Stochastic optimal control of open quantum systems using path integral methods

Bert Kappen
Faculty of Science
Radboud University Nijmegen
The Netherlands

The quantum picture

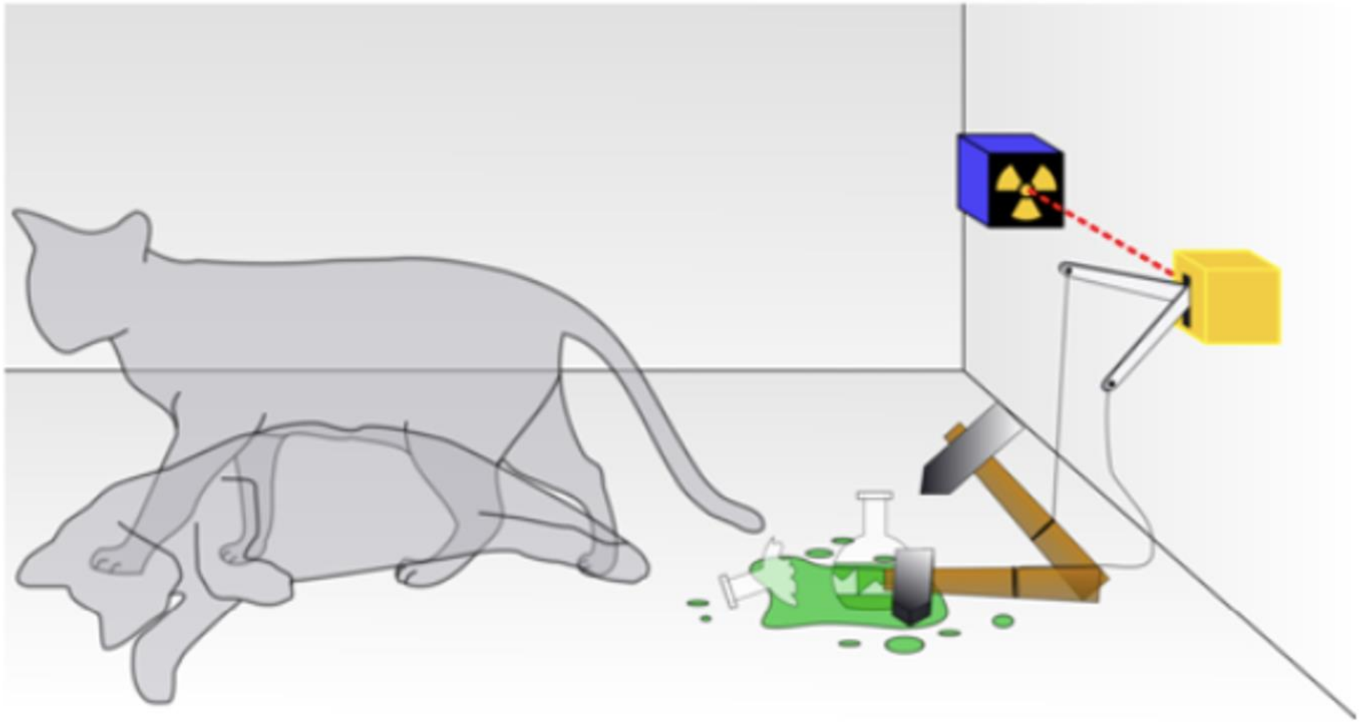

A quantum state $\psi(s)$ represents possible outcomes s simultaneously.

The quantum picture

A quantum state $\psi(s)$ represents possible outcomes s simultaneously.

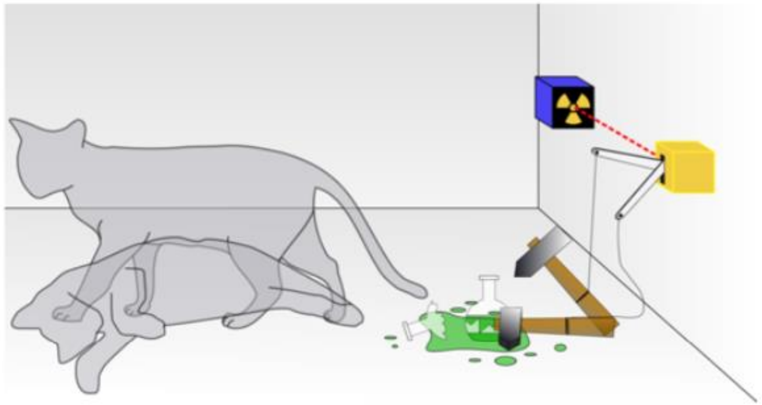

Define a procedure to map a probability distribution to a quantum state:

$$q(s) \qquad \leftrightarrow \qquad \psi(s)$$

Estimate expected values by performing repeated measurement on the same quantum state.

Quantum advantage

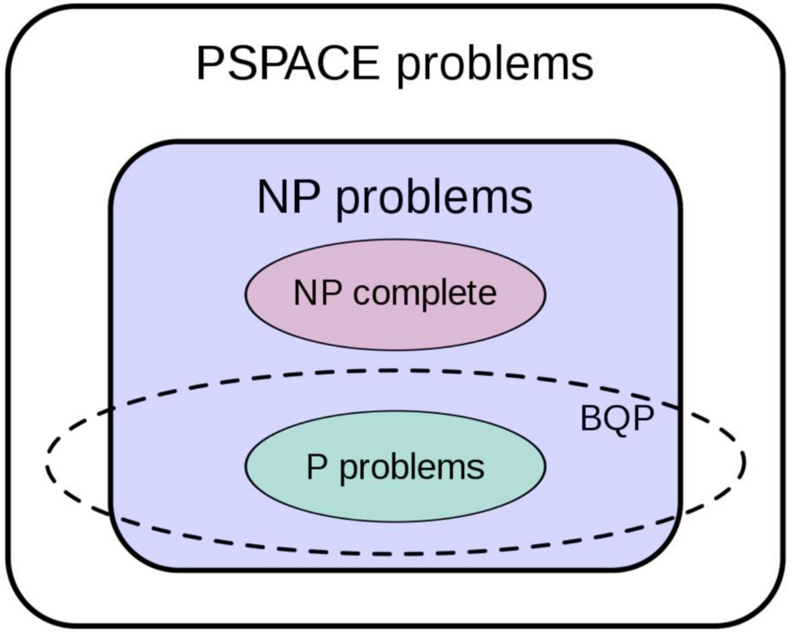

Image credits: wikipedia.org

P: solvable in poly time
NP: solution verifiable in poly time
Pspace: solvable with poly memory
BQP: Bounded Quantum Polynomial

It is conjectured that BQP solves hard problems outside of P, specifically, problems in NP. Examples are

- Integer factorization (Shor's algorithm) $\mathcal{O}\left(e^{N^{\frac{1}{3}}}\right) \rightarrow \mathcal{O}\left(N^2\right)$
- Solving sparse linear system (HHL) $\mathcal{O}\left(N\right) \rightarrow \mathcal{O}\left(\log N\right)$

IBM **Quantum**

2025	2026	2027	2028	2029	2033+
Enhance quantum execution speed and parallelization with partitioning and quantum modularity	Improve quantum circuit quality to allow 7.5K gates	Improve quantum circuit quality to allow 10K gates	Improve quantum circuit quality to allow 15K gates	Improve quantum circuit quality to allow 100M gates	Beyond 2033, quantum-centric supercomputers will include 1000's of logical qubits unlocking the full power of quantum computing
Functions	Mapping collections	Specific libraries			General purpose QC libraries
Resource management	Circuit knitting x p	Intelligent orchestration			Circuit libraries
Flamingo (5K)	Flamingo (7.5K)	Flamingo (10K)	Flamingo (15K)	Starling (100M)	Blue Jay (1B)
Error mitigation	Error mitigation	Error mitigation	Error mitigation	Error correction	Error correction
5k gates 156 qubits	7.5k gates 156 qubits	10k gates 156 qubits	15k gates 156 qubits	100M gates 200 qubits	1B gates 2000 qubits
Quantum modular	Quantum modular	Quantum modular	Quantum modular	Error corrected modularity	Error corrected modularity
Up to 156x7 = 1092 qubits	Up to 156x7 = 1092 qubits	Up to 156x7 = 1092 qubits	Up to 156x7 = 1092 qubits		

Quantum Computing in the NISQ era and beyond

John Preskill

Institute for Quantum Information and Matter and Walter Burke Institute for Theoretical Physics,
California Institute of Technology, Pasadena CA 91125, USA
30 July 2018

Noisy Intermediate-Scale Quantum (NISQ) technology will be available in the near future. Quantum computers with 50-100 qubits may be able to perform tasks which surpass the capabilities of today's classical digital computers, but noise in quantum gates will limit the size of quantum circuits that can be executed reliably. NISQ devices will be useful tools for exploring many-body quantum physics, and may have other useful applications, but the 100-qubit quantum computer will not change the world right away — we should regard it as a significant step toward the more powerful quantum technologies of the future. Quantum technologists should continue to strive for more accurate quantum gates and, eventually, fully fault-tolerant quantum computing.

