


ApHMM: Accelerating Profile Hidden Markov Models for Fast and Energy-Efficient Genome Analysis

Can Firtina¹, Kamlesh Pillai², Gurpreet S. Kalsi², Bharathwaj Suresh², Damla Senol Cali³, Jeremie S. Kim¹, Taha Shahroodi⁴, Meryem Banu Cavlak¹, Joel Lindegger¹, Mohammed Alser¹, Juan Gómez Luna¹, Sreenivas Subramoney², and Onur Mutlu^{1,3}



arXiv Preprint


SAFARI

¹ETH zürich

²intel

³Carnegie Mellon

⁴TU Delft
Delft University of Technology



Source Code


1: Profile Hidden Markov Models (pHMMs)

PHMM Sequence: ...ACTT...

Sequence #1: ...AGGGCTT...

Sequence #2: ...ATT... (Deleted C)

Sequence #3: ...ACTG...



Insertions (I)

Deletions (D)

Substitution or Match

PHMMs are useful to identify **differences and similarities** between sequences using:

Transitions between states

Recognition (**emission**) of each character within states

Probabilities assigned to transitions and emissions

2: The Baum-Welch Algorithm

1. Forward Calculation

$$F_t(i) = \sum_{j \in V} F_{t-1}(j) \alpha_{ji} e(S[t], v_i) \quad i \in V, \quad 1 < t \leq n_S$$

3.1. Updating Transitions

$$\alpha_{ij}^* = \frac{\sum_{t=1}^{n_S-1} \alpha_{ij} e_{S[t+1]}(v_j) F_t(i) B_{t+1}(j)}{\sum_{t=1}^{n_S-1} \sum_{x \in V} \alpha_{ix} e_{S[t+1]}(v_x) F_t(i) B_{t+1}(x)} \quad \forall \alpha_{ij} \in A$$

2. Backward Calculation

$$B_t(i) = \sum_{j \in V} B_{t+1}(j) \alpha_{ij} e(S[t+1], v_j) \quad i \in V, \quad 1 \leq t < n_S$$



3.2. Updating Emissions

$$e_X^*(v_i) = \frac{\sum_{t=1}^{n_S} F_t(i) B_t(i) [S[t] = X]}{\sum_{t=1}^{n_S} F_t(i) B_t(i)} \quad \forall X \in \Sigma, \forall i \in V$$


Massive Parallelism

Data Dependency

3: Sources of Inefficiencies in the Baum-Welch Algorithm



a) Forward Calculation in pHMMs



b) Forward Calculation in HMMs

The Baum-Welch algorithm is the **main computational overhead**

Standard HMM accelerators are **oblivious to the data dependency pattern** in pHMMs

4: Problem and Goal

- The Baum-Welch algorithm is useful for many applications but **remains computationally costly due to several sources of inefficiencies**

