Towards a quantum LINPACK benchmark

Lin Lin

Department of Mathematics, UC Berkeley
Lawrence Berkeley National Laboratory

Joint work with Yulong Dong (Berkeley)

Google Summer Symposium 2020,

arXiv: 2006.04010

https://github.com/qsppack/racbem

How many quantum computers will we need?

Finally we have a few quantum computers

Google AI Quantum

IBM

rigetti

Berkeley
UNIVERSITY OF CALIFORNIA

...

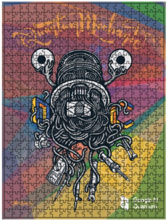

I think there is a world market for maybe five computers.
–Thomas Watson, president of IBM, 1943

In 2019, there are around 2 billion computers in the world (estimate).

Prediction is very difficult, especially if it's about the future.
–Niels Bohr (also attributed to others)

Let us make some prediction: there will be say, 10000 quantum computers

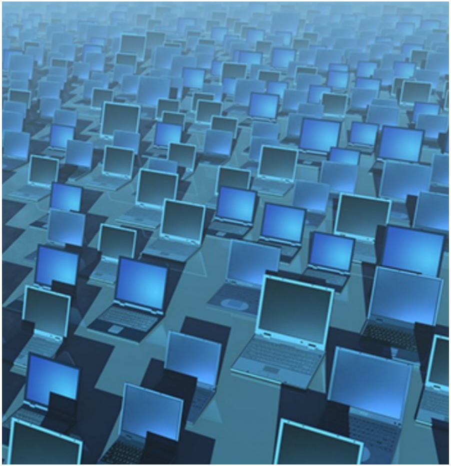

How to select the TOP500 quantum computers?

First, how to do the job for classical supercomputers?

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	RIKEN Center for Computational Science Japan	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D Fujitsu	7,299,072	415,530.0	513,854.7	28,335
2	DOE/SC/Oak Ridge National Laboratory United States	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,414,592	148,600.0	200,794.9	10,096
3	DOE/NNSA/LLNL United States	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM / NVIDIA / Mellanox	1,572,480	94,640.0	125,712.0	7,438
4	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
5	National Super Computer Center in Guangzhou China	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482

NEWS

Japan Captures TOP500 Crown with Arm-Powered Supercomputer

June 22, 2020

FRANKFURT, Germany; BERKELEY, Calif.; and KNOXVILLE, Tenn.—The 55th edition of the TOP500 saw some significant additions to the list, spearheaded by a new number one system from Japan. The latest rankings also reflect a steady growth in aggregate performance and power efficiency.

The new top system, Fugaku, turned in a High Performance Linpack (HPL) result of 415.5 petaflops, besting the now second-place Summit system by a factor of 2.8x. Fugaku, is powered by Fujitsu's 48-core A64FX SoC, becoming the first number one system on the list to be powered by ARM processors. In single or further reduced precision, which are often used in machine learning and AI applications, Fugaku's peak performance is over 1,000 petaflops (1 exaflops). The new system is installed at RIKEN Center for Computational Science (R-CCS) in Kobe, Japan.

What is LINPACK? Why LINPACK?

[1] https://www.top500.org/, 55th edition of the TOP500, June 2020

LINPACK benchmark

- Interested in using supercomputers for scientific computing (instead of e.g. recognizing cats, but maybe now it is difficult to distinguish the two..)
- Solving linear systems: building block for numerous scientific computing applications
- LINPACK: Solve $Ax = b$ with dense, random matrices. No obvious applications. Supremacy?
- Controversial over its effectiveness since early days. Alternative benchmarks have been proposed[1]. Still solving linear systems with some random sparsity patterns.
- Has been used to define TOP500 since its debut in 1993.
- Quantum LINPACK benchmark?

