

BITSAD: A Domain-Specific Language for Bitstream Computing

Your Brain is a Unary Computer '19

Kyle Daruwalla, Heng Zhuo, and Mikko Lipasti

22 June 2019

University of Wisconsin - Madison, Dept. of Elec. and Comp. Eng.



Motivation





Figure 1: The resource (LUTs + FFs) consumption of Bitstream Computing (BC) implementations are much lower than floating point (FP) and fixed point (FXP) designs.

1. Why do we need BITSAD
2. What is BITSAD
3. What does it bring
4. What is coming next

Motivation and Background

Limitations of Current Approaches

RoboBee has ultra-low resources constraints:
(< 35 mW compute power)¹



Traditional computing paradigms cannot meet these constraints!

- Prior system connected to desktop computer
- Current on-device chip only supports basic stationary flight³

¹Duhamel et al. 2011

²Ma 2015

³Zhang et al. 2017

What is Bitstream Computing?

Stochastic Bitstreams:



$$\begin{array}{l} \mathbb{E}[S_1] = 0.5 \quad \underline{0, 1, 0, 1, 1, 1, 0, 0} \\ \mathbb{E}[S_2] = 0.75 \quad \underline{1, 1, 0, 1, 1, 0, 1, 1} \end{array} \bigg\} \underline{0, 1, 0, 1, 1, 0, 0, 0} \quad \mathbb{E}[S_1 S_2] = \mathbb{E}[S_1] \mathbb{E}[S_2] = 0.375$$
$$\begin{array}{l} \mathbb{E}[S_1] = 0.5 \quad \underline{0, 1, 0, 1, 1, 1, 0, 0} \\ \mathbb{E}[S_2] = 0.25 \quad \underline{1, 0, 0, 0, 1, 0, 0, 0} \end{array} \bigg\} \underline{1, 1, 0, 1, 1, 1, 0, 0} \quad \mathbb{E}[S_1 + S_2] = \mathbb{E}[S_1] + \mathbb{E}[S_2] = 0.75$$

Deterministic Bitstreams:



Density of "1" \Rightarrow Higher
amplitude


- Sequence is *deterministic*
- Oversampled audio data
- Leads to efficient filters

Traditional design flow turnover is slow:



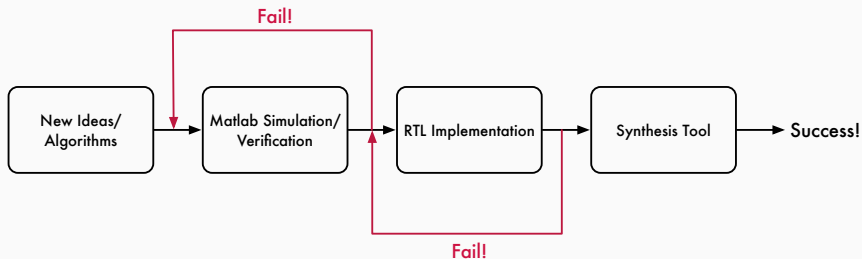
Design flow

Traditional design flow turnover is slow:



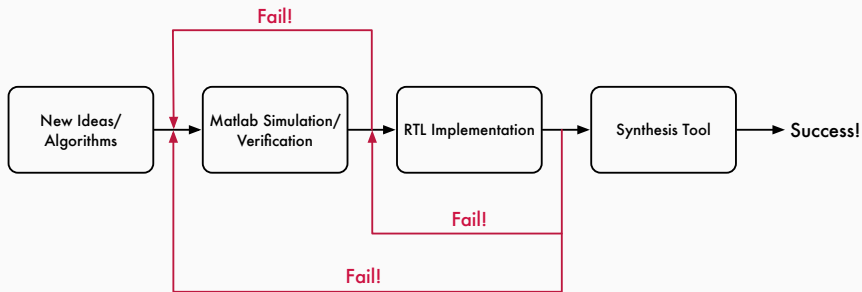
Design flow

Traditional design flow turnover is slow:




Design flow

Traditional design flow turnover is slow:



Design flow

Traditional design flow turnover is slow:



Where is the problem coming from:



What can we do about it:



