

NISTIR 7896

Third-Round Report of the SHA-3 Cryptographic Hash Algorithm Competition

Shu-jen Chang
Ray Perlner
William E. Burr
Meltem Sönmez Turan
John M. Kelsey
Souradyuti Paul
Lawrence E. Bassham

<http://dx.doi.org/10.6028/NIST.IR.7896>

NISTIR 7896

Third-Round Report of the SHA-3 Cryptographic Hash Algorithm Competition

Shu-jen Chang
Ray Perlner
William E. Burr
Meltem Sönmez Turan
John M. Kelsey
Souradyuti Paul
Lawrence E. Bassham
*Computer Security Division
Information Technology Laboratory*

<http://dx.doi.org/10.6028/NIST.IR.7896>

November 2012



U.S. Department of Commerce
Rebecca M. Blank, Acting Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Under Secretary of Commerce for Standards and Technology and Director

Abstract

The National Institute of Standards and Technology (NIST) opened a public competition on November 2, 2007, to develop a new cryptographic hash algorithm – SHA-3, which will augment the hash algorithms specified in the Federal Information Processing Standard (FIPS) 180-4, *Secure Hash Standard (SHS)*. The competition was NIST’s response to advances in the cryptanalysis of hash algorithms.

NIST received sixty-four submissions in October 2008, and selected fifty-one first-round candidates on December 10, 2008; fourteen second-round candidates on July 24, 2009; and five third-round candidates – BLAKE, Grøstl, JH, Keccak and Skein, on December 9, 2010, to advance to the final round of the competition. Eighteen months were provided for the public review of the finalists, and on October 2, 2012, NIST announced the winning algorithm of the SHA-3 competition – Keccak. This report summarizes the evaluation of the five finalists and the selection of the SHA-3 winner.

KEY WORDS: Cryptographic hash algorithm; Cryptographic hash function; Cryptography; Cryptographic hash competition; SHA-3 competition.

Acknowledgements

NIST thanks the submitters of all the candidate algorithms, especially the submitters of the SHA-3 finalists for their continued diligence and support. NIST is also grateful for the efforts of those in the cryptographic community that provided security, implementation, and performance analyses of the candidate algorithms throughout the competition, and those who provided feedback on the hash forum or published papers on the various technical aspects of the candidates. Specifically, NIST thanks the organizers of the following projects:

SHA-3 Zoo:

Christian Rechberger, Jean-Phillipe Aumasson, Florian Mendel, Tomislav Nad, Martin Schl affer, Gilles van Assche;

ECRYPT Benchmarking of All Submitted Hashes (eBASH):

Daniel J. Bernstein, Tanja Lange;

eXternal Benchmarking eXtension (XBX):

Christian Wenzel-Benner, Jens Gr af;

George Mason University Department of Electrical and Computer Engineering Hardware Benchmarking:

Kris Gaj, Jens-Peter Kaps;

Virginia Tech Department of Electrical and Computer Engineering ASIC Benchmarking:

Patrick Schaumont, Leyla Nazhand-Ali;

Eidgen ossische Technische Hochschule Z urich (ETHZ) ASIC Implementation:

Frank K. G urkaynak;

And the authors of the following reports:

ECRYPT II SHA-3 Design and Cryptanalysis Report:

Christian Rechberger, Tor E. Bj rstad, Joan Daemen, Christophe De Canni ere, Praveen Gauravaram, Dmitry Khovratovich, Willi Meier, Tomislav Nad, Ivica Nikoli c, Matt Robshaw, Martin Schl affer, S oren S. Thomsen, Elmar Tischhauser, Deniz Toz, Gilles Van Assche, Kerem Varicı;

ECRYPT II Intermediate Status Report:

Praveen Gauravaram, Florian Mendel, Mar a Naya-Plasencia, Vincent Rijmen, Deniz Toz;

for their outstanding contributions and support to the SHA-3 competition.

In addition, NIST extends its appreciation to the KU Leuven Department Elektrotechniek-ESAT/COSIC team led by Bart Preneel, and Sebastiaan Indesteege for their outstanding support to the First SHA-3 Candidate Conference.

The authors of this report also thank NIST's Hirofumi Sakane and Caroline Scace for their support in conducting power analysis of the finalists. Last but not least, the authors thank the other members of NIST's SHA-3 team, who reviewed the candidate algorithms and the public comments, performed testing, provided technical input and administrative support, and participated in numerous meetings during the five-year competition. They are: Elaine B. Barker, Sara J. Caswell, Donghoon Chang, Lily Chen, Quynh Dang, Morris J. Dworkin, James R. Nechvatal, Rene Peralta, William T. Polk, and Andrew Regenscheid.

TABLE OF CONTENTS

1. Introduction	1
1.1 Purpose of this Document.....	1
1.2 Background	1
1.3 Organization of this Document.....	3
2. Evaluation Criteria.....	4
2.1 Security	4
2.2 Cost and Performance.....	4
2.3 Algorithm and Implementation Characteristics	4
3. Selection Process.....	5
3.1 Security	5
3.2 Performance.....	6
3.3 Other Algorithm and Implementation Characteristics	6
3.4 Complementing SHA-2.....	7
3.5 Selection Conclusion	7
4. Security Analysis of the Finalists.....	9
4.1 Security Overview.....	9
4.1.1 Overview of Security Resources.....	10
4.1.2 Domain Extenders and Proofs.....	10
4.1.3 Cryptanalysis and Security Margin	12
4.1.4 Distinguishing Attacks and Differential Properties.....	15
4.1.5 Depth of Analysis and Understandability of Algorithms.....	16
4.1.6 Tweak History of the Finalists.....	16
4.1.7 Side Channel Attacks and Countermeasures	18
4.2 Finalist Profiles and Cryptanalysis	19
4.2.1 BLAKE	20
4.2.2 Grøstl	23
4.2.3 JH.....	26
4.2.4 Keccak	28
4.2.5 Skein	31
4.3 Security Summary	33
5. Performance Comparison of the SHA-3 Finalists	35
5.1 Software Performance	35
5.1.1 Computer Systems – the Current Playing Field	35
5.1.2 Candidate Software Performance Studies	36
5.1.3 Beyond The Superscalar	42
5.1.4 Software Performance Summary.....	45
5.2 Hardware Performance.....	46
5.2.1 High-Performance Implementations	48
5.2.2 Compact Implementations.....	52
5.2.3 Discussion of SHA-2 and the SHA-3 Finalists	55
5.2.4 Hardware Performance Summary	57
6. Other Considerations.....	58

6.1 Intellectual Property	58
6.2 Other Features	58
7. Conclusion.....	59
References.....	60
Appendix A – eBASH Shootout Plots	71

List of Acronyms

AES-NI	Advanced Encryption Standard – New Instructions
ALU	Arithmetic Logic Unit
AMD	Advanced Micro Devices
ARM	Advanced RISC Machine
ARX	Modular Addition, Rotation, and eXclusive OR
ASIC	Application Specific Integrated Circuit
AVX	Advanced Vector eXtensions
CF	Compression Function
CISC	Complex Instruction Set Computer
CMOS	Complementary Metal–Oxide–Semiconductor
CPU	Central Processing Unit
DPA	Differential Power Analysis
eBASH	ECRYPT Benchmarking of All Submitted Hashes
EM	Electromagnetic
ETHZ	Eidgenössische Technische Hochschule Zürich
FIPS	Federal Information Processing Standard
FPGA	Field Programmable Gate Array
FRN	Federal Register Notice
GMU	George Mason University
GPU	Graphics Processor Unit
HAIFA	HAsh Iterative FrAmework
HMAC	Keyed-Hash Message Authentication Code
ISA	Instruction Set Architecture
IV	Initialization Value
kGE	kilo Gate Equivalents
MAC	Message Authentication Code
MD	Merkle-Damgård
MIPS	Microprocessor without Interlocked Pipeline Stages
NIST	National Institute of Standards and Technology
OC-48	Optical Carrier-48 (Transmission Rate)
PPC	PowerPC
PRF	Pseudo Random Function
RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
RSA	Rivest, Shamir, and Adelman
RSA-FDH	RSA Full Domain Hash
SCA	Side Channel Analysis
SFS	Semi Free Start
SHA-3	Secure Hash Algorithm-3
SHS	Secure Hash Standard
SIMD	Single Instruction Multiple Data
SSE	Streaming SIMD Extensions
UBI	Unique Block Iteration
VT	Virginia Tech
XBX	eXternal Benchmarking eXtension
XOR	Exclusive OR

1. Introduction

1.1 Purpose of this Document

This report summarizes the third and final round of the SHA-3 (Secure Hash Algorithm-3) Cryptographic Hash Algorithm Competition. The third round began on January 31, 2011, when the National Institute of Standards and Technology (NIST) posted the finalists' submissions for the third round, and ended on October 2, 2012, when NIST announced the SHA-3 winner – Keccak. This report explains the evaluation and selection process.

1.2 Background

NIST opened a public competition on November 2, 2007, to develop a new cryptographic hash algorithm (referred to as SHA-3) to augment the hash algorithms specified in Federal Information Processing Standard (FIPS) 180-2, *Secure Hash Standard*¹ [1]. The competition was NIST's response to advances in the cryptanalysis of hash algorithms in recent years. An attack by Wang et al. [4], and extended by many others, introduced serious concerns about the security of the SHA-1 government standard hash algorithm when used for the generation of digital signatures and in other security applications that require collision resistance.

NIST proposed a *Draft Minimum Acceptability Requirements, Submission Requirements, and Evaluation Criteria for Candidate (Hash) Algorithms* for public comment in a Federal Register Notice in January 2007 (FRN-Jan07) [5]. These requirements and evaluation criteria were updated, based on public feedback, and included in a later, second Federal Register Notice published on November 2, 2007 (FRN-Nov07) [6], which called for the submission of candidate algorithms and launched the "SHA-3" competition.

The competition, modeled after the Advanced Encryption Standard (AES) Competition [7] in 1997, generated great interest in the cryptographic community, and inspired numerous SHA-3 entries from around the world. NIST received sixty-four entries by October 2008, and selected fifty-one first-round candidates in December 2008, fourteen second-round candidates in July 2009, and five third-round candidates – BLAKE, Grøstl, JH, Keccak, and Skein, on Dec. 9, 2010, to advance to the final round of the competition. Table 1 lists these finalists and their submitters.

Algorithm	Designer
BLAKE [8]	Jean-Philippe Aumasson, Luca Henzen, Willi Meier, Raphael C.-W. Phan
Grøstl [9]	Praveen Gauravaram, Lars R. Knudsen, Krystian Matusiewicz, Florian Mendel, Christian Rechberger, Martin Schl�affer, S�oren S. Thomsen
JH [10]	Hongjun Wu
Keccak [11]	Guido Bertoni, Joan Daemen, Micha�el Peeters, Gilles Van Assche
Skein [12]	Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, Jesse Walker

Table 1. SHA-3 Finalists

¹ FIPS 180-2 [1], *Secure Hash Standard*, was the government hash standard when the SHA-3 competition began; it was superseded by FIPS 180-3[2], *Secure Hash Standard*, in October 2008, and again by FIPS 180-4 [3], *Secure Hash Standard*, in March 2012.

NIST has published a report for the first and second round of the competition. These reports [13, 14], and all the relevant information about the SHA-3 competition, are available at NIST’s hash competition website: <http://www.nist.gov/hash-competition>.

For the third round of the competition, NIST also allowed the submitters of the remaining candidates to make minor modifications to their algorithms by Jan. 16, 2011. Such modifications are often referred to as “tweaks” by the cryptographic community. NIST posted the final tweaked submissions on January 31, 2011, and started the third round of the competition.

The review process of the finalists spanned nineteen months, and included a final SHA-3 Candidate Conference [15] held in March 2012. Throughout the competition, the cryptographic community provided an enormous amount of feedback to NIST. Most of the comments were sent to NIST or to the Hash Forum [16], which is a public email list, hash-forum@nist.gov, established to facilitate dialogue about the hash workshops that NIST held prior to the competition, and the SHA-3 competition. In addition, some organizers held hash workshops [17, 18, 19, 20, 21, 22, 23], or set up public websites to facilitate cryptanalysis or performance benchmarking of the candidates, such as the SHA-3 Zoo [24], the ECRYPT Benchmarking of All Submitted Hashes (eBASH) [25], and the eXternal Benchmarking eXtension (XBX) [26] sites. Many of the cryptanalysis and performance studies were published as papers in major cryptographic conferences or leading cryptographic journals. ECRYPT II has also published two reports [27, 28] on the SHA-3 candidates prior to the finalists’ third-round tweaks.

Based on the public feedback and internal review of the finalists, NIST selected Keccak as the SHA-3 winner. NIST made the selection announcement on October 2, 2012 [29], and officially ended the SHA-3 competition. Table 2 shows the competition timeline, including major events leading to the start of the SHA-3 competition.

Date	Event
10/31-11/1/2005	Cryptographic Hash Workshop [30], NIST, Gaithersburg, Maryland.
8/24-8/25/2006	Second Cryptographic Hash Workshop [31], UCSB, California.
1/23/2007	Issued the Federal Register Notice <i>Announcing the Development of New Hash Algorithm(s) for the Revision of Federal Information Processing Standard (FIPS) 180-2, Secure Hash Standard</i> [5].
11/2/2007	Issued the Federal Register Notice <i>Announcing a Request for Candidate Algorithm Nominations for a New Cryptographic Hash Algorithm (SHA-3) Family</i> [6]. SHA-3 competition began.
10/31/2008	SHA-3 Submission Deadline.
12/10/2008	First-round candidates announced. The first round began.
2/25-2/28/2009	First SHA-3 Candidate Conference [32], KU Leuven, Belgium.
7/24/2009	Second-round candidates announced. The first round ended.
9/28/2009	Second-round candidates posted for public review. The second round began.
8/23-8/24/2010	Second SHA-3 Candidate Conference [33], UCSB, California.
12/9/2010	SHA-3 finalists announced. The second round ended.
1/31/2011	Third-round candidates posted for public review. The third round began.
3/22-3/23/2012	Third SHA-3 Candidate Conference [15], Washington, D.C.
10/2/2012	Keccak announced as the SHA-3 winner. SHA-3 competition ended.

Table 2. SHA-3 Competition Timeline

1.3 Organization of this Document

This document is organized as follows. Section 2 summarizes the evaluation criteria specified in FRN-Nov07. Section 3 describes the selection process of the SHA-3 winner. Section 4 summarizes the five finalists and their security analyses that are publically available. Section 5 compares the software and hardware performance of the finalists. Section 6 discusses other consideration factors associated with the finalists. Section 7 concludes the report.

2. Evaluation Criteria

Throughout the competition, NIST used the evaluation criteria specified in FRN-Nov07 [6] to evaluate the candidates. These criteria were discussed and further clarified at the First SHA-3 Candidate Conference in KU Leuven, Belgium, in February 2009. In order of importance, NIST considered the security, cost, and algorithm and implementation characteristics of a candidate in its selection decision.

2.1 Security

FRN-Nov07 identified security as the first evaluation criterion, and listed the following as the major evaluation factors:

- (i) applications of hash functions;²
- (ii) specific requirements when hash functions are used to support Keyed-Hash Message Authentication Code (HMAC) [34], Pseudo Random Functions (PRFs), or Randomized Hashing;
- (iii) additional security requirements of hash functions;
- (iv) evaluations relating to attack resistance; and
- (v) other consideration factors.

NIST studied the large amount of feedback received from the cryptographic community, and discusses the security analysis of the finalists in Section 4.

2.2 Cost and Performance

FRN-Nov07 identified cost as the second evaluation criterion, which includes:

- (i) *computational efficiency*, which refers to the speed of the algorithm; and
- (ii) *memory requirements*, which include the code size and the random-access memory (RAM) requirements for software implementations, as well as the gate-counts for hardware implementations.

For the cost and performance evaluations of the finalists, NIST studied the results obtained from several major benchmarking efforts. The candidates' performance in both software and hardware implementations are compared in Section 5.

2.3 Algorithm and Implementation Characteristics

FRN-Nov07 also identified an evaluation criterion related to the *flexibility* and *simplicity* of a candidate's design. Candidates with greater flexibility are preferable, such as flexibility in running efficiently on a wide variety of platforms, or flexibility in using parallelism, or instruction-set extensions, to achieve higher performance. Furthermore, a candidate would be judged on its relative design simplicity, with the intent to encourage easy-to-understand and easy-to-analyze designs to provide more confidence in its security.

These characteristics of the finalists are addressed in Sections 3, 4, 5, and 6, as appropriate.

² In this report, the terms "hash function" and "hash algorithm" are used interchangeably.

3. Selection Process

The selection of Keccak as the SHA-3 algorithm was a very difficult decision – all five finalists would have made acceptable choices for SHA-3. NIST’s decision was based on a number of considerations, particularly security, performance, and the desire for SHA-3 to complement the existing SHA-2 algorithms. In addition to rating highly with respect to these primary considerations, Keccak also benefited from a simple and elegant design, a great deal of flexibility in choosing parameters, and other features, such as a built-in authenticated-encryption mode.

3.1 Security

NIST compared the security properties of the finalists, based on the published design documents and cryptanalysis produced during the competition, and on the security proofs provided by both the designers and other researchers. This data was useful in making the decision, but it required some interpretation; each algorithm has its own mix of published attacks and design logic, and each has received a somewhat different depth of cryptanalysis. However, NIST was able to extract the following points from this information:

- a. No finalist has a published attack that, in any real sense, threatens its practical security, and each of the finalists has a security proof for its *domain extender*³ that shows that the security of the whole hash function is ultimately based on the security of its underlying components.
- b. NIST was able to estimate the security margins of the candidates, based on the largest number of rounds of the hash functions or underlying components that has been successfully attacked using collision-type attacks. By this metric, which is discussed in more detail in Section 4.1.3, Grøstl and JH have the smallest security margins, Skein has a somewhat larger security margin than Grøstl and JH, and BLAKE and Keccak have very large security margins. None of the candidates has an absolutely unacceptable security margin, but it would be difficult for NIST to select Grøstl or JH without seriously considering adding more rounds.
- c. There are also a number of distinguishing attacks and differential properties published on the finalists and their underlying component functions. It was not entirely clear how to interpret these results, but they largely left the same impression as the collision-type attacks: if distinguishers costing more than the claimed security strength of a 512-bit hash function are omitted, Grøstl and JH have differential properties that extend through most or all of the rounds, while the remaining three candidates only have distinguishers that extend through about half of their rounds.
- d. The cryptanalysis performed on BLAKE, Grøstl, and Skein appears to have a great deal of depth, while the cryptanalysis on Keccak has somewhat less depth, and that on JH has relatively little depth.
- e. While the tweaks on BLAKE, Keccak, and JH were sufficiently minor in that all the cryptanalysis performed throughout the competition applies to their final versions, Grøstl and Skein were tweaked more significantly to address attacks on their first- and second-round versions. That said, even in the cases of Grøstl and Skein, most of the earlier cryptanalysis carries over to the final version with only minor modifications; only the

³ A definition for this term is provided in Section 4.1.2.

attacks targeted by the third-round tweaks are significantly affected, and in both cases, the tweaks appear to be effective in eliminating the relevant vulnerabilities.

- f. NIST also considered the security of the finalists against side channel analysis, but there was little available to distinguish between the finalists. While a case can be made that ARX-based (Modular Addition, Rotation, and eXclusive OR-based) designs like BLAKE and Skein may be more expensive to secure against side channel analysis than logic-based designs like Keccak and JH, any of the candidates can probably be made secure against side channel analysis with sufficient effort on the implementation end.

Keccak received a significant amount of cryptanalysis, although not quite the depth of analysis applied to BLAKE, Grøstl, or Skein. This cryptanalysis left Keccak with a huge security margin – only five out of its 24 rounds has been broken by a near-collision attack. Keccak’s large security margin, after significant cryptanalytic effort, suggests that Keccak will remain secure in the future.

3.2 Performance

NIST was fortunate to have a great depth of performance data on the five finalists that could also be compared with the performance data of the SHA-2 algorithms. This data included software implementations on many different kinds of Central Processing Units (CPUs), and hardware implementations in both Field Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs). All this data made simple comparisons very difficult; most algorithms excelled on some platforms and lagged on others. However, a few patterns emerged from the performance data, which affected NIST’s decision:

- a. All five finalists perform well enough to be usable in most applications.
- b. None of the five finalists is the best for every application, and none offers really compelling improvements over the SHA-2 algorithms.
- c. The ARX-based algorithms, BLAKE and Skein, perform extremely well in software.
- d. Keccak has a clear advantage in throughput/area performance in hardware implementations.
- e. Grøstl and JH are considerably slower than the other three algorithms in most software implementations.

Keccak provides an excellent set of performance trade-offs – it is broadly competitive with SHA-2 in software, while providing much better throughput/area performance in hardware.

3.3 Other Algorithm and Implementation Characteristics

In addition to the bare essentials of security and performance, NIST also considered *other features*, which was another evaluation criterion identified in FRN-Nov07.

Keccak benefits from a very clean and elegant design, which is adaptable to a wide variety of circumstances. Its sponge domain extender makes it easy to analyze the security implications of changing the message block size to improve performance, and provides a mechanism for extending the output size to fit applications like RSA Full Domain Hash (RSA-FDH). Keccak

also provides an 800-bit version of its internal permutation, which, if studied further from a security perspective, may be vetted and approved by NIST at some time in the future. If this possibility is realized, the 800-bit permutation will increase the flexibility of Keccak even further.

While several of the candidates provide additional functionality above and beyond the simple drop-in replacement for existing hash functions, Keccak's built-in authenticated-encryption mode is one of the better "extras" offered and is a very natural extension of the sponge construction.

Overall, Keccak's elegant design, flexibility, and "extras," like its authenticated-encryption mode, were points in its favor as the SHA-3 winner.

3.4 Complementing SHA-2

One additional consideration contributed to the selection of Keccak as the SHA-3 winner: SHA-3 will be deployed into a world in which the use of SHA-2 is widespread, and will coexist with SHA-2 as an approved hash function. This makes a strong argument for an SHA-3 design that complements that of SHA-2.

- a. SHA-2 is a software-oriented ARX design, as are BLAKE and Skein. In terms of performance and implementation properties, these two algorithms are rather similar to SHA-2.
- b. Because SHA-2 is an ARX-based design with a key schedule, it has some important design elements in common with BLAKE and Skein, although neither is closely related to SHA-2. However, cryptanalytic tools that apply to SHA-2 in the future seem more likely to apply to BLAKE or Skein than to the other three finalists.
- c. Keccak is a hardware-oriented design that is based entirely on simple bit-oriented operations and moving bits around. Its performance and implementation properties are significantly different from those of SHA-2.

Among all the finalists, NIST believes that Keccak provides the most advantages as an addition to NIST's toolbox of cryptographic algorithms. Keccak's software performance is broadly comparable to that of the SHA-2 algorithms, but its hardware performance is much better, and Keccak does not need the 32- or 64-bit integer additions or key schedule that are required by SHA-2.

3.5 Selection Conclusion

NIST's decision rested almost entirely on the generous contributions of time and effort by the cryptographic community. Multiple implementations of the algorithms in hardware and software made meaningful performance comparisons possible, and most of this work was done by volunteers. Analyzing the security of a new cryptographic hash function is extremely demanding, requiring hard work from people with rare expertise and talents. This analysis was provided by the members of the cryptographic community, including, but not limited to, the submitters of the SHA-3 candidates. Without this huge volume of published information, representing a great many years of skilled effort, NIST could never have run this competition.

Based on the data available, all five finalists appear to be excellent hash algorithms, providing acceptable performance and security properties. Any of the five would have made an acceptable choice for SHA-3.

Keccak offers acceptable performance in software, and excellent performance in hardware. It has a large security margin, suggesting a good chance of surviving without a practical attack during its working lifetime. Keccak is also a fundamentally new and different algorithm that is entirely unrelated to the SHA-2 algorithms. For all these reasons, NIST chose Keccak as the new SHA-3 algorithm. Keccak's bit-oriented design may facilitate better understanding in the future, but at present the newness of the design means it is somewhat less well-understood than the ARX or S-box-based designs.

4. Security Analysis of the Finalists

This section describes the security of the five finalists. It is organized as follows: Section 4.1 gives an overview of the criteria that NIST used in comparing the depth and significance of published cryptanalysis results on the five finalists and includes several tables comparing the finalists according to these criteria. Section 4.2 summarizes the design and lists the published cryptanalysis results for each of the finalists. Section 4.3 gives NIST's assessment of what conclusions can be drawn from the results described in the first two subsections.

4.1 Security Overview

Security is the most important criterion for a new hash function; performance or implementation advantages of a broken hash function are irrelevant. This section of the report discusses how NIST weighed the evidence regarding the security of each of the five finalists, and how this weighing helped determine the final selection of SHA-3.

The submission packages of the five finalists, and the published analysis of the finalists during the four years of public review, provide significant information on their security. This information falls into six categories:

- a. *Proofs of security relating to the hash functions' structure.* These proofs typically show that the hash function has various good properties, assuming some good properties of the underlying compression function, block cipher, or fixed permutation.
- b. *Cryptanalysis results on the hash function or its components that relate directly to the core security properties of a hash function: preimage-resistance, second-preimage-resistance, and especially, collision-resistance.* Even when these attacks involve giving an attacker more power than they would have in an attack on the full hash function, as with semi-free-start collisions, these results are often the best indication available of how well the hash function will survive future attacks. The fraction of the hash function (or underlying component function) broken by these attacks can be used to calculate a *security margin* - the fraction of the hash function's rounds that are still unbroken. For the SHA-3 finalists and SHA-2, collision and near-collision attacks are the most informative source of information about the hash functions' security; the published preimage attacks have much higher required work,⁴ and are, therefore, less worrisome.
- c. *Theoretical distinguishers on the hash function or its components.* These involve some property that would not be present in an ideal hash function or component, but which has no obvious relevance to any practical security concern about the hash function. Sometimes, using these properties even requires more work than the claimed security level of the hash function.
- d. *Depth and maturity of cryptanalysis results.* Some algorithms have received more cryptanalysis than others; some have analysis that seems to be noticeably more mature (for example, many different papers building on one another would indicate maturity). This is necessarily somewhat subjective; different cryptographers might come to somewhat different conclusions. A naïve approach might merely count the cryptanalysis papers. This is important in evaluating the meaningfulness of the security margins of the hash function; a hash function

⁴ This reflects the fact that a brute-force collision attack on a hash function requires only the square root of the work required for a brute-force preimage attack.

that has had little analysis may have a large security margin simply because few cryptanalysts have spent any time studying it.

- e. *Tweak history of the finalists.* Prior to the start of the second and third rounds of the competition, the designers of the remaining candidates were allowed to make minor changes, referred to as “tweaks,” to their designs. Some designers simply tweaked their algorithms by increasing the number of rounds, which does not invalidate previous analysis results; others chose to modify the internal structure of a component function to address a particular attack, or to improve some property of the hash function. Tweak history affects the meaningfulness of the security margin of a hash function in a way similar to that of the depth of cryptanalysis. If an algorithm has been tweaked in ways that invalidate previous cryptanalysis results, then the cryptanalysis on the final version of the algorithm is correspondingly less mature.
- f. *Side channel analysis.* Different algorithms have different implementation properties, and thus may have different susceptibility to side channel attacks, or may be more or less expensive to defend against such attacks. While most applications of hash functions do not need resistance against side channel attacks, some important ones (most notably MAC algorithms and key-derivation functions) do require this resistance, so differences in vulnerability to side channel attacks, or in the cost of defending against them, matter for the selection of SHA-3.

No finalist was broken or seriously threatened by any attack published so far. All the finalists have proofs relating their security to ideal properties of their underlying compression function, permutation, or block cipher. Two of the finalists (Keccak and JH) have distinguishing attacks on their underlying primitives. However, the distinguishing attack on Keccak's permutation is far beyond the security level claimed by a 512-bit hash function. The security results by themselves did not disqualify any finalists, but they offered the best-available evidence for assessing the likelihood that future cryptanalysis may threaten these algorithms.

4.1.1 Overview of Security Resources

NIST's understanding of the security of the finalist hash functions is based on the work of the cryptographic community, as it appears in published cryptanalysis and security-proof papers. Some of these papers were presented at the NIST-sponsored SHA-3 conferences; others were presented at other conferences. In general, the results have had some level of peer review, though it is always possible that incorrect results could have made it past the review. Some results provide concrete findings, such as collisions for reduced-round versions of the hash functions, which can be verified; others are too expensive to be implemented in practice.

4.1.2 Domain Extenders and Proofs

Hash function design can be broken into two parts: building a fixed-size component like a compression function or fixed permutation, and then using that fixed-size component to build a hash function that can process variable-length strings up to some huge maximum length. The construction used to build the full hash function is called a *domain extender*. The next two subsections discuss the analysis of the fixed-size components used in the finalist hash functions. This section discusses the analysis of the domain extenders of the finalists. Unlike the cryptanalysis results, the results on domain extenders are mainly built on proofs relating the properties of the full hash function to the properties of the underlying fixed-length components.

BLAKE, Grøstl, and Skein build their full, arbitrary-input-length hash function from a fixed-input-length hash function called a *compression function*. At a lower level, the compression functions of BLAKE and Skein are based on block ciphers, whereas that of Grøstl is based on a pair of fixed permutations. The security properties of these hash functions can either be viewed as being inherited from their compression functions or directly from the underlying components of their compression functions. JH and Keccak build their full hash functions directly from a fixed permutation. In each case, the full hash function inherits its security properties from the properties of this permutation.

Published proofs show that the security of all five finalists can be inferred from properties of their underlying building blocks. This kind of proof cannot guarantee that the hash function is secure, but it does show that certain kinds of weaknesses in the full hash function can only happen as a result of corresponding weaknesses in its underlying building blocks. Because of this, proofs of this kind guarantee security against generic attacks – attacks that only exploit the domain extender and not the internals of the underlying building blocks. Since cryptanalysts and designers spend most of their time considering the underlying building blocks, this is an important guarantee.

