

Massively Multi-Author Hybrid Artificial Intelligence

by John Pendlebury, Mark Humphrys and Ray Walshe

There is an emerging consensus in much of AI and cognitive science that “intelligence” is most likely the product of thousands of highly specialised subsystems collaborating in some kind of ‘Network of Mind’. In 2001, Mark Humphrys proposed that if Artificial Intelligence (AI) is to “scale up”, it will require a collaborative effort involving researchers from diverse disciplines, across multiple laboratories (<http://computing.dcu.ie/~humphrys/WWM/>). Until now there has never been an easy system to facilitate the construction of hybrid AI from the work of multiple laboratories. The World-Wide-Mind is the latest in a series of prototype systems, which enables the construction of hybrid AI systems from multiple laboratories.

The World-Wide-Mind server (<http://w2mind.computing.dcu.ie>) allows developers to pose software problems, (whether related to AI or not), for others to pose solutions to. Problems, such as a game of chess, or maze to be solved, are known as “worlds”. Solutions to these problems are known as “minds”. Both worlds and minds can be developed off-line and uploaded to the World-Wide-Mind server. As facilitated by many video hosting websites, such as Youtube, authors can upload their work and be assured that it will be hosted indefinitely.

Any web user can run a mind by selecting it from a list of minds displayed in a web browser. A new instance of the world and a new instance of the mind are created and run together, after which the world will assign a score to the mind.

This score can be used by mind authors to choose the most successful minds as components in their own hybrid minds, without the need to consult the original mind’s author, or install anything. It then becomes possible to create entire




Figure 1: Some worlds currently hosted on the World-Wide-Mind server, including a space simulation, a tennis game, a mining game and a battle simulation.

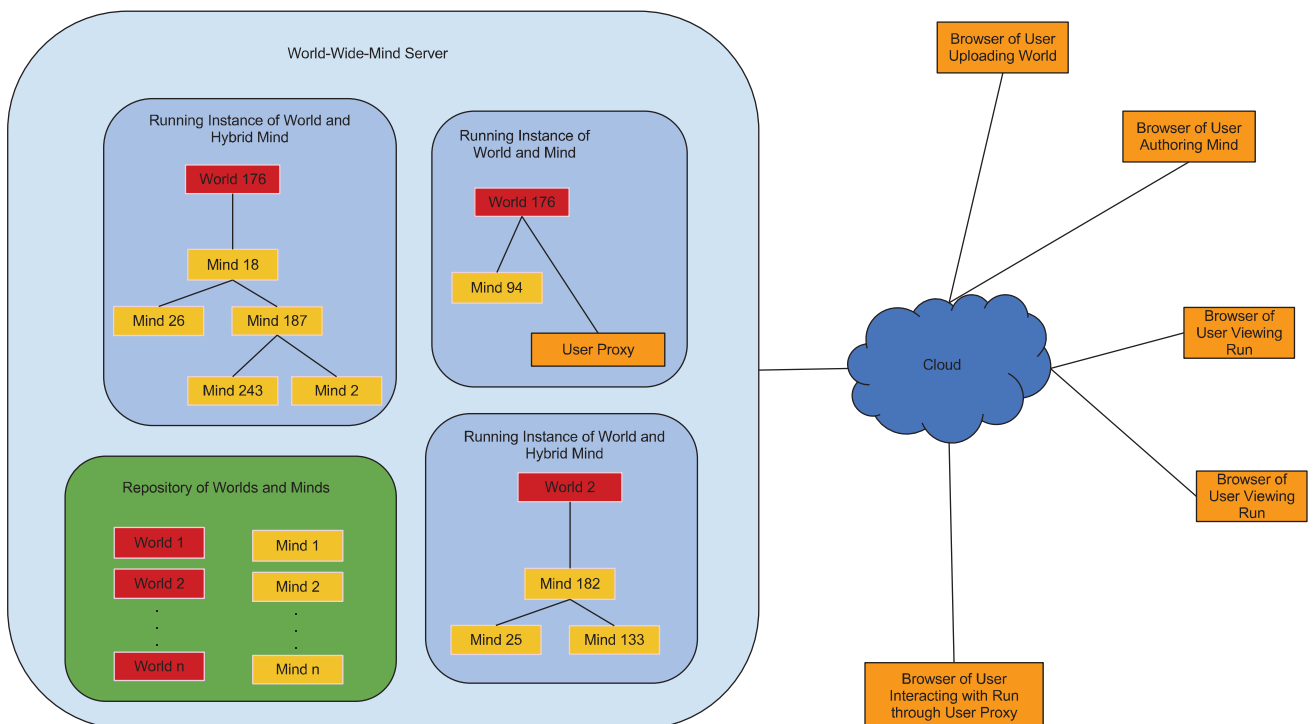


Figure 2: An Architecture for the World-Wide-Mind showing the World-Wide-Mind server running three instances of worlds with minds; two of which are running with hybrid minds and one running with an individual mind and a user proxy, allowing a user to interact with the world and mind.

hierarchies of minds with one mind at the top of the hierarchy arbitrating between the actions of minds below it, which might themselves be arbitrators of minds below them, and so on. During a run worlds can opt to output images. The system also has the facility to generate a video using these images.

