

More Data, More Science and... Moore's Law

Kathy Yelick

**Professor of Electrical Engineering and Computer Sciences
University of California at Berkeley
Associate Laboratory Director for Computing Sciences
Lawrence Berkeley National Laboratory**

Three parts

- **Science opportunities**
- **Technology challenges**
- **Productivity approaches**

Part I

Science is poised for transformation


and Science in the broadest sense



Medicine



Basic Science



Engineering



Environment



Security

Experimental Science is Changing

By using our website you agree to our use of cookies in accordance with our cookie policy.

OK

PRIVACY POLICY +

JAX[®] MICE & SERVICES



JAX[®] Mice are the highest quality and most-published mouse models in the world. Take advantage of our large inventories of common inbred strains and the convenience of having your breeding and drug efficacy needs met by the leading experts in mouse modeling.

Search for Mice

Advanced Mice Search

Search for mice by strain, stock, gene, allele and synonyms




Breed Your Mouse

Test Your Drug

Cryopreserve Your Mouse

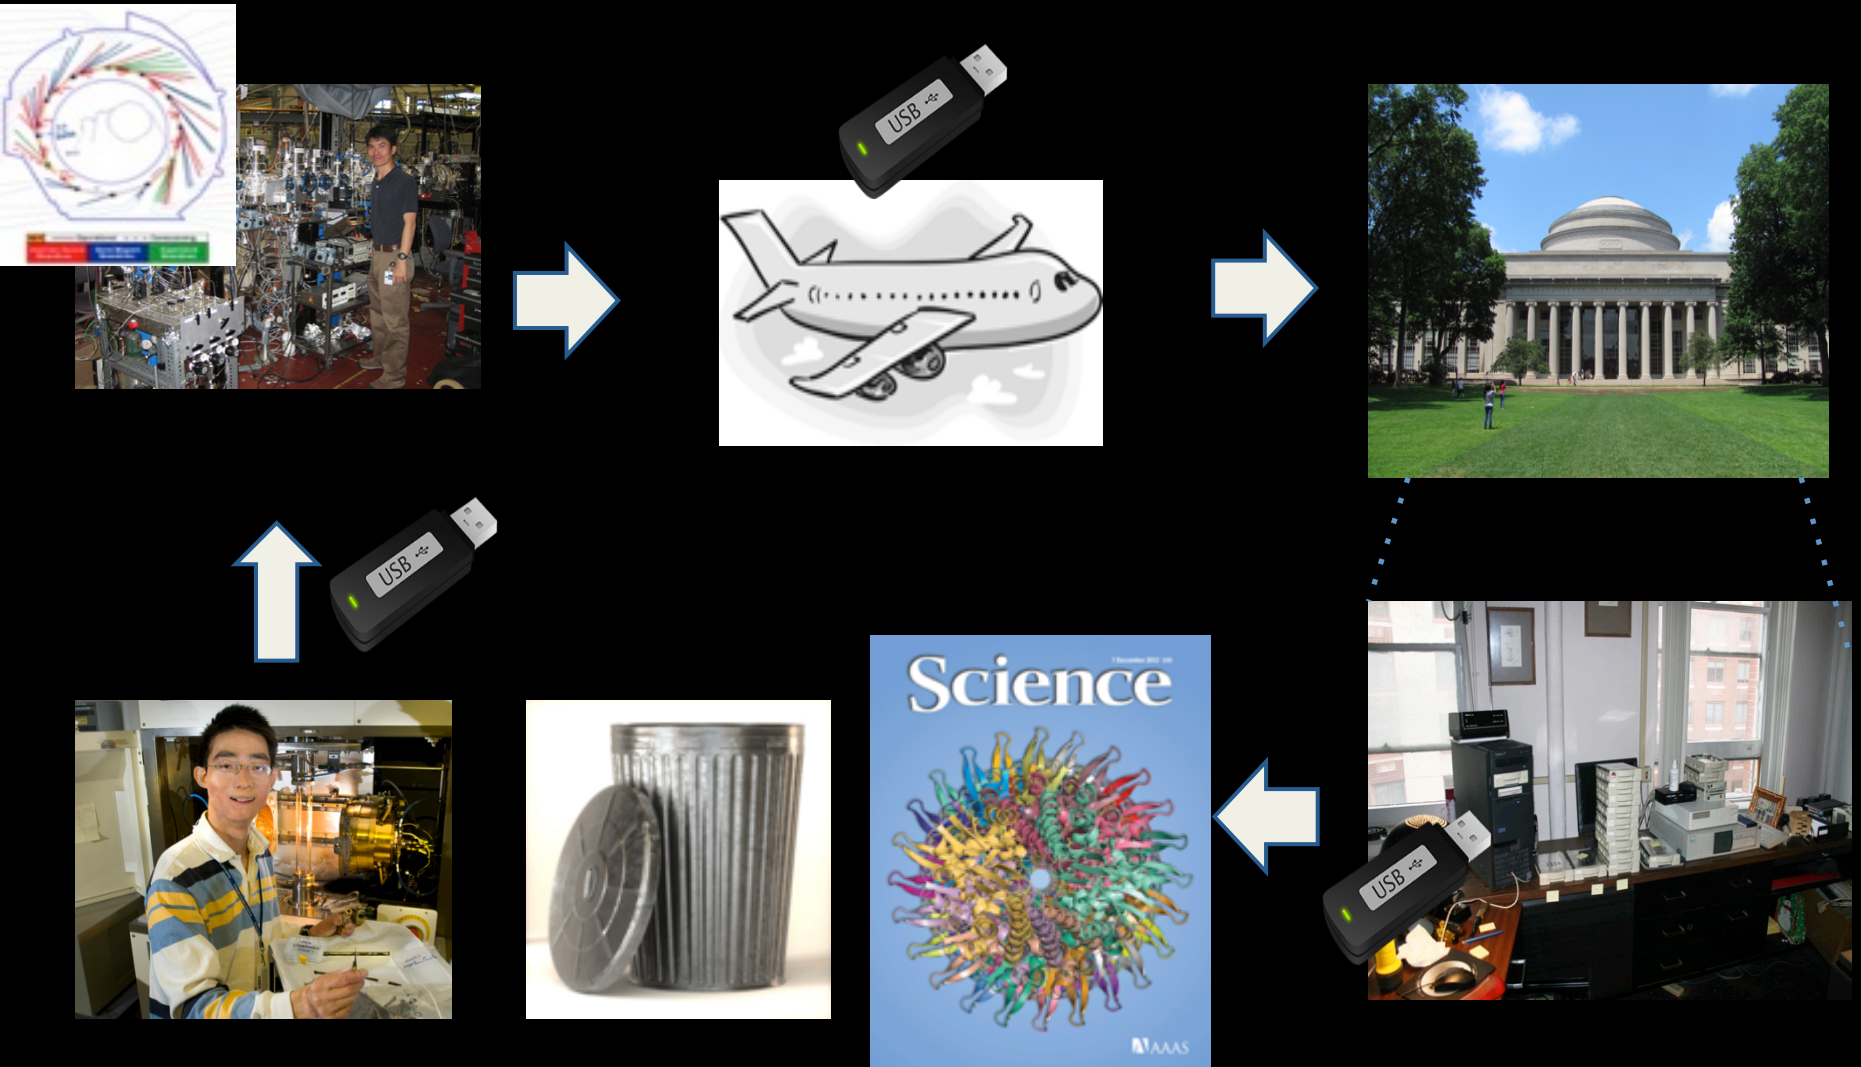
Empowering New Scientists



17-year-old Brittany Wegner creates breast cancer detection tool that is 99% accurate on a minimally invasive, previously inaccurate test.

