
QNETSIM for memory assisted design of satellite and hybrid quantum networks

Seid Koudia*, Francesco Vista, Heyang Peng, Ilora Maity, Intidhar Bedhief, Juan Carlos Hernandez Hernandez, Symeon Chatzinotas

Signal Processing and Communications (SIGCOM),
Interdisciplinary Centre for Security, Reliability and Trust (SnT),
University of Luxembourg, L-1855 Luxembourg, Luxembourg

*Corresponding author: seid.koudia@uni.lu

Abstract. Space-based quantum information networks require simulation tools that capture orbital dynamics, probabilistic quantum operations, and classical coordination. We introduce QNETSIM, a modular discrete-event simulator for the design and performance evaluation of satellite and hybrid space–ground quantum networks. The framework integrates orbital geometry, free-space optical and fiber channels, quantum memories, entanglement generation, and Bell-state measurements within a common event-driven architecture. Physical models account for optical loss, atmospheric turbulence, detector imperfections, finite memory capacity, and storage-induced decoherence. Space–time routing and segmented nesting support entanglement distribution across changing topologies, while a software-defined control plane coordinates generation, storage, and swapping. Configurable scenarios encompass onboard repeaters, satellite constellations, ground-based repeaters, and direct satellite distribution. Through explicit interfaces between physical models, network protocols, and orchestration, QNETSIM links component characteristics and operational policies to entanglement fidelity, throughput, latency, availability, and memory utilization, enabling systematic exploration of architectural and resource-management trade-offs in future space-based quantum information networks.

1 Introduction

Quantum information networks aim to connect quantum processors, memories, sensors, and photonic devices through the distribution of quantum states and entanglement. By extending quantum resources beyond individual laboratories, these networks can support quantum key distribution (QKD), distributed quantum computation, and entanglement-enhanced sensing and metrology [1,2]. Their development requires more than the transmission of isolated photons: useful quantum correlations must be established between selected endpoints, preserved until they are needed, and consumed by applications under appropriate timing and quality constraints. Space-based quantum information networks offer a route to extending these capabilities across continental and intercontinental distances, complementing terrestrial optical-fiber infrastructure with satellite links.

The motivation for a space segment originates in the physical limitations of long-distance quantum transmission. Optical-fiber transmissivity decreases exponentially with distance, imposing a fundamental rate–loss trade-off on repeaterless quantum communication [3]. Unknown quantum states cannot be noiselessly copied or regenerated using conventional optical amplification. Satellite links offer a different propagation geometry: most of the path lies outside the atmosphere, avoiding the material absorption accumulated along a comparable fiber connection. Their performance is instead governed by diffraction, finite receiving apertures, atmospheric propagation, and terminal pointing [4]. The satellite distribution of entanglement demonstrated by Yin et al. established an experimental foundation for this approach [5]. Its extension to network services requires connecting satellite-generated resources to ground-based users and, where appropriate, to quantum memories and repeater nodes.

This hybrid perspective is particularly relevant to the Eu-

ropean Quantum Communication Infrastructure (EuroQCI). The European Commission and the European Space Agency envisage an infrastructure combining national terrestrial quantum networks with a space segment, supporting the protection of sensitive communications and critical infrastructure across Europe [6]. EAGLE-1 is being developed to demonstrate satellite-to-ground QKD and provide operational experience for this wider infrastructure [7]. Such initiatives make the integration of optical ground stations, terrestrial networks, and their control systems an immediate engineering concern. QKD delivers classical secret keys, whereas general entanglement services additionally require the distribution and manipulation of shared quantum states. Moving toward memory-assisted networking therefore raises further questions about storage, swapping, resource allocation, and interfaces between independently operated network segments.

Space-based quantum links face several coupled physical limitations. Orbital motion changes slant range, elevation, propagation delay, and the duration of useful contact, while atmospheric extinction, turbulence, clouds, and background illumination influence the optical opportunities within that contact [4, 8]. Finite apertures and pointing errors determine how much light is collected; wavefront distortion and adaptive-optics performance influence coupling into the accepted receiving modes [9]. Source brightness, multipair emission, detector efficiency, dark counts, timing jitter, and dead time then determine which arrivals contribute to usable correlations [10–12]. Consequently, increasing the emitted photon rate or extending geometrical visibility does not necessarily improve application performance. Wavelength, receiver acceptance, detector technology, and operating conditions must be assessed together.

Quantum memories and Bell-state measurements introduce a further set of constraints. Memories permit asynchronous elementary links to be joined, but their usefulness depends

jointly on write and retrieval efficiencies, coherence time, bandwidth, and the number of independently addressable modes. Ensemble optical depth and control-field dynamics influence storage and retrieval, while environmental noise degrades the stored logical state [13, 14]. Entanglement swapping additionally requires compatible optical modes and a successful measurement, with passive linear-optical Bell analyzers subject to an intrinsic discrimination limit under their standard resource assumptions [15]. Repeater placement consequently trades reduced synchronization demands against storage loss, decoherence, and additional probabilistic operations. For onboard memories, these trade-offs also involve payload integration and space qualification, motivating quantitative capability studies rather than assumptions of ideal storage [16].

The same physical constraints shape network control and management. A geometrically available route is useful only if its elementary resources can be generated, stored, and joined before their quality or the contact opportunities expire. Routing must therefore account for time-dependent connectivity, memory occupancy, and the ordering of swaps. Classical communication is integral to this process: heralding, resource identification, reservation, measurement outcomes, and completion notifications must remain associated with the correct quantum resources [17]. Signaling and processing delays can consume part of the available storage interval. A software-defined architecture separates network-wide scheduling and admission policies from local execution, while fast pointing, synchronization, and optical correction remain terminal functions. Management must also reconcile heterogeneous device capabilities, service requirements, and control interfaces across the space and terrestrial segments [2, 17].

These interdependencies make system-level simulation an essential instrument for mission planning. Before committing to payload and ground-segment designs, mission studies must assess orbit selection, ground-station placement, optical apertures, wavelengths, memory provisioning, and operating policies against a common set of service objectives. A longer contact can be offset by higher propagation loss; additional memory capacity can remain underused when loading is inefficient; and a shorter route can perform poorly when its resource-generation times are incompatible. Link-budget calculations provide necessary physical inputs, but do not by themselves resolve stochastic waiting, finite-capacity contention, failed swaps, or delayed control actions. Discrete-event simulation provides a natural means of representing these dependencies and relating component assumptions to delivered-pair rate, conditional fidelity, latency, availability, and resource utilization [18–20]. Configuration-controlled studies can then support requirement allocation and sensitivity analysis before experimental deployment.

Existing simulation and architecture studies provide important foundations. NetSquid [18], QuNetSim [21], QNET [22], SeQUeNCe [19], and QuISP [23] offer complementary approaches to device modeling, event execution, and quantum-network protocols. Satellite-specific work has examined ground-memory architectures, orbiting repeaters, constellation coverage, and space-borne memory requirements [8, 16, 24, 25]. Event-based multi-satellite studies, satellite-connected quantum cities, and coupled orbital–optical simulations further

demonstrate the importance of temporal and physical modeling [20, 26, 27]. The QUASAR simulator also addresses dynamic satellite architectures and routing [28]. Table 1 summarizes the simulator abstractions relevant to the present work. These contributions motivate an integrated mission-design environment in which physical assumptions, finite quantum resources, and programmable policies can be varied through explicit interfaces.

The roadmap of de Forges de Parvy et al. distinguishes a direct entanglement-resource demonstrator from a networking stage that assembles elementary links through storage and swapping [2]. QNETSIM addresses the system-level design questions raised by this transition. We introduce it as a modular discrete-event simulator for space-based and hybrid space–ground quantum information networks. Its architecture combines orbital geometry, free-space optical and fiber channels, quantum memories, entanglement generation and swapping, network algorithms, and software-defined orchestration on a common event timeline. Physical models are separated from the scheduling core and expose operation probabilities and conditional quality descriptors to node-level resource management. The accompanying physical formulation relates matter–light interactions and open-system evolution to effective storage and measurement channels, providing a basis for selecting and refining the network-level descriptions. Space–time-based distribution and segmented nesting build on the framework of Koudia et al. [29]. Within QNETSIM, these mechanisms connect evolving link opportunities to the scheduling and consumption of stored entanglement resources.

Three scenarios organize the evaluation around complementary mission-design questions. In S1, an onboard repeater connects two optical ground stations, exposing the interaction between geometry, asynchronous loading, and memory coherence. In S2, a low-Earth-orbit constellation introduces time-dependent routing and segmented swapping. In S3, a ground-based repeater at Windhof joins a satellite downlink to a fiber connection toward Kirchberg, examining heterogeneous resource supply. A separate direct Matera–Hermos calculation provides a memory-free optical reference for receiver and source operating-point studies. Together, these cases connect architectural choices to observable network behavior and establish a framework for investigating the transition from individual satellite quantum links to managed entanglement services. The Results present the service indicators, architecture, and scenario findings; the Methods detail the physical models, algorithms, configurations, and statistical treatment, including the derivations in Sections 4.2–4.5.

2 Results

2.1 Service requirements and performance measures

The memory-assisted service is a pair shared between specified endpoints, with a quality descriptor and the classical information required to interpret its measurement frame. A direct detected coincidence is a different product: detection consumes the arriving photons. Table 2 maps the service indicators of Ref. [2] to quantities recorded by the simulator

Table 1: Representative simulator abstractions. Entries identify modeling and interface scope, rather than an accuracy or speed ranking.

Tool	Execution / interface	State and component representation	Protocol and satellite scope
QuNetSim [21]	Python, event-driven	Backend quantum states; configurable connections	Quantum/classical messages, routing, teleportation, and entanglement protocols.
NetSquid [18, 30]	Python; optimized C/Cython DES	Ket, stabilizer, and density matrices; components with time-dependent noise	Asynchronous protocols and classical control; satellite applications include [27].
QNET [22]	Python DES; QCompute-related interfaces	Modular devices, channels, and physical models	Resource/protocol simulation and computation interfaces; documented Micius example.
SeQUeNCe [19]	Python DES	Hardware and quantum-resource components	Layered entanglement, resource, network, and application management.
QuISP [23]	C++ / OMNeT++	Scalable error-based resource representation	Heterogeneous repeater-network protocols.
QNETSIM	Native token DES; Python workflows; REST control	Analytical FSO/fiber and platform-specific memories; finite modes; field-to-instrument reference models	Single repeater, constellation, and hybrid workflows; STBD, nesting, and request/QoS control.

Table 2: Service requirements and corresponding observables. Observation time, contact time, and event count are kept distinct when reporting rates.

Requirement	Observable	Design decision
Resource rate	Delivered pairs per observation second; separately pairs per pass or contact second	Source allocation and optical/memory provisioning.
State quality	Conditional Bell fidelity of a declared delivery population; accepted counts at a target	Memory coherence, age cutoff, and swap depth.
Latency	Request-to-delivery time, separating cold start from stocked inventory and physical readiness from known frame	Proactive storage and classical signaling.
Availability	Fraction of requests meeting quality and deadline targets, distinct from geometrical visibility	Temporal routing and coverage.
Resource cost	Occupied memory-mode seconds per delivered pair; overflow and expiry counts	Bank capacity and buffer policy.
Timing and control	Relative photon-arrival uncertainty, gate acceptance, control traffic and delay	Synchronization and controller placement.

and to the associated design decisions.

A usable-pair rate counts only deliveries meeting the specified fidelity and deadline. Failed requests remain in the reliability denominator. Conditional fidelity is undefined when there are no deliveries, and a bin average cannot determine how many individual pairs exceed a new quality threshold. These distinctions prevent a high raw count from being mistaken for a high-quality service. Full KPI definitions and the retained bin-level statistics are given in Sections 4.6 and 2.4.1.

2.2 A common execution and resource architecture


Figure 1a shows the separation of application requests, network decisions, local agents, and physical models. The event core supplies one ordered simulation timeline. Source, storage, and repeater agents evaluate physical operations when their events execute. A resource ledger records which excitations exist and where they are stored. A controller authorizes future actions using information received through classical messages; it does not directly modify a physical state.

A pair has a unique identity, a request identifier, two endpoint records, a state representation, and a known or pending Pauli frame. Each endpoint retains its memory bank, mode, original write time, last evolution time, and allocation version. A successful swap consumes the two local operands and


creates a new pair identity for the two surviving endpoints. Their original write times are preserved. Capacity counts both occupied and committed slots, including operations awaiting a herald. Multiple labels referring to the same physical bank share its capacity.

The physical interface separates outcome probability from the conditional output state. Compact configurations represent quality by a scalar fidelity or Werner parameter. Bell-probability vectors preserve independent Pauli errors at almost the same state-storage cost; general state-dependent filtering requires a density matrix or a more general quantum instrument. This choice is recorded with the simulation configuration. Sections 4.2 and 4.3 derive the accepted memory and measurement operations underlying those representations.

Events distinguish emission, arrival, memory loading, herald reception, swapping, and final notification. Time ordering uses the scheduled time followed by declared priority and insertion order. Source rates are allocated across modes rather than multiplied by memory capacity. Poisson thinning can skip unsuccessful attempts when those failures have no effect on the system state; periodic sources instead require geometric skipping. Detector busy periods, failed-attempt reservations, and source multipair statistics require explicit treatment where relevant. These conditions define the acceleration boundary without changing the resource budget.



(a) Layered execution and resource architecture.



(b) Event/action information flow.

Figure 1: QNETSIM architecture and orchestration. (a) Physical models are evaluated at discrete-event boundaries, node agents execute local operations, and the resource ledger records endpoint occupancy and state descriptors. Solid arrows denote software interactions, not quantum-state copying; the instrument-valued formulation specifies the meaning of compact physical interfaces. (b) Request, generation, storage, swap, and delivery messages are scheduled operations rather than instantaneous transfers. Physical outcome, controller knowledge, and endpoint frame readiness have distinct timestamps. Local BSM operands are consumed while the original ages of surviving remote memories are preserved. The external-control transport carries telemetry and commands; effective execution times remain simulation events.

2.3 Orchestration and information flow

QNETSIM supports autonomous local execution, an in-process controller, and an external Ryu-hosted REST controller. Node proxies report timestamped storage and swap outcomes and retrieve pending actions. The controller maintains observed inventories, topology, request deadlines, and swap dependencies. Figure 1b follows the request, generation, storage, swap, and notification sequence; Table 11 gives the message fields.

The separation between control and execution matters when messages are delayed. A decision based on a valid storage report can arrive after an operand has expired or been consumed. Local validation of the allocation version and request epoch rejects such stale instructions. Swaps reserve and consume their operands atomically. Quality estimates may be propagated from original write times, while future random channel realizations remain outside the controller’s information set.


Classical propagation, serialization, queueing, processing, and polling are separate modeled delays. Their sum determines when an instruction can become eligible on the simulation timeline; wall-clock HTTP execution time is a software

measurement rather than automatically a simulated network delay. Optical wavefront correction, pointing, and source synchronization remain local control functions. The network controller sets slower scheduling objectives and responds to telemetry.

Ryu is retained as a compatibility adapter, while the scientific message contract is transport independent. Its upstream maintenance status motivates this separation [31]. The scenario results below characterize the selected execution configurations; an external-controller equivalence would additionally require matched input and action traces. The resource and event contracts provide the common comparison boundary.

2.4 S1: Onboard storage and geometry-dependent memory requirements

Scenario S1 uses two elementary resources whose local halves occupy separate memories on a satellite; their remote halves belong to two OGSs. The onboard BSM consumes the local excitations and, on success, connects the ground endpoints. In a satellite-source realization, one photon from each elementary pair travels down its corresponding arm and its partner is



4 independent seeds per setting; no temporal smoothing. Fidelity undefined without deliveries.

Figure 2: Complete S1 conditional-fidelity sweep. Rows are cases A–D; the left column varies T_2 , and the right varies σ_R^2 . Each curve is the median of supported run-bin means over four repeats, with no temporal smoothing. Empty-support bins remain undefined. The first two rows have co-located stations, not a 1000 km separation; the lower rows have 500 km separation and an Alice-anchored pass.

retained locally. A stored-pair service also requires heralded loading at the remote endpoint. Optical conversion and loading are therefore distinct from photon detection and must be included in a hardware realization.


The controlled campaign comprises four local-plane geometries at 500 km altitude and 7.2 km/s projected speed. Cases A and B have co-located terminals and form limiting cases with 0° and 90° track directions. Cases C and D have 500 km station separation and an Alice-anchored pass at 90° and 45° . Their observation windows are 300, 440, 340, and 440 s. These definitions prevent the co-located presets from being interpreted as separated-city links or their different-duration

yields as a paired crossing-angle comparison.

The recorded cold-atom settings are $d = 20$, dephasing exponent $n = 1.3$, initial mode efficiency 0.9, and retrieval-timescale parameter $T_{\text{mem}} = 1000$ s. The sweeps allocate 8192 modes per label, use a nominal source rate $5 \times 10^5 \text{ s}^{-1}$, and record counts in 2 s bins. Optical defaults include 810 nm wavelength, a 0.15 m transmitter, a 0.50 m receiver, and $12 \mu\text{rad}$ divergence. Table 6 records the numerical configuration fields.

Each geometry has four coherence settings at fixed Rytov variance $\sigma_R^2 = 0.08$ and four turbulence settings at the default $T_2 = 100$ s. Four runs per setting give 32 settings and 128 runs. Table 3 reports the mean integrated count and sample standard

QNESIM Single Repeater (historical local-plane preset): DES only



4 independent seeds per setting; no temporal smoothing. Fidelity undefined without deliveries.

Figure 3: Complete S1 rate sweep for the same 32 settings and 128 runs as Fig. 2. The ordinate is the mean recorded end-to-end delivery rate in pairs/s, obtained from 2 s bins; no rate multiplier or temporal smoothing is applied. The geometrical and observation-window definitions are given in Section 2.4.

Table 3: S1 integrated delivered pairs across the coherence sweep. Entries are mean $\pm$ sample standard deviation over four runs, not confidence intervals. A and B are co-located limits; C and D have 500 km station separation. Observation durations are 300, 440, 340, and 440 s, respectively. All cases use $\sigma_R^2 = 0.08$.

Case	$T_2 = 0.3$ s	$T_2 = 3$ s	$T_2 = 30$ s	$T_2 = 300$ s
A	446.2 $\pm$ 383.7	15,423.5 $\pm$ 4,614.9	27,201.5 $\pm$ 242.6	27,570.8 $\pm$ 204.2
B	362.5 $\pm$ 335.3	18,580.5 $\pm$ 5,049.0	28,838.8 $\pm$ 621.7	29,842.5 $\pm$ 217.4
C	1.2 $\pm$ 1.0	284.8 $\pm$ 106.2	4,750.2 $\pm$ 267.6	15,921.8 $\pm$ 41.5
D	4.2 $\pm$ 3.3	213.2 $\pm$ 17.4	5,264.2 $\pm$ 137.2	18,689.0 $\pm$ 148.7

deviation between runs. Figures 2 and 3 show the complete conditional-fidelity and delivery-rate sweeps for cases A–D, including the co-located limiting cases.

Coherence requirements depend strongly on loading geometry. In case A, increasing T_2 from 30 to 300 s increases mean delivery by only 1.36%. In separated cases C and D, the

Table 4: All 32 S1 settings. $\bar{N}$ and s_N are the mean and sample standard deviation of integrated deliveries over four runs. R_{peak} is the maximum ensemble mean rate; t_R is its bin-end time relative to the preset origin. F_{peak} is the maximum median of supported run-bin mean fidelities; $\bar{F}_{\text{pair}}$ is the success-count-weighted fidelity over all bins and repeats. These are reanalyses of the recorded bin tables.

