

BITBENCH: A Benchmark for Bitstream Computing

Languages, Compilers, and Toolchains for Embedded Systems '19

Kyle Daruwalla, Heng Zhuo, Carly Schulz, and Mikko Lipasti

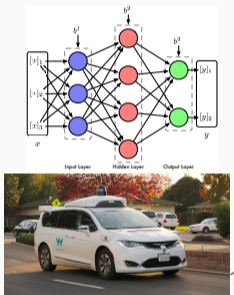
23 June 2019

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



Motivation

CV/ML algorithms enable powerful new applications



Fabrication techniques enabling pico-aerial vehicles (PAVs)

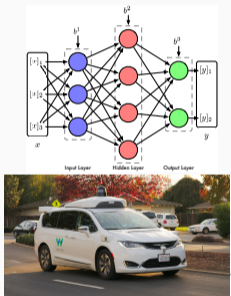


¹Dllu 2017

²Ma 2015

Motivation

CV/ML algorithms enable powerful new applications



Fabrication techniques enabling pico-aerial vehicles (PAVs)



How do we enable advanced CV/ML algorithms on ultra-low power platforms?

Roadmap for talk:

1. Highlight the need for bitstream computing
2. Discuss typical applications of bitstream computing
3. Present a series of benchmarks that cover typical applications (BITBENCH)
4. Show some initial results on the benchmarks
5. Discuss potential future directions

Motivation and Background

Limitations of Current Approaches

PAVs have ultra-low resources constraints
(< 35 mW compute power)³

Traditional computing paradigms cannot meet these constraints!

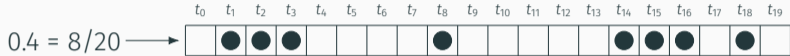
- Prior system tethered to desktop computer
- Current on-device chip can only support basic stationary flight⁴

³Duhamel et al. 2011.

⁴Zhang et al. 2017.

What is Bitstream Computing?

Stochastic Bitstreams:



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

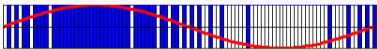
What is Bitstream Computing?

Stochastic Bitstreams:



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

Deterministic Bitstreams:



Density of "1" \Rightarrow Higher amplitude

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

Applications of Bitstream Computing

Looking at RoboBee⁵

A flying robot should be able to:

1. Mimic biological bees
2. Perform aerial surveillance
3. Collect weather data
4. Assist search and rescue
5. Communicate with other robotic bees

⁵Floreano and Wood 2015.

A flying robot should be able to:

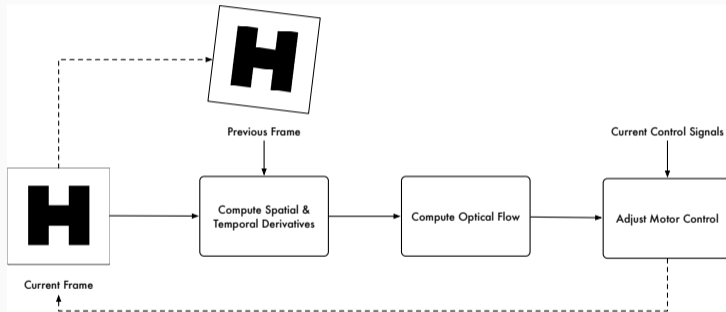
1. Mimic biological bees (object tracking, optical flow, inverse kinematics)
2. Perform aerial surveillance (optical flow, inverse kinematics)
3. Collect weather data (optical flow, inverse kinematics)
4. Assist search and rescue (object tracking, homography decomposition)
5. Communicate with other robotic bees (audio filtering)

⁵Floreano and Wood 2015.