[1] Why Linpack no longer works as well as it once did. (link)

Climbing the Quantum Mount Everest

Quantum advantage

Quantum LINPACK benchmark

STEPS away ↕

Quantum supremacy

Quantum linear system problem (QLSP)

- Use quantum computers to solve

$$|x\rangle \propto A^{-1} |b\rangle .$$

- How to get the information in $A, |b\rangle$ into a quantum computer? read-in problem.
- κ: condition number of $A = \|A\| \left\|A^{-1}\right\|$; ϵ: target accuracy.
- Proper assumptions on A (e.g. d-sparse) so that oracles cost $\mathrm{poly}(n)$, while $A \in \mathbb{C}^{2^n \times 2^n}$. (Potentially) exponential speedup.

Compare the complexities of QLSP solvers

Significant progress in the past few years: Near-optimal complexity matching lower bounds.

Algorithm	Query complexity	Remark
HHL,(Harrow-Hassidim-Lloyd, 2009)	$\widetilde{\mathcal{O}}(\kappa^2/\epsilon)$	w. VTAA, complexity becomes $\widetilde{\mathcal{O}}(\kappa/\epsilon^3)$ (Ambainis 2010)
Linear combination of unitaries (LCU),(Childs-Kothari-Somma, 2017)	$\widetilde{\mathcal{O}}(\kappa^2 \mathrm{polylog}(1/\epsilon))$	w. VTAA, complexity becomes $\widetilde{\mathcal{O}}(\kappa\, \mathrm{poly}\log(1/\epsilon))$
Quantum singular value transformation (QSVT) (Gilyén-Su-Low-Wiebe, 2019)	$\widetilde{\mathcal{O}}(\kappa^2 \log(1/\epsilon))$	Queries the RHS only $\widetilde{\mathcal{O}}(\kappa)$ times
Randomization method (RM) (Subasi-Somma-Orsucci, 2019)	$\widetilde{\mathcal{O}}(\kappa/\epsilon)$	Prepares a mixed state; w. repeated phase estimation, complexity becomes $\widetilde{\mathcal{O}}(\kappa\, \mathrm{poly}\log(1/\epsilon))$
Time-optimal adiabatic quantum computing (AQC(exp)) (An-**L.**, 2019, 1909.05500)	$\widetilde{\mathcal{O}}(\kappa\, \mathrm{poly}\log(1/\epsilon))$	No need for any amplitude amplification. Use time-dependent Hamiltonian simulation.
Eigenstate filtering (**L.**-Tong, 2019, 1910.14596)	$\widetilde{\mathcal{O}}(\kappa \log(1/\epsilon))$	No need for any amplitude amplification. Does not rely on any complex subroutines.

Quantum benchmark problem

- All these algorithms require full fault-tolerant computers to get anywhere.
- Getting the matrix into quantum computer alone (using e.g. LCU) can be prohibitively expensive on near-term devices.
- Will likely remain so for some time for real applications.
- Quantum LINPACK benchmark is different: do we really need / want to generate random numbers classically and get them into quantum computers say via QRAM?

Idea: from the success of Google's supremacy circuit

- A big random, unitary matrix, almost drawn from Haar measure.
- Linear algebra usually works with non-unitary matrices.
- How about taking the upper-left diagonal block n-qubit matrix of the $(n+1)$-qubit unitary matrix (can be random)?

$$U_A = \begin{pmatrix} A & \cdot \\ \cdot & \cdot \end{pmatrix}$$

- FACT: It can represent in principle any n-qubit matrix (up to scaling)!
- Block-encoding (Gilyén-Su-Low-Wiebe, 2019)
- RAndom Circuit Block-Encoded Matrix (RACBEM)

