

EXPLORING INSTRUCTION FUSION OPPORTUNITIES IN GENERAL PURPOSE PROCESSORS

Sawan Singh¹ Arthur Perais² Alexandra Jimborean¹
Alberto Ros¹

¹Computer Engineering Department
University of Murcia, Spain



²TIMA, Univ. Grenoble Alpes, CNRS,
Grenoble INP, Grenoble, France



MICRO'55, Session 1B, October 3, 2022

OVERVIEW

→ INSTRUCTION FUSION:


- For example **CMP + branch** instructions in x86



OVERVIEW

→ INSTRUCTION FUSION:

- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions




OVERVIEW

→ INSTRUCTION FUSION:

- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream




OVERVIEW

→ INSTRUCTION FUSION:

- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream
- Accessing **Contiguous** memory locations (memory μ -ops)




OVERVIEW

→ INSTRUCTION FUSION:

- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream
- Accessing **Contiguous** memory locations (memory μ -ops)
- **Same** base register (memory μ -ops)



OVERVIEW

→ INSTRUCTION FUSION:

- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

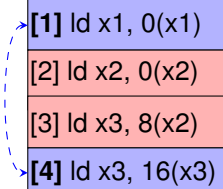
→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream
- Accessing **Contiguous** memory locations (memory μ -ops)
- **Same** base register (memory μ -ops)

→ HELIOS: Predict and fuse distant memory instructions:

- In **code**, non-consecutive instructions

| | |
|-----|---------------|
| [1] | ld x1, 0(x1) |
| [2] | ld x2, 0(x2) |
| [3] | ld x3, 8(x2) |
| [4] | ld x3, 16(x3) |



OVERVIEW

→ INSTRUCTION FUSION:


- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream
- Accessing **Contiguous** memory locations (memory μ -ops)
- **Same** base register (memory μ -ops)

→ HELIOS: Predict and fuse distant memory instructions:

- In **code**, non-consecutive instructions
- In **memory**, non-contiguous locations



OVERVIEW

→ INSTRUCTION FUSION:


- For example **CMP + branch** instructions in x86
- **One** Scheduler/RoB entry for **two** instructions

→ STATE OF THE ART:

- **Consecutive instructions** in the instruction stream
- Accessing **Contiguous** memory locations (memory μ -ops)
- **Same** base register (memory μ -ops)

→ HELIOS: Predict and fuse distant memory instructions:

- In **code**, non-consecutive instructions
- In **memory**, non-contiguous locations
- In **similarity**, different base register



OUTLINE

- 1 Overview
- 2 **Background & Motivation**
- 3 Helios
- 4 Methodology & Results
- 5 Conclusions

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- Simple ISA = Numerous fusion opportunities

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- **Simple** ISA = **Numerous** fusion opportunities
- Various **consecutive** pairs proposed by Celio *et al.* [1]
 - ★ **Contiguous** memory pairs (*load-load* & *store-store*)
 - ★ Other pairs (*slli-srli*, *add-ld* ...)

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- **Simple** ISA = **Numerous** fusion opportunities
- Various **consecutive** pairs proposed by Celio *et al.* [1]
 - ★ **Contiguous** memory pairs (*load-load* & *store-store*) → **5.6%**
 - ★ Other pairs (*slli-srli*, *add-ld* ...) → **1.1%**

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- Simple ISA = Numerous fusion opportunities
- Various consecutive pairs proposed by Celio *et al.* [1]
 - ★ Contiguous memory pairs (*load-load* & *store-store*) → 5.6%
 - ★ Other pairs (*slli-srli*, *add-ld* ...) → 1.1%

→ OBSERVATIONS:

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- **Simple ISA = Numerous** fusion opportunities
- Various **consecutive** pairs proposed by Celio *et al.* [1]
 - ★ **Contiguous** memory pairs (*load-load & store-store*) → **5.6%**
 - ★ Other pairs (*slli-srli, add-ld ...*) → **1.1%**

