

Cache Coherence Scaling on Manycore Systems

Oliver Türk
`oliver.tuerk@fau.de`

Friedrich-Alexander-Universität Erlangen-Nürnberg

17.01.2017

Talbe Of Contents

Motivation

Mainstream cache implementation today

Scaling

- The Naive Way

- Inexact Tracking of Sharers

- Cache Hierarchy

References


- ▶ Complexity will increase

- ▶ Complexity will increase
- ▶ Pollacks Rule - problematic

- ▶ Complexity will increase
- ▶ Pollacks Rule - problematic
- ▶ Solution: Split processor into cores

- ▶ Complexity will increase
- ▶ Pollacks Rule - problematic
- ▶ Solution: Split processor into cores
- ▶ More cores - also problematic

Pollack's Rule Inverted



Motivation

Mainstream cache implementation today

Scaling


- The Naive Way

- Inexact Tracking of Sharers

- Cache Hierarchy

References


Mainstream cache implementation today



- hardware coherent

Intel i7 6950X processor with
25Mb L3 cache (1)


Mainstream cache implementation today



- ▶ hardware coherent
- ▶ inclusive

Intel i7 6950X processor with
25Mb L3 cache (1)


Mainstream cache implementation today



- ▶ **hardware coherent**
- ▶ **inclusive**
- ▶ **last level shared**

Intel i7 6950X processor with
25Mb L3 cache (1)


Mainstream cache implementation today



Intel i7 6950X processor with
25Mb L3 cache (1)

- ▶ **hardware coherent**
- ▶ inclusive
- ▶ last level shared
- ▶ size of last level greater than aggregate below


Mainstream cache implementation today



Intel i7 6950X processor with
25Mb L3 cache (1)

- ▶ **hardware coherent**
- ▶ inclusive
- ▶ last level shared
- ▶ size of last level greater than aggregate below
- ▶ exact tracking of sharers with **1** bit for each core

Mainstream cache implementation today



Cache implementation on Intel i7 CPU (2)

Motivation

Mainstream cache implementation today

Scaling

- The Naive Way

- Inexact Tracking of Sharers

- Cache Hierarchy

References

Today we have typically..

- ▶ 64Byte per cache line
- ▶ 4 cores
- ▶ 1 Bit per core for tracking
- ▶ overhead below **5%**

Lets assume..

- ▶ 64Byte per cache line
- ▶ 1024 cores
- ▶ 1 Bit per core for tracking

Lets assume..

- ▶ 64Byte per cache line
- ▶ 1024 cores
- ▶ 1 Bit per core for tracking

This results in..

- ▶ 1024 bit (**64Byte**) for tracking
- ▶ overhead of ???

Lets assume..

- ▶ 64Byte per cache line
- ▶ 1024 cores
- ▶ 1 Bit per core for tracking

This results in..

- ▶ 1024 bit (**64Byte**) for tracking
- ▶ overhead of **50%**

not acceptable!

Motivation

Mainstream cache implementation today

Scaling

The Naive Way

Inexact Tracking of Sharers

Cache Hierarchy

References

Lets assume..

- ▶ 64Byte per cache line
- ▶ 1024 cores
- ▶ **32** Bits for tracking in total

Lets assume..

- ▶ 64Byte per cache line
- ▶ 1024 cores
- ▶ 32 Bits for tracking in total

This results in..

- ▶ overhead of 6.25%

and is acceptable.

How it works..

- ▶ inclusive shared cache

How it works..

- ▶ inclusive shared cache
- ▶ single level shared cache

How it works..

- ▶ inclusive shared cache
- ▶ single level shared cache
- ▶ exact tracking for up to **32** sharers

How it works..

