

Conditional Queries with Finite Historical States and Periodic Replacement

Guillaume Bolduc

Abstract—We study a deterministic conditional-query problem in which a sampled index is encoded before a later query, the current index may change, and every surviving input-dependent distinction belongs to a finite historical alphabet. On a logarithmic probability grid with P points, displacement bound A , spacing h , and at most K historical states, we obtain the exact worst Euclidean query error $\tanh(hD_*/4)/\sqrt{2}$. The optimal expanded-fiber diameter is $D_* = P - 1$ when $2A \geq P - 1$, and $D_* = 2A + \lceil (P - 2A)/K \rceil - 1$ otherwise. Consecutive encoder blocks and query-dependent midpoint responses attain the bound. Rational geometric parameters give a bit-machine implementation with separately counted workspace and output-precision error. A two-slot periodic construction retains bounded-age codes under correct fresh acquisition, bounded acknowledged service and initialized-only scoring. After a finite attack on stored symbols or forward payload fields, clean overwrite restores current-task accuracy while acquisition, executable encoder/decoder and control remain correct. A separate one-hop replicated crash-erasure model admits an exact period-endpoint law under all-copy lifetime, full-drain and no-side-channel contracts. Quadratic objectives transfer the query law to optimizer-estimation error. A normalized family expresses the conditional error in terms of retained mass. The results have distinct information, timing and fault assumptions; they do not establish a general network resource optimum or Byzantine source/control recovery.

Index Terms—information constraints, membership changes, communication delay, deterministic minimax estimation.

I. INTRODUCTION

Conditioning a probability row on surviving coordinates is nonlinear. The resources for a later conditional query depend on its declared information interface. The relevant resources therefore depend on the specified query, the order in which information becomes available, and the distinctions that remain accessible to the decoder.

This article analyzes one finite family. An encoder observes a sampled index before a query scale is revealed. The current index may differ from the sample. The target is the probability row obtained by admitting one coordinate and deleting another. Every surviving input-dependent register, packet, metadata distinction or external record accessible to the answer counts as historical information. We derive an exact deterministic minimax law at this information cut. We then realize its upper construction in specified replacement and crash-erasure regimes, keeping historical state, peak workspace, communication, timing, trust and control resources separate.

The principal calculations are the single-cut minimax law (Theorem 1) and its matching converse and construction in the specified one-hop crash model (Theorem 5). Periodic replacement and clean overwrite are realizations with fresh

correct acquisition. The arithmetic result is a workspace upper construction; the optimizer and retained-mass results are transfers of the same query law.

The probability geometry is related to the Hilbert-metric/total-variation estimate of Cohen and Fausti [1]. Under their convention the distance is an ℓ^1 distance; our binary first-coordinate diameter is half that distance. The admissible query set contains a scale attaining the diameter. The hyperbolic function by itself is not asserted as a new metric identity.

Maximum-distortion quantization is a direct predecessor. Tuncel *et al.* [9, Secs. 1–2] treat finite as well as asymptotic fixed-length codes, and express scalar feasibility as a covering problem. Their interval-graph result [9, Theorem 1] also supplies constructive interval-covering machinery. The response-vector reduction below places our problem within finite maximum-distortion coding. A later-query interpretation alone is therefore not a distinct coding framework. Our calculation evaluates the robust distortion for the specified logistic grid and its clipped displacement neighborhoods, then derives (4)–(5); no non-derivability or priority claim follows merely from that evaluation.

Muresan and Effros [10, Secs. II–III and VIII] optimize expected rate and distortion through scalar-code partitions and shortest paths. Their conventional fixed-rate contiguity result does not justify contiguity for every network or decoder-side-information model; their Section VIII gives failures of that extension. Here the criterion is a maximum over hidden displacement and every query. Lemma 2 proves intervalization for this particular expanded-fiber cost, then evaluates its cover. It is not a general contiguity theorem or a claim that interval methods are new.

Kaspi and Merhav [11, Secs. II–III, Theorem 1] study sequential instantaneous coding with correlated decoder side information, expected code length and average distortion, obtaining scalar encoder/decoder structures with time sharing. Here every query is admissible after every sample, the query does not rule out a sampled index, the current index remains hidden, and the resource is a hard alphabet cardinality. Their scalar optimality statement is not substituted for our worst-case diameter calculation. Functional distortion coding is also established [2]. Table I records the inspected formulations; it does not certify absence of an earlier equivalent evaluation.

Open-network information limits for stochastic current-mean estimation [3], distributed sliding-window monitoring [4], and dynamic atomic memory [5] address different interfaces and resources. Delayed and indefinitely open quantized coordination also has direct predecessors [6]. Uniform and nonuniform delays in priority consensus are treated in [7]. Here communication delays enter an observation-age and

TABLE I
CLAIM-LOCAL COMPARISON OF THE INSPECTED CODING FORMULATIONS.

Formulation	Information and criterion	Relation to the present calculation
Tuncel <i>et al.</i> [9]	Finite/asymptotic fixed-length coding; maximum distortion; finite reproduction alphabet	Our response-vector reduction is in this general class; the logistic radius and clipped cover are evaluated explicitly here
Muresan–Effros [10]	Finite scalar partitions; expected rate/distortion; convex-codecell design	Our maximum expanded-fiber cost requires its own intervalization proof; no universal contiguity transfer
Kaspi–Merhav [11]	Sequential source/decoder side information; expected instantaneous length; average distortion	Same scalar encoder/decoder pattern, different optimization criterion and information constraints
Present (2)	Encoder observes j ; decoder sees code and k ; hidden z ; at most K states	Exact finite P, A, K radius, attaining query responses, and separate one-hop endpoint reduction

service contract, rather than a delay-invariant conservation law. Randomized quantized consensus can bound state/message representations with high probability while its reconstruction uses a separate precision assumption [8]; that statement is not a deterministic hard bound on all arithmetic. No superiority comparison follows without matching targets, information, resources and scoring times.

The mathematical claims below are self-contained. The comparisons identify prior mechanisms and differences of scope; no first-publication claim is made. General changing graphs, cumulative historical objectives, arbitrary Byzantine control, total network resource optimality and physical realization remain outside the guarantees.

II. CONDITIONAL-QUERY MODEL

Let $P \geq 2$, $A \geq 0$, and $K \geq 1$ be integers, and set $N = P - 1$. Fix $h > 0$ and $0 < t_{\min} < e^{-Nh}$. Public constants are independent of the task input. The encoder observes

$$j \in \{0, \dots, N\}$$

and stores $s = e(j)$ in an alphabet of cardinality at most K . Subsequently the environment chooses

$$z \in \{0, \dots, N\}, \quad |z - j| \leq A, \quad k \in \{0, \dots, 2N\}.$$

Every such query is admissible after every allowed pair (j, z) . The encoder is chosen before the query and does not observe k . The decoder receives only s, k and public constants.

