To appear in: *Proceedings of the 2003 IEEE/RSJ International Conference on Robotics and Automation (ICRA 2003)*
Taipei, Taiwan, September 2003

Ada – Intelligent Space: An artificial creature for the Swiss Expo.02

Kynan Eng, Andreas Bäbler, Ulysses Bernardet, Mark Blanchard, Marcio Costa[1], Tobi Delbrück, Rodney J Douglas, Klaus Hepp, David Klein, Jonatas Manzolli[1], Matti Mintz[2], Fabian Roth, Ueli Rutishauser, Klaus Wassermann, Adrian M Whatley, Aaron Wittmann, Reto Wyss, Paul F M J Verschure

Institute of Neuroinformatics, University/ETH Zurich, Switzerland
[1]Nucleo Interdisciplinar de Comunicacao Sonora, University of Campinas, Brazil
[2]Psychobiology Research Unit, Tel-Aviv University, Tel-Aviv, Israel

Abstract

Ada is an entertainment exhibit that is able to interact with many people simultaneously, using a language of light and sound. "She" received 553,700 visitors over 5 months during the Swiss Expo.02 in 2002. In this paper we present the broad motivations, design and technologies behind Ada, and a first overview of the outcomes of the exhibit.

1. The Ada Project

Ada is an interactive space developed for the Swiss national exhibition Expo.02 located in Neuchâtel. Conceptually, "she" can be seen as an inside-out robot with visual, audio and tactile input, and non-contact light and sound effectors. Visitors to Ada are immersed in an environment where their only sensory stimulation comes from Ada herself (and other visitors). Like an organism, Ada's output is designed to have a certain level of coherence and convey an impression of a basic unitary sentience to her visitors. She can communicate with them collectively by using global lighting and background music to express overall internal states, or on an individual basis through the use of local light and sound effects.

To realise such a space, several simultaneous lines of research and development were pursued. Topics under investigation include:

- Audio processing and localisation
- Multi-modal tracking
- Automatic music composition
- Real-time neuromorphic control systems
- Human-machine interaction via whole-body locomotion

From a more application-oriented perspective, Ada was used to gain practical experience in handling large-scale behavioural integration issues in autonomous systems. Certain components of the final system are also being evaluated for technology transfer projects.

Development of Ada commenced in late 1998 and ramped up to a maximum team size of about 25 people. A total of over 100 people were directly involved in the construction and running of Ada. The exhibit ran continuously for up to 12 hours a day over 5 months from 15 May to 20 October 2002. During this period 550,000 visitors entered the space.

Diversity was a defining characteristic of the project. The technical development team came from many different nations and disciplines, ranging from biological sciences through engineering to musical composition. On top of this was a team of architects, artists, publicists, scenographers, on-site managers and guides for handling the production and operation of the exhibit. The financial stakeholders represented most of the main sectors of Swiss society: government (via the Expo.02 organisation), education (ETH Zurich), private industry (Manor, a department store chain) and two private foundations. Correspondingly, the project goals had to address many issues simultaneously. These issues included: contractually defined requirements for system uptime, visitor numbers and standard of entertainment experience, special events for sponsors and media, and encouragement of discussion of the societal impact of future autonomous technologies. Positive publicity for all stakeholders also had to be assured as far as possible. Their requirements had to be balanced against the primary research goals of the project.

This paper presents an overview of Ada's design: her sensors and effectors, system architecture and behaviours. An overview of the operational results is also given, followed by a brief description of the data collection strategies.

2. Sensors, Effectors and Core Services

In total (including auxiliary exhibition areas), Ada has 15 video inputs, 367x3 tactile inputs, 9 audio input

Table 1: *Descriptions of functionality and types of software found at different levels in Ada*

Level	Functionality	Software
4: Behavioural modulation	Goal function evaluation, behaviour mode selection, emotional model	Simulated neurons (IQR421 [2])
3: Behavioural modules	Coordinated high-level interactions	Simulated neurons, software agents
2: Sensorimotor processes	Filtering of raw input data	Procedural or object-oriented code
1: Device I/O drivers	Interface to hardware	Procedural or object-oriented code
0: Hardware devices	Motor control, sound production, sensor reading, light setting	On-device logic

