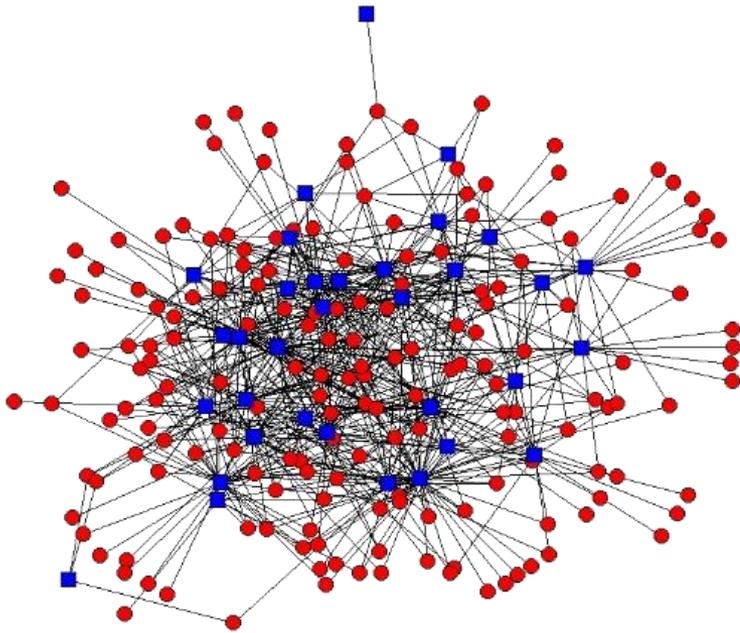


# Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages

MACIEJ BESTA, TORSTEN HOEFLER

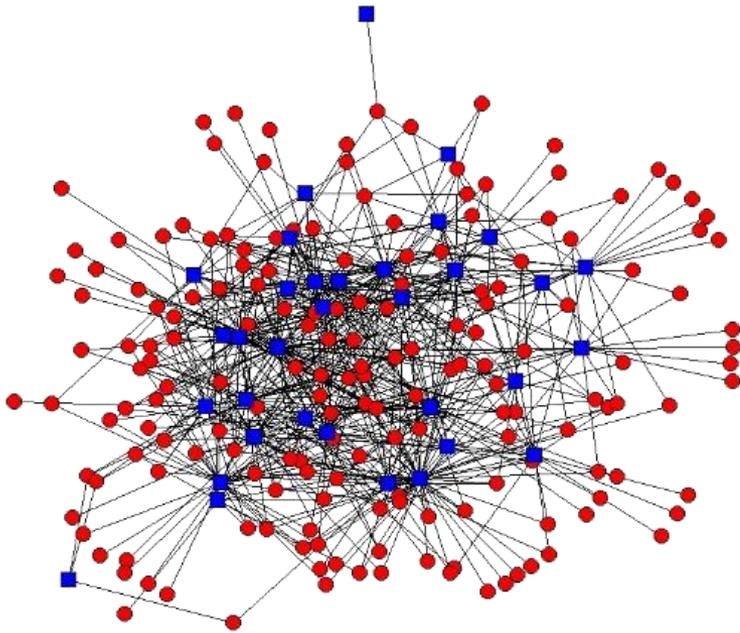


# LARGE-SCALE IRREGULAR GRAPH PROCESSING



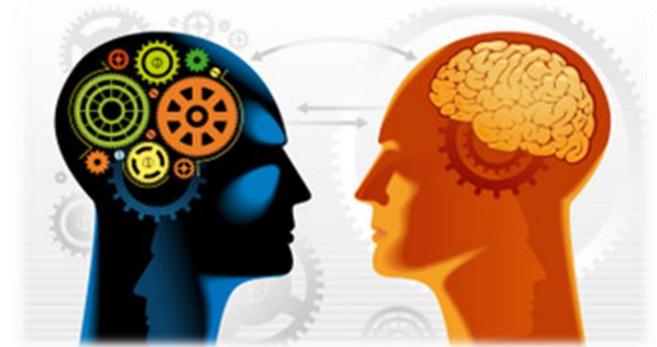
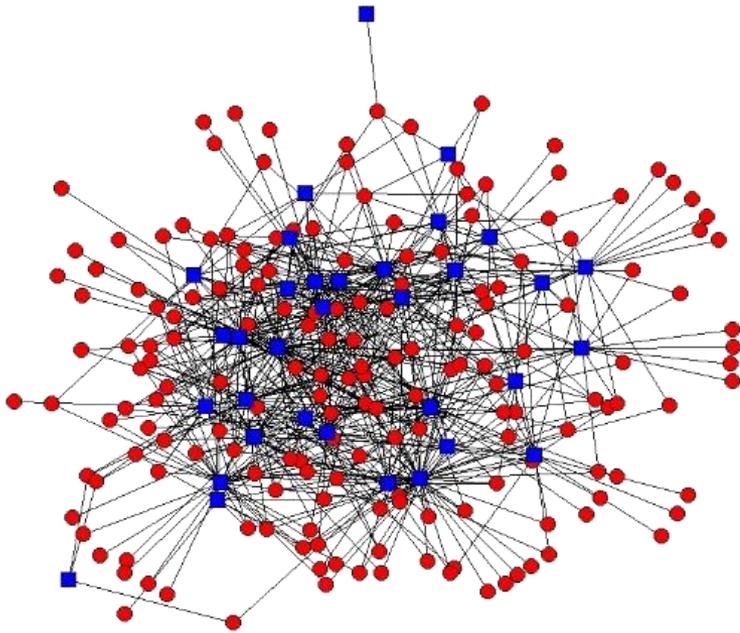
# LARGE-SCALE IRREGULAR GRAPH PROCESSING

- Becoming more important [1]



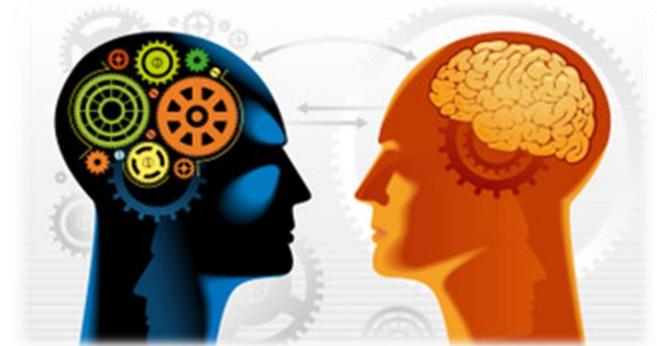
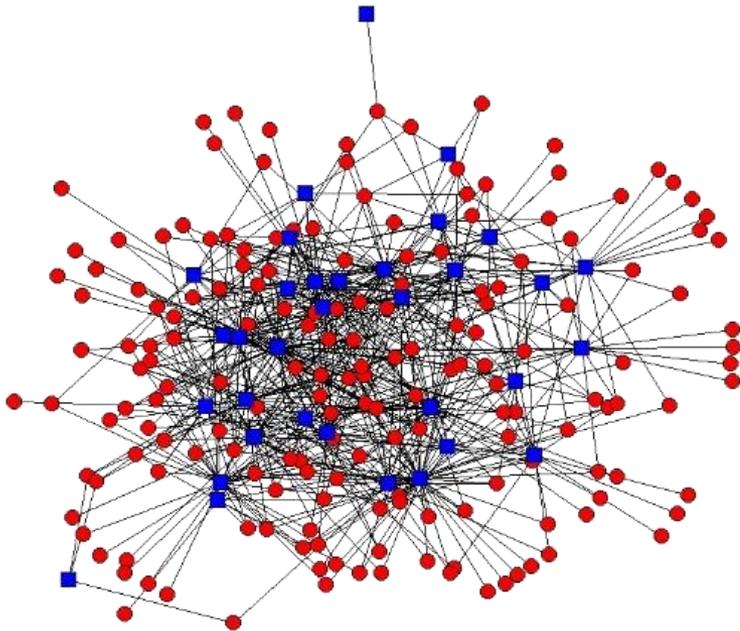
# LARGE-SCALE IRREGULAR GRAPH PROCESSING

- Becoming more important [1]
  - Machine learning



# LARGE-SCALE IRREGULAR GRAPH PROCESSING

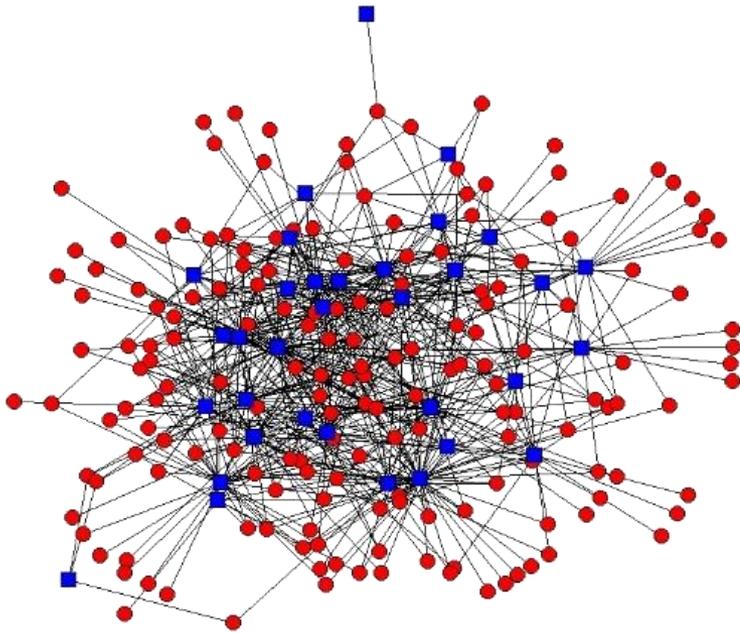
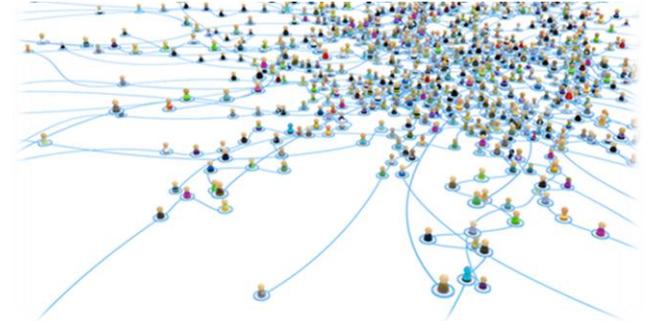
- Becoming more important [1]
  - Machine learning
  - Computational science



$$\frac{1}{\sqrt{2}} |\text{cat}\rangle + \frac{1}{\sqrt{2}} |\text{dog}\rangle$$

# LARGE-SCALE IRREGULAR GRAPH PROCESSING

- Becoming more important [1]
  - Machine learning
  - Computational science
  - Social network analysis



$$\frac{1}{\sqrt{2}} |\text{cat}\rangle + \frac{1}{\sqrt{2}} |\text{dog}\rangle$$

# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS

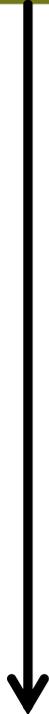
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS

Proc p

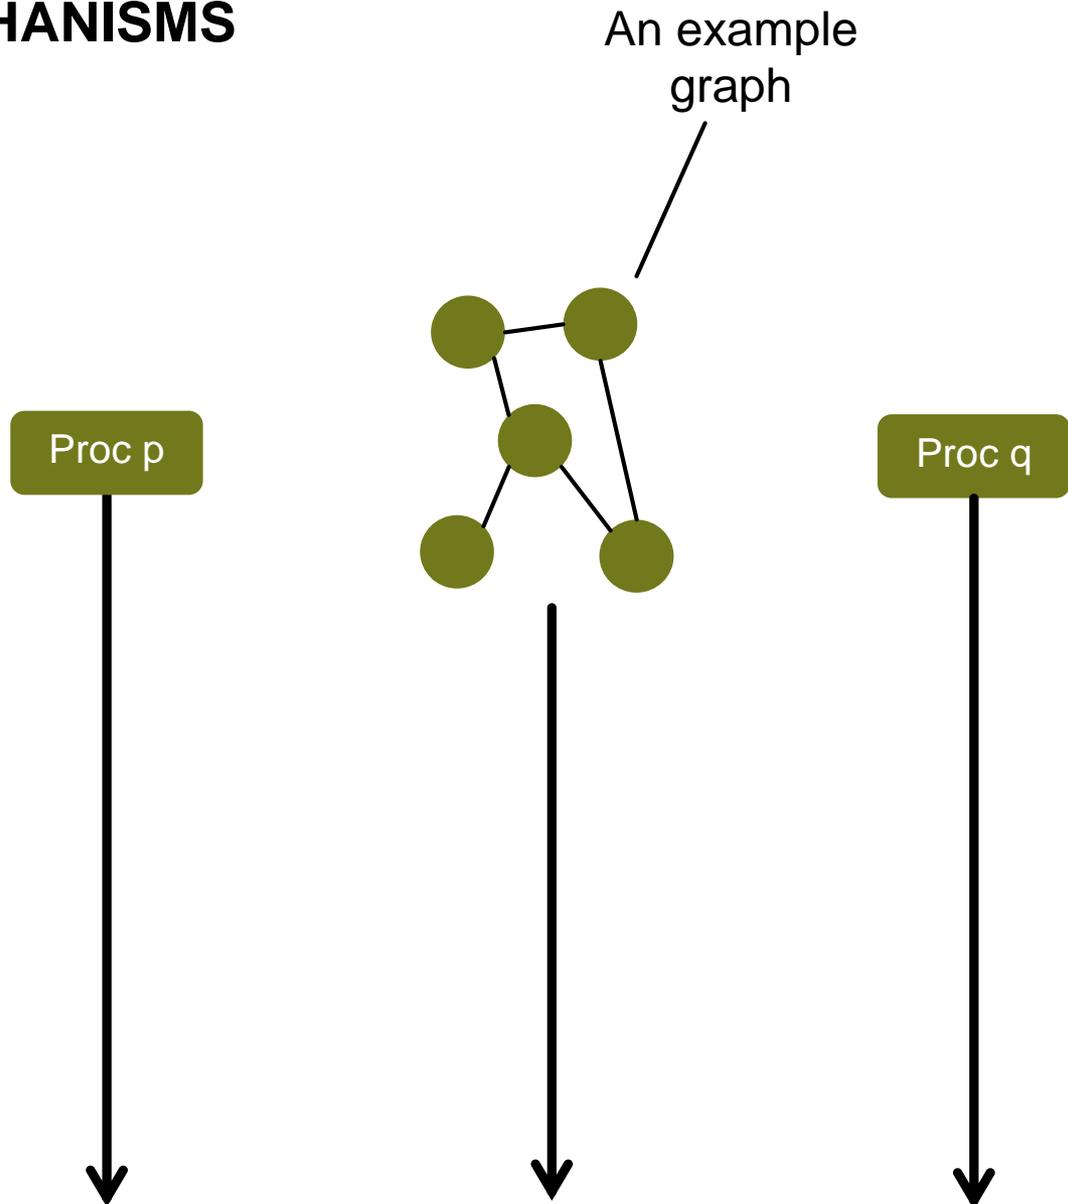


Proc q



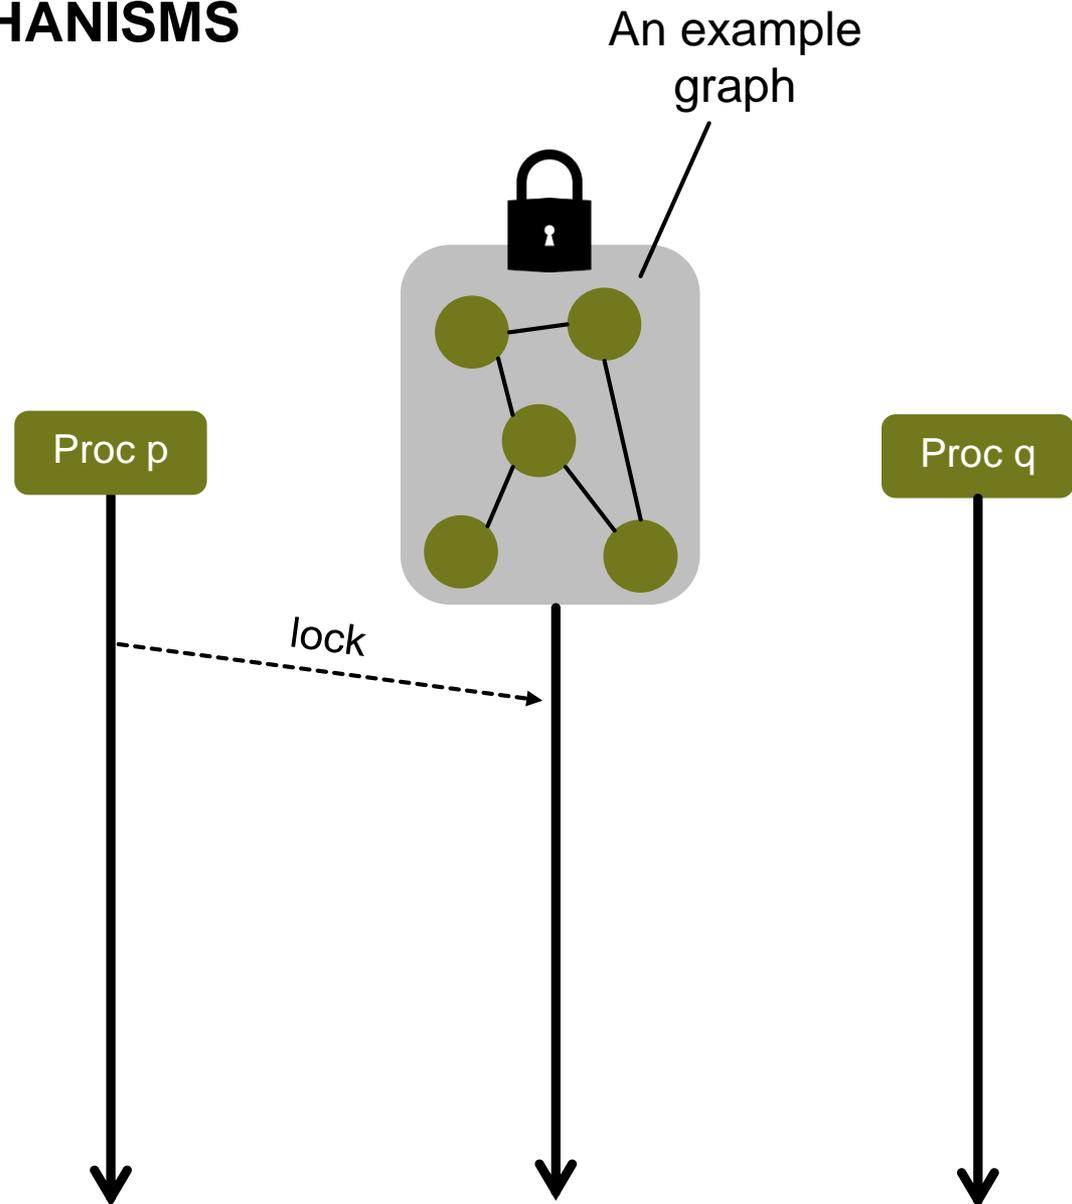
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS



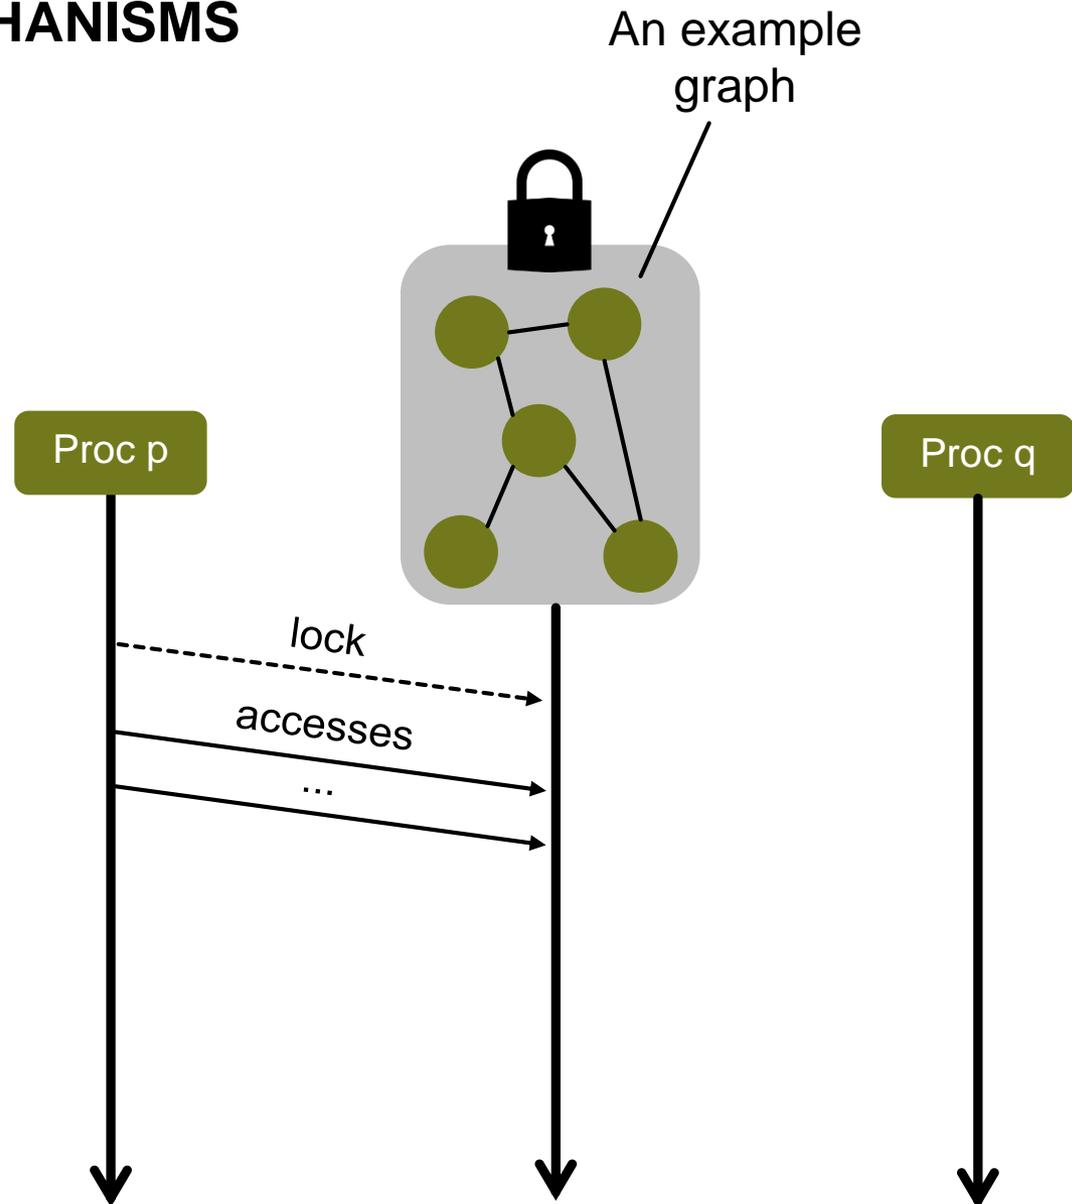
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS



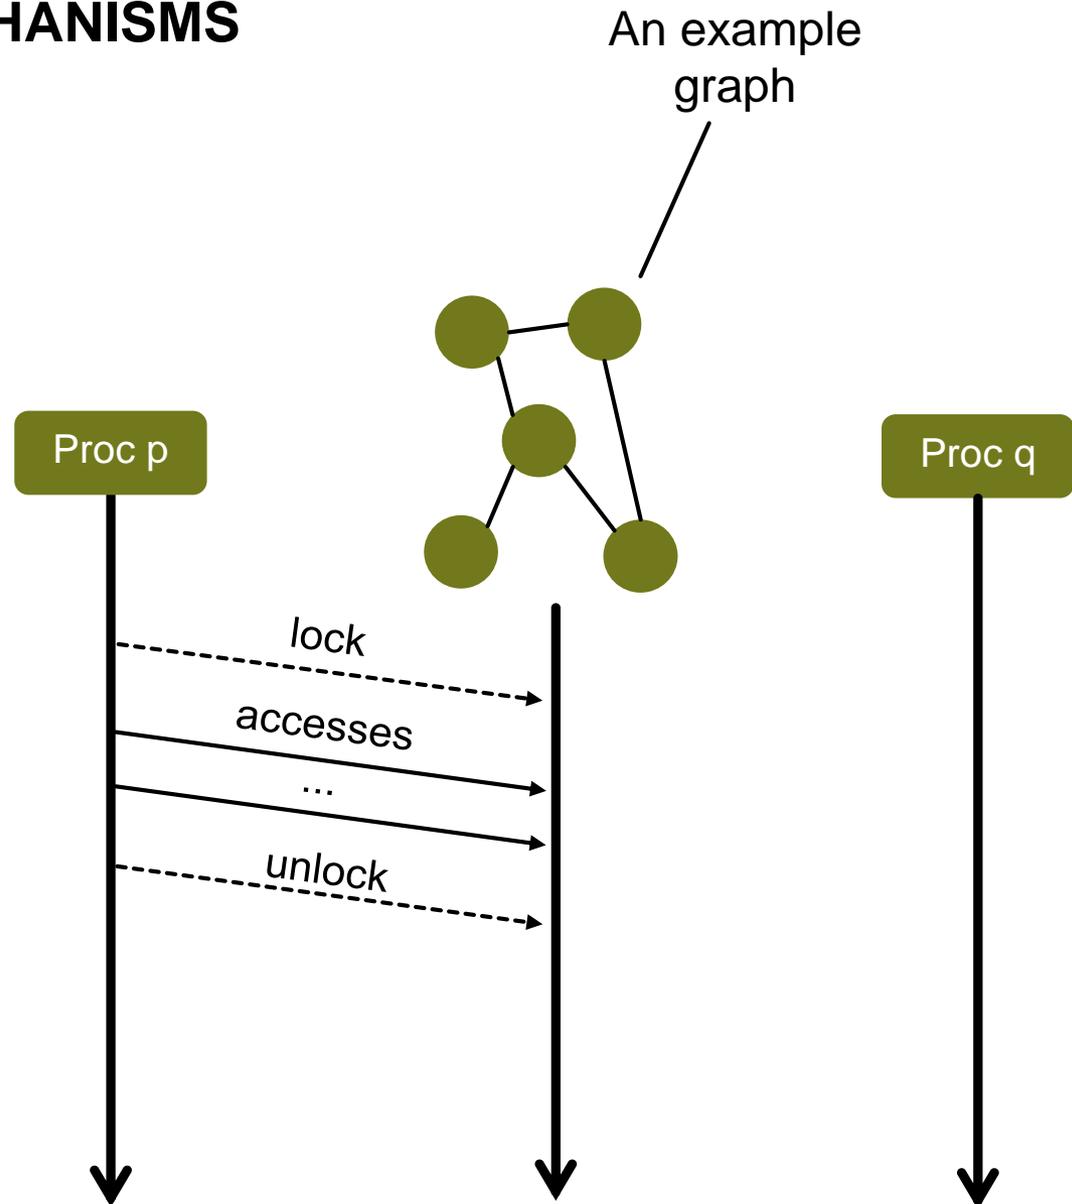
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS



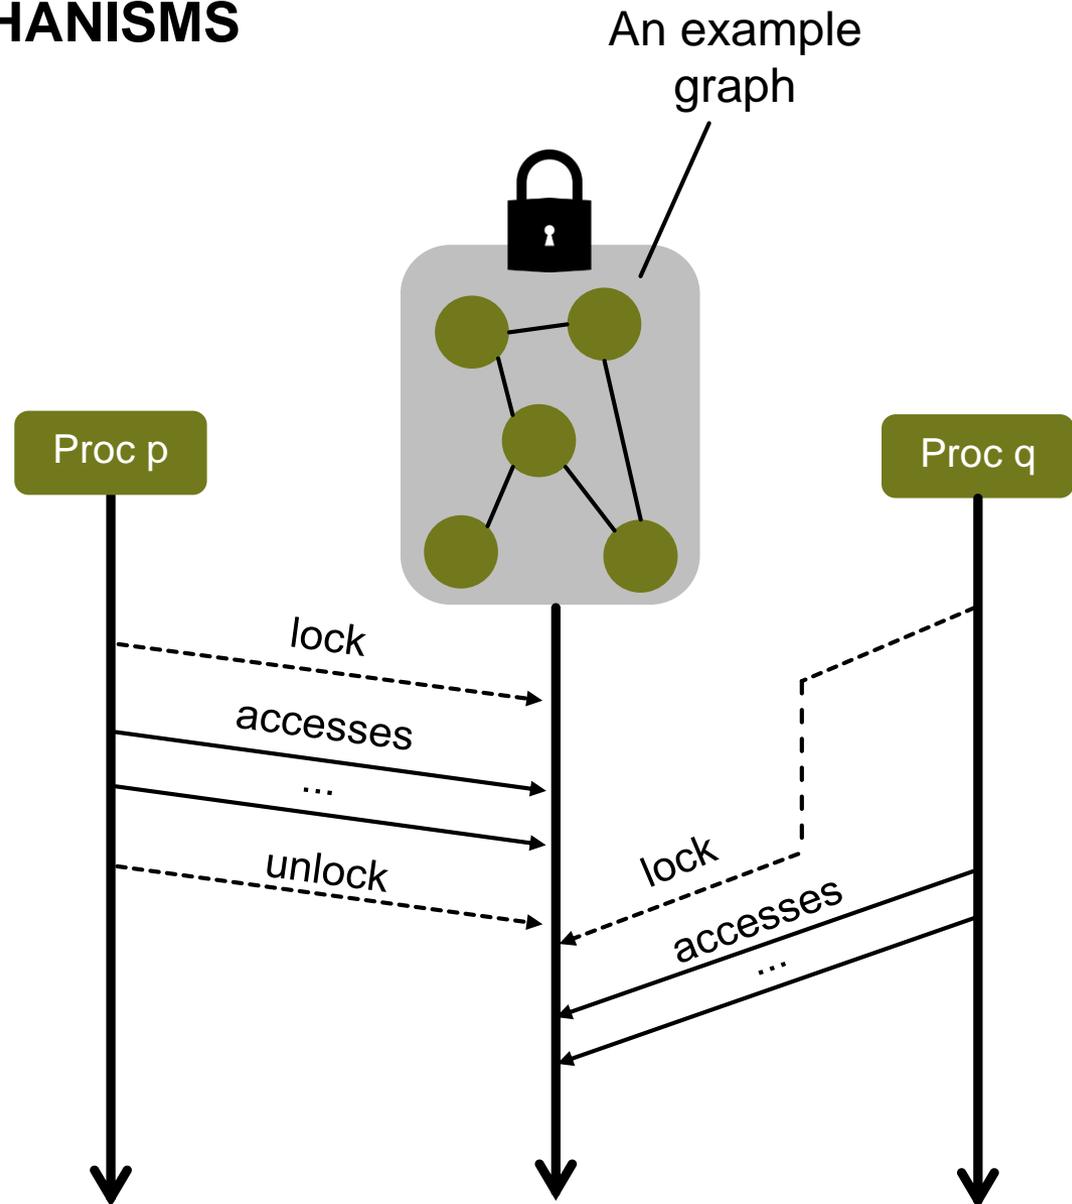
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS



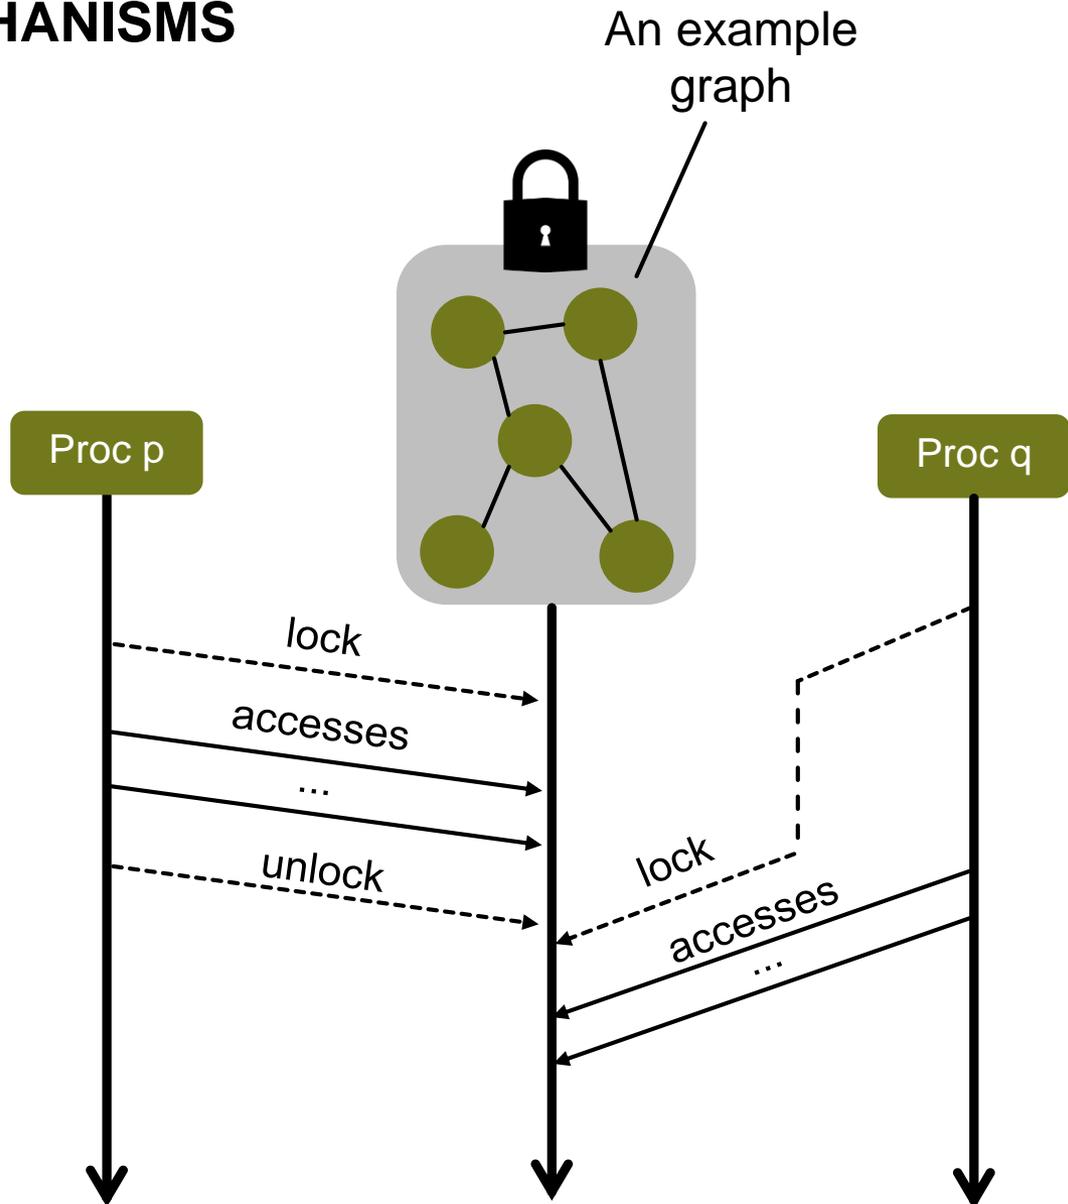
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS



# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS

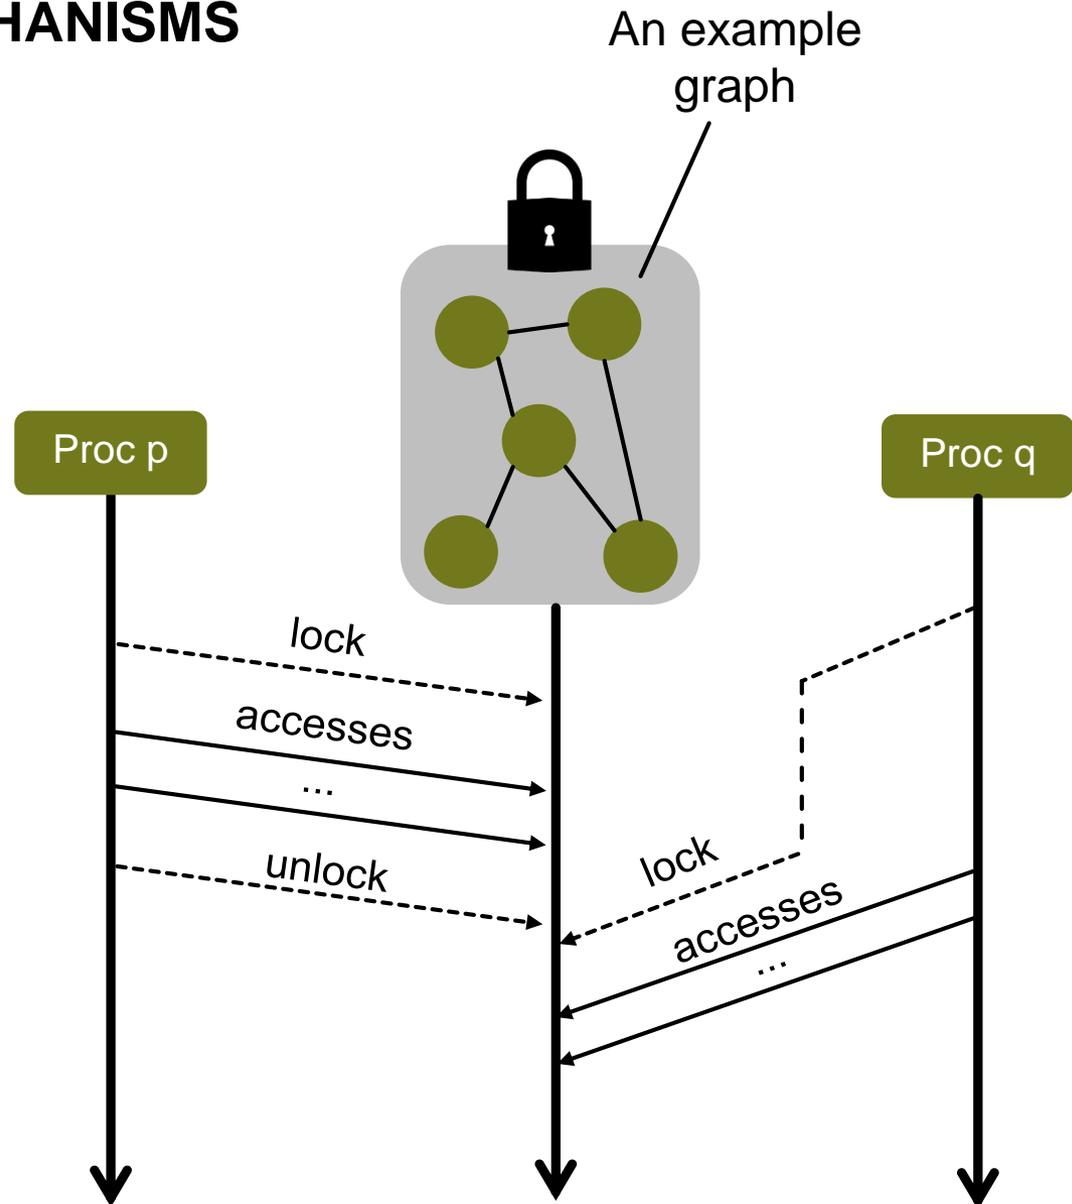


# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS

✓ Simple protocols

✗ Serialization



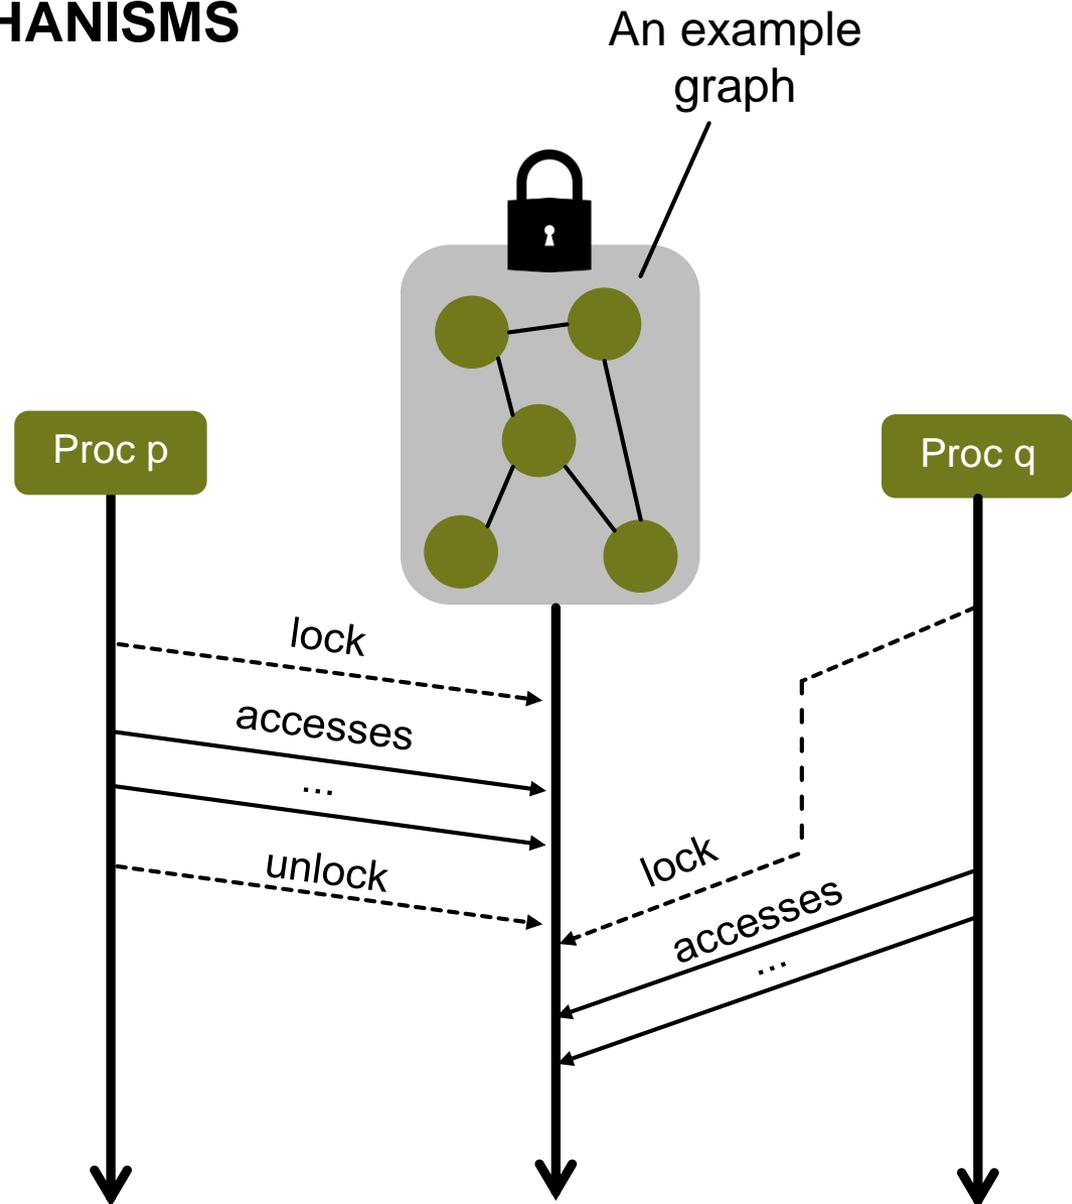
# SYNCHRONIZATION MECHANISMS

## COARSE LOCKS

✓ Simple protocols

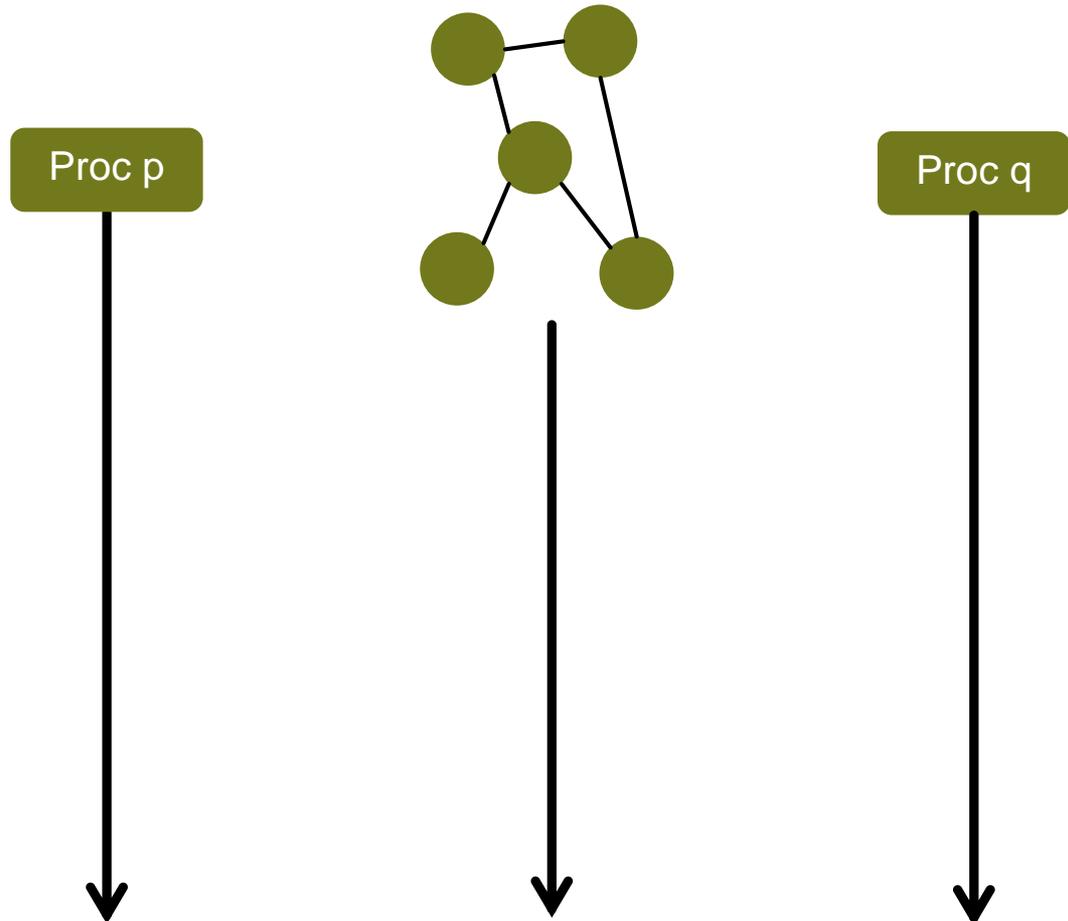
✗ Serialization

✗ Detrimental performance



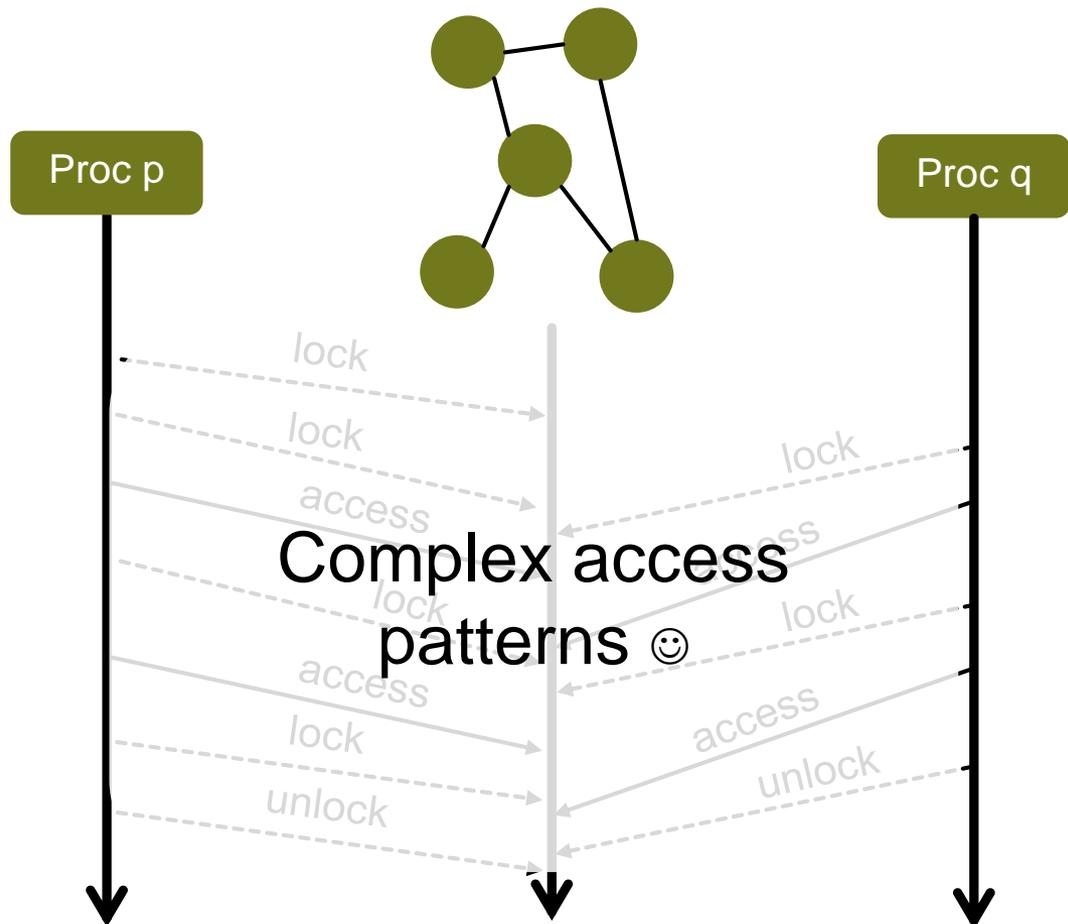
# SYNCHRONIZATION MECHANISMS

## FINE LOCKS



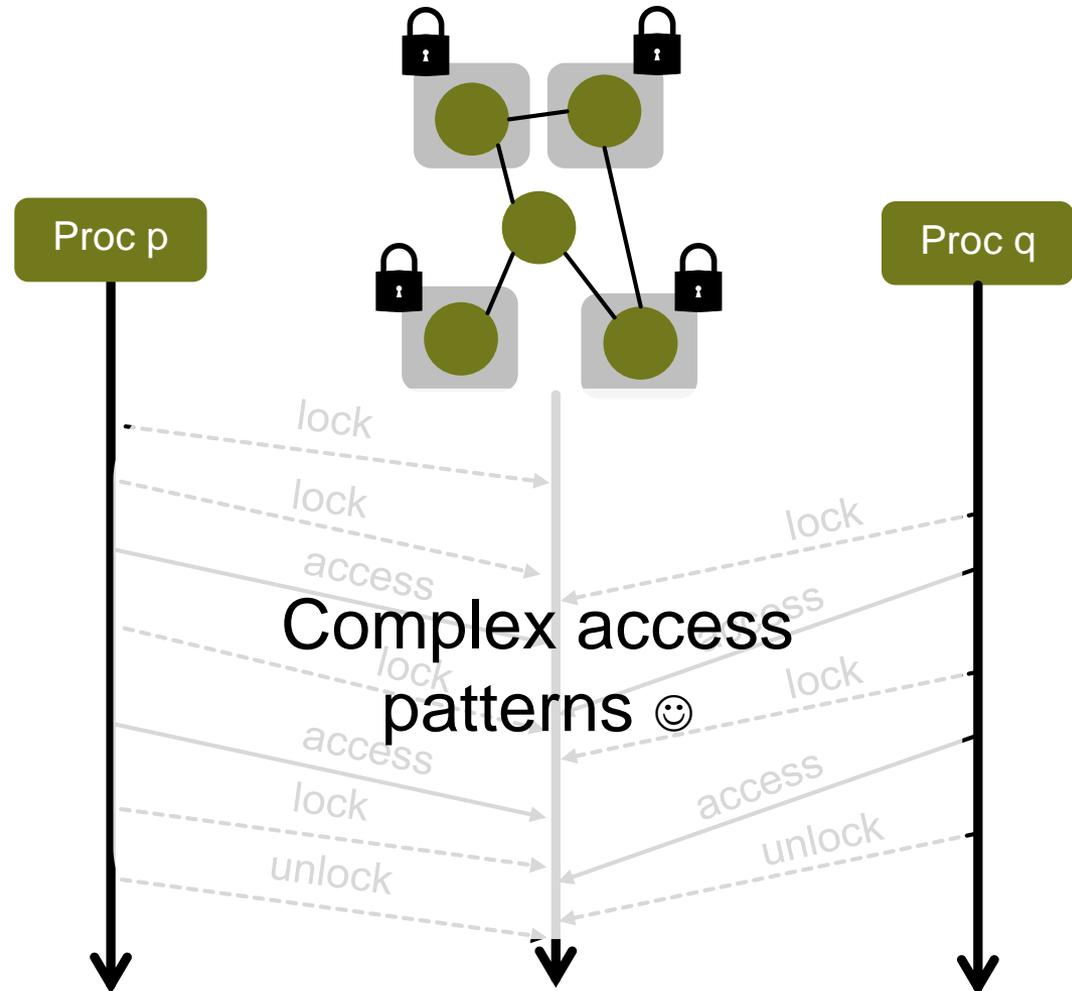
# SYNCHRONIZATION MECHANISMS

## FINE LOCKS



# SYNCHRONIZATION MECHANISMS

## FINE LOCKS

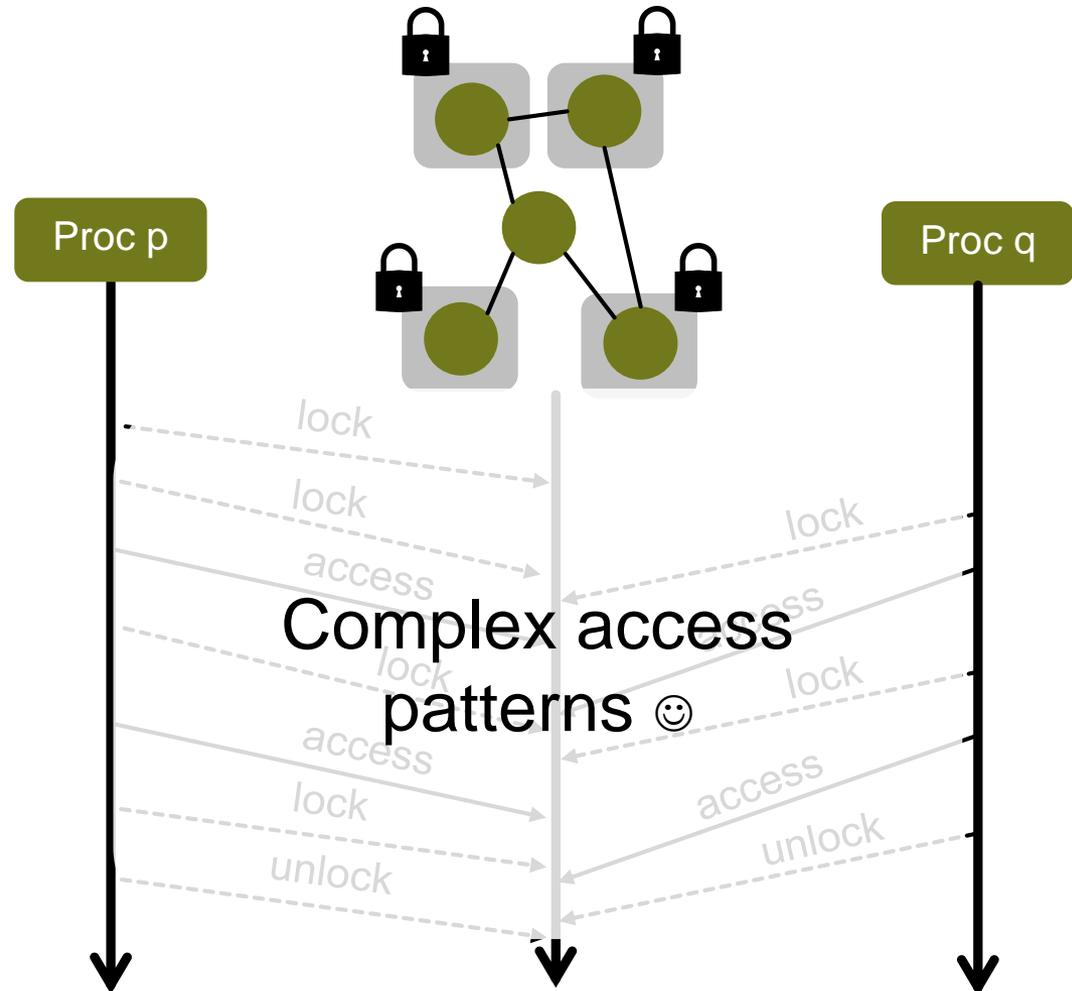


# SYNCHRONIZATION MECHANISMS

## FINE LOCKS



Higher performance possible



# SYNCHRONIZATION MECHANISMS

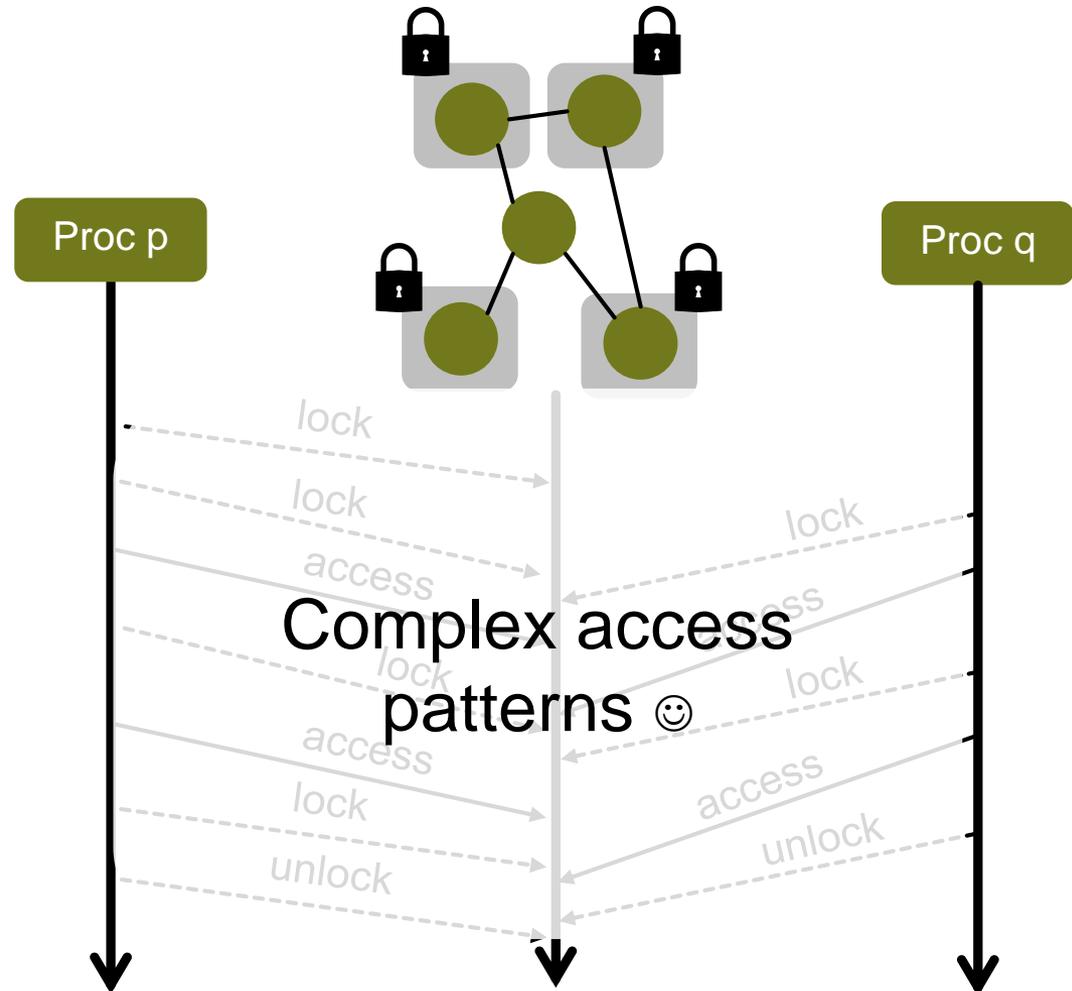
## FINE LOCKS



Higher performance possible



Complex protocols



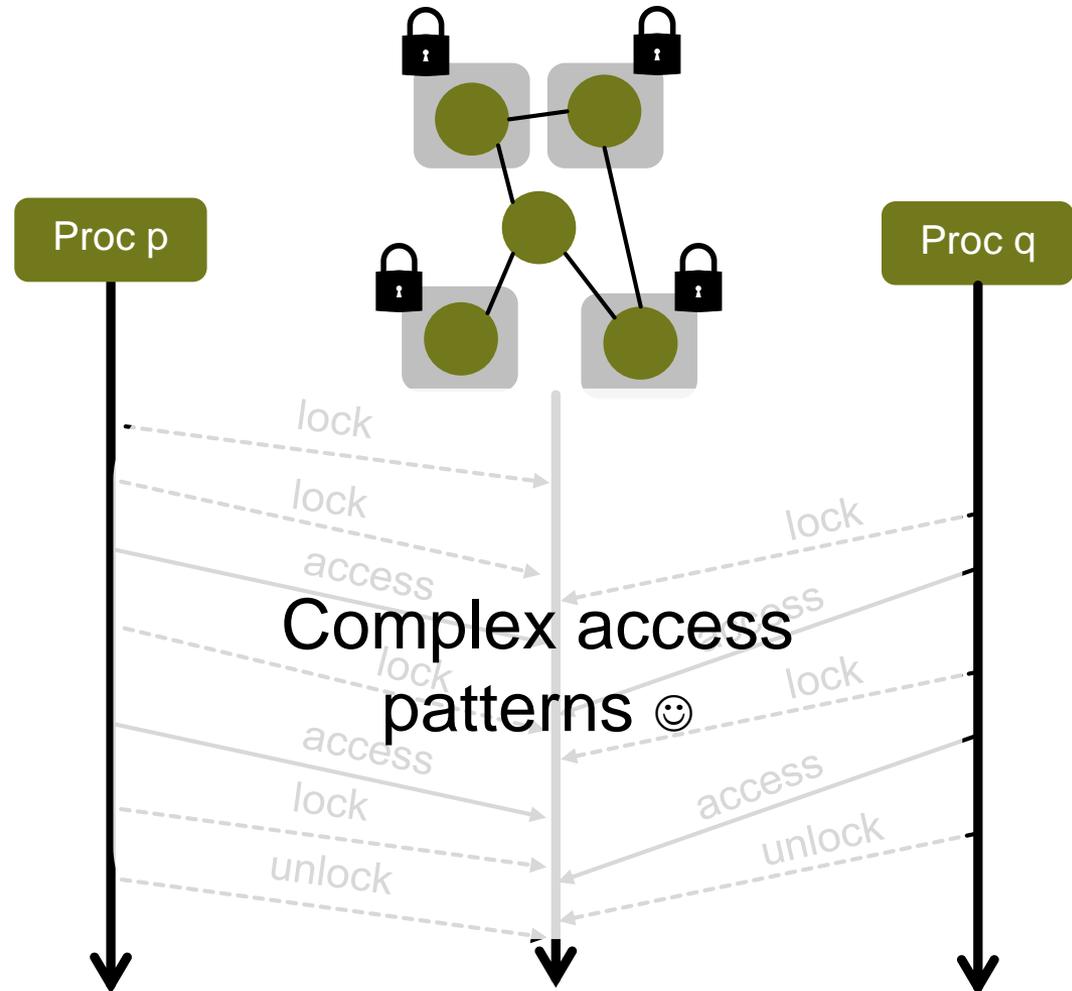
# SYNCHRONIZATION MECHANISMS

## FINE LOCKS

✓ Higher performance possible

✗ Complex protocols

✗ Risk of deadlocks



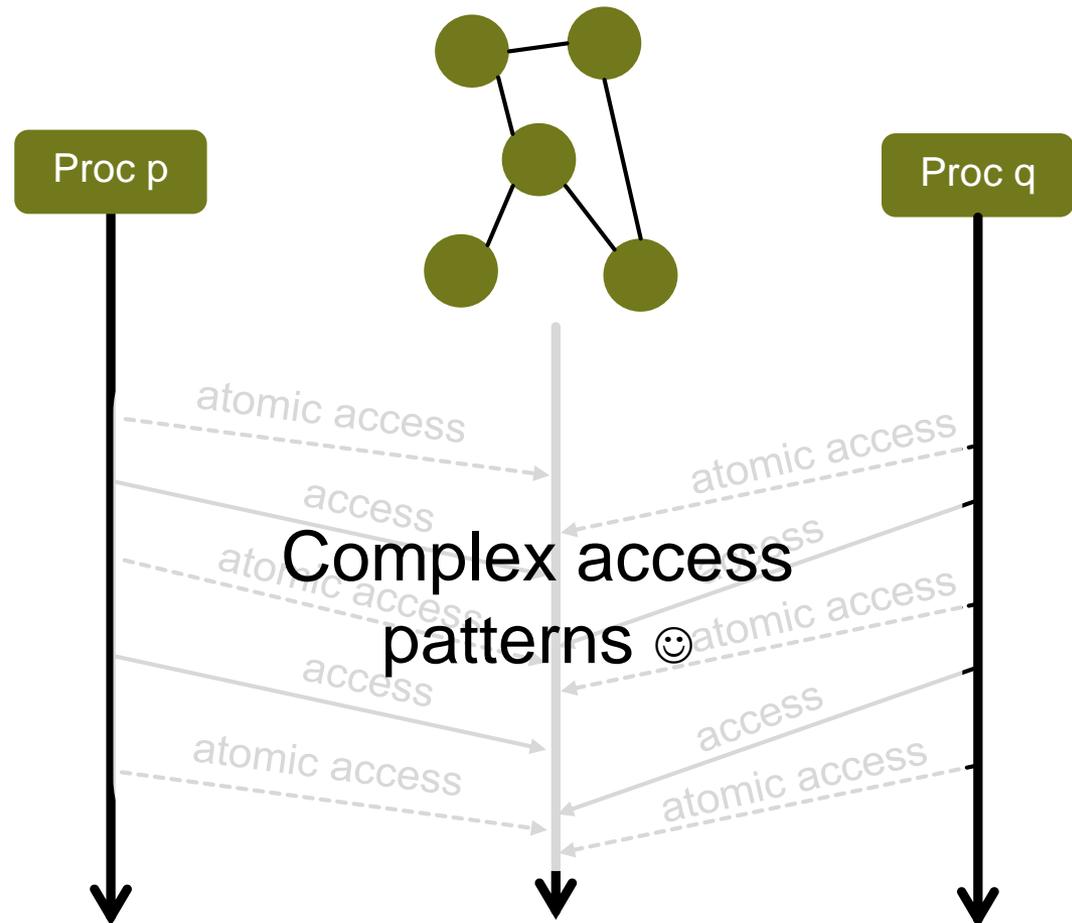


# SYNCHRONIZATION MECHANISMS

## ATOMIC OPERATIONS



High performance (may be challenging to get)



# SYNCHRONIZATION MECHANISMS

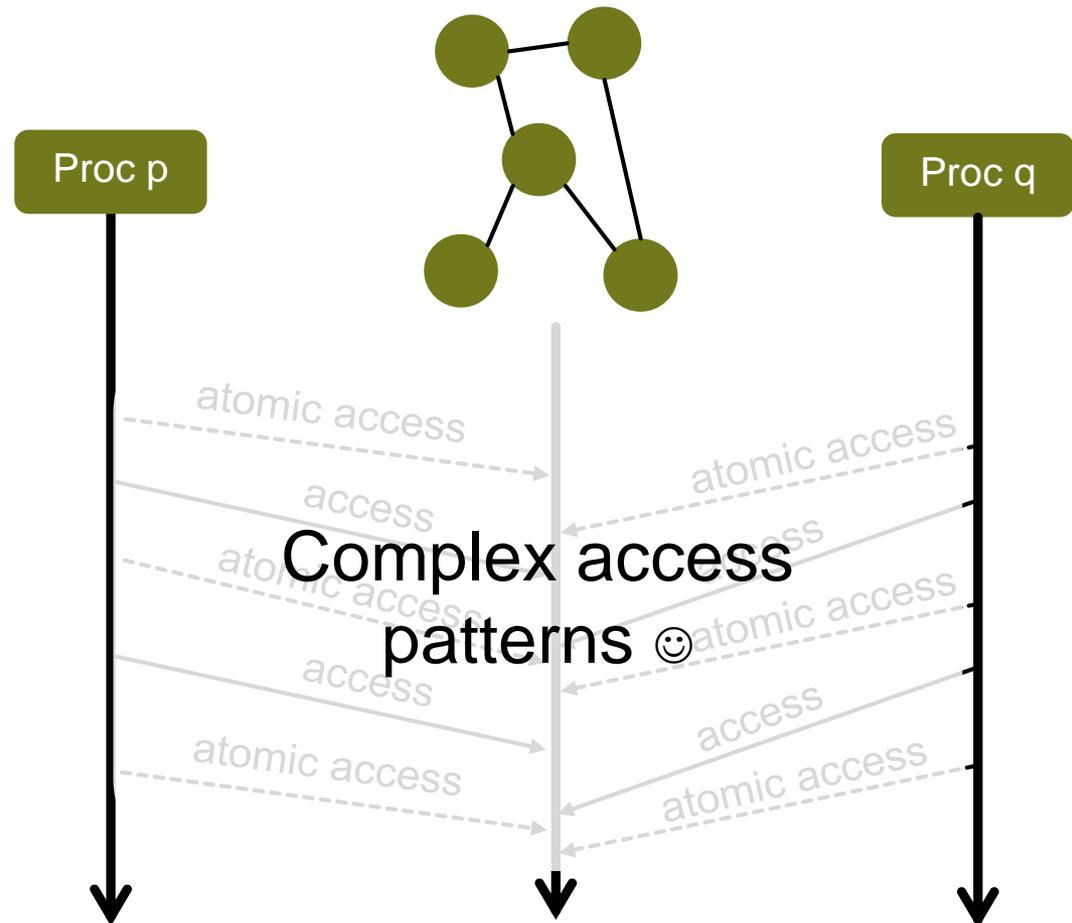
## ATOMIC OPERATIONS



High performance (may be challenging to get)



Complex protocols



# SYNCHRONIZATION MECHANISMS

## ATOMIC OPERATIONS



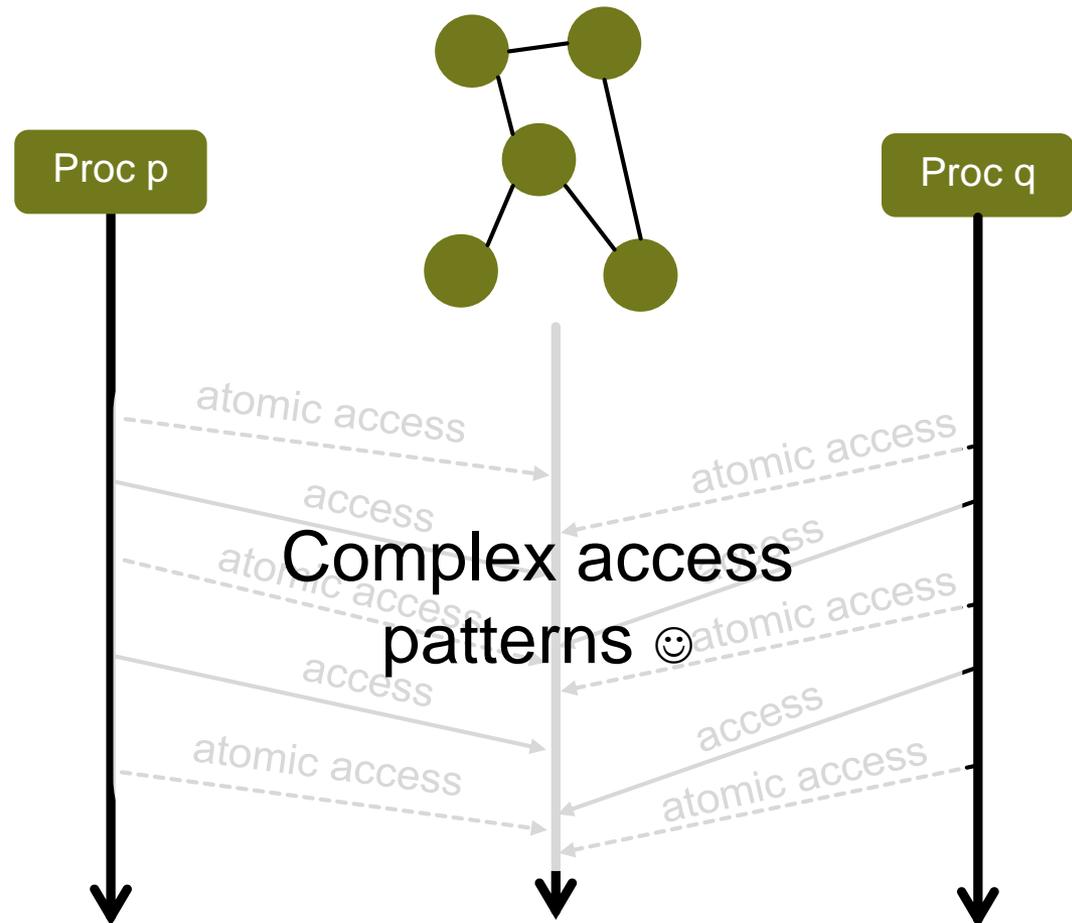
High performance (may be challenging to get)



Complex protocols

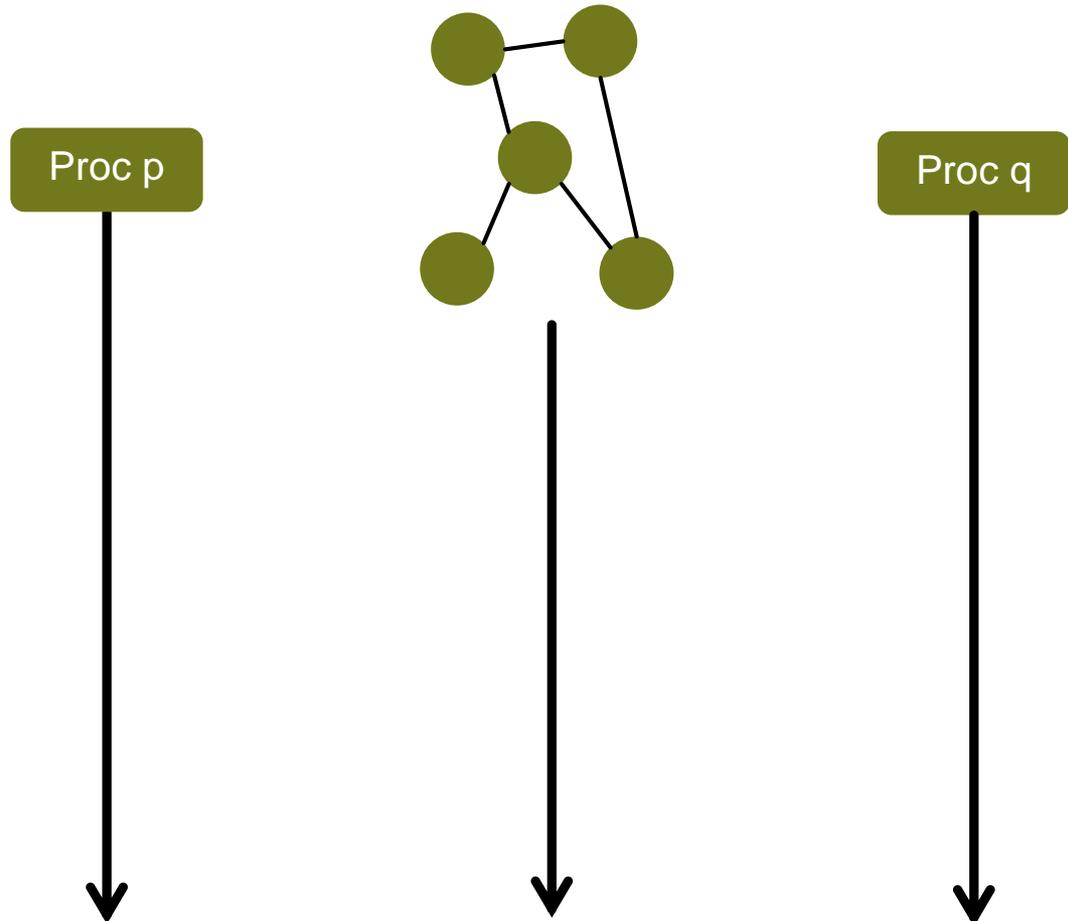


Subtle issues (ABA, ...)



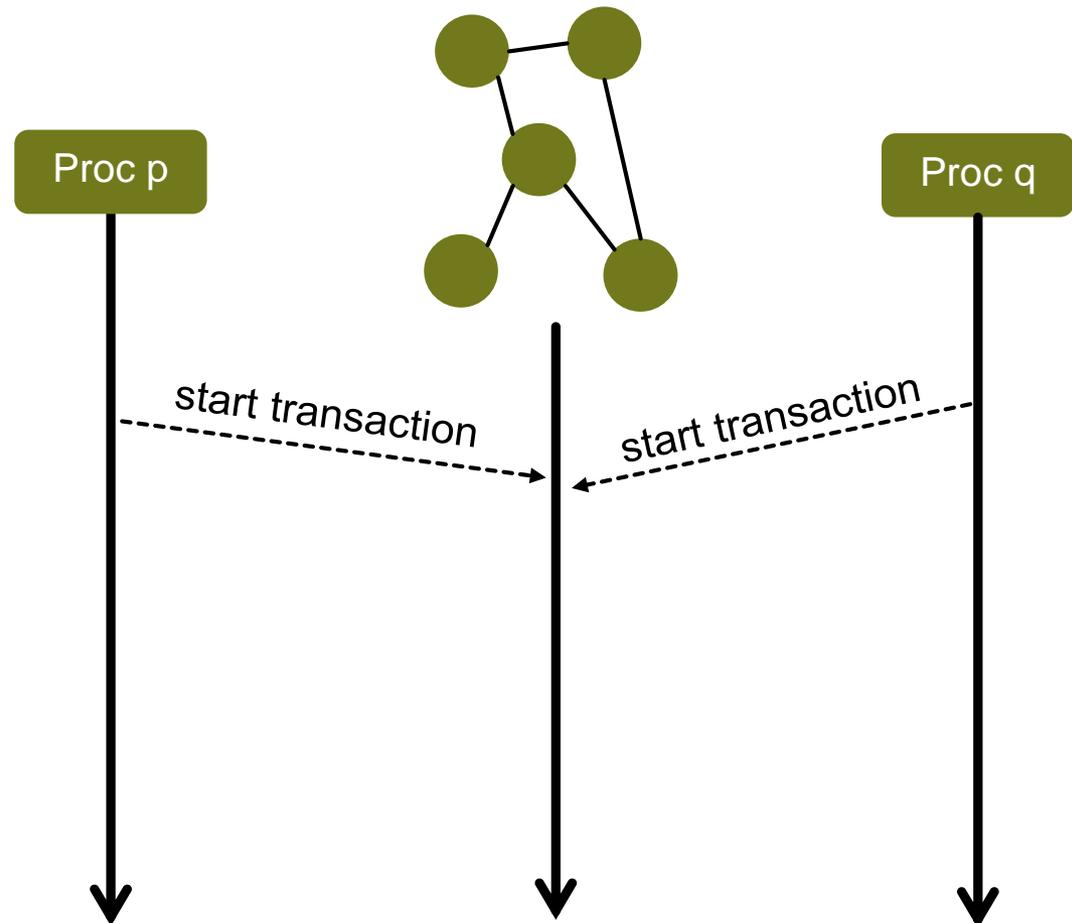
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



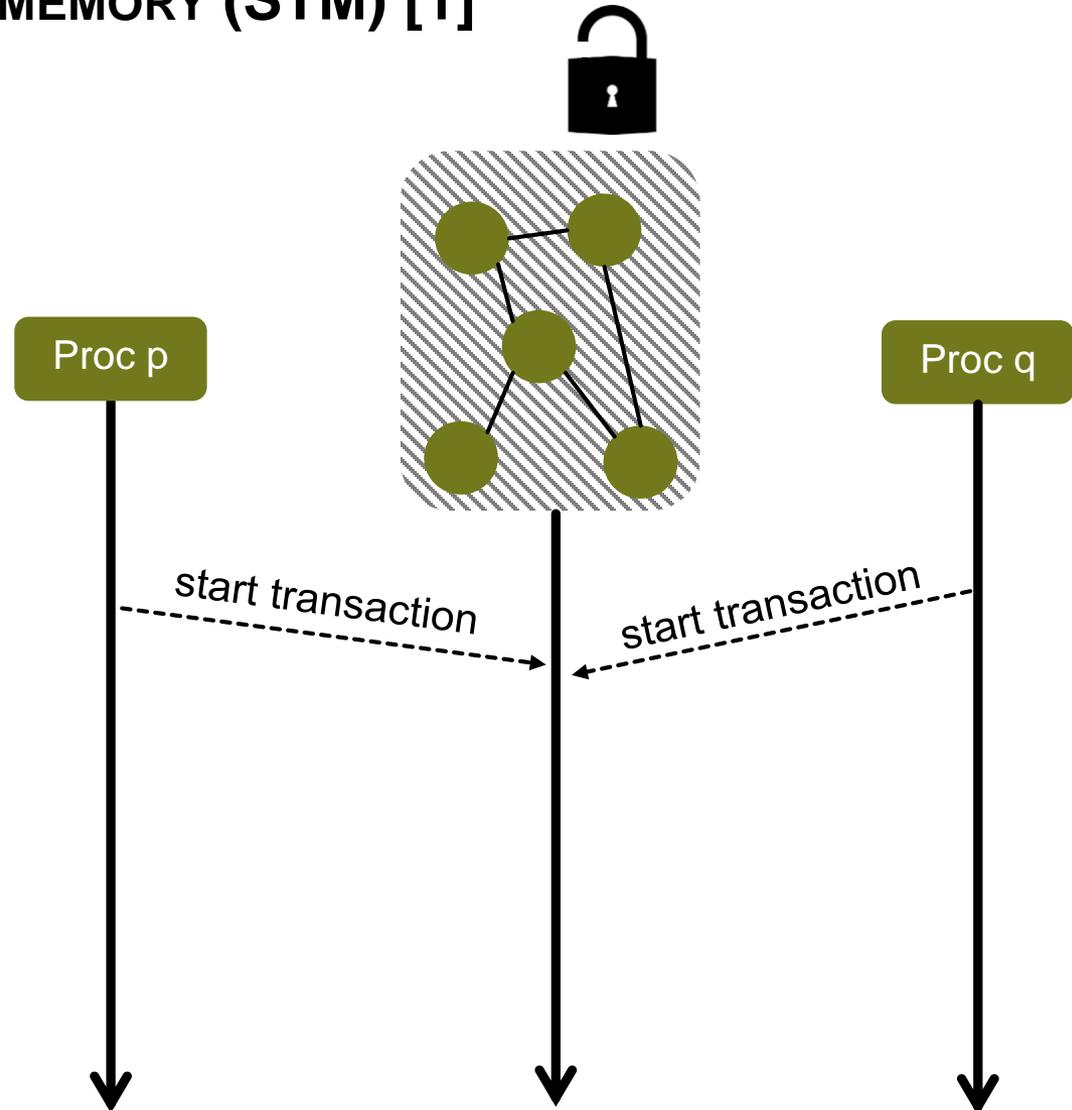
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



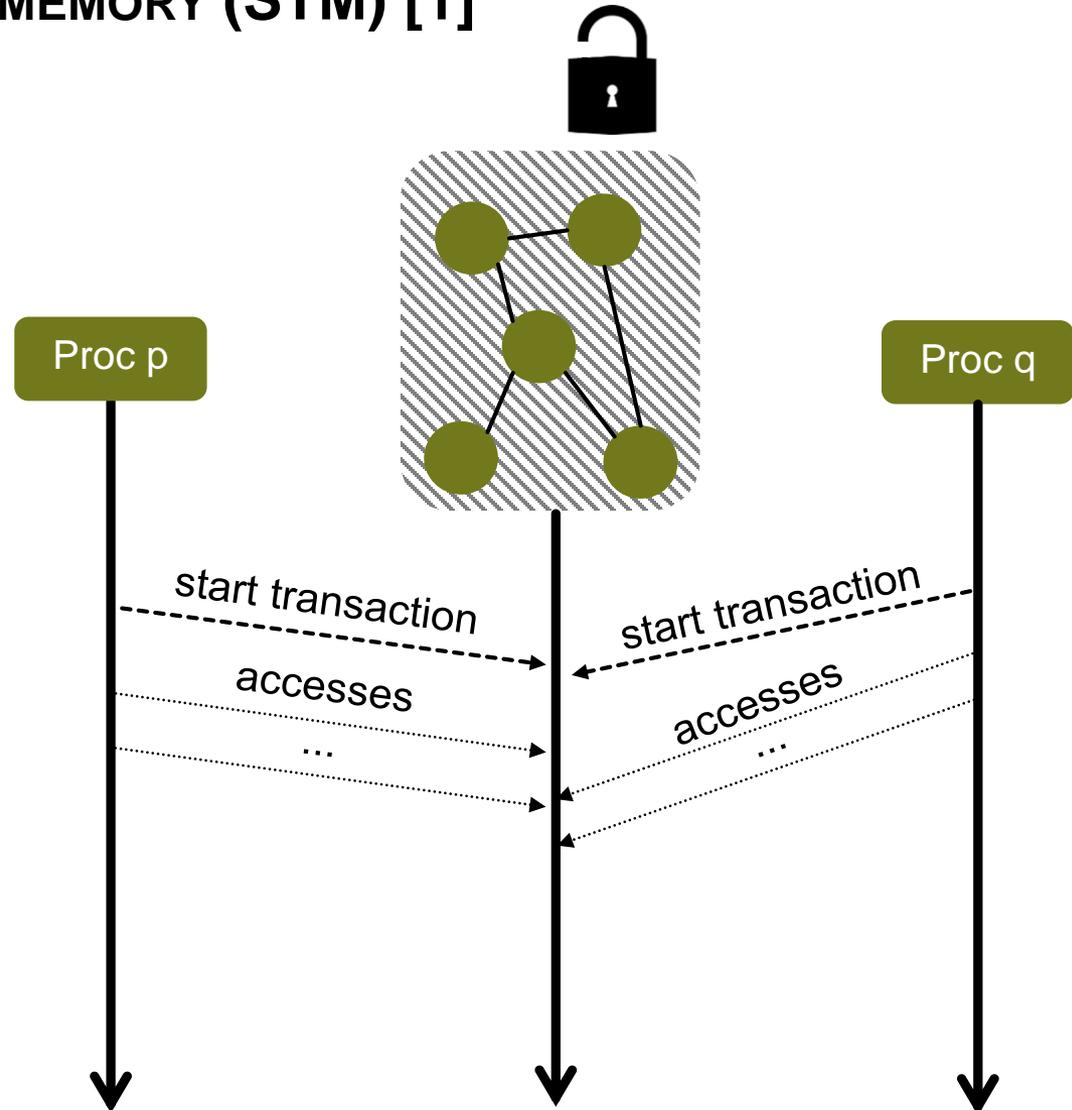
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



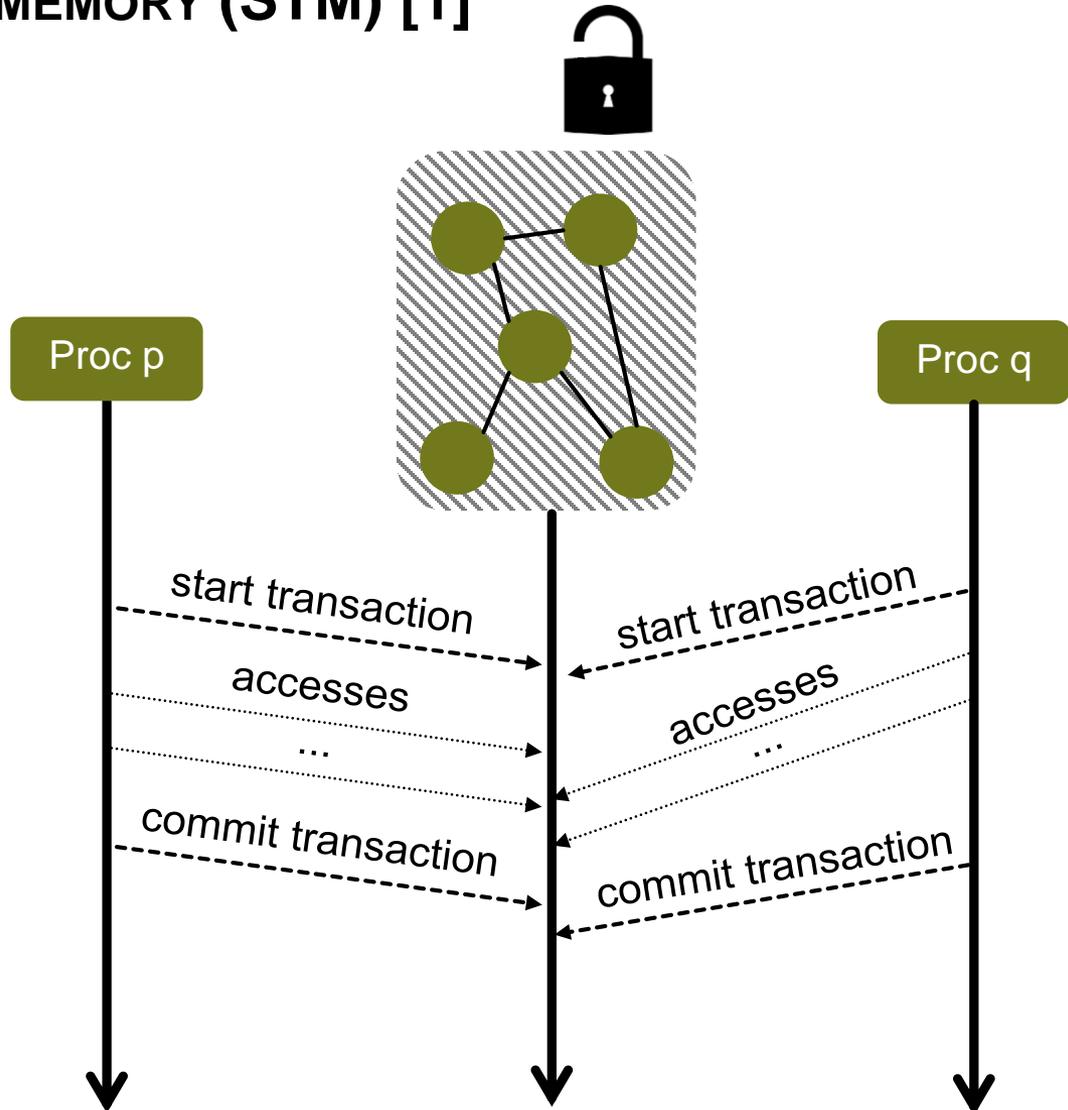
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



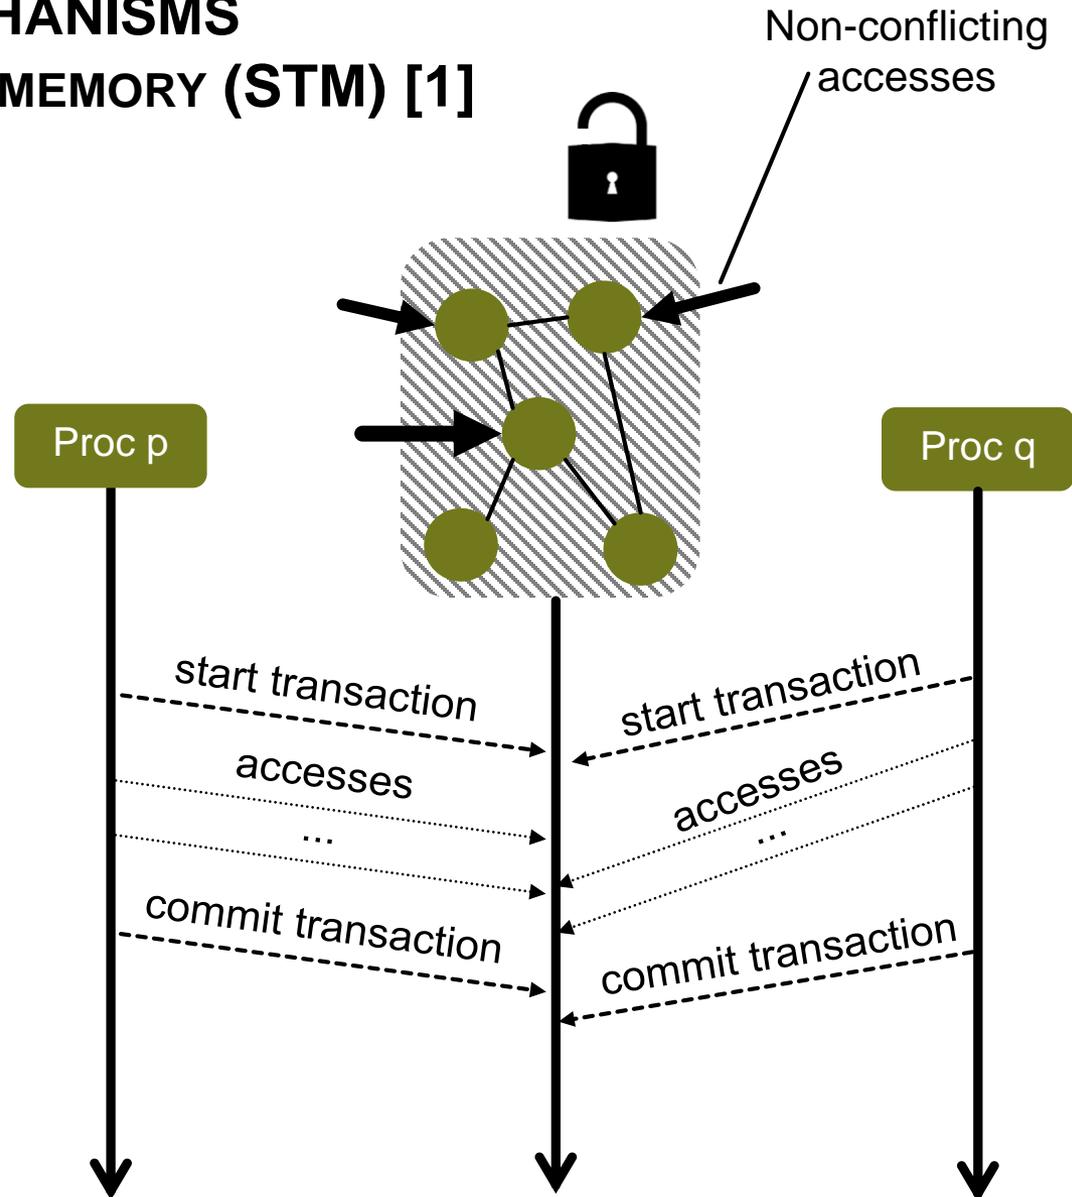
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



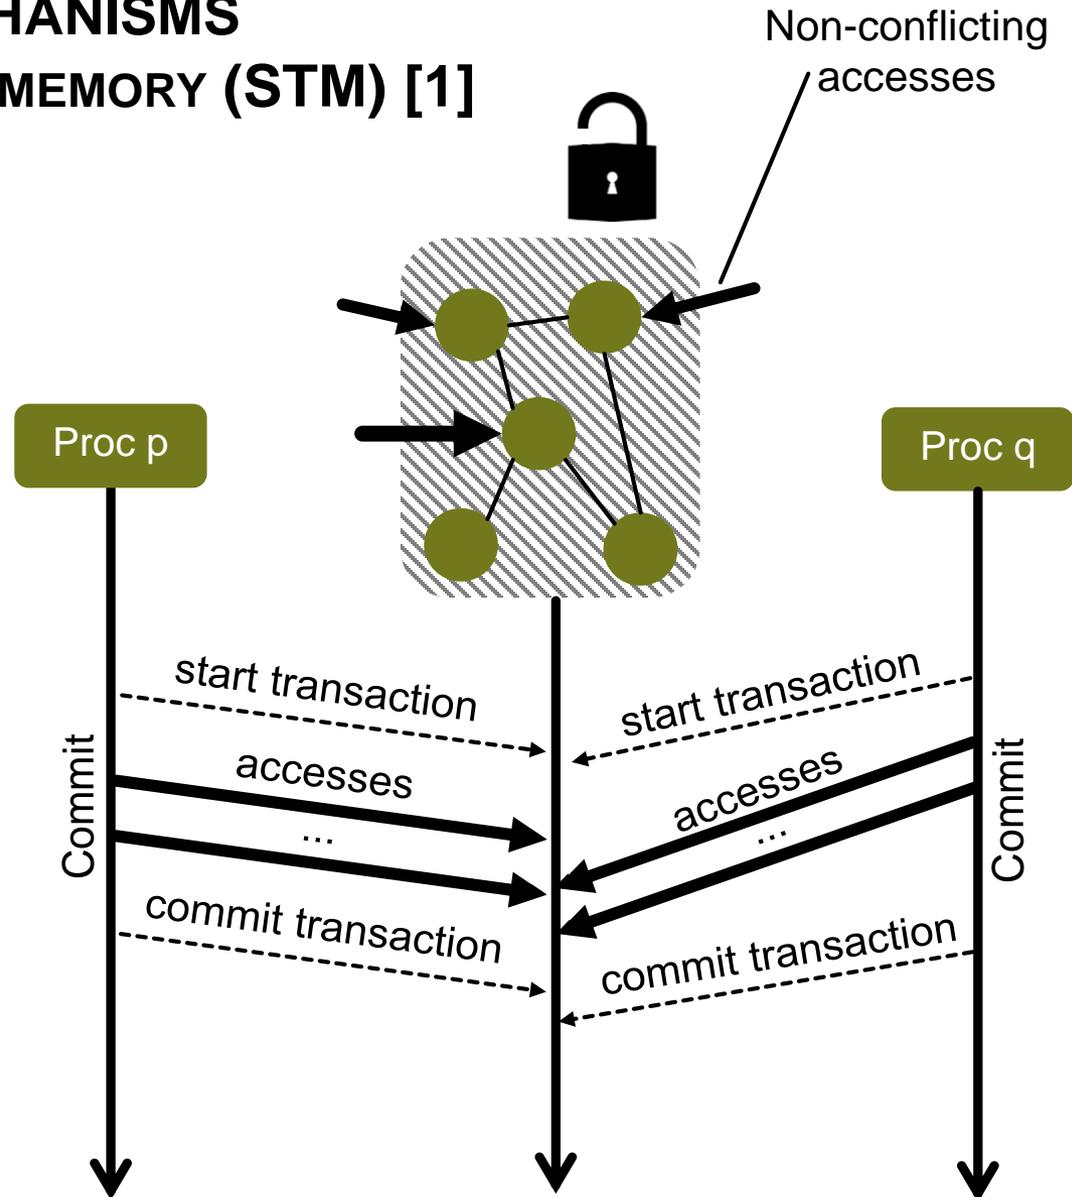
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



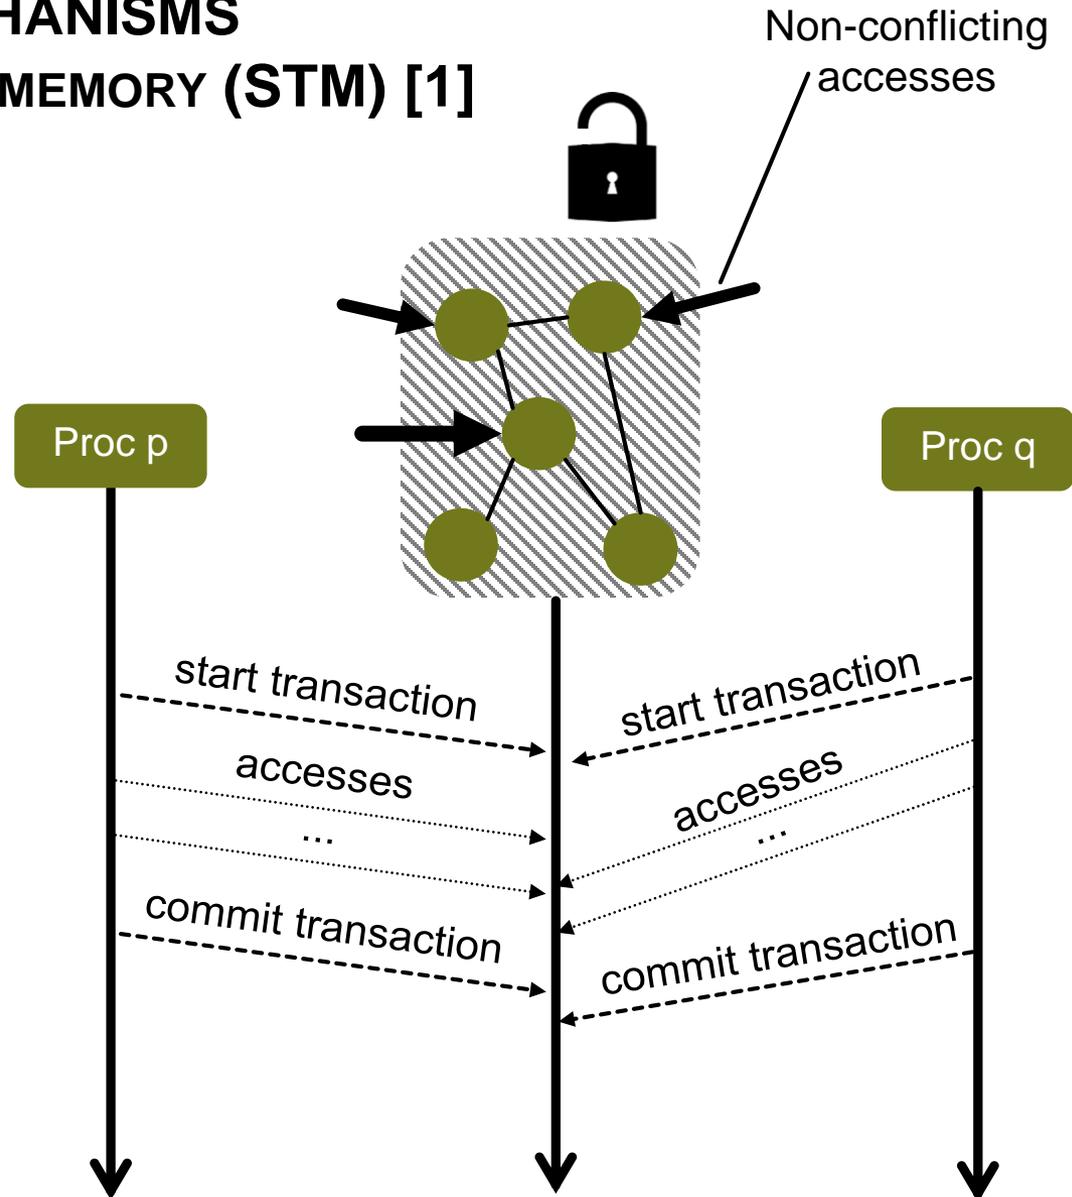
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



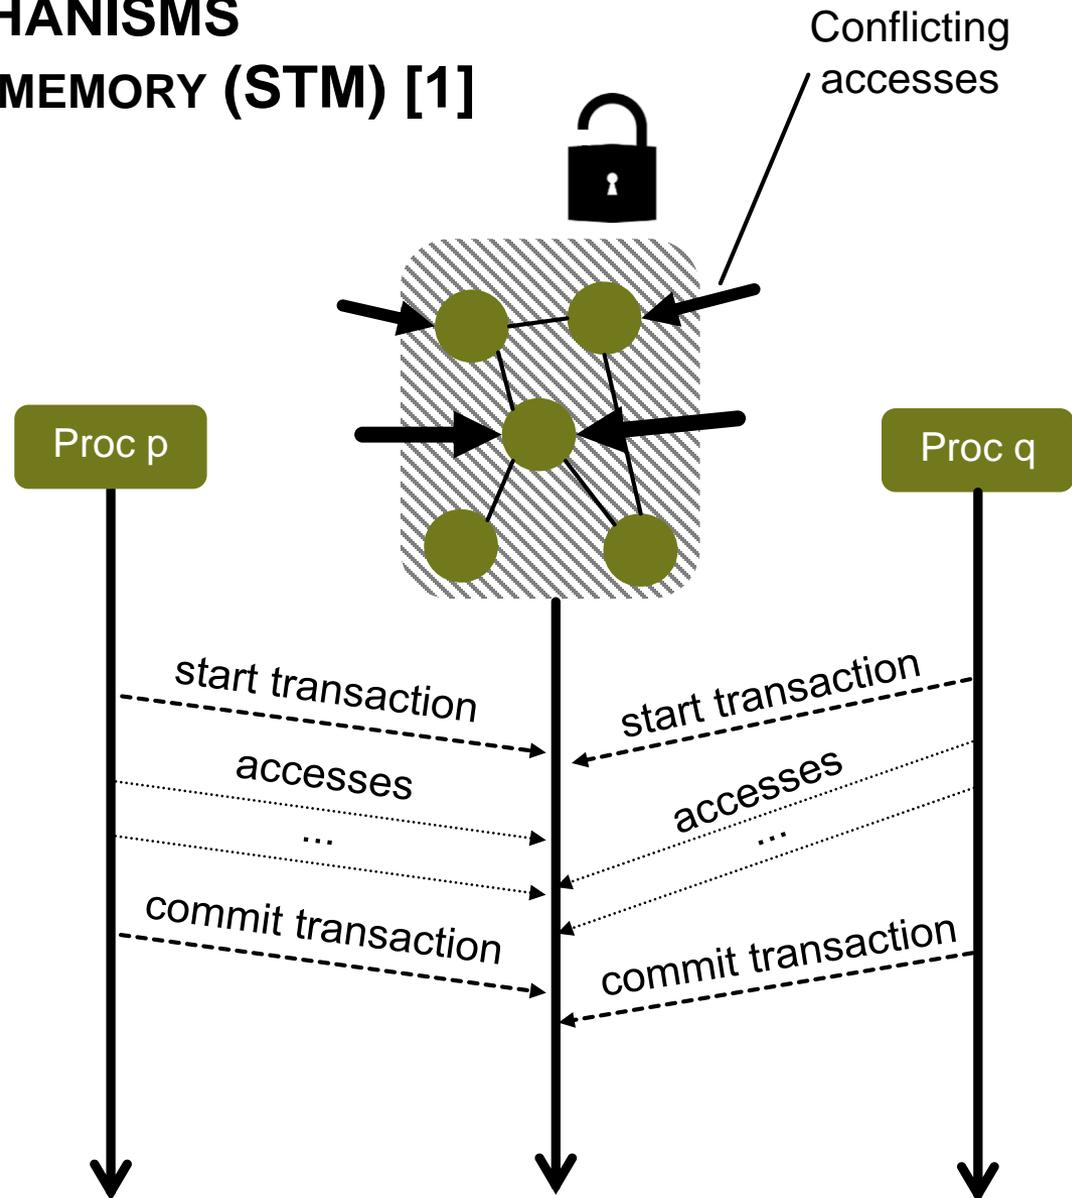
# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



