



Top Searches: Apple • Android • Windows 8 • iPad      Trending: Windows 8 • Batteries • 3D • Automobiles • Slideshow • Cloud


## Intel’s 50-core champion: In-depth on Xeon Phi

By Joel Hruska on July 30, 2012 at 12:06 pm

We recently had a chance to sit down with Intel, talk about the company’s upcoming Xeon Phi/Many Integrated Core (MIC) hardware, and discuss how the card is a vital step along the path to exascale computing. Intel isn’t ready to disclose the card’s full technical specifications, but a number of scientific and research institutions have released papers detailing their experience with Knights Ferry, Xeon Phi’s predecessor. Just so we’re clear on terminology, MIC is an architecture, Knights Ferry and Knights Corner are specific implementations of that architecture, and Xeon Phi is the brand name under which KNC will be sold.

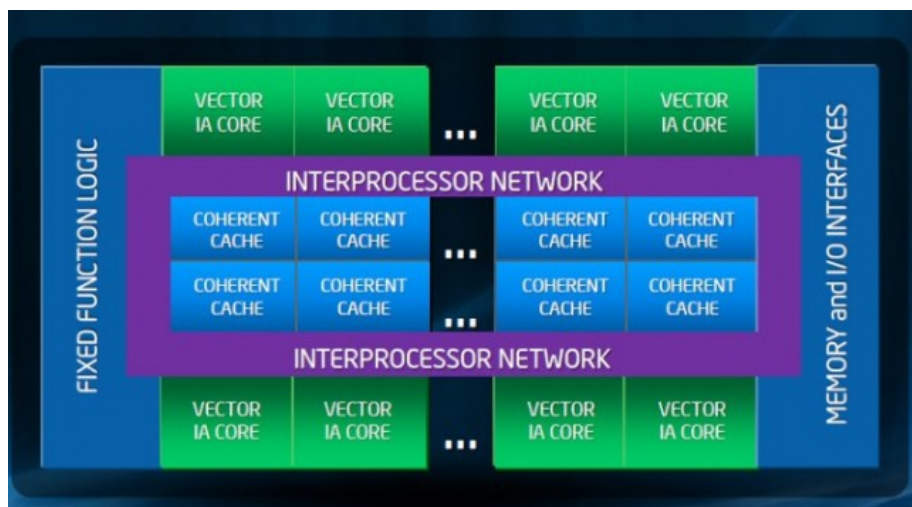
### From Larrabee to Knights Ferry

Intel’s MIC (pronounced “Mike”) began as a hybrid GPU/HPC (high-performance computing) product known as Larrabee. Intel officially announced Larrabee in 2007 and soon claimed that the card would usher in a new era of ray-traced video games and incredible performance. Intel eventually shelved its GPU ambitions once it became clear that Larrabee wasn’t going to be able to match the high-end hardware then available from Teams Green and Red, and rebranded Larrabee as an HPC-only part. The new design was dubbed Knights Ferry (KNF), and Intel began shipping it out to HPC developers in 2010.



[1]  
Larrabee's original design

So how much of Larrabee is left in Intel’s MIC? It depends on where you look. All of the display hardware and integrated logic necessary to drive a GPU is gone, but the number-crunching capabilities of the cores themselves appear largely unchanged. One difference we know of is that while Larrabee and KNF focused on single-precision floating point math, the upcoming Knights Corner will offer strong double-precision support as well. Compare Larrabee’s block diagram, above, with Knights Ferry, below.




