

K-Cipher: A Low Latency, Bit Length Parameterizable Cipher

Michael Kounavis, Sergej Deutsch, Santosh Ghosh, and David Durham

Intel Labs, Intel Corporation, 2111, NE 25th Avenue, Hillsboro, OR 97124
Email: {michael.e.kounavis, sergej.deutsch, santosh.ghosh, david.durham}@intel.com

November 2019

Abstract

We present the design of a novel low latency, bit length parameterizable cipher, called the “K-Cipher”. K-Cipher is particularly useful to applications that need to support ultra low latency encryption at arbitrary ciphertext lengths. We can think of a range of networking, gaming and computing applications that may require encrypting data at unusual block lengths for many different reasons, such as to make space for other unencrypted state values. Furthermore, in modern applications, encryption is typically required to complete inside stringent time frames in order not to affect performance. K-Cipher has been designed to meet these requirements. In the paper we present the K-Cipher design and discuss its rationale.

1 Introduction

The paper presents a novel block cipher design which is lightweight, hardware efficient, as well as bit length parameterizable. Our cipher is called the “K-Cipher”. In the K-Cipher design, the block length is not fixed (e.g., fixed to 128 bits), neither takes values from a small set of options (e.g., 64 or 128 bits). Instead the block length can take any arbitrary value between an upper and a lower bound and is an input parameter passed into the cipher.

The need for such block cipher comes from the requirements of modern applications. To address a wide range of vulnerabilities, applications employ cryptographic mechanisms to provide confidentiality and integrity. A common characteristic of applications is that their features have been designed to work well with some existing or specially designed cipher (e.g., AES [2] or QARMA [6]).

In the paper we argue for the need of a new block cipher family that is more agile and easier to tune to the needs of a particular application or hardware. A new requirement we introduce is that the block length of the cipher should be an input parameter to the encryption operation. Furthermore, the specification and security analysis of block ciphers should be, to some degree, block length independent. Such independence should substantially surpass what is accomplished today, where the block length is selected from a small number of options. Our design supports encryption at arbitrary block lengths which, in our prototype take all values from 24 to 1024 at increments of 1.

The reason why we believe such requirement is important, is because applications operate on a variety of data of different lengths. The encrypted portions of such data may be of arbitrary lengths as well. Rather than designing an application with a fixed block length cipher in mind, we argue for the opposite. That is, to have the ability to arbitrarily tune the

block length of a block cipher to meet the needs of a particular application or hardware. This is something that is not possible today.

Another property of our cipher family is that it supports ultra low latency encryption in hardware. It is fair to state that the landscape of today’s block ciphers, which are either standard (e.g., AES [1]), or are to be standardized though the current NIST lightweight cryptography competition [7], does not include any cipher that simultaneously meets the two requirements stated: (i) support for full confusion and diffusion over arbitrary block lengths, as part of the encryption and decryption operations; and (ii) support for ultra low latency encryption and decryption operations in hardware.

For example, past lightweight cipher designs such as NSA’s “Simon” and “Speck” [8], or designs like “PRINCE” [9] are not bit length parameterizable. Furthermore, such ciphers support critical paths, which can get even smaller with the design proposed here. For example, the PRINCE rounds, even though contain simpler Sbox transformations, and simpler Mix Columns matrices when compared to AES, still require several clocks in the critical path, in typical client and server frequencies. Others ciphers like Simon employ a simple Feistel structure, which includes only elementary logical AND, XOR and rotation operations over 32-72 rounds. Furthermore, Simon supports only 5 fixed lengths: i.e., 32, 48, 64, 96 and 128 bits. Last, some submissions to the NIST lightweight cryptography competition, such as TinyJAMBU [11] or Xoodyak [12] support encryption at both very high speeds and arbitrary plaintext lengths. However, they do not fully diffuse the bits of their plaintext input into the bits of the ciphertext output. Instead, for a wide range of plaintext inputs, they merely add in GF(2) the input plaintext bits into the bits of some sponge state.

