Presented at the 5th International Workshop on Artificial Intelligence Techniques for Ambient Intelligence (AITAmI'10), Kuala Lumpur, Malaysia, 18th July 2010

Quantum Computing: Non-deterministic controllers for Artificial Intelligent Agents

Alexander BANNIKOV[a], Simon EGERTON[1,b], Vic CALLAGHAN[b], Brian DAVID JOHNSON[c], Muneeb SHAUKAT[a]
[a]University Of Essex,Colchester, United Kingdom
[b] Monash University. Kulua Lumpur, Malaysia
Jalan Lagoon Selatan, Bandar Sunway, 46150, Selangor Darul Ehsan, Malaysia
[c] Intel Corparation, Portland, USA
abanni@essex.ac.uk, Simon.Egerton@infotech.monash.edu.my, vic@essex.ac.uk, brian.david.johnson@intel.com, mshaukat@essex.ac.uk

Abstract. This paper describes work that seeks to take the first steps towards demonstrating the application of Sci-Fi prototyping methodology to developing futuristic products. Science-Fiction Prototyping is based on an iterative interaction between science fiction and fact. We begin the paper by explaining the science-fiction prototyping process. We then identify a futuristic product to act as an exemplar of this process (free willed domestic robots) before describing virtual-reality systems and how they can provide a suitable visualisation environment for conceptualizing products. Next we identify a paper that provides a promising approach to emulating intelligent control and free-will, quantum computing. We then provide an introduction to quantum computing and argue the need for a special quantum development environment before presenting an architecture for the quantum development tools and the robot. Finally we present an evaluation methodology before summarising our work. We are in the process of implementing this system and we hope that by the time the workshop occurs, we will be able to report some initial results.

Keywords. autonomous agents, intelligent environments, quantum controllers, robotics, programming toolkits, prototyping, creative science, free-will.

Introduction

The nature and level of intelligence of embedded agents is a fundamental issue in the development of future intelligent environments. To date, the best agents for controlling intelligent environments are people. Therefore one line of research seeks to understand and mimic the intelligence of people better. At the heart of human intelligence is the brain but this is still largely a black box, with several theories on how it may work. It is clear that people are highly individualistic, capable of seemingly independent thought and decision making. At times people are highly logical, even somewhat mechanistic, but at other times people can seem irrational or emotional is decision making. Assuming evolution hones our intelligent mechanisms, then it has sculptured both our irrational and rational sides to produce an intelligence that seemingly works well in controlling natural living environments. Thus, understanding these mechanisms and how to incorporate any beneficial aspects into intelligent environment agents is potentially useful to designing agents that interface with people and control natural environments. One of the aspects of human minds that have always puzzled agent developers is what is the role and root of free-will in human intelligence and personality? Free will is a tricky thing; for people it can be argued to be at the centre of what makes us human. It is fundamental for us to function in our cultures, societies and governments. We are interested in these questions as we are interested in designing better agents. In earlier work [5] we formulated these questions and in this paper we layout the first steps towards building an experimental system to take this work forward.

1. Sci-fi Prototyping

According to Julian Blecker "*Productively confusing science fact and science fiction may be the only way for the science of fact to reach beyond itself and achieve more than incremental forms of innovation*".

