Six Passes Are Enough: Security Proofs for NIST AES Key Wrap

Yuwen Zhou
University of British Columbia Okanagan
zhouyw07@student.ubc.ca

Abstract

AES Key Wrap protects keys entering and leaving hardware security modules and keys exchanged through S/MIME, JOSE and PKCS#11. NIST standardized it as KW and KWP in SP 800-38F and is now revising that document. In twenty-five years the algorithm has received no security proof, and apart from a remark on very long messages, nothing published explains its six passes of AES. Its design breaks the assumptions of existing proofs in three ways. One AES key serves every step, whereas proofs for generalized Feistel networks assume independent round functions; the chaining value has only 64 bits, so collisions on it occur long before the birthday bound of AES and a proof cannot exclude them; and unwrapping runs the algorithm backward on adversarial ciphertexts, which the usual arguments handle only with independently keyed halves. We prove that if AES is a strong pseudorandom permutation (PRP), the six-pass wrapping function at the core of the standard is also a strong PRP, with main term $88\sigma^2/2^{128}$ for σ semiblocks of data. KW and KWP are deterministic authenticated encryption schemes, with an additional term for guessing the 64-bit check value. Our main step is to sample the intermediate values of each message from both ends, forward from the input and backward from the output, and to join the two sides in the middle, which treats forward and inverse queries alike. Every bad event then involves at most two messages and needs two independent 64-bit collisions, which we count rather than exclude. The count needs only two passes on each side of the join and one for the join, so five passes give a bound of the same form and the sixth is to spare. Two passes, by contrast, can be distinguished from random with 2^{32} short queries. All bounds extend to many independent keys without a loss in the number of keys.

Declaration of generative AI use (IACR category: Substantial AI contribution). The proof presented in this paper was developed through an interactive dialogue with Claude Opus 5.5 (Anthropic) and ChatGPT-6 (OpenAI). The author used them to explore multiple possible proof strategies and directions, and the AI generated the initial draft under the author's guidance. The author then critically examined the generated proof, verified each step for correctness and rigor, and revised it where necessary. The author takes full responsibility for the final content.

1. Introduction

Key wrapping is the encryption of one cryptographic key under another. It is how keys enter and leave hardware security modules, how customers import keys into cloud key-management services, and how the content key of an encrypted message reaches its recipients. The dominant algorithm for the job is AES Key Wrap. NIST posted it in 2001 [1], the IETF published it as RFC 3394 and, with padding, as RFC 5649 [2], [3], and NIST standardized the two forms as KW

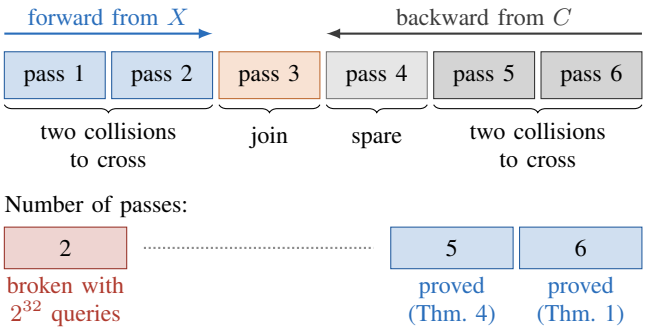

Figure 1. Two-sided sampling and the number of passes. The proof samples each message forward from its input X and backward from its output C, and joins the two sides in pass 3. One 64-bit collision lets a pair of messages cross one pass, so reaching the join costs two collisions.

and KWP in SP 800-38F [4]. They appear in CMS and S/MIME [5], JOSE [6], PKCS#11 [7], KMIP [8], the Web Cryptography API [9] and key import in AWS KMS [10], and KWP is required by the draft S/MIME profile for the CNSA 2.0 suite [11]. The keys that these systems wrap protect everything else, so a weakness in this one algorithm would be a weakness everywhere. NIST is now revising SP 800-38F [12].

After twenty-five years of use, no one has proved these algorithms secure. Rogaway and Shrimpton, who gave key wrapping its formal goal of deterministic authenticated encryption (DAE), examined this design in 2006. They wrote that it had no proof and was "so complex that providing one would be difficult," and they conjectured that the Feistel network at its core is a strong pseudorandom permutation [13]. SP 800-38F repeats the remark that no proof is known [4], and so does the most recent analysis, written for the NIST revision in October 2026 [14]. That analysis points to a second gap. The design runs six passes of AES over the data, twelve AES calls for every 128 bits wrapped, and apart from one observation about very long messages nothing published says why six; whether fewer would do has been a matter of heuristics [14].

The difficulty lies in the core of the standard, a function W that runs a shift register through six passes of AES (Figure 2). Each step encrypts a 64-bit chaining value together with one 64-bit word of the register, and the features that make W simple to build make it hard to analyze. AES under one key serves every step of every message, whereas proofs for generalized Feistel networks assume independent round functions [16], [17]. The chaining value has only 64 bits, so two computations collide on it long before the birthday bound of AES. Proofs for enciphering modes such as CMC and EME exclude such collisions outright [18], [19], but in W they happen and must be survived. And unwrapping runs W backward on a ciphertext that the adversary chose, while the usual way to lift one-directional security to two directions needs independently keyed halves [16], which W does not have.

1.1. Our results

The standard is secure. We prove W, KW and KWP secure. The model replaces AES under the secret key by a uniform random permutation, and returning to AES costs its strong-PRP advantage. For an adversary that submits σ semiblocks (64-bit words) in forward and inverse queries of its choice, W is a strong pseudorandom permutation with advantage at most

$$\frac{88\,\sigma^2}{2^{128}}$$

plus a term of higher order (Theorem 1). KW is a DAE scheme at an extra cost of $q_v/2^{63}$ for q_v unwrap attempts (Theorem 2), and KWP, including its one-block path, is a DAE scheme with main terms of order $\sigma^2/2^{128} + q_v/2^{64}$ (Theorem 3). A key that has wrapped 2^{48} keys of 256 bits is still within a bound of about 2^{-21}.

Five passes suffice, and two do not. With five passes instead of six, W meets the same theorem with the constant 76 (Theorem 4). With two passes it fails at the 64-bit level: 2^{32} chosen inputs of four semiblocks distinguish it from random with advantage above 0.39 (Proposition 5), a short-message case of a collision attack known for long messages [15]. So six passes are enough, and the sixth is to spare (Figure 1).

Many keys. A deployment has many key-encryption keys, each wrapping a modest amount of data. For any number of independent keys our bounds add up key by key (Corollary 6). The main term for W becomes $88\,\sigma\sigma_{\max}/2^{128}$, where $\sigma_{\max}$ is the largest amount of data under one key, and no factor for the number of keys appears.

1.2. Our technique: two-sided sampling

A collision on the 64-bit chaining value is cheap, and the proof does not try to exclude it. What a collision buys is small: two messages that collide run together for at most one pass and then part again. To do damage, a pair has to stay together across two passes, which takes two independent collisions and has probability 2^{-128}. The proof counts the ways in which this can happen.

The count becomes rigorous through one observation. The probability of a transcript, in the real and in the ideal world, depends only on its set of input–output pairs, and not on which side the adversary queried [20]. So the analysis may sample the permutation in whatever order it likes, and we sample each message from both ends: the first two passes forward from the plaintext, the last passes backward from the ciphertext, joined in the third pass. Forward and inverse queries become the same object, the steps that must take constrained values sit between internal random values that the adversary never sees, and every bad event is an event about at most two messages. The layout needs two sampled passes on each side of the join and one pass to join them, which is five; the standard's sixth pass sits next to the join and makes the count on one side easier. KWP needs a further argument, because its one-block path exposes the bare permutation under the same key, and we prove its security directly.

Section 2 recalls the algorithms and goals, Section 3 states the theorems, and Section 4 explains the proof on small examples. Sections 5–7 give the proofs, with the case analyses in the appendices and the complete derivations in Appendices E–I.

2. AES Key Wrap and Its Security Goals

The wrapping function. A semiblock is a 64-bit string. Let π be a permutation of $\{0,1\}^{128}$; in the standard, π is AES under the key-encryption key. For a number of passes r, the function W_r takes $n = L+1 \geq 3$ semiblocks, written $X = (X_A; X_1, \ldots, X_L)$, loads them into a chaining value A and a register $(R_1, \ldots, R_L)$, and runs rL steps (Figure 2, left). Step t encrypts the chaining value together with the first register word, $(\gamma, \delta) = \pi(A \,\|\, R_1)$. The left half γ, xored with the 64-bit encoding $\langle t \rangle$ of the step number, becomes the new chaining value, and the right half δ goes to the back of the register. One pass over the register takes L steps, and the output is the final state $C = (C_A; C_1, \ldots, C_L)$. The standard fixes $r = 6$, and we write $W = W_6$; Rogaway and

procedure $W_r(X_A; X_1, \ldots, X_L)$
  $A \leftarrow X_A$; $(R_1, \ldots, R_L) \leftarrow (X_1, \ldots, X_L)$
  for $t = 1$ **to** rL **do**
    $(\gamma, \delta) \leftarrow \pi(A \,\|\, R_1)$; $A \leftarrow \gamma \oplus \langle t \rangle$
    $(R_1, \ldots, R_L) \leftarrow (R_2, \ldots, R_L, \delta)$
  return $(A; R_1, \ldots, R_L)$

procedure KW.Wrap(P) P: $L \geq 2$ semiblocks
  return $W_6(\mathrm{ICV1}; P)$

procedure KW.Unwrap(C)
  $(A; P) \leftarrow W_6^{-1}(C)$
  if $A = \mathrm{ICV1}$ **then return** P **else return** $\perp$

procedure KWP.Wrap(P) P: m bytes
  $P' \leftarrow P$ zero-padded to $n' = \lceil m/8 \rceil$ semiblocks
  if $n' = 1$ **then return** $\pi(\Lambda_m \,\|\, P')$
  return $W_6(\Lambda_m; P')$

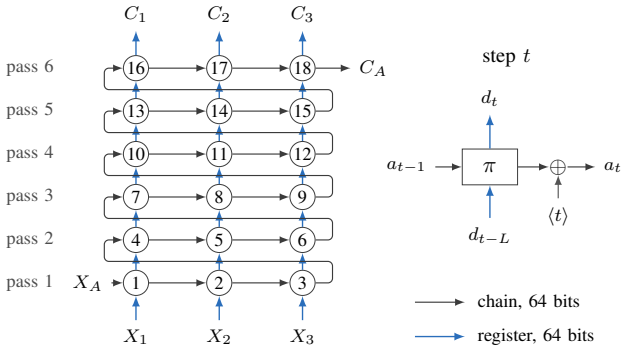

Figure 2. AES Key Wrap. Left: the wrapping function W_r with r passes and the two schemes of SP 800-38F, which use $W = W_6$; π is AES under the key-encryption key. Right: the computation of W unrolled for pass length $L = 3$. Step t reads the chaining value a_{t-1} written by step $t-1$ and the register word d_{t-L} written one pass earlier, and writes a_t and d_t.

Shrimpton call it a six-round Feistel network [13]. We call L the *pass length*, and pass k is the set of steps $(k-1)L+1, \ldots, kL$. The inverse W_r^{-1} runs the steps backward with π^{-1}.

The unrolled view. It helps to unroll the register (Figure 2, right). Write a_t for the value of A after step t and d_t for the word written to the register at step t. Then step t has

$$\text{input } (a_{t-1},\, d_{t-L}), \qquad \text{output } (a_t \oplus \langle t \rangle,\, d_t).$$

So every step reads two values and writes two values. The *chain* carries a value from step t to step $t+1$, and the *register* carries a value from step t to step $t+L$; steps $t-1$ and $t-L$ are the *suppliers* of step t. The two never mix, so every equation in this paper is either a chain equation or a register equation. The input fixes the values at the bottom of the computation, $a_0 = X_A$ and $d_{j-L} = X_j$ for $1 \leq j \leq L$, and the output fixes those at the top, $a_{rL} = C_A$ and $d_{(r-1)L+j} = C_j$.

KW and KWP. The standard builds two schemes from W. KW wraps a plaintext P of $L \geq 2$ semiblocks as $W(\mathrm{ICV1} \,\|\, P)$ with $\mathrm{ICV1} = \texttt{A6A6A6A6A6A6A6A6}$, and unwrapping applies W^{-1} and rejects unless the first semiblock is ICV1. KWP wraps a plaintext of m bytes, $1 \leq m < 2^{32}$. It sets $\Lambda_m = \texttt{A65959A6} \,\|\, [m]_{32}$ and pads P with zero bytes to $n' = \lceil m/8 \rceil$ semiblocks, giving P'. If $n' = 1$ the ciphertext is $\pi(\Lambda_m \,\|\, P')$, a single call of AES, and if $n' \geq 2$ it is $W(\Lambda_m \,\|\, P')$. Unwrapping inverts accordingly and accepts only if the first semiblock is Λ_m with $8(n'-1) < m \leq 8n'$ and the padding bytes are zero. The one-block path shares π with W. A long wrap has $m \geq 9$, so its first semiblock is never that of a one-block wrap.

Security goals. We analyze these schemes in the ideal-permutation model, replacing AES under the secret key by a uniform permutation π. An adversary against W_r makes q forward or inverse queries, of any lengths $n_i \geq 3$, in any order, and tries to distinguish (W_r, W_r^{-1}) from independent uniform permutations P_n of $(\{0,1\}^{64})^n$, one for each length n. Its advantage is written $\mathrm{Adv}^{\pm\mathrm{prp}}_{W_r}$. We write $\sigma = \sum_i n_i$ for the total number of semiblocks and $L_{\max}$ for the largest pass length. An adversary against a wrapping scheme makes q_e wrap queries and q_v unwrap queries, and tries to distinguish the real oracles from an oracle pair that answers each new wrap query with a uniform string of the right length and rejects every unwrap [13]. It may not unwrap an earlier

TABLE 1. ADVANTAGE BOUNDS FOR WRAPPING 256-BIT KEYS ($L = 4$), AS $\log_2$. THE UNWRAP COUNT IS q_v.

wrapped keys q	q_v	KW (Thm. 2)	KWP (Thm. 3)
2^{32}	2^{20}	-43.0	-37.8
2^{40}	2^{30}	-32.9	-27.6
2^{48}	2^{30}	-20.9	-14.9

wrap answer, and its advantage is written $\mathrm{Adv}^{\mathrm{dae}}$. We count every semiblock that goes through W or π, including the leading semiblock added by the scheme, and all theorems assume $\sigma \le 2^{60}$.

3. Results

3.1. The standard

Theorem 1. *For every adversary against W that makes q forward or inverse queries in any mix, each of at least 3 semiblocks, with $\sigma \le 2^{60}$ semiblocks in total,*

$$\mathrm{Adv}_{W}^{\pm\mathrm{prp}} \;\le\; \frac{88\,\sigma^2}{2^{128}} + \frac{3.1 \cdot 10^6\, q\,\sigma\, L_{\max}^5}{2^{192}}.$$

In words, an adversary who may run W in both directions cannot tell it from a family of random permutations, except with probability about $\sigma^2/2^{128}$ times a small constant.

Theorem 2. *If the key is used only through KW, then for every adversary making q_e wrap and q_v unwrap queries, with $q = q_e + q_v$ and $\sigma \le 2^{60}$,*

$$\mathrm{Adv}_{\mathrm{KW}}^{\mathrm{dae}} \;\le\; \frac{88\,\sigma^2}{2^{128}} + \frac{3.1 \cdot 10^6\, q\,\sigma\, L_{\max}^5}{2^{192}} + \frac{q_v}{2^{63}} + \frac{q_e^2}{2^{193}}.$$

The term $q_v/2^{63}$ is the chance of guessing the 64-bit check value.

Theorem 3. *If the key is used only through KWP, then for every adversary making at most q wrap queries of more than 8 bytes, any number of one-block wraps, and at most q_v unwrap queries, with $\sigma \le 2^{60}$,*

$$\mathrm{Adv}_{\mathrm{KWP}}^{\mathrm{dae}} \le \frac{5758\,\sigma^2}{2^{128}} + \frac{35868\,L_{\max}\,\sigma}{2^{128}} + \frac{73\,q_v}{2^{64}} + \frac{7721568\,(q+q_v)\,\sigma\,L_{\max}^5}{2^{192}}.$$

One-block queries enter Theorem 3 only through σ. All resources are deterministic caps, and all three theorems hold for adaptive adversaries.

Reading the bounds. The term in $L_{\max}^5$ is of higher order. It is below the main term whenever $L_{\max} \le 2^{10}$, a range that contains every symmetric key, at most 8 semiblocks, and common private keys such as RSA-4096 and ML-DSA. The bound of Theorem 1 is meaningful for one message of up to about 2^{28} semiblocks, and at $\sigma = 2^{60}$ for messages of up to about $2^{12.6}$ semiblocks. Table 1 evaluates the bounds for wrapping 256-bit keys. The KW bound stays below 1 up to $\sigma = 2^{60}$, the limit of the theorem, and the KWP bound up to about $\sigma = 2^{57}$.

Published attacks. Table 2 sets the bounds next to the published attacks. Each attack lies within them, and the one-block path of KWP shows that a term of shape $\sigma^2/2^{128}$ is necessary for KWP.

TABLE 2. PUBLISHED ATTACKS AND OUR BOUNDS. HERE m IS THE PLAINTEXT LENGTH IN SEMIBLOCKS.

Attack	Relation to our bounds
Kelsey forgery on KW, advantage $2^{-64} + (m/2^{64})^6/6!$ [4], [14]	Below the terms $q_v/2^{63}$ and $88\sigma^2/2^{128}$ since $\sigma \geq 2m$
Encryption-only distinguisher on W, advantage $\frac{q^2}{2}(m/2^{64})^6/6!$ [14]	Below $88q^2(m/2^{64})^2$; at its best parameters the data is $2^{88.75}$ semiblocks, beyond $\sigma = 2^{60}$
KWP one-block path versus random strings, advantage $1 - e^{-q^2/2^{129}}$ [14]	The $\sigma^2/2^{128}$ shape is necessary for KWP
KWP forgery, $2^{-63.994}$ per attempt [14]	Within $73\,q_v/2^{64}$

3.2. The number of passes

Theorem 4. *For every adversary against W_5 that makes q forward or inverse queries in any mix, each of at least 3 semiblocks, with $\sigma \leq 2^{60}$ semiblocks in total,*

$$\mathrm{Adv}^{\pm\mathrm{prp}}_{W_5} \;\leq\; \frac{76\,\sigma^2}{2^{128}} + \frac{1.01 \cdot 10^6\, q\,\sigma\, L^5_{\max}}{2^{192}}.$$

In words, removing one pass from the standard leaves a strong pseudorandom permutation with a bound of the same form.

Proposition 5. *Let $B = 2^{64}$. For every $4 \leq q \leq B$ there is an adversary against the two-pass function W_2 that makes q forward queries of four semiblocks and has*

$$\mathrm{Adv}^{\pm\mathrm{prp}}_{W_2} \;\geq\; 1 - \exp\Big(-\frac{(B-3)(q-2)(q-3)}{2B^2}\Big) - \frac{\binom{q}{2}}{B^2+1}.$$

For $q = 2^{32}$ the right side exceeds 0.39.

In words, two passes are distinguished by a chosen-plaintext attack at the birthday bound of the 64-bit chaining value, with messages as short as a 192-bit key and its check value. With the first semiblock set to ICV1, the same queries distinguish two-pass KW from the ideal DAE scheme, with the last term replaced by $\binom{q}{2}/B^2$. The attack is the two-pass, short-message case of a chosen-plaintext distinguisher that the Wycheproof documentation describes for long messages and attributes to an idea of Ferguson [15].

3.3. Many keys

In the multi-user setting [21] there are u users with independent keys, the adversary addresses each query to a user of its choice, and all users are real or all are ideal. User i has its own deterministic caps: q_i queries, $\sigma_i \leq 2^{60}$ semiblocks, largest pass length L_i, and for the schemes $q_{e,i}$ wraps and $q_{v,i}$ unwraps. Write $\sigma = \sum_i \sigma_i$ and $\sigma_{\max} = \max_i \sigma_i$.

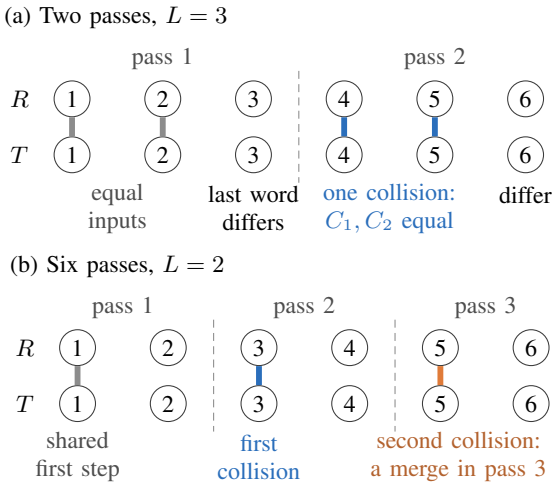

Figure 3. One collision buys one pass. Each circle is a step, and a bar joins two steps with equal inputs. (a) With two passes, one 64-bit collision after step 3 makes steps 4 and 5 equal, and the ciphertext shows it. (b) With more passes the two messages part again, and a pair with a common first step needs two collisions before its inputs meet in pass 3, an event of probability about 2^{-128}.

Corollary 6. *Let ϵ_i^W, ϵ_i^{KW} and $\epsilon_i^{\mathrm{KWP}}$ be the bounds of Theorems 1, 2 and 3 evaluated at the caps of user i. Then*

$$\mathrm{Adv}_W^{\text{mu-}\pm\text{prp}} \le \sum_i \epsilon_i^W, \qquad \mathrm{Adv}_{\mathrm{KW}}^{\text{mu-dae}} \le \sum_i \epsilon_i^{\mathrm{KW}},$$
$$\mathrm{Adv}_{\mathrm{KWP}}^{\text{mu-dae}} \le \sum_i \epsilon_i^{\mathrm{KWP}}.$$

In particular the main terms are at most $88\,\sigma\sigma_{\max}/2^{128}$ for W and KW and $5758\,\sigma\sigma_{\max}/2^{128}$ for KWP. If each user holds a uniform k-bit AES key and AES is modeled as an ideal cipher to which the adversary makes p queries, each bound grows by $up/2^k + \binom{u}{2}/2^k$.

In words, what limits security is the data under each key and not the data in the system. As an example, take 2^{30} key-encryption keys that each wrap 2^{20} keys of 256 bits. The main term of the bound is $2^{-46.9}$, against $2^{-16.9}$ if one key wrapped everything.

4. The Proof on Small Examples

One collision buys one pass. Take two passes, $L = 3$, and two inputs R and T that differ only in their last word (Figure 3a). Steps 1 and 2 are the same in both computations. Step 3 reads the differing word, so its two outputs differ; suppose that their left halves collide, which is one 64-bit equation. Then steps 4 and 5 read equal chaining values and the register words written by steps 1 and 2, so their inputs are equal, and the words they write are the ciphertext words C_1 and C_2. With only two passes one collision is visible in the output. This is Proposition 5.

With more passes the same collision does no harm, because the agreement ends: step 6 reads the two different words written at step 3, and to agree in the next pass the computations must collide again. Figure 3b follows a pair through the standard's W in the smallest case, $L = 2$, with $X^R = (A; x_1, x_2)$ and $X^T = (A; x_1, x_2')$. Step 1 has the same input in both computations, and step 2 differs because the second register word differs. At step 3 the register halves are equal,

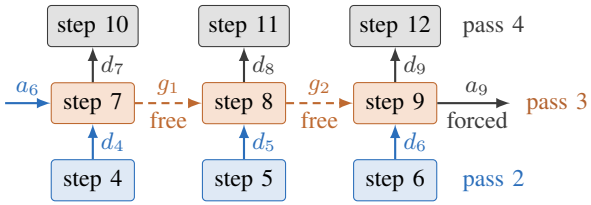

Figure 4. The joining pass for $L = 3$. A joining step knows its input and the right half of its output: blue values were sampled on the forward side and gray values on the backward side. The left halves g_1, g_2 are free and drawn uniformly, and the last output is forced. The process fails only if a joining input is already defined or a joining output is already used.

since both are d_1, and the chain halves are equal only if the two step-2 outputs agree in their left halves: one 64-bit equation. Suppose this happens. Step 4 differs again, but at step 5 the register halves are equal, since both are d_3, and the chain halves are equal if the two step-4 outputs agree in their left halves: a second 64-bit equation. Now the two messages have merged inside pass 3, which is the kind of event a proof must rule out, and it has cost two independent equations.

A collision on the 64-bit chain buys only one pass of agreement, so a pair of messages that reaches the middle of the computation has paid twice, for a total of 2^{-128}, and there are at most σ^2 ways, up to a constant, to place such a pair of equations.

Why the order of sampling matters. The H-coefficient technique fixes a transcript and lower-bounds its real probability by choosing π step by step, and the natural order processes each query in its own direction. Forward queries are then constrained at the end of the last pass, where the output must match C, and inverse queries at the end of pass 1. Those constrained points have halves equal to constants that the adversary picked, so a later inverse query can name them with one equation or with none. The event that a computation visits a constrained point then has probability 2^{-64} or worse, and the count above is lost.

Two-sided sampling. The way out is that the transcript probability does not depend on the direction of each query: an inverse query with answer X imposes the same condition $W(X) = C$ as a forward query. So direction is a choice of the proof. We sample each message forward through passes 1–2 starting from X, and backward from C down to pass 4, and we join the two sides in pass 3 (Figures 1 and 4). This two-sided sampling is the idea behind the Luby–Rackoff and Naor–Reingold proofs [22], [23], used here as a framework. Every value at a joining step is now internal, either a half of a fresh draw from one of the two sides or a fresh independent value, and no transcript constant touches a joining step. Three things follow. A joining step cannot be reached for free or with one equation. A computation that lands on another message's ordinary point of π reads exactly what π would give, so its path does not change; only the joining steps can change a path, and therefore the first bad event is an event about the paths of two messages alone. And forward and inverse queries are treated identically.

Counting. For a pair of messages we replay the two computations and write every unexpected equality as a linear equation over 64-bit values. In the language of structure graphs for CBC-MAC [24], a collision implied by earlier ones is *induced*, and a surprising one is an *accident*. Each independent equation costs a factor of about 2^{-64}. No bad event happens with fewer than two independent equations, and the bad events with exactly two number at most a constant times $L_R L_T$. Three or more equations give the higher-order term of Theorem 1; its factor $L_{\max}^5$ counts the positions of three equations.

Five passes, and the sixth. The two collisions are paid in the two sampled passes on each side

of the join, and nothing else needs a pass. With five passes the layout is passes 1–2 forward, pass 3 joining and passes 5–4 backward. It is symmetric, so one count serves both sides, and this gives Theorem 4. The standard has a third pass on the backward side, and the proof of Theorem 1 uses pass 4 as a buffer: it stops as soon as a step of pass 4 lands on an existing point, so every value handed down to a joining step is fresh.

5. Proof of Theorem 1

A two-sided sampling process bounds the adversary's advantage by the probability that the process goes wrong, going wrong is always an event about at most two messages, and such an event needs accidents that can be counted. Throughout, $B = 2^{64}$, $M = 2^{128}$, and $\rho_N = B/(M - N)$, which bounds the probability of fixing one 64-bit half of a fresh value of π when fewer than N values have been used. A *message* is a query together with its answer, $T = (L_T, X^T, C^T)$, and it has $6L_T$ steps. We write $\ell = \sum_T L_T = \sigma - q$ and $S = 6\ell \le 6\sigma$. A *point* is a pair (I, O) of 128-bit values, with halves $I = (\alpha, \beta)$ and $O = (\gamma, \delta)$.

5.1. From the adversary to a sampling process

Fix a deterministic adversary, which we may assume never repeats a query or asks for the inverse of an earlier answer. Its *transcript* τ is the list of its messages. For an attainable τ,

$$\Pr_{\text{real}}[\tau] = \Pr_{\pi}[\, W(X^T) = C^T \ \forall T \,],$$
$$\Pr_{\text{ideal}}[\tau] = \prod_n \frac{1}{(2^{64n})_{q_n}},$$

where q_n counts the messages of n semiblocks. Query directions do not appear in either expression, and the rest of the proof rests on this fact. By the H-coefficient technique in its expectation form [20], [25], if $\Pr_{\text{real}}[\tau] \ge (1 - \phi(\tau)) \Pr_{\text{ideal}}[\tau]$ for every attainable τ, then $\mathrm{Adv} \le \max_\tau \phi(\tau)$. So we need a lower bound on the probability that a uniform permutation is consistent with all messages of a fixed transcript.

The process. We obtain the bound by choosing π on the fly, message by message, in the order that suits the analysis. The process $\Gamma(\tau)$ takes the messages one at a time. For a message of pass length L it starts with the *forward steps* $t = 1, \ldots, 2L$. The input (a_{t-1}, d_{t-L}) of such a step is known. If π is already defined there, the step is *old* and reads the output; otherwise it is *fresh* and draws its output uniformly from the unused outputs. The process continues with the *backward steps* $t = 6L, \ldots, 3L+1$, which work the same way in the other direction: the output $(a_t \oplus \langle t \rangle, d_t)$ is known, and a fresh step draws its input uniformly from the unused inputs. Only pass 3 is left, and its *joining steps* $t = 2L + j$, $j = 1, \ldots, L$, connect the two sides (Figure 4). The input of a joining step is known, and so is the right half d_{2L+j} of its output, because that value is the β-half of backward step $3L + j$. For $j < L$ the left half of the output is unconstrained, and the step is *free*: the process draws g uniformly from all 2^{64} values and proposes the output (g, d_{2L+j}). For $j = L$ the whole output $(a_{3L} \oplus \langle 3L \rangle, d_{3L})$ is determined.

Three things can go wrong, and we call each a *bad event*: an *input clash*, where a joining step finds its input already defined; an *output clash*, where a joining step finds its proposed or determined output already used; and an old backward step in pass 4. A clashing output is not inserted, and the process stops at the first bad event. The third event is there for convenience: with it, the right half of every joining output is half of a fresh draw. We let Good be the event that a

run has no bad event. A good run defines all $6L_T$ points of every message, and comparing the probability of its point set under π and under Γ gives the first lemma.

Lemma 7 (Ratio)**.** *For every attainable* τ, $\Pr_{\text{real}}[\tau] \geq (1 - \phi(\tau)) \Pr_{\text{ideal}}[\tau]$ *with*

$$\phi(\tau) = \frac{q^2}{2^{193}} + \Pr_{\Gamma}[\neg \mathsf{Good}].$$

Proof. On Good the defined points form a partial injection containing the whole computation of every X^T, ending in C^T, so a permutation containing them agrees with τ. Two distinct draw sequences first differ at a draw with the same argument, so the corresponding events are disjoint. Take a fixed good point set and let D be the number of points created before a given one. A fresh forward or backward draw contributes $1/(M - D)$ both to the probability that π contains the set and to the probability that Γ produces it. A free joining step contributes $1/(M - D)$ to the first and $1/B$ to the second, and the determined step $1/(M - D)$ and 1. The ratio of the two probabilities is therefore at least $\prod_T (B/M)^{L_T - 1} M^{-1} = \prod_T B^{-(L_T+1)}$, and summing over good point sets gives $\Pr_{\text{real}}[\tau] \geq \prod_T B^{-(L_T+1)} \Pr_{\Gamma}[\mathsf{Good}]$. Finally $(2^{64n})_{q_n} \geq 2^{64nq_n}(1 - q_n^2/2^{64n+1})$ and $n \geq 3$ compare this with $\Pr_{\text{ideal}}[\tau]$. □

The first term of ϕ is the gap between sampling with and without replacement in the ideal world, and the second is the probability that the process goes wrong. Output clashes at free steps are the one bad event that we bound without counting, because the right half of such an output is itself uniform. Let Clash be the event that the first bad event is an output clash at a free joining step.

Lemma 8 (Free output clashes)**.** $\Pr_{\Gamma}[\mathsf{Clash}] \leq \ell S/(M - S)$.

Proof. Let h be step $2L + j$ of T, $j < L$, reached before any bad event, and let U_h be the number of used outputs with right half $d = d_{2L+j}$. Since g is uniform, the proposed output is used with probability U_h/B, so $\Pr_{\Gamma}[\mathsf{Clash}] = \mathbb{E}_{\Gamma}[\sum_h U_h]/B$. The supplier of h, backward step $3L + j$, lies in pass 4, so it was fresh; otherwise the process would have stopped. Thus d is the β-half of a fresh draw without replacement. Every output used when h is reached has a right half that was determined before that draw, or is the right half of an earlier joining output of T, supplied by another, later fresh pass-4 draw. In either case it equals d with probability at most ρ_S. There are at most S used outputs and at most ℓ such steps, and $\rho_S/B = 1/(M - S)$. □

For the other bad events we move from lazy to eager sampling, which makes probability statements easier. In $\Psi(\tau)$ we sample a uniform permutation π and an independent tape of uniform 64-bit values, one per free joining step, and run Γ with three changes: a fresh forward step at x gets output $\pi(x)$, a fresh backward step at y gets input $\pi^{-1}(y)$, and a free joining step takes g from the tape. The permutation π is never evaluated at a joining step, and the bad events and the stopping rule are those of Γ. Let Dup be the event that two created points share an input or an output, which never happens in Γ.

Lemma 9 (Coupling)**.** $\Pr_{\Gamma}[\neg \mathsf{Good}] \leq \Pr_{\Gamma}[\mathsf{Clash}] + \frac{S\ell}{M-S} + \Pr_{\Psi}[\neg \mathsf{Good} \wedge \neg \mathsf{Clash} \wedge \neg \mathsf{Dup}]$.

Proof. Couple the two processes step by step up to the first bad event, maximally at each draw. Free joining draws have the same distribution in both. A fresh forward draw $\pi(x)$ is uniform on the complement of the revealed outputs of π, while Γ draws on the complement of all used outputs. The extra used outputs are joining outputs, at most ℓ of them, so the distance is at most $\ell/(M - S)$, and backward draws are the same with inputs; there are at most S such draws. When

the runs agree up to the first bad event, Ψ has the same points as Γ, so Dup does not occur, and the first bad event is the same in both. □

5.2. Every bad event involves at most two messages

Now fix π and the tape. The *path* of a message T is its computation under π and its own tape alone, with no reference to any other message. Its forward part runs π from (X_A, X_1) through steps $1, \ldots, 2L$, and its backward part runs π^{-1} from $(C_A \oplus \langle 6L \rangle, C_L)$ through steps $6L, \ldots, 3L+1$. Its joining points take their inputs from the forward part and the tape, and have outputs made of the tape value and the known right half for $j < L$, and the determined output for $j = L$. The bad events of the process translate into events on paths.

Definition 10 (Bad events on paths)**.** *For R before T, the event $\mathsf{Bad}_{R,T}$ holds if one of the following holds on the paths of R and T. (i) An input event: a joining input of T equals the input of a point of R, or a forward-step input of T equals the input of a joining point of R. (ii) An output event: the output of a backward step of T in pass 4 equals the output of a point of R, or the output of a backward step of T equals the output of a point of R in R's pass 4 or of a joining point of R. (iii) The determined output of T equals the output of a point of R. The event Bad_T for one message holds if a joining input of T equals the input of a forward or backward point of T or of an earlier joining point of T, or a pass-4 output of T equals the output of a forward point of T or of a backward point of T at a higher index, or the determined output equals the output of another point of T.*

Lemma 11 (Exact reduction)**.** *On* $\neg\mathsf{Dup}$, $\{\neg\mathsf{Good} \wedge \neg\mathsf{Clash}\} \subseteq \bigcup_{R<T} \mathsf{Bad}_{R,T} \cup \bigcup_T \mathsf{Bad}_T$.

Proof. Let x be the first step in processing order at which the actual computation in Ψ differs from the path, or at which a bad event occurs. Every earlier step follows its path, so x's computed side does too. If x is an old forward or backward step that lands on a point z, then z lies on a path. If z was created by a fresh forward or backward step, the value read is that of π, and there is no difference. If z is a joining point, this is an input or output event; z belongs to an earlier message, since a message's joining steps come after its other steps. A fresh forward or backward step reads π or π^{-1}, with no difference. An old pass-4 step gives an output event. At a joining step, a defined input is an input event and a used determined output is event (iii), while a free joining step draws its own tape value and causes no difference, and its output clash is excluded by $\neg\mathsf{Clash}$. On $\neg\mathsf{Dup}$ lookups are unique. So the first such x is one of the listed events. □

5.3. Accidents and what they cost

Fix a pair (R, T) or a single message, and replay the steps of its paths in the order: R's forward steps, R's backward steps, R's joining steps, and then the same for T. Let N be the number of steps. We treat each drawn half as a formal variable: a fresh forward step has output variables (γ, δ), a fresh backward step has input variables (α, β), and a free joining step has its tape variable. Every half is then a variable xor a constant, or a constant. Chain halves use the variables γ, α and the tape, register halves use δ and β, and these two families are disjoint. A *class* is a set of forward or backward steps that share one point, where a forward step joins a class by its input and a backward step by its output. Joining points are classes of their own. A step that joins an existing class *hits* it. We write $x \to y$ when step x hits the class of an earlier step y, $x \sim y$ when two steps share a class, and $c(x)$ for the class of x. A *root* is step 1 or step $6L$

of a message, whose compared side consists of constants. A run of merged pairs at consecutive indices is a *diagonal*, and the pair one pass away from a merged pair, which shares its register value, is its *stride successor*.

At each step x we compare, against every existing class, the input for a forward step, the output for a backward step, the input for a joining step, and both for the last joining step. Following the terminology of structure graphs [24], a comparison is an *induced* collision if both of its half-equations follow from the equations so far, it is *impossible* if one contradicts them, and it is *open* otherwise. A true open comparison is an *accident*. Its new half-equations are added to the system E, which raises $\mathrm{rank}(E)$, the number of independent 64-bit equations, by one or two. The *accident sequence* of a run is the list of its accidents.

A comparison is *final* if, when true, it is one of the bad events above, and the sequence ends there. For a pair these are every comparison at a joining step of T, a forward-step input of T against a joining input of R, every output comparison at a pass-4 step of T, and a backward-step output of T against the output of a pass-4 point or a joining point of R. For a single message they are the comparisons of Bad_T. A forward step of T that lands on a pass-4 point of R is not final. The next lemma makes precise the idea that each equation costs a factor of about 2^{-64}.

