Secure and Fast Implementations of Two Involution Ciphers

Billy Bob Brumley*

Aalto University School of Science and Technology, Finland
billy.brumley@tkk.fi

Abstract. Anubis and Khazad are closely related involution block ciphers. Building on two recent AES software results, this work presents a number of constant-time software implementations of Anubis and Khazad for processors with a byte-vector shuffle instruction, such as those that support SSSE3. For Anubis, the first is serial in the sense that it employs only one cipher instance and is compatible with all standard block cipher modes. Efficiency is largely due to the S-box construction that is simple to realize using a byte shuffler. The equivalent for Khazad runs two parallel instances in counter mode. The second for each cipher is a parallel bit-slice implementation in counter mode.

Key words: Anubis, Khazad, involution ciphers, block ciphers, software implementation, timing attacks.

1 Introduction

Anubis and Khazad are two block ciphers by Barreto and Rijmen submitted during the NESSIE project (see [Pre02] for a summary). Anubis [BR01a] works on 128-bit blocks and is quite similar in many respects to AES. Khazad [BR01b] is a "legacy-level" cipher working on 64-bit blocks and is closely related to Anubis. These are both involution ciphers: decryption differs from encryption only in the key schedule.

The motivation for this work comes largely from cache-timing attacks, where an attacker attempts to recover parts of the cryptosystem state by observing the variance in timing measurements due to processor data caching effects. These attacks can be time-driven and carried out remotely by measuring the latency of a high level operation, or trace-driven and locally by exploiting the cache structure to determine the sequence of lookups the cryptosystem performs. The vulnerability exists when part of the state is used as an index into a memory-resident table.

A high-speed table-based implementation of AES unrolls lower level operations such as SubBytes, ShiftRows, and MixColumns into four tables of size 256 containing 32-bit values. Lookups into these tables, indexed by state values, are combined with XOR to carry out AES rounds in a more software-friendly manner, relaxing the need to manipulate a large number of single byte values and bits within those bytes. Similar versions of both Anubis and Khazad exist, in fact provided as the C reference implementations and discussed in both specifications [BR01a,BR01b, Sect. 7.1].

Cache-timing attacks are a serious threat and can easily lead to leakage of key material. Although there are numerous published attacks on such implementations, a practical noteworthy one is Bernstein's AES time-driven attack [Ber04]. Anubis and Khazad are presumably susceptible to this and other timing attacks. In light of these attacks, a reasonable security requirement for any cipher is that it can be implemented to use a constant amount of time. In this context, Bernstein defines constant as "independent of the AES key and input" [Ber04, Sect. 8]. The concept of security within this paper is with respect to timing attacks.

To this end, this work shows that constant-time and efficient implementations of both Anubis and Khazad are possible. Four such implementations appear herein, summarized as follows.

* Supported in part by the European Commission's Seventh Framework Programme (FP7) under contract number ICT-2007-216499 (CACE).

– The first Anubis implementation runs only one instance of the cipher, compatible with all standard block cipher modes. This is efficient due to a byte-vector shuffle instruction, allowing elegant realization of the nonlinear layer in constant-time. The Khazad implementation is otherwise analogous but with a smaller state runs two parallel cipher instances, here in counter mode under the same key.
– The second Anubis implementation bit-slices the state and runs eight parallel instances, here in counter mode. Not surprisingly, this is faster but requires a parallel block cipher mode. Analogously, the Khazad implementation runs 16 parallel instances.

This work builds upon two recent results on AES software implementations that remarkably manage to achieve constant-time and exceptional performance at one stroke.

– A common hardware technique to compute the AES S-box uses an isomorphism $\mathbb{F}_{2^8} \to \mathbb{F}_{2^4}^2$ and subsequently reduces the problem of inversion in the latter field to that of one in in the ground field. Armed with a byte-vector shuffle instruction and using a novel field element representation, Hamburg presents techniques for fast and constant-time software implementation of AES [Ham09]. Running only a single instance of the cipher, the implementation is compatible with all standard block cipher modes.
– Käsper and Schwabe present AES bit-slice techniques, aligning individual bits of state bytes in distinct registers [KS09]. The implementation runs eight parallel streams in counter mode under the same key. Not only does this provide a constant-time implementation, but also is currently the fastest published AES counter mode implementation in software. Table-based AES implementations on common platforms are inherently limited to 10 cycles per byte. The authors show that bit-slicing circumvents this limit.