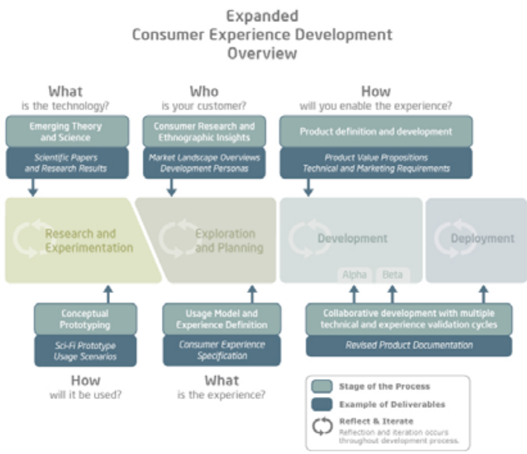

Figure 1. SF prototyping tool applied to a long term development process

Science-fiction and science-fact have a long history together. It is no secret that generations of scientists have been inspired to go into science because of the awe and wonder they experienced reading science fiction. There is a long list of inventions that first saw the light of day in the pages of science fiction and were later brought into reality. The web site technovelgy.com has captured nearly 2000 examples like Arthur C. Clarke's communication satellite and Star Trek's communicator. Conversely, science fact has done its fair share of influencing fiction. There is an entire subgenre of SF known as "Hard SF" that only concerns itself with scientifically plausible futures.

The practice of science fiction (SF) prototyping capitalizes on this synergistic relationship by creating SF stories based specifically on current research and development taking place in universities and commercial labs [9]. By placing the technology that is being developed into a fictional world and exploring its implications, both positive and negative, the scientists and research team gain a wholly new perspective on their work [2]. Understanding the broader view of a technology has been seen as an essential component in the development process. This process is described in figure 1 and illustrates how product design starts as concepts on the left of the diagram, moving to a commercial product on the right of the diagram. During its passage from left to right, it is subject of various influences and process shown on the figure.

2. Motivating Scenarios

By way of a set of scenarios we developed a series of robot stories; the Dr. Simon Egerton stories. These have taken current scientific writings and used a fictional world to examine various implications of the theory as well as the situations they might bring about. It was science fiction based science fact; specifically emerging scientific theory from computer science, robotics and neuroscience.

The scientific theories at play in *Brain Machines* come from two recent works. The first is a chapter from Michael Brooks' exceptional book 13 Things That Don't Make Sense. Chapter eleven is entitled, Free Will – Your decisions are not your own. In it Brooks does a brisk work of moving though a history of free will experimentation and the latest advances in neuroscience research. Ultimately he shows that science is proving that humans really don't have free will but that "for all practical proposes, it makes sense to retain the illusion. Human consciousness, our sense of self and intention, may be nothing more than a by-product of being enormously complex machines that are our big-brained bodies, but it is a useful one, enabling us to deal with a complex environment" [1].

The second work is a paper from Italian astrophysicist Paola A. Zizzi called I, Quantum Robot: Quantum Mind control on a Quantum Computer. In the paper Zizzi explores using quantum metathought and metalanguage as a way to control robots or computers that could become self aware. Simply put, metathought is "the mental process of thinking about our own thought…the process of thinking about thinking" [13]. Zizzi uses metalanguage to keep a robot from attaining free will. "With opportune boundary conditions, an apparently self-aware quantum robot reaches a level of thought. In this case the robot can still be controlled by a metalanuage which prevents him to reach the level of metathought." [13]. The goal of Zizzi's theory is to keep a robot from attaining free will.

Finally the SF prototype also incorporates work by Simon Egerton, Victor Callaghan and Graham Clarke, "*Instability and Irrationality: Destructive and Constructive Services within Intelligent Environments*" which provides a quantum persona transfer scheme for intelligent artificial agents in the form of a robot.

The prototype we propose is based on one of the Egerton stories; Brian Machines. The Egerton robot science fiction prototypes really are experiments in extremes. They look for the worst case scenario, go right for the heart of the debate, and search for the nastiest problems that might arise in science and culture. Brain Machines examines the worst legal and psychological effect of a society coming to grips with the terror of non-determinism. It also puts forward the idea that free will in robots may have positive effects. If humans must retain a delusion that we have free will to survive in a complex environment, then why not apply the same principle to artificial intelligence? If thousands of years of human evolution have taught us anything it's that adaptability is crucial for survival. Why wouldn't a non-deterministic approach to robots and artificial intelligence not increase its chances of survival in a complex environment? More specifically, the challenge concerns the taking of decisions, that at one level are seemingly simple (and therefore implementable in a simulated robot) whereas at another level, when aggregated together are revealing as to the cognitive nature of the basic decision making mechanisms. At the highest level the test concerns executing simple a simple instruction, repeatedly; where a quantum based versus non-quantum robot is requested to bring a drink to the owner, which is not used (and at a higher-level, when taken collectively, is seemingly aimless) [9]. The lower-level test concerns the individual decisions that make a path-finding robots aggregated behaviour more or less intelligent; with quantum control, or without. As our work is at an early stage, in this paper we deal with the latter test, which is described in section (4).

