Enabling Software Resilience in GPGPU Applications via Partial Thread Protection

Lishan Yang, Bin Nie, Adwait Jog, and Evgenia Smirni

William & Mary

GPUs & Soft Errors

Supercomputing

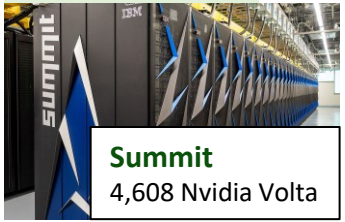

Deep learning

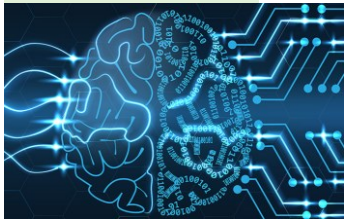

Self-driving cars

SDC → Misclassification

- GPUs are commonly deployed
- GPUs are prone to **soft errors**
  - High-energy radioactive particles (i.e., cosmic rays) cause **bit flips**
  - Commonly observed
  - Impact on long-running applications can be tremendous
    - ***Masked*** output: Correct
    - ***Crash, hang***, …
    - ***Silent Data Corruption (SDC)*** output: Incorrect
  - SDCs in critical applications can be dangerous
- Protection:
  - Error correction code (ECC) 5%~40% overhead
  - Software solution: re-computation
    - Detection: Duplication x**2** Computation
    - Correction: Triplication x**3** Computation

RQ: How to protect GPGPU applications **selectively**?

*Li, Guanpeng, et al. "Understanding error propagation in deep learning neural network (DNN) accelerators and applications." SC 2017.