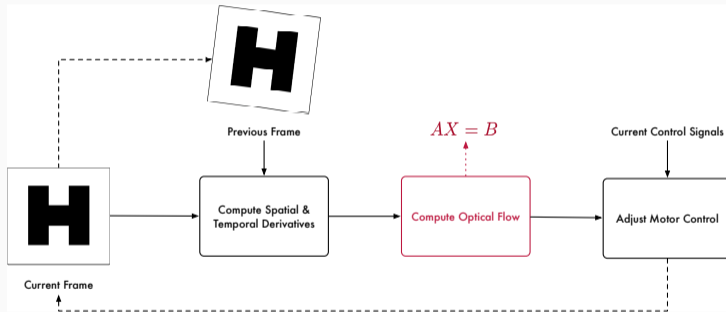
Introduce kernels with three examples:

1. Drone stabilization
2. Navigation and path planning
3. Channel shaping

An Example: Drone Stabilization



An Example: Drone Stabilization



Algorithm 1 Linear Solver (Pseudoinverse)⁶

Require: Input matrices $A \in \mathbb{R}^{m \times p}$, $B \in \mathbb{R}^{m \times n}$, Step-size $\alpha > 0$

- 1: **for** $k = 1 \dots T$ **do**
 - 2: $X_{k+1} \leftarrow (I - \alpha A^T A)X_k + \alpha A^T B$
 - 3: **end for**
 - 4: **return** X_T , final solution
-

⁶Shukla et al. 2018.

Algorithm 2 Linear Solver (Pseudoinverse)⁶

Require: Input matrices $A \in \mathbb{R}^{m \times p}$, $B \in \mathbb{R}^{m \times n}$, Step-size $\alpha > 0$

- 1: **for** $k = 1 \dots T$ **do**
 - 2: $X_{k+1} \leftarrow (I - \alpha A^T A)X_k + \alpha A^T B$
 - 3: **end for**
 - 4: **return** X_T , final solution
-

Used to implement:

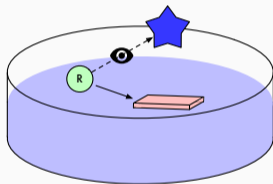
- Optical flow
- Object tracking
- Inverse kinematics

Algorithm	Size of A	Size of B
Optical Flow	4×2	4×1
Object Tracking	3×3	3×3
Inverse Kinematics	2×2	2×1

⁶Shukla et al. 2018.

Navigation and Path Planning

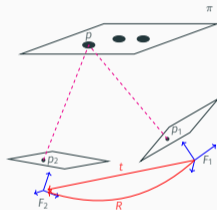
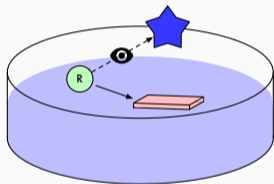
Consider a robot (denoted “R”) that needs to navigate an unknown environment



⁷Malis and Vargas 2007.

Navigation and Path Planning

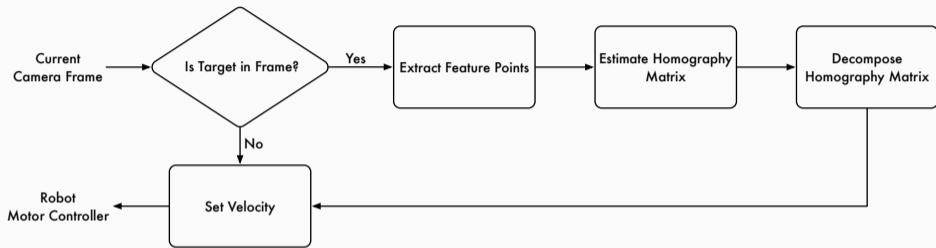
Consider a robot (denoted “R”) that needs to navigate an unknown environment



Navigation by visual cues — homography estimation and decomposition⁷

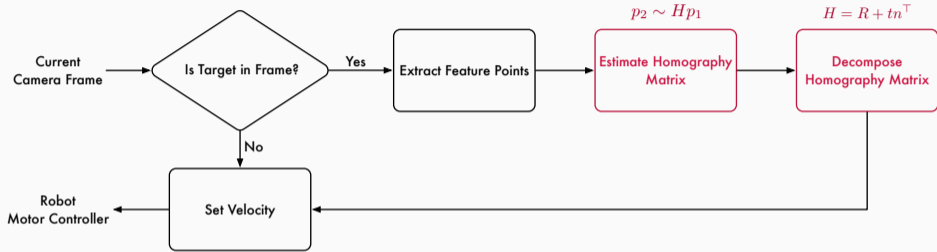
⁷Malis and Vargas 2007.