Lemma 12 (Cost of a sequence)**.** *For every sequence w,* $\Pr[$*the accident sequence starts with* $w] \le \rho_N^{\mathrm{rank}(w)}$. *Also* $\Pr[\mathrm{rank} \ge 3] \le 2N^6\rho_N^3$ *for* $N \ge 12$.

Proof. Given w, the replay is a fixed symbolic object: it fixes the classes and the list of draws $g_1, \dots, g_m$ made before the last accident of w. Reduce the chain and the register parts of $E(w)$ separately to echelon form, ordering variables by draw, so that each basis equation has its latest variable as a pivot and pivots are distinct. Let P_i be the number of pivots in g_i, so that $\sum_i P_i = \mathrm{rank}(w)$. The two variables of a permutation draw lie on different lines, so a basis equation with a pivot in g_i contains no other variable of g_i.

Let C_i be the event that the execution agrees with the symbolic replay on every comparison before g_{i+1} and that every basis equation with its pivot in $g_1, \dots, g_i$ holds. Then $C_i \subseteq C_{i-1}$, and C_i depends only on the history through g_i. If the accident sequence starts with w, then all equations of $E(w)$ hold, hence all basis equations, so C_m holds.

Fix a history in C_{i-1} with $D < N$ revealed points. If the argument of g_i is already revealed, C_i fails. Otherwise the permutation is uniform over the extensions of the revealed points, so the answer is uniform on at least $M - D$ values. Each pivot in g_i fixes its half to a value determined by the history. So $\Pr[C_i \mid \text{history}]$ is at most $B/(M-D) \le \rho_N$ for one pivot and $1/(M-D) \le \rho_N^2$ for two, and a tape pivot costs $1/B \le \rho_N$. Multiplying gives $\rho_N^{\mathrm{rank}(w)}$.

For the second claim, every accident adds rank at least one, so the first prefix of rank at least 3 has length at most 3. An accident is a choice of step, class and kind, with at most $N^2 + 2N$ options. Thus the probability is at most $\big((N^2+2N) + (N^2+2N)^2 + (N^2+2N)^3\big)\rho_N^3$. □

We must show that one equation never suffices for a bad event, and count the sequences of rank exactly two.

5.4. One equation is not enough

Lemma 13. *No sequence of rank at most 1 produces* $\mathsf{Bad}_{R,T}$ *or* Bad_T.

Proof. The three parts of the proof are the picture of Figure 3b: the forward side, the backward side, and the two together. Before any accident, forward and backward variables are disjoint. A forward input compared with a backward point's input has two non-identity half-equations, and

likewise a backward output against a forward point's output, so a first rank-one accident is not between a forward and a backward step.

On the forward side, the possible rank-one accidents are a prefix extension, two pass-1 steps with equal constants, or two steps with merged register suppliers (Appendix A), and their induced consequences stay in passes 1–2 of both messages. A joining input uses forward right halves and either the last forward step's left half or a tape value. Replacing the forward left halves beyond pass 2 by independent tape values introduces no identity, so no input event occurs. On the backward side with equal lengths, the index reversal $r = 6L + 1 - t$ maps backward steps to forward steps with the same structure, since the counter at reverse index r is $\langle 6L + 1 - r \rangle$, which is injective and equal at equal r for both messages. The forward argument applies, and the consequences of one equation stay out of pass 4.

With unequal lengths $a = L_R \neq b = L_T$, the only shared point at no cost is the last one, $T_{6b} \sim R_{6a}$, and its backward continuation is ruled out by unequal counters. A rank-one accident with equal top constants has both occurrences in pass 6. A pair off the diagonal has no induced chain continuation. A same-index pair can start a diagonal only inside the overlap of the two top passes, and at the boundary of the overlap the next register comparison sets a variable equal to a constant. The stride successor $T_{5b} \to R_{5a}$ of the shared last point costs one chain equation, and its own stride successor $T_{4b} \to R_{4a}$ costs another, so no pass-4 event occurs.

Across the two sides, a forward accident alone leaves the backward input variables independent, and a backward accident alone leaves the forward outputs and tapes independent. One equation can fix one half of an equality between the two sides, but not both. With B_t the point of backward step t, the determined output is $(\alpha(B_{3L+1}) \oplus \langle 3L \rangle,\ \beta(B_{4L}))$, whose halves come from two pass-4 points that are distinct because $L \geq 2$, so it cannot be matched at rank at most one either. □

Counting at rank two relies on a statement about where a variable can occur. Each half at step x is written by one earlier step, its supplier, and if nothing has touched the variable since then, the comparison at x must pay for it.

Lemma 14 (Untouched supplier)**.** *Let v be a variable drawn by a fresh supplier y, and let x be the step that reads it. Suppose that between y and x no true comparison has v in a compared half. Then for every comparison of x's computed side with an existing class, the half-equation containing v is new.*

Proof. If $h = v \oplus k$ and v occurs neither in E nor in h', then $h = h'$ is not implied by E. The hypothesis gives the two exclusions. For a fresh forward supplier, a later forward hit on the supplier's class copies v without comparing it, but the copy is read later than step x; a backward hit on the class compares its output and so touches v. A fresh backward supplier is the mirror image. A tape variable is read only by the next joining step, and the two suppliers of the determined output are read by that output. □

Before the first bad event every pass-4 step of the current message is fresh, and Lemma 14 then has three consequences. First, a comparison of the determined output of T adds exactly two equations, so a rank-two sequence ending there is that single accident. Second, every rank-two sequence ending against a joining output of R is one direct accident (Appendix B). Third, a joining input at $j \geq 2$ contains the previous joining step's tape value, and an output comparison at a pass-4 step $t < 4L$ contains the variable of the fresh step $t + 1$, so these comparisons cannot be reached with no new equation once the budget is spent.

For the first consequence, the two variables of the determined output are $\alpha(B_{3L+1})$, drawn at the last backward step, and $\beta(B_{4L})$. Every later step is a joining step or a fresh pass-4 step, and any

TABLE 3. COVER OF THE RANK-TWO SEQUENCES OF A PAIR, IN THE PROOF FOR SIX PASSES (THIS SECTION) AND FOR FIVE PASSES (SECTION 6). THE ROWS OVERLAP.

Sequences	Six	Five
All accidents forward against forward (including joining)	$40ab$	$40ab$
All accidents backward against backward (unequal lengths)	$8ab+1$	$9ab+1$
One direct accident between the two sides	$8ab$	$4ab$
Prefixes with an explicit accident between the two sides	$14ab$	$10ab$
Prefixes that fully pin the target of an induced hit between the two sides	$49ab$	$49ab$
Total used	$151ab$	$143ab$

true comparison there would already be a bad event, so neither variable is touched. Neither occurs in a competing output either, since earlier backward outputs use suppliers with higher indices, and earlier joining outputs use tapes and suppliers $3L+j<4L$. So Lemma 14 applies to both halves.

This leaves four interfaces to analyze: T's first joining input, T's first pass-4 output, R's pass-4 outputs as targets, and T's forward steps against R's first joining input.

5.5. Counting the sequences of rank two

We write $\rho=\rho_S$, which is at least ρ_N for every replay, and take single messages first. Let H_P be the union of Bad_P, an accident of rank increment two in P's replay stopped at its first bad event, and the mixed family $P_j\to P_1$, $P_{6L}\to c(P_{j-1})$ for $2\le j\le L$, which pins one left half and one right half of the same class by two rank-one accidents.

Lemma 15. $\Pr[H_P]\le 20L_P^2\rho^2+2(6L_P)^6\rho^3$.

Proof. A single-accident sequence has at most $\binom{6L}{2}+(6L-1)=18L^2+3L-1$ choices, the extra $6L-1$ being comparisons of the determined output. The mixed family has at most $L-1$ sequences, and is charged as a prefix, so that all its continuations are covered once. Since $18L^2+4L-2\le 20L^2$, Lemmas 12 and 13 give the claim once we check that two internal rank-one accidents outside the mixed family produce no bad event. After one internal accident on one side there is no further induced internal hit on that side, because adjacent occurrences have unequal counters. So a second rank-one accident on that side adds another chain equation and has a register identity, which leaves the register unconstrained, and such accidents cannot force a first hit across the two sides. Two forward accidents are covered by the forward analysis and two backward accidents by its reversal. With one of each, each side has only the consequences of its one accident, and a comparison across sides still needs a new register equation. □

For a pair we exclude H_R and H_T, write $a=L_R$ and $b=L_T$, and cover the rank-two sequences that produce a first bad event of (R,T) by the families of Table 3. Two of the rows count prefixes of rank two, whether or not they lead to a bad event, so their continuations need no classification: $\Pr[\text{prefix } w \text{ and a later event}]\le\Pr[\text{sequence starts with } w]\le\rho^2$.

Lemma 16. *After excluding H_R and H_T, the number of rank-two sequences or prefixes covering a first bad event of (R, T) is at most $151\, L_R L_T$.*

Proof. If all accidents and hits are forward against forward, this is the forward analysis of Appendix A, applied to the forward steps and the joining steps. It remains valid through the first input event, because replacing the forward left halves beyond pass 2 by tape values preserves every input equation, and six families arise with at most $40ab$ sequences. For backward against backward with equal lengths, the index reversal maps pass 4 to reversed pass 3 and R's joining outputs to the reversed joining steps, so the same six families arise. With unequal lengths the consequences of one accident stay in the top pass except for one stride pair at the top, and Lemma 14 shows that the only non-direct sequence is the chain of two stride pairs descending from a shared last step (Appendix B); the direct sequences number $5ab$ for comparisons final on either side and $3ab$ against joining outputs of R, at most $8ab + 1$ in all.

We turn to accidents between the two sides. The direct ones are T's joining input against a backward point of R ($3ab$ choices), T's pass-4 output against a forward point of R ($2ab$), and the determined output against any point of R ($6a$), at most $8ab$ together. A single rank-two accident between a forward step and a backward point has at most $(2b)(3a)+(3b)(2a) = 12ab$ choices. For two rank-one accidents with an explicit one between the sides, the first must pin a free coordinate to a constant at a root of the other side, and the accident between the sides then lands on the class of that coordinate; there are at most $(a-1)+(a-1)+(a+b-2)+b = 3a+2b-4 \le 2ab$ such prefixes.

Until the first hit between the two sides, their variables are disjoint. So an induced hit between the sides needs both free coordinates of its target class to reduce to constants, which takes one chain pin and one register pin from two rank-one accidents. There are at most $48ab + 4 \le 49ab$ of these (Appendix B). We charge all such prefixes, including those that lead to no bad event, so nothing needs to be said about how far constants propagate.

Two remaining configurations contribute nothing. One forward accident together with one backward accident, and no hit between the sides, gives each side only the consequences of its one accident, so Lemma 13 rules out a bad event on one side, and an induced hit between the sides would need two equations on one side. And a rank-one accident on one side followed by a final comparison between the sides is not a rank-two sequence: each half of that comparison has one variable from each side, so it is implied only if both are pinned, which needs two equations on that line, and the sequence has rank three. □

5.6. Putting it together

By Lemmas 11, 12, 15 and 16, with $\rho = \rho_S$,

$$\begin{aligned}
&\Pr_\Psi[\neg\mathsf{Good} \wedge \neg\mathsf{Clash} \wedge \neg\mathsf{Dup}] \\
&\quad \le \rho^2\Big(151\sum_{R<T} L_R L_T + 20\sum_T L_T^2\Big) \\
&\qquad + 2\cdot 6^6\rho^3\Big[\sum_T L_T^6 + \sum_{R<T}(L_R+L_T)^6\Big].
\end{aligned}$$

The first bracket equals $75.5\,\ell^2 - 55.5\sum_T L_T^2 \le 75.5\,\sigma^2$. In the second, $(L_R+L_T)^6 \le 32(L_R^6+L_T^6)$, so it is at most $32q\sigma L_{\max}^5$, and $64\cdot 6^6 = 2{,}985{,}984$. Lemmas 8 and 9 add $2\ell S/(M-S) \le 12\,\sigma^2/(M-S)$, and Lemma 7 adds $q^2/2^{193} \le \sigma^2/(18BM)$. The main terms sum to $(75.5+6+$

$6)\,\sigma^2/M = 87.5\,\sigma^2/M$, and for $\sigma \le 2^{60}$ the corrections in the denominators keep the total below $88\,\sigma^2/M$ and the higher-order coefficient below $3.1 \cdot 10^6$. This proves Theorem 1. □

6. How Many Passes?

6.1. Two passes: proof of Proposition 5

Take $L = 3$. Fix A, x_1, x_2, choose q distinct words $v_1, \ldots, v_q$, and query $X^{(j)} = (A; x_1, x_2, v_j)$ for $j = 1, \ldots, q$. The adversary outputs 1 if two of the answers agree in their first two data words C_1, C_2 (Figure 3a).

In the real world, steps 1 and 2 are identical for every query. Expose this common computation first; it reveals at most two points of π. The step-3 inputs (a_2, v_j) are distinct, and at most two of them are among the exposed inputs, so at least $m = q - 2$ are fresh. They depend only on the common prefix, so their outputs can be exposed before any step of pass 2, and they are sampled uniformly without replacement from the unused outputs. Suppose j of these outputs have been exposed and have distinct left halves. Among all outputs whose left half equals one of those j halves, at least $j(B-1) - 2 \ge j(B-3)$ are unused, so the next output repeats a left half with probability at least $j(B-3)/B^2$. Hence two step-3 outputs share a left half with probability at least

$$1 - \prod_{j=1}^{m-1} \Big(1 - \frac{j(B-3)}{B^2}\Big) \;\ge\; 1 - \exp\Big(-\frac{(B-3)\,m(m-1)}{2B^2}\Big).$$

For such a pair of queries the chaining values after step 3 agree, and so do the first two register words after pass 1, which the common prefix wrote. So steps 4 and 5 have identical inputs in the two computations and write identical C_1 and C_2, whatever happens at step 6.

In the ideal world, two distinct inputs of four semiblocks receive distinct uniform outputs. Given the first output, exactly $B^2 - 1$ other outputs agree with it in C_1, C_2, so a pair agrees with probability $(B^2-1)/(B^4-1) = 1/(B^2+1)$, and a union bound over the $\binom{q}{2}$ pairs completes the proof. □

6.2. Five passes: proof of Theorem 4

The proof follows Section 5 with the buffer removed. A message of pass length L now has $5L$ steps. The process Γ_5 samples the forward steps $1, \ldots, 2L$ and the backward steps $5L, \ldots, 3L+1$ as before and joins the two sides in pass 3, with free joining draws uniform on all 2^{64} values. It stops at an input clash or an output clash, and an old backward step is not a bad event. Here $\ell = \sigma - q$ and $S = 5\ell$.

Sampling. The proof of Lemma 7 never used pass 4, so $\phi(\tau) = q^2/2^{193} + \Pr_{\Gamma_5}[\neg\mathsf{Good}]$ as before. In the coupling with the eager process Ψ_5, joining draws agree in distribution, and there are at most 4ℓ forward and backward draws, each at distance at most $\ell/(M-S)$. So

$$\Pr_{\Gamma_5}[\neg\mathsf{Good}] \;\le\; \Pr_{\Psi_5}[\neg\mathsf{Good} \wedge \neg\mathsf{Dup}] + \frac{4\ell^2}{M - 5\ell}.$$

Lemma 8 has no counterpart, because the right half of a joining output may now come from an old step. Output clashes at free joining steps are counted as accidents, like every other bad event.

Paths. The exact reduction holds with a shorter list of events. An old backward step that lands on a point created by a forward or backward step reads π^{-1} and stays on its path, in pass 4 as

anywhere else. For R before T the events on paths are these: a joining input of T equals the input of a point of R; a forward-step input of T equals a joining input of R; a joining output of T, free or determined, equals the output of a point of R; and a backward-step output of T equals a joining output of R. The events of one message are the comparisons of its joining inputs and joining outputs with its own earlier points. The first-difference argument of Lemma 11 shows that on $\neg$Dup a run that is not good has one of these events.

Costs. Lemma 12 holds as stated, since comparing a free joining output is one more comparison of two half-expressions and respects the separation into chain and register equations. A replay of N steps has at most N^2 comparisons of the former kinds and $N^2/5$ comparisons of joining outputs, which gives $\Pr[\text{rank} \geq 3] \leq 2N^6\rho_N^3$ for $N \geq 10$.

One message. Define H_P as in Section 5, with the mixed family $P_j \to P_1$, $P_{5L} \to c(P_{j-1})$. Then

$$\Pr[H_P] \leq 17L_P^2\rho^2 + 2(5L_P)^6\rho^3.$$

The single-accident sequences number $\binom{5L}{2} + \sum_{j=1}^{L}(4L + j - 1) = 17L^2 - 3L$, where the sum counts the output comparisons of the joining steps, and the mixed family adds $L-1$. Two rank-one accidents outside the mixed family produce no event, by the argument of Lemma 15, which needs only that counters are injective within one message.

Pairs. The right column of Table 3 lists the cover. The forward count of Appendix A uses only the two forward passes and the comparisons at joining inputs, so it applies unchanged. On the backward side the consequences of Lemma 14 are no longer available, since they rested on fresh pass-4 steps. For equal lengths $a = b = L$ a symmetry replaces them. Put $r = 5L + 1 - t$. Read in the index r, the backward side runs π^{-1} from the chaining value $C_A \oplus \langle 5L \rangle$ with the step counter $\langle 5L - r \rangle$, which is injective and the same for both messages, and the register still has stride L. The reversed tapes $\bar{g}_k = g_{L-k} \oplus \langle 3L - k \rangle$, $1 \leq k < L$, are again independent and uniform. Joining outputs become joining inputs, and a backward output that meets a joining output of R becomes a forward input that meets a joining input of R. Final comparisons do not change the paths before them, so their order may be reversed when bounding the union of first events. The backward side is therefore a forward side, and its count is the $40ab$ of Appendix A.

For unequal lengths the two reversed counter schedules differ, and the count is done directly.

Lemma 17. *Let $a = L_R \neq b = L_T$. After excluding H_R, H_T and the prefixes that fully pin a target, the rank-two sequences and prefixes that cover a first bad event on the backward sides of (R, T) number at most $9ab + 1$.*

The $9ab + 1$ are $5ab$ direct final comparisons, $4ab$ single accidents of rank two between backward steps of the two messages, and one chain: a shared last step $T_{5b} \sim R_{5a}$, its stride successor $T_{4b} \to R_{4a}$, and equal determined outputs. Appendix C gives the proof. Between the two sides, the direct accidents are the joining inputs of T against backward points of R and the joining outputs of T against forward points of R, $b(2a) + b(2a) = 4ab$ in all. A single rank-two accident between a forward step and a backward point has $(2b)(2a) + (2b)(2a) = 8ab$ choices, the four families of prefixes with an explicit rank-one accident between the sides are those of Section 5 and number at most $2ab$, and the count of full pins in Appendix B does not depend on the number of backward passes. The total is $143ab$.

Putting it together. With $\rho = \rho_S$,

$$143\sum_{R<T} L_R L_T + 17\sum_{T} L_T^2 = 71.5\,\ell^2 - 54.5\sum_{T} L_T^2,$$

which is at most $71.5\,\sigma^2$, and the terms of rank three sum to at most $2 \cdot 5^6 \cdot 32\,q\sigma L_{\max}^5\rho^3 = 10^6\,q\sigma L_{\max}^5\rho^3$. Adding the coupling term $4\ell^2/(M-S)$ and $q^2/2^{193}$ gives main terms that sum to $75.5\,\sigma^2/M$, and for $\sigma \le 2^{60}$ the corrections in the denominators keep the total below $76\,\sigma^2/M$ and the higher-order coefficient below $1.01 \cdot 10^6$. □

7. KW, KWP and Many Keys

7.1. Proof of Theorem 2

KW is encode-then-encipher with 64 bits of fixed redundancy, over an enciphering scheme whose tweak is the length [13], [26], so Theorem 2 follows from Theorem 1 by the standard reduction. Given a DAE adversary with q_e wrap and q_v unwrap queries, an adversary $\mathcal{A}$ against W answers a wrap query P by the forward query $W(\mathrm{ICV1} \,\|\, P)$ and an unwrap query C by the inverse query $W^{-1}(C)$ followed by the format check. It uses the same total number of semiblocks, counting each wrap's ICV1 semiblock. In the ideal world of Theorem 1, an unwrap of a new C returns a uniform unused string of its length, whose first semiblock equals ICV1 with probability at most $2^{-64} \cdot 2^{64n}/(2^{64n} - q)$, and summing with $q \le 2^{63}$ gives $q_v/2^{63}$. The ideal encryption oracle of the DAE game returns independent random strings, whereas the permutation world has no collision, which costs $q_e^2/2^{193}$ since $n \ge 3$. Theorem 1 covers the inverse queries of $\mathcal{A}$ and reveals the preimage to it, which is all the reduction needs. □

7.2. Proof of Theorem 3

KWP adds one difficulty to W, and it is the reason the argument of Section 5 does not carry over unchanged. A one-block wrap returns $\pi(\Lambda_m \,\|\, P')$ for $m \le 8$, so the adversary sees π on inputs of its choice, inside a format with about 2^{64} members. We call these *exposed points*. Both halves of their inputs are chosen, both halves of their outputs are seen, and they sit in the same π as all of W. In the two-sided process, an unwrap query with a chosen ciphertext could land on an exposed output at no cost at all. So we do not process unwraps as messages. We prove the DAE bound directly, processing wraps forward only, and letting an unwrap enter only through the probability that it would be accepted. The accident counting of Section 5 remains the main tool.

Let p be the number of one-block wraps, q the number of long wraps with pass lengths $L_T \ge 2$, $\ell = \sum_T L_T$, $S = 6\ell$, and $D = S + p$. Let r be the number of two-semiblock unwraps, and $q_L = q_v - r$ the number of long unwraps, with pass lengths L_k and $s = \sum_k L_k$. The semiblock budget gives $\ell + q + 2p + \sum_k (L_k + 1) + 2r \le \sigma$.

A transcript τ lists the wraps with answers, the exposed points with answers, and the unwraps with answer "reject"; transcripts with an accepted unwrap have ideal probability 0. The process $\Gamma(\tau)$ takes the exposed points first, as a prescribed injection, and then the long wraps in query order, forward. Pass-6 positions other than the last draw γ only, with δ set by the ciphertext, and the last step is determined. Let $\mathcal{G}$ be the good event: no collision of created points, no old step in passes 3–6, and no 18 points sharing a γ or a δ. For a good point set P, let $V_k(P)$ be the probability, over a uniform permutation extending P, that unwrap k is accepted.

Lemma 18. *For every transcript with no accepted unwrap,*

$$\frac{\Pr_{\mathrm{real}}[\tau]}{\Pr_{\mathrm{ideal}}[\tau]} \;\ge\; 1 - \Pr_{\Gamma}[\neg\mathcal{G}] - \mathbb{E}_{\Gamma}\Big[1_{\mathcal{G}} \sum_k V_k(P)\Big].$$

TABLE 4. EVERY LOSS IN THEOREM 3. EACH ROW IS AVERAGED OVER IDEAL TRANSCRIPTS.

Source	Bound
Collisions	$19\,\sigma^2/M$
Cap on equal halves	$0.032\,\sigma^2/M$
Freshness, rank two	$[170\sum L_T^2 + 1720\sum_{R<T} L_R L_T + 18\,p\ell]/M$
Freshness, messages at risk	$17600\,p\ell/M$
Freshness, higher order	$7721568\,q\sigma L_{\max}^5/B^3$
Two-semiblock unwraps	$(1+2^{-7})\,r/B + 7r\sigma/M$
Long unwrap, step 1 fresh	$9/B$ each
Long unwrap, step 1 old	$64/B + 14L_kD/M + 35868L_k^2/M + 91728\,p/M + 7\cdot 10^6 L_k^6/B^3$ each

Proof. Distinct good histories assign different outputs to the same input at their first differing forward draw, so their completion events are disjoint, and every completion agrees with all wraps. Given P, the probability that all unwraps are rejected is at least $(1-\sum_k V_k(P))_+$. Taking the exposed points first costs nothing: they have probability $1/(M)_p \geq M^{-p}$, while the ideal world gives each a uniform string. The forward draws compare as in the ordinary H-coefficient argument. □

The right side splits into the terms of Table 4.

Collisions and the cap. Charge each sampling collision to its creation occurrence: a draw that is not a pass-6 position costs $1/M$, and a pass-6 position costs $1/B$ for its fresh γ and, on averaging over the ideal answers, $1/B$ for the equal δ. With the ideal collision charges this totals at most $19\sigma^2/M$. Each output coordinate of a created point is a uniform draw or an ideal answer, so 18 labels sharing a value cost at most B^{-17}, and the union bound on both coordinates is at most $2D(D/B)^{17}/18! \leq 0.032\sigma^2/M$. The cap ensures that a known half has at most $\mu = 17$ matching points, which we use below.

Two-semiblock unwraps. Fix a rejected ciphertext C of two semiblocks. If C equals no determined wrap output, its preimage in a uniform completion is valid with probability at most $v/(M-D)$, where $v = |\mathsf{Valid}_2| = 256(B-1)/255$ counts the valid one-block inputs. A determined wrap output equal to C is unique on a good history, and its input halves come from two distinct fresh draws, so it costs at most v/M. These two cases share one leading term because which of them holds is fixed by τ. A fresh output equals C with probability at most $1/M$, and a pass-6 position costs at most v/B^3. Adding the case that a future one-block answer equals an earlier rejected ciphertext gives $(1+2^{-7})r/B + 7r\sigma/M$.

Freshness with exposed points. The main work is freshness in passes 3–6. We use eager sampling: the exposed map is prescribed by τ and the rest of π is an independent uniform function. A *visit* is a step whose input is an exposed point. Step 1 of a long wrap is never a visit, because its first semiblock has $m \geq 9$ and exposed inputs have $m \leq 8$. A first visit at a step $x \leq L$ needs one equation, a *pin* of the predecessor's γ to a format value. Its right half is the known input semiblock, so at most eight exposed points are possible, and the pin has at most 8 targets rather than p. A visit above pass 1 needs a second equation, with up to p targets, and we count it as a final event. Let a *cluster* be a maximal run of consecutive visits, and let $\mathsf{BadVisit}$ be the event that

some message has a visit above pass 1 or two clusters. Without BadVisit every message has at most one cluster, inside pass 1, and every accident has at most $N^2 + 8N$ choices, which is what removes the factor p from the higher-order term.

One rank-one mechanism does reach a second cluster: a pin at step 2 of a message with plaintext length m, followed by induced returns, when some exposed answer has $\gamma = \Lambda_m \oplus \langle b \rangle$ with $2 \le b \le \min(14, L-2)$. We say such a message is *at risk*. Weighting the messages by $1/L$ gives $\mathbb{E}[\#\text{at risk}] \le 2200\,\ell p/B$, since the sum of the weights over all values of γ is at most $104\,H_{2^{29}} < 2200$, where H_n is the nth harmonic number. We charge the whole branch to the event Risk that some message at risk has its step 2 on an exposed input.

Lemma 19. *Averaged over ideal transcripts,*

$$\mathbb{E}_{\text{ideal}} \Pr[\mathsf{Risk}] \;\le\; \frac{8\,\mathbb{E}[\#\textit{at risk}]}{B} \;\le\; \frac{17600\,p\ell}{M}.$$

Proof. Fix τ. Step 1 of a long wrap is never a visit, so its output is uniform and independent of τ. The input of step 2 is $(\gamma \oplus \langle 1 \rangle, X_2)$. At most eight exposed inputs have right half X_2, one per value of $m \le 8$, so the step lands on one with probability at most $8/B$. □

Lemma 20. *The probability of an old step in passes 3–6, averaged over ideal transcripts, is at most*

$$\frac{170\sum_T L_T^2 + 1720\sum_{R<T} L_R L_T + 18\,p\ell}{M} + \frac{17600\,p\ell}{M} + \frac{7721568\,q\sigma L_{\max}^5}{B^3}.$$

The proof, in Appendix D, runs the accident sequences of Section 5 with the exposed classes held fixed as constants. A single equation produces no event beyond the at-risk branch, the rank-two sequences fall into ten families, and rank three gives the last term.

Long unwraps. Fix a long unwrap of pass length $L \ge 2$ and a good P, and run its chain backward through steps $t = 6L, \dots, 1$ with requested outputs $Y_t = (\alpha_{t+1} \oplus \langle t \rangle, z_t)$, where $z_t = \beta_{t+L}$ or a ciphertext block. Step t is *fresh* if Y_t is new, and otherwise it is a *hit*: a *P-hit* if Y_t is an output of P, and a *self-hit* if it repeats the request of a higher fresh step. The *parents* of t are $t+1$ and $t+L$. A hit with exactly one fresh parent is a *paid hit*: the fresh half must equal a value determined by the target, which costs about $1/B$ per target, and by the cap there are at most μ targets in P. An unwrap is accepted only if it recovers a legal first semiblock and zero padding.

Lemma 21. *For a long unwrap with $L \le 2^{29}$, $\Pr[\textit{valid} \wedge \textit{step 1 fresh}] \le 9/B$, and*

$$\Pr[\textit{valid} \wedge \textit{step 1 old}] \le \frac{64}{B} + \frac{14LD}{M} + \frac{35868\,L^2}{M} + \frac{91728\,p}{M} + \frac{7 \cdot 10^6 L^6}{B^3}.$$

A fresh step 1 supplies one of eight legal first semiblocks. When step 1 is old, the proof in Appendix D counts paid hits: with none, the unwrap walks through old points and pays for a comparison with a hidden left half; with one or two, the set of fresh steps is rigid and a further payment is found; and three paid hits give the last term.

Adding up. Using $\ell^2 = \sum L_T^2 + 2\sum_{R<T} L_R L_T$, the rank-two sequences cost at most $860\ell^2/M$. For the long unwraps, $\sum_k 14 L_k D = 14s(6\ell + p)$, $\sum_k L_k^2 \le L_{\max} s$, $\sum_k 91728p = 91728 q_L p \le 45864sp$, and $\sum_k L_k^6 \le s L_{\max}^5$. The terms of order σ^2/M other than the one with $L_{\max}$ are

$$f = 860\ell^2 + 17618\,p\ell + 84\,\ell s + 14\,ps + 45864\,ps + 7r\sigma.$$

Put $R = \sigma - \ell$, so that $2p + s \le R$. Then $17618p\ell + 84\ell s \le 8809\ell R$ and $45878ps \le 5734.75R^2$, hence $f - 7r\sigma \le 860\ell^2 + 8809\ell R + 5734.75R^2 \le 5734.75\sigma^2$. With $7r\sigma \le 3.5\sigma^2$ and $19.032\sigma^2$ from the first two rows, the coefficient is at most 5758. The $1/B$ terms give $(1 + 2^{-7})r + 73q_L \le 73q_v$. Since $s \le \sigma$ when $s > 0$, the higher-order terms give $7721568(q + q_v)\sigma L_{\max}^5/B^3$. This proves Theorem 3. □

7.3. Many keys: proof of Corollary 6

The wrapping function. Fix a complete transcript τ of all users. The users' oracles are independent, so the probability that all of them are consistent with τ is a product over users, in the real world and in the ideal world. Theorem 1 holds transcript by transcript, so the real factor of user i is at least $(1 - \epsilon_i^W)_+$ times its ideal factor, and $\prod_i (1 - \epsilon_i^W)_+ \ge 1 - \sum_i \epsilon_i^W$. The main term follows from $\sum_i \sigma_i^2 \le \sigma\sigma_{\max}$.

KW. By the first part, replace every user's W by an independent family of random permutations. Each new inverse answer then has a valid first semiblock with probability at most $2/B$, and wrapping differs from independent random strings only through collisions within one user and one length, at most $q_{e,i}^2/(2B^3)$ for user i.

KWP. The rows of Table 4 are averaged over ideal transcripts, so Theorem 3 does not bound each transcript, and we use a hybrid: in hybrid j the users $1, \ldots, j$ are ideal and the others real. An adversary that separates hybrids $j - 1$ and j gives a single-user adversary that sends the queries of user j to its own oracle and simulates the other users itself, within the caps of user j. So the gap is at most $\epsilon_j^{\mathrm{KWP}}$.

AES keys. Compare with the experiment in which the users' permutations are independent of the cipher that the adversary queries. There the key of each of the p queries equals some user key with probability at most $u/2^k$, and two user keys coincide with probability at most $\binom{u}{2}/2^k$. Outside these events the two experiments coincide. □

8. Related Work

AES Key Wrap. Rogaway and Shrimpton defined DAE and gave SIV as a provably secure design [13]. Dworkin asked for review of key-wrap algorithms [27], and SP 800-38F records Kelsey's observation on long messages [4]. The Wycheproof documentation turns the same collision idea into a chosen-plaintext distinguisher and points out the one-block path of KWP [15]. Mattsson corrected the forgery bounds, gave an encryption-only distinguisher, analyzed the one-block path and the role of the step counter, and estimated how the forgery probability would change with the number of passes [14]. Khovratovich proved security for a different design built on a fixed wide permutation [28]. Our theorems give the corresponding upper bounds and place the number of passes between a proof and an attack.

Feistel networks and one permutation. Structurally W is a generalized Feistel network. Hoang and Rogaway proved the security of such networks by coupling [16], Shen, Guo and Wang improved the bounds [17], and Nakaya and Iwata treated round functions built from tweakable block ciphers by the H-coefficient method [29]. All of these assume independent round functions, whereas in W the round function is one permutation of the whole state of two halves. The problem of one permutation used many times is familiar from key-alternating ciphers, which are analyzed by counting chains of collisions [30]–[33]. Those proofs rely on secret keys xored around a public

permutation; in W the only secret is the permutation, the added constants are public counters, and there is no key to blind a collision.

Modes and collision counting. Structure graphs count accidents across CBC-MAC chains [24], [34], and our accident sequences serve the same purpose for pairs of computations that run in two directions. Two-sided sampling goes back to Luby–Rackoff and Naor–Reingold [22], [23]. In those proofs the two sides belong to different keyed functions, whereas here the two sides of every message use the same π. Multi-user security [21] is established for GCM in TLS 1.3 with bounds in the data per key [35], [36], and Corollary 6 gives bounds of that form for key wrapping.

9. Discussion and Conclusion

AES Key Wrap now has a security proof: W is a strong pseudorandom permutation, and KW and KWP are deterministic authenticated encryption schemes. Sampling each message from both ends reduces the analysis to pairs of messages, and the 64-bit chain is handled by counting collisions rather than excluding them.

For the revision of SP 800-38F. When the algorithm is used to wrap keys, the parameters of the standard have a proof behind them, and the bounds give the revision a quantitative basis for its limits. First, the bounds are in the data under each key and add up across keys (Corollary 6), so a usage limit belongs on each key-encryption key and not on the system. Second, the theorems take the key to serve one scheme only, which supports the recommendation to keep separate keys for KW, for KWP and for any other mode [14], the lesson also taught by attacks that reuse a key across modes [37], [38]. Third, the six passes now have a rationale: five are provably enough and two are not, so the standard has a margin of one pass.

Open problems. The smallest sufficient number of passes is open. Three and four passes lie between Proposition 5 and Theorem 4, and in our framework they call for an argument that works with a single sampled pass on one side of the join. The other natural question is the beyond-birthday security that Rogaway and Shrimpton suggested for this design [13].

References

[1] National Institute of Standards and Technology, "AES key wrap specification," Nov. 2001.

[2] J. Schaad and R. Housley, "Advanced Encryption Standard (AES) key wrap algorithm," RFC 3394, Sep. 2002.

[3] R. Housley and M. Dworkin, "Advanced Encryption Standard (AES) key wrap with padding algorithm," RFC 5649, 2009.

[4] M. Dworkin, "Recommendation for block cipher modes of operation: Methods for key wrapping," NIST SP 800-38F, Dec. 2012.

[5] J. Schaad, "Use of the Advanced Encryption Standard (AES) encryption algorithm in Cryptographic Message Syntax (CMS)," RFC 3565, Jul. 2003.

[6] M. Jones, "JSON Web Algorithms (JWA)," RFC 7518, May 2015.

[7] OASIS, "PKCS #11 specification version 3.1," 2023.

[8] OASIS, "Key Management Interoperability Protocol specification version 2.1," 2020.

[9] M. Watson, Ed., "Web Cryptography API," W3C Recommendation, Jan. 2017.

[10] Amazon Web Services, "AWS Key Management Service developer guide: Importing key material," 2026.

[11] A. Becker and M. Jenkins, "Commercial National Security Algorithm (CNSA) suite 2.0 profile for S/MIME," IETF Internet-Draft, 2026.

[12] National Institute of Standards and Technology, "Pre-draft call for comments: Methods for key wrapping (SP 800-38F Rev. 1)," May 2026.

[13] P. Rogaway and T. Shrimpton, "A provable-security treatment of the key-wrap problem," in *EUROCRYPT*, LNCS 4004, 2006, pp. 373–390. Full version: IACR ePrint 2006/221.

[14] J. Preuß Mattsson, "Revisiting AES Key Wrap: Forgery bounds, structural observations, and recommendations for the revision of NIST SP 800-38F," IACR Cryptology ePrint Archive, Paper 2026/2636, Oct. 2026.

[15] C2SP, "Project Wycheproof: KeyWrap (NIST SP 800-38F)," documentation file `doc/key_wrap.md`, https://github.com/C2SP/wycheproof, accessed Nov. 2026.

[16] V. T. Hoang and P. Rogaway, "On generalized Feistel networks," in *CRYPTO*, LNCS 6223, 2010, pp. 613–630.

[17] Y. Shen, C. Guo, and L. Wang, "Improved security bounds for generalized Feistel networks," *IACR Trans. Symmetric Cryptol.*, vol. 2020, no. 1, 2020.

[18] S. Halevi and P. Rogaway, "A tweakable enciphering mode," in *CRYPTO*, LNCS 2729, 2003, pp. 482–499.

[19] S. Halevi and P. Rogaway, "A parallelizable enciphering mode," in *CT-RSA*, LNCS 2964, 2004, pp. 292–304.

[20] J. Patarin, "The 'coefficients H' technique," in *SAC*, LNCS 5381, 2008, pp. 328–345.

[21] M. Bellare, A. Boldyreva, and S. Micali, "Public-key encryption in a multi-user setting: Security proofs and improvements," in *EUROCRYPT*, LNCS 1807, 2000, pp. 259–274.

[22] M. Luby and C. Rackoff, "How to construct pseudorandom permutations from pseudorandom functions," *SIAM J. Comput.*, vol. 17, no. 2, pp. 373–386, 1988.

[23] M. Naor and O. Reingold, "On the construction of pseudorandom permutations: Luby–Rackoff revisited," *J. Cryptology*, vol. 12, no. 1, pp. 29–66, 1999.

