

Bitwise primitives proposal for Plutus

Koz Ross

February 18, 2022

1 Motivation

Bitwise operations are one of the most fundamental building blocks of algorithms and data structures. They can be used for a wide variety of applications, ranging from representing and manipulating sets of integers efficiently, to implementations of cryptographic primitives, to fast searches. Their wide availability, law-abiding behaviour and efficiency are the key reasons why they are widely used, and widely depended on.

At present, Plutus lacks meaningful support for bitwise operations, which significantly limits what can be usefully done on-chain. While it is possible to mimic some of these capabilities with what currently exists, and it is always possible to introduce new primitives for any task, this is extremely unsustainable, and often leads to significant inefficiencies and duplication of effort.

We describe a list of bitwise operations, as well as their intended semantics, designed to address this problem.

2 Goals

To ensure a focused and meaningful proposal, we specify our goals below.

2.1 Useful primitives

The primitives provided should enable implementations of algorithms and data structures that are currently impossible or impractical. Furthermore, the primitives provided should have a high power-to-weight ratio: having them should enable as much as possible to be implemented.

2.2 Maintaining as many algebraic laws as possible

Bitwise operations, via Boolean algebras, have a long and storied history of algebraic laws, dating back to important results by the like of de Morgan, Post and many others. These algebraic laws are useful for a range of reasons: they guide implementations, enable easier testing (especially property testing) and in some cases much more efficient implementations. To some extent, they also formalize our intuition about how these operations ‘should work’. Thus, maintaining as many of these laws in our implementation, and being clear about them, is important.

2.3 Allowing efficient, portable implementations

Providing primitives alone is not enough: they should also be efficient. This is not least of all because many would associate ‘primitive operation’ with a notion of being ‘close to the machine’, and therefore fast. Thus, it is on us to ensure that the implementations of the primitives we provide have to be implementable in an efficient way, across a range of hardware.

3 Non-goals

We also specify some specific non-goals of this proposal.

3.1 No partial operations

Partiality is an undesirable property in any language and any functionality. This is especially true of bitwise operations, which are fairly low-level and need highly-predictable semantics to allow implementers to use them with confidence and know exactly what answers to expect. Furthermore, especially off-chain, dealing with partiality is annoying and difficult to track down issues with: MLabs in particular have lost endless hours with division-by-zero bugs.

3.2 No metaphor-mixing between numbers and bits

A widespread legacy of C is the mixing of treatment of numbers and blobs of bits: specifically, the allowing of logical operations on representations of numbers. This applies to Haskell as much as any other language: according to the Haskell Report, it is in fact *required* that any type implementing `Bits` implement `Num` first. While GHC Haskell only mandates `Eq`, it still defines `Bits` instances for types clearly meant to represent numbers. This is a bad choice, as it creates complex situations and partiality in several cases, for arguably no real gain other than C-like bit twiddling code.

Even if two types share a representation, their type distinctness is meant to be a semantic or abstraction boundary: just because a number is represented as a blob of bits does not necessarily mean that arbitrary bit manipulations are sensible. However, by defining such a capability, we create several semantic problems:

- Some operations end up needing multiple definitions to take this into account. A good example are shifts: instead of simply having left or right shifts, we now have to distinguish *arithmetic* versus *logical* shifts, simply to take into account that a shift can be used on something which is meant to be a number, which could be signed. This creates unnecessary complexity and duplication of operations.
- As Plutus `BuiltinIntegers` are of arbitrary precision, certain bitwise operations are not well-defined on them. A good example is bitwise complement: the bitwise complement of 0 cannot be defined sensibly, and in fact, is partial in its `Bits` instance.
- Certain bitwise operations on `BuiltinInteger` would have quite undesirable semantic changes in order to be implementable. A good example are bitwise rotations: we should be able to ‘decompose’ a rotation left or right by n into two rotations (by m_1 and m_2 such that $m_1 + m_2 = n$) without changing the outcome. However, because trailing zeroes are not tracked by the implementation, this can fail depending on the choice of decomposition, which seems needlessly annoying for no good reason.
- Certain bitwise operations on `BuiltinInteger` would require additional arguments and padding to define them sensibly. Consider bitwise logical AND: in order to perform this sensibly on `BuiltinIntegers` we would need to specify what ‘length’ we assume they have, and some policy of ‘padding’ when the length requested is longer than one, or both, arguments. This feels unnecessary, and isn’t even clear exactly how this should be done: for example, how would negative numbers be padded?