2 Vector Operations

In 64-bit mode, processors with Streaming SIMD Extensions 3 (SSE3) can operate on 16 128-bit SIMD registers `xmm0` through `xmm15`. SSE3 and predecessors contain a wealth of instructions for parallel computation amongst these registers. Cryptosystem implementations usually restrict to a smaller subset of these instructions dealing with integer values. Supplemental SSE3 (SSSE3) introduces a handful of new instructions, the most interesting for this work being a byte shuffler `pshufb`. Note that recent AMD processors implement SSE3 but not SSSE3, although a related instruction is slated for the eXtended Operations (XOP) extension.

2.1 Byte Shuffling

Since the implementations in this work make heavy use of `pshufb`, a brief description of the instruction is in order. The name already implies the ability to shuffle bytes around in a vector, but perhaps hides an important aspect of the instruction. Aranha, López, and Hankerson note its versatility [ALH10, Sect. 2.1]:

> "A powerful use of this instruction is to perform 16 simultaneous lookups in a 16-byte lookup table."

Formally, given 16-signed-byte vector operands a and b, components of the 16-byte vector output r of `pshufb` satisfy

$$r_i = \begin{cases} a_{b_i \bmod 16} & \text{if } b_i \geq 0, \\ 0 & \text{otherwise,} \end{cases}$$

so b holds the indices into the table and a the values. Indeed, this allows to implement any $\mathbb{F}_2^4 \to \mathbb{F}_2^8$ function in parallel: this is a constant-time hardware lookup table, shuffling the

values in a based on the indices in b. To summarize, the use of `pshufb` is either that of shuffling bytes around in a fixed manner or implementing lookups into a fixed table, and the distinction is in the operand order.

2.2 Linear Maps

Given the above, one can implement a linear map $\phi : \mathbb{F}_2^8 \to \mathbb{F}_2^8$ on 16 bytes in parallel. Denote $\alpha \in \mathbb{F}_{2^8}$ by $\alpha = \alpha_H x^4 + \alpha_L$ where α_i are the 4-bit nibbles. Linearity ensures $\phi(\alpha) = \phi(\alpha_H x^4) + \phi(\alpha_L)$ and each input on the right is effectively only four bits. Denote 16-byte vectors $t_{\phi H}$ and $t_{\phi L}$ that map the corresponding input to the output; these are the a from the previous section. The following steps realize ϕ in parallel:

1. Mask the lower nibble (α_L) of each byte in the input vector. (`pand`)
2. Bit-shift the input four positions towards LSB and mask again (α_H). (`psrlq`, `pand`)
3. Shuffle $t_{\phi L}$ and $t_{\phi H}$ with their respective indices from the above steps. (`pshufb` $\times$ 2)
4. Bitwise XOR the two outputs together. (`pxor`)

The second mask is a minor inconvenience due to the lack of an instruction to shift bits of individual bytes in a 16-byte vector (there are no `psllb` and `psrlb` instructions). The following implementations uses this strategy often. Note that when applying multiple maps to the same input, the first two steps are needed only once.

3 The Anubis Cipher and Implementation

Although Anubis supports variable length keys, this work only explicitly considers 16-byte keys; generalizations are straightforward. Analogous to AES-128, Anubis consists of a 16-byte state. The state is either viewed as a vector in $\mathbb{F}_{2^8}^{16}$ or a 4×4 matrix with entries in $\mathbb{F}_{2^8}$ depending on the context. The specification denotes this by a map μ, but this work omits this formalization.

This section describes the components of the cipher following the style of the specification, and at the same time presents the serial SSSE3 implementation of said components. See the specification for a formal treatment [BR01a].

The nonlinear layer γ This layer is otherwise analogous to the AES SubBytes step, but with a different S-box. It applies an S-box $S : \mathbb{F}_{2^8} \to \mathbb{F}_{2^8}$ to each byte of the input. To facilitate efficient hardware implementation, the designers chose to build S using a three layer substitution-permutation network (SPN), where each layer includes two S-boxes $P, Q : \mathbb{F}_{2^4} \to \mathbb{F}_{2^4}$ termed "mini-boxes". Fig. 1 depicts this structure.

