WADE SHEN

PROGRAM MANAGER
DARPA/MTO

HIERARCHICAL IDENTIFY VERIFY EXPLOIT (HIVE)

THE HIVE PROGRAM

Program Objective: *A graph processing stack yielding a 1,000x increase in computational efficiency over GPU solutions for DoD graph analytics.*

HIVE

Static and streaming analytics for trillion edge graphs

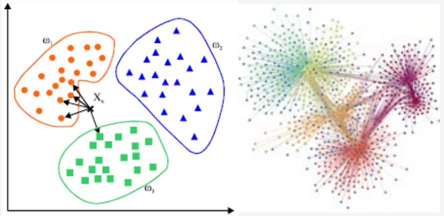

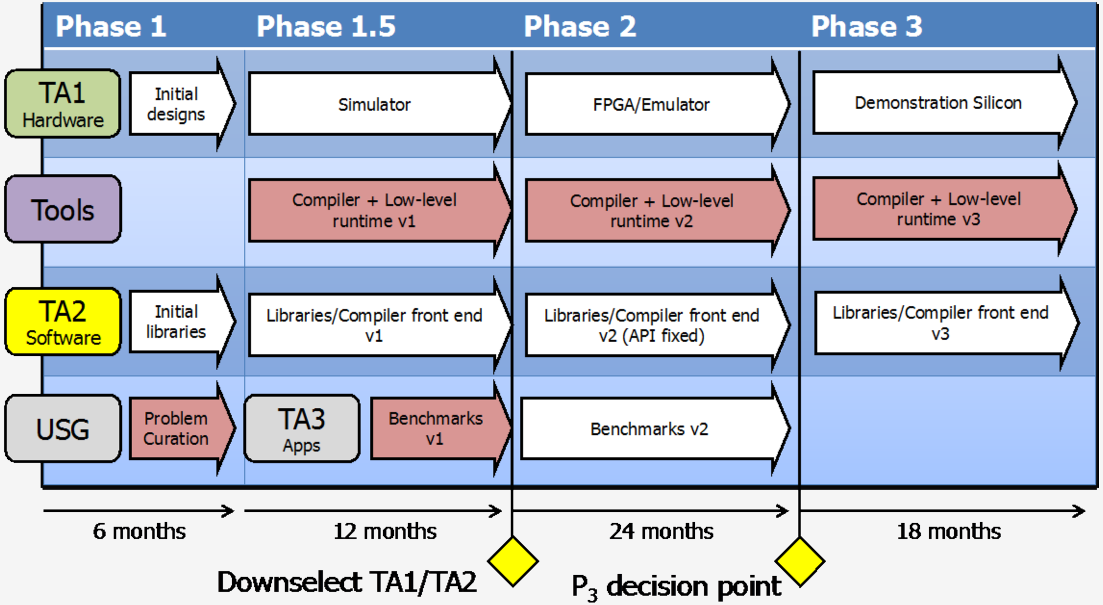

MANY DOD PROBLEMS ARE GRAPH PROBLEMS

Intelligence	Graph algorithm
Geolocation inference	Label Propagation
Persona de-aliasing	Stochastic Graph Matching
Target prioritization	Personalized PageRank
Seeded target discovery	Vertex Nomination
Organization discovery	Local Community Detection
Detection of money laundering	Query by Example
Leadership detection	Role prediction

Operations	Graph algorithm
Target audience discovery	Snowball Sampling
Network mapping	Community Detection
Network infrastructure discovery	Graph Projection
Cyber attack detection	Anomaly Detection

Support	Graph algorithm
Logistics/route plan	Hierarchical Hub Labeling
HR selection	K-nearest neighbors
Ops planning	Trellis search

GRAPH VS. NUMERIC WORKLOADS

- The graph processing fallacy
  - GPUs/supercomputers designed for matrix math
  - All graphs = sparse matrices
  - ∴ GPUs/supercomputers process graphs well

Graph problems driven by "random accesses"

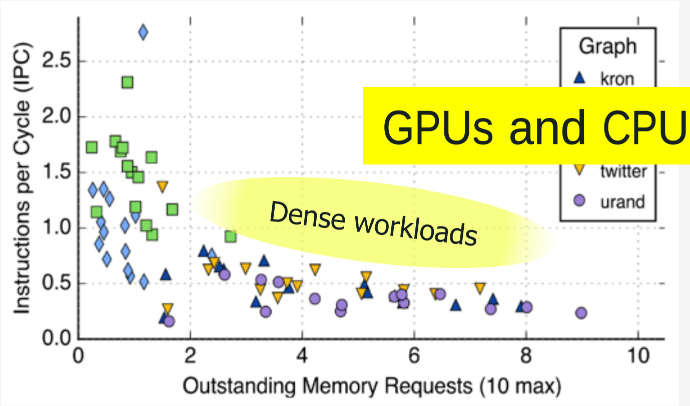

CPU/GPUs use caches to optimize locality

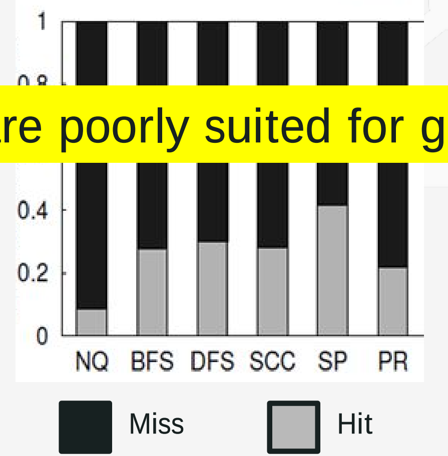

Scale poorly for data movement problems

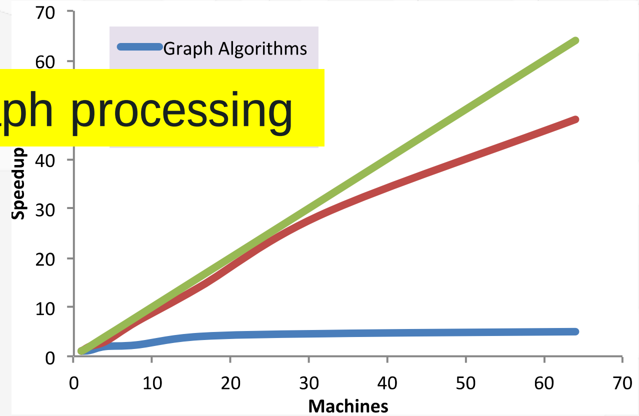

GPUs and CPUs are poorly suited for graph processing

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is Unlimited.

