




Intel's Knights Landing – what's old, what's new?

April 2nd 2014, Annual Concurrency Forum Meeting
Andrzej Nowak, CERN openlab CTO office



old die shot

KNC 2012	KNL 2015+
Core	
61 cores	72 cores
P54C (Pentium Pro)	Silvermont (Atom)
1-1.2 GHz	?
Memory	
16 GB GDDR	16 GB eDRAM (3D)
0 GB DDR3	Up to 384 GB DDR4
Package	
22nm	14nm
1 TF DP	3 TF DP
PCIe gen2	PCIe gen3, Socket

Core progression

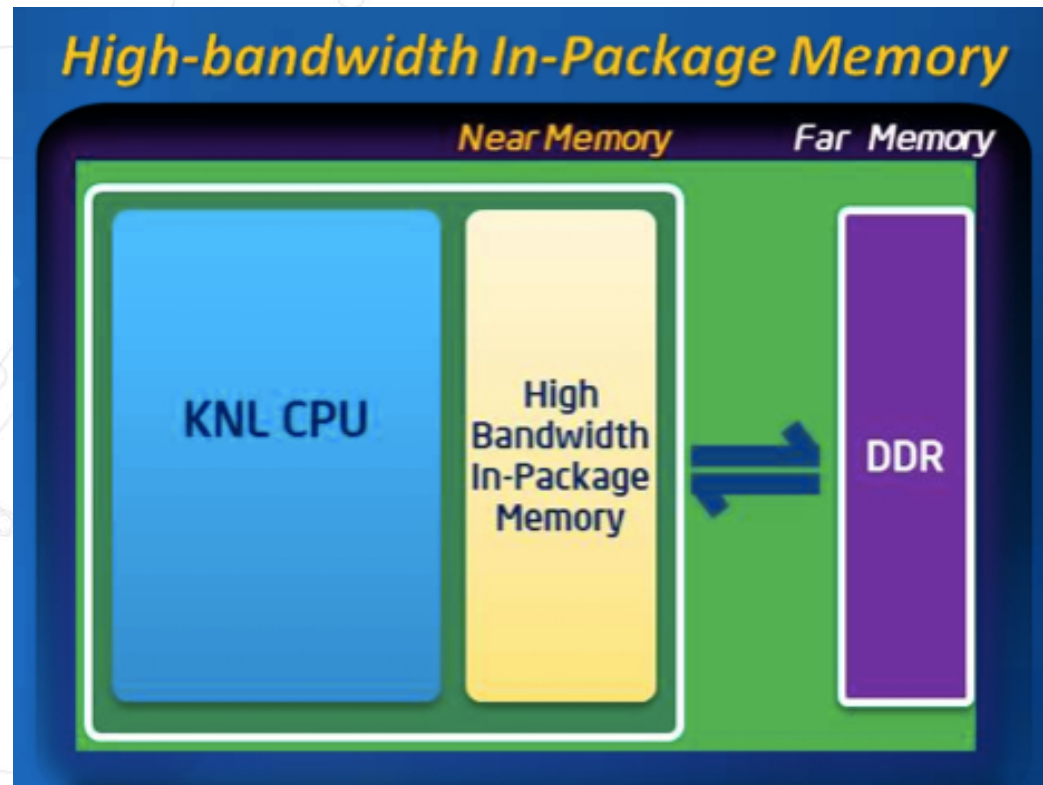
- Move from a beefed-up P54C core to a Silvermont (modern Atom)
 - Out of order!
 - Still 4 threads/core
 - 1 MB L2 shared across two cores – no change in cache size
- Mesh fabric interconnect rather than ring bus
- AVX-512 (3.1) implemented – convergence with Xeon
- The promise: 3-4x more GF/Watt than current supercomputers. Will it hold?

Programmability (1)

- Bootability implies wide-spread OS support
 - Fully native applications made easy (i.e. normal)
 - Kernel patches already inbound
- AVX 512 convergence will improve the quality of software tools and compilers
 - Less maintenance, more performance
 - No guarantee of binary compatibility so far
- Parallelization, vectorization
 - No escape

Programmability (2)

- A new layer of memory
 - High-BW eDRAM on chip, DDR further away
 - Can we afford to remain agnostic?



“Crazy” stuff

- >100 PF system target
- 3D memory on board, 500GB/s
- Socket option (single)
- Integrated next-gen 100Gbit interconnect: Storm Lake
- 200W TDP only?



What does this mean for us?

- No fundamental shift
 - Still need to parallelize, vectorize
- Better software support
 - Better chance to get good performance with GNU and open source stacks
- Faster cores
 - Always better
- Not many more cores
 - Reduces scaling issues
- Socket option and bootability
 - Easier management, finally lots of memory
- Integrated interconnect
 - Fantastic device for data taking and HPC

Thank you



Andrzej.Nowak@cern.ch