


FPGA-based Multicore Microprocessor Emulation

Written by The Hitex Design Group

Hitex has enabled Infineon to expedite development of its brand-new AURIX multicore 32-bit microcontroller (MCU) architecture by developing a complete FPGA real-time microprocessor simulation platform, Meridian. Stretching existing FPGA technology to its limits, the Meridian platform has allowed Infineon to test its multicore IP months ahead of real silicon tape-out, enabling processor and toolchain bring-up well in advance of normal timelines.

This recently-announced new multicore architecture is the foundation of Infineon's next generation MCU family for automotive powertrain and safety applications. The multicore architecture features up to three processor cores to share the application load, includes lockstep cores and contains further enhanced safety mechanisms to support applications up to ISO 26262 ASIL-D.

As this is a major step in a fast-moving market, it was vital for Infineon to be able to prove as much as possible before real silicon was available. We thought it would be interesting for engineers to share our experiences in such an interesting challenge – from both hardware and software perspectives.



The challenge

Our brief was to design a board around the high-performance Virtex6 550T FPGA device, which has nearly 550,000 configurable gates and a 1760-ball BGA package. The board needed to add peripherals to the FPGA to extend its usefulness for evaluating the “soft” microprocessor, including an LCD screen, CAN transceivers, FlexRay interfacing, Ethernet and ADCs. The board also needed a “System Manager”, which would in itself be a powerful Infineon TC1797 TriCore device – a predecessor of the FPGA “soft” micro – to manage bring-up of the FPGA image. The system would then be used for development of the debuggers and flash programmers themselves.

From the software side, the System Manager would need to access multiple FPGA images stored locally on an SD card and upload an image directly using the FPGA's JTAG connection. It also has to offer the user the opportunity to select the FPGA image (before uploading) and set configuration options such as the processor speed that should be emulated.


If all that weren't enough, the System Manager should also emulate an Infineon CIC61508 Safety Monitor device – an intelligent watchdog that forms part of the Infineon PRO-SIL SafeTcore system in conjunction with existing TriCores.

So not much then!

The approach: hardware

Much of the design was fairly standard – interfacing for communications peripherals, multiple power supplies of varying voltages and outputs etc. The more interesting parts centred around choosing the correct FPGA banks of inputs/outputs (i/os) to allow high-speed signals and clocks to be routed correctly, together with handling multiple reset signals from multiple sources – always a headache for multiprocessor systems, let alone one where one device (the FPGA) has two operation modes, i.e. programming and running the emulation.

There was also the challenge of managing multiple signal voltage levels, at high speeds and sometimes bi-directional, for interfacing a 1.0V FPGA to 2.5V, 3.3V and 5V systems. Selection of bus translators was therefore crucial to the design.



An interesting and novel system was required to allow control of the FPGA for more unusual tasks, such as introducing disturbances into the grid array, starting/stopping the simulation etc. This meant the System Manager had to emulate the FPGA's JTAG signals as well as providing the clock to the simulated core. This allows all manner of new possibilities which aren't possible on normal target hardware, such as simulating drifting or jittery clocks, "flash" value changes etc. Essentially, it's possible to speed up, slow down or stop time!

One twist to the system is that it is somewhat of a schizophrenic – before programming, the FPGA must be handled in one way, with considerations of isolating clocks and signals from the i/o pins while in this state. However, once programmed, the device effectively becomes a multicore microprocessor, which means the board itself must be able to handle both "personas" transparently.

The approach: software

Hardware is nothing without software – all the more so in this case where there is virtual hardware. The System Manager software could be developed separately on a standard Infineon TC1797 TriBoard development kit, which allowed us to implement and test the SD card file system, CIC61508 Safety Monitor emulation and Virtex FPGA JTAG connection using a low-end FPGA system that was compatible.



The trickiest software to implement was the programming and JTAG interface to the FPGA, as this is normally handled in a debug/programming tool, or CPLD. Thanks to the processing power and real-time performance of the TC1797, relatively high clock rates could be achieved while retaining a good measure of flexibility.

