Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Parallel Schedule of Pumping Purification in LEO Satellite Environment

Yifan Gao

yifan.gao@vuw.ac.nz

Victoria University of Wellington

Alvin Valera

alvin.valera@ecs.vuw.ac.nz

Victoria University of Wellington

Wuyungerile li

Inner Mongolia University

Winston Seah

winston.seah@ecs.vuw.ac.nz

Victoria University of Wellington

Method Article

Keywords: LEO Satellite, Classical Communication, Quan- tum Purification, Fidelity Upper Bound

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11804865/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Parallel Schedule of Pumping Purification in LEO Satellite Environment

Yifan Gao[1] Alvin Valera[1] Wuyungerile Li[2] Winston K.G. Seah[1]
[1]*School of Engineering and Computer Science* [2]*College of Computer Science*
Victoria University of Wellington *Inner Mongolia University*
Wellington, New Zealand Hohhot, China
yifan.gao@vuw.ac.nz, {alvin.valera, winston.seah}@ecs.vuw.ac.nz

***Abstract*—On low Earth orbit (LEO) inter-satellite links (ISLs), each purification round incurs a classical communication (CC) delay that cannot be ignored relative to the coherence time of the quantum memories. When each satellite has only a few qubits, the retained Einstein-Podolsky-Rosen (EPR) pair decoheres while every auxiliary EPR pair is generated, which limits the fidelity that pumping purification can reach. In this paper, we analyze a four-qubit node generates the next auxiliary EPR pair while the current purification round is still in progress. We show that this parallel schedule falls into two cases, separated by a critical number of link generation attempts. Then, we derive closed forms for its fidelity upper bound, convergence time and residual lifetime. Compared with sequential purification, the parallel schedule raises the fidelity upper bound by 0.004, saves up to eight purification rounds and reduces purification time by 80.7%. With two link-generation attempts per auxiliary EPR pair, the parallel schedule keeps the retained EPR pair alive for 361 slots at 40 km and extends the largest distance where purification is beneficial to 47.8 km.**

***Index Terms*—LEO Satellite, Classical Communication, Quantum Purification, Fidelity Upper Bound**

I. Introduction

Entanglement purification is conventionally regarded as a method of increasing fidelity without an intrinsic upper bound. However, on a free-space inter-satellite link (ISL), the attainable fidelity is bounded with very restricted node capacity. The classical communication (CC) delay is not negligible to the memory coherence time, so the target Einstein-Podolsky-Rosen (EPR) pair decoheres while each auxiliary EPR pair is generated and heralded. Moreover, every EPR pair is distributed over the same channel, so the auxiliary EPR pair is generated at a fixed initial fidelity that does not follow the target EPR pair upward. With the minimum capacity of four qubits, the model proposed by Gao et al. [1] generates each auxiliary pair only after the current round has ended, so the target pair also decoheres throughout the generation of every auxiliary pair. However, the same four qubits allow the next auxiliary pair to be generated while the current round is still in progress, since a round holds only one qubit of the auxiliary pair it consumes. One questions therefore arise, viz.,

1) At the same node capacity, how much does parallel scheduling raise the purification efficiency and extend the maximum ISL distance?
2) Under which conditions is the purification limited by the generation of the auxiliary pairs rather than by the duration of a purification round?

This work was supported in part by the China Scholarship Council (CSC), Grant No. 202508250019.

We consider two LEO satellites connected by a single ISL, each equipped with a four-qubit quantum repeater. A retained target EPR pair is purified repeatedly by the pumping BBPSSW protocol [2], with auxiliary pairs generated over the same free-space channel. The objective is to maximize the fidelity of the target EPR pair until it approaches its upper bound within a tolerance, subject to the limited node capacity, the cutoff time of the quantum memories, the CC delay of every round and the unheralded error of the local CNOT gates.

We model pumping purification on a four-qubit node where the next auxiliary EPR pair is generated during the current purification round, tracking the timing and fidelity of every EPR pair, so that one round takes the maximum of the link generation duration and a purification round duration, instead of the sum of both durations. We show that the parallel schedule falls into two cases decided by a critical number of link generation attempts, limited by the waiting auxiliary EPR pairs in Case 1 ($T_{\mathrm{gen}} \leq T_R$) and by the waiting retained EPR pair E_0 in Case 2 ($T_{\mathrm{gen}} > T_R$). Then, we derive the maximum number of link generation attempts and the feasibility condition of later rounds. For each case, we derive closed forms for the fidelity upper bound $F_\infty(d, g)$, the number of rounds $N_\zeta(d, g)$ and the time $T_\zeta(d, g)$ to approach it. We also derive their the residual lifetime $R^{(\infty)}(d, g)$ of E_0. Simulation results show that although the parallel schedule can only raise the fidelity upper bound slightly compared to the sequential schedule, it substantially reduces resource consumption, requiring up to seven fewer purification rounds and reducing purification time by up to 80.7%. Furthermore, it extends the maximum feasible purification distance in the LEO environment from 46.17 km to 47.8 km.

II. Related Work

Existing purification models are developed for fiber networks and either ignore the CC delay or treat it as a fixed overhead that does not affect the final fidelity [3]–[7]. Under this assumption, a limited node capacity can only limit the generation of auxiliary pairs and is therefore treated as a throughput or scheduling constraint. Existing models that

account for finite CC delays do not simultaneously consider free-space channels and finite node capacities [4], [7]–[9].

Entanglement buffering lets fresh EPR pairs be generated in parallel with a stored EPR pair that they pump [10], [11], but models generation as a rate process and treat purification as instantaneous, so the duration of a purification round and its CC never appear. Conversely, the works that account for the CC delay process the pairs sequentially. The idle periods imposed by CC delays expose stored EPR pairs to further decoherence [8] and the first pair decoheres while the second one is generated, so the two EPR pairs enter purification at different fidelities [1]. Mobayenjarihani et al. [12] remove the per-round CC by differing heralding and purification outcomes. None of these works does the generation of the next auxiliary pair parallel a purification round in progress. When both take finite time, the attainable fidelity depends on which of the two is longer. To the best of our knowledge, this effect of node capacity on fidelity has not been analysed.

III. System Model

A. Scenario

We consider a one-hop ISL with satellite *A* (source) and satellite *B* (destination) at distance d apart connected by a free space link [13]. Distance d is constant within the observation window since the cutoff time of a link is ≈ 0.3 s, which is far shorter than the ≈ 5 min contact window of inter-plane LEO satellites [13], [14]. Time is slotted with slot duration $\Delta t = 1$ ms. Each satellite has quantum memory for $N_q = 4$ qubits. Stored qubits undergo amplitude damping at rate Γ, so an entangled pair of qubits with fidelity F stored for i slots has fidelity $FD(i)$, where $D(i) = e^{-2\Gamma i \Delta t}$

An EPR pair is established by repeated generation attempts. In each attempt, *A* prepares an EPR pair locally with probability p in $\tau_{\text{prep}} = \lceil \frac{1}{R_{\text{src}}\Delta t} \rceil$ slots, holding both of its qubits, where R_{src} is the EPR pair generation rate. One qubit is then emitted as a photon to be captured by *B* with probability $q_{\text{error}}(d)$ and the outcome is heralded back to *A*, which takes $\tau_{\text{dist}} = \left\lceil \frac{2d}{c} + t_{\text{proc}}^{(\text{Classical})} \right\rceil \cdot \frac{1}{\Delta t}$ slots, including the heralding message of $\tau_{\text{herald}} = \left\lceil \frac{d}{c} + t_{\text{proc}}^{(\text{Classical})} \right\rceil \cdot \frac{1}{\Delta t}$ slots, where $t_{\text{proc}}^{(\text{Classical})}$ is the classical processing time. A pair is accepted $p_{\text{suc}}(d) = 1$ if its fidelity meets the threshold F_{th} and an attempt succeeds with probability $p'(d) = p \cdot q_{\text{error}}(d) \cdot p_{\text{suc}}(d)$. A failed attempt resets its qubits for $\tau_{\text{rst}} = 1$ slot. Assuming every EPR pair is established at its g-th attempt, generating an EPR pair takes

$$T_{\text{gen}}(g) = g(\tau_{\text{prep}} + \tau_{\text{dist}} + \tau_{\text{rst}})$$

slots. The first pair established is the target pair (E_0), retained throughout. The auxiliary pairs $E_1, E_2, \cdots$ are numbered in order of establishment and round k consumes E_k. We write q_{kA} and q_{kB} for the qubits of E_k at *A* and *B*.

