A Roadmap for Reconfigurable Computing in the DataCenter.

FPGA'14, Feb. 26, 2014

H. Peter Hofstee, PhD
Workload-Optimized and Hybrid Systems
IBM Research

- AGENDA
-
  - Current uses of Reconfigurable Logic at IBM
  - Limitations
  - Capabilities with Power8 / CAPI

- Current and Recent uses of Reconfigurable Logic at IBM
-
  - Datapower
  - Storage Controller
  - Netezza
  - Texas Memory Systems (Enterprise FLASH)
  - Z Enterprise Data Compression

Ideal World?

"LegUp.org" High-Level Synthesis

LegUp accepts a standard C program as input and automatically compiles the program to a hybrid architecture containing an FPGA-based MIPS soft processor and custom hardware accelerators that communicate through a standard bus interface. In the hybrid processor/accelerator architecture, program segments that are unsuitable for hardware implementation can execute in software on the processor.

"Our long-term vision is to fully automate the flow in Fig. 1, thereby creating a self-accelerating adaptive processor in which profiling, hardware synthesis and acceleration happen transparently without user awareness."

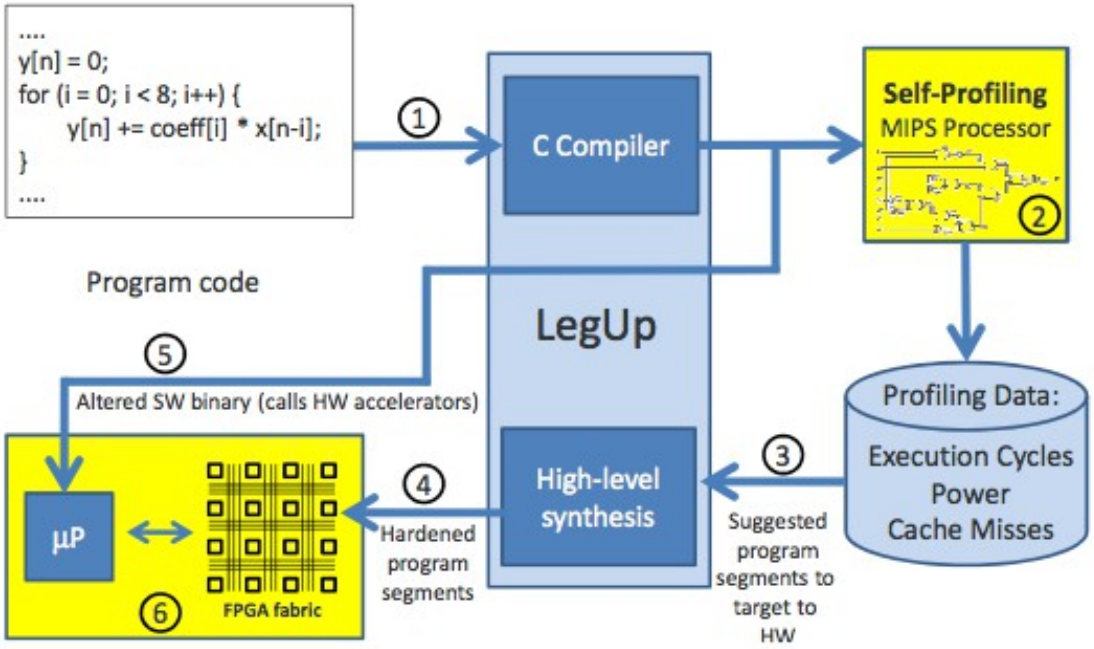

Figure 1: Design flow with LegUp.

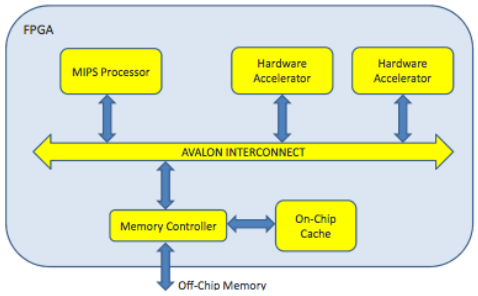

Figure 2: Target system architecture.