Quantum state preparation

- Prepare a quantum state that
  - Represents the ground state of a complex molecule (drug design, material design)
  - Solves a complex optimization problem (traveling salesman problem)
  - Sample from a complex distribution (statistical problems)

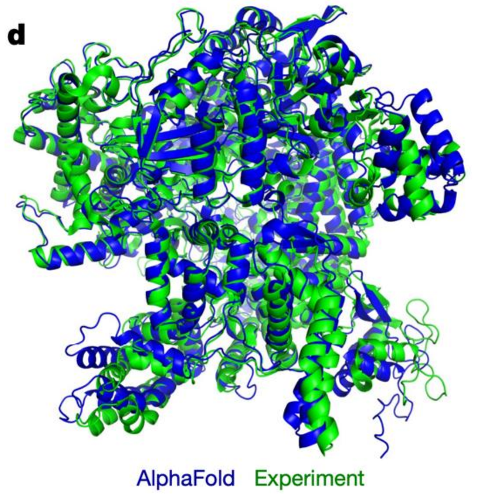

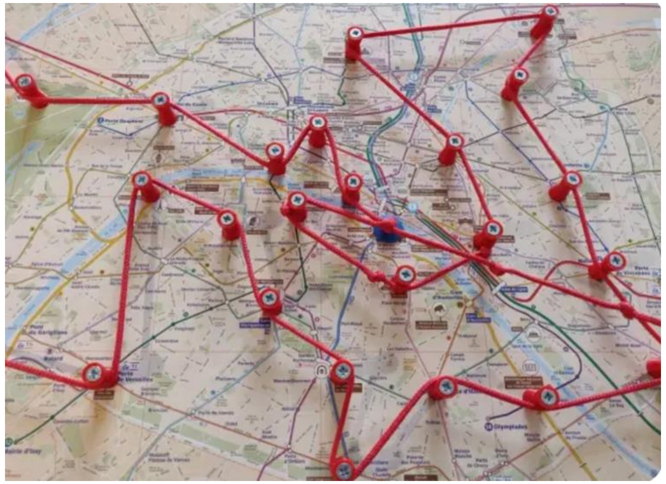

Quantum variational algorithms

Most famous example is the variational quantum eigensolver VQE to find the ground state of a Hamiltonian

$$\min_{\theta} \frac{\langle \psi | H \psi \rangle}{\langle \psi | \psi \rangle} \qquad |\psi\rangle = U(\theta)\,|0\rangle$$

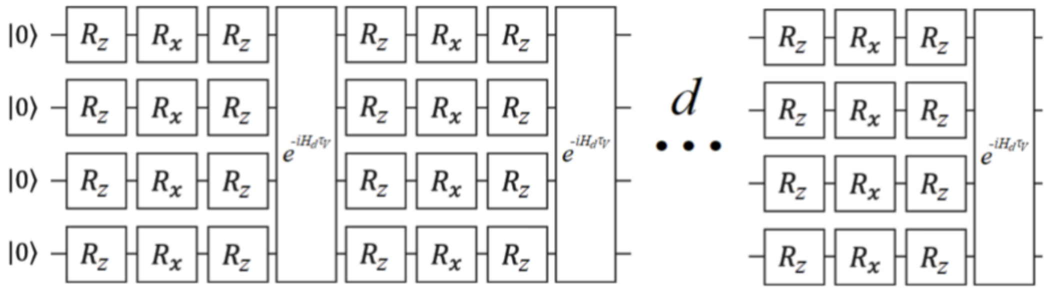

Figure 1a: Gate-based quantum circuit

θ is found by minimizing $R(\theta) = \frac{\langle \psi | H \psi \rangle}{\langle \psi | \psi \rangle}$ with $|\psi\rangle = U(\theta)\,|0\rangle$ with respect to θ using gradient descend.

Barren plateaus

θ is found by minimizing $R(\theta) = \frac{\langle\psi|H\psi\rangle}{\langle\psi|\psi\rangle}$ with $|\psi\rangle = U(\theta)|0\rangle$ with respect to θ using gradient descend.

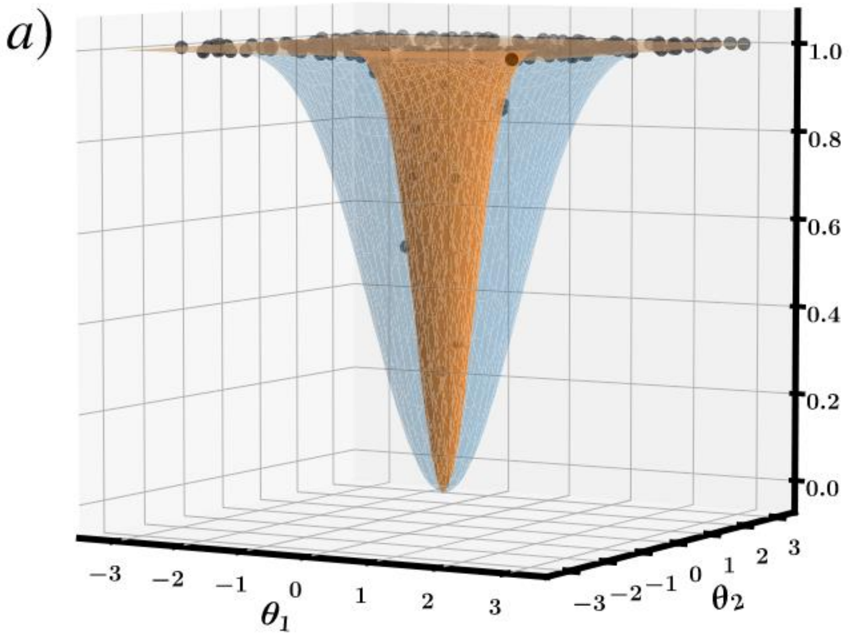

Gradients vanish almost everywhere for large problems

Noise improves convergence

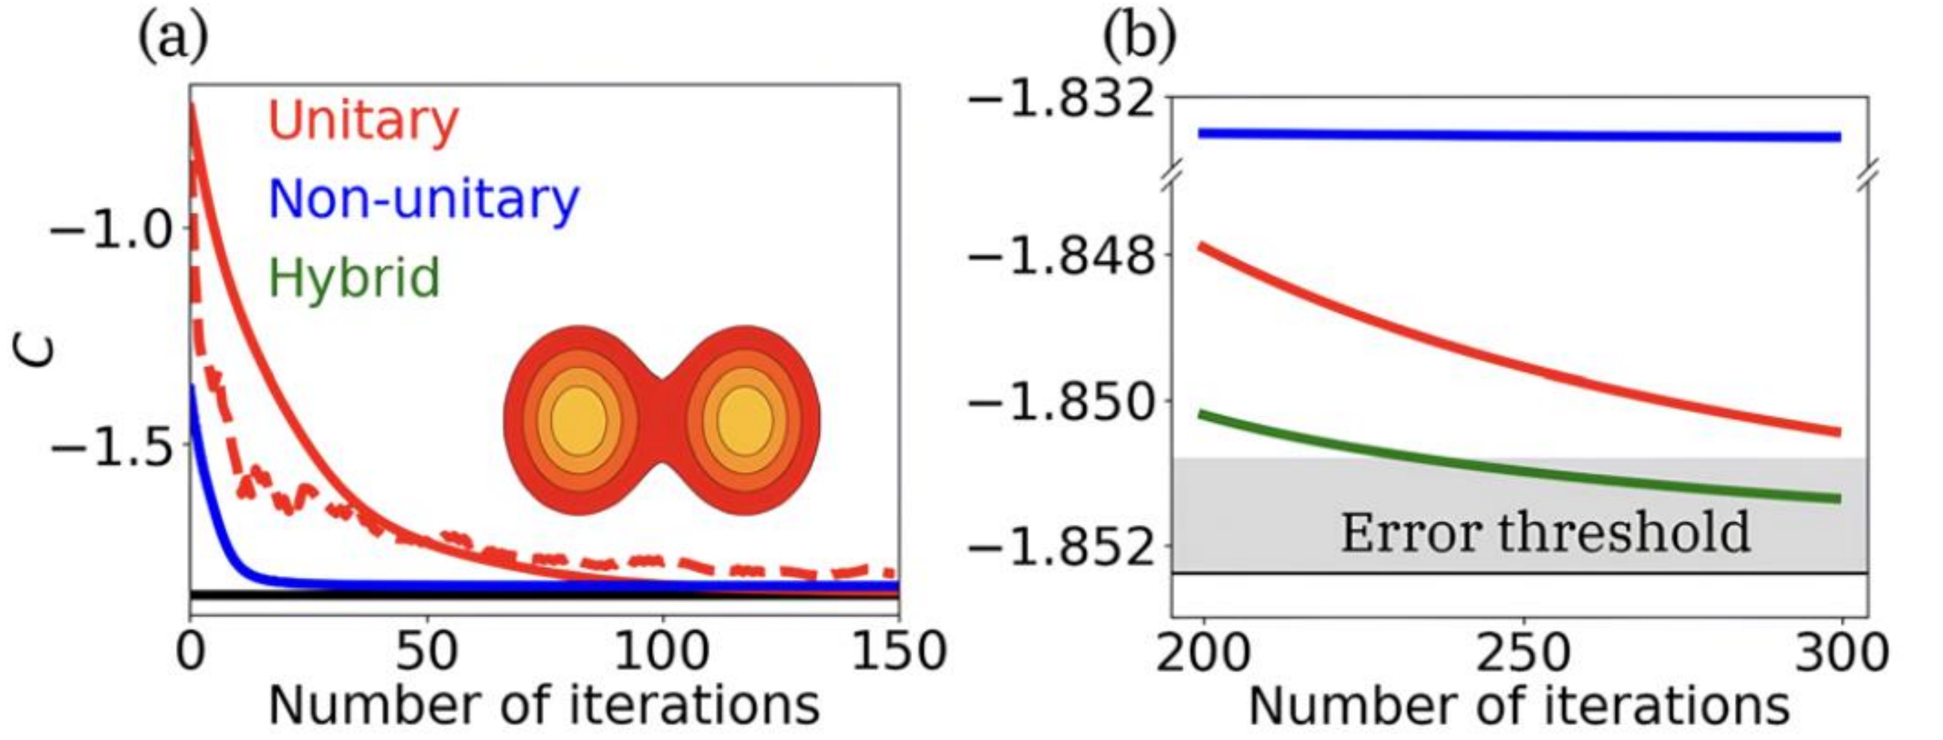

Engineered dissipation to mitigate barren plateaus
Sannia et al. 2024

Optimal control for analog quantum circuits

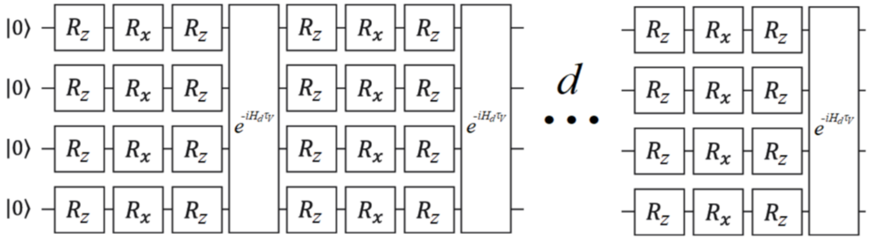

Figure 1a: Gate-based quantum circuit

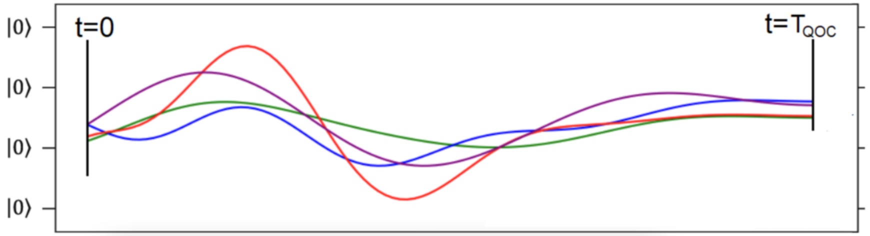

Figure 1b: Pulse-based quantum circuit

$$|\psi_0\rangle \quad \rightarrow \quad \frac{d}{dt}|\psi\rangle = -iH(u)\,|\psi\rangle \quad \rightarrow \quad |\psi_f\rangle$$

with $H(u) = H_0 + \sum_k u_k H_k$. Find optimal u_k.