A purification round consists of a synchronization message ($\tau_{\text{sync}} = \tau_{\text{herald}}$ slots), bilateral CNOT gates with the qubit of E_0 as control and that of E_k as target (typically for CNOT), which takes τ_{gate} slots. Each gate acts ideally with probability p_g and otherwise replaces its two qubits with the maximally mixed state. Then, a round has a second synchronization (τ_{sync} slots), measurement of E_k in the computational basis (τ_{BSM} slots) and an acknowledgment exchange of the outcome ($\tau_{\text{ack}} = \left\lceil 2\left(\frac{d}{c} + t_{\text{proc}}^{(\text{Classical})}\right) \right\rceil \cdot \frac{1}{\Delta t}$ slots). The steps before the BSM and the whole purification round therefore last

$$w_0 = 2\tau_{\text{sync}} + \tau_{\text{gate}} \tag{1}$$

$$T_R = w_0 + \tau_{\text{BSM}} + \tau_{\text{ack}} \tag{2}$$

slots. If the outcomes agree, E_0 is retained with increased fidelity, otherwise, both pairs are discarded and the schedule restarts with a new target pair ($F_{\text{fail}} < \frac{1}{2}$. We write $F_{E_0}^{(k)}(d, 0; g)$ for the fidelity of E_0 after k successful rounds. E_k is held only for the first w_0 slots and released once its readout begins, whereas E_0 is held for the whole round, because its state is determined at the BSM, but the next round cannot begin before the ACK. Between its heralding and the start of the round that consumes it, an auxiliary EPR pair keeps one qubit at each satellite and decoheres, we use δ to denote this interval, which depends on the schedule rather than on the protocol.

B. Timeline

A new EPR pair can be generated only when at least two qubits of *A* are free. Both stay occupied until its heralding. If success, one of them is released, if reports failure, both are reset during τ_{rst} and reused. Generating an EPR pair therefore occupies two qubits of *A* throughout its whole duration $T_{\text{gen}}(g)$, including its failed attempts, during which the retained EPR pair keeps decohering. Since E_0 holds one qubit of *A* throughout, let $n_1(t)$ be the number of auxiliary EPR pairs under generation (two qubits each) and $n_2(t)$ be the number of heralded auxiliary EPR pairs (one qubit each), the capacity constraint at *A* after E_0 is heralded is

$$2n_1(t) + n_2(t) + 1 \leq N_q$$

For $N_q = 4$, at most one EPR pair can be under generation and $n_2(t) \leq 1$ whenever $n_1(t) = 1$. Hence, while one auxiliary EPR pair is held by a purification round in progress and another one has been heralded, only one qubit of *A* is free, so no new generation can start, this is what delays the whole pumping schedule and lowers the fidelity upper bound.

Every EPR pair has fidelity F_0 when its local preparation is complete and $F_0'(d)$ when its photon is captured by *B*. An EPR pair not yet used for purification whose photon was captured for t slots has fidelity

$$F(d, t) = F_0'(d)D(t) \tag{3}$$

so it is heralded at *A* with $F(d, \tau_{\text{herald}})$ and an auxiliary pair E_k that waits δ_k slots enters the BSM with $F(d, \tau_{\text{herald}} + \delta_k + w_0)$. After k rounds, E_0 stored for t slots has $F_{E_0}^{(k)}(d, t; g) = F_{E_0}^{(k)}(d, 0; g)D(t)$, where $F_{E_0}^{(0)}(d, t) = F(d, \tau_{\text{herald}} + t)$ is counted from the heralding of E_0. By Eq. (3), E_0 stays above $\frac{1}{2}$ after k rounds with its residual lifetime [1] as

$$R^{(k)}(d, g) = \frac{1}{2\Gamma\Delta t} \ln 2F_{E_0}^{(k)}(d, 0; g). \tag{4}$$

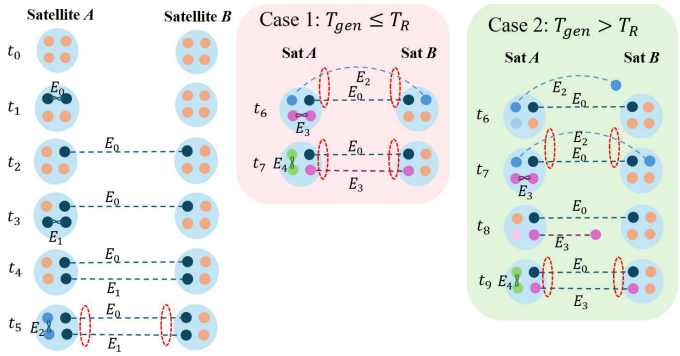


Fig. 1: Pumping purification with $N_q = 4$ qubits per satellite. The instants $t_0, \cdots, t_5$ are common to both cases. The pink part ($T_{\text{gen}}(g) \leq T_R$) and the green part ($T_{\text{gen}}(g) > T_R$) continue from t_5. The orange positions are free and colored positions hold the qubits of the indicated pair. A light-colored position is a qubit of A that is still served for a pair whose photon has been emitted but which has not yet been heralded. The red dashed lines mark the purification rounds in progress.

Let $u_k, h_k = u_k - \tau_{\text{rst}} + T_{\text{gen}}(g)$ and s_k be the instants when the link generation of E_k begins, E_k is heralded and the k-th purification round begins. Then,

$$\begin{aligned} u_1 &= t_2 + \tau_{\text{rst}}, u_2 = h_1 + \tau_{\text{rst}}, \\ u_{k+1} &= \max\{h_k, s_{k-1} + w_0\} + \tau_{\text{rst}}, k \geq 2 \\ s_1 &= h_1, s_k = \max\{h_k, s_{k-1} + T_R\}, k \geq 2 \end{aligned}$$

Generating E_{k+1} needs two free qubits of A, which exist only once E_k is heralded and E_{k-1} is released at the start of the BSM of round $k-1$, the first link generation attempt then begins after a reset. A round begins once its auxiliary EPR pair is heralded and the previous round has ended. Between rounds, E_0 idles for $T_{\text{idle},k} = s_k - s_{k-1} - T_R \geq 0$ slots, $k \geq 2$. We call the steady-state value of $s_k - s_{k-1}$ the cycle time T_{cyc}. If $T_{\text{gen}}(g) > T_R$, E_0 waits for E_k, ($T_{\text{idle},k} > 0, \delta_k = 0$). Otherwise, E_k waits for E_0 ($T_{\text{idle},k} = 0, \delta_k > 0$).

