

The MAC In The Box (MITB) Project


Mike Gordon, Robert Künnemann, Graham Steel¹

Abstract

The “Mac in the Box” (MITB) project was conceived at a workshop [3] where Graham Steel gave a talk showing that API insecurities in crypto tokens could be found using the Tookan tool [2] (this work built on earlier research by Clulow, Bond and others [4]). At the workshop Steel and Gordon started a research collaboration to see if would be feasible to design token implementations that could be formally verified to have security properties establishing their trustworthiness (in particular, to be unbreakable using Tookan-like tools). The project aims to complement research on breaking existing tokens by designing, verifying and building a family of simple, low cost, open source hardware devices for hashing passwords and other data. The research goal is to explore removing the assumption “that the implementation is correct with respect to the documentation” [15]. Such a device can be used, for example, to store passwords more securely. By protecting a secret that is used in the computation of password hashes, a so-called *local parameter*, an attacker performing a brute-force on the password database is considerably slowed down, being only able to compute hashes at a rate the device dictates. A proof-of-concept toy prototype MITB specification and implementation have been formalised and the implementation verified to meet the specification using the HOL4 proof assistant [1].

1 Introduction to the MITB device

MITB is a standalone device that computes a MAC using the KECCAK sponge function [5]. It has two 1-bit control inputs `skip_inp`, `move_inp`, two data inputs `block_inp` and `size_inp`, a 1-bit control output `ready_out` and a data output `digest_out`.



MITB is parametrised on three numbers (r, c, n) and a permutation function f . These are part of the KECCAK specification (see Section 3). An actual device would be manufactured with specific values for the parameters.

The input `block_inp` is r -bits wide and the output `digest_out` is n -bits wide. The input `size_inp` has sufficient bits to represent a number of size r or less. For convenience it is modelled as a number rather than a bitstring. Truth-values T, F model bits 1, 0, respectively.

MITB runs continuously after being switched on. It is implemented as a state-machine using combinational logic and registers (see Section 5). All a user can observe (assuming tamper-resistant manufacture) are the sequences of values appearing on the outputs `ready_out` and `digest_out`, which depend on the values input via `skip_inp`, `move_inp`, `block_inp` and `size_inp`.

From a user’s point of view MITB can be in either of two states: *ready* or *absorbing*. It powers up into state *ready*. The 1-bit output `ready_out` indicates whether the state is *ready* (T output) or *absorbing* (F output).

¹Listed in alphabetical order.

The input `skip_inp` ‘freezes’ MITB: holding it `T` stops the state changing on successive cycles. The input `move_inp` causes the state to change on the next cycle; in particular it is used to signal that MITB should start absorbing a message.

The MAC of a message M is specified as the KECCAK hash of the result of concatenating a secret key onto the front of the message, i. e. the hash of $key||M$, where $||$ denotes bitstring concatenation. The hash algorithm is defined in Section 3.2. The protocol for using MITB to compute the MAC of a message is described below. The main correctness property of the device is that if the specified protocol is used to input a message then its MAC will appear on `digest_out`. The main security property is that no matter what inputs are supplied, the secret key cannot be revealed. These properties will be expressed as constraints on what sequences of inputs and outputs are possible using a temporal logic notation.

MITB has a permanent memory for holding an r -bit secret key. The key can be set or changed by holding both `skip_inp` and `move_inp` `F` in the *ready* state. The data being input on `block_inp` then overwrites the stored key. As long as `skip_inp` and `move_inp` are held `F`, the stored key is updated on each cycle (discussed in Section 6).

MITB is ready to compute the MAC of a message in state *ready*. The protocol for computing the MAC of M is as follows ($|B|$ denotes the number of bits in B):

1. The user splits M into a sequence of blocks, $M = B_1||B_2||\dots||B_{m-1}||B_m$, such that all blocks except the last one are r -bits wide, i. e. $|B_i| = r$ for $1 \leq i < m$ and $|B_m| < r$. If r divides exactly into $|M|$, then B_m is taken to be the empty block (so $|B_m| = 0$).
2. When `ready_out` is `T` the user puts MITB into the *absorbing* state by inputting `F` on `skip_inp` and `T` on `move_inp` (`block_inp` and `size_inp` are ignored during this step).
3. Starting on the next cycle, and continuing for m cycles, the user inputs `F` on both `move_inp` and `skip_inp`, B_i on `block_inp` and $|B_i|$ on `size_inp`, where $1 \leq i \leq m$. During this time `F` will be output on `ready_out`.
4. After inputting B_m , the user keeps inputting `F` on `skip_inp` and `move_inp` until `ready_out` becomes `T`. On the cycle when this happens the hash of $key||M$ will appear on `digest_out`. The number of cycles taken depends on $|B_m|$. If $|B_m| \neq r-1$ then `ready_out` will become `T` on the cycle after B_m is input. If $|B_m| = r-1$ then `ready_out` will become `T` the cycle after the cycle after B_m is input.

The timing diagrams below illustrate this protocol. The MAC computation starts at cycle t and `X` means ‘don’t care’ (if `skip_inp` is `T` then other inputs are ignored). The first line shows the cycle count. The message M is split by the user into blocks B_1, \dots, B_m as described above. When MITB is started the volatile memory holds zeros, so `digest_out` is all zeros too.

If the size of the last block B_m is not $r-1$ (i. e. $|M| \bmod r \neq r-1$).

| | | | | | | | | | | |
|-------------------------|----------------|----------------|----------------|-----|----------------|----------------|----------------|-----|----------------|-----------------------|
| Cycles: | 0 | 1 | 2 | ... | t | t+1 | t+2 | ... | t+m | t+(m+1) |
| <code>skip_inp</code> | <code>T</code> | <code>T</code> | <code>T</code> | ... | <code>F</code> | <code>F</code> | <code>F</code> | ... | <code>F</code> | <code>X</code> |
| <code>move_inp</code> | <code>X</code> | <code>X</code> | <code>X</code> | ... | <code>T</code> | <code>F</code> | <code>F</code> | ... | <code>F</code> | <code>X</code> |
| <code>block_inp</code> | <code>X</code> | <code>X</code> | <code>X</code> | ... | <code>X</code> | B_1 | B_2 | ... | B_m | <code>X</code> |
| <code>size_inp</code> | <code>X</code> | <code>X</code> | <code>X</code> | ... | <code>X</code> | r | r | ... | $ B_m $ | <code>X</code> |
| <code>ready_out</code> | <code>T</code> | <code>T</code> | <code>T</code> | ... | <code>T</code> | <code>F</code> | <code>F</code> | ... | <code>F</code> | <code>T</code> |
| <code>digest_out</code> | 0 | 0 | 0 | ... | 0 | 0 | 0 | ... | 0 | $\text{Hash}(key M)$ |

If the last block has size $r-1$.

| Cycles: | 0 | 1 | 2 | ... | t | t+1 | t+2 | ... | t+m | t+(m+1) | t+(m+2) |
|------------|---|---|---|-----|---|-----|-----|-----|-----|---------|--------------|
| skip_inp | T | T | T | ... | F | F | F | ... | F | F | F |
| move_inp | X | X | X | ... | T | F | F | ... | F | F | F |
| block_inp | X | X | X | ... | X | B1 | B2 | ... | Bn | X | X |
| size_inp | X | X | X | ... | X | r | r | ... | r-1 | X | X |
| ready_out | T | T | T | ... | T | F | F | ... | F | F | T |
| digest_out | 0 | 0 | 0 | ... | 0 | 0 | 0 | ... | 0 | 0 | Hash(key M) |

Once the MAC, $\text{Hash}(key||M)$, has been computed, it can be forced to persist on the output `digest_out` by holding T on `skip_inp`.

