Upper bounds for the clock speeds of fault-tolerant distributed quantum computation using satellites to supply entangled photon pairs

Hudson Leone,[*] S. Srikara,[†] Peter P. Rohde,[‡] and Simon Devitt[§]

Center for Quantum Software and Information, University of Technology Sydney, Sydney, New South Wales 2007, Australia

(Received 3 October 2022; accepted 7 November 2023; published 29 December 2023)

Despite recent advances in quantum repeater networks, entanglement distribution on a continental scale remains prohibitively difficult and resource intensive. Using satellites to distribute maximally entangled photons (Bell pairs) between distant stations is an intriguing alternative. Quantum satellite networks are known to be viable for quantum key distribution, but the question of if such a network is feasible for fault tolerant distributed quantum computation (FTDQC) has so far been unaddressed. In this paper we determine a closed form expression for the rate at which logical Bell pairs can be produced between distant surface code encoded qubits using a satellite network to supply imperfect physical Bell pairs. With generous parameter assumptions, our results show that FTDQC with satellite networks over statewide distances (500–999 km) is possible up to a collective clock rate on the order of 1 MHz, while continental (1000–4999 km) and transcontinental (5000+ km) distances run on the order of 10 kHz and 100 Hz, respectively.

DOI: 10.1103/PhysRevResearch.5.043302

I. INTRODUCTION

It is well known that the Hilbert space of a quantum system grows exponentially with the number of qubits processed. This fact motivates research into quantum computer networking since multiple such devices working together is thought to have more computational power than the sum of its parts. One complication of communicating quantum information though is that conventional repeater networks cannot be used to amplify a quantum signal in transit. This is because unknown quantum information cannot be perfectly copied [1]. It is therefore necessary to use ultrareliable transmission strategies which ensure that a state is delivered with near certainty. One well established strategy is quantum state teleportation, which uses a bipartite entangled state distributed between two distant parties as a resource for one party to communicate a qubit of information to the other [2].

The quantum repeater was the first technology proposed for long distance entanglement distribution [3] but is presently considered infeasible for continental distances since it requires expensive repeater stations every 100 km or so [4–10]. A more viable alternative may be to use satellites to distribute maximally entangled photons (Bell pairs) between distant ground stations. The seminal Quantum Experiments at Space Scale (QUESS) project demonstrated that a quantum satellite today can distribute Bell pairs over distances of 1200 km at a rate around one kilohertz [11]. Additional theoretic work indicates that satellite networks perform suitably well for quantum key distribution (QKD) [12,13]. Unlike QKD, however, a fault-tolerant distributed quantum computation requires a continuous, high-volume supply of high-fidelity Bell pairs.

In this paper, we determine a closed form expression for the rate at which satellites can produce logical Bell pairs between distant error-corrected qubits. This in turn is the rate at which fault-tolerant quantum state teleportation can be performed and equivalently is the clock speed of the distributed quantum computer. Using generous parameter assumptions, we find that this clock speed is upper bounded on the orders of 1 MHz for state distances, 10 kHz for continental distances, and 100 Hz for transcontinental distances. Since the power available to a satellite naturally limits the rate at which entanglement can be supplied, this suggests long-term scalability issues for satellite based FTDQC. The choice of computational problem is incidental to our consideration of resource estimation, but, for the sake of completeness, we choose to consider RSA public key factorization using Shor's algorithm. We chose the surface code as our logical qubit encoding due to its high physical error tolerance, inexpensive two qubit operations via lattice surgery [14], and because it is the preferred encoding method for current hardware [15,16]. Similarly, we believe this choice of encoding is incidental to our final result since most of the Bell pairs are expended—not on the codes themselves, but on free-space attenuation and atmospheric effects. Based on the logical error tolerance required for this algorithm, we compute the required rate of high-fidelity Bell pairs needed to perform a lattice surgery operation (the fault tolerant parity check that is used to produce the logical Bell state). Working backwards, we calculate the rate of lower fidelity Bell pairs incident on the ground needed to perform entanglement purification with sufficiently high confidence. Finally, we incorporate free space and atmospheric

[*]leoneht0@gmail.com
[†]Srikara.Shankara@student.uts.edu.au
[‡]dr.rohde@gmail.com
[§]Simon.Devitt@uts.edu.au

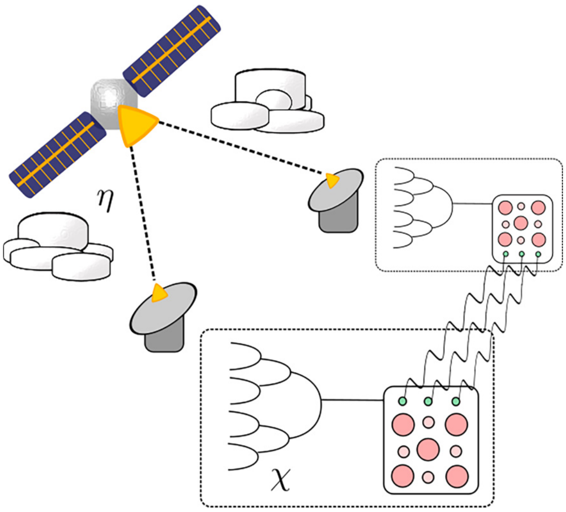

FIG. 1. Basic schematic of our scenario: a satellite in a constellation distributes entangled photon pairs between two distant ground stations. We assume an average double down-link attenuation η due to atmospheric and free-space effects with ideal weather conditions, and we assume perfect photon capture and conversion at the ground stations. To compensate for entanglement degradation, a nondeterministic recursion protocol is used to purify χ pairs into a single pair of sufficient quality. Finally, the purified pair is used to implement a fault-tolerant lattice surgery operation to create a logical Bell pair between the code patches.