The multicore device's software was, of course, impossible to write before the Meridian board was ready, since that is the very reason for its existence – to develop the software and test the hardware before the real silicon comes along. It was also only planned that basic drivers be developed at this stage, as the board was to be initially only used by Infineon for proving the silicon IP, and so higher level drivers could be developed later.


Putting theory into practice

All of the above makes good design sense... but what actually happened on the project? What would we have done differently? Well, like all clean-sheet projects of this complexity, pinning down the specification was key, as it would have been possible to “just add this or that” to both hardware and software, given the very high level of flexibility the platform offers. One aspect of hardware design which is never given enough attention is the power supply, and the highly-fluctuating power requirements of the multitude of on-board SRAMs caused problems for our power budget, resulting in a minor change for later boards. A golden rule – never scrimp on the power budget!

Minimising PCB layers is always sensible practice where applicable, but here the key was everything working and with such a high-value board, it was prudent to use multiple grounding layers. Likewise, matching track lengths and impedances was vital for the memory connections in particular. This kept the signals clean and separated where necessary.

The end result


The Meridian platform has allowed Infineon to write software drivers and application benchmarks to evaluate the architectural building blocks of a microprocessor – which is normally not feasible due to the modelling required to simulate the external environment. Simulation of the internal IP has been around for many years but requires immense computing power, with entire clusters of dedicated machines running for many hours to produce a result.



Using the Meridian board, it has been possible to supply the multicore IP to key partners who would not normally have the access to the requisite computing farms to simulate their specific use case. The early availability of the IP in an electrically representative environment has allowed real external sensors, actuators and communication channels to be integrated, thus reducing the simulation runtime from hours to real time, since the FPGA-based simulation will run at up to 30MHz. This is sufficient to allow accurate evaluation of external bus interfaces, communications and signal generation peripherals in a real environment. Multicore code is driving CAN networks, motors, LEDs and TFT screens before production silicon is available!

Debugging the chicken and the egg

Traditionally, it has only been possible to develop the debugger once the silicon was available, but without the debugger, it wasn't possible to test the debug interface! The Meridian microprocessor simulator platform allowed a new approach to debugging: where the debugging could previously only be carried out once the silicon was available, the Hitex HiTOP debugger was developed using the FPGA multicore simulation before arrival of the real silicon, so that Hitex could ensure that HiTOP could meet the challenge of debugging multiple processing cores on a single chip. For instance, in



a single core system there is normally only one entity to be debugged via one JTAG connection; the Meridian platform, however, also supports multicore configurations, so every JTAG command needs to

specify the target core as well, implying some routing and arbitration. The debugger needs to be able to connect to the specific core under test and differentiate between them.

These first versions of multicore compilers and debuggers, targeted at multicore applications, are now available to selected partners for evaluation.

Safety first

A vital part of Infineon's automotive microcontroller strategy is to incorporate advanced features needed for safety-critical systems, such as ECC on internal buses and RAM together with new IP blocks to monitor bus accesses, memory violations etc. The Meridian board has the ability to test these, and also includes a CIC61508 Safety Monitor intelligent watchdog device, so that the dedicated safety IP can be proven and developed (hardware and software) before the real silicon is available.

Copy protection

Obviously, with something as valuable as an entire microcontroller's IP, it was necessary to protect the FPGA image from being reverse-engineered. This means that every image, stored on a simple SD card, is encrypted using an AES 128-bit encryption key; a matching key is blown into the FPGA as a one-time programmable operation. Losing the key is not an option – it would render the £8000 FPGA useless for future images!

What next?

Hitex and Infineon are still exploring the benefits of the Meridian platform – it is already being used for EU-funded research projects into new safety-critical architectures and power control systems. The approach is complementary to other software based simulation tools, so provides Infineon with a generic platform to further tune and experiment with future architectures, all in pursuit of maximising performance and building architectures that are right first time.

Find out more

To find out more about our expertise, please contact The Design Team directly via the details at the bottom of the page.