BRUTUS: Identifying Cryptanalytic Weaknesses in CAESAR First Round Candidates

Markku-Juhani O. Saarinen

Norwegian University of Science and Technology †
mjos@iki.fi

Abstract. This "half-year" report summarizes our results from security analysis covering all 57 CAESAR first round candidates. We have manually identified security issues with three candidates, two of which are more serious, and these ciphers been withdrawn from the competition. We have developed a testing framework, BRUTUS, to facilitate automatic detection of simple security lapses and susceptible statistical structures across all ciphers. From this testing we have security usage notes on four submissions and statistical notes on a further four. We highlight that some of the CAESAR algorithms pose an elevated risk if employed in real-life protocols due to a class of adaptive chosen plaintext attacks. Although AEADs are often defined (and are best used) as discrete primitives that authenticate and transmit only complete messages, in practice these algorithms are easily implemented in a fashion that outputs observable ciphertext data when the algorithm has not received all of the (attacker-controlled) plaintext. For an implementor, this strategy appears to offer seemingly harmless and compliant storage and latency advantages. If the algorithm uses the same state for secret keying information, encryption, and integrity protection, and the internal mixing permutation is not cryptographically strong, an attacker can exploit the ciphertext-plaintext feedback loop to to reveal secret state information or even keying material. We conclude that the main advantages of exhaustive, automated cryptanalysis is that it acts as a very necessary sanity check for implementations and gives the cryptanalyst insights that can be used to focus more specific attack methods on given candidates.

Keywords: Adaptive Chosen Plaintext Attacks, Authenticated Encryption, CAESAR, BRUTUS, Automated Cryptanalysis.

1 Introduction

Authenticated Encryption with Associated Data (AEAD) algorithms provide message confidentiality and integrity protection with a single cryptographic primitive. As such, they offer functionality similar to combining a stream or block cipher with a Message Authentication Code (MAC) on protocol level.

† This research was carried out during the tenure of an ERCIM "Alain Bensoussan" Fellowship Programme, hosted at Department of Telematics, NTNU, Norway.

This two-algorithm approach has been the predominant way of securing messages in popular Internet security protocols since mid-1990's. Its potential problems were identified early by H. Krawczyk and others [21]. Still, current TLS 1.2 [10] mandates support only for the `TLS_RSA_WITH_AES_128_CBC_SHA` cipher suite, which combines AES [27] in CBC [11] Confidentiality Mode with SHA-1 [30] hash algorithm in HMAC [29] Message Authentication mode and a TLS-specific padding scheme. Similar approaches have been taken by other popular security protocols such as IPSec [19, 20] and SSH [51]. This separation has been exploited by numerous real-life attacks [9, 32, 44].

When Authenticated Encryption techniques such as GCM [28] are used, most problems related to intermixing of two separate algorithms (such as padding) are removed. Furthermore, AES-GCM works in a single pass, resulting in increased throughput and a decreased implementation footprint. AES-GCM has rapidly replaced older methods in practical usage. It is endorsed and effectively enforced for U.S. and Allied National Security Systems [8]. AES-GCM has been adopted for use in many protocols, including TLS, SSH, and IPSec [6, 17, 40].

However, GCM is widely seen as an unsatisfactory standard with brittle security assurances [31] and therefore a new NIST-sponsored competition, CAESAR (Competition for Authenticated Encryption: Security, Applicability, and Robustness) was launched in 2014 [7]. The CAESAR competition has multiple stages or "elimination rounds." The call for algorithms resulted in 57 first round candidates.

Structure of this paper and our contributions. We give a description of AEADs that most CAESAR candidates conform to in Section 2. We started our evaluation by getting to know the voluminous supplied documentation, which led to cryptanalytic results on three candidates (Section 3). We then describe the development of our framework for automated cryptanalysis, BRUTUS, in Section 4, together with security usage notes obtained. A key observation which may not have been fully considered by all submitters is the non-atomic nature of AEADs in real life, which is captured in the notion of adaptive chosen plaintext attacks (Section 5). The candidates can be classified according to their robustness against adaptive chosen plaintext attacks, which generally do not apply to AES-GCM. This is done in Section 6, followed by conclusions in Section 7.