Machine Learning + Online Data + Cloud Computing

Old School Experimental Science



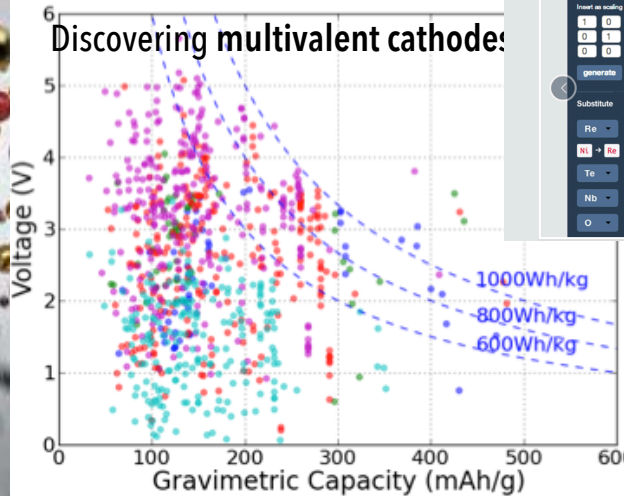
Computing, experiments, networking and expertise in a “Superfacility” for Science



Computational Scientists May Never Login



Materials Project





The screenshot shows the Materials Project website. The header has the title 'The Materials Project' and a description: 'Harnessing the power of supercomputing and state of the art electronic structure methods, the Materials Project provides open web-based access to computed information on known and predicted materials as well as powerful analysis tools to inspire and design novel materials.' There are buttons for 'Learn more', 'YouTube Tutorials', and 'Sign In or Register to start using'. Below the header, there is a section for 'Supercell Dimensions' with a table of values (a, b, c) and a 'generate' button. To the right, there is a 3D visualization of a crystal structure and a table of 'Lattice Parameters' and 'Fractional Coordinates'.

Novel computing technology here?



- Over 10,000 users of the Materials Project since 2011
- 1000x more computing needed for more materials, environments, and details

Simulations both consume and produce massive data




-
- **Simulation:** Resolve clouds, identify tipping points, quantify confidence, quantify agriculture and economics
 - **Observation:** Use environmental sensors and samples to refine models
 - **Analysis:** Machine learning identifies extreme events; Reanalysis of 100+ years of weather to improve
-


Breakthrough science will occur at the interface of observation and simulation




Materials and Chemistry





Subsurface science




Environment and Climate



Cosmology



Combustion



And it will need capable exascale systems, computer science, statistical machine learning and new mathematical models and algorithms

Challenge: Finding Scientific Data

Safari File Edit View History Bookmarks Window Help

www.google.com/search?tbs=sbi:AMhZZIu-Ft1o4xXlJhVjcLUv_1GtY_1M9gV_1hy

Google Google Maps Amazon News Popular

Berkeley Lab (...) TeamSnap :: M... Google CalMail - You... Search Results...

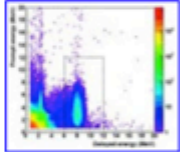
+You Search Images Mail Drive Calendar Sites Groups More

CalMail - You must be logged in to a page.

Google Antineutrinos.jpg

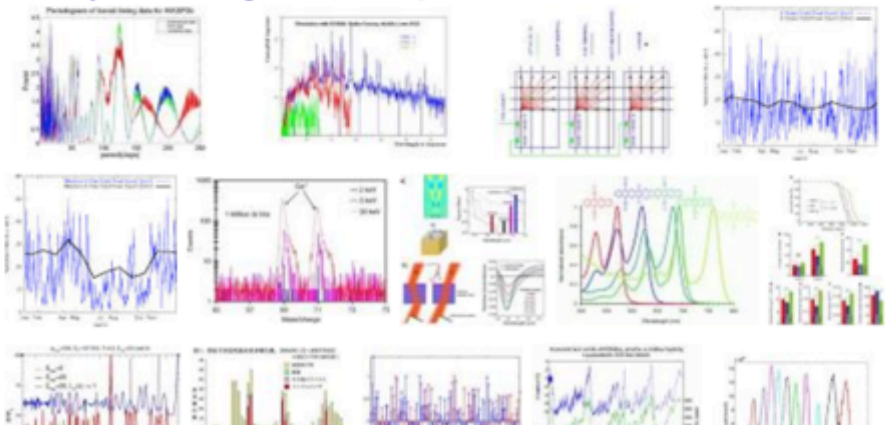
Web Images Maps Shopping More Search tools

Tip: Try entering a descriptive word in the search box.


 Image size:
153 × 133

No other sizes of this image found.

