is paper describes objective technical results and analysis. Any subjective views or opinions that might be expres
the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Governm

Past, Present, and Future

SAND2020-0926C

Steve Plimpton
Sandia National Laboratories

SIAM Conference on Parallel Processing
for Scientific Computing (PP20)

Seattle - Feb 2020

CCR
Center for Computing Research

Sandia National Laboratories is a multi-mission laboratory managed and operated by
National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of

oratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia,
neywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract

Thanks to collaborators

- **LAMMPS** (molecular dynamics)
  Aidan Thompson, Dan Bolintineanu, Mike Brown, Paul Crozier, Gary Grest, Stephen Foiles, Bruce Hendrickson, Pieter in't Veld, Jeremy Lechman, Stan Moore, Ray Shan, Leo Silbert, Mark Stevens, Julian Tranchida, Christian Trott, Mitch Wood
- **SPARTA** (rarefied gas dynamics via DSMC)
  Michael Gallis, Tim Koehler, Stan Moore, John Torcynski, Alan Stagg
- **SPPARKS** (materials processing via on-lattice Monte Carlo)
  Aidan Thompson, Corbett Battaille, Liz Holm, John Mitchell, Veena Tikare, Ed Webb
- **MR-MPI** (MapReduce on top of MPI)
  Karen Devine, Jon Berry, Todd Plantenga
- **PHISH** (processing of streaming data)
  Tim Shead, Jon Berry, Cindy Phillips, Alexandra Porter

Thanks to collaborators

- **LAMMPS** (molecular dynamics)
  Aidan Thompson, Dan Bolintineanu, Mike Brown, Paul Crozier, Gary Grest, Stephen Foiles, Bruce Hendrickson, Pieter in't Veld, Jeremy Lechman, Stan Moore, Ray Shan, Leo Silbert, Mark Stevens, Julian Tranchida, Christian Trott, Mitch Wood
- **SPARTA** (rarefied gas dynamics via DSMC)
  Michael Gallis, Tim Koehler, Stan Moore, John Torcynski, Alan Stagg
- **SPPARKS** (materials processing via on-lattice Monte Carlo)
  Aidan Thompson, Corbett Battaille, Liz Holm, John Mitchell, Veena Tikare, Ed Webb
- **MR-MPI** (MapReduce on top of MPI)
  Karen Devine, Jon Berry, Todd Plantenga
- **PHISH** (processing of streaming data)
  Tim Shead, Jon Berry, Cindy Phillips, Alexandra Porter
- **Disciplines**:
  physics, chem eng, materials sci, biophysics, shock physics, polymers, soft matter & solid-state & granalar materials, computer sci, math, data analytics, combinatorics

Thanks to collaborators

- **LAMMPS** (molecular dynamics)
  Aidan Thompson, Dan Bolintineanu, Mike Brown, Paul Crozier, Gary Grest, Stephen Foiles, Bruce Hendrickson, Pieter in't Veld, Jeremy Lechman, Stan Moore, Ray Shan, Leo Silbert, Mark Stevens, Julian Tranchida, Christian Trott, Mitch Wood
- **SPARTA** (rarefied gas dynamics via DSMC)
  Michael Gallis, Tim Koehler, Stan Moore, John Torcynski, Alan Stagg
- **SPPARKS** (materials processing via on-lattice Monte Carlo)
  Aidan Thompson, Corbett Battaille, Liz Holm, John Mitchell, Veena Tikare, Ed Webb
- **MR-MPI** (MapReduce on top of MPI)
  Karen Devine, Jon Berry, Todd Plantenga
- **PHISH** (processing of streaming data)
  Tim Shead, Jon Berry, Cindy Phillips, Alexandra Porter
- **Disciplines**:
  physics, chem eng, materials sci, biophysics, shock physics, polymers, soft matter & solid-state & granalar materials, computer sci, math, data analytics, combinatorics

- DOE labs are great place for inter-disciplinary collaborations!

Past ghost of parallel computing

Past ghost of parallel computing

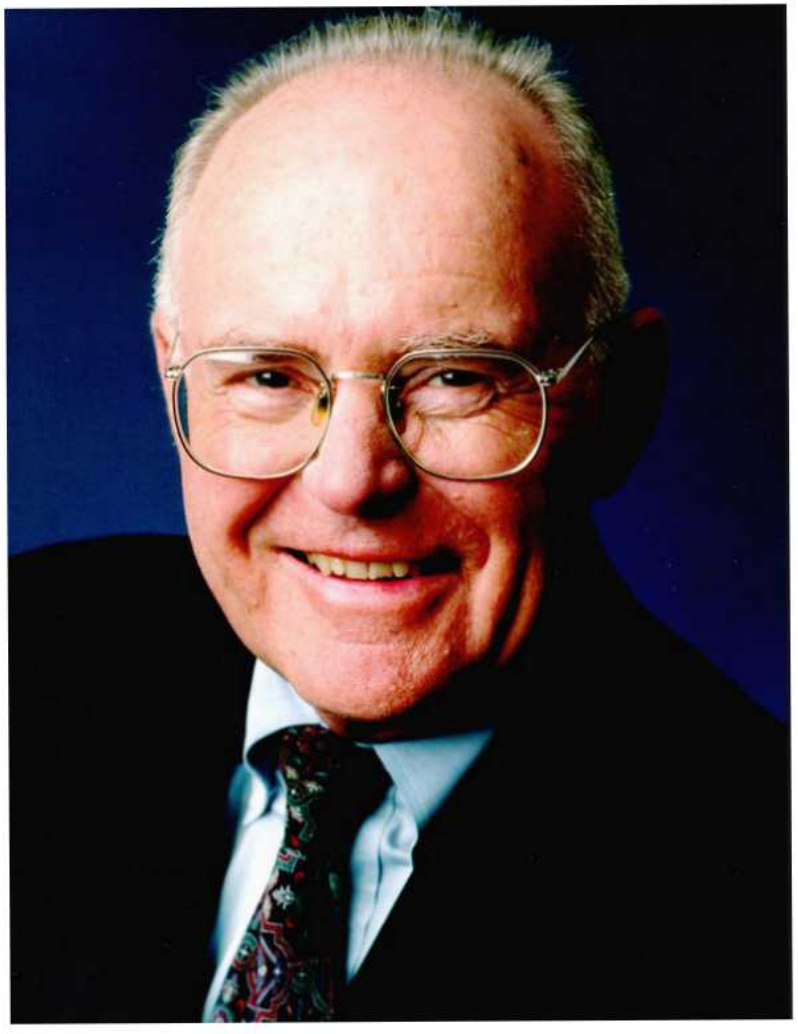

Gordon Moore

Past ghost of parallel computing

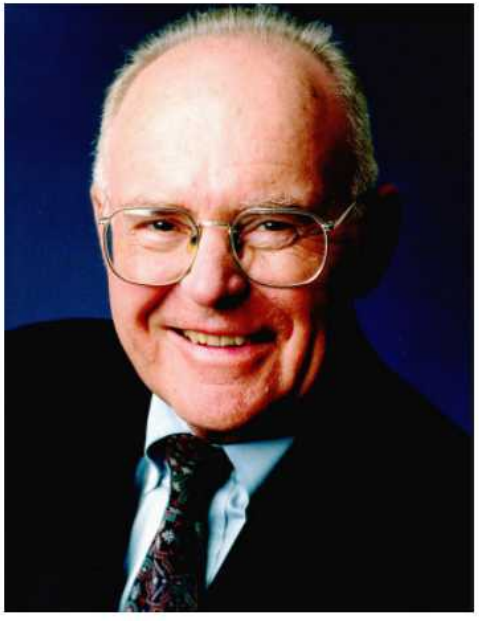

Gordon Moore

Moore's Law Law (due to Peter Lee):