All other input-dependent distinctions accessible to the answer are included in s : clocks, state length, packet order, addresses, labels, physical state, advice, random seeds, pending copies and external storage. Strategies are deterministic. This is a cardinality contract at a declared historical cut, not a peak-memory contract.

Define

$$t_z = t_{\min} e^{hz}, \quad v_k = t_{\min} e^{hk/2}, \quad \beta_k = \frac{v_k}{1 + v_k}.$$

The original row is $p(z) = (1 - t_z, t_z)$ on ordered labels (a, b) . Admission of c gives

$$J_k(p) = ((1 - \beta_k)p_a, (1 - \beta_k)p_b, \beta_k).$$

Deleting a and normalizing the retained coordinates gives

$$\Phi(z, k) = (F_k(z), 1 - F_k(z)), \quad F_k(z) = \frac{t_z}{t_z + v_k}. \quad (1)$$

Each query is evaluated on its declared current row; queries do not cumulatively mutate a separate historical table. The

task specifies this operation order. Every intermediate row is positive normalized: $0 < t_z < 1$, $0 < \beta_k < 1$, and the retained mass is positive. Reversing the order can change the result; for $p = (3/4, 1/4)$, $\beta = 1/2$, admission then deletion gives $(1/5, 4/5)$, whereas deletion then admission gives $(1/2, 1/2)$.

For a deterministic probability-row decoder d , set

$$\mathcal{E}(P, A, K, h) = \inf_{e, d} \max_{j, z, k \text{ admissible}} \|d(e(j), k) - \Phi(z, k)\|_2. \quad (2)$$

The maximum is over inputs, not a probability distribution. The task target is not defined by the algorithm's schedule.

A. Reduction to maximum-distortion quantization

A reproduction may be the complete response vector $y = (y_k)_{k=0}^{2N}$, each y_k a binary probability row. Define

$$\delta_A(j, y) = \max_{\substack{z \in \{0, \dots, N\}, \\ 0 \leq k \leq 2N}} \|y_k - \Phi(z, k)\|_2.$$

Then (2) is exactly $\inf_{e, y_1, \dots, y_K} \max_j \delta_A(j, y_{e(j)})$: storing a symbol selects a public response vector and answering reveals its component k . Conversely every such quantizer defines a decoder. Lemma 1 shows that the vector of query midpoints for the extreme sampled indices of each nonempty fiber minimizes that fiber's distortion. Thus the finite menu indexed by $0 \leq l \leq u \leq N$, with at most $P(P + 1)/2$ vectors, contains an optimum. This finite-alphabet reduction uses the proved fiber geometry, not a new general coding theorem.

III. EXACT HISTORICAL-STATE MINIMAX LAW

A. Radius of one encoder fiber

Lemma 1. For a nonempty encoder fiber $C \subseteq \{0, \dots, N\}$, define

$$\begin{aligned} l &= \min C, & u &= \max C, \\ L &= \max(0, l - A), & U &= \min(N, u + A), \\ D_C &= U - L. \end{aligned}$$

Its optimal worst-case query error is

$$R(C) = \frac{\tanh(hD_C/4)}{\sqrt{2}}. \quad (3)$$

Proof. Compatible current indices form a union of clipped integer intervals; this union need not be an interval. Its extremes are L, U , both attainable. Since F_k is increasing, the extreme first coordinates are $F_k(L), F_k(U)$. The triangle inequality bounds the worst error of any output below by half the

distance between the extreme rows. Their midpoint is a probability row and attains that bound for all compatible indices. Thus the optimal error at fixed k is $(F_k(U) - F_k(L))/\sqrt{2}$.

For $a = t_L$, $b = t_U$, $v > 0$, the scalar difference is

$$\frac{b}{b+v} - \frac{a}{a+v} = \frac{(b-a)v}{(a+v)(b+v)}.$$

If $a = b$, it vanishes. Otherwise its derivative has the sign of $ab - v^2$, so its maximum is at $v = \sqrt{ab}$, with value

$$\frac{\sqrt{b} - \sqrt{a}}{\sqrt{b} + \sqrt{a}} = \tanh\left(\frac{\log(b/a)}{4}\right) = \tanh(hD_C/4).$$

The admissible integer query $k = L + U$ supplies this exact scale. A fixed query-dependent midpoint decoder attains the bound. $\square$

B. Covering arbitrary fibers

Lemma 2. The smallest possible maximum expanded-fiber diameter among partitions into at most K nonempty fibers is

$$D_* = \begin{cases} N, & 2A \geq N, \\ 2A + \lceil (P - 2A)/K \rceil - 1, & 2A < N. \end{cases} \quad (4)$$

Proof. If $2A \geq N$, an integer $j \in [0, N]$ with $N - A \leq j \leq A$ exists. Its compatible current indices include both endpoints. Its fiber therefore has diameter N , and no fiber exceeds N .

Suppose $2A < N$. Sample $j = A$ forces a diameter at least $2A$. Arbitrary fibers can be reduced to consecutive blocks without increasing their maximum expanded diameter. Replace each fiber by its integer hull $[l, u]$. These hulls cover the domain. Starting at the first unassigned index x , choose a hull containing x with greatest right endpoint u , assign $[x, u]$, and continue at $u+1$. For the selected hull $[l, u]$, the new block has left expanded endpoint $\max(0, x - A) \geq \max(0, l - A)$ and the same right expanded endpoint $\min(N, u + A)$. Its diameter cannot increase. That hull cannot be selected again after its endpoint is passed. This constructs a consecutive partition with at most the original number of fibers and no larger diameter.

Consider an integer candidate $2A \leq D < N$, and write $q = D - 2A + 1 \geq 1$. The longest admissible first block is $[0, D - A]$, leaving next start $A + q > A$. A later block starting at $l > A$ can finish at N exactly when $l \geq N + A - D$. If it cannot finish, its longest admissible endpoint is $l + D - 2A = l + q - 1$, covering q indices.