2 Authenticated Encryption with Associated Data

Most CAESAR Authenticated Encryption Algorithm with Associated Data (AEAD) algorithms have the following inputs:

K : A secret, shared confidentiality and integrity key.
N : Public Nonce or Initialization Vector. Optionally transmitted.
P : Message payload, for which both confidentiality and integrity is protected.
A : Associated "header" data. This data is only authenticated.[1]

[1] Associated Data A may be transmitted unencrypted or implicitly known to both parties (meta information such message sequence numbers, endpoint identities).

The AEAD transform will output a single binary string C that contains additional entropy bits for detection of modifications:

$$\mathsf{AEAD}(K, N, P, A) \rightarrow C. \tag{1}$$

The inverse transform will only return the original message payload P if correct values for K, N, A, and C are supplied:

$$\mathsf{AEAD}^{-1}(K, N, C, A) \rightarrow P \textit{ or } \mathsf{FAIL}. \tag{2}$$

We may semi-formally characterize the security requirements for AEAD and AEAD^{-1} which are relevant to this work as follows:

Confidentiality. Even if a large number of chosen (N, P, A) (with non-repeating N) can be supplied by an attacking algorithm to an encryption oracle $\mathsf{AEAD}(?, N, P, A) \rightarrow C$, it should be infeasible to distinguish the corresponding outputs C from equal-length random strings.

Integrity. It should be infeasible to create any new set (N, C, A) that would *not* output $\mathsf{AEAD}^{-1}(?, N, C, A) = \mathsf{FAIL}$ for an unknown key, even if a very large number of valid (N, P, C, A) sets for that secret key are available.

More trivial security properties follow from these requirements. Each submission was free to define what "infeasible" in their particular case means. For the confidentiality requirement this is traditionally expected to mean effort commensurate with an exhaustive search for the secret key K. The forgery effort (integrity goal) depends on the size of authentication variable (message expansion from P to C), but can be defined to be lower. For example AES-GCM archives a significantly lower level of integrity protection than information theoretically expected [33, 36]. As CAESAR is a cryptographic *competition*, we may consider all such suboptimal features to be relative weaknesses.

3 Manual Cryptanalysis

CAESAR candidates came in many shapes and sizes. We refer to the Authenticated Encryption Zoo web site for one classification and current status of each one of the candidates.[2] Here's a rough breakdown:

- 8 are clearly based on the Sponge construction.
- 9 are somehow constructed from AES components.
- 19 are AES modes of operation.
- 21 are based on other design paradigms or are entirely ad hoc.

A group of proposals cannot be even evaluated according to established cryptologic criteria and we sidestep those in this report. We trust that the CAESAR selection committee will arrive at the same conclusion for the second round.

[2] `https://aezoo.compute.dtu.dk/`

We spent some time familiarizing ourselves with the substantial amount of technical documentation after it was released in March 2014. Based on the specifications alone, we identified clear cryptanalytic problems with three candidates:

1. **PAES** [50] suffered from rotational cryptanalytic flaws as round constants were not used. Similar observations were made simultaneously by Sasaki-Wang [41] and Jean-Nikolić [18] teams. PAES has been withdrawn from the CAESAR competition.
2. **HKC** [15] was found to suffer from an almost linear authentication function, which could be used for high-probability message forgeries. HKC has been withdrawn from the CAESAR competition.
3. **iFeed[AES]** [52]. We offered criticism towards this proposal the authentication tag depends only on the last block of the plaintext. [3]

In addition to these, we have identified more minor problems with other proposals which have been addressed privately or via the CAESAR mailing list.

4 Exhaustive Methodology: The BRUTUS Framework

By June 2014, most of the 57 teams had submitted reference implementations for their candidates. Many of these candidates had multiple parameter choices, bringing the total number of implementations to over 160.

The implementations were integrated into the SUPERCOP[4] speed testing framework by D. Bernstein. In addition to very rudimentary coherence testing, the sole functionality of SUPERCOP is in performance measurement. SUPERCOP is not very well suited for statistical testing or other experimental work.

