

WHIRLBOB, the Whirlpool based Variant of STRIBOB

Lighter, Faster, and Constant Time

Markku–Juhani O. Saarinen¹ and Billy Bob Brumley²

¹ Queen’s University Belfast, UK *

mjos@iki.fi

² Tampere University of Technology, Finland

billy.brumley@tut.fi

Abstract. WHIRLBOB, also known as STRIBOB₂, is an AEAD (Authenticated Encryption with Associated Data) algorithm derived from STRIBOB₁ and the Whirlpool hash algorithm. WHIRLBOB/STRIBOB₂ is a second round candidate in the CAESAR competition. As with STRIBOB₁, the reduced-size Sponge design has a strong provable security link with a standardized hash algorithm. The new design utilizes only the LPS or ρ component of Whirlpool in flexibly domain-separated BLNK Sponge mode. The number of rounds is increased from 10 to 12 as a countermeasure against Rebound Distinguishing attacks. The 8×8 -bit S-Box used by Whirlpool and WHIRLBOB is constructed from 4×4 -bit “MiniBoxes”. We report on fast constant-time Intel SSSE3 and ARM NEON SIMD WHIRLBOB implementations that keep full miniboxes in registers and access them via SIMD shuffles. This is an efficient countermeasure against AES-style cache timing side-channel attacks. Another main advantage of WHIRLBOB over STRIBOB₁ (and most other AEADs) is its greatly reduced implementation footprint on lightweight platforms. On many lower-end microcontrollers the total software footprint of π +BLNK = WHIRLBOB AEAD is less than half a kilobyte. We also report an FPGA implementation that requires 4,946 logic units for a single round of WHIRLBOB, which compares favorably to 7,972 required for Keccak / Keyak on the same target platform. The relatively small S-Box gate count also enables efficient 64-bit bitsliced straight-line implementations. We finally present some discussion and analysis on the relationships between WHIRLBOB, Whirlpool, the Russian GOST Streebog hash, and the recent draft Russian Encryption Standard Kuznyechik.

Keywords: WHIRLBOB, STRIBOB₁, Authenticated Encryption, Sponge Designs, Timing Attacks, Whirlpool, Streebog, CAESAR Competition.

1 Introduction

STRIBOB₂, or WHIRLBOB, is an Authenticated Encryption with Associated Data (AEAD) algorithm and a CAESAR (“Competition for Authenticated Encryption: Security, Applicability, and Robustness”) [21] competition Second Round Candidate.

* Much of this work was performed while the author was visiting Norwegian University of Science and Technology (Trondheim, Norway) and Tampere University of Technology (Finland).




Fig. 1. A simplified view of a Sponge-based AEAD. Padded Secret Key, Nonce, and Associated Authenticated Data - all represented by d_u words - are first “absorbed” into the state. The π permutation is then also used to encrypt data p_i into ciphertext c_i (or vice versa) and finally to “squeeze” out a Message Authentication Code h_i .

AEAD algorithms and modes such as GCM [45] provide both confidentiality and integrity protection for messages in a single step, thus eliminating the need for a separate MAC algorithm such as HMAC [46]. This has clear advantages for performance and implementation footprint.

WHIRLBOB uses STRIBOBr1’s BLNK Sponge AEAD mode and parameters without modification. Outside the CAESAR context, BLNK can be also used in a wider set of applications, even to build entire secure lightweight protocol suites [54]. A sponge mode requires only a single cryptographic component; an unkeyed cryptographic permutation π (See Figure 1). As with other provable Sponge modes, we assume that π is indistinguishable from a random permutation. This work focuses on π permutation design and implementation – for BLNK padding details and analysis we refer to [32,53,55,56].

The STRIBOBr1 CAESAR [55] candidate was derived from the Russian GOST hash standard Streebog [26]. In close examination Streebog appears to be modeled after the Whirlpool hash [4], with substantial modifications. However, STRIBOBr1 and WHIRLBOB only differ in the particular numerical selections for tables C , S , and L . These components, $L \circ P \circ S$ or the “LPS permutation” is derived directly from that of Whirlpool for WHIRLBOB. The program code of 64-bit reference implementations is essentially equivalent for both algorithms. Both STRIBOBr1 and WHIRLBOB have 12 rounds.

We show that the particular structure of the Whirlpool components allows WHIRLBOB to have much more efficient SIMD, constant-time, lightweight, and hardware implementations. One of our aims is to allow the same implementation core (such as a special instruction or coprocessor of a SoC [55]) to be also used for unkeyed hashing according to the Whirlpool standard. This is useful in applications that also require efficient standards-compliant certificate signature processing.

The corresponding standardized, Miyagushi-Preneel hash functions Streebog and Whirlpool require two (or more) times as much as state and processes data in bigger chunks when compared to STRIBOBr1 and WHIRLBOB. Our BLNK Sponge mode also supports randomized hashing and MACing without encryption. Our Sponge variants are slightly faster than the original hashes, yet have a provable security relation. All security parameters remain unmodified from STRIBOBr1.

2 WHIRLBOB

As with STRIBOBr1, the state consists of $b = 512$ bits, split in our BLNK mode as:

- $r = 256$ - bit *rate* “block size”, which directly interacts with output and input.
- $c \approx 254$ - bit *capacity*, which is the secret state. Some bits are lost to padding.

These two halves are mixed together by a keyless, random-indistinguishable permutation π . According to theorems such as those given in [32,55] this is sufficient for $k = 192$ - bit secret key security level when less than 2^{64} bits are processed under same key and nonce pair. For the CAESAR variant the nonce length is fixed at $n = 128$ bits.

Despite having almost equivalent speed and size on generic 64-bit platforms, the size and performance characteristics of STRIBOBr1 and WHIRLBOB differ significantly on various implementation platforms such as FPGA, low-end microcontrollers, SIMD systems, and in bitslicing implementations.

WHIRLBOB’s permutation π is indeed highly similar to AES. In case of STRIBOBr1, the “Russian 512-bit block AES” permutation had to be somewhat laboriously uncovered from the structure (see Section 5.3), but the particularities and history of Whirlpool make the connection immediately clear.

