Data Processing on Modern Hardware

Jens Teubner, TU Dortmund, DBIS Group
`jens.teubner@cs.tu-dortmund.de`

Summer 2016

Part V

Execution on Multiple Cores

Example: Star Joins

Task: run parallel instances of the query ($\nearrow$ introduction)

dimension

```
SELECT SUM(lo_revenue)
  FROM part, lineorder      fact table
 WHERE p_partkey = lo_partkey
   AND p_category <= 5
```

```
        |
        ⋈
       / \
      σ   lineorder
      |
    part
```

To implement $\bowtie$ use either

- a **hash join** or
- an **index nested loops join**.

Execution on "Independent" CPU Cores

Co-run independent instances on different CPU cores.

Concurrent queries may seriously affect each other's performance.

Shared Caches

In Intel Core 2 Quad systems, two cores **share** an L2 Cache:

What we saw was **cache pollution**.

$\rightarrow$ How can we avoid this cache pollution?

Cache Sensitivity

Dependence on cache sizes for some TPC-H queries:

(a) L2 Miss Rate

(b) CPI

Some queries are more sensitive to cache sizes than others.

- **cache sensitive:** hash joins
- **cache insensitive:** index nested loops joins; hash joins with very small or very large hash table

Locality Strength

This behavior is related to the **locality strength** of execution plans:

Strong Locality
  small data structure; reused very frequently

  - *e.g.*, small hash table

Moderate Locality
  frequently reused data structure; data structure $\approx$ cache size

  - *e.g.*, moderate-sized hash table

Weak Locality
  data not reused frequently or data structure $\gg$ cache size

  - *e.g.*, large hash table; index lookups

Execution Plan Characteristics

Locality effects how caches are used:

	strong	moderate	weak
amount of cache used	small	large	large
amount of cache needed	small	large	small

cache pollution

Plans with **weak locality** have most severe impact on co-running queries.

Impact of co-runner on query:

	strong	moderate	weak
strong	low	moderate	high
moderate	moderate	high	high
weak	low	low	low

Experiments: Locality Strength

Index Join to Index Join —+—
Index Join to Hash Join —×—
Hash Join to Index Join —*—
Hash Join to Hash Join ⋯□⋯
Index Join to Index Join (bitmap scan) —■—

Locality-Aware Scheduling

An optimizer could use knowledge about localities to **schedule** queries.

- **Estimate** locality during query analysis.
    - Index nested loops join $\rightarrow$ weak locality
    - Hash join:

        $$\begin{array}{rcl} \text{hash table} \ll \text{cache size} & \rightarrow & \text{strong locality} \\ \text{hash table} \approx \text{cache size} & \rightarrow & \text{moderate locality} \\ \text{hash table} \gg \text{cache size} & \rightarrow & \text{weak locality} \end{array}$$

- **Co-schedule** queries to minimize (the impact of) cache pollution.

✎ **Which queries should be co-scheduled, which ones not?**
- Only run weak-locality queries next to weak-locality queries.
    - $\rightarrow$ They cause high pollution, but are not affected by pollution.
- Try to co-schedule queries with small hash tables.

PostgreSQL; 4 queries (different `p_category`s); for each query: $2 \times$ hash join plan, $2 \times$ INLJ plan; impact reported for hash joins:


Source: Lee *et al.* VLDB 2009.

Cache Pollution

Weak-locality plans cause cache pollution, because they **use** much cache space even though they do not strictly **need** it.

By **partitioning** the cache we could reduce pollution with little impact on the weak-locality plan.


But:

- Cache allocation controlled by **hardware**.

Cache Organization

Remember how caches are organized:

- The **physical address** of a memory block determines the **cache set** into which it could be loaded.


Thus,

- We can **influence hardware behavior** by the **choice of physical memory allocation**.

Page Coloring

The address ↔ cache set relationship inspired the idea of **page colors**.

- Each memory page is assigned a **color**.[5]
- Pages that map to the **same cache sets** get the **same color**.

cache set

cache

memory page

memory

✎ **How many colors are there in a typical system?**

---

[5]Memory is organized in **pages**. A typical **page size** is **4 kB**.

Page Coloring

- By using memory only of certain colors, we can effectively restrict the cache region that a query plan uses.