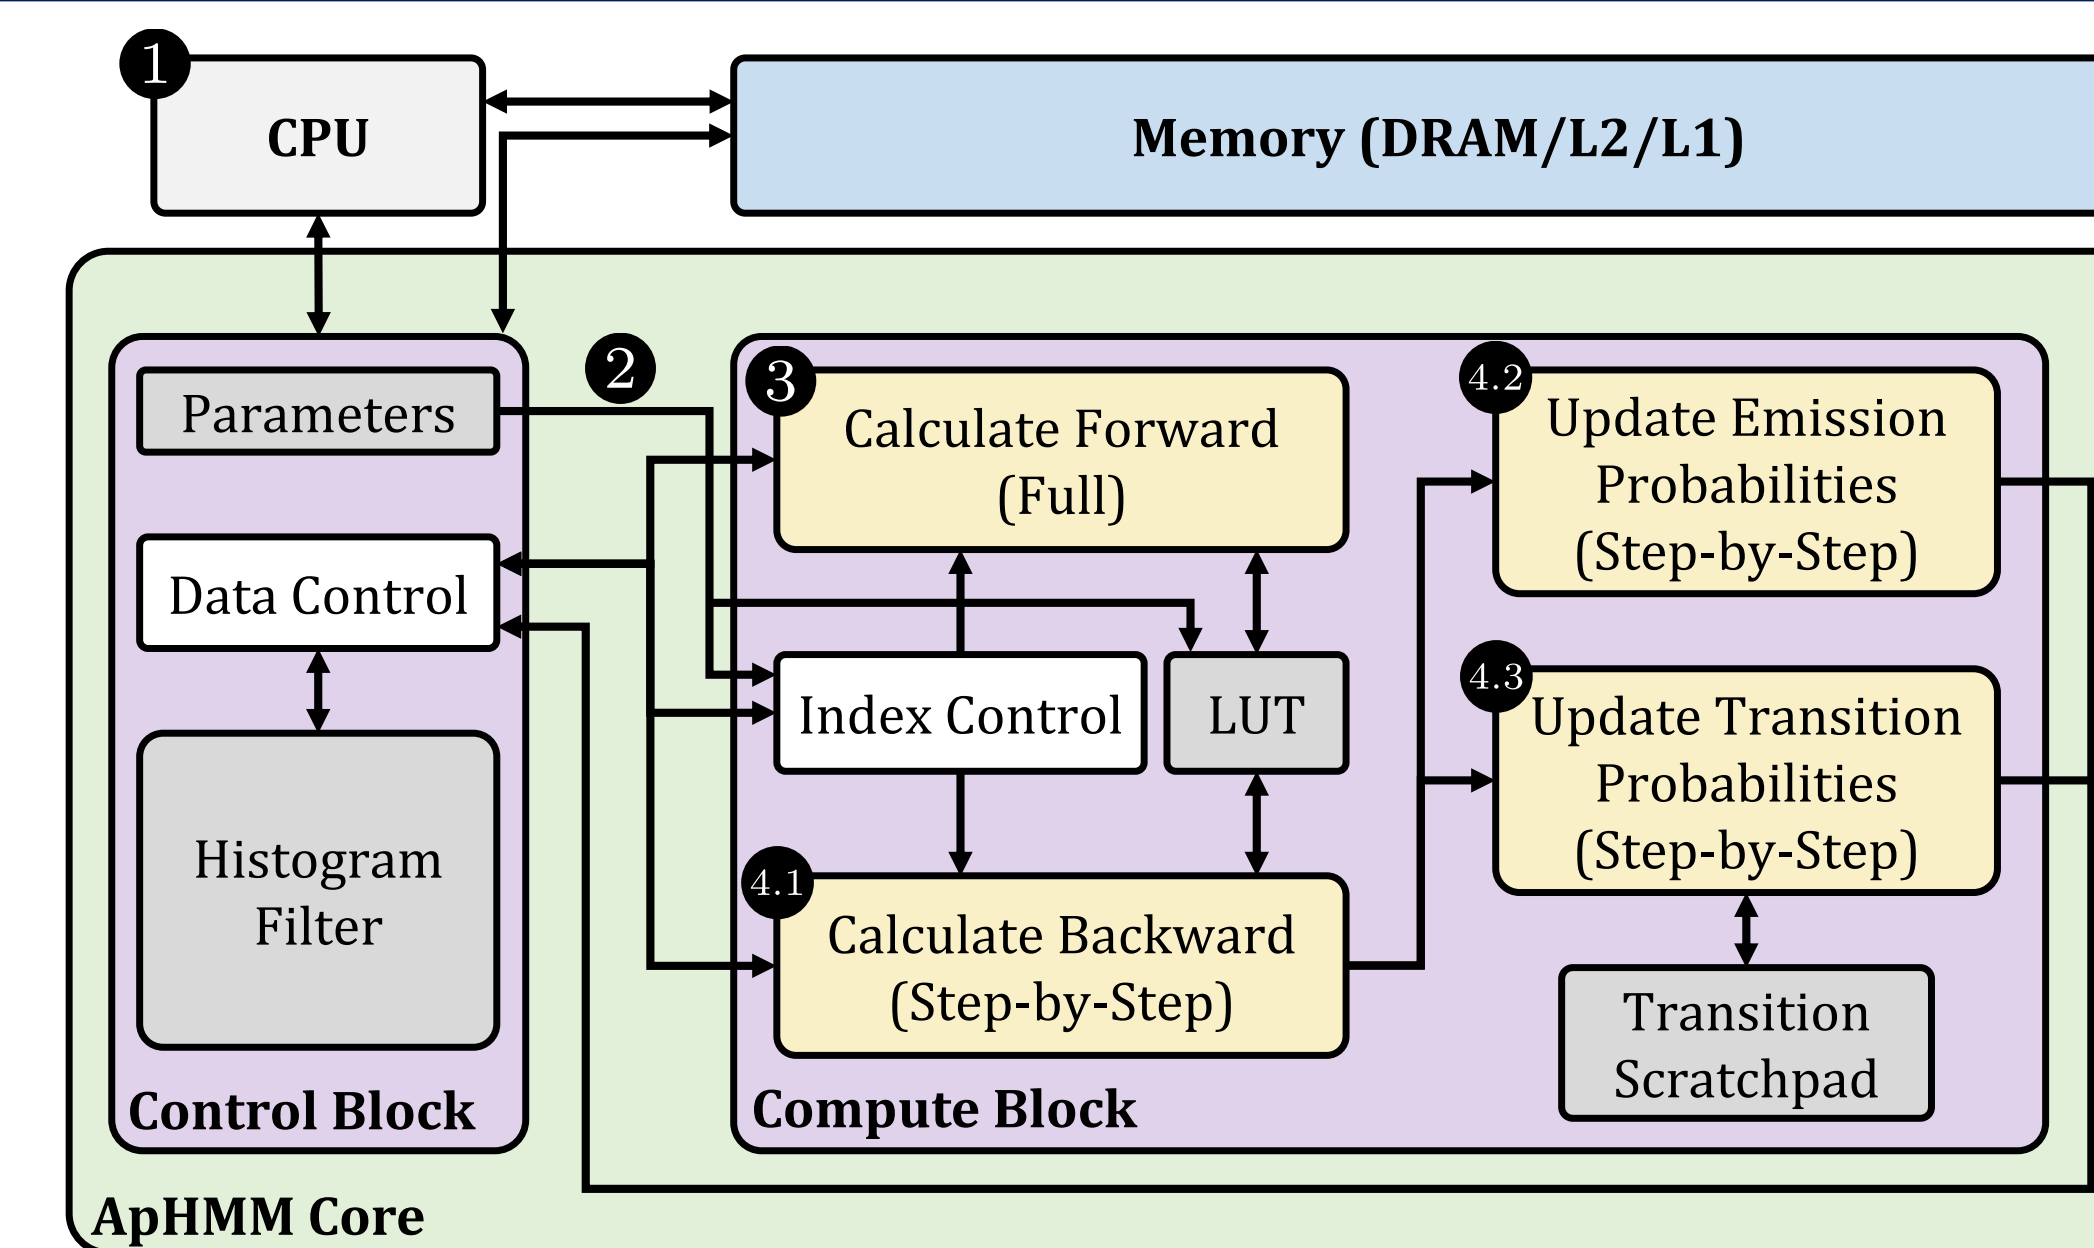
- Our goal is to
 - Accelerate the Baum-Welch algorithm** while