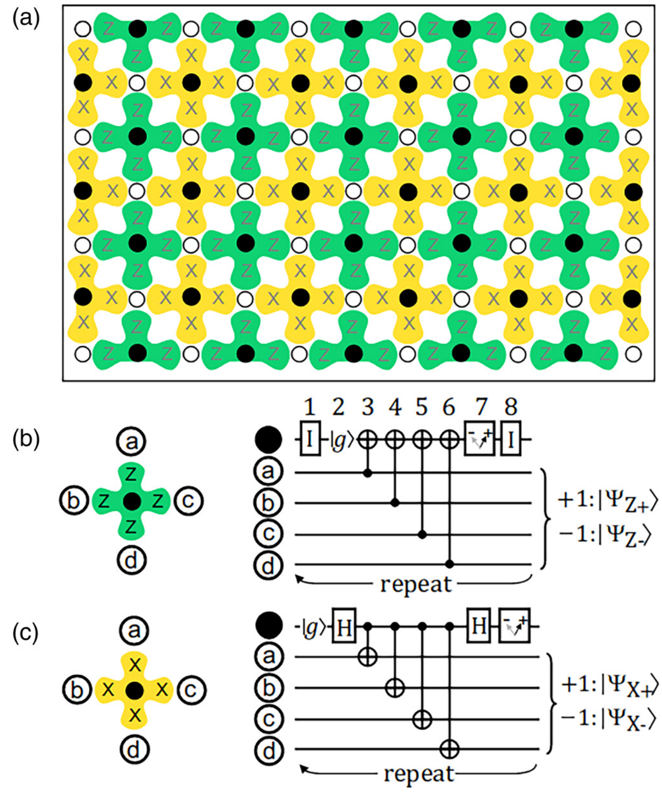

FIG. 2. (a) Top-down schematic of the surface code. The white (black) circles are the data (syndrome) qubits and the green (yellow) clover structures are the vertex (face) plaquettes. (b), (c) Syndrome extraction circuits for the two types of stabilizer generators for the surface code, respectively. Figure reproduced from Horsman *et al.* [15].

attenuation to estimate the rate of Bell pairs that a given satellite must be able to produce. Figure 1 shows a diagrammatic overview of our scenario. We convert this rate to the required satellite power assuming the use of the brightest available Bell source and with an estimate of the maximum power of a commercially available satellite. Finally, we calculate the creation rate of the logical Bell state.

II. INTRODUCING THE SURFACE CODE AND LATTICE SURGERY

We begin with a brief overview of the surface code and the lattice surgery operation. For a more complete description, we advise the reader to consult [14]. A surface code is a stabilizer code consisting of two types of qubits (differing only in their respective function) arranged in a checkerboard pattern [Fig. 2(a)]. Each lattice can encode one logical qubit and the larger the lattice is, the less likely the logical qubit is to suffer an undetectable error. The *data qubits* of the lattice encode the quantum information of the logical qubit, while the *syndrome qubits* are used to periodically measure the stabilizer generators of the code. This process of measuring the stabilizer generators is also known as *syndrome extraction*. By performing syndrome extraction, one gains partial information about any accumulated errors on the data qubits, which allows one to correct the encoded state with high probability. There are two types of syndrome extraction for the surface code, each of which can be represented as a five qubit circuit [Figs. 2(b) and 2(c)]. These circuits differ only with respect to their measurement bases and are shown diagrammatically as cloverlike tiles that cover the lattice. Importantly, all syndrome extraction circuits can be run in parallel, meaning the time it takes to implement a single *syndrome extraction cycle* is $6T$—the depth of the circuit multiplied by the average gate time of the architecture.

We couple two surface code qubits using a technique called *lattice surgery*. This is done by performing a syndrome extraction cycle over two code patches as if they were merged together (Fig. 3). We uncouple the state by measuring the qubits along the seam connecting the patches. This *merging* and *splitting* is equivalent to performing an XX parity measurement between the two logical qubits. When both codes are initialized as $|0\rangle_L$, the resulting state is a maximally entangled Bell pair. Notably, lattice surgery can be performed on surface codes that are *not directly adjacent*. Distributed Bell pairs can be used to teleport the two-qubit gates needed for the syndrome extraction circuits along the seam of the code (Fig. 3). What this means is that two parties can establish error-corrected entanglement between a pair of surface codes provided they have sufficiently many physical Bell pairs to perform the lattice surgery operation.

How many entangled pairs are required for lattice surgery and what is the rate at which they are required? These will depend on the size of the surface codes. It was previously mentioned that larger codes are more fault tolerant than smaller ones. This is quantified with *code distance*, which for a square lattice is roughly equal to the number of qubits on

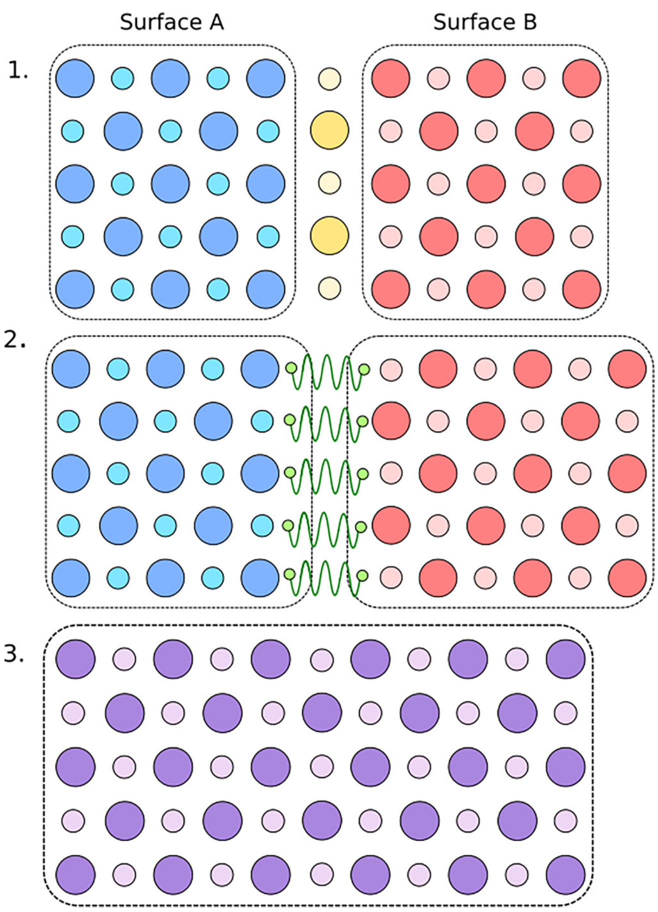

FIG. 3. Illustration of lattice surgery between two spatially separated surface codes A and B. (1) A buffer of syndrome and data qubits is initialized between the two surfaces for continuity of the qubit pattern. (2) The buffer qubits are merged into surface B and Bell pairs are delivered for each qubit pair on the boundary. (3) The syndrome extraction cycle of the code proceeds as normal. This joins the two patches into a single code.