[24] M. Bellare, K. Pietrzak, and P. Rogaway, "Improved security analyses for CBC MACs," in *CRYPTO*, LNCS 3621, 2005, pp. 527–545.

[25] V. T. Hoang and S. Tessaro, "Key-alternating ciphers and key-length extension: Exact bounds and multi-user security," in *CRYPTO*, LNCS 9814, 2016, pp. 3–32.

[26] M. Bellare and P. Rogaway, "Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography," in *ASIACRYPT*, LNCS 1976, 2000, pp. 317–330.

[27] M. Dworkin, "Request for review of key wrap algorithms," IACR Cryptology ePrint Archive, Paper 2004/340.

[28] D. Khovratovich, "Key wrapping with a fixed permutation," IACR Cryptology ePrint Archive, Paper 2013/145.

[29] Y. Nakaya and T. Iwata, "Generalized Feistel structures based on tweakable block ciphers," *IACR Trans. Symmetric Cryptol.*, vol. 2022, no. 4, pp. 24–91, 2022.

[30] S. Chen and J. Steinberger, "Tight security bounds for key-alternating ciphers," in *EUROCRYPT*, LNCS 8441, 2014, pp. 327–350.

[31] S. Chen, R. Lampe, J. Lee, Y. Seurin, and J. Steinberger, "Minimizing the two-round Even–Mansour cipher," *J. Cryptology*, vol. 31, no. 4, 2018.

[32] Y. Wu, L. Yu, Z. Cao, and X. Dong, "Tight security analysis of 3-round key-alternating cipher with a single permutation," in *ASIACRYPT*, LNCS 12491, 2020, pp. 662–693.

[33] L. Yu, Y. Wu, Y. Yu, Z. Cao, and X. Dong, "Security proofs for key-alternating ciphers with non-independent round permutations," in *TCC*, 2023.

[34] A. Jha and M. Nandi, "Revisiting structure graphs: Applications to CBC-MAC and EMAC," *J. Math. Cryptol.*, vol. 10, no. 3–4, pp. 157–180, 2016.

[35] M. Bellare and B. Tackmann, "The multi-user security of authenticated encryption: AES-GCM in TLS 1.3," in *CRYPTO*, LNCS 9814, 2016, pp. 247–276.

[36] V. T. Hoang, S. Tessaro, and A. Thiruvengadam, "The multi-user security of GCM, revisited: Tight bounds for nonce randomization," in *ACM CCS*, 2018, pp. 1429–1440.

[37] T. Jager, K. G. Paterson, and J. Somorovsky, "One bad apple: Backwards compatibility attacks on state-of-the-art cryptography," in *NDSS*, 2013.

[38] F. Strenzke and J. Roth, "Legacy encryption downgrade attacks against LibrePGP and CMS," *IACR Commun. Cryptol.*, vol. 2, no. 1, 2025.

Appendix A.
The Forward Pair Analysis

This appendix supplies the forward analysis used in Lemmas 13 and 16 and in Section 6. Fix messages R before T with pass lengths L, L' and write R_t, T_t for their inputs at step t. Let p be the common prefix depth, the largest $t \le \min(L, L')$ with $X_A = X'_A$ and $X_u = X'_u$ for $u \le t$. A chain equation involves γ-variables and constants, and a register equation involves δ-variables and constants. For two steps in one class, the *chain equation* of the pair is the equality of their left halves and the *register equation* is that of their right halves. A half-equation is an *identity* if it holds syntactically.

With no equations, the induced hits are exactly $T_t \to R_t$ for $t \le p$, and three kinds of open comparison have rank one: a *prefix extension* at index $\min(L, L')+1$, when $p = \min(L, L')$; *equal constants*, two pass-1 steps with equal X-constants; and a *stride hit* $T_{u+L'} \to R_{u+L}$ with $u \le p$, the pair of steps that read the register values written by the merged steps T_u and R_u. All lie in passes 1–2. A chain identity needs equal counters, hence equal indices, so it never holds inside one message. Across the pair it needs the predecessors merged, and induction on the index gives exactly the prefix.

Rank one cannot reach a bad event. One accident lies on otherwise unconstrained classes, so its only induced consequences are prefix copies and the following. After a prefix extension, the next index needs a register equation on a new variable. An equal-constants accident at a common index t can start a diagonal run, which ends at index $L + p$ or earlier, at most $2L - 1$. A stride hit has a run up to $L + p$. All merged pairs therefore lie in passes 1–2.

At rank two, write r_a and r_d for the numbers of chain and register equations, so that (r_a, r_d) is $(0, 2)$, $(1, 1)$ or $(2, 0)$. Let (T_k, R_m) be the event pair, with $k > 2L'$ or $m > 2L$, and put $d = k - m$.

TABLE 5. FORWARD RANK-TWO FAMILIES PRODUCING A PASS-3 EVENT.

Family	Sequences
One accident of rank two	$36LL'$
Prefix extension, then its stride successor	1
Stride hit at $u \le p$, then its stride successor	p
Re-merge after a single differing block ($L = L'$)	1
Stride hit, its diagonal run, then one more accident ($L = L'$)	$p(p+1)/2$
Two same-index accidents with a run between ($L = L'$)	LL'
Total, using $p \le \min(L, L')$ and $L, L' \ge 2$	$40LL'$

When $(r_a, r_d) = (0, 2)$, every same-class pair is a chain identity, so the merged indices form an interval starting at 1. Past the prefix, each new index needs a register equation on a new variable, which requires rank at least three.

When $(r_a, r_d) = (1, 1)$, there is one chain edge and one register edge, so there is at most one diagonal run with a bottom. If $d = 0$ and $L = L'$, the event lies in a run $[b + 1, k']$ with $k' \ge 2L+1$. The register equations on the run that are not identities are distinct, and counting them forces $b = p + 1$, or $b = L$ with $p = L - 1$ (the *re-merge*), or a single pair carrying both edges. If $d = 0$ and $L \ne L'$, a run of length at least two needs two register edges, so the event pair carries both. If $d \ne 0$, the chain equation of the event pair is the chain edge. Either its register equation is the register edge, which is a single accident, or its register-parent pair lies on a diagonal starting right after the prefix, which is a prefix extension followed by its stride successor.

When $(r_a, r_d) = (2, 0)$, every same-class pair is a register identity. Following register suppliers shows that $(T_{k-L'}, R_{m-L})$ and $(T_{k-2L'}, R_{m-2L})$ are merged. Both accidents are stride-type hits, with the first one in passes 1–2. Index $p + 1$ is never merged, so no run passes through index $L + p + 1$. This leaves the stride chains and, for $L = L'$, the run chains and the double re-merge of Table 5.

Accidents inside one message produce no event between the two messages at rank two beyond those charged to that message. Let K be the set of indices k such that T_k shares a class with a step of R. Then K is an interval $[1, K_{\max}]$ with $K_{\max} \le p + 1$: an induced hit at k between the messages has its chain equation an identity or the recorded edge, in either case $k - 1 \in K$, and the register budget allows at most one index beyond the prefix. So every pair across the messages has $k \le p + 1 \le 2L'$. An event would need $m > 2L$, and then R_m shares a class with R_k, an internal event of R.

Appendix B.
Accidents Between the Two Sides

This appendix proves the second consequence of Lemma 14 and supplies the backward and two-sided counts used in Lemmas 15 and 16.

Joining outputs as targets. Let $a = L_R$. Every pass-4 backward step i of R is fresh before the first event of R. Its input halves occur in the input of R_i and, on the output side, in steps $i - 1$ and $i - a$, whose comparisons are final, or in a joining output. So the first non-final access by T to a pass-4 input point of R must be a forward hit on it, which adds two equations. At total

rank two, a prefix reaches at most one such point. The determined joining output of R has halves from two different pass-4 points, so one of them is untouched. Hence a rank-two sequence ending against it is one direct accident. A joining output with $j < L$ has a tape value that is new in any first final comparison.

One message. Suppose the first two accidents each have rank one and one of them is between a forward and a backward step. A first accident between two backward steps constrains only backward α-variables, so a later backward hit on a forward class costs two equations. After a first accident between two forward steps the register is unconstrained, so the first backward hit needs a new register equation and its chain equation must already follow; that forces step $6L$ and a γ pinned to a constant, which happens exactly when the target is step 1. So the sequences are $P_j \to P_1$, $P_{6L} \to c(P_{j-1})$, $2 \le j \le L$. Charging this prefix once per message covers all its later continuations.

Backward steps of unequal lengths. Let $a \neq b$. With at most one backward accident of rank one, every backward hit has both occurrences in the top pass, except possibly $T_{5b} \to R_{5a}$. Suppose a rank-two sequence ending in a first bad event is not one direct accident. A final second accident has a new chain equation, so its register equation is an identity, and the only sharing in pass 5 between the messages is $T_{5b} \sim R_{5a}$. So the sequence is the chain $T_{6b} \sim R_{6a}$, $T_{5b} \to R_{5a}$, $T_{4b} \to R_{4a}$. An induced final comparison reduces to $T_{4b} \to R_{4a}$; with $v = \alpha(R_{4a+1})$ the comparison either adds an equation or would need $v \oplus \langle 4b \rangle = v \oplus \langle 4a \rangle$, which fails since $a \neq b$. The direct count is $b(3a) + (3b)a - ba = 5ab$ for comparisons final on either side, plus $3ab$ against joining outputs of R, plus one chain, for a total of $8ab + 1$.

An induced first hit between the sides. Equations on the two sides have disjoint supports, so both free coordinates of the target class reduce to constants. A single accident of rank two on one side cannot supply both to one class, because its two suppliers are at indices separated by $L - 1$. So there are two rank-one accidents. Since $r_a \le 1$, each class on one side has at most two occurrences per message, hence at most four across the pair. A first chain pin needs a root as one endpoint, which gives at most $4(a + b - 2)$ choices over both sides. The second accident pins the register coordinate of the same class. There are at most four consumers of that coordinate, and at most six occurrences have the required chain expression, four for an unpinned variable plus two roots. This gives $24 \cdot 4(a + b - 2) \le 48ab$ sequences. If the register pin comes first it is a prefix extension (one sequence), and the chain pin has at most four choices. The total is at most $49ab$. The count uses class multiplicity, the positions of the roots and the injectivity of the counters, and not the number of backward passes.

An explicit accident between the sides. For two rank-one accidents with an explicit one between the two sides, the first accident and the one between the sides are, up to symmetry: $R_j \to R_1$ then $T_{6b} \to c(R_{j-1})$; $R_j \to R_{6a}$ then $T_1 \to c(R_{j+1})$; a root accident of T or R pinning one γ, then T_{6b} on that class; and a prefix extension pinning a shared δ, then a pass-6 backward step of T on the root class. The counts are $a - 1$, $a - 1$, $a + b - 2$ and b. With five passes the same four families arise, with $5a$ and $5b$ in place of $6a$ and $6b$ and the last pass in place of pass 6.

Appendix C.
Five Passes: Backward Steps of Unequal Lengths

This appendix proves Lemma 17. Let $a = L_R \neq b = L_T$. For a message P of pass length L, the backward steps of pass 5 are $4L < t \le 5L$ and those of pass 4 are $3L < t \le 4L$. For a

backward class c write (α_c, β_c) for its input variables. The output requested at a backward step other than the root is

$$Y_t^P = \begin{cases} (\alpha_{c(P_{t+1})} \oplus \langle t \rangle,\ C_{t-4L}^P), & 4L < t < 5L, \\ (\alpha_{c(P_{t+1})} \oplus \langle t \rangle,\ \beta_{c(P_{t+L})}), & 3L < t \le 4L, \end{cases}$$

the root output is the constant pair (h_P, C_L^P) with $h_P = C_A^P \oplus \langle 5L \rangle$, and the determined joining output is $F_P = (\alpha_{c(P_{3L+1})} \oplus \langle 3L \rangle,\ \beta_{c(P_{4L})})$. All counters below are distinct 64-bit encodings, since $\sigma \le 2^{60}$.

Fresh continuation. If a backward step of the current message is fresh after the last accident that is not final, its new α-variable is absent from the equations and from every existing output, so the next backward step cannot be an induced hit. Repeating this, all later backward steps of the message are fresh, and the last fresh α also prevents an induced comparison of the message's determined output with an output of the earlier message.

One accident. Suppose that, before a true final comparison, the backward replay has exactly one accident, of rank one. We claim: (i) its equation is a chain equation; (ii) every pass-4 step is fresh, except that the accident itself can be the stride pair $T_{4b} \to R_{4a}$ after the free sharing $T_{5b} \sim R_{5a}$; (iii) apart from a comparison already made, a comparison that involves a pass-4 step has its chain equality implied only if a same-index diagonal of pass 5 reaches the first pass-4 step of the longer message, or if an adjacent root alias in a message of pass length two has pinned the α of its root class. In the second case a register equation that completes the comparison pins the β of the same class, which is a full pin.

With no equation, the only class that the two messages can share is the root class, and its two consumers have the different counters $5a-1$ and $5b-1$. So the accident compares two pass-5 steps with equal register constants or is the stride pair, and both add a chain equation. Let the accident be in pass 5 with target class C. If α_C is not an endpoint of its equation, a chain equality with another output is an identity, with the same supplier class and the same counter. That is a same-index diagonal between the messages, which continues while the register constants agree and stops at the boundary of the overlap of the two top passes. If α_C is an endpoint, the accident is an adjacent alias $P_x \to c(P_{x+1})$. For a target other than the root it gives $\alpha_{c(P_{x+1})} \oplus \alpha_{c(P_{x+2})} = \langle x \rangle \oplus \langle x+1 \rangle$, and the next consumer in the same message has counter $x-1$, so it does not repeat the equality. A parallel copy through a root shared with the other message, of pass length K, needs $x+2 = 5L$ and $5K-1 = (5L-3) \oplus (5L-2) \oplus (5L-1)$. Since $(n-3) \oplus (n-2) \oplus (n-1) \in \{n, n-4\}$ for every n, this gives $5K = 5L+1$ or $5K = 5L-3$, which is impossible. For an adjacent root alias, $x = 5L-1$, with root class G, the equation is $\alpha_G = h_P \oplus \langle 5L-1 \rangle$, and the output at step $5L-2$ gets the constant chain half $h_P \oplus \langle 5L-1 \rangle \oplus \langle 5L-2 \rangle$, which a root of either message can match. If $L > 2$ this step is in pass 5, and a further chain identity would need two root indices that differ by 1 or 2, whereas both are multiples of 5. If $L = 2$ the step is $4L$, in pass 4, with register half β_G, and completing its comparison with a root pins β_G.

It remains to check pass 4. An output at a pass-4 step $t < 4L$ uses the α of an earlier pass-4 step, which is absent from a pass-5 equation, and after the stride pair the shared pass-4 class has consumers with counters $4a-1 \neq 4b-1$. The first pass-4 step $4L$ has the last pass-5 class $c(P_{4L+1})$ as its chain supplier. Its chain equality is implied only if the supplier variable is shared unchanged, which needs equal counters and is the diagonal case; or the supplier is pinned, which is the root alias with $L = 2$; or the supplier occurs in the alias equation of $P_{4L+1} \to P_{4L+2}$, whose offset is $\langle 4L+1 \rangle \oplus \langle 4L+2 \rangle = 3$. Two first pass-4 outputs differ by $\langle 4a \rangle \oplus \langle 4b \rangle$, whose two low bits are zero, and

an output that reads the other endpoint would need the counter $4L \oplus 3 = 4L+3$, while its consumer has counter $4L+2$ and a shared root would need $L = 3$ and $5K - 1 = 15$. This proves the claim.

Rank at most one. With no equation there is no final comparison that holds. After one accident, pass-4 classes are separate from pass-5 classes, the only shared pass-4 pair is the stride pair, and there is no register equation. A joining output compared with a backward output, in either role, has a register equality that is not implied, since one supplier is in pass 4 and the other is a constant or in pass 5. Two joining outputs have the same register supplier only through the stride pair, and then both are the determined outputs. A free tape always needs a new chain equation, and $F_T = F_R$ needs $\alpha_{c(T_{3b+1})} \oplus \alpha_{c(R_{3a+1})} = \langle 3b \rangle \oplus \langle 3a \rangle$ on fresh pass-4 suppliers that are not those of the stride pair. So no bad event has rank at most one, and a rank-two sequence whose second accident is final and adds one equation is exactly the chain $T_{5b} \sim R_{5a}$, $T_{4b} \to R_{4a}$, $F_T = F_R$.

Two accidents, then an induced final comparison. The first equation is a chain equation. Suppose the second is too. With no register equation every backward class lies inside pass 5 or inside pass 4, so an induced final comparison involves no free tape and no backward output, and only $F_T = F_R$ remains. If R_{3a+1} is fresh, its α is in no earlier output of R and in neither equation, because before the second accident the claim allows no hit on R_{3a+1}. Later hits copy the variable without constraining it, so $F_T = F_R$ compares unrelated variables, or the same variable with counters $3a \neq 3b$. If R_{3a+1} is old, both accidents lie inside R and R_{3a+1} is the second, with a target R_y whose α is in neither equation unless $y = 3a + 2$. In that corner the register identity forces $R_{4a+1} \sim R_{4a+2}$, so the accidents are $R_{4a+1} \to R_{4a+2}$ and $R_{3a+1} \to R_{3a+2}$. For $a \geq 3$ the root of T can share only the root of R, and its continuation is the parallel copy excluded above. For $a = 2$ the first equation pins the class of R_9 and R_{10}, and a root of T that enters it continues onto the root output only if $5b - 1 = 9$, that is $b = a$. In both cases the next step of T is fresh, and fresh continuation excludes $F_T = F_R$.

Suppose instead that the second equation is a register equation. Its chain equality is implied, so by the claim the prefix is a full pin, which is excluded, or it extends a same-index diagonal of pass 5 across the first pass-4 step of the longer message P. With $m = \max(a, b)$, it identifies step $4m$ of P with step $4m$ of the shorter message and pins $\beta_{c(P_{5m})}$ to a register constant. The next pair of the diagonal, at $4m - 1$, needs a register equation on $\beta_{c(P_{5m-1})}$, a different class: the diagonal begins no higher than $5\min(a, b) \leq 5m - 5$, so the one chain accident cannot also alias P_{5m-1} with the root. The α of the class just hit is outside the chain equation, and its only matching output has the same counter $4m - 1$, so the next step has no other induced target, and it does not meet F_R, whose chain supplier R_{3a+1} is another, untouched pass-4 class. The next step is fresh, and fresh continuation applies.

The count. A single final accident of rank two is one of the direct comparisons: the joining outputs of T against the backward and joining outputs of R, and the backward outputs of T against the joining outputs of R, $b(2a+a)+(2b)a = 5ab$ in all. A single accident of rank two that is not final is charged as a prefix; inside one message it belongs to H_R or H_T, and between the messages there are at most $(2b)(2a) = 4ab$. Two rank-one accidents of which the second is final give the one chain, and two rank-one accidents followed by an induced final comparison give nothing outside the full pins. The total is $9ab + 1$. □

Appendix D.
The Case Analysis for KWP

This appendix proves Lemmas 20 and 21. A *W-accident* is an accident whose target is not exposed, and every rank-two sequence costs at most $1/M$, by the argument of Lemma 12.

TABLE 6. RANK-TWO SEQUENCES FOR FRESHNESS WITH EXPOSED POINTS, FOR A SINGLE MESSAGE AND FOR A PAIR (COEFFICIENT OF $L_R L_T$).

Family	Single	Pair
No visit	$18L^2$	40
Final BadVisit, no target factor	$128L^2$	–
Final BadVisit, exposed target	$18Lp$	–
Two accidents that are pins	–	64
Pin, then W-accident with a register identity	–	96
Pin, then W-accident with a register constant	$24L^2$	224
Pin, then W-accident, first merge above pass 2	–	48
Pin, induced merge, same-index W-accident	–	96
Pin, induced merge, boundary accident after pin	–	768
Induced merge, boundary accident before the pin	–	384
Coefficient of L^2 / of $L_R L_T$	170	1720

Proof of Lemma 20. The accident sequences of Section 5 apply, with the exposed classes held fixed as constants, and we stop each replay at its first BadVisit occurrence or at its first old step in passes 3–6. A single equation produces no event. Two distinct exposed points cannot supply the same full output, since the exposed map is injective, so a step above pass 2 cannot be matched for free, and the only rank-one route to a second cluster is the at-risk branch that Lemma 19 charges.

At rank two, sequences with no visit are the forward ones, $18L^2$ for one message and $40L_R L_T$ for a pair, and a final BadVisit costs at most $128L^2+18Lp$ sequences over one message. Sequences with a visit that survive BadVisit fall into the families of Table 6. The hard family is one pin $\gamma(c_0) = \lambda$ and one further chain equation. When the chain equations have rank exactly one, a class has occurrences at two different indices only near a visit: a recurrence at index y needs an exposed occurrence at $y-1$ or $y-2$. Let $P_0 = (R_k, T_m)$ be the first merge involving a step above pass 2. Its register-parent pair is merged before it, and, with no register equation, each single-variable right half can equal only the same parent variable. Consider the linear map sending a class to its register-parent class. It is nilpotent, because following register parents decreases the index. The reduced γ-support of every merged pair is 0 or that of the second equation w. Since a nilpotent map cannot fix w, the parent pair has support 0.

If the W-accident is P_0 itself, this gives a shared pin and a same-index accident. If P_0 is induced, then $k = 2L_R + 1$ and $m = 2L_T + 1$, and the chain of register parents gives $L_R = L_T$ and equal first steps, so there is no sequence inside one message. Take the last diagonal run ending at $2L+1$. It starts after $L+1$, since otherwise the largest unmerged index in the first pass would give two distinct exposed points with equal outputs, contradicting injectivity. At the bottom of the run the two predecessor classes are distinct and not exposed, and the equation is $\gamma(U) = \gamma(V)$. Matching support and constants forces the W-accident to be a same-index accident between the messages, or to have an index-1 or cluster-exit side. The counts for a fixed pair are $96L^2$, $768L^2$ and $384L^2$.

At rank three or more, take a replay with N occurrences of steps of W and put $A_N = N^2+8N$. Before BadVisit every accident has at most A_N choices. An accident whose target is a step of W has at most N^2 choices, and a hit on an exposed point starts a first-pass cluster whose right half is a fixed input semiblock, which leaves at most eight points; visits inside the cluster are induced

TABLE 7. LONG UNWRAP WITH STEP 1 OLD.

Family	Cost
No fresh step	$18/B + 2LD/M$
Double hit	$12LD/M + 72L^2/M$
Step 1 a P-hit with a fresh parent	$36/B$
Two paid hits on P	$10404L^2/M$
One fresh step in the top pass, with a self-hit and a paid hit on P	$17L^2/M$
Self-hit with a fresh and a paid-hit parent	$2040L^2/M$
Step 1 a self-hit on a legal fresh step, with a fresh parent	$9/B + 96L^2/M$
Self-hit, hidden comparison above	$19845L^2/M$
Exposed top, one paid hit on a target in P	$331L/M + 73728\,p/M$
Exposed top, one paid hit on a self-target	$10L^2/M + 18000\,p/M$
Two paid hits, one self-target	$2040L^2/M$
Two paid hits, targets fixed by the fresh runs	$818L^2/M$
Three paid hits, or two with a third payment	$7 \cdot 10^6 L^6/B^3$

lookups. The first prefix reaching rank three has at most three accidents, so its probability is at most $(A_N + A_N^2 + A_N^3)/B^3 \le 5N^6/B^3$ for $N \ge 12$. Summing over messages and pairs, with $N = 6L_T$ or $N = 6(L_R + L_T)$, gives $5 \cdot 6^6[\sum_T L_T^6 + \sum_{R<T}(L_R + L_T)^6]/B^3$. A first BadVisit whose last visit has up to p targets adds at most $[5N^6 + 3pN^5]/B^3$, summed over single messages only. Bounding each sum by $q\sigma L_{\max}^5$ gives the last term, with no factor of p. □

Proof of Lemma 21. All probabilities retain the expectation over ideal transcripts and Γ, and the indicator of $\mathcal{G}$. We write $\rho = B/(M - D - 6L)$ for the cost of fixing one half of a fresh value. A fresh step 1 supplies one of eight legal first semiblocks, at cost $8\rho \le 9/B$. Suppose step 1 is old, and count paid hits; Table 7 lists the families and their costs. A basic ingredient is that the right half of a pass-6 point is visible but its left half is not. So a request whose Y is fixed by τ matches a created point at a cost of at most $17/B + D/M$: at most μ pass-6 points share the requested right half, each costing $1/B$ for its hidden left half, and any other created point costs $1/M$. We call such a match a *hidden comparison*.

If no step is fresh, the computation is a walk through P. A computation made entirely of exposed points is invalid, since an exposed first input carries a length field $m \le 8$, which is not legal for a long message, and any other all-old chain pays a hidden comparison or a legality check on an accidentally legal point, at most $18/B + 2LD/M$ in all. A hit with both parents fresh, a *double hit*, has both requested halves fresh, so the union over positions and targets costs $12LD/M + 72L^2/M$. Two paid hits on points of P cost $(\mu\rho)^2$ per pair of positions, up to a density factor below 2 that accounts for drawing without replacement, and $\mu = 17$ gives $10404L^2/M$.

Let F be the set of fresh steps. Fresh intervals begin at an index at least 2, and each interval's first index minus one is a paid hit. If $f > L$ is fresh and $f - L$ is a hit, then $f - L$ is a paid hit. So with exactly one paid hit, F is one interval in pass 1, and the steps above it are hits. The upper chain is then a walk through old points, and at least $4L + 1$ of its steps form an exposed prefix unless a hidden comparison occurs. An initial run of steps on exposed points of length at

least $4l+1$ needs at least $\lceil l/2 \rceil$ distinct exposed answers with γ in a set of at most $4l+30$ format values, since one exposed answer serves at most eight counters. With $x = p/[B(1-p/M)] \le 1/31$ this gives $\sum_{l \ge 2} l \Pr[\text{prefix} \ge 4l+1] \le 4096\, p/B$. This prefix bound combines with the single payment, so no standalone p/B term appears, and gives $331L/M + 73728\, p/M$ for a payment on a target in P and $10L^2/M + 18000\, p/M$ for one on a self-target.

With exactly two paid hits, F is one of: a single step above pass 1; an interval $[a, L+1]$; two intervals in pass 1; or an interval in pass 1 together with $[c, h_2]$ with $c \ge L+2$ and $h_2 \le 2L$. When a self-target is pinned by the selected positions, the number of targets drops and the cost is quadratic. In the last case every step above h_2 is a P-hit, and each hit in the gap has a right half fixed by the top run, which yields a third payment through a copied α. There are four templates for this payment: an upper self-target above c; an upper self-target at c with more fresh steps above; a lower self-target with a target in P above; and a singleton upper interval. Each reveals a copied left half whose first true comparison fixes it, and two comparisons on the same copied value would need two different counters to agree, so no equation is paid for twice. In the last two templates three input variables are fixed in turn. With a and c the starts of the two fresh intervals and t_1 the fresh target of the self-hit at $a-1$, the hit at $c-1$ fixes α_c among at most μ candidates, the hit at $a-1$ fixes α_a, and the first true comparison that involves α_{t_1} fixes it among at most $6L(6L+\mu)$ values. The three equations are triangular. The first two templates are similar, and Section I.6 gives all four cases.

With three or more paid hits, each has a private half, α_{u+1} or β_{u+L}, used for the first time at that hit u. So for positions $u_1 > u_2 > u_3$, each paid hit costs at most Q/B, with $Q = 6L+\mu$, up to the density factor $(M/(M-D-6L))^{6L} < 2$, and $2(6L)^3 Q^3/B^3 \le 1317006\, L^6/B^3$. The four two-paid-hit templates obey the same count, for a total of $6585030\, L^6/B^3 \le 7 \cdot 10^6 L^6/B^3$. □

Appendix E.
Complete Proof of Theorem 1

This appendix gives the proof of Theorem 1 with every step written out. Section 5 states the lemmas and gives short proofs; the numbering of those lemmas is kept, and the new statements here are auxiliary. The forward pair analysis, summarized in Appendix A, is written out in Appendix F. Throughout, $B = 2^{64}$, $M = B^2$ and $\rho_N = B/(M-N)$.

E.1. Links

Fix a message T of pass length L, with $6L$ steps. In the unrolled view of Section 2, step t reads the two *links* a_{t-1} and d_{t-L} and writes the two links a_t and d_t: its input is (a_{t-1}, d_{t-L}) and its output is $(a_t \oplus \langle t \rangle, d_t)$. The constant links are $a_0 = X_A$ and $d_{j-L} = X_j$ at the bottom and $a_{6L} = C_A$ and $d_{5L+j} = C_j$ at the top, $1 \le j \le L$. The chain links a_t and the register links d_t never meet in one equation.

Every non-constant link is created by exactly one step of the process of Section 5, which we call its *supplier*:

- a fresh forward step $t \le 2L$ draws its output and creates a_t and d_t;
- a fresh backward step $t \ge 3L+1$ draws its input and creates a_{t-1} and d_{t-L};
- a free joining step $2L+j$, $j < L$, draws g_j and creates $a_{2L+j} = g_j \oplus \langle 2L+j \rangle$, while its register link d_{2L+j} is the β-half of backward step $3L+j$;
- the determined joining step $3L$ creates nothing: its output $(a_{3L} \oplus \langle 3L \rangle, d_{3L})$ has $a_{3L} = \alpha(B_{3L+1})$ and $d_{3L} = \beta(B_{4L})$, where B_t denotes the point of backward step t.

An old step creates nothing; the links it writes (forward) or reads (backward) are halves of the point it hits. The suppliers of the halves of each computed side are:

- the input of forward step x: chain half from step $x-1$, or the constant X_A if $x=1$; register half from step $x-L$, or the constant X_x if $x \le L$;
- the input of joining step $2L+j$: chain half from step $2L$ if $j=1$, and the tape value g_{j-1} otherwise; register half from step $L+j$;
- the output of backward step x: chain half from step $x+1$, or C_A if $x=6L$; register half from step $x+L$, or C_{x-5L} if $x>5L$;
- the determined output: chain half from step $3L+1$, register half from step $4L$.

The supplied value is a variable of the supplier's class only if the supplier was fresh. If the supplier was old, the value is the corresponding half of the class it hit.

E.2. The process and the ratio

The process $\Gamma(\tau)$ is that of Section 5. We write out the proofs of Lemmas 7, 8 and 9.

Proof of Lemma 7. Call a draw sequence of Γ *good* if it has no bad event. Along a good sequence the defined points form a partial injection. Fresh forward outputs and fresh backward inputs are drawn from unused values, and fresh steps have unused arguments by definition. Joining inputs are unused because there is no input clash, and joining outputs, free or determined, are unused because there is no output clash. The point set contains, for every message T, all $6L_T$ points of the computation of W on X^T, and this computation ends in C^T. So every permutation that contains the point set agrees with τ.

Two distinct good draw sequences first differ at a draw made at the same step with the same history, hence with the same argument, and they assign it different values. So the events "π contains the point set of the sequence" are disjoint for distinct good sequences.

Fix a good point set and list its points in creation order; let D_j be the number of points created before the jth one. A uniform permutation contains the set with probability $\prod_j 1/(M-D_j)$. The process produces it with probability $\prod_j 1/(M-D_j)$ over the fresh forward and backward steps, times $1/B$ for each free joining step, times 1 for each determined step. The ratio of the first probability to the second is

$$\prod_{\text{free } h} \frac{B}{M-D_h} \prod_{\text{determined } h} \frac{1}{M-D_h} \ge \prod_T \Big(\frac{B}{M}\Big)^{L_T-1} \frac{1}{M} = \prod_T B^{-(L_T+1)}.$$

Summing over good point sets gives $\Pr_{\text{real}}[\tau] \ge \prod_T B^{-(L_T+1)} \Pr_\Gamma[\mathsf{Good}]$. In the ideal world, $(N)_k = N^k \prod_{i<k}(1-i/N) \ge N^k(1-k^2/(2N))$ with $N=2^{64n}$, and $n \ge 3$, give

$$\Pr_{\text{ideal}}[\tau] \le \frac{\prod_T B^{-(L_T+1)}}{1-q^2/2^{193}}.$$

Hence $\Pr_{\text{real}}[\tau]/\Pr_{\text{ideal}}[\tau] \ge (1-q^2/2^{193})(1-\Pr_\Gamma[\neg\mathsf{Good}]) \ge 1-\phi(\tau)$. □

No term for the number of used outputs appears in the ratio, because a free joining step draws its left half from all B values and an output that is already used is a bad event rather than an excluded value.

Proof of Lemma 8. Let h be the free joining step $2L+j$ of T, $j < L$, and let Reach_h be the event that h is reached before any bad event. Given the history up to h, the proposed output (g, d) with $d = d_{2L+j}$ is used with probability exactly U_h/B, where U_h counts the used outputs with right half d. Since the process stops at its first bad event,

$$\Pr_\Gamma[\mathsf{Clash}] = \frac{1}{B} \sum_h \mathbb{E}_\Gamma\big[1_{\mathsf{Reach}_h} U_h\big].$$

Fix h. Its register supplier, backward step $3L + j$ of T, lies in pass 4. On Reach_h this step was fresh, since an old pass-4 step is a bad event. So d is the β-half of an input drawn uniformly from at least $M - S$ unused inputs. Between that draw and h the process runs only the backward steps $3L + j - 1, \ldots, 3L + 1$ of T and its joining steps $2L + 1, \ldots, 2L + j - 1$, because messages are processed one at a time. Every output that is used when h is reached falls into one of two cases.

- Its right half was determined before the draw of d. This covers all outputs of earlier messages, the outputs of the forward steps of T and of its backward steps above $3L + j$, and also the outputs of the backward steps $t \in [3L+1, 3L+j-1]$ of T: their right halves are d_t, the β-halves of backward steps $t + L \geq 4L + 1$, processed before step $3L + j$. Such an output has right half d with probability at most ρ_S over the draw of d.
- It is an earlier joining output of T, at a step $2L + j'$ with $j' < j$. Its right half is the β-half of backward step $3L + j'$, a fresh draw made after the draw of d. It equals the fixed value d with probability at most ρ_S over that later draw.

The event Reach_h is only used as an indicator that we drop when bounding from above; we never condition on later success. There are at most S used outputs, so $\mathbb{E}[1_{\mathsf{Reach}_h} U_h] \leq S\rho_S$. There are at most ℓ free joining steps, and $\rho_S/B = 1/(M - S)$, which gives the claim. □

The bound is stated for the stopped process. Without the stopping rule it fails: an old pass-4 step can copy the β-half of a fresh point into a later joining output, and the copied value is no longer a fresh draw. That configuration is a bad event, and the process has stopped before it matters.

Proof of Lemma 9. Run Γ and Ψ side by side up to the first bad event and couple each pair of corresponding draws maximally. A free joining draw is uniform on all B values in both processes, so the two agree. A fresh forward draw of Ψ at input x is $\pi(x)$, which is uniform on the complement of the outputs of π revealed so far; the corresponding draw of Γ is uniform on the complement of all used outputs. The two complements differ only by the joining outputs, at most ℓ values, and both have at least $M - S$ elements, so the total variation distance is at most $\ell/(M - S)$. A fresh backward draw is the same with inputs in place of outputs. There are at most S fresh forward and backward draws, so the coupling fails before the first bad event with probability at most $S\ell/(M - S)$.

When the coupling succeeds, Ψ has created exactly the points of Γ up to the first bad event, so Dup does not occur, and the first bad event is the same in both runs. Therefore

$$\Pr_\Gamma[\neg\mathsf{Good} \wedge \neg\mathsf{Clash}] \leq \Pr_\Psi[\neg\mathsf{Good} \wedge \neg\mathsf{Clash} \wedge \neg\mathsf{Dup}] + \frac{S\ell}{M - S},$$

and adding $\Pr_\Gamma[\mathsf{Clash}]$ gives the lemma. □

E.3. Paths and the exact reduction

The paths and the events $\mathsf{Bad}_{R,T}$ and Bad_T are those of Section 5. In the notation of the links, the bad events on paths are as follows; B_u^R is the point of backward step u of R and J_j^R the joining point $2L_R + j$.

- Input events: the input of a joining step of T equals the input of any point of R; or the input of a forward step of T equals the input of some J_j^R.
- Output events: the output of a backward step of T in pass 4 equals the output of any point of R; or the output of any backward step of T equals the output of some B_u^R with $u \in [3L_R + 1, 4L_R]$ or of some J_j^R.
- The determined output of T equals the output of any point of R.

A forward step of T whose input equals the input of a pass-4 point of R is not an event: it reads π at that point and stays on its path.

The proof of Lemma 11 in Section 5 is complete as given; we add the observation on which it rests. In Ψ, an old forward or backward step that hits a point created by a fresh forward or backward step reads the value of π at that point, which is the value its path reads. Only joining points carry values that are not values of π, and a joining step is processed after all forward and backward steps of its own message. So a message leaves its path only by hitting a joining point of an earlier message, which is an input or output event, or at one of its own joining steps, which is an input clash or an output clash. Each event is a statement about the paths of at most two messages, and no third message enters.

E.4. The replay and the cost of a sequence

For a pair (R, T) the replay processes the steps of the two paths in the order of Section 5: the forward steps of R, its backward steps from $6L_R$ down, its joining steps, and then the same for T. For a single message it processes that message alone. Let N be the number of steps, $N = 6(L_R + L_T)$ or $N = 6L_T$.

Every half that the replay handles is a variable xor a constant, or a constant. The variables are the output halves (γ, δ) of the class created by a fresh forward step, the input halves (α, β) of the class created by a fresh backward step, and the tape value of a free joining step. Chain variables are the γ, α and tape variables, and register variables are the δ and β variables. Joining points are classes of their own and are never merged with another class.

At step x the replay compares x's computed side with every existing class: the input for a forward step and for a joining step, the output for a backward step, and both for the determined joining step. With E the current set of half-equations, a comparison is induced if both half-equations follow from E and the identities, impossible if one contradicts E, and open otherwise. A half-equation is an *identity* if it holds syntactically: the same constant, or the same variable with the same constant. A true open comparison is an accident; its half-equations that do not follow from E are added to E. Since chain and register equations use disjoint variables, an accident raises $\mathrm{rank}(E)$ by one or by two.

The final comparisons for a pair are: every comparison at a joining step of T; a comparison of the input of a forward step of T with the input of a joining point of R; every output comparison at a pass-4 step of T; and a comparison of the output of a backward step of T with the output of a pass-4 point or a joining point of R. For a single message they are the comparisons of Bad_T. A true final comparison is a bad event on paths, and the accident sequence ends there.