→ OBSERVATIONS:

★ FOCUS ON MEMORY FUSION:

- **Majority** of fused μ -ops are memory instructions
- **IPC** improvement is similar when fusing all pairs

BACKGROUND & MOTIVATION

→ RISC-V FUSION:

- **Simple** ISA = **Numerous** fusion opportunities
- Various **consecutive** pairs proposed by Celio *et al.* [1]
 - ★ **Contiguous** memory pairs (*load-load* & *store-store*) → **5.6%**
 - ★ Other pairs (*slli-srli*, *add-ld* ...) → **1.1%**

→ OBSERVATIONS:

★ FOCUS ON MEMORY FUSION:

- **Majority** of fused μ -ops are memory instructions
- **IPC** improvement is similar when fusing all pairs


★ MISSING OPPORTUNITIES:

- **Non-consecutive** & **non-contiguous** memory fusion
- **Doubles** the number of fused memory pairs (**+7.58%**)

BACKGROUND & MOTIVATION

→ NON-CONSECUTIVE FUSION CHALLENGES



- FINDING PAIRS: Non-consecutive fusion does not **scale**



BACKGROUND & MOTIVATION

→ NON-CONSECUTIVE FUSION CHALLENGES


- **FINDING PAIRS:** Non-consecutive fusion does not **scale**
- **MEMORY ACCESS:** Base register might be **modified** by another instruction




OUTLINE

- 1 Overview
- 2 Background & Motivation
- 3 **Helios**
- 4 Methodology & Results
- 5 Conclusions

- Fusion **distance** predictor
- Fuse only **memory instructions**
- Implement a *tournament-like* fusion **predictor**



- Fusion **distance** predictor
- Fuse only **memory instructions**
- Implement a **tournament-like** fusion **predictor**
- **Predict** instructions:
 - ★ Non-Consecutive
 - ★ Non-Contiguous
 - ★ Different base register



- Helios consists of 2 structures: the Unfused Committed History (UCH) and the Fusion Predictor (FP)

HELIOS

→ Helios consists of 2 structures: the Unfused Committed History (UCH) and the Fusion Predictor (FP)

- UNFUSED COMMITTED HISTORY:

- ★ Save the **sequence number** of not-already-fused μ -ops
- ★ Separate history for Loads (**6** entries) and Stores (**1** entry)
- ★ **Fills** the Fusion Predictor

HELIOS


→ Helios consists of 2 structures: the Unfused Committed History (UCH) and the Fusion Predictor (FP)


- UNFUSED COMMITTED HISTORY:

- ★ Save the **sequence number** of not-already-fused μ -ops
- ★ Separate history for Loads (6 entries) and Stores (1 entry)
- ★ **Fills** the Fusion Predictor

- FUSION PREDICTOR:


- ★ Provide **predicted distance** to μ -ops at decode
- ★ The predictor's confidence is **updated** at execute





| UCH(Load) | |
|-----------|------------|
| Seq No. | Line Addr. |
| 3 | 0x14 |
| 7 | 0x08 |
| 9 | 0x24 |


| FP | |
|-------|-----|
| Dist. | Tag |
| | |



Search @Commit

| UCH(Load) | |
|-----------|------------|
| Seq No. | Line Addr. |
| 3 | 0x14 |
| 7 | 0x08 |
| 9 | 0x24 |


| FP | |
|-------|-----|
| Dist. | Tag |
| | |



Add @Commit

| UCH(Load) | |
|-----------|------------|
| Seq No. | Line Addr. |
| 3 | 0x14 |
| 7 | 0x08 |
| 9 | 0x24 |
| 13 | 0x250 |


| FP | |
|-------|-----|
| Dist. | Tag |
| | |



Search @Commit

| UCH(Load) | |
|-----------|------------|
| Seq No. | Line Addr. |
| 3 | 0x14 |
| 7 | 0x08 |
| 9 | 0x24 |
| 13 | 0x250 |



| FP | |
|-------|-----|
| Dist. | Tag |
| | |



Match found!