Development Process. We decided to build our own testing framework which would allow more rapid experimentation. We lifted the reference implementations from the SUPERCOP framework as we had no use for it. Our BRUTUS[5] toolkit compiles each reference implementation into a dynamically linked library that can be loaded "on the fly" into an arbitrary experimentation program. The standard test module performs coherence testing, speed tests, and generates test vectors known as Known Answer Tests (KATs). Interfacing with arbitrary languages can be archived via small native components.

Due to the disappointingly poor quality of some of the code (even from some prominent cryptologic security teams), many implementations had to be corrected to fix memory leaks and other elementary errors that affected stability of experimentation. We avoided modifying the mathematical structure of the implementations even when it appeared to contradict the supplied documentation. BRUTUS is intended purely as a research and experimentation tool.

[3] Similar issues apply to some other proposals such as OCB[23] and OTR[24], which restricts their usage in protocols where some level of collision resistance is expected. Furthermore iFeed decryption cannot be parallelized (thereby forming a bottleneck) and hence we see no real advantages in its use over GCM.

[4] `http://bench.cr.yp.to/supercop.html`

[5] Framework only, no statistical tools: `https://github.com/mjosaarinen/brutus/`

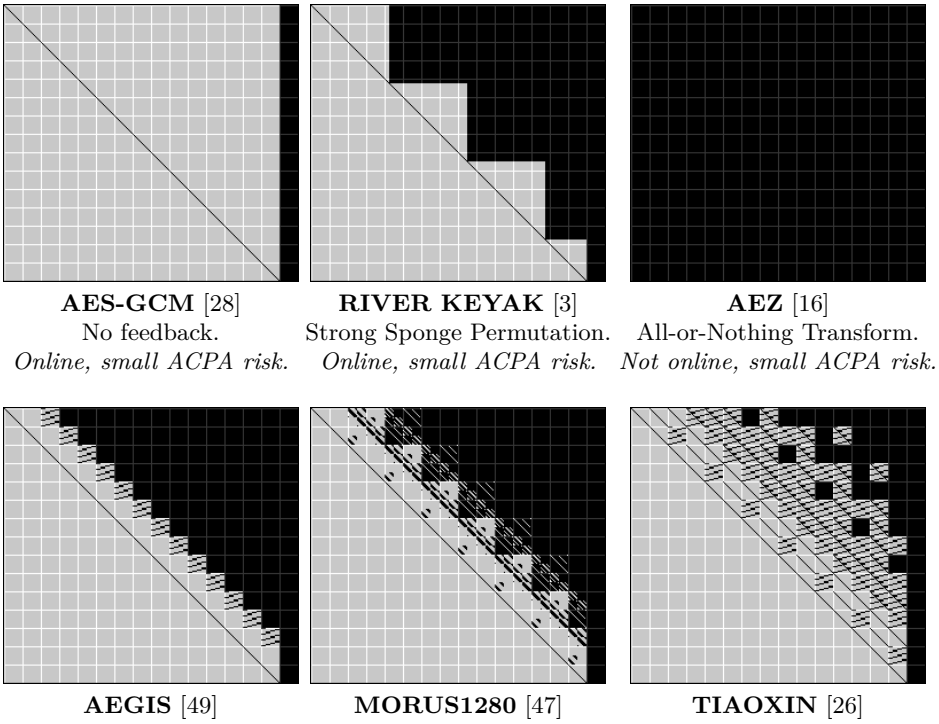

Fig. 1. Visualization of feedback properties of some CAESAR candidates. Here each pixel represents a single byte. Grid lines are every 16 bytes (128 bits). The Y coordinate is the single plaintext byte change location offset. Each pixel line represents 256 bytes of ciphertext difference, with affected ciphertext bytes darkened. The authentication tag is usually seen as a bar on the right side; those bytes are affected by any change. The "ripples" on the lower three diagrams are one indication of inconsistent mixing.

Identifying Ciphers and Modes. An interesting advantage gained from having a coherent and easy interface for all ciphers is that an "identifying gallery"[6] of proposed modes and ciphers can be constructed. This allows black-box identification of ciphers in some cases. The diagrams are independent of secret keying information. Figure 1 shows some members of this gallery.