Navigation Pipeline



⁸Dubrofsky 2009.

Navigation Pipeline



Requires linear solver and singular value decomposition⁸

⁸Dubrofsky 2009.

Singular Value Decomposition (SVD)

Algorithm 3 Iterative SVD⁹

Require: Input matrix $A \in \mathbb{R}^{m \times n}$ and initial guess $v_0 \in \mathbb{R}^n$

1: **for** $k = 1, 2, \dots$ (until convergence) **do**

2: $w_k = Av_{k-1}$

3: $\alpha_k = \|w_k\|_2 = \sqrt{w_k^T w_k}$

4: $u_k = w_k / \alpha_k$

5: $z_k = A^T u_k$

6: $\sigma_k = \|z_k\|_2 = \sqrt{z_k^T z_k}$

7: $v_k = z_k / \sigma_k$

8: **end for**

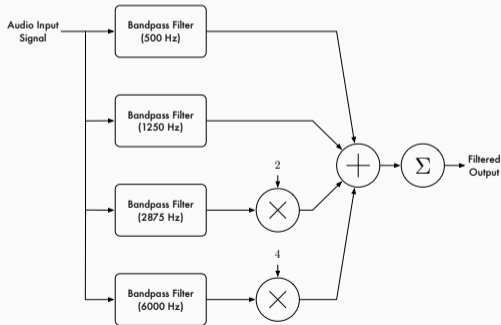
9: **return** First left/right singular vectors, u_k & v_k , and first singular value, σ_k

⁹Bentbib and Kanber 2015.

Channel Shaping

Audio information is sent across several frequency channels

Need to shape input signal to focus on channels of interest

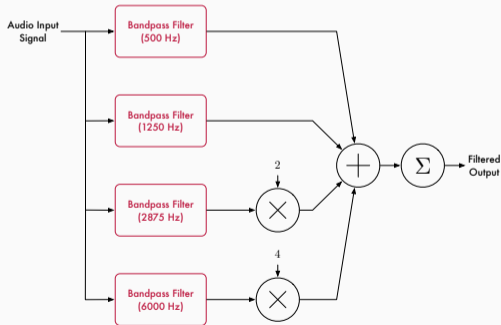


Channel Shaping

Audio information is sent across several frequency channels

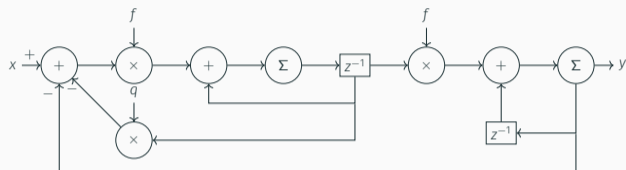
Need to shape input signal to focus on channels of interest

How do we create efficient filters?



Traditional Filters

Consider the digital state-variable filter (SVF)¹⁰



Filter coefficients given by:

$$f = 2 \sin \left(\frac{\pi F_c}{F_s} \right)$$

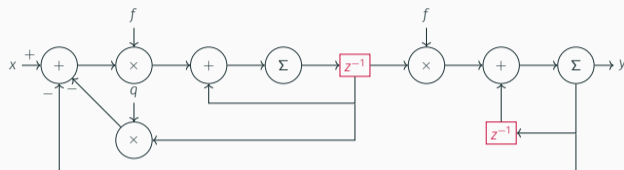
$$q = \frac{1}{Q}$$

Parameter	Symbol
Center Frequency	F_c
Sample Frequency	F_s
Quality Factor	Q

¹⁰Redmond 2003.