Note that

- Applications (usually) have **no control** over physical memory.
- Memory allocation and virtual $\leftrightarrow$ physical mapping are handled by the **operating system**.
- We need **OS support** to achieve our desired **cache partitioning**.

MCC-DB: Kernel-Assisted Cache Sharing

MCC-DB ("Minimizing Cache Conflicts"):

- Modified Linux 2.6.20 kernel
    - Support for **32 page colors** (4 MB L2 Cache: 128 kB per color)
    - **Color specification** file for each process (may be modified by application at any time)

- Modified instance of PostgreSQL
    - **Four colors** for regular buffer pool

        ✎ **Implications on buffer pool size (16 GB main memory)?**

    - For **strong- and moderate-locality** queries, allocate colors as needed (*i.e.,* as estimated by query optimizer)

Experiments

Moderate-locality hash join and weak-locality co-runner (INLJ):

Experiments

Moderate-locality hash join and weak-locality co-runner (INLJ):

Source: Lee *et al.* *VLDB 2009.*

Experiments: MCC-DB

PostgreSQL; 4 queries (different `p_category`s); for each query: 2 × hash join plan, 2 × INLJ plan; impact reported for hash joins:

hash table size

Source: Lee *et al.* VLDB 2009.

Building a Shared-Memory Multiprocessor

What the programmer likes to think of...

CPU core	CPU core	CPU core	CPU core

shared main-memory

✎ **Scalability? Moore's Law?**

Centralized Shared-Memory Multiprocessor

Caches help mitigate the bandwidth bottleneck(s).

- A **shared bus** connects CPU cores and memory.
    - $\rightarrow$ the shared bus may or may not be shared physically.
- The Intel Core architecture, *e.g.*, implemented this design.

Centralized Shared-Memory Multiprocessor

The shared bus design with caches makes sense:

+ **symmetric design**; uniform access time for every memory item from every processor
+ **private data** gets **cached locally**
  - $\rightarrow$ behavior identical to that of a uniprocessor
? **shared data** will be **replicated to private caches**
  - $\rightarrow$ Okay for parallel **reads**.
  - $\rightarrow$ But what about **writes** to the replicated data?
  - $\rightarrow$ In fact, we'll want to use memory as a mechanism to communicate between processors.

The approach does have **limitations**, too:

- For **large core counts**, shared bus may still be a (bandwidth) bottleneck.

Caches and Shared Memory

Caching/replicating shared data can cause problems:

read x (4)
$x := 42$ (42)

CPU   CPU

read x (4)
read x (4) ⚡

cache   cache
$x = 42$ $x = 4$

shared main memory $x = 42$

Challenges:

- Need **well-defined semantics** for such scenarios.
- Must **efficiently implement** that semantics.

Cache Coherence

The desired property (semantics) is **cache coherence**.

Most importantly:[6]

> *Writes to the **same location** are **serialized**; two writes to the same location (by any two processors) are seen in the same order by all processors.*

Note:

- We did not specify **which** order will be seen by the processors.
    - → ✎ **Why?**

---

[6]We also demand that a read by processor P will return P's most recent write, provided that no other processor has written to the same location meanwhile. Also, every write must be visible by other processors after some time.

Cache Coherence Protocol

Multiprocessor (or multicore) systems maintain coherence through a **cache coherence protocol**.

Idea:

- Know **which cache/memory** holds the **current value** of the item.
- Other replicas might be stale.

Two alternatives:

1. **Snooping-Based Coherence**
   - $\rightarrow$ All processors communicate to agree on item states.

2. **Directory-Based Coherence**
   - $\rightarrow$ A centralized **directory** holds information about state/whereabouts of data items.

Snooping-Based Cache Coherence

Rationale:

- All processors have access to a **shared bus**.
- Can snoop on the bus to track other processors' activities.

Use to track the **sharing state** of each cached item:

Meta data for each **cache block**:

- (sharing) state
- block identification (tag)

(sharing) state   tag   data

✎ **Ignoring Multiprocessors for a moment, which state information might make sense to keep?**

Strategy 1: Write Update Protocol

Idea:

- On **every write**, propagate the write to **every copy**.
  - $\rightarrow$ Use bus to **broadcast writes**.[7]

✎ **Pros/Cons of this strategy?**

---

[7]The protocol is thus also called *write broadcast* protocol.