The number of people predicting the **death** of Moore's Law doubles every two years

Past ghost of parallel computing

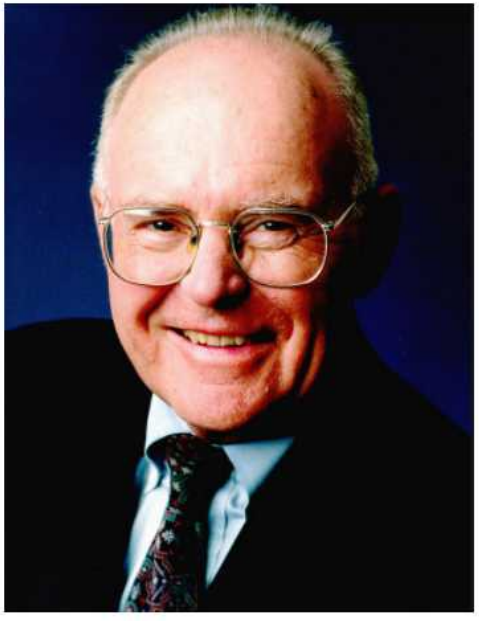

Gordon Moore

Moore's Law Law (due to Peter Lee):

The number of people predicting the **death** of Moore's Law doubles every two years

Moore's Law + parallelism have **dominated** HPC for 30 years:

Past ghost of parallel computing

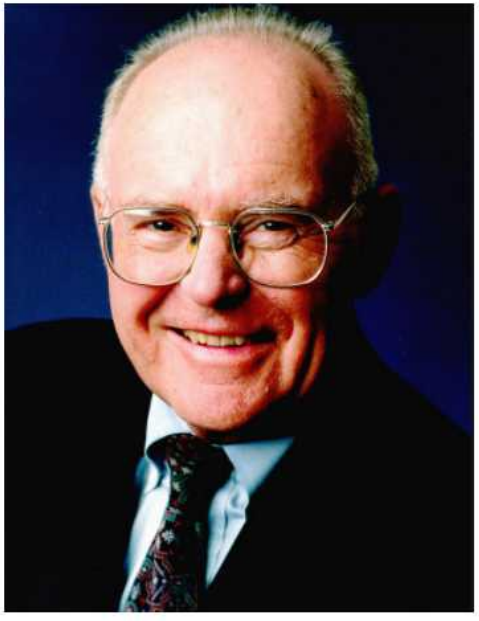

Gordon Moore

Moore's Law Law (due to Peter Lee):

The number of people predicting the **death** of Moore's Law doubles every two years

Moore's Law + parallelism have **dominated** HPC for 30 years:

- Gigaflop in 1988: Cray YMP
- Teraflop in 1997: ASCI Red (Cray)
- Petaflop in 2011: RoadRunner (IBM)
- Exaflop in $\sim$2021: several !

Past ghost of parallel computing

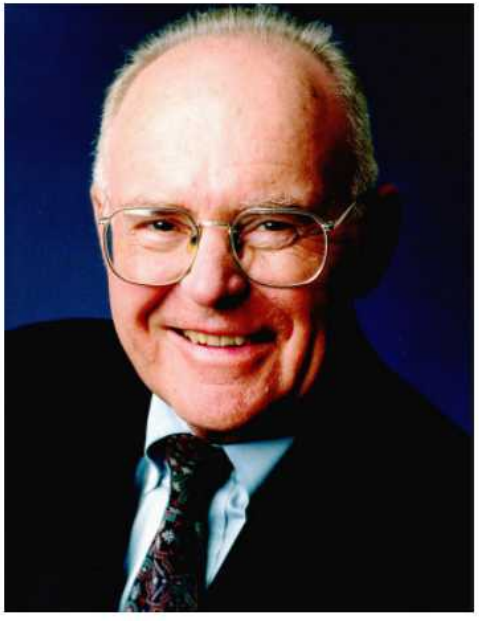

Gordon Moore

Moore's Law Law (due to Peter Lee):

The number of people predicting the **death** of Moore's Law doubles every two years

Moore's Law + parallelism have **dominated** HPC for 30 years:

- Gigaflop in 1988: Cray YMP
- Teraflop in 1997: ASCI Red (Cray)
- Petaflop in 2011: RoadRunner (IBM)
- Exaflop in $\sim$2021: several !

Doubling Linpack flops every 14 months for 33 years !

