

#### A Hierarchical Task Scheduler for Heterogeneous Computing

Architecture & Performance Group

25-May-2021

#### ISC High Performance (The HPC Event) June 24-July 2, 2021 Virtual

ORNL is managed by UT-Battelle, LLC for the US Department of Energy



ORNL is managed by UT-Battelle for the US Department of Energy

https://csmd.ornl.gov/profile/narasinga-rao-miniskar miniskarnr@ornl.gov

## About

- Future of HPC: Heterogenous Computing
- Task Scheduling
  - Balance application performance and programmability
  - Varying data communication time between tasks
  - Varying computation time across different variants of accelerators
- Proposal

- RANGER Platform: Hardware assisted task scheduling framework
- RISC-V cores with Accelerators
- Hierarchical Task Scheduler (TS)
  - Local level: Task into fine grained sub-tasks ; Accelerator Specific TS
  - Global level: Coarse grain task specification with programming portability; Coarse grain TS
- Extended GEM5 simulator for RANGER platform



# State of the art for Heterogenous Task Schedulers



CCM<sup>\*</sup> for each task on 3 devices

- Popular Static Task Scheduler: PEFT (Predict Earliest Finish Time)
  - Considers impact of current scheduling decision on subsequent decisions
- Task: Atomic unit for data transfer and computation
  - Can be portioned into finer sub-tasks
  - Smaller fine grain sub-tasks (streaming)
    - Lesser scratch pad memory
    - More scheduling opportunity but increases complexity
    - Less efficient data transfers due to large number of starting latencies
    - Accelerator friendly
    - Data transfers overlaps with computation <a href="#">#OAKI</a>



3 \* CCM: Cost Computation Matrix

# Vanilla Platform with Heterogenous Accelerators

- Centralized Heterogenous Scheduler: PEFT
- Complexity depends on the number of tasks
- Accelerators and DMAs are controlled through





### **RANGER Platform with Heterogenous Accelerators**

- Accelerator Interface: Memory-mapped IO
- RISC-V for host and accelerator devices

Device

**RISC-V (ASCS Scheduler)** 

MEMCTRL

SRAM (Buffers)

Ctrl Data Ctrl Data

DMA

DMA

Out



Memory Mapped IO

Ctrl

Accelerator

Wide bus

L: Task initialization time C<sub>i</sub>: Configuration of DMAs & Accelerator

# **Hierarchical Task Scheduler Flow**

- Why Accelerator Specific?
  - Configurable Tile size
  - Configurable Data reuse
  - Should support multiple accelerators
  - Programmable is a best solution



#### **Experimental Setup**

• Simulator: GEM5

7



- Benchmarks: DNNs (Inception-V3, ResNet-50, UNet, VGG16)
- Accelerators: Convolution (CONV), Batch Normalization (BN), Dense

|    | Accelerator | Functionality              | MAC Units | SRAM Size (KB) | Area (mm <sup>2</sup> ) |
|----|-------------|----------------------------|-----------|----------------|-------------------------|
| 1  | CONV1024    | 2D Convolution             | 1024      | 256.0          | 1.81                    |
| 2  | CONV512     | 2D Convolution             | 512       | 256.0          | 1.28                    |
| 3  | CONV256     | 2D Convolution             | 256       | 256.0          | 1.01                    |
| 4  | CONV128     | 2D Convolution             | 128       | 256.0          | 0.88                    |
| 5  | CONV64      | 2D Convolution             | 64        | 128.0          | 0.59                    |
| 6  | BN1024      | Batch Normilization        | 1024      | 8.0            | 1.09                    |
| 7  | BN512       | Batch Normilization        | 512       | 4.0            | 0.55                    |
| 8  | BN256       | Batch Normilization        | 256       | 2.0            | 0.28                    |
| 9  | BN128       | <b>Batch Normilization</b> | 128       | 1.0            | 0.14                    |
| 10 | BN64        | <b>Batch Normilization</b> | 64        | 0.5            | 0.08                    |
| 11 | DENSE1024   | Dense                      | 1024      | 128.0          | 1.30                    |
| 12 | DENSE512    | Dense                      | 512       | 128.0          | 0.76                    |
| 13 | DENSE256    | Dense                      | 256       | 128.0          | 0.50                    |
| 14 | DENSE128    | Dense                      | 128       | 128.0          | 0.37                    |
| 15 | DENSE64     | Dense                      | 64        | 128.0          | 0.30                    |