Strategy 2: Write Invalidate Protocol

Idea:

- **Before writing** an item, **invalidate all other copies**.

Activity	Bus	Cache A	Cache B	Memory
				$x = 4$
A reads x	cache miss for x	$x = 4$		$x = 4$
B reads x	cache miss for x	$x = 4$	$x = 4$	$x = 4$
A reads x	– (cache hit)	$x = 4$	$x = 4$	$x = 4$
B writes x	invalidate x	$\cancel{x = 4}$	$x = 42$	$x = 4$[8]
A reads x	cache miss for x	$x = 42$	$x = 42$	$x = 42$

$\rightarrow$ Caches will **re-fetch** invalidated items automatically.

- Since the bus is shared, other caches may answer "cache miss" messages ($\rightsquigarrow$ necessary for write-back caches).

---

[8]With write-through caches, memory will be updated immediately.

Write Invalidate—Realization

Realization:

- To **invalidate**, **broadcast** address on bus.
- All processors continuously **snoop on bus**:
    - *invalidate* message for address held in own cache
        - $\rightarrow$ Invalidate own copy
    - *miss* message for address held in own cache
        - $\rightarrow$ Reply with own copy (for write-back caches)
        - $\rightarrow$ Memory will see this and abort its own read

✎ **What if two processors try to write at the same time?**

Write Invalidate—Tracking Sharing States

Through snooping, can monitor all bus activities by all processors.
  $\rightarrow$ Track **sharing state**.

Idea:

- Sending an *invalidate* will make local copy the only one valid.
    - $\rightarrow$ Mark local cache line as *modified* ($\approx$ *exclusive*).
- If a local cache line is already *modified*, writes need **not** be announced on the bus (no *invalidate* message).
- Upon read request by other processor:
    - $\rightarrow$ If local cache line has state *modified*, **answer** the request by sending local version.
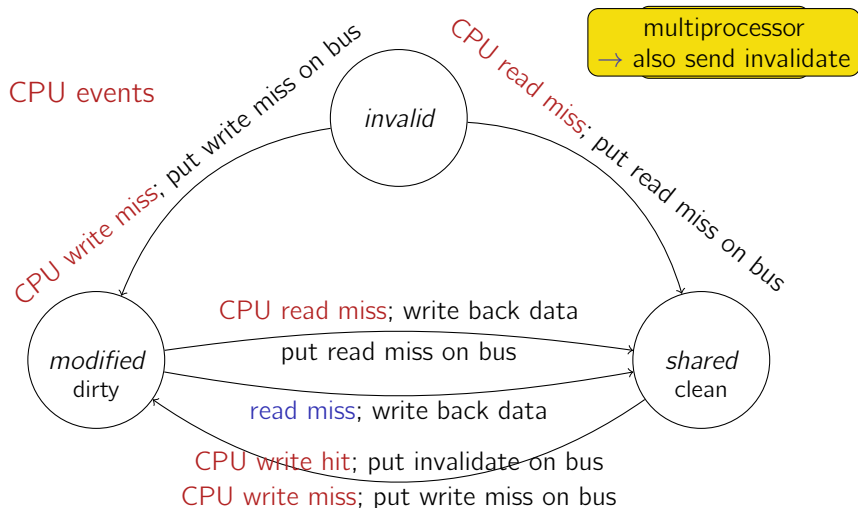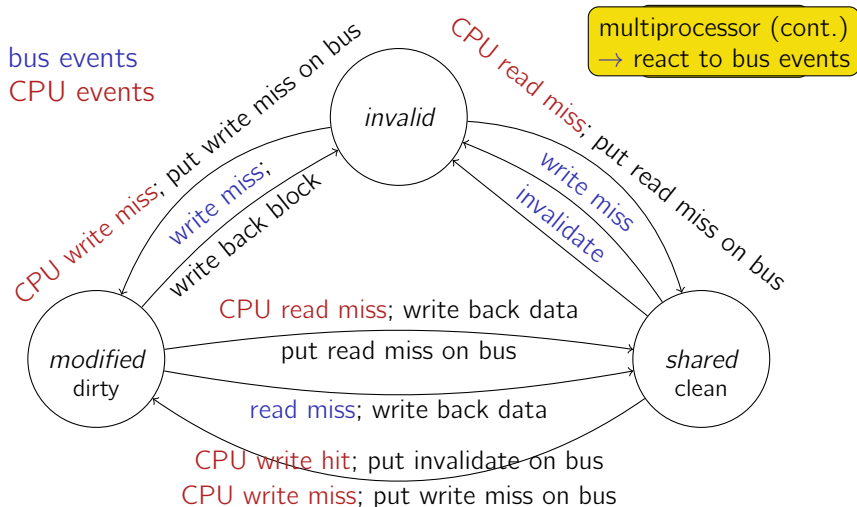   - $\rightarrow$ Change local cache state to *shared*.