either edge of the lattice. For two codes of distance D, each lattice surgery would ideally require D Bell pairs, as indicated in Fig. 3. One subtlety that must be addressed, however, is that measurement errors make syndrome extraction unreliable. In practice, we require D syndrome extraction cycles for each logical operation of the surface codes. This means that D^2 Bell pairs are required in total. Given that the time to implement one round of syndrome extraction is $6T$, we require $6TD$ units of time for the D rounds of fault-tolerant surgery. The corresponding production rate of logical pairs is therefore

$$R_{LP} \equiv \frac{1}{6TD}, \tag{1}$$

which, when multiplied by the required number of Bell pairs, gives us the rate of ideal physical pairs needed to sustain logical pair production at a rate of R_{LP}:

$$R_{IP} = D^2 R_{LP} = \frac{D}{6T}. \tag{2}$$

We advise the reader at this stage that Table V catalogs the definitions of constants used throughout the paper for easy reference.

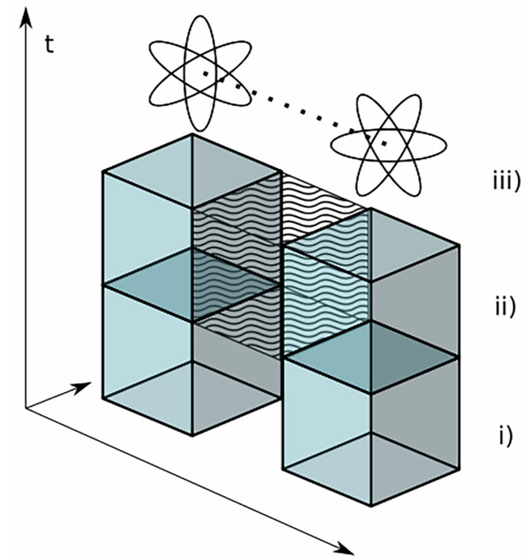

FIG. 4. Simplified space-time diagram for generating a logical Bell state between two surface codes. Each solid cube represents D cycles of syndrome extraction where D is the code distance. (i) Initialization of surface code qubits in the $|00\rangle_L$ state. (ii) Lattice merging with shared entanglement represented by the wavy cube connecting the two code patches. (iii) Lattice splitting, which requires a single round of syndrome extraction.

III. ESTABLISHING CODE DISTANCE

In general, we want to reduce the code distance in order to minimize the number of physical pairs needed for lattice surgery. We therefore aim to find the smallest possible code distance that is still sufficiently large to protect logical qubits in a practical instance of quantum computing. This first requires an understanding of how the code distance relates to the error rate of the encoded qubit. That relationship is given as

$$P_L = \alpha(\beta p)^{\frac{D+1}{2}}, \tag{3}$$

where p is the physical error rate and P_L is the logical error rate *per syndrome extraction* cycle. Devitt *et al.* [17] propose parameter values $\alpha = 0.3$ and $\beta = 70$ based on their numerical data which we will adopt as well. Recall that one round of syndrome extraction is unreliable due to measurement errors and that all surface code operations require D rounds of syndrome extraction for fault-tolerant execution. Assuming that P_L is small, we approximate the overall success rate for a surface code operation to the first order as

$$(1 - P_L)^D \approx (1 - DP_L). \tag{4}$$

The process of generating a logical pair from start to finish requires *four surface code operations* in total. This is most easily visualized with a space-time diagram as shown in Fig. 4. Here, each solid cube represents the D counts of syndrome extraction required for each operation. The two disjoint cubes at the base represent the preparation of the $|00\rangle_L$ state. The middle two cubes with the wavy cube in between represent a lattice surgery with shared entanglement on the code boundary. Finally, a splitting operation is done which (unlike the other operations) requires only one syndrome cycle and is therefore considered negligible. The total success rate of the logical pair preparation is therefore

$$(1 - DP_L)^4 \approx 1 - 4DP_L, \tag{5}$$

which means the failure rate of logical pair production as a function of code distance is

$$P_{LB} \equiv 4DP_L = 4D\alpha(\beta p)^{\frac{D+1}{2}}. \tag{6}$$

To determine a reasonable value for D, we need to substitute P_{LB} with the tolerable error rate for some nontrivial quantum circuit. For our calculations, we consider Shor's prime-factorization algorithm for RSA public key breaking. This is a well established benchmark with implications in cybersecurity, though itself has little bearing in the context of this paper. The circuit we consider for implementing this factorization is given by Beauregard [18], which in the case of 2048 bit factorization can tolerate a logical error rate of 4.28×10^{-21} [19]. Given $P_{LB} = 4.28 \times 10^{-21}$ and, assuming a physical error rate $p = 0.001$, we find that $D \approx 37$.

IV. ACCOUNTING FOR PURIFICATION

An entangled pair distributed through a noisy channel naturally loses some of its entanglement through decoherence. The same is true for an entangled photon pair passing through the atmosphere. This decay must be rectified with an entanglement distillation (purification) protocol, which takes a number of low-quality pairs and produces a high-quality pair with some overall success probability using local operations and classical communications [20]. For lattice surgery to succeed with satellite based entanglement communication, we require an additional resource overhead to account for the losses due to purification—not only because n pairs are converted into one, but also because the protocol is nondeterministic. On this latter point, we require the entire purification process to succeed with a rate S close to one. This means that given some initial quantity of imperfect pairs, we need to be sure up to confidence S that we can purify the required number of pairs needed for lattice surgery to succeed. This is done by *circuit multiplexing*, wherein many instances of a protocol are performed in parallel in order to improve the probability of a successful outcome. Circuit multiplexing is a common technique in linear-optical quantum computing where most operations are nondeterministic.