Present ghost of parallel computing

Present ghost of parallel computing

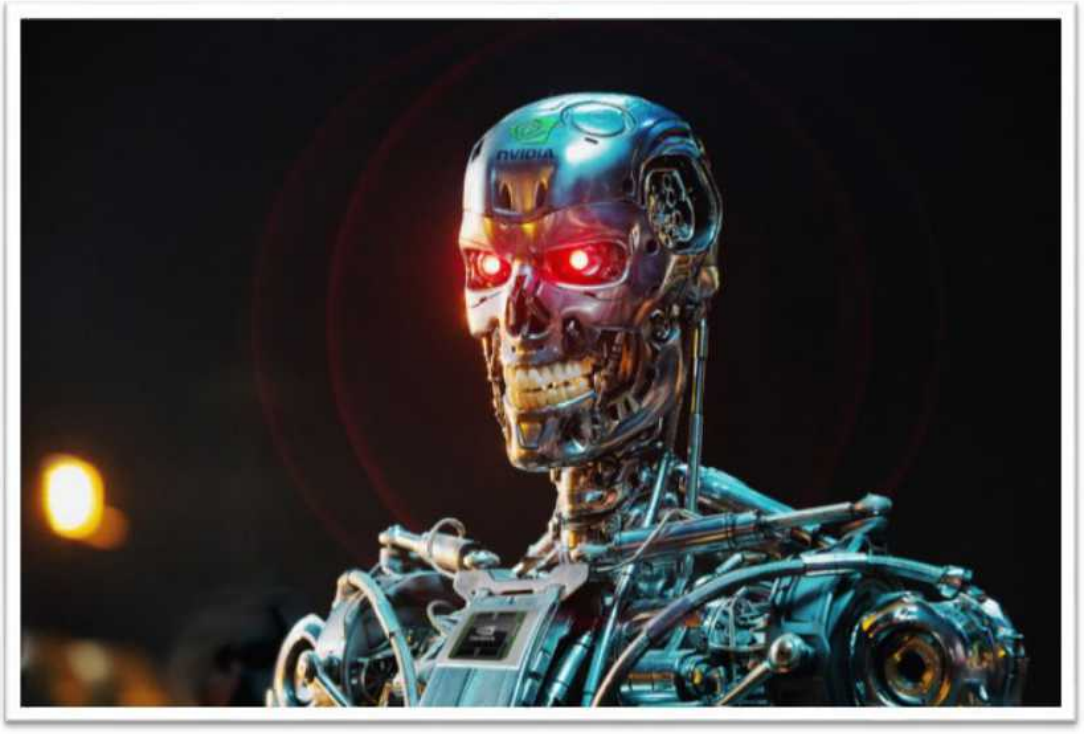

NVIDIA GPUs

Present ghost of parallel computing

NVIDIA GPUs

GPUs are **terminating** the competition

2007: GP-GPU + Cuda $\Rightarrow$
scientific computing

Then: double precision, FP error correction,
multi-GPU nodes, NVLink, etc ...

Present ghost of parallel computing

NVIDIA GPUs

GPUs are **terminating** the competition
2007: GP-GPU + Cuda $\Rightarrow$
scientific computing
Then: double precision, FP error correction,
multi-GPU nodes, NVLink, etc ...

- Largest HPC platforms are **increasingly** GPU-based
- All 4 next big DOE machines: ALCF, ORNL, NERSC, LLNL

Present ghost of parallel computing

NVIDIA GPUs

GPUs are **terminating** the competition
2007: GP-GPU + Cuda $\Rightarrow$
scientific computing
Then: double precision, FP error correction,
multi-GPU nodes, NVLink, etc ...

- Largest HPC platforms are **increasingly** GPU-based
- All 4 next big DOE machines: ALCF, ORNL, NERSC, LLNL
- Intel will now speak the word GPU in public !

Present ghost of parallel computing

NVIDIA GPUs

GPUs are **terminating** the competition
2007: GP-GPU + Cuda $\Rightarrow$
scientific computing
Then: double precision, FP error correction,
multi-GPU nodes, NVLink, etc ...

- Largest HPC platforms are **increasingly** GPU-based
- All 4 next big DOE machines: ALCF, ORNL, NERSC, LLNL
- Intel will now speak the word GPU in public !

- Countless SIAM PP and CSE talks in last 10 years on GPUs
- How many of you have worked on GPU code or algorithms ?

What have 2 ghosts meant for computational science?

Good news:

What have 2 ghosts meant for computational science?

Good news:

- Free speed-up!
- More science even if don't change algorithms or code
- Haven't had to think **hard** about new algorithms

What have 2 ghosts meant for computational science?

Good news:

- Free speed-up!
- More science even if don't change algorithms or code
- Haven't had to think **hard** about new algorithms

- Linpack: 1 BG/Q core / 1 Cray YMP proc = **41x** !

What have 2 ghosts meant for computational science?

Good news:

- Free speed-up!
- More science even if don't change algorithms or code
- Haven't had to think **hard** about new algorithms

- Linpack: 1 BG/Q core / 1 Cray YMP proc = **41x** !
- That's just 1988 to 2012 ... GPUs have gone further

What have 2 ghosts meant for computational science?

Bad news: (good news was partially fake news)

What have 2 ghosts meant for computational science?

Bad news: (good news was partially fake news)

- Lots of work to adapt algorithms and codes (MPI, GPUs)
- Creating machines that are harder to use and program

What have 2 ghosts meant for computational science?

Bad news: (good news was partially fake news)

- Lots of work to adapt algorithms and codes (MPI, GPUs)
- Creating machines that are harder to use and program
- Darwinian selection:
  Down-selecting for apps that run well
  on less general-purpose machines
- Opportunity cost:
  This is all work **not spent**
  devising new algorithms or doing science

Imbalance ratios over 30 years

- **Local balance** = flops to pay for 1 on-node word (8 bytes)
- **Remote balance** = flops to pay for 1 off-node word

Imbalance ratios over 30 years

- **Local balance** = flops to pay for 1 on-node word (8 bytes)
- **Remote balance** = flops to pay for 1 off-node word

Year	Machine	Linpack	Flops/ local	Flops/ remote
1988	Cray YMP	2.1 Gflops	0.5	0.5
1997	ASCI Red (SNL)	1.6 Tflops	8.3	20
2011	RoadRunner (LANL)	1.0 Pflops	6.7	170
2011	K-Computer (Japan)	11 Pflops	15	95
2012	Sequoia (LLNL)	17 Pflops	32	160
2013	Titan (ORNL)	18 Pflops	29	490
2016	Sunway TaihuLight (China)	93 Pflops	130	1500
2018	Summit (ORNL)	149 Pflops	45	1300
2021	Aurora (ANL)	1.0 Eflops	120	1600
2022	Frontier (ORNL)	1+ Eflops	200	3500

Most impact: faster machines or better algorithms ?

What has impacted your discipline's science more in last 30 yrs?
Billion X more compute power or **better algorithms** ?

Most impact: faster machines or better algorithms ?

What has impacted your discipline's science more in last 30 yrs?
Billion X more compute power or **better algorithms** ?

- Algorithms **is not** adapting to run MPI parallel or on GPUs
- Algorithms **is** new models, new numerical methods,
  lower complexity methods

Most impact: faster machines or better algorithms ?

What has impacted your discipline's science more in last 30 yrs?
Billion X more compute power or **better algorithms** ?

- Algorithms **is not** adapting to run MPI parallel or on GPUs
- Algorithms **is** new models, new numerical methods, lower complexity methods

- Many big algorithmic breakthroughs $>$ 30 years ago
  - **1986**: Nlog(N) N-body tree codes, Barnes & Hut
  - **1988**: Nlog(N) particle/mesh FFT solvers, Hockney & Eastwood
  - **1987**: O(N) fast multipole, Greengard & Rokhlin
  - **1977**: O(N) multigrid, Achi Brandt

Most impact: faster machines or better algorithms ?

What has impacted your discipline's science more in last 30 yrs?
Billion X more compute power or **better algorithms** ?

- Algorithms **is not** adapting to run MPI parallel or on GPUs
- Algorithms **is** new models, new numerical methods, lower complexity methods

- Many big algorithmic breakthroughs $>$ 30 years ago
  - **1986**: Nlog(N) N-body tree codes, Barnes & Hut
  - **1988**: Nlog(N) particle/mesh FFT solvers, Hockney & Eastwood
  - **1987**: O(N) fast multipole, Greengard & Rokhlin
  - **1977**: O(N) multigrid, Achi Brandt

Asked experts in climate, PIC, DFT, CFD, MD ...

