

Fast iteration and prototyping in high-performance computing medical applications: a case study with Mentor Vista

8th INFIERI WORKSHOP

Alberto Garcia-Fernandez

10/21/2016



The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement No. [317446] INFIERI "Intelligent Fast Interconnected and Efficient Devices for Frontier Exploitation in Research and Industry".



Thales R&T HPC lab topics.

High-performance computing architectures.

- Finding the right architecture for a given application domain.

Tooled-up flows for high-performance cyber-physical systems.

- Efficiently programming adaptive architectures.


Mission.

- Evaluate available vendor tools.
- Show the gap between our needs and the tools capabilities.
- Propose, develop and evaluate solutions to bridge this gap

Objectives


1. Use of tools for fast exploration on new heterogeneous architectures for high performance computing applications

- Fast performance evaluation of new chips and platforms.
- Architecture-adequacy discovery for domain representative algorithms.
 - Image processing.
 - Pixel tracking.
 - Big data.



2. Architecture design and prototyping on low-power image processing platforms

- Improve performance per watt: peak by 4x and average by 10x by the end of 2018.
 - 100x and 200x by 2023.
- Medical X-Ray Imaging.
- Reduce the level of radiation by 75%.
 - More powerful image processing.
 - The device needs to be small and mobile.
 - Processing close to the sensor to help reduce extraneous wires.






Case study

Image processing kernel chain

- Reference implementation in sequential C code.
- Requirements.
 - **Must:** 1024x1024 , 25 fps 420 Mbits/sec
 - **Want:** 1344x1344, 30 fps 870 Mbits/sec

Different data access patterns.

- Pixel by pixel.
 - Good spatial locality.
 - Poor temporal locality.
- Sliding window.
 - Poor spatial locality.
 - Good temporal locality.
- Chessboard.
 - Poor spatial locality.
 - Poor temporal locality.



Tools available

- HW/SW co-designers.
 - Integration of compiler and synthesizer.
 - Reduced set of architectures for exploration.



➤ Virtual prototyping.

- HW simulators or emulators.
- Design the software and test it on a VP.



➤ SW parallelizers.

- Adapt SW for heterogeneous devices.
- OpenMP, OpenCL, CUDA.
- No custom custom HW (i.e. FPGA).



Mentor Vista

- SystemC TLM integrated development environment.
- Architecture design exploration.
- Virtual prototyping.
- Model the hardware before register-transfer level is available.

Sourcery Codebench

- C/C++ integrated development environment.
- Development for embedded devices.
 - ARM, IA32, MIPS, Power.
- Possibility to link with Mentor Vista.
 - Develop the SW before having the HW.



Development in three use cases

Model the hardware.

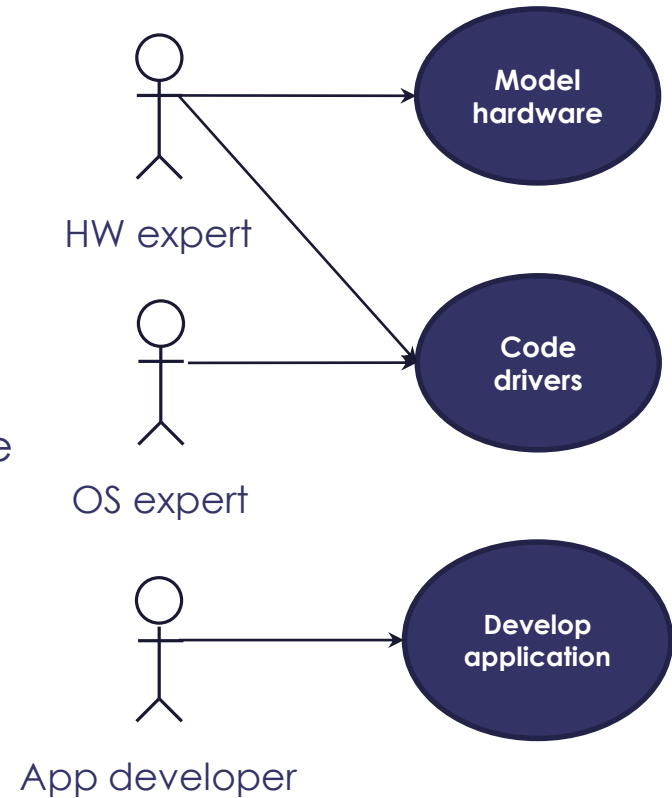
- Model IPs within FPGA fabric.
- Provide a bus mapping.
- Hardware knowledge required.