3. The Evaluation Environment

3.1. The Virtual Environment

Before using any kind of expensive (or unavailable) devices, the algorithms, that may be used, should be tested in a virtual environment. It is not necessary to have the precise equivalent of a real world 'inside' a computer program. Only the objects and/or effects that really affect the processing protocols should be involved in the simulation. The results that are gained from running tests in a virtual environment would then be used for designing future products.

Virtual Reality and Augmented Reality environments have matured over the decades. Advances in computer graphics and network technologies have empowered the VR industry. The 'gaming industry' has contributed greatly to the development of such environments. The latest gaming engines try to mimic physical characteristics of the real world to provide a realistic gaming experience. Multiplayer capabilities allow user interaction and collaboration across the globe. Research projects are emerging which utilise the power of these gaming systems for social and educational causes [3]. However, even with these advances, VR environments are far from providing a realistic natural simulation. The sheer number of different physical interactions and variables in a real physical environment renders any notion of useable virtual emulations system, for now, as primitive.

Consider the example of virtual prototyping, which seeks to illustrate physical, or software product ideas in a simulated digital world. Rapid iterations on these virtual prototypes promises to provide a quick and low-cost method for potential customers to experience what a product is like, before it exists in reality. In general, existing virtual prototypes are either too complex for the end-user involvement (eg simulators for engineers`) or too simple to give any real experience (paper designs). Even complex prototypes fail to mimic the real world operation because of performance constraints. In addition, current tools and techniques are standalone in nature and do not provide support of for collaboration and community involvement, which is especially important is getting a wider variety of views on a product concept.

We believe that virtual environments can provide a solution for this problem and we propose to build a simulation based on RealXtend, [12], an open-source virtual world software development (C# and Python) derived from *OpenSim* and based on the code from *Second Life* allowing us to benefit from their realistic avatars and landscaping. To create the graphics for our world we use *Google SketchUp* editing suite, plus *Google 3D Warehouse*, a vast online repository of three-dimensional models created by people using *SketchUp*, most of which are also free to use [6]. Some examples of Google 3D Warehouse objects are shown in Figure 2.

Figure 2. Example of Google 3D Warehouse Android

3.2. The Real Environment

The robot we plan to use for our real world tests is shown in figure 3. It uses an ARM based processor with 8 infrared proximity sensors; these sensors are the primary sensory inputs to our robot controller.

Figure 3. The experimental Quantum Based Path Finding Robot (courtesy of Creative Science Systems)

4. Quantum Principles

Quantum theory has variously been described as “spooky” and “too random” (sic Einstein) and it is certainly not any easy theory to get to grips with, even its greatest advocates say of it “if you think you understand quantum mechanics, you don't understand quantum mechanics” (Richard Feynman) with others labelling the theory as fundamentally incomprehensible. At best the theory itself is controversial and at the very heart of this controversy is the notion of *entanglement*, or *superposition*, whereby objects within a system become entangled and measuring the state of one entangled object will instantaneously affect the state of the other object(s), even when the objects are at a distance. This phenomenon would seem to break the universal speed limit of light. Moreover, the state of the measured object cannot be predicted or established before the measurement, with all possible outcomes being as equally likely. [8] Probabilistic *interference* is the other controversial issue, where, under certain circumstances, quantum events are said to have negative probabilities, leading to uncertainty through probability cancelation. Unlike classical probability where the results of probability sums are always greater than or equal to the summed probabilities, quantum probabilities behave very differently [11]. Quantum state probabilities are described by complex numbers, for example, with