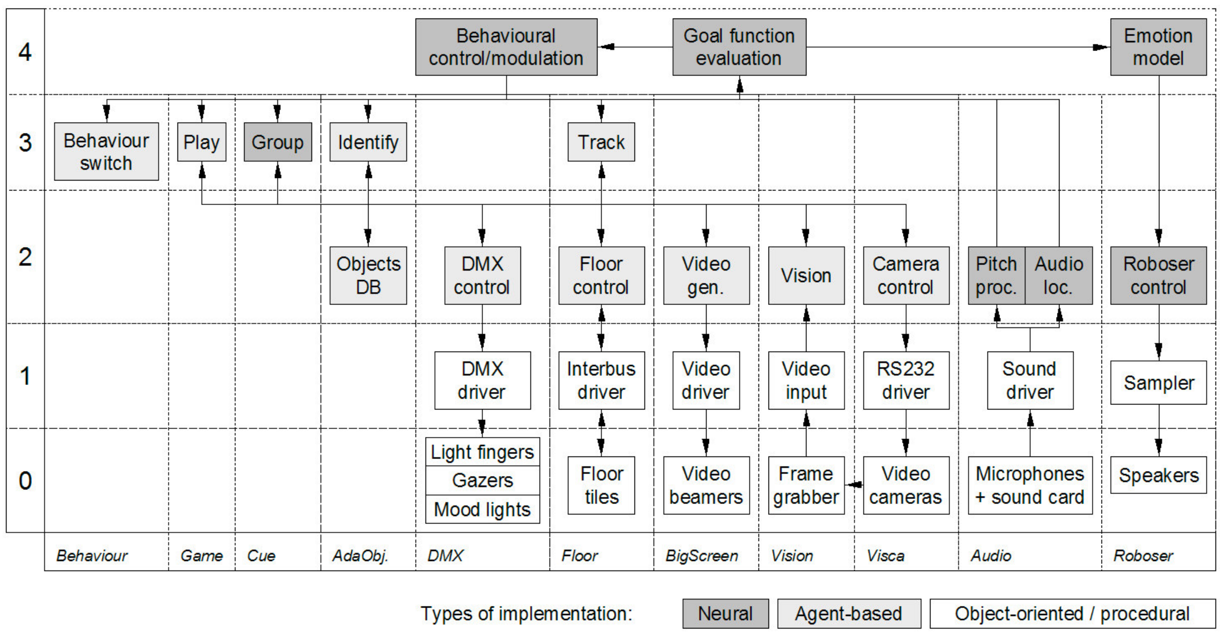

Figure 1: *Overview of Ada system architecture, organised into conceptual layers.*

channels, 46 mechanical degrees of freedom, 17 output audio channels, 367x3 floor tile lights, 30 ambient lights and 20 full-screen video outputs. All of these inputs and outputs can be addressed independently, giving a rich array of sensory modalities and output possibilities. Ada has the following sensory capabilities:

- **Vision:** Pan-tilt cameras called *gazers* are available to Ada for focused interactions with specific visitors. The cameras have zoom and digital filtering capabilities that are controlled on-line.

- **Hearing:** There are clusters of three fixed microphones each in the ceiling plane, with which Ada is able to localise sound sources by triangulation. Some basic forms of sound and word recognition and pitch extraction are available.

- **Touch:** Ada has a “skin” of 0.66 m wide hexagonal pressure-sensitive floor tiles [3] that can detect the presence of visitors by their weight. Each contains a microcontroller and sits on a serial bus running an industrial automation protocol called Interbus.

As well as sensing, Ada can also express herself and act upon her environment in the following ways:

- **Visual:** Ada uses a 360° ring of 12 LCD projectors to express her internal states and visitor interaction dynamics. These projectors collectively show a single, unified display of 3D objects covering multiple screens in real-time, as well as live video that can move with smooth transitions between screens. There is also a ring of ambient lights for setting the overall visual tone of the space. Local visual effects can be created using the red, green and blue neon lights in each floor tile in Ada’s skin.