Table 3 summarizes the security proofs for the finalists' and SHA-2's domain extenders. All numbers in this table are in bits. The sizes of the key, block, and tweak are denoted by k , b and t , respectively, and L denotes the length of the target message in blocks. “Coll,” “Pre,” “2nd Pre,” and “Indiff” denote the approximate proven generic security against collision, preimage, second preimage, and indifferenciability attacks, respectively. References for the finalists' security bounds are presented in Subsections 4.2.1.1 – 4.2.5.1.

For the purpose of the SHA-3 decision, these results can be summarized as follows:

- a. All five finalists have indifferenciability proofs, which guarantee that the hash function resists generic attacks up to at least the complexity of a brute-force collision search, assuming that the underlying primitive is ideal.⁵
- b. Under the same assumption, all of the candidates except JH are proven to have near-optimal security relating to the core properties of collision, preimage, and second-preimage resistance. Even in the case of JH, the indifferenciability proof guarantees some generic security for these properties (the optimal security, in the case of collision search), and there is no known way for an attacker to exploit the JH domain extension to find preimages, or second preimages more efficiently than a brute-force search.

Nothing in these results provides a strong justification for selecting or discarding any of the five finalists. However, these results are quite important for evaluating the finalists' security, as they show that the security of the full hash functions depends only on the security of the underlying primitives, which is where the overwhelming majority of the cryptanalytic effort has been focused so far.

⁵ Due to the length extension property, the indifferenciability bound for SHA-2 is 1. According to FRN-Nov07, SHA-3 candidates were required to resist attacks exploiting this property, but this was not considered a requirement for earlier hash functions (i.e., those developed prior to the SHA-3 competition). Essentially all Internet protocols, in particular those using HMAC [34], are designed to be secure even if they use a hash function that has the length extension property.

Algorithm	Domain Extender	Underlying Primitive	Primitive size	Hash size	Security			
					Coll	Pre	2 nd Pre	Indiff
BLAKE	HAIFA	Block cipher	$k=512$ $b=512$	224 256	112 128	224 256	224 256	128 128
			$k=1024$ $b=1024$	384 512	192 256	384 512	384 512	256 256
Grøstl	Grøstl	A pair of permutations	512 512	224 256	112 128	224 256	$256 - \log_2 L$	128 128
			1024 1024	384 512	192 256	384 512		$512 - \log_2 L$ 256
JH	JH	Permutation	1024	224 256 384 512	112 128 192 256	224 256 256 256	224 256 256 256	256 256 256 256
				224 256 384 512	112 128 192 256	224 256 384 512	224 256 384 512	224 256 384 512
Keccak	Sponge	Permutation	1600	224 256 384 512	112 128 192 256	224 256 384 512	224 256 384 512	224 256 384 512
				224 256 384 512	112 128 192 256	224 256 384 512	224 256 384 512	224 256 384 512
Skein	UBI	Tweakable block cipher	$k=512$ $b=512$ $t=128$	224 256 384 512	112 128 192 256	224 256 384 512	224 256 384 512	256 256 256 256
			$k=512$ $b=256$	224 256	112 128	224 256	$256 - \log_2 L$	1
SHA-2 ⁶	MD	Block cipher	$k=1024$ $b=512$	384 512	192 256	384 512	$512 - \log_2 L$	1

Table 3. Proven Generic Security⁷ of the SHA-3 Finalists and SHA-2 in Bits

4.1.3 Cryptanalysis and Security Margin

The most important security requirements for a hash function are collision and preimage resistance. Informally, collision-resistance means that it is very difficult to find two different input strings that yield the same message digest, while preimage-resistance means that it is very difficult to find an input string that yields a particular desired message digest. This subsection discusses the published attacks on the hash functions that relate to these core security properties of a hash function. In particular, the subsection focuses on the security against collision and related attacks, rather than on preimage attacks.

Given the attacks published on the SHA-3 finalists, collision attacks are more informative about the ultimate security of the hash functions for three reasons. First, it is often possible to speed up the brute-force search of a preimage by a small number of bits using a clever representation of the internal state of the hash function and its computations – there are a number of results along these lines on the finalists, and while the techniques used are clever, it is difficult to see any implications of these attacks for the security of the hash functions. Second, brute-force preimage

⁶ SHA-2's domain extender is optimally collision and preimage resistant under the ideal cipher model [35]. The security of the Merkle-Damgård (MD) domain extender against second-preimage attacks is $n - \log_2 L$, where L is the length of the message in blocks [36], assuming the compression function behaves like a random oracle.

⁷ Note that these security estimates ignore any constant factors in the proofs that are unrelated to the input and output sizes of the hash function and its components.

attacks on the SHA-3 finalists are extremely computationally demanding, requiring at least 2^{224} compression function computations, and the published cryptanalytic preimage attacks typically have much higher complexities than the published collision and related attacks for this reason. And finally, for the SHA-3 finalists, the preimage attacks were not the attacks that covered the most rounds with more than a tiny advantage over a brute-force search.

Table 4 shows the collision-type attacks that cover the largest number of rounds for each of the finalists, while maintaining an attack complexity less than that of a brute-force collision search. The best-known collision attack on SHA-2 is also shown, for reference. The column marked “CF Call” gives the time complexity of the best attack, relative to the computational cost of one execution of the hash algorithm’s compression function.

Algorithm	Attack Type	Target	Rounds	Percent Broken	CF Call	Reference
BLAKE	Semi-free-start near collision	Compression function	4/14	29 %	2^{21}	[37]
Grøstl	Semi-free-start collision	Compression function	6/10	60 %	2^{120}	[38]
JH	Semi-free-start near collision	Compression function	26/42	62 %	2^{112}	[39]
Keccak	Near collision	Hash function	5/24	21 %	Example given	[40]
Skein	Semi-free-start near collision	Compression function	32/72	44 %	2^{105}	[41]
SHA-2	Collision	Hash function	24/64	38 %	Example given	[42]

Table 4. Best Known Collision-type Attacks against the SHA-3 Finalists and SHA-2

In many ways, these attacks are hard to compare with one another. For example, the best attacks on BLAKE, Keccak, and SHA-2 can be implemented in practice, whereas the best attacks on Grøstl, JH, and Skein would require huge computational resources to implement, and cannot be fully verified. The best attacks on Keccak and SHA-2 apply to the hash function as a whole, while the best attacks on BLAKE, Grøstl, JH, and Skein apply to their underlying compression functions, and give the attacker more control over the inputs to those functions than they would have in an attack on the full hash function. More fundamentally, comparing cryptanalytic results can favor algorithms that are not heavily analyzed, or algorithms that require the development of new techniques to analyze them.

Despite these limitations, however, these results are what is available; they, along with some security arguments offered by the designers, are the basis for comparing the finalists' security against cryptanalysis. It is a truism that attacks only get better. The SHA-3 competition has brought a great deal of analysis to bear on the finalists, but the winner will inevitably be the focus of even more cryptanalytic effort, and so it must be assumed that the existing published attacks will improve over time.

The first piece of information that can be taken from these results is that, after four years of intense analysis, no attack has come anywhere close to threatening the core security properties of any of the five finalists. This provides some reassurance about the SHA-3 process. If two or three of the five finalists had been broken or seriously threatened by now, this would raise questions about the community's ability to develop secure hash functions.

The second piece of information from these results is that there are differences in the size of the security margins of the different finalists, where the security margin is defined as the fraction of the hash or compression function that has not been successfully attacked. (For example, an attack on six rounds of a ten-round hash function would give a 40 % security margin.) Keccak has the largest security margin, with 79 % of its hash function still unbroken; JH has the smallest security margin, with 38 % unbroken. Comparing these security margins must be done carefully, given the different kinds of attacks, work factors, and amounts of cryptanalysis performed on the different candidates. Table 5 tries to summarize this picture. The depth of analysis is a rather subjective measure, based on NIST’s reading of the literature. This measure takes into account not only the number of cryptanalysis papers published, but also the depth and maturity of the analysis and the tools used to cryptanalyze the algorithm. The “depth of cryptanalysis” measure is described in somewhat more detail in Section 4.1.5.

Algorithm	Best Attack	Security Margin	Work	Depth of Analysis
BLAKE	Semi-free-start near collision	71 %	Practical	High
Grøstl	Semi-free-start collision	40 %	Impractical	Very High ⁸
JH	Semi-free-start near collision	38 %	Impractical	Low
Keccak	Near collision	79 %	Practical	Medium
Skein	Semi-free-start near collision	56 %	Impractical	High ⁸
SHA-2	Collision	62 %	Practical	Medium

Table 5. Security Margins for the Five Finalists Based on Collision-type Attacks

This way of looking at the cryptanalysis results is quite unfavorable to JH, which received relatively limited cryptanalysis, and still seems to have a rather small security margin. It is relatively favorable to BLAKE and Keccak (with large security margins, despite a reasonable depth of analysis), and also to Skein (with a relatively large security margin, despite a reasonable depth of analysis, and a best attack with an impractically high amount of work required). The picture for Grøstl is less clear; Grøstl has a relatively low security margin, but is the most deeply analyzed finalist, and its best attack requires an impractically high amount of work and applies only to its compression function.

Considering this table alongside the performance of the different algorithms makes the picture somewhat clearer: JH and Grøstl are generally the two slowest algorithms in software. An extremely fast algorithm whose security margin is uncomfortably narrow most likely can simply be tweaked by adding rounds, and NIST explicitly stated in its call for submissions that it might adjust the number of rounds for the winner. This is not an appealing option for an algorithm whose performance is already marginal. If JH or Grøstl were selected, they would likely have to keep their relatively narrow security margins in order to maintain acceptable performance.

⁸ These two algorithms were significantly tweaked for the third round. Some of the cryptanalysis performed on the earlier versions of the algorithms no longer applies.

Considering the 512-bit versions of the hash functions gives very similar results; there are somewhat fewer results on 512-bit versions, but because brute-force attacks against a 512-bit hash function are more expensive than the corresponding brute-force attacks against a 256-bit hash function, more expensive analyses can be considered to be valid attacks, and this sometimes allows an attack to extend another round or two.

4.1.4 Distinguishing Attacks and Differential Properties

Many of the cryptanalytic attacks on the finalists do not directly attack the core hash function properties of collision- and preimage-resistance. Instead, they attempt to find some unexpected property of the underlying primitives (block ciphers and fixed permutations) on which the finalists are based. For example, there may be a predictable relationship between the XOR differences of two input messages and their corresponding message digests – i.e., a differential property. It is not at all clear how relevant these results are in evaluating the security of the finalists. These distinguishing attacks have no practical relevance in attacking the hash functions. However, there is a good reason to give them at least some weight: all five finalists have proofs of security on their domain extenders, showing that if their underlying component functions behave in an ideal way, the full hash function will be secure according to some formal definition. A theoretical distinguishing attack is a way of showing that these components do not behave in an ideal way, and thus potentially undermining these proofs. Table 6 shows the best-known distinguishing attacks and differential-property analyses.

Algorithm	Target	Rounds	Fraction of Target Analyzed	CF Call	Reference
BLAKE-256	Block Cipher	7/14	50 %	2^{232}	[43]
Grøstl-256	Permutation	9/10	90 %	2^{368}	[44]
JH	Compression Function	42/42	100 %	2^{304}	[39]
Keccak	Permutation	24/24	100 %	2^{1579}	[45]
	Permutation	14/24	58 %	$2^{255.77}$	[46]
Skein	Compression Function	37/72	52 %	$2^{511.2}$	[47]
	Compression Function	36/72	50 %	2^{454}	[48]
SHA-2	Hash Function	46/64	72 %	2^{46}	[49]

Table 6. Best Known Distinguishing Attacks and Differential Properties for the Finalists and SHA-2

Note that for BLAKE, Grøstl, and SHA-2, which have both 256- and 512-bit functions, this table applies only to the 256-bit functions.

4.1.5 Depth of Analysis and Understandability of Algorithms

Another aspect of the published security results that NIST considered in the selection of SHA-3 was how well understood the finalists seemed to be, and how understandable they ultimately appear to be.

Evaluating the depth of analysis on each of the finalists is inherently somewhat subjective; different cryptographers may come to different conclusions when reading the literature. However, this evaluation is also necessary, because the security margins discussed above can be misleading if some algorithms have received little or no analysis, while others have been carefully scrutinized for weaknesses. More importantly, choosing a SHA-3 winner involves making a prediction that this winner is very unlikely to suffer some practical attack during its working lifetime. Unfortunately, the state-of-the-art in hash function cryptanalysis and design doesn't allow this prediction to be made with certainty. NIST had to rely on its own evaluation of the best attacks, and also of how much further the best published attacks might be pushed, with further analysis.

It is easy, but rather misleading, to quantify the amount of analysis by counting the number of cryptanalysis papers published on each finalist. Certainly, NIST took notice of how many papers were published on each algorithm, but this metric can be misleading on many levels: finalists with substantial tweaks often had more papers published simply to apply old attacks to the tweaked algorithm; some algorithms may have attracted more relatively small results as opposed to a smaller number of deeper results; and algorithms with more weaknesses may even have attracted more papers attacking them. Instead, NIST considered the depth of cryptanalysis on the different finalists.

The only finalist NIST believed to be seriously called into question by this analysis was JH, which has received much less analysis than the other finalists, and yet has a relatively small security margin. NIST also considered the depth of analysis in looking at security margins; the relatively large security margin of Keccak and the relatively small security margin of Grøstl partly reflect the difference in the intensity of cryptanalysis they received during the competition. Grøstl, Skein, and BLAKE received more analysis than Keccak, probably because of Grøstl's similarity with AES and several other SHA-3 candidates, and Skein's and BLAKE's similarity with the large number of ARX-based hash functions meant that many existing techniques and tools existed for beginning the analysis.

One common bit of feedback that NIST received from the community involved ARX vs. non-ARX designs. Recall that an ARX design gets its security from the combination of addition, rotation, and exclusive OR operations. (Many ARX-based hash functions and ciphers include other elements, such as bitwise logical functions and variable rotates.) BLAKE and Skein are ARX designs, as are the existing SHA-2 hash functions. Some people in the cryptographic community feel that this kind of design is inherently hard to understand, and expressed a preference for non-ARX designs that may ultimately be better understood. It was not clear how much weight to give this feedback, given that at present, all widely used hash functions are ARX designs, and the two ARX-based finalists appear to be more extensively analyzed than Keccak or JH.

4.1.6 Tweak History of the Finalists

In the Federal Register Notice [6] that called for the SHA-3 algorithm submissions, NIST indicated that a submission package could include a tunable security parameter, such as the

number of rounds, which would allow the selection of a range of possible security/performance trade-offs. That parameter, *the number of rounds*, was indeed utilized by all the designing teams of the finalists, and was specified in their submission packages. Prior to the start of the second and third rounds of the competition, the designers of the remaining candidates were also allowed to make minor changes, referred to as “tweaks,” to their designs. Some designers simply tweaked their algorithms by increasing the number of rounds, which does not invalidate previous analysis results; others chose to modify the internal structures of a component function. These tweaks are summarized in Table 7.

Algorithm	Second-Round		Third-Round	
	Names	Tweaks	Names	Tweaks
BLAKE	BLAKE-28 BLAKE-32 BLAKE-48 BLAKE-64	None	BLAKE-224 BLAKE-256 BLAKE-384 BLAKE-512	Number of rounds increased from 10 to 14 for BLAKE-224 and BLAKE-256; and from 14 to 16 for BLAKE-384 and BLAKE-512.
Grøstl	Grøstl-0	None	Grøstl	Shift values in Q changed. Round constants in P and Q changed.
JH	JH	None	JH42	Number of rounds increased from 35.5 to 42.
Keccak	Keccak	Message block size (rate) increased. Number of rounds increased from 18 to 24.	Keccak	Padding rule simplified. Diversifier parameter removed. The restriction on the supported values of r removed.
Skein	Skein	Rotation constants changed.	Skein	Key schedule parity constant changed.

Table 7. Tweak History of the SHA-3 Finalists

Of the finalists, BLAKE and JH were only tweaked to increase the number of rounds. This does not affect any cryptanalysis prior to the tweak and is, therefore, the most innocuous type of tweak.

Keccak changed the block sizes and padding rules, which affects generic security, but left the underlying 1600-bit permutation unaltered. Keccak also increased the number of rounds from 18 to 24. These tweaks are also fairly innocuous. All of the cryptanalysis on Keccak’s 1600-bit permutation remains valid, and the changes in generic security caused by the tweaks are well understood.

Skein and Grøstl were both tweaked in ways that invalidated some cryptanalysis made prior to the final round. In each case, the third-round tweak was targeted at specific cryptanalysis. Skein’s tweak targeted attacks exploiting the rotational symmetry of the key-schedule parity constant, and Grøstl’s tweak targeted attacks exploiting the similarity of the P and Q permutations to one another. In both cases, cryptanalysis similar to the targeted attacks ceased to be a major threat, and attacks that were not targeted by the tweak needed some modification, but were, for the most part, readily adapted to apply to the final version. Skein was also modified in the second round in

a way that could potentially affect analyses done on the first-round version of Skein, but this tweak was not in response to any specific attack.

4.1.7 Side Channel Attacks and Countermeasures

4.1.7.1 Overview of Side Channel Attacks on Hash Functions

Side channel attacks exploit the physical properties of an implementation of a cryptographic algorithm to learn some information about the algorithm's internal state. For example, a given software implementation of a hash function might take slightly more or less time to process an input string, depending on the contents of the string. The side channels commonly considered in the literature are timing, power, and electromagnetic (EM) emissions.

4.1.7.2 Applicability of Side Channel Attacks

Side channel attacks are relevant only in a narrow subset of hashing applications. In many applications, such as digital signatures, the input to the hash function is not a secret, and therefore, side-channels on the hash function are not relevant. Side channel attacks can be used to reveal some secret information that has been processed by a hash function, but cannot be used to violate the core collision- and preimage-resistance properties of a hash function. Side channels are relevant when there is a secret value being processed by the hash function, such as a key-derivation key or the key in a MAC computation, and when the device implementing the hash function can be observed carefully by an attacker. For example, a power-analysis attack is only practical when an attacker can precisely measure the power used during cryptographic computations – such attacks pose a lot more threat to smartcard applications than to server applications.

One finalist, Keccak, has an interesting property that was first noticed internally at NIST and described in [50] (and presented in [51]): because the hash function does not have any non-invertible steps, an attacker who learns the entire intermediate state for any HMAC-Keccak computation can use this to determine the original key used for the HMAC computation, and can forge arbitrary messages for this key. By contrast, the other four finalists and the SHA-2 algorithms are somewhat less vulnerable; determining some arbitrary intermediate state of the HMAC computation can allow some forgeries (extensions of previously observed messages), but does not reveal the HMAC key or allow arbitrary forgeries. This property does not appear to NIST to have a major security impact, however.

4.1.7.3 Published Side Channel Attacks and Analysis on the SHA-3 Finalists

Side channel attacks work by finding some statistical relationship between things that an attacker can measure, such as the power or time required to do some computation, and some secret internal state of the algorithm. In general, a naïve implementation of a cryptographic algorithm is very likely to be vulnerable to a side channel attack, and indeed, side channel attacks were reported on straightforward implementations of all five SHA-3 finalists [52, 53]. Devices that need to resist these attacks introduce countermeasures to limit the amount of information leaked. The countermeasures can exist at various levels, from the design of transistors all the way up to the design of protocols. Some countermeasures are more or less algorithm-independent; others are specific to a particular kind of cryptographic algorithm.

The reported side channel attacks on the SHA-3 finalists used power, EM analysis, and timing side channels. In broad terms, side channel attacks can be targeted at hardware (ASIC or FPGA)

implementations, low-end software implementations (on small microprocessors without memory caches), and high-end software implementations. In general, EM and power analysis are more difficult on hardware implementations than on software implementations [52].

Based on [52, 53, 54, 55, 56], the SHA-3 finalists can be broken into three groups with respect to side channel attacks: the S-box-based design, the logical-operations-based designs, and the ARX-based designs.

Grøstl, which uses the eight-bit S-box from AES, is the only S-box-based design. There are countermeasures against side channel attacks that can be used to protect S-box accesses. Naïve high-end software implementations of Grøstl are vulnerable to cache-timing attacks related to those found against AES [57]. However, the use of the AES-NI instructions eliminates the timing channel, and there are efficient bit- and byte-slice implementations of Grøstl for high-end software [58] that prevent timing attacks, even without the AES-NI instructions. (These implementations would move Grøstl into the logical-operations category.)

JH⁹ and Keccak use logical operations – these are probably the easiest algorithms to protect against side channel attacks, and the Keccak team proposed some techniques for protecting Keccak implementations from side channel attacks in [55].

BLAKE and Skein, like the SHA-2 algorithms, are ARX-based designs, which require 32- or 64-bit integer additions and rotates. These are probably the most expensive designs to protect from side channel analysis.

Unfortunately, the published analyses offered very limited information about the difficulty of implementing each of the SHA-3 finalists in a way that is resistant to side channel attacks.

One general technique for protecting an algorithm from side channel attacks is called *masking* [55], which involves dividing the secret values into “shares” and then operating on the shares independently. The more shares the variables are divided into, the higher the order of differential power analysis¹⁰ (DPA) the algorithm can be protected against. The type of operations that an algorithm requires determines how the shares are derived and recombined. The operations can be divided into two categories: logical and arithmetic. Creating shares for either of the operation types is straightforward; however, when the two operation types are mixed, a conversion must be performed to convert the shares from one representation to the other. Since Keccak and the bit-slice implementation of JH only use logical operations, creating shares and using them is straightforward. However, the ARX-based algorithms (BLAKE and Skein) use a combination of logic and arithmetic operations; these algorithms require a great deal of computation for the masking to protect against SCA, which results in severe performance degradation.

4.2 Finalist Profiles and Cryptanalysis

This section provides a brief description of each finalist, including its basic building blocks, domain extender, and tunable parameters. Each description is followed by a summary of the existing security analysis, including the analyses of any previous versions of the algorithm. The inclusion of these analyses is intended to provide references as to why certain candidates were tweaked in order to address previously unknown vulnerabilities. As much as is feasible, the analyses are presented in chronological order.

⁹ While JH uses two 4x4-bit S-boxes, the suggested bit-sliced implementation of JH uses logical operations to implement these S-boxes rather than a table look-up.

¹⁰ A powerful category of side channel attacks.

4.2.1 BLAKE

BLAKE uses HAIFA [59] as its domain extender. The BLAKE compression function is based on a wide-pipe block cipher using a modified Davies-Meyer construction. The block cipher operates on an inner state that can be represented as a four-by-four matrix of words. BLAKE-224 and BLAKE-256 use a 512-bit-wide block cipher consisting of 32-bit words, while BLAKE-384 and BLAKE-512 use a 1024-bit-wide block cipher consisting of 64-bit words. The internal state is initialized using an Initialization Value (IV), a salt and a counter, and is updated using the G function, which is based on the *ChaCha* stream cipher [60]. The G function updates the columns and the disjoint diagonals of the state using ARX operations. The input-message words and constants are selected using round-dependent fixed permutations. The designers defined four toy versions for analysis purposes, namely: BLOKE, FLAKE, BLAZE, and BRAKE.

To generate 224-, 256-, 384- and 512-bit message digests, four instances were proposed: BLAKE-28, -32, -48 and -64. For the final round of the competition, these functions were renamed to BLAKE-224, -256, -384 and -512; and the tunable parameter of BLAKE – the number of rounds – was increased from 10 to 14 rounds for BLAKE-224 and -256, and from 14 to 16 rounds for BLAKE-384 and -512.

4.2.1.1 Analysis of the BLAKE Domain Extender

Andreeva et al. [61] proved that BLAKE-256 is secure against preimage, second-preimage and collision attacks up to 2^{256} , 2^{256} , and 2^{128} queries, respectively; similarly, BLAKE-512 is secure up to 2^{512} , 2^{512} , and 2^{256} queries against the above attacks if the underlying block cipher is structurally strong. Andreeva et al. [61] and Chang et al. [62] observed that the compression function of BLAKE does not behave ideally, even if its underlying block cipher is ideal. Then, Andreeva et al. [61] and Chang et al. [62] proved that, under the assumption that the block cipher is ideal, BLAKE-256 and BLAKE-512 are indistinguishable from random oracles up to 2^{128} and 2^{256} queries, respectively; this means that they are secure against all nontrivial generic attacks up to 2^{128} and 2^{256} queries, respectively.

4.2.1.2 Pre-Round 3 Analysis

The first public analysis of BLAKE was done by Li and Xu [63]. The authors found a method to control some of the intermediate hash words using message modification techniques, and presented free-start collision and preimage attacks for BLAKE, with the compression function reduced to 2.5 rounds. The time complexities are 2^{96} , 2^{112} , 2^{160} and 2^{224} for the collision attack, and 2^{209} , 2^{241} , 2^{355} and 2^{481} for the preimage attack, for BLAKE-28, -32, -48, and -64, respectively.

Guo and Matusiewicz [64] presented a near-collision attack for the compression function of BLAKE-32, reduced to four middle rounds, starting from round three, with complexity 2^{56} . The G function is linearized by replacing the modular addition with the XOR operation, and then rotation-invariant differences are introduced to the chaining value, the salt, the counter, and the message to obtain 24-bit near collisions.

Aumasson et al. [65] proved that one round of BLAKE is a permutation of the message and presented an efficient inversion algorithm. The authors presented an improved preimage attack on 1.5 rounds, with complexity 2^{128} , and impossible differentials for the permutation with five and

six rounds for BLAKE-32 and -64, respectively. The authors also described near collisions for four of the middle rounds of the compression function of BLAKE-32.

Sönmez Turan and Uyan [66] used a hill-climbing approach that introduced random differences to message blocks and showed practical near collisions with 47 and 72 bits of difference for the compression function of BLAKE-32, reduced to 1.5 and two rounds, respectively, with complexity 2^{26} .

Vidali et al. [67] presented an efficient method for an arbitrary number of collisions for the full-round BLOKE and an internal collision attack on BRAKE.

Su et al. [37] used linear differential techniques to analyze the near-collision resistance of the compression function of BLAKE and showed near-collision attacks on 152, 396, and 306 bits for the compression functions of BLAKE-32, -64 and -64, reduced to four, four, and five middle rounds, with complexity 2^{21} , 2^{16} , and 2^{216} , respectively.

Ming et al. [68] studied the reversibility properties of the internal permutation of BLAKE. The authors presented some differential properties of G and G^{-1} , and showed that BLAKE has strong resistance against differential attacks.

Biryukov et al. [43] applied the boomerang attack to BLAKE-32 and obtained distinguishers on the compression function, reduced to seven rounds, and on the keyed permutation, reduced to eight rounds, with complexity of 2^{232} and 2^{242} , respectively. The authors argued that the attack is applicable to the other versions of BLAKE as well. Leurent [69] noted that the characteristics used for the eight-round attack are incompatible.

Table 8 summarizes the cryptanalysis results on BLAKE before its Round 3 submission. In this table and the summary-of-cryptanalysis tables to follow, a “SFS Collision” denotes a “Semi-Free-Start Collision,” and “CF” stands for “Compression Function.” Attacks with an empty table entry for “Memory” have negligible memory complexity.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Hash function	Preimage	224	2.5 rounds	2^{209}		[63]
Hash function	Preimage	256	2.5 rounds	2^{241}		[63]
Hash function	Preimage	384	2.5 rounds	2^{355}		[63]
Hash function	Preimage	512	2.5 rounds	2^{481}		[63]
Hash function	Collision	224	2.5 rounds	2^{96}		[63]
Hash function	Collision	256	2.5 rounds	2^{112}		[63]
Hash function	Collision	384	2.5 rounds	2^{160}		[63]
Hash function	Collision	512	2.5 rounds	2^{224}		[63]
CF	Near collision	256	4 rounds	2^{56}		[64]
Hash function	Preimage	256	1.5 rounds	2^{128}		[65]
Block Cipher	Impossible differential	224,256	5 rounds			[65]
Block Cipher	Impossible differential	384,512	6 rounds			[65]
CF	SFS near collision	256	2 rounds	example		[66]
Hash function	Collision	all	BLOKE	example		[67]
CF	SFS collision	all	BRAKE	example		[67]
CF	Near collision	256	4 rounds	2^{21}		[37]
CF	Near collision	512	4 rounds	2^{16}		[37]
CF	Near collision	512	5 rounds	2^{216}		[37]
CF	Distinguisher	256	7 rounds	2^{232}		[43]
Block Cipher	Distinguisher	256	8 rounds	2^{242}		[43]

Table 8. Summary of Cryptanalysis Results on BLAKE Prior to its Round 3 Submission

4.2.1.3 Round 3 Analysis

Only a few cryptanalysis papers were published on the final version of BLAKE. However, since the only substantive change made for the third round of the competition was an increase in the number of rounds, previous analyses of BLAKE are still valid.

Tuple cryptanalysis is a variant of structural cryptanalysis that uses ordered multisets. Aumasson et al. [70] presented a distinguisher, using Tuple cryptanalysis, on the four-round permutation of BLAKE-256, with time complexity of 2^{64} .

Dunkelman and Khovratovich [71] used an iterative differential characteristic based on two different characteristics of the internal permutation, and combined it with a rebound process to find colliding pairs in the compression function, reduced to three rounds, with a complexity of 2^{60} . The attack was extended to six rounds, with a complexity of 2^{456} and a memory requirement of 2^{354} .

Table 9 summarizes the cryptanalysis results on BLAKE's Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Block Cipher	Distinguisher	256	4 rounds	2^{64}		[70]
Block Cipher	Distinguisher	256	6 rounds	2^{456}	2^{354}	[71]

Table 9. Summary of Cryptanalysis Results on BLAKE's Round 3 Submission

4.2.2 Grøstl

Grøstl is a wide-pipe Merkle-Damgård (MD) hash algorithm with an output transformation. The compression function is a novel construction, using two fixed $2n$ -bit permutations – denoted by P and Q – in parallel to produce a $2n$ -bit compression function. Intermediate chaining values are 512 bits wide for Grøstl-256, and 1024 bits for Grøstl-512. The output transformation processes the final chaining state and discards half the bits of the result to yield an n -bit message digest. Grøstl-256 has ten rounds, while Grøstl-512 has fourteen rounds.