The security property that MITB guarantees is that no matter what inputs are supplied, the value output on `digest_out` is always either an n -bit bitstring representing 0 or the hash of some message. Assuming KECCAK is secure, then MITB does not reveal any information about the stored key.

2 Formal specification using temporal logic

An implementation of MITB is modelled with a next-state function `MITB` (defined in Section 5) that gives the next state s' when input i is received in state s . The observable outputs are determined by the current state, so the model is a Moore machine. MITB is parametrised on the KECCAK parameters (r, c, n) and permutation function f , so the next-state function `MITB` takes these as arguments, hence $s' = \text{MITB}(r, c, n) f(i, s)$.²

A user of MITB can supply inputs on `skip_inp`, `move_inp`, `block_inp` and `size_inp` and observe the resulting outputs on `ready_out` and `digest_out`, where `ready_out` is a boolean value showing which state the device is in and `digest_out` is an n -bit word consisting of zeros in the *absorbing* state and the bottom n bits of the volatile memory in the *ready* state.

Using a standard method [12], the implementation of MITB will be represented by a formula:

$$\text{MITB_IMP } key(r, c, n) f \\ (\text{cntl_sig}, \text{pmem_sig}, \text{vmem_sig}) \\ (\text{skip_inp}, \text{move_inp}, \text{block_inp}, \text{size_inp}, \text{ready_out}, \text{digest_out})$$

where `MITB_IMP` is a predicate defined in terms of `MITB`.

The 6-tuple $(\text{skip_inp}, \text{move_inp}, \text{block_inp}, \text{size_inp}, \text{ready_out}, \text{digest_out})$ has components that are functions from time to values representing the sequence of values on the four user-supplied inputs and two device-generated outputs. The triple $(\text{cntl_sig}, \text{pmem_sig}, \text{vmem_sig})$ contains the function `cntl_sig` representing a sequence of values that the control register can take, `pmem_sig` representing values of the permanent memory and `vmem_sig` representing the values of the volatile memory. The notation $\text{MITB_IMP}(r, c, n) f \models \phi$ means that all runs of MITB satisfy the property ϕ .

$$\text{MITB_IMP}(r, c, n) f \models \phi \Leftrightarrow \forall \sigma_1 \sigma_2. \text{MITB_IMP}(r, c, n) f \sigma_1 \sigma_2 \Rightarrow \phi(\sigma_1, \sigma_2)$$

where the quantified variable σ_1 ranges of triples of state functions and σ_2 ranges over 6-tuples of port functions. Properties ϕ are expressed as Linear Temporal Logic (LTL) formulas. The specification given in Section 4 consists of a conjunction of temporal logic formulas that all executions of MITB are required to satisfy.

²Following standard λ -calculus notation, functions may be curried and brackets around arguments omitted.

2.1 Atomic formulas

The *atomic formulas* `skiplnp(b)`, `movelnp(b)`, `blocklnp(bs)`, `sizelnp(len)`, `readyOut(b)`, `digestOut(bs)` are true of an execution σ if and only if the argument is the value of the corresponding input or output at cycle 0 of σ . Here b ranges over bits (modelled as truth-values), bs ranges over bitstrings (modelled as lists of truth-values) and len ranges over natural numbers. For example, `MITB_IMP (r, c, n) f \models readyOut T` specifies that any execution of MITB must output T on `ready_out` during the first cycle.

Non-atomic properties specify relationships between values input and output at times other than the first cycle; for this the following temporal operators are used.

2.2 Temporal specification operators

The operators below are mostly standard concepts from LTL, though sometimes with different names from those commonly used. Syntactically these operators combine temporal formulas ϕ , ϕ_1 , ϕ_2 etc. Semantically they are predicates on pairs of tuples of functions as in $\phi(\sigma_1, \sigma_2)$ as occurring in the definition of `MITB_IMP (r, c, n) f \models ϕ` .

The first table defines ‘lifted’ logical operators that are useful for combining temporal formulas built using the temporal operator in the later tables.

Lifted logical operators

| | |
|--|--|
| <code>Bool <i>b</i></code> | $= \lambda\sigma. b$ |
| <code>Not ϕ</code> | $= \lambda\sigma. \neg(\phi \sigma)$ |
| <code>(ϕ_1 And ϕ_2)</code> | $= \lambda\sigma. (\phi_1 \sigma) \wedge (\phi_2 \sigma)$ |
| <code>(ϕ_1 Or ϕ_2)</code> | $= \lambda\sigma. (\phi_1 \sigma) \vee (\phi_2 \sigma)$ |
| <code>(ϕ_1 Implies ϕ_2)</code> | $= \lambda\sigma. (\phi_1 \sigma) \Rightarrow (\phi_2 \sigma)$ |
| <code>(Forall <i>x</i>. ϕ)</code> | $= \lambda\sigma. \forall x. \phi \sigma$ |
| <code>(Exists <i>x</i>. ϕ)</code> | $= \lambda\sigma. \exists x. \phi \sigma$ |

If π is a function representing a sequence of values, then $\pi \downarrow n = \lambda t. \pi(t + n)$, which is π with n elements chopped off the front. If σ is a tuple of functions, then $\sigma \downarrow n$ is defined recursively through its tuple structure: $(\sigma_1, \dots, \sigma_m) \downarrow n = (\sigma_1 \downarrow n, \dots, \sigma_m \downarrow n)$. `Next` (see below) uses this.

Standard LTL operators

| | |
|--|---|
| <code>Next ϕ</code> | $= \lambda\sigma. \phi(\sigma \downarrow 1)$ |
| <code>Always ϕ</code> | $= \lambda\sigma. \forall t. \phi(\sigma \downarrow t)$ |
| <code>Sometime ϕ</code> | $= \lambda\sigma. \exists t. \phi(\sigma \downarrow t)$ |
| <code>ϕ_1 Until ϕ_2</code> | $= \lambda\sigma. \exists t_1. \phi_2(\sigma \downarrow t_1) \wedge \forall t_2. t_2 < t_1 \Rightarrow \phi_1(\sigma \downarrow t_2)$ |

An additional operator `UntilN(t)`, parametrised on a number t , is defined recursively in terms of the standard operators above: `ϕ_1 UntilN(t) ϕ_2` is like `ϕ_1 Until ϕ_2` , but with t specifying the exact number of cycles taken to reach a state in which ϕ_2 is true.

Cycle counting Until operator

| | |
|---|--|
| <code>ϕ_1 UntilN(0) ϕ_2</code> | $= \phi_2$ |
| <code>ϕ_1 UntilN(<i>t</i>+1) ϕ_2</code> | $= \phi_1$ And Next(<code>ϕ_1 UntilN(<i>t</i>) ϕ_2</code>) |

It is easy to show by induction on t that `ϕ_1 Until ϕ_2` = `Exists t. ϕ_1 UntilN(t) ϕ_2` .

3 The KECCAK sponge function

In this section relevant parts of the KECCAK algorithm are described and then formalised (see Section 3.2). In order to make it easy to check that the KECCAK algorithm is being correctly formalised, extracts from the specification will be cut-and-pasted from the official KECCAK reference document [5]. These, and other imported extracts, are included in boxes such as:

KECCAK (pronounced [kɛtʃak]) is a family of sponge functions [8] that use as a building block a permutation from a set of 7 permutations. In this chapter, we introduce our conventions and notation, specify the 7 permutations underlying KECCAK and the KECCAK sponge functions. We also give conventions for naming parts of the KECCAK state.