[Visually similar images](#) - Report images




Challenge: Filtering, De-Noising and Curating Data



AmeriFlux & FLUXNET: 750 users access carbon sensor data from 960 carbon flux data years

Arno Penzias and Robert Wilson discover Cosmic Microwave Background in 1965


The Black Swans of Computing with 1990s Technology



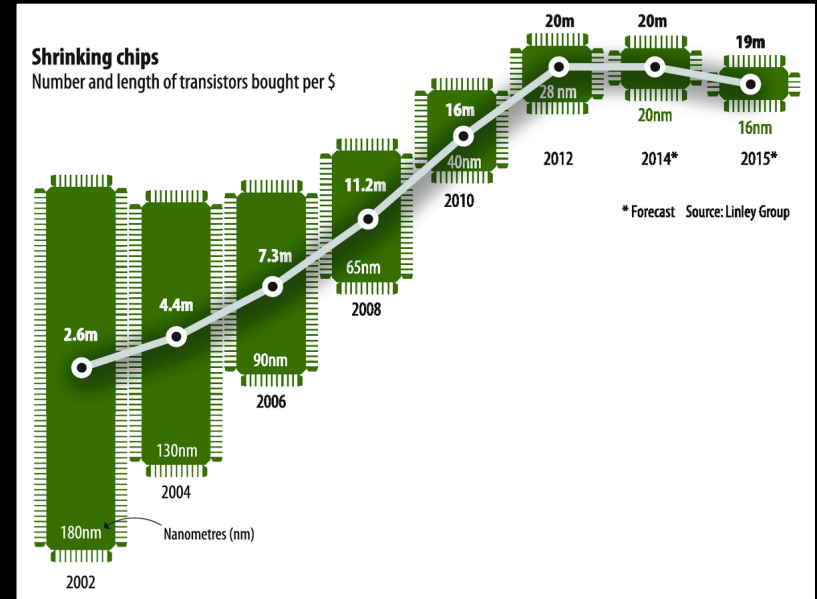
Part II

- **Science is poised for transformation**
- **Computing performance growth is slowing**

Computing Performance Stalls



2004: End of clock speed scaling



202?: End of transistor density scaling

We need an end game for Moore's Law and a new game after that

A Generic 2020+ Node Architecture

Memory Stacks on Package

Low Capacity, High Bandwidth, Software Control?

Lightweight Cores
(tiny, simple, massively parallel)
Throughput-Optimized

Bulky Cores

*Latency
Optimized*

DRAM

DRAM

DRAM

DRAM


NVRAM: Burst
Buffers / rack-local
storage (software
control)

Part III


- **Science is poised for transformation**
- **Computing performance growth is slowing**
- **How more productive tools can help**

(What productivity means to me)


