Delivering Innovations for HPC Systems

The Project 38 Journey

Eric Cheng

Quick Facts

Goals

- Technical collaboration between NSA, DOE Office of Science, NNSA
- Deliver architecture innovations for high-performance computing (HPC) systems
- Develop enduring partnerships

Value of Investment

- Buy: broaden capabilities in commercial ecosystem
- Build: reduce cost of purpose-built designs
- Use: increase user/design productivity

More Information

- https://www.nitrd.gov/documents/HPC-Performance-Improvements-Project-38.pdf

Timeline

US Government Investment Crucial

Unique Requirements

- Computation vs. communication balance
- Software productivity challenges

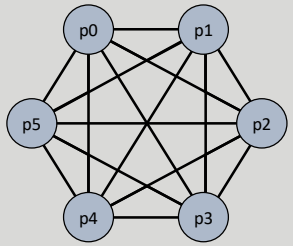

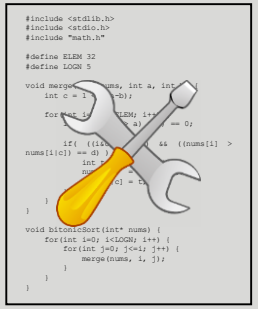

Market Pressures

- Commercial drivers (e.g., Google, Facebook, etc.)
- Narrowing domains (e.g., Cerebras)

[Google]

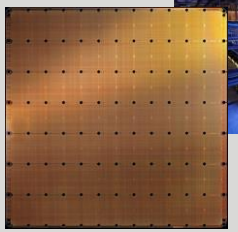

[Cerebras]

Foreign Investment

- Fugaku
- TaihuLight

[RIKEN]

[Xinhua: NSC]

Timeline

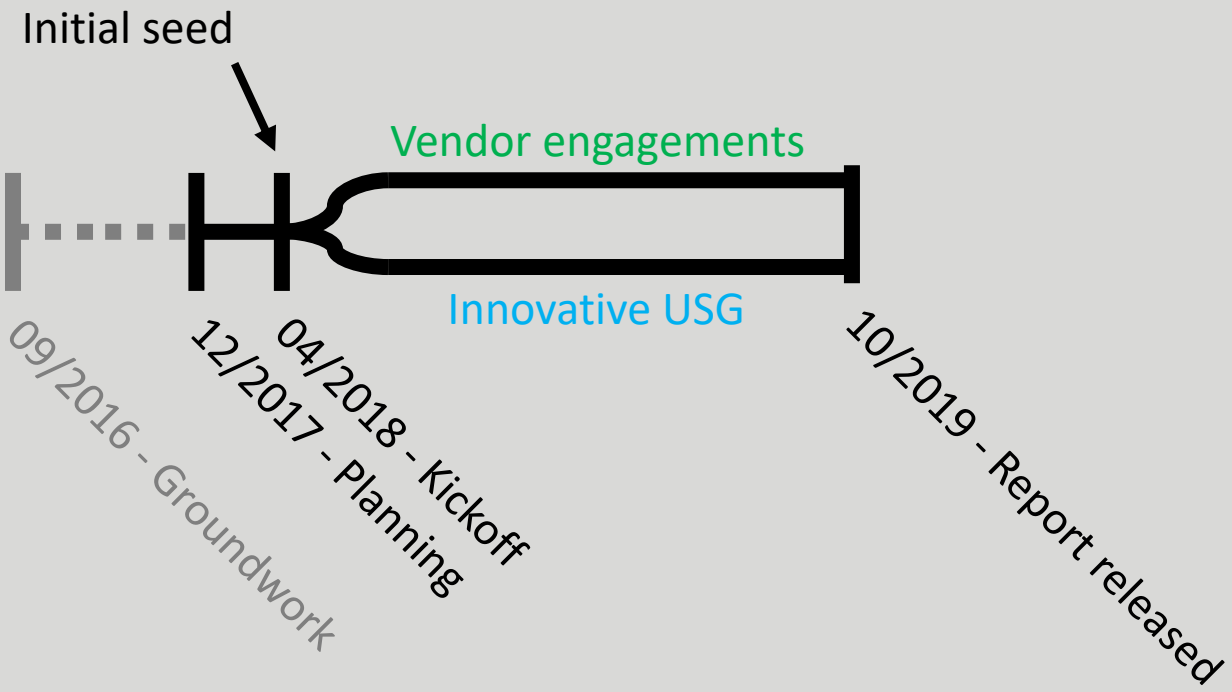
Initial seed
Vendor engagements
Innovative USG
09/2016 - Groundwork
12/2017 - Planning
04/2018 - Kickoff
10/2019 - Report released

“Org Chart”

U.S. DEPARTMENT OF ENERGY | Office of Science

National Nuclear Security Administration

Vendor 1

- Argonne National Laboratory
- LPS The Laboratory for Physical Sciences
- Berkeley Lab
- National Security Agency

Vendor 2

- Argonne National Laboratory
- LPS The Laboratory for Physical Sciences
- Berkeley Lab
- Brookhaven National Laboratory
- Lawrence Livermore National Laboratory
- Sandia National Laboratories
- National Security Agency
- DARPA

Vendor 3

- Lawrence Livermore National Laboratory
- LPS The Laboratory for Physical Sciences
- Los Alamos National Laboratory
- Oak Ridge National Laboratory
- Pacific Northwest National Laboratory
- Sandia National Laboratories
- National Security Agency
- DARPA

Vendor 4

- Argonne National Laboratory
- LPS The Laboratory for Physical Sciences
- Berkeley Lab
- Pacific Northwest National Laboratory
- National Security Agency

