Enabling Seamless Execution on Hybrid CPU/FPGA Systems: Challenges & Directions

Meena Belwal, Madhura Purnaprajna, Sudarshan TSB
Department of Computer Science and Engineering,
Amrita School of Engineering,
Bangalore Campus, India

AMRITA
VISHWA VIDYAPEETHAM
U N I V E R S I T Y
Established u/s - 3 of the UGC Act 1956
BENGALURU CAMPUS

FPGA

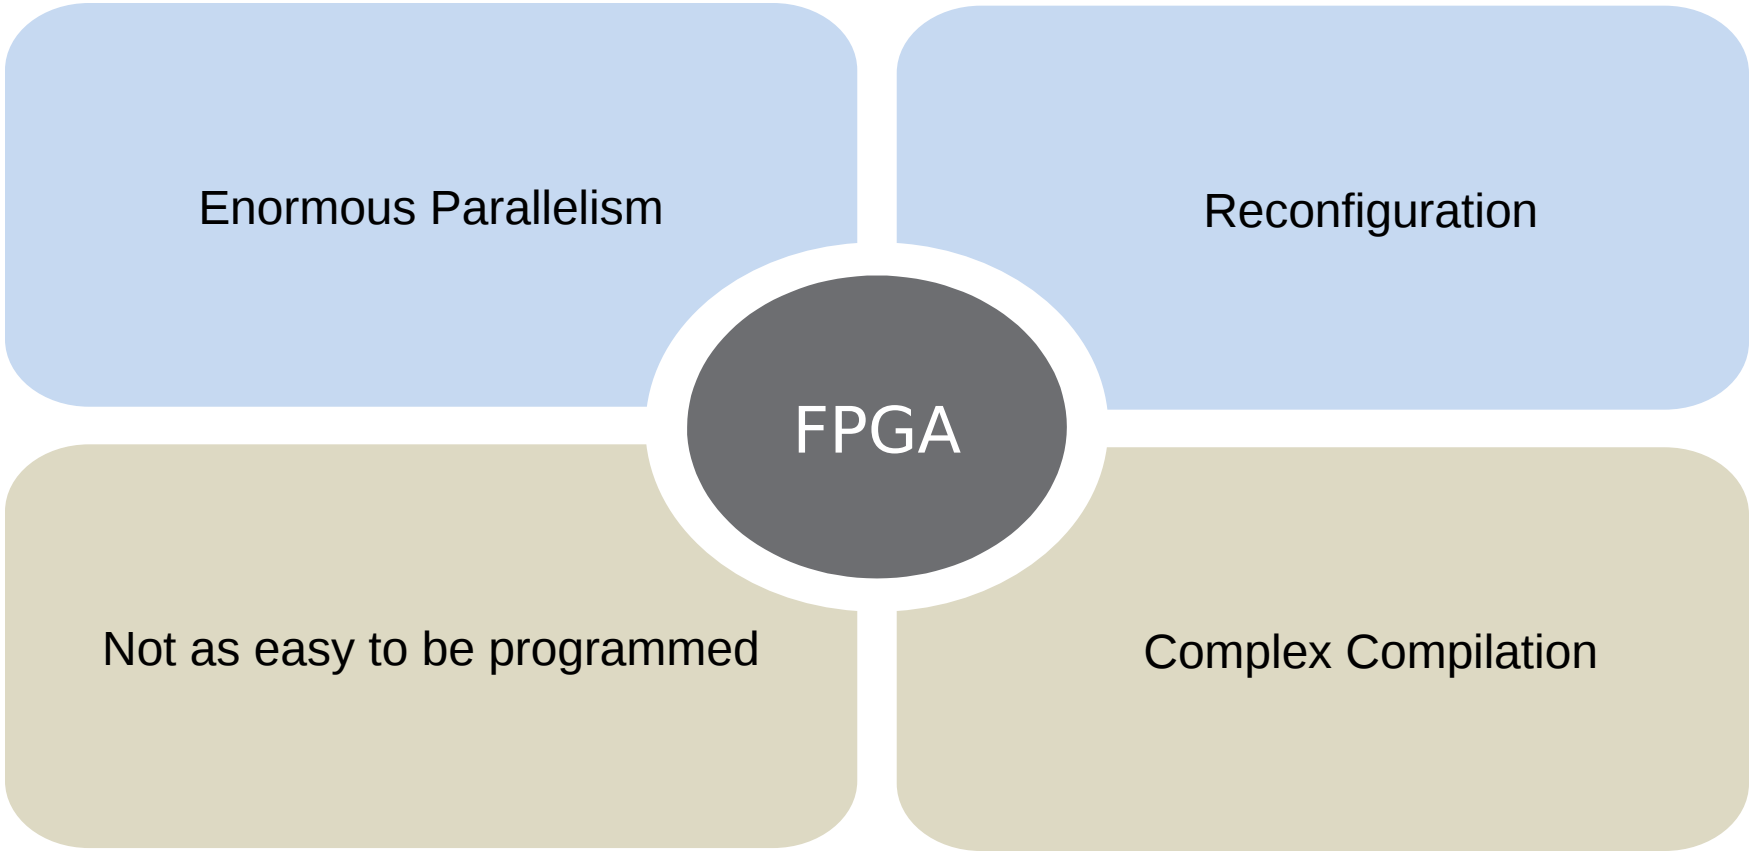

Applications defined for conventional processors cannot be used directly on FPGAs.