Proof of Lemma 12. (a) *The symbolic replay.* Given w and the transcript constants, the replay is a deterministic symbolic object. It says, for every comparison made before the last accident of w, whether that comparison is induced, impossible, open and false, or an accident. So it fixes the classes and the list $g_1, \dots, g_m$ of draws made before the last accident of w, in processing order. A draw is the creation of a class by a fresh forward step (variables γ, δ), by a fresh backward step (variables α, β), or a free joining step (one tape variable). Joining steps query no permutation.

(b) *The basis.* $E(w)$ is an affine system over $\{0,1\}^{64}$ whose coefficients on the variables are in $\{0,1\}$. It splits into its chain part and its register part. Order the variables by draw and reduce each part separately to echelon form, so that every basis equation has its latest variable as pivot and distinct basis equations have distinct pivots. If $E(w)$ is inconsistent the event is empty. Let P_i be the number of pivots in draw g_i, so that $\sum_i P_i = \operatorname{rank}(w)$. The two variables of a permutation draw lie on different lines, so a basis equation with its pivot in g_i contains no other variable of g_i, and all its other variables belong to $g_1, \dots, g_{i-1}$. Also $P_i \le 2$, and $P_i \le 1$ for a tape draw.

(c) *The events.* Let A_i be the event that the execution agrees with the symbolic replay on every comparison made before draw g_{i+1}, and, for $i = m$, on every comparison up to and including the last accident of w. Let C_i be the event that A_i holds and every basis equation with its pivot in $g_1, \dots, g_i$ holds for the actual values. Both are determined by the history through g_i, that is, by the revealed points of π and the tape values through g_i, and $C_i \subseteq C_{i-1}$. If the accident sequence starts with w, every accident of w is a true comparison, so all equations of $E(w)$ hold, and so does every basis equation, being a combination of them. Thus $\{\text{the sequence starts with } w\} \subseteq A_m \subseteq C_m$. The conditions that are not pivots are bounded by 1 below.

(d) *One draw.* Fix a history in C_{i-1} with $D < N$ revealed points. If g_i is a permutation draw whose argument, the input of a forward step or the output of a backward step, is already revealed, then A_i fails. Otherwise, given the history, π is uniform over the extensions of the revealed partial injection, because the queries made so far are functions of the earlier answers and the tape. So the new answer is uniform on at least $M - D$ values. Each basis equation with its pivot in g_i fixes that half to a value determined by the history, and at most B^{2-P_i} of the 128-bit values satisfy all of them. Hence

$$\Pr[C_i \mid \text{history}] \le \begin{cases} B/(M-D) \le \rho_N, & P_i = 1, \\ 1/(M-D) \le \rho_N^2, & P_i = 2, \end{cases}$$

using $\rho_N^2 = M/(M-N)^2 \ge 1/(M-D)$. For $P_i = 0$ we use the bound 1, and a tape draw with $P_i = 1$ has probability $1/B \le \rho_N$.

(e) By the chain rule, $\Pr[C_m] \le \prod_i \rho_N^{P_i} = \rho_N^{\operatorname{rank}(w)}$.

For the second claim, every accident raises the rank by at least one, so the shortest prefix of rank at least 3 has at most three accidents. An accident is a step, a class and a kind of comparison; the second kind occurs only at the determined step, so there are at most $N^2 + 2N$ choices. The union bound over these prefixes gives

$$\Pr[\text{rank} \ge 3] \le \big(A + A^2 + A^3\big)\rho_N^3, \qquad A = N^2 + 2N.$$

For $N \ge 12$, $A^3 \le (7/6)^3 N^6 < 1.59 N^6$ and $A + A^2 < 0.02 N^6$, so the bracket is below $2N^6$. □

E.5. One equation is not enough

We write out the proof of Lemma 13.

Between the two sides. Before any accident the forward variables and the backward variables are disjoint sets. The input of a forward step compared with the input of a backward point has

chain halves $\gamma \oplus k$ or a constant against a variable α, and register halves δ or a constant against a variable β; neither half-equation is an identity, so the comparison has rank two. The same holds for the output of a backward step against the output of a forward point. So a first accident of rank one is not between a forward and a backward step.

The forward side. Accidents between forward steps of rank one are, by Lemma 26 of Appendix F, a prefix extension, two pass-1 steps with equal constants, or a stride hit, and by Lemma 27 their induced consequences stay inside passes 1–2 of both messages. A joining input uses the forward register links d_{L+j} and either the forward chain link a_{2L} ($j = 1$) or a tape value. Replace, in the forward analysis, the chain variables of the steps beyond pass 2 by independent tape values. This introduces no identity, because a tape value occurs in exactly one joining input. So no input event occurs with one forward accident.

The backward side, equal lengths. Put $r = 6L + 1 - t$. A backward step at index t becomes a step at reverse index r that reads its output and computes its input with π^{-1}. The counter at reverse index r is $\langle 6L+1-r\rangle$, which is injective in r and the same for both messages when their lengths agree. The top constants become bottom constants, and the register stride is still L. The forward analysis uses only these properties, so it applies to the reversed backward side, and the consequences of one equation stay in reversed passes 1–2, that is, out of pass 4.

The backward side, unequal lengths. Let $a = L_R \neq b = L_T$. With no equation, the only shared point is the last one, $T_{6b} \sim R_{6a}$, when the two top outputs agree; its immediate backward continuation compares chain halves $\alpha \oplus \langle 6b - 1\rangle$ and $\alpha \oplus \langle 6a - 1\rangle$ of the same variable, which is impossible. A rank-one accident with equal top register constants has both occurrences in pass 6. An accident between steps of different indices has no induced chain continuation, because the chain identity needs equal counters. A same-index accident can start a diagonal only inside the overlap of the two top passes, and at the boundary of that overlap the next register comparison sets a variable equal to a constant, which needs a new equation. The stride successor $T_{5b} \to R_{5a}$ of the shared last point has a register identity and costs one chain equation, and its own stride successor $T_{4b} \to R_{4a}$ costs another. So no pass-4 event occurs at rank one. Internal accidents in the top passes and their copies through a shared last point obey the same restriction. This includes top passes that overlap, as for $(a, b) = (6, 7)$.

Both sides together. With one forward accident, the backward input variables are unconstrained; with one backward accident, the forward output variables and the tapes are unconstrained. One equation can fix one half of an equality between the two sides, but not both. Finally the determined output is $(\alpha(B_{3L+1}) \oplus \langle 3L\rangle,\ \beta(B_{4L}))$, whose two halves come from the distinct pass-4 points B_{3L+1} and B_{4L} ($L \geq 2$). Neither half has been identified with anything by a rank-one closure, so its comparison with another output is not induced at rank at most one. This proves Lemma 13.

E.6. The untouched supplier and its consequences

We prove Lemma 14 in the following precise form. "Untouched" refers to true comparisons, accidents or induced; a comparison that turns out false adds nothing to E.

Proof of Lemma 14. Algebra. Let $h = v \oplus k$ be the half of x's computed side that contains v, and h' the corresponding half of the compared class. If v occurs neither in E nor in h', then $h = h'$ is not implied by E: its coefficient at v is one, while that coefficient is zero in every equation of E.

Occurrences. We show the two exclusions. v is not in E, because before its draw it did not exist and afterwards no true comparison had it in a compared half. For h', consider the kinds of supplier.

- A fresh forward supplier y. Its γ is read by step $y+1$ and its δ by step $y+L$. A later forward step that hits $c(y)$ copies the output of $c(y)$ without comparing it, but the steps that read the copy come later in processing order than the native reader. Indeed $y+1$ is processed right after y when $y < 2L$, and when $y = 2L$ only backward steps of the same message lie between; and a copy at a step z between y and $y+L$ passes its δ to $z+L > y+L$. A backward step that hits $c(y)$ compares its output, which contains v, and so touches v.
- A fresh backward supplier y is the mirror image: its α is read by $y-1$ and its β by $y-L$, in decreasing processing order, a later backward hit on $c(y)$ copies its input to readers that come later, and a forward hit on $c(y)$ compares its input and touches v.
- A tape value g_j is read only by the next joining step, which is processed right after the step that drew it.
- The two suppliers of the determined output are read by that output.

So before x, v cannot appear in the compared side of another class without a true comparison that touches it, and $v \notin h'$. □

The lemma is about the comparisons made at the reader x itself. It does not apply to a later comparison against a point of R that already contains v; those targets need the separate argument of the second consequence below.

Before the first bad event every pass-4 step of the message being processed is fresh. This gives the three consequences used in Section 5.

Lemma 22. *Every true comparison of the determined output of T with an existing output adds exactly two equations. Hence a rank-two sequence ending at such a comparison is that single accident.*

Proof. Write $A_T = \alpha(B_{3L+1})$ and $D_T = \beta(B_{4L})$. A_T is drawn at the last backward step $3L+1$. After it come only joining steps, and a true comparison at a joining step is a bad event, so before the comparison at the determined step A_T occurs in no equation. D_T is drawn at backward step $4L$; every later backward step of T is in pass 4 and is fresh, and a true comparison at a pass-4 step is a bad event, so D_T occurs in no equation either. Neither variable occurs in a competing output. Earlier backward outputs of T use chain suppliers with indices at least $3L+2$ and register suppliers with indices at least $4L+1$; earlier joining outputs of T use tapes and register suppliers $3L+j < 4L$; outputs of earlier messages and of the forward steps of T were created before these draws. So each half-equation of the comparison contains a variable absent from E and from the other side. The two halves lie on different lines, so the rank grows by two. □

Lemma 23. *Every rank-two sequence that ends with a comparison of an output of T against a joining output of R is a single accident.*

Proof. Put $a = L_R$. Before the first bad event of R, every pass-4 step $i \in [3a+1, 4a]$ of R is fresh; write its input as (α_i, β_i). In the path of R these variables occur in the input of B_i^R and in one output: α_i in the chain half of the output of step $i-1$, and β_i in the register half of the output of step $i-a$. That output is always a pass-4 output or a joining output of R, so a backward comparison of T against it is final. Points of R created earlier do not contain these variables, and the later pass-4 steps of R are fresh. So, before the first bad event of the pair, T cannot reach α_i or β_i by a backward hit on B_i^R, which is final, nor through their native output, which is final as well, nor by a forward hit on a joining point of R, which is final. The first access by T that is not final must be a forward hit on the input of B_i^R, and it adds two equations. After it, a different

pass-4 point of R still has an untouched β, whose native output is a joining output, so the first hit cannot bring it in through its far side. Hence a prefix of rank two that is not final reaches at most one pass-4 input point of R.

The determined output of R is $(\alpha_{3a+1} \oplus \langle 3a \rangle, \beta_{4a})$, whose halves belong to two different pass-4 points ($a \geq 2$), and whose native output is that joining output itself. With preceding rank at most one, both halves are untouched and the comparison adds two equations. With preceding rank two, at least one is untouched and the comparison adds at least one more. A free joining output of R is (g_j^R, β_{3a+j}); its tape value is new in any first final comparison, and with preceding rank at most one its β is untouched as well. In all cases a rank-two sequence ending there is one accident. $\square$

Comparisons with a private variable. By Lemma 14, the following comparisons always contain a variable that is absent from E and from the compared class, and so they add at least one equation: a joining input of T at $j \geq 2$, which contains g_{j-1}; an input comparison against the joining input of R at $j \geq 2$, which contains R's tape value g_{j-1}^R; an output comparison against a free joining output of R, which contains g_j^R; and an output comparison at a pass-4 step $t < 4L$ of T, whose chain half contains $\alpha(B_{t+1})$ of the fresh pass-4 step $t+1$. Such a comparison cannot end a sequence whose rank is already two.

This leaves the four interfaces named in Section 5. With $a = L_R$ and $b = L_T$, their halves are:

- the first joining input of T: $\gamma(c(T_{2b})) \oplus \langle 2b \rangle$ and $\delta(c(T_{b+1}))$;
- the output of T_{4b}: $\alpha(c(T_{4b+1})) \oplus \langle 4b \rangle$ and $\beta(c(T_{5b}))$;
- the output of R_u, $3a < u \leq 4a$, as a target: $\alpha(c(R_{u+1})) \oplus \langle u \rangle$ and $\beta(c(R_{u+a}))$;
- the first joining input of R, met by a forward step of T: $\gamma(c(R_{2a})) \oplus \langle 2a \rangle$ and $\delta(c(R_{a+1}))$.

The two suppliers in each row are distinct because the lengths are at least 2. That alone does not settle these interfaces, since an earlier hit can make two links equal; they are covered by the counts of Sections E.7 and E.8.

E.7. One message

For a message P of pass length L, consider the replay of P alone, stopped at its first event of Bad_P. Let IR_P be the event that this replay has an accident of rank increment two, and let M_P be the event that its accident sequence starts with one of the prefixes

$$P_j \to P_1, \qquad P_{6L} \to c(P_{j-1}), \qquad 2 \leq j \leq L.$$

The event H_P of Section 5 is $\mathsf{Bad}_P \cup \mathsf{IR}_P \cup \mathsf{M}_P$.

Lemma 24. *If the first two accidents of the replay of P alone each have rank increment one and one of them is between a forward and a backward step, then they are $P_j \to P_1$ and $P_{6L} \to c(P_{j-1})$ for some $2 \leq j \leq L$. Their equations are $\gamma(c(P_{j-1})) = X_A \oplus \langle j-1 \rangle$ and $\delta(c(P_{j-1})) = C_L$, and they need $X_j = X_1$ and $C_A \oplus \langle 6L \rangle = X_A \oplus \langle j-1 \rangle$.*

Proof. Before the first hit between the two sides their variables are disjoint, so by Lemma 13 the first rank-one accident is between two forward steps in pass 1 or two backward steps in pass 6. If it is between backward steps, it constrains only backward chain variables, both free output halves of every forward class are unconstrained, and a backward hit on a forward class costs two equations. If it is between forward steps, the register line is still unconstrained, so the first backward hit on a forward class adds a register equation and its chain equation must already follow. Before that hit, the output of a backward step other than the root has a backward chain variable that is absent

from the forward equation, so the hit is at step $6L$, whose chain half is the constant $C_A \oplus \langle 6L \rangle$. The forward equation must therefore pin a single γ to a constant, which happens exactly when its target is step 1: $P_j \to P_1$ gives $\gamma(c(P_{j-1})) \oplus \langle j-1 \rangle = X_A$, and its register half needs $X_j = X_1$. Then $P_{6L} \to c(P_{j-1})$ adds $\delta(c(P_{j-1})) = C_L$. The argument keeps every occurrence of $c(P_1)$, including the case $j = 2$, where P_2 joins that class. □

Proof of Lemma 15. By Lemma 13, every sequence in H_P has rank at least two. Let w be the shortest prefix of the accident sequence of rank at least two at which H_P is decided. If $\mathrm{rank}(w) \geq 3$, Lemma 12 with $N = 6L$ charges it within $2(6L)^6\rho^3$. Otherwise $\mathrm{rank}(w) = 2$ and w is one of three kinds.

- A single accident of rank two. Any step can be compared with any earlier class, and the determined step also compares its output, so there are at most $\binom{6L}{2} + (6L-1) = 18L^2 + 3L - 1$ such prefixes. This covers IR_P and every event of Bad_P that is a single accident.
- Two accidents of rank one, one of them between the two sides. By Lemma 24 these are the at most $L-1$ prefixes of M_P. We charge the prefix and not a continuation, so every later event of this realization of P is covered once, however many later messages reuse its classes.
- Two accidents of rank one within the two sides. We show that they produce no event. After one accident inside one side there is no further induced hit inside that side, because the possible reuse through adjacent occurrences has unequal counters. So a second rank-one accident inside the sides adds another chain equation and has a register identity, and the register line stays unconstrained. Two forward accidents are covered by the single-message part of the forward analysis, Lemma 28(a), and two backward accidents by its reversal. With one forward and one backward accident, each side carries only the consequences of its one accident, which by Lemma 13 stay away from the joining pass and pass 4. A final comparison between the two sides then still has one half on each line that contains an unconstrained variable from the side without an accident on that line, so it needs two equations.

Each prefix of rank two has probability at most ρ^2 by Lemma 12. In total $18L^2+3L-1+(L-1) = 18L^2 + 4L - 2 \leq 20L^2$. □

The singleton charge is used through the deterministic inclusion

$$\bigcup_{R<T} \mathsf{Bad}_{R,T} \subseteq \bigcup_{P} H_P \;\cup\; \bigcup_{R<T} \big(\mathsf{Bad}_{R,T} \cap \neg H_R \cap \neg H_T\big),$$

which needs no independence. In particular a pair event that follows an accident of rank two inside R is charged to R once.

E.8. Pairs

Fix R before T, put $a = L_R$ and $b = L_T$, and exclude H_R and H_T. Consider a first bad event of the pair whose accident sequence has rank two; by Lemma 13 rank at most one is impossible. Call a hit *mixed* if it is a forward step hitting a class created by a backward step, or a backward step hitting a class created by a forward step. Exactly one of the following holds.

1) There is no mixed hit before the event. Then every accident is between two forward steps (including joining inputs), or between two backward steps (including outputs against

joining outputs), or the event is a direct comparison between a step of one side and a point of the other side.

2) There is a mixed accident.
3) There is a mixed hit, and every mixed hit is induced.

We bound each case by the rows of Table 3. In cases 2 and 3 we count prefixes of rank two, whether or not a bad event follows, and use

$$\Pr[\text{prefix } w \text{ and a later event}] \le \Pr[\text{the sequence starts with } w] \le \rho^2,$$

so their continuations need no classification.

Forward against forward: $40ab$**.** If all accidents and hits are between forward steps, the replay of the forward steps and the joining inputs is the forward pair analysis of Appendix F. That analysis bounds the sequences producing a first equality of an input of a step beyond pass 2 with another input, and the input events of the present process are such equalities at the joining pass. Replacing the forward chain variables beyond pass 2 by independent tape values preserves every input equation up to the first input event, because each tape value occurs in exactly one joining input and a comparison involving it is final. So Lemma 28(b) applies, with at most $40ab$ sequences.

Backward against backward, equal lengths: $40ab$**.** With $a = b$, the reversal $r = 6L + 1 - t$ of Section E.5 maps the backward steps to the steps of reversed passes 1–3, pass 4 to reversed pass 3, and the outputs of the joining points of R to inputs at the reversed joining pass. The output events become input events at reversed pass 3, and the same six families arise, at most $40ab$.

Backward against backward, unequal lengths: $8ab + 1$**.** Let $a \neq b$ and suppose there is no mixed hit. Joining outputs of R as targets are direct by Lemma 23. We first show:

> (B1) With at most one rank-one accident between backward steps, every backward hit has both occurrences in pass 6, except possibly $T_{5b} \to R_{5a}$, which is then the accident itself and follows the shared last point $T_{6b} \sim R_{6a}$.

A first rank-one accident between backward steps adds a chain equation, and it compares two pass-6 steps with equal register constants or is the stride successor $T_{5b} \to R_{5a}$; the register line stays unconstrained. Consider a first induced hit outside pass 6 after a pass-6 accident. A variable never equals a top register constant, so both occurrences lie outside pass 6. A supplier below the top boundary is fresh and untouched before this hit, so the only candidate pair is (T_{5b}, R_{5a}). Put $v = \alpha(c(R_{5a+1}))$. If v is absent from the equation, the comparison either adds an equation containing v or reads v on both sides and needs $v \oplus \langle 5b\rangle = v \oplus \langle 5a\rangle$, which is impossible. The equation contains v only through the internal accident $R_{5a+1} \to R_{5a+2}$, which makes the lowest pass-6 class its own chain supplier. Even then the boundary pair is not induced: for $a > 3$ the input α of the last point of R is absent from that equation, and after a shared last point the next step of T is fresh because its counter differs; for $a = 3$ and for $a = 2$ the equation touches that α, but a continuation onto the adjacent class needs $6b - 1 = 6a - 1$. In all cases, once T takes a fresh step, its next chain supplier belongs to T alone and an equation inside R does not constrain it. Finally, if the first accident is $T_{5b} \to R_{5a}$, the α on its far side is absent from its equation, its immediate backward continuation has unequal counters, and no further hit outside pass 6 is induced. This proves (B1).

Now take a rank-two sequence of backward accidents ending in a first bad event.

- *A single accident of rank two* is a direct final comparison: a pass-4 output of T against any of the $3a$ backward points of R, or any of the $3b$ backward outputs of T against a

pass-4 point of R, together $b(3a) + (3b)a - ba = 5ab$; or a backward output of T against one of the a joining outputs of R, $3ab$.

- *A final second accident of rank one.* By (B1) the first accident gives no access to the protected chain suppliers, so the chain equation of the second is new and its register equation must be an identity. The only sharing in pass 5 between the two messages is $T_{5b} \sim R_{5a}$, so the second accident is its stride successor $T_{4b} \to R_{4a}$, and the sequence is the chain $T_{6b} \sim R_{6a}$, $T_{5b} \to R_{5a}$, $T_{4b} \to R_{4a}$: one sequence.
- *An induced final comparison after rank two.* By the comparisons with a private variable above, an induced final comparison is $T_{4b} \to R_{4a}$. Suppose first that every pass-5 step of R is fresh, and put $v = \alpha(R_{4a+1})$. Before the first bad event its native output R_{4a} cannot be met, so the only way for T to read v is a backward hit on R_{4a+1}. By (B1) that hit is not in a rank-one closure, so it happens only after the budget is spent, and it does not constrain v, which lies on the far side of the hit. Then either the chain half of T_{4b} does not contain v, and the comparison adds an equation, or it does and needs $v \oplus \langle 4b \rangle = v \oplus \langle 4a \rangle$, which is impossible. If a pass-5 step of R is old, both accidents lie inside R, because below the top boundary the first old pass-5 step of R has a fresh chain supplier, and at the boundary the adjacent case treated in (B1) rules out an implied chain comparison. So the first is a pass-6 accident $R_j \to R_i$, $5a < j < i \leq 6a$, and the second its stride successor $R_{j-a} \to R_{i-a}$, which contains a fresh pass-5 chain variable and adds no relation among the pass-6 chain variables. So the replay of the top of T is the same as after one accident. At the first pass-5 step $5b$ of T, either the last step of T is fresh and its register half has no identity with a pass-5 output of R, or the last point of T is that of R and its targets are the stride children of that class. If that class repeats, the second accident merges exactly those two stride children, so the only candidate is $c(R_{5a})$, whose output is unchanged and whose chain comparison with T_{5b} involves only pass-6 variables; by (B1) it is not induced. So T_{5b} is fresh, every later backward step of T has a private fresh chain supplier, and no event follows.

A direct final hit can itself create a repeated occurrence, for instance by joining a class already reached at the last step of T; such sequences are counted among the direct ones. The total is at most $5ab + 3ab + 1 = 8ab + 1$, which is below $40ab$, so the uniform bound $40ab$ for backward against backward holds for all lengths.

Direct comparisons between the two sides: $8ab$**.** A single final accident between the sides is a joining input of T against the input of one of the $3a$ backward points of R ($3ab$), a pass-4 output of T against the output of one of the $2a$ forward points of R ($2ab$), or the determined output of T against the output of any of the $6a$ points of R ($6a$). Together this is $5ab + 6a \leq 8ab$, since $b \geq 2$. By Lemma 22 every sequence ending at the determined output is one of these. The count $5ab + 6a$ is attained: for $a = b = 2$ the 32 comparisons, 12 joining inputs of T against inputs of $B_7^R, \ldots, B_{12}^R$, 8 pass-4 outputs of T against outputs of $R_1, \ldots, R_4$, and 12 comparisons of the determined output, can each be the first event with two new equations.

Prefixes with a mixed accident: $14ab$**.** A single mixed accident of rank two is a forward step of T against one of the $3a$ backward points of R, or a backward step of T against one of the $2a$ forward points of R: at most $(2b)(3a) + (3b)(2a) = 12ab$ prefixes. Mixed accidents of rank two inside one message are charged to H_R or H_T. For two accidents of rank one with a mixed one among them, before the first mixed hit the two sides have disjoint variables. With one chain equation, a rank-one mixed accident needs a γ or α that is already pinned to a constant, and the other side must supply that constant at its root. With one register equation, the first accident must

be a prefix extension; it pins the register half of the shared first class of the forward side, and a pass-6 output of T can supply the matching register constant. A first accident between backward steps of T cannot prepare a rank-one hit of a backward step on a forward class, because neither free half of a forward class occurs in its equation and T has no forward step left. Up to the mixed prefixes of one message, which are M_R and M_T, the possibilities are:

First accident	Mixed accident	Count
$R_j \to R_1$	$T_{6b} \to c(R_{j-1})$	$a-1$
$R_j \to R_{6a}$	$T_1 \to c(R_{j+1})$	$a-1$
$T_1 \to R_i$ or $T_j \to R_1$	T_{6b} on the pinned class	$a+b-2$
prefix extension	a pass-6 step of T on the first class	b

Their sum is $3a + 2b - 4 \le 2ab$, so there are at most $14ab$ prefixes in this case.

Prefixes that pin a target fully: $49ab$**.** Suppose every mixed hit is induced and consider the first one. Before it, the equations of the two sides have disjoint variables, so both free halves of the target class must reduce to constants: a *full pin*, one chain pin and one register pin. A single accident of rank two inside one side cannot supply both pins to one class, because its chain and register suppliers are at indices that differ by $L - 1 \ge 1$ and are distinct in the free state. So there are two accidents of rank one. We use:

> $(*)$ If the chain rank is at most one, every class inside one side has at most two occurrences in each message.

Indeed, two occurrences in one message have different counters, so their chain equation is not an identity and connects their distinct chain suppliers; three occurrences would need chain rank two. This is Lemma 25 and its mirror. Across the pair a class has at most four occurrences.

If the chain pin comes first, it is a chain equation with a root as one endpoint: inside R ($a-1$ choices), inside T ($b-1$), the root of T against an occurrence of R ($a-1$), or an occurrence of T against the root of R ($b-1$), on either side, so at most $4(a+b-2)$ choices. Fix it, pinning the chain half of a class c. The second accident must pin the register half of the same class. By $(*)$, c has at most four occurrences in the rank-one closure, hence at most four steps that read its register half. Choose one such reader; its chain comparison must already be implied. For a given chain expression, at most six occurrences have an equal chain expression after the first pin: four readers of an unpinned variable and the two roots for a pinned variable or a constant. So there are at most $4 \cdot 6 = 24$ second accidents per first pin. This counts comparisons with every occurrence kept, and does not assume that induced comparisons are identities. The total is $24 \cdot 4(a+b-2) = 96(a+b-2) \le 48ab$, using $2(a+b-2) \le ab$ for $a, b \ge 2$.

If the register pin comes first, the free state leaves only the prefix extension, at most one. Its pinned first class has one occurrence per message, since the chain rank is zero, and a following chain pin compares one of its at most two chain readers with one of the two roots, at most four choices. In all, at most $48ab + 4 \le 49ab$ prefixes. We charge all of them, including those that are followed by no event, so nothing has to be said about how far constants propagate. The count uses the multiplicity of classes, the positions of the roots and the injectivity of the counters, and not the number of backward passes.

Two configurations that contribute nothing. One forward accident and one backward accident, with no mixed hit, give each side only the consequences of its one accident, so by Lemma 13

there is no event on either side, and an induced hit between the sides would need two equations on one side. And a rank-one accident on one side followed by a final comparison between the sides has rank three: each half of that comparison has one variable from each side, so it is implied only if both are pinned, which needs two equations on that line.

This completes the proof of Lemma 16: $40 + 40 + 8 + 14 + 49 = 151$, where the row for unequal backward lengths is absorbed in the second 40.

E.9. Assembly

By Lemmas 11, 12, 15 and 16 and the inclusion of Section E.7, with $\rho = \rho_S \geq \rho_N$ for every replay,

$$\begin{aligned}
&\Pr_\Psi[\neg\mathsf{Good} \wedge \neg\mathsf{Clash} \wedge \neg\mathsf{Dup}] \\
&\quad \leq \rho^2\Big(151 \sum_{R<T} L_R L_T + 20 \sum_T L_T^2\Big) \\
&\qquad + 2\rho^3\Big[\sum_T (6L_T)^6 + \sum_{R<T} \big(6(L_R + L_T)\big)^6\Big].
\end{aligned}$$

Since $\ell^2 = \sum_T L_T^2 + 2\sum_{R<T} L_R L_T$, the first bracket equals $75.5\,\ell^2 - 55.5\sum_T L_T^2 \leq 75.5\,\sigma^2$. In the second, $(L_R + L_T)^6 \leq 32(L_R^6 + L_T^6)$ and $\sum_{R<T}(L_R^6 + L_T^6) = (q-1)\sum_T L_T^6$, so the bracket is at most $6^6[1 + 32(q-1)]\sum_T L_T^6 \leq 6^6 \cdot 32\, q\, \ell L_{\max}^5$, and $2 \cdot 6^6 \cdot 32 = 2{,}985{,}984$. Lemmas 8 and 9 add $2\ell S/(M-S) \leq 12\sigma^2/(M-S)$, and Lemma 7 adds $q^2/2^{193} \leq \sigma^2/(18BM)$ since $q \leq \sigma/3$. With $\delta = S/M \leq 6 \cdot 2^{-68}$ for $\sigma \leq 2^{60}$,

$$\frac{75.5}{(1-\delta)^2} + \frac{12}{1-\delta} + \frac{1}{18B} < 88, \qquad \frac{2{,}985{,}984}{(1-\delta)^3} < 3.1 \cdot 10^6.$$

This proves Theorem 1.

Appendix F.
The Forward Pair Analysis in Full

This appendix proves the statements of Appendix A. Fix messages R before T with pass lengths $L = L_R$ and $L' = L_T$, both at least 2, and write $X = X^R$ and $X' = X^T$. The replay processes the steps $1, \ldots, 3L$ of R and then the steps $1, \ldots, 3L'$ of T, comparing the input of each step with the input of every existing class; for a single message it processes that message alone. Write R_t and T_t both for the step and for its input.

The joining steps. In this appendix we treat every step up to $3L$ as a forward step whose class has two independent output variables. The input of joining step $2L + j$ with $j \geq 2$ then contains the chain variable of the class of step $2L + j - 1$ in place of the tape value g_{j-1}. Before the first event the class of a joining step is a singleton, since a hit on it is an event, so that chain variable occurs in exactly one input, as the tape value does. The two replays therefore make the same comparisons with the same equations up to the first event, and every statement below holds for the process of Section 5.

Events. The single-message event E_R is that the input of a step of R in pass 3 equals the input of another step of R. The pair event $E_{R,T}$ is that $T_k = R_m$ with $k > 2L'$ or $m > 2L$; let $E^*_{R,T} = E_{R,T} \cap \neg E_R \cap \neg E_T$. These are the input events of the process of Section 5 restricted to forward steps and joining inputs.

F.1. Classes, variables and identities

Let p be the common prefix depth, the largest $t \le \min(L, L')$ with $X_A = X'_A$ and $X_u = X'_u$ for $u \le t$, or $p = 0$ if there is none. Since distinct messages have distinct inputs, $p < \max(L, L')$, and $p < L$ if $L = L'$.

A class is a set of steps with equal inputs, and a class c has output variables $\gamma(c)$ and $\delta(c)$. The halves of the input of step t of a message are

$$a(t) = \begin{cases} X_A, & t = 1, \\ \gamma(c(t-1)) \oplus \langle t-1 \rangle, & t \ge 2, \end{cases}$$

$$z(t) = \begin{cases} X_t, & t \le L, \\ \delta(c(t-L)), & t > L. \end{cases}$$

A chain equation involves only γ-variables and constants, and a register equation only δ-variables and constants. For two steps x, y in one class, the equality $a(x) = a(y)$ is the chain equation of the pair and $z(x) = z(y)$ its register equation. A half-equation is an *identity* if it holds syntactically: the same constant, or for chain halves the same class and the same counter, or for register halves the same class.

The comparisons and accidents are those of Section 5. Since the chain and register variables are disjoint, an accident adds [a-equation new]+[z-equation new] $\in \{1, 2\}$ to the rank, so $\mathrm{rank}(w)$ is at least the number of accidents in w. Write r_a and r_d for the chain and register ranks.

The same-class principle. If two steps end in one class, both half-equations of the pair are identities or follow from E: chain the hits that put them into one class.

Edge graphs. For every pair (u, v) in one class whose chain equation is not an identity, put the edge $\{c(u-1), c(v-1)\}$ into a graph G_a, with a vertex $*$ for a step of index 1; similarly put $\{c(u-L_u), c(v-L_v)\}$ into G_d, with $*$ for a constant register half. By the principle each edge is a two-variable equation in E, so the rank of the spanning forest of G_a is at most r_a and that of G_d at most r_d.

Lemma 25. *A class contains at most $r_a + 1$ steps of one message.*

Proof. Two steps of one message have different counters, so their chain equation is not an identity and gives an edge of G_a. The predecessor classes of j such steps are j distinct vertices, since equal classes with different counters would give an inconsistent equation, and they lie in one component of G_a. So $j - 1 \le r_a$. □

F.2. The free state

Write $c_u = \{R_u, T_u\}$ for $u \le p$.

Lemma 26. *With E empty:*

(a) *the induced hits are exactly $T_t \to R_t$ for $t \le p$;*
(b) *the open comparisons of rank one are exactly:*

- *a* prefix extension*: $p = \min(L, L')$, and $T_{L+1} \to R_{L+1}$ with the register equation $\delta(c_1) = X'_{L+1}$ if $L < L'$, or $T_{L'+1} \to R_{L'+1}$ with $\delta(c_1) = X_{L'+1}$ if $L' < L$;*
- equal constants*: two pass-1 steps, of R or of T and not a prefix pair, whose register halves are equal constants; the chain equation on their predecessor classes is new;*

- *a* stride hit*:* $T_{u+L'} \to R_{u+L}$ *with* $u \le p$*, whose register halves are both* $\delta(c_u)$*, with the new chain equation* $\gamma(c(T_{u+L'-1})) \oplus \gamma(c(R_{u+L-1})) = \langle u+L'-1\rangle \oplus \langle u+L-1\rangle$*;*

(c) *no comparison of rank at most one involves a step beyond pass 2.*

Proof. (a) A chain identity needs equal counters, hence equal indices, so it never holds inside one message. Across the pair it needs index t with $t-1$ merged, or $t=1$ and $X_A = X'_A$, and for $t \le \min(L, L')$ the register halves are a pair of constants. Induction on t gives exactly the prefix.

(b) Rank one needs one identity half. A chain identity needs index t with $t-1 \le p$; for $t \le \min(L, L')$ the register constants differ, which makes the comparison impossible, or agree, which makes it induced. So $t = \min(L, L')+1$ with $p = \min(L, L')$: a prefix extension. A register identity needs equal constants, which for two pass-1 steps is the second case, or the δ of one class read twice; the only classes with two steps are the c_u, $u \le p$, which gives the stride hit. The chain equation of a stride hit is never an identity: for $L = L'$ it would need $u + L - 1 \le p < L$.

(c) A prefix extension, an equal-constants comparison and a stride hit lie in passes 1–2 of both messages, since $u + L \le 2L$ and $u + L' \le 2L'$, and prefix hits lie in pass 1. □

F.3. Induced continuations

After accidents with equations E, an induced hit $x \to y$ needs its chain equation to be an identity or in the span of E, and the same for its register equation. Three mechanisms produce induced hits.

(M1) *Diagonal runs.* If (T_t, R_t) share a class, the chain equation of (T_{t+1}, R_{t+1}) is an identity, and the pair is induced if and only if its register equation is an identity or follows from E. A maximal run $[b+1, k]$ of merged same-index pairs either starts right after the prefix, $b = p$, or has a *bottom*: (T_b, R_b) is not merged, and the chain equation at $b+1$ is the edge $\{c(T_b), c(R_b)\}$, with value 0, which must lie in the span of E.

(M2) *Stride successors.* If $x \sim y$, then the steps that read their register values, $x+L_x$ and $y+L_y$, have identical register halves.

(M3) *Parallel copies.* A comparison whose predecessor classes, or whose register-supplier classes, are those of an edge already in E, with the same counter difference.

If $p < \min(L, L')$, index $p+1$ of a diagonal is never merged. If $X_A = X'_A$, the constants X_{p+1} and X'_{p+1} differ, so the register equation is impossible. If $X_A \ne X'_A$, then $p = 0$ and the chain constants at index 1 differ. So index $p+1$ can be merged only if $p = \min(L, L')$, and then its register equation is the prefix-extension equation or needs an edge of G_d.

F.4. Rank at most one

Lemma 27. *No accident sequence of rank at most one produces* E_R *or* $E_{R,T}$*.*

Proof. Rank 0 is Lemma 26(a). At rank one there is one accident r_1, a prefix extension, an equal-constants pair or a stride hit by Lemma 26(b). Its edge lies on free classes, singletons or prefix pairs, so (M3) gives only prefix copies of r_1 at the same indices. By (M2), the stride successor of r_1 has a chain equation that is not an identity, since its counter difference is $L - L'$ plus the index difference of r_1 or its predecessors are distinct singletons, and it is not the one equation in E, which lies on the predecessors of r_1; so it is not induced. By (M1) a run exists only if r_1 is a same-index pair.

- Equal constants at a common index $t \le \min(L, L')$: the run continues while $X'_u = X_u$ for $u \le \min(L, L')$ and, if $L = L'$, through $L+1, \dots, L+p$, whose register suppliers are prefix pairs. At $L+p+1$ the register equation compares $\delta(c(T_{p+1}))$ with $\delta(c(R_{p+1}))$, a new equation, and the run stops. If $L \ne L'$ it stops at $\min(L, L')$, where a constant meets a δ. Its last index is at most $L + p \le 2L - 1$.
- A stride hit with $L = L'$, at index $u+L$: the run is $[u+L, L+p]$, whose register suppliers are prefix pairs, and then it stops, at most $2L - 1$.
- A prefix extension at index $\min(L, L') + 1$: the next index needs a register equation on a new variable, and the run stops.