- ▶ inclusive shared cache
- ▶ single level shared cache
- ▶ exact tracking for up to **32** sharers
- ▶ inexact tracking for more sharers

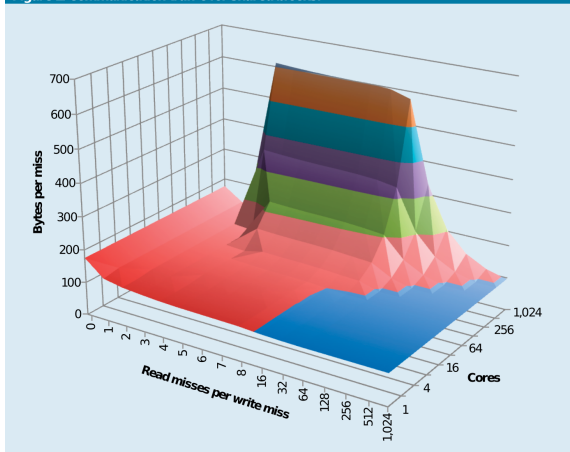
How it works..

- ▶ inclusive shared cache
- ▶ single level shared cache
- ▶ exact tracking for up to **32** sharers
- ▶ inexact tracking for more sharers

High traffic with more than 32 sharers (threads)

Scaling - inexact tracking of sharers

Figure 1. Communication traffic c for shared blocks.



Traffic with inexact tracking of sharers (3)

Motivation

Mainstream cache implementation today

Scaling

The Naive Way

Inexact Tracking of Sharers

Cache Hierarchy

References

Advantages..

- ▶ tracking of subsequent clusters only

Advantages..

- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent

Advantages..

- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent
- ▶ cache coherent

Advantages..

- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent
- ▶ cache coherent
- ▶ software backwards compatibility

Advantages..

- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent
- ▶ cache coherent
- ▶ software backwards compatibility
- ▶ invariable traffic

Advantages..

- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent
- ▶ cache coherent
- ▶ software backwards compatibility
- ▶ invariable traffic

Disadvantages..

Advantages..


- ▶ tracking of subsequent clusters only
- ▶ access to memory stays transparent
- ▶ cache coherent
- ▶ software backwards compatibility
- ▶ invariable traffic

Disadvantages..

- ▶ higher latency between distant cores

Scaling - cache hierarchy

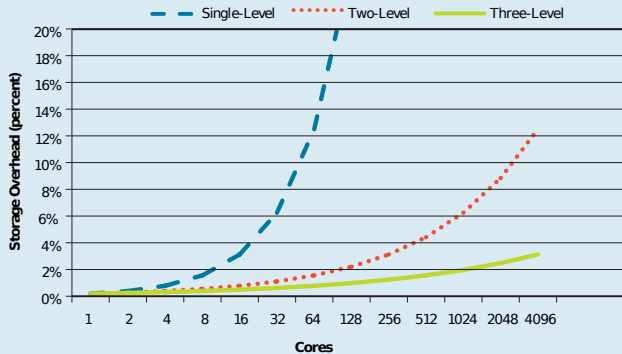
Figure 1. Communication traffic for shared blocks.



Traffic with exact tracking of sharers (3)


Scaling - cache hierarchy

Figure 3. storage overhead in shared caches.



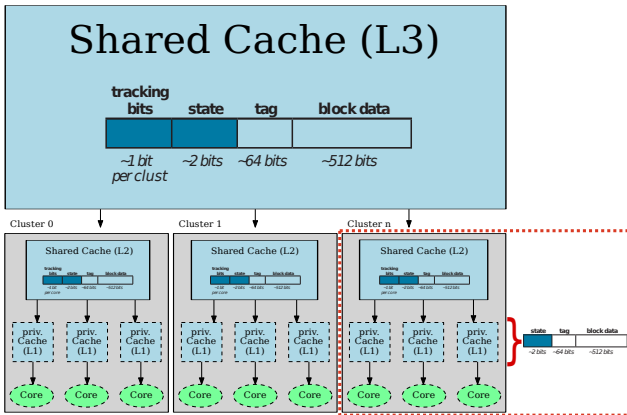
Storage Overhead for multiple hierarchy levels (3)