Note that the citation “[8]” in the box above refers to the citations in the KECCAK reference document, not to the references at the end of this paper.

KECCAK is based on a ‘sponge construction’ in which an arbitrary length message is iteratively ‘absorbed’ into a finite state. The number of iterations depend on the length of the message. Once all of the message has been absorbed, the resulting state can be ‘squeezed’ to extract a digest. For the general KECCAK algorithm this squeezing can also be iterative; the number of iterations depending on the desired length of the digest. However, for the SHA-3 application, the digest length is such that no squeeze iterations are actually needed (see also Section 3.2.4).

KECCAK is a family of algorithms. Each algorithm in the family corresponds to a choice of values for parameters r , called the *bitrate* and c , called the *capacity*. $\text{KECCAK}[r, c]$ denotes the specific algorithm for the indicated parameter values. SHA-3 recognises four instances of KECCAK, namely $\text{KECCAK}[1152, 448]$, $\text{KECCAK}[1088, 512]$, $\text{KECCAK}[832, 768]$, $\text{KECCAK}[576, 1024]$. For each of these instances, a length n of digest is specified, namely: 224, 256, 384, 512, respectively. $[\text{KECCAK}[r, c]]_n$ denotes $\text{KECCAK}[r, c]$ with a digest of length n . Thus:

$$(r, c, n) \in \{ (1152, 448, 224), (1088, 512, 256), (832, 768, 384), (576, 1024, 512) \}$$

SHA-3 recommends using the strongest parameter values; smaller values of the parameters are for testing and, possibly, lightweight hashing. The state of the SHA-3 sponge algorithm thus consists of $r + c = 1600$ bits, which is called the *width* and denoted by b . Note that $b = 25 \times 2^6$. More generally, the width b is defined to be 25×2^l , where l is another parameter that can take on one of the seven values in $\{0, 1, 2, 3, 4, 5, 6\}$. The seven choices for l correspond to the “7 permutations” mentioned in the box above. The SHA-3 instance of KECCAK fixes the value of l to be 6. The initial value of the state before a message has been absorbed by the sponge algorithm is 0^b , i.e. each of the 1600 state bits is 0.

3.1 The sponge algorithm

KECCAK is based on a function $f : \mathbb{Z}_2^b \rightarrow \mathbb{Z}_2^b$ that permutes the state, where $\mathbb{Z}_2 = \{0, 1\}$ is the set of the two bits 0 and 1, and \mathbb{Z}_2^b is the set of bitstrings of length b ($b = 1600$ for SHA-3).

The sponge algorithm applies the state permutation f on each iteration of absorbing a message M into the state. The absorption algorithm is outlined below. A detailed formal specification of f is not given here; f is treated as an uninterpreted parameter in the specification and proofs. This is discussed in Section 6.

3.1.1 Padding

The algorithm is parametrised on numbers r and c , as discussed above. The first step is to pad the message so that it can be split into an exact number, k say, of blocks of length r . This is done by appending “a single bit 1 followed by the minimum number of bits 0 followed by a single bit 1 such that the length of the result is a multiple of the block length”. If the message length is already a multiple of the block length r , then it is still padded by appending $10^{r-2}1$, where 0^p denotes a bitstring of length p with each bit being 0. If the message needs only one bit added to make its length a multiple of r , then $10^{r-1}1$, which has length $r+1$, is appended. Thus padding appends “at least 2 bits and at most the number of bits in a block plus one”.

In the KECCAK reference, the result of padding a message M using a block size x is denoted by $M||\text{pad}[x](|M|)$ and the specific padding described above is called *multi-rate* padding and denoted by $\text{pad}10^*1$. Here is the actual text from the reference:

For the padding rule we use the following notation: the padding of a message M to a sequence of x -bit blocks is denoted by $M||\text{pad}[x](|M|)$. This notation highlights that we only consider padding rules that append a bitstring that is fully determined by the bitlength of M and the block length x . We may omit $[x]$, $(|M|)$ or both if their value is clear from the context.

KECCAK makes use of the *multi-rate* padding.

Definition 1. Multi-rate padding, denoted by $\text{pad}10^*1$, appends a single bit 1 followed by the minimum number of bits 0 followed by a single bit 1 such that the length of the result is a multiple of the block length.

Multi-rate padding appends at least 2 bits and at most the number of bits in a block plus one.

3.1.2 Absorbing

The absorption algorithm takes a message M as input and returns a digest h . It consists of three steps: padding the input, iteratively computing a sequence of b -bit states $s_0 \dots s_m$ ($b = 1600$), extracting the digest from the final state s_m . In more details the three steps are:

1. Apply padding to M . Let the resulting blocks be B_1, \dots, B_m ; each of length r .
2. $s_0 = 0^b$ and for $i = 1, \dots, m$ iteratively compute $s_i \leftarrow f(s_{i-1} \oplus (B_i || 0^{b-r}))$, where $||$ is bitstring concatenation, \oplus is bitwise XOR and f is the permutation function.
3. Output $h = \lfloor s_m \rfloor_n$, where $\lfloor s_m \rfloor_n$ is the first n bits of s_m and it is guaranteed that $n < b$ as $b = 1600$ and $n \in \{224, 256, 384, 512\}$.

The Keccak-reference [5] description of the algorithm is in the box below, where the parameters for the sponge construction are the permutation f of width b , a padding rule “pad” and the bitrate $r < b$. The input is M and the output (h in step 3 of the pseudo-code above) is $\lfloor Z \rfloor_l$.

The individual blocks are named P_i in the box below, rather than Bi as above.

Algorithm 1 The sponge construction $\text{SPONGE}[f, \text{pad}, r]$

Require: $r < b$

Interface: $Z = \text{sponge}(M, \ell)$ with $M \in \mathbb{Z}_2^*$, integer $\ell > 0$ and $Z \in \mathbb{Z}_2^\ell$

$P = M \parallel \text{pad}[r](|M|)$

$s = 0^b$

for $i = 0$ to $|P|_r - 1$ **do**

$s = s \oplus (P_i \parallel 0^{b-r})$

$s = f(s)$

end for

$Z = \lfloor s \rfloor_r$

while $|Z|_{rr} < \ell$ **do**

$s = f(s)$

$Z = Z \parallel \lfloor s \rfloor_r$

end while

return $\lfloor Z \rfloor_\ell$

3.2 Padding and absorbing

MITB (r, c, n) f is designed to compute $\text{Hash}(r, c, n) f s_0(\text{key} \parallel M)$, where s_0 is the initial state, M is the message whose MAC is required and Hash is the KECCAK hash function:

$\text{Hash}(r, c, n) f s m = \text{Squeeze } n(\text{Absorb } f c s(\text{Split } r(\text{Pad } r m)))$

The function Hash is defined for arbitrary values of the parameters r, c, n, f and arbitrary states s and bitstrings m (representing messages). The auxiliary functions Pad , Split , Absorb and Squeeze are defined below. The notation $[b_1, \dots, b_n]$ denotes a bitstring consisting of the n bits b_1, \dots, b_n . In particular $[T]$ is the bitstring consisting of exactly one bit (T representing 0). The function $\text{Zeros } u$ maps a number to a bitstring consisting of u zeros, e.g. $\text{Zeros } 3 = [F, F, F]$. $\text{Zeros } u$ is another notation for 0^u .