# SYNCHRONIZATION MECHANISMS

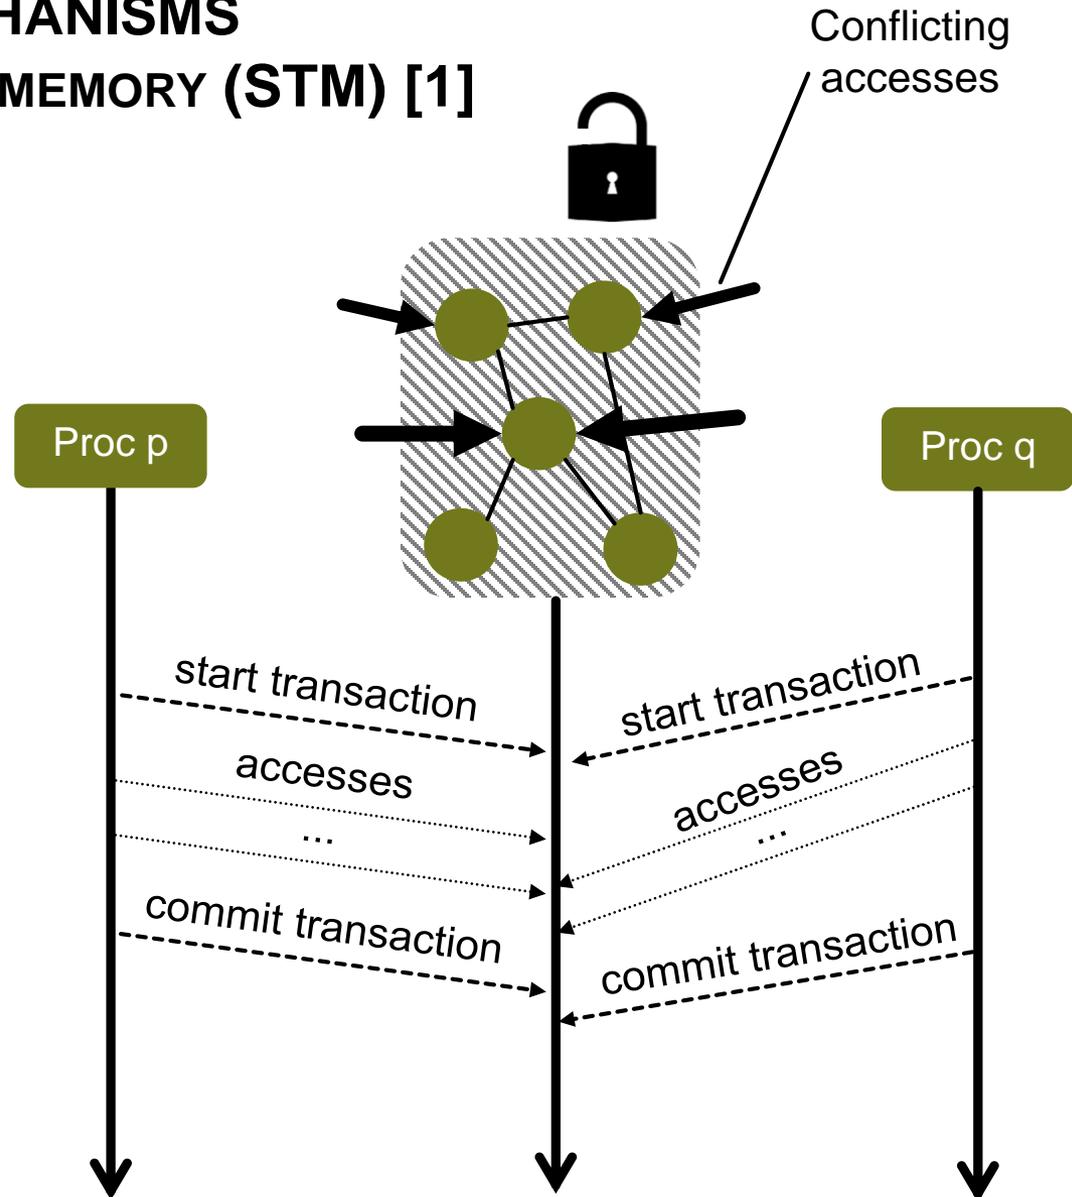
## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]



# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]

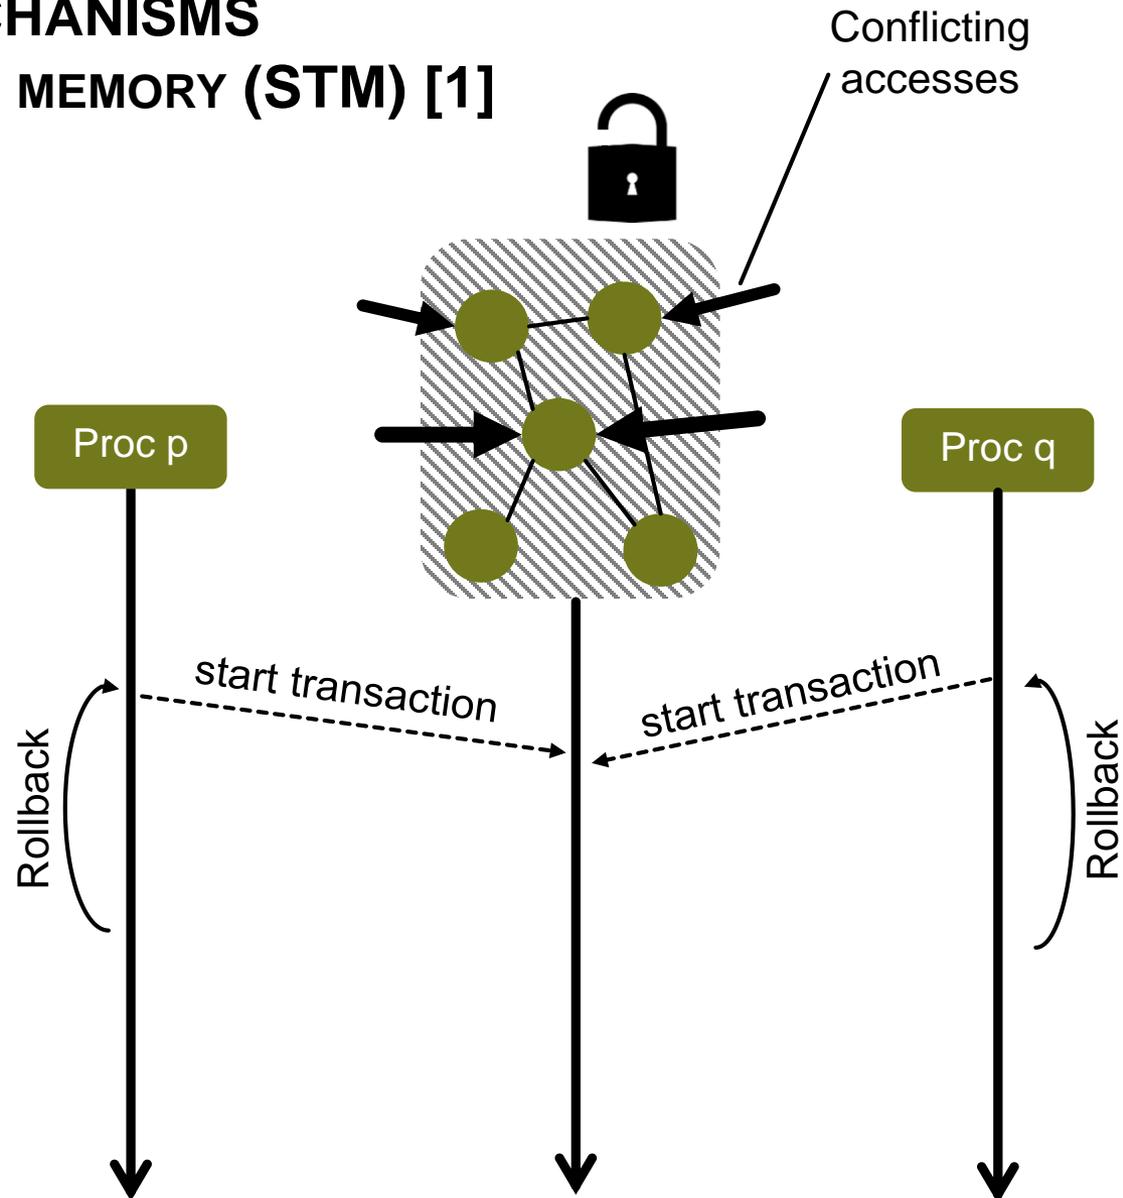
Conflicts solved with rollbacks and/or serialization.



# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]

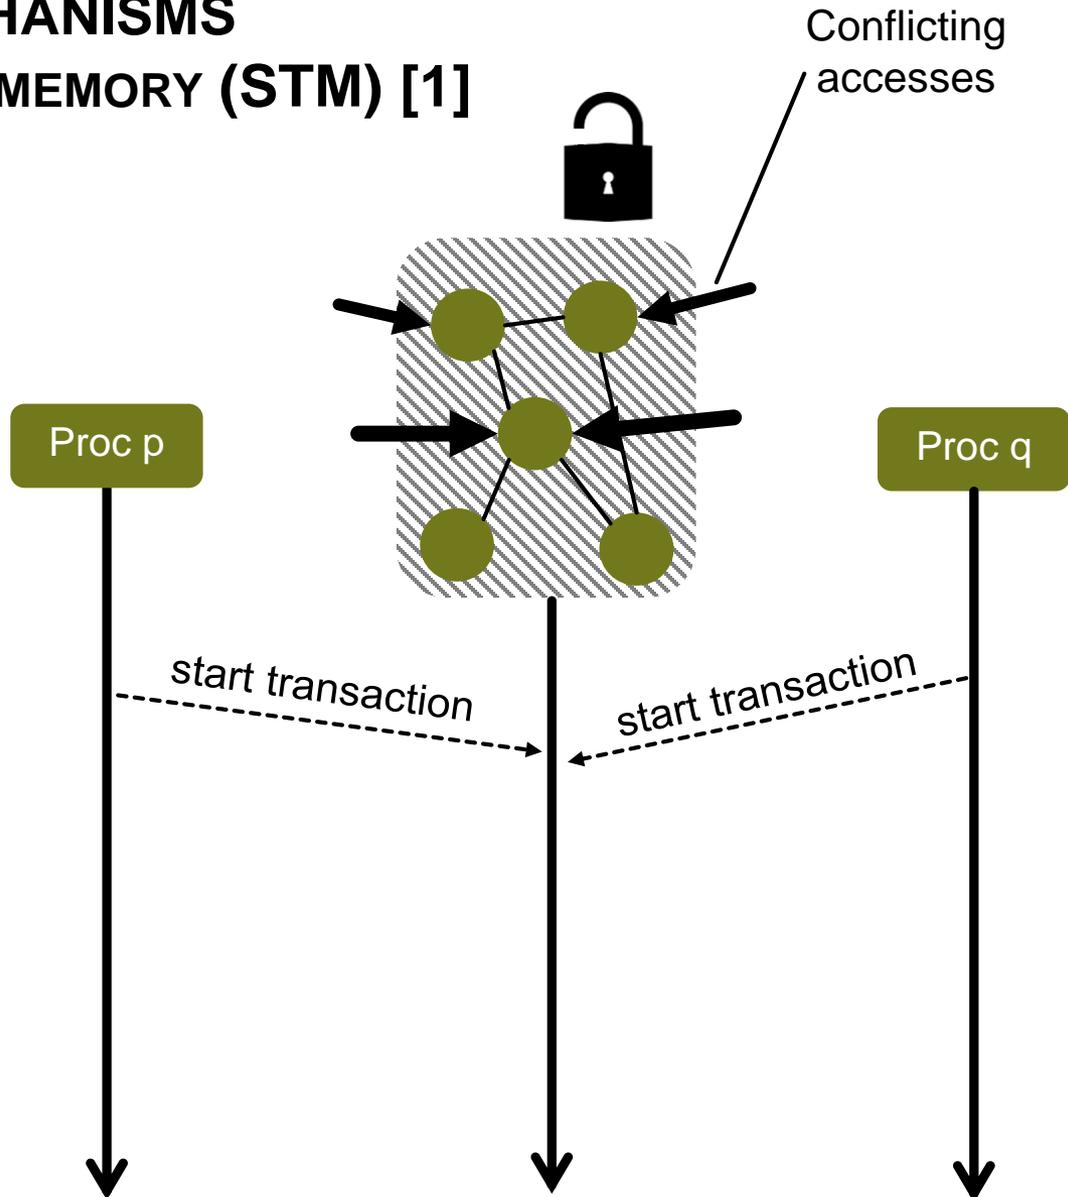
Conflicts solved with rollbacks and/or serialization.



# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]

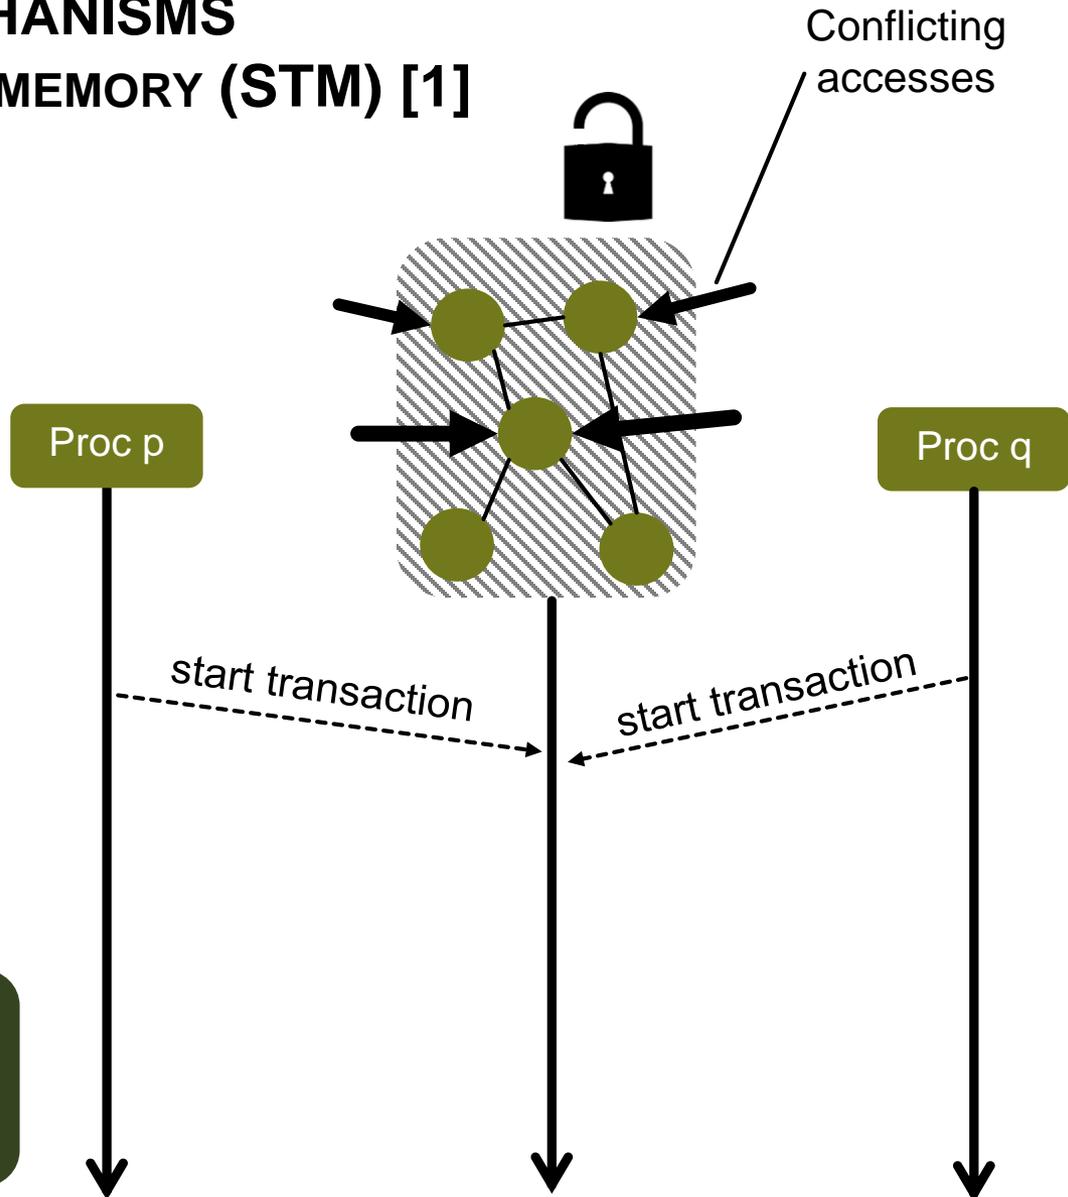
Conflicts solved with rollbacks and/or serialization.



# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]

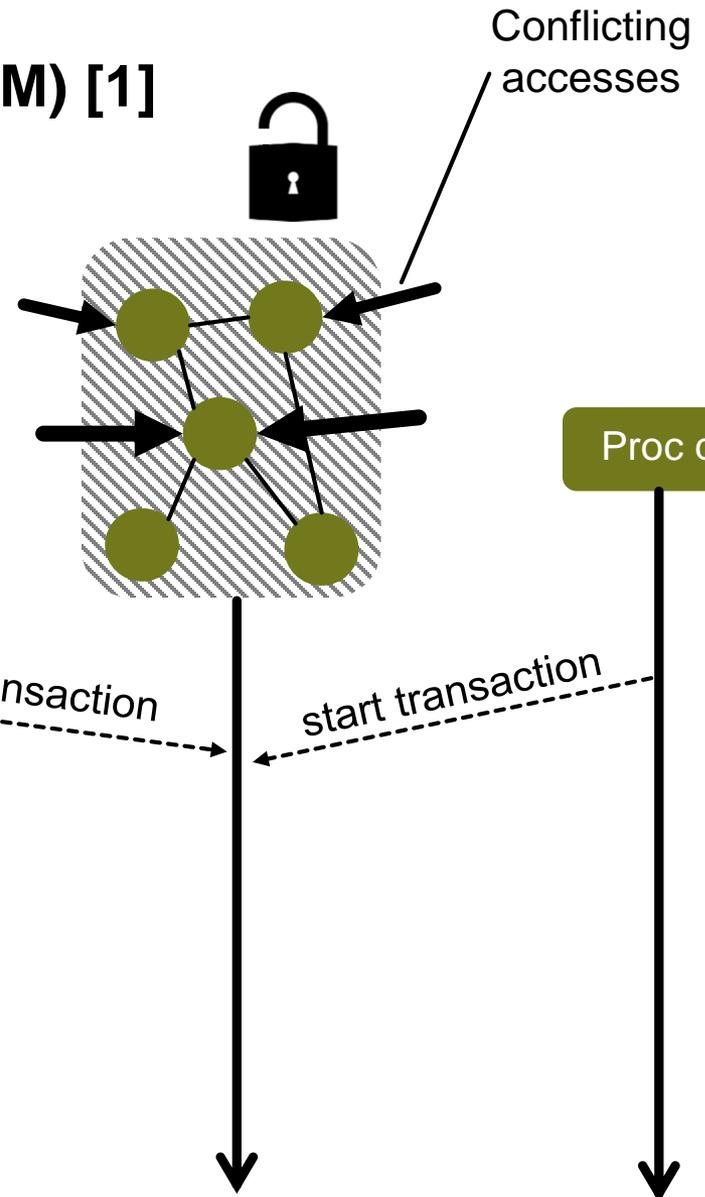
Conflicts solved with rollbacks and/or serialization.



# SYNCHRONIZATION MECHANISMS

## SOFTWARE TRANSACTIONAL MEMORY (STM) [1]

Conflicts solved with rollbacks and/or serialization.



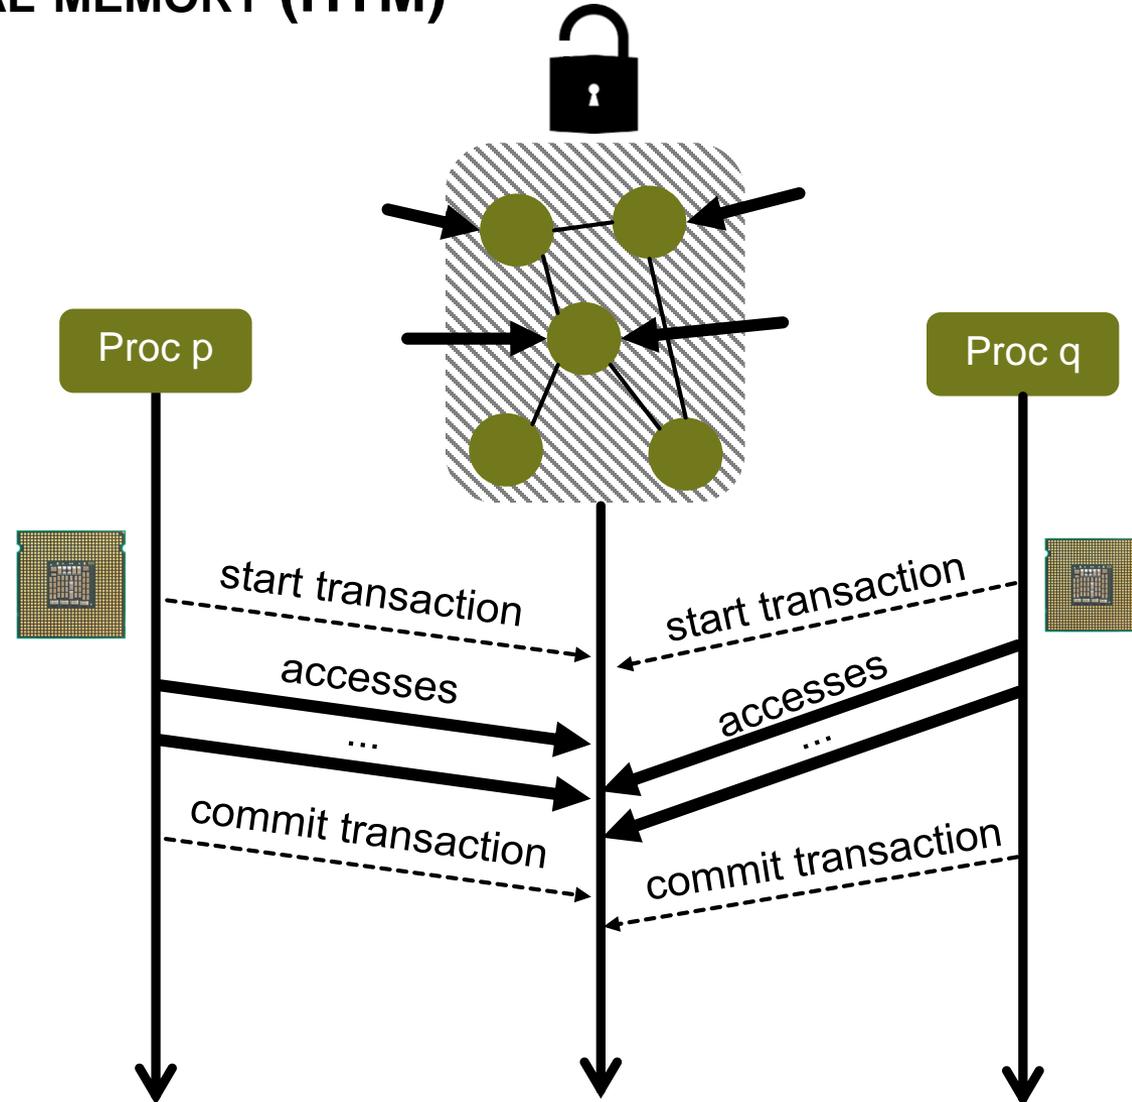
✘ Software overheads

✔ Simple protocols

# SYNCHRONIZATION MECHANISMS

## HARDWARE TRANSACTIONAL MEMORY (HTM)

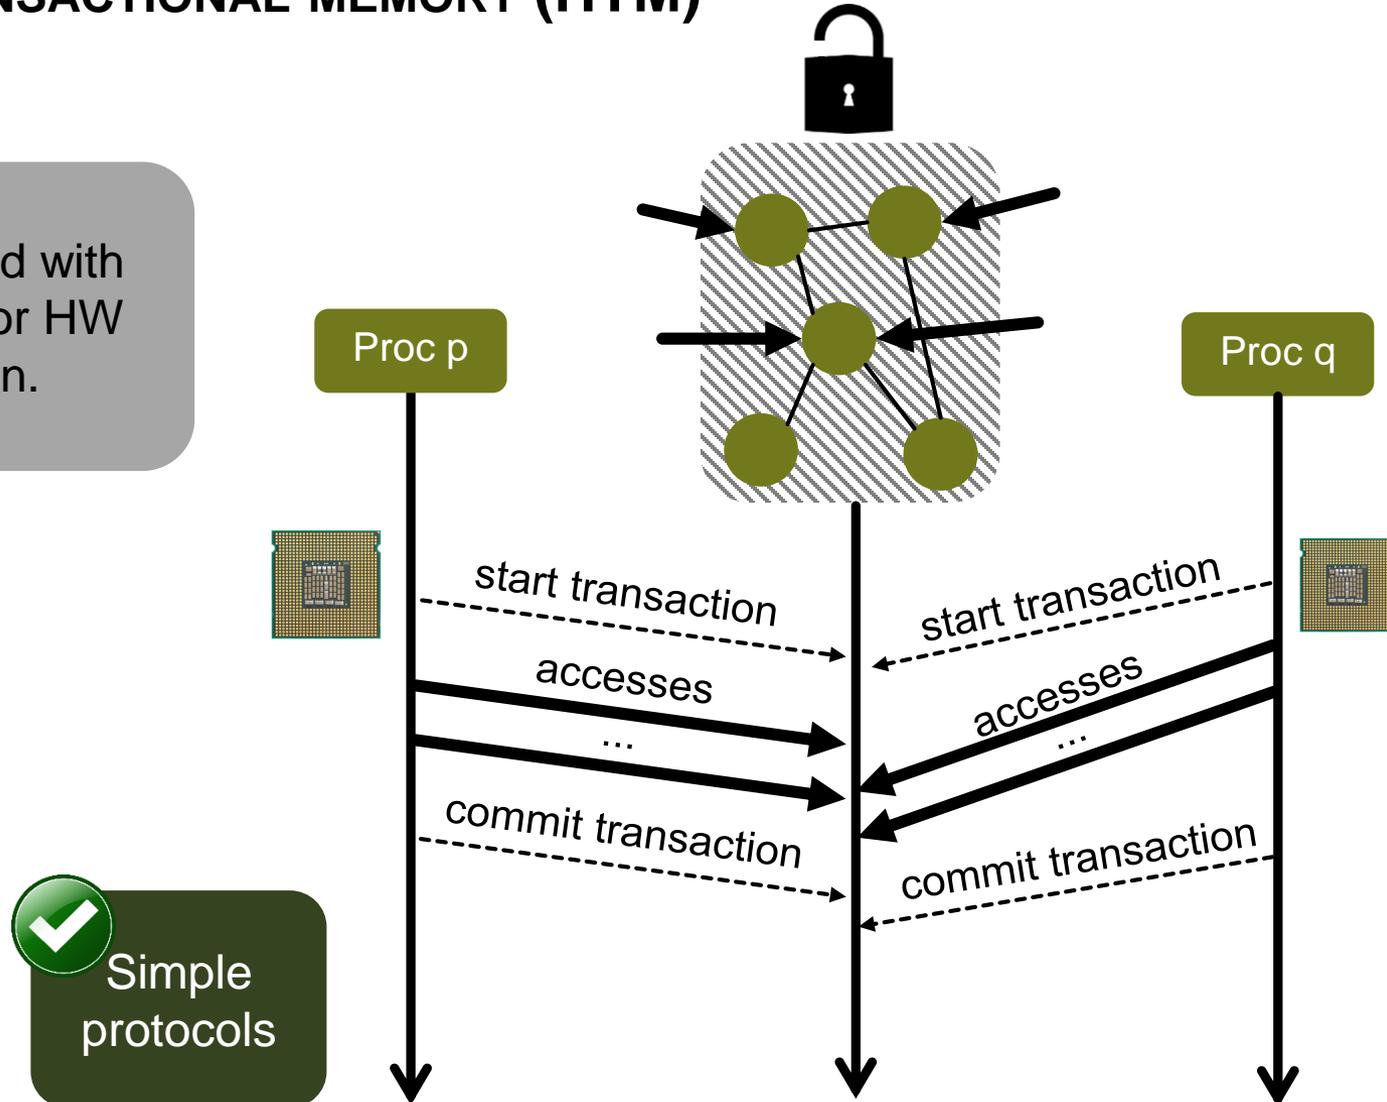
Conflicts solved with rollbacks and/or HW serialization.



# SYNCHRONIZATION MECHANISMS

## HARDWARE TRANSACTIONAL MEMORY (HTM)

Conflicts solved with rollbacks and/or HW serialization.



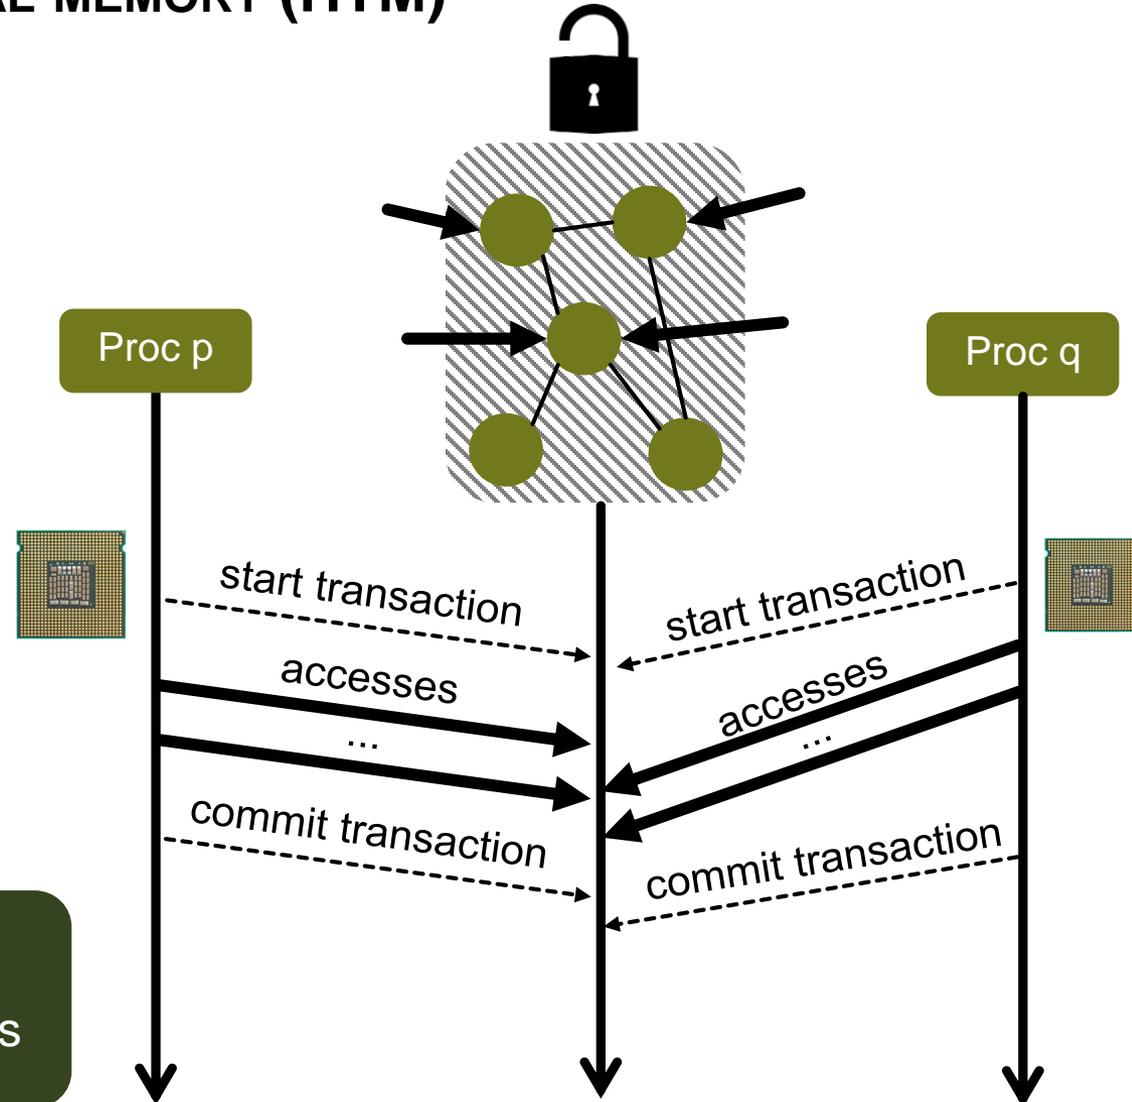
# SYNCHRONIZATION MECHANISMS

## HARDWARE TRANSACTIONAL MEMORY (HTM)

Conflicts solved with rollbacks and/or HW serialization.

? High performance?  
For graphs?

✓ Simple protocols



# HARDWARE TRANSACTIONAL MEMORY

# HARDWARE TRANSACTIONAL MEMORY



Vega

# HARDWARE TRANSACTIONAL MEMORY



Rock

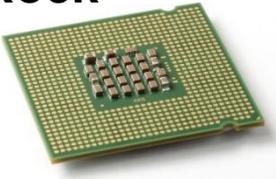


Vega

# HARDWARE TRANSACTIONAL MEMORY



Rock



Vega



BlueGene/Q



POWER8

# HARDWARE TRANSACTIONAL MEMORY



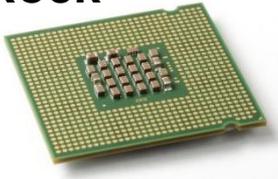
Vega



BlueGene/Q



Rock



Haswell



POWER8

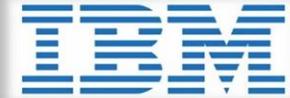
# HARDWARE TRANSACTIONAL MEMORY



Rock



Vega



BlueGene/Q

Haswell

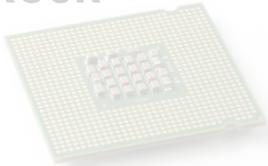


POWER8

# HARDWARE TRANSACTIONAL MEMORY



Rock



Vega



They offer programmability, how about performance?



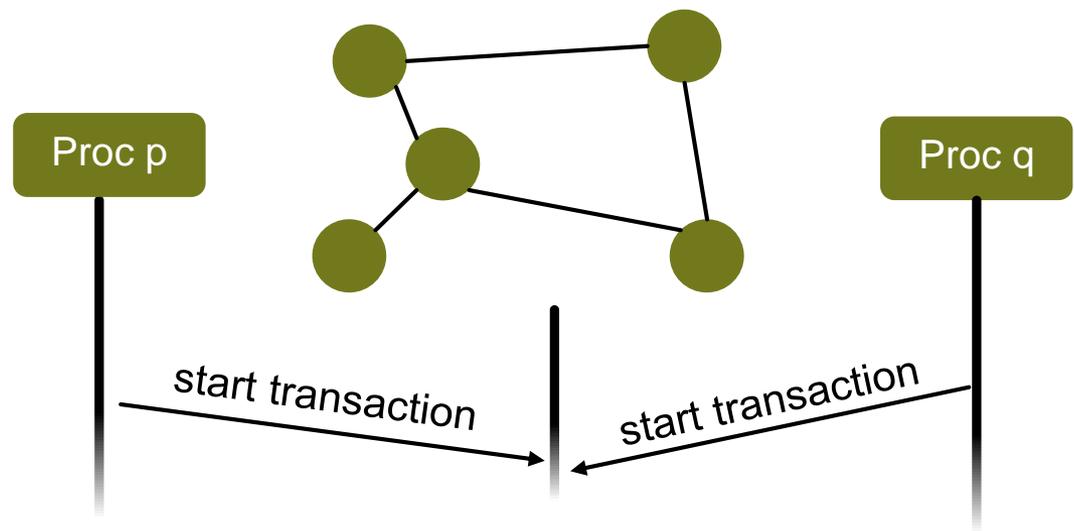
BlueGene/Q

Haswell



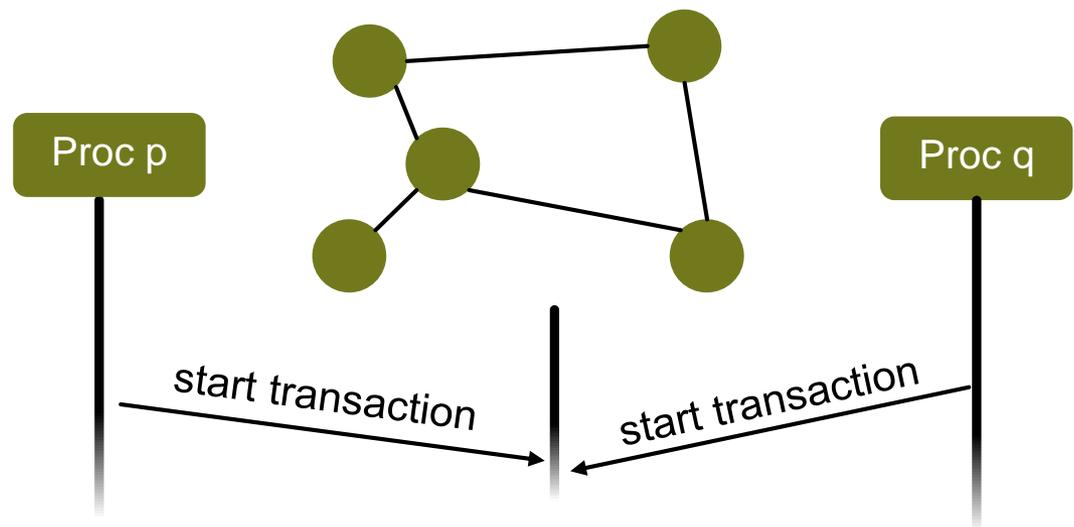
POWER8

# SHARED- & DISTRIBUTED-MEMORY MACHINES



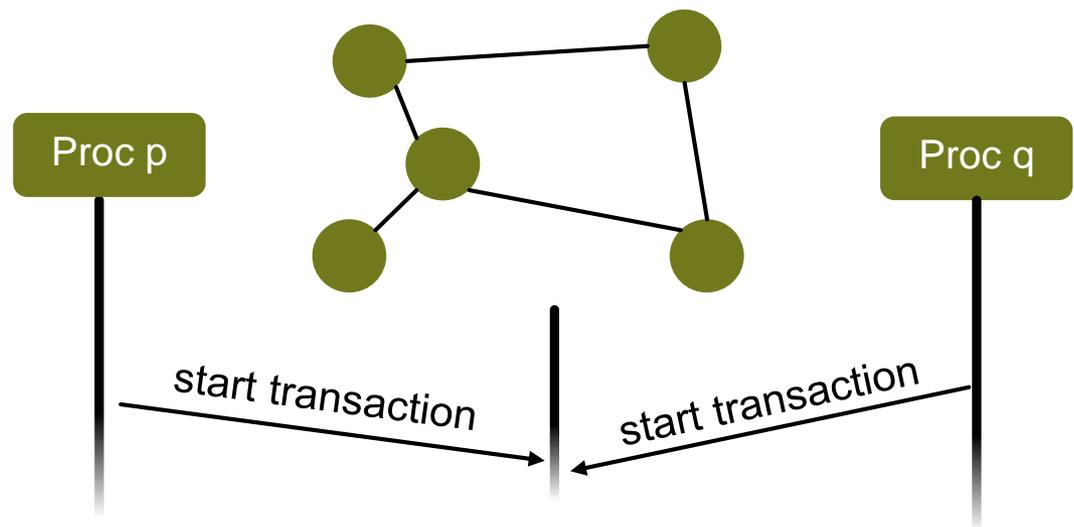
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains



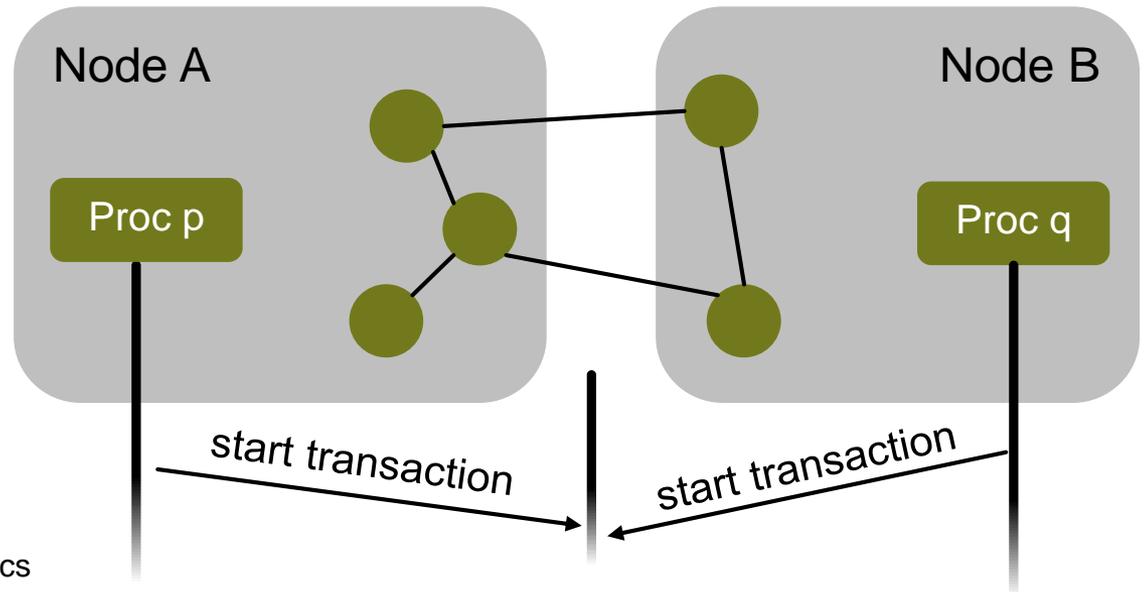
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]



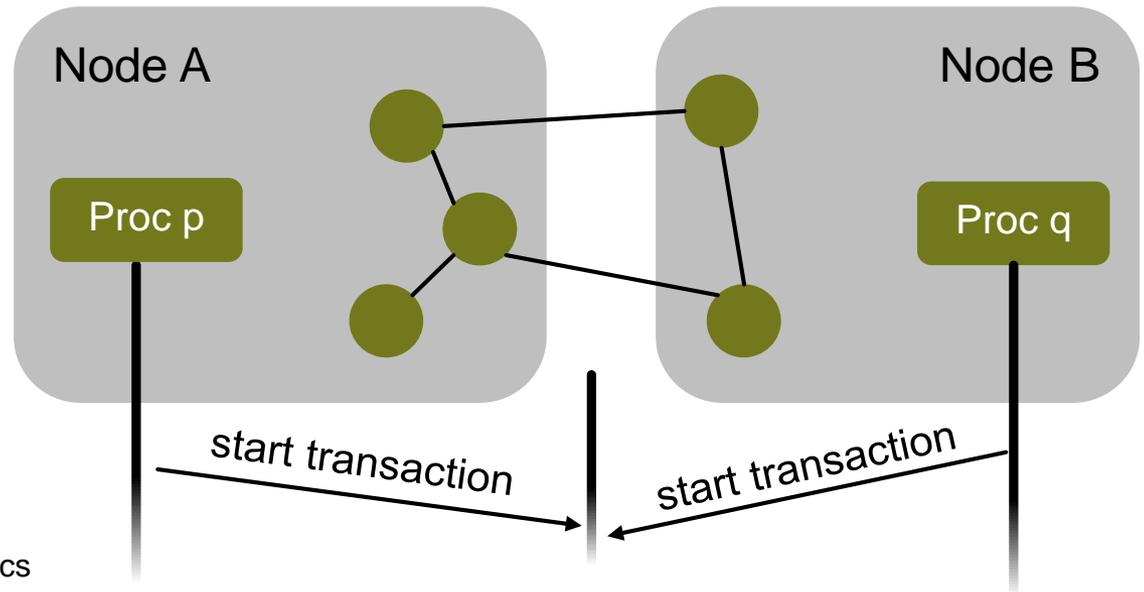
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]



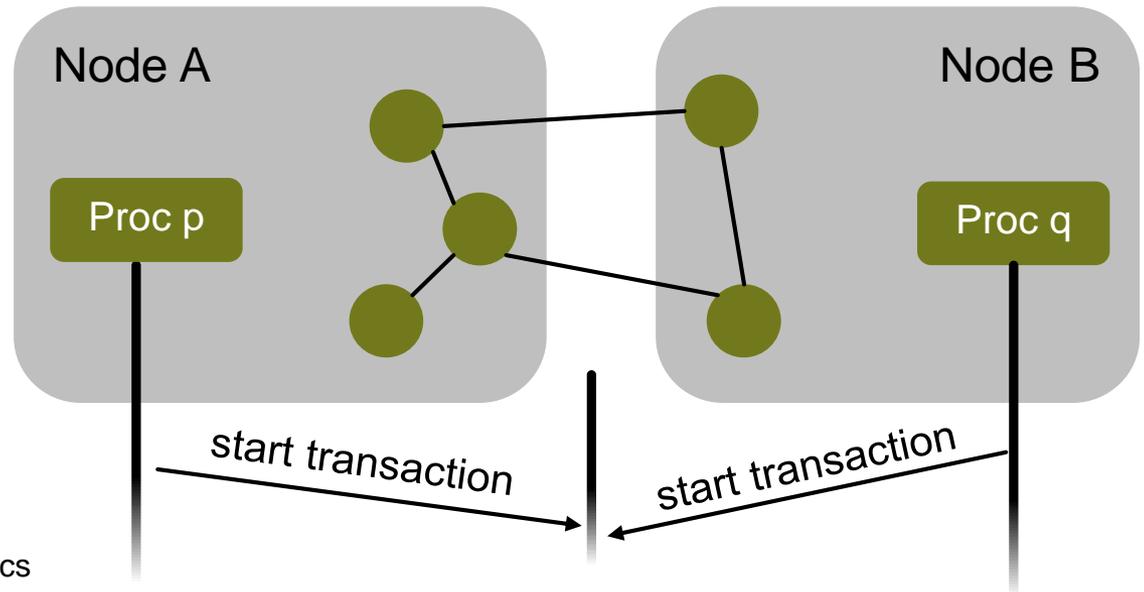
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]
- However, some do not:



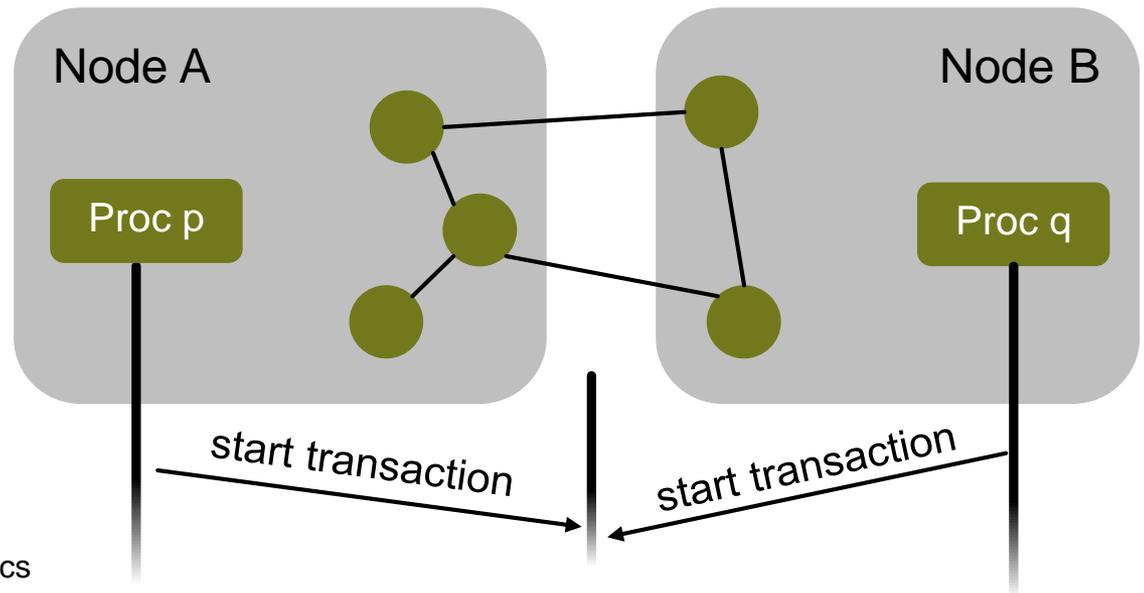
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]
- However, some do not:
  - Very large instances



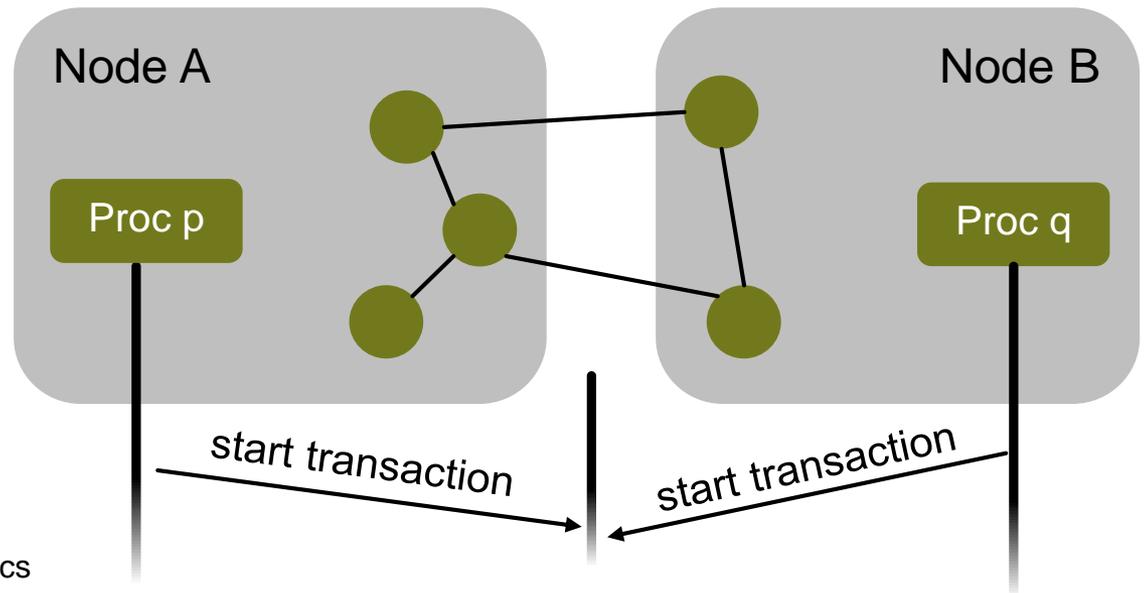
# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]
- However, some do not:
  - Very large instances
  - Rich vertex/edge data



# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]
- However, some do not:
  - Very large instances
  - Rich vertex/edge data
- Fat nodes with lots of RAM still expensive (\$35K for a machine with 1TB of RAM [1])

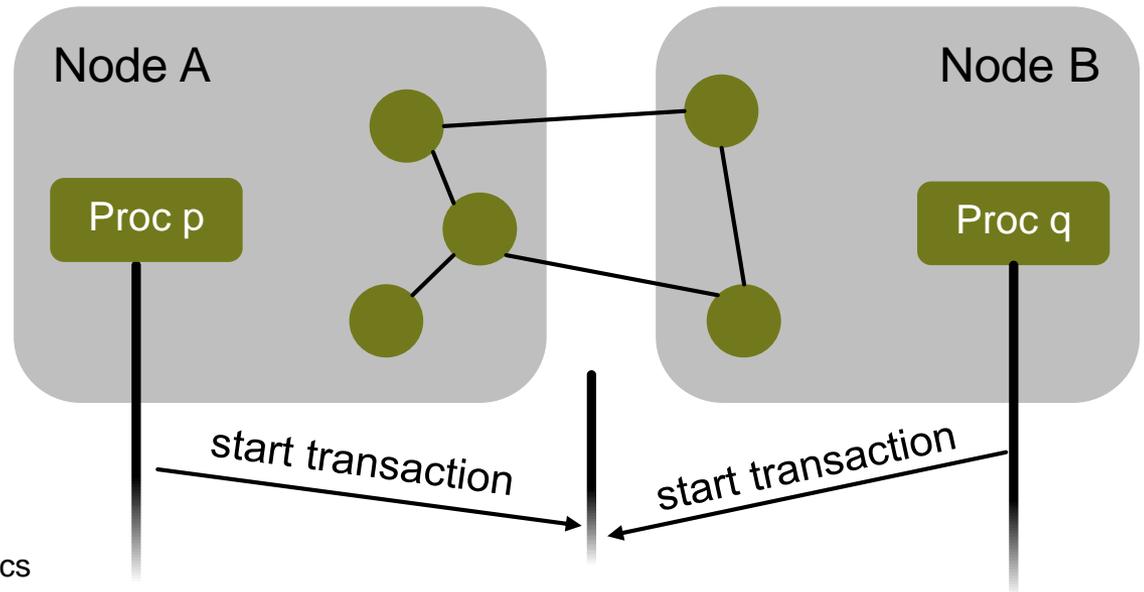


# SHARED- & DISTRIBUTED-MEMORY MACHINES

- HTM works fine for single shared-memory domains
  - Most graphs fit in such machines [1]
- However, some do not:
  - Very large instances
  - Rich vertex/edge data
- Fat nodes with lots of RAM still expensive (\$35K for a machine with 1TB of RAM [1])



How to apply  
HTM in such a  
setting?