2 Overview

We envision that K-Cipher will be a useful tool for the encryption and decryption of data of varying lengths, performed by many different types of applications such as running in networking devices, gaming consoles, servers, low power clients and so on. K-Cipher supports fast encryption based on a novel confusion-diffusion network, which we discuss in this paper. The primitives it employs are: (i) block wide addition with carries. In this operation, the carry out bit is ignored. The operation is invertible, its inverse being subtraction with borrows. For the subtraction operation, K-Cipher just ignores the borrow out bit; (ii) block wide bit level reordering; and finally (iii) wide Sbox substitution, which is realized as inversion in a binary Galois Field.

K-Cipher is designed to support confidentiality at desired security levels by employing the least possible number of rounds, which, in the proposed design is equal to 2. We can also consider this number to be a configurable parameter. Future cryptanalysis will determine the exact number of rounds which will be necessary to support desired security levels.

Substitution box lengths are determined by the block length. For example, to diffuse across 32 bits, K-Cipher uses 8-bit substitution boxes, as $32 \text{ bits} < (8 \text{ bits})^2 \text{ rounds}$. To diffuse across 128 bits, K-Cipher uses 16-bit substitution boxes, as $128 \text{ bits} < (16 \text{ bits})^2 \text{ rounds}$. All primitives of the K-Cipher apply to a wide range of block lengths. Arbitrary ciphertext lengths are supported using Galois field inverters of varying lengths, the sum of which is equal to the requested input and ciphertext lengths. In our implementation the employed Galois field inverters have fixed lengths but the last one, the length of which is determined by the block length.

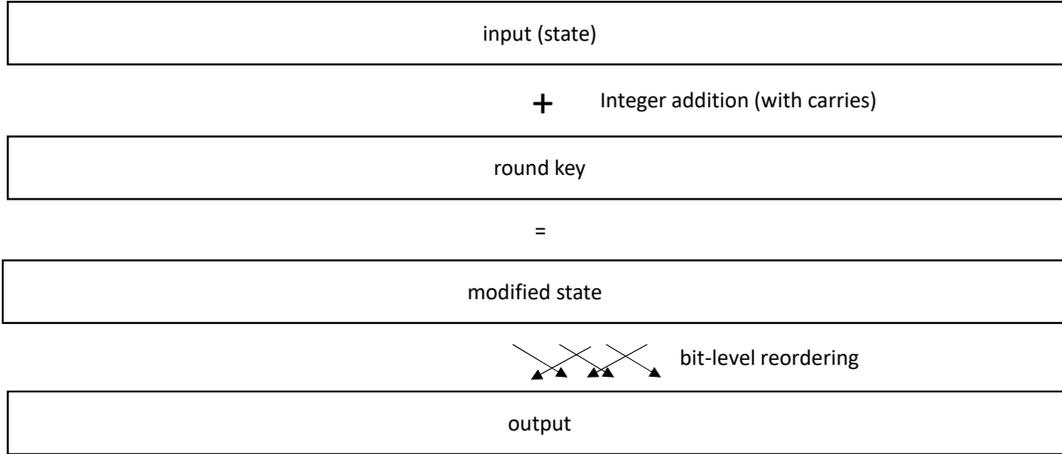


Figure 1: The Aggressive Adder Component of the K-Cipher Round

3 K-Cipher Design

3.1 Notation

We will be denoting as $\langle a_n \dots a_1 a_0 \rangle$ a bit string of length n consisting of bits a_0, \dots, a_n , where a_0 is the least significant bit of the string and a_n is its most significant bit. Similarly, we will be denoting as $N(\langle a_n \dots a_1 a_0 \rangle)$ the binary number which is represented by the string $\langle a_n \dots a_1 a_0 \rangle$.