Greedy longest blocks use the fewest blocks. Indeed their greatest admissible endpoint $g(l)$ is nondecreasing in l , since increasing the start cannot increase the expanded diameter of a retained endpoint. For precision, let greedy and competing partitions have last covered indices G_r, H_r after r blocks, with $G_0 = H_0 = -1$. Inductively, if $G_{r-1} \geq H_{r-1}$ and greedy has not finished, monotonicity gives $G_r = g(G_{r-1} + 1) \geq g(H_{r-1} + 1) \geq H_r$. Once greedy finishes the comparison is automatic. Thus greedy uses no more blocks than any feasible consecutive partition. After $K-1$ nonfinal blocks the next start is $A + (K-1)q$. The last block can finish precisely when

$$A + (K-1)q \geq N + A - D \iff Kq \geq P - 2A.$$

Earlier greedy completion also satisfies this inequality. For $K = 1$, it would force $D \geq N$, so no $D < N$ suffices. The least feasible diameter is (4). If it equals N , the one-block partition attains it. $\square$

Theorem 1. In model (2),

$$\mathcal{E}(P, A, K, h) = \frac{\tanh(hD_*/4)}{\sqrt{2}}. \quad (5)$$

The infimum is attained.

Proof. Lemma 1 gives each encoder's optimized radius as the maximum of $\tanh(hD_C/4)/\sqrt{2}$. Monotonicity and Lemma 2 give the lower bound over every partition, including nonconsecutive fibers. For the upper bound use the greedy partition with diameter D_* , or one block if $D_* = N$. Store the block ordinal. Its public endpoints determine L, U . Return

$$d(s, k) = (m, 1 - m), \quad m = \frac{F_k(L) + F_k(U)}{2}. \quad (6)$$

No unknown j or z is used after the cut. Lemma 1 proves attainment. $\square$

The quantifier order is one encoder before every input and future query. The result is not an encoder selected separately for each query. If $A = 0$, then $D_* = \lceil P/K \rceil - 1$; in particular $K \geq P$ gives zero ideal error. If $K = 1$, then $D_* = N$. If $2A \geq N$, historical storage cannot remove the missing-current-index uncertainty floor. The excluded singleton $P = 1$ has one known target and zero error.

C. Historical bits and fixed-length communication

Corollary 1. For integer historical budget $M \geq 0$, a cut with at most 2^M states has (5) with $K = 2^M$. For a single fixed-length C -bit one-way payload followed by M historical bits, with no additional input-dependent channel, use

$$K = 2^{\min(M, C)}, \quad M, C \in \mathbb{Z}_{\geq 0}. \quad (7)$$

Proof. The retained state is a deterministic function of the payload. Its cardinality is at most both 2^C and 2^M . Conversely the attaining block ordinal fits in both budgets. Apply Theorem 1. $\square$

These are single-cut historical budgets. Variable lengths, interactive transcripts, timing channels, aggregate network memory and peak workspace require different contracts. Source inputs, query registers, public constants, packet replicas, identifiers and arithmetic remain separately counted resources.

D. Output precision

Proposition 1. For integer $g \geq 1$, replacing (6) by the nearest midpoint of one of 2^g equal cells in $[0, 1]$ gives a positive probability row with error at most

$$\frac{\tanh(hD_*/4)}{\sqrt{2}} + \frac{1}{\sqrt{2} 2^g}. \quad (8)$$

A scalar evaluation error certified at most η , followed by clipping to $[0, 1]$ and the same rounding, adds at most $\sqrt{2}\eta$.

Proof. Midpoint-cell scalar error is at most 2^{-g-1} . Pair Euclidean error is $\sqrt{2}$ times scalar error. Clipping cannot increase error from the ideal $m \in [0, 1]$. Add these increments to Theorem 1's ideal radius by triangle inequality. $\square$

Equation (8) is an upper bound, not a sharp finite-output minimax equality.

IV. FINITE ARITHMETIC AND INFORMATION SURVIVAL

A. Rational geometric grids

Fix integers $a > b \geq 1$, set $r = a/b$, $h = 2 \log r$, and choose rational $0 < t_{\min} < r^{-2N}$. Then

$$F_k(z) = \frac{r^{2z}}{r^{2z} + r^k}.$$

For $d = 2z - k$,

$$F_k(z) = \begin{cases} a^d / (a^d + b^d), & d \geq 0, \\ b^{-d} / (a^{-d} + b^{-d}), & d < 0. \end{cases} \quad (9)$$

No numerical evaluation of h is needed by this decoder.

Use a deterministic multitape bit machine with a constant number of binary work tapes, read-only tapes for public parameters and decoder inputs (s, k) , and a write-only output stream. Workspace is the peak number of work-tape cells visited during one answer, including operand copies, products, carries, quotients, remainders, rounding buffers and explicit counters. Fixed finite control, read-only input/public storage, the retained historical code, emitted output and protocol metadata are separate resources; any copy of those data onto a work tape counts as workspace. This is a bit-space construction for the decoder, not a bound on a particular language runtime, total network storage or service time. Public A, K may be replaced by $\min(A, N), \min(K, P)$ for this finite family.

Proposition 2 (decoder bit workspace). In this machine model, let B be the bit length of a , $\ell = 2NB + 2$, and $g \geq 1$. Each endpoint fraction has positive numerator and denominator of fewer than ℓ bits. A midpoint fraction has numerator and denominator of at most $2\ell + 1$ bits. Given a canonical block ordinal and query, its g -bit cell index and positive midpoint-cell probability row can be computed exactly with $O(NB + g + \log P)$ peak work-tape bits. Historical code, public/read-only inputs, emitted output and protocol metadata are counted separately; working copies are included in this bound.

Proof. The complete integer-arithmetic construction and storage bound are in Appendix A.

The common $t_{\min}$ cancels in (9). Realizing the original probability rows or admission masses still requires its separate representation. Noncomputable arbitrary public reals do not receive a finite-computation guarantee.

B. Complete erasure

A complete erasure cut removes every sampled-input distinction from the future decoder's causal past. New query inputs and scheduling are input-independent and there is no fresh distinguishing observation.

Proposition 3. At a complete erasure cut with $A = 0$, every deterministic answer scheme has worst error at least $\tanh(hN/4)/\sqrt{2}$, attained by the query-dependent midpoint of the endpoint answers.

Proof. Decoder information is identical for sampled indices $0, N$. Query $k = N$ is admissible in both. Lemma 1 supplies half their distance and an attaining midpoint for every query. Additional workers with identical information cannot change this argument. $\square$

A packet, durable store or fresh measurement can carry the missing distinction if counted. The result does not require a particular worker to survive.

C. Controlled handoff

Consider $H \geq 1$ successive overlapping handoffs. Incarnation I_0 has a supplied, counted initial code $s = e(j)$; each newcomer initially has no task information. A correct route service binds intended source and successor incarnations. A payload carries its phase, endpoints and code. The receiver accepts its expected transaction once, validates the range, stores the code atomically and sends a matching acknowledgement. Duplicates are idempotent, and unexpected or retired incarnations cannot accept it. The sender retains its code and retires only after processing that acknowledgement. Delivery, storage and acknowledgement processing have finite bounds. No false-data or arbitrary replay is permitted. Queries are scored only after coherent initialization and do not change the code. All metadata and timing are sampled-input independent.