Hybrid(CPU+FPGA) Computing Bottleneck

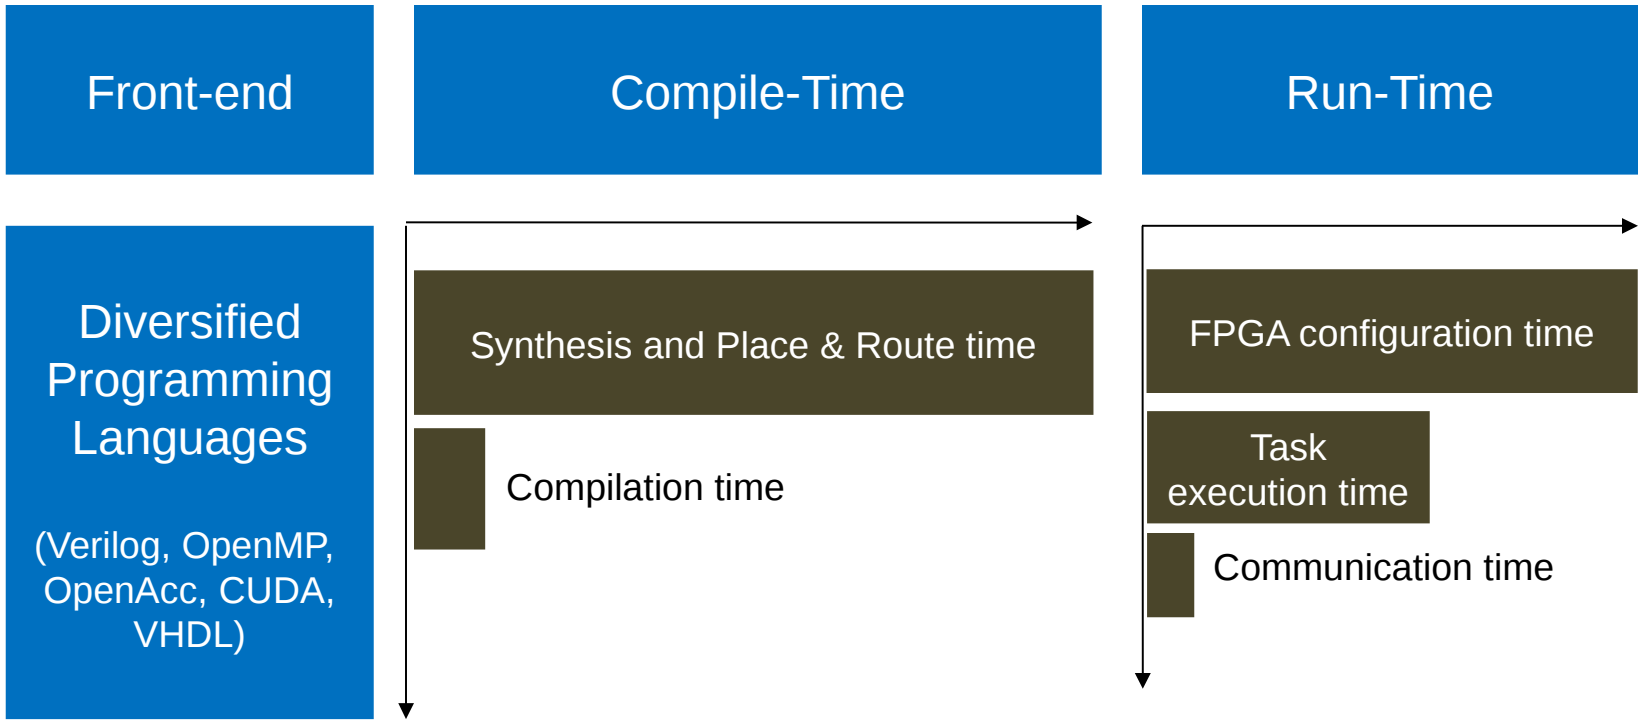

Ideally, dynamic automatic code partitioning and application mapping can simplify the process of introducing FPGAs in legacy CPU-only systems.