Code the kernel drivers.

- Implement the drivers of the fabric IPs in the Linux device tree.
- Hardware and OS knowledge required.

Develop the application.

- Adapt the reference application to the model of computation required.
- Offload computation to the FPGA IPs.
- Deploy the application in the prototype.
- Software development knowledge required.



Modeling the hardware

SystemC and TLM for the hardware IP models.

Easy and straightforward to develop.

➤ Accuracy depending on the level of detail.

Vista for the IP modeling:

- HW registers,
- Memory addresses.
- IRQs.
- Energy consumptions.
- Physical bus connections.

Define the bus mapping.

➤ Design the memory layout of the FPGA.

TLM Library: fpga_models Type: Bus Model Name: FPGA_AXI_BUS

Name	Protocol	Side	Width (bits)	Bind Mode	Reset value	Vector size
FPGAmaster	/TLM/protocols/axi	master	128	mandatory		
FPGAslave	/TLM/protocols/axi	slave	128	mandatory		
ImgExregisters	/TLM/protocols/axi	master	128	mandatory		
ImgExdma	/TLM/protocols/axi	slave	128	mandatory		
FrameBuffer	/TLM/protocols/axi	master	128	mandatory		
FPGAmem	/TLM/protocols/axi	master	128	mandatory		
Ker1registers	/TLM/protocols/axi	master	128	mandatory		
Ker1dma	/TLM/protocols/axi	slave	128	mandatory		
Ker2reg						
Ker2dm						
Ker3reg						
Ker3dm						
Ker4reg						
Ker4dm						

```
$fpga_top.fpga_axi_bus.ImgExregisters_base_address = 0x00000000
$fpga_top.fpga_axi_bus.ImgExregisters_size       = 0x00001000

$fpga_top.fpga_axi_bus.FrameBuffer_base_address = 0x00001000
$fpga_top.fpga_axi_bus.FrameBuffer_size        = 0x10000000

$fpga_top.fpga_axi_bus.Ker1registers_base_address = 0x10001000
$fpga_top.fpga_axi_bus.Ker1registers_size       = 0x00001000

$fpga_top.fpga_axi_bus.Ker2registers_base_address = 0x10002000
$fpga_top.fpga_axi_bus.Ker2registers_size       = 0x00001000

$fpga_top.fpga_axi_bus.Ker3registers_base_address = 0x10003000
$fpga_top.fpga_axi_bus.Ker3registers_size       = 0x00001000

$fpga_top.fpga_axi_bus.Ker4registers_base_address = 0x10004000
$fpga_top.fpga_axi_bus.Ker4registers_size       = 0x00001000

$fpga_top.fpga_axi_bus.FPGAmem_base_address     = 0x20000000
$fpga_top.fpga_axi_bus.FPGAmem_size            = 0x10000000
```

OPEN

Coding the kernel drivers

UIO-based approach used to provide the interface to the FPGA fabric IPs.

- Gap between HW modelling and SW application development.

Modify the linux device tree and rebuild the root.

- Not covered by the tools.
- Not a trivial task.
- Add one node per IP.

Forward the memory, bus, and IRQ mapping defined in the previous stage.

- Calculate corresponding memory addresses and IRQs.
- Documentation provided by Mentor.