PROGRAM STRUCTURE

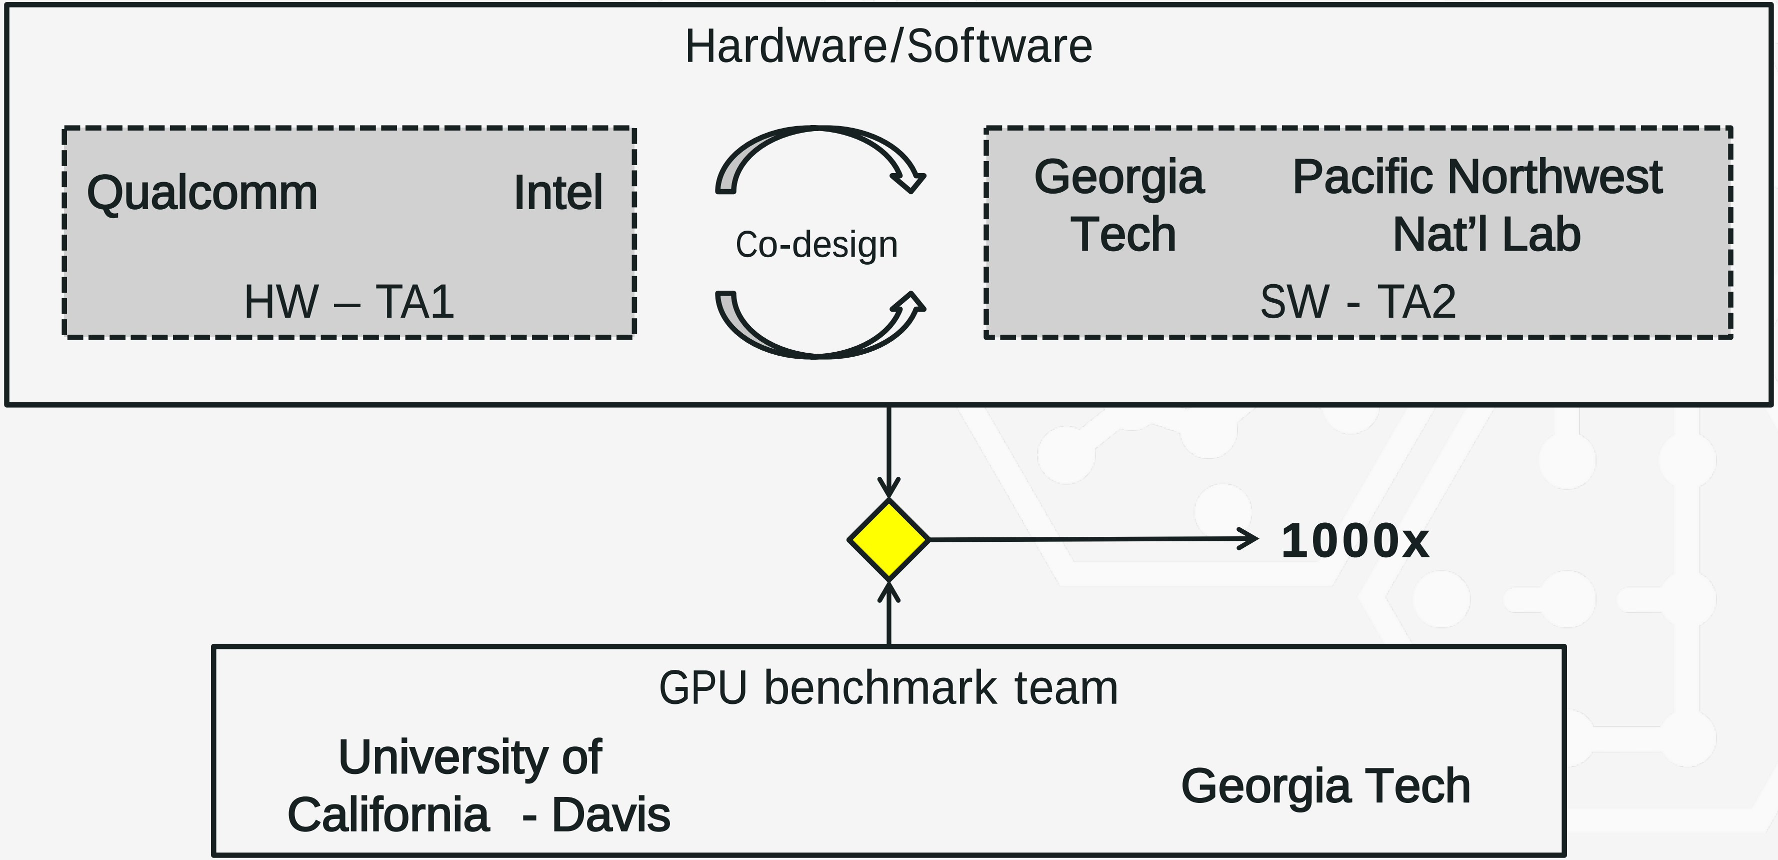

ERI ELECTRONICS RESURGENCE INITIATIVE

SUMMIT

2018 | SAN FRANCISCO, CA | JULY 23-25

JOSHUA FRYMAN

SENIOR PRINCIPAL ENGINEER, PHD

INTEL DCG, DARPA HIVE PI

THE ELECTRONICS RESURGENCE INITIATIVE

INTEL'S HIVE: FUTURE GRAPH ANALYTICS

JOSHUA FRYMAN, PHD

DARPA HIVE PI

SENIOR PRINCIPAL ENGINEER

INTEL DCG

WHAT'S SO HARD ABOUT GRAPHS . . . ?

Behavior	Dense Compute	Graph Compute
Compute intensity / arithmetic properties	Lots of computationally intensive math ops; some data massaging that's spatio-temporal friendly	Not computationally (math) intensive; mostly scheduling memory accesses and control flow problems
Cacheline access behavior	Worker threads use ~95% of their full cacheline; control threads are complicated	~50% of cachelines evicted with ≤ 16B used and ~75% with ≤ 32B used
Control flow inter-arrival behavior	Workers have long runs between branches; control threads are sufficiently predictable	~80% of branches occur inside dependent memory chains; extreme stress on pipelines and structures
Memory flow inter-arrival behavior	~65% of memory references back-to-back; excellent locality effectiveness	~65% of memory references back-to-back; nested dependent pointer chains

WHAT'S SO HARD ABOUT GRAPHS . . . ?

