Policies for elementary link generation in quantum networks

Sumeet Khatri

Department of Physics & Astronomy, Louisiana State University

arXiv:2007.03193

October 7, 2020

Quantum networks

- Nodes are sending/receiving stations.
- Elementary links represent:
  - Entanglement → undirected graph;
  - Quantum channels → directed graph.
- Link creation is probabilistic.
- The graph is dynamic: links disappear and reappear over time.
- Main questions: end-to-end rates, end-to-end fidelities, how to find paths in the network (i.e., routing) for multi-user requests.

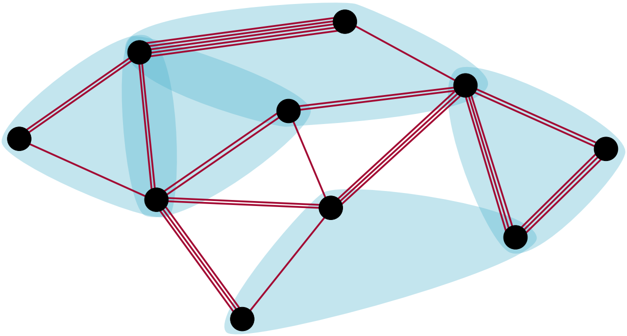

Red lines indicate bipartite entanglement; blue bubbles indicate multipartite entanglement.

Quantum networks

- Results from classical networking don't always directly carry over to quantum networks.

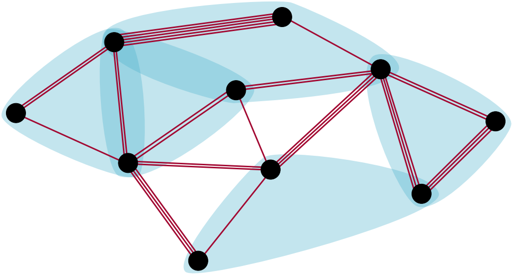

Quantum networks

- Results from classical networking don't always directly carry over to quantum networks.
- Elements of quantum network protocols:
  - Creating elementary links
  - Entanglement distillation/error correction
  - Entanglement swapping/joining measurements
  - Routing

  Need some way to "stitch" together these elements via a sequence of actions over time.

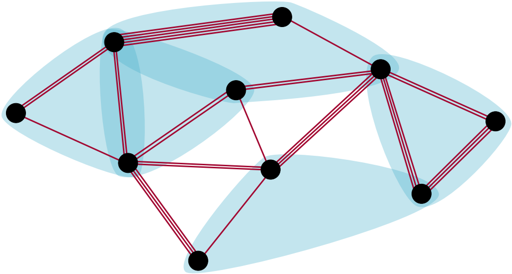

Application	
Transport	Qubit transmission
Network	Long distance entanglement
Link	Robust entanglement generation
Physical	Attempt entanglement generation

[arXiv:1903.09778]

Quantum networks

- Results from classical networking don't always directly carry over to quantum networks.
- Elements of quantum network protocols:
  - Creating elementary links
  - Entanglement distillation/error correction
  - Entanglement swapping/joining measurements
  - Routing

  Need some way to "stitch" together these elements via a sequence of actions over time.
- One route: **decision processes**; related to
  - Rule-sets
  - Finite state machines
  - Scheduling

  Provides a route to performing **reinforcement learning** in a quantum network.

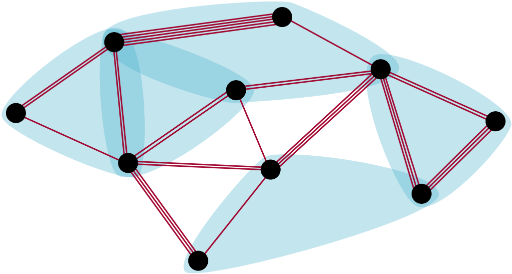

Application	
Transport	Qubit transmission
Network	Long distance entanglement
Link	Robust entanglement generation
Physical	Attempt entanglement generation

[arXiv:1903.09778]

This work

- Cast elementary link generation formally as a quantum decision process.
  - Average quantum state of the link at any time.
  - Fidelity of the link and link activity probability as a function of time.
- Look at the memory cutoff policy in the finite-horizon and infinite-horizon cases.
- Policy optimization in the finite-horizon case.

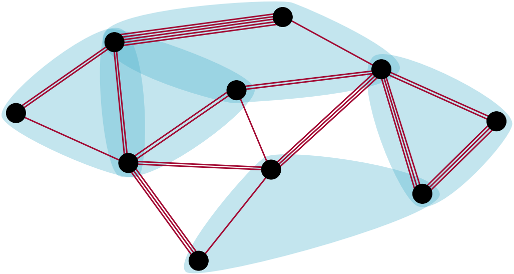

Application	
Transport	Qubit transmission
Network	Long distance entanglement
Link	Robust entanglement generation
Physical	Attempt entanglement generation

[arXiv:1903.09778]

This work

★ We focus on elementary links only as a first step to illustrate the method. Extensions to higher levels of the stack is a potential direction for future work.