Proposition 4. Every initialized successor after a completed handoff holds exactly the original code. It can use Theorem 1's decoder wherever the stipulated uncertainty $|z - j| \leq A$ holds, for every finite H .

Proof. The initial source has s . Correct delivery stores that exact s before the matching acknowledgement permits source retirement. Idempotence and endpoint/phase checks preserve it. Induction proves the claim for each handoff. The decoder then receives the same historical code and allowed query as in Theorem 1. $\square$

Stop-and-wait limits outstanding logical transactions. For the physical count here, additionally assume one delivery per emitted message, no channel replication and no transport retry: at most one data copy and one acknowledgement are outstanding per transaction. Idempotent acceptance alone does not bound physical copies under replication. Overlap stores up to two historical codes; source inputs, arithmetic and packet copies have separate costs. Finite-horizon identifiers require $O(\log(H+1))$ bits. Infinite monotone identifiers grow. Bounded labels require an additional proved packet-lifetime, no-replay and rebinding contract.

Handoff does not create a new observation. If a latent signal satisfies

$$|x(t) - x(s)| \leq \rho|t - s|, \quad x(s) = hj, \quad x(t) = hz,$$

then

$$|z - j| \leq \left\lfloor \frac{\rho(t - s)}{h} \right\rfloor. \quad (10)$$

Only the scored endpoints are grid-valued. A continuous signal constrained to a finite grid at every intermediate time could not change. After H handoffs of duration at most δ and additional waiting w , an age bound $H\delta + w$ suffices. A uniform per-link delay does not imply a replacement-independent observation age.

V. PERIODIC REPLACEMENT WITH BOUNDED SOURCE AGE

This section gives an achievability construction, not a min-max theorem over general distributed strategies.

Two endpoint slots hold at most one incarnation each. At $t_q = qT$, a fresh receiver I_{q+1} arrives in the empty alternating slot; incumbent I_q is its source. A correct acquisition channel samples the same exogenous ρ -Lipschitz task at t_q , returning grid index j_q . Acquisition and encoding finish within σ . During computation the source retains its prior code. It then atomically replaces that code by $e(j_q)$, clears the raw observation and input-dependent workspace, and sends one payload.

Forward delivery, receiver validation/storage, acknowledgement return and source acknowledgement processing have nonnegative finite bounds d_f, p_r, d_b, p_s . Put

$$B_0 = \sigma + d_f + p_r, \quad C_0 = \sigma + d_f + p_r + d_b + p_s < T. \quad (11)$$

Uniform directional delay means $d_f = d_b$; unequal directional bounds are allowed. Realized times may vary below the bounds.

At every scored query, the current signal has the grid value $x(t) = hz$ for an integer z ; intermediate signal values need not be grid-valued. Routes bind current source/destination incarnations. The receiver validates expected phase, endpoints and code range before becoming ready. The source retains its code until the matched acknowledgement. Correct channels produce no conflicting payload. Equal-time primitives use the causal order birth/rebinding, sample, encoding/overwrite, send, deliver/store, acknowledge, return/process and retire; queries occur between completed primitives. Waiting receivers are not scored. Retired sources stop all priority operations. The selected construction has one delivery per emitted message, no channel replication, and no transport retry or arbitrary retained replay. There is one physical data word and one acknowledgement per transaction.

All encoders and decoders use the fixed uncertainty

$$\bar{A} = \left\lceil \frac{\rho(T + \sigma)}{h} \right\rceil. \quad (12)$$

Lemma 3. Each transaction finishes before the next period. The predecessor remains available until receiver storage. At most two incarnations are present, and each I_q , $q \geq 1$, retires within $T + C_0$ of birth.

Proof. Storage finishes by $t_q + B_0$, and acknowledgement processing by $t_q + C_0 < t_{q+1}$. Retirement follows storage, freeing the old slot before the next arrival. The initialized successor becomes the next source. Since I_q is born at t_{q-1} and retires by $t_q + C_0$, its lifetime has the stated bound. After coherent startup initialization, induction gives at least one ready incarnation and the slot count. A waiting initial source is not scored. $\square$

Lemma 4. After the declared startup convention, a code used by an initialized worker at time t originated at some sample time s with $0 \leq t - s \leq T + \sigma$.

Proof. Code for j_q originates at t_q . Its source overwrites by $t_q + \sigma$, so its age at overwrite is at most σ . The source can then use it until retirement by $t_q + C_0$, giving age at most $C_0 < T$. The receiver stores it by $t_q + B_0$, uses it through the next period boundary t_{q+1} , and retains it during its next computation until overwrite by $t_{q+1} + \sigma$. Its age throughout that interval is at most $T + \sigma$. Initially I_0 either has a sample

at most T old at t_0 , whose age remains at most $T + \sigma$ until first overwrite, or waits for its first new code before answering. These phases cover every scored use. Equation (10) gives (12). $\square$

Theorem 2 (periodic bounded-age construction). Assume the ordered two-slot protocol above: a correct fresh ρ -Lipschitz grid observation each period; valid source/destination bindings; acquisition, encoding, delivery, storage and acknowledgement satisfying $C_0 < T$; retention until matched acknowledgement; the stated startup convention; and no scoring of an uninitialized receiver. The answer uses only the retained code, allowed query and public constants, with no additional sampled-input channel through control or timing. Then every scored query under the fixed encoder and decoder (6), in every finite or infinite execution satisfying these contracts, has error at most

$$\frac{\tanh(hD_*(P, \bar{A}, K)/4)}{\sqrt{2}}, \quad (13)$$

plus the finite-output increment in (8) in the rational subfamily.

Proof. Lemmas 3–4 identify the actual stored code, its initialization and the precise uncertainty premise of Theorem 1. The query is revealed after encoding and remains in the allowed family. Apply (6) at each scored service event. Proposition 2 supplies finite arithmetic and rounding. The argument uses no final membership event, stationary task or permanently surviving incarnation. $\square$

The age-family bound (13) need not be sharp for the reachable periodic task family. Take $P = K = 2$, $T = 1$, $\sigma = 2/5$, $\rho/h = 4/5$, and $d_f = p_r = d_b = p_s = 1/10$, with clean startup waiting for its first fresh code. Then $C_0 = 4/5 < T$ and $\bar{A} = 1$, giving bound $1/(3\sqrt{2})$ when $h = \log 4$. Yet adjacent grid-valued samples differ by less than h , so they have the same index. At a grid-valued query between samples, its distance from the preceding sample is also less than h ; the queried index is therefore the same. The identity encoder with two states has zero error on this actual family. Equation (13) is an upper bound obtained from the larger age family, not an exact periodic minimax law.