Implementability and side channels. It is clear that some proposals are poorly suited for hardware-only implementation. For example, any algorithm actually requiring `malloc()` dynamic memory allocation – which in itself is a side channel security headache – is difficult to implement in hardware. How this

[6] `https://mjos.fi/aead_feedback/`

will be addressed is left to the CAESAR committee as hardware implementations are not expected before the second round. We have implemented some proposals in FPGA and proposed SÆHI, a sufficiently generic API that allows hybrid software-hardware implementations and is therefore able to cover almost all candidates [39]. BRUTUS is capable of supporting this API.

Efficient implementation of parallelized modes in plain ANSI C is nontrivial. We refer to SUPERCOP results for software performance metrics across a number of implementation targets. Speed-optimized implementations were not even expected for first round candidates, so such comparisons would be unfair (the call was for "readable" implementations).

Security usage notes on various ciphers. From our automated testing we arrived at the following notes:

1. **CMCC** [43] does not use all of its keying material for short messages and therefore a trivial forgery can be made even if a part of the secret key is not known. The author has proposed a tweak.
2. **CALICO** [42] had an extraordinarily long key (32+16 = 48 bytes), which consists of a 32-byte decryption key and a 16-byte MAC key. If you have a false key (with something else in the first key 32 bytes), CALICO will not detect it and will just output nonsense. This can be circumvented in implementations but does violate basic AEAD security expectations. The author withdrew CALICO from the competition earlier.
3. **PAEQ** [4] implementations exhibited a property in which authentication of associated data only (i.e. no payload) did not depend on the supplied nonce at all, leading to replay forgery attacks in case a protocol is sending A only. The authors noted that the specification forbids such messages (but were allowed in actual implementation for compatibility), but are working on a tweak. We encourage such a tweak as this would make the proposal plug-in compatible with AES-GCM in security protocols where signaling frequently demands authentication of metadata only.
4. **YEASv2** [5]. Although it is mentioned the specification, the nonce has only 127 effective bits. The ignored bit is bit 0 of the last of byte of the 16-byte IV sequence. This is an unfortunate selection; if we are using network (big endian) byte order, this is the least significant bit of the nonce. If running sequence numbers are used, every two consecutive messages will have equivalent nonces and security will break.

All of these issues are fairly easy to address. Again we sidestep proposals that do not meet the basic CAESAR call compliance and sanity criteria.

Implementation Security. Based on our cursory code review of the 160+ implementations, our general advice is strongly against using CAESAR reference ciphers as a part of any real-life application requiring stability or security at this stage of competition.

5 Most AEADs are not Atomic

When described in the fashion of Equations 1 and 2, an AEAD transform appears to be an atomic, indivisible operation. Two-pass CAESAR candidates can essentially only be implemented this way. The AEZ [16] and SIV [22] candidates are examples of such "All-or-nothing Transforms" [34].

Due to efficiency and memory conservation reasons, most CAESAR candidates can work in "online" mode where the full plaintext block P is not required for the encryption algorithm to be able to produce some of the ciphertext. This is generally done by dividing the message to uniform-sized message blocks $\mathsf{pad}(M) = M_1 \mid\mid M_2 \mid\mid \cdots \mid\mid M_n$. The AEAD maintains an internal state X which is initialized with some value derived from K and N. This is then iterated over blocks M_i and the final state is subjected to another transformation to produce a MAC tag T.

$X_0 = \mathsf{key}(K, N)$	Initialize state from Key and Nonce.
$X_i = \mathsf{mix}(M_i, X_{i-1})$	Mix message blocks with state, $1 \leq i \leq n$.
$C_i = \mathsf{out}(X_i)$	Ciphertext block derived from state, $1 \leq i \leq n$.
$T = \mathsf{fin}(X_n)$	Finalization – compute the authentication tag.

The ciphertext is constructed as $C = C_1 \mid\mid C_2 \mid\mid \cdots \mid\mid C_n \mid\mid T$. This type of construction allows C_i to be output immediately after M_i is fed into the mixing transform. All Sponge-based [2] constructions and many proposed block cipher modes of operation fall into this category.