RACBEM

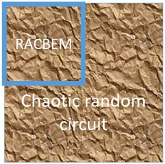

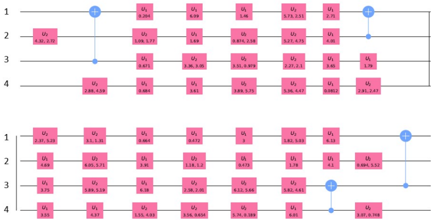

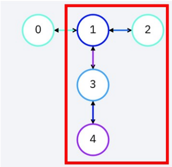

Qubit 1 takes $|0\rangle$
Get a 3-qubit RACBEM

$$A = \begin{pmatrix} -0.066-0.055i & -0.301-0.093i & -0.133-0.118i & 0.046+0.016i & -0.024-0.027i & -0.119-0.056i & 0.393+0.403i & -0.141-0.061i \\ -0.301-0.099i & -0.067-0.041i & 0.042+0.014i & -0.152-0.095i & -0.119-0.058i & -0.025-0.021i & -0.130-0.054i & 0.459+0.333i \\ -0.079-0.188i & -0.105+0.155i & -0.178+0.097i & -0.195-0.104i & 0.396-0.069i & -0.253-0.270i & 0.028+0.124i & -0.096+0.099i \\ -0.109+0.149i & -0.107-0.176i & -0.197-0.104i & -0.154+0.128i & -0.239-0.276i & 0.386-0.128i & -0.097+0.101i & 0.050+0.115i \\ -0.135+0.033i & -0.412+0.299i & 0.037+0.272i & -0.038-0.065i & -0.003+0.045i & 0.053+0.156i & -0.278+0.162i & 0.053-0.070i \\ -0.419+0.294i & -0.118+0.047i & -0.035-0.059i & 0.083+0.263i & 0.051+0.158i & 0.003+0.041i & 0.048-0.065i & -0.248+0.208i \\ -0.071-0.125i & -0.060+0.118i & 0.029+0.367i & -0.316+0.245i & -0.150+0.247i & 0.265-0.009i & 0.038+0.066i & -0.038+0.074i \\ -0.064+0.114i & -0.090-0.114i & -0.319+0.249i & 0.097+0.350i & 0.262+0.001i & -0.113+0.269i & -0.038+0.075i & 0.049+0.057i \end{pmatrix}$$

- A very flexible way to construct a non-unitary matrix with respect to any coupling map of the quantum architecture.
- Take upper-left diagonal block: measure one-qubit.
  $A = (\langle 0| \otimes I_n)\, U_A\, (|0\rangle \otimes I_n)$

Error of RACBEM on IBM Q

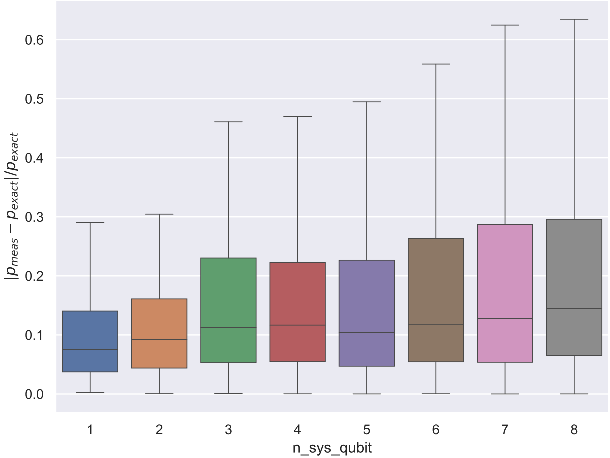

Relative error in success probability of obtaining $A|0^n\rangle$ for different number of system qubits, on the 5-qubit backend ibmq_burlington, and 15-qubit backend ibmq_16_melbourne.

Representing matrix functions, say $f(A) = A^{-1}$

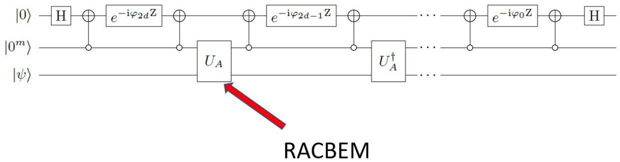

- General quantum singular value transformation circuit (QSVT) (Gilyén-Su-Low-Wiebe, 2019). Always even order polynomial, and 1 extra ancilla qubit.

- Follow the natural layout of the quantum circuit. Can run without even calling a transpiler.

- Applications: Linear systems, time series, spectral measure, thermal energy ... with Hermitian-RACBEM.

