

Satellite-Based Entanglement Distribution between Ground-Based Memories for Coordination Applications

Sergio Loarte^{*a,b,c}, Janice van Dam^{a,b,c}, and Stephanie Wehner^{a,b,c}

^aQuTech, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

^bEEMCS, Quantum Computer Science, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands

^cKavli Institute of Nanoscience, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

ABSTRACT

Long-distance entanglement distribution between ground-based quantum memories enables coordination applications, in which distant users must produce correlated responses without communicating once the task has started. They are valuable where classical signalling latency makes real-time coordination impractical: examples include distributed scheduling, frequency-band rendezvous against jammers, and latency-bound trading. We study a satellite distributing entanglement to two ground users over continental distances. As a benchmark for these applications we evaluate the Clauser-Horne-Shimony-Holt (CHSH) game, where measuring sufficiently high-fidelity shared entanglement wins with higher probability than any classical strategy achieves. We present an end-to-end simulation framework, from entanglement distribution to the CHSH game played at the users' memories, coupling a satellite link model to NetSquid, a discrete-event quantum network simulator. Stored pairs are discarded past a cutoff age, since decohered entanglement no longer provides an advantage. Requests to play arrive at random times and are served with entanglement only if a stored pair is available, otherwise by the optimal classical strategy. Increasing the cutoff raises availability but admits older, lower-quality pairs, a trade-off that determines how large the advantage is. We showcase a European scenario: trapped-ion memories in Delft and Innsbruck served by a low-Earth-orbit satellite. Currently demonstrated memory coherence times limit entanglement availability, resulting in a negligible quantum advantage. A follow-up study at an optimistic coherence time maps the advantage over cutoff times and request rates. The advantage peaks at an intermediate balance of cutoff and request rate, reaching ~ 0.4 percentage points above the classical bound consistently across simulated passes.

Keywords: satellite quantum communication, entanglement distribution, coordination applications, CHSH game, quantum memory, NetSquid, link budget, quantum networks

1. INTRODUCTION

Imagine a scenario where two parties at distant sites must each choose an action, each seeing only its own local input, and the two choices are useful only if they are correlated in a prescribed way. Once the task has started the parties cannot communicate, or do not have time to, so neither learns what the other saw or what it intends to do. Shared entanglement improves what the two parties can achieve, as local measurement of high-fidelity entanglement produces correlated responses that beat any classical strategy, without real-time communication, obtaining a *quantum advantage*. These tasks are well suited to early-stage quantum networks, whose entanglement supply is limited: one shared pair per instance is enough, and that pair only needs to survive from its storage into memory until the task runs. Entanglement-assisted formulations for this family of low-latency coordination applications are found in distributed scheduling and routing under latency constraints,¹ frequency-band rendezvous against distributed jammers,² and latency-bound coordination such as high-frequency trading.³

Further author information:

*Corresponding author: Sergio Loarte, E-mail: s.loarte@tudelft.nl

These applications therefore require shared entanglement between distant ground-based quantum memories, which buffer entangled states until an application request consumes them.⁴ For users located far apart, the no-communication condition stems naturally from the latency introduced by the distance, of the order of milliseconds at continental separations. Distributing entanglement over such distances can be achieved using a satellite: free-space link loss grows only polynomially with distance, whereas fibre loss grows exponentially. This was demonstrated by the Micius satellite, with dual-downlink entanglement distribution over more than a thousand kilometres.^{5,6}

As a benchmark for the coordination class we adopt the Clauser–Horne–Shimony–Holt (CHSH) game.⁷ Two players, Alice and Bob, receive uniformly random bits x and y and must output bits a and b , without communicating, such that $a \oplus b = x \wedge y$. The best classical strategy wins with probability $3/4$: for example, both players always output 0, which wins whenever $x \wedge y = 0$, i.e. on three of the four input pairs. Measurements on a shared unit-fidelity entangled pair win with probability up to $\cos^2(\pi/8) \approx 0.854$, the Tsirelson bound.⁸ The CHSH game is therefore a valid proxy to evaluate the presence of a quantum advantage for this family of non-local coordination applications. The users need the delivered entanglement to be *present* when requested and *good enough* when used.

We capture this dual requirement in a single figure of merit. When a game request finds a usable stored pair, the players measure it; otherwise they fall back to the optimal classical strategy. The global win rate is

$$P_{\text{win}} = A P_q + (1 - A) \frac{3}{4}, \quad (1)$$

where the availability A is the fraction of requests met with a usable stored pair and P_q is the win probability when the players measure the delivered noisy pair. We count only those requests issued while the satellite is co-visible from both stations. Continuous service coverage could be achieved using multiple satellites, which is beyond the scope of this work.

Two mechanisms limit A : the link delivers pairs at a low rate, so a memory is often empty when a request arrives, and a stored pair decoheres while it waits. Once a pair has aged too far, its win probability falls to $3/4$, where measuring it is no better than the classical fallback. Because of this, a pair older than a chosen cutoff age t_c is discarded. This exposes a tradeoff: a larger t_c raises A , because a stored pair stays eligible for longer, but it also admits older and more decohered pairs, which lowers P_q . Our goal is to reach $P_{\text{win}} > \frac{3}{4}$, i.e. to gain a quantum advantage. We quantify whether, and under which operating choices, satellite-delivered entanglement between two memory-equipped users achieves this.

