Title: On random circuits and their uses in compilation

Speakers: Earl Campbell

Series: Perimeter Institute Quantum Discussions

Date: October 28, 2020 - 11:00 AM

URL: http://pirsa.org/20100062

Abstract: I will review work by myself and others in recent years on the use of randomization in quantum circuit optimization. I will present general results showing that any deterministic compiler for an approximate synthesis problem can be lifted to a better random compiler. I will discuss the subtle issue of what "better" means and how it is sensitive to the metric and computation task at hand. I will then review specific randomized algorithms for quantum simulations, including randomized Trotter (Su & Childs) and my group's work on the qDRIFT and SPARSTO algorithms. The qDRIFT algorithm is of particular interest as it gave the first proof that Hamiltonian simulation is possible with a gate complexity that is independent of the number of terms in the Hamiltonian. Since quantum chemistry Hamiltonians have a very large (~N^4) number of terms, randomization is especially useful in that setting. I will conclude by commenting on a recent Caltech paper with interesting results on the derandomization of random algorithms! Some of the relevant preprints include:

https://arxiv.org/abs/1910.06255

https://arxiv.org/abs/1811.08017

https://arxiv.org/abs/1612.02689

Earl Campbell

On random quantum circuits and their uses in compilation

an overview of work by myself and others

All work done at University of Sheffield
talk includes work with collaborators: Yingkai Ouyang and David White

Currently a contractor for AWS Center for Quantum Computing,

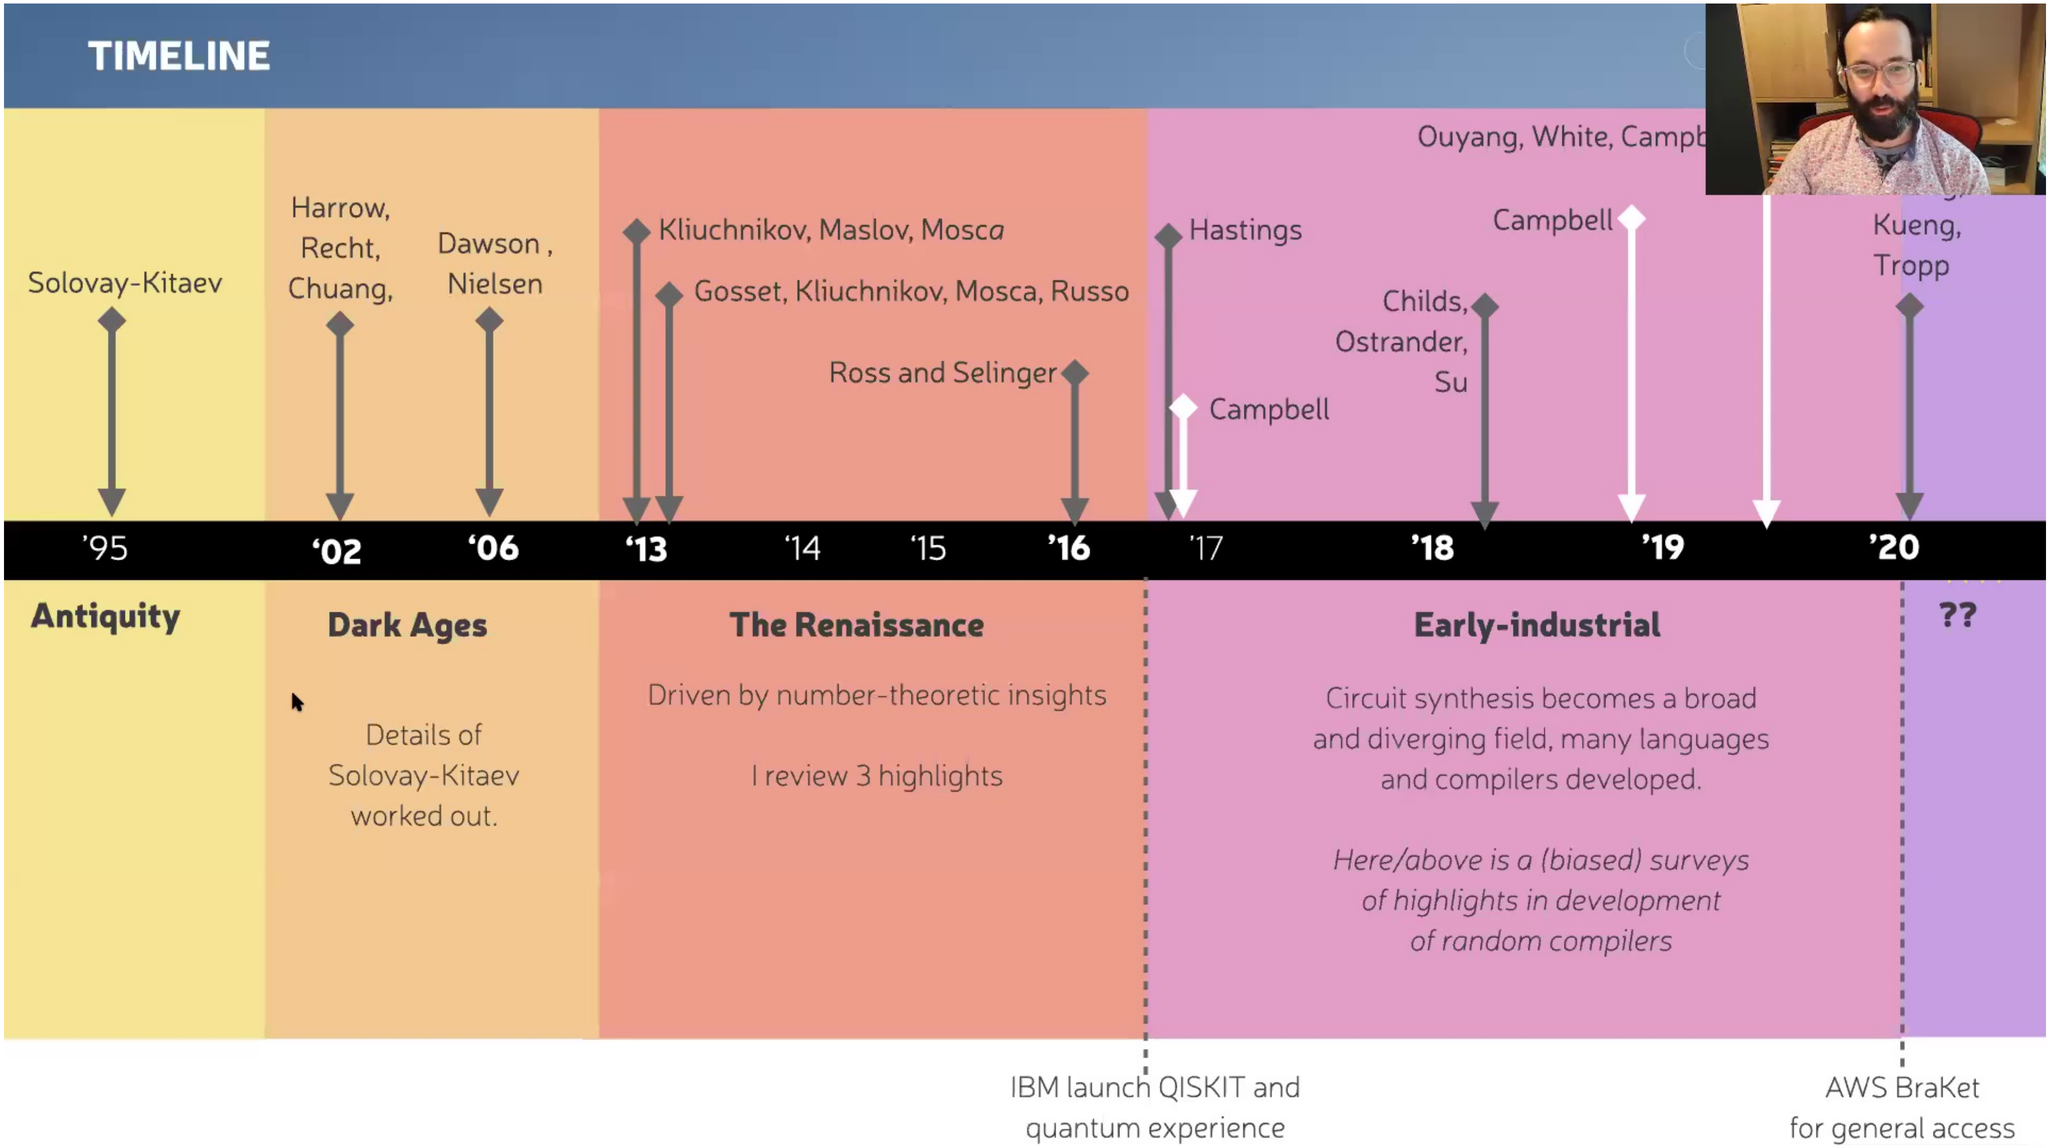
TIMELINE
Solovay-Kitaev
Harrow, Recht, Chuang,
Dawson , Nielsen
Kliuchnikov, Maslov, Mosca
Gosset, Kliuchnikov, Mosca, Russo
Ross and Selinger
Hastings
Campbell
Childs, Ostrander, Su
Campbell
Ouyang, White, Campb
Kueng, Tropp
'95
'02
'06
'13
'14
'15
'16
'17
'18
'19
'20
Antiquity
Dark Ages
Details of Solovay-Kitaev worked out.
The Renaissance
Driven by number-theoretic insights
I review 3 highlights
Early-industrial
Circuit synthesis becomes a broad and diverging field, many languages and compilers developed.
Here/above is a (biased) surveys of highlights in development of random compilers
??
IBM launch QISKIT and quantum experience
AWS BraKet for general access

Antiquity & middle ages

GATE SYNTHESIS / COMPILING

Gate set $\mathcal{G}$

A collection of unitaries used to build circuits

e.g. from $\mathcal{G} = \{A, B, C\}$ we can build $ABC, ABBC, CBC,$ etc

GATE SYNTHESIS / COMPILING

Gate set $\mathcal{G}$

A collection of unitaries used to build circuits

e.g. from $\mathcal{G} = \{A, B, C\}$ we can build $ABC, ABBC, CBC,$ etc

Universality - Informal statement

$\mathcal{G}$ is **universal** if it can implement any unitary (upto finite precision)

Universality - Formal statement

$\mathcal{G}$ is **universal** if for any target unitary V and $\epsilon > 0$ there exists a finite circuit $U \in \langle \mathcal{G} \rangle$ such that $d(U, V) \leq \epsilon$

Example

Clifford+T or Clifford+Toffoli

GATE SYNTHESIS / COMPILING

Cost model $\mathfrak{C} : \mathcal{G} \to \mathbb{R}_+$

Each elementary gate given a positive valued "cost"

This induces a circuit cost

if $\quad U = \prod_i G_i \quad$ then $\quad \mathfrak{C}(U) = \sum_i \mathfrak{C}(G_i)$

Example

Uniform cost model:

$\mathfrak{C}(G) = 1$ for all $G \in \mathcal{G}$

Magic state cost model / T-count:

$\mathfrak{C}(T) = 1$ and $\mathfrak{C}(C) = 0$ for all C in the Clifford group.