$$C_{P1} = 2 + 5i$$

$$C_{p2} = 1 - 2i$$

The probabilities are resolved by taking the modulus squared of both probabilities,

$$P(C_{P1}) = |\, C_{P1}|^2 = 2^2 + 5^2 = 29$$

$$P(C_{P2}) = |\, C_{P2}|^2 = 1^2 + (-2)^2 = 5$$

And the joint probabilities of both,

$$P(C_{P1} + C_{P2}) = |(2 + 5i) + (1 - 2i)|^2 = |(3 + 3i)|^2 = 3^3 + 3^2 = 18$$

Unlike classical probabilities the joint is not 29 + 5 as one would expect. The two probabilities have interfered with each other leading to cancelation.

Although quantum theory accurately predicts experimental observations, the debate continues. Philosophers have attributed these quantum effects to an absolute universal randomness (i.e. randomness not borne from measurement error) and argue these effects underpin and justify arguments for *free-will* and also provide the mechanism for knowledge generation, or creativity. Although quantum mechanics remains somewhat controversial, we would like to exploit two of its core theoretical properties of entanglement and probabilistic interference, properties which no other probabilistic system possesses. The hopeful expectation is that our quantum based control models will introduce noise into the system allowing beneficial emergent, or creative, behaviours to develop, quite apart from standard deterministic control models.

4.1. Quantum Algorithms

Quantum systems are constructed from quantum bits, or *qbits,* and can be grouped to for quantum registers, or *qregisters.* Quantum bits in superposition simultaneously exist in all possible states. In conventional logic, an N-bit register of base 2 will have 2^N distinct states; whereas an N-bit binary qregister will have 2^N states in superposition i.e. all 2^N states will exist simultaneously [4]. The Grover algorithm (developed by Grover in 1996) [7] is able to manipulate qregisters without collapsing the superpositions and can effectively search through the 2^N superpositions to locate an optimal solution state in $O(\sqrt{N})$ time. This is a significant advantage over classical linear search methods with O(N) times [11]. We plan to exploit the Grover algorithm to demonstrate how a quantum based robot controller can optimally solve path finding problems. These experiments will be the first of a set leading to our implementation of our multiple-persona controller model, as detailed in our earlier paper [5].

5. Quantum Programming and Development Toolkit

Programming languages are being developed so programmers can make use of the quantum hardware being developed in research labs around the world. For those without access to quantum hardware, a number of software emulations have emerged in recent years, simulating quantum effects in Hilbert Space [10].

We aim to collect together existing quantum software libraries and integrate them into a full quantum development toolkit (QDT). Our Quantum development environment will allow users to directly implement quantum controllers, or input standard deterministic controllers based on classical or fuzzy logic and experiment to "quamtum-fying" the system and experiment with potential new emergent behaviours. Our conceptual outline for the quantum controller tool kit is diagrammed in figure4.

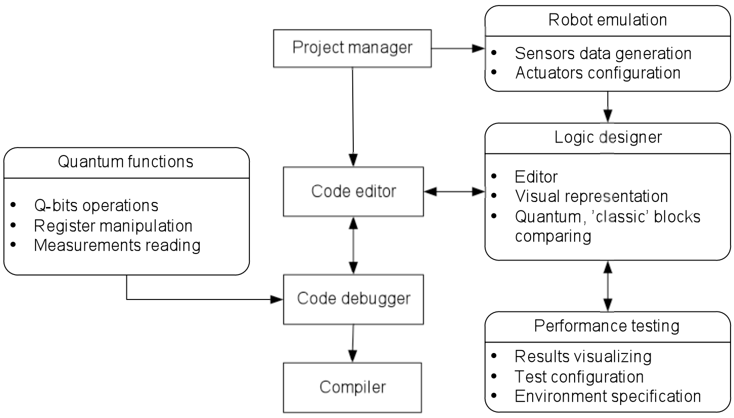

Figure 4. Block Diagram of Quantum Controller Toolkit

6. Evaluation and Benchmark Framework

To evaluate the quantum controllers developed with our quantum developer toolkit we have chosen to build a small simulated mobile robot that will aim solve path finding challenges. The general architecture for the robot is illustrated in figure 5.