Traditional Filters

Consider the digital state-variable filter (SVF)¹⁰



Filter coefficients given by:

$$f = 2 \sin \left(\frac{\pi F_c}{F_s} \right)$$

$$q = \frac{1}{Q}$$

Choose F_s by adjusting delay blocks to get better coefficients!

¹⁰Redmond 2003.

Bitstream Filtering

First choose f and q to be binary-friendly (e.g. $f = 0.125$ and $q = 1.875$)

Set effective sample frequency according to¹¹:

$$F_s = \frac{\pi F_c}{\arcsin(f/2)}$$

$$d = \left\lceil \frac{F_{\text{PDM}}}{F_s} \right\rceil$$

Parameter	Symbol
Center Frequency	F_c
Sample Frequency	F_s
PDM Sample Rate	F_{PDM}
Delay	d

¹¹Lipasti and Schulz 2017.

Bitstream Filtering

First choose f and q to be binary-friendly (e.g. $f = 0.125$ and $q = 1.875$)

Set effective sample frequency according to¹¹:

$$F_s = \frac{\pi F_c}{\arcsin(f/2)}$$
$$d = \left\lceil \frac{F_{\text{PDM}}}{F_s} \right\rceil$$

Parameter	Symbol
Center Frequency	F_c
Sample Frequency	F_s
PDM Sample Rate	F_{PDM}
Delay	d

Using binary-friendly coefficients results in 4-bit FXP units instead of 16-bit units

¹¹Lipasti and Schulz 2017.

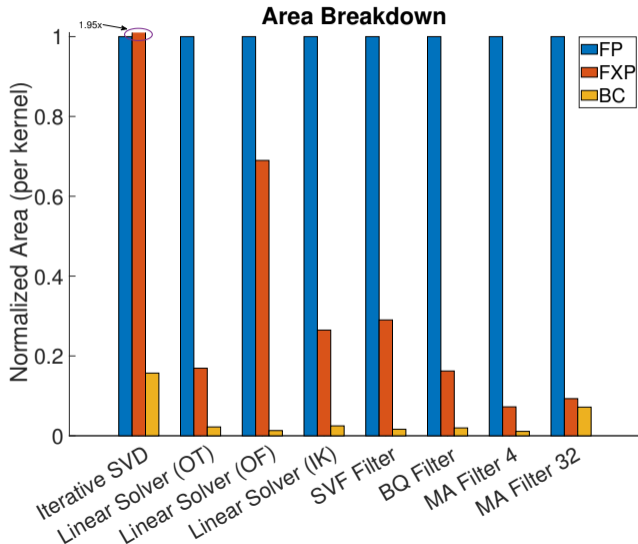
BITBENCH **Kernels**

Summary of Kernels

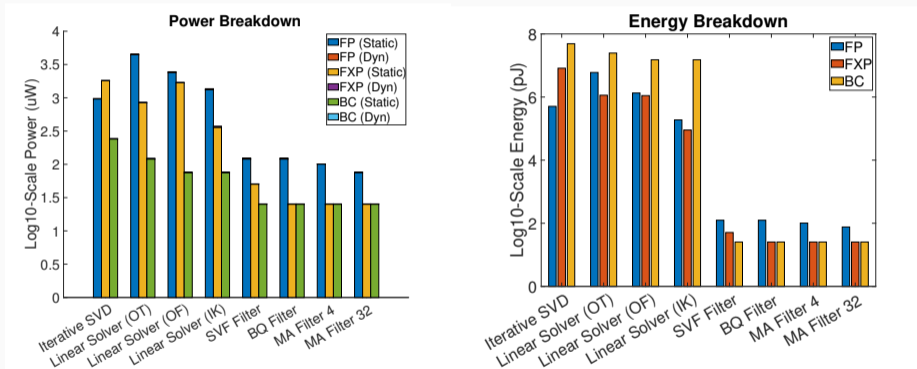
We present the following kernels:

- Linear Solver (Optical Flow)
- Linear Solver (Object Tracking)
- Linear Solver (Inverse Kinematics)
- Iterative SVD
- State-Variable Filter
- Bi-Quad Filter
- Moving Average Filter (Width of 4)
- Moving Average Filter (Width of 32)

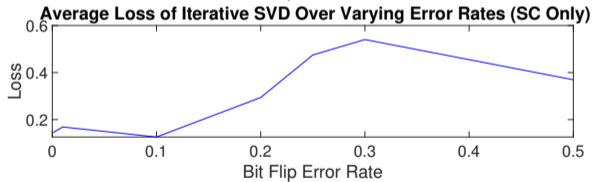
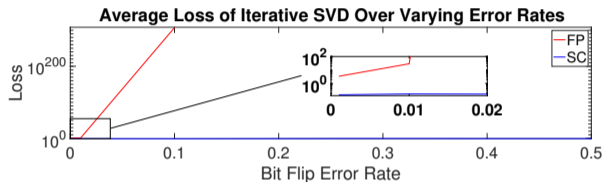
Area Results



Power and Energy Results



Fault Tolerance



Conclusions

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

Bitstream computing is:

- Low resource
- Ultra-low power
- Under-researched (opportunity for novel work)

BITBENCH introduces 8 kernels centered on:

- Linear solvers
- Singular value decomposition
- Filtering

Conclusions

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

Bitstream computing is:





• Low resource

Check out the benchmarks at: <https://github.com/UW-PHARM/BitBench>
Built using our domain-specific language for bitstream computing
(<https://github.com/UW-PHARM/BitSAD>)




- Linear solvers
- Singular value decomposition
- Filtering

Questions?




References i

-  Bentbib, A. H. and A. Kanber (2015). “Block power method for SVD decomposition”. In: *Analele Stiintifice ale Universitatii Ovidius Constanta, Seria Matematica* 23.2, pp. 45–58. ISSN: 18440835. DOI: [10.1515/auom-2015-0024](https://doi.org/10.1515/auom-2015-0024).
-  Dllu (2017). URL: https://commons.wikimedia.org/wiki/File:Waymo_Chrysler_Pacifica_in_Los_Altos,_2017.jpg.
-  Dubrofsky, Elan (2009). “Homography Estimation”. PhD thesis. Carleton University.
-  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).

References ii

-  Floreano, Dario and Robert J. Wood (2015). “Science, technology and the future of small autonomous drones”. In: *Nature* 521.7553, pp. 460–466. ISSN: 14764687. DOI: [10.1038/nature14542](https://doi.org/10.1038/nature14542).
-  Lipasti, Mikko and Carly Schulz (2017). *End-to-End Stochastic Computing*. Austin, TX, USA. URL: http://pharm.ece.wisc.edu/papers/wp3_2017.pdf.
-  Ma, Kevin Y. (2015). *RoboBee*. URL: <http://www.aboutkevinma.com/index.html#publications> (visited on 04/01/2018).

References iii

-  Malis, Ezio and Manuel Vargas (2007). “Deeper understanding of the homography decomposition for vision-based control”. In: *Sophia* 6303.6303, p. 90. ISSN: 0036-8075. DOI: 10.1126/science.318.5857.1691b. URL: <http://hal.archives-ouvertes.fr/inria-00174036/>.
-  Redmond, Nigel (2003). *The digital state variable filter*. URL: <http://www.earlevel.com/main/2003/03/02/the-digital-state-variable-filter/>.
-  Shukla, Rohit et al. (2018). “Computing Generalized Matrix Inverse on Spiking Neural Substrate”. In: *Frontiers in Neuroscience* 12, p. 115. ISSN: 1662-453X. DOI: 10.3389/fnins.2018.00115. URL: <https://www.frontiersin.org/article/10.3389/fnins.2018.00115>.



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).