# OVERVIEW OF OUR RESEARCH

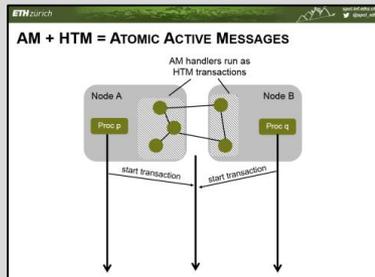
# OVERVIEW OF OUR RESEARCH

**HTM for graphs in SM & DM environments**



# OVERVIEW OF OUR RESEARCH

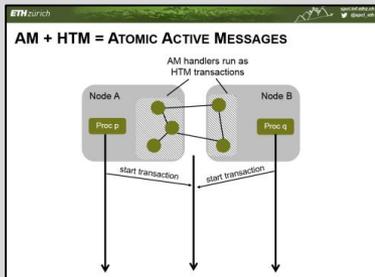
## HTM for graphs in SM & DM environments



HTM + Active Messages  
= Atomic Active Messages

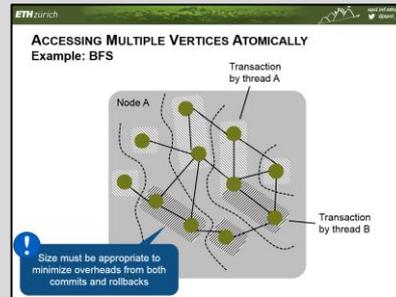
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



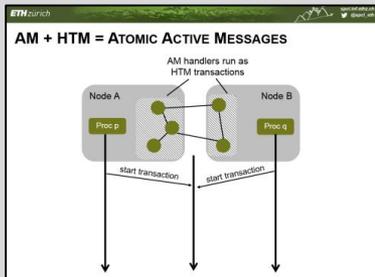
HTM + Active Messages  
= Atomic Active Messages

## Coarsening & coalescing



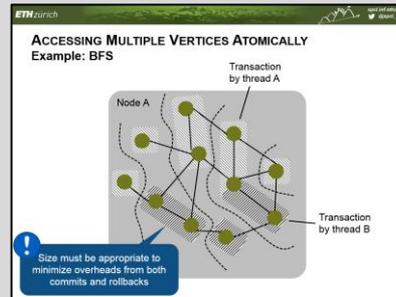
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



HTM + Active Messages  
= Atomic Active Messages

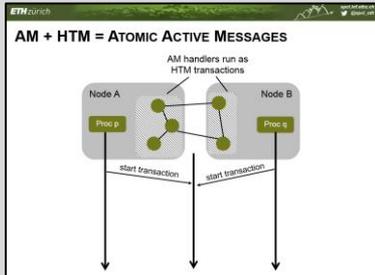
## Coarsening & coalescing



## Performance Modeling & Analysis

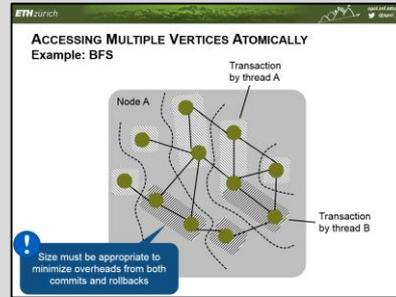
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



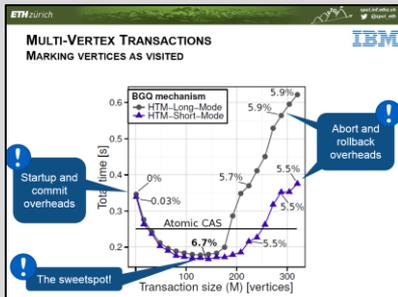
HTM + Active Messages  
= Atomic Active Messages

## Coarsening & coalescing



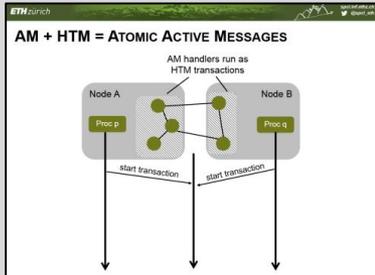
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis



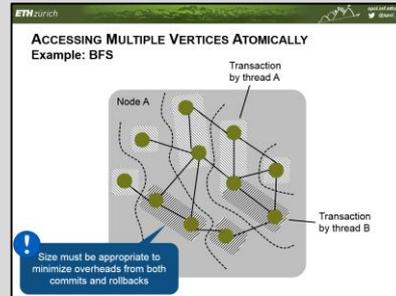
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



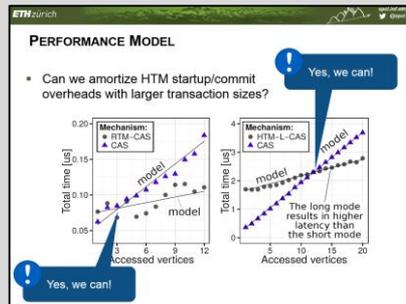
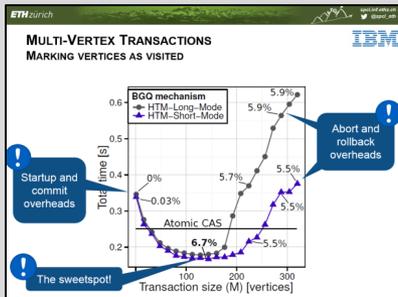
HTM + Active Messages  
= Atomic Active Messages

## Coarsening & coalescing



## Performance Modeling & Analysis

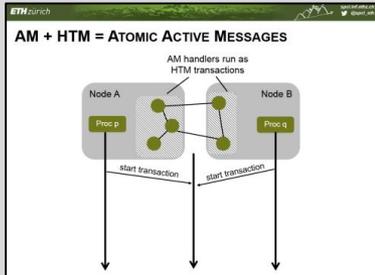
### Haswell & BG/Q Analysis



## Performance model

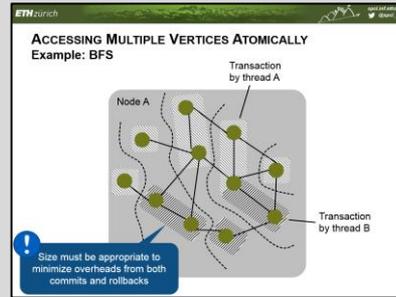
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



HTM + Active Messages  
= Atomic Active Messages

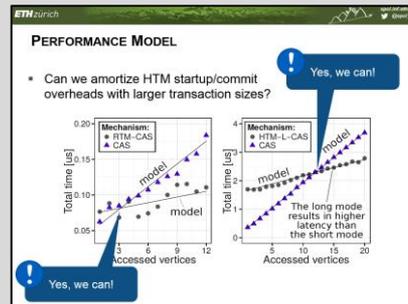
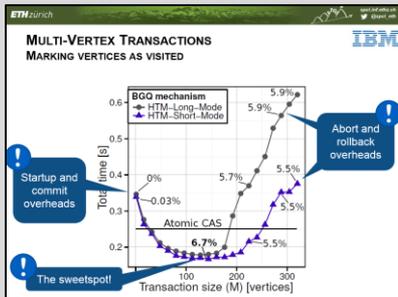
## Coarsening & coalescing



## Evaluation

## Performance Modeling & Analysis

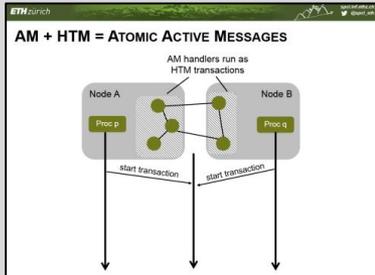
### Haswell & BG/Q Analysis



## Performance model

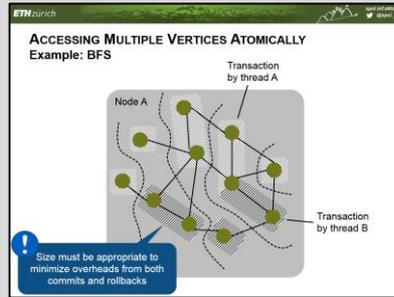
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



HTM + Active Messages  
= Atomic Active Messages

## Coarsening & coalescing

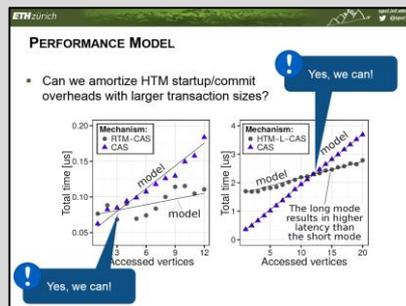
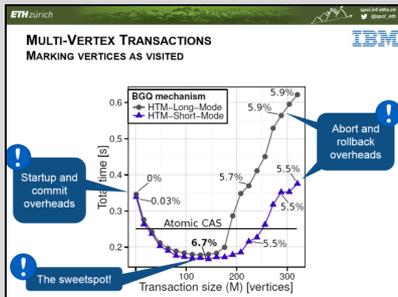


## Evaluation

Considered engines and graphs

## Performance Modeling & Analysis

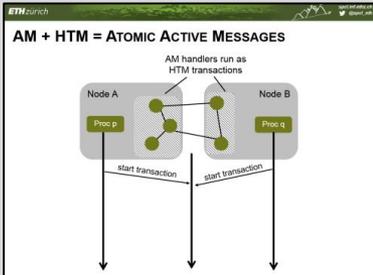
### Haswell & BG/Q Analysis



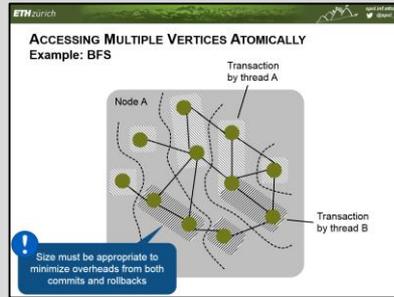
Performance model

# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



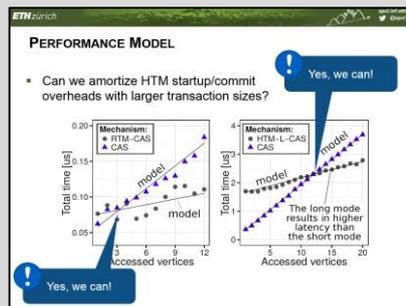
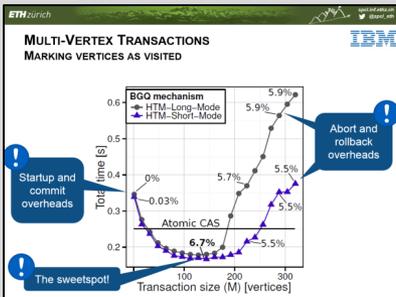
## Coarsening & coalescing



HTM + Active Messages  
= Atomic Active Messages

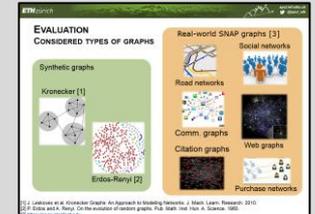
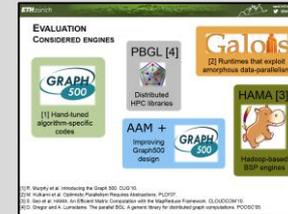
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

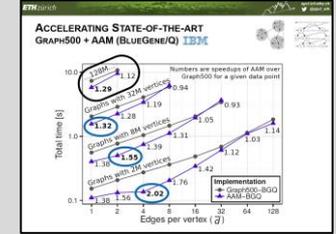
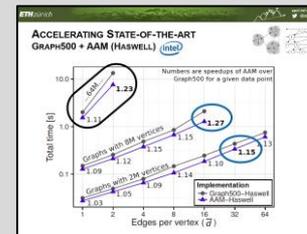


## Performance model

## Evaluation



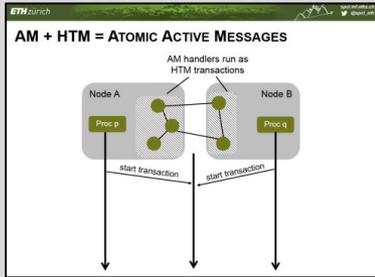
## Considered engines and graphs



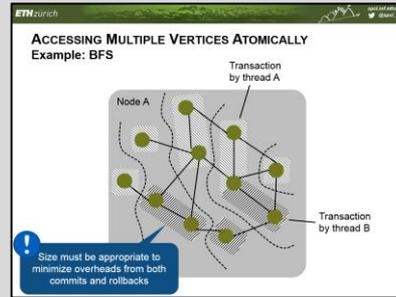
## Accelerating state-of-the-art

# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



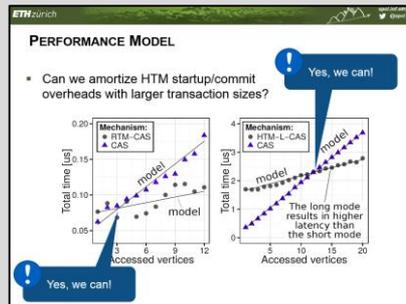
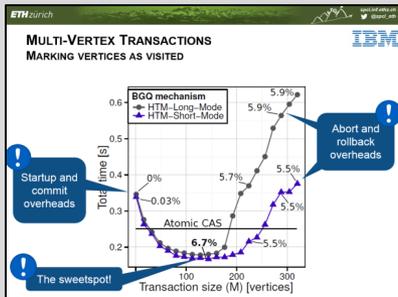
## Coarsening & coalescing



HTM + Active Messages  
= Atomic Active Messages

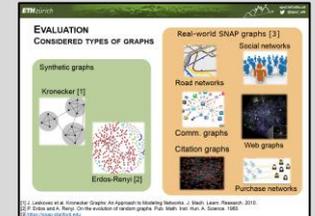
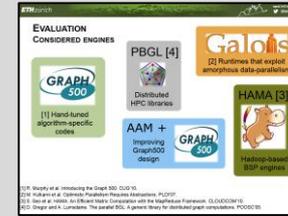
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

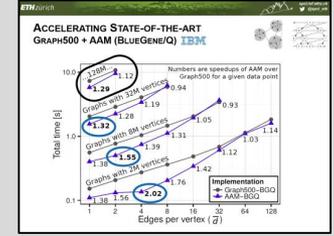
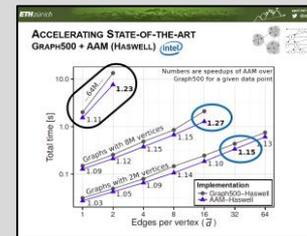


## Performance model

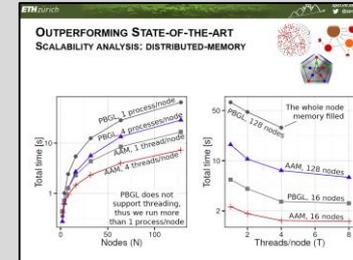
## Evaluation



## Considered engines and graphs



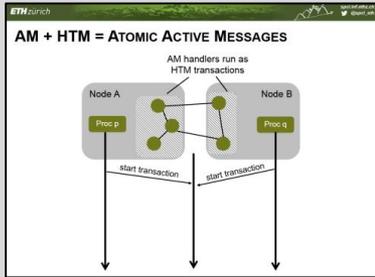
## Accelerating state-of-the-art



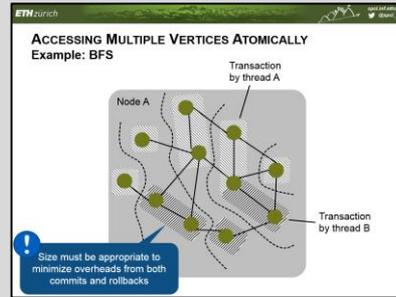
## Scalability

# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



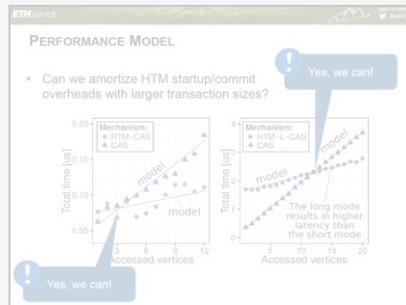
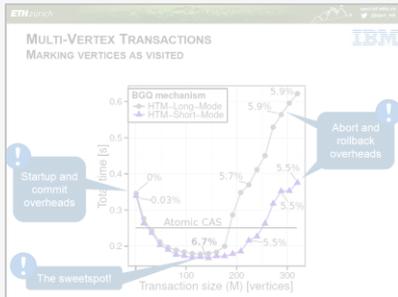
## Coarsening & coalescing



HTM + Active Messages = Atomic Active Messages

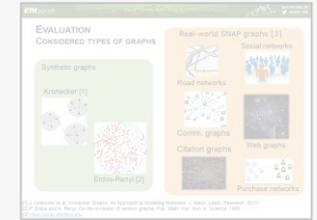
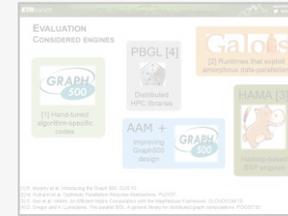
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

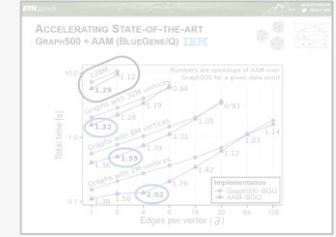
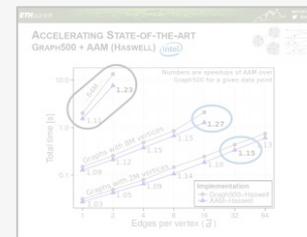


Performance model

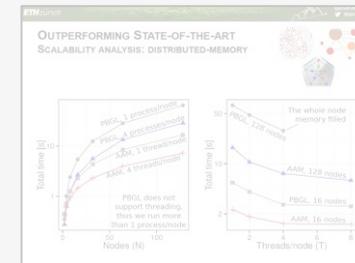
## Evaluation



Considered engines and graphs



Accelerating state-of-the-art



Scalability

# ACTIVE MESSAGES (AM)

# ACTIVE MESSAGES (AM)

Process p

Process q

# ACTIVE MESSAGES (AM)

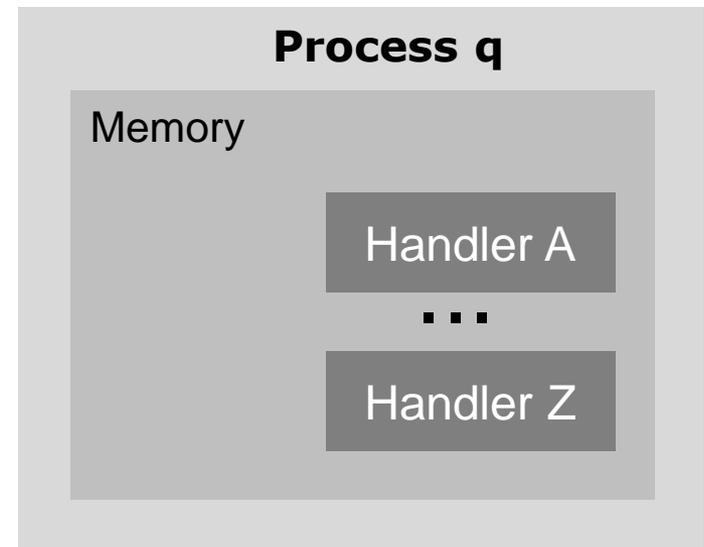
Process p

Process q

Memory

# ACTIVE MESSAGES (AM)

Process p



# ACTIVE MESSAGES (AM)

Process p

Process q

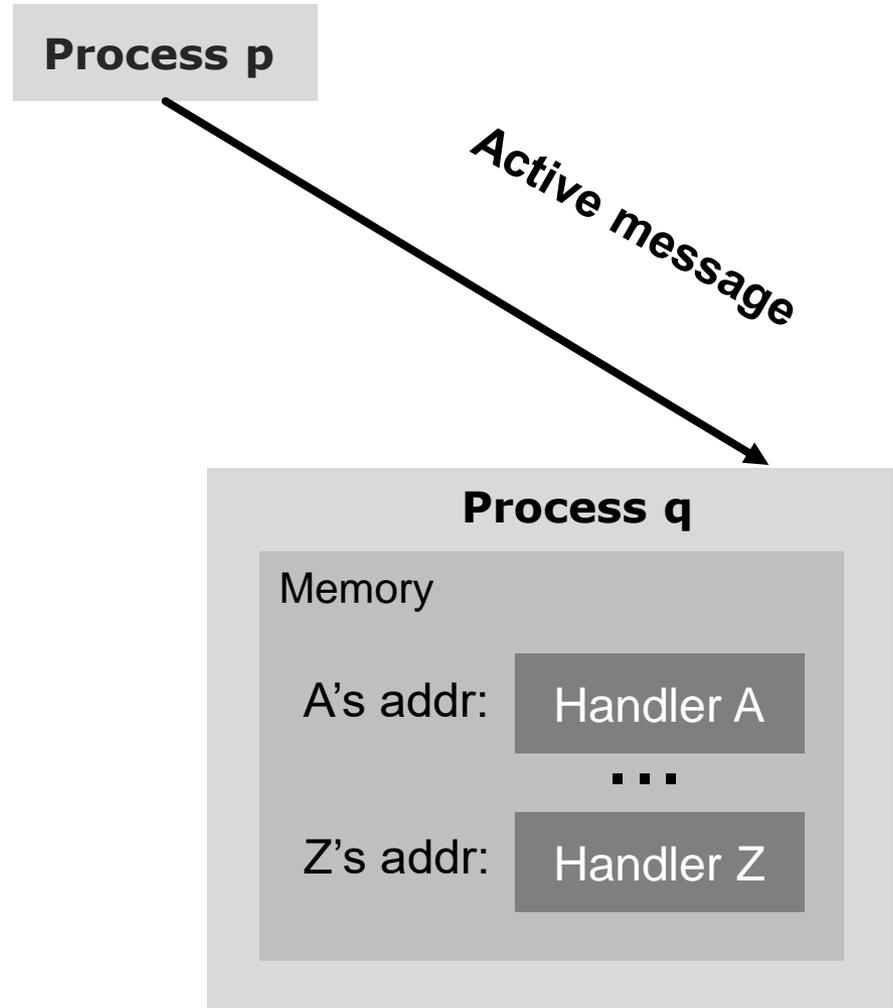
Memory

A's addr: Handler A

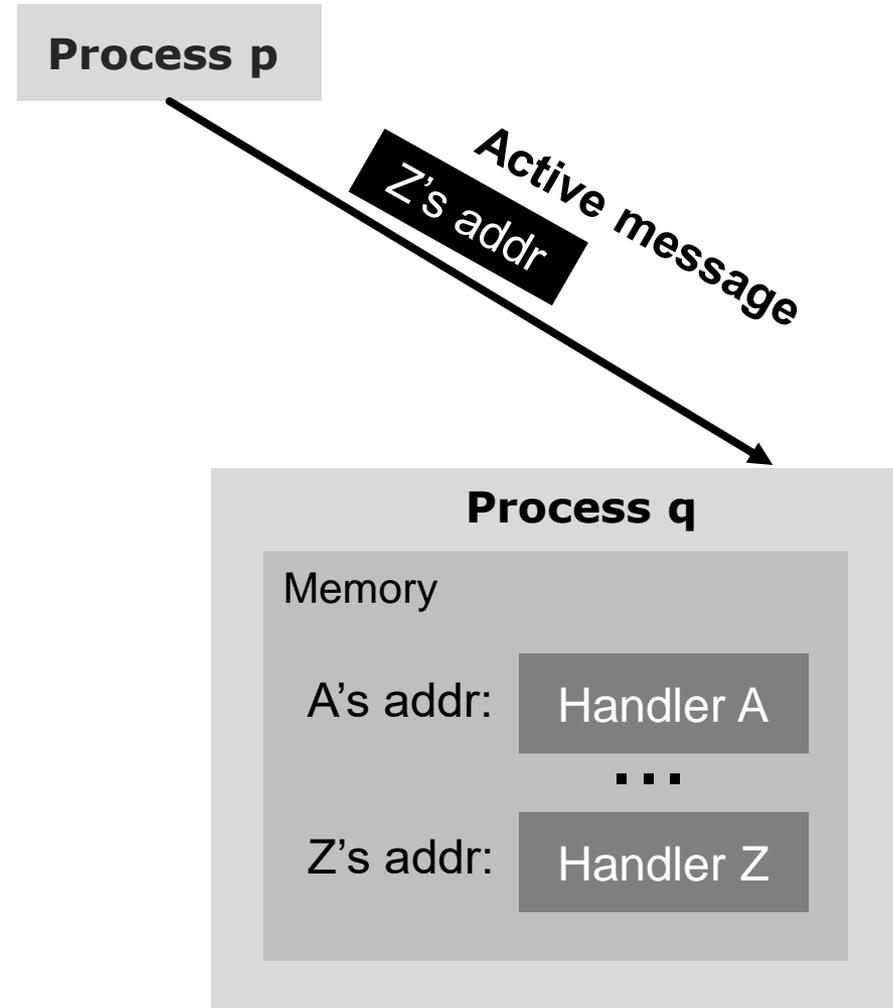
...

Z's addr: Handler Z

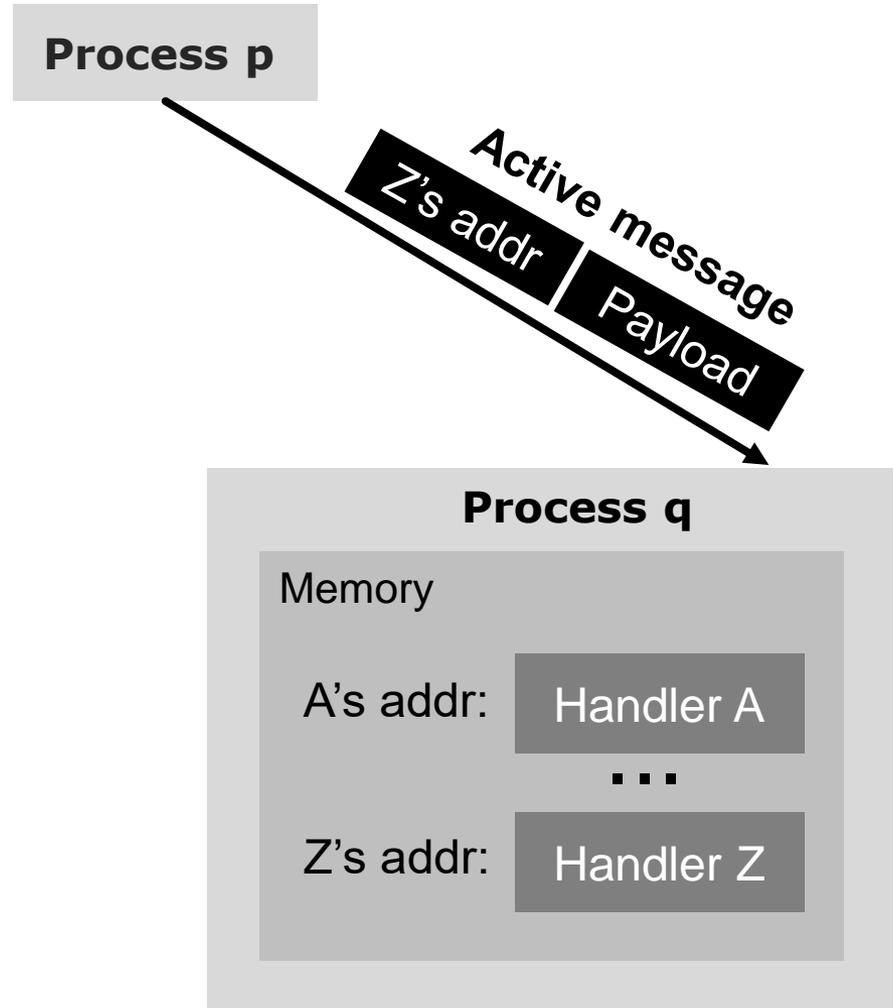
# ACTIVE MESSAGES (AM)



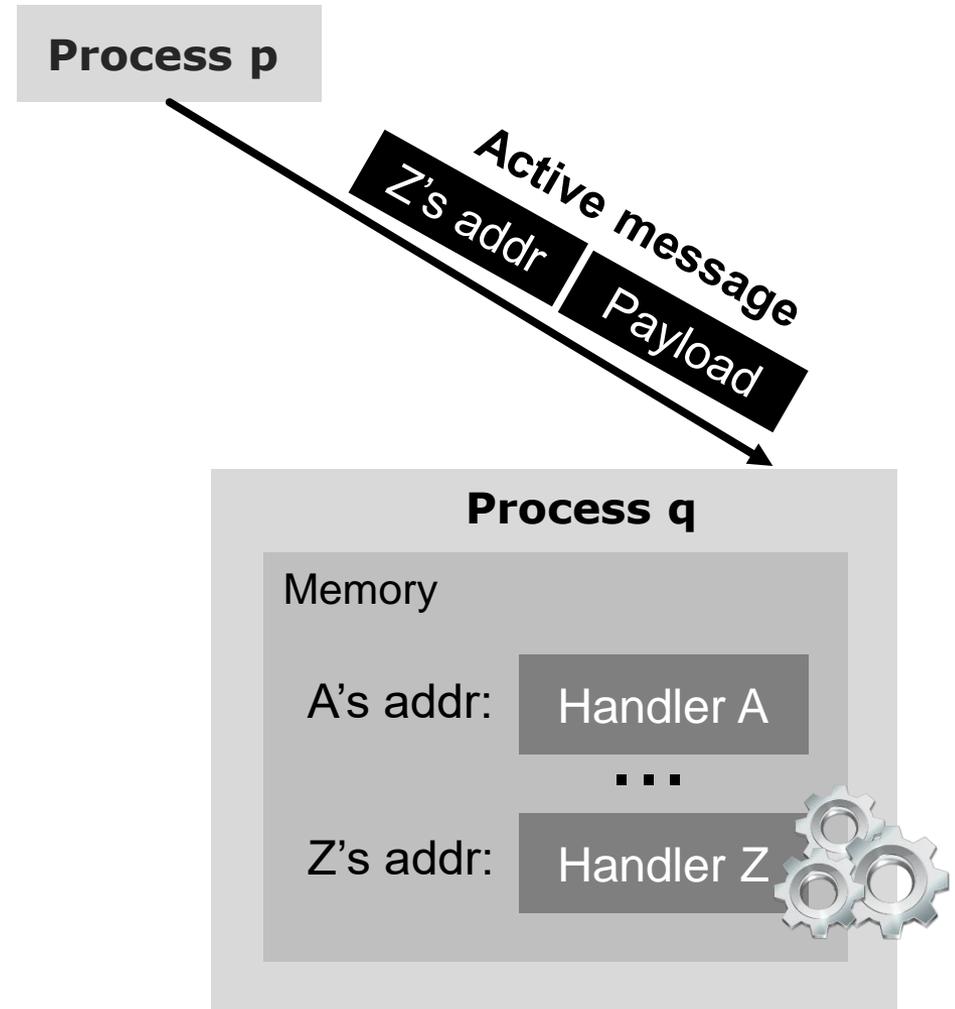
# ACTIVE MESSAGES (AM)



# ACTIVE MESSAGES (AM)



# ACTIVE MESSAGES (AM)



# ACTIVE MESSAGES (AM)



Process p

Active message  
Z's addr | Payload

Process q

Memory

A's addr: Handler A

...

Z's addr: Handler Z



# ACTIVE MESSAGES (AM)





AM++ [1]

GASNet [2]

Process p

Active message  
Z's addr | Payload

Process q

Memory

A's addr: Handler A

...

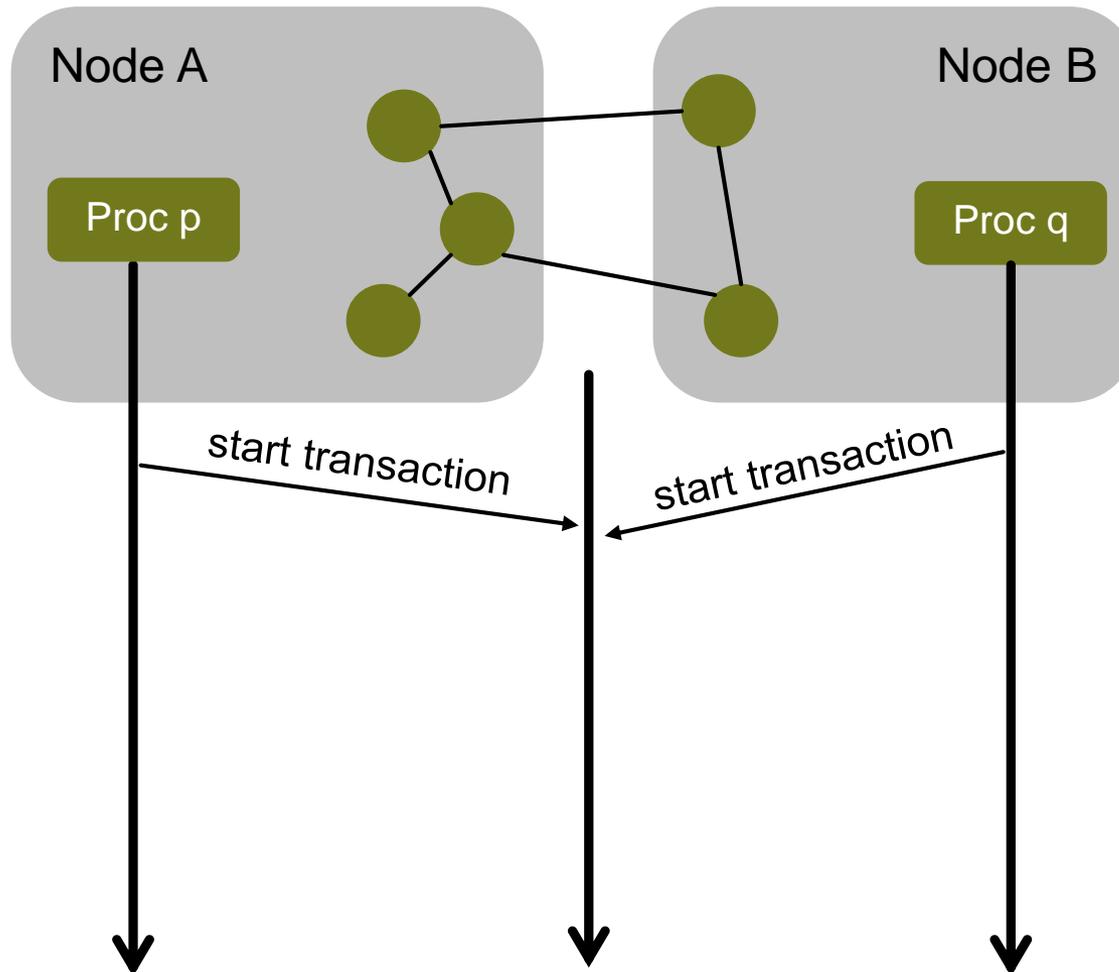
Z's addr: Handler Z



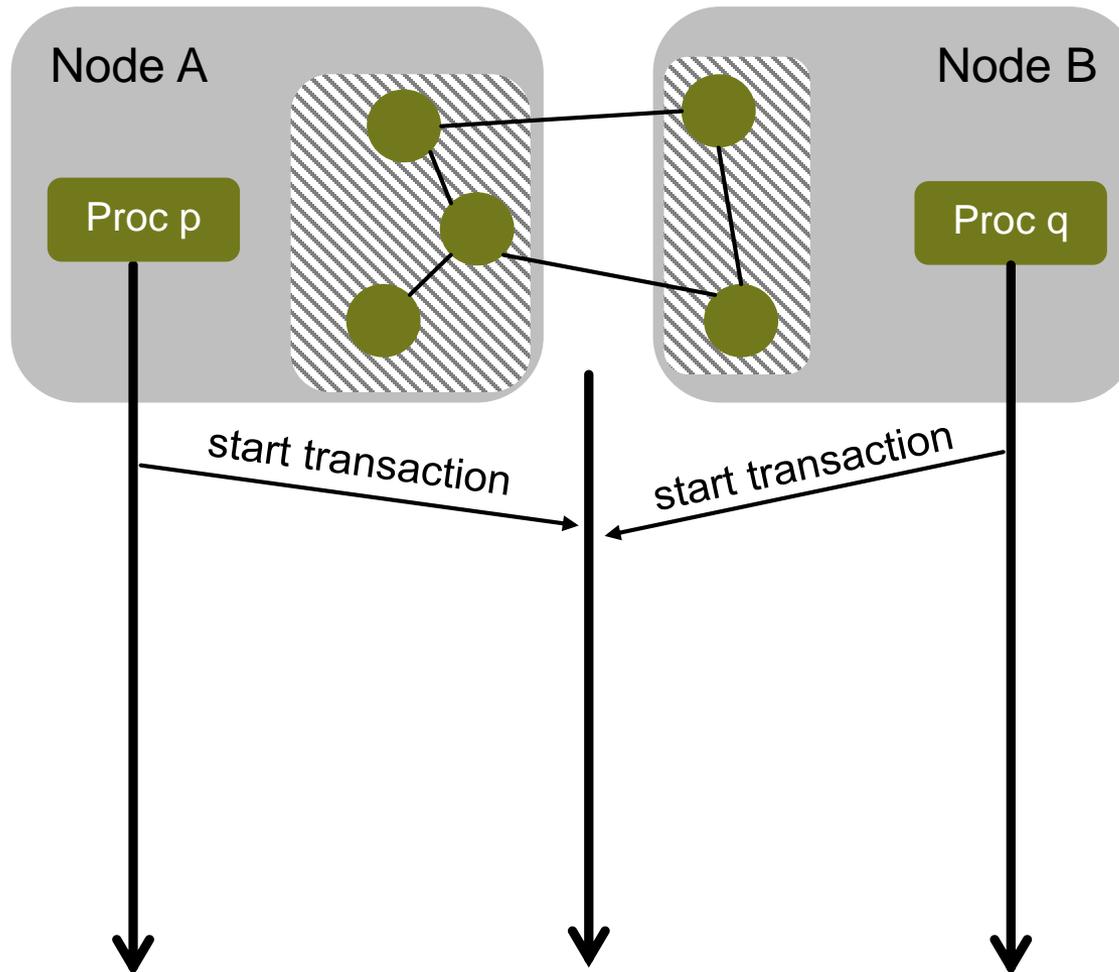
[1] J. J. Willcock et al. AM++: A generalized active message framework. PACT'10.

[2] D. Bonachea, GASNet Specification, v1.1. Berkeley Technical Report. 2002.

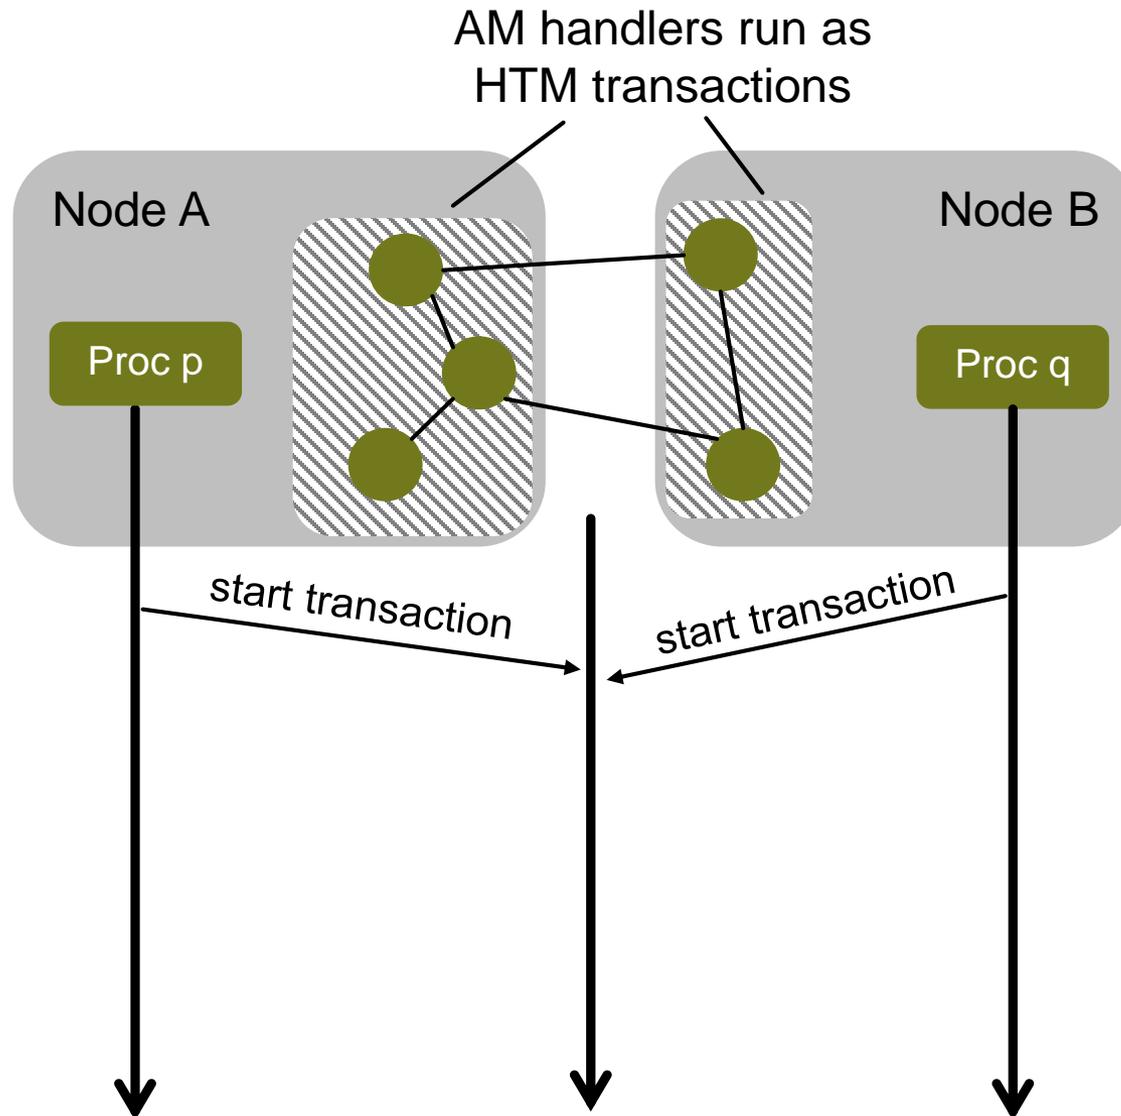
# AM + HTM = ...



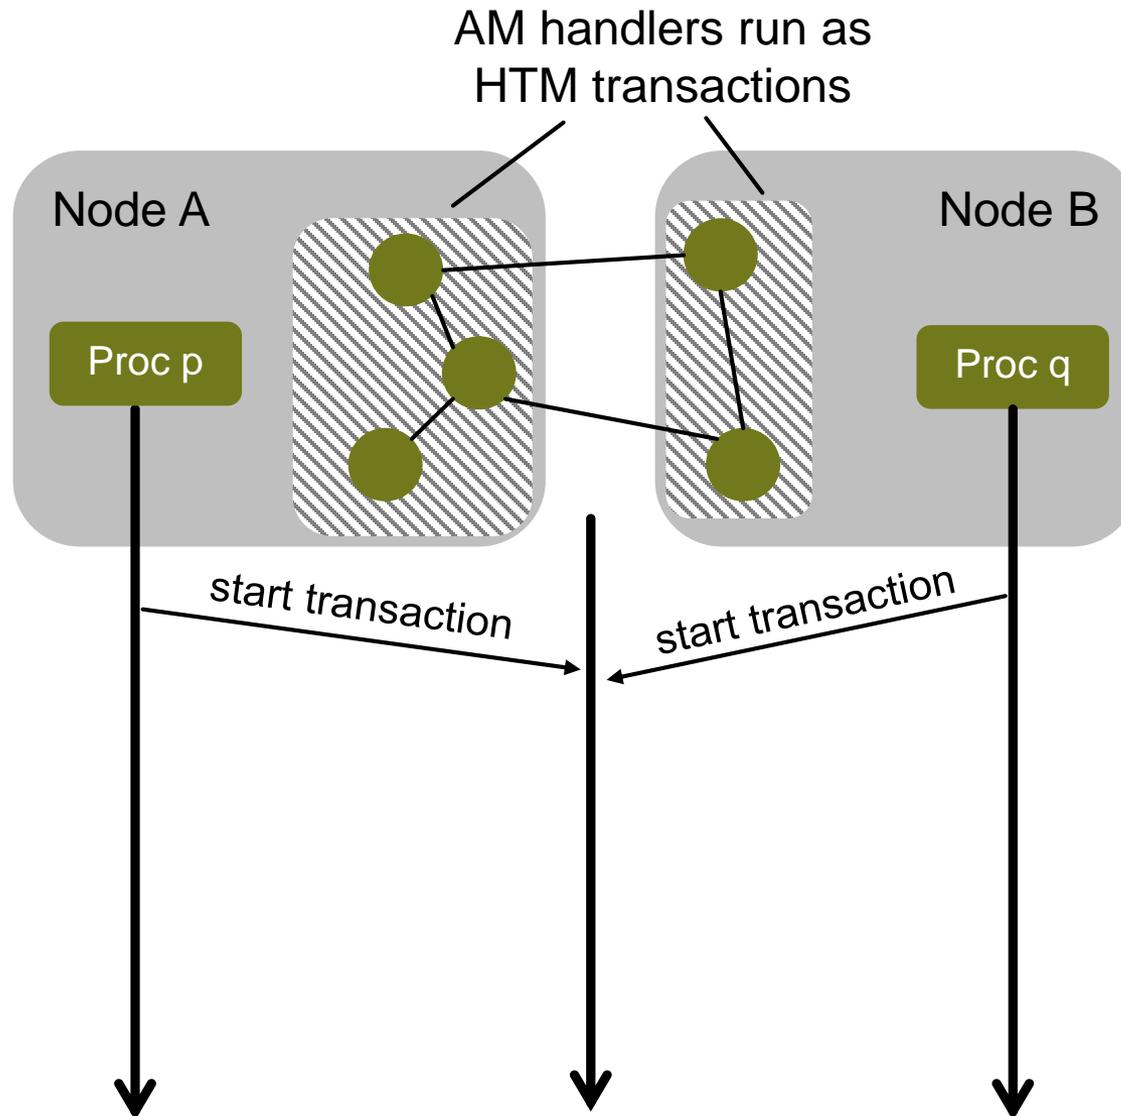
# AM + HTM = ...