The permutations, P and Q , are based on the structure of AES, reusing the AES S-box, but expanding the size of the block to 512 bits for Grøstl-224 and Grøstl-256, and to 1024 bits for Grøstl-384 and Grøstl-512 in a straightforward way.

Grøstl was tweaked for the third round of the competition to eliminate the attacks on the original submission.

4.2.2.1 Analysis of the Grøstl Domain Extender

Andreeva et al. [73] showed that, if the underlying permutations P and Q are free from all structural weaknesses and are independent, Grøstl-256 is secure against preimage, second-preimage and collision attacks up to 2^{256} , $2^{256-\log_2 L}$, and 2^{128} queries (where L denotes the length of the message in blocks). Similarly, Grøstl-512 is secure against the above attacks up to 2^{512} , $2^{512-\log_2 L}$, and 2^{256} queries. Under the same assumptions, Grøstl-256 and Grøstl-512 are indistinguishable from random oracles up to 2^{256} and 2^{512} queries, respectively [74]; this guarantees the security of Grøstl-256 and Grøstl-512 against all nontrivial generic attacks up to 2^{256} and 2^{512} queries.

In [75] Kelsey noted the importance of truncation in the output transformation of the Grøstl hash function.

4.2.2.2 Pre-Round 3 Analysis

Prior to the third round, Mendel et al. [76] showed a six-round (out of ten) semi-free-start collision on the Grøstl-256 compression function, with 2^{120} time and 2^{64} memory complexity. Mendel et al. [77] improved this result and showed a six-round semi-free-start collision on the Grøstl-256 compression function, with 2^{64} time and 2^{64} memory complexity; a seven-round distinguisher on the permutation component of the Grøstl-256 compression function, with 2^{55} time complexity; and a seven-round distinguisher on the output transformation of Grøstl-256, with 2^{56} time complexity.

Gilbert and Peyrin [78] provided a seven-round semi-free-start collision on the Grøstl-256 compression function that requires 2^{120} time and 2^{64} memory; and an eight-round distinguisher for the Grøstl-256 compression function that requires 2^{112} time and 2^{64} memory.

Mendel et al. [79] also presented a four-round collision on the Grøstl-256 hash algorithm, requiring 2^{64} time and memory; a seven-round semi-free-start collision on the Grøstl-256 compression function, requiring 2^{120} time and 2^{64} memory; a five-round collision on the Grøstl-512 compression function, requiring 2^{176} time and 2^{64} memory; and a seven-round semi-free-start collision on the Grøstl-512 compression function, requiring 2^{152} time and 2^{64} memory.

Peyrin [80] demonstrated a five-round collision on the Grøstl-256 hash algorithm that requires 2^{79} time and 2^{64} memory; a six-round collision on the Grøstl-512 hash algorithm, requiring 2^{177} time

and 2^{64} memory; a full-round distinguisher on the Grøstl-256 compression function, with 2^{192} time and 2^{64} memory complexity; a full-round distinguisher on the Grøstl-256 component permutations, with 2^{192} time and 2^{64} memory complexity; an eleven-round distinguisher on the Grøstl-512 compression function, with 2^{640} time and 2^{64} memory complexity; and an eleven-round distinguisher on the Grøstl-512 component permutations, with 2^{640} time and 2^{64} memory complexity.

In addition, Ideguchi et al. [81] showed a six-round collision on the Grøstl-256 hash algorithm, requiring 2^{112} time and 2^{32} memory; a seven-round semi-free-start collision on the Grøstl-256 compression function, requiring 2^{80} time and 2^{64} memory; an eight-round semi-free-start collision on the Grøstl-256 compression function, requiring 2^{192} time and 2^{64} memory¹¹; a seven-round distinguisher on the Grøstl-256 component permutations, requiring 2^{19} time, with no significant memory requirement; and an eight-round distinguisher on the Grøstl-256 component permutations, requiring 2^{64} time and 2^{64} memory.

The attacks have improved over time, as the basic techniques (rebound attacks, super-S-box attacks, and differentials considered between the corresponding rounds of the P and Q permutations) are extended and refined. The time complexities of the rebound attacks on Grøstl by Peyrin [80] have been improved by a factor of about 2^{10} by Naya-Plasencia [82]; however, this attack does not apply to the tweaked submission for the final round of the competition.

Sasaki et al. [83] found semi-free-start collisions on Grøstl-512, reduced to seven rounds, with 2^{152} time and 2^{56} memory complexity. They also found a distinguishing attack on the internal permutations of Grøstl-256, reduced to eight rounds, with 2^{48} time and 2^8 memory complexity.

In a presentation [84] at FSE 2011, Boura et al. mentioned that techniques similar to those described in their paper [85] could be used to produce a zero-sum distinguisher with time complexity 2^{509} for the 512-bit permutations of Grøstl-256.

Table 10 summarizes the cryptanalysis results on Grøstl before its Round 3 submission.

¹¹ This is above the collision bound, and therefore, is listed as a distinguisher in Table 10.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
CF	SFS collision	256	6 rounds	2^{120}	2^{64}	[76]
CF	SFS collision	224,256	6 rounds	2^{64}	2^{64}	[77]
Permutation	Distinguisher	224,256	7 rounds	2^{55}		[77]
Output transform	Distinguisher	224,256	7 rounds	2^{56}		[77]
Permutation	Distinguisher	256	7 rounds	2^{56}		[78]
Permutation	Distinguisher	256	8 rounds	2^{112}	2^{64}	[78]
CF	SFS collision	256	7 rounds	2^{120}	2^{64}	[78]
CF	Distinguisher	256	8 rounds	2^{112}	2^{64}	[78]
Hash function	Collision	224,256	4 rounds	2^{64}	2^{64}	[79]
Hash function	Collision	384,512	5 rounds	2^{176}	2^{64}	[79]
CF	SFS collision	256	7 rounds	2^{120}	2^{64}	[79]
CF	SFS collision	384,512	7 rounds	2^{152}	2^{64}	[79]
Permutation	Distinguisher	256	9 rounds	2^{80}	2^{64}	[80]
Permutation	Distinguisher	256	10 rounds	2^{192}	2^{64}	[80]
Permutation	Distinguisher	512	11 rounds	2^{640}	2^{64}	[80]
CF	Distinguisher	256	9 rounds	2^{80}	2^{64}	[80]
CF	Distinguisher	256	10 rounds	2^{192}	2^{64}	[80]
CF	Distinguisher	512	11 rounds	2^{640}	2^{64}	[80]
Hash function	Collision	256	5 rounds	2^{79}	2^{64}	[80]
Hash function	Collision	512	6 rounds	2^{177}	2^{64}	[80]
Hash function	Collision	224	5 rounds	2^{48}	2^{32}	[81]
Hash function	Collision	256	5 rounds	2^{48}	2^{32}	[81]
Hash function	Collision	256	6 rounds	2^{112}	2^{32}	[81]
CF	SFS collision	224,256	7 rounds	2^{80}	2^{64}	[81]
CF	Distinguisher	224,256	8 rounds	2^{192}	2^{64}	[81]
Permutation	Distinguisher	224,256	7 rounds	2^{19}		[81]
Permutation	Distinguisher	224,256	8 rounds	2^{64}	2^{64}	[81]
CF	Distinguisher	256	10 rounds	2^{182}	2^{64}	[82]
Permutation	Distinguisher	256	10 rounds	2^{175}	2^{64}	[82]
CF	Distinguisher	512	11 rounds	2^{630}	2^{64}	[82]
CF	SFS collision	512	7 rounds	2^{152}	2^{56}	[83]
Permutation	distinguisher	256	8 rounds	2^{48}	2^8	[83]
Permutation	Distinguisher	256	10 rounds	2^{509}		[84]

Table 10. Summary of Cryptanalysis Results on Grøstl Prior to its Round 3 Submission

4.2.2.3 Round 3 Analysis

In response to the differential attacks mentioned above [79, 80, 81], Grøstl was tweaked [72] for the final round to introduce more structural differences between the P and Q permutations, and to change the round constants and the shift values in Q . These changes have rendered the previous attacks ineffective.

Schläffer [38] published the first cryptanalysis on the tweaked Grøstl and showed a semi-free-start collision on the Grøstl-256 compression function, reduced to six rounds, with 2^{120} time and 2^{64} memory complexity. This paper also showed a semi-free-start collision on the six-round Grøstl-512, with 2^{180} time and 2^{64} memory complexity. For collisions on the full hash function,

only three rounds have been successfully attacked, with 2^{64} , 2^{64} and 2^{192} time complexities and negligible memory for Grøstl-224, Grøstl-256, and Grøstl-512, respectively.

Using rebound cryptanalysis techniques, Jean et al. [44] discovered differential properties for up to nine rounds of the Grøstl-256 permutation, with 2^{368} time and 2^{64} memory complexity; and up to ten rounds of the Grøstl-512 permutation, with 2^{392} time and 2^{64} memory complexity.

Khovratovich [86] studied the preimage resistance of the output transformation function of Grøstl-256, and found a *biclique attack* [47] for up to six rounds, with 2^{251} time complexity.

Wu et al. [87] also published pseudo-preimage attacks on up to five rounds of Grøstl-256, with 2^{245} time and 2^{230} memory complexity, and up to eight rounds of Grøstl-512, with 2^{507} time and 2^{507} memory complexity.

Table 11 summarizes the cryptanalysis results on Grøstl’s final-round submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Hash function	Collision	224,256	3 rounds	2^{64}		[38]
Hash function	Collision	512	3 rounds	2^{192}		[38]
CF	SFS collision	256	6 rounds	2^{120}	2^{64}	[38]
CF	SFS collision	384	6 rounds	2^{180}	2^{64}	[38]
CF	SFS collision	512	6 rounds	2^{180}	2^{64}	[38]
Permutation	Differential Property	256	9 rounds	2^{368}	2^{64}	[44]
Permutation	Differential property	512	8 rounds	2^{280}	2^{64}	[44]
Permutation	Differential property	512	9 rounds	2^{328}	2^{64}	[44]
Permutation	Differential property	512	10 rounds	2^{392}	2^{64}	[44]
Output transform	Preimage	256	6 rounds	2^{251}		[86]
Output transform	Preimage	256	5 rounds	2^{206}	2^{48}	[87]
Output transform	Preimage	512	8 rounds	2^{495}	2^{16}	[87]
Hash function	Pseudo-preimage	256	5 rounds	2^{245}	2^{230}	[87]
Hash function	Pseudo-preimage	512	8 rounds	2^{507}	2^{507}	[87]

Table 11. Summary of Cryptanalysis Results on Grøstl’s Final-Round Submission

4.2.3 JH

The JH family of hash algorithms for different hash sizes is based on a single compression function, F_8 , which uses a fixed 1024-bit permutation, E_8 . In F_8 , a message block of 512 bits is XORed with the first 512 bits of the input of E_8 , and the last 512 bits of the output of E_8 . The permutation E_8 uses two 4x4-bit S-boxes, S_0 and S_1 , an eight-bit linear permutation L , and a permutation P_8 . Different members of the JH family are distinguished by the different IVs used, and the message digests are obtained by truncating the final output to the desired hash sizes. JH was tweaked after the second round by increasing the number of rounds from 35.5 to 42.

4.2.3.1 Analysis of the JH Domain Extender

Several papers have analyzed the domain extender of JH. Bhattacharyya et al. [88], and Mendel and Thomsen [89] analyzed the preimage resistance of the JH domain extender from an information-theoretic point of view; however, the actual complexities of the attacks are no better than a brute-force search [90]. Bagheri [91] showed that a pseudo-collision of JH can easily be found, since JH’s compression function is invertible; however, since the IV of JH is a fixed constant that is determined by the message-digest size, the pseudo-collision attack is not considered to be a serious threat to JH.

Bhattacharyya et al. [88] also showed that the indistinguishability security of the JH domain extender is guaranteed up to approximately 172 bits.

Lee and Hong [92] showed that the JH domain extender achieved optimal collision resistance (256 bits for the 512-bit output size, and 192, 128, and 112 bits for the other three sizes).

Andreeva et al. [73] have improved the second-preimage resistance of the JH domain extender from 172 bits to approximately 256 bits.

Moody et al. [93] improved the indistinguishability security bound of the JH domain extender from 172 bits to approximately 256 bits.

4.2.3.2 Pre-Round 3 Analysis

Turan and Uyan [66] used hill-climbing techniques to produce a semi-free-start near collision on ten rounds of JH’s compression function, with $2^{23.24}$ time complexity. Rijmen et al. [94] proposed a semi-free-start collision attack for sixteen rounds of the JH hash function, with $2^{178.24}$ time and $2^{101.12}$ memory complexity, and a semi-free-start near-collision attack for twenty-two rounds of JH’s compression function, with $2^{156.56}$ time and $2^{143.70}$ memory complexity. Naya-Plasencia [82] proposed a semi-free-start collision attack for sixteen rounds of JH’s compression function, with $2^{96.12}$ time and $2^{96.12}$ memory complexity, and a semi-free-start near-collision attack for twenty-two rounds of JH’s compression function, with $2^{95.63}$ time and $2^{95.63}$ memory complexity.

Table 12 summarizes the cryptanalysis results on JH before its Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
CF	SFS near collision	all	10 rounds	$2^{23.24}$		[66]
Hash function	SFS collision	256	16 rounds	$2^{178.24}$	$2^{101.12}$	[94]
CF	SFS near collision	256	19 rounds	$2^{156.77}$	$2^{143.70}$	[94]
CF	SFS near collision	256	22 rounds	$2^{156.56}$	$2^{143.70}$	[94]
CF	SFS collision	256	16 rounds	$2^{96.12}$	$2^{96.12}$	[82]
CF	SFS near collision	256	22 rounds	$2^{95.63}$	$2^{95.63}$	[82]

Table 12. Summary of Cryptanalysis Results on JH Prior to its Round 3 Submission

4.2.3.3 Round 3 Analysis

Naya-Plasencia et al. [39] extended their previous analysis to produce semi-free-start internal near collisions on 37 rounds of the JH compression function, with 2^{352} time and $2^{57.6}$ memory

complexity, and used similar techniques to find differential characteristics for all 42 rounds of the E_8 permutation, with 2^{304} time and $2^{57.6}$ memory complexity.

Table 13 summarizes the cryptanalysis results on JH’s Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Permutation	Differential Property	all	42 rounds	2^{304}	$2^{57.6}$	[39]
Permutation	Differential Property	all	42 rounds	2^{352}	$2^{57.6}$	[39]
CF	SFS near collision	all	37 rounds	2^{352}	$2^{57.6}$	[39]
CF	SFS near collision	all	26 rounds	2^{112}	$2^{57.6}$	[39]

Table 13. Summary of Cryptanalysis Results on JH’s Round 3 Submission

4.2.4 Keccak

Keccak follows the sponge construction model [95], with a 1600-bit permutation. The message block size varies according to the output size: Keccak-512 has a block size of 576 bits, Keccak-384 has 832 bits, Keccak-256 has 1088 bits, and Keccak-224 has 1152 bits. In keeping with the standard terminology for a sponge construction, the message block size is referred to as r , for rate, and the difference between the permutation size and the message block size is referred to as c , for capacity.

The Keccak team also defined a number of additional variants of Keccak. First, they defined a one-size-fits-all variant “Keccak,” with rate 1024 and capacity 576, that they believe can be used, in practice, for any output size. The Keccak team does not consider this the official submitted version for producing 512-bit and 384-bit outputs, since it does not provide the expected preimage-resistance. However, the Keccak team argues that all the attacks that arise from the reduced capacity have an unrealistically high complexity of at least 2^{288} operations, and “Keccak” with an output of 384 or 512 bits would be significantly faster than the submitted Keccak-384 and Keccak-512. The Keccak team also produced scaled-down versions of the Keccak permutation, with 25, 50, 100, 200, 400, and 800 bits, instead of 1600, and set up a cryptanalysis contest [96] challenging cryptanalysts to attack versions of Keccak with the capacity set to as few as 160 bits.

The permutation can be considered as a substitution-permutation network with five-bit-wide S-boxes, or as a combination of a linear mixing operation and a very simple nonlinear mixing operation.

The recommended security parameter for Keccak is twenty-four rounds, which was tweaked, in the second round, from the eighteen rounds specified in the original submission. Additionally, the message block sizes of the 224-bit and 256-bit versions were increased from 1024 bits to 1152 and 1088 bits, respectively, as part of the second-round tweak. Keccak was also tweaked in the third round, by simplifying the padding rule.

4.2.4.1 Analysis of the Keccak Domain Extender

Bertoni et al. [97] proved that Keccak- n is secure against preimage, second-preimage, and collision attacks up to 2^n , 2^n , and $2^{n/2}$ queries if the underlying permutation is ideal. Bertoni et al. [97] proved that, under the same assumption, Keccak- n is indifferentiable from a random oracle up to 2^n queries, which means that Keccak- n is secure against all nontrivial generic attacks up to 2^n queries.

Gligoroski et al. [98] observed that if a particular second-preimage is found (most likely requiring a one-time effort equivalent to brute-force), then an arbitrary number of second-preimages can be found very easily for any message.

4.2.4.2 Pre-Round 3 Analyses

Aumasson and Khovratovich [99] attempted algebraic attacks against significantly reduced-round versions of Keccak by applying automatic cryptanalysis tools (a triangulation tool, and cube testers) to Keccak’s permutation. The analysis was limited to three- and four-round versions of Keccak, but the authors conjectured that structures may be observed in the permutation for up to ten rounds. The authors of [99] concluded that Keccak has a comfortable security margin to prevent structural distinguishers.

Morawiecki and Srebrny [100] used SAT-solver techniques to find preimages for three rounds of Keccak, with 40 unknown message bits.

Aumasson and Meier [46] first showed practical zero-sum distinguishers on nine rounds of the Keccak permutation, and impractical distinguishers on a sixteen-round variant. Boura and Canteaut [101] later extended this to eighteen rounds. However, at a complexity of 2^{1370} , this attack is only of theoretical interest. The Keccak team increased the number of rounds to twenty-four after the initial distinguisher was announced. Responses to [46] and [101] from the Keccak team were provided in [102, 103]. Very high-complexity distinguishers continued to be published: Boura and Canteaut [104] attacked 20 rounds, with a complexity of 2^{1586} . Boura et al. [85] attacked all 24 rounds, with a complexity of 2^{1590} .

Lathrop [105] was able to use a cube attack to conduct a key-recovery attack on a secret prefix Message Authentication Code (MAC), with a four-round Keccak variant.

Table 14 summarizes the cryptanalysis results on Keccak before its Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Reference
Permutation	Observations	all			[99]
Hash function	Preimage	1024	3 rounds	2^{34}	[100]
Permutation	Distinguisher	all	9 rounds	$2^{29.83}$	[46]
Permutation	Distinguisher	all	14 rounds	$2^{255.77}$	[46]
Permutation	Distinguisher	all	16 rounds	$2^{1023.88}$	[46]
Permutation	Distinguisher	all	18 rounds	2^{1370}	[101]
Permutation	Distinguisher	all	20 rounds	2^{1586}	[104]
Permutation	Distinguisher	all	24 rounds	2^{1590}	[85]
Secret-prefix MAC	Key recovery	224	4 rounds	2^{19}	[105]

Table 14. Summary of Cryptanalysis Results on Keccak Prior to its Round 3 Submission

4.2.4.3 Round 3 Analyses

Morawiecki [106] used SAT-solver techniques to find preimages for two rounds of Keccak, with 80 unknown message bits, and to find collisions on two rounds of Keccak, with a 160-bit output. Both were contest versions, with capacity reduced to 160 bits. Morawiecki demonstrated the low attack complexity by actually performing the attack, rather than by computing a theoretical complexity.

Duan and Lai [45] improved the time complexity of the 24-round zero-sum distinguisher attack to 2^{1579} .

Duc et al. [107] applied rebound techniques to find a rebound differential property for eight rounds of the Keccak permutation, with a time complexity of $2^{491.47}$.

Bernstein [108] used algebraic techniques to speed up a brute-force second-preimage search for up to eight rounds of Keccak.

Naya-Plasencia et al. [109] presented a preimage attack on two rounds of Keccak-224 and Keccak-256, a collision attack on two rounds, and a near-collision attack on three rounds of Keccak-224 and Keccak-256.

Dinur et al. [40] produced example collisions for Keccak-224 and Keccak-256, reduced to four rounds, and near collisions on five rounds.

Table 15 summarizes the cryptanalysis results on Keccak's Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Hash function	Collision	160	$r=\{240,640,1440\}$, $c=160$, 1, 2 rounds	Example		[106]
Hash function	Preimage	80	$r=\{240,640,1440\}$, $c=160$, 1, 2 rounds	Example		[106]
Permutation	Distinguisher	all	24 rounds	2^{1579}		[45]
Permutation	Differential Property	all	8 rounds	$2^{491.47}$		[107]
Hash function	Distinguisher	224,256	4 rounds	2^{25}		[109]
Hash function	Near collision	224,256	3 rounds	2^{25}		[109]
Hash function	Collision	224,256	2 rounds	2^{33}		[109]
Hash function	2 nd Preimage	224,256	2 rounds	2^{33}	2^{29}	[109]
Hash function	2 nd Preimage	512	6 rounds	2^{506}	2^{176}	[108]
Hash function	2 nd Preimage	512	7 rounds	2^{507}	2^{320}	[108]
Hash function	2 nd Preimage	512	8 rounds	$2^{511.5}$	2^{508}	[108]
Hash function	Collision	224,256	4 rounds	Example		[40]
Hash function	Near collision	224,256	5 rounds	Example		[40]

Table 15. Summary of Cryptanalysis Results on Keccak's Round 3 Submission

4.2.5 Skein

Skein is an iterative hash algorithm that is built on a tweakable block cipher – Threefish. Threefish is used to build the compression function of Skein using a modified Matyas-Meyer-Oseas construction, which is then iterated using a domain extender similar to the HAIFA domain extender used by BLAKE. The designers refer to the whole construction as a Unique Block Iteration (UBI). The internal structure of Threefish is a 72-round substitution-permutation network using a 128-bit MIX function consisting of 64-bit addition, rotate and XOR operations, each used exactly once per MIX function.

For the second round, Skein was tweaked by modifying the rotation constants, which the submitters felt had not been fully optimized for the first round of the competition. For the final round of the competition, Skein was tweaked by modifying a constant in the key schedule of Threefish to harden the function against rotational cryptanalysis.

4.2.5.1 Analysis of the Skein Domain Extender

Bellare et al. [110] proved that Skein- n is secure against preimage and collision attacks up to approximately 2^n and $2^{n/2}$ queries, respectively, if the underlying tweakable block cipher (i.e., Threefish) is ideal. They also proved that, under the same assumption, Skein is indistinguishable from a random oracle up to 2^{256} queries, which means that Skein is secure against all nontrivial generic attacks up to 2^{256} queries. Andreeva et al. [73] proved that Skein-512 is second-preimage-resistant up to 2^{512} queries, under the same assumption.

4.2.5.2 Pre-Round 3 Analysis

Early attacks on Threefish included a 33-round key-recovery attack by Chen and Jia [111], and a 35-round distinguisher by Aumasson et al. [112]. McKay and Vora [113] published an attack on Threefish, using pseudo-linear functions to analyze the block cipher, yielding a 15-round key-recovery attack. In addition, linear differential analysis by Su et al. [37] produced near collisions on 24 rounds of the Skein compression function.

The most severe attacks on earlier versions of Skein exploited the near-rotational symmetries of its MIX function (where only the lack of a carry bit at either end of the 64-bit adder breaks the symmetry). The first of these attacks was the related-key, key-recovery attack of Khovratovich and Nikolic [114] on 39 rounds of Threefish-256, and 42 rounds of Threefish-512. Later, Khovratovich et al. [115] provided a rotational distinguisher, based on the concept of rotational collisions. By using what the authors described as a rebound attack, they extended the distinguisher to 53 and 57 rounds of the Skein-256 and Skein-512 compression functions, respectively. However, the Round 3 tweak, changing the key schedule constant from 0x5555555555555555 to 0x1BD11BDAA9FC1A22, renders the rotational properties of the MIX function far less useful to the attacker. No comparable rotational attacks have been proposed on the final round version of Skein.

Table 16 summarizes the cryptanalysis results on Skein before its Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Block cipher	Key recovery	512	32 rounds	2^{195}	2^{12}	[111]
Block cipher	Key recovery	512	33 rounds	$2^{324.6}$		[111]
Block cipher	Key recovery	512	34 rounds	$2^{474.4}$		[111]
CF	Near collision	512	17 rounds	2^{24}		[112]
Block cipher	Impossible differential	512	21 rounds	-		[112]
Block cipher	Key recovery	512	26 rounds	$2^{507.8}$		[112]
Block cipher	Key recovery	512	32 rounds	2^{312}	2^{71}	[112]
Block cipher	Distinguisher	512	35 rounds	2^{478}		[112]
Block cipher	Key recovery	256	15 rounds	2^{232}		[113]
CF	Near collision	256	24 rounds	2^{60}		[37]
CF	Near collision	512	24 rounds	2^{230}		[37]
CF	Near collision	1024	24 rounds	2^{395}		[37]
Block cipher	Key recovery	256	39 rounds	$2^{254.1}$		[114]
Block cipher	Key recovery	512	42 rounds	2^{507}		[114]
CF	Distinguisher	256	53 rounds	2^{251}		[115]
CF	Distinguisher	512	57 rounds	2^{503}		[115]

Table 16. Summary of Cryptanalysis Results on Skein prior to its Round 3 Submission

4.2.5.3 Round 3 Analysis

Several cryptanalysis papers have been published regarding the Round 3 version of Skein. Khovratovich et al. [47] published attacks using a splice-and-cut attack, augmented by a differential structure called a *biclique*. This resulted in a 22-round preimage attack on Skein-512, requiring 2^{511} compression-function calls, and a 37-round pseudo-preimage attack, requiring $2^{511.2}$ compression-function calls. Khovratovich [86] also used bicliques to find collisions on 14 rounds of Skein-512, and 12 rounds of Skein-256, using $2^{254.5}$ and $2^{126.5}$ compression-function calls, respectively; and preimages on 14 rounds of Skein-256, using $2^{251.4}$ compression-function calls. In another attack paper, Yu et al. [41] used differential cryptanalysis to find semi-free-start near collisions on 32 rounds of Skein-256, using 2^{105} compression-function calls. Yu et al. [48] also published a 36-round distinguisher on the Skein-512 compression function, using 2^{454} compression-function calls, and a 34-round key-recovery attack on Threefish-512, using 2^{424} compression-function calls.

Aumasson et al. [70] used Tuple cryptanalysis, which was also applied to BLAKE, to find a distinguisher on a 17-round version of Threefish.

Table 17 summarizes the cryptanalysis results on Skein's Round 3 submission.

Target	Attack type	Output	Variant	CF Call	Memory	Reference
Hash function	Preimage	512	22 rounds	2^{511}	2^6	[47]
Hash function	Pseudo-preimage	512	22 rounds	2^{508}	2^6	[47]
Hash function	Pseudo-preimage	512	37 rounds	$2^{511.2}$	2^{64}	[47]
Hash function	Collision	512	14 rounds	$2^{254.5}$		[86]
Hash function	Collision	256	12 rounds	$2^{126.5}$		[86]
Hash Function	Preimage	256	14 rounds	$2^{251.4}$		[86]
CF	SFS near collision	256	32 rounds	2^{105}		[41]
CF	Distinguisher	512	32 rounds	$2^{104.5}$		[48]
CF	Distinguisher	512	36 rounds	2^{454}		[48]
Block Cipher	Key recovery	512	34 rounds	2^{424}		[48]
Block Cipher	Distinguisher	512	17 rounds	2^{64}	2^{64}	[70]

Table 17. Summary of Cryptanalysis Results on Skein's Round 3 Submission

4.3 Security Summary

No finalist was considered to be a clear winner or loser from the security analysis. While the details of the security proofs on domain extenders differ somewhat between finalists, all have proofs that relate the security of their full hash functions to the security of their underlying compression functions and fixed permutations. None of the published cryptanalytic attacks come anywhere close to threatening the security of any of the five finalists. The published distinguishing attacks are interesting, but do not have obvious implications for the security of the finalists.

However, there are important differences between the finalists. JH and Grøstl have relatively small security margins that, when combined with their performance in software, probably cannot be addressed by increasing their numbers of rounds. Skein's security margin is a little larger, and could be further expanded by adding rounds, if necessary, given its excellent software performance. BLAKE and Keccak have very large security margins. Furthermore, JH and Grøstl have distinguishing attacks against them that reflect this difference – the best distinguishing attack on the JH permutation distinguishes all 42 rounds, and the best attack, on the underlying permutations of Grøstl, distinguishes 90 % of its rounds. By contrast, the only distinguishing attack on more than 14 of the 24 rounds of Keccak has a complexity that is far beyond any security bound expected of a 512-bit or shorter hash function, and Skein and BLAKE have no known distinguishing attacks that come close to threatening their full-round versions.

Grøstl, Skein, and BLAKE have a large number of attack papers reflecting considerable depth of analysis, and Keccak has somewhat less, but still received considerable cryptanalytic attention. JH received significantly less analysis, and still has a relatively small security margin, suggesting that more analysis might further erode its security margin. Based on the available literature, JH does not appear to be very well-understood yet, and this influenced NIST's selection of a winner.

Another area where there has been some preliminary analysis of the security of the candidates is their resistance to side channel attacks. This area is somewhat difficult to analyze, since these attacks depend on the implementation of the algorithm, and not just the algorithm itself. As such, any algorithm can be implemented to be secure against side channel analysis if the implementer is willing to pay a large enough performance penalty. The challenges involved in defending against side channel analysis, however, vary among the finalists. Grøstl may be difficult to defend against cache-timing attacks on software platforms that do not have native support for the AES S-box in hardware. While the other finalists are unlikely to be vulnerable to timing attacks, they may still

need countermeasures to defend against power-analysis attacks, and in this respect, BLAKE and Skein may suffer somewhat, since it is somewhat harder to mask the modular addition instructions used in those algorithms.

5. Performance Comparison of the SHA-3 Finalists

This section discusses the performance of the finalist candidates when implemented in software and hardware.

