K-PACT: Kernel Planning for Adaptive Context Switching - A Framework for Clustering, Placement, and Prefetching in Spectrum Sensing

H. Umut Suluhan[1], Jiahao Lin[2], Serhan Gener[1],
Chaitali Chakrabarti[3], Umit Ogras[2] and Ali Akoglu[1]

[1]Department of Electrical and Computer Engineering, University of Arizona
{**suluhan**, gener, akoglu}@arizona.edu

[2]Department of Electrical and Computer Engineering, University of Wisconsin-Madison
{jlin445, uogras}@wisc.edu

[3]School of Electrical, Computer and Energy Engineering, Arizona State University
chaitali@asu.edu

COLLEGE OF ENGINEERING
Electrical & Computer Engineering

Reconfigurable Computing Lab

IEEE/ACM 2025 INTERNATIONAL CONFERENCE ON COMPUTER-AIDED DESIGN
ICCAD
44th Edition

Motivation and Contribution

Advanced spectrum sensing classify signal types within an environment

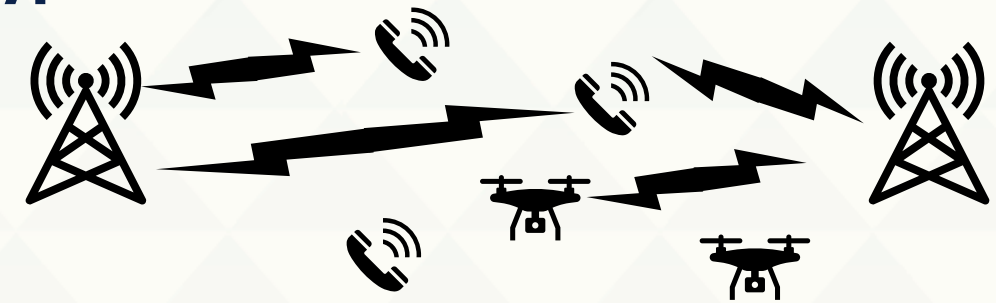

Seamless reconfiguration and parallel workload processing are key enabler for real-time spectrum sensing

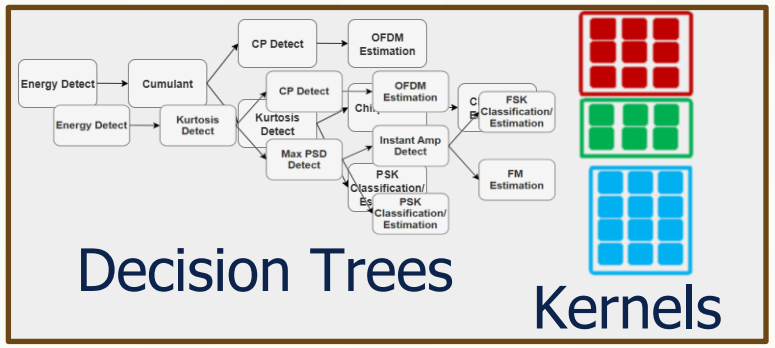

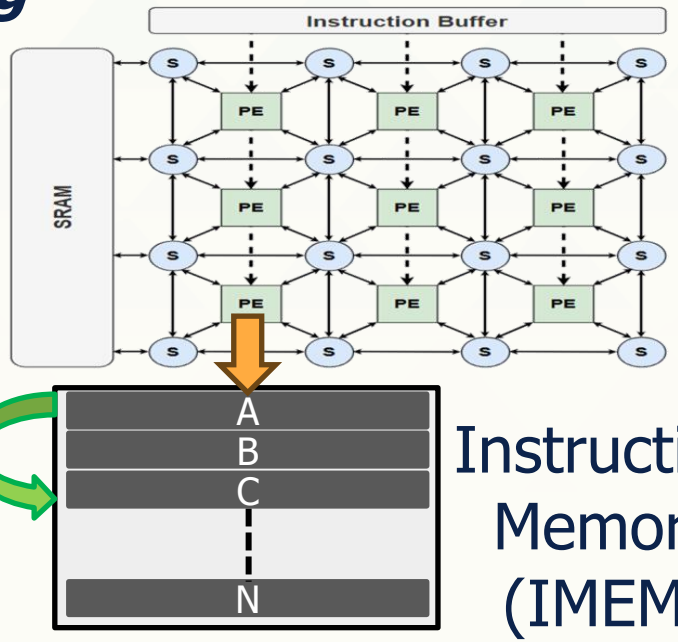

Intelligent PE allocation and IMEM management strategies are needed to mitigate costly overheads