3.2.1 Pad

$\text{Pad } r m$ pads bitstring m , using the KECCAK rules described in Section 3.1.1. It can be concisely defined by:

$\text{Pad } r m = m \parallel [T] \parallel \text{Zeros}((r - (|m| + 2) \text{ MOD } r) \text{ MOD } r) \parallel [T]$

Whilst it is clear that $\text{Pad } r m$ appends a bitstring $[T] \parallel \text{Zeros}(x) \parallel [T]$ to m , it may not be obvious that x should be $(r - (|m| + 2) \text{ MOD } r) \text{ MOD } r$. However, for all $r > 1$:

$(r - (|m| + 2) \text{ MOD } r) \text{ MOD } r =$

if $|m| \text{ MOD } r = r-1$ **then** $r-1$ **else** $r - (|m| \text{ MOD } r) - 2$

so the complicated formula x is correct.

3.2.2 Split

Split $r m$ splits m into a list of blocks of length r , except for the last one, which has length $|m| \bmod r$. This has already been described in Section 1. To define it formally the list-processing functions **Cons**, **Take** and **Drop** are used.

Cons $e l$ adds an element e to the front of list l , for example **Cons** 0 [1, 2, 3] = [0, 1, 2, 3]. **Take** $u l$ returns the first u elements of l , for example **Take** 3 [0, 1, 2, 3, 4, 5] = [0, 1, 2]. **Take** $u l$ is the same as the $[l]_u$, as used in Section 3.1.2. **Drop** $u l$ removes the first u elements of l , for example **Drop** 3 [0, 1, 2, 3, 4, 5] = [3, 4, 5]. If $u \leq |l|$ then $|\mathbf{Take} \ u \ l| = u$, $|\mathbf{Drop} \ u \ l| = |l| - u$ and $\mathbf{Take} \ u \ l \parallel \mathbf{Drop} \ u \ l = l$.

The function **Split** is then defined recursively by:

$$\mathbf{Split} \ r \ m = \text{if } (r = 0) \vee |m| \leq r \text{ then } [m] \text{ else } \mathbf{Cons} \ (\mathbf{Take} \ r \ m) \ (\mathbf{Split} \ r \ (\mathbf{Drop} \ r \ m))$$

It is straightforward to verify that, if $r > 0$, then the result of concatenating all the blocks in **Split** $r m$ is m , that all blocks in **Split** $r m$, except the last one, have size r and that the last block in **Split** $r m$ has size $|m| \bmod r$.

3.2.3 Absorb

Absorb $f \ c \ s \ bkl$ absorbs the blocks in a list of blocks bkl starting from a state s as described in Section 3.1.2. It is defined recursively on bkl by:

$$\begin{aligned} \mathbf{Absorb} \ f \ c \ s \ [] &= s \\ \mathbf{Absorb} \ f \ c \ s \ (\mathbf{Cons} \ bk \ bkl) &= \mathbf{Absorb} \ f \ c \ (f(s \oplus (bk \parallel \mathbf{Zeros} \ c))) \ bkl \end{aligned}$$

An equivalent alternative definition of **Absorb** uses the standard iteration combinator **Foldl** that is widely used in functional programming:

$$\mathbf{Absorb} \ f \ c = \mathbf{Foldl} \ (\lambda s \ bk. f(s \oplus (bk \parallel \mathbf{Zeros} \ c)))$$

where: $\mathbf{Foldl} \ fn \ e \ [] = e$ and $\mathbf{Foldl} \ fn \ e \ (\mathbf{Cons} \ x \ l) = \mathbf{Foldl} \ fn \ (fn \ e \ x) \ l$.

3.2.4 Squeeze

An unusual feature of the general KECCAK hash algorithm is that it can generate output digests of arbitrary length using an iterative ‘squeezing’ algorithm. However, for SHA-3 the recommended digest size n is 224. This is specified to be the bottom n bits in the final state after absorbing the message. Thus: **Squeeze** $n \ s = \mathbf{Take} \ n \ s$ or just **Squeeze** = **Take**.

3.3 MACs based on Keccak

The KECCAK designers claim that a secure message authentication code (MAC) can be computed by simply hashing the result of concatenating a key onto the front of a message. The KECCAK specification [5] says:

Unlike SHA-1 and SHA-2, Keccak does not have the length-extension weakness, hence does not need the HMAC nested construction. Instead, MAC computation can be performed by simply prepending the message with the key.

This KECCAK MAC of message M using key key is $\mathbf{Hash} \ (r, c, n) \ f \ (\mathbf{Zeros}(r+c)) \ (key \parallel M)$.

4 Formal specification of MITB

The specification of MITB requires that $\text{MITB}(r, c, n) f \models \phi$, where ϕ is a conjunction of properties. The particular properties in the conjunction are given mnemonic names and explained and defined separately in the boxes below. Some of these properties need to refer to the secret key *key* which is stored as $f(\text{key} \parallel \text{Zeros}(c))$ in MITB's permanent memory. The atomic property `pmemState` is used to state such properties: $\text{MITB}(r, c, n) f \models \text{pmemState}(s)$ is true if and only if MITB is storing *s* in its permanent memory.

4.1 Initialisation: Init

The property `Init` specifies that the 'power up' state of MITB is *ready*, that zeros are being output on `digest_out` and that $f(\text{key} \parallel \text{Zeros}(c))$ is stored in permanent memory.

```
Init key c n f = readyOut(T) And digestOut(Zeros n) And pmemState(f(key||Zeros(c)))
```

4.2 Freezing the state: Freeze

The property `Freeze` specifies that the state and outputs of MITB remains unchanged as long as T is input on `skip_inp`.

```
Freeze =  
Always  
(Forall s b1 b2.  
  skipInp(T) And pmemState(s) And readyOut(b1) And digestOut(b2)  
  Implies  
  Next(pmemState(s) And readyOut(b1) And digestOut(b2)))
```

4.3 Resetting: Reset

The property `Reset` specifies that inputting T on `movelnP` and F on `skipInP` in an *absorbing* state, i. e. when `ready_out` is F, results in a return to the *ready* state on the next cycle, with permanent memory unchanged and `Zeros n` being output at `ready_out`.

```
Reset n =  
Always  
(Forall s.  
  (movelnP(T) And skipInP(F) And readyOut(F) And pmemState(s))  
  Implies  
  Next(readyOut(T) And pmemState(s) And digestOut(Zeros n)))
```

4.4 Installing a new key: KeyUpdate

The property `KeyUpdate` entails that when in state *ready*, inputting `F` on both `skip_inp` and `move_inp` and inputting *key* (where $|key| = r$) on `block_inp`, results in $f(key||Zeros(c))$ being stored in the permanent memory and then remaining in the *ready* state on the next cycle.

```

KeyUpdate r c f =
  Always
  ((readyOut(T) And skipInp(F) And moveInp(F))
   Implies
   (Forall key.
    blockInp(key) And Bool(|key| = r)
    Implies Next(readyOut(T) And pmemState(f(key||Zeros(c))))))

```

The key installed by `KeyUpdate` overwrites the existing one in the permanent memory.

4.5 Computing a MAC: ComputeMAC

The definition of `ComputeMAC` below specifies that if the user follows the correct protocol for inputting a message then its MAC is computed.