Write Invalidate—State Machine
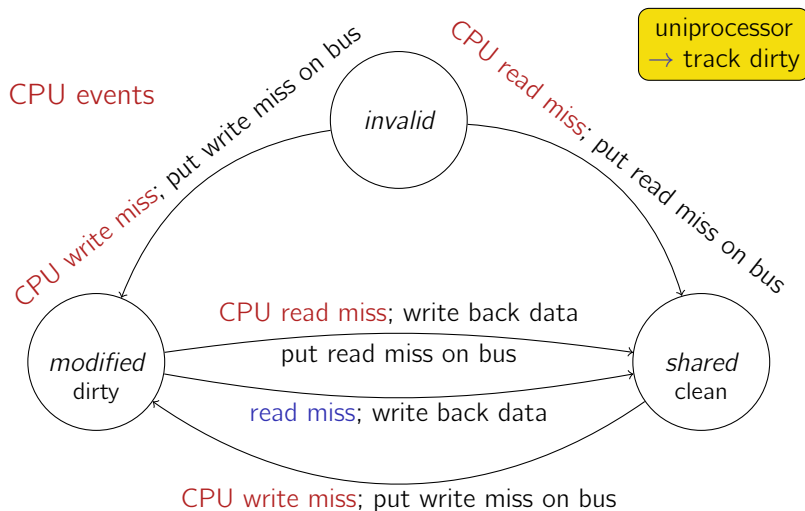
Local caches track sharing states using a **state machine**.

Write Invalidate—State Machine

Local caches track sharing states using a **state machine**.

CPU events

multiprocessor
→ also send invalidate

CPU write miss; put write miss on bus

CPU read miss; put read miss on bus

invalid

modified
dirty

shared
clean

CPU read miss; write back data
put read miss on bus

read miss; write back data

CPU write hit; put invalidate on bus
CPU write miss; put write miss on bus

Write Invalidate—State Machine

Local caches track sharing states using a **state machine**.

multiprocessor (cont.)
→ react to bus events

bus events
CPU events

- CPU write miss; put write miss on bus
- write miss; write back block
- CPU read miss; put read miss on bus
- write miss
- invalidate
- CPU read miss; write back data — put read miss on bus
- read miss; write back data
- CPU write hit; put invalidate on bus
- CPU write miss; put write miss on bus

States: *invalid*, *modified* / dirty, *shared* / clean

Write Invalidate—Notes

Notes:

- Because of the three states *modified*, *shared*, and *invalid*, the protocol on the previous slide is also called **MSI protocol**.

- The *Write Invalidate* protocol ensures that any valid cache block is either

    - in the **shared state in one or more caches** or
    - in the **modified state in exactly one cache**.
      (Any transition to the *modified* state invalidates all other copies of the block; whenever another cache fetches a copy of the block, the *modified* state is left.)

- The *MSI* protocol also ensures that every *shared* item has also been written back to memory. ◁

MSI Protocol—Extensions

Actual systems often use **extensions** to the MSI protocol, e.g.,

MESI (E for exclusive)

- Distinguish between exclusive (but clean) and modified (which implies that the copy is exclusive).
- Optimizes the (common) case when an item is first read ($\leadsto$ exclusive) then modified ($\leadsto$ modified).