To achieve this, we present an end-to-end simulation framework that spans this chain from orbital physics to application outcome. A free-space link model computes the elevation-dependent losses and the per-pair state quality along a satellite pass, and emits a timestamped stream of stored entangled pairs. The NetSquid discrete-event simulator⁹ then evolves the in-memory decoherence of those pairs and plays the CHSH game against a Poisson request process: requests to play arrive at random times, which the players do not know in advance and cannot influence. We validate the link model against published Micius dual-downlink measurements, and we showcase the full chain on a European use case: a Micius-like satellite serving ground-based trapped-ion memories in Delft and Innsbruck over a single pass. For this scenario we characterise the delivered pair rate, contrast the published trapped-ion coherence time with an optimistic near-term value, and map the request-rate and cutoff plane, which shows that the advantage peaks at an intermediate balance of both quantities.

The paper is organised as follows. Section 2 describes the simulation model: the system architecture, the satellite link budget and its validation against Micius, the ground-segment memory interface and delivered state, and the event-driven CHSH application. Section 3 defines the baseline European scenario and characterises the delivered pair rate. Section 4 reports the coherence-limited CHSH results and the request-rate and cutoff policy map. Section 5 concludes. Appendix A collects the detailed loss and fidelity terms, Appendix B the Micius validation, and Appendix C the CHSH measurement protocol.

2. SIMULATION MODEL

2.1 System Architecture and Dataflow

The modelled system is a single low-Earth-orbit satellite carrying an entangled-photon-pair source based on spontaneous parametric down-conversion (SPDC). The satellite distributes the two photons of each pair through two simultaneous free-space downlinks to two distant optical ground stations. Each ground station then couples its downlink photon into a quantum memory.

We adopt the downlink rather than the uplink configuration. It keeps the demanding hardware (the single-photon detectors and the quantum memories) on the ground and leaves the satellite carrying only a comparatively simple and radiation-tolerant source. It also avoids the loss penalty of an uplink, in which turbulence disturbs the beam at the start of its path and the resulting spread grows over the full propagation distance, whereas a downlink beam meets the atmosphere only in the last few tens of kilometres.⁶




Figure 1. Main modules of the simulation package (blocks) and their dataflow (arrows).

The simulator represents this architecture as a modular pipeline (Fig. 1). Configuration files, two-line elements (TLEs), and station coordinates drive an orbit propagator. The propagator supplies the elevation and the slant range (the line-of-sight distance from a ground station to the satellite) to the link budget (Sec. 2.2). The fidelity module then combines the link budget, the receiver-coupling factor, and the photon-memory-scheme factors (Sec. 2.3) into the per-pair delivered state, modelled as a Werner state. The satellite-link module emits timestamped entangled-pair events, each carrying the full density matrix of the delivered pair, to the on-ground CHSH application. Everything from the moment of storage onwards runs inside a discrete-event NetSquid simulation (Sec. 2.4), which evolves the decoherence of each stored pair until it is consumed or discarded. A generator produces the application's requests, and the runner schedules, collects, and post-processes the outcomes.

Execution is split into two stages. The orbital and optical physics (propagation, the elevation-dependent link budget, and the per-pair fidelity) is evaluated once per pass, resulting in the timestamped event list passed to NetSquid. Separating the deterministic physics pre-computation from the stochastic event-level run keeps every run reproducible under a fixed seed. It also makes parameter sweeps over the application and memory layer

inexpensive, because a single orbital pass is replayed under many request rates, cutoff policies, and memory configurations without recomputing the physics.

2.2 Satellite Link Budget

The free-space optical channel sets how many of the emitted photons reach the ground stations. We model each downlink arm as a chain of multiplicative loss factors that vary along the pass:

- **Geometric (diffraction) loss.** We treat each downlink as a Gaussian beam leaving the transmitter aperture at its waist w_0 . For a downlink wavelength λ , the beam radius grows with slant range z as $w(z) = w_0 \sqrt{1 + (z/z_R)^2}$, with Rayleigh range $z_R = \pi w_0^2/\lambda$. The ground telescope collects the power falling inside its aperture, a fraction $1 - \exp(-2r^2/w(z)^2)$ for aperture radius r on a centred beam. Since $w(z)$ far exceeds the aperture over the whole pass, this is the dominant loss term.
- **Pointing loss.** Residual satellite pointing jitter, relative to the beam divergence, further reduces the collected fraction. We use the far-field jitter model of Titterton,¹⁰ $\eta_p = [1 + 16(\sigma_p/\theta_{\text{div}})^2]^{-1}$, where σ_p is the RMS pointing error of the satellite terminal and $\theta_{\text{div}} = 2\lambda/(\pi w_0)$ the full diffraction-limited divergence of the beam.
- **Atmospheric extinction.** Molecular absorption and Rayleigh/Mie scattering are integrated along the slant path, from layerwise extinction profiles generated with the libRadtran radiative-transfer package.¹¹ libRadtran is open source and runs inside the model, allowing us to regenerate the profiles for the downlink wavelength and for the observing conditions of each ground station location.

Together these factors give the per-arm free-space efficiency η_{link} , to which we add the optical throughput of the satellite and ground-station telescopes. The ground receiver enters in two stages, each a fixed per-arm loss: the telescope and its adaptive-optics relay optics,¹² followed by coupling of the corrected wavefront into single-mode fibre.¹³ As our focus is a system-level budget rather than a real-time one, we do not model the turbulence-induced fluctuations in the received power either.