GATE SYNTHESIS / COMPILING

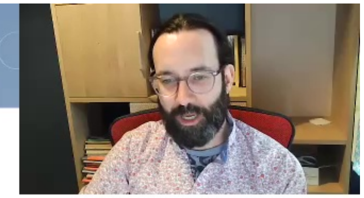

blackbox software

V

$\epsilon > 0$

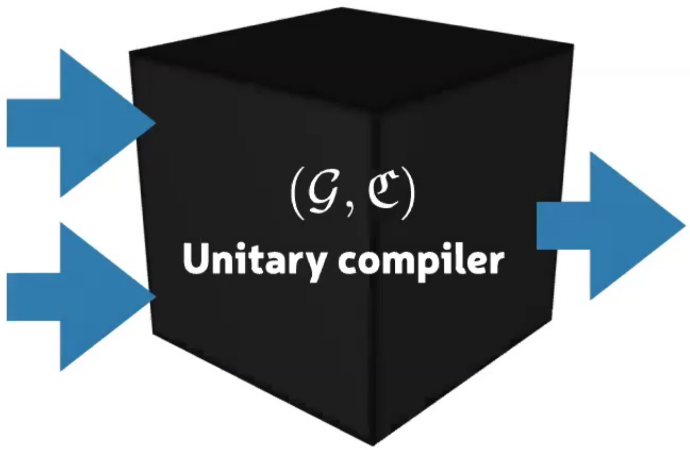

$U \in \langle \mathcal{G} \rangle$

with $||U - V||_\infty \leq \epsilon$

usually with some promise

$\mathfrak{C}(U) \leq f(\epsilon)$

For ***efficient*** compilers The promise function $f(\epsilon)$ is often polylog $f(\epsilon) \leq A \log(1/\epsilon)^\gamma$

An ***optimal*** compiler will have the lowest possible $\mathfrak{C}(U)$ and $f(\epsilon)$

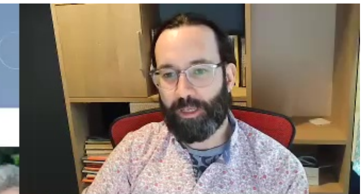

Solovay ('97 email claim)-Kitaev ('97 paper)

Consider any universal gate set $\mathcal{G}$ (generating a group) with uniform cost.

An algorithm to solve the compiling problem using $\mathfrak{C}(U) \leq O(\log(1/\epsilon)^{\gamma})$ where γ is between 3 and 4.

Runtime is efficient in Hilbert space dimension and $\log(1/\epsilon)$

Robert Solovay (1983)

Alexei Kitaev

But what is the constant γ ?

Solovay ('97 email claim)-Kitaev ('97 paper)

Consider any universal gate set $\mathcal{G}$ (generating a group) with uniform cost.

An algorithm to solve the compiling problem using $\mathfrak{C}(U) \leq O(\log(1/\epsilon)^\gamma)$ where γ is between 3 and 4.

Runtime is efficient in Hilbert space dimension and $\log(1/\epsilon)$

Robert Solovay (1983)

Alexei Kitaev

The Solovay-Kitaev algorithm
Dawson and Nielsen QIC 2006

Michael Nielsen

?

Christopher Dawson

undergrad thesis

Aram Harrow

The exponent constant is

$$\gamma = 3.97$$

The prefactor grows with the Hilbert space dimension

GATE SYNTHESIS / COMPILING

Efficient discrete approximations of quantum gates

Harrow, Recht, Chuang

Journal of Mathematical Physics (2002)

Consider any universal gate set $\mathcal{G}$ (generating a group) with uniform cost.

There exists a solution to the compiling problem with $\mathfrak{C}(U) \leq O(\log(1/\epsilon))$ and this is optimal!

No construction algorithm given (except an inefficient brute force search)

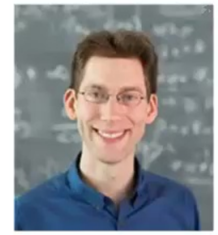

Aram Harrow

Ben Recht

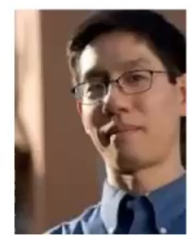

Issac Chuang

Theoretically possible
(Harrow, Recht, Chuang)

$$\gamma = 1$$

Practically feasible for modest Hilbert space dimension
(Solovay,Kitaev) see also (Harrow and Nielsen,Dawson)

