What's New in Supercomputing - 2022 Edition

Calvin W. Johnson, San Diego State University

2022 Long Range Plan Town Hall @ Argonne, Nov 15, 2022

What's New in Supercomputing - 2022 Edition

Before we talk about where we want to go,

We need to remember where we have been.

My first personal computer – a Commodore VIC 20
(early 1980's)

now seen in museums!

2022 Long Range Plan Town Hall @ Argonne, Nov 15, 2022

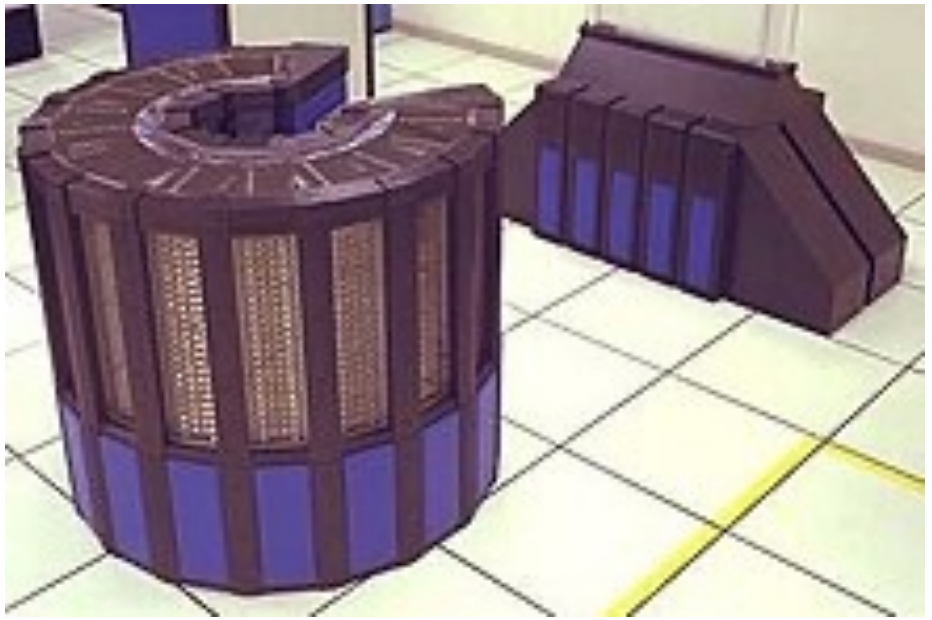

Peak speed:
1.9 Gflops!

My first **super**computer – a Cray 2
(1985)

now seen in museums!

Jump ahead 30 years…

Top supercomputers **2015** (***last long-range plan***)

#1: Tianhe-2A (China):	3.1 Mcores	33.9 Pflops max	(17.8 MW)
#2: Titan (ORNL) :	0.6 Mcores	17.6 Pflops	(8.2 MW)
#3: Sequoia (LLNL) :	1.5 Mcores	17.1 Pflops	
#4: K (RIKEN, Japan):	0.7 Mcores	10.5 Pflops	
#5: Mira (ANL) :	0.8 Mcores	8.6 Pflops	
#7: Stampede (TACC):	0.5 Mcores	5.2 Pflops	
#10: Vulcan (LLNL) :	0.4 Mcores	4.3 Pflops	
…			
#25: Edison (NERSC):	0.13 Mcores	1.6 Pflops	

x 1-10 million speed over Cray-2!

Top supercomputers 2022 (*today*)

= 1.1 Eflops!

#1: Frontier (ORNL)	8.7 Mcores	1,100 Pflops max	(21 MW)
#2: Fugaku (RIKEN)	7.6 Mcores	440 Pflops	(30 MW)
#4: Summit (ORNL)	2.4 Mcores	150 Pflops	
#5: Sierra (LLNL)	1.6 Mcores	94 Pflops	
#7: Perlmutter (NERSC)	0.8 Mcores	70 Pflops	

- approx x 3 cores
- x 10-30 flops
- < x2 in power consumption (cf. "Green500" list)

Let's look in detail

2015 Titan : 18,688 nodes (16-core CPU + 2688 core GPU)
Invidia GPUs: core clock 0.7 GHz, **6 GB memory**, memory clock 5.2 GT/s, **bus bandwidth 250 Gb/s**
cost: ~$100M

2022 Frontier: 9,742 nodes (64-core CPU+ 4 x 220 core GPUs)
Radeon GPUs: core clock 1 GHz, **128 GB memory**, memory clock 3.2 GT/s, **bus bandwidth 2x 1638 Gb/s**

cost: $600M (compare FRIB: $760M)

2022 Summit: 4,608 nodes (2 x 22-core CPUs + 6 x 5120 GPUs)
Invidia GPUs: core clock 1? GHz, **16GB memory**, memory clock 1.6 GT/s, **bus bandwidth 900 Gb/s**

Over the next 5 years we can expect….

A few more exaflop machines

A lot of sub-exaflop (0.1 Eflop) machines

My takeaway

CPUs becoming more like GPUs (or *manycore*)

GPUs becoming more like CPUs (tradeoff cores for more bandwidth, more memory)

Mid-size resources

Not everyone needs a million cores. Many problems addressable by 1,000-10,000 cores. (While many universities have computing facilities, these are often oversubscribed.)

NSF-sponsored consortium ACCESS "Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support" www.access-ci.org
formerly XSEDE "eXtreme Science & Engineering Discovery Environment"

Generally open to US-based researchers & students

Mid-size resources

Examples:

Stampede2 @ Texas Advanced Computing Center: estimated 18PF peak performance
4204 KNL nodes x 68 cores CPU, and 1736 Skylake (SKX) nodes x 48 cores CPU

Expanse @ San Diego Supercomputer Center:
728 AMD EPYC nodes x 64 cores CPU, + 52 nodes 2 x 20 core CPU + 4 x 5102 core CPU

Anvil @ Purdue University
1000 nodes x 128 cores CPU, + 16 nodes 128 cores CPU+ 4 x NVIDIA Tensor GPUs

Computational Science

It’s not enough to just have the machines. We need to use them well.
In particular in effective use of advanced architectures such as GPUs.

Need to couple physicists with computer science & applied math.
We have had great successes: SciDAC and INCITE projects (e.g., UNEDF), uncertainty quantification/. Bayesian techniques/machine learning.
Major facilities such as NERSC also offer aid.

Nonetheless we should not rest on these laurels but encourage further efforts, especially percolating down to smaller groups.

Computational Science @ SDSU

SDSU has a Ph.D in computational science (joint with UC Irvine)
(currently carrying out a search for new faculty member—I'm chairing the search!)

Students take courses in computer science, applied math, statistics…
but carry out research in domain science such as physics, chemistry, biology, etc.

Hence expert in computing, but can also 'talk to' domain scientists fluently.

Very attractive both to national labs and to students as a career path.

GPUs are not great for every problem

For example, *sparse matrix multiplication* (typical in, say, no-core shell model), does not translate easily to GPUs.