Note that entanglement purification cannot produce maximally entangled pairs in practice because of experimental uncertainty. For this reason, we define an *ideal pair* to be a state ρ_{AB} such that its *fidelity* with respect to a maximally entangled pair (for example $|\phi^+\rangle\langle\phi^+|$) is close to one:

$$F(\rho, |\phi^+\rangle\langle\phi^+|) \geqslant 1 - \epsilon, \tag{7}$$

$$F(\rho, \sigma) = [\mathrm{Tr}(\sqrt{\sqrt{\sigma}\rho\sqrt{\sigma}})]^2. \tag{8}$$

Let the *purification factor* χ be the number of nonideal pairs required to generate one ideal pair with confidence S. The rate of physical pairs required for lattice surgery at rate R_{LP} given some purification process is therefore

$$R_{IP+P} = R_{IP}\ \chi. \tag{9}$$

In general, the purification factor will depend on the initial state ρ_{in}, the required output fidelity F_{id}, the required success rate S, and the choice of purification protocol. Determining optimal purification protocols for arbitrary mixed states remains an open challenge so, for the sake of argument, we choose to consider the well-established *parity-check recurrence protocol* by Bennett *et al.* [21]. In the most optimistic scenario, this protocol takes two pairs of the form

$$\rho_0 = F|\phi^+\rangle\langle\phi^+| + (1 - F)|\phi^-\rangle\langle\phi^-| \tag{10}$$

and with probability $F^2 + (1 - F)^2$ returns a state of the same form with a new fidelity of

$$f(F) = \frac{F^2}{F^2 + (1 - F)^2}. \tag{11}$$

The reason the protocol is said to be a *recurrence protocol* is because the output pairs can be used as inputs for a subsequent round of purification. In this way, the protocol can be repeated until the target fidelity is reached. Let N be the minimum number of purification rounds needed to reach the threshold fidelity F_{id} from an initial fidelity of F_0.

Given the required number of purification rounds, what is the overall likelihood of the multiround protocol succeeding? Let $p(F)$ denote the success probability of a single $2 \to 1$ purification block

$$p(F) = F^2 + (1 - F)^2 \tag{12}$$

and let $F_0, F_1, \ldots, F_{N-1}$ denote the input pair fidelities for the respective purification rounds. Note that the kth purification round contains 2^{k-1} many $2 \to 1$ purification blocks that each need to succeed in order for the next round to go ahead. The overall success probability for the protocol is therefore

$$P \equiv p(F_0)^{2^{N-1}} \times p(F_1)^{2^{N-2}} \times \cdots \times p(F_{N-1})^{2^0}. \tag{13}$$

How many times must we multiplex a protocol with overall success probability P in order to guarantee one success with a confidence of at least S? Let $\mathcal{B}(P, k)$ be a binomial distribution where k is some number of trials and let $\Phi(B(P, k), x)$ denote the cumulative distribution function of the binomial up to x. The probability that at least one purification is successful is then

$$P_{\geqslant 1} \equiv 1 - \Phi(B(P, k), 1). \tag{14}$$

The minimum number of circuits needed to achieve an overall confidence S is therefore

$$K \equiv \min_k(P_{\geqslant 1} \geqslant S). \tag{15}$$

The purification factor is now the number of multiplexed circuits times the number of pairs needed for each circuit:

$$\chi = K \times 2^N. \tag{16}$$

Another important consideration for satellite based entanglement distribution is how the entanglement ought to be encoded into the photon pairs. One established and robust option is the polarization basis. Corroborating theoretical and experimental results indicate that polarization mode errors are comparatively small for atmospheric transmission [12]. For our study we let $F_0 = 0.87$, which is the collection fidelity reported by the Quantum Experiments at Space Scale

(QUESS) group [11]. For the sake of argument we set our target fidelity to be $F_{\mathrm{id}} = 0.999$ and our confidence rating as $S = 0.999$. With these parameters we find that $N = 2$ rounds of purification are sufficient to meet the target fidelity. This has a corresponding success rate of $P = 0.573$, which indicates $K = 9$ and $\chi = 36$.

V. ATMOSPHERIC AND FREE-SPACE ATTENUATION

Let η be the double down-link attenuation of a photon pair from a satellite or, in other words, the success rate of pair transmission. Our objective now is to determine the rate of photon pairs required to meet the R_{IP+P} rate needed for purification and lattice surgery, again with some confidence S. The number of pairs that reach the ground after k attempts is a random variable that follows a binomial distribution; however, since the attempt rate is very large ($k \gg 1$), we approximate this as a normal distribution with mean $k\eta$ and variance $k\eta(1-\eta)$:

$$\mathcal{N}(k\eta,\ k\eta(1-\eta)). \tag{17}$$

Similar to our calculations for the purification factor, the probability that at least R_{IP+P} photon pairs are transmitted given k attempts is the upper area of the probability density function from the point R_{IP+P}:

$$P_{\geqslant R_{IP+P}} \equiv 1 - \Phi(\mathcal{N}(k\eta,\ k\eta(1-\eta)), R_{IP+P}). \tag{18}$$

The required pair generation rate of the satellite, R_{PG}, is therefore the minimum value of k such that the probability of exceeding R_{IP+P} is greater than or equal to the success rate:

$$R_{PG} = \min_k (P_{\geqslant R_{IP+P}}). \tag{19}$$

A further simplification is possible using Markov's inequality, which gives an upper bound for the probability that a random variable X of some distribution is greater than a constant a:

$$P(X \geqslant a) \leqslant \frac{E(X)}{a}. \tag{20}$$

Adapting this inequality for our case by setting $X \to R_{PG}$, $a \to R_{IP+P}$, and $E(X) \to R_{PG} \times \eta$, we find that

$$S = P(R_{PG} \geqslant R_{IP+P}) \leqslant \frac{R_{PG}\,\eta}{R_{IP+P}}. \tag{21}$$