# AM + HTM = ...

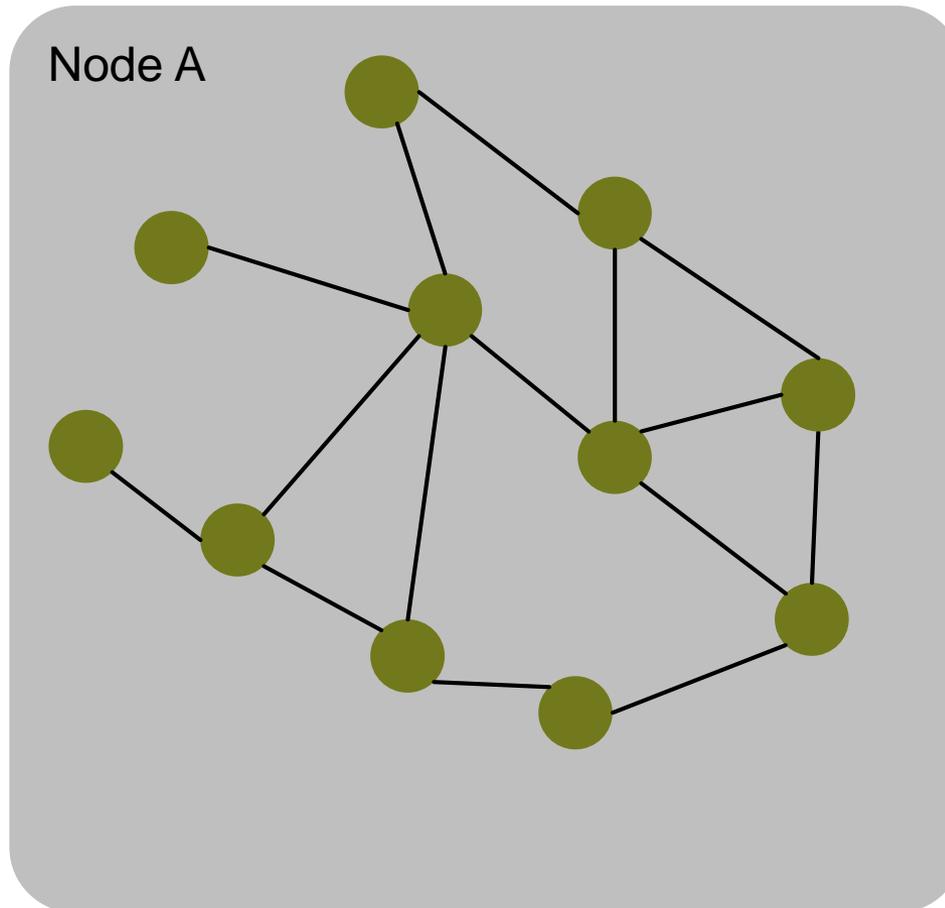


# AM + HTM = ATOMIC ACTIVE MESSAGES



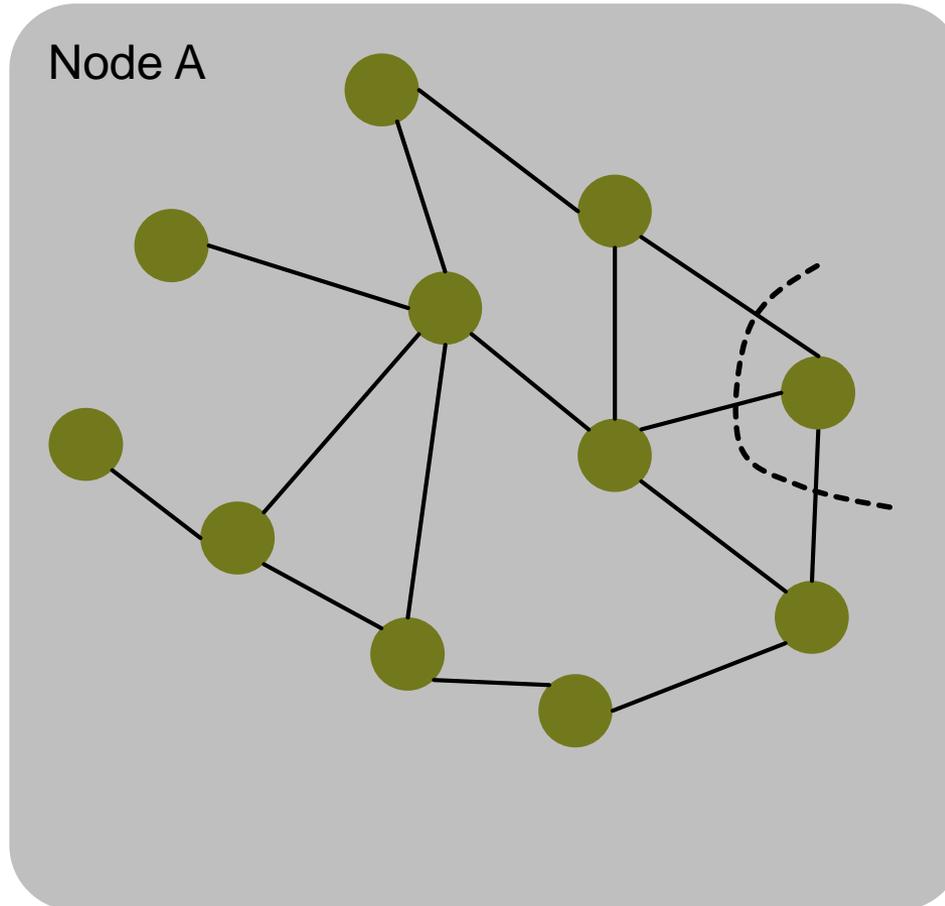
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



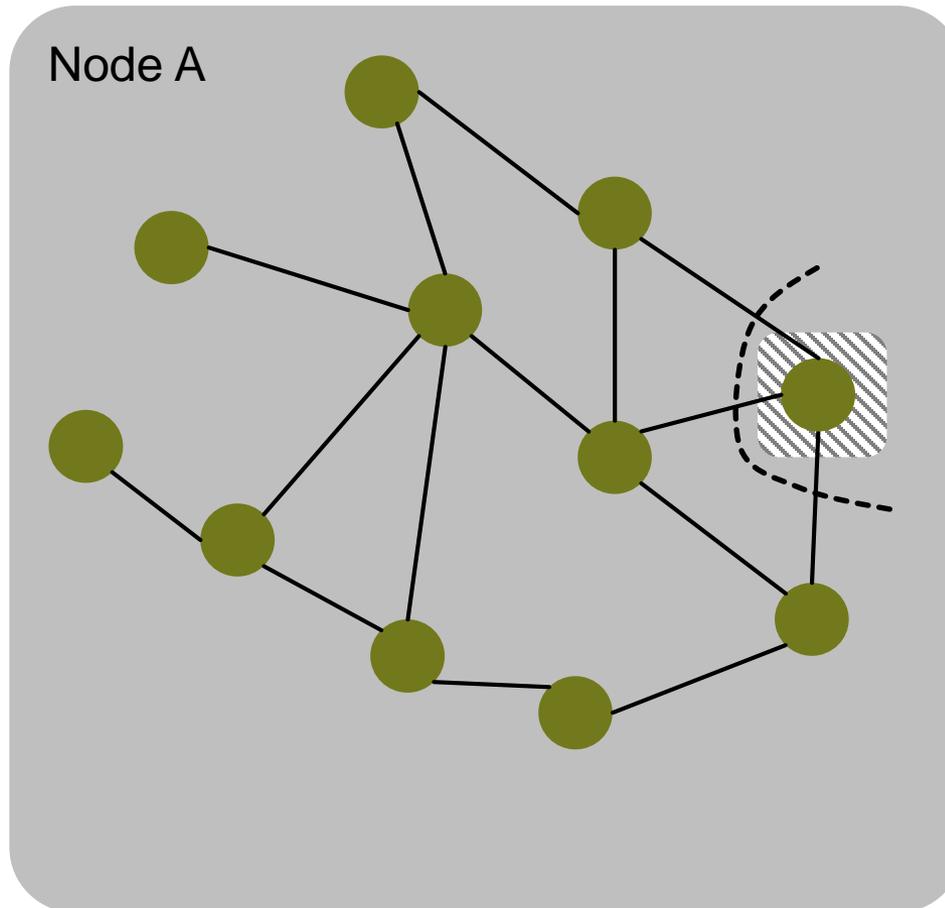
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



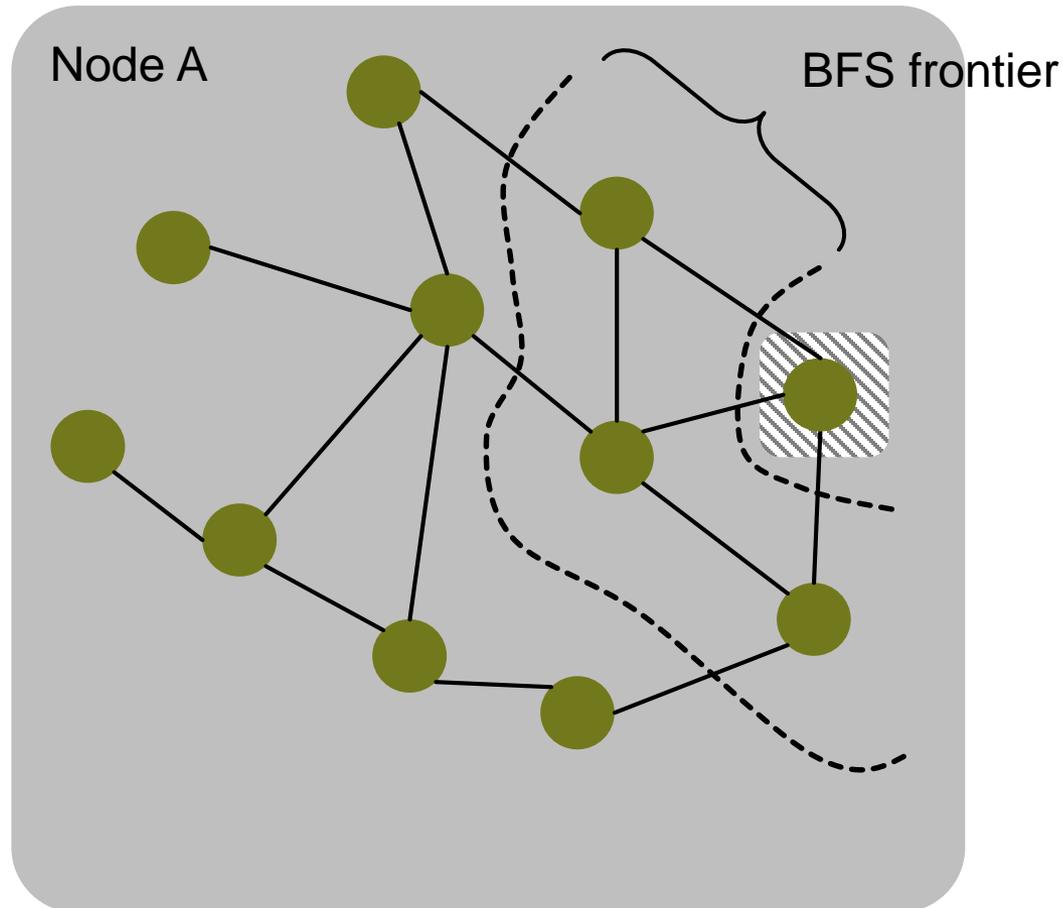
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



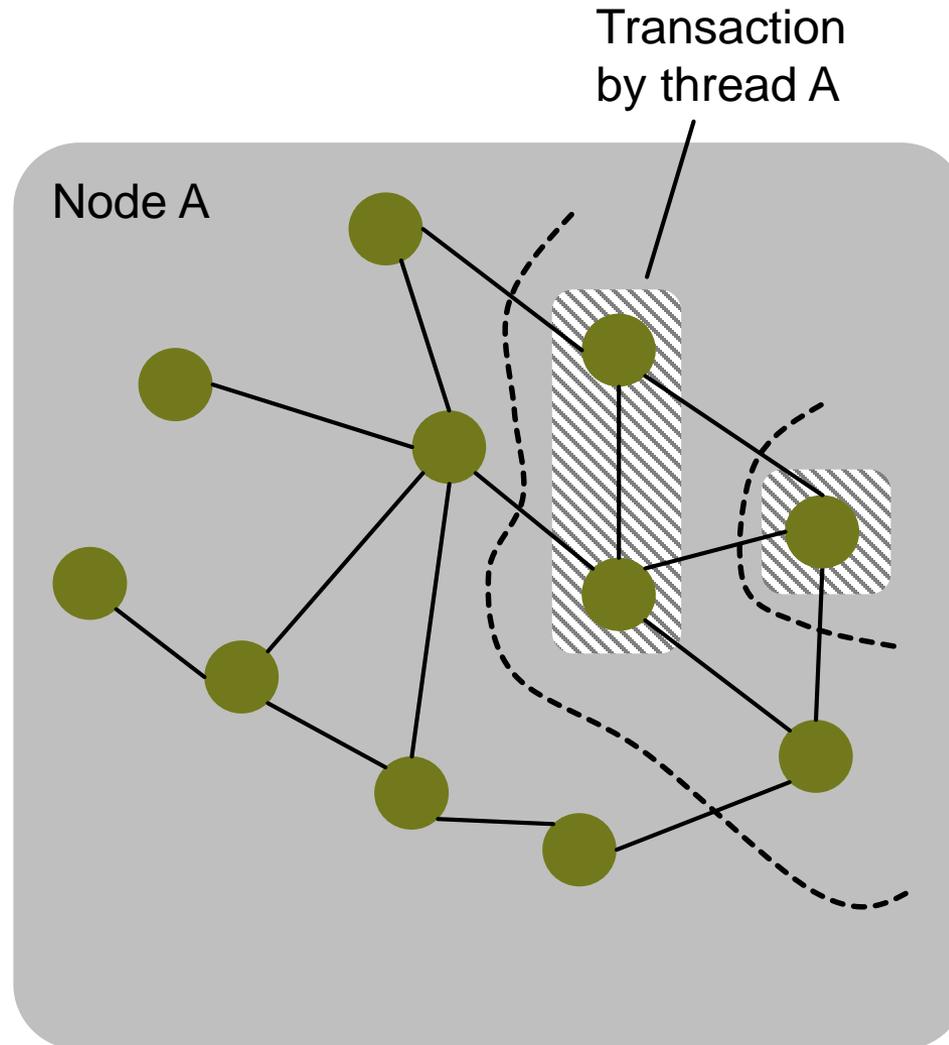
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



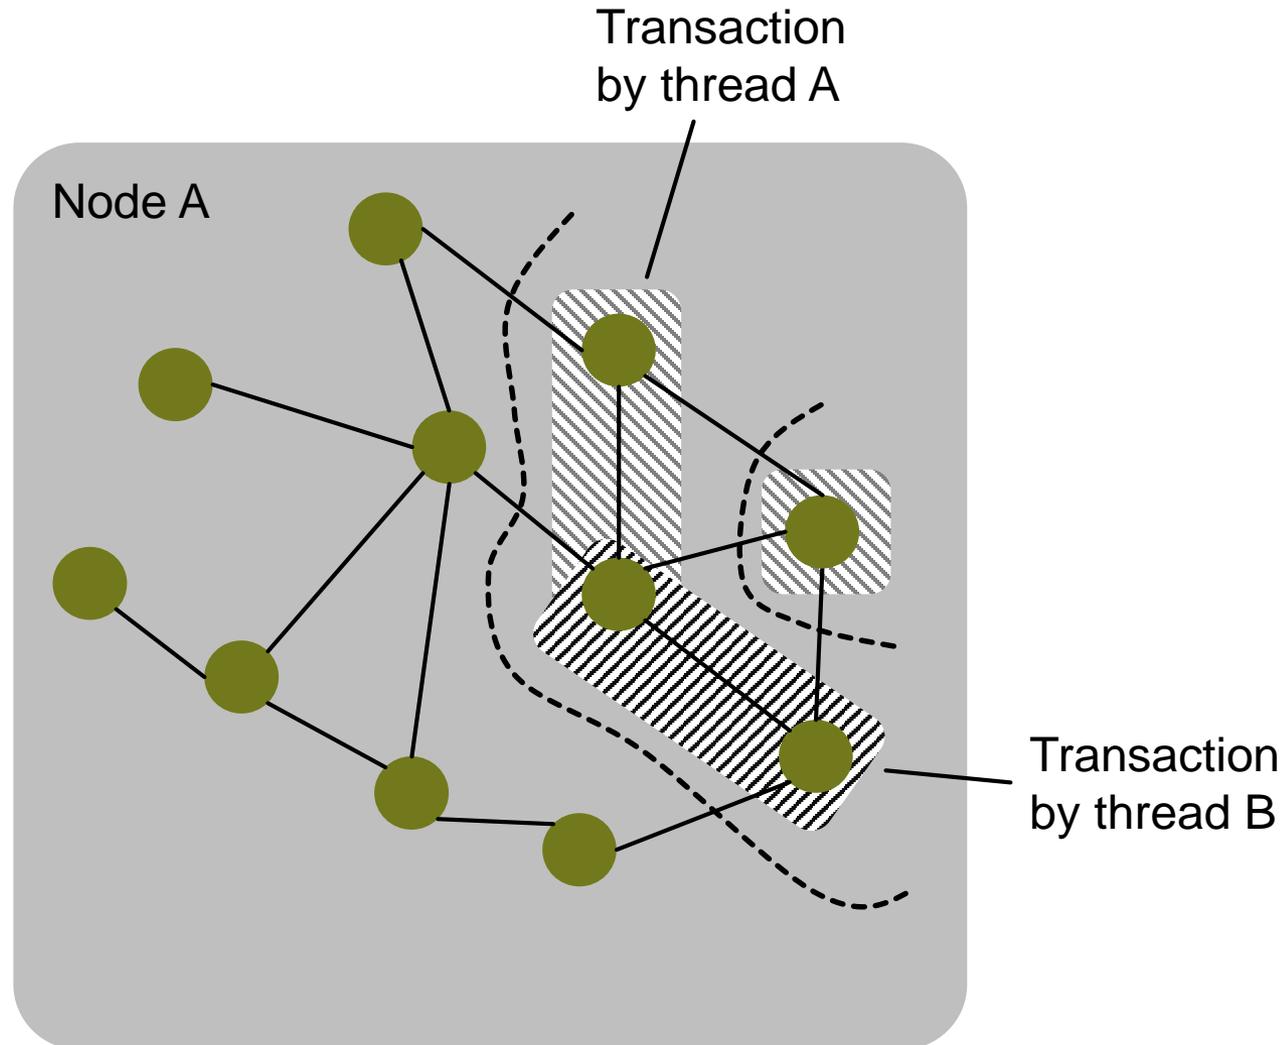
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



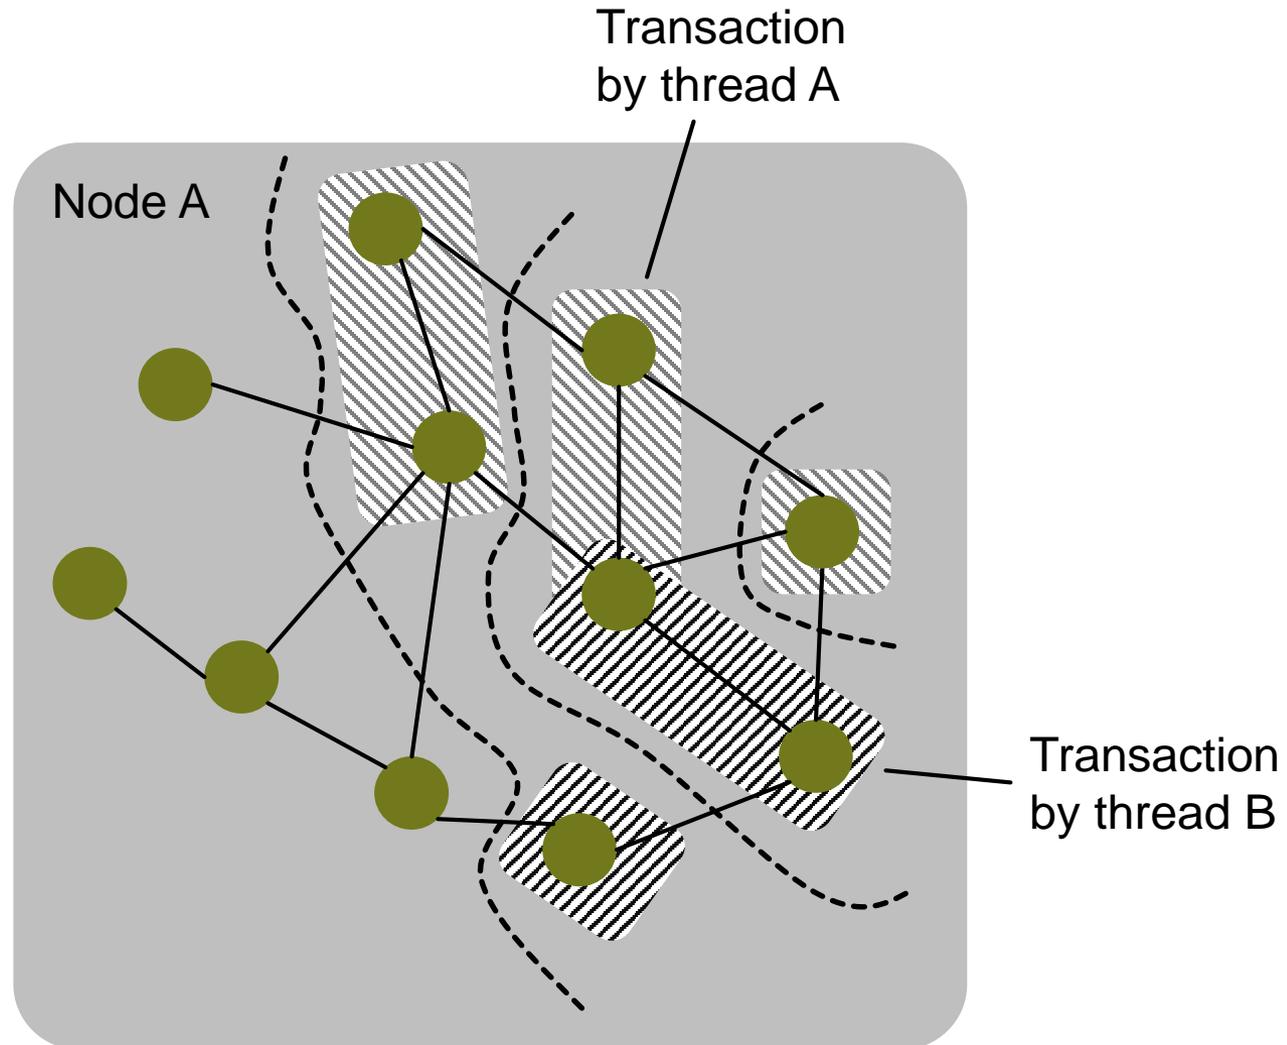
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



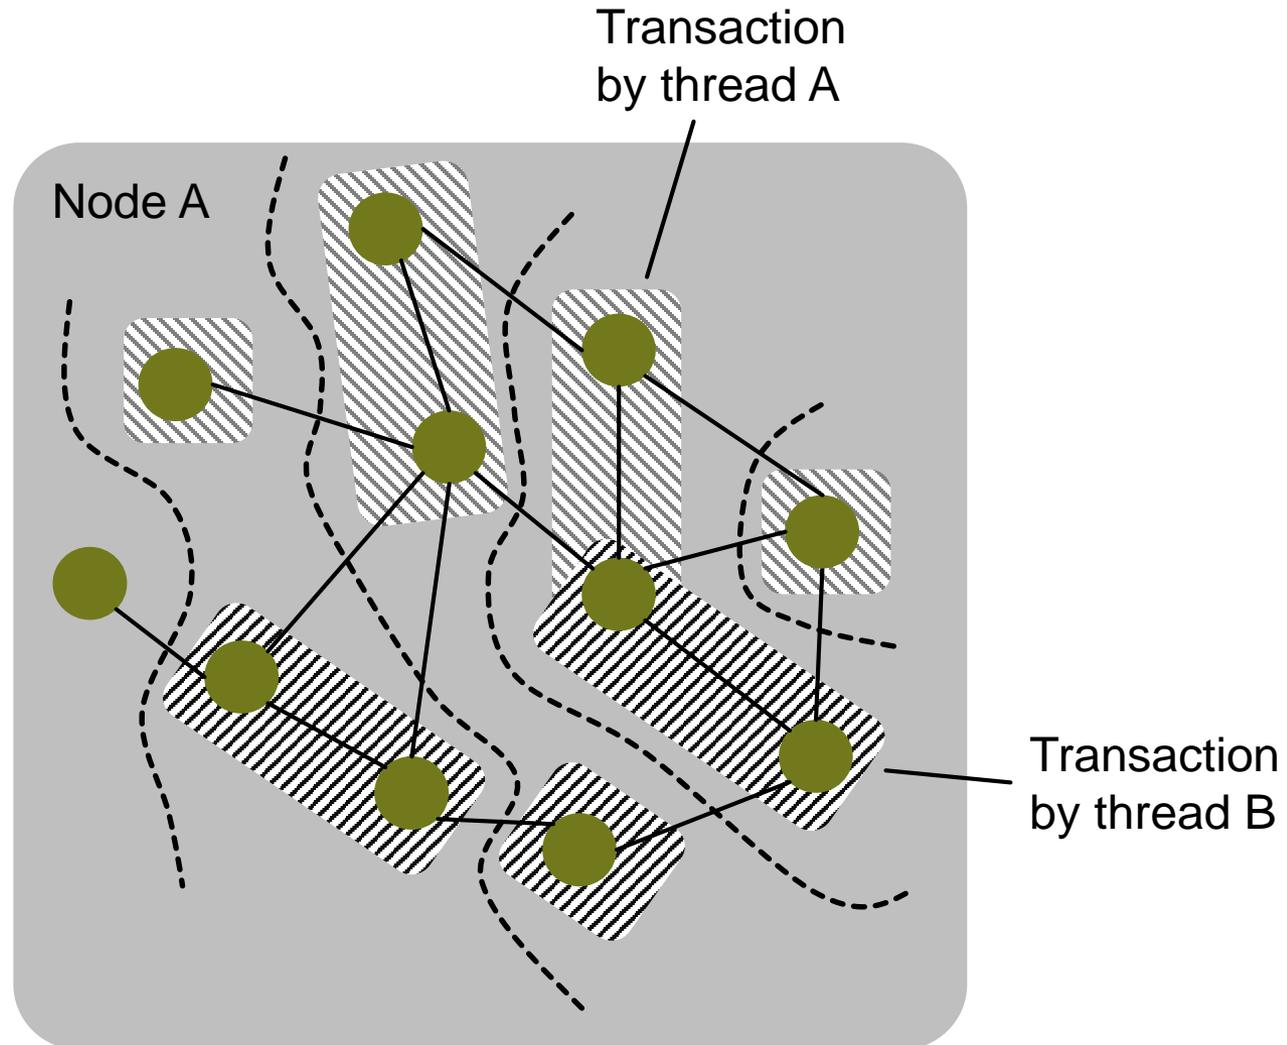
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



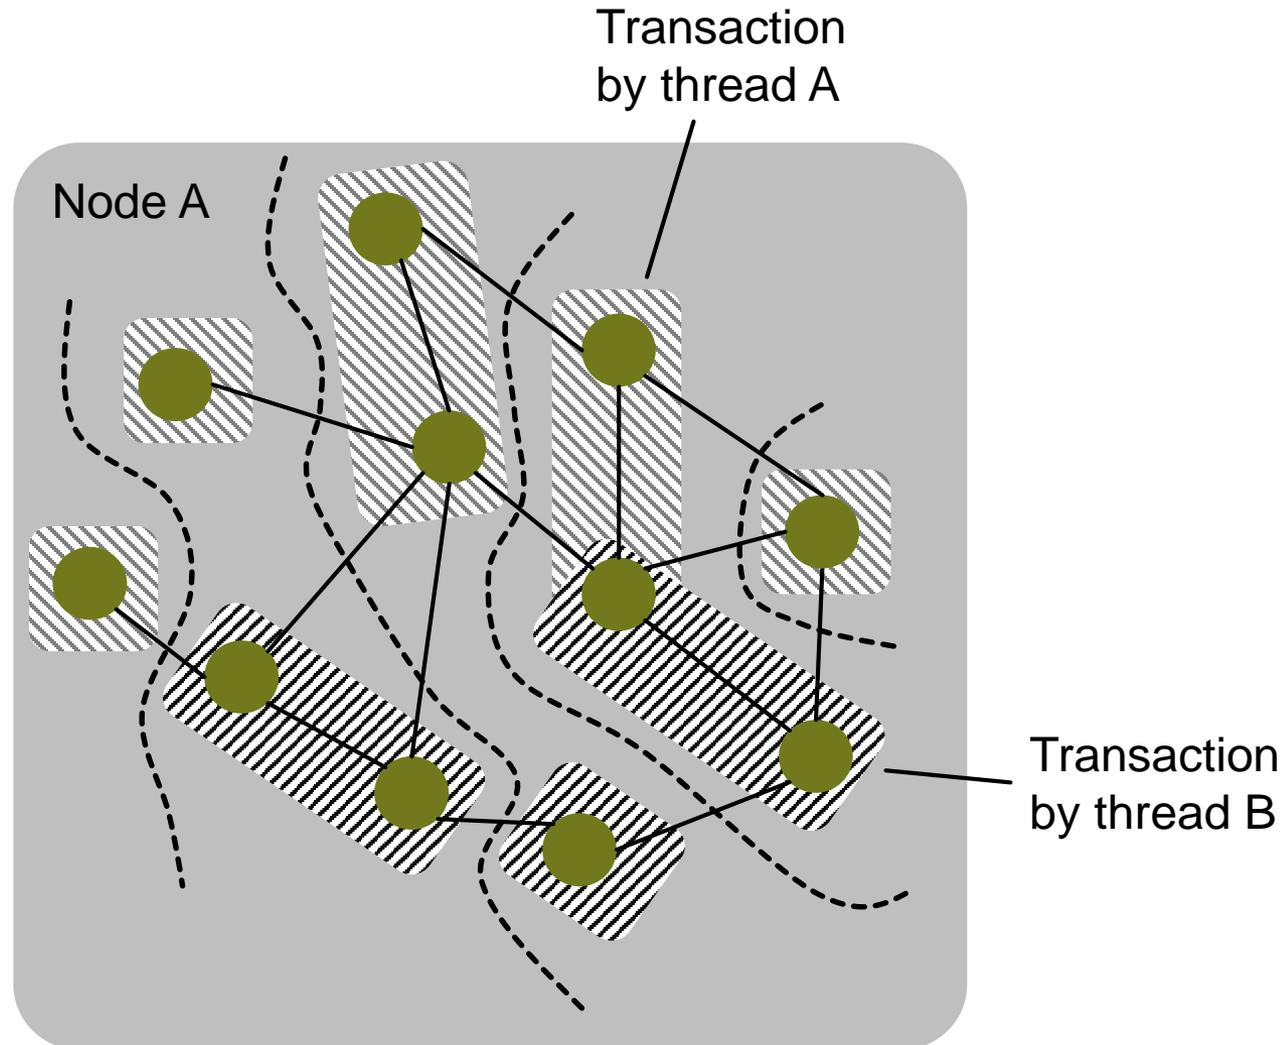
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



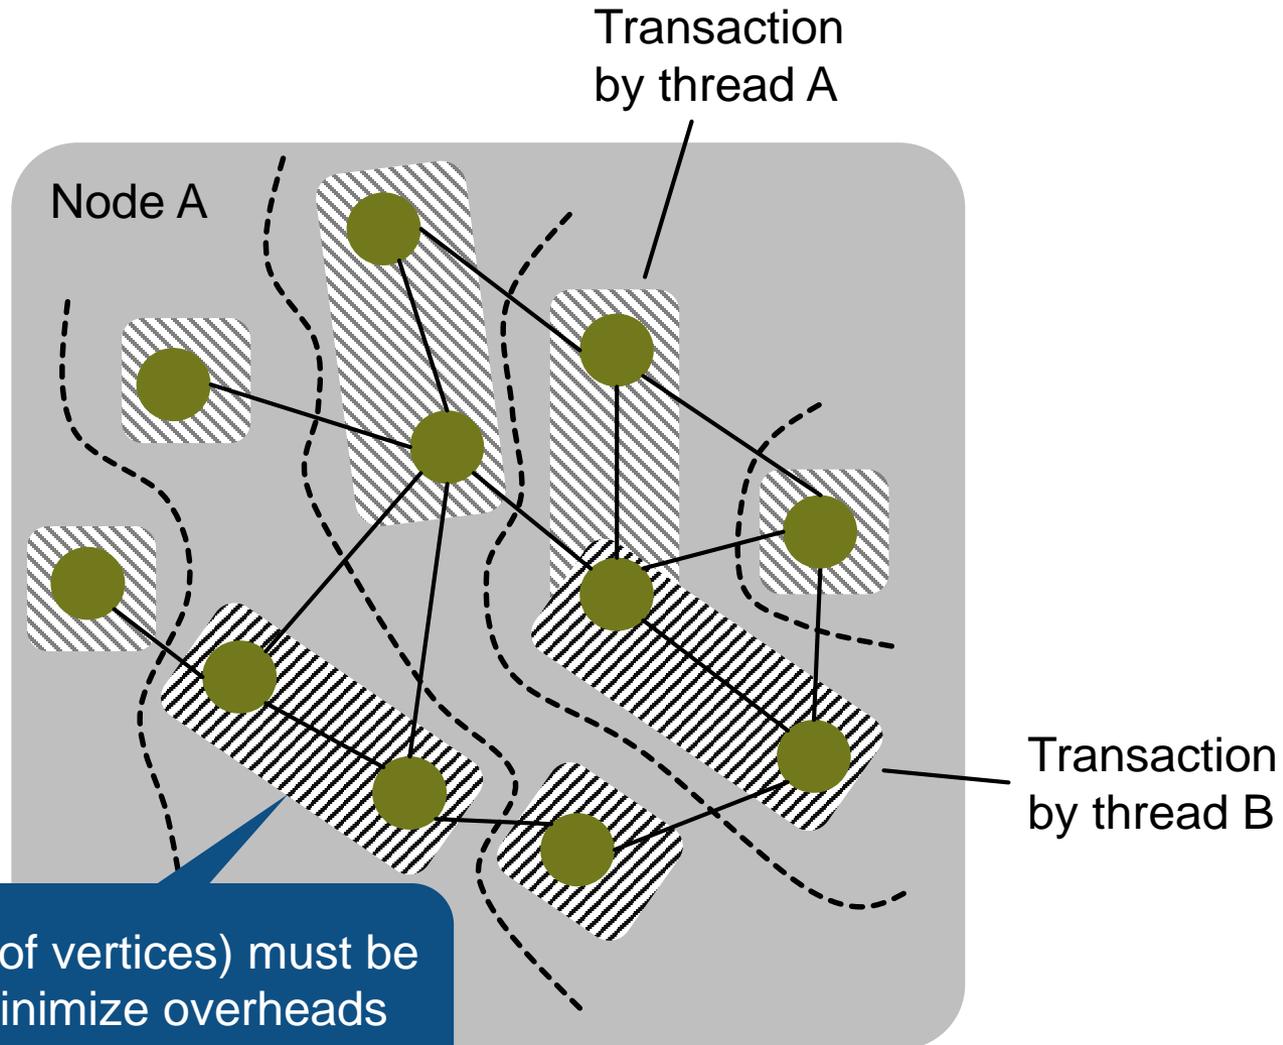
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS



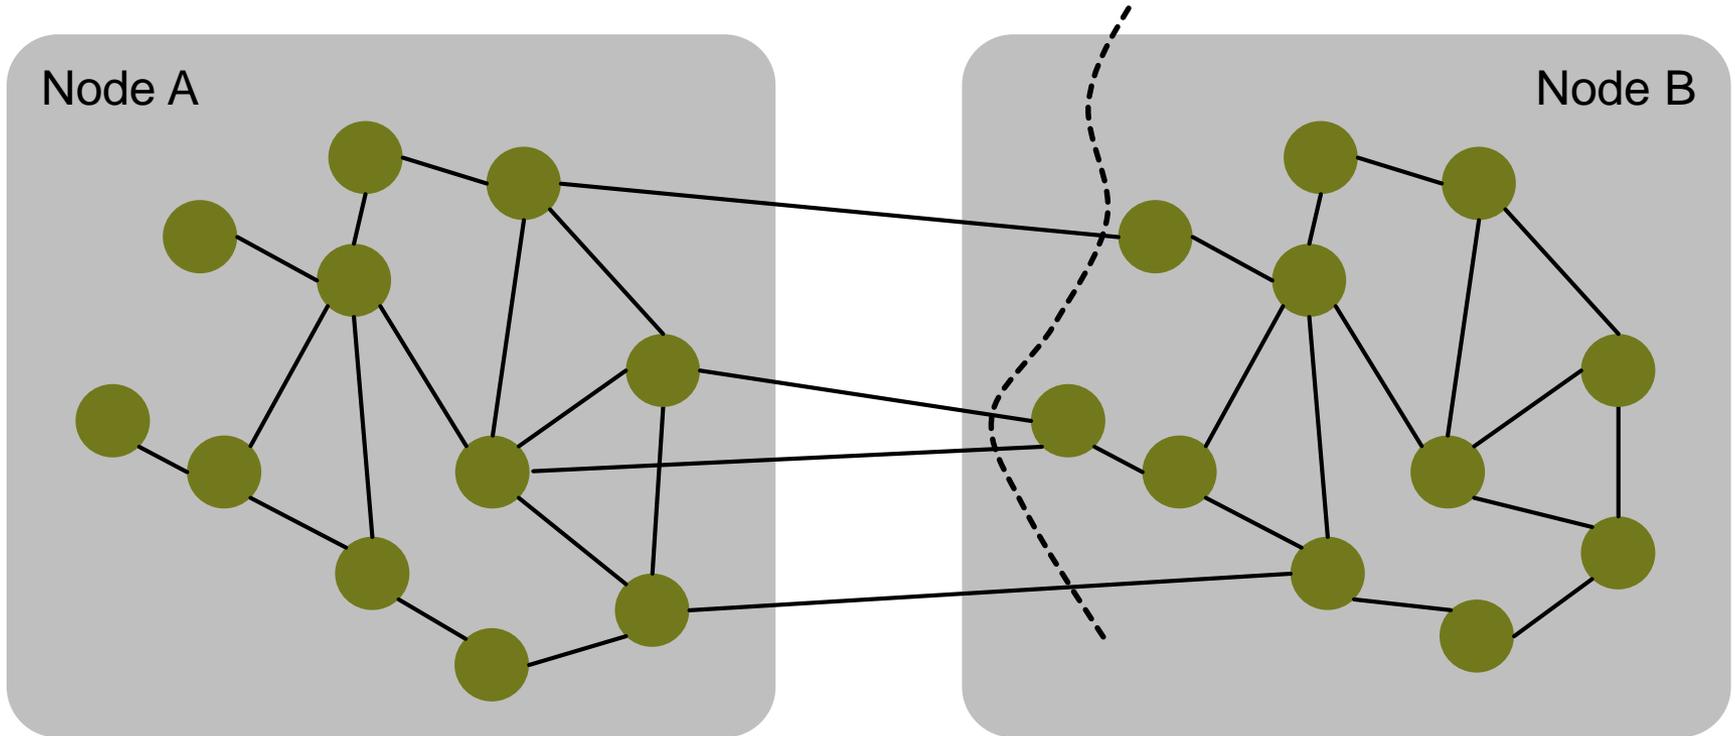
# ACCESSING MULTIPLE VERTICES ATOMICALLY

## Example: BFS

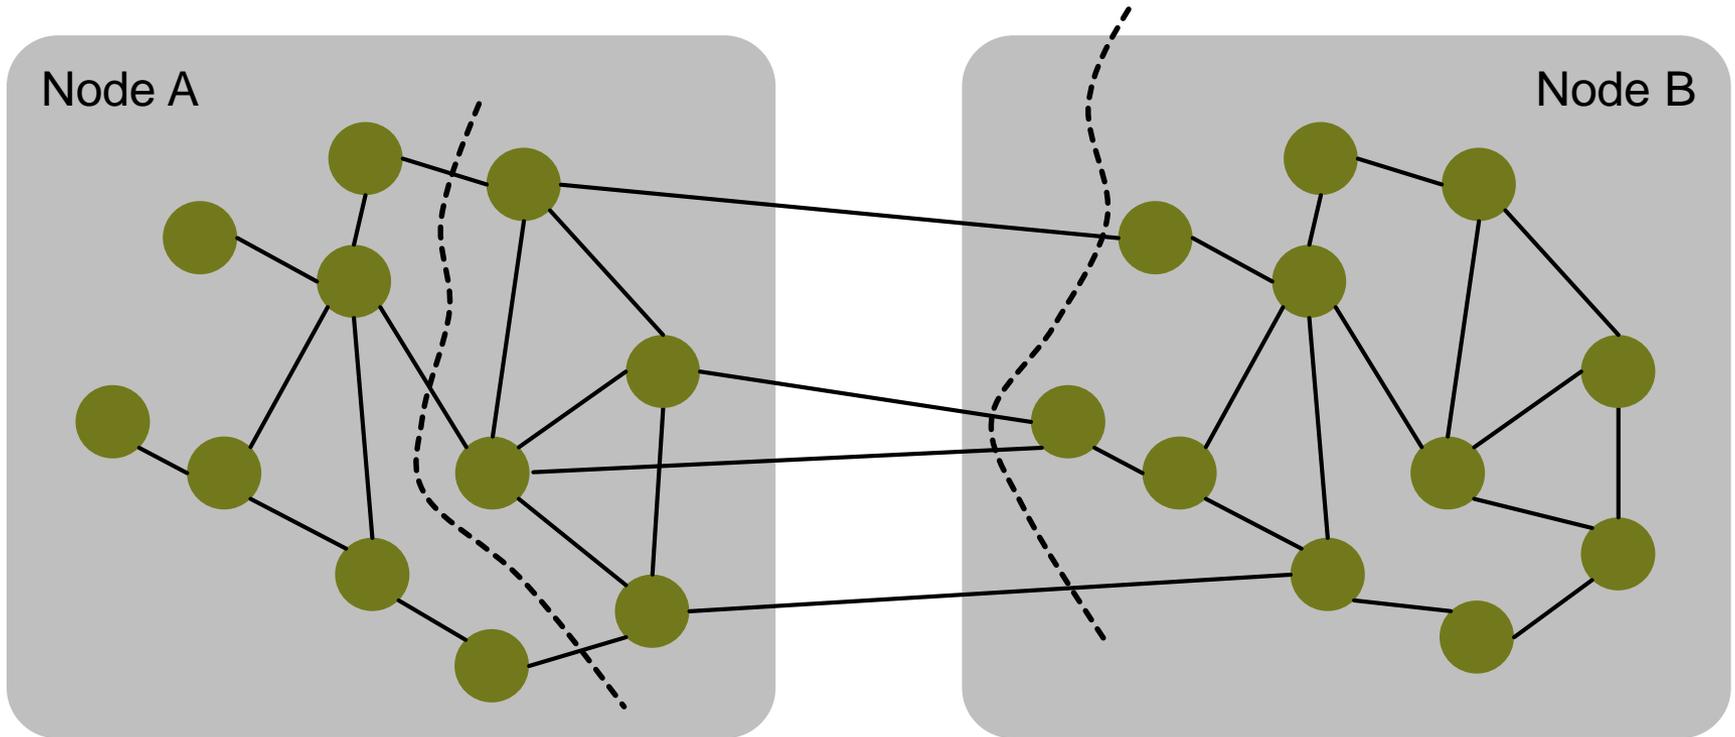


! Size (the number of vertices) must be appropriate to minimize overheads from both commits and rollbacks

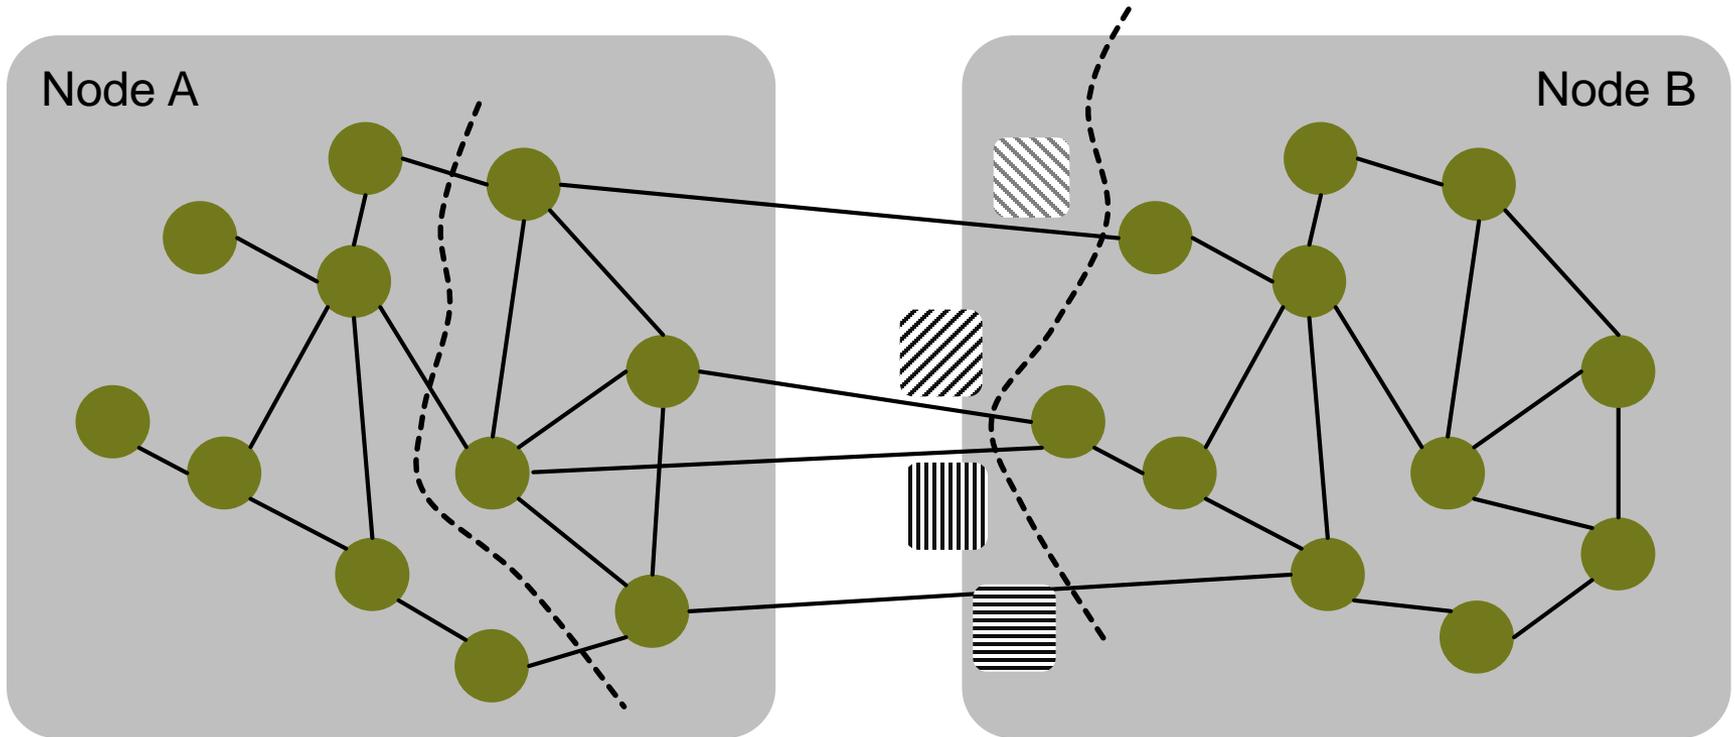
# TRANSFERRING TRANSACTIONS ACROSS NODES



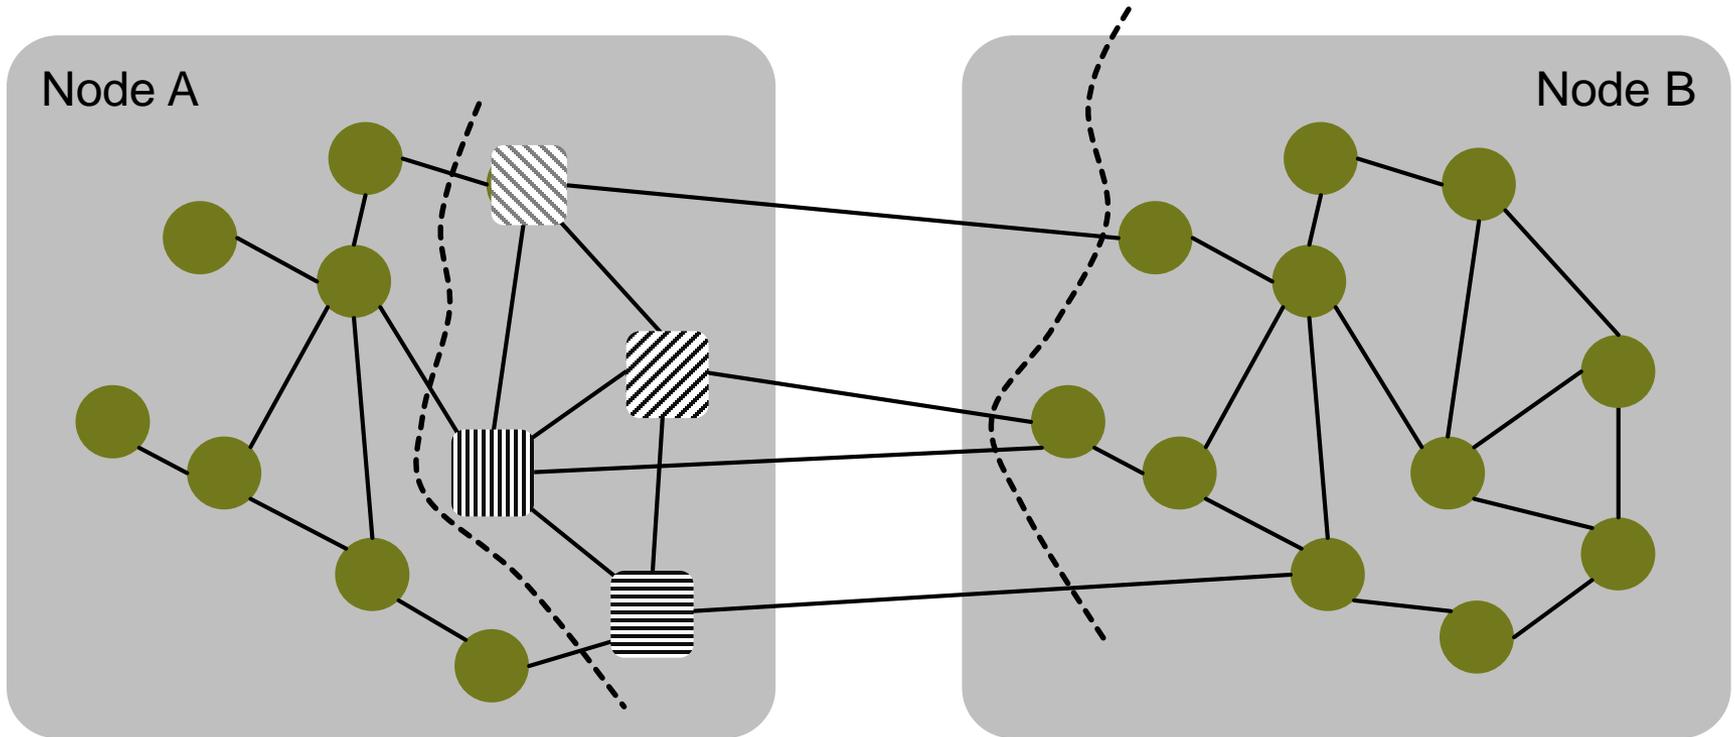
# TRANSFERRING TRANSACTIONS ACROSS NODES



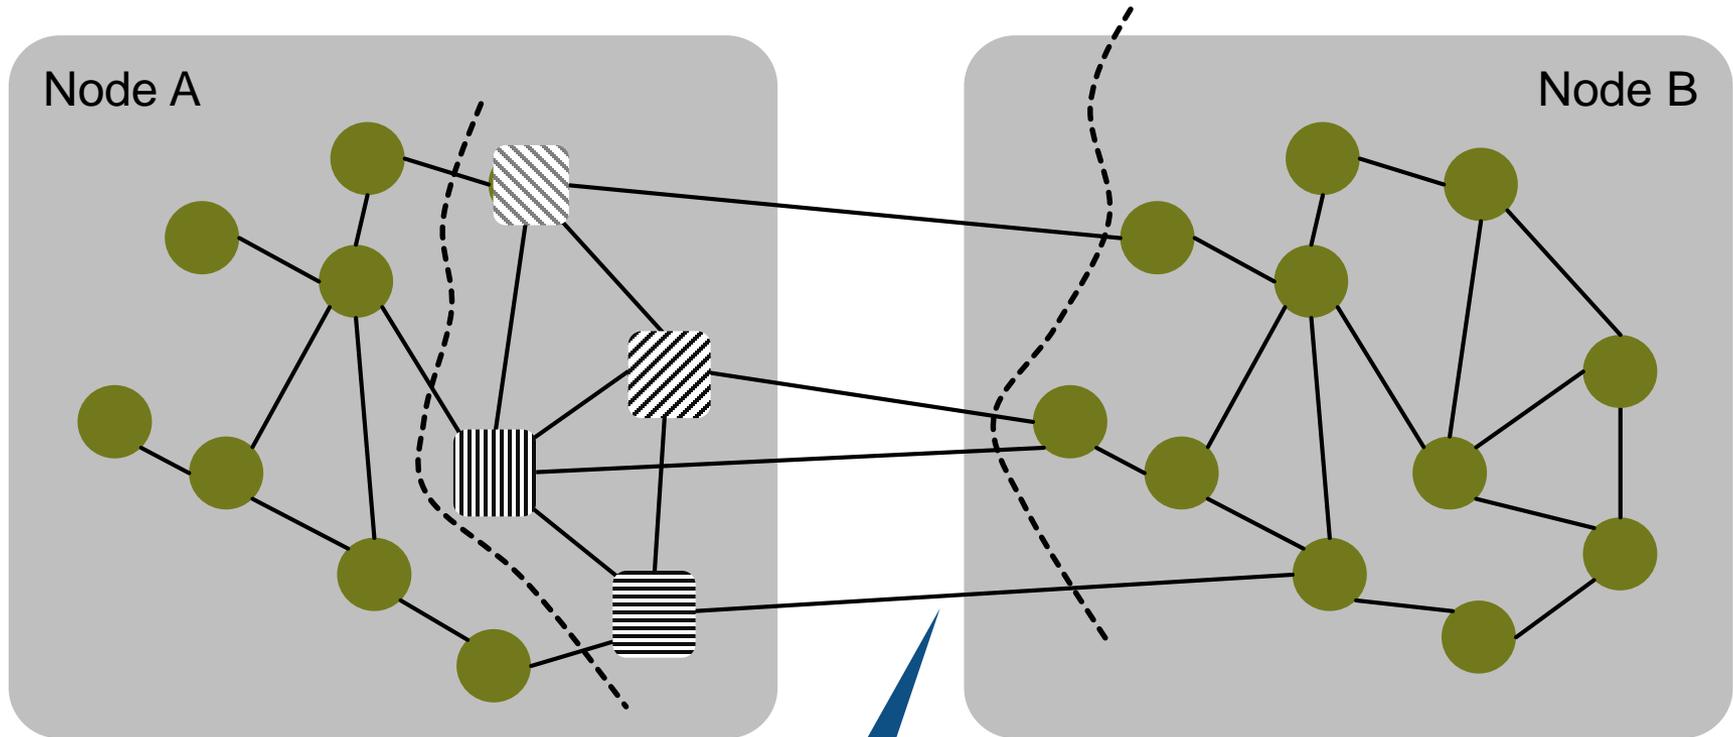
# TRANSFERRING TRANSACTIONS ACROSS NODES



# TRANSFERRING TRANSACTIONS ACROSS NODES

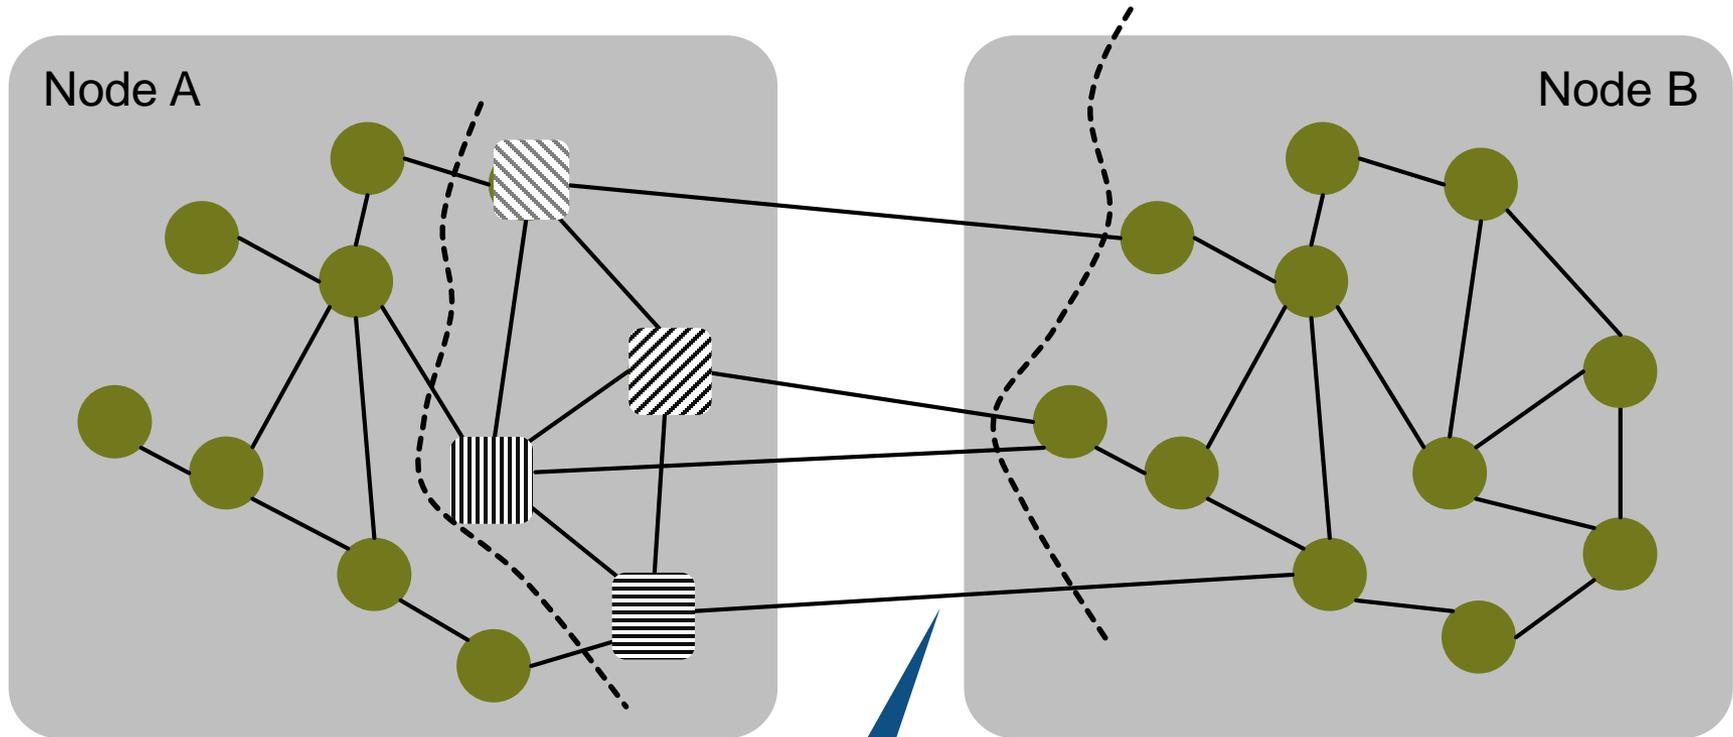


# TRANSFERRING TRANSACTIONS ACROSS NODES



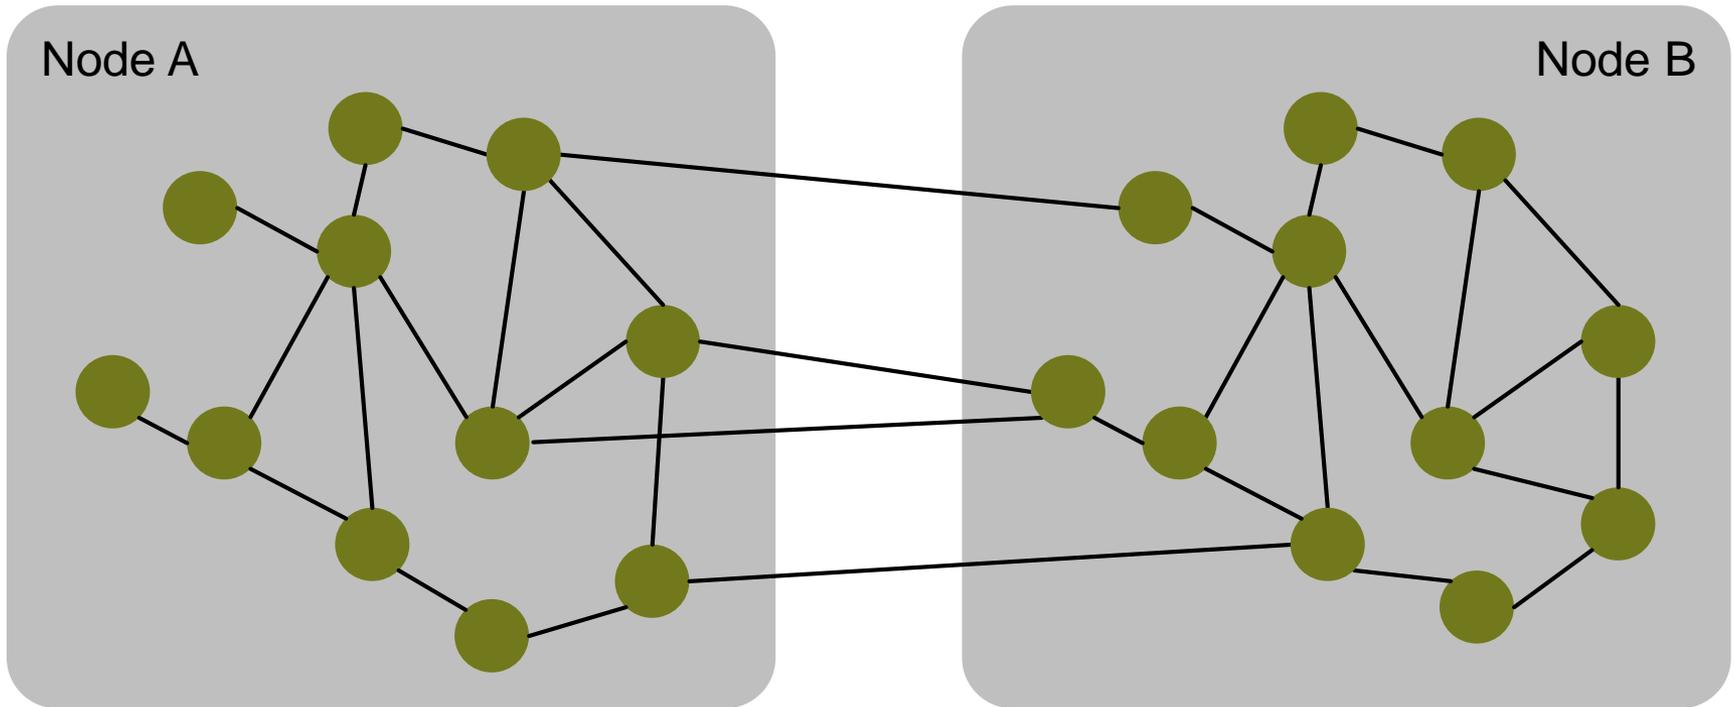
Transactions must be appropriately coalesced to minimize communication overheads