Lanczos algorithm for eigenpairs depends upon *matrix-vector multiplication*

(and enforced (re)orthogonalization of vectors because of round-off error)

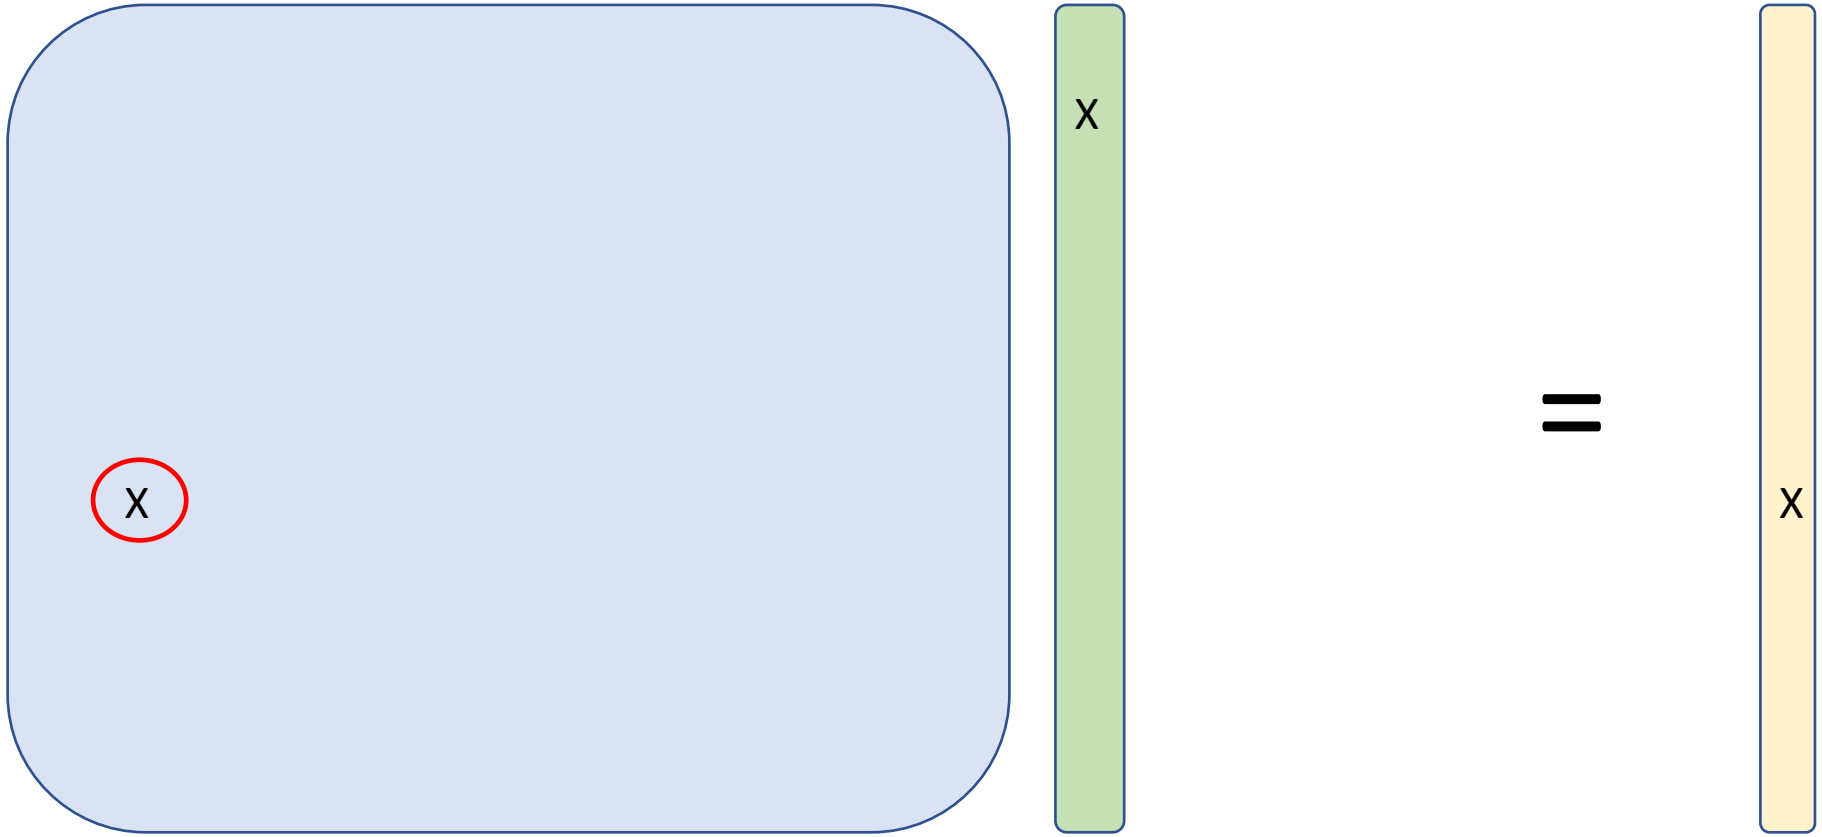

Sparse matrix-vector multiplication violates data locality

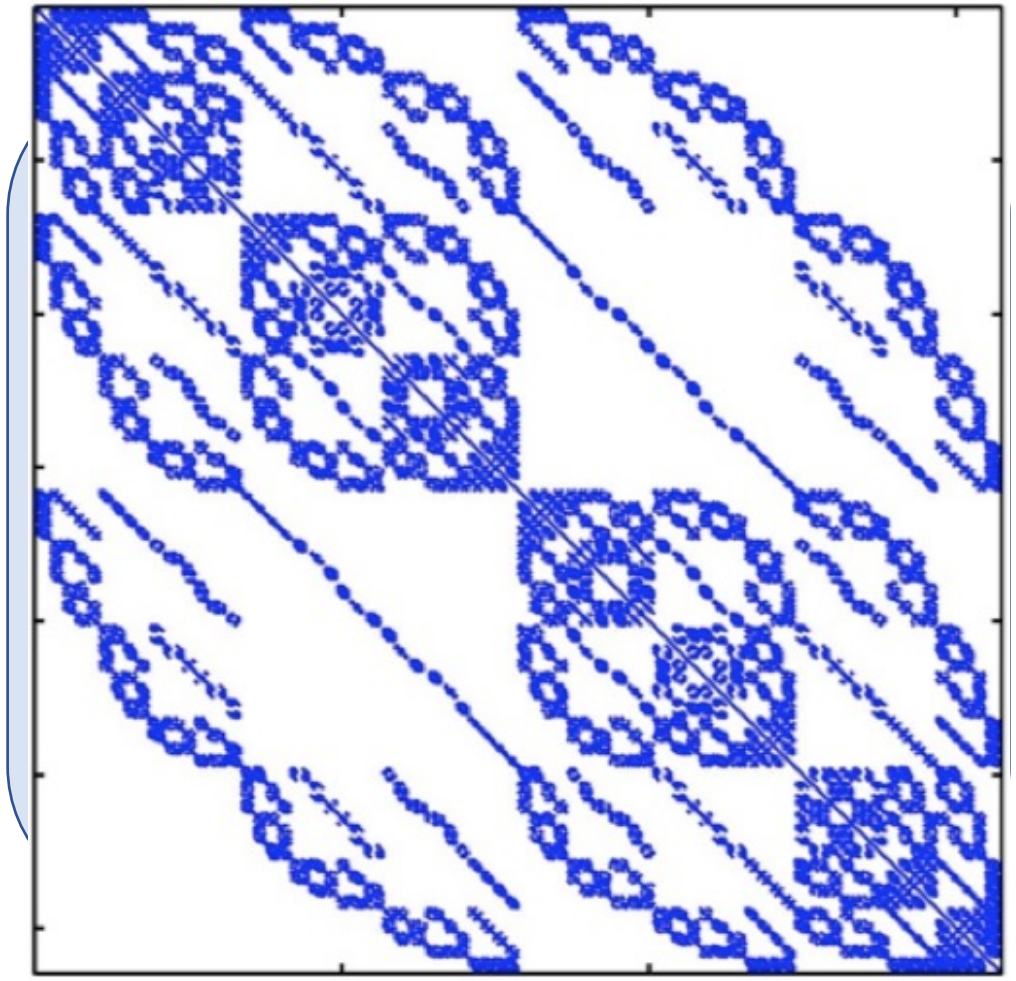

x

=

x

Loss of data locality makes it hard to balance work (some are more local than others)

GPUs are not great for every problem

Sparse matrix multiplication does not translate easily to GPUs.