Possible UIO driver modification required.

```
imgext@40000000 {
    interrupt-parent = <&ps7_scugic_0>;
    compatible = "fpgafabric";
    status = "okay";
    reg = <0x40000000 0x00001000>;
    interrupts = <0 29 4>;
};

framebuffer@40001000 {
    compatible = "simple-framebuffer";
    reg = <0x40001000 0x10000000>;
    width = <1024>;
    height = <1024>;
    stride = <(1024 * 4)>;
    format = "a8r8g8b8";
};

ker1@50001000 {
    interrupt-parent = <&ps7_scugic_0>;
    compatible = "fpgafabric";
    status = "okay";
    reg = <0x50001000 0x00001000>;
    interrupts = <0x0 30 0x4>;
};

ker2@50002000 {
    interrupt-parent = <&ps7_scugic_0>;
    compatible = "fpgafabric";
    status = "okay";
    reg = <0x50002000 0x00001000>;
    interrupts = <0x0 31 0x4>;
};

ker3@50003000 {
    interrupt-parent = <&ps7_scugic_0>;
    compatible = "fpgafabric";
    status = "okay";
    reg = <0x50003000 0x00001000>;
    interrupts = <0x0 32 0x4>;
};
```


Developing the software

Compile and execute the reference code in the Virtual Prototype.

- Simple and easy.
- Only recompile dependent libraries (if any).

Low level programming in order to use the hardware modeled.

- Easy once examples are provided.

Most of the time invested adapting the reference code (sequential) to other models of computation.

- Finite state machines.
- Concurrent execution.
- Pipeline.

Mentor has several projects using GStreamer.

- Not checked in this case study, but interesting for the future.

The screenshot shows an IDE window with a project explorer on the left and a code editor on the right. The project explorer shows a tree structure with folders like 'Factorial', 'imagedisplay', and 'imagedisplay_a10'. The code editor displays the file 'kernel_management.h' with the following C++ code:

```
66     }
67 }
68
69
70 void memoryMappings() {
71     /* mmap the UIO devices */
72     imgext_reg = (imgext_reg_t *) mmap(NULL, IMGEXT_REG_SIZE,
73         PROT_READ | PROT_WRITE, MAP_SHARED, imgExtDevice, 0);
74     if (MAP_FAILED == imgext_reg) {
75         perror("mmapn imgext_reg failed");
76         exit(EXIT_FAILURE);
77     }
78
79     kernell_reg = (kernel_reg_t *) mmap(NULL, KERNEL_REG_SIZE,
80         PROT_READ | PROT_WRITE, MAP_SHARED, kernel1Device, 0);
81     if (MAP_FAILED == kernell_reg) {
82         perror("mmapn kernell_reg failed");
83         exit(EXIT_FAILURE);
84     }
85
86     kernel2_reg = (kernel_reg_t *) mmap(NULL, KERNEL_REG_SIZE,
87         PROT_READ | PROT_WRITE, MAP_SHARED, kernel2Device, 0);
88     if (MAP_FAILED == kernel2_reg) {
89         perror("mmapn kernel2_reg failed");
90         exit(EXIT_FAILURE);
91     }
92
93     kernel3_reg = (kernel_reg_t *) mmap(NULL, KERNEL_REG_SIZE,
94         PROT_READ | PROT_WRITE, MAP_SHARED, kernel3Device, 0);
95     if (MAP_FAILED == kernel3_reg) {
96         perror("mmapn kernel3_reg failed");
97     }
98 }
```