Parallel tracks: vendor engagements + innovative USG design

Initial Drivers (Manageable Scope)

Mix of applications

- Fast Fourier Transform (FFT)
- High-performance Geometric Multigrid (HPGMG)
- HipMER (bioinformatics)
- Kripke (radiation transport)
- Mini-tensor Contraction (MTC)
- Sparse matrix trisolve

Mix of hardware features

- Fixed function accelerators
- Message queues
- Recode engine
- Scatter/gather memory controller
- Word-granularity scratchpads

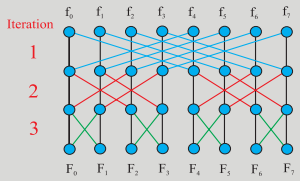

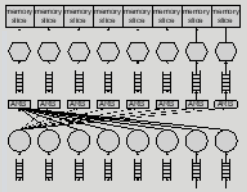

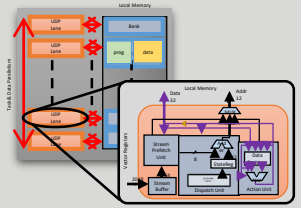

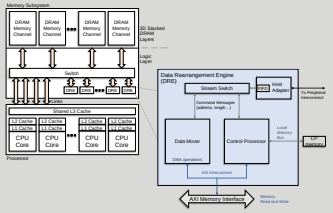

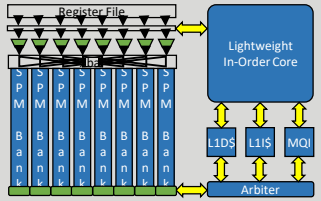

Impact: 40% to over 10× performance improvement

Timeline

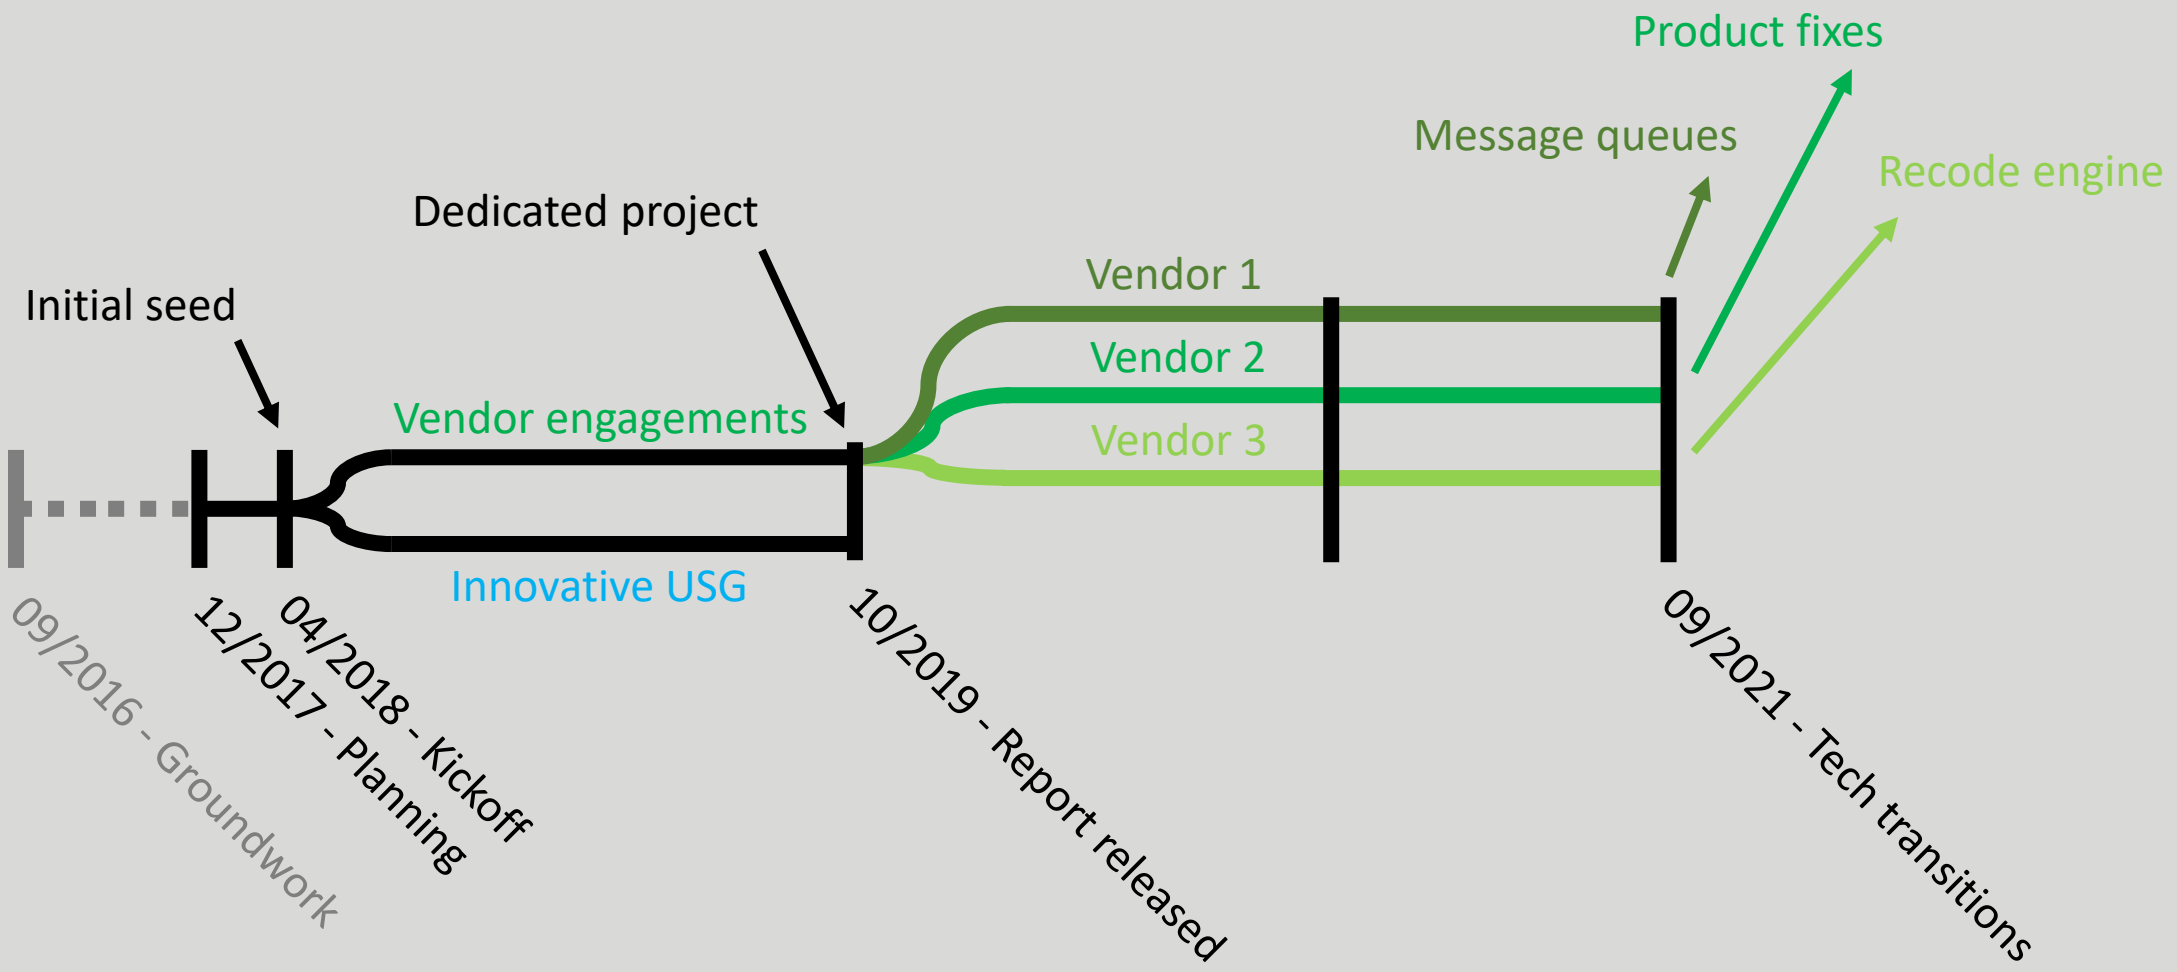
Product fixes
Message queues
Recode engine
Dedicated project
Vendor 1
Initial seed
Vendor 2
Vendor engagements
Vendor 3
Innovative USG
09/2016 - Groundwork
12/2017 - Planning
04/2018 - Kickoff
10/2019 - Report released
09/2021 - Tech transitions

