SAND2011-8668P

Arthur: Sandia's NNSA/ASC Experimental Architecture Testbed with 84 Intel® Knights Ferry Cards

November 15, 2011

James Ang, Ph.D.
Manager, Scalable Computer Architectures
Sandia National Laboratories
Albuquerque, NM 87185

Acknowledgements:
James Laros III, Matthew Bohnsack, Victor Kuhns, Jason Repik

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Intel® Many Integrated Core (MIC) Architecture Testbed

- *Arthur* integrated by Appro International and accepted by Sandia on 9/30/11
- *Arthur* is a "first of a kind" 42 node experimental Knights Ferry (KNF) cluster
  - Node has two 6-core Intel® Xeon® processor 5600 @ 3.46GHz and 24 GB DDR3-1600MHz
  - Node also has two 30-core Intel® Knights Ferry software development cards @ 1.05GHz and each card has 2GB GDDR5-1800Mhz
  - Node has one 80GB Intel® SSD SATA 3Gb/s, MLC NAND Flash drive
- Interconnection Network: Mellanox Infiniscale IV QDR Infiniband
- Separate Ethernet system management network
- Planned Upgrades
- Early 2012 – Upgrade with Future Intel® Xeon® processor E5 family
- 2012 – Replace KNF with pre-production Knights Corner (KNC) co-processors