List of kernel accelerators and their area estimations in a TSMC 16nm technology

|           | Accel | erator | s    |     |    |       |      |         |     |    |       |     |     |     |       |     | Area mm <sup>2</sup> |          |          |
|-----------|-------|--------|------|-----|----|-------|------|---------|-----|----|-------|-----|-----|-----|-------|-----|----------------------|----------|----------|
|           | 2D Co | onvolu | tion |     |    | Batch | Norm | ilizati | on  |    | Dense |     |     |     | Total |     |                      |          |          |
|           | 1024  | 512    | 256  | 128 | 64 | 1024  | 512  | 256     | 128 | 64 | 1024  | 512 | 256 | 128 | 64    |     | RANGER               | Baseline | Overhead |
| Design    |       |        |      |     |    |       |      |         |     |    |       |     |     |     |       |     |                      |          |          |
| Design 1  | 1     | 1      | 1    | 1   | 1  | 1     | 1    | 1       | 1   | 1  | 1     | 1   | 1   | 1   | 1     | 15  | 11.290               | 10.945   | 3.15 %   |
| Design 2  | 2     | 2      | 2    | 2   | 2  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 30  | 22.581               | 21.891   | 3.15 %   |
| Design 3  | 3     | 3      | 3    | 3   | 3  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 35  | 28.266               | 27.461   | 2.93 %   |
| Design 4  | 4     | 4      | 4    | 4   | 4  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 40  | 33.951               | 33.031   | 2.79 %   |
| Design 5  | 5     | 5      | 5    | 5   | 5  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 45  | 39.636               | 38.601   | 2.68 %   |
| Design 6  | 6     | 6      | 6    | 6   | 6  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 50  | 45.321               | 44.171   | 2.60 %   |
| Design 7  | 7     | 7      | 7    | 7   | 7  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 55  | 51.006               | 49.741   | 2.54 %   |
| Design 8  | 8     | 8      | 8    | 8   | 8  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 60  | 56.691               | 55.311   | 2.49 %   |
| Design 9  | 9     | 9      | 9    | 9   | 9  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 65  | 62.376               | 60.881   | 2.46 %   |
| Design 10 | 10    | 10     | 10   | 10  | 10 | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 70  | 68.061               | 66.451   | 2.42 %   |
| Design A  | 1     | 1      | 1    | 1   | 1  | 1     | 1    | 1       | 1   | 1  | 1     | 1   | 1   | 1   | 1     | 15  | 11.290               | 10.945   | 3.15 %   |
| Design B  | 2     | 2      | 2    | 2   | 2  | 2     | 2    | 2       | 2   | 2  | 2     | 2   | 2   | 2   | 2     | 30  | 22.581               | 21.891   | 3.15 %   |
| Design C  | 3     | 3      | 3    | 3   | 3  | 3     | 3    | 3       | 3   | 3  | 3     | 3   | 3   | 3   | 3     | 45  | 33.871               | 32.836   | 3.15 %   |
| Design D  | 4     | 4      | 4    | 4   | 4  | 4     | 4    | 4       | 4   | 4  | 4     | 4   | 4   | 4   | 4     | 60  | 45.161               | 43.781   | 3.15 %   |
| Design E  | 5     | 5      | 5    | 5   | 5  | 5     | 5    | 5       | 5   | 5  | 5     | 5   | 5   | 5   | 5     | 75  | 56.452               | 54.727   | 3.15 %   |
| Design F  | 6     | 6      | 6    | 6   | 6  | 6     | 6    | 6       | 6   | 6  | 6     | 6   | 6   | 6   | 6     | 90  | 67.742               | 65.672   | 3.15 %   |
| Design G  | 7     | 7      | 7    | 7   | 7  | 7     | 7    | 7       | 7   | 7  | 7     | 7   | 7   | 7   | 7     | 105 | 79.032               | 76.617   | 3.15 %   |
| Design H  | 8     | 8      | 8    | 8   | 8  | 8     | 8    | 8       | 8   | 8  | 8     | 8   | 8   | 8   | 8     | 120 | 90.323               | 87.563   | 3.15 %   |
| Average   |       |        |      |     |    |       |      |         |     |    |       |     |     |     |       |     |                      |          | 2.74 %   |

