Future of HPC

John Urbanic

Parallel Computing Scientist

Pittsburgh Supercomputing Center

Moore's Law abandoned serial programming around 2004

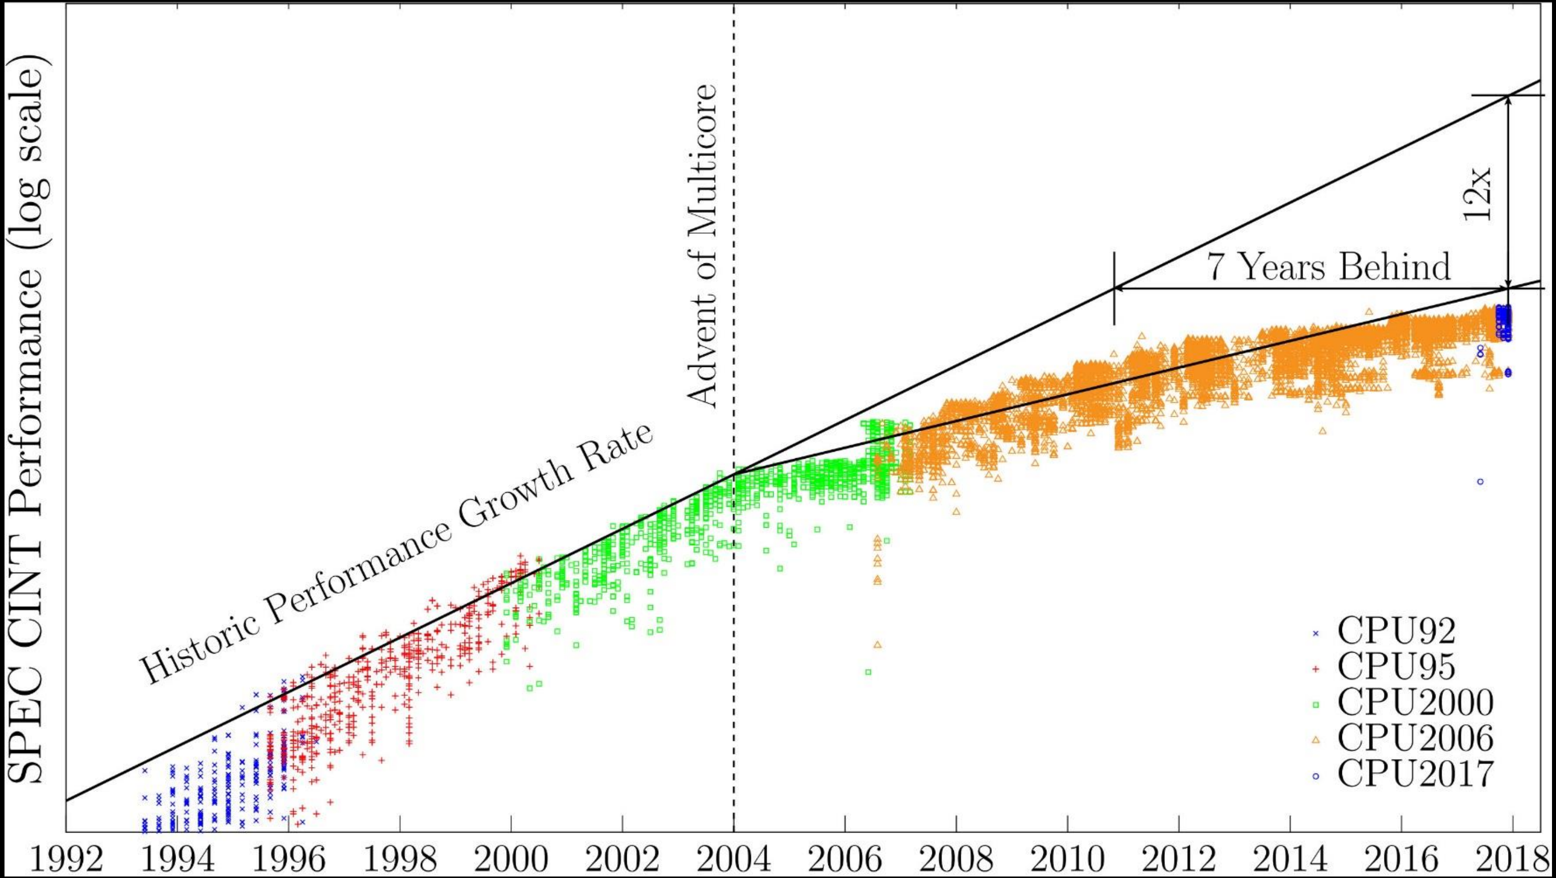

Courtesy Liberty Computer Architecture Research Group

But Moore's Law is only beginning to stumble now.

Intel process technology capabilities

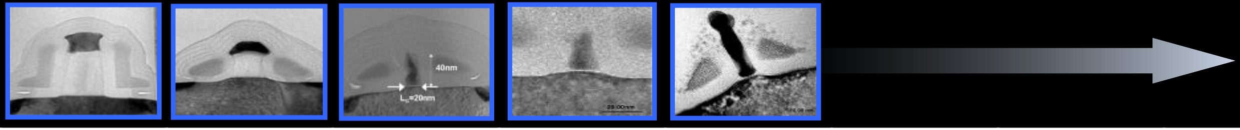

High Volume Manufacturing	2004	2006	2008	2010	2012	2014	2018	2021
Feature Size	90nm	65nm	45nm	32nm	22nm	14nm	10nm	7nm
Integration Capacity (Billions of Transistors)	2	4	8	16	32	64	128	256

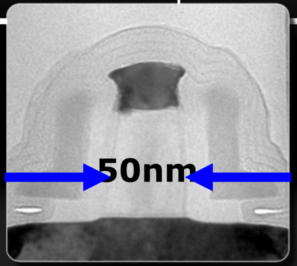

Transistor for 90nm Process

Source: Intel

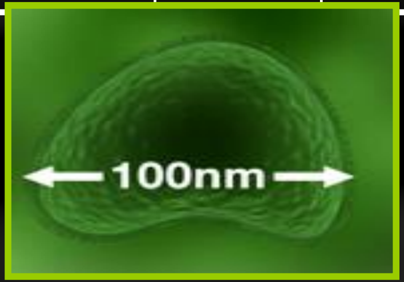

Influenza Virus

Source: CDC

…but our metrics are less clear.

After a while, “there was no one design rule that people could point to and say, ‘That defines the node name’ … The minimum dimensions are getting smaller, but I’m the first to admit that I can’t point to the one dimension that’s 32 nm or 22 nm or 14 nm. Some dimensions are smaller than the stated node name, and others are larger.”

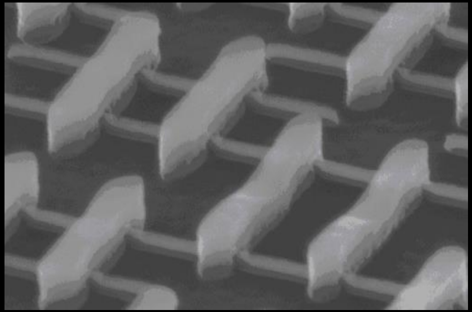

Mark Bohr, Senior fellow at Intel.

From *The Status of Moore's Law: It's Complicated* (IEEE Spectrum)

Now tradeoffs are stealing these gains.

The density and power levels on a state-of-the-art chip have forced designers to compensate by adding:

- **error-correction circuitry**
- **redundancy**
- **read- and write-boosting circuitry for failing static RAM cells**
- **circuits to track and adapt to performance variations**
- **complicated memory hierarchies to handle multicore architectures.**

All of those extra circuits add area. Some analysts have concluded that when you factor those circuits in, chips are no longer twice as dense from generation to generation. One such analysis suggests, the density improvement over the past three generations, from 2007 on, has been closer to 1.6 than 2.

And cost per transistor has gone *up* for the first time ever:

- **2012 20M 28nm transistors/dollar**
- **2015 19M 16nm transistors/dollar**

At end of day, we keep using all those new transistors.

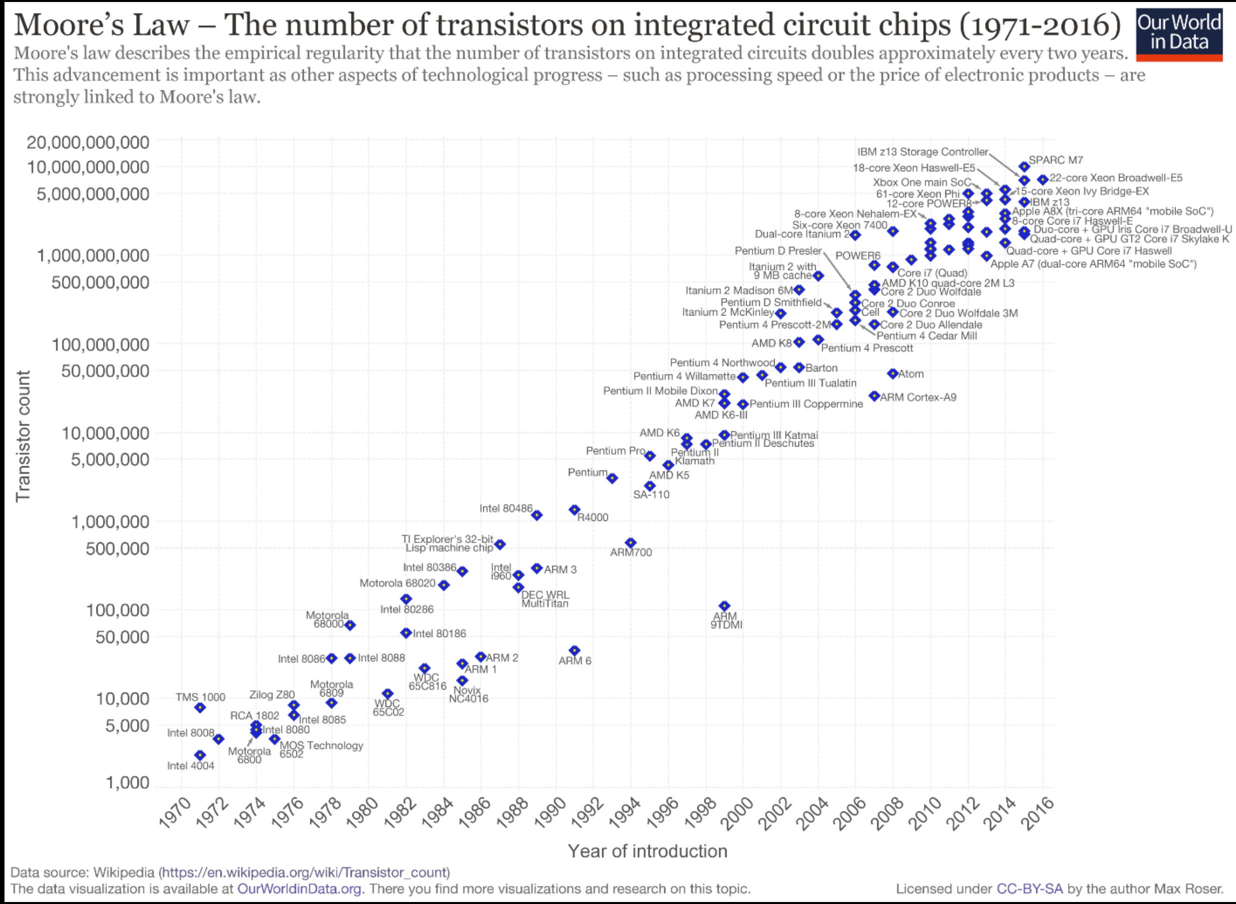

That Power and Clock Inflection Point in 2004…
didn't get better.