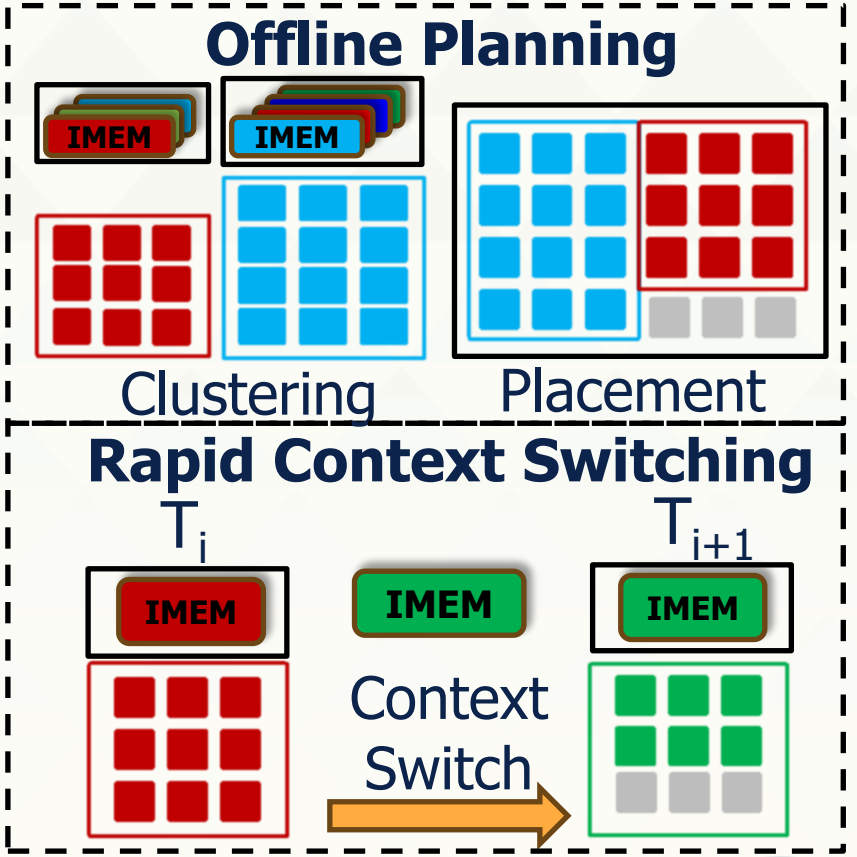

Clustering and Placement Approach

$$\mathcal{K}_i \in \mathcal{C}_j \iff (e_{i-1} < s_i) \wedge \sum_{\mathcal{K}_k \in \mathcal{C}_j} \mathrm{I}_{\mathrm{req}}(\mathcal{K}_k) < \mathrm{I}_{\mathrm{lim}},$$

$$\forall \mathcal{K}_i^{(s_i,e_i)} \in \{\mathcal{K}_1, \ldots, \mathcal{K}_N\}, \quad \forall \mathcal{C}_j \in \{\mathcal{C}_1, \ldots, \mathcal{C}_M\}$$

Objective Statement:

- Group N kernels into M clusters
- Satisfy an IMEM size constraint
- Minimize total number of clusters