The user is required to split the message into blocks and then input them and their lengths. An individual block *bk* is input by putting it on `block_inp` and its size $|bk|$ on `size_inp` whilst holding both `skip_inp` and `move_inp` at `F`. This is represented by the temporal formula `InputBlock bk` defined by:

```
InputBlock bk = blockInp(bk) And sizeInp |bk| And skipInp(F) And moveInp(F)
```

The recursively defined property `InputBlocks r [bk1, ..., bkm]` uses `InputBlock` and specifies the procedure for inputting the sequence of blocks bk_1, \dots, bk_m . It is a temporal formula which the user must ensure holds. The complexity in the definition is (i) to ensure that an extra empty block is added when the last block has size *r* and (ii) to drive the device for an extra cycle when the size of the last block is *r*−1.

```

InputBlocks r [] = Bool(F)      (This case should not arise in practice.)
InputBlocks r (Cons bk bkl) =
  if bkl = []
  then if |bk| = r
        then InputBlock r bk And Next(InputBlock r [])
        else if |bk| = r - 1
              then InputBlock r bk And Next(skipInp(F) And moveInp(F))
              else InputBlock r bk
  else InputBlock r bk And Next(InputBlocks r bkl) :

```

In the definition of `ComputeMAC` in the box below, the lines are numbered for use in the detailed explanation given after the definition.

```

1 ComputeMAC ( $r, c, n$ )  $f$  =
2 Always
3 (forall  $key\ m$ .
4   (readyOut(T) And skipInp(F) And moveInp(T)
5     And Bool( $|key| = r$ ) And pmemState( $f(key||Zeros(c))$ ))
6   Implies
7   Next
8     (InputBlocks  $r$  (Split  $r\ m$ )
9     Implies
10    (readyOut(F) And digestOut(Zeros  $n$ ))
11    UntilN(if  $|m| \text{ MOD } r = r-1$  then  $|m| \text{ DIV } r + 2$  else  $|m| \text{ DIV } r + 1$ )
12    (readyOut(T) And digestOut(MAC  $key\ (r, c, n)\ m$ ) And pmemState( $f(key||Zeros(c))$ ))))

```

The definition of `ComputeMAC` should be compared to the informal description on page 2. The following detailed description of the formula refers to the lines in the box above.

- Line 1:** The property `ComputeMAC` is parametrised on the KECCAK parameters (r, c, n) and a permutation function f (see Section 3).
- Line 2:** `Always` specifies that the property holds at all times; without this `ComputeMAC` would only be true at time 0.
- Line 3:** The `forall` quantification is inside the scope of the enclosing `Always`, therefore it binds the key key and message m at the time the transaction holds.
- Line 4:** A MAC computation can only be started when MITB is in state *ready*. To start the computation in this state the user inputs F on input `skip_inp` and T on input `move_inp`. This causes MITB to go into the *absorbing* state on the next cycle.
- Line 5:** The key key used for the computation is assumed to be of size r (i. e. 1152 bits for the SHA-3 recommended instance of KECCAK); $f(key||Zeros(c))$ is assumed to be stored in the permanent memory of the device (see Section 4.4).
- Lines 6-7:** The body of the `forall` is of the form ‘*precondition Implies Next absorb*’ which should be read ‘if *precondition* holds then on the next cycle *absorb* holds’, where *absorb* has the form ‘*input Implies invariant UntilN(number-of-steps) result*’. and occupies lines 8-12.
- Line 8:** The message is split into blocks and these are input on successive cycles.
- Line 10:** During the computation `ready_out` shows F and 0s are output on `digest_out`.
- Line 11:** The computation takes $|m| \text{ DIV } r + 2$ steps if $|m| \text{ MOD } r$ is $r-1$, otherwise it takes $|m| \text{ DIV } r + 1$ steps.
- Line 12:** The output `ready_out` changes to T and the MAC appears on `digest_out`. The value in the permanent memory is unchanged from the start of the computation.

4.6 All reachable state are secure: Secure

The property **Secure** entails that in all reachable states the value output at `digestOut` is either `Zeros n` or `Hash (r, c, n) f m` for some message m . See Section 6 for a discussion of why this may be too weak to be significant.

$$\begin{aligned} \text{Secure } (r, c, n) f = \\ \text{Always}(\text{digestOut}(\text{Zeros } n) \text{ Or Exists } m. \text{digestOut}(\text{Hash } (r, c, n) f (\text{Zeros}(r+c) m)) \end{aligned}$$

4.7 Complete specification

The specification of MITB is the conjunction of the properties defined in the preceding sections.

$$\begin{aligned} (\text{MITB } (r, c, n) f, s_0) \models & \text{Init } key \ c \ n \ f & \text{And} \\ & \text{Freeze} & \text{And} \\ & \text{Reset } n & \text{And} \\ & \text{KeyUpdate } r \ c \ f & \text{And} \\ & \text{ComputeMAC } (r, c, n) \ f & \text{And} \\ & \text{Secure } (r, c, n) \ f \end{aligned}$$

Whether or not this is either a correct or sufficient specification is discussed in Section 6.

5 Implementation

Recall from Section 2 that for any temporal formula ϕ :

$$\text{MITB_IMP } (r, c, n) f \models \phi \Leftrightarrow \forall \sigma_1 \sigma_2. \text{MITB_IMP } (r, c, n) f \sigma_1 \sigma_2 \Rightarrow \phi(\sigma_1, \sigma_2)$$

where `MITB_IMP` is a predicate defined in terms of a next-state function `MITB` representing a Moore machine and parametrised on the KECCAK parameters (r, c, n) and the KECCAK permutation function f .

This section describes a concrete definition of `MITB_IMP` that has been proved to implement the specification given in Section 4.7.³ The description is presented in two stages: first a curried function `MITB_FUN` specifies the behaviour abstractly, and then `MITB_FUN` is refined to `MITB`, which is then used to define `MITB_IMP` that models a high level register transfer level (RTL) implementation. A diagram of the states and transitions of `MITB_FUN` is on page 14.

5.1 Behavioural specification: MITB_FUN

`MITB_FUN` takes an abstract state, which is a triple $(\text{cnt1}, \text{pmem}, \text{vmem})$, and an input i (elaborated below) and returns the next state $(\text{cnt1}', \text{pmem}', \text{vmem}')$. The first component, `cnt1`, can have one of three values: `Ready`, `Absorbing` and `AbsorbEnd`. `Ready` corresponds to the *ready* state described in Section 1 and both `Absorbing` and `AbsorbEnd` correspond to the *absorbing* state. The second and third components of an abstract state, `pmem` and `vmem`, are bit-strings of length $r+c$ and represent the values of the permanent and volatile memory.

³See <http://www.cl.cam.ac.uk/~mjcg/MITB/> for details.

An input i can either be `Move`, `Skip` or `Input bk len` , where bk is a bitstring of size r and len is the number of bits of bk that constitutes the block being input (thus $len \leq r$). The bitstring bk and number len represent values being input on `block_inp` and `size_inp`.