The construction consumes correct observation, period timers, routing and service resources. It does not implement membership agreement. Historical codes occupy up to two copies; raw indices, arithmetic, one data/ack pair and control metadata remain separate. A finite horizon uses growing identifiers. If all copies finish within their transaction window $C_0 < T$, no arbitrary replay exists, and old route bindings retire before reuse, phase modulo three and two slot identities suffice. Those extra premises are essential to bounded metadata.

VI. RESTRICTED DATA-CODE RECOVERY AND OPTIMIZER ESTIMATION

A. Recovery after finite code corruption

Use Section 5's task and protocol. An attacker active on $[t_a, t_b]$, with $0 \leq t_a \leq t_b$, may replace historical code registers and forward data code fields by arbitrary valid alphabet elements. The number of such replacements during the episode is unrestricted. It cannot alter task values, fresh observations,

encoder/decoder code, parameters, queries, control metadata, clocks, routes, lifecycle, acknowledgement service or packet lifetime. It cannot create packets or replay beyond a bounded transaction. This is corruption of stored historical symbols and forward data fields under trusted acquisition, executable code and control. It does not permit malicious modification of the encoder, decoder or controller. Before recovery no smaller certified error than the probability-space diameter is asserted.

Theorem 3 (recovery after finite stored-symbol corruption). Assume Theorem 2's acquisition, service, initialization and answer-interface contracts, with task values, fresh observations, executable encoder/decoder, clocks, routes, lifecycle, acknowledgement service and packet lifetime protected. Let an attack on $[t_a, t_b]$ be restricted to valid historical symbols and forward data fields, without packet creation or out-of-window replay. Let q_* be the first transaction with $t_{q_*} > t_b$, and κ_* the event cut just after its source atomically writes the clean fresh code. Every scored query serviced after that cut satisfies (13). The *overwrite cut's timestamp* is at most $t_b + T + \sigma$. Queries at a strictly later physical time obey the bound; a query at that timestamp is covered only when it follows κ_* in event order.

Proof. A strictly later periodic transaction starts within T of t_b . Earlier transactions and their copies finish before it starts by (11). All predecessors have retired. Correct acquisition and encoding depend only on the fresh observation, not the corrupted code. By $t_{q_*} + \sigma$ atomic replacement removes the old value. Its new receiver either waits or stores the same clean payload. These are the only extant incarnations, and later acquisitions/transactions remain clean. Lemma 4 supplies the age bound. A query preceding overwrite at the same timestamp is excluded; no right-continuity convention is presumed. No node detects t_b , and recovery uses ordinary periodic acquisition. $\square$

The guarantee restores accuracy for the current task, not an unknown preattack historical target or a conserved mean. The next acquisition occurs within T and records the task value at that acquisition time. Its code is computed within σ and stored by the receiver within B_0 of acquisition. A transient change between acquisitions need not be observed. For example, with $T = 1$, the signal $x(t) = 2h \min(t, 1 - t)$ on $[0, 1]$, extended by zero, is $2h$ -Lipschitz: every integer-time sample is zero, although $x(1/2) = h$. The age-error guarantee does not imply detection of every excursion.

The strict boundary matters. Set $P = K = 2$, $\bar{A} = 0$, $h = \log 4$, $t_{\min} = 1/16$, $\rho = 0$, $T = 1$, and every service duration to zero. With true index zero, corrupt the incumbent code to ordinal one during $[1/4, 1/2]$. At time 1, query $k = 1$ before clean overwrite gives $(2/3, 1/3)$ instead of $(1/3, 2/3)$, squared error $2/9$, although (13) is zero. All clean primitives may then occur at that same time. This refutes replacing the event-cut condition by a service-time comparison with the actual overwrite timestamp.

Proposition 5 (indistinguishable targets). Suppose two admissible executions give the decoder identical complete information at a common scored query, but require distinct clean targets a, b . Every deterministic answer has worst error at least $\|a - b\|_2/2$ on this pair. Worker agreement or sender identity does not remove this bound if it supplies no distinguishing

information.

Proof. The common decoder information yields one answer y . The triangle inequality gives $\|a - b\|_2 \leq \|a - y\|_2 + \|y - b\|_2$, proving the bound. If the admitted observation channel can make sampled indices $0, N$ indistinguishable, with $A = 0$ and query $k = N$, Lemma 1 gives separation $\sqrt{2} \tanh(hN/4)$. That specialization requires this particular pair to be indistinguishable; noninjectivity alone does not imply it. Distinct past histories with the same current target do not satisfy the hypothesis. $\square$

B. Exact quadratic transfer

For fixed $u, v \in \mathbb{R}^d$, define

$$\begin{aligned} f_y(x) &= \frac{1}{2}y\|x - u\|^2 + \frac{1}{2}(1 - y)\|x - v\|^2, \\ x_*(y) &= yu + (1 - y)v. \end{aligned} \quad (14)$$

Theorem 4. In model (2), the exact ideal minimax error for estimating $x_*(F_k(z))$ is

$$\frac{\|u - v\|}{2} \tanh(hD_*/4). \quad (15)$$

Proof. The gradient of f_y is $x - x_*(y)$, with Hessian I , so the minimizer is unique. For a fiber and query, extreme minimizers differ by $(F_k(U) - F_k(L))(u - v)$. Any estimate has worst error at least half their distance; the midpoint attains it for every compatible index because all targets lie on the segment. Lemmas 1–2 give (15). If $u = v$, every target coincides and error is zero. $\square$

This is minimizer estimation, not optimizer convergence. For (13) and Theorem 3, the achieved optimizer error is $\|u - v\|/\sqrt{2}$ times achieved row error. Rational-output rounding adds at most $\|u - v\|/(2 \cdot 2^9)$; anchor representation and arithmetic have their own errors. In $\mathbb{R}^3$ this is a segment projection, not formation or physical rendezvous.

Proposition 6. Let p, q be probability rows, $F_p = \sum_j p_j f_j$ be m -strongly convex with $m > 0$, and F_p, F_q have attained zero-gradient minimizers x_p, x_q . Assume the f_j are differentiable and $\|\nabla f_j(x_q)\| \leq G$. Then

$$\|x_p - x_q\| \leq \frac{G}{m} \|p - q\|_1. \quad (16)$$

Proof. Since $\nabla F_p(x_q) = \sum_j (p_j - q_j) \nabla f_j(x_q)$, strong monotonicity, $\nabla F_p(x_p) = \nabla F_q(x_q) = 0$, and Cauchy–Schwarz give

$$m\|x_p - x_q\|^2 \leq \nabla F_p(x_q) \cdot (x_q - x_p) \leq G\|p - q\|_1 \|x_p - x_q\|.$$

The zero-displacement case is immediate; otherwise divide. $\square$

The gradient bound is at the relevant minimizer, not a global bound. After actual support changes the labels, functions, attained minima and convexity hypotheses must be re-established. No nonconvex, constrained or algorithmic convergence claim follows.