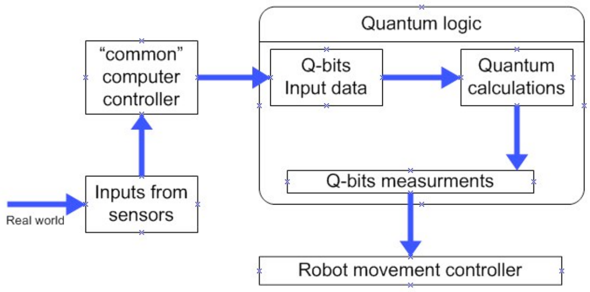

Figure 5 Architectural model of the quantum controller for the path finding robot

When defining a benchmark, a clear vision of required outcomes and internal states should be presented. But when talking about the quantum mechanisms, the certainty evolves into a bit different statement: the results must lead to some actions that are required for achieving the better (or different) outcomes along with measuring the probabilities of not getting that result. All the calculation must be processed as fast as possible.

Several testing methodologies are used: a black box testing for 'high-level' behaviour of a system in general terms. That will demonstrate whether a project solves the issues and gets the correct result, which was required in the beginning of the test (or simulation).

The more in-depth view of the controller is required in order to illustrate its working process. There is an issue that, in real life, quantum calculations of the internal state may be measured only after the result is found, otherwise the whole chain of manipulations may be lost. To solve this problem a simulation is designed in which similar algorithmic calculations take place, the only difference is that, when making measurements, the q-bits data it is not lost and execution of a command can proceed.

This 'white-box' test would demonstrate the performance of the quantum controller itself (when the controller has several input signals and, as a result, gives some output commands with respect to quantum calculations). Also, it may determine probable issues with combining the 'classical computational' blocks with quantum ones.

To illustrate the operation of the quantum controller we have designed a benchmark that will operate at two levels; high-level robot behaviour controller performance. The high-level behaviour is based on path-finding within an unknown environment for a mobile robot whereas the low level performance is based on comparing a quantum based controller to a regular PID controller. These are described in the following section.

6.1. A Path Finding Benchmark

Figure 6 illustrates the simulated environment, where a virtual robot with a quantum controller performs a path-finding solution to approach a destination area in a most efficient manner.

In that case, external information is received by the robot's sensors (all the environment and sensor data are simulated) and the general goal is pre-set into the system (eg approach a certain coordinate or/and a virtual marker).

A robot has several degrees of freedom – it may move forward, backward and turn left or right depending on its internal quantum controller.

When using a classical controller, a robot may archive only a certain aim (like avoid an obstacle, using 5 meters "detection" distance) and every time the behaviour will be same, but using quantum technique the next time an experiment is done, the behaviour may be slightly different (and the probability of that difference may be measured), so it will let a robot sensors as well time change the velocity, angle or rotation for the current situation or even "ignore" the boundaries of the collision avoidance which, as a result, may result in more (or less) efficient trajectory movement.

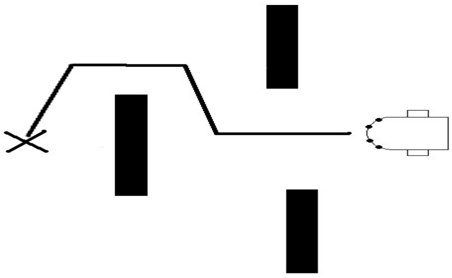

Figure 6. Path-finding task using emulator.