- **Audio:** Ada is able to generate a wide range of sound effects. These sounds can be distributed across the entire space or localised using a matrix mixer. She expresses herself using sound and music composed in real-time on the basis of her internal states and sensory input. She can also change the pitch of her output depending on what she hears from her visitors. The composition is generated using a system called Roboser [13].

- **Touch:** Ada has twenty 16-bit pan-tilt *light fingers* for pointing at visitors or indicating different locations in the space. They are standard theatre lights on a serial bus called DMX, which is also used to control the ambient lights and the gazers.

The core services of Ada support her higher-level functions. A *tracking system* uses information from the floor tile pressure sensors to determine the location, speed, direction and weight of visitors. The limited resolution of the tiles means that it is not always possible to distinguish individual paths, so in some cases Ada only knows about the presence of groups of people at certain locations. To obtain more information about individual visitors, a *vision system* deploys gazers to collect images of people who have been localised on the floor. The *audio system* localises and recognises basic sounds (e.g. the word "Ada", pitch, note and key) to help identify salient individuals. On the output side, the *Roboser* system composes real-time music and sound effects, a *video server* allows the visualisation of saved and live images, and a *DMX server* controls the light fingers, gazers and ambient lights.

3. Behaviours and Interactions

The degree of success with which visitors can be convinced that Ada is an artificial organism depends on the nature of their interactions. The operation of the space needs to be coherent, real-time, and reliable enough to work for extended periods of time. As well as this, it must be understandable to visitors and sufficiently rich in the depth of interactions so that visitors feel the presence of a basic unitary intelligence.

To provide for a natural progression in visitor interaction, Ada incorporates at least four basic behavioural functions. First, she can *track* individual visitors or groups of visitors, possibly (but not necessarily) giving them an indication that they are being tracked. At the same time, she can *identify* those visitors who are more "interesting" than others because of their responsiveness to simple cues that Ada uses to probe their reactions. These people are encouraged to form a *group* in part of the space through the use of various light and sound cues. When the conditions are appropriate, Ada rewards a group of visitors by *playing* one of a number of games with them. She continuously evaluates the results of her actions and expresses emotional states accordingly, and tries to regulate the distribution and flow of visitors. These four behavioural functions are decomposed into smaller behaviours that call on the core services as needed.

4. System Architecture

Ada's architecture can be roughly sketched out as a series of levels (see Figure 1), with a gradient of decreasing biological plausibility as the proportion of traditional procedural code increases. The types of data stored or "learned" at each level also follows a similar gradient of biological plausibility. Each level contains modules that communicate with other modules in the same layer, as well as with modules in adjacent layers. The metaphor used is that of distributed brain-like computation, characterised by tight coupling within individual modules and loose coupling between modules. The underlying software is a mixture of simulated neural networks, agent-based systems and conventional procedural or object-oriented software. By far the largest amount of technical effort in the project was expended on developing this hybrid system and tuning it so that the different computing paradigms used could co-exist and co-operate usefully. The types of computations performed at each of the different levels are summarised in Table 1.

Different communication protocols are used to connect the components of the system: a socket-based protocol for the simulated neural networks, and an asynchronous message-based middleware for data transfer between agents. Data exchanged in this way includes floor data (input), visitor tracking data (internal), behavioural states (internal) and DMX device control (output).

5. System Goals and Action Selection

Conceptually, Ada is an artificial organism that tries to maximise its own goal functions, which we interpret as her "happiness". This means that the system as a whole must implicitly or explicitly compute its level of happiness, which can then be used to determine if certain actions contribute to this goal. As a first approximation we can write:

$$H = f(g_s, g_r, g_i)$$

H = overall goal or "happiness"
g_s = survival
g_r = recognition
g_i = interaction

Survival is a measure of how well Ada satisfies her basic requirements, which are to maintain a certain flow of visitors over time and to keep these people moving with a certain average speed. *Recognition* quantifies how well Ada has been able to track and collect data about people, as a pre-condition for more advanced interactions. This process can be seen as Ada "carving" objects out of the world of her sensory data, which is implemented as a progressive filtering of the sensory data and the creation of objects in an internal database once certain criteria of persistence and coherence have been satisfied. *Interaction* measures the number of successful human interactions that Ada has been involved in, with more complex interactions such as games being weighted more highly.