VII. UNANNOUNCED CRASH ERASURES AT PERIOD ENDPOINTS

This section changes the availability model. It does not combine crash erasures with Section 6's code corruption into one fault theorem.

At each qT , new groups of $f+1$ observers and $f+1$ recipients arrive, for integer $f \geq 0$. All correct observers obtain the same fresh grid snapshot j_q from a correct acquisition channel. Encoding finishes within σ . At most f observers and at most f recipients crash at arbitrary unannounced cuts in the period, including before sending. Crashes stop computation and future sends; no false values or arbitrary replay occur. Thus at least one incarnation per role survives the period.

Each observer may send one fixed-length C -bit task word to each recipient, with $C \geq 0$ integer. Routes, phases, destinations, packet presence, lengths, timings, ordering and crashes cannot convey additional sampled-input information. Observers have only that period's snapshot as task input. Recipients obtain task-dependent information only through the permitted words, without additional observations, raw-observer state access or another data/query channel. At the scored endpoint every accessible surviving input-dependent register, buffer, packet or side channel is counted in $M \geq 0$ historical bits.

Every emitted word, including one whose sender/destination crashes, is delivered or terminally discarded within d of generation. Correct words to continuing recipients are delivered, including post-send source crashes. Each word has one-delivery semantics: no transport replication or retry. Copies to crashed/retired incarnations cannot be redirected. Terminal discard makes the copy unavailable to future decoders. Continuing recipients validate and clear the complete applicable delivery workload within a further p , not merely the first copy. Require

$$\sigma + d + p < T. \quad (17)$$

For unequal link bounds use their maximum d ; a reverse delay is irrelevant to this no-ack construction.

All queries are revealed at $(q+1)T$, after encoding, with any allowed k . They precede group retirement. Old retirement precedes the next group's birth/sample primitives at the same timestamp. The intermediate empty cut is inactive and unscored. Each incarnation lasts at most one period. The task signal is ρ -Lipschitz with sample/endpoint grid values. Set

$$A_T = \lfloor \rho T / h \rfloor.$$

Every pair with $|z-j| \leq A_T$ is realized by a linear signal over the period, remaining in $[0, hN]$; future periods can extend it by holding or another admissible interpolation.

Theorem 5 (one-hop period-endpoint crash-erasure law). Assume a correct common grid snapshot for $f+1$ observers; at most f crashes per observer and recipient role; and the ρ -Lipschitz task and endpoint query order above. Recipients obtain sampled-input information only through the permitted fixed-length C -bit words, and all accessible endpoint task distinctions count in M historical bits. Control, timing and crashes provide no extra input channel. Every emitted

copy, including a post-send-crash copy, has the stated one-delivery/lifetime contract, and the full applicable workload is cleared within p , with $\sigma + d + p < T$. Let $\mathcal{S}$ contain total deterministic strategies satisfying the stated information, resource and service contracts for every admitted task path, role-crash pattern and channel/service scenario. Each strategy answers every scored query. The class cannot exclude an admitted environmental scenario because its chosen strategy fails to implement a response. For each $S \in \mathcal{S}$, let $\Omega(S)$ contain only complete executions generated by S under the stated admissible task paths, crash patterns and channel/service schedules. The admitted execution set $\Omega(S)$ is nonempty and includes the execution induced by each such environmental scenario. Every emitted word and stored update must be that produced by S ; packet contents from another encoder are not independent environmental choices. The stated input-independence of control and scheduling remains required. An execution is infinite or reaches at least one scored period endpoint. Let $Q(\omega)$ be its represented scored periods and $R_q(\omega)$ the recipients surviving period q . Write $\hat{\Phi}_{i,q}^{S,\omega}(k)$ for that endpoint's answer. With $z_q(\omega) \in \{0, \dots, N\}$, define

$$\mathcal{E}_{\text{crash}} = \inf_{S \in \mathcal{S}} \sup_{\omega \in \Omega(S)} \sup_{q \in Q(\omega)} \max_{\substack{i \in R_q(\omega) \\ 0 \leq k \leq 2N}} \|\hat{\Phi}_{i,q}^{S,\omega}(k) - \Phi(z_q(\omega), k)\|_2.$$

Its exact ideal value is

$$\frac{\tanh(hD_*(P, A_T, 2^{\min(M,C)})/4)}{\sqrt{2}}. \quad (18)$$

The upper construction works for every finite or infinite number of group replacements.

Proof of the lower bound. Crash all but one observer and one recipient before any payload is sent, with identities and timing fixed independently of inputs. For each admitted S , this scenario with every allowed initial sampled index induces an execution in $\Omega(S)$; the strategy determines the word and answer in that execution. Removed observers leave no task distinction. The surviving recipient's full historical state is a deterministic function of the surviving observer's C -bit word and has at most both 2^C and 2^M states. Every allowed pair and query is realized by the specified signal. Theorem 1 and Corollary 1 give (18). Other executions cannot lower a worst-case error. Different observers may have different encoders; the selected surviving observer still has the same cardinality restriction.

Proof of the upper bound. Every observer uses the same public greedy encoder with $K_0 = 2^{\min(M,C)}$ and sends its ordinal in a padded C -bit word to every recipient. Correct simultaneous snapshots are equal, so valid words agree. Each recipient stores the first validated code, clears its receive buffer and treats equal later copies idempotently. At least one surviving observer sends to every surviving recipient. Every emitted copy is delivered or discarded by $qT + \sigma + d$, regardless of source crash. The whole-workload drain by $qT + \sigma + d + p < (q+1)T$ leaves precisely the retained ordinal as accessible task-dependent endpoint storage. No raw source state or additional data/query channel is available to

the endpoint answer. Its ordinal fits in M bits. The age displacement is bounded by A_T , and (6) attains (18). Repeating the construction independently in each period does not require a final stationary phase. The lower-bound subfamily remains realizable in a selected period. $\square$

Historical cut storage is not peak storage. There are at most $(f+1)^2$ words in flight, all in the current period, and up to $(f+1)C$ incoming task bits per recipient before clearing, plus processing workspace and metadata. The data bound is per edge, not an aggregate bandwidth optimum. Fresh acquisition consists of $f+1$ copies of one correct snapshot; these are not independent evidence of truth. Finite-horizon metadata grows logarithmically; bounded labels additionally require the stated packet-lifetime, no-replay and rebinding conditions. Rational finite-output evaluation adds (8)'s increment, without turning (18) into a finite-output equality. Actual computation must meet σ, p ; computability alone supplies no deadline.