GPGPU Application Parallelization

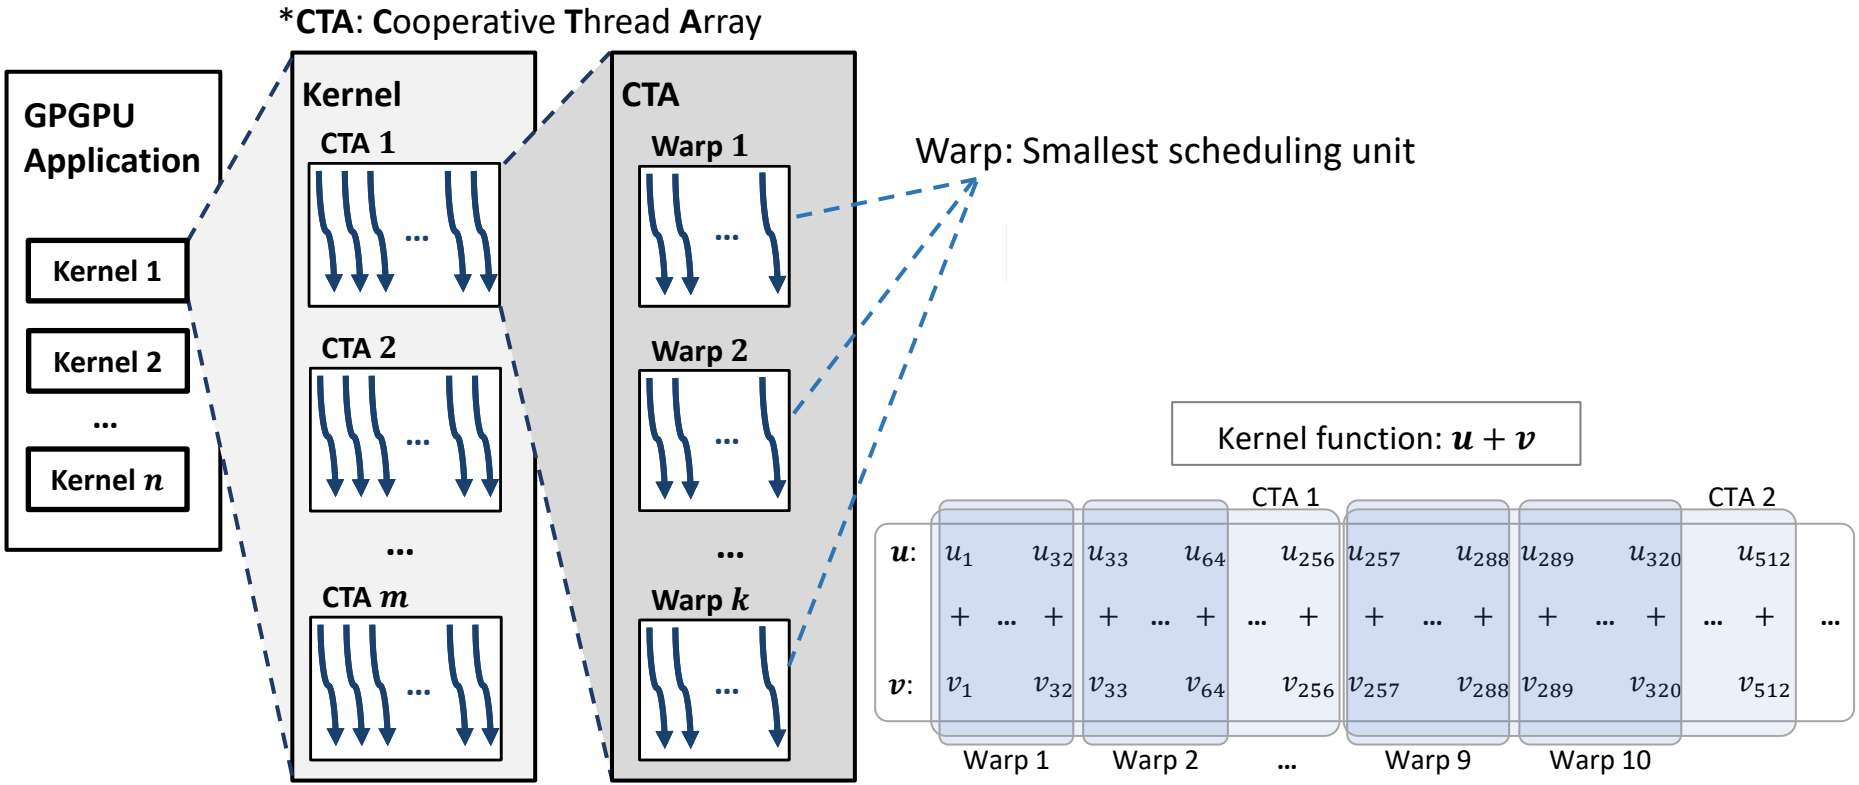

GPGPU Application Parallelization

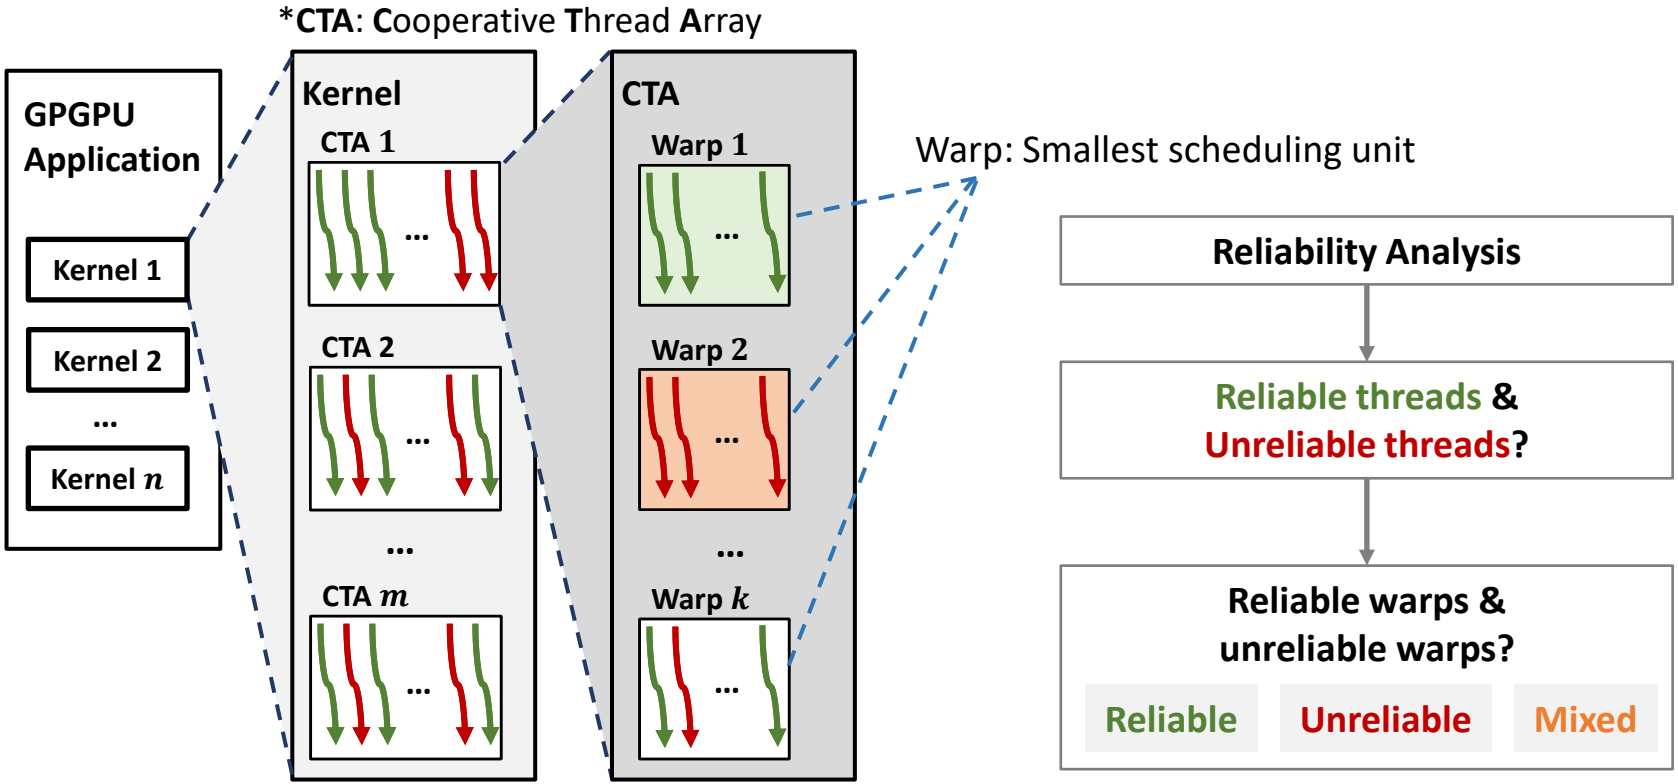

Characterization

- All threads are **reliable**

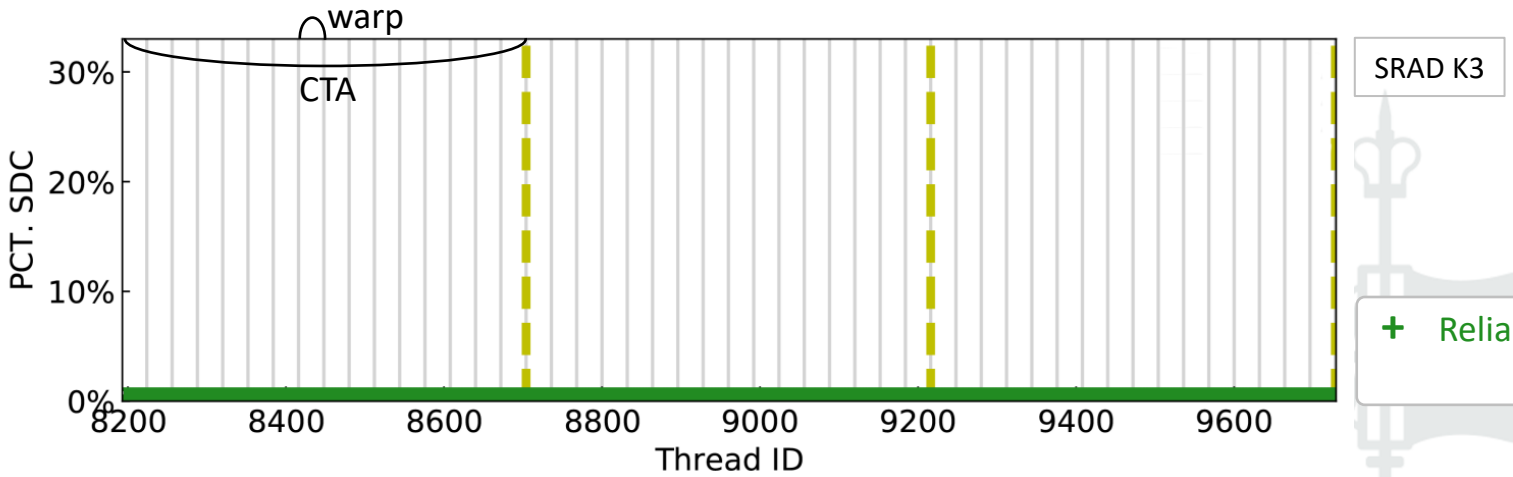