Case	Setting	$\bar{N}$	s_N	R_{peak} (s^{-1})	t_R (s)	F_{peak}	$\bar{F}_{\text{pair}}$
A	$T_2 = 0.3$ s	446.25	383.74	4.750	-36	0.4053	0.3665
A	$T_2 = 3$ s	15,423.50	4,614.93	108.000	+16	0.8246	0.7538
A	$T_2 = 30$ s	27,201.50	242.63	182.875	-4	0.8889	0.8690
A	$T_2 = 300$ s	27,570.75	204.23	179.375	+4	0.8915	0.8741
A	$\sigma_R^2 = 0.02$	27,673.50	54.76	181.875	+2	0.8935	0.8761
A	$\sigma_R^2 = 0.08$	27,583.50	135.74	177.750	+8	0.8911	0.8729
A	$\sigma_R^2 = 0.2$	27,212.25	173.48	178.000	+8	0.8855	0.8671
A	$\sigma_R^2 = 0.5$	26,047.75	73.68	174.000	-4	0.8740	0.8535
B	$T_2 = 0.3$ s	362.50	335.31	6.000	+0	0.4043	0.3603
B	$T_2 = 3$ s	18,580.50	5,048.96	121.375	+6	0.8182	0.7697
B	$T_2 = 30$ s	28,838.75	621.74	178.625	-10	0.8878	0.8611
B	$T_2 = 300$ s	29,842.50	217.43	183.625	+6	0.8915	0.8682
B	$\sigma_R^2 = 0.02$	29,653.50	242.08	181.875	+10	0.8932	0.8694
B	$\sigma_R^2 = 0.08$	29,574.50	221.04	189.875	+4	0.8901	0.8666
B	$\sigma_R^2 = 0.2$	29,217.50	169.95	177.000	+4	0.8858	0.8605
B	$\sigma_R^2 = 0.5$	28,031.00	182.33	175.375	-2	0.8738	0.8468
C	$T_2 = 0.3$ s	1.25	0.96	0.500	-168	0.2500	0.2735
C	$T_2 = 3$ s	284.75	106.19	18.375	-166	0.7397	0.6237
C	$T_2 = 30$ s	4,750.25	267.65	29.500	-90	0.7786	0.6428
C	$T_2 = 300$ s	15,921.75	41.55	83.250	+0	0.8526	0.8327
C	$\sigma_R^2 = 0.02$	12,474.75	180.66	62.875	-30	0.8158	0.7708
C	$\sigma_R^2 = 0.08$	12,255.75	176.35	62.750	-30	0.8115	0.7676
C	$\sigma_R^2 = 0.2$	12,020.00	121.91	60.125	-34	0.8045	0.7606
C	$\sigma_R^2 = 0.5$	11,382.00	80.03	60.750	-24	0.7897	0.7450
D	$T_2 = 0.3$ s	4.25	3.30	0.750	+156	0.3819	0.3157
D	$T_2 = 3$ s	213.25	17.44	28.125	+160	0.7412	0.6310
D	$T_2 = 30$ s	5,264.25	137.15	43.625	+154	0.8130	0.6298
D	$T_2 = 300$ s	18,689.00	148.65	110.625	+44	0.8651	0.8343
D	$\sigma_R^2 = 0.02$	14,342.25	151.59	93.375	+66	0.8243	0.7737
D	$\sigma_R^2 = 0.08$	14,142.25	128.11	93.625	+52	0.8182	0.7699
D	$\sigma_R^2 = 0.2$	13,798.50	101.16	91.625	+52	0.8102	0.7630
D	$\sigma_R^2 = 0.5$	13,075.25	67.70	87.250	+58	0.7886	0.7462

corresponding factors are 3.35 and 3.55. Thus a coherence value that is already near a delivery plateau in a co-located test remains limiting when different link histories must be bridged. The large run-to-run spread at short coherence also shows why one favorable realization is an inadequate basis for a provisioning threshold.

Turbulence produces a smaller change within the particular configured range. For case A, integrated mean delivery decreases from 27 673.5 to 26 047.8 as σ_R^2 increases from 0.02 to 0.5; in case D it decreases from 14 342.3 to 13 075.3. These within-campaign sensitivities do not establish a universal ordering of turbulence and memory impairments. The complete numerical grid, including fidelity aggregation and rate maxima, is retained in Table 4.

The onboard case is a capability study. The synchronization concern in Ref. [2] concerns independent incoming photons without suitable space-qualified memories. Controlled recall from local storage changes the synchronization problem, but introduces memory, cooling, vacuum, optical-control, and qualification requirements. Space-borne memory proposals motivate assessing those requirements explicitly [16]. The cold-atom numerical settings are not a flight-qualified payload specification; selecting a rare-earth crystal would instead require its own cryogenic and transition-matching assumptions.

2.4.1 Complete single-repeater parameter grid

The full parameter grid consists of cases A–D, both control variables, four settings per variable, and four repetitions per

setting. Table 4 places integrated yield beside maximum ensemble rate, maximum plotted median-bin fidelity, and success-count-weighted fidelity. The maximum rate and fidelity need not occur in the same time bin. Figures 2 and 3 preserve every panel of the original grids.

Case windows are 300, 440, 340, and 440 s, respectively. Differences between cases consequently include both loading geometry and exposure duration. The two co-located cases also use distinct seeds and windows and are not a paired test of crossing angle. The table reports between-run sample dispersion, not confidence intervals estimated from adjacent bins. Neither the median-bin fidelity nor its count-weighted counterpart reconstructs a new individual-pair threshold count.

2.4.2 A dated Paris–Vienna pass: from optical geometry to delivered pairs

A companion S1 run uses Paris and Vienna city-center proxies, a circular 360 km design orbit, and an assumed midpoint of 25 September 2026 at 17:00 UTC. A 10° elevation mask gives approximately 312.7 s of common visibility during a 600 s observation. Both stations reach approximately 31.75° elevation. This run uses 256 modes per label and $T_2 = 100$ s, separating finite-bank behavior from the high-capacity controlled sweep.

Figure 4 combines the two downlink budgets with the delivery records generated for the same geometry. The per-arm mean loss near the midpoint is approximately 30.98 dB, including source, collection, atmosphere, configured mean




Figure 4: Dated S1 optical conditions and delivery dynamics. (a,b) Paris and Vienna downlink budgets. (c,d) Common visibility and slant range from the geometry used for the run. (e,f) Delivery rate and conditional mean fidelity from the 2 s counting bins. Time zero is the assumed pass midpoint. Rate and fidelity use bin-end times; attached optical geometry is sampled at bin midpoints. Fidelity gaps contain no deliveries. Bank occupancy is summarized in Section 2.4.2; the component budget and illumination context are given in Section 4.7.

turbulence, static pointing, coupling, and detection. Memory recall and BSM are separate operations. Of the 300 two-second bins, 157 contain deliveries, totaling 19 502 pairs: 32.503 pairs/s over the full observation, or 62.375 pairs/s when normalized by common visibility. The peak bin rate is 116.5 pairs/s, and success-count weighting of the bin mean fidelities gives 0.8550.

The peak Paris- and Vienna-side occupancies are 251 and 256 records, respectively, and the maximum combined occupancy is 321. The Vienna-side bank reaches its 256-mode capacity. At the end, the Paris-side bank is empty and the Vienna-side bank still holds 256 records. Those residuals are excluded from delivered counts. Capacity and expiry are separate controls: retained records need not be retrievable, and coherence decay does not itself release a slot. The successful-wait statistic averages 0.529 s after count weighting.

This case shows why bank-level monitoring complements an optical link budget. Nearly balanced pass geometry does not guarantee balanced inventory throughout generation and consumption. Admission, consumption, and clearing decisions must operate at each bank, rather than only on total memory capacity.

2.5 S2: Time-windowed service through a LEO constellation

Scenario S2 connects Luxembourg/Windhof to Singapore through time-varying satellite paths. The timing configuration uses 150 s snapshots, requests every 300 s, a 600 s planning horizon, and 30 s delivery cadence. Figure 5 shows the resulting service windows, fidelity, route chronology, and individual counted events.

The route timeline spans launches from 0 to 6900 s, mostly with three-hop routes and a final four-hop route. Quality remains near 0.78 during the plotted delivery windows. Between them, the original quality trace uses zero to indicate no delivery; it does not describe a zero-fidelity quantum state. Both count panels retain pairs-per-event units. A continuous-time throughput must instead sum these counts over a stated observation interval.

Connectivity becomes a service through both the contact plan and execution schedule. The four views expose this transformation more clearly than a path-availability value alone. A comparison of routing policies should hold the planning horizon, optical resources, source budgets, and traffic fixed. In particular, changing the graph horizon together with T_2 would mix the benefit of a larger search space with that of longer physical storage. The present trace demonstrates




Figure 5: S2 space-time distribution between Luxembourg/Windhof and Singapore. (a) Mean delivered pairs per event. (b) Quality during service windows; zero intervals mark no delivery. (c) Route timeline for successive requests. (d) Recorded delivery events. Snapshot duration, request spacing, planning horizon, and delivery cadence are 150, 300, 600, and 30 s, respectively. Counts remain in their original per-event units.

temporal service organization; the common policy-comparison contract is specified in Section 4.5.

2.6 S3: Ground-based swapping between a satellite downlink and fiber

Scenario S3 places the repeater at Windhof, between a satellite endpoint and a terrestrial endpoint at Kirchberg. Figure 6 defines the topology and its pass geometry. An 810 nm downlink loads the satellite-facing memory, while a 16.14 km fiber model connects the second elementary segment. A local Bell-state measurement joins the two resources. The numerical geometry has a near-zenith passage with peak elevation 88.7° and an approximately 760 s contact. Times are measured from simulation start.

This is a heterogeneous supply problem: the satellite segment is intermittent and strongly time dependent, whereas the fiber segment remains available with approximately constant attenuation. The two memory banks allow their elementary resources to become ready at different times before local swapping. Locating the repeater at the OGS also places receiving optics, storage controls, and spectral matching equipment on the ground. The fiber length is a simulation input, not a surveyed infrastructure length.

Figure 7 separates raw endpoint-pair delivery from delivery qualified by $F \geq 0.8$. The downlink loss falls from approximately 24 dB to 11 dB near closest approach. The raw delivery envelope follows this changing link opportunity, with its maximum near 380 s and an across-realization mean of roughly 50–60 pairs/s near the peak. The 1 s counting

realization resolves fluctuations around that envelope rather than a contact-wide constant rate.

The coherence sweep uses $T_2 = 0.6, 6, 60$, and 600 s. At 0.6 s, the conditional endpoint fidelity remains close to 0.49 and the displayed qualified yield is essentially zero. For 6 s, fidelity exceeds 0.8 only around the middle of the pass, producing a narrow qualified-delivery window. The 60 and 600 s cases stay above the threshold throughout the plotted contact, although their conditional fidelities differ away from closest approach. Their cumulative qualified counts nearly overlap at approximately 7.9×10^3 pairs, compared with approximately 2.0×10^3 for 6 s. These are rounded readings of the plotted outputs.

2.7 Matera–Helmos: Direct distribution and a fidelity-constrained optical operating point

The Matera–Helmos case sends the two photons of one on-board entangled-pair source through separate 0.25 m transmitters to 1.5 m and 2.3 m ground receivers. It provides the memory-free resource-production reference. A circular two-body orbit with semimajor axis 6939.13 km and inclination 97.64° gives about 293 s of shared visibility above 20° in the selected pass. Table 7 states the epoch, coordinate conventions, obscurations, and receiver assumptions.

The optical calculation uses an equivalent Gaussian beam, $M^2 = 1$, $1 \mu\text{rad}$ per-axis pointing uncertainty, and equal detection efficiency 0.60 at 810 and 1550 nm. Each detector has 50 ps RMS timing jitter. This compares wavelength-dependent propagation under equal receiver assumptions,




Figure 6: S3 architecture and contact geometry. (a) An 810 nm satellite downlink and a 16.14 km modeled fiber link meet at the Windhof repeater, which has separate memories and a Bell-state measurement (BSM). (b) Satellite ground track over the simulated contact. (c) Windhof elevation, reaching 88.7°. The track and times are scenario coordinates, rather than an operational acquisition schedule.

rather than particular detector technologies. At the joint-rate peak the two-arm loss including detection is approximately 22.78 dB at 810 nm and 28.80 dB at 1550 nm. These optimistic optical assumptions are stated separately from the $12 \mu\text{rad}$ divergence used in S1.

The conditional peak rates at a literal 10^9 emitted pairs/s are 5.27×10^6 and 1.32×10^6 genuine coincidences/s for ideal linear detectors. That high-brightness reference is not the selected service operating point: different-pair accidentals at a 1 ns coincidence window reduce the modeled mixture fidelity to approximately 0.62. Figure 8 shows how brightness, finite detector live time, and the fidelity target jointly restrict operation.

At target fidelity 0.90 and fixed 1 ns window, the upper allowed brightness is approximately 138.4 million pairs/s. With the four-detector, 50 ns dead-time approximation, the genuine-coincidence contributions are approximately 570 000/s at 810 nm and 160 000/s at 1550 nm. At target 0.95, they fall to approximately 271 000/s and 71 000/s. Joint brightness/window optimization gives conditional 0.90-target operating points near 490 million pairs/s and 266 ps at 810 nm, and 534 million pairs/s and 234 ps at 1550 nm. The corresponding genuine contributions are approximately 1.14×10^6 /s and 4.02×10^5 /s. Tables 9 and 10 retain the numerical values and assumptions.

The receiver field stop must affect signal and background consistently. In the receiver sensitivity in Methods, the same angular radius defines the encircled signal fraction and accepted sky solid angle. At 810 nm and 20° elevation, a

$10 \mu\text{rad}$ half-angle produces a combined field-stop penalty of approximately 0.50 dB for the prescribed sea-level-referenced turbulence profile, but 5.17 dB when its surface exponential is rebased to each station. This is a conditional, phase-only receiver sensitivity, not a new turbulent pass simulation. It shows why a field of view cannot be selected solely from a background count estimate; its signal acceptance must also be evaluated. Section 4.8.4 gives the calculation.

The comparison with experiment is likewise conditional. Yin et al. measured dual-downlink attenuation of 64–82 dB [5]. Their reported fine-tracking uncertainty was smaller than the pointing uncertainty assumed here, so improved pointing cannot explain the low reference loss. Divergence, aperture illumination, receiver throughput, and accepted optical modes must be compared at the same budget boundaries. For fixed small receivers in the far field, two-arm collection scales approximately as the inverse fourth power of divergence; changing 2 to $10 \mu\text{rad}$ alone costs about 28 dB under that approximation. The direct case therefore illustrates a coupled operating-point calculation, rather than a mission-performance advantage.

3 Discussion

The scenarios reveal different design constraints rather than one universal memory requirement. In S3, the 60 and 600 s coherence settings yield nearly the same number of pairs sat-




Figure 7: S3 causal loading and endpoint delivery, Part I. (a) Sampled and reference Windhof downlink loss. (b) Raw endpoint-pair delivery in 1 s bins, showing one counting realization and the across-realization mean. (c) Delivery satisfying $F \geq 0.8$ for $T_2 = 0.6, 6, 60$, and 600 s. (d) Count-weighted conditional endpoint Bell fidelity; the dotted line marks the qualification threshold. (e) Mean actual storage ages at BSM readout for the satellite-link and fiber-link OGS memories, on a logarithmic age axis. (f) Reserved and stored OGS modes in each bank. Part II continues with cumulative qualified counts, age distributions, and resource/policy controls.




Figure 8: Direct Matera-Helmos operating-point trade-off. Left: accepted-mixture fidelity against source brightness for a full 1 ns relative-time window. Right: genuine-coincidence contribution with four detectors per station and a 50 ns nonparalyzable dead-time approximation. Optical probabilities are frozen at the joint-rate peak. Equal detection efficiency and jitter are used at both wavelengths. The displayed rates are conditional receiver calculations, not stored-pair network delivery rates.

isfying $F \geq 0.8$, despite different conditional fidelities away from closest approach. The 6 s case meets the target only near the pass peak, while the subsecond case produces essentially no qualified deliveries. By contrast, the separated-terminal S1 geometries continue to benefit strongly between 30 and 300 s. A memory selected from a balanced, co-located test can therefore be underprovisioned for asynchronous link histories. Beyond a target-dependent count plateau, retrieval efficiency, loading balance, and BSM service time deserve comparison with further coherence improvements; S3's 50 ms control explicitly exposes a measurement-service bottleneck. These results favor joint dimensioning of optical links, mem-

ory banks, and operating policies rather than selecting each component independently.

The dated S1 pass connects optical budgets to counted deliveries and exposes a second constraint: per-bank inventory can remain imbalanced even when the pass geometry is nearly symmetric. In S3, millisecond satellite-link storage contrasts with second-scale fiber-link storage, while the fiber-facing OGS bank remains heavily occupied. The relevant control variables are not only total capacity and generation rate, but also which bank accepts, consumes, or clears each excitation and the quality of the resources still held. At constellation scale, the timing of a viable route and the cadence of its

execution jointly determine when the endpoints receive a service. Evaluating resource production and consumption on a common timeline is therefore necessary to distinguish geometrical connectivity from usable network availability.

The architecture supports this evaluation through separate physical, execution, and policy interfaces. Loss determines which resources survive, storage histories determine their quality, swaps consume their operands, and classical messages determine when an outcome is known. A richer physical model can replace a compact one at this interface without changing the definition of the delivered service. The reference instruments and detailed derivations in Methods specify this model boundary and support progressive calibration against experimental observations.

From simulation to testbed orchestration

The existing APIs provide a concrete integration point for orchestration protocols in quantum-network testbeds. QNETSIM exposes an HTTP/JSON event-action interface through which external controllers receive resource and execution reports and issue scheduling commands. The implemented messages include `ENT_REQUEST`, `REPORT_STORAGE`, and `SWAP_RESULT`, with corresponding actions for generation, storage, swapping, delivery, and request termination. Request and resource identifiers associate commands with their outcomes, while timestamps preserve the distinction between physical completion and controller knowledge.

This separation allows orchestration logic to be developed against simulated nodes and subsequently connected to physical-node adapters through the same logical service contract. The adapter translates device-specific observations and commands without placing hardware details inside the routing policy. A staged integration can retain simulated satellite links while incorporating a real controller or ground-node management stack, enabling repeatable tests of admission, scheduling, delayed acknowledgements, and failure recovery. Hardware-in-the-loop operation additionally requires explicit mapping between simulation time and laboratory clocks, authenticated interfaces, and device-specific timing validation. Fast optical synchronization, pointing, and wavefront correction remain local functions; the orchestration interface coordinates service-level operations rather than replacing those control loops.

KMS integration for satellite-assisted QKD networks

An important extension is coupling QNETSIM to a key management system (KMS) for satellite-assisted quantum key distribution (QKD). This direction is relevant to EuroQCI's integration of terrestrial and space segments and to the ground-service interfaces required by initiatives such as EAGLE-1 [6, 7]. The integration objective is to connect mission-dependent key production to key availability at authorized applications, while preserving the separation between quantum-network simulation, the satellite provider's service, and cryptographic key management.

The proposed coupling uses a provider-service adapter between QNETSIM and the KMS. QNETSIM supplies con-

tact opportunities, link conditions, and time-resolved key-production estimates from an explicitly selected QKD and post-processing model. In emulation, the adapter represents processing delays, the availability of matching key blocks at the intended endpoints, and the provider's key-identifier semantics. The KMS then manages key pooling, authorization, reservation, retrieval, and consumption. In a live deployment, the satellite or terrestrial QKD service remains authoritative for actual final-key availability; QNETSIM supplies predictions and receives operational metadata for model calibration. Coincidence counts or entanglement fidelities are not themselves secret-key blocks: secret-key accounting requires protocol-specific parameter estimation, reconciliation, privacy amplification, and the associated security assumptions [10].

Standardized interfaces provide complementary integration targets. ETSI GS QKD 014 specifies application-facing key delivery, ETSI GS QKD 015 specifies the interface between an SDN-enabled QKD node and its controller, and ETSI GS QKD 020 supports horizontal key transfer between interoperating key-management entities within a common trusted node [32–34]. Device and service adapters can map QNETSIM's orchestration abstractions to these interfaces while respecting their different scopes. The orchestrator requires availability, demand, allocation status, and completion metadata, rather than access to secret key values. Test key material must remain isolated from operational cryptographic services, and provider-specific internal key-relay mechanisms can remain behind their service interfaces.

This coupling would extend mission planning from estimating optical yield to evaluating sustained key-service availability. Predicted satellite contacts can be compared with application demand and the inventories of the participating KMSs to assess reserve requirements, depletion risk between passes, request latency, and prioritization under limited supply. Delayed post-processing, weather-related contact loss, authentication costs, and trusted-relay consumption can then be included in the same operational study. Unlike stored entanglement, classical key inventories do not undergo quantum decoherence; their availability is governed by production, allocation, consumption, and security policy. These distinctions make the integration useful both for EuroQCI-oriented testbeds and for other heterogeneous satellite-terrestrial QKD networks.

Scope and limitations

The numerical campaigns characterize their declared configurations. The S1 runs use compact state evolution; the exact Bell-vector and general instrument formulations are physical reference models, not retrospective reinterpretations of those curves. The S3 presentation distinguishes raw and fidelity-qualified delivery, storage ages, memory occupancy, and BSM resource/policy controls. It includes one counting realization and an across-realization mean, but no confidence interval or replication count is specified for these panels. Routing-policy gains, native-engine versus independent-simulator agreement, and local versus external-controller equivalence require matched experiments beyond these outputs. The direct-link study covers one pass with assumed receiver parameters; its frozen-peak optimization

and separate field-stop sensitivity are not a seasonal, detector-specific, or finite-key mission prediction. Space-qualified hardware, spectral conversion, and heralded loading require platform-specific calibration. The event-action APIs establish the software integration points; operation with physical testbed devices and a production KMS, including clock synchronization, interface conformance, and security controls, requires separate integration and validation.