[2]  
The relabeled Xeon Phi MIC block diagram

### A Knights Tale

So let’s talk about Knights Corner/Xeon Phi. Xeon Phi ramps up Knights Ferry; Intel isn’t giving many details yet, but we know the architecture will pack 50 or more

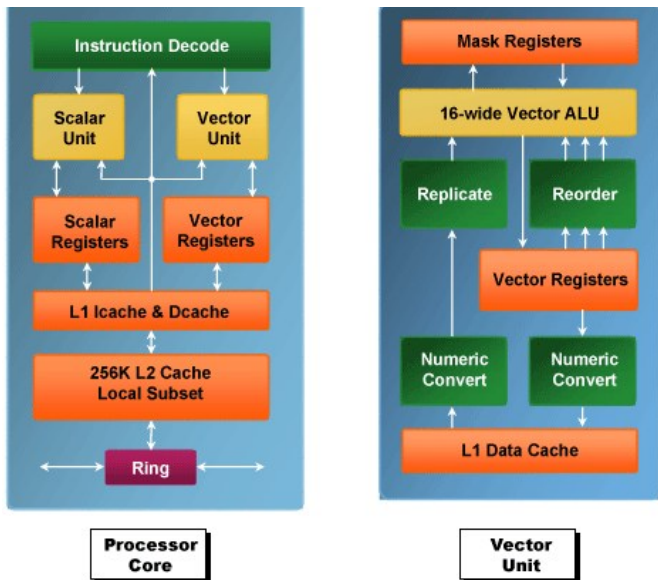
cores and at least 8GB of RAM. In this space, total available memory is an important feature. Knights Ferry, with its 32 cores and max of 2GB of RAM, could only offer 64MB of RAM per core; a 50-core Xeon Phi with 8-16GB of RAM would offer between 163-327MB per core.



[3]

The Knights Ferry die, Aubrey Isle. Die size on KNF, at 45nm, was rumored to be roughly 700mm sq.

It's logical to think Intel's core counts and RAM load will vary depending on yields and customer needs. Customers with large regular datasets might see better scaling from a 50-core chip with 16GB of RAM, while small data sets might do best with an 8GB card and 64 cores. The layout of the Aubrey Isle die at the heart of Knights Ferry, pictured above, makes a 64-core target chip a strong possibility, with varying numbers of cores disabled to improve yields.



[4]

The overall Pentium core is on the left, the vector engine is on the right. Diagrams are not to scale

The cores at the heart of Intel’s first Xeon Phi are based on the P54C revision of the original Pentium and appear largely unchanged from the design Intel planned to use for Larrabee. Despite some squabbling from Team Green, we recommend not conflating the phrase “based on,” with “hampered by.” Intel returned to the P5 microarchitecture for Larrabee because it made good sense to do so — but Knights Corner isn’t a bunch of early 1990s hardware glued on a PCB. Intel has added 64-bit support, larger on-die caches (the Pentium Classic never had an on-die L2, or an L1 with 1-cycle latency), a 512-bit, bi-directional ring bus that ties the entire architecture together, advanced power consumption management circuitry, and 32 512-bit vector registers. It’s the latter that give Xeon Phi its oomph; A top-end Core i7 today has 16 256-bit AVX registers.

From a computational perspective, calling Knights Corner a “modified Pentium” is like calling the starship Enterprise a modified Space Shuttle. The updated P54C core is better thought of as a launch gantry; it’s the framework Intel used for creating something new, not the vehicle itself.

Is Knights Corner x86 compatible? Mostly — or, perhaps more accurately, it’s x86-compatible enough [5]. Intel’s software blog states the following: “Programs written in high-level languages (C, C++, Fortran, etc.) can easily remain portable despite any ISA or ABI [application binary interface] differences. Programming efforts will center on exploiting the high degree of parallelism through vectorization and scaling: Vectorization to utilize Knights Corner vector instructions and scaling to use more than 50 cores. This has the familiarity of optimizing to use a highly-parallel SMP system based on CPUs.”


There are a handful of x86/x86-64 instructions, including a few fairly common ones, that KNC won’t support. The vectorization and scalar instructions that KNC/KNF introduced are also unique — KNC doesn’t support traditional SIMDs like MMX, SSE, or AVX... yet. That “yet” is important, because it’s virtually guaranteed that Intel will cross-pollinate its instruction sets at some point in the future. The Transactional Synchronization Extensions [6] (TSX) set to debut in Haswell might be extremely useful for Knights Corner’s successor.