Appendix A lists the parameter values assumed. The final rate of successfully stored entangled pairs is determined by combining this free-space link budget with the ground-segment chain discussed in Sec. 2.3.

Validation against Micius. We validated the satellite link model against the simultaneous Micius downlinks to Delingha and Nanshan reported by Yin et al.⁵ Their Fig. 2 provides the satellite–station slant ranges and the measured end-to-end attenuation of the two downlinks. We recreated the same pass and reproduced both metrics. The comparison and details are in Appendix B.

2.3 Ground Segment: Memory Interface and Delivered State

The downlink photon enters the memory through cavity-assisted photon scattering (CAPS)¹⁴, in which it reflects off a single atom in an optical cavity and becomes entangled with it deterministically. The reflection efficiency follows from the cavity cooperativity C alone.¹⁵ We choose cavity reflection over emission-based interference schemes,¹⁶ which require a probabilistic Bell-state measurement and photon indistinguishability on each arm, and therefore often deliver lower rate and fidelity in this setting.

The application and the cavity scheme jointly set the memory requirements: deterministic high-efficiency readout, since a player who fails to read out cannot signal the failure in time and the round is played blind; the longest possible coherence time, since a stored pair must survive until a request arrives; compatibility with a cavity resonance; and an efficient, frequency-convertible photon interface.

The resulting per-arm photon–memory-scheme efficiency is $\eta_{\text{sch}} = \eta_{\text{CAPS}} \eta_{\text{det}} \eta_{\text{QFC}}$, the product of the CAPS reflection, the detector efficiency, and the QFC chain efficiency.^{15,17} The heralding detector sits *after* the cavity, so it operates at the ion transition rather than at the downlink wavelength; we take its efficiency and its dark-count rate from the same detector-bias point.¹⁸

Delivered state fidelity. The dominant downlink fidelity-loss mechanism is noise photons collected by the receiving telescope and spurious detector clicks, which degrade the pair fidelity. The per-arm signal weight w is

$$w = \frac{1}{1 + \bar{n}/\eta_{\text{tr}}}, \quad \bar{n} = n_{\text{sky}} + n_{\text{dc}}, \quad (2)$$

where $\eta_{\text{tr}} = \eta_{\text{link}} \eta_{\text{rx}} \eta_{\text{cpl}}$ is the per-arm transmission, combining the free-space link budget of Sec. 2.2 with the receiver optics and fibre-coupling losses, and $\bar{n}$ is the total mean noise-photon number per detection window. It has two contributions. The sky-background term is¹⁹

$$n_{\text{sky}} = \frac{H \Omega \pi (D_{RX}/2)^2 \Delta \lambda \Delta T}{hc/\lambda}, \quad (3)$$

with sky spectral radiance H , field of view Ω , receiver aperture D_{RX} , filter bandwidth $\Delta \lambda$, and coincidence window ΔT . The dark-count term converts the detector's spontaneous click probability $p_{\text{dc}} = r_{\text{dc}} \Delta T$ into an equivalent noise-photon number, $n_{\text{dc}} = N_{\text{det}} p_{\text{dc}} / (1 - p_{\text{dc}})$, with $N_{\text{det}} = 1$ for the cavity-reflection interface, which heralds on one detector per arm. Appendix A gives the baseline values.

The delivered pair is then modelled as a Werner state whose parameter combines the source quality with both arms' signal weights and the memory-scheme Werner factors,

$$w_{\text{final}} = w_{\text{src}} (w_A w_{\text{sch},A})(w_B w_{\text{sch},B}), \quad (4)$$

where w_{src} is set by the source fidelity, w_{sch} collects the CAPS gate and readout fidelities, and w_A and w_B are the per-arm free-space signal weights of Eq. (2). With w_{final} , the full 4×4 density matrix is constructed in the $|\Psi^-\rangle$ basis, the Bell state produced by the satellite links.

2.4 Event Generation and the CHSH Application

The link budget and the memory interface together determine how many pairs are stored, and when. The orbit is sampled on a one-second grid, and the elevation-dependent attenuation is held constant over bins of duration Δt_{bin} ; a binning study over both dynamic link models found the delivered-pair statistics unbiased to within 0.2% for bins up to 60 s. Per bin, the number of generated pairs is drawn from a Poisson distribution with mean

$$\mu = R_{\text{src}} \eta_{\text{eff},A} \eta_{\text{eff},B} \Delta t_{\text{bin}}, \quad (5)$$

where R_{src} is the on-board pair-generation rate and $\eta_{\text{eff}} = \eta_{\text{tr}} \eta_{\text{sch}}$ is the per-arm end-to-end efficiency, combining the transmission link budget (Sec. 2.2) with the photon-memory-scheme efficiency (Sec. 2.3). Each pair is assigned a delivered state, and the result is a timestamped entanglement-event list handed to the discrete-event stage.

Discrete-event stage. The timestamped event stream feeds the state insertion in NetSquid,⁹ simulating the delivery of each pre-computed pair state into the two memory nodes at its scheduled time. From that moment the stored qubits decohere under the memory's T_1/T_2 noise model until they are consumed or discarded. Each node can host k memory slots. In this work we study a single memory per node, $k = 1$, under the assumption that a pair generated while the slot is occupied is not stored.