2.1 Structure of the π Permutation

The computation of π follows almost exactly the operation of the internal key schedule of Whirlpool 3.0 [4]. The only modification is that the number of rounds is increased from $R = 10$ to $R = 12$ for extra security margin against Rebound Distinguisher attacks [35,36].³

To compute $\pi(x_0) = x_{12}$ we iterate the LPS = $L \circ P \circ S$ composite mixing function with round constants C_r . For rounds $0 \leq r < 12$:

$$x_{r+1} = L(P(S(x_r))) \oplus C_r. \quad (1)$$

If we use the AES-style notation of Whirlpool, S is equivalent to `SubBytes`, P corresponds to `ShiftColumns`, L to `MixRows`, followed by `AddRoundKey`.

We write the 512-bit state as a matrix $M[0 \dots 7][0 \dots 7]$ of 8×8 bytes, which can be serialized to a byte vector as $\text{vec}[8i + j] = M[i][j]$.

S : `SubBytes`: Each one of the 64 bytes in the state is substituted using the (singular) 8×8 - bit S-Box described in Section 2.2. For $0 \leq i, j < 8$

$$M'[i][j] \leftarrow S(M[i][j]). \quad (2)$$

P : `ShiftColumns`: A byte shuffle. For $0 \leq i, j < 8$

$$M'[(i + j) \bmod 8][j] \leftarrow M[i][j]. \quad (3)$$

³ These attacks would not be directly applicable with the BLNK mode anyway. This is because the attacker can never access more than r bits of the internal state.

L : MixRows: Each of the 8 row vectors

$$V_i = (M[i][0], M[i][1], \dots, M[i][7]) \quad (4)$$

is individually multiplied by a circulant, low-weight 8×8 MDS matrix in the finite field $\text{GF}(2^8)$ characterized by primitive polynomial $p(x) = x^8 + x^4 + x^3 + x^2 + 1$.

$$V'_i = V_i \cdot \begin{pmatrix} 01 & 01 & 04 & 01 & 08 & 05 & 02 & 09 \\ 09 & 01 & 01 & 04 & 01 & 08 & 05 & 02 \\ 02 & 09 & 01 & 01 & 04 & 01 & 08 & 05 \\ 05 & 02 & 09 & 01 & 01 & 04 & 01 & 08 \\ 08 & 05 & 02 & 09 & 01 & 01 & 04 & 01 \\ 01 & 08 & 05 & 02 & 09 & 01 & 01 & 04 \\ 04 & 01 & 08 & 05 & 02 & 09 & 01 & 01 \\ 01 & 04 & 01 & 08 & 05 & 02 & 09 & 01 \end{pmatrix}. \quad (5)$$

C_r : AddRoundKey: The key schedule operation is effectively equivalent to the one used by Whirlpool’s “internal block cipher” W . Blocks of eight bytes from the S-Box are used round keys C_r for the first row. For round $0 \leq r < 12$, $0 \leq j < 8$:

$$M'[0][j] \leftarrow M[0][j] \oplus S(8r + j) \quad (6)$$

Rest of the rows are unaffected by C_r . For $1 \leq i < 8$, $0 \leq j < 8$:

$$M'[i][j] \leftarrow M[i][j]. \quad (7)$$

We offer the listing of Appendix A as WHIRLBOB v1.0 π reference implementation. Whirlpool ISO Standard trace test vectors can be used to verify the correctness of this π implementation up to $r = 10$ [30]. One simply observes the keying “line” of these traces and ignores the encryption “line”.

2.2 The S-Box

The Whirlpool and WHIRLBOB 8×8 -bit S-Box design utilizes three 4×4 -bit “mini-boxes” given in Table 1: E , E^{-1} , and R . Figure 2 shows how these are used to construct the 8×8 -bit S-Box.

This computation can even be performed on the fly on 4-bit microcontrollers. FPGA implementations save a significant number of LUTs by explicitly utilizing the 4-bit structure rather than implementing a general 8×8 lookup table. These small S-Boxes can often fit into SIMD registers and accessed via constant-time shifts or shuffles, thus enabling implementations resistant to timing attacks.

2.3 BLNK Mode

The padding details and operation of BLNK Sponge mode for WHIRLBOB and STRIBOBr1 are equivalent. Please see [55] for details. Appendix B.2 contains test vectors for running the WHIRLBOB π in CAESAR-compliant AEAD mode.

Table 1. Three 4×4 miniboxes that are used to build the 8×8 S-Box in Whirlpool 3.0 and WHIRLBOB 1.0.

x	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
$E(x)$	1	B	9	C	D	6	F	3	E	8	7	4	A	2	5	0
$E^{-1}(x)$	F	0	D	7	B	E	5	A	9	2	C	1	3	4	8	6
$R(x)$	7	C	B	D	E	4	9	F	6	3	8	A	2	5	1	0

3 Implementation

The entire byte-oriented implementation of π fits onto a single page; See Appendix A. Remarkably, in addition to π , only the S-Box `wbob_sbox[256]` (See Section 2.2) together with minimal BLNK logic are required for full AEAD implementation. On many microcontrollers WHIRLBOB’s entire software footprint is in the 500B range. Slightly more is required for a shared secret handshake protocol and two-way secure BLINKER protocol [54].

This is a significant improvement over STRIBOBr1, which typically needs almost 2kB. STRIBOBr1 is also much slower and larger on low-end microcontrollers due to the “heavy” MDS matrix. The reference implementation is written for compactness and clarity; it is not optimal when it comes to speed or size. We refer to section 7.3 of [4] for techniques that greatly reduce the number of XORs required, resulting in increased processing speed. Additional tables will be required, however, and this will increase the overall implementation size.

3.1 Constant-Time SIMD Implementation

Due largely to Whirlpool’s S-Box structure and generous parallelism, it is well suited for high speed, constant-time implementation on Single Instruction Multiple Data (SIMD) architectures. Here we focus on ARM’s NEON as the reference architecture since the state layout fits the registers nicely, but also consider Intel’s SSSE3 as another explicit example. The goal is to improve performance, while at the same time avoiding memory-resident table lookups that cause execution time to depend on the data cache state and thus algorithm state (the crux of cache timing attacks).

