

Formalizing Data Locality in Task-Parallel Applications

Germán Ceballos

Erik Hagersten

David Black-Schaffer

TAPEMS@ICA3PP '16

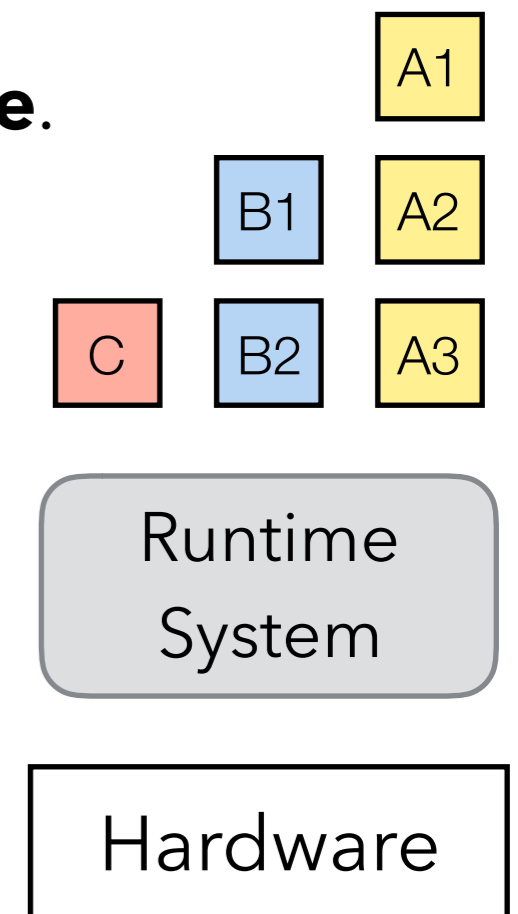
UppsalaUniversity



Tasks in the Multicore Era

Why Task-Based Programming?

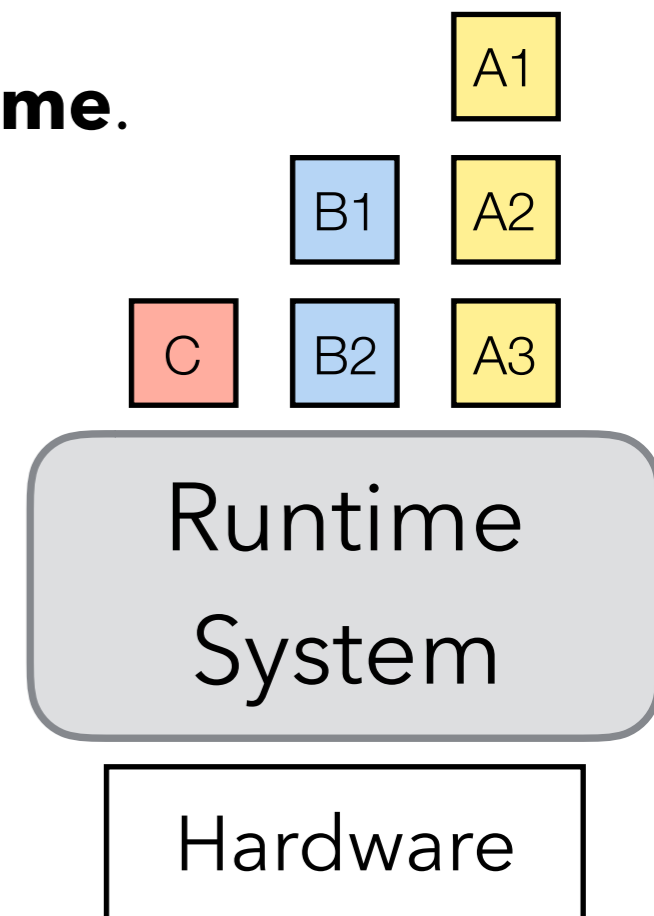
- Easy to express parallelism (vs threads).
- Delegates many **programmer** decisions to the **runtime**.



Tasks in the Multicore Era

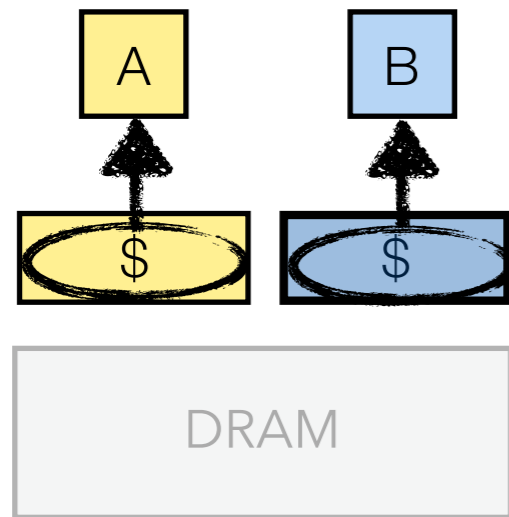
Why Task-Based Programming?

- Easy to express parallelism (vs **threads**).
- Delegates many **programmer** decisions to the **runtime**.
- Allows runtime optimizations:
 - Data Locality
 - Memory Footprint
 - Energy Efficiency

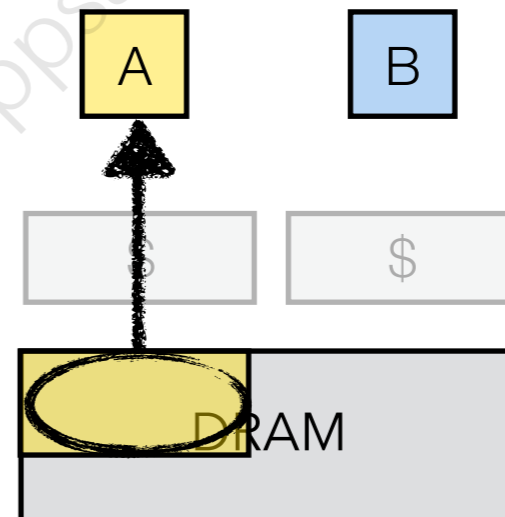


The Inconvenient Truths

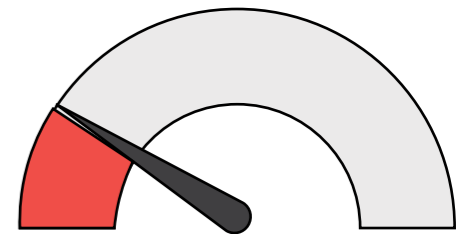
- 1 Caches are **critical** for **performance**



Faster!

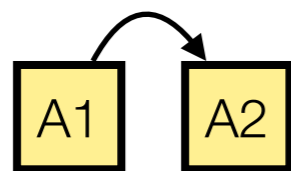
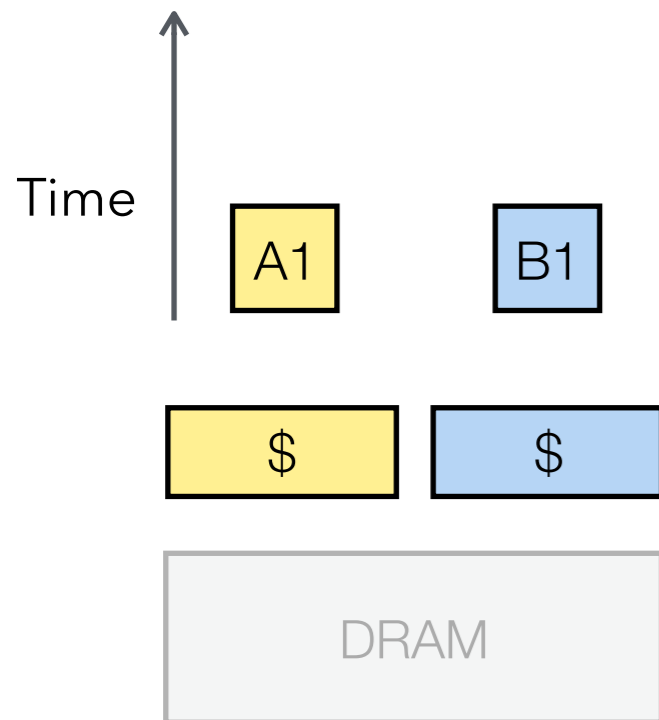


Slower



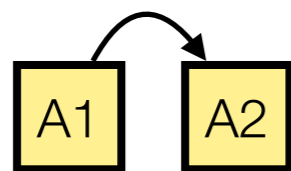
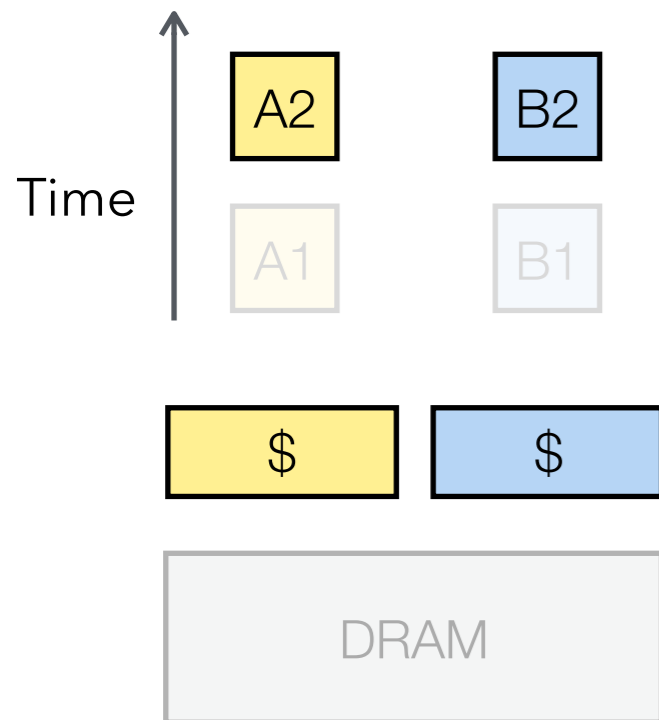
The Inconvenient Truths

- 2 Schedule **changes** how tasks **reuse** data



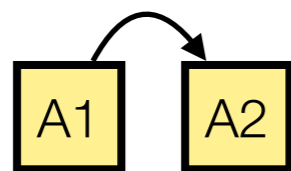
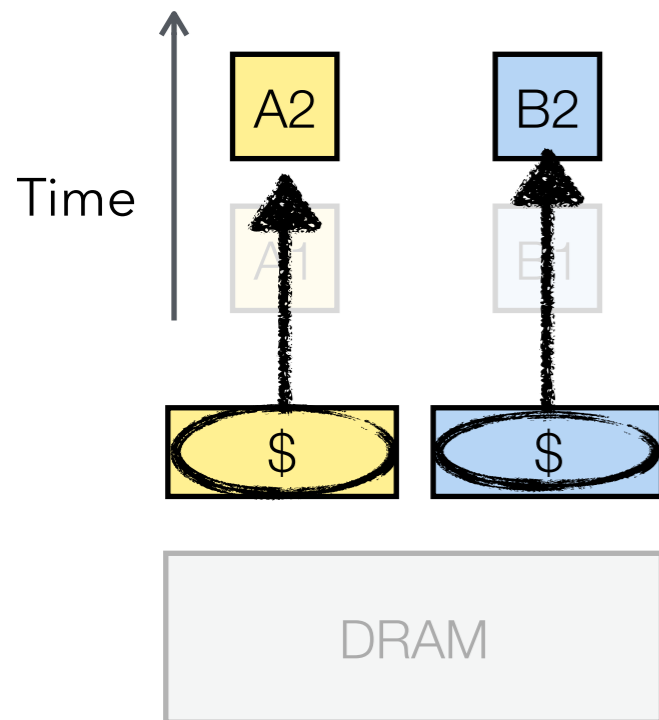
The Inconvenient Truths

- 2 Schedule **changes** how tasks **reuse** data



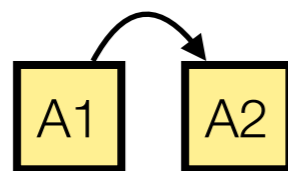
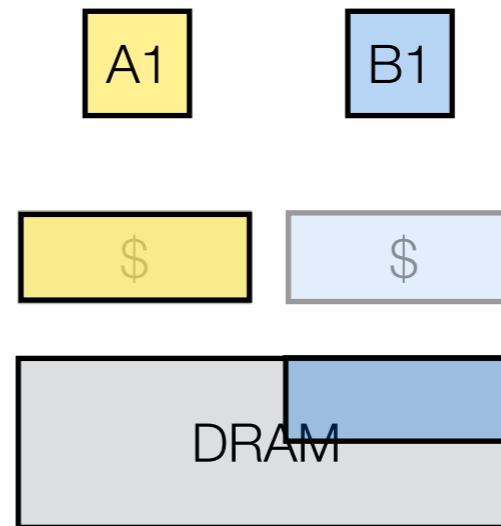
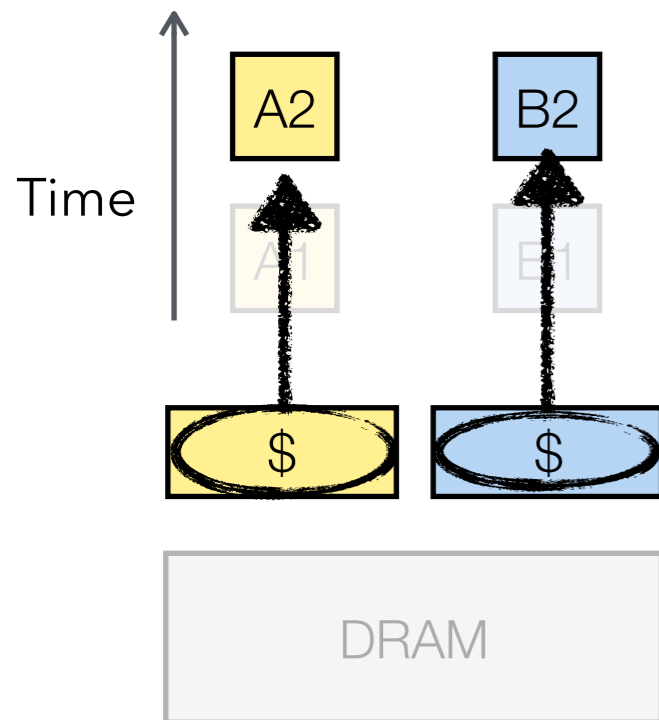
The Inconvenient Truths

- 2 Schedule **changes** how tasks **reuse** data



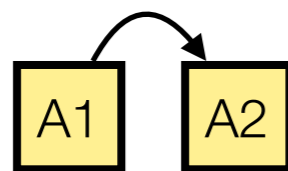
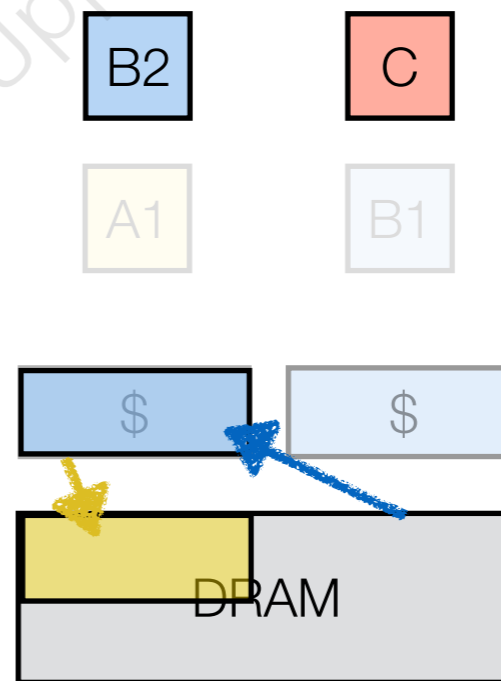
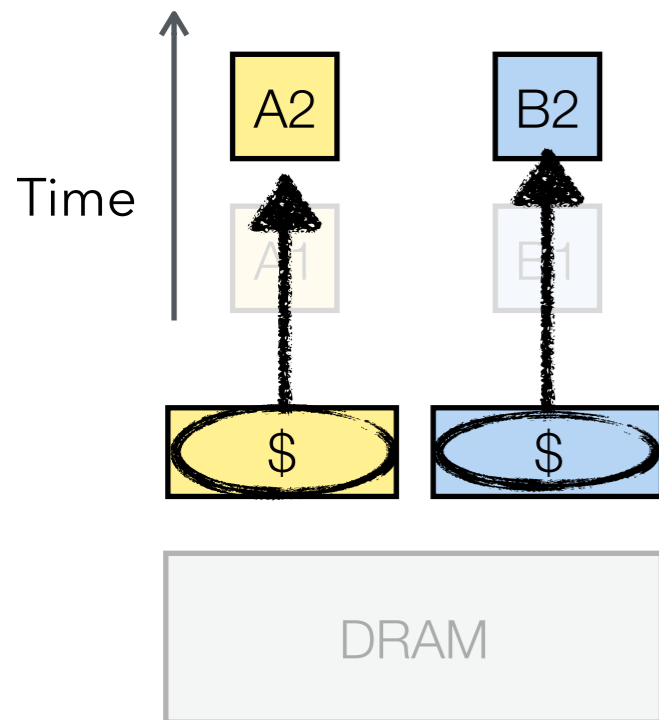
The Inconvenient Truths

- 2 Schedule **changes** how tasks **reuse** data



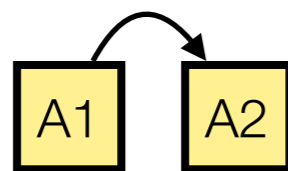
The Inconvenient Truths

2 Schedule **changes** how tasks **reuse** data



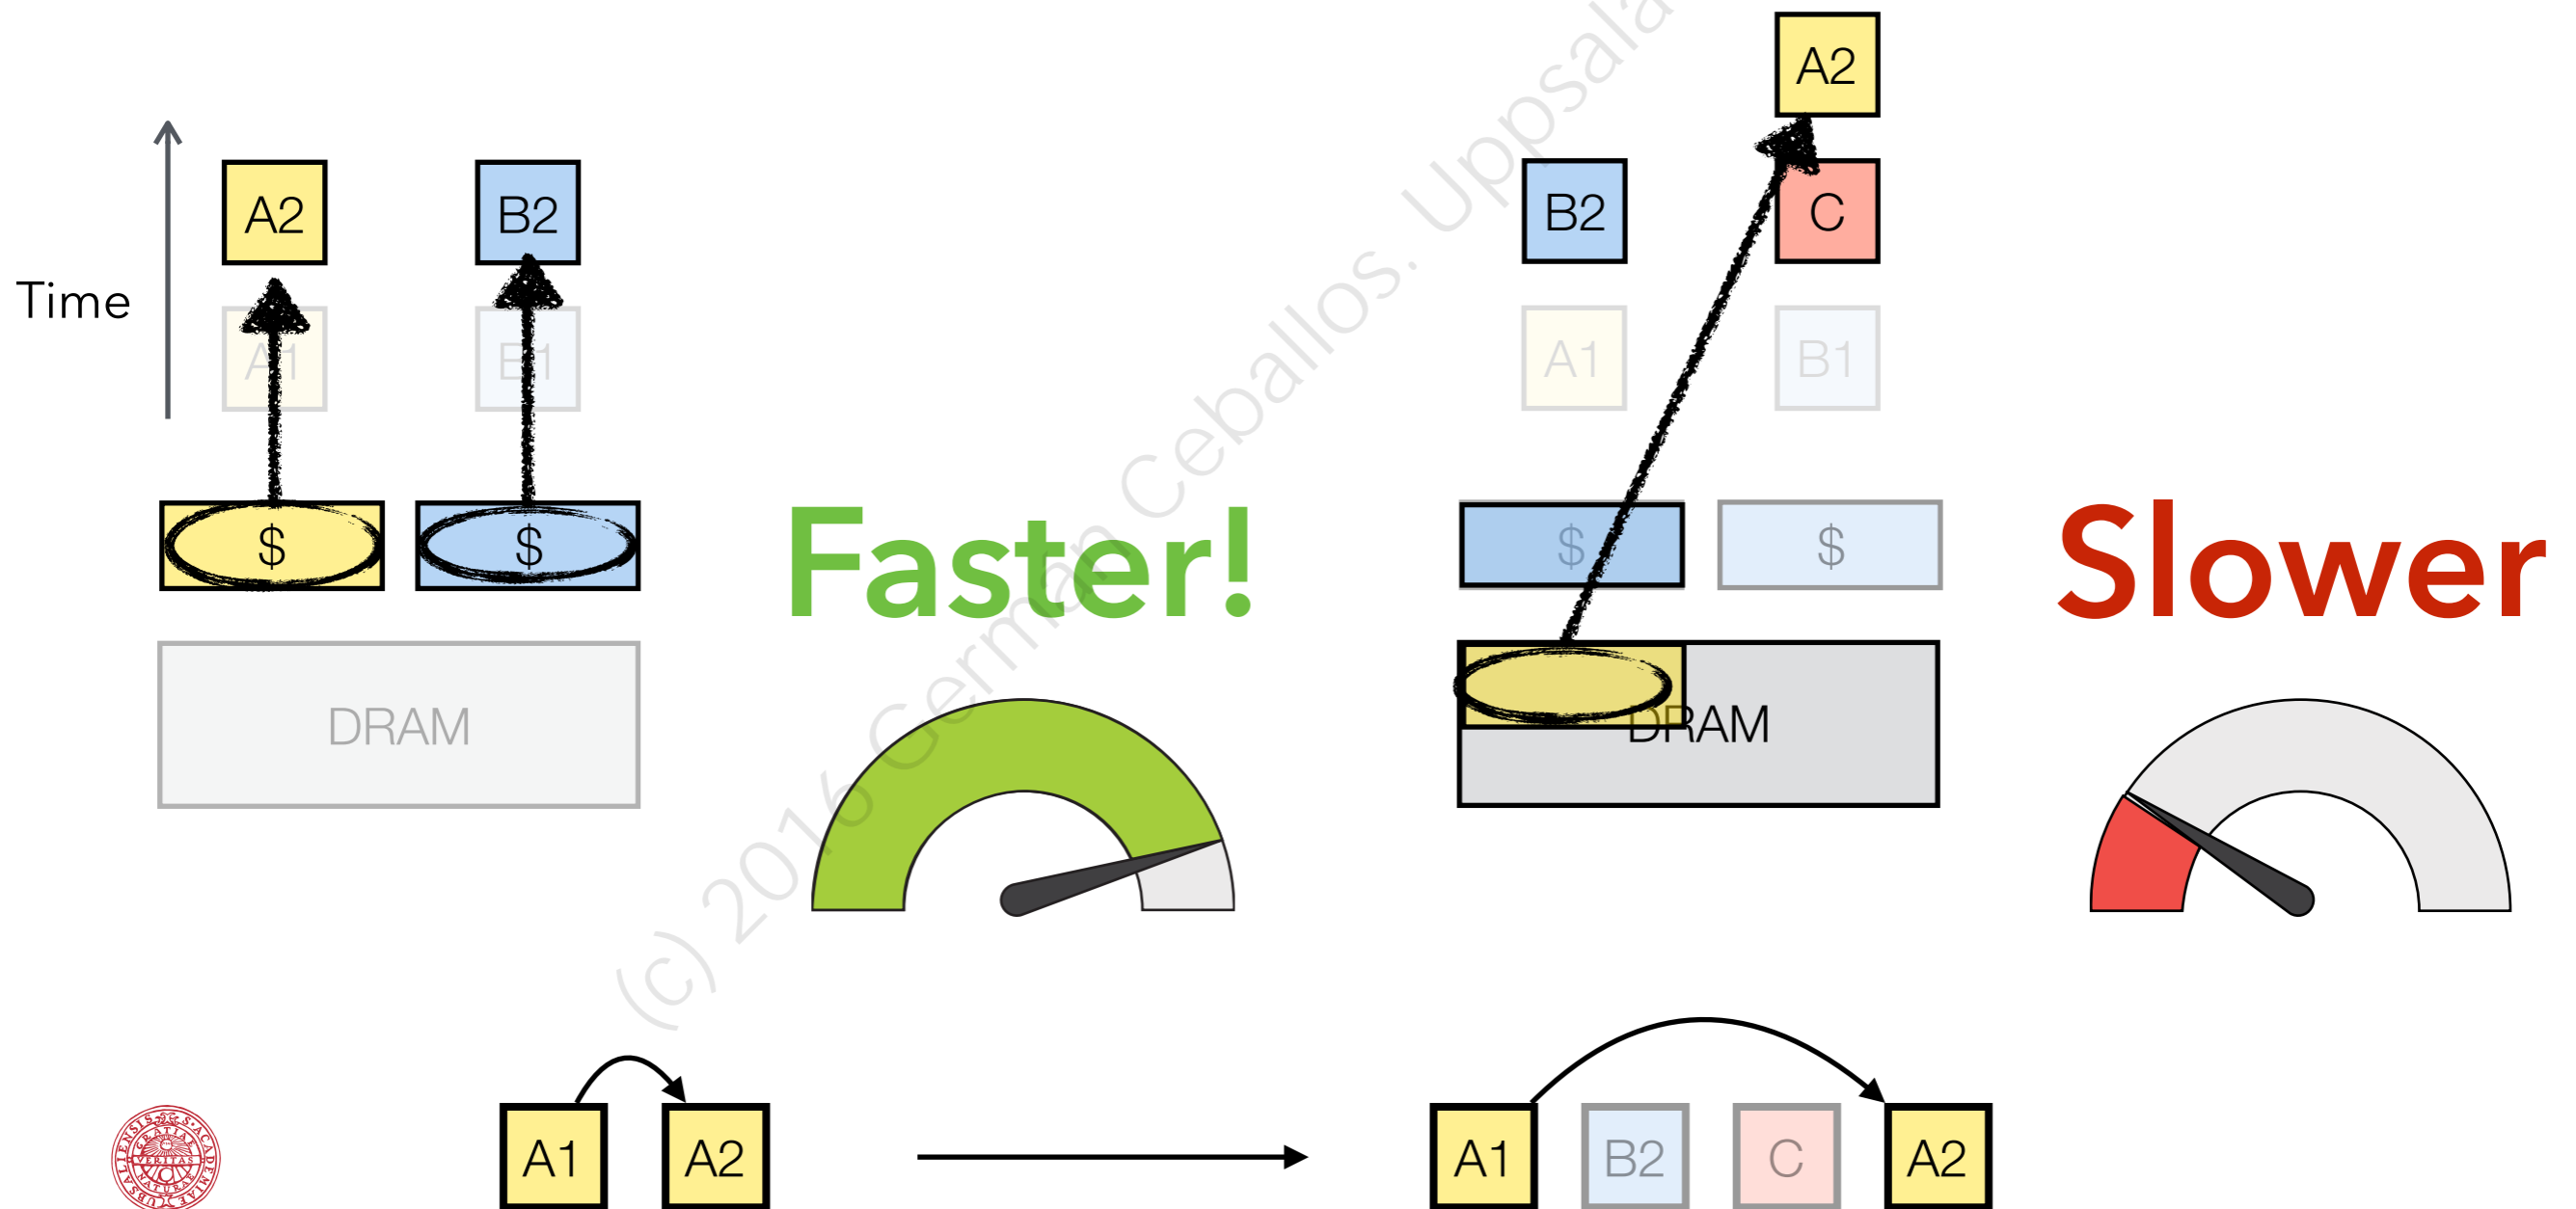
The Inconvenient Truths

2 Schedule **changes** how tasks **reuse** data



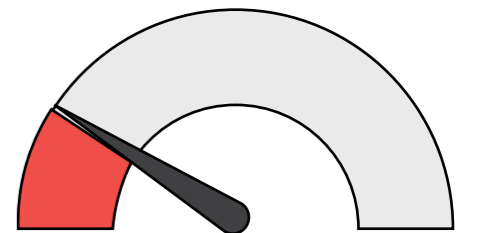
The Inconvenient Truths

2 Schedule **changes** how tasks **reuse** data



Faster!

Slower



The Inconvenient Truths

- 1 Caches are **critical** for **performance**
- 2 Schedule **changes** how tasks **reuse** data

≠ schedules → **≠ reuses** → **≠ performance**

What if we could **predict** effects of **scheduling** in **memory**?

(just from looking at one schedule)



The StatTask Model

Our Contributions

- A **formal description** of the task-based execution model both as:
 - Sequence of memory accesses.
 - Sequence of tasks.
 - Equivalence between them.
- **StatTask: new** Statistical Cache Model
 - **What?: Predicts** memory (cache) behaviour for arbitrary schedules
 - **Accurate:** compared to measured results
 - **Flexible:** predicts from a **single** profile
 - **Robust:** predicts similar behaviour for inputs of roughly same size



The StatTask Model

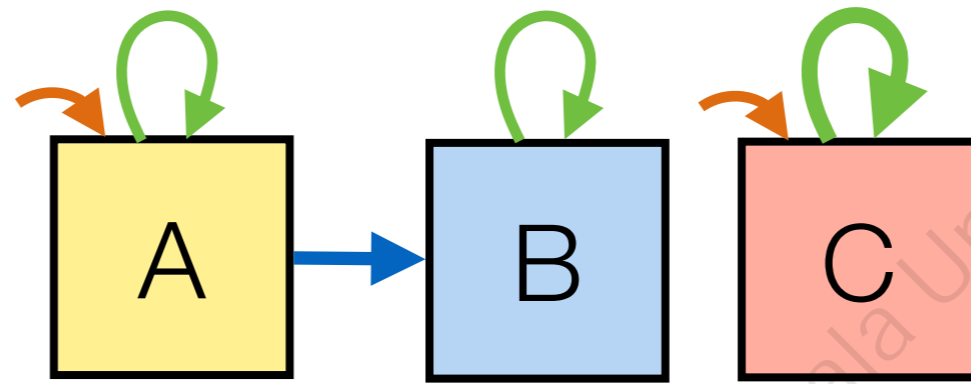
Our Contributions

- A **formal description** of the task-based execution model both as:
 - Sequence of memory accesses.
 - Sequence of tasks.
 - Equivalence between them.
- **StatTask: new** Statistical Cache Model
 - **What?: Predicts** memory (cache) behaviour for arbitrary schedules
 - **Accurate:** compared to measured results
 - **Flexible:** predicts from a **single** profile
 - **Robust:** predicts similar behaviour for inputs of any size



Why?

Potential: Reuse Classification



Type

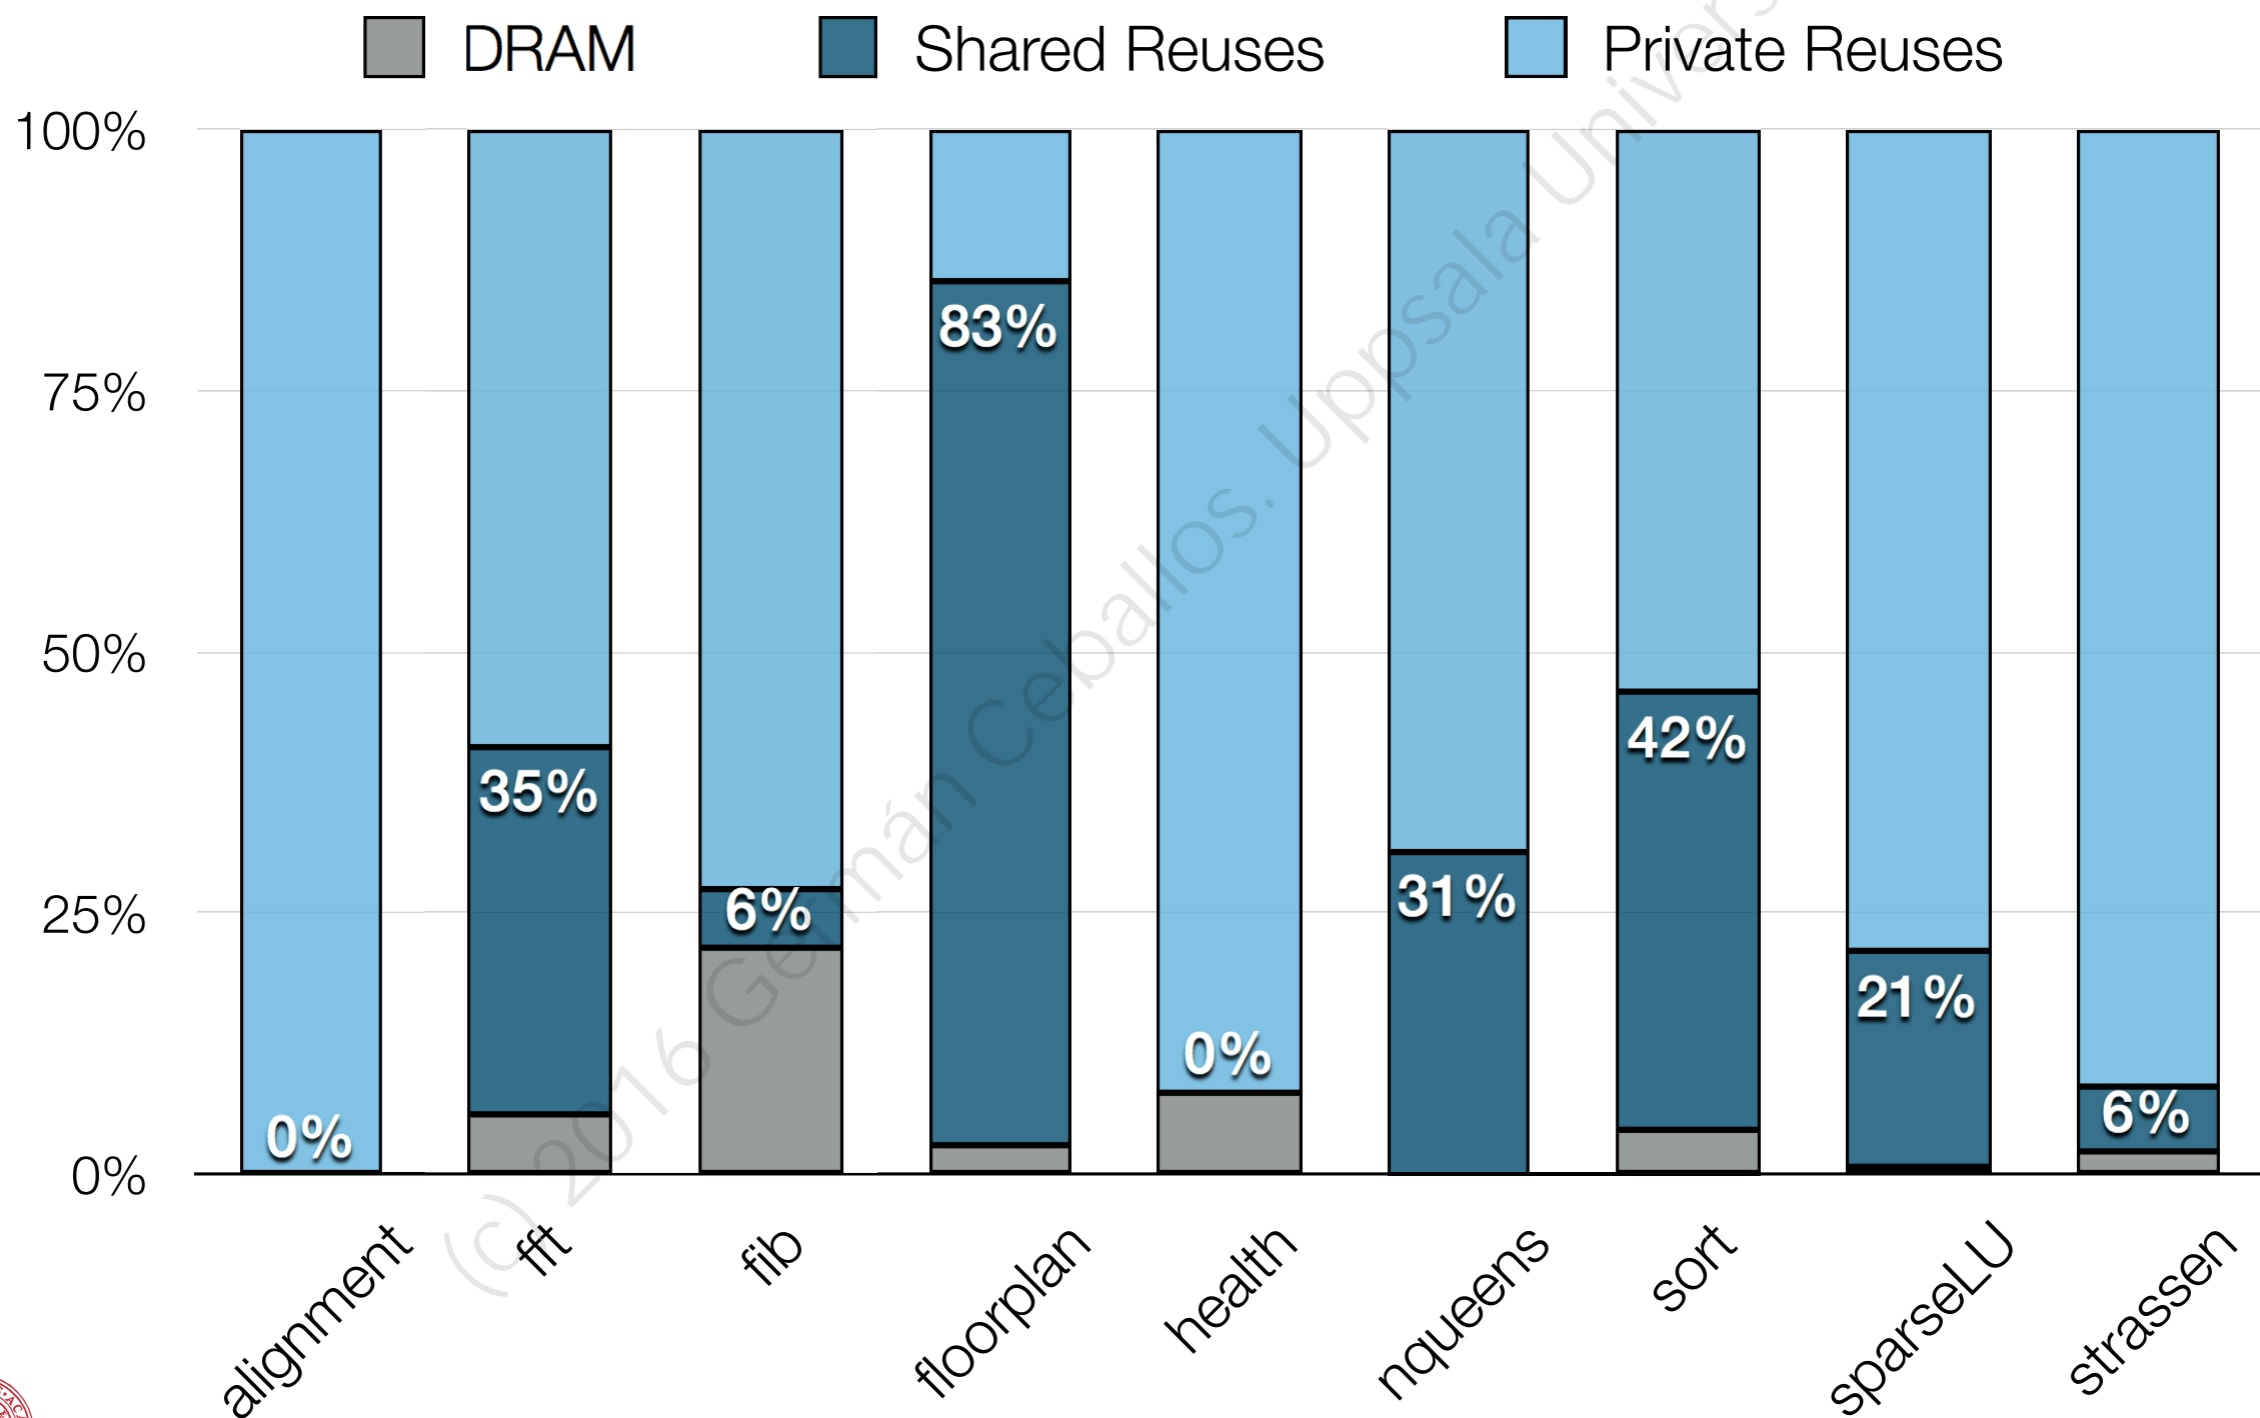
Where?

Depending on

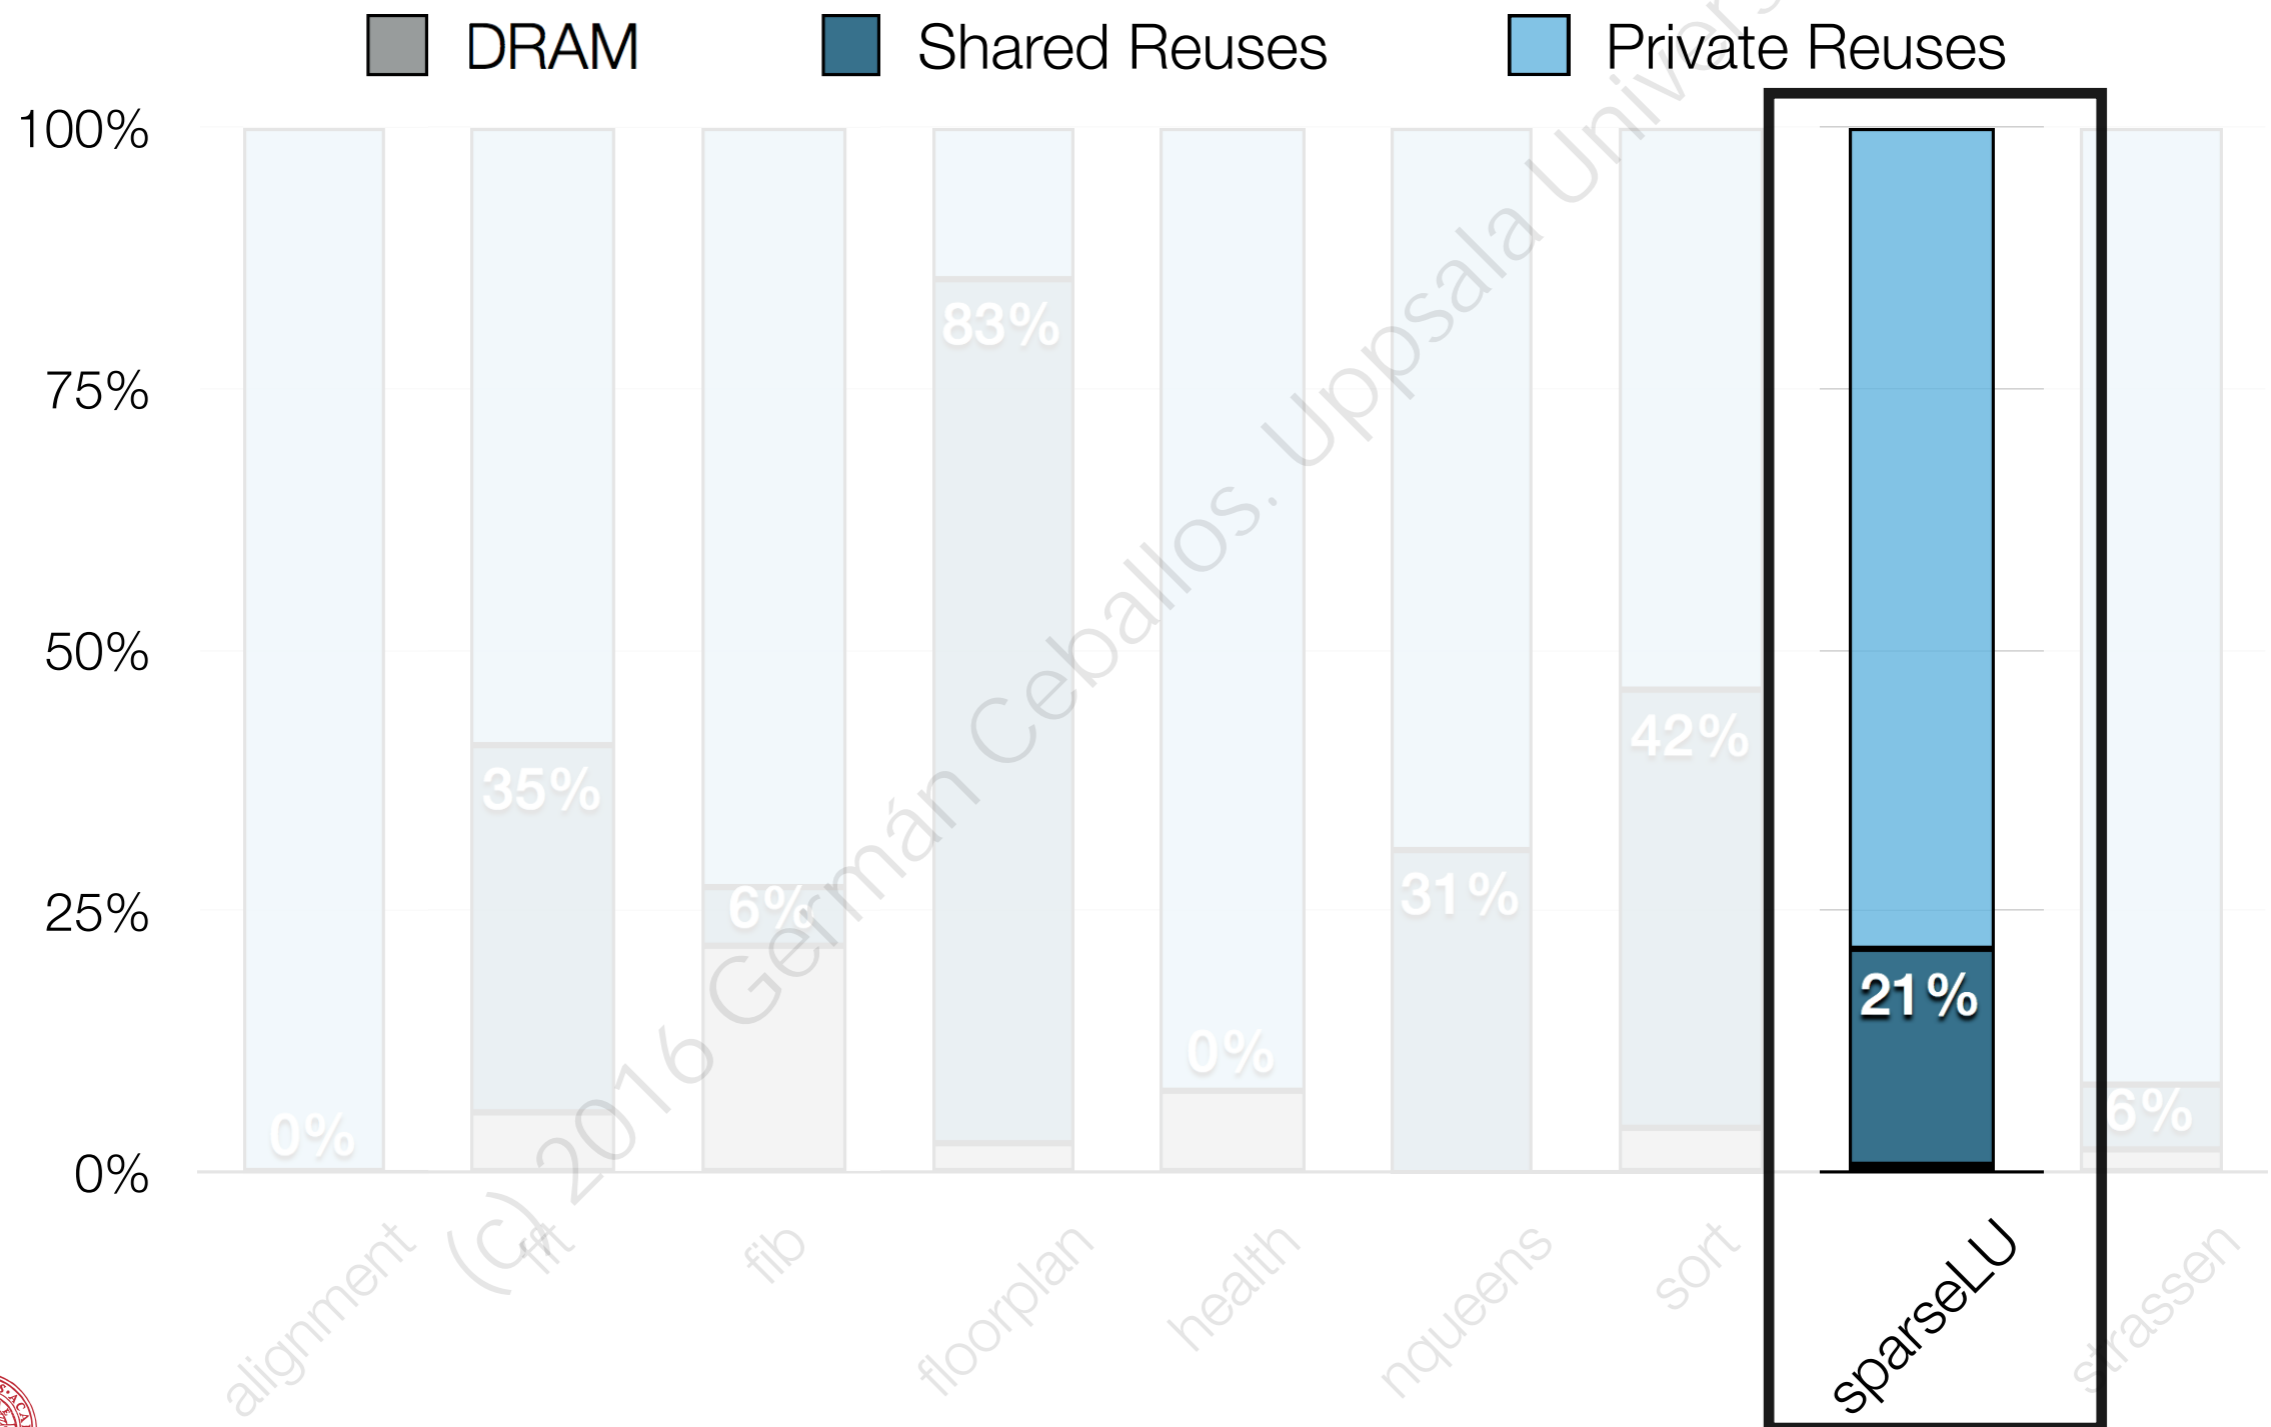
DRAM Access	DRAM	Data Set Size
Private Reuse	Cache/DRAM	Memory Access Pattern Cache Size
Shared Reuse	Cache/DRAM	Schedule