3.2 The Aggressive Adder

Starting with Figure 1, we illustrate one of the basic components of the K-Cipher round called the “aggressive adder”. The aggressive adder accepts as input some state, and then adds to this state a round key. The addition performed is not in the typical $\text{GF}(2)$ arithmetic, but is in the integer arithmetic. Integer addition, if seen as a bit-logical operation, performs strong mixing of its input bits, in order to produce the bits of the output. The mixing performed demonstrates regularity due to the use of carry values. By “mixing” in this document we mean computations on single bit values that involve a plurality of AND, OR, NAND, NOR or XOR operations.

For example lets consider that we add the numbers $N(\langle a_3 a_2 a_1 a_0 \rangle)$ and $N(\langle b_3 b_2 b_1 b_0 \rangle)$ with each other and with some input carry value c_0 . The first bit of the result is equal to $a_0 \oplus b_0 \oplus c_0$. The carry produced from the addition of the first two bits is equal to $a_0 b_0 \oplus b_0 c_0 \oplus a_0 c_0$. Similarly, the second bit of the result is $a_1 \oplus b_1 \oplus a_0 b_0 \oplus b_0 c_0 \oplus a_0 c_0$ and the carry produced from the addition of the second two bits is equal to $a_1 b_1 \oplus a_1 a_0 b_0 \oplus a_1 b_0 c_0 \oplus a_1 a_0 c_0 \oplus b_1 a_0 b_0 \oplus b_1 b_0 c_0 \oplus b_1 a_0 c_0$. Moving on to the addition of the third least significant bits of the input, the same pattern of computation is repeated. The input bits are XOR-ed with each other and with the input carry, in order to produce the output bit. Furthermore, the input bits are multiplied with each other in $\text{GF}(2)$ arithmetic (i.e., undergo a logical AND operation) and with the input carry and, subsequently, the products are XOR-ed with each other in order to produce the output carry. The third least significant bit of the result, as computed using this pattern, is $a_2 \oplus b_2 \oplus a_1 b_1 \oplus a_1 a_0 b_0 \oplus a_1 b_0 c_0 \oplus a_1 a_0 c_0 \oplus b_1 a_0 b_0 \oplus b_1 b_0 c_0 \oplus b_1 a_0 c_0$. The third output carry

