



---

## **Using DDT to debug applications on the IBM Cell BE processor**

August 2007

---

## Using DDT to debug applications on the IBM Cell BE processor

The Cell Architecture is the result of collaboration between IBM, Sony and Toshiba to design a high-performance and power-efficient processor that could drive applications in the fields as diverse as gaming, HDTV and supercomputing.<sup>1</sup>

The Cell is an innovative heterogeneous multiprocessor consisting of:

- The PPU – The Power Processing Unit containing an IBM 64-bit Power Architecture™ core, the so-called Power Processing Element (PPE).
- Eight SPUs – Specialized co-processor units each containing a Synergistic Processing Element (SPE).

with a coherent on-chip bus for communication between the elements.

Whilst the PPU has a familiar processor, with ordinarily at least 256Mb RAM (global memory) available, each SPU has a large register set and a local store of 256Kb. Access to the global memory from the SPUs is performed using DMA via the bus, or by exchanging messages with the PPU through a “mailbox” mechanism.

### *Programming the Cell*

Standard PowerPC programs will run unmodified on a Cell system such as the IBM QS20 BladeCenter or the Sony PS3™ which can both run the Fedora™ or Yellow Dog™ distributions of the Linux operating system – using just the PPU. However high-performance computing users of the Cell will be keen to exploit the capability of the SPUs fully. The choices for this fit broadly into three models:

- Transparently – using libraries that have been optimized for the Cell architecture. Many HPC codes use proprietary or open-source libraries for the computational kernel of the application, and these may have been ported to the Cell. Simple re-linking will enable the application to use the new

libraries.

- Intermediately – using advanced languages or compiler directives. A number of third parties provide compilers and/or libraries that can optimize the transfer of data and the computation between the SPUs and the PPU. This can involve rewriting the computational kernel of an application, or just adding compiler directives rather similar to parallelizing using OpenMP.<sup>2</sup>
- Directly – the developer writes applications for both the SPUs and the PPU and has responsibility for data transfer and synchronization. This is typically used when wanting to directly control the behaviour of the Cell, or to hand-optimize the performance of the computational kernel, or for code that does not fit well with the patterns supported by the advanced languages.

### *The Need to Debug*

Inevitably, where there is programming, there is debugging. This white paper introduces Allinea's DDT, a debugger for multi-threaded and parallel applications which has enhanced support for the Cell.

Like programming the Cell, the process of debugging can also differ from a user's previous experience. Historically trying to find bugs by placing “print” statements in code has been popular, but this has always led to a repeated cycle of modify-compile-run that can be rather slow to achieve results.

This method does not work well in multi-threaded environments, and this is particularly true with the Cell where a print issued by the SPE will lead to communication between the PPU and SPU to produce the output which can also change the behaviour of the bug.

Graphical debugging tools, like Allinea's DDT, represent a leap forward beyond the print solution. At runtime the user can control progress of a program, see all the values of its variables, its




The Parallel Stack View is another very popular feature for showing divergent behaviour across processors and threads. This window displays the call stacks of every thread – optionally including every thread of every process for a program with multiple Cells – in a single tree view.

Where threads share a common stack – this meaning they are in the same part of a program – then they will be on the same branch of the tree. The number of threads at each part of the tree is shown, which means it is simple to find the threads that are behaving differently from the norm just by looking at those branches containing only one thread.

## Controlling the PPU and SPUs

Across the top of DDT's main window, the SPU and PPU threads are shown – they are identified by the “PPU” and “SPU” suffixes underneath the thread number. By clicking on a thread to select it, this thread becomes current – which means that the variables being shown are those present on that particular thread.



DDT Showing SPU and PPU Threads

Breakpoints can be added to program to stop PPU and SPU threads when a particular point in the program is reached. DDT also allows breakpoints to be set that apply to all, or just one of the threads.

DDT can control all the threads, or just one at a time, this allows the user to focus on an individual PPU or SPU thread – stepping one line of code at a time whilst the other threads are paused. Following individual threads through a computation is one of the most useful capabilities of a debugger.