BITSAD Bitstream Synthesizer and Designer

Based on Scala, which:

1. is a general-purpose, functional programming language
2. runs on Java virtual machine
3. uses IDE or command line build tool (`sbt`)
4. allows plugin in arbitrary compiler phases



Listing 1: Example code on `SBitstreams`.

```
1var a = SBitstream(0.5)
2var b = SBitstream(0.25)
3var c = SBitstream(-0.5)
4
5var d = a + b
6var e = a * c // handles sign
```

Listing 2: Example code on `DBitstreams`.

```
1var buff = DelayBuffer(32) // delay buffer of length 32
2var sdm = SDM() // Sigma-Delta modulator FXP -> DBitstream
3
4var y = 2 * buff.pop + x.pop // x is a pre-loaded DBitstream
5var z = sdm.evaluate(y)
```

Basic Operators

Operator	Description	SBitstream	DBitstream
+	Addition	Y	Y
-	Subtraction	Y	Y
*	Multiplication	Y	Y
/	Division	Y	N
:/	Fixed-Gain Div.	Y	N

Table 1: Operators defined on the `SBitstream` and `DBitstream` data type.

Listing 3: Operator Examples with mixed types.

```
1var x = SBitstream(0.5)
2var y = x - 0.1
3var z = y * 0.2 + 1
```

Similar to Matlab code:

Listing 4: Creating a `Matrix[A]`.


```
1 var a = Matrix(Array(  
2   Array(0.1, 0.2),  
3   Array(0.3, 0.4)  
4 )
```

Listing 5: Working with `Matrix[A]`.

```
1 var a = rand[SBitstream](2, 2) // generate some random matrices  
2 var b = rand[SBitstream](2, 1) // could be any Numeric type  
3 var c = a * b                 // c is a 2x1 vector  
4 var d = norm(0.25 * c)        // takes L2-norm of vector c
```


Listing 6: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 7: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 8: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w :/ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z :/ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 9: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 10: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 11: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```




Listing 12: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```



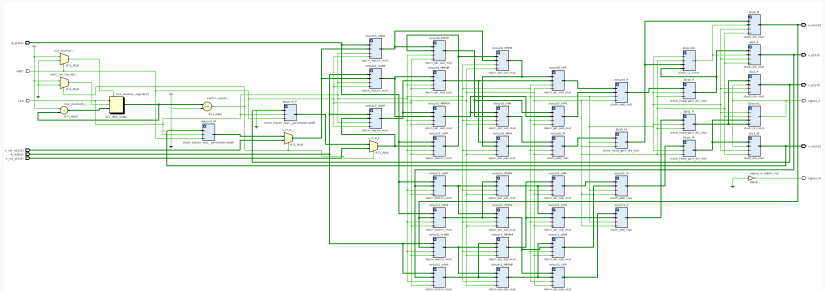
Listing 13: Example BITSAD program.

```
1 case class Module (params: Parameters) {
2
3   // Define outputs
4   val outputList = List(("v", params.n, 1),
5                         ("u", params.m, 1),
6                         ("sigma", 1, 1))
7
8   def loop(A: Matrix[SBitstream], v: Matrix[SBitstream]):
9     (Matrix[SBitstream], Matrix[SBitstream], SBitstream) = {
10    // Update right singular vector
11    var w = A * v
12    var wScaled = w ./ math.sqrt(params.m)
13    var u = wScaled / Matrix.norm(wScaled)
14
15    // Update left singular vector
16    var z = A.T * u
17    var zScaled = z ./ math.sqrt(params.n)
18    var sigma = Matrix.norm(zScaled)
19    var _v = zScaled / sigma
20
21    (u, _v, sigma)
22  }
23
24 } //Singular value decomposition
```



But, wait, why not just write Verilog then?

But, wait, why not just write Verilog then?




Warning

It is a bad idea!