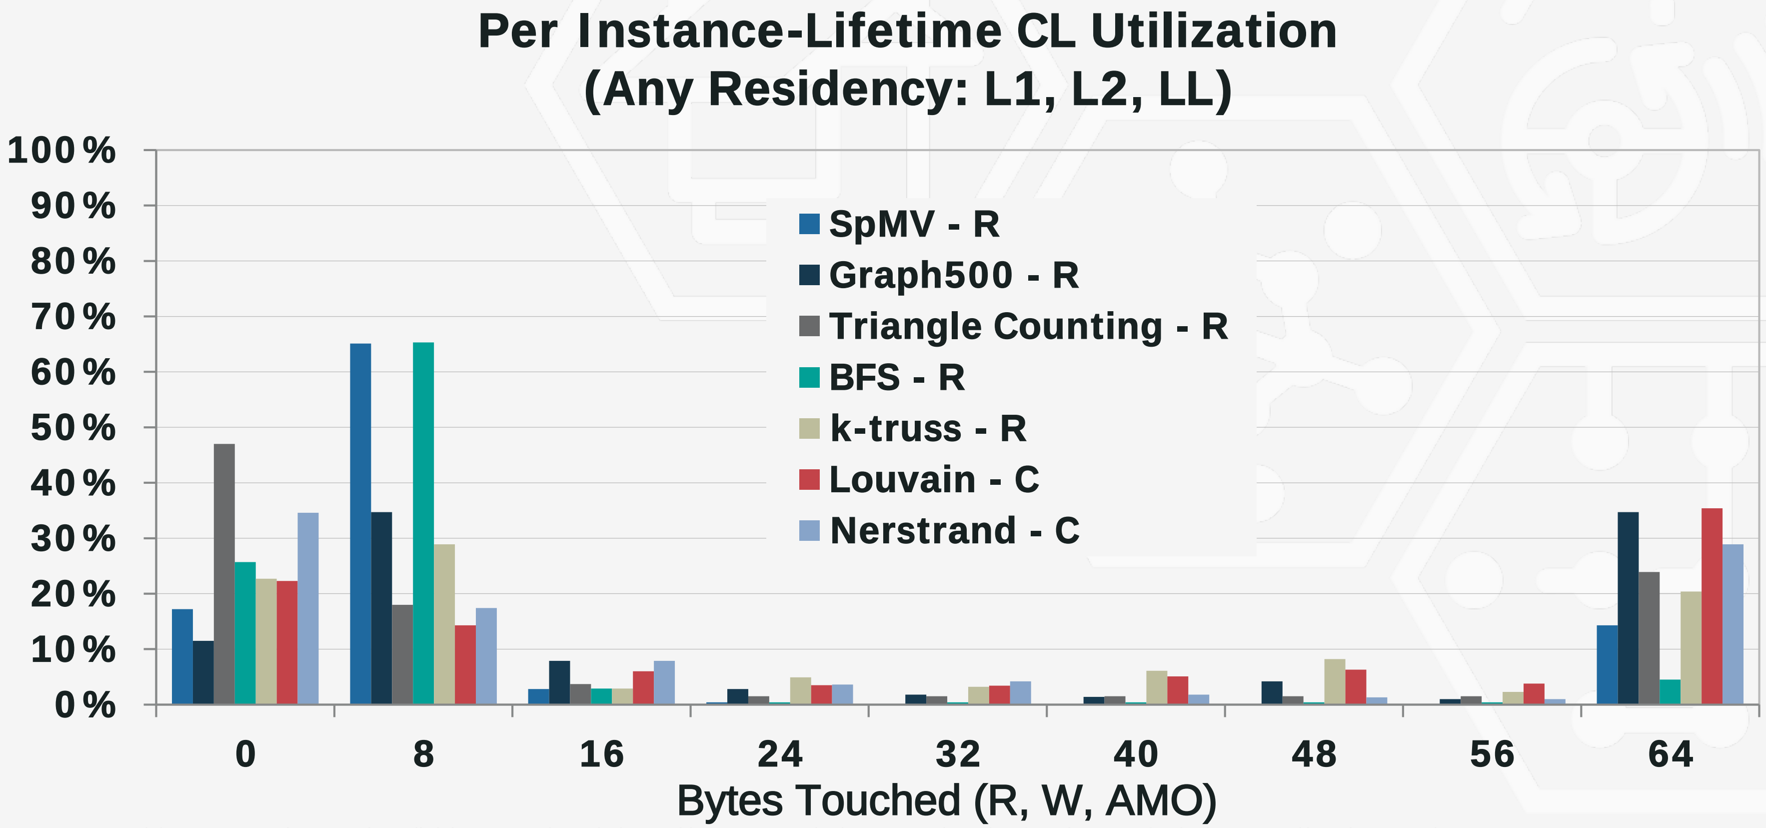

WHERE DOES "BUSINESS AS USUAL" TAKE US . . . ?