The Adaptive Chosen Plaintext Attack also applies AEAD designs which are *not* necessarily based on block ciphers at all. We assume that an attacker can adaptively feed a plaintext block M_i to the cipher as a function of previously observed ciphertext blocks

$$M_i = f_{\mathsf{atk}}(C_1, C_2, \cdots, C_{i-1}). \tag{3}$$

The attacker function f_{atk} can perform some reasonable amount of computation for the feedback operation.

We argue that this is a relevant model offering insights especially to smart card applications and other lightweight applications where an attacker has full control over the communication channel.

The goal of the attacker is to derive information about the internal state X_i. This information can be used in attacks of various degrees of severity:

1. Distinguish or partially predict C_{i+1}.
2. Fully derive X_i; predict all future C_i and T.
3. Derive information about K.

Note that message authentication is not an issue in an adaptive chosen plaintext attack on an AEAD as encryption cannot really fail. The inverse scenario of Equation 3, a chosen ciphertext attack, is less realistic as it would seem to automatically break the definition given by Equation 2. However, this scenario has been considered in the literature [1].

6 CAESAR Candidates and Real-life Protocols: Susceptibility to Adaptive Chosen Plaintext Attacks

In order to integrate a CAESAR AEAD into a real-life protocol such as TLS, SSH, or IPSec, one has to not only define the appropriate ciphersuite identifiers but also usage and formatting mechanisms.

In case of all AEADs, an obvious path of integration is to adopt the mechanisms used for AES-GCM in relevant RFCs: TLS in [40], SSH in [17], and IPSec in [6]. This will allow implementors to essentially "plug in" the algorithms into existing protocol implementation frameworks. In many protocol instances, the ciphers are subjected to adaptive chosen plaintext attacks with relative ease.

Even though the CAESAR call for algorithms[7] was careful to require concrete security claims for full AEAD transforms, the security claims related to this type of attack are not explicitly stated for many ciphers. However, internal mixing qualities of a design offer a direct insight into the robustness of a cipher against adaptive chosen plaintext attacks.

Based on our automated analysis, at least ACORN [46], AEGIS [49], MORUS [47], and TIAOXIN [26] represent significantly elevated adaptive chosen plaintext attack risk. We are currently formalizing our observations, but we note that – as an example – the effective internal state can be trivially forced to be smaller, helping birthday attacks. These proposals have a single state without separation between authentication, confidentiality, or keying state. In this, they are similar to Sponge designs. Indeed, if these had been *labeled* "sponge designs" they could be declared "broken" due to the weakness of their mixing functions. This illustrates the difficulty of security comparisons among candidates.

In many ways these ciphers resemble Helix [14] and Phelix [45], which were proposed as an authenticated stream ciphers a decade ago. These ciphers were attacked in under various assumptions [25, 48]. Another earlier, similar (but lightweight) authenticated design is the Hummingbird cipher [12, 13], which was successfully cryptanalyzed [35, 37].

These ciphers seem to have been created with ad hoc design methods and offer no provable security assurances. This by no means indicates that they cannot be used securely and use of these candidates may be highly justified in many cases as they are among fastest (or, in case of ACORN, smallest) candidates.

In comparison, we offer the following proof sketches for resistance of certain other essential classes of algorithms to adaptive chosen plaintext attacks of this type.

Theorem 1. *AES-GCM is not vulnerable to adaptive chosen plaintext attacks.*

Proof. The Galois/Counter Mode has an essentially independent counter mode and a polynomial-based authentication mechanism. Since the counter mode keystream can be generated *a priori* to encryption, any ciphertext-plaintext feedback will not yield useful information about the internal state of the mode. □

[7] `http://competitions.cr.yp.to/caesar-call.html`

Theorem 2. *Sponge modes with strong permutations such as* DuplexWrap *[3] or* BLNK *[38] are not vulnerable to adaptive chosen plaintext attacks.*

Proof. These modes utilize a cryptographically strong permutation between any two blocks of data and therefore the adaptive attacker has no access to capacity beyond that barrier. □

As there are some proposals that employ various stronger notions of provable security, we make the following general observation:

Observation 1 *Provably secure modes that have two or more passes over data are not vulnerable to adaptive chosen plaintext attacks.*

Figure 1 offers a visualization of Theorems 1 and 2 and the final observation.