Therefore, if $S = 1 - \epsilon$ where $0 < \epsilon \ll 1$, we can approximate the pair generation rate as

$$R_{PG} \approx \frac{R_{IP+P}}{\eta}. \tag{22}$$

Substituting with Eqs. (9) and (2), respectively, we momentarily conclude by obtaining an expression that relates the photon pair rate to the logical pair rate:

$$R_{PG} = \frac{D^2 R_{LP}\ \chi}{\eta}. \tag{23}$$

Let us now focus our attention on choosing a suitable value for η. Much work has been done to characterize atmospheric attenuation for satellite based entanglement distribution. Bonato *et al.* developed a model for η in the context of quantum key distribution (QKD) [12] as did Mazzarella *et al.* [22] and Khatri *et al.* [13]. The results of Khatri *et al.* are especially relevant for our discussion, as they demonstrated the feasibility of a quantum satellite network for QKD and conducted extensive simulations on a 400 satellite constellation to find average down-link attenuations between major cities. We elected to use their three most optimistic attenuation rates for state, continental, and transcontinental distances, the values of which are presented in Table I. Due to the importance of these results in our calculations, a brief overview of their attenuation model and satellite constellation is provided for completeness.

The major contribution in down-link attenuation is attributed to beam widening. This has contributions from natural free-space widening and atmospheric diffraction. Although atmospheric diffraction contributes a higher widening rate, the team counterintuitively demonstrates that free-space widening is the more significant of the two contributions. This is because the atmospheric depth over which diffraction takes place is considerably smaller than the free-space distance. Beam wandering, where the median of the Gaussian profile shifts in the (x, y) plane, is known to be negligible for down-link transmission and so is not considered. Perfect weather conditions are assumed and the atmosphere is taken to be a homogeneous layer of constant density. The effects of ambient light are considered on the daytime regions of the globe since sunlight contributes a significantly higher signal to noise ratio and therefore decreases transmission efficiency.

The design objective of the constellation is to supply continuous global coverage with as few satellites as possible such that no double down-link channel ever exceeds 90 dB. The constellation consists of a number of equally spaced rings of satellites in polar orbits with an identical number of satellites in each ring. The optimal satellite configuration was decided as follows: in the worst case scenario, two ground stations each located on the equator are separated by some distance d. A constellation of 400 satellites was proposed which optimizes their figure of merit: the ratio of average entanglement distribution to the total number of satellites.

VI. REQUIRED PAIR GENERATION RATE AND SATELLITE POWER

In this section, we relate the pair generation rate R_{PG} to the required satellite power P_s. We assume all power is exclusively allocated to the task of pair production. The rate at which a satellite can generate pairs (S_{PG}) depends on the brightness of the source (N_p), the power consumption of each source (P_r), and the power available to the satellite (P_s):

$$S_{PG} = \frac{P_s N_p}{P_r}. \tag{24}$$

Rearranging this and setting $S_{PG} \to R_{PG}$ gives us the required satellite power for a particular generation rate:

$$P_s = \frac{R_{PG} P_r}{N_p}. \tag{25}$$

Substituting R_{PG} with Eq. (22), we obtain a closed form expression relating the satellite power to the rate of logical pair

TABLE I. Most optimistic average double down-link losses for three distance classifications over a 24 h period for a 400 satellite constellation [13].

Classification	City pairs	Average loss (dB)
State (500–999 km)	Toronto - New York City	45.1
Continental (1000–4999 km)	Sydney - Auckland	65.6
Transcontinental (5000 km+)	New York City - London	79.1

production:

$$P_s = \frac{D^2 R_{LP}\ \chi P_r}{N_p \eta}. \tag{26}$$

To determine a realistic upper bound for the maximum satellite power, we present a brief survey of the power ratings of Indian communication satellites in Table II. The most powerful of these, the GSAT-11, has a considerably larger power rating than the others but at a comparable budget. We therefore estimate that a satellite with a power rating of 10 kW is on the order of the most powerful commercial satellite possible with current technology.

The brightest available Bell-pair source reported at the time of writing is a waveguide integrated AlGaAson microresonator with a brightness of 20×10^9 Bell pairs $\mathrm{s}^{-1}\mathrm{mW}^{-2}$ [23]. Its high output combined with a micrometer scale form factor makes it a promising candidate as a satellite-based entanglement source. From the experimental data, the highest attainable rate reported was 4×10^6 pairs per second at a power of 15 μW. According to the team, increasing the power beyond this point would exceed the lasing threshold of the microresonator, which would reduce the overall entanglement visibility. With this information, we set N_P and P_r to the aforementioned values of brightness and power per source, respectively.

VII. RESULTS AND DISCUSSION

Let us begin by discussing the significance of estimating maximum achievable logical pair rates in the context of distributed quantum computation. The rate at which logical Bell pairs can be generated is essentially the *global clock speed* of a distributed quantum computer. More precisely, the logical operations of state teleportation, gate teleportation, and nonlocal two qubit operations are all rate limited by logical pair production. By estimating the maximum possible rate of logical pairs, we indicate the rate at which *nonlocal* operations between distant qubits can be performed.

TABLE II. Survey of Indian communication satellites launched between 2018 and 2019 with power ratings and costs.