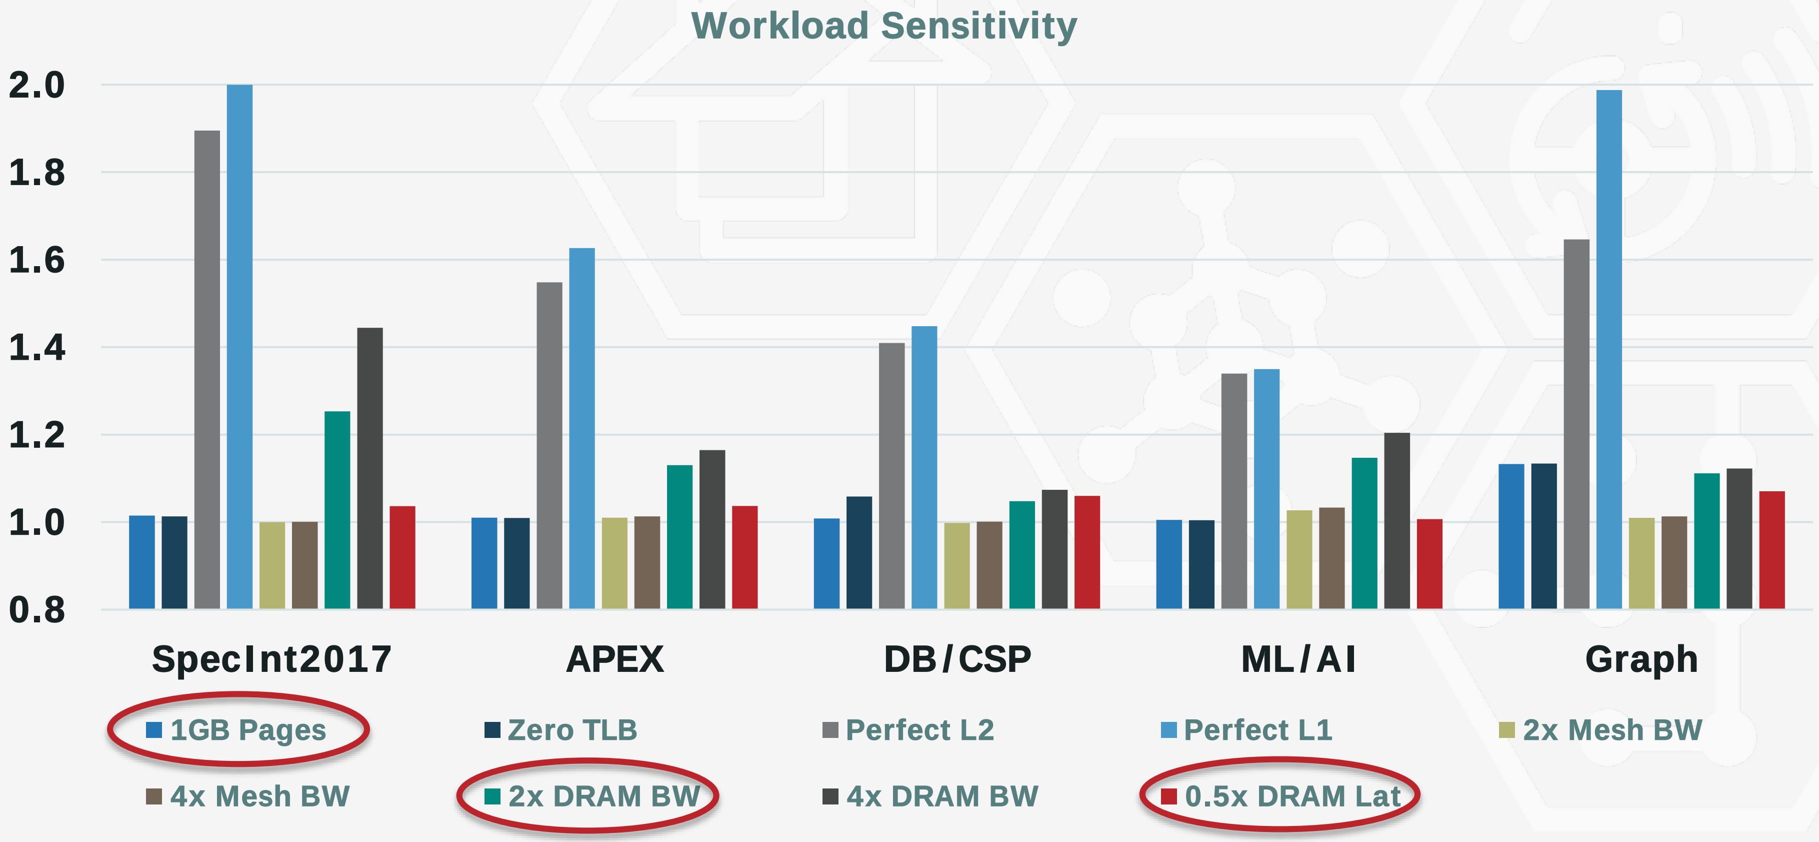

BACK TO THE BASICS – SYSTEM CO-DESIGN

Xeon Phi – *k*-Truss Graph Algorithm

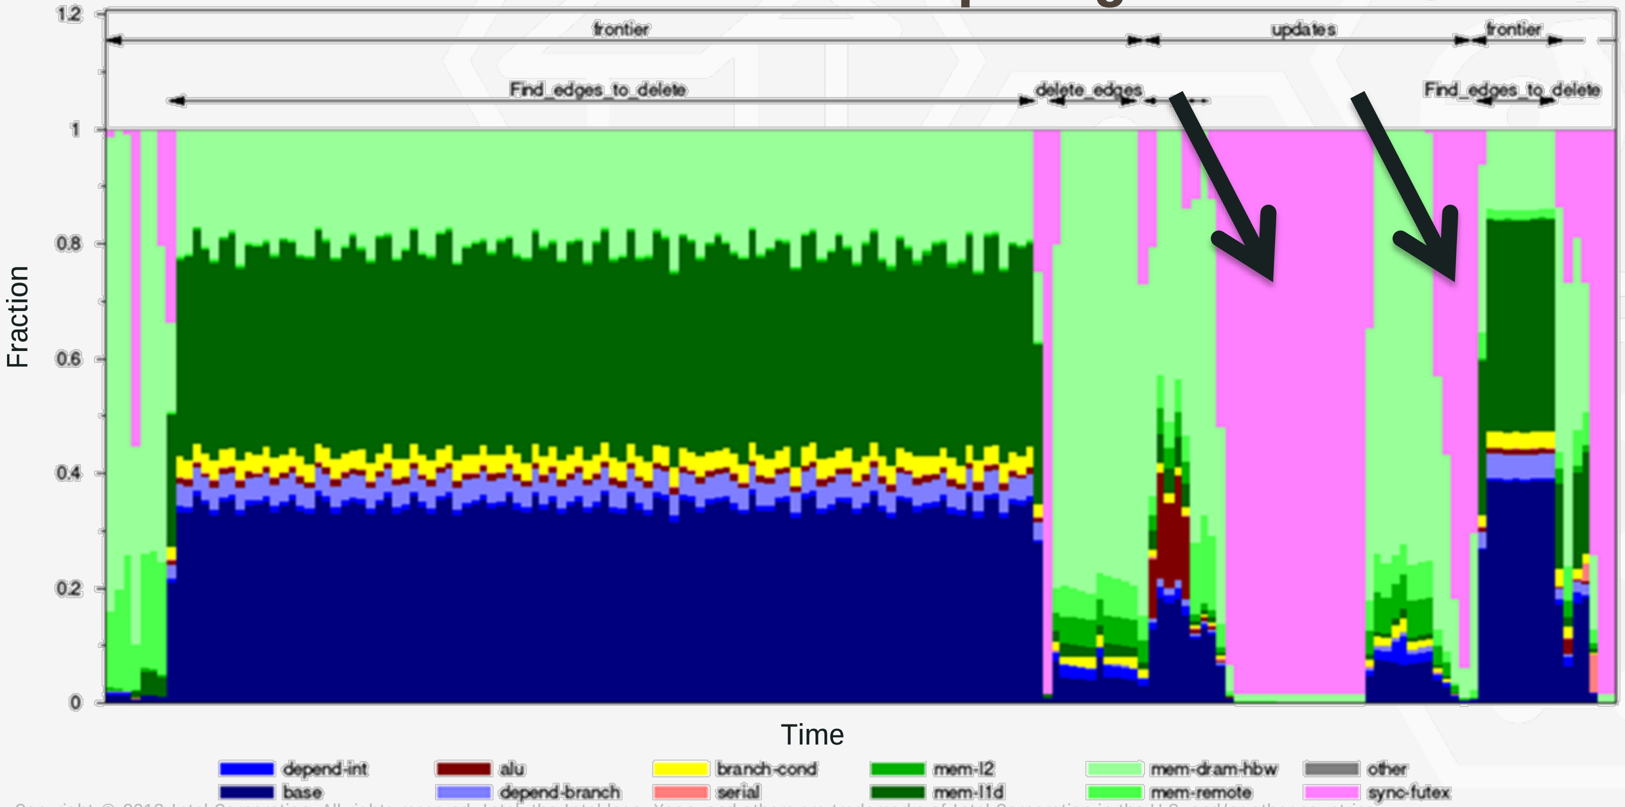

BACK TO THE BASICS – SYSTEM CO-DESIGN

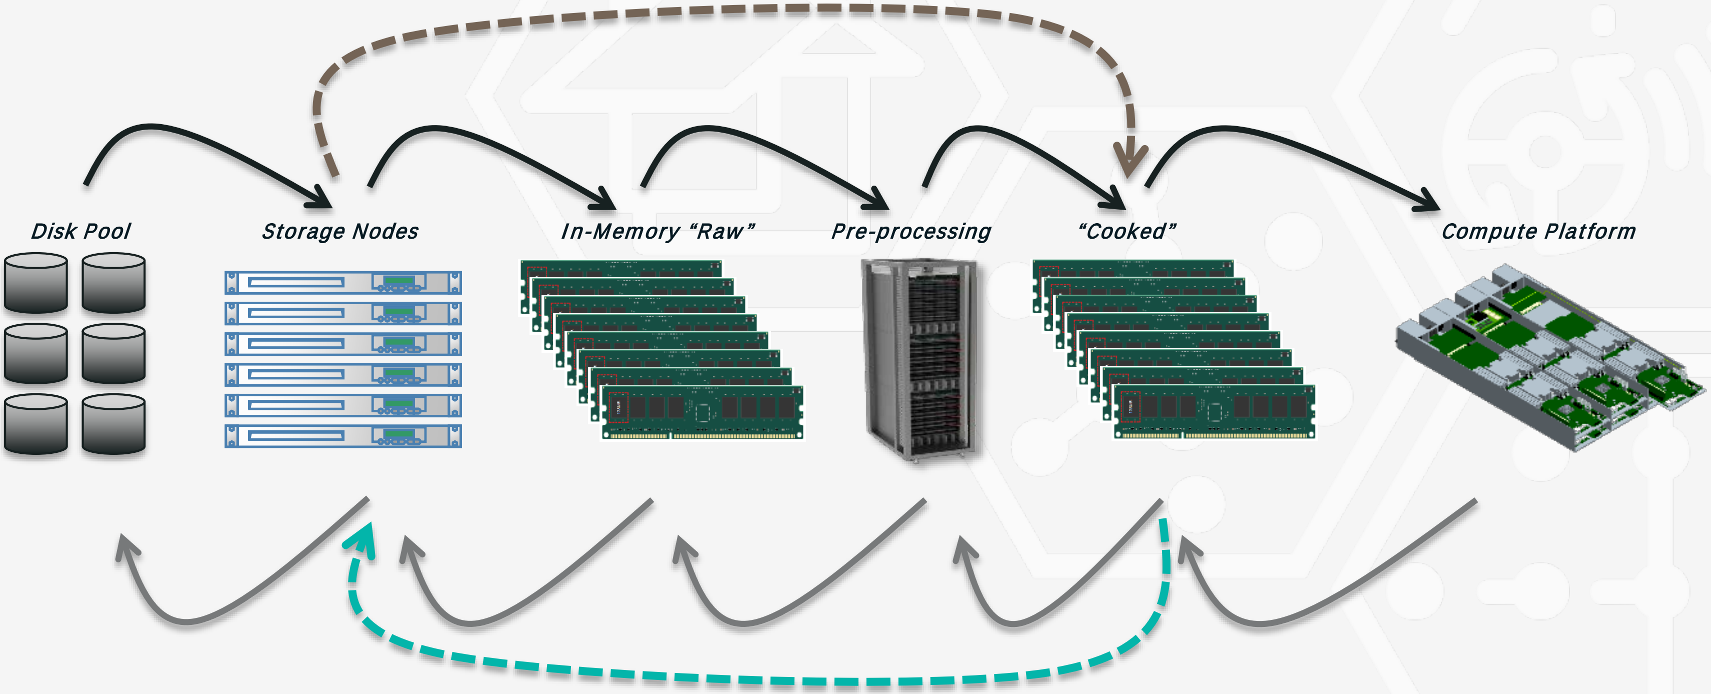

Question: What’s the ***right*** problem to optimize here?
[The Laws of {Amdahl, Gustafson, Little} Are Not Forgiving]

BUILDING A GRAPH ANALYTICS SYSTEM

- Intel is developing a HIVE solution
  - 1,000x Perf/W gain target on 100+TB
- Locality will be problematic
  - Divide and Conquer has imbalance
  - Dynamic graphs warp partitions
- Focus on a scalable platform at all levels
  - Memory, Network, and Compute
  - Target O(seconds) for 100+TB *kernels*
- Support multiple representations equally