The right machine for the problem




**Massive
Independent
Jobs for
Analysis and
Simulations**



**Nearest
Neighbor
Simulations**



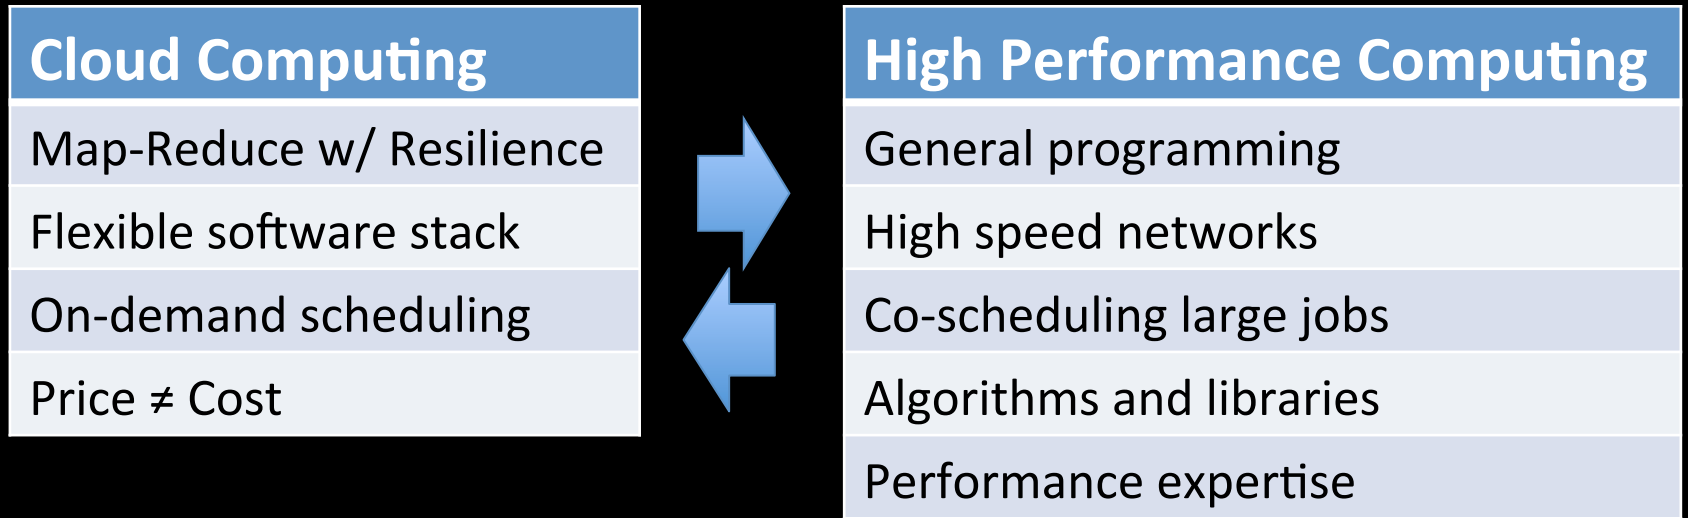
**All-to-All
Simulations**



**Random
access, large
data Analysis**

Applications often fit in multiple categories

The right systems software for the problem



Ideas and technologies have transferred in both directions

Ultimately, cloud computing is business model not a technology


A convenient programming model for the problem



De Novo Genome assembly for the environment and health

“Scaffolding” step on shared memory

- Human: 44 hours to 20 secs
- Wheat: “doesn’t run” to 32 secs



All 2015 human sequences could have been assembled on 10% of annual NERSC resources; metagenomes are much harder

A programming model (UPC) that encourages the right algorithmic thinking (hash tables, atomics,...)

