

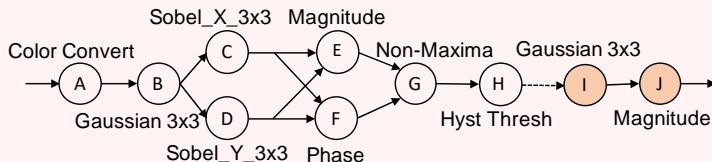
## Overview

- OpenVX (by Khronos Group, same that makes OpenCL)
  - C library and run-time system for Computer Vision
  - Portable, targets many hardware accelerators
  - Think: cache + pipeline friendly version of OpenCV
- OpenVX application is dynamic
  - Arbitrary graph of pipelined “kernels” or nodes
  - Application not known in advance, cannot use HLS
- Overlay for OpenVX (Processor + SVP + CVIs)
  - C language API run on ARM or MicroBlaze
  - Kernels accelerated on soft vector processor (SVP)
  - Many kernels synthesized in advance using Vivado HLS into a Custom Vector Instruction (CVI)
  - Multiple CVIs selected, merged and configured into one Partially Reconfigurable Region (PRR)
- Single-kernel performance example (Magnitude)
  - Baseline before acceleration with CVIs

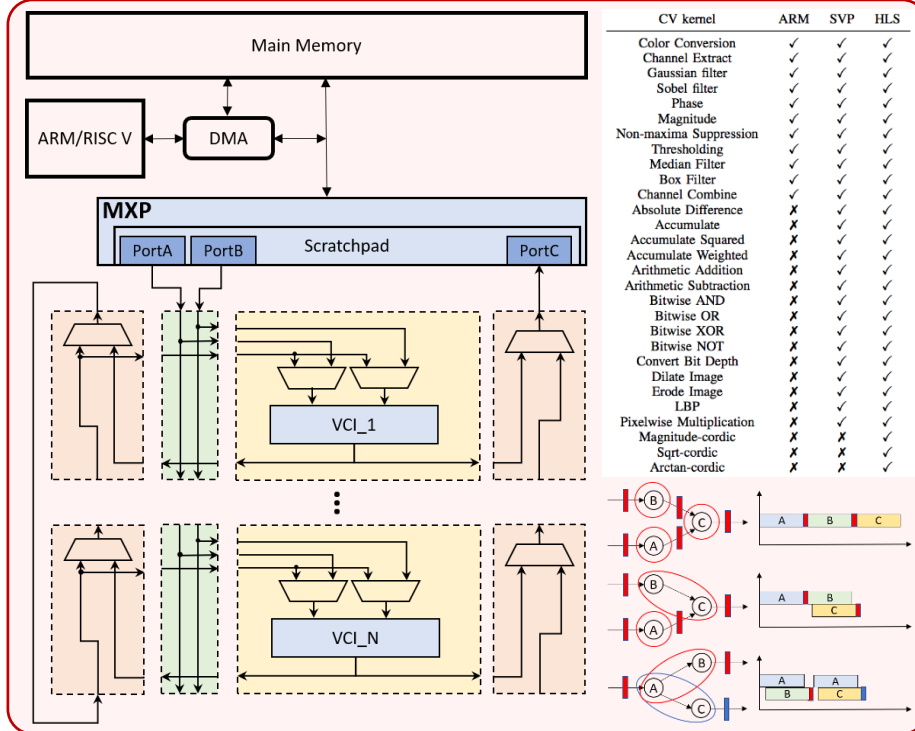
Running Platform	Throughput (MegaPixel/Sec)	Speedup vs ARM
ARM Cortex-A9 (667MHz)	10.31	1.0
VectorBlox SVP-V4 (100MHz / 12,989 LUTs)	65.54	6.3
VectorBlox SVP-V8 (100MHz / 22,517 LUTs)	128.92	12.5
Custom hardware (100MHz / 3,000 LUTs)	1176.04	114

## Example OpenVX Application Graph

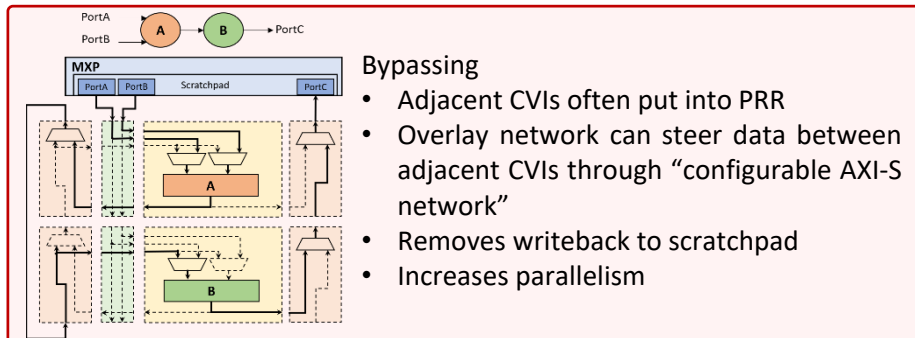
### Canny-blur Application



## System Architecture



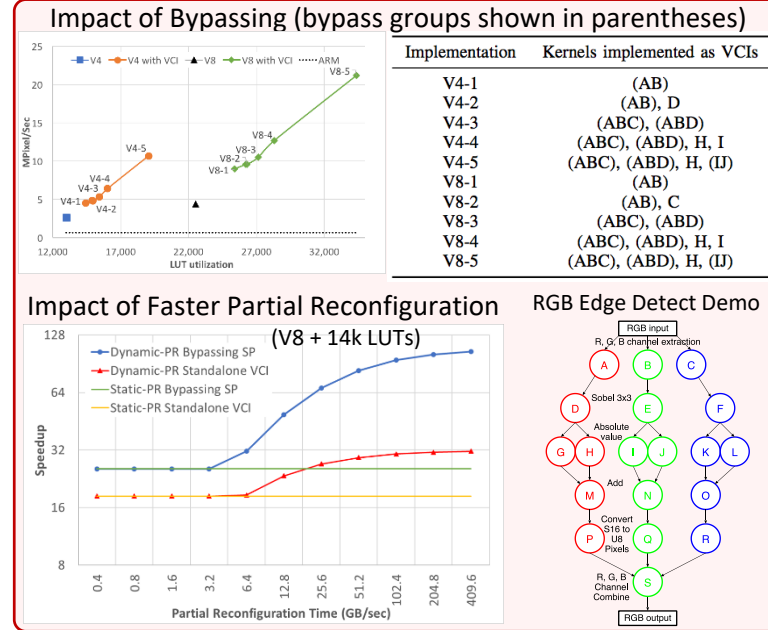
## Bypassing (Pipelining)



### Bypassing

- Adjacent CVIs often put into PRR
- Overlay network can steer data between adjacent CVIs through “configurable AXI-S network”
- Removes writeback to scratchpad
- Increases parallelism

## Performance (on Canny-blur)



## Conclusions/Future Work

- OpenVX can be implemented as a pure software API
- Basic acceleration using soft vector processor
- Further acceleration using custom vector instructions
  - Bypassing improves pipeline parallelism
  - Faster PR allows more CVIs when hardware limited
- OpenVX Overlay can be part of a larger CNN / ML Overlay

## Acknowledgements/References

- NSERC provided funding
- Xilinx, CMC Microsystems provided CAD tools
- VectorBlox provided soft vector processor