- Eliminating its inefficiencies**

- Using a hardware-software co-design**

- ApHMM is the first work that accelerates the Baum-Welch algorithm for pHMMs**

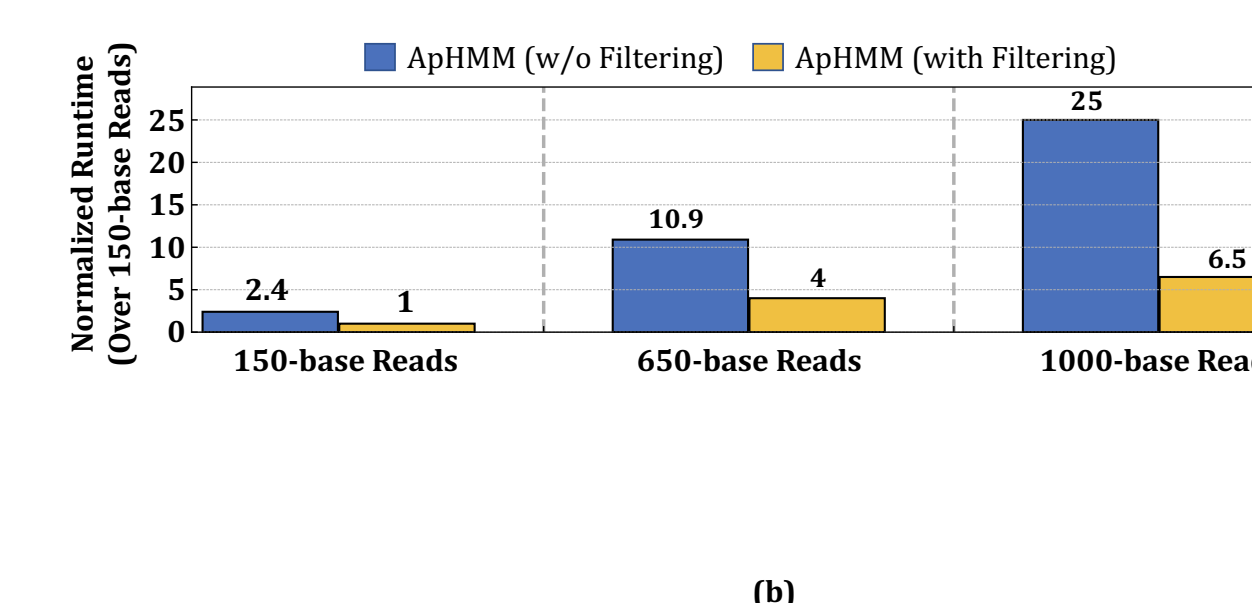
5: ApHMM Overview & Filtering Mechanism



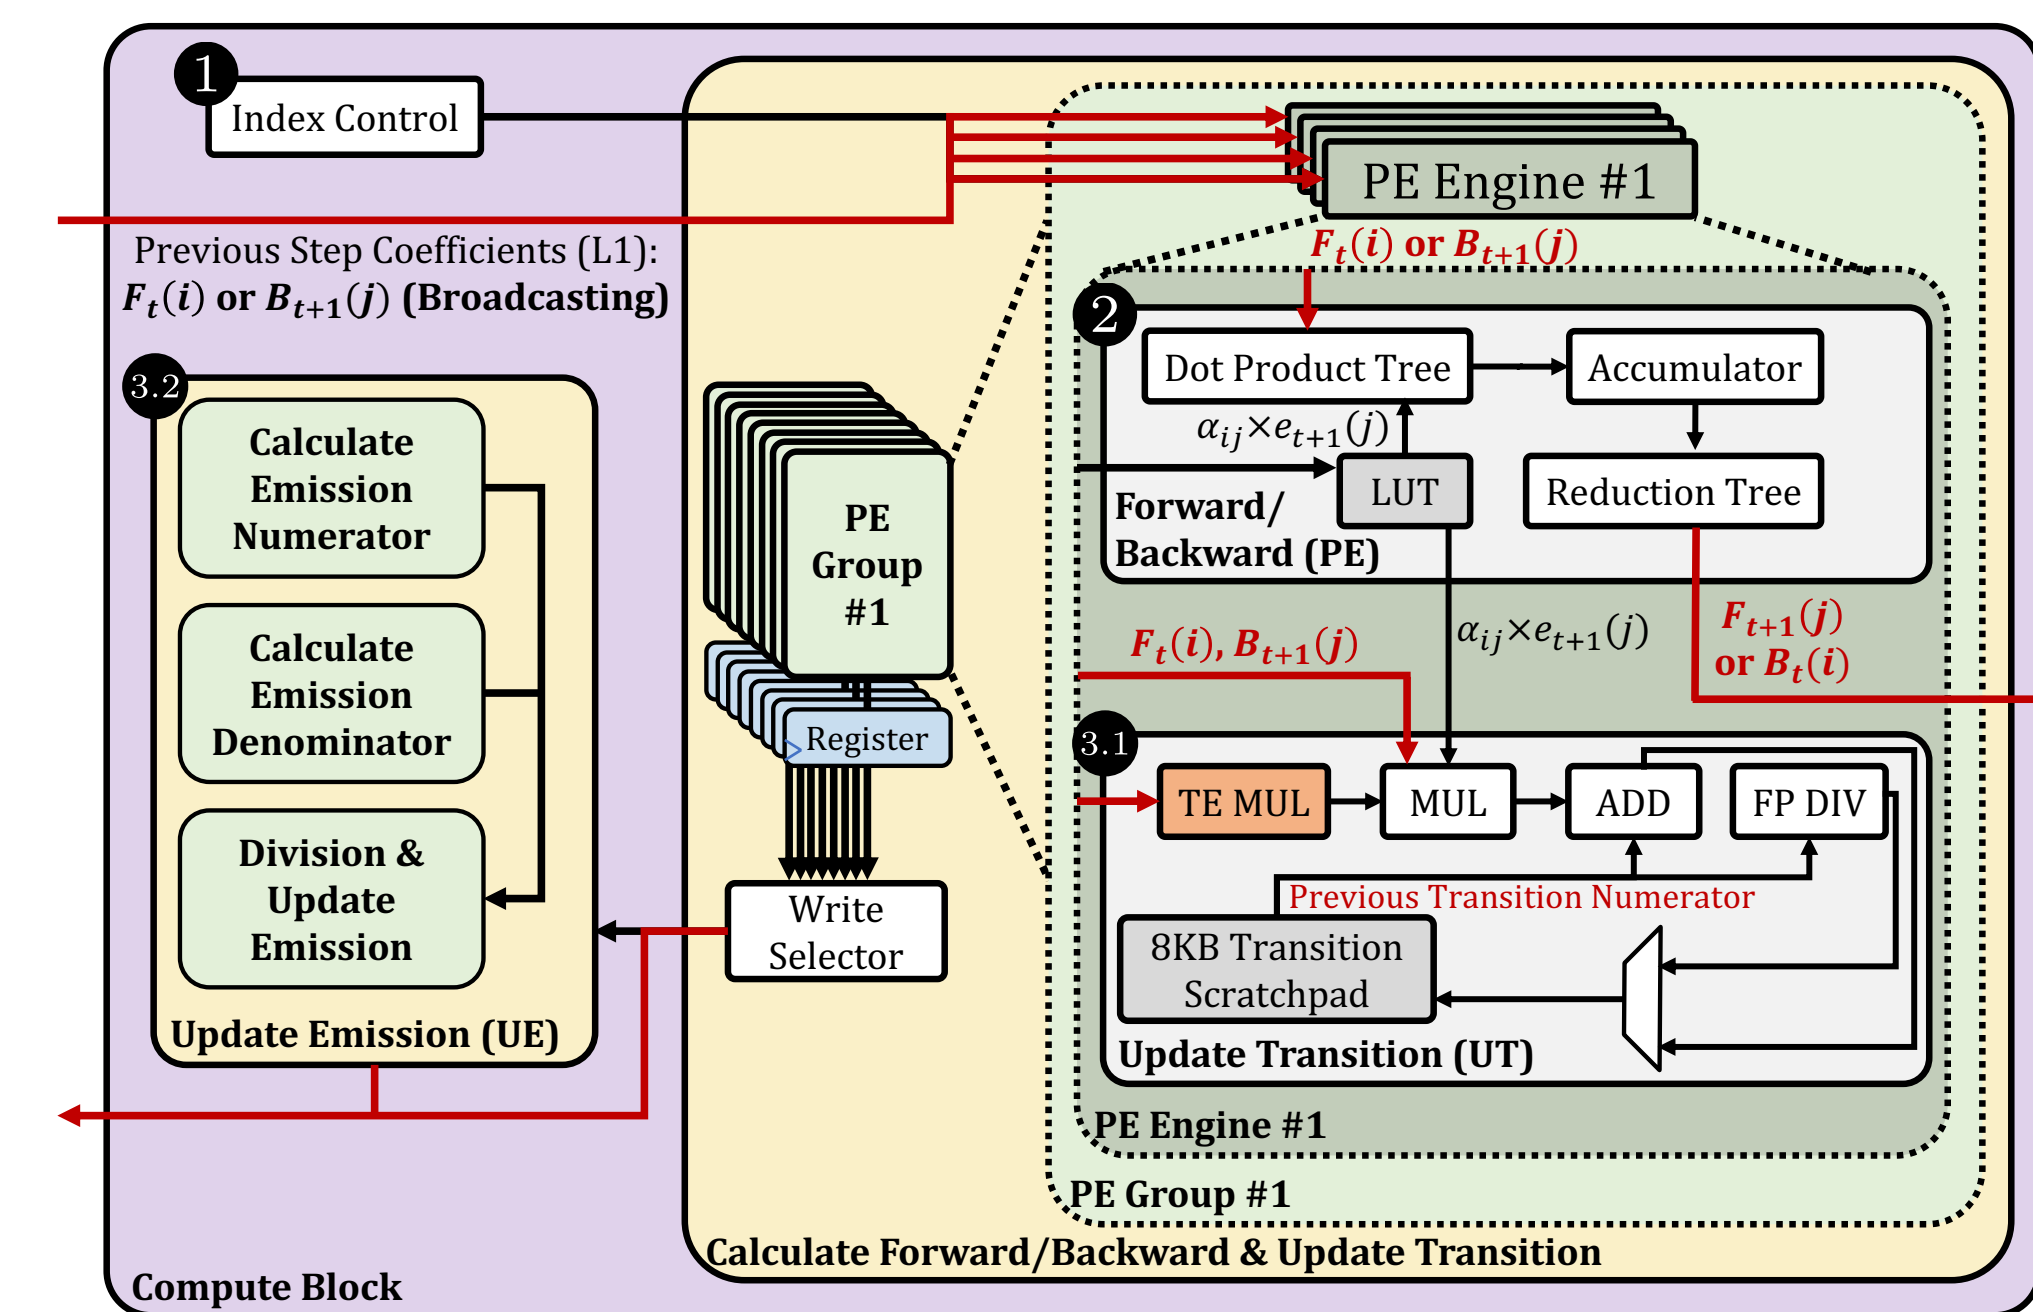
State IDs (i)	Max. Value
8, 9	1.00
10, 14	0.94
15, 16, 18	0.88
11, 20, 21, ...	0.82
13, 17, 19, ...	0.76
...	...
...	...
...	0.06