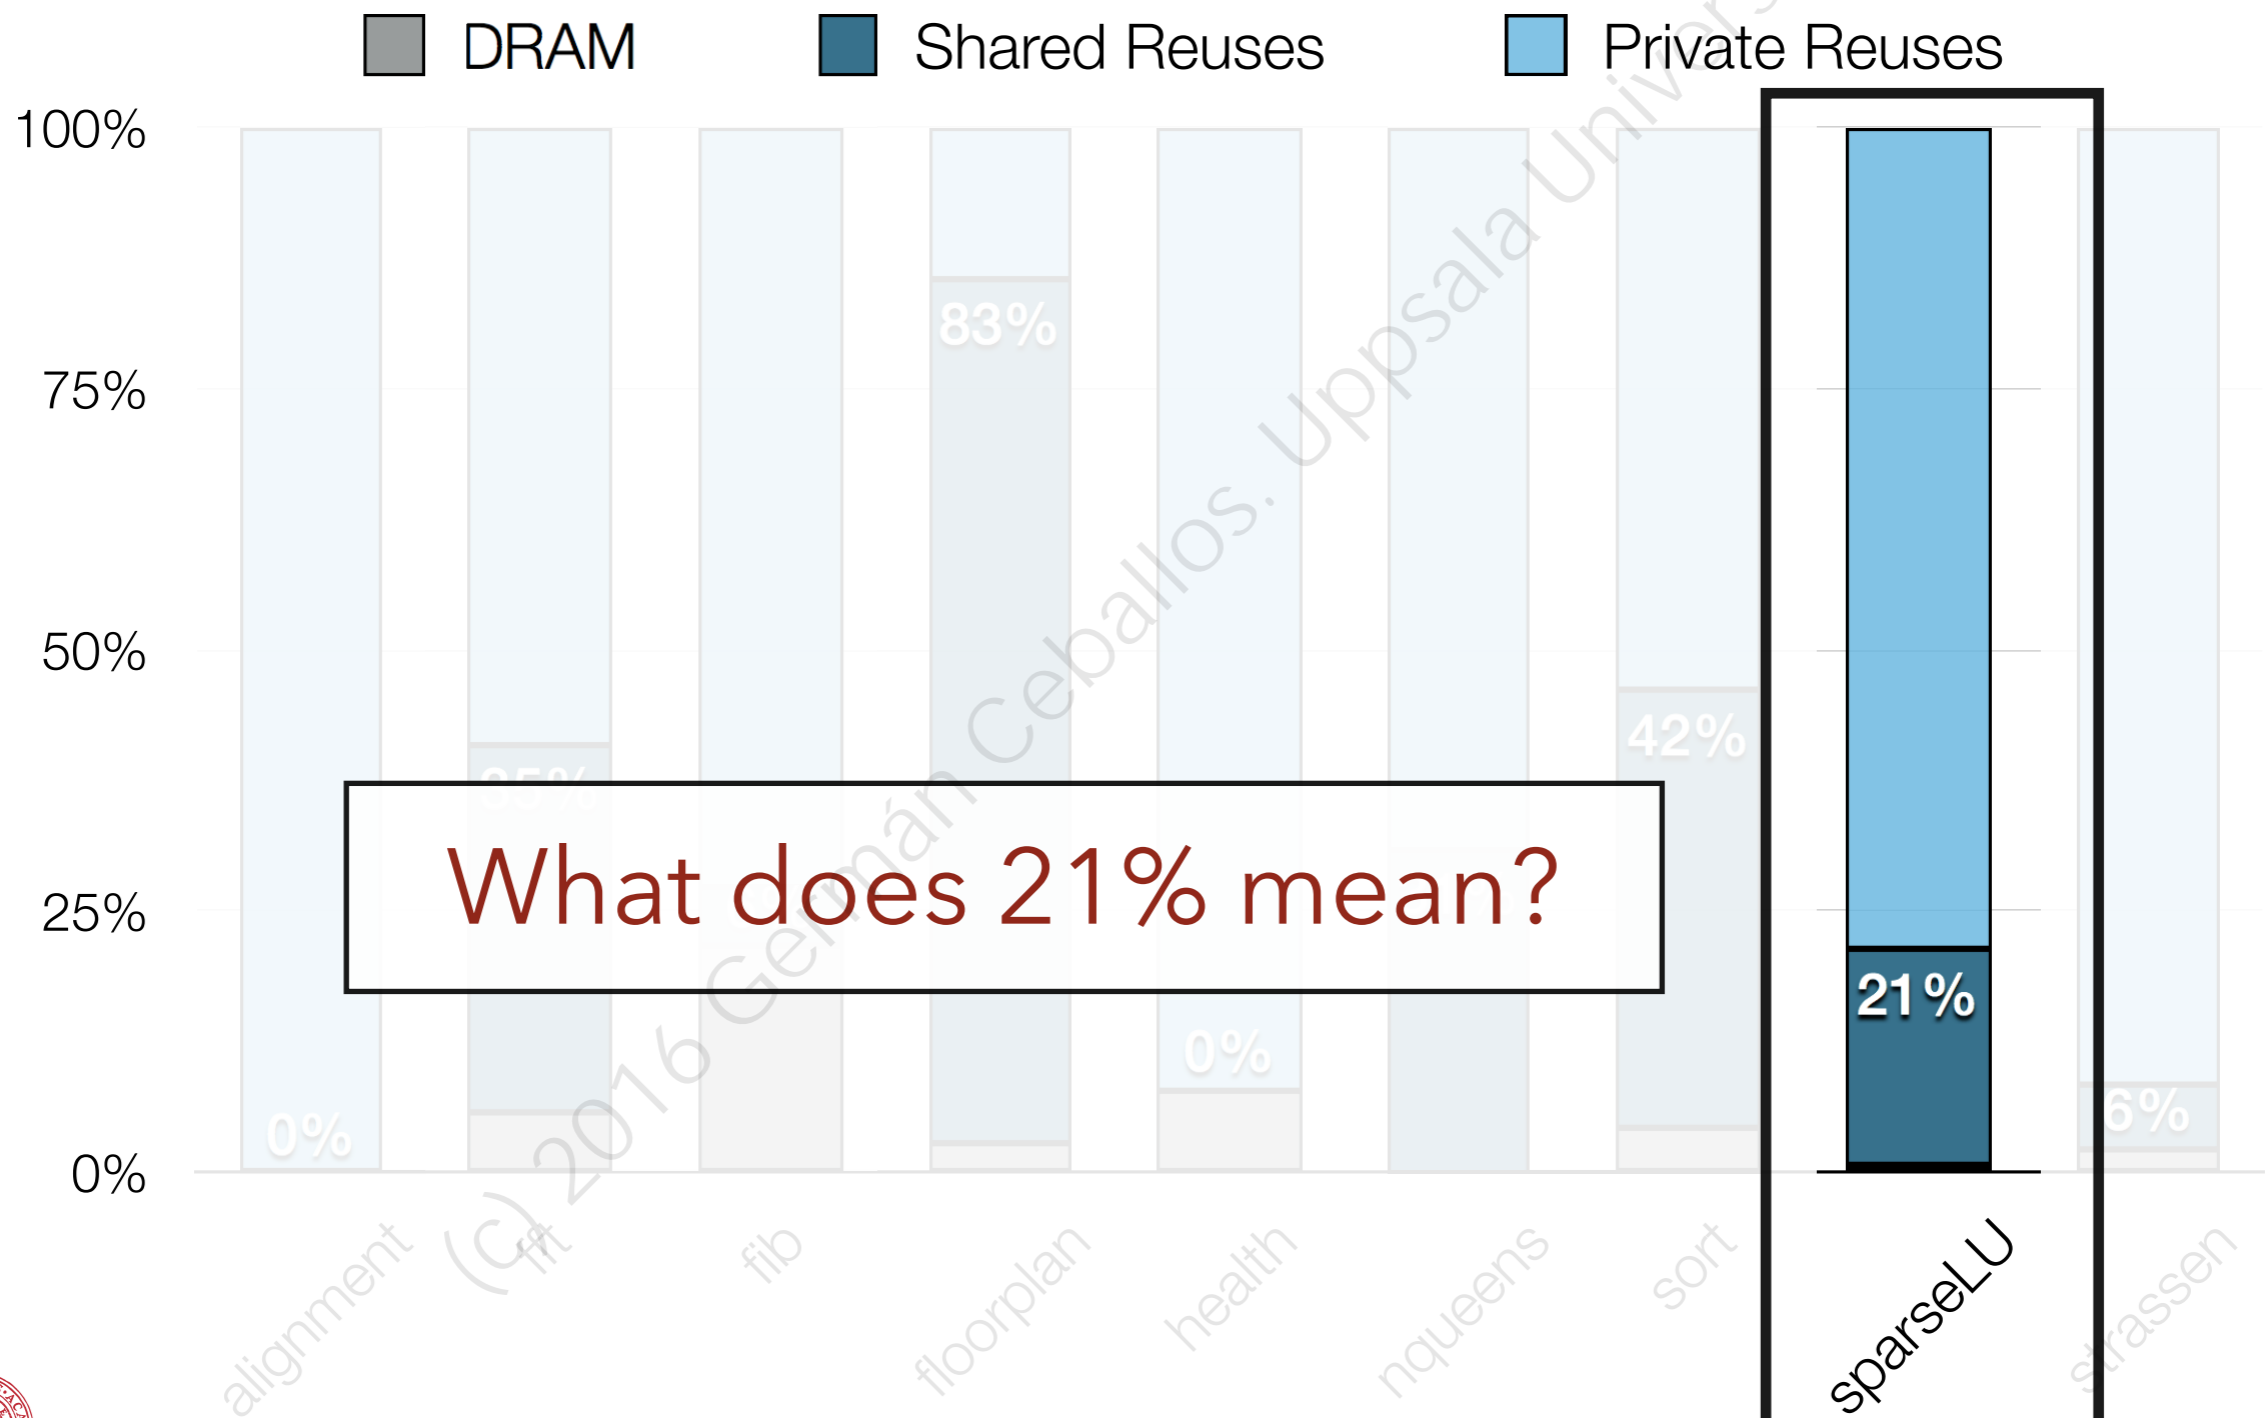
BOTS Reuse Potential



BOTS Reuse Potential

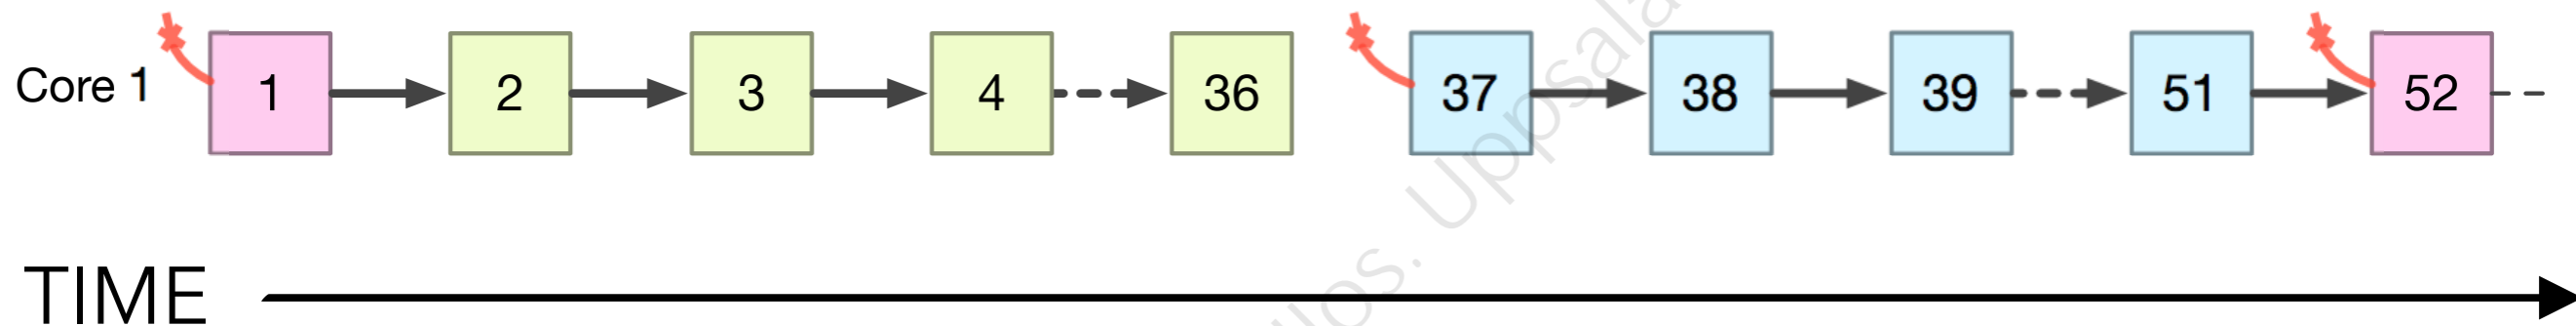


BOTS Reuse Potential

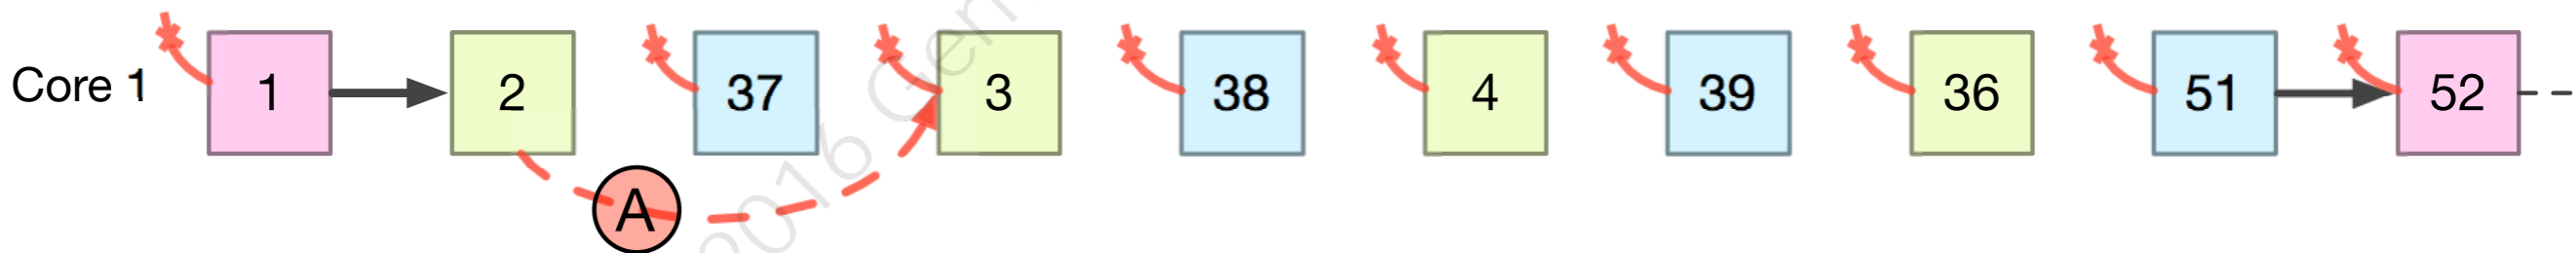


SparseLU Schedules (ST)

Good Schedule



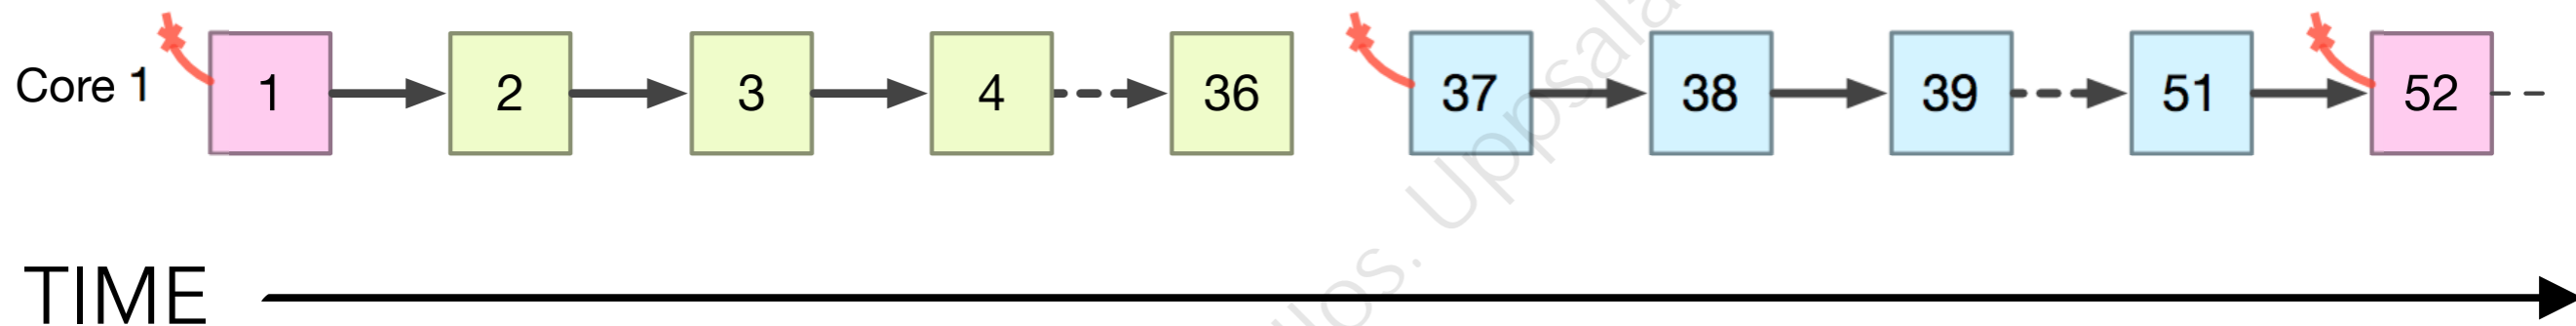
Bad Schedule



SparseLU Schedules (ST)

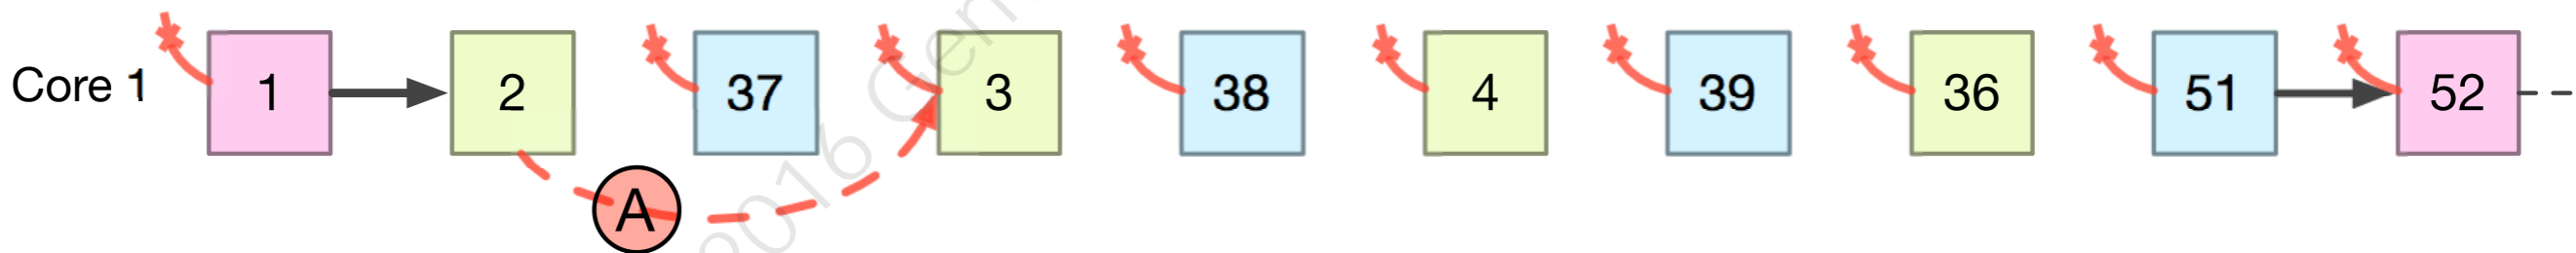
Good Schedule

Miss Ratio: 0,27%



Bad Schedule

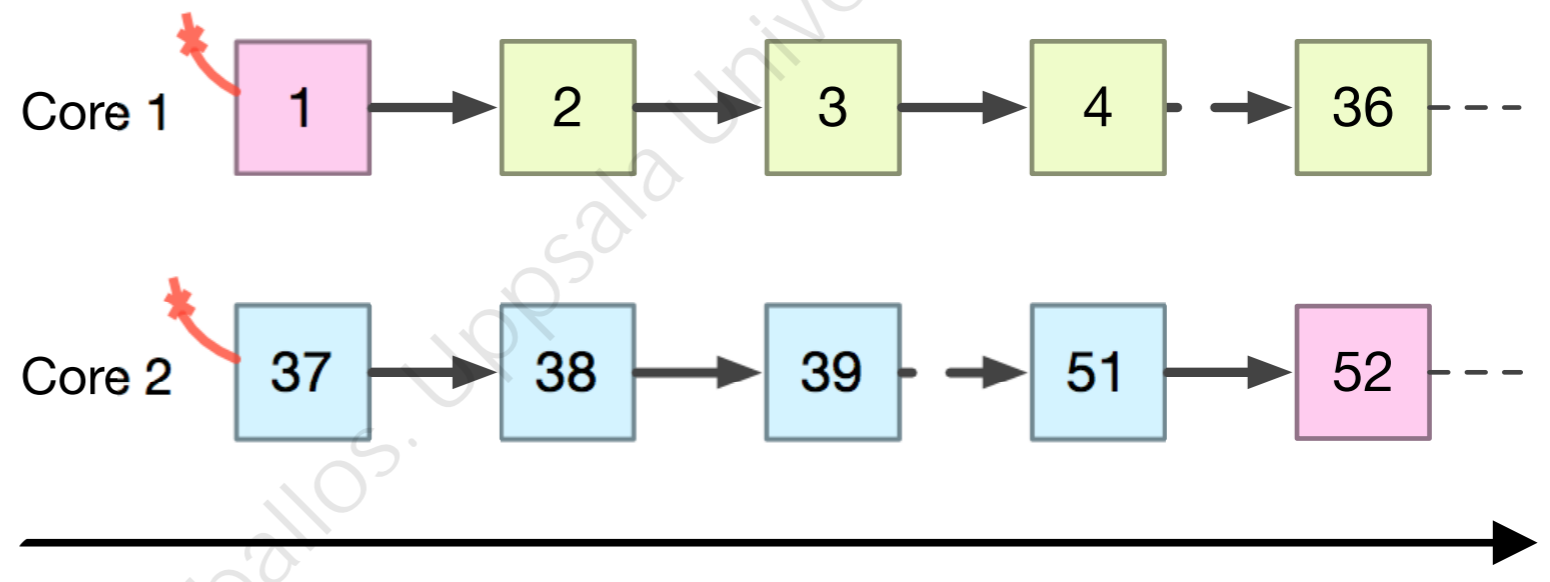
Miss Ratio: 8,01%



SparseLU Schedules (MT)

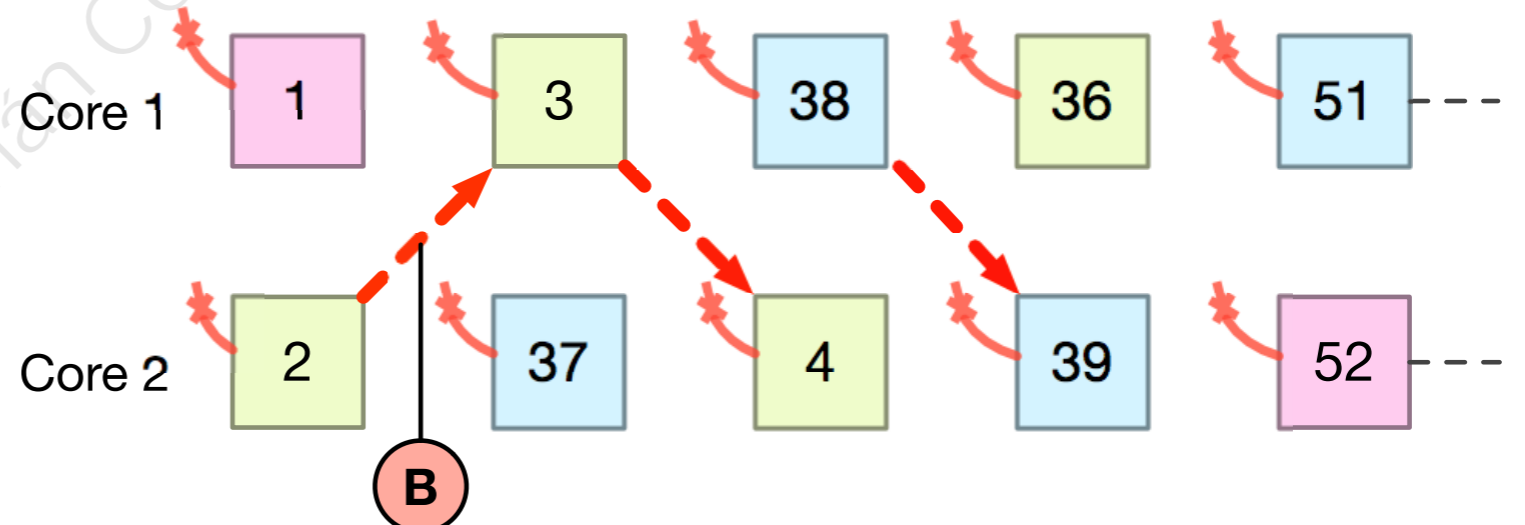
Good Schedule

Miss Ratio: 0,04%



Bad Schedule

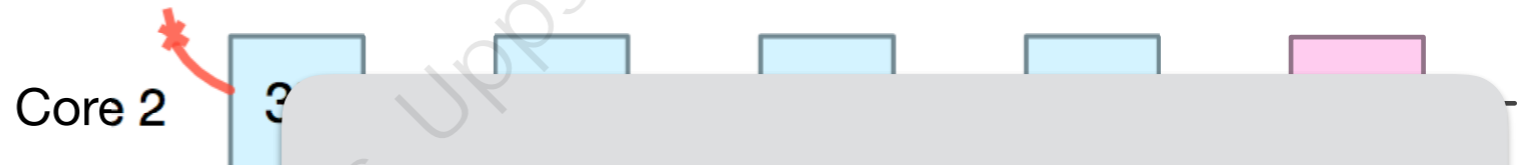
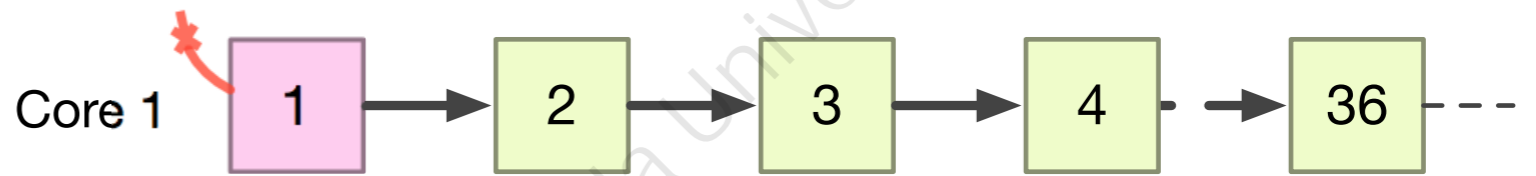
Miss Ratio: 8,01%



SparseLU Schedules (MT)

Good Schedule

Miss Ratio: 0,04%

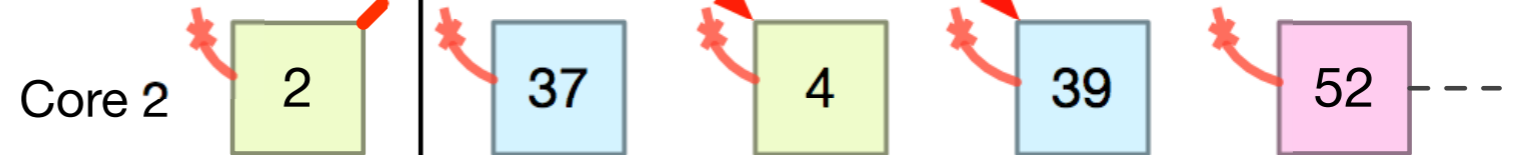
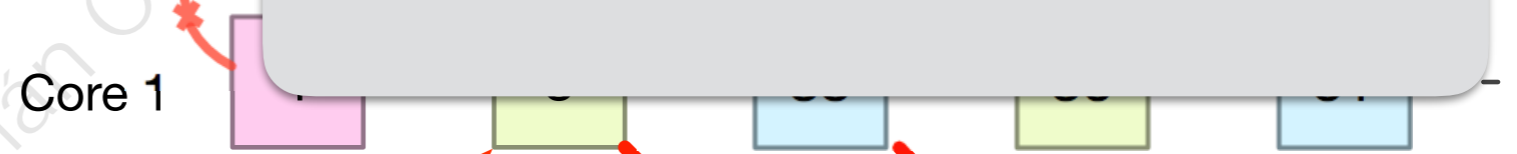


TIME

StatTask can predict **good** or **bad** cache behaviour

Bad Schedule

Miss Ratio: 8,01%



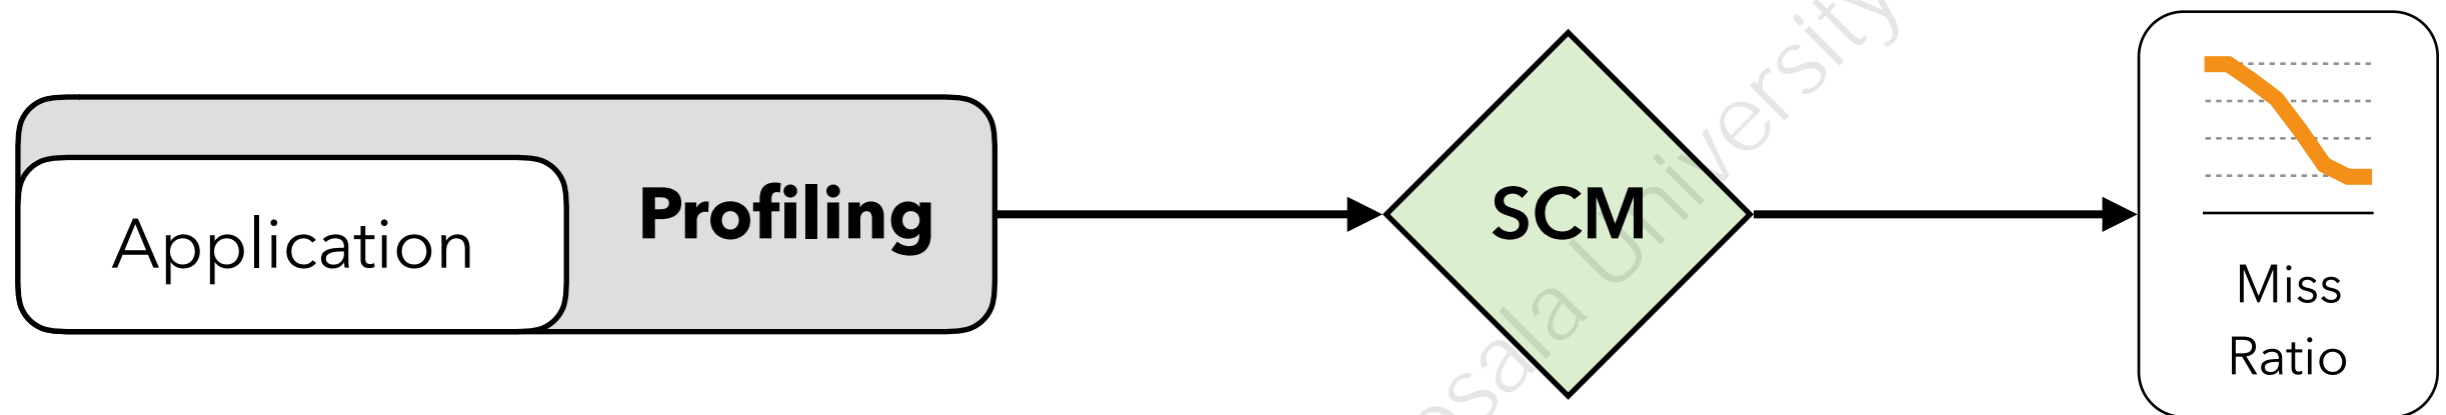
B

(c) 2016 Germán Ceballos. Uppsala University

How can we study data locality?

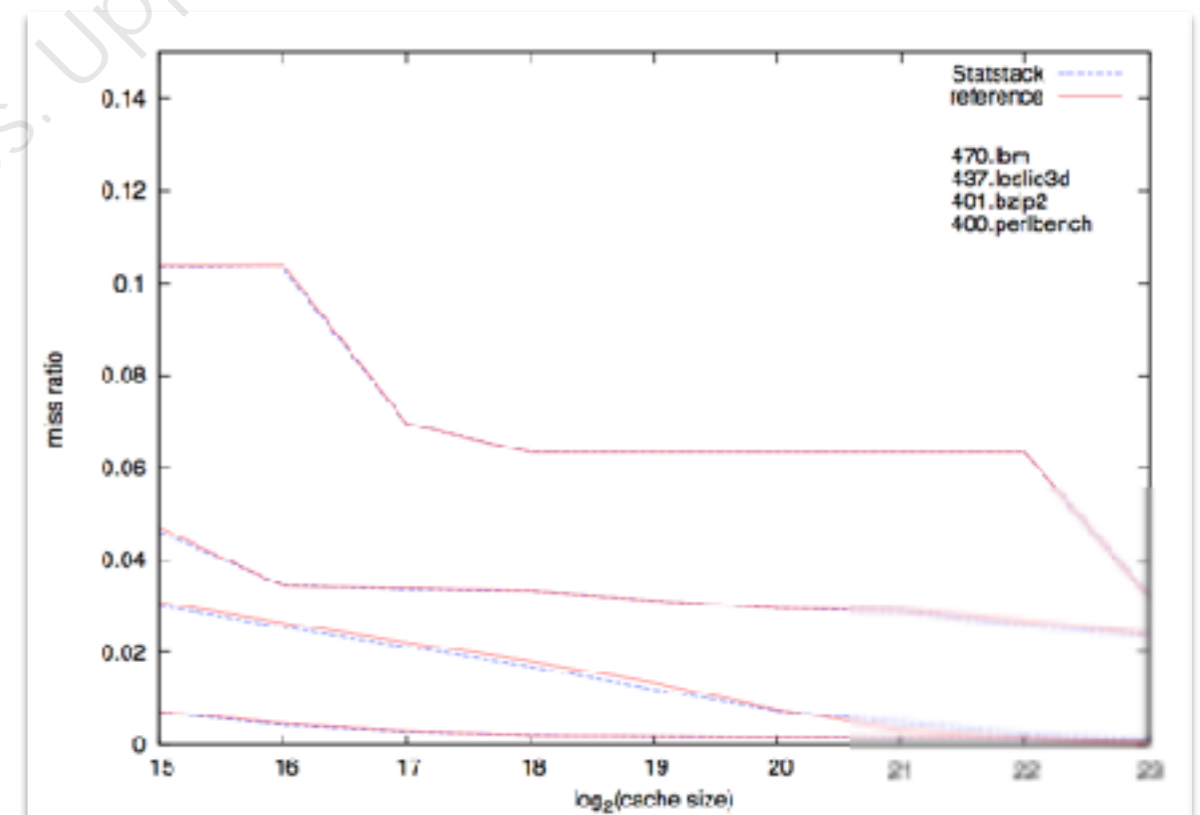


Statistical Cache Models*



Why SCM?

- Fast + Low Overhead
- Accurate
- Just a single profile needed
- Flexible (different cache sizes)



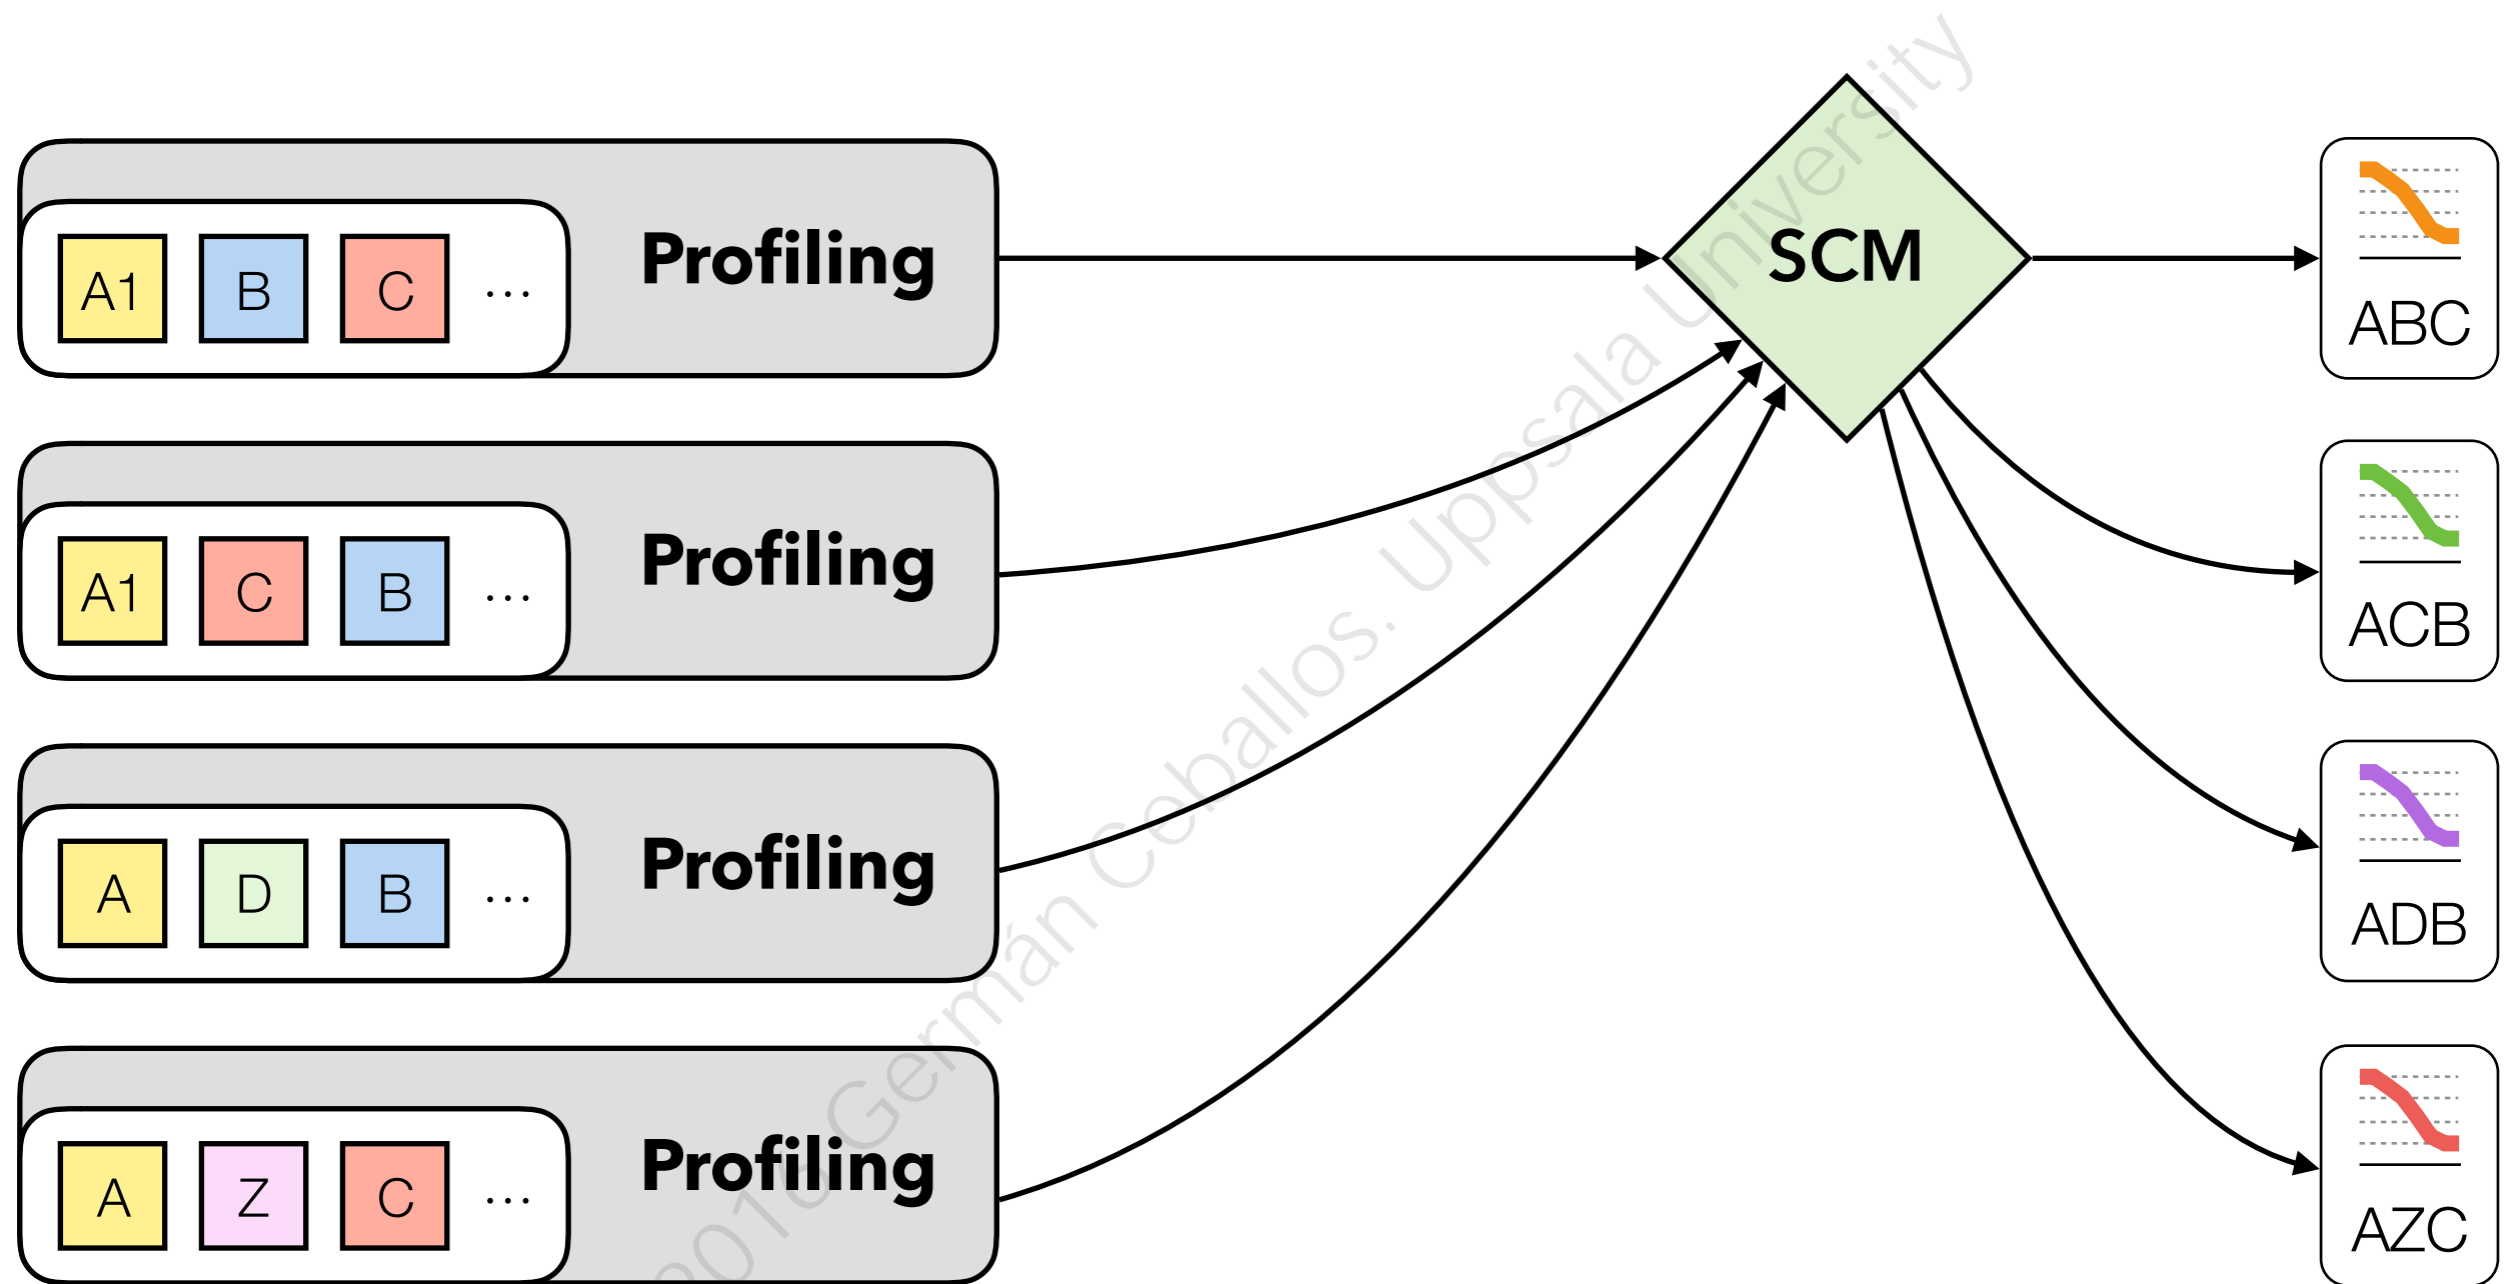
* E. Berg and E. Hagersten. 2004.

StatCache: a probabilistic approach to efficient and accurate data locality analysis.

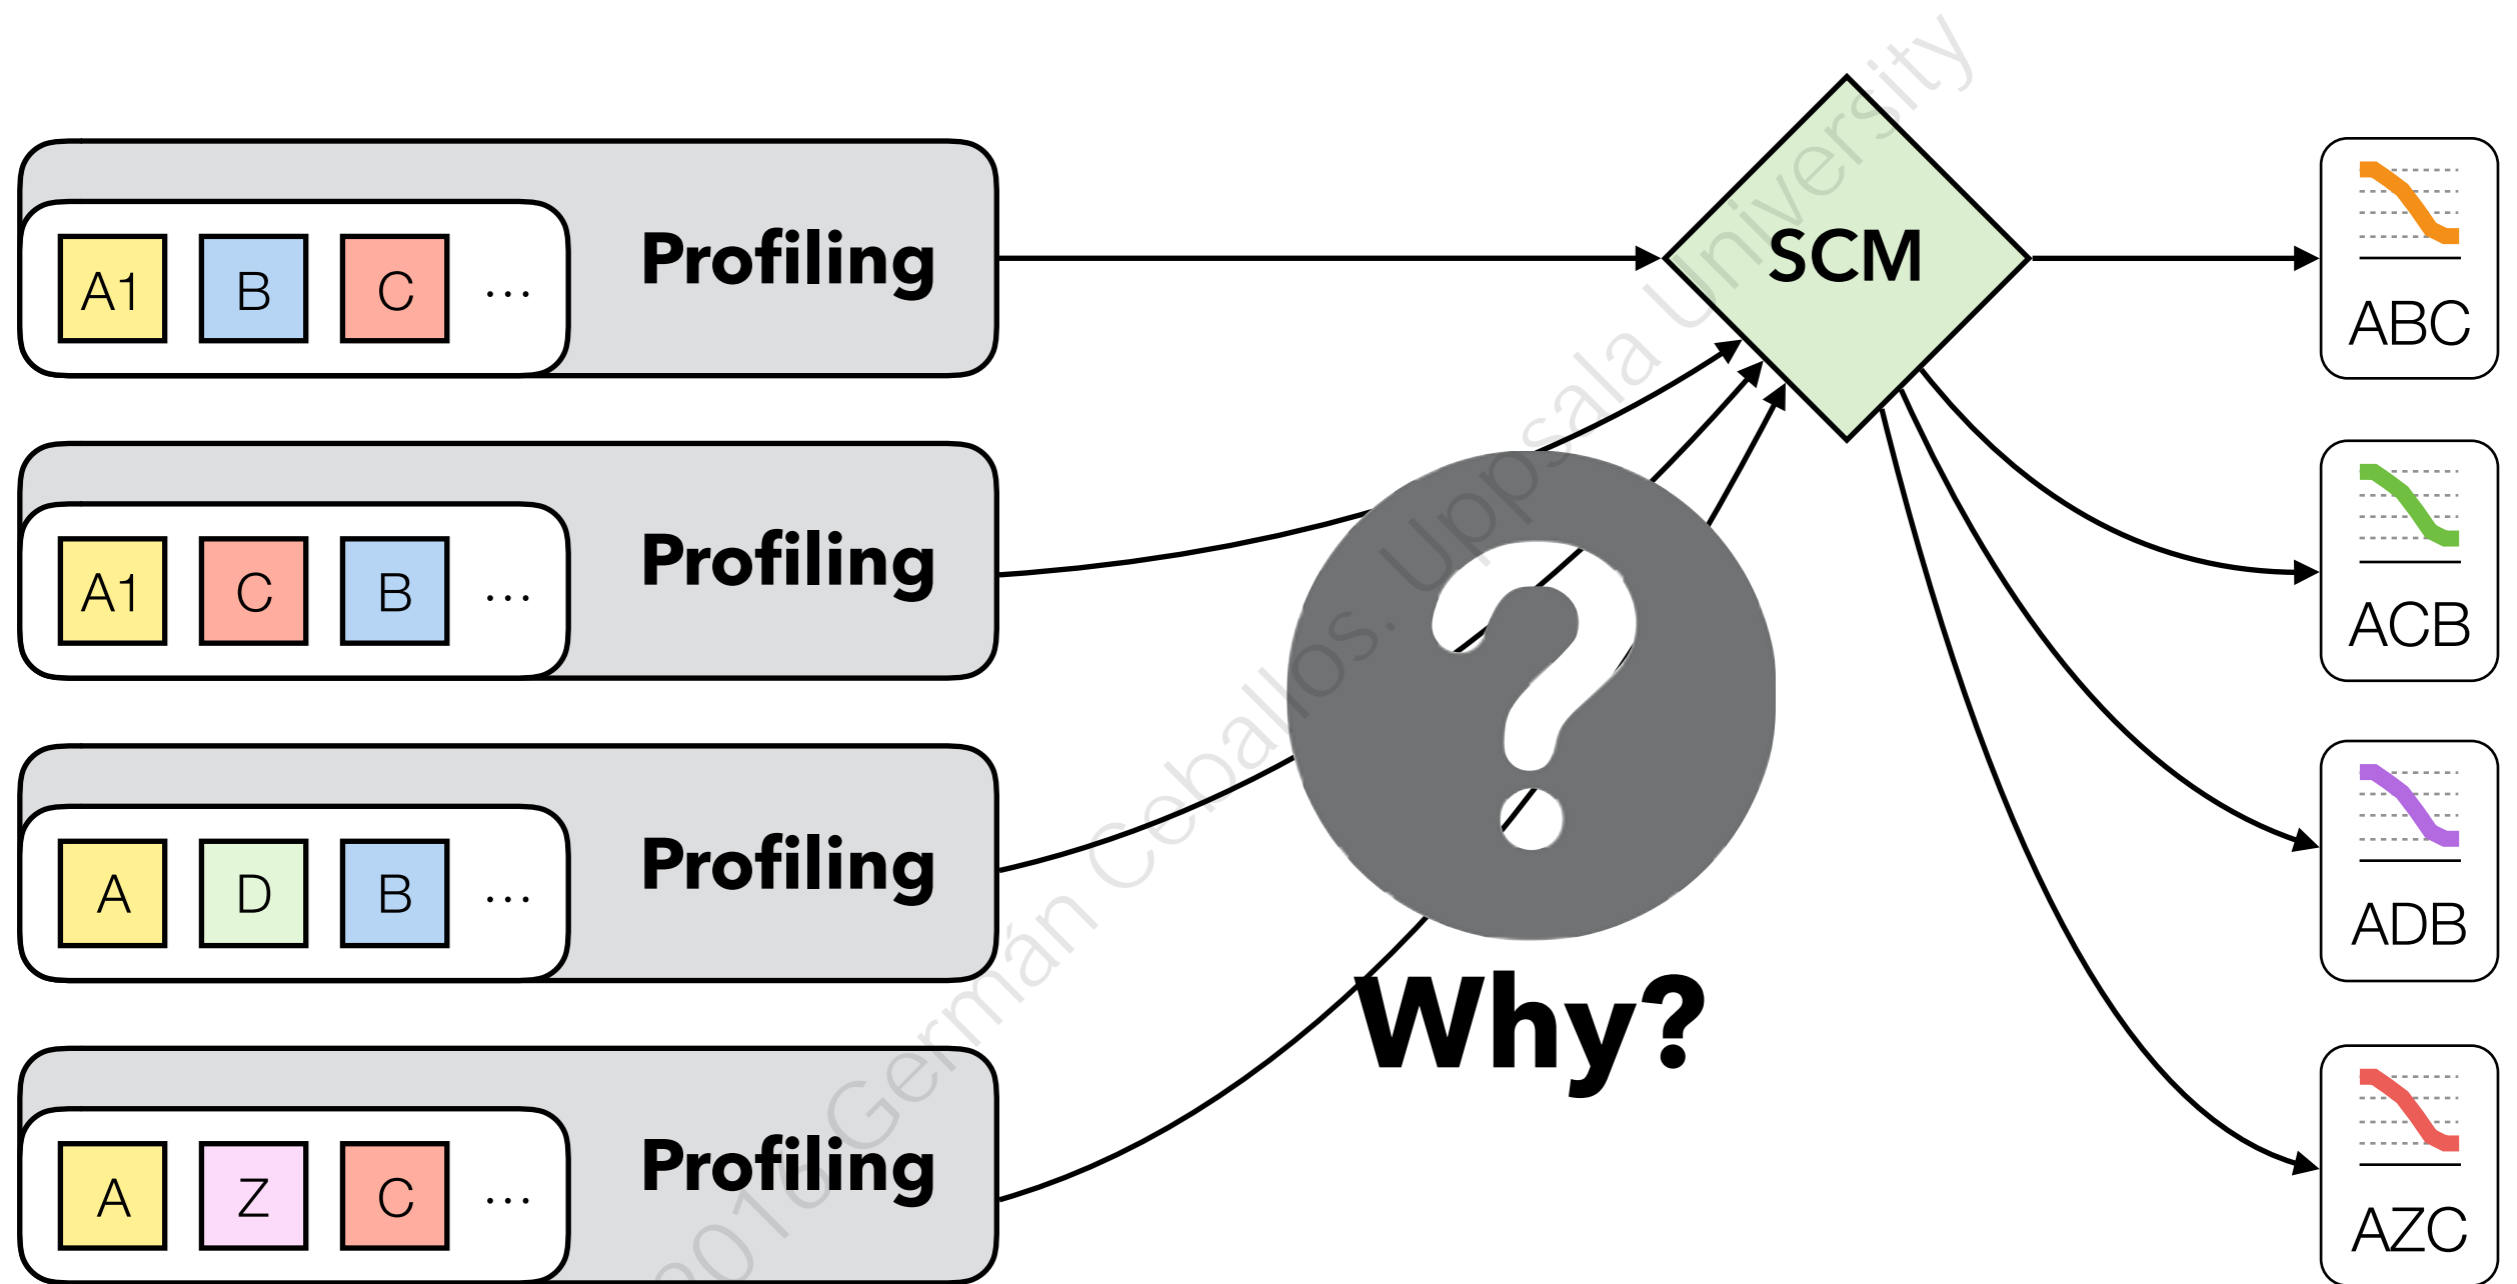
In *Proceedings of the 2004 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS '04)*. USA