$$\gamma = 3.97$$

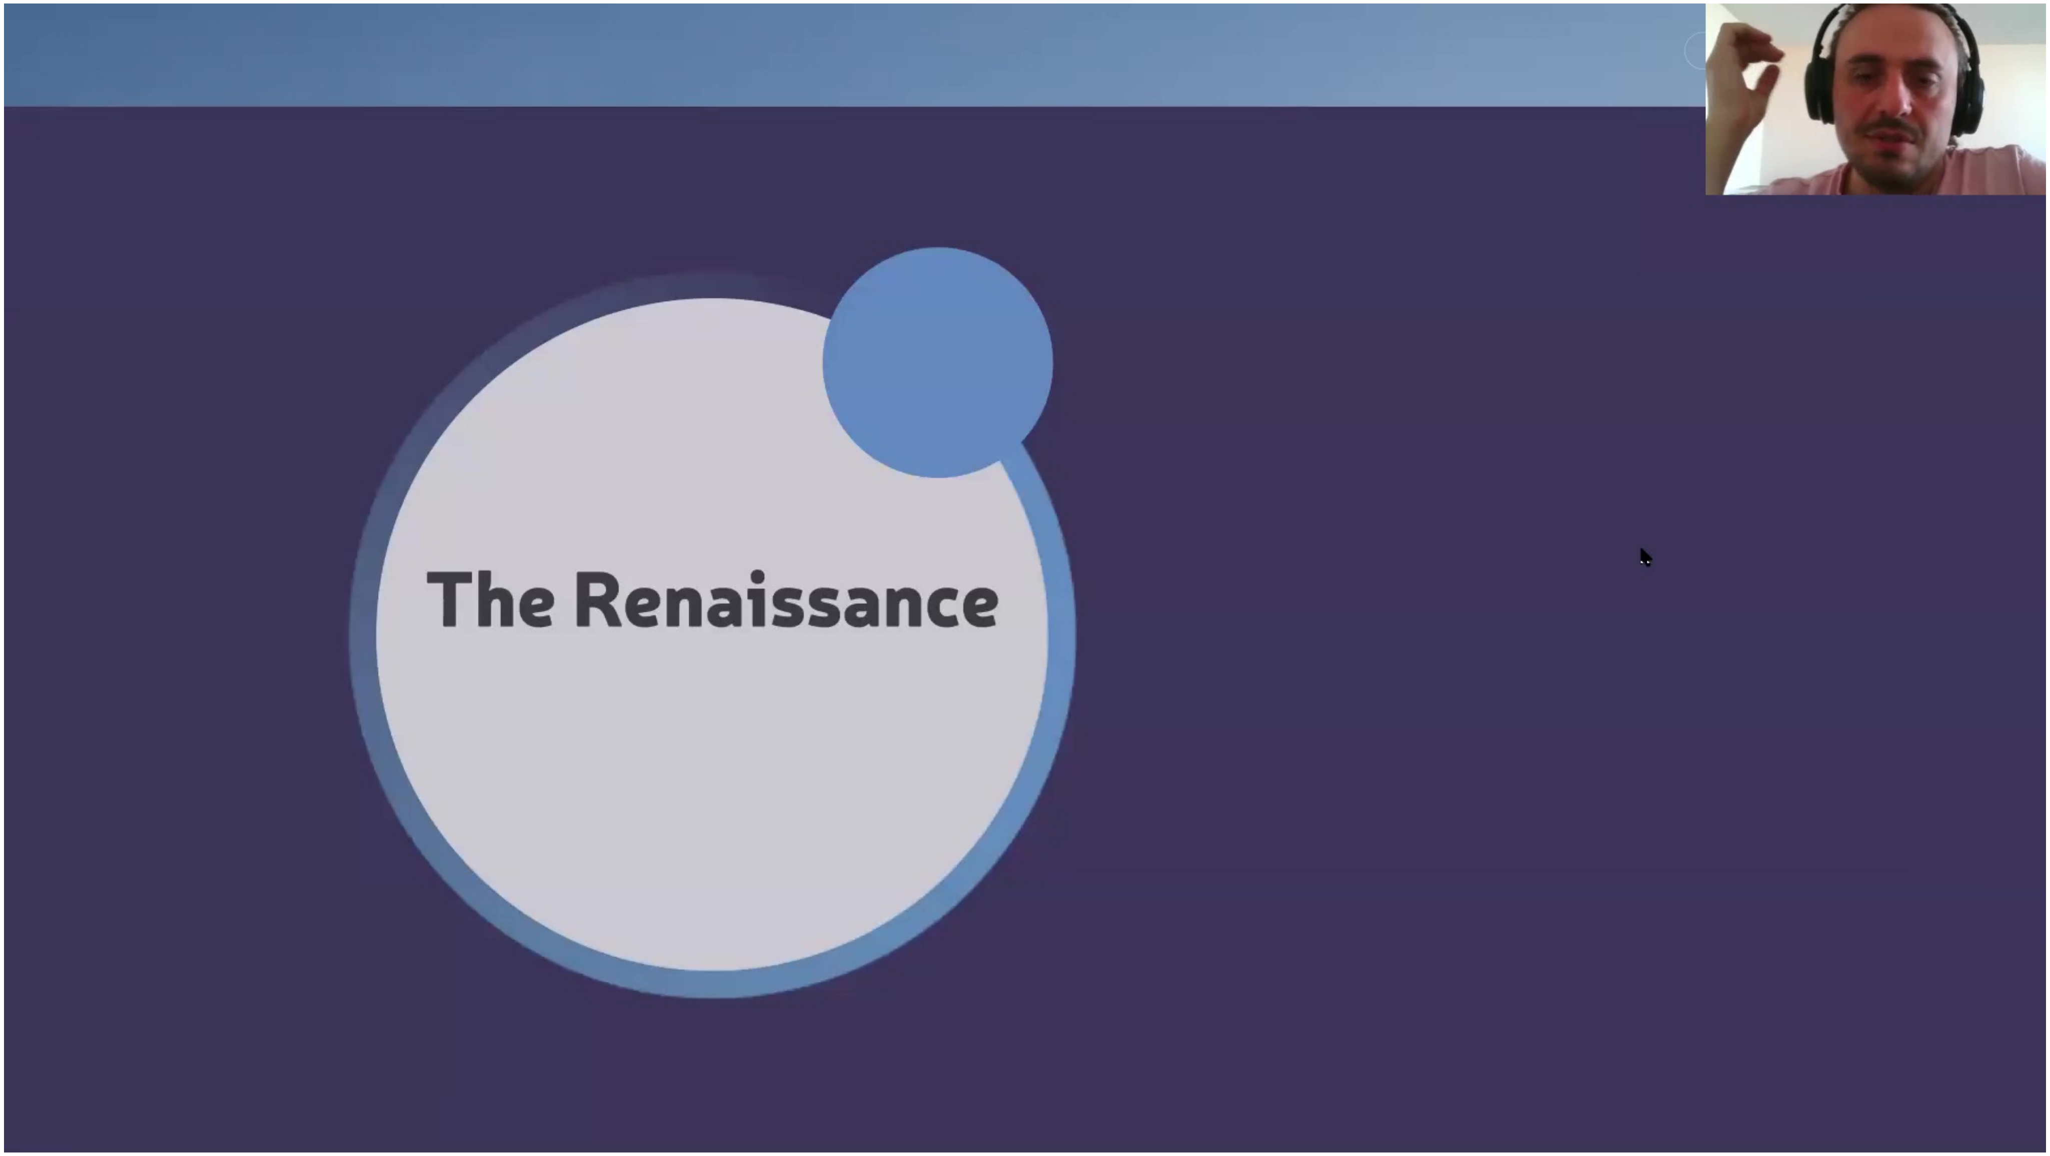
The Renaissance

FOCUS ON SPECIFIC GATE SETS

Progress driven by focusing on a specific gate set important in fault-tolerant quantum computing

Gate set generated by $\mathcal{G} = \{T, CNOT, S, H\}$

Clifford group

$$\{CNOT, S, H\}$$

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

$$S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$$

$CNOT$ 2-qubit control-X

non-Clifford element

$$T = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{i} \end{pmatrix}$$

Also called pi over 8 phase gate

$$\mathfrak{C}(C) = 0 \qquad \mathfrak{C}(T) = 1$$

THE RENAISSANCE

Fast and efficient exact synthesis of single qubit unitaries generated by Clifford and T gates
Kliuchnikov, Maslov, Mosca,
QIC **13 607 (2013) — arXiv:1206.5236**

Vadym Kliuchnikov

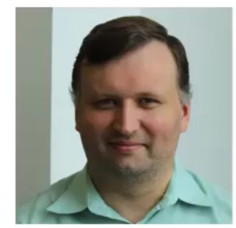

Dmitri Maslov

Mike Mosca

Let $\mathcal{G}$ be single qubit Clifford+T and $U \in \langle \mathcal{G} \rangle$
then we can efficiently compute a circuit achieving the optimal T-count. $\mathfrak{C}(U)$

Proof uses the ring $\mathbb{Z}[\frac{1}{\sqrt{2}}, i]$

In bullet points:

- Only **single qubit** gates
- Only **exact** synthesis & not all unitaries (no epsilon)
- **Efficient**
- **Optimal**

THE RENAISSANCE

An algorithm for the T-count

Gosset, Kliuchnikov, Mosca, Russo

preprint arXiv:1308.4134

Let $\mathcal{G}$ be n-qubit Clifford+T and $U \in \langle \mathcal{G} \rangle$
then we can compute a circuit achieving the optimal T-count $\mathfrak{C}(U)$
in time $\mathcal{O}(2^{n\mathfrak{C}(U)}\mathrm{poly}(n, \mathfrak{C}(U)))$

Proof uses the ring $\mathbb{Z}[\frac{1}{\sqrt{2}}, i]$

In bullet points:

- Any number of qubits!
- Only **exact** synthesis & not all unitaries (no epsilon)
- **Inefficient!**
- **Optimal**

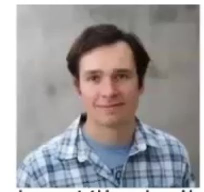

Vadym Kliuchnikov

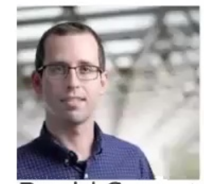

David Gosset

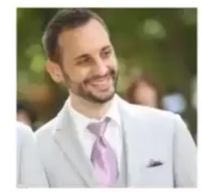

Vincent Russo

Mike Mosca

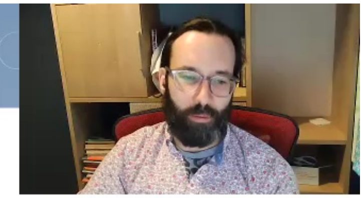

THE RENAISSANCE

Optimal ancilla-free Clifford+T approximation of z-rotations
Ross and Selinger
QIC ***16*** 901 (2016) — arXiv:1403.2975

Let $\mathcal{G}$ be single qubit Clifford+T and $U = \exp(-i\theta Z)$
then for any $\epsilon > 0$ we can find approximation V with $d(U,V) \leq \epsilon$
and optimal T-count. $\mathfrak{C}(V) \leq 4\log(1/\epsilon) + O(\log\log(1/\epsilon))$

In bullet points:

- Only single-qubit Z-axis rotations
- Inexact synthesis
- **Efficient!**
- **Optimal** (under mild assumptions)
- Practically means 100s of gates rather than 10,000s of gates for Solovay-Kitaev

$\gamma = 1$

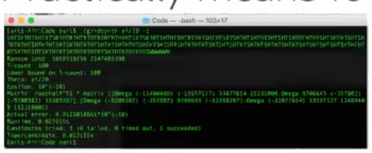

Convenient command line tool

Neil “Julien” Ross

Peter Selinger

THE RENAISSANCE

Still leaves open:

- Heuristics for many-qubit synthesis
  - *too hard to hope for practical, optimal solutions*
- Approaches for specific tasks
  - *e.g. Hamiltonian simulation - more on this later*
- Benefits of ancilla and measurements
  - *Many nice/partial results but little known regarding optimality.*
- ***Benefits of randomness - remainder of this talk***