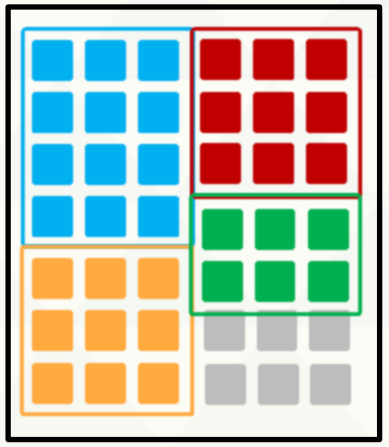

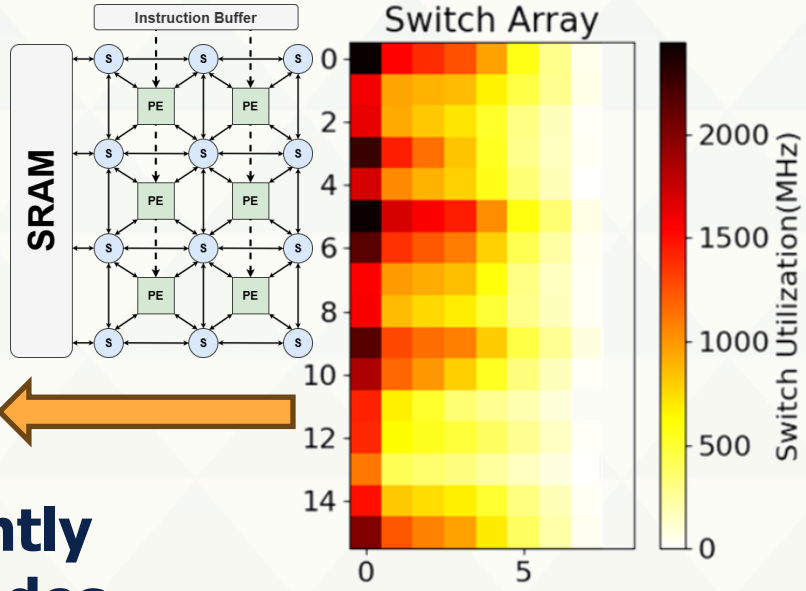

Prioritize frequently accessed head nodes

Routing Congestion on Switch Array

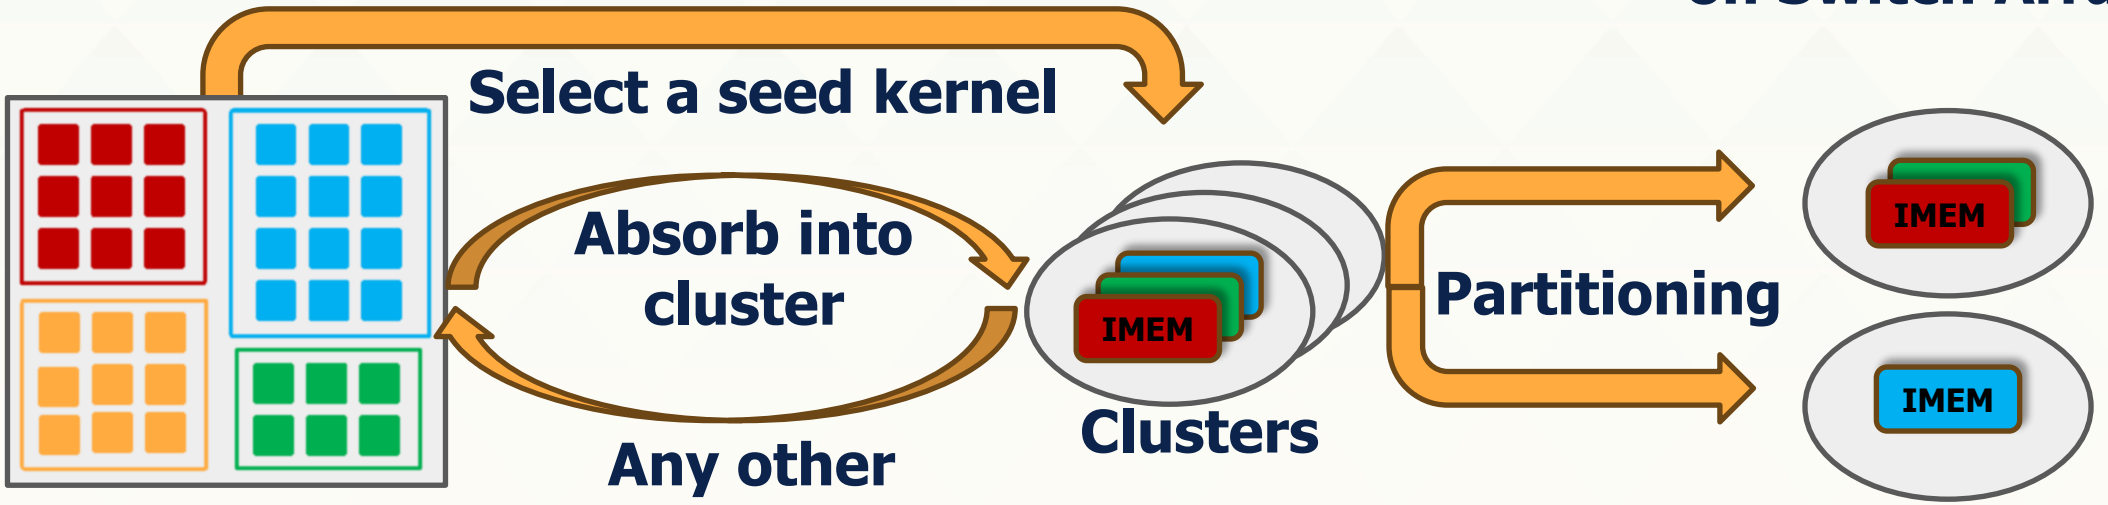

Pre-Initialization Strategies

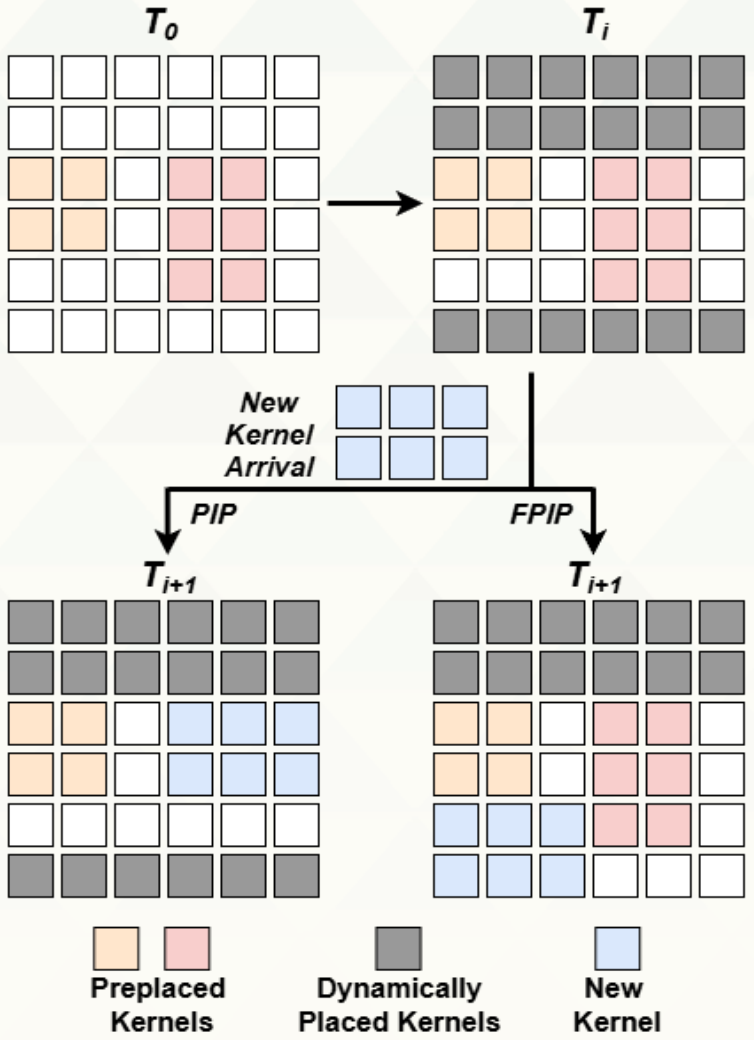

- **Pre-initialized Placement (PIP)**
  - dynamic resource manager can evict the pre-initialized kernels
  - diminishes the benefits of profiling-guided pre-initialization
- **Fixed Pre-initialized Placement (FPIP)**
  - dynamic resource manager cannot evict the pre-initialized kernels
  - maintains pre-initialization and fully leverages rapid switching capability

No/Soft/Hard Switch

No Switch
- Kernel is already loaded in IMEM
- Instantaneous switch

Soft Switch
- New kernel binary exists in IMEM
- *Cost:* Cluster search
- Low switching latency

Hard Switch
- New kernel binary is not in IMEM
- *Cost:* Cluster search, allocation, eviction
- High switching latency