1) Common Instants: Fig. 1 shows the instants $t_0, \cdots, t_5$ shared by both cases (only the successful link generation attempt of E_0 is drawn). t_1 and t_3 mark the end of the local preparation of E_0 and E_1. The generation of E_0 begins at t_0 and E_0 is heralded at t_2 with $F_{E_0}^{(0)}(d, 0) = F(d, \tau_{\text{herald}})$. One qubit of A is then released and the generation of E_1 begins on two of the three free qubits after a reset. E_1 is heralded at $t_4 = t_2 + T_{\text{gen}}(g)$, when E_0 has fidelity $F_{E_0}^{(0)}(d, T_{\text{gen}}(g))$. As A now holds only q_{0A} and q_{1A}, the first purification round and the generation of E_2 begins together. At $t_5 = t_4 + w_0$, the BSM begins with E_1 with $F(d, \tau_{\text{herald}} + w_0)$, After BSM, E_1 is released. After the ACK, at $t_4 + T_R$, E_0 has fidelity $F_{E_0}^{(1)}(d, 0; g) = \beta F_{\text{post,ng}}(F_{E_0}^{(0)}(d, T_{\text{gen}}(g) + w_0 + \tau_{\text{BSM}}), F(d, \tau_{\text{herald}} + w_0))$ and E_2 is heralded at $h_2 = t_4 + T_{\text{gen}}(g)$. Whether E_2 is heralded before or after the first round ends at $t_4 + T_R$ separates as two cases, the cycle time is $T_{\text{cyc}} = \max\{T_{\text{gen}}(g), T_R\}$.

2) Case 1 ($T_{\text{gen}}(g) \leq T_R$): In the pink part of Fig. 1, E_2 is heralded before the first purification round ended and waits $\delta_2 = T_R - T_{\text{gen}}(g)$ slots for E_0. The second purification round begins at once at $t_6 = t_4 + T_R$.

Therefore, E_0 never idles, $T_{\text{idle},k} = 0, s_k = s_1 + (k-1)T_R$, but the auxiliary pairs always wait. At first, the generation of E_{k+1} starts right after h_k, so the waiting time grows by $T_R - T_{\text{gen}}(g)$ per purification round, until E_k is heralded before E_{k-1} is released at $s_{k-1} + w_0$. From then on, generation start after this release, the waiting time of the auxiliary pairs is

$$\delta_k = \min\{(k-1)(T_R - T_{\text{gen}}(g)), 2T_R - w_0 - T_{\text{gen}}(g)\}, k \geq 2 \tag{5}$$

And the post-fidelity of E_0 after BSMs is $F_{E_0}^{(k)}(d, 0; g) = \beta F_{\text{post,ng}}(F_{E_0}^{(k-1)}(d, w_0 + \tau_{\text{BSM}}; g), F(d, \tau_{\text{herald}} + \delta_k + w_0)), k \geq 2$.

In the steady state, δ_k reaches

$$\delta_\infty = \lim_{k \to \infty} \delta_k = \begin{cases} 2T_R - w_0 - T_{\text{gen}}(g), & T_{\text{gen}}(g) < T_R \\ 0, & T_{\text{gen}}(g) = T_R \end{cases} \tag{6}$$

and every auxiliary pair enters the BSM with fidelity $F_{\text{aux}} = F(d, \tau_{\text{herald}} + \delta_\infty + w_0)$, which lowers the fidelity upper bound compared with consuming every auxiliary EPR pair as soon as it is heralded. Since E_0 is stored for only $w_0 + \tau_{\text{BSM}}$ slots between an ACK and the next projection, with $\alpha = D(w_0 + \tau_{\text{BSM}})$ and this F_{aux}, the fidelity upper bound is [1]

$$\begin{aligned} F_\infty(d, g) &= \frac{-B + \sqrt{B^2 - 4AC}}{2A}, \\ A = m_{21}, B &= m_{22} - m_{11}, C = -m_{12} \end{aligned} \tag{7}$$

where $m_{11} = 8\beta\alpha p_g^2(10F_{\text{aux}} - 1), m_{12} = \beta(9 - p_g^2 - 8F_{\text{aux}}p_g^2), m_{21} = 16\alpha p_g^2(4F_{\text{aux}} - 1)$ and $m_{22} = 4(p_g^2 - 4p_g^2 F_{\text{aux}} + 9)$ are the entries of the pumping map

$$M = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix}. \tag{8}$$

Also, $T_{\text{cyc}} = T_R$. The schedule further requires $F_{\text{aux}} > \frac{1}{2} \Rightarrow \delta_\infty < \frac{\ln\left(2F(d, \tau_{\text{herald}} + w_0)\right)}{2\Gamma\Delta t}$ as $\delta_k \leq 2T_R - w_0 - T_{\text{gen}}(g)$. Otherwise, the auxiliary pair has to be discarded before it is used.

3) Case 2 ($T_{\text{gen}}(g) > T_R$): In the green part of Fig. 1, when the first purification round ends at $t_6 = t_4 + T_R$, E_2 is still be distributed and holds two qubits of A. With one free qubit, A cannot start a new link generation and E_0 idles until $h_2 = t_4 + T_{\text{gen}}(g)$, where the second purification round and the generation of E_3 begin together. The same pattern repeats, $h_k = s_k = t_2 + kT_{\text{gen}}(g), \delta_k = 0$ and $T_{\text{idle},k} = T_{\text{gen}}(g) - T_R$ for $k \geq 2$, so $F_{E_0}^{(k)}(d, 0; g) = \beta F_{\text{post,ng}}\left(F_{E_0}^{(k-1)}(d, w'; g), F(d, \tau_{\text{herald}} + w_0)\right)$, where the total waiting time of E_0 here is $w' = T_{\text{gen}}(g) - T_R + w_0 + \tau_{\text{BSM}}$. Then, $F_\infty(d, g)$ follows Eq. (7) with $\alpha = D(T_{\text{gen}}(g) - T_R + w_0 + \tau_{\text{BSM}})$ and $F_{\text{aux}} = F(d, \tau_{\text{herald}} + w_0)$ and $T_{\text{cyc}} = T_{\text{gen}}(g)$.

C. Fidelity Variation

In any case, E_0 decays in purification round k in two parts, accumulated from the ACK of the previous round until the projection α_k and accumulated after it during the ACK $\beta =$

$D(\tau_{\text{ack}})$, while the auxiliary pair is frozen once its readout begins. If without idle time ($T_{\text{gen}}(g) \leq T_R$), $\alpha_k = \alpha_0 = D(w_0 + \tau_{\text{BSM}})$. The fidelity of E_0 then follows the recursion as

$$
\begin{aligned}
&F_{E_0}^{(k)}(d,0;g) \\
&= \beta F_{\text{post,ng}}\left(\alpha_k F_{E_0}^{(k-1)}(d,0;g), F(d,\tau_{\text{herald}}+\delta_k+w_0)\right), \\
&k \geq 1 \\
&F_{E_0}^{(0)}(d,0) = F(d,\tau_{\text{herald}}), \alpha_0 = D(w_0+\tau_{\text{BSM}}), \\
&\alpha_1 = D(T_{\text{gen}}(g)+w_0+\tau_{\text{BSM}}), \\
&\alpha_k = \alpha_0 D(T_{\text{idle},,k}), k \geq 2
\end{aligned}
\tag{9}
$$

where the first purification round contains the whole generation duration of E_1 and $F(d,\tau_{\text{herald}}+\delta_k+w_0) = F(d,\tau_{\text{herald}}+w_0)D(\delta_k)$. The schedule thus enters the recursion only through $T_{\text{idle},,k} = \max\{h_k - s_{k-1} - T_R, 0\}$ and $\delta_k = \max\{0, s_{k-1}+T_R-h_k\}$, at most one of them is positive in any round, that either E_0 waits for the auxiliary pair, or the auxiliary pair waits for E_0.

IV. Steady-state Analysis

The parallel schedule changes the fidelity recursion only through the timing of the purification rounds, like how long E_0 waits idle between purification rounds and how long the auxiliary pairs wait, both determined by the relation between $T_{\text{gen}}(g)$ and T_R. This section analyses the boundary between the two cases decided by a critical number of generation attempts, then we derive the fidelity upper bound, the number of rounds and the residual lifetime of E_0 after k rounds in closed-form.