This composition can be implemented elegantly using `pshufb`. Since P and Q are four bits to four bits but the instruction allows a parallel four bit to eight bit lookup, the bit permutations following P and Q can be unrolled for each layer to provide shifted and spread versions of their output. For example, $Q(\texttt{0x1}) = \texttt{0xE} = \texttt{1110}$ but following the first layer the upper two bits get shifted two positions towards the MSB: here the lookup provides $Q_0(\texttt{0x1}) = \texttt{0x32} = \texttt{110010}$. This unrolling yields the following six lookup tables for the corresponding layers:

$$
\begin{aligned}
t_{Q0} &= \texttt{0x20012313311000333003022212113221} \\
t_{P0} &= \texttt{0x0408804C488488C4C08C404400C8CC0C} \\
t_{P1} &= \texttt{0x01022013122122313023101100323303} \\
t_{Q1} &= \texttt{0x80048C4CC44000CCC00C08884844C884} \\
t_{Q2} &= \texttt{0x08010B070D04000F0C03020A06050E09} \\
t_{P2} &= \texttt{0x102080706090A0D0C0B0405000E0F030}.
\end{aligned}
$$

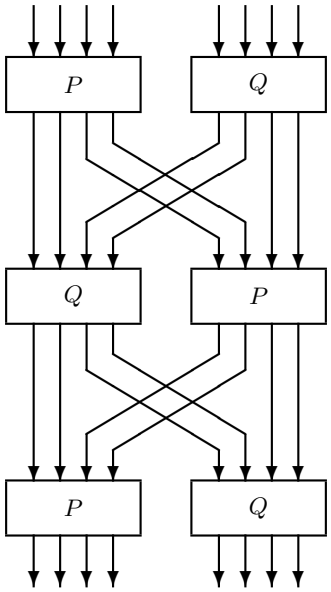

Fig. 1. S-box S as a three layer SPN with mini-boxes P and Q.

As the last layer does not permute the bits, note t_{P2} and t_{Q2} are simply the nibble-shifted and original contents as bytes, respectively, of P and Q.

With these tables in hand, the following steps implement layer i of S:

1. Mask the lower nibble of each byte in the input vector. (`pand`)
2. Bit-shift the input four positions towards LSB and mask again. (`psrlq`, `pand`)
3. Shuffle t_{Pi} and t_{Qi} with their respective indices from the above steps. (`pshufb` $\times$ 2)
4. Bitwise OR the two outputs together. (`por`)

Iterating this concept for all layers shows that S can be realized in parallel on all 16 input bytes using six `pand`, three `psrlq`, six `pshufb`, and three `por`.

The transposition τ Viewing the input as a 4×4 matrix, this mapping outputs the transpose. To illustrate:

$$\begin{bmatrix} \texttt{0 1 2 3} \\ \texttt{4 5 6 7} \\ \texttt{8 9 A B} \\ \texttt{C D E F} \end{bmatrix} \mapsto \begin{bmatrix} \texttt{0 4 8 C} \\ \texttt{1 5 9 D} \\ \texttt{2 6 A E} \\ \texttt{3 7 B F} \end{bmatrix}.$$

This requires only one `pshufb` instruction with indices defined as

$$t_\tau = \texttt{0x0F0B07030E0A06020D0905010C080400}.$$

The linear diffusion layer θ This layer shares some similarities with the AES MixColumns step. It multiplies the input in matrix form by the symmetric matrix

$$H = \begin{bmatrix} \texttt{01 02 04 06} \\ \texttt{02 01 06 04} \\ \texttt{04 06 01 02} \\ \texttt{06 04 02 01} \end{bmatrix} = \begin{bmatrix} 1 & x & x^2 & x^2+x \\ x & 1 & x^2+x & x^2 \\ x^2 & x^2+x & 1 & x \\ x^2+x & x^2 & x & 1 \end{bmatrix}$$

and $\theta : a \mapsto a \cdot H$ with all operations done in $\mathbb{F}_{2^8} = \mathbb{F}_2[x]/(x^8 + x^4 + x^3 + x^2 + 1)$. Viewing the input vector components as $a_i \in \mathbb{F}_{2^8}$, examining this matrix product reveals we need $a_i b_j$ for all i and all $b_j \in \{1, x, x^2, x^2 + x\}$. That is, we need the result of three distinct linear maps applied to the input. Applying the machinery from Sect. 2.2 yields $t_2 = ax$ and $t_4 = ax^2$, then the final product is $t_6 = t_4 + t_2$. It remains to shuffle these vectors using the following indices corresponding to their positions in the columns of H:

$$t_{\theta 2} = \texttt{0x0E0F0C0D0A0B0809060704050203 0001}$$
$$t_{\theta 4} = \texttt{0x0D0C0F0E09080B0A0504070601000302}$$
$$t_{\theta 6} = \texttt{0x0C0D0E0F08090A0B0405060700010203}$$

and the output is the XOR of these three shuffled vectors with the input. This strategy realizes the entire layer using seven `pshufb`, six `pxor`, two `pand`, and one `psrlq`.

The cyclical permutation π This operation is otherwise analogous to the AES ShiftRows step, but cyclically shifts column i of the matrix downward i positions instead. This map only appears in the key schedule. To illustrate:

$$\begin{bmatrix} \texttt{0} & \texttt{1} & \texttt{2} & \texttt{3} \\ \texttt{4} & \texttt{5} & \texttt{6} & \texttt{7} \\ \texttt{8} & \texttt{9} & \texttt{A} & \texttt{B} \\ \texttt{C} & \texttt{D} & \texttt{E} & \texttt{F} \end{bmatrix} \mapsto \begin{bmatrix} \texttt{0} & \texttt{D} & \texttt{A} & \texttt{7} \\ \texttt{4} & \texttt{1} & \texttt{E} & \texttt{B} \\ \texttt{8} & \texttt{5} & \texttt{2} & \texttt{F} \\ \texttt{C} & \texttt{9} & \texttt{6} & \texttt{3} \end{bmatrix}.$$

Again this requires only one `pshufb` instruction with indices defined as

$$t_\pi = \texttt{0x0306090C0F0205080B0E0104070A0D00}.$$

The key extraction ω This is a linear mapping involving the Vandermonde matrix

$$V = \begin{bmatrix} \texttt{01} & \texttt{01} & \texttt{01} & \texttt{01} \\ \texttt{01} & \texttt{02} & \texttt{02}^2 & \texttt{02}^3 \\ \texttt{01} & \texttt{06} & \texttt{06}^2 & \texttt{06}^3 \\ \texttt{01} & \texttt{08} & \texttt{08}^2 & \texttt{08}^3 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & x & x^2 & x^3 \\ 1 & x^2 + x & x^4 + x^2 & x^6 + x^5 + x^4 + x^3 \\ 1 & x^3 & x^6 & x^5 + x^4 + x^3 + x \end{bmatrix} = \begin{bmatrix} \texttt{01} & \texttt{01} & \texttt{01} & \texttt{01} \\ \texttt{01} & \texttt{02} & \texttt{04} & \texttt{08} \\ \texttt{01} & \texttt{06} & \texttt{14} & \texttt{78} \\ \texttt{01} & \texttt{08} & \texttt{40} & \texttt{3A} \end{bmatrix}$$

and $\omega : a \mapsto V \cdot a$. This map also only appears in the key schedule.

This is quite a different situation compared to θ, where we needed the product of every entry in the matrix with every component of the input vector a. For example, here we need $(x^2 + x)a_i$ only for $4 \leq i < 8$. When computing with 16-component vectors, this kind of selective computation is difficult to accomplish in an elegant fashion.

On the other hand, realizing distinct linear maps as in Sect. 2.2 with the same input amortizes the cost of the first two steps: the nibbles (indices into tables) need be produced only once. In light of this, one strategy is overcomputation by producing $a_i b_j$ for all i and all b_j as distinct entries in V. Computing six of the maps (`02`, `04`, `08`, `14`, `3A`, and `40`) is enough to reach the remaining two with XOR chains (`06` and `78`). This strategy uses 12 `pshufb`, nine `pxor`, two `pand`, and one `psrlq`.