Next page: But will Xeon Phi perform? [7]

### Will it blend perform?

Intel hasn’t allowed anyone to release hard data on Knights Ferry performance, and there’s absolutely no info available on Xeon Phi — but when we spoke with Dr. Glenn Brook, a research engineer with NICS (National Institute for Computational Sciences), he indicated that Intel’s KNF delivered a solid combination of strong scaling, power efficiency, and overall performance. The various papers [8] highlighted at the TACC-Intel Highly Parallel Computing Symposium this past spring are generally optimistic regarding KNC’s ability to deliver capabilities that the HPC industry will find useful.

It’s important to note that the gains weren’t universal. A team from the National Center for Atmospheric Research presented a paper on KNF’s performance in their climate models [9] and reported that: “In general, we observed excellent scalability with both models on KNF. However, single-thread performance is poor with or without vectorization enabled, which may indeed be the cause of the exceedingly favorable scalability.”





[10]

Intel’s support for OpenMP [11]/MPI, and the fine-grained control of processor resources available on KNF, were also highlighted as significant features. One of the major challenges in HPC research is apparently the need to juggle small jobs with large ones in order to make the most efficient use of available computing resources. Knights Corner is designed to address that concern; the card can be deployed as part of an HPC cluster, or treated as a stand-alone co-processor for crunching specific tasks. It’s also designed to be flexible in terms of software and hardware. We expect, for example, to see some form of Turbo Boost debut on KNC, though likely under a different name. Scaling clock speeds depending on total co-processor workload is one way Intel can improve performance, and it’s one way of offsetting relatively low single-thread efficiency.

KNC can communicate with other MIC cards across the PCIe bus and can even kick off jobs on other boards, but the PCIe bus only offers a fraction of the bandwidth the card has access to internally. KNF was a PCIe 2.0 solution and it wouldn’t be surprising if Knights Corner keeps that bus rate — moving to PCIe 3.0 would double effective bandwidth but that’s not nearly enough to offset the higher access latency and minimal bandwidth compared to keeping data local.

As important as benchmarks and hard performance figures are, KNC’s real advantage probably won’t come from these metrics.

**Software: Intel’s ace in the hole**


Most supercomputing articles focus almost entirely on hardware; software support isn’t as sexy as racks of big iron jammed with high-end components [12]. After reading several papers released at TACC-Intel and discussing the matter with Dr. Brooks, I’m increasingly convinced that this hardware-centric approach to coverage misses the larger issues of the field.

One of our questions for Dr. Brook was how much Intel’s vaunted x86 compatibility actually mattered to scientists working in the real world [13]. His response was as follows: “It is also important to realize that many NSF-funded researchers are funded to conduct fundamental science *not to port, optimize, or develop scientific codes...* As such, many researchers view code development as an overhead cost that should be avoided or minimized when possible, and most are extremely reluctant to abandon years (or decades) of existing code development to transition to a significantly different programming model.” (Emphasis added).

An institution’s decision to adopt MIC or CUDA may be driven less by horsepower and more by accessibility, even in cases where one option offers substantially better real-world performance. Raw compute performance has increased by an order of magnitude in the past seven years. The fastest computer from the June 2005 Top500 list is ranked #125 today — but the third fastest machine from 2005 doesn’t even make the list. Humans, unfortunately, haven’t gotten any upgrades in the intervening period. Efficiently managing the distribution of all that compute performance remains an extremely challenging task, which is part of why the CPU remains the basic currency of the HPC realm.

This is where Intel’s biggest advantage over Nvidia will come into play, and it’s going to be a damned hard card for the GPU designer to counter. If you program on x86 and you need a performance monitor, thread analyzer, optimized library, or compiler, Intel has you covered. The company’s tools aren’t just numerous; they’re sophisticated, and they offer substantial under-the-hood visibility. That visibility is important when it comes to optimizing workloads for cluster computing. Latencies and chip-to-chip communication levels that are entirely manageable in a 4S (quad-socket) system can overwhelm a network when scaled across nodes.