Analog controlled superconducting qubits have significantly shorter circuit times:
H2, HeH+, 2 qubits: 9 ns versus 500-800 ns
LiH, 4 qubits: 40 ns versus 3.500-82.000 ns
T1,T2 times of IBMQ are order 70.000 ns on average

(Meitei et al. 2021)

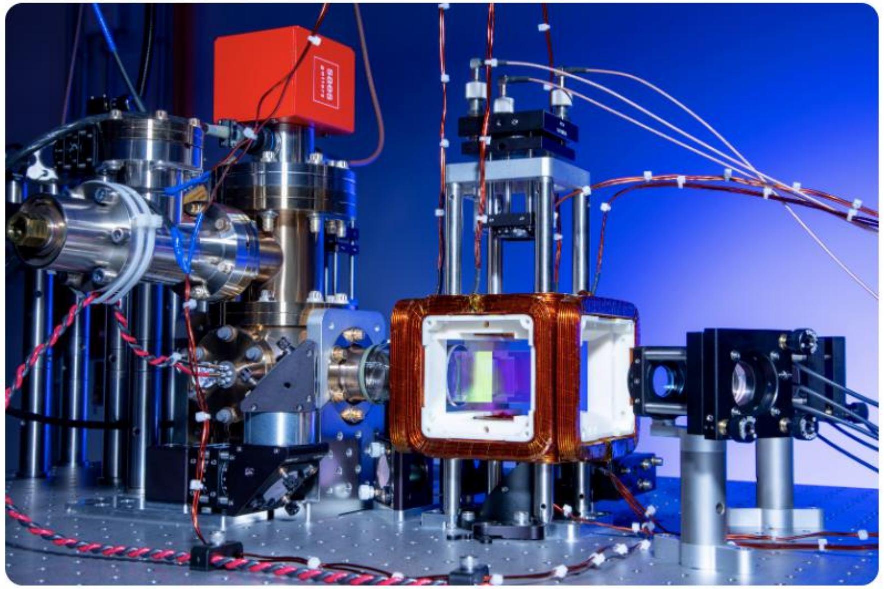

Quantum computing with neutral atoms

QuEra's quantum computing technology uses lasers to arrange and excite individual neutral atoms into highly energetic states. These excited-atom qubits naturally interact at a distance, enabling entanglement and a multi-qubit connectivity that can be turned on and off at will. As atomic positions can be rearranged from one calculation to the next, these processors present extremely flexible and programmable layouts for their users. The ease of assembly and control, and the strong quantum coherence properties of neutral atoms, uniquely positions the technology to access new frontiers in simulating large quantum systems, exploring quantum optimization, and sampling.

QuEra's Aquila processor

Aquila is QuEra's first generation of quantum processing units (QPU) available on Amazon Braket. It operates up to 256 qubits in analog mode. The qubits have long lifetimes, supp
tens of qubit flips before decoherence sets in.

Approach

- Stochastic optimal control of analog open quantum systems
  - Analog yields shorter circuits
  - Stochastic yields better optimization
- Unravelings aka quantum trajectories
- Quantum state preparation as a stochastic optimal control problem
  - path integral control formulation

Outline

- Open quantum systems
  - Unravelings
- Stochastic optimal control
  - Path integral control
- Stochastic optimal control of open quantum systems
  - Examples of quantum state preparation