Denote the resulting vectors by r_i; these need to combined at different indices before XOR-summing them to arrive at the result (three `pxor`). For column j of V with entries $[v_{0j}, v_{1j}, v_{2j}, v_{3j}]$ the needed vector is

$$[v_{0j}[a_{4j}, \ldots, a_{4j+3}], v_{1j}[a_{4j}, \ldots, a_{4j+3}], v_{2j}[a_{4j}, \ldots, a_{4j+3}], v_{3j}[a_{4j}, \ldots, a_{4j+3}]].$$

One way to achieve this is through a series of interleaves: `punpckldq` interleaves the lower two 4-byte values in the first operand with those in the second, and `punpckhqdq` the high 8-byte value.

The following illustrates this concept with $j = 1$ where vectors $\{r_1 = a, r_2 = ax, r_6 = a(x^2 + x), r_8 = ax^3\}$ facilitate constructing the vector

$$[\texttt{01}a_4, \texttt{01}a_5, \texttt{01}a_6, \texttt{01}a_7, \texttt{02}a_4, \texttt{02}a_5, \texttt{02}a_6, \texttt{02}a_7, \texttt{06}a_4, \texttt{06}a_5, \texttt{06}a_6, \texttt{06}a_7, \texttt{08}a_4, \texttt{08}a_5, \texttt{08}a_6, \texttt{08}a_7].$$

Here the r_i are filled with dummy data to help observe the interleaving action:

$$\begin{aligned}
r_1 &= \texttt{0x33333333222222221111111100000000}\\
r_2 &= \texttt{0x77777777666666665555555544444444}\\
r_6 &= \texttt{0xBBBBBBBBAAAAAAAA9999999988888888}\\
r_8 &= \texttt{0xFFFFFFFFEEEEEEEEDDDDDDDDCCCCCCCC}\\
t_0 &= \texttt{0x55555551111111144444444400000000 (punpckldq)}\\
t_1 &= \texttt{0xDDDDDDDD99999999CCCCCCCC88888888 (punpckldq)}\\
t_2 &= \texttt{0xDDDDDDDD99999999555555551111111 (punpckhqdq)}.
\end{aligned}$$

These operations accomplish the goal of extracting bytes $v_4, \ldots, v_7$ from each of the given $v = r_i$ to a vector in a specific order corresponding to column j of V. The vectors for other j are obtained similarly with three instructions, but different interleaves. The exception being $j = 0$, using only one `pshufd` to broadcast the lower 4-byte value of the input across the vector (the instruction takes an immediate argument).

The key schedule Given the cipher key K, round keys K^i for $0 \leq i \leq 12$ satisfy $K^r = (\tau \circ \omega \circ \gamma)(\kappa^r)$ where $\kappa^0 = K$ and $\kappa^r = (\sigma[c^r] \circ \theta \circ \pi \circ \gamma)(\kappa^{r-1})$ for $r > 0$, σ is addition in $\mathbb{F}_{2^8}^{16}$, and c^r are vector constants dependent only on S. Note the shared application of γ.

The complete cipher Anubis initializes the state as $\sigma[K^0]$ applied to the input. This gets iteratively transformed through 12 rounds by $\sigma[K^r] \circ \theta \circ \tau \circ \gamma$ where the last round omits θ.

4 A Bit-slice Implementation

Käsper and Schwabe use the SIMD registers to represent eight AES instances running in parallel [KS09]. While these can be unrelated instances with distinct keys, parallel block cipher modes such as counter mode are where this method is particularly interesting: encrypting the next eight counter values under one distinct key in parallel. Eight SIMD registers hold the entire state for these eight instances, but each register represents one bit-slice of the state bytes for all instances.

Naturally, the same approach can be used to implement Anubis in counter mode. Denote 128-bit SIMD registers r_i for $0 \leq i < 7$ each holding bit i of all state bytes. Byte j of r_i holds bit i of the jth state byte for all eight instances, each instance at a fixed offset within these bytes.

With this representation, some of the components from the previous section remain unchanged and are simply iterated for each r_i. For example, τ and the byte shuffles at the end of θ. As this counter mode implementation uses only a single distinct key, the key schedule components stay the same, but the resulting round keys must be subsequently converted into bit-slice format using eight times the storage.

See [KS09, Sect. 4.1] for a brief discussion on data conversion to and from bit-slice format. This implementation uses the same code for said conversion.

The two components that differ significantly in implementation compared to the serial case are γ and θ, the only layers where any time consuming operations are carried out during encryption. A discussion follows.