Related work in this area includes simulated ISA extensions to a RISC architecture for parallel table lookups to speed up Whirlpool [28]. These extensions are then used to build essentially a hardware-assisted analogue of the traditional T tables software implementation – storing the state in rows and issuing a single instruction to perform 8 parallel lookups from the 8-bit S-Box input to the 64-bit linear layer output and XOR-summing the results, repeated for each row. AES [27] and Anubis [18] can also take advantage of SSSE3’s variable byte shuffle instruction for fast and secure implementations.

NEON has 32×64 - bit SIMD registers and SSSE3 16×128 - bit. We store the state column-wise (one column per NEON register, two columns per SSSE3 register), i.e. byte position j of register i contains the state byte in column i and row j . The `SubBytes` step is not sensitive to this ordering, but both `ShiftColumns` and




Fig. 2. The WHIRLBOB 8×8 - bit S-Box is constructed from three 4×4 - bit “miniboxes”. In this diagram the most significant bits are on the left: E operates on the higher nybble.

MixRows are. Since both of these architectures feature variable byte shuffle instructions (`vtbl.u8` for NEON and `pshufb` for SSSE3), implementing `SubBytes` is a direct translation of Figure 2 to these instructions. This amounts to 40 NEON shuffles and half as many SSSE3 shuffles. For `ShiftColumns`, NEON uses `vext` for byte-wise register rotation and SSSE3 `pshufb` with constant rotation distances since each register holds two columns. For `MixRows` we use the row formula from the Whirlpool specification [4, Sec. 7.3] where the multiplications by x are a simple left shift (native on NEON, integer addition on SSSE3) and conditional XOR (operand masked by signed right shift on NEON, comparison on SSSE3). The formula is fairly symmetric around even and odd byte positions – while NEON implements it as written with 24 multiplications, SSSE3 slightly rearranges a few registers to parallelize across the full 128-bit register width and use half as many multiplications.

3.2 Generic Constant-Time Bitsliced Implementation

The byte-oriented $8 \times 64 = 512$ - bit state can be rapidly split into eight 64-bit registers. The parallelism evident in Figure 2 helps to speed up bitsliced implementation. We see that for $2/3$ of the time, the S-Box has effectively two independent 4-bit execution paths. Interleaving these may greatly reduce wait states due to the superscalar architecture employed by most modern CPUs.

Appendix B of the current 2003 Whirlpool specification [4] gives listings with 14-16 instructions/gates for each of the miniboxes (if `ANDN` instruction is allowed). Those were used in our reference bitsliced implementations.




Fig. 3. WHIRLBOB was implemented on the FPGA logic fabric of Xilinx Zynq 7010. The implementation integrates with the AXI bus of ARM Cortex A9 on the SoC chip.

3.3 Implementation Summary

We currently have six implementations of WHIRLBOB. They mainly differ in the implementation technique used for the π cryptographic permutation.

- **C 8-bit:** This is the minimal reference implementation which is optimized for clarity and low-resource platforms, corresponding to Appendix A.
- **C 64-bit:** Standard speed-optimized implementation for most platforms, utilizing large lookup tables. Apart from Whirlpool-derived tables, equivalent to the implementation of Sa14E.
- **C Bitsliced:** Straight-line, fully bitsliced implementation without data-dependent branches or lookups. Resistant to timing attacks.
- **NEON Intrinsics:** Fast constant-time version that avoids table lookups by storing 4×4 - bit miniboxes in SIMD registers.
- **SSSE3 Intrinsics:** Similar but for 128-bit SIMD registers.
- **Verilog 12-cycle:** This is the hardware reference implementation. Source code is about 350 lines. Additional logic is required for AXI Bus integration.

Software Implementations. The first three implementations use only C99, and are hence easily portable. The 64-bit reference implementation is almost exactly as fast as OpenSSL’s Whirlpool on the same platform. See Table 2 for implementation metrics. We also have various embedded implementations.

Hardware Implementation. The hardware implementation has been proven on FPGA (Figure 3). The SÆHI proposal reports total post place-and-route utilization on Artix-7 of 4,946 logic units for a single round of WHIRLBOB, which compares favorably to 7,972 required for Keccak/Keyak [55]. Throughput is roughly 2 MB/s for each MHz.

4 Comparison with Other AEAD Schemes

At the time of writing (Q1/2015) the dominant AEAD scheme is the Galois / Counter Mode (GCM) for the AES block cipher [43,45], which is recommended for use with

Table 2. Comparing software implementations of WHIRLBOB’s π .

Metric and Target	Speed MB/sec	Footprint		Source C lines
		Code	Data	
<u>Single Core of 2.8GHz Core i7 860</u>				
8-bit C99 Reference	4.6311	326	256	97
64-bit C99 Reference	95.368	1942	16512	128
Bitsliced C99 Reference	30.856	4592	768	345
SSSE3 (Constant-Time)	123.15	1290	1152	256
<u>BeagleBone Black 1.0GHz Cortex-A8</u>				
8-bit C99 Reference	0.8288	352	256	97
64-bit C99 Reference	3.3435	6524	16512	128
Bitsliced C99 Reference	1.4353	15704	768	345
NEON (Constant-Time)	9.2084	1528	1072	320

TLS, SSH and IPSec protocols by NSA as part of “Suite B” [19,29,48,58]. GCM message authentication is based on polynomial evaluation in the finite field $GF(2^{128})$. The required multiplication can be exceedingly slow on lightweight platforms. An LFSR-style implementation of a 128×128 - bit multiplication will require thousands of cycles on 8-bit targets.