(There are many papers, but the executive summary is: *not easy*.)

This goes back to the issues of bandwidth (data transfer) and storage.

Block Lanczos (with matrix-matrix multiplication) can improve things

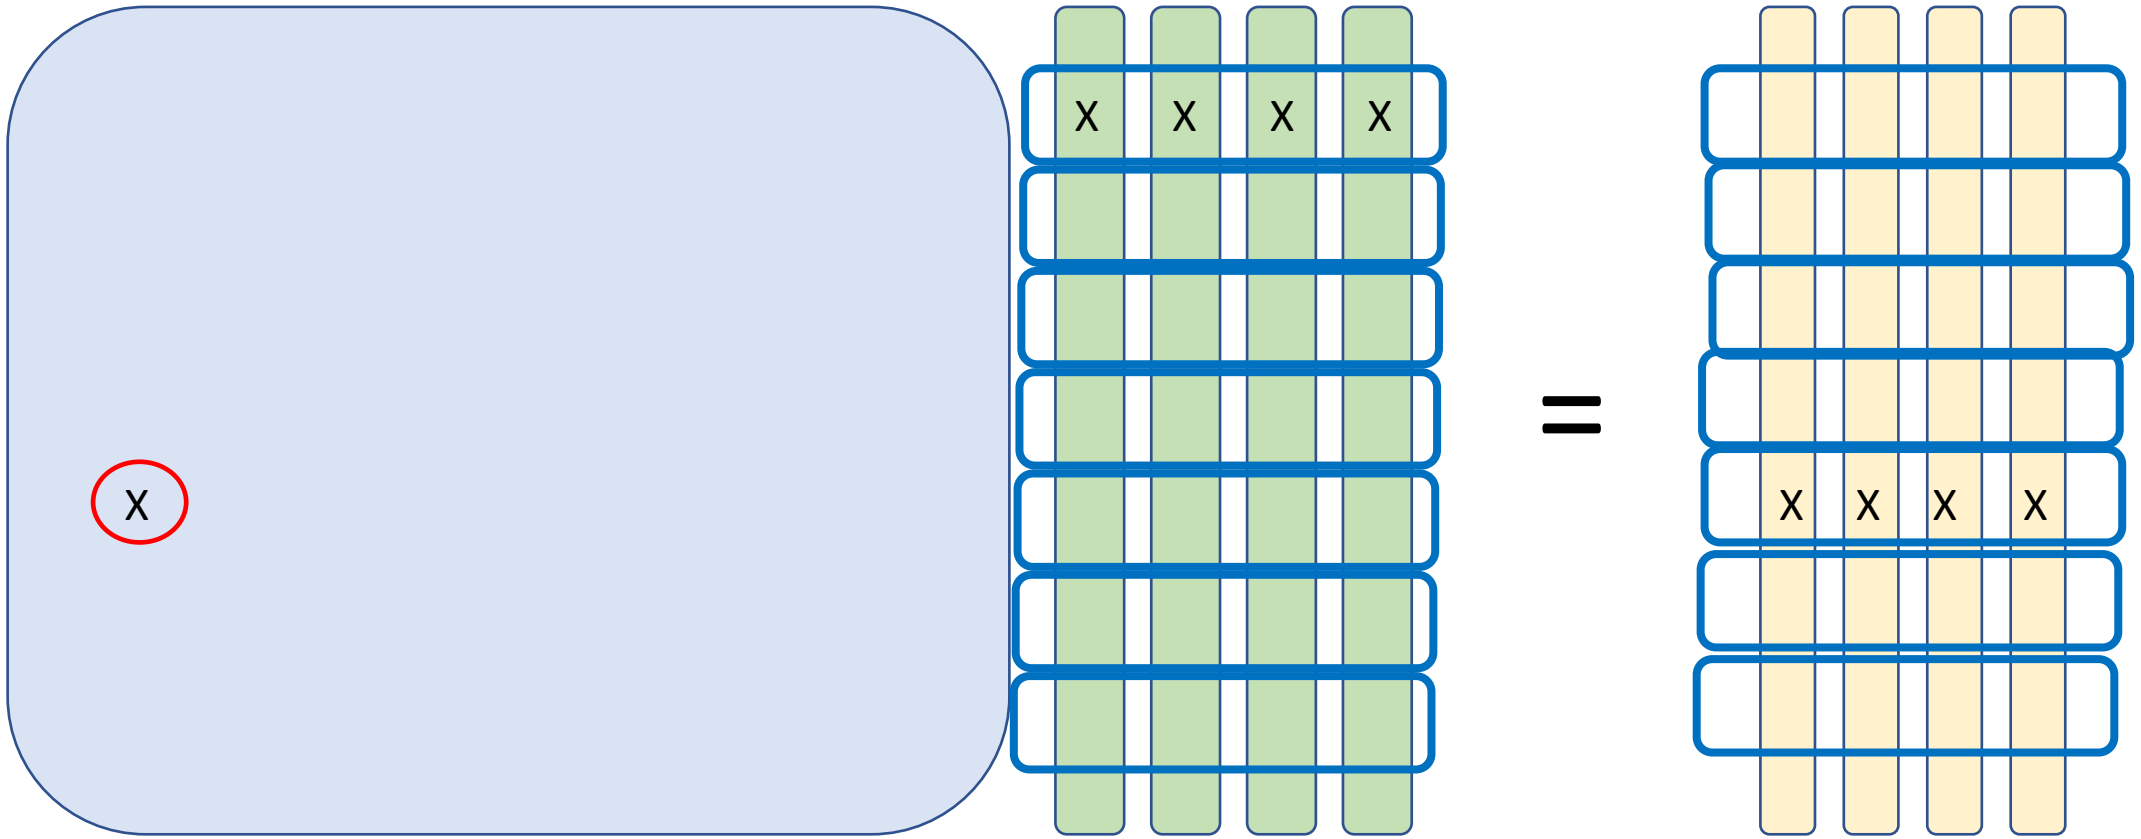

GPUs are not great for every problem

For example, *sparse matrix multiplication* (typical in, say, no-core shell model), does not translate easily to GPUs.