The nonlinear layer γ The previous serial implementation relies heavily on `pshufb` as a lookup table to realize γ. In contrast, bit-slicing relies on boolean expressions alone to evaluate the S-box, facilitated by access to individual bits of all state bytes collected in one register. Indeed, this is the allure of bit-slicing.

The specification gives boolean expressions for P and Q with 18 gates each, implementing S with 108 gates [BR01a, Appx. B]. This is not significantly lighter than the current smallest published AES S-box with 115 gates [BP09], although the former appeared at inception while the later took roughly a decade of research to whittle down. SSE instructions `pand`, `por`, and `pxor` directly implement said expressions.

The simple construction of S as an SPN using smaller P and Q easily allows the implementation to remain entirely within the working register set: the stack is not required.

The linear diffusion layer θ This component also does not require the stack. Similar to MixColumns, viewing the input and output of θ as matrices one can derive a formula for each byte of the output:

$$b_{ij} = a_{ij} + x(a_{i1-j} + a_{i3-j} + x(a_{i2+j} + a_{i3-j}))$$

where all the subscripts are modulo 4. Each multiplication by x implies three XOR gates for reduction. This leads to a cost of 38 `pxor` and 24 `pshufb`, noteably heavier than the 27 `pxor` and 16 `pshufd` of MixColumns [KS09, Sect. 4.4]. This is compounded by the fact that `pshufd` takes a destination register operand, saving register to register moves; unfortunately `pshufb` does not allow this.

5 The Khazad Cipher

The Khazad block cipher [BR01b] works on 8-byte blocks and uses a 16-byte key. The state is viewed as an element of $\mathbb{F}_{2^8}^8$. The components are almost identical to those of Anubis in many respects. A description of the components follows alongside a description of the first implementation. This sections presents two Khazad implementations, analogous to the two Anubis implementations. Both require a parallel block cipher mode when only one distinct key is used.

5.1 Two Parallel Instances

As the SIMD registers are 16-byte, this section considers implementation of stateful-decryption counter mode (SDCTR) [BK06], although one can also run two unrelated instances in parallel. In this case, the counter is initially set to the 8-byte IV. The implementation here uses most of the machinery developed in Sect. 3. The nonlinear layer γ remains the same and the previous implementation is reused here verbatim. The permutations π and τ in Anubis have no equivalent in Khazad.

The linear diffusion layer θ This linear layer multiplies the input vector by the symmetric matrix

$$H = \begin{bmatrix} 01 & 03 & 04 & 05 & 06 & 08 & 0B & 07 \\ 03 & 01 & 05 & 04 & 08 & 06 & 07 & 0B \\ 04 & 05 & 01 & 03 & 0B & 07 & 06 & 08 \\ 05 & 04 & 03 & 01 & 07 & 0B & 08 & 06 \\ 06 & 08 & 0B & 07 & 01 & 03 & 04 & 05 \\ 08 & 06 & 07 & 0B & 03 & 01 & 05 & 04 \\ 0B & 07 & 06 & 08 & 04 & 05 & 01 & 03 \\ 07 & 0B & 08 & 06 & 05 & 04 & 03 & 01 \end{bmatrix}$$

and $\theta : a \mapsto a \cdot H$. The strategy is the same as the corresponding layer in Anubis. First compute three linear maps (`02`, `04`, and `08`) and derive the remaining maps with XOR chains. The output is the XOR-sum of the input and the seven shuffled vectors resulting from the linear maps. This implementation uses 15 `pxor`, 13 `pshufb`, two `pand` and one `psrlq`.

The key schedule Round keys satisfy $K^r = (\sigma[K^{r-2}] \circ \sigma[c^r] \circ \theta \circ \gamma)(K^{r-1})$ where $0 \leq r \leq 8$ and K^{-2} and K^{-1} are the first and second eight bytes of the key K, respectively. There is no component corresponding to the key extraction ω in Anubis.

The complete cipher Khazad initializes the state as $\sigma[K^0]$ applied to the input. This gets iteratively transformed through eight rounds by $\sigma[K^r] \circ \theta \circ \gamma$ where the last round omits θ.