QNETSIM thus connects the resource and networking stages of the satellite-quantum-internet roadmap [2], while providing an interface-based path from mission studies to testbed orchestration and key-service evaluation. Its memory-assisted scenarios show why temporal resource accounting matters, and the direct case illustrates how optical performance must be qualified by receiver statistics and state quality. Releasing configurations, counting conventions, interface definitions, and figure-generation data alongside the simulator makes these design decisions inspectable and supports their subsequent assessment in hybrid experimental networks.

4 Methods

The physical interfaces, derivations, numerical configurations, and verification procedures are collected here. The overview establishes the common calculation boundaries; the subsequent subsections provide the detailed memory, measurement, optical, and routing models and the records underlying the scenario results.

4.1 Overview of physical interfaces and calculation conventions

4.1.1 Quantum memory operations and measurement models

The physical reference begins with an optical encoding and an open matter–light model. Projection onto accepted output modes defines a completely positive, trace-nonincreasing memory operation

$$\mathcal{J}_s(\rho) = \sum_{\mu} B_{\mu}\rho B_{\mu}^{\dagger}, \quad \mathcal{Q}_s = \sum_{\mu} B_{\mu}^{\dagger}B_{\mu} \leq \mathbb{I}. \quad (1)$$

Its trace is the probability of successful recall. An orthogonal erasure outcome completes the operation into a trace-preserving channel. Balanced mode efficiencies and conditional pure dephasing give the reduced channel

$$\mathcal{M}_{\tau}(\rho) = \eta(\tau)\mathcal{D}_{\lambda_L(\tau)}(\rho) + [1 - \eta(\tau)]\text{Tr}(\rho)|\emptyset\rangle\langle\emptyset|, \quad (2)$$

$$\mathcal{D}_{\lambda_L}(\rho) = \frac{1 + \lambda_L}{2}\rho + \frac{1 - \lambda_L}{2}Z\rho Z. \quad (3)$$

Here η describes non-erasure probability and λ_L describes surviving logical coherence. The Kraus operators are the appropriately embedded $\sqrt{\eta(1 + \lambda_L)}/2\mathbb{I}$, $\sqrt{\eta(1 - \lambda_L)}/2Z$, and $\sqrt{1 - \eta}|\emptyset\rangle\langle j|$ for $j = 0, 1$. After compensating a deterministic unitary, the conditional entanglement fidelity is $(1 + \lambda_L)/2$, whereas the unconditional overlap is $\eta(1 + \lambda_L)/2$.

Section 4.2 develops separate reductions for cold atoms, rare-earth-ion-doped crystals, and trapped ions. Cold-atom storage

starts from Maxwell–Bloch equations for the signal field, optical polarization, and spin wave, with optical depth and control pulses determining the write/read kernels [13, 14, 35]. The ideal backward-CRIB expression $(1 - e^{-d})^2$ belongs to its specified echo model [36, 37]. Dynamical decoupling modifies the spin-noise filter and its coherence function [38]. Trapped-ion interfaces retain finite-dimensional internal dynamics, optical coupling, and effective dissipative operators [39, 40]. A fitted T_2 law is distinguished from separately modeled population loss or retrieval leakage.

For monotone $\lambda_L(\tau) = e^{-\chi(\tau)}$, an already evolved state aged from a to b receives

$$\lambda_L(b | a) = e^{-\chi(b) + \chi(a)}. \quad (4)$$

This preserves the physical age when events subdivide storage. Replacing it by $e^{-\chi(b-a)}$ resets a nonexponential history.

In the Bell-diagonal representation, let p_g denote probabilities labeled by $g \in \mathbb{Z}_2^2$. Ideal swapping of independent elementary resources, followed by the known Pauli correction, gives

$$p_g^{\text{out}} = \sum_h p_h^{(1)} p_{g \oplus h}^{(2)}. \quad (5)$$

For Werner inputs this reduces to $F_{\text{out}} = F_1 F_2 + (1 - F_1)(1 - F_2)/3$. Pure phase damping remains distinct from isotropic depolarization. Section 4.3 derives the click-pattern instrument, its conditioning, and the scope of compact alternatives.

For ideal indistinguishability and the stated two-of-four passive analyzer, a reduced genuine-herald probability is

$$p_b = \frac{1}{2}\eta_{r,1}\eta_{r,2}\eta_{o,1}\eta_{o,2}\eta_{d,1}\eta_{d,2}A_t, \quad (6)$$

where recall, post-recall optics, detection, and timing acceptance have separate boundaries. The factor $1/2$ reflects the restricted linear-optical analyzer [15, 41]; any loss already sampled upstream is omitted here. General distinguishability and false-herald effects require the accepted measurement instrument rather than assigning one visibility parameter to both success and conditional fidelity.

4.1.2 Optical channels, receiver coupling, and orbital geometry

For a Gaussian beam of radius w_L at a circular receiver of diameter D_r , geometrical collection is $1 - \exp[-D_r^2/(2w_L^2)]$. Extinction, pointing, aperture collection, mode coupling, and detection define distinct accounting boundaries. Random pointing can be integrated into an effective spot size or sampled as an offset, but not counted twice. Annular receivers require integration over the clear annulus rather than a uniform-area correction to a centrally peaked beam [4].

The detailed atmospheric model uses the Hufnagel–Valley profile, Fried coherence length, and Rytov diagnostics with explicit elevation and altitude conventions [42, 43]. The Gamma–Gamma description follows its normalized-irradiance formulation [44]. A unit-mean fading factor alone does not imply an additional mean attenuation. Finite-aperture propagation, wavefront correction, and accepted-mode coupling determine that loss; the field-based construction and reduced model choices are given in Section 4.4 [9, 45].

For receiver angular acceptance θ , the conditional signal fraction is $E_j(\theta) = \int_{|\alpha| \leq \theta} P_j(\alpha) d^2\alpha$, with a normalized receiving point-spread function P_j . The same angle sets the sky solid angle $2\pi(1 - \cos \theta)$. Fiber intensity transmission is $10^{-[\alpha(\lambda_0)L + \ell_{\text{ins}}]/10}$, with attenuation and insertion loss in dB. Its spectral and polarization transfer, conversion noise, and matching to the memory transition are additional interface properties.

Geometry supplies position, range, elevation, visibility, and propagation time. The dated cases state epoch, coordinate and time conventions, Earth model, propagator, and station positions. A circular design orbit and a controlled local-plane pass answer different questions; neither is relabeled as a reconstructed mission ephemeris [46, 47]. Original memory-write times are independent of the snapshot bins used for routing.

4.1.3 Space-time routing and causal resource execution

STBD represents changing connectivity by node-time vertices, communication opportunities, and memory-waiting edges [29]. For declared fidelity and memory utilities,

$$U_e = (U_e^{\text{fid}})^{\omega_f} (U_e^{\text{mem}})^{\omega_m}, \quad \omega_f + \omega_m = 1, \quad (7)$$

$$U(P) = \prod_{e \in P} U_e, \quad c_e = -\ln U_e. \quad (8)$$

The additive cost ranks routes. A realizable plan must additionally reserve both endpoints of each stored pair, respect optical-terminal and source budgets, provide two local operands for every BSM, and propagate final quality. A communication edge creates a two-ended resource; it does not move an unknown qubit through the graph as if it were a classical packet. Heralded storage must preserve that excitation [48].

Binary nesting joins adjacent ready segments. Segmented nesting completes snapshot-local segments early and carries surviving endpoints across later opportunities. An ℓ -link route needs $\ell - 1$ successful swaps, with failures consuming the measured operands. The full algorithm, resource-state updates, and computational boundaries appear in Section 4.5. The usable-pair estimator is

$$R_u = \frac{1}{T_{\text{obs}}} \sum_j \mathbf{1}\{j \text{ delivered}, F_j \geq F_{\text{min}}, t_j \leq D_j\}. \quad (9)$$

Per-link threshold checks do not replace this final boundary: four independent Werner links of fidelity 0.9 yield approximately 0.6731 after ideal swapping.

4.1.4 Direct-link coincidence and live-time calculation

For CW Poisson pair brightness R , per-arm detection probabilities η_M, η_H , station noise rates N_M, N_H , and a full relative window τ ,

$$g(\tau) = \text{erf}\left[\frac{\tau}{2\sqrt{2}\sigma_\Delta}\right], \quad (10)$$

$$C = R\eta_M\eta_H g(\tau), \quad (11)$$

$$A = (R\eta_M + N_M)(R\eta_H + N_H)\tau, \quad (12)$$

$$F_{\text{mix}} = \frac{F_0 C + A/4}{C + A}. \quad (13)$$

The accidental model assumes all-pairs association and a completely mixed effective accidental state [11]. Sky counts follow entrance radiance times clear area, accepted solid angle, filter bandwidth, optical/detection efficiencies, and inverse photon energy [12]. For target $1/4 < F_\star < F_0$, the feasible brightness interval is obtained from $R^2 - KR + b_M b_H \leq 0$, where $b_j = N_j/\eta_j$ and $K = [(F_0 - F_\star)/(F_\star - 1/4)]g/\tau - b_M - b_H$.

With four evenly loaded nonparalyzable detectors per station, the independent live fractions are

$$\ell_j = [1 + (R\eta_j + N_j)t_d/4]^{-1}, \quad C_{\text{live}} = C\ell_M\ell_H. \quad (14)$$

The argument is detected singles before dead-time suppression. The optimization maximizes C_{live} subject to the mixture-fidelity constraint. Section 4.8 specifies the numerical roots, timing range, equal-efficiency assumptions, and the separate field-stop sensitivity.

4.2 Quantum memory models: field modes, open dynamics, and Kraus operators

Scope and conventions

This subsection derives quantum-memory instruments for REID crystals, cold atomic ensembles, and trapped ions. Starting from the optical encoding and open matter-light dynamics, it defines accepted output modes, non-erasure probability, conditional state, and Kraus operators. Analytical retrieval and coherence laws are then identified as reduced models with stated assumptions, not unique microscopic solutions. Optical-depth conventions, independent-loss and coupled-recall alternatives, and the noise observable defining T_2 are kept explicit.

The optical carrier wavelength is denoted λ_0 ; the dimensionless logical coherence factor is $\lambda_L(\tau)$. The storage duration is τ , not the orbital time. Optical depth is d ; the number of optical interface modes is M_{int} and the number of simultaneously occupied memory slots is M_{cap} . These are different quantities. A platform's native T_2 is the $1/e$ time of the specified coherence observable, not an automatic erasure lifetime or request timeout. The different native optical and memory parameters are highlighted in Table.5

Quantized modes and the encoded input

For orthonormal transverse/polarization modes u_μ , introduce frequency operators satisfying

$$[\hat{a}_\mu(\omega), \hat{a}_\nu^\dagger(\omega')] = \delta_{\mu\nu}\delta(\omega - \omega'). \quad (15)$$

A normalized wavepacket and its one-photon state are

$$\hat{a}_\mu^\dagger[f] = \int d\omega f(\omega) \hat{a}_\mu^\dagger(\omega), \quad \int |f(\omega)|^2 d\omega = 1. \quad (16)$$

The positive-frequency field is an expansion in these modes,

$$\hat{E}^{(+)}(\mathbf{r}, t) = \sum_\mu \mathbf{e}_\mu \int d\omega \mathcal{E}_\omega u_\mu(\mathbf{r}_\perp, \omega) \hat{a}_\mu(\omega) e^{i[k(\omega)z - \omega t]}. \quad (17)$$

The normalization factor $\mathcal{E}_\omega$ depends on the chosen continuum convention; its product with the transition dipole determines the coupling g_μ . The logical input is encoded by the isometry

$$V_{\text{in}}|j\rangle = \hat{a}_j^\dagger[f_j]|\text{vac}\rangle, \quad j = 0, 1, \quad |\psi\rangle = \alpha|0\rangle + \beta|1\rangle. \quad (18)$$

Two polarization components, two time bins, or two rails require an interface that preserves their relative amplitude and phase. A single three-level transition, by itself, specifies only one optical coupling. An additional rail or a suitable multilevel implementation must be stated to obtain the full polarization-qubit memory. This is an encoding requirement, not extra multiplexing capacity.

Open-system evolution and the exact reduction

Let S include the accepted optical modes and the material degrees of freedom. Initially assume $\rho_{SB} = V_{\text{in}}\rho_{\text{in}}^\dagger \otimes \rho_B$ and write

$$H_{\text{tot}}(t) = H_S(t) + H_B + H_{SB}, \quad (19)$$

$$H_B = \sum_\alpha \hbar\omega_\alpha b_\alpha^\dagger b_\alpha, \quad H_{SB} = \sum_\alpha Q_\alpha \otimes B_\alpha. \quad (20)$$

The exact propagator is time ordered. For $\rho_B = \sum_\nu p_\nu |\nu\rangle\langle\nu|$, a Kraus family for the reduced system evolution is

$$E_{\alpha\nu}(\tau) = \sqrt{p_\nu} \langle\alpha|\mathcal{T} e^{-i \int H_{\text{tot}}(t) dt/\hbar} |\nu\rangle, \quad \sum_{\alpha\nu} E_{\alpha\nu}^\dagger E_{\alpha\nu} = \mathbb{I}. \quad (21)$$

Preparation of the material ground state and the write/read intervals are included in this evolution. This equation is the formal field-plus-bath to Kraus connection.

When weak system–bath coupling, short bath correlations, and the secular approximation are appropriate, the reduced dynamics take the time-dependent master-equation form

$$\dot{\varrho} = -\frac{i}{\hbar}[H_S + H_{\text{LS}}, \varrho] + \sum_\alpha \mathcal{D}[L_\alpha]\varrho, \quad (22)$$

$$\mathcal{D}[L]\varrho = L\varrho L^\dagger - \frac{1}{2}\{L^\dagger L, \varrho\}. \quad (23)$$

Classical colored detuning noise can instead be sampled in $H_S(t)$, followed by ensemble averaging. The same noise contribution must not also be included in a fitted Lindblad rate. A stretched-exponential coherence law is not evidence that the underlying bath is Markovian.

Let V_{out} embed the accepted logical output. Projection of the reduced evolution gives successful operators $B_r = V_{\text{out}}^\dagger E_r V_{\text{in}}$ and

$$\mathcal{J}_s(\rho) = \sum_r B_r \rho B_r^\dagger, \quad Q_s = \sum_r B_r^\dagger B_r \leq \mathbb{I}_2. \quad (24)$$

All rejected output sectors can be deliberately coarse-grained into an orthogonal flag $|\emptyset\rangle$. With V embedding the two-dimensional accepted space into the three-dimensional flagged space, one complete Kraus family is

$$K_{s,r} = V B_r, \quad K_{f,j} = |\emptyset\rangle\langle j| \sqrt{\mathbb{I}_2 - Q_s}, \quad j = 0, 1. \quad (25)$$

Indeed, the successful effects sum to Q_s and the failure effects to $\mathbb{I}_2 - Q_s$. The operational quantities are

$$p_s(\rho) = \text{Tr}(Q_s \rho), \quad \rho_s = \mathcal{J}_s(\rho)/p_s(\rho). \quad (26)$$

Conditioning is undefined when $p_s = 0$. The normalization is generally input dependent; the conditional transformation is a linear trace-preserving channel only in special cases, including $Q_s = \eta \mathbb{I}_2$.

From the field transfer kernel to a survival operator

For a fixed classical realization ξ , a passive linearized, vacuum-noise memory can be written

$$\hat{a}_\mu^{\text{out}}(t) = \sum_\nu \int dt' \mathcal{K}_{\mu\nu}^{(\xi)}(t, t'; \tau) \hat{a}_\nu^{\text{in}}(t') + \hat{v}_\mu(t). \quad (27)$$

The Langevin/environment term $\hat{v}$ restores the output commutators and accounts for lost excitations. Projecting onto normalized incoming modes f_ν and accepted outgoing modes g_μ yields

$$A_{\mu\nu}^{(\xi)}(\tau) = \iint dt dt' g_\mu^*(t) \mathcal{K}_{\mu\nu}^{(\xi)}(t, t'; \tau) f_\nu(t'). \quad (28)$$

For a passive, number-preserving system–environment dilation with vacuum environment and a single input excitation, A_ξ is a contraction. The successful map is

$$\mathcal{J}_s(\rho) = \mathbb{E}_\xi[A_\xi \rho A_\xi^\dagger], \quad Q_s = \mathbb{E}_\xi[A_\xi^\dagger A_\xi]. \quad (29)$$

One averages the unnormalized operation and then conditions. In general $\mathbb{E}[A \rho A^\dagger] \neq \mathbb{E}[A] \rho \mathbb{E}[A]^\dagger$.

This reduction requires more than linear equations: the noise state and commutators must be specified. A mixture of Gaussian channels need not itself be Gaussian. A single ion is not a bosonic ensemble merely because a single-excitation calculation is linear. Added photons from four-wave mixing, fluorescence, thermal excitation, or imperfect optical pumping require further operators and detector-conditioned sectors; they are not captured by a vacuum-erasure model. The passive, vacuum-noise reduction is used only under these conditions.

Balanced-rail dephasing plus erasure: explicit Kraus construction

Suppose the accepted operation has no rail mixing and write

$$A_\xi = U_{\text{det}} \text{diag}(a_{0,\xi}, a_{1,\xi}), \quad (30)$$

$$\eta_j = \mathbb{E}|a_{j,\xi}|^2, \quad \zeta = \mathbb{E}[a_{0,\xi} a_{1,\xi}^*]. \quad (31)$$

For balanced efficiencies $\eta_0 = \eta_1 = \eta > 0$ and a compensated mean phase, set $\lambda_L = \zeta/\eta \in [0, 1]$. Then

$$\mathcal{J}_s(\rho) = \eta U_{\text{det}} \begin{pmatrix} \rho_{00} & \lambda_L \rho_{01} \\ \lambda_L \rho_{10} & \rho_{11} \end{pmatrix} U_{\text{det}}^\dagger. \quad (32)$$

Since $Z\rho Z$ changes the sign of the off-diagonal elements, this equals $\eta U_{\text{det}} \mathcal{D}_{\lambda_L}(\rho) U_{\text{det}}^\dagger$, where

$$\mathcal{D}_{\lambda_L}(\rho) = \frac{1 + \lambda_L}{2} \rho + \frac{1 - \lambda_L}{2} Z \rho Z. \quad (33)$$

Thus the four operators used in the balanced-rail reduced channel are

$$K_0 = \sqrt{\eta(1 + \lambda_L)/2} V U_{\text{det}}, \quad K_1 = \sqrt{\eta(1 - \lambda_L)/2} V U_{\text{det}} Z, \quad (34)$$

$$K_2 = \sqrt{1 - \eta} |\emptyset\rangle\langle 0|, \quad K_3 = \sqrt{1 - \eta} |\emptyset\rangle\langle 1|. \quad (35)$$

Their completeness is $\sum_j K_j^\dagger K_j = \eta \mathbb{I}_2 + (1 - \eta) \mathbb{I}_2 = \mathbb{I}_2$. Identifying $\lambda_L = e^{-\chi}$ gives $p_\phi = (1 - e^{-\chi})/2$.

For a maximally entangled reference input, after correcting U_{det} ,

$$F_{e|s} = \frac{1 + \lambda_L}{2}, \quad F_e = \eta \frac{1 + \lambda_L}{2}, \quad p_s = \eta. \quad (36)$$

The conditional average pure-state fidelity, a different quantity, is $F_{\text{av}|s} = (2F_{e|s} + 1)/3$. A phase-dephased ideal Bell pair approaches $F_{e|s} = 1/2$, not $1/4$. The latter is the isotropic depolarization limit and belongs to a different channel.

Unequal rails. Cauchy–Schwarz gives $|\zeta| \leq \sqrt{\eta_0 \eta_1}$. After removing the phase of ζ , let $\lambda_u = |\zeta|/\sqrt{\eta_0 \eta_1}$. Replace the successful operators in Eq. (35) by

$$K_0 = \sqrt{(1 + \lambda_u)/2} V \text{diag}(\sqrt{\eta_0}, \sqrt{\eta_1}), \quad (37)$$

$$K_1 = \sqrt{(1 - \lambda_u)/2} V \text{diag}(\sqrt{\eta_0}, -\sqrt{\eta_1}), \quad (38)$$

and the failure factors by $\sqrt{1 - \eta_j}$. Known output phases/unitaries can be appended. Now

$$p_s(\rho) = \eta_0 \rho_{00} + \eta_1 \rho_{11}, \quad (39)$$

$$F_{e|s} = \frac{\eta_0 + \eta_1 + 2|\zeta|}{2(\eta_0 + \eta_1)} \quad (40)$$

for the maximally entangled reference. A scalar efficiency is insufficient when the rails have unequal survival or couple differently to the bath.

Common attenuation is not logical dephasing. If $a_{0,\xi} = a_{1,\xi} = a_\xi$, then $\zeta = \eta$ and $\lambda_L = 1$, even when $\eta \ll 1$ or the common phase fluctuates. Conversely, a random relative phase can give $\lambda_L < 1$ with no loss. This establishes why the retrieval envelope and conditional coherence must be separately identified before imposing any relation between them.