Message Queues

Enabling Technology

- Efficient inter-processor communication
- Reduce need for mutex

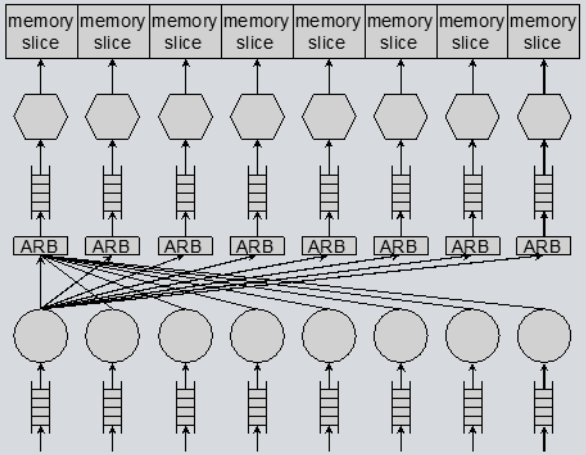

Impact

- 2-8× performance
- 5-12× latency reduction
- Vendor integration

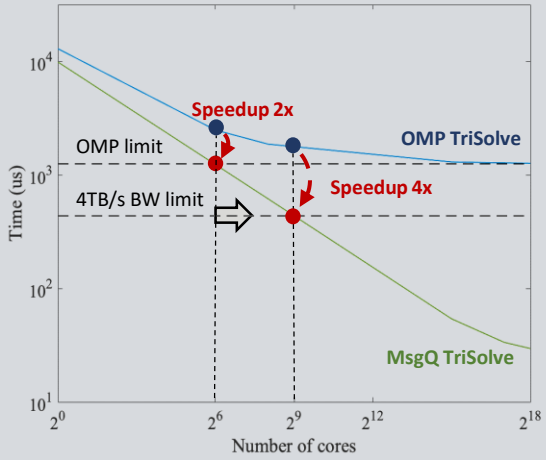

Integration

- Hardware IP available
- Straightforward intrinsics interface

1. Initialization: $x := b$; $lsum := 0$
2. **while** (I have work to do) **do**
4. **if** $(pollQ() > 0)$
5. $getQ(\&vfunc, \&K)$
3. $fmod(K) = fmod(K) - 1$ or $fmod(K) = -1$ by $\&vfunc$
26. **endif**
4. **if** $(fmod(K) = -1$ and I own this column)
5. **for** (each of my $L(I,K) \neq 0, I > K$)
6. $lsum(I, myid) = lsum(I, myid) + L(I,K)x(K)$
7. $putQ(\&id, \&vfunc, \&I)$ for id on diagonal block
18. **end for**
19. **else if** $(fmod(K) = 0$ and I own diagonal block)
20. $x(K) = x(K) - lsum(K, id)$ for each id in row K
23. $x(K) = L(K,K)^{-1}x(K)$
7. $putQ(\&id, \&vfunc, \&K)$ for each id in column K
26. **endif**
27. **end while**

Algorithm 2: Lower triangular solve $Lx = b$ with TaskQ.