It is often more efficient to use the CCM [44,62] double-mode of operation on lightweight platforms, since implementing a full extra AES operation can be faster than the finite field multiplication operation. CCM and GCM are currently the only two FIPS - standardized authenticated modes. The performance characteristics of AES-CCM AEAD can be expected to be very similar to WHIRLBOB due to their structural similarities and relative data bandwidth:

- WHIRLBOB: 12 rounds with 64 S-Boxes for 256 bits of data.
- AES-192-CCM: 2×12 rounds with 16 S-Boxes for 128 bits of data.

There are additional (patented) AES modes which will be faster on 8-bit platforms – such as AES-OTR [40] and AES-OCB [34], and dozens of others. Virtually all AEAD block cipher modes offer lower levels of integrity protection (2^{64} level even for 128-bit tags) and are not directly usable in wider Sponge applications such as non-randomized hashing.

Currently only unoptimized reference implementations are available for most CAESAR candidates [20], making fair performance comparisons difficult. Furthermore, no other CAESAR candidate is targeted at 192-bit security level (apart from AES modes) and little attention has been paid to 8-bit or hardware implementations.

We note that leading full-featured Sponge candidates, directly SHA3 / Keccak - based Ketje Keyak [12] and [6] have significantly slower reference implementation than STRIBOBr1 and WHIRLBOB (Table 3). WHIRLBOB falls very significantly from candidates such as NORX [3] and MORUS [63], which have been designed specifically with 64-bit targets in mind. Our proposal can claim a more conservative security margin when compared to these candidates, however.

Table 3. Relative performance of some CAESAR candidates on a AMD64 reference system in SUPERCOP testing (smaller number indicates faster speed).

MORUS 1280 - 128 [63]	0.09
NORX 64-4-1 [3]	0.19
ASCON-128n [24]	0.89
WHIRLBOB Intel SSSE3 Constant-Time	1.00
WHIRLBOB and STRIBOBr1 64-bit Reference	1.26
Lake Keyak [12]	2.23
Ketje Sr. [6]	4.25
PRIMATES (HANUMAN, GIBBON, APE) [2]	50+

bench.cr.yp.to/web-impl/amd64-titan0-crypto_aead.html

5 Security Analysis and Design Notes

For analysis of the round function against classical Differential and Linear cryptanalysis we refer to Whirlpool literature [4]. Two additional rounds increase WHIRLBOB’s resistance against best known attacks [35,36].

Most of the security arguments and proofs offered for STRIBOBr1 and BLNK in [55] also apply to the new proposal. These are based on indistinguishability arguments for the π permutation and a simple theorem (Thm. 1, Sec. 3.3. in [55]) that loosely ties the compression function in Miyagushi-Preneel mode [41,51] with the indistinguishability of π . A random-indistinguishable π and appropriate padding rules are sufficient to construct Sponge-based hashes [7], Tree Hashes [11], MACs [10], Authenticated Encryption (AE) algorithms [9], and pseudorandom extractors (SHAKEs, PRFs, and PRNGs) [8,47].

5.1 Side-Channel and Implementation Attacks

Due to the minibox structure, we may load the 4×4 - bit tables in registers and access them via constant-time shuffles on Intel SSSE3 and ARM NEON SIMD targets as noted in Section 3.1. WHIRLBOB is also relatively well suited for bitsliced implementation due to its particular S-Box and MDS design as noted in Section 2.2.

Being unconditional straight-line code without data-dependent table lookups, bitsliced and byte shuffling implementations are effective countermeasures against cache timing attacks, which can be mounted against cryptographic primitives with large tables such as AES [1,5,50,61].

A non-constant-time implementation of the S-Box on Whirlpool, Streebog, or STRIBOBr1 on 64-bit platforms typically requires lookup tables of up to $8 \times 256 \times 8 = 16384B$. Even though this size easily fits into the Level 2 cache of any 64-bit system, one may see that timing attacks are possible as L2 caches are not always shared even between different execution cores within a single CPU unit. This is due to the process switching operation of most 64-bit operating systems.

5.2 Historical Modifications to Whirlpool

Whirlpool has received a significant amount of analysis in the almost 15 years since its original publication. Whirlpool was the only hash function in the final NESSIE portfolio in addition to SHA-2 hashes [42]. Whirlpool has also been standardized by ISO as part of ISO/IEC 10118-3:2004 [30].

The amended MDS matrix used by current ('03) Whirlpool is also used by WHIRLBOB as a countermeasure to the structural observations given in [59]. Our design is based on Whirlpool 3.0.

Whirlpool was found to be vulnerable to a Rebound Distinguisher [35,36,39]. That 2^{188} attack applies to the 10-round variant; our 12-round version should offer a comfortable security margin, especially as our security target is 2^{192} . The way the round constants are derived from the S-Box allows this change to be made in a straightforward manner.

5.3 Notes on the origins of Streebog, Kuznyechik, and STRIBOBr1

The GOST R 34.11-2012 “Streebog” standard text [26] does not describe the linear step as a 8×8 matrix-vector multiplication with $GF(2^8)$ elements like the STRIBOBr1 spec [55], but as a 64×64 binary matrix multiplication. One can see that $8 \times 8 \times 8 = 512$ bits are required to describe the former, but $64 \times 64 = 4096$ bits are required for the latter. The more effective description was discovered by Kazymorov and Kazymorova in [33] by exhaustively testing all 30 irreducible polynomial bases, revealing an AES-like MDS structure. The origin of the particular numerical values of that MDS matrix and round constants is still a mystery. They do not appear to offer similar avenues for size or performance optimization like those in Whirlpool 3.0 and STRIBOBr2 do.

The 8-bit S-Box used by STRIBOBr1 was directly lifted from Streebog so that hardware and software components developed for Streebog could be shared or recycled when implementing STRIBOBr1. The same S-Box is also used by the recently proposed Russian Encryption Standard “Kuznyechik” [25,60].

