

Intel Fast Forward Processor and Memory

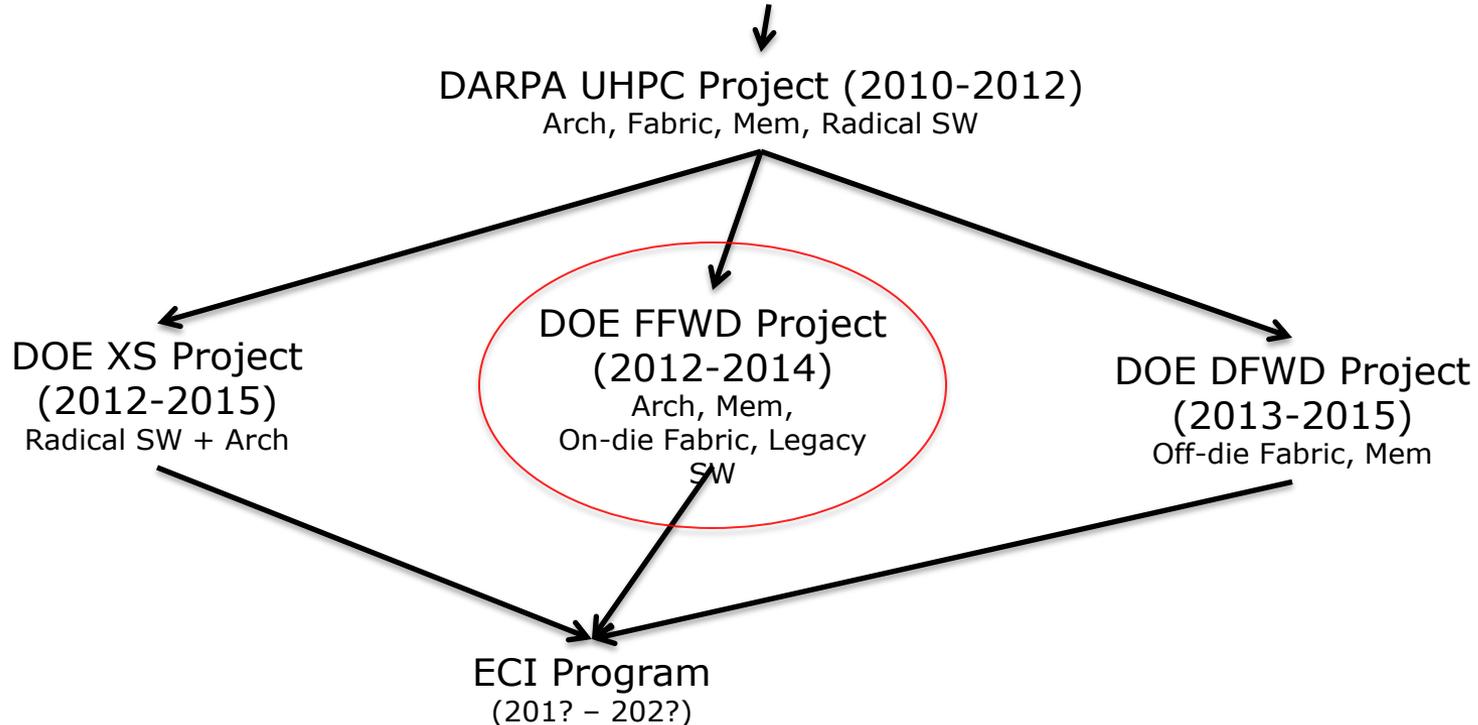
Bill Feiereisen, Intel Corporation

February 2014

The Challenge

- 1 DP ExaFLOP (10¹⁸ DP FLOP/sec) in 20 MW
 - This is ~20 pJ/Op at system level with all taxes; compute ~5 pJ

DARPA and DOE ExaScale Studies (pre-2010):



Top Exascale Challenges

- System Power & Energy
- Efficient, memory subsystems
- Extreme parallelism, data locality, gentle-slope programmability
- New execution model comprehending self awareness with introspection
- Resiliency to provide system reliability

