# OCC Pinning: Optimizing Concurrent Computations through Thread Pinning

# Brahmaiah Gandham

Dr. Praveen Alapati

▲□▶ ▲□▶ ▲□▶ ▲□▶ ■ ●の00

Department of Computer Science and Engineering

**Ecole Centrale School of Engineering** 

Mahindra University



#### Introduction

OCC Pinning

Experimental Evaluation

Conclusion

#### References

< 行

æ

# Uniform Memory Access (UMA) Architecture



# Non-uniform Memory Access (NUMA) Architecture



| -   |        | ~ .     |      |
|-----|--------|---------|------|
| Bra | hmaiah | n (sand | ham  |
| Dia | maiai  | i Gana  | nann |

June 19, 2024

∃ →

< 1<sup>™</sup> >

# 32 - Core Server



June 19, 2024

## Execution of a Sequential Program



## Execution of a Sequential Program



## Execution of a Sequential Program



# Execution of a Concurrent Program



# Execution of a Concurrent Program



# **NUMA** Pinning



Brahmaiah Gandham

June 19, 2024

# **NUMA** Pinning



Brahmaiah Gandham

June 19, 2024

Increases average memory access time.

- Increases average memory access time.
- Increases inter-NUMA communication overhead.

- Increases average memory access time.
- Increases inter-NUMA communication overhead.
- Increases contention resources.

- Increases average memory access time.
- Increases inter-NUMA communication overhead.
- Increases contention resources.
- Load imbalance.



June 19, 2024





Brahmaiah Gandham

June 19, 2024

#### Algorithm 1 OCC Pinning

- 1: pid ← getCurrentProcessID();
- 2: threadsInfo  $\leftarrow$  getThreadsInfoForPID(pid);
- 3: spids  $\leftarrow$  extractSPIDs(threadsInfo);
- 4: for spid in spids do
- 5: cpuIndex  $\leftarrow$  getAndIncrement() % totalCPUs;
- 6: setThreadAffinity(spid, cpuIndex);
- 7: end for

Thread pinning assigns specific software threads to dedicated processor cores.

Image: Image:

э

- Thread pinning assigns specific software threads to dedicated processor cores.
- It prevents thread thrashing (threads constantly switch cores) and boosts efficiency.

- Thread pinning assigns specific software threads to dedicated processor cores.
- It prevents thread thrashing (threads constantly switch cores) and boosts efficiency.
- Leverage operating system mechanisms to identify processes and threads.

- Thread pinning assigns specific software threads to dedicated processor cores.
- It prevents thread thrashing (threads constantly switch cores) and boosts efficiency.
- Leverage operating system mechanisms to identify processes and threads.
- Leverage the taskset command to enforce desired thread-to-core mappings.

- Thread pinning assigns specific software threads to dedicated processor cores.
- It prevents thread thrashing (threads constantly switch cores) and boosts efficiency.
- Leverage operating system mechanisms to identify processes and threads.
- Leverage the taskset command to enforce desired thread-to-core mappings.
- This effectively overrides the dynamic behavior of the operating system scheduler.

# **Experimental System**

Hardware:

- Intel(R) Xeon(R) Platinum 8260 CPU @ 2.40GHz
- RAM: 512GB RAM
- CPU(s): 96
- Thread(s) per core: 2
- Core(s) per socket: 24
- Socket(s): 2
- L1d cache: 32K, L1i cache: 32K, L2 cache: 1MB, L3 cache: 36MB

#### Software:

- Ubuntu 22.04 LTS
- JDK Runtime version 19.02

#### **Pinning configurations:**

- OCC Pinning
- NUMA Pinning
- No Pinning

#### **Concurrent Data structures:**

- Contention Adapting Binary Search Tree [1]
- Striped Hash Set [2]
- Lazy List [3]

#### Workload representation:

xC-yl-zD

Read-only workload (100C-0I-0D). 

э

- Read-only workload (100C-0I-0D).
- Read dominant workload (70C-20I-10D).

- Read-only workload (100C-0I-0D).
- Read dominant workload (70C-20I-10D).
- Balanced workload (50C-25I-25D).

- Read-only workload (100C-0I-0D).
- Read dominant workload (70C-20I-10D).
- Balanced workload (50C-25I-25D).
- Update dominant workload (30C-35I-35D).

- Read-only workload (100C-0I-0D).
- Read dominant workload (70C-20I-10D).
- Balanced workload (50C-25I-25D).
- Update dominant workload (30C-35I-35D).
- Update only workload (0C-50I-50D).

## Experimental Evaluation



Figure 1: Average throughput of a CA tree, striped hash set, and lazy list implementations.

| <b>D</b> |       |         |     |
|----------|-------|---------|-----|
| ' Kra    | hmaia | h (sand | ham |
|          | mana  |         |     |

June 19, 2024

< 行

#### CA Tree

 OCC pinning achieved an average speedup of 1.10× compared to both NUMA pinning and No pinning.

#### Striped Hash Set

OCC pinning achieved an average speedup of 1.25× compared to both NUMA pinning and No pinning.

#### Lazy List

OCC pinning achieved an average speedup of 1.26× compared to both NUMA pinning and No pinning.

In the NUMA world, the OCC pinning strategy helps for efficient mapping between software threads and hardware threads.

- In the NUMA world, the OCC pinning strategy helps for efficient mapping between software threads and hardware threads.
- This approach achieves better resource utilization, minimizes interference, and cache misses, and outperforms traditional scheduling policies.

- In the NUMA world, the OCC pinning strategy helps for efficient mapping between software threads and hardware threads.
- This approach achieves better resource utilization, minimizes interference, and cache misses, and outperforms traditional scheduling policies.
- It optimizes performance by utilizing OS integration, granular affinity binding, and distributing threads across CPU cores.

- In the NUMA world, the OCC pinning strategy helps for efficient mapping between software threads and hardware threads.
- This approach achieves better resource utilization, minimizes interference, and cache misses, and outperforms traditional scheduling policies.
- It optimizes performance by utilizing OS integration, granular affinity binding, and distributing threads across CPU cores.
- It paves the way for realizing the full potential of concurrent applications.

- In the NUMA world, the OCC pinning strategy helps for efficient mapping between software threads and hardware threads.
- This approach achieves better resource utilization, minimizes interference, and cache misses, and outperforms traditional scheduling policies.
- It optimizes performance by utilizing OS integration, granular affinity binding, and distributing threads across CPU cores.
- It paves the way for realizing the full potential of concurrent applications.
- Future work includes the potential extension of this strategy to GPUs.

#### References

- Konstantinos Sagonas and Kjell Winblad. 2015. Contention Adapting Search Trees. In 2015 14th International Symposium on Parallel and Distributed Computing. 215–224.
- Steve Heller, Maurice Herlihy, Victor Luchangco, Mark Moir, William N. Scherer, and Nir Shavit. 2006. A Lazy Concurrent List-Based Set Algorithm.
- Maurice Herlihy and Nir Shavit. 2012. The Art of Multiprocessor Programming, Revised Reprint (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- The Linux Foundation. 2022. numactl(8) Linux man page. https://man7.org/linux/man-pages/man8/numactl.8.html.

# Thank You ③

Brahmaiah Gandham

June 19, 2024

< 4 P→ <

→ < Ξ →</p>

25 / 25

э