4.2.1 REID spin-wave CRIB and dynamical decoupling

Material platform and matter–light Hamiltonian. Consider $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ with an optical interface near 605.977 nm, controlled reversible inhomogeneous broadening (CRIB), and a long-lived hyperfine spin wave protected by a clock/ZEFOZ operating point and dynamical decoupling. The storage level $|s\rangle$ and transfer control field connect optical excitation to the spin-wave stage:

$$H/\hbar = H_{\text{field}}/\hbar + \sum_j [\Delta_j(t) \sigma_{ee}^{(j)} + \delta_j(t) \sigma_{ss}^{(j)}] - \sum_j \left[\sum_\mu g_{\mu j} \hat{a}_\mu(z_j) \sigma_{eg}^{(j)} + \Omega_j(t) \sigma_{es}^{(j)} + \text{h.c.} \right], \quad (41)$$

$$\Delta_j(t) = \Delta_j^{(0)} + \Delta_{\text{CRIB}}(z_j, t) + \delta \Delta_j(t). \quad (42)$$

Here $H_{\text{field}}/\hbar = \sum_\mu \int dz \hat{a}_\mu^\dagger (-ic \partial_z) \hat{a}_\mu$. The sign of the coupling is a phase convention. Optical spontaneous decay, spin leakage, and fluctuating detunings enter through Eq. (23) or

an explicit bath. With most ions in $|g\rangle$, the optical polarization obeys the linear optical Maxwell–Bloch equations; retaining Ω adds the spin-wave equation used below for ensembles. Immobility removes the atomic-diffusion term but not spectral diffusion or magnetic noise.

Optical depth and backward-echo amplitude. For ideal, phase-matched backward CRIB, complete detuning reversal, adequate prepared bandwidth, and lossless spin-transfer pulses, an excitation created at depth z has input and return amplitude weights $e^{-\alpha z/2}$. The coherent source contribution is proportional to αdz . Consequently the normalized retrieval amplitude is

$$A_{\text{CRIB}}^{(0)} = \int_0^L \alpha e^{-\alpha z} dz = 1 - e^{-d}, \quad d = \alpha L, \quad (43)$$

and

$$\eta_{\text{CRIB}}^{(0)}(d) = (1 - e^{-d})^2. \quad (44)$$

This is the ideal backward-CRIB limit under the stated assumptions [36, 37]. Imperfect preparation, irreversible optical dephasing, insufficient bandwidth, imperfect transfer pulses, and collection must enter the actual kernel or additional efficiencies. Increasing d does not remove these errors. Atomic-frequency-comb storage is a distinct echo protocol, not another name for CRIB [49].

From the spin bath to the DD coherence factor. In a logical toggling frame, use $H_\phi(t) = \hbar y(t) \delta \omega(t) Z/2$, where ideal π pulses change $y(t)$ between $+1$ and -1 . For one realization,

$$\rho_{01}(\tau) = e^{-i\phi(\tau)} \rho_{01}(0), \quad \phi(\tau) = \int_0^\tau y(t) \delta \omega(t) dt. \quad (45)$$

For zero-mean Gaussian noise with correlation $C(t - t')$, the Gaussian characteristic function gives

$$\lambda_L(\tau) = e^{-\chi(\tau)}, \quad \chi = \frac{1}{2} \int_0^\tau \int_0^\tau y(t) y(t') C(t - t') dt dt'. \quad (46)$$

With the two-sided spectrum $S_2(\omega) = \int_{-\infty}^\infty C(u) e^{-i\omega u} du$ and $Y(\omega) = \int_0^\tau y(t) e^{i\omega t} dt$,

$$\chi = \frac{1}{4\pi} \int_{-\infty}^\infty S_2(\omega) |Y(\omega)|^2 d\omega. \quad (47)$$

For a classical real process, the positive-frequency form is $\chi = (2\pi)^{-1} \int_0^\infty S_2(\omega) |Y(\omega)|^2 d\omega$. An equivalent convention defines $F(\omega\tau) = \omega Y(\omega)$ and $S_F = S_2/2$, giving $\chi = \pi^{-1} \int_0^\infty S_F |F|^2 / \omega^2 d\omega$. The spectral definition must accompany the prefactor [38].

At a ZEFOZ point the first derivative of the transition frequency with respect to magnetic field vanishes. Residual quadratic magnetic shifts, pulse errors, and technical noise remain. A fitted coherence and efficiency model

$$\chi_R(\tau) = \left(\frac{\tau}{T_{2,R}^{\text{DD}}} \right)^{n_R}, \quad \eta_{\text{DD}} = e^{-(\tau/T_{\text{mem}})^{m_R}} \quad (48)$$

therefore replaces a specified noise spectrum and pulse sequence over a declared fitting interval. m_R is the efficiency-envelope exponent, not the mode capacity. Non-Gaussian noise may require higher cumulants than Eq. (46).

Independent-loss and coupled-recall variants. The independent-loss variant uses

$$\eta_R^{(0)}(\tau) = \eta_{\text{CRIB}}^{(0)}(d)\eta_{\text{DD}}(\tau), \quad (49)$$

whereas the coupled-recall variant uses

$$\eta_R^{(2)}(\tau) = \eta_R^{(0)}(\tau)e^{-2\chi_R(\tau)}. \quad (50)$$

Both are retained as explicitly different effective models. Each gives a CPTP instrument when inserted with $\lambda_L = e^{-\chi_R}$ into Eq. (35). Only the first is implied by the scalar erasure-plus-dephasing form with independent envelopes. The second additionally assumes that recoverable collective-mode population decays with the square of the same fitted coherence amplitude. That extra identification needs a microscopic correlation model or calibration; the general field reduction does not impose it. Optional population loss adds $e^{-\tau/T_{\text{loss}}}$ when it is genuinely leakage out of both logical rails. The same loss must not be charged again at recall.

4.2.2 Cold-atom EIT/Raman memory

Hamiltonian, linearization, and noise. For an atomic Λ system (^{87}Rb as its example), replace the REID detunings by a one-photon detuning Δ and two-photon detunings $\delta_j(t)$, and retain the quantum coupling and classical control $\Omega(t)$. Define coarse-grained collective polarization $P \propto \sum_j \sigma_{ge}^{(j)}$ and spin wave $S \propto \sum_j \sigma_{gs}^{(j)}$. Weak excitation lets the population differences be replaced by their ground-state values; an appropriate collective normalization gives

$$(\partial_t + c\partial_z)\hat{\mathcal{E}} = iG\hat{P}, \quad (51)$$

$$\partial_t\hat{P} = -(\gamma + i\Delta)\hat{P} + iG\hat{\mathcal{E}} + i\Omega\hat{S} + \hat{F}_P, \quad (52)$$

$$\partial_t\hat{S} = -(\gamma_s + i\delta)\hat{S} + i\Omega^*\hat{P} + \hat{F}_S. \quad (53)$$

Here γ is an *amplitude* decay rate, and G is the collective coupling in the selected normalization. The reservoir correlations, not the mere presence of $\hat{F}_P$, determine whether added noise is negligible. EIT and off-resonant Raman storage are different operating limits of these equations, although an ideal optimization can relate their maximum efficiencies [13, 35].

A factor-of-two optical-depth convention. Set $\Omega = \Delta = 0$ and eliminate the stationary polarization: $P = iG\mathcal{E}/\gamma$. Then

$$\partial_z\mathcal{E} = -\frac{G^2}{c\gamma}\mathcal{E}, \quad \frac{I(L)}{I(0)} = e^{-2G^2L/(c\gamma)}. \quad (54)$$

Matching the intensity optical-depth convention $I(L)/I(0) = e^{-d}$ gives $d = 2G^2L/(c\gamma)$. The free-space solution of Gorshkov *et al.* uses $d_G = G^2L/(c\gamma)$, so $d_G = d/2$ for these normalizations [14]. This conversion must precede importing a storage kernel. Otherwise even a numerically exact kernel can be evaluated at the wrong physical optical depth.

EIT storage as a change of field character. On two-photon resonance with slowly varying, fixed-phase control, the dark excitation has negligible optical polarization:

$$\hat{\Psi} = \cos\vartheta\hat{\mathcal{E}} - \sin\vartheta\hat{S}, \quad \tan\vartheta = G/|\Omega|, \quad v_g = c\cos^2\vartheta. \quad (55)$$

Reducing $|\Omega|$ converts the propagating dark excitation to a material spin wave. Reversing the control converts it back. Pulse bandwidth, finite optical depth, nonadiabatic coupling to bright states, and spin decoherence limit the efficiency; a dark-state expression alone does not determine the write/read probability.

Explicit kernel derivation in the adiabatic limit. Introduce retarded dimensionless time $s = \gamma(t - z/c)$, coordinate $x = z/L$ for uniform density, $\bar{\Delta} = \Delta/\gamma$, and $\bar{\Omega} = \Omega/\gamma$. Neglect spin decay only during the short write/read step. With normalized mode amplitudes the equations become

$$\partial_x\mathcal{E} = i\sqrt{d_G}P, \quad (56)$$

$$\partial_sP = -(1 + i\bar{\Delta})P + i\sqrt{d_G}\mathcal{E} + i\bar{\Omega}S, \quad (57)$$

$$\partial_sS = i\bar{\Omega}^*P. \quad (58)$$

Laplace transforming $x \mapsto u$ yields

$$\mathcal{E}(u, s) = i\sqrt{d_G}P(u, s)/u + \mathcal{E}_{\text{in}}(s)/u, \quad (59)$$

$$D(u) = 1 + i\bar{\Delta} + d_G/u, \quad (60)$$

$$\partial_sS = -\frac{|\bar{\Omega}|^2}{D(u)}S - \frac{\bar{\Omega}^*\sqrt{d_G}}{uD(u)}\mathcal{E}_{\text{in}}, \quad (61)$$

where the last equation uses adiabatic elimination of P . This scalar inhomogeneous equation is solved with an integrating factor. Define $h(s_1, s_2) = \int_{s_1}^{s_2} |\bar{\Omega}(v)|^2 dv$ and $q = 1 + i\bar{\Delta}$. The inverse Laplace transform gives the write kernel

$$S(x, s_w) = \int_0^{s_w} W(x, s)\mathcal{E}_{\text{in}}(s)ds, \quad (62)$$

$$W(x, s) = -\frac{\sqrt{d_G}\bar{\Omega}^*(s)}{q}e^{-[h(s, s_w) + d_Gx]/q} \times I_0\left(\frac{2\sqrt{h(s, s_w)d_Gx}}{q}\right). \quad (63)$$

Here I_0 is the modified Bessel function. The corresponding forward read kernel, with read time reset to zero and $S_r(x)$ the initial spin-wave profile, is

$$\mathcal{E}_{\text{out}}(s) = \int_0^1 R(s, x)S_r(x)dx, \quad (64)$$

$$R(s, x) = -\frac{\sqrt{d_G}\bar{\Omega}(s)}{q}e^{-[h(0, s) + d_G(1-x)]/q} \times I_0\left(\frac{2\sqrt{h(0, s)d_G(1-x)}}{q}\right). \quad (65)$$

A global phase convention is immaterial to efficiencies. Backward readout reverses the spin-wave coordinate and must also satisfy the wavevector-matching condition. These are free-space write and read kernels under the boundary and adiabatic assumptions just stated [14].

Let $D_{\text{sw}}(x, x'; \tau)$ describe storage evolution. The full write-store-read mode kernel is now explicitly

$$\mathcal{K}(s, s'; \tau) = \int_0^1 \int_0^1 R(s, x)D_{\text{sw}}(x, x'; \tau)W(x', s')dx dx'. \quad (66)$$

Projection by Eq. (28) and the open-system average by Eq. (29) complete the path to the Kraus operators. A solver can precompute $\mathcal{K}$ or its singular modes offline; DES events need only the reduced instrument and its parameter identifier.

For a normalized initial spin wave, completed retrieval without spin decay also admits the independent efficiency check

$$\eta_r = \int_0^1 \int_0^1 k_{d_G}(x, y) S_r(1-x) S_r^*(1-y) dx dy, \quad (67)$$

$$k_{d_G}(x, y) = \frac{d_G}{2} e^{-d_G(x+y)/2} I_0(d_G \sqrt{xy}). \quad (68)$$

It follows by integrating the polarization/spin excitation-balance equation over readout time. For a flat unit-norm spin wave,

$$\eta_{r,\text{flat}} = 1 - e^{-d_G} [I_0(d_G) + I_1(d_G)]. \quad (69)$$

This is retrieval from a prescribed spin wave, not the total write-read efficiency. The total efficiency includes $\eta_w = \int_0^1 |S(x, s_w)|^2 dx$ and the overlap of the retrieved field with the accepted temporal mode. A tabulated d alone cannot set all three.

Parametric efficiency and kernel-based efficiency. The compact law $\eta_{\text{EIT}}(d) \simeq (1 - e^{-d})^2$ is a *parametric* backward-recall surrogate, not a general consequence of EIT. Equations (63)–(69) define a kernel-based alternative. A simulation must select one efficiency model; the two are not interchangeable at fixed d .

Motion, magnetic inhomogeneity, and T_2 . The collective spin wave has wavevector $\mathbf{q} = \mathbf{k}_s - \mathbf{k}_c$. A displacement $\Delta \mathbf{r}$ contributes $e^{i\mathbf{q} \cdot \Delta \mathbf{r}}$ to its retrieval overlap. Under Gaussian diffusion with $\mathbb{E}[\Delta r_a \Delta r_b] = 2D\tau \delta_{ab}$,

$$C_{\text{mot}}(\tau) = \mathbb{E} e^{i\mathbf{q} \cdot \Delta \mathbf{r}} = e^{-D|\mathbf{q}|^2 \tau}. \quad (70)$$

Independent quasistatic two-photon detuning of variance σ_δ^2 contributes $e^{-\sigma_\delta^2 \tau^2/2}$; homogeneous spin decay contributes $e^{-\gamma_0 \tau}$. The combined functional is therefore

$$\chi_C(\tau) = a\tau + b\tau^2, \quad a = D|\mathbf{q}|^2 + \gamma_0, \quad b = \sigma_\delta^2/2. \quad (71)$$

Its $1/e$ time obeys $aT_2 + bT_2^2 = 1$, giving the stable expression

$$T_{2,C} = \frac{2}{a + \sqrt{a^2 + 4b}}. \quad (72)$$

For $a = b = 0$, the coherence time is infinite. For ballistic thermal motion, the different model $C_{\text{mot}} = \exp[-|\mathbf{q}|^2 (k_B T / m_a) \tau^2 / 2]$ replaces the diffusion factor. A trapping model may require neither limit. These alternatives must not be superimposed for the same motional mechanism.

An alternative fitted model sets $\chi_C = (\tau/T_{2,C})^{n_C}$ and, for the coupled-recall variant,

$$\eta_C^{(2)}(\tau) = \eta_{\text{EIT}}(d) \eta_{\text{mode}}(\tau) e^{-2\chi_C(\tau)}, \quad \lambda_{L,C} = e^{-\chi_C(\tau)}. \quad (73)$$

This coupled model identifies optical-mode survival and logical dephasing through the same functional. For two physical rails, that identification is an assumption about their noise correlations. In a kernel backend, motional loss of collective retrieval and differential logical noise are extracted from η_j and ζ . The mode factor $\eta_{\text{mode}} = \text{Tr}(M^\dagger M) / M_{\text{int}}$ is an average over a specified equally weighted input ensemble, not a universal state-independent efficiency when $M^\dagger M$ is non-scalar.

Native-parameter consistency test. Microscopic parameters and a fitted coherence time cannot be varied independently for the same mechanism. For example, $D = 10^{-4} - 10^{-3} \text{ m}^2/\text{s}$ and $|\mathbf{q}| = 10^4 - 10^6 \text{ m}^{-1}$ would place $(D|\mathbf{q}|^2)^{-1}$ between 10^{-9} and 10^{-4} s under the stated diffusion law. They cannot simultaneously describe millisecond-to-second motional coherence without changing the motion or signal-control geometry model. A run therefore uses either a consistent microscopic parameter set or a calibrated effective T_2, n_C pair.

4.2.3 Trapped-ion Raman interface and internal-state memory

Finite-dimensional matter dynamics. A single-ion Λ interface couples optical excitation to a long-lived internal state through a Raman interaction. For a cavity-assisted realization, the optical input/output boundary is specified by

$$H/\hbar = \delta_c a^\dagger a + \Delta |e\rangle\langle e| + \delta |s\rangle\langle s| + (ga|e\rangle\langle g| + \Omega(t)|e\rangle\langle s| + \text{h.c.}), \quad (74)$$

$$L_{\text{ex}} = \sqrt{\kappa_{\text{ex}}} a, \quad L_{\text{loss}} = \sqrt{\kappa_{\text{loss}}} a, \quad (75)$$

$$L_g = \sqrt{\Gamma_g} |g\rangle\langle e|, \quad L_s = \sqrt{\Gamma_s} |s\rangle\langle e|. \quad (76)$$

Here $\kappa = \kappa_{\text{ex}} + \kappa_{\text{loss}}$ and $\Gamma = \Gamma_g + \Gamma_s$ are *population* decay rates. The associated amplitude linewidths are $\kappa/2$ and $\Gamma/2$. Additional dark-state noise and leakage operators are specified separately. A polarization memory needs the multilevel/dual-rail extension identified above. No collective bosonic approximation is made for the ion.

An explicit single-photon capture and readout kernel.

The incoming continuum wavepacket can be reduced to a single-excitation amplitude equation for a fixed classical-noise realization. Let $\mathbf{v} = (c, e, s)^T$ denote the cavity-photon, optical-excited-state, and stored-state amplitudes. The coherent/no-jump branch obeys

$$\dot{\mathbf{v}} = \mathbf{A}(t)\mathbf{v} + \sqrt{\kappa_{\text{ex}}} \mathbf{e}_c f_{\text{in}}(t), \quad (77)$$

where $\mathbf{e}_c = (1, 0, 0)^T$, $\mathbf{e}_s = (0, 0, 1)^T$, and

$$\mathbf{A}(t) = \begin{pmatrix} -\kappa/2 - i\delta_c & -ig^* & 0 \\ -ig & -\Gamma/2 - i\Delta & -i\Omega(t) \\ 0 & -i\Omega^*(t) & -\gamma_\ell/2 - i\delta(t) \end{pmatrix}. \quad (78)$$

Here γ_ℓ is population leakage out of the selected storage state, not a relabeling of logical $1/T_2$. Input/output normalization gives $f_{\text{out}} = f_{\text{in}} - \sqrt{\kappa_{\text{ex}}} c$. For a normalized incident photon and initially empty cavity/unexcited ion, let $U_w(t, t')$ be the

time-ordered propagator of $\mathbf{A}$ during writing. The coherent storage amplitude is

$$s_w = \sqrt{\kappa_{\text{ex}}} \int_0^{t_w} \mathbf{e}_s^T \mathbf{U}_w(t_w, t') \mathbf{e}_c f_{\text{in}}(t') dt', \quad \eta_w^{\text{coh}} = |s_w|^2. \quad (79)$$

For retrieval from the unit storage excitation with vacuum optical input,

$$f_r(t) = -\sqrt{\kappa_{\text{ex}}} \mathbf{e}_c^T \mathbf{U}_r(t, 0) \mathbf{e}_s, \quad \eta_r^{\text{coh}} = \int |f_r(t)|^2 dt. \quad (80)$$

The accepted read-mode amplitude is $\int g_{\text{out}}^*(t) f_r(t) dt$, rather than $\sqrt{\eta_r^{\text{coh}}}$ unless the modes match. Multiplying the write amplitude, the storage propagator, and this read amplitude supplies $a_{j,\xi}$ for each logical rail and hence the quantities η_j, ζ in the Kraus reduction. This makes the dependence on $g, \Omega, \Delta, \kappa, \Gamma$, and the field mode explicit instead of treating η_{TI} as a function of T_2 alone.

The amplitude equations describe the coherent/no-jump contribution. Spontaneous jumps that return population to an accepted internal state can produce incoherent, unheralded errors; their trajectories or full master-equation contributions must be added through Eq. (24). They may be classified as erasure only when the specified herald rejects them. The equation is therefore a physical kernel construction.