Tweaks in algorithms can improve things. On CPUs, at least, going from vector Lanczos to block Lanczos speeds up by x 2 (improvement in data locality/ reduction of cache misses). How to scale to GPUs not clear, though

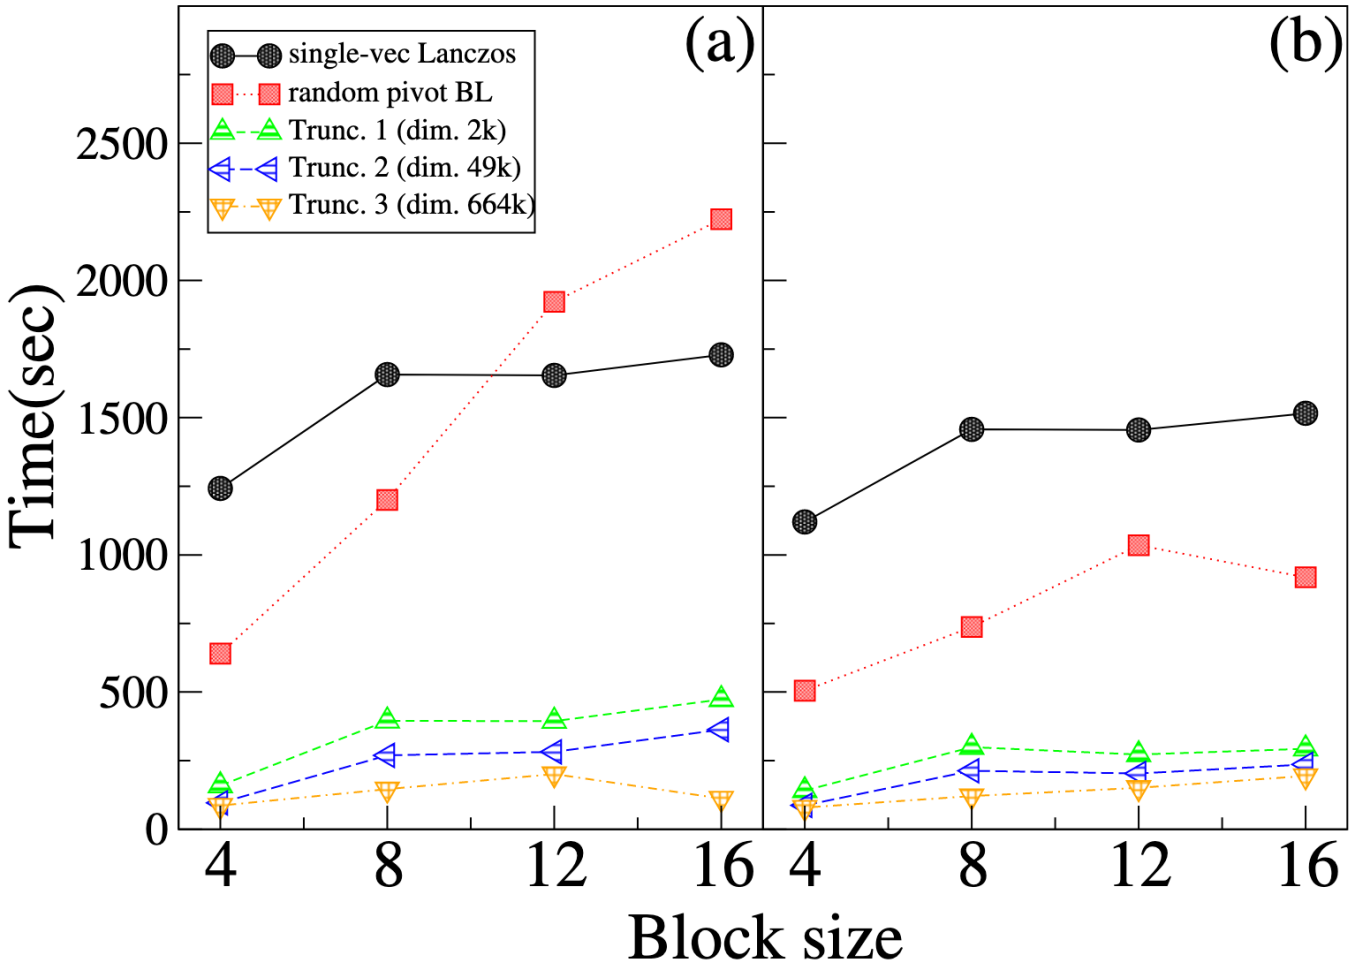

Zbikowski & Johnson,
arXiv: 2210.15763

Block Lanczos speeds up matrix application + 'bootstrapping' from truncated solution reduces # iterations

Time-to-solution for ^{7}Be NCSM using bootstrapped block Lanczos (including reorthogonalization)

Total matrix-multiplication time for ^{7}Be NCSM using bootstrapped block Lanczos

HPC summary

We are at the dawn of the ‘exascale’ age.

Convergence of *manycore* CPUs and GPUs (with more memory+bandwidth) seems to lead to better performance *and* will make migrating codes easier.

Resources exist for aid in migration to GPUs—are they enough? (Discuss)

We still need a healthy ‘ecosystem’ of mid- to extreme-scale computers, equally accessible to researchers.

We should also **highlight** and **advocate** for continuing and increased coupling of physicists with computer science/applied math

Comment on ‘data management’ and ‘code management’

As a community, we should articulate ‘**best practices**’ for data and code management.

Just slapping it on GitHub should not be enough.

Data and code should be extensively curated and documented.

e.g. The `BIGSTICK` shell-model manual is over 120 pages!

Comments on Machine Learning and Quantum Computing

At Physical Review C (and elsewhere) we are seeing an explosion of papers on machine learning/artificial intelligence and on quantum computing.

(*Need referees—please volunteer!*
write to: cjohnson at aps.org)

The community, including but not only journals, should agree on '**best practices**' for these topics

Comments on Machine Learning and Quantum Computing

Nonetheless, there are good ML/AI papers and bad ML/AI papers

Good papers:
-- Must be reproducible (duh). Not only training data, but network topology/hyperparameters must be included.
-- Understand how many parameters are used
-- How is overfitting handled? How well does it extrapolate/predict data not in training?
-- *Physics-informed* ML seems to lead to improved performance and leads to understanding.

Bad papers:
-- Use off-the-shelf software without understanding the guts (i.e., # of parameters)
-- Simply calculate the rms error and declare victory!
-- Deliberately disdain any need for understanding

Comments on Machine Learning and Quantum Computing

Quantum computing is not longer in its infancy, but still a toddler.