★ Prior work on network protocols:

- Aparicio et al., AINTEC 2011; Rod Van Meter & Joe Touch, IEEE 2013.
- Pirker et al.: NJP 20 053054 (2018), NJP 21 033003 (2019).
- Dahlberg et al., arXiv:1903.09778.

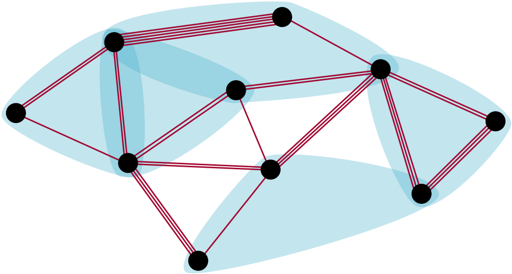

Application	
Transport	Qubit transmission
Network	Long distance entanglement
Link	Robust entanglement generation
Physical	Attempt entanglement generation

[arXiv:1903.09778]

Decision processes

Classical decision process (see textbook by Martin Puterman)

- Discrete sets $\mathcal{X}$ (states) and $\mathcal{A}$ (actions).
- Transition function $T(x_t, a_t, x_{t+1})$ governs state transition (can be probabilistic).
- *Decisions*: mappings d_t from histories to actions (can be probabilistic).

★ Goal is to learn the *policy*:

$$\pi = (d_1, d_2, \dots)$$

that maximizes expected reward.

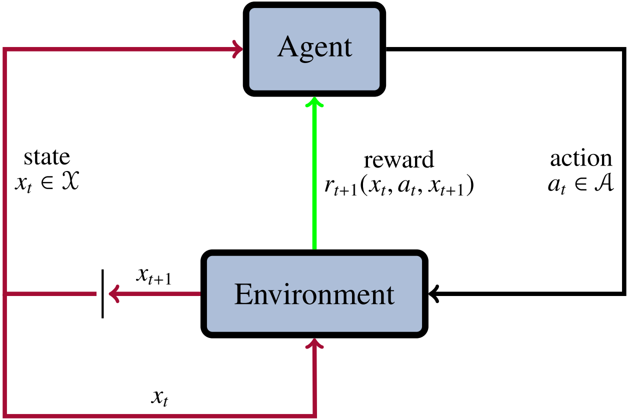

Histories are $h^t = (x_1, a_1, x_2, a_2, \dots, a_{t-1}, x_t)$. Then $d_t(h^t)$ is an action or probability distribution over actions.

Decision processes

Quantum decision process (as defined in [PRA 90, 032311 (2014)])

- ★ Agent classical, but environment quantum.
- ★ $\mathcal{X}$ contains *classical information* about the quantum state of the environment.
- ★ Transition functions are completely positive trace non-increasing maps:

$$\mathcal{T}^{x_t, a_t, x_{t+1}}$$

with $\sum_{x_{t+1}} \mathcal{T}^{x_t, a_t, x_{t+1}}$ being trace preserving.

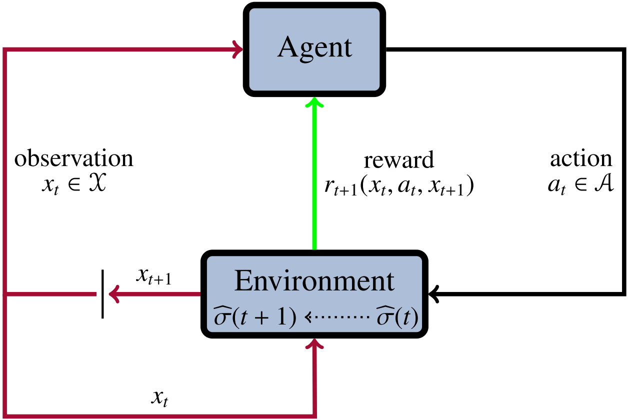

Elementary link generation in a quantum network

- Source produces state ρ^S
- Source state transmitted to nodes: $\rho^S_{\text{out}} := \mathcal{L}(\rho^S)$
- Nodes perform heralding procedure:
  - Success: $\rho^S_{\text{out}} \mapsto \widetilde{\rho}_0 := \mathcal{M}_1(\rho^S_{\text{out}})$
  - Failure: $\rho^S_{\text{out}} \mapsto \widetilde{\tau}^{\varnothing} := \mathcal{M}_0(\rho^S_{\text{out}})$.

★ This process occurs independently for all elementary links.