Not much about the particular design criteria of the Streebog S-Box has been published. That S-Box was apparently selected at least 5 years ago as Streebog already appeared in RusCrypto '10 proceedings [38]. We can easily observe that it offers reasonable resistance against basic methods of cryptanalysis. Its differential bound [13] is $P = \frac{8}{256}$ and best linear approximation [37] holds with $P = \frac{28}{128}$. There does not seem to be any exploitable algebraic weaknesses. These are the exact same bounds as can be found for the Whirlpool and STRIBOBr2 S-Box, but fall short from the bounds of the AES S-Box. The Streebog S-Box does not seem to allow decomposition to 4×4 -bit miniboxes like Whirlpool (Figure 2). However, very recent ongoing work has revealed it to also have an optimized representation [16].

The Rijndael AES S-Box is constructed from finite field inversion x^{-1} operation in $GF(2^8)$ (inspired by the Nyberg construction [49]) and an affine bit transform that serves as a countermeasure against, among other things, Interpolation Attacks [31] on the AES predecessor SHARK [52]. We refer to [23] for more information about the AES design process.

We had brief informal discussions with some members of the Streebog and design team at the CTCrypt '14 workshop (05-06 June 2014, Moscow RU). Their recollection was that the aim was to choose a “randomized” S-Box that meets the basic differential, linear, and algebraic requirements. Randomization using various building blocks was simply iterated until a “good enough” permutation was found. This was seen as an effective countermeasure against yet-unknown attacks. At the time of Streebog S-Box selection (before 2010’s) the emergence of allegedly effective AES Algebraic Attacks such as [22] was a major concern for much of the symmetric cryptographic community. Hence it was felt appropriate to avoid too much algebraic structure in either the S-Box or MDS matrix while also ensuring necessary resistance against known attacks such as DC and LC. Algebraic attack attempts of this type against AES have since largely fizzled out, so we feel confident that the Whirlpool S-Box should be sufficient for our claimed security level, especially as it offers significantly better speeds in bitsliced implementations.

One is left with the impression that Streebog is a “whitened” or randomized copy of the original Whirlpool design. Despite its partially unknown origins and relative shortcomings on some implementation targets, we consider STRIBOBr1 to be at least as secure as STRIBOBr2 if appropriately implemented. Indeed some of the more successful attacks on AES and Whirlpool have been based on their deep structural self-similarities and simplistic key schedules [14,15,17].

6 Conclusions

We have introduced WHIRLBOB, an algorithm for Authenticated encryption with Associated Data. WHIRLBOB is a variant of the STRIBOBr1 first round CAESAR candidate but borrows its main components from the Whirlpool 3.0 hash. WHIRLBOB, also known as STRIBOBr2, is a CAESAR [21] second round candidate [57].

WHIRLBOB has extremely small implementation footprint on resource-limited software platforms – typically under half a kilobyte. Its particular S-Box and MDS design allows WHIRLBOB to have efficient constant-time bitsliced and SIMD byte shuffling implementations. This is an effective countermeasure against cache timing attacks, which are a concern against AES. The $b = 8 \times 64$ - bit state size is particularly suitable for bitslicing of a byte-oriented algorithm on 64-bit platforms and byte slicing for SIMD platforms.

WHIRLBOB has superb implementation characteristics on FPGA (ASIC), SIMD and lightweight embedded platforms. We recommend WHIRLBOB especially for those platforms. Furthermore WHIRLBOB offers provable security assurance through its security relationship with the well-analyzed Whirlpool hash.

We have also discussed the design choices for the STRIBOB S-Box and other components used in the Streebog hash and Kuznyechik cipher, which are becoming standards for the Russian security market.

7 Acknowledgements

We thank Oleksandr Kazymyrov, Vasily Shishkin, Bart Preneel, and Paulo Barreto for their helpful comments. Dr. Saarinen carried out much his research during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme at Norwegian University of Technology, Trondheim.

References

1. Onur Aciicmez, Werner Schindler, and Çetin Kaya Koç. Cache based remote timing attack on the AES. In Masayuki Abe, editor, *CT-RSA 2007*, volume 4377 of *LNCS*, pages 271–286. Springer, 2007. doi:10.1007/11967668_18.
2. Elena Andreeva, Begül Bilgin, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. PRIMATES v1 - Submission to the CAESAR Competition. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/primatesv1.pdf>.
3. J.-P. Aumasson, Philipp Jovanovic, and S. Neves. CAESAR submission: NORX v1. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/norxv1.pdf>.
4. Paulo S. L. M. Barreto and Vincent Rijmen. The Whirlpool hashing function. NESSIE Algorithm Specification, 2000, Revised May 2003. URL: <http://www.larc.usp.br/~pbarreto/WhirlpoolPage.html>.
5. D. J. Bernstein. Cache-timing attacks on AES. Technical report, University of Chigaco, 2005. URL: <http://cr.yp.to/antiforgery/cachetiming-20050414.pdf>.
6. G. Bertoni, Joan Daemen, M. Peeters, G. Van Assche, and R. Van Keer. CAESAR submission: Ketje v1. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/ketjev1.pdf>.
7. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge functions. In *Ecrypt Hash Workshop 2007*, May 2007. URL: <http://events.iaik.tugraz.at/HashWorkshop07/program.html>.
8. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge-based pseudo-random number generators. In Stefan Mangard and François-Xavier Standaert, editors, *CHES 2010*, volume 6225 of *LNCS*, pages 33–47. Springer, 2010. doi:10.1007/978-3-642-15031-9_3.
9. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Duplexing the sponge: Single-pass authenticated encryption and other applications. In A. Miri and S. Vaudenay, editors, *SAC 2011*, volume 7118 of *LNCS*, pages 320–337. Springer, 2011. doi:10.1007/978-3-642-28496-0_19.
10. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. The Keccak reference, version 3.0. NIST SHA3 Submission Document, January 2011. URL: <http://keccak.noekeon.org/Keccak-reference-3.0.pdf>.
11. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sakura: a flexible coding for tree hashing. IACR ePrint 2013/231, April 2013. URL: <https://eprint.iacr.org/2013/231>.
12. Guido Bertoni, Joan Daemen, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. CAESAR submission: Keyak v1. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/keyakv1.pdf>.
13. Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer, 1993. doi:10.1007/978-1-4613-9314-6.