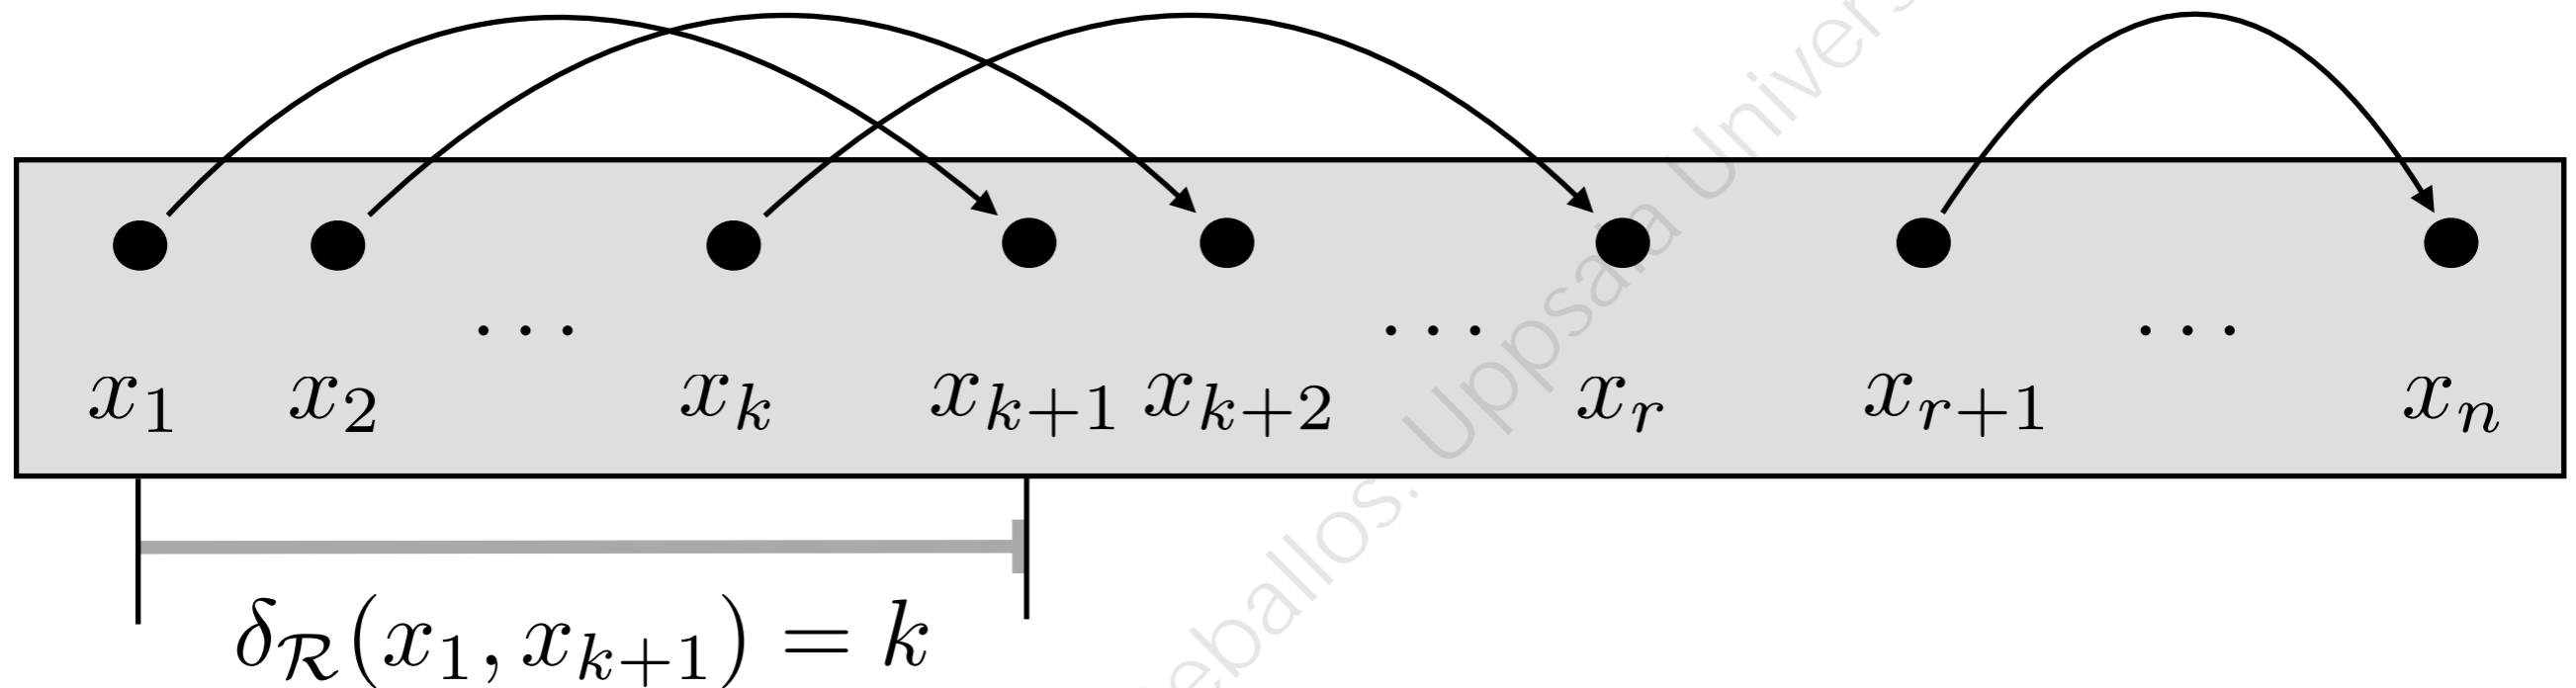
The Problem



The Problem



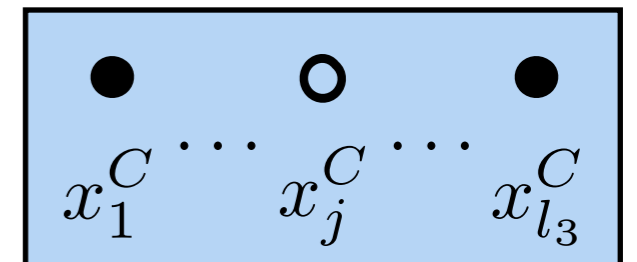
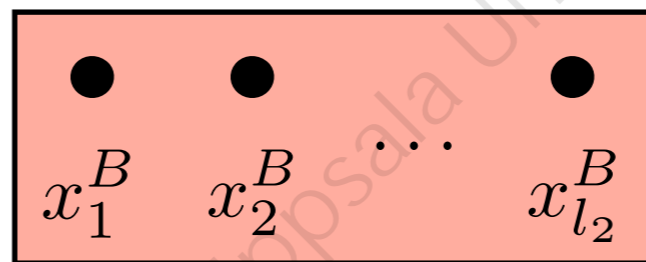
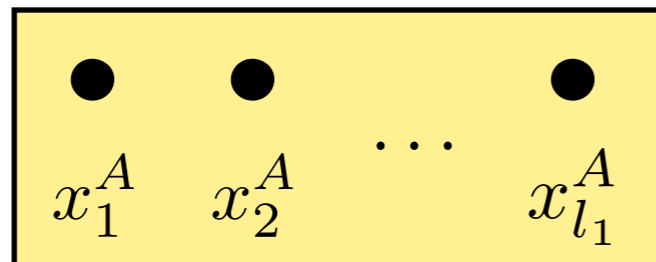
Statistical Cache Models



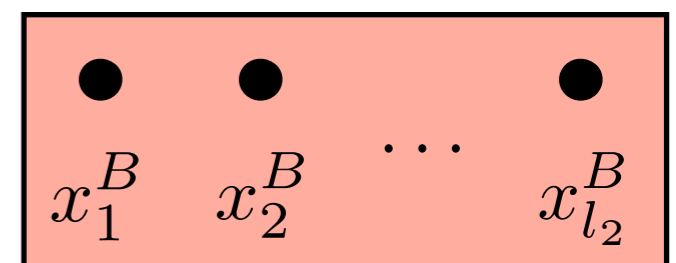
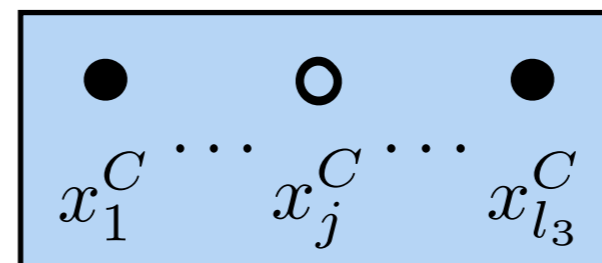
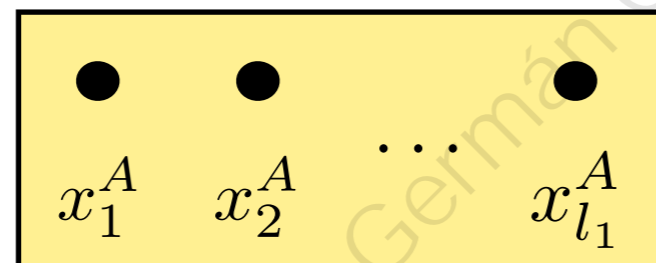
1. Sample memory accesses
2. Identifying consecutive reuses between them
3. Counting intervening accesses between reuses
4. Use statistical inference to estimate cache behaviour

SCM with Tasks

S

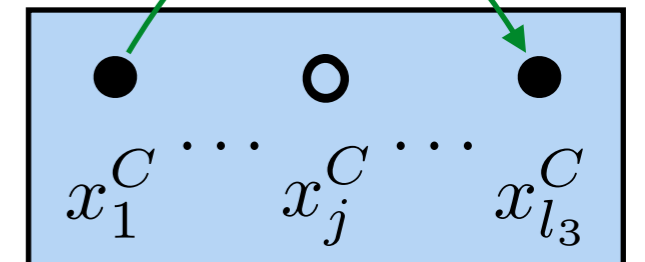
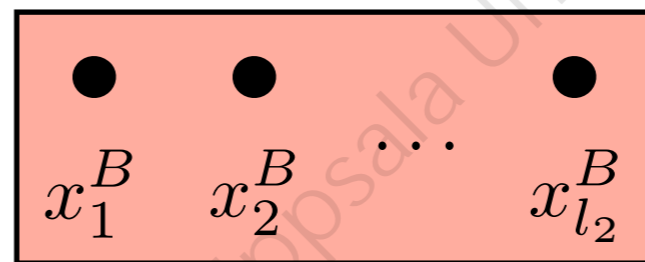
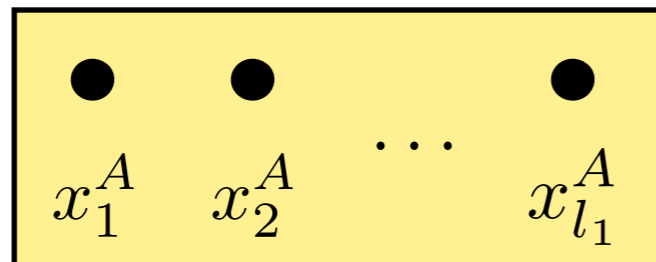


S'

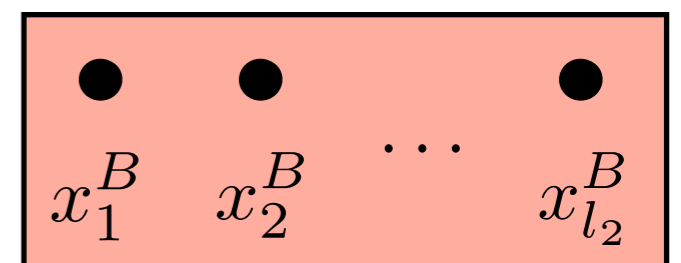
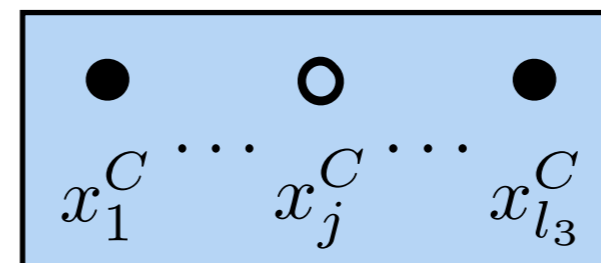
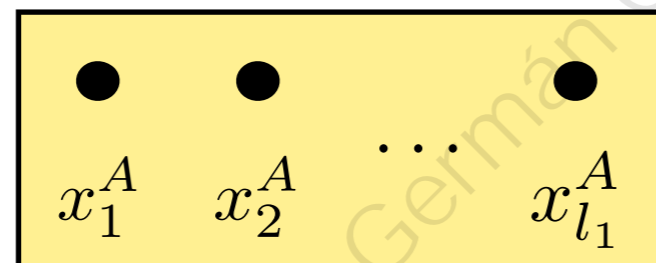


SCM with Tasks

S

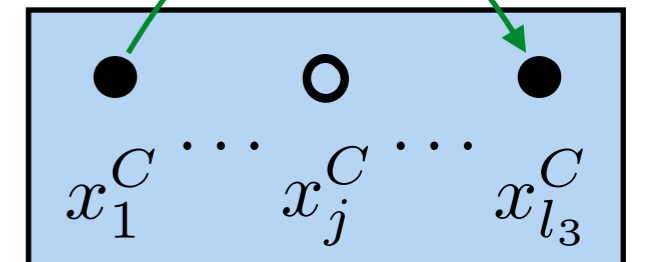
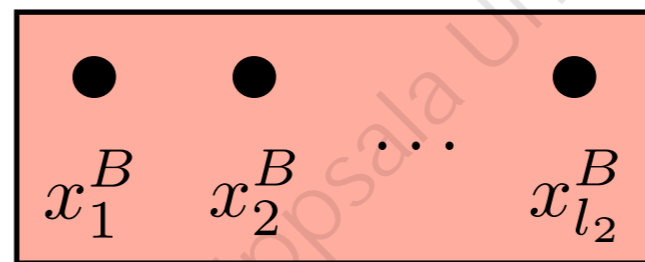
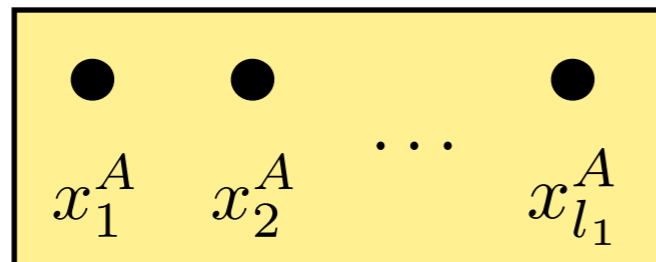


S'

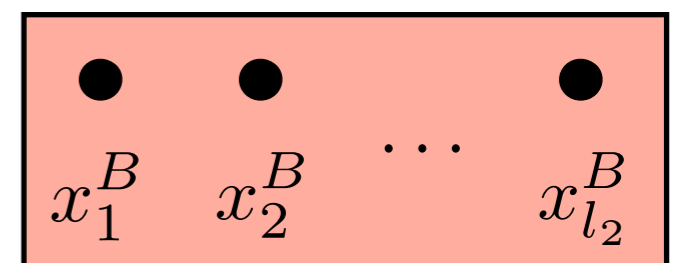
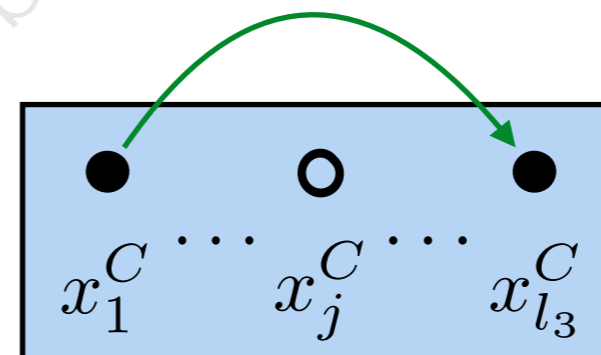
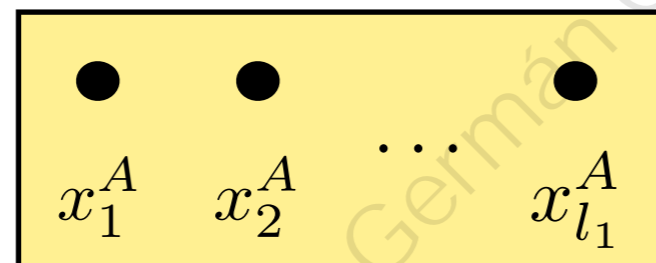


SCM with Tasks

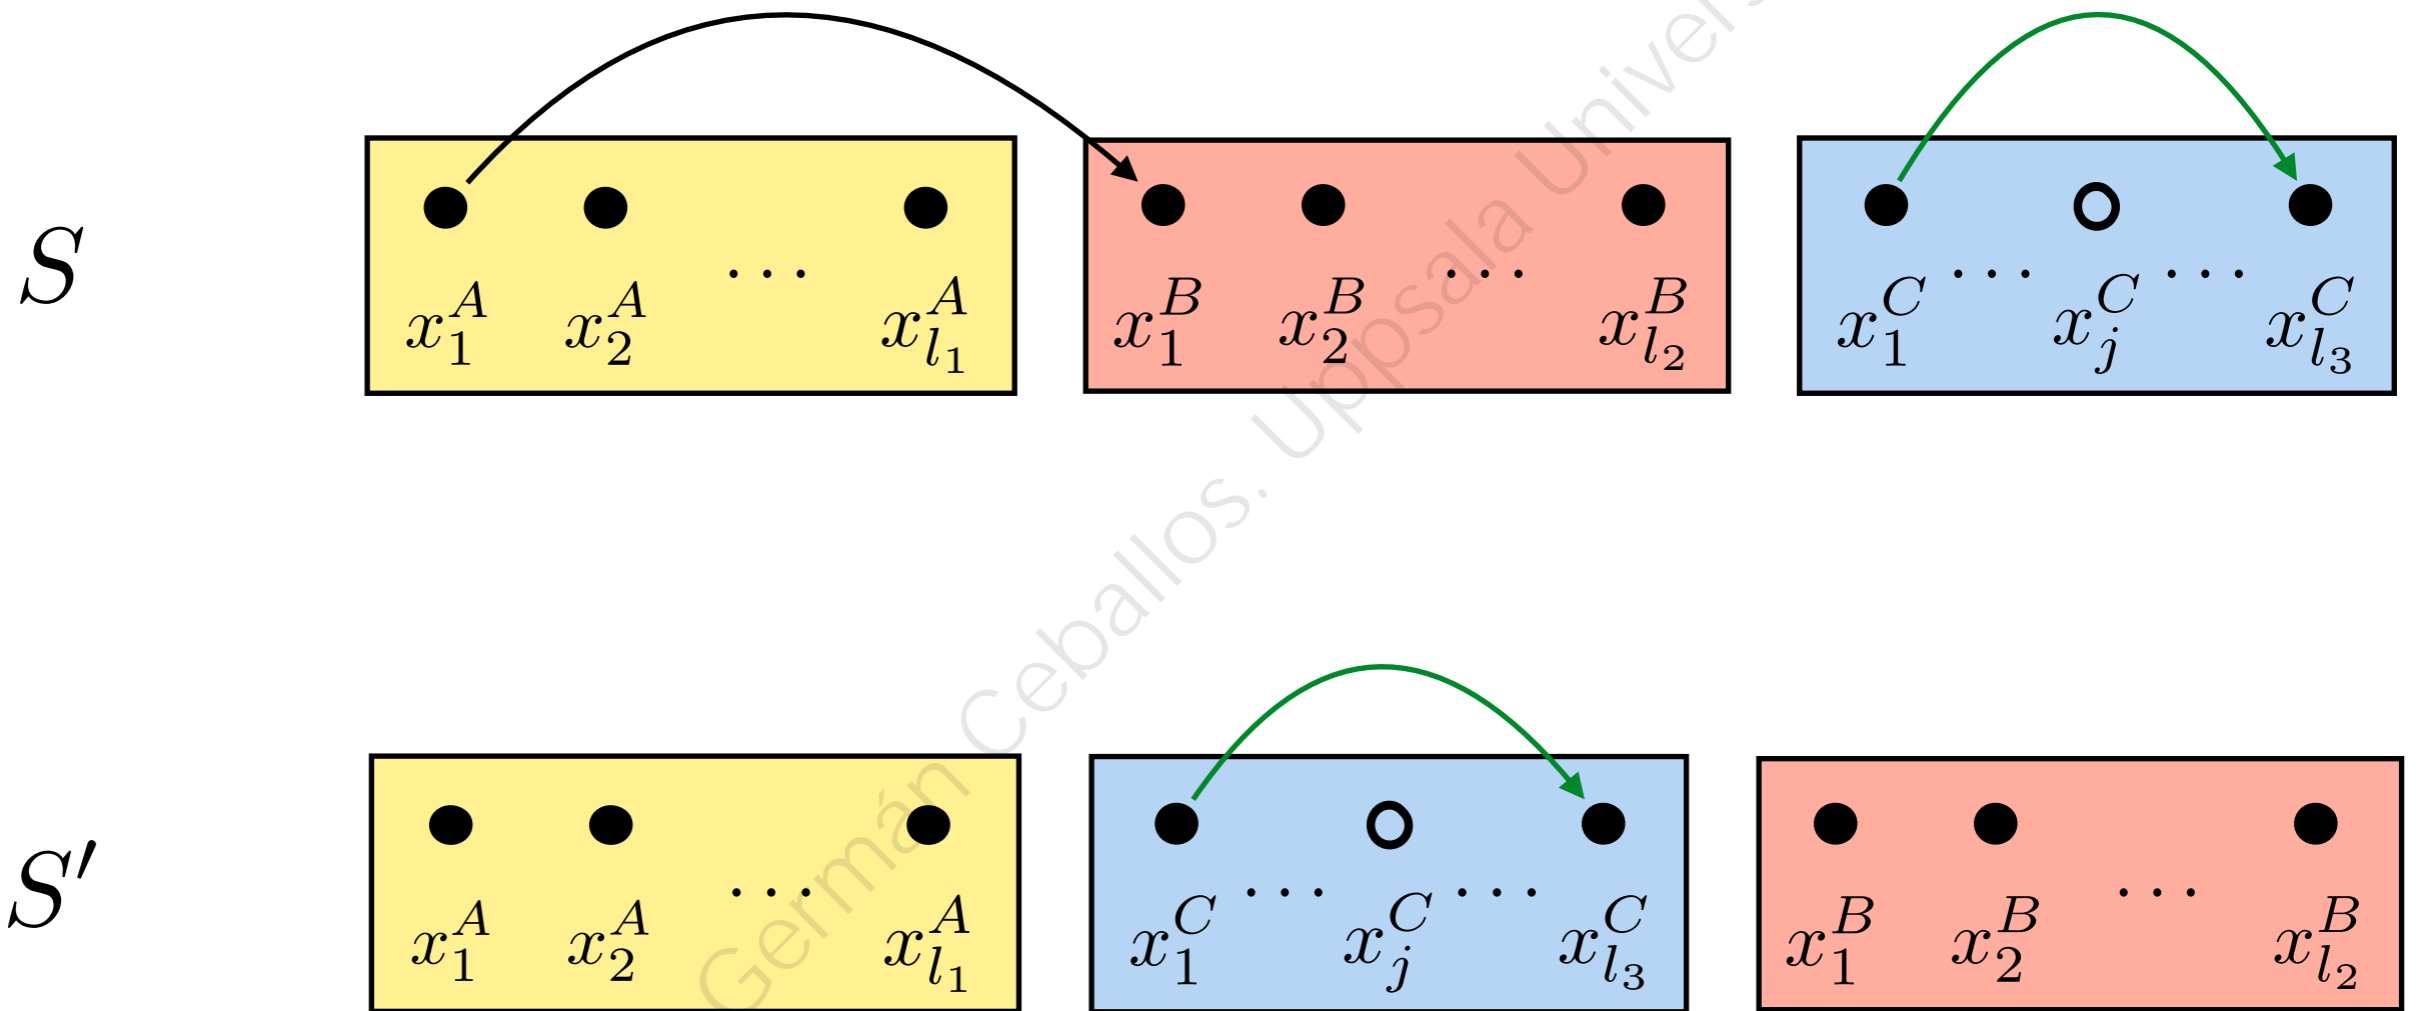
S



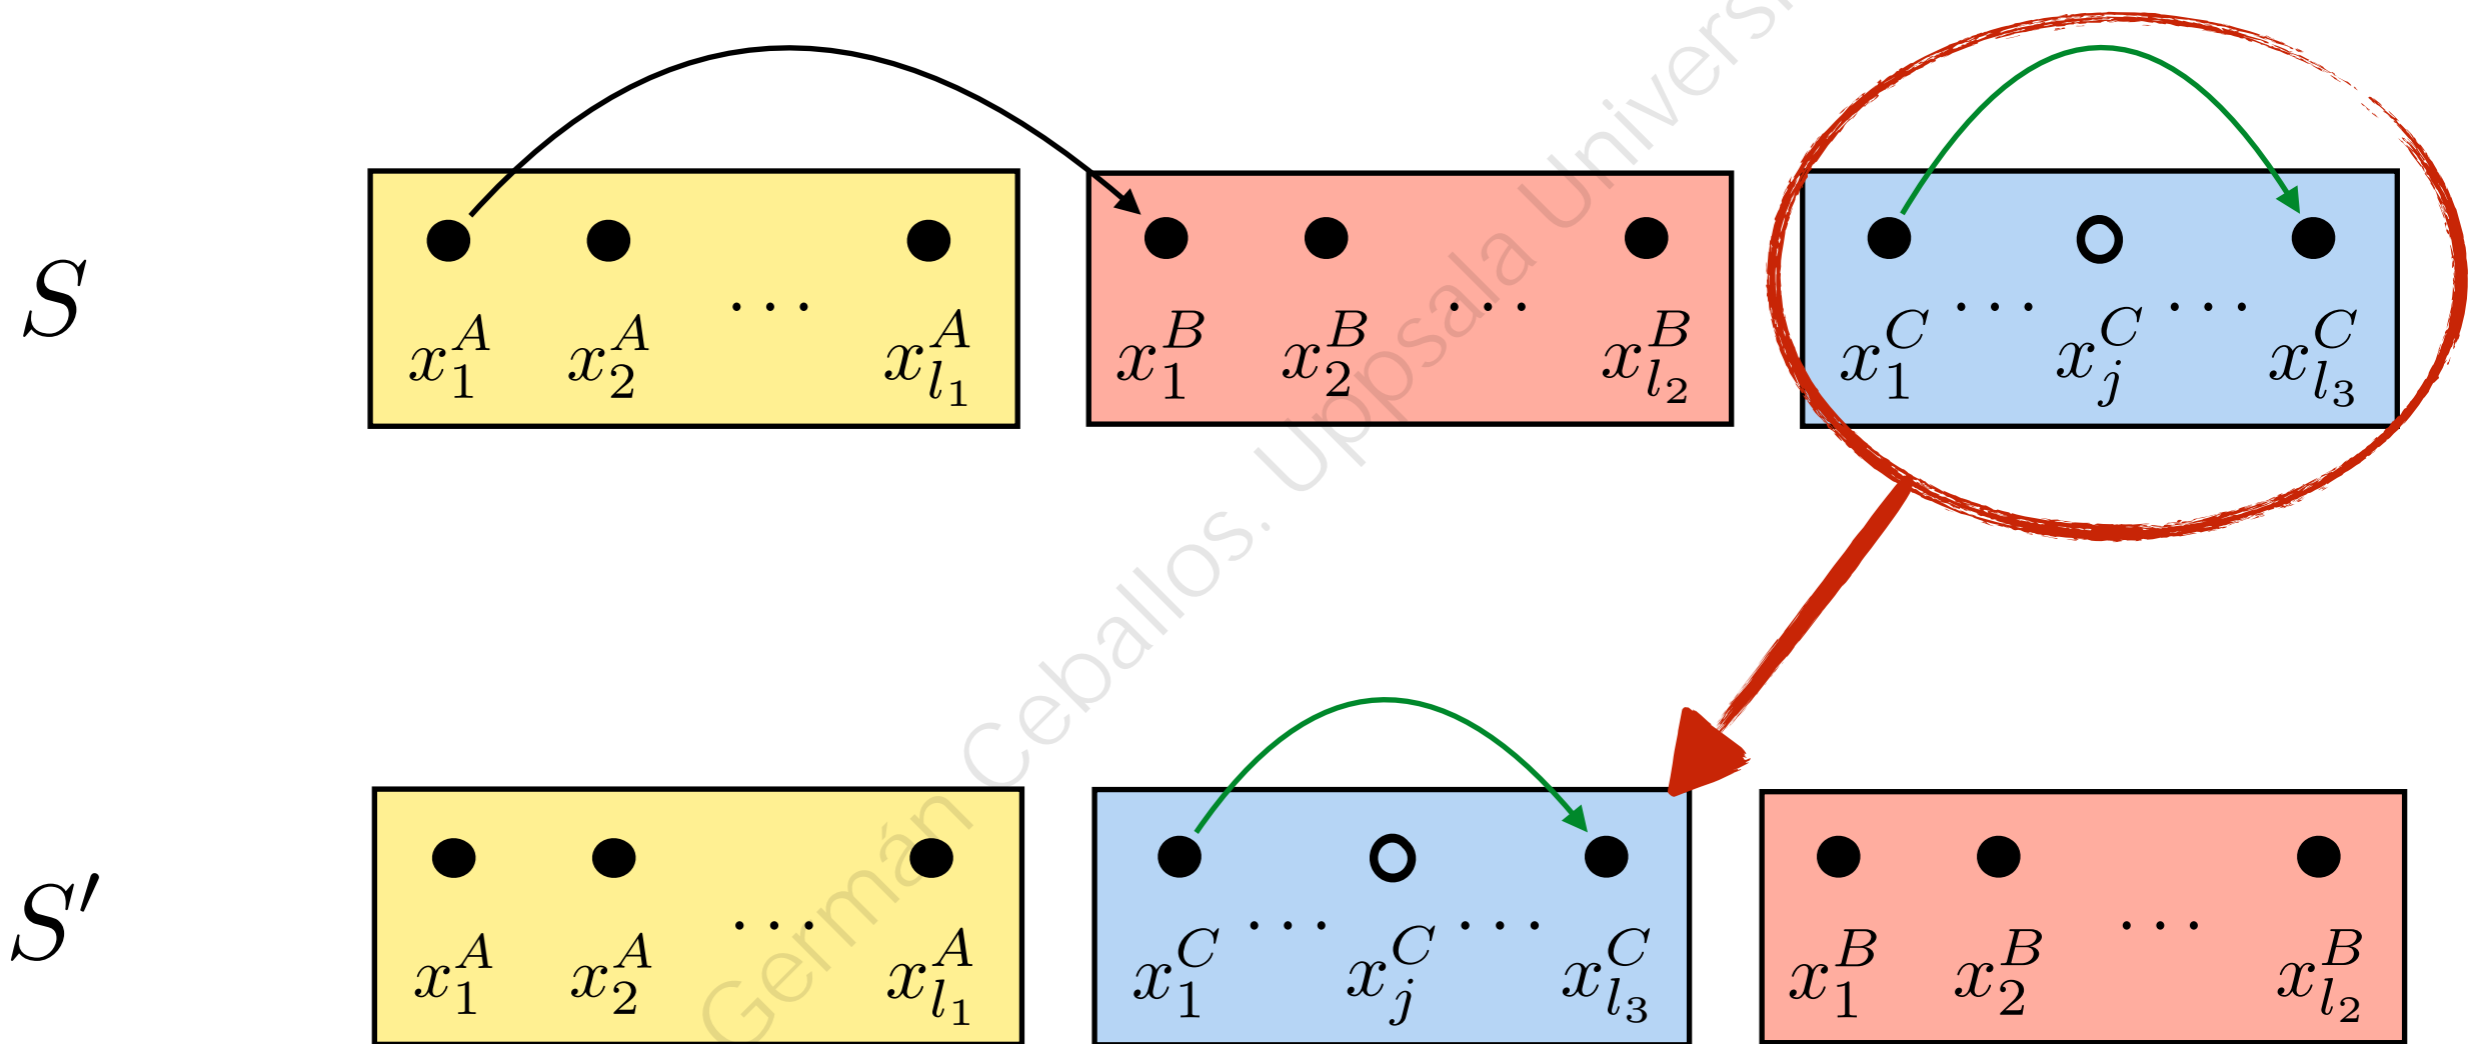
S'