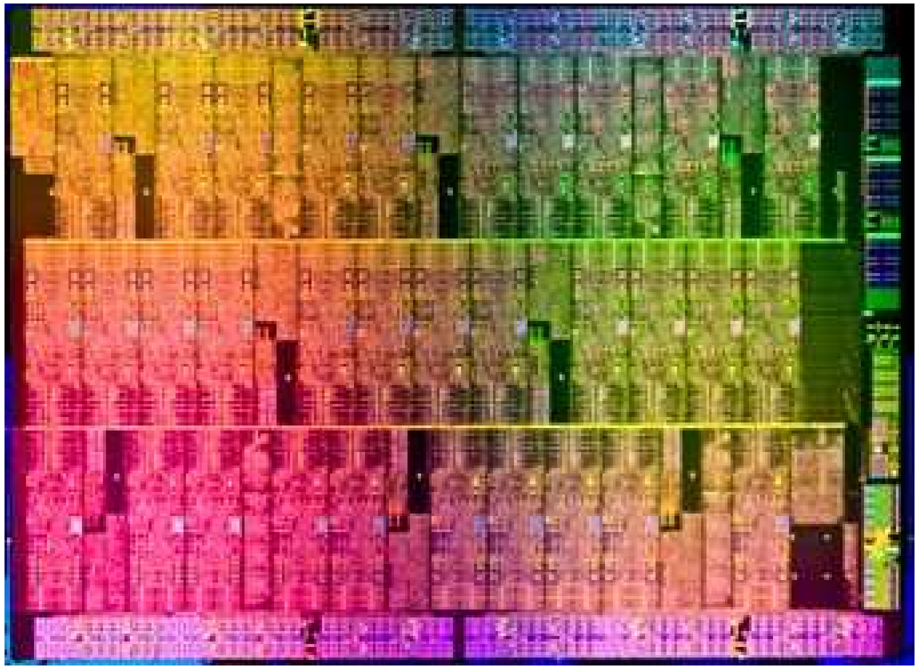

Aubrey Isle*

*Intel coprocessor on the KNF Card

Intel MIC Architecture Testbed

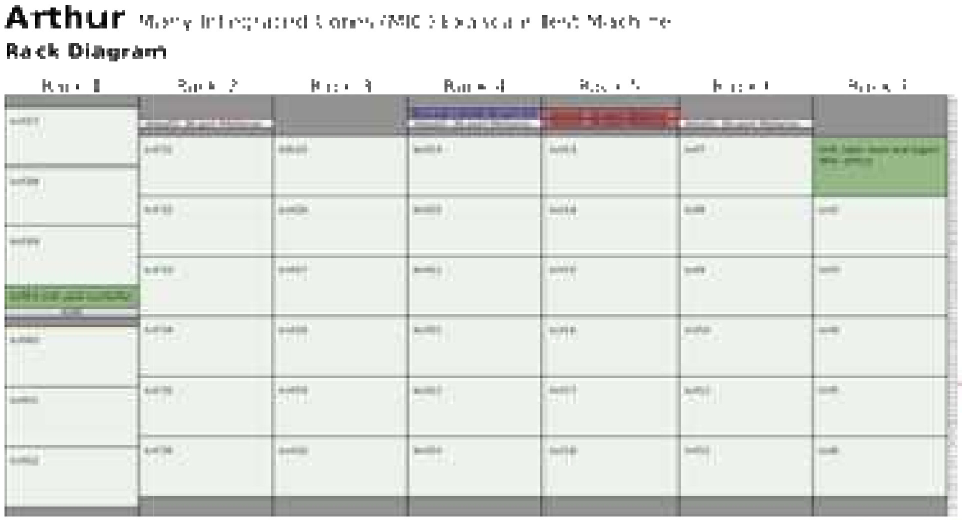

Arthur Diagrams

– Matt Bohnsack

Arthur Photographs

– Victor Kuhns

Intel MIC Architecture Testbed

Arthur Many Integrated Cores (MIC) Exascale Test Machine

InfiniBand Diagram - QDR Fat Tree with 5 36-port Switches

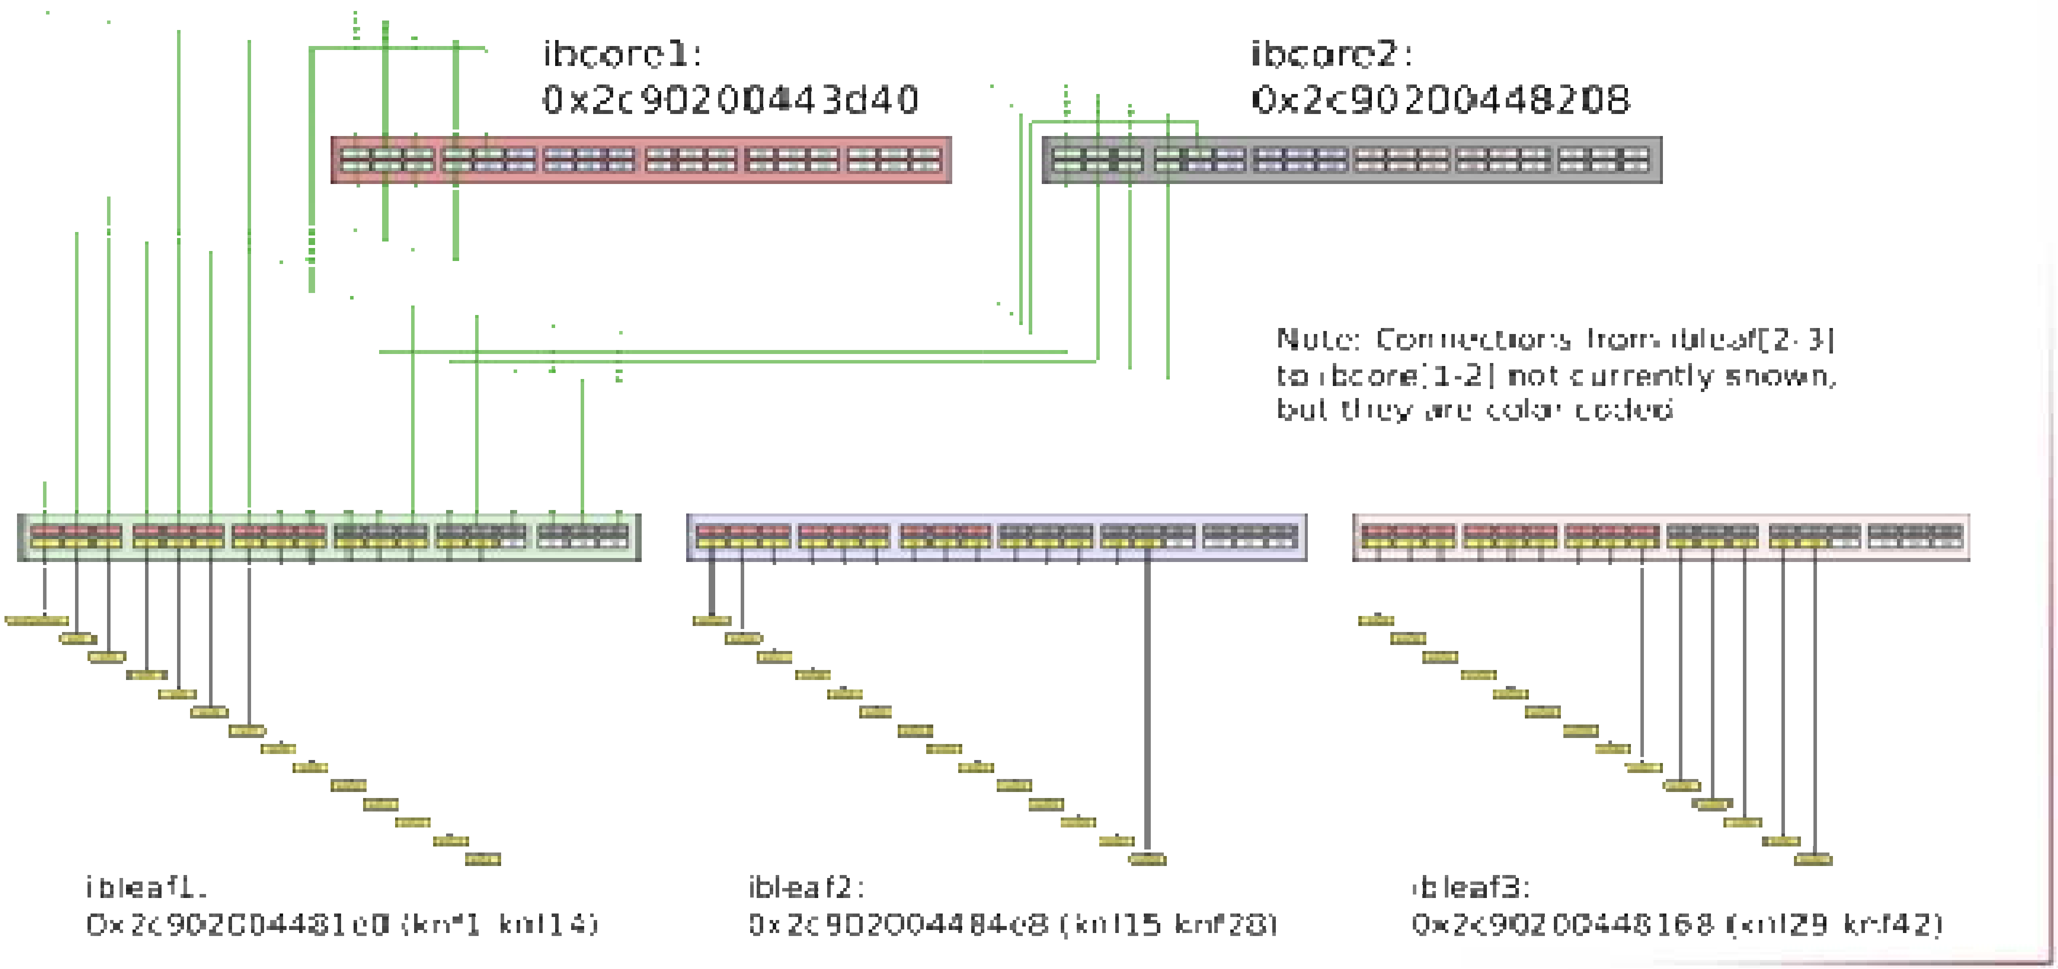

Testbed Experiments

- Run Mantevo proxy applications on Arthur to understand the data movement options
  - Initial testing with miniFE, miniMD, and miniGhost
    see http://mantevo.org
  - Investigate and Evaluate Coding of Mantevo miniapps with Intel® MIC Programming Models: Intel® TBB, ArBB, and Cilk™ Plus
- Work with Intel® on the University of Minnesota's PPM turbulent, compressible CFD simulation
- Validation of SST architectural simulation results

  see http://code.google.com/p/sst-simulator/
- System Software R&D
  - Portals4 Implementations

    see http://code.google.com/p/portals4
  - Kitten Lightweight Kernel and Runtime software
  - Runtime support for power management

Initial Simulation Results

- [see an image from Paul Woodward's PPM simulation]:
  - Piece-wise Parabolic Method, compressible, turbulent fluid flow simulation results
  - Visualize 3D Arthur PPM results at NNSA/ASC exhibit – #803

Thank You

Drawing for all attendees

Chance to win an iPod Nano
Touch Screen – 8GB

Questions for Appro Give-Away

Participants who answer the following questions correctly will receive an Intel® Hat!

#1: What is the name of Sandia's Appro, Intel MIC experimental architecture testbed?

#2: How many Knights Ferry cards are in the experimental architecture testbed?

#3: What is the interconnection network in the experimental architecture testbed?

See me after for any questions you might have!!!!