These complexities, and many more besides, are poor choices, owing more to the legacy of C than any real useful functionality. Furthermore, they feel like a casual and senseless undermining of type safety and its guarantees for very small and questionable gains. Therefore, defining bitwise operations on `BuiltinInteger` is not something we wish to support.

There are legitimate cases where a conversion from `BuiltinInteger` to `BuiltinByteString` is desirable; this conversion should be provided, and be both explicit and specified in a way that is independent of the machine or the implementation of `BuiltinInteger`, as well as total and round-tripping. Arguably, it is also desirable to provide built-in support for `BuiltinByteString` literals specified in a way convenient to their treatment as blobs of bytes (for example, hexadecimal or binary notation), but this is outside the scope of this proposal.

4 Proposed operations

We propose several classes of operations. Firstly, we propose two operations for inter-conversion between `BuiltinByteString` and `BuiltinInteger`, whose semantics are specified in Subsection 5.2:

`integerToByteString :: BuiltinInteger -> BuiltinByteString`: Convert a number to a bitwise representation.

`byteStringToInteger :: BuiltinByteString -> BuiltinInteger`: Reinterpret a bitwise representation as a number.

We also propose several logical operations on `BuiltinByteStrings`, whose semantics are specified in Subsection 5.3:

`andByteString :: BuiltinByteString -> BuiltinByteString -> BuiltinByteString`: Perform a bitwise logical AND on the arguments, producing a result whose size is the maximum of the sizes of its arguments.

`iorByteString :: BuiltinByteString -> BuiltinByteString -> BuiltinByteString`: Perform a bitwise logical IOR on the arguments, producing a result whose size is the maximum of the sizes of its arguments.

`xorByteString :: BuiltinByteString -> BuiltinByteString -> BuiltinByteString`: Perform a bitwise logical XOR on the arguments, producing a result whose size is the maximum of the sizes of its arguments.

`complementByteString :: BuiltinByteString -> BuiltinByteString`: Complement all the bits in the argument, producing a result whose size is the same as the argument.

Lastly, we define the following additional operations, whose semantics are specified in Subsection 5.4:

`shiftByteString :: BuiltinByteString -> BuiltinInteger -> BuiltinByteString`: Performs a bitwise shift of the first argument by the absolute value of the second argument, with padding, the direction being indicated by the sign of the second argument.

`rotateByteString :: BuiltinByteString -> BuiltinInteger -> BuiltinByteString`: Performs a bitwise rotation of the first argument by the absolute value of the second argument, the direction being indicated by the sign of the second argument.

`popCountByteString :: BuiltinByteString -> BuiltinInteger`: Returns the number of 1 bits in the argument.

`isBitSetByteString :: BuiltinByteString -> BuiltinInteger -> BuiltinBool`: If the bit at the position given by the second argument in the first argument is 1, return `True`, and `False` otherwise.

`isBitClearByteString :: BuiltinByteString -> BuiltinInteger -> BuiltinBool`: If the bit at the position given by the second argument in the first argument is 0, return `True`, and `False` otherwise.

4.1 Why these operations?

As stated in Subsection 3.2, there needs to be a well-defined interface between the ‘world’ of `BuiltinInteger` and `BuiltinByteString`. To provide this, we require `integerToByteString` and `byteStringToInteger`, which is designed to roundtrip (that is, describe an isomorphism). Furthermore, by spelling out a precise description of the conversions in Subsection 5.2, we make this predictable and portable.

Our choice of logical AND, IOR, XOR and complement as the primary logical operations is driven by a mixture of prior art, utility and convenience. These are the typical bitwise logical operations provided in hardware, and in most programming languages, leading to most algorithm descriptions relying on bitwise implementations assuming them as primitive. While we could reduce this number (and, in fact, due to Post, we know that there exist two *sole* sufficient operators), this would be both inconvenient and inefficient. As an example, consider implementing XOR using AND, IOR and complement: this would translate $x \text{ XOR } y$ into

$$(\text{COMPLEMENT } x \text{ AND } y) \text{ IOR } (x \text{ AND } \text{COMPLEMENT } y)$$

This is both needlessly complex and also inefficient, as it requires copying the arguments twice, only to throw away both copies.

Like our ‘baseline’ bitwise operations above, shifts and rotations are widely used, and considered as primitive. While `popCountByteString` could theoretically be simulated using `isBitSetByteString` and a fold, this is quite inefficient: the most efficient method for this operation would involve using something like the Harley-Seal algorithm, which requires a large lookup table. Furthermore, population counting is so commonly needed that at least one major hardware platform (x86) provides it as a hardware primitive.