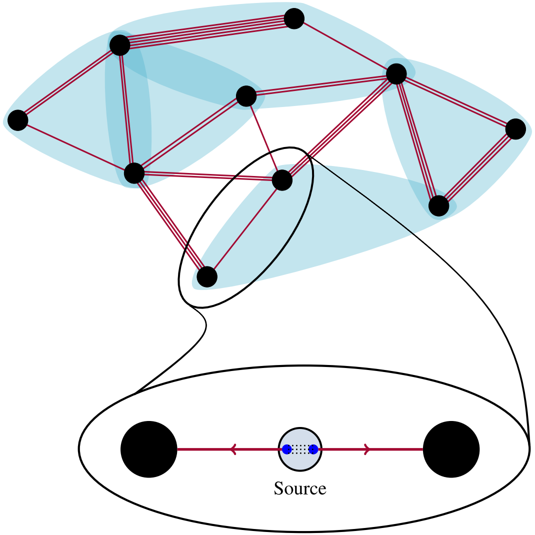

Elementary link generation in a quantum network

- After heralding success: $\widetilde{\rho}_0$
- After heralding failure: $\widetilde{\tau}^{\varnothing}$

Heralding success probability $p \coloneqq \mathrm{Tr}[\widetilde{\rho}_0]$.

$$\rho_0 \coloneqq \frac{\widetilde{\rho}_0}{p}, \quad \tau^{\varnothing} \coloneqq \frac{\widetilde{\tau}^{\varnothing}}{1-p}.$$

After success, the state is held in memories at the nodes, which undergo decoherence processes given by quantum channel $\widehat{\mathcal{N}} = \mathcal{N}_1 \otimes \cdots \otimes \mathcal{N}_k$.

State after m time steps in memory:

$$\rho(m) \coloneqq \widehat{\mathcal{N}}^{\circ m}(\rho_0).$$

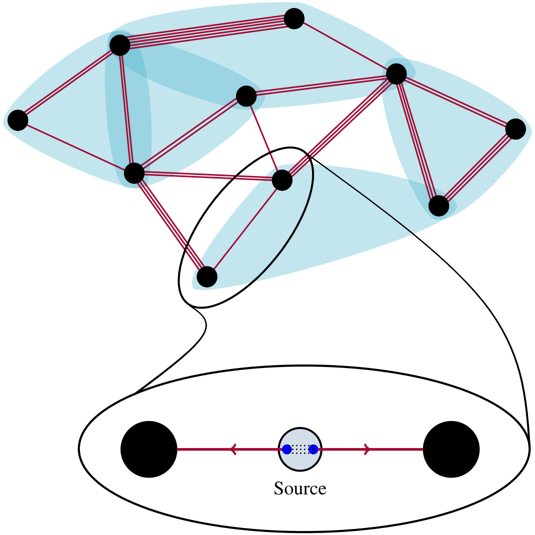

Link evolution as a decision process

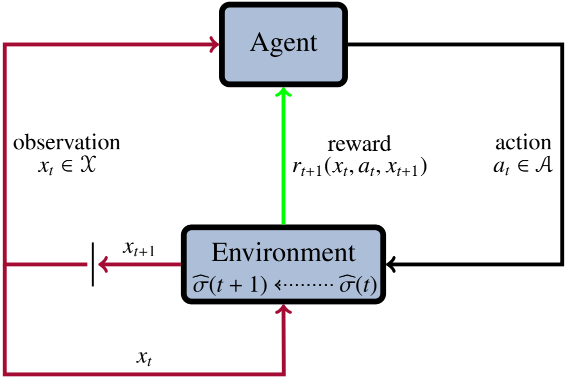

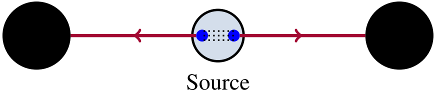

Link evolution as a decision process

Agents are the nodes.

Environment is quantum system from the source.

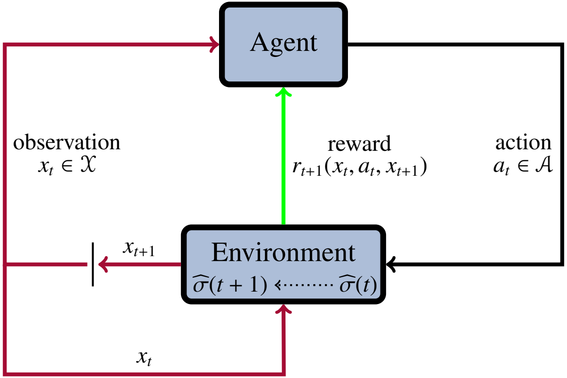

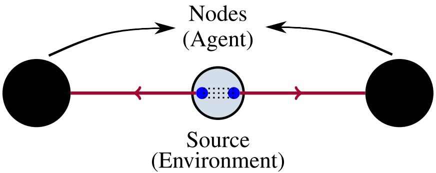

Link evolution as a decision process

Agents are the nodes.

Environment is quantum system from the source.