# TRANSFERRING TRANSACTIONS ACROSS NODES

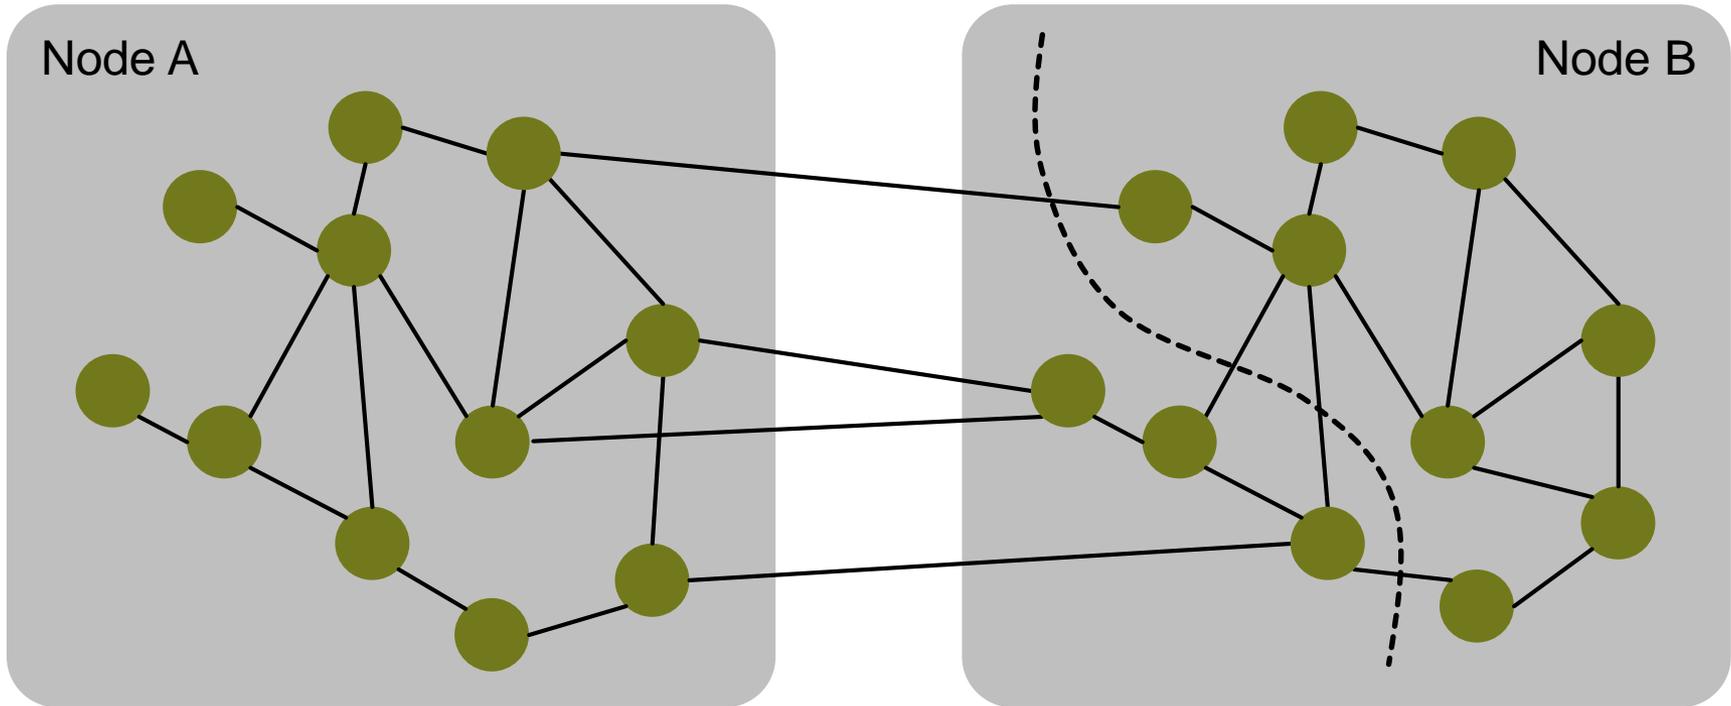


Transactions must be appropriately coalesced to minimize communication overheads

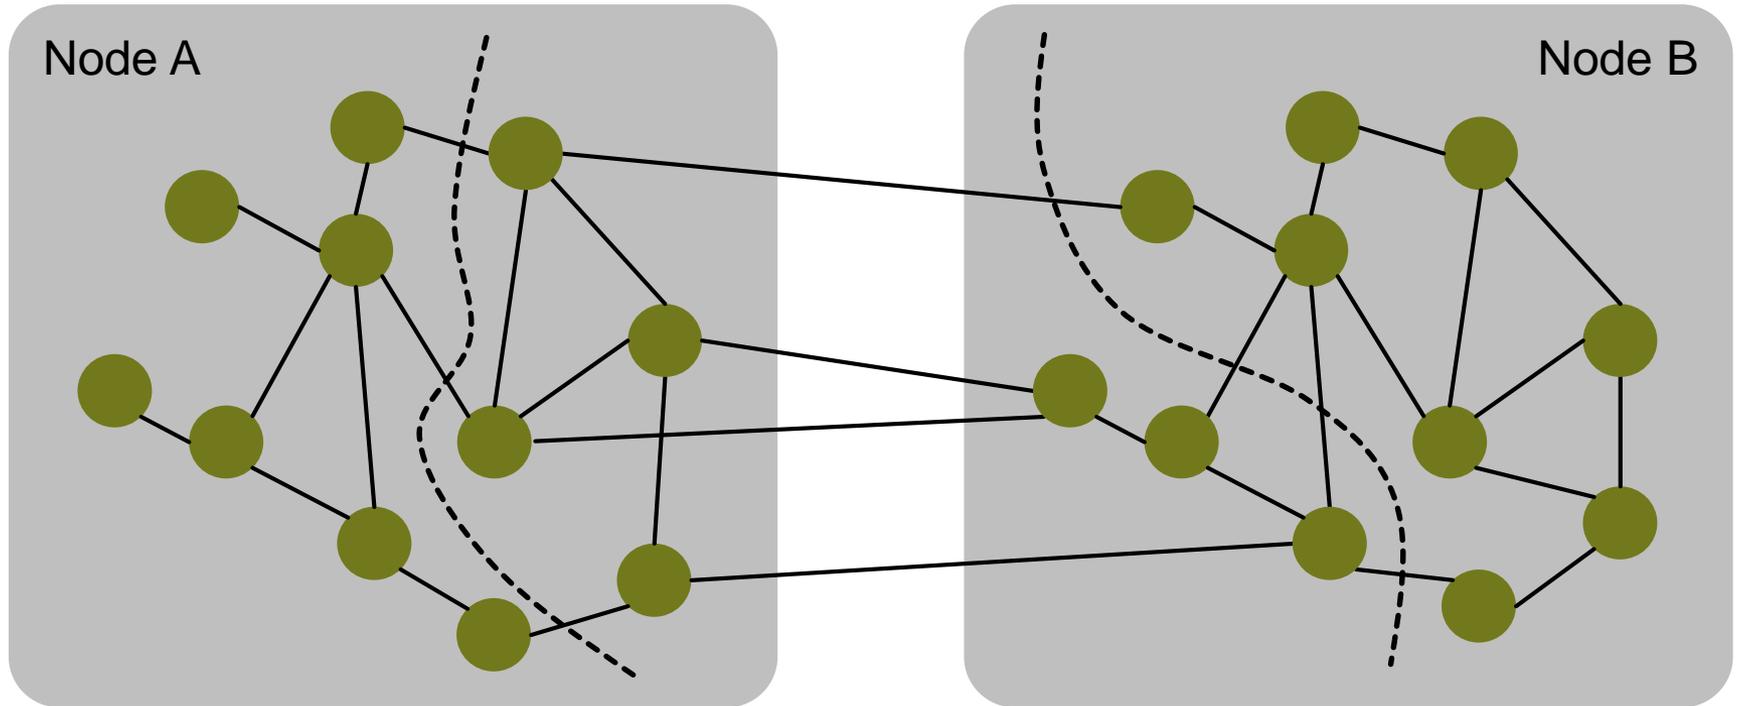
# EXECUTING TRANSACTIONS ON MULTIPLE NODES



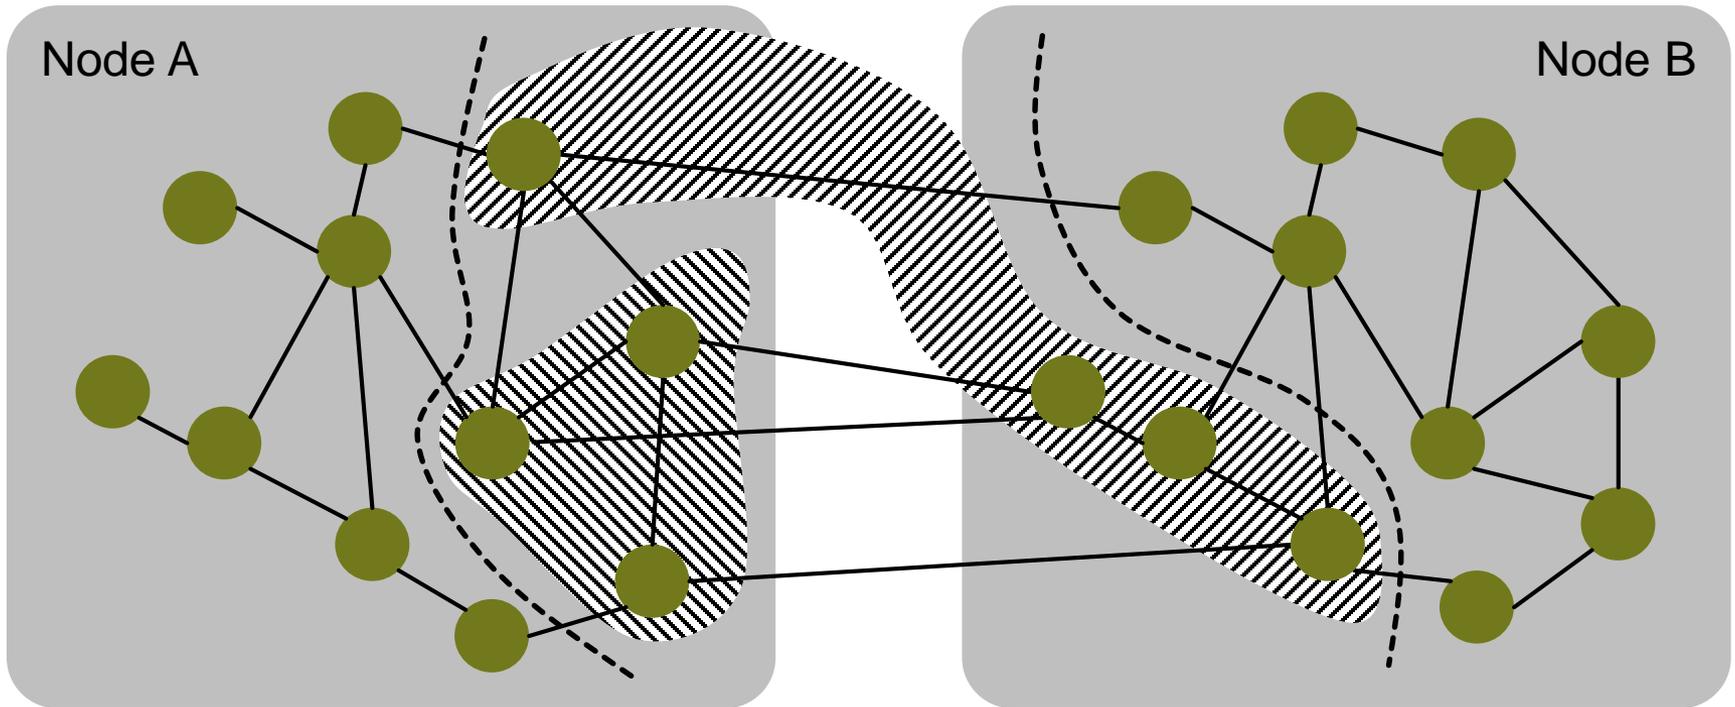
# EXECUTING TRANSACTIONS ON MULTIPLE NODES



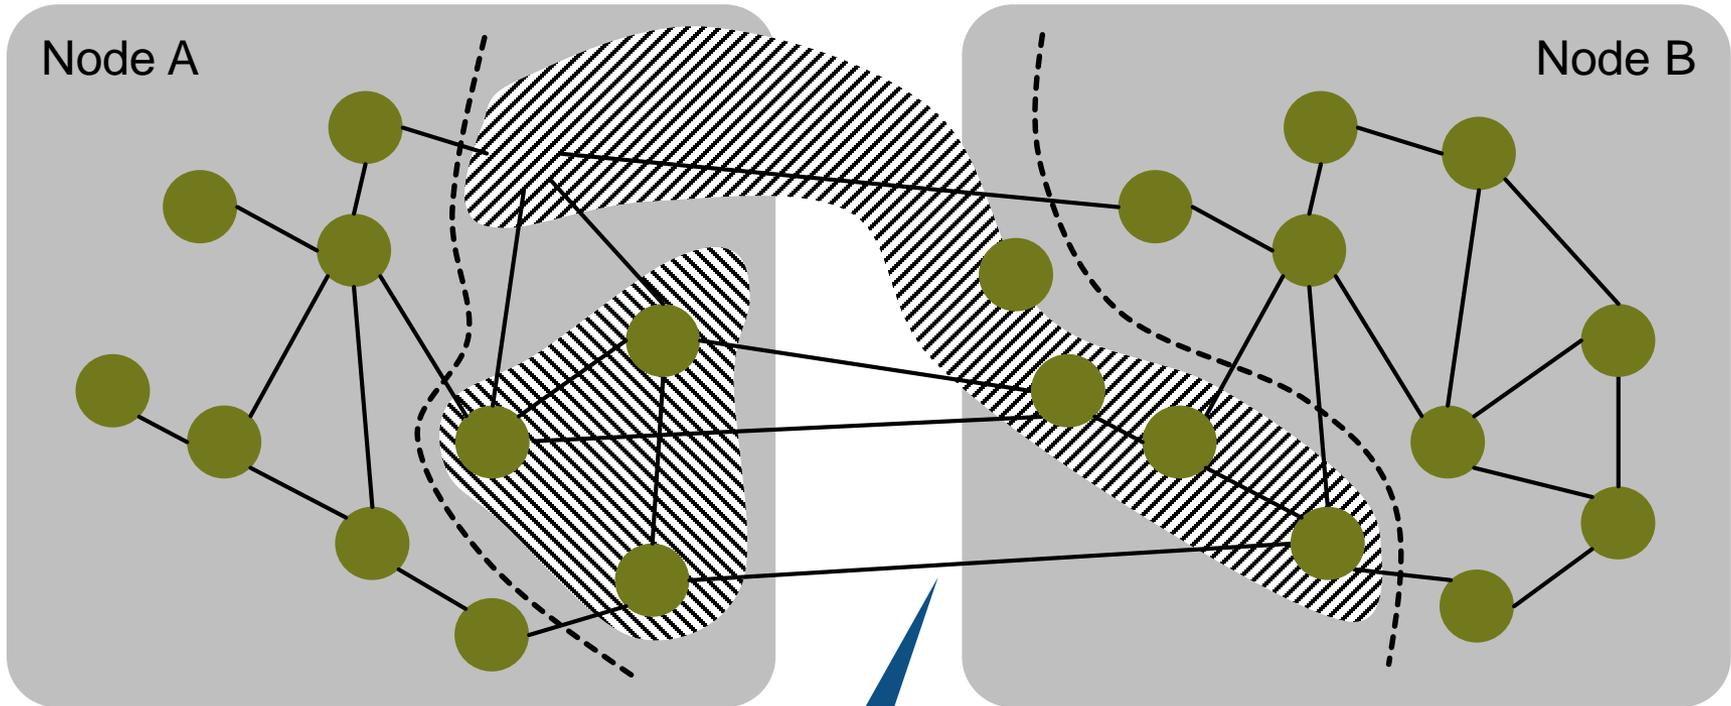
# EXECUTING TRANSACTIONS ON MULTIPLE NODES



# EXECUTING TRANSACTIONS ON MULTIPLE NODES

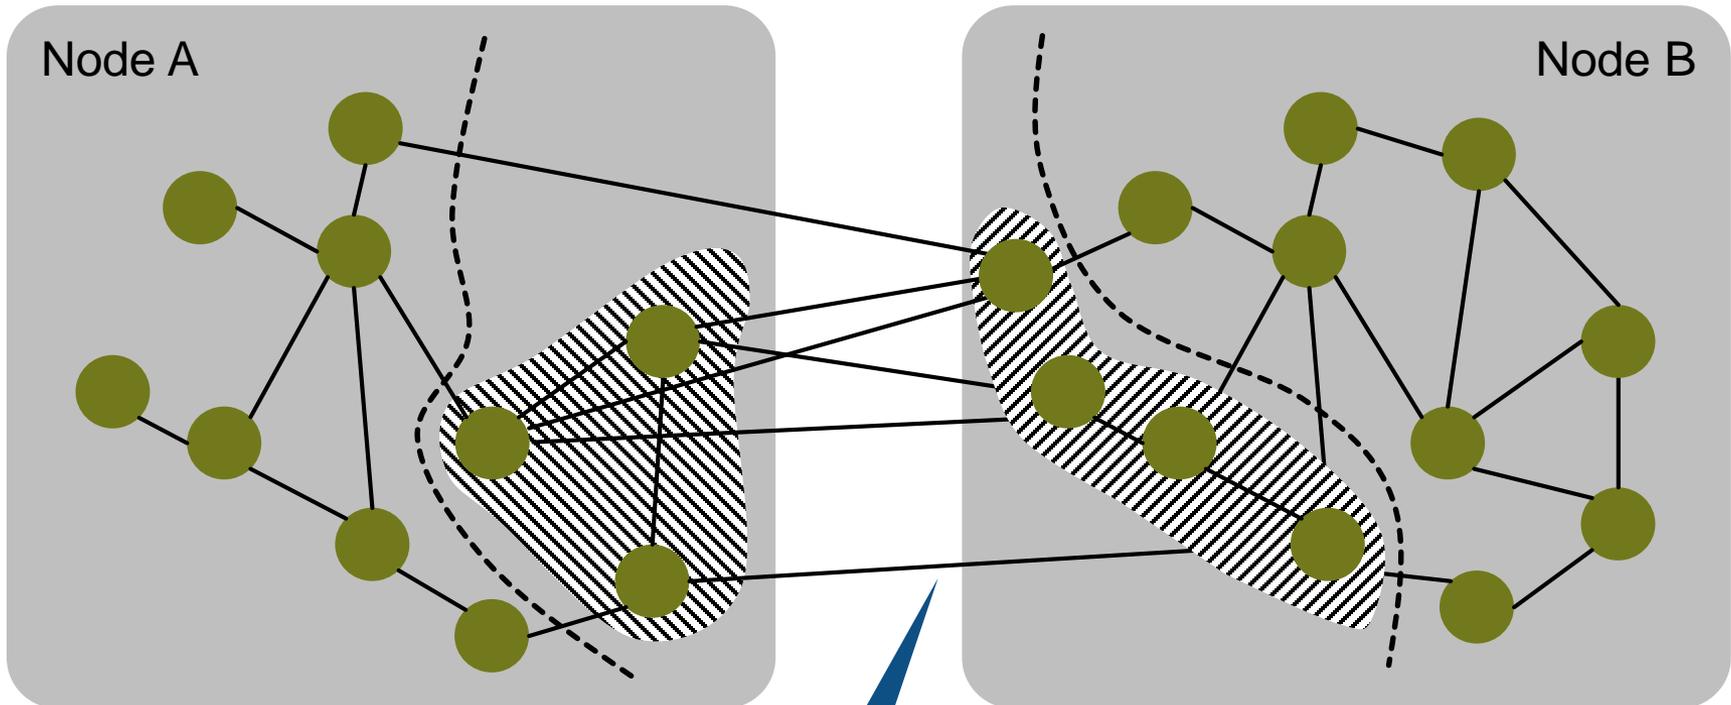


# EXECUTING TRANSACTIONS ON MULTIPLE NODES



Vertices must be appropriately relocated to enable execution of a hardware transaction

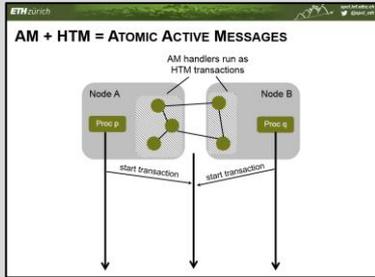
# EXECUTING TRANSACTIONS ON MULTIPLE NODES



Vertices must be appropriately relocated to enable execution of a hardware transaction

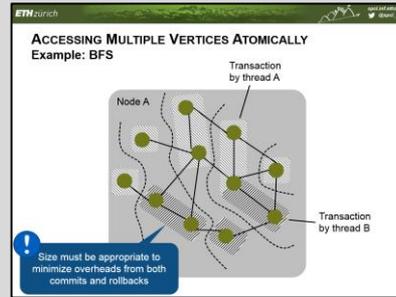
# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



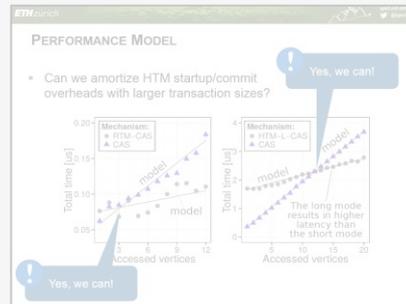
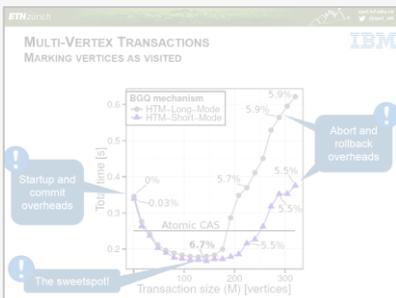
HTM + Active Messages  
= Atomic Active Messages

## Coarsening & coalescing



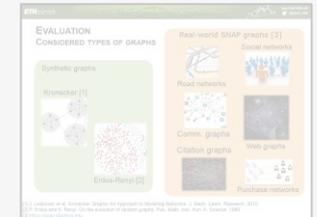
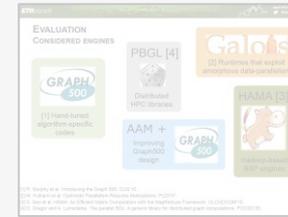
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

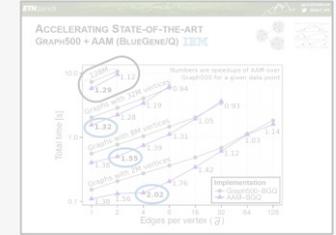
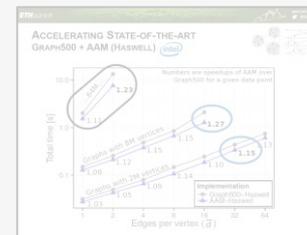


## Performance model

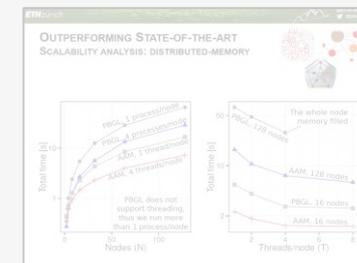
## Evaluation



## Considered engines and graphs



## Accelerating state-of-the-art



## Scalability





# PERFORMANCE ANALYSIS

## RESEARCH QUESTIONS

# PERFORMANCE ANALYSIS

## RESEARCH QUESTIONS



How can we implement AAM handlers to run most efficiently?

# PERFORMANCE ANALYSIS

## RESEARCH QUESTIONS



How can we implement AAM handlers to run most efficiently?



What are performance tradeoffs related to HTM?

# PERFORMANCE ANALYSIS

## RESEARCH QUESTIONS



How can we implement AAM handlers to run most efficiently?



What are performance tradeoffs related to HTM?



What are advantages of HTM over atomics for AAM?

# PERFORMANCE ANALYSIS

## RESEARCH QUESTIONS



How can we implement AAM handlers to run most efficiently?



What are performance tradeoffs related to HTM?



What are advantages of HTM over atomics for AAM?



What are the optimal transaction sizes?  
Can we amortize transaction overheads?

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

- Evaluation on 3 machines

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

- Evaluation on 3 machines
  - Intel Haswell server



Commodity machines

# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

- Evaluation on 3 machines
  - Intel Haswell server
  - InfiniBand cluster



# PERFORMANCE ANALYSIS

## TYPES OF MACHINES

- Evaluation on 3 machines
  - Intel Haswell server
  - InfiniBand cluster
  - IBM BlueGene/Q



# PERFORMANCE ANALYSIS CONSIDERED MECHANISMS

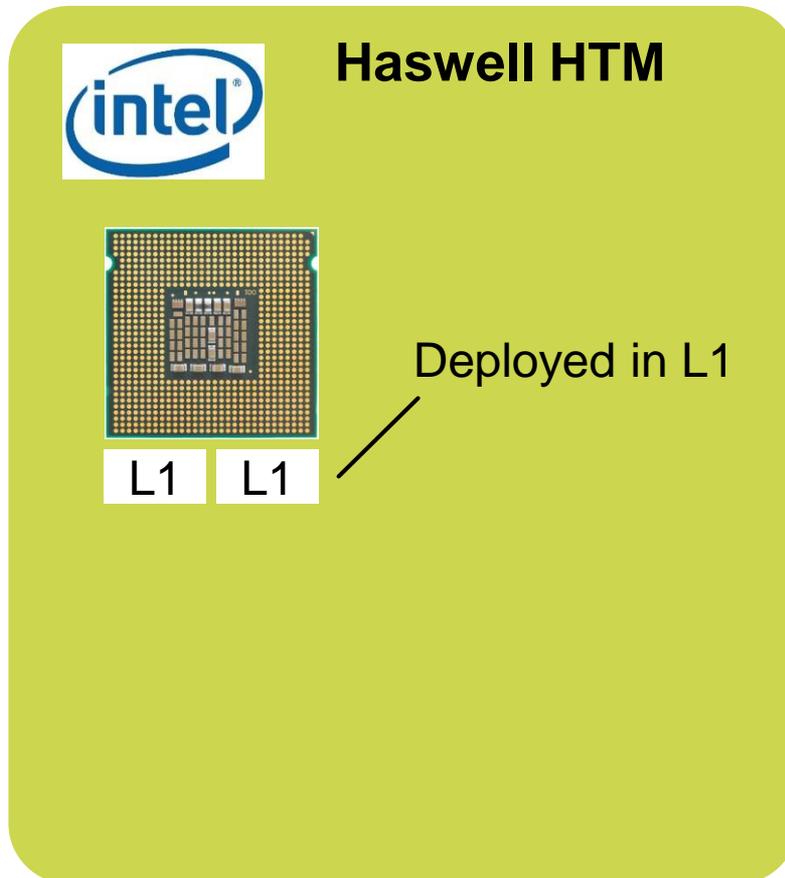
# PERFORMANCE ANALYSIS CONSIDERED MECHANISMS



Haswell HTM

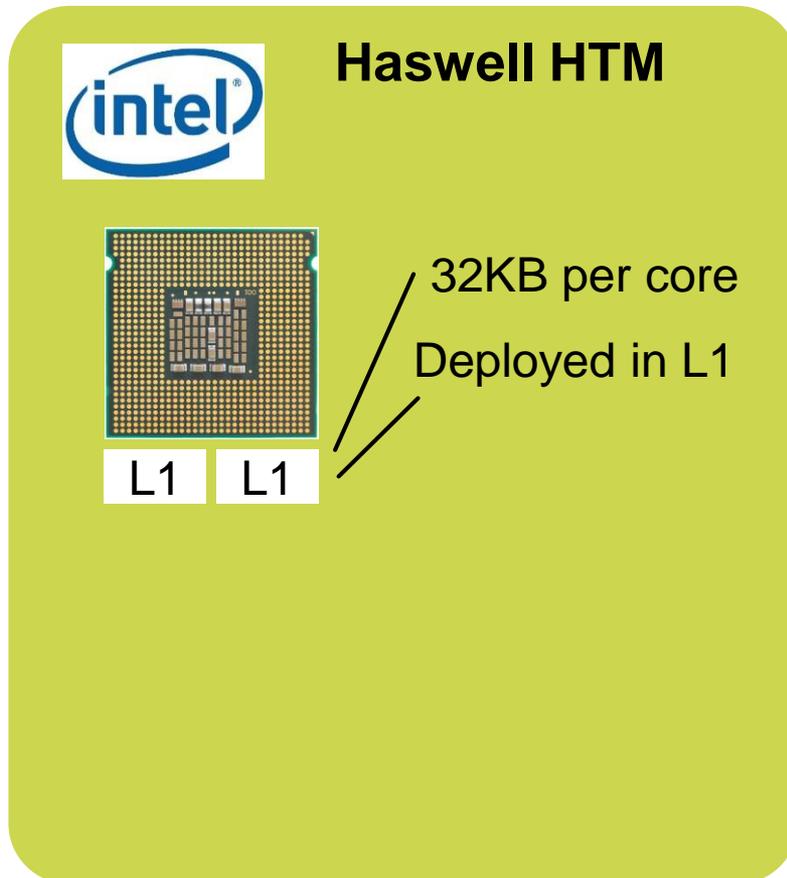
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



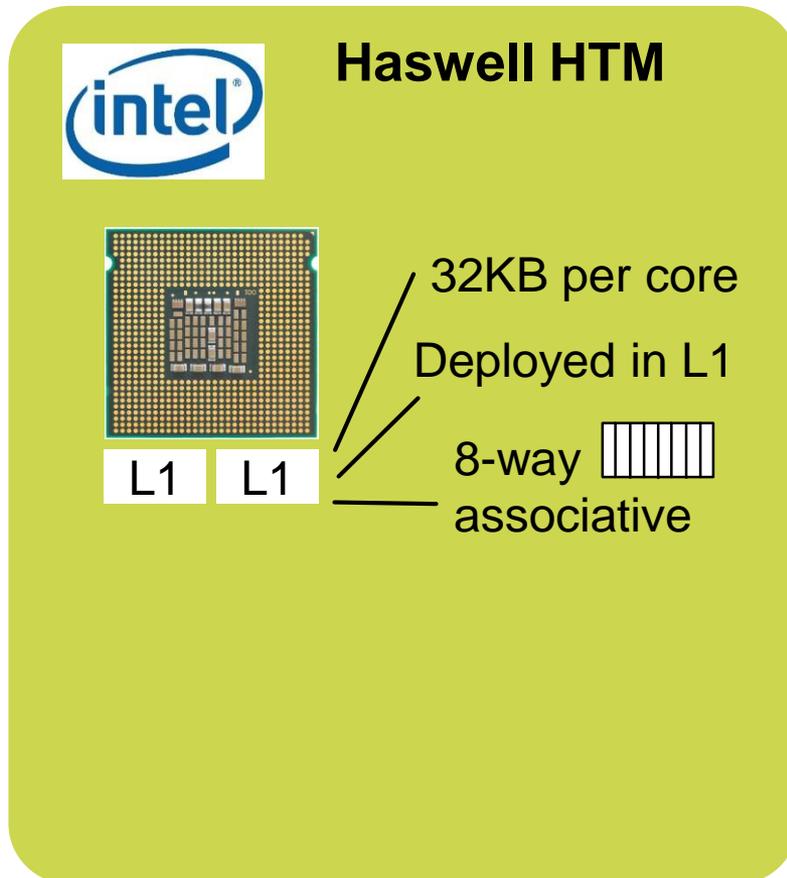
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



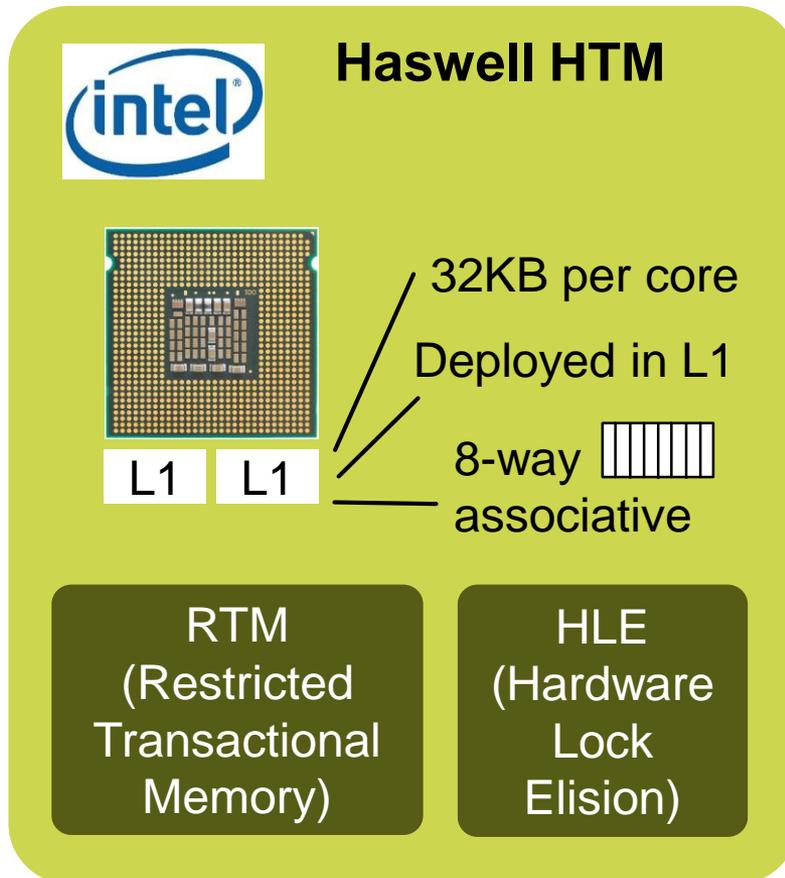
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



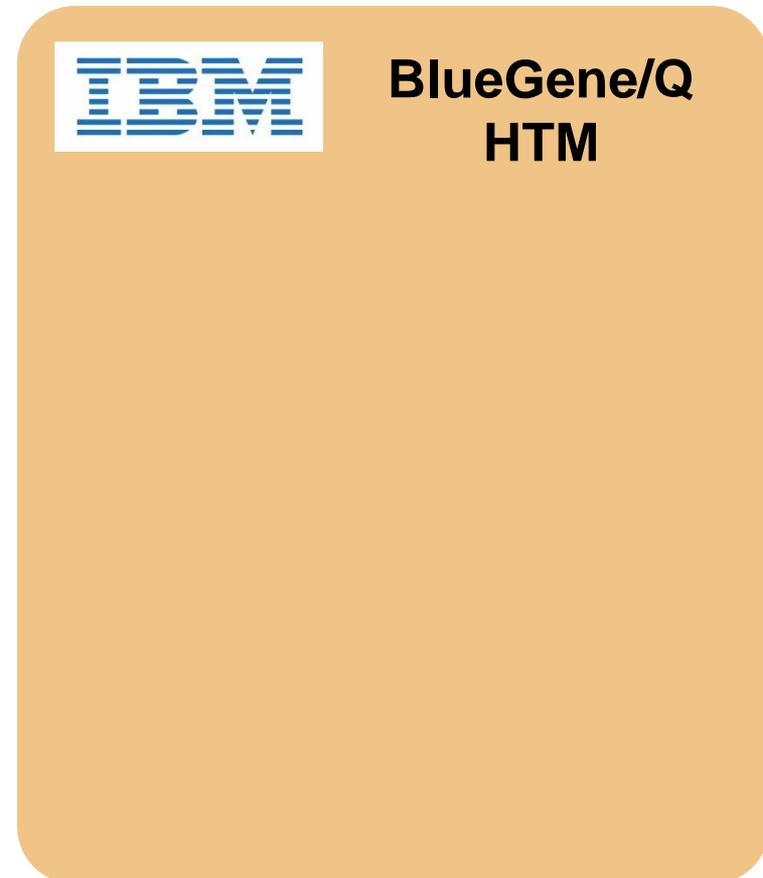
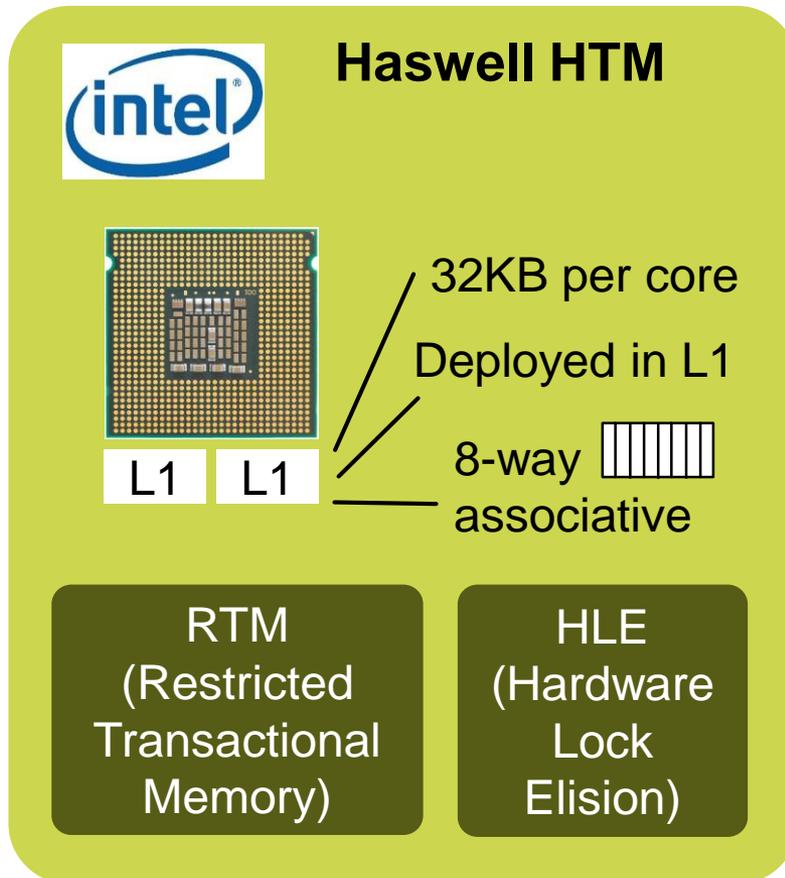
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



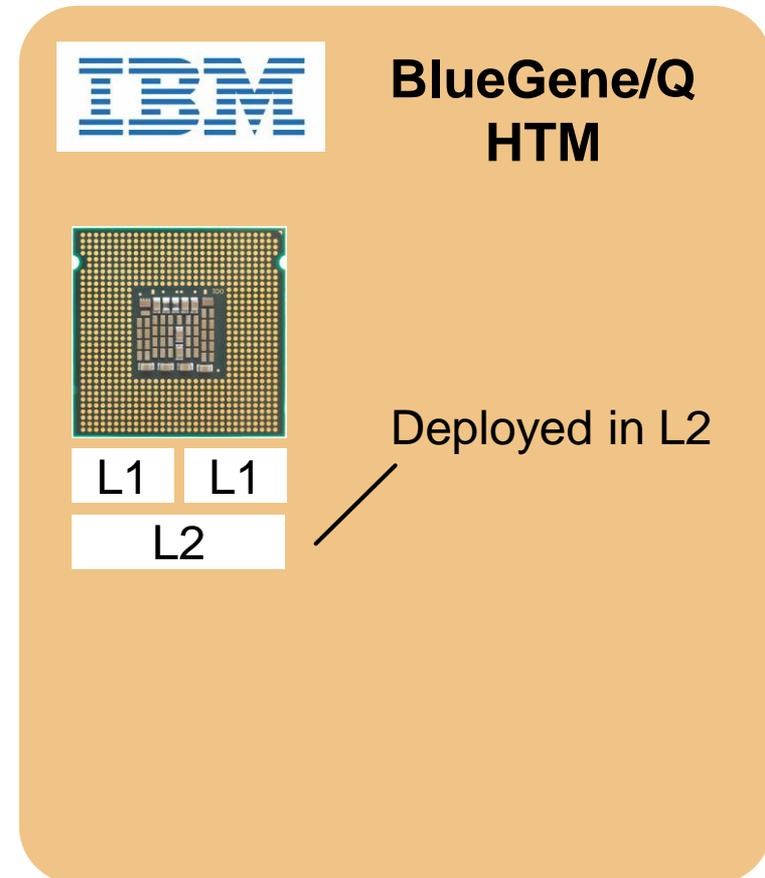
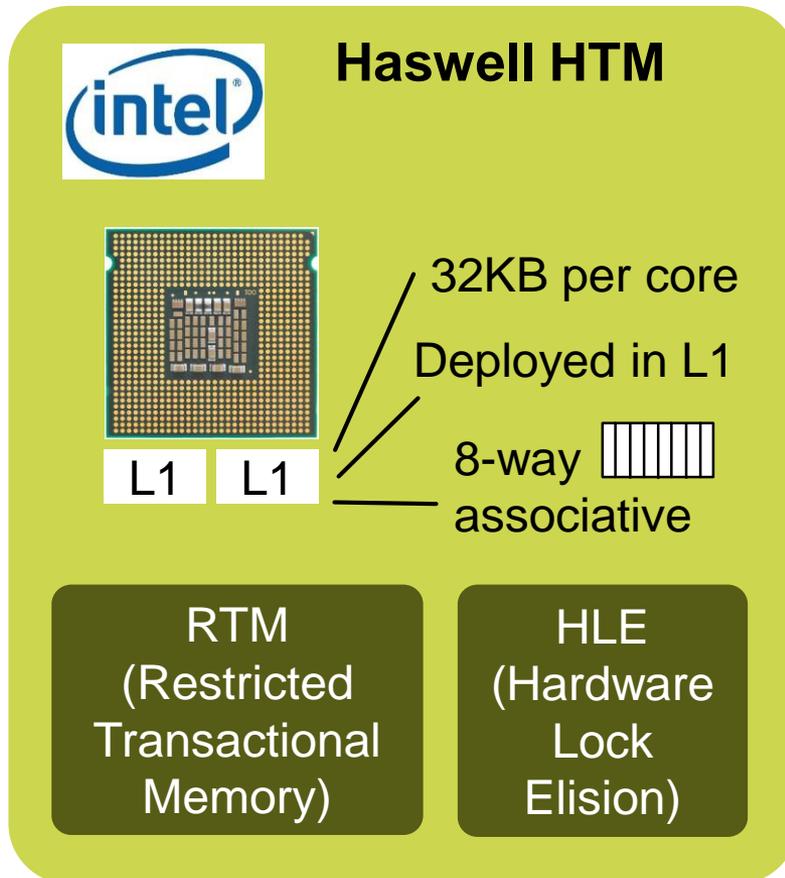
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



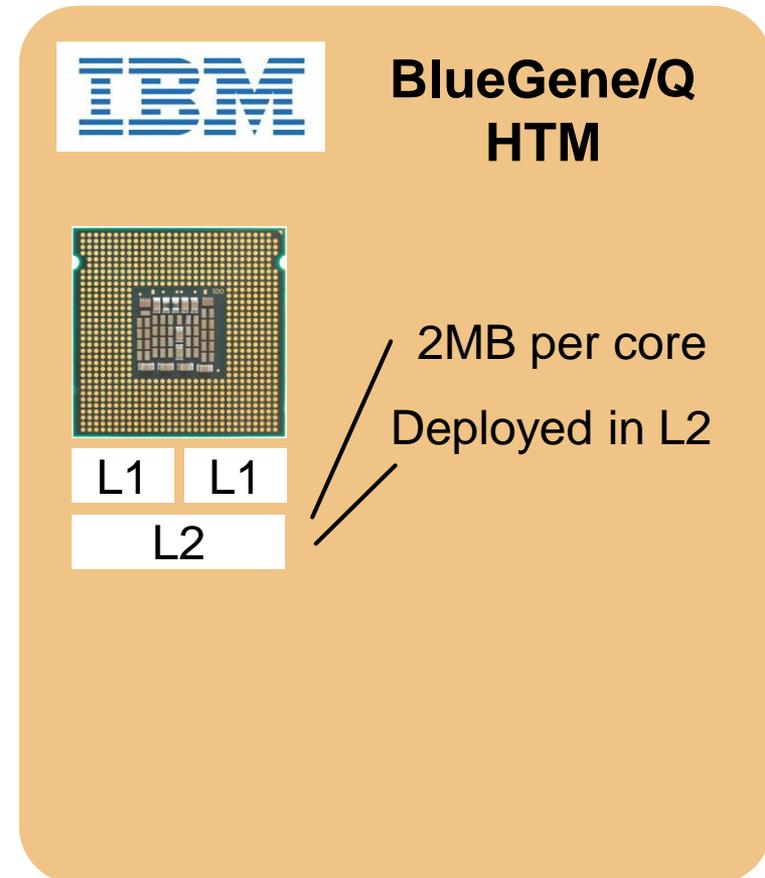
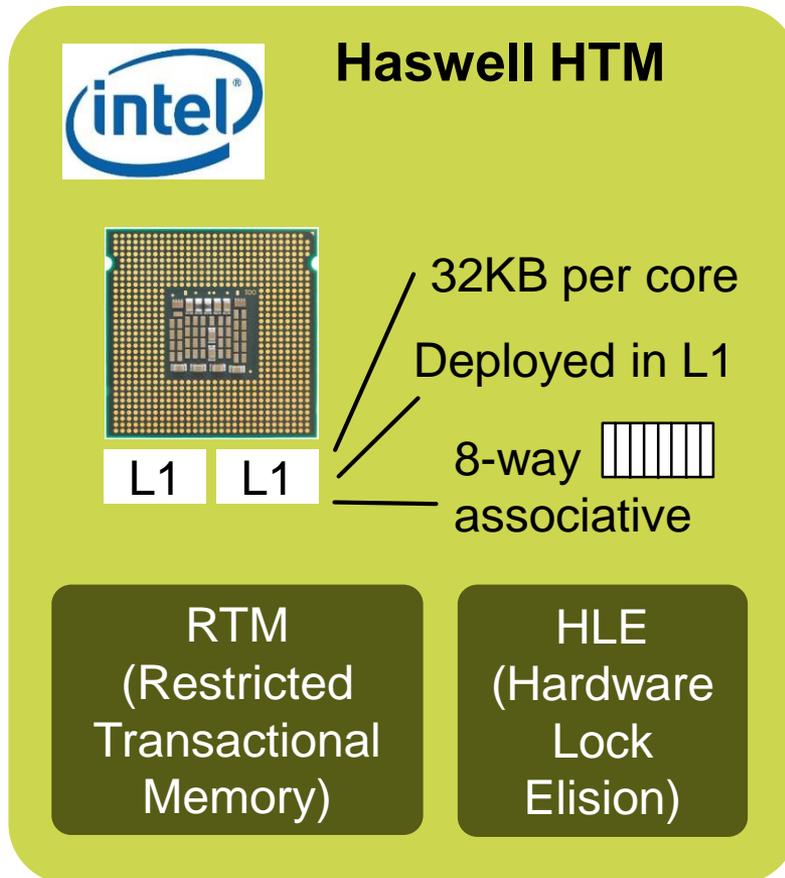
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



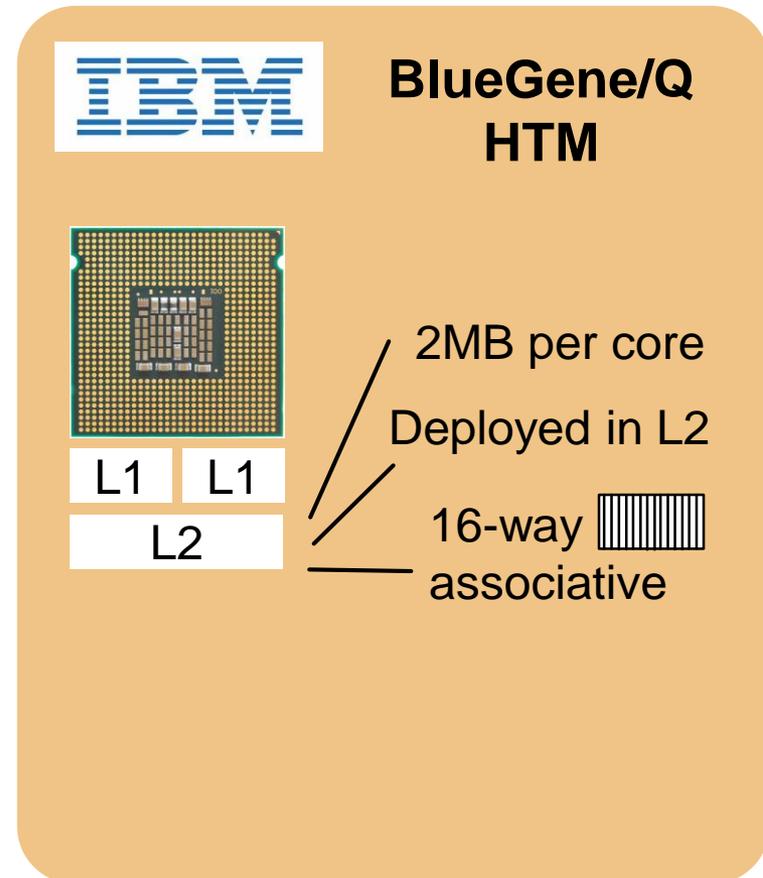
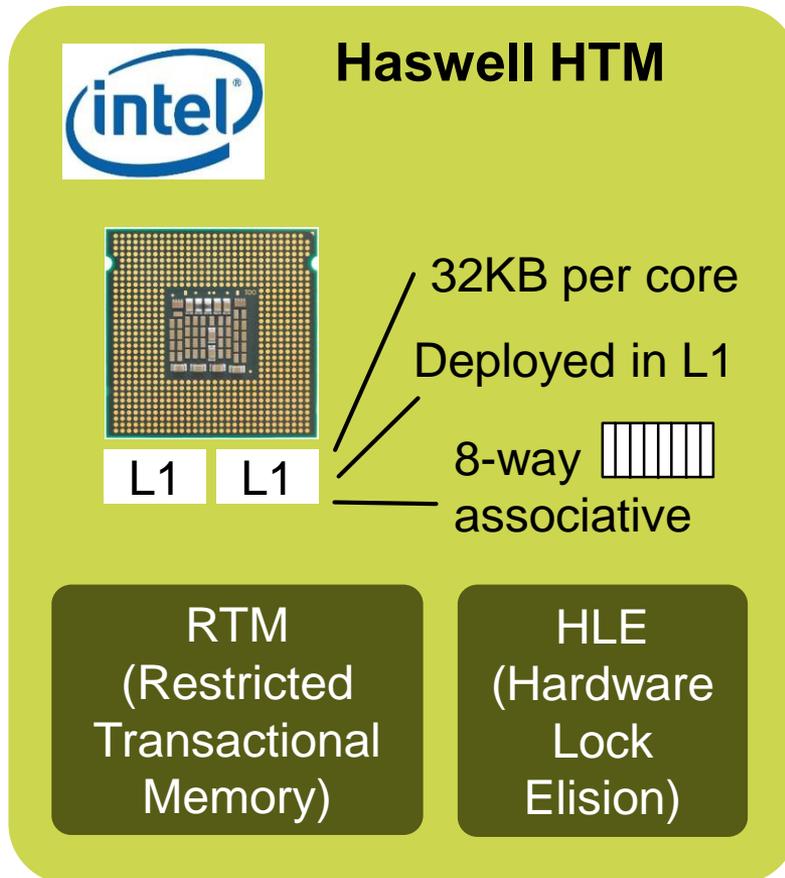
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