My answer for classical molecular dynamics

Materials and bioscience modeling at the atomic scale

My answer for classical molecular dynamics

Materials and bioscience modeling at the atomic scale

Faster machines have had an **overwhelmingly** greater impact

My answer for classical molecular dynamics

Materials and bioscience modeling at the atomic scale

Faster machines have had an **overwhelmingly** greater impact

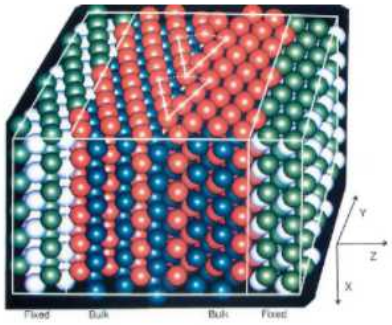

- (1989) my thesis
- 1000 atoms, 50K steps,
  single grain boundary
- Few hours on a YMP proc

My answer for classical molecular dynamics

Materials and bioscience modeling at the atomic scale

Faster machines have had an **overwhelmingly** greater impact

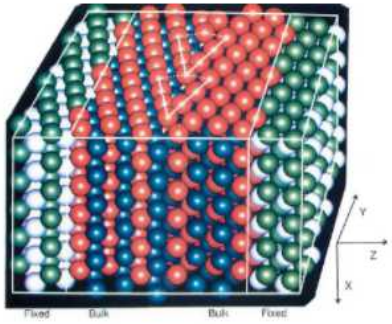

- (1989) my thesis
- 1000 atoms, 50K steps, single grain boundary
- Few hours on a YMP proc

- (2017) LLNL Nature paper
- 2B atoms, 0.5B steps, 700K grain boundaries
- Weeks on 1/3 of Sequoia

My answer for classical molecular dynamics

Materials and bioscience modeling at the atomic scale

Faster machines have had an **overwhelmingly** greater impact

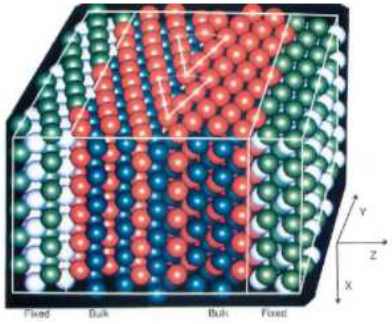

- (1989) my thesis
- 1000 atoms, 50K steps, single grain boundary
- Few hours on a YMP proc

- (2017) LLNL Nature paper
- 2B atoms, 0.5B steps, 700K grain boundaries
- Weeks on 1/3 of Sequoia

same MD algorithm, **same** material model (1984)

What will the near future be like?

- **Specialization of SoCs** = systems on a chip
- Enabled by easy design tools in growing ARM ecosystem
- CPUs with **multiple IP blocks** for scientific computing
  - each IP block = 1 kernel, runs at ASIC speed, e.g. TPU
  - traditional: matmul, matvec, FFT, MPI protocols, etc
  - exotic: neuromorphic, quantum computing ?
  - latest iPhone chip has 40+ IP blocks
    - processing speech, video, images, GPS, etc

What will the near future be like?

- **Specialization of SoCs** = systems on a chip
- Enabled by easy design tools in growing ARM ecosystem
- CPUs with **multiple IP blocks** for scientific computing
  - each IP block = 1 kernel, runs at ASIC speed, e.g. TPU
  - traditional: matmul, matvec, FFT, MPI protocols, etc
  - exotic: neuromorphic, quantum computing ?
  - latest iPhone chip has 40+ IP blocks
    - processing speech, video, images, GPS, etc
- **Specialization** of HPC machines **for individual disciplines**
  - fusion science builds an HPC to model ITER (MHD or PIC)
  - machines for MD (Anton), for PIC, for CFD, etc
  - machines for particles, for sparse FE, etc
  - trading off speed versus generality

What will the near future be like?

- **Specialization of SoCs** = systems on a chip
- Enabled by easy design tools in growing ARM ecosystem
- CPUs with **multiple IP blocks** for scientific computing
  - each IP block = 1 kernel, runs at ASIC speed, e.g. TPU
  - traditional: matmul, matvec, FFT, MPI protocols, etc
  - exotic: neuromorphic, quantum computing ?
  - latest iPhone chip has 40+ IP blocks
    - processing speech, video, images, GPS, etc
- **Specialization** of HPC machines **for individual disciplines**
  - fusion science builds an HPC to model ITER (MHD or PIC)
  - machines for MD (Anton), for PIC, for CFD, etc
  - machines for particles, for sparse FE, etc
  - trading off speed versus generality
- Where are **experts** who can program/use these machines?

What will the near future be like?

- **Specialization of SoCs** = systems on a chip
- Enabled by easy design tools in growing ARM ecosystem
- CPUs with **multiple IP blocks** for scientific computing
  - each IP block = 1 kernel, runs at ASIC speed, e.g. TPU
  - traditional: matmul, matvec, FFT, MPI protocols, etc
  - exotic: neuromorphic, quantum computing ?
  - latest iPhone chip has 40+ IP blocks
    - processing speech, video, images, GPS, etc
- **Specialization** of HPC machines **for individual disciplines**
  - fusion science builds an HPC to model ITER (MHD or PIC)
  - machines for MD (Anton), for PIC, for CFD, etc
  - machines for particles, for sparse FE, etc
  - trading off speed versus generality
- Where are **experts** who can program/use these machines?
- In this room ! Good news for **SIAM** members !

But will the next 30 years be like the last 30 ?

exa $= 10^{18}$ flops
zetta $= 10^{21}$
yotta $= 10^{24}$
ronna $= 10^{27}$ (proposed)

But will the next 30 years be like the last 30 ?

exa $= 10^{18}$ flops
zetta $= 10^{21}$
yotta $= 10^{24}$
ronna $= 10^{27}$ (proposed)