2022 Long Range Plan Town Hall @ Argonne, Nov 15, 2022

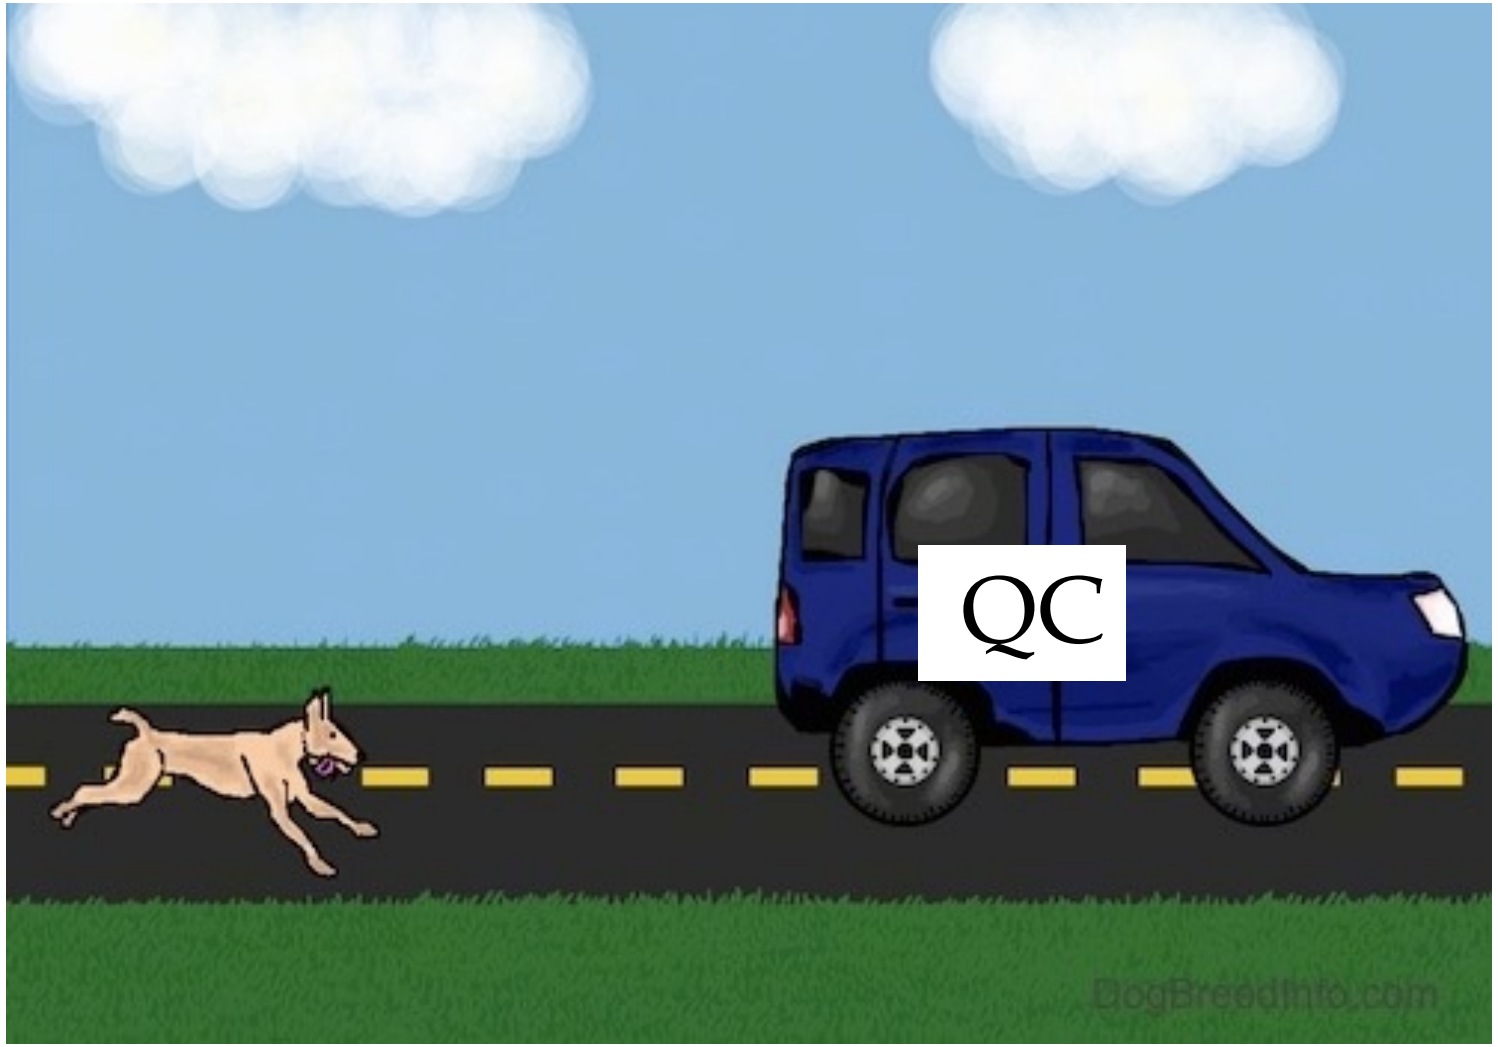
QC
DogBreedInfo.com

2022 Long Range Plan Town Hall @ Argonne, Nov 15, 2022

Let's look at the data requirements
in more detail

Consider ^{12}C, $N_{max}=8$

M-scheme dimension 0.6 billion

55 single-particle orbits (n l j)
440 single particle states (n l j m) | 0 1 1 0 0 1 … >

Let's look at the data requirements
in more detail

Consider ^{12}C, $N_{max}=8$

M-scheme dimension 0.6 billion

55 single-particle orbits (n l j)
440 single particle states (n l j m) | 0 1 1 0 0 1 … >

= minimum # of qubits needed (probably more)

Let's look at the data requirements
in more detail

Consider ^{12}C, $N_{max}=8$

M-scheme dimension 0.6 billion *by superposition*

~~*# J-coupled 2-body matrix elements*: ~ 1.5 million~~
~~< a b J | H | c d J >~~ *input*

***# uncoupled 2-body matrix elements* ~ 10 million!**
$\mathbf{V_{ijkl}\ a^{+}_{i}\ a^{+}_{j}\ a_{l}\ a_{k}}$ = # 'Pauli strings'
= # of terms to be evaluated in a quantum circuit
(or, # of separate quantum circuits to be evaluated!)

~~*# many-body matrix elements:* ~ 1.2 trillion~~
~~(or 5 Tb storage)~~ *not relevant?*

We're still a long ways from catching the car we want!

Noise and noise mitigation remains a huge issue, even in small cases

To tackle problems our community cares about, we will need on the order of $> 10^{5-7}$ logical qubits (with error correction, 10^{6-9} physical qubits)