Top Exascale Challenges

- Intel's FFWD:
- System Power & Energy
- Efficient, memory subsystems
- Extreme parallelism, data locality, gentle-slope programmability
- New execution model comprehending self awareness with introspection
- Resiliency to provide system reliability

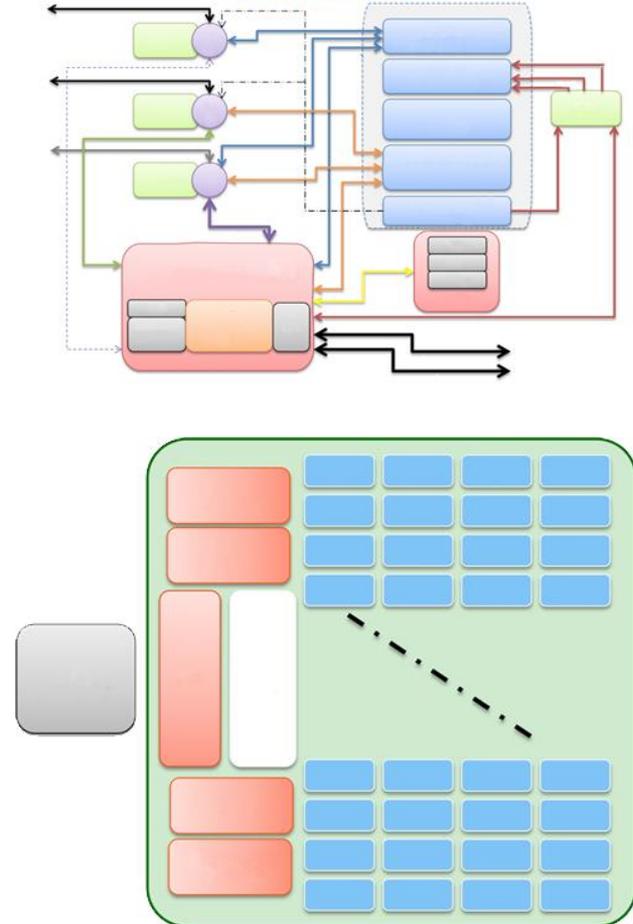
Exascale Architecture Thesis Position

- System view: Energy is king, communication is queen
 - These will make or break the system
- Small simple cores
 - Efficiency and scalability vs IPC limitations
- State management, relaxed ordering, hybrid coherence
 - Limit coherency support, no large-regions of ccSM
 - Inject relaxed memory models, totally visible state
- Prioritize small comms/synch instead of ST/vector perf
 - Strong scaling is a power and cost imperative
 - Dedicate area to efficient (power and latency) comms helpers
- Pay the cost of resiliency on remote operations
 - Confirm comms to recover from remote faults/failures
 - Develop protocols (SW and HW) to quarantine errors, limit ripples
- Make it fractal
 - Same trade-offs apply in hierarchal fashion to entire system
 - Tapering is required in the hierarchy even with over-provisioning

Intel's FFWD Elements

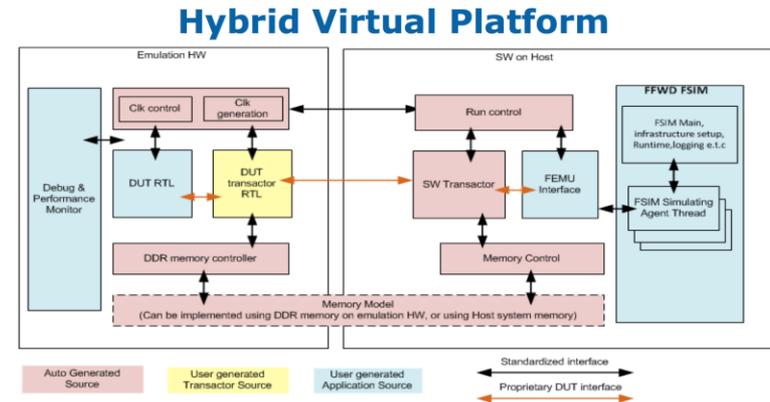
- Core Architecture Processor defined
 - Large number of energy-efficient cores (processing elements PE)
 - With dedicated support units to off-chip interfaces
- Software Tools and Applications are provided:
 - Architecture is exercised in a full execution-driven parallel functional simulator that itself is a parallel application
 - Allows functionality, timing, and energy estimation
 - Compiler toolchain and runtime environment
 - Focus on the six 6 DOE FFWD mini-apps, but prepare for more complex applications
- Design and Implementation:
Architecture is captured in RTL:
 - Provides area and power estimates
 - Is a key software development vehicle for ongoing research
 - Prepares for the fabrication of a future silicon test chip