University of Toronto / Altera
http://www.eecg.toronto.edu/~janders/fpga60-legup.pdf

- Steps in the right direction ...
-
  - Shared Memory Heterogeneous Architecture
    - No difference in memory model between compute accelerator and CPU core
    - Accelerators are always optional
  - Programming with more Locality Awareness
    - Optimization, not a requirement
    - Usually a good thing to do anyway
  - Programming for predictability
    - Good thing to do anyway
  - Interoperable runtimes and runtime optimization
    - e.g. LLVM
  - Standards-based
    - e.g. OpenCL → OpenMP

POWER8 Integrated PCIe Gen 3

POWER7

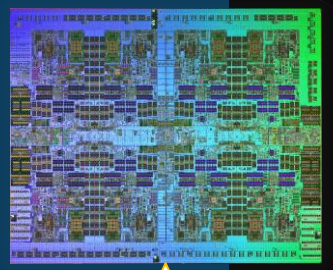

POWER8

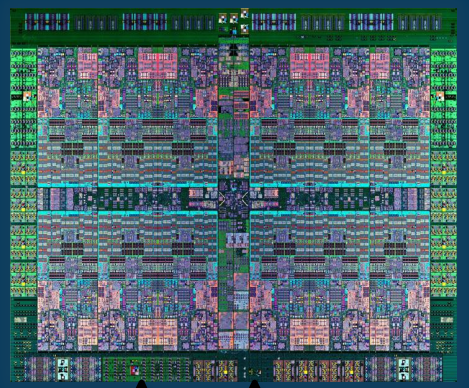

Native PCIe Gen 3 Support

- Direct processor integration
- Replaces proprietary GX/Bridge
- Low latency
- Gen3 x16 bandwidth (16 Gb/s)

Transport Layer for CAPI Protocol

- Coherently Attach Devices connect to processor via PCIe
- Protocol encapsulated in PCIe

GX Bus

I/O Bridge

PCIe G2

PCI Device

PCIe G3

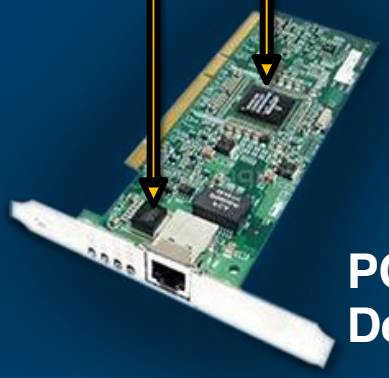

PCI Device

POWER8 CAPI Coherent Accelerator Processor Interface

Virtual Addressing

- Accelerator can work with same memory addresses that the processors use
- Pointers de-referenced same as the host application
- Removes OS & device driver overhead

Hardware Managed Cache Coherence

- Enables the accelerator to participate in “Locks” as a normal thread Lowers Latency over IO communication model

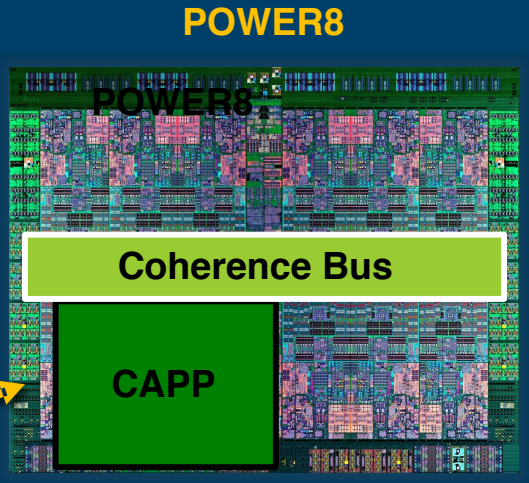

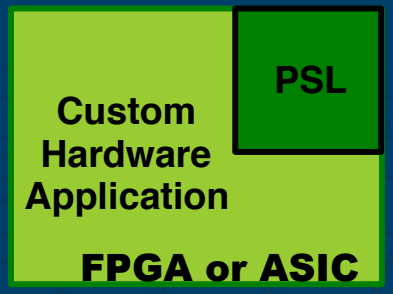

PCIe Gen 3

Transport for encapsulated messages

Processor Service Layer (PSL)

- Present robust, durable interfaces to applications
- Offload complexity / content from CAPP

Customizable Hardware Application Accelerator

- Specific system SW, middleware, or user application
- Written to durable interface provided by PSL

Advantages of Shared Memory Accelerator

- **Virtual Addressing**
  - Removes the requirement for pinning system memory for PCIe transfers
    - Eliminates the copying of data into and out of the pinned DMA buffers
    - Eliminates the operating system call overhead to pin memory for DMA
  - Accelerator can work with same addresses that the processors use
    - Pointers can be de-referenced same as the host application
      - Example: Enables the ability to traverse data structures
- **Elimination of Device Driver**
  - Direct communication with Application
  - No requirement to call an OS device driver or Hypervisor function for mainline processing
- **Enables Accelerator Features not possible with PCIe**
  - Enables efficient Hybrid Applications
    - Applications partially implemented in the accelerator and partially on the host CPU
  - Visibility to full system memory
  - Simpler programming model for Application Modules

- Where we hope this takes us.
-
  - More flexible infrastructure
    - E.g. Cloud
  - Innovation pipeline
    - More DataPowers, Netezzas, etc.
  - More computer architecture research

Thank You!

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)