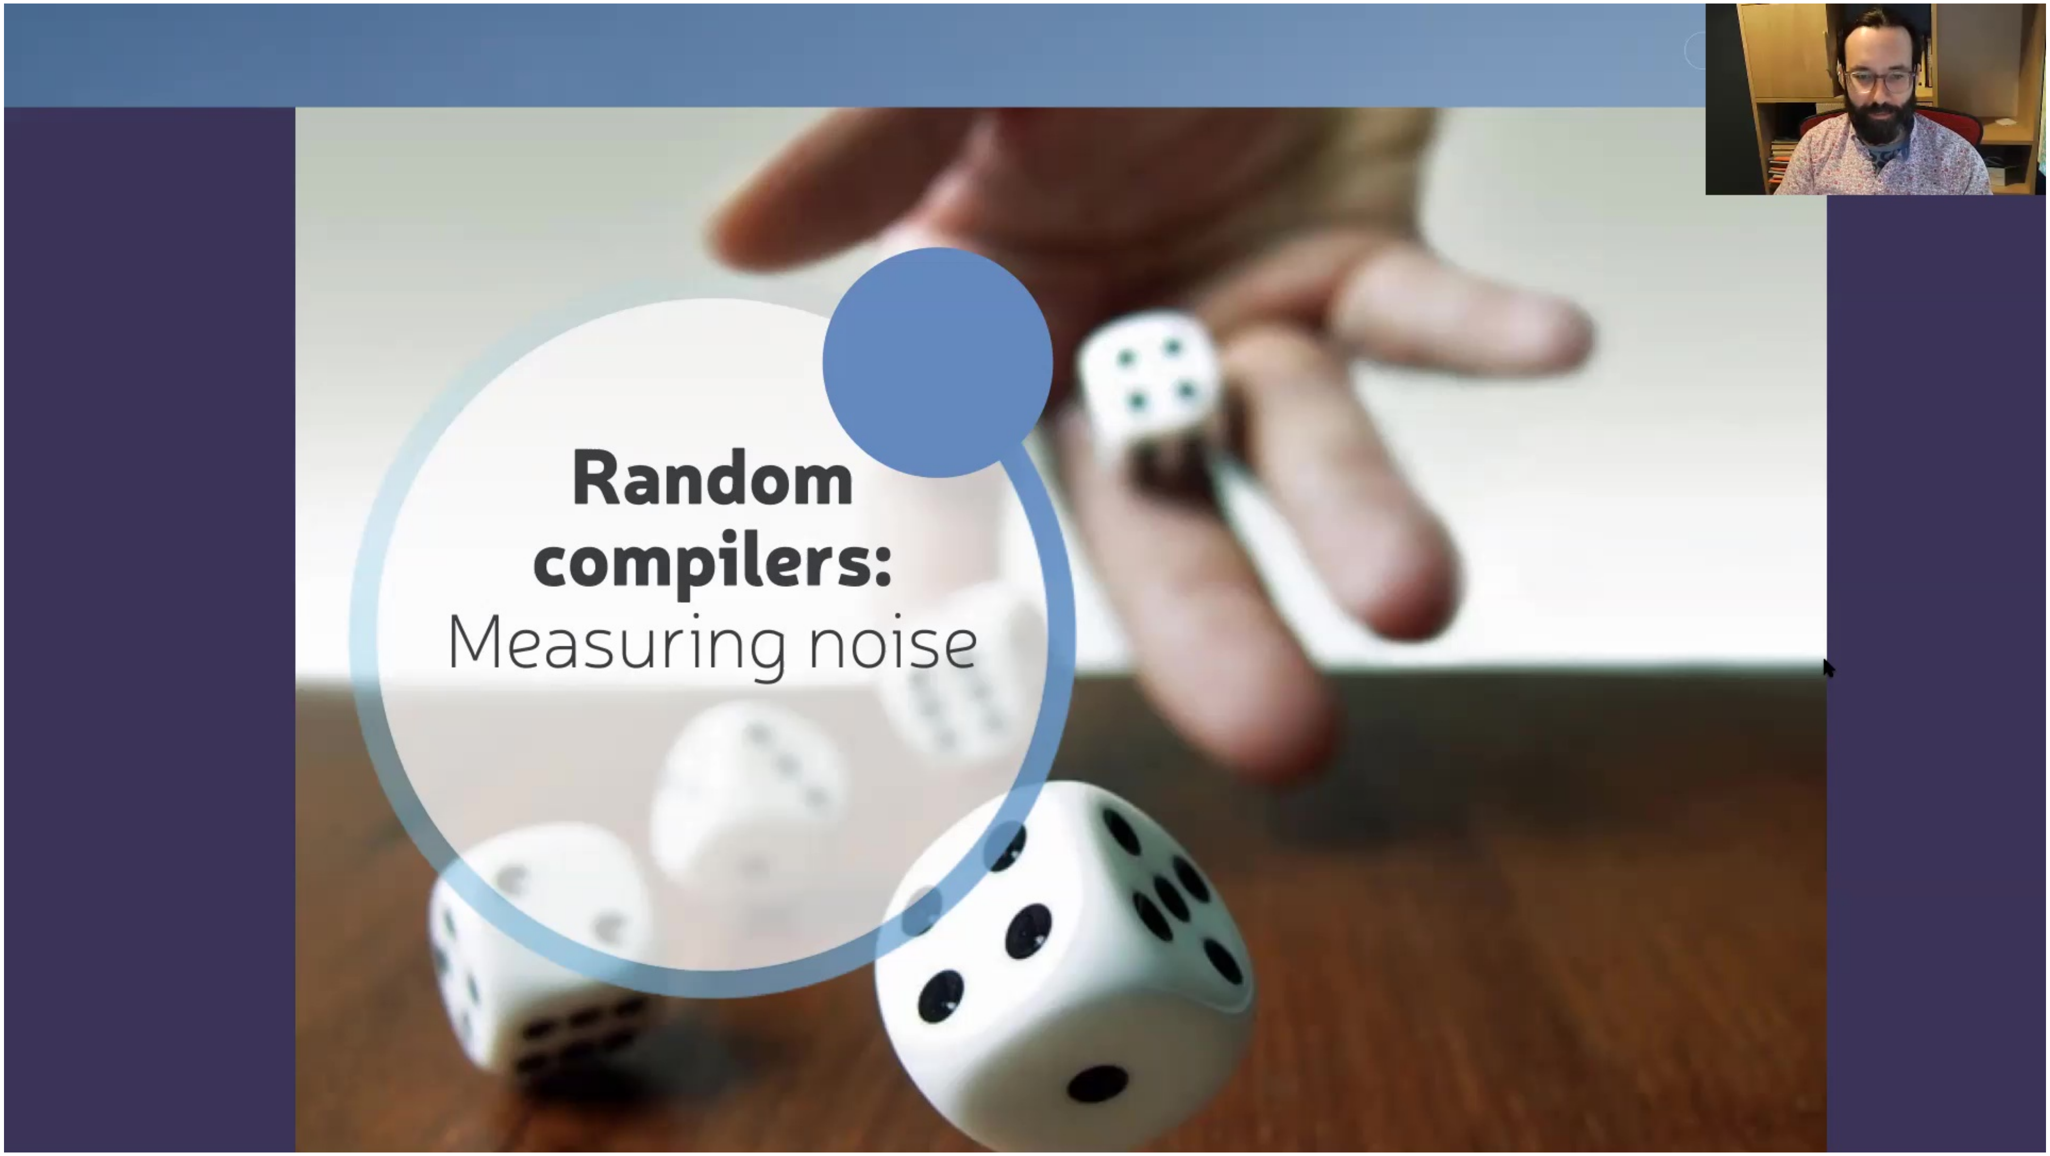
Random compilers:
Measuring noise

MEASURING ERRORS

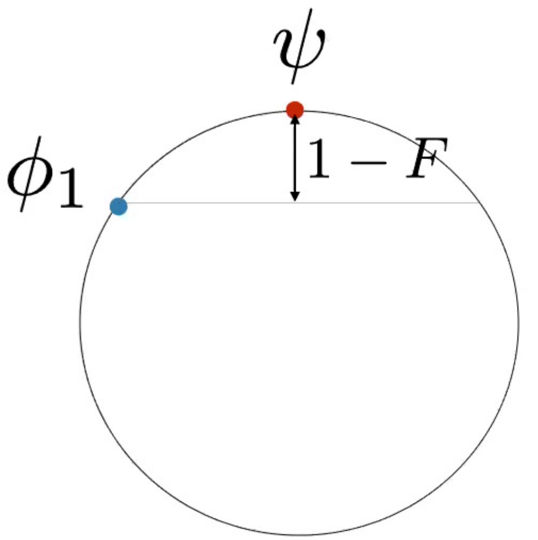

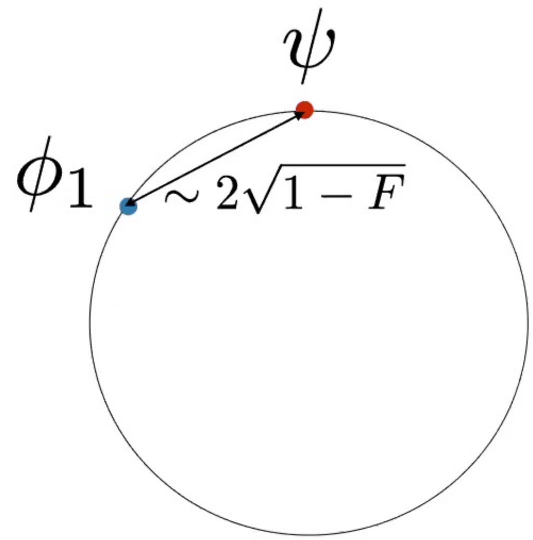

Fidelity

$$F(\phi_1, \psi) = \langle \phi_1 | \psi \rangle \langle \psi | \phi_1 \rangle$$

1-norm error $||M||_1 = \mathrm{Tr}[\sqrt{MM^\dagger}]$

where $M = |\phi_1\rangle\langle\phi_1| - |\psi\rangle\langle\psi|$

MEASURING ERRORS

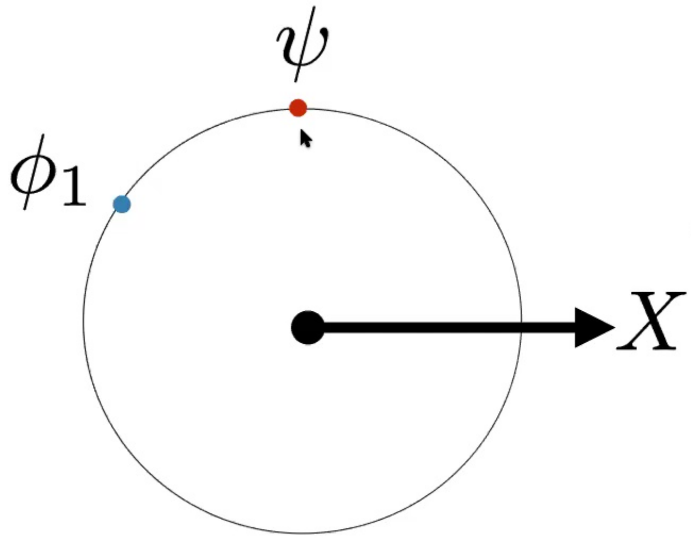

1-norm error $||M||_1 = \mathrm{Tr}[\sqrt{MM^\dagger}]$

where $M = |\phi_1\rangle\langle\phi_1| - |\psi\rangle\langle\psi|$

Measure some observable X, then

Error in probability of measurement outcomes $\leq$ **1-norm error**

Not fidelity!!!

MEASURING ERRORS

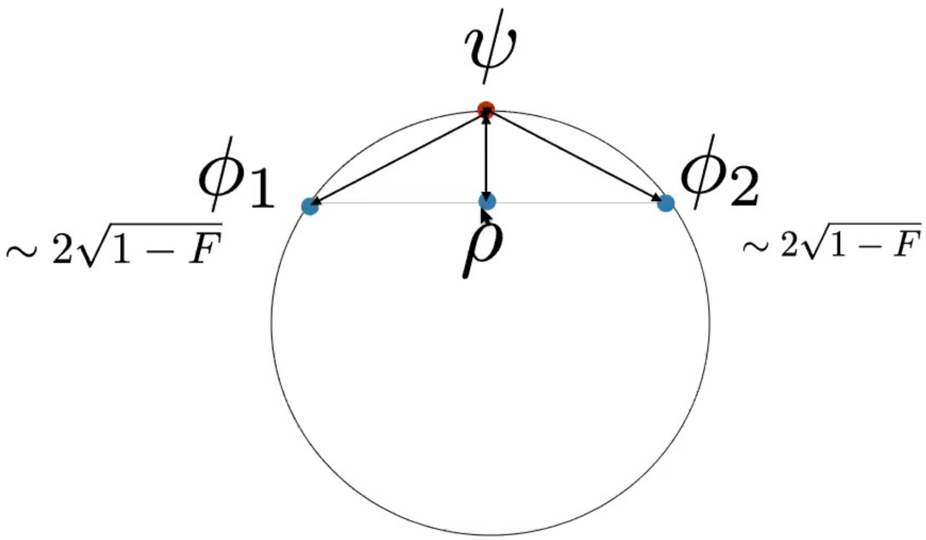

$\rho =$ **50% probability** ϕ_1
50% probability ϕ_2

Randomness quadratically reduced
1-norm error!

Although (average) fidelity is unchanged

GATE SYNTHESIS / COMPILING

Random compiling problem

Given a V and $\epsilon > 0$ output a probability distribution of circuits, realising

$$\mathcal{E}(\rho) = \sum_i p_i U_i \rho U_i^\dagger$$

such that $d(\mathcal{E}, \mathcal{V}) \leq \epsilon$ and minimise $\mathfrak{C}(\mathcal{E})$

Def:
$\mathcal{U}(\cdot) := U \cdot U^\dagger$

GATE SYNTHESIS / COMPILING

Random compiling problem

Given a V and $\epsilon > 0$ output a probability distribution of circuits, realising

$$\mathcal{E}(\rho) = \sum_i p_i U_i \rho U_i^\dagger$$

such that $d(\mathcal{E}, \mathcal{V}) \leq \epsilon$ and minimise $\mathfrak{C}(\mathcal{E})$

Def:
$\mathcal{U}(\cdot) := U \cdot U^\dagger$

What distance we use matters!