Per-pass operation couples the physics to the application. The two stations sit at different ranges, so the two photons of a pair have different times of flight: the first-arriving qubit begins to age in its memory while the other one is still in flight, and the pair becomes usable only once both photons are heralded. Every request is answered: if both stations hold an entangled pair within the cutoff time, they play the round with it; if either qubit is missing or has expired, they fall back to the optimal classical strategy. A stored pair that reaches its cutoff before a request arrives is discarded. Details on the CHSH game are given in Appendix C.

3. BASELINE EUROPEAN SCENARIO

As the baseline use case we take a Micius-like satellite distributing entanglement to ground stations in Delft and Innsbruck, each equipped with a ground-based quantum memory. These locations are 731 km apart, beyond the reach of direct-fibre entanglement distribution, but close enough that a single low-Earth-orbit satellite sees both stations at once on regular passes, a condition that any single satellite equipped only with a source must satisfy to deliver entanglement.

The satellite follows the Micius orbit (TLE epoch June 2020), and we simulate the best common-elevation pass found for this station pair, beginning 2020-06-16 at 23:44 UTC and lasting 420 s. Figure 2 shows the slant range from the satellite to each station over the pass. The shaded band marks the co-visibility window, the interval in which both stations see the satellite above the minimum elevation, chosen here as 10° . Entanglement can be delivered during this window, with slant ranges varying between roughly 600 and 2000 km. On the ground, each memory is a $^{40}\text{Ca}^+$ ion in an optical cavity.²⁰ Trapped-ion memories meet all four requirements of Sec. 2.3, and $^{40}\text{Ca}^+$ ions in particular already serve as nodes in practical quantum-network experiments: light-matter entanglement has been distributed over 50 km of optical fibre²¹ and between ion qubits 230 m apart,²² telecom conversion preserves the entangled state,^{17,23} and readout reaches fidelities above 0.999.²⁴ Table 1 summarises the scenario and hardware parameters. Table 3 in Appendix A gives the extended model-parameter set.




Figure 2. Slant range from the satellite to the Delft and Innsbruck ground stations over the selected pass (2020-06-16, 23:44 UTC).

Figure 3 shows the delivered-pair rate across the pass: the per-second number of pairs that survive both downlinks and both memory interfaces. This is the supply reaching the memories before the application starts consuming it. The rate follows the elevation-driven link budget, rising from zero as the co-visibility window opens and peaking at 0.77 pairs per second near the satellite’s closest approach. Over 1000 seeds, one pass delivers on average 97.2 pairs.




Figure 3. Delivered-pair supply over the Delft–Innsbruck pass. Top: the link-budget rate per 10s attenuation bin. Bottom: cumulative delivered pairs; the line is the mean over 1000 seeded passes and the band spans their 10–90% range.

Table 1. Baseline scenario parameters

Parameter	Value
Stations	Delft, NL (52.01°N, 4.36°E, 0 m); Innsbruck, AT (47.27°N, 11.40°E, 574 m)
Satellite	Micius-like LEO (TLE epoch June 2020)
Pass	2020-06-16, 23:44 UTC, 420 s window
Min. elevation	10°
Slant range	$\sim 600\text{--}2000 \text{ km}$
Conditions	Clear weather, moonless night ²⁵
Downlink wavelength	1550 nm (telecom C-band)
Source rate / fidelity	$1.0 \times 10^7 \text{ s}^{-1}$ / 0.963 ²⁶
Memory	$^{40}\text{Ca}^+$ trapped ion
Cooperativity	$C = 10^{27}$
Coherence T_2	62 ms ²²

4. RESULTS

4.1 Coherence-Limited Quantum Advantage

We begin with the published $^{40}\text{Ca}^+$ coherence, $T_2 = 62\text{ ms}$.²² As a starting point we propose a conservative cutoff of $t_c = 0.10 T_2$, keeping used pairs well inside the coherence window, since the pair already undergoes decoherence before being stored in the memories. As this results in memories holding valid pairs for only a few milliseconds, the memories will often be empty when a request arrives. We therefore select a high request rate with respect to the entangled pair supply, 5 Hz. In this regime, with a conservative cutoff time and a saturated request rate, we can expect most pairs to be used to play a game, and to do so with a higher-than-classical win rate. We evaluate the figure of merit of Eq. (1), repeated here for convenience: $P_{\text{win}} = \frac{3}{4} + A(P_q - \frac{3}{4})$. The availability A is measured as the fraction of requests served by stored entanglement, and P_q as the average win rate of those games.

Table 2 reports the outcome. There is no resolvable quantum advantage in this case: the mean global win rate is 75.001%, statistically consistent with the classical bound. The main issue is the low availability: only 0.050% of incoming requests are served with entanglement, and of those, the mean win rate is 76.45% even with a conservative $t_c = 0.10 T_2$. The percentage of stored pairs actually used to serve requests (0.97%) gives us more information: most of the pairs are actually discarded, even with a high request rate. This highlights the importance of the quantum memory decoherence value.

Once the first qubit is stored in memory, it already starts decohering. It waits for the other half of the pair to be stored ($\sim 1\text{ ms}$ of time-of-flight difference, averaged over the pass), and both further decohere until the heralding confirmation has crossed the classical channel between the stations, an assumed 750 km fibre path taking 3.75 ms. Of the 6.2 ms cutoff, $\sim 4.6\text{ ms}$ are therefore spent before the pair can serve any request. These effects are highly noticeable when the memory T_2 is on the order of milliseconds and explain the values obtained in this analysis.