The all-copy premise cannot be replaced by a bound only for survivor pairs. For $f = 1$, $P = 2$, $M = C = 1$, $A_T = 0$, $T = 1$, $\sigma = d = 1/10$, $p = 1/5$, let one observer send at $1/10$, crash at $3/20$, and have its copy arrive at $99/100$, with validation ending $109/100$. A survivor-pair-only premise does not exclude that copy, leaving task-dependent endpoint buffer state. It adds no independent index information but invalidates the claimed physical code-only representation. The all-emitted-copy bound excludes it because $99/100 - 1/10 = 89/100 > d$.

With only f observers, the permitted crash-before-send schedule can erase all observation carriers; hence $f+1$ observers are necessary for this direct one-hop cohort, not for every possible sensing/storage architecture. Scoring before delivery, permitting input-dependent timing, granting recipients another task channel, or replacing the current snapshot by an erased historical target changes the theorem's model.

VIII. RETAINED MASS AND NORMALIZED FAMILIES

Before deletion, the mass retained on (b, c) is

$$s(z, k) = \frac{t_z + v_k}{1 + v_k}.$$

Lemma 5. Its minimum over the declared task/query family is

$$s_{\min} = \frac{2t_{\min}}{1 + t_{\min}}. \quad (19)$$

Proof. The expression increases in t ; its derivative in v is $(1-t)/(1+v)^2 > 0$. Both minima occur at zero indices. $\square$

The conditional row in (1) cancels $t_{\min}$. Thus equal retained mass need not determine conditional-query error. For $P = 2$, $K = 1$, $A = 0$, $t_{\min} = 1/16$, the same $s_{\min} = 2/17$ gives errors $1/(3\sqrt{2})$ at $h = \log 4$ and $1/(2\sqrt{2})$ at $h = \log 9$. Both families remain positive. A universal error formula depending only on $s_{\min}, P, A, K$ is therefore false.

Impose instead the additional normalization $t_{\max} = t_{\min}e^{hN} = 1/2$. Writing $s = s_{\min}$,

$$t_{\min} = \frac{s}{2-s}, \quad h = \frac{1}{N} \log \frac{2-s}{2s}, \quad 0 < s < 2/3.$$

Theorem 6. In this normalized family the exact single-cut ideal error is

$$\mathcal{E}_s = \frac{1}{\sqrt{2}} \tanh \left(\frac{D_*(P, A, K)}{4N} \log \frac{2-s}{2s} \right). \quad (20)$$

For a log-signal age bound δ , use

$$A_s = \left\lfloor \frac{\rho\delta N}{\log((2-s)/(2s))} \right\rfloor$$

provided every grid pair within that bound is admissible.

Proof. Solve (19) for $t_{\min}$ and substitute the normalization to obtain h . The domain ensures $h > 0$ and all original-row probabilities below one. Apply Theorem 1. Equation (10) gives A_s . Exactness in a timing interpretation needs full pair reachability; Section 7's linear paths provide it at period endpoints. A smaller admissible family can have smaller minimax error. $\square$

At $A = 0, K \geq P$, the error is zero even for arbitrarily small s . For fixed $P > K, A = 0$, it tends to $1/\sqrt{2}$ as $s \downarrow 0$. Holding physical $\rho\delta$ fixed requires recomputing A_s , rather than pretending it stays constant. Realizing small original masses numerically has a separately counted parameter cost. This is a normalized relative-query law, not a general inverse-mass sensitivity theorem.

IX. ASSUMPTIONS, VERIFICATION AND SCOPE

The proof package contains deterministic partition enumeration, rational arithmetic and ordered event fixtures. Their domains and tolerances are explicit. Exhaustive partitions for $2 \leq P \leq 8$ test arbitrary fibers rather than only interval encoders. Rational fixtures check admission/deletion, endpoint evaluation and midpoint rounding. The two-slot event traces check listed readiness, age, retirement and post-attack query instances. The main crash checker checks first-copy availability, equal codes and finite query errors; it does not check aggregate receive-buffer drainage. A separate one-copy trace checks its stated clearing events. The correction fixtures distinguish early-code availability from all-copy lifetime and full-workload clearing. The assertion-level claim map records these different scopes; none is a complete protocol implementation. These finite computations corroborate their listed instances and preserve regressions; they do not prove unrestricted theorem statements, hardware operation or empirical benefits.

Additional local replays checked every set partition for $2 \leq P \leq 14$, including 190,899,322 partitions at $P = 14$, against (4) for every allowed A, K . An exact-integer Python counterpart checked 4,194,152 nonempty-fiber/uncertainty pairs through $P = 17$ in the subfamily $r = 2$. It evaluated every query scale for each distinct expanded range and reused those verified range results. These are finite checks of the partition and query formulas, not a proof for every P or real h . A separately supplied timing script reproduces arithmetic bounds and static role-survivor counts; it does not execute packet buffers or corrupted-register transitions, so its output is not attributed as a protocol-state recovery test.

The total peak cost includes historical codes, raw inputs, public parameter representations, arithmetic, payload buffers,

TABLE II
RESOURCES AND RESTRICTIONS OF THE DECLARED RESULTS.

Result	Required resource/restriction	Consequence when removed
Exact cut law	Encoder precedes query; full admissible index/query family; all side channels counted	Query-dependent encoding or extra information changes the problem
Finite arithmetic	Finitely specified rational geometric parameters; declared workspace/output precision	Arbitrary exact reals need not be computable; ideal output is not finite output
Handoff	Correct binding, bounded service, retain-until-ack	Premature retirement can erase the sole code
Bounded-age replacement	Fresh correct observation each period; explicit timers/routes; (11)	Forwarding an unchanged old code has growing source age
Code recovery	Trusted acquisition/control; finite attack; bounded old-copy lifetime	Indistinguishable executions with distinct targets preclude exact recovery; arbitrary Byzantine recovery is unproved
Crash endpoint law	Role-wise crash budgets; common correct snapshot; all-copy lifetime and full drain; (17)	Late copies invalidate the endpoint bit representation; complete erasure removes information
Optimizer transfer	Attained minima, specified objective regularity and representations	No algorithmic or physical convergence follows
Mass parameterization	Stated normalization and pair reachability	Equal mass alone does not determine error

control envelopes and identifiers. Neither a single-cut cardinality bound nor a repeated common code equals a total-network storage optimum. Acquisition and synchronization are external resources even without a permanent worker. A timestamp deadline differs from a causal event cut, and a service-time target differs from a later received response; no unproved response-age bridge is assumed.