Diamond norm distance

$$d_\diamond(\mathcal{E}, \mathcal{U}) := \frac{1}{2}||\mathcal{E} - \mathcal{U}||_\diamond = \frac{1}{2}\max_\rho \frac{||(\mathcal{E} \otimes \mathbb{I})(\rho) - (\mathcal{U} \otimes \mathbb{I})(\rho)||_1}{||\rho||_1}$$

where $||X||_1 := \mathrm{Tr}[\sqrt{X^\dagger X}]$ is the Schatten 1-norm

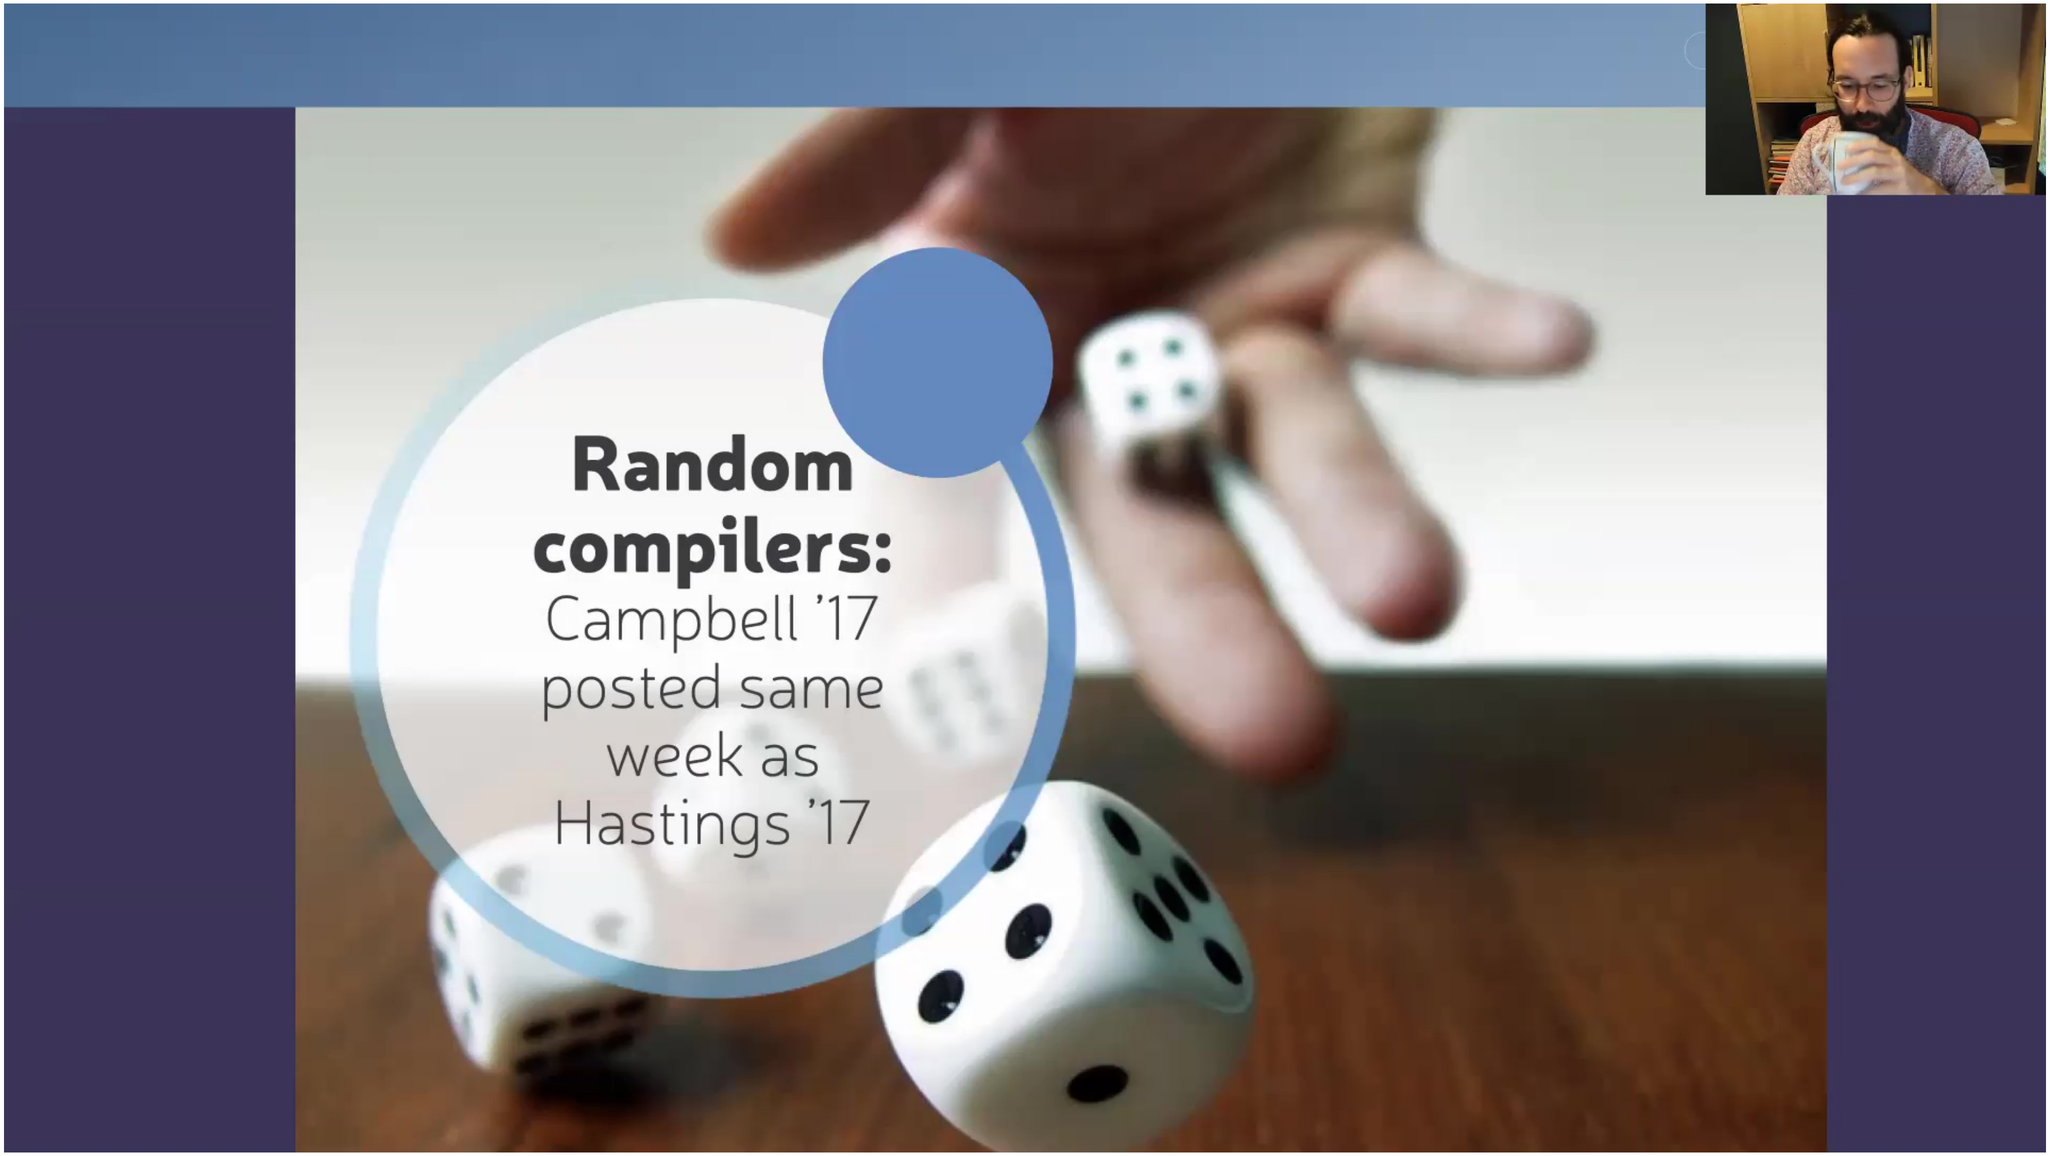
Random
compilers:
Campbell '17
posted same
week as
Hastings '17

Shorter gate sequences for quantum computing by mixing unitaries

Earl Campbell

Phys. Rev. A 95, 042306 (2017) **— arXiv:1612.02689**

Turning Gate Synthesis Errors into Incoherent Errors

Matthew B. Hastings[1,2]

[1] *Station Q, Microsoft Research, Santa Barbara, CA 93106-6105, USA*

[2] *Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA 98052, USA*

Using error correcting codes and fault tolerant techniques, it is possible, at least in theory, to produce logical qubits with significantly lower error rates than the underlying physical qubits. Suppose, however, that the gates that act on these logical qubits are only approximation of the desired gate. This can arise, for example, in synthesizing a single qubit unitary from a set of Clifford and T gates; for a generic such unitary, any finite sequence of gates only approximates the desired target. In this case, errors in the gate can add coherently so that, roughly, the error ϵ in the unitary of each gate must scale as $\epsilon \lesssim 1/N$, where N is the number of gates. If, however, one has the option of synthesizing one of several unitaries near the desired target, and if an average of these options is closer to the target, we give some elementary bounds showing cases in which the errors can be made to add incoherently by averaging over random choices, so that, roughly, one needs $\epsilon \lesssim 1/\sqrt{N}$. We remark on one particular application to distilling magic states where this effect happens automatically in the usual circuits.

THE RESULT

Preamble

Let $\mathcal{G}$ be a universal gate set with cost measure $\mathfrak{C}$

assume some blackbox unitary compiler

V

$0 < \epsilon \leq 0.01$

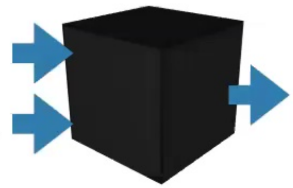

equipped with promise $\mathfrak{C}(U) \leq f(\epsilon)$

Theorem

...*then* there exists a random sequence $\mathcal{E}(\rho) = \sum_j p_j U_j \rho U_j^\dagger$

with $d_\diamond(\mathcal{V}, \mathcal{E}) \leq 10\epsilon^2$ and cost $\mathfrak{C}(\mathcal{E}) \leq f(\epsilon)$

2nd Theorem (paraphrased)

For single qubit axial rotations: *the assumptions can be relaxed and inequalities tightened slightly.*

THE RESULT

"same cost gets you better error suppression"

Theorem

...then there exists a random sequence $\mathcal{E}(\rho) = \sum_j p_j U_j \rho U_j^\dagger$

with $d_\diamond(\mathcal{V}, \mathcal{E}) \leq 10\epsilon^2$ and cost $\mathfrak{C}(\mathcal{E}) \leq f(\epsilon)$

"same error suppression for lower cost?"

