AICR · OPEN PROBLEM

Security of the Merkle-Damgård Combiner Under Constant-Length Salts

Closing the gap below the source's own $|Z_i| > |M| + \lambda$ requirement

Status. Statement: AI-written from Yevgeniy Dodis, Eli Goldin and Peter Hall, *Random Oracle Combiners: Merkle-Damgård Style*, IACR ePrint 2025/609. Not yet checked by a human. Open in both directions: security is proved only for salts longer than the message by at least the security parameter, and the source has neither an attack nor a security proof once the salts shrink to a constant number of blocks.

Category. *Idealized Models & Non-Uniformity.*

Conjecture (informal). *The source's main construction XORs the Merkle-Damgård extensions of two salted hash functions together, and proves it is a secure random oracle combiner as long as each salt is longer than the message being hashed by at least the security parameter. The proof genuinely needs that margin to work, but the authors have not found any attack that actually exploits shorter salts; in particular they have no attack even when the salt length is a small constant number of blocks, independent of how long the message is. They explicitly leave open which of the two is really true: that short salts really do break the construction, or that the construction is secure with much shorter salts than its current proof requires.*

1 The setting

The source's main construction is $\widetilde{C}^{h_1,h_2}_{Z_1,Z_2}(M) := h_1^*(M,Z_1) \oplus h_2^*(M,Z_2)$, where $h_1, h_2 : \{0,1\}^{n+\delta} \to \{0,1\}^n$ are compression functions and h_1^*, h_2^* are their Merkle-Damgård extensions. The source's Theorem 3.1 proves $\widetilde{C}$ is a secure random oracle combiner (Definition 1) whenever $|Z_1|, |Z_2| > |M| + \lambda$; the proof relies on this margin via a compression argument bounding the simulator's recursion. No attack against $\widetilde{C}$ is known for shorter salts, including salts of constant length, independent of $|M|$.

The source's own conclusion states the position exactly: "While our analysis of $\widetilde{C}$ critically uses the fact that $|Z_i| > |M| + \lambda$, we do not have any attacks against this construction, once the $|Z_i|$ is a constant number of blocks, independent of $|M|$. In particular, it is possible our construction $\widetilde{C}$ provides a solution to" the source's broader, separately-recorded open problem about the existence of any combiner with $O(\lambda)$-length, message-independent salts. It continues: "Thus, it would be very interesting to either show that $\widetilde{C}$ is insecure unless $|Z_i| > |M|$, or find a supporting proof of security for much shorter salts (ideally, constant number of blocks)."

This is a narrower and more concrete question than that broader one: it is about this one fixed, already-specified construction, not about the existence of some combiner in general, so settling it requires only exhibiting an attack or a security proof against $\widetilde{C}$, not a new construction. A proof of security for $\widetilde{C}$ at constant-length salts would answer the broader existence question as well; an attack on $\widetilde{C}$ at constant-length salts would not, since some other combiner might still achieve short, message-independent salts.

2 Notation and parameters

- λ is the security parameter.
- $M \in \{0,1\}^*$ is the combiner's message input.
- $h_1, h_2 : \{0,1\}^{n+\delta} \to \{0,1\}^n$ are compression functions, with h_1^*, h_2^* their Merkle-Damgård extensions.
- Z_1, Z_2 are independent random salts.

3 The conjecture

Definition 1 (Random oracle combiner, informal; source's framing of Dodis-Ferguson-Goldin-Hall-Pietrzak). A combiner $C_Z^{h_1,h_2} : \{0,1\}^m \to \{0,1\}^n$ is a *secure random oracle combiner* if, for every oracle circuit g, both $C_Z^{\mathcal{H},g^{\mathcal{H}}}$ and $C_Z^{g^{\mathcal{H}},\mathcal{H}}$ are indifferentiable from a fresh random oracle $\mathcal{G} : \{0,1\}^m \to \{0,1\}^n$, over the random choice of the salt Z, against every polynomial-query distinguisher.

Conjecture 1 (Security of $\widetilde{C}$ under short salts). *Let $\widetilde{C}^{h_1,h_2}_{Z_1,Z_2}(M) := h_1^*(M,Z_1) \oplus h_2^*(M,Z_2)$ be as above, proved a secure random oracle combiner (Definition 1) whenever $|Z_1|,|Z_2| > |M| + \lambda$. Resolve which of the following holds:*

(i) there is a choice of $|Z_1|,|Z_2| = O(1)$ blocks, independent of $|M|$, together with an efficient adversary breaking the indifferentiability of $\widetilde{C}$ from a random oracle; or

(ii) $\widetilde{C}$ remains a secure random oracle combiner even when $|Z_1|,|Z_2|$ are a constant number of blocks, independent of $|M|$.

Remark 1 (The source's dichotomy is not perfectly symmetric). The source phrases the "insecure" alternative more broadly, as showing $\widetilde{C}$ is insecure unless $|Z_i| > |M|$ — i.e., over the whole sub-threshold range $|Z_i| \le |M|$ — while it narrows the "secure" alternative specifically to the constant-blocks case, the regime where it reports no known attack. Resolving Conjecture 1 at the constant-blocks extreme, in either direction, need not by itself settle the source's broader-phrased alternative for every $|Z_i|$ strictly between a constant and $|M|$.

4 Bibliography

[DGH25] Yevgeniy Dodis, Eli Goldin and Peter Hall. *Random oracle combiners: Merkle-Damgård style.* IACR ePrint 2025/609. The source. The construction and its security theorem are Section 3 (Theorem 3.1); the open problem is Section 4 ("Conclusion and Open Problems"), page 27.