Subtleties of Bitstream Computing

Example: Moving Average Filter

Consider the following equivalent expressions for a moving average filter of length 4:



Example: Moving Average Filter


Consider the following equivalent expressions for a moving average filter of length 4:

Factored:


$$0.25 * (d1 + d2 + d3 + x)$$

Distributed:

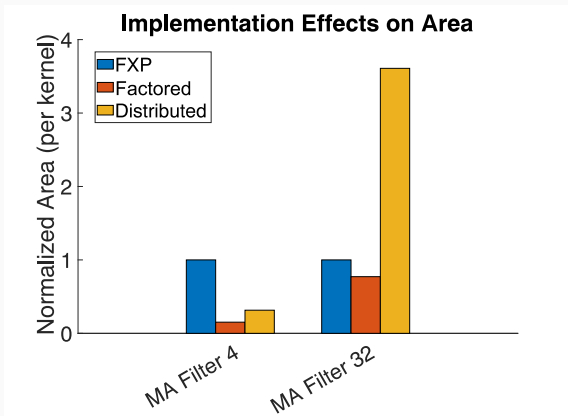
$$0.25 * d1 + 0.25 * d2 + 0.25 * d3 + 0.25 * x$$



Example: Moving Average Filter



Example: Moving Average Filter




Remark

BITSAD allows users to effectively explore these trade-offs

Work in Progress

Example: State Variable Filter

Constant coefficient: F , Q , input: x , output: y




Baseline:

$$d1 = F * (x - d2 - Q * d1_old) + d1_old$$

Example: State Variable Filter

Constant coefficient: F , Q , input: x , output: y



Baseline:


$$d1 = F * (x - d2 - q * d1_old) + d1_old$$

with strength reduction:

$$d1 = F * x - F * d2 - F * Q * d1_old + d1_old$$

Example: State Variable Filter

Constant coefficient: F , Q , input: x , output: y



Baseline:

$$d1 = F * (x - d2 - q * d1_old) + d1_old$$

with strength reduction:


$$d1 = F * x - F * d2 - F * Q * d1_old + d1_old$$

with strength reduction and algebraic simplification:

$$d1 = F * x - F * d2 + (1 - F * Q) * d1_old$$

Example: State Variable Filter

Constant coefficient: F , Q , input: x , output: y




Let BITSAD handle this

- try different combinations
- compare synthesized results
- choose the best implementation

Latency Issue

Bitstreaming Computing latency is high!

$$p = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T X_t$$




Error Bound for varying $p \in [0.01, 1]$ and varying T as well

Latency Issue

Bitstreaming Computing latency is high!

$$p = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T X_t$$




How do we trade off between latency and accuracy?

Latency Issue

Bitstreaming Computing latency is high!

$$p = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T X_t$$




How do we trade off between latency and accuracy?

Bit-level, cycle-accurate software simulation in BITSAD

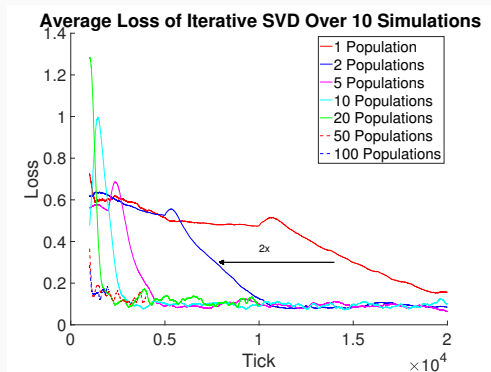
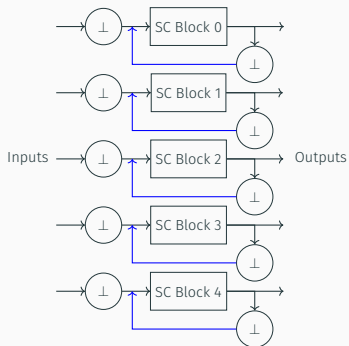
Latency Issue