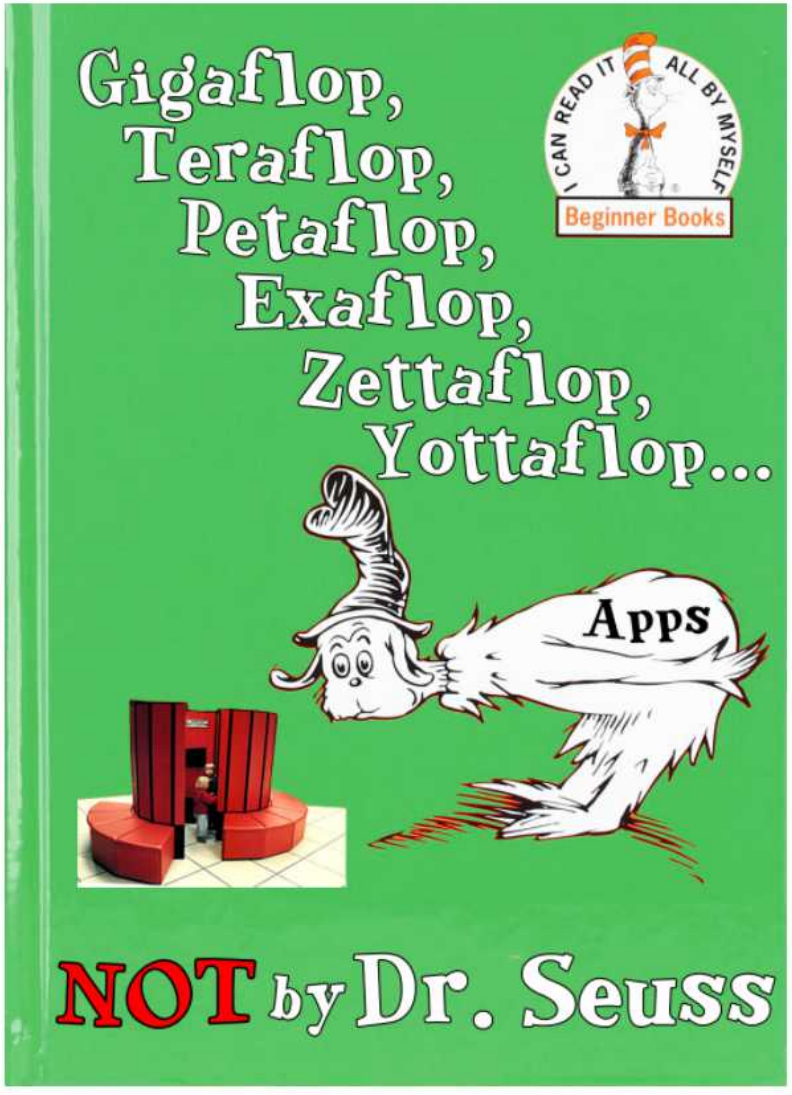

But will the next 30 years be like the last 30 ?

exa $= 10^{18}$ flops
zetta $= 10^{21}$
yotta $= 10^{24}$
ronna $= 10^{27}$ (proposed)

Abbott and Costello routine:

DOE: We'd like funding for a new yottaflop machine.
Congress: I know it's a lotta flops, but how many exactly ?
DOE: I told you, it's a yottaflops.

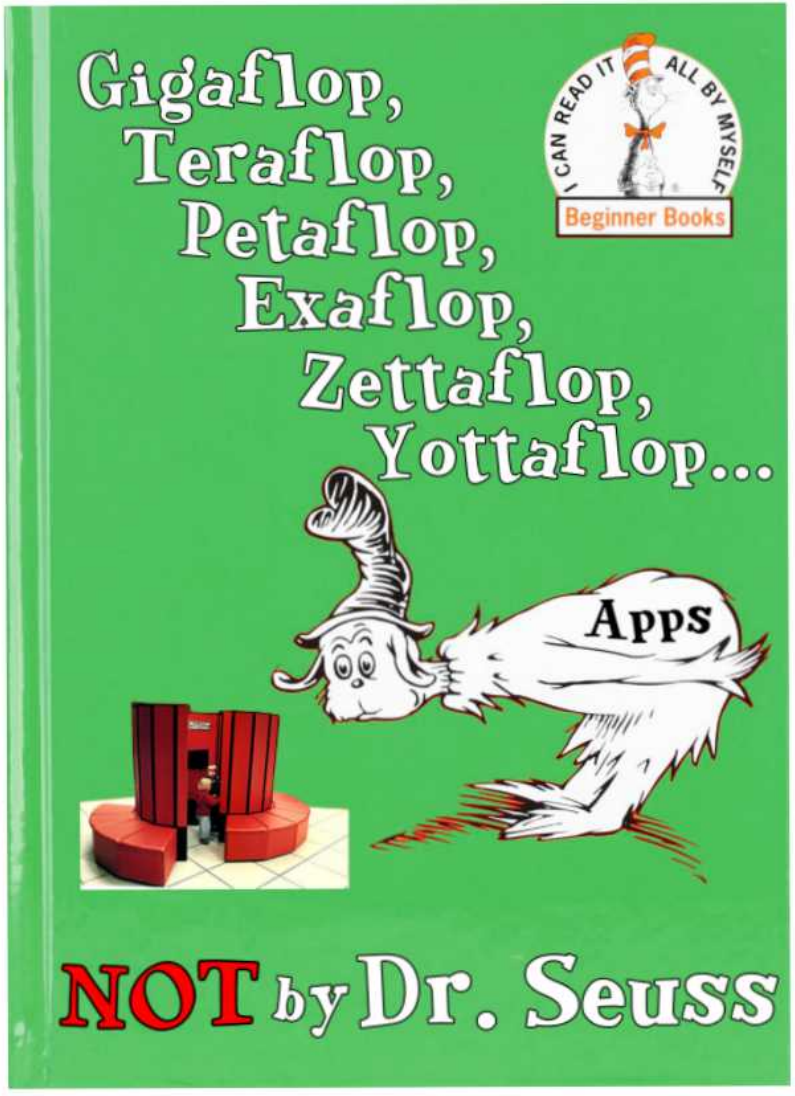

But will the next 30 years be like the last 30 ?

exa $= 10^{18}$ flops
zetta $= 10^{21}$
yotta $= 10^{24}$
ronna $= 10^{27}$ (proposed)

Abbott and Costello routine:

DOE: We'd like funding for a
new yottaflop machine.
Congress: I know it's a lotta flops,
but how many exactly ?
DOE: I told you, it's a yottaflops.
Congress: OK, let me guess,
will it also cost a yotta $$?
DOE: No, not a yotta $$,
but yes, a whole lotta $$