Various heterogeneous designs, the number of kernel accelerators, and the estimated area. Design A and B are the aliases of Design 1 and 2, respectively.

| Experimental Evaluation: Makespan |              |             |                |            |             |               |             |               | Model          |             |       |          |               |  |  |
|-----------------------------------|--------------|-------------|----------------|------------|-------------|---------------|-------------|---------------|----------------|-------------|-------|----------|---------------|--|--|
|                                   |              |             |                |            |             |               |             |               |                | ion-v3      | 189   | 20,469   | 108×          |  |  |
|                                   |              |             |                |            |             |               |             |               | Resnet         | t-50        | 107   | 6,824    | $64 \times$   |  |  |
|                                   |              |             |                |            |             |               |             |               | UNet           |             | 17    | 12,372   | $728 \times$  |  |  |
|                                   |              |             |                |            |             |               |             |               | VGG1           | 6           | 16    | 38,064   | 2,379×        |  |  |
|                                   | Makespan     |             |                |            |             |               |             |               |                |             |       |          |               |  |  |
| Model                             | Inception-v3 | 3           |                | Resnet-50  |             |               | VGG16       |               |                | UNet        |       |          |               |  |  |
| Architecture<br>Design            | RANGER       | Baseline    | Speedup        | RANGER     | Baseline    | Speedup       | RANGER      | Baseline      | Speedup        | RANGER      | Basel | ine      | Speedup       |  |  |
| Design 1                          | 94,788,333   | 762,320,311 | $8.04 \times$  | 88,346,488 | 504,482,689 | 5.71×         | 389,202,487 | 3,193,921,550 | 8.21×          | 336,496,490 | 1,233 | ,341,834 | 3.67×         |  |  |
| Design 2                          | 53,422,061   | 752,853,720 | $14.09 \times$ | 57,531,844 | 500,829,541 | $8.71 \times$ | 209,315,117 | 3,149,123,579 | $15.04 \times$ | 201,002,387 | 1,177 | ,865,323 | $5.86 \times$ |  |  |
| Design 3                          | 41,881,254   | 752,544,223 | $17.97 \times$ | 52,782,400 | 500,433,710 | $9.48 \times$ | 191,533,321 | 3,147,951,787 | $16.44 \times$ | 173,986,414 | 1,171 | ,882,097 | $6.74 \times$ |  |  |
| Design 4                          | 38,749,897   | 752,450,123 | $19.42 \times$ | 52,195,186 | 500,340,314 | $9.59 \times$ | 181,632,752 | 3,147,767,996 | $17.33 \times$ | 157,652,345 | 1,169 | ,892,512 | $7.42 \times$ |  |  |
| Design 5                          | 37,999,017   | 752,407,810 | $19.80 \times$ | 52,191,088 | 500,295,143 | $9.59 \times$ | 180,193,518 | 3,147,649,984 | $17.47 \times$ | 150,250,684 | 1,169 | ,590,055 | $7.78 \times$ |  |  |
| Design 6                          | 37,988,551   | 752,393,040 | $19.81 \times$ | 52,191,088 | 500,288,595 | $9.59 \times$ | 180,193,518 | 3,147,724,843 | $17.47 \times$ | 150,250,684 | 1,169 | ,578,215 | $7.78 \times$ |  |  |
| Design 7                          | 37,914,589   | 752,387,551 | $19.84 \times$ | 52,190,081 | 500,282,221 | $9.59 \times$ | 180,193,518 | 3,147,642,459 | $17.47 \times$ | 150,250,684 | 1,169 | ,575,296 | $7.78 \times$ |  |  |
| Design 8                          | 37,898,912   | 752,382,660 | $19.85 \times$ | 52,190,081 | 500,275,555 | $9.59 \times$ | 180,193,518 | 3,147,724,163 | $17.47 \times$ | 150,250,684 | 1,169 | ,574,480 | $7.78 \times$ |  |  |
| Design 9                          | 37,891,335   | 752,382,660 | $19.86 \times$ | 52,190,081 | 500,272,549 | $9.59 \times$ | 180,193,518 | 3,147,669,094 | $17.47 \times$ | 150,250,684 | 1,169 | ,574,480 | $7.78 \times$ |  |  |
| Design 10                         | 37,618,903   | 752,367,084 | $20.00 \times$ | 52,190,081 | 500,264,005 | 9.59×         | 178,630,640 | 3,147,724,117 | 17.62×         | 147,870,807 | 1,169 | ,574,480 | 7.91×         |  |  |
| Design A                          | 94,788,333   | 762,320,311 | $8.04 \times$  | 88,346,488 | 504,482,689 | $5.71 \times$ | 389,202,487 | 3,193,921,550 | $8.21 \times$  | 336,496,490 | 1,233 | ,341,834 | 3.67×         |  |  |
| Design B                          | 53,422,061   | 752,853,720 | $14.09 \times$ | 57,531,844 | 500,829,541 | $8.71 \times$ | 209,315,117 | 3,149,123,579 | $15.04 \times$ | 201,002,387 | 1,177 | ,865,323 | $5.86 \times$ |  |  |
| Design C                          | 42,241,318   | 752,541,894 | $17.82 \times$ | 52,447,643 | 500,420,375 | $9.54 \times$ | 174,801,184 | 3,144,474,922 | $17.99 \times$ | 173,986,414 | 1,171 | ,816,637 | $6.74 \times$ |  |  |
| Design D                          | 38,570,566   | 752,444,207 | $19.51 \times$ | 51,672,158 | 500,332,941 | $9.68 \times$ | 155,344,163 | 3,144,196,072 | $20.24 \times$ | 157,652,123 | 1,169 | ,797,015 | $7.42 \times$ |  |  |
| Design E                          | 37,793,805   | 752,406,134 | $19.91 \times$ | 51,362,009 | 500,254,081 | $9.74 \times$ | 153,596,950 | 3,144,134,148 | $20.47 \times$ | 150,247,240 | 1,169 | ,471,954 | $7.78 \times$ |  |  |
| Design F                          | 37,793,452   | 752,401,577 | $19.91 \times$ | 51,162,596 | 500,254,081 | $9.78 \times$ | 153,596,950 | 3,144,134,148 | $20.47 \times$ | 150,247,240 | 1,169 | ,470,736 | $7.78 \times$ |  |  |
| Design G                          | 37,708,203   | 752,386,686 | $19.95 \times$ | 51,162,596 | 500,254,081 | $9.78 \times$ | 153,596,950 | 3,144,134,148 | $20.47 \times$ | 150,247,240 | 1,169 | ,468,346 | $7.78 \times$ |  |  |
| Design H                          | 37,704,900   | 752,386,686 | 19.95×         | 51,161,299 | 500,254,081 | 9.78×         | 153,596,950 | 3,144,134,148 | 20.47×         | 150,247,240 | 1,169 | ,468,346 | 7.78×         |  |  |
| Average                           |              |             | 17.66×         |            |             | 9.10×         |             |               | 16.96×         |             |       |          | 6.96×         |  |  |