**Accelerator/Co-Processor System Share**



[14]

Of the 58 Top500 computers with co-processor/accelerators installed, 53 of them use Nvidia hardware. That’s a huge jump from June 2011, when just 12 systems used the company’s CUDA cards, but it pales in comparison to the 374 systems using Intel CPUs. According to Intel, optimizing for Xeon Phi can often improve performance on traditional Xeons, even though the two processors don’t share vectorization capabilities. Unlike Intel, Nvidia can’t build a “pure” Tesla box in which both CPUs and GPUs carry their own brand, and that gives Intel a natural advantage.

None of this, however, means that Nvidia is doomed. Far from it. I wouldn’t be surprised if the upcoming Kepler-based K20 Tesla is capable of beating Knights Corner’s performance in some real-world tests. NV’s upcoming Tesla might also have a size advantage; Xeon Phi’s die is going to be absolutely enormous. Visual estimates on Intel’s Knights Ferry pegged its die size at close to 700 mm<sup>2</sup>; Knights Corner is built on a much smaller process (22nm vs. 45nm) but it also adds cores. Yield management will be critical for both companies, which is probably part of the reason Nvidia is waiting until Q4 [15] to launch the part.

The coming war between Intel and Nvidia for the supercomputing market will have a real impact on consumer products, even if it takes several years for the research to trickle down. The HPC industry is struggling to deal with problems of power efficiency, interconnect scaling, storage speed, processor utilization, and communication latency. The mobile phone and tablet markets, meanwhile, are fighting with the very same problems, with the added headache of battery life thrown in. Advances at the top of the market will increasingly shape the bottom (and vice versa). The battle to be the company whose hardware powers both spheres of influence is about to kick off in earnest.

Read: The history of supercomputers [16], or What are supercomputers actually used for? [17]

**Endnotes**

1. : <http://www.extremetech.com/wp-content/uploads/2012/07/Larrabee-GPU1.jpg>
2. : <http://www.extremetech.com/wp-content/uploads/2012/07/Intel-Mic.jpg>

3. : [http://www.extremetech.com/wp-content/uploads/2012/07/Aubrey\\_Isle\\_die.jpg](http://www.extremetech.com/wp-content/uploads/2012/07/Aubrey_Isle_die.jpg)
4. : <http://www.extremetech.com/wp-content/uploads/2012/07/KNF-Core.jpg>
5. x86-compatible enough: <http://software.intel.com/en-us/blogs/2012/06/05/knights-corner-micro-architecture-support/>
6. Transactional Synchronization Extensions: <http://www.extremetech.com/computing/117501-intels-haswell-will-include-new-multi-core-enhancements>
7. But will Xeon Phi perform?: <http://www.extremetech.com/extreme/133541-intels-64-core-champion-in-depth-on-xeon-phi/2>
8. various papers: <http://www.tacc.utexas.edu/ti-hpcs12/program>
9. climate models: <http://www.tacc.utexas.edu/documents/13601/2ecbc5f2-2519-4650-8b60-298cb035247e>
10. : <http://www.extremetech.com/wp-content/uploads/2012/07/XeonPhiCard.jpg>
11. OpenMP: <http://en.wikipedia.org/wiki/OpenMP>
12. big iron jammed with high-end components: <http://www.extremetech.com/extreme/125271-the-history-of-supercomputers>
13. in the real world: <http://www.olcf.ornl.gov/wp-content/training/electronic-structure-2012/20120206-Brook-AACE-MICAApps.pdf>
14. : <http://www.extremetech.com/wp-content/uploads/2012/07/Top500.png>
15. waiting until Q4: <http://www.extremetech.com/computing/130937-tsmc-still-struggling-with-28nm-qualcomm-and-nvidia-threaten-to-jump-ship>
16. The history of supercomputers: <http://www.extremetech.com/extreme/125271-the-history-of-supercomputers>
17. What are supercomputers actually used for?: <http://www.extremetech.com/extreme/122159-what-can-you-do-with-a-supercomputer>

---

Printed from <http://www.extremetech.com/extreme/133541-intels-64-core-champion-in-depth-on-xeon-phi>. Copyright ©2012 **ExtremeTech** unless otherwise noted.