- How to obtain the phase factors: optimization based approach to get > 10000 phase factors [1]

[1] (Dong-Meng-Whaley-**L.**, 2002.11649). https://github.com/qsppack/qsppack

Solving linear system on IBM Q and QVM

Compute $\left\|\mathfrak{H}^{-1}\left|0^n\right\rangle\right\|_2^2$. (sigma: noise level on QVM)

Well conditioned linear system

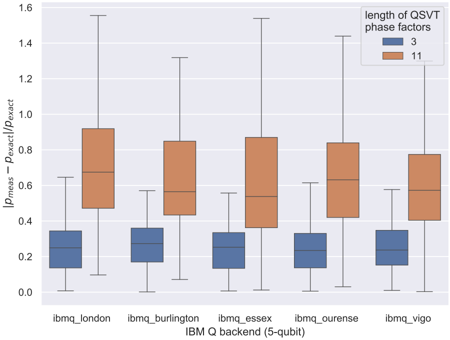

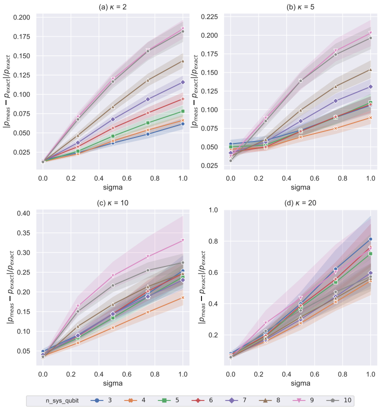

Time series (no Trotter)

$$s(t) = \langle \psi | e^{i\mathfrak{H}t} | \psi \rangle$$

QVM:

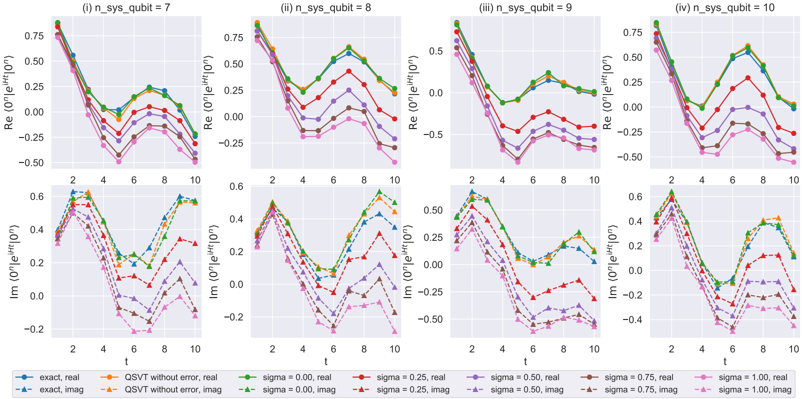

Spectral measure

$$s(E) = \langle \psi | \delta(\mathfrak{H} - E) | \psi \rangle \approx \frac{1}{\pi} \operatorname{Im} \langle \psi | (\mathfrak{H} - E - \mathrm{i}\eta)^{-1} | \psi \rangle$$

QVM:

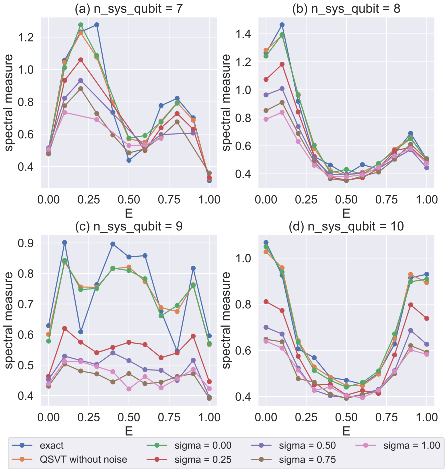

Thermal energy

$$E(\beta) = \frac{\mathrm{Tr}[\mathfrak{H} e^{-\beta \mathfrak{H}}]}{\mathrm{Tr}[e^{-\beta \mathfrak{H}}]}$$