Recode Engine

Enabling Technology

- High-speed, fine-grained data recoding
- Programmable transforms

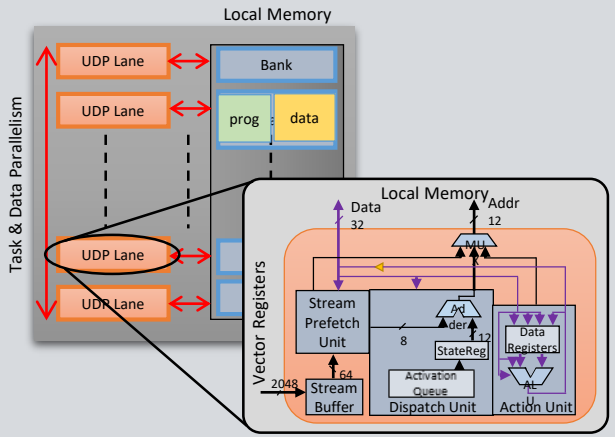

Impact

- SpMV evaluation
  - 5× performance increase
  - 8× bandwidth reduction
- Vendor integration

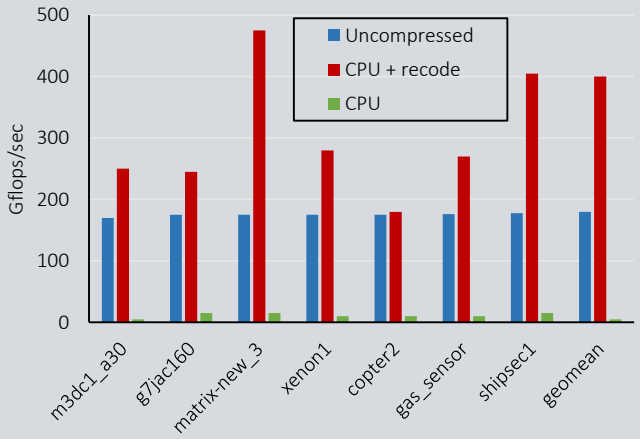

Integration

- Synthesizable RTL
- Emulation infrastructure
- LLVM compiler support
- Rich library support

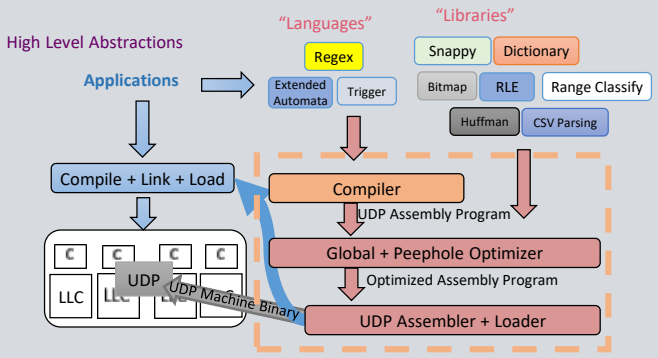

Product Design Fixes

Tensor Contraction Evaluation

	Small input	Medium input	Large input
Accelerator	2.5	3.1	3.4
x86 openMP	1.0	1.0	1.0
Speedup	**0.41×**	**0.32×**	**0.29×**