Representative benchmarks: *SRAD K3~K4, NeuralNetwork K1~K4*

Characterization

❖ All threads are **reliable**
(SRAD K3~K4, NeuralNetwork K1~K4)

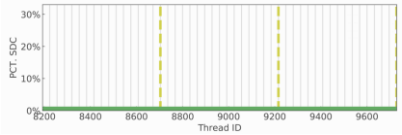

➢ All threads are **unreliable**

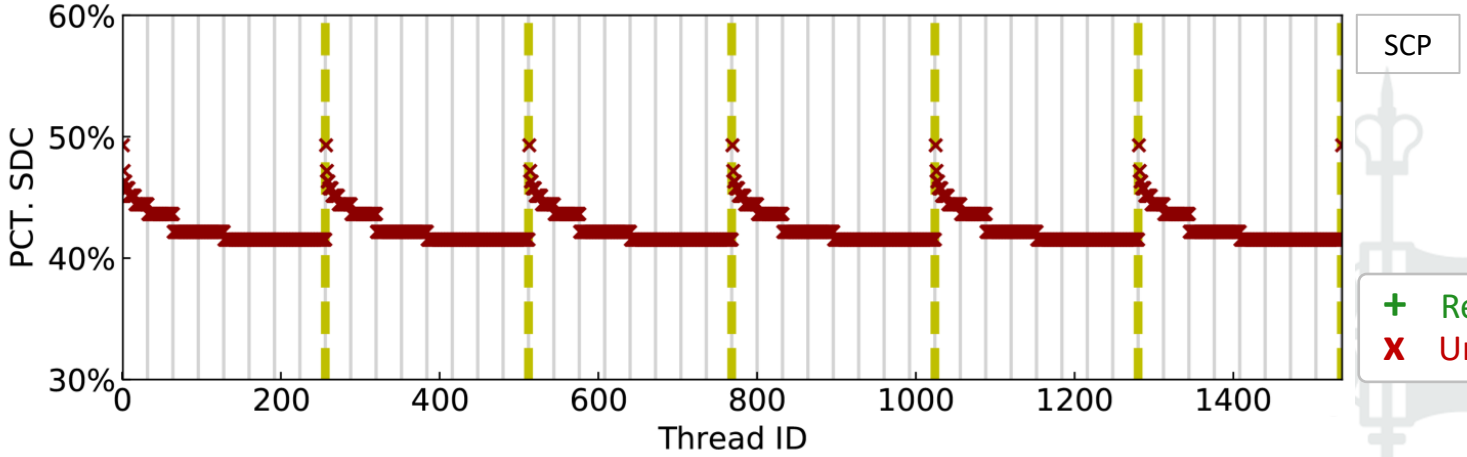

Representative benchmarks: *SCP, MVT*

Characterization

❖ All threads are **reliable**
(SRAD K3~K4, NeuralNetwork K1~K4)

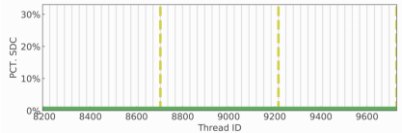

❖ All threads are **unreliable**
(SCP, MVT)