Eliminating the excited state without discarding dissipation. Write all coherent operators in frequency units, $h = H/\hbar$. Split the excitation and de-excitation couplings as v_+ and $v_- = v_+^\dagger$, and define $h_{\text{NH}} = h_e - i \sum_j L_j^\dagger L_j / 2$. Adiabatic elimination gives

$$h_{\text{eff}} = h_g - \frac{1}{2} v_- [h_{\text{NH}}^{-1} + (h_{\text{NH}}^{-1})^\dagger] v_+, \quad (81)$$

$$L_{\text{eff},j} = L_j h_{\text{NH}}^{-1} v_+. \quad (82)$$

For $h_{\text{NH}} = (\Delta - i\Gamma/2)|e\rangle\langle e|$ and $v_+ = |e\rangle\langle g| + \Omega|s\rangle$, the coherent part contains

$$H_{\text{eff}}/\hbar \supset -\frac{\Delta}{\Delta^2 + (\Gamma/2)^2} [|g|^2 a^\dagger a |g\rangle\langle g| + |\Omega|^2 |s\rangle\langle s| + g^* \Omega a^\dagger |g\rangle\langle s| + \text{h.c.}], \quad (83)$$

where the Hermitian-conjugate term applies to the off-diagonal coupling. The effective spontaneous-emission operators include

$$L_{\text{eff},g} = \frac{\sqrt{\Gamma_g}}{\Delta - i\Gamma/2} |g\rangle\langle g| + \Omega|s\rangle, \quad (84)$$

and the analogous $|s\rangle$ channel. Thus the large-detuning limit gives the Raman coupling magnitude $|g\Omega/\Delta|$, together with AC Stark shifts and scattering rates proportional to $\Gamma|\Omega/\Delta|^2$. Retaining only the coherent coupling would omit precisely the open-system errors that the memory model is meant to capture [40].

Cooperativity rather than ensemble optical depth. With the above linewidth convention, define

$$C = \frac{4|g|^2}{\kappa\Gamma}. \quad (85)$$

In a resonant bad-cavity limit with the required bandwidth assumptions, cavity elimination gives a desired-mode emission rate $\Gamma_c = 4|g|^2/\kappa$. The branching fraction into the external port is then

$$\beta_{\text{ex}} = \frac{\kappa_{\text{ex}}}{\kappa} \frac{\Gamma_c}{\Gamma + \Gamma_c} = \frac{\kappa_{\text{ex}}}{\kappa} \frac{C}{1 + C}. \quad (86)$$

This is a limit derived for a specified cavity, not a universal write-read efficiency. The actual η_{TI} is obtained from capture, control-pulse, retrieval, spectral, and spatial-mode matching. The incoming photon can be treated as a one-photon wavepacket in a continuum or a source-cavity dilation, and Eq. (21) then produces the corresponding instrument. Free-space collection uses its optical-mode overlap instead of cavity C . An effective optical coupling parameter must not be interpreted simultaneously as a dimensionless cooperativity and a collection probability, and an ensemble optical depth d is not an intrinsic single-ion parameter.

Dark-storage coherence and effective retrieval. During storage the control is off unless the protocol explicitly applies optical fields. For an additive effective rate model,

$$\frac{1}{T_{2,I}} = \gamma_0 + \Gamma_B + |\Delta k|^2 D_{\text{mot}} + \alpha_\phi \gamma_{\text{sc}}, \quad \lambda_{L,I} = e^{-\tau/T_{2,I}}. \quad (87)$$

Off-resonant optical scattering contributes only over intervals in which the relevant light is present. A Ramsey/echo/DD measurement can instead provide the effective coherence law [39]. A coupled-recall closure additionally assumes

$$\eta_I^{(2)}(\tau) = \eta_{\text{TI}} e^{-2\tau/T_{2,I}}. \quad (88)$$

This is a phenomenological retrieval envelope. Pure phase noise of an internal ion qubit does not alone imply that envelope; an independent optical survival measurement or microscopic derivation is needed to identify it with a physical device.

Population relaxation within the qubit is not erasure. When relaxation takes $|1\rangle \rightarrow |0\rangle$ inside the logical space, its Kraus operators are

$$A_0 = \text{diag}(1, \sqrt{u}), \quad A_1 = \sqrt{1-u}|0\rangle\langle 1|, \quad u = e^{-\tau/T_1}. \quad (89)$$

Compose these with pure-dephasing operators $D_0 = \sqrt{(1+\nu)/2}\mathbb{I}$ and $D_1 = \sqrt{(1-\nu)/2}Z$, $\nu = e^{-\tau/T_\phi}$. Then

$$\rho_{01}(\tau) = \sqrt{u\nu} \rho_{01}(0), \quad \frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}. \quad (90)$$

An independent optical erasure efficiency η gives successful operators $\sqrt{\eta} V D_j A_k$ and the same two erasure operators as Eq. (35). Its conditional entanglement fidelity is

$$F_{e|s} = \frac{1 + u + 2\sqrt{u\nu}}{4}. \quad (91)$$

This additional channel is distinct from the pure-dephasing-plus-erasure closure. Replacing internal relaxation by a flag would incorrectly make an unheralded error heralded.

Table 5: Memory-platform parameters and reduced-model alternatives. Efficiencies refer to the full declared write–store–read primitive. All rows use Eq. (35) only under balanced-rail and conditional pure-dephasing assumptions. Splitting total efficiency into write and conditional recall requires an explicit factorization.

Platform	Optical interface and native controls	Conditional coherence	Effective retrieval efficiency
REID: spin-wave CRIB + DD	d , prepared spectrum, transfer pulses, T_{mem} , m_R , DD sequence	$\lambda_L = e^{-(\tau/T_{2,R}^{\text{DD}})^{n_R}}$	Independent loss: $(1 - e^{-d})^2 \eta_{\text{DD}}$; coupled recall: multiply by $e^{-2\chi_R}$.
Cold atoms: EIT/Raman	d , $\Omega(t)$, detuning, bandwidth, M_{int} , D , Δk , σ_δ , γ_0	$e^{-a\tau - b\tau^2}$ or fitted $e^{-(\tau/T_{2,C})^{n_C}}$	Coupled recall: $\eta_{\text{EIT}}(d)\eta_{\text{mode}}(\tau)e^{-2\chi_C}$; distinguish parametric η_{EIT} from kernel-derived efficiency.
Trapped ion: Raman	g , Ω , Δ , cavity linewidth/collection, C , internal-state noise	$e^{-\tau/T_{2,I}}$	Coupled recall: $\eta_{\text{TI}}e^{-2\tau/T_{2,I}}$.

4.2.4 From memory instruments to DES parameters

The reduced model identifier must specify the optical-depth convention, logical encoding, recall variant, noise observable defining T_2 , and the location of every efficiency. A write event samples p_w ; recall samples $p_{r|w}(\tau)$, with $p_w p_{r|w}(\tau) = \eta(\tau)$ for the chosen primitive. It is not valid to sample the total η both at write and at recall.

For a token already aged to a and a monotone fitted $\lambda_L(a) > 0$, an event at age $a + \Delta$ applies $\lambda_L(a + \Delta)/\lambda_L(a)$. The original write time must survive swaps at other nodes. Non-Markovian bath correlations between tokens or interventions require an augmented noise history; one-time Kraus maps alone do not determine multi-time process statistics. The F/W metadata can reproduce its chosen isotropic approximation, whereas exact composition of the above dephasing channels requires at least Bell-diagonal probabilities for Pauli-diagonal pairs. The richer state record defines the instrument-valued reference model and the state representation required by the native validation contract.

4.3 Bell-state measurements: optical modes, detector instruments, and swapping

4.3.1 Model scope

The model proceeds from polarization-encoded elementary pairs through photon retrieval, beam-splitter interference, detector response, accepted click patterns, and the remote conditional state. The accepted click-pattern instrument determines both herald probability and conditional output state. Bell-diagonal closure gives an efficient exact reduction within the independent Pauli-noise domain.

4.3.2 The two incoming photons and the interferometer

Label the remote endpoints A, B and the two local repeater qubits R_1, R_2 . Before swapping the state is $\rho_{AR_1} \otimes \rho_{R_2B}$ when the elementary resources are independent. An optical realization first retrieves R_1, R_2 into input ports a, b . With normalized auxiliary wavepackets f, g , a balanced beam splitter acts, up to known phases, as

$$\hat{a}_p^\dagger[f] \mapsto (\hat{c}_p^\dagger[f] + \hat{d}_p^\dagger[f])/\sqrt{2}, \quad (92)$$

$$\hat{b}_p^\dagger[g] \mapsto (\hat{c}_p^\dagger[g] - \hat{d}_p^\dagger[g])/\sqrt{2}, \quad p \in \{H, V\}. \quad (93)$$

For identical auxiliary modes, the antisymmetric polarization Bell state produces the cross-port, orthogonal-polarization

signature. A suitable polarization-resolved analyzer also distinguishes one further Bell state. The standard ancilla-free passive analyzer cannot unambiguously resolve all four equiprobable Bell states; the usual two-of-four device has success probability 1/2 for independent Bell-pair inputs [15, 41]. Detector architecture and the accepted click patterns must be specified to justify accepting both states; accepting only one gives 1/4 in this ideal input case.

HOM overlap is an auxiliary-mode quantity. For two independent single photons with the same tested polarization and auxiliary density operators σ_1, σ_2 , beam-splitter expansion gives

$$P_{\text{cc}} = \frac{1}{2}(1 - V_{\text{HOM}}), \quad V_{\text{HOM}} = \text{Tr}(\sigma_1 \sigma_2). \quad (94)$$

For pure auxiliary states this is $|\langle f|g \rangle|^2$. The density operators in this expression describe the unresolved temporal, spectral, or spatial states involved in the interference test. Inserting the logical polarization marginals of perfect Bell pairs would instead give $\text{Tr}[(\mathbb{I}_2/2)^2] = 1/2$ and misidentify perfect indistinguishability as 50%. The auxiliary-mode and logical-polarization spaces must therefore be distinguished.

When spatial, spectral/temporal, and residual alignment variables are separable, their overlaps can factorize as $V_s V_t V_p$. Spectral–spatial correlations or frequency-dependent polarization transformations require the joint density operators. A mode overlap cannot be inferred from received intensity alone.

4.3.3 Temporal modes, physical delay, and timestamp jitter

For normalized spectral envelopes the field overlap is

$$V_t = \left| \int d\omega f^*(\omega)g(\omega)e^{-i\omega\delta t} \right|^2. \quad (95)$$

As a derived example, let each temporal intensity have variance σ_p^2 , with a relative delay δt and carrier offset $\delta\omega$. Gaussian integration gives

$$V_t = \exp[-\delta t^2/(4\sigma_p^2) - \sigma_p^2\delta\omega^2]. \quad (96)$$

Physical emission-time or synchronization fluctuations randomize δt and can reduce the averaged field overlap. Detector timestamp jitter occurs after interference: it does not retroactively change the incident wavepackets. It changes the accepted

coincidence ensemble. Time-resolved postselection can couple these effects, but that coupling must be evaluated with a joint detection model rather than counted twice.

For a full relative-time coincidence window w , suppose the registered difference is Gaussian with mean $\bar{\delta}$ and variance σ_{rel}^2 . Then

$$A_t(w) = \Phi\left(\frac{w/2 - \bar{\delta}}{\sigma_{\text{rel}}}\right) - \Phi\left(\frac{-w/2 - \bar{\delta}}{\sigma_{\text{rel}}}\right), \quad (97)$$

$$\sigma_{\text{rel}}^2 = \sigma_{d,1}^2 + \sigma_{d,2}^2 + \sigma_{\text{sync}}^2 + \sigma_{\text{intrinsic}}^2, \quad (98)$$

where the final term is present only when an intrinsic arrival-difference spread is modeled. For independent absolute gates about two trigger times, the acceptance is instead the product of two single-detector probabilities. A squared single-detector acceptance corresponds to the latter interpretation. It must not silently be reused for a single relative coincidence window.

For example, two independent 50 ps detector RMS jitters, zero other spread, and $w = 200$ ps give $A_t = \text{erf}(1) \simeq 0.842701$ for a relative gate. The product of two independent absolute-gate acceptances is $\text{erf}(\sqrt{2})^2 \simeq 0.911070$. This is an algebraic timing-convention check, not a detector calibration.

4.3.4 Detector POVMs and the accepted-outcome instrument

An on/off detector with efficiency η_d and independent dark-click probability p_d in its specified gate has effects

$$\Pi_0 = (1 - p_d) \sum_{n=0}^{\infty} (1 - \eta_d)^n |n\rangle\langle n|, \quad \Pi_1 = \mathbb{I} - \Pi_0. \quad (99)$$

The probability $p_d = 1 - e^{-D_d w}$ follows from a Poisson dark process of rate D_d when the detector gate is w . A coincidence window and the physical detector-active gate need not be identical. Photon-number-resolving detectors, dead time, afterpulsing, and saturation require their own response models.

For a resolved click pattern c , tensor products of detector effects and the time-acceptance rule define Π_c . Let U_{opt} include beam splitting, polarization analysis, optical loss dilations, and all auxiliary modes, and let ϱ include the two input resources, retrieved fields, and environment states. The unnormalized remote operation is

$$\mathcal{J}_c(\rho) = \text{Tr}_{R,\text{opt},E} \left[(\mathbb{I}_{AB} \otimes \sqrt{\Pi_c}) U_{\text{opt}} \varrho U_{\text{opt}}^\dagger (\mathbb{I}_{AB} \otimes \sqrt{\Pi_c}) \right]. \quad (100)$$

A bath/detector basis expansion supplies Kraus operators, exactly as in Eq. (21). Its trace gives the accepted-pattern probability. The sum over accepted patterns, with the appropriate remote Pauli corrections, gives the heralded output. This formulation includes false heralds from dark counts, multi-pair emission, or one missing photon without confusing them with successful two-photon Bell projections. Empty endpoint sectors must remain flags unless a justified model replaces them with an effective qubit state.

In the restricted no-background, single-photon, indistinguishable-auxiliary-mode, balanced-loss case with two accepted Bell states and independent Bell-diagonal input resources, the genuine-herald probability reduces to

$$p_{\text{true,BSM}} = \frac{1}{2} \eta_{r,1} \eta_{r,2} \eta_{o,1} \eta_{o,2} \eta_{d,1} \eta_{d,2} A_t. \quad (101)$$

Here $\eta_{r,j}$ are conditional memory-recall probabilities, and $\eta_{o,j}$ are optical efficiencies from the recall boundary to the detectors. Efficiencies already sampled by the caller are set to unity inside this primitive. HOM visibility affects the accepted-state operation and, for a particular pattern rule, may also affect its probability. It is not generally interchangeable with p_{BSM} , nor does $V_{\text{HOM}} = 1$ remove the $1/2$ analyzer factor. A phenomenological probability $p_{\text{BSM}} \simeq V_{\text{HOM}}$ does not represent this two-of-four optical analyzer.

4.3.5 Exact swapping and the Bell-diagonal reduction

For ideal resolved Bell outcome r , define $M_r = \langle B_r |_{R_1 R_2}$. The unnormalized endpoint state is

$$\tilde{\rho}_{AB}^{(r)} = (\mathbb{I}_A \otimes M_r \otimes \mathbb{I}_B)(\rho_{AR_1} \otimes \rho_{R_2 B})(\mathbb{I}_A \otimes M_r^\dagger \otimes \mathbb{I}_B). \quad (102)$$

Its trace is the outcome probability, and the known Pauli correction P_r acts on an endpoint before comparing to the target state. Optical loss and accepted-pattern weights multiply or generalize this trace-decreasing operation.

Let $G = \mathbb{Z}_2^2$ label Bell states $|B_{xz}\rangle = (\mathbb{I} \otimes X^x Z^z)|\Phi^+\rangle$, and let p_g, q_g be independent input Bell probabilities. Pauli labels add by bitwise XOR after the outcome correction. Therefore

$$r_g = \sum_{h \in G} p_h q_{g \oplus h}. \quad (103)$$

This equation follows either by inserting the Bell expansion into Eq. (102), or by moving each input Pauli through the Bell projection; unobservable Pauli phases cancel in the density operator. A local Pauli-noise channel updates p by the same convolution. The four probabilities are consequently a closed representation for independent Pauli noise and ideal corrected swaps. Correlated resources and non-Pauli operations are outside that closure.

For the Werner/isotropic parametrization,

$$p = (F_1, (1 - F_1)/3, (1 - F_1)/3, (1 - F_1)/3), \quad (104)$$

and similarly for q , Eq. (103) gives

$$F_{\text{ideal}} = F_1 F_2 + \frac{(1 - F_1)(1 - F_2)}{3}, \quad W_{\text{out}} = W_1 W_2, \quad (105)$$

with $W_j = (4F_j - 1)/3$. The second term represents matching nonidentity Bell errors that cancel under swapping.

Scalar comparator and failure of isotropic closure. For the diagnostic comparison in Section 4.1.1, define

$$F_{\text{sc}} = \frac{1}{2} [1 + V(2F_1 F_2 - 1)]. \quad (106)$$

This expression is excluded from the physical reference backend. At $V = 1$, two inputs with $F_1 = F_2 = 1/2$ give $F_{\text{sc}} = 1/4$, whereas Eq. (105) gives $1/3$. Two completely mixed inputs give $F_{\text{sc}} = 1/16$ instead of $1/4$. The inconsistency concerns the stated isotropic input family and local swapping operation; a fidelity below $1/4$ is possible for other two-qubit states. Independent phase damping is represented by its Bell-error kernel rather than by forcing the state into a Werner family.

4.3.6 Partial distinguishability and false-herald mixtures

A mixture of indistinguishable and distinguishable optical sectors can be written at the unnormalized-instrument level as $\mathcal{J} = V\mathcal{J}_{\text{ind}} + (1 - V)\mathcal{J}_{\text{dis}}$ when that decomposition is physically justified. Its conditional fidelity is

$$F_s = \frac{Vp_{\text{ind}}F_{\text{ind}} + (1 - V)p_{\text{dis}}F_{\text{dis}}}{Vp_{\text{ind}} + (1 - V)p_{\text{dis}}}. \quad (107)$$

A convex interpolation with weights $V, 1 - V$ is valid only when the branch success probabilities coincide or when V already denotes the conditional weight. A universal distinguishable-branch fidelity of $1/2$ is not determined by mode overlap alone.

Similarly, for true and accidental accepted-event rates,

$$F_{\text{obs}} = \frac{C_{\text{true}}F_{\text{true}} + C_{\text{acc}}F_{\text{acc}}}{C_{\text{true}} + C_{\text{acc}}}. \quad (108)$$

For an explicitly assumed completely mixed false-qubit state, $F_{\text{acc}} = 1/4$. Different false-event mechanisms can yield other states. The scalar expression $C_{\text{true}}F_{\text{swap}}/(C_{\text{true}} + C_{\text{acc}})$ corresponds instead to assigning zero target overlap to the false branch. That choice is not white qubit noise, and the assumed false-event state must accompany any fidelity comparison.

4.3.7 Event semantics at the repeater

A swap reserves two distinct local qubits, evaluates their accumulated storage evolution, attempts recall, and applies the accepted detector instrument. Both operands are consumed by the destructive attempt, including an inconclusive outcome. A successful resource joins the two surviving remote endpoints and carries the outcome/frame information; it does not create a new material qubit at the repeater. Remote qubits retain their original memory-write times. Physical completion, arrival of the herald, and availability of a known Pauli frame are separate DES events. The SDN controller receives the reported outcome only after the modeled classical delay.

4.4 FSO, adaptive optics, fiber, and orbital geometry

The analytical turbulence descriptions use the propagation conventions of Andrews and Phillips [42]; the Gamma–Gamma irradiance model additionally follows Al-Habash, Andrews, and Phillips [44]. Receiver long-exposure phase statistics are treated using Fried’s formulation [43].

4.4.1 From propagation of field modes to a quantum channel

An end-to-end optical trial efficiency separates source preparation, geometric collection, atmospheric propagation, pointing, coupling, and detection. These categories define accounting boundaries, not necessarily independent random variables. A propagation channel is distinguished from the complete trial in which it is used. The field-mode construction below connects a mode-resolved backend to analytical, event-level approximations [4, 9, 45].

For a scalar paraxial mode in a medium with a small refractive-index perturbation δn , a representative propagation equation is

$$\partial_z u = \frac{i}{2k} \nabla_{\perp}^2 u + ik\delta n(\mathbf{r}_{\perp}, z, t)u - \frac{\alpha_{\text{ext}}(z, \lambda_0)}{2}u, \quad k = 2\pi/\lambda_0. \quad (109)$$

The time argument denotes slow environmental evolution relative to optical flight; dynamic propagation requires its retarded value when that approximation fails. A split-step backend alternates Fresnel propagation with random phase screens and extinction. The screen spectrum, outer/inner scales, sampling resolution, aperture truncation, and temporal correlation determine its validity. A label such as “Rytov variance” alone does not define that backend.