Backup Slides from Intel®

Aubrey Isle Co-Processor Architecture

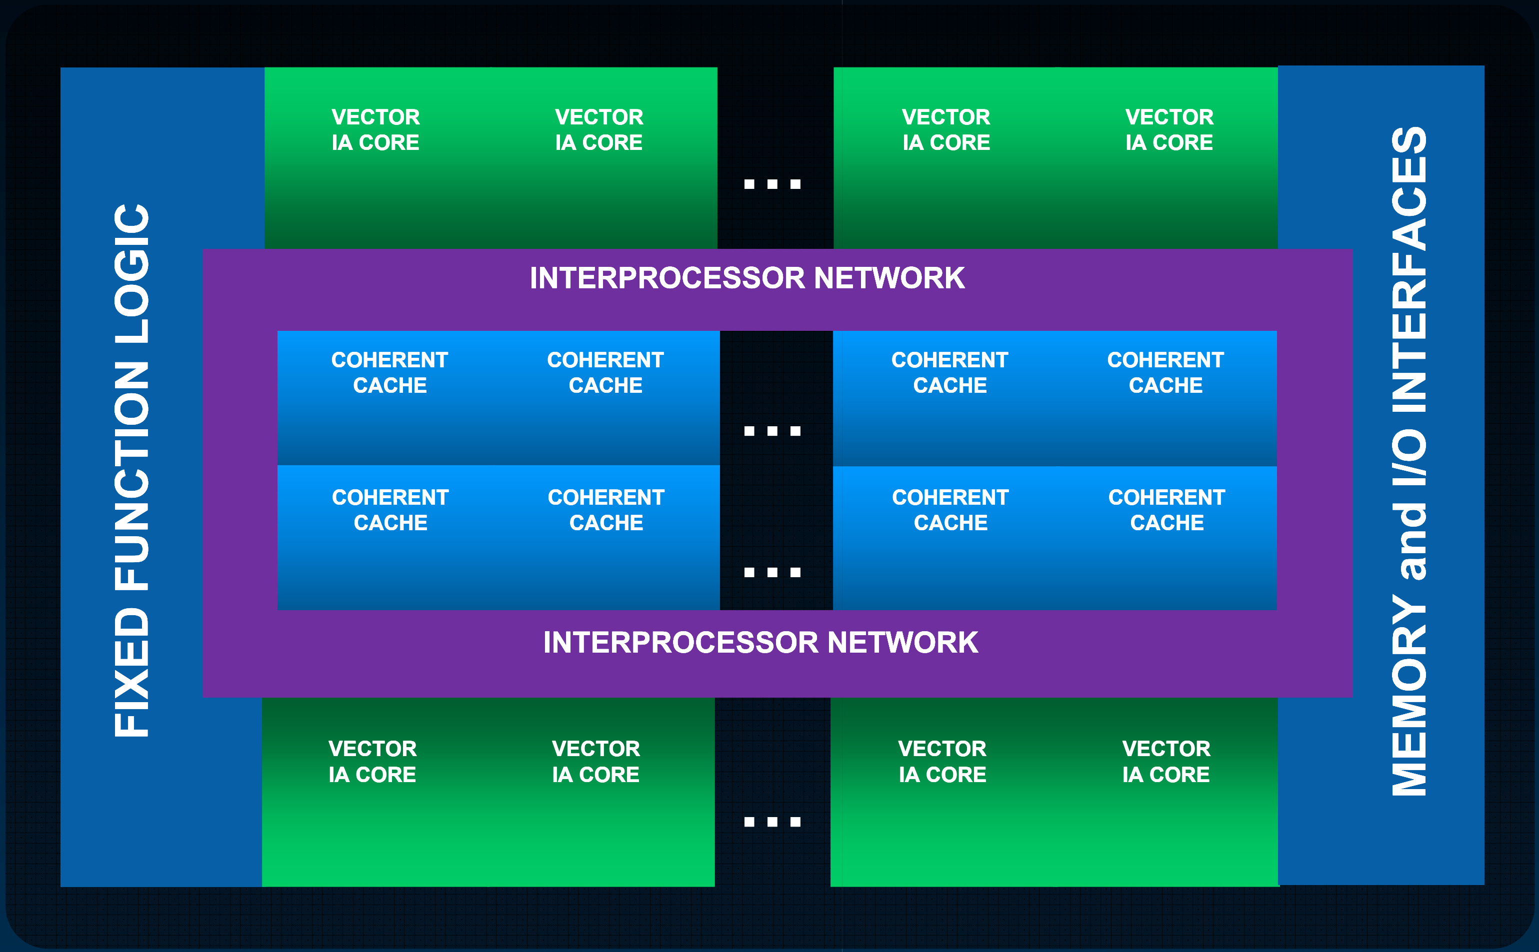

Multiple x86 cores
- In-order, short pipeline
- Multi-thread support

Supports virtual memory

16-wide vector units (512b)
Extended instruction set

Fully coherent caches

1024-bit ring bus
GDDR5 memory

Standard Intel Architecture Programming and Memory Model

For illustration only.
Future options subject to change without notice.

Aubrey Isle Core

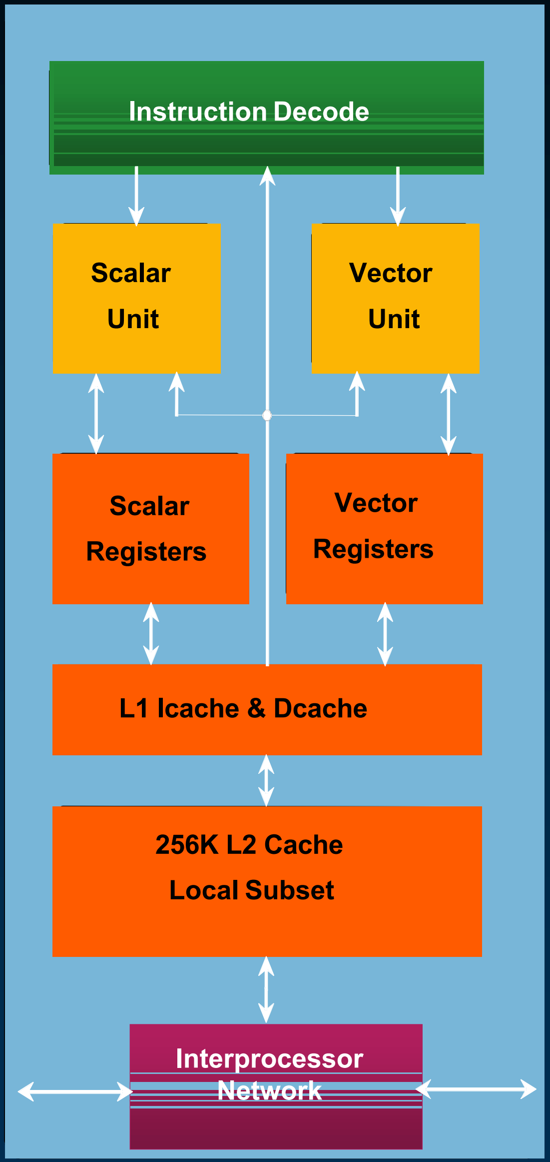

The Aubrey Isle co-processor core:

- Scalar pipeline derived from the dual-issue Pentium® processor
- Short execution pipeline
- Fully coherent cache structure
- Significant modern enhancements such as multi-threading, 64-bit extensions, and sophisticated pre-fetching.
- 4 execution threads per core
- Separate register sets per thread
- Supports IEEE standards for floating point arithmetic
- Fast access to its 256KB local subset of a coherent L2 cache.
- 32KB instruction cache per core
- 32KB data cache for each core.

Enhanced x86 instructions set with:

- Over 100 new instructions,
- Wide vector processing operations
- Some specialized scalar instructions
- 3-operand, 16-wide vector processing unit (VPU)
- VPU executes integer, single-precision float, and double precision float instructions

Interprocessor Network

1024 bits wide, bi-directional (512 bits in each direction)

“Knights Ferry” Software Development Platform

Software Development Platform

Growing availability through 2011

Aubrey Isle Co-Processor

Up to 32 cores, up to 1.2 GHz

Up to 128 threads at 4 threads / core

Up to 8MB shared coherent cache

Up to 2 GB GDDR5

Bundled with Intel HPC SW tools

The “Knights” Family

Future Knights Products

Knights Corner

1st Intel® MIC product

22nm process

>50 Intel Architecture Cores

Knights Ferry

Software Development Platform

Future options subject to change without notice.

Co-design and Changing the HPC Paradigm

- 5+ years ago dual-core microprocessors arrived
  - Moore's Law is powering multicore processors
  - Exacerbate data movement problem for HPC
  - Growing performance gap
- Co-design – an implicit statement that multi-core processors need redesign to address HPC performance gaps
- We assume new hardware capabilities will also benefit mainstream computing
  - Sandia can play a key role in Crossing the Chasm . . .