Open extensions include general temporal graph dissemination, cumulative historical survivor functionals, independently verifiable source/control truth, aggregate interactive resource bounds, response-time tracking and physical implementation. These are different research obligations and do not support the present claims.

X. CONCLUSION

The exact cut law follows from a binary conditional diameter and an expanded-fiber covering problem. The attaining encoder precedes the query and its decoder is query dependent. Finite arithmetic, periodic replacement, code recovery and crash erasures require separate resource and timing contracts. The two-slot and one-hop constructions preserve the stated conditional accuracy with fresh correct acquisition and bounded delivery/processing, without a permanently surviving worker. This does not give fixed total resources for unrestricted indefinite operation: bounded identifiers additionally require the declared lifetime, no-replay and rebinding contracts. Quadratic optimizer estimation and retained-mass normalization transfer the law without establishing algorithmic convergence or a general mass-only principle.

ACKNOWLEDGMENT

AI agents assisted drafting and refinement throughout Sections I–X, Appendix A, the bibliography and this acknowledgment, including mathematical formulation, proof construction and adversarial examination. It also assisted development of the accompanying verification code. It is not an author. No simulation result is presented. Human final review and acceptance remain pending; the human author is responsible for the manuscript. The previously displayed future acceptance paragraph is preserved in the revision archive and is not asserted as a completed event in this preprint.

APPENDIX A ARITHMETIC CONSTRUCTION

Proof. For $0 \leq m \leq 2N$, $a^m, b^m < 2^{2NB}$, and their sum is less than $2^{\ell-1}$. A bit machine can compute each power by repeated grade-school multiplication, retaining a constant number of binary operands, product and carry/scratch buffers of $O(\ell)$ bits plus an $O(\log P)$ -bit exponent counter. A product accumulator is cleared before reuse. Input copies and arithmetic scratch remain included in peak workspace. No unit-cost unbounded integer instruction or uncounted heap is assumed. If the endpoints are $n_0/s_0, n_1/s_1$, their midpoint is

$$m = \frac{n_0 s_1 + n_1 s_0}{2 s_0 s_1} = \frac{n}{s}, \quad 0 < n < s < 2^{2\ell+1}.$$

Compute $q = \lfloor 2^g n / s \rfloor$ and return first coordinate $(2q + 1)/2^{g+1}$. It is the midpoint of the containing cell, with scalar error at most 2^{-g-1} . The shifted numerator $2^g n$, divisor, quotient, remainder and division scratch fit in a constant number of $O(\ell + g)$ -bit buffers. Bitwise long division and schoolbook multiplication therefore use $O(\ell + g)$ peak work cells; writing the result does not hide these working buffers in the output stream. The public greedy partition can be reconstructed to its stored ordinal using $O(\log P)$ -bit counters; at most P fibers are used, so a K -entry table is unnecessary. The canonical ordinal needs at most $\lceil \log_2 P \rceil$ bits; transport padding/validation buffers are outside this decoder model and remain protocol costs. This proves peak bit space, not a runtime or deadline bound. $\square$

REFERENCES

- [1] S. N. Cohen and E. Fausti, “Hyperbolic contractivity and the Hilbert metric on probability measures”, arXiv:2309.02413v2, 2024. [Online]. Available: <https://arxiv.org/abs/2309.02413v2>
- [2] V. Doshi, D. Shah, and M. Médard, “Source coding with distortion through graph coloring”, author-hosted manuscript, n.d. Accessed: Oct. 3, 2026. [Online]. Available: <https://www.mit.edu/~medard/page2/papers08/doshi.pdf>
- [3] C. Monnoyer de Galland and J. M. Hendrickx, “Fundamental performance limitations for average consensus in open multi-agent systems”, IEEE Trans. Autom. Control, vol. 68, no. 2, pp. 646–659, 2023, doi: 10.1109/TAC.2022.3141685. Inspected preprint: arXiv:2004.06533v4. [Online]. Available: <https://arxiv.org/abs/2004.06533v4>
- [4] H.-L. Chan, T.-W. Lam, L.-K. Lee, and H.-F. Ting, “Continuous monitoring of distributed data streams over a time-based sliding window”, in Proc. STACS, 2010. Inspected preprint: arXiv:0912.4569v2. [Online]. Available: <https://arxiv.org/abs/0912.4569v2>

- [5] S. Gilbert, N. A. Lynch, and A. A. Shvartsman, “RAMBO: A robust, re-configurable atomic memory service for dynamic networks”, *Distributed Computing*, vol. 23, no. 4, pp. 225–272, 2010, doi: 10.1007/s00446-010-0117-1. [Online]. Available: <https://groups.csail.mit.edu/tds/papers/Gilbert/rambo-journal3.pdf>.
- [6] J. Hu, K. H. Johansson, and A. I. Rikos, “Distributed coordination algorithms with efficient communication for open multi-agent systems with dynamic communication links and processing delays”, *arXiv:2603.08038v1*, 2026. [Online]. Available: <https://arxiv.org/abs/2603.08038v1>.
- [7] S. Naderi, M. Blondin, and S. Roy, “Consensus in multi-agent systems with uniform and nonuniform communication delays”, *arXiv:2603.16523v1*, 2026. [Online]. Available: <https://arxiv.org/abs/2603.16523v1>.
- [8] B. Charron-Bost and P. Lambein-Monette, “Randomization and quantization for average consensus”, in *Proc. IEEE Conf. Decision Control*, 2018, pp. 3716–3721, doi: 10.1109/CDC.2018.8619817. Inspected preprint: *arXiv:1804.10919v1*. [Online]. Available: <https://arxiv.org/abs/1804.10919v1>.
- [9] E. Tuncel, P. Koulgi, S. L. Regunathan, and K. Rose, “Zero-error source coding with maximum distortion criterion”, in *Proc. Data Compression Conf.*, 2002, pp. 92–101, doi: 10.1109/DCC.2002.999947. [Online]. Available: https://web.ece.ucsb.edu/scl/pubs/pubs_D/d02_1.pdf.
- [10] D. Muresan and M. Effros, “Quantization as histogram segmentation: Optimal scalar quantizer design in network systems”, *IEEE Trans. Inf. Theory*, vol. 54, no. 1, pp. 344–366, 2008, doi: 10.1109/TIT.2007.911170. [Online]. Available: <https://authors.library.caltech.edu/records/vn1nc-rtv93>.
- [11] Y. Kaspi and N. Merhav, “Zero-delay and causal single-user and multi-user lossy source coding with decoder side information”, *IEEE Trans. Inf. Theory*, vol. 60, no. 11, pp. 6931–6942, 2014, doi: 10.1109/TIT.2014.2358238. Inspected preprint: *arXiv:1301.0079*. [Online]. Available: <https://arxiv.org/abs/1301.0079>.