For environmental realization ξ , let $\mathcal{U}_{\xi}$ propagate the input field to the receiver, P_A restrict it to the aperture, and $C_{\text{AO}, \xi}$ apply the receiver correction. In accepted input/output mode bases,

$$T_{mn}^{(\xi)} = \langle g_m | C_{\text{AO}, \xi} P_A \mathcal{U}_{\xi} | f_n \rangle. \quad (110)$$

Passivity implies $T_{\xi}^{\dagger} T_{\xi} \leq \mathbb{I}$. A vacuum-environment dilation can be written $\hat{\mathbf{b}} = T_{\xi} \hat{\mathbf{a}} + L_{\xi} \hat{\mathbf{e}}$ with $T_{\xi} T_{\xi}^{\dagger} + L_{\xi} L_{\xi}^{\dagger} = \mathbb{I}$ on the output modes. Restricting to a single photon gives the successful map $\rho \mapsto T_{\xi} \rho T_{\xi}^{\dagger}$ and the erasure completion of Eq. (25). For fading, average the unnormalized maps before conditioning. A polarization-dependent transfer needs a joint polarization–spatial matrix; scalar turbulence propagation alone does not establish polarization decoherence.

An offline lookup can expose $p(\eta, \mathcal{E}_{\text{pol}}, \delta t \mid \text{geometry, weather, AO})$ to the DES. The joint distribution matters: conditioning on received photons favors high-transmission realizations, and averaging a conditional fidelity independently of transmission generally gives a different ensemble.

4.4.2 Gaussian spreading and aperture collection

The equivalent-Gaussian model uses the $1/e^2$ intensity radius

$$w_L = \sqrt{w_0^2 + (\theta_{\text{div}} L)^2}, \quad w_0 = D_t/2. \quad (111)$$

The relation $w_0 = D_t/2$ is a chosen launch filling convention, not a universal identity for a truncated telescope beam. For a diffraction-limited Gaussian, $\theta_{\text{div}} = \lambda_0/(\pi w_0)$; an M^2 factor or a measured divergence can replace this relation. Consequently an 810/1550 nm comparison must declare whether divergence is fixed by optics, measured separately, or held artificially equal.

For normalized intensity $I(r) = 2e^{-2r^2/w_L^2}/(\pi w_L^2)$, a centered aperture of radius $a = D_r/2$ collects

$$\eta_{\text{geo}} = \int_0^a 2\pi r I(r) dr = 1 - e^{-2a^2/w_L^2} = 1 - e^{-D_r^2/(2w_L^2)}. \quad (112)$$

This gives the centered-aperture collection law. It approaches unity for a large aperture and scales as $D_r^2/(2w_L^2)$ when the aperture is small.

4.4.3 Pointing: instantaneous displacement versus mean jitter loss

For beam-center displacement $\mathbf{r}_0$ relative to the aperture center, the exact Gaussian collection is

$$\eta_{\text{ap}}(r_0) = \int_{|r| \leq a} \frac{2e^{-2|r-r_0|^2/w_L^2}}{\pi w_L^2} d^2r = F_{\chi^2_2(4r_0^2/w_L^2)}(4a^2/w_L^2). \quad (113)$$

The last form is the noncentral chi-square cumulative distribution with two degrees of freedom. It follows because each transverse coordinate of the Gaussian intensity distribution has variance $w_L^2/4$. It is directly usable for numerical sampling or quadrature.

In the small-aperture approximation, $\eta_{\text{ap}}(r_0)/\eta_{\text{ap}}(0) \approx e^{-2r_0^2/w_L^2}$. If pointing displacement has independent coordinate variance $\sigma_r^2 = L^2\sigma_\theta^2$ and mean offset of magnitude b , Gaussian integration gives

$$\mathbb{E} \left[e^{-2r_0^2/w_L^2} \right] = \frac{\exp[-2b^2/(w_L^2 + 4\sigma_r^2)]}{1 + 4\sigma_r^2/w_L^2}. \quad (114)$$

For zero mean and a *radial* RMS displacement $r_{\text{rms}} = \sqrt{2}\sigma_r$, the denominator is $1 + 2r_{\text{rms}}^2/w_L^2$. The expression $e^{-2(\sigma_\theta L/w_L)^2}$ is instead a representative-offset penalty obtained by substituting one displacement into the instantaneous loss law. Using the displaced-aperture integral already accounts for geometric collection and pointing, so a second geometric or pointing factor would double count them.

4.4.4 Extinction, elevation, and the turbulence profile

Absorption and scattering attenuate power according to

$$T_{\text{ext}} = T_{\text{abs}}T_{\text{sca}} = \exp \left[- \int_{\text{ray}} (\alpha_{\text{abs}} + \alpha_{\text{sca}}) ds \right]. \quad (115)$$

For an approximately plane-parallel atmosphere of effective vertical thickness h_0 and *elevation* ϵ , $L_{\text{atm}} \approx h_0/\sin \epsilon$. A curved atmosphere and a minimum elevation mask are preferable near the horizon. Clouds are a separate availability/extinction process; a clear-air turbulence parameter does not model cloud blockage.

A Hufnagel–Valley profile is

$$C_n^2(h) = 0.00594(\nu/27)^2(10^{-5}h)^{10}e^{-h/1000} + 2.7 \times 10^{-16}e^{-h/1500} + A_0e^{-h/100}, \quad (116)$$

with h in meters and the conventional profile coefficients. A_0 and ν parameterize a selected atmospheric profile, not a measurement unless calibrated. In the plane-wave slant-path approximation, the Fried coherence diameter is

$$r_0 = \left[0.423k^2 \frac{1}{\sin \epsilon} \int C_n^2(h) dh \right]^{-3/5}. \quad (117)$$

The factor is cosecant when the argument is elevation. Equivalently, it is secant for the zenith angle $z = \pi/2 - \epsilon$. At zenith, Eq. (117) remains finite. An implementation must not interchange these angle conventions.

For a homogeneous plane wave in weak turbulence,

$$\sigma_R^2 = 1.23C_n^2k^{7/6}L^{11/6} \quad (118)$$

is a standard limiting parameterization. Slant paths, spherical/Gaussian waves, finite apertures, and uplink versus downlink propagation require the corresponding path weighting. The same value of σ_R^2 does not specify beam wander, wavefront residuals, or temporal correlations uniquely. In particular, a scalar intensity distribution is not a substitute for a spatial random field [9, 45].

4.4.5 Normalized fading and empirical mean collection

Weak-fluctuation surrogate. A normalized lognormal irradiance is

$$H = e^{-s_\ell^2/2 + s_\ell Z}, \quad Z \sim \mathcal{N}(0, 1), \quad (119)$$

$$\mathbb{E}H = 1, \quad \text{Var}(H) = e^{s_\ell^2} - 1. \quad (120)$$

Hence $s_\ell^2 = \ln(1 + \sigma_I^2)$ if σ_I^2 is the scintillation index. In the weak plane-wave limit $\sigma_I^2 \approx \sigma_R^2$. A separately fitted relation such as $\sigma_I^2 \approx 1.23\sigma_R^{12/5}$ is an alternative surrogate, not that weak-limit identity. Likewise, a switch near $\sigma_R^2 = 0.3$ is a model-selection rule, not a universal atmospheric phase boundary.

Gamma–Gamma surrogate. Let $X \sim \text{Gamma}(\alpha, 1/\alpha)$ and $Y \sim \text{Gamma}(\beta, 1/\beta)$ independently, and $H = XY$. Integrating $p_X(x)p_Y(h/x)/x$ gives

$$p_H(h) = \frac{2(\alpha\beta)^{(\alpha+\beta)/2}}{\Gamma(\alpha)\Gamma(\beta)} h^{(\alpha+\beta)/2-1} K_{\alpha-\beta}(2\sqrt{\alpha\beta h}), \quad h > 0. \quad (121)$$

The normalization follows from the product construction; $\mathbb{E}H = 1$ and $\text{Var}(H) = \alpha^{-1} + \beta^{-1} + (\alpha\beta)^{-1}$. In terms of $s_R = \sigma_R^2$, one analytical shape prescription is

$$\alpha^{-1} = \exp \left[\frac{0.49s_R}{(1 + 1.11s_R^{6/5})^{7/6}} \right] - 1, \quad (122)$$

$$\beta^{-1} = \exp \left[\frac{0.51s_R}{(1 + 0.69s_R^{6/5})^{5/6}} \right] - 1. \quad (123)$$

These expressions specify the shape parameters of this surrogate. Their use for a particular finite-aperture slant link still requires justification. H is relative irradiance and can exceed unity. If $\eta = \bar{\eta}H$ is clipped to $[0, 1]$, the clipping changes the mean and must be recorded. A bounded distribution of collected power or a wave-optics lookup avoids interpreting unbounded irradiance as transmissivity.

An explicit limitation of the mean-loss equation. For fixed extinction and a normalized irradiance fluctuation,

$$\mathbb{E}[T_{\text{ext}}H] = T_{\text{ext}}. \quad (124)$$

A turbulence-dependent exponential reduction of the mean therefore does not follow from this expectation. Such a

reduction can instead represent a separately calibrated *mean collected-mode loss*, for example

$$\bar{\eta}_{\text{turb}} = e^{-C_{\text{atm}} S_R^{6/5}}, \quad \eta \approx \eta_{\text{fixed}} \bar{\eta}_{\text{turb}} H, \quad (125)$$

subject to boundedness and a declared joint model. Beam broadening and finite-aperture loss can change mean collected power even though normalized scintillation has mean one. We retain Eq. (125) as an empirical backend, not as a derived consequence of Eq. (120). The single-repeater example uses $C_{\text{atm}} = 0.06$. The QUASAR mean-signal calculation does not apply an additional mean turbulence penalty; its turbulence quantities are separate diagnostics.

4.4.6 Adaptive optics and conditional mode coupling

Normalize the field before aperture loss so $\int |u|^2 d^2r = 1$, and the accepted receiver mode so $\int_A |g|^2 d^2r = 1$. Then

$$\eta_{\text{ap}} = \int_A |u|^2 d^2r, \quad (126)$$

$$\eta_{\text{cpl|ap}} = \frac{|\int_A g^*(\mathbf{r}) u(\mathbf{r}) e^{-i\phi_{\text{AO}}(\mathbf{r})} d^2r|^2}{\eta_{\text{ap}}} \leq 1. \quad (127)$$

The overall accepted fraction is $\eta_{\text{ap}} \eta_{\text{cpl|ap}}$, with fixed mode-matching optics included consistently. This expression directly connects the corrected field modes to the single-photon transfer in Eq. (110).

For a matched pupil with sufficiently small residual phase, a useful approximation is $\eta_{\text{cpl}} \simeq \eta_{\text{ideal}} e^{-\sigma_{\phi, \text{res}}^2}$. Fitting error, servo lag, wavefront-sensor noise, anisoplanatism, and uncorrected modes enter $\sigma_{\phi, \text{res}}^2$ according to the specified AO design. This is a regime-limited reduction, not a complete AO model. A relative gain $G_{\text{AO}} = \eta_{\text{AO on}} / \eta_{\text{AO off}}$ may exceed one; a physical coupling probability may not. Factors η_{AO} and G_{AO} must be distinguished on that basis. Direct bucket detection and single-mode fiber injection do not have identical wavefront sensitivity [9].

The empirical factors $\eta_{\text{HV}} = e^{-(D_r/r_0)^{5/3}}$, $G_{\text{ap}} = 1 - e^{-(D_r/D_{\text{sat}})^2}$, and $G_{\text{bw}} = 1/(1 + 2\sigma_{\text{bw}}^2/w_e^2)$ can be retained as named calibrated surrogates. They are not automatically independent of the aperture, pointing, and turbulence terms already evaluated. An independent product of mean efficiencies also differs from the mean of their correlated product.

4.4.7 Fiber propagation and wavelength conversion

For fiber length L_f in kilometers, power-attenuation coefficient α_f in dB/km, and insertion loss ℓ_{ins} in dB,

$$\eta_f = 10^{-(\alpha_f L_f + \ell_{\text{ins}})/10}. \quad (128)$$

A mode-resolved narrowband propagation matrix can be written

$$T_f(\omega) = \sqrt{\eta_f(\omega)} J_f(\omega) e^{i\varphi_f(\omega)}, \quad (129)$$

$$\varphi_f(\omega) = L_f [\beta_0 + \beta_1(\omega - \omega_0) + \frac{1}{2}\beta_2(\omega - \omega_0)^2 + \dots], \quad (130)$$

using consistent length units in the phase. J_f represents birefringence or polarization-dependent transfer. Group delay is $n_g L_f / c$; dispersion changes the outgoing temporal mode and therefore BSM overlap. Fiber attenuation alone is an erasure process, not a depolarizing map. Random birefringence, polarization-mode dispersion, and detector-conditioned background can reduce conditional fidelity through distinct mechanisms.

A downlink at 810 nm and a telecom fiber or a Pr/alkali/ion memory are not automatically spectrally compatible. An explicit frequency-conversion interface must specify conversion efficiency, temporal/spectral transfer, added noise, and polarization preservation. Merely assigning the same token to those devices bypasses a physical interface that can dominate the experiment.

4.4.8 Orbital state, visibility, and time of flight

For a Keplerian benchmark with semimajor axis a_o , eccentricity e_o , and mean motion $n_o = \sqrt{\mu/a_o^3}$, solve

$$M(t) = M_0 + n_o(t - t_0) = E_o - e_o \sin E_o. \quad (131)$$

The inertial position follows from rotations of

$$\mathbf{r}_{\text{orb}} = \begin{pmatrix} a_o(\cos E_o - e_o) \\ a_o\sqrt{1 - e_o^2} \sin E_o \\ 0 \end{pmatrix}, \quad (132)$$

$$\mathbf{r}_I = R_3(\Omega_o) R_1(i_o) R_3(\omega_o) \mathbf{r}_{\text{orb}}. \quad (133)$$

This ideal propagation is not a replacement for TLE/OMM propagation with its appropriate SGP4 conventions. Epoch, reference frame, time scale, Earth rotation, and the propagator must be recorded [46]. The single-repeater local tangent-plane trajectory is a separate controlled geometry used to isolate offset and crossing-angle effects; its labels must not be mistaken for a mission ephemeris.

For ground position $\mathbf{r}_g$ and local up unit vector $\mathbf{u}_g$ in the same Earth-fixed frame as the satellite,

$$L_g = \|\mathbf{r}_s - \mathbf{r}_g\|, \quad \epsilon_g = \arcsin [\mathbf{u}_g \cdot (\mathbf{r}_s - \mathbf{r}_g) / L_g]. \quad (134)$$

A link is usable only if the elevation mask, tracking constraints, optical-terminal allocation, weather model, and hardware readiness all permit it. For spherical Earth radius R_E and altitude h , the geometric slant-range check is

$$L_g(\epsilon) = \sqrt{(R_E + h)^2 - R_E^2 \cos^2 \epsilon} - R_E \sin \epsilon. \quad (135)$$

It gives $L_g(\pi/2) = h$ and $L_g(0) = \sqrt{(R_E + h)^2 - R_E^2}$. An inter-satellite line segment must remain outside the selected Earth/atmosphere exclusion radius. Find the closest point $\mathbf{r}_1 + u(\mathbf{r}_2 - \mathbf{r}_1)$ with u clamped to $[0, 1]$ and test its norm, rather than applying a ground-elevation mask to an inter-satellite link.

Photon arrival time solves $t_a - t_e = \|\mathbf{r}_r(t_a) - \mathbf{r}_t(t_e)\|/c$ in vacuum, with optical-path corrections when required. The frozen-geometry approximation uses $L(t_e)/c$ and should be checked at the desired resolution. First-order Doppler is $\delta\omega/\omega \simeq -\dot{L}/c$; spectral compensation and synchronization affect memory acceptance and Eq. (96). The DES tracks propagation, storage, heralding, and correction delays separately.

4.4.9 From source and count processes to state quality

Source repetition f_{rep} , mean pairs per excitation μ , pair flux R_p , and the simulator's per-mode attempt rate have different units and meanings. For an ideal single two-mode-squeezed pair mode, $P_n = \mu^n / (1 + \mu)^{n+1}$ and $R_p = f_{\text{rep}}\mu$. A multimode polarization-entangled SPDC source needs its appropriate photon-number and joint spectral model. A “1000 MHz emitted” label must therefore be linked to either a pulse clock or an emitted pair flux before inferring multi-pair noise.

For a direct two-downlink coincidence benchmark in the low-occupancy regime,

$$C_{\text{true}}(t) = R_p(t)\eta_A(t)\eta_B(t)A_t(w), \quad (136)$$

$$C_{\text{acc}}(t) \simeq S_A(t)S_B(t)w, \quad (137)$$

where S_j are the relevant independent singles processes, including background and dark counts, and the accidental expression assumes a full-width relative gate with negligible saturation. Genuine paired counts must not be counted again as independent accidentals. At high brightness or appreciable $S_j w$, an explicit point-process/detector model replaces this approximation.

A dimensionally defined signal-to-noise ratio is $C_{\text{true}}/C_{\text{acc}}$, with both in s^{-1} or both counted in the same window. If genuine pairs are ideal and accidental qubit states are $\mathbb{I}_4/4$, then

$$W = \frac{C_{\text{true}}}{C_{\text{true}} + C_{\text{acc}}}, \quad F = \frac{1 + 3W}{4}. \quad (138)$$

This derives the SNR-to-Werner relation only for that mixture model. It does not say that pure propagation loss itself depolarizes a surviving photon. Source, background, and dark-count parameters must use an explicit counts-or-rates convention before evaluation. At a delivered-service boundary, memory and swapping are further operations and cannot be inferred from a direct coincidence rate.

4.5 Space–time-based distribution and segmented nesting

STBD and segmented nesting are adopted from the space–time distribution framework of Koudia et al. [29]. This subsection formalizes their use with explicit two-ended resources, age-dependent noise, and execution constraints.

4.5.1 Planning and physical execution

STBD uses a space–time graph, product utilities, destination-time search, binary nesting, and early nesting of snapshot-consecutive segments. Causal edge durations, endpoint memory accounting, final-state quality, and failure handling connect that routing structure to an executable entanglement plan. Optimality of the selected additive route cost is distinct from global optimality for a stochastic, multi-request network.

Let the management horizon be divided by $t_0 < t_1 < \dots < t_K$. A snapshot $G_k = (V, E_k)$ records candidate physical contacts. A lifted vertex (v, k) denotes node v at time t_k , not a qubit that has been copied into every layer. Two different edge types are required. A communication edge $(u, k) \rightarrow (v, k')$ represents a physical elementary-link

opportunity whose generation and required heralding can complete no earlier than $t_{k'}$. A waiting edge $(v, k) \rightarrow (v, k+1)$ represents keeping a physical excitation in an allocated memory slot. Waiting edges require nonzero memory capacity and an appropriate survival/quality model.

The snapshot-to-space–time construction is described in Ref. [29]. Communication edges across layers specify link opportunities, not complete storage reservations or physical swaps. A communication edge creates a *two-ended* entanglement resource; it does not forward an unknown qubit like a classical packet. Memory waiting, endpoint inventories, and operation prerequisites complete the execution model.

4.5.2 Weights, horizons, and the shortest-path transformation

For each eligible edge, define

$$U_e = (U_e^{\text{fid}})^{\omega_f} (U_e^{\text{mem}})^{\omega_m}, \quad (139)$$

$$U_e^{\text{mem}} = e^{-\Delta t_e/t_c}, \quad \omega_f, \omega_m \geq 0, \quad \omega_f + \omega_m = 1, \quad (140)$$

$$U(P) = \prod_{e \in P} U_e. \quad (141)$$

The function $U_e^{\text{fid}} \in (0, 1]$ is part of the policy configuration, not a unique consequence of the physical backend. Choosing a monotone function of predicted fidelity assigns a preference among routes.

For positive utilities,

$$\arg \max_P U(P) = \arg \min_P \sum_{e \in P} c_e, \quad c_e = -\log U_e \geq 0. \quad (142)$$

Zero-utility edges are removed. This proves equivalence between multiplicative-utility routing and an additive shortest path for the chosen graph and weights. It does not include resource contention, correlated generation failures, or output-state constraints unless they are encoded in the state and feasible edge set. All target time layers can connect to a terminal sink; a single shortest-path search then replaces an explicit loop over destination times. If edges advance time, the graph is acyclic and a topological dynamic program suffices. Otherwise nonnegative-weight shortest-path methods are appropriate.

The planning horizon may be tied to a coherence scale t_c or fixed independently. Automatically setting it to the swept physical T_2 changes both memory physics and planning foresight. A memory-sensitivity study therefore fixes the horizon, while a separate horizon sweep measures the routing-policy effect. A stretched-exponential T_2 is an operational coherence parameter, not a sharp deterministic failure boundary.

For independent Poisson opportunities of intensity $\lambda_e(t)$ with no resource blocking, an optional planning predictor is

$$p_{e, \text{window}} = 1 - \exp \left[- \int_{t_a}^{t_b} \lambda_e(t) dt \right]. \quad (143)$$

This is the probability of at least one generation opportunity over that interval, conditional on the declared link model.

4.5.3 The physical plan and resource constraints

Represent an elementary or swapped pair by

$$P = (\text{id}, u, m_u, t_{w,u}, v, m_v, t_{w,v}, \mathbf{p}, \text{frame}, \text{request}), \quad (144)$$

where $\mathbf{p}$ is a Bell-probability vector in the refined Pauli-diagonal model, and the two original write times are retained. A scalar F/W record remains a compact approximation-specific adapter. A generation action declares the source, propagation direction, retained local half, remote loading operation, and successful-herald condition.