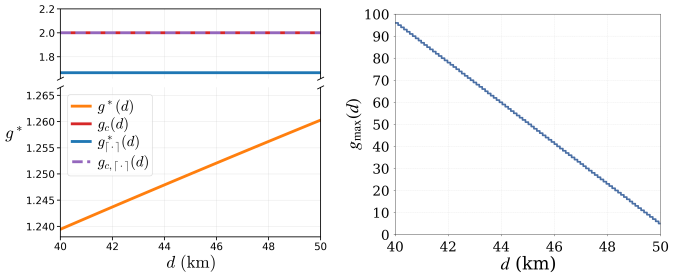
A. The Number of Link Generation Attempts

The two cases share the same purification round duration T_R as Eqs. (1)(2) and differ only in $T_{\text{gen}}(g) = g(d)(\tau_{\text{prep}} + \tau_{\text{dist}}(d) + \tau_{\text{rst}})$, so their boundary is a number of attempts. Setting $T_{\text{gen}}(g) = T_R(d)$ gives the critical value

$$
g^*(d) = \frac{2\tau_{\text{sync}}(d) + \tau_{\text{gate}} + \tau_{\text{BSM}} + \tau_{\text{ack}}(d)}{\tau_{\text{prep}} + \tau_{\text{dist}}(d) + \tau_{\text{rst}}} \tag{10}
$$

Therefore, if $g(d) > g^*(d)$, E_0 waits, if $g(d) < g^*(d)$, the auxiliary pairs wait. At $g(d) = g^*(d)$, neither waits. Since $g(d)$ is an integer, the smallest g for E_0 waits is $g_c(d) = \lceil g^*(d) \rceil$. We assume all EPR pairs of a purification schedule are established at the same g-th attempt, so a schedule stays in one case in every round.

Fig. 2(a) shows $g^*(d)$ and $g_c(d)$ in $d \in [40, 50]$ km, also with every part of the duration rounded up to whole slots $(g^*_{\lceil\cdot\rceil}(d), g_{c,\lceil\cdot\rceil}(d))$. Since $d/c = 0.13$–0.17 ms $\ll \Delta t = 1$ ms, $g^*(d)$ increases only from 1.24 to 1.26 without rounding and stays at ≈ 1.67 with rounding. So, both give $g_c(d) = g_{c,\lceil\cdot\rceil}(d) = 2$ for all $d \in [40, 50]$ km. Therefore, the parallel is in Case 1 only for $g = 1$.



(a) No. of Link-Gen Attempts (b) $g_{\max}(d)$ theoretically

Fig. 2: Critical $g^*(d)$ and Maximum $g_{\max}(d)$

1) Maximum Number of Attempts: While $g^*(d)$ separates the two cases, it does not limit g. An upper limit of g should follow from the survival of E_0 before the first purification round. In both cases, E_0 waits for the whole $T_{\text{gen}}(g)$ slots and $w_0 + \tau_{\text{BSM}}$ slots of the first purification round before E_1 is projected, so the first purification round requires $F(d,\tau_{\text{herald}})e^{-2\Gamma(T_{\text{gen}}(g)+w_0+\tau_{\text{BSM}})\Delta t} \geq \frac{1}{2}$. With $R^{(0)}(d)$ from Eq. (4), the largest number of attempts is

$$
g_{\max}(d) = \left\lfloor \frac{R^{(0)}(d) - w_0 - \tau_{\text{BSM}}}{\tau_{\text{prep}} + \tau_{\text{dist}} + \tau_{\text{rst}}} \right\rfloor \tag{11}
$$

Fig. 2(b) shows $g_{\max}(d)$ from for $d \in [40, 50]$ km. It decreases from 96 at 40 km to 5 at 50 km, a drop of about 9 per km. In this range, the decrease comes from $R^{(0)}(d)$, which becomes shorter as $F_{E_0}^{(0)}(d)$ drops with distance. However, $g_{\max}(d)$ only ensures that the fidelity of E_0 stays above $\frac{1}{2}$ in the first purification round, it does not ensure a gain from purification even if it is successful.

B. Pumping Map

After k_0 initial rounds, every purification of each schedule applies the same map to E_0, that E_0 is stored for w slots from an ACK to the next projection, so $\alpha_k = D(w)$ in Eq. (9) and the auxiliary pair enters the BSM with a constant fidelity F_{aux}. For $k > k_0$, Eq. (9) reduces to [1]

$$
\begin{aligned}
F_{E_0}^{(k)}(d,0;g) &= G(F_{E_0}^{(k-1)}(d,0;g)) \\
&= \frac{m_{11}F_{E_0}^{(k-1)}(d,0;g) + m_{12}}{m_{21}F_{E_0}^{(k-1)}(d,0;g) + m_{22}}
\end{aligned}
$$

where m_{ij} is the entries of the pumping metric as Eq. (8). The schedules differ only in w, F_{aux} and k_0 (Tab. I). In Case 2, only the first purification round differs, so $k_0 = 1$. In Case 1, α is constant from the second round on, but F_{aux} changes while $\delta_k = \min\{(k-1)(T_R - T_{\text{gen}}(g)), \delta_\infty\}$ grows (Eqs. (5)(6)), depending on the waiting time δ_k of the auxiliary pairs, which changes over the first purification round. By Eqs. (5)(6), $\delta_k = \min\{(k-1)(T_R - T_{\text{gen}}(g)), \delta_\infty\}$ with $\delta_\infty = 2T_R - w_0 - T_{\text{gen}}(g)$, where δ_k grows by $T_R - T_{\text{gen}}(g)$ per round until the generation of the next auxiliary pair has to wait for the release of the previous one, then it stays at δ_∞. $F_{\text{aux}} = F(d,\tau_{\text{herald}} + \delta_k + w_0)$ is therefore constant from round $k_0 + 1$ on, where

TABLE I: Parameters of the pumping map G

Schedule	w	$\alpha = D(w)$	β	F_{aux}	k_0	T_{cyc}
Case 1, $g \le g^*(d)$	$w_0 + \tau_{\text{BSM}}$	$D(w_0 + \tau_{\text{BSM}})$	$D(\tau_{\text{ack}})$	$F(d, \tau_{\text{herald}} + \delta_\infty + w_0)$	Eq. (12)	T_R
Case 2, $g > g^*(d)$	$T_{\text{gen}} - T_R + w_0 + \tau_{\text{BSM}}$	$D(T_{\text{gen}} - T_R + w_0 + \tau_{\text{BSM}})$	$D(\tau_{\text{ack}})$	$F(d, \tau_{\text{herald}} + w_0)$	1	T_{gen}
Sequential [1]	$T_{\text{gen}} + w_0 + \tau_{\text{BSM}}$	$D(T_{\text{gen}} + w_0 + \tau_{\text{BSM}})$	$D(\tau_{\text{ack}})$	$F(d, \tau_{\text{herald}} + w_0)$	0	$T_{\text{gen}} + T_R$

so k_0 is the smallest integer when $\delta_{k_0+1} = \delta_\infty \Rightarrow (k_0 + 1 - 1)(T_R - T_{\text{gen}}(g)) = 2T_R - w_0 - T_{\text{gen}}(g)$, so

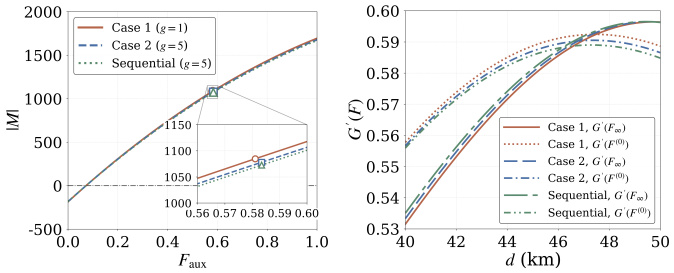
$$k_0 = \left\lceil \frac{2T_R - w_0 - T_{\text{gen}}(g)}{T_R - T_{\text{gen}}(g)} \right\rceil, \tag{12}$$

where $T_R - T_{\text{gen}}(g) > 0$ in Case 1. When $T_R = T_{\text{gen}}(g), k_0 = 1$, because $\delta_k = 0$ in every purification round. In the sequential schedule [1], by contrast, E_0 waits the whole $T_{\text{gen}}(g)$ before every purification round, so $k_0 = 0$.