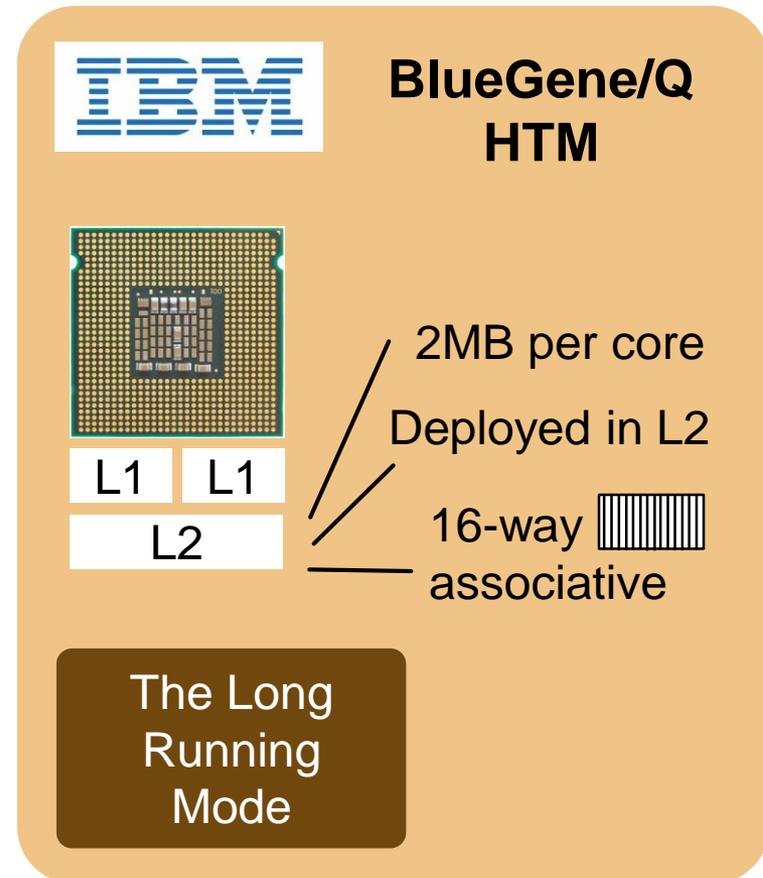
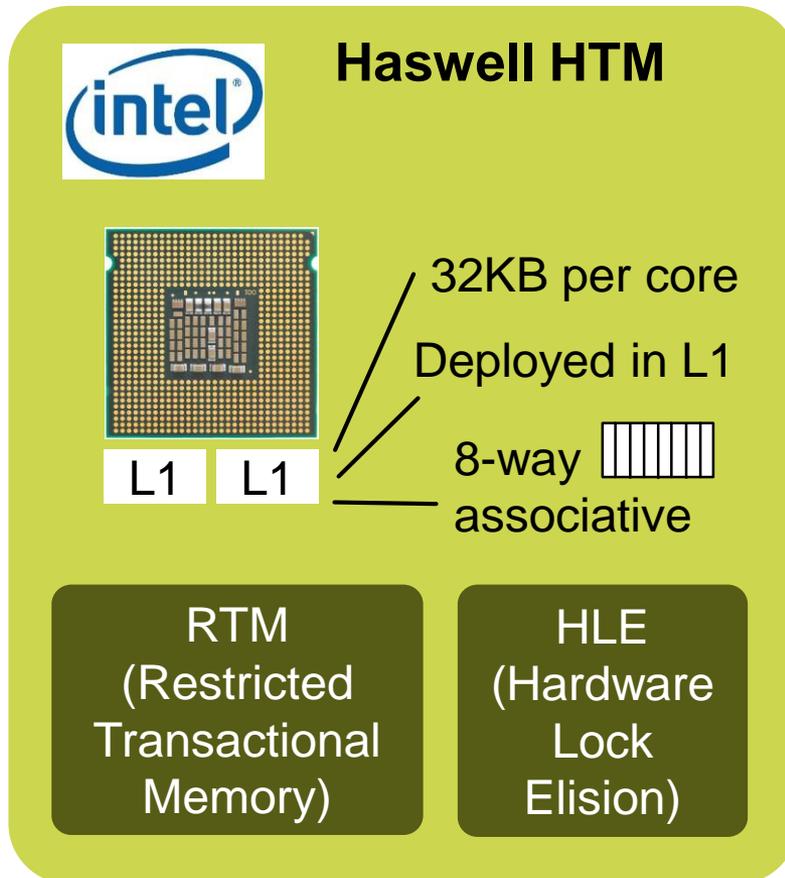
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS



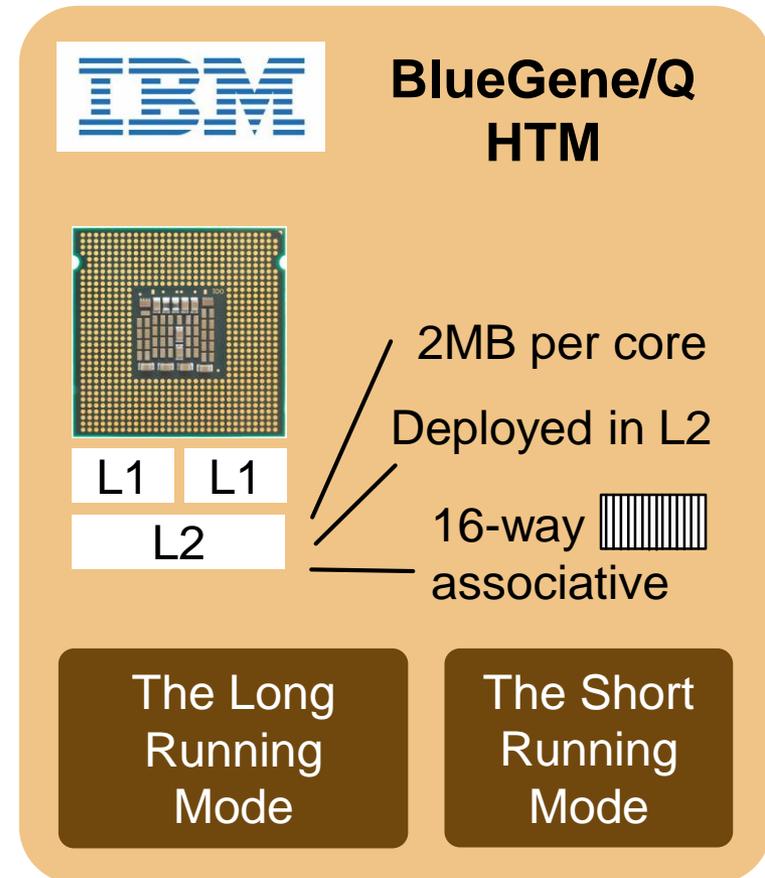
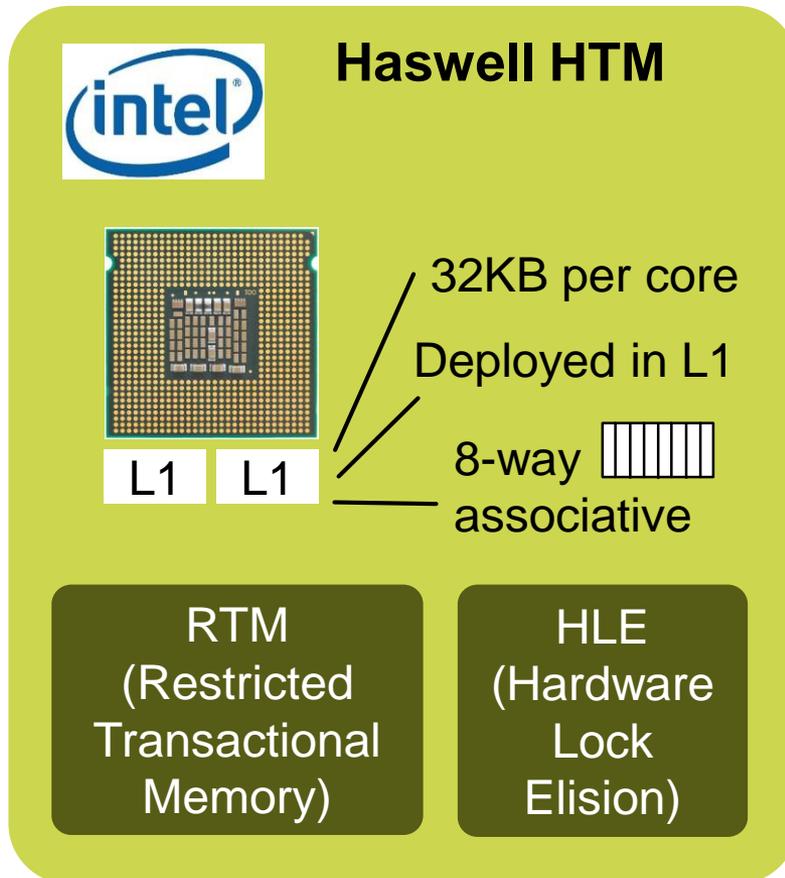
# PERFORMANCE ANALYSIS

## CONSIDERED MECHANISMS

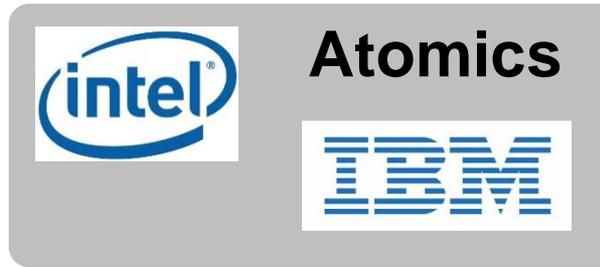


# PERFORMANCE ANALYSIS

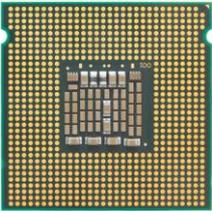
## CONSIDERED MECHANISMS



# PERFORMANCE ANALYSIS CONSIDERED MECHANISMS



### Haswell HTM



32KB per core  
Deployed in L1

L1 | L1

8-way associative 

RTM  
(Restricted Transactional Memory)

HLE  
(Hardware Lock Elision)

### BlueGene/Q HTM



2MB per core  
Deployed in L2

L1 | L1  
L2

16-way associative 

The Long Running Mode

The Short Running Mode

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

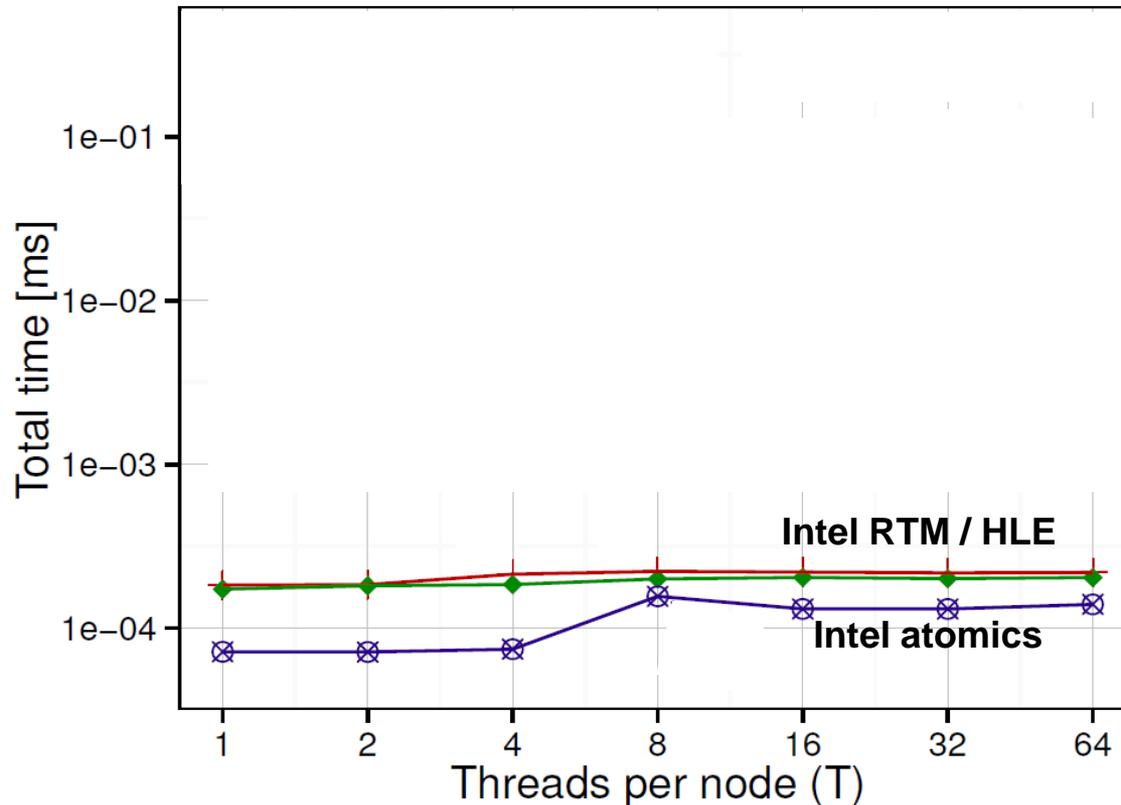
```
// start handler  
if(!v.visited) {  
    v.visited = 1;  
}  
// finish handler
```

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 racing accesses/vertex)



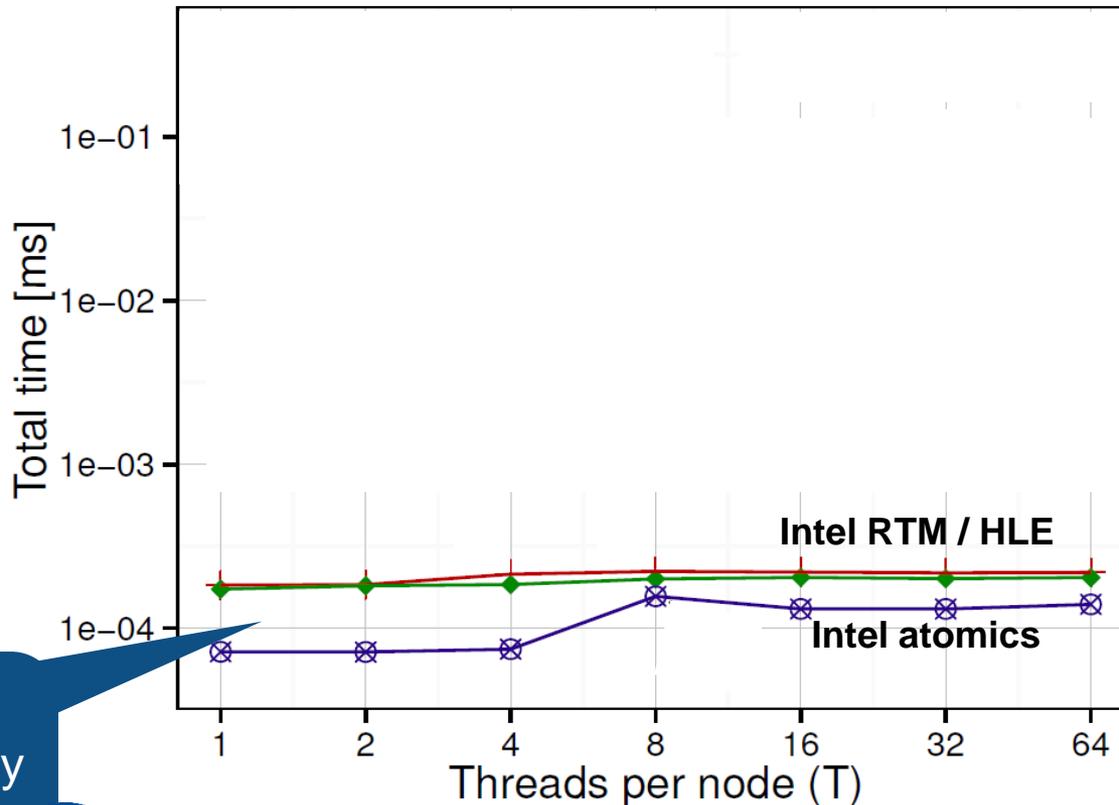
```
// start handler  
if(!v.visited) {  
    v.visited = 1;  
}  
// finish handler
```

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 racing accesses/vertex)



```
// start handler  
if(!v.visited) {  
    v.visited = 1;  
}  
// finish handler
```

Atomics  
(CAS) slightly  
faster than  
HTM

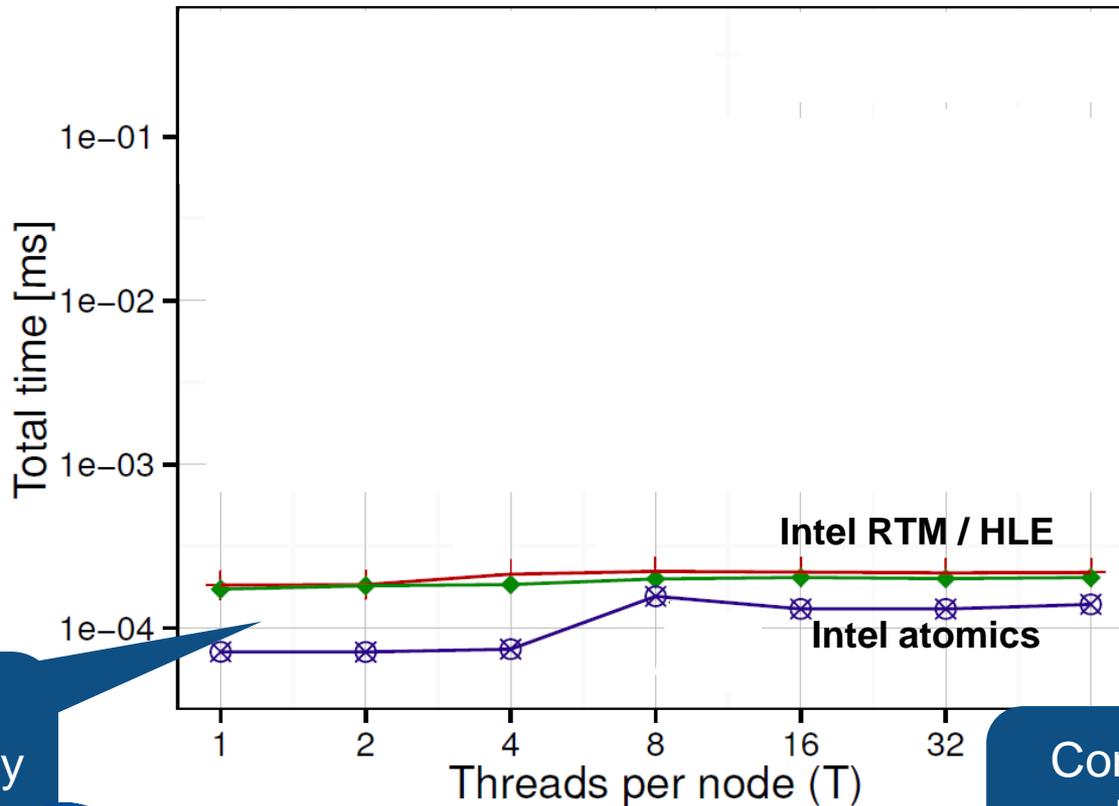


# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 racing accesses/vertex)



```
// start handler  
if(!v.visited) {  
    v.visited = 1;  
}  
// finish handler
```

Atomics  
(CAS) slightly  
faster than  
HTM

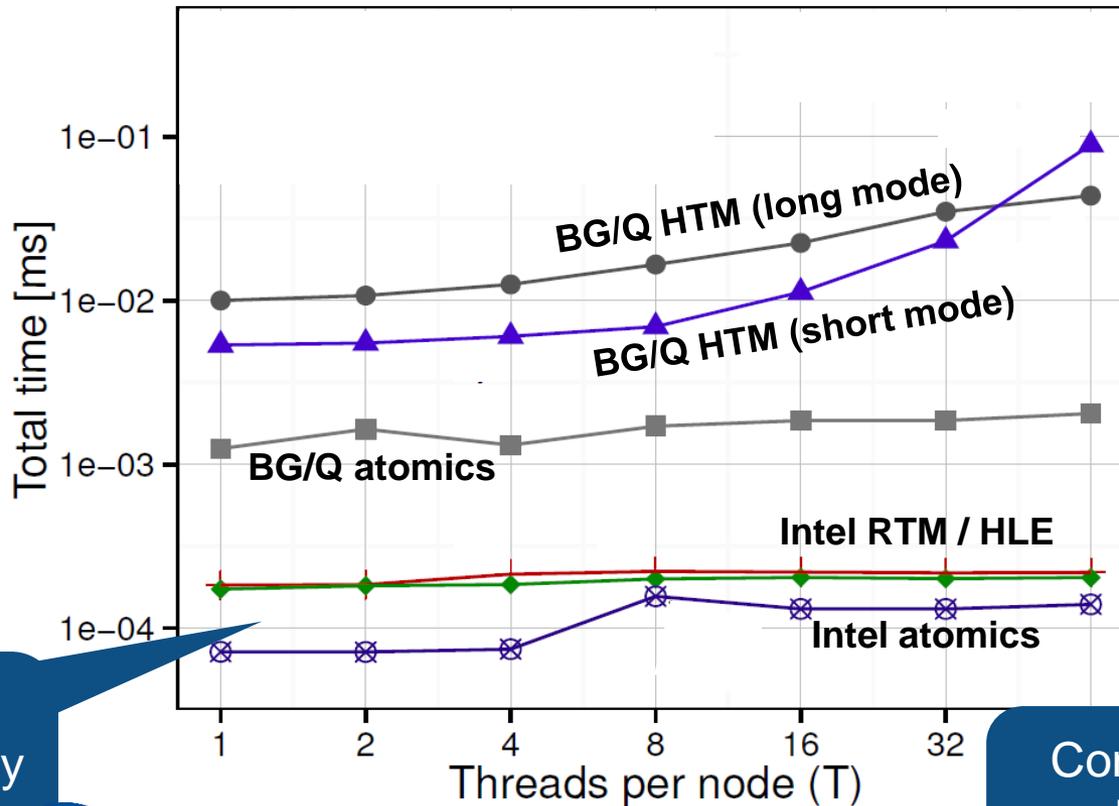
Commit  
overheads  
dominate

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 racing accesses/vertex)



```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```

Atomics  
(CAS) slightly  
faster than  
HTM

Commit  
overheads  
dominate

# SINGLE-VERTEX TRANSACTIONS

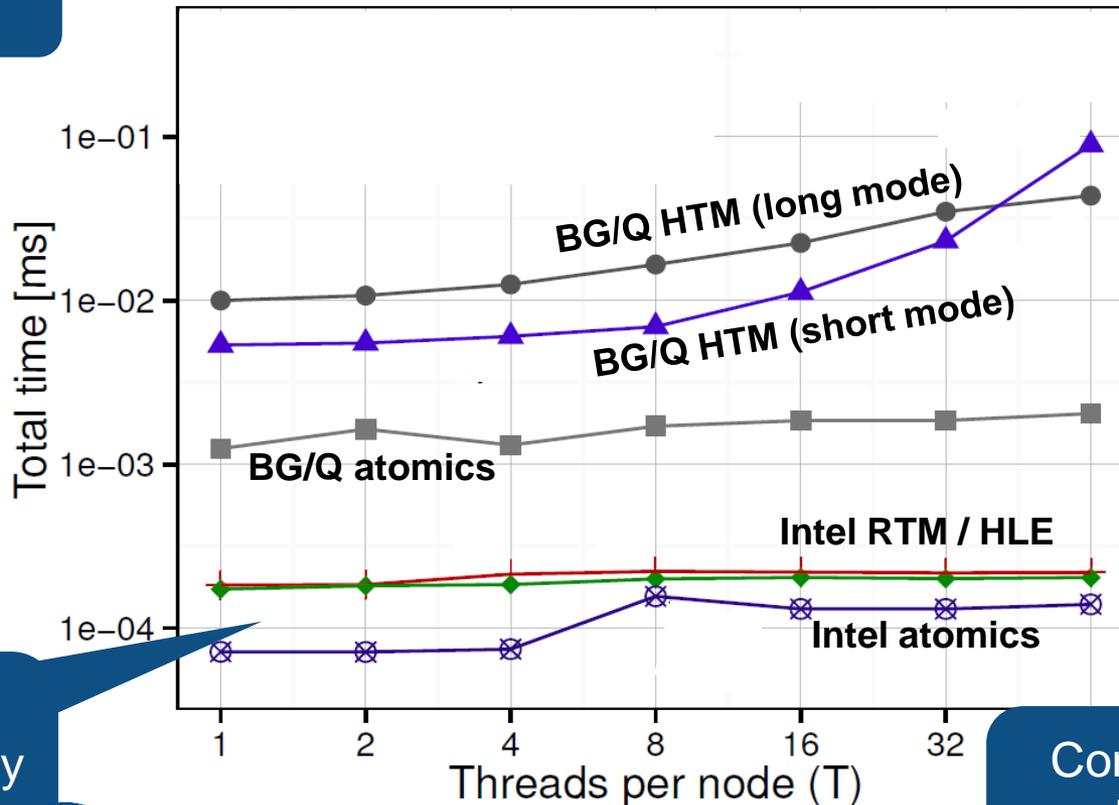
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Very few  
aborts

Lower contention  
(10 racing accesses/vertex)

```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```



Atomics  
(CAS) slightly  
faster than  
HTM

Commit  
overheads  
dominate

# SINGLE-VERTEX TRANSACTIONS

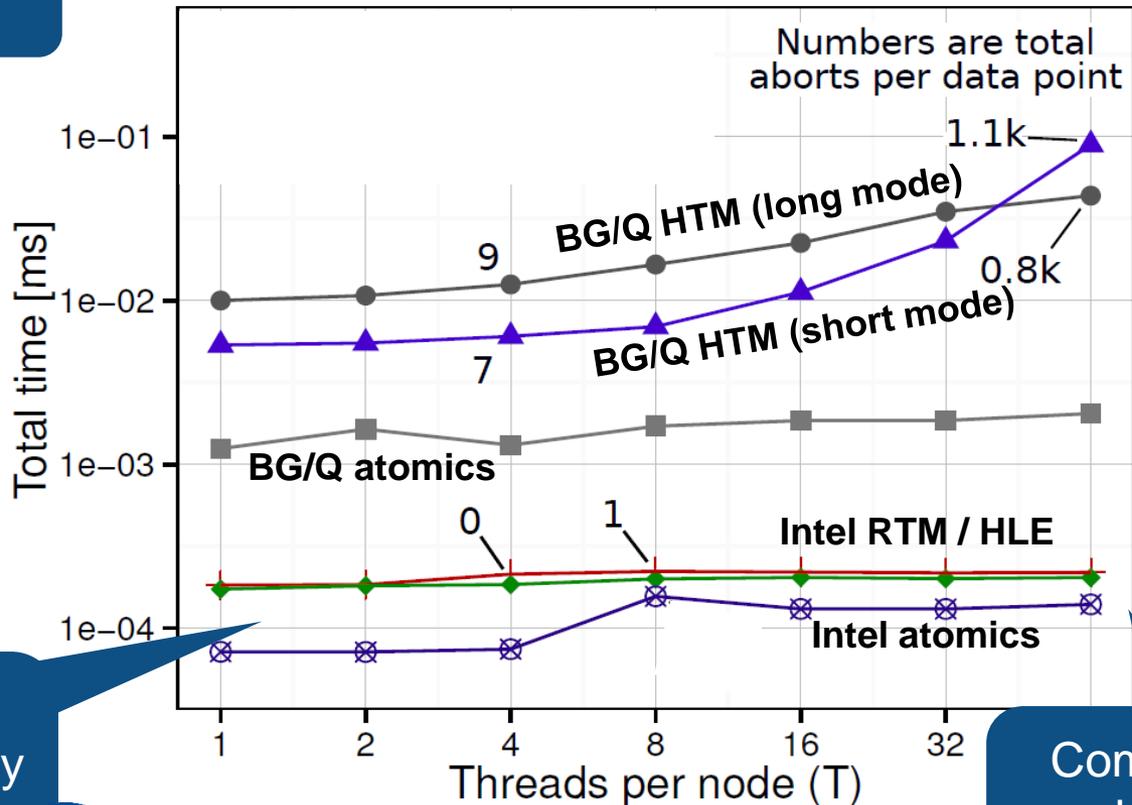
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Very few  
aborts

Lower contention  
(10 racing accesses/vertex)

```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```



Atomics  
(CAS) slightly  
faster than  
HTM

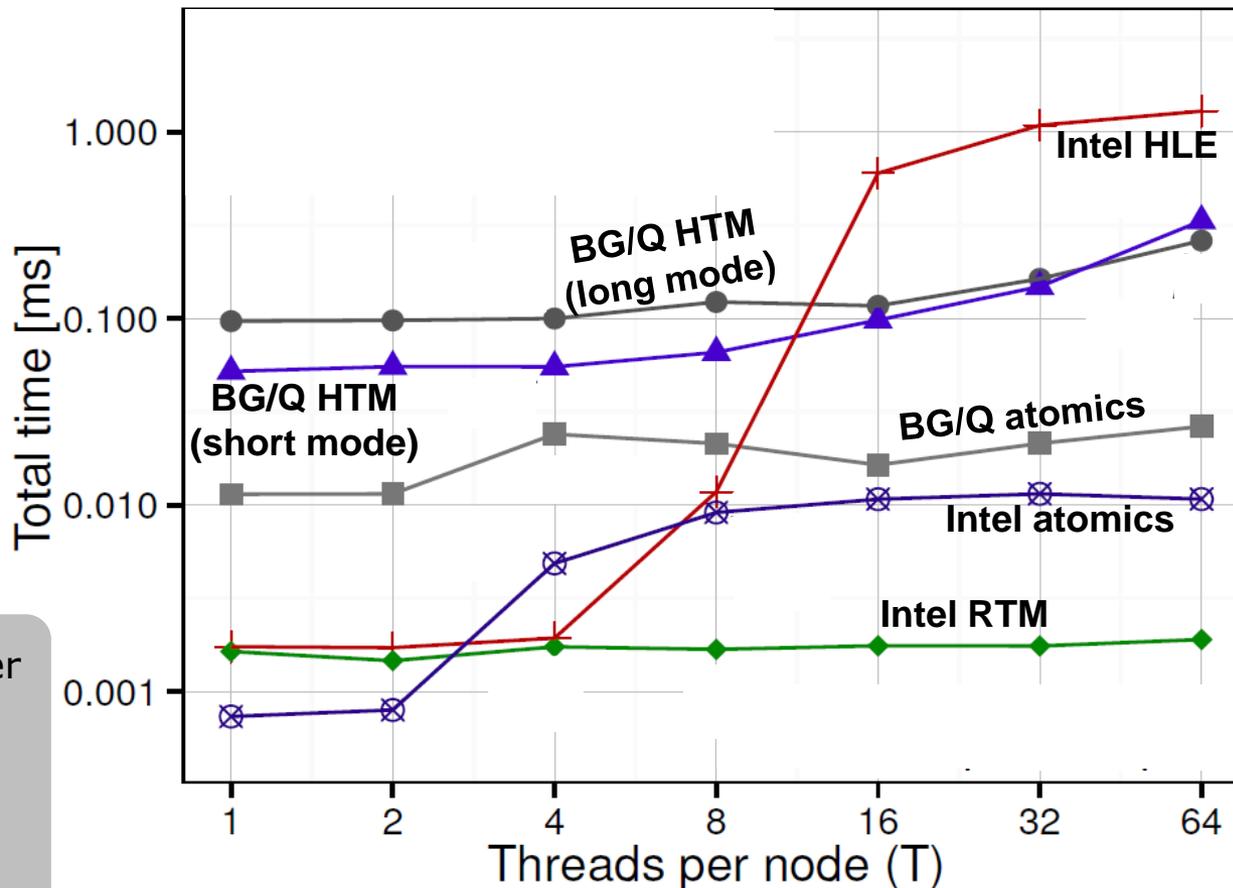
Commit  
overheads  
dominate

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Higher contention  
(100 racing accesses/vertex)



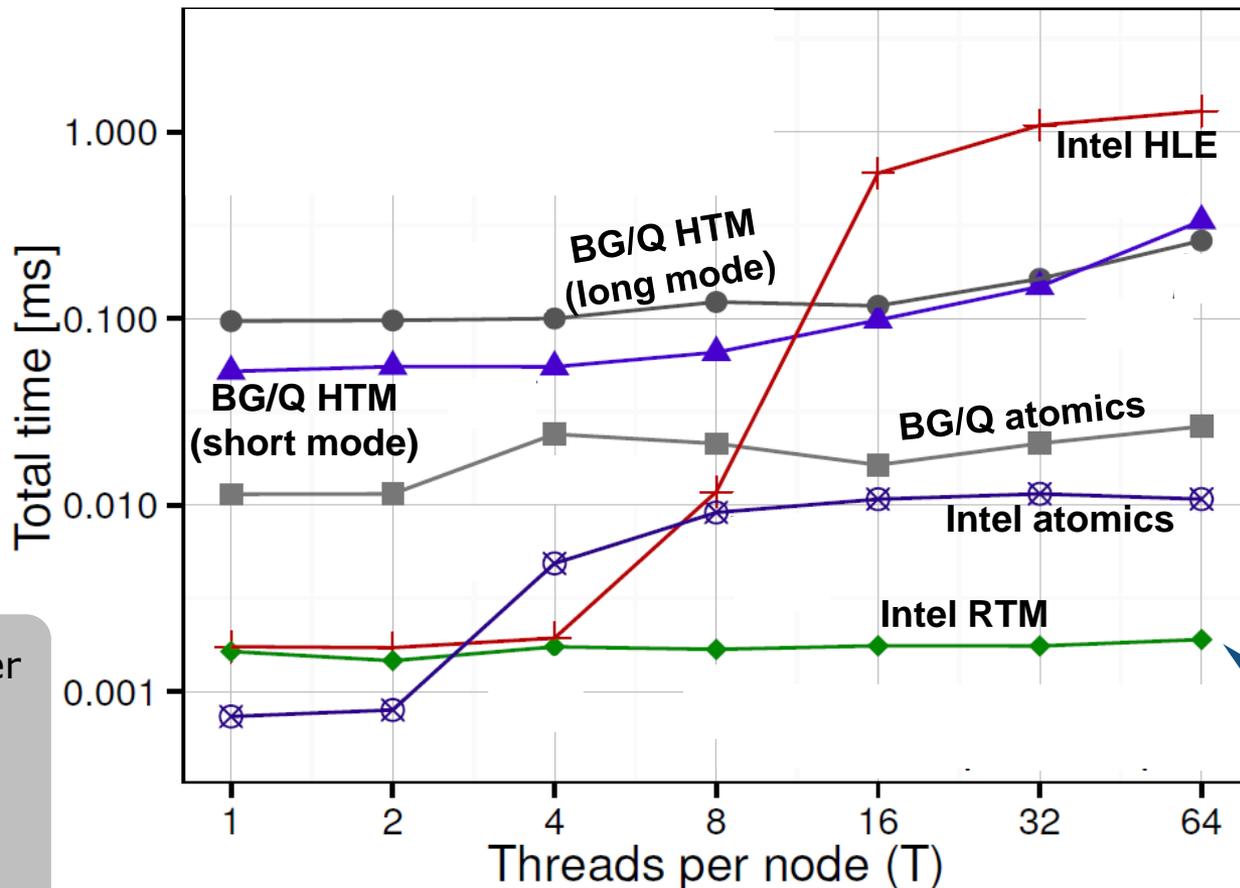
```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish
handler
```

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Higher contention  
(100 racing accesses/vertex)



```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```

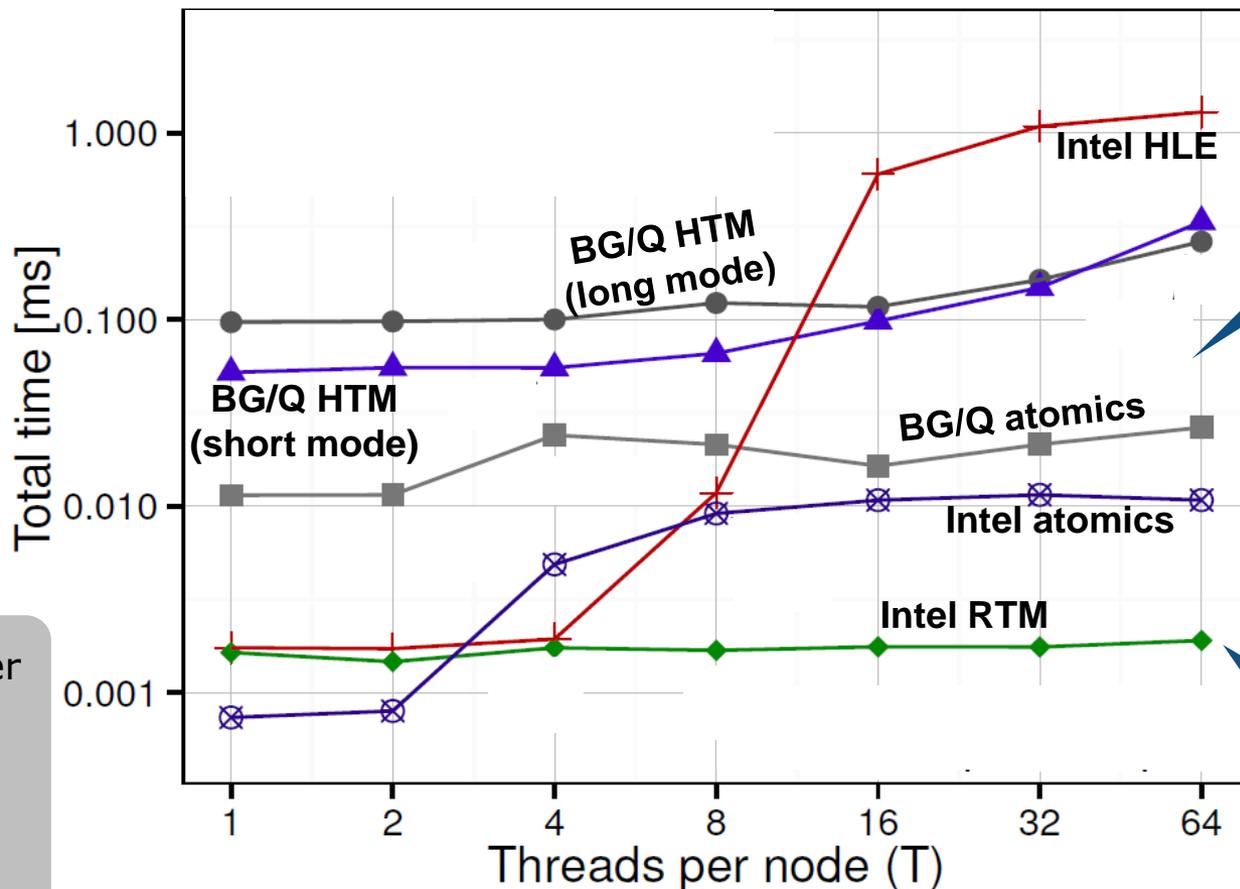
!  
RTM  
better  
than  
atomics

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Higher contention  
(100 racing accesses/vertex)



! BG/Q HTM still worse (L1 vs L2 matters!)

! RTM better than atomics

```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```

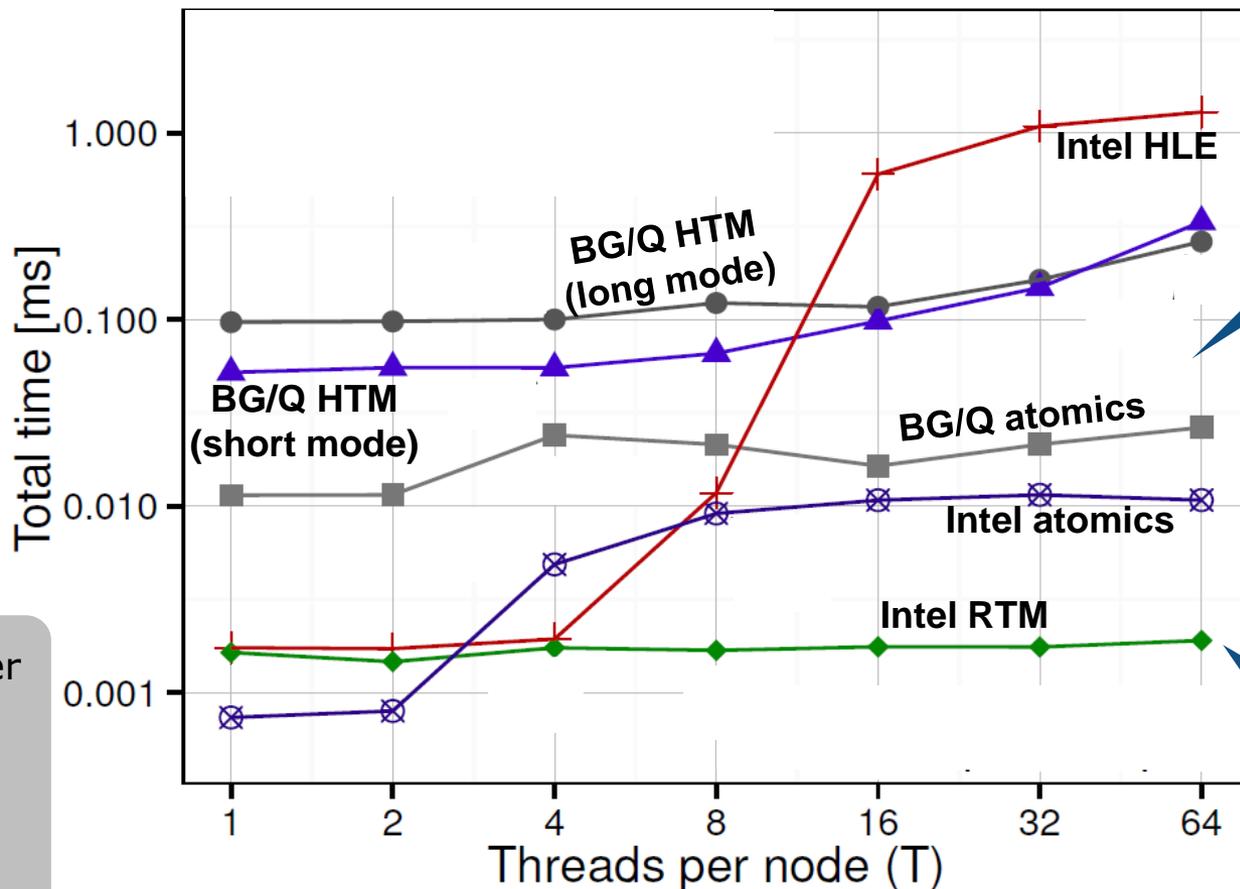
# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Higher contention  
(100 racing accesses/vertex)

Still very  
few  
aborts



! BG/Q  
HTM still  
worse (L1  
vs L2  
matters!)

! RTM  
better  
than  
atomics

```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish
handler
```

# SINGLE-VERTEX TRANSACTIONS

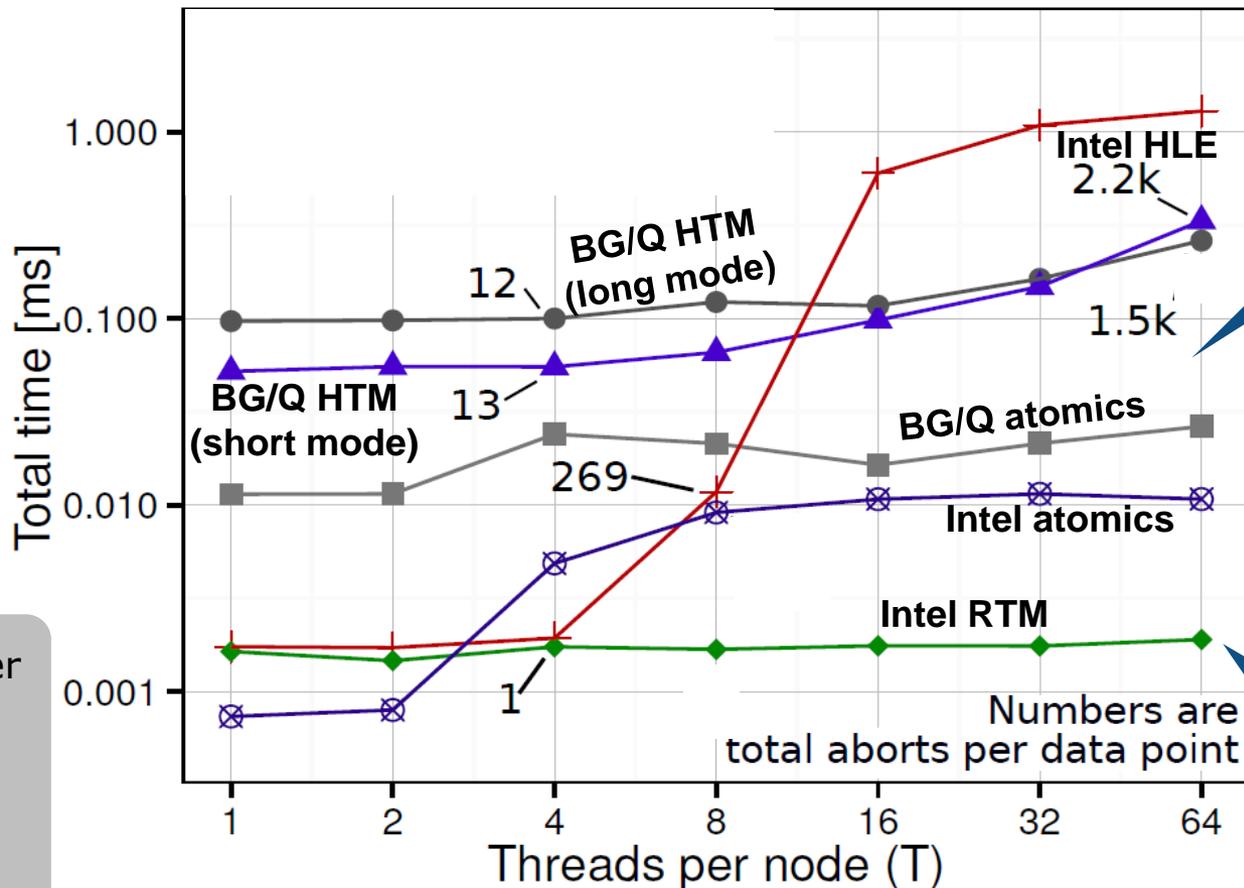
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Still very few aborts

Higher contention  
(100 racing accesses/vertex)

BG/Q HTM still worse (L1 vs L2 matters!)

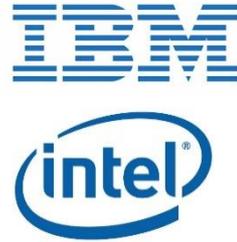


```
// start handler
if(!v.visited) {
  v.visited = 1;
}
// finish handler
```

RTM better than atomics

# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

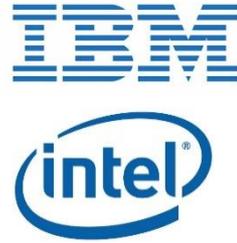
Used in  
PageRank



```
// start handler  
v.rank++;  
// finish handler
```

# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

Used in  
PageRank



```
// start handler  
v.rank++;  
// finish handler
```



Atomics always  
outperform HTM

# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

Used in  
PageRank