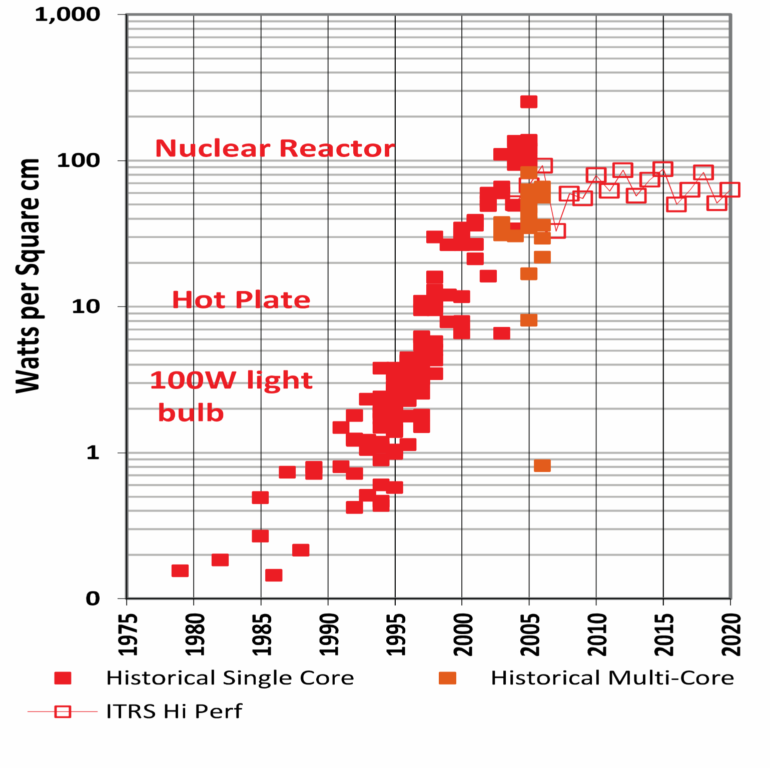

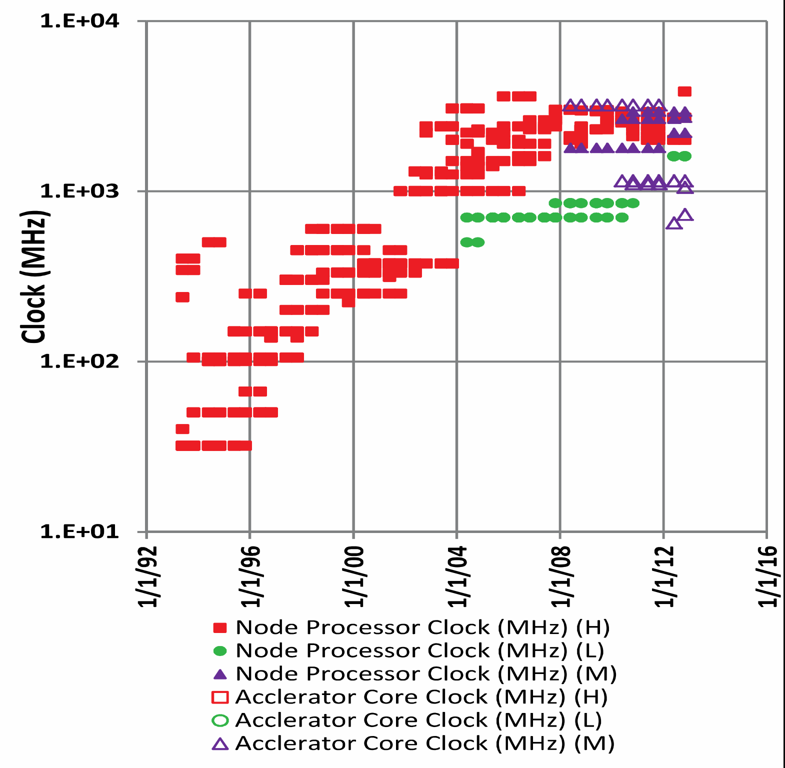

Fun fact: At 100+ Watts and <1V, currents are beginning to exceed 100A at the point of load!

Courtesy Horst Simon, LBNL

Not a new problem, just a new scale…

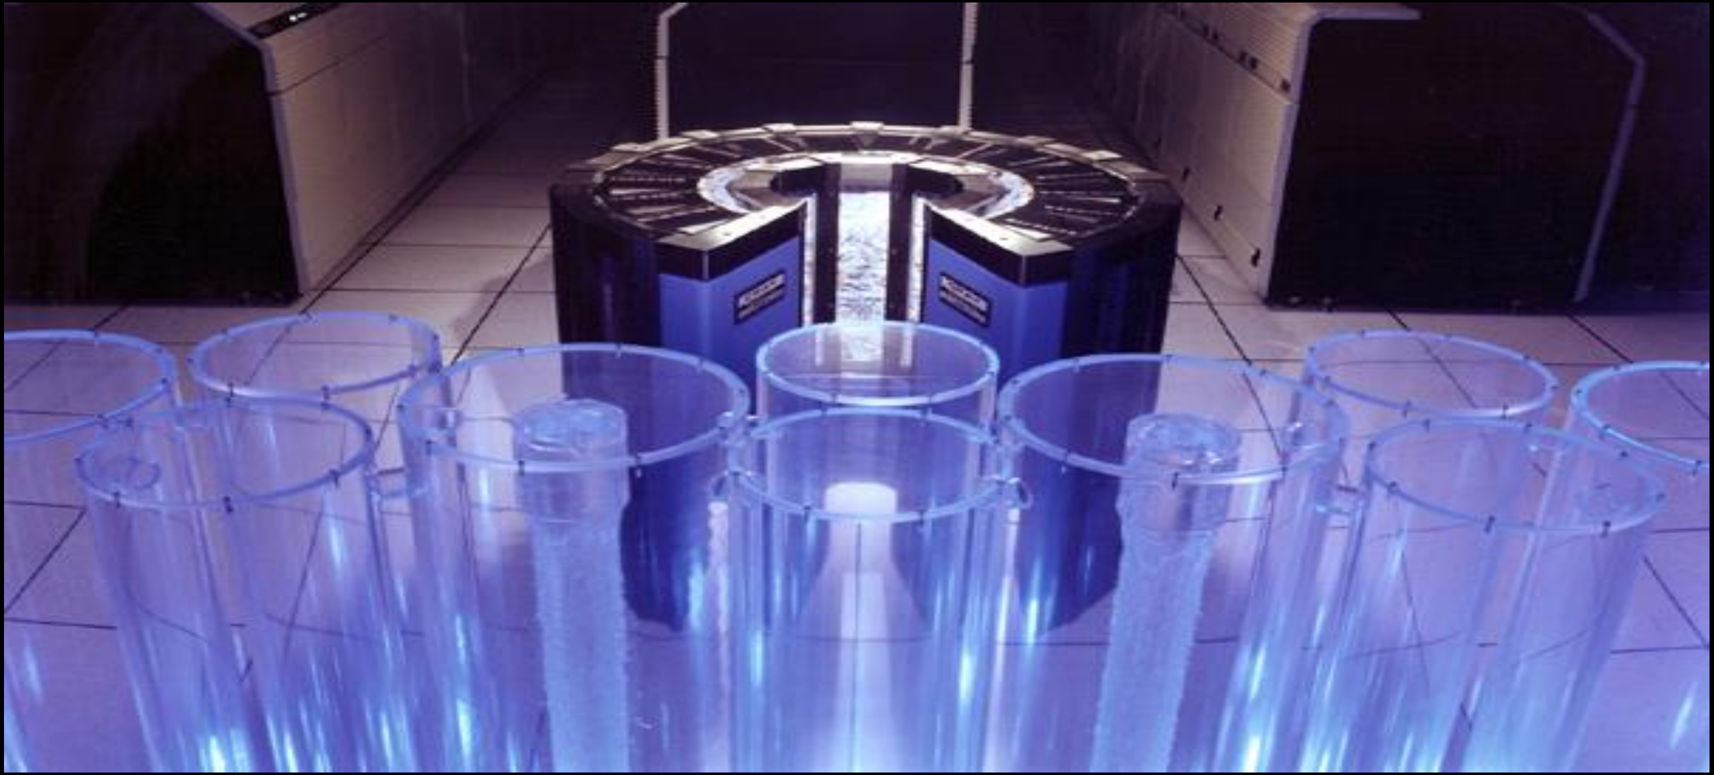

Cray-2 with cooling tower in foreground, circa 1985

How to get same number of transistors to give us more performance without cranking up power?

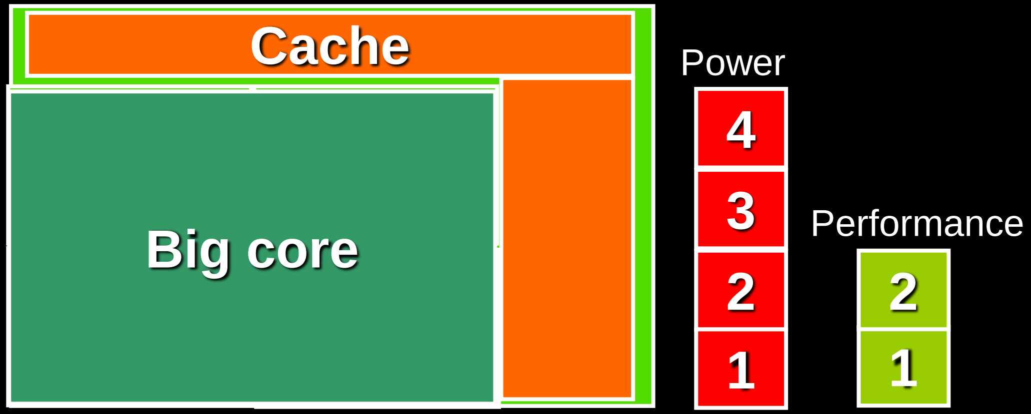

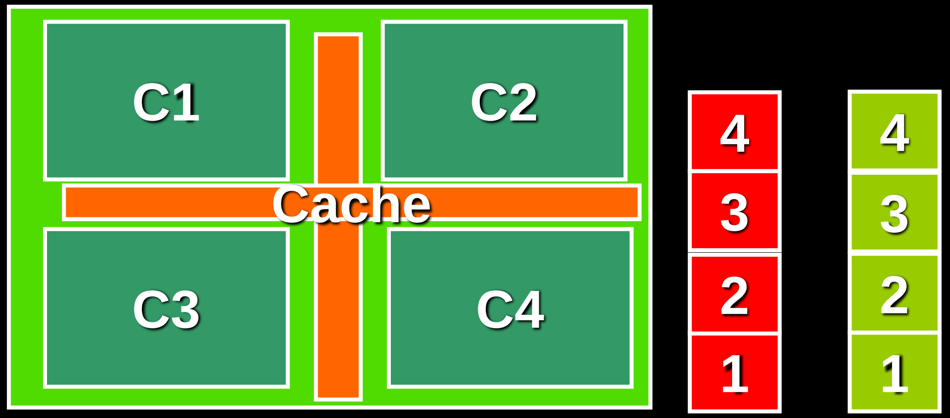

Key is that

Performance ≈ √ area

Power = ¼

Performance = 1/2

And how to get more performance from more transistors with the same power.

A 15% Reduction In Voltage Yields

RULE OF THUMB

Frequency Reduction	Power Reduction	Performance Reduction
15%	45%	10%

SINGLE CORE

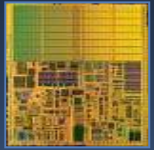

Area = 1
Voltage = 1
Freq = 1
Power = 1
Perf = 1

DUAL CORE

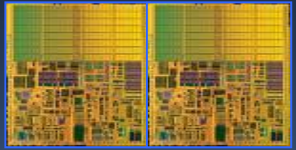

Area = 2
Voltage = 0.85
Freq = 0.85
Power = 1
Perf = ~1.8

Single Socket Parallelism

Processor	Year	Vector	Bits	SP FLOPs / core / cycle	Cores	FLOPs/cycle
Pentium III	1999	SSE	128	3	1	3
Pentium IV	2001	SSE2	128	4	1	4
Core	2006	SSE3	128	8	2	16
Nehalem	2008	SSE4	128	8	10	80
Sandybridge	2011	AVX	256	16	12	192
Haswell	2013	AVX2	256	32	18	576
KNC	2012	AVX512	512	32	64	2048
KNL	2016	AVX512	512	64	72	4608
Skylake	2017	AVX512	512	96	28	2688