In both cases, the fidelities $F_{E_0}^{(k)}(d, 0; g)$ of the initial rounds $k \le k_0$ do not follow the fixed map G, but obtained directly from the general recursion with Eq. (9). Fig. 3 shows $|M|$ and G' for Case 1 with $g = 1$, Case 2 with $g = 5$ and the sequential schedule with $g = 5$. Fig. 3(a) plots the determinant of of the matrix M in Eq. (8) against F_{aux} at $d = 45$ km and the markers show the value of F_{aux} obtained by each schedule. From Eq. (8), $|M| = -48\alpha\beta p_g^2 Q(F_{\text{aux}})$, where $Q(F_{\text{aux}}) = 16p_g^2 F_{\text{aux}}^2 - (8p_g^2 + 48)F_{\text{aux}} + p_g^2 + 3$ depends only on p_g^2. Since Case 1 and Case 2 differ from the sequential schedule only in α and F_{aux}, the three curves have the same shape and differ only in the factor $\alpha\beta$, which is 0.995, 0.985 and 0.980 for Case 1, Case 2 and the sequential schedule, respectively. All three curves cross zero at the same $F_{\text{aux}} = 0.073$ and then, $|M| > 0$ above it. As shown in Fig. 3(a), G is increasing for $F_{\text{aux}} \ge 0.1$ and every $p_g \in (0, 1]$, since $F_{\text{aux}} \ge 0.5$ unless useless, so every purification round moves $F_{E_0}^{(k)}(d, 0; g)$ toward $F_\infty(d, g)$ without passing it. The inset shows Case 1 has the largest $|M|$, followed by Case 2 and the sequential schedule. This is because E_0 waits for a different time before each purification round. It does not wait in Case 1, while it waits $T_{\text{gen}}(g) - T_R$ slots in Case 2 and $T_{\text{gen}}(g)$ slots in the sequential schedule. The value of $F_{\text{aux}} = 0.581$ in Case 1 is lower than in the other two schedules (0.583), since its auxiliary pairs wait δ_∞ slots before they are used. Moreover, $F_{\text{post,ng}}(F_{\text{targ}}, F_{\text{aux}}) = \frac{p_g^2\left[F_{\text{targ}}F_{\text{aux}} + \frac{(1-F_{\text{targ}})(1-F_{\text{aux}})}{9}\right] + \frac{1-p_g^2}{8}}{p_{s,\text{ng}}(F_{\text{targ}}, F_{\text{aux}})}$ is increasing in both of its arguments, so G is also increasing in α and F_{aux}.

Fig. 3(b) shows $G'(F_\infty(d, g))$ and $G'(F_{E_0}^{(0)}(d))$ against d. From 40 to 50 km, $G'(F_\infty(d, g))$ rises from about 0.53 to 0.60, so G convergence slows with distance, while $G'(F_{E_0}^{(0)}(d))$ peaks at about 47.4 km. The two sets of curves for each schedule cross where $F_\infty(d, g) = F_{E_0}^{(0)}(d)$, at 46.77 km for Case 1, 47.8 km for Case 2 and 46.17 km for the sequential schedule, purification can increase the fidelity of E_0 below these distances.

1) Fidelity Upper Bound: The fidelity upper bound is the fixed point when $F_\infty(d, g) = G(F_\infty(d, g))$, which is given by Eqs. (7) (8) with different α, β and F_{aux} of Tab. I. Since G increases with α and with F_{aux}, a larger α or a larger F_{aux} can let $F_\infty(d, g)$ highers. The two cases limit $F_\infty(d, g)$ in different ways.



(a) $F_2 > 0.1, \det M > 0, G \uparrow$ (b) $G'(F) < 1 \Rightarrow G$ converges

Fig. 3: Monotonicity and Convergence Rate of Pumping Map

In Case 1, E_0 never waits in idle, so $w = w_0 + \tau_{\text{BSM}}$ is the shortest possible storage time of E_0 per purification round and α takes its largest value. The fidelity bound is limited by the auxiliary pairs, which wait $\delta_\infty = 2T_R - w_0 - T_{\text{gen}}(g)$ slots before they are consumed and enter the BSM with $F_{\text{aux}} = F(d, \tau_{\text{herald}} + w_0)D(\delta_\infty)$. The earlier an auxiliary pair is heralded, the worse the final fidelity, so the shorter $T_{\text{gen}}(g)$, the longer it waits and the lower $F_\infty(d, g)$ is.

In Case 2, every auxiliary pair is consumed as soon as it is heralded, so $F_{\text{aux}} = F(d, \tau_{\text{herald}} + w_0)$ takes its largest value. The fidelity bound is limited by E_0, which waits in idle for $T_{\text{gen}}(g) - T_R$ slots in every purification round, so $\alpha = D(T_{\text{gen}}(g) - T_R + w_0 + \tau_{\text{BSM}})$. Since the waiting time grows with $T_{\text{gen}}(g)$, $F_\infty(d, g)$ decreases with g in this case.

Both limitations vanish at $T_{\text{gen}}(g) = T_R$, where $F_\infty(d, g)$ attains its largest value $F_\infty^{\text{best}}(d)$, the fixed point of G with $\alpha = D(w_0 + \tau_{\text{BSM}})$ and $F_{\text{aux}} = F(d, \tau_{\text{herald}} + w_0)$. Since g is an integer, so the fidelity upper bound is attained at $g = g_c(d)$. For each g, $d_{\max}(g)$ is the largest distance with $F_\infty(d, g) > F_{E_0}^{(0)}(d, 0)$.

2) Number of Purification Rounds: The remaining infidelity after k rounds is split as $1 - F_{E_0}^{(k)}(d, 0; g) = [1 - F_\infty(d, g)] + [F_\infty(d, g) - F_{E_0}^{(k)}(d, 0; g)]$, further rounds cannot reduce the first term and the second determines how long pumping is worth continuing. Once the gap $e_k = F_\infty(d, g) - F_{E_0}^{(k)}(d, 0; g)$ is small, each purification round shrinks by the factor $\lambda = G'(F_\infty(d, g)) \in (0, 1)$ and the number of rounds required to bring it below a tolerance ζ is $N_\zeta(d, g) \simeq \left\lceil \frac{\ln \frac{\zeta}{|F_\infty(d,g) - F_{E_0}^{(0)}(d,0)|}}{\ln \lambda} \right\rceil$ in [1].

The same works in both cases for $k > k_0$, with λ evaluated

at α, β and F_{aux} of Tab. I. Counting the first k_0 rounds, which do not follow the fixed map G, gives us the total number of rounds as

$$N_\zeta(d,g) \simeq k_0 + \left\lceil \frac{\ln \frac{\zeta}{|F_\infty(d,g) - F_{E_0}^{(k_0)}(d,0;g)|}}{\ln \lambda} \right\rceil. \tag{13}$$

$N_\zeta(d,g)$ thus depends on the case by k_0, λ and the initial reducible gap $|F_\infty(d,g) - F_{E_0}^{(k)}(d,0;g)|$.