Satellite name	Power	Budget (10 million USD)
GSAT-11	13.6 kW	7.43
GSAT-31	4.7 kW	6.46
GSAT-7A	3.3 kW	6.32–10.11
GSAT-29	4.6 kW	2.08
GSAT-30	6 kW	6.46

How does the global clock speed of a distributed quantum computer relate to its overall utility? This is a difficult (if not impossible) question to answer in the absolute sense, but in general we understand that a fast quantum computer is preferable over a slow one. Let us propose a thought experiment to resolve this insight with greater resolution. Suppose we have a quantum computer q and a quantum algorithm j that takes an integer as input. Let J be a fictitious oracle that takes a clock speed as input and returns the smallest integer such that the quantum algorithm j achieves supremacy on q (i.e., completes the computation faster than any existing classical computer). We expect that J is a continuous function since in principle supremacy is possible at any clock rate (provided one chooses a sufficiently big input) and we expect that J is monotonically increasing since it would be absurd for a slow quantum computer to achieve supremacy before a fast one. From these properties of J, we see that *fast quantum computers can reach supremacy with smaller computational problems than slow ones.* This means that fast quantum computers are likely to be useful for a *broad range* of problems, whereas slow quantum computers will only realize an advantage for very large calculations. Additionally, a slow quantum computer will require more physical resources than a fast one in order to achieve supremacy. It is for these two reasons that fast quantum computers are *strongly preferred* over slow ones. Although we cannot quantify the utility of a distributed quantum computer given its global clock speed, we can at least compare our estimates to the average gate times of a variety of *physical qubits* (Table III). In this way, we get an idea of how powerful our hypothetical distributed system is given our current understanding of what is possible with these contemporary systems.

Our numerical results are presented in Fig. 5. Here, we plot the generation rate of logical Bell pairs [as given by Eq. (26)] versus the required satellite power for state,

TABLE III. Sample of average gate times for common qubit architectures.

Architecture	Average gate time	Rate
Superconducting qubits [24]	50 ns	2×10^7 Hz
NV diamond [25]	0.05 μs	2×10^7 Hz
Ion trap [26]	1.6 μs	6.25×10^5 Hz
NMR spins [27]	1 ms	1×10^3 Hz
Distance category		Rate
State		2×10^6
Continental		1×10^4
Transcontinental		6×10^2

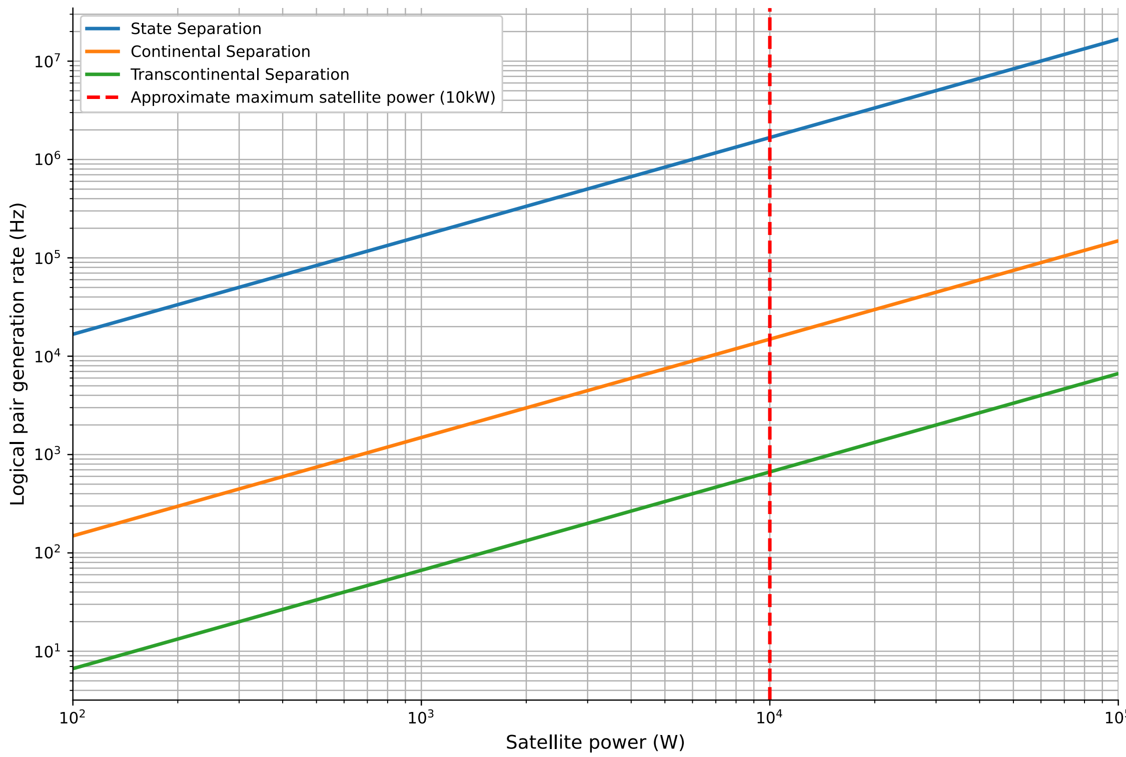

FIG. 5. Rate at which logical surface code Bell pairs can be generated versus the available satellite power [Eq. (26)] for three different distance ratings (Table I). The vertical dashed line indicates the approximate maximum power of a commercial satellite [29].

continental, and transcontinental distance ratings (Table I). Additionally, we plot a vertical dashed line indicating the approximate maximum power of a commercial satellite (10 kW). We estimate the fastest possible logical pair rates for each distance category by looking at where the three curves intersect this vertical. For state distances, this rate is around 2×10^6 s^{-1}. For continental distances, $\approx 1 \times 10^4$ s^{-1}, and for transcontinental distances, $\approx 6 \times 10^2$ s^{-1}. Let us compare these estimates with the average gate times of common qubit technologies (Table III). Here we see that the maximum achievable rate at the statewide distance is comparable to the rate of trapped-ion systems. Continental and transcontinental distances in turn are comparable to the speeds of an NMR-spin quantum computer.

With this in mind, we stress that these upper bounds are *far from realistically achievable* due to the numerous highly optimistic parameter assumptions and simplifications we made throughout this work. We treated photon capture and conversion as a lossless process and assumed that quantum memories and local operations are effectively noiseless. We treated incoming pairs with a special noise model to improve the purification rate. Our adapted model for double-down link attenuation assumes ideal weather conditions and we assume that 100% of a satellite's power can be allocated exclusively to photon pair production. We also point out that all resource estimation up to this point has been done with respect to a single pair of distributed surface codes, which suggests the communication infrastructure would need to scale in proportion to the number of distributed qubits.