Pairs inside one message never have a chain identity, so no other merges occur. All merged pairs lie in passes 1–2, so there is no event. □

F.5. Rank two

Lemma 28. *(a) Every rank-two accident sequence that produces E_R is a single accident of rank two inside R, which is itself the event.*

(b) Every rank-two accident sequence that produces $E^*_{R,T}$ is in one of the following families:

- *a single accident $x \to c$ of rank two, with x a step of T and c a class containing a step of R, fewer than $36LL'$ choices;*
- *a prefix extension followed by its stride successor, which has a register identity and a new chain equation, with the event at R_{2L+1} if $L < L'$ or at $T_{2L'+1}$ if $L' < L$: at most 1;*
- *a stride hit $(T_{u+L'}, R_{u+L})$ followed by its stride successor $(T_{u+2L'}, R_{u+2L})$ in pass 3, $u \le p$: at most p;*
- *for $L = L'$, the* re-merge*: $T_{p+2} \to R_{p+2}$ (equal constants, bottom $p+1$), or if $p = L-1$ the stride hit $T_{L+1} \to R_{L+1}$ (bottom L), with $X' = X$ on $[p+2, L]$, followed by $T_{L+p+1} \to R_{L+p+1}$, which has a chain identity and the register equation $\delta(c(T_{p+1})) = \delta(c(R_{p+1}))$; the run then covers every index above and reaches $2L+1$: at most 1;*
- *for $L = L'$, a stride hit at index $u + L$ with its run $[u + L, L + p]$, followed by $T_k \to R_k$ with $k \in [u + 2L, 2L + p]$, whose register half is an identity because $(k - L, k - L)$ lies in the run, and whose chain equation is new at $k - 1$: at most $p(p + 1)/2$;*
- *for $L = L'$ and $p \ge 1$, two same-index accidents with a run between: equal constants at $t \in [p + 2, L]$ with $X' = X$ on $[t, L]$, giving the run $A = [t, L + p]$, then $T_j \to R_j$ with $j \in [\max(L+p+2, t+L), 2L+p]$, whose register half is an identity because $(j-L, j-L)$ lies in A and whose chain equation is new at the bottom $j - 1 > L + p$; the run B from j reaches $2L+1$, since its register suppliers lie in A up to $L+p$: at most $L(L-p-1) \le LL'$.*

The total is at most $36LL' + [L = L']\, LL' + 2 + p + p(p+1)/2 \le 40LL'$, using $p \le \min(L, L')$ and $L, L' \ge 2$.

Proof. We split by $(r_a, r_d) \in \{(0,2), (1,1), (2,0)\}$. Let (T_k, R_m) be the event pair, with $k > 2L'$ or $m > 2L$, and put $d = k - m$.

Case $(0, 2)$. Every same-class pair has a chain identity, so it has equal indices and merged predecessors, and the merged indices form an interval $[1, k']$. Having $k' > p$ needs $p = \mu := \min(L, L')$ and, for $u \in [\mu + 1, k']$, the register equations $\delta(c_{u-L}) = X'_u$, or $\delta(c_{u-L}) = \delta(c_{u-L'})$ (for $L < L'$; the other case is symmetric), each with the new variable $\delta(c_{u-\mu})$. An event needs $k' \ge 2\mu + 1$, so $r_d \ge k' - \mu \ge \mu + 1 \ge 3$, which is impossible.

Case $(1,1)$. There is one chain edge e_a and one register edge e_d, so at most one diagonal run has a bottom, and every half-equation that is not an identity is e_a or e_d.

- $d = 0$ and $L = L'$. The event pair lies in a run $[b+1, k']$ with $k' \geq 2L+1$. For u in the run with $u \in [L+1, 2L]$, the register equation is an identity if and only if $u - L \leq p$ or $u - L$ lies in the run. Each register equation that is not an identity is a distinct edge $\{c(T_{u-L}), c(R_{u-L})\}$, so there is at most one. If $b \leq L-1$ there are $b-p$ of them, which forces $b = p+1$: the re-merge. If $b = L$ there are $L-p$, which forces $p = L-1$: the re-merge. If $L+1 \leq b \leq 2L-1$ there are at least two. If $b \geq 2L$, a single pair carries both edges: a single accident. The case $b = p$ is impossible, since index $p+1 < L+1$ has different constants.
- $d = 0$ and $L \neq L'$. At indices in $(\min(L, L'), \max(L, L')]$ the register equations compare a constant with a δ, and above $\max(L, L')$ the register suppliers $(u - L', u - L)$ have different indices, so an identity would need a second chain edge. A run of length at least two therefore needs two register edges, and only a single pair remains: a single accident.
- $d \neq 0$. The chain equation of the event pair is the edge $e_a = \{c(T_{k-1}), c(R_{m-1})\}$. If its register equation is e_d, both edges lie on the event pair, and no free rank-one option carries either of them (their edges lie in passes 1–2 by Lemma 26), so the event pair is itself the accident: a single accident. If its register equation is an identity, then $(T_{k-L'}, R_{m-L})$ is merged with index difference $d' = d + L - L'$. $d' \neq 0$ would need a second chain edge, so $d' = 0$ and that pair lies on a diagonal. A bottom would need a chain edge equal to e_a, forcing $d = 0$; so the diagonal starts after the prefix, $p = \mu$, and by the register budget the run is the prefix-extension pair alone. This is a prefix extension followed by its stride successor.

Case $(2,0)$. Every same-class pair has a register identity. If $k > 2L'$ then $z(T_k)$ is a δ, so $m > L$ and $(T_{k-L'}, R_{m-L})$ is merged, and again $(T_{k-2L'}, R_{m-2L})$; the case $m > 2L$ is symmetric. Both accidents are register-identity options with chain edges, and r_1 is an equal-constants pair or a stride hit by Lemma 26, with its edge in passes 1–2. The event pair P_0 is not induced by the edge of r_1, since its predecessors lie at index at least $2L'$ or $2L$. So $P_0 = r_2$, or P_0 lies on a diagonal with bottom r_2. A diagonal of length at least two above $\max(L, L')$ needs at each index a merged register-supplier pair (M2) with index difference $L - L'$, which for $L \neq L'$ needs a new chain edge each time; so $P_0 = r_2$ when $L \neq L'$. Then $P_1 = (T_{k-L'}, R_{m-L})$ is merged before r_2. Either $P_1 = r_1$, so that r_1 is a stride hit and P_2 is a prefix pair: a stride hit followed by its stride successor. Or P_1 lies on the run of r_1, which needs r_1 to be a same-index pair; for $L \neq L'$ such a run never leaves pass 1, since at $\min(L, L') + 1$ a constant meets a δ, a contradiction.

For $L = L'$ we have $d = 0$, since with $d \neq 0$ the three levels need distinct chain edges. P_0 lies on a run B with bottom r_2, and P_1 on a run. The merged same-index pairs are the prefix, the run A with bottom r_1, and B. Index $p+1$ is never merged, so no run passes through index $L+p+1$, whose register suppliers are $(p+1, p+1)$. Hence B starts at some $j \geq L+p+2$, A ends at $L+p$, and every register supplier $(w-L, w-L)$ of B, $w \in [j,k]$, lies in A; so $j - L$ is at least the start of A and $k \leq 2L+p$. If r_1 is a stride hit, with $A = [u+L, L+p]$, this is a stride hit with its run followed by one more accident. If r_1 is an equal-constants pair at t, with $A = [t, L+p]$, which needs $X' = X$ on $[t, L]$ and $p \geq 1$ for $L+1$ to lie in A, this gives two same-index accidents with a run between.

Accidents inside one message. It remains to show that an accident inside one message produces no event $E^*_{R,T}$ at rank two beyond the single accidents already counted. Let $r_1 = (x \to c)$ with x and the earliest step of c in the same message; if c contains a step of R when r_1 occurs and x

is a step of T, r_1 is counted among the single accidents. Let K be the set of indices k such that T_k shares a class with a step of R.

1) K is an interval $[1, K_{\max}]$ containing $[1, p]$. An induced hit $T_k \to c'$ with $y \in c'$ a step of R has its chain equation an identity, which puts T_{k-1} with R_{k-1}, so that $k-1 \in K$, or equal to e_a. Suppose it is e_a. Then $\{c(T_{k-1}), c(y-1)\}$ is the pair of predecessor classes of r_1. If r_1 is inside R, both classes are classes of R or prefix classes, and $k-1 \in K$. If $r_1 = (T_j \to T_i)$ is inside T, then $c(y-1)$ must be the one predecessor class of r_1 that contains a step of R. If $i \le p$, the target of r_1 already contains R_i and r_1 is a single accident counted above. So $i = p+1$, $c(y-1) = c_p$, $y = R_{p+1}$ and $k = j$. Its register equation, together with e_d ($z(T_j) = z(T_{p+1})$), reduces to $z(T_{p+1}) = z(R_{p+1})$. If $p < \min(L, L')$, this is $X'_{p+1} = X_{p+1}$, which is impossible. If $L' < L$ and $p = L'$, it is $\delta(c_1) = X_{L'+1}$, which is not in the span of e_d, a $\delta = \delta$ equation. If $L < L'$, $p = L$ and $j = L'+1$, then e_d is the prefix-extension equation, and the distinct earlier classes of T_{L+1} and R_{L+1} would be forced equal, so the sequence is inconsistent and has probability 0. In all cases $k-1 \in K$ or $k = 1$.
2) $K_{\max} \le p+1$. At $k = p+1$ the chain equation is $\gamma(c_p) \oplus \langle p \rangle = a(y)$. It is an identity only for $y = R_{p+1}$, and it equals e_a only if c_p and $c(y-1)$ are the predecessor classes of r_1 with the right counter, which again means that T_{p+1} joins the class of R_{p+1}. Then the register equation at $p+1$ is impossible, or $p = \mu$ and it is the prefix-extension equation; for instance $L = 3$, $L' = 2$, $X = (0; 10, 11, 12)$, $X' = (0; 10, 11)$, $r_1 = R_4 \to R_3$ and $e_d : \delta(c_1) = 12$ make $T_3 \to R_3$ induced. Index $p+2$ is then needed next. Its register equation compares $\delta(c_2)$ with a constant or with a δ of another class, an equation on a variable other than $\delta(c_1)$, because $c_1 \ne c_2$ and an accident merging them would be a pass-1 pair of rank one. But r_1 supplies only one register equation, used at $p+1$. Hence $K_{\max} \le p+1$.
3) So every pair (T_k, R_m) across the messages has $k \le p+1 \le 2L'$, and $E^*_{R,T}$ needs $m > 2L$. If $k \le p$, then R_m lies in the prefix class c_k, so $R_m \sim R_k$ is a pair inside R with an index in pass 3: this is E_R. If $k = p+1$, then $R_m \sim R_{p+1}$, again E_R. So $E^*_{R,T}$ fails.

An accident inside T is symmetric. Part (a) is the version for one message: there are no prefix pairs, (M1) never applies, two equal-constants accidents cannot reach pass 3 because the register chain would need three chain edges, and a rank-two accident induces nothing but itself.

The bound on the total uses $p \le \min(L, L')$, $L, L' \ge 2$: $2 + p + p(p+1)/2 \le 3\min(L, L')^2 \le 3LL'$. □

Appendix G.
Complete Proof of Theorem 4

Section 6 gives the proof of Theorem 4 as a list of changes to Section 5, and Appendix C proves Lemma 17. This appendix writes out the remaining components. A message T of pass length L has $5L$ steps, $\ell = \sum_T L_T = \sigma - q$ and $S = 5\ell$.

G.1. The process

The process $\Gamma_5(\tau)$ takes the messages one at a time. For a message of pass length L it runs the forward steps $1, \dots, 2L$ and the backward steps $5L, \dots, 3L+1$ as in Section 5, and then the

joining steps $2L+j$, $1 \le j \le L$. With B_t the point of backward step t, the joining points are

$$\begin{aligned} I_{2L+1} &= (a_{2L},\, d_{L+1}), \\ I_{2L+j} &= (g_{j-1} \oplus \langle 2L+j-1 \rangle,\, d_{L+j}) \quad (j \ge 2), \\ O_{2L+j} &= (g_j,\, \beta(B_{3L+j})) \qquad\qquad\quad (j < L), \\ O_{3L} &= (\alpha(B_{3L+1}) \oplus \langle 3L \rangle,\, \beta(B_{4L})), \end{aligned}$$

where each g_j is drawn uniformly from all 2^{64} values. The process stops at the first bad event: an *input clash*, a joining input already defined; or an *output clash*, a free or determined joining output already used. A rejected joining candidate is not inserted. An old backward step is not a bad event. Good_5 is the event that no bad event occurs.

G.2. Ratio and coupling

Lemma 29. *For every attainable τ, $\Pr_{\mathrm{real}}[\tau] / \Pr_{\mathrm{ideal}}[\tau] \ge 1 - q^2/2^{193} - \Pr_{\Gamma_5}[\neg \mathsf{Good}_5]$.*

Proof. Fix a good draw history. Its points form a partial injection that contains the whole five-pass computation of every X^T, ending in C^T. Let D be the number of points created before a given creation. A fresh forward or backward draw contributes $1/(M-D)$ to the probability of the history under Γ_5 and to the probability that a uniform permutation contains its point set, so these factors cancel. A free joining step contributes $1/B$ and $1/(M-D)$, and the determined step 1 and $1/(M-D)$. So the ratio of the second probability to the first is at least

$$\prod_T \Big(\frac{B}{M}\Big)^{L_T - 1} \frac{1}{M} = \prod_T B^{-(L_T+1)}.$$

Distinct good histories assign, at their first differing draw, different values to the same argument, so their completion events are disjoint. Summing gives $\Pr_{\mathrm{real}}[\tau] \ge \prod_T B^{-(L_T+1)} \Pr_{\Gamma_5}[\mathsf{Good}_5]$. The ideal correction is that of Lemma 7: $\Pr_{\mathrm{ideal}}[\tau] \le \prod_T B^{-(L_T+1)} / (1 - q^2/2^{193})$. Multiply and use $(1-x)(1-y) \ge 1 - x - y$. □

The argument uses neither the number of passes nor the freshness of any backward step, and the ratio has no term for occupied outputs.

Lemma 30. *Let Ψ_5 run Γ_5 with a fresh forward step at x given output $\pi(x)$, a fresh backward step at y given input $\pi^{-1}(y)$, for a uniform permutation π, and the free joining steps given values from an independent uniform tape, with the same stopping rule. Then*

$$\Pr_{\Gamma_5}[\neg \mathsf{Good}_5] \le \Pr_{\Psi_5}[\neg \mathsf{Good}_5 \wedge \neg \mathsf{Dup}] + \frac{4\ell^2}{M - 5\ell}.$$

Proof. Couple corresponding draws maximally up to the first bad event. A fresh forward draw of Ψ_5 avoids the outputs of π revealed so far, while that of Γ_5 also avoids the joining outputs, at most ℓ further values; both are uniform on at least $M - S$ values, so their distance is at most $\ell/(M-S)$. Backward draws are the same with inputs. There are at most 4ℓ forward and backward draws, two passes on each side. Joining draws, including the check for an output clash, have the same distribution in both processes. When the coupling succeeds, the two runs have the same points and the same first bad event, and there is no Dup. □

G.3. Paths and the reduction

Paths are defined as in Section 5, from π, the constants of the message and its own tape. For R before T, let $\mathsf{Bad}^{(5)}_{R,T}$ be the union of:

- a joining input of T equals the input of a point of R;
- the input of a forward step of T equals a joining input of R;
- a joining output of T, free or determined, equals the output of a point of R;
- the output of a backward step of T equals a joining output of R.

$\mathsf{Bad}^{(5)}_T$ is the event that a joining input or joining output of T equals the input, respectively output, of a forward or backward point of T or of an earlier joining point of T.

Lemma 31. *On* $\neg\mathsf{Dup}$, $\{\neg\mathsf{Good}_5\} \subseteq \bigcup_T \mathsf{Bad}^{(5)}_T \cup \bigcup_{R<T} \mathsf{Bad}^{(5)}_{R,T}$.

Proof. Let x be the first step, in processing order, at which the run of Ψ_5 differs from the path or a bad event occurs. Every earlier step follows its path, so x's computed side does too. An old forward or backward step that hits a point created by a fresh forward or backward step reads π or π^{-1} there, which is what its path reads; this includes an old step in pass 4. A hit on a joining point of an earlier message is the second or the fourth event. A fresh forward or backward step reads π or π^{-1}. At a joining step, a defined input is the first event or a single-message event, and a used proposed output is the third event or a single-message event; otherwise the step uses its own tape value and follows its path. On $\neg\mathsf{Dup}$ every lookup is unique. No number of passes and no freshness assumption enters. □

G.4. Costs

Lemma 12 holds as stated: comparing a free joining output is one more comparison of two halves, each a variable xor a constant or a constant, and it respects the separation into chain and register equations. For the second claim we recount the comparisons. A replay of N steps makes at most N^2 comparisons of the ordinary kinds, and each of its at most $N/5$ joining steps also compares its output with at most N classes, so an accident has at most $A_N = \frac{6}{5}N^2$ choices. For $N \geq 10$,

$$A_N + A_N^2 + A_N^3 \leq \Big(1.728 + \frac{1.44}{N^2} + \frac{1.2}{N^4}\Big) N^6 < 2N^6,$$

so $\Pr[\text{rank} \geq 3] \leq 2N^6\rho_N^3$.

G.5. One message

For a message P of pass length L, let $H^{(5)}_P$ be the union of $\mathsf{Bad}^{(5)}_P$, an accident of rank increment two in the replay of P alone, stopped at its first event, and a prefix $P_j \to P_1$, $P_{5L} \to c(P_{j-1})$ with $2 \leq j \leq L$.

Lemma 32. $\Pr[H^{(5)}_P] \leq 17L^2\rho^2 + 2(5L)^6\rho^3$.

Proof. Give each of the $5L$ steps its ordinary comparison: the input for forward and joining steps, the output for backward steps. These give at most $\binom{5L}{2}$ single accidents. Each joining step $2L+j$ also compares its output with at most $4L+j-1$ existing points, which adds $\sum_{j=1}^{L}(4L+j-1) = 4L^2 + L(L-1)/2$. Together this is $\binom{5L}{2} + 4L^2 + \frac{L(L-1)}{2} = 17L^2 - 3L$, and the mixed prefixes add

at most $L-1$, for at most $17L^2-2L-1 \le 17L^2$ prefixes of rank two. Rank three or more costs $2(5L)^6\rho^3$. It remains to show that two rank-one accidents outside the mixed prefixes produce no event of P. Before the first hit between forward and backward points the two sides have disjoint variables.

- Two accidents between forward steps: the single-message part of the forward analysis, Lemma 28(a), excludes a first joining-input event; two rank-one accidents cannot carry the needed register sharing through two passes. No register half of a forward class is pinned to a constant, so a first hit induced from the backward side is impossible too.
- Two accidents between backward steps: reverse the computation of P. The same argument excludes a joining-output event. The reversal needs only that the counters are injective within one message.
- One accident on each side, with no explicit hit between the sides: each accident is a chain equation, the register halves stay unpinned, so no hit between the sides is induced, and each side separately carries only the consequences of one accident, which do not reach its joining interface.
- The second accident is the first explicit hit between the sides. In the replay of P alone all forward steps come first, so it is a backward step hitting a forward class. If the first accident was between backward steps, neither free half of a forward class is constrained and the hit costs two more equations. If it was between forward steps, the register line is unconstrained, so for the hit to cost one equation its chain equality must already hold: the backward step is the root P_{5L} and the chain half of the target must have been pinned by $P_j \to P_1$. That is the mixed prefix.

A joining comparison at rank one is excluded by the argument for one accident. A tape value takes part in no earlier equation that is not final. A final comparison between the two sides after one accident inside one side still has two half-equations that are not implied, because the side without an accident supplies an independent variable on each line. □

None of the three consequences of Lemma 14 is used; they rested on fresh pass-4 steps.

G.6. Pairs

Fix R before T, $a = L_R$, $b = L_T$, and exclude $H_R^{(5)}$ and $H_T^{(5)}$. The rows of the right column of Table 3 are as follows.

Forward against forward: $40ab$**.** The forward analysis of Appendix F uses only the two forward passes and the comparisons at joining inputs. Beyond pass 2, replacing the fresh chain halves by independent tape values preserves the expressions used in these comparisons. So Lemma 28(b) gives at most $40ab$.

Backward against backward, equal lengths: $40ab$**.** For $a = b = L$ put $r = 5L + 1 - t$. Read in r, the backward side runs π^{-1} from the chaining value $C_A \oplus \langle 5L \rangle$, with step counter $\langle 5L - r \rangle$, which is injective and the same for both messages at equal r; the register stride is still L. The reversed tapes

$$\bar{g}_k = g_{L-k} \oplus \langle 3L - k \rangle, \qquad 1 \le k < L,$$

are a bijective image of independent uniform tapes, hence again independent and uniform. Joining-output comparisons become joining-input comparisons, and a backward output that meets a joining output of R becomes a forward input that meets a joining input of R. Final comparisons do not

change the paths before them, so their order may be reversed when bounding the union of first final events. The backward side is therefore a forward side, and the count is $40ab$.

Backward against backward, unequal lengths: $9ab + 1 \le 40ab$**.** The two reversed counter schedules differ, and the reversal does not apply. Lemma 17, proved in Appendix C, gives $9ab+1$.

Example. The full-pin exclusion in Lemma 17 is needed. Take $a = 2$, $b = 3$, let the last output of R be (h, c) and its other pass-5 register constant also be c, and write the input of R_{10} as (u, v). Then the output of R_9 is $(u \oplus 9, c)$, and the rank-one accident $R_9 \to R_{10}$ imposes $u = h \oplus 9$. Consequently the output of R_8 is $(u \oplus 8, v) = (h \oplus 1, v)$. If the last output of T is $(h \oplus 1, e)$, the comparison $T_{15} \to R_8$ needs only $v = e$. So an alias at the top followed by a register equation lets the last step of the longer message hit pass 4 of the shorter one. The two equations pin both u and v of $c(R_{10})$, and the prefix is charged among the full pins below.

Direct comparisons between the sides: $4ab$**.** A joining input of T against the input of one of the $2a$ backward points of R, $b(2a)$, and a joining output of T, free or determined, against the output of one of the $2a$ forward points of R, $b(2a)$: together $4ab$.

Prefixes with an explicit accident between the sides: $10ab$**.** A single rank-two accident between a forward step of T and a backward point of R, or a backward step of T and a forward point of R, has at most $(2b)(2a) + (2b)(2a) = 8ab$ choices. For two rank-one accidents with an explicit one between the sides, the first must pin one free half of the target. A chain pin lets a root of the other side supply the implied half, and a register pin is a forward prefix extension and lets a constant of the outer pass of the other side supply it. The four families of Appendix B remain, with $5a$ and $5b$ in place of $6a$ and $6b$, and their counts $a - 1$, $a - 1$, $a + b - 2$ and b add up to $3a + 2b - 4 \le 2ab$. These families use roots and steps of the outer passes, and neither the removed backward pass nor fresh pass-4 steps. Together, $10ab$.

Prefixes that pin a target fully: $49ab$**.** The count of Section E.8 applies verbatim: before the first hit between the sides their equations have disjoint supports, so an induced hit needs both free halves of its target class to reduce to constants; a single rank-two accident on one side cannot supply both, since its chain and register suppliers are $L - 1 \ge 1$ positions apart; with chain rank at most one a class has at most two occurrences per message (Lemma 25 and its mirror) and four over the pair; a first chain pin has one of at most $4(a + b - 2)$ root positions and the second pin at most $4 \cdot 6 = 24$ choices, which gives $96(a + b - 2) \le 48ab$; and a first register pin is the one prefix extension, followed by at most four chain pins. The total is $48ab + 4 \le 49ab$. The count uses multiplicities, root positions and counter injectivity, and not the number of backward passes.

The paragraph of Appendix B on joining outputs as targets is not used for five passes, because its freshness hypothesis is gone; the equal-length case is covered by the reversal and the unequal-length case by Lemma 17. The total is $40ab + 40ab + 4ab + 10ab + 49ab = 143ab$.

G.7. Assembly

With $\rho = \rho_S$, Lemmas 29, 30, 31, Lemma 12 with the count above, Lemma 32 and the cover give

$$\begin{aligned}\mathrm{Adv}^{\pm\mathrm{prp}}_{W_5} \le \frac{q^2}{2^{193}} + 71.5\,\ell^2\rho^2 + \frac{4\ell^2}{M-S} \\ + 2\cdot 5^6\Big[\sum_T L_T^6 + \sum_{R<T}(L_R + L_T)^6\Big]\rho^3,\end{aligned}$$

because $143\sum_{R<T} L_R L_T + 17\sum_T L_T^2 = 71.5\,\ell^2 - 54.5\sum_T L_T^2 \le 71.5\,\ell^2$. The bracket is at most $32q\sigma L_{\max}^5$, as in Section E.9, so its coefficient is $64 \cdot 5^6 = 1{,}000{,}000$. For $\sigma \le 2^{60}$ put $\delta = S/M \le 5 \cdot 2^{-68}$. Since $q \le \sigma/3$,

$$\frac{71.5}{(1-\delta)^2} + \frac{4}{1-\delta} + \frac{1}{18B} < 76, \qquad \frac{10^6}{(1-\delta)^3} < 1.01 \cdot 10^6.$$

This proves Theorem 4.

Appendix H.
Complete Proof of Theorem 3, Part 1: Wraps, Two-Semiblock Unwraps and Freshness

This appendix and the next give the complete derivations behind the proof of Theorem 3 in Section 7.2 and Appendix D. This one treats the wraps and the two-semiblock unwraps: the setting (Section H.1), Lemma 18 (Section H.2), the first two rows of Table 4 (Section H.3), the row of the two-semiblock unwraps (Section H.4), and Lemmas 19 and 20 with every family of Table 6 (Sections H.5–H.13). Throughout, $B = 2^{64}$, $M = B^2$ and $\sigma \le 2^{60}$.

H.1. The scheme and the game

Format values and counters. For $0 \le j < 2^{32}$ we write $\Lambda_j = \texttt{A65959A6} \,\|\, [j]_{32}$, as in Section 2, and call Λ_j a *short* format value if $1 \le j \le 8$ and a *long* one if $j \ge 9$. All step indices below are smaller than 2^{32}, so the counter $\langle t \rangle$ has a zero left half, and

$$\Lambda_j \oplus \langle t \rangle = \Lambda_{j \oplus t},$$

where $j \oplus t$ is the xor of the 32-bit encodings.

Wraps. A wrap of $m \le 8$ bytes is a *one-block wrap*. It returns $\pi(\Lambda_m \,\|\, P')$, and its input together with its answer is an *exposed point*. We number the exposed points $i = 1, \ldots, p$, with inputs (Λ_{m_i}, P_i), $m_i \le 8$, and answers (γ_i, δ_i). A wrap of $m \ge 9$ bytes is a *long wrap*, which we also call a message. Its pass length is $L = \lceil m/8 \rceil \ge 2$, and it returns $W(X)$ for $X = (\Lambda_m; X_1, \ldots, X_L)$. Two properties of the format are used throughout.

(F1) The first input (Λ_m, X_1) of a long wrap is never an exposed input, since its left half is long.

(F2) The first semiblock Λ_m of a long wrap determines its pass length, so two long wraps with equal first semiblocks have equal pass lengths.

Since $m < 2^{32}$, every long wrap has $L \le 2^{29}$, and its step indices are at most $6L < 2^{32}$.

Unwraps. An unwrap of a ciphertext C of two semiblocks computes $\pi^{-1}(C)$ and accepts if and only if the result lies in

$$\mathsf{Valid}_2 = \{(\Lambda_m, P') : 1 \le m \le 8,\ P' \text{ zero after byte } m\}.$$

Hence

$$v := |\mathsf{Valid}_2| = \sum_{m=1}^{8} 2^{8m} = \frac{256(B-1)}{255} < \frac{256B}{255}.$$

An unwrap of $L+1 \ge 3$ semiblocks computes W^{-1} and accepts if and only if the result lies in the set Valid_{L+1} of inputs whose first semiblock is Λ_m with $8L - 7 \le m \le 8L$ and whose

padding bytes are zero. The first $L-1$ plaintext semiblocks are arbitrary and the last one has $m-8(L-1) \in \{1,\ldots,8\}$ free bytes, so $|\mathsf{Valid}_{L+1}| = B^{L-1}v$.

The game and the resources. In the ideal world every wrap query is answered by an independent uniform string of its length, and every unwrap is rejected. The adversary is deterministic, repeats no query and never unwraps the answer of a previous wrap; repeated queries can be answered from a table, and unwraps of a malformed length are rejected in both worlds. Let p be the number of one-block wraps, q the number of long wraps, L_T the pass length of the long wrap T, $\ell = \sum_T L_T$, $S = 6\ell$ and $D = S + p$. Let r be the number of two-semiblock unwraps and $q_L = q_v - r$ the number of long unwraps, with pass lengths L_k, and $s = \sum_k L_k$. All of these are deterministic caps, and

$$\ell + q + 2p + \sum_k (L_k + 1) + 2r \le \sigma.$$

In particular $\ell + 2p \le \sigma$, $\ell + p \le \sigma$ and $D \le 6\sigma < B$.

The process and the good event. A transcript τ lists the exposed points, the long wraps with their answers, and the unwraps with the answer "reject"; a transcript with an accepted unwrap has ideal probability 0 and is not considered. The process $\Gamma(\tau)$ first inserts the p exposed points, and then takes the long wraps in query order. For a long wrap T with answer $C = (C_A; C_1, \ldots, C_L)$, it runs the steps $t = 1, \ldots, 6L$ forward. The input of step t is known when the step is reached. If the current point set contains a point with that input, the step is *old* and reads the output of that point. Otherwise the step is *fresh* and creates a point: at a pass-6 position $t = 5L + j$ with $j < L$ it draws γ uniformly from $\{0,1\}^{64}$ and sets $\delta = C_j$; at $t = 6L$ the output $(C_A \oplus \langle 6L \rangle, C_L)$ is determined; at every other step it draws the output uniformly from $\{0,1\}^{128}$. All draws are independent, and a value may be drawn twice. Let P be the final point set, made of the exposed points and the created points. The good event $\mathcal{G}$ requires three conditions: no two points of P have the same output; no step in passes 3–6 of a long wrap is old; and no 18 points of P share their γ, and no 18 share their δ (the *cap*). For a good P and an unwrap k, $V_k(P)$ is the probability, over a uniform permutation extending P, that unwrap k is accepted.

H.2. Proof of Lemma 18

Let τ be a transcript with no accepted unwrap. The adversary is deterministic, so $\Pr_{\text{ideal}}[\tau]$ is the probability that the ideal oracles give the answers of τ, and the same holds in the real world. The ideal answers are independent uniform strings and the ideal unwraps reject with probability 1, so

$$\Pr_{\text{ideal}}[\tau] = M^{-p} \prod_T B^{-(L_T+1)}.$$

In the real world, $\Pr_{\text{real}}[\tau]$ is the probability, over a uniform permutation π, that every wrap of τ has the answer listed in τ and every unwrap of τ is rejected.

A *history* of Γ is the sequence of its draws; it determines the point set P. Call a history good if $\mathcal{G}$ holds. For a good history, P is a partial injection: its inputs are distinct because a fresh step has an input not in the current set, and its outputs are distinct by $\mathcal{G}$. It contains the whole computation of every long wrap of τ, ending in its answer, and every exposed point. Hence $\{\pi \supseteq P\}$ implies that π agrees with all wraps of τ. Two distinct good histories first differ at a draw made at the same input, with different outputs, so the events $\{\pi \supseteq P\}$ of distinct good histories are disjoint.

Given $\pi \supseteq P$, the probability that all unwraps are rejected is at least $(1 - \sum_k V_k(P))_+$ by the union bound. Therefore

$$\Pr_{\text{real}}[\tau] \geq \sum_{\text{good}} \Pr[\pi \supseteq P] \Big(1 - \sum_k V_k(P)\Big)_+,$$

the sum running over good histories. Consider a good history with n created points. On $\mathcal{G}$ every pass-6 position is fresh, so the history consists of $n - \ell$ full draws, $\ell - q$ draws of a γ alone, and q determined outputs, and $\Pr_\Gamma[\text{history}] = M^{-(n-\ell)} B^{-(\ell-q)}$. On the other side, $\Pr[\pi \supseteq P] = 1/(M)_{p+n} \geq M^{-p-n}$; in particular, inserting the exposed points first costs nothing, since their factor $1/(M)_p$ is at least the ideal factor M^{-p}. Thus

$$\begin{aligned} \frac{\Pr[\pi \supseteq P]}{\Pr_\Gamma[\text{history}]} &\geq M^{-p} M^{-\ell} B^{\ell-q} \\ &= M^{-p} \prod_T B^{-(L_T+1)} = \Pr_{\text{ideal}}[\tau]. \end{aligned}$$

Summing over good histories,

$$\begin{aligned} \frac{\Pr_{\text{real}}[\tau]}{\Pr_{\text{ideal}}[\tau]} &\geq \mathbb{E}_\Gamma \Big[1_{\mathcal{G}} \Big(1 - \sum_k V_k(P)\Big)_+ \Big] \\ &\geq 1 - \Pr_\Gamma[\neg\mathcal{G}] - \mathbb{E}_\Gamma \Big[1_{\mathcal{G}} \sum_k V_k(P) \Big]. \end{aligned}$$

This is Lemma 18. □

By the H-coefficient technique in its expectation form [20], [25], the advantage is at most the ideal expectation of $\Pr_\Gamma[\neg\mathcal{G}] + \mathbb{E}_\Gamma[1_{\mathcal{G}} \sum_k V_k(P)]$. Call τ *colliding* if two exposed answers are equal, or two long wraps of the same length have equal answers. On a colliding transcript $\mathcal{G}$ fails with probability 1: either two exposed points share an output, or the last steps of the two long wraps have the same determined output, and then they collide or one of them is an old step in pass 6. On a non-colliding transcript, $\neg\mathcal{G}$ is the union of three events: Coll, two points of P with the same output; Cap, 18 points of P sharing a γ or a δ; and Old, an old step in passes 3–6. Hence

$$\begin{aligned} \mathbb{E}_{\text{ideal}} \Pr_\Gamma[\neg\mathcal{G}] \leq\ & \Pr_{\text{ideal}}[\text{colliding}] \\ &+ \mathbb{E}_{\text{ideal}} \big[1_{\neg\text{coll}} \Pr_\Gamma[\mathsf{Coll}] \big] \\ &+ \mathbb{E}_{\text{ideal}} \big[1_{\neg\text{coll}} \Pr_\Gamma[\mathsf{Cap}] \big] \\ &+ \mathbb{E}_{\text{ideal}} \big[1_{\neg\text{coll}} \Pr_\Gamma[\mathsf{Old}] \big], \end{aligned}$$

where $1_{\neg\text{coll}}$ is the indicator of a non-colliding transcript. The first two terms are the collision row of Table 4 (Lemma 33), the third is the cap (Lemma 34), and the last is the freshness bound of Lemma 20, proved in Section H.13.

H.3. Collisions and the cap

We use *labels* for the points that Γ may create: one label for each step of each long wrap, S in all, and one for each exposed point, p in all. We realize the draws of Γ by an independent tape that holds one uniform value per long-wrap step and is read only if the step is fresh. Every label then has a well-defined output value: for a step outside the pass-6 positions, its tape value; for a pass-6 position other than the last, the tape value as γ and the answer semiblock C_j as δ; for the

last step, $(C_A \oplus \langle 6L\rangle, C_L)$; for an exposed point, its answer. A created point has the output value of its label. Distinct labels never use the same answer semiblock in the same coordinate. Each tape value is uniform and independent of everything else, and each answer semiblock is uniform given all previous answers and independent of the other semiblocks of its answer.

Lemma 33 (Collisions)**.**

$$\Pr_{\text{ideal}}[\textit{colliding}] + \mathbb{E}_{\text{ideal}}\big[1_{\neg\text{coll}}\Pr_\Gamma[\mathsf{Coll}]\big] \le \frac{19\,\sigma^2}{M}.$$

Proof. Two exposed answers are equal with probability $1/M$, and two answers of long wraps of the same length L with probability $B^{-(L+1)} \le B^{-3}$. Hence

$$\Pr_{\text{ideal}}[\text{colliding}] \le \binom{p}{2}\frac{1}{M} + \binom{q}{2}\frac{1}{B^3} \le \frac{p^2}{2M} + \frac{q^2}{2B^3}.$$

For Coll we charge each collision to the later of the two labels, which is a long-wrap step x, and to the other label u, and show that the output values of x and u are equal with probability at most $1/M$, averaged over the ideal answers. If x is outside the pass-6 positions, its output value is a full tape value, equal to that of u with probability $1/M$. If x is a pass-6 position other than the last, its γ is a tape value, equal to the γ of u with probability $1/B$, and its δ is the answer semiblock C_j, which differs from the source of the δ of u and is therefore equal to it with probability $1/B$, independently. If x is the last step, its output value is made of two answer semiblocks, and the output value of u is a full tape value, or a tape value and another answer semiblock, or two other answer semiblocks; in each case equality has probability $1/M$. There are $\binom{S}{2}$ pairs of long-wrap labels and pS pairs of a long-wrap label with an exposed label, so

$$\mathbb{E}_{\text{ideal}}\big[1_{\neg\text{coll}}\Pr_\Gamma[\mathsf{Coll}]\big] \le \frac{\binom{S}{2} + pS}{M} \le \frac{18\ell^2 + 6p\ell}{M}.$$

Adding up, and using $\ell + 2p \le \sigma$ and $q \le \sigma$,

$$\begin{aligned} 18\ell^2 + 6p\ell + \frac{p^2}{2} &= 18\ell^2 + 3\ell(2p) + \frac{(2p)^2}{8} \\ &\le 18(\ell + 2p)^2 \le 18\sigma^2, \end{aligned}$$

and $q^2/(2B^3) \le \sigma^2/M$. The total is at most $19\sigma^2/M$. □

Lemma 34 (The cap)**.**

$$\mathbb{E}_{\text{ideal}}\big[1_{\neg\text{coll}}\Pr_\Gamma[\mathsf{Cap}]\big] \le \frac{2D(D/B)^{17}}{18!} \le \frac{0.032\,\sigma^2}{M}.$$

Consequently, on $\mathcal{G}$, for every $u \in \{0,1\}^{64}$ at most $\mu = 17$ points of P have $\gamma = u$, and at most μ have $\delta = u$.