Contents

Design Approaches

श्रद्धावान् लभते ज्ञानम्

FPGA-only model

Processor-only model

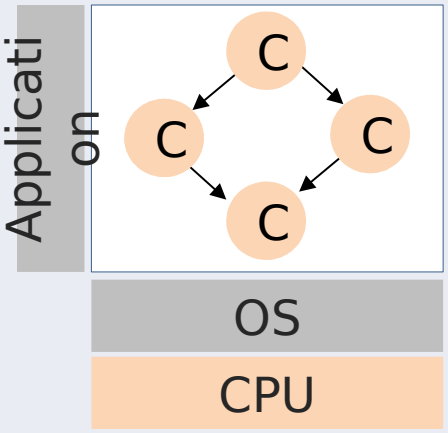

Processor + FPGA (offload) model

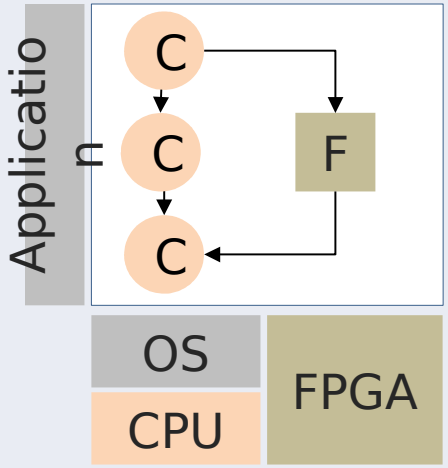

Processor + FPGA (Unified OS) model