The Unfair Advantage

- As the first driver of the 917 race car, Donohue proved to Porsche that his team was not like other race teams
- The Unfair Advantage he enjoyed was based on his ability to communicate with Porsche engineers on their terms
  - Not just a race car driver, Donohue was also a Mechanical Engineer
  - Donohue was directly involved in the development of the Porsche 917
- Sandia's interest in serial #1 HPC systems is to help develop Intel® MIC architecture for our applications

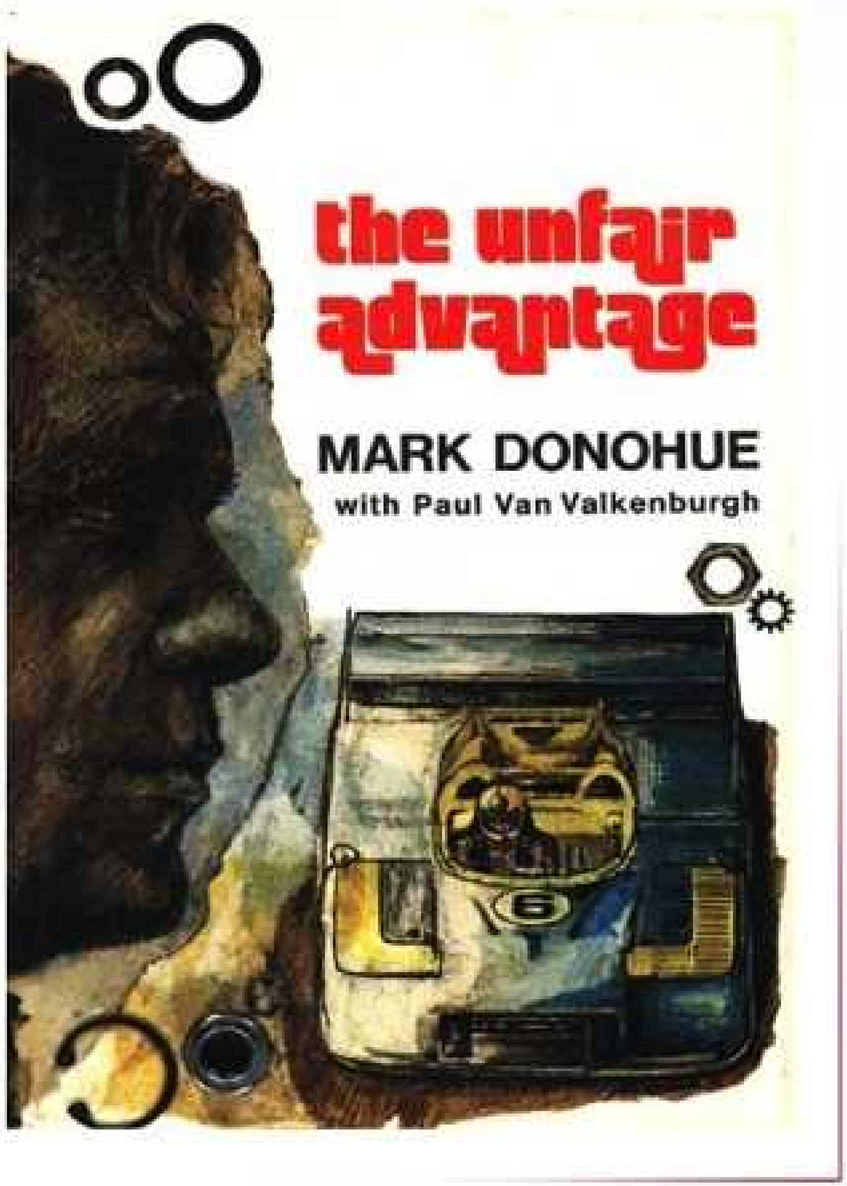

The Issue / Our Challenge:
Commodity processor adoption of capabilities for HPC

- **The MPP HPC paradigm, while based on X86 processor designs, <u>never influenced those designs</u>**
- **How can HPC co-design innovations be integrated into future X86 processor designs?**
- **Collaboration to help develop Intel® MIC architecture for scalability of Sandia and NNSA/ASC applications**
  - ***Arthur* is also a testbed to understand how HPC requirements can influence commodity processor designs**