MESIF (F for forward)

- In M(E)SI, if shared items are served by caches (not only by memory), **all** caches might answer miss requests.
- MESIF extends the protocol, so at most one shared copy of an item is marked as forward. Only this cache will respond to misses on the bus.
- Intel i7 employs the MESIF protocol.

MSI Protocol—Extensions

MOESI (*O* for *owned*)

- *owned* marks an item that might be outdated in memory; the owner cache is responsible for the item.
- The owner **must** respond to data requests (since main memory might be outdated).
- *MOESI* allows moving around dirty data between caches.
- The AMD Opteron uses the *MOESI* protocol.
- MOESI avoids the need to write every shared cache block back to memory ($\rightsquigarrow \lhd$).

Limitations of a Shared Bus

Limitations of a shared bus:

- Large core counts $\rightarrow$ **high bandwidth**.
- Shared buses cannot satisfy bandwidth demands of modern multiprocessor systems.

Therefore:

- **Distribute** memory
- Communicate through **interconnection network**

Consequence:

- **Non-uniform memory access (NUMA)** characteristics

Bandwidth Demand

E.g., Intel Xeon E7-8880 v3:

- 2.3 GHz clock rate
- 18 cores per chip (36 threads)
- Up to 8 processors per system

Back-of-the-envelope calculation:

- 1 byte per cycle per core $\rightarrow$ 331 GB/s
- Data-intensive applications might demand much more!

Shared memory bus?

- Modern bus standards can deliver at most a few ten GB/s.
- Switching very high bandwidths is a challenge.

Distributed Shared Memory

Idea: Distribute memory

$\rightarrow$ Attach to individual compute nodes

Example: 8-Way Intel Nehalem-EX

- Interconnect: Intel Quick Path Interconnect (QPI)[9]
- Memory may be local, one hop away, or two hops away.
    - $\rightarrow$ Non-uniform memory access (NUMA)

---

[9]The AMD counterpart is HyperTransport.

Distributed Memory and Snooping

Idea:

- Extend snooping to distributed memory.
- **Broadcast** coherence traffic, send data **point-to-point**.

✎ **Problem solved?**
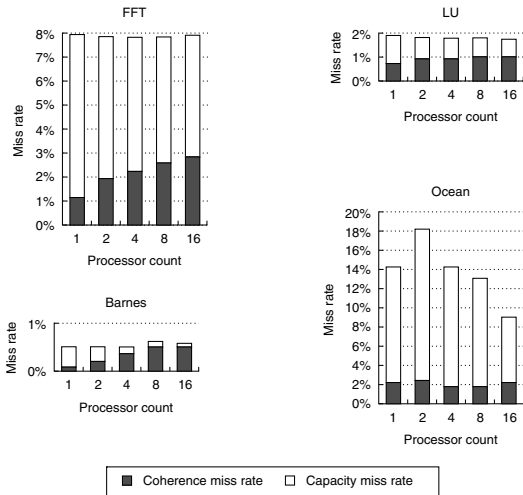
Example:

- Scientific Applications
- ↗ Hennessy & Patterson, Sect. I.5

→ AMD Opteron is a system that still uses the approach.

Directory-Based Cache Coherence

To avoid all-broadcast coherence protocol:

- Use a **directory** to keep track of which item is replicated where.
- Direct coherence messages only to those nodes that actually need them.

Directory:

- Either keep a **global directory** ($\sim$ scalability?).
- Or define a **home node** for each memory address.
  - $\rightarrow$ Home node holds directory for that item.
  - $\rightarrow$ Typically: distribute directory along with memory.

Protocol now involves

- **directory/-ies** (at item home node(s)),
- **individual caches** (local to processors).

Parties communicate **point-to-point** (no broadcasts).

Directory-Based Cache Coherence

Messages sent by individual nodes:

Message type	Source	Destination	Message contents	Function of this message
Read miss	Local cache	Home directory	P, A	Node P has a read miss at address A; request data and make P a read sharer.
Write miss	Local cache	Home directory	P, A	Node P has a write miss at address A; request data and make P the exclusive owner.
Invalidate	Local cache	Home directory	A	Request to send invalidates to all remote caches that are caching the block at address A.
Invalidate	Home directory	Remote cache	A	Invalidate a shared copy of data at address A.
Fetch	Home directory	Remote cache	A	Fetch the block at address A and send it to its home directory; change the state of A in the remote cache to shared.
Fetch/invalidate	Home directory	Remote cache	A	Fetch the block at address A and send it to its home directory; invalidate the block in the cache.
Data value reply	Home directory	Local cache	D	Return a data value from the home memory.
Data write-back	Remote cache	Home directory	A, D	Write-back a data value for address A.