$\mathcal{X} = \{0, 1\}$ and $\mathcal{A} = \{0, 1\}$

- $A(t) = 1$: request from source
- $A(t) = 0$: wait/keep link currently in memory

- $X(t) = 1$: heralding succeeded/link established
- $X(t) = 0$: heralding failed/link not established

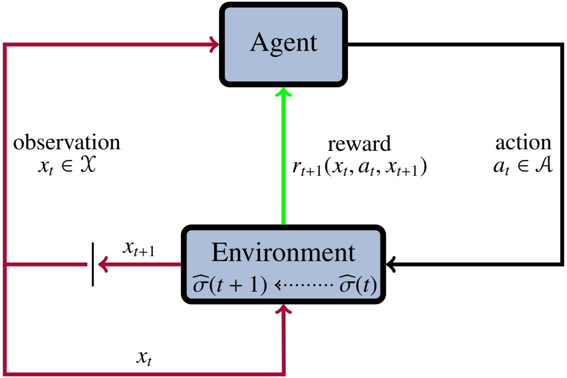

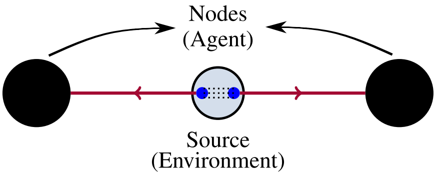

Link evolution as a decision process

Agents are the nodes.

Environment is quantum system from the source.

$\mathcal{X} = \{0, 1\}$ and $\mathcal{A} = \{0, 1\}$

- $A(t) = 1$: request from source
- $A(t) = 0$: wait/keep link currently in memory

- $X(t) = 1$: heralding succeeded/link established
- $X(t) = 0$: heralding failed/link not established

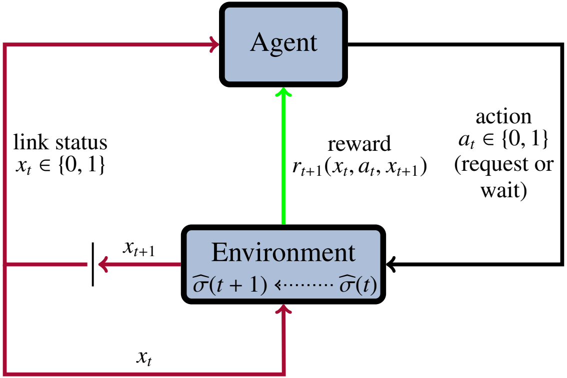

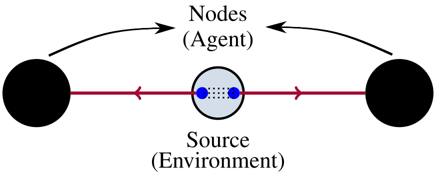

Link evolution as a decision process

Transition functions $\mathcal{T}^{x_t,a_t,x_{t+1}}$:

$$\mathcal{T}^{x_t,1,1}(\sigma) := \mathrm{Tr}[\sigma]\widetilde{\rho}_0 \quad \forall\ x_t \in \{0,1\},$$

$$\mathcal{T}^{x_t,1,0}(\sigma) := \mathrm{Tr}[\sigma]\widetilde{\tau}^{\varnothing} \quad \forall\ x_t \in \{0,1\},$$

$$\mathcal{T}^{1,0,1}(\sigma) := \widehat{\mathcal{N}}(\sigma) \quad \text{(decoherence)},$$

$$\mathcal{T}^{0,0,0}(\sigma) := \sigma.$$

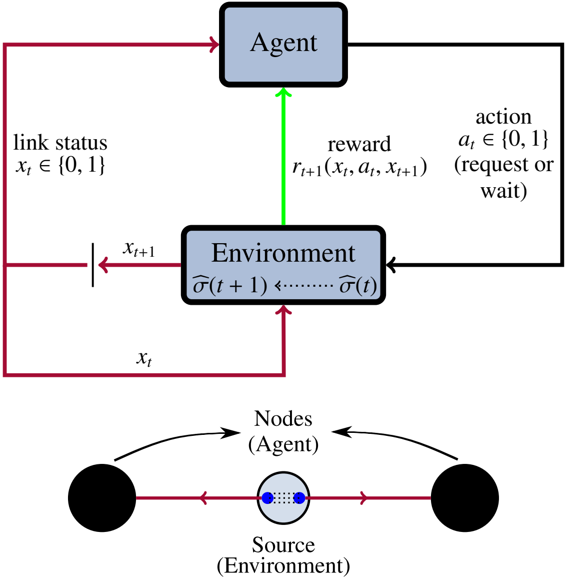

Link evolution as a decision process

Reward:

$$r_{t+1}(x_t, a_t, x_{t+1}) = \begin{cases} 0 & \text{if } t+1 < T \\ \delta_{x_{t+1},1} \langle \psi | \sigma(t+1|h^{t+1}) | \psi \rangle & \text{if } t+1 = T \end{cases}$$

Time T is the *horizon time* of the decision process (assumed finite).

