From: Watanabe Dai [dai.watanabe.td@hitachi.com]
Sent: Monday, October 05, 2009 11:14 PM
To: hash-function@nist.gov
Cc: hash-forum@nist.gov
Subject: OFFICIAL COMMENT: Luffa (Round 2)

Dear All,

Great thanks to Stefan Tillich who pointed out an error in the specification document of Luffa v2.
The description of the finalization in Section 3.3 remains to be old and we fix the error for the consistency to the figure and other documents.
The updated package consisting of specification document and the errata is now available at http://www.sdl.hitachi.co.jp/crypto/luffa/Luffa_v2_update_October2009.zip

Regards,
Dai Watanabe

From: hash-forum@nist.gov on behalf of Watanabe Dai [dai.watanabe.td@hitachi.com]
Sent: Tuesday, August 10, 2010 3:42 AM
To: Multiple recipients of list
Subject: OFFICIAL COMMENT: Luffa (Round 2)

Dear All,

What interesting theoretical discussions we had for the last several days!

For the meanwhile, Luffa team uploaded the hardware implementation report on our web site.

http://www.sdl.hitachi.co.jp/crypto/luffa/ACompactHardwareImplementationOfSHA-3CandidateLuffa_20100810.pdf

Abstract.
In this document, the hardware performance of Luffa-256 are reported.
Our implementations mainly target size optimized implementations
in ASIC and the smallest architecture can be implemented with
only 10.3 KGE while it achieves about 500 Mbps.

Our implementations are fully autonomous.
At this moment, the above mentioned architecture is the second smallest in the known implementations of SHA-3 candidates.

Besides, it may be valueable to remark that two slides presented at ESC 2010 and FSE 2010 are available on-line, which deal with the Sbox design of Luffa and HOD distinguishing attack on Luffa-v1 respectively.
The FSE paper is the updated version of the document which was included in the second round package and there is no difference in their technical contents.

Regards,
Dai

From: hash-forum@nist.gov on behalf of Hirotaka Yoshida [hirotaka.yoshida.qv@hitachi.com]
Sent: Monday, November 08, 2010 9:14 PM
To: Multiple recipients of list
Subject: OFFICIAL COMMENT: Luffa (Round 2)

Dear All,

We uploaded a report on the security analysis of Luffa on our web site:
http://www.sdl.hitachi.co.jp/crypto/luffa/FindingCollisionsForReducedLuffa-256v2_20101108.pdf

Abstract.
This paper presents the first collision finding analysis
of Luffa-256 v2. We show that collisions for 4 out of 8 steps
of Luffa-256 v2 can be found with complexity 2^{90}
using sophisticated message modification techniques.
We also show that applying the same approach to find collisions to
Luffa-256 v2 reduced to 5 steps would require complexity 2^{224}.

Our results are on the hash function itself while the previous ones are on its building blocks (e.g. Zero-sum distinguishers) or attacks mounted under the assumptions which are considerably optimistic for the attacker (e.g. Rebound attacks).

Therefore our analysis can be seen as an indication that the full Luffa-256 v2 has a large security margin against attacks mounted under the realistic scenario.

Best regards,
Hirotaka Yoshida

From: hash-forum@nist.gov on behalf of Watanabe Dai [dai.watanabe.td@hitachi.com]
Sent: Tuesday, November 09, 2010 1:06 AM
To: Multiple recipients of list
Subject: OFFICIAL COMMENT: Luffa (Round 2)

Dear all,

Recently, we got some new implementation results of Luffa.
This is just their summary and guide to the references.

* PC platforms
Thanks to Thomaz Oliveira and Julio Lopez of University of Campinas, Luffa's throughputs on the latest Intel proccessors (45nm Intel Core micro architecutre and Nehalem architecutre) in 64-bit mode gain about 30% pervormance improvement.
(e.g. 9.5 cpb on Berlekamp in amd64 mode).

* Embedded Platforms
Recently, XBX team announced the performances of all SHA-3 candidates and we have nothing to add.
Luffa requires small RAM and ROM.

* FPGA
Our verilog codes for area optimized ASIC implementations are also synthesized by Xilinx ISE.
The following results are picked up from our updated report:

  Platform: Xilinx Virtex-5
  1. 548 slices, 1660   Mbps, 162 MHz
  2. 355 slices,   33.3 Mbps, 50 MHz

Note that they are not optimized for FPGA yet, and we did not use a block RAM nor a shift register.

Please see the following paper for the detail:
http://www.sdl.hitachi.co.jp/crypto/luffa/ACompactHardwareImplementationOfSHA-3CandidateLuffa_20100810.pdf

Regards,
Dai

From: hash-forum@nist.gov on behalf of Watanabe Dai [dai.watanabe.td@hitachi.com]
Sent: Tuesday, November 09, 2010 6:32 AM
To: Multiple recipients of list
Subject: Re: OFFICIAL COMMENT: Luffa (Round 2)

Dear all,

I am sorry for the wrong address for the reference.
The correct one is

http://www.sdl.hitachi.co.jp/crypto/luffa/ACompactHardwareImplementationOfSHA-3CandidateLuffa_20101105.pdf

And I appreciate Stefan Tillich who kindly pointed out my mistake.

Regards,
Dai

(2010/11/09 14:54), Watanabe Dai wrote:
> Dear all,
>
>
> Recently, we got some new implementation results of Luffa.
> This is just their summary and guide to the references.
>
> * PC platforms
> Thanks to Thomaz Oliveira and Julio Lopez of University of Campinas,
> Luffa's throughputs on the latest Intel proccessors (45nm Intel Core
> micro architecutre and Nehalem architecutre) in 64-bit mode gain about
> 30% pervormance improvement.
> (e.g. 9.5 cpb on Berlekamp in amd64 mode).
>
>
> * Embedded Platforms
> Recently, XBX team announced the performances of all SHA-3 candidates
> and we have nothing to add.
> Luffa requires small RAM and ROM.
>
>
> * FPGA
> Our verilog codes for area optimized ASIC implementations are also
> synthesized by Xilinx ISE.
> The following results are picked up from our updated report:
>
>   Platform: Xilinx Virtex-5
>   1. 548 slices, 1660   Mbps, 162 MHz
>   2. 355 slices,   33.3 Mbps, 50 MHz
>
> Note that they are not optimized for FPGA yet, and we did not use a
> block RAM nor a shift register.
>
> Please see the following paper for the detail:
> http://www.sdl.hitachi.co.jp/crypto/luffa/ACompactHardwareImplementati
> onOfSHA-3CandidateLuffa_20100810.pdf
>
>
> Regards,
> Dai

From: hash-forum@nist.gov on behalf of Watanabe Dai [dai.watanabe.td@hitachi.com]
Sent: Friday, December 03, 2010 3:52 AM
To: Multiple recipients of list
Subject: OFFICIAL COMMENT: Luffa (Round 2) implementation update

Dear all,

We ported the fastest x64 assembly codes to the x86 platform.
The code for Luffa-256 is based on UNICAMP code and the others are based on Hitachi codes.
The throughputs evaluated by SUPERCOP on my PC are as follows:

Luffa-256	9.76 cycles/B
Luffa-384	15.43
Luffa-512	20.51

Evaluation environment:
* CPU: Core2Duo E8400 3GHz
* OS: Ubuntu Linux 10.10 Desktop (32-bit)
* Data length: 4096 bytes

Regards,
Dai