Same Memory Block
Filter is full
Ignore rest when the filter is full

Histogram Filter




6: Computing the Baum-Welch Algorithm



7: Evaluation Methodology

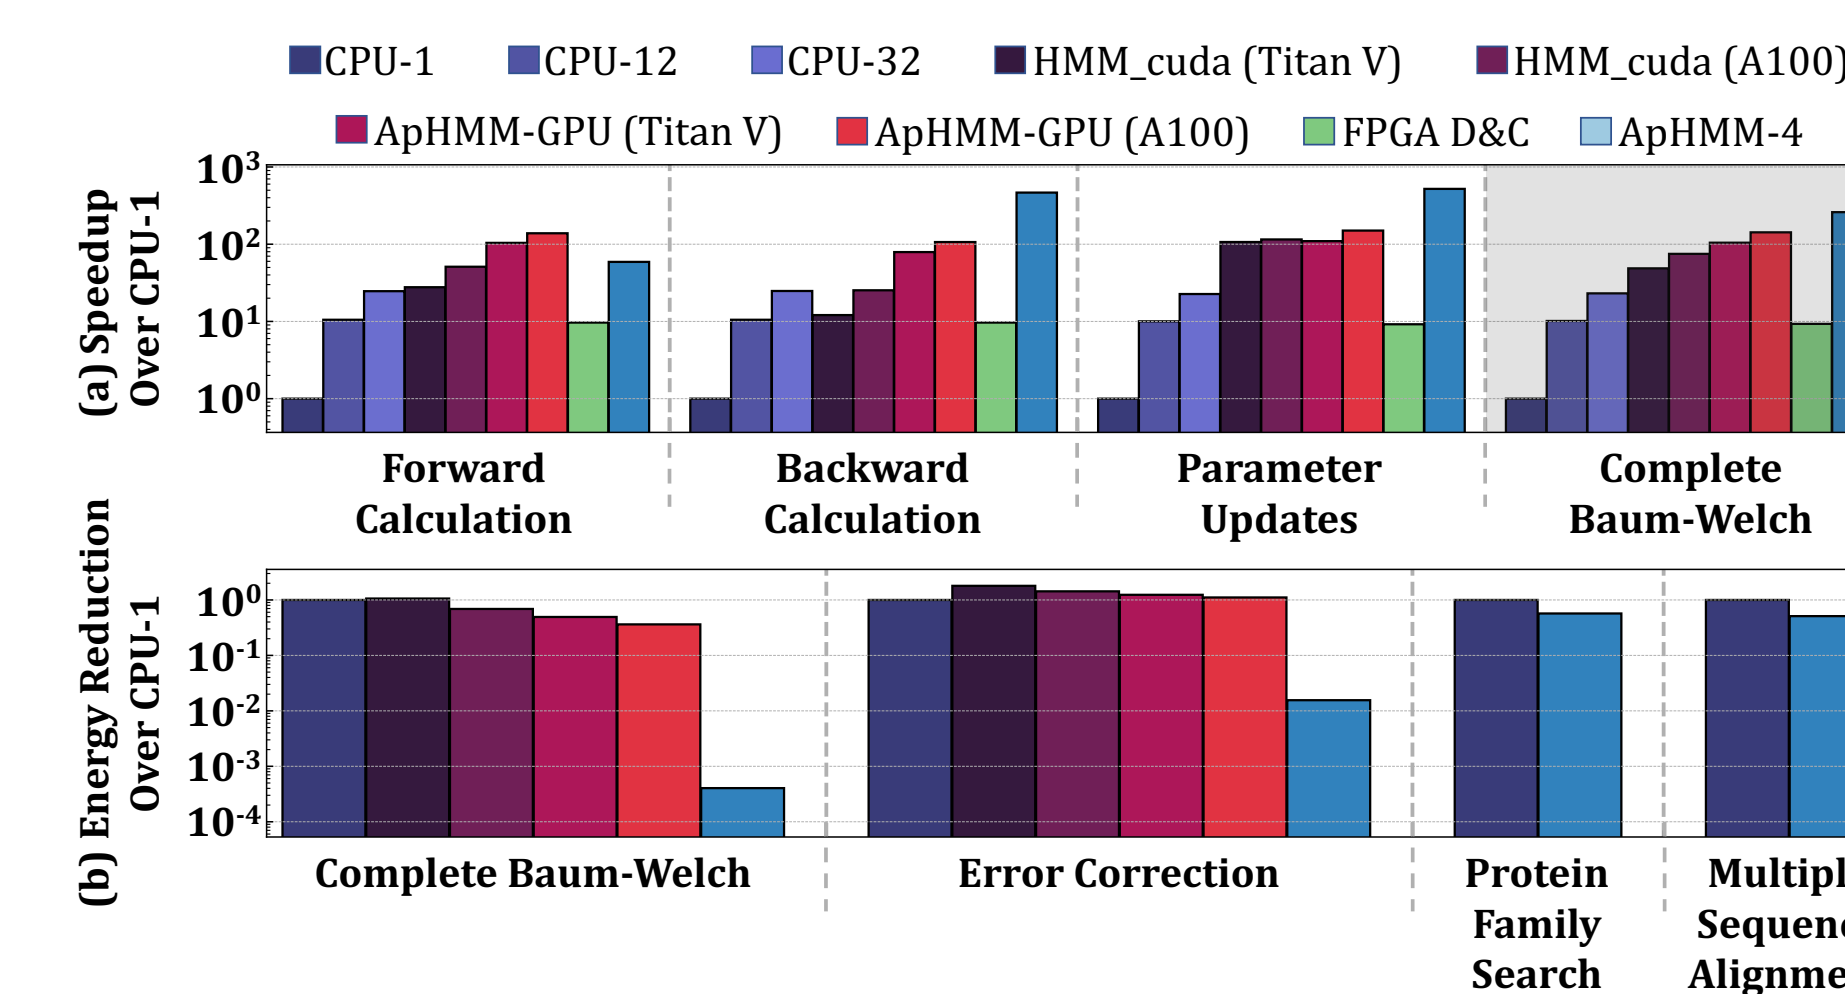
- ApHMM (**ASIC**) and ApHMM-**GPU** implementations
- Compared with **CPU, GPU, and FPGA** baselines
- Use cases
 - Error correction and polishing
 - Multiple sequence alignment
 - Protein Family Search
- Evaluating
 - Area and power
 - Performance
- Datasets
 - Error correction: 50,000 random E. coli reads
 - Protein family search: Mapping Mitochondrial carrier to the entire Pfam database
 - Multiple sequence alignment: Aligning commonly used protein families to each other

8: Results




✓ 4-core ApHMM provides the **best overall speedup**

✗ **Limited scalability** for using many cores with ApHMM due to the **data movement overhead**



✓ **Significant improvements** in terms of **performance and energy consumption**

✗ **Limited performance improvements** for the Forward Calculation due to the **data movement overhead**



✓ **ApHMM significantly accelerates the end-to-end execution times** of important applications