CubeHash attack analysis (2.B.5)

Daniel J. Bernstein *

Department of Computer Science
University of Illinois at Chicago
Chicago, IL 60607–7045
cubehash@box.cr.yp.to

This document is an analysis of CubeHash with respect to known attacks (e.g., differential cryptanalysis) and their results. CubeHash is a new algorithm; there are no previously published materials describing or analyzing the security of CubeHash.

CubeHashr/b becomes very easy to break if b is very large. For example, CubeHash$r/112$ allows an attacker to use the last 112-byte block of input to freely adjust 112 of the 128 state bytes. This last-block "message modification" effectively reduces the influence of the previous input blocks to a "narrow pipe" of only 16 bytes, so the attacker has a reasonable chance of finding a collision after cycling through just 2^{64} input messages. Note that this attack is independent of the state-transformation details, and in particular is independent of r.

However, the attacker loses control over the state as b decreases. Generic attacks have success probability dropping exponentially with the "pipe size" $1024 - 8b$. For example, CubeHash$r/32$ has a 768-bit pipe, putting generic attacks—including advanced attacks such as "herding"—far out of reach.

An attacker can try to do better with non-generic attacks. A single round of the CubeHash transformation is not terribly complicated. The extreme case CubeHash$1/b$ has a single-bit input change affecting only about ten bits of the state at the beginning of the next block. However, the hypercube structure of CubeHash distributes those bits widely, and if b is small then I don't see how an attacker can prevent the differential from spreading into the entire state over the next several blocks. I don't know whether differential probabilities can be rigorously *proven* to be small for sufficiently small b.

The $10r$ final rounds in CubeHash are an output filter, thoroughly mixing the state bits before the final truncation to the desired output size. As far as I can tell, 10 rounds are already overkill, providing full protection against linear attacks, differential attacks, etc. By flipping an additional state bit before these final rounds, CubeHash distinguishes the final block of the message from the previous blocks, apparently making slide attacks, length-extension attacks, etc. as difficult as standard differential attacks.

Protection against trap doors. CubeHash has a few constants that could be modified, but as far as I know there is no way that any design of this type could have a hidden vulnerability. See the CubeHash specification for discussion of the rotation distances, the hypercube structure, etc.

* The author was supported by the National Science Foundation under grant ITR–0716498. Date of this document: 2008.10.28.