Proof. If 18 created points share their γ, then their 18 labels have equal γ-values. For a given set of 18 labels, order them so that every tape value comes after all answers; each γ-value is then uniform given the previous ones, so all 18 are equal with probability at most B^{-17}. The same holds for δ. There are at most $D^{18}/18!$ sets of 18 labels, so the probability is at most $2D^{18}/(18!\,B^{17}) = 2D(D/B)^{17}/18!$. With $D \le 6\sigma$ and $\sigma \le 2^{60}$, we have $D^{18} \le 6^{18}\sigma^{16}\sigma^2 \le 6^{18}\,2^{960}\sigma^2$ and $B^{17} = 2^{960}M$, so

$$\frac{2D^{18}}{18!\,B^{17}} \le \frac{2\cdot 6^{18}}{18!}\cdot\frac{\sigma^2}{M} < \frac{0.0318\,\sigma^2}{M},$$

since $6^{18} = 101{,}559{,}956{,}668{,}416$ and $18! = 6{,}402{,}373{,}705{,}728{,}000$. □

H.4. Two-semiblock unwraps

Lemma 35 (Two-semiblock unwraps)**.** *Summed over the two-semiblock unwraps k,*

$$\sum_k \mathbb{E}_{\text{ideal}}\mathbb{E}_\Gamma\big[1_{\mathcal{G}}V_k(P)\big] \le \frac{(1+2^{-7})\,r}{B} + \frac{7r\sigma}{M}.$$

Proof. Consider one rejected ciphertext C of two semiblocks. It is a function of the answers given before it was asked, and it is not the answer of a previous one-block wrap, since the adversary never unwraps a wrap answer. For a good P, if C is the output of a point of P, then $V_k(P)$ is 1 or 0 according to whether that point has a valid input; otherwise $\pi^{-1}(C)$ is uniform on the at least $M-D$ unused inputs, and $V_k(P) \le v/(M-D)$. The points of P whose output can equal C are of four kinds.

Determined outputs. The last step of a long wrap R has the output $Y_R = (C_A^R \oplus \langle 6L_R\rangle, C_{L_R}^R)$, which is determined by τ. If $C = Y_R$ for some R, this point is unique on a good history, and its input halves are $\gamma \oplus \langle 6L_R - 1\rangle$ from the fresh step $6L_R - 1$ in pass 6 and δ from the fresh step $5L_R$ in pass 5. These are two distinct fresh draws, so the input is valid with probability at most v/M. Whether C equals some Y_R is determined by τ, so a ciphertext pays either this term or $v/(M-D)$, and never both.

Full draws. A fresh step outside pass 6 has a uniform output, equal to C with probability $1/M$. There are at most 5ℓ such steps.

Pass-6 positions. A fresh pass-6 position $t = 5L + j$, $j < L$, has input halves $\gamma \oplus \langle t-1\rangle$ from step $t-1$ and δ from step $t-L$ in pass 5, both fresh on $\mathcal{G}$, and its own drawn γ. For its point to have output C with a valid input, the first two of these three independent uniform halves must form a valid input and the third must equal the left half of C, which has probability at most $(v/M)(1/B) = v/B^3$. There are $\ell - q \le \ell$ such positions.

Future exposed points. A one-block wrap asked after C has a uniform answer, equal to C with probability $1/M$, and its input is valid. Over all pairs this costs at most pr/M.

Per ciphertext this gives at most $v/(M-D) + (5\ell + (v/B)\ell)/M$, and in all

$$\sum_k \mathbb{E}\big[1_{\mathcal{G}}V_k(P)\big] \le r\Big(\frac{v}{M-D} + \frac{5\ell + (v/B)\ell + p}{M}\Big).$$

Since $D \le B$ and $v < (256/255)B$,

$$\frac{v}{M-D} \le \frac{256}{255}\cdot\frac{1}{B(1-1/B)} \le \frac{1+2^{-7}}{B},$$

and $5\ell + (256/255)\ell + p \le 7(\ell + p) \le 7\sigma$. □

H.5. Freshness: eager sampling and pure paths

From here on, τ is a non-colliding transcript, so the exposed map $(\Lambda_{m_i}, P_i) \mapsto (\gamma_i, \delta_i)$ is injective. Let F be the random function on $\{0,1\}^{128}$ that takes the exposed answers on the exposed inputs and independent uniform values everywhere else. The *eager form* of Γ runs Γ with $F(x)$ in place of every draw at a fresh step with input x: a full draw becomes $F(x)$, and a pass-6 position other than the last takes γ from $F(x)$ and keeps $\delta = C_j$.

The *pure path* of a long wrap U is its computation with F at every step, without lookups: its pure inputs are $U_1 = (X_A^U, X_1^U)$ and, for $t \ge 2$,

$$U_t = \big(F_\gamma(U_{t-1}) \oplus \langle t-1\rangle,\ \beta_t\big),$$

where $\beta_t = X_t^U$ for $t \le L_U$ and $\beta_t = F_\delta(U_{t-L_U})$ for $t > L_U$. It depends only on F and X^U. We say that U_t is *above pass 2* if $t > 2L_U$, and *above pass 1* if $t > L_U$.

Lemma 36 (Pure paths)**.** *The eager form of* Γ *has the distribution of* Γ*. If* Old *occurs, then one of the following events holds on the pure paths:*

(i) *a pure input above pass 2 is an exposed input;*
(ii) *for some message* T*, a pure input of* T *above pass 2 equals another pure input of* T *(the event* E_T*);*
(iii) *for some* $R \neq T$*, a pure input of* R *equals a pure input of* T*, and one of the two is above pass 2 (the event* $E_{R,T}$*).*

Proof. A fresh step has an input outside the current point set, which contains the exposed inputs and every input at which F has been used, so F is evaluated there for the first time and its value is uniform and independent of the past, as in Γ.

Let H be the event that some pure input equals the pure input of a different step at a pass-6 index; it is contained in (ii)$\cup$(iii). Assume $\neg H$. We show by induction over the processing order that every step of the eager form has its pure input. A step at input x uses the output of the point at x. If the step is fresh, that output is $F(x)$, except that a pass-6 position keeps $\delta = C_j$; but the δ of a pass-6 point is never used inside its own message, since $t + L > 6L$. If the step is old and the point is exposed, its output is $F(x)$ by definition of F. If the point was created by a previous step, that step had input x by induction, and by $\neg H$ it is not at a pass-6 index, so its output is $F(x)$. In all cases the next inputs are pure. So an old step in passes 3–6 of T has a pure input above pass 2 that equals an exposed input, or a pure input of T at a lower index, or a pure input of a message processed before T. □

Write $E^*_{R,T}$ for $E_{R,T} \wedge \neg E_R \wedge \neg E_T$; then $E_{R,T} \subseteq E^*_{R,T} \cup E_R \cup E_T$. Event (i) is a visit above pass 1 in the sense of Section H.7.

H.6. Accidents with exposed classes

Replays. A *single replay* runs through the pure inputs of one message T in the order $T_1, \dots, T_{6L_T}$; a *pair replay* runs through those of R and then those of T, for R before T. Let N be the number of steps. A *class* is a set of steps with equal pure inputs. The exposed points are classes from the start, with constant input and constant output. Every other class c carries two variables, the halves $\gamma(c)$ and $\delta(c)$ of F at its input; for an exposed class c of the point i we put $\gamma(c) = \gamma_i$ and $\delta(c) = \delta_i$. The input halves of step U_t are

$$\alpha(U_t) = \begin{cases} X_A^U, & t = 1, \\ \gamma(c(U_{t-1})) \oplus \langle t-1 \rangle, & t \ge 2, \end{cases}$$

and $\beta(U_t) = X_t^U$ for $t \le L_U$, $\beta(U_t) = \delta(c(U_{t-L_U}))$ for $t > L_U$. Every half is a variable xor a constant, or a constant. The step U_{t-1} is the *predecessor* of U_t, and for $t > L_U$ the step U_{t-L_U} is its *register parent*. Chain equations involve only γ-variables and register equations only δ-variables. As in Section 5, each step is compared with every existing class, exposed or not, on both halves; a comparison is induced, impossible or open, a true open comparison is an accident, and its new half-equations are added to the system E, whose rank splits as (r_a, r_d), the ranks of its chain and register parts. A variable is *determined* if E implies that it equals a constant. A step that joins an exposed class is a *visit*. An accident whose target is exposed is a *visit accident*, and an accident

whose target is not exposed is a *W-accident*. Two steps in one class form a *merged pair*; it is *protected* if one of its steps is above pass 2. The chain equation of a merged pair is an *identity* if it holds syntactically: the two steps have the same index and their predecessors are in one class, or both have index 1 and equal first semiblocks. Classes are never merged after they are formed, so steps placed in different classes stay in different classes.

Lemma 37 (Cost of a sequence)**.** *For every accident sequence w of a single or pair replay,* $\Pr_F[$*the accident sequence starts with* $w] \le B^{-\mathrm{rank}(w)}$.

Proof. Given w and τ, the replay is a determined symbolic object: it says, for every comparison before the last accident of w, whether it is induced, impossible, open and false, or an accident, and hence gives the non-exposed classes $c_1, c_2, \dots$ in order of creation. Let $\Phi_1, \Phi_2, \dots$ be independent uniform 128-bit values, and realize F lazily by giving the input of the i-th non-exposed class created by the execution the value Φ_i. This is a valid realization of F, since F is uniform and independent at every non-exposed input. As long as the execution follows w, its classes are those of the replay, so its variables are the halves of the Φ_i and every equation of $E(w)$ holds for them. So the event lies in the event that the Φ_i satisfy $E(w)$, an affine system of rank $\mathrm{rank}(w)$ in 64-bit blocks with coefficients in $\{0, 1\}$, which has probability $B^{-\mathrm{rank}(w)}$. □

So every sequence of rank two costs at most $1/M$, and the proof of Lemma 20 consists in counting sequences.

H.7. Visits, clusters and messages at risk

Entering an exposed point. By (F1), step 1 of a long wrap is never a visit. Consider a visit at U_x whose predecessor U_{x-1} is not exposed. If $x \le L_U$, then $\beta(U_x) = X_x^U$ is a constant, and the exposed inputs with right half X_x^U are at most eight, one for each $m \le 8$, since the exposed inputs are distinct. Its chain half $\gamma(c(U_{x-1})) \oplus \langle x-1 \rangle$ must equal the short value Λ_j of the target. Unless this is already implied, the visit is an accident whose only new equation is

$$\gamma(c(U_{x-1})) = \lambda, \qquad \lambda = \Lambda_j \oplus \langle x-1 \rangle,$$

which we call a *pin*; it has rank one and at most eight targets. If $x > L_U$, then $\beta(U_x) = \delta(c(U_{x-L_U}))$ is a variable unless the register parent is exposed, and the visit needs a register equation in addition. After a visit at U_t with $t < L_U$, the input of U_{t+1} is $(\gamma_i \oplus \langle t \rangle, X_{t+1}^U)$, a constant, so whether U_{t+1} is a visit is determined by τ: visits inside a run of visits are induced.

Clusters. A *cluster* of U is a maximal run $[x, b]$ of consecutive visits. Its first step x is at least 2, and if $b < 6L_U$ its *exit* U_{b+1} is not exposed and has the constant chain half $K = \gamma_{i_b} \oplus \langle b \rangle$, where i_b is the exposed point visited at b. Let $\mathsf{BadVisit}_U$ be the event that U has a visit above pass 1 or two clusters, and $\mathsf{BadVisit}$ the union over all messages. These events are defined on pure paths, so $\mathsf{BadVisit}_U$ is decided by the single replay of U. Without $\mathsf{BadVisit}_U$, the message U has at most one cluster, and it lies in $[2, L_U]$; every visit accident of U is then a pin.

Messages at risk. A long wrap T of m_T bytes is *at risk* if some exposed answer satisfies

$$\gamma_i = \Lambda_{m_T} \oplus \langle b \rangle \quad \text{for some } 2 \le b \le \min(14, L_T - 2).$$

This is determined by τ. Let Risk be the event that the pure step-2 input of some message at risk is an exposed input. Section H.9 shows that one equation can produce a second cluster only in a message at risk whose step 2 is a visit, that is, only inside Risk.

Lemma 38 (Weight of the messages at risk)**.** $\mathbb{E}_{\text{ideal}}[\#\textit{at risk}] \le 2200\,\ell p/B$.

Proof. For $g \in \{0,1\}^{64}$ let $w(g)$ be the largest value of $1/L$ over the representations

$$g = \Lambda_m \oplus \langle b\rangle, \quad L = \left\lceil \frac{m}{8} \right\rceil, \quad 2 \le b \le \min(14, L-2),$$

with $9 \le m < 2^{32}$, and $w(g) = 0$ if there is none. This function does not depend on the adversary. If T is at risk through the answer i, then $w(\gamma_i) \ge 1/L_T$, so $1 \le L_T \sum_i w(\gamma_i)$. Summing over the messages at risk gives, pointwise,

$$\#\text{at risk} \le \ell \sum_{i=1}^{p} w(\gamma_i).$$

Bounding the maximum by a sum over the pairs (m,b) with $9 \le m < 2^{32}$ and $2 \le b \le 14$, each of which gives one value of g, and using that at most eight values of m have $\lceil m/8\rceil = L$,

$$\begin{aligned}\sum_g w(g) &\le 13 \sum_{m=9}^{2^{32}-1} \frac{1}{\lceil m/8\rceil} \le 104 \sum_{L=2}^{2^{29}} \frac{1}{L} \\ &\le 104\, H_{2^{29}} \le 104\,(1 + 29\ln 2) < 2200,\end{aligned}$$

where H_n is the nth harmonic number. Each γ_i is the left half of a uniform answer, uniform on $\{0,1\}^{64}$ whatever the previous queries, so $\mathbb{E}[w(\gamma_i)] = B^{-1}\sum_g w(g) < 2200/B$, and $\mathbb{E}[\#\text{at risk}] \le 2200\,\ell p/B$, even when the lengths are chosen after the answers. □

Proof of Lemma 19. Take any τ; the messages at risk are then determined. Let T be one of them. Its first pure input (Λ_{m_T}, X_1) is not exposed by (F1), so $F(T_1)$ is uniform and independent of τ. The pure input of step 2 is $(F_\gamma(T_1) \oplus \langle 1\rangle, X_2)$. At most eight exposed inputs have right half X_2, one for each $m \le 8$, so T_2 is exposed with probability at most $8/B$. The union over the messages at risk gives $\Pr_F[\mathsf{Risk}] \le 8\,\#\text{at risk}/B$, and Lemma 38 gives

$$\mathbb{E}_{\text{ideal}} \Pr_F[\mathsf{Risk}] \le \frac{8 \cdot 2200\,\ell p}{B^2} = \frac{17600\,p\ell}{M}. \qquad \square$$

H.8. A single pin

The analysis of one and two equations rests on a description of what one pin can do. Consider a replay and a prefix of its accident sequence whose chain part is spanned by one pin $\theta : \gamma(c_0) = \lambda$, made at the visit Y_x, with $x \le L_Y$ and $c_0 = c(Y_{x-1})$.

Lemma 39 (Recurrence)**.** *In this situation:*

(a) *there is no visit before θ, and before θ no class contains steps at two different indices;*
(b) *if some step U_y lies in c_0 and $y \ne x-1$, then $x = 2$, and U_{y-1} is exposed, or U_{y-2} is exposed and $U_{y-1} \in c_0$; without $\mathsf{BadVisit}_U$, $y \le L_U + 2$;*
(c) *a message U that has a visit has its first visit at index x, shares the first semiblock and $X_1, \dots, X_{x-1}$ with Y, and has $L_U = L_Y$. The same holds for every message with a step in c_0 at index $x-1$.*

If, moreover, the register part is zero and there is no visit above pass 1, then:

(d) *c_0 has no step above pass 1;*

(e) *a non-exposed step U_y with $y > L_U$ whose chain half is constant after substituting λ for $\gamma(c_0)$ has $y = L_U + 1$, its predecessor U_{L_U} is exposed or lies in c_0, and U has a visit.*

Proof. (a) Before θ the chain part of E is zero. A first visit cannot be induced: step 1 has a long first semiblock, and every other chain half before the first visit contains a variable that is not determined. So Y_x is the first visit. Two steps at different indices have chain halves with different counters, or one of them has index 1 and a long constant against a variable; with no chain equation, they cannot be in one class.

(b) Suppose $x > 2$. The chain half of c_0 is $\gamma(c(Y_{x-2})) \oplus \langle x-2 \rangle$, where Y_{x-2} is not exposed and, by (a), $c(Y_{x-2}) \neq c_0$; so this variable is not involved in θ. The chain equation of U_y with Y_{x-1} lies in the span of θ, so $\alpha(U_y)$ must contain the same variable with the same counter, which gives $y = x - 1$. Hence $x = 2$, and the chain half of c_0 is the long constant X_A^Y. Now $y \neq 1$, and $y = 2$ is impossible: either $U_1 \in c_0$, and then $\alpha(U_2) = \lambda \oplus \langle 1 \rangle = \Lambda_j$ is short, or U_1 would be exposed, against (F1). So $y \geq 3$, and $\alpha(U_y) = X_A^Y$ must be implied by θ: either U_{y-1} is exposed, or $U_{y-1} \in c_0$. In the second case, apply the same to U_{y-1}: U_{y-2} is exposed or lies in c_0, and both $U_{y-2}, U_{y-1} \in c_0$ would give $\gamma(c_0) \oplus \langle y-2 \rangle = X_A^Y = \gamma(c_0) \oplus \langle y-1 \rangle$. Without $\mathsf{BadVisit}_U$, exposed steps of U lie in pass 1, so $y - 2 \leq L_U$.

(c) Let U_y be the first visit of U. Its predecessor is not exposed, and its chain half equals a short constant, so $U_{y-1} \in c_0$. If $y - 1 \neq x - 1$, then (b) gives an exposed step of U before U_y, which is impossible. So $y = x$ and $U_{x-1} \in c_0$. For $2 \leq t \leq x - 1$, if (U_t, Y_t) is in one class, then, as the chain half of Y_t contains $\gamma(c(Y_{t-1}))$, which by (a) is not involved in θ, also $U_{t-1} \in c(Y_{t-1})$. Going down from $t = x - 1$ to $t = 1$ gives $X_A^U = X_A^Y$ and $X_1^U = X_1^Y$, so $L_U = L_Y$ by (F2), and then the right halves at $t \leq x - 1$ give $X_t^U = X_t^Y$. The last sentence follows from the same argument started at $U_{x-1} \in c_0$.

(d) An occurrence of c_0 at index $x - 1$ lies in pass 1 by (c). A different-index occurrence U_y with $y > L_U$ has, by (b), $x = 2$ and $y \leq L_U + 2$. The class c_0 has constant right halves, while $\beta(U_{L_U+1}) = \delta(c(U_1))$ is a variable, so $y \neq L_U + 1$. If $y = L_U + 2$, then $\alpha(U_y) = \gamma(c(U_{L_U+1})) \oplus \langle L_U + 1 \rangle$ must equal X_A^Y; the predecessor is not exposed, so it lies in c_0, which was just excluded.

(e) The predecessor U_{y-1} is exposed or lies in c_0. An exposed step lies in pass 1, and by (d) so does a step of c_0; hence $y - 1 \leq L_U$ and $y = L_U + 1$. If U_{L_U} is exposed, U has a visit. If $U_{L_U} \in c_0$, its index is not $x - 1$, since by (c) that would give $L_U = L_Y \geq x$; so (b) gives an exposed step of U. □

Lemma 40 (No register equation)**.** *Suppose the register part of E is zero and no message of the replay has a visit above pass 1.*

(a) *Every non-exposed class is of* constant type*, all its steps having constant right halves, or of* variable type*, all its steps having the right half $\delta(h)$ for one non-exposed class h, its* register parent class*.*

(b) *If (U_k, U'_m) is merged and $\beta(U_k)$ is a variable, then $\beta(U'_m)$ is a variable and the register parents $(U_{k-L_U}, U'_{m-L_{U'}})$ are merged. Every step above pass 2 has a variable right half.*

(c) *A class with a step in pass 1 is of constant type, and all its steps lie in passes 1–2.*

Proof. (a) An equality between $\delta(h)$ and a constant, or between $\delta(h)$ and $\delta(h')$ with $h \neq h'$, is a register equation that is not an identity. With none in E, two steps in one class have identical right halves. (b) follows from (a). A step U_y with $y > 2L_U$ has the register parent U_{y-L_U} above pass 1, which is not exposed, so its right half is a variable. (c) A step in pass 1 has a constant

right half. A step U'_y of the class with $y > L_{U'}$ then has a constant right half, so its register parent is exposed, hence in pass 1, and $y \le 2L_{U'}$. □

H.9. One equation

Lemma 41 (Rank one)**.** *Let τ be non-colliding.*

(a) *In a single or pair replay, no accident sequence of rank at most one produces a protected merged pair or a visit above pass 1.*

(b) *In the single replay of U, if a sequence of rank at most one gives U two clusters, then U is at risk and U_2 is a visit.*

Proof. If no visit occurs, the exposed classes are never joined, the replay is the forward replay of Appendix A, and one equation keeps all merged pairs in passes 1–2. Suppose a visit occurs. A visit above pass 1 needs a constant right half, hence an exposed register parent, that is, a previous visit; so the first visit is in pass 1. It is not induced by a W-accident. Indeed, before any visit only the first steps have constant chain halves, so a W-accident determines a γ only if one of its two sides is a first step, and then the accident's step, and every step at the same index whose predecessor lies in the determined class, has a long chain half. The determined class has a step at another index only in the pattern $U_2 \to c(U_1)$, and then $\alpha(U_3) = \Lambda_{m\oplus 3}$, which is long for $L_U \ge 3$, where $m \ge 17$; for $L_U = 2$ only $m = 11$ gives a short value, at U_3, above pass 1, whose right half $\delta(c(U_1))$ is a variable. By Lemma 39(a) the first visit is not induced with no equation either. So the single accident is a pin $\theta : \gamma(c_0) = \lambda$ at a step Y_x, $x \le L_Y$, with $c_0 = c(Y_{x-1})$, and $(r_a, r_d) = (1, 0)$. The class c_0 has the pass-1 step Y_{x-1}, so all its steps have the constant right half X^Y_{x-1}.

(a) We claim that, for every message U of the replay and every $t \ge L_U + 1$, the step U_t is not exposed and not in c_0. We argue by induction over the processing order. For $t = L_U + 1$, $\beta(U_t) = \delta(c(U_1))$ is a variable, since U_1 is not exposed, and there is no register equation, while exposed inputs and the steps of c_0 have constant right halves. For $t \ge L_U + 2$, the step U_{t-1} is not exposed and not in c_0 by induction, so $\alpha(U_t)$ contains the variable $\gamma(c(U_{t-1}))$, which is not determined; hence U_t is not a visit. It is not in c_0 either: if $x = 2$, the chain half of c_0 is the constant X^Y_A; if $x > 2$, it is $\gamma(c(Y_{x-2})) \oplus \langle x - 2\rangle$, and equality would need $t - 1 = x - 2$.

The same variable shows that every merged pair containing a step U_t with $t \ge L_U + 2$ is a same-index pair (U_t, U'_t) of two different messages whose predecessors are merged: the implied chain equation must cancel $\gamma(c(U_{t-1}))$ against the same variable with the same counter. Now let (U_k, U'_m) be a protected merged pair, with $k > 2L_U \ge L_U + 2$. Then $m = k$ and $U \ne U'$, and, applying the same fact to whichever step has index at least 2 above its pass length, the pairs (U_t, U'_t) are merged for all t from k down to $\min(L_U, L_{U'}) + 1$. If $L_U \ne L_{U'}$, the pair at $\min(L_U, L_{U'}) + 1$ has the variable right half $\delta(c(\cdot_1))$ on the shorter side and a constant on the other, which needs a register equation. So $L_U = L_{U'} = L$, and (U_t, U'_t) is merged for every $t \in [L + 1, 2L + 1]$. For $u \in [1, L]$, the right halves at $u + L$ agree, so either (U_u, U'_u) is merged, or both steps are exposed with equal δ: a variable never equals a constant here, and distinct non-exposed classes have distinct δ-variables. Not all of $[1, L]$ is merged, since the two queries would then be equal. Let u be the largest unmerged index in $[1, L]$. Then U_u and U'_u are distinct exposed points with equal δ, and the pair at $u + 1$ is merged, so its chain halves give $\gamma(U_u) \oplus \langle u\rangle = \gamma(U'_u) \oplus \langle u\rangle$. The two exposed points would have the same output, against injectivity. Visits above pass 1 were excluded by the claim.

(b) In the single replay of U, the single accident is the pin θ at U_x, whose cluster $[x, b]$ is induced, and Lemma 39 applies. Let U_y be the first visit of a second cluster; by (a), $y \le L_U$. Its predecessor is not exposed and its chain half is short, so $U_{y-1} \in c_0$, at an index other than $x-1$. Lemma 39(b), applied to U_{y-1}, gives $x = 2$, so U_2 is a visit, and U_{y-2} is exposed, or U_{y-3} is exposed and $U_{y-2} \in c_0$. As the steps strictly between b and y are not exposed, $y-2 = b$ or $y-3 = b$, and in both cases the exit U_{b+1} lies in $c_0 = c(U_1)$. Its chain half then equals X_A^U, so $\gamma_{i_b} \oplus \langle b \rangle = \Lambda_{m_U}$. With the pin $\gamma(c_0) = \Lambda_j \oplus \langle 1 \rangle$ and the second visit at $(\Lambda_{j'}, X_y)$,

$$\Lambda_{j'} = \alpha(U_y) = \Lambda_j \oplus \langle 1 \rangle \oplus \langle y-1 \rangle,$$

so $y - 1 = 1 \oplus j \oplus j' \le 15$ since $j, j' \le 8$. Hence $2 \le b \le y-2 \le \min(14, L_U - 2)$, and $\gamma_{i_b} = \Lambda_{m_U} \oplus \langle b \rangle$: the message U is at risk. □

So, outside Risk, every event of Lemma 36 and every occurrence of BadVisit needs a sequence of rank at least two.

H.10. Rank two: a final visit

We first count the sequences that make $\mathsf{BadVisit}_T$ happen. The event $\mathsf{BadVisit}_T$ is decided by the single replay of T, which we stop at its first occurrence of $\mathsf{BadVisit}_T$. The visit that completes $\mathsf{BadVisit}_T$, a visit above pass 1 or the first visit of a second cluster, is the *final visit* T_y. Its chain half is right because (A) the final comparison pins $\gamma(c(T_{y-1}))$, or (B) $\gamma(c(T_{y-1}))$ is already determined, or (C) T_{y-1} is exposed. Its right half is (Z1) the constant X_y, for $y \le L$; (Z2) a constant δ_i, for $y > L$ with T_{y-L} exposed; or (Z3) the variable $\delta(c(T_{y-L}))$, which the final comparison sets equal to the right half of one of up to p exposed inputs. Case (C) with (Z1) continues a cluster and is not final, and (C) with $y > L$ forces $y - 1 = L$, the end of a pass-1 cluster, so that $\beta(T_{L+1}) = \delta(c(T_1))$ is a variable: (C) occurs only with (Z3) at $y = L + 1$.

Lemma 42 (Final visit)**.** *Outside* Risk*, the sequences of rank two of the single replay of T that reach* $\mathsf{BadVisit}_T$ *are covered by the families of Table 8, at most $128L_T^2 + 18L_Tp$ sequences.*

Proof. Write $L = L_T$ and let c_0 be the class pinned by the first pin, at x, with cluster $[x, b]$. A sequence of rank two has one accident of rank two, or two accidents of rank one. We first justify each row and then show that the list is complete. A row may count sequences that do not actually reach $\mathsf{BadVisit}_T$.

Row 1, case (A) with (Z3). The final comparison adds a chain and a register equation, so it is the whole sequence. There are at most $5L$ positions above pass 1 and p targets.

Row 2, case (A) with (Z1) or (Z2). A first pin at x with at most eight targets, and a second pin at the final visit. A second cluster in pass 1 starts at some $y \in [x+2, L]$, at most $(L-x-1)_+$ positions; a visit above pass 1 whose right half comes from the first cluster has $y - L \in [x, L]$, at most $L - x + 1$ positions. Each final position has at most eight targets, so the count is

$$64 \sum_{x=2}^{L} \big[(L-x-1)_+ + (L-x+1)\big] = 64\,(L^2 - 3L + 3) \le 64L^2.$$

Row 3, case (C) with (Z3). The first pin determines its induced cluster, which runs up to L, and the constant chain half of T_{L+1}; the final register pin has at most p targets: $8(L-1)p \le 8Lp$.

TABLE 8. RANK-TWO SEQUENCES ENDING AT THE FIRST $\mathsf{BadVisit}_T$ OF ONE MESSAGE OF PASS LENGTH L. THE ROWS OVERLAP.

Family	Sequences
1. First visit above pass 1, both halves new	$5Lp$
2. A second pin	$64L^2$
3. Cluster up to L, then a register pin at $L+1$	$8Lp$
4. $T_2 \to T_1$, then a register pin at T_3 ($L=2$)	p
5. Pin at step 2, return $T_L \to T_1$, register pin at $L+1$	$8p$
6. Pin, then a W-accident hitting c_0	$48L^2$
7. Pin, then a W-accident hitting the class of its own predecessor, a constant-input class	$16(L-1)$
8. $T_v \to T_1$, then a pin at $v+1$	$8L$
9. $T_{u+1} \to T_u$ with $u \geq 2$, then a pin at $u+2$	$8L$
10. $T_2 \to T_1$, then a pin at T_4	8
11. Pin at step 2, cluster up to L, register accident $T_{L+1} \to T_1$	8
Total	$\leq 128L^2 + 18Lp$

Row 4, case (B) with (Z3), the chain half determined by a W-accident. Before any visit, a rank-one W-accident determines a γ only by hitting T_1, and the determined class has a step at another index only for $T_2 \to T_1$; then $\alpha(T_3) = \Lambda_{m\oplus 3}$ is short only for $L = 2$ and $m = 11$, and T_3 is above pass 1. The final register pin has at most p targets.

Row 5, case (B) with (Z3), the chain half determined by the pin. Before the final accident the chain part is θ, the register part is zero, and there is no visit above pass 1, so Lemma 39 applies. The final visit T_y, $y > L$, has $T_{y-1} \in c_0$ at an index other than $x-1$, so $x = 2$, and by Lemma 39(d) $y - 1 \leq L$: hence $y = L+1$ and $T_L \in c_0 = c(T_1)$, a return to T_1. The pin is at step 2 with eight targets, the final position is $L+1$, and the final register pin has p targets: $8p$.

Row 6, case (B) with (Z1) or (Z2), a W-accident returning to c_0. Choose the first pin and the position of the W-accident; its target is c_0: $8(L-1)\cdot 6L \leq 48L^2$.

Row 7, case (B) with (Z1) or (Z2), a W-accident that hits the class of its own predecessor and determines its γ. After the first pin, a non-exposed step with a constant chain half is the exit T_{b+1}, or, when $x = 2$ and the exit returns to T_1, the step T_{b+2} after the return; outside Risk there is at most one such further step. An accident $T_{b+2} \to T_{b+1}$ (or $T_{b+3} \to T_{b+2}$) imposes $\gamma(c(T_{b+1})) = K \oplus \langle b+1\rangle$, and the next chain half, $K \oplus \langle b+1\rangle \oplus \langle b+2\rangle$, is a constant that can be a short value. For example, with $L = 5$, $X = (\Lambda_{40}; 0, 0, 1, 1, 0)$, exposed points $(\Lambda_1, 0) \mapsto (\Lambda_2, 7)$ and $(\Lambda_7, 0) \mapsto (100, 9)$, the pin $T_2 \to (\Lambda_1, 0)$ and the accident $T_4 \to T_3$, with $\gamma(c(T_3)) = \Lambda_3$, make $T_5 = (\Lambda_7, 0)$ an induced visit. Each first pin gives at most two such locations: $16(L-1)$.

Rows 8–10, a W-accident before the first pin. The patterns are $T_v \to T_1$ followed by the first pin at $v+1$, at most L positions and eight targets; $T_{u+1} \to T_u$ with $u \geq 2$ followed by the first pin at $u+2$, which pins one end of the chain edge and hence both ends, again at most $8L$; and $T_2 \to T_1$ followed by the first pin at T_4, where the first accident makes the chain half of T_3 the constant $X_A \oplus \langle 3\rangle$ and the pin determines $\gamma(c(T_3))$, at most 8.

Row 11, a register W-accident creating the return to c_0. For a later induced final visit the first pin is at step 2, the cluster runs up to L, and the register accident is $T_{L+1} \to T_1$, setting

$\delta(c(T_1)) = X_1$; then T_{L+2} has the constant chain half $\gamma(c_0) \oplus \langle L+1 \rangle$ and the right half δ_i of the exposed point visited at step 2. Both positions and the target are determined by the first pin: 8.

Completeness. A single accident of rank two with a final visit is row 1: a W-accident of rank two cannot make a first visit induced, since any γ it determines belongs to a predecessor class that does not recur. Now let the two accidents have rank one.

If the first is a pin and the second is a W-accident at the step z with target class c, a later induced visit must read a determined γ, and the only new step that reads one is $z+1$, which reads $\gamma(c)$. If the new chain equation determines a γ, then $\gamma(c)$ is determined only when $c = c_0$ or when c is the class of $z-1$; these are rows 6 and 7. A different newly determined class has no new step from which its γ is read. If the new chain equation is an edge between two classes, its only possible additional copies come from the rank-one return to T_1 and from the new step created by the accident. With one repeated end, the counters disagree. With both ends repeated, the only remaining configuration would need $w \oplus (w+1) = 1 \oplus (w+2)$ for some w, which is impossible: $w \oplus (w+1) \oplus (w+2)$ is $w+3$ for even w and $w-1$ for odd w, never 1. If the second accident adds only a register equation, its chain halves are constants, and for a later induced visit its target must be c_0, which is row 11. If the second accident is the final register pin, its chain half comes from the cluster (row 3) or from c_0 (row 5). If it is a pin, this is row 2.

If the first accident is a W-accident and the second is the first pin, the predecessor classes of the first accident give rows 8, 9 and 10; outside these patterns the newly determined class has no usable recurrence after the cluster, a return to T_1 with $\gamma(c(T_1))$ not determined gives no constant, and the value determined by $T_2 \to T_1$ is a long format value at every relevant pass-1 position unless row 10 supplies the second determined class. If the second accident is the final register pin, this is row 4. Two W-accidents cannot make a first visit induced: before the first visit every right half above pass 1 is a variable, so the first visit is in pass 1, and before it the values determined by W-accidents alone have the form $X_A \oplus$ an xor of pass-1 counters, which is long for $L \geq 3$; for $L = 2$, step 2 is already the W-accident and no pass-1 position remains. The branches through a pin at step 2 of a message at risk lie in Risk.

Total. The rows without p give $64L^2 + 48L^2 + 16(L-1) + 8L + 8L + 8 + 8 = 112L^2 + 32L \leq 128L^2$, and the rows with p give $(13L+9)p \leq 18Lp$, using $L \geq 2$. □

H.11. Rank two: events with a visit

Now consider a replay, single or pair, stopped at its first protected merged pair, with no BadVisit up to that point and outside Risk. If no visit occurs, the replay is the forward replay of Appendix A. For a pair this gives at most $40L_RL_T$ sequences (Table 5). For a single message it gives at most $\binom{6L}{2} \leq 18L^2$: there are no shared prefixes, a chain identity never holds inside one message, two equal-constants accidents cannot reach pass 3 since the chain of register parents would need three chain edges, and an accident of rank two induces nothing but itself.

Lemma 43 (Structure)**.** *Without* BadVisit*, an accident sequence of rank two with a visit contains a pin, and it has one of three forms: two pins, in different messages; a pin θ and a W-accident ω that adds a register equation, with ω after θ; or a pin θ and a W-accident ω that adds a chain equation e.*

Proof. Without BadVisit, every visit is in pass 1, and every visit accident is a pin, of rank one. The first visit is not induced by W-accidents. With one W-accident of rank one this was shown in the proof of Lemma 41. A W-accident of rank two determines γ-values only of predecessor classes that do not recur. With two W-accidents, every right half above pass 1 is a variable before

the first visit, so the first visit is in pass 1, and the values determined by W-accidents alone have the form $X_A \oplus$ an xor of pass-1 counters, which is long for $L \geq 3$, while for $L = 2$ step 2 is already a W-accident and no pass-1 position remains. So the sequence contains a pin, and its other accident has rank one. Two pins in one message would start two clusters, so two pins are in different messages. A W-accident that adds only a register equation cannot come first: in the state with no equation, the only open comparison of rank one with a chain identity is the prefix extension of Appendix A, which needs a common prefix of depth $\min(L_R, L_T)$ with $L_R \neq L_T$; but two long wraps with a common first semiblock have equal pass lengths by (F2), and equal lengths with a full common prefix would make the queries equal. □

Two pins. Two pins in different messages have at most $(8L_R)(8L_T) = 64L_RL_T$ choices. We charge all of them, whatever follows.

Lemma 44 (A pin and a register equation)**.** *The sequences consisting of a pin θ and a later W-accident ω that adds a register equation are at most $24L^2$ for a single message and at most $(96 + 224)L_RL_T$ for a pair.*

Proof. The chain equation of ω is implied by θ. After substituting λ for $\gamma(c_0)$ and the constants of the exposed classes, it is either an identity, with the same variable and the same counter on both sides, or an equality between two constants.

Identity type. The two steps have the same index and predecessors in one class. Inside one message this is impossible, so ω is a same-index accident between the two messages. With lengths a, b there are at most $8(a+b)$ pins and $6\min(a, b)$ positions, so at most $48(a+b)\min(a, b) \leq 96ab$ sequences.

Constant type. Both chain halves are constant after θ. Immediately before ω the state has rank $(1, 0)$, and each message has at most four non-exposed steps with a constant chain half. Such a step U_y has $y = 1$, or an exposed predecessor, or its predecessor in c_0. If U has a cluster $[x, b]$ (its first visit is at x by Lemma 39(c)), the exposed predecessor gives the exit $b+1$; a predecessor in c_0 at index $x - 1$ gives the exposed step x; and a predecessor in c_0 at another index gives, by Lemma 39(b) and the single cluster, the indices $b+2$ and $b+3$. If U has no visit, only $y = 1$ and $y = x$ remain, by Lemma 39(b),(c). So the candidates are $\{1, b+1, b+2, b+3\}$ or $\{1, x\}$, and they are determined by the replay with the pin alone. For one message there are at most $\binom{4}{2} = 6$ pairs of candidates and $8L$ pins, at most $48L \leq 24L^2$ sequences; for a pair, at most $\binom{8}{2} = 28$ pairs and $8(a + b)$ pins, at most $224(a + b) \leq 224ab$, since $a + b \leq ab$ for $a, b \geq 2$. □

It remains to treat a pin and a chain equation, the case $(r_a, r_d) = (2, 0)$. In the rest of this section the chain part is spanned by $\theta : \gamma(c_0) = \lambda$, with $c_0 = c(Y_{x-1})$, and the equation e of the W-accident ω; the register part is zero, and Lemma 40 applies. For a merged pair Q, let $\nu(Q)$ be the *reduced support* of its chain equation: the formal sum, over $\mathbb{F}_2$, of the predecessor classes of its two steps that are neither exposed nor c_0, with nothing for a step of index 1. Since the chain equation of Q lies in the span of θ and e, after substituting λ for $\gamma(c_0)$ it is trivial or equal to the reduction of e, so $\nu(Q) \in \{0, \nu_e\}$ with $\nu_e \neq 0$ the reduced support of e. A merged pair with $\nu(Q) = 0$ is of *identity type* or of *constant type*, the latter meaning that both chain halves are constants after substituting λ; a merged pair with $\nu(Q) = \nu_e$ is of *equation type*.