  - Sparse matrix operations & GraphBLAS
  - Meta-data laden graph abstractions
- Opportunities to engage and partner

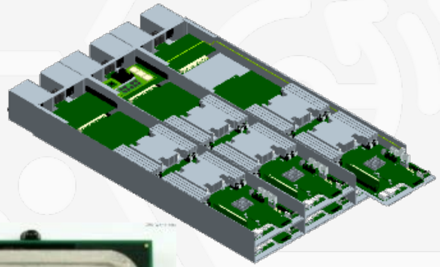

IMPACT AND OPPORTUNITIES

- Open-source Graph primitives and tools
- Actively seeking workloads and datasets
- Co-design targets with customer input

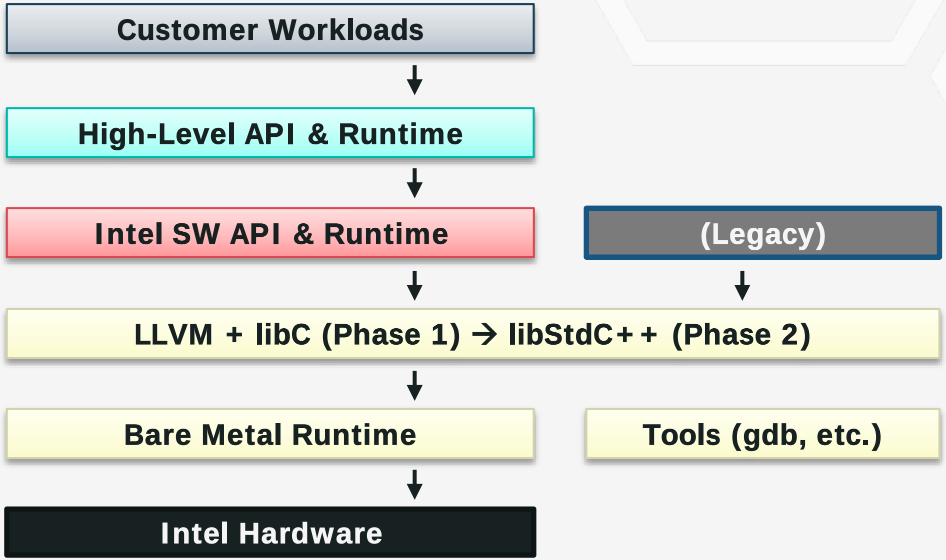

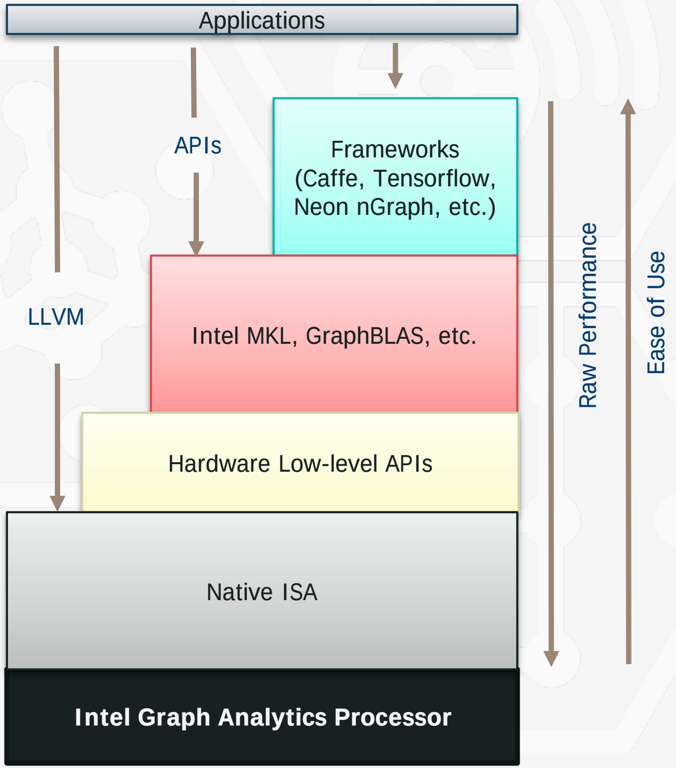

ERI ELECTRONICS RESURGENCE INITIATIVE

SUMMIT

2018 | SAN FRANCISCO, CA | **JULY 23-25**

SHEKAR BORKAR

SENIOR DIRECTOR OF TECHNOLOGY

QUALCOMM

HONEYCOMB

A GRAPH ANALYTICS PROCESSOR WITH HIGH EFFICIENCY

SHEKHAR BORKAR (PI)
MATT RADECIC (PM)

QUALCOMM INTELLIGENT SOLUTIONS, INC.
JULY 2018

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA).