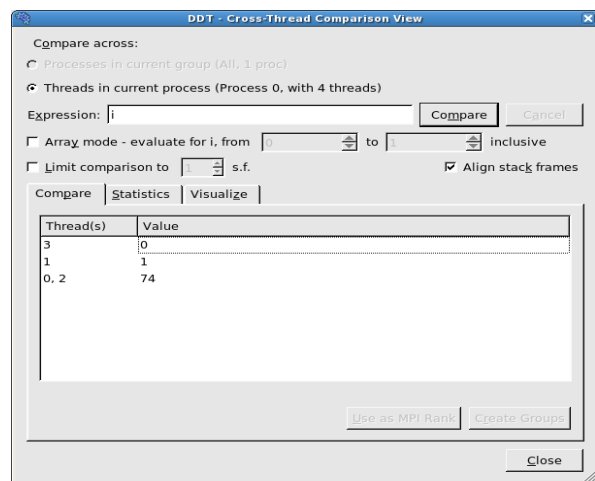
Often a Cell code will exhibit SIMD (Single Instruction Multiple Data) parallelism – each SPU thread executing the same part of a code at the same time. In these situations it can be instructive to step each thread ahead a line at a time – DDT has a special mode for this, selected by toggling the “step

threads together” button. When threads diverge, it can be the source of a problem.

## Examining Variables

Each SPU and PPU has its own addressable memory region. For the SPU, the local store holds variables – such as automatic stack variables, or global or static variables – and also local heap memory allocated using the malloc() system call within the SPU.


There are some traps to SPU programming – a global variable is only “global” on the same SPU – each SPU and PPU will have a different variable, unlike conventional multi-threaded programming where processes share the same memory.



Comparing Variables to Detect Bugs

DDT can help spot problems like this – it has tools such as the Cross Thread Comparison window which will examine the variables with the same identifier across each SPU and PPU thread. This is ideal for spotting rogue values, the common values are also grouped together so that the developer can really focus on the differences.

If an individual SPU or PPU thread is selected, the values of its variables are shown in the Local Variables tab. Dragging the current line markers in the source code will select the variables shown on each selected line, and evaluate these too in the Current Line tab. For user defined data types – structures, classes and derived types – the contents can be opened up by just clicking on the variables.

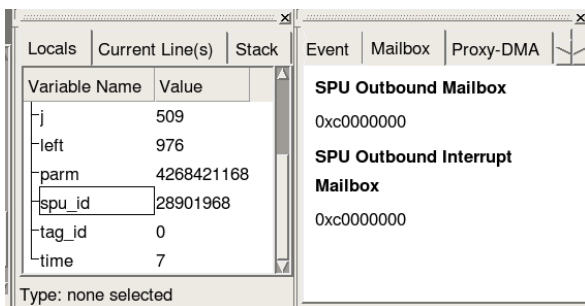


*Examining Local Variables with DDT*

Variables can also be dropped into the Evaluations box for more detailed probing. Pointers can be followed from the evaluation window, by right clicking and selecting “dereference”.

## Low-Level Cell Status

There are additional tabs in DDT that show the current status of DMA requests and mailbox events – this information is directly provided by the kernel and can show important data such as the current contents of the mailboxes that are used to communicate between the PPU and SPUs. SPUs and PPUs have, in total, three mailboxes – two for SPU



*Directly Displaying Mailbox Status*

to PPU communication (one interrupting the PPU, the other must be polled) and one for PPU to SPU communication (polled).

## Summary

The IBM Cell brings extensive benefits to demanding users in terms of performance, but to exploit these fully requires a new way of programming. This paper has shown how DDT can be used to simplify the transition to Cell programming by providing an easy to use and highly capable debugger which is able to debug complete Cell applications. Not only can DDT debug multiple

threads within a single Cell, DDT can also be used to debug clusters of Cells.

DDT for Cell is available from <http://www.allinea.com>. Supported platforms are the IBM QS20 Blade Center and Sony PS3™ running Fedora Core 6 and IBM Cell SDK 2.1.

## About Allinea Software

Allinea Software is a leading supplier of development tools for parallel and high performance computing. Allinea was founded by computer scientists from Warwick and Oxford Universities, giving the company unrivalled expertise in scientific and parallel computing and an insight into the challenges of exploiting parallelism as it enters the mainstream.

## References

1. The Cell Project at IBM Research (<http://www.research.ibm.com/cell>)
2. RapidMind (<http://www.rapidmind.net>) and Cell Superscalar (<http://www.bsc.es/cellsuperscalar>)
3. The IBM Cell Software Development Kit (<http://www.ibm.com/developerworks/power/cell/>)