To quantify the effect of a longer T_2 , we repeat the run at an optimistic value, $T_2 = 4\text{ s}$. We keep the cutoff and the request rate at the same values, $t_c = 0.10 T_2$ and 5 Hz respectively.

This analysis results in an availability roughly 80 times larger than the one before, at 4.1%. Since now the entangled pairs remain in memory much longer, the fraction of pairs used to serve requests dramatically increased, to 86.1%. Additionally, the longer coherence time provides a higher win rate amongst the served requests, at 78.87%. All of these factors together translate into a global win of 75.159%, showing a quantum advantage of more than 0.1 percentage point across all seeds.

The remaining question is whether we can do better with this memory: most pairs are used due to the large request rate, but this also means that more requests will be served classically, bringing the global win rate down. At the same time, the cutoff will affect both the availability and the fidelity of the pairs being consumed. We therefore sweep both parameters in the next section.

Table 2. CHSH game played over the Delft–Innsbruck pass, for a 5 Hz request rate, and cutoff $t_c = 0.10 T_2$.

Memory (T_2)	Availability	Pairs used	Played win	Global win
$^{40}\text{Ca}^+$, 62 ms ²²	0.050%	0.97%	76.45%	75.001%
4 s (optimistic)	4.10%	86.1%	78.87%	75.159%

4.2 Request-Rate and Cutoff Policy Map

With the same hardware and pass conditions, we can choose to change the cutoff policy and to study different request rates. Eq. (1) shows that they act in opposite directions. A longer cutoff raises availability, because a stored pair stays eligible for longer, but lowers played-game quality, because the pair a request consumes is older and more decohered. A higher request rate does the opposite: it lowers availability, since demand outruns what the link supplies, but raises played-game quality, because a pair consumed on arrival has had no time to decohere.

We sweep the request rate from 0.1 to 100 Hz against cutoff ratios up to $0.75 T_2$ at $T_2 = 4$ s, and map four quantities: the global advantage in percentage points (Fig. 4), availability (Fig. 5), the win rate on played games (Fig. 6), and the fraction of stored pairs a game actually consumes (Fig. 7).

Availability behaves as the supply-limited picture predicts. It rises with the cutoff, from about 5% at $0.05 T_2$ to about 31% at $0.75 T_2$ at the lowest demand, and falls with the request rate at every cutoff. Once demand outruns supply, from 10 Hz upwards, essentially every stored pair is consumed. Furthermore, in the saturated regime represented by a request rate of 100 Hz, most pairs are consumed before they have had time to decohere in memory, yielding the same played win rate across different cutoff times.

In lower request regimes, we can see that played-game quality runs the other way. At 0.1 Hz it falls from about 79% at $0.05 T_2$ to 72% at $0.75 T_2$, dropping below the classical $3/4$ by $0.5 T_2$ and significantly below it at $0.75 T_2$. Past that crossing a longer cutoff is actively harmful: it buys availability by admitting pairs that are no longer worth playing, and the global advantage turns negative in the low-demand, long-cutoff corner.

The global advantage peaks between the extremes, as it benefits from both high availability and high played-game win rate. The best region is found at cutoffs of $0.1\text{--}0.2 T_2$ and demand between 0.2 and 0.5 Hz, where the advantage reaches about 0.4 percentage points above the classical bound.

It drops in three directions, each for a different reason: towards high demand because availability collapses even though quality is at its ceiling, giving only 0.01 points at 100 Hz; towards long cutoffs at low demand because quality falls below the bound; and towards the shortest cutoff because availability starves faster than quality improves: $0.05 T_2$ is worse than $0.1 T_2$ at every rate.

The map therefore highlights an operational conclusion: a balance between the cutoff and the request rate is optimal for the overall win rate of our application.




Figure 4. The global advantage, in percentage points above the classical $3/4$ bound.




Figure 5. Availability over the request-rate/cutoff plane at $T_2 = 4$ s: the fraction of requests met with a stored pair.




Figure 6. Win rate on the games played with entanglement.




Figure 7. Fraction of stored pairs consumed by the application.

5. DISCUSSION AND CONCLUSION

We have presented an integrated satellite-link and NetSquid simulation framework that spans the end-to-end chain from orbital physics to application execution, resolving how orbital geometry, optics and memory hardware combine into an application-level win rate. The model can recreate any orbital pass for any pair of co-visible ground stations. We showcased the full chain on a European scenario: a Micius-like satellite serving trapped-ion memories in Delft and Innsbruck over a single pass.

An initial study with a published $^{40}\text{Ca}^+$ coherence of 62 ms highlights the relevance of this parameter. At this value, entanglement is often not present at the users' memories when a request to run the application arrives (0.050% of requests find stored entanglement and under 1% of the pairs stored are actually consumed, with the rest being discarded). As a result, the global win rate is statistically consistent with the classical bound. At $T_2 = 4$ s, the same link serves 4.10% of requests, consumes 86% of its stored pairs and clears the bound at 75.159%.

The policy map presented showcases the effect of the operating knobs one can choose from, cutoff time and application request rate, given a fixed set of hardware and pass conditions (Sec. 4.2). An optimal operating region is identified, where the availability of stored entangled pairs and their win rate combine in a maximum global win rate. This quantum advantage, ~ 0.4 percentage points, is modest but consistent over all simulated passes and shows how with near-term hardware we can use satellite-distributed entanglement to beat any classical strategy for coordination applications.

