AHA: Agile Hardware-Compiler Co-design and Verification

Priyanka Raina
Assistant Professor of Electrical Engineering
Stanford University
praina@stanford.edu

August 30, 2023

AHA Research Thrusts

Domain-Specific Accelerator Architectures

AHA Research Thrusts

Domain-Specific Accelerator Architectures

Focus on programmability for an application domain

AHA Research Thrusts

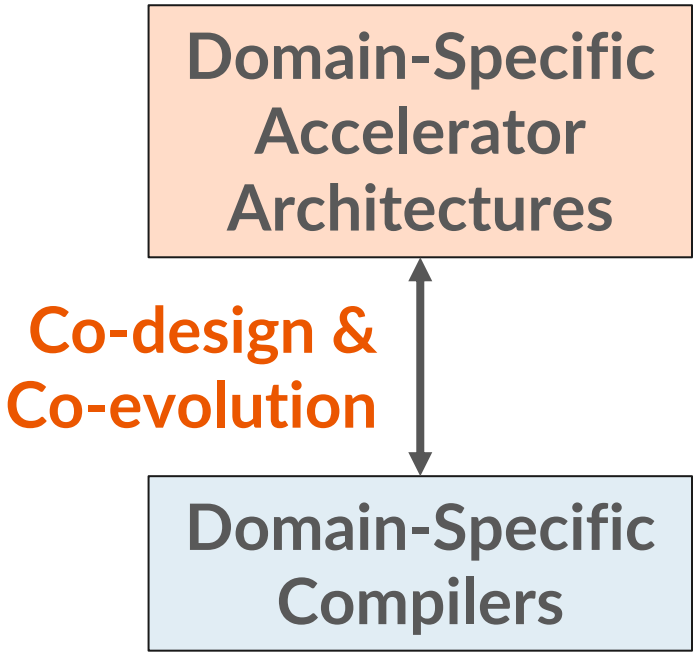

Focus on programmability for an application domain

AHA Research Thrusts