Architecture RANGER Baseline Increase

Comparison of makespans for various RANGER and baseline designs. On average, RANGER achieves a 12.7×speedup

#### **RANGER Speedup** (Fixed 10 parallel instances of each application)



Measured speedup of RANGER by running 10 parallel instances of each application with respect to Designs 1–10, which contain an increasing number of kernel accelerators. The speedup plateaus at Design 3 due to an insufficient number of tasks for the available kernel accelerators



Measured speedup of RANGER by running an increasing number of instances of the same application on Designs A–H. RANGER demonstrates excellent scalability with 100 instances of application running in parallel

### Area and Makespan Compared to Reference

- Reference: (Hypothetical) Identical number of Accelerators as its RANGER counterpart
- Has large scratchpad to hold complete task I/O
- Requires no local task scheduler

|                                 | Area mm <sup>2</sup> |           |                |
|---------------------------------|----------------------|-----------|----------------|
| Model<br>Architecture<br>Design | RANGER               | Reference | Difference     |
| Design 1                        | 11                   | 165       | 14.61×         |
| Design 2                        | 23                   | 275       | $12.18 \times$ |
| Design 3                        | 28                   | 337       | 11.93×         |
| Design 4                        | 34                   | 440       | 12.96×         |
| Design 5                        | 40                   | 445       | 11.23×         |
| Design 6                        | 45                   | 484       | $10.67 \times$ |
| Design 7                        | 51                   | 512       | $10.04 \times$ |
| Design 8                        | 57                   | 546       | 9.63×          |
| Design 9                        | 62                   | 603       | $9.67 \times$  |
| Design 10                       | 68                   | 662       | 9.73×          |
|                                 |                      |           |                |