Lemma 45 (Parent support)**.** *Let Θ be the linear map on formal sums of classes that are neither exposed nor c_0, with $\Theta(c) = h$ if c is of variable type with register parent class $h \neq c_0$, and $\Theta(c) = 0$ otherwise. Then Θ is nilpotent. If $Q_0 = (U_k, U'_m)$ is merged and $\beta(U_k)$ is a variable, then $Q_1 = (U_{k-L_U}, U'_{m-L_{U'}})$ is merged, $\nu(Q_1) = \Theta(\nu(Q_0))$, and $\nu(Q_1) = 0$.*

Proof. If $\Theta(c) = h \neq 0$ and y is the step of c with the smallest index, then h contains the register parent of y, of smaller index; so Θ strictly decreases the smallest index of a class, and some power of Θ is zero. Q_1 is merged by Lemma 40(b). Consider the contribution of U_k; that of U'_m is the same with primes. Let $c = c(U_{k-1})$. If c is of variable type with register parent class h, then $U_{k-1-L_U} \in h$, and the predecessor of U_{k-L_U} contributes h to $\nu(Q_1)$, or nothing when $h = c_0$; in both cases this is $\Theta(c)$. If c is exposed, then $k - 1 \leq L_U$, so $U_{k-L_U} = U_1$ contributes nothing. If c is of constant type, in particular if $c = c_0$ (Lemma 40(c)), then either $k - 1 \leq L_U$, so again $U_{k-L_U} = U_1$, or U_{k-1-L_U} is exposed and contributes nothing. Equal classes contribute equally, so cancellations agree, and $\nu(Q_1) = \Theta(\nu(Q_0))$. If $\nu(Q_0) = 0$ then $\nu(Q_1) = 0$. If $\nu(Q_0) = \nu_e$ then $\Theta(\nu_e) \in \{0, \nu_e\}$, and $\Theta(\nu_e) = \nu_e$ would give $\Theta^j(\nu_e) = \nu_e \neq 0$ for every j. □

Lemma 46 (A pin and a chain equation)**.** *The sequences consisting of a pin θ and a W-accident ω adding a chain equation, and producing a first protected merged pair, are none for a single message, none for a pair of unequal lengths, and at most $(48 + 96 + 768 + 384)L^2$ for a pair of equal lengths L.*

Proof. Let Q_0 be the first protected merged pair in processing order; if its newly processed step has an accident, Q_0 is the pair of that accident. Before Q_0, no step above pass 2 shares a class with another step. We orient $Q_0 = (U_k, U'_m)$ so that $k > 2L_U$; then $\beta(U_k)$ is a variable (Lemma 40(b)), and $Q_1 = (U_{k-L_U}, U'_{m-L_{U'}})$ is merged before Q_0.

Case 1: $\omega = Q_0$. By Lemma 41, the sequence up to Q_0 has rank two, so θ precedes ω. By Lemma 45, Q_1 is of identity or constant type, and it is merged before ω, when the chain part is θ alone. If Q_1 is of constant type, its step U_{k-L_U} is above pass 1 with a constant chain half, so Lemma 39(e) gives $k - L_U = L_U + 1$ and a visit in U. The right half of U_{L_U+1} is the variable $\delta(c(U_1))$, so $U'_{m-L_{U'}}$ is above pass 1 as well, and the same argument gives $m = 2L_{U'} + 1$ and a visit in U'. If Q_1 is of identity type, follow its diagonal down; every pair on it is merged while the chain part is θ alone, so it is of identity or constant type. It cannot reach index 1: for unequal lengths it would cross index $\min(L_U, L_{U'}) + 1$, where one right half is a variable $\delta(c(\cdot_1))$ and the other a constant; for equal lengths it would cover $[1, L]$ and make the two queries equal. So it ends at a merged pair of constant type at some index $j > 1$ whose predecessors are not merged. They cannot both lie in c_0, so one of them is exposed, at an index at least x by Lemma 39(c); the other is exposed or lies in c_0 at an index other than $x - 1$, as $j - 1 \geq x$, and Lemma 39(b) gives a visit in its message too. In both cases both messages have a visit, so by Lemma 39(c) they share the pin and have equal lengths L, they are different, since $k = m$ for two steps of one message would be the same step, and Q_0 is a same-index accident. As R is processed first and both messages visit before ω, the pin is in R. So for a given pair there are at most $8(L - 1)$ pins and $6L$ positions: $48L^2$.

Case 2: $\omega \neq Q_0$, so Q_0 is induced. Suppose $k - 1 > 2L_U$. Then $q = c(U_{k-1})$ is not exposed and has no other step before Q_0. If $\gamma(q)$ cancels against the other side of the chain equation of Q_0, the other predecessor lies in q, a protected merged pair before Q_0. Otherwise $\gamma(q)$ occurs in the span of θ and e; not in θ, by Lemma 40(c); so in e. If the step supplying it in e is not U_{k-1}, then q had two steps before Q_0; if it is U_{k-1}, then ω involves U_k, so it is a protected merged pair, and by the choice of Q_0 it would be Q_0. Both are excluded, so $k = 2L_U + 1$. Then Q_1 contains U_{L_U+1}, whose right half $\delta(c(U_1))$ is a variable, so by Lemma 40(b) $m > 2L_{U'}$ and $Q_2 = (U_1, U'_{m-2L_{U'}})$ is merged. The same argument applied to U'_m gives $m = 2L_{U'} + 1$, so $Q_2 = (U_1, U'_1)$, and (F2) gives $L_U = L_{U'} = L$. A single message would merge a step with itself, so $U \neq U'$, and the two messages have equal first semiblocks and equal X_1.

Let J be the set of indices t with (U_t, U'_t) merged; it contains 1, $L+1$ and $2L+1$, and misses some $d \in [2, L]$, since the queries differ. Let $[s, 2L+1]$ be the maximal run of J ending at $2L+1$. Then $s > L+1$. Otherwise all of $[L+1, 2L+1]$ lies in J, and by Lemma 40(a),(b) every $u \in [1, L]$ is in J or carries two distinct exposed points with equal δ. The largest $u \in [1, L]$ outside J then has $u+1 \in J$ (using $L+1 \in J$ if $u = L$), and the chain halves at $u+1$ give equal γ for the two exposed points, against injectivity. So at the bottom s the predecessor classes $V_1 = c(U_{s-1})$ and $V_2 = c(U'_{s-1})$ are distinct and not exposed, since $s - 1 > L$, and the chain equation at s is

$$\gamma(V_1) = \gamma(V_2).$$

At most one of V_1, V_2 is c_0, and the other carries a variable that θ does not involve, so this equation is not implied by θ alone; after substituting λ it is the reduction of e. If neither class is c_0, the reduction has two variables and constant 0; then both sides of the chain equation of ω have variable predecessors with equal counters, and ω is a same-index accident between the messages. If one of them is c_0, the reduction is $\gamma(V) = \lambda$ for the other class V, and the other side of the chain equation of ω is constant after θ: if it comes from c_0, matching λ again forces equal counters and a same-index accident; otherwise that side is a step of index 1, or a step with an exposed predecessor, which, as ω is a W-accident, is a cluster exit. So ω is a same-index accident, or a *boundary accident*: one with a side of index 1 or a side that is a cluster exit. This covers accidents inside one message and accidents at different indices.

For a given pair of equal lengths, the pin has at most $16L$ choices, in either message. Same-index accidents ω, in either order with the pin, give $16L \cdot 6L = 96L^2$. If θ precedes ω and ω has a side of index 1 or a cluster exit, then, for a given pin, the replay up to ω has no further accident, each message has at most one cluster, and there are at most four such steps, two first steps and two exits, each against at most $12L$ steps: $16L \cdot 4 \cdot 12L = 768L^2$. If ω precedes θ, there is no visit before ω, so its constant side has index 1: $2 \cdot 12L \cdot 16L = 384L^2$. □

H.12. Rank three and more

Lemma 47 (Higher order)**.** *For a replay with $N \geq 12$ steps, the probability that a prefix of the accident sequence reaches rank three before any* BadVisit *is at most $5N^6/B^3$. For the single replay of a message stopped at its first* BadVisit*, the probability that the sequence up to the stopping point has rank at least three is at most $[5N^6 + 3pN^5]/B^3$. Summed over all single and pair replays, the contributions of rank at least three are at most $7721568\, q\sigma L_{\max}^5/B^3$.*

Proof. Put $A_N = N^2 + 8N$. Before BadVisit every accident has at most A_N choices: one with a non-exposed target is a step and a class, at most N^2 choices, and a visit accident is a pin in pass 1 whose right half is a given input semiblock, which leaves at most eight targets per step. Visits inside a cluster are induced, and a comparison with a non-exposed class whose halves involve a cluster constant is a W-accident. Every accident adds rank at least one, so the first prefix of rank at least three has at most three accidents, and by Lemma 37 its probability is at most B^{-3}. For $N \geq 12$, $A_N \leq \frac{5}{3}N^2$, so

$$\begin{aligned}\frac{A_N + A_N^2 + A_N^3}{B^3} &\leq \Big(\frac{125}{27} + \frac{25}{9N^2} + \frac{5}{3N^4}\Big)\frac{N^6}{B^3} \\ &\leq \frac{5N^6}{B^3}.\end{aligned}$$

In a single replay stopped at its first BadVisit, either rank three is reached before, or the final accident, a visit above pass 1 with at most Np choices, completes rank three after at most two accidents. This adds

$$\frac{Np(A_N + A_N^2)}{B^3} \le \Big(\frac{25}{9} + \frac{5}{3N^2}\Big)\frac{pN^5}{B^3} \le \frac{3pN^5}{B^3}.$$

With $N = 6L_T$ for single replays and $N = 6(L_R + L_T)$ for pairs, $5 \cdot 6^6 = 233280$, $(L_R + L_T)^6 \le 32(L_R^6 + L_T^6)$, $\sum_T L_T^6 \le \ell L_{\max}^5$ and $\sum_{R<T}(L_R^6 + L_T^6) \le (q-1)\sum_T L_T^6$, the replays reaching an event before BadVisit contribute

$$\begin{aligned}\frac{233280}{B^3}\Big[\sum_T L_T^6 + \sum_{R<T}(L_R + L_T)^6\Big] &\le \frac{233280\,(1 + 32(q-1))\,\ell L_{\max}^5}{B^3} \\ &\le \frac{7464960\,q\sigma L_{\max}^5}{B^3}.\end{aligned}$$

The single replays stopped at BadVisit contribute, with $3 \cdot 6^5 = 23328$,

$$\begin{aligned}\sum_T \frac{233280\,L_T^6 + 23328\,pL_T^5}{B^3} &\le \frac{(233280\,\ell + 23328\,pq)L_{\max}^5}{B^3},\end{aligned}$$

at most $256608\,q\sigma L_{\max}^5/B^3$ since $\ell \le \sigma$ and $p \le \sigma$. The sum is

$$\frac{(7464960 + 256608)\,q\sigma L_{\max}^5}{B^3} = \frac{7721568\,q\sigma L_{\max}^5}{B^3}.$$

□

H.13. Proof of Lemma 20

Let τ be non-colliding. By Lemma 36, Old implies one of the events (i)–(iii) on the pure paths. Event (i) is a visit above pass 1, contained in BadVisit, and $E_{R,T} \subseteq E^*_{R,T} \cup E_R \cup E_T$. Hence

$$\begin{aligned}\mathsf{Old} \subseteq \mathsf{Risk} \;\cup\; \bigcup_T \big(\mathsf{BadVisit}_T \cup E_T\big) \setminus \mathsf{Risk} \\ \cup \bigcup_{R<T} \big(E^*_{R,T} \setminus (\mathsf{BadVisit}_R \cup \mathsf{BadVisit}_T)\big) \setminus \mathsf{Risk}.\end{aligned}$$

The first event is bounded, in the ideal average, by Lemma 19. For the union over T, run the single replay of T up to its first occurrence of $\mathsf{BadVisit}_T$ or of a protected merged pair inside T; for the union over pairs, run the pair replay up to its first protected merged pair, which on $E^*_{R,T}$ is a pair between R and T, with no BadVisit in R or T. By Lemma 41, outside Risk the accident sequence up to the stopping point has rank at least two. The sequences of rank exactly two cost $1/M$ each by Lemma 37, and they are covered by the families of Table 9: for a single message, $18L^2$ with no visit, $128L^2 + 18Lp$ reaching $\mathsf{BadVisit}_T$ (Lemma 42), and $24L^2$ with a visit (Lemmas 43, 44 and 46); for a pair, $40L_RL_T$ with no visit and $(64 + 96 + 224 + 48 + 96 + 768 + 384)L_RL_T$ with a visit, where the families of Lemma 46 exist only for equal lengths, so that $L^2 = L_RL_T$. The sequences of rank at least three are bounded by Lemma 47.

TABLE 9. THE FAMILIES OF TABLE 6 AND WHERE THEY ARE DERIVED.

Family	Single	Pair	From
No visit	$18L^2$	40	App. A
Final BadVisit, no target factor	$128L^2$	–	L. 42
Final BadVisit, exposed target	$18Lp$	–	L. 42
Two accidents that are pins	–	64	L. 43
Pin, then a register W-accident, chain halves identical	–	96	L. 44
Pin, then a register W-accident, chain halves constant	$24L^2$	224	L. 44
Pin, then W-accident, first merge above pass 2	–	48	L. 46
Pin, induced merge, same-index W-accident	–	96	L. 46
Pin, induced merge, boundary accident after pin	–	768	L. 46
Induced merge, boundary accident before the pin	–	384	L. 46
Coefficient of L^2 / of $L_R L_T$	170	1720	

The sum. The coefficients are $18 + 128 + 24 = 170$ for a single message and $40 + 64 + 96 + 224 + 48 + 96 + 768 + 384 = 1720$ for a pair. Summing over messages and pairs, averaging over the non-colliding ideal transcripts, and adding Lemma 19 and Lemma 47,

$$\begin{aligned}&\mathbb{E}_{\mathrm{ideal}}\big[1_{\neg\mathrm{coll}}\,\Pr_\Gamma[\mathsf{Old}]\big]\\&\quad\le \frac{170\sum_T L_T^2 + 1720\sum_{R<T} L_R L_T + 18\,p\ell}{M}\\&\qquad + \frac{17600\,p\ell}{M} + \frac{7721568\,q\sigma L_{\max}^5}{B^3}.\end{aligned}$$

This is Lemma 20, the colliding transcripts being charged to the collision row of Table 4. □

Appendix I.
Complete Proof of Theorem 3, Part 2: Long Unwraps and the Sum

This appendix proves Lemma 21, derives every row of Table 7, and adds up the rows of Table 4, which proves Theorem 3. The scheme, the game, the process Γ, the good event $\mathcal{G}$ and the validity probabilities $V_k(P)$ are those of Section H.1. We use the cap of Lemma 34: on $\mathcal{G}$ no 18 points of P share a left output half or a right output half, so a known half is the corresponding half of at most $\mu = 17$ points of P. Throughout, $B = 2^{64}$, $M = B^2$ and $\sigma \le 2^{60}$.

I.1. The backward run of a long unwrap

Fix an ideal transcript τ, an outcome of Γ in $\mathcal{G}$ with point set P, and a long unwrap query with ciphertext $C = (C_A; C_1, \ldots, C_L)$, $L \ge 2$. The adversary does not unwrap the answer of a preceding wrap, so C is not the ciphertext of a wrap made before this query. Probabilities are over a uniform permutation π extending P. Where a bound involves hidden values of P or the ideal answers, it is also averaged over ideal transcripts and over Γ with the indicator of $\mathcal{G}$, as in Table 4; all other bounds hold for every good P.

The run. The unwrap computes $W^{-1}(C)$ through the steps $t = 6L, \ldots, 1$. In the unrolled view of Section 2, step t has input (a_{t-1}, d_{t-L}) and output $(a_t \oplus \langle t \rangle, d_t)$, and we write $(\alpha_t, \beta_t) = (a_{t-1}, d_{t-L})$ for its input. Then $\alpha_{6L+1} = C_A$, and step t requests the output

$$Y_t = (\alpha_{t+1} \oplus \langle t \rangle,\, z_t), \qquad z_t = \begin{cases} \beta_{t+L}, & t \le 5L, \\ C_{t-5L}, & t > 5L, \end{cases}$$

and obtains $(\alpha_t, \beta_t) = \pi^{-1}(Y_t)$. The unwrap recovers the first semiblock α_1 and the data semiblocks $\beta_1, \ldots, \beta_L$. It is *valid* only if $\alpha_1 = \Lambda_m$ with $8(L-1) < m \le 8L$, one of eight *legal* first semiblocks, and the padding bytes are zero. Since $m < 2^{32}$, a valid unwrap has $L \le 2^{29}$, so Lemma 21 covers every valid long unwrap, and we assume $L \le 2^{29}$ from now on.

Fresh steps and hits. Step t is *fresh* if Y_t is neither an output of P nor the request of a higher fresh step; its input is then drawn uniformly from the inputs used neither by P nor by higher fresh steps. Otherwise step t is a *hit*: a *P-hit* if Y_t is an output of P, and a *self-hit on t'* if $Y_t = Y_{t'}$ for a higher fresh step t', its *target*. Fresh requests are distinct and are not outputs of P, so every hit is of exactly one kind and has one target, and a self-hit reads the input of its target. A hit with no fresh step above it is a P-hit. Let F be the set of fresh steps, and $a = \min F$ when $F \ne \emptyset$. Since P has at most D points, a fresh step draws from at least $M - \bar{D}$ inputs, where $\bar{D} = D + 6L$, and fixing one half of its input to a given value costs at most

$$\rho = \frac{B}{M - \bar{D}}.$$

Parents and paid hits. The *parents* of step t are $t+1$, which supplies α_{t+1}, and $t+L$, which supplies β_{t+L}. A parent above $6L$ does not exist, and the corresponding half is then a ciphertext value. A hit with both parents fresh is a *double hit*. A hit with exactly one fresh parent is a *paid hit*, and its *private half* is α_{t+1} if the fresh parent is $t+1$ and β_{t+L} if it is $t+L$. A target of a paid hit is a point of P whose other output half equals the other half of the request, of which there are at most μ by the cap, or a higher fresh step, of which there are at most $6L$. We put

$$Q = \mu + 6L \le \tfrac{29}{2} L \qquad (L \ge 2).$$

Resources. The semiblock budget gives $\ell + 2p + L + 1 \le \sigma$, so $\bar{D} = 6\ell + p + 6L \le 6\sigma \le 3B/8$, $D/B \le 3/8$ and $p \le B/32$. Hence

$$\mu\rho \le \frac{18}{B}, \quad 8\rho \le \frac{9}{B}, \quad \rho^2 \le \frac{2}{M}, \quad \rho \le \frac{1 + 2^{-60}}{B},$$

and $L^2 \le 2^{58} = B/64$.

Visible and hidden values. The transcript fixes the exposed points, inputs and outputs, and the ciphertext C^R of every long wrap R, hence the output of the last step of R and the right halves of the pass-6 points of R. We call these values *visible*. We write R_s for the point of step s of a wrap R, and $(\alpha(w), \beta(w))$ and $(\gamma(w), \delta(w))$ for the input and output of a point w. Every point of P is exposed or was created by a step of a wrap. A point created in passes 1–5 has a uniformly drawn output, and a point created at a pass-6 position other than the last has a uniformly drawn left half and a visible right half. These *hidden* values are independent of each other and of τ. On $\mathcal{G}$ passes 3–6 are fresh, so every pass-6 position creates its own point.

I.2. Three tools for paid hits

Lemma 48 (First use)**.** *Let f be a fresh step. The half α_f enters no request before step $f-1$, and β_f enters no request before step $f-L$. Consequently the private half of a paid hit u enters a request for the first time at u; the other half of the request of u is a value of P, a ciphertext value, or a half that entered a request before u; and distinct paid hits have distinct private halves.*

Proof. The request of step t uses the left input half of step $t+1$ and the right input half of step $t+L$. Step $t+1$ has left input half α_f only if it is f or a self-hit on f, and a self-hit on f is a step $r<f$; so α_f is used at $f-1$, or at $r-1<f-1$. The same argument with step $t+L$ gives the claim for β_f. For a paid hit u with fresh parent $u+1$, the other half z_u is a ciphertext value, or the right input half of the hit $u+L$, which is a value of P or the half β_f of the target $f>u+L$ of a self-hit; β_f entered the request of step $f-L>u$. The case of the fresh parent $u+L$ is symmetric, with α_f entering the request of step $f-1>u$. Finally, the private half α_{u+1} or β_{u+L} determines u. □

Lemma 49 (Array model)**.** *Let the* array *consist of independent uniform 128-bit values, one* entry *for each step $t\le 6L$, and let the* array run *be the run in which a fresh step t takes entry t as its input, whether or not this input is in use. For every set E of runs,*

$$\Pr[E]\le\kappa\,\Pr_{\text{array}}[E],\qquad \kappa=\Big(\frac{M}{M-\bar D}\Big)^{6L}<1+10^{-9}.$$

Proof. Given P and C, a run is determined by the inputs of its fresh steps, and entry t is read only at step t, and only if that step is fresh. A run with k fresh steps has probability at most $(M-\bar D)^{-k}$ in the uniform completion, since each fresh input is uniform on at least $M-\bar D$ values. It has probability M^{-k} in the array run, which follows it whenever the k entries at its fresh steps equal its inputs. The ratio is at most $(M/(M-\bar D))^k\le\kappa$, and summing over the runs in E gives the inequality. Finally $\kappa\le\exp(6L\bar D/(M-\bar D))$ and $6L\bar D\le 6\cdot 2^{29}\cdot\frac{3}{8}B<2^{95}$, so the exponent is below 2^{-32}. □

In the array run we keep writing (α_t,β_t) for entry t; it is the input of step t when that step is fresh. All bounds of the array run carry over with the factor $\kappa<2$, the density factor of Appendix D, which accounts for drawing without replacement.

Lemma 50 (Paid hits at given steps)**.** *Fix steps $u_1>\cdots>u_h$, and for each i a kind of target with a number k_i: a point of P ($k_i=\mu$), a given step $t_i>u_i$ ($k_i=1$), any higher fresh step ($k_i=6L$), or any target ($k_i=Q$). The probability that every u_i is a paid hit with a target of its kind is at most $\prod_i k_i/B^h$ in the array run, and at most $2\prod_i k_i/B^h$ in the uniform completion.*

Proof. In the array run, reveal the halves of the entries in the order in which they first enter a request. The run up to the comparisons of a step is a function of the halves revealed before that step, and a half not yet revealed is uniform given them. Consider u_i on the event. By Lemma 48, its private half is not yet revealed, while the other half of its request and the requests of all higher fresh steps are. A target of the given kind is matched only if the private half equals one of at most k_i values determined by them: one for each point of P whose other output half equals the other half of the request, and one for each admissible higher fresh step. So the conditional probability is at most k_i/B, and the chain rule gives the product. Lemma 49 adds the factor $\kappa<2$. □

I.3. Hidden comparisons and the top walk

Lemma 51 (Hidden comparisons)**.** *Let y and z be determined by τ, and average over Γ with the indicator of $\mathcal{G}$.*

(a) *The probability that (y, z) is the output of a point of P that is neither exposed nor the last point of a wrap is at most $17/B + D/M$.*

(b) *Let R and s be determined by τ, with s in passes 3–6 of R and $s < 6L_R$, and let $w_0 = R_s$, which on $\mathcal{G}$ is created at step s. The probability that $(\gamma(w_0) \oplus y, z)$ is the output of a point of P other than w_0 is at most $17/B + D/M$. If $s \le 5L_R$, then $\Pr[\delta(w_0) = z] \le 1/B$.*

Proof. (a) A point created in passes 1–5 has a uniform output independent of τ, so it equals (y, z) with probability $1/M$, and there are at most D such points. A point created at a pass-6 position other than the last has a visible right half and a uniform hidden left half, so it can match only if its right half is z. The number of such points is determined by τ, and on $\mathcal{G}$ it is at most $\mu = 17$, since they are distinct points of P with equal right halves. Each matches with probability $1/B$.

(b) A point $w \ne w_0$ with a visible right half, that is, an exposed point, a last point, or a pass-6 point, matches only if $\delta(w) = z$, and on $\mathcal{G}$ at most 17 of them do. Each then needs $\gamma(w) = \gamma(w_0) \oplus y$, which has probability $1/B$, as $\gamma(w_0)$ is uniform and independent of $\gamma(w)$. A point created in passes 1–5 other than w_0 has an output independent of $\gamma(w_0)$ and matches with probability $1/M$. If $s \le 5L_R$, the right output half of w_0 is uniform. □

A *hidden comparison* is a comparison of the kind described in (a) or (b), including the comparison $\delta(w_0) = z$ of (b).

For $0 \le j < 6L$ we say that the *top above j is old* if steps $6L, \ldots, j + 1$ are all hits. They are then P-hits, their requests are computed from C and P alone, and this is a property of P. We call the points $w_{6L}, \ldots, w_{j+1}$ that these steps hit the *top walk*.

Lemma 52 (Top walk)**.** *Let $0 \le j < 6L$. Averaged over Γ with the indicator of $\mathcal{G}$, the probability that the top above j is old and the top walk makes a successful hidden comparison at a step above j is at most $18/B + D/M$. If moreover $j < 5L$ and the top above j is old, then one of the following holds.*

(W1) *The top walk makes a successful hidden comparison at a step above j.*

(W2) *Every w_t with $t > j$ is exposed.*

(W3) *For a wrap R with $L_R < L$ and $5L_R \le j$, the point w_t is exposed for $t > 6L_R$, and $w_t = R_t$ for $j < t \le 6L_R$.*

(W4) *The point w_t is exposed for $t \ge j + 2$, and w_{j+1} is the last point of a wrap.*

(W5) *C is the ciphertext of a wrap of pass length L made after the unwrap.*

The events (W2)–(W4) are determined by τ, and (W5) has ideal probability at most D/M.

Proof. Let u be the largest $t > j$ such that w_t is not exposed; if there is none, (W2) holds. Every request Y_t with $t \ge u$ is determined by τ: its left half is $C_A \oplus \langle 6L \rangle$ or $\alpha(w_{t+1}) \oplus \langle t \rangle$ with w_{t+1} exposed, and its right half is a ciphertext block or $\beta(w_{t+L})$ with $t + L > u$, so that w_{t+L} is exposed. If w_u is not the last point of a wrap, step u is a successful hidden comparison of kind (a), and (W1) holds.

Let $w_u = R_s$ be the last point of a wrap R, so $s = 6L_R$; note $L_R \ge 2$. If $u = j + 1$, (W4) holds. Otherwise step $u - 1$ is above j. Since $\alpha(R_s) = \gamma(R_{s-1}) \oplus \langle s - 1 \rangle$,

$$Y_{u-1} = \big(\gamma(R_{s-1}) \oplus \langle s - 1 \rangle \oplus \langle u - 1 \rangle,\ z_{u-1}\big),$$

where z_{u-1} is a ciphertext block or the right input half of an exposed point, and R_{s-1} is a pass-6 point other than the last. If $u \neq s$, the left half differs from $\gamma(R_{s-1})$, so $w_{u-1} \neq R_{s-1}$, and step $u-1$ is a successful hidden comparison of kind (b) with $w_0 = R_{s-1}$: (W1).

Let $u = s$; then $L_R \leq L$, since $u \leq 6L$. The counters of the unwrap and of R now agree: as long as $w_t = R_t$, the request of step $t-1$ is $(\gamma(R_{t-1}), z_{t-1})$, because $\alpha(R_t) = \gamma(R_{t-1}) \oplus \langle t-1 \rangle$. Suppose first that $L_R = L$. Then $u = 6L$, and $w_{6L} = R_{6L}$ gives $C_A = C_A^R$ and $C_L = C_L^R$. If $C = C^R$, then R does not precede the unwrap, and (W5) holds. Otherwise let t^* be the largest step with $5L < t^* < 6L$ and $C_{t^*-5L} \neq C^R_{t^*-5L}$. The walk follows R down to step t^*+1, and at step $t^* > 5L > j$ the request $(\gamma(R_{t^*}), C_{t^*-5L})$ differs from the output of R_{t^*}, so step t^* is a successful hidden comparison of kind (b): (W1).

Suppose now that $L_R < L$. For $5L_R \leq t-1 < 6L_R$ we have $t-1+L > 6L_R = u$, so z_{t-1} is a ciphertext block or the right input half of an exposed point, and it is determined by τ. Going down from u, at each step $t-1$ with $t-1 > j$ and $t-1 > 5L_R$ the walk either continues to R_{t-1}, when z_{t-1} equals the visible right half of R_{t-1}, or makes a successful hidden comparison of kind (b) with $w_0 = R_{t-1}$, which is (W1). If it continues down to step $j+1$ and $5L_R \leq j$, then (W3) holds. Otherwise it reaches R_{5L_R+1} with $5L_R > j$, and step $5L_R$ requests $(\gamma(R_{5L_R}), z_{5L_R})$, where R_{5L_R} was created in pass 5. This step hits R_{5L_R} only if $\delta(R_{5L_R}) = z_{5L_R}$, and another point only by a hidden comparison of kind (b); both are hidden comparisons, so (W1) holds. This proves the list of alternatives for $j < 5L$.

For the probability, fix τ and any $j < 6L$. Before its first hidden comparison, every comparison of the walk is decided by visible values: hits on exposed points and on last points compare visible outputs, and a hit on R_{t-1} in the aligned case compares the visible right half of R_{t-1} with z_{t-1}, the left halves being equal. So the walk up to its first hidden comparison, and that comparison itself, are determined by τ; if they lie above j, a successful hidden comparison above j requires this one to succeed. By Lemma 51 it succeeds with probability at most $17/B + D/M$, or at most $1/B + 17/B + D/M$ at the step $5L_R$. Under (W2)–(W4) the walk above j consists of visible steps only, so these events are determined by τ. For (W5), at most $q \leq D$ wraps follow the unwrap, and the ideal answer of each is a uniform string of $L+1 \geq 3$ semiblocks chosen after C, so (W5) has probability at most $D/B^3 \leq D/M$. □

Lemma 53 (Exposed prefix)**.** *For $2 \leq l \leq 2^{29}$ let*

$$\Phi_l = \{\Lambda_m \oplus \langle t \rangle : \ 1 \leq m \leq 8, \ 2l \leq t \leq 6l-1\},$$

let $\mathcal{E}_l$ be the event that at least $\lceil l/2 \rceil$ distinct exposed points have their left output half in Φ_l, and put $x = p/[B(1-p/M)]$. Sums over l run over $2 \leq l \leq 2^{29}$.

(a) *If the top above L is old and (W2), (W3) or (W4) holds with $j = L$, then $\mathcal{E}_L$ holds.*
(b) *$\sum_{l \geq 2} l \Pr[\mathcal{E}_l] \leq 4096\, p/B$ and $\sum_{l \geq 2} l^2 \Pr[\mathcal{E}_l] \leq 34570\, x$.*

Proof. (a) In each case the steps $6L, \ldots, 2L$ hit exposed points: (W2) gives $5L \geq 4L+1$ exposed steps at the top; (W3) gives $L_R \leq L/5$ and $6L - 6L_R \geq \frac{24}{5}L \geq 4L+1$; and (W4) gives $5L - 1 \geq 4L+1$. For $2L < t \leq 6L$ the input of w_t is $(\Lambda_{m_t}, \cdot)$ with $m_t \leq 8$, so $\gamma(w_{t-1}) = \Lambda_{m_t} \oplus \langle t-1 \rangle \in \Phi_L$, which gives $4L$ values of $t-1$. For a fixed exposed point and each m, at most one t solves $\gamma = \Lambda_m \oplus \langle t \rangle$, so one exposed point occurs at most eight times among $w_{6L-1}, \ldots, w_{2L}$, and at least $\lceil L/2 \rceil$ distinct exposed points occur.

(b) Since $t < 6 \cdot 2^{29} < 2^{32}$, the value $\Lambda_m \oplus \langle t \rangle$ has left 32 bits `A65959A6` and right 32 bits $[m]_{32} \oplus [t]_{32}$, which lies in the aligned block of 16 integers containing t, because $m < 16$. The $4l$ values of t meet blocks that cover at most $4l+30$ integers, so $|\Phi_l| \leq 4l+30$. Given the

preceding answers, each one-block answer has its left half in a fixed set H with probability at most $|H|/[B(1-p/M)]$. With $\nu = \lceil l/2 \rceil$, a union bound over the sets of ν one-block queries gives

$$\begin{aligned}\Pr[\mathcal{E}_l] &\le \binom{p}{\nu}\Big(\frac{|\Phi_l|}{B(1-p/M)}\Big)^{\nu} \\ &\le \frac{(|\Phi_l|x)^{\nu}}{\nu!} \le \frac{(8\nu+32)^{\nu}x^{\nu}}{\nu!},\end{aligned}$$

since $l \le 2\nu$ gives $|\Phi_l| \le 8\nu + 30$. The value $\nu = 1$ occurs only for $l = 2$, and $\nu \ge 2$ for $l \in \{2\nu-1, 2\nu\}$. Since $p \le B/32$, we have $x \le 1/31$ and $x^{\nu} \le x\,31^{1-\nu}$. With the weights 2 and $(2\nu-1)+2\nu = 4\nu-1$,

$$\sum_{l\ge 2} l\Pr[\mathcal{E}_l] \le x\Big[80 + \sum_{\nu\ge 2}\frac{(4\nu-1)(8\nu+32)^{\nu}}{31^{\nu-1}\,\nu!}\Big].$$

The terms up to $\nu = 16$, together with 80, add up to less than 2714; the term $\nu = 16$ is below 24, and from there on each term is less than $15/16$ of the previous one. So the bracket is below $2714 + 15\cdot 24 = 3074$, and with $x \le \frac{32}{31}\frac{p}{B}$ and $3074 \cdot \frac{32}{31} < 4096$ the first bound follows. With the weights 4 and $(2\nu-1)^2 + (2\nu)^2 = 8\nu^2 - 4\nu + 1$,

$$\sum_{l\ge 2} l^2\Pr[\mathcal{E}_l] \le x\Big[160 + \sum_{\nu\ge 2}\frac{(8\nu^2-4\nu+1)(8\nu+32)^{\nu}}{31^{\nu-1}\,\nu!}\Big].$$

The terms up to $\nu = 40$, together with 160, add up to less than 34567.7; the term $\nu = 40$ is below 0.78, and from there on each term is less than 0.74 of the previous one. So the bracket is below $34567.7 + 0.78\cdot 0.74/0.26 < 34570$. □

I.4. Step 1 fresh

If step 1 is fresh, its input is uniform on at least $M - \bar{D}$ values, and a valid unwrap needs its left half to be one of the eight legal first semiblocks. So, for every good P,

$$\Pr[\text{valid} \wedge \text{step 1 fresh}] \le 8\rho \le \frac{9}{B}.$$

This is the first claim of Lemma 21 and the row "Long unwrap, step 1 fresh" of Table 4.

I.5. Step 1 old: the families

Table 10 restates the rows of Table 7, in the same order, as events $\mathcal{F}_1, \ldots, \mathcal{F}_{13}$, which are defined in the paragraphs below. We bound each of them here, and Section I.7 shows that they cover the event that the unwrap is valid and step 1 is old.

$\mathcal{F}_1$: no fresh step. Let $\mathcal{F}_1$ be the event that the unwrap is valid and $F = \emptyset$. Every step is then a P-hit, so the top above 0 is old, and Lemma 52 applies with $j = 0$. Under (W2), α_1 is the left input half of an exposed point, which is Λ_m with $m \le 8$ and is not legal for $L \ge 2$. (W3) needs $5L_R \le 0$ and does not occur. Under (W4), step 1 hits the last point R_s of a wrap, and $\alpha_1 = \gamma(R_{s-1}) \oplus \langle s-1 \rangle$ is legal with probability at most $8/B$, since $\gamma(R_{s-1})$ is a hidden uniform

TABLE 10. THE ROWS OF TABLE 7 AS EVENTS. HERE $a = \min F$, AND c IS A FRESH STEP ABOVE L.