Reward at the horizon time is the fidelity with respect to a target state (if link is active).

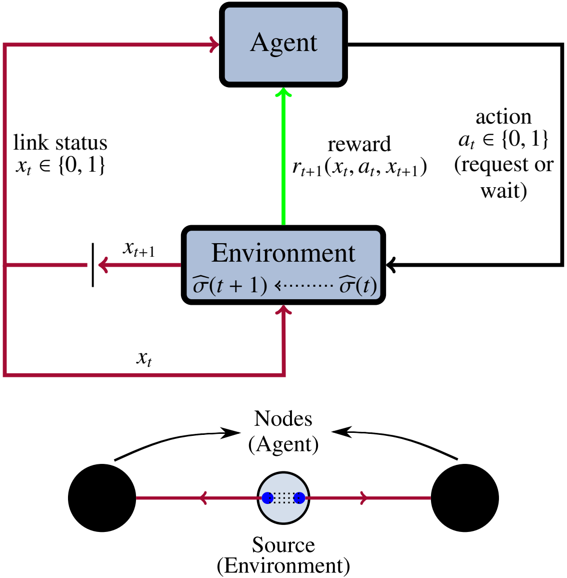

Average quantum state of a link

$$\widehat{\sigma}(1) = |0\rangle\langle 0| \otimes \widetilde{\tau}^{\varnothing} + |1\rangle\langle 1| \otimes \widetilde{\rho}_0,$$

$$\widehat{\sigma}(t) = \sum_{h^t} |h^t\rangle\langle h^t| \otimes \widetilde{\sigma}(t; h^t) \Rightarrow \sigma(t) = \sum_{h^t} \widetilde{\sigma}(t; h^t).$$

In general, for any $h^t = (x_1, a_1, x_2, a_2, \ldots, a_{t-1}, x_t)$:

$$\widetilde{\sigma}(t; h^t) = \left(\prod_{j=1}^{t-1} d_j(h_j^t)(a_j)\right)(\mathcal{T}^{x_{t-1}, a_{t-1}, x_t} \circ \cdots \circ \mathcal{T}^{x_1, a_1, x_2})(\widetilde{\sigma}(1; x_1)),$$

$$\widetilde{\sigma}(1; 1) := \widetilde{\rho}_0 \quad \text{(link active)},$$

$$\widetilde{\sigma}(1; 0) := \widetilde{\tau}^{\varnothing} \quad \text{(link inactive)}.$$

Conditional states:

$$\sigma(t|h^t) = \frac{\widetilde{\sigma}(t; h^t)}{\Pr[H(t) = h^t]}, \quad \Pr[H(t) = h^t] = \mathrm{Tr}[\widetilde{\sigma}(t; h^t)]$$

Average quantum state of a link

Recall transition functions:

$$\mathcal{T}^{x_t,1,1}(\sigma) := \mathrm{Tr}[\sigma]\widetilde{\rho}_0, \quad \mathcal{T}^{x_t,1,0}(\sigma) := \mathrm{Tr}[\sigma]\widetilde{\tau}^{\varnothing} \quad \forall\ x_t \in \{0,1\},$$

$$\mathcal{T}^{1,0,1}(\sigma) := \widehat{\mathcal{N}}(\sigma), \qquad \mathcal{T}^{0,0,0}(\sigma) := \sigma.$$

Then,

$$\sigma(t|h^t) = x_t \rho(M(t)(h^t)) + (1-x_t)\tau^{\varnothing} \quad (\rho(m) = \widehat{\mathcal{N}}^{\circ m}(\rho_0)),$$

where $M(t)(h^t)$ is the number of time steps that the state is held in memory for the history h^t at time t. Also,

$$\Pr[H(t) = h^t] = \left(\prod_{j=1}^{t-1} d_j(h_j^t)(a_j)\right) p^{N_{\mathrm{succ}}(t)(h^t)}(1-p)^{N_{\mathrm{req}}(t)(h^t) - N_{\mathrm{succ}}(t)(h^t)}.$$

Average quantum state of a link

$$\sigma(t|h^t) = x_t \rho(M(t)(h^t)) + (1 - x_t)\tau^{\varnothing},$$

$$\Rightarrow \sigma(t) = \sum_{h^t} \widetilde{\sigma}(t; h^t) = \sum_{h^t} \Pr[H(t) = h^t]\sigma(t|h^t)$$

$$= \sum_{h^t : x_t = 0} \Pr[H(t) = h^t]\tau^{\varnothing} + \sum_{h^t : x_t = 1} \Pr[H(t) = h^t]\rho(M(t)(h^t))$$

$$\Rightarrow \boxed{\sigma(t) = (1 - \Pr[X(t) = 1])\tau^{\varnothing} + \sum_{m=0}^{t-1} \Pr[M(t) = m, X(t) = 1]\rho(m).}$$