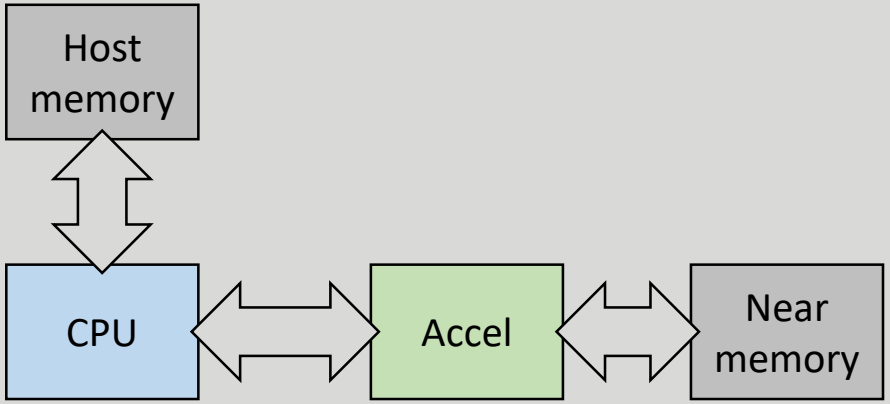

USG identifies cache structure bottleneck

USG proposes data layout fix

Result: 3x performance over x86

Timeline

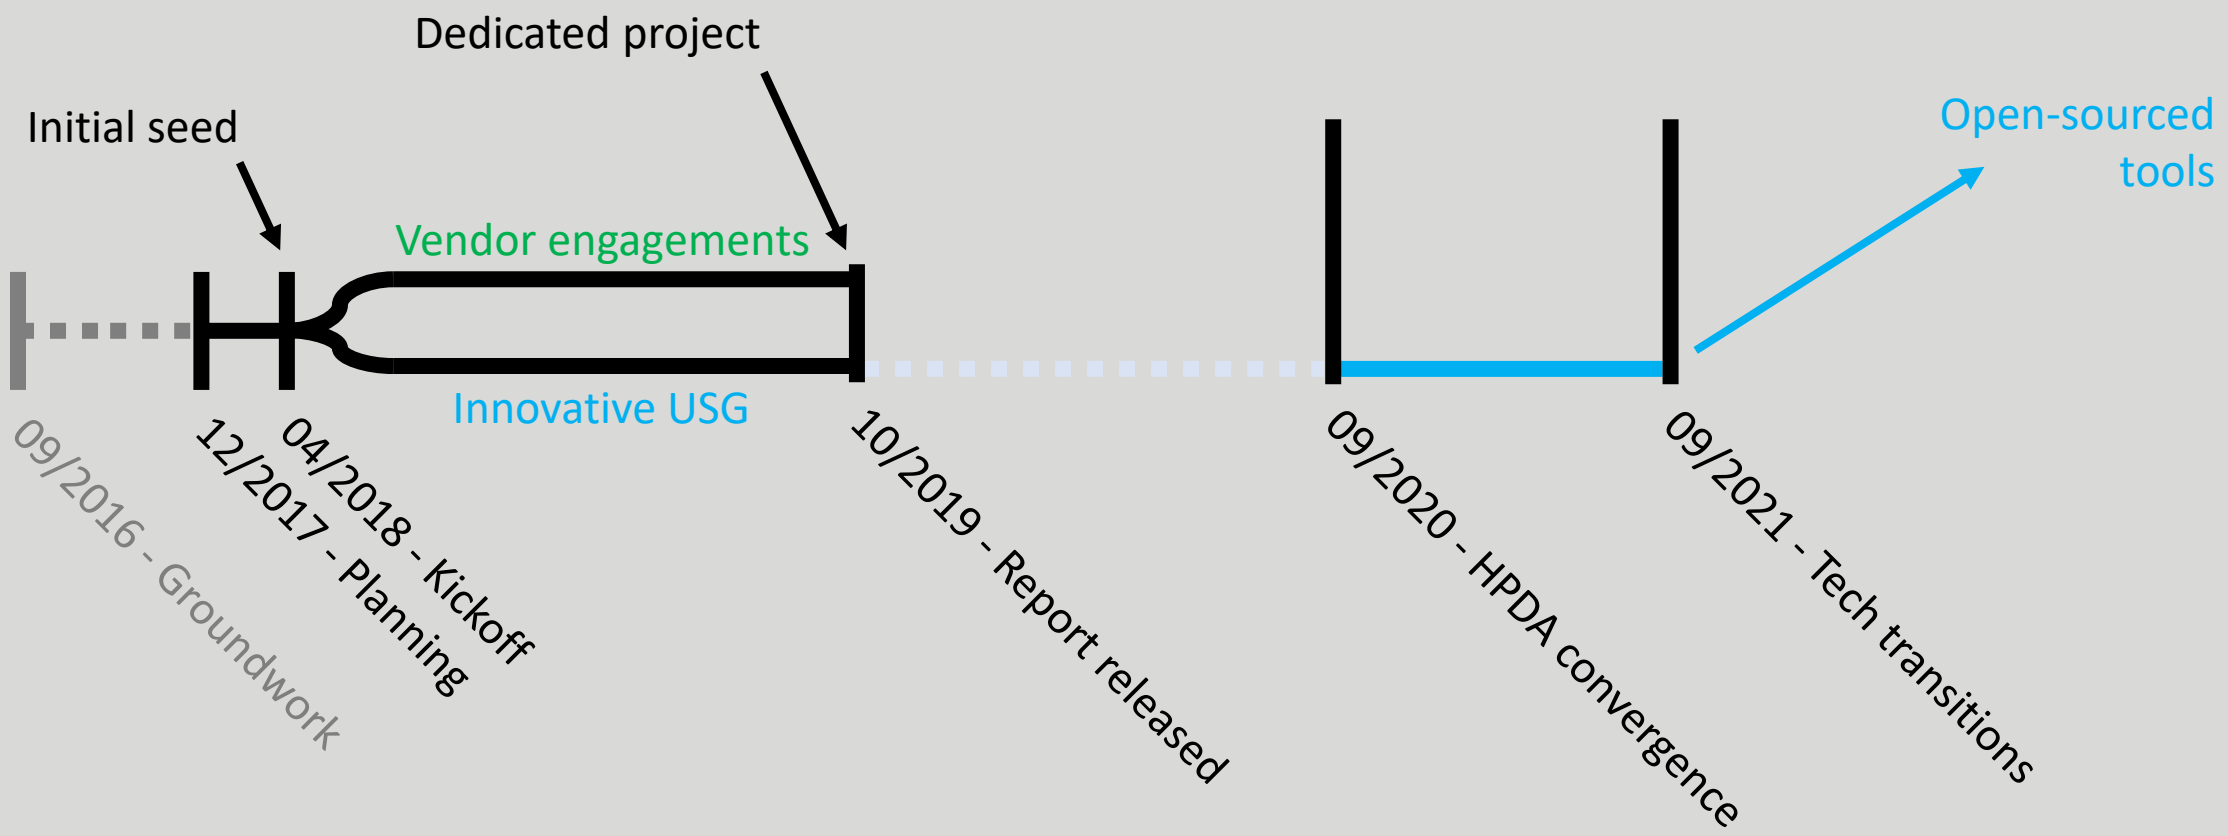
Dedicated project
Initial seed
Vendor engagements
Open-sourced tools
Innovative USG
09/2016 - Groundwork
12/2017 - Planning
04/2018 - Kickoff
10/2019 - Report released
09/2020 - HPDA convergence
09/2021 - Tech transitions