In late 2011 funding was secured from the Irish Research Council for Science, Engineering & Technology (IRCSET) to enhance the World-Wide-Mind over three years. Current work in scaling up this platform is to move image generation to the client. This will allow the system to update the user's view of a run in real-time.

Distributed games frequently use artificial intelligent agents to enhance the user experience. As the user will be viewing the state of a world in real-time, there is no technical reason why they cannot interact with minds in real-time during a run. The advantage to our system would be to allow minds to learn from direct interaction with human agents. With this enhancement the system will resemble a generic MMOG (Massively Mutli-player Online Game) platform capable of running any game that a user may upload.

Many worlds and minds have been written for the World-Wide-Mind. A selection of the best of these can be found at <http://w2mind.computing.dcu.ie>. Some minds on the current system are individuals; others are hybrids consisting of several, or even dozens of individual minds. In the future, instead of developing solutions consisting of dozens of specialised minds, potentially we could develop solutions consisting of thousands, or hundreds of thousands of minds.

The World-Wide-Mind already resembles an ecology, where unsuccessful minds are ignored and successful minds are reused constantly. If this is indeed the case then what kinds of problems will these new hybrids be capable of solving? We hope that this system will harness an unexploited creativity for creating hybrid AIs that until now has been almost entirely dormant.

Links:

<http://computing.dcu.ie/~humphrys/WWM/>
<http://w2mind.computing.dcu.ie>
[http://computing.dcu.ie/~humphrys/wwmdev/
selected.worlds.html](http://computing.dcu.ie/~humphrys/wwmdev/selected.worlds.html)

Please contact:

John Pendlebury, School of Computing, Dublin City University, Ireland
Tel: +353 1 7005616
E-mail: jpendlebury@computing.dcu.ie

Bionic Packaging: A Promising Paradigm for Future Computing

by Patrick Ruch, Thomas Brunschwiler, Werner Escher, Stephan Paredes and Bruno Michel

The spectacular progress in the development of computers has been following Moore's Law for at least six decades. Now we are hitting a barrier. While neuromorphic computing architectures have been touted as a promising basis for low-power bio-inspired microprocessors down the road, imitating the packaging of mammalian brains is a new concept which may open new horizons independent of novel transistor technologies or non-Van Neumann architectures.

If one looks at the transistor count in microprocessors over the last four decades we have gone from 2500 to 2,500,000,000, a gain of six decades.

But today, two major roadblocks to further progress have arisen: power density and communication delays. The currently fastest computer was built by Fujitsu and is at the Riken Institute in Japan; it has a capacity of 8 petaflops and a power consumption of more than twelve megawatts --- enough for some ten thousand households. As to communication, while pure on-chip processing tasks can be completed in shorter and shorter times as technology improves, the transmission delays between processors and memory or other processors grow as a percentage of total time and severely limit overall task completion.

Energy considerations

A processor architecture that imitates the mammalian brain promises a revolution. Compared with the mammalian brain, today's computers are terribly inefficient and wasteful of energy, with the number of computing operations per unit energy of the best man-made machines being in the order of ~0.01% of the human brain depending on the workload. This inefficiency occurs not only in the processing chips themselves but also in the energy-hungry air conditioners that are needed to remove the heat generated by the processors.

In the October 2009, No. 79, issue of ERCIM News the Aquasar project was described. In Aquasar, the individual semiconductor chips are water cooled, with water being 4,000 times more effective than air in removing heat. Aquasar uses micro-fluidic channels in chip-mounted water coolers to transport the heat. This result is achieved with relatively hot coolant so that the recovered thermal energy is used for heating a building. Aquasar is installed and running at the ETH in Zurich; a successor machine with a three petaflop capacity is currently being installed in Munich. The integration of microchannels into the semiconductor chips themselves promises to sustain very high power dissipation as the industry strives to increase integration density further and also move to 3D chip stacks.

Going further in the new paradigm, another cause of energy inefficiency is the loss (in power and in space) in delivering