All of the SHA-3 finalist candidates, as well as SHA-2, have four variants with 224-, 256-, 384- and 512-bit message-digest outputs. Skein and JH generate all four message-digest sizes with the same compression function, and therefore, all four run at about the same rate¹². Keccak also uses a single compression function to generate all four output sizes, but the bigger the output, the smaller the message block that is processed by each compression-function call, and therefore, the bigger the message digest, the slower Keccak runs. BLAKE, Grøstl and SHA-2 use two different compression functions, one for 384- and 512-bit message digests, and another for 224- and 256-bit message digests; these two hash-function variants usually run at different speeds on the same platform.

Therefore, in the performance discussions, these algorithms will be referred to as:

- BLAKE-256 and BLAKE-512
- Grøstl-256 and Grøstl-512
- JH
- Keccak-224, Keccak-256, Keccak-384 and Keccak-512
- Skein
- SHA-256 and SHA-512

In the performance discussion, whenever the name of an algorithm is used without a specific digest size attached, then the statement applies for all four digest sizes.

5.1 Software Performance

This section compares the performance of the SHA-3 candidate algorithms and SHA-2 when implemented in software on many of the most widely used current computing platforms. The general categories of computers considered are discussed below.

5.1.1 Computer Systems – the Current Playing Field

The software performance of the SHA-3 finalists has been measured in the following classes of commercially available computers:

AMD64¹³: These machines are general-purpose Complex Instruction Set Computers (CISC) with a 64-bit-word orientation and Single Instruction Multiple Data (SIMD) vector units. They all run the 64-bit-oriented AMD64 instruction set architecture (ISA). AMD64 machines predominate today in desktop, laptop and netbook computers, as well as servers of all sorts. They typically have very large RAM memories, and increasingly have two or more independent “cores,” each capable of independently executing a program thread. Each core typically has superscalar, pipelined execution units supporting 64-bit-word integer, logic and floating-point operations and is able to simultaneously launch two instructions at each clock cycle. Typically, there is also a

¹² This is perhaps not as good an assumption for hardware as for software, because different digest sizes have an effect on the area of the circuit, and possibly the clock rate as well, but the effect is usually small.

¹³ This report adopts the “AMD64” nomenclature of the eBASH website. Other common designations for this instruction-set architecture are “x86-64” and “Intel-64”. This choice of names is for ease of comparison with the wealth of data on the eBASH website and does not indicate any NIST endorsement of any vendor or product for any purpose.

vector unit in each core. Since vector units were introduced for the x86 architecture in the 1990s, they have gone in stages from eight 64-bit registers to sixteen 256-bit registers for the latest-announced products. The most recent machines have vector units that can simultaneously execute eight 32-bit operations or four 64-bit operations in a 256-bit vector register. Most desktop, laptop and server computers sold today are AMD64 machines.

X86: These are the 32-bit predecessors of the AMD64 computers. Many legacy systems run on X86 computers, which now usually are AMD64 machines operating in the X86 mode. Most fairly recent examples include a vector unit and super-scalar execution units.

ARM-NEON¹⁴: These machines run a relatively high-end implementation of the ARM ISA with a vector unit. The NEON vector instruction set uses registers that can be viewed as thirty-two 64-bit registers or sixteen 128-bit registers. Many of the SHA-3 finalists benefit significantly from 64-bit instructions or bit-slice implementations on wider words, and run markedly faster on the NEON-equipped ARM machines.

32-bit RISC: These Reduced Instruction Set Computers (RISC) are scalar or super-scalar machines that are typically used today in a wide range of applications from smart phones, tablet computers, and appliances, such as GPS units and music players, to controllers and sensors embedded in many products. By far the most widely used RISC ISA is the ARM, which is widely licensed in a variety of “cores” that are incorporated in ASIC. Other legacy 32-bit RISC ISAs studied in the SHA-3 competition include the PowerPC (PPC), and MIPS (Microprocessor without Interlocked Pipeline Stages) ISAs.

Embedded Microcontrollers: These are small computers, typically included on an ASIC with memory and other application-specific logic. In this category, the primary constraint is usually the RAM memory, although power may be another constraint. This is the most diverse category, and there are a number of ISAs, including low-end implementations of the 32-bit ARM, as well as sixteen and eight-bit microcontrollers. Applications for such computers include smart cards, sensors, smart meters, servo controllers, some RFID tags and a plethora of potentially networked appliances.

5.1.2 Candidate Software Performance Studies

Several studies have been done that compare SHA-2 and the SHA-3 finalists using programs written by a single programmer or a small team for a specific platform or language, or with a specific design goal. These include studies of optimized Java code running on a current AMD64 computer [116], and optimized assembler code on the ARM11 processor [117], which is widely used in smart phones and tablet computers, but does not include the NEON vector engine. However, the vast bulk of the available SHA-3 finalist performance data was provided by two cooperative projects, eBASH (ECRYPT Benchmarking of All Submitted Hashes) [25], and XBX (eXternal Benchmarking eXtension) [26], both are part of the ECRYPT Benchmarking of Cryptographic System (eBACS) [118]. These two projects consolidated algorithm implementations from many different sources, ran them on many different computers and reported on their performance. NIST appreciates and is grateful for the significant efforts of all involved.

¹⁴ NEON is a 128-bit SIMD architecture extension for the ARM Cortex™-A series processors.

5.1.2.1 eBASH: General-Purpose Computers

During the course of the SHA-3 competition, a large number of hash functions were benchmarked on general-purpose computers, and a significant amount of data on all of them can be found on the eBASH site (<http://bench.cr.yp.to/ebash.html>). Many C-language and some assembly implementations of all the finalists were submitted by different programmers to eBASH, then compiled with a range of different options and run on a large number of different computers. The best comparative presentation of the data for the SHA-3 finalists and SHA-2 is the “shootout” graphs [119] found at: <http://bench.cr.yp.to/results-sha3.html>, and NIST made extensive use of this data, which includes six graphs that summarize algorithm performance for various message lengths. In the shootout presentation, only the program with the best performance for each SHA-3 finalist on each computer is represented in the summary graphs. These summaries include data from fifteen AMD64 processor models, two X86 processor implementations, the PPC G4, four different ARM core designs and one MIPS32 implementation. In many cases, the eBASH benchmarks were run on several different computers with the same general processor model.

Figures A1-A6 in Appendix A show the six eBASH summary graphs, as of September 10, 2012. Only the 512-bit and 256-bit variants are plotted. Throughput is stated in machine cycles-per-byte, where fewer cycles indicate better performance. The fastest performance of the best algorithms on the latest machines is about six cycles per byte.

One other feature of the eBASH website deserves special attention here. At <http://bench.cr.yp.to/primitives-sha3.html>, there is a table labeled “Which hash functions are measured? (SHA-2/SHA-3 excerpt),” which is a table of all the variants of all the SHA-3 finalists, plus SHA-512 and SHA-256, and three “tree mode” implementations of BLAKE and Keccak that implement *i*-thread parallel implementations of the hash algorithms. The implementations that depend heavily on vector units are often given names that identify the type of vector unit: SSE (for Streaming SIMD Extensions), AVX (for Advanced Vector eXtensions), MMX¹⁵ for AMD64 or x86 machines, or NEON for ARM machines.

The computers used in eBASH are classified into four of the five groups described above:

- AMD64: use the AMD64 ISA and generally include a vector unit.
- X86: use the 32-bit X86 ISA and may include a vector unit.
- ARM-NEON: use the 32-bit ARM ISA with the NEON vector unit.
- 32-bit RISC: use the following 32-bit RISC ISAs: ARM, MIPS or PPC. A vector unit is not used.

To try to visually distill the complex graphs of the eBASH website into a simpler presentation, the performance of each algorithm is categorized into *high*, *medium* or *low* levels (shown as green, aqua, and red, respectively) for the four different classes of computers in Tables 18 and 19. A candidate’s performance on these platforms is marked at the specific cell with the platform identified, such that a reader can quickly grasp the candidate’s performance by reading across the algorithm row.

Table 18 is for long messages (greater than 4096 bytes), while Table 19 illustrates shorter messages of 64 bytes. These are based on the eBASH “shootout” graphs for six different sizes of messages (long, 4096-bytes, 1536-bytes, 576-bytes, 64-bytes and 8-bytes). These simpler presentations are essentially “eyeball” judgments of relatively noisy data; the reader can study

¹⁵ MMX is a single instruction, multiple data (SIMD) instruction set designed by Intel and introduced in 1996.

Figures A-1 through A-6 in Appendix A (the actual detailed eBASH shootout plots) to see how well the categorizations fit the data. The data for JH, in particular, seems erratic; this is apparently because the fastest JH vector code was completed very late, and had been run on some, but not all, the benchmark machines. This, however, makes little comparative difference, because even the fastest JH implementations are among the slowest of the algorithms. For each class of machine, the performance of the fastest algorithms ranged between four and eight times the speed of the slowest. When data points seem spread across both high and medium, or medium and low, crosshatching is used to draw attention to this.

5.1.2.1.1 Long Messages

Platform Algorithm	High Performance				Medium Performance				Low Performance			
	AMD64	X86	ARM-NEON	32-bit RISC	AMD64	X86	ARM-NEON	32-bit RISC	AMD64	X86	ARM-NEON	32-bit RISC
BLAKE-512	AMD64		ARM-NEON			X86		32-bit RISC				
BLAKE-256	AMD64	X86	ARM-NEON	32-bit RISC	AMD64							
Grøstl-512					AMD64 (AES-NI)				AMD64	X86	ARM-NEON	32-bit RISC
Grøstl-256					AMD64 (AES-NI)				AMD64	X86	ARM-NEON	32-bit RISC
JH					AMD64	X86	ARM-NEON		AMD64	X86		32-bit RISC
Keccak-512									AMD64	X86	ARM-NEON	32-bit RISC
Keccak-256					AMD64	X86	ARM-NEON	32-bit RISC		X86		
Skein	AMD64	X86	ARM-NEON			X86		32-bit RISC				
SHA-512					AMD64	X86		32-bit RISC			ARM-NEON	
SHA-256		X86	ARM-NEON	32-bit RISC	AMD64							

	High Performance
	Medium Performance
	Low Performance
	Performance straddling two categories

Table 18. eBASH Performance Comparison for Long (> 4096-byte) Messages

AMD64: Skein (all sizes) and BLAKE-512 are consistently the fastest algorithms and the only two algorithms that generally are faster than SHA-512 on AMD64 platforms. BLAKE-256 is also fast on the newer AMD64 platforms of one vendor with larger vector-register files [120], but is slower on most of the machines of the other major vendor. This may reflect differences in the vector units of the two vendors. Keccak-256, SHA-512 and SHA-256 are in the medium-level group, averaging about half the speed of the high-level group, along with Grøstl-512 and Grøstl-256 on very new machines with the AES New Instructions (AES-NI). JH is also in the medium group for very recent AMD64 processors with 256-bit vector registers, but not for older

processors with 128-bit vector registers. Grøstl-512, Grøstl-256 (no AES-NI), JH and Keccak-512 are in the low group, with performance only around ¼ of that of the high-performance group.

X86: The high group includes BLAKE-256, and SHA-256, all of which are ARX-type 32-bit-word algorithms. Skein does fairly well, straddling the high-medium groups, despite using 64-bit operations on a 32-bit machine. BLAKE-512, JH, Keccak-256 and SHA-512 have medium performance. The low group consists of Grøstl and Keccak-512.

ARM - NEON: The high group consists of BLAKE, Skein and SHA-256, the medium group contains JH and Keccak-256, and the low group consists of Grøstl, Keccak-512 and SHA-512. The ability of the NEON vector unit to perform 64-bit operations may help Skein and BLAKE-512 on this processor.

32-bit RISC: The high-performance algorithms are BLAKE-256 and SHA-256; the medium group contains BLAKE-512, Keccak-256, Skein and SHA-512. Grøstl, JH and Keccak-512 are in the low group.

5.1.2.1.2 64-byte Messages

Platform Algorithm	High Performance				Medium Performance				Low Performance			
	AMD64	X86	ARM-NEON	32-bit RISC	AMD64	X86	ARM-NEON	32-bit RISC	AMD64	X86	ARM-NEON	32-bit RISC
BLAKE-512	AMD64					X86	ARM-NEON	32-bit RISC				
BLAKE-256	AMD64	X86	ARM-NEON	32-bit RISC	AMD64							
Grøstl-512									AMD64	X86	ARM-NEON	32-bit RISC
Grøstl-256					AMD64 (AES NI)				AMD64	X86	ARM-NEON	32-bit RISC
JH					AMD64	X86	ARM-NEON				ARM-NEON	32-bit RISC
Keccak-512					AMD64		ARM-NEON	32-bit RISC		X86		
Keccak-256					AMD64	X86	ARM-NEON	32-bit RISC				
Skein	AMD64	X86					ARM-NEON	32-bit RISC				
SHA-512					AMD64		ARM-NEON	32-bit RISC		X86		
SHA-256		X86	ARM-NEON	32-bit RISC	AMD64							

	High Performance
	Medium Performance
	Low Performance
	Performance straddling two categories

Table 19. eBASH Performance Comparison for 64-byte Messages

This comparatively small message size starts to show the effects of both fixed-per-message overhead and of different input block sizes, which range from 512 to 1088 bits. Average cycles-

per-byte seem to be about twice that of long messages, but the ordering of algorithms has changed remarkably little. At this message size, Keccak-512 looks about as good in comparison to others as it ever will, since 64 bytes just fits, with padding, in a single message block and Keccak has no extra finalization round, while other algorithms either need two blocks, process a significantly larger-than-needed 128-byte message block or have an extra finalization round that has a bigger relative effect on short messages.

AMD64: As with long messages, BLAKE and Skein lead all others. BLAKE-256 is in the high group on new machines with large vector-register sizes, but drops down into the medium group for older AMD64 machines. Grøstl-256 stays in the medium group with JH, Keccak-512, Keccak-256, SHA-512 and SHA-256, for machines that implement the AES-NI instructions, but falls into the low group on machines that do not implement the AES-NI instructions. Keccak-512 is in the medium group with Keccak-256, because its relatively small 576-bit input block inputs the entire message in one block. Grøstl-512 is in the low group even with the AES-NI instructions, probably because of the extra overhead of its final “blank” round.

X86: BLAKE-256, Skein and SHA-256 are in the high group on this 32-bit-word ISA. The medium group consists of BLAKE-512, JH and Keccak-256. Grøstl-256, Grøstl-512, Keccak-512 and SHA-512 show low performance.

ARM-NEON: The high-performance group consists of BLAKE-256 and SHA-256; the medium group consists of BLAKE-512, JH, Keccak-512, Keccak-256, Skein and SHA-512; and the low group consists of Grøstl-256 and Grøstl-512.

32-bit RISC: BLAKE-256 and SHA-256, the two algorithms that use 32-bit modular addition extensively, are in the high group; BLAKE-512, Keccak-512, Keccak-256, Skein and SHA-512 are in the medium-performance group, while Grøstl-256, Grøstl-512 and JH make up the low group.

5.1.2.2 **XBX: Embedded Microcontrollers**

Most of the data on embedded microcontrollers comes from the XBX effort; its homepage is at <http://xbx.das-labor.org/trac>. XBX [26] focused on embedded computers, where it is usual to compile or assemble code on some development system, and then run it on a small computer or “microcontroller.” The XBX team collected and measured SHA-3 implementations on eight embedded platforms. Some of the 32-bit ARM processors benchmarked on XBX were as powerful as some of the machines tested by eBASH, and included vector units, while others were 8- and 16-bit machines, characteristically used with very little memory.

In contrast to eBASH, XBX gives the RAM and ROM requirements of the fastest implementations for each processor tested, and “area” versus speed plots, for each of the implementations measured. The area metric used in the XBX is: $\text{area} = 4 \times \text{RAM} + \text{ROM}$. This is a heuristic based on the general observation that RAM memory requires about four times the area on a chip as ROM memory. In the larger systems used in eBASH, the memory required by hash algorithms is rarely an issue, but memory use is often an issue and sometimes the major constraint with embedded microcontrollers. In addition to XBX, [117] gives optimized implementations for all five finalists plus SHA-2 on the ARM-11 processor, which has been commonly used in cell phones, but has been supplanted by newer processors in high-end cell phones.

Eight-bit: Only one 8-bit microcontroller was tested, the Atmel ATmega128P. The smallest area and the best performance was attained by Grøstl-256, perhaps the only software case where Grøstl was the fastest algorithm. This probably is because of the eight-bit orientation of the AES operations used by Grøstl. BLAKE-256 was nearly as fast as Grøstl, but took somewhat more area. Keccak-256 was also compact and third in throughput. Interestingly, SHA-256 was comparatively big and slow, and SHA-512 had the worst overall performance of any algorithm.

Sixteen-bit: Only one 16-bit microcontroller, the Texas Instruments' MSP430FG4618, was tested. Overall, BLAKE-256 was fast and small, and SHA-256 was second. Grøstl-256 was small but fairly slow. Keccak-256 was a bit bigger than BLAKE-256 and somewhat faster than Grøstl-256, but no match for either BLAKE-256 or SHA-256. Keccak-256 requires a relatively small area and has reasonably good throughput.

ARM (thumb instructions): This version of the ARM processor is a low-end microcontroller that only implements the 16-bit thumb instructions. The XBX authors consider small area to be the primary goal for this platform, and BLAKE and JH-256 are the two smallest algorithms. BLAKE-256 and SHA-256 have the highest (very similar) throughputs, but the area requirement for SHA-256 is higher than for BLAKE-256. Keccak has a fairly small area requirement, and Keccak-256 is third in overall performance. Skein has the largest area requirement, and Grøstl-512 and JH are simply slow.

32-bit RISC: This class includes several ARM processors running the regular 32-bit instructions (but without the NEON vector processor) and a 32-bit MIPS processor. Both XBX and [117] provide data for the 32-bit RISC processor. SHA-256 is the fastest algorithm overall, followed by BLAKE-256. Since these are both 32-bit-oriented ARX algorithms, it is expected that SHA-256 and BLAKE-256 should be the fastest algorithms. Skein, BLAKE-512, Keccak-256 and SHA-512 vie for third place, depending on the specific processor, while Grøstl-256, Grøstl-512 and JH are much (typically three to eight times) slower than SHA-256. BLAKE-256 and Grøstl-256 are usually the smallest algorithms in area.

ARM with NEON: This is a relatively fast ARM core with the addition of a vector unit that supports 64-bit operations. With the addition of the vector unit, Skein becomes the fastest algorithm, followed fairly closely by BLAKE-256, BLAKE-512 and SHA-256, Skein is about twice as fast as Keccak-256, three times as fast as Keccak-512, and four to nine times the speed of JH and Grøstl. The overall results are similar to those on eBASH for ARMs with the NEON vector unit, although JH seems to do somewhat worse in the XBX benchmarks than is shown in the eBASH results.

The XBX investigators give an overall ranking for embedded computers that factors in their understanding of whether the usual goal for the particular computer is throughput or minimum area. In that ranking, BLAKE is first, while Skein is in second place, but only slightly better than Keccak and Grøstl. The investigators did not include SHA-2 in this ranking; however, considering the good performance of SHA-256 on 32-bit RISC machines, if SHA-2 were included, it might beat Skein for second place, since there are four 32-bit platforms where SHA-2 is the fastest algorithm, although rarely by a large margin. However, the smallest SHA-256 implementations nearly always require somewhat more area than the smallest BLAKE-256 implementations, so BLAKE-256 has the overall advantage over SHA-2.

5.1.3 Beyond The Superscalar

The sophisticated superscalar arithmetic logic units (ALUs) of large-scale general-purpose computers attempt to extract as much parallel execution as they can from programs written as a single-thread sequence of instructions. Additional parallel execution may be explicitly programmed by taking advantage of multiple core processors that can run independent threads, or by using vector units that perform the same operation on a number of words in parallel.

Most large-scale general-purpose processor chips today have between two and eight cores, each capable of executing an independent program thread. Since all of the SHA-3 finalists can be run efficiently as a thread on a single core, it seems obvious that very large files can be broken into n pieces and each piece hashed in parallel on separate cores in an n -core processor. The message digest of each separate piece may then be hashed in tree fashion. However, operating-system-level interprocess synchronization may be expensive enough to limit the utility of this to very large files, and it is hard to see why one of the finalists would be better suited to parallel execution in separate process threads than any other finalist.

Graphics Processor Units (GPUs) were not considered in the competition. GPUs are common on processors for laptop, desktop and even tablet computers. They are arrays of fairly specialized processors that are organized for stream processing, and are most commonly used for image and video rendering. Gaming and High-Definition video applications have driven GPU development. GPUs have been applied to cryptography, but more for cryptanalysis than for protecting data. GPUs may be well suited to applications that can be divided into many independent processes, such as password cracking, but it is hard to envision much near-term use of GPUs for data protection.

Vector units with eight 64-bit registers became available in X86 machines in 1996. These machines could simultaneously perform one 64-bit or two 32-bit operations in a register. In 2012, the current generation of AMD64 machines has sixteen 256-bit registers. Current machines can simultaneously do four 64-bit operations or eight 32-bit operations in a register. The total size of the vector-register file has increased by a factor of eight in about 16 years.

Vector units are both evolving rapidly and spreading from large-scale general-purpose computers to embedded computers. They are often the vehicle for the addition of new instructions, such as the AES extensions to the AMD64 instruction set. Vector units facilitate more efficient implementations of bit-slice designs, such as JH; of AES and algorithms that use the AES S-boxes (e.g., Grøstl); and of BLAKE-256 on the more recent AMD-64 machines. On ARM machines that implement the Neon vector unit, they speed up most of the finalists. Some new instructions will be useful for hash functions; for example, rotate instructions are not available in most current vector processors, but are announced for inclusion in the next generation of AMD-64 machines.

Figures 1-3 graph the eBASH data for three recent AMD64-architecture processors that implement recent vector units operating on 256-bit vectors. In these graphs, each bar represents the mean cycles-per-byte for the fastest implementation of each algorithm on one of the three processors. Smaller bars mean faster throughput. The five finalist algorithms plus SHA-2 are plotted; however, in the case of Skein and JH, only one bar is plotted, because all four output sizes use the same compression function, and run at about the same rate. In these graphs, one designation is used for two or four message-digest sizes when they run at the same rate; one bar is plotted for each version with a different rate. For example, since BLAKE-256 and BLAKE-224 run at the same speed, only BLAKE-256 is plotted, but all four digest sizes are plotted for

Keccak, because each runs at a different speed. For consistency with other algorithms, the nomenclature for Keccak has also been changed in this report from that used for eBASH, which generally labels the Keccak variants by their capacity; for example, the eBASH “keccak512,” which outputs a 256-bit message digest, becomes “keccak-256” in Figures 1-3.

It is reasonable to expect vector-register files to continue to grow, more instructions to be added, and to find more use of vector units, even in embedded systems. It seems too big an assumption to expect that because AES has inspired additional support instructions, that SHA-3 would also do so. If it did, that would likely be very far in the future, unless the instructions, like vector rotates, have more general uses than implementing hash functions alone.

Figure 1 plots the cycles-per-byte of long messages for “h6sandy,” an Intel® Core i3-2310M¹⁶, which is a “Sandy Bridge” internal architecture processor with two 2100 MHz cores that is intended for laptop computers. Figure 2 plots cycles-per-byte of long messages for “sandy0,” an Intel Core i7-2600K, which is a Sandy Bridge four-core 3400 MHz processor that is intended for relatively high-end desktop computers. Figure 3 plots cycles-per-byte of long messages for “hydra6,” an AMD FX-8120, which is a “Bulldozer” internal architecture processor with four 3100 MHz cores, also intended for desktop applications. The Intel Sandy Bridge processors implement the AVX vector instruction set, with sixteen 256-bit registers; the AMD Bulldozer processors implement the XOP vector instruction set, which is similar to the AVX instruction set, with some additional extensions, including, in particular, the 32- and 64-bit rotate operations. There is one significant difference between the two Sandy Bridge processors: sandy0 implements the AES-NI instructions, and h6sandy does not. The AES-NI instructions affect only the performance of Grøstl.

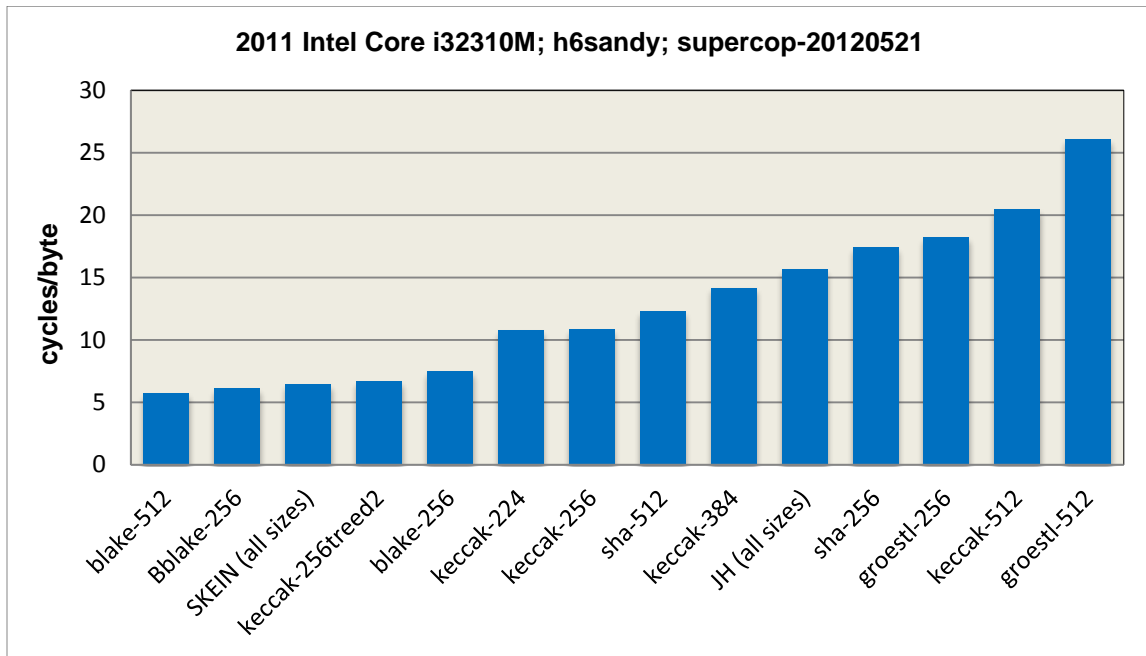


Figure 1. SHA-3 Finalists Cycles/Byte on Sandy Bridge Laptop Processor with Current Vector Unit

¹⁶ The identification of certain products in this document does not imply recommendation by the US National Institute of Standards and Technology (NIST) or other agencies of the US Government, nor does it imply that the products identified are necessarily the best available for the purpose.

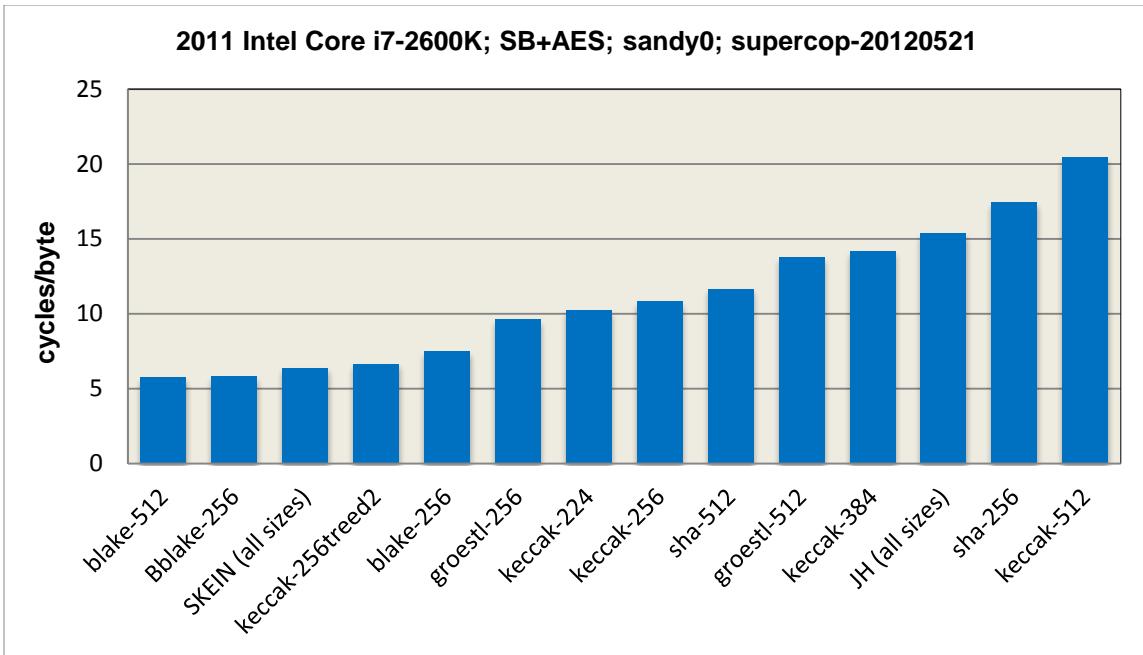


Figure 2. SHA-3 Finalists Cycles/Byte on Sandy Bridge Desktop Processor with Current Vector Unit

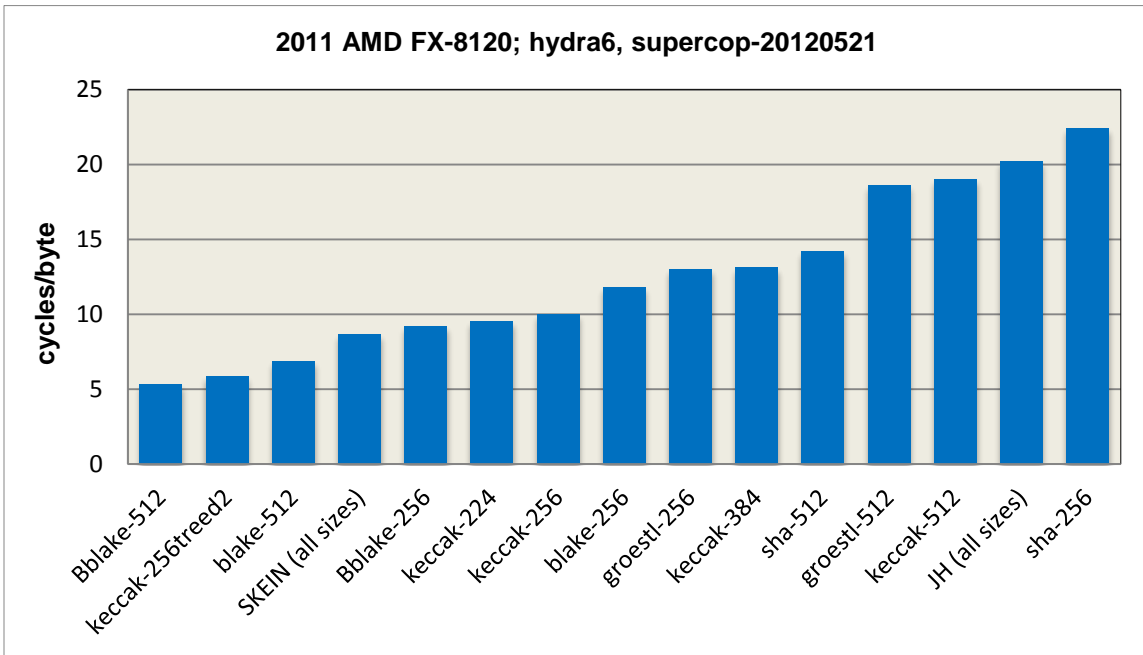


Figure 3. SHA-3 Finalists Cycles/Byte on the Bulldozer Desktop Processor with Current Vector Unit

Skein does not seem to benefit from the vector units, probably because of the different rotation constants used in its MIX operations, even though four parallel MIX operations occur in each round. Skein, however, is one of the fastest algorithms on all three of these machines, relying on their superscalar general-purpose ALUs.