Nonetheless, science has stared down seemingly insurmountable challenges before

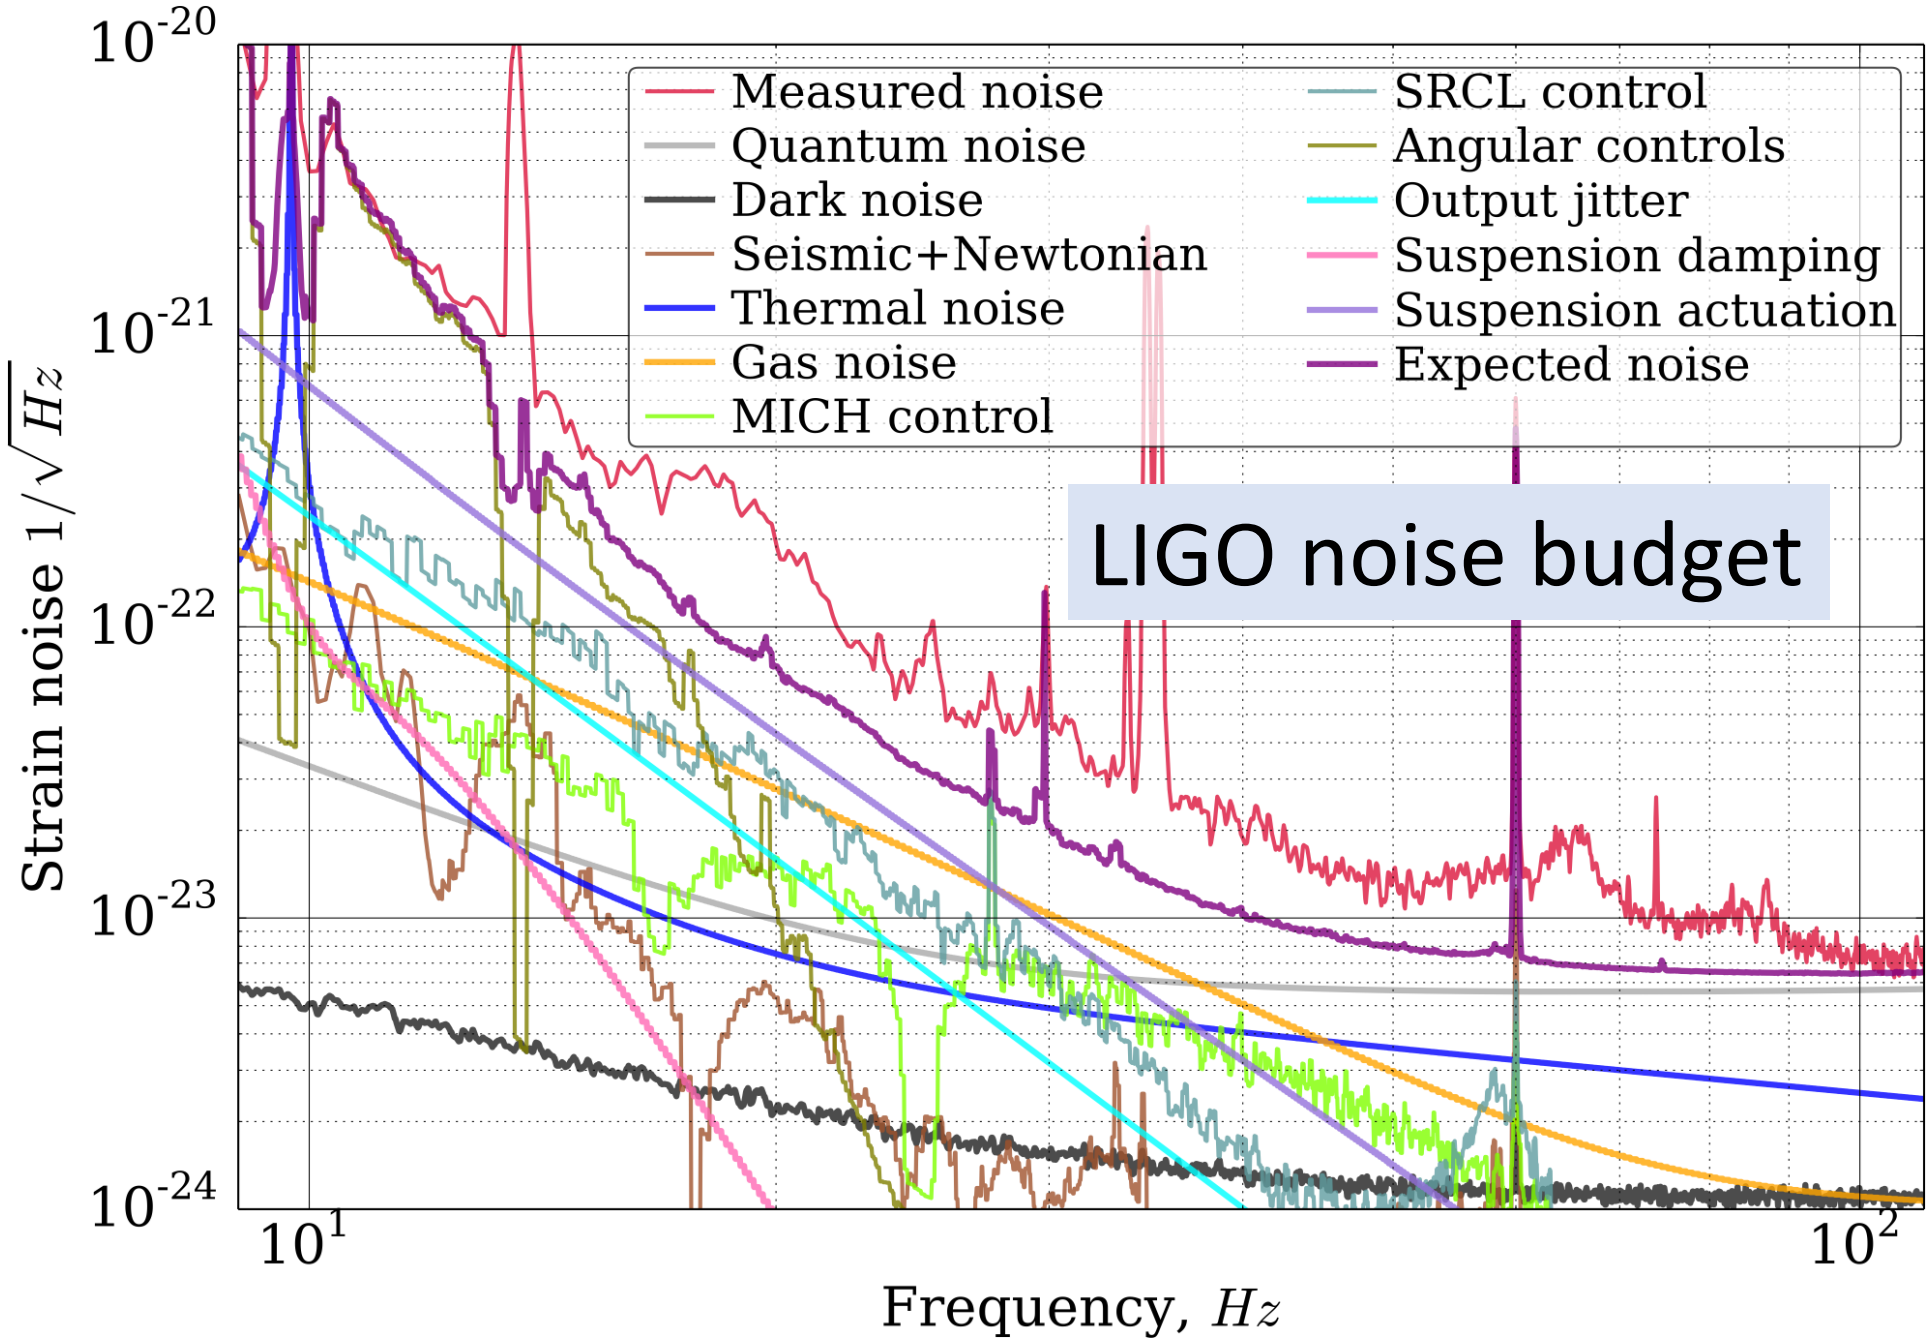
Measured noise
Quantum noise
Dark noise
Seismic+Newtonian
Thermal noise
Gas noise
MICH control
SRCL control
Angular controls
Output jitter
Suspension damping
Suspension actuation
Expected noise
LIGO noise budget
Strain noise $1/\sqrt{Hz}$
10^{-20}
10^{-21}
10^{-22}
10^{-23}
10^{-24}
10^{1}
10^{2}
Frequency, Hz

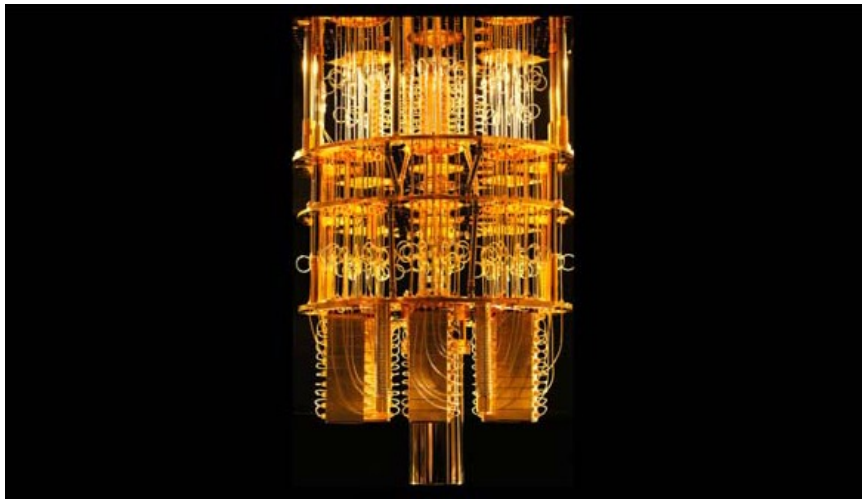

=?