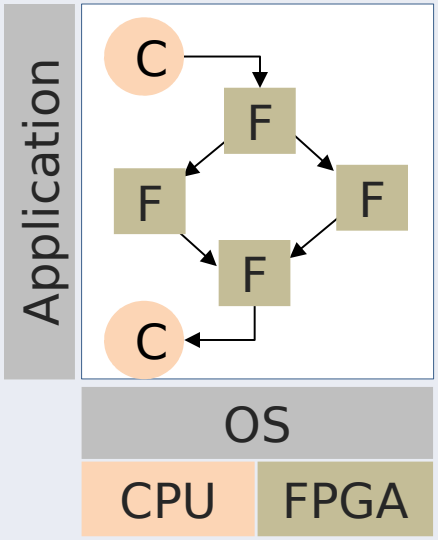

Design Approaches

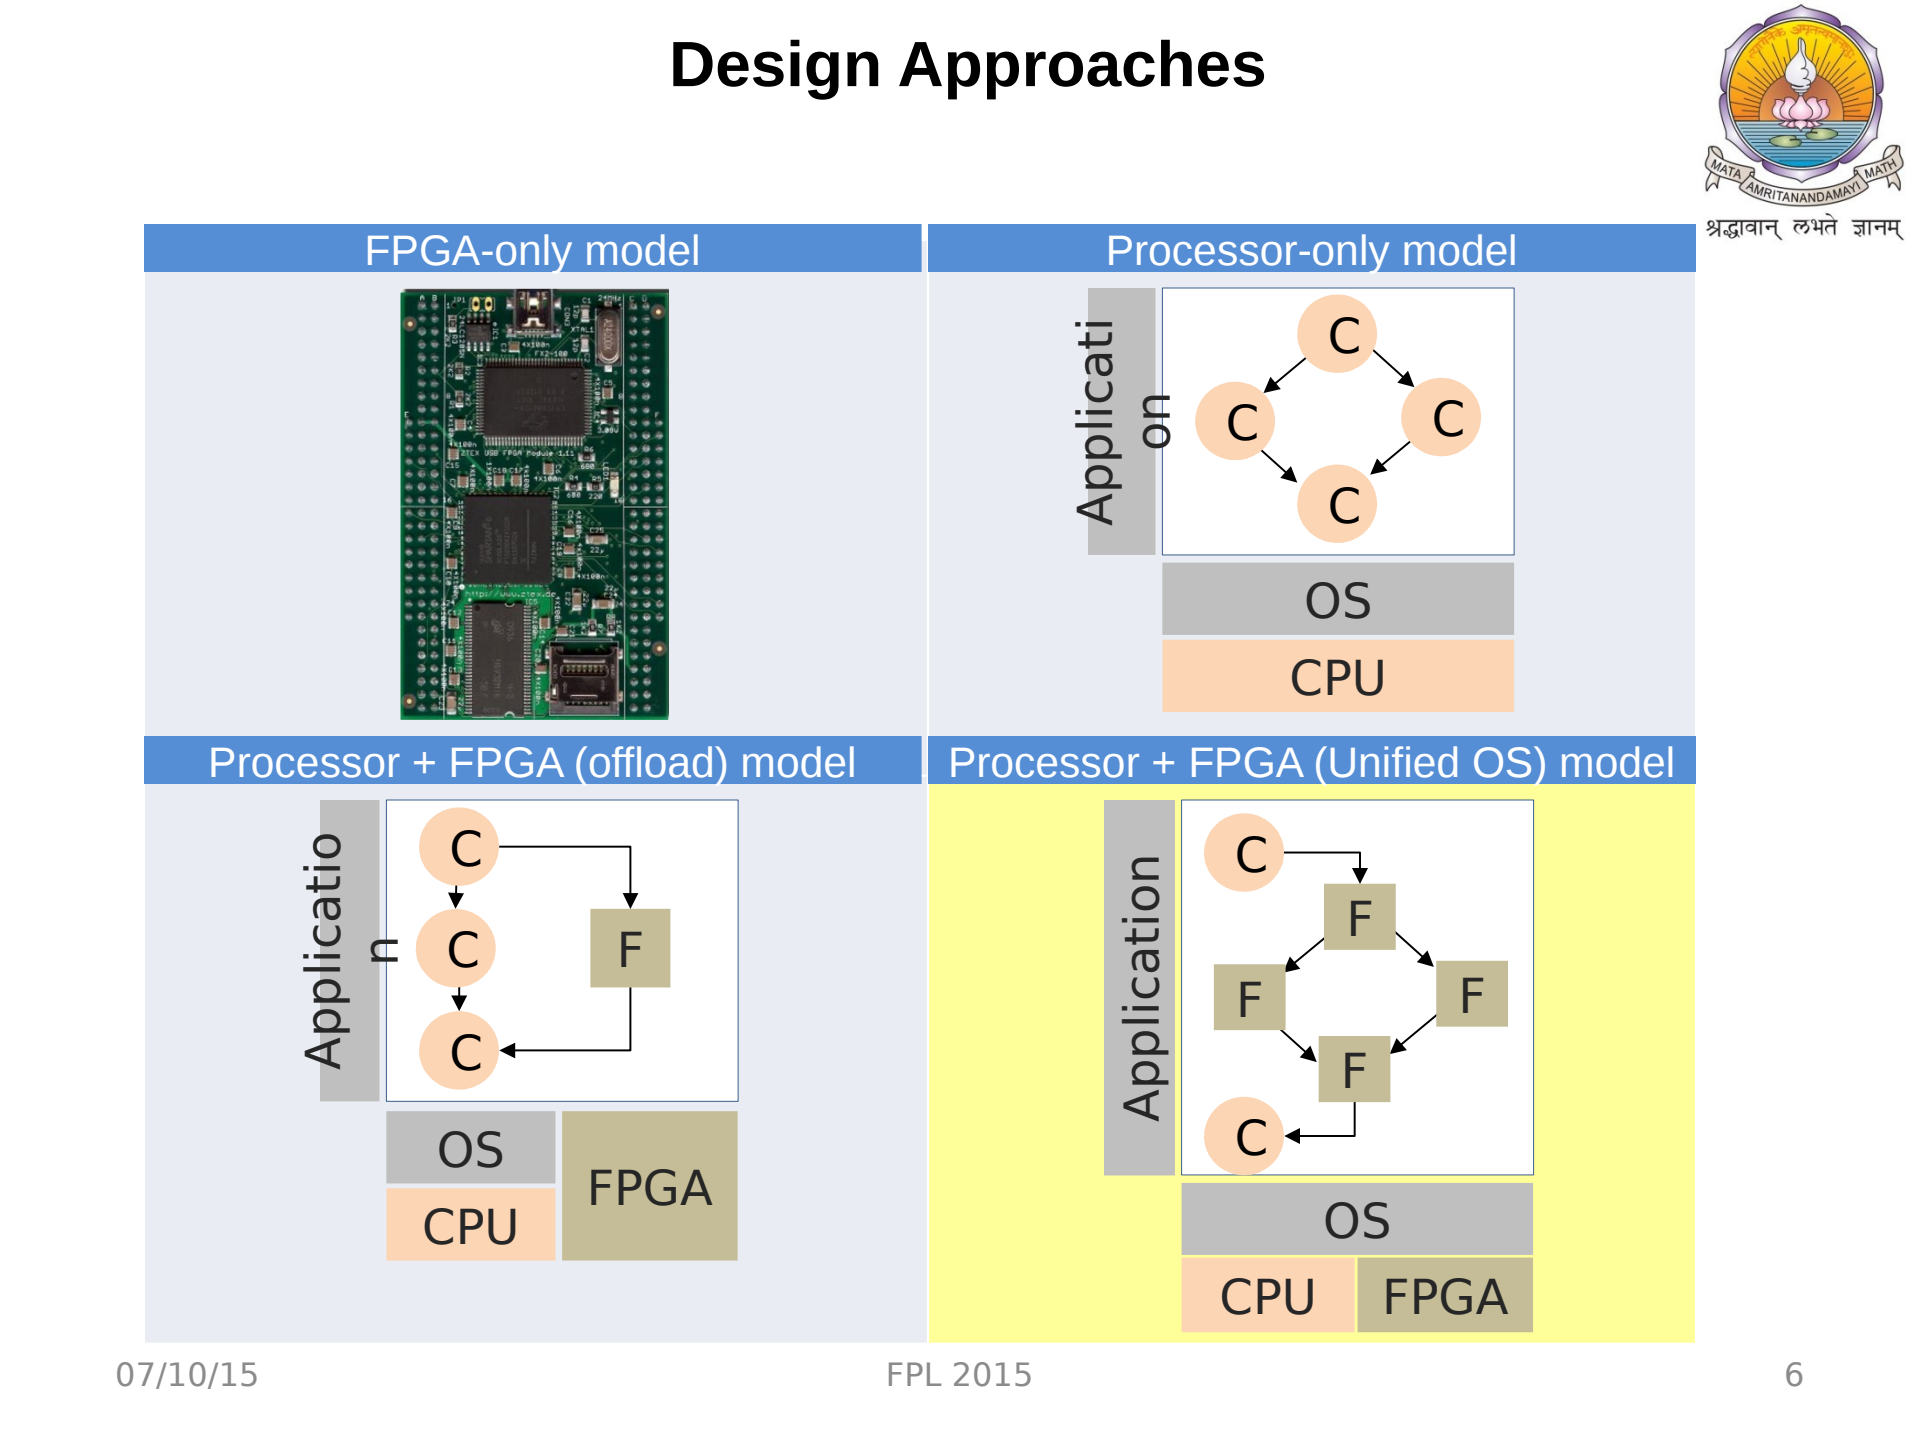

Processor + FPGA (Unified OS Model)