Putting It All Together

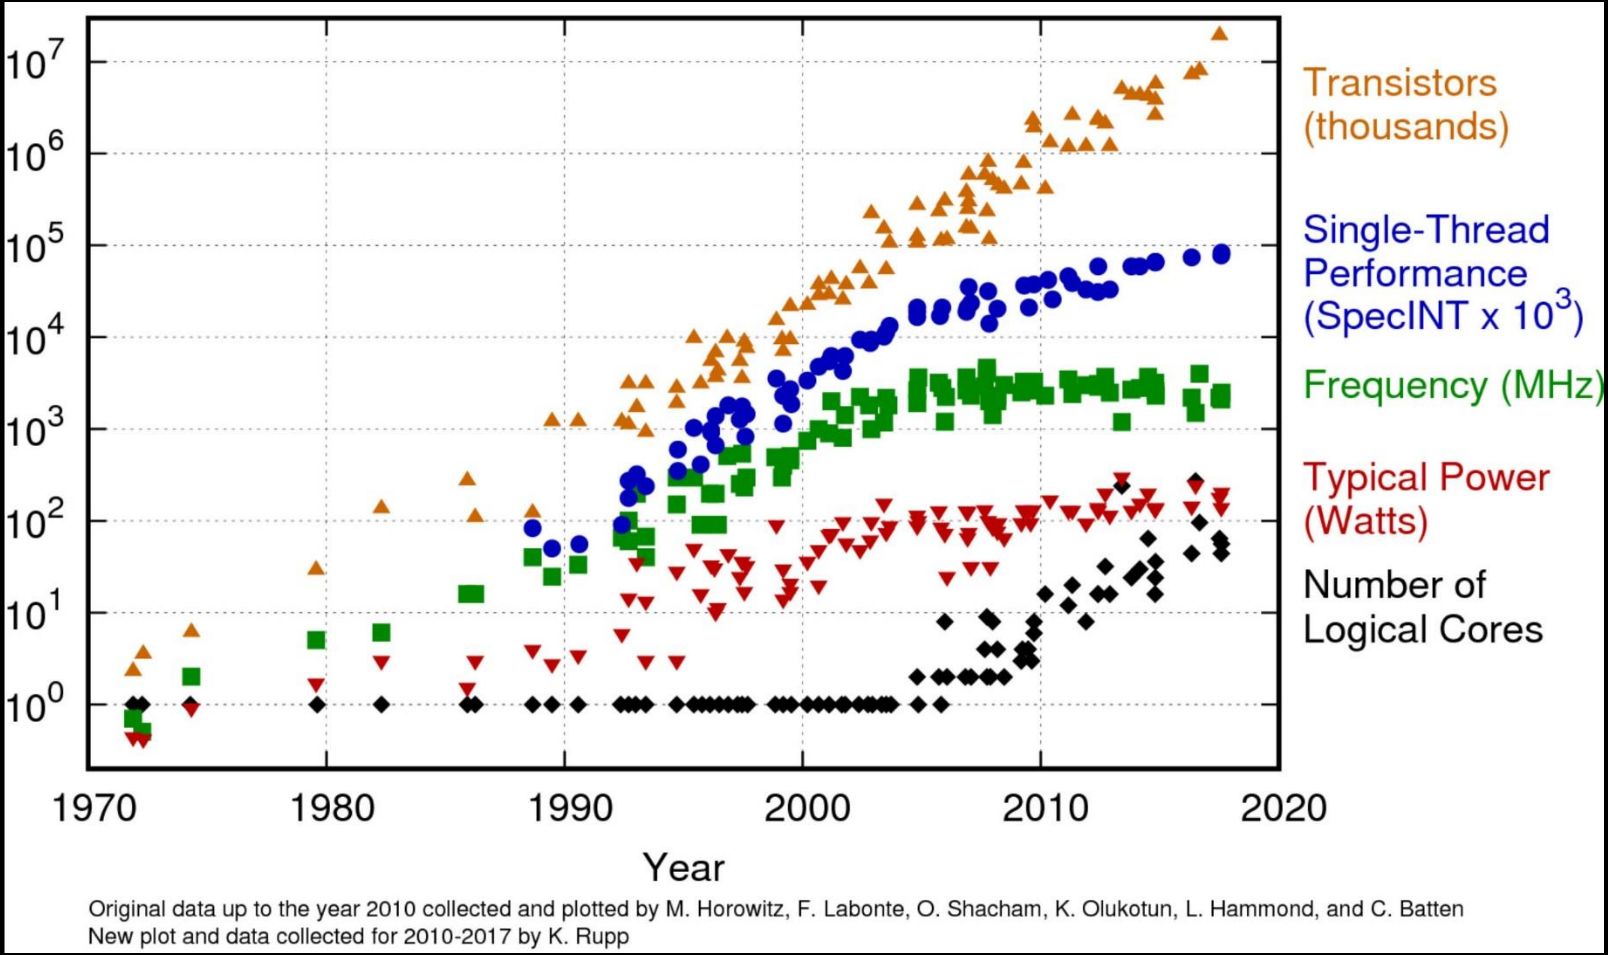

Many Levels and Types of Parallelism

- Vector (SIMD)
- Instruction Level (ILP)
  - Instruction pipelining
  - Superscaler (multiple instruction units)
  - Out-of-order
  - Register renaming
  - Speculative execution
  - Branch prediction

} Compiler (not your problem)

OpenMP {
- Multi-Core (Threads)
- SMP/Multi-socket

OpenACC {
- Accelerators: GPU & MIC

MPI {
- Clusters
- MPPs

Also Important

- ASIC/FPGA/DSP
- RAID/IO

The pieces fit like this…

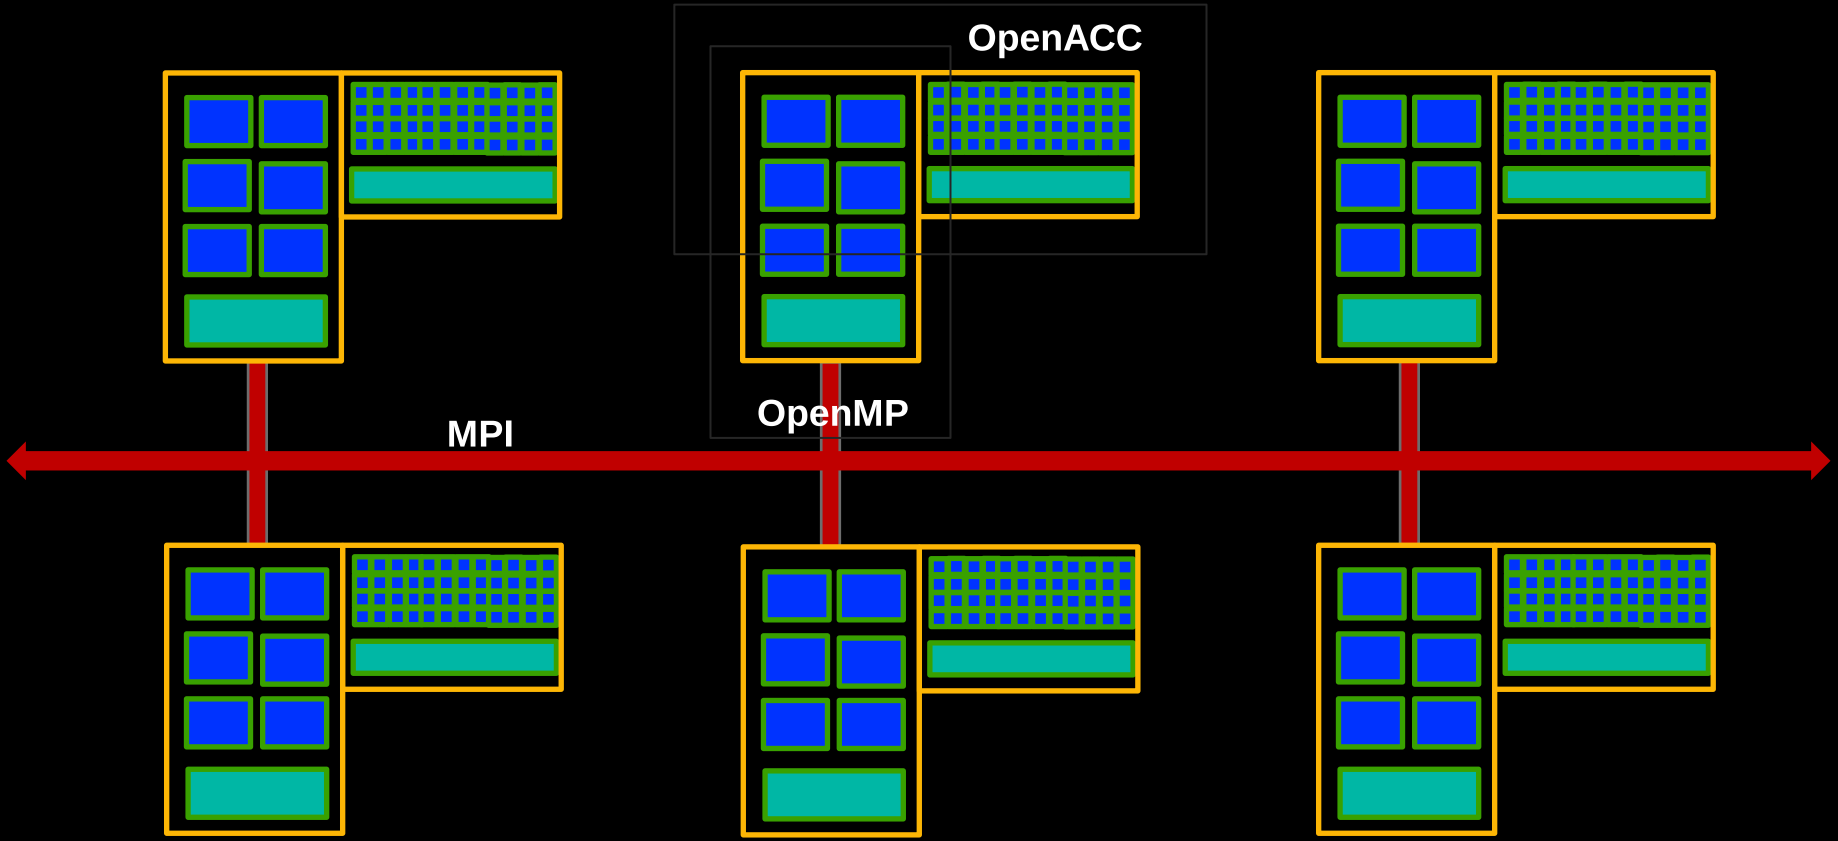

The Long-awaited Exascale - This year!

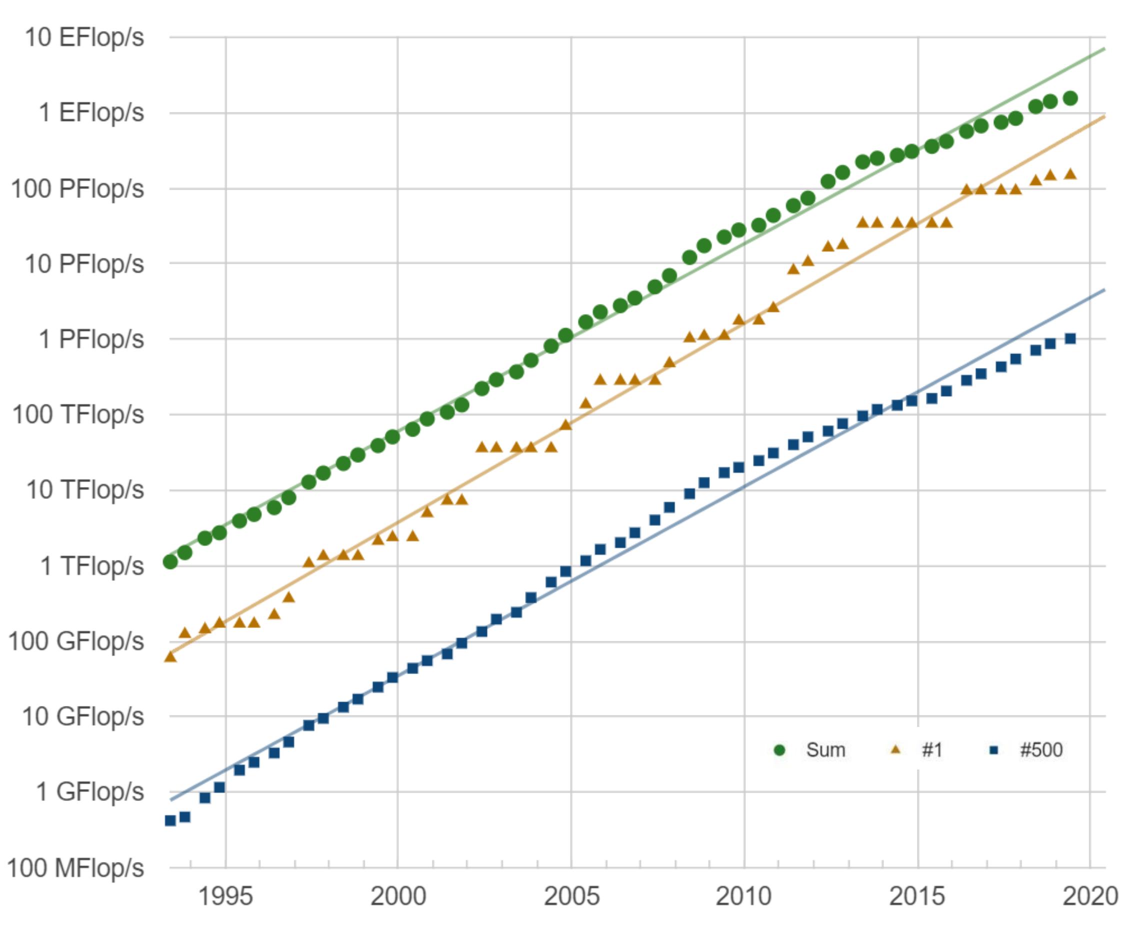

Staying on track to Exascale

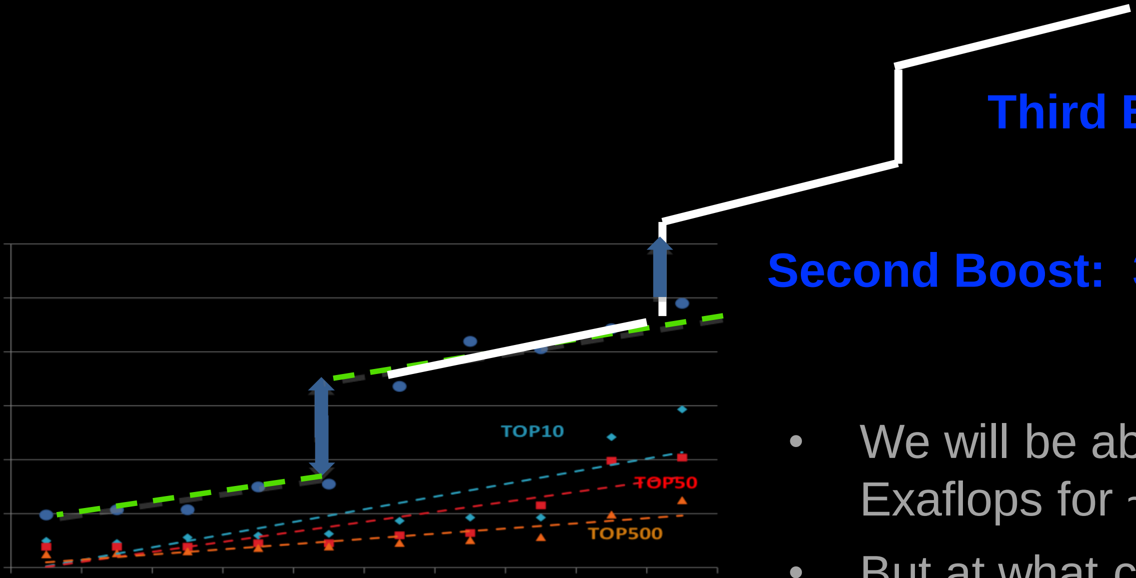

Third Boost: SiPh (2020 – 2024)

Second Boost: 3D (2016 – 2020)

First boost: many-core/accelerator

- We will be able to reach usable Exaflops for ~20 MW by 2021
- But at what cost?
- Will any of the other technologies give additional boosts after 2025?