A Network that Makes Location Irrelevant



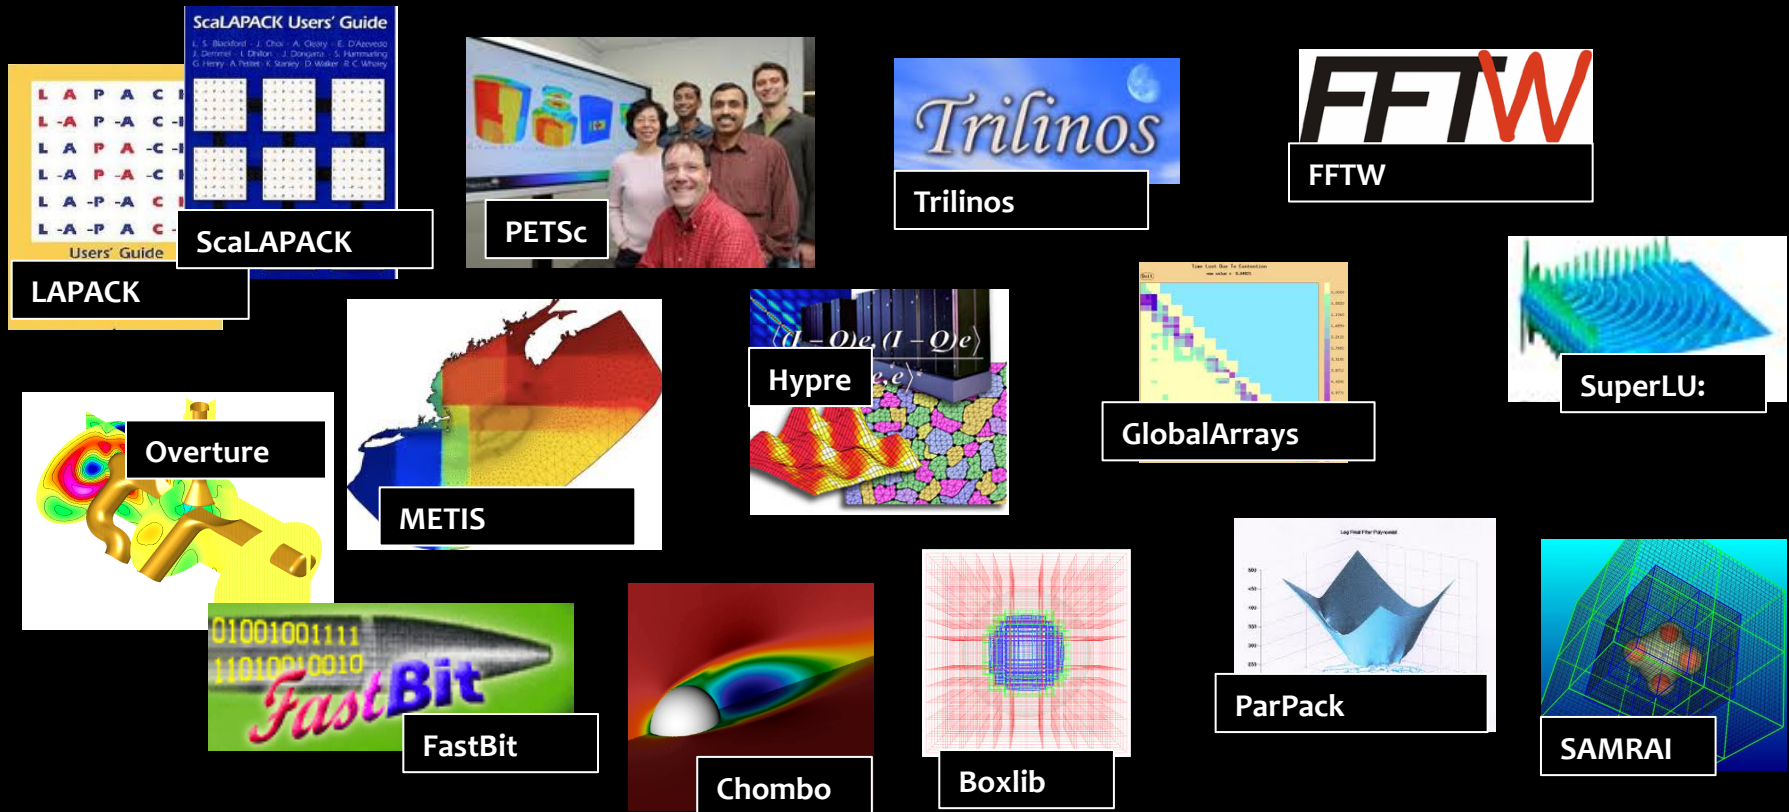
Technology, policy, and social challenges across:

- Networks and countries
- Facilities
- Universities (Science DMZ)
- Systems (SDN)

ESnet Goal: Scientific progress **completely unconstrained by the location** of people, instruments, computational resources, or data.