7 Conclusions

We have presented a summary of our initial examination and analysis covering all 57 CAESAR first round proposals (we are only presenting results that we have obtained ourselves). As an executive note, we currently strongly recommend against using any of the CAESAR ciphers in real-life applications despite their novelty and often famous authorship.

During manual examination we have identified cryptographic problems with three proposals, two of which have been withdrawn from the competition.

We have described our development of the BRUTUS testing framework which allows tests to be made that automatically cover all candidates. As performance testing was not even required in the first round (and is adequately addressed by the SUPERCOP toolkit), we focused on the structural differences of various candidates. We offer security usage notes for four candidates.

From the BRUTUS automated tests we observe that some candidates offer less than convincing resistance against adaptive chosen plaintext attacks. This is significant since one of the main motivations for the CAESAR competition is to seek secure replacements for the AES-GCM algorithm which is provably secure against this type of attack. Sponge permutation designs and two-pass provably secure modes are also resistant. Such an attack can be mounted with relative ease in conceivable instances of real-life protocols such as TLS, SSH, and IPSec.

Based on our experience, the most valuable output from exhaustive, automated testing across actual cipher implementations is that it catches implementation errors and possible errors in *security usage* – discrepancies between the assumptions of the users of the algorithm and its designers. These often break real-life protocols and applications that utilize encryption algorithms. The insights obtained from statistical testing of (internal) quantities can be used by a cryptanalyst to focus more specific analysis efforts against those candidates that are expected to be vulnerable to a particular method of attack.

References

1. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Mouha, N., and Yasuda, K. How to securely release unverified plaintext in authenticated encryption. IACR ePrint 2014/144, `eprint.iacr.org/2014/144`, February 2014.
2. Bertoni, G., Daemen, J., Peeters, M., and Assche, G. V. Duplexing the sponge: Single-pass authenticated encryption and other applications. In *SAC 2011* (2011), A. Miri and S. Vaudenay, Eds., vol. 7118 of *LNCS*, Springer, pp. 320–337.
3. Bertoni, G., Daemen, J., Peeters, M., Assche, G. V., and Keer, R. V. CAESAR submission: Keyak v1. CAESAR First Round Submission, `competitions.cr.yp.to/round1/keyakv1.pdf`, March 2014.
4. Biryukov, A., and Khovratovich, D. PAEQv1. CAESAR First Round Submission, `competitions.cr.yp.to/round1/paeqv1.pdf`, March 2014.
5. Bosselaers, A., and Vercauteren, F. YAES v2. CAESAR First Round Submission, `competitions.cr.yp.to/round1/yaesv2.pdf`, May 2014.
6. Burgin, K., and Peck, M. Suite B Profile for Internet Protocol Security (IPsec). IETF RFC 6380, October 2011.
7. CAESAR. CAESAR call for submissions. `competitions.cr.yp.to/caesar-call.html`, January 2014.
8. CNSSP. National information assurance policy on the use of public standards for the secure sharing of information among national security systems. CNSS Policy No. 15, October 2012.
9. Degabriele, J. P., and Paterson, K. G. On the (in)security of IPsec in MAC-then-encrypt configurations. In *ACM Conference on Computer and Communications Security* (2010), E. Al-Shaer, A. D. Keromytis, and V. Shmatikov, Eds., ACM, pp. 493–504.
10. Dierks, T., and Rescorla, E. The Transport Layer Security (TLS) Protocol Version 1.2. IETF RFC 5246, August 2008.
11. Dworkin, M. Recommendation for block cipher modes of operation. NIST Special Publication 800-38A, December 2001.
12. Engels, D., Fan, X., Gong, G., Hu, H., and Smith, E. M. Hummingbird: Ultra-lightweight cryptography for resource-constrained devices. In *Financial Cryptography and Data Security, FC 2010 Workshops* (2010), R. Sion, R. Curtmola, S. Dietrich, A. Kiayias, J. M. Miret, K. Sako, and F. Sebé, Eds., vol. 6054 of *LNCS*, Springer, pp. 3–18.
13. Engels, D., Saarinen, M.-J. O., Schweitzer, P., and Smith, E. M. The Hummingbird-2 lightweight authenticated encryption algorithm. In *RFIDSec '11* (2011), A. Juels and C. Paar, Eds., vol. 7055 of *LNCS*, Springer, pp. 19–31.
14. Ferguson, N., Whiting, D., Schneier, B., Kelsey, J., Lucks, S., and Kohno, T. Helix: Fast encryption and authentication in a single cryptographic primitive. In *FSE 2003* (2003), T. Johansson, Ed., vol. 2887 of *LNCS*, Springer, pp. 330–346.
15. Henricksen, M., Kiyomoto, S., and Lu, J. The HKC authenticated stream cipher (ver.1). CAESAR First Round Submission, `competitions.cr.yp.to/round1/hkcv1.pdf`, March 2014.
16. Hoang, V. T., Krovetz, T., and Rogaway, P. AEZ v1: Authenticated-Encryption by Enciphering. CAESAR First Round Submission, `competitions.cr.yp.to/round1/aezv1.pdfl`, March 2014.
17. Igoe, K. Suite B Cryptographic Suites for Secure Shell (SSH). IETF RFC 6239, May 2011.