Courtesy Horst Simon, LBNL

Top 10 Systems as of June 2020

#	Site	Manufacturer	Computer	CPU Interconnect [Accelerator]	Cores	Rmax (Tflops)	Rpeak (Tflops)	Power (MW)
1	RIKEN Center for Computational Science **Japan**	Fujitsu	Fugaku	ARM 8.2A+ 48C 2.2GHz Torus Fusion Interconnect	**7,299,072**	**415,530**	**513,854**	**28.3**
2	DOE/SC/ORNL **United States**	IBM	Summit	Power9 22C 3.0 GHz Dual-rail Infiniband EDR NVIDIA V100	**2,414,592**	**148,600**	**200,794**	**10.1**
3	DOE/NNSA/LLNL **United States**	IBM	Sierra	Power9 3.1 GHz 22C Infiniband EDR NVIDIA V100	**1,572,480**	**94,640**	**125,712**	**7.4**
4	National Super Computer Center in Wuxi **China**	NRCPC	Sunway TaihuLight	Sunway SW26010 260C 1.45GHz	**10,649,600**	**93,014**	**125,435**	**15.3**
5	National Super Computer Center in Guangzhou **China**	NUDT	Tianhe-2 (MilkyWay-2)	Intel Xeon E5-2692 2.2 GHz TH Express-2 Intel Xeon Phi 31S1P	**4,981,760**	**61,444**	**100,678**	**18.4**
6	Eni S.p.A **Italy**	Dell	HPc5	Xeon 24C 2.1 GHz Infiniband HDR NVIDIA V100	**669,760**	**35,450**	**51,720**	**2.2**
7	Eni S.p.A **Italy**	NVIDIA	Selene	EPYC 64C 2.25GHz Infiniband HDR NVIDIA A100	**272,800**	**27,580**	**34,568**	**1.3**
8	Texas Advanced Computing Center/Univ. of Texas **United States**	Dell	Frontera	Intel Xeon 8280 28C 2.7 GHz InfiniBand HDR	**448,448**	**23,516**	**38,745**	
9	Cineca **Italy**	IBM	Marconi100	Power9 16C 3.0 GHz Infiniband EDR NVIDIA V100	**347,776**	**21,640**	**29,354**	**1.5**
10	Swiss National Supercomputing Centre (CSCS) **Switzerland**	Cray	Piz Daint Cray XC50	Xeon E5-2690 2.6 GHz Aries NVIDIA P100	**387,872**	**21,230**	**27,154**	**2.4**

Today

- Pflops computing fully established with more than 500 machines
- The field is thriving
- Interest in supercomputing is now worldwide, and growing in many new markets
- Exascale projects in many countries and regions

USA: ECP by the Numbers

7 YEARS

$1.7B

A seven-year, $1.7 B R&D effort that launched in 2016

6 CORE DOE LABS

Six core DOE National Laboratories: Argonne, Lawrence Berkeley, Lawrence Livermore, Oak Ridge, Sandia, Los Alamos

- Staff from most of the 17 DOE national laboratories take part in the project

3 FOCUS AREAS

Three technical focus areas: Hardware and Integration, Software Technology, Application Development supported by a Project Management Office

100 R&D TEAMS

1000 RESEARCHERS

More than 100 top-notch R&D teams

Hundreds of consequential milestones delivered on schedule and within budget since project inception

The Plan

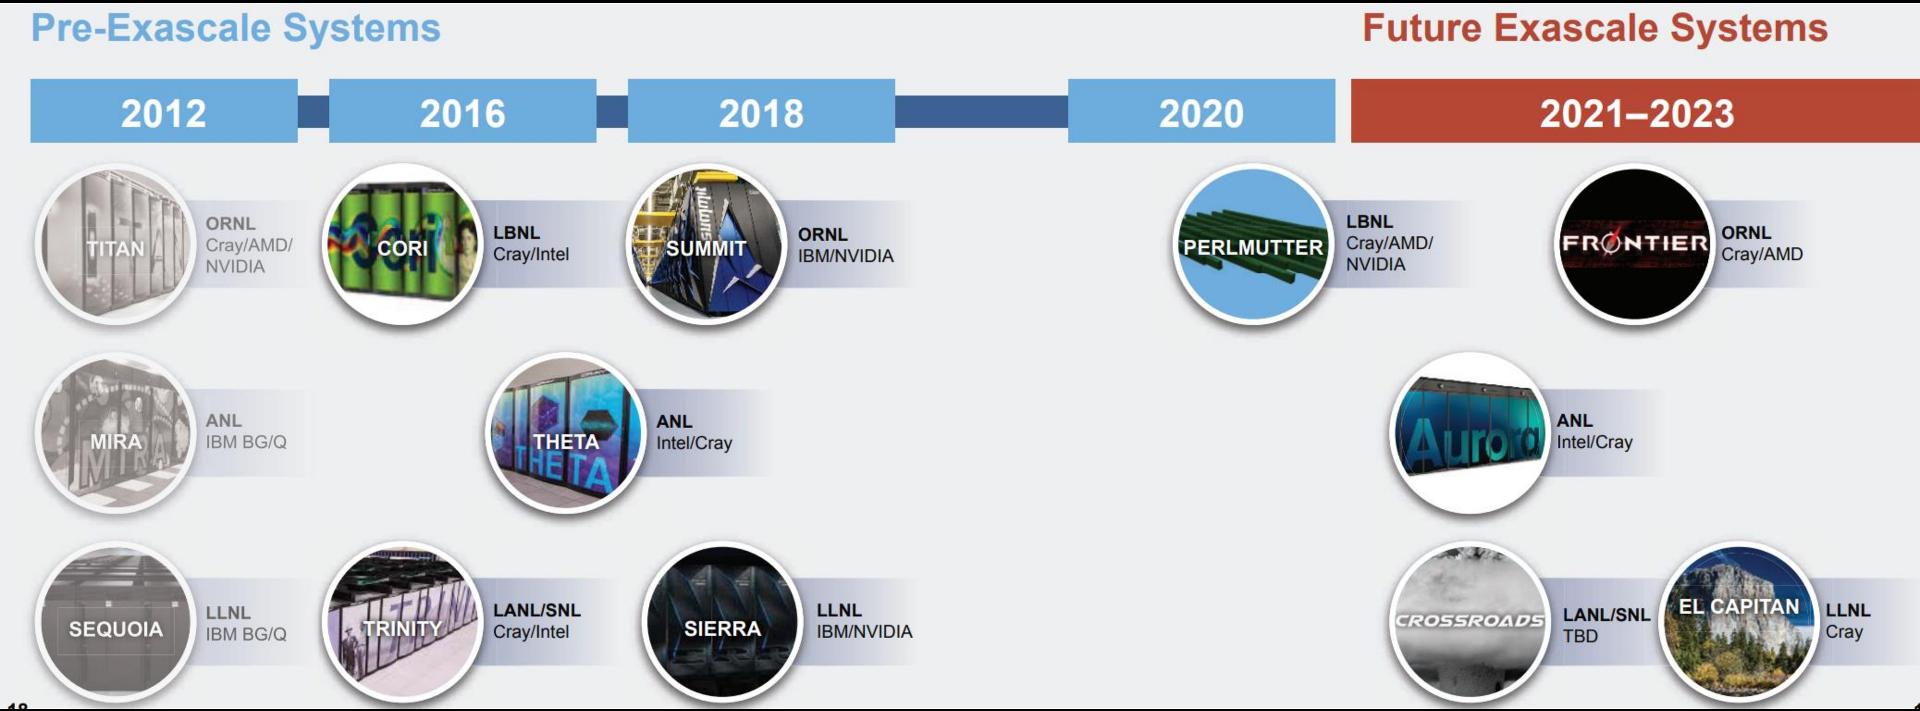
Pre-Exascale Systems
Future Exascale Systems
2012
2016
2018
2020
2021–2023
TITAN
ORNL
Cray/AMD/
NVIDIA
CORI
LBNL
Cray/Intel
SUMMIT
ORNL
IBM/NVIDIA
PERLMUTTER
LBNL
Cray/AMD/
NVIDIA
FRONTIER
ORNL
Cray/AMD
MIRA
ANL
IBM BG/Q
THETA
ANL
Intel/Cray
Aurora
ANL
Intel/Cray
SEQUOIA
LLNL
IBM BG/Q
TRINITY
LANL/SNL
Cray/Intel
SIERRA
LLNL
IBM/NVIDIA
CROSSROADS
LANL/SNL
TBD
EL CAPITAN
LLNL
Cray

System Designs

System	Performance	Power	Interconnect	Node
Aurora (ANL)	> 1 EF		100 GB/s Cray Slingshot Dragonfly	2 Intel Xeon CPU + 6 Intel Xe GPUs
El Capitan (LLNL)	> 1.5 EF	30-40 MW	100 GB/s Cray Slingshot Dragonfly	AMD Epyc CPU + 4 Radeon GPUs
Frontier (ORNL)	> 1.5 EF		100 GB/s Cray Slingshot Dragonfly	AMD Epyc CPU + 4 Radeon GPUs
Perlmutter (LBNL)			Cray Slingshot Dragonfly	2 AMD Epyc CPU + 4 Volta GPUs

Obstacles?

One of the many groups established to enable this outcome (the Advanced Scientific Computing Advisory Committee) puts forward this list of 10 technical challenges:

- Energy efficient circuit, power and cooling technologies.
- High performance interconnect technologies.
- Advanced memory technologies to dramatically improve capacity and bandwidth.
- Scalable system software that is power and resilience aware.
- Data management software that can handle the volume, velocity and diversity of data-storage
- Programming environments to express massive parallelism, data locality, and resilience.
- Reformulating science problems and refactoring solution algorithms for exascale.
- Ensuring correctness in the face of faults, reproducibility, and algorithm verification.
- Mathematical optimization and uncertainty quantification for discovery, design, and decision.
- Software engineering and supporting structures to enable scientific productivity.

It is not just "exaflops" – we are changing the whole computational model
Current programming systems have WRONG optimization targets

Old Constraints

- **Peak clock frequency *as primary limiter for performance improvement***
- **Cost: *FLOPs* are biggest cost for system: *optimize for compute***
- **Concurrency: Modest growth of parallelism by adding nodes**
- **Memory scaling*: maintain byte per flop capacity and bandwidth***
- **Locality*: MPI+X model (uniform costs within node & between nodes)***
- **Uniformity: *Assume uniform system performance***
- **Reliability*: It's the hardware's problem***

New Constraints

- **Power** *is primary design constraint for future HPC system design*
- **Cost:** *Data movement dominates: optimize to minimize data movement*
- **Concurrency:** *Exponential growth of parallelism within chips*
- **Memory Scaling:** *Compute growing 2x faster than capacity or bandwidth*
- **Locality**: *must reason about data locality and possibly topology*
- **Heterogeneity**: *Architectural and performance non-uniformity increase*
- **Reliability**: *Cannot count on hardware protection alone*

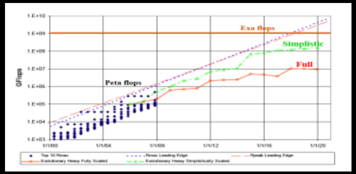

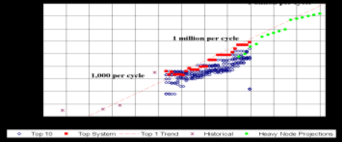

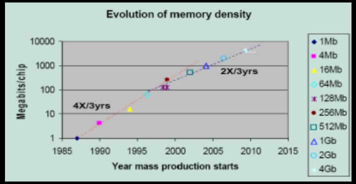

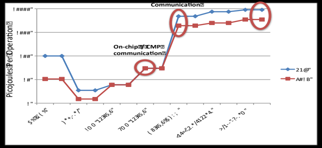

Fundamentally breaks our current programming paradigm and computing ecosystem

Adapted from John Shalf

End of Moore's Law Will Lead to New Architectures

Non-von
Neumann

ARCHITECT

von Neu

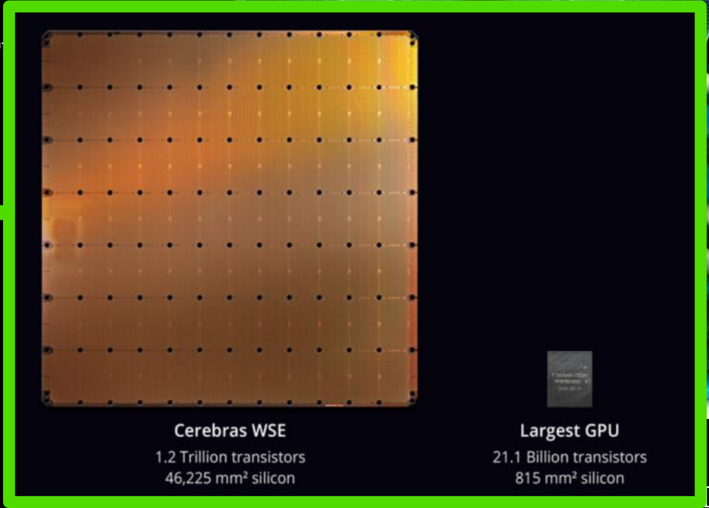

B

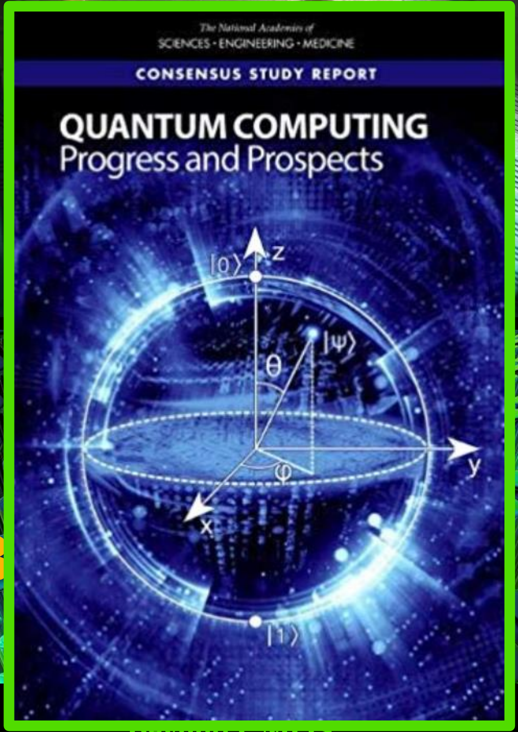

Beyond CMOS

TECHNOLOGY

It would only be the 6th paradigm.