Use scientific libraries and frameworks

- *Encapsulate the best algorithms*
- *Hide machine details*

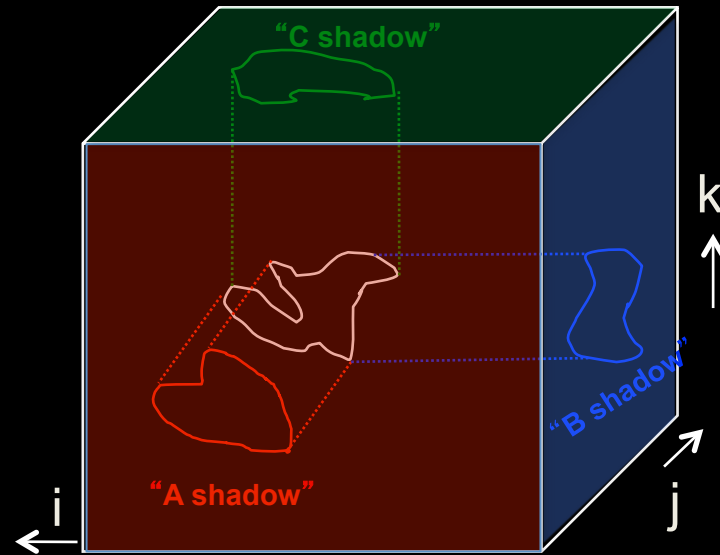



Enabled by portable programming models like MPI

The right libraries for the problem

7 Motifs of Simulation	7 Motifs of Data
Monte Carlo methods	Basic statistics
Particle methods	Generalized N-Body
Unstructured meshes	Graph-theory
Dense Linear Algebra	Linear algebra Optimizations
Sparse Linear Algebra	
Spectral methods	Integrations
Structured Meshes	Alignment

The right algorithms, models and methods



Optimize for data movement as first-class concern

- Change numerical methods (e.g., high order)
- Design algorithms that minimize data movement

High level programming abstractions



Where is Performance Portability?

- **Multiple generations of GPUs, Xeon Phi, ...**
 - Not performance portable across systems
- **How do applications address this problem?**

`#ifdef`

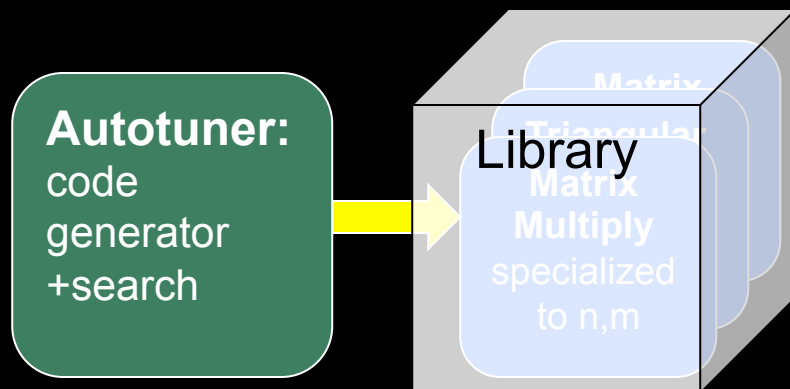
Approach 1: Use of libraries, templates and preprocessors

Approach 2: OpenMP4 absorbs necessary accelerator features (likely), but code still diverges

Approach 3: Architectures converge “enough” by 2023, perhaps with co-design help

Write code generators and autotuners

What code generators do we have?




Two unsolved compiler problems

- Dependence analysis
- Optimization selection

Dense Linear Algebra	Atlas
Spectral Algorithms	FFTW, Spiral
Particle	TCE
Sparse Linear Algebra	OSKI
Structured grids	TBD
Unstructured Grids	
Monte Carlo	

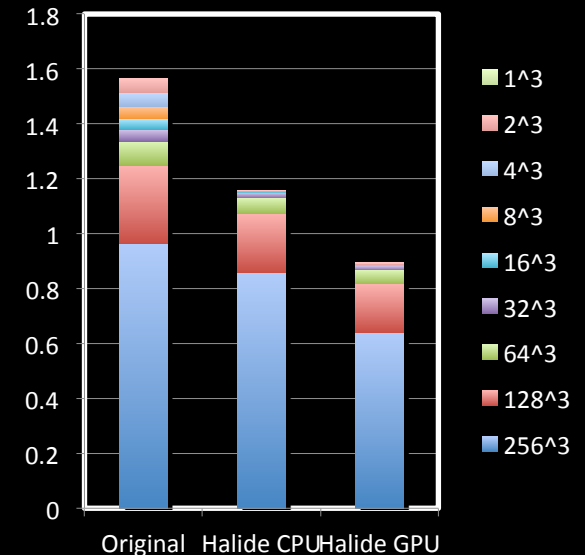
Develop and influence languages and compilers

Rose
compiler
input




Rose compiler output

Halide DSL has 10+ developers
and 50+ users in industry



Productivity will be key to dealing with complexity



Conclusions

NSCI is **not**:

- limited to the high end
- about integrating enough chips to reach exaflop

It **is** about:

- The cost of computing
- Power requirements, size, and speed of computers
- Solving real problems that require more computing