As a system, Ada has the goal of maximising the value of *H*. There are multiple strategies for achieving this: for example, Ada could encourage high visitor throughput, but in doing so have very few possibilities for recognition and interaction (g_s high, g_r and g_i low). Alternatively, Ada could also achieve an equivalent value of H with only a few visitors in the space, but with high recognition and interaction with each visitor (g_s low, g_r and g_i high). The actual computation of *H* occurs over multiple levels: an explicit top-level calculation is done using simulated neurons, and in parallel individual behaviours also calculate their own contributions to the parameters for *H*.

The results of the *H* calculation are combined with other high-level inputs and the state history to select the most appropriate behavioural state for Ada. Behaviour selection occurs at multiple levels – for example, the floor tiles display colours that depend on the local effects in use as well as the overall state of the space. At the top level, a neural modulation scheme activates and inhibits the underlying behaviours. This modulation can take a variety of forms, including a "hard" winner-take-all (WTA) scheme, a "soft" multiple-winner WTA, or a scheme where the behaviours run freely. The extent to which the behaviour modulation needs to be "hard" depends on the subjective evaluation of how the behaviours interact and/or interfere with each other. In practices, because of the constraints imposed by a high visitor flow rate, the behavioural control is run in a "hard" mode.

6. Computational Infrastructure

Ada runs on a 100 Mbit network of 31 PCs (AMD Athlon XP 1800+, 1.0 Gb RAM, Linux). Driver cards are used for DMX and Interbus communications. In addition, 40 frame grabbers, 4 sound cards and 11 dual-headed accelerated graphics cards are installed. Laptops on a wireless LAN enable system testing and tuning to occur while walking around in the main space.

7. Operational Issues

Due to the extremely large number of visitors that wanted to see Ada, it was necessary to control their flow very rigidly to avoid problems with overcrowding. This was necessary both for safety reasons and to ensure that each visitor had a certain minimum amount of space with which to interact with Ada. Table 2 summarises a typical visitor experience in the space and Figure 2 shows a typical scene in the main space.

During normal operation, the main space received about 25 visitors at a time, giving a nominal capacity of about 300 visitors per hour and an instantaneous occupancy of 125 visitors at any one time. The tracking system worked as reliably for wheelchairs and children weighing more than about 20 kg as it did for adults.

Figure 2: *A typical live user interaction scene within Ada. Visible are floor tiles, a visitor being highlighted by a light finger (centre left – the light finger itself is in the ceiling frame), a dynamic 3D visualisation (top) and a live gazer video on the screens (top left).*

Table 2: *Typical visitor experience in Ada exhibit*

Region	Visitor experience	Time (min)
Queue (outside)	"Brainworkers" neuroscience video on big screen (10 min)	0-90
Condition-ing tunnel	Sequential introduction to individual sensors and effectors	5
Voyeur corridor	View group of visitors interacting with Ada, listen to guide's explanation	5
Main space	Interact with Ada	5
Brainarium	View "control room" screens and look back into main space	5
Explana-torium	Art by H R Giger depicting future technology, guest book, videos with statements by scientists, credits poster	5

Ada ran for over 1700 hours on 159 consecutive days with an uptime of better than 98.3%, where uptime was defined as having a system that functioned well enough to enable a normal flow of visitors through the exhibit. Discounting outages due to deficiencies in building services that were beyond the control of the project team, the overall system uptime was over 99.1%. Nine stable versions of the Ada software were released during the Expo, incorporating incremental improvements in user functionality and data logging facilities. On any given day, either the latest development version or the stable version could be run, depending on the demands for testing and experimentation.

The public reaction to the exhibit was overwhelmingly positive. Surveys conducted by the Expo.02 organisation indicated that Ada was one of the 5 most popular attractions out of over 60 at the Expo. An online poll [8] also indicated that Ada was the most popular of the IT-related exhibits at the Expo. The visitor queue length of >30 minutes outside the exhibit for the entire duration of the Expo indicated that the final attendance of 553,700 could have been even higher if the capacity of the space had been larger.