Since the purification rounds start every T_{cyc} slots for $k > k_0$, the time needed to complete $N_\zeta(d,g)$ rounds from the heralding of E_0, is approximately

$$T_\zeta(d,g) \simeq T_{\text{gen}}(g) + (N_\zeta(d,g) - 1)T_{\text{cyc}} + T_R \tag{14}$$

with $T_{\text{cyc}} = T_R$ in Case 1 and $T_{\text{cyc}} = T_{\text{gen}}(g)$ in Case 2.

$T_\zeta(d,g)$ is thus conditioned on the success of every round, since a failed round leaves the fidelity of E_0 below $\frac{1}{2}$ and terminates the purification. Round k succeeds with probability

$$\begin{aligned} &p_{s,\text{ng}}(F_{\text{targ},k}, F_{\text{aux},k}) \\ &= p_g^2 \frac{8F_{\text{targ},k}F_{\text{aux},k} - 2F_{\text{targ},k} - 2F_{\text{aux},k} + 5}{9} + \frac{1-p_g^2}{2}, \end{aligned} \tag{15}$$

where the fidelities of the two EPR pairs entering its BSM are $F_{\text{targ},k} = F_{E_0}^{(k-1)}(d, w_k; g)$, where w_k is the storage time of E_0 from the previous ACK to the projection and $F_{\text{aux},k} = F(d, \tau_{\text{herald}} + \delta_k + w_0)$. With $w_1 = T_{\text{gen}} + w_0 + \tau_{\text{BSM}}$ and $w_k = w$ of Tab. I for $k \geq 2$, the probability that all $N_\zeta(d,g)$ purification rounds succeed is $P_\zeta(d,g) = \Pi_{k=1}^{N_\zeta(d,g)} p_{s,\text{ng}}\left(F_{E_0}^{(k-1)}(d,w_k;g), F(d,\tau_{\text{herald}} + \delta_k + w_0)\right)$.

3) Residual Lifetime: The residual lifetime $R^{(k)}(d,g)$ of E_0 after k rounds is given by Eq. (4) and round $k+1$ is admissible if E_0 is stored for at most $R^{(k)}(d,g)$ slots before its projection in that purification round. For $k \geq k_0$, this storage time is w of Tab. I. Since G is increasing, $R^{(k)}(d,g)$ varies monotonically with $k \geq k_0$, increasing if $F_{E_0}^{(k_0)}(d,0;g) < F_\infty(d,g)$ and decreasing otherwise, so $R^{(k)}(d,g)$ converges to

$$R^{(\infty)}(d,g) = \frac{1}{2\Gamma\Delta t} \ln 2F_\infty(d,g), \tag{16}$$

All purification after k_0 is therefore admissible if

$$w \leq \min\{R^{(k_0)}(d,g), R^{(\infty)}(d,g)\}. \tag{17}$$

Proposition 1. *If the first purification round is admissible*

$$T_{\text{gen}}(g) + w_0 + \tau_{\text{BSM}} \leq R^{(0)}(d),$$

all later purification rounds are admissible as well.

Proof. In Case 1, $w = w_0 + \tau_{\text{BSM}}$ is the shortest possible storage time of E_0, so this condition is the easiest to satisfy, although the lower F_{aux} may reduce $R^{(k)}(d,g)$. In Case 2, $w = T_{\text{gen}}(g) - T_R + w_0 + \tau_{\text{BSM}}$ grows with g, so the condition becomes more demanding as more attempts are needed. For the initial rounds $k \leq k_0$, it needs to check the first purification round. Before it, E_0 is heralded at t_2, waits for the whole generation of E_1 and E_1 is projected only after the subsequent $w_0 + \tau_{\text{BSM}}$ slots, so E_0 is stored for $T_{\text{gen}}(g) + w_0 + \tau_{\text{BSM}}$ slots, the same as the sequential schedule before every purification. This is the longest storage time of E_0 in both cases. In Case 2, every later round requires only $T_{\text{gen}}(g) - T_R + w_0 + \tau_{\text{BSM}}$ slots, T_R fewer. In Case 1, E_0 never waits idle after the first round, so every later round, including the initial rounds ($2 \leq k \leq k_0$), requires only $w_0 + \tau_{\text{BSM}}$ slots. Moreover, every successful round raises $F_{E_0}^{(k)}(d,0;g)$ when purification is beneficial and therefore raises $R^{(k)}(d,g)$. Therefore, if the first round is admissible with $T_{\text{gen}}(g) + w_0 + \tau_{\text{BSM}} \leq R^{(0)}(d)$, all later rounds are admissible as well. □

V. Evaluation

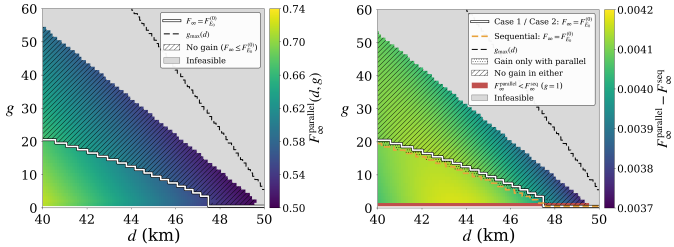
We compare the parallel schedule (Case 1 and Case 2) and the sequential schedule of in terms of $F_\infty(d,g)$ (Eq. (7)), N_ζ (Eq. (13)), T_ζ (Eq. (14)) and $R^{(k)}$ (Eqs. (16)), where N_ζ and T_ζ are computed conditioned on every purification round being successful. We evaluate the models over $\Gamma = 0.5s^{-1}, p_g = 0.99$ [13], $\Delta t = 1$ ms, $\tau_{\text{gate}} = \tau_{\text{BSM}} = \tau_{\text{rst}} = 1$ slot. $g^* = 2$ as Fig. 2(a), so, parallel schedule is in Case 1 if $g = 1$, is in Case 2 if $g \geq 2$. Figs. (4)(5)(6) show each metric over $d \in [40, 50]$ km and $g \in \{1, \cdots, 60\}$, with $\zeta \in \{10^{-2}, 10^{-3}, 10^{-4}\}$ in Fig. (5). Among them, the white solid line and the orange dashed line mark where $F_\infty(d,g) = F_{E_0}^{(0)}(d)$ for the parallel and the sequential schedule, respectively. The black dashed line is $g_{\max}(d)$ (Eq. (11)), the gray part is infeasible because it violates Eq. (11) or Eq. (17), the hatched area is feasible but provides no gain because $F_\infty(d,g) \leq F_{E_0}^{(0)}(d)$ and the dotted part means only the parallel schedule gains.

A. Fidelity Boundary

Fig. 4(a) shows the fidelity upper bound $F_\infty^{\text{parallel}}(d,g)$ of the parallel schedule. The color scale ranges from 0.5 (dark purple) to 0.74 (yellow). At every distance, $F_\infty^{\text{parallel}}(d,g)$ reaches its largest value at $g = 2$, which is slightly higher than its value at $g = 1$, because the auxiliary pairs in Case 1 wait δ_∞ slots before they are used, which lowers the post-fidelity after each purification. At 40 km, the highest fidelity is 0.718 at $g = 2$, compared with 0.715 at $g = 1$ and at 45 km, the highest fidelity is 0.606 at $g = 2$, compared with 0.602 at $g = 1$. For $g > 2$, $F_\infty^{\text{parallel}}(d,g)$ decreases with g, because the idle time of E_0 in each round $T_{\text{gen}}(g) - T_R$ grows with g. $F_\infty^{\text{parallel}}(d,g)$ also decreases with d, because both $F_{E_0}^{(0)}(d)$ and F_{aux} become lower and it falls to 0.50 at the feasibility boundary. The white line drops from $g = 21$ at 40 km to $g = 1$ above 47.5 km, so the part below it where purification raises the fidelity of E_0 shrinks quickly with distance. Above the white line, the parallel schedule can still work up to $g = 54$ at 40 km and $g = 27$ at 45 km, about half of $g_{\max}(d)$, but $F_\infty^{\text{parallel}}(d,g) \leq F_{E_0}^{(0)}(d)$, so purification lowers the fidelity instead of raising it.