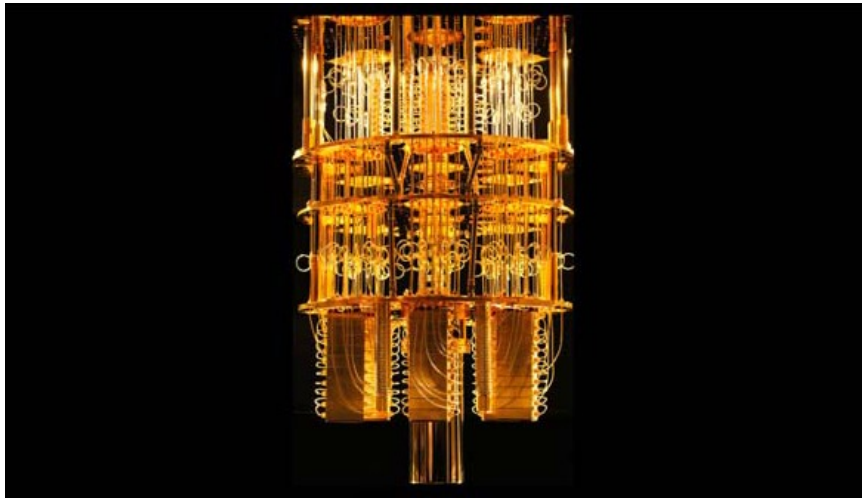

=?

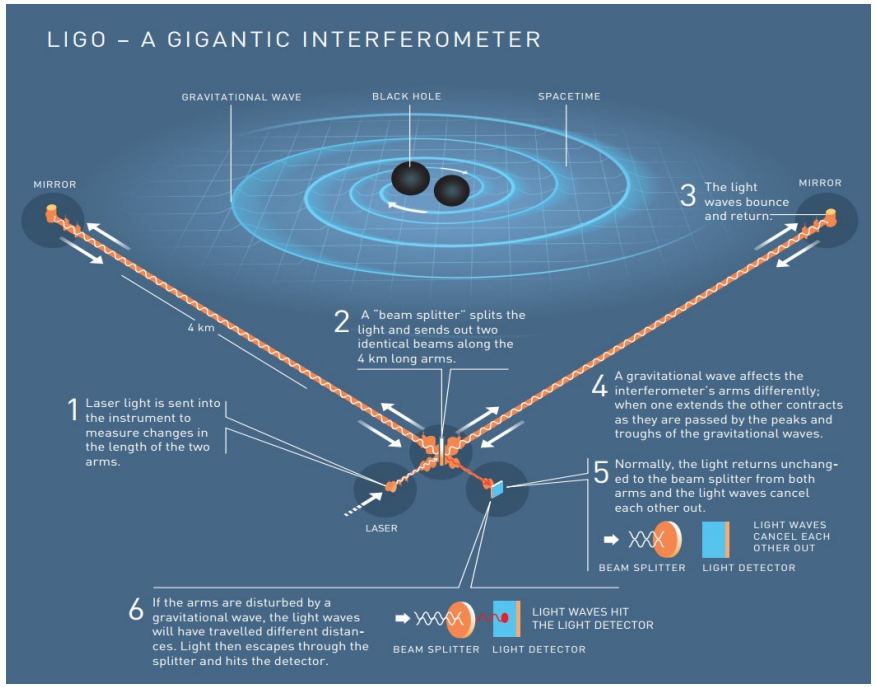
LIGO – A GIGANTIC INTERFEROMETER
GRAVITATIONAL WAVE
BLACK HOLE
SPACETIME
MIRROR
MIRROR
3 The light waves bounce and return.
4 km
2 A "beam splitter" splits the light and sends out two identical beams along the 4 km long arms.
1 Laser light is sent into the instrument to measure changes in the length of the two arms.
4 A gravitational wave affects the interferometer's arms differently; when one extends the other contracts as they are passed by the peaks and troughs of the gravitational waves.
5 Normally, the light returns unchanged to the beam splitter from both arms and the light waves cancel each other out.
LASER
LIGHT WAVES CANCEL EACH OTHER OUT
BEAM SPLITTER
LIGHT DETECTOR
6 If the arms are disturbed by a gravitational wave, the light waves will have travelled different distances. Light then escapes through the splitter and hits the detector.
LIGHT WAVES HIT THE LIGHT DETECTOR
BEAM SPLITTER
LIGHT DETECTOR

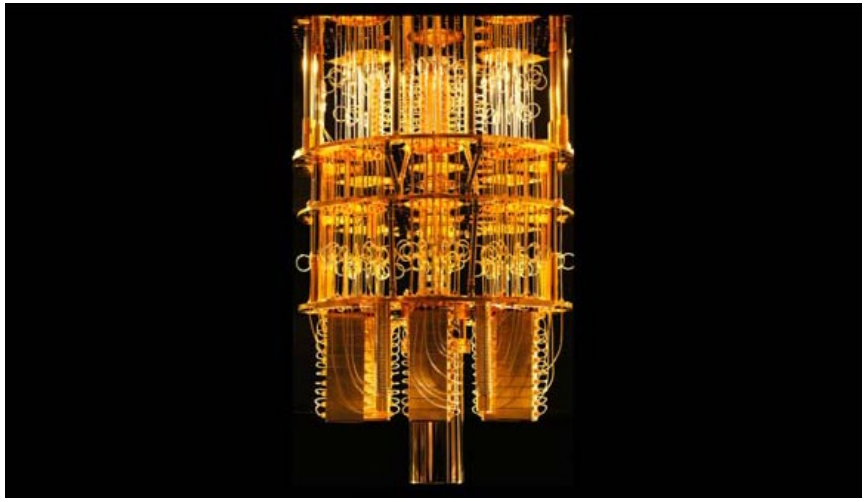

=?

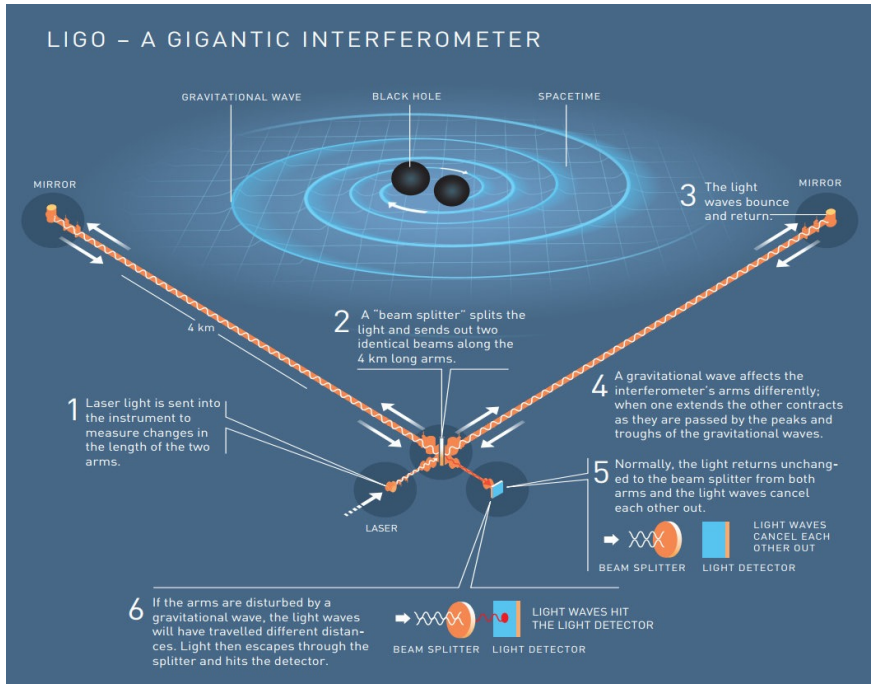
LIGO – A GIGANTIC INTERFEROMETER
GRAVITATIONAL WAVE
BLACK HOLE
SPACETIME
MIRROR
MIRROR
3 The light waves bounce and return.
4 km
2 A "beam splitter" splits the light and sends out two identical beams along the 4 km long arms.
1 Laser light is sent into the instrument to measure changes in the length of the two arms.
4 A gravitational wave affects the interferometer's arms differently; when one extends the other contracts as they are passed by the peaks and troughs of the gravitational waves.
5 Normally, the light returns unchanged to the beam splitter from both arms and the light waves cancel each other out.
LASER
BEAM SPLITTER
LIGHT DETECTOR
LIGHT WAVES CANCEL EACH OTHER OUT
6 If the arms are disturbed by a gravitational wave, the light waves will have travelled different distances. Light then escapes through the splitter and hits the detector.
BEAM SPLITTER
LIGHT DETECTOR
LIGHT WAVES HIT THE LIGHT DETECTOR