FFWD Processor microarchitecture overview

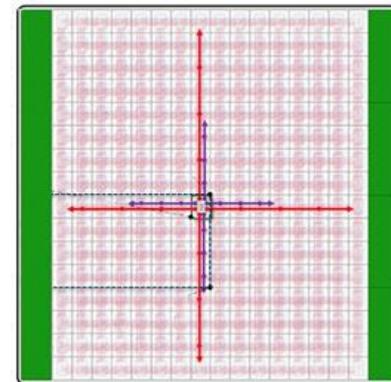


Intel's FFWD Elements

- RTL is implemented in an FPGA Emulator:
 - Enables hardware/software co-design by using a Hybrid Virtual Platform for emulation
 - Improves RTL validation coverage, including comprehensive fault resiliency analysis.
 - Allows performance estimation by running mini-apps on the emulator
- Alternate memory architectures and on-die interconnect topologies are explored:
 - Hierarchical and heterogeneous memory technologies
 - Various approaches to connect memories and cores on die
- Co-Design
 - The design is exercised in the simulator and the emulator through the porting and simulation of the six mini-applications LULESH, CoMD, miniFE, SMC, SNAP, Nekbone ... and CLOMP



Memory and Network Interconnect



Co-design

- ... has been very useful to us ... (but we do receive mixed messages) 😊
 - We have adjusted the ISA and improved workload scaling
 - We are able to show the limits that the physical designs allow
 - Opened up the conversation to adjust “algorithmic thinking”
 - Open and honest dialog over MPI and OpenMP
- Hack-a-thons bring Labs together with Intel researchers
 - to exercise the design and alter it collaboratively
 - to alter the software stack and mini-apps
 - two sessions in 2013 and one in 2014 at an Intel facility with active Intel engineers in the room
 - Hands-on with internal simulators, tools

Summary

- First revision of next-gen Arch and models established
- NTV computations to reduce energy/op.
- Data movement limited by power and energy
 - Observed bandwidth tapers as the distance grows
- Data locality (programming system) is paramount
- Increase the on-die memory to compute ratio
- Rethink chip and chip design
 - Wires and transistors are cheap, feeding them expensive
- Sensors for introspection for dynamic resource management

Legal Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel, Xeon, Look Inside and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright ©2014 Intel Corporation.

Risk Factors

The above statements and any others in this document that refer to plans and expectations for the third quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as “anticipates,” “expects,” “intends,” “plans,” “believes,” “seeks,” “estimates,” “may,” “will,” “should” and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel’s actual results, and variances from Intel’s current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the company’s expectations. Demand could be different from Intel’s expectations due to factors including changes in business and economic conditions; customer acceptance of Intel’s and competitors’ products; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Uncertainty in global economic and financial conditions poses a risk that consumers and businesses may defer purchases in response to negative financial events, which could negatively affect product demand and other related matters. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel’s products; actions taken by Intel’s competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel’s response to such actions; and Intel’s ability to respond quickly to technological developments and to incorporate new features into its products. The gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; segment product mix; the timing and execution of the manufacturing ramp and associated costs; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. Intel’s results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel’s products and the level of revenue and profits. Intel’s results could be affected by the timing of closing of acquisitions and divestitures. Intel’s results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues, such as the litigation and regulatory matters described in Intel’s SEC reports. An unfavorable ruling could include monetary damages or an injunction prohibiting Intel from manufacturing or selling one or more products, precluding particular business practices, impacting Intel’s ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. A detailed discussion of these and other factors that could affect Intel’s results is included in Intel’s SEC filings, including the company’s most recent reports on Form 10-Q, Form 10-K and earnings release.