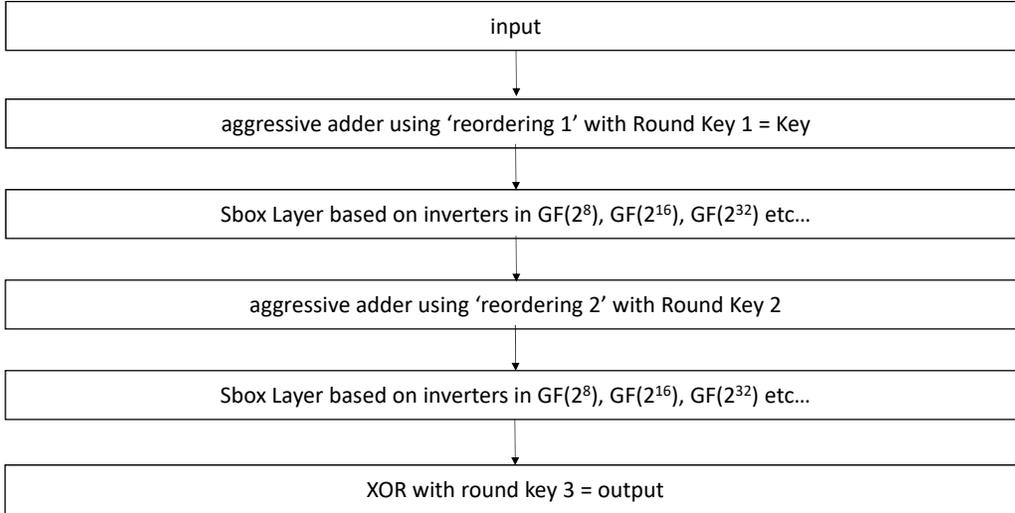


Figure 2: Two Round K-Cipher Specification

is $a_2b_2 \oplus a_2a_1b_1 \oplus a_2a_1a_0b_0 \oplus a_2a_1b_0c_0 \oplus a_2a_1a_0c_0 \oplus a_2b_1a_0b_0 \oplus a_2b_1b_0c_0 \oplus a_2b_1a_0c_0 \oplus b_2a_1b_1 \oplus b_2a_1a_0b_0 \oplus b_2a_1b_0c_0 \oplus b_2a_1a_0c_0 \oplus b_2b_1a_0b_0 \oplus b_2b_1b_0c_0 \oplus b_2b_1a_0c_0$.

From the logical expressions above, it becomes evident that the mixing performed by the addition with carries stage, as measured by the number of GF(2) products which are XOR-ed with each other, gets only stronger as one moves from the least significant bit of the result toward the most significant bit. In fact, it grows stronger exponentially. It is easy to show that the n -th output bit for the result is produced by XOR-ing $2^n + 1$ products.

To destroy the regularity which characterizes the addition with the carries stage, the aggressive adder performs a bit level reordering operation on the addition output. Such reordering operation places the output bits coming from the integer adder in a seemingly random order, so that the number of GF(2) products of the logic equation of the result no longer increases monotonically but instead increases and decreases in a pseudorandom manner. Furthermore the bit level reordering operation aids the subsequent wide substitution stage, shown in Figures 2 and 3, ensuring that each bit of the output of the K-Cipher results from mixing all bits of the input with all bits of the key. The addition with carries is a bit length independent operation. Its specification is independent of the length of the inputs. It is also invertible, its inverse being the subtraction with borrows. Any final carry out or borrow out signals produced from such operations are ignored.

3.3 Two Round K-Cipher Specification

Moving onto Figure 2, the processing steps of an instance of the K-Cipher employing two rounds are shown. The first round consists of an aggressive adder stage that performs reordering using a first index sequence denoted as “reordering 1” and a first round key denoted as “Round Key 1”. The aggressive adder of the first round is followed by a wide substitution stage denoted as “Sbox Layer” in the figure. The second round consists of an aggressive adder stage also performing reordering using a second index sequence denoted as “reordering 2” and a second round key denoted as “Round Key 2”. The second aggressive adder is followed by a second

wide substitution stage. The processing steps of the K-Cipher conclude with an XOR operation performed between the cipher state and a third round key denoted as “round key 3”.

The Sbox layer performs the following steps. It first divides its input N bits into blocks of M bits. Let’s assume for now that N is a multiple of M . The cases where N is not a multiple of M are discussed further below. If N is a multiple of M , the Sbox layer employs an array of N/M inverters in $\text{GF}(2^M)$ arithmetic which replace their input bits with the bits of the inverse in $\text{GF}(2^M)$. Inversion in the Galois Field arithmetic $\text{GF}(2^M)$ is another operation supporting strong bit mixing. The mixing performed by the Galois Field inverters employed by the K-Cipher does not demonstrate the regularity of addition with carries and is in fact pseudo-random. K-Cipher is designed to support strong encryption security by employing additions and inversions in two unrelated types of arithmetic (i.e., Galois Field and integer) and by combining those into sequences of few rounds. K-Cipher rounds, despite the fact that they are few, succeed in strongly mixing their input and key bits potentially thwarting differential, algebraic and other types of attacks, as suggested by analysis, which we currently carry out and will discuss in a future paper.

The Sbox layer, as defined so far, is bit length independent provided that the length of the state of the cipher N is a multiple of the width of the inverters employed M . In this case the specification of the cipher is generic and each wide substitution stage employs N/M inverters. If N is not a multiple M , then these situations can be handled as shown in Figure 3. In the figure there are m substitution boxes of width M which are employed, plus one more of width $K = N - m \cdot M$, where K is non-zero. The substitution stage employs m inverters in the in $\text{GF}(2^M)$ arithmetic and one inverter in the $\text{GF}(2^K)$ arithmetic handling the last K bits of the cipher state.