5.2 Sixteen Parallel Instances

Lastly, the bit-slice implementation of Khazad in counter mode. Khazad works on 8-byte blocks and with 128-bit SIMD registers aligning the bits of bytes in the state, this implies 16 parallel streams. With all the machinery from the previous implementations, the only component to implement is θ. The approach is exactly the same as with the bit-slice Anubis implementation: derive a formula for the output bytes and accumulate the result in output bits iteratively. For each of the eight input bits, this works out to 14 `pxor` and seven `pshufb` to produce a degree-10 polynomial. Similarly the reduction uses a total of 12 `pxor` to clear the three top bits.

6 Results

Table 1 contains the timings for all of the implementations described in this paper. All SSSE3 implementations are in assembly, while the "Table" entry denotes the C reference implementations available with the cipher specifications.

7 Conclusion

This paper presents a number of constant-time and implementations of the Anubis and Khazad block ciphers. The results show that constant-time and efficient are not mutually exclusive with respect to their software implementation. Irrespective of their current usage or lack thereof, the work here also further showcases the potential of a vector-byte shuffle instruction to provide both secure and fast software implementations of cryptosystems. It is surprising that the compact S-box used in Anubis and Khazad has not managed to make its way into other recent cipher designs. Its particularly efficient software implementation here using `pshufb` greatly encourages further use: within these two ciphers themselves or even as a building block for other ciphers.

Table 1. Timing results in cycles per byte.

Cipher	Method	Mode	Instances	Core 2 Quad Q9550
Anubis	SSSE3	CTR	1	23.1
Anubis	SSSE3	CTR	8	11.2
Anubis	Table	CTR	1	20.7
Khazad	SSSE3	CTR	2	18.6
Khazad	SSSE3	CTR	16	10.3
Khazad	Table	CTR	1	19.8

Realizing the threat that timing attacks pose to software implementations, more recent trends in cipher design are away from the rather traditional view of an S-box as a lookup table towards methods that better suit constant-time implementations using native instructions supported by common processors. For example, the Threefish block cipher explicitly states this as a design criteria [FLS+09, Sect. 8.1], using an extremely simple nonlinear function MIX consisting of a rotation, XOR, and addition modulo 2^{64} iterated during a large number of rounds. However, equipped with powerful instructions like `pshufb` it will be interesting to see how cryptologists harness this machinery and what the future holds for cipher design.

References

[ALH10] Diego F. Aranha, Julio López, and Darrel Hankerson. High-speed parallel software implementation of the η_T pairing. In Josef Pieprzyk, editor, *CT-RSA*, volume 5985 of *Lecture Notes in Computer Science*, pages 89–105. Springer, 2010.

[Ber04] Daniel J. Bernstein. Cache-timing attacks on AES. `http://cr.yp.to/papers.html#cachetiming`, 2004.

[BK06] M. Bellare and T. Kohno. The secure shell (SSH) transport layer encryption modes. RFC 4344, January 2006.

[BP09] Joan Boyar and Rene Peralta. New logic minimization techniques with applications to cryptology. Cryptology ePrint Archive, Report 2009/191, 2009. `http://eprint.iacr.org/`.

[BR01a] Paulo S. L. M. Barreto and Vincent Rijmen. The Anubis block cipher. `http://www.larc.usp.br/~pbarreto/anubis-tweak.zip`, 2001.

[BR01b] Paulo S. L. M. Barreto and Vincent Rijmen. The Khazad legacy-level block cipher. `http://www.larc.usp.br/~pbarreto/khazad-tweak.zip`, 2001.

[CG09] Christophe Clavier and Kris Gaj, editors. *Cryptographic Hardware and Embedded Systems - CHES 2009, 11th International Workshop, Lausanne, Switzerland, September 6-9, 2009, Proceedings*, volume 5747 of *Lecture Notes in Computer Science*. Springer, 2009.

[FLS+09] Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, and Jesse Walker. The Skein hash function family. Submission to NIST (Round 2), 2009.

[Ham09] Mike Hamburg. Accelerating AES with vector permute instructions. In Clavier and Gaj [CG09], pages 18–32.

[KS09] Emilia Käsper and Peter Schwabe. Faster and timing-attack resistant AES-GCM. In Clavier and Gaj [CG09], pages 1–17.

[Pre02] B. Preneel. The NESSIE project: towards new cryptographic algorithms. In *Information Security Applications, 3rd International Workshop, WISA 2002*, pages 16–33, 2002.