2

|              | Makespan     |            |            |            |            |            |             |             |            |             |             |            |
|--------------|--------------|------------|------------|------------|------------|------------|-------------|-------------|------------|-------------|-------------|------------|
| Model        | Inception-v3 |            |            | Resnet-50  |            |            | VGG16       |             |            | UNet        |             |            |
| Architecture | RANGER       | Reference  | Difference | RANGER     | Reference  | Difference | RANGER      | Reference   | Difference | RANGER      | Reference   | Difference |
| Design       |              |            |            |            |            |            |             |             |            |             |             |            |
| Design 1     | 94,788,333   | 95,707,644 | -0.96%     | 88,346,488 | 88,621,038 | -0.31%     | 389,202,487 | 214,684,328 | 81.29%     | 336,496,490 | 216,796,889 | 55.21%     |
| Design 2     | 53,422,061   | 56,538,771 | -5.51%     | 57,531,844 | 63,717,300 | -9.71%     | 209,315,117 | 165,634,098 | 26.37%     | 201,002,387 | 156,806,696 | 28.18%     |
| Design 3     | 41,881,254   | 45,388,469 | -7.73%     | 52,782,400 | 58,538,562 | -9.83%     | 191,533,321 | 150,559,308 | 27.21%     | 173,986,414 | 145,184,318 | 19.84%     |
| Design 4     | 38,749,897   | 40,891,867 | -5.24%     | 52,195,186 | 56,076,426 | -6.92%     | 181,632,752 | 144,136,809 | 26.01%     | 157,652,345 | 134,955,240 | 16.82%     |
| Design 5     | 37,999,017   | 39,601,255 | -4.05%     | 52,191,088 | 55,193,132 | -5.44%     | 180,193,518 | 141,111,694 | 27.70%     | 150,250,684 | 129,117,096 | 16.37%     |
| Design 6     | 37,988,551   | 38,921,698 | -2.40%     | 52,191,088 | 54,871,507 | -4.88%     | 180,193,518 | 139,536,510 | 29.14%     | 150,250,684 | 128,893,489 | 16.57%     |
| Design 7     | 37,914,589   | 38,412,723 | -1.30%     | 52,190,081 | 54,847,500 | -4.85%     | 180,193,518 | 138,354,082 | 30.24%     | 150,250,684 | 127,945,051 | 17.43%     |
| Design 8     | 37,898,912   | 38,260,460 | -0.94%     | 52,190,081 | 54,561,879 | -4.35%     | 180,193,518 | 137,473,689 | 31.07%     | 150,250,684 | 127,915,851 | 17.46%     |
| Design 9     | 37,891,335   | 38,260,436 | -0.96%     | 52,190,081 | 54,561,879 | -4.35%     | 180,193,518 | 137,473,689 | 31.07%     | 150,250,684 | 127,915,851 | 17.46%     |
| Design 10    | 37,618,903   | 37,782,552 | -0.43%     | 52,190,081 | 53,711,730 | -2.83%     | 178,630,640 | 134,278,895 | 33.03%     | 147,870,807 | 121,184,411 | 22.02%     |
| Average      |              |            | -2.95%     |            |            | -5.06%     |             |             | 31.01%     |             |             | 20.53%     |

Comparison of makespans for RANGER and reference. On average, RANGER shows only 10.88% of penalty, which is the measurement of performance overhead of the local ASCS.

# Conclusion

- Presents RANGER framework (Extremely heterogenous computing)
  - An architecture design for hierarchical task scheduling
- Hierarchical Task Scheduling
  - Only requires coarse grained task dependency specification
  - Fine grain accelerator specific scheduling at lower level
- RANGER uses customized RISC-V cores
- Achieves 12.7x performance gain in terms of makespan
- Area Overhead: +2.7% in a 16nm technology



# Thank you