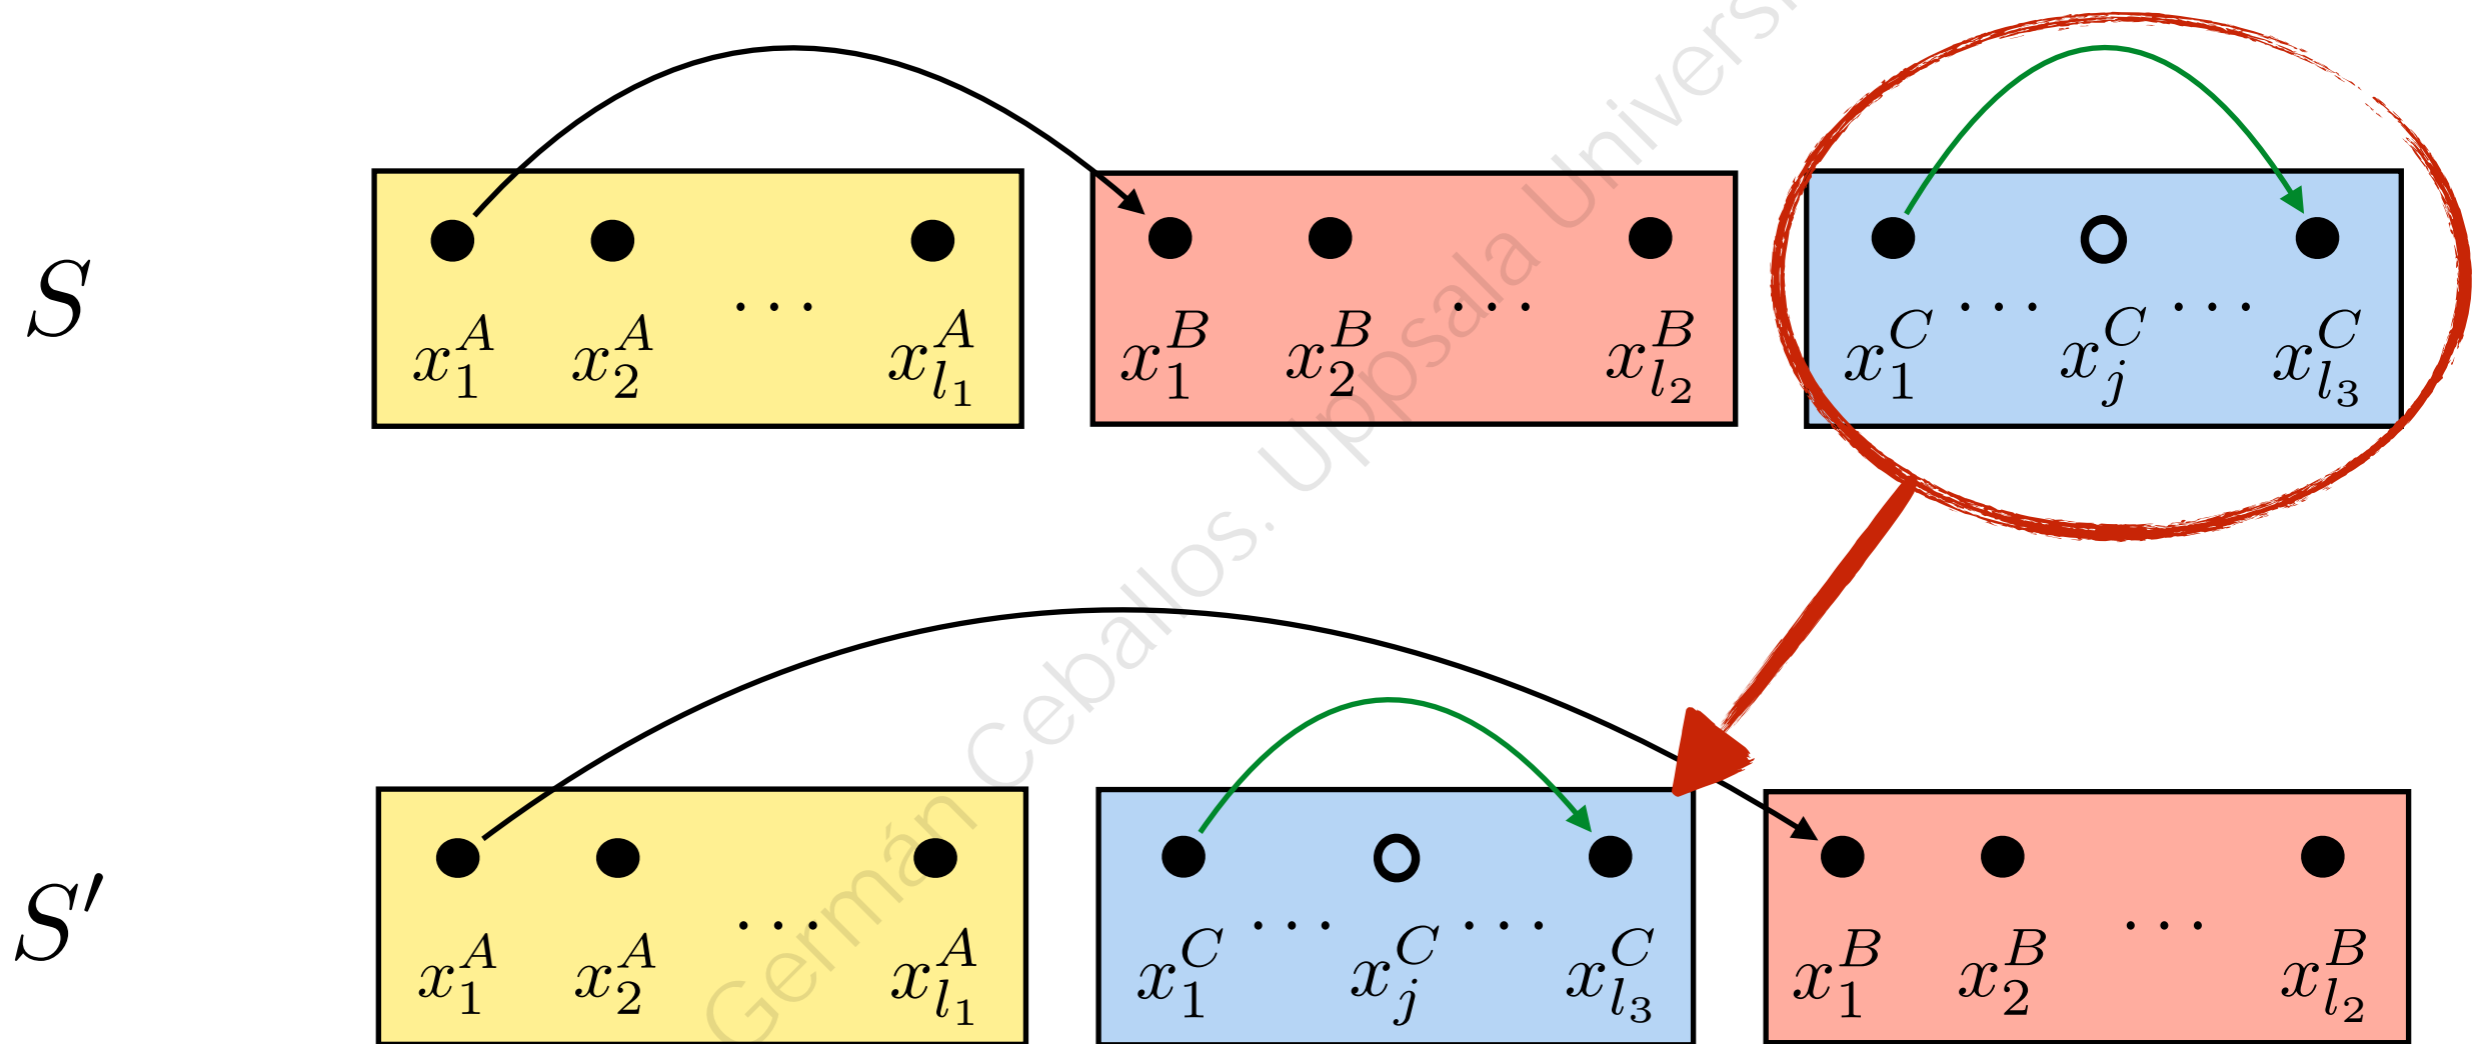
SCM with Tasks



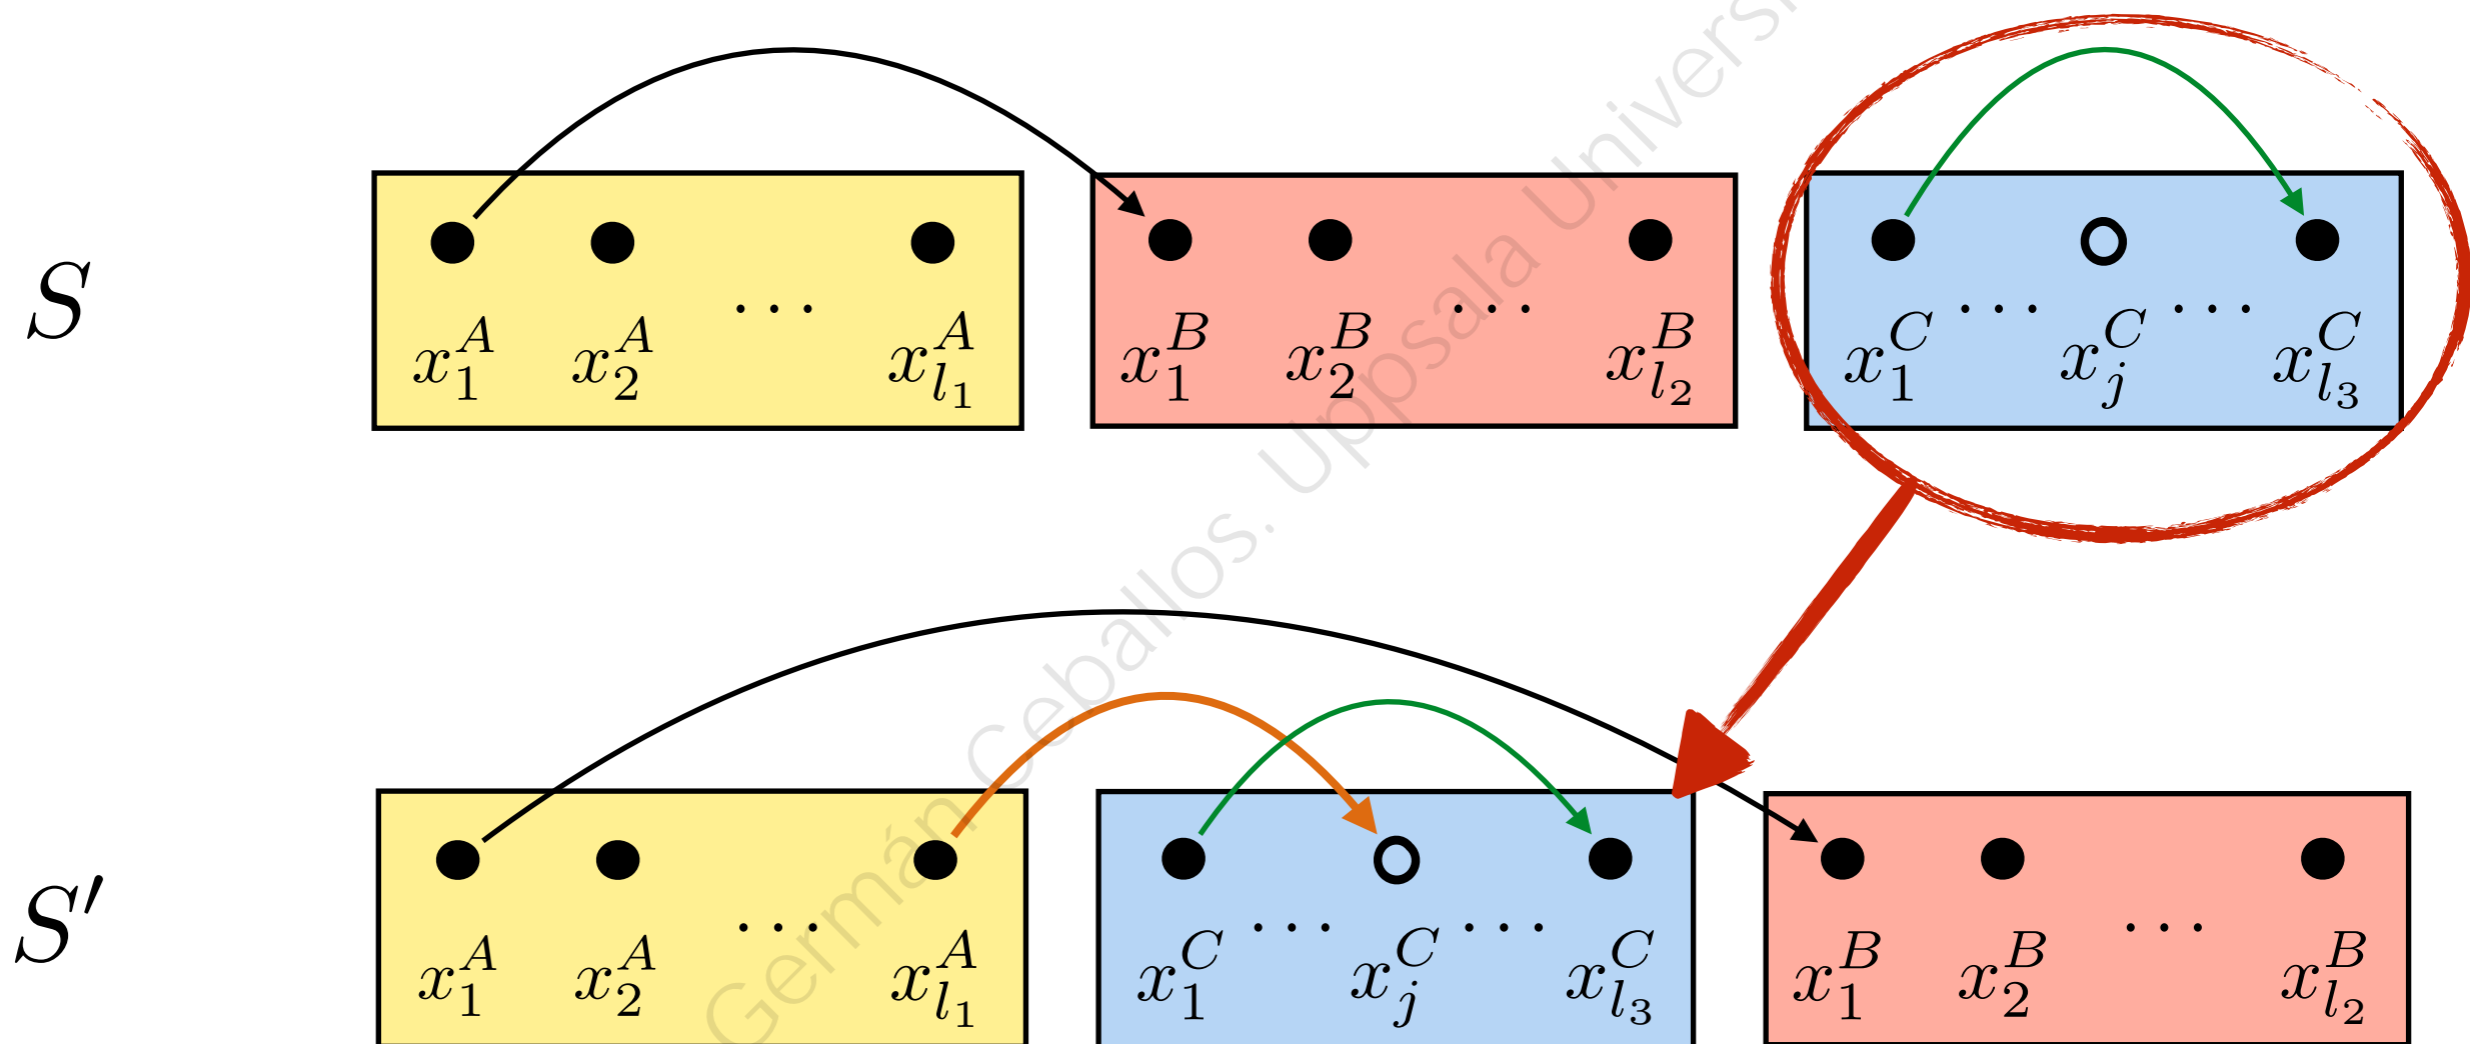
SCM with Tasks



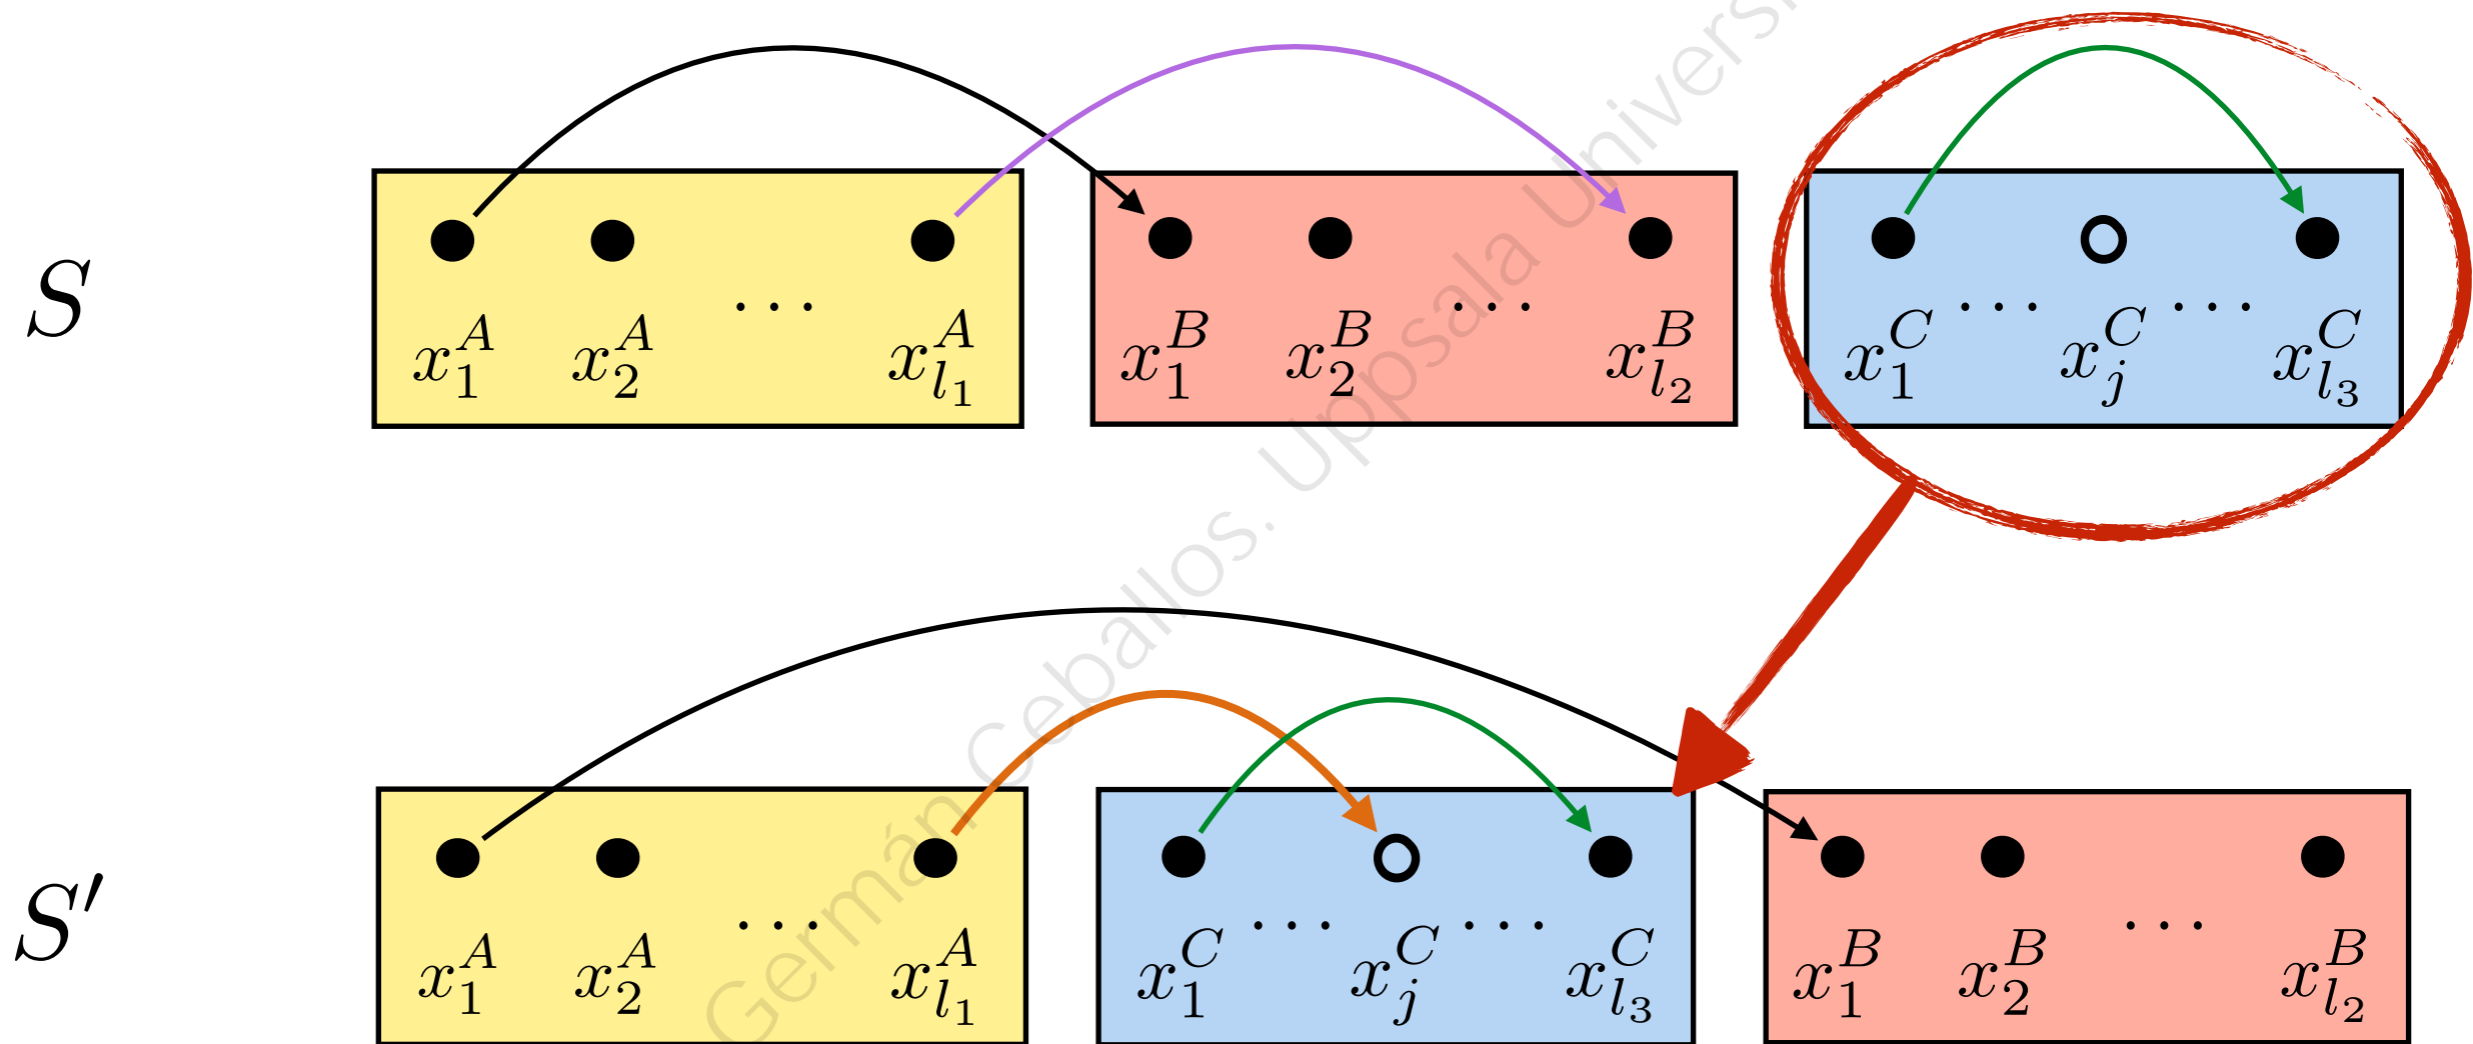
SCM with Tasks



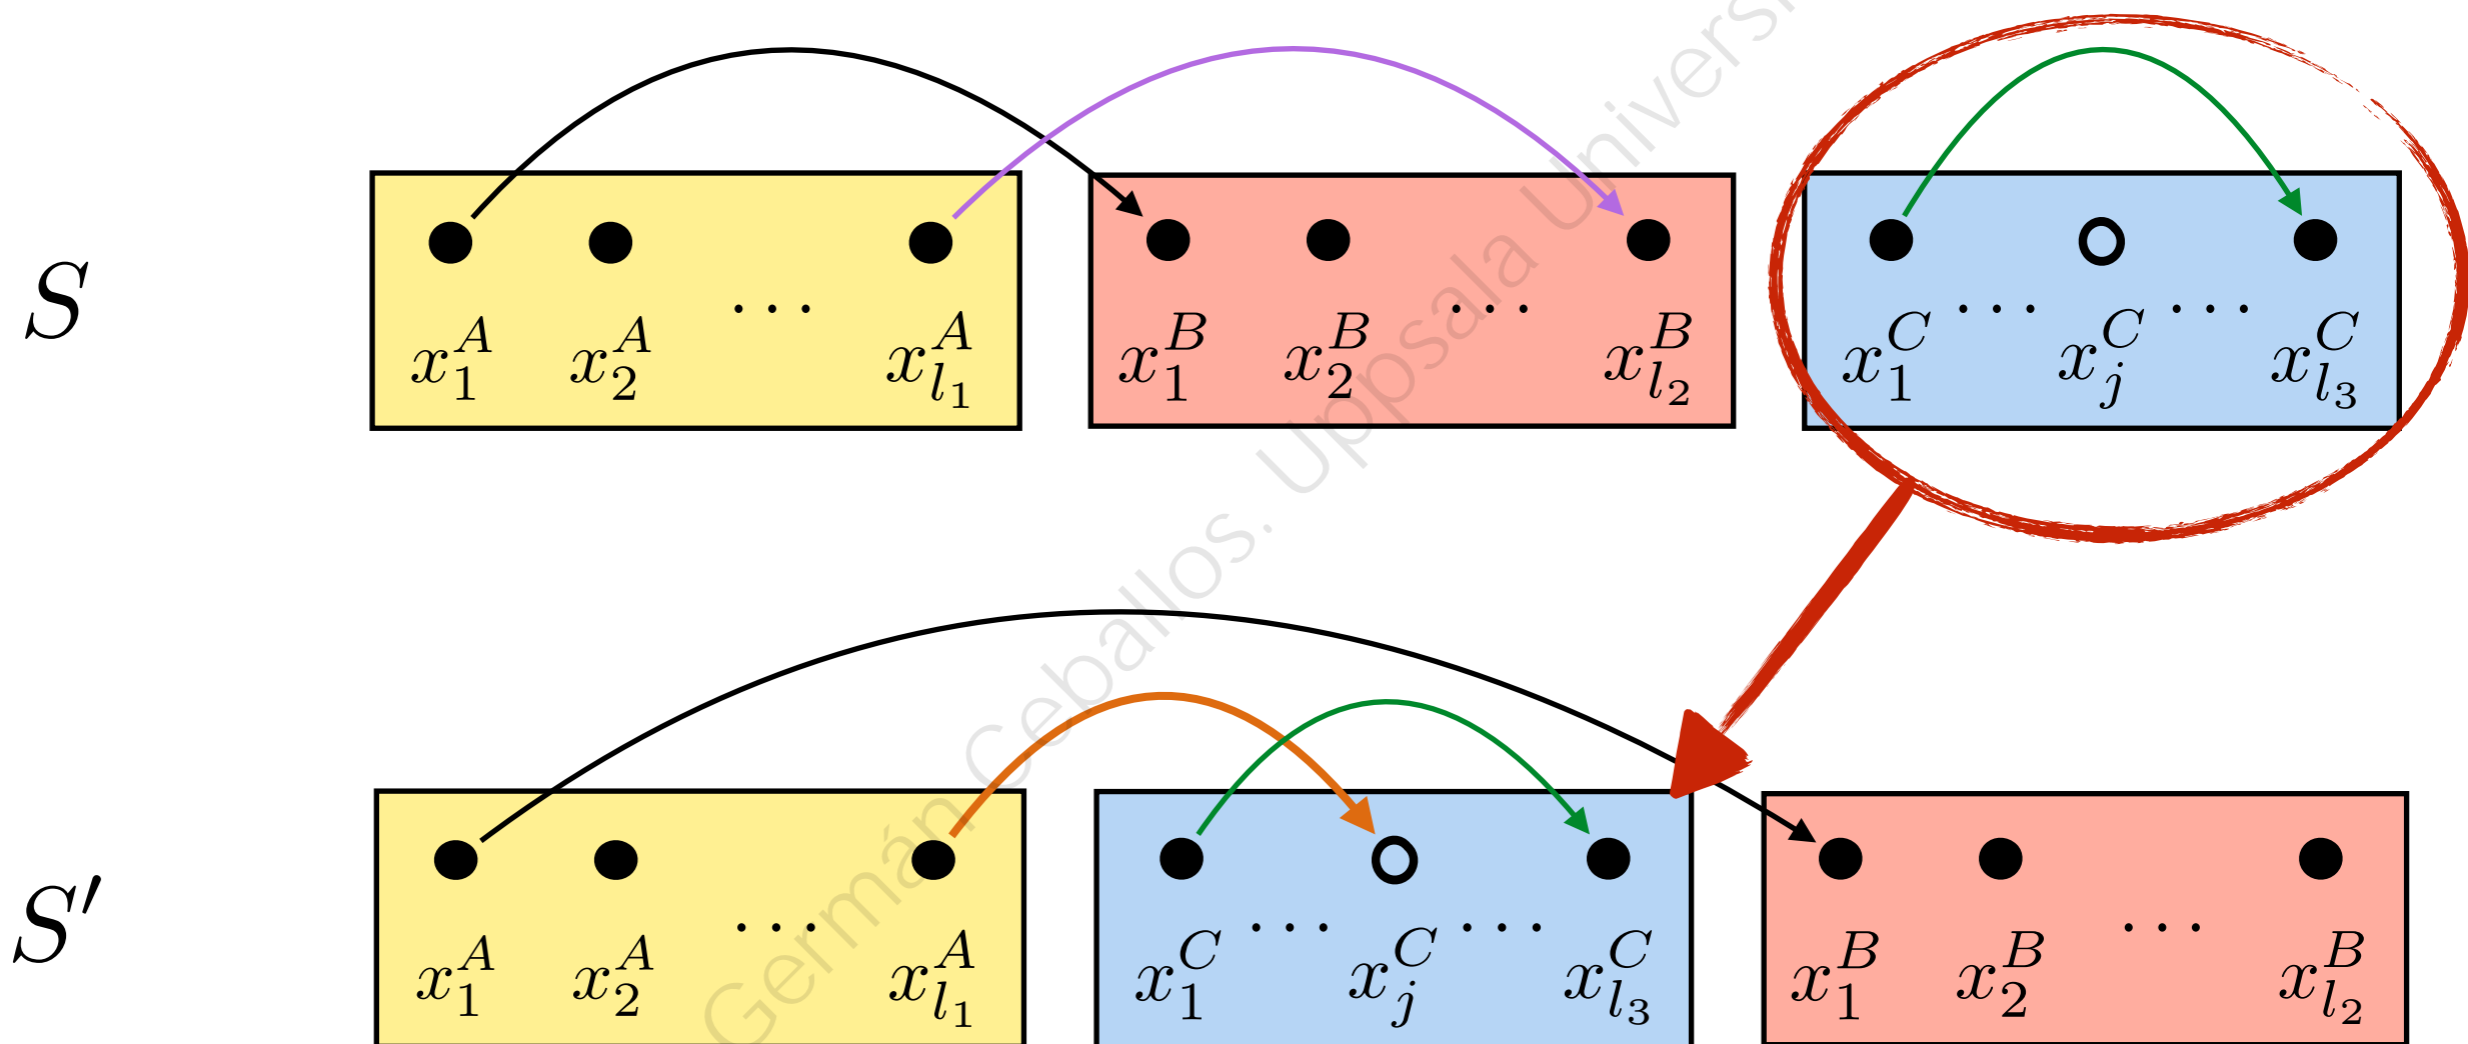
SCM with Tasks



SCM with Tasks



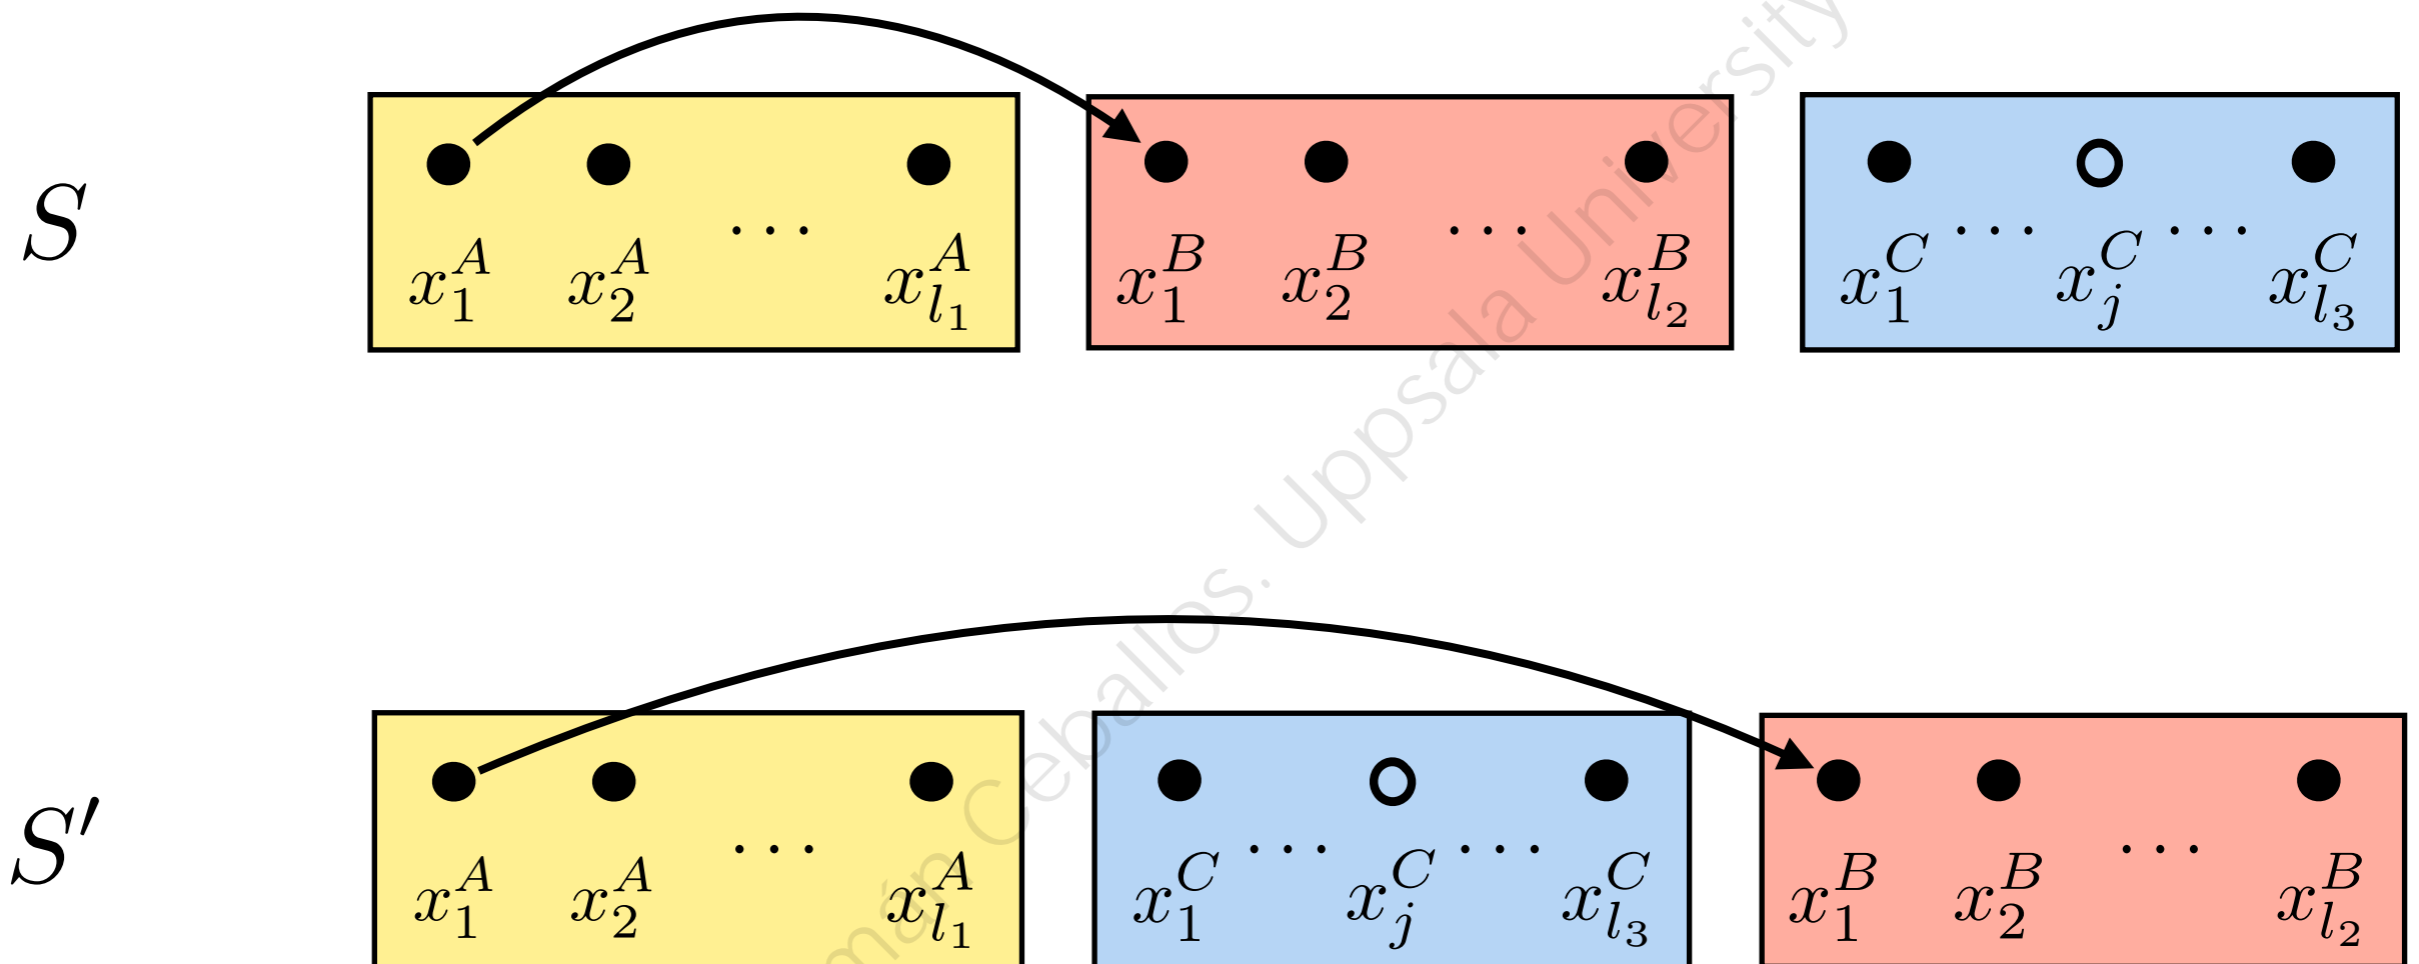
SCM with Tasks



Reuses can change!