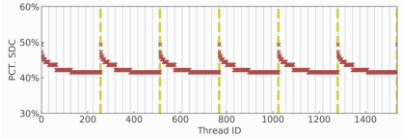

➢ **Mixed** warps: reliable threads + unreliable threads

- Well-organized

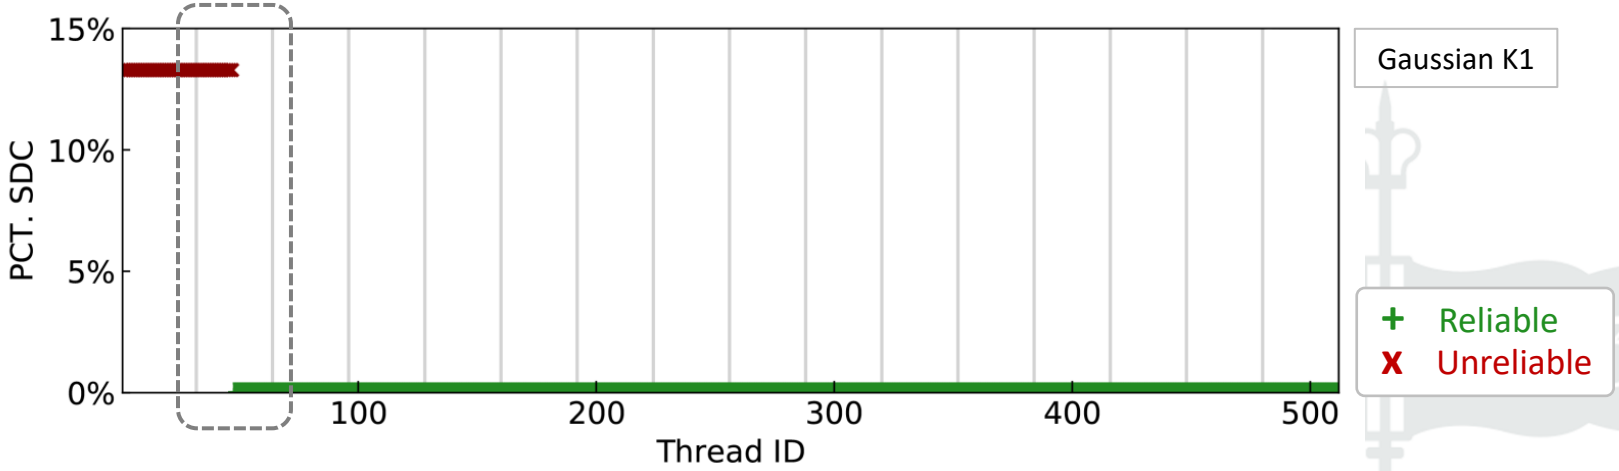

Representative benchmarks: *Gaussian K1, NearestNeighbor*

Characterization

❖ All threads are **reliable**
(SRAD K3~K4, NeuralNetwork K1~K4)

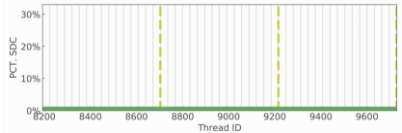

❖ All threads are **unreliable**
(SCP, MVT)

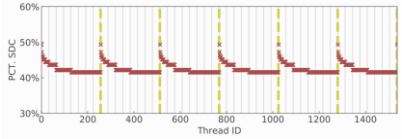

➢ **Mixed** warps: reliable threads + unreliable threads

❖ Well-organized
(Gaussian K1, NearestNeighbor)

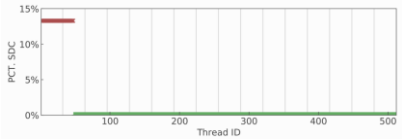

- Badly-organized (Need remapping)

