CONJURA · OPEN PROBLEM

Salt Length vs. Merkle-Damgård Call Count in Random Oracle Combiners

Existence, or impossibility and a trade-off, for message-independent salts and calls

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. Settled only for salts strictly longer than the message; whether $O(\lambda)$-length salts and $O(1)$ Merkle-Damgård calls can coexist, or are provably incompatible, is fully open.

Category. *Idealized Models & Non-Uniformity.*

Conjecture (informal). *A random oracle combiner takes two hash functions and builds a new one that is provably a fresh random oracle as long as at least one of the two inputs already is one. The source builds such a combiner out of Merkle-Damgård hash functions by salting and XORing their outputs, but its proof needs the random salts to be strictly longer than the message being hashed, and as a result the combiner must make roughly one Merkle-Damgård call per message block. The source asks whether this is inherent: is there any secure combiner of this kind whose salt length no longer needs to grow with the message, while still calling the underlying hash functions only a constant number of times? Or, if no such combiner can exist, what is the best trade-off achievable between how short the salt can be and how many times the underlying functions must be called?*

1 The setting

Dodis, Ferguson, Goldin, Hall and Pietrzak [DFGHP23] introduced *random oracle combiners* (ROCs): given two hash functions h_1, h_2, a ROC $C_Z^{h_1,h_2}$ is a new hash function guaranteed to be indifferentiable from a fresh random oracle as long as at least one of h_1, h_2 is itself a random oracle, while the other may depend on it arbitrarily. They built a length-preserving ROC $h_1(M, Z_1) \oplus h_2(M, Z_2)$ treating h_1, h_2 as monolithic random oracles, with salts Z_1, Z_2 strictly longer than

the message M, and left open whether the same holds when h_1, h_2 are instead the compression functions of a Merkle-Damgård hash.

The source resolves that question: it proves that $\widetilde{C}^{h_1,h_2}_{Z_1,Z_2}(M) := h^*_1(M, Z_1) \oplus h^*_2(M, Z_2)$, where h_1, h_2 are fixed-length compression functions and h^*_1, h^*_2 their Merkle-Damgård extensions, is a secure ROC whenever $|Z_1|, |Z_2| > |M| + \lambda$. This closes the Dodis-Ferguson-Goldin-Hall-Pietrzak conjecture for this construction, but the salts must still grow with $|M|$, and the resulting combiner makes on the order of $|M|$ Merkle-Damgård calls to h^*_1, h^*_2. The source's own conclusion records this as the main remaining inefficiency and states the question left open: whether some secure combiner of this shape can instead use only $O(\lambda)$ bits of salt while still calling the underlying compression functions a constant number of times, independent of the message length — and if not, what trade-off between salt length and call count is the true one.

2 Notation and parameters

- λ is the security parameter.
- n, δ are the output length and the per-block input length of the compression functions, i.e. $h_1, h_2 : \{0,1\}^{n+\delta} \to \{0,1\}^n$.
- $M \in \{0,1\}^*$ is the combiner's message input, of length $m = |M|$.
- h^*_1, h^*_2 are the Merkle-Damgård extensions of h_1, h_2 to arbitrary-length inputs.
- Z_1, Z_2 are independent random salts appended to M before hashing.

3 The conjecture

Definition 1 (Random oracle combiner, informal; source's framing of Dodis-Ferguson-Goldin-Hall-Pietrzak [DFGHP23]). Let $C^{h_1,h_2}_Z : \{0,1\}^m \to \{0,1\}^n$ take a message M, a salt $Z = (Z_1, Z_2)$, and oracle access to h_1, h_2. Say C is a *secure random oracle combiner* if, for every oracle circuit g, both $C^{\mathcal{H}, g^{\mathcal{H}}}_Z$ and $C^{g^{\mathcal{H}}, \mathcal{H}}_Z$ are indifferentiable, in the sense of Maurer, Renner and Holenstein, from a fresh random

oracle $\mathcal{G} : \{0,1\}^m \to \{0,1\}^n$, over the random choice of Z, against every polynomial-query distinguisher.

Definition 2 (Message-independent salts and calls). A secure random oracle combiner $C^{h_1,h_2}_{Z_1,Z_2} : \{0,1\}^m \to \{0,1\}^n$ (Definition 1) has *message-independent salts and calls* if

(a) $|Z_1|, |Z_2| = O(\lambda)$, independent of $m = |M|$; and

(b) C makes only $O(1)$ calls to h_1^* and h_2^*, independent of m.

Conjecture 1 (Salt length vs. call count). *Either exhibit a secure random oracle combiner with message-independent salts and calls (Definition 2), or prove that no secure random oracle combiner can satisfy both (a) and (b) simultaneously and, in that case, determine the optimal trade-off between the salt length $|Z_1|, |Z_2|$ and the number of calls made to h_1^*, h_2^*.*

Remark 1 (What is already known). The source's own construction $\widetilde{C}$ shows that $|Z_1|, |Z_2| > |M| + \lambda$ suffices for security, at the cost of making the number of Merkle-Damgård calls grow with $|M|$; no construction achieving both (a) and (b) of Definition 2 is known, and no impossibility result is given either. A closely related, narrower question — whether $\widetilde{C}$ itself remains secure, or is instead attackable, once its own salts shrink to a constant number of blocks — is recorded separately on this site; a security proof for $\widetilde{C}$ in that regime would witness this conjecture's first alternative, but resolving this conjecture does not require working with $\widetilde{C}$ specifically.

4 Bibliography

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

[DFGHP23] Yevgeniy Dodis, Niels Ferguson, Eli Goldin, Peter Hall and Krzysztof Pietrzak. *Random oracle combiners: Breaking the concatenation barrier for collision-resistance.* CRYPTO 2023, Part II. Introduces random oracle combiners and the length-preserving construction whose Merkle-

Damgård-friendly analogue the source studies; cited by the source as the origin of the underlying question.