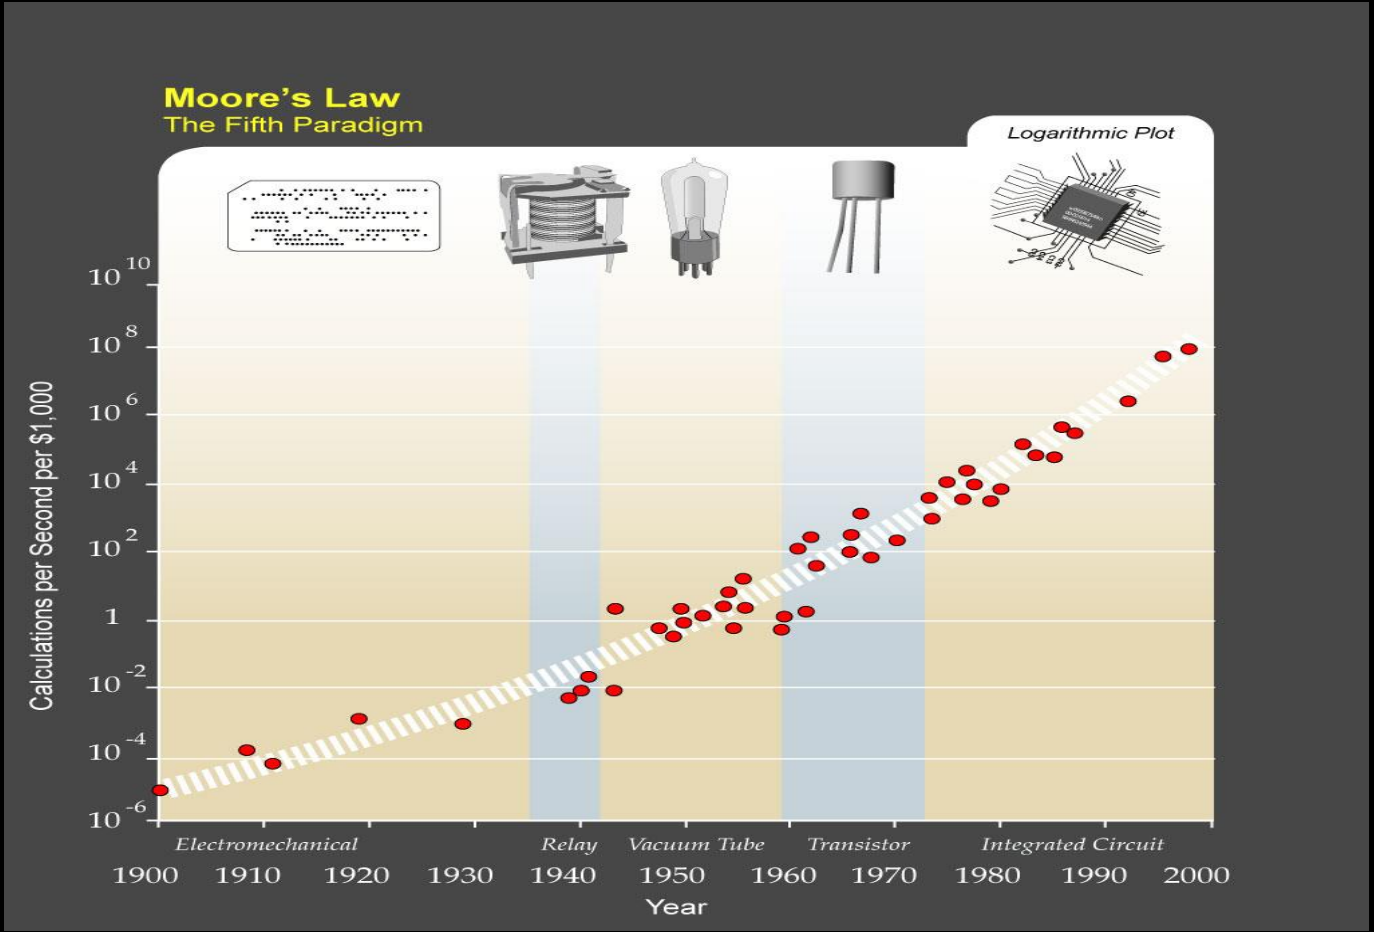

We can do better. We have a role model.

- **Straight forward extrapolation results in a real-time human brain scale simulation at about 1 - 10 Exaflop/s with 4 PB of memory**
- **Current predictions envision Exascale computers in 2021 with a power consumption of at best 20 - 30 MW**
- **The human brain takes 20W**
- **Even under best assumptions in 2020 our brain will still be a million times more power efficient**

Courtesy Horst Simon, LBNL

As a last resort, we ~~could~~ will learn to program again.

It has become a mantra of contemporary programming philosophy that developer hours are so much more valuable than hardware, that the best design compromise is to throw more hardware at slower code.

This might well be valid for some Java dashboard app used twice a week by the CEO. But this has spread and results in…

The common observation that a modern PC (or phone) seems to be more laggy than one from a few generations ago that had literally one thousandth the processing power.

Moore's Law has been the biggest enabler (or more accurately rationalization) for this trend. If Moore's Law does indeed end, then progress will require good programming.

No more garbage collecting, script languages. I am looking at you, Java, Python, Matlab.

Do you really care about software?

Of course you should. Here are a few reassuring words that software at exascale is not an afterthought, followed by more than a few application examples.

ECP application domains.

National security	Energy security	Economic security	Scientific discovery	Earth system	Health care
Stockpile stewardship Next generation simulation tools for assessing nuclear weapons performance Response to hostile threat environments and reentry conditions	Turbine wind plant efficiency High-efficiency, low-emission combustion engine and gas turbine design Materials design for extreme environments of nuclear fission and fusion reactors Design and commercialization of Small Modular Reactors Subsurface use for carbon capture, petroleum extraction, waste disposal Scale-up of clean fossil fuel combustion Biofuel catalyst design	Additive manufacturing of qualifiable metal parts Reliable and efficient planning of the power grid Seismic hazard risk assessment	Find, predict, and control materials and properties Cosmological probe of the standard model of particle physics Validate fundamental laws of nature Demystify origin of chemical elements Light source-enabled analysis of protein and molecular structure and design Whole-device model of magnetically confined fusion plasmas	Accurate regional impact assessments in Earth system models Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols Metagenomics for analysis of biogeochemical cycles, climate change, environmental remediation	Accelerate and translate cancer research

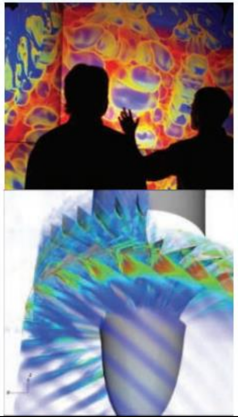

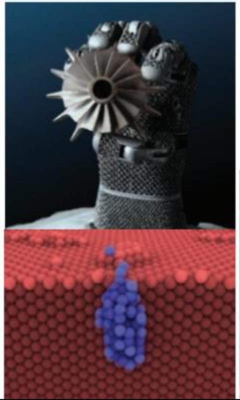

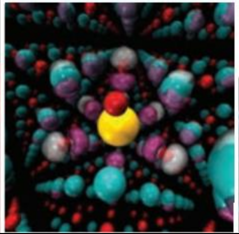

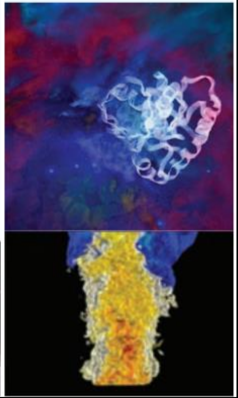

xSDK Version 0.4.0: December 2018 (even better today)

https://xsdk.info

Each xSDK member package uses or can be used with one or more xSDK packages, and the connecting interface is regularly tested for regressions.

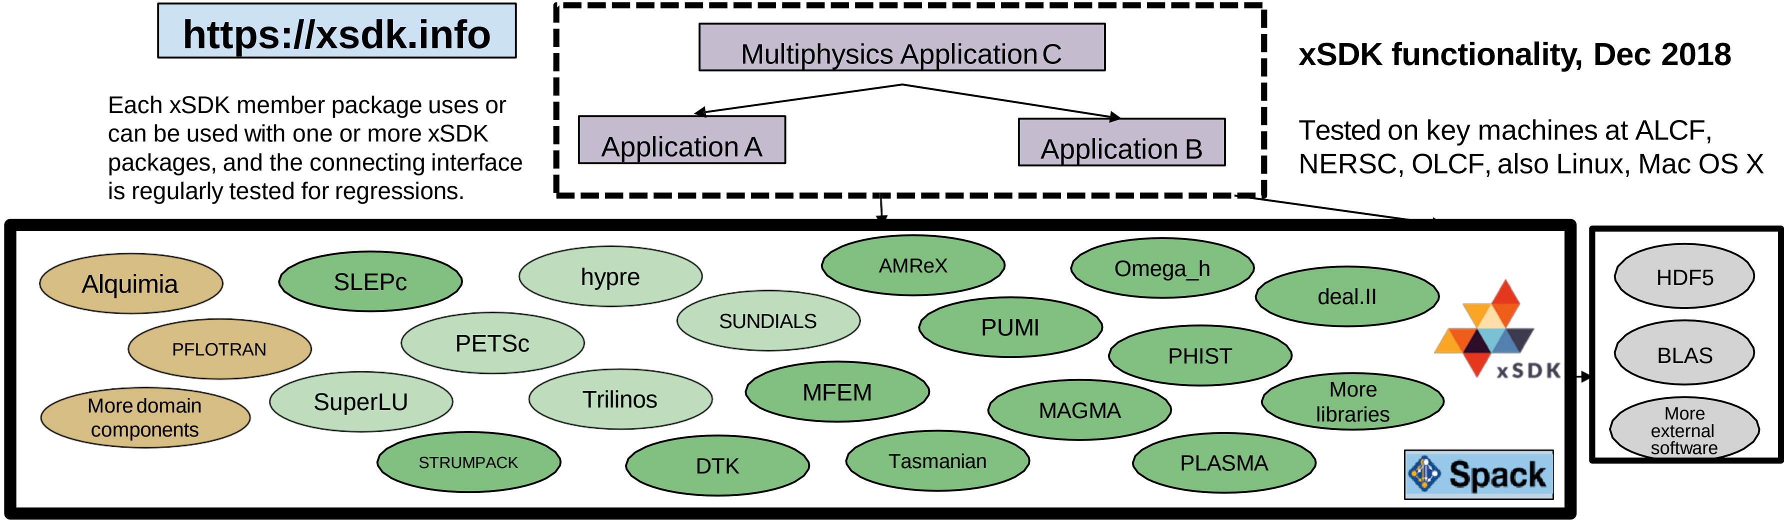

xSDK functionality, Dec 2018

Tested on key machines at ALCF, NERSC, OLCF, also Linux, Mac OS X

December 2018

- 17 math libraries
- 2 domain components
- 16 mandatory xSDK community policies
- Spack xSDK installer

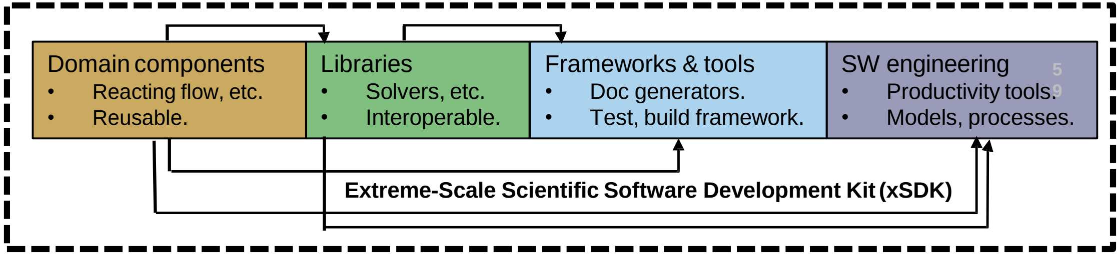

Impact: Improved code quality, usability, access, sustainability

Foundation for work on performance portability, deeper levels of package interoperability