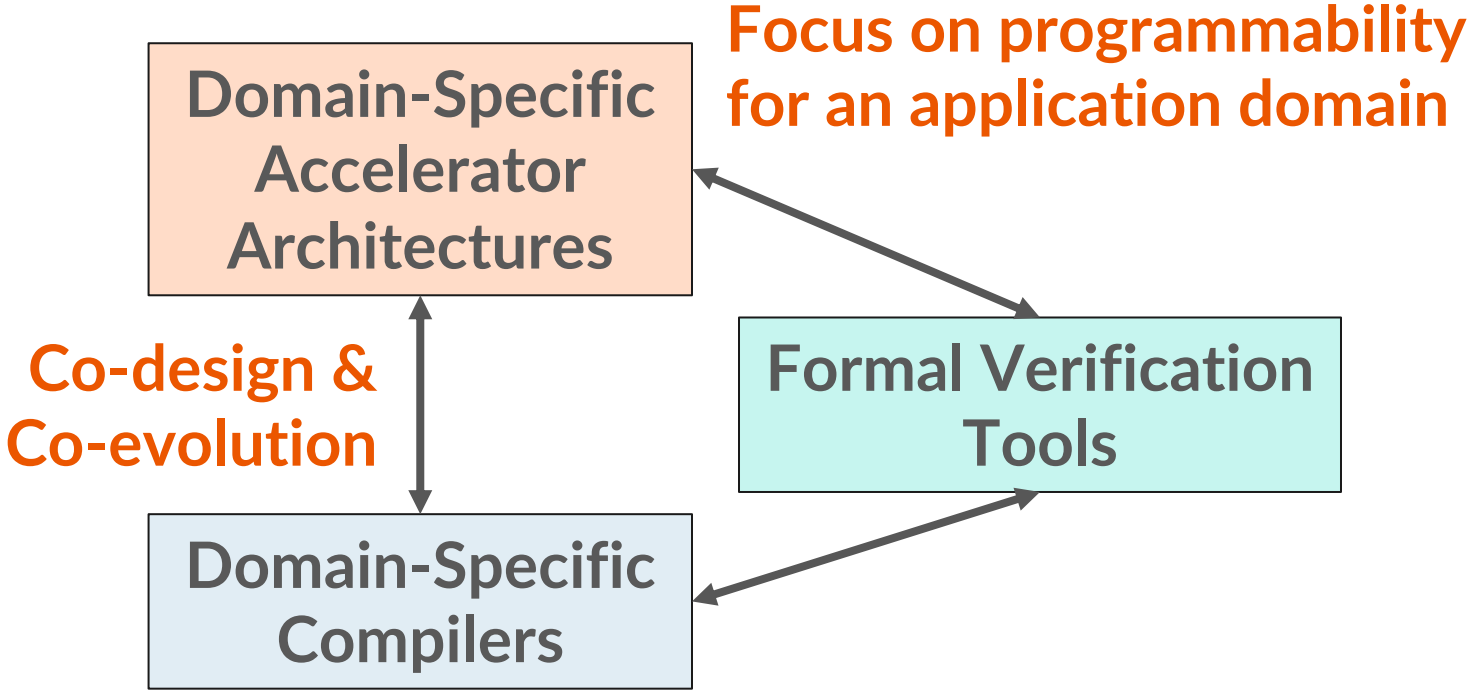

CGRAs as Accelerator Templates

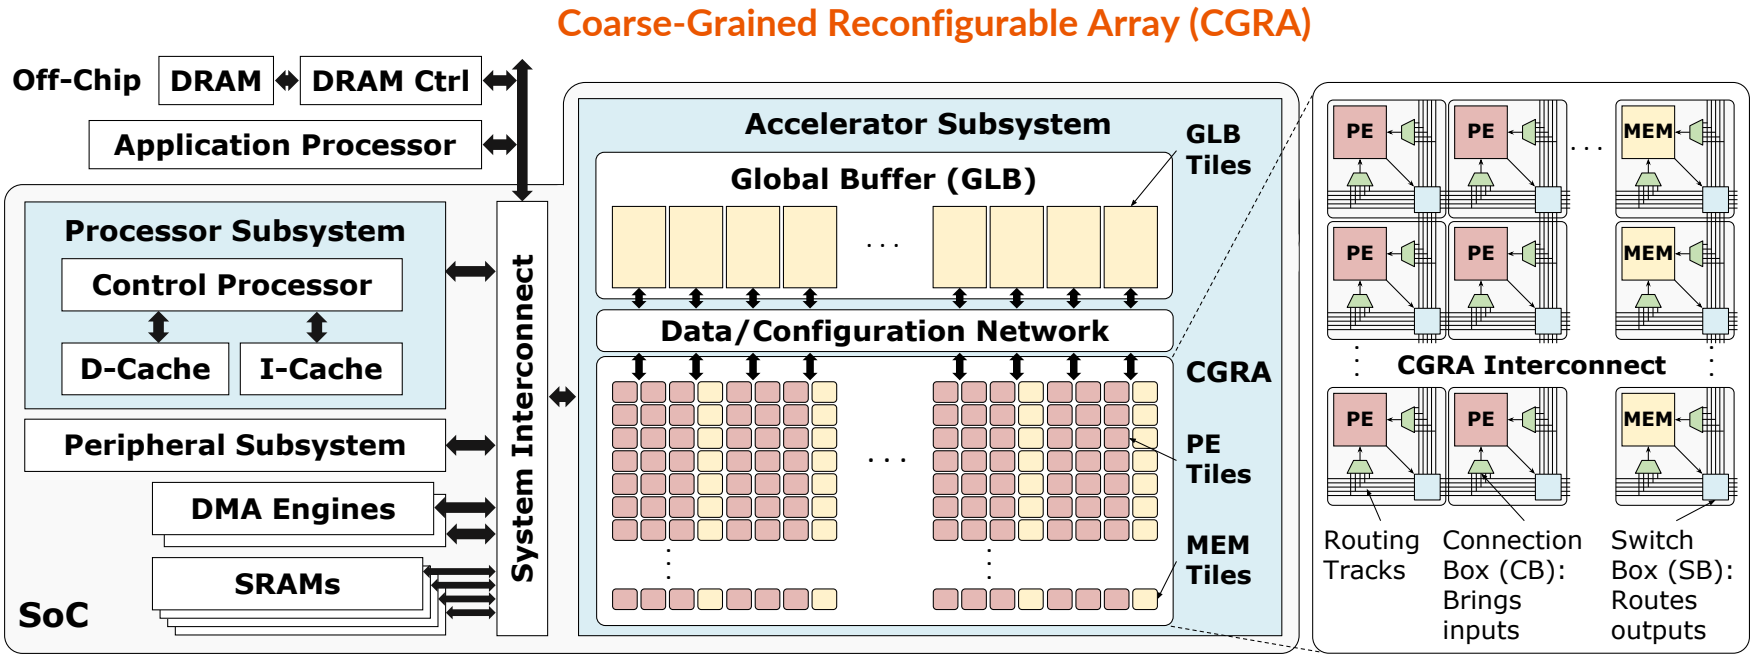

The Initial AHA System: Dense Imaging and ML

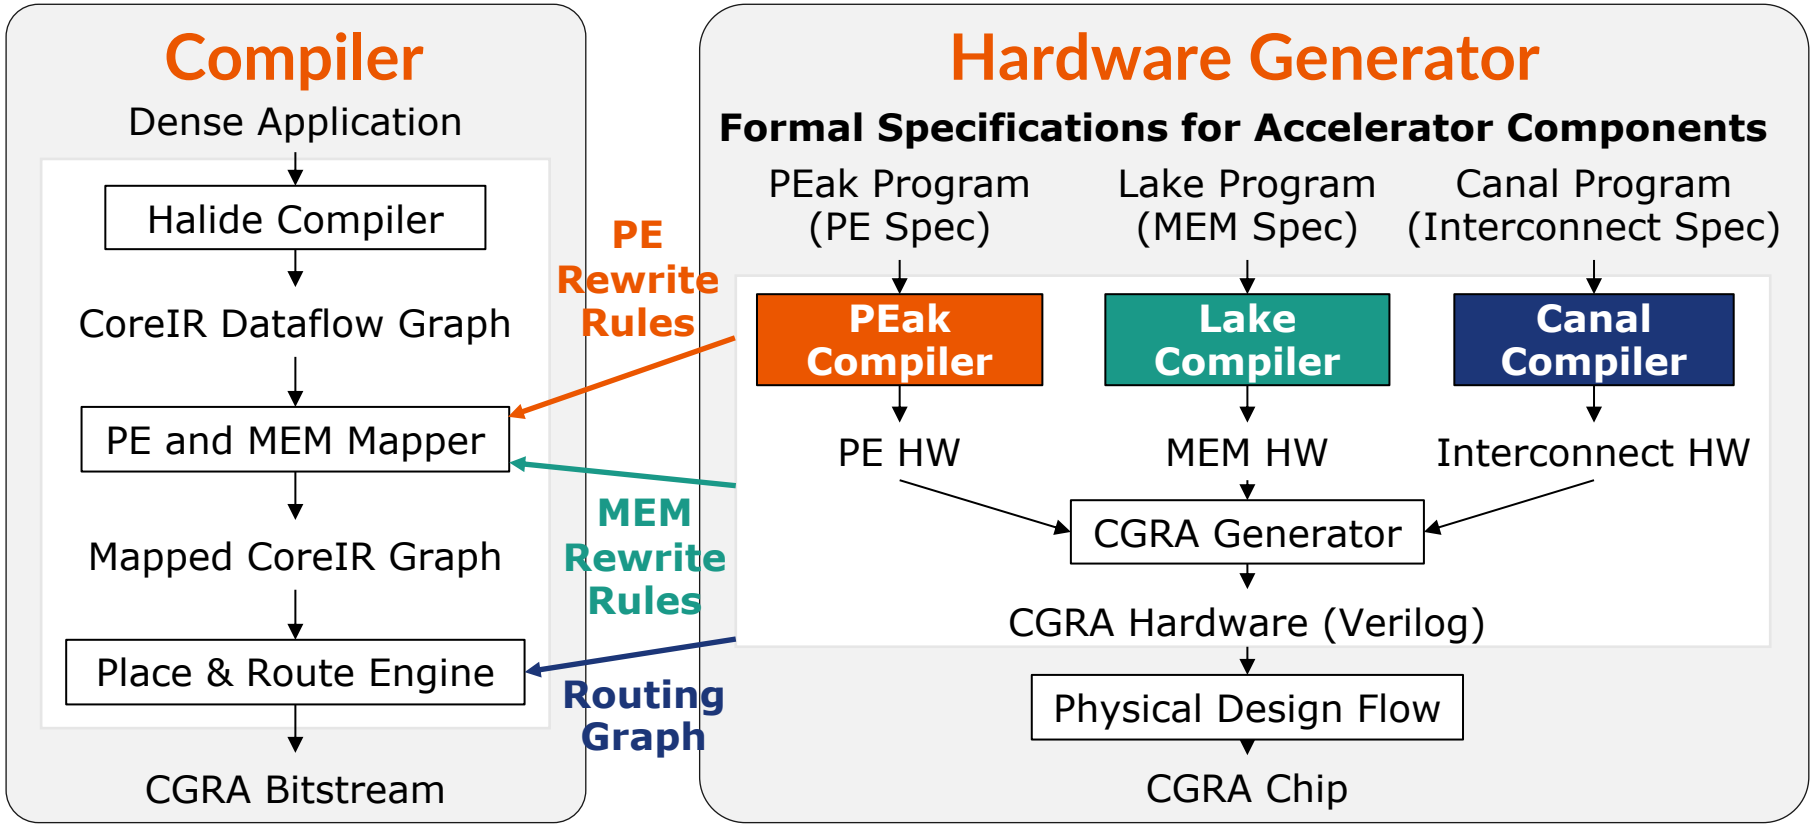

Extended to Sparse Applications