PROGRAM GOALS

Performance	10 GTEPs / Node
Energy Efficiency	0.5 GTEPs / W 2 nJ / TE
Processing Efficiency	100x in Hardware 10x in Software 1000x Total
Memory Efficiency	90% both, random & sequential accesses
Demonstration	16 Nodes 160 GTEPs system
Scalability	Beyond 16 nodes to Tera TEPs

SCALABLE 160 GTEPS SYSTEM, CONSUMING < 320 WATTS

HIVE GOALS COMPARED TO GRAPH-500 (Q4-2016)

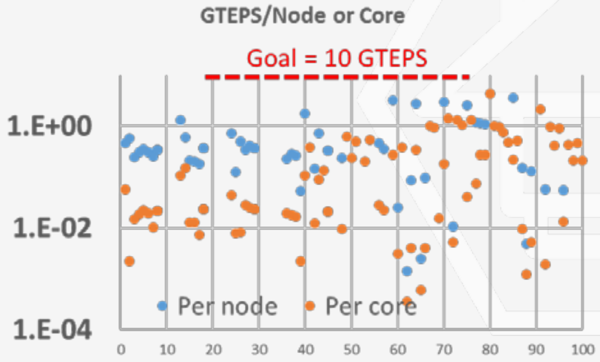

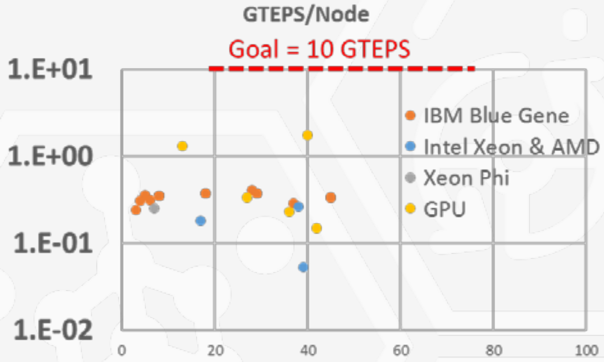

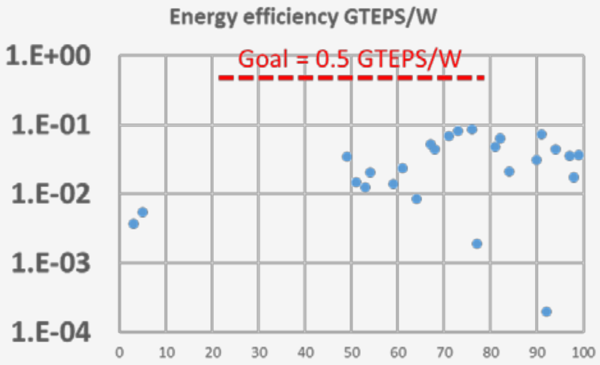

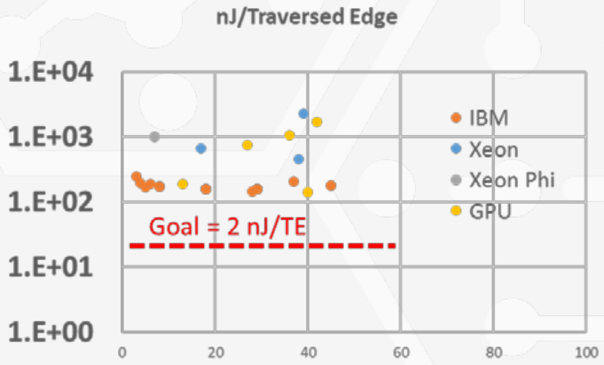

GOALS ARE EVEN HARDER CONSIDERING GRAPH-500 IS PROBABLY NOT A GOOD REPRESENTATIVE

CHALLENGES: 160 GTEPS @ 320 W

Compute

DRAM

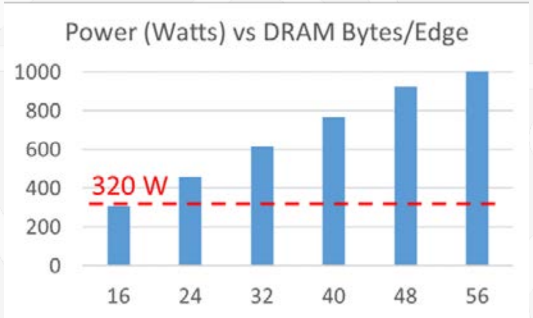

Interconnect

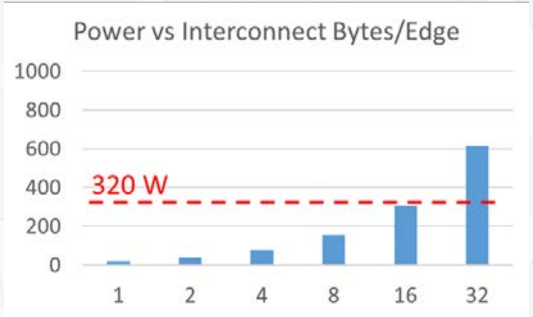

INVESTIGATION PRIORITIES: (1) INTRA-NODE DATA MOVEMENT, (2) INTERCONNECT, (3) COMPUTE

Memory Subsystem	Interconnects
Intelligent memory controller	Hierarchical & heterogeneous
Fine-grain data movement management	Simple, high-radix interconnects
Optimized data layout	Right balance of Electrical and Optical

SYSTEM ARCHITECTURE

Node

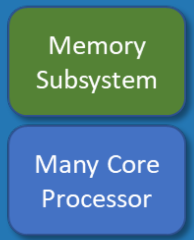

Multi-node Board

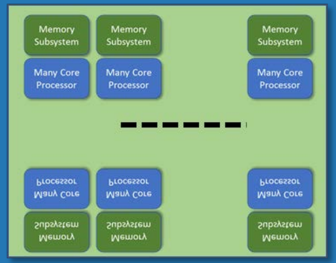

Scalable system

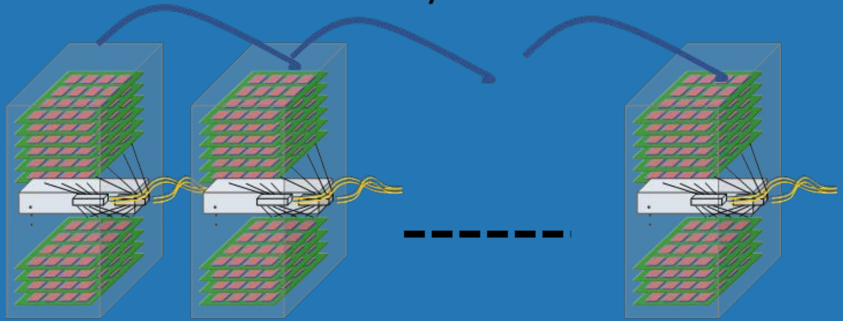

Captured in Functional Simulator

System Simulator

Fully functional and being used for analysis

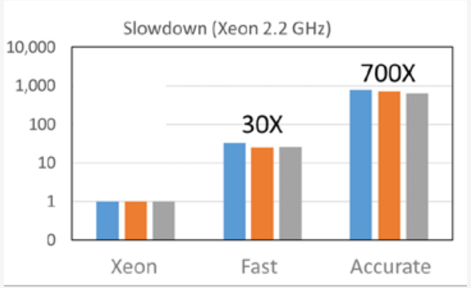

MIPS comparison		
	Xeon	HC
Kernel 1	1973	946
Kernel 2	1478	724
Kernel 3	2716	1107