Unitary evolution is rotation in Hilbert Space

$$\frac{\partial \psi}{\partial t} = -i H \psi$$

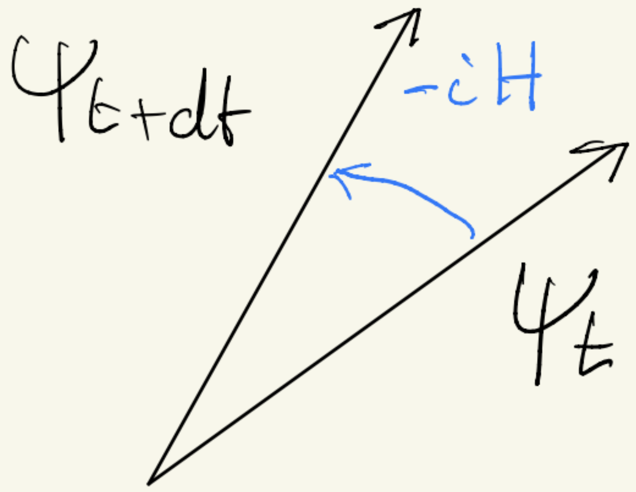

Linear and deterministic

2. Measurement

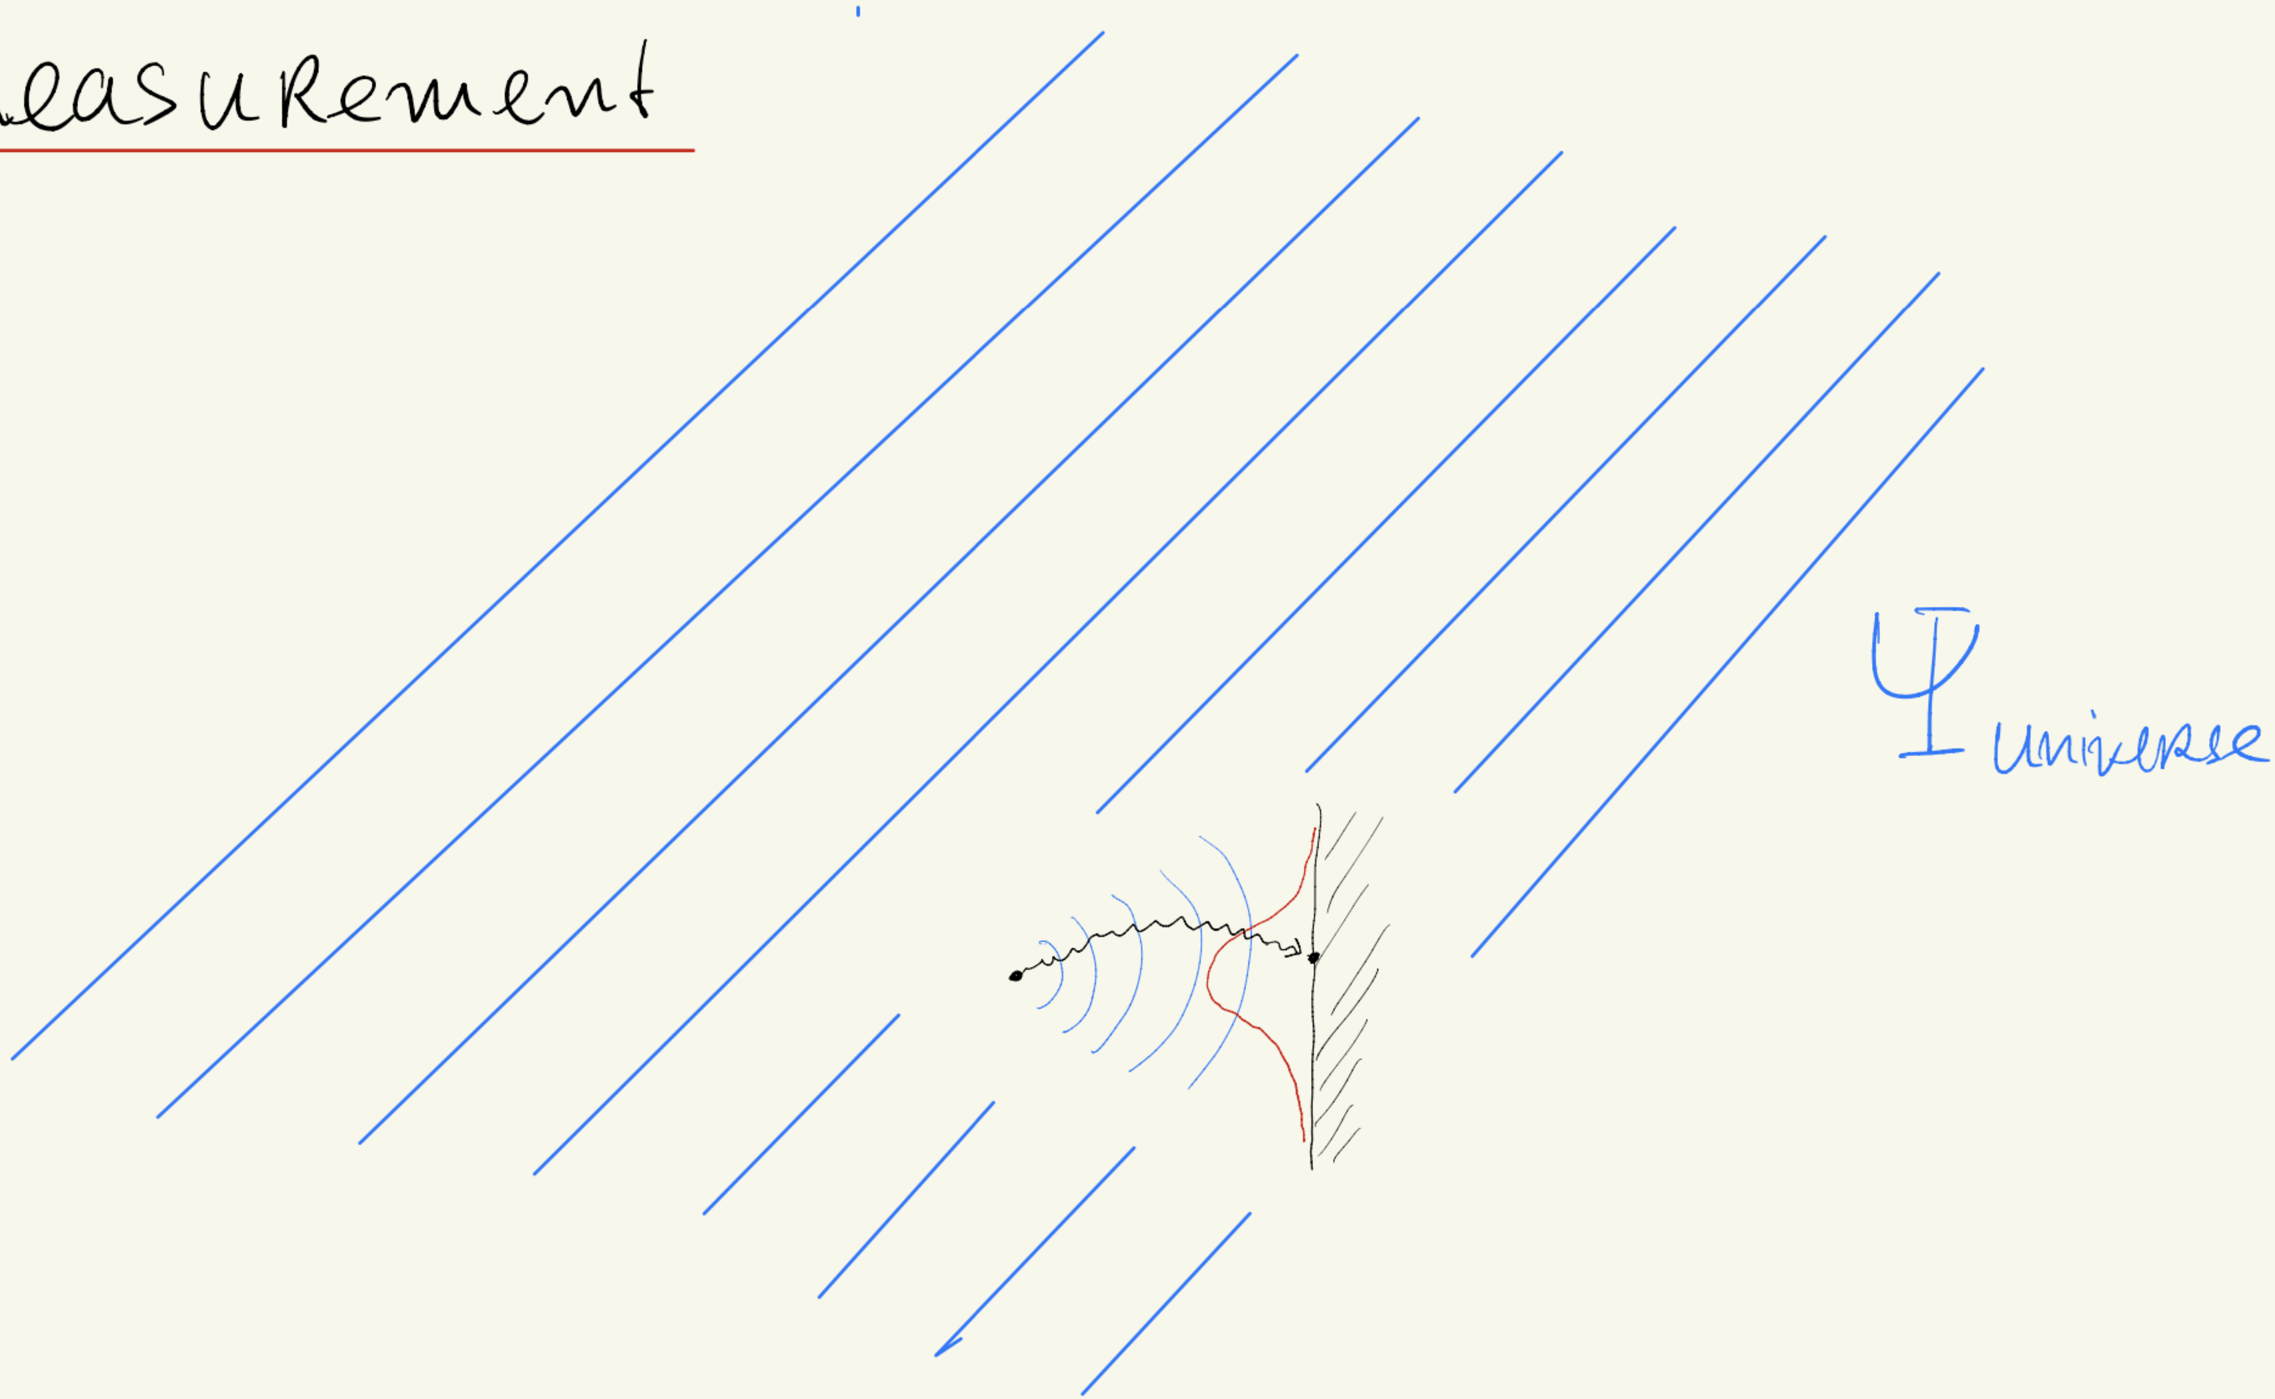

Density matrix

$$\frac{d}{dt}|\psi\rangle = -iH|\psi\rangle$$

$$\rho = |\psi\rangle\langle\psi|$$

$$\frac{d}{dt}\rho = -i[H,\rho]$$

ρ is called a pure state

Mixed State

ω_1, b_1 $\qquad\qquad$ ω_2, b_2

$$|\psi\rangle = |\omega_1 ; \omega_2\rangle + |b_1 ; b_2\rangle$$

$$\rho = |\psi\rangle\langle\psi|$$

Total system is in a pure state ρ

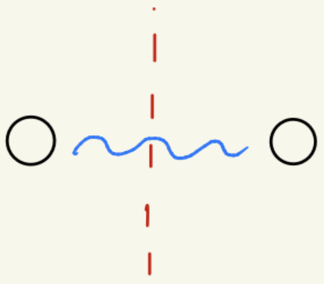

$$\rho_1 = \text{"average"}(\rho) = \frac{1}{2}|\omega_1\rangle\langle\omega_1| + \frac{1}{2}|b_1\rangle\langle b_1|$$

Subsystem is in a mixed state ρ_1

Pure v.s mixed state

pure state

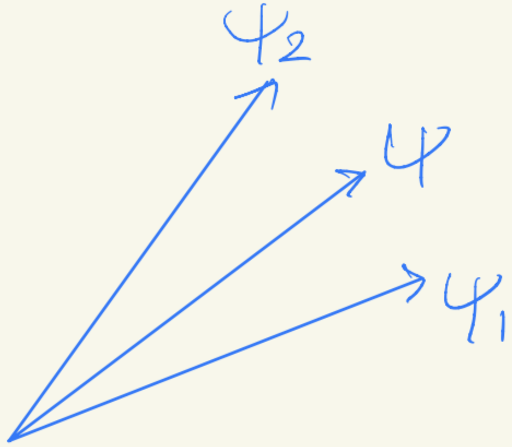

$$\psi = \frac{1}{\sqrt{2}}(\psi_1 + \psi_2)$$

$$\rho = \psi \psi^{\dagger}$$

mixed state

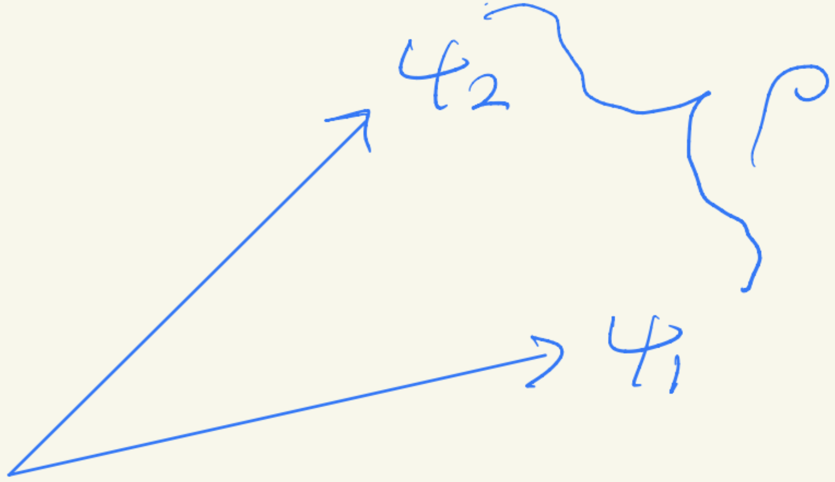

$$\rho = \frac{1}{2}\psi_1 \psi_1^{\dagger} + \frac{1}{2}\psi_2 \psi_2^{\dagger}$$

The Lindblad equation

When coupling between system and environment is weak and environment is sufficiently large[1]:

$$\dot{\rho} = \underbrace{-i[H,\rho]}_{\text{unitary part}} + \underbrace{D_{kl}\left(C_k\rho C_l^\dagger - \frac{1}{2}\{C_l^\dagger C_k, \rho\}\right)}_{\text{coupling to environment}}$$

Examples of Lindblad operators are measurement operators in which case C_k is Hermitian, or dissipation operators such as $\sigma_\pm$ for a single qubit.

Unravelings

Define the stochastic Schrödinger equation

$$d\psi \quad = \quad -iH\psi dt - \frac{1}{2}D_{kl}\left(C_l^\dagger C_k - 2c_k C_l + c_k c_l\right)\psi dt + (C_k - c_k)\psi d\xi_k$$

$d\xi_k$ is a real-valued Wiener process with $\langle d\xi_k \rangle = 0$ and $\langle d\xi_k d\xi_l \rangle = D_{kl} dt$.

$c_k = \psi^\dagger C_k^{(h)} \psi$ real-valued make SSE non-linear and $d||\psi||^2 = 0$.

This defines an unraveling of the Lindblad equation, meaning that the expectation $\rho = \left\langle \psi\psi^\dagger \right\rangle$ satisfies the Lindblad equation.

$$\dot{\rho} = -i[H, \rho] + D_{kl}\left(C_k \rho C_l^\dagger - \frac{1}{2}\{C_l^\dagger C_k, \rho\}\right)$$

Measuring a single qubit

Consider a single qubit with $H = 0$ and $C = \hat{\sigma}_z$ and $D = 1$.