18. Jean, J., and Nikolić, I. Using AES round symmetries to distinguish PAES. `www1.spms.ntu.edu.sg/~syllab/m/images/6/6e/Using_AES_Round_Symmetries_to_Distinguish_PAES.pdf`, March 2014.
19. Kent, S. IP authentication header. IETF RFC 4302, December 2005.
20. Kent, S. IP encapsulating security payload (ESP). IETF RFC 4303, December 2005.
21. Krawczyk, H. The order of encryption and authentication for protecting communications (or: How secure is SSL?). In *CRYPTO 2001* (2001), J. Kilian, Ed., vol. 2139 of *LNCS*, Springer, pp. 310–331.
22. Krovetz, T. HS1-SIV (v1). CAESAR 1st Round, `competitions.cr.yp.to/round1/hs1sivv1.pdf`, March 2014.
23. Krovetz, T., and Rogaway, P. OCB (v1). CAESAR First Round Submission, `competitions.cr.yp.to/round1/ocbv1.pdf`, March 2014.
24. Minematsu, K. AES-OTR v1. CAESAR First Round Submission, `competitions.cr.yp.to/round1/aesotrv1.pdf`, March 2014.
25. Muller, F. Differential attacks against the helix stream cipher. In *FSE 04* (2004), B. Roy and W. Meier, Eds., vol. 3017 of *LNCS*, Springer, pp. 94–108.
26. Nikolić, I. Tiaoxin-346, Version 1.0. CAESAR First Round Submission, `competitions.cr.yp.to/round1/tiaoxinv1.pdf`, March 2014.
27. NIST. Advanced Encryption Standard (AES). Federal Information Processing Standards Publication FIPS 197, November 2001.
28. NIST. Recommendation for block cipher modes of operation: Galois/counter mode (GCM) and GMAC. NIST Special Publication 800-38D, 2007.
29. NIST. The Keyed-Hash Message Authentication Code (HMAC). Federal Information Processing Standards Publication FIPS 198-1, July 2008.
30. NIST. Secure Hash Standard (SHS). Federal Information Processing Standards Publication FIPS 180-4, March 2012.
31. NIST VCAT. NIST Cryptographic Standards and Guidelines Development Process: Report and Recommendations of the Visiting Committee on Advanced Technology of the National Institute of Standards and Technology, July 2014.
32. Paterson, K. G., and Yau, A. K. L. Cryptography in theory and practice: The case of encryption in IPsec. In *EUROCRYPT 2006* (2006), S. Vaudenay, Ed., vol. 4004 of *LNCS*, Springer, pp. 12–29.
33. Procter, G., and Cid, C. On weak keys and forgery attacks against polynomial-based MAC schemes. In *FSE '13* (2013), S. Moriai, Ed., vol. 8424 of *LNCS*, Springer, pp. 287–304.
34. Rivest, R. All-or-nothing encryption and the package transform. In *FSE 97* (1997), E. Biham, Ed., vol. 1267 of *LNCS*, Springer, pp. 210–218.
35. Saarinen, M.-J. O. Cryptanalysis of Hummingbird-1. In *FSE 2011* (2011), A. Joux, Ed., vol. 6733 of *LNCS*, Springer, pp. 328–341.
36. Saarinen, M.-J. O. Cycling attacks on GCM, GHASH and other polynomial MACs and hashes. In *FSE 2012* (2012), A. Canteaut, Ed., vol. 7549 of *LNCS*, Springer, pp. 216–225.
37. Saarinen, M.-J. O. Related-key attacks against full Hummingbird-2. In *FSE 2013* (2013), S. Moriai, Ed., vol. 8424 of *LNCS*, Springer, pp. 467–482.
38. Saarinen, M.-J. O. Beyond modes: Building a secure record protocol from a cryptographic sponge permutation. In *CT-RSA 2014* (2014), J. Benaloh, Ed., vol. 8366 of *LNCS*, Springer, pp. 270–285.
39. Saarinen, M.-J. O. Simple AEAD Hardware Interface (SÆHI) in a SoC: Implementing an On-Chip Keyak/WhirlBob Coprocessor. In *TrustED '14, 03 November 2014 Scottsdale AZ, USA* (November 2014), ACM (to appear).