or

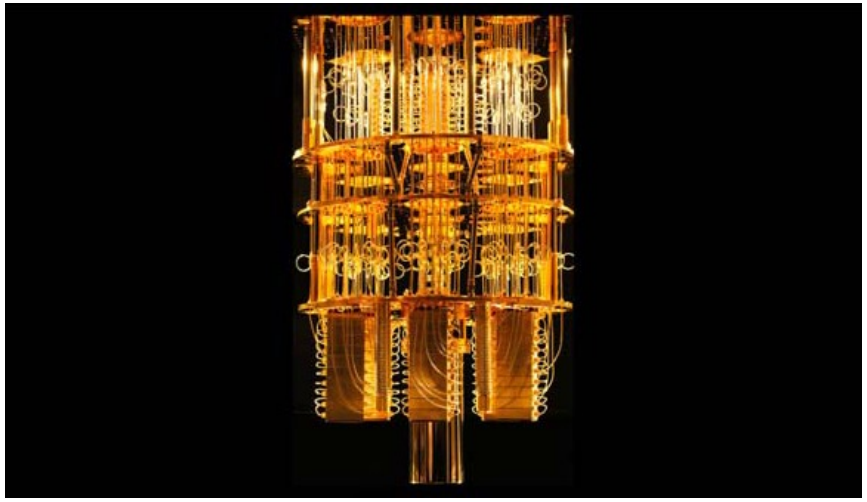

=?

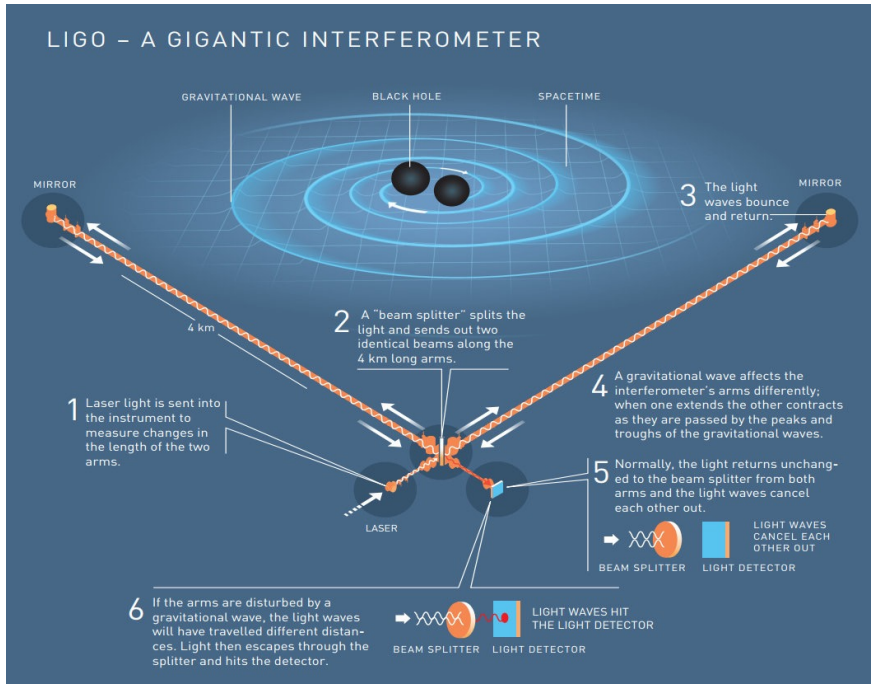
LIGO – A GIGANTIC INTERFEROMETER
GRAVITATIONAL WAVE
BLACK HOLE
SPACETIME
MIRROR
MIRROR
3 The light waves bounce and return.
4 km
2 A "beam splitter" splits the light and sends out two identical beams along the 4 km long arms.
1 Laser light is sent into the instrument to measure changes in the length of the two arms.
4 A gravitational wave affects the interferometer's arms differently; when one extends the other contracts as they are passed by the peaks and troughs of the gravitational waves.
5 Normally, the light returns unchanged to the beam splitter from both arms and the light waves cancel each other out.
LASER
LIGHT WAVES CANCEL EACH OTHER OUT
BEAM SPLITTER
LIGHT DETECTOR
6 If the arms are disturbed by a gravitational wave, the light waves will have travelled different distances. Light then escapes through the splitter and hits the detector.
LIGHT WAVES HIT THE LIGHT DETECTOR
BEAM SPLITTER
LIGHT DETECTOR

or

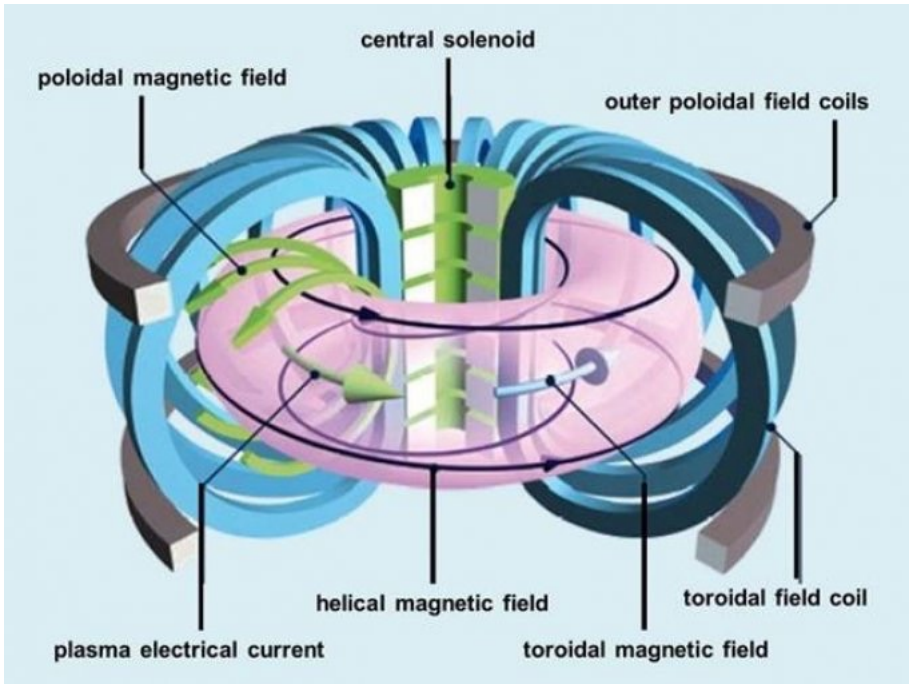
central solenoid
poloidal magnetic field
outer poloidal field coils
helical magnetic field
toroidal field coil
plasma electrical current
toroidal magnetic field