory project is a case study in designing and using a mixed-reality infrastructure in a working factory with existing control software.

### 1.2 The Virtual Factory concept

Rather than simply build a virtual world replicating the factory, our aim was to create an electronic infrastructure that captured data from many of the company’s processes, and then find useful ways of using and displaying that information for different sets of people in the company. In some cases, the system offered machine control as well.

So far, we have developed several applications based on this infrastructure, optimized for virtual and mobile teleoperation, simulation, data visualization, and lightweight collaboration.

The information environment itself is a complex set of interlinked real-time video, 3D and text-based layers of information about the chocolate factory, with lightweight collaborative/social network functionality for multiple users. Mixed reality and mobile applications are part of the information stack; so is a virtual “mirror world” 3D representation of the factory floor. In this Virtual Factory world, data layers from the underlying database show up in a realtime, multi-user 3D world (or worlds). These visualization layers each focus on different data, different functions and different users. Because the Virtual Factory world is multi-user, just like the popular virtual worlds of Second Life or There.com, people can meet and talk with each other inworld. Each layer uses appropriate representation for its purpose: multi-camera video streams, or representational “mirror world” graphics, or more abstract data visualizations representing the current or past state of the sensor fusion network in/around the TCHO factory and associated processes.

For example, inside the virtual factory world, you can visit a 3D model of a machine to read its sensors' status, or move closer to it to trigger an inworld video overview of its function. This multi-user collaborative space can be used for tasks like remote factory observation, client tours, employee training, and data visualization.

On the mobile side, one experimental iPhone app provides lab techs with mobile monitoring and control, while another gives factory managers realtime, mobile access to sensitive temperature control systems. In the chocolate development lab, where intricate processes are developed to bring out the flavors in each bean, accurate tracking of time and temperature are essential. The app allows a real-time view into the lab (via pan-tilt-zoom steerable camera) and individual control over machines and sensors. Data from the lab is also represented in the 3D virtual world. Finally, a network of video cameras installed around the chocolate factory and inside factory machines streams live video to web, virtual world, or mobile web apps.

Mashups between the data representations in different applications can call out “hot spots” or hitches in production, or suggest a need for maintenance. As is common with technology mashups, it's likely that some useful qualities will emerge with continued use.

### 1.3 Choosing what to build

One of the ideas behind the Virtual Factory project is to create 3D information environments that enable a number of different applications, depending on the type of user in the virtual world. These collaborative applications are aimed at specific tasks performed by specific types of users. At the TCHO factory, we have identified seven potential types of users, with subsets of each:

- Managers (though the CEO and the operations manager need different data and level of detail)



Fig. 1. An avatar in the Multiverse-based Virtual Factory.



Fig. 2. A photo of the factory floor under construction.

- Operators (technical, mechanical, maintenance – some with expertise on different machines)
- Trainees (learning each machine’s processes and getting an overview of the whole multi-day chocolate process)
- Tour guides (for real-life tours, and virtual world tours – or both at the same time)
- B2B clients (technically knowledgeable, more interested in process and product reliability)
- Suppliers and sources (cacao farmers and producers, other materiel)
- Consumer customers (on-site tours, at point-of-purchase in the TCHO store, or via Web-based virtual world online)

Using the same virtual world for all these types of users at one time may not be viable. However there are logical and even necessary crossovers between some user types. For example:

- Managers and operators need to discuss daily operations with each other
- Managers and B2B clients look at the same data in discussing business
- Operators and trainees need to share an environment for teaching
- Tour guides and clients or customers need to share the same environment

Some users and applications are better supported by augmented reality or mobile remote control systems, rather than a virtual world. For these users, we can design collaborative applications for tasks such as:

- remote factory observation, machine monitoring
- process/workflow monitoring and analysis
- augmented reality-based inspections
- mobile teleoperation
- b2b customer visits (augmented reality)
- education and training of employees
- visitor tours
- inventory tracking

We began the research collaboration by identifying who to work with inside the chocolate factory, and who would best be served by the technologies we planned to develop. At the same time, given the company's start-up status, TCHO staff were very busy; the project needed to be primarily designed and built by the research lab's staff, based on observation and interviews. The first two applications we chose to build on top of the mixed-reality infrastructure (a virtual world and a mobile monitoring and control iPhone app) were agreed upon as being both immediately useful and generally extensible for future applications.