	Event	Cost
$\mathcal{F}_1$	$F = \emptyset$, valid	$18/B + 2LD/M$
$\mathcal{F}_2$	a double hit	$12LD/M + 72L^2/M$
$\mathcal{F}_3$	step 1 a paid hit on P	$36/B$
$\mathcal{F}_4$	two paid hits on P	$10404L^2/M$
$\mathcal{F}_5$	$f > 5L$ fresh, $f-1$ a self-hit on f, $f-L$ a P-hit	$17L^2/M$
$\mathcal{F}_6$	a self-hit with a fresh parent and a paid hit on P as parents	$2040L^2/M$
$\mathcal{F}_7$	step 1 a paid self-hit, valid	$9/B + 96L^2/M$
$\mathcal{F}_8$	a self-hit below an old top with a hidden comparison	$19845L^2/M$
$\mathcal{F}_9$	old top above L, $a-1$ a paid hit on P	$331L/M$ $+ 73728\,p/M$
$\mathcal{F}_{10}$	old top above L, $a-1$ a paid self-hit	$10L^2/M$ $+ 18000\,p/M$
$\mathcal{F}_{11}$	$a > L$; $a-1$, $a-L$ with one self-target	$2040L^2/M$
$\mathcal{F}_{12}$	two paid hits, targets fixed by a and c	$818L^2/M$
$\mathcal{F}_{13}$	three paid hits, or a template	$6585030\,L^6/B^3$

left half. When (W2) or (W4) holds, the walk makes no hidden comparison, and these events are determined by τ. So for each τ only one of the bounds of (W1) and (W4) applies, and

$$\begin{aligned}\Pr[\mathcal{F}_1] &\le \max\Big(\frac{18}{B} + \frac{D}{M}, \frac{8}{B}\Big) + \frac{D}{M} \\ &= \frac{18}{B} + \frac{2D}{M} \le \frac{18}{B} + \frac{2LD}{M}.\end{aligned}$$

$\mathcal{F}_2$: double hit. Let $\mathcal{F}_2$ be the event that some hit has both parents fresh. Fix a step t and a potential target: a point of P or a higher step t'. The two halves of $Y_t = (\alpha_{t+1} \oplus \langle t \rangle, \beta_{t+L})$ are drawn at the different steps $t+L$ and $t+1$, in this order. When step $t+L$ draws, the right half of the target is determined: for a step $t' > t$ it is a ciphertext block or $\beta_{t'+L}$ with $t'+L > t+L$. When step $t+1$ draws, the left half $\alpha_{t'+1} \oplus \langle t' \rangle$ is determined, since $t'+1 > t+1$. Each draw then fixes one half at cost ρ. With $6L$ steps and at most $\bar{D} = D + 6L$ targets,

$$\Pr[\mathcal{F}_2] \le 6L\bar{D}\rho^2 \le \frac{2 \cdot 6L(D+6L)}{M} = \frac{12LD}{M} + \frac{72L^2}{M}.$$

$\mathcal{F}_3$: step 1 a paid hit on P. Let $\mathcal{F}_3$ be the event that step 1 is a P-hit with exactly one fresh parent. If the fresh parent is step 2, its input is drawn at step 2, and step 1 compares $\alpha_2 \oplus \langle 1 \rangle$ with the left halves of the points of P whose right half is $z_1 = \beta_{L+1}$, the right input half of the hit $L+1$, known at that time. There are at most μ such points, and the cost is at most $\mu\rho \le 18/B$. If the fresh parent is $L+1$, the private half is β_{L+1}; Lemma 50 with $h = 1$ gives the array probability μ/B, and Lemma 49 gives at most $\kappa\mu/B \le 18/B$. So $\Pr[\mathcal{F}_3] \le 36/B$.

$\mathcal{F}_4$**: two paid hits on** P**.** Let $\mathcal{F}_4$ be the event that two steps are paid hits on points of P. Lemma 50 with $h = 2$ and $k_1 = k_2 = \mu$, over at most $\binom{6L}{2} \le 18L^2$ pairs of steps, gives

$$\Pr[\mathcal{F}_4] \le \frac{2 \cdot 18L^2\mu^2}{M} = \frac{10404\, L^2}{M}.$$

The factor 2 is the density factor $\kappa < 2$ of Lemma 49, which accounts for drawing without replacement.

$\mathcal{F}_5$**: one fresh step in the top pass.** Let $\mathcal{F}_5$ be the event that some fresh step $f > 5L$ has a self-hit on f at step $f - 1$ and a P-hit at step $f - L$. Fix f and, in the array run, all entries other than entry f. The request Y_f does not involve entry f, and the self-hit at $f - 1$ requires $\alpha_f \oplus \langle f{-}1\rangle = \gamma(Y_f)$, which fixes α_f. Given α_f, the run down to step $f{-}L{+}1$ does not involve β_f, which first enters a request at $f - L$ by Lemma 48. So the left half of Y_{f-L} is determined, and $\beta_f = z_{f-L}$ must be the right half of one of at most μ points of P with that left half. At most μ values of entry f qualify, and with L choices of f and Lemma 49,

$$\Pr[\mathcal{F}_5] \le \frac{2\mu L}{M} = \frac{34L}{M} \le \frac{17L^2}{M}.$$

$\mathcal{F}_6$**: self-hit with a fresh and a paid-hit parent.** Let $\mathcal{F}_6$ be the event that some self-hit $u \le 5L$, on a target t, has one fresh parent, and its other parent is a paid hit on P. Then u is a paid hit with the given target t, and its other parent is a paid hit on P at a step determined by u and by which parent is fresh. Lemma 50 with $h = 2$ and the numbers 1 and μ, over the $5L \cdot 6L$ choices of (u, t) and the two choices of the fresh parent, gives

$$\Pr[\mathcal{F}_6] \le \frac{2 \cdot 2 \cdot 30L^2\mu}{M} = \frac{2040\, L^2}{M}.$$

$\mathcal{F}_7$**: step 1 a self-hit on a legal fresh step.** Let $\mathcal{F}_7$ be the event that the unwrap is valid and step 1 is a self-hit, on a target d, with exactly one fresh parent. Step 1 reads the input of d, so $\alpha_d = \alpha_1$ is legal. If $d = 2$, step 2 is fresh, and its input has a legal left half with probability at most $8\rho \le 9/B$; here legality and the self-hit concern the same half, and only legality is counted. If $d \ge 3$, fix the fresh parent, step 2 or step $L{+}1$, and in the array run fix all halves other than α_d and the private half of step 1. The equation $Y_1 = Y_d$ fixes the private half, because the corresponding half of Y_d is formed from inputs of steps above d and, by Lemma 48, involves neither the private half nor α_d. Moreover α_d takes one of eight legal values, so the array probability is at most $8/M$. With at most $6L$ choices of d, two choices of the fresh parent, Lemma 49 and $192L \le 96L^2$,

$$\Pr[\mathcal{F}_7] \le \frac{9}{B} + \frac{2 \cdot 6L \cdot 2 \cdot 8}{M} \le \frac{9}{B} + \frac{96L^2}{M}.$$

$\mathcal{F}_8$**: self-hit, hidden comparison above.** Let $\mathcal{F}_8$ be the event that some self-hit $u \le 5L$, on a target t, has as parent $u + L$ a P-hit with no fresh parent and as parent $u + 1$ a fresh step or a paid hit on P, while the top above $u + L$ is old and the top walk makes a successful hidden comparison above $u + L$. The last two conditions concern P only, and by Lemma 52 they have probability at most $18/B + D/M$ for each u; for $u = 5L$ there is no step above $u + L = 6L$, and they cannot hold. For a fixed P, the remaining condition costs at most $\mu\rho$. If $u + 1$ is fresh, the self-hit requires $\alpha_{u+1} \oplus \langle u\rangle = \gamma(Y_t)$, with Y_t known when α_{u+1} is drawn, at cost ρ. If $u + 1$ is a paid hit on P, its fresh parent is $u + 2$, since $u + 1 + L$ lies above $u + L$ and is a P-hit, and the cost is $\mu\rho$. With the $30L^2$ choices of (u, t) and the two patterns,

$$\begin{aligned}\Pr[\mathcal{F}_8] &\le 60L^2\mu\rho\Big(\frac{18}{B} + \frac{D}{M}\Big) \\ &\le 60L^2 \cdot \frac{18}{B}\Big(\frac{18}{B} + \frac{3}{8B}\Big) = \frac{19845\, L^2}{M}.\end{aligned}$$

$\mathcal{F}_9$ and $\mathcal{F}_{10}$: exposed top. Let $\mathcal{F}_9$ (respectively $\mathcal{F}_{10}$) be the event that $F \neq \emptyset$, $a = \min F \geq 2$, no fresh step lies above L, and step $a-1$ is a P-hit (respectively a self-hit). Every step above L is then a hit with no fresh step above it, so a P-hit, and the top above L is old. Fix P with this property and a step $2 \leq a \leq L$, and suppose that a is fresh. The parent $a-1+L > L$ of step $a-1$ is a P-hit, so the request of $a-1$ is $(\alpha_a \oplus \langle a-1 \rangle, z_{a-1})$, with z_{a-1} known when α_a is drawn. A P-hit at $a-1$ then costs at most $\mu\rho$, and a self-hit on a given $t \in [a, L]$ costs at most ρ, since Y_t is also known. Summing over the at most L values of a, and over the $L(L-1)/2 \leq L^2/2$ pairs (a,t) with $2 \leq a \leq t \leq L$, the conditional probabilities of $\mathcal{F}_9$ and $\mathcal{F}_{10}$ given P are at most $\mu L\rho$ and $L^2\rho/2$.

By Lemma 52 with $j = L < 5L$, an old top above L implies one of (W1)–(W5). Under (W5) the unwrap repeats the computation of a wrap R of pass length L: if steps above t hit $R_{t+1}, R_{t+2}, \ldots$, then Y_t is the output of R_t, because $\alpha(R_{t+1}) \oplus \langle t \rangle = \gamma(R_t)$ and z_t is C^R_{t-5L} or $\beta(R_{t+L}) = \delta(R_t)$. So every step is a P-hit, $F = \emptyset$, and neither family occurs. Under (W2)–(W4), Lemma 53(a) gives the event $\mathcal{E}_L$ of the ideal transcript. Since L may depend on τ, we use $L\,1_{\mathcal{E}_L} \leq \sum_{l \geq 2} l\,1_{\mathcal{E}_l}$ and $L^2\,1_{\mathcal{E}_L} \leq \sum_{l \geq 2} l^2\,1_{\mathcal{E}_l}$. With Lemma 53(b),

$$\begin{aligned}\Pr[\mathcal{F}_9] &\leq \mu L\rho\Big(\frac{18}{B} + \frac{D}{M}\Big) + \mu\rho \sum_{l \geq 2} l \Pr[\mathcal{E}_l] \\ &\leq \frac{18}{B}\Big(\frac{18L}{B} + \frac{3L}{8B} + \frac{4096\,p}{B}\Big) \leq \frac{331L}{M} + \frac{73728\,p}{M},\end{aligned}$$

using $18\,(18 + \frac{3}{8}) = 330.75$, and

$$\begin{aligned}\Pr[\mathcal{F}_{10}] &\leq \frac{L^2\rho}{2}\Big(\frac{18}{B} + \frac{D}{M}\Big) + \frac{\rho}{2} \sum_{l \geq 2} l^2 \Pr[\mathcal{E}_l] \\ &\leq \frac{10L^2}{M} + \frac{18000\,p}{M},\end{aligned}$$

using $\rho \leq (1 + 2^{-60})/B$, $\frac{1}{2}(18 + \frac{3}{8}) < 10$, $x \leq \frac{32}{31}\frac{p}{B}$ and $\frac{1}{2} \cdot 34570 \cdot \frac{32}{31} < 17843$. The single payment below L multiplies the prefix bound, so no term of order p/B arises.

Two selected paid hits. The remaining families concern two paid hits chosen as follows. If $a > L$ is fresh, these are $a-1$ and $a-L$, with the fresh parent a, the first through α_a and the second through β_a. If $2 \leq a \leq L$ and $L+2 \leq c \leq 6L$ are fresh, these are the *lower* step $a-1$ and the *upper* step $c-1$, with the fresh parents a and c. When these steps are self-hits, we write t_1 for the target of $a-1$ and t_2 for the target of the other step, $a-L$ or $c-1$.

$\mathcal{F}_{11}$: two paid hits, one self-target. Let $\mathcal{F}_{11}$ be the event that for some fresh $a > L$ both $a-1$ and $a-L$ are paid hits, one with a target in P and the other with a self-target. Lemma 50 with the numbers μ and $6L$, the two choices of the step with the self-target, and the at most $5L$ choices of a give

$$\Pr[\mathcal{F}_{11}] \leq \frac{2 \cdot 5L \cdot 2\mu \cdot 6L}{M} = \frac{2040\,L^2}{M}.$$

$\mathcal{F}_{12}$: targets fixed by the fresh runs. Let $\mathcal{F}_{12}$ be the union of the following six events, in each of which the named steps are paid hits with the stated targets. There are at most $5L$ choices of $a > L$, at most $5L^2$ pairs (a,c) with $2 \leq a \leq L$ and $L+2 \leq c \leq 6L$, and at most L choices of a when $c = L+2$, and each bound is Lemma 50 with $h = 2$.

(i) $a > L$; $a-1$ on a, and $a-L$ on a self-target: $2 \cdot 5L \cdot 6L/M = 60L^2/M$.
(ii) $a-1$ on a step in $\{a, c\}$, and $c-1$ on P: $2 \cdot 5L^2 \cdot 2\mu/M = 340L^2/M$.

(iii) $a-1$ on P, and $c-1$ on c: $2 \cdot 5L^2\mu/M = 170L^2/M$.
(iv) $a-1$ on a step in $\{a,c\}$, and $c-1$ on c: $2 \cdot 5L^2 \cdot 2/M = 20L^2/M$.
(v) $c = L+2$; $a-1$ on P, and $c-1$ on a self-target: $2 \cdot L\mu \cdot 6L/M = 204L^2/M$.
(vi) $c = L+2$; $a-1$ on a step in $\{a,c\}$, and $c-1$ on a self-target: $2 \cdot L \cdot 2 \cdot 6L/M = 24L^2/M$.

In total, $\Pr[\mathcal{F}_{12}] \le (60+340+170+20+204+24)L^2/M = 818L^2/M$. A self-target of $c-1$ in $\{a,c\}$ is c, since it lies above $c-1$.

$\mathcal{F}_{13}$: three paid hits, or two with a third payment. Let $\mathcal{F}_{13}$ be the union of the event that at least three steps are paid hits and of the four template events T_X, T_C, T_{YP} and T_{YS} of Section I.6. Lemmas 55, 57, 58 and 59 give

$$\Pr[\mathcal{F}_{13}] \le \frac{5 \cdot 1317006\, L^6}{B^3} = \frac{6585030\, L^6}{B^3} \le \frac{7 \cdot 10^6 L^6}{B^3}.$$

I.6. The shape of the fresh steps and the templates

Suppose that step 1 is old, $F \neq \emptyset$, and there is no double hit. Each maximal interval $[g,h]$ of consecutive fresh steps has $g \ge 2$, and $g-1$ is a hit with the fresh parent g, hence a paid hit; distinct intervals give distinct paid hits. If $f \in F$ and $f > L$, then $f - L$ is fresh or a paid hit with the fresh parent f.

Lemma 54 (Shape of F)**.** *Suppose that step 1 is old, $F \neq \emptyset$, and there is no double hit.*

(a) *If exactly one step is a paid hit, then F is an interval $[a,h] \subseteq [2,L]$, and every step above L is a P-hit.*
(b) *If exactly two steps are paid hits, then F is one of the following: (i) $\{a\}$ with $a > L$; (ii) $[a, L+1]$ with $3 \le a \le L$; (iii) the union of two intervals in $[2,L]$; (iv) $[a,h_1] \cup [c,h_2]$ with $h_1 \le L$, $c \ge L+2$ and $[c-L, h_2-L] \subseteq [a,h_1]$, so that $h_2 \le h_1 + L \le 2L$. In case (iv) the paid hits are $a-1$ and $c-1$, every step above h_2 is a P-hit, and every step in $[h_1+1, c-2]$ is a hit. If moreover step 2 is a hit, then $a \ge 3$ and $c \ge h_1 + 3$.*

Proof. (a) There is one interval $[a,h]$. If $a > L$, steps $a-1$ and $a-L$ are hits with the fresh parent a, hence two paid hits, which is excluded. If $a \le L < h$, step 1 is a hit with the fresh parent $L+1$, hence a paid hit, so $1 = a-1$; but then step 1 has both parents fresh, which is excluded. So $F \subseteq [2,L]$, and the steps above L, being hits with no fresh step above them, are P-hits.

(b) Let F be one interval $[a,h]$. If $a > L$, the paid hits are $a-1$ and $a-L$. If $a+1 \in F$, then $a+1-L \le a-1$ is a hit with the fresh parent $a+1$; it is a third paid hit unless $a+1-L = a-1$, and then $a-1$ has both parents fresh. So $F = \{a\}$. If $a \le L < h$, step 1 is a hit with the fresh parent $L+1$; for $a = 2$ it would be a double hit, so $a \ge 3$, step 1 is a paid hit, and the paid hits are $a-1$ and 1. If $L+2 \in F$, step 2 is a hit with the fresh parent $L+2$, a third paid hit unless $a = 3$, and then it is a double hit. So $h = L+1$. One interval inside $[2,L]$ gives only one paid hit, since a hit with a fresh parent $f \le L$ is $f-1$.

So let F have two maximal intervals $I_1 = [a,h_1]$ and $I_2 = [c,h_2]$, $h_1 < c$; the paid hits are $a-1$ and $c-1$. If both intervals lie in $[2,L]$, this is (iii). Otherwise, for every fresh $f > L$ the step $f-L$ is fresh: if not, it is a paid hit with the fresh parent f, hence $a-1$ or $c-1$, which already has the fresh parent a or c, and it would be a double hit. If $a > L$, then $f = a$ gives $a - L \in F$, below $\min F$; so $a \le L$. If $h_1 > L$, then $L+1 \in I_1$, and $f = L+1$ gives $1 \in F$; so $h_1 \le L$, and some fresh step above L lies in I_2. If $c \le L+1$, then $L+1 \in I_2$ and again $1 \in F$; so $c \ge L+2$. The steps $f-L$, $f \in I_2$, form the interval $[c-L, h_2-L]$ of fresh steps, which

lies in one maximal interval, not in I_2 since $c - L < c$; so it lies in I_1. The steps above h_2 are hits with no fresh step above them, hence P-hits, and the steps in $[h_1 + 1, c - 2]$ are not fresh. If step 2 is a hit, then $a \geq 3$. Finally $c \leq h_1 + 2$ would force $c = h_1 + 2$, hence $h_1 = L$ and $c = L + 2$, and then $c - L = 2 \in I_1$ gives $a \leq 2$. □

Lemma 55 (Three paid hits)**.** *The probability that at least three steps are paid hits is at most* $2(6L)^3Q^3/B^3 \leq 1317006\, L^6/B^3$.

Proof. Each paid hit has a private half that enters a request for the first time at that hit, and distinct paid hits have distinct private halves (Lemma 48). Lemma 50 with $h = 3$ and $k_i = Q$, over at most $(6L)^3$ triples $u_1 > u_2 > u_3$, gives $2(6L)^3Q^3/B^3$. With $Q \leq \frac{29}{2}L$, $2 \cdot 216 \cdot (\frac{29}{2})^3 = 1317006$. □

With exactly two paid hits, the remaining configurations have the shape (iv). We write (S) for the conditions

> $F = [a, h_1] \cup [c, h_2]$ with $3 \leq a \leq h_1 \leq L$ and $h_1 + 3 \leq c \leq h_2 \leq h_1 + L$, every step above h_2 is a P-hit, and $a - 1$ and $c - 1$ are paid hits.

Under (S), the requests of the steps above h_2 are computed from C and P alone; we call their values the values of the top walk. Since $c - 2 + L \geq h_1 + L + 1 > h_2$ and $c \leq 2L$, the right halves $z_{c-1} = \beta_{c-1+L}$ and $z_{c-2} = \beta_{c-2+L}$ are values of the top walk. Steps $c - 2$ and $a - 2$ are hits, the first because it lies in $[h_1 + 1, c - 2]$, the second because $a = \min F$ and $a \geq 3$. The four template events are:

> T_X: (S), and $c - 1$ is a self-hit on a step t_2 with $c < t_2 \leq h_2$;
> T_C: (S), $h_2 > c$, and $c - 1$ is a self-hit on c;
> T_{YP}: (S), $c - 1$ is a P-hit, and $a - 1$ is a self-hit on a step $t_1 \notin \{a, c\}$;
> T_{YS}: (S), $h_2 = c$, $c - 1$ is a self-hit on c, and $a - 1$ is a self-hit on a step $t_1 \notin \{a, c\}$.

In each template the two paid hits fix the left halves $\theta_1 = \alpha_a$ and $\theta_2 = \alpha_c$, and a third left half θ_3, copied by a self-hit, is fixed by its first true comparison. The next lemma bounds that comparison.

Lemma 56 (A copied left half)**.** *In the array run, let θ be the left half of entry e, and fix all other halves of the array. Declare false every comparison one of whose sides involves θ; the resulting* default run *has comparison outcomes that do not depend on θ. For a step j of the default run, let $\mathcal{A}_j$ be the set of values of θ for which some comparison of step j involving θ is true. Then $|\mathcal{A}_j| \leq Q$, and the array run agrees with the default run up to the first step at which a comparison involving θ is true; if that step is j, then $\theta \in \mathcal{A}_j$.*

Proof. In the default run every left input half is θ itself or free of θ: a self-hit copies the left half of a fresh step's entry, a P-hit takes that of a point of P, and a fresh step other than e takes its own entry. Every right half is free of θ, since θ is a left half. So every request is free of θ or equals $(\theta \oplus \langle j \rangle, z)$ with z free of θ, where j is its step. At step j there are four kinds of comparison involving θ. (1) The request $(\theta \oplus \langle j \rangle, z)$ against an output (γ, δ) of P is true only if $\delta = z$ and $\theta = \gamma \oplus \langle j \rangle$; by the cap this gives at most μ values. (2) The request $(\theta \oplus \langle j \rangle, z)$ against the request of a higher fresh step that is free of θ gives one value for each such step. (3) The request $(\theta \oplus \langle j \rangle, z)$ against a request $(\theta \oplus \langle f \rangle, z')$ of a higher fresh step f requires $\langle j \rangle = \langle f \rangle$ and is never true. (4) A request free of θ against a request $(\theta \oplus \langle f \rangle, z')$ gives one value for each f. So $|\mathcal{A}_j| \leq \mu + 6L = Q$. Until a comparison involving θ is true, the array run makes the same comparisons with the same outcomes as the default run. □

Lemma 57 (Upper self-target above c)**.** *In the array run,* $\Pr[T_X] \leq (6L)^3Q^3/B^3$*, so in the uniform completion* $\Pr[T_X] \leq 1317006\, L^6/B^3$.

Proof. Fix a, c and t_2, and put $\theta_1 = \alpha_a$, $\theta_2 = \alpha_c$ and $\theta_3 = \alpha_{t_2}$, three different entries since $a < c < t_2$. Fix all other halves of the array. We show that at most Q^2 triples $(\theta_1, \theta_2, \theta_3)$ give T_X with these a, c, t_2.

The half θ_2. Step c is fresh, so the left half of Y_{c-1} is $\theta_2 \oplus \langle c-1 \rangle$, and the self-hit on t_2 makes it equal to $\alpha_{t_2+1} \oplus \langle t_2 \rangle$, the left half of Y_{t_2}. Here α_{t_2+1} is an entry if $t_2 < h_2$ and a value of the top walk if $t_2 = h_2$. The request Y_{t_2} is formed before θ_3 and θ_1 enter any request (Lemma 48) and does not involve θ_2. So θ_2 is determined.

The half θ_3. Apply Lemma 56 to θ_3, with θ_2 fixed; down to step $a-1$ the default run does not involve θ_1. On T_X, θ_3 enters the request of the fresh step $t_2 - 1 \geq c$, and the comparisons of the fresh steps $t_2 - 1, \ldots, c$ are all false. Step $c-1$ hits t_2, and its comparison with Y_{t_2-1} is false, since $Y_{t_2} \neq Y_{t_2-1}$ when $t_2 - 1$ is fresh. Step $c-2$ is a hit, and its request $(\theta_3 \oplus \langle c-2 \rangle, z_{c-2})$ involves θ_3, because $c-1$ reads the input of t_2. So the first true comparison involving θ_3 is at step $c-2$, and $\theta_3 \in \mathcal{A}_{c-2}$, which has at most Q values. The excluded kind (3) is the comparison with the request of $t_2 - 1$, which would need $\langle c-2 \rangle = \langle t_2 - 1 \rangle$.

The half θ_1. Now the run down to step $a-1$ is determined. Step $a-1$ is a paid hit with private half θ_1 and request $(\theta_1 \oplus \langle a-1 \rangle, z_{a-1})$, where z_{a-1} is determined. Its target is one of at most μ points of P with right half z_{a-1} or one of at most $6L$ higher fresh steps, so at most Q values of θ_1 remain.

With at most $(6L)^3$ choices of (a, c, t_2), $\Pr[T_X] \leq (6L)^3 Q^2/B^3 \leq (6L)^3 Q^3/B^3$, and Lemma 49 and $Q \leq \frac{29}{2} L$ give the second bound. □

Lemma 58 (Upper self-target at c with more fresh steps above)**.** *In the array run,* $\Pr[T_C] \leq (6L)^3 Q^3/B^3$*, so in the uniform completion* $\Pr[T_C] \leq 1317006\, L^6/B^3$.

Proof. Fix a and c, and put $\theta_1 = \alpha_a$, $\theta_2 = \alpha_c$ and $\theta_3 = \alpha_{c+1}$; step $c+1 \leq h_2$ is fresh, and the three entries are different. Fix all other halves. By Lemma 48, the run above step c involves none of $\theta_1, \theta_2, \theta_3$. Step c is fresh with request $(\theta_3 \oplus \langle c \rangle, z_c)$, and the self-hit of $c-1$ on c gives

$$\theta_2 \oplus \langle c-1 \rangle = \theta_3 \oplus \langle c \rangle. \tag{E}$$

Step $c-2$ is a hit, and since $c-1$ reads the input of c, its request is $(\theta_2 \oplus \langle c-2 \rangle, z_{c-2})$, where z_{c-2} is a value of the top walk. Its target is a point of P with right half z_{c-2}, at most μ of them, or a fresh step above $c-2$. The target c is impossible: $\theta_2 \oplus \langle c-2 \rangle = \theta_3 \oplus \langle c \rangle$ together with (E) would give $\langle c-2 \rangle = \langle c-1 \rangle$. A target $f \in [c+1, h_2]$ has the left half $\alpha_{f+1} \oplus \langle f \rangle$ of Y_f, where α_{f+1} is an entry above $c+1$ or a value of the top walk. So the target is one of at most Q candidates determined by the run above c, each candidate determines θ_2, and (E) then determines θ_3. Finally θ_1 takes at most Q values, as in the proof of Lemma 57. These are two independent equations for (θ_2, θ_3), solved directly, and a third for θ_1. So at most Q^2 triples remain for each (a, c), and $\Pr[T_C] \leq (6L)^2 Q^2/B^3 \leq (6L)^3 Q^3/B^3$. □

Lemma 59 (Lower self-target)**.** *In the array run,* $\Pr[T_{YP}] \leq (6L)^3 Q^3/B^3$ *and* $\Pr[T_{YS}] \leq (6L)^3 Q^3/B^3$*, so in the uniform completion each is at most* $1317006\, L^6/B^3$.

Proof. Fix a, c, t_1 and, for T_{YP}, the point w of P hit by $c-1$; we call this the outer choice. Put $\theta_1 = \alpha_a$, $\theta_2 = \alpha_c$ and $\theta_3 = \alpha_{t_1}$, three different entries since $t_1 \notin \{a, c\}$. Fix all other halves. We show that for each outer choice at most $6LQ$ triples give the event.

Since z_{c-1} is a value of the top walk, in T_{YP} there are at most μ candidates w with $\delta(w) = z_{c-1}$.

The half θ_2. Step c is fresh, so the left half of Y_{c-1} is $\theta_2 \oplus \langle c-1 \rangle$. In T_{YP} it equals $\gamma(w)$. In T_{YS} it equals the left half $\alpha_{c+1} \oplus \langle c \rangle$ of Y_c, where α_{c+1} is a value of the top walk since $c = h_2$. In both cases θ_2 is determined by the outer choice.

The half θ_1. The self-hit of $a-1$ on t_1 gives $\theta_1 \oplus \langle a-1\rangle = \alpha_{t_1+1} \oplus \langle t_1\rangle$, the left half of Y_{t_1}. The run above t_1 involves neither θ_3, which first enters a request at $t_1 - 1$, nor θ_1, which first enters one at $a - 1 < t_1$; it may involve θ_2 when $c > t_1$, and θ_2 is already determined. So θ_1 is determined.

The half θ_3. Now every half except θ_3 is fixed, and we apply Lemma 56. On the event, step $a-1$ reads the input of t_1, so its left input half is θ_3, and step $a-2 \geq 1$ is a hit whose request $(\theta_3 \oplus \langle a-2\rangle, z_{a-2})$ involves θ_3. So some comparison involving θ_3 is true, at step $a-2$ or above, and if the first one is at step j, then $\theta_3 \in \mathcal{A}_j$. Hence θ_3 lies in a union of at most $6L$ sets of at most Q values each.

There are at most $(6L)^2$ pairs (a, c), at most $6L \leq Q$ choices of t_1, and at most $\mu \leq Q$ choices of w (one in T_{YS}). So the probability is at most $(6L)^2 Q^2 \cdot 6LQ/B^3 = (6L)^3 Q^3/B^3$. □

The three equations of Lemma 59 are triangular: the equation for θ_2 involves neither θ_1 nor θ_3, the equation for θ_1 may involve θ_2 but not θ_3, and θ_3 enters its equation with coefficient one while every other half in it is already fixed. So the third equation is not implied by the first two, and no equation is paid for twice; by kind (3) of Lemma 56, two comparisons on the same copied value would need two different counters to agree. The same holds for the equations used in Lemmas 57 and 58.

I.7. Proof of Lemma 21

Lemma 60 (Cover)**.** *If the unwrap is valid and step 1 is old, then one of $\mathcal{F}_1$, $\mathcal{F}_2$, $\mathcal{F}_3$, $\mathcal{F}_4$, $\mathcal{F}_7$, $\mathcal{F}_9$, $\mathcal{F}_{10}$, $\mathcal{F}_{11}$, $\mathcal{F}_{12}$, $\mathcal{F}_{13}$ occurs.*

Proof. If $F = \emptyset$, $\mathcal{F}_1$ occurs, and a double hit is $\mathcal{F}_2$. So let $F \neq \emptyset$ with no double hit, and let $a = \min F \geq 2$. Step $a-1$ is a hit with the fresh parent a, hence a paid hit.

If step 2 or step $L+1$ is fresh, step 1 is a hit with a fresh parent, hence a paid hit. If its target is in P, this is $\mathcal{F}_3$. Otherwise it is a self-hit on a fresh d, and validity makes $\alpha_1 = \alpha_d$ legal: $\mathcal{F}_7$. From now on steps 2 and $L+1$ are hits, so $a \geq 3$.

If no step above L is fresh, then $F \subseteq [2, L]$, and step $a-1$ is a P-hit ($\mathcal{F}_9$) or a self-hit ($\mathcal{F}_{10}$).

Otherwise let c be the least fresh step above L. If $a > L$, then $c = a$, and steps $a-1$ and $a-L$ are hits with the fresh parent a, hence paid hits. If $a \leq L$, then $c \geq L+2$ since $L+1$ is a hit, and $c-1 \geq L+1$ is a hit with the fresh parent c, hence a paid hit. These are the two selected paid hits. If both have targets in P, this is $\mathcal{F}_4$. For $a > L$, one self-target is $\mathcal{F}_{11}$, and two self-targets with $t_1 = a$ are $\mathcal{F}_{12}$(i). For $a \leq L$, the configurations of $\mathcal{F}_{12}$(ii)–(vi) remain to be excluded, and with them the following five remain:

(r1) $a > L$, both targets are self-targets, and $t_1 \neq a$;
(r2) $a \leq L$, $a-1$ on $t_1 \notin \{a, c\}$, and $c-1$ on P;
(r3) $a \leq L$, $a-1$ on P, $c-1$ on $t_2 \neq c$, and $c \geq L+3$;
(r4) $a \leq L$, $a-1$ on $t_1 \notin \{a, c\}$, and $c-1$ on a self-target;
(r5) $a \leq L$, $a-1$ on $t_1 \in \{a, c\}$, $c-1$ on $t_2 \neq c$, and $c \geq L+3$.

Indeed, for $a \leq L$ with a self-target: a lower self-target with an upper target in P is (ii) or (r2); a lower target in P with $t_2 = c$ is (iii), with $t_2 \neq c$ and $c = L+2$ it is (v), and otherwise (r3); two self-targets with $t_1 \notin \{a, c\}$ give (r4), and with $t_1 \in \{a, c\}$ they give (iv) if $t_2 = c$, (vi) if $t_2 \neq c$ and $c = L+2$, and otherwise (r5).

In (r1)–(r5), three or more paid hits give $\mathcal{F}_{13}$. With exactly two, Lemma 54(b) applies. Since a fresh step lies above L and $L+1$ is a hit, cases (ii) and (iii) are excluded. In case (i), $F = \{a\}$

and every self-target is a, so (r1) is impossible. In case (iv), since step 2 is a hit, (S) holds with the a and c above, the upper paid hit is $c-1$, and $t_2 \in [c, h_2]$. Then (r2) is T_{YP}; in (r3) and (r5), $t_2 > c$ and this is T_X; in (r4), $t_2 > c$ gives T_X, and $t_2 = c$ gives T_C if $h_2 > c$ and T_{YS} if $h_2 = c$. All of these lie in $\mathcal{F}_{13}$. □

The proof does not use $\mathcal{F}_5$, $\mathcal{F}_6$ and $\mathcal{F}_8$; their bounds are rows of Table 7 and are included in the sum below.

The sum. By Lemma 60, $\Pr[\text{valid} \wedge \text{step 1 old}]$ is at most the sum of the bounds of Section I.5. The terms in $1/B$ add up to $(18{+}36{+}9)/B = 63/B \le 64/B$, and those in LD/M to $(2{+}12)LD/M = 14LD/M$. The terms in L^2/M are

$$\begin{aligned} &72 + 10404 + 17 + 2040 + 96 \\ &\quad + 19845 + 10 + 2040 + 818 = 35342, \end{aligned}$$

and $\mathcal{F}_9$ adds $331L \le 165.5L^2$, so the coefficient of L^2/M is below $35508 \le 35868$. The terms in p/M add up to $(73728{+}18000)p/M = 91728\,p/M$, and $\mathcal{F}_{13}$ gives $6585030L^6/B^3 \le 7{\cdot}10^6 L^6/B^3$. Hence

$$\begin{aligned} \Pr[\text{valid} \wedge \text{step 1 old}] \le \frac{64}{B} + \frac{14LD}{M} + \frac{35868\,L^2}{M} \\ + \frac{91728\,p}{M} + \frac{7 \cdot 10^6 L^6}{B^3}. \end{aligned}$$

Together with Section I.4 this proves Lemma 21. When the adversary chooses L from the transcript, the bounds hold with L inside the expectation; the weighted sums of Lemma 53 are what allow this for $\mathcal{F}_9$ and $\mathcal{F}_{10}$. □

I.8. Adding up

By Lemma 18 (Section H.2), the DAE advantage is at most the sum of the rows of Table 4, each averaged over ideal transcripts. The rows are bounded in Lemma 33 (collisions, $19\sigma^2/M$), Lemma 34 (the cap, $0.032\sigma^2/M$), Lemmas 20 and 19 with Section H.13 (the three rows on freshness), Lemma 35 (two-semiblock unwraps), and Lemma 21 (one row for each long unwrap k, with pass length L_k). Ciphertexts of other lengths are rejected without a call of π, so the number of long unwraps is $q_L \le q_v - r$. Write $s = \sum_k L_k$ over the long unwraps. The semiblock budget gives

$$\ell + 2p + s + 2r \le \sigma,$$

and $q_L \le s/2$, since every $L_k \ge 2$.

Rank two. Since $\ell^2 = \sum_T L_T^2 + 2\sum_{R<T} L_R L_T$ and $170 \le 860$,

$$170 \sum_T L_T^2 + 1720 \sum_{R<T} L_R L_T \le 860\,\ell^2.$$

Long unwraps. Summing the per-unwrap bound over k,

$$\begin{aligned} \textstyle\sum_k 14 L_k D &= 14s(6\ell + p) = 84\,\ell s + 14\,ps, \\ \textstyle\sum_k 35868 L_k^2 &\le 35868\, L_{\max} s, \\ \textstyle\sum_k 91728\,p &= 91728\, q_L p \le 45864\, ps, \\ \textstyle\sum_k 7 \cdot 10^6 L_k^6 &\le 7 \cdot 10^6\, L_{\max}^5 s. \end{aligned}$$

The coefficient of σ^2/M. The rows on messages at risk and on rank two contribute $17600\,p\ell$ and $18\,p\ell$. Apart from the term $35868 L_{\max} s$, the terms of order σ^2/M outside the first two rows are f/M with

$$f = 860\ell^2 + 17618\,p\ell + 84\,\ell s + 14\,ps + 45864\,ps + 7r\sigma.$$

Put $R = \sigma - \ell$, so that $2p + s \le R$. Then

$$17618\,p\ell + 84\,\ell s = \ell\,(8809 \cdot 2p + 84s) \le 8809\,\ell R,$$

and, since $(2p)s \le (2p+s)^2/4$,

$$45878\,ps = 22939\,(2p)s \le \frac{22939}{4} R^2 = 5734.75\,R^2.$$

As $860 \le 5734.75$ and $8809 \le 2 \cdot 5734.75 = 11469.5$,

$$\begin{aligned} f - 7r\sigma &\le 860\ell^2 + 8809\,\ell R + 5734.75\,R^2 \\ &\le 5734.75\,(\ell + R)^2 = 5734.75\,\sigma^2. \end{aligned}$$

Also $7r\sigma \le 3.5\sigma^2$, since $2r \le \sigma$. With $19\sigma^2/M$ and $0.032\sigma^2/M$ from the first two rows, the coefficient of σ^2/M is at most

$$19.032 + 5734.75 + 3.5 = 5757.282 \le 5758.$$

The term $35868 L_{\max} s$ is at most $35868 L_{\max}\sigma$.

The terms in $1/B$. The two-semiblock unwraps give $(1 + 2^{-7})r/B$, and each long unwrap gives $9/B$ with step 1 fresh and $64/B$ with step 1 old. So these terms are at most

$$\frac{(1 + 2^{-7})\,r + 73\,q_L}{B} \le \frac{73\,q_v}{B}.$$

Higher order. Freshness gives $7721568\,q\sigma L_{\max}^5/B^3$, and the long unwraps give at most $7 \cdot 10^6 L_{\max}^5 s/B^3$. If $s > 0$, then $q_v \ge 1$ and $s \le \sigma$, so $7 \cdot 10^6 L_{\max}^5 s \le 7721568\,q_v \sigma L_{\max}^5$, and the two terms together are at most $7721568\,(q + q_v)\,\sigma L_{\max}^5/B^3$.

Theorem 3. Adding the four parts, for every adversary with $\sigma \le 2^{60}$,

$$\begin{aligned} \mathrm{Adv}^{\mathrm{dae}}_{\mathrm{KWP}} \le \frac{5758\,\sigma^2}{2^{128}} + \frac{35868\,L_{\max}\,\sigma}{2^{128}} + \frac{73\,q_v}{2^{64}} \\ + \frac{7721568\,(q + q_v)\,\sigma\,L_{\max}^5}{2^{192}}. \end{aligned}$$

All resources are deterministic caps, so the bound holds for adaptive adversaries. □