The planned ECP ST SDKs will span all technology areas

xSDK (16)	PMR Core (17)	Tools and Technology (11)	Compilers & Support (7)	Visualization Analysis & Reduction (9)	Data Mgmt, I/O Services, & Checkpoint restart (12)	Ecosystem/E4S at-large (12)
hypre	Legion	TAU	openarc	ParaView	FAODEL	BEE
FleSCI	Kokkos (Support)	HPCToolkit	Kitsune	Catalyst	ROMIO	FSEFI
MFEM	RAJA	Dyninst Binary Tools	LLVM	VTK-m	Mercury (part of Mochi suite)	Kitten Lightweight Kernel
Kokkoskernels	CHAI	Gotcha	CHiLL Autotuning Compiler	SZ	HDF5	COOLR
Trilinos	PaRSEC*	Caliper	LLVM OpenMP compiler	zfp	Parallel netCDF	NRM
SUNDIALS	DARMA	PAPI	OpenMP V & V	VisIt	ADIOS	ArgoContainers
PETSc/TAO	GASNet-EX	Program Database Toolkit	Flang/LLVM Fortran compiler	ASCENT	Darshan	Spack
libEnsemble	Qthreads	Search using Random Forests		Cinema	UnifyCR	MarFS
STRUMPACK	BOLT	Siboka		ROVER	VeloC	GUFI
SuperLU	UPC++	C2C			IOSS	Intel GEOPM
ForTrilinos	MPICH	Sonar			HXHIM	mpiFileUtils
SLATE	Open MPI				SCR	TriBITS
MAGMA	Umpire					
DTK	QUO					
Tasmanian	Papyrus					
TuckerMPI	SICM					
	AML					

Key

- PMR
- Tools
- Math Libraries
- Data and Vis
- Ecosystems and Delivery

Appendix

Endless apps…

CEED is targeting several ECP applications

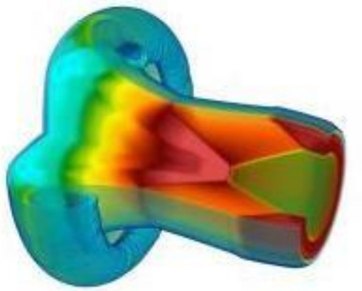

Compressible flow (MARBL)

Climate (E3SM)

Urban systems (Urban)

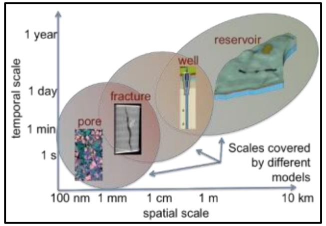

Subsurface (GEOS)

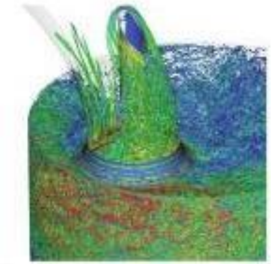

Combustion (Nek5000)

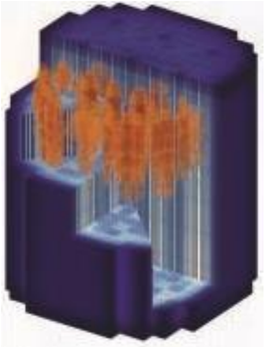

Modular Nuclear Reactors (ExaSMR)

Wind Energy (ExaWind)

Additive Manufacturing (ExaAM)

Magnetic Fusion (WDMApp)

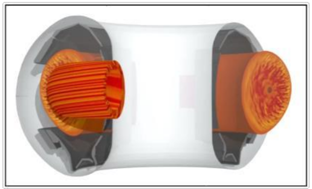

ECP EXASCALE COMPUTING PROJECT

PI: Tzanio Kolev (LLNL)

ECP's Adaptive Mesh Refinement Co-Design Center: AMReX

- Develop and deploy software to support block-structured adaptive mesh refinement on exascale architectures
  - Core AMR functionality
  - Particles coupled to AMR meshes
  - Embedded boundary (EB) representation of complex geometry
  - Linear solvers
  - Supports two modalities of use
    - Library support for AMR
    - Framework for constructing AMR applications
- Provide direct support to ECP applications that need AMR for their application
- Evaluate software technologies and integrate with AMReXwhen appropriate
- Interact with hardware technologies / vendors

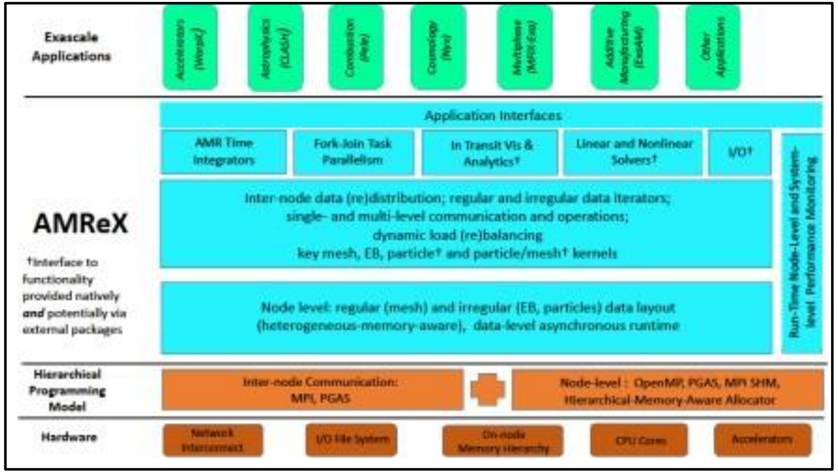

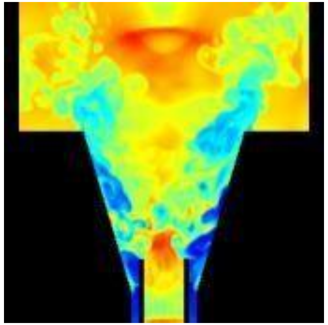

Application	Particles	ODEs	Linear Solvers	EB
Combustion	X	X	X	X
Multiphase	X		X	X
Cosmology	X	X	X	
Astrophysics	X	X	X	
Accelerators	X			

PI: John Bell (LBNL)

ECP's Co-Design Center for Online Data Analysis and Reduction

CODAR

Traditional approach:
Compute…output…analyze [offline]

Write simulation output to secondary storage; read back for analysis

Decimate in time when simulation output rate exceeds output rate of computer

New approach:
Online data analysis and reduction

Co-optimize simulation, analysis, reduction for performance and information output

Substitute CPU cycles for I/O, via online data (de)compression and/or online data analysis

Provide the right information at the right time and place to accelerate discovery!

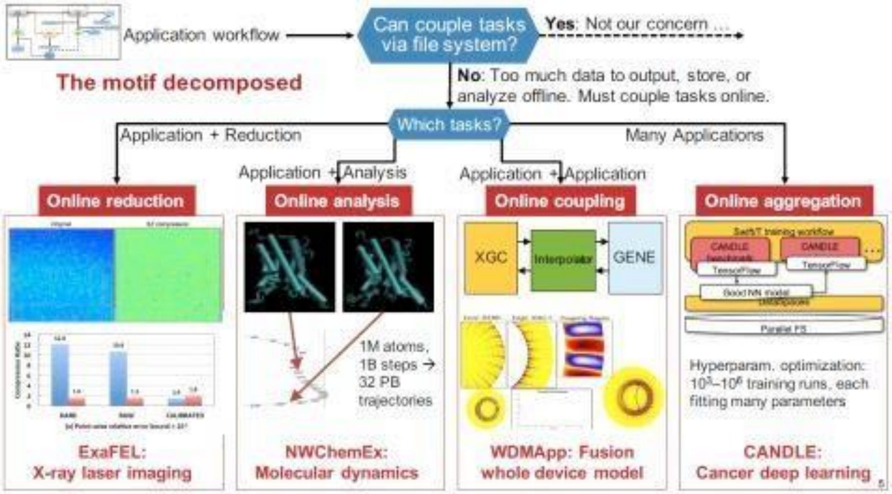

<u>Goal:</u> Replace the activities in HPC workflow that have been mediated through file I/O with in-situ methods / workflows. data reduction, analysis, code coupling, aggregation (e.g. parameter studies).

<u>Cross-cutting tools:</u>

- Workflow setup, manager (Cheetah, Savanna); Data coupler (ADIOS-SST); Compression methods (MGARD, FTK, SZ), compression checker (Z-checker)
- Performance tools (TAU, Chimbuco, SOSFlow)

EXASCALE COMPUTING PROJECT

PI: Ian Foster (ANL)

ECP's Co-Design Center for Particle Applications: CoPA

Goal: Develop algorithms and software for particle methods,

Cross-cutting capabilities:

- Specialized solvers for quantum molecular dynamics (Progress / BML).
- Performance-portable libraries for classical particle methods in MD, PDE (Cabana).
- FFT-based Poisson solvers for long-range forces.

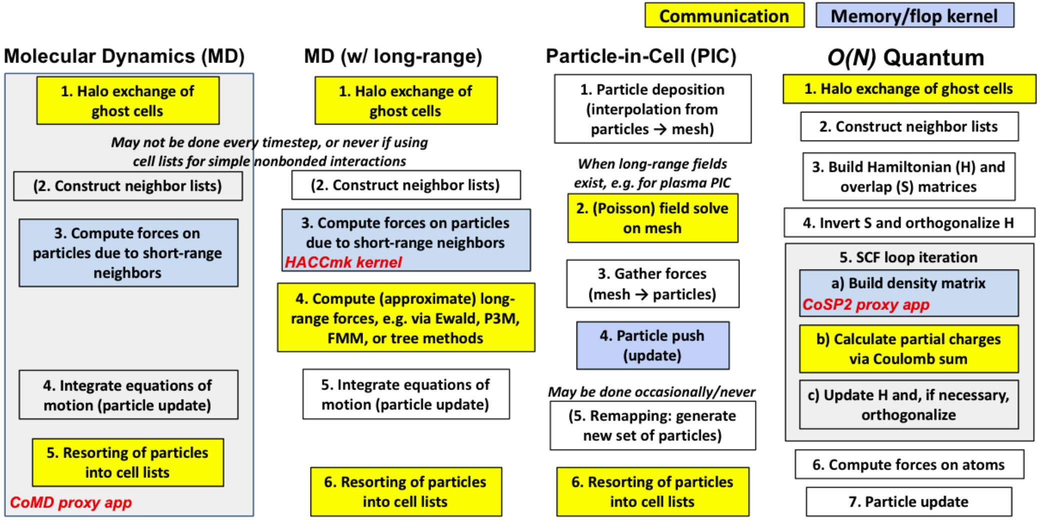

Technical approach:

- High-level C++ APIs, plus a Fortran interface (Cabana).
- Leverage existing / planned FFT software.
- Extensive use of miniapps / proxy apps as part of the development process.

PI: Sue Mniszewski (LANL) recently replacing Tim Germann (LANL), who is taking on a larger role in ECP

ECP’s Co-Design Center for Machine Learning: ExaLearn

Bringing together experts from 8 DOE Laboratories

- AI has the potential to accelerate scientific discovery or enable prediction in areas currently too complex for direct simulation (ML for HPC and HPC for ML)
- AI use cases of interest to ECP:
  - *Classification and regression*, including but not limited to image classification and analysis, e.g. scientific data output from DOE experimental facilities or from national security programs.
  - *Surrogate models* in high-fidelity and multiscale simulations, including uncertainty quantification and error estimation.
  - *Structure-to-function relationships*, including genome-to-phenome, the prediction of materials performance based on atomistic structures, or the prediction of performance margins based on manufacturing data.
  - *Control systems*, e.g., for wind plants, nuclear power plants, experimental steering and autonomous vehicles.
  - *Inverse problems* and optimization. This area would include, for example, inverse imaging and materials design.
- Areas in need of research
  - Data quality and statistics
  - Learning algorithms
  - Physics-Informed AI
  - Verification and Validation
  - Performance and scalability
  - Workflow and deployment

Expected Work Product: A Toolset That . . .