#### 1.4 Building the virtual factory

To create the mixed-reality system, we needed both a virtual world and a source of real-world data to match it with. The TCHO factory is run by commercial industrial control software called Wonderware InTouch, which is in widespread use in factories around the world [14]. We worked with Wonderware engineers to access the underlying MySQL database for the system that controls most of the factory's machines and sensors (see system architecture section below).


The resulting data stream is parsed via XML and distributed into the Virtual Factory world. The data appears in floating text against a transparent color cloud, called a "Data Spot." This design allows sensor data to be easily read inworld, while not blocking visibility of the machine model it's attached to. The Data Spots are animated sensor display indicators showing such data as hot water temperature, chocolate temperature, or machine state.

The machine models and building model were developed by a digital artist, from measurements and photographs, not from CAD models. Live video is available inworld from ten Axis network cameras around the factory floor, some of them actually inside machines (Fig. 3). The video from all cameras is recorded to and stored on an onsite digital video recorder, for more manageable access to multiple streams.

#### 1.5 Virtual Factory deployment

As a first deployment of the virtual factory, a large display (a Samsung 56" LCD screen) and a client/server computer running the Multiverse virtual world platform were installed on the real factory floor.

For ease of use, the machine boots automatically into the Virtual Factory world (Figs. 1 and 4). To simplify matters for factory personnel giving live tours, we created an automatic animated tour in the virtual world that moves an



**Fig. 3.** A network camera inside the chocolate molding line lets factory operators see whether any chocolate is stuck in the molding tray.



**Fig. 4.** The Virtual Factory on the Multiverse platform, installed on the real factory floor.

avatar from one machine to the next, showing HD video on the function of each machine. However, this is still a full-featured virtual world – a user can choose to break away from the automated tour, and to use normal inworld features such as text chat, audio chat, etc.

This system is currently in daily use by TCHO personnel for tours for B2B clients and the general public.

## 2 PLATFORM CHOICES

One of the biggest barriers to entry for using virtual world platforms in the enterprise is that they have a reputation for being expensive, costly to develop, and difficult to maintain on a day-to-day basis. Other concerns include security, ability to import and view office documents, and video handling. We developed a list of requirements and then experimented with several virtual world platforms; this is an ongoing process. Virtual world platform features we looked for included:

- Open source; GPL license preferred

- Enterprise-quality security: allow running servers locally (behind our firewall)
- Very stable, able to support many users and large graphics
- Good API
- Good server/client messaging communication
- Customizable: allow python scripting or C++ or C# or Java plug-ins
- Efficient content import pipeline  
Direct file import (drag-n-drop) is the easiest; or support import through external tools external to the world
- Ability to import content in multiple formats.
- Preferably in SketchUp, Collada, VRML, or possibly other formats like 3DS, Wave front OBJ, or mesh, for which there are tools to enable conversion and importation
- Uses physics engines; ability to import physics content.
- Good documentation, community forums, 3rd party tools
- Sophisticated avatars, customizable
- Sophisticated 3D graphics handling, good lighting
- In-world features:
  - Allows world-editing using a GUI tool including drag and drop features
  - Has collision detection and physics engine
  - Allows playing/building with primitives
  - Allows application of textures/materials to primitives
  - Allows import of animations (bvh, other) and attaching the animations to objects
  - Client should run cross-platform: Windows, Mac and Linux
  - Supports live video, embedded video, MJPEG, audio, Open Office document viewing (and preferably editing), app sharing, VNC, and audio and text chat