Average state conditioned on link being active:

$$\boxed{\sigma(t|X(t) = 1) = \sum_{m=0}^{t-1} \Pr[M(t) = m|X(t) = 1]\rho(m), \quad \rho(m) = \widehat{\mathcal{N}}^{\circ m}(\rho_0).}$$

Fidelity of the link

$$f_m(\rho_0;\psi) := \langle\psi|\rho(m)|\psi\rangle = \langle\psi|\widehat{\mathcal{N}}^{\circ m}(\rho_0)|\psi\rangle.$$

Fidelity random variables:

$$\widetilde{F}(t;\psi) := X(t) f_{M(t)}(\rho_0;\psi), \quad F(t;\psi) := \frac{\widetilde{F}(t;\psi)}{\Pr[X(t)=1]}.$$

Average fidelities:

$$\boxed{\begin{aligned} \mathbb{E}[\widetilde{F}(t;\psi)] &= \sum_{m=0}^{t-1} f_m(\rho_0;\psi)\Pr[M(t)=m, X(t)=1], \\ \mathbb{E}[F(t;\psi)] &= \sum_{m=0}^{t-1} f_m(\rho_0;\psi)\Pr[M(t)=m|X(t)=1]. \end{aligned}}$$

$\mathbb{E}[\widetilde{F}(t;\psi)]$ is the expected reward after horizon time t.

Summary so far

- ✓ Cast elementary link generation formally as a quantum decision process.
  - Average quantum state of the link at any time.
  - Fidelity of the link and link activity probability as a function of time.
- Look at the memory cutoff policy.
- Policy optimization in the finite-horizon case.

The memory cutoff policy

$\star$ Once the link is established, keep it for $t^\star \geq 0$ time steps, then discard and request new link.

Deterministic policy:

$$d_t(h^t) = \begin{cases} 0 & \text{if } M(t)(h^t) < t^\star \text{ (wait)}, \\ 1 & \text{if } M(t)(h^t) = t^\star \text{ (request)}. \end{cases}$$

Two special cases:

- $t^\star = 0$: request new link at every time step.
- $t^\star = \infty$: once link is established, keep it in memory indefinitely.

To get the average quantum state and average fidelity, need $\Pr[M(t) = m, X(t) = 1]$ for this policy.

The memory cutoff policy

Look at the steady-state/infinite-horizon ($t \to \infty$) behavior.

$$\lim_{t \to \infty} \Pr[M(t) = m, X(t) = 1] = \frac{p}{1 + t^\star p}, \quad t^\star < \infty, \quad m \in \{0, 1, \dots, t^\star\}.$$

The memory cutoff policy

Look at the steady-state/infinite-horizon ($t \to \infty$) behavior.

$$\lim_{t\to\infty} \Pr[M(t) = m, X(t) = 1] = \frac{p}{1 + t^\star p}, \quad t^\star < \infty, \quad m \in \{0, 1, \dots, t^\star\}.$$

Probability that the link is active:

$$\lim_{t\to\infty} \Pr[X(t) = 1] = \lim_{t\to\infty} \mathbb{E}[X(t)]$$

$$= \frac{(t^\star + 1)p}{1 + t^\star p}, \quad t^\star \geq 0.$$

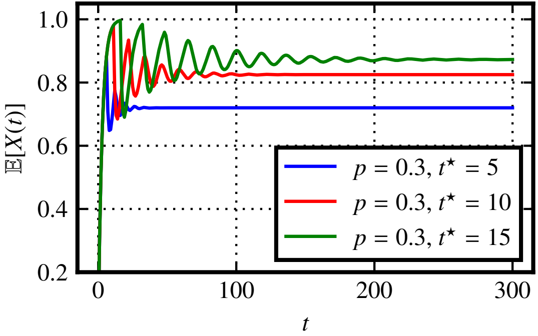

The memory cutoff policy

$$\lim_{t\to\infty} \Pr[M(t)=m, X(t)=1] = \frac{p}{1+t^\star p}, \quad \lim_{t\to\infty} \Pr[X(t)=1] = \frac{(t^\star+1)p}{1+t^\star p}$$

Average quantum state:

$$\sigma(t) = (1 - \Pr[X(t)=1])\tau^{\varnothing} + \sum_{m=0}^{t^\star} \Pr[M(t)=m, X(t)=1]\rho(m)$$

$$\Rightarrow \lim_{t\to\infty} \sigma(t) = \frac{1-p}{1+t^\star p}\tau^{\varnothing} + \frac{p}{1+t^\star p}\sum_{m=0}^{t^\star}\rho(m) \quad (t^\star < \infty),$$

$$\lim_{t\to\infty} \sigma(t|X(t)=1) = \sum_{m=0}^{t^\star} \Pr[M(t)=m|X(t)=1]\rho(m) = \frac{1}{t^\star+1}\sum_{m=0}^{t^\star}\rho(m), \quad (t^\star < \infty).$$

★ These results can be used to calculate/estimate entanglement distillation rates.