IBM Q (left) and QVM (right)

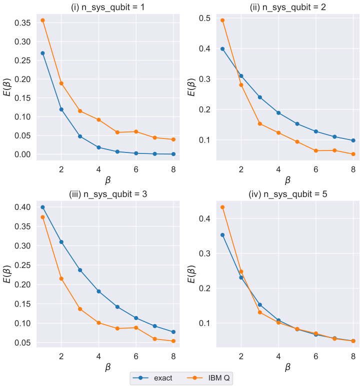

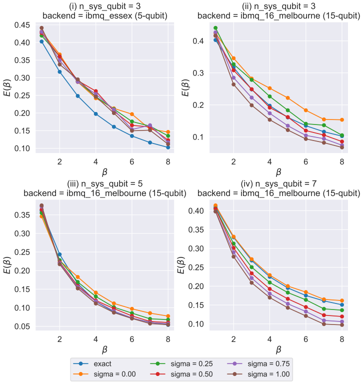

[1]Use the minimally entangled typical thermal state (METTS) algorithm (White, 2009) (Motta et al, 2020)

Conclusion

- Linear system with (Hermitian-)RACBEM: Quantum LINPACK benchmark
- Uses a supremacy circuit as building block.
- Can be easily engineered w.r.t. almost any architecture.
- Maybe only steps away from the supremacy test.
- Hardness? Each U_A is already hard..
- More quantitative ways to measure success and classical hardness: cross-entropy test etc.

References

https://github.com/qsppack/qsppack
https://github.com/qsppack/racbem

- Y. Dong and L. Lin, Random circuit block-encoded matrix and a proposal of quantum LINPACK benchmark [arXiv:2006.04010]
- Y. Dong, X. Meng, K. B. Whaley, L. Lin, Efficient Phase Factor Evaluation in Quantum Signal Processing [arXiv:2002.11649]
- L. Lin and Y. Tong, Optimal quantum eigenstate filtering with application to solving quantum linear systems [arXiv:1910.14596]
- D. An and L. Lin, Quantum linear system solver based on time-optimal adiabatic quantum computing and quantum approximate optimization algorithm [arXiv:1909.05500]

Acknowledgement

Thank you for your attention!

Lin Lin
https://math.berkeley.edu/~linlin/

U.S. DEPARTMENT OF ENERGY | Office of Science
QAT4Chem
NSF
Google AI

Two issues of RACBEM

- Not Hermitian.
- Matrix becomes increasingly singular as n increases.

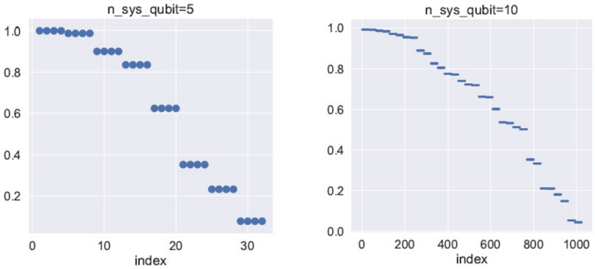

Fixing both issues: Hermitian-RACBEM

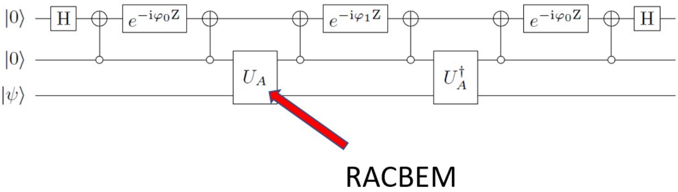

- Simple quantum singular value transformation circuit (QSVT) (Gilyén-Su-Low-Wiebe, 2019) of degree 2.

- Explicit formulation of the Hermitian matrix by tracing out 2 ancilla (the extra ancilla can be reused later).

$$\mathfrak{H} = [-2\sin(2\varphi_0)\sin\varphi_1]\, A^\dagger A + \cos(2\varphi_0 - \varphi_1).$$

- Fully adjustable condition number.

- This is the really useful thing.