40. Salter, M., and Housley, R. Suite B Profile for Transport Layer Security (TLS). IETF RFC 6460, January 2012.
41. Sasaki, Y., and Wang, L. A practical universal forgery attack against paes-8. IACR ePrint 2014/218, eprint.iacr.org/2014/218, March 2014.
42. Taylor, C. The Calico family of authenticated ciphers, version 8. CAESAR First Round Submission, competitions.cr.yp.to/round1/calicov8.pdf, March 2014.
43. Trostle, J. AES-CMCC v1.1. CAESAR First Round Submission, competitions.cr.yp.to/round1/aescmccv11.pdf, March 2014.
44. Vaudenay, S. Security flaws induced by CBC padding - applications to SSL, IPSEC, WTLS ... In *EUROCRYPT 2002* (2002), L. R. Knudsen, Ed., vol. 2332 of *LNCS*, Springer, pp. 534–546.
45. Whiting, D., Schneier, B., Lucks, S., and Muller, F. Phelix - fast encryption and authentication in a single cryptographic primitive. ECRYPT Stream Cipher Project Report 2005/027, 2005. www.schneier.com/paper-phelix.html.
46. Wu, H. ACORN: A Lightweight Authenticated Cipher (v1). CAESAR First Round Submission, competitions.cr.yp.to/round1/acornv1.pdf, March 2014.
47. Wu, H., and Huang, T. The Authenticated Cipher MORUS (v1). CAESAR First Round Submission, competitions.cr.yp.to/round1/morusv1.pdf, March 2014.
48. Wu, H., and Preneel, B. Differential-linear attacks against the stream cipher phelix. eSTREAM preprint http://www.ecrypt.eu.org/stream/papersdir/2006/056.pdf, December 2006.
49. Wu, H., and Preneel, B. AEGIS: A Fast Authenticated Encryption Algorithm (v1). CAESAR First Round Submission, competitions.cr.yp.to/round1/aegisv1.pdf, March 2014.
50. Ye, D., Wang, P., Hu, L., Wang, L., Xie, Y., Sun, S., and Wang, P. PAES v1: Parallelizable authenticated encryption schemes based on AES round function. CAESAR First Round Submission, competitions.cr.yp.to/round1/paesv1.pdf, March 2014.
51. Ylönen, T., and Lonvick, C. The secure shell (SSH) transport layer protocol. IETF RFC 4253, January 2006.
52. Zhang, L., Wu, W., Sui, H., and Wang, P. iFeed[AES] v1. CAESAR First Round Submission, competitions.cr.yp.to/round1/ifeedaesv1.pdf, March 2014.

Document version 20141023075700.