The satisfactory operational result was partly the result of accumulated experience during the development process. From 1998 onwards, ten increasingly large public tests were run to evaluate the feasibility and scalability of the underlying technologies, gauge visitor impressions, and test different interaction scenarios. The two key issues that stood out from the results of the tests were the need for effective visitor flow control, and the importance of communicating Ada's intentions clearly through the use of effective cues and visitor pre-conditioning sequences. One direct consequence of this experience was the decision to employ guides to actively inform Ada's visitors as much as possible about what they would see in the exhibit.

8. Data Logging and Experiments

Data can be logged at several different parts of the system simultaneously. A brief description of the types of data logged is given in Table 3.

Table 3: *Types of data logged in Ada*

Description	Type
Floor raw load data: sensor values and calibrated RGB neon output colour	ASCII Text
Floor server data: positions of all currently loaded floor tiles at every time step	ASCII Text
Tracking data: Onset, path and endpoint of all visitor tracks on the floor	ASCII Text
Neural network internal states for behaviour selection	ASCII Text
Camera view (up to 4 gazers or overhead cameras simultaneously)	DV
Gazer view of tracked visitors	MPEG
Visitor questionnaires gauging their responses to their experience in Ada	ASCII Text
Ada automatically generated music	MIDI
Sound in space	WAV

Data is accumulated in multiple locations at ~5 Gb per hour during active logging, not including DV or WAV data. The data is automatically sent to a central repository each evening for backup on to DVD-R. A timeserver keeps all timestamps across the cluster synchronised to within a few tenths of a second, which is sufficient for most types of analysis.

As of this writing, some preliminary analysis of the ~150 Gb of collected data had been performed. Development, experiments and analysis have been focussing on the following areas:

- Verification of correct system operation with regard to its internal world-view representation
- Automatic calibration of gazers and multi-modal visual/tactile tracking
- Audio localisation and recognition in noisy spaces
- Assessment of visitor reactions to the exhibit based on demographic measures, and an investigation of the effects of various manipulations of the functionality of the space on visitor perception
- Automatic neural behaviour and action selection
- Automatic music generation based on an internal emotional model
- The ability of the space to actively affect the speed and position of visitors, and the effect of boundary conditions (entry/exit placement, pre-conditioning sequences) on their distribution in the space

9. Related Projects

Ada's closest relative from a project perspective was also her direct physical neighbour – the EPFL Robotics exhibit at the Expo.02. This project dealt with different technical content to Ada (autonomous cooperating museum guide robots), but both projects were of similar size and operated under similar conditions.

Several research projects deal with issues related to home automation and "intelligent rooms", and many companies offer commercial home automation systems such as the GE Smart series from GE Industrial Systems [7]. This system offers a substrate for connecting electrical devices and home network services with a common software interface. The control system software is based on rule sets or driven directly by end users, either within the building or via remote links. In this sort of system, the design emphasis is on ease of end-user installation, operation and customisation, rather than advanced behavioural functionality.

More advanced control systems exist in projects such as the Intelligent Room at MIT [10]. The Intelligent Room project aims to develop systems that support human activities in a seamless, flexible way. To date, work has been done on context-aware speech and gesture recognition, flexible resource allocation [6] and an agent-based extension to Java called MetaGlue. Ada has a similar set of functionalities, but with three main differences. Firstly, Ada is a completed product and is much larger than the Intelligent Room, in terms of physical size, number of components and degree of behavioural integration. Secondly, the design of the user interaction with the space is immersive rather than invisible – the building does not serve its users' needs in the background, but is an active participant in their experiences. Thirdly, the space actively tries to achieve its own goals by engaging its users.