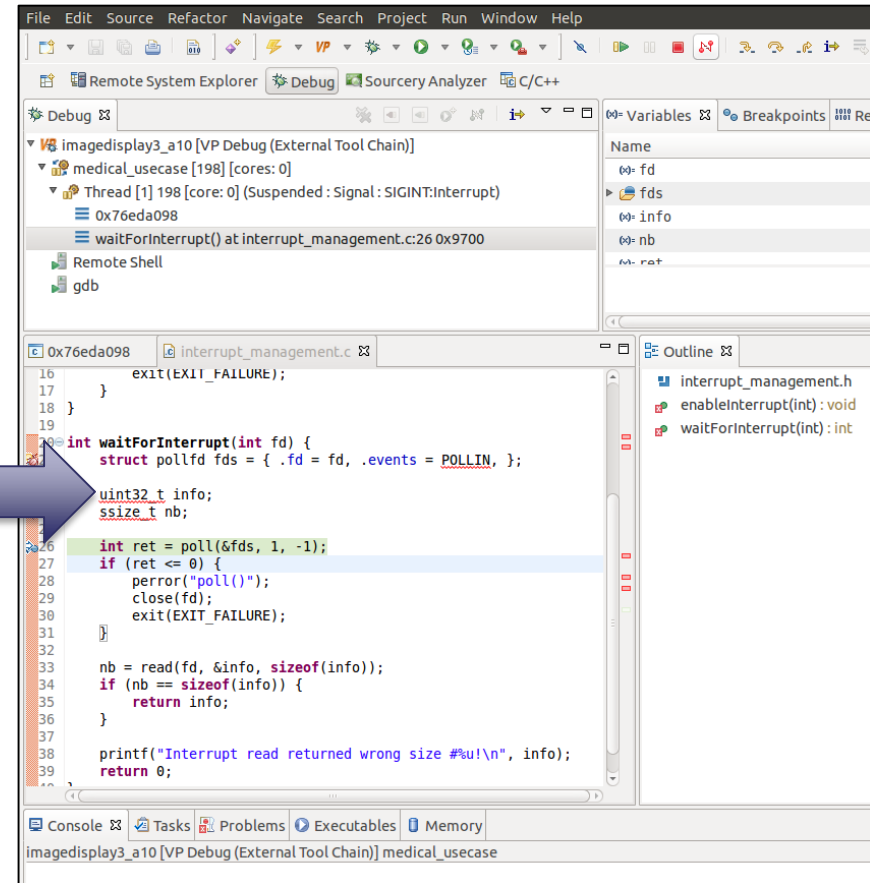
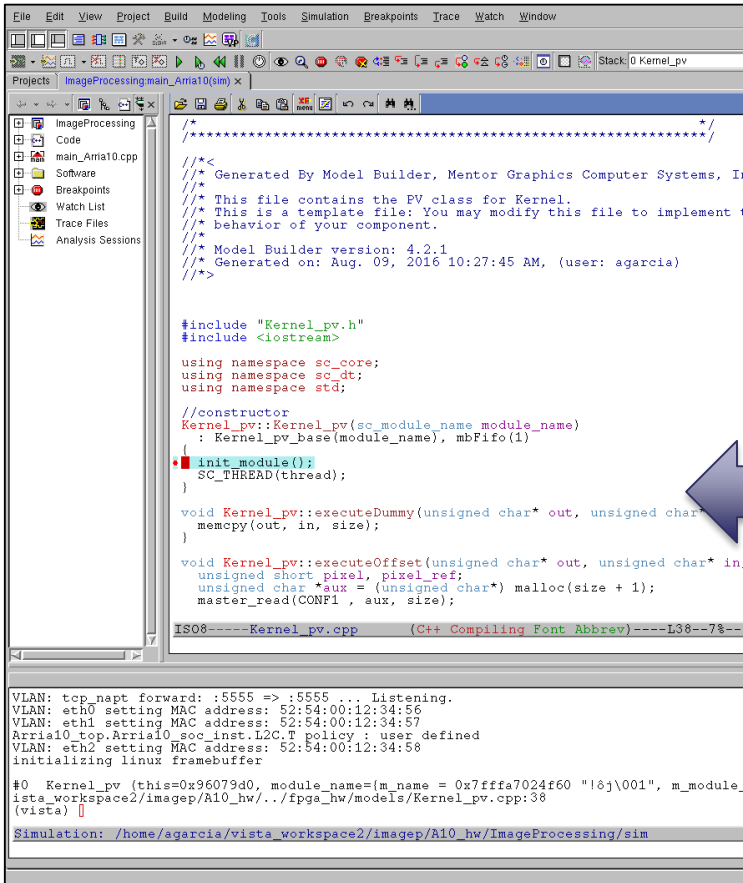
The terminal window at the bottom shows the execution output:

```
<terminated> imagedisplay3_zynq [VP Debug (External Tool Chain)] Remote Shell
Executing kernel 1...
Executing kernel 2...
Executing kernel 3...
Executing kernel 4...
Processed frame 71.
Executing kernel 1...
Executing kernel 2...
Executing kernel 3...
Executing kernel 4...
Processed frame 72.
Processed 73 frames. Closing.
Frames Per Second - 14.141399
Child exited with status 0
GDBserver exiting
```