The definition of `MITB_FUN`, and other functions that follow, have been cut-and-pasted from the files input to the HOL4 proof assistant⁴ used for the verification and then lightly edited to improve readability and to make them compatible with the notation used elsewhere in this document. The definition of `MITB_FUN` uses ML-style pattern matching, so there are separate equations for the various combinations of values of `cntl` and the input i . These equations are numbered for easy reference (the numbers are not in the original source).

```

1: (MITB_FUN (r,c,n) f (cntl,pmem,vmem) Skip =
    (cntl,pmem,vmem))
^
2: (MITB_FUN (r,c,n) f (Ready,pmem,vmem) (Input key len) =
    (Ready,f(key || Zeros c),Zeros(r+c)))
^
3: (MITB_FUN (r,c,n) f (Ready,pmem,vmem) Move =
    (Absorbing,pmem,pmem))
^
4: (MITB_FUN (r,c,n) f (Absorbing,pmem,vmem) Move =
    (Ready,pmem,Zeros(r+c)))
^
5: (MITB_FUN (r,c,n) f (Absorbing,pmem,vmem) (Input blk len) =
    if len ≤ r-2 then
      (Ready,pmem,
        f(vmem ⊕ (Take len blk || [T] || Zeros((r-len)-2) || [T] || Zeros c)))
    else if len = r-1 then
      (AbsorbEnd,pmem,f (vmem ⊕ (Take len blk || [T] || Zeros c)))
    else (Absorbing,pmem,f(vmem ⊕ (blk || Zeros c))))
^
6: (MITB_FUN (r,c,n) f (AbsorbEnd,pmem,vmem) Move =
    (Ready,pmem,Zeros(r+c)))
^
7: (MITB_FUN (r,c,n) f (AbsorbEnd,pmem,vmem) (Input blk len) =
    (Ready,pmem,f (vmem ⊕ (Zeros(r-1) || [T] || Zeros c))))

```

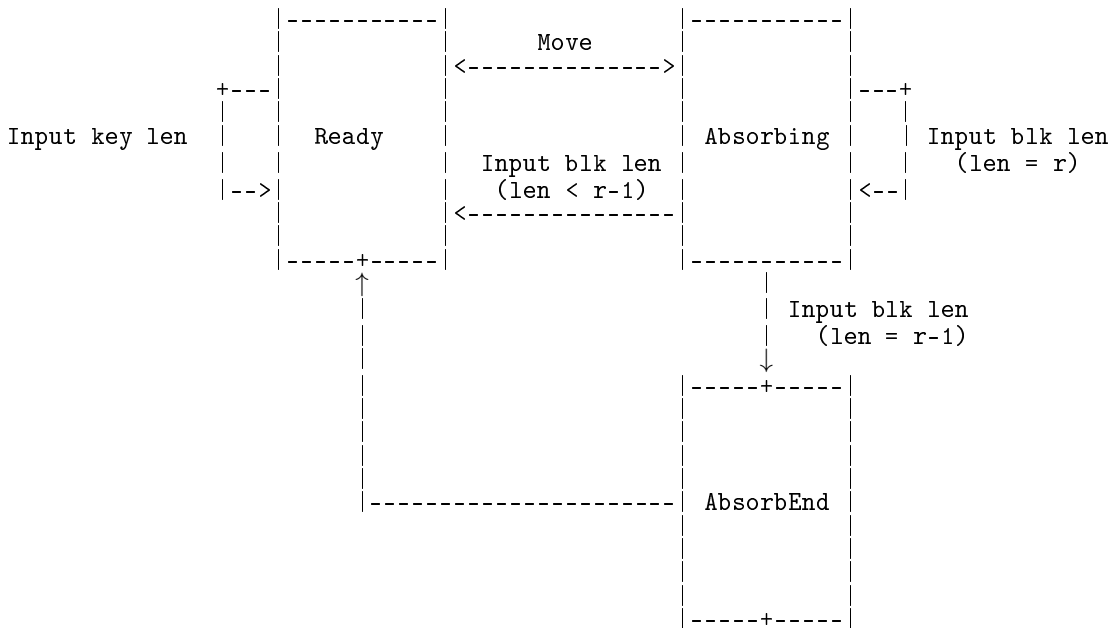
Equation 1 says that if `Skip` is input, then the state stays the same. Equations 2 and 3 describe what happens in the *ready* state (i.e. `cntl = Ready`). If the input is `Input key len` then the permanent memory `pmem` is set to $f(key || Zeros c)$, the volatile memory `vmem` is set to $Zeros(r+c)$, and the state remains *ready*. If the input is `Move` then the next state is *absorbing* (`cntl = Absorbing`) with the permanent memory unchanged and the volatile memory set to the value of the permanent memory. Equations 4 and 6 specify that if `Move` is input whilst *absorbing*, then the volatile memory is reset to zeros and the device returns to the *ready* state.

The most complex equation is 5, which specifies the state transition corresponding to absorbing a block. What happens depends on whether the block is the last one, which is signalled by the input length being less than r and corresponds to the first two branches of the conditional. The complexity here is because the device does the padding, as described in Section 3.1.1. If the last block is one bit short of being a full block of length r ($len = r-1$) then one bit

⁴<http://hol.sourceforge.net/>

is added and the device enters the sub-state of *absorbing* with `cntl = AbsorbEnd`, then on the next cycle, described in equation 7, the remaining padding (i.e. $r-1$ zeros and a final `T`) is added and the permutation f applied before transitioning back to the *ready* state. The final `else`-clause in equation 5 specifies the absorption of a non-final block, as described in Section 3.1.2. Such a block must have size exactly r and is absorbed by: (i) appending c zeros to it, (ii) then XOR-ing the result with the current value, `vmem`, of the volatile memory, then (iii) applying the KECCAK permutation f to the result of the XOR-ing, and finally (iv) the volatile memory is updated to the result of this application of f .

Here is an overview of MITB_FUN in the form of an ASCII art state transition diagram.



The function MITB is similar MITB_FUN except that it decodes the inputs into abstract commands `Skip`, `Move` and `Input blk len`.

```

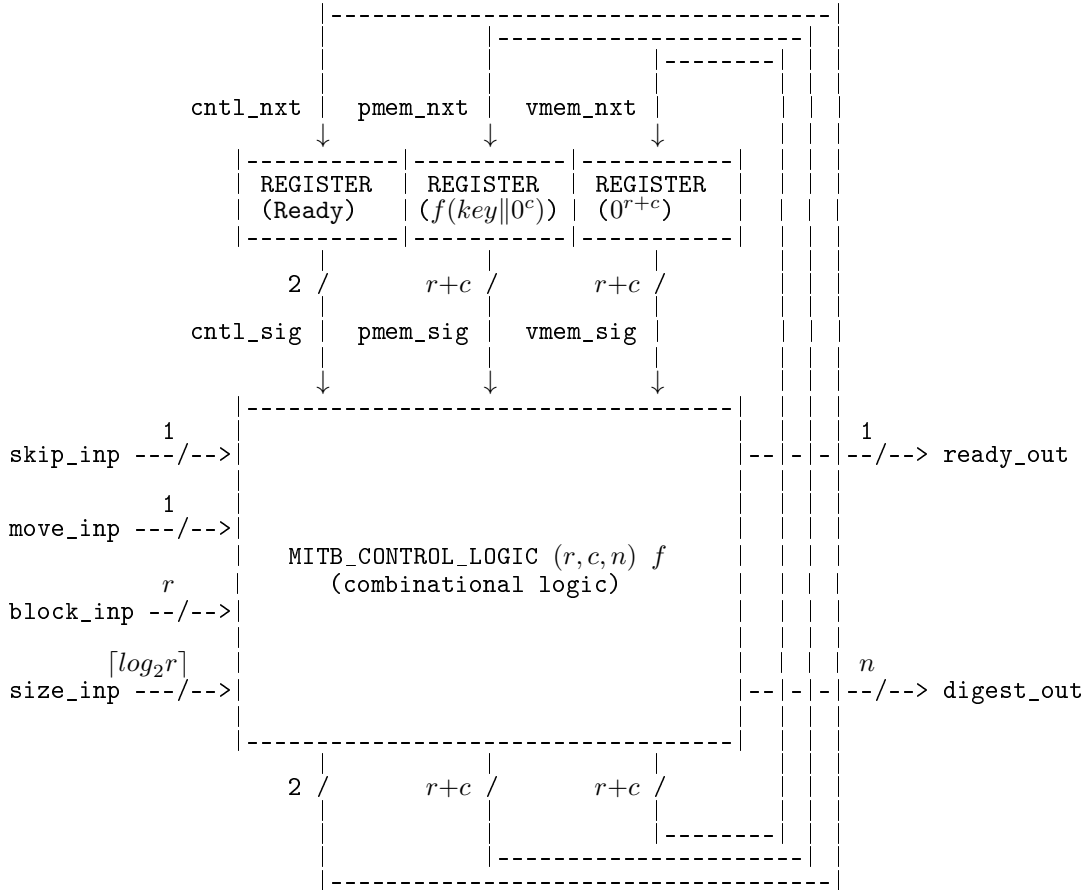
MITB (r,c,n) f ((skip,move,block,size),(cntl,pmem,vmem)) =
  MITB_FUN (r,c,n) f (cntl,pmem,vmem)
  (if skip = [T] then Skip
   else if move = [T] then Move
   else Input block size)
  