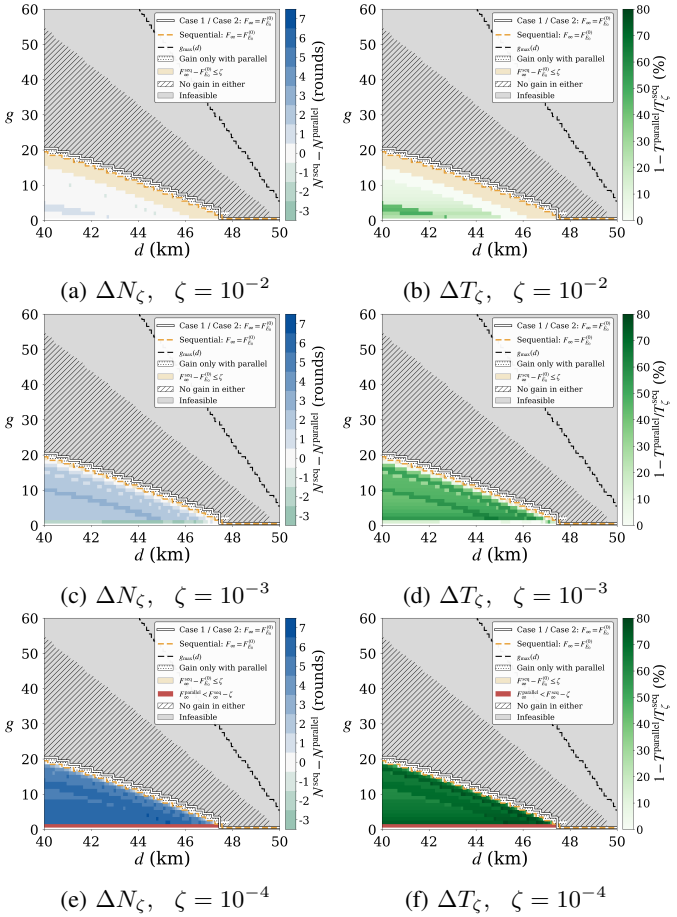
Fig. 4(b) compares the fidelity upper bounds of parallel and sequential schedules through $\Delta F_\infty(d,g) = F_\infty^{\text{parallel}}(d,g) - F_\infty^{\text{seq}}(d,g)$ for every (d,g) where both schedules are feasible. For $g \geq 2$, $\Delta F_\infty(d,g)$ is positive everywhere and lies between 0.0037 (dark blue) and 0.0042 (yellow), because E_0 is stored

for T_R fewer slots in each purification round of Case 2, so α is larger by the fixed factor $D(T_R)$ while β and F_{aux} are the same as in the sequential schedule. $\Delta F_\infty(d,g)$ is largest for small g at $d \in [42, 45]$ km. It decreases as (d,g) grows towards the feasible boundary, because the bounds of both schedules approach $\frac{1}{2}$ there, which leaves little room for them to differ. The red row at $g = 1$ marks the only place with $\Delta F_\infty(d,g) = -0.0008$, because F_{aux} needs to wait for another δ_∞ before it enters the BSMs. The parallel schedule has $\Delta g \in \{1, 2\}$ more than the sequential schedule at every d. At 40 km, the maximum available g is 21 for the parallel schedule and $g = 20$ for the sequential schedule. At 45 km, the maximum available g is 10 for the parallel schedule and $g = 9$ for the sequential schedule. The dotted band between the white and orange lines marks the (d,g) where only the parallel schedule gives a gain. At $g = 2$, $d_{\max} = 47.2$ km for the sequential schedule and $d_{\max} = 47.8$ km for the parallel schedule.



(a) $F_\infty(d,g)$ of Parallel (b) $\Delta F_\infty(d,g)$

Fig. 4: Parallel has Higher Fidelity Upper Bound



(a) ΔN_ζ, $\zeta = 10^{-2}$ (b) ΔT_ζ, $\zeta = 10^{-2}$

(c) ΔN_ζ, $\zeta = 10^{-3}$ (d) ΔT_ζ, $\zeta = 10^{-3}$

(e) ΔN_ζ, $\zeta = 10^{-4}$ (f) ΔT_ζ, $\zeta = 10^{-4}$

Fig. 5: Parallel needs Fewer Purification Rounds and Time

B. Number of Rounds and Convergence Time

Since $F_\infty(d,g)$ of the parallel schedule is higher, it thus has a larger gap to close and usually needs one more round than the sequential schedule to come within ζ of its own fidelity bound. Figs. (5) therefore compares both schedules at the same target $F_\infty^{\text{seq}} - \zeta$. Figs. (5(a))(5(c))(5(e)) show $\Delta N = N^{\text{seq}} - N^{\text{parallel}}$, where the blue parts mean the parallel schedule needs fewer rounds than the sequential schedule. At the light yellow part in all figures, $F_\infty^{\text{seq}}(d,g) - F_{E_0}^{(0)}(d) \le \zeta$, so both schedules reach the target without purification, because the gain of these areas can not exceed those tolerances. The light yellow band lies just below the orange line, covers $g \in [16, 19]$ at 40 km and $g \in [5, 8]$ at 45 km for $\zeta = 10^{-2}$ and almost disappears for $\zeta = 10^{-4}$. At the red part in Figs. (5(e))(5(f)) where appear only at $g = 1$, $F_\infty^{\text{parallel}}(d,g) < F_\infty^{\text{seq}}(d,g) - \zeta$, because Case 1 is the only case whose fidelity bound is lower than the sequential schedule, by about 0.0008 as shown in Fig. 4(b), when ζ is smaller than this gap, the target $F_\infty^{\text{seq}}(d,g) - \zeta$ lies above the fidelity bound of Case 1 and therefore cannot be reached.

For $\zeta = 10^{-2}$ in Fig. 5(a), both schedules need the same number of rounds, at most (d,g) and the parallel schedule saves only 1 round at small g for $d \in [40, 42]$ km. The purification round saving of the parallel schedule reaches up to 3 and 7 rounds for $\zeta = 10^{-3}$ and 10^{-4}, respectively, shown in Figs. (5(c))(5(e)). ΔN_ζ grows as ζ decreases, because the higher the fidelity bound, the more likely the parallel schedule is to pass the target while the sequential schedule still approaches its own fidelity bound slowly. The row $g = 1$ is green for $\zeta = 10^{-3}$ in Fig. 5(c), because the fidelity loss of Case 1 is smaller than ζ, so Case 1 still reaches the target but needs 1 to 3 more purification rounds to come close to its own fidelity bound. The same row is red for $\zeta = 10^{-4}$ up to 47.8 km in Fig. 5(e), because the loss of Case 1 now exceeds ζ.

Figs. (5(b))(5(d))(5(f)) show the time saving of the parallel schedule by $1 - \frac{T^{\text{parallel}}}{T_\zeta^{\text{seq}}}$. The largest time saving of the parallel schedule is 48.5%, 65.2% and 80.7% for $\zeta = 10^{-2}, 10^{-3}$ and 10^{-4}, respectively. For $\zeta = 10^{-3}$ in Fig. 5(d), all rows with $g \ge 2$ are dark green, while the row $g = 1$ is light green, because the extra purification rounds of Case 1 cancel most of the time saved by its shorter duration per round.

Together, Fig. (5) shows that the parallel schedule reaches the same fidelity target much faster than the sequential schedule when $g = 2$, this advantage grows with the required accuracy.