Scaling - cache hierarchy: single-level shared

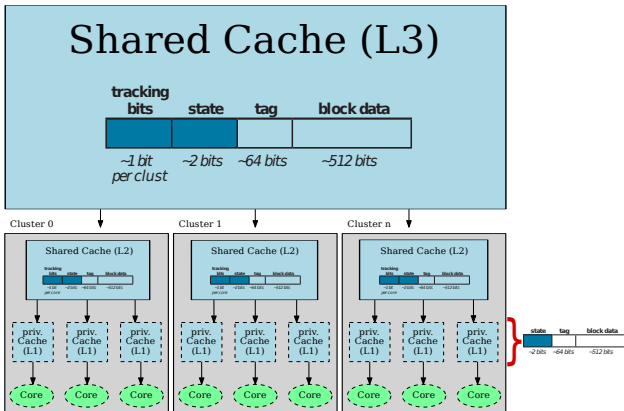


Shared and private caches with 2 levels

Scaling - cache hierarchy: two-level shared



Scaling - cache hierarchy: two-level shared



Shared and private caches with 3 levels

Mainstream processor: i7 6700K..

- ▶ 4 cores
- ▶ *one-level* hierarchy
- ▶ 256 KiB per L2
- ▶ 8 times of aggregate
- ▶ **$1.75 * 10^9$** transistors

Lets assume..

- ▶ 1024 cores
- ▶ *two-level* hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

Scaling - cache memory in transistors

Lets assume..

- ▶ 1024 cores
- ▶ *two-level* hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

Scaling - cache memory in transistors

Lets assume..

- ▶ 1024 cores
- ▶ *two-level* hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

- ▶ **32** cores per cluster

Scaling - cache memory in transistors

Lets assume..

- ▶ 1024 cores
- ▶ *two-level* hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

- ▶ **32** cores per cluster
- ▶ **32** cluster

Lets assume..

- ▶ 1024 cores
- ▶ *two-level* hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

- ▶ **32** cores per cluster
- ▶ **32** cluster
- ▶ L2 with $256\text{KiB} * 32\text{cores} * 8 = \mathbf{64\ MiB}$

Lets assume..

- ▶ 1024 cores
- ▶ two-level hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

- ▶ **32** cores per cluster
- ▶ **32** cluster
- ▶ L2 with $256\text{KiB} * 32\text{cores} * 8 = \mathbf{64\ MiB}$
- ▶ L3 with $64\text{MiB} * 32\text{clusters} * 8 = \mathbf{16\ GiB}$

Lets assume..

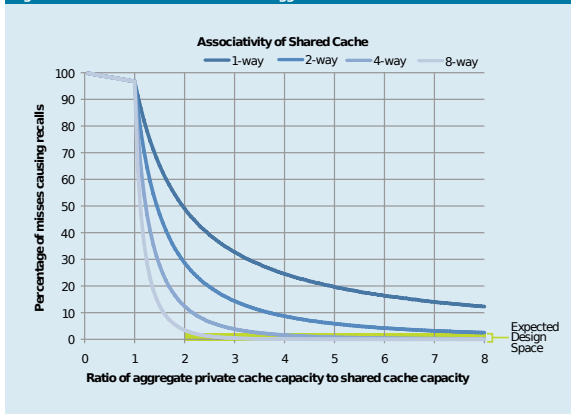
- ▶ 1024 cores
- ▶ two-level hierarchy
- ▶ 6 transistors per flip-flop
- ▶ 256 KiB per L1
- ▶ 8 times of aggregate

This results in..

- ▶ 32 cores per cluster
- ▶ 32 cluster
- ▶ L2 with $256\text{KiB} * 32\text{cores} * 8 = 64 \text{ MiB}$
- ▶ L3 with $64\text{MiB} * 32\text{clusters} * 8 = 16 \text{ GiB}$
- ▶ $824 * 10^9$ transistors

Scaling - cache memory

Figure 4. Likelihood a shared cache miss triggers a recall.



Associativity VS capacity (3)



I. Cutress, *Intel i7 6950X DIE*,
<http://www.anandtech.com/show/10337/the-intel-broadwell-e-review-core-i7-6950x-6900k-6850k-and-6800k-tested-up-to-10-cores>, [Online, accessed 15-01-2017], 2016.



J. Treibig, G. Hager, presented at the International Conference on Parallel Processing and Applied Mathematics, pp. 615–624.



M. M. K. Martin, M. D. Hill, D. J. Sorin, *Commun. ACM* **55**, 78–89, ISSN: 0001-0782 (July 2012).