The Lindblad equation is

$$\dot{\rho} = \hat{\sigma}_z \rho \hat{\sigma}_z - \rho$$

The SSE is

$$\begin{aligned} d\psi &= -\frac{1}{2}(\hat{\sigma}_z - m_z)^2 \psi dt + (\hat{\sigma}_z - m_z)\psi d\xi \\ dm_z &= 2(1 - m_z^2) d\xi \end{aligned}$$

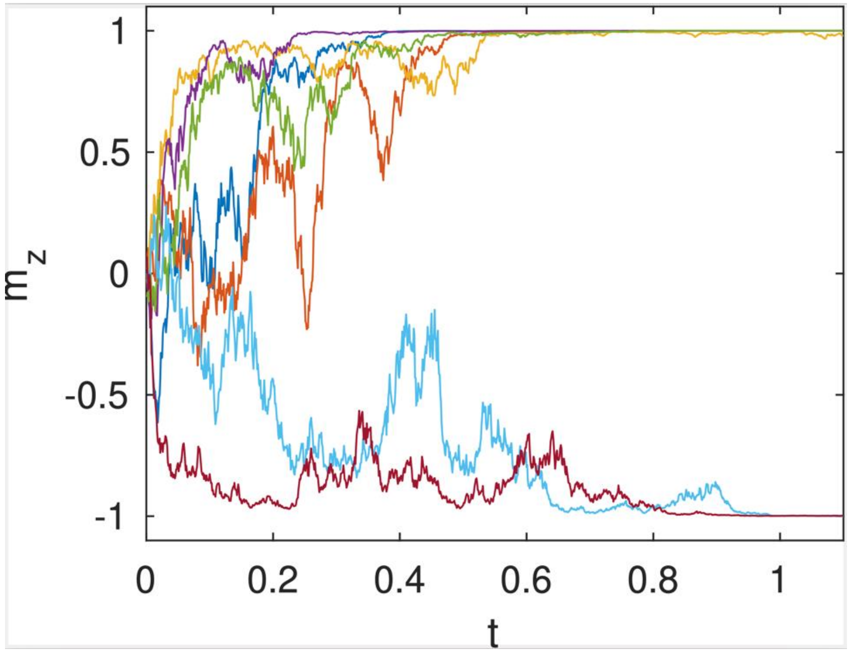

with $m_z = \psi^\dagger \hat{\sigma}_z \psi$ and $\left\langle d\xi^2 \right\rangle = dt$

Note, that $\hat{\sigma}_z \to i\hat{\sigma}_z$ leaves Lindblad equation invariant but makes the unraveling linear

$$d\psi = -\frac{1}{2}\psi dt + i\hat{\sigma}_z \psi d\xi \qquad dm_z = 0$$

Optimal control theory

Given a current state and a future desired state, what is the best/cheapest/fastest way to get there.

Stochastic optimal control

Consider a stochastic dynamical system

$$dx = f(t, x, u)dt + g(t, x, u)d\xi$$

$d\xi$ Gaussian noise $\langle d\xi d\xi' \rangle = \nu dt$.

The cost is an expectation:

$$C_u(t, x) \quad = \quad \left\langle \phi(x(T)) + \int_t^T d\tau V(t, x(t), u(t, x)) \right\rangle_u$$

over all stochastic trajectories with control function u starting at t, x.

The optimal cost to go: $J(t, x) = \min_u C_u(t, x)$ satisfies the Bellman equation

$$-\partial_t J = \min_u \left(V + f \nabla J + \frac{1}{2} \nu \mathrm{Tr}(g \nu g' \nabla^2) J \right)$$

is a PDE with boundary condition $J(\cdot, T) = \phi(\cdot)$.

Delayed choice

Compute the optimal control in x, t when

$$dx = udt + d\xi \qquad C_u(t, x) = \left\langle \phi(x_T) + \int_t^T ds \frac{1}{2} u^2 \right\rangle_u$$

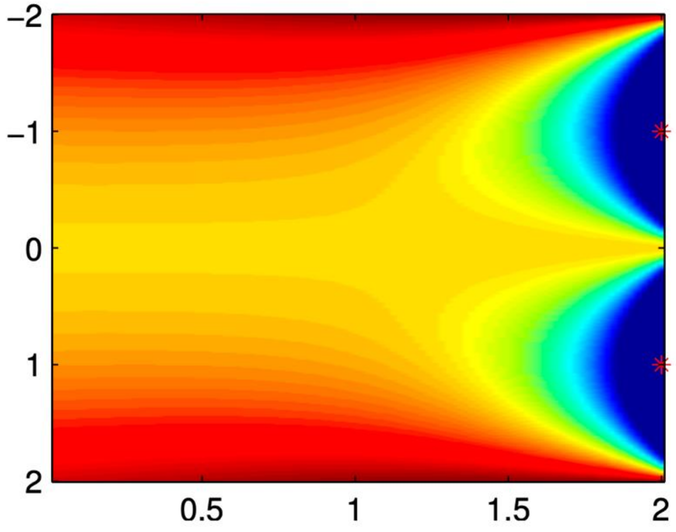

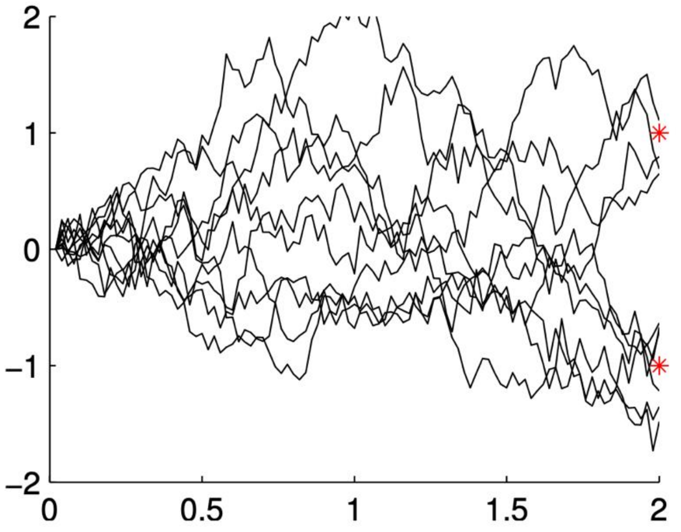

Decision is made at $t = T - \frac{1}{\nu}$.

Path integral (PI) control methods

$$
\begin{aligned}
dx &= f(t,x)dt + g(t,x)(u(t,x)dt + d\xi) \qquad \langle d\xi \rangle = \nu dt \\
C_u(t,x) &= \langle S_u(\tau) \rangle_u \qquad S_u(\tau) = \phi(x(T)) + \int_t^T ds V(t,x) + \frac{1}{2} u' R u + u' R d\xi
\end{aligned}
$$

with $\tau = x_{t:T}$ a trajectory starting at t, x.

When $R = \lambda \nu^{-1}$ the Bellman equation can be linearized by a log transform. The solution is given as a path integral

$$
\begin{aligned}
J(t,x) &= -\lambda \log \left\langle e^{-S_u(\tau)/\lambda} \right\rangle_u \\
u^*(t,x) &= u(t,x) + \lim_{dt \downarrow 0} \frac{1}{dt} \frac{\left\langle d\xi_t e^{-S_u(\tau)/\lambda} \right\rangle}{\left\langle e^{-S_u(\tau)/\lambda} \right\rangle}
\end{aligned}
$$

The optimal control solution can be obtained by sampling [Kappen, 2005, Thijssen and Kappen, 2015].

Delayed choice

Compute the optimal control in x, t when

$$dx = udt + d\xi \qquad C_u(t,x) = \left\langle \phi(x_T) + \int_t^T ds \frac{1}{2} u^2 \right\rangle_u$$

Solution

$$J(t,x) = -\nu \log \left\langle e^{-\phi(x_T)/\nu} \right\rangle_{u=0} \qquad u(t,x) = -\nabla J(t,x)$$

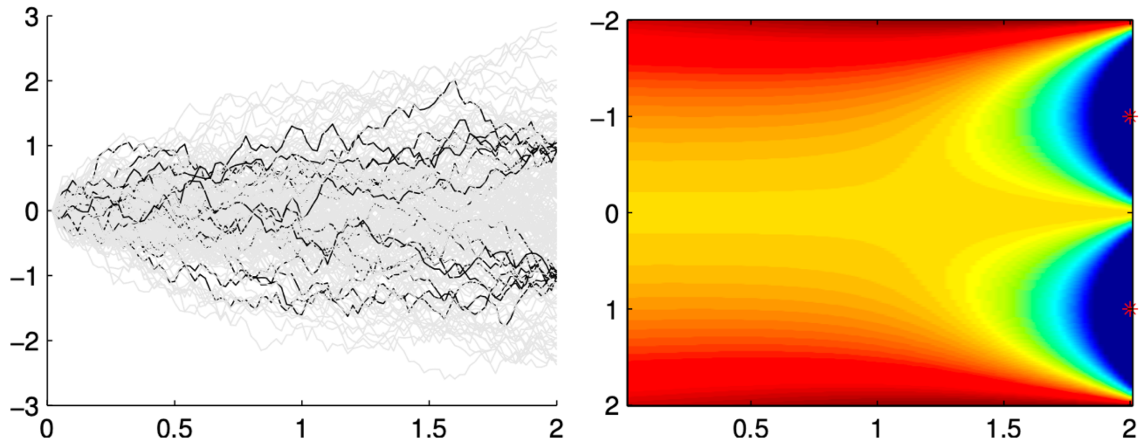

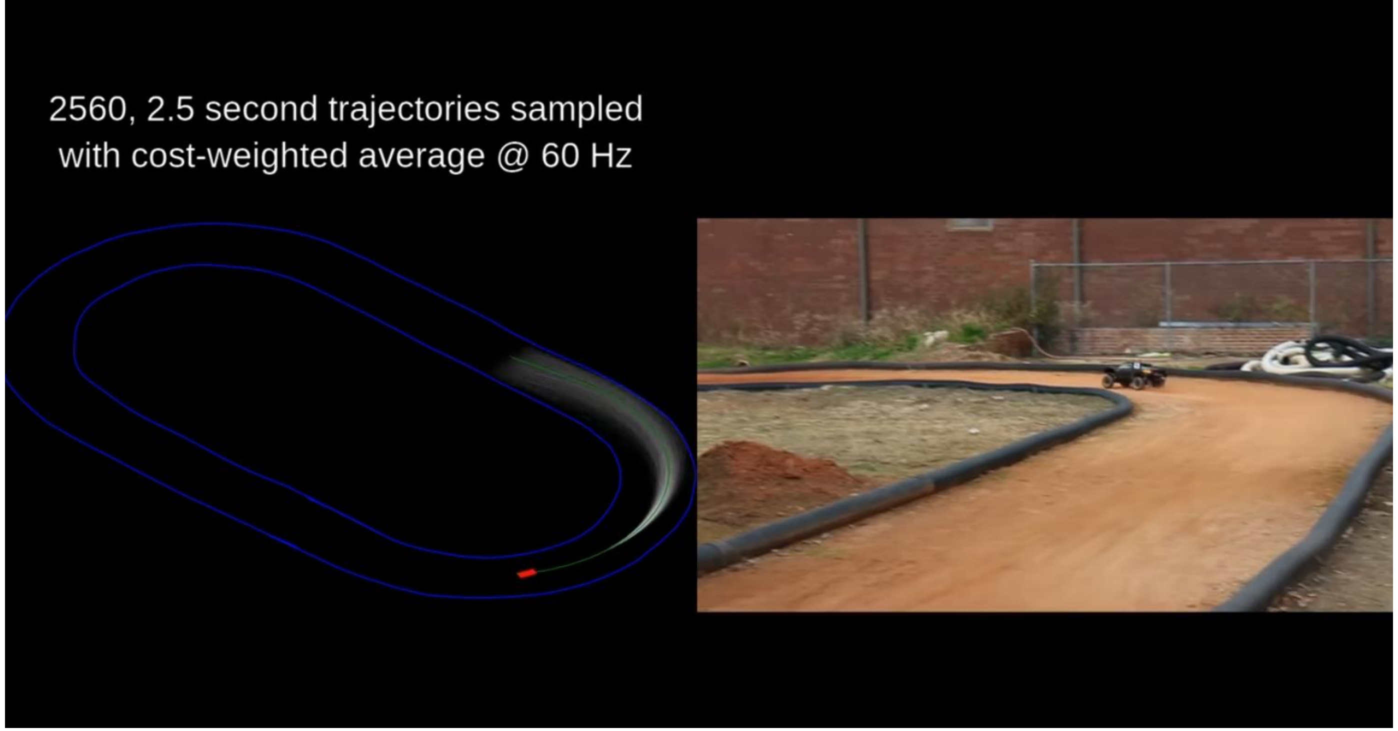

Non-linear stochastic optimal control solution computed in real time (Williams et al. 2016)