Corollary

if the unitary cost is polylog $\mathfrak{C}(U) \leq f(\epsilon) = A \log_2(1/\epsilon)^\gamma$

...then there exists a random circuit $\mathcal{E}(\rho) = \sum_j p_j U_j \rho U_j^\dagger$

with $d_\diamond(\mathcal{V}, \mathcal{E}) \leq \epsilon$ and cost $\mathfrak{C}(\mathcal{E}) \leq C^\gamma f(\epsilon) \sim (1/2)^\gamma f(\epsilon)$

$$C = \left(\frac{1}{2}\right)\left(1 + \frac{\log(A)}{\log(1/\epsilon)}\right) \to_{\epsilon \to 0} \left(\frac{1}{2}\right)$$

APPLICATIONS TO SPECIFIC BLACK BOXES

WE SAVE ~ 2^γ

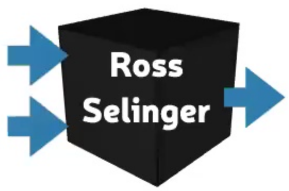

$$\gamma = 1$$

For single qubit Clifford+T gate set
and T-count cost metric

Resource saving 2X

	optimal unitary	random compiling
$\epsilon = 10^{-20}$	600	314 or less
$\epsilon = 10^{-10}$	300	165 or less
$\epsilon = 10^{-5}$	150	90 or less

APPLICATIONS TO SPECIFIC BLACK BOXES

WE SAVE ~ 2^{γ}

$\gamma = 1$

For single qubit Clifford+T gate set
and T-count cost metric

Resource saving 2X

	optimal unitary	random compiling
$\epsilon = 10^{-20}$	600	314 or less
$\epsilon = 10^{-10}$	300	165 or less
$\epsilon = 10^{-5}$	150	90 or less

$\gamma \sim 3.97$

For any gate-set / cost metric

Resource saving $2^{3.97} \sim 15.7$

Still a good option for T-gate synthesis
over a modest number of qubits (more than 1, less than ~10)