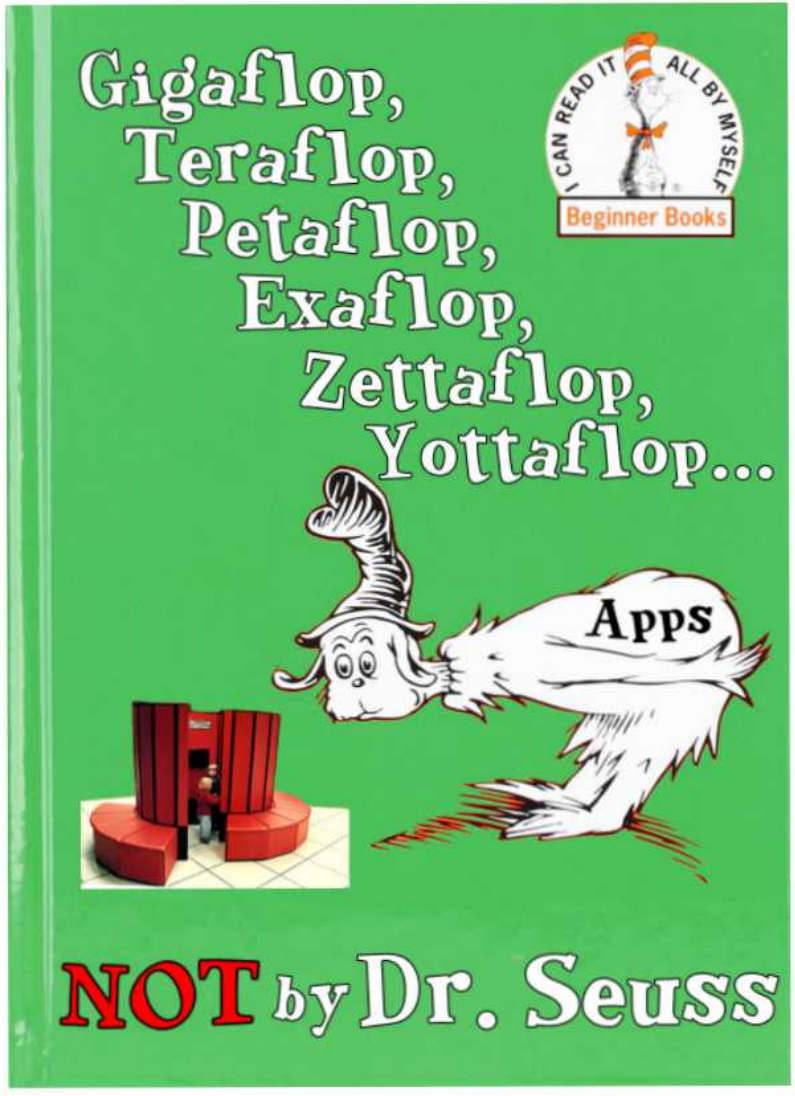

Can we afford HPC machines beyond exaflops?

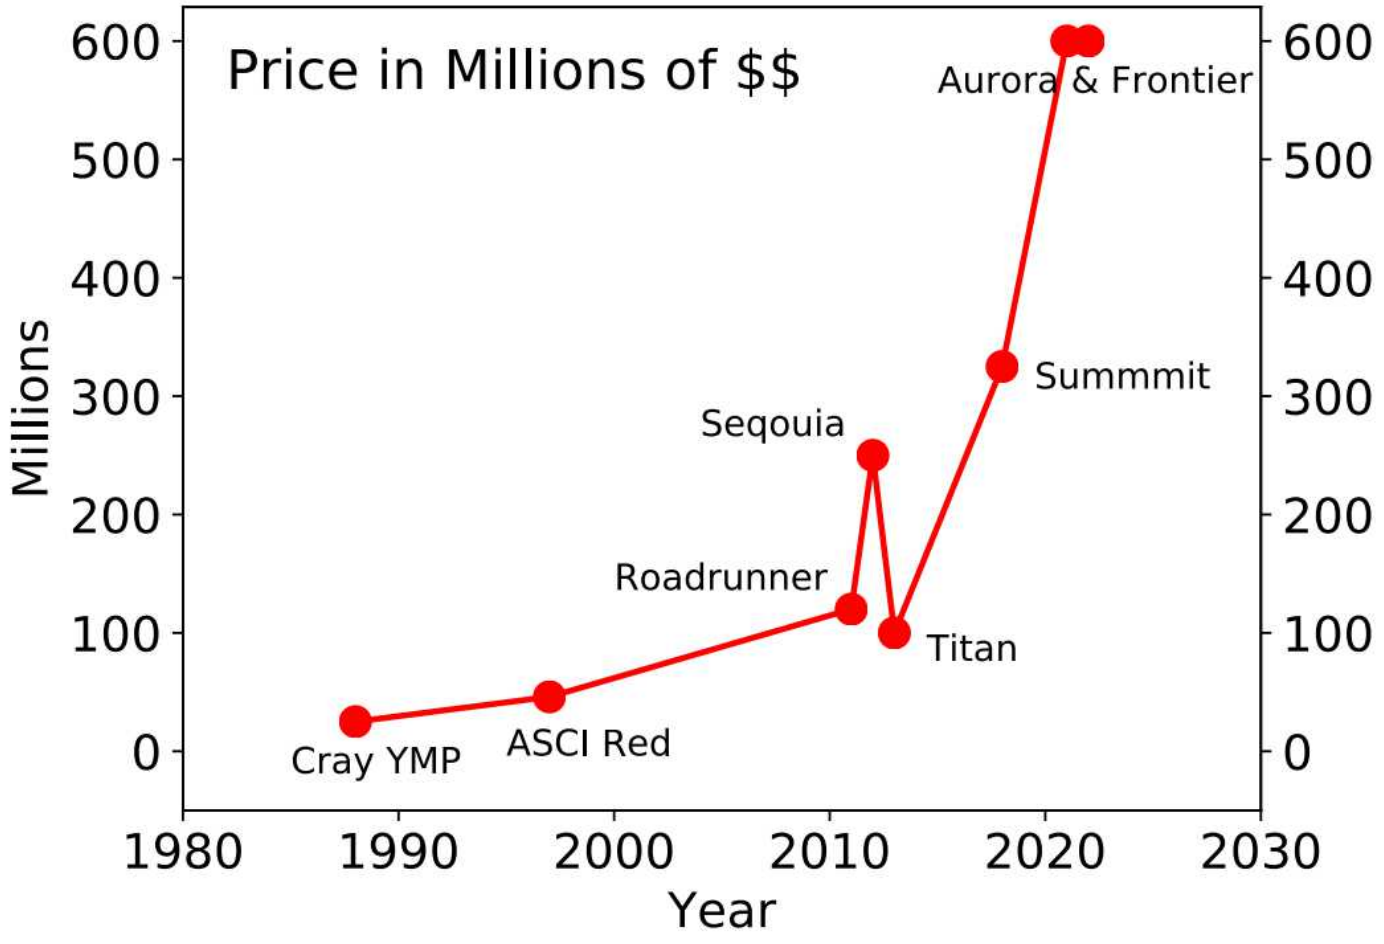

Can we afford HPC machines beyond exaflops?