We have not yet found a virtual world platform that meets all these requirements, though several come close. We experimented with a number of platforms, spending a good deal of time working with Sun's Wonderland platform and the Smalltalk-based Qwaq Forums (now Teleplace). [13, 11] Currently we are using a game platform called Multiverse, and exploring several other platforms (OpenSim, ExitReality).[10, 2]

## 2.1 The Multiverse platform

The Multiverse virtual world platform is designed to support MMOG (massively multiplayer online gaming). It has an interesting client/server architecture that allows us to customize or add new custom servers. The messaging architecture is also very efficient. From a researcher's standpoint, its licensing structure is also convenient: the development environment for both client and server is free to use. Multiverse's business model is to charge a small license fee only once a company begins to make money

using their platform. Though Multiverse is not open source, this business model meant that we could experiment with this platform without too much upfront cost (in contrast to some other platforms who do charge a licensing fee).

## 3 SYSTEM ARCHITECTURE

The factory setup at TCHO uses PLCs (Programmable Logic Circuits), alarms, sensors and cameras for capturing data from the machines and the factory. The data is accessed and stored by custom machine controllers, Wonderware InTouch factory automation and video servers; then it is processed and distributed to allow various client applications to access it. We have successfully displayed the real-time factory information in the Virtual Factory and on mobile devices running in intra- and internets.

The data path from the factory control system to the virtual world is as follows (see also Fig. 5):

- A Java program on the virtual world (VW) server machine reads the PLC data (from the Programmable Logic Circuits that run many of the factory machines) from a WonderWare DA Server through an FSGateway using DDE protocol. This updates a database table every five seconds.
- Then a Web app (servlet) on the VW machine reads the local database (instead of pinging the Wonderware system) to send the data to the request from here.
- Finally, the Multiverse server on the VW machine reads the data from the web servlet and broadcasts it to all Multiverse clients.

The Multiverse server is implemented in such a way that it sends the data to the clients based on the user's permissions level, currently defined as either "manager" or "operations personnel". These views differ in that operators have a more cluttered view; they see more level of detail on each machine, including alarms and maintenance history, while managers see the only high-level data on current machine state.

### 3.1 Multiverse client customizations

We added many customizations in the Multiverse client through the Iron Python interface. The customizations can be categorized as follows:

*Document support:* This feature is important for enterprise applications; a good deal of corporate communication is handled in office documents. The Mutiverse client now can support documents such as pdf, word, excel, Flash, or html. The documents are embedded in the world in the Flash format, but when the user clicks the Flash panel it opens up the real document in an overlay with which the user can edit, print, save a local copy, or any other actions that would normally be performed with those documents. (Both Wonderland and Teleplace also have robust document support.)