14. Alex Biryukov and Dmitry Khovratovich. Related-key cryptanalysis of the full AES-192 and AES-256. In Mitsuru Matsui, editor, *ASIACRYPT '09*, volume 5912 of *LNCS*, pages 1–18. Springer, 2009. doi:10.1007/978-3-642-10366-7_1.
15. Alex Biryukov, Dmitry Khovratovich, and Ivica Nikolić. Distinguisher and related-key attack on the full AES-256. In Sjai Halevi, editor, *CRYPTO '09*, volume 5677 of *LNCS*, pages 231–249. Springer, 2009. doi:10.1007/978-3-642-03356-8_14.
16. Alex Biryukov, Lèò Perrin, and Aleksei Udovenko. The secret structure of the s-box of streebog, kuznechik and stribob. IACR ePrint 2015/812, August 2015. URL: <https://eprint.iacr.org/2015/812>.
17. Andrey Bogdanov, Dmitry Khovratovich, and Christian Rechberger. Biclique cryptanalysis of the full AES. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT '11*, volume 7073 of *LNCS*, pages 344–371. Springer, 2011. doi:10.1007/978-3-642-25385-0_19.
18. Billy B. Brumley. Secure and fast implementations of two involution ciphers. In T. Aura, K. Järvinen, and K. Nyberg, editors, *NordSec '10*, volume 7127 of *LNCS*, pages 269–282. Springer, 2012. doi:10.1007/978-3-642-27937-9_19.
19. K. Burgin and M. Peck. Suite B Profile for Internet Protocol Security (IPsec). IETF RFC 6380, October 2011.
20. CAESAR. CAESAR: Competition for authenticated encryption: Security, applicability, and robustness, January 2014. URL: <http://competitions.cr.yyp.to/caesar.html>.
21. CAESAR. CAESAR first and second round submissions, July 2015. URL: <http://competitions.cr.yyp.to/caesar-submissions.html>.
22. N. Courtois. How fast can be algebraic attacks on block ciphers? IACR ePrint 2006/168, May 2006. URL: <https://eprint.iacr.org/2006/168>.
23. Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - the Advanced Encryption Standard*. Springer, 2002. doi:10.1007/978-3-662-04722-4.
24. C. Dobraunig, M. Eichlseder, Florian Mendel, and M. Schläffer. Ascon v1 - Submission to the CAESAR Competition. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yyp.to/round1/asconv1.pdf>.
25. Denis M. Dygin, Ivan V. Lavrikov, Grigory B. Marshalko, Vladimir I. Rudskoy, Dmitry I. Trifonov, and Vasily A. Shishkin. On a new Russian Encryption Standard. *Mathematical Aspects of Cryptography*, 6(2):29–34, 2015. (Abstract In Russian). URL: http://www.mathnet.ru/php/archive.phtml?wshow=paper&jrnid=mvk&paperid=142&option_lang=eng.
26. GOST. Information technology. cryptographic protection of information, hash function. GOST R 34.11-2012, 2012. (In Russian). URL: <http://protect.gost.ru/v.aspx?control=7&id=180209>.
27. M. Hamburg. Accelerating AES with vector permute instructions. In C. Clavier and K. Gaj, editors, *CHES '09*, volume 5747 of *LNCS*, pages 18–32. Springer, 2009. doi:10.1007/978-3-642-04138-9_2.
28. Y. Hilewitz, Y. L. Yin, and R. B. Lee. Accelerating the Whirlpool hash function using parallel table lookup and fast cyclical permutation. In K. Nyberg, editor, *FSE '08*, volume 5086 of *LNCS*, pages 173–188. Springer, 2008. doi:10.1007/978-3-540-71039-4_11.
29. K. Igoe. Suite B Cryptographic Suites for Secure Shell (SSH). IETF RFC 6239, May 2011. URL: <https://tools.ietf.org/html/rfc6239>.
30. ISO/IEC. Information technology - security techniques - hash-functions - part 3: Dedicated hash-functions. ISO/IEC 10118-3:2004, 2004. URL: <https://www.iso.org/obp/ui/#iso:std:iso-iec:10118:-3:ed-3:v1:en>.
31. T. Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In Eli Biham, editor, *FSE '97*, volume 1267 of *LNCS*, pages 99–112. Springer, 1996.

32. Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond $2^{c/2}$ security in sponge-based authenticated encryption modes. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014*, volume 8873 of *LNCS*, pages 85–104. Springer, 2014. doi:10.1007/978-3-662-45611-8_5.
33. Oleksandr Kazymyrov and Valentyna Kazymyrova. Algebraic aspects of the Russian hash standard GOST R 34.11-2012. In *CTCrypt '13, June 23-24, 2013, Ekaterinburg, Russia*, 2013. IACR ePrint 2013/556. URL: <https://eprint.iacr.org/2013/556>.
34. Ted Krovetz and Phillip Rogaway. OCB (v1). CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/ocbv1.pdf>.
35. Mario Lamberger, Florian Mendel, Christian Rechberger, Vincent Rijmen, and Martin Schl affer. Rebound distinguishers: Results on the full whirlpool compression function. In Mitsuru Matsui, editor, *ASIACRYPT '09*, volume 5912 of *LNCS*, pages 126–143. Springer, 2009. doi:10.1007/978-3-642-10366-7_8.
36. Mario Lamberger, Florian Mendel, Martin Schl affer, Christian Rechberger, and Vincent Rijmen. The rebound attack and subspace distinguishers: Application to Whirlpool. *J. Cryptology*, 28:257–296, 2015. doi:10.1007/s00145-013-9166-5.
37. Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In T. Helleseeth, editor, *EUROCRYPT '93*, volume 765 of *LNCS*, pages 386–397. Springer, 1994. doi:10.1007/3-540-48285-7_33.
38. D. V. Matyuhin, V. I. Rudskoy, and V. A. Shishkin. Promising hashing algorithm. RusCrypto '10 Workshop, 02 April 2010, 2010. (In Russian).
39. Florian Mendel, Christian Rechberger, Martin Schl affer, and S oren S. Thomsen. The rebound attack: Cryptanalysis of reduced Whirlpool and Gr ostl. In Orr Dunkelman, editor, *FSE 2009*, volume 5665 of *LNCS*, pages 260–276. Springer, 2009. doi:10.1007/978-3-642-03317-9_16.
40. Kazuhiko Minematsu. AES-OTR v1. CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.yp.to/round1/aesotrv1.pdf>.
41. S. Miyaguchi, K. Ohta, and M. Iwata. 128-bit hash function (n -hash). *NTT Review*, (2):128–132, 1990.
42. NESSIE. *Final report of European project number IST-1999-12324, named New European Schemes for Signatures, Integrity, and Encryption*. NESSIE, April 2004. URL: <https://www.cosic.esat.kuleuven.be/nessie/Bookv015.pdf>.
43. NIST. Advanced Encryption Standard (AES). Federal Information Processing Standards Publication FIPS 197, November 2001. URL: <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.
44. NIST. Counter with Cipher Block Chaining - Message Authentication Code (CCM). NIST Special Publication 800-38C, May 2004.
45. NIST. Recommendation for block cipher modes of operation: Galois/counter mode (GCM) and GMAC. NIST Special Publication 800-38D, 2007. URL: <http://csrc.nist.gov/publications/nistpubs/800-38D/SP-800-38D.pdf>.
46. NIST. The Keyed-Hash Message Authentication Code (HMAC). Federal Information Processing Standards Publication FIPS 198-1, July 2008.
47. 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.
48. NSA. Suite B Cryptography, 2005. URL: http://www.nsa.gov/ia/programs/suiteb_cryptography.
49. Kaisa Nyberg. Differentially uniform mappings for cryptography. In Tor Helleseeth, editor, *EUROCRYPT '93*, volume 765 of *LNCS*, pages 55–64. Springer, 1994. doi:10.1007/3-540-48285-7_6.