PI needs importance sampling

$$J(x,t) = -\lambda \log \left\langle e^{-S_u(\tau)/\lambda} \right\rangle_u \qquad w_i = \frac{e^{-S_u(\tau_i)/\lambda}}{\sum_{j=1}^{N} e^{-S_u(\tau_j)/\lambda}} \qquad N_{ESS} = \left(\sum_{i=1}^{N} w_i^2 \right)^{-1}$$

All u are unbiased estimators, but with some are more effective (smaller variance):
- Better u (smaller C) is better sampler (smaller variance, higher N_{ESS})
- Optimal u (minimal C) is optimal sampler (zero variance, $N_{ESS} = N$)

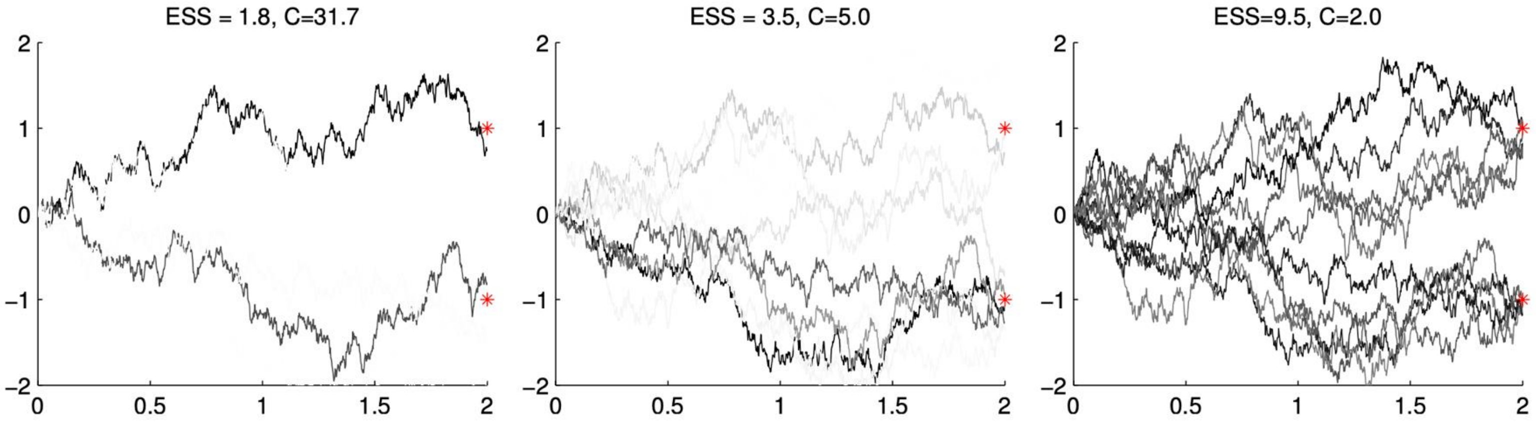

Thijssen Kappen 2015

Adaptive importance sampling (PICE)

With parametrized control $u(x, t|\theta)$, adaptive importance sampling improves the control solution $\theta^p \to \theta^{p+1}$, starting with some inital control solution θ^0.

$$\theta^{p+1} = \theta^p + \eta \frac{\left\langle e^{-S^u(t_0)/\lambda} \int_0^T \frac{\partial u(X_t^u, t|\theta^p)}{\partial \theta} d\xi_t \right\rangle_{\theta^p}}{\left\langle e^{-S^u(t_0)/\lambda} \right\rangle_{\theta^p}}$$

Stochastic integral

Importance weights

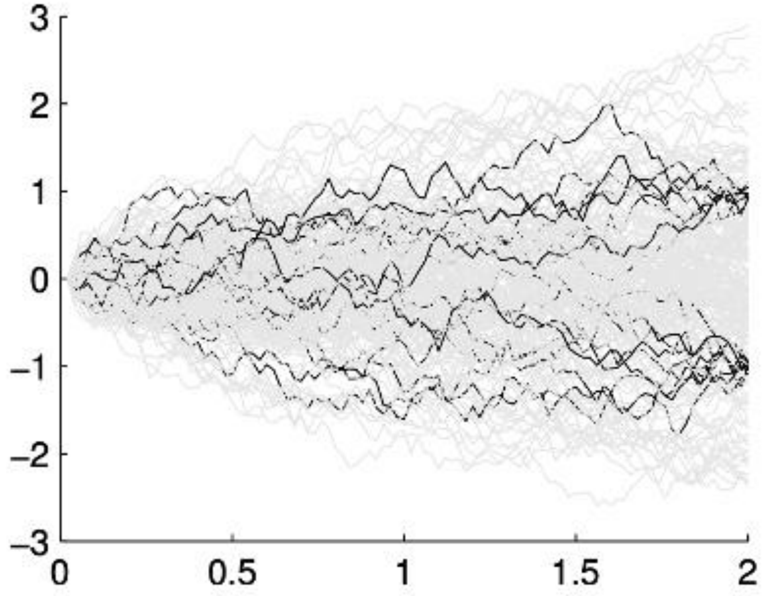

Kappen et al. 2015

Control of the Lindblad equation

Define $H = H_0 + u_k H_k$. The (deterministic) control problem is

$$\begin{aligned}\dot{\rho}_t &= -i[H,\rho_t] + D_{kl}\left(C_k\rho_t C_l^\dagger - \frac{1}{2}\{C_l^\dagger C_k, \rho_t\}\right)\\ C(\rho_0,u) &= \mathrm{Tr}(G\rho_T) + \int_0^T \frac{1}{2}u_t' R u_t\end{aligned}$$

with G some targer Hamiltonian.

When the C_k can be transformed into anti-Hermitian operators $\tilde{C}_k = -iH_k$ and $\tilde{D} > 0$, the stochastic optimal control problem is of the path integral form:[3]

$$\begin{aligned}d\psi_t &= -iH_0\psi_t dt - \frac{1}{2}\tilde{D}_{kl}H_l H_k \psi_t dt - i(u_k dt + d\tilde{\xi}_k)H_k\psi_t\\ C(\psi_0,u) &= \left\langle \psi_T^\dagger G\psi_T + \frac{1}{2}\int_0^T u_t' R u_t dt\right\rangle_u\end{aligned}$$

[3] In the case of feedback control, the control problem depends on the unraveling.

Control of one qubit

Unitary control $H = u_x\hat{\sigma}_x + u_u\hat{\sigma}_y$ and dissipation $\hat{\sigma}_+$ and $\hat{\sigma}_-$ has Lindblad equation

$$\dot{\rho} = \underbrace{-i[H,\rho]}_{\text{control}} + \underbrace{D(\hat{\sigma}_+\rho\hat{\sigma}_- + \hat{\sigma}_-\rho\hat{\sigma}_+ - \rho)}_{\text{dissipation}}$$

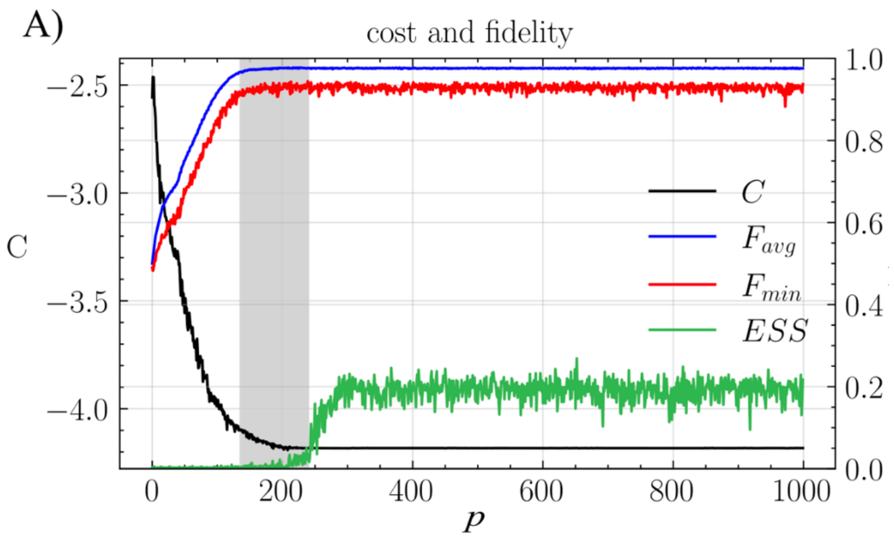

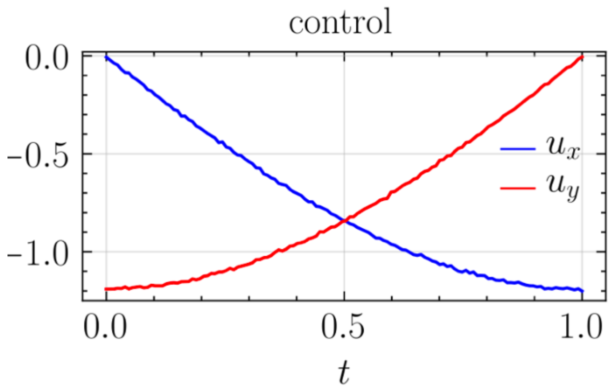

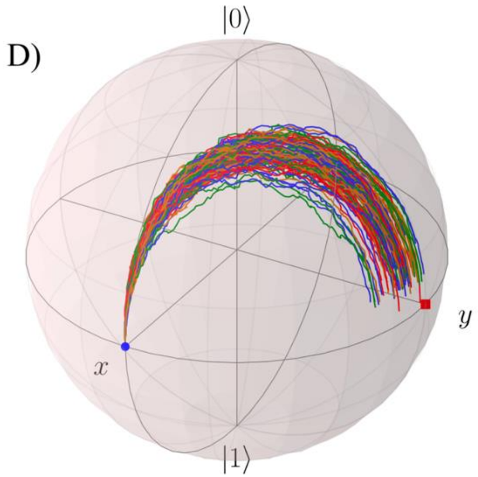

K=128, n_traj= 400. mean asymptotic fidelity = 0.9759

QDC versus Open GRAPE

(Open) GRAPE (Boutin 2017) solves the deterministic control problem using the Lindblad equation.
It requires large number of pulses for accurate approximation of gradient