The memory cutoff policy

Total quantum state of a network:

$$\bigotimes_{e \in E} \bigotimes_{j=1}^{N_e^{\max}} \left(\frac{1 - p_{e,j}}{1 + t_{e,j}^{\star} p_{e,j}} \tau_{e,j}^{\varnothing} + \frac{p_{e,j}}{1 + t_{e,j}^{\star} p_{e,j}} \sum_{m=0}^{t_{e,j}^{\star}} \rho_{e,j}(m) \right)$$

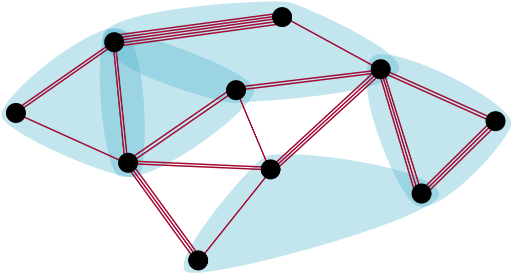

Edge capacity: $N_e^{\max}$

Edge flow:

$$N_e(t) = \sum_{j=1}^{N_e^{\max}} X_{e,j}(t) \Rightarrow \lim_{t \to \infty} \mathbb{E}[N_e(t)] = \sum_{j=1}^{N_e^{\max}} \frac{(t_{e,j}^{\star} + 1) p_{e,j}}{1 + t_{e,j}^{\star} p_{e,j}}$$

The memory cutoff policy

Total quantum state of a network:

$$\bigotimes_{e\in E}\bigotimes_{j=1}^{N_e^{\max}}\left(\frac{1-p_{e,j}}{1+t^\star_{e,j}p_{e,j}}\tau^{\varnothing}_{e,j}+\frac{p_{e,j}}{1+t^\star_{e,j}p_{e,j}}\sum_{m=0}^{t^\star_{e,j}}\rho_{e,j}(m)\right)$$

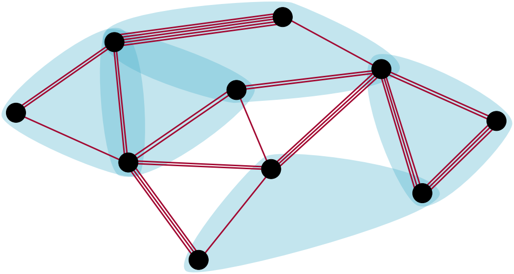

Edge capacity: $N_e^{\max}$

Edge flow:

$$N_e(t)=\sum_{j=1}^{N_e^{\max}}X_{e,j}(t)\Rightarrow\lim_{t\to\infty}\mathbb{E}[N_e(t)]=\sum_{j=1}^{N_e^{\max}}\frac{(t^\star_{e,j}+1)p_{e,j}}{1+t^\star_{e,j}p_{e,j}}$$

End-to-end flow for a chain (number of edge-disjoint paths):

$$N_{AB}(t)=\min\{N_1(t),N_2(t),N_3(t),N_4(t)\}.$$

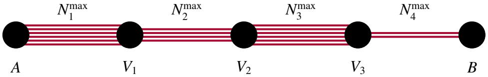

Summary so far

- ✓ Cast elementary link generation formally as a quantum decision process.
  - Average quantum state of the link at any time.
  - Fidelity of the link and link activity probability as a function of time.
- ✓ Look at the memory cutoff policy.
- Policy optimization in the finite-horizon case.

Policy optimization

- ⋆ For simplicity, consider the finite-horizon case; infinite-horizon ($t \to \infty$) for future work.
- ⋆ Goal is to determine policy that maximizes a particular figure of merit after a given, finite number of time steps.

Policy optimization

- ★ For simplicity, consider the finite-horizon case; infinite-horizon ($t \to \infty$) for future work.
- ★ Goal is to determine policy that maximizes a particular figure of merit after a given, finite number of time steps.

We maximize the expected reward, i.e., $\mathbb{E}[\widetilde{F}(t)]$. Other possibilities:

- ► $\mathbb{E}[X(t)]$ (probability that link is active): optimal policy is to keep the link indefinitely once it is established $\Rightarrow$ reduced fidelity over time.
- ► $\mathbb{E}[F(t)]$ (fidelity of the link, given that it is active): optimal policy is to request a link at every time step $\Rightarrow$ keeps the link active probability low.

Policy optimization

- ⋆ For simplicity, consider the finite-horizon case; infinite-horizon ($t \to \infty$) for future work.
- ⋆ Goal is to determine policy that maximizes a particular figure of merit after a given, finite number of time steps.