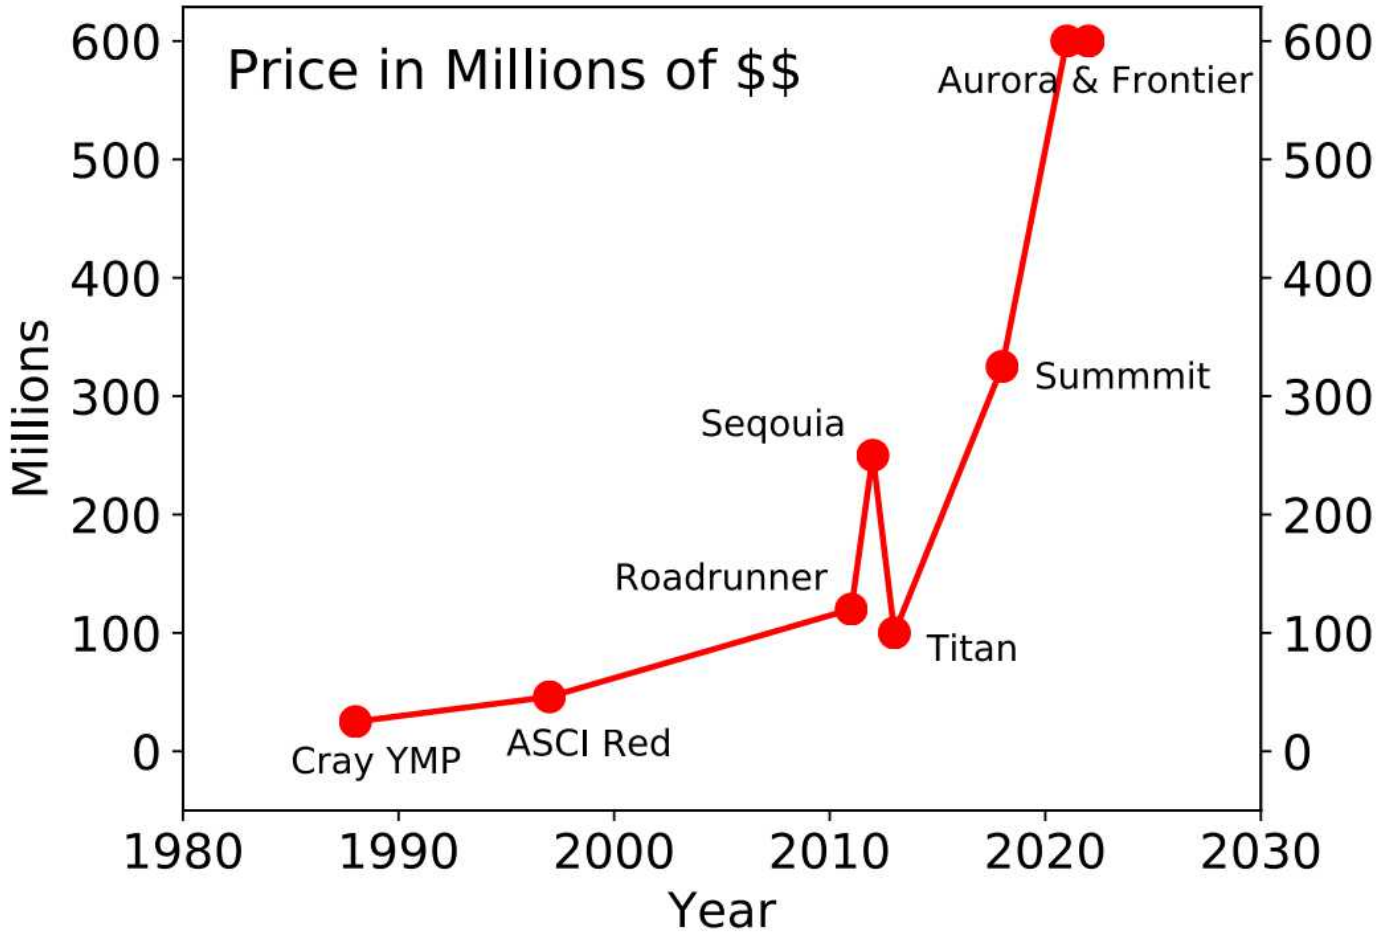

Will anyone pay a billion $$ for an HPC machine ?

Candidate #1 for **Future** ghost of parallel computing

Neuromorphic computing

Carver Mead

Data

Candidate #1 for **Future** ghost of parallel computing

Neuromorphic computing

Carver Mead

Data

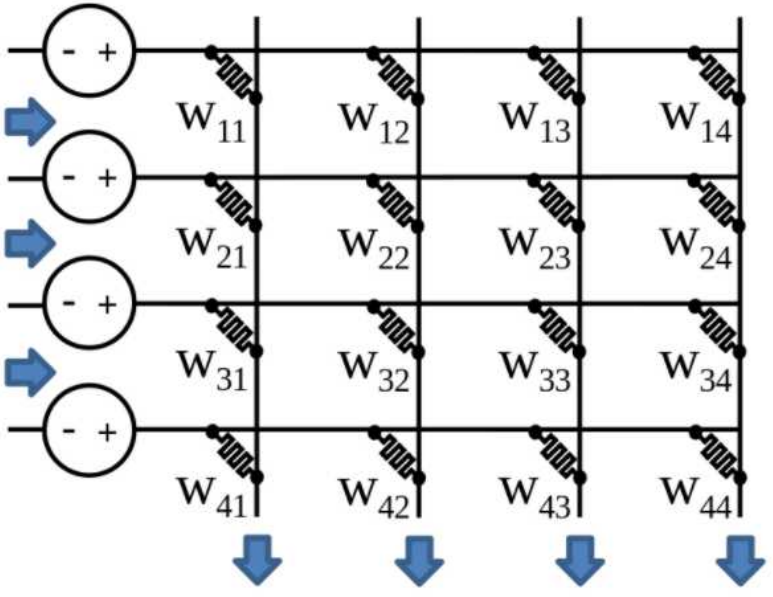

Candidate #1 for **Future** ghost of parallel computing

Neuromorphic computing

Carver Mead

Data

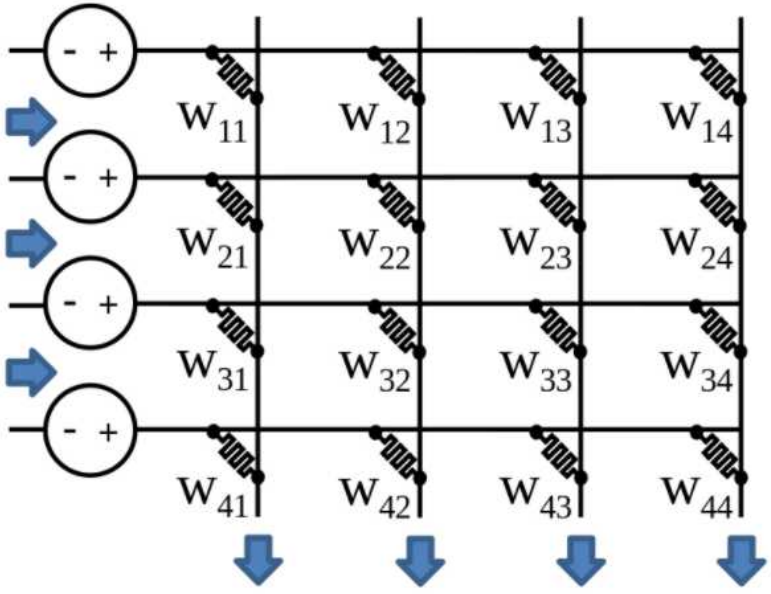

- Memristor-based **matvec** in one cycle !
- ML: high voltages to **train**, lower voltages to **classify**
- Promise: low power, very fast (N^2 ops per clock), small size
- Challenges: materials, reproducibility, matrix size, sparsity
- Low-precision: maybe **4-8 bits** at best

Candidate #2 for **Future** ghost of parallel computing

Quantum computing

R Feynman

Peter Shor

Candidate #2 for **Future** ghost of parallel computing

Quantum computing

R Feynman

Peter Shor

Shor's algorithm:

- integer factorization
- polynomial in log(N)
- (nearly) exponentially faster than classical factoring

Candidate #2 for **Future** ghost of parallel computing

Quantum computing

R Feynman

Peter Shor

Shor's algorithm:

- integer factorization
- polynomial in log(N)
- (nearly) exponentially faster than classical factoring

- Promise: maybe **awesome** in 2 domains ⇒ crypto, QM calcs
- Challenges: materials, scale up, how to program or do I/O
- PR problem: How do you count flops ?