Let us now consider the ways in which we might improve the performance of a quantum satellite network. Our options are to increase the throughput of the satellites, decrease the required pair generation rate, improve the efficiency of our purification, or reduce the relative attenuation of the downlink channel. In the first case, the only possibilities are to increase the satellite power or to improve the brightness of the photon pair source. It is unlikely that the power available to a commercial satellite will dramatically increase, though we suspect the brightness of entanglement sources will be improved at least an order of magnitude in the near-distant future. With respect to the pair generation rate [Eq. (22)], we find that the code distance contributes one order of magnitude and is almost certainly irreducible for the surface code since error-corrected logical qubits are required for distributed quantum computation. It is possible however that choosing another code type may yield a small advantage. In our work, we considered a parity check recurrence purification and showed the corresponding purification factor contributes one order of magnitude to the required pair rate. Given our optimistic noise assumptions, we do not consider it likely that our estimate of $\chi = 36$ will be dramatically improved, though this remains an open question.

It was brought to our attention that it may be possible to avoid nondeterministic recurrence purification by using a pair generation protocol that is different from the one depicted in Fig. 4. Here, the idea is to implement an error correcting code over multiple *imperfect* logical Bell states, which would eliminate the need to purify initial entanglement resources. For a specific example, we extend our gratitude to Gidney for providing us with a Stim implementation of a surface-code purification using a 5-qubit code [28]. One drawback with this approach, however, is that concatenating error-correcting

TABLE IV. Summary of selected parameter values with justifications.

Parameter	Value	Justification
D	37	Required code distance for FT RSA
F_0	0.87	QUESS Satellite data [11]
$F_{\rm id}$	0.999	Required fidelity rating
S	0.999	Requirement of surface code
χ	36	Best estimate
P_s	10 kW	Highest power rating for comms sat.
P_r	15 μW	Power of brightest Bell source [23]
N_P	4×10^6 s^{-1}	Throughput of Bell source [23]

codes increases the length of the space-time circuit which in turn decreases the rate at which logical pairs can be generated. A promising direction for further research is whether or not such a strategy could yield an advantage for logical pair production.

By far the most significant contribution to the required rate of photon pairs is the double down-link attenuation. The optimistic results we selected from Khatri *et al.* contribute between five and eight orders of magnitude depending on the distance between stations. As this loss results from free-space transmission and atmospheric diffraction, there are few options for mitigating this effect. One notable strategy is to store half of a pair on the satellite as the other transmits rather than sending both halves at the same time. We credit Shaw with this idea. At first this seems like only a superficial difference, but it turns out that this trick can effectively *halve the attenuation rate*. This is because when half of a pair is established on the ground, the other half can be used for *entanglement swapping* with another established pair (albeit with a success probability of 50%). A significant disadvantage with this approach though is that the satellite must reliably control, process, and measure an enormous number of pairs, which is presently infeasible for a satellite.

Up to this point we considered a down-link transmission model where entanglement is generated in satellites and distributed between ground stations. The reverse case is up-link transmission where entanglement pairs are prepared on the ground and fired up to a satellite which performs entanglement swapping to project the pair between the stations. The main advantage of this approach is that ground stations can generate significantly more photon pairs since power is no longer a major limiting factor. The downside, however, is that the attenuation rate for up link is significantly higher than for down link. This is because beam wandering effects from atmospheric turbulence are more significant when the Gaussian profile is small. Engineering challenges are another difficulty for these hypothetical networks. Unlike down-link satellites, which are relatively passive, the up-link satellite would be required to control and measure incoming photon pairs with quantum precision. This is a demanding task even for a laboratory on Earth. Whether or not these trade-offs are overall beneficial remains to be seen, though our group is currently investigating this in greater detail.

TABLE V. Table of the most significant constants with definitions.

Constant	Definition
D	Surface code distance
R_{LP}	Logical pair generation rate
R_{IP}	Rate of ideal Bell pairs required to produce logical pairs at rate R_{LP}
P_{LB}	Tolerable error rate of logical pair production
S	The likelihood that a given routine succeeds (also called the confidence)
χ	Number of imperfect pairs needed to purify down to one ideal pair with confidence S
R_{PG}	Required photon pair rate to generate logical pairs at rate R_{LP}
η	Double down-link attenuation rate (the success rate of a photon-pair transmitting)
P_s	Satellite power
P_r	Power consumption per photon pair source
N_P	Brightness of Bell source (pairs per unit time)

We believe our results indicate a need to reconsider the problem of long-distance entanglement distribution. One understated consideration of quantum networks is that they have *no latency requirements*. Unlike classical data networks, entanglement can be stored as a physical resource and transported by moving error corrected memory units. This is the motivating principle of the *quantum sneakernet* [17], which may be a more viable long-term alternative to quantum satellite networks. We note that the quantum sneakernet has its own significant engineering challenges since it is predicated on the aspiration it is possible to make a sufficiently portable and scalable quantum memory unit that can be transported long distances. Given the alternatives, however, we feel that the sneakernet has a greater potential for scalability than any quantum satellite network. This is because satellites will always be rate limited by power consumption, which is not a problem in the sneakernet model.

In this work, we determined upper bounds for the rates of logical pair generation between surface code qubits at a variety of distances when the physical entanglement is supplied by a quantum satellite network. We began by establishing a reasonable code distance that was predicated on an arbitrarily "hard" computational problem. We accounted for losses in entanglement purification by considering a $2 \rightarrow 1$ parity check recurrence protocol under an optimal noise model. In order to calculate our estimates for attainable pair rates, we also developed a closed form equation relating the available satellite power to the achievable logical pair production rate.

ACKNOWLEDGMENTS

The authors wish to thank K. Kuntz, T. Jennewein, and P. Avila for their feedback and insights—in particular for introducing us to the possibility of up-link satellite networks as an avenue for further research. We thank C. Gidney for his suggestion on avoiding nondeterministic recurrence purification and providing us with a proof of concept [28]. Finally, the student authors wish to thank our colleagues and supervisors for attending our seminar regarding this paper at the UTS Centre for Quantum Software and Engineering. This research was conducted by the Australian Research Council Centre of Excellence for Engineered Quantum Systems (Project No.