0.5X performance
@
100X lower power

SENSITIVITY TO PROCESSOR FREQUENCY

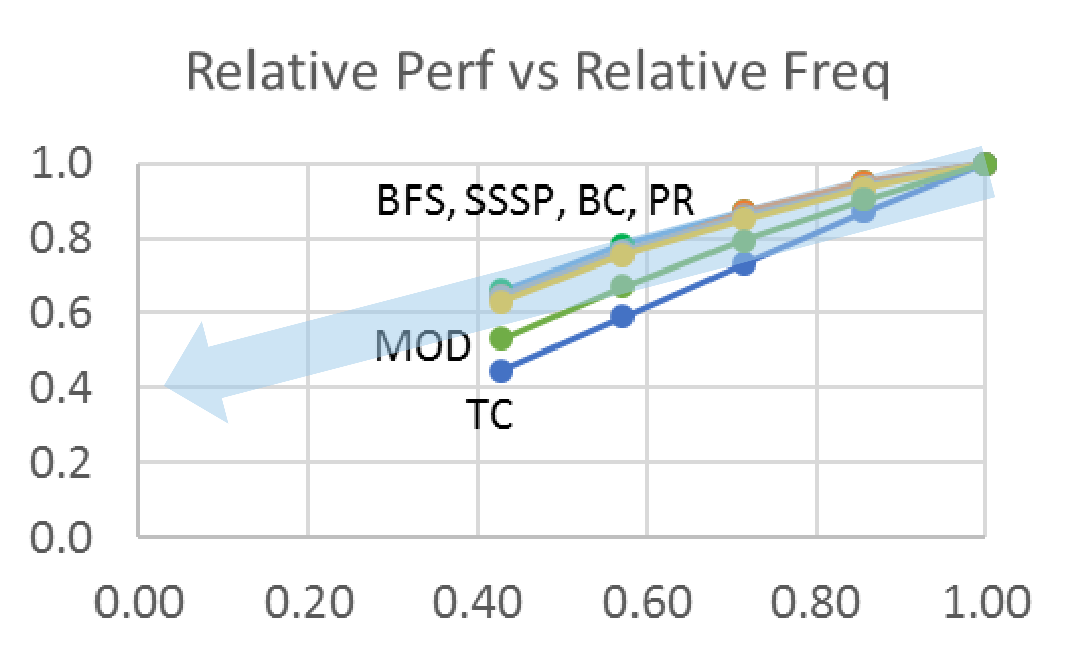

WORKLOADS ARE NOT VERY SENSITIVE TO PROCESSOR FREQUENCY

SENSITIVITY TO DATA-MOVEMENT PERFORMANCE

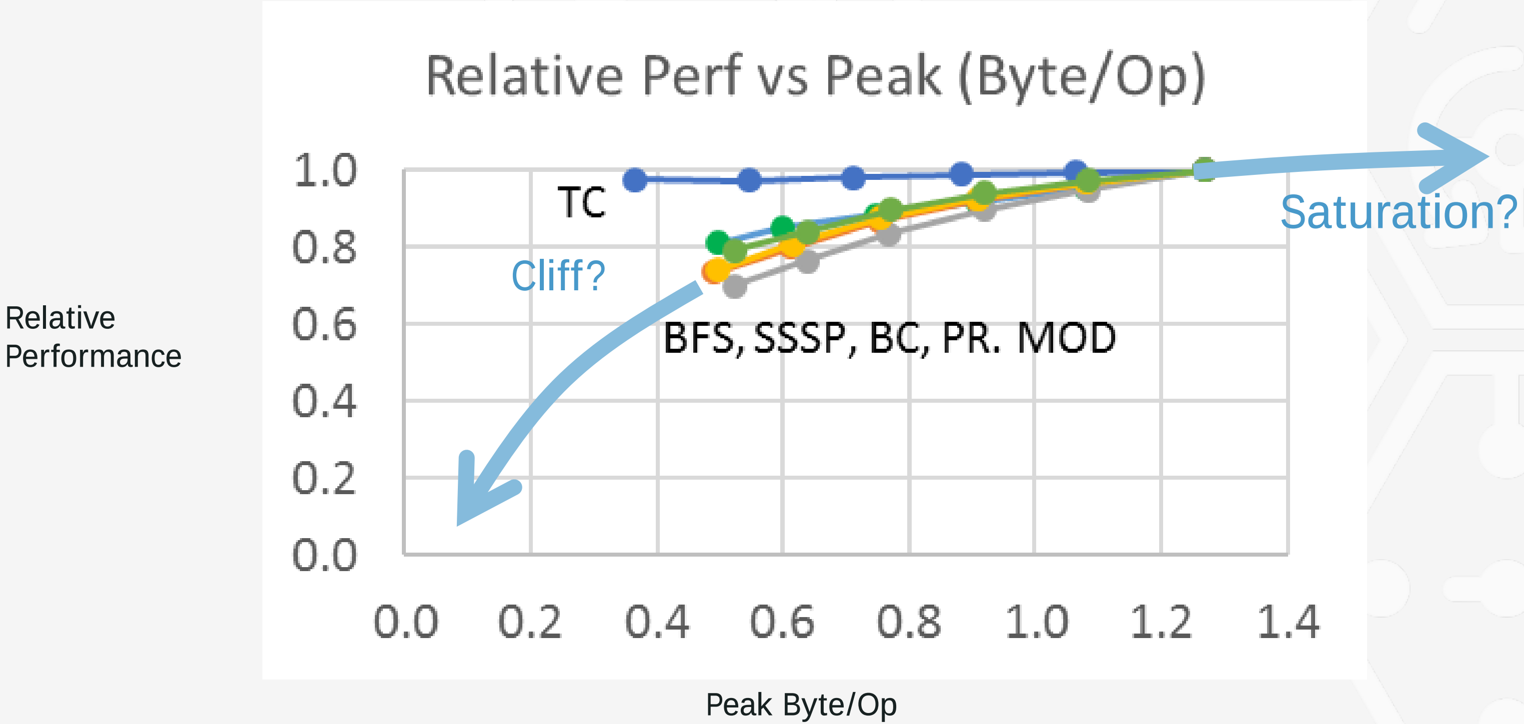

WORKLOADS ARE MORE SENSITIVE TO DATA-MOVEMENT PERFORMANCE

MULTI-THREADED WORKLOAD BEHAVIOR (NODE)