BLAKE logically allows parallel execution of four of its building-block *G* functions, with their rotates; unlike Skein, BLAKE's four parallel rotates all use the same rotation values, facilitating

vectorization. On these machines with the most advanced vector units, BLAKE-256 (which uses 32-bit words) becomes almost as fast as Skein-256, which uses 64-bit words and, in older AMD64 machines, has a bigger advantage over BLAKE-256. BLAKE-512 is slightly faster than Skein-512 on these newer machines, although Skein-512 has the advantage on older AMD64 machines.

Word rotation, used extensively by BLAKE, Skein and Keccak, is not usually a directly implemented instruction in current vector instruction sets, but it can be implemented by shifts and XORs. Keccak does relatively better on the Bulldozer, because only the Bulldozer implements vector-rotate instructions, giving about a factor of 1.7 speed improvement over older vector units [121]. BLAKE also benefits from the Bulldozer rotate instruction [120]. A single Skein thread probably will not be helped very much by a simple rotate instruction that rotates all words by the same amount. Grøstl is significantly helped by the AES-NI instructions and corresponding AES extension on the Bulldozer. JH clearly benefits from vector implementations, but still is among the slowest on all three processors.

In addition to the five finalists and the SHA-2 algorithms described above, two “tree-modes,” Bblake-256 and keccak-256treed2, were implemented on all three processors, while another tree-mode, Bblake-512, was implemented for the AMD FX-8120. These modes generate a single 512- or 256-bit message digest from two parallel streams using the regular compression function, and are examples of “Processing Multiple Buffers in Parallel” [122] or “multi-buffering” to increase performance; in these particular cases, this is done by dividing a single file into two parts to be hashed in parallel. The same technique can also be applied to two independent hash streams, as might occur on a server with many simultaneously active TLS connections.

Gueron and Krasnov [123] have investigated “multi-lane” parallel hashing modes for SHA-256, and for a four-lane version of SHA-256 with 8192-byte messages on the Intel® Core™ i7-3770 “Ivy Bridge”¹⁶ processor with the AVX vector instruction set, they obtained 5.46 clocks per byte. This is a little (5 % to 8 %) better than the results that they obtained for multi-lane versions of Blake-512 and Blake-256, and about 22 % better than a multi-lane version of Keccak. From their work it seems fair to conclude that, for large messages with a multi-lane hashing mode, or where multi-buffering allows the parallel execution of multiple hash operations, SHA-256 is very competitive in throughput, and that 32-bit-word-oriented algorithms (SHA-256 or Blake-256) can be as fast on big 64-bit-oriented processors as 64-bit-word algorithms, wherever multi-lane or multi-buffered execution is practical. These may be the bulk of cases where hash function performance is an issue.

All of the candidates and SHA-2 potentially can benefit from a tree mode or multi-buffer implementations, although perhaps to varying degrees. The different rotation values that complicate vectorizing a single Skein thread should not matter in a multi-buffer implementation. However, there has not been enough work done on this to draw strong conclusions about which algorithms benefit most. In the future, NIST expects to consider a hashing recommendation for a standardized parallel hashing mode.

5.1.4 Software Performance Summary

Skein and BLAKE, the two ARX finalist candidates, have the best overall software performance. Only Skein and BLAKE seem to be faster than SHA-2 in most cases.

Skein has the advantage of a single compression function for all four message-digest sizes, which reduces code size if all four message-digest sizes are required. Skein-256 has a small to moderate

performance advantage over BLAKE-256 on AMD64 platforms, which may, or may not, carry over to future 64-bit ARMv8 [124] processors if the processors adopt superscalar ALUs similar to the ones used in AMD64 processors. BLAKE-512 seems to gain a modest performance advantage over Skein-512 on more recent AMD64 machines by better fitting their vector units.

On 32-bit machines (mainly ARM processors) without vector units, BLAKE-256 is the overall leader, although it has only a small throughput advantage over SHA-256. On ARMs with the NEON vector unit, Skein seems the fastest algorithm, followed fairly closely by BLAKE.

On small embedded computers, BLAKE-256 has the best overall performance. BLAKE-256's throughput is usually at or near the top, although occasionally slightly less than SHA-256, while BLAKE-256 generally has smaller memory requirements than SHA-2 and most of the other candidates.

Keccak is reasonably fast for 224- and 256-bit message-digest sizes on a range of machines, but Keccak-512 is among the slower algorithms on most machines. Keccak benefits significantly from vector rotate instructions, which may become common on future processors. Like Skein, Keccak uses one compression function for all four digest sizes, and is amenable to compact implementations.

Grøstl-256 seems to need special instruction-set support for AES S-boxes to achieve even the medium level of throughput performance on computers with a 64- or 32-bit word orientation. Even with such support, Grøstl-512 ranks near the bottom in throughput. However, on the one eight-bit microcontroller tested, Grøstl-256 was the fastest finalist and one of the most compact.

JH seems to benefit from larger vector sizes, but even with vector instruction sets that allow 128-bit operations, it is difficult to achieve a medium-level performance. JH may be more competitive on vector units with 256-bit or larger vector registers.

5.2 Hardware Performance

There are two major implementation technologies for cryptographic circuits: Field Programmable Gate Arrays and Application Specific Integrated Circuits. FPGAs are programmable devices that allow the quick implementation of logic circuits that can easily be tested and modified. They are mainly used for relatively low-volume products. ASICs are custom integrated circuits that are fabricated on a silicon wafer. The initial fixed costs for the design and fabrication of ASICs is higher and takes longer than for FPGAs, but the performance (*throughput* and *energy-required-per-message-bit*, or *energy-per-bit* for short, to compute the hash of a message) of ASICs is much better than FPGAs. ASICs are also much less expensive than comparable FPGA implementations when a large number of ASICs are produced. The modern paradigm for implementing high-volume electronics is a system on a chip: an ASIC with a processor, some memory and various support and external interface circuits, possibly including cryptographic hardware. The ASIC share of the market is considerably larger than the FPGA share, but there is still a significant FPGA market, since FPGAs reduce fixed costs, get products to market faster and changes are much easier during the design process.

It was practical to conduct FPGA design studies of the SHA-3 finalists, where many variations on a design were implemented and their performance measured, but was much too expensive to do the same thing with ASICs. However, the relative performance of the finalist algorithms in ASICs and FPGAs are often, but not always, similar [125, 126]. In particular, a good correlation has been shown between results obtained using standard-cell CMOS 65nm UMC ASIC technology

and a 65nm high-performance Altera FPGA family Stratix III [125]. The similar correlation may not exist between a different ASIC technology and a different FPGA family, due to significant differences among various ASIC standard-cell libraries, and various internal architectures of modern FPGA families.

The data obtained from ASIC implementations during the SHA-3 competition is particularly valuable, because ASICs will be used for high-volume applications. However, the performance data from FPGA studies gave more detailed insight into the effects of design trade-offs. The existing ASIC implementations of the SHA-3 finalists implement a full round of the algorithm, or even unroll several rounds; FPGA studies also provided insight into folded partial-round implementations and various unrolled or pipelined implementations. Some finalist algorithms use different compression functions for their 256- and 512-bit variants; in order to reduce the implementation cost, only the 256-bit ASIC implementations of the finalists were built.

There are many hardware design trade-offs. While the memory required by a hash function in software is not usually a concern on any but very small processors, the area needed by a hardware hash-function implementation affects the space available for other functions on the integrated circuit and the price of the final product. Computers running software can exploit only certain patterns of parallel execution – for example, a word addition and a rotate at the same time, or adding or XORing the same value to several words in parallel. On the other hand, a circuit can be designed to simultaneously execute every operation where parallel execution is logically possible. However, parallel execution requires additional area on the chip, and many designs sacrifice parallel execution or do not unroll logic, when it would speed up the circuit, in order to save area. Reusing complex circuit elements saves area, but usually slows throughput and adds additional switching and control logic, so at some point, little or no area savings may result. The usual overall performance metric for a hardware hash-function implementation is throughput/area.

The compression function of each of the SHA-3 finalists repeatedly executes a fairly simple round function. This is inherently a serial process. When the entire round is implemented in a single clock cycle, the result often gives the best throughput-to-area ratio, and this appears to be the case for SHA-2 and several of the finalists. However, the round functions themselves usually consist of repetitions of similar operations, so a round can be “folded” and only a part of the round implemented in a single clock cycle. At the other extreme, the round functions can be unrolled, and more than one round can be executed in a single clock cycle, but, while the circuit gets bigger because of the unrolling, the throughput usually does not get proportionately higher. Skein, in particular, benefits from unrolling because unrolling allows the rotation circuitry to be simplified, and because the delay of k consecutive additions is much smaller than k times the delay of a single addition, where k is an unrolling factor (typically 4).

When several rounds or operations within a round are unrolled, it is possible to pipeline several independent hash-function streams by adding extra memory to the circuit, and this may improve the overall throughput/area ratio of the circuit, but not the maximum throughput of a single hash function stream. This is analogous to “multi-buffering” [122] software hash functions with vector units. Skein benefits from pipelining [127], as can BLAKE [125].

All the SHA-3 finalist candidates can be efficiently implemented with 32- and 64-bit logic or arithmetic operations in order to run well on contemporary, general-purpose computers. The time needed for a modular addition of the natural word size of a general-purpose computer is often as fast as bitwise logic operations. In hardware, this is not the case: word-length bitwise logic is faster than addition with its carry propagation delays. Therefore, the fastest hash function on a general-purpose computer is not necessarily the fastest hash function in hardware.

Of the five finalists, BLAKE and Skein are designed with multiple modular addition operations, while Keccak, Grøstl and JH are not. Not surprisingly, of these five algorithms, BLAKE and Skein have the best throughput on general-purpose computers, while Keccak, Grøstl and JH usually give better throughput/area performance on hardware implementations. SHA-2 also uses modular addition, but with more non-linear bitwise operations, and requires fewer adder circuits than either BLAKE or Skein.

The George Mason University Cryptographic Engineering Research Group (GMU CERG) Database provided a website resource similar to the eBASH and XBX software benchmarking efforts. The GMU CERG Database consolidated results from all groups that had published results of the SHA-3 finalist hardware implementations, and had views of the ASIC Rankings [128] and the FPGA rankings [129].

5.2.1 High-Performance Implementations

In the final round of the competition, there were three complete ASIC implementations of the 256-bit versions of all five finalists and SHA-2. The studies were done by teams at GMU [125], The Eidgenössische Technische Hochschule Zürich (ETHZ) [130], and Virginia Tech (VT) [131]. The GMU designs were optimized for maximum throughput/area, while the ETHZ designs were optimized for throughput/area at a data rate of 2.488 Gbps (the rate of an OC-48 channel). Both the GMU and ETHZ designs were implemented on the same ASIC with a 65 nm CMOS technology, and the results were reported by Gürkaynak et al. [130]. The VT designs were optimized for throughput/area, implemented in a 130 nm CMOS technology and the results were reported by Gao et al. [131, 132] and summarized in [128].

It is difficult to directly compare the throughput of implementations with such different technologies, but areas are stated in kGate Equivalents (kGE), which are roughly comparable. To make a rough comparison of all three sets of designs, the method employed by Gaj et al. [125] in their “shootout” section was used to normalize the throughput/area and energy/bit results for the candidate implementations, relative to the SHA-256 implementation by the same group using the same technology. High throughput/area designs of SHA-2 seem to be well understood, and all three implementations seem to have roughly the same area in kGEs, so NIST considered that this approach was justified and allows a rough comparison of the performance of 65 nm GMU and ETHZ implementations with the 130 nm VT ASIC implementations on the same normalized scale. *This comparison for throughput/area is somewhat unfair to the ETH implementations, because they were optimized for a specific throughput, not, for maximum throughput/area.*

Figure 4 below compares the areas for the designs, and Figure 5 compares their throughputs. Figures 6 and 7 give the normalized throughput/area and energy per bit. For these designs it is clear that only Keccak-256 is better than SHA-256 in throughput/area, by a factor between two and four. SHA-2 gives its best throughput/area performance with a smaller circuit than any of the finalists. Grøstl and Skein usually require significantly larger areas to achieve their best throughput/area ratios. In energy-per-bit consumed, Keccak and SHA-256 are roughly comparable, and better than the four remaining finalists.

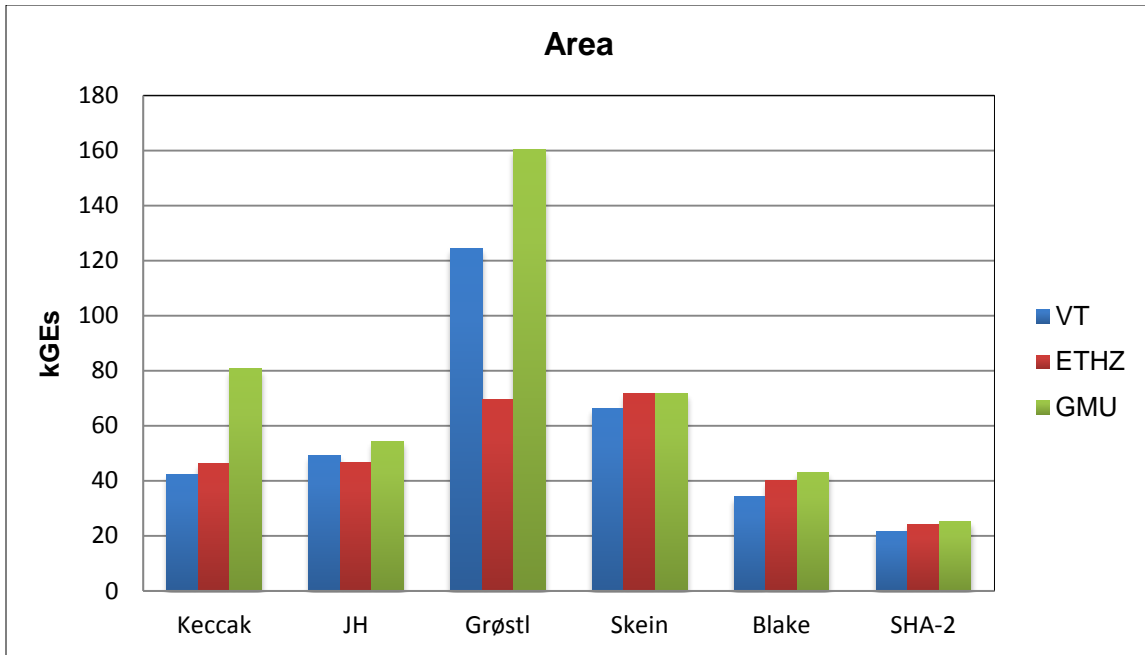


Figure 4. Area in kGate Equivalents for Three 256-bit ASIC Implementations

Comparable ASIC performance data is not available for the 512-bit finalists. It is obvious that BLAKE-512 and Grøstl-512 will roughly double in area; that the area of Keccak-512 will stay about the same, but the throughput will decrease to a bit more than half the rate of Keccak-256; and that the throughput of JH and Skein will stay about the same. The throughput of BLAKE-512 and Grøstl-512 will be reduced by their increase over the number of rounds required for the 256-bit variants, but increased by their larger message-block sizes.

However, there is FPGA data to directly compare the speed differences between 256-bit and 512-bit hash functions. Gaj et al. reported on a broad examination of high- and medium- speed FPGA implementations of the finalists [125], including folded, unrolled and pipelined implementations, and both 256-bit and 512-bit variants for four different FPGA devices. In their “shootout” section, they combined their data with that of implementations from all other published implementations, including Latif et al. [133], and selected the design with the best throughput/area for each of the algorithms and FPGA devices. Figures 8 and 9 below plot the throughput/area for the best implementations of each of the 256-bit and 512-bit finalists, normalized by the throughput/area of SHA-256 or SHA-512 to adjust for the performance differences among the FPGA devices. The best Keccak designs appear to have about a factor of two advantage over SHA-2 for both the 256-bit and 512-bit hash functions. JH-256 and Grøstl-256 have very similar maximum throughput/area ratios that are very close to that of SHA-256. Skein-256 and BLAKE-256 have maximum throughput/area ratios of about half that of SHA-256. However, JH-512 is now clearly in second place behind Keccak-512, and is better than SHA-512, while Grøstl-512 and Skein-512 are similar and a bit below SHA-512. BLAKE-512 is in last place, with a little better than half the throughput/area ratio of SHA-512.

Clearly, in ASICs or FPGAs, and at 256-bits or 512-bits, Keccak has the best throughput/area ratio for medium- to high-performance implementations of the finalists, and is the only finalist that is generally better than SHA-2 in throughput/area. Keccak also has the most compact implementation of any of the finalists when they are optimized for throughput/area. In addition,

Keccak exhibits the lowest energy-consumption-per-message-bit of any of the finalists, and has about the same energy-per-bit requirement as SHA-2.

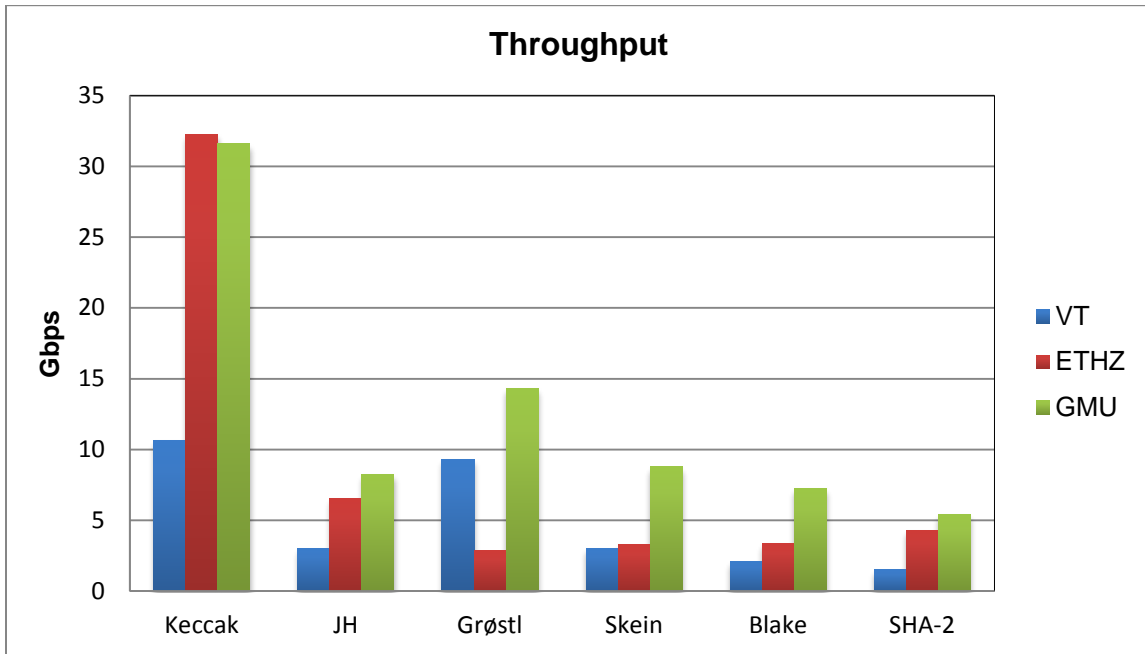


Figure 5. Throughputs of Three 256-bit ASIC Implementations

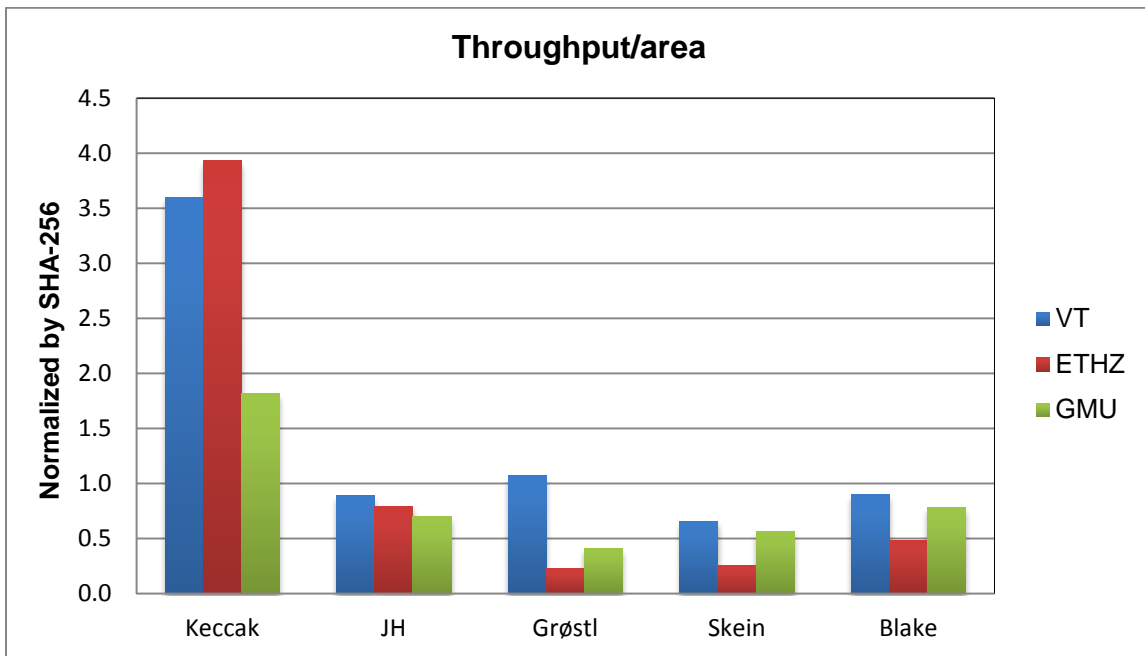


Figure 6. Normalized Throughput/area for Three 256-bit ASIC Implementations

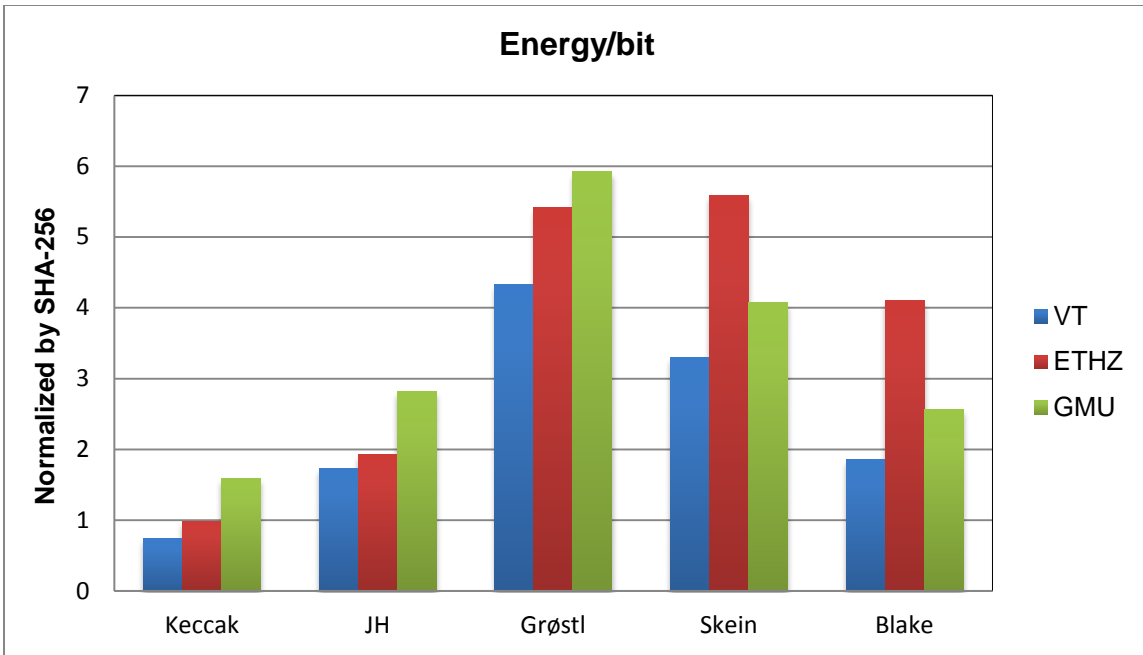


Figure 7. Normalized Energy per Bit for Three 256-bit ASIC Implementations

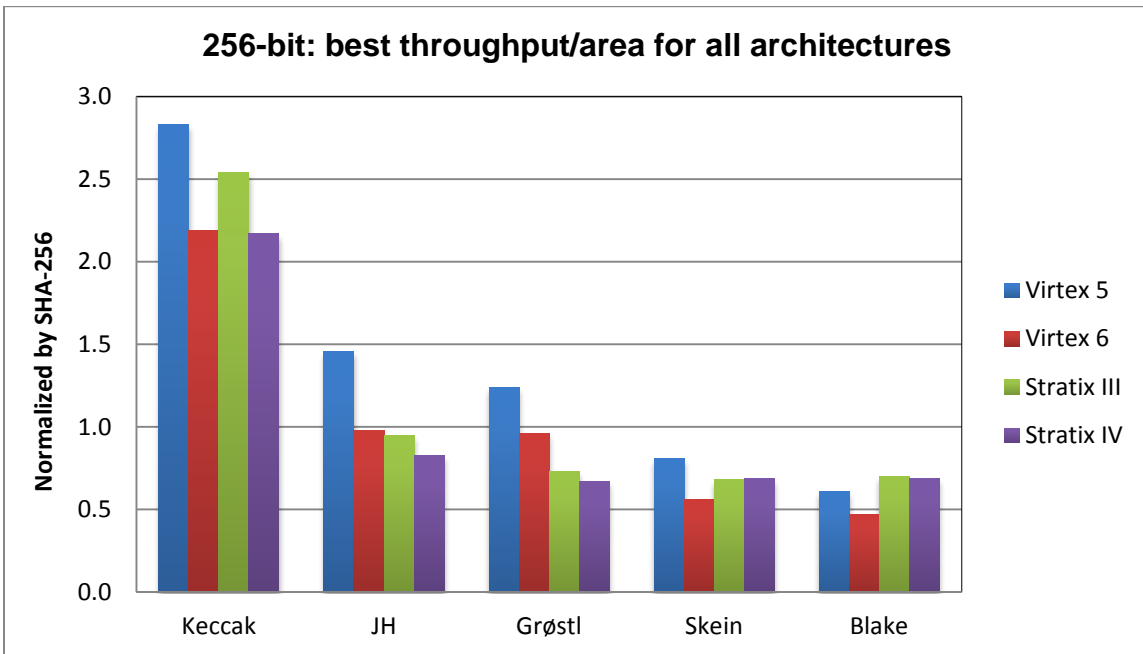


Figure 8. (Normalized) Best Throughput/area for 256-bit FPGA Implementations

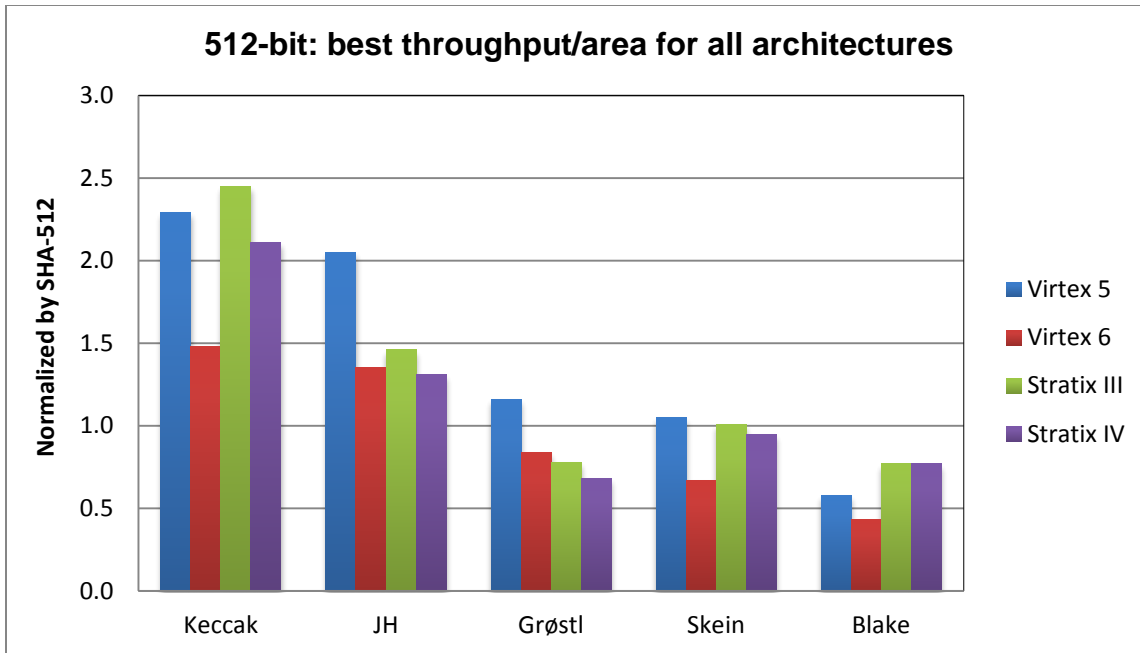


Figure 9. (Normalized) Best Throughput/area for 512-bit FPGA Implementations

5.2.2 Compact Implementations

Extremely small or “lightweight” implementations perform differently from the high-performance implementations discussed above, in that they do not implement a full round, rather they find some way to “fold” the round or to reuse the same circuitry several times to execute a round, and are harder to evaluate. For compact implementations, the first concern is with the area of the circuit: the application designer may want the smallest implementation that meets some throughput target, or may want the smallest implementation possible, and be willing to live with the performance that results. Several papers described lightweight implementations of the 256-bit versions of the finalists, and one lightweight study featured the 512-bit versions of the SHA-3 finalists.