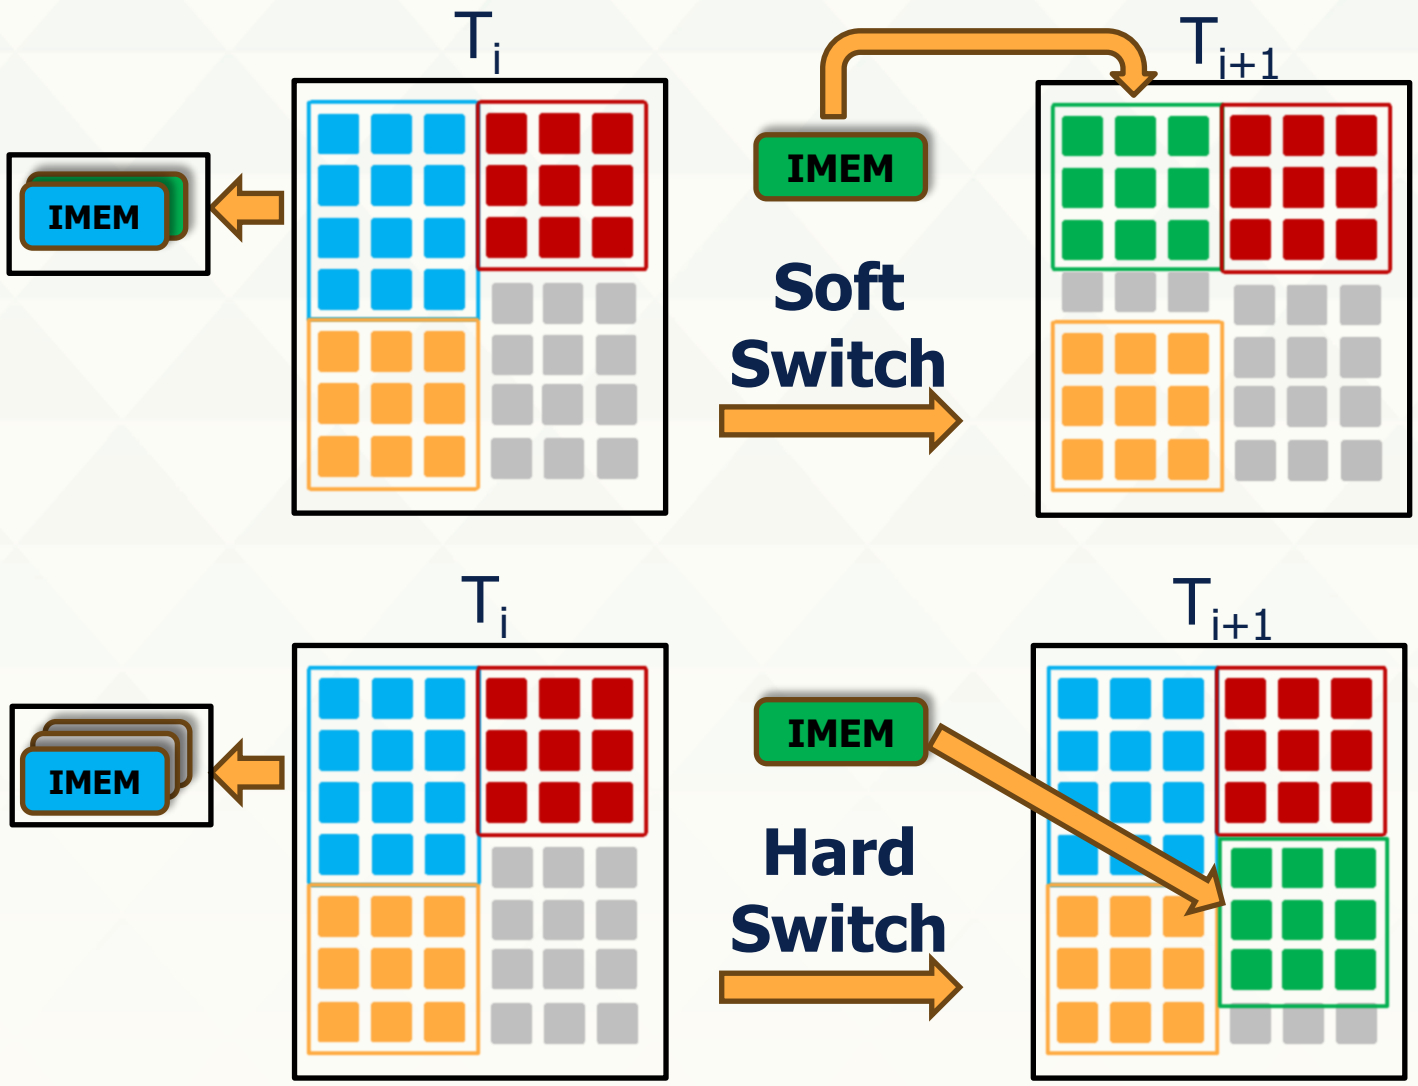

Experimental Setup

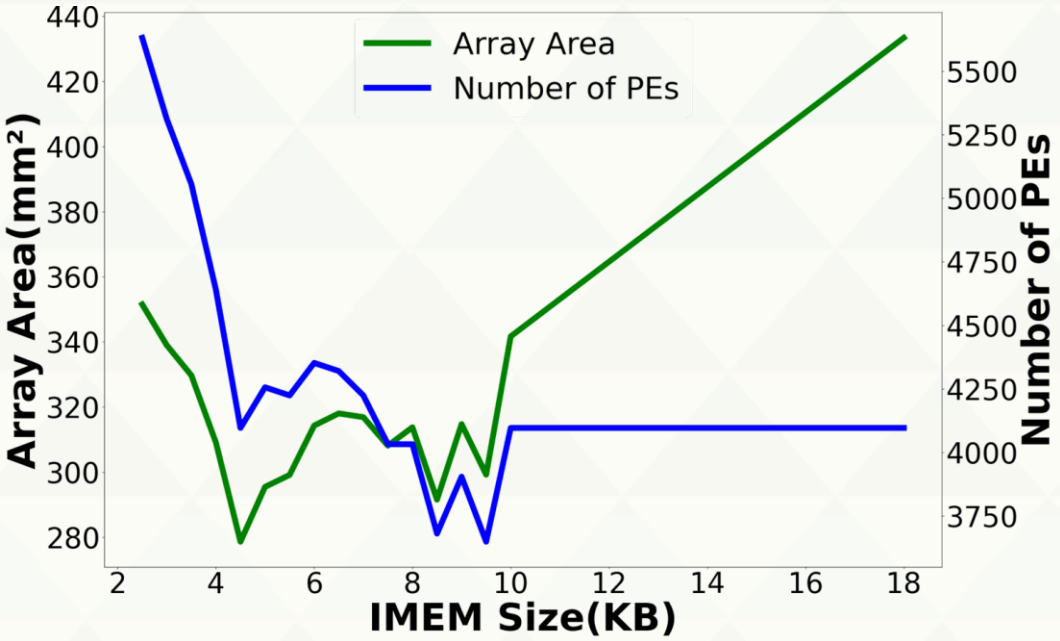

Optimal Configuration

- 4.5 KB IMEM size
- 4096 PEs
- 278.56 mm^2

Simulation Platform

- Open-source domain-specific system-on-chip simulator (DS3)*

Scheduling Modes

- **Baseline:** performs only hard switches
- **Dynamic Placer:** forms clusters and execute soft switches
- **Pre-initialized Placement+Dynamic Placer:** evicts kernels from the preplaced regions
- **Fixed Pre-initialized Placement+Dynamic Placer:** preserves preplaced kernels

Metrics

- Switch Count/Time
- Scheduling Time
- Execution Time per Tree

* S. E. Arda *et al.*, "DS3: A system-level domain-specific system-on-chip simulation framework," *IEEE Transactions on Computers*, vol. 69, no. 8, pp. 1248–1262, 2020.