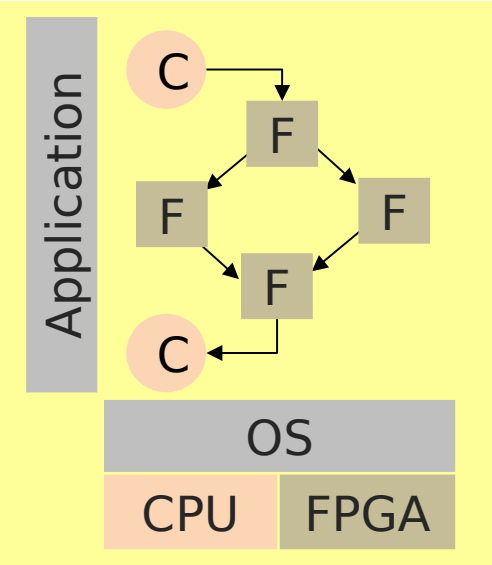

- CPU, FPGA at par
- Static techniques are application specific that leads to non-portability, and limits the designer's productivity.

To tap the full potential of CPU/FPGA hybrid systems a **computational model** with modern **operating services** is needed that hides platform specific CPU/FPGA distinction from the programmers.

Programming Model	OS and Run-time	Overhead Minimization

Programming Models and Hardware Compilation

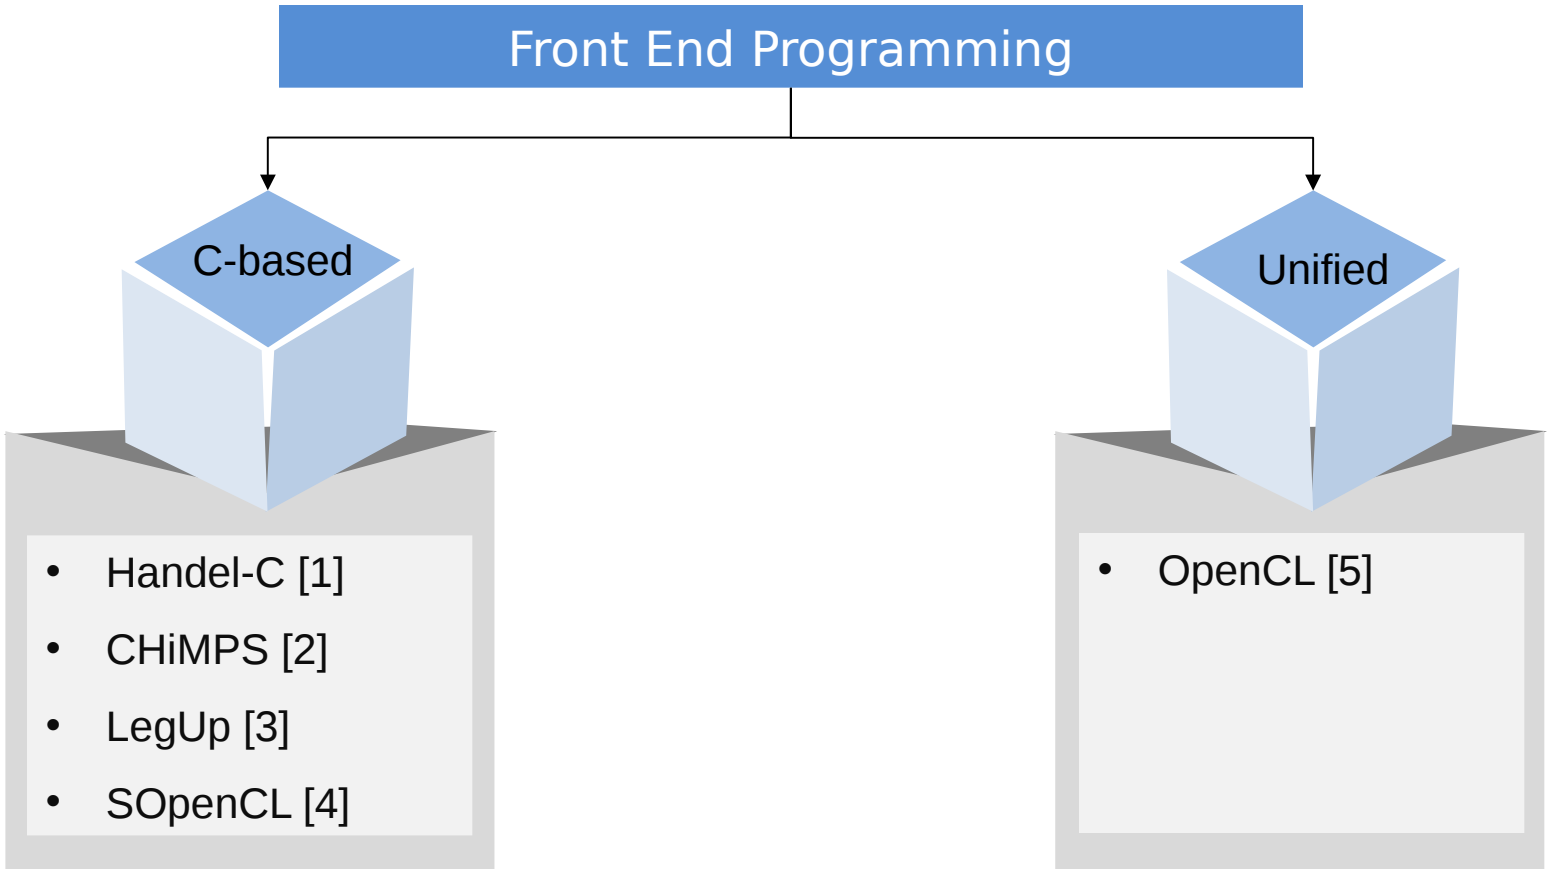

- Static Code Compilation
- Source-to-Source Translation

OS and Run-Time Systems

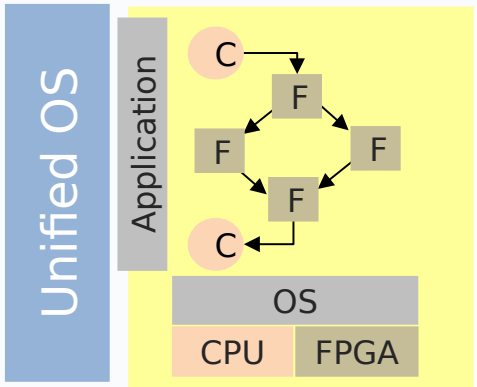

- Hthreads [6]
- ReconOS [7]
- Rainbow OS [8]

+

- Uniform APIs
- Application profiling
- Multitasking
- Dynamic Frequency Scaling

−

- Overhead and jitter
- Static configuration of HW threads
- Application specific

Overhead Minimization

Soft processors implemented on FPGA fabric

- Ease FPGA usability (Programmed using HLL)

GPU Overlay [9]	Floating and vector operations
VDR Overlay [10]	DFG (Data Flow Graphs) mapped to FPGA
Other Overlays	IF[11], FSM[12]

Overhead Minimization

Overlays

Accelerated Backend