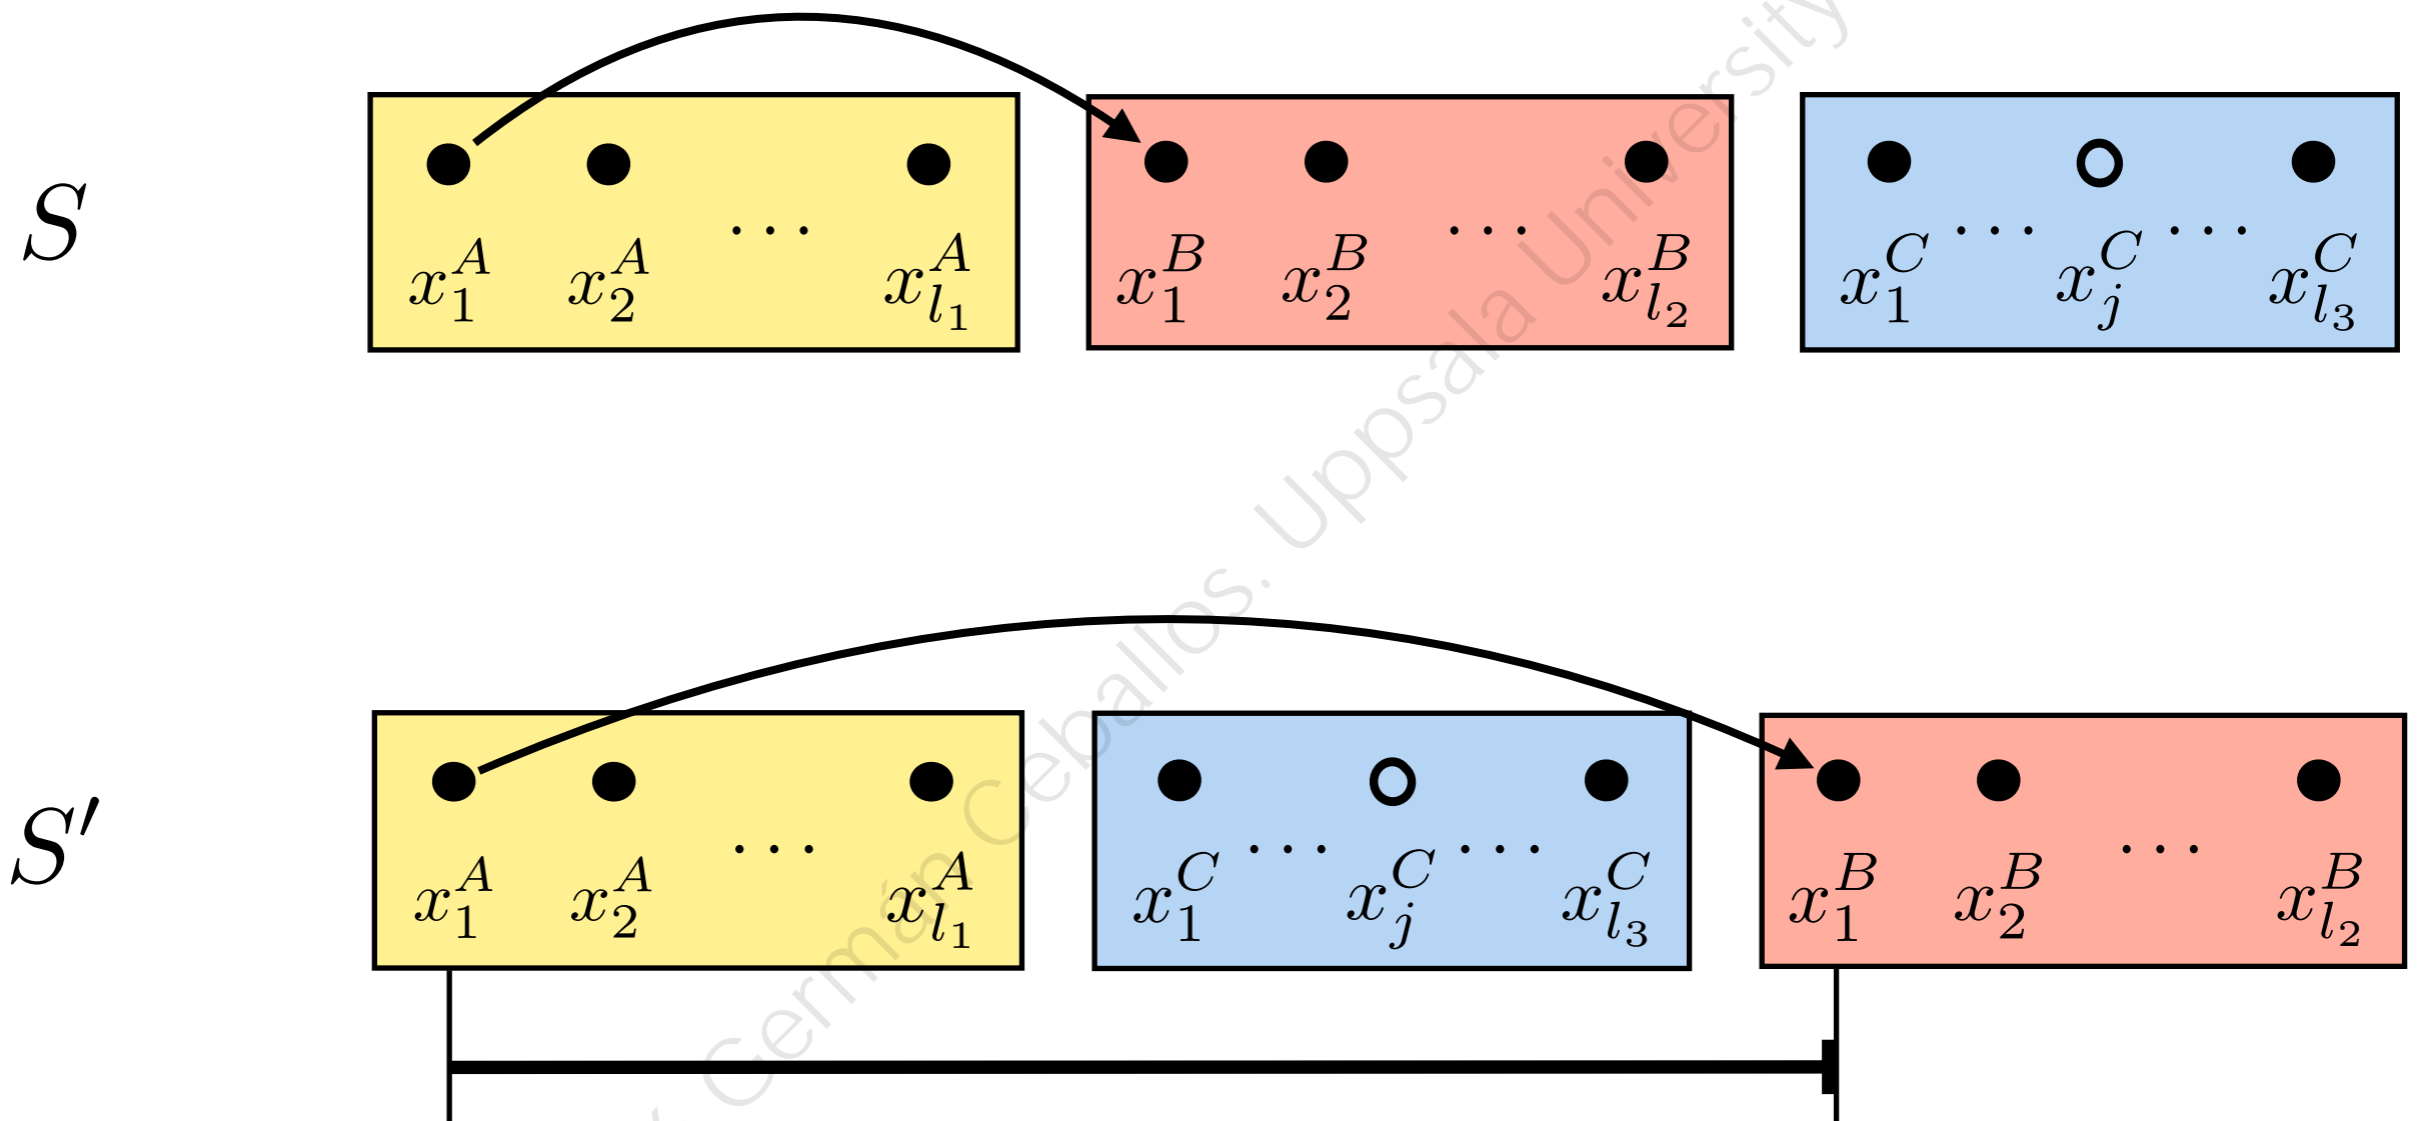


SCM with Tasks

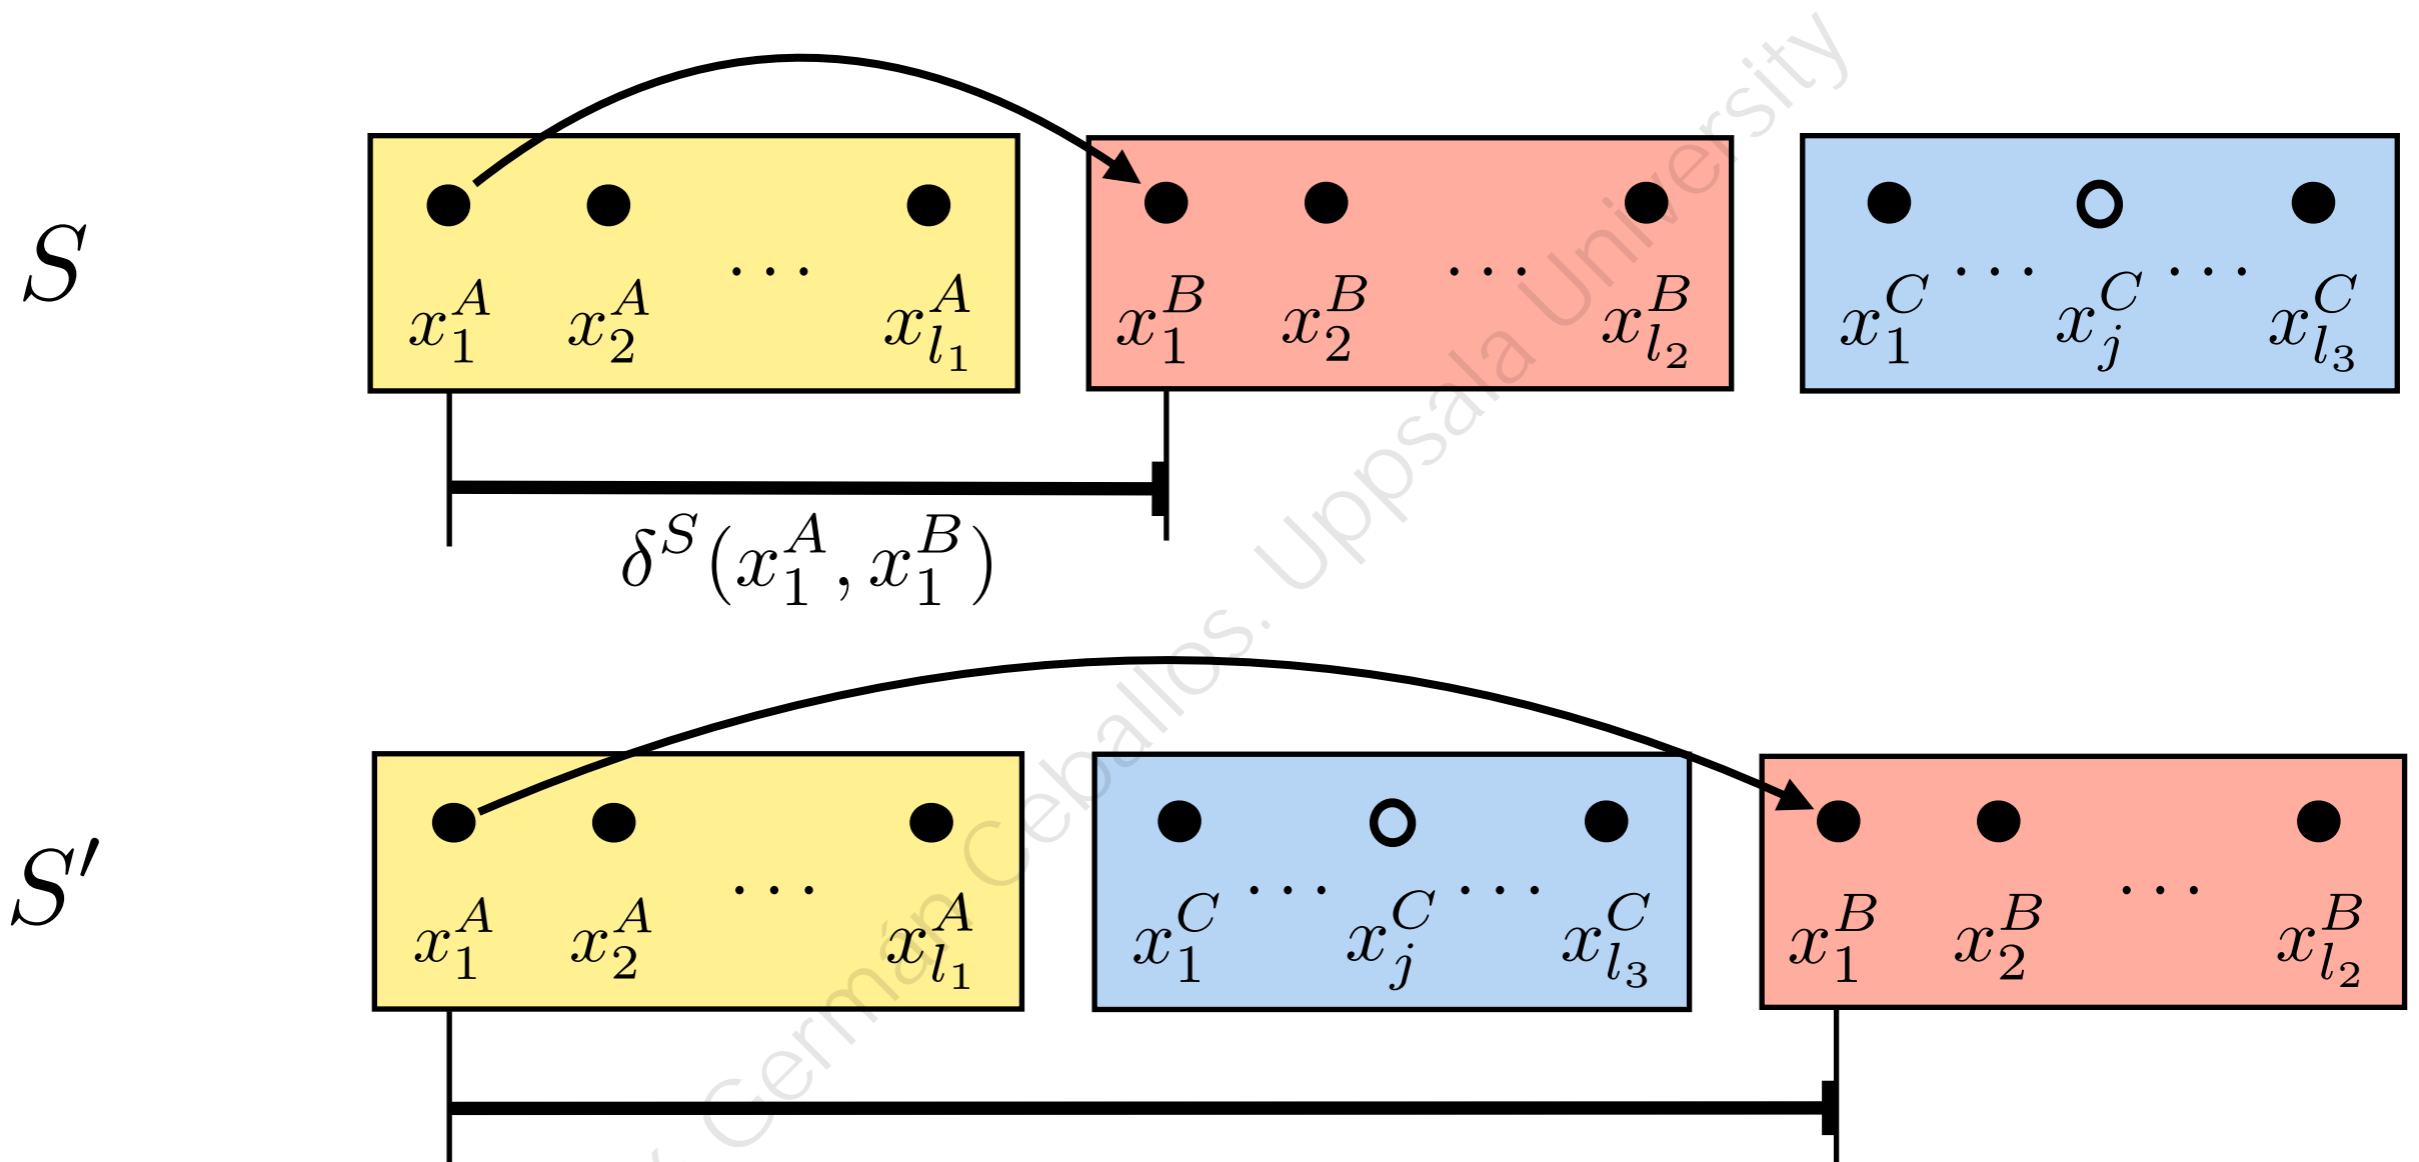


(c) 2016 German Ceballos, Uppsala University

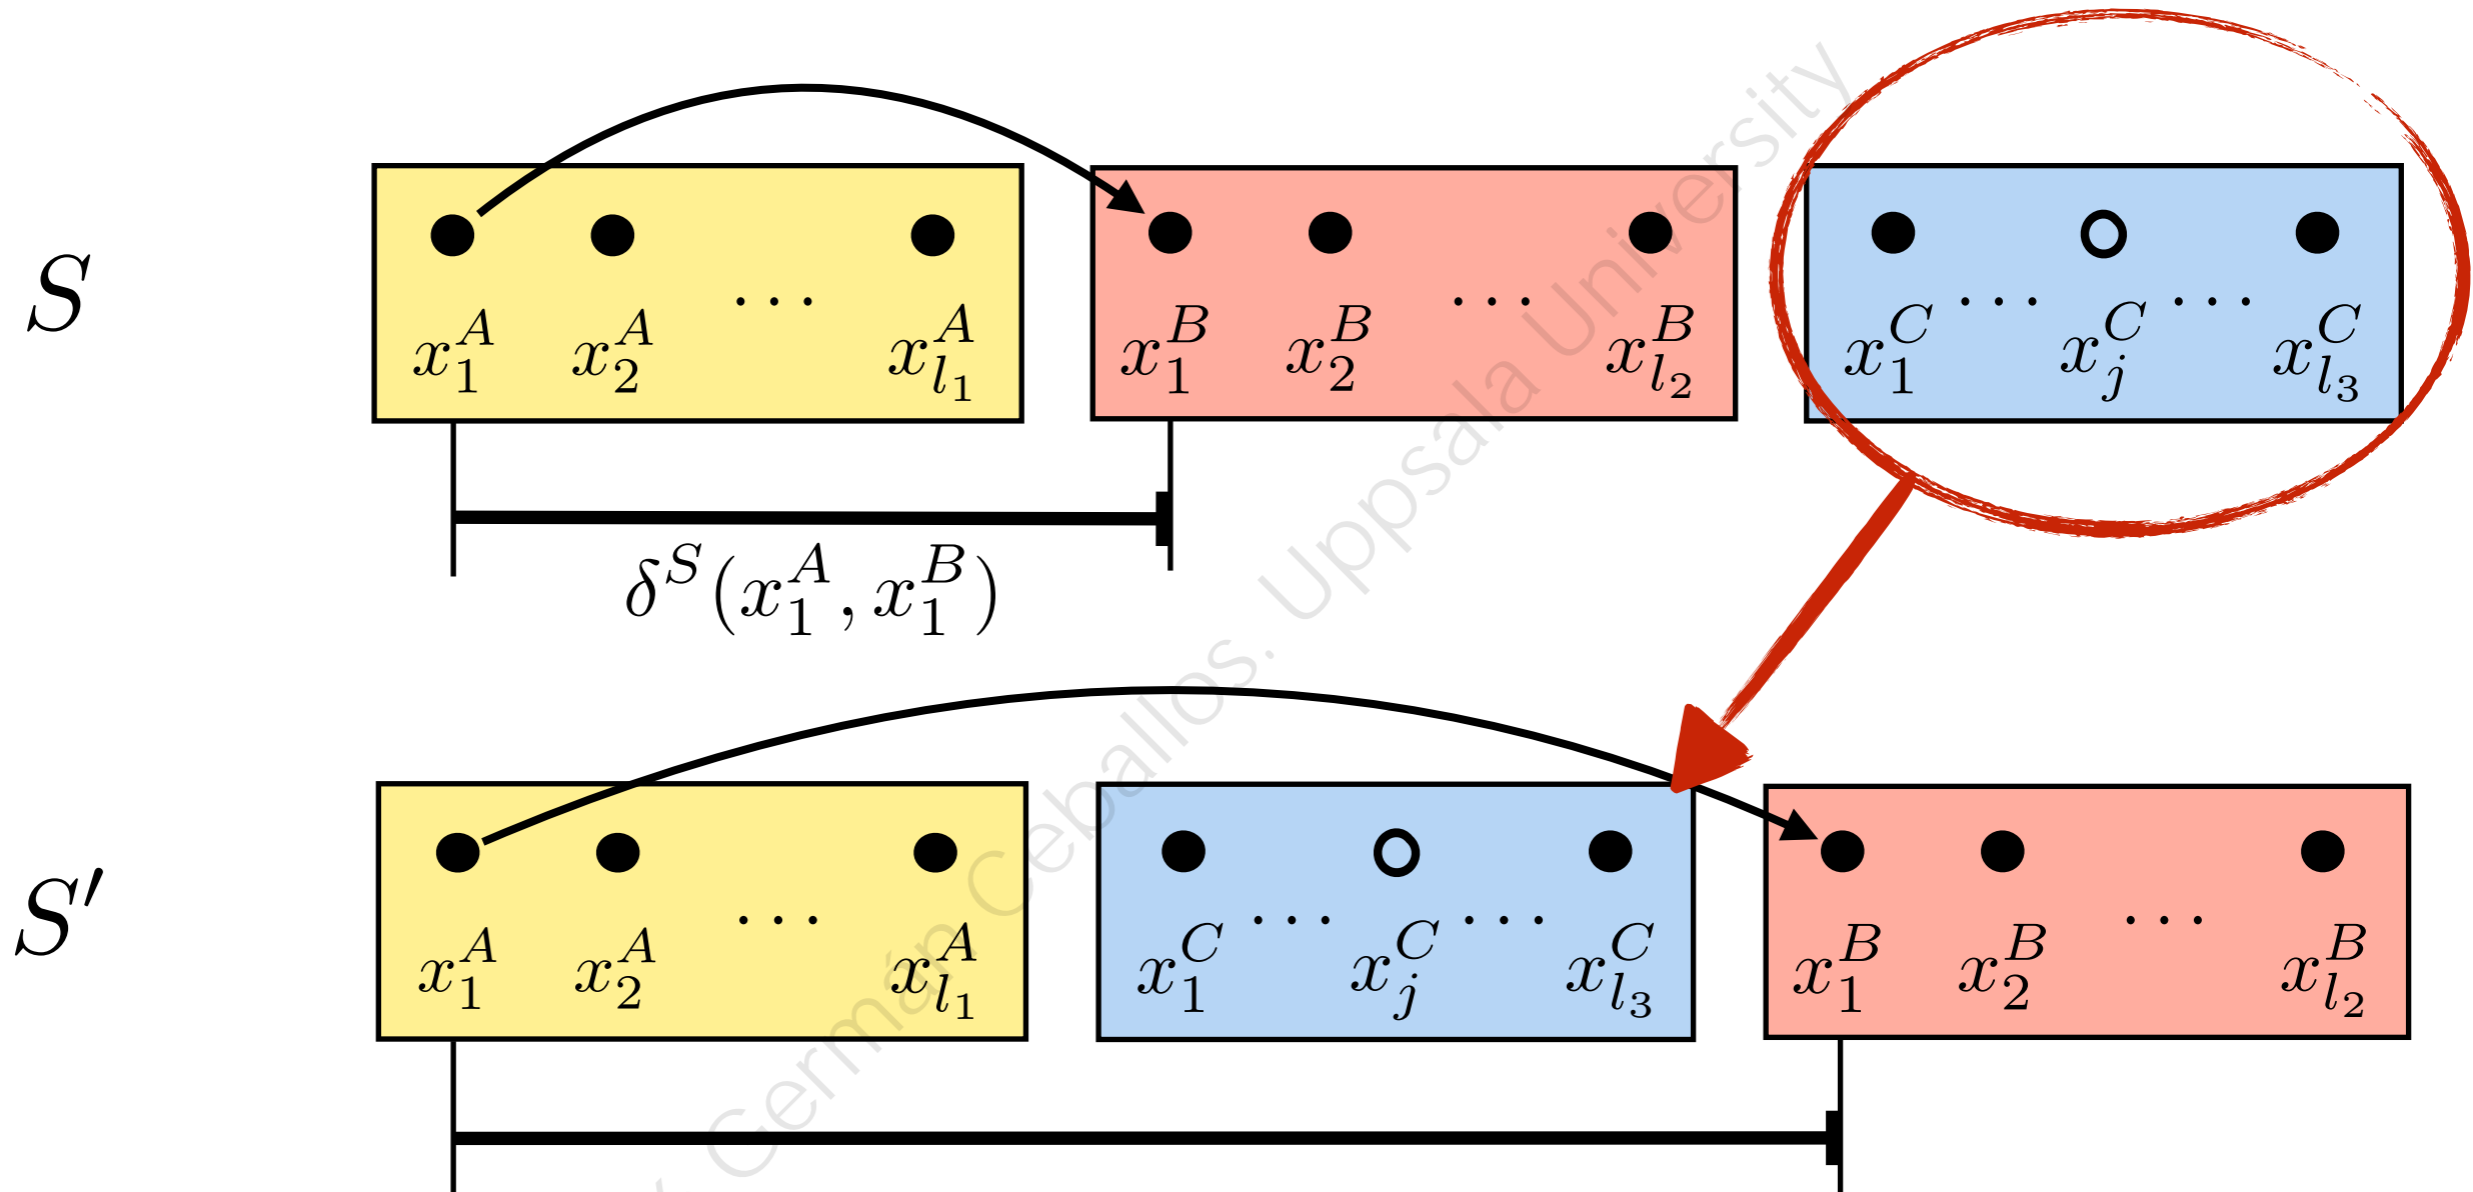
SCM with Tasks



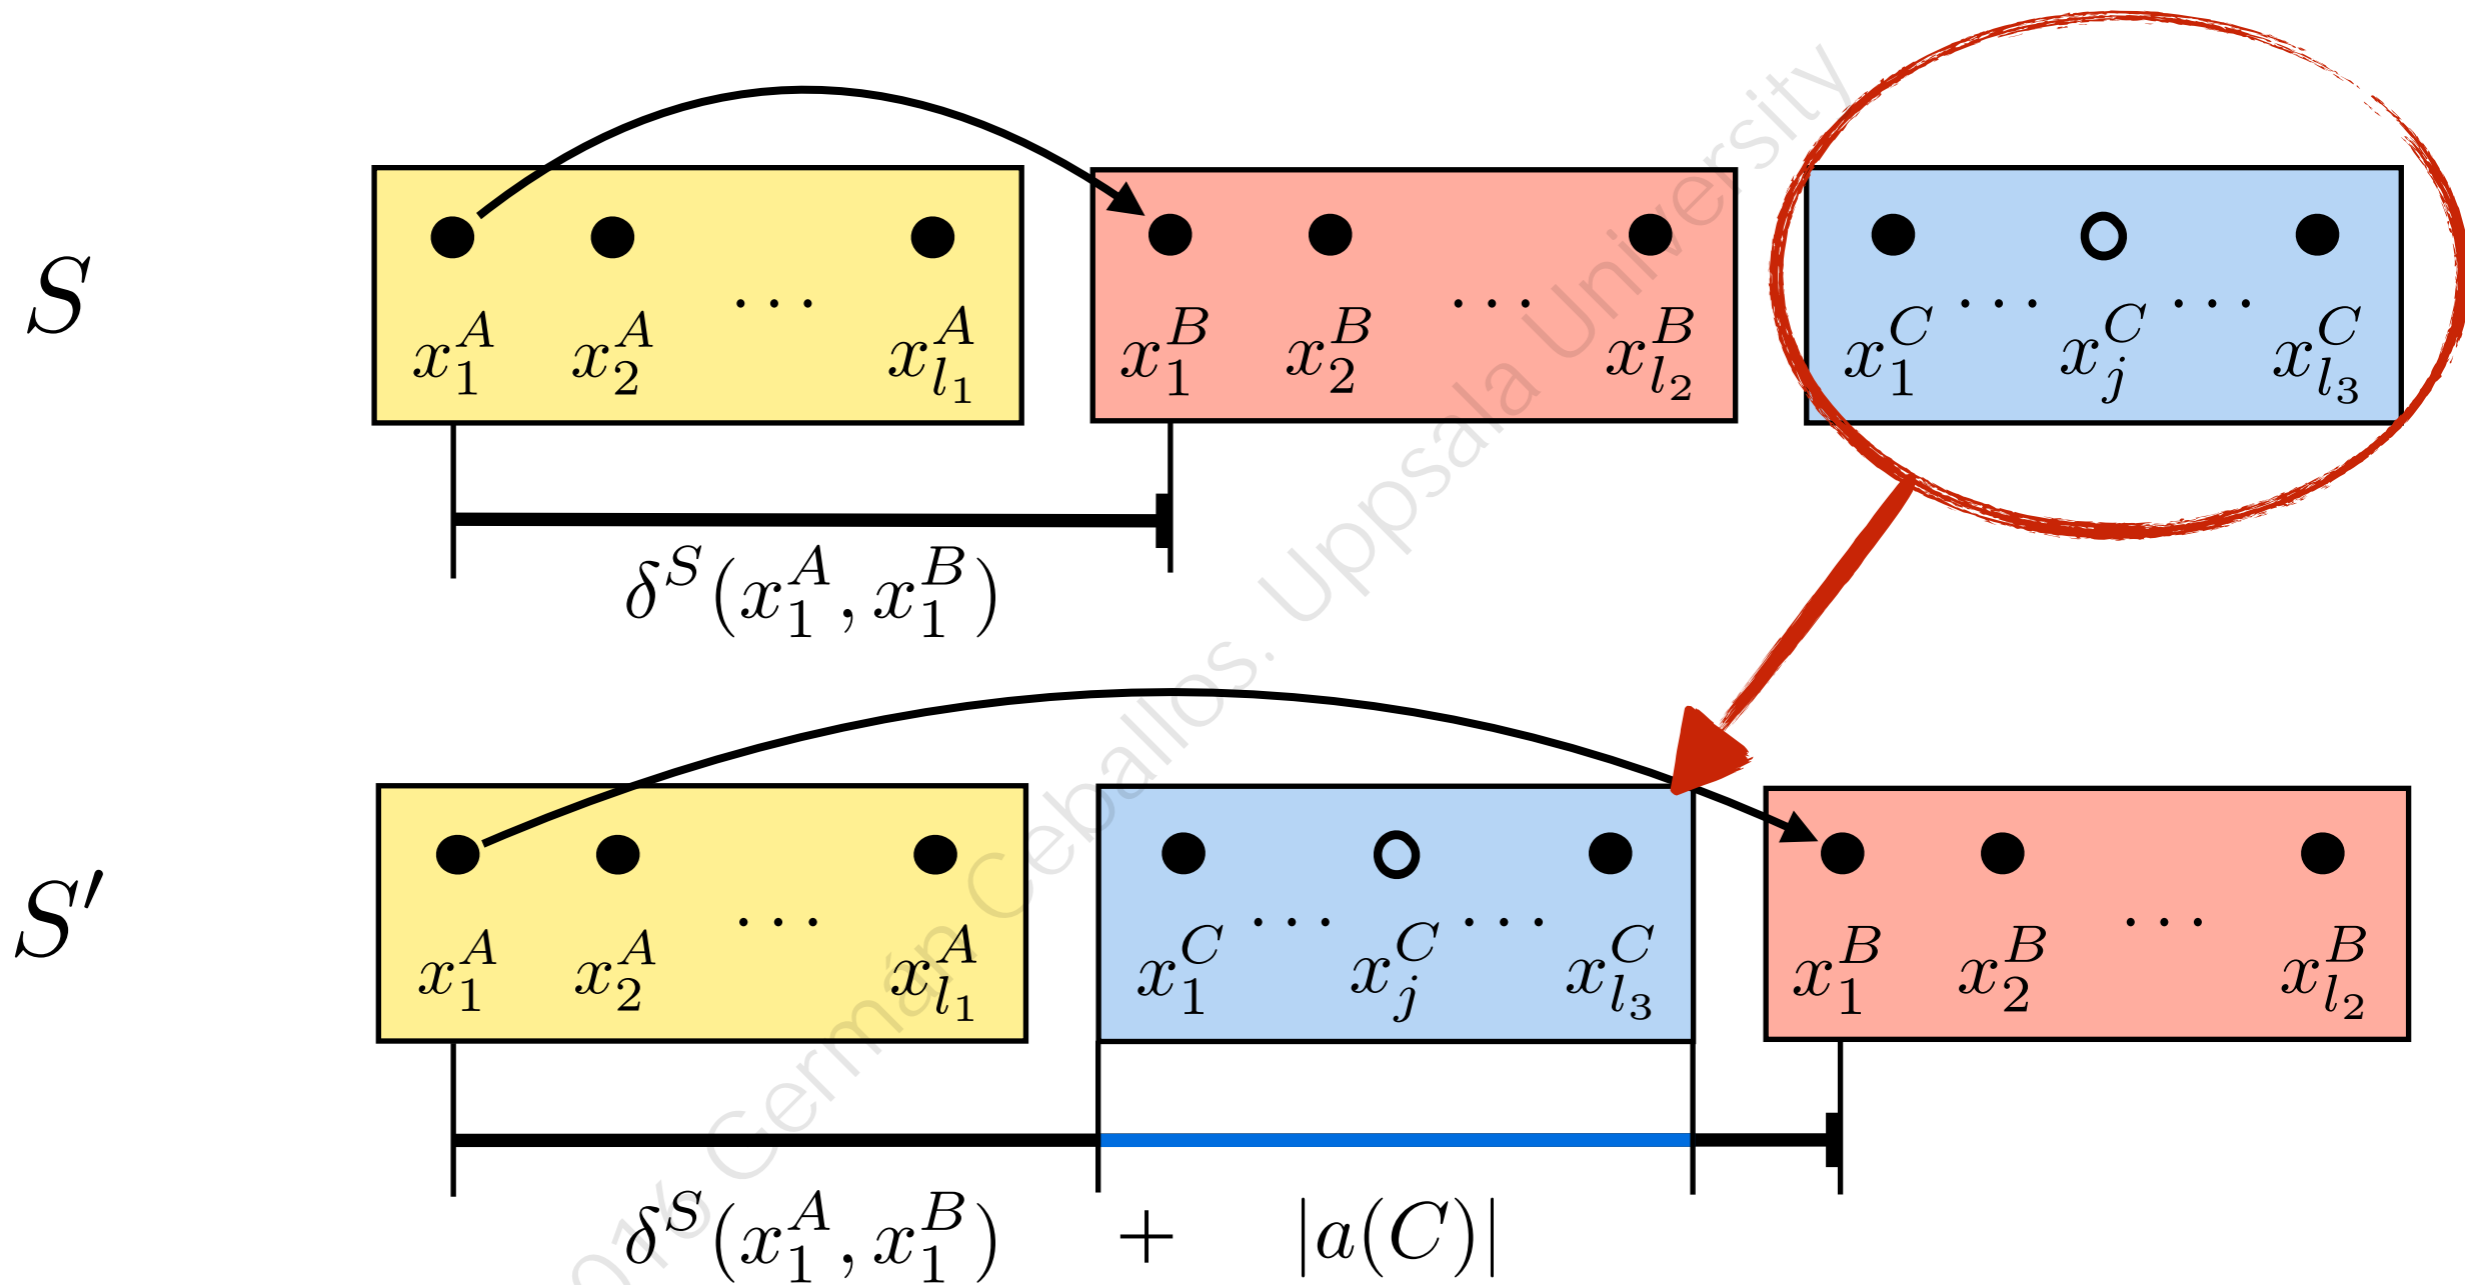
SCM with Tasks



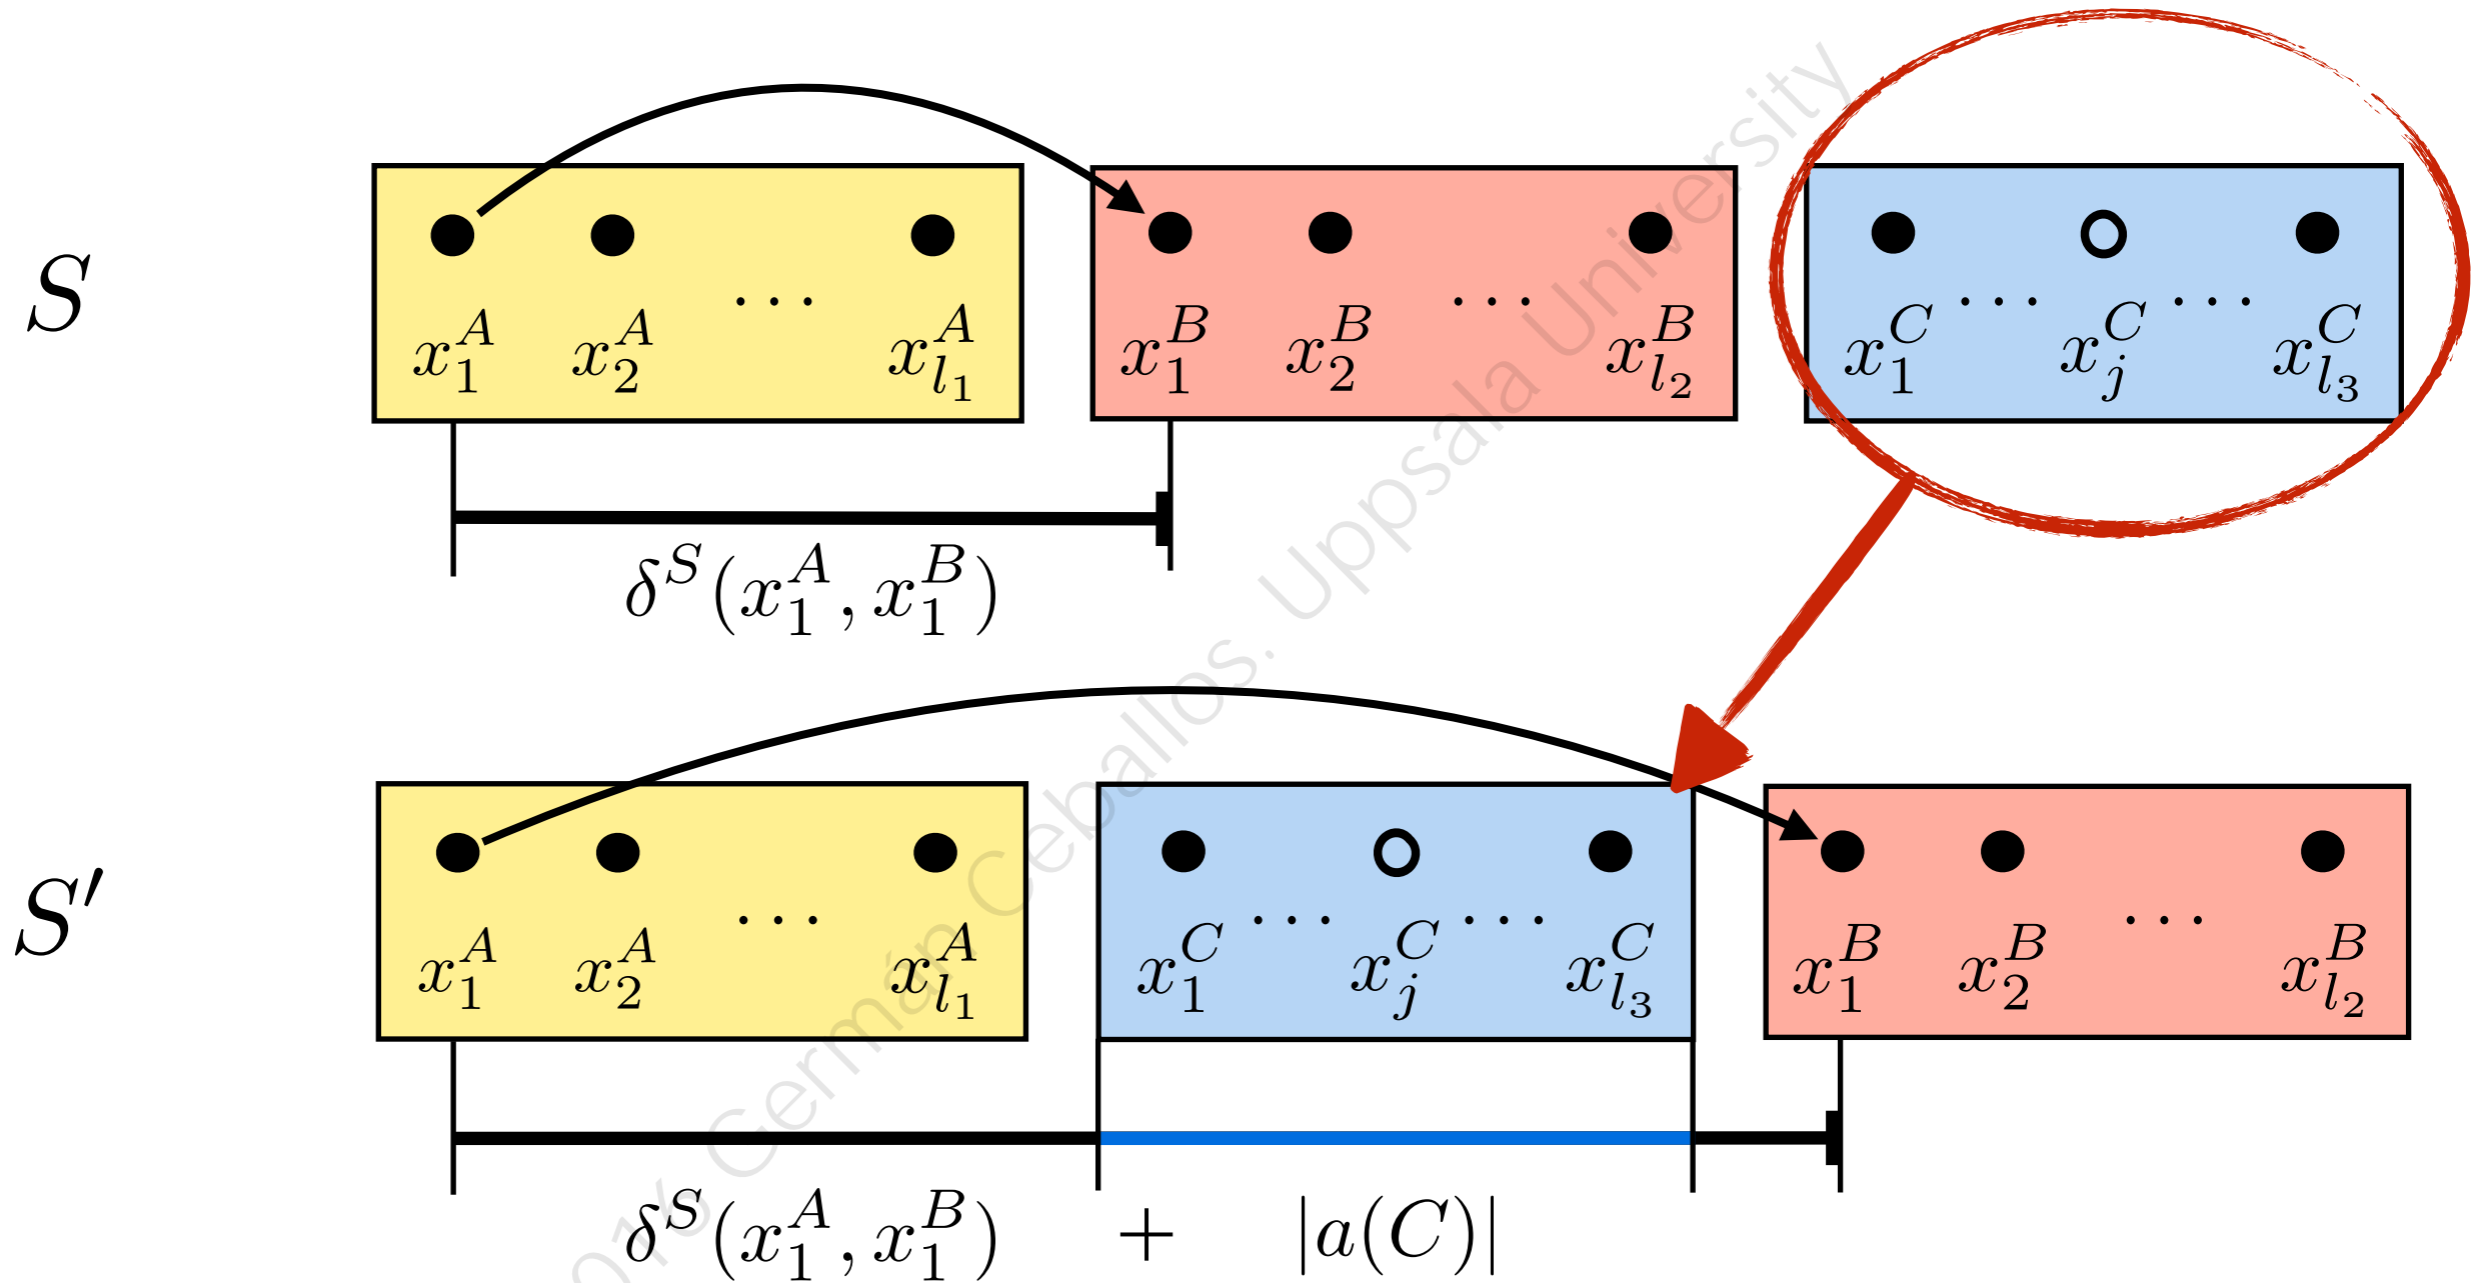
SCM with Tasks



SCM with Tasks



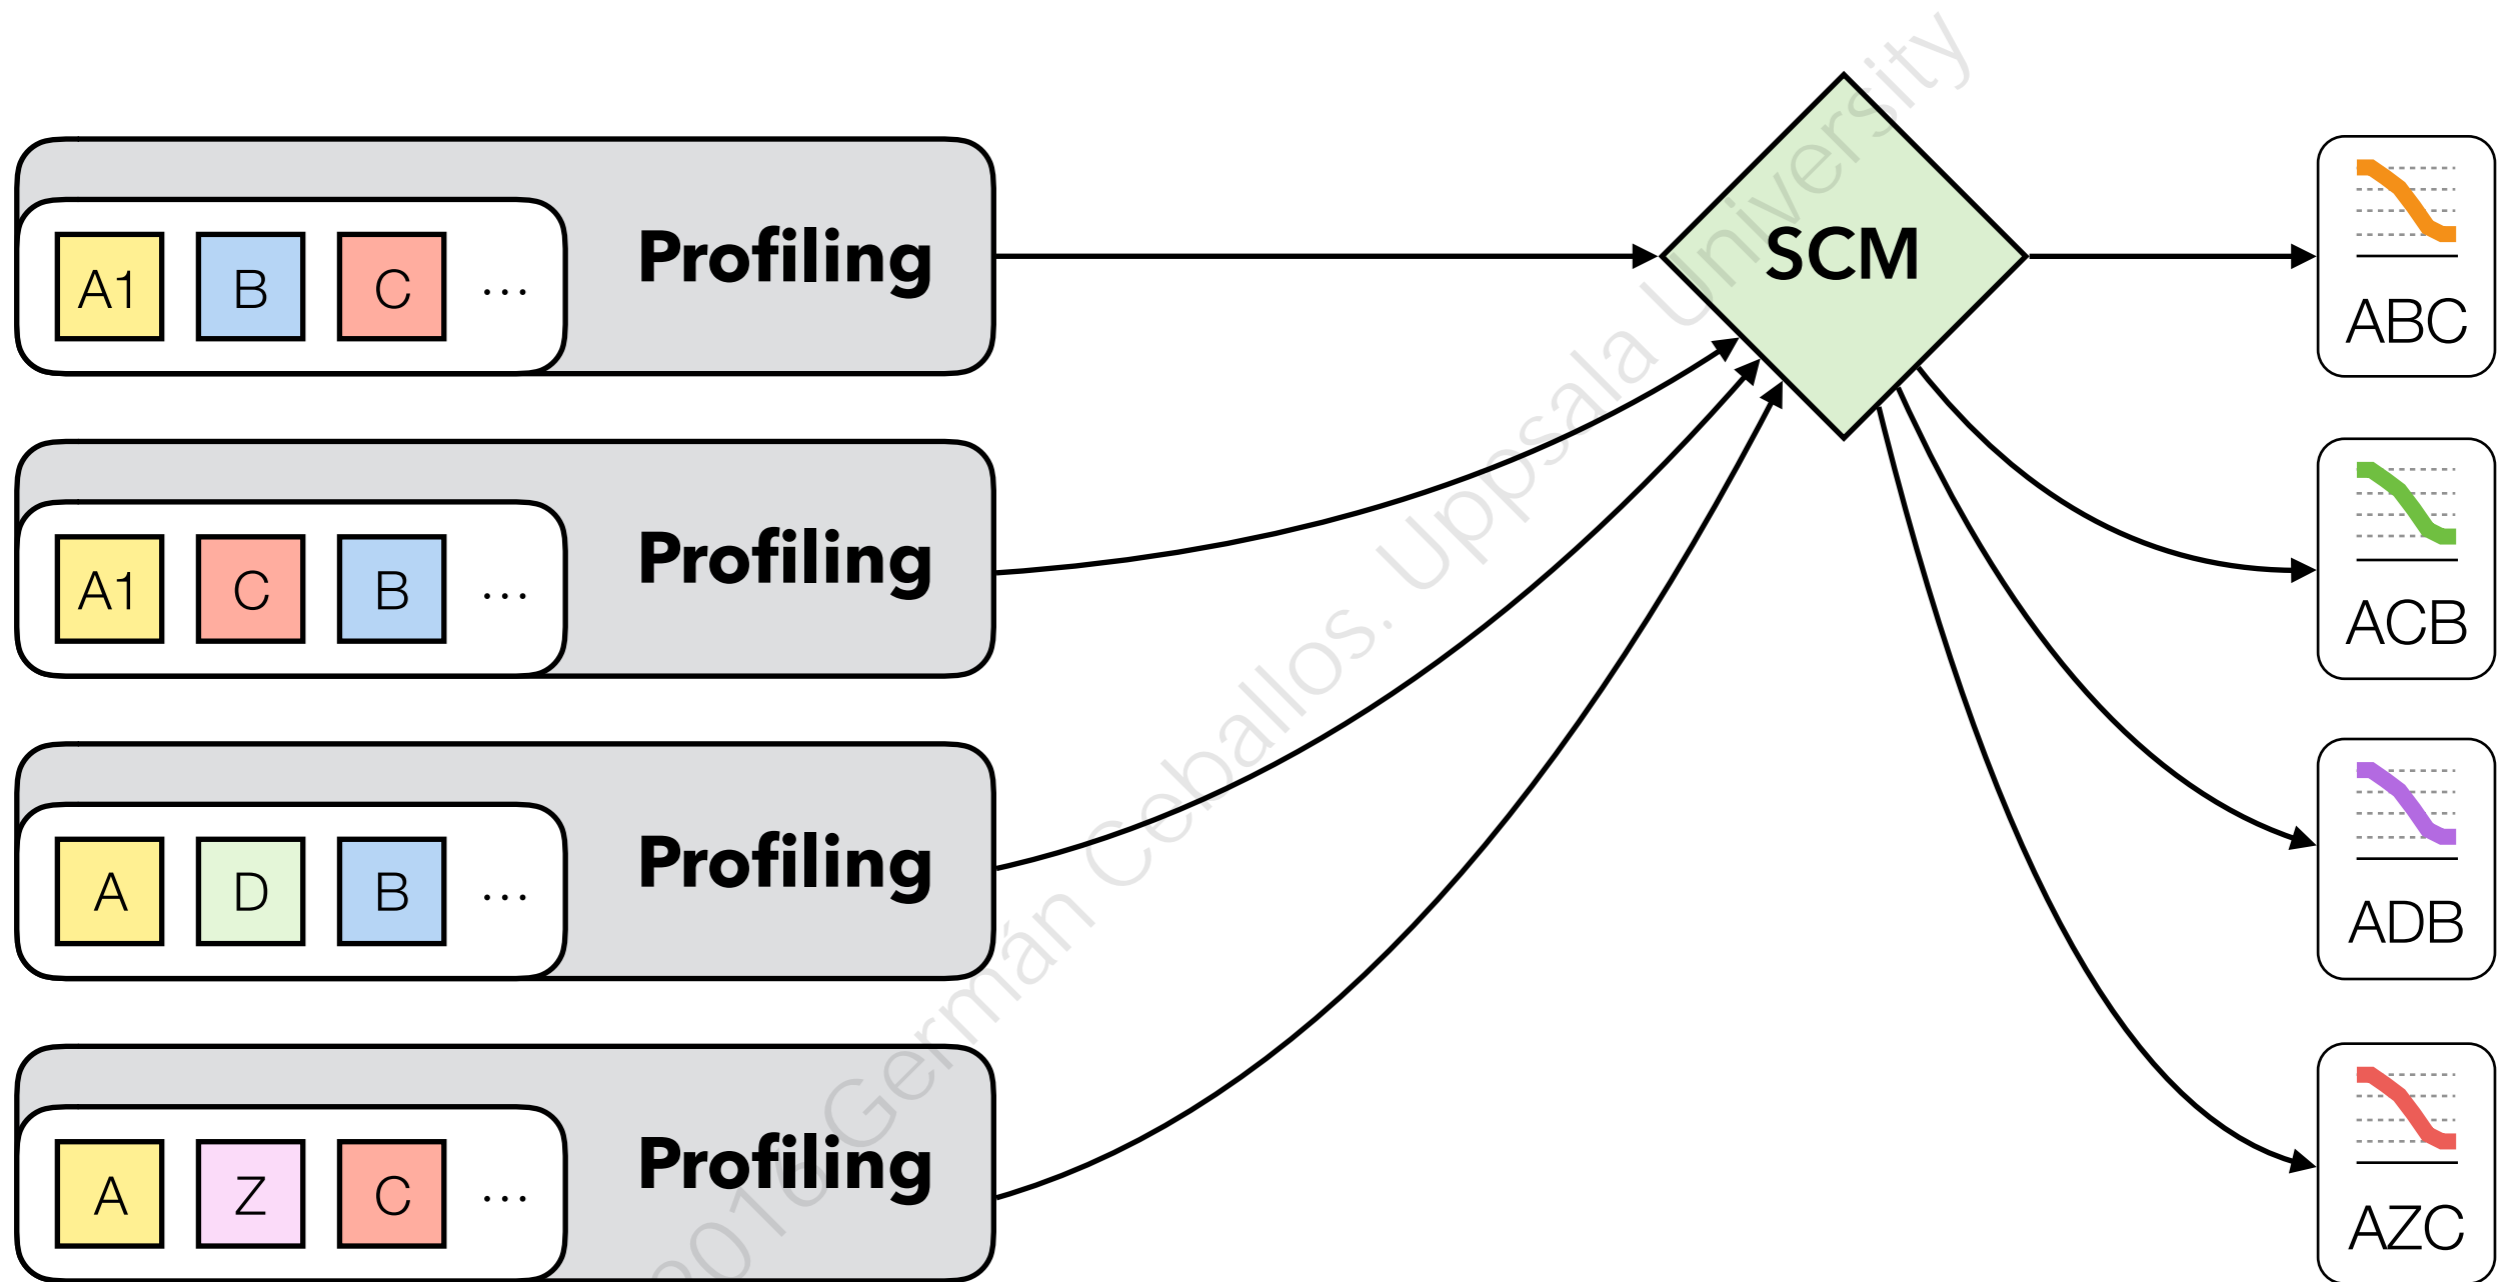
SCM with Tasks



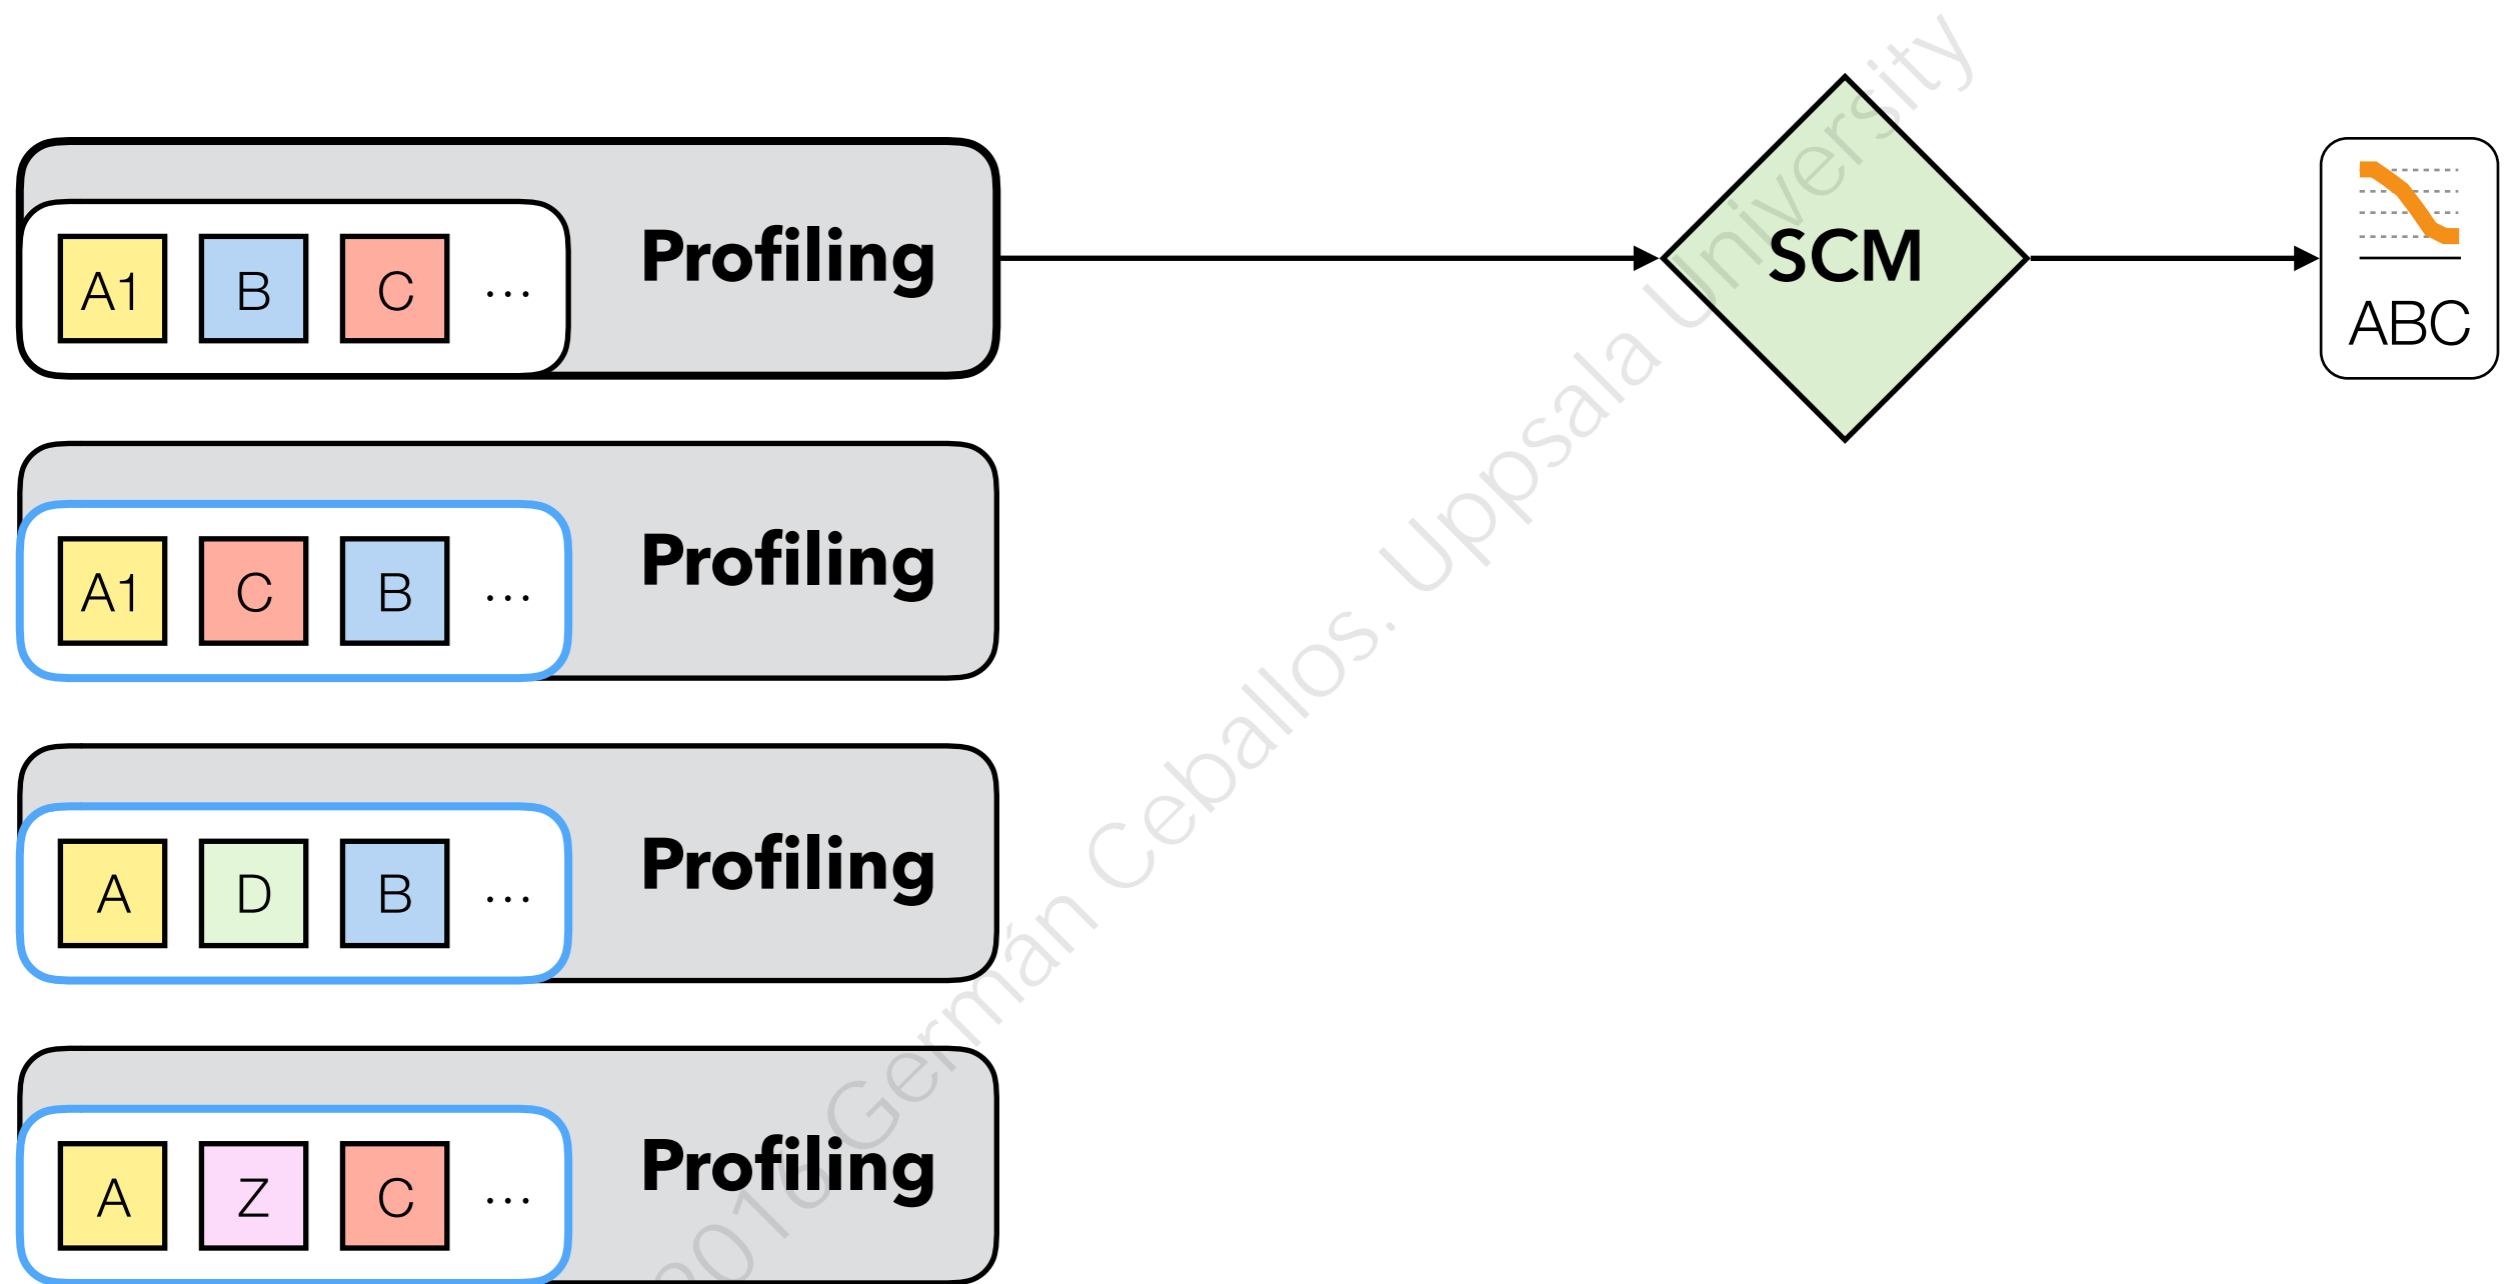
Distances can change!