High-performance Data Analytics (HPDA) System

Mission Focused

- High-performance Data Analytics (HPDA)
- Graph analytics at scale
- Data science at scale

Key Drivers

- User productivity
- Scale, performance, energy-efficiency, cost
- Emerging technologies

Synergies

- IARPA AGILE
- PNNL Data Model Convergence

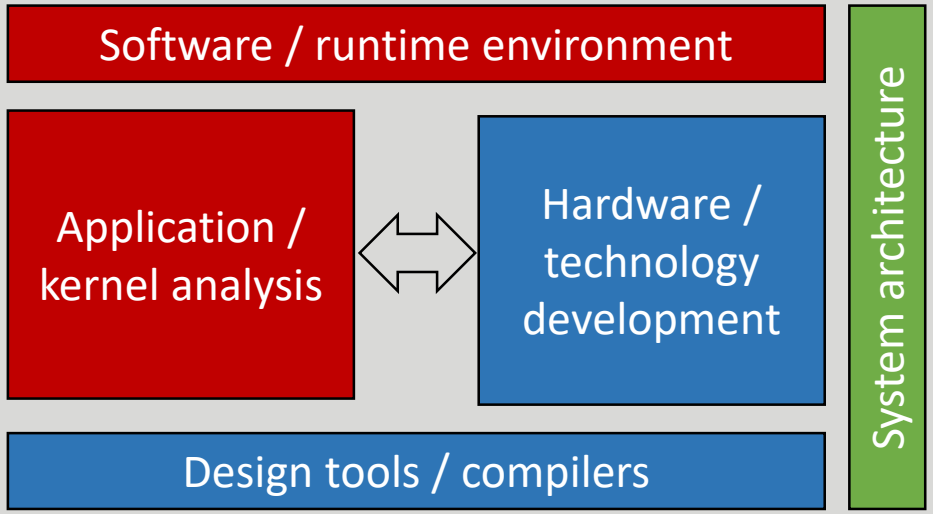

Goal: system-level design for complete end-to-end workflows

Convergence of Applications

Challenge Problems

- Breadth First Search (BFS)
- Triangle counting
- PageRank
- Global RandomAccess (GUPS)
- …

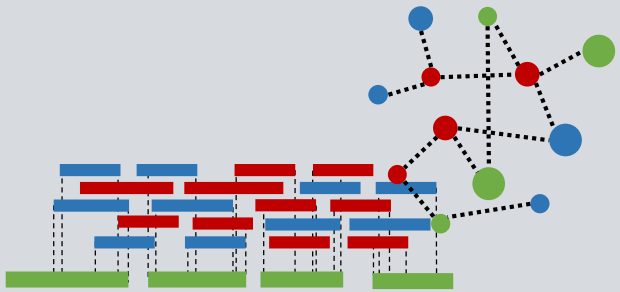

Impact

- Workload fusion (leverage frameworks): GraphBLAS, NWGraph
- Bottleneck identification and technology solutions

  SPGEMM (sparse lookup): CAM acceleration (2.4-5.1x speedup)

  SpMV (data movement): recode engine (1.1-6.3x speedup)

  SpMV (atomics): message queue (1000 cycle access reduction)

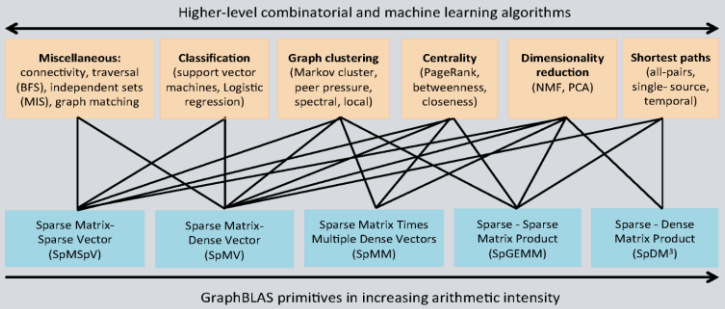

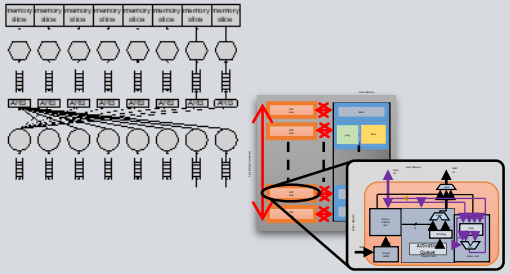

Convergence of Architectures / Technologies

Technologies

- Integration: chiplets, 3D
- Fundamental: emerging memory, photonics, PIM
- Open-source: ISA, runtimes, managers

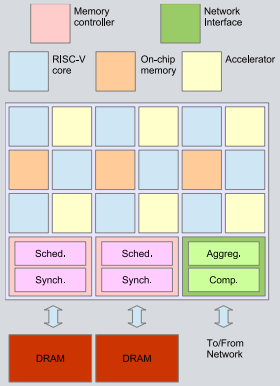

Impact

- Hotspot Aware Manager: 30% speedup (stacked memory solution)
- Minos Computing Library: 10% of hand-optimized CUDA, cross-platform port

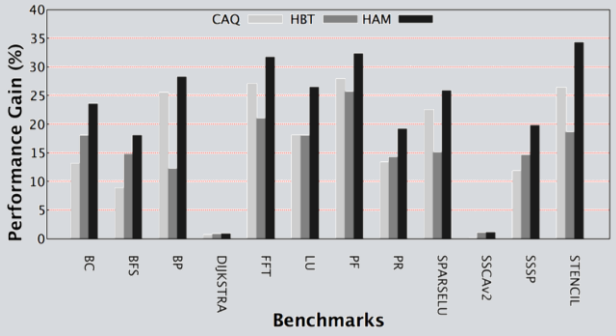

Ecosystem Impact / Transitions

- Seeded IARPA AGILE workflow development
- Informed DOE LDRD programs
- Drove development of new DOE ASCR programs
- Influenced community designs