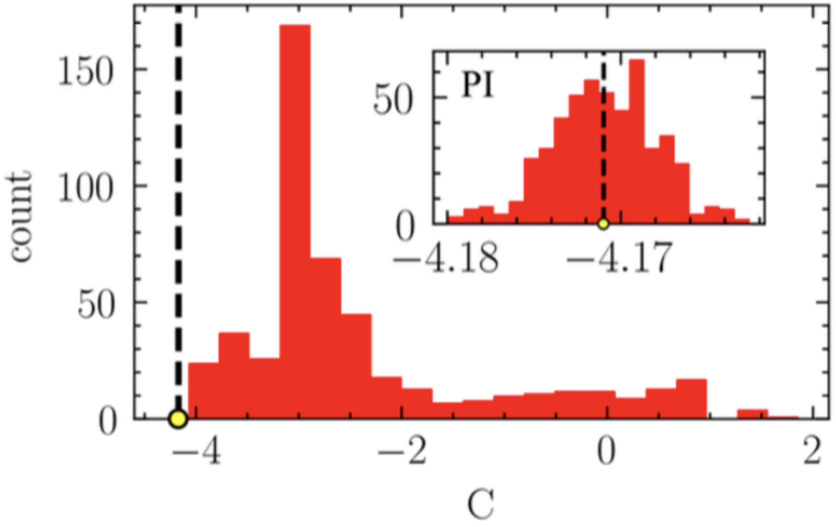

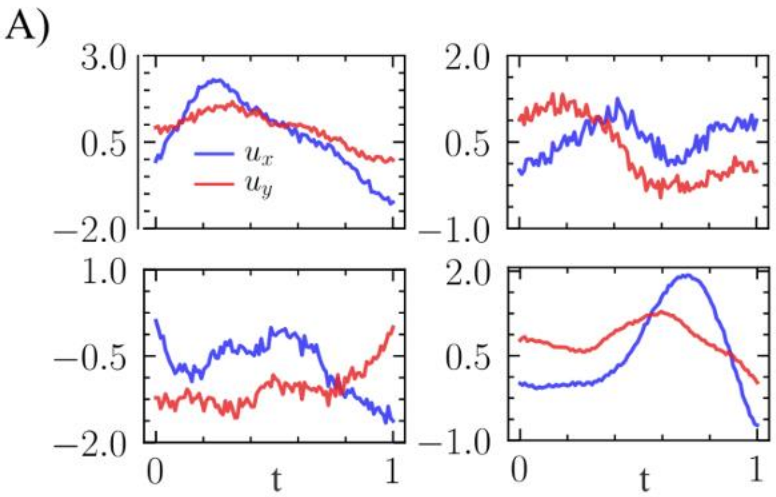

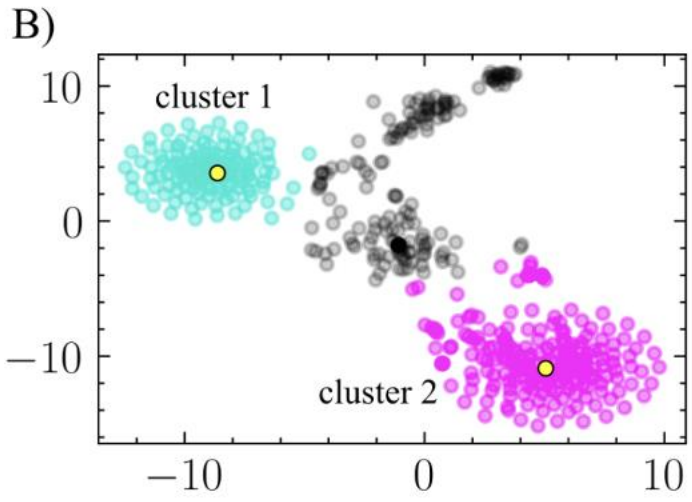

Histogram of 505 different initializations. Some Open GRAPE solutions. Clustering of solutions

The sweet spot: Stochastic optimal control solution as a proxy for deterministic optimal control solution

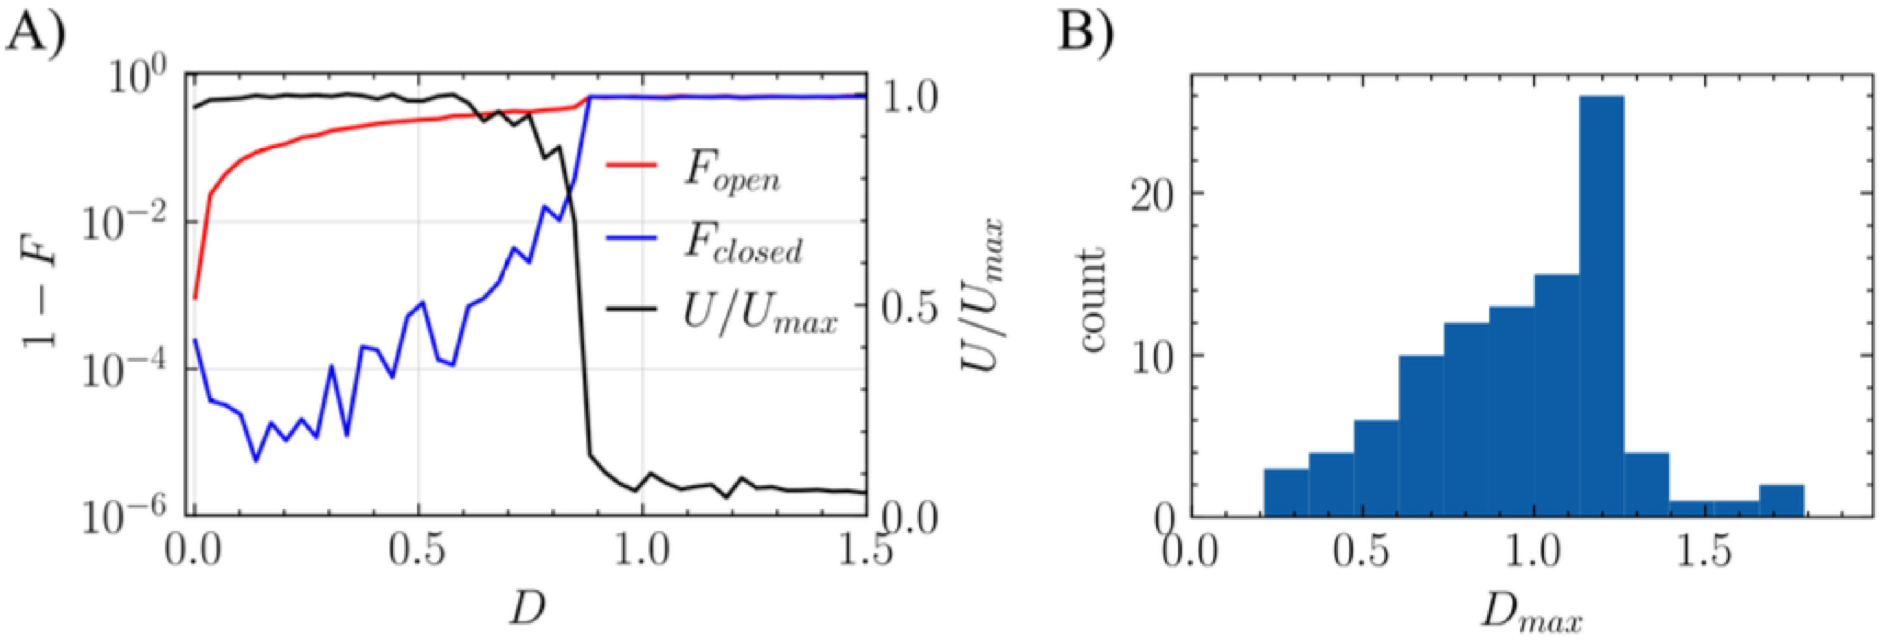

Left: X -> Y. Reaches F>0.98 for D< 0.8.
Right: X -> Haar. Reaches F > 0.98 for all instances, often with large noise

NMR physics preparation of n=4 GHZ state

Use nuclear spins of small molecules as coupled qubits.

Carbon-13-iodotrifluroethylene

$$H_{\mathrm{NMR}} = \sum_j \left(\pi \nu_j \hat{\sigma}_z^{(j)}\right) + \sum_{i<j} \left(\frac{\pi}{2} J_{ij} \hat{\sigma}_z^{(i)} \hat{\sigma}_z^{(j)}\right),$$

$$H_{\mathrm{c}}(t) = \sum_j \left(\pi u_{xj}(t) \hat{\sigma}_x^{(j)} + \pi u_{yj}(t) \hat{\sigma}_y^{(j)}\right).$$