Acceleration of synthesis and place and route process

Megablocks[13]	RPU generated statically
Warp Processors[14]	Dynamic synthesis, Lean P&R
Virtual FPGA[15]	Lean synthesis and P & R algorithms
HMFlow[16]	Hard macros
BPR[17]	Functional reuse
qFlow[18]	Split design (invariant & evolving logic set)

Future Research Directions

Alternatives to Overlays

- Modifying FPGA fabric to map soft processors [19]
- Easing code portability on FPGAs by making the reconfigurable fabric more processor-centric

Accelerating Hardware Compilation

Drawing Parallels from CPU/GPU Hybrid Systems

Future Research Directions

- Alternatives to Overlays
- Accelerating Hardware Compilation
- Drawing Parallels from CPU/GPU Hybrid Systems

- Lazy Man's logic synthesis (to perform logic synthesis offline) [20]
- The optimal structure for a Boolean function is retrieved from a precomputed library of structures.
- The library is created by an exhaustive synthesis of a large set of benchmarks

Future Research Directions

- Alternatives to Overlays
  - SnuCL
  - FluidCL
- Accelerating Hardware Compilation
  - Static Schemes: OpenMP, CUDA, Unified Memory
  - APU from AMD
- Drawing Parallels from CPU/GPU Hybrid Systems

Thank You