Convergence of Tools

Hardware / Software

- EDA (design) tools
- Design productivity
- Compilers
- Software frameworks
- …

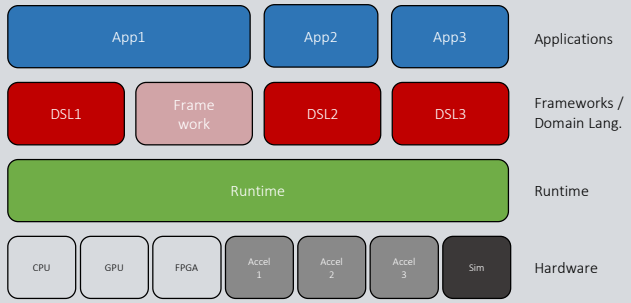

Community Impact / Transitions

- https://github.com/anikau31/systemc-clang

  SystemC to Verilog translator

- https://gitlab.pnnl.gov/sodalite/soda-frontend

  HLS (high-level synthesis) compiler for ML frameworks

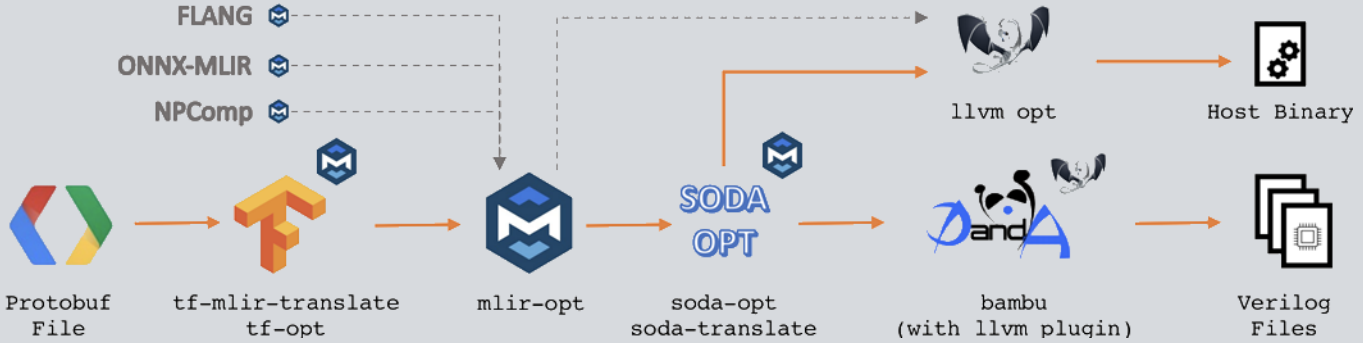

Timeline

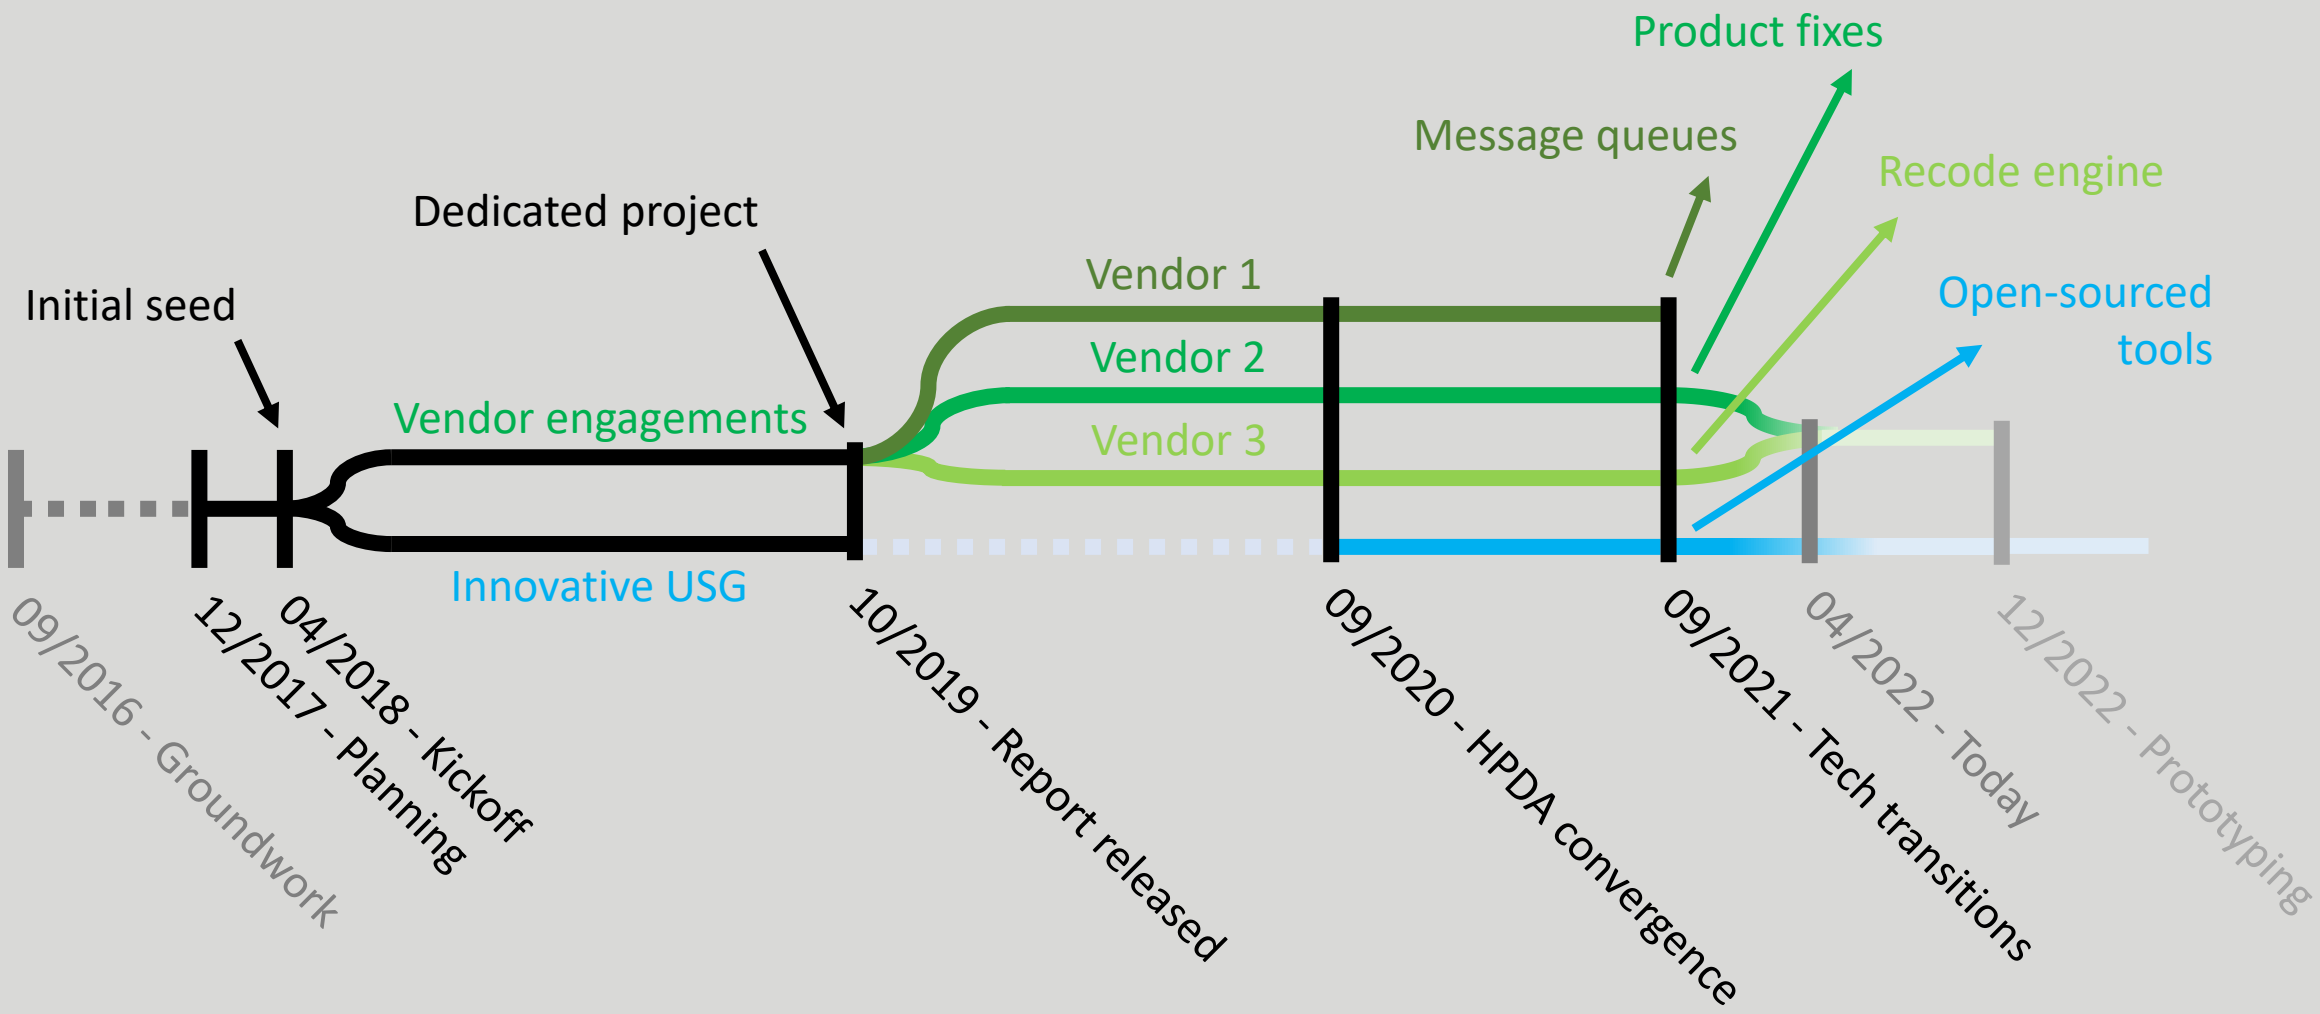
Product fixes
Message queues
Recode engine
Dedicated project
Vendor 1
Open-sourced tools
Initial seed
Vendor 2
Vendor engagements
Vendor 3
Innovative USG
09/2016 - Groundwork
12/2017 - Planning
04/2018 - Kickoff
10/2019 - Report released
09/2020 - HPDA convergence
09/2021 - Tech transitions
04/2022 - Today
12/2022 - Prototyping

Parting Thoughts

What

- Enduring cross-agency effort (NSA, DOE SC, NNSA)
- Architecture innovations for HPC

Why

- Buy: broaden capabilities
- Build: reduce cost
- Use: increase productivity

More Information

- https://www.nitrd.gov/documents/HPC-Performance-Improvements-Project-38.pdf