HW/SW Co-debugging

Dual approach: Vista for debugging the HW and Sourcery for the SW.

Quite simple to perform co-debugging if some insights are provided.



This document may not be reproduced, modified, adapted, published, in any way, in whole or in part or disclosed to a third party without the prior written consent of Thales - © Thales 2016 All rights reserved.

Platform profiling

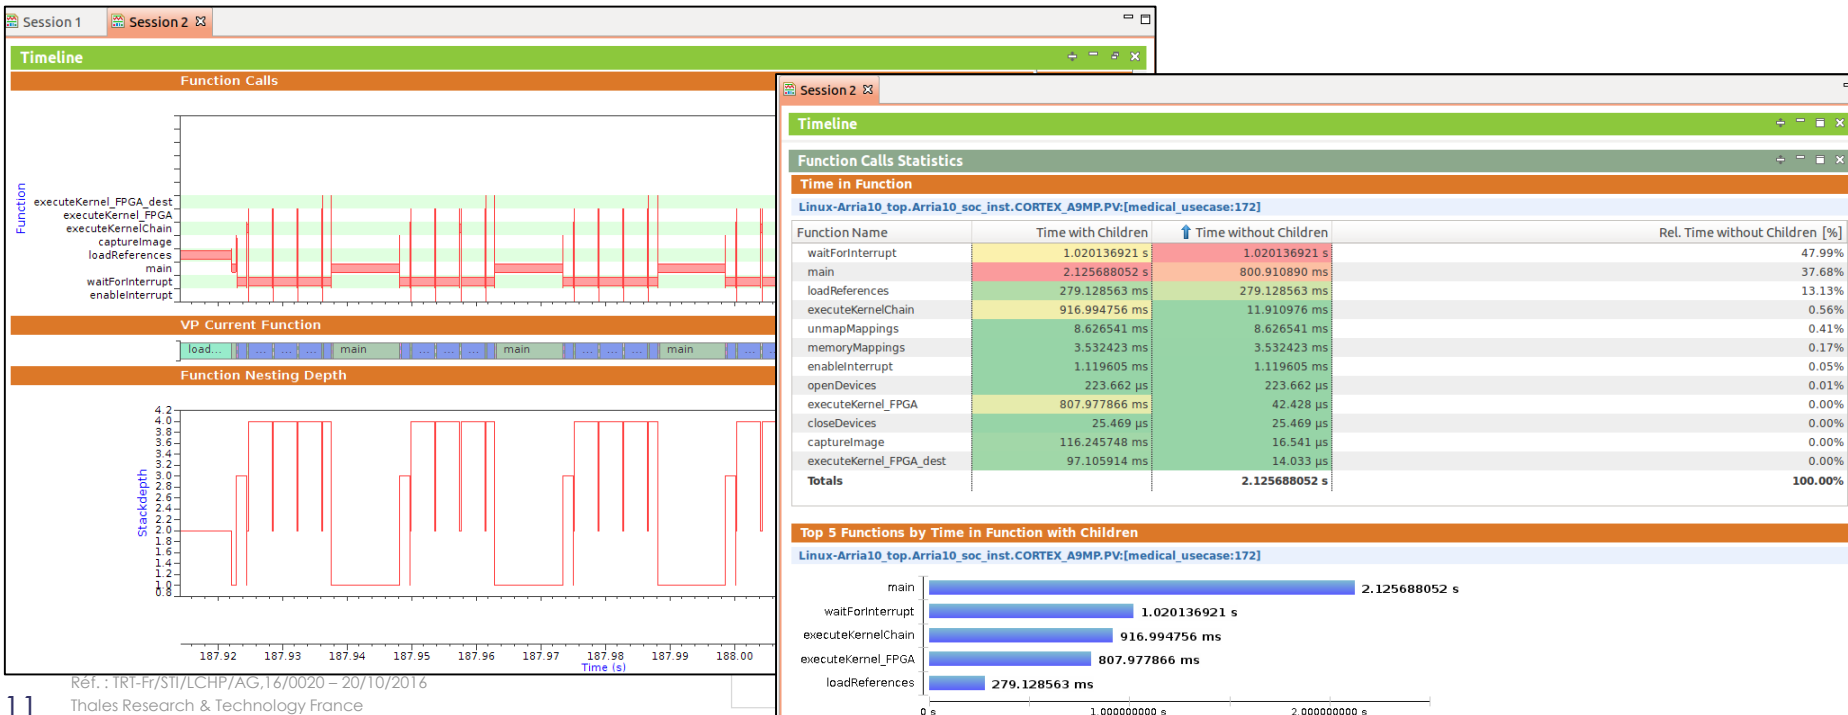
Entities and functions to be traced defined through a trace file.