```

5.2 Register transfer behaviour and structure: MITB and MITB_DEV

To make a concrete device based on the behaviour specified in MITB, a hardware structure implementing the state transitions needs to be designed and verified. This will consist of registers for storing the state (`cntl`, `pmem`, `vmem`) – probably using different memory technologies for `pmem` and for `cntl` and `vmem` – and a specification of the control logic.

The diagram that follows on page 15 shows such a design, albeit one that is still quite abstract and missing many details (see Section 6 for some discussion). In this diagram the datapaths are labelled with the names used in the formal specification given later (e.g. `cntl_sig`, `cntl_nxt`)

and a number followed by “/” indicates the width of the labelled datapath in bits. Note, however, that in the model the values on `cntl_sig`, `cntl_nxt` and `size_inp` have not been coded as bitstring and, for simplicity, are kept abstract. The `cntl` component holds one of three values `Ready`, `Absorbing` or `AbsorbEnd`, so two bits are sufficient to encode these. The input `size_inp` is a number less than or equal to r , so can be encoded in $\lceil \log_2 r \rceil$ bits.



This diagram is expressed in logic in a standard way using relations.⁵ Melham’s book is a comprehensive reference [12]. The definition of MITB below uses the relation `REGISTER` to model registers as a unit delay:

```
REGISTER init_state (inp,out) ⇔
  (out 0 = init_state)
  ∧
  ∀t. out (t+1) = inp t
```

Three instances of `REGISTER` are used in the definition of MITB: to store `cntl`, `pmem` and `vmem`. Their initial values are indicated in brackets (using 0^c to abbreviate `Zeros c` and 0^{r+c} to abbreviate `Zeros(r + c)`).

The relation `MITB_CONTROL_LOGIC` packages up MITB as a relation and adds some logic to drive the outputs `ready_out` and `digest_out`.

⁵See <http://www.cl.cam.ac.uk/~mjcg/WhyHOL.pdf>.

```

MITB_CONTROL_LOGIC (r,c,n) f
  (cntl_sig,pmem_sig,vmem_sig,skip_inp,move_inp,block_inp,size_inp,
   cntl_nxt,pmem_nxt,vmem_nxt,ready_out,digest_out)
⇔
(∀t.
  (cntl_nxt t,pmem_nxt t,vmem_nxt t) =
  MITB (r,c,n) f
    ((skip_inp t,move_inp t,block_inp t,size_inp t),cntl_sig t,
     pmem_sig t,vmem_sig t))
^
(∀t. ready_out t = [cntl_sig t = Ready])
^
(∀t. digest_out t =
  if cntl_sig t = Ready then Take n (vmem_sig t) else Zeros n)

```

The diagram on page 15 shows widths of the various datapaths. In an HDL like Verilog these would be expressed in the type system, but here the predicate WIDTH is used.

```
Width sig n ⇔ ∀t. |sig t| = n
```

Also the verification requires the KECCAK parameters to satisfy $2 < r$, $0 < c$ and $n \leq r$, which is clearly satisfied by the particular values used by SHA-3. This constraint is enforced using the predicate GoodParameters.

```
GoodParameters (r,c,n) ⇔ 2 < r ∧ 0 < c ∧ n ≤ r
```

The implementation MITB_IMP combines the registers for holding the various state components with the combinational logic MITB_CONTROL_LOGIC.

```

MITB_IMP key (r,c,n) f
  (cntl_sig,pmem_sig,vmem_sig)
  (skip_inp,move_inp,block_inp,size_inp,ready_out,digest_out) ⇔
∃cntl_nxt pmem_nxt vmem_nxt.
  GoodParameters (r,c,n) ∧ (∀s. |f s| = |s| ∧
  (|key| = r) ∧ Width pmem_sig (r+c) ∧
  Width vmem_sig (r+c) ∧ Width pmem_nxt (r+c) ∧
  Width vmem_nxt (r+c) ∧ Width skip_inp 1 ∧ Width move_inp 1 ∧
  Width block_inp r ∧ (∀t. size_inp t ≤ r) ∧
  Width ready_out 1 ∧ Width digest_out n ∧
  REGISTER Ready (cntl_nxt,cntl_sig) ∧
  REGISTER (f(key||Zeros c)) (pmem_nxt,pmem_sig) ∧
  REGISTER (Zeros(r+c)) (vmem_nxt,vmem_sig) ∧
  MITB_CONTROL_LOGIC (r,c,n) f
    (cntl_sig,pmem_sig,vmem_sig,skip_inp,move_inp,block_inp,
     size_inp,cntl_nxt,pmem_nxt,vmem_nxt,ready_out,digest_out)

```

A non-vacuity theorem verifying that every sequence of inputs on skip_inp, move_inp, block_inp and size_inp determines sequences of outputs on ready_out and digest_out has been proved:

$$\begin{aligned}
&\vdash \forall \text{key } r \ c \ n \ f \ \text{skip_inp} \ \text{move_inp} \ \text{block_inp} \ \text{size_inp}. \\
&\quad (\forall s. (|f \ s| = |s|) \wedge \text{GoodParameters } (r,c,n) \wedge \\
&\quad (|\text{key}| = r) \wedge \text{Width } \text{skip_inp} \ 1 \wedge \text{Width } \text{move_inp} \ 1 \wedge \\
&\quad \text{Width } \text{block_inp} \ r \wedge (\forall t. \text{size_inp} \ t \leq r) \\
&\quad \Rightarrow \\
&\quad \exists \text{cntl_sig} \ \text{pmem_sig} \ \text{vmem_sig} \ \text{ready_out} \ \text{digest_out}. \\
&\quad \text{MITB_IMP } \text{key} \ (r,c,n) \ f \\
&\quad \quad (\text{cntl_sig}, \text{pmem_sig}, \text{vmem_sig}) \\
&\quad \quad (\text{skip_inp}, \text{move_inp}, \text{block_inp}, \text{size_inp}, \text{ready_out}, \text{digest_out})
\end{aligned}$$

6 Discussion and further work

The methods used here, which combine temporal logic specifications with mechanised proofs about state machines represented in higher order logic, are not new and date from the 1980s [10, 12, 7]. The contribution of this case study is to show how these old ideas may possibly be useful on a timely example. Although this study is both trivial and incomplete, it is hoped it may be a first step towards something significant.