For an executable schedule π , physical memory occupancy obeys

$$\sum_{q: q \text{ stored at } v \text{ during } t} 1 \leq M_v(t). \quad (145)$$

Optical terminals, source flux, detector dead time, and BSM processors impose analogous constraints. A memory bank created dynamically for every link label must not create unlimited physical capacity: per-label reservations must sum to the hardware budget. If a source of total attempt rate r_v serves several modes/links, their assigned rates sum to at most r_v , unless independent additional sources are explicitly modeled.

A swap at node r can consume P_{ur} and P_{rv} only if their local qubits are distinct, present, unconsumed, not reserved by another operation, and causally known to the executor. Their remote endpoints become u, v on success. The pair identity changes, but $t_{w,u}, t_{w,v}$ do not. A failure consumes both local inputs and invalidates their pair correlations; cancellation/heralds reach the remote records according to the modeled classical delay.

4.5.4 Exact age bookkeeping and the end-to-end quality constraint

For a fitted coherence $\lambda(a) = e^{-\chi(a)}$ and an already applied age a , a later event at b must apply

$$\lambda_{a \rightarrow b} = \frac{\lambda(b)}{\lambda(a)} = e^{-[\chi(b) - \chi(a)]}, \quad b \geq a. \quad (146)$$

This telescopes over arbitrary event subdivisions. Replacing it by $e^{-\chi(b-a)}$ is generally wrong for $\chi(a) = (a/T_2)^n$ with $n \neq 1$. Consequently a memory-wait edge has an age-dependent quality cost unless exponential decay is assumed. Exact optimization may require age-augmented states; a simpler planner must verify actual ages during execution. This bookkeeping describes the chosen monotone effective channel, not general multi-time non-Markovian noise under arbitrary interventions.

Under an isotropic independent-noise approximation, define $\mathcal{H}(\pi)$ as the set of all actual local-qubit storage histories, from write to local consumption or final delivery. Then

$$W_{\text{end}}(\pi) = \prod_{e \in P} W_e \prod_{q \in \mathcal{H}(\pi)} \lambda_q^{\text{iso}}(\tau_q) \prod_{s \in \pi} v_s, \quad (147)$$

where v_s is included only if the BSM error is explicitly modeled as an independent isotropic channel. Every elementary link factor and every storage history appears once. Transition

waiting is included through the actual surviving endpoint histories, not through an unspecified constant.

For pure-dephasing memory models, use the Bell convolution in Eq. (103), with each local phase-error distribution $[(1 + \lambda_L)/2, (1 - \lambda_L)/2, 0, 0]$ in the declared Bell ordering, instead of multiplying a Werner parameter by λ_L . Local amplitude damping or state-dependent loss requires a more general representation. The physical channel and the metadata update must agree.

When Eq. (147) is valid and $F_{\min} > 1/4$, the final-quality condition is

$$-\ln W_{\text{end}} \leq -\ln[(4F_{\min} - 1)/3]. \quad (148)$$

Checking $F_e \geq F_{\min}$ for every elementary edge is insufficient. Four independent Werner links with $F_e = 0.9$ have, even with perfect memories and swaps,

$$F_{\text{end}} = \frac{1 + 3[(4 \times 0.9 - 1)/3]^4}{4} \approx 0.673126. \quad (149)$$

This explicit counterexample explains the need for final-state gating beyond per-link threshold filtering.

4.5.5 Binary and segmented nesting

A path of n nodes has $\ell = n - 1$ elementary links and $\ell - 1$ repeater swaps. In a balanced binary schedule, adjacent available *segments* are joined at their shared repeater, with an unmatched segment carried to the next level. The logical depth is

$$L_{\text{nest}} = \lceil \log_2 \ell \rceil \quad (\ell \geq 1), \quad (150)$$

so a direct elementary link needs zero rounds and a two-link path needs one. A BSM acts on two local qubits, not on two different repeater nodes. Parallel operations must use disjoint local operands and obey processor constraints.

For segmented nesting, split the selected route into snapshot-consecutive subpaths $P^{(1)}, \dots, P^{(K')}$. Establish elementary links in each segment, join ready local subsegments, store the remaining segment endpoints across the required transitions, and stitch neighboring segment resources when the prerequisites are available. The DES, not the coarse snapshot index, sets every generation, recall, BSM, and herald timestamp.

Proposition 1 (Restricted memory-quality comparison)

Suppose a segmented and a late-nesting schedule use the same successful elementary pairs and the same accepted ideal swaps. Suppose their final delivery time and endpoint memories are fixed, and segmentation weakly shortens every internal-qubit storage history without introducing any extra waiting, noise, or resource conflict. Under independent monotone isotropic storage factors, $F_{\text{seg}} \geq F_{\text{late}}$. The inequality is strict only if a changed history has strictly decaying coherence and the remaining visibility product is nonzero.

Proof. Each unchanged factor in Eq. (147) cancels in the ratio. Every shortened history increases or preserves its nonnegative coherence factor. Hence $W_{\text{seg}} \geq W_{\text{late}}$ and $F = (1 + 3W)/4$ preserves the ordering. The stated conditions determine strictness. $\square$

This proposition identifies the conditions under which early nesting preserves or improves memory-induced quality.

Algorithm 1 STBD with resource-constrained physical execution

Require: Endpoints (A, B) , $F_{\min}$, deadline D , horizon H , physical budgets, model identifiers

```
1: Accept request; create request/plan epoch and deadline events
2: Build time-expanded contacts and feasible waiting edges from available or forecast information
3: Rank candidate paths by Eq. (142); declare  $U^{\text{fid}}$ 
4: for candidate path in the configured search budget do
5:   Build segmented swap dependencies and endpoint reservations
6:   if predicted final quality, timing, and hardware budgets are feasible then
7:     Issue generation schedule; enter event-driven execution
8:   break
9: end if
10: end for
11: while request active and simulation time  $< D$  do
12:   Process arrived storage/swap reports, not future outcomes
13:   Revalidate token identities, original ages, and reservations
14:   if two prerequisites for a permitted local swap are ready then
15:     Schedule recall, detector instrument, and result herald
16:   end if
17:   if a pair reaches  $(A, B)$  with an available correction frame then
18:     Evaluate final-state quality at the delivery boundary
19:     if  $F \geq F_{\min}$  then
20:       Deliver and release unused reservations
21:     end if
22:   end if
23:   if failure, expiry, or invalidated contact prevents completion then
24:     Apply declared local retry, cycle reset, or replanning policy
25:   end if
26: end while
27: Abort unresolved request; notify endpoints and clear its resources
```

4.5.6 Executable request and retry procedure

Algorithm 1 combines path selection, segmented nesting, feedback, and retry with explicit resource validation. Candidate enumeration, reservation method, replanning trigger, and retry locality are policy parameters.

With K layers and $|V|$ physical nodes, the basic lifted graph has $O(K|V|)$ vertices and $O(\sum_k |E_k| + K|V|)$ edges. A nonnegative-weight heap-based shortest path costs $O((|V| + |\mathcal{E}|) \log |V|)$; an acyclic graph can be processed in $O(|V| + |\mathcal{E}|)$. Age augmentation, multi-commodity reservations, and candidate backtracking add cost. Planning time and event-execution time should be reported separately.

4.5.7 Throughput and control-plane KPIs

A product of source, memory, swap, and transmission probabilities multiplied by a mode count is a per-trial yield scaling, not a rate. A trial-frequency or elapsed-time factor is required for units of pairs per second. The delivered-rate estimator is

$$R_{\text{usable}} = \frac{1}{T_{\text{obs}}} \sum_j \mathbf{1}\{\text{delivered}_j, F_j \geq F_{\min}, t_j \leq D_j\}. \quad (151)$$

In a stationary two-link reset toy model with independent exponential readiness times $X_i \sim \text{Exp}(r_i)$, ideal unlimited waiting, and each cycle consuming both links,

$$\mathbb{E} \max(X_1, X_2) = \frac{1}{r_1} + \frac{1}{r_2} - \frac{1}{r_1 + r_2}. \quad (152)$$

Adding a fixed swap/notification overhead t_o and independent accepted success probability p_b gives $R = p_b / [\mathbb{E} \max(X_1, X_2) + t_o]$. This is a useful limiting-case check, not an approximation for a fading, finite-buffer constellation without further analysis.

Memory-mode time per usable pair is $\int_0^{T_{\text{obs}}} \sum_v B_v(t) dt / N_{\text{usable}}$; retries and drops are separated by cause. Controller telemetry carries request/plan epoch, pair ID, endpoints, mode IDs, original write times, quality/model metadata, and event timestamp. Classical messages are scheduled with propagation, serialization, processing, queuing, and optional polling delay. Ryu HTTP wall-clock time is a runtime measurement, not automatically simulated network latency. Controller comparisons must hold information availability, policy, and physical random streams fixed when testing only the transport boundary.

4.6 Scenario configurations and counting conventions

4.6.1 Single-repeater configurations

The controlled S1 sweep and dated Paris–Vienna run have distinct geometric and capacity settings. Section 2.4 defines the two co-located limits and the two 500 km separated-terminal cases; Table 6 records their numerical defaults. The satellite stores one local half of each elementary pair and performs the BSM, while the remote halves belong to the two OGSs. Hardware implementation requires its own loading and spectral-interface specification.

4.6.2 Run-bin statistics and model representation

The complete controlled campaign contains 32 settings, four repeats per setting, 24 320 run-bin rows, and 6080 ensemble rows. Every rate is a counted number of successes divided by its 2 s bin width, followed by aggregation over repeats. Fidelity in the plotted sweep is the median of supported repeat-bin mean fidelities. Empty-support bins are undefined. The integrated count for each run is the sum of its bin counts. Means and sample standard deviations are calculated across the four full runs.

The supplied compact-state trajectories retain the state-update conventions of their numerical configuration. They

Table 6: Recorded numerical settings for the controlled S1 sweep and dated Paris–Vienna run. The sweep and geographical experiment use the same recorded optical and memory defaults, with the listed capacity and elevation-mask differences. All efficiencies are dimensionless; angular jitter and the static pointing factor are separate configuration fields.

Parameter	Four-archetype sweep	Dated Paris–Vienna case
Execution / source rate	Local event execution; $5 \times 10^5 \text{ s}^{-1}$	Local event execution; $5 \times 10^5 \text{ s}^{-1}$
Memory platform / depth / exponent	Cold atom; $d = 20$; $n = 1.3$	Cold atom; $d = 20$; $n = 1.3$
T_2 and retrieval envelope	$T_2 = 100 \text{ s}$ default; $T_{\text{mem}} = 1000 \text{ s}$, $m = 1$	$T_2 = 100 \text{ s}$; $T_{\text{mem}} = 1000 \text{ s}$, $m = 1$
Initial mode efficiency / capacity	0.9; 8192 modes per label	0.9; 256 modes per label
Wavelength / apertures	810 nm; $D_t = 0.15 \text{ m}$, $D_r = 0.50 \text{ m}$	Same
Divergence / pointing	$12 \mu\text{rad}$; static factor 0.97; jitter $0.01 \mu\text{rad}$	Same
Source / coupling / detection factors	0.85/0.70/0.75	Same
Atmospheric fields	$t_{\text{abs}} = t_{\text{sca}} = 0.94$, $c_{\text{atm}} = 0.06$	Same
Turbulence / AO	$\sigma_R^2 = 0.08$ default; loss/quality coefficients 0.08/0.03; AO gain 1	Same
Stochastic fading / jitter	Both enabled	Both enabled
Polarization / phase / temporal mismatch	0.02 rad / 0.1 rad / 20 ps; coherence-width field 100 ps	Same
Count-model fields	$p_t = 0.5$, $n_0 = 10^{-5}$, $n_{\text{dc}} = 120 \text{ s}^{-1}$, $n_{\text{conv}} = 0$, gate 1 ns	Same
BSM configuration	$V_{\text{HOM}} = 0.97$, $\eta_d = 0.82$, jitter 50 ps, window 200 ps, $p_{\text{dc}} = 10^{-6}$	Same
Elevation mask / bin width	$20^\circ / 2 \text{ s}$	$10^\circ / 2 \text{ s}$
Replication / seed	Four repeats per setting; base seed 62345	One run; seed 103445

are distinguished from the exact Bell-vector and general instrument constructions in Sections 4.2 and 4.3. In particular, a low-fidelity scalar trajectory is not retrospectively assigned the pure-dephasing limit. The full-grid plots, original run-bin records, and parameter definitions remain together in the source-data package.

4.6.3 S2 constellation counting boundary

The S2 figure retains the four views of the same timed distribution experiment. Its 150 s snapshots, 300 s request spacing, 600 s horizon, and 30 s delivery cadence are separately specified. Counts are per event. Zero quality values between plotted windows are a no-delivery display convention. No controller-equivalence or routing-baseline gain is inferred from this single service trace.

4.6.4 S3 observation and reporting boundary

The S3 topology is specified in Fig. 6, and its replacement delivery and control results occupy the two parts of Fig. 7. Raw delivery counts endpoint pairs before fidelity qualification; qualified delivery applies $F \geq 0.8$ to the delivered resources. Raw rates use 1 s bins. The conditional endpoint fidelity is count weighted, whereas the cumulative qualified curves accumulate only accepted deliveries. A count-weighted bin mean is not substituted for the individual qualification decision.

The reported diagnostics distinguish actual storage ages at BSM readout, storage-age distributions at the OGS and remote endpoints, reserved/stored OGS slots, and BSM busy-time fraction. Reserved/stored occupancy is not interpreted as the number of ready, high-fidelity pairs. The resource/policy cases are labeled FIFO with 100 μs service, freshest-first control, 50 ms BSM service, and 0.1 s optical interpolation. The 100 μs and 50 ms values are service durations; 0.1 s is an optical interpolation setting, not a delivery-count bin or a mandatory swap interval. The panel sets contain one counting


realization and an across-realization mean but do not specify their replication count or confidence intervals; cumulative totals and other numerical summaries are therefore quoted only at plot-reading precision. The curves do not by themselves identify a microscopic memory or BSM instrument, whose selection belongs to the simulation configuration.

4.7 Dated Paris–Vienna configuration, records, and geographical context

The dated design case uses city-center proxies at Paris (48.8566°N, 2.3522°E) and Vienna (48.2082°N, 16.3738°E), with zero station altitude. The assumed midpoint is 25 September 2026 at 17:00 UTC. The circular 360 km orbit has inclination approximately 85.77° and spans a 600 s observation. The shared interval above 10° lasts approximately 312.656 s. Geometric, propagation, and counting timestamps are separately retained in the numerical archive.


The mean midpoint budget per arm includes about 26.770 dB geometrical loss, 0.539 dB atmospheric loss, 0.038 dB configured mean turbulence loss, 0.132 dB static pointing loss, 1.549 dB coupling loss, 0.706 dB source loss, and 1.249 dB detection loss. The total is about 30.984 dB per arm. The deterministic budget excludes memory recall and BSM.

For the same 300 two-second bins, the total failure column closes to 313 914 582 stage-specific increments: 313 853 542 link failures, 14 937 recall failures, 45 771 BSM failures, and 332 full-storage rejections. These categories are not divided by one common denominator to infer a composite success probability. The final retained inventory is excluded from delivered counts.



(a) Global illumination and European ground-track context at the labeled midpoint.

Paris-Vienna common-visibility pass | 2026-09-25
Time-resolved ground illumination and projected downlink connections



(b) Snapshots near the beginning, midpoint, and end of common visibility.

Figure 9: Dated S1 geographical and illumination context. (a) The colored ground track spans different times, with common visibility highlighted, while illumination is evaluated at the labeled midpoint. (b) Ground illumination is recomputed for each of the three labeled instants; colored lines are projected connections to Paris and Vienna, not physical ray paths. The stations are city-center proxies and the orbit is a circular design case rather than a TLE or reconstructed spacecraft pass. Solar positions follow the NOAA approximate geometric equations; the base map is Natural Earth public-domain land data. Second-rounded UTC stamps are the timestamp references, while offset annotations are coarsely rounded. Illumination is diagnostic and does not modify the simulated background.

4.7.1 Geographical and illumination overlays

The global and European maps and all three pass snapshots are retained in Fig. 9(a,b). Illumination is a diagnostic overlay: it does not gate the fixed-background channel in this design case. At the midpoint, modeled solar elevation is approximately $+6.45^\circ$ at Paris and -2.78° at Vienna. Terrain, atmospheric refraction, and a satellite-eclipse calculation are outside these overlays.

4.8 Direct Matera–Hermos inputs and extended receiver analysis

4.8.1 Geometry, optical assumptions, and reference point

The direct-distribution case uses one onboard entangled-pair source and two simultaneous downlinks, with one photon sent to Matera and the other to Hermos. It isolates optical collection and receiver count statistics from memory synchronization and swapping. Table 7 defines the geometry and optical assumptions. This case corresponds to the memory-free resource stage. The 810 and 1550 nm cases use equal wavelengths on their two arms, equal detector efficiency, and equal timing jitter; they compare wavelength-dependent propagation under a common receiver model.

Table 7: Direct-downlink configuration. Geographic positions use WGS84. The two-body elements have a J2000-like frame and an assumed UTC epoch. Additional optical and receiver values are stated modeling assumptions.

Parameter	Matera / common	Helmos / interpretation
Latitude, longitude	40.6493°, 16.7041°	37.9854°, 22.1983°
Station altitude; receiver diameter	536.9 m; 1.5 m	2340 m; 2.3 m
Receiver blocked-area fraction	0.1000	0.1035
Minimum elevation	20°	Both stations
Semimajor axis; eccentricity	6939.1284 km; 0	Circular two-body reference
Inclination; RAAN	97.6355°; 284.5°	Inertial element convention
Argument of periapsis; initial anomaly	0°; 54.0783°	Phase at the stated epoch
Epoch; calculation spacing	1 Jan. 2032, 22:59:28 UTC; 1 s	Assumed time/frame interpretation
Transmitter diameter; waist radius	0.25 m; 0.125 m	Each arm; equivalent untruncated Gaussian
Beam quality; pointing RMS	$M^2 = 1$; 1 μ rad per axis	Independent, zero-mean Gaussian pointing
Wavelengths; source fidelity	810 and 1550 nm; $F_0 = 0.99$	One equal-wavelength case at a time
TX efficiency; RX efficiency; PDE	0.80; 0.50; 0.60	Each applied once, both wavelengths
Zenith atmospheric transmission	$T_z = 0.90$	Slant transmission $T_z^{1/\sin \epsilon}$
Detector jitter; detectors per station	50 ps RMS each; 4	Independent jitter; evenly loaded outputs
Dark rate; spectral bandwidth	100 s^{-1} per detector; 1 nm	400 s^{-1} per station
Field-stop half-angle; sky radiance	10 μ rad; $10^{-6} W m^{-2} sr^{-1} nm^{-1}$	Same reference radiance at both sites/wavelengths
HV surface coefficient; wind	$1.7 \times 10^{-14} m^{-2/3}$; 21 m/s	Atmospheric-profile sensitivities in Section 4.8.4




Figure 10: Direct baseline: Matera–Helmos pass geometry. Elevations and slant ranges for the supplied circular-orbit interpretation. Shading identifies simultaneous visibility above the 20° mask. The per-pass optical analysis and the frozen-peak operating-point calculation use different observation boundaries, stated in the text.

The reference two-body period is 95.88 min. Maximum elevations are 62.64° at Matera and 63.21° at Helmos. Shared samples above the elevation mask run from 23:20:00 to 23:24:52 on the assumed UTC date, corresponding to an approximately 293 s exposure on the 1 s grid. This is shown in Fig. 10.

The equivalent Gaussian model includes pointing through

$$\theta_{\text{div}} = \frac{M^2 \lambda_0}{\pi w_0}, \quad (153)$$

$$w_{\text{eff}}^2 = w_0^2 + L^2(\theta_{\text{div}}^2 + 4\sigma_\theta^2), \quad (154)$$

$$\eta_{\text{ann}} = e^{-2r_i^2/w_{\text{eff}}^2} - e^{-2r_o^2/w_{\text{eff}}^2}, \quad (155)$$

where $r_o = D_r/2$ and $r_i = r_o \sqrt{o}$ for blocked-area fraction o . The per-arm detection probability before the separate field-stop sensitivity is

$$\eta_j = \eta_{\text{TX}} \eta_{\text{ann},j} T_z^{1/\sin \epsilon_j} \eta_{\text{RX}} \eta_d. \quad (156)$$

Gaussian pointing is already included in w_{eff} [4]. Table 8 separates the incremental pointing contribution diagnostically.