- HW entities (models, buses, signals, CPUs, etc.).
- SW entities (OS, functions, etc.).

Problem with very large dense traces or very long simulations.

- Actually it's a problem common to all tracing tools.

Very good visualization of results and tracing analysis tools.



Design space exploration

Migration from Arria10 to Zynq.

- HW FPGA models almost the same.
- SW application code totally portable.

Change the bus sizes of the FPGA HW models.

- Information provided by Mentor Vista's compiler.
- Change bus sizes of the fabric or use adapters.

Find the corresponding bus between FPGA fabric and SoC.



- Learn about the underlying architecture.
- Choose a new bus to connect.
- Recalculate memory mapping.

Redo the DTS file for the new platform.

- Requires deep understanding of the platforms and their DTSs.

Not very different devices (A10 vs Zynq).


- Both of them SoC.
- Waiting for the Zynq Ultrascale.
- Still untested real heterogeneity.
- GPGPUs unexplored.



- Zynq Ultrascale+ MPSoC
- ARM Cortex-A53 APU
- ARM Cortex-R5 RPU
- Mali-400MP2 GPU
- H.264/H.265 VCU



OPEN



THALES

■ Mentor Vista and Sourcery Codebench are very mature tools.

- Ready to be used in real developments.
- Still some level of expertise is required.

■ Intended to help the HW/SW co-integration using a virtual prototype.

- Not focused on fast iteration and early (real) prototyping.
- Best use case: software application on SoCs.
- No real heterogeneity (for now). No GPGPU, many-cores, or DSPs.

■ Fast prototyping on FPGAs: check out HW/SW co-designers.

- Reduced set of supported HW platforms.

■ Explore real heterogeneity: check out SW parallelizers.

- High-level interfaces: OpenCL, OpenMP.
- HW design left aside. No FPGA (for now).

Future works

Share the expertise about the tool in order to speed-up the learning curve.

- Accelerate the transition between functional specification and architecture implementation.
- Reduce specific knowledge required.

Provide reference implementations ready to be used.

- Generic HW models.
 - Common kernels or data access patterns.
 - Information about timing and power consumption.
 - Bus mappings.
- Application templates implementing different models of computation.
 - Bradford Nichols, Dick Buttlar and Jacqueline Proulx Farrell. *Pthreads Programming*. 1996. O'Reilly.
 - *Boss/Worker*.
 - *Peer model*.
 - *Pipeline model*.