A similar project, also named the Intelligent Space, is being pursued at the University of Tokyo [1]. The general approach is to design a platform to facilitate communication between the entities that inhabit it – whether they be humans, robots, or components of the space itself. The concept of a Distributed Intelligent Network Device (DIND) is proposed for connecting devices in the space. Each DIND has sensors, processing and communications components. In this way the space is seen not as an explicit entity like Ada, but as a common networking medium in a physical area. Another group at the University of Tokyo has developed a system for accumulating human behaviour in a small prototypical apartment [12] using mainly tactile sensors. Two noteworthy developments are a pressure-sensitive bed and a high-resolution pressure-sensitive floor [11] for use in the invalid care industry where 24-hour monitoring of patients is desirable.

An animal-like analogue to Ada is the *Mutant* dog robot [5] and its commercially available successor *Aibo* from Sony. Ada and Aibo are both complete systems designed to interact with the general public, and both integrate visual, audio and tactile information to produce behaviour. They both have an internal emotional model and layered system architectures: Aibo's architecture is agent-based, while Ada has a hybrid of simulated neural networks and agent-based software components. Sony has formalised its system architecture in the OPENR model for building robots [4]. The main differences between Aibo and Ada are the obvious ones of appearance and size. By looking like a dog, Aibo has an inherent advantage over Ada for human interactions. A decision made in designing Ada was to explore the limits of human interactions that could be supported without the use of pre-existing metaphors, and to discourage visitors from anthropomorphising the system. On the engineering front, Aibo has the dual challenges of miniaturisation and minimising power consumption, whereas Ada faces power consumption constraints on a much larger scale.

10. Outlook

Ada is one of the largest-scale artificial organisms yet created. During five months of successful active operation, Ada successfully entertained over half a million visitors, while also serving as a platform for research into several different topics. While definite plans were not available at the time of writing, it is expected that Ada will be reassembled in a new location to allow an ongoing mix of research and public performance.

11. Acknowledgments

Ada is supported by: ETH/University Zurich, Expo.02, Manor AG, Velux Stiftung and Gebert Rüf Stiftung.

12. References

[1] Appenzeller, G, Lee, J-H and Hashimoto, H, "Building Topological Maps by Looking at People: An Example of Cooperation between Intelligent Spaces and Robots", *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 1997)*

[2] Bernardet, U, Blanchard, M and Verschure, PFMJ, "IQR: a distributed system for real-time real-world neuronal simulation", *Neurocomputing*, 2002, 44-46: 1043-1048

[3] Delbrück, T, Douglas, R J, Marchal, P, Verschure, P and Whatley, A M, "A device for controlling a physical system", EU patent application 99120136.9-2215 (1999)

[4] Fujita, M and Kageyama, K, "An Open Architecture for Robot Entertainment" (1997), *Proceedings of the First International Conference on Autonomous Agents*, pp 435-442, ACM Press

[5] Fujita, M and Kitano, H, "Development of an Autonomous Quadruped Robot for Robot Entertainment" (1998), *Autonomous Robots* 5: 7-20

[6] Gajos, K, Weisman, L and Shrobe, H, "Design Principles For Resource Management For Intelligent Spaces", *Proceedings of the Second International Workshop on Self-Adaptive Software (IWSAS'01)*

[7] GE Industrial Systems, GE Smart web page, www.ge-smart.com

[8] Infoweek online poll, Sep 2002, www.infoweek.ch

[9] Lee, J-H & Hashimoto, H; "Intelligent Space – Its concept and contents" (2002), *Advanced Robotics* Vol. 16 No. 3: 265-280

[10] MIT Intelligent Room project web page, www.ai.mit.edu/projects/iroom

[11] Morishita, H, Fukui, R, Sato, T, "High Resolution Pressure Sensor Distributed Floor for Future Human-Robot Symbiosis Environments", *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 2002)*

[12] Noguchi, H, Mori, T, Sato, T, "Construction of Accumulation System for Human Behavior Information in Room", *Proceedings of the International Conference on Intelligent Robots and Systems (IROS 2002)*

[13] Wasserman, K C, Blanchard, M, Bernadet, U, Manzolli, J M and Verschure, P F M J (2000), "Roboser: An Autonomous Interactive Composition System", In: I Zannos (Ed.), *Proceedings of the International Computer Music Conference (ICMC),* pp. 531-534. San Francisco, CA, USA: The International Computer Music Association.