The frozen probabilities are $(\eta_M, \eta_H) = (0.053163022, 0.099111603)$ at 810 nm and

(0.024818782, 0.053066425) at 1550 nm. These numerical inputs retain guard digits for exact reproduction; they do not imply corresponding physical precision. The ideal linear-detector, 10^9 pairs/s stress reference gives approximately 5.269×10^6 and 1.317×10^6 peak genuine coincidences/s. Its station singles loading reaches 53.16 and 99.11 Mcps at 810 nm, motivating the fidelity-qualified and finite-live-time operating points below.

4.8.2 Background, coincidence acceptance, and a fidelity constraint

For entrance-plane sky spectral radiance $\mathcal{L}_\lambda$, clear area A_{clear} , field solid angle Ω , and bandwidth $\Delta\lambda$, the station sky rate is [12]

$$B_j = \mathcal{L}_\lambda A_{\text{clear},j} \Omega \Delta\lambda \eta_{\text{RX}} \eta_d \frac{\lambda_0}{h\nu c}. \quad (157)$$

Use wavelength in meters in the photon-energy factor and nm for the bandwidth when radiance is per nm. The reference sky rates are $(611.2, 1431.4) s^{-1}$ at 810 nm and $(1169.6, 2739.2) s^{-1}$ at 1550 nm. With the station dark rate, the total noise inputs are $(1011.2, 1831.4) s^{-1}$ and $(1569.6, 3139.2) s^{-1}$.

Table 8: Peak reference loss budget in dB. Collection and the incremental pointing contribution sum to the pointing-broadened annular collection. The total detection probability is used only once in the count equations.

Contribution	810 Matera	810 Helmos	1550 Matera	1550 Helmos
Gaussian annular collection	3.86	1.97	8.43	5.24
Incremental pointing contribution	2.17	1.35	0.91	0.80
Slant atmospheric loss	0.52	0.52	0.52	0.52
TX and RX fixed loss	3.98	3.98	3.98	3.98
Total optical loss	10.53	7.82	13.83	10.53
Detection loss	2.22	2.22	2.22	2.22
Total including detection	12.74	10.04	16.05	12.75

For CW Poisson pair brightness R , all-pairs time-tag association, and a full relative window τ , define [11]

$$g(\tau) = \text{erf}\left(\frac{\tau}{2\sqrt{2}\sigma_\Delta}\right), \quad \sigma_\Delta = \sqrt{2} \times 50 \text{ ps}, \quad (158)$$

$$C = R\eta_M\eta_H g(\tau), \quad (159)$$

$$A = (R\eta_M + N_M)(R\eta_H + N_H)\tau, \quad (160)$$

$$F_{\text{mix}} = \frac{F_0 C + A/4}{C + A}. \quad (161)$$

The accidental term includes different-pair coincidences and independent sky/dark contributions. The 1/4 overlap declares a completely mixed effective accidental-qubit state. The statistic is the fidelity of that accepted mixture model.

For $1/4 < F_\star < F_0$, write

$$\zeta = \frac{F_0 - F_\star}{F_\star - 1/4}, \quad b_j = \frac{N_j}{\eta_j}, \quad K = \frac{\zeta g(\tau)}{\tau} - b_M - b_H. \quad (162)$$

Then $F_{\text{mix}} \geq F_\star$ is equivalent to $R^2 - KR + b_M b_H \leq 0$. If $K \geq 2\sqrt{b_M b_H}$, the feasible brightness interval is

$$R_- \leq R \leq R_+, \quad R_\pm = \frac{K \pm \sqrt{K^2 - 4b_M b_H}}{2}. \quad (163)$$

Thus background sets a lower brightness limit while different-pair accidentals set an upper limit. At $R = 10^9 \text{ s}^{-1}$ and $\tau = 1 \text{ ns}$, $R\tau = 1$ and $F_{\text{mix}} \approx 0.62$; this stress point fails the selected $F_\star = 0.90$ target.

4.8.3 Detected singles, live time, and optimized operation

For four evenly loaded nonparalyzable detectors per station and dead time t_d , the independent live-fraction approximation gives

$$\ell_j = \left[1 + \frac{(R\eta_j + N_j)t_d}{4}\right]^{-1}, \quad C_{\text{live}} = C\ell_M\ell_H. \quad (164)$$

The argument is *detected singles before dead-time suppression*. Applying the same live fractions to independent accidentals preserves Eq. (161) within this approximation. Table 9 reports fidelity-qualified brightness and genuine-coincidence contributions.

Narrowing the window alone has a finite limit because

$$\frac{\tau}{g(\tau)} \rightarrow \sqrt{2\pi}\sigma_\Delta \quad (\tau \rightarrow 0). \quad (165)$$

At 10^9 pairs/s and the specified jitter, the zero-background limiting fidelity is 0.8786. A 0.90 target therefore requires a

brightness reduction as well as timing selection. Optimizing brightness and a 20–5000 ps window, subject to $R \leq 10^9 \text{ s}^{-1}$, gives the values in Table 10. These are peak stationary optima of the declared model; pass-integrated adaptive operation is a different optimization.

4.8.4 Signal and background through the same field stop

The receiver sensitivity treats field acceptance consistently on the signal and noise sides. For a normalized angular point-spread function $P_j(\alpha)$ conditional on entrance-pupil collection,

$$E_j(\theta) = \int_{|\alpha| \leq \theta} P_j(\alpha) d^2\alpha, \quad \eta_j(\theta) = \eta_{j,\text{pre}} E_j(\theta). \quad (166)$$

The same half-angle sets $\Omega = 2\pi(1 - \cos \theta)$ in Eq. (157). A uniform sky is integrated over that accepted solid angle.

For a uniformly illuminated annular pupil and no AO, the phase-only Kolmogorov long-exposure optical transfer function is [42, 43]

$$T(\nu) = T_{\text{pupil}}(\lambda_0 \nu) \exp[-3.44(\lambda_0 \nu / r_0)^{5/3}], \quad (167)$$

where ν is angular spatial frequency in cycles/radian and T_{pupil} is the normalized annular-pupil autocorrelation. Fourier-Bessel inversion gives

$$E(\theta) = 2\pi\theta \int_0^{D/\lambda_0} T(\nu) J_1(2\pi\theta\nu) d\nu. \quad (168)$$

The Fried parameter uses the HV profile integrated from station height to 30 km, with $r_0 = [0.423k^2 \csc \epsilon \int C_n^2(h) dh]^{-3/5}$.

At 810 nm, 20° elevation, and a 10 μrad half-angle, a sea-level-referenced surface term yields $E_M = 0.9251$, $E_H = 0.9643$, and 0.50 dB combined field-stop loss. Re-basing only the surface exponential to each station height yields 0.5314, 0.5723, and 5.17 dB. Equation (168) defines the angular sensitivity. These two prescribed profiles are separate physical assumptions, not fits to one dataset. Doubling the half-angle approximately quadruples the sky term, while the signal gain follows $E(\theta)$.

The receiver calculation is separate from Tables 9–10, which hold the original optical probabilities fixed. Multiplying their probabilities by E_j and recomputing the bounds gives the corresponding combined sensitivity. The present calculation omits amplitude scintillation, nonuniform finite-beam pupil illumination, receiver tracking residuals, and measured AO errors; those model boundaries are consolidated in Section 3.

Table 9: Frozen-peak operating points with full window $\tau = 1$ ns. The last two columns are genuine coincidences/s at the stated mixture-fidelity target. The live-time column assumes four detectors per station and $t_d = 50$ ns.

Wavelength (nm)	$F_\star$	Brightness R_+ (million pairs/s)	C (pairs/s)	C_{live} (pairs/s)
810	0.90	138.42	729,366	570,131
1550	0.90	138.34	182,199	160,008
810	0.95	57.11	300,892	270,728
1550	0.95	57.02	75,098	71,100

Table 10: Joint brightness/window optimization at $F_\star = 0.90$ and $t_d = 50$ ns, retaining the frozen optical probabilities.

λ (nm)	R (10^6 pairs/s)	τ (ps)	C_{live} (pairs/s)
810	489.92	265.51	1.139×10^6
1550	534.14	233.63	4.017×10^5

4.9 Controller messages and service-accounting conventions

Table 11 specifies the messages exchanged between the controller and execution layer. The full service definitions additionally include network access area, auxiliary-channel traffic, and synchronization precision. Stocked-service latency starts with suitable elementary inventory already available; cold-start latency includes generation. Both physical-ready and frame-known completion can be retained. Contact availability is geometrical; service availability includes quality, timing, and resource requirements.

Data availability

The numerical data supporting the findings of this study will be made publicly available upon publication through the QNETSIM repository at <https://github.com/seidkoudia-max/QNETSIM>. The release will include the simulation outputs underlying the figures and tables, scenario configurations, physical-model parameters, random seeds, and analysis and plotting scripts. Software-version and execution-environment information will accompany the data to support reproduction of the reported results.

Code availability

The QNETSIM repository is <https://github.com/seidkoudia-max/QNETSIM>. The reproducibility release will contain the simulator, scenario drivers, physical-model configurations, environment information, and commands for regenerating the results. The source package includes the numerical reanalysis and reference-calculation scripts used in this article.

Author contributions

S.K. developed the quantum-memory models. S.K. and H.P. developed the free-space optical (FSO) channel models, and H.P. modeled the Hong–Ou–Mandel (HOM) interference. F.V. developed the models for segmented entanglement swapping and scheduling. J.C. developed the application programming

interfaces (APIs). I.M. and I.B. supervised the development of the software-defined networking (SDN) component. All authors contributed to the discussion of the QNETSIM results. S.K. and S.C. supervised the overall work.

Acknowledgements

This work was carried out as part of the European Space Agency (ESA) project *Quantum information NETWORKS SIMulator* (QNETSIM), under ESA Contract No. 40001 49440125/NL/FGL/lf.

References

- [1] S. Wehner, D. Elkouss, and R. Hanson, “Quantum internet: A vision for the road ahead,” *Science*, vol. 362, no. 6412, p. eaam9288, 2018.
- [2] L. de Forges de Parny *et al.*, “Satellite-based quantum information networks: use cases, architecture, and roadmap,” *Communications Physics*, vol. 6, p. 12, 2023.
- [3] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi, “Fundamental limits of repeaterless quantum communications,” *Nature Communications*, vol. 8, p. 15043, 2017.
- [4] S. Pirandola, “Satellite Quantum Communications: Fundamental Bounds and Practical Security,” *Physical Review Research*, vol. 3, p. 023130, 2021.
- [5] J. Yin *et al.*, “Satellite-based entanglement distribution over 1200 kilometers,” *Science*, vol. 356, pp. 1140–1144, 2017.
- [6] European Commission, “Commission and European Space Agency sign EuroQCI implementation agreement,” Shaping Europe’s Digital Future, Jan. 2025. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/news/commission-and-european-space-agency-sign-euroqci-implementation-agreement>
- [7] European Space Agency, “Eagle-1,” Official mission overview. [Online]. Available: https://www.esa.int/Applications/Connectivity_and_Secure_Communications/Eagle-1
- [8] S. Khatiri, 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 Information*, vol. 7, p. 4, 2021.
- [9] V. Marulanda Acosta *et al.*, “Analysis of satellite-to-ground quantum key distribution with adaptive optics,” *New Journal of Physics*, vol. 26, p. 023039, 2024, arXiv:2111.06747.
- [10] X. Ma, C.-H. F. Fung, and H.-K. Lo, “Quantum key distribution with entangled photon sources,” *Physical Review A*, vol. 76, p. 012307, 2007.
- [11] J. Szlachetka, K. Joarder, and P. Kolenderski, “Ultrabright source of non-degenerate polarization-entangled photon pairs based on off-the-shelf polarization optics,” *Applied Physics Letters*, vol. 123, p. 144001, 2023.

Table 11: Controller–execution contract. Every message carries request/plan epoch and simulation timestamps; resource references include allocation versions.

Message/action	Payload	Resource meaning
ENT_REQUEST	Endpoints, service, $F_{\min}$, deadline	Creates request state and admission decision.
StartSchedule	Link/source ID, interval, rate/modes	Authorizes attempts within source and terminal budgets.
REPORT_STORAGE	Endpoint/mode IDs, write and herald times, state tag	Updates the controller after signaling delay.
InstructSwap	BSM ID, two operands/versions, expiry	Atomically validates and reserves two local qubits.
SWAP_RESULT	Outcome, consumed IDs, surviving endpoints, frame	Records destructive execution and surviving histories.
StoreToken	Derived pair identity and surviving endpoints	Associates metadata; an extra physical write is explicit.
DeliverPair	Pair ID, frame, quality, service-ready time	Announces the usable service at the endpoints.
ClearRequest	Request ID and epoch	Clears residual resources and invalidates stale actions.

- [12] C. Bonato, A. Tomaello, V. Da Deppo, G. Naletto, and P. Villoresi, “Feasibility of satellite quantum key distribution,” *New Journal of Physics*, vol. 11, p. 045017, 2009.
- [13] A. V. Gorshkov, A. André, M. Fleischhauer, A. S. Sørensen, and M. D. Lukin, “Universal approach to optimal photon storage in atomic media,” *Physical Review Letters*, vol. 98, p. 123601, 2007.
- [14] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, “Photon storage in Λ -type optically dense atomic media. II. free-space model,” *Physical Review A*, vol. 76, p. 033805, 2007, arXiv:quant-ph/0612083.
- [15] J. Calsamiglia and N. Lütkenhaus, “Maximum efficiency of a linear-optical Bell-state analyzer,” *Applied Physics B*, vol. 72, pp. 67–71, 2001.
- [16] M. Gündoğan, J. S. Sidhu, V. Henderson, L. Mazzarella, J. Wolters, D. K. L. Oi, and M. Krutzik, “Proposal for space-borne quantum memories for global quantum networking,” *npj Quantum Information*, vol. 7, p. 128, 2021.
- [17] W. Kozłowski, S. Wehner, R. Van Meter, B. Rijsman, A. S. Cacciapuoti, M. Caleffi, and S. Nagayama, “Architectural principles for a quantum internet,” RFC Editor, RFC 9340, Mar. 2023, informational; Internet Research Task Force. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc9340.html>
- [18] T. Coopmans *et al.*, “NetSquid, a NETwork Simulator for QUantum Information using Discrete events,” *Communications Physics*, vol. 4, p. 164, 2021.
- [19] X. Wu, A. Kolar, J. Chung, D. Jin, T. Zhong, R. Kettimuthu, and M. Suchara, “SeQeNCE: A Customizable Discrete-Event Simulator of Quantum Networks,” arXiv:2009.12000, 2020, revised 2021.
- [20] J. Meister, P. Kleinpaß, and D. Orsucci, “Simulation of satellite and optical link dynamics in a quantum repeater constellation,” *EPJ Quantum Technology*, vol. 12, p. 5, 2025.
- [21] QuNetSim authors, “Introduction,” Official QuNetSim documentation. [Online]. Available: <https://tqsd.github.io/QuNetSim/intro.html>
- [22] Baidu Quantum, “qcompute-qnet,” PyPI, official-maintainer package description and tutorial inventory. [Online]. Available: <https://pypi.org/project/qcompute-qnet/>
- [23] R. Satoh *et al.*, “QuISP: a Quantum Internet Simulation Package,” in *IEEE International Conference on Quantum Computing and Engineering*, 2022, pp. 353–364, preprint: arXiv:2112.07093 (2021).
- [24] K. Boone, J.-P. Bourgoin, E. Meyer-Scott, K. Heshami, T. Jennewein, and C. Simon, “Entanglement over global distances via quantum repeaters with satellite links,” *Physical Review A*, vol. 91, p. 052325, 2015.
- [25] C. Liorni, H. Kampermann, and D. Bruß, “Quantum repeaters in space,” *New Journal of Physics*, vol. 23, p. 053021, 2021.
- [26] J. Wallnöfer *et al.*, “Simulating quantum repeater strategies for multiple satellites,” *Communications Physics*, vol. 5, p. 169, 2022.
- [27] R. Yehia, M. Schiavon, V. Marulanda Acosta, T. Coopmans, I. Kerenidis, D. Elkouss, and E. Diamanti, “Connecting quantum cities: simulation of a satellite-based quantum network,” *New Journal of Physics*, vol. 26, p. 073015, 2024.
- [28] Y. Shi, Z. Wang, B. Yuan, G. Wu, and Z. Zhao, “QUASAR: Quantum satellite architecture and routing simulator,” arXiv:2609.04920v1, 2026, preprint.
- [29] S. Koudia, J. ur Rehman, and S. Chatzinotas, “Space-Based Quantum Internet: Entanglement Distribution in Time-Varying LEO Constellations,” arXiv:2409.17032, 2024.
- [30] NetSquid authors, “Official simulator overview and interface examples,” Official NetSquid website. [Online]. Available: <https://netsquid.org/>
- [31] Ryu maintainers, “Ryu project README,” 2026. [Online]. Available: <https://github.com/faucetsdn/ryu>
- [32] European Telecommunications Standards Institute, “Quantum key distribution (QKD); protocol and data format of REST-based key delivery API,” ETSI, Group Specification ETSI GS QKD 014 V1.1.1, Feb. 2019. [Online]. Available: https://www.etsi.org/deliver/etsi_gs/QKD/001_099/014/01.01.01_60/gs_qkd014v010101p.pdf
- [33] —, “Quantum key distribution (QKD); control interface for software defined networks,” ETSI, Group Specification ETSI GS QKD 015 V2.1.1, Apr. 2022. [Online]. Available: https://www.etsi.org/deliver/etsi_gs/QKD/001_099/015/02.01.01_60/gs_QKD015v020101p.pdf
- [34] —, “Quantum key distribution (QKD); protocol and data format of REST-based interoperable key management system API,” ETSI, Group Specification ETSI GS QKD 020 V1.1.1, Jun. 2026. [Online]. Available: https://www.etsi.org/deliver/etsi_gs/QKD/001_099/020/01.01.01_60/gs_QKD020v010101p.pdf
- [35] M. Fleischhauer and M. D. Lukin, “Quantum memory for photons: Dark-state polaritons,” *Physical Review A*, vol. 65, p. 022314, 2002.

-
- [36] B. Kraus, W. Tittel, N. Gisin, M. Nilsson, S. Kröll, and J. I. Cirac, “Quantum memory for nonstationary light fields based on controlled reversible inhomogeneous broadening,” *Physical Review A*, vol. 73, p. 020302(R), 2006.
- [37] N. Sangouard, C. Simon, M. Afzelius, and N. Gisin, “Analysis of a quantum memory for photons based on controlled reversible inhomogeneous broadening,” *Physical Review A*, vol. 75, p. 032327, 2007.
- [38] J. Bylander *et al.*, “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit,” *Nature Physics*, vol. 7, pp. 565–570, 2011.
- [39] P. Drmota *et al.*, “Robust quantum memory in a trapped-ion quantum network node,” *Physical Review Letters*, vol. 130, p. 090803, 2023.
- [40] F. Reiter and A. S. Sørensen, “Effective operator formalism for open quantum systems,” *Physical Review A*, vol. 85, p. 032111, 2012, arXiv:1112.2806.
- [41] N. Lütkenhaus, J. Calsamiglia, and K.-A. Suominen, “Bell measurements for teleportation,” *Physical Review A*, vol. 59, p. 3295, 1999.
- [42] L. C. Andrews and R. L. Phillips, *Laser Beam Propagation through Random Media*, 2nd ed. Bellingham, WA: SPIE Press, 2005.
- [43] D. L. Fried, “Optical resolution through a randomly inhomogeneous medium for very long and very short exposures,” *Journal of the Optical Society of America*, vol. 56, no. 10, pp. 1372–1379, 1966.
- [44] M. A. Al-Habash, L. C. Andrews, and R. L. Phillips, “Mathematical model for the irradiance probability density function of a laser beam propagating through turbulent media,” *Optical Engineering*, vol. 40, no. 8, pp. 1554–1562, 2001.
- [45] D. Vasylyev, A. A. Semenov, and W. Vogel, “Atmospheric quantum channels with weak and strong turbulence,” *Physical Review Letters*, vol. 117, p. 090501, 2016.
- [46] D. A. Vallado, P. Crawford, R. Hujsak, and T. S. Kelso, “Revisiting Spacetrack Report #3,” AIAA, AIAA Paper 2006-6753, 2006, reference implementation and test cases maintained through CelesTrak.
- [47] B. Rhodes, “Skyfield: high precision research-grade positions for planets and Earth satellites,” Software documentation, “Coordinates” and “Positions”. [Online]. Available: <https://rhodesmill.org/skyfield/>
- [48] H. Tanji, S. Ghosh, J. Simon, B. Bloom, and V. Vuletić, “Heralded single-magnon quantum memory for photon polarization states,” *Physical Review Letters*, vol. 103, p. 043601, 2009.
- [49] M. Afzelius, C. Simon, H. de Riedmatten, and N. Gisin, “Multimode quantum memory based on atomic frequency combs,” *Physical Review A*, vol. 79, p. 052329, 2009.