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

IBM responds: actually only 2.5 days on Summit (worst case), and Summit could do it more accurately

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

IBM responds: actually only 2.5 days on Summit (worst case), and Summit could do it more accurately

What is X ?? generating "certifiable" **random numbers**

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

IBM responds: actually only 2.5 days on Summit (worst case), and Summit could do it more accurately

What is X ?? generating "certifiable" **random numbers**

Linpack
Benchmark

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

IBM responds: actually only 2.5 days on Summit (worst case), and Summit could do it more accurately

What is X ?? generating "certifiable" **random numbers**

Linpack
Benchmark

TOP 500
The List.

Has quantum supremacy already been achieved ?

Quantum supremacy: when a quantum computer performs some computation that is intractable on a classical HPC machine

Oct 2019: **Google** claims QS, in a Nature paper
53 cubits computed X in 200 secs, versus 10000 years

IBM responds: actually only 2.5 days on Summit (worst case), and Summit could do it more accurately

What is X ?? generating "certifiable" **random numbers**

Linpack Benchmark

TOP 500 The List.

https://www.top500.org

Candidate #3 for **Future** ghost of parallel computing

Machine learning and AI - 1st ghost that isn't hardware !

John Hopfield

Geoffrey Hinton

Yann LeCun

Candidate #3 for **Future** ghost of parallel computing

Machine learning and AI - 1st ghost that isn't hardware !

John Hopfield

Geoffrey Hinton

Yann LeCun

- Some say it's the **next revolution**:
  - as of 2019, DOE has an AI technology office
  - "AI won't replace scientists and engineers,
    but scientists and engineers who use AI
    will replace those who don't" (D Womble & Microsoft)
- Some say it's already **over-hyped**:
  - it's just glorified fitting
  - it's re-packaged linear algebra and optimization methods

Machine learning for materials modeling

- **Holy grail**: classical potentials with quantum accuracy
  - $O(N)$ cost for $O(N^3)$ accuracy

Machine learning for materials modeling

- **Holy grail**: classical potentials with quantum accuracy
  - $O(N)$ cost for $O(N^3)$ accuracy
- **Idea**: train NN potentials on quantum data, i.e. Ax = b

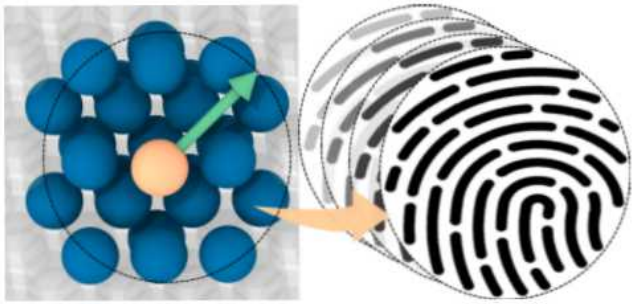

Input: geometric descriptors of an atom's neighborhood
Output: energy/forces on atom
Truth: quantum DFT forces

- Issues: Configs to train on, quantify errors, transferable ?

Machine learning for materials modeling

- **Holy grail**: classical potentials with quantum accuracy
  - $O(N)$ cost for $O(N^3)$ accuracy
- **Idea**: train NN potentials on quantum data, i.e. Ax = b

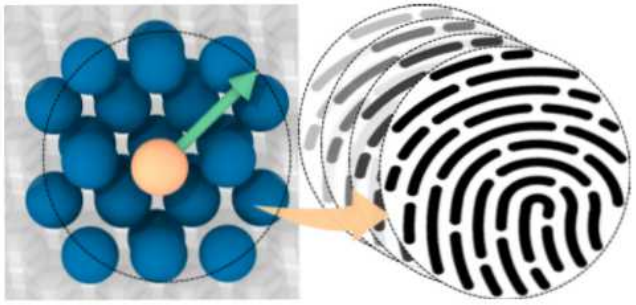

Input: geometric descriptors of an atom's neighborhood
Output: energy/forces on atom
Truth: quantum DFT forces

- Issues: Configs to train on, quantify errors, transferable ?
- Lots of **excitement**, but jury still out on how useful it will be

Machine learning for materials modeling

- **Holy grail**: classical potentials with quantum accuracy
  - $O(N)$ cost for $O(N^3)$ accuracy
- **Idea**: train NN potentials on quantum data, i.e. Ax = b

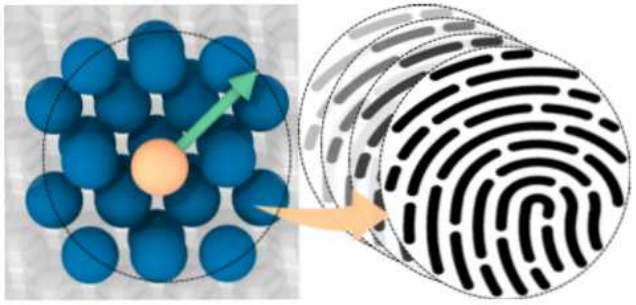

Input: geometric descriptors of an atom's neighborhood
Output: energy/forces on atom
Truth: quantum DFT forces

- Issues: Configs to train on, quantify errors, transferable ?
- Lots of **excitement**, but jury still out on how useful it will be
- **ML challenge**: what can it really do for scientific computing and modeling/simulation

Who will be the Future ghost of parallel computing?







Who will be the Future ghost of parallel computing?








Who will be the Future ghost of parallel computing?







- Future has not yet been written
- Could be **one of you** !

Who will be the Future ghost of parallel computing?







- Future has not yet been written
- Could be **one of you** !

- What will have **greatest impact** on comp science of future?
- Is ML a **distraction** from the hard algorithmic & hard science problems mod/sim faces?
- Or is ML a **silver bullet** for some of those problems?

Summary and thanks

For scientific computing, ever-faster computers have made the last 30 years **easy**. The next 30 will be different, harder, and **more interesting**. Good luck, next generation!

Summary and thanks

For scientific computing, ever-faster computers have made the last 30 years **easy**. The next 30 will be different, harder, and **more interesting**. Good luck, next generation!

Thanks:

Photoshop wizard: Brad Carvey
Hardware gurus: Si Hammond, Ron Brightwell
App experts: Doug Kothe (PIC), Mark Taylor (climate),
Francois Gygi (DFT), Guilia Galli (DFT), Paul Fischer (CFD)

Summary and thanks

For scientific computing, ever-faster computers have made the last 30 years **easy**. The next 30 will be different, harder, and **more interesting**. Good luck, next generation!

Thanks:

Photoshop wizard: Brad Carvey
Hardware gurus: Si Hammond, Ron Brightwell
App experts: Doug Kothe (PIC), Mark Taylor (climate),
Francois Gygi (DFT), Guilia Galli (DFT), Paul Fischer (CFD)

Disclaimer: no ghosts were harmed in the making of these slides