50. Dag Arne Osvik, Adi Shamir, and Eran Tromer. Cache attacks and countermeasures: The case of AES. In David Pointcheval, editor, *CT-RSA 2006*, volume 3860 of *LNCS*, pages 1–20. Springer, 2006. doi:10.1007/11605805_1.
51. B. Preneel. *Analysis and Design of Cryptographic Hash Functions*. PhD thesis, K. U. Leuven (Belgium), 1993. URL: http://homes.esat.kuleuven.be/~preneel/phd_preneel_feb1993.pdf.
52. Vincent Rijmen, Joan Daemen, Bart Preneel, Antoon Bosselaers, and Erik De Win. The cipher SHARK. In Dieter Gollmann, editor, *FSE '96*, volume 1039 of *LNCS*, pages 99–111. Springer, 1996. doi:10.1007/3-540-60865-6_47.
53. Markku-Juhani O. Saarinen. Simple AEAD hardware interface (SÆHI) in a SoC: Implementing an on-chip Keyak/WhirlBob coprocessor. In *TrustED '14 Proceedings of the 4th International Workshop on Trustworthy Embedded Device*, pages 51–56. ACM, 2014. doi:10.1145/2666141.2666144.
54. Markku-Juhani O. Saarinen. StriBob: Authenticated encryption from GOST R 34.11-2012 LPS permutation. In *CTCrypt '14, 05-06 June 2014, Moscow, Russia. Preproceedings.*, pages 170–182, June 2014. URL: <https://eprint.iacr.org/2014/271>.
55. Markku-Juhani O. Saarinen. The STRIBOBr1 authenticated encryption algorithm. CAESAR, 1st Round Candidate, March 2014. URL: <http://www.stribob.com>.
56. Markku-Juhani O. Saarinen. StriBob: Authenticated encryption from GOST R 34.11-2012 LPS permutation. *Mathematical Aspects of Cryptography*, 6(2):67–78, 2015. (Abstract In Russian). URL: http://www.mathnet.ru/php/archive.phtml?wshow=paper&jrnid=mvk&paperid=146&option_lang=eng.
57. Markku-Juhani O. Saarinen and Billy B. Brumley. STRIBOBr2: “WHIRLBOB”, second round caesar algorithm tweak specification. CAESAR 2nd Round Candidate, August 2015. URL: <http://www.stribob.com>.
58. M. Salter and R. Housley. Suite B Profile for Transport Layer Security (TLS). IETF RFC 6460, January 2012. URL: <https://tools.ietf.org/html/rfc6460>.
59. T. Shirai and K. Shibutani. On the diffusion matrix employed in the Whirlpool hashing function. NESSIE Public Report, 2003. URL: <http://www.cosic.esat.kuleuven.be/nessie/reports/phase2/whirlpool-20030311.pdf>.
60. Vasily Shishkin, Denis Dygin, Ivan Lavrikov, Grigory Marshalko, Vladimir Rudskoy, and Dmitry Trifonov. Low-weight and hi-end: Draft Russian Encryption Standard. In *CTCrypt '14, 05-06 June 2014, Moscow, Russia. Preproceedings.*, pages 183–188, June 2014.
61. Michael Weiß, Benedikt Heinz, and Frederic Stumpf. A cache timing attack on AES in virtualization environments. In Angelos D. Keromytis, editor, *FC 2012*, volume 7397 of *LNCS*, pages 314–328. Springer, 2013. doi:10.1007/978-3-642-32946-3_23.
62. D. Whiting, R. Housley, and Niels Ferguson. Counter with CBC-MAC (CCM). IETF RFC 3610, September 2003. URL: <https://tools.ietf.org/html/rfc3610>.
63. Hongjun Wu and Tao Huang. The Authenticated Cipher MORUS (v1). CAESAR First Round Submission, March 2014. URL: <http://competitions.cr.ypt.to/round1/morusv1.pdf>.