The generation of the index sequences employed by the K-Cipher, which support bit level reordering, is accomplished by the following algorithm: The algorithm first determines the number of times d it needs to iterate over the bits of a substitution box in order to distribute these bits over all substitution boxes. We refer to these bits of a substitution box as “bits-to-be-reordered”. The parameter d is equal to $\text{ceil}(M/b)$. Then, for each of the d iterations, the algorithm generates a random sequence of numbers. These are the indexes of the substitution boxes where the “bits-to-be-reordered”, associated with the current iteration, will be placed. Subsequently, for each “bit-to-be-reordered”, the algorithm picks a bit position at random from among the empty bit positions in the input bit’s target substitution box and assigns this position to the bit. This last step is repeated for all iterations of a triply nested loop executed by the algorithm.

One can show that the algorithm produces sequences of indexes that are both correct and proper. By correct we mean that every bit of the input is placed in a different bit position of the output and there is no input which is omitted from the output. By proper we mean that if such reordering operations are combined with wide substitution operations, then, after $\log_M N$ rounds all bits of the input have been fully mixed with each other, even if additions with carries are absent.

3.4 Key Schedule and Use of Tweaks

The key schedule algorithm of the K-Cipher has the same structure as the cipher itself shown in Figure 2. However, it uses a different set of reordering sequences. Furthermore, instead of adding key bits into its state, it adds a plurality of constants. This is to expand a single key into a sequence of three round keys used by the cipher rounds.

Last we discuss, how the K-Cipher can be turned into a tweakable block cipher. This is done as follows: As said above, the key schedule of the K-Cipher involves 3 round keys. A

- [4] *AMD Secure Encrypted Virtualization (SEV)*, <https://developer.amd.com/sev/>, 2016.
- [5] M. Rutland, *ARM v8.3 Pointer Authentication*, presentation, available online at https://events.static.linuxfound.org/sites/events/files/slides/slides_23.pdf, 2017.
- [6] R. Avanzi, *The QARMA Block Cipher Family*, Cryptology ePrint Archive: Report 2016/444.
- [7] *NIST Lightweight Cryptography Competition*, available online at <https://csrc.nist.gov/projects/lightweight-cryptography>
- [8] R. Beaulieu, D. Shors, J. Smith, S. Treatman-Clark, B. Weeks, and L. Wingers *The Simon and Speck Families of Lightweight Block Ciphers*, Cryptology ePrint Archive: Report 2013/404.
- [9] J. Borghoff, A. Canteaut, T. Guneysu, E. B. Kavun, M. Knezevic, L. R. Knudsen, G. Leander, V. Nikov, C. Paar, C. Rechberger, P. Rombouts, S. S. Thomsen, and T. Yalcin, *PRINCE: a low-latency block cipher for pervasive computing applications*, ASIACRYPT 2012, Proceedings of the 18th international conference on The Theory and Application of Cryptology and Information Security, Pages 208-225, Beijing, China December 02 - 06, 2012.
- [10] Y. Dodis, T. Liu, M. Stam, J. Steinberger, *Indifferentiability of Confusion-Diffusion Networks*, hskip 1em plus 0.5em minus 0.4emCryptology ePrint Archive: Report 2015/680.
- [11] H. Wu, and T. Huang, *TinyJAMBU: A Family of Lightweight Authenticated Encryption Algorithms*, Submission to the NIST Lightweight Cryptography Competition, available online at <https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/round-1/spec-doc/TinyJAMBU-spec.pdf>.
- [12] J. Daemen, S. Hoffert, M. Peeters, G. Van Assche, and R. Van Keer, *Xoodyak, a lightweight cryptographic scheme*, Submission to the NIST Lightweight Cryptography Competition, available online at <https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/round-1/spec-doc/Xoodyak-spec.pdf>.