Kavun and Yalçın [134] reported lightweight “word-serial” implementations for the five finalists in a 90 nm CMOS ASIC, with areas between 9.2 and 15.5 kGEs. The word-serial designs implemented a small number of word-sized logic circuits (such as 32-bit XOR, AND, modular addition or rotation) that were reused in a serial fashion, supplemented, where convenient or advantageous, by a few specialized logic circuits for parallel operations, such as the Skein permutation of 64-bit words or the Keccak π permutation. Their results are summarized in Table 20. Their most compact implementation was for Grøstl-256, and the best throughput/area ratio was for BLAKE-256. The authors estimated that for an additional 2 kGE, they could have produced a Grøstl-256 implementation with about the same area and throughput as BLAKE-256, so BLAKE-256 and Grøstl-256 seem best suited to this kind of implementation. For comparison, the optimized throughput/area implementations of SHA-256 provided by Gürkaynak et al. [130] used around 25 kGE in area, so these implementations were around half the area of a high-performance SHA-256 implementation.

While most lightweight-implementation studies focused on 256-bit versions of the algorithms, Kerckhof et al. [135] implemented lightweight versions of the 512-bit variants of the finalists. All of these 512-bit implementations used a 64-bit internal path, which is a natural choice, as most of

the algorithms are designed to operate well on 64-bit processors. Grøstl-512 had the highest throughput/area ratio of any of the finalists, about five times that of Keccak and at least twice that of any other algorithm. The authors claimed that “the clear advantage of Keccak in a high-throughput FPGA implementation context vanishes in a low area one.”

Algorithm	Word (bits)	Area (kGE)	Throughput (kbps@100kHz)	Throughput/Area (Bps/GE)
BLAKE-256	32	11.3	213	18.88
Grøstl-256	32	9.2	40	4.32
JH-256	32	13.6	36	2.61
Keccak-256	64	15.2	91	5.96
Skein-256	64	15.5	86	5.58

Table 20. Compact ASIC “Word-serial” Implementations of the 256-bit Finalists

Jungk [136] reported a study of constrained 256-bit implementations on the Virtex-5 FPGAs that made extensive use of distributed RAM in his implementations. The Keccak design is not folded along 64-bit Keccak “lanes,” but rather folded into units of eight 5-bit \times 5-bit “slices.” Jungk concluded that the results hinted that Grøstl is the best overall performer for compact implementations when the throughput/area ratio is the most important consideration. Grøstl is followed by JH, Keccak, and BLAKE, which are close together, while Skein trails behind. When area is the most important consideration, JH appears to be the best, followed by BLAKE, Grøstl, Keccak and Skein.

Kaps et al. [137] reported a GMU study on area-constrained FPGA implementations. The GMU study fixed two targets: one of 768 slices and another of 450 to 650 slices, plus one block of RAM for the Xilinx Spartan-3 FPGAs. GMU designed the highest throughput implementations of the 256-bit finalists and SHA-256 that they could under those constraints. Figure 10 shows the areas used for the GMU Spartan-3 implementations, while Figure 11 gives the throughput/area for the implementations. BLAKE-256 appears to be the clear throughput/area winner, followed by Grøstl-256 and then JH-256. Only BLAKE-256 and Grøstl-256 are better than SHA-256. Figure 12, from the Kaps presentation [137], plots the GMU designs (ported to a Virtex-5) and the Jungk designs in a throughput versus area chart. In this chart, the relative performance changes as the area doubles and triples. At the smaller sizes (around 500 slices), BLAKE is the fastest, but by 1200 slices, Grøstl-256 becomes the overall throughput/area leader. By about 1700 slices, Keccak-256 passes Blake, but is still slower than Grøstl-256. In this range, Skein is the slowest algorithm, and JH, while better than Skein, seems to always require more area than the other three to achieve the same throughput.

At *et. al.* [138] implemented Grøstl-256 and Grøstl-512 as well as AES-128, AES-192 and AES-256 with an 8-bit data path that requires only 169 slices on a Virtex-6 FPGA. Because Grøstl uses the AES S-boxes and similar operations, much of the circuitry is shared, and including AES adds only 67 slices.

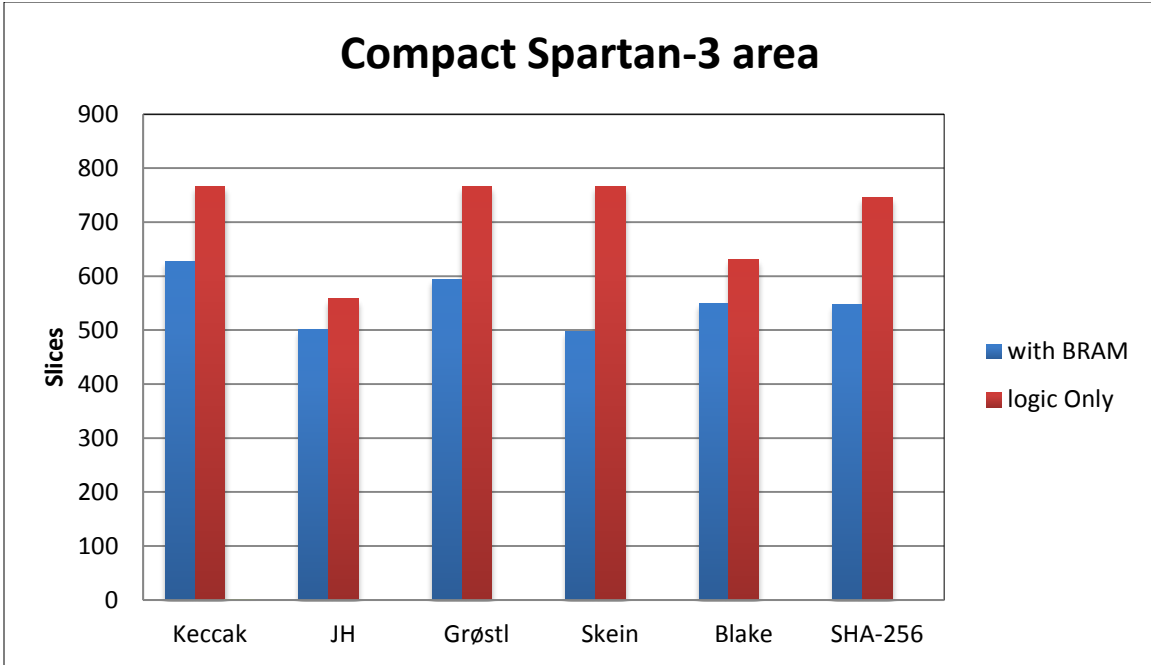


Figure 10. Area of GMU Compact Spartan-3 Implementations of the 256-bit Variants

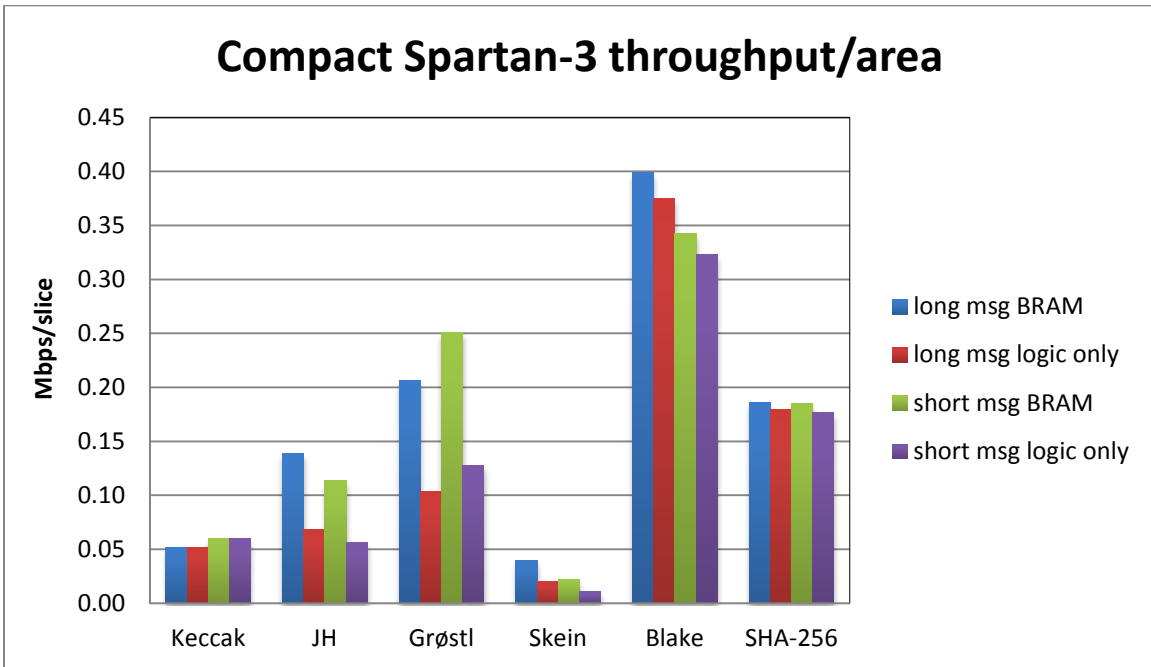


Figure 11. Throughput/Area for GMU Spartan-3 Implementations of the 256-bit Variants

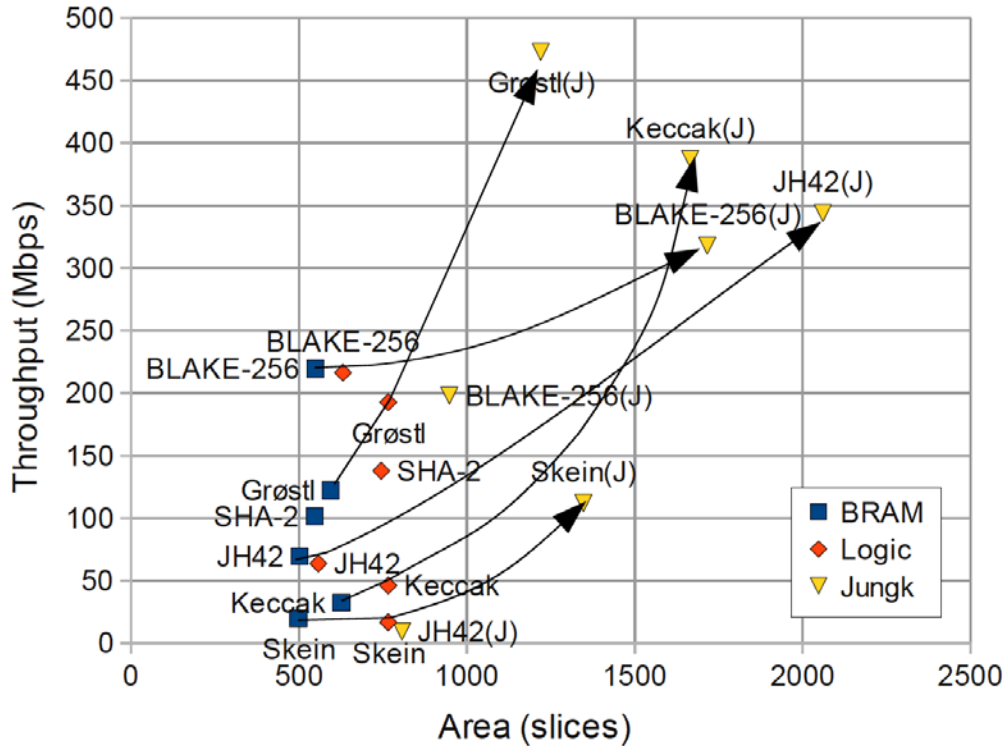


Figure 12. Comparison of GMU Constrained Implementations Ported to Virtex-5 with Jungk [136] Implementations

5.2.3 Discussion of SHA-2 and the SHA-3 Finalists

The implementations of SHA-2 and each of the finalists is discussed below.

5.2.3.1 SHA-2

The SHA-2 algorithms use a conventional Davies-Meyer construction with a chaining variable of eight words and a message block input of sixteen words. SHA-256 uses 32-bit words, while SHA-512 uses 64-bit words in its design. Except for the word size and the number of rounds that is called for, the two compression functions are similar and rely on bitwise AND operations in the round function and modular addition, in the message expansion and the round function, for nonlinearity. Only two words of the chaining variable require recomputation during each round, and part of round $i+1$ can be pre-computed during round i . SHA-2's round function has no symmetries that permit straightforward folding, but the whole round function is still quite compact, usually resulting in a fairly good throughput/area ratio, – as good as or better than any of the finalists except for Keccak. A “full” SHA-2 implementation is only a little bigger than many folded implementations of the SHA-3 finalists. SHA-2 provides a reasonably good benchmark for comparison with the finalists, because hardware implementation of SHA-2 has been well studied [139, 140, 141] and is apparently mature. Consequently, there seems to be less variation in the size and throughput of various implementations of SHA-2 by different investigators using the same devices, than for the SHA-3 finalists. Many of the investigators exploring hardware performance included SHA-2 for comparative purposes in their studies. To compare implementations by different investigators using different technologies or FPGA devices in a single chart, Figures 6 through 9 are normalized using the performance of SHA-2 for the same implementer and technology.

5.2.3.2 BLAKE

In hardware, BLAKE is one of the most flexible SHA-3 finalists, since it can be folded vertically and horizontally by two or four, and pipelines readily within a single round. BLAKE gives the best performance of any algorithm for very compact FPGA implementations, and the same would probably be true for ASIC implementations. Like SHA-2, BLAKE has the advantage of two different compression functions to aid compact 256-bit implementations, one with 32-bit words and the other with 64-bit words. High-performance implementations of BLAKE in FPGAs or ASICs typically require about twice the size of SHA-2, with about the same throughput, so BLAKE's throughput/area ratio is roughly half that of SHA-2. This ratio seems to hold for both BLAKE variants. When FPGA implementations of BLAKE are folded to be smaller than SHA-2 (roughly half the size of SHA-2), they seem, at best, to give about half the throughput/area result of full SHA-2 implementations on the same FPGA device, and sometimes much less.

5.2.3.3 Grøstl

Grøstl relies on the eight-bit AES S-boxes for nonlinearity, and a Grøstl implementation can share circuitry with an AES implementation where both algorithms are needed [138, 142, 143]. Grøstl uses two very similar fixed permutations of either 512 or 1024 bits to build the round function for 256-bit and 512-bit variants, respectively. There are many possible design variations for Grøstl, more, perhaps, than any other finalist. Most designs take several times the area of SHA-2. However, Jungk [136] reported on a Grøstl-256 implementation that was folded eight ways on a Virtex-5 FPGA, with a size similar to a full SHA-256 implementation. This implementation gave about $\frac{3}{4}$ the throughput/area performance of the GMU SHA-2 implementation on the Virtex-5, and a somewhat better performance than Jungk's folded Keccak implementation on the Virtex-5. While high-performance Grøstl implementations seem to be quite large and nearly always have a lower throughput/area ratio than SHA-2, Grøstl is generally better than BLAKE and Skein in throughput/area.

5.2.3.4 JH

Like Skein, JH has a single compression function that is used by both 256-bit and 512-bit variants and, like Skein, the 256- and 512-bit hash functions run at about the same throughput rate. This means that 512-bit JH implementations are more competitive in throughput/area, compared to the other 512-bit algorithms than is the JH 256-bit implementation. JH can be readily folded and pipelined, and Jungk [136] and Kerckhof et al. [135] both reported highly folded FPGA implementations that were compact, but slow (compared to the BLAKE implementations with similar area). In the GMU FPGA throughput/area "shootout," JH-256 has about the same result as SHA-256, while JH-512 is better than SHA-512, but worse than Keccak-512.

5.2.3.5 Keccak

The Keccak permutation has a 1600-bit state that can be viewed as a 5×5 array of 64-bit "lanes" or as sixty-four 5×5 -bit "slices". In software, the algorithm is conveniently implemented as 64-bit logic operations, rotations, loads and stores. However, general-purpose computers cannot exploit most of the parallelism latent in the algorithm, and a full hardware implementation of Keccak is naturally highly parallelizable, resulting in a very good throughput/area ratio, typically about twice or more of the throughput/area ratio of a full-round SHA-2 implementation. Keccak is the only finalist that is consistently better than SHA-2 in throughput/area for high-performance hardware implementations. A full-round implementation of Keccak-256 in either ASICs or

FPGAs seems to be significantly (two to four times) bigger, but much faster (four to ten times) than SHA-256. Keccak-512 uses the same compression function as Keccak-256, and has about half the throughput rate, and SHA-512 gives a somewhat lower throughput/area measurement, so that Keccak-512 retains a reduced, but still substantial, throughput/area advantage over SHA-512.

The results for compact implementations are more complex. Keccak can be folded either along the lanes or in units of 1, 2, 4, 8, 16 or 32 slices. Kerckhof et al. [135] implemented a lane-oriented, folded version of Keccak with five 64-bit registers and a 64-bit barrel rotator, but got lower performance than the other four finalists and concluded that Keccak is inherently less suitable for compact FPGA implementation than the other SHA-3 finalists. However Jungk [136] studied compact FPGA implementations of the finalists and used an eight- (of 64) slice FPGA implementation of Keccak, getting a somewhat larger, but much faster, implementation that is second in his study only to Grøstl in throughput/area. This suggests that reasonably efficient, folded implementations of Keccak are possible, but such folded Keccak implementations have only about average performance.

5.2.3.6 Skein

A Skein round consists of four parallel MIX operations, each with one 64-bit modular addition operation, one 64-bit rotation and one 64-bit XOR operation. Skein has 72 rounds, with a three-tier round structure: subkeys are added every four rounds, and the rotation pattern repeats every eight rounds. When the eight rounds are unrolled, no logic is required for the rotations. The maximum throughput/area performance of a non-pipelined Skein implementation seems to occur when four rounds are implemented in a single cycle. Walker et al. [127] reported on an ASIC design with two pipelined stages in eight rounds as a good hardware trade-off for a router or server processing multiple hash streams. The Skein round can be folded vertically, and Jungk [136] reported a Skein-256 implementation that used only three 32-bit adders, but this design was about the same size and five to seven times slower than the full-round implementations of SHA-256 reported by Gaj et al. [125] on the same FPGA device. Overall, Skein lends itself to unrolled, pipelined implementations that can achieve high overall throughput; however, this approach requires a comparatively large area. Since the same compression function is used for both Skein-256 and Skein-512, Skein-512 is more competitive in throughput/area. Skein is not a good choice for compact 256-bit hash function implementations or when it is implemented to maximize throughput/area (typically by unrolling four rounds), since it requires around three times the area of SHA-256.

5.2.4 Hardware Performance Summary

In high-performance implementations, Keccak is the clear throughput/area winner, and the only candidate that is better than SHA-2. Skein and BLAKE have the lowest maximum throughput to area ratio.

In very compact implementations, if the assumption that a 256-bit variant would be preferred over a 512-bit variant is correct, then it would be expected that the two finalists with different 256-bit and 512-bit compression functions, BLAKE-256 and Grøstl-256, would have an advantage; this appears to be the case for hardware implementations. However, the advantage is not entirely because of the separate compression functions – they are also very amenable to folding and pipelining, and it is probably easier to tailor their implementations to a specific size because of this flexibility. Reasonably efficient, folded or serial implementations of all of the algorithms seem achievable, but BLAKE and Grøstl seem the most flexible.

6. Other Considerations

6.1 Intellectual Property

One of the SHA-3 submission requirements is the Intellectual Property Statement. Each submission package needed to contain a set of signed statements to indicate that the submitted algorithm would be available worldwide on a royalty free basis during the competition, and remain so after the competition if the algorithm is chosen as the SHA-3 winner. These statements included:

- Statement by the Submitter,
- Statement by Patent (and Patent Application) Owner(s) (if applicable), and
- Statement by Reference/Optimized Implementations' Owner(s).

For the last two statements, separate statements were required if multiple individuals are involved. FRN-Nov07 [6] provides full details of these requirements.

All of the SHA-3 finalist packages satisfied this requirement; and the designers have reconfirmed their commitment to this agreement in the last stage of the selection process.

6.2 Other Features

A number of candidate algorithms suggested ways that their submissions could be modified to produce added functionality above and beyond the ways that current hash functions are generally used. BLAKE's domain extender incorporates a salt field, which could be used for more efficient randomized hashing, for example, and its compression function is built from a 512- or 1024-bit wide block cipher that could potentially be used on its own. Keccak's sponge domain extender also provides a significant amount of flexibility. It can be fairly easily modified to produce a larger or smaller message digest, or to trade off bits of security against generic attacks for increased performance in a controlled way. The Keccak team also published an efficient way of using the Keccak permutation in an authenticated-encryption mode [144].

The Skein team probably devoted the most effort of any of the finalists towards developing extra functionality. In addition to the tweakable block-cipher, Threefish, which Skein is based upon, the Skein team also suggested ways that Skein could be used for a variety of applications including tree-mode hashing, a custom MAC construction, hashing the public key into message digests for signatures, a stream cipher, and a way of "personalizing" Skein for specific applications to avoid two different algorithms hashing the same data in ways that could be exploited by an attacker.

The Skein team also published a paper [145] strongly advocating that NIST standardize similar extras, whatever the final SHA-3 winner is. The highest priority applications, according to the paper, were support for more efficient MAC constructions than HMAC, and support for nonstandard message-digest sizes. Other features listed as important were tree-based hashing, and hash-based signatures. Finally, the paper listed two more applications as "nice to have:" a standardized technique for "personalization," and a hash-function-based authenticated-encryption mode.

7. Conclusion

After a very long and complex evaluation process, NIST has selected Keccak as the winner of the SHA-3 Competition. It was chosen from a field of five very strong candidates, each of which provides some outstanding performance characteristics and design features. All five finalists appear to provide an adequate security margin, as well.

Contrary to the fears leading up to the SHA-3 Competition, SHA-2 has held up well in the face of continued cryptanalysis. The new SHA-3 will need to compete with an existing algorithm (SHA-2) that also offers very strong security and performance. Keccak was chosen, not just for its very strong overall security and performance, but because it offers exceptional performance in areas where SHA-2 does not, and because it relies on completely different architectural principles from those of SHA-2 for its security.

Keccak has a large security margin, good general performance, excellent efficiency in hardware implementations, and a flexible design. Keccak uses a new “sponge construction” domain extender, which is based on a fixed permutation, that can be readily adjusted to trade generic security strength for throughput, and can generate larger or smaller hash outputs, as required. The Keccak designers have also defined a modified chaining mode for Keccak that provides authenticated encryption.

NIST plans to augment the current hash standard, FIPS 180-4, to include the new SHA-3 algorithm, and publish a draft FIPS 180-5 for public review. After the close of the public comment period, NIST will revise the draft standard, as appropriate, in response to the public comments that NIST receives. A final review, approval, and promulgation process will then follow. Additionally, NIST may consider standardizing additional constructions based on the Keccak permutation, such as an authenticated-encryption mode, in the future. Any such future standardization efforts will be conducted in consultation with the public, the Keccak design team and the larger cryptographic research community.

References¹⁷

- [1] FIPS PUB 180-2, Secure Hash Standard (SHS), Information Technology Laboratory, National Institute of Standards and Technology, August 2002, <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>
- [2] FIPS PUB 180-3, Secure Hash Standard (SHS), Information Technology Laboratory, National Institute of Standards and Technology, October 2008, http://csrc.nist.gov/publications/fips/fips180-3/fips180-3_final.pdf
- [3] FIPS PUB 180-4, Secure Hash Standard (SHS), Information Technology Laboratory, National Institute of Standards and Technology, March 2012, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>
- [4] X. Wang, Y. L. Yin, H. Yu, “Collision Search Attacks on SHA-1,” 2005, <http://www.c4i.org/erehwon/shanote.pdf>
- [5] Announcing the Development of New Hash Algorithm(s) for the Revision of Federal Information Processing Standard (FIPS) 180-2, Secure Hash Standard, Federal Register / Vol. 72, No. 14 / Tuesday, January 23, 2007 / Notices 2861, http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Jan07.pdf
- [6] Announcing Request for Candidate Algorithm Nominations for a New Cryptographic Hash Algorithm (SHA-3) Family, Federal Register/ Vol. 72, No. 212 / Friday, November 2, 2007 / Notices 62212, http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Nov07.pdf
- [7] J. Nechvatal, E. Barker, L. Bassham, W. Burr, M. Dworkin, J. Foti, E. Roback, “Report on the Development of the Advanced Encryption Standard (AES),” Information Technology Laboratory, National Institute of Standards and Technology, October 2, 2000, <http://csrc.nist.gov/archive/aes/round2/r2report.pdf>
- [8] J.-P. Aumasson, L. Henzen, W. Meier, R. C.-W. Phan, “SHA-3 proposal BLAKE,” Submission to NIST (Round 3), 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/submissions_rnd3.html
- [9] P. Gauravaram, L. R. Knudsen, K. Matusiewicz, F. Mendel, C. Rechberger, M. Schl affer, S. Thomsen, “Gr ostl A SHA-3 Candidate,” Submission to NIST (Round 3), 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/submissions_rnd3.html
- [10] H. Wu, “The Hash Function JH,” Submission to NIST (Round 3), 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/submissions_rnd3.html
- [11] G. Bertoni, J. Daemen, M. Peeters, G. V. Assche, “Keccak Specifications,” Submission to NIST (Round 3), 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/submissions_rnd3.html
- [12] N. Ferguson, S. Lucks, B. Schneier, D. Whiting, M. Bellare, T. Kohno, J. Callas, J. Walker, “The Skein Hash Function Family,” Submission to NIST (Round 3), 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/submissions_rnd3.html

¹⁷ A number of the references below are to online resources provided by individuals or organizations that are not affiliated with NIST. While some effort was made to provide links to resources that NIST expects to remain available over the long term, this was not always possible. In these cases, NIST felt that even temporarily available online resources were of value.