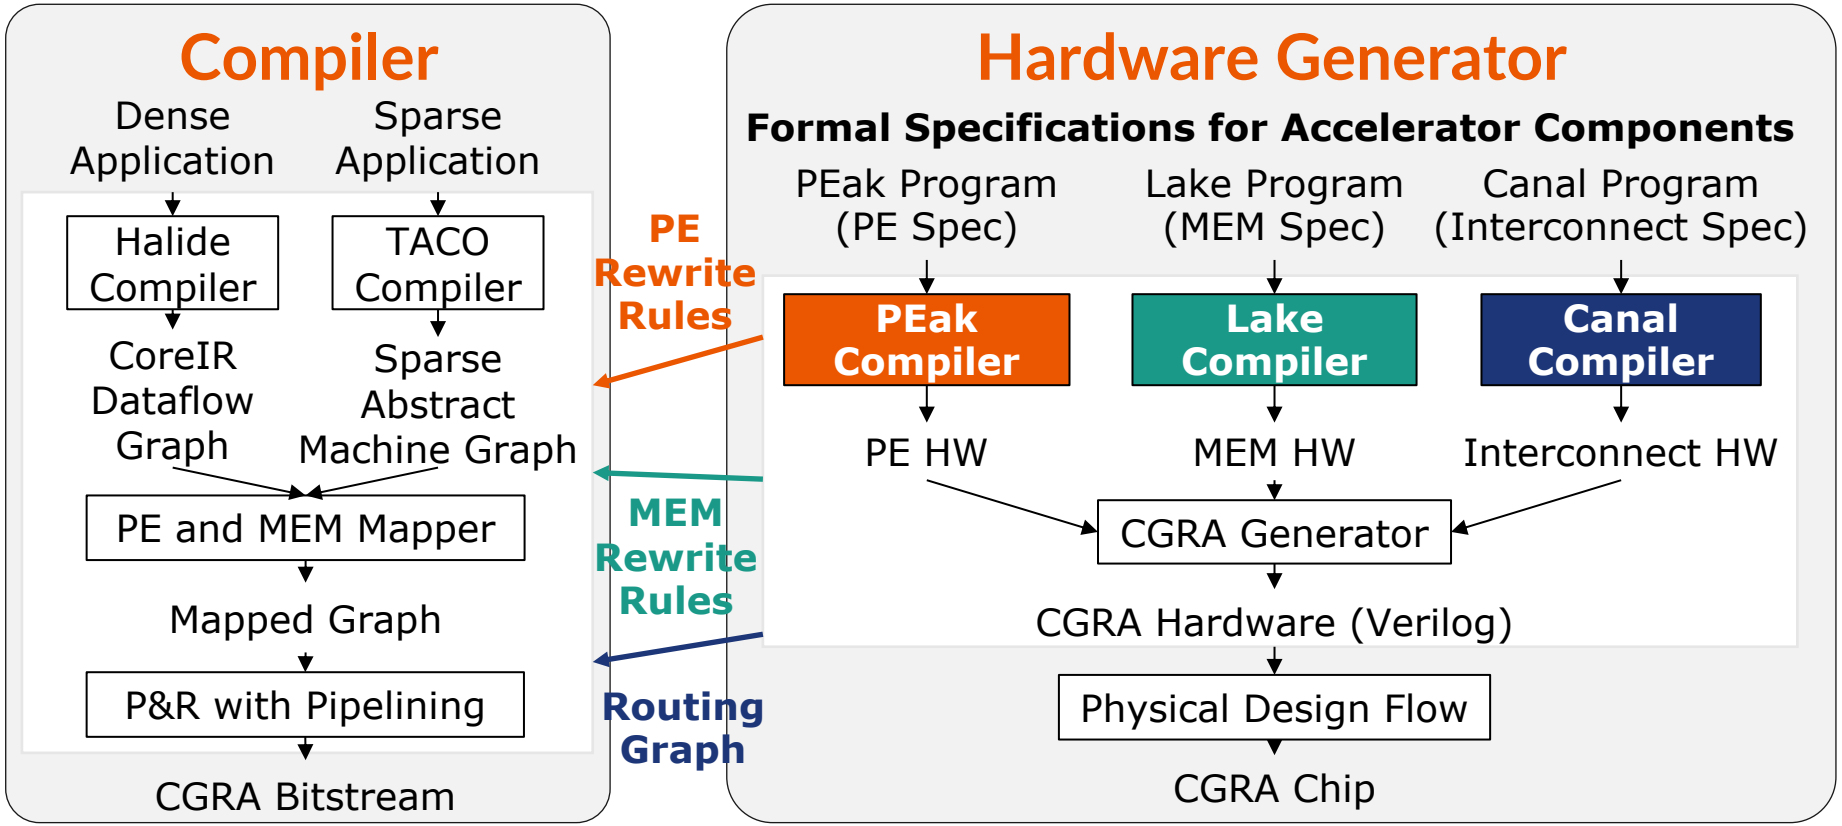

Accelerators Designed Using Our Agile Flow

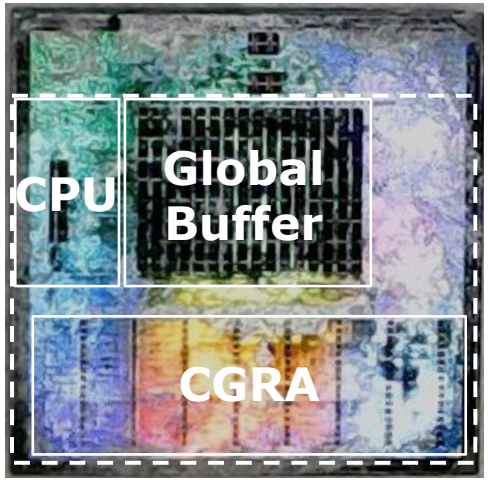

Amber SoC

TSMC 16

Statically scheduled dense data processing e.g. image processing and ML

VLSI 2022, JSSC 2023