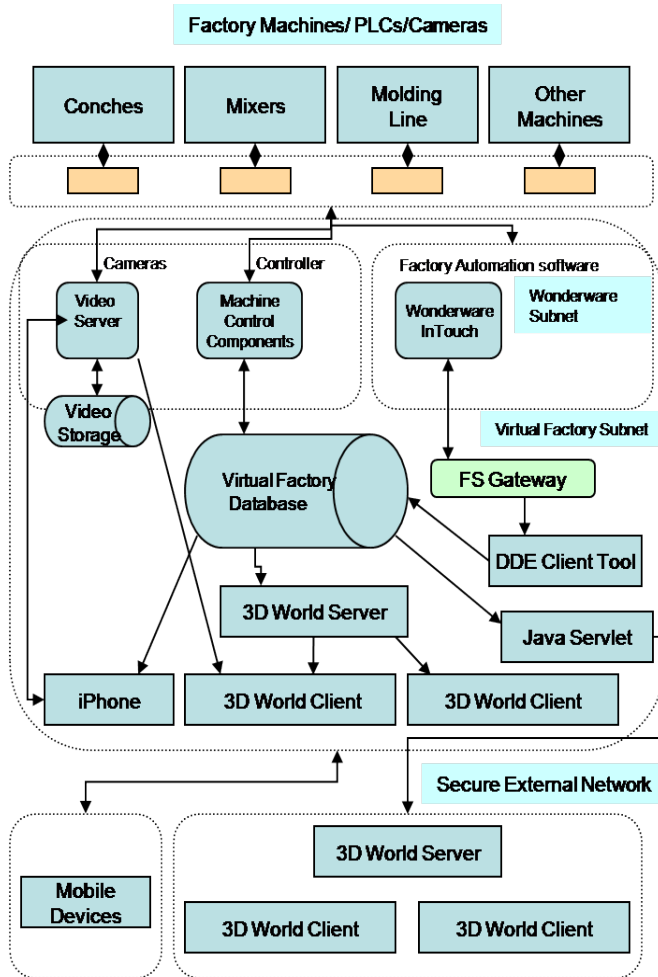


Fig. 5. Mixed reality factory connection architecture

*Server customizations for live factory sensor data import:* the sensor data is read from the MySQL database by the server and sends the data to the user's client, if they have the correct privileges to see this sensitive information. These privileges are configured on the server.

*User input devices:* support for Wiimote and for 3dConnexion devices

*User keyboard customizations* for avatar control for standard user actions like increasing or decreasing avatar speed, controlling rotation speed, and switching between walk, run, or fly modes; or controlling the automated tour.

*Machine context information:* an in-world video panel, live machine sensor and status data, and an animated pop-up machine showing machine interiors.

*Video Support* for networked Axis cameras, and the capability to play remote or local videos.

Integrating mobile with virtual and mixed reality systems is a logical extension for industrial smart environments [7]. With this in mind, we have built an experimental iPhone app that provides mobile laboratory monitoring and control.

The tasks on the main factory floor differ from the tasks in the smaller TCHO development lab, where cacao beans from different growers and geographic areas are analyzed and new recipes developed for them. Each new harvest or type of cacao bean may require different treatment: a longer, slower roast, or a higher grind temperature. Accurate tracking and control of time and temperature are essential. Using the iPhone, machines can be controlled remotely: power turned on/off, temperatures and timer settings can be read and changed, and lab environmental controls can be changed.

Six people regularly use the TCHO Lab. One of the most difficult aspects of the job is that the chocolate melangers (heated grinders) may be set to run from 14 – 26 hours, often with five or six simultaneous runs, and each with different timing. This extended time period meant at least two people needed to be involved, usually more. However, with the remote control app, lab personnel can start the melangers and then control temperature and power-down time from wherever they are, via iPhone. Alarms can be set for push notifications to remind users when needed.

The mobile app provides greater accuracy and transparency for the people operating the TCHO Lab, and runs via a database that allows accurate tracking over months of the myriad combinations of cacao bean selection and treatment.

The iPhone application gives the user a real-time view into this lab (via PTZ or Pan-Tilt-Zoom steerable camera). The mobile system is integrated with the factory's multi-camera surveillance system and the database underlying the virtual factory worlds.

#### 4.1 Mixed-reality architecture for mobile/virtual

Underlying the lab application is the central database that provides data to all devices: mobile, virtual, and web via python scripts. The current rule for interaction precedence is atomic interaction – the most recent command is executed. (If the system is required to scale up, we will revisit this, adding stricter permissions and an awareness of how many other people are currently logged in, and who they are.)

The hardware setup in the TCHO lab that is driven by the iPhone app consists of off-the-shelf components: an Ethernet 10/100 feed to a Control Ethernet-to-serial (16-port) box. This communicates with eight Omega temperature sense/set serial devices for the melangers and roaster ovens. It also controls a serial controlled isolated I/O 8-relay module controls AC power switches on melangers and lab light. A set of custom Python scripts communicate with the main database and web server.

## 5 CONCLUSION

The Virtual Factory mixed-reality infrastructure, along with the two applications of it we report on in this paper, form a case study in the uses of mixed-reality systems for the enterprise. Both the Virtual Factory world and the iPhone




Fig. 6. Screen shots from the iPhone app for controlling and monitoring machines in the chocolate development lab.