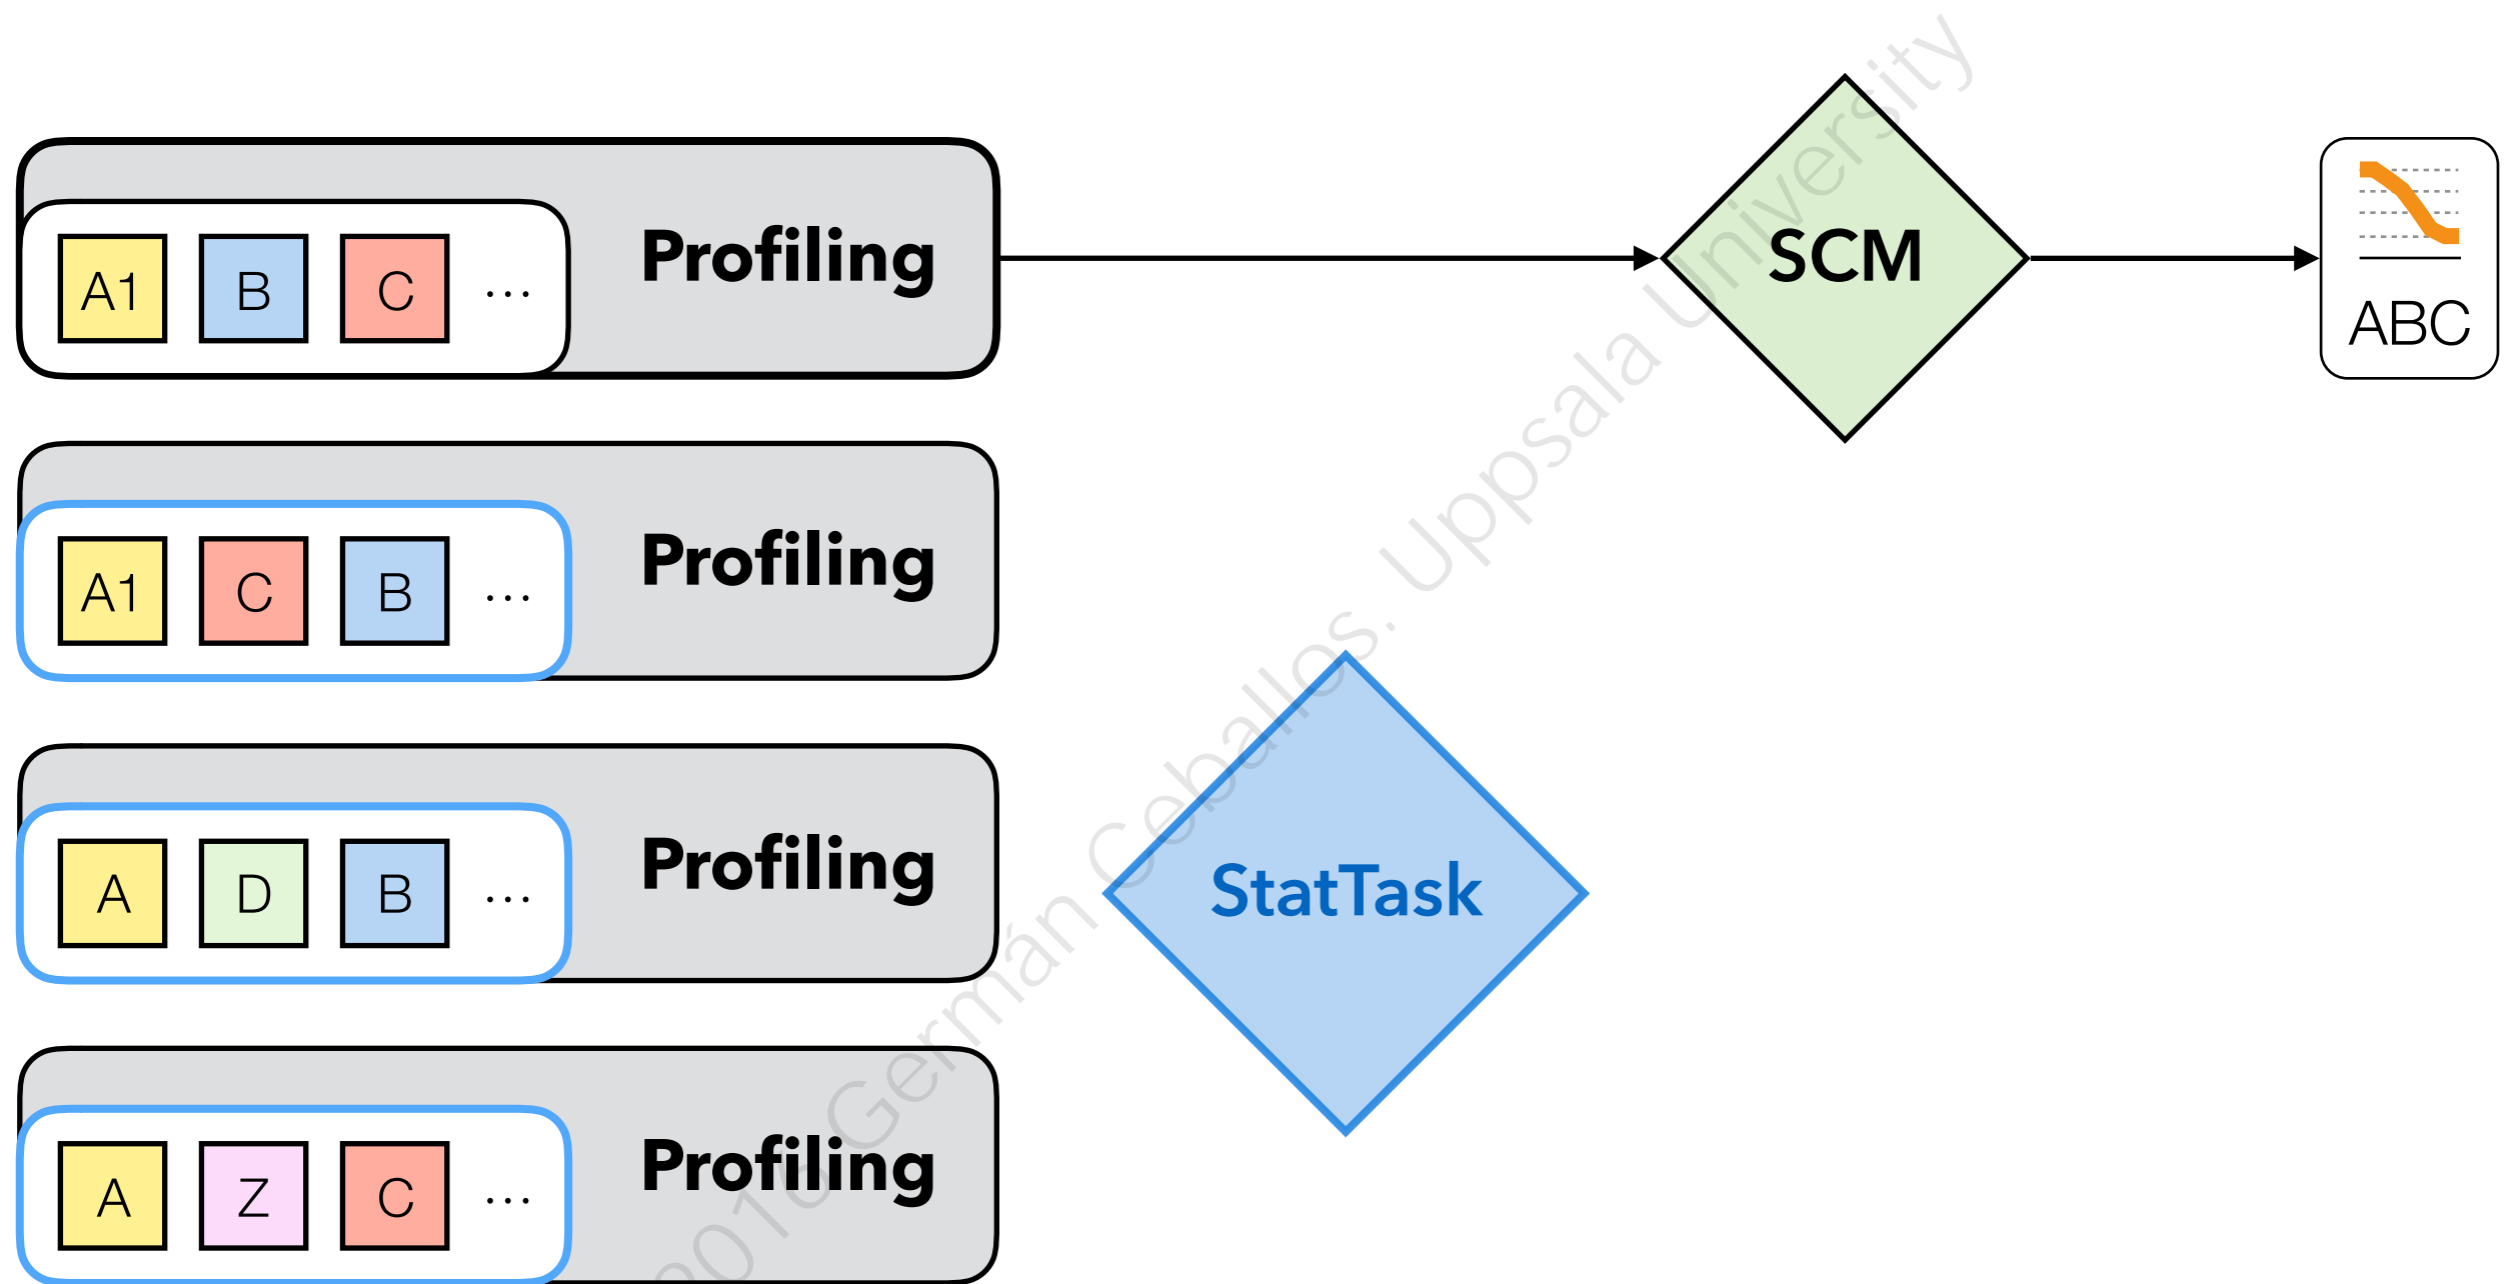
The Problem



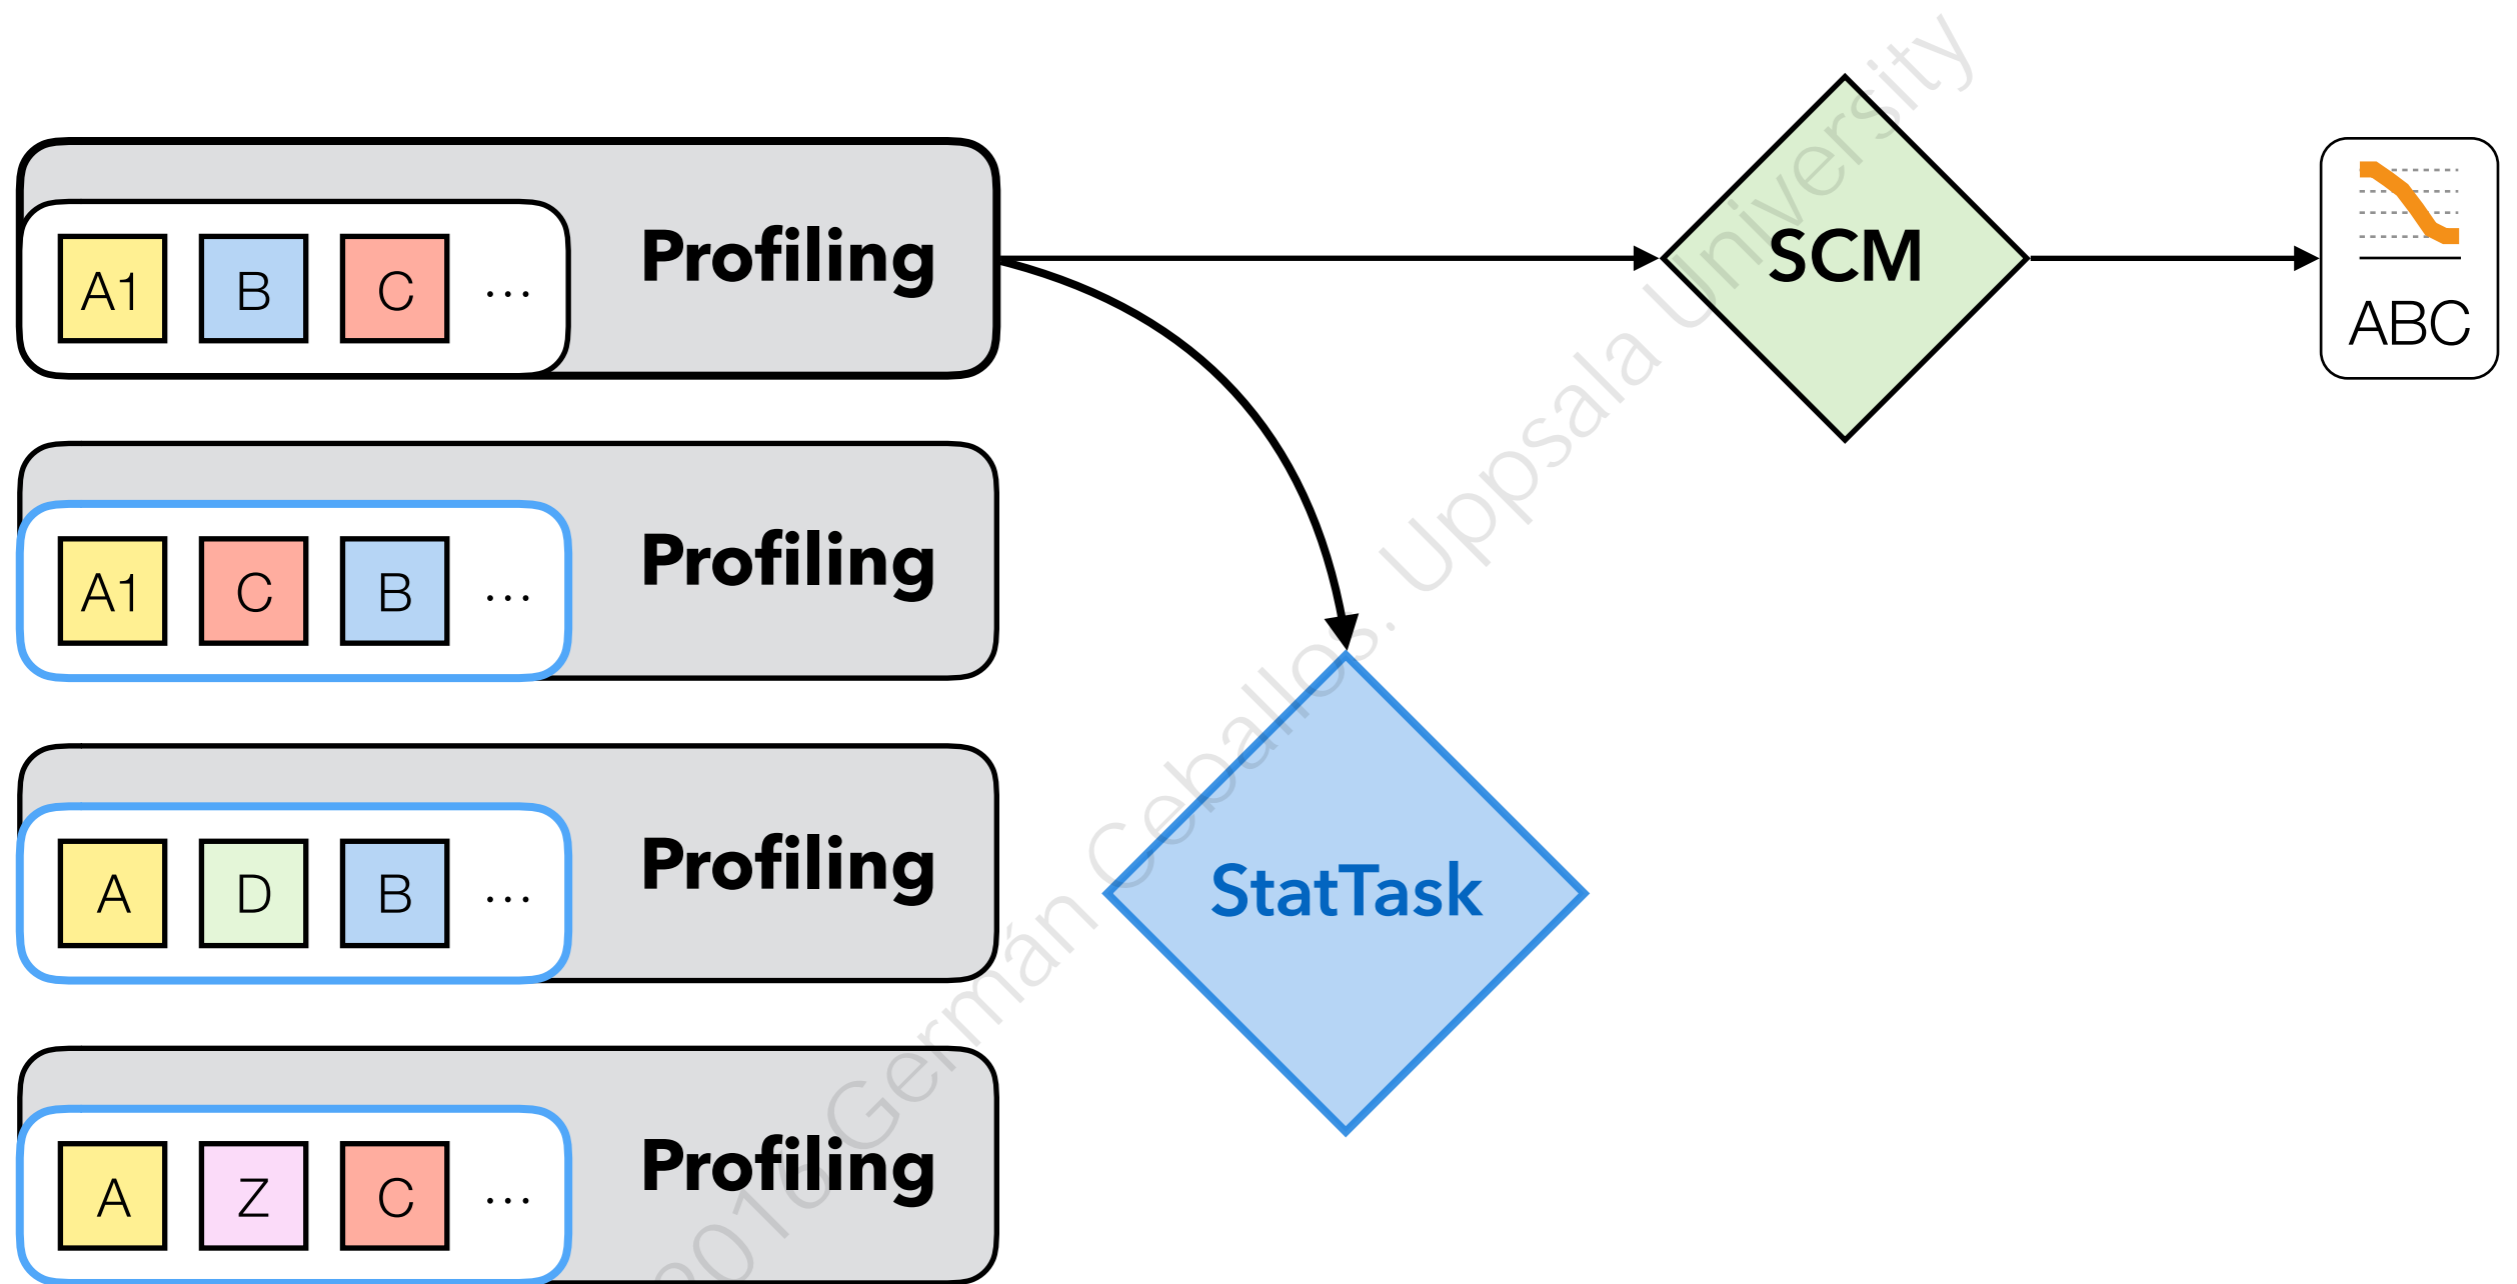
The Solution



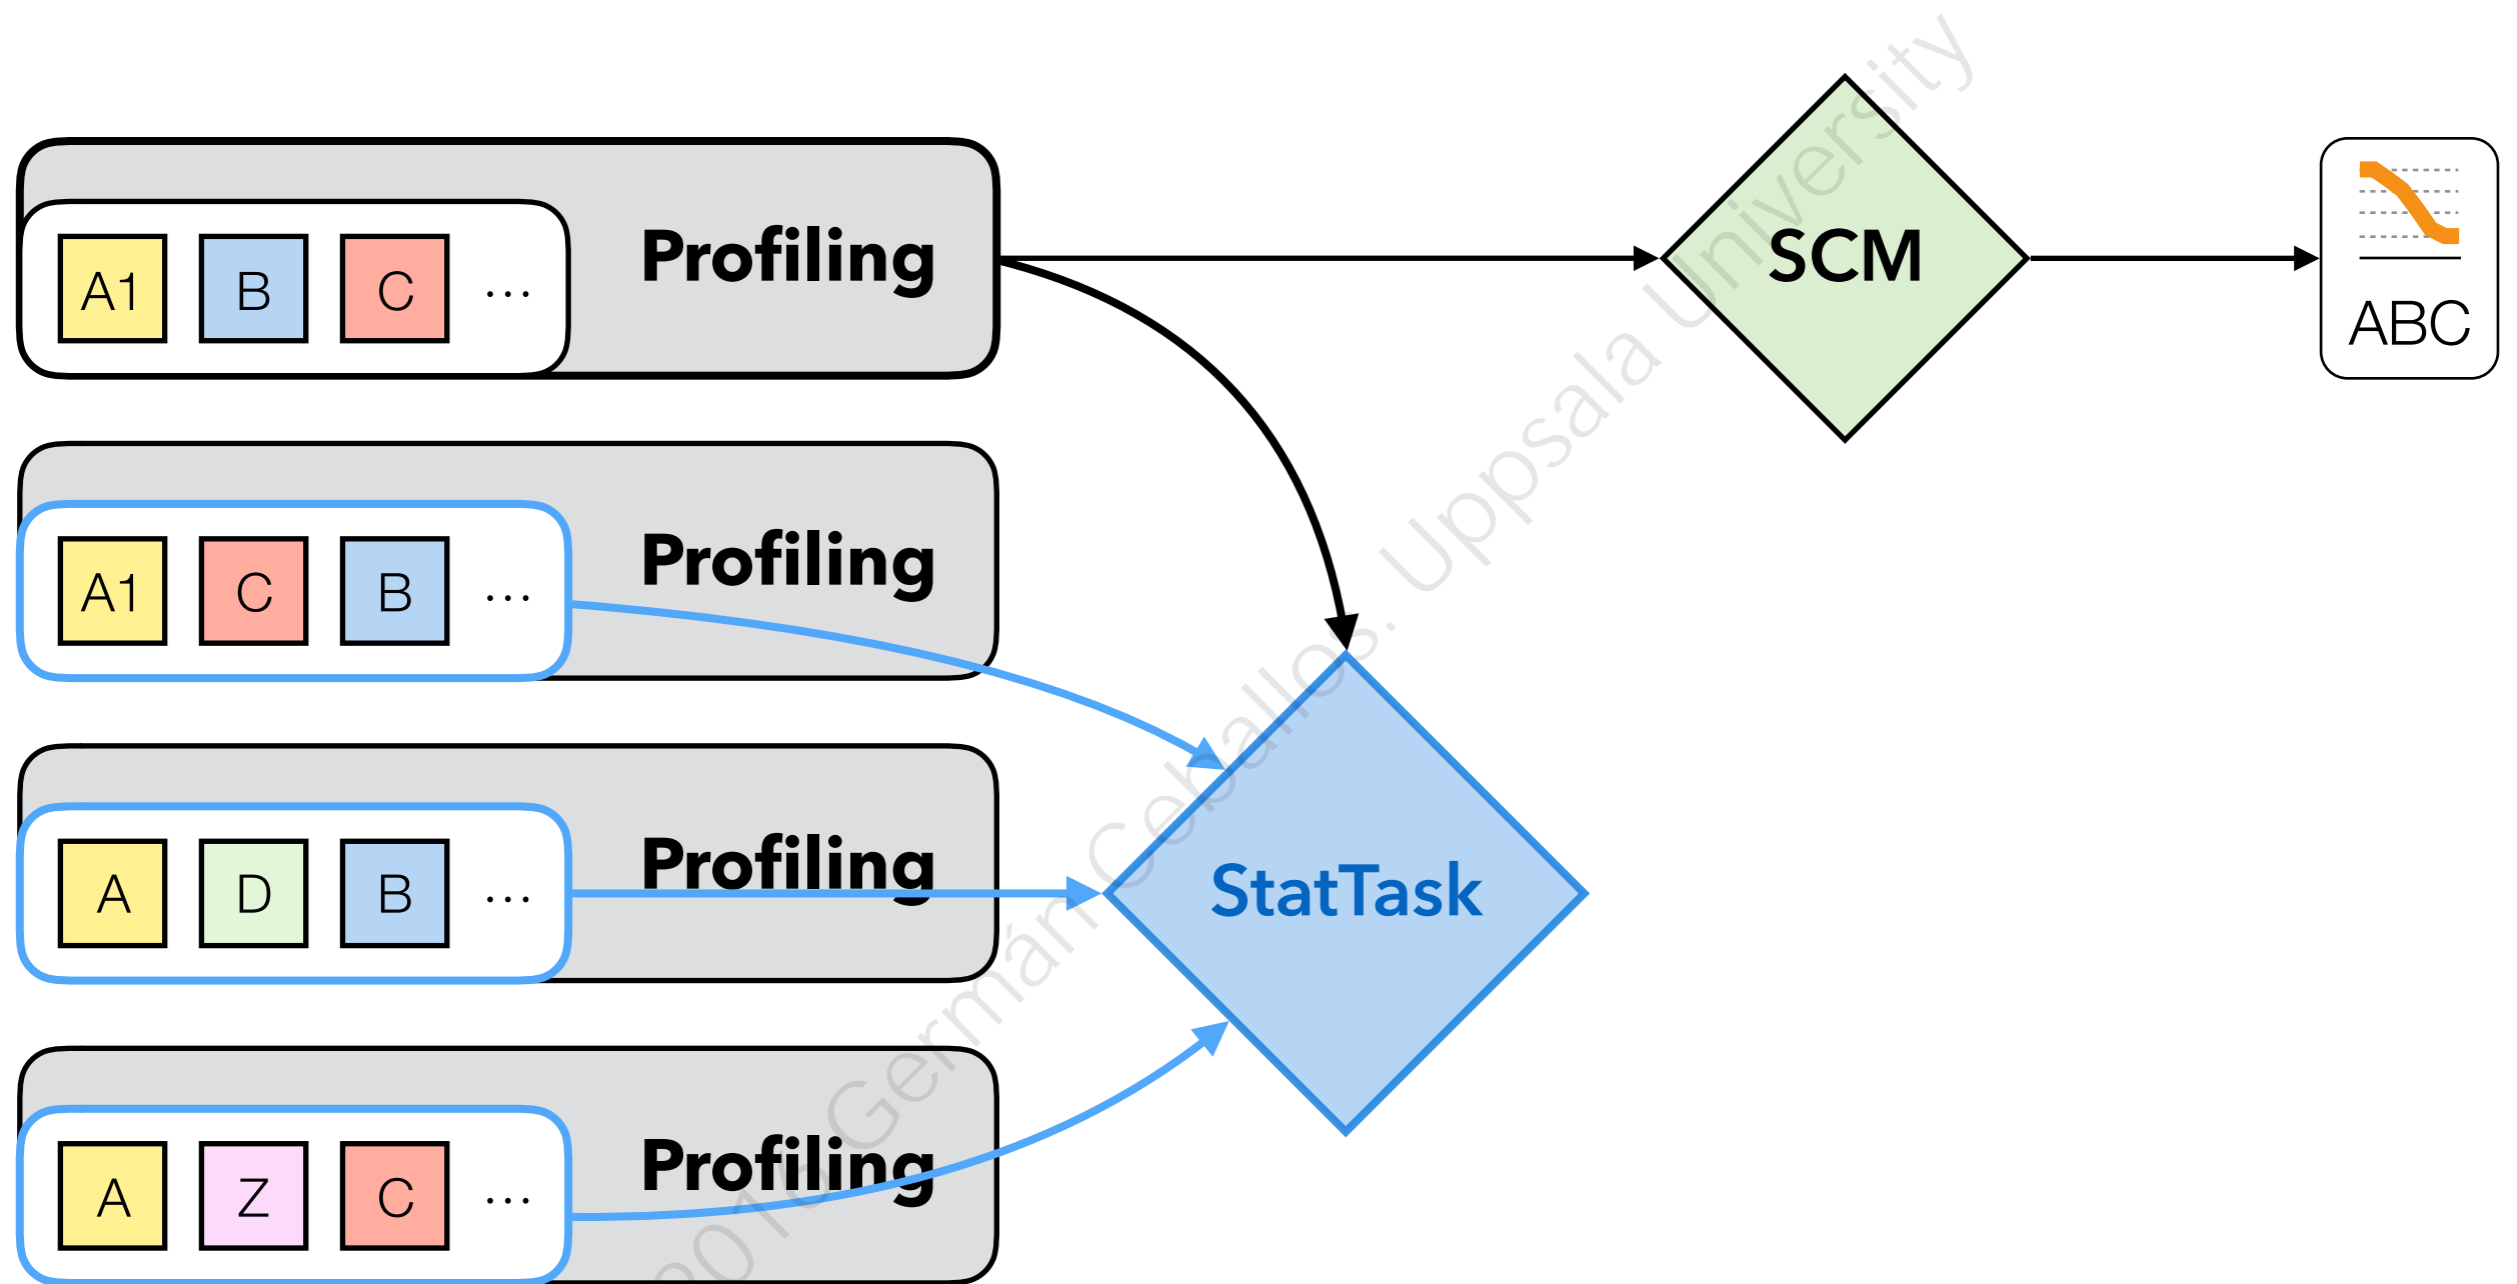
The Solution



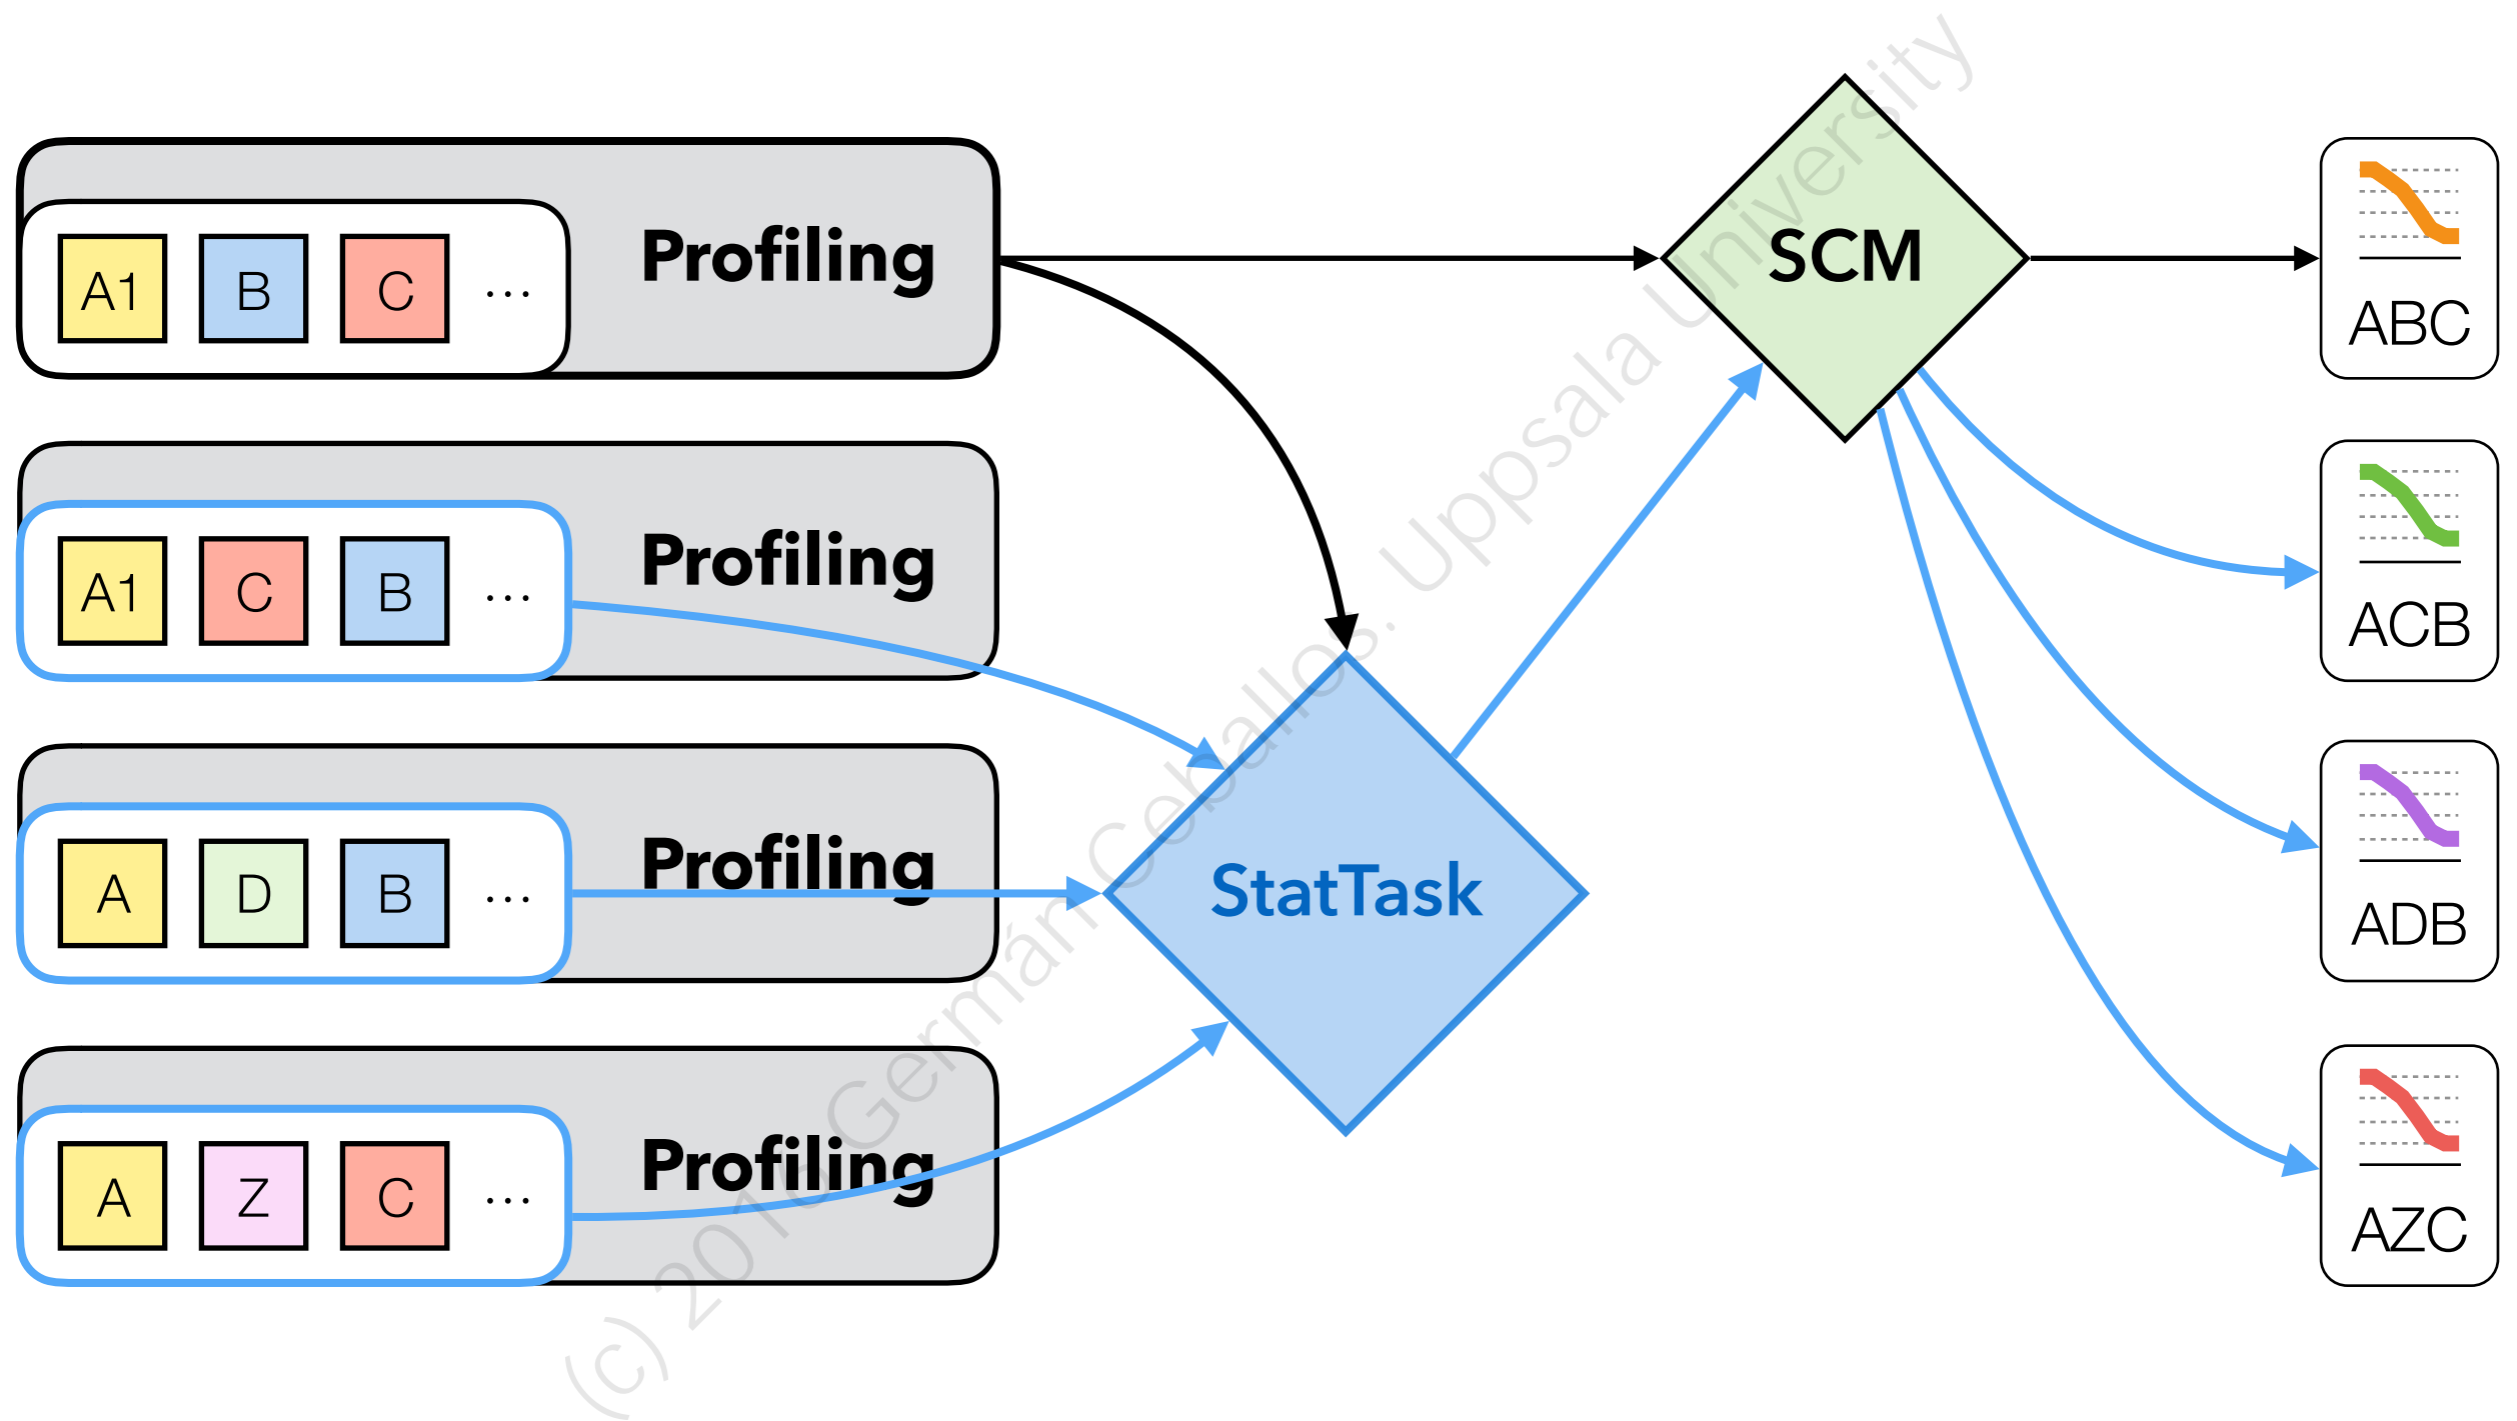
The Solution



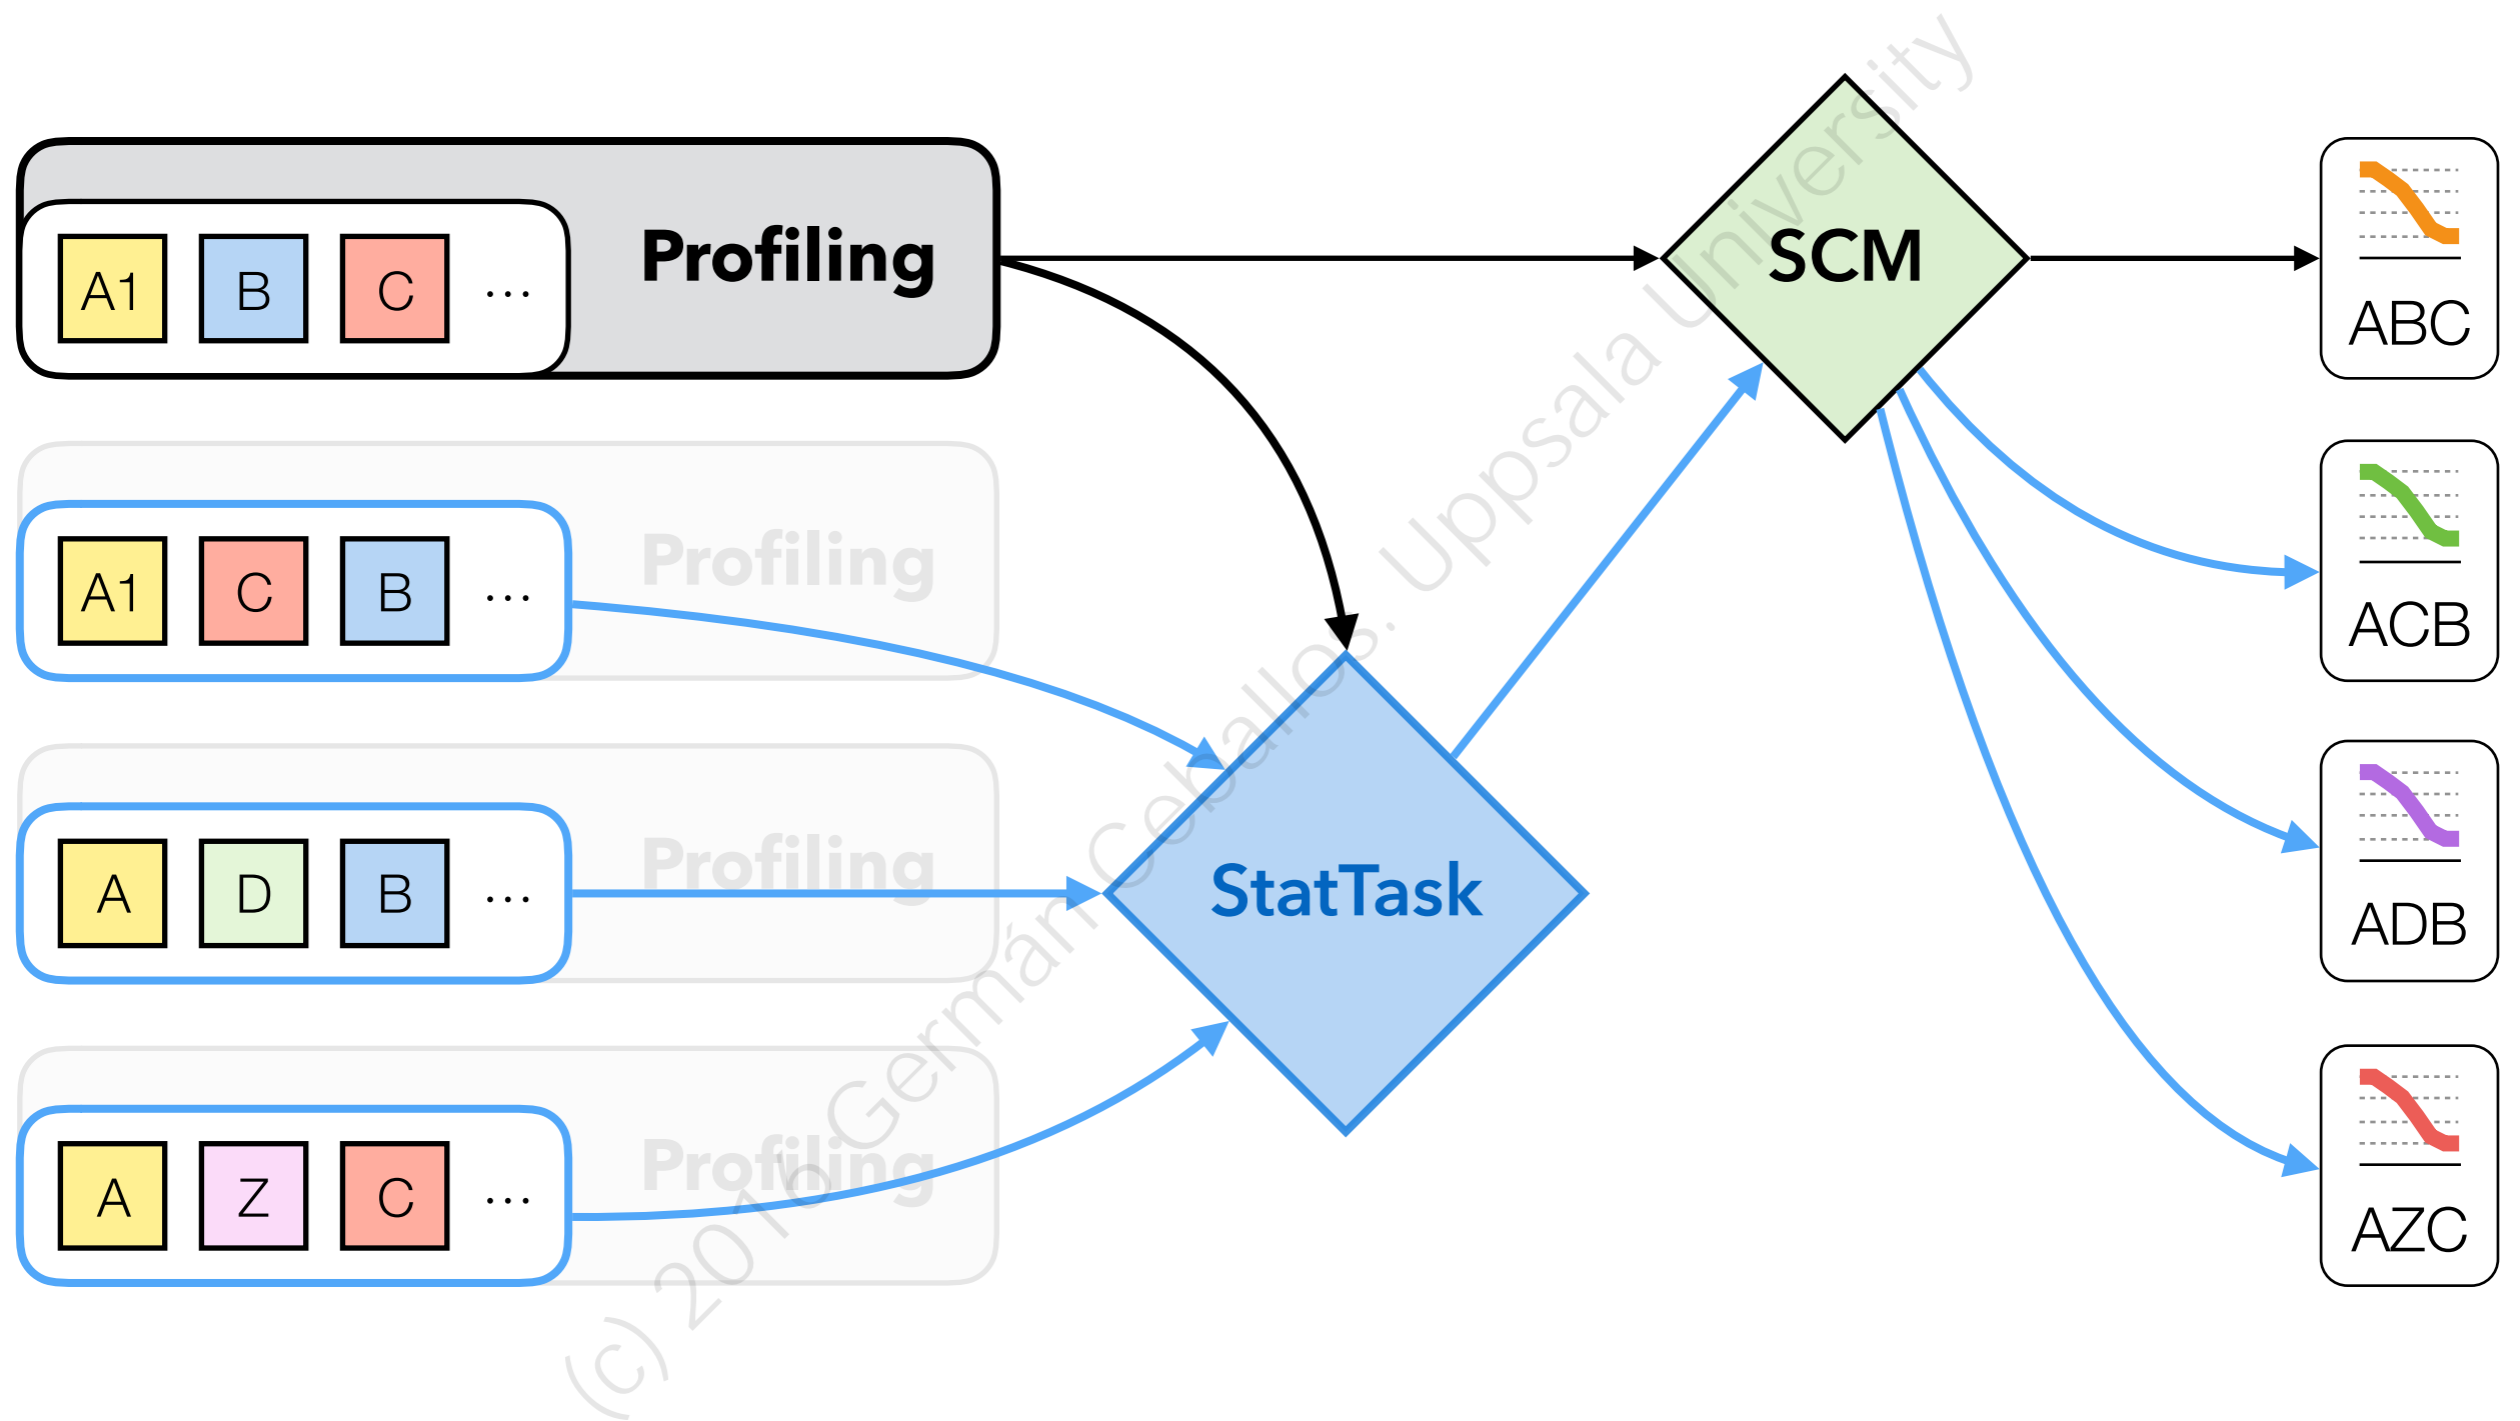
The Solution



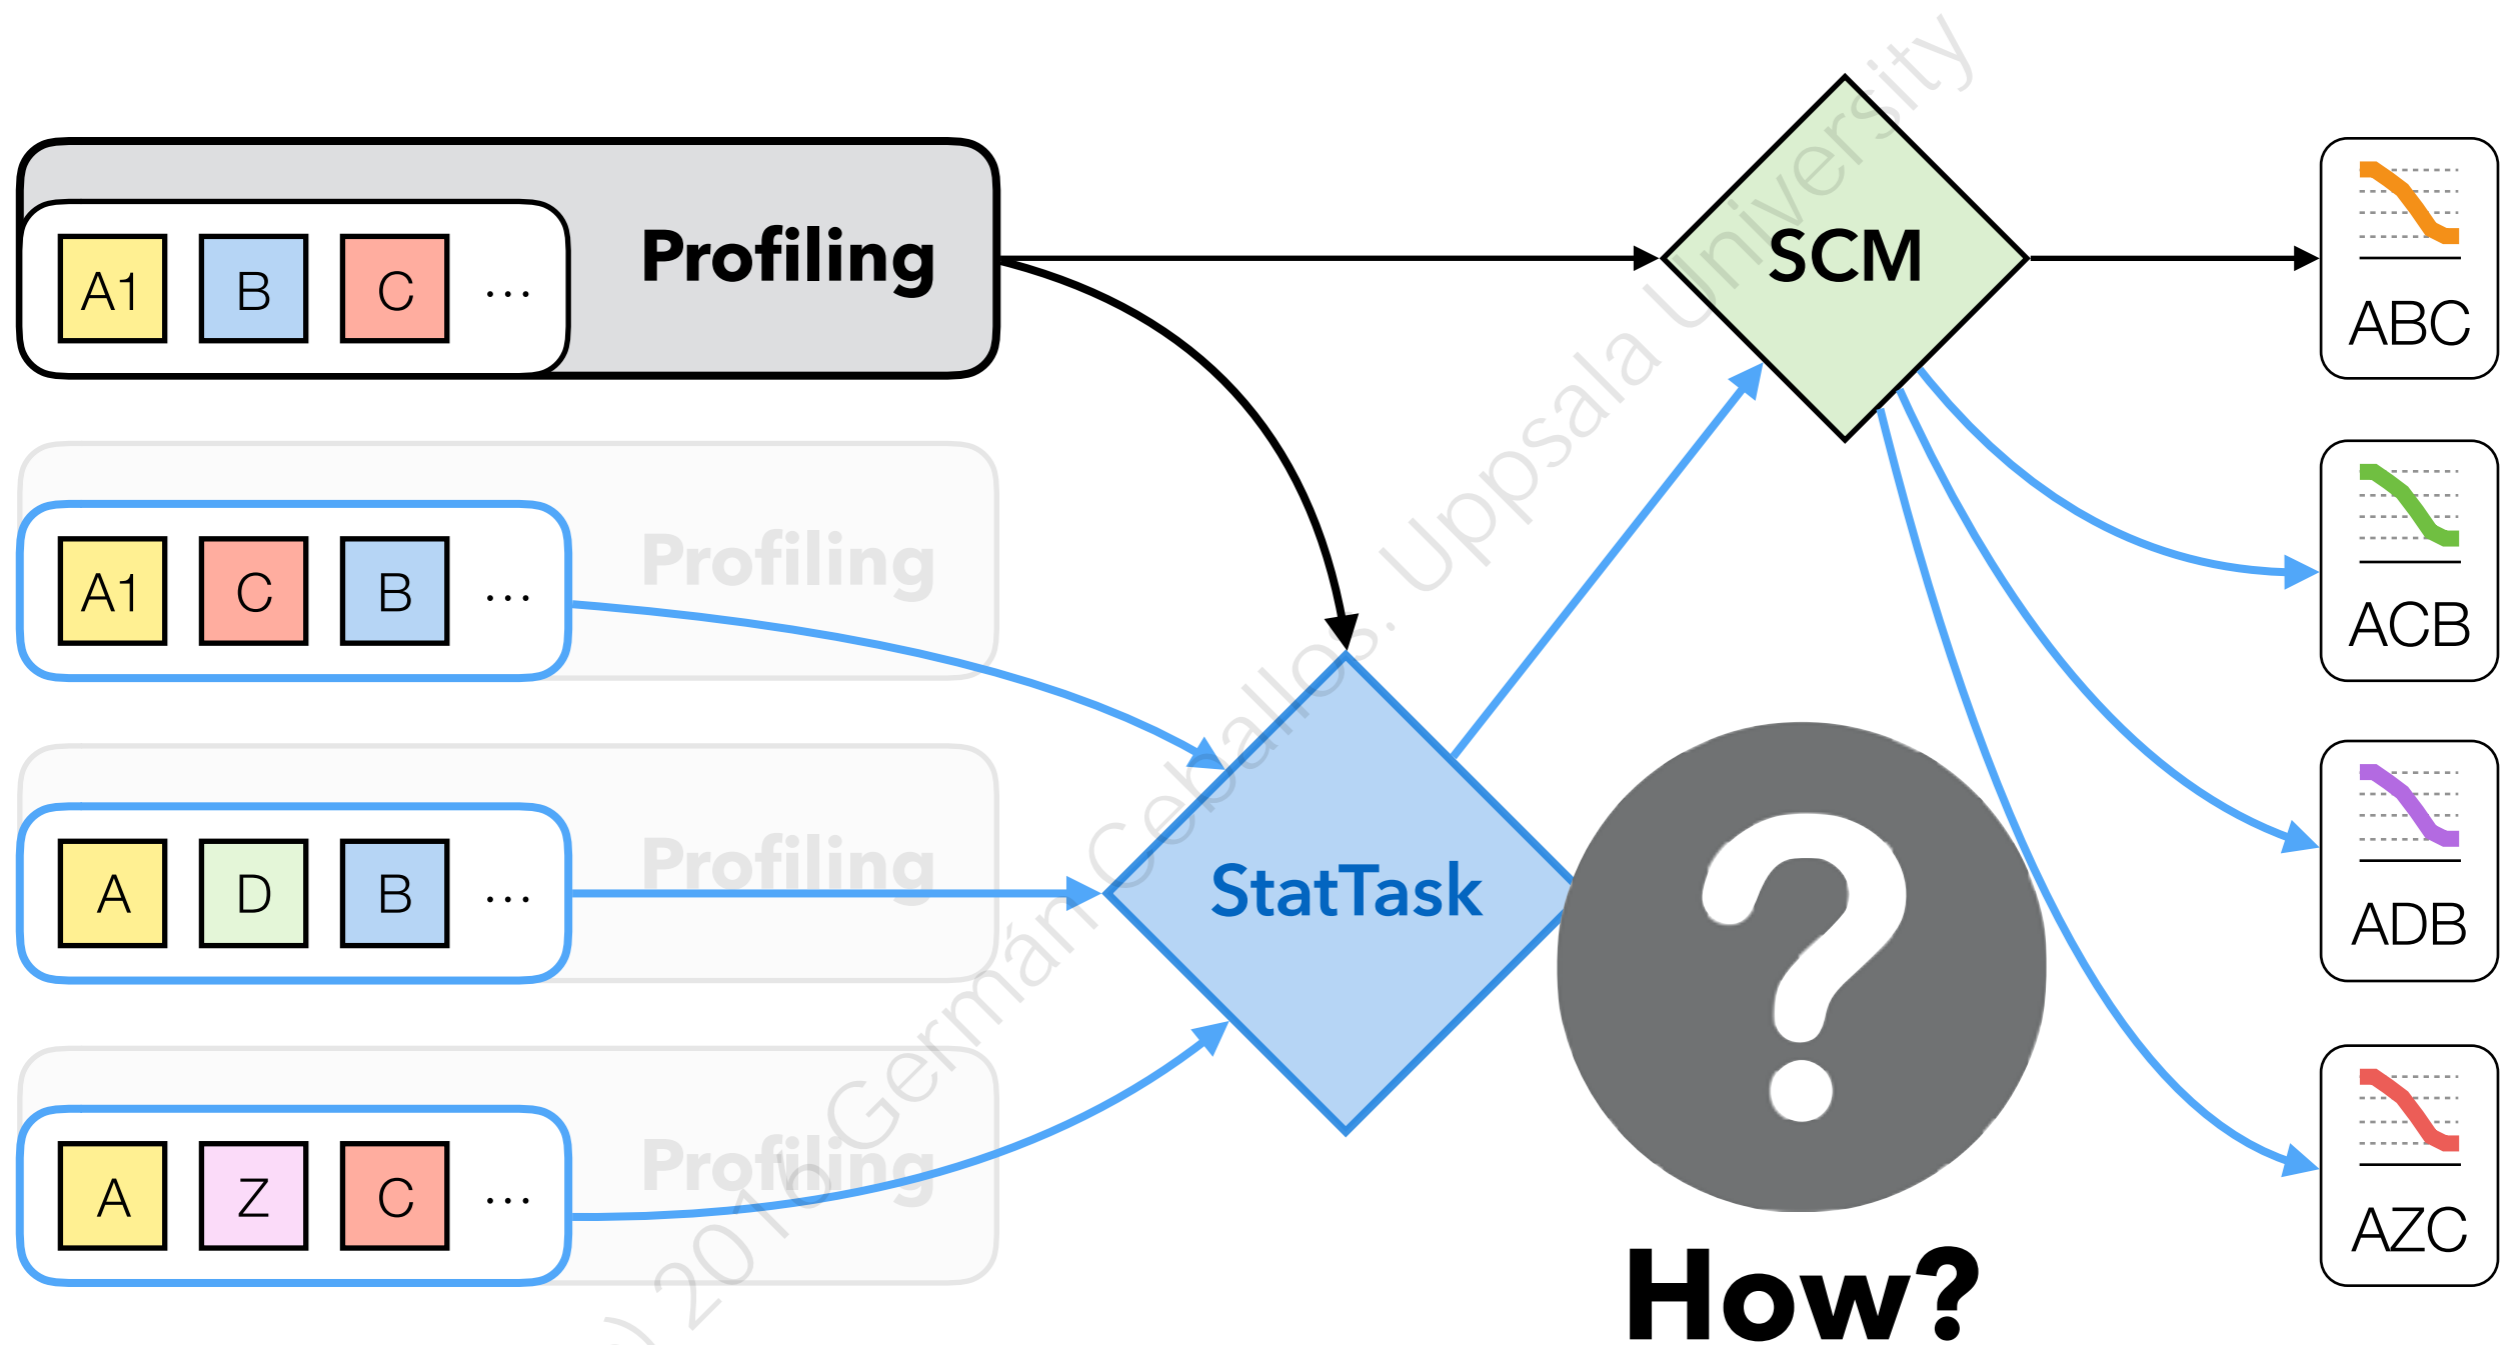
The Solution



The Solution



The Solution



StatTask Overview

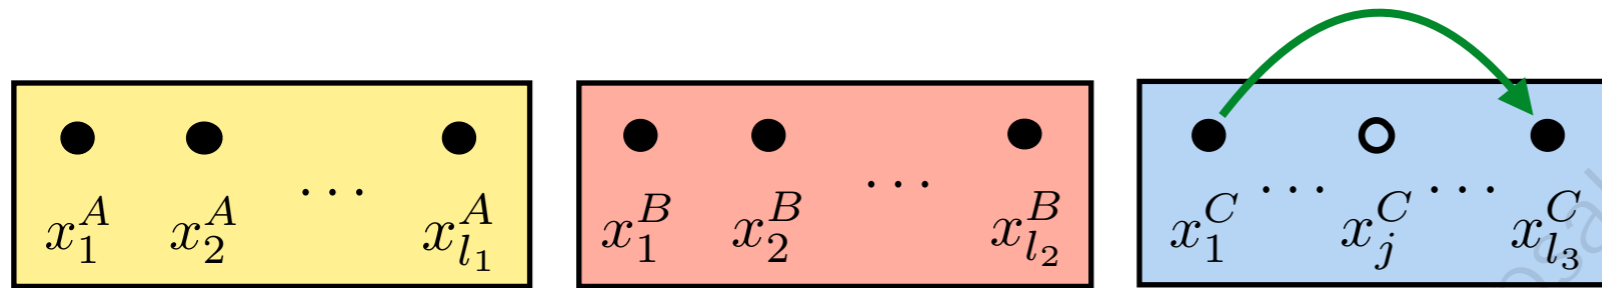


1. Profile **one** schedule for the application
2. Build the pair of reuses
3. Input a **new** schedule
4. Rebuild the pair of reuses for the new schedule
5. Recalculate the reuse distances
6. Use statistical cache models to estimate miss ratios