6.2. A PID Verses Quantum Controller Benchmark

At a lower level of evaluation we are benchmarking a Quantum controller to a regular PID controller. For example, if a robot receives data from a number sensors with a PID controller (ie without the quantum controller) its operation will follow a number of rules or equations giving the robot a deterministic behaviour in response to any input. However, when using quantum mechanisms, the certainty of the next action of a robot cannot be predicted with the same certainty. To quantify this facet of quantum behaviour its probability may be measured and the emergent behaviour recorded. Thus, to compare the approaches, the behaviours (for the same conditions) are recorded and the differences measured.

7. Summary

In this conceptual paper we have described how we are extending our earlier work on science fiction prototyping. For this, we have chose to investigate the how a quantum computer based domestic robot, operating with in an intelligent environment, might be designed and visualized in a virtual world. The science fiction prototyping elements arose from research first published in IE08, which proposed that the benefits (or not) of mirroring non-deterministic aspects of human behaviour together with offering a more efficient means of processing might usefully be investigated by replacing fuzzy controllers with quantum versions. This science was then fictionalized and extended by a series of stories by Brian Johnson (the Egerton stories), which was in turn used as the template for the conceptual scientific ideas presented in this paper. The motivation for doing this arose from two areas first, the need to provide a case study of Sci-Fi prototyping and to explore the application of quantum computing to robotic. However, to advance this work and build a quantum controller, we first need a quantum development toolkit (QDT). This short paper describes our work-in-progress to these ends. Thus in this paper we have introduced Sci-Fi prototyping, described virtual and real test-beds, reviewed the state of the art in quantum development software, proposed an architecture for both the quantum robots and the QDT. The originality of this work lies in both the nature of the proposed development system (there are no other integrated virtual/quantum toolkit architecture for robotics), the quantum robot agent benchmarks and the quantum control architecture. Clearly this work is at an early phase (and is ongoing) but we hope that it will prove to be a thought-provoking topic for the workshop audience and we look forward to reporting on the use of these concepts in later conferences.

References

[1] M. Brook, “13 Things That Don’t Make Sense”, Vintage, August 11, 2009.

[2] V. Callaghan, “*Tales From a Pod*”, Creative Science 2010 (CS’10), Kuala Lumpur, Malaysia, 19th July, 2010

[3] M. Davies, V. Callaghan, L. Shen, "Modelling Pervasive Environments Using Bespoke & Commercial Game-Based Simulators" 2nd International Conference on Life System Modelling and Simulation (LSMS'07) Shanghai, China, September 14-17 2007

[4] D. Dong, C. Chen, C. Zhang and Z. Chen, “Quantum robot: structure, algorithms and applications,” Hefei, Anhui, University of Science and Technology of China, Robotica , vol. 24 issue 4, pp. 513-521, July 2006.

[5] S. Egerton, V. Zamudio, V. Callaghan and G. Clarke “Instability and Irrationality: Destructive and Constructive Services within Intelligent Environments”. Intelligent Environment 09, Barcelona, Spain, 20-21 July 2009

[6] Google, 3D Warehouse, http://sketchup.google.com/3dwarehouse, Retrieved: 22nd February 2010.

[7] L.K. Grover, “A fast quantum mechanical algorithm for database search. Proceedings of the 28th ACM Symposium on Theory of Computing”, Philadelphia: 212, 1996

[8] D. Janzing, J.D. Decker, T, “How much is a quantum controller controlled by the controlled system?” vol. 19; Number 3, pp. 241-258, 2008.

[9] B. Johnson “Brain Machines” Intelligent Environment 09, Barcelona, Spain, 20-21 July 2009.

[10] Libquantum, http://www.libquantum.de/, Retrieved: 10th April 2010.

[11] M.A. Nielsen “Quantum Computation and Quantum Information” Cambridge University Press, September, 2000.

[12] RealXtend, RealXtend – Open source platform for interconnected virtual worlds, http://www.realxtend.org, Retrieved: 1st March 2010.

[13] P.A. Zizzi, “I, Quantum Robot: Quantum Mind Control on a Quantum Computer, Journal of Physics Conf. Ser., 67,. 012045, 2008.