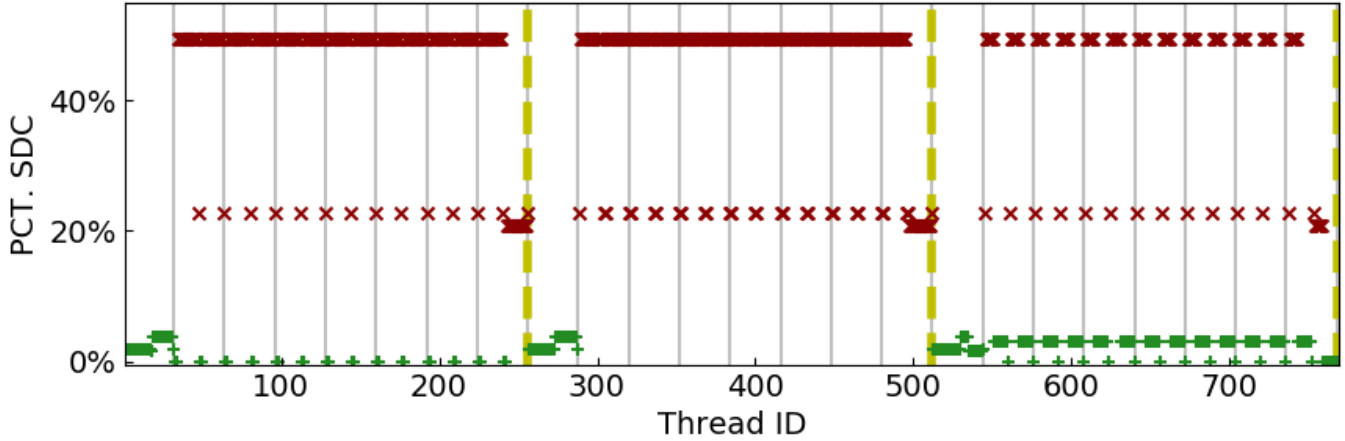

Hotspot
(First 768 threads / 9216 in total)

+ Reliable
x Unreliable

Representative benchmarks: *Gaussian K2, PathFinder, MeanFilter, Laplacian, 2DCONV, HotSpot, Jmeint*