Fig. 7. Mixed-reality examples of data visualization: a control panel for a temperature sensor and controller, built in Qwaq Forums (now Teleplace) and communicating with sensor network database via Python. Control operations performed on the iPhone show up here, and moving the sliders inworld here will update the data on the iPhone app.

monitoring/control app are in daily use at TCHO, and we continue to iterate the design of each, as well as take on new ways to leverage the rich mixed-reality infrastructure we have built. Next steps include a mobile display of the Wonderware data from the factory floor, for TCHO's engineers; and a tablet-based augmented reality maintenance log for factory machines.

The most difficult part of this work was choosing and working with a virtual world platform. Though we were able to build reasonable representations of the Virtual Factory in all three virtual platforms we tried, we have not yet found a virtual world platform that completely meets all our requirements for a robust, secure, programmable or scriptable, yet inexpensive enterprise application. Developing 3D assets is also still a time-consuming process, although cross-platform formats like Collada help with the platform testing process. However this is a rapidly developing area; we are continuing to test and compare new virtual world platforms as they come out.

In particular we see promise in some browser-based 3D platforms such as ExitReality. We are also exploring the RealXtend/OpenSim virtual world platform, a true open source platform with excellent functionality; and we are interested in the idea of 3D on mobile devices.

We believe that our real-world, hands-on approach to this mixed reality/virtual world project is unique in this developing field, and we believe there is value in this space for enterprise companies, particularly in a versatile infrastructure that allows ongoing mixed-reality integration of mobile, ubiquitous, augmented and virtual applications.

## 5. REFERENCES

- [1] Arduino: <http://www.arduino.cc/> (link: 1/15/2010)
- [2] ExitReality: <http://www.exitreality.com> (link: 1/15/2010)
- [3] S. Greenberg , C. Fitchett, Phidgets: easy development of physical interfaces through physical widgets, Proceedings of the 14th annual ACM symposium on Tangible and embedded interaction, 2007. pp.209 - 218.
- [4] "IBM Launches 3-D Data Center in OpenSim," Virtual Worlds News, 21 February 2008; [www.virtualworldsnews.com/2008/02/ibm-launches-3.html](http://www.virtualworldsnews.com/2008/02/ibm-launches-3.html) .
- [5] J. Lifton , M. Laibowitz , D. Harry , N.Gong , M. Mittal , J. Paradiso, "Metaphor and Manifestation—Cross-Reality with Ubiquitous Sensor/Actuator Networks", IEEE Pervasive Computing, vol. 8, no. 3, pp. 24-33.
- [6] P. Milgram and F. Kishino, "A Taxonomy of Mixed Reality Visual Displays," IEICE Trans. Information and Systems, vol. E77-D, no. 12, 1994, pp. 1321–1329.
- [7] M. Musolesi et al., "The Second Life of a Sensor: Integrating Real-World Experience in Virtual Worlds Using Mobile Phones," Proc. 5th ACM Workshop Embedded Networked Sensors (HotEmNets 08), ACM Press, 2008, pp. 12–16.
- [8] Y. Ohta and H. Tamura, Mixed Reality: Merging Real and Virtual Worlds, Springer-Verlag, 1999.
- [9] J. Paradiso and J. Landay, "Guest Editors' Introduction: Cross-Reality Environments", IEEE Pervasive Computing, vol. 8, no. 3, pp. 14-15, July-Sept. 2009.
- [10] OpenSim: <http://www.realxtend.org/> (link: 1/15/2010)
- [11] Teleplace.: <http://www.teleplace.com>
- [12] TCHO Ventures, Inc. <http://www.tcho.com/> (link: 1/15/2010)
- [13] Wonderland: <http://www.projectwonderland.com/> (link: 1/15/2010)
- [14] Wonderware: <http://www.wonderware.com/> (link: 1/15/2010)