- Has a line-of-sight to exascale computing, e.g. through using exascale platforms directly, or providing essential components to an exascale workflow
- Does not replicate capabilities easily obtainable from existing, widely-available packages
- Builds in domain knowledge where possible “Physics”-based ML and AI
- Quantifies uncertainty in predictive capacity
- Is interpretable
- Is reproducible
- Tracks provenance

EXASCALE COMPUTING PROJECT

PI: Frank Alexander (BNL)

Machine Learning in the Light Source Workflow

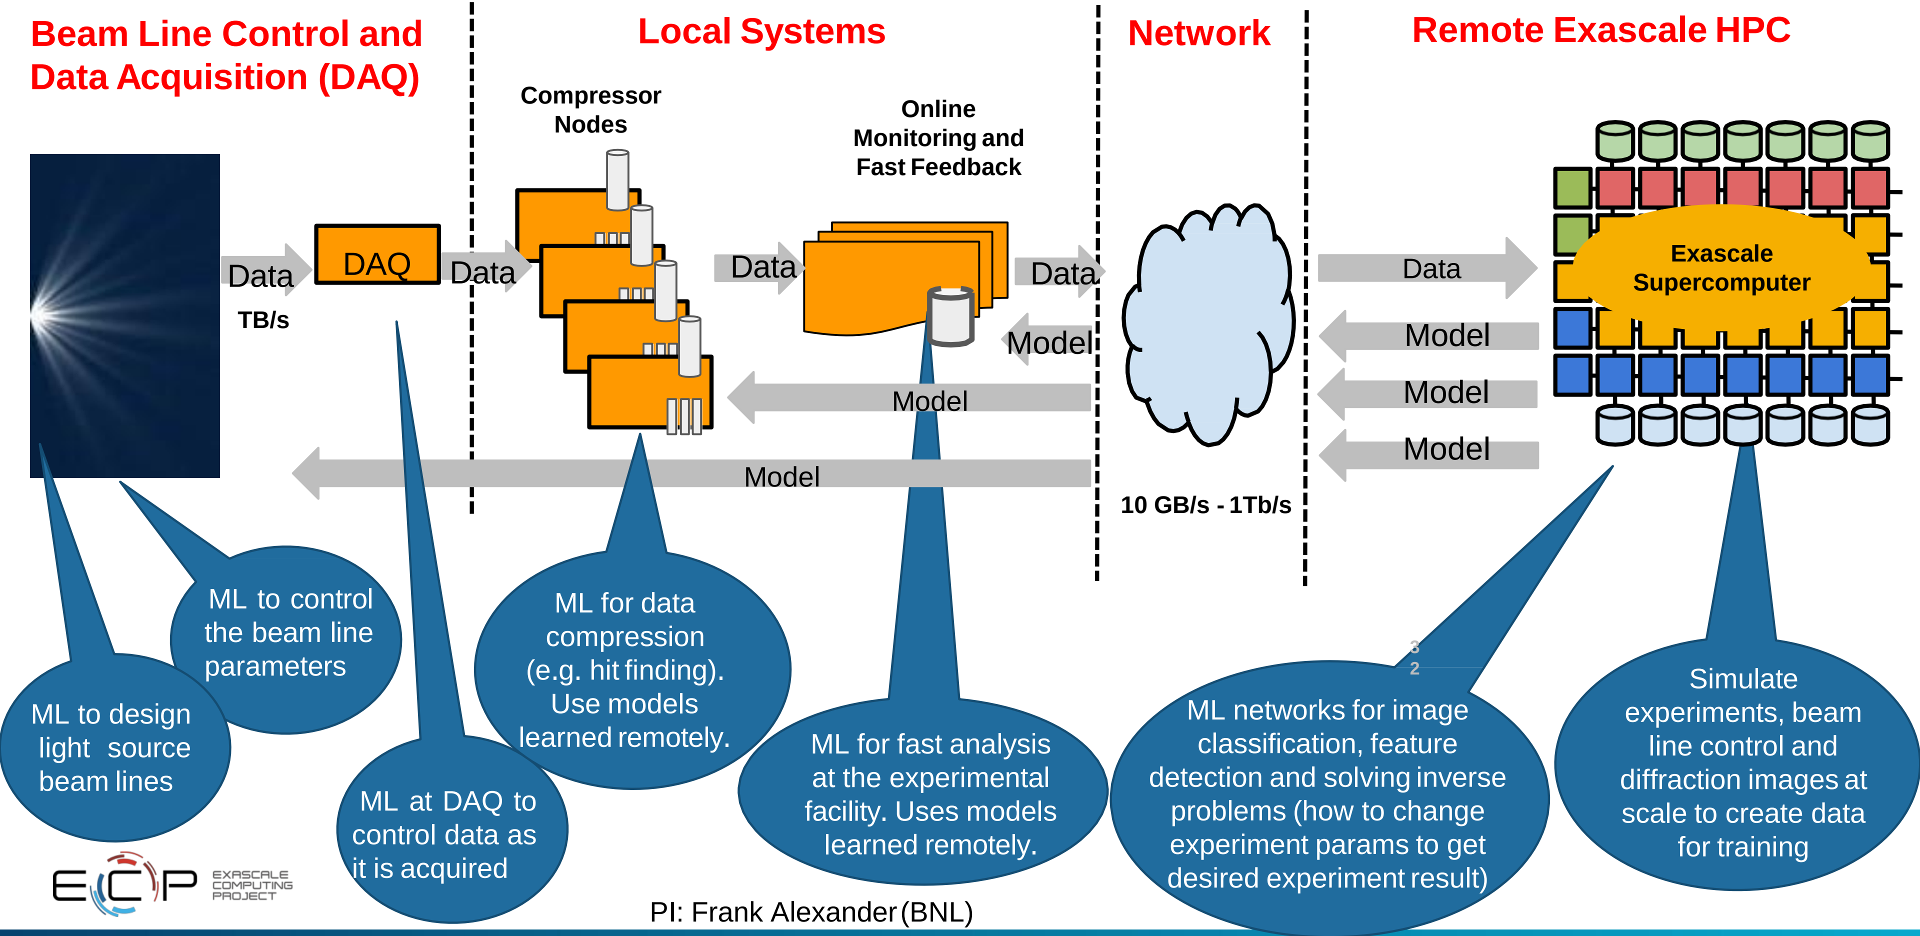

EXASCALE COMPUTING PROJECT

PI: Frank Alexander (BNL)

Exascale apps can deliver transformative products and solutions

ExaWind

Turbine Wind Plant Efficiency
(Mike Sprague, NREL)

- Harden wind plant design and layout against energy loss susceptibility
- Increase penetration of wind energy

Challenges: linear solver perf in strong scale limit; manipulation of large meshes; overset of structured & unstructured grids; communication-avoiding linear solvers

ExaAM

Additive Manufacturing (AM) of Qualifiable Metal Parts
(John Turner, ORNL)

- Accelerate the widespread adoption of AM by enabling routine fabrication of qualifiable metal parts

Challenges: capturing unresolved physics; multi-grid linear solver performance; coupled physics

EQSIM

Earthquake Hazard Risk Assessment
(David McCallen, LBNL)

- Replace conservative and costly earthquake retrofits with safe purpose-fit retrofits and designs

Challenges: full waveform inversion algorithms

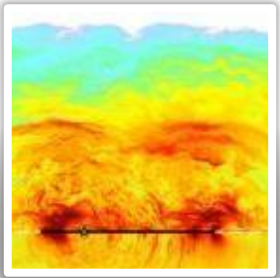

EQSIM: Understanding and predicting earthquake phenomenon

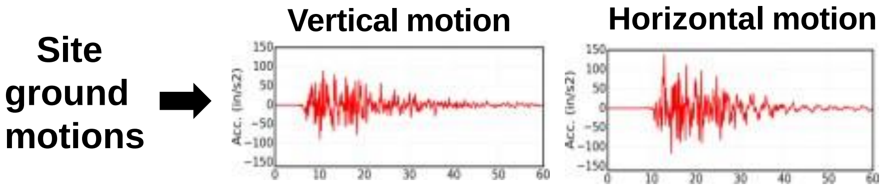

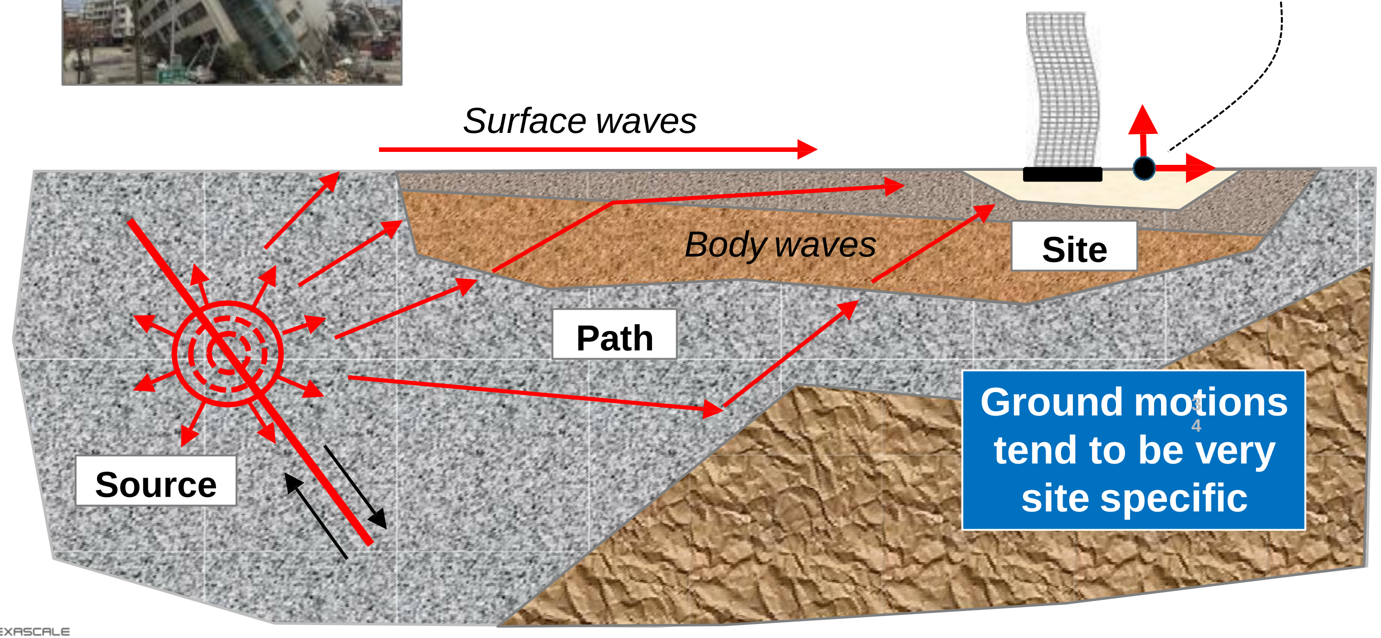

PI: David McCallen (LBNL)

EQSIM: The Exascale “Big Lift” – regional ground motion simulations at engineering frequencies

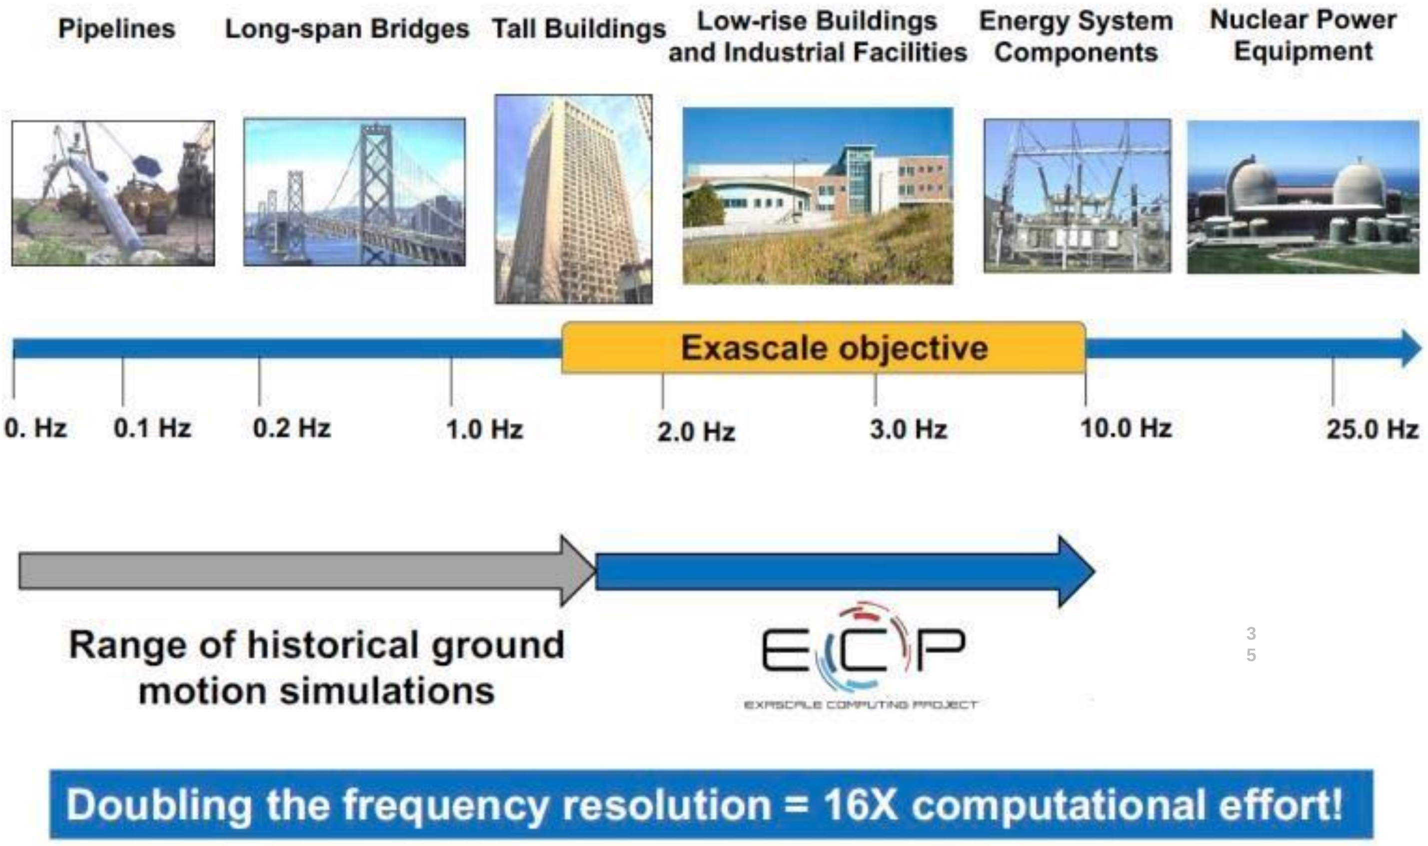

PI: David McCallen (LBNL)