The satellite link model is sufficient for system-level budget analysis. However, it could be refined to account for specific local-time conditions, to study in real time whether a deployed system is able to use a given satellite pass to gain a quantum advantage for a coordination application.

Follow-up research can perform a systematic sweep over hardware and operating parameters and identify the system requirements for a given advantage and request demand. One may also consider multiplexing to increase the rate, as well as more than one memory slot per user, which the current model already supports.

Finally, the current tool is bounded to a single low-Earth-orbit satellite delivering entanglement only to ground-station pairs it can see simultaneously. Multi-satellite and intercontinental architectures would lift this dual-visibility constraint. Proposed families include orbiting and hybrid repeater designs,^{28–30} satellites carrying on-board quantum memories,^{31,32} routing and scheduling across time-varying constellations,³³ and comparative protocol evaluations under free-space channels.³⁴ Pursuing them would extend earlier space-segment design work for trusted-node QKD constellations³⁵ from key delivery to end-to-end entanglement distribution, to eventually enable a quantum advantage for coordination applications at global scale.

APPENDIX A. EXTENDED MODEL PARAMETERS

Table 3 lists the parameter values assumed by the link budget of Sec. 2.2, the ground-segment chain of Sec. 2.3, and the EU scenario presented in Sec. 3.

Quantum memory. We assume an optimistic cooperativity $C = 10^{27}$ for the $^{40}\text{Ca}^+$ trapped-ion memory, which gives a CAPS reflection efficiency $\eta_{\text{CAPS}} = 0.41$. The qubit is encoded in two Zeeman sublevels of the $S_{1/2}$ electronic ground state, so decoherence is modelled as $T_1 \rightarrow \infty$, $T_2 = 62$ ms.²²

Dark counts. The baseline takes $r_{\text{dc}} = 10$ Hz at the 854 nm ion transition, read at the same detector bias point as the efficiency $\eta_{\text{det}} = 0.67$.¹⁸ With $\Delta T = 1$ ns the conversion of Sec. 2.3 gives $n_{\text{dc}} = 1.0 \times 10^{-8}$, dominating over the moonless-night sky term $\bar{n} = 1.4 \times 10^{-10}$ that Eq. (2) returns at the diffraction-limited field of view of Table 3.

Table 3. Extended model parameters

Wavelength λ (EU downlink)	1550 nm
Receiver aperture D_{rx}	1.2 m ⁵
Transmitter aperture D_{tx}	0.3 m ³⁶
Pointing jitter σ_p	2 μrad ⁵
Transmitter-optics loss	3.0 dB per arm ³⁷
Receiver optics / AO-to-fibre coupling	2.8 dB ¹² / 3.2 dB ¹³
Source rate R_{src} / fidelity F_{src}	$1.0 \times 10^7 \text{ s}^{-1}$ / 0.963 ²⁶
Coincidence window ΔT / filter $\Delta\lambda$	1 ns / 1 nm
Field of view $\Omega = \pi(0.45\lambda/D_{\text{rx}})^2$ (SMF) ¹²	$1.06 \times 10^{-12} \text{ sr}$
Sky radiance H (moonless night, 1550 nm)	$15 \text{ W m}^{-2} \text{ m}^{-1} \text{ sr}^{-1}$ ²⁵
CAPS reflection η_{CAPS} ($C=10$)	0.412 ^{15,27}
Detector efficiency η_{det} / dark-count rate (854 nm)	0.67 / 10 Hz ¹⁸
QFC chain η_{QFC} (1550 nm $\leftrightarrow$ 854 nm)	0.572 per arm ¹⁷
CAPS gate fidelity / readout fidelity	0.958 ²⁰ / 0.99991 ²⁴
Memory T_1	∞ (ground-state Zeeman qubit ²²)
Coherence T_2 (published / optimistic)	62 ms ²² / 4 s

APPENDIX B. MICIUS VALIDATION

This appendix gives the detail of the Micius validation of Sec. 2.2: how the candidate pass was identified, which system-budget terms were held fixed, and how much each atmospheric input contributes. Figure 8 shows the comparison between the results published in Yin et al.⁵ and our simulation output.




Figure 8. Simulated and measured Micius dual-downlink performance over the common visibility window of the Delingha and Nanshan stations.⁵ (a) Slant ranges for the geometry-matched candidate pass, against values digitised from their Fig. 2a. (b) Total two-link attenuation predicted by the weather-constrained model, against values digitised from their Fig. 2b.

Candidate pass. Yin et al.⁵ report measurements collected between 6 September and 22 October 2018, but they do not identify the calendar date of the pass plotted in their Fig. 2. We propagated the Micius orbital elements for candidate passes in that interval and found the best match on 1 October 2018 at 18:11:11 UTC.