We maximize the expected reward, i.e., $\mathbb{E}[\widetilde{F}(t)]$. Other possibilities:

- ► $\mathbb{E}[X(t)]$ (probability that link is active): optimal policy is to keep the link indefinitely once it is established $\Rightarrow$ reduced fidelity over time.
- ► $\mathbb{E}[F(t)]$ (fidelity of the link, given that it is active): optimal policy is to request a link at every time step $\Rightarrow$ keeps the link active probability low.

The quantity $\widetilde{F}(t) = X(t) f_{M(t)}(\rho_0)$ balances both and leads to a non-trivial optimal policy. So we maximize $\mathbb{E}[\widetilde{F}(t)]$.

Policy optimization

- ⋆ Classical decision processes: exact optimal policy can be obtained using backward recursion (obtaining the optimal action at the *last* time step and working backwards).
- ⋆ Analogous result holds in the quantum case.

Policy optimization

- ⋆ Classical decision processes: exact optimal policy can be obtained using backward recursion (obtaining the optimal action at the *last* time step and working backwards).
- ⋆ Analogous result holds in the quantum case.

For all $t \geq 1$ and for any pure state ψ,

$$\max_{\pi} \mathbb{E}[\widetilde{F}(t+1;\psi)] = \sum_{x_1=0}^{1} \max_{a_1 \in \{0,1\}} v_1^*(t; x_1, a_1),$$

$$v_j^*(t; h^j, a_j) = \sum_{x_{j+1}=0}^{1} \max_{a_{j+1} \in \{0,1\}} v_{j+1}^*(t; h^{j+1}, a_{j+1}), \quad 1 \leq j \leq t-2,$$

$$v_{t-1}^*(t; h^{t-1}, a_{t-1}) = \sum_{x_t=0}^{1} \max_{a_t \in \{0,1\}} \langle \psi | \widetilde{\sigma}'(t+1; h^t, a_t, 1) | \psi \rangle.$$

Policy optimization

For all $t \geq 1$ and for any pure state ψ,

$$\max_{\pi} \mathbb{E}[\widetilde{F}(t+1;\psi)] = \sum_{x_1=0}^{1} \max_{a_1 \in \{0,1\}} v_1^*(t; x_1, a_1),$$

$$v_j^*(t; h^j, a_j) = \sum_{x_{j+1}=0}^{1} \max_{a_{j+1} \in \{0,1\}} v_{j+1}^*(t; h^{j+1}, a_{j+1}), \quad 1 \leq j \leq t-2,$$

$$v_{t-1}^*(t; h^{t-1}, a_{t-1}) = \sum_{x_t=0}^{1} \max_{a_t \in \{0,1\}} \langle \psi | \widetilde{\sigma}'(t+1; h^t, a_t, 1) | \psi \rangle.$$

Furthermore, the optimal policy is deterministic: $\pi^* = (d_j^* : 1 \leq j \leq t)$, where

$$d_j^*(h^j) = \arg\max_{a_j \in \{0,1\}} v_j^*(t; h^j, a_j), \quad 1 \leq j \leq t-1,$$

$$d_t^*(h^t) = \arg\max_{a_t \in \{0,1\}} \langle \psi | \widetilde{\sigma}'(t+1; h^t, a_t, 1) | \psi \rangle.$$

Summary

- ✓ Cast elementary link generation as a quantum decision process.
- ✓ Looked at the memory cutoff policy in the finite-horizon and infinite-horizon cases. Obtained analytic expressions for:
    - Average quantum state of the link at any time.
    - Fidelity of the link and link activity probability as a function of time.
    - Steady-state/infinite-horizon expressions for the fidelity and link activity probability.
- ✓ Policy optimization in the finite-horizon case.

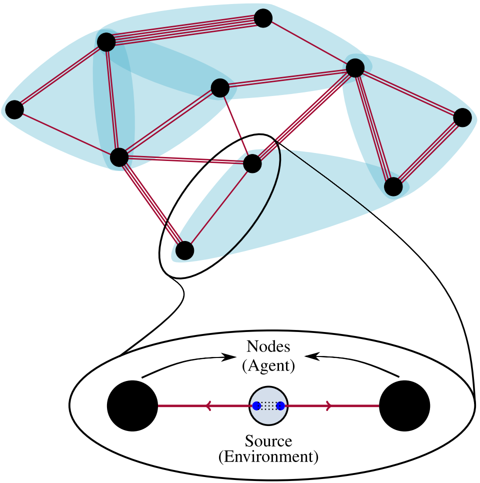

Paper available at arXiv:2007.03193.

Directions for future work

- Go beyond elementary link level: include entanglement distillation and joining measurements into the decision process.
- Use the decision process as a basis for reinforcement learning of (near-optimal) protocols.
- Perform policy optimization in the infinite-horizon setting.
- Extend results to “one-way” repeaters.
- More general question: develop *quantum algorithm* to solve quantum decision processes (some work in this direction already exists); idea is to see whether exponential speed-up can be obtained over the usual backward recursion method, which is exponentially slow in the horizon time.

Thank you!