Exascale apps can deliver transformative products and solutions

MFIX-Exa

Scale-up of Clean Fossil Fuel Combustion
(Madhava Syamlal, NETL)

- Commercial-scale demonstration of transformational energy technologies – curbing CO_2 emissions at fossil fuel power plants by 2030

Challenges: load balancing; strong scaling thru transients

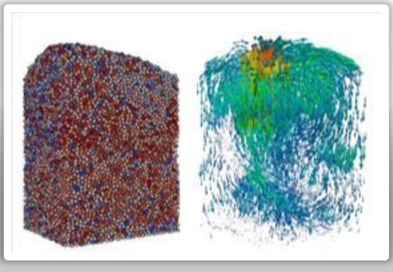

GAMESS

Biofuel Catalyst Design
(Mark Gordon, Ames)

- Design more robust and selective catalysts orders of magnitude more efficient at temperatures hundreds of degrees lower

Challenges: weak scaling of overall problem; on-node performance of molecular fragments

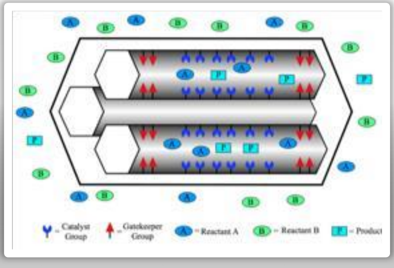

EXAALT

Materials for Extreme Environments
(Danny Perez, LANL)

- Simultaneously address time, length, and accuracy requirements for predictive microstructural evolution of materials

Challenges: SNAP kernel efficiency on accelerators; efficiency of DFTB application on accelerators

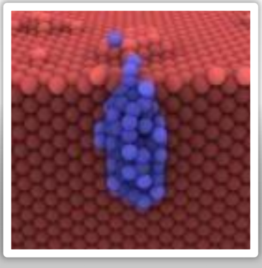

EXASCALE COMPUTING PROJECT

Exascale apps can deliver transformative products and solutions

ExaSMR

Design and Commercialization of Small Modular Reactors
(Steve Hamilton, ORNL)

- Virtual test reactor for advanced designs via experimental-quality simulations of reactor behavior

Challenges: existing GPU-based MC algorithms require rework for hardware less performant for latency-bound algorithms with thread divergence; performance portability with OCCA & OpenACC not achievable; insufficient node memory for adequate CFD + MC coupling

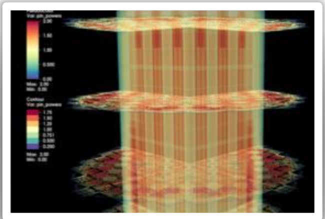

Subsurface

Carbon Capture, Fossil Fuel Extraction, Waste Disposal
(Carl Steefel, LBNL)

- Reliably guide safe long-term consequential decisions about storage, sequestration, and exploration

Challenges: performance of Lagrangian geomechanics; adequacy of Lagrangian crack mechanics) + Eulerian (reaction, advection, diffusion) models; parallel HDF5 for coupling

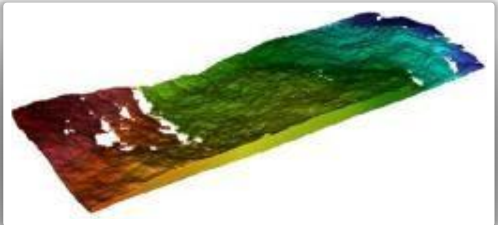

QMCPACK

Materials for Extreme Environments
(Paul Kent, ORNL)

- Find, predict and control materials and properties at the quantum level with unprecedented accuracy for the design novel materials that rely on metal to insulator transitions for high performance electronics, sensing, storage

Challenges: minimizing on-node memory usage; parallel on-node performance of Markov-chain Monte Carlo

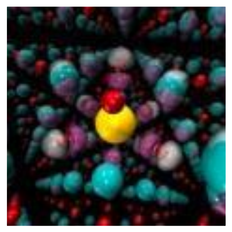

Exascale apps can deliver transformative products and solutions

ExaSGD

Reliable and Efficient Planning of the Power Grid
(Henry Huang, PNNL)

- Optimize power grid planning, operation, control and improve reliability and efficiency

Challenges: parallel performance of nonlinear optimization based on discrete algebraic equations and possible mixed-integer programming

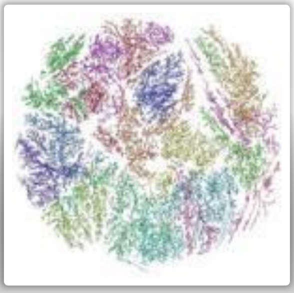

Combustion-PELE

High-Efficiency, Low-Emission Combustion Engine Design
(Jackie Chen, SNL)

- Reduce or eliminate current cut-and-try approaches for combustion system design

Challenges: performance of chemistry ODE integration on accelerated architectures; linear solver performance for low-Mach algorithm; explicit LES/DNS algorithm not stable

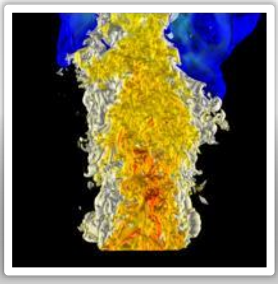

EXASCALE COMPUTING PROJECT

Exascale apps can deliver transformative products and solutions

E3SM-MMF

Accurate Regional Impact Assessment in Earth Systems
(Mark Taylor, SNL)

- Forecast water resources and severe weather with increased confidence; address food supply changes

Challenges: MMF approach for cloud-resolving model has large biases; adequacy of Fortran MPI+OpenMP for some architectures; Support for OpenMP and OpenACC

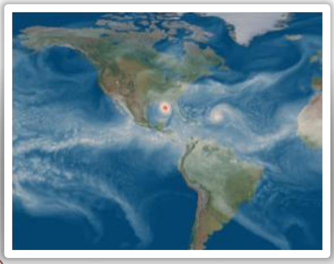

NWChemEx

Catalytic Conversion of Biomass-Derived Alcohols
(Thom Dunning, PNNL)

- Develop new optimal catalysts while changing the current design processes that remain costly, time consuming, and dominated by trial-and-error

Challenges: computation of energy gradients for coupled-cluster implementation; on- and off-node performance

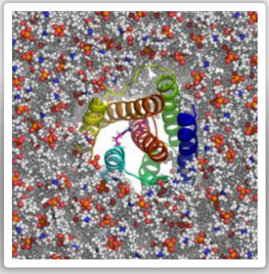

ExaBiome

Metagenomics for Analysis of Biogeochemical Cycles
(Kathy Yelick, LBNL)

- Discover knowledge useful for environmental remediation and the manufacture of novel chemicals and medicines

Challenges: Inability of message injection rates to keep up with core counts; efficient and performant implementation of UPC, UPC++, GASNet; GPU performance; I/O performance

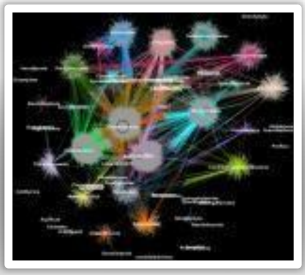

EXASCALE COMPUTING PROJECT

E3SM-Multiscale Modeling Framework (MMF)

Cloud Resolving Climate Model for E3SM

- Develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.

- Cloud resolving climate model is needed to reduce major systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection – such as convective storm systems

- Challenge: cloud resolving climate model using traditional approaches requires *zettascale* resources

- E3SM "conventional" approach:
  - Run the E3SM model with a global cloud resolving atmosphere and eddy resolving ocean.
    - 3 km atmosphere/land (7B grid points) and 15-5 km ocean/ice (1B grid points)
  - Achieve throughput rate of 5 SYPD to perform climate simulation campaigns including a 500 year control simulation
  - Detailed benchmarks on KNL and v100 GPUs show negligible speedups compared to conventional CPUs
    - Low arithmetic intensity of most of the code; throughput requirements lead to strong scaling and low work per node.

- E3SM-MMF: Use a multiscale approach ideal for new architectures to achieve cloud resolving convection on Exascale
  - Exascale will make "conventional" cloud resolving simulations routine for shorter simulations (process studies, weather prediction)
  - For cloud resolving climate simulations, we need fundamentally new approaches to take advantage of exascale resources

Convective storm system nearing the Chicago metropolitan area
http://www.spc.noaa.gov/misc/AbtDerechos/derechofacts.htm

PI: Mark Taylor (SNL)

Exascale apps can deliver transformative products and solutions

ExaSky

Cosmological Probe of the Standard Model of Particle Physics
(Salman Habib, ANL)

- Unravel key unknowns in the dynamics of the Universe: dark energy, dark matter, and inflation

Challenges: subgrid model accuracy; OpenMP performance on GPUs; file system stability and availabilty

LatticeQCD

Validate Fundamental Laws of Nature
(Andreas Kronfeld, FNAL)

- Correct light quark masses; properties of light nuclei from first principles; <1% uncertainty in simple quantities

Challenges: performance of critical slowing down; reducing network traffic to reduce system interconnect contention; strong scaling performance to mitigate reliance on checkpointing

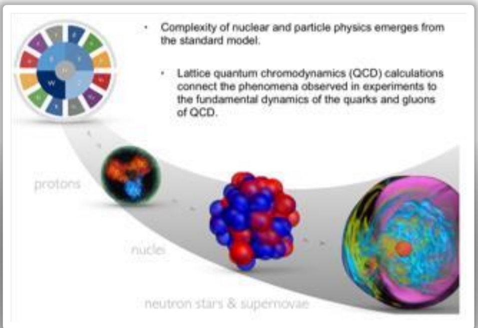

WarpX

Plasma Wakefield Accelerator Design
(Jean-Luc Vay, LBNL)

- Virtual design of 100-stage 1 TeV collider; dramatically cut accelerator size and design cost

Challenges: scaling of Maxwell FFT-based solver; maintaining efficiency of large timestep algorithm; load balancing

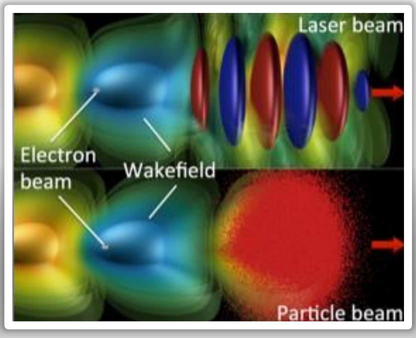

EXASCALE COMPUTING PROJECT

Exascale apps can deliver transformative products and solutions

WDMApp

High-Fidelity Whole Device Modeling of Magnetically Confined Fusion Plasmas
(Amitava Bhattacharjee, PPPL)

- Prepare for ITER exps and increase ROI of validation data and understanding
- Prepare for beyond-ITER devices

Challenges: robust, accurate, and efficient code-coupling algorithm; reduction in memory and I/O usage

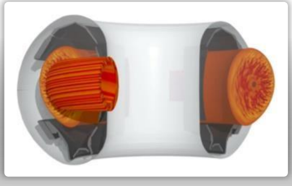

ExaStar

Demystify Origin of Chemical Elements
(Dan Kasen, LBNL)

- What is the origin of the elements?
- How does matter behave at extreme densities?
- What are the sources of gravity waves?

Challenges: delivering performance on accelerators; delivering fidelity for general relativity implementation

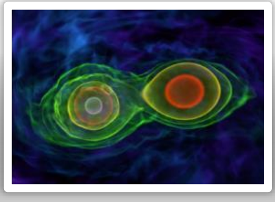

ExaFEL

Light Source-Enabled Analysis of Protein and Molecular Structure and Design
(Amadeo Perazzo, SLAC)

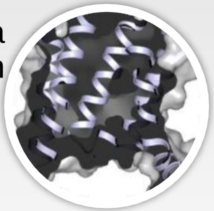

- Process data without beam time loss
- Determine nanoparticle size and shape changes
- Engineer functional properties in biology and materials science

Challenges: improving the strong scaling (one event processed over many cores) of compute-intensive algorithms (ray tracing, M-TIP) on accelerators

CANDLE

Accelerate and Translate Cancer Research
(Rick Stevens, ANL)

- Develop predictive preclinical models and accelerate diagnostic and targeted therapy through predicting mechanisms of RAS/RAF driven cancers

Challenges: increasing accelerator utilization for model search; effectively exploiting HP16; preparing for any data management or communication bottlenecks

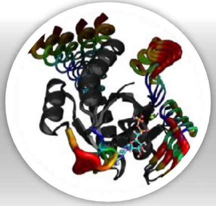

ECP EXASCALE COMPUTING PROJECT