On the pseudorandomness of Shabal's keyed permutation

Jean-Philippe Aumasson

FHNW, Windisch, Switzerland

Abstract. We report observations suggesting that the permutation used in Shabal does not behave pseudorandomly. This does not affect the security of Shabal as submitted to the NIST Hash Competition.

Shabal's core function $\mathcal{P}$ maps a $(384 + 512 + 512)$-bit chain value (A, B, C) to another (A, B, C) chain value, letting C invariant and using a 512-bit message block M. The set $\mathcal{P} = \{\mathcal{P}_{M,C}(\cdot,\cdot)\}_{M,C}$ is a family of permutations of $\{0,1\}^{896}$, claimed in [4, §4.4.2] to be indistinguishable from a pseudorandom permutation (PRP).

The classical adversary makes (adaptive) queries to $\mathcal{P}_{M,C}$, for some random unknown M and C, and tries to distinguish $\mathcal{P}_{M,C}$ from a random permutation. The relaxed notion of weak PRP (wPRP) only considers "known plaintexts", i.e. random tuples $(A, B, \mathcal{P}_{M,C}(A, B))$, rather than "adaptive chosen plaintexts".

We present an algorithm that distinguishes $\mathcal{P}$ from a random permutation family, in a variant of the wPRP setting: the adversary gets tuples $(A, B, \mathcal{P}_{M,C}(A, B))$ for some *fixed unknown* A, B, and C (e.g., the IV of Shabal), and for some random M, such that the last 64 bits of M can be chosen for each tuple. Ideally, distinguishing $\mathcal{P}$ from a random permutation family in this setting would require about $2^{512-64} = 2^{448}$ computations of a $\mathcal{P}$ permutation.

Our algorithm exploits the fact that in the 3-round permutation of Shabal, the value of some variables after the second round don't depend on all the key bits. We use cube testers [1] to build a statistical distinguisher. In the terminology of [1], we use 7 cube variables in $M[14]$ and 5 superpoly variables in $M[15]$. The Boolean components tested correspond to the bits of $B[5]$ after the second round, which can be observed as follows:

- invert the finalization loop of $\mathcal{P}_{M,C}$
- guess $M[5], \ldots, M[13]$
- invert the last 11 loops of the third round, observe $B[5]$

Using a neutrality test, we observe that for the cube formed by variable bits $10, \ldots, 16$ of $M[14]$, the variable bits $25, \ldots, 30$ of $M[15]$ are (almost) neutral in the Boolean function corresponding to the bits $25, 26, 27$ of $B[5]$ (after round 2). The complexity of the cube tester is 2^{12} queries, which gives in total a cost of $2^{9\times 32+16} = 2^{300}$ (against 2^{448} ideally).

These results suggest that Shabal's keyed permutation is not pseudorandom, but don't affect the security of Shabal as a hash function (its iteration mode precludes any application of our algorithm). Note that other SHA-3 submissions, e.g. CubeHash [3], also rely on a nonpseudorandom permutation [2].

References

1. Jean-Philippe Aumasson, Itai Dinur, Willi Meier, and Adi Shamir. Cube testers and key recovery attacks on reduced-round MD6 and Trivium. In *FSE 2009*. to appear.
2. Jean-Philippe Aumasson, Willi Meier, María Naya-Plasencia, and Thomas Peyrin. Inside the hypercube. Cryptology ePrint Archive, Report 2008/486, 2008.
3. Daniel J. Bernstein. Cubehash specification (2.B.1). Submission to NIST, 2008.
4. Emmanuel Bresson, Anne Canteaut, Benoît Chevallier-Mames, Christophe Clavier, Thomas Fuhr, Aline Gouget, Thomas Icart, Jean-François Misarsky, María Naya-Plasencia, Pascal Paillier, Thomas Pornin, Jean-René Reinhard, Céline Thuillet, and Marion Videau. Shabal, a submission to NIST's cryptographic hash algorithm competition. Submission to NIST, 2008.