- [13] A. Regenscheid, R. Perlner, S. Chang, J. Kelsey, M. Nandi, S. Paul, “*Status Report on the First Round of the SHA-3 Cryptographic Hash Algorithm Competition*,” NIST IR 7620, September 2009, http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/documents/sha3_NISTIR7620.pdf
- [14] M. Sönmez Turan, R. Perlner, L. Bassham, W. Burr, D. Chang, S. Chang, M. Dworkin, J. Kelsey, S. Paul, R. Peralta, “*Status Report on the Second Round of the SHA-3 Cryptographic Hash Algorithm Competition*,” NIST IR 7764, February 2011, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/documents/Round2_Report_NISTIR_7764.pdf
- [15] The Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, <http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/index.html>
- [16] Hash Forum, http://csrc.nist.gov/groups/ST/hash/email_list.html
- [17] ECRYPT Hash Workshop 2007, Barcelona, Spain, May 24-25, 2007, <http://events.iaik.tugraz.at/HashWorkshop07/>
- [18] Hash Functions in Cryptology: Theory and Practice, Leiden, The Netherlands, June 2-6, 2008, <http://www.lorentzcenter.nl/lc/web/2008/309/info.php3?wsid=309>
- [19] International Workshop on Security 2009 (IWSEC09), Toyama, Japan, October 28-30, 2009, <http://www.iwsec.org/2009/index.html>
- [20] Hash³: Proofs, Analysis, and Implementation, Tenerife, Spain, November 16-20, 2009, <https://www.cosic.esat.kuleuven.be/ecrypt/courses/tenerife09/index.shtml>
- [21] ECRYPT II Hash Workshop, Tallinn, Estonia, May 19-20, 2011, <http://www.ecrypt.eu.org/hash2011/>
- [22] Quo Vadis Cryptology? SHA-3 Contest - 7th International Workshop on the State of the Art in Cryptology and New Challenges Ahead, Warsaw, Poland, May 23-24, 2011, <http://cryptography.gmu.edu/quovadis/>
- [23] Hash Workshop on SHA3 Finalists, Kolkata, India, December 9-10, 2011, <http://www.isical.ac.in/~coec/workshop/hash2011.html>
- [24] ECRYPT SHA-3 Zoo, http://ehash.iaik.tugraz.at/wiki/The_SHA-3_Zoo
- [25] D. Bernstein and T. Lange (editors), eBASH: ECRYPT Benchmarking of All Submitted Hashes, <http://bench.cr.yp.to/ebash.html>
- [26] External Benchmarking Extension (XBX), <http://xbx.das-labor.org/trac>
- [27] C. Rechberger, T. Bjørstad, J. Daemen, C. De Cannière, P. Gauravaram, D. Khovratovich, W. Meier, T. Nad, I. Nikolić, M. Robshaw, M. Schläffer, S. Thomsen, E. Tischhauser, D. Toz, G. Van Assche, K. Varıcı, “*ECRYPT II SHA-3 Design and Cryptanalysis Report*,” August 2010, <http://www.ecrypt.eu.org/documents/D.SYM.4.pdf>
- [28] P. Gauravaram, F. Mendel, M. Naya-Plasencia, V. Rijmen, D. Toz, “*ECRYPT II Intermediate Status Report*,” July 2011, <http://www.ecrypt.eu.org/documents/D.SYM.7.pdf>
- [29] Information Technology Laboratory, National Institute of Standards and Technology, “*SHA-3 Selection Announcement*,” October 2, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/sha-3_selection_announcement.pdf
- [30] Cryptographic Hash Workshop, NIST Gaithersburg, MD, October 31-November 1, 2005, http://csrc.nist.gov/groups/ST/hash/first_workshop.html

- [31] The Second Cryptographic Hash Workshop, UCSB, CA, August 24-25, 2006, http://csrc.nist.gov/groups/ST/hash/second_workshop.html
- [32] The First SHA-3 Candidate Conference, KU Leuven, Belgium, February 25-28, 2009, <http://csrc.nist.gov/groups/ST/hash/sha-3/Round1/Feb2009/index.html>
- [33] The Second SHA-3 Candidate Conference, UCSB, CA, August 23-24, 2010, <http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/>
- [34] FIPS PUB 198-1, The Keyed-Hash Message Authentication Code (HMAC), Information Technology Laboratory, National Institute of Standards and Technology, July 2008, http://csrc.nist.gov/publications/fips/fips198-1/FIPS-198-1_final.pdf
- [35] J. Black, P. Rogaway, T. Shrimpton, “*Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV*,” CRYPTO 2002, pp. 320-335, <http://www.cs.ucdavis.edu/~rogaway/papers/hash.pdf>
- [36] C. Bouillaguet, P. Fouque, “*Practical Hash Functions Constructions Resistant to Generic Second Preimage Attacks Beyond the Birthday Bound*,” Soumis dans Information Processing Letters (IPL), 2011, <http://www.di.ens.fr/~fouque/pub/ipl11.pdf>
- [37] B. Su, W. Wu, S. Wu, L. Dong, “*Near-Collisions on the Reduced-Round Compression Functions of Skein and BLAKE*,” Cryptology ePrint Archive, Report 2010/355, 2010, <http://eprint.iacr.org/2010/355.pdf>
- [38] M. Schl affer, “*Updated Differential Analysis of Gr stl*,” January 2011, <http://www.groestl.info/groestl-analysis.pdf>
- [39] M. Naya-Plasencia, D. Toz, K. Varıcı, “*Rebound Attack on JH42*,” Advances in Cryptology – ASIACRYPT 2011, LNCS 7073, pp. 252-269, Springer, 2011, http://rd.springer.com/chapter/10.1007/978-3-642-25385-0_14
- [40] I. Dinur, O. Dunkelman, A. Shamir, “*New Attacks on Keccak 224 and Keccak 256*,” Cryptology ePrint Archive, Report 2011/624, 2011, <http://eprint.iacr.org/2011/624.pdf>
- [41] H. Yu, J. Chen, K. Wang, X. Wang, “*Near-Collision Attack on the Step-Reduced Compression Function of Skein-256*,” Cryptology ePrint Archive, Report 2011/148, 2011, <http://eprint.iacr.org/2011/148.pdf>
- [42] S. Sanadhya, P. Sarkar, “*New Collision Attacks against Up to 24-Step SHA-2*,” Cryptology ePrint Archive, Report 2008/270, 2008, <http://eprint.iacr.org/2008/270.pdf>
- [43] A. Biryukov, I. Nikolić, A. Roy, “*Boomerang attacks on BLAKE-32*,” Fast Software Encryption 2011, LNCS 6733, pp. 218-237, Springer, 2011, <http://www.springerlink.com/content/r72k542v10330131/>
- [44] J. Jean, M. Naya-Plasencia, and T. Peyrin, “*Improved Rebound Attack on the Finalist Gr stl*,” FSE 2012, Washington DC, March 19-21, 2012, <http://www.di.ens.fr/~jean/pub/fse2012.pdf>
- [45] M. Duan, X. Lai, “*Improved Zero-Sum Distinguisher for Full Round Keccak-f Permutation*,” Cryptology ePrint Archive, Report 2011/023, 2011, <http://eprint.iacr.org/2011/023.pdf>
- [46] J.P. Aumasson, W. Meier, “*Zero-sum Distinguishers for Reduced Keccak-f and for the Core Functions of Luffa and Hamsi*,” Email to Hash Forum, September 9, 2009, <http://www.131002.net/data/papers/AM09.pdf>

- [47] D. Khovratovich, C. Rechberger, A. Savaliev, “*Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 family*,” Cryptology ePrint Archive, Report 2011/286, 2011, <http://eprint.iacr.org/2011/286.pdf>
- [48] H. Yu, J. Chen, K. Wang, X. Wang, “*The Boomerang Attacks on the Round Reduced Skein-512*,” Cryptology ePrint Archive, Report 2012/238, 2012, <http://eprint.iacr.org/2012/238.pdf>
- [49] M. Lamberg, F. Mendel, “*Higher-Order Differential Attack on Reduced SHA-256*,” Cryptology ePrint Archive, Report 2011/037, 2011, <http://eprint.iacr.org/2011/037.pdf>
- [50] D. Chang, “*Midgame Attacks and Defense Against Them*,” unpublished manuscripts, April 2012.
- [51] D. Chang and M. Yung, “*Midgame Attacks (and their consequences)*,” Crypto 2012 Rump Session, August 21, 2012, <http://crypto.2012.rump.cr.yp.to/008b781ca9928f2c0d20b91f768047fc.pdf>
- [52] M. Zohner, M. Kasper, M. Stottinger, “*Side Channel Analysis of the SHA-3 Finalists*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/ZOENER_paper.pdf
- [53] O. Benoit, T. Peyrin, “*Side channel Analysis of Six SHA-3 Candidates*,” CHES 2010, LNCS 6225, pp. 140-157, Springer, 2010, <http://www.springerlink.com/content/822377q22h78420u/?MUD=MP>
- [54] G. Bertoni, J. Daemen, M. Peeters, G. Van Assche, “*Note on Side Channel Attacks and Their Countermeasures*,” <http://keccak.noekeon.org/NoteSideChannelAttacks.pdf>. Submitted to the hash forum on May, 15, 2009, <http://cio.nist.gov/esd/emaildir/lists/hash-forum/msg01511.html>
- [55] G. Bertoni, J. Daemen, M. Peeters, G. Van Assche, “*Building Power Analysis Resistant Implementations of Keccak*,” the Second SHA-3 Candidate Conference, UCSB, CA, 2010, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/documents/papers/BERTONI_KeccakAntiDPA.pdf
- [56] S. Hoerder, M. Wojcik, S. Tillich, D. Page, “*An Evaluation of Hash Functions on a Power Analysis Resistant Processor Architecture*,” Cryptology ePrint Archive, Report 2010/614, 2010, <http://eprint.iacr.org/2010/614.pdf>
- [57] D. Osvik, A. Shamir and E. Tromer, “*Cache Attacks and Countermeasures: the Case of AES*,” Cryptology ePrint Archive, Report 2005/271, 2005, http://reference.kfupm.edu.sa/content/c/a/cache_attacks_and_countermeasures_the_c_3_469711.pdf
- [58] M. Schl affer, “*Cryptanalysis of AES-Based Hash Functions*,” PhD Thesis, March 2011, Graz University of Technology, Austria, https://online.tugraz.at/tug_online/voe_main2.getvolltext?pCurrPk=58178
- [59] E. Biham, O. Dunkelman, “*A Framework for Iterative Hash Functions — HAIFA*,” NIST Second Cryptographic Hash Workshop, UCSB, CA, August 2006, http://csrc.nist.gov/groups/ST/hash/documents/DUNKELMAN_NIST3.pdf
- [60] D. Bernstein, “*ChaCha, A Variant of Salsa20*,” <http://cr.yp.to/chacha.html>
- [61] E. Andreeva, A. Luykx, and B. Mennink, “*Provable Security of BLAKE with Non-Ideal Compression Function*,” Third SHA-3 Candidate Conference, Washington DC, March

- 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/MENNINK_paper.pdf
- [62] D. Chang, M. Nandi, and M. Yung, “Indifferentiability of the hash algorithm BLAKE,” Cryptology ePrint Archive, Report 2011/623, 2011, <http://eprint.iacr.org/2011/623>
- [63] J. Li, L. Xu, “Attacks on Round-reduced BLAKE,” Cryptology ePrint Archive, Report 2009/238, 2009, <http://eprint.iacr.org/2009/238.pdf>
- [64] J. Guo, K. Matusiewicz, “Round-reduced Near-Collisions of BLAKE-32,” WEWoRC 2009, <http://www1.spms.ntu.edu.sg/~guojian/doc/blake-col.pdf>
- [65] J.-P. Aumasson, J. Guo, S. Knellwolf, K. Matusiewicz, W. Meier, “Differential and Invertibility Properties of BLAKE,” Cryptology ePrint Archive, report 2010/043, 2010, <http://eprint.iacr.org/2010/043.pdf>
- [66] M. Sönmez Turan, E. Uyan, “Practical Near-Collisions for Reduced Round Blake, Fugue, Hamsi and JH,” the Second SHA-3 Candidate Conference, UCSB, CA, 2010, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/documents/papers/TURAN_Paper_Erdener.pdf
- [67] J. Vidali, P. Nose, E. Pašalic, “Collisions for Variants of the BLAKE Hash Function,” Proceedings of Information Processing Letters, vol. 110, Issue 14-15, 2010, <http://lkrv.fri.uni-lj.si/~janos/blake/collisions.pdf>
- [68] M. Ming, H. Qiang, S. Zeng, “Security Analysis of BLAKE-32 Based on Differential Properties,” 2010 International Conference on Computational and Information Sciences, pp. 783-786, 2010, <http://dl.acm.org/citation.cfm?id=1933979>
- [69] G. Leurent, “ARXtools: A toolkit for ARX Analysis,” the Third SHA-3 Candidate Conference, Washington DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/LEURENT_paper.pdf
- [70] J.-P. Aumasson, G. Leurent, W. Meier, F. Mendel, N. Mouha, R.C.W. Phan, Y. Sasaki, P. Susil: “Tuple cryptanalysis of ARX with application to BLAKE and Skein,” ECRYPT II Hash Workshop 2011, May 19-20, 2011, http://www.ecrypt.eu.org/hash2011/proceedings/hash2011_01.pdf
- [71] O. Dunkelman, D. Khovratovich, “Iterative differentials, symmetries and message modification in BLAKE-256,” ECRYPT II Hash Workshop 2011, May 19-20, 2011, http://www.ecrypt.eu.org/hash2011/proceedings/hash2011_02.pdf
- [72] P. Gauravaram, L. R. Knudsen, K. Matusiewicz, F. Mendel, C. Rechberger, M. Schläffer, S. Thomsen, “Tweaks on Grøstl,” Submission to NIST (Round 3), 2011, <http://www.groestl.info/Round3Mods.pdf>
- [73] E. Andreeva, B. Mennink, B. Preneel and M. Skrobot, “Security Analysis and Comparison of the SHA-3 Finalists BLAKE, Grøstl, JH, Keccak, and Skein,” the Third SHA-3 Candidate Conference, Washington DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/ANDREEVA_paper.pdf
- [74] E. Andreeva, B. Mennink, B. Preneel, “On the Indifferentiability of the Grøstl Hash Function,” LNCS 6280, pp. 88-105, Springer, 2010, http://rd.springer.com/chapter/10.1007/978-3-642-15317-4_7
- [75] J. Kelsey, “Some notes on Grøstl,” 2009, <http://ehash.iaik.tugraz.at/uploads/d/d0/Groestl-comment-april28.pdf>
- [76] F. Mendel, C. Rechberger, M. Schläffer, S. Thomsen, “The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl,” Proceedings of FSE, LNCS 5665, pp.

- 260-276, Springer, 2009, <http://www2.mat.dtu.dk/people/S.Thomsen/MendelRST-fse09.pdf>
- [77] F. Mendel, T. Peyrin, C. Rechberger, M. Schl affer, “Improved Cryptanalysis of the Reduced Gr stl Compression Function, ECHO Permutation and AES Block Cipher,” Proceedings of SAC, 5867, pp. 16-35, 2009, https://online.tugraz.at/tug_online/voe_main2.getvolltext?pCurrPk=44420
- [78] H. Gilbert, T. Peyrin, “Super-Sbox Cryptanalysis: Improved Attacks for AES-like permutations,” Cryptology ePrint Archive, Report 2009/531, 2009, <http://eprint.iacr.org/2009/531.pdf>
- [79] F. Mendel, C. Rechberger, M. Schl affer, S. Thomsen, “Rebound Attacks on the Reduced Gr stl Hash Function,” Proceedings of CT-RSA, LNCS 5985, pp. 350-365, Springer, 2010, <http://www.springerlink.com/content/p016144684v473r6/>
- [80] T. Peyrin, “Improved Differential Attacks for ECHO and Gr stl,” Cryptology ePrint Archive, Report 2010/223, 2010, <http://eprint.iacr.org/2010/223.pdf>
- [81] K. Ideguchi, E. Tischhauser, B. Preneel, “Improved Collision Attacks on the Reduced-Round Gr stl Hash Function,” Cryptology ePrint Archive, Report 2010/375, 2010, <http://eprint.iacr.org/2010/375.pdf>
- [82] M. Naya-Plasencia, “How to Improve Rebound Attacks,” Cryptology ePrint Archive, Report 2010/607, 2010, <http://eprint.iacr.org/2010/607.pdf>. Also published in Advances in Cryptology - CRYPTO 2011, LNCS 6841, pp. 188-205, Springer, 2011, http://rd.springer.com/chapter/10.1007/978-3-642-22792-9_11
- [83] Y. Sasaki, Y. Li, L. Weng, K. Sakiyama, K. Ohta, “Non-full-active Super-Sbox Analysis: Applications to ECHO and Gr stl,” Proceedings of Asiacrypt, LNCS 6477, pp. 38-55, 2010, <http://www.springerlink.com/content/b288848681x52214/>
- [84] C. Boura, A. Canteaut, C. De Canniere, “Higher Order Differential Properties of Keccak and Luffa” (Presentation), FSE 2011, Feb. 15, 2011, <http://fse2011.mat.dtu.dk/slides/Higher-order%20differential%20properties%20of%20Keccak%20and%20Luffa.pdf>
- [85] C. Boura, A. Canteaut, C. De Canniere, “Higher Order Differential Properties of Keccak and Luffa,” Cryptology ePrint Archive, Report 2010/589, 2010, <http://eprint.iacr.org/2010/589.pdf>
- [86] D. Khovratovich, “Bicliques for permutations: collision and preimage attacks in stronger settings,” Cryptology ePrint Archive, Report 2012/141, 2012, <http://eprint.iacr.org/2012/141.pdf>
- [87] S. Wu, D. Feng, W. Wu, J. Guo, L. Dong, and J. Zou, “(Pseudo) Preimage Attack on Round-Reduced Gr stl Hash Function and Others (Extended Version),” Cryptology ePrint Archive, Report 2012/206, 2012, <http://eprint.iacr.org/2012/206>
- [88] R. Bhattacharyya, A. Mandal, M. Nandi, “Security Analysis of the Mode of JH Hash Function,” FSE 2010, LNCS 6147, pp. 168-191, 2010, http://www.isical.ac.in/~rishi_r/FSE2010-146.pdf
- [89] F. Mendel, S. Thomsen, “An Observation on JH-512,” 2008, http://ehash.iaik.tugraz.at/uploads/d/da/Jh_preimage.pdf
- [90] H. Wu, “The Complexity of Mendel and Thomsen's Preimage Attack on JH-512,” 2009, http://ehash.iaik.tugraz.at/uploads/6/6f/Jh_mt_complexity.pdf

- [91] N. Bagheri, “Re: Free start collision, second preimage, and Preimage on JH,” Email to Hash Forum, November 29, 2008, <http://cio.nist.gov/esd/emailldir/lists/hash-forum/msg00981.html>
- [92] J. Lee, D. Hong, “Collision resistance of the JH hash function,” Cryptology ePrint Archive, Report 2011/019, 2011, <http://eprint.iacr.org/2011/019.pdf>
- [93] D. Moody, S. Paul, D. Smith-Tone, “Improved indifferentiability security bound for the JH mode,” the Third SHA-3 Candidate Conference, Washington, DC, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/PAUL_paper.pdf
- [94] V. Rijmen, D. Toz, and K. Varıcı, “Rebound Attack on Reduced-Round Versions of JH,” FSE 2010, LNCS 6147, pp. 286-303, 2010, <https://www.cosic.esat.kuleuven.be/publications/article-1431.pdf>
- [95] G. Bertoni, J. Daemen, M. Peeters, G. Van Assche, “Sponge Functions,” ECRYPT Hash Workshop, 2007, <http://sponge.noekeon.org/SpongeFunctions.pdf>
- [96] Keccak Crunchy Crypto Collision and Pre-image Contest, http://keccak.noekeon.org/crunchy_contest.html
- [97] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, “On the Indifferentiability of the Sponge Construction,” EUROCRYPT 2008, LNCS 4965, pp. 181-197, Springer-Verlag, Berlin, 2008, http://rd.springer.com/chapter/10.1007/978-3-540-78967-3_11
- [98] D. Gligoroski, R. Ødeård, and R. Jensen, “OBSERVATION: An explicit form for a class of second preimages for any message M for the SHA-3 candidate Keccak,” Cryptology ePrint Archive, Report 2011/261, 2011, <http://eprint.iacr.org/2011/261>
- [99] J.P. Aumasson, D. Khovratovich, “First Analysis of Keccak,” 2009, <http://www.131002.net/data/papers/AK09.pdf>
- [100] P. Morawiecki, M. Srebrny, “A SAT-based Preimage Analysis of Reduced KECCAK Hash Functions,” Cryptology ePrint Archive, Report 2010/285, 2010, <http://eprint.iacr.org/2010/285.pdf>
- [101] C. Boura, A. Canteaut, “A Zero-Sum Property for the Keccak-f Permutation with 18 Rounds,” Email to Hash Forum, January 14, 2010, http://www-rocq.inria.fr/secret/Anne.Canteaut/Publications/zero_sum.pdf
- [102] G. Bertoni, J. Daemen, M. Peeters, G. Van Assche, “Note on Zero-sum Distinguishers of Keccak-f,” Email to Hash Forum, January 15, 2010, <http://keccak.noekeon.org/NoteZeroSum.pdf>
- [103] G. Bertoni, J. Daemen, M. Peeters, G. Van Assche, “Note on Keccak Parameters and Usage,” Email to Hash Forum, February 23, 2010, <http://keccak.noekeon.org/NoteOnKeccakParametersAndUsage.pdf>
- [104] C. Boura, A. Canteaut, “Zero-Sum Distinguishers for Iterated Permutations and Applications to Keccak-f and Hamsi-256,” Selected Areas in Cryptography, 2010, <https://www.rocq.inria.fr/secret/Christina.Boura/data/sac.pdf>
- [105] J. Lathrop, “Cube Attacks on Cryptographic Hash Functions,” Master’s Thesis, Rochester Institute of Technology, May 21, 2009, https://ritdml.rit.edu/bitstream/handle/1850/10821/20782_pdf_C9905D10-60E8-11DE-AE1A-0950D352ABB1.pdf?sequence=1

- [106] P. Morawiecki, First Solutions to Keccak Crunchy Crypto Contest, Submitted on July 29, 2011, http://keccak.noekeon.org/crunchy_contest.html
- [107] A. Duc, J. Guo, T. Peyrin, L. Wei, “*Unaligned Rebound Attack: Application to Keccak*,” Cryptology ePrint Archive, Report 2011/420, 2011, <http://eprint.iacr.org/2011/420.pdf>
- [108] D. Bernstein, “*Second Preimages for 6 (7? (8??)) Rounds of Keccak?*,” Email to Hash Forum, November 27, 2010, <http://cio.nist.gov/esd/emaildir/lists/hash-forum/msg02232.html>, or http://ehash.iaik.tugraz.at/uploads/6/65/NIST-mailing-list_Bernstein-Daemen.txt
- [109] M. Naya-Plasencia, A. Röck, and W. Meier, “*Practical Analysis of Reduced-Round Keccak*,” INDOCRYPT 2011, LNCS 7107, pp. 236-254, Springer, 2011, http://naya.plasencia.free.fr/Maria/papers/Keccak_differential.pdf
- [110] M. Bellare, T. Kohno, S. Lucks, N. Ferguson, B. Schneier, D. Whiting, J. Callas, and J. Walker, “*Provable Security Support for the Skein Hash Family*,” Version 1.0, Apr 2009, <http://www.skein-hash.info/sites/default/files/skein-proofs.pdf>
- [111] J. Chen, K. Jia, “*Improved Related-Key Boomerang Attacks on Round-Reduced Threefish-512*,” Cryptology ePrint Archive, Report 2009/526, 2009, <http://eprint.iacr.org/2009/526.pdf>
- [112] J.P. Aumasson, Ç. Çalık, W. Meier, O. Özen, R. Phan, K. Varıcı, “*Improved Cryptanalysis of Skein*,” Cryptology ePrint Archive, Report 2009/438, 2009, <http://eprint.iacr.org/2009/438.pdf>
- [113] K. McKay, P. Vora, “*Pseudo-Linear Approximations for ARX Ciphers: With Application to Threefish*,” Cryptology ePrint Archive, Report 2010/282, 2010, <http://eprint.iacr.org/2010/282.pdf>
- [114] D. Khovratovich, I. Nikolic, “*Rotational Cryptanalysis of ARX*,” proceedings of FSE2010, pp. 333-346, 2010, <http://www.skein-hash.info/sites/default/files/axr.pdf>
- [115] D. Khovratovich, I. Nikolic, C. Rechberger, “*Rotational Rebound Attacks on Reduced Skein*,” Cryptology ePrint Archive, Report 2010/538, 2010, <http://eprint.iacr.org/2010/538.pdf>
- [116] C. Hanser, “*Performance of the SHA-3 Candidates in Java*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/HANSER_paper.pdf
- [117] P. Schwabe, B. Yang, and S. Yang, “*SHA-3 on ARM11 processors*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/YANG_paper.pdf
- [118] D. Bernstein and T. Lange (editors), eBACS: ECRYPT Benchmarking of Cryptographic Systems, <http://bench.cr.yp.to/>
- [119] D. Bernstein and T. Lange, “*The New SHA-3 Software Shootout*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/BERNSTEIN_paper.pdf
- [120] S. Neves and J.-P. Aumasson, “*BLAKE and 256-bit advanced Vector Extensions*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012,

- http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/NEVES_paper.pdf
- [121] G. Van Assche, “*Keccak, tree hashing and rotation instructions*,” Email to Hash Forum, May 30, 2012, <http://cio.nist.gov/esd/emaildir/lists/hash-forum/msg02431.html>
- [122] V. Gopal, J. Guilford, W. Feghali, E. Ozturk, G. Wolrich, and M. Dixon, “*Processing Multiple Buffers in Parallel to Increase Performance on Intel Architecture Processors*,” Intel Corporation White Paper, July 2010, <http://download.intel.com/design/intarch/papers/324101.pdf>
- [123] S. Gueron and V. Krasnov, “*Parallelizing message schedules to accelerate the computations of hash functions*,” Cryptology ePrint Archive, Report 2012/067, 2012, <http://eprint.iacr.org/2012/067.pdf>
- [124] R. Grisenthwaite, “*ARMv8 Technology Preview*,” Presentation at ARM 2011 TechCon, http://www.arm.com/files/downloads/ARMv8_Architecture.pdf
- [125] K. Gaj, E. Homsirikamol, M. Rogawski, R. Shahid, and M. Sharif, “*Comprehensive Evaluation of High-Speed and Medium-Speed Implementations of Five SHA-3 Finalists Using Xilinx and Altera FPGAs*,” Cryptology ePrint Archive, Report 2012/368, 2012, <http://eprint.iacr.org/2012/368>
- [126] X. Guo, S. Huang, L. Nazhandali and P. Schaumont, “*On The Impact of Target Technology in SHA-3 Hardware Benchmark Rankings*,” Cryptology ePrint Archive, Report 2010/536, 2010, <http://eprint.iacr.org/2010/536>
- [127] J. Walker, F. Sheikh, S. Mathew and R. Krishnamurthy, “*A Skein-512 Hardware Implementation*,” the Second SHA-3 Candidate Conference, UCSB, CA, 2010, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/documents/papers/WALKER_skein-intel-hwd.pdf
- [128] GMU Cryptographic Engineering Research Group, ATHENa Database of ASIC Results - ASIC Rankings, available at http://cryptography.gmu.edu/athenadb/asic_hash/rankings_view
- [129] GMU Cryptographic Engineering Research Group, ATHENa Database of FPGA Results - FPGA Rankings, available at http://cryptography.gmu.edu/athenadb/fpga_hash/rankings_view
- [130] F. Gürkaynak, K. Gaj, B. Muheim, E. Homsirikamol, C. Keller, M. Rogawski, H. Kaeslin, and J. Kaps, “*Lessons Learned from Designing a 65 nm ASIC for Evaluating Third Round SHA-3 Candidates*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/GURKAYNAK_paper.pdf
- [131] X. Guo, M. Srivastav, S. Huang, D. Ganta, M. Henry, L. Nazhandali, and P. Schaumont, “*ASIC Implementations of Five SHA-3 Finalists*,” Proceedings of 2012 Design Automation and Test in Europe Conference, DATE 2012, track A5 on Secure Systems, Dresden, Germany, Mar. 12-16, 2012, <http://rijndael.ece.vt.edu/sha3/publications/DATE2012SHA3.pdf>
- [132] M. Srivastav, X. Guo, S. Huang, D. Ganta, M. Henry, L. Nazhandali, and P. Schaumont, “*Design and Benchmarking of an ASIC with Five SHA-3 Finalist Candidates*,” Elsevier Microprocessors and Microsystems: Embedded Hardware Design (MICPRO),

- Special Issue on Digital System Security and Safety (in press), preprint available at <http://rijndael.ece.vt.edu/schaum/papers/2012micpro.pdf>
- [133] K. Latif, M. Rao, A. Aziz and A. Mahboob, “*Efficient Hardware Implementations and Hardware Performance,*” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/KASHIF_paper.pdf
- [134] E. Kavun, T. Yalçın, “*On the Suitability of SHA-3 Finalists for Lightweight Applications,*” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/KAVUN_paper.pdf
- [135] S. Kerckhof, F. Durvaux, N. Veyrat-Charvillon, F. Regazzoni, G. de Dormale, and F. Standaert, “*Compact FPGA Implementations of the Five SHA-3 Finalists,*” ECRYPT II Hash Workshop 2011, May 19-20, 2011, http://www.ecrypt.eu.org/hash2011/proceedings/hash2011_10.pdf
- [136] B. Jungk, “*Evaluation Of Compact FPGA Implementations For All SHA-3 Finalists,*” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/JUNGK_paper.pdf
- [137] J. Kaps, P. Yalla, K. Surapathi, B. Habib, S. Vadlamudi and S. Gurung, “*Lightweight Implementations of SHA-3 Finalists on FPGAs,*” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/KAPS_paper.pdf, Presentation at http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/presentations/KAPS_presentation.pdf
- [138] N. At, J. Beuchat, E. Okamoto, I. San, and T. Yamazaki, “*A Low-Area Unified Hardware Architecture for the AES and the Cryptographic Hash Function Grøstl,*” Cryptology ePrint Archive, Report 2012/535, 2012, <http://eprint.iacr.org/2012/535.pdf>
- [139] R. Chaves, G. Kuzmanov, L. Sousa, S. Vassiliadis, “*Improving SHA-2 Hardware Implementations,*” Cryptology ePrint Archive, Report 2006/24, 2006, <http://www.iacr.org/archive/ches2006/24/24.pdf>
- [140] R. Chaves, G. Kuzmanov, L. Sousa, S. Vassiliadis, “*Cost-Efficient SHA Hardware Accelerators,*” IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 16, No. 8, pp. 999-1008, August 2008, <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4560238>
- [141] M. Rogawski, X. Xin, E. Homsirikamol, D. Hwang and Kris Gaj, “*Implementing SHA-1 and SHA-2 Standards on the Eve of SHA-3 Competition,*” CryptArchi 2009, Jun 25-27, 2009, Prague, http://cryptography.gmu.edu/athena/presentations/CryptArchi_2009_rogawski.pdf
- [142] K. Järvinen, “*Sharing resources between AES and the SHA-3 second round candidates Fugue and Grøstl,*” the Second SHA-3 Candidate Conference, UCSB, CA, 2010, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/documents/papers/Jarvinen_paper_20100709.pdf
- [143] M. Rogawski and K. Gaj, “*A High-Speed Unified Hardware Architecture for the AES and SHA-3 Candidate Grøstl,*” Proceedings of the 15th Euromicro Conference on Digital

- System Design, Special Session on Architectures and Hardware for Security Applications (AHSA), Cesme, Izmir, Turkey, September 5-8, 2012, <http://mason.gmu.edu/~mrogawsk/arch/dsd2012.pdf>
- [144] G. Bertoni, J. Daemen, M. Peeters, G. V. Assche, “*Duplexing the sponge: single-pass authenticated encryption and other applications*,” the Second SHA-3 Candidate Conference, UCSB, CA, 2010, http://csrc.nist.gov/groups/ST/hash/sha-3/Round2/Aug2010/documents/papers/DAEMEN_DuplexSponge.pdf
- [145] S. Lucks, D. McGrew, D. Whiting, “*Batteries Included - Features and Modes for Next Generation Hash Functions*,” the Third SHA-3 Candidate Conference, Washington, DC, March 22-23, 2012, http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/March2012/documents/papers/LUCKS_paper.pdf

Appendix A – eBASH Shootout Plots

http://bench.cr.yp.to
20120910

crypto_sha3
Long messages
skein512256 sha512 keccak1024round3jh512 groestl1512
skein512256 blake256 keccak512 sha256 groestl256 round3jh256