Fixed system budget. For each arm, the simulated attenuation is the sum of diffraction, pointing, atmospheric extinction, and a fixed 12.8 dB system budget. That budget contains detector (2.8 dB), transmitter telescope (3.0 dB), ground optics (3.0 dB) and coupling (4.0 dB) losses, from the published Micius experiments.^{5,37,38}

Atmospheric inputs. We infer the local atmospheric conditions at the two observatories during the candidate pass from NASA POWER hourly data.³⁹ We then recreate those conditions in our libRadtran implementation.¹¹

APPENDIX C. CHSH MEASUREMENT PROTOCOL

This appendix specifies the measurement protocol used by the application layer. Alice receives input $x \in \{0, 1\}$, applies $R_y(-\pi/2)$ if $x = 1$, and measures in the Z basis to obtain outcome a . Bob first converts the delivered $B_{11} = |\Psi^-\rangle$ pair to $|\Phi^+\rangle$ by applying X followed by Z . He then applies $R_y(-\pi/4)$ for $y = 0$ or $R_y(+\pi/4)$ for $y = 1$ and measures in the Z basis to obtain outcome b . The players win when $a \oplus b = x \wedge y$. For an ideal pair these angles achieve the Tsirelson value $S = 2\sqrt{2}$, which corresponds to the quantum win probability $P_q = \cos^2(\pi/8) \approx 0.854$.

ACKNOWLEDGMENTS

The authors would like to thank Conor Bradley and Michał van Hooft for useful discussions. This project has received funding from an NWO VICI grant. Claude Opus by Anthropic has been used to audit the simulation code and to assist with figure generation and text edits.

REFERENCES

- [1] da Silva, F. F. and Wehner, S., “Entanglement improves coordination in distributed systems,” in [*Proceedings of the 2nd Workshop on Quantum Networks and Distributed Quantum Computing (QuNet ’25)*], 14–20, Association for Computing Machinery, New York, NY, USA (2025).
- [2] Wehner, S., “Quantum advantage for coordinated frequency selection against distributed jammers,” *arXiv preprint arXiv:2604.20647* (2026).
- [3] Ding, D. and Jiang, L., “Coordinating decisions via quantum telepathy,” *arXiv preprint arXiv:2407.21723* (2024).
- [4] Wehner, S., Elkouss, D., and Hanson, R., “Quantum internet: A vision for the road ahead,” *Science* **362**(6412), eaam9288 (2018).
- [5] Yin, J., Li, Y.-H., Liao, S.-K., Yang, M., Cao, Y., others, Peng, C.-Z., Ekert, A. K., and Pan, J.-W., “Entanglement-based secure quantum cryptography over 1,120 kilometres,” *Nature* **582**, 501–505 (2020).
- [6] Khatri, S., Brady, A. J., Desporte, R. A., Bart, M. P., and Dowling, J. P., “Spooky action at a global distance: analysis of space-based entanglement distribution for the quantum internet,” *npj Quantum Information* **7**, 4 (2021).
- [7] Clauser, J. F., Horne, M. A., Shimony, A., and Holt, R. A., “Proposed experiment to test local hidden-variable theories,” *Physical Review Letters* **23**(15), 880–884 (1969).
- [8] Cirel’son, B. S., “Quantum generalizations of Bell’s inequality,” *Letters in Mathematical Physics* **4**, 93–100 (1980).
- [9] Coopmans, T., Knegjens, R., Dahlberg, A., Maier, D., Nijsten, L., others, and Wehner, S., “NetSquid, a NETwork Simulator for QUantum Information using Discrete events,” *Communications Physics* **4**, 164 (2021).
- [10] Titterton, P. J., “Power reduction and fluctuations caused by narrow laser beam motion in the far field,” *Applied Optics* **12**(2), 423–425 (1973).
- [11] Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., et al., “The libRadtran software package for radiative transfer calculations (version 2.0.1),” *Geoscientific Model Development* **9**(5), 1647–1672 (2016).