↗ Hennessy & Patterson, Computer Architecture, 5th edition, page 381.

Individual caches use a state machine similar to the one on slide 208.

messages from
home directory
CPU events

invalid

write miss; send write miss msg

invalidate;
write back block

read miss; send read miss msg

invalidate

modified

read miss; write back data
send read miss msg

fetch; write back data

write hit; send invalidate msg
write miss; send write miss msg

shared
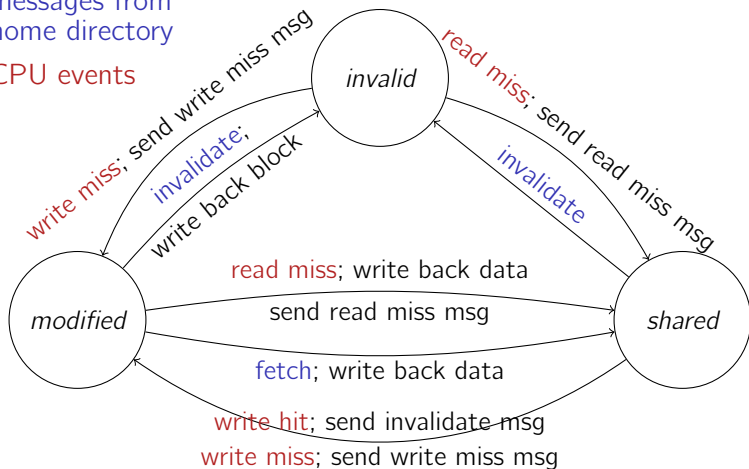
Directory-Based Coherence—State Machine

The **directory** has its own state machine.
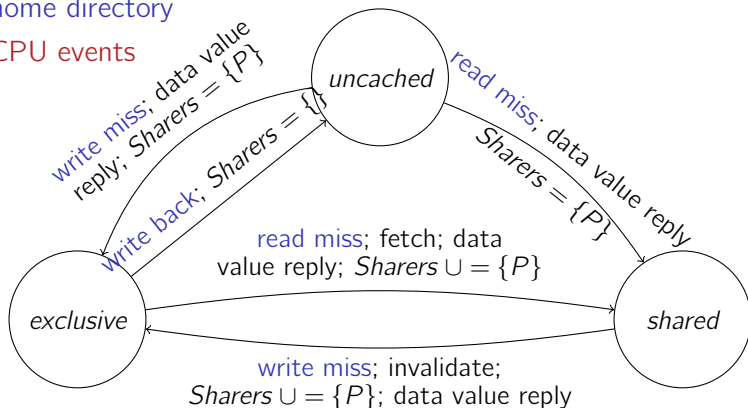
Cache Coherence Cost

Experiment:

- Several threads randomly increment elements of an integer array; Zipfian probability distribution, no synchronization[10].

Intel Nehalem EX; 1.87 GHz; 2 CPUs, 8 cores/CPU.

---

[10]In general, this will yield incorrect counter values.

Cache Coherence Cost

Two types of **coherence misses**:

true sharing miss

- $\rightarrow$ Data shared among processors.
- $\rightarrow$ Often-used mechanism to **communicate** between threads.
- $\rightarrow$ These misses are **unavoidable**.

false sharing miss

- $\rightarrow$ Processors use **different data items**, but the items reside in the **same cache line**.
- $\rightarrow$ Items get invalidated/migrated, even though no data is actually shared.

✎ **How can false sharing misses be avoided?**

NUMA—Non-Uniform Memory Access

Distribution makes memory access **locality-sensitive**.