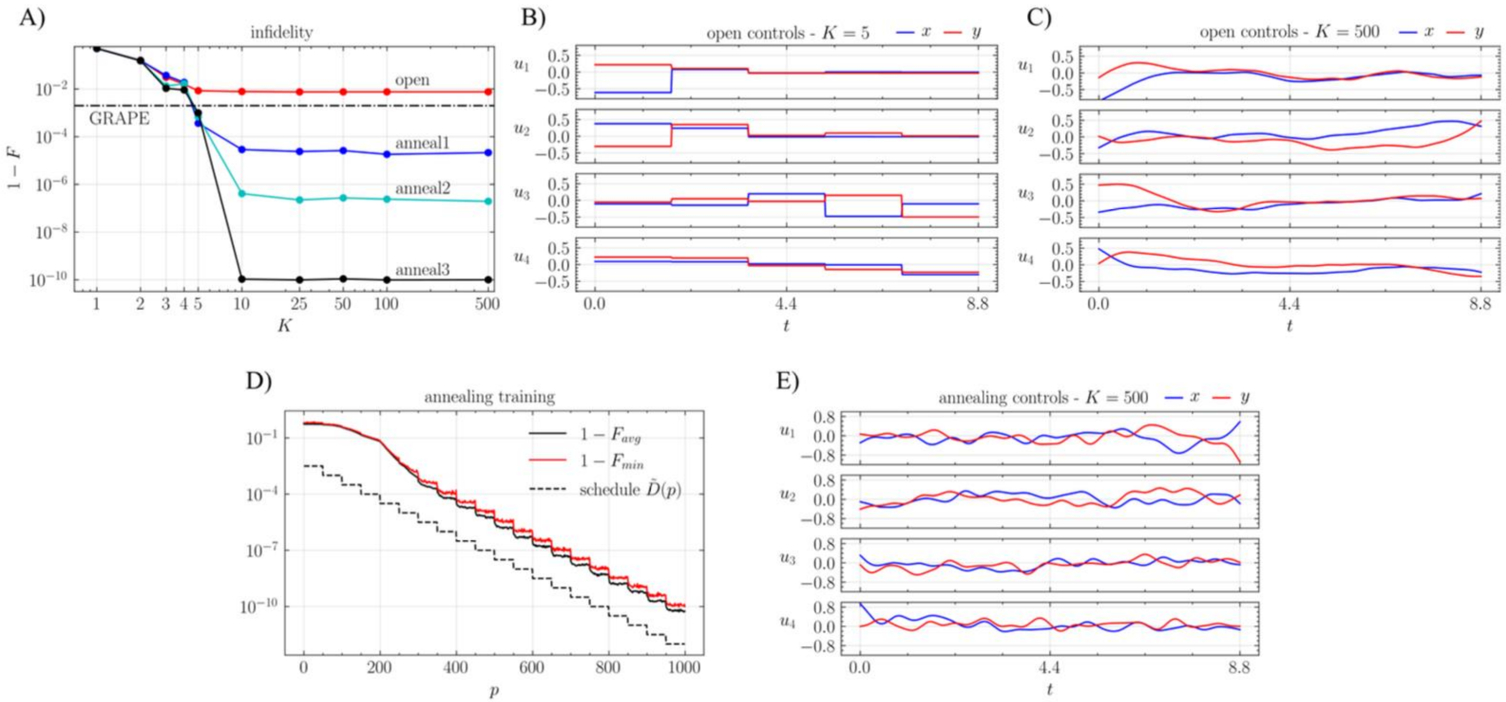

GRAPE needs large number of pulses for accuracy (1760) and Fidelity 0.998 (Chen et al. 2023).
QDC has no such requirement (32). Annealing yields infidelity 1e-10.

Scaling up: classical computing

$$u_k^{(p+1)} = u_k^{(p)} + \sum_{i=1}^{N} w_i \int_{\tau_k}^{\tau_{k+1}} d\xi_t^i$$

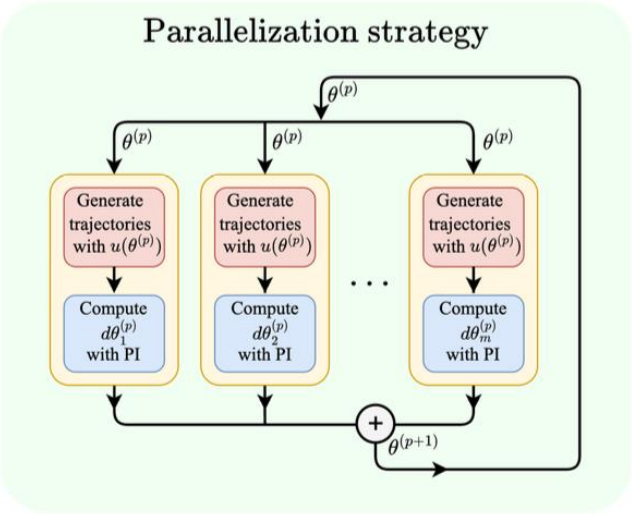

Parallelization of importance sampling is very efficient

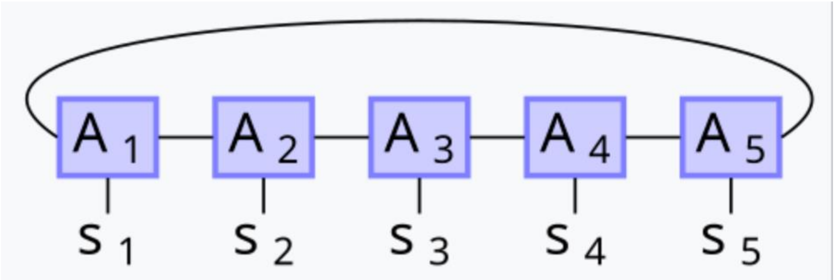

Tensor networks for efficient representation of large wave functions

Scaling up: quantum computing

- Scalability requires that the parameter optimization is executed on a quantum device.
  - This holds for VQE for digital circuits
  - This does not hold for optimal control methods for analog quantum devices (Grap, Crab)
  - QDC can be optimized on quantum hardware

Scaling up: quantum computing

$$u_k^{(p+1)} = u_k^{(p)} + \sum_{i=1}^{N} w_i \int_{\tau_k}^{\tau_{k+1}} d\xi_t^i$$

- QDC on quantum device:
  - Weights depend on noise trajectory and end cost
  - End cost obtainable from measurement
  - The noise realization is unknown
- Approach
  - Assume noiseless device
  - provide the noise externally
  - Exploit sweet spot idea

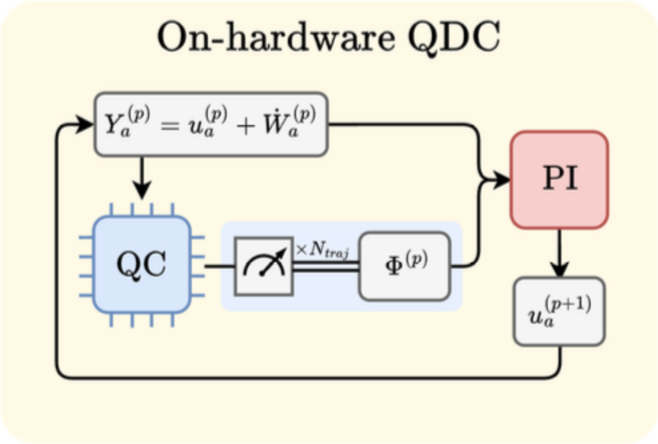

Summary

- Analog quantum computing is promising alternative for digital quantum circuits
  - Shorter circuit times
- Open quantum system approach is promising to mitigate barren plateaus

- New control framework for open analog quantum systems based on unravelings and path integral control.
  - Provides quadratic speed-up
  - Exploits quantum-noise sweet spot
  - Outperforms deterministic control methods
  - Scalable on quantum computers

Future work

- Application on analog quantum devices
- Scale up quantum simulation
    - Tensor networks
    - Parallel hardware, GPUs

- Benchmark studies comparing analog and digital circuits (quantum chemistry, QUBO)
- (Optimization of) quantum annealing

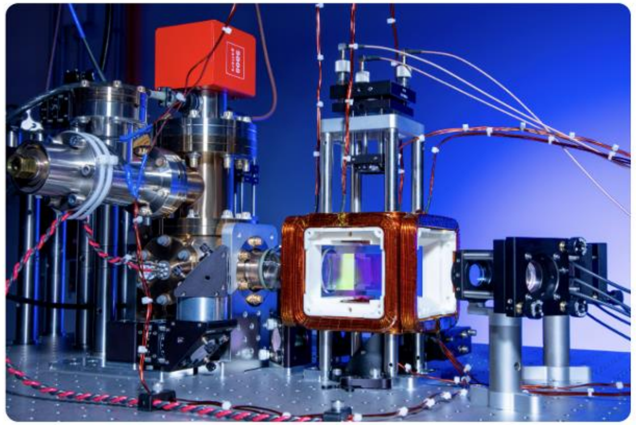

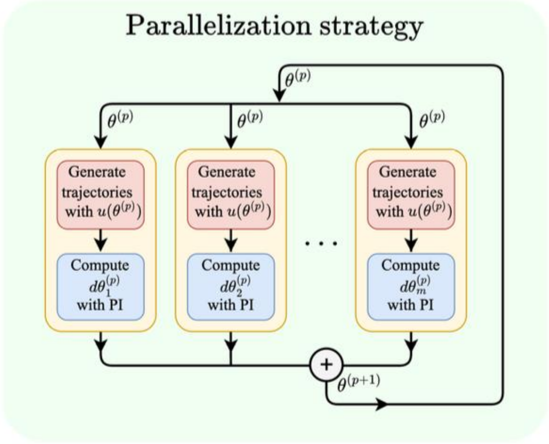

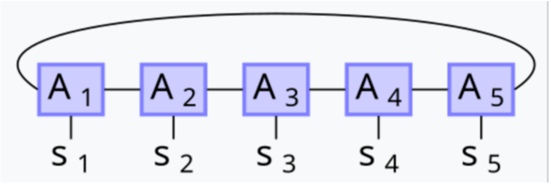

References

[Diósi, 2023] Diósi, L. (2023). Hybrid completely positive Markovian quantum-classical dynamics. *Physical Review A*, 107(6):062206.

[Kappen, 2005] Kappen, H. (2005). Linear theory for control of non-linear stochastic systems. *Physical Review letters*, 95:200201.

[Kappen and Ruiz, 2016] Kappen, H. and Ruiz, H. (2016). Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, pages 10.1007/s10955–016–1446–7.

[Kappen et al., 2012] Kappen, H. J., Gómez, V., and Opper, M. (2012). Optimal control as a graphical model inference problem. *Machine learning*, 87:159–182.

[Layton et al., 2023] Layton, I., Oppenheim, J., and Weller-Davies, Z. (2023). A healthier semi-classical dynamics. arXiv:2208.11722 [gr-qc, physics:hep-th, physics:quant-ph].

[Meitei et al., 2021] Meitei, O. R., Gard, B. T., Barron, G. S., Pappas, D. P., Economou, S. E., Barnes, E., and Mayhall, N. J. (2021). Gate-free state preparation for fast variational quantum eigensolver simulations. *npj Quantum Information*, 7(1):155.

[Thijssen and Kappen, 2015] Thijssen, S. and Kappen, H. J. (2015). Path integral control and state-dependent feedback. *Phys. Rev. E*, 91:032104. http://arxiv.org/abs/1406.4026.

[Villanueva and Kappen, 2024] Villanueva, A. and Kappen, H. J. (2024). Stochastic optimal control of open quantum systems. arxiv.org/abs/2410.18635.

Peyman Najafi

Aaron Villanueva

Eduardo Dominguez