Resilient Software Protection via Remapping

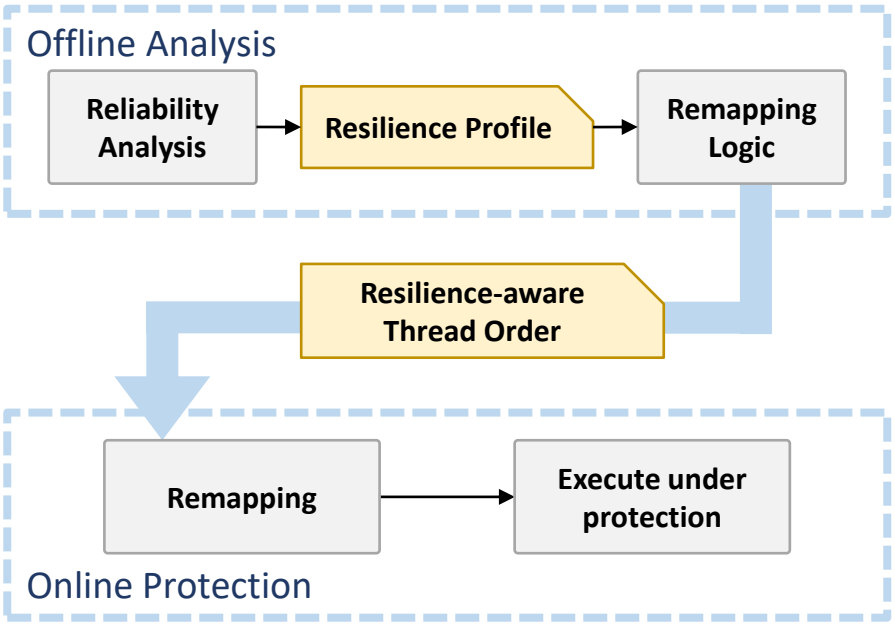

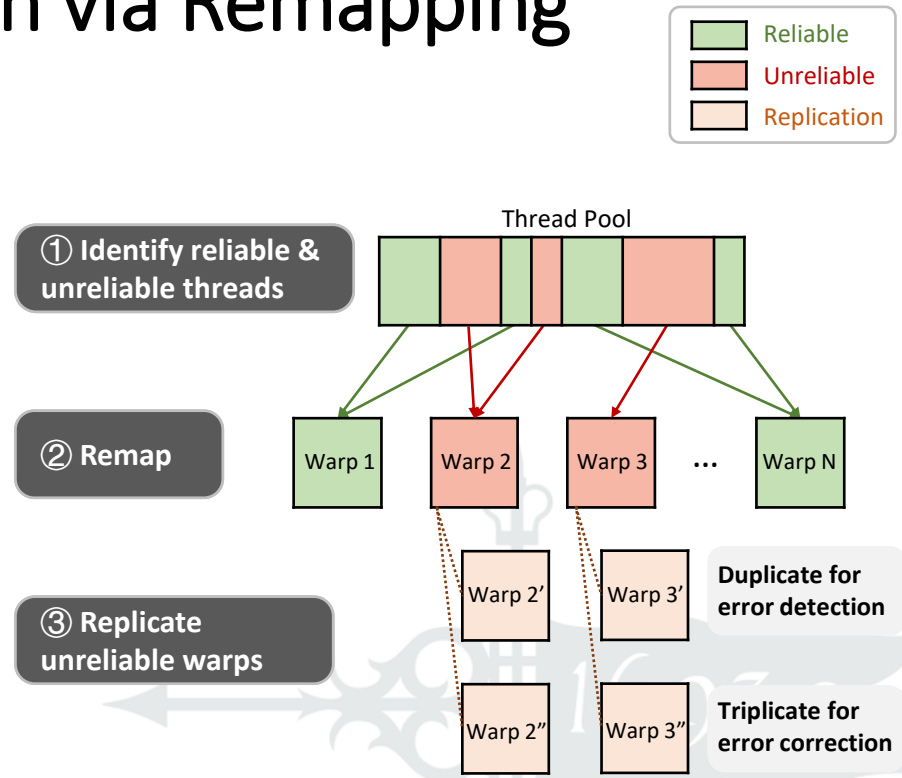

Evaluation: Effectiveness

➢ Example pattern:

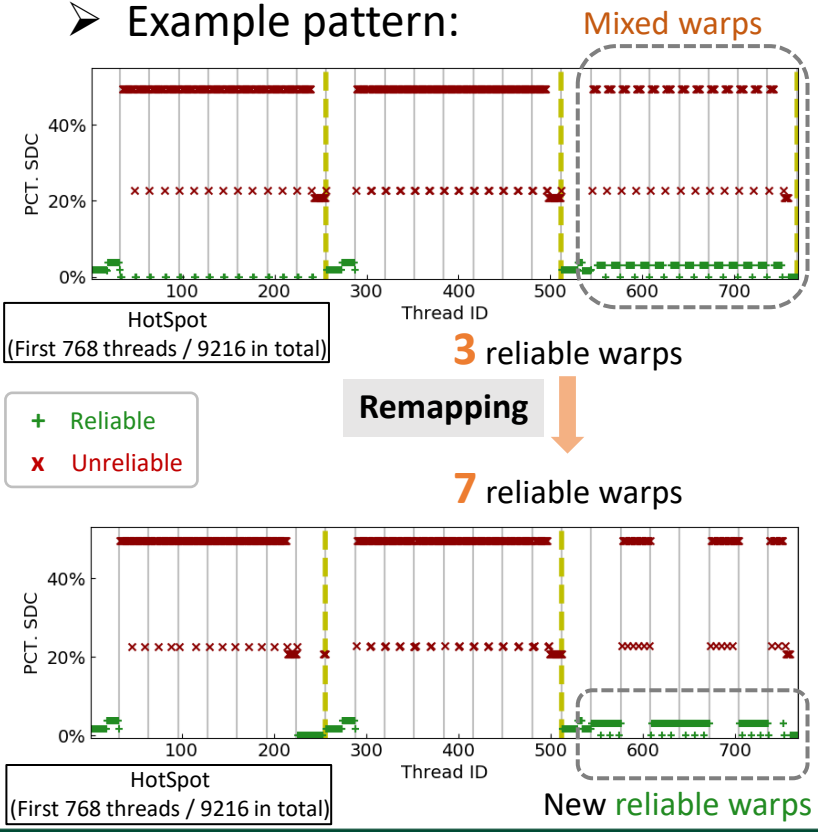

➢ Percentage of reliable warps:

23.40% → 42.08%

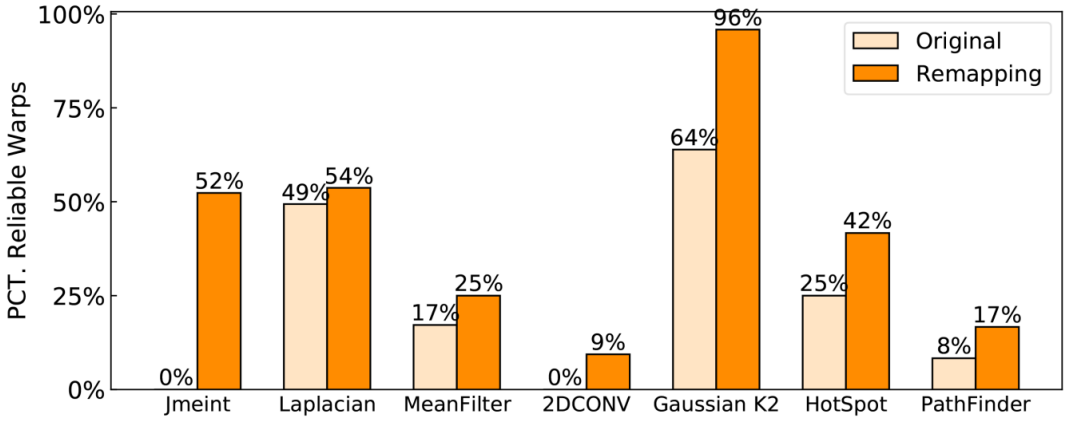

Evaluation: Execution Savings

- Error detection: *Remapping* vs. *Duplication* (***R**edundant **M**ulti-**T**hreading*)

  Average Saving: **20.61%**

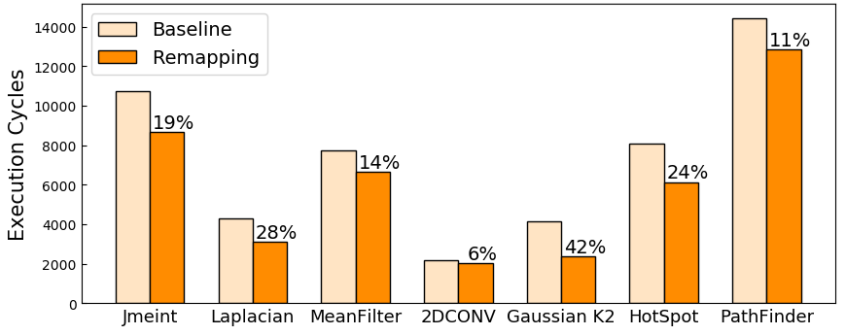

- Error correction: *Remapping vs. Triplication* (***T**riple **M**odular **R**edundancy*)

  Average Saving: **27.15%**

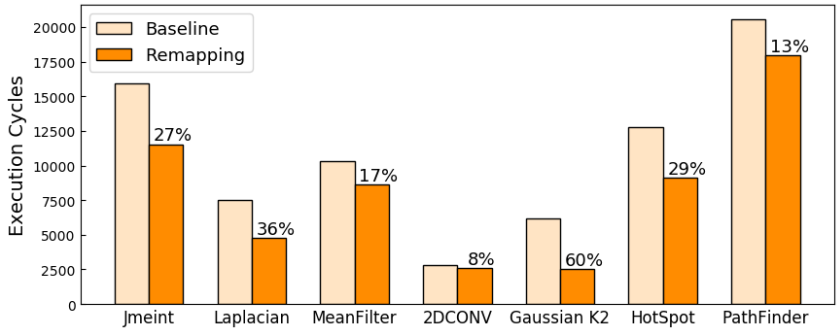

RQ: How to protect GPGPU applications selectively?

Answer:

- Reliability characterization at warp level
  - Identify reliable/unreliable/mixed warps
- Remap threads to CTAs: more reliable warps!
- Partial protection

- Low overhead: 1.63%
- Significant execution savings:
  - Error detection: 20.61%
  - Error correction: 27.15%

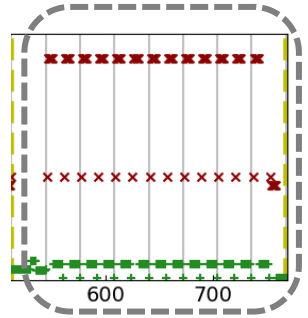

Mixed warps

Remapping

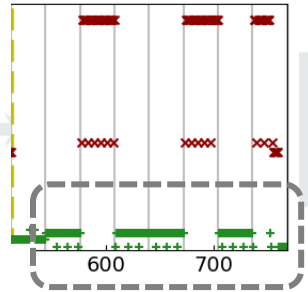

More reliable warps

WILLIAM & MARY
CHARTERED 1693

Thank you :)

Enabling Software Resilience in GPGPU Applications via Partial Thread Protection

Lishan Yang, Bin Nie, Adwait Jog, and Evgenia Smirni
William & Mary

This work is supported by NSF grant CCF-1717532