| UCH(Load) | |
|-----------|------------|
| Seq No. | Line Addr. |
| 3 | 0x14 |
| 7 | 0x08 |
| 9 | 0x24 |
| 13 | 0x250 |

| FP | |
|-------|-----|
| Dist. | Tag |
| | |




HELIOS

OVERVIEW



HELIOS

OVERVIEW




Get Prediction @Decode

| FP | |
|-------|--------|
| Dist. | Tag |
| 2 | 0xFFFC |

HELIOS

OVERVIEW




Get Prediction @Decode

| FP | |
|-------|--------|
| Dist. | Tag |
| → 2 | 0xFFFC |

HELIOS


OVERVIEW



| FP | |
|-------|--------|
| Dist. | Tag |
| 2 | 0xFFFC |

HELIOS

OVERVIEW




Update Predictor @Execute

| FP | |
|-------|--------|
| Dist. | Tag |
| 2 | 0xFFFC |

HELIOS


OVERVIEW



If mispredicted

| FP | |
|-------|--------|
| Dist. | Tag |
| 2 | 0xFFFC |

→ **DEADLOCKS:** Cyclic dependencies can result in **deadlock**



- **DEADLOCKS**: Cyclic dependencies can result in **deadlock**
- **CONSISTENCY**: Fusing non-consecutive memory μ -ops may **break memory order**

- **DEADLOCKS**: Cyclic dependencies can result in **deadlock**
- **CONSISTENCY**: Fusing non-consecutive memory μ -ops may **break memory order**

Details in the paper!

OUTLINE

- 1 Overview
- 2 Background & Motivation
- 3 Helios
- 4 **Methodology & Results**
- 5 Conclusions

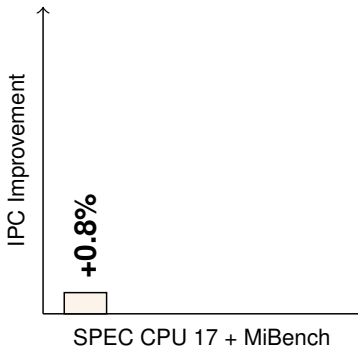
METHODOLOGY & RESULTS

- Spike + in house OoO core (**Intel Icelake**)
- SPEC CPU 2017 + MiBench applications
 - *SPEC CPU 2017*: Skip **10B**, collect results for **500M**
 - *MiBench*: Run from begin to end