Switch Time Analysis

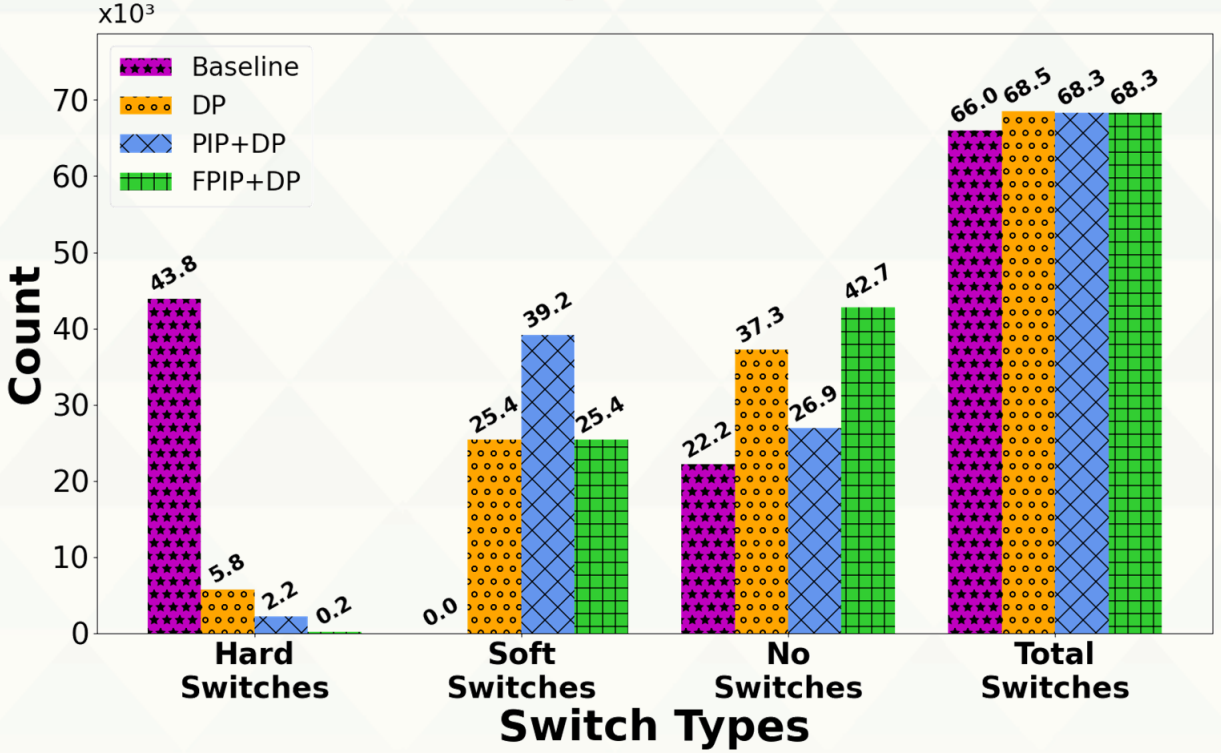

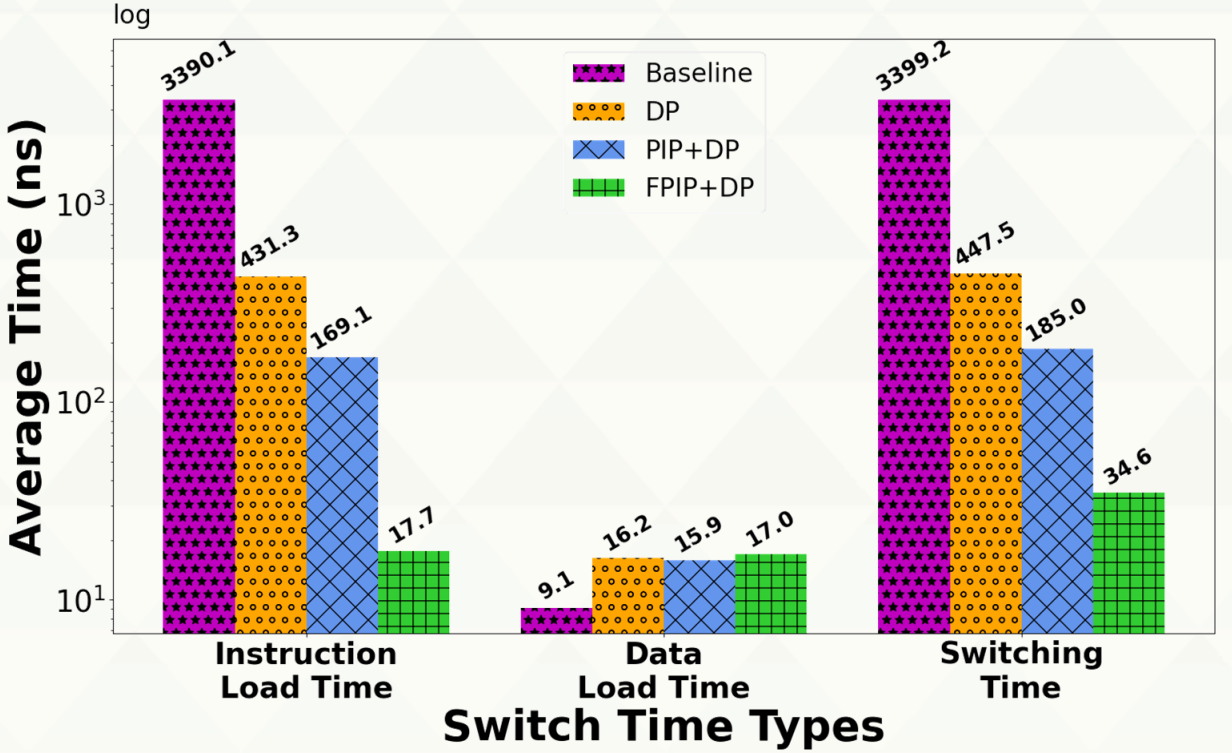

- 207.81x reduction in hard switch count
- PIP+DP -> high utilization of soft switches
- FPIP+DP -> extensive use of no switches

- 191.51x reduction in instruction load time
- Data load time becomes the bottleneck
- 98.24x reduction in overall switching time

Pre-initialized placement exploits rapid switching, reducing expensive off-chip fetches