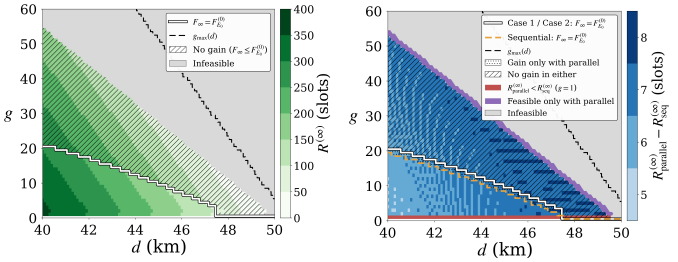
C. Residual Lifetime

We finally evaluate the residual lifetime rounded down to integer slots. Fig. 6(a) shows $R^{(k)}$ of the parallel schedule, which ranges from 7 to 361 slots. At every distance, $R^{(k)}$ is largest at $g = 2$, with 361, 189 and 112 slots at 40, 45 and 47 km, which is 3 or 4 slots more than at $g = 1$. For $g > 2$, $R^{(k)}$ decreases with g, to 333, 298 and 261 slots for $g = 10, 20$ and 30 at $d = 40$ km. It also decreases with d as $F_\infty(d, g)$ drops. Without purification, E_0 lives 294, 157 and 102 slots at 40, 45 and 47 km, respectively. Therefore, purification extends the lifetime of E_0 by up to about 70 slots below the while line and shortens it above. At the upper edge of the feasible region, $R^{(\infty)}$ exceeds the storage time w of one purification round by only a few slots, with 165 against 161 slots at $(d, g) = (40, 54)$ and 83 against 80 slots at $(d, g) = (45, 27)$, so this region ends because E_0 can no longer survive one more purification round under Eq. (17).

Fig. 6(b) shows $\Delta R = R^{(\infty)}_{\text{parallel}} - R^{(\infty)}_{\text{seq}}$ where both schedules are feasible. For $g \geq 2$, the parallel schedule keeps E_0 alive mostly more than the sequential schedule by 6 or 7 slots. The gap drops to 5 slots only at a few (d, g), with $g \leq 12$ and $d \in [40, 41]$ km, then it rises to 8 slots near the feasibility boundary for $d > 45$ km. The red row at $g = 1$ holds the (d, g) where the parallel schedule lives shorter than the sequential schedule for 1 or 2 slots. This is because the parallel schedule is in Case 1 at $g = 1$, where each auxiliary EPR pair waits δ_∞ slots before it is used and therefore has a lower F_{aux}, so Case 1 lives a bit shorter than the sequential schedule. In the purple band, the sequential schedule violates Eq. (17), so only the parallel schedule can run there, although without a gain. Eq. (17) requires the storage time w of E_0 in each purification round to be no longer than $R^{(\infty)}$. Compared with the sequential schedule, the parallel schedule stores E_0 for $T_R = 5$ fewer slots in each round and keeps it alive about 7 slots longer, so it can tolerate about 12 more slots of storage before this condition fails. Each additional attempt adds 3 slots to w and shortens $R^{(\infty)}_{\text{parallel}}$ by a few slots, so the margin between them shrinks by several slots as per g increases. The extra 12 slots therefore allow one or two more link generation attempt(s) before the parallel schedule becomes infeasible. Hence, for $g \geq 2$, the parallel schedule keeps E_0 alive about 7 slots longer and remains feasible for slightly larger g than the sequential schedule.

VI. Conclusion

We analyze the pumping purification over a LEO ISL with four qubits per satellite, where the next auxiliary EPR pair is generated while the current purification happens. We model this parallel schedule in two cases separated by $g^*(d)$ and compare its fidelity upper bound, the number of purification rounds, the purification time and residual lifetime with the sequential schedule. For $g \geq 2$, the parallel schedule raises the fidelity upper bound by about 0.004, saves up to 80.7% of the purification time for $\zeta = 10^{-4}$ and keeps E_0 alive about 6 or 7 slots longer. However, purification helps only for small g, since at 40 km it raises the fidelity only for $g \leq 21$, far below $g_{\max}(d) = 96$. The best choice is $g = g_c = 2$, which gives the longest residual lifetime of 361 slots at 40 km and extends the largest distance with a gain to 47.8 km.



(a) $R^{(k)}$ of Parallel (b) Parallel Lives Longer

Fig. 6: Parallel Schedule Lives Longer in Case 2

Our future work will be to analyze the consequence if g varies from pair to pair, so that different purification rounds of the same schedule may fall in different cases.

References

[1] Y. Gao, A. Valera, W. Li, and W. K. G. Seah, "Modeling Entanglement Purification in LEO Satellite Environment," *arXiv (under review)*, 2026.

[2] C. H. Bennett *et al.*, "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels," *Physical review letters*, vol. 76, no. 5, p. 722, 1996.

[3] N. K. Panigrahy, T. Vasantam, D. Towsley, and L. Tassiulas, "On the Capacity Region of a Quantum Switch With Entanglement Purification," in *IEEE INFOCOM*, New York City, NY, USA, 17-20 May 2023, pp. 1–10.

[4] M. Victora, S. Tserkis, S. Krastanov, A. S. de la Cerda, S. Willis, and P. Narang, "Entanglement Purification on Quantum Networks," *Physical Review Research*, vol. 5, no. 3, p. 033171, 2023.

[5] J. Liu *et al.*, "Joint Swapping and Purification With Failures for Entanglement Distribution in Quantum Networks," in *IEEE/ACM 33rd International Symposium on Quality of Service (IWQoS)*, Gold Coast, Australia, 2-4 July 2025, pp. 1–10.

[6] A. Zang, X.-A. Chen, E. Chitambar, M. Suchara, and T. Zhong, "Entanglement Purification in Quantum Networks: Guaranteed Improvement and Optimal Time," *arXiv preprint arXiv:2505.02286*, 2025.

[7] W. Li and S. Zhang, "Simulation of Quantum Repeater Networks under Decoherence and Purification Constraints," *arXiv preprint arXiv:2510.07471*, 2025.

[8] V. Vasan, A. Nico-Katz, B. A. Bash, D. C. Kilper, and M. Ruffini, "Entanglement Purification With Finite Latency Classical Communication in Quantum Networks," *arXiv preprint arXiv:2509.03667*, 2025.

[9] A. Fayyaz, P. Krishnamurthy, K. Seshadreesan, A. Babay, and D. Tipper, "To Purify or Not to Purify: Entanglement Purification under Input Fidelity Asymmetry in Quantum Networks," *arXiv preprint arXiv:2605.08771*, 2026.

[10] B. Davies, Á. G. Iñesta, and S. Wehner, "Entanglement Buffering With Two Quantum Memories," *Quantum*, vol. 8, p. 1458, 2024.

[11] Á. G. Iñesta, B. Davies, S. Kar, and S. Wehner, "Entanglement Buffering With Multiple Quantum Memories," *npj Quantum Information*, vol. 12, no. 1, p. 64, 2026.

[12] M. Mobayenjarihani, G. Vardoyan, and D. Towsley, "Optimistic Entanglement Purification in Quantum Networks," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 1143–1153.

[13] Y. Gao, A. Valera, and W. K. G. Seah, "Markov Chain Model of Entanglement Setup in Noisy Dynamic LEO Satellite Networks," *arXiv preprint arXiv:2512.20047*, 2025.

[14] A. Mollakhani, J. Tiamraj, S.-J. Cao, and D. Guo, "Inter-Satellite Link Configuration for Fast Delivery in Low-Earth-Orbit Constellations," in *IEEE Aerospace Conference*, Big Sky, MT, USA, 7-14 Mar 2026, pp. 1–9.