METHODOLOGY & RESULTS




OtherPairs



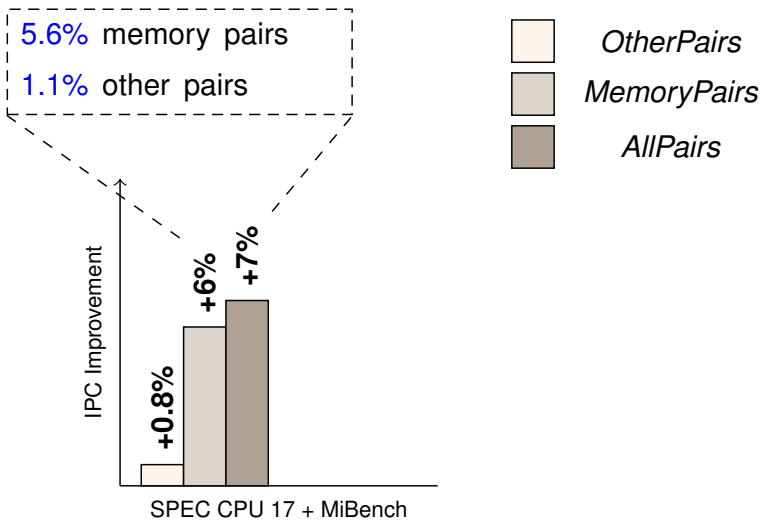
METHODOLOGY & RESULTS



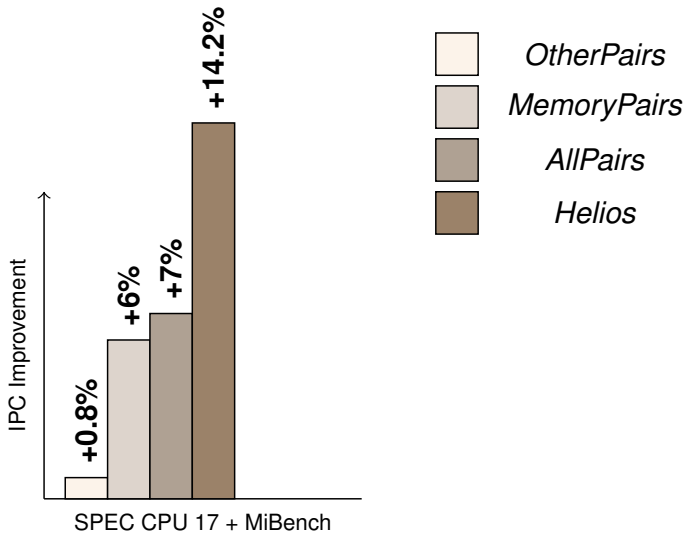
METHODOLOGY & RESULTS



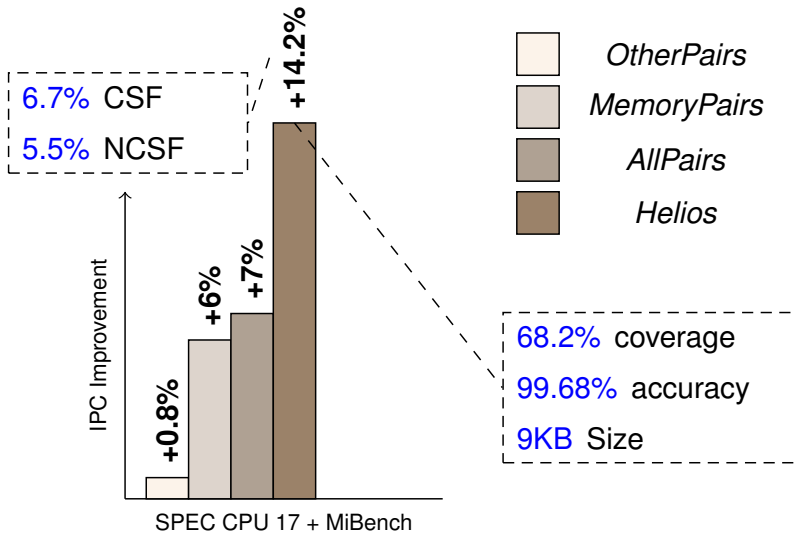
METHODOLOGY & RESULTS




METHODOLOGY & RESULTS



METHODOLOGY & RESULTS



METHODOLOGY & RESULTS



OUTLINE

- 1 Overview
- 2 Background & Motivation
- 3 Helios
- 4 Methodology & Results
- 5 **Conclusions**

CONCLUSIONS

- HELIOS: First detailed design proposal to fuse memory instructions:-
 - ★ Non-consecutive
 - ★ Non-contiguous
 - ★ Different base register
- Easy to scale

EXPLORING INSTRUCTION FUSION OPPORTUNITIES IN GENERAL PURPOSE PROCESSORS

Sawan Singh¹ Arthur Perais² Alexandra Jimborean¹
Alberto Ros¹

singh.sawan@um.es

Thank you for your attention!




European Research Council
Established by the European Commission

ECHO, ERC Consolidator Grant (No 819134)

This presentation and recording belong to the authors. No distribution is allowed without the authors' permission. 15/15

BACKUP SLIDES!

HELIOS VS VECTORIZATION



- Loop dependencies can make code not vectorizable
- Helios can predict and fuse `load(ai-1)` with `load(bi)`
- Helios can also try to fuse `load(ai-1)`, `load(bi)` & `store(ai, reg2)` with itself from the previous or next iteration over the branch