- [12] Marulanda Acosta, V., Dequal, D., Schiavon, M., Montmerle-Bonnefois, A., Lim, C. B., Conan, J.-M., and Diamanti, E., “Analysis of satellite-to-ground quantum key distribution with adaptive optics,” *New Journal of Physics* **26**(2), 023039 (2024).
- [13] Canuet, L., Védrenne, N., Conan, J.-M., Petit, C., Artaud, G., Rissons, A., and Lacan, J., “Statistical properties of single-mode fiber coupling of satellite-to-ground laser links partially corrected by adaptive optics,” *J. Opt. Soc. Am. A* **35**(1), 148–162 (2018).
- [14] Duan, L.-M. and Kimble, H. J., “Scalable photonic quantum computation through cavity-assisted interactions,” *Physical Review Letters* **92**, 127902 (2004).
- [15] Ji, J.-W., Sunami, S., Kikura, S., Goban, A., and Simon, C., “A global quantum network with ground-based single-atom memories in optical cavities and satellite links,” *arXiv preprint arXiv:2507.02333* (2025).
- [16] Bernien, H., Hensen, B., Pfaff, W., Koolstra, G., Blok, M. S., Robledo, L., Taminiau, T. H., Markham, M., et al., “Heralded entanglement between solid-state qubits separated by three metres,” *Nature* **497**, 86–90 (2013).
- [17] Arenskötter, E., Bauer, T., Kucera, S., Bock, M., Eschner, J., and Becher, C., “Telecom quantum photonic interface for a $^{40}\text{Ca}^+$ single-ion quantum memory,” *npj Quantum Information* **9**(1), 34 (2023).
- [18] Li, H., Zhang, L., You, L., Yang, X., Zhang, W., Liu, X., Chen, S., Wang, Z., and Xie, X., “Large-sensitive-area superconducting nanowire single-photon detector at 850 nm with high detection efficiency,” *Optics Express* **23**(13), 17301–17308 (2015).
- [19] Gruneisen, M. T., Flanagan, M. B., Sickmiller, B. A., Black, J. P., Stoltenberg, K. E., and Duchane, A. W., “Modeling daytime sky access for a satellite quantum key distribution downlink,” *Optics Express* **23**(18), 23924 (2015).
- [20] Borne, A., Northup, T. E., Blatt, R., and Dayan, B., “Efficient ion-photon qubit swap gate in realistic ion cavity-qed systems without strong coupling,” *Optics Express* **28**(8), 11822 (2020).
- [21] Krutyanskiy, V., Meraner, M., Schupp, J., Krcmarsky, V., Hainzer, H., and Lanyon, B. P., “Light-matter entanglement over 50 km of optical fibre,” *npj Quantum Information* **5**, 72 (2019).
- [22] Krutyanskiy, V., Galli, M., Krcmarsky, V., Baier, S., Fioretto, D. A., Pu, Y., Mazloom, A., Sekatski, P., Canteri, M., Teller, M., Schupp, J., Bate, J., Meraner, M., Sangouard, N., Lanyon, B. P., and Northup, T. E., “Entanglement of trapped-ion qubits separated by 230 meters,” *Physical Review Letters* **130**(5), 050803 (2023).
- [23] Bock, M., Eich, P., Kucera, S., Kreis, M., Lenhard, A., Becher, C., and Eschner, J., “High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion,” *Nature Communications* **9**, 1998 (2018).
- [24] Myerson, A. H., Szwer, D. J., Webster, S. C., Allcock, D. T. C., Curtis, M. J., Imreh, G., Sherman, J. A., Stacey, D. N., et al., “High-fidelity readout of trapped-ion qubits,” *Physical Review Letters* **100**(20), 200502 (2008).
- [25] Stathis, A., Ntanos, A., Lyras, N. K., Giannoulis, G., Panagopoulos, A. D., and Avramopoulos, H., “Toward converged satellite/fiber 1550 nm DS-BB84 QKD networks: Feasibility analysis and system requirements,” *Photonics* **11**(7), 609 (2024).
- [26] Spellmeyer, N. W., Lee, C., Scheinbart, M., and Hamilton, S. A., “Precision synchronization over large Doppler ranges for satellite-based free-space quantum networking,” in [*Quantum Computing, Communication, and Simulation IV*], Hemmer, P. R. and Migdall, A. L., eds., *Proc. SPIE* **12911**, 12911–11 (2024). The 0.963 entangled-state fidelity is quoted from the presentation version of this work, NASA Technical Reports Server 20240000030, slide 3.
- [27] Kassa, E., Gao, S., Teh, S., van Dinter, D., and Takahashi, H., “How to integrate a miniature optical cavity in a linear ion trap: shielding dielectrics and trap symmetry,” *Physical Review Applied* **23**, 024038 (2025).
- [28] Liorni, C., Kampermann, H., and Bruß, D., “Quantum repeaters in space,” *New Journal of Physics* **23**, 053021 (2021).
- [29] Wallnöfer, J., Hahn, F., Gündoğan, M., Sidhu, J. S., Wiesner, F., Walk, N., Eisert, J., and Wolters, J., “Simulating quantum repeater strategies for multiple satellites,” *npj Quantum Information* **8**, 129 (2022).
- [30] Dawar, M., Paillet, C., Vyas, N., Thain, A., Guilherme, R. H., and Riedinger, R., “Feasibility of satellite-augmented global quantum repeater networks.” *arXiv:2603.11127* (2026).

- [31] Gündoğan, M., Sidhu, J. S., Henderson, V., Mazzarella, L., Wolters, J., Oi, D. K. L., and Krutzik, M., “Space-borne quantum memories for global quantum communication,” *npj Quantum Information* **7**, 128 (2021).
- [32] Tubío, V. D. et al., “Satellite-assisted quantum communication with single photon sources and atomic memories.” arXiv:2411.09533 (2024).
- [33] Koudia, S., ur Rehman, J., and Chatzinotas, S., “Space-based quantum internet: Entanglement distribution in time-varying LEO constellations.” arXiv:2409.17032 (2024).
- [34] Zaunders, N. and Ralph, T. C., “Entanglement distribution via satellite: an evaluation of competing protocols assuming realistic free-space optical channels.” arXiv:2510.01633 (2025).
- [35] Loarte Castro, S., *Towards a global space-based QKD network*, Master’s thesis, Delft University of Technology (2019).
- [36] Lu, C.-Y., Cao, Y., Peng, C.-Z., and Pan, J.-W., “Micius quantum experiments in space,” *Reviews of Modern Physics* **94**, 035001 (2022).
- [37] Yin, J., Cao, Y., Li, Y.-H., Liao, S.-K., Zhang, L., Ren, J.-G., Cai, W.-Q., Liu, W.-Y., et al., “Satellite-based entanglement distribution over 1200 kilometers,” *Science* **356**, 1140–1144 (2017).
- [38] Chen, Y.-A., Zhang, Q., Chen, T.-Y., Cai, W.-Q., Liao, S.-K., Zhang, J., Chen, K., Yin, J., et al., “An integrated space-to-ground quantum communication network over 4,600 kilometres,” *Nature* **589**, 214–219 (2021).
- [39] NASA Langley Research Center, “NASA prediction of worldwide energy resources (POWER) project: Hourly time series data.” NASA Langley Research Center (LaRC) POWER Project, funded through the NASA Earth Science Division (2026). Hourly 2.x.x version, source MERRA-2; data accessed July 2026, <https://power.larc.nasa.gov/>.