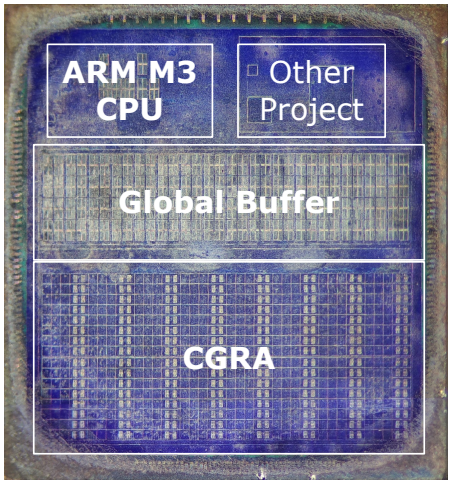

Onyx SoC

GF 12 (Taped out: Oct 2022)

+ Higher compute density
+ Improved pipelining
+ Sparse tensor algebra e.g. graph analytics

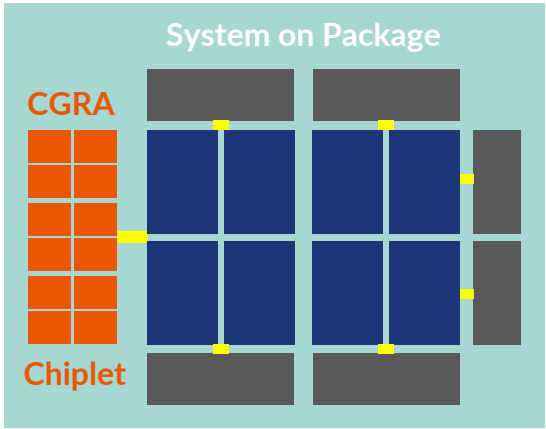

Opal SoC

Intel 16 (Tape out: Nov 2023)