In the rest of this section the incompleteness of the current work is described. Following that, there is a first partial attempt at a security assessment of the MITB project. Finally, possible future work is outlined.

6.1 Adequacy and incompleteness

The current specification has not been validated as a usable API description. All that has been done is to prove that an implementation of MITB as described in Section 5 implements this specification. The list of properties in Section 4 are ad hoc. Whilst they are consistent, since the MITB implementation is a model, how can one know if they are adequate? For example, the property `Secure` defined in Section 4.6 is weak because “`Exists m`” doesn’t say anything about what m might be and there might be a value in the range of the hash function that reveals something about the secret key. Rather than an existential quantification, perhaps the specification `Secure` should be more explicit, for example say that the digest output is either zero or is the hash of some previously input message. Such security properties are subtle and we are not sure what are the ones we should require of MITB. We are not even sure if the temporal logic notation used here is a good choice of property language; understanding exactly what the definition of `ComputeMAC` means is quite challenging.

The design of MITB is also ad hoc in that the effects of the various inputs may well be badly designed. For example, having the stored key reset whenever `move_inp` is `F` in the `ready` state may be dangerous (e.g. an attacker could surreptitiously reset the key). Adjusting the API would be easy (e.g. just removing the ability to reset the key).

There are two major omissions in the work so far. The first major omission is that the MITB design is not realistic hardware. The input `block_inp` is 1600-bits wide and the input `size_inp` is a number. Whilst the latter is trivial to fix by encoding the numbers as bitstrings, the former will require sequential buffering, say to accumulate blocks 16-bits at a time over a 100 cycles. Adapting the state machine to do this should be straightforward, and the temporal logic specifications should be easy to adapt to work with the additional data acquisition cycles,

and most of the existing proofs should be reusable. The work involved is classical data and temporal abstraction [12]. However, this is all work which has not been done.

The second major omission is that the multi-round KECCAK- f permutation function has not been implemented. It is treated as an uninterpreted function. This nicely separates the API aspects of MITB from the cryptographic computation concerns, but a complete implementation would need hardware implementing f , which might need several cycles (e.g. one for each sub-round), so also requiring temporal refinement as discussed in the previous paragraph. There is a discussion in Chapter 4 (Hardware) of the KECCAK implementation overview [6]. Although the KECCAK- f is quite complicated [5] creating a verified hardware designs implementing it should be straightforward, though possibly a lot of work due to all the intricate details. An approach using verifying synthesis might be appropriate.

6.2 Security assessment

Imagine the MITB design and verification have been completed, say down to synthesisable RTL represented in logic, but resembling a standard HDL like Verilog. What would have been achieved? All that would have been shown is that the API functionality specified in LTL is realised by the HDL model. However:

1. Maybe the verifier is lying about having completed the proof?
 - This could be mitigated by replaying the proof.
2. Maybe the proof tool (HOL4) is unsound?
 - This could be mitigated by using another independent tool (e.g. HOL Light, ProofPower, Isabelle/HOL) together with an expert audit of the tools used, which should have a strong soundness pedigree.
3. How can the design model be securely manufactured?
 - There are many challenges here ranging from unsound synthesis tools, to unsafe implementation technologies [9].
4. What threats does the model ignore?
 - It is assumed KECCAK has good cryptographic properties. Being an NIST approved standard, and considering previous allegations of the NSA subverting NIST standards[13, 14], this may raise ‘Snowden worries’. There is no modelling of side-channels or tampering attacks which could extract $f(key||0^c)$ from the permanent memory.

Cohn [8] give an early discussion of issues related to some of those above that provoked a controversy [11].

Finally, although MITB was conceived as part of a secure password secrecy system, it has not been discussed how such a complete system would work, so there is no analysis of what actual contribution to security an MITB device might make.

6.3 Next steps: short and long term

Completing the design and implementation would seem to be the essential first next step, but maybe it would be better to step back now and decide where the MITB project is going and what is most critical from a total system security perspective.

Careful thinking is needed to evaluate what contribution a formal verification of MITB could make to enhancing the security of using hashing in a real world setting.

References

- [1] The MITB project. <http://www.cl.cam.ac.uk/~mjcg/MITB/Autumn2013Summary.html>.
- [2] Graham Steel's talk on Tookan at FMATS1, 2011. <http://www.cl.cam.ac.uk/~mjcg/Meeting.SecurityTools.2011/papers/Steel-fmats-tookan.pdf>.
- [3] First workshop on Formal Methods And Tools for Security (FMATS1), Cambridge, 2011. <http://www.cl.cam.ac.uk/~mjcg/Meeting.SecurityTools.2011/FMATS1.html>.
- [4] R. Anderson, M. Bond, J. Clulow, and S. Skorobogatov. Cryptographic processors – a survey. Technical Report UCAM-CL-TR-641, University of Cambridge, Computer Laboratory, Aug. 2005. <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-641.pdf>.
- [5] G. Bertoni, J. Daemen, M. Peeters, and G. V. Assche. Online keccak specifications, 2009. <http://keccak.noekeon.org/>.
- [6] G. Bertoni, J. Daemen, M. Peeters, G. V. Assche, and R. V. Kleer. Keccak implementation overview, 2012. <http://keccak.noekeon.org/Keccak-implementation-3.2.pdf>.
- [7] C.-T. Chou. Predicates, temporal logic, and simulations. In J. Joyce and C.-J. Seger, editors, *Higher Order Logic Theorem Proving and Its Applications*, volume 780 of *Lecture Notes in Computer Science*, pages 310–323. Springer Berlin Heidelberg, 1994.
- [8] A. Cohn. The notion of proof in hardware verification. *J. Autom. Reasoning*, 5(2):127–139, 1989. <http://www.cl.cam.ac.uk/~mjcg/papers/AvraProofPaper.pdf>.
- [9] S. Drimer. Security for volatile FPGAs. Technical Report UCAM-CL-TR-763, University of Cambridge, Computer Laboratory, Nov. 2009. <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-763.pdf>.
- [10] J. J. Joyce. Formal specification and verification of asynchronous processes in higher-order logic. Technical Report UCAM-CL-TR-136, University of Cambridge, Computer Laboratory, June 1988.
- [11] D. MacKenzie. The fangs of the VIPER. *Nature*, 352:467–468, 1991.
- [12] T. Melham. *Higher Order Logic and Hardware Verification*, volume 31 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1993.
- [13] Perlroth, Nicole. Government Announces Steps to Restore Confidence on Encryption Standards. “internal memos leaked by a former N.S.A. contractor, Edward Snowden, suggest that the N.S.A. generated one of the random number generators used in a 2006 N.I.S.T. standard — called the Dual EC DRBG standard — which contains a back door for the N.S.A.”.
- [14] D. Shumow and N. Ferguson. On the Possibility of a Back Door in the NIST SP800-90 Dual Ec Prng .
- [15] Solar Designer. Password hashing at scale. Slides for a talk at YaC 2012 (see Slide 9). <http://www.openwall.com/presentations/YaC2012-Password-Hashing-At-Scale/>.