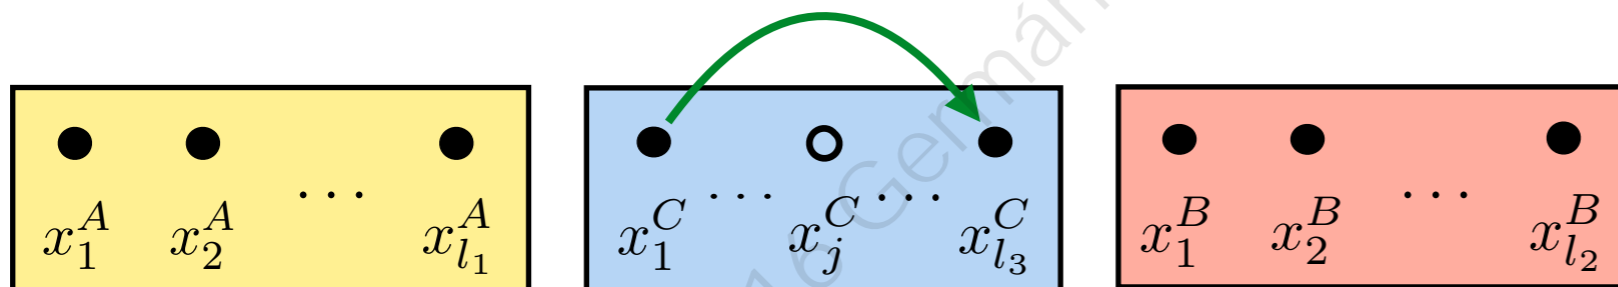


Rebuild Reuses



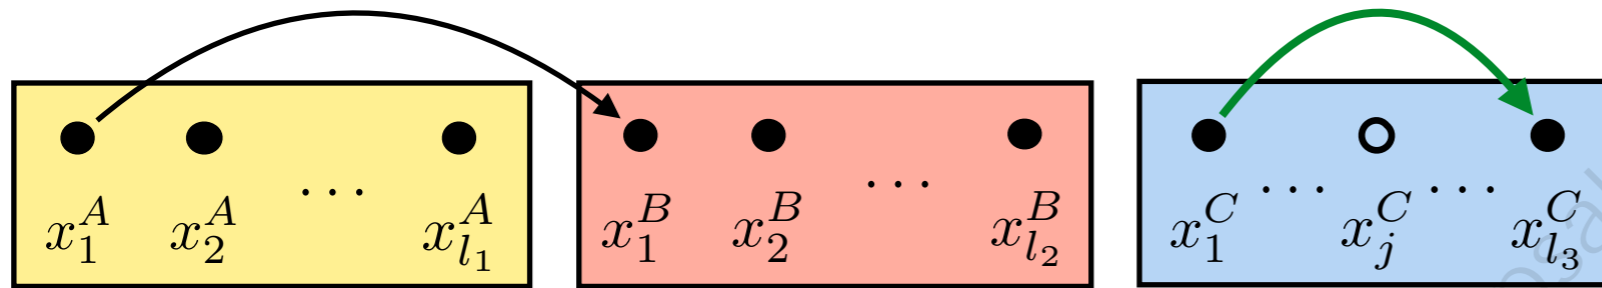
1. Private Reuses

- All are kept





Rebuild Reuses

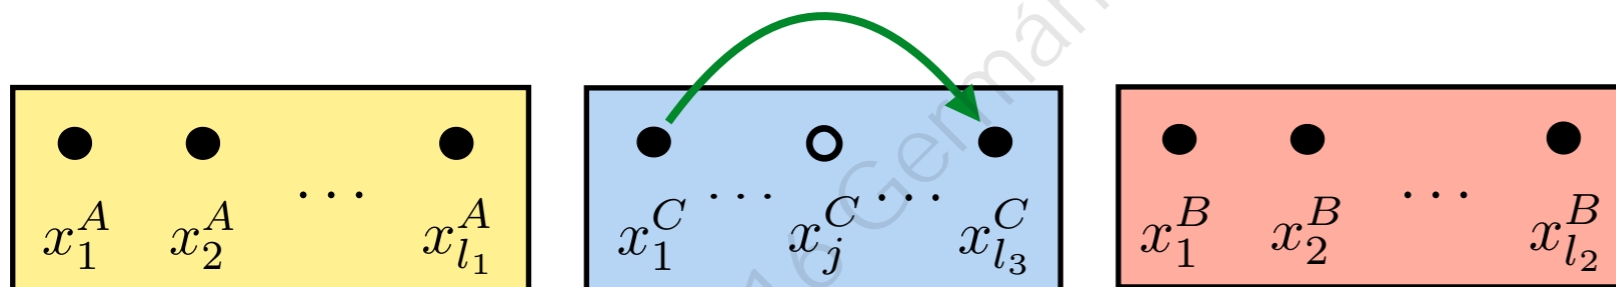


1. Private Reuses

- All are kept

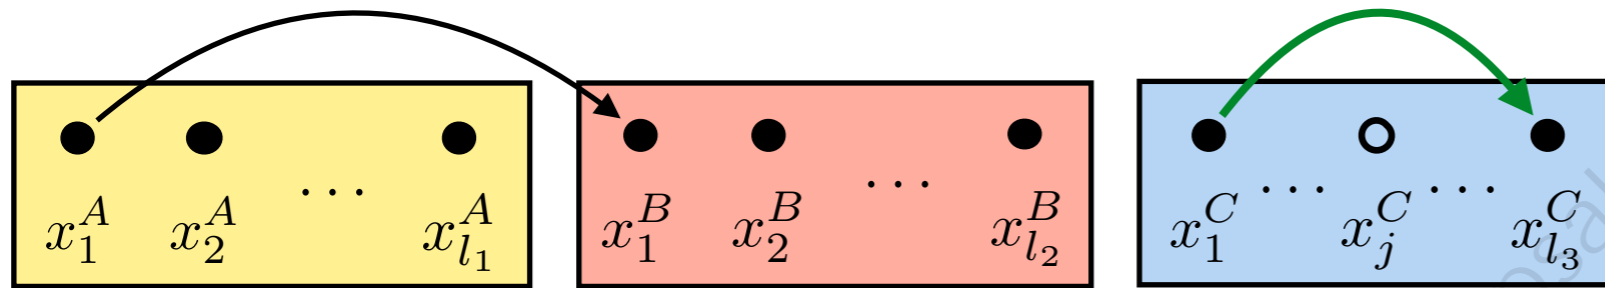
2. Shared Reuses

- Some are **kept**





Rebuild Reuses

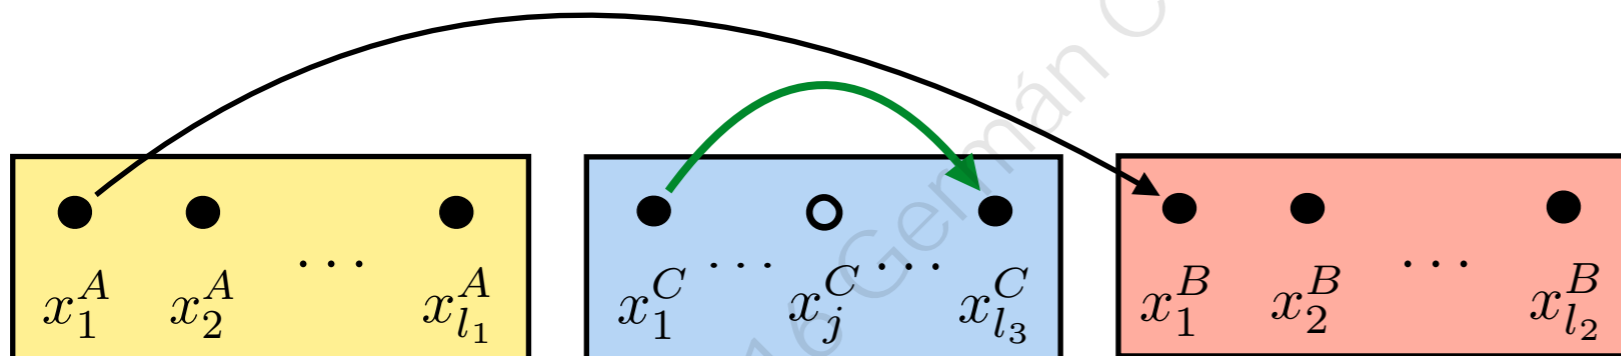


1. Private Reuses

- All are kept

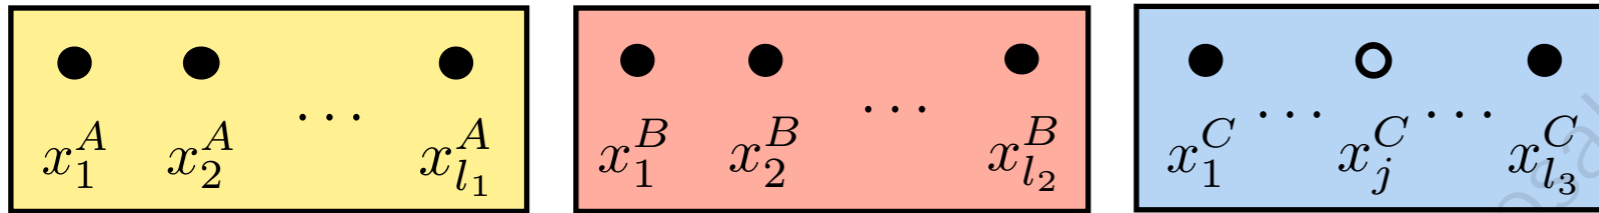
2. Shared Reuses

- Some are **kept**





Rebuild Reuses

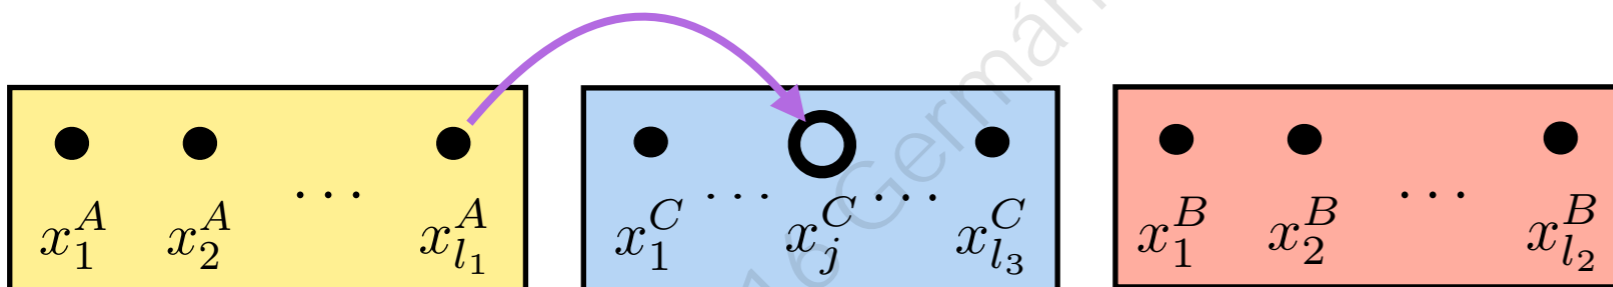


1. Private Reuses

- All are kept

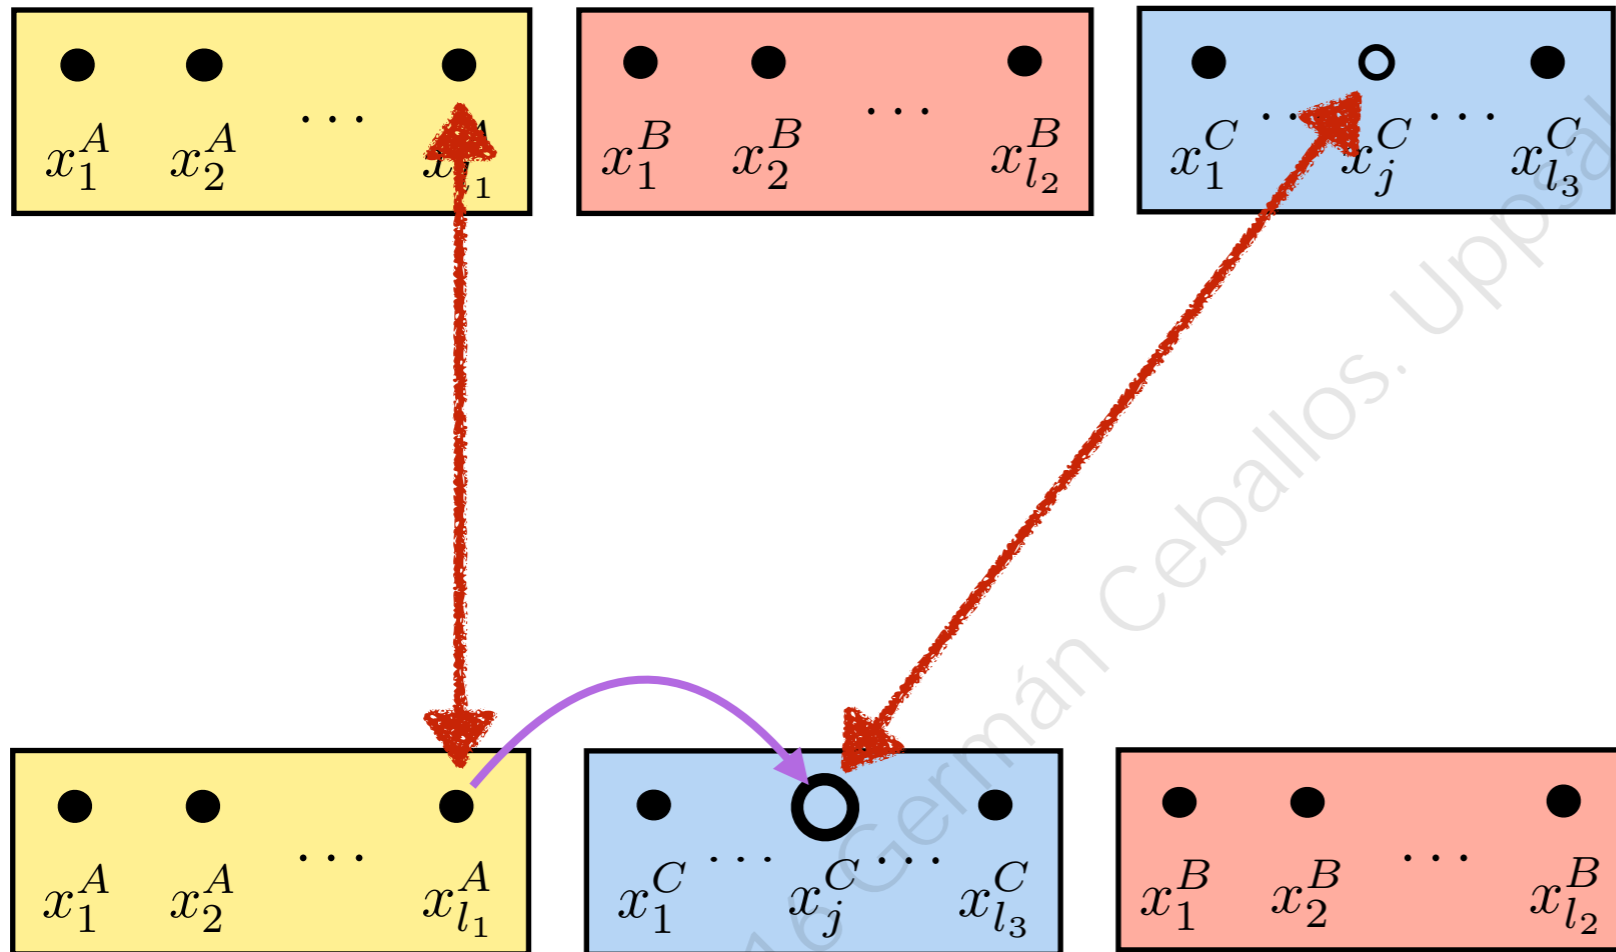
2. Shared Reuses

- Some are kept
- Some appear





Rebuild Reuses

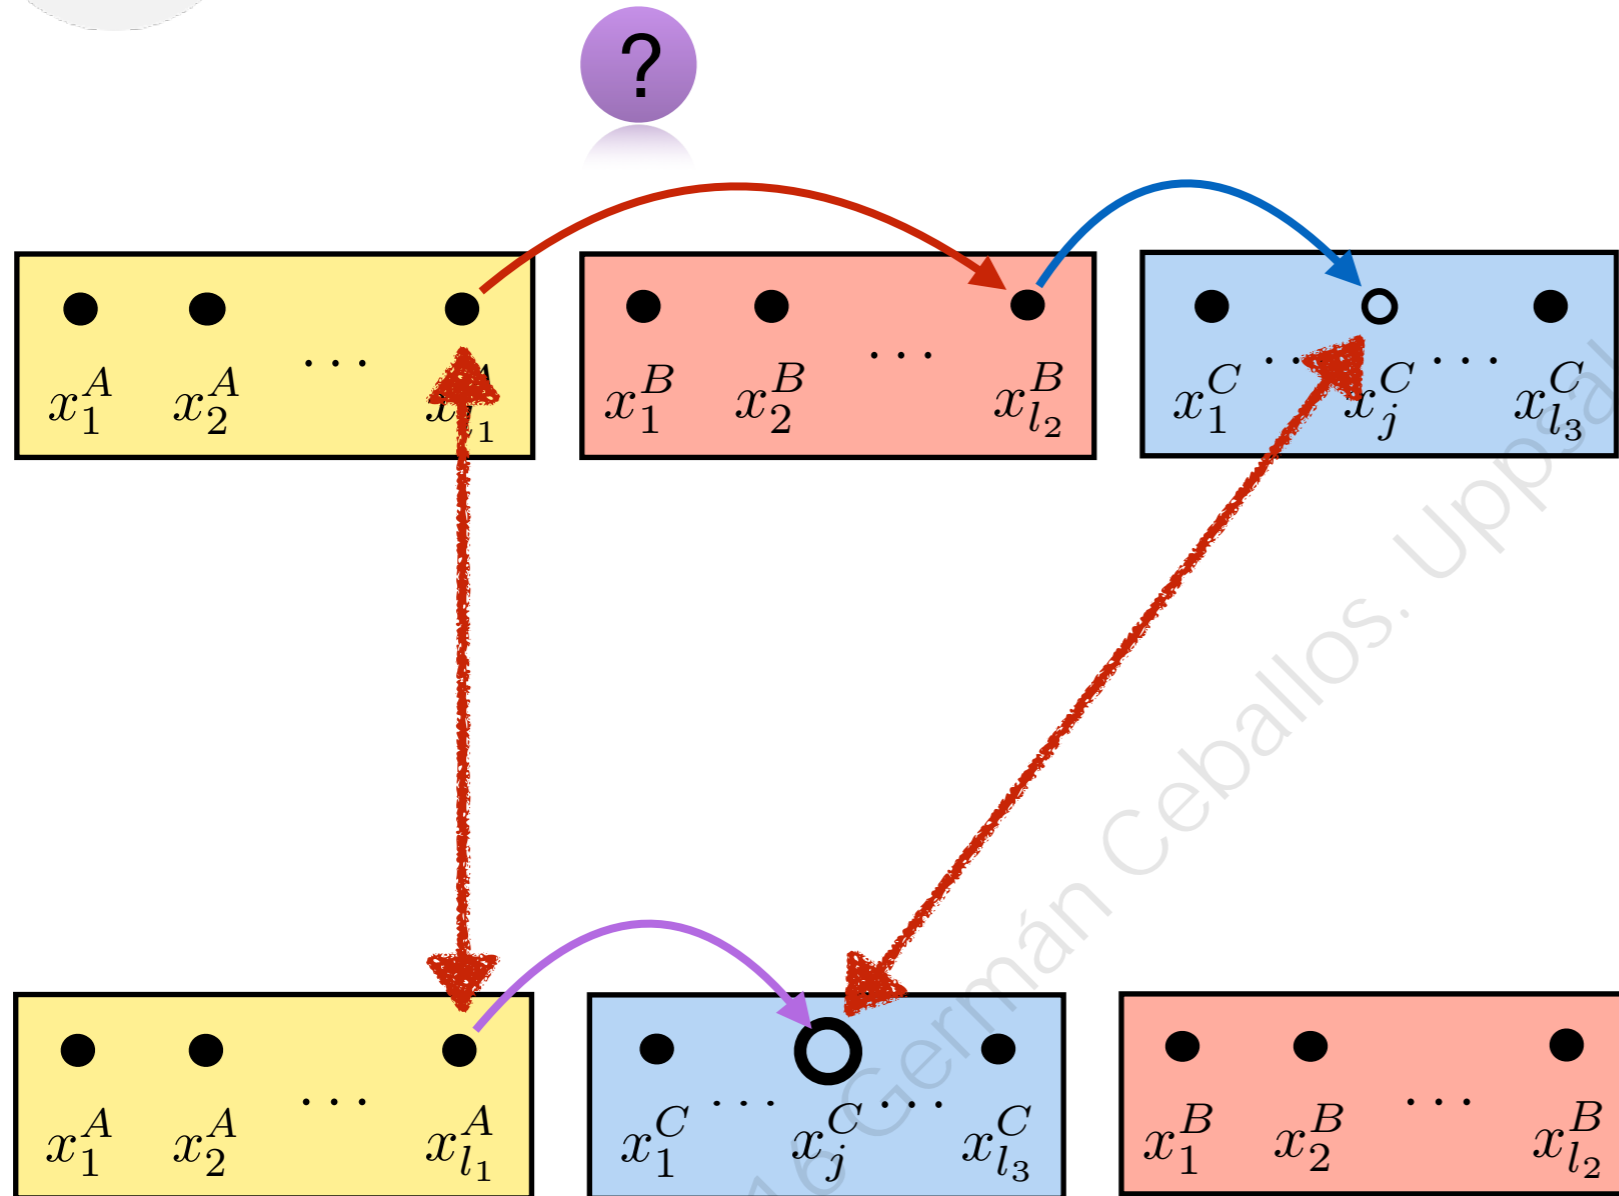


1. Private Reuses
 - All are kept
2. Shared Reuses
 - Some are kept
 - Some appear





Rebuild Reuses

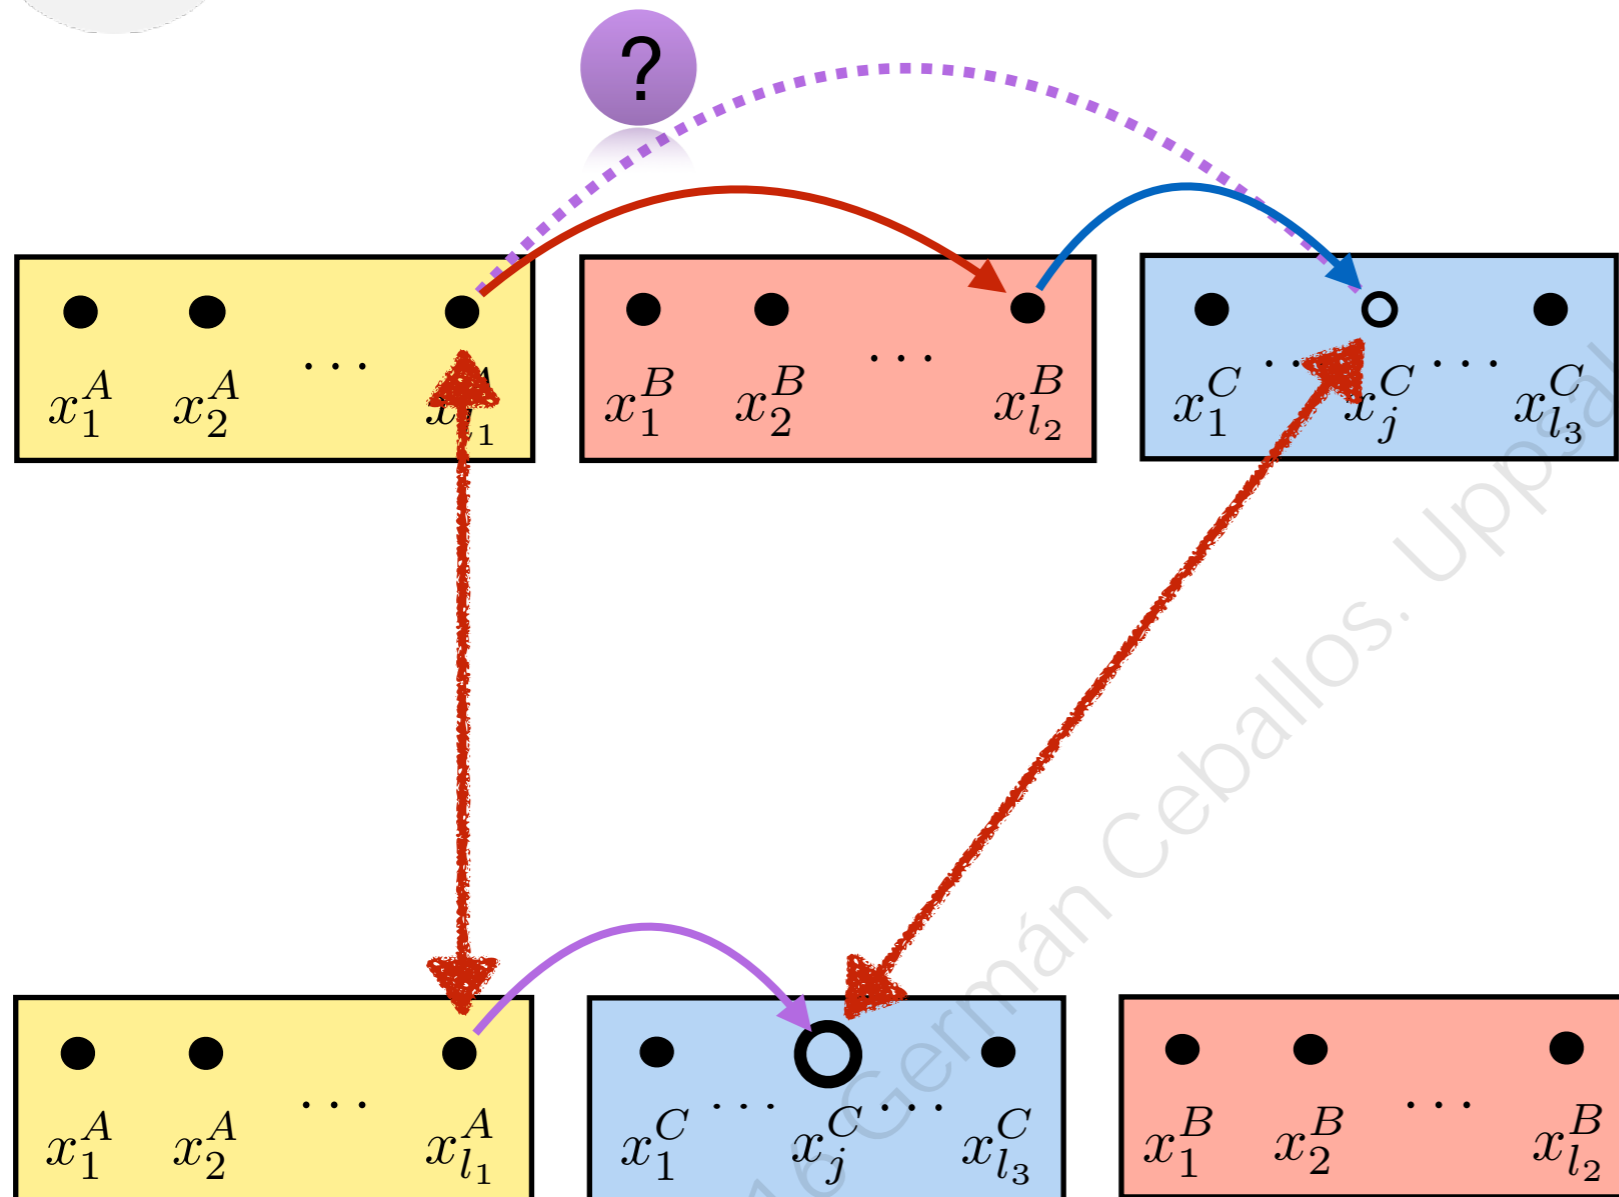


1. Private Reuses
 - All are kept
2. Shared Reuses
 - Some are kept
 - Some appear





Rebuild Reuses



1. Private Reuses

- All are kept

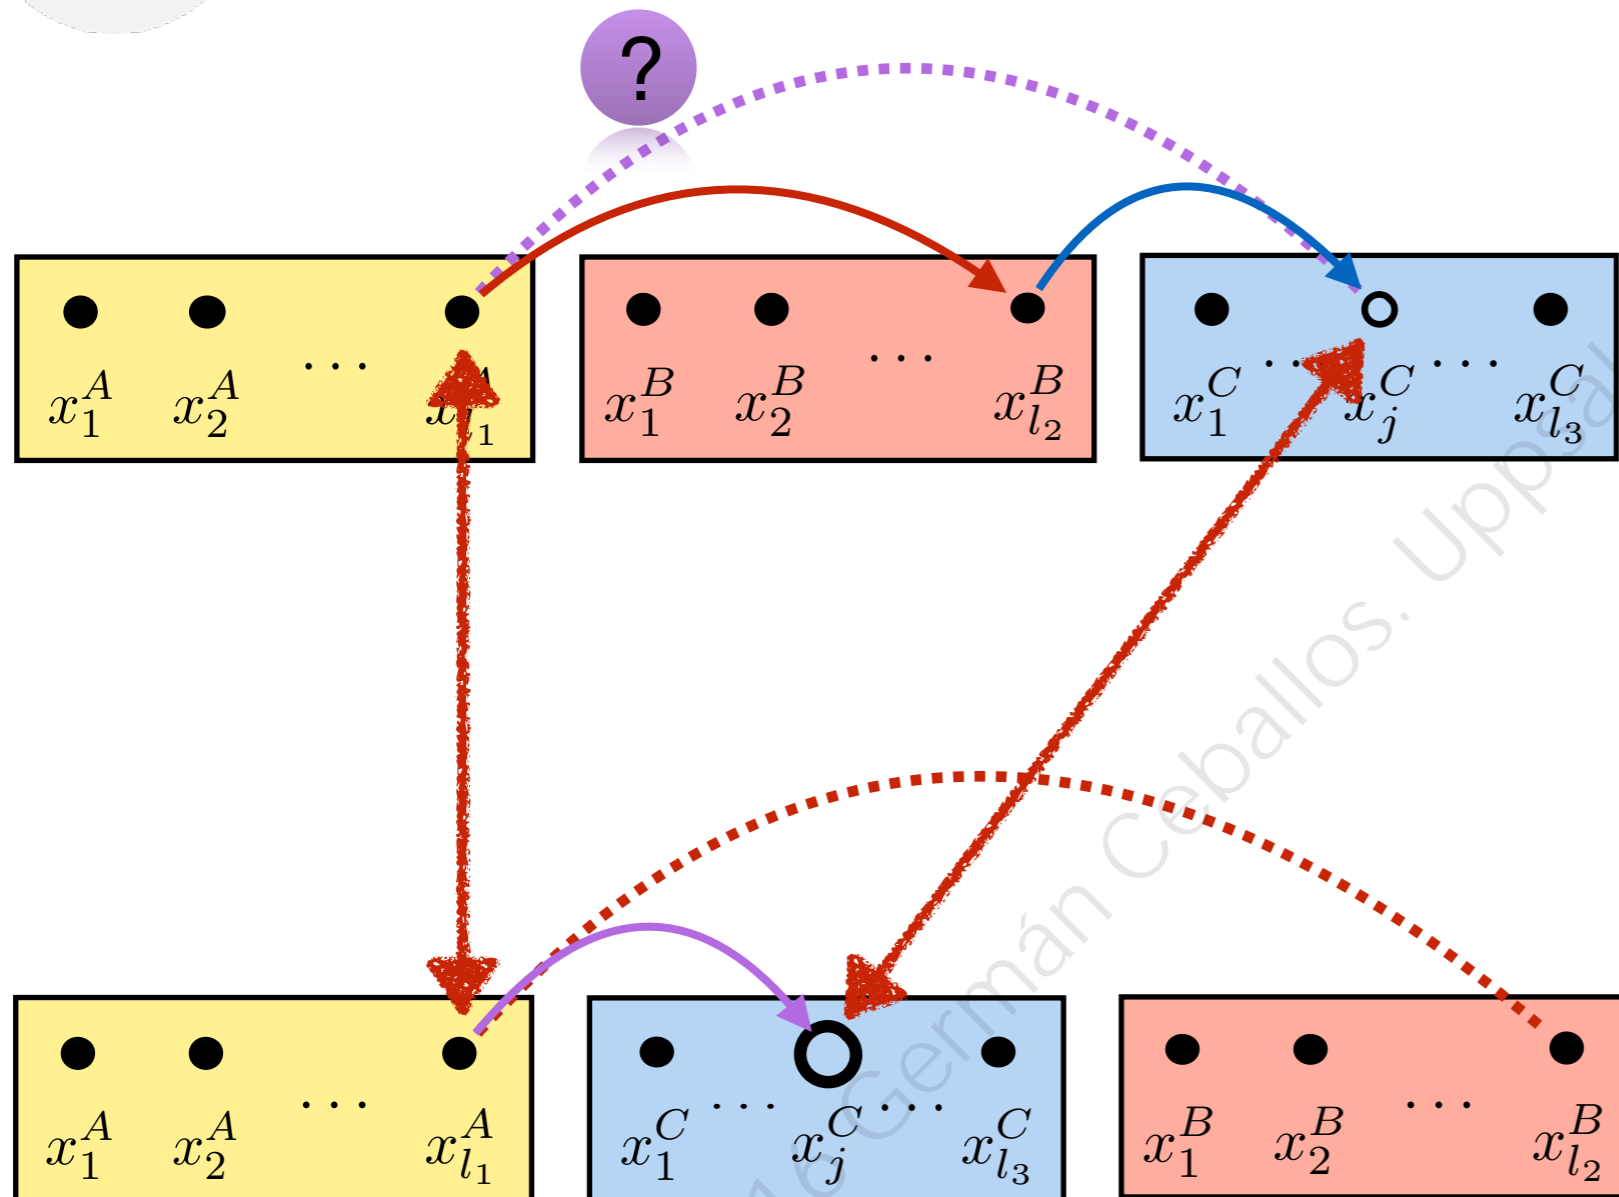
2. Shared Reuses

- Some are kept
- Some appear
- Some disappear





Rebuild Reuses



1. Private Reuses

- All are kept

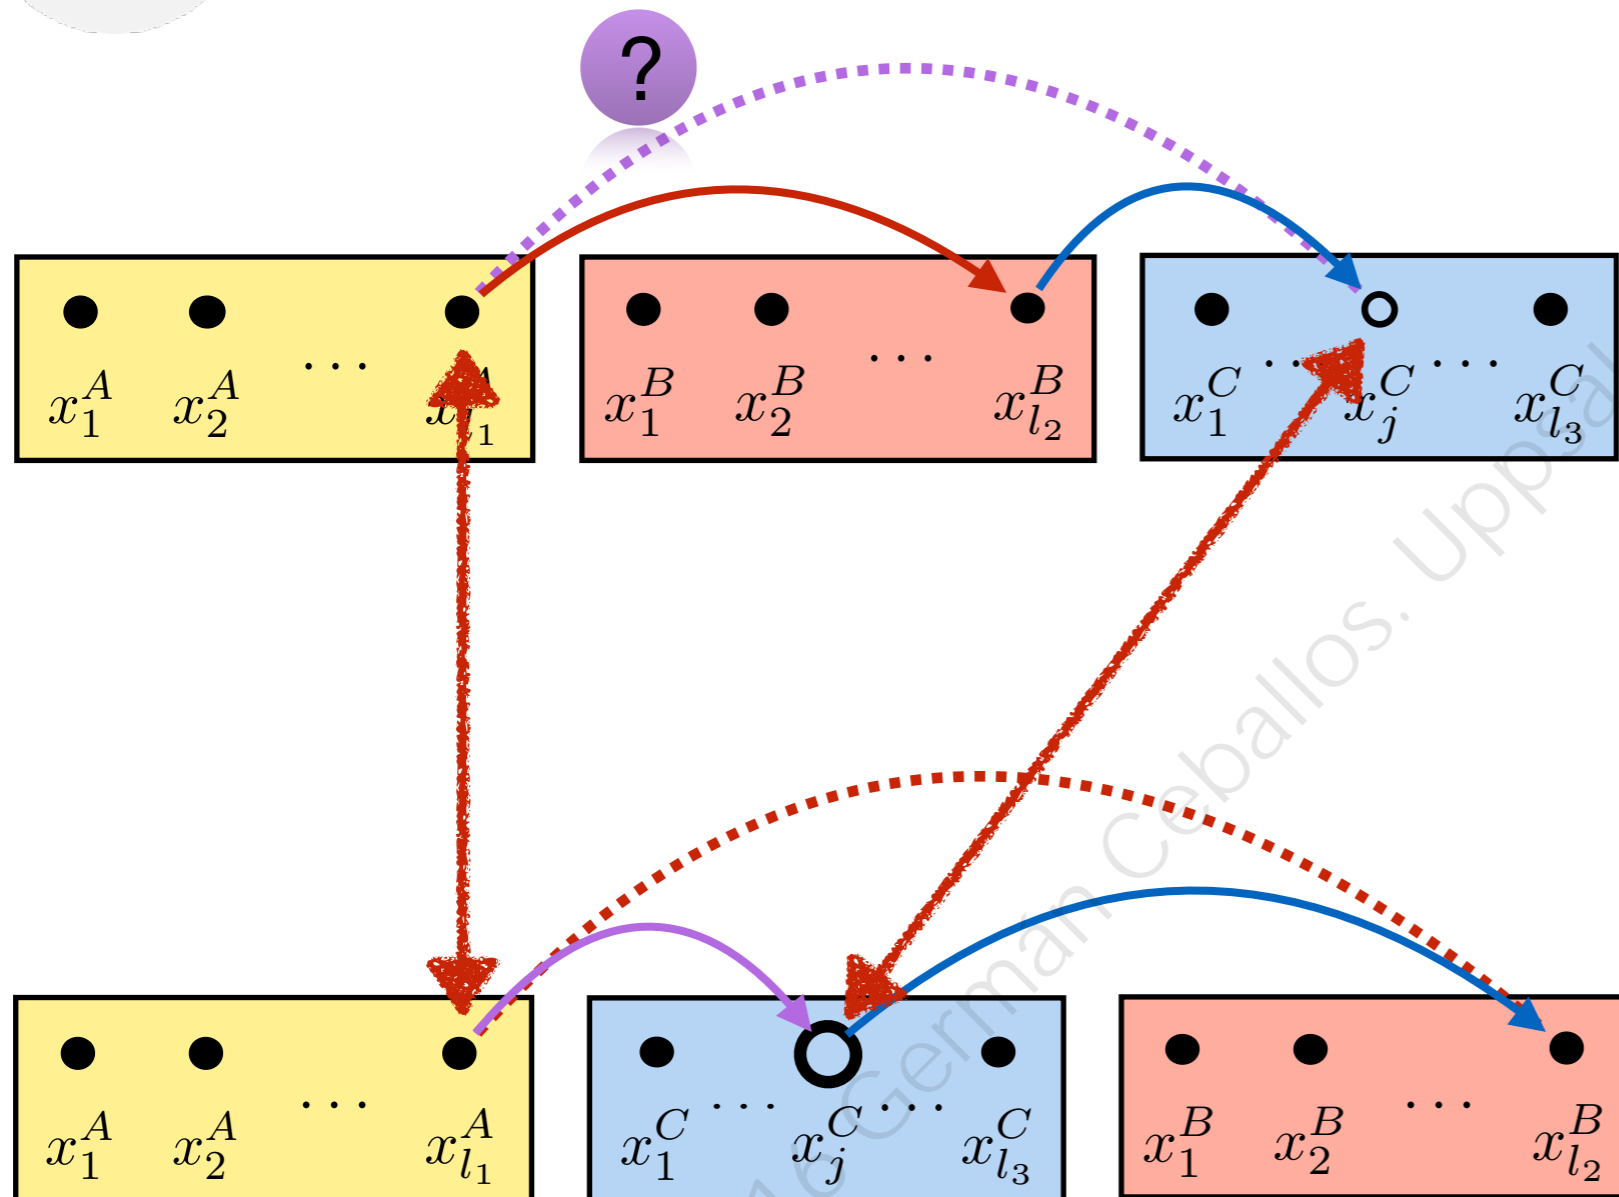
2. Shared Reuses

- Some are kept
- Some appear
- Some disappear





Rebuild Reuses



1. Private Reuses

- All are kept

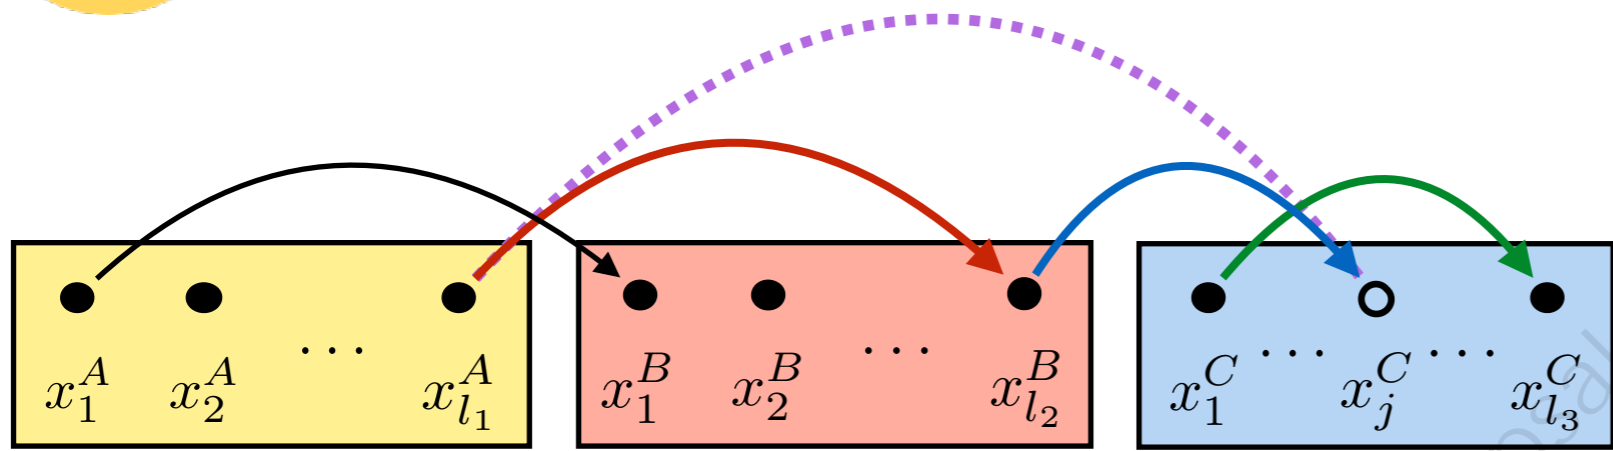
2. Shared Reuses

- Some are kept
- Some appear
- Some disappear



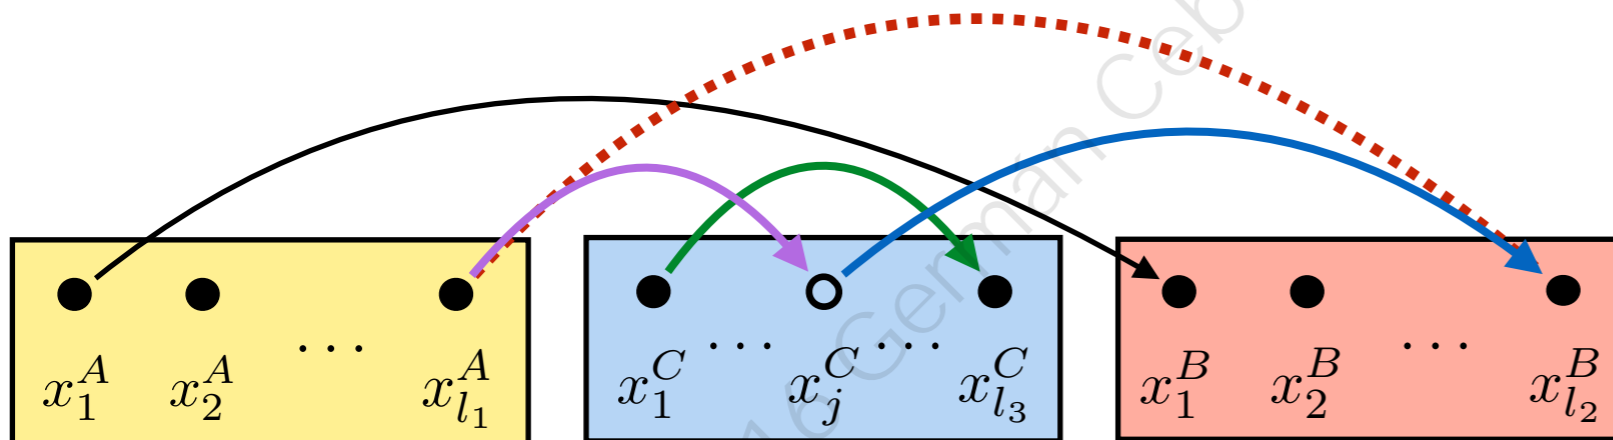


Recalculate Distances



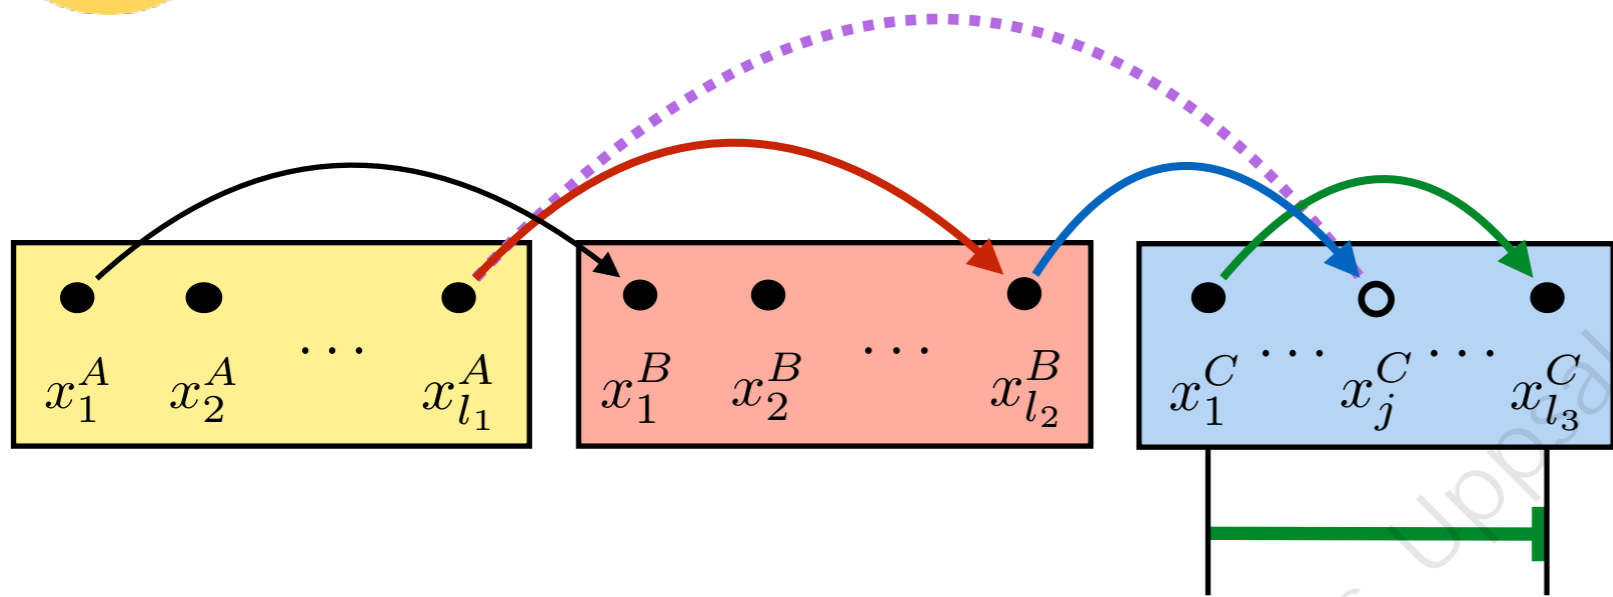
1. Private Reuses

- Same distance



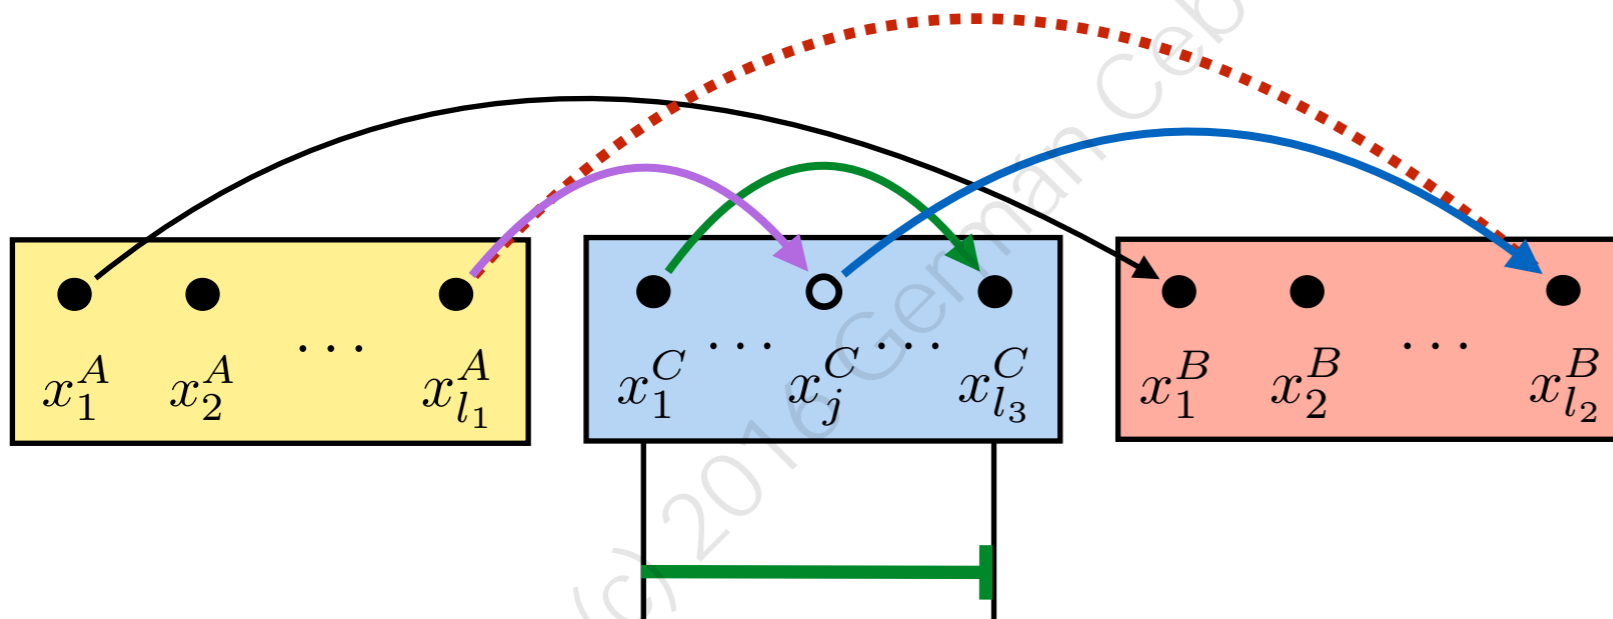


Recalculate Distances



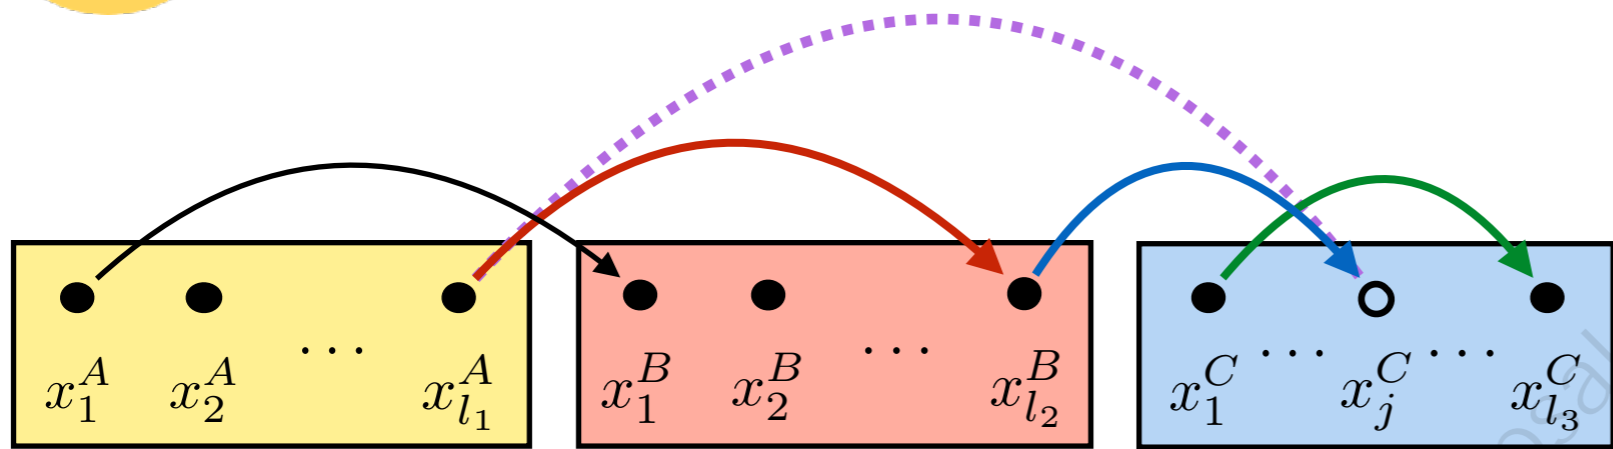
1. Private Reuses

- Same distance



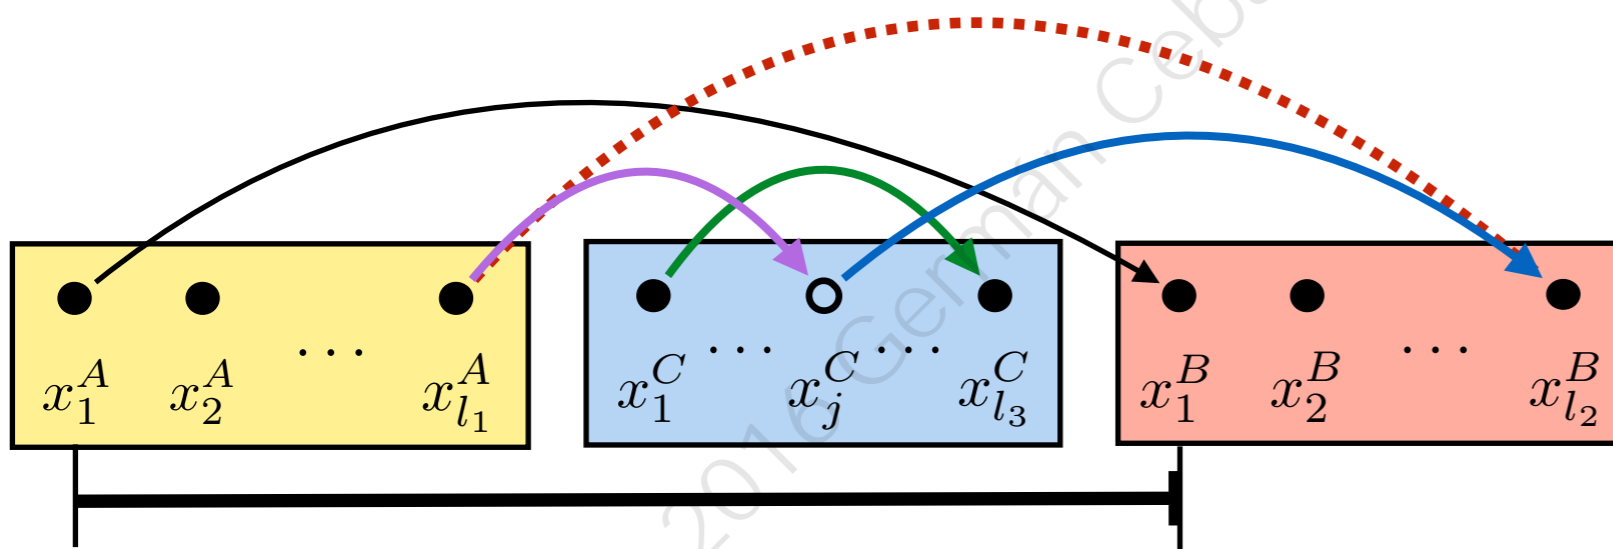


Recalculate Distances



1. Private Reuses

- Same distance



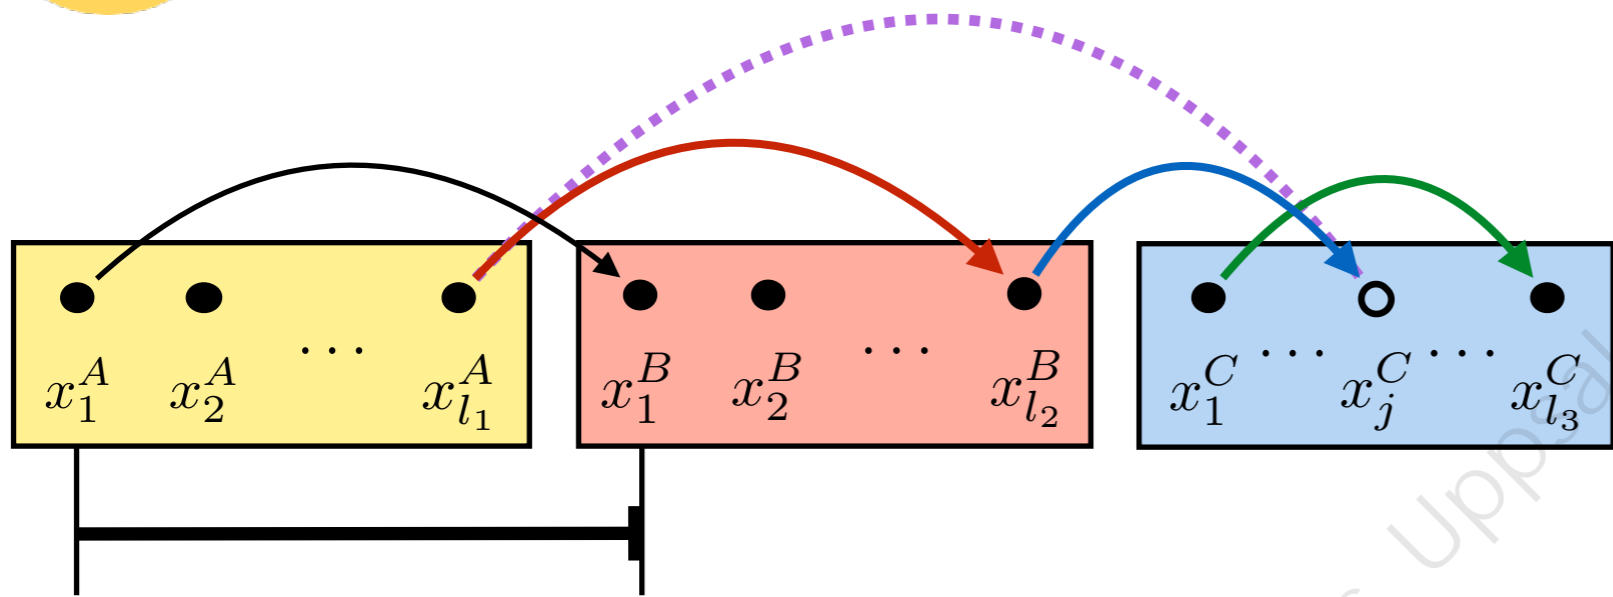
2. Shared Reuses

- Kept reuses



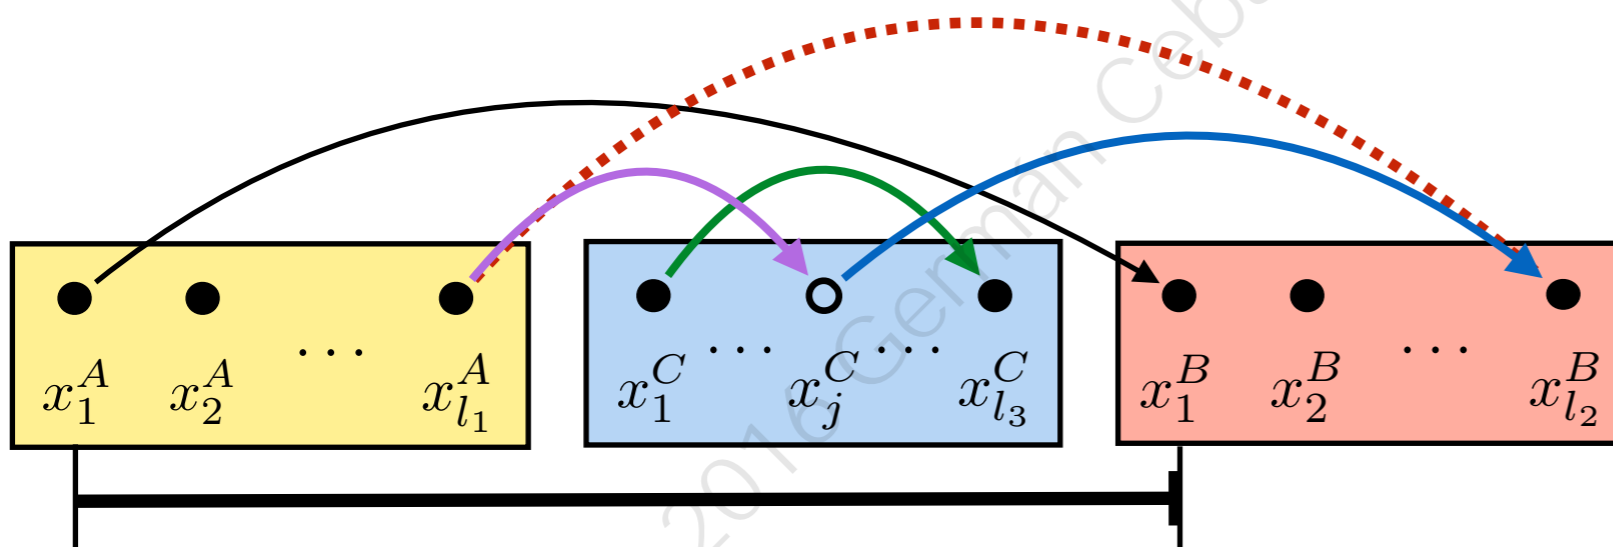


Recalculate Distances



1. Private Reuses

- Same distance



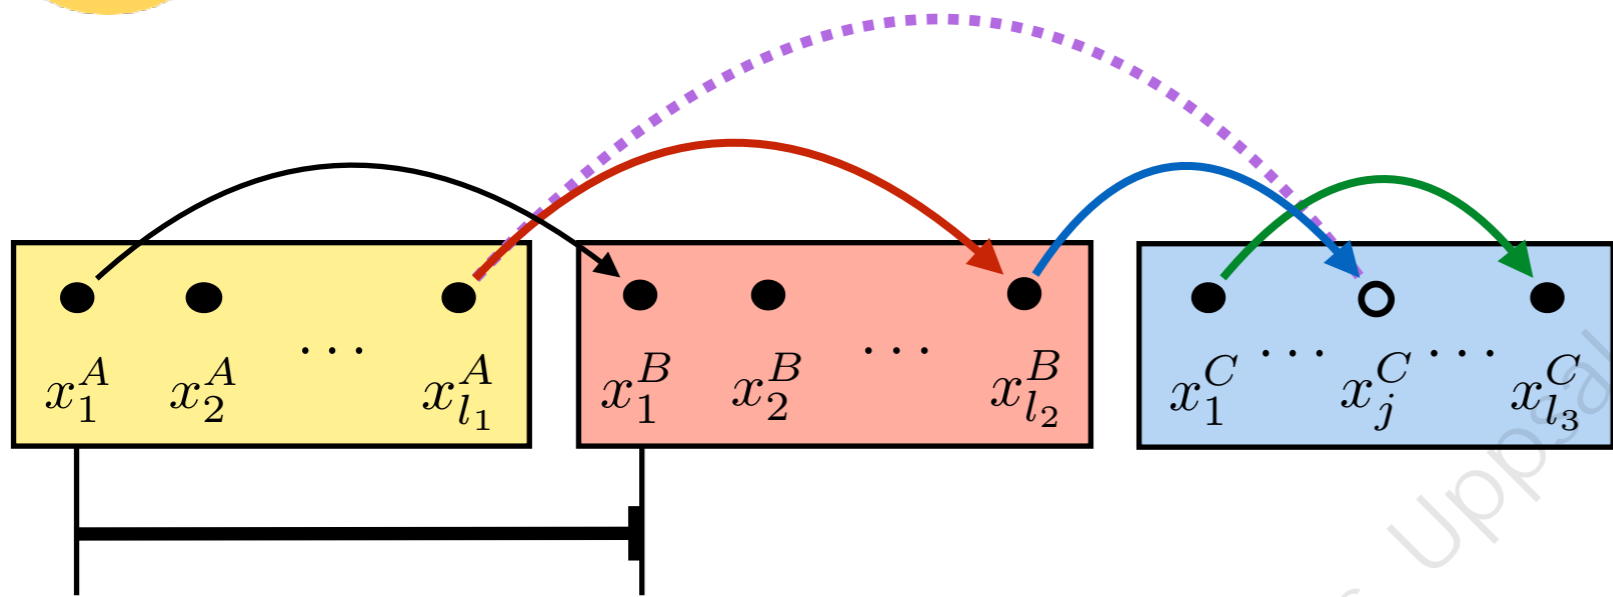
2. Shared Reuses

- Kept reuses



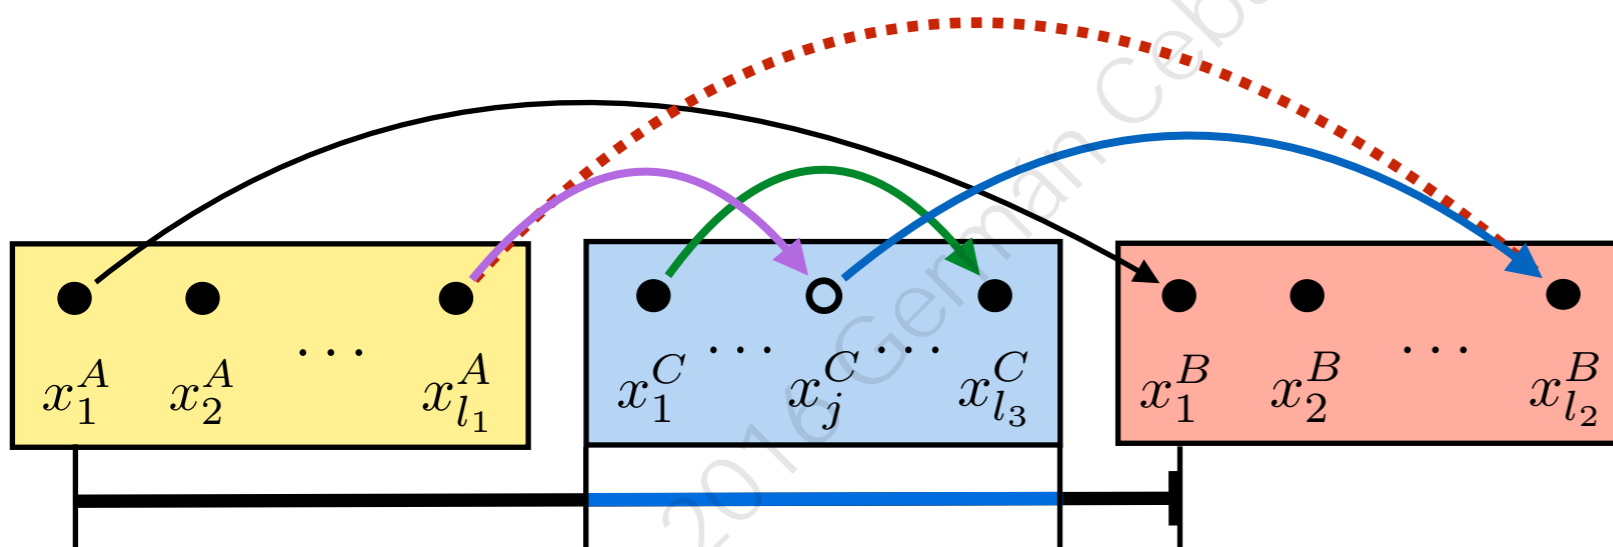


Recalculate Distances



1. Private Reuses

- Same distance



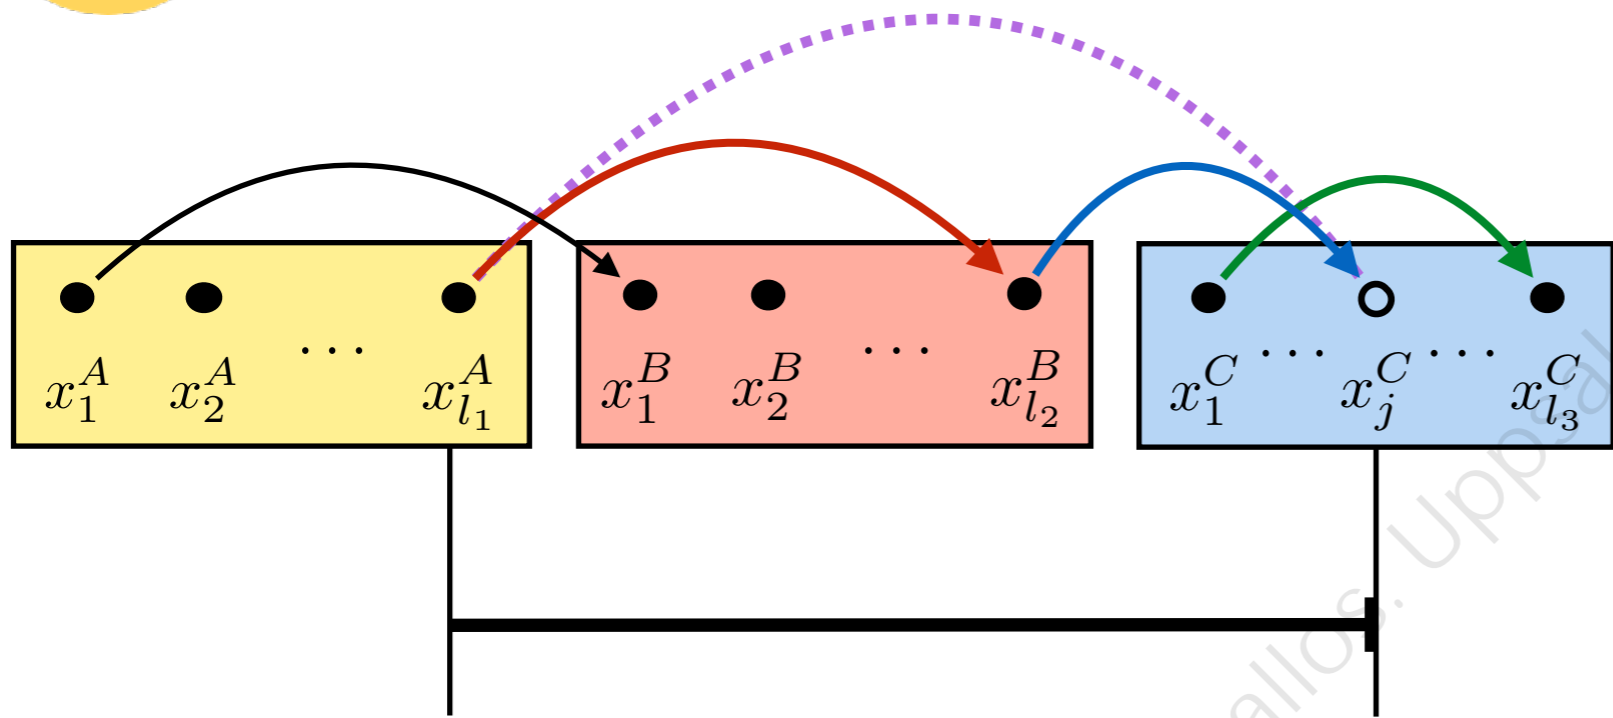
2. Shared Reuses

- Kept reuses



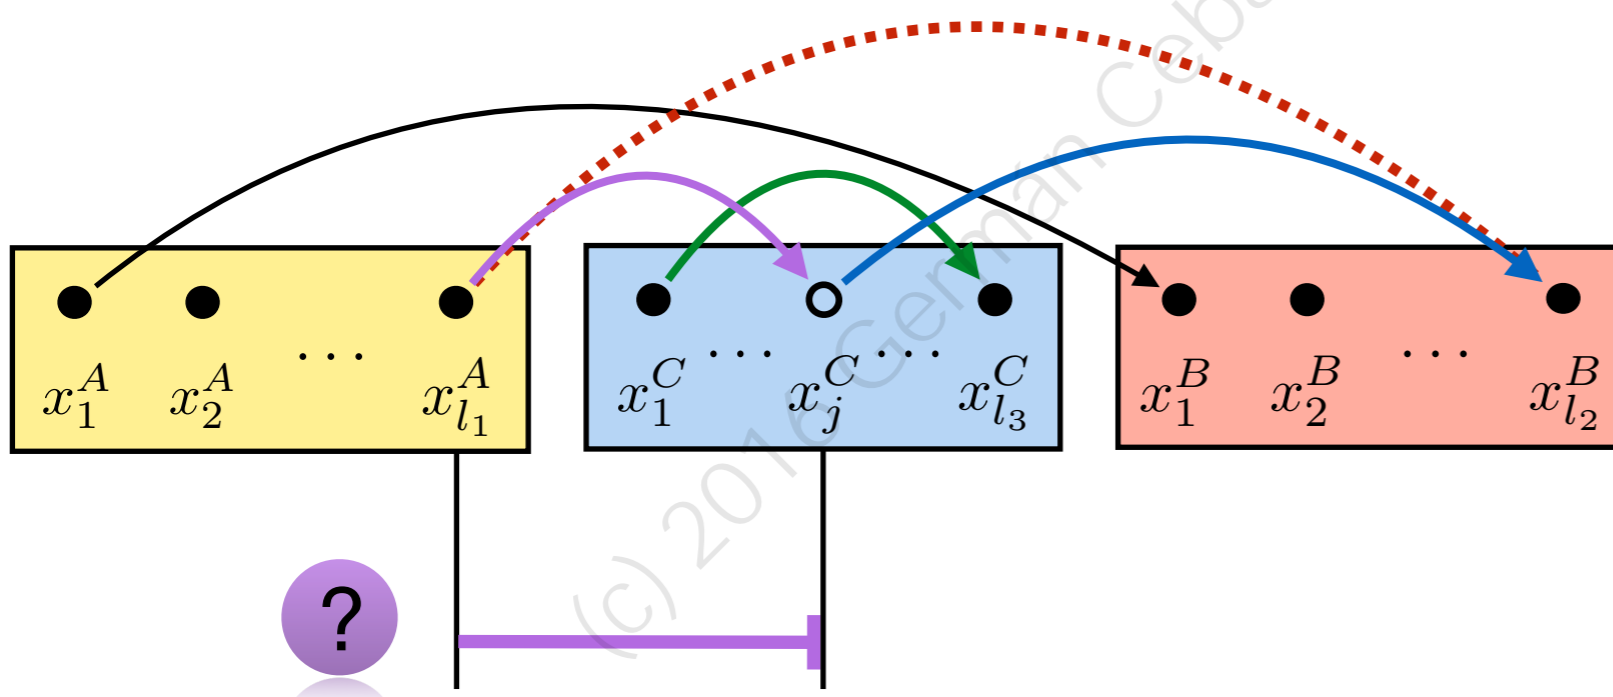