Having both `isBitSetByteString` and `isBitClearByteString` is arguably somewhat redundant. However, it allows us the benefit of totality without confounding answers: you can never get `True` from either of them on an invalid index, and for any valid index, only one of them will return `True`.

Primitive	Linear in
<code>integerToByteString</code>	Argument (only one)
<code>byteStringToInteger</code>	Argument (only one)
<code>andByteString</code>	Longer argument
<code>iorByteString</code>	Longer argument
<code>xorByteString</code>	Longer argument
<code>complementByteString</code>	Argument (only one)
<code>shiftByteString</code>	<code>BuiltinByteString</code> argument
<code>rotateByteString</code>	<code>BuiltinByteString</code> argument
<code>popCountByteString</code>	Argument (only one)
<code>isBitSetByteString</code>	<code>BuiltinByteString</code> argument
<code>isBitClearByteString</code>	<code>BuiltinByteString</code> argument

Table 1: Primitives and which argument they are linear in

4.2 Costing

All of the primitives we describe are linear in one of their arguments. For a more precise description, see Table 1.

5 Semantics

5.1 Preliminaries

We define $\mathbb{N}^+ = \{x \in \mathbb{N} \mid x \neq 0\}$. We assume that `BuiltinInteger` is a faithful representation of \mathbb{Z} . A *bit sequence* $s = s_n s_{n-1} \dots s_0$ is a sequence such that for all $i \in \{0, 1, \dots, n\}$, $s_i \in \{0, 1\}$. A bit sequence $s = s_n s_{n-1} \dots s_0$ is a *byte sequence* if $n = 8k - 1$ for some $k \in \mathbb{N}$. We denote the *empty bit sequence* (and, indeed, byte sequence as well) by \emptyset .

We intend that `BuiltinByteStrings` represent byte sequences, with the sequence of bits being exactly as the description above. For example, given the byte sequence `0110111100001100`, the `BuiltinByteString` corresponding to it would be `"o\lf"`.

Let $i \in \mathbb{N}^+$. We define the sequence `binary`(i) = $(d_0, m_0), (d_1, m_1), \dots$ as

- $m_0 = i \bmod 2$, $d_0 = \frac{i}{2}$ if i is even, and $\frac{i-1}{2}$ if it is odd.
- $m_j = d_{j-1} \bmod 2$, $d_j = \frac{d_{j-1}}{2}$ if d_{j-1} is even, and $\frac{d_{j-1}-1}{2}$ if it is odd.

5.2 Representation of `BuiltinInteger` as `BuiltinByteString` and conversions

We describe the translation of `BuiltinInteger` into `BuiltinByteString` which is implemented as the `integerToByteString` primitive. Informally, we represent `BuiltinIntegers` with the least significant bit at bit position 0, using a twos-complement representation. More precisely, let $i \in \mathbb{N}^+$. We represent i as the bit sequence $s = s_n s_{n-1} \dots s_0$, such that:

- $\sum_{j \in \{0, 1, \dots, n\}} s_j \cdot 2^j = i$; and
- $s_n = 0$.
- Let `binary`(j) = $(d_0, m_0), (d_1, m_1), \dots$. For any $j \in \{0, 1, \dots, n-1\}$, $s_j = m_j$; and
- $n + 1 = 8k$ for the smallest $k \in \mathbb{N}^+$ consistent with the previous requirements.

For 0, we represent it as the sequence `00000000` (one zero byte). We represent any $i \in \{x \in \mathbb{Z} \mid x < 0\}$ as the twos-complement of the representation of its additive inverse. We observe that any such sequence is by definition a byte sequence.

To interpret a byte sequence $s = s_n s_{n-1} \dots s_0$ as a `BuiltinInteger`, we use the following process:

- If s is `00000000`, then the result is 0.

2. Otherwise, if $s_n = 1$, let s' be the twos-complement of s . Then the result is the additive inverse of the result of interpreting s' .
3. Otherwise, the result is $\sum_{i \in \{0,1,\dots,n\}} s_i \cdot 2^i$.

The above interpretation is implemented as the `byteStringToInteger` primitive. We observe that `byteStringToInteger` and `integerToByteString` form an isomorphism. More specifically:

```
byteStringToInteger . integerToByteString =
integerToByteString . byteStringToInteger =
id
```

5.3 Bitwise logical operations on `BuiltinByteString`

Throughout, let $s = s_n s_{n-1} \dots s_0$ and $t = t_m t_{m-1} \dots t_0$ be two byte sequences.