CE170100009) and funded by the Australian Government. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. This research was developed with funding from the Defense Advanced Research Projects Agency [under the Quantum Benchmarking (QB) program under Awards No. HR00112230007 and No. HR001121S0026 contracts]. H.L. was supported by the ARC Centre of Excellence for Quantum Computation and Communication Technology (CQC2T), Project No. CE170100012.

APPENDIX: TABLES OF CONSTANTS

This appendix contains Tables IV and V which, respectively, summarize our experimental parameters and catalogue the most important constants used throughout the paper.

[1] W. K. Wootters and W. H. Zurek, A single quantum cannot be cloned, Nature (London) **299**, 802 (1982).

[2] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels, Phys. Rev. Lett. **70**, 1895 (1993).

[3] Q. Ruihong and M. Ying, Research progress of quantum repeaters, J. Phys.: Conf. Ser. **1237**, 052032 (2019).

[4] A. G. Fowler, D. S. Wang, C. D. Hill, T. D. Ladd, R. Van Meter, and L. C. L. Hollenberg, Surface code quantum communication, Phys. Rev. Lett. **104**, 180503 (2010).

[5] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, Quantum repeater with encoding, Phys. Rev. A **79**, 032325 (2009).

[6] Y. Li, S. D. Barrett, T. M. Stace, and S. C. Benjamin, Long range failure-tolerant entanglement distribution, New J. Phys. **15**, 023012 (2013).

[7] S. Muralidharan, J. Kim, N. Lütkenhaus, M. D. Lukin, and L. Jiang, Ultrafast and fault-tolerant quantum communication across long distances, Phys. Rev. Lett. **112**, 250501 (2014).

[8] K. Azuma, K. Tamaki, and H.-K. Lo, All-photonic quantum repeaters, Nat. Commun. **6**, 6787 (2015).

[9] W. J. Munro, K. A. Harrison, A. M. Stephens, S. J. Devitt, and K. Nemoto, From quantum multiplexing to high-performance quantum networking, Nat. Photon. **4**, 792 (2010).

[10] S. J. Devitt, W. J. Munro, and K. Nemoto, High performance quantum computing, Prog. Inform. **8**, 49 (2011).

[11] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai, G.-B. Li, Q.-M. Lu, Y.-H. Gong, Y. Xu, S.-L. Li, F.-Z. Li, Y.-Y. Yin, Z.-Q. Jiang, M. Li, J.-J. Jia *et al.*, Satellite-based entanglement distribution over 1200 kilometers, Science **356**, 1140 (2017).

[12] C. Bonato, A. Tomaello, V. Da Deppo, G. Naletto, and P. Villoresi, Feasibility of satellite quantum key distribution, New J. Phys. **11**, 045017 (2009).

[13] S. Khatri, A. J. Brady, R. A. Desporte, M. P. Bart, and J. P. Dowling, Spooky action at a global distance: Analysis of space-based entanglement distribution for the quantum internet, npj Quantum Inf. **7**, 4 (2021).

[14] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, Surface code quantum computing by lattice surgery, New J. Phys. **14**, 123011 (2012).

[15] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).

[16] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, Nature **605**, 669 (2022).

[17] S. J. Devitt, A. D. Greentree, A. M. Stephens, and R. Van Meter, High-speed quantum networking by ship, Sci. Rep. **6**, 36163 (2016).

[18] S. Beauregard, Circuit for Shor's algorithm using 2n+3 qubits, Quantum Inf. Comput. **3**, 175 (2003).

[19] J. Li, J. Lee, and J. Heo, Resource analysis of quantum computing with noisy qubits for shor's factoring algorithms, Quantum Inf. Process. **21**, 2 (2022).

[20] W. Dür and H. J. Briegel, Entanglement purification and quantum error correction, Rep. Prog. Phys. **70**, 1381 (2007).

[21] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, Purification of noisy entanglement and faithful teleportation via noisy channels, Phys. Rev. Lett. **76**, 722 (1996).

[22] L. Mazzarella, C. Lowe, D. Lowndes, S. K. Joshi, S. Greenland, D. McNeil, C. Mercury, M. Macdonald, J. Rarity, and D. K. Li Oi, Quarc: Quantum research cubesat—a constellation for quantum communication, Cryptography **4**, 7 (2020).

[23] T. J. Steiner, J. E. Castro, L. Chang, Q. Dang, W. Xie, J. Norman, J. E. Bowers, and G. Moody, Ultrabright entangled-photon-pair generation from an AlGaAs-on-insulator microring resonator, PRX Quantum **2**, 010337 (2021).

[24] A. Noguchi, A. Osada, S. Masuda, S. Kono, K. Heya, S. P. Wolski, H. Takahashi, T. Sugiyama, D. Lachance-Quirion, and Y. Nakamura, Fast parametric two-qubit gates with suppressed residual interaction using the second-order nonlinearity of a cubic transmon, Phys. Rev. A **102**, 062408 (2020).

[25] Y. Chou, S.-Y. Huang, and H.-S. Goan, Optimal control of fast and high-fidelity quantum gates with electron and nuclear spins of a nitrogen-vacancy center in diamond, Phys. Rev. A **91**, 052315 (2015).

[26] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, Fast quantum logic gates with trapped-ion qubits, Nature (London) **555**, 75 (2018).

[27] D. Lu, A. Brodutch, J. Park, H. Katiyar, T. Jochym-O'Connor, and R. Laflamme, NMR quantum information processing, *Electron Spin Resonance (ESR) Based Quantum Computing*, edited by T. Takui, L. Berliner, and G. Hanson (Springer New York, 2016), pp. 193–226.

[28] C. Gidney, 5 qubit stim distillation circuit, https://gist.github.com/Strilanc/40c4607b8b95165c92734eeced27d450.

[29] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.5.043302 for a Python script for producing Fig. 5.