SKETCH

ALGORTIHM SKETCH

assume some blackbox unitary compiler

V

ϵ

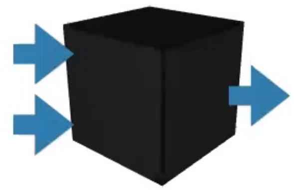

U equipped with promise

$$\mathfrak{C}(U) \leq f(\epsilon)$$

1) Use black-box to find a net of solutions "nearby" to target

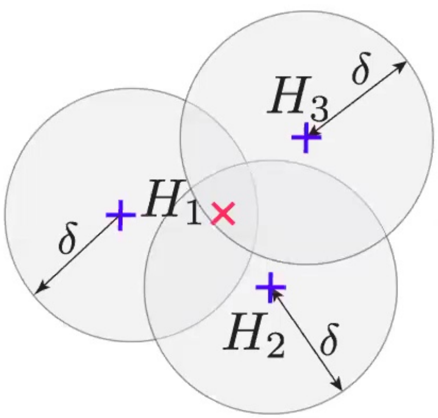

2) Lift to find "nearby" Hamiltonian generators

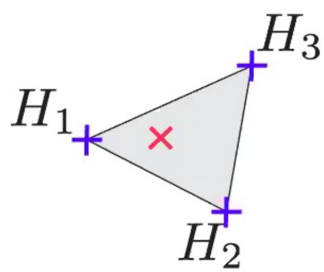

3**)** Solve to find probability weights assigned to each circuit

$$\sum_j p_j H_j = H_{\text{target}}$$

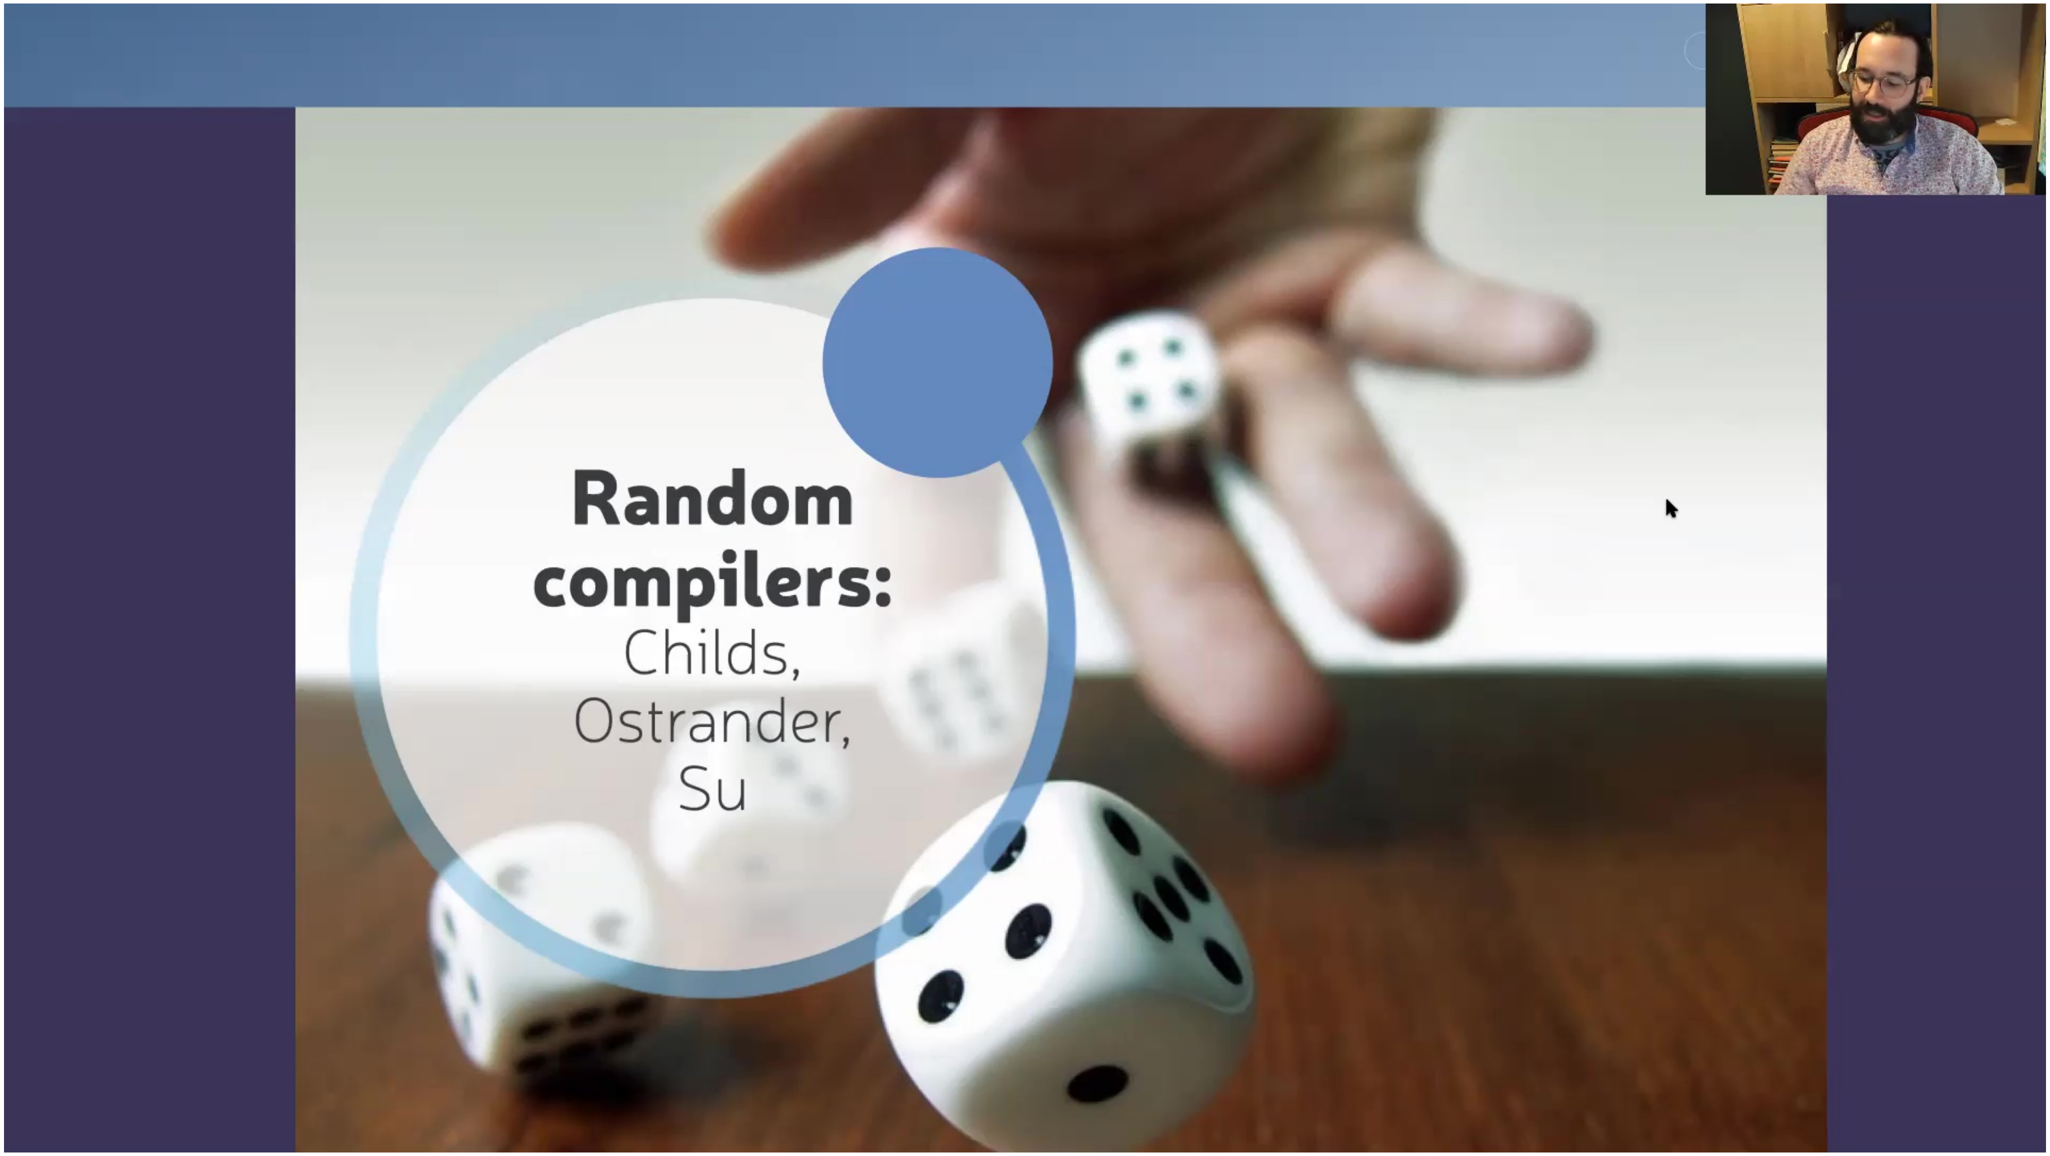
Random compilers:
Childs, Ostrander, Su

RANDOM PERMUTATIONS

Faster Quantum simulation by randomisation

Childs, Ostrander, Su

Quantum 3, 182 (2019). — arXiv:1805.08385

Andrew Childs

Aaron Ostrander

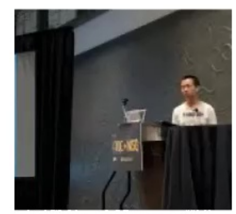

Yuan Su

Task: Hamiltonian simulation

Given Hamiltonian with fixed decomposition

$$H = \sum_{j=1}^{L} c_j H_j \quad \text{with} \quad ||H_j|| = 1$$

Approximate $U = \exp(iHt)$ for some t

Using gate set $\mathcal{G} = \{\exp(iH_j\tau)\}$ and cost model is number of such gates

When H_j are Pauli operators further synthesis possible by Ross-Selinger

Old (deterministic) solution

Trotter-Suzuki Product Formulae

1st order

$$\exp(iHt) \approx \left(\prod_{j=1}^{L} \exp(iH_j t/r) \right)^r$$

e.g. One 1st order step

$$H = 5H_1 + 2H_2 + H_3 + H_4 + H_5$$

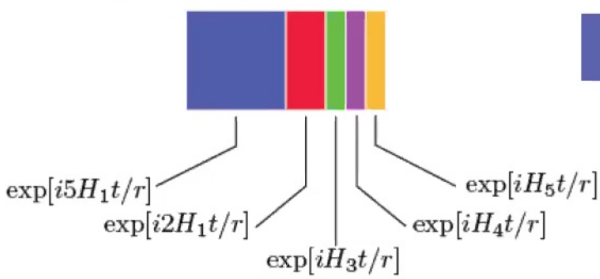

e.g. Ten first order steps

RANDOM PERMUTATIONS

Each permutation
has same error bound.

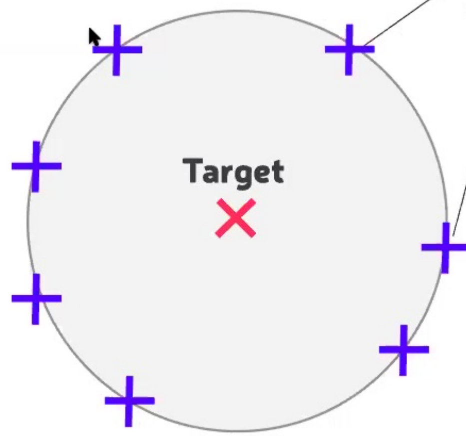

Target unitary is surrounded
by many different nearby solutions
that we can randomly selected from

Can use Campbell/Hastings results!

RANDOM PERMUTATIONS

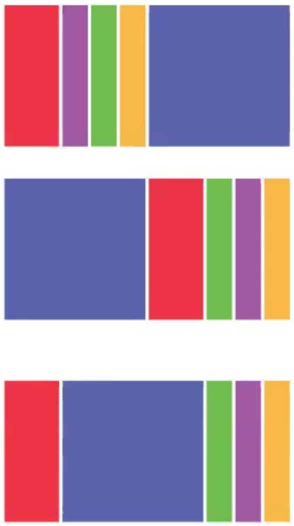

Each permutation
has same error bound.

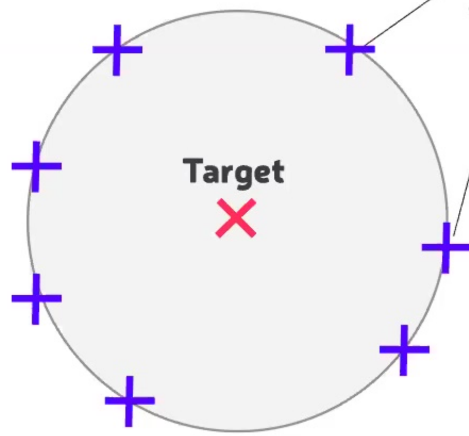

Some permutation of 1 Trotter step

Another permutation of 1 Trotter step

Target unitary is surrounded
by many different nearby solutions
that we can randomly selected from

Can use Campbell/Hastings results!

Ten deterministic first order steps

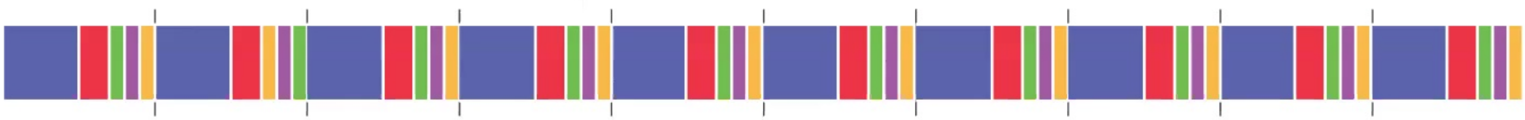

Ten randomised first order steps (each choice must be i.i.d)

Gate complexity
to obtain ϵ error

$$G_{\text{det}} = O(L^4 t^2/\epsilon)$$

$$G_{\text{rand}} = O(L^{2.5} t^{1.5}/\epsilon^{0.5})$$

RANDOM PERMUTATIONS

Paper contains theorem statements & proofs showing improvement

Numerical results for 1D Heisenberg Spin chain with a random magnetic field

r = Number of Trotter steps, so lower is better

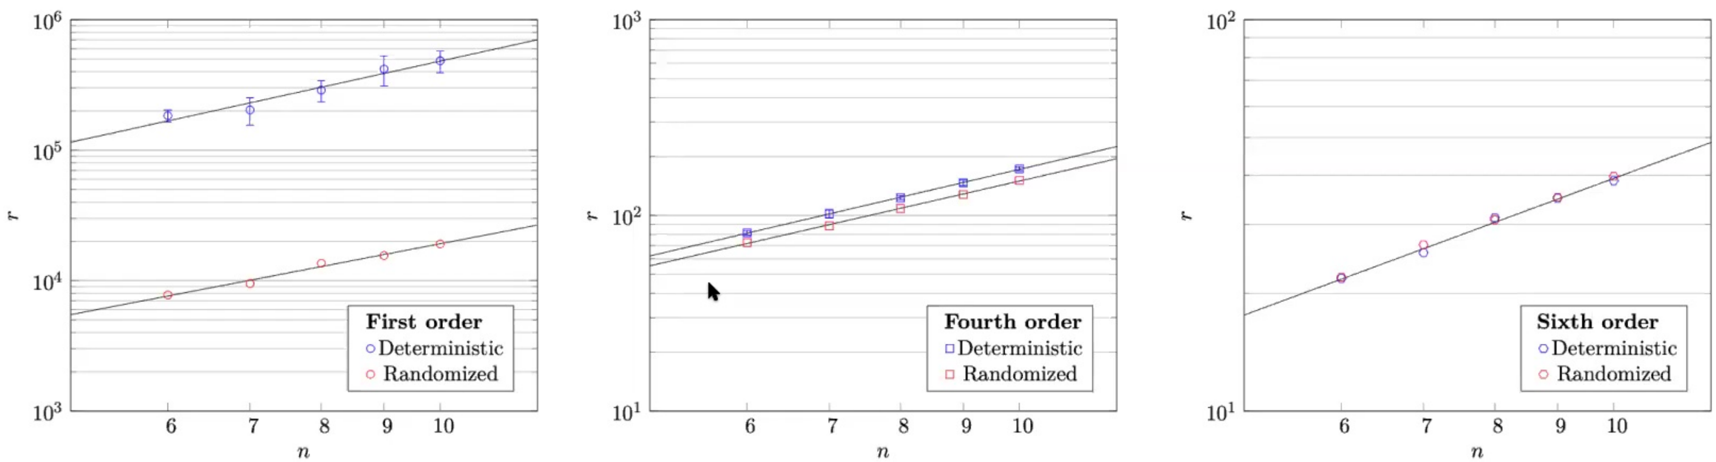

Figure 1: Comparison of the values of r between deterministic and randomized product formulas. Error bars are omitted when they are negligibly small on the plot. Straight lines show power-law fits to the data.


Random compilers:
qDRIFT &
SPARSTO

Task: Hamiltonian simulation

Given Hamiltonian with fixed decomposition

$$H = \sum_{j=1}^{L} c_j H_j \quad \text{with} \quad ||H_j|| = 1$$

Approximate $U = \exp(iHt)$ for some t

Using gate set $\mathcal{G} = \{\exp(iH_j\tau)\}$ and cost model is number of such gates

When H_j are Pauli operators further synthesis possible by Ross-Selinger

BUT, what if L is <u>very large</u> (e.g. quantum chemistry $L = O(N^4)$)
then standard Trotter approaches have large complexity

QDRIFT

A random compiler for fast Hamiltonian simulation
Earl Campbell
Phys. Rev. Lett. **123** 070503 (2019)

Hamiltonian & assumptions

$$H = \sum_{j=1}^{L} c_j H_j$$

w.l.o.g

$$||H_j|| = 1$$

$$c_j \geq 0$$

letting

$$\lambda = \sum_j c_j$$

qDRIFT = quantum drift

With probability

$$p_j = c_j / \lambda$$

Do $\exp(i\lambda t H_j / G)$

and repeat G times

so sequence looks visually like

rather than

Results

Gate complexity

$$G_{qDRIFT} = 2\lambda^2 t^2 / \epsilon$$

No L dependence

PROOF SKETCH

Given $H = \sum_{j=1}^{L} c_j H_j$ and $\mathcal{L}(\rho) = i[H, \rho]$

Want the unitary "channel" $\exp(\mathcal{L}t) = \mathbb{I} + t\mathcal{L} + O(t^2)$ [*]

Define $\mathcal{L}_j(\rho) = i[H_j, \rho]$ so that $\mathcal{L} = \sum_j c_j \mathcal{L}_j$

Randomly selecting 1 gate (and forgetting) gives the channel $\mathcal{E} = \sum_j p_j \exp(\tau \mathcal{L}_j)$

Setting $p_j = \frac{c_j}{\lambda}$ we have $\mathcal{E} = \frac{1}{\lambda} \sum_j c_j \exp(\tau \mathcal{L}_j) = \mathbb{I} + \sum_j \frac{c_j \tau}{\lambda} \mathcal{L}_j + \dots$ [**]

Setting $\tau = \lambda t$ then [*] and [**] match to 1st order

One 1 gate per "step" so no L dependence. Do some carefully bounding of higher order effects and done.