Scheduling and Execution Time Analysis

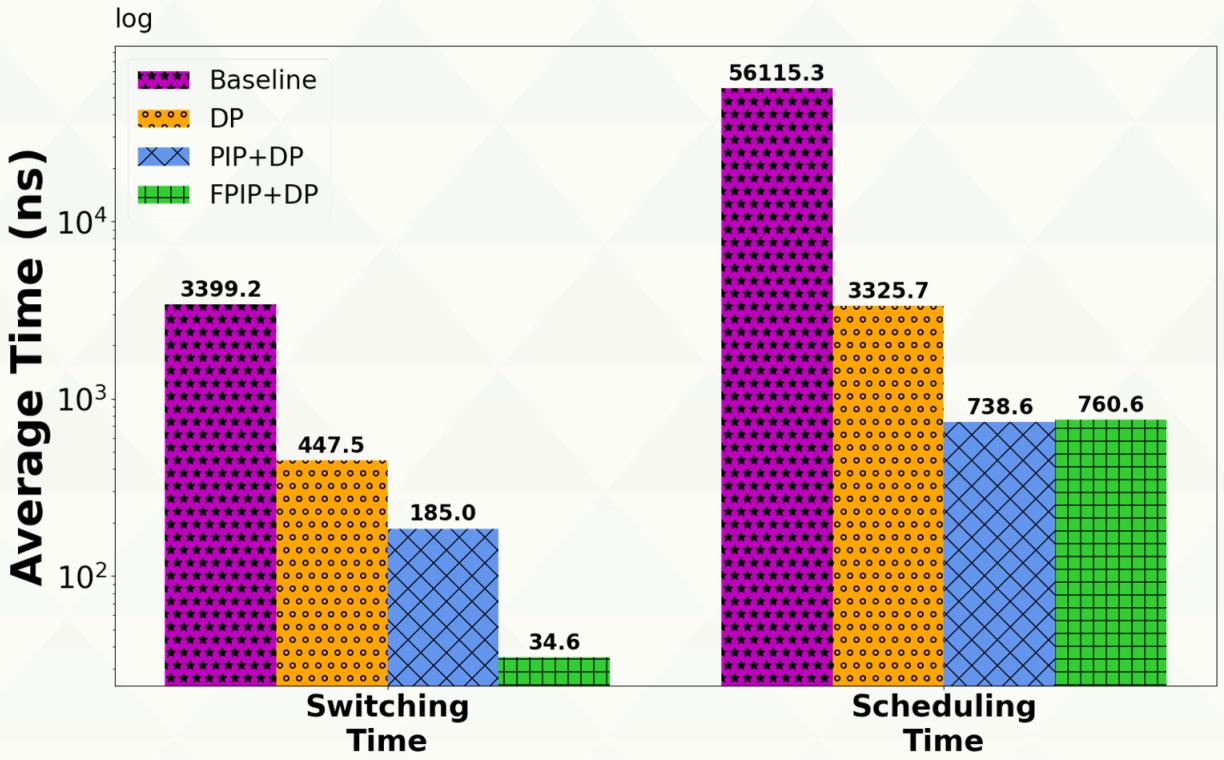

Scheduling Time:

- Cluster Search, Placement/Routing Decisions
- 75.97x scheduling time improvement
- High eviction cost for FPIP+DP

Average Execution Time Per Tree

Execution Approach	Execution Time (us)	Speedup w.r.t *Baseline*	Speedup w.r.t *DP*
Baseline	431.622	1x	N/A
DP	9.943	43.40x	1x
PIP+DP	3.865	111.67x	2.57x
FPIP+DP	3.247	132.92x	3.06x

- 132.92x execution time speedup
- Significant scheduling overhead

Systems with rapid switching capabilities need efficient scheduling strategies to maximize the throughput

Conclusion and Future Work

- Spectrum sensing requires high throughput and instantaneous detection capability to handle highly dynamic spectrum environments.
- We present K-PACT, a planning tool that assigns independent kernels to PEs with separate IMEM banks, enabling fast, conflict-free context switching for kernels that do not execute concurrently.
- Simulation results show that K-PACT reduces the
  - ***instruction load time by 191.53×***
  - ***overall execution time of a subband by 132.92×***

- This approach can be generalized beyond spectrum sensing, meeting the needs of multi-context and rapidly changing workloads.
- As future work, we aim to design routing-aware clustering and placement algorithms with efficient scheduling heuristics to further enhance performance through faster context switching.

Thank You!

UA - RCL Webpage

Contact

- H. Umut Suluhan - suluhan@arizona.edu
- Serhan Gener – gener@arizona.edu
- Ali Akoglu – akoglu@arizona.edu

Questions?