$\rightarrow$ **Non-Uniform Memory Access (NUMA)**

	bandwidth	latency
①	24.7 GB/s	150 ns
②	10.9 GB/s	185 ns
③	10.9 GB/s	230 ns
③/④[11]	5.3 GB/s	235 ns

↗ Li *et al.* NUMA-Aware Algorithms: The Case of Data Shuffling. *CIDR 2013*

---

[11] ③ with cross traffic along ④.
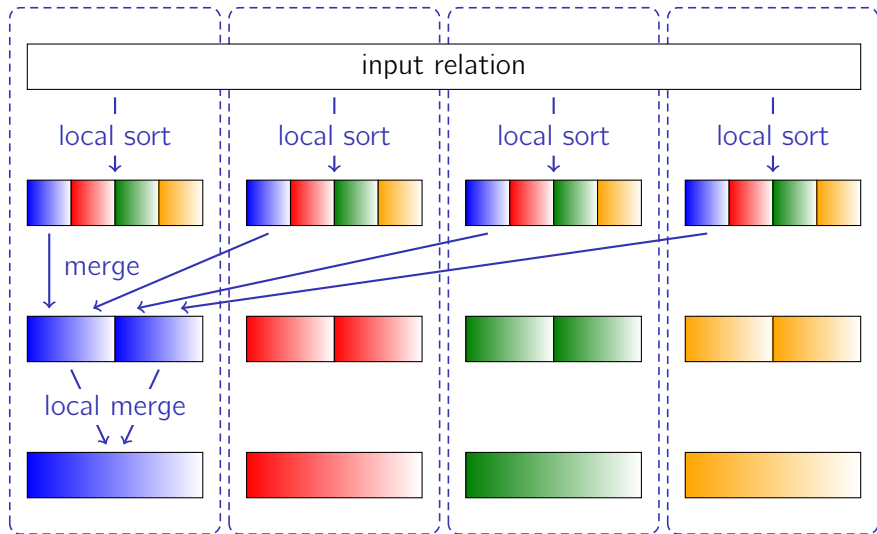
Problem: Merging is **bandwidth-bound**.

$\rightarrow$ Merge multiple runs (from NUMA regions) at once
(Two-way merging would be more CPU-efficient because of SIMD.)

$\rightarrow$ Might need **more instructions**, but brings bandwidth and compute **into balance**.

NUMA Effects in Detail

Bandwidth:

- Single links have **lower bandwidth** than memory controllers.


| Intel Nehalem EX | Intel Sandy Bridge EP |

Task: Find all $\langle r, s \rangle$ in w_R, w_S that satisfy $p(r, s)$.

1. **scan** window, 2. **insert** new tuple, 3. **invalidate** old
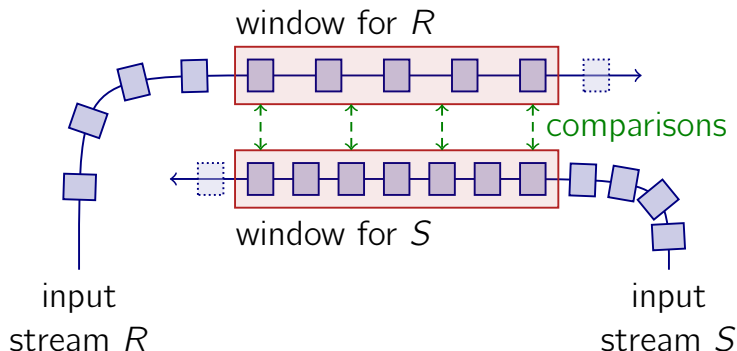
NUMA-Aware Execution?

❶ **bandwidth** bottlenecks

❷ **long-distance** communication

❸ **centralized** coordination and memory

$\rightarrow$ Parallel, but not NUMA-aware.

Handshake Join:

Streams flow by in **opposite directions**
Compare tuples when they **meet**

Data flow representation $\rightarrow$ **parallelization**:

- **No bandwidth bottleneck** ❶ ✓
- Communication/synchronization stays **local** ❷ ✓

Synchronization

Coordination can now be done **autonomously**

$R \rightarrow$ ... core 1, core 2, core 3, core 4, core 5 ... $\leftarrow S$

- no more centralized coordination ❸ ✓
- Autonomous **load balancing**
- **Lock-free message queues** between neighbors

Beyond 48 Cores... (FPGA-based simulation)