Recalculate Distances



1. Private Reuses

- Same distance

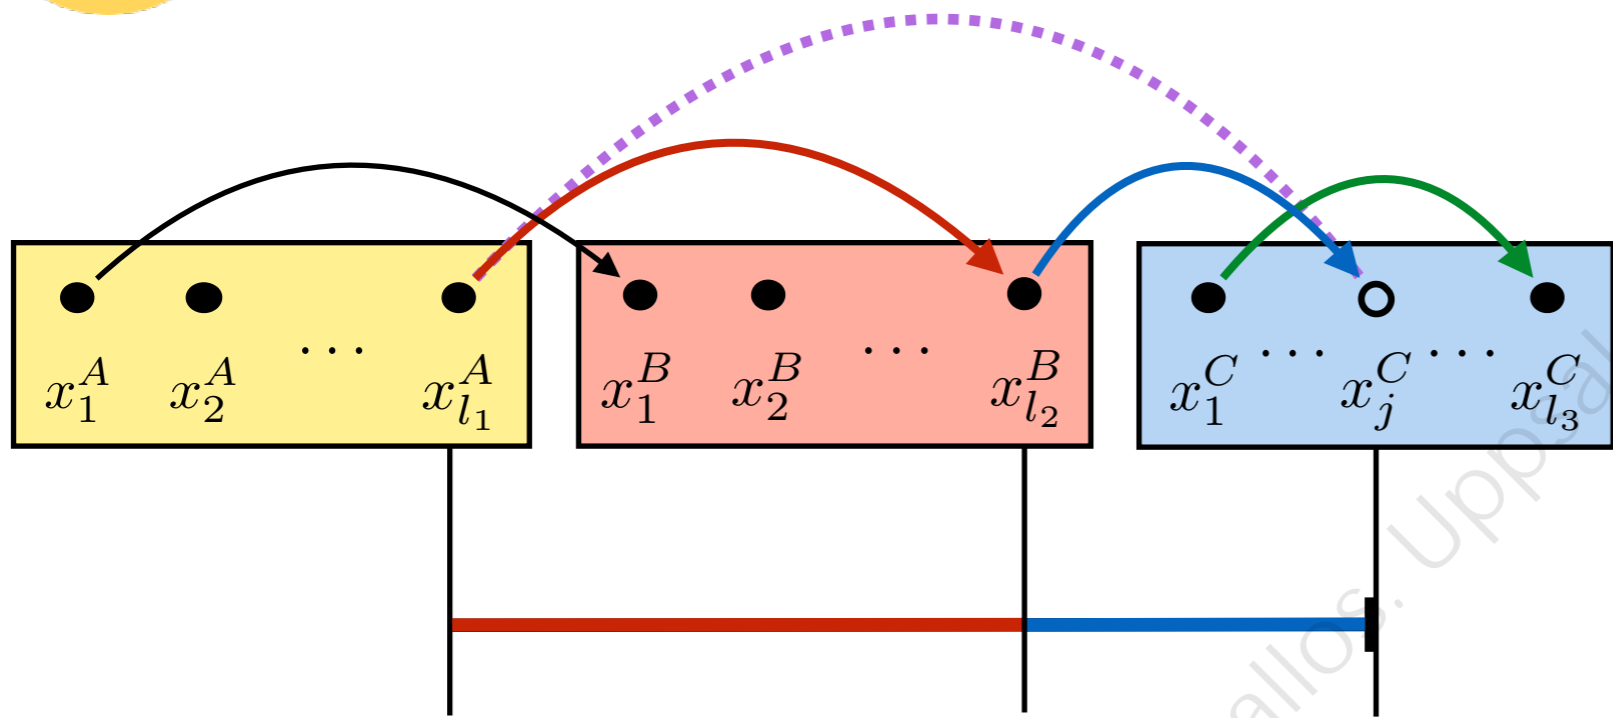


2. Shared Reuses

- Kept reuses
- New reuses

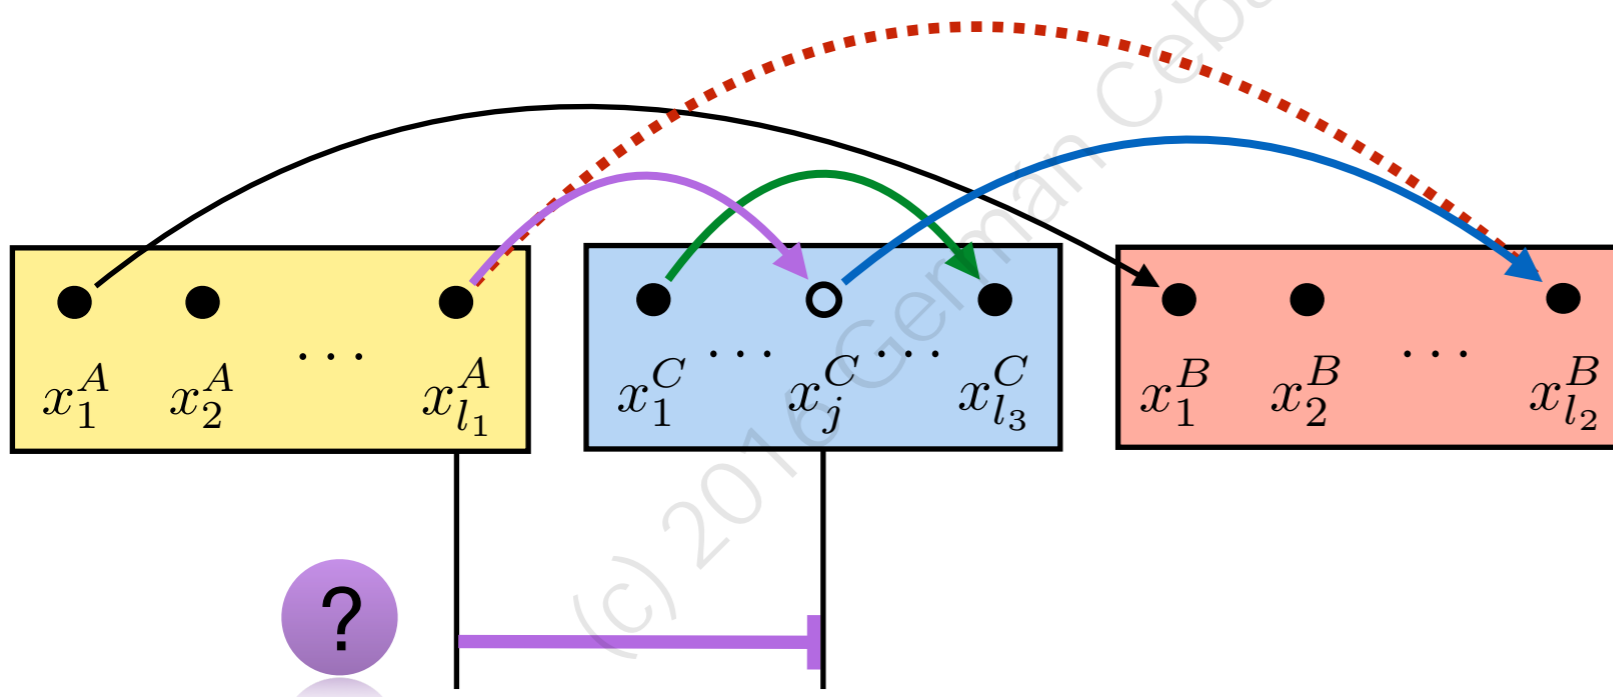


Recalculate Distances



1. Private Reuses

- Same distance

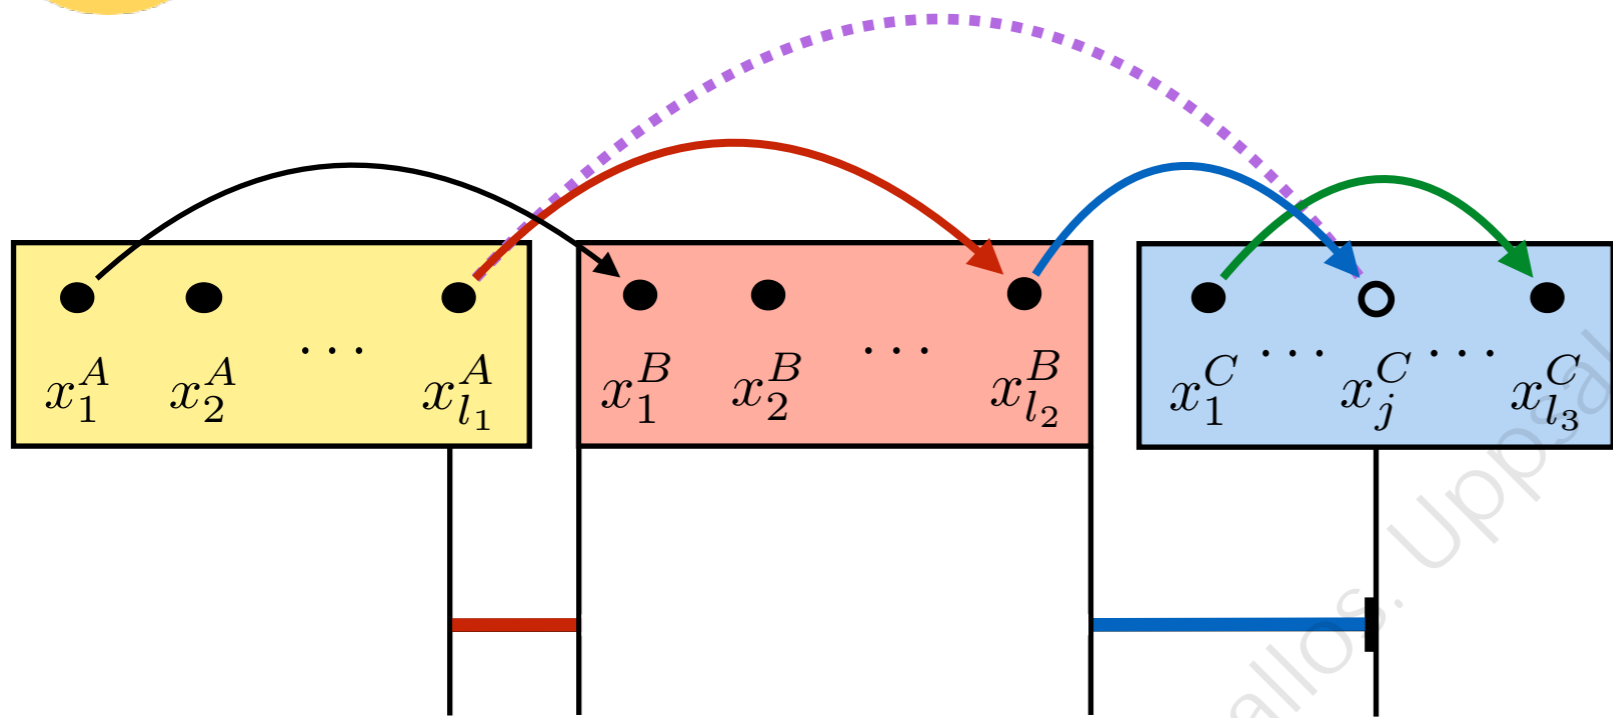


2. Shared Reuses

- Kept reuses
- New reuses

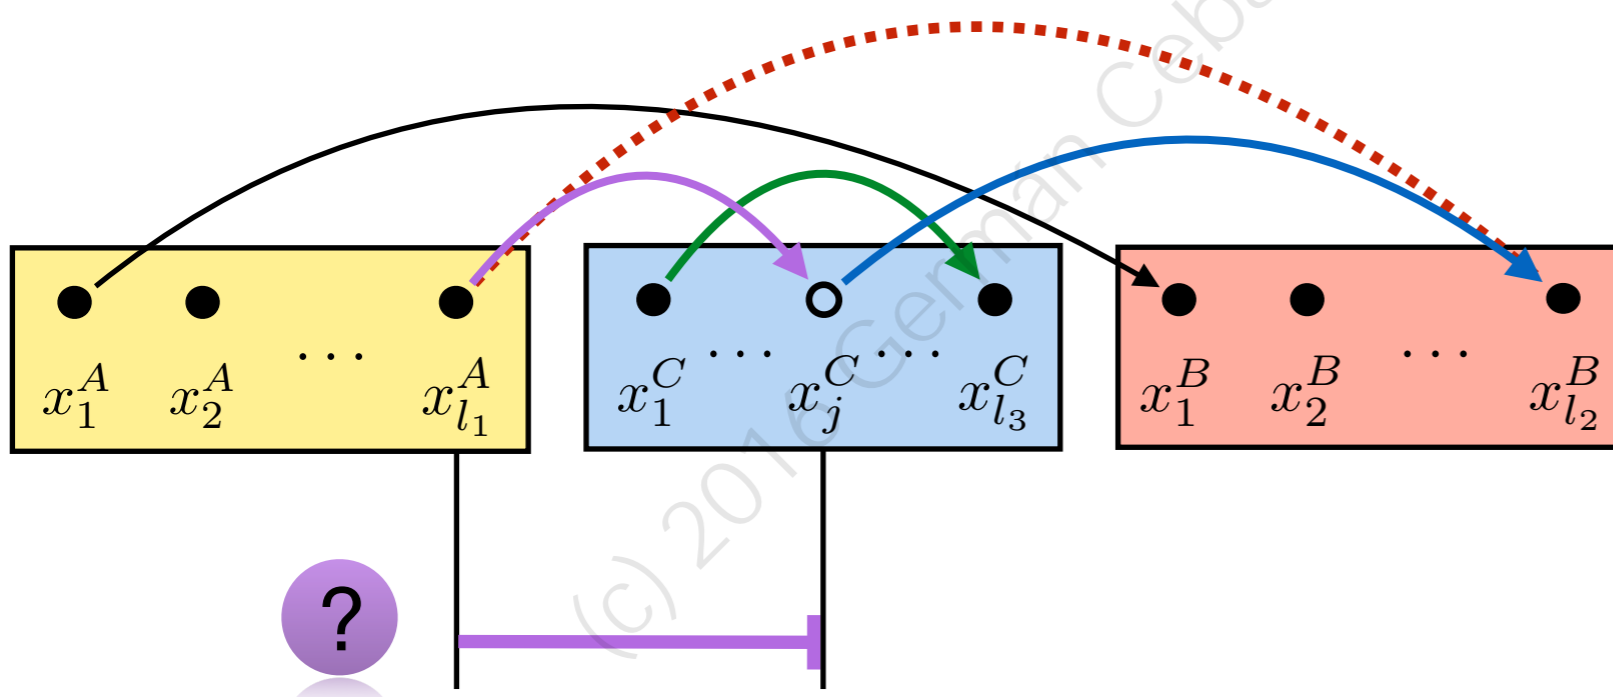


Recalculate Distances



1. Private Reuses

- Same distance



2. Shared Reuses

- Kept reuses
- New reuses

(c) 2016 German Ceballos. Uppsala University

Evaluation & **Results**



Results

- Barcelona OpenMP Tasks Suite (BOTS)*
- Quad-core Intel Sandy Bridge.
 - 8KB private L1
 - 256KB private L2.
 - 6MB shared L3.
 - 32GB RAM
- **Evaluation**
 - **Accuracy**: predicted vs measured miss ratios for **same schedule**
 - **Flexibility**: predicted miss ratio from **different schedules**
 - **Robustness**: sensitivity to different data set of same size (10% difference)

*Alejandro Duran, Xavier Teruel, Roger Ferrer, Xavier Martorell, and Eduard Ayguade. 2009.

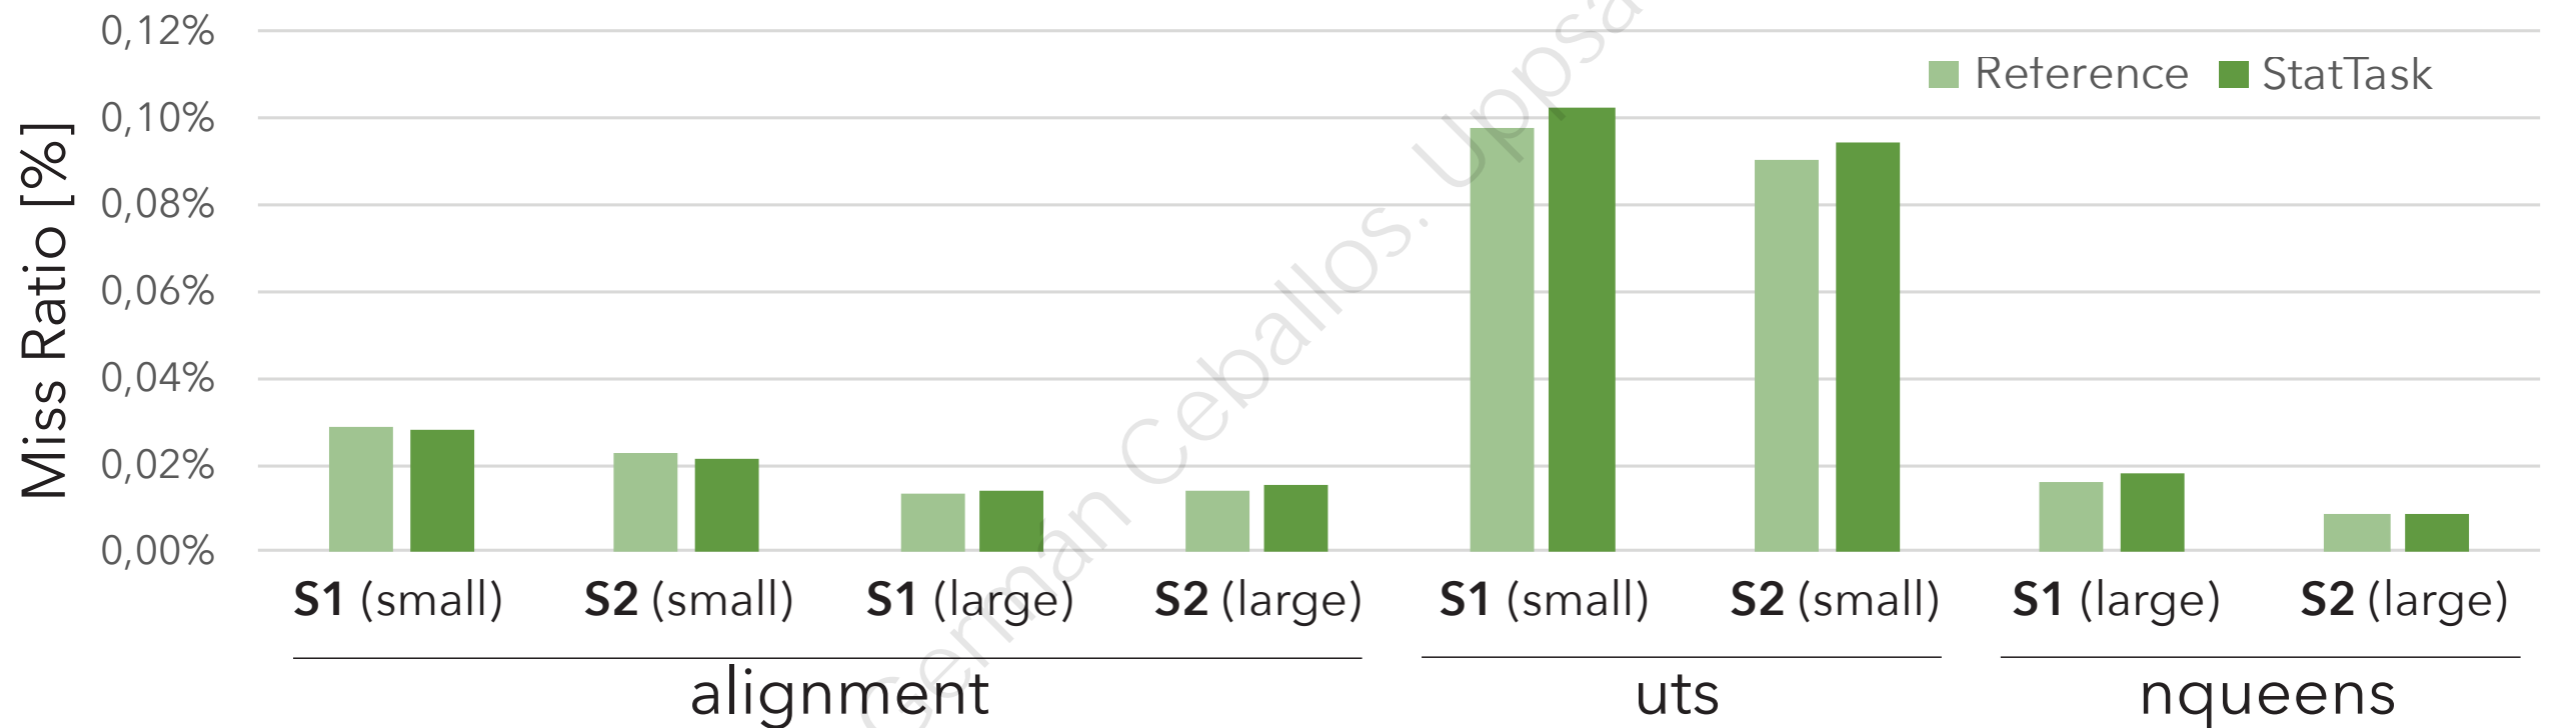
Barcelona OpenMP Tasks Suite: A Set of Benchmarks Targeting the Exploitation of Task Parallelism in OpenMP.

In *Proceedings of the 2009 International Conference on Parallel Processing (ICPP '09)*. Washington, DC, USA.



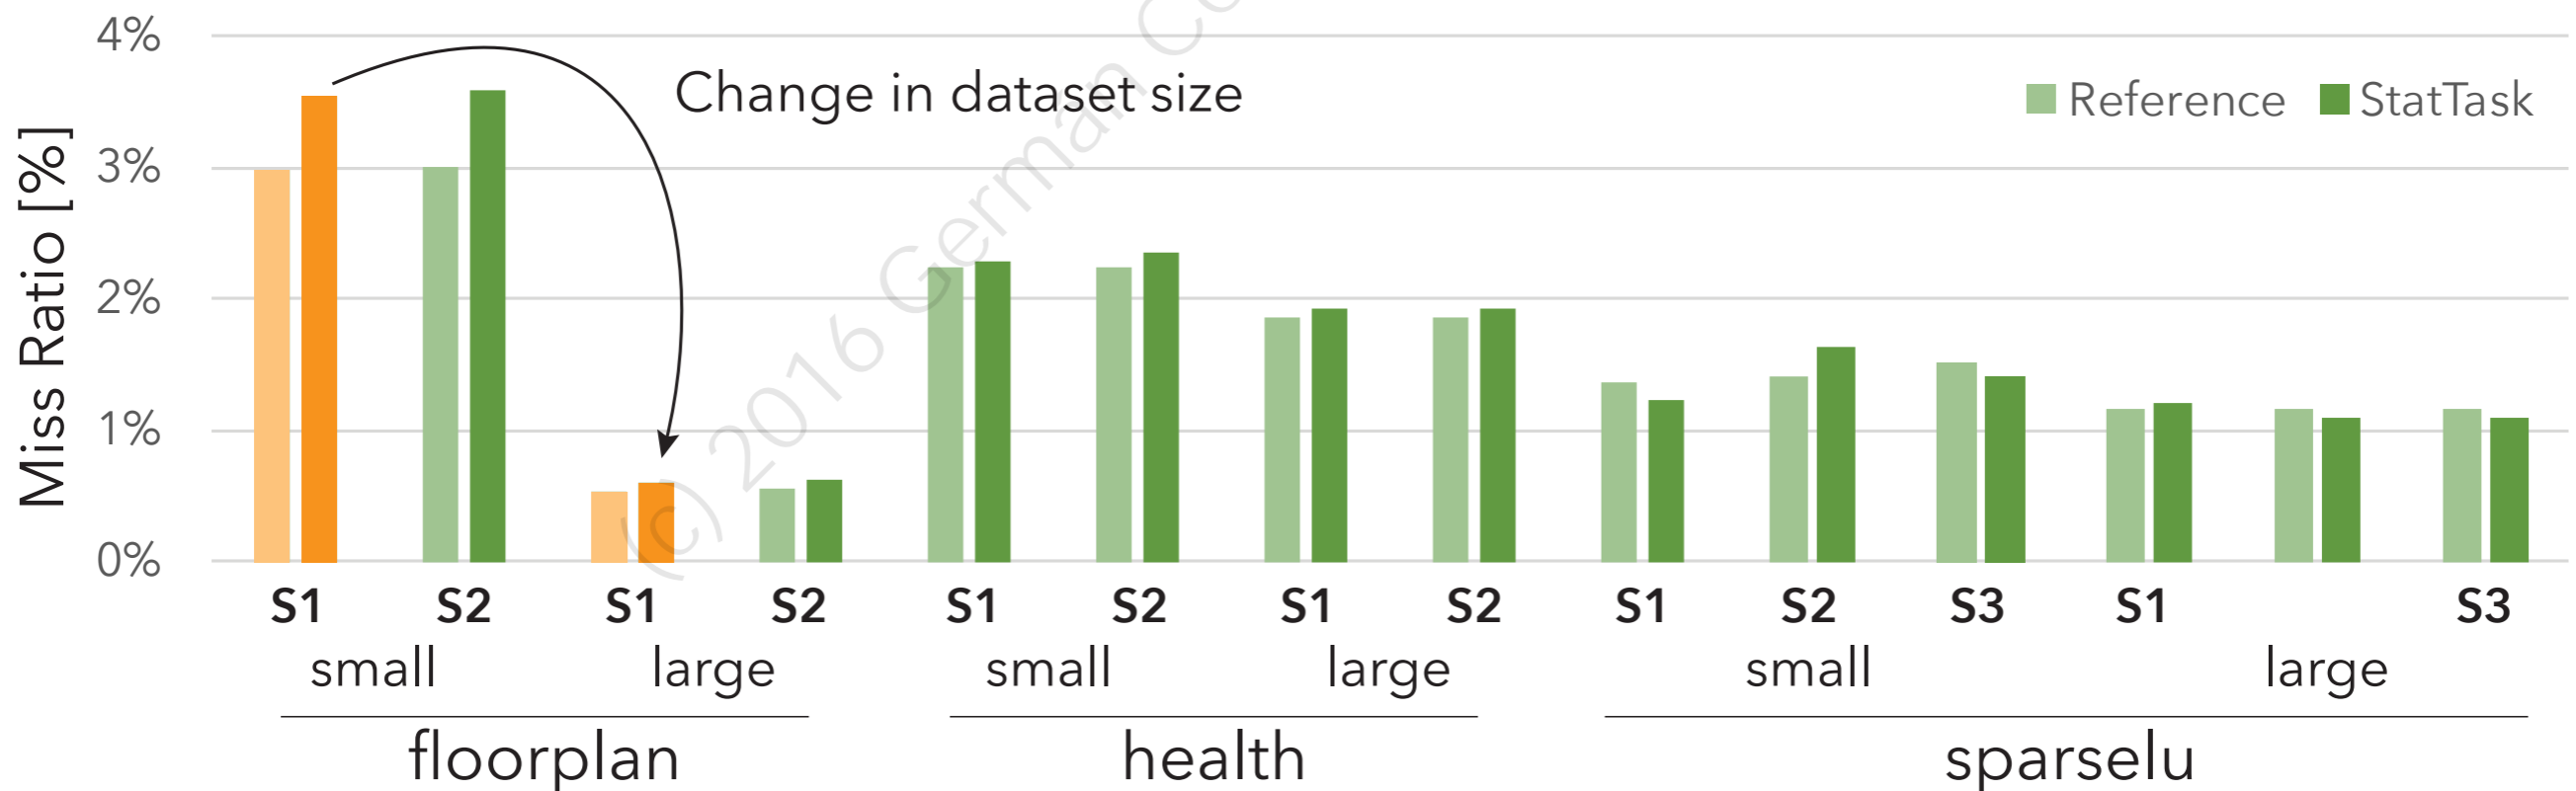
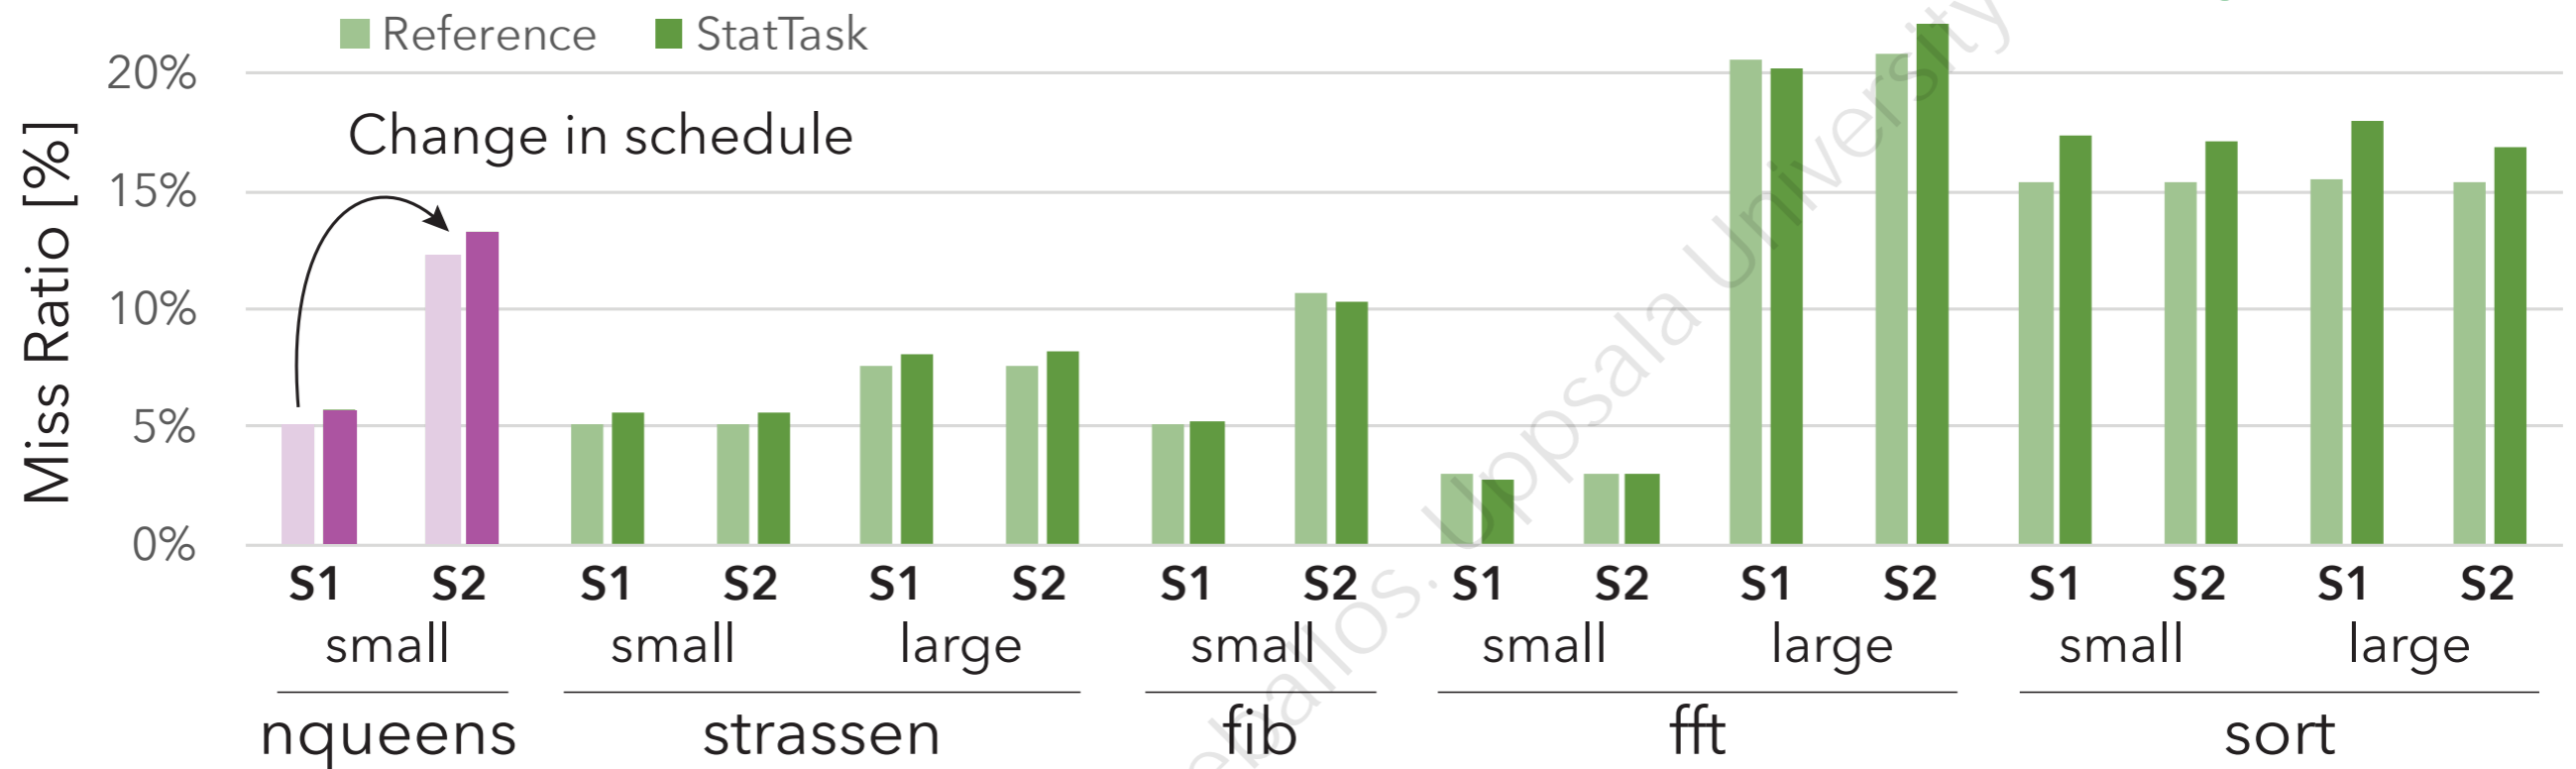
StatTask's Accuracy

Low Miss Ratios

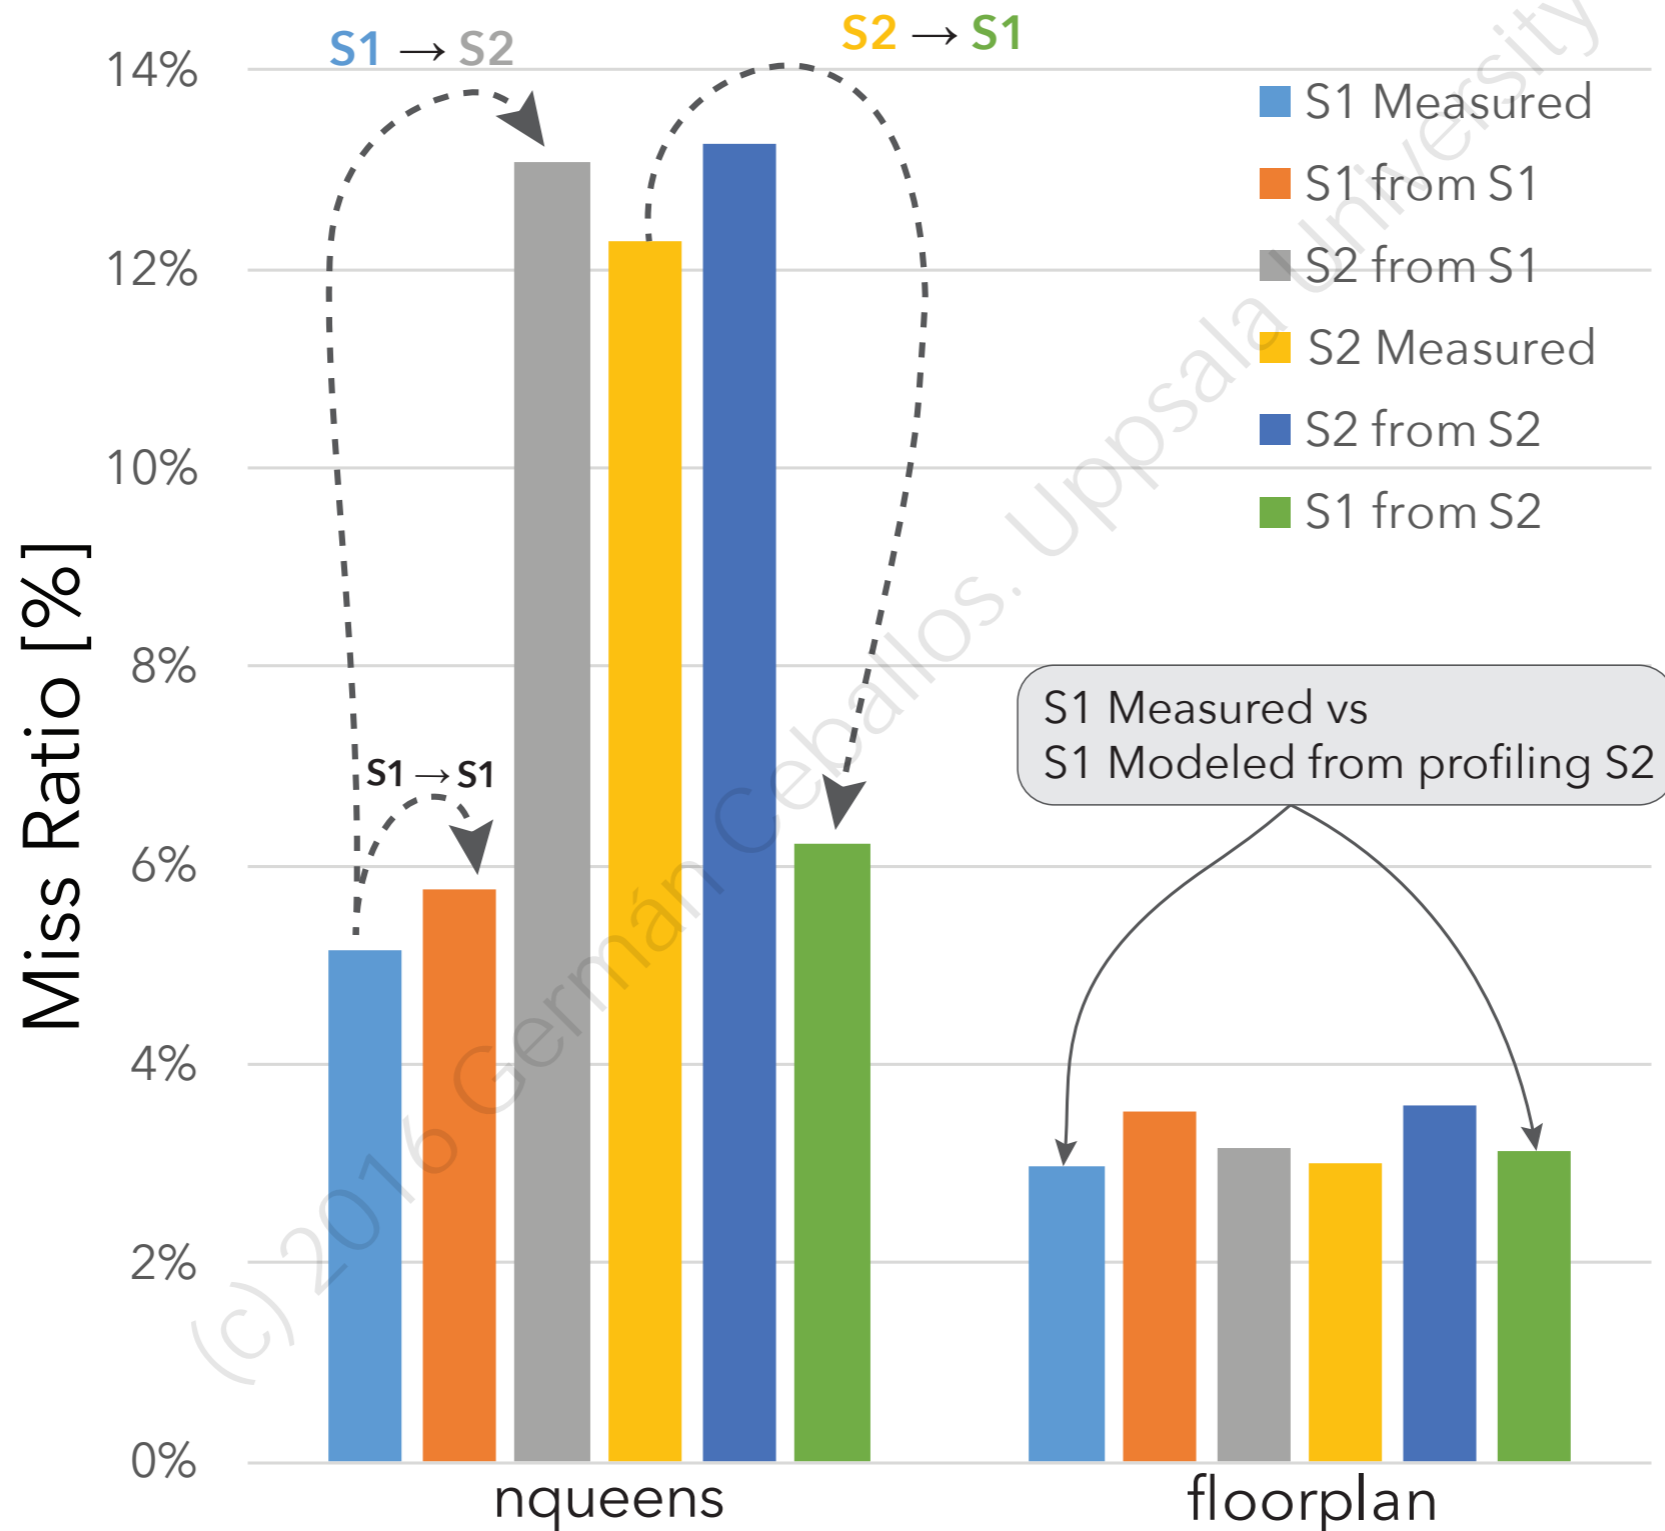


Reference: HW Performance Counters for Last Level Cache (L3)

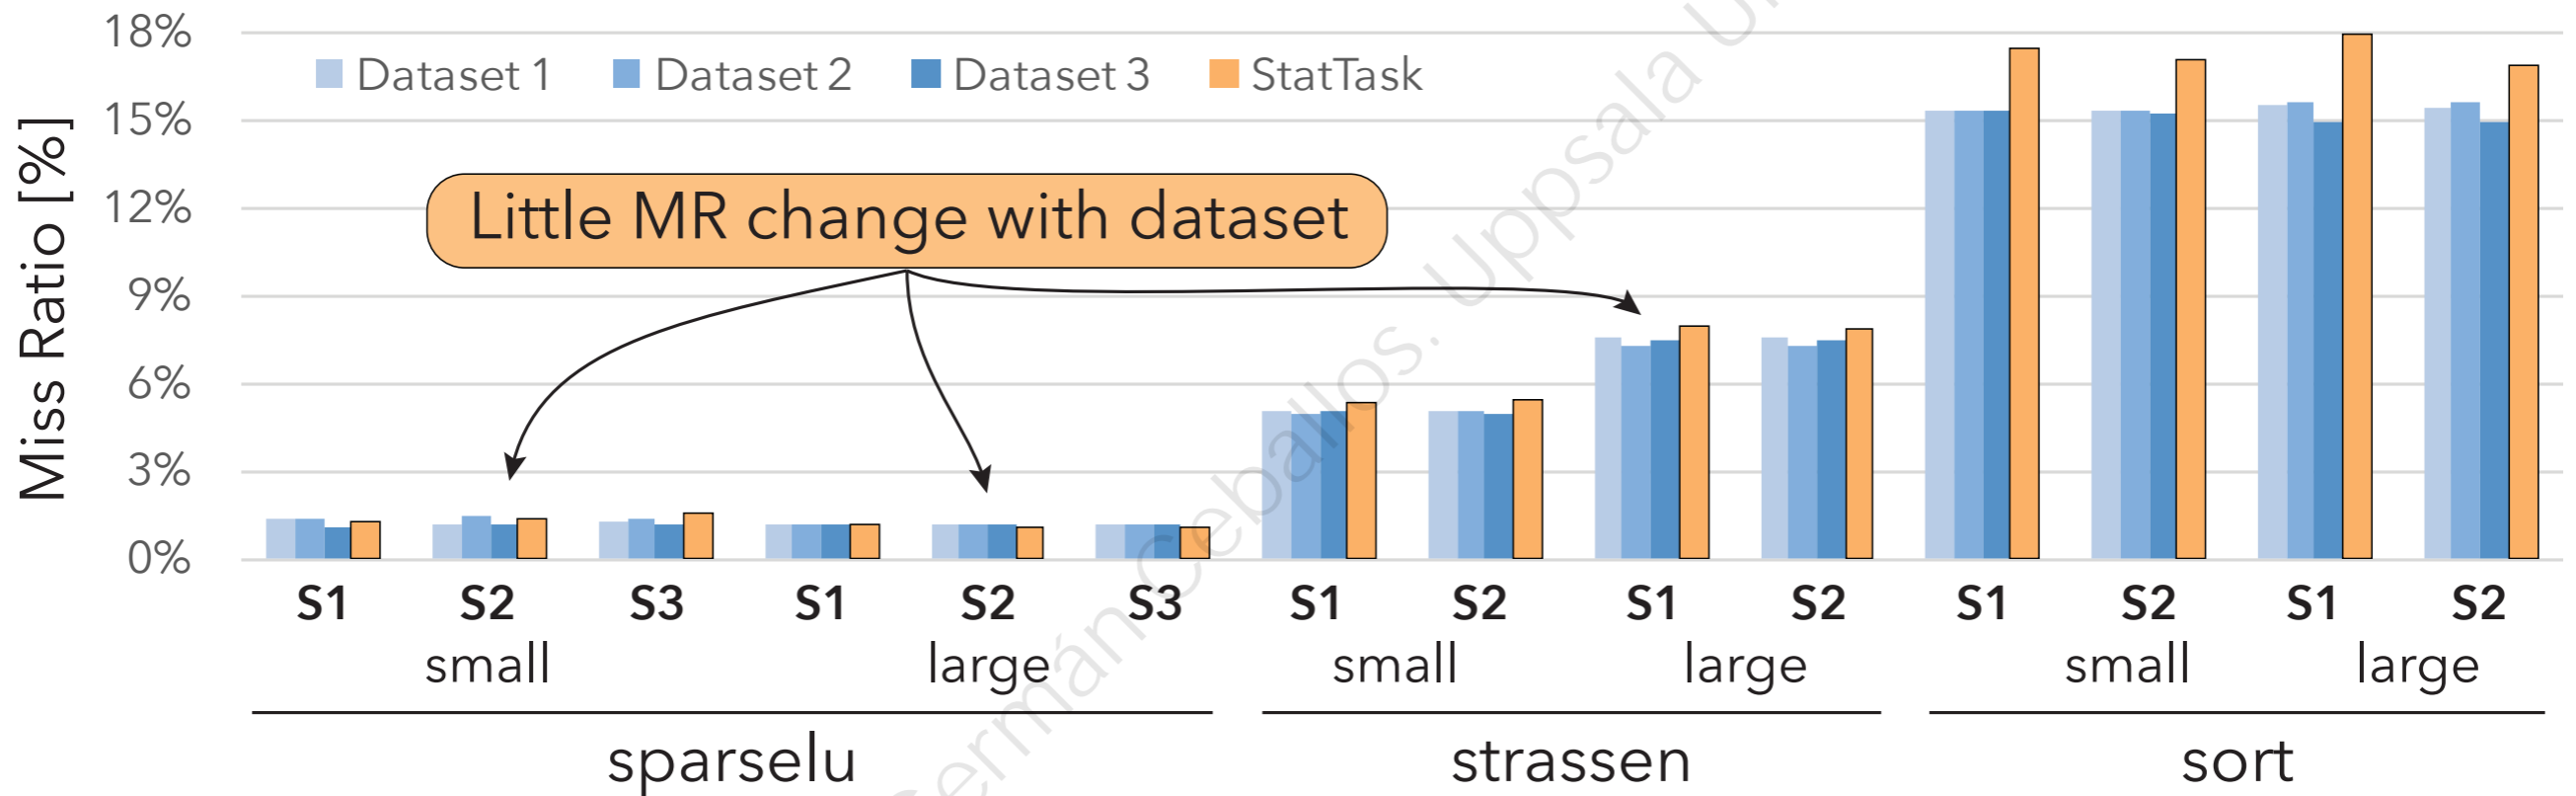
StatTask's Accuracy



StatTask's Flexibility



StatTask's Robustness



Conclusion

- Different task **schedules** have different **cache behaviour**
 - ≠ schedules** → **≠ reuses** → **≠ performance**
- **Statistical Cache Models** can be used to predict different schedules
 - Need to be leveraged to **recompute reuses**
- **StatTask**
 - Formalizes the task-based execution model
 - **Accurate**: compared to measured results
 - **Flexible**: predicts from a **single** profile
 - **Robust**: predicts similar behaviour for inputs of roughly same size



Formalizing Data Locality in Task-Parallel Applications

Germán Ceballos

Erik Hagersten

David Black-Schaffer

Thank You!

ICA3PP '16

Uppsala University

(c) 2016 Germán Ceballos, Uppsala University