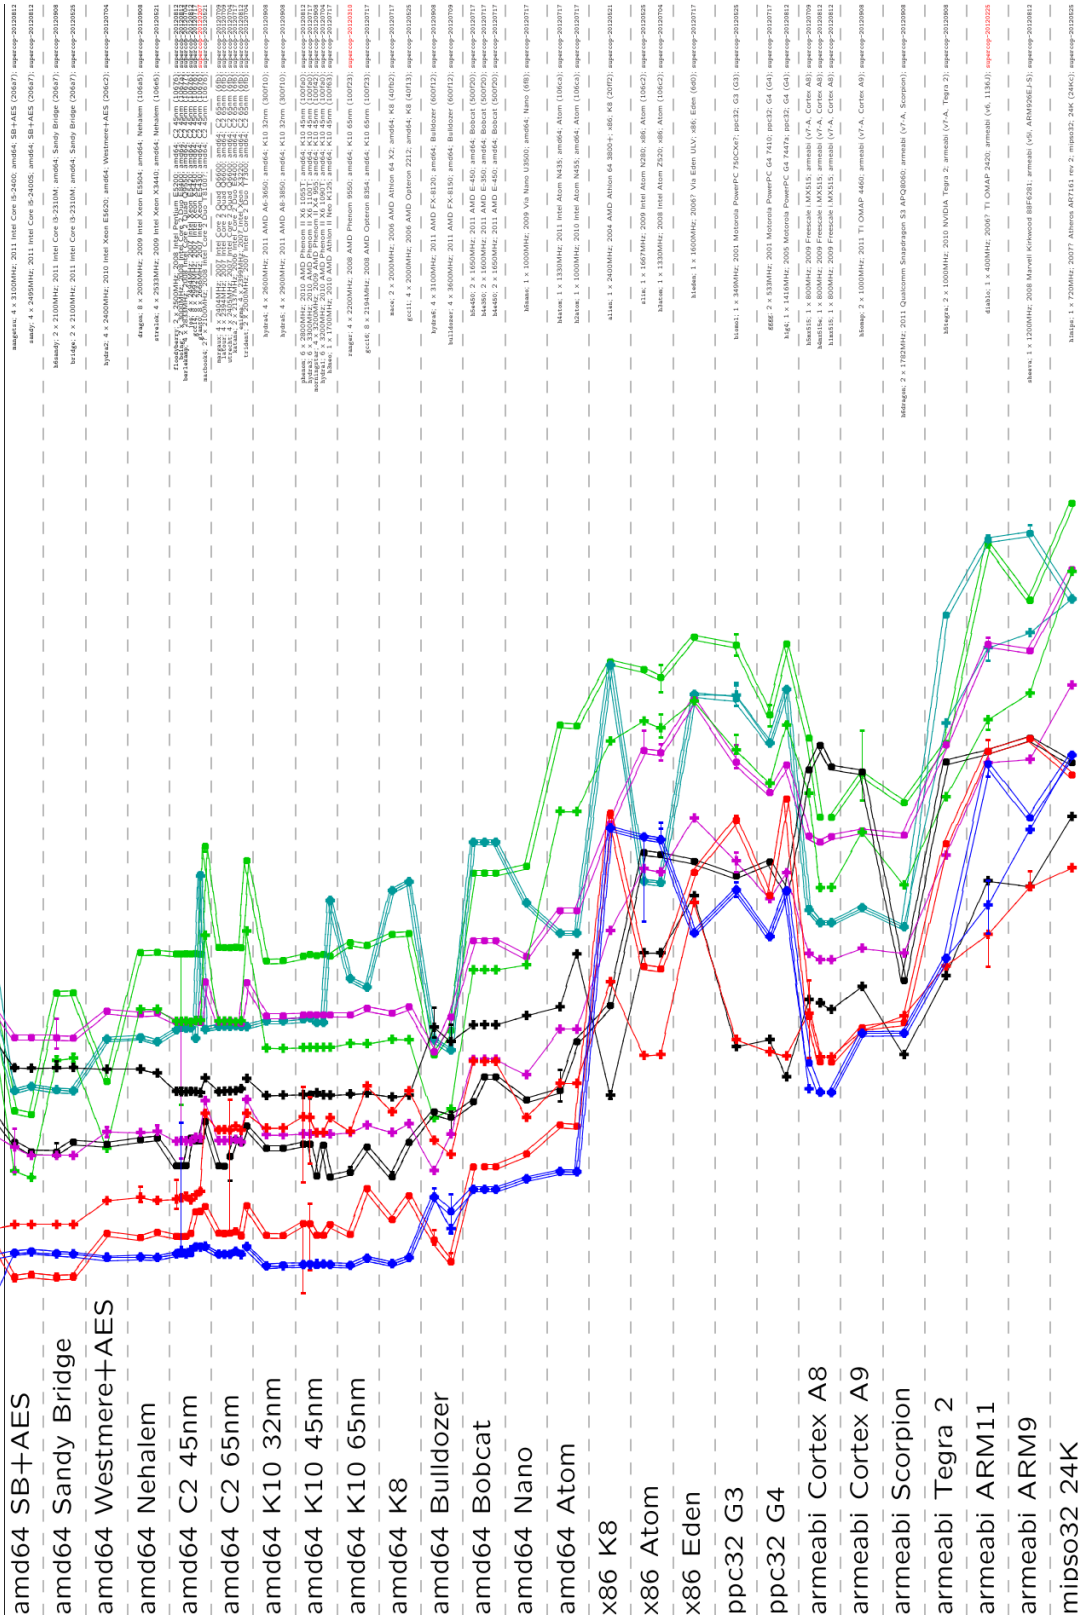
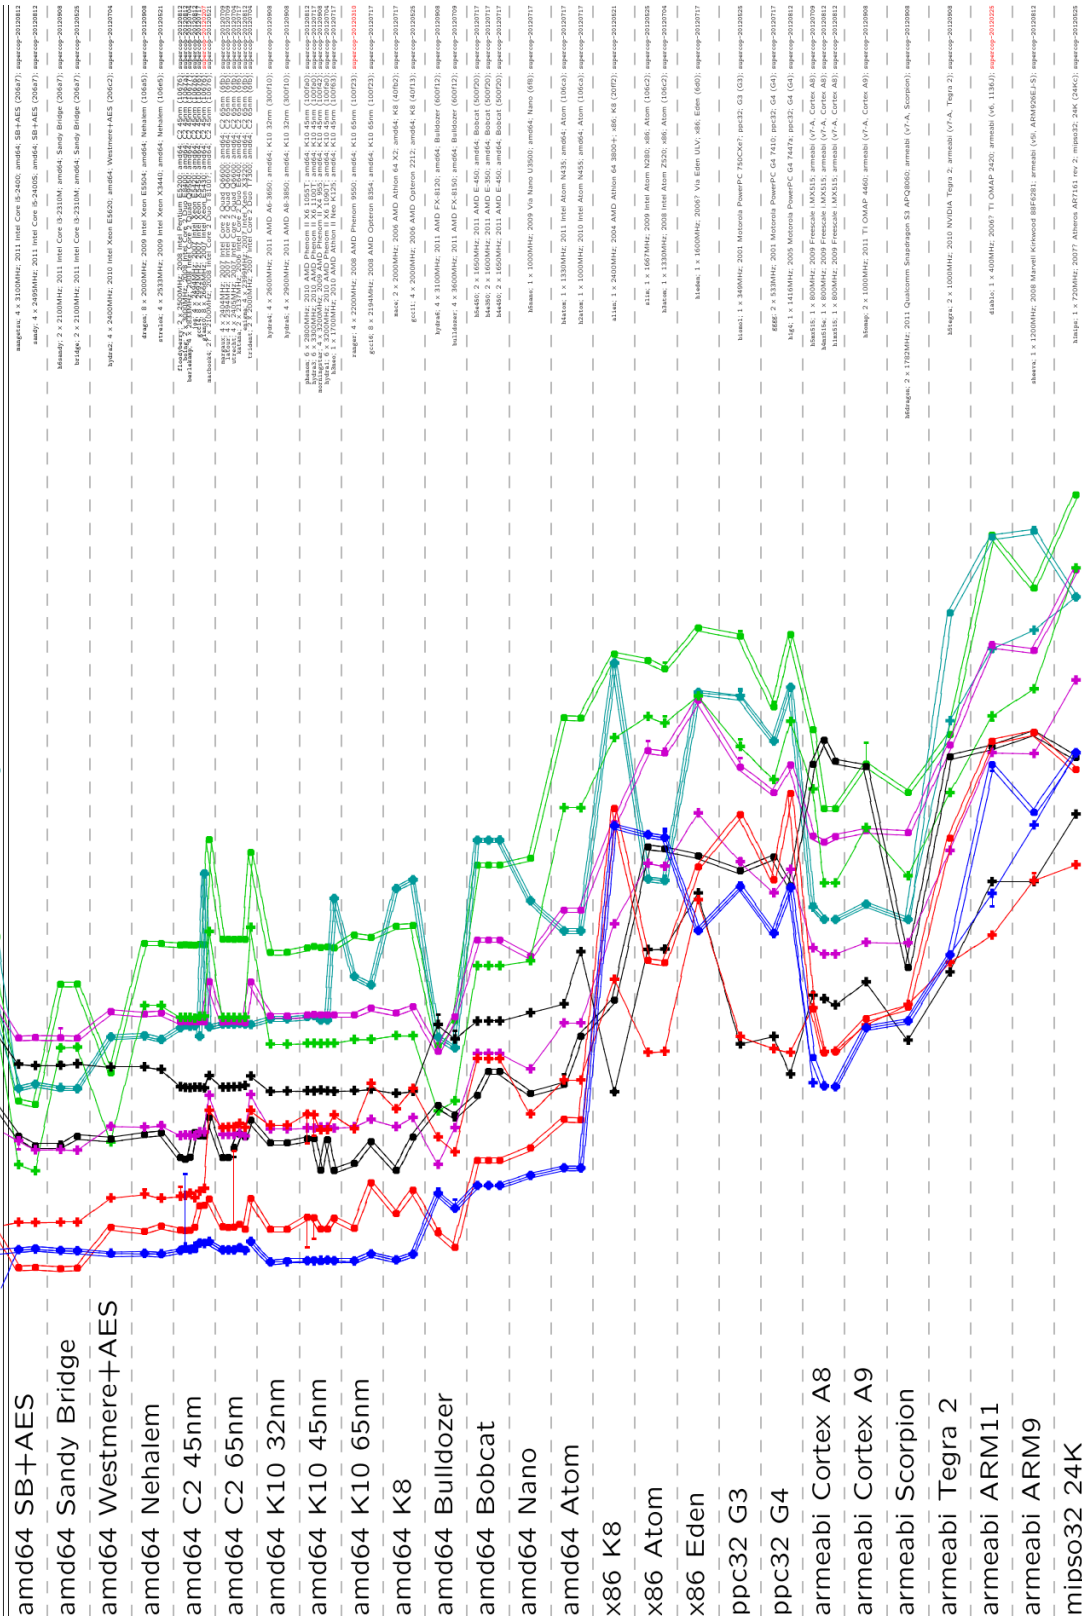


Figure A-1. eBash SHA-3 Finalist Performance for Long Messages

http://bench.cr.yp.to/20120910

crypto_sha3
4096 bytes

sha512 keccak1024round3jh51z groestl1512
sha256 groestl256 round3jh256



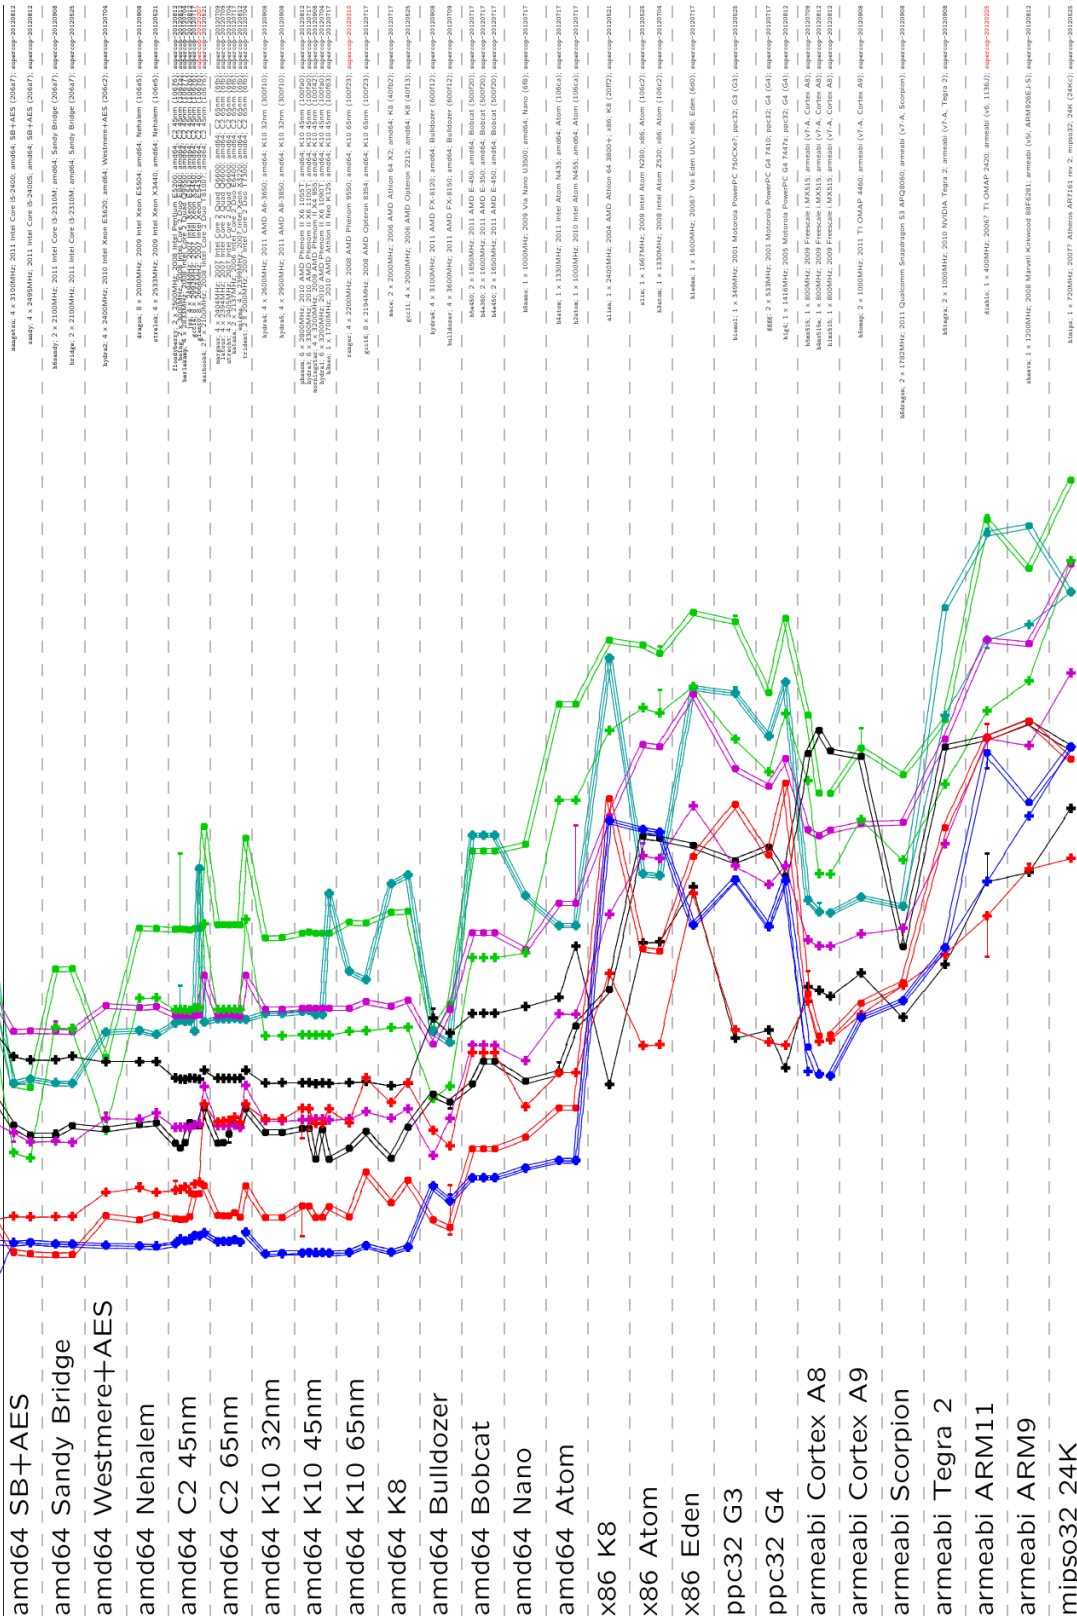
Cycles per byte 4 8 16 32 64 128 256 512 1024 2048 4096

Figure A-2. eBash SHA-3 Finalist Performance for 4096 Byte Messages

crypto_sha3
1536 bytes

sha512
keccak1024
round3jh512
groestl1512

sha256
keccak256
groestl1256
round3jh256



Cycles per byte 4 8 16 32 64 128 256 512 1024 2048 4096

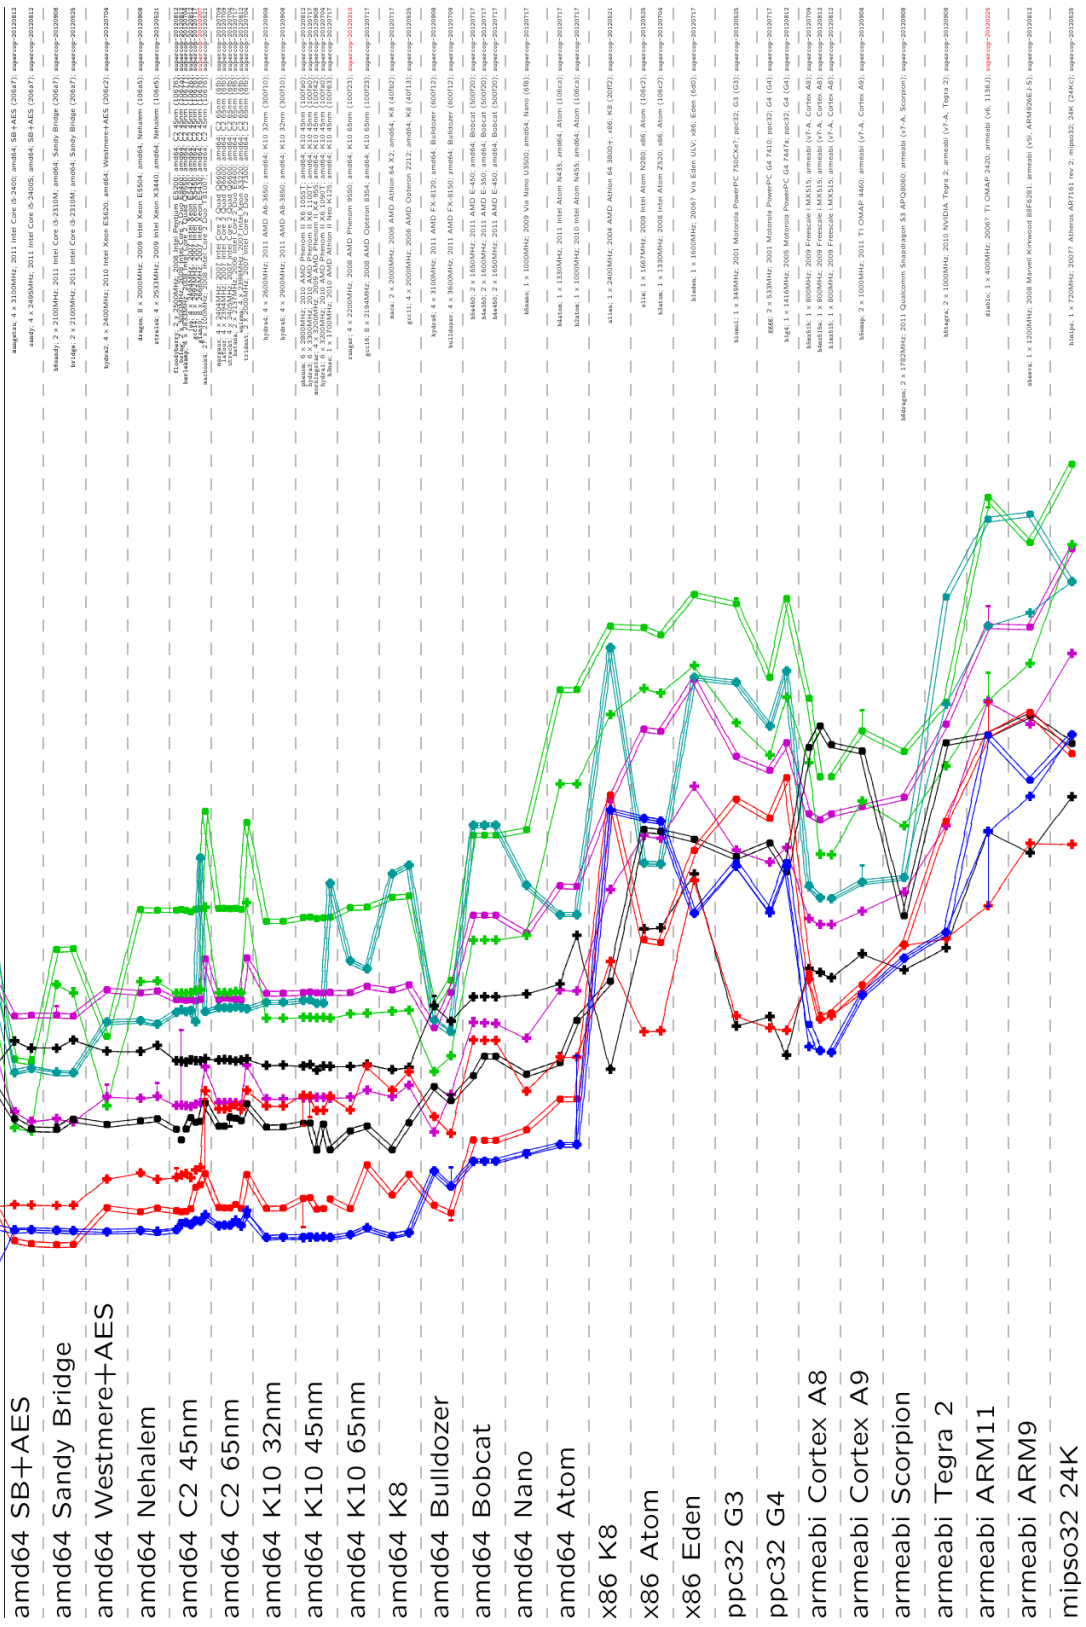
Figure A-3. eBash SHA-3 Finalist Performance for 1536 Byte Messages

http://bench.cr.yp.to
20120910

crypto_sha3
576 bytes

skein512512 skein512256 sha512 sha256 keccak512 keccak256
blake2b blake2s

sha512 sha256 keccak512 keccak256
groestl1512 groestl1256 round3jh256 round3jh256



amd64 SB+AES
amd64 Sandy Bridge
amd64 Westmere+AES
amd64 Nehalem
amd64 C2 45nm
amd64 C2 65nm
amd64 K10 32nm
amd64 K10 45nm
amd64 K10 65nm
amd64 K8
amd64 Bulldozer
amd64 Bobcat
amd64 Nano
amd64 Atom
x86 K8
x86 Atom
x86 Eden
ppc32 G3
ppc32 G4
armeabi Cortex A8
armeabi Cortex A9
armeabi Scorpion
armeabi Tegra 2
armeabi ARM11
armeabi ARM9
mipso32 24K

Cycles per byte 4 8 16 32 64 128 256 512 1024 2048 4096

Figure A-4. eBash SHA-3 Finalist Performance for 576 Byte Messages

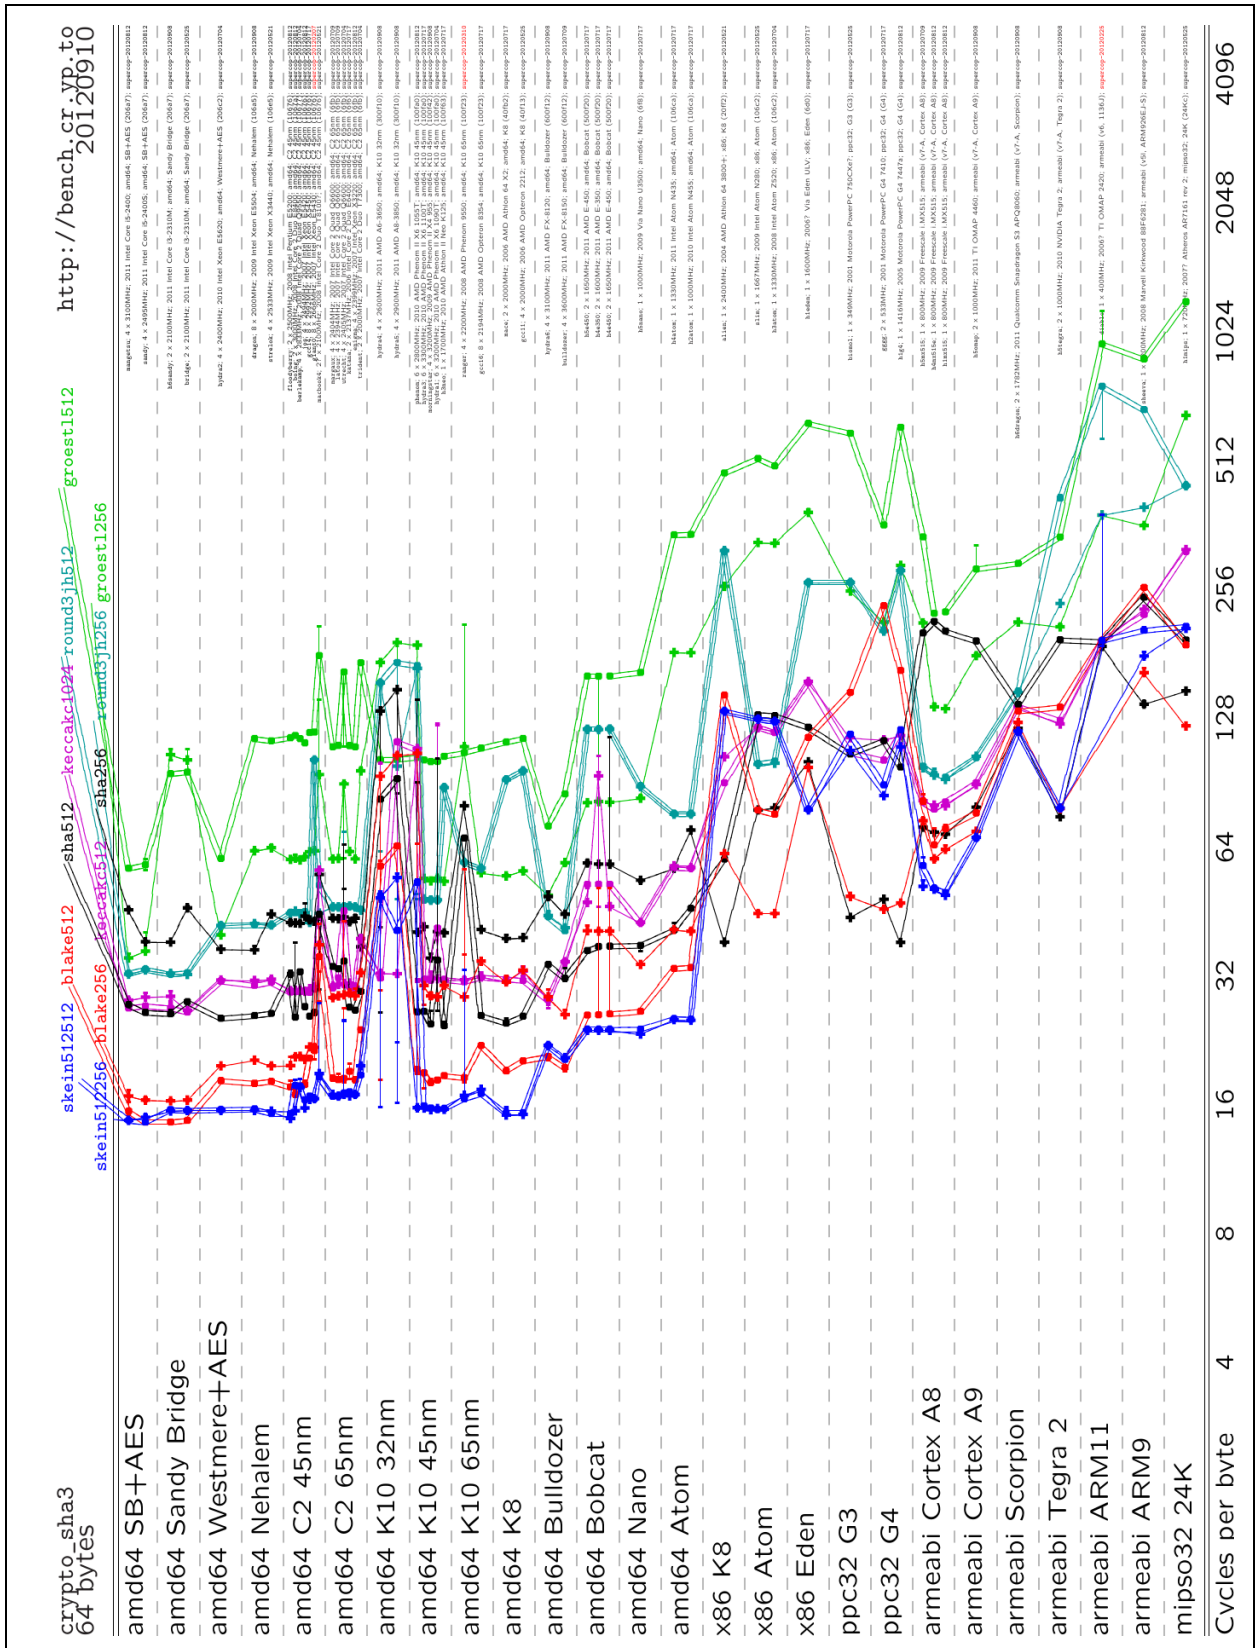


Figure A-5. eBash SHA-3 Finalist Performance for 64 Byte Messages