Population coding,
inspired from biology:

$$\bar{p} = \frac{1}{T/N} \sum_{t=1}^{T/N} \frac{1}{N} \sum_{i=1}^N X_{i,t}$$

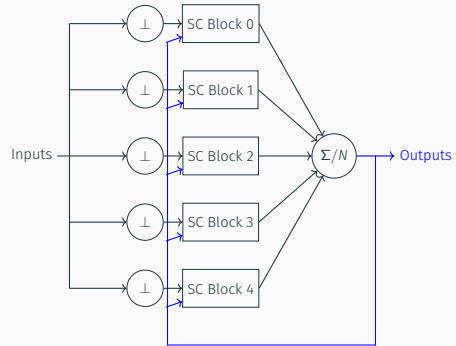
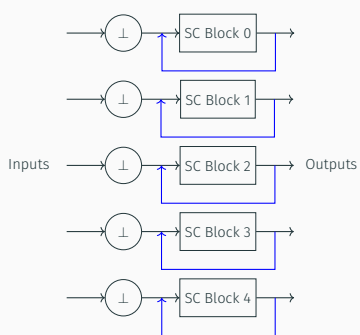


Population Coding: good?

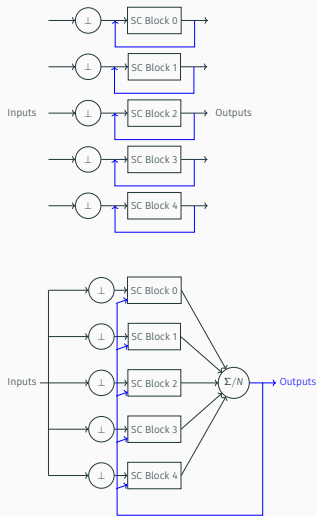


Experiments done with iterative SVD.

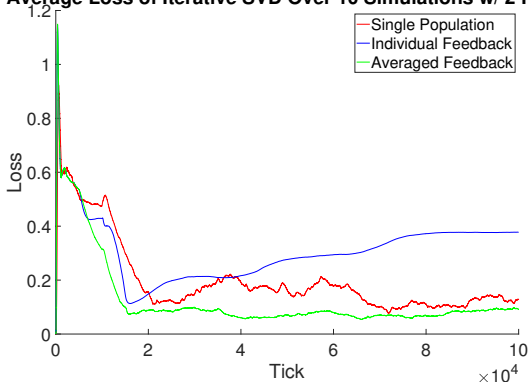
Population Coding: how to do better?



Population Coding: better!



Average Loss of Iterative SVD Over 10 Simulations w/ 2 Pops



Design optimization based on:

- Design details (RTL)
- Design requirements (Timing, Area, Power)

Design optimization iteration:

- Strength reduction and algebraic simplification
- Population coding
- more to explore

Conclusions

Ultra-low power/resource constrained applications require new computing paradigm

BITSAD allows:




- software algorithm testing
- Verilog generation automation
- fast design turnovers

What is more:

- BITBENCH, tomorrow at LCTES'19
(<https://github.com/UW-PHARM/BitBench>)
- BITSAD v2, in progress
(<https://github.com/UW-PHARM/BitSAD>)

Dont be bit sad, use BITSAD!
Questions?

References i

-  Duhamel, Pierre Emile et al. (2011). “Hardware in the loop for optical flow sensing in a robotic bee”. In: *IEEE International Conference on Intelligent Robots and Systems*, pp. 1099–1106. ISSN: 2153-0858. DOI: [10.1109/IROS.2011.6048759](https://doi.org/10.1109/IROS.2011.6048759).
-  Ma, Kevin Y. (2015). *RoboBee*. URL: <http://www.aboutkevinma.com/index.html#publications> (visited on 04/01/2018).
-  Zhang, Xuan et al. (2017). “A Fully Integrated Battery-Powered System-on-Chip in 40-nm CMOS for Closed-Loop Control of”. In: *IEEE Journal of Solid-State Circuits* 52.9, pp. 2374–2387. DOI: [10.1109/JSSC.2017.2705170](https://doi.org/10.1109/JSSC.2017.2705170).

What is Bitstream Computing?

Given a floating point number, p , as the mean of a Bernoulli distribution, can be encoded with a stochastic bitstream, the value of which, X is:

$$\mathbb{P}(X_t = 1) = p \quad \mathbb{P}(X_t = 0) = 1 - p \quad (1)$$

With timesteps T , p can be estimated as:

$$p = \mathbb{E}X_t \approx \frac{1}{T} \sum_{t=1}^T X_t \quad (2)$$