+ Chiplet co-packaged with application processor
+ Sparse machine learning

Today's Posters

1: Onyx, 2: Onyx Evaluation
4: Lake Memory Generator
5: Dynamic Reconfiguration

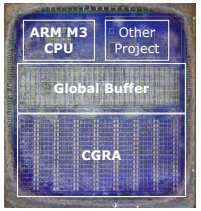

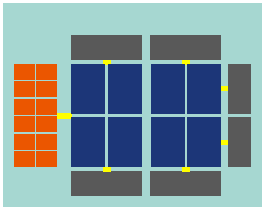

3: Opal

Domain-Specific Accelerator Architectures

Formal Verification Tools

Domain-Specific Compilers

Today's Posters

1: Onyx, 2: Onyx Evaluation
4: Lake Memory Generator
5: Dynamic Reconfiguration

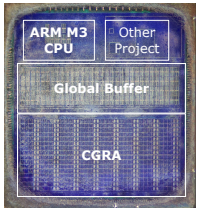

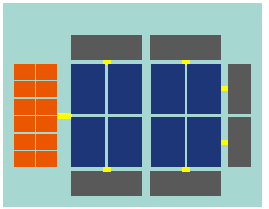

3: Opal

Domain-Specific Accelerator Architectures

Formal Verification Tools

Domain-Specific Compilers

6: Sparse PyTorch
7: Sparse Recurrences & Linear Solvers
8: Sparse ML Compiler
9: Sparse Tiling
10: Sparse Reshape

Today's Posters

1: Onyx, 2: Onyx Evaluation
4: Lake Memory Generator
5: Dynamic Reconfiguration

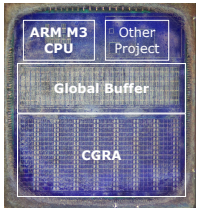

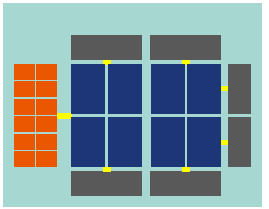

3: Opal

Domain-Specific Accelerator Architectures

Formal Verification Tools

Domain-Specific Compilers

11: Rewrite Rule Generation
12: Application Pipelining
13: CNN P&R
14: P&R Visualization

6: Sparse PyTorch
7: Sparse Recurrences & Linear Solvers
8: Sparse ML Compiler
9: Sparse Tiling
10: Sparse Reshape

Today's Posters

1: Onyx, 2: Onyx Evaluation
4: Lake Memory Generator
5: Dynamic Reconfiguration

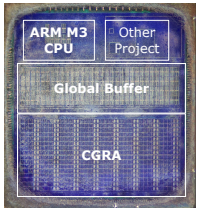

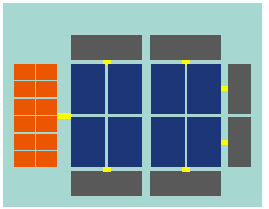

3: Opal

Domain-Specific Accelerator Architectures

11: Rewrite Rule Generation
12: Application Pipelining
13: CNN P&R
14: P&R Visualization

Formal Verification Tools

16: Formal Translation Validation
17: Online Learning for Bounded Model Checking
18: Optimization Modulo Theories

6: Sparse PyTorch
7: Sparse Recurrences & Linear Solvers
8: Sparse ML Compiler
9: Sparse Tiling
10: Sparse Reshape

Domain-Specific Compilers

Today's Posters

1: Onyx, 2: Onyx Evaluation
4: Lake Memory Generator
5: Dynamic Reconfiguration

19: Agile PD Tools
20: Energy-Aware DSE
21: Timing Bug Discovery

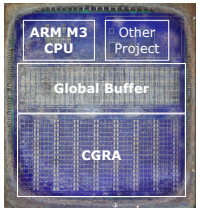

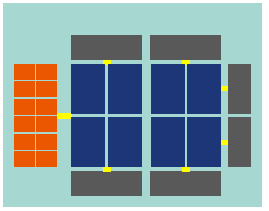

3: Opal

22: XR
23: Crypto
24: Viterbi for Links
25: ML Training

Domain-Specific Accelerator Architectures

Formal Verification Tools

Domain-Specific Compilers

11: Rewrite Rule Generation
12: Application Pipelining
13: CNN P&R
14: P&R Visualization

16: Formal Translation Validation
17: Online Learning for Bounded Model Checking
18: Optimization Modulo Theories

6: Sparse PyTorch
7: Sparse Recurrences & Linear Solvers
8: Sparse ML Compiler
9: Sparse Tiling
10: Sparse Reshape