A WHIRLBOB 1.0 π “8-bit” Reference Implementation

This ANSI C function implements the WHIRLBOB 512×512 -bit π permutation.

```
void wbob_pi(uint8_t st[64]) // WHIRLBOB Pi
{
    int r, i, j;
    uint8_t t[64], x, *pt;

    for (r = 0; r < 12; r++) { // 12 rounds
        for (i = 0; i < 64; i++) {
            t[(i & 7) + ((i + (i << 3)) & 070)] = // P
                wbob_sbox[st[i]]; // S
        }

        // The round constants C come from the S-box
        pt = (uint8_t *) &wbob_sbox[8 * r];
        for (i = 0; i < 8; i++)
            st[i] = pt[i]; // C in first 8
        for (i = 8; i < 64; i++)
            st[i] = 0; // zero the rest

        // Apply the circular, low weight MDS matrix
        for (i = 0; i < 64; i += 8) {
            pt = &st[i]; // start of row
            for (j = 0; j < 8; j++) {
                x = t[i + j]; // Circular MDS
                pt[j & 7] ^= x; // 01
                pt[(j + 1) & 7] ^= x; // 01
                pt[(j + 3) & 7] ^= x; // 01
                pt[(j + 5) & 7] ^= x;
                pt[(j + 7) & 7] ^= x;

                // x <- 02
                x = (x << 1) ^ (x & 0x80 ? 0x1D : 0x00);
                pt[(j + 6) & 7] ^= x; // 02

                // x <- 04
                x = (x << 1) ^ (x & 0x80 ? 0x1D : 0x00);
                pt[(j + 2) & 7] ^= x; // 04
                pt[(j + 5) & 7] ^= x; // 01 + 04 = 05

                // x <- 08
                x = (x << 1) ^ (x & 0x80 ? 0x1D : 0x00);
                pt[(j + 4) & 7] ^= x; // 08
                pt[(j + 7) & 7] ^= x; // 01 + 08 = 09
            }
        }
    }
}
```


B Test Vectors

B.1 The 12-round π transform

These are derived from ISO Test vectors of Whirlpool 3.0. We give the input x_0 and results after 1, 10 (as in Whirlpool), and full 12 rounds of processing.

$$\begin{array}{l}
 x_0 = \begin{pmatrix} 77\ 38\ E1\ B5\ 41\ A0\ 36\ EA \\ 45\ 8D\ 50\ F8\ 0F\ A0\ 1C\ 44 \\ 72\ 88\ CE\ 97\ D1\ A0\ DC\ F0 \\ 16\ 95\ FF\ D6\ E7\ 1D\ 09\ 25 \\ 33\ BE\ 30\ 9F\ 01\ 2A\ 59\ 09 \\ 72\ 91\ 14\ 59\ 5F\ 08\ 6E\ 76 \\ 07\ 18\ AF\ E3\ 65\ BC\ 09\ DE \\ B6\ AF\ A1\ 80\ BC\ EC\ 2A\ 98 \end{pmatrix} \\
 \\
 x_{10} = \begin{pmatrix} B4\ 74\ E1\ 56\ 96\ 31\ B9\ 6C \\ 21\ A1\ B6\ 33\ CC\ 89\ 68\ 1A \\ B1\ 97\ 25\ 86\ 7B\ 2B\ 3F\ 09 \\ 4C\ 73\ C7\ 62\ 93\ A8\ 15\ CF \\ 55\ 15\ C0\ C0\ 9A\ 05\ 05\ 16 \\ 23\ 44\ 8D\ 8D\ D3\ 5F\ B3\ 6E \\ 7E\ 6C\ 2D\ 37\ 12\ D0\ F3\ 3E \\ CE\ B8\ 04\ F2\ 8D\ 9F\ C9\ 99 \end{pmatrix} \\
 \\
 x_1 = \begin{pmatrix} 1A\ 78\ 4D\ 7D\ BD\ 4C\ 17\ E6 \\ 27\ 31\ 10\ AA\ 63\ C5\ 9E\ 25 \\ 7A\ 2E\ B7\ 48\ C4\ 5D\ E0\ 23 \\ 6D\ 0D\ 61\ 9F\ 6C\ 1D\ 80\ AE \\ 01\ A2\ D5\ 6E\ DB\ 41\ D9\ A0 \\ E9\ 06\ 4C\ D1\ 27\ 95\ FA\ 86 \\ 77\ 62\ 31\ BC\ B4\ 4E\ C6\ 01 \\ 6F\ CD\ BC\ 98\ 10\ 78\ 6F\ EC \end{pmatrix} \\
 \\
 x_{12} = \begin{pmatrix} 3F\ 72\ C2\ 60\ EE\ 28\ EF\ EA \\ 42\ 8E\ B5\ 3A\ FB\ 8A\ 33\ A2 \\ 03\ E4\ 72\ 31\ 90\ A5\ 1A\ D3 \\ 3E\ 68\ E6\ 46\ FC\ 94\ 3C\ C7 \\ 80\ 42\ 9E\ 2E\ CB\ 32\ 75\ 93 \\ 30\ AA\ E2\ 21\ 21\ C8\ 99\ ED \\ 86\ 1E\ 06\ 9E\ 91\ 1F\ 89\ 6C \\ D2\ 99\ EC\ 7E\ E9\ 0B\ 01\ 10 \end{pmatrix}
 \end{array}$$

The last entry corresponds to the final output $\pi(x_0) = x_{12}$.

B.2 Authenticated Encryption

Inputs are plain ASCII.

K = "192-bit Secret Key value" (24 B)
 N = "Nonces Used Once" (16 B)
 A = "AAD Test Vector Exact Block 32 B" (32 B)
 P = "2 Block Test Vector for STRIBOB192r2d2" (38 B)

Authenticated ciphertext has 38 message bytes + 16 for MAC = 54 (0x36) bytes:

	x0	x1	x2	x3	x4	x5	x6	x7	x8	x9	xA	xB	xC	xD	xE	xF
0x	59	9C	5F	69	7F	16	30	07	B4	D5	52	30	24	0C	2B	7B
1x	0A	93	4E	4C	63	19	4F	AC	EA	2D	D5	4E	BD	05	61	2C
2x	19	92	47	FC	A1	97	AE	AE	71	0F	0D	ED	3E	56	5B	D0
3x	26	FE	20	F6	4A	4F										

For BLNK padding technical implementation details, see [57].