We describe the semantics of `andByteString`. For inputs s and t , the result is the byte sequence $u = u_{\max\{n,m\}} u_{\max\{n,m\}-1} \dots u_0$ such that for all $i \in \{0, 1, \dots, \max\{n, m\}\}$, we have

$$u_i = \begin{cases} 1 & i \leq \min\{n, m\} \text{ and } s_i = t_i = 1 \\ t_i & n < i \leq m \\ s_i & m < i \leq n \\ 0 & \text{otherwise} \end{cases}$$

For `iorByteString`, for inputs s and t , the result is the byte sequence $u = u_{\max\{n,m\}} u_{\max\{n,m\}-1} \dots u_0$ such that for all $i \in \{0, 1, \dots, \max\{n, m\}\}$, we have

$$u_i = \begin{cases} 1 & i \leq \min\{n, m\} \text{ and } s_i = 1 \\ 1 & i \leq \min\{n, m\} \text{ and } t_i = 1 \\ t_i & n < i \leq m \\ s_i & m < i \leq n \\ 0 & \text{otherwise} \end{cases}$$

For `xorByteString`, for inputs s and t , the result is the byte sequence $u = u_{\max\{n,m\}} u_{\max\{n,m\}-1} \dots u_0$ such that for all $i \in \{0, 1, \dots, \max\{n, m\}\}$, we have

$$u_i = \begin{cases} 1 & i \leq \min\{n, m\} \text{ and } s_i \neq t_i \\ t_i & n < i \leq m \\ s_i & m < i \leq n \\ 0 & \text{otherwise} \end{cases}$$

We observe that each of `andByteString`, `iorByteString` and `xorByteString` describes a commutative and associative operation. Furthermore, for each of these operations, (at least) \emptyset is an identity element.

We now describe the semantics of `complementByteString`. For input s , the result is the byte sequence $u = u_n u_{n-1} \dots u_0$ such that for all $i \in \{0, 1, \dots, n\}$ we have

$$u_i = \begin{cases} 1 & s_i = 0 \\ 0 & \text{otherwise} \end{cases}$$

We observe that `complementByteString` is self-inverting. We also note the following equivalences hold; these are the DeMorgan laws:

```
complementByteString (andByteString b b') =
iorByteString (complementByteString b) (complementByteString b')
```

```
complementByteString (iorByteString b b') =
andByteString (complementByteString b) (complementByteString b')
```

5.4 Mixed operations

Throughout this section, let $s = s_n s_{n-1} \dots s_0$ and $t = t_m t_{m-1} \dots t_0$ be byte sequences, and let $i \in \mathbb{Z}$.

We describe the semantics of `shiftByteString`. Informally, these are logical shifts, with negative shifts ‘moving’ away from bit index 0, and positive shifts ‘moving’ towards bit index 0. More precisely, given the argument s and i , the result of `shiftByteString` is the byte sequence $u_n u_{n-1} \dots u_0$, such that for all $j \in \{0, 1, \dots, n\}$, we have

$$u_j = \begin{cases} s_{j+i} & j - i \in \{0, 1, \dots, n\} \\ 0 & \text{otherwise} \end{cases}$$

We observe that for k, ℓ with the same sign and any `bs`, we have

```
shiftByteString (shiftByteString bs k) l = shiftByteString bs (k + l)
```

We now describe `rotateByteString`, assuming the same inputs as the description of `shiftByteString` above. Informally, the ‘direction’ of the rotations matches that of `shiftByteString` above. More precisely, the result of `rotateByteString` on the given inputs is the byte sequence $u_n u_{n-1} \dots u_0$ such that for all $j \in \{0, 1, \dots, n\}$, we have $u_j = s_{j+i \bmod (n+1)}$. We observe that for any k, ℓ , and any `bs`, we have

```
rotateByteString (rotateByteString bs k) l = rotateByteString bs (k + l)
```

We also note that

```
rotateByteString bs 0 = shiftByteString bs 0 = bs
```

For `popCountByteString` with argument s , the result is

$$\sum_{j \in \{0, 1, \dots, n\}} s_j$$

Informally, this is just the total count of 1 bits. We observe that for any `bs` and `bs'`, we have

```
popCountByteString bs + popCountByteString bs' =
popCountByteString (appendByteString bs bs')
```

For `isBitSetByteString` on s and i , we return `True` if $0 \leq i \leq n$ and $s_i = 1$, and `False` otherwise. For `isBitClearByteString` on the same arguments, we return `True` if $0 \leq i \leq n$ and $s_i = 0$, and `False` otherwise.