QDRIFT

A random compiler for fast Hamiltonian simulation
Earl Campbell
Phys. Rev. Lett. **123** 070503 (2019)

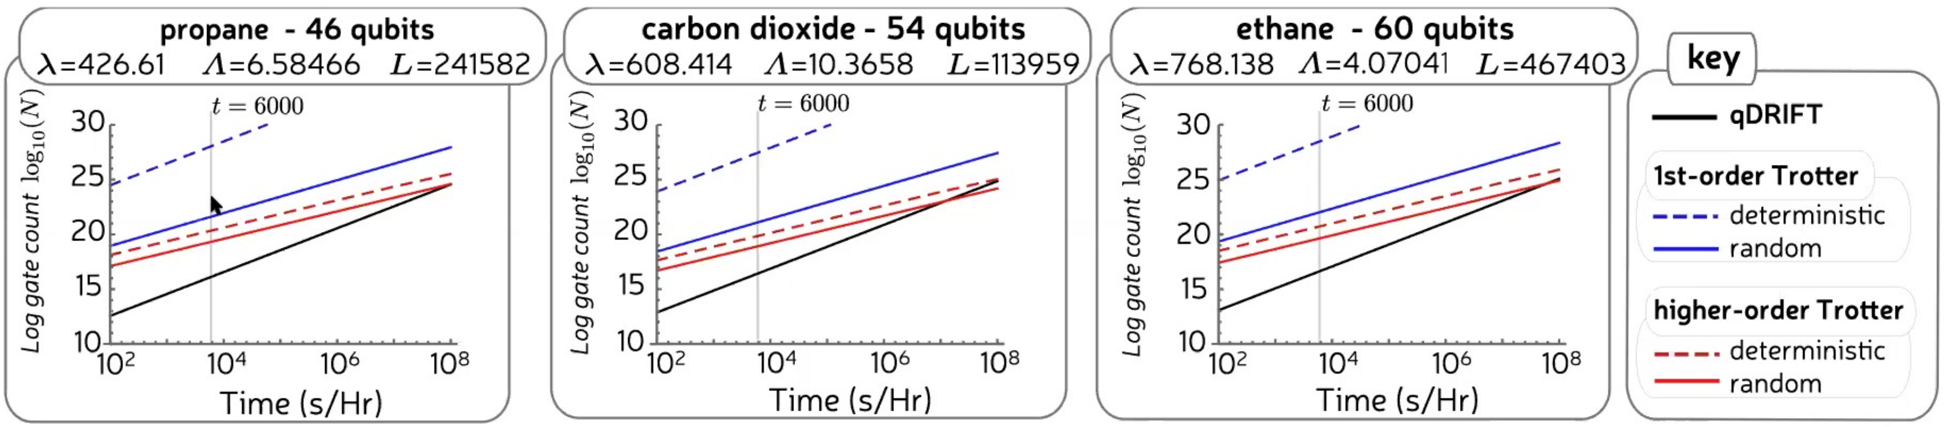

Caveats:

- bound in terms of $\Lambda = \max(c_j)$ would be nice to redo in terms of tighter "commutator" bounds;
- Works well for chemistry, but not spin chain Hamiltonians;
- Error in terms of diamond norm

SPARSTO

Compilation by stochastic Hamiltonian sparsification

Ouyang, White, Campbell

***Quantum* 4 235 (2020)**

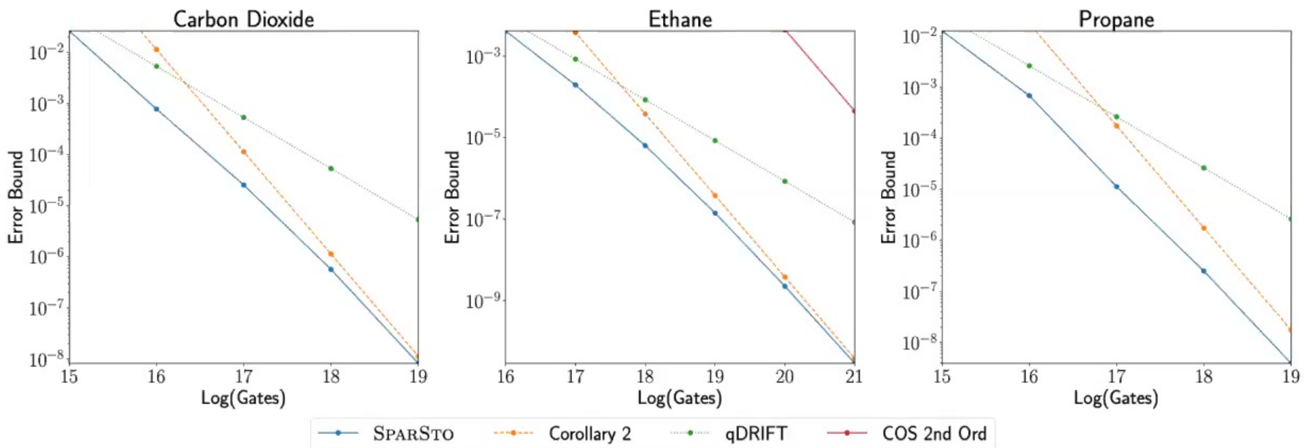

SPARSTO: Sparsify the Hamiltonian to reduce effective ***L*** interpolating between qDRIFT and 2nd order Trotter with some improvements in-between

Yingkai Ouyang

David White


Derandomizing!

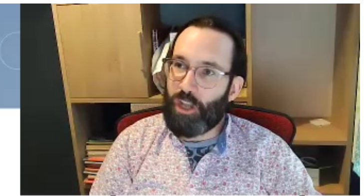

Quantum simulation via randomized product formulas:
Low gate complexity with accuracy guarantees
Chen, (Robert) Huang, Kueng and Tropp
arXiv:2008.11751

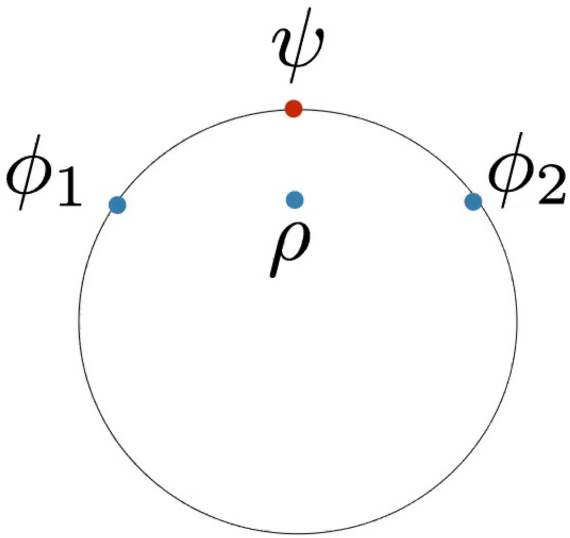

Given a randomised algorithm
(e.g. qDRIFT or SPARSTO)
with some mixed state/channel,
how good are the individual samples?

Chi-Fang Chen

(Robert) Huang

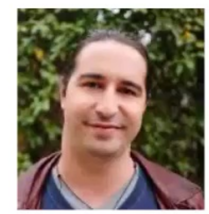

Richard Kueng

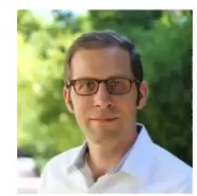

Joel Tropp

DERANDOMIZED

Quantum simulation via randomized product formulas:
Low gate complexity with accuracy guarantees
Chen, (Robert) Huang, Kueng and Tropp
arXiv:2008.11751

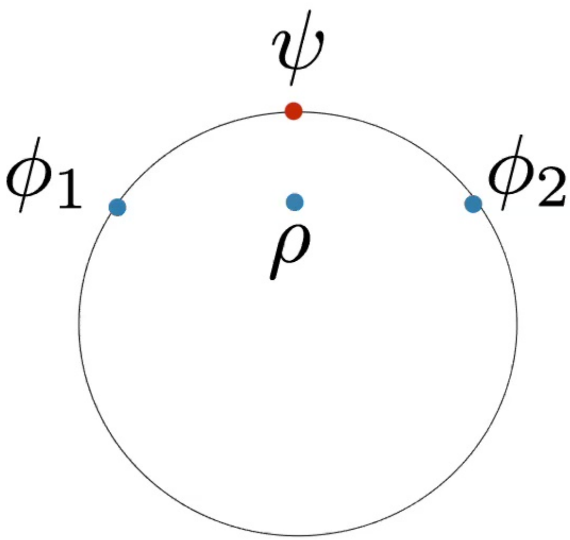

Given a randomised algorithm
(e.g. qDRIFT or SPARSTO)
with some mixed state/channel,
how good are the individual samples?

Derandomized qDRIFT: with prob δ the sampled **n**-qubit circuit has diamond norm error less than ϵ when we use G gates:

$$G = O(nt^2\lambda^2/\epsilon^2) + O(\log(1/\delta)t^2\lambda^2/\epsilon^2)$$

Compare $G_{qDRIFT} = 2\lambda^2 t^2/\epsilon$

Chi-Fang Chen

(Robert) Huang

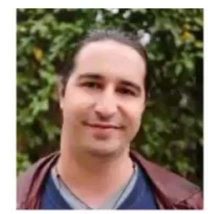

Richard Kueng

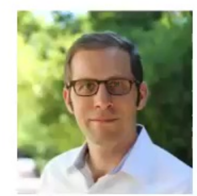

Joel Tropp

OUTLOOK

- Not much in terms of software implementation & experimentation of random protocols, especially w.r.t multi-qubit circuits
- Improved randomised algorithms using convex optimisation.
- Understanding roles of randomness in some cases like phase estimation where we are interested in energy error and not unitary error!
- Randomness in post-Trotter (e.g. LCU) circuits?
- Beyond i.i.d probability distributions?

THANK YOU!