12 threads:
Performance improves with BW

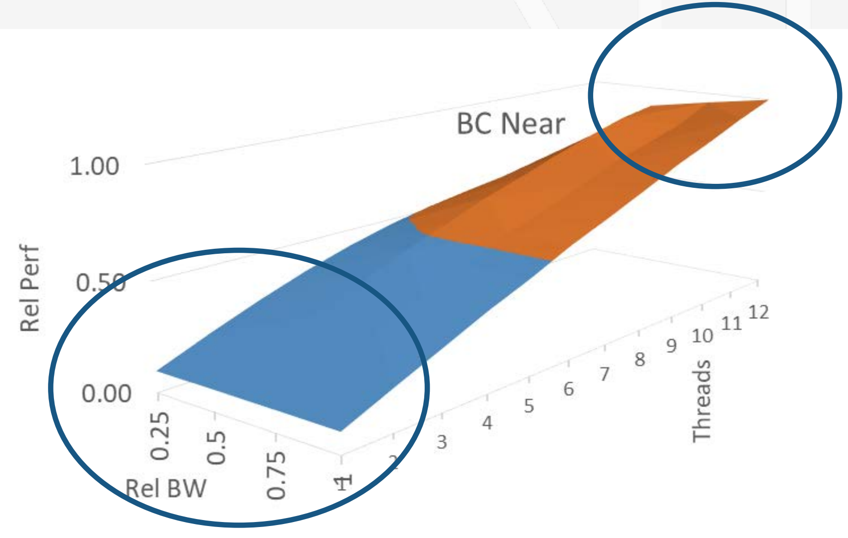

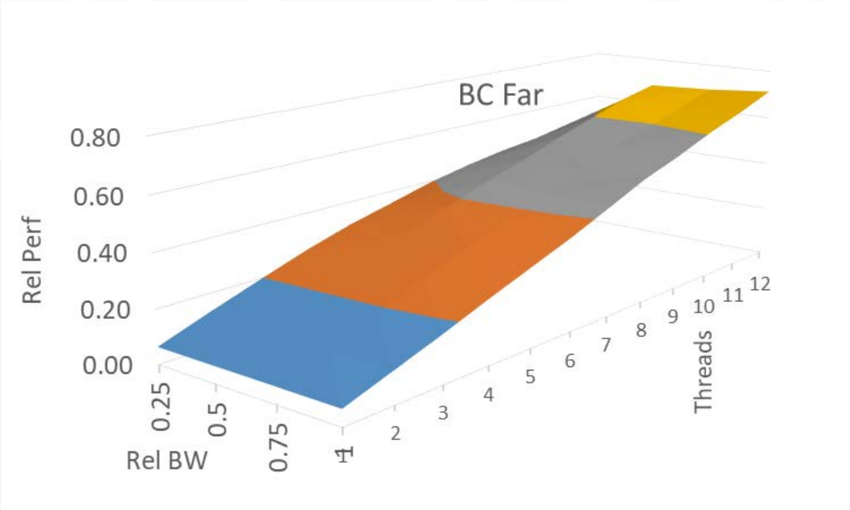

Single thread:
Minimal performance change with BW
BW overprovisioned by the platform

Same behavior with 50% higher latency

SIMULATED PERFORMANCE, ENERGY, POWER

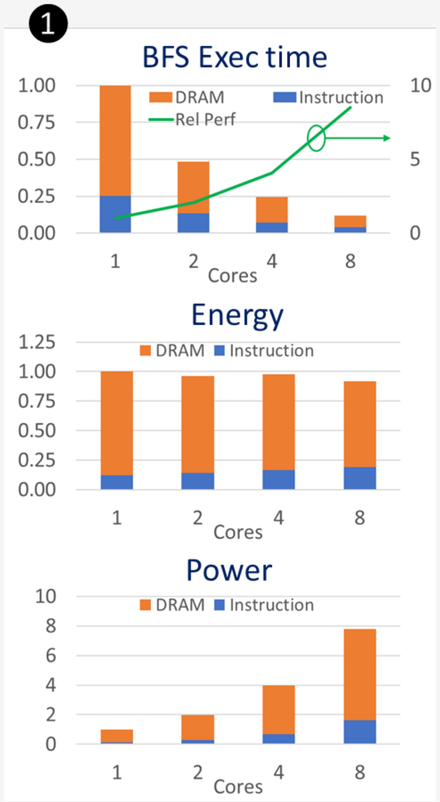

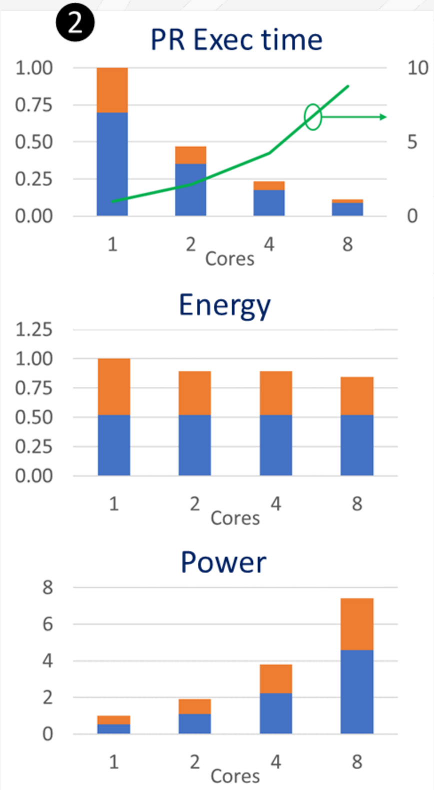

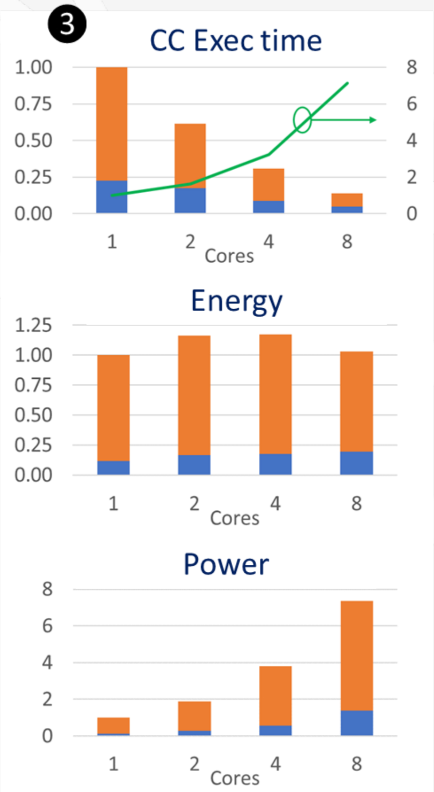

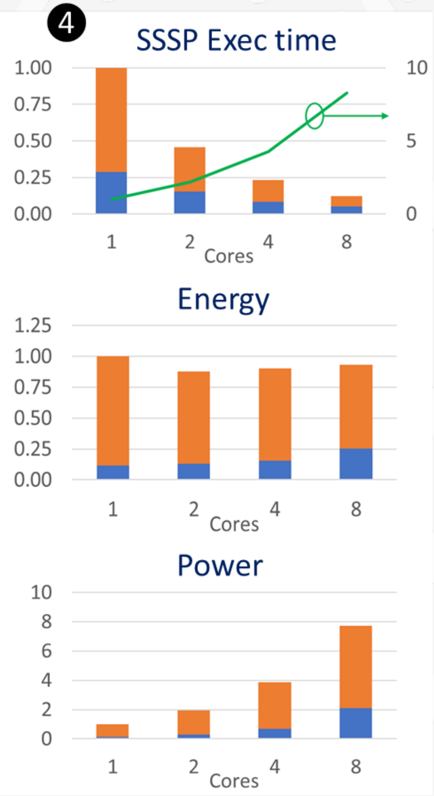

SUMMARY

- DARPA-hard goals, yet achievable!
  - Simulation based workload analysis shows data movement dominates
    - Not much by compute

Therefore...

SYSTEM DESIGN MUST BE OPTIMIZED FOR DATA MOVEMENT!

ERI ELECTRONICS RESURGENCE INITIATIVE

SUMMIT

2018 | SAN FRANCISCO, CA | **JULY 23-25**