```
// start handler  
v.rank++;  
// finish handler
```



Atomics always  
outperform HTM



The reason: each transaction always modifies some  
memory cell, increasing the number of conflicts

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$   
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$   
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Startup  
overheads

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$   
vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead  
per vertex

Startup  
overheads

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$  vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead  
per vertex

Startup  
overheads

Time to modify  $N$  vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$  vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead  
per vertex

Startup  
overheads

Time to modify  $N$  vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead  
per vertex

Startup  
overheads

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$  vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead per vertex

Startup overheads

Time to modify  $N$  vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead per vertex

Startup overheads

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$  vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead per vertex

Startup overheads

Time to modify  $N$  vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead per vertex

Startup overheads

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

**!** Transaction startup overheads dominate

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Time to modify  $N$  vertices with atomics:

$$T_{AT}(N) = A_{AT}N + B_{AT}$$

Overhead per vertex

Startup overheads

Time to modify  $N$  vertices with a transaction

$$T_{HTM}(N) = A_{HTM}N + B_{HTM}$$

Overhead per vertex

Startup overheads

We predict that:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

Transactions' cost grows slower

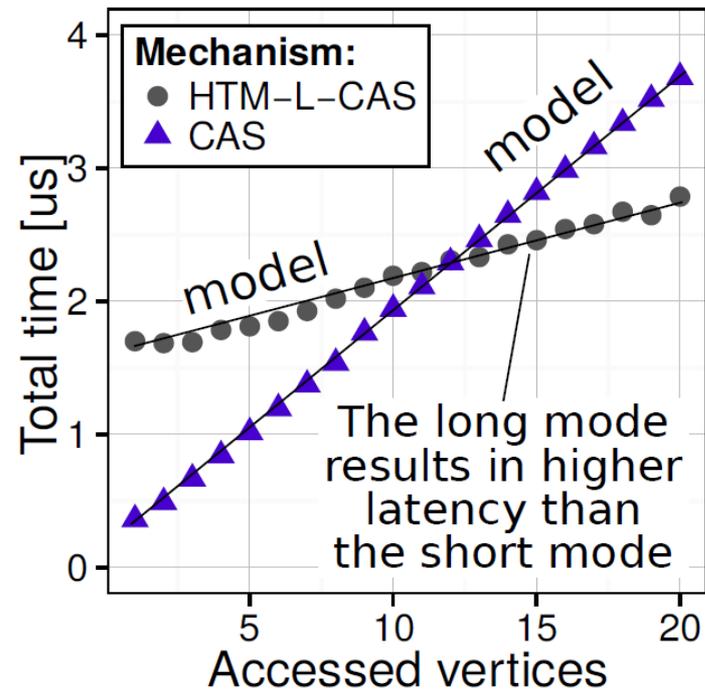
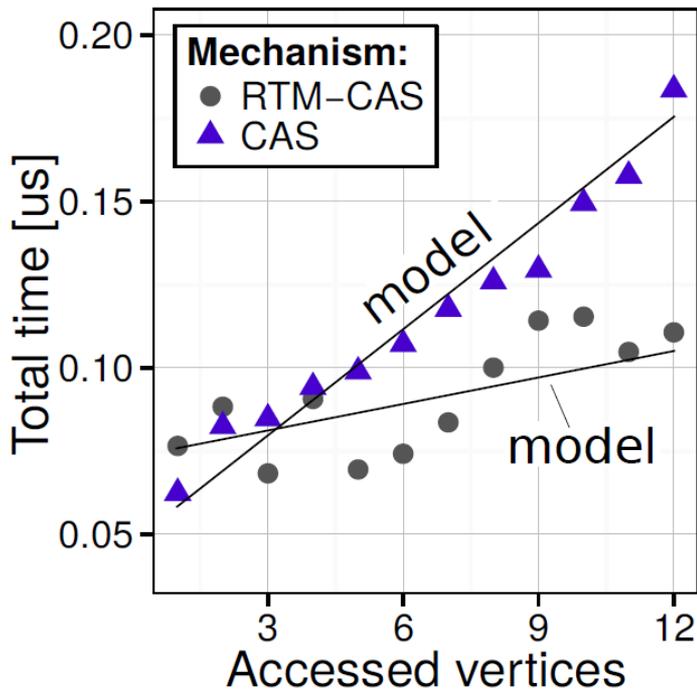
Transaction startup overheads dominate

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS



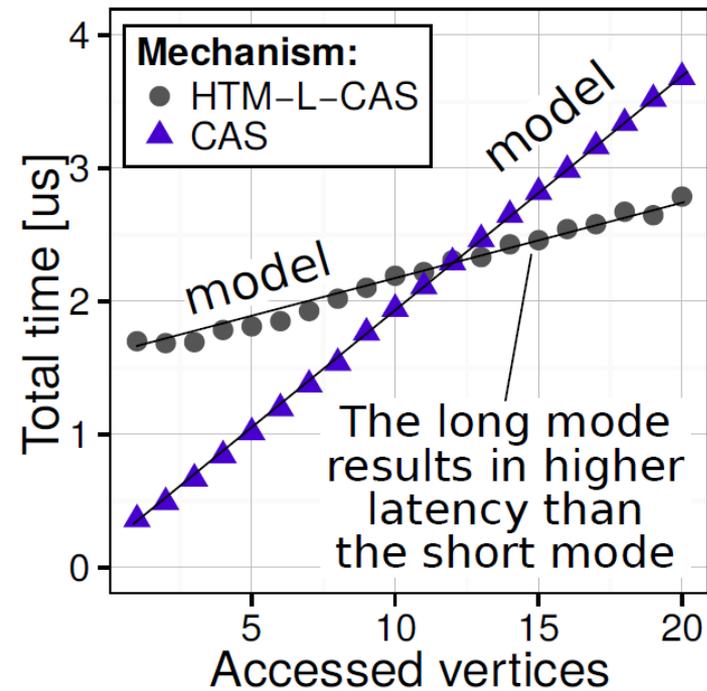
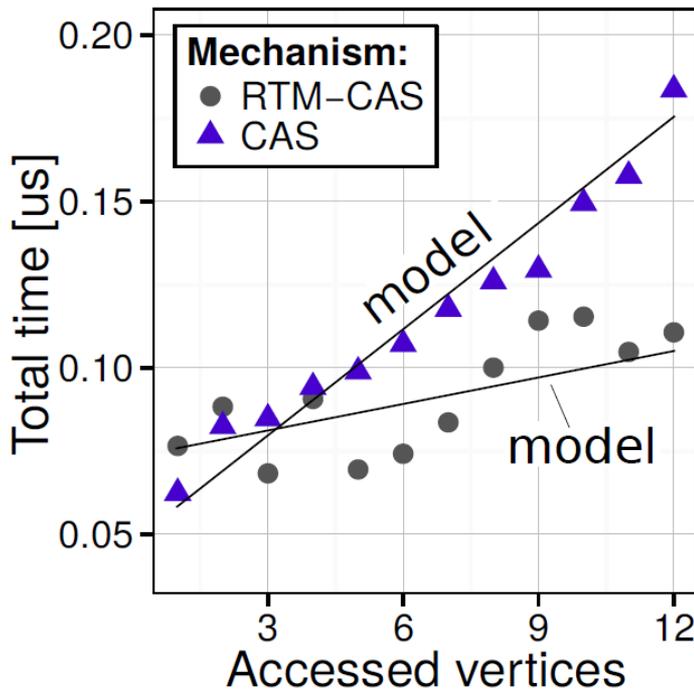
# PERFORMANCE MODEL

## ATOMICS VS TRANSACTIONS

Indeed:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$



# PERFORMANCE MODEL

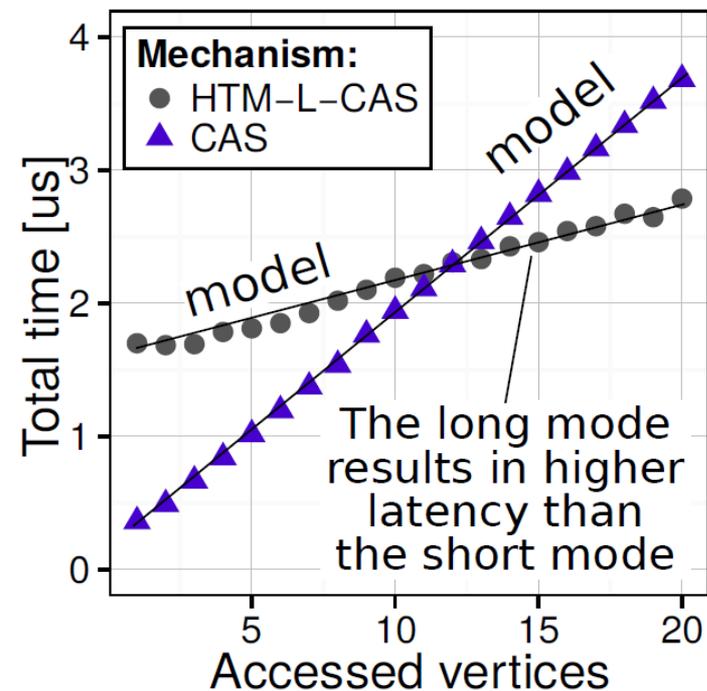
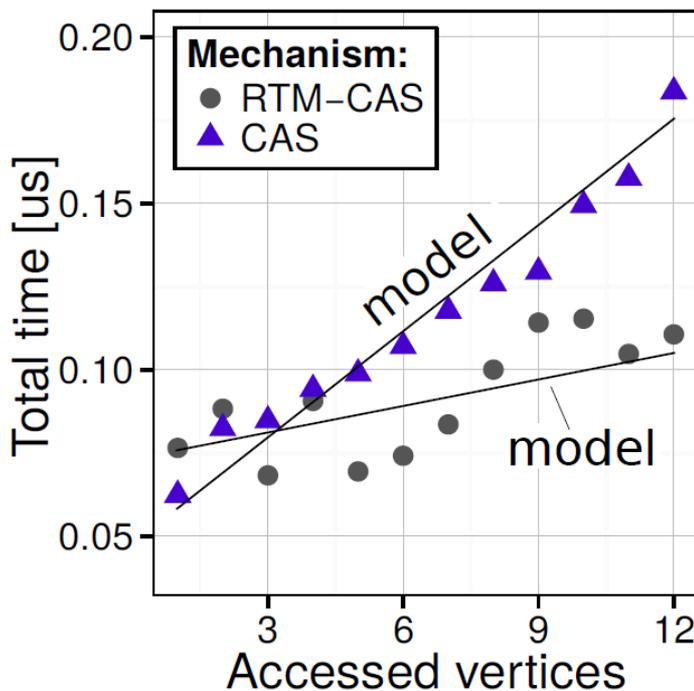
## ATOMICS VS TRANSACTIONS

Indeed:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$

- Can we amortize HTM startup/commit overheads with larger transaction sizes?



# PERFORMANCE MODEL

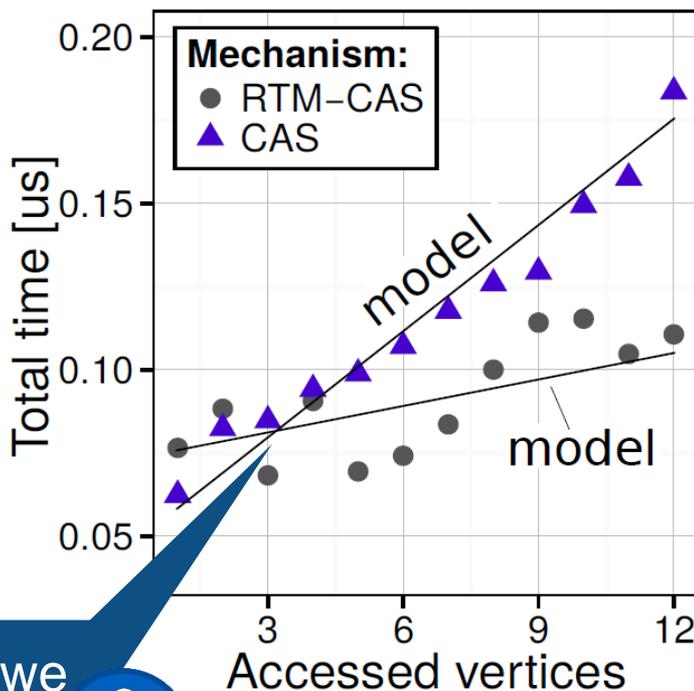
## ATOMICS VS TRANSACTIONS

- Can we amortize HTM startup/commit overheads with larger transaction sizes?

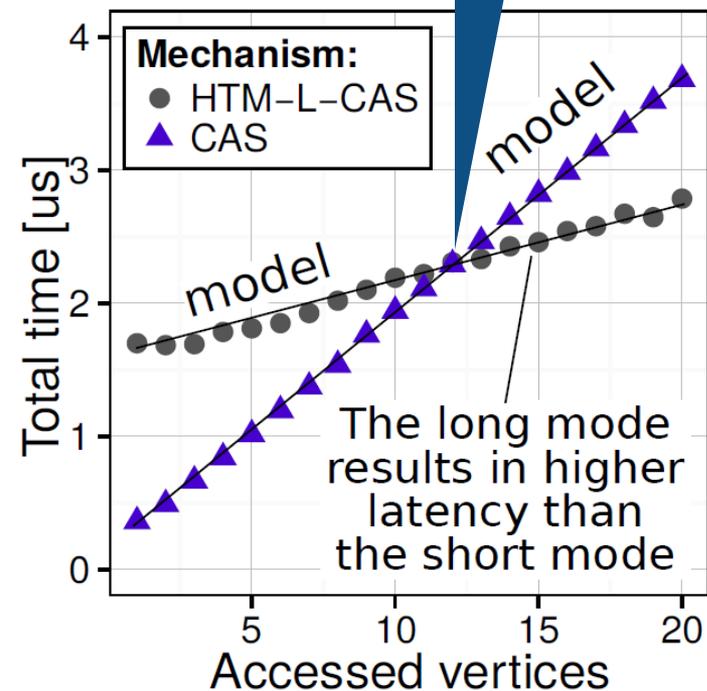
Indeed:

$$B_{AT} < B_{HTM}$$

$$A_{AT} > A_{HTM}$$



Yes, we can!



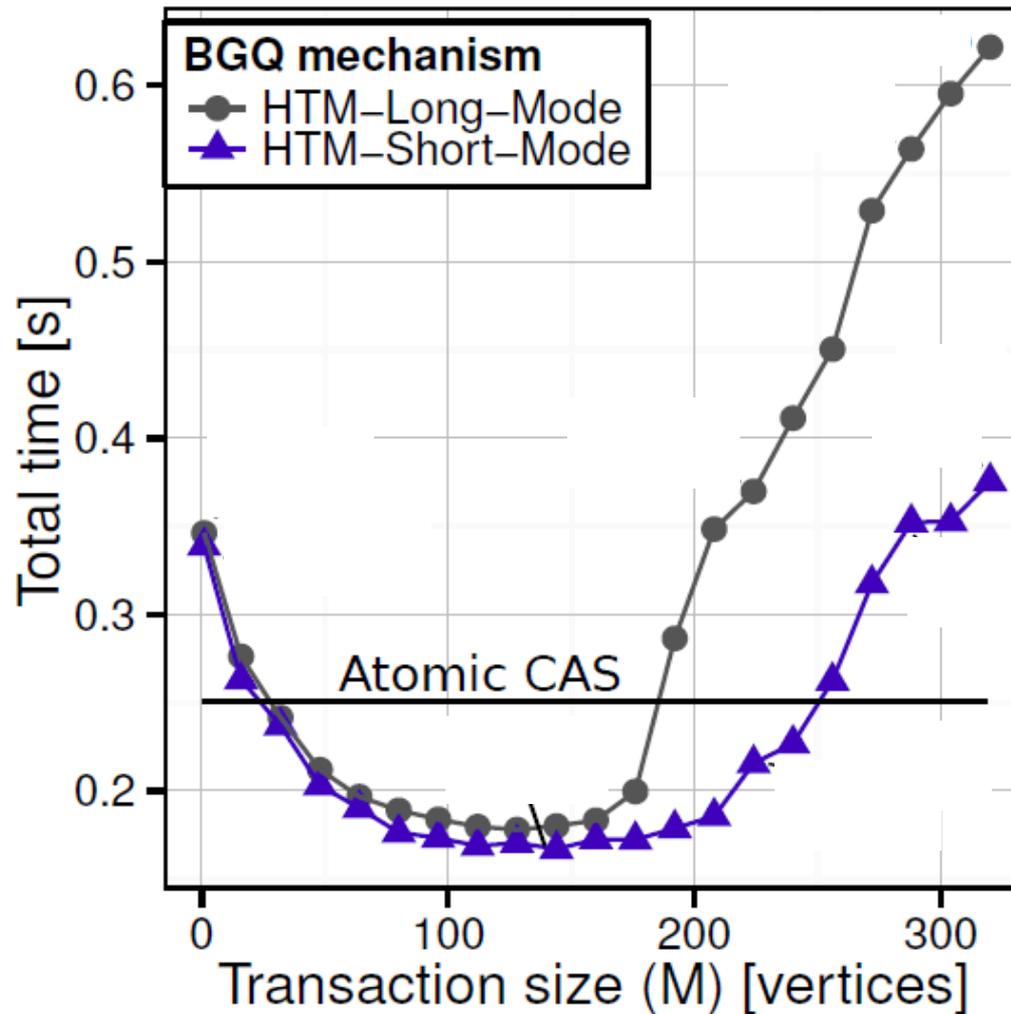
Yes, we can!

# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED

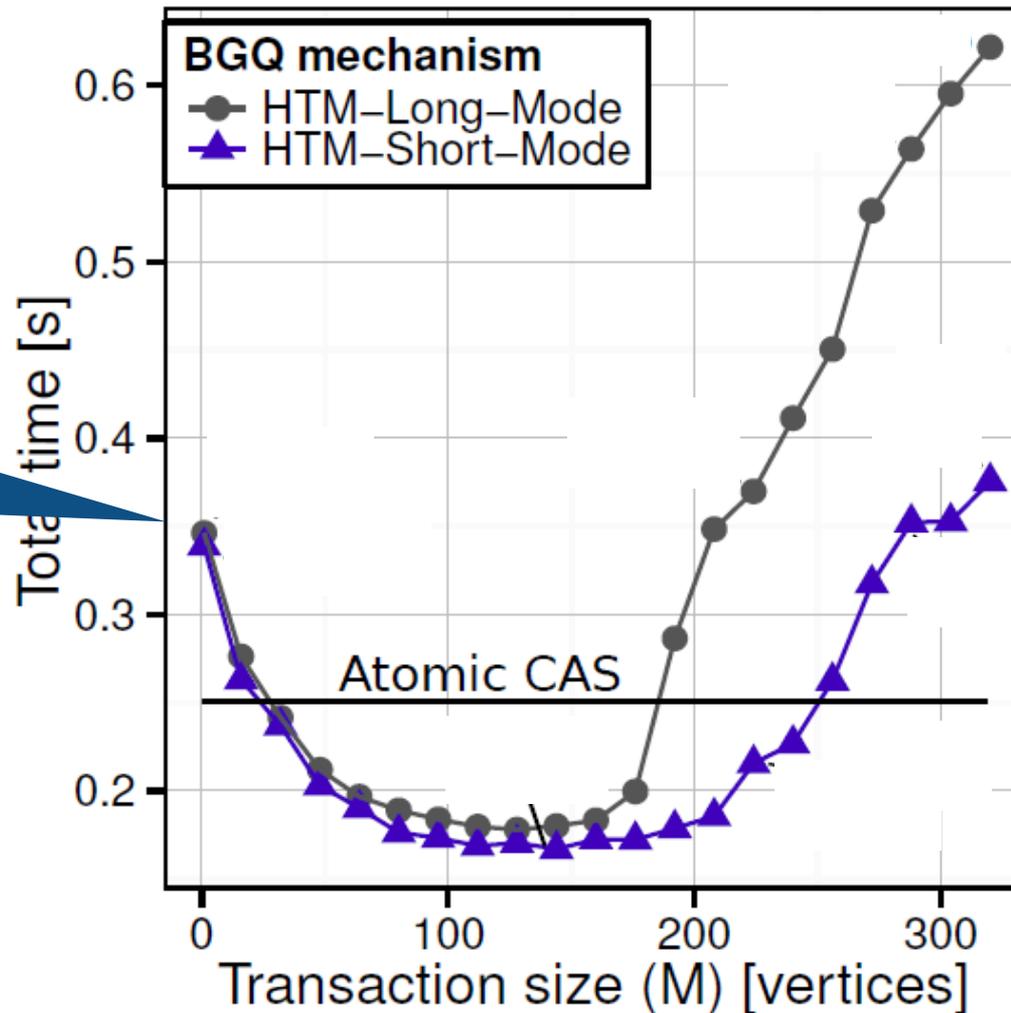
# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED



# MULTI-VERTEX TRANSACTIONS

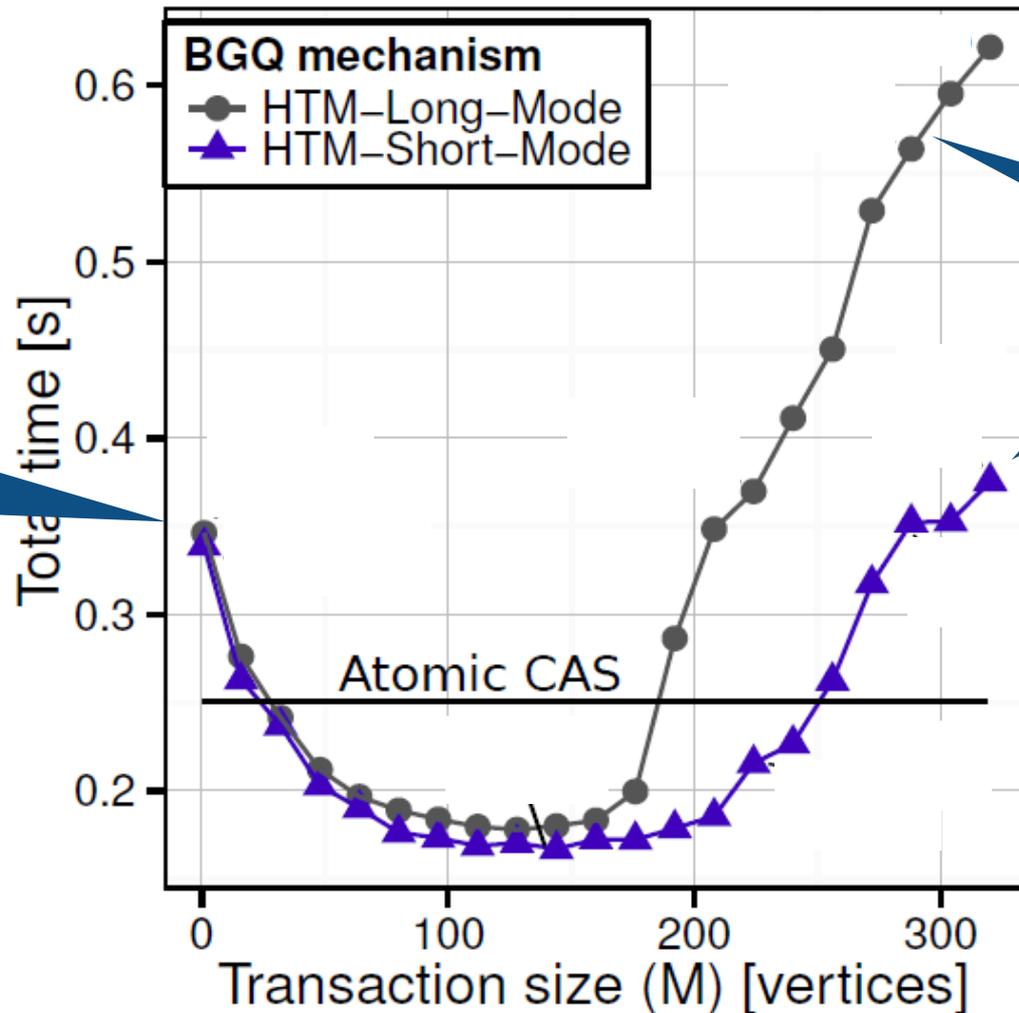
## MARKING VERTICES AS VISITED



! Startup and commit overheads

# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED

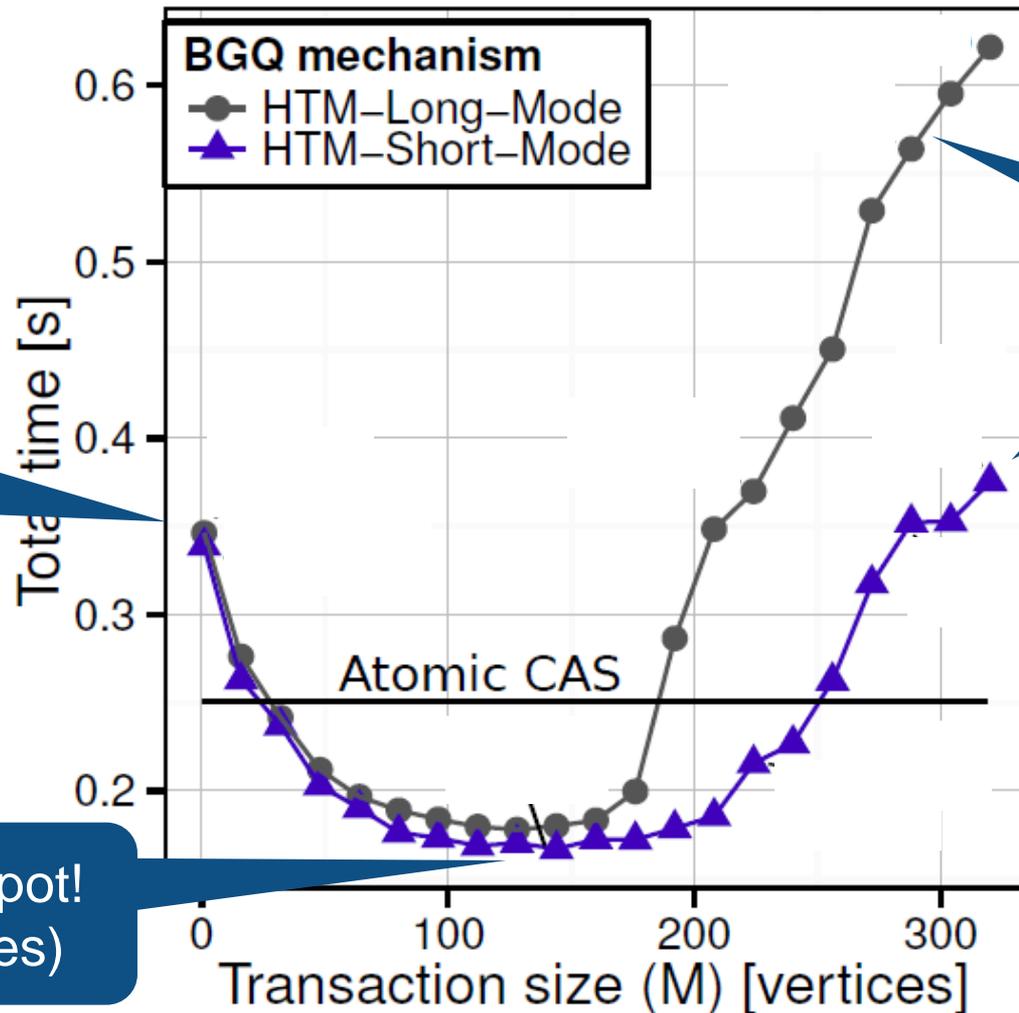


Startup and  
commit  
overheads

Abort and  
rollback  
overheads

# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED



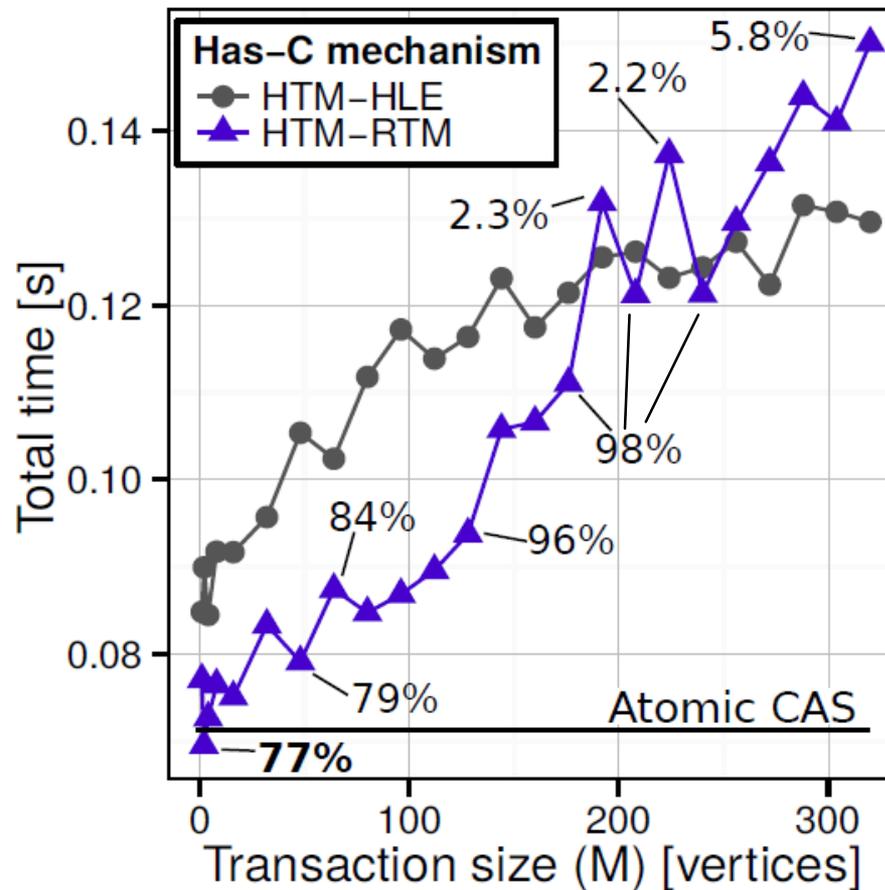
! Startup and commit overheads

! The sweetspot!  
(144 vertices)

! Abort and rollback overheads

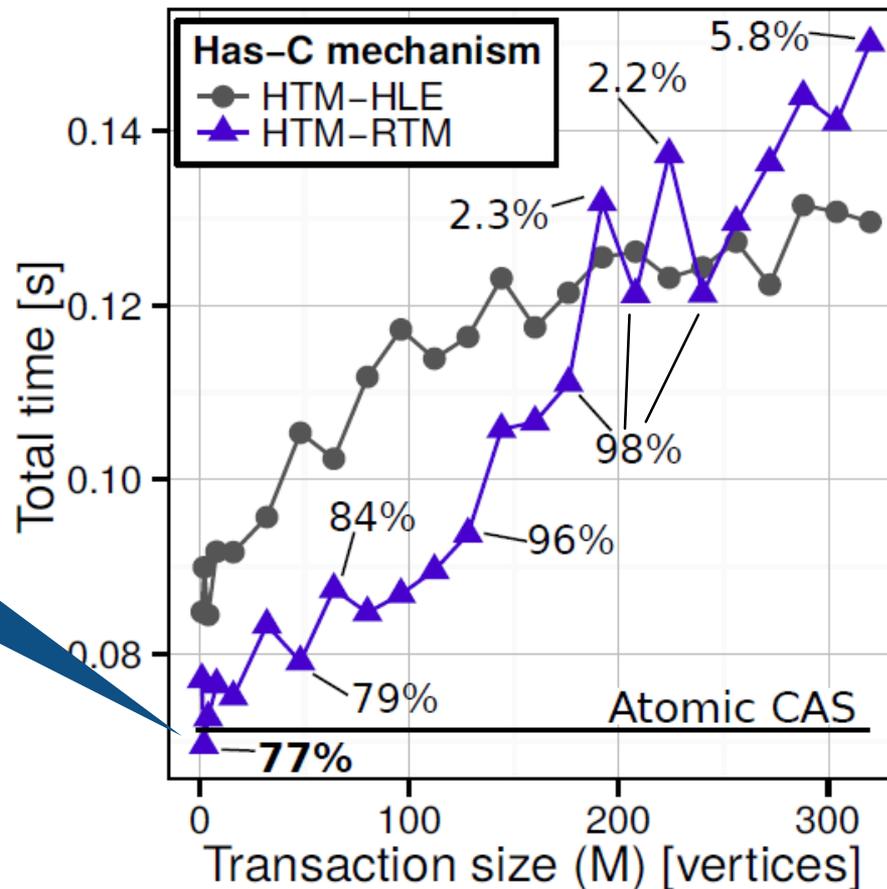
# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED



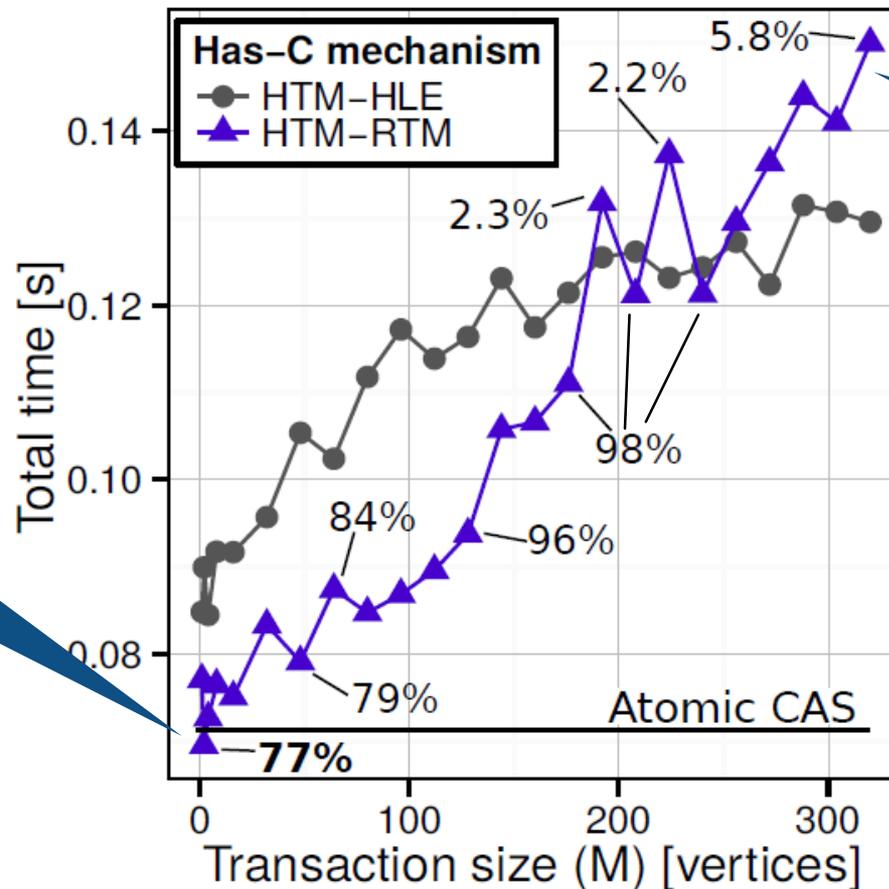
# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED



# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED



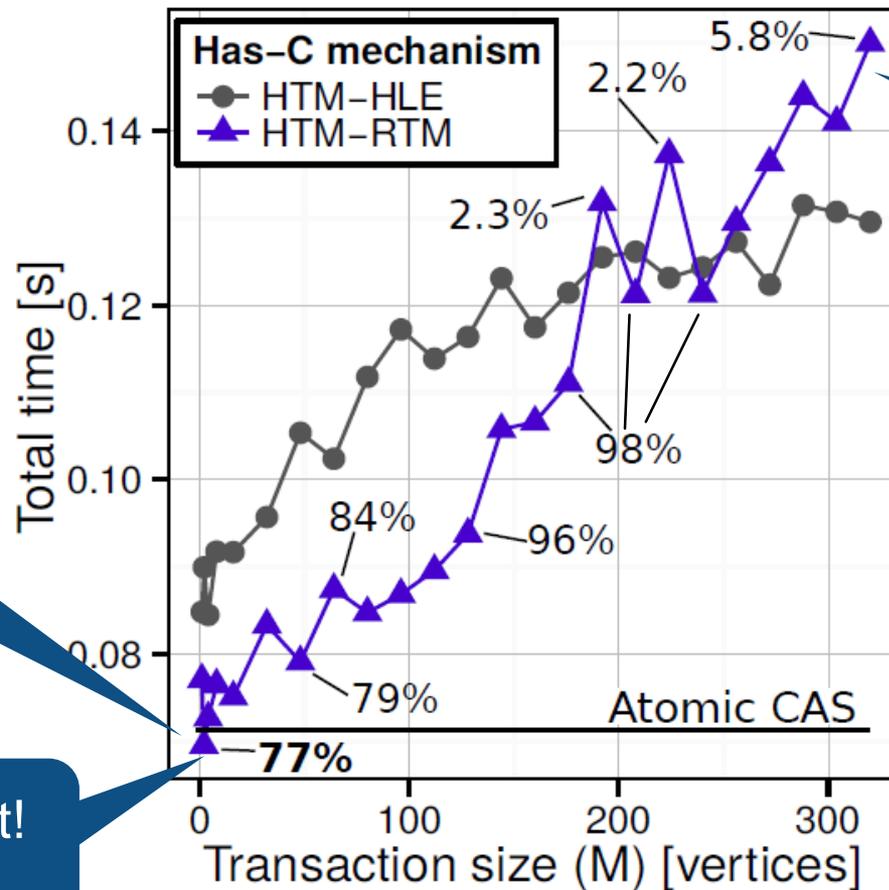
Startup and  
commit  
overheads

Abort and  
rollback  
overheads

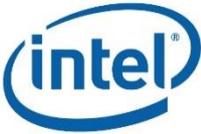
Atomic CAS

# MULTI-VERTEX TRANSACTIONS

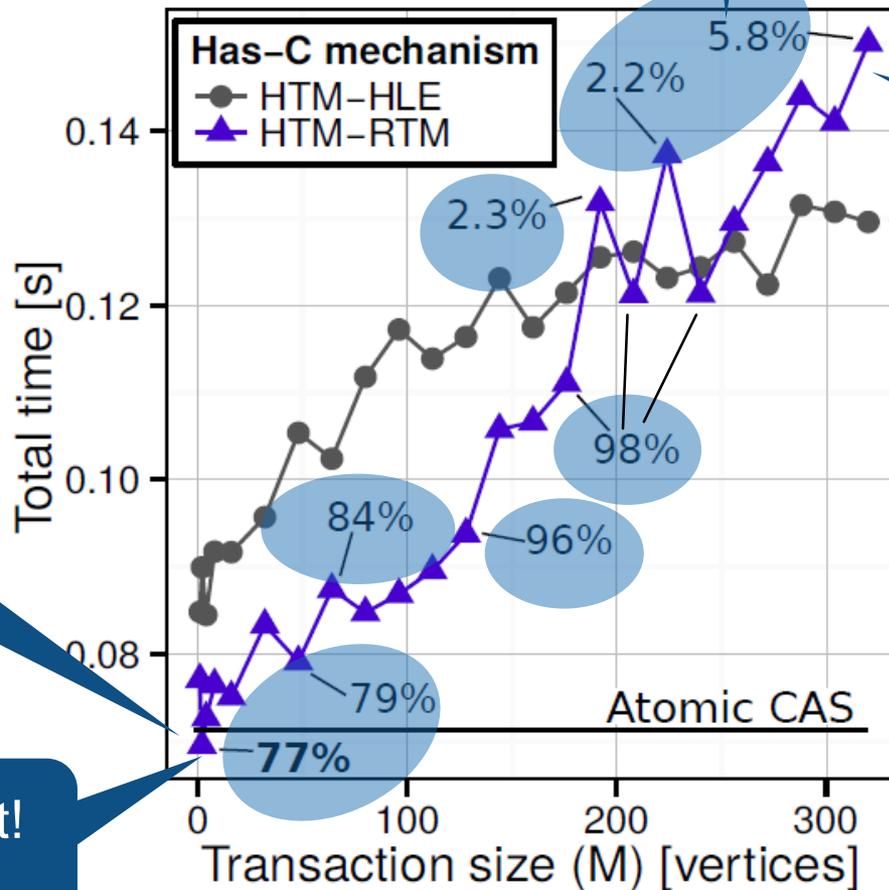
## MARKING VERTICES AS VISITED



# MULTI-VERTEX TRANSACTIONS MARKING VERTICES AS VISITED



Numbers: % of aborts due to HTM capacity overflows



Startup and commit overheads

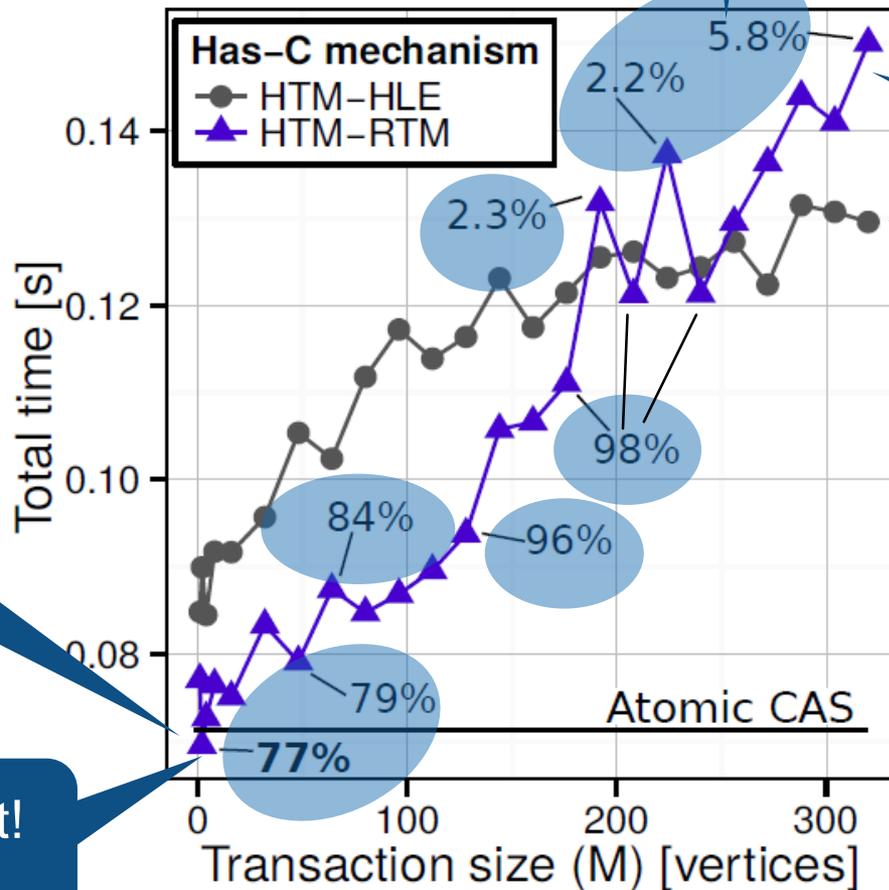
The sweetspot! (2 vertices)

Abort and rollback overheads

# MULTI-VERTEX TRANSACTIONS MARKING VERTICES AS VISITED



Numbers: % of aborts due to HTM capacity overflows



Startup and commit overheads

The sweetspot! (2 vertices)

Abort and rollback overheads

Majority of aborts are due to HTM capacity overflows (small cache size & associativity)

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



How can we implement AAM handlers most effectively?



What are performance tradeoffs related to HTM?



What are advantages of HTM over atomics for AAM?



What are the optimal transaction sizes?  
Can we amortize transaction overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



What are  
advantages of  
HTM over  
atomics for  
AAM?



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



For some algorithms  
(BFS) HTM is better



What are  
performance  
tradeoffs related  
to HTM?



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



For some algorithms  
(BFS) HTM is better



For others  
(PageRank)  
atomics give more  
performance



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



For some algorithms  
(BFS) HTM is better

„May fail”



For others  
(PageRank)  
atomics give more  
performance



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



For some algorithms  
(BFS) HTM is better

„May fail”



For others  
(PageRank)  
atomics give more  
performance

„Always  
succeed”



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



For some algorithms  
(BFS) HTM is better

„May fail”



For others  
(PageRank)  
atomics give more  
performance

„Always  
succeed”



AAM establishes a  
whole hierarchy of  
algorithms; check  
the paper 😊



What are the optimal  
transaction sizes?  
Can we amortize  
transaction  
overheads?

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED



„It really depends” 😊.  
But... There are some  
regularities



What are  
performance  
tradeoffs related  
to HTM?



For some algorithms  
(BFS) HTM is better

„May fail”



AAM establishes a  
whole hierarchy of  
algorithms; check  
the paper 😊



For others  
(PageRank)  
atomics give more  
performance

„Always  
succeed”



Size for BG/Q ~100  
>  
Size for Haswell ~10



Yes, we can

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED

! „It really depends” 😊.  
But... There are some regularities

? What are performance tradeoffs related to HTM?

! For some algorithms (BFS) HTM is better

„May fail”

AAM establishes a whole hierarchy of algorithms; check the paper 😊

! For others (PageRank) atomics give more performance

„Always succeed”

Same for other graphs

! Size for BG/Q ~100  
>  
Size for Haswell ~10

! Yes, we can

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED

! „It really depends” 😊.  
But... There are some regularities

! Larger cache & associativity → fewer aborts & more coarsening

! For some algorithms (BFS) HTM is better

„May fail”

! For others (PageRank) atomics give more performance

„Always succeed”

AAM establishes a whole hierarchy of algorithms; check the paper 😊

Same for other graphs

! Size for BG/Q ~100  
>  
Size for Haswell ~10

! Yes, we can

# PERFORMANCE ANALYSIS

## QUESTIONS ANSWERED

! „It really depends” 😊.  
But... There are some regularities

! Larger cache & associativity → fewer aborts & more coarsening

! Larger (L2) cache → higher latency

! For some algorithms (BFS) HTM is better

„May fail”

! For others (PageRank) atomics give more performance

„Always succeed”

AAM establishes a whole hierarchy of algorithms; check the paper 😊

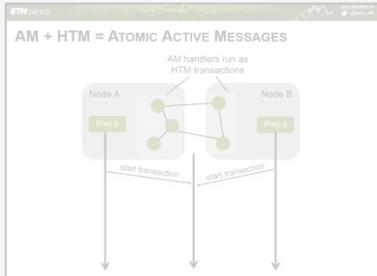
Same for other graphs

! Size for BG/Q ~100  
>  
Size for Haswell ~10

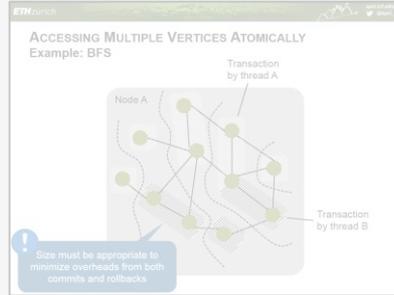
! Yes, we can

# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



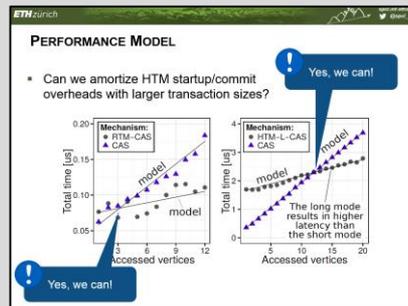
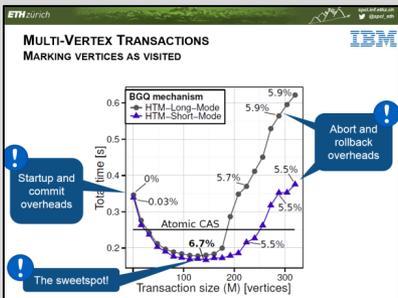
## Coarsening & coalescing



HTM + Active Messages  
= Atomic Active Messages

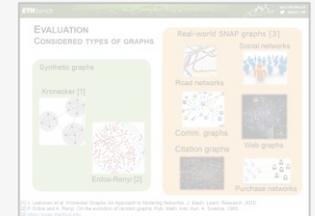
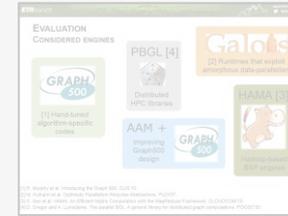
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

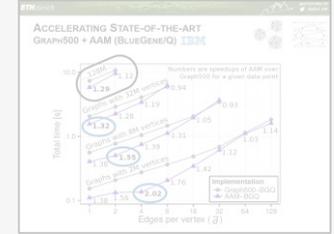
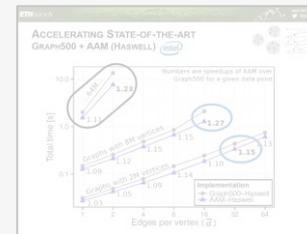


### Performance model

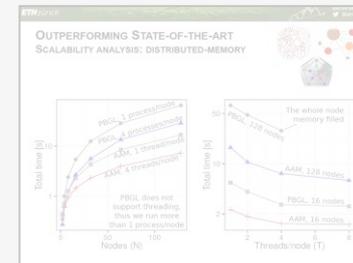
## Evaluation



## Considered engines and graphs



## Accelerating state-of-the-art

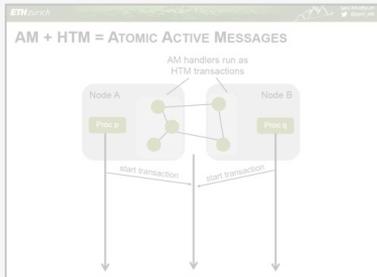


## Scalability

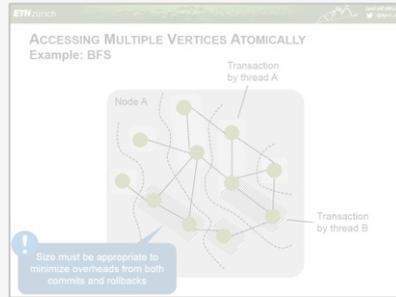


# OVERVIEW OF OUR RESEARCH

## HTM for graphs in SM & DM environments



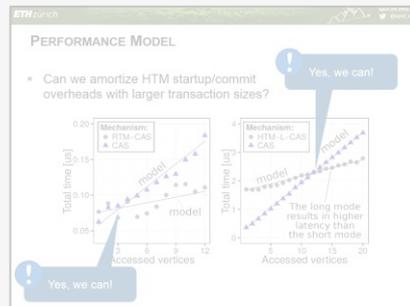
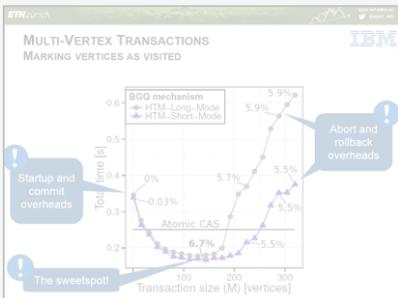
## Coarsening & coalescing



HTM + Active Messages  
= Atomic Active Messages

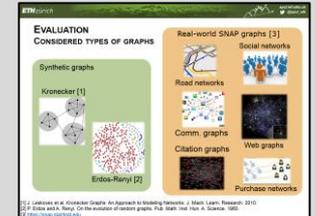
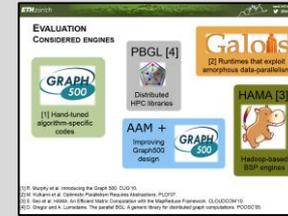
## Performance Modeling & Analysis

### Haswell & BG/Q Analysis

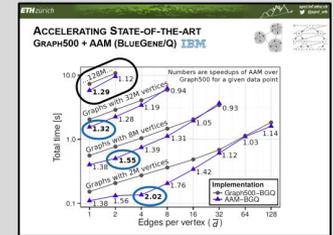
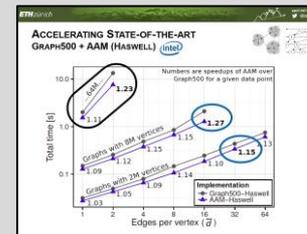


### Performance model

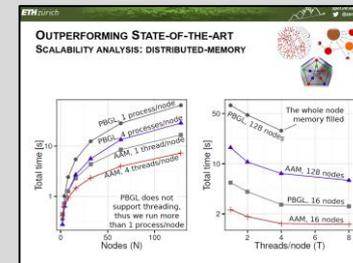
## Evaluation



## Considered engines and graphs



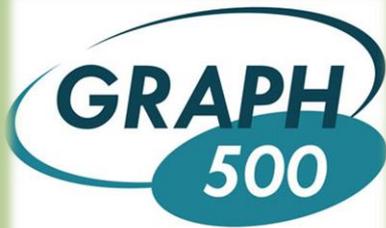
## Accelerating state-of-the-art



## Scalability

# EVALUATION CONSIDERED ENGINES

# EVALUATION CONSIDERED ENGINES



[1] Hand-tuned  
algorithm-specific  
codes

# EVALUATION

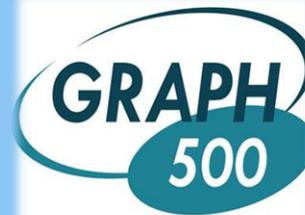
## CONSIDERED ENGINES



[1] Hand-tuned  
algorithm-specific  
codes

AAM +

Improving  
Graph500  
design



[1] R. Murphy et al. Introducing the Graph 500. CUG'10.

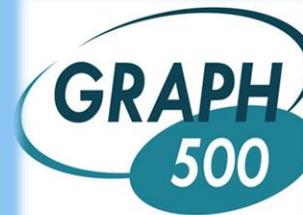
# EVALUATION CONSIDERED ENGINES



[1] Hand-tuned  
algorithm-specific  
codes

AAM +

Improving  
Graph500  
design



# Galois

[2] Runtimes that exploit  
amorphous data-parallelism

[1] R. Murphy et al. Introducing the Graph 500. CUG'10.

[2] M. Kulkarni et al. Optimistic Parallelism Requires Abstractions. PLDI'07.

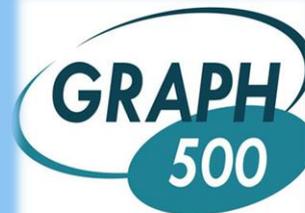
# EVALUATION CONSIDERED ENGINES



[1] Hand-tuned  
algorithm-specific  
codes

AAM +

Improving  
Graph500  
design



# Galois

[2] Runtimes that exploit  
amorphous data-parallelism

# HAMA [3]



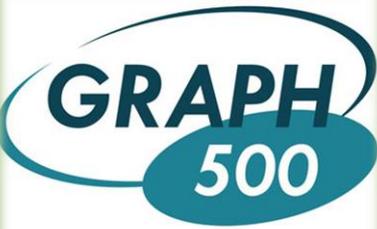
Hadoop-based  
BSP engines

[1] R. Murphy et al. Introducing the Graph 500. CUG'10.

[2] M. Kulkarni et al. Optimistic Parallelism Requires Abstractions. PLDI'07.

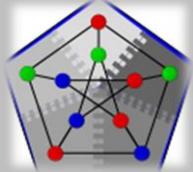
[3] S. Seo et al. HAMA: An Efficient Matrix Computation with the MapReduce Framework. CLOUDCOM'10.

# EVALUATION CONSIDERED ENGINES



[1] Hand-tuned  
algorithm-specific  
codes

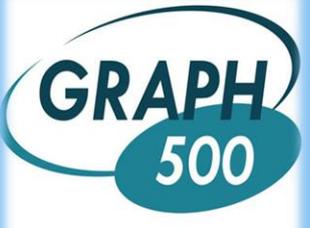
PBGL [4]



Distributed  
HPC libraries

AAM +

Improving  
Graph500  
design



Galois

[2] Runtimes that exploit  
amorphous data-parallelism

HAMA [3]



Hadoop-based  
BSP engines

[1] R. Murphy et al. Introducing the Graph 500. CUG'10.

[2] M. Kulkarni et al. Optimistic Parallelism Requires Abstractions. PLDI'07.

[3] S. Seo et al. HAMA: An Efficient Matrix Computation with the MapReduce Framework. CLOUDCOM'10.

[4] D. Gregor and A. Lumsdaine. The parallel BGL: A generic library for distributed graph computations. POOSC'05.

# EVALUATION

## CONSIDERED TYPES OF GRAPHS

# EVALUATION

## CONSIDERED TYPES OF GRAPHS

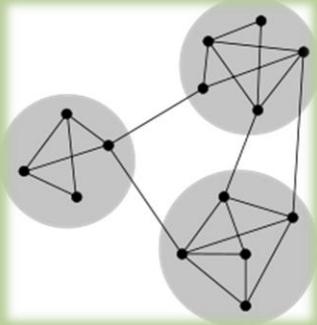
Synthetic graphs

# EVALUATION

## CONSIDERED TYPES OF GRAPHS

Synthetic graphs

Kronecker [1]

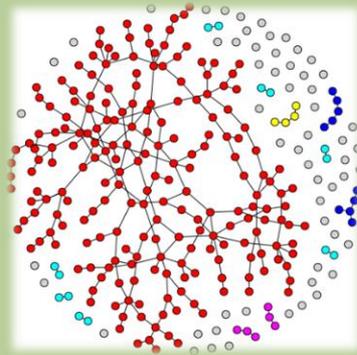
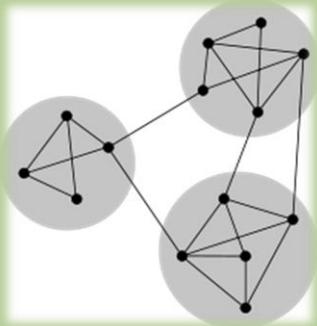


# EVALUATION

## CONSIDERED TYPES OF GRAPHS

### Synthetic graphs

#### Kronecker [1]



#### Erdős-Rényi [2]

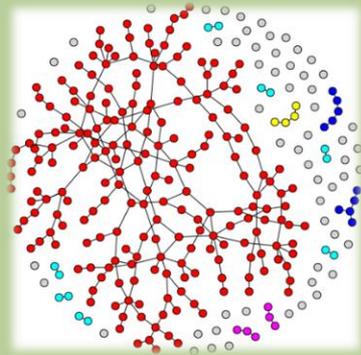
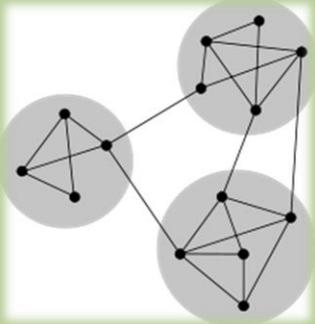
- [1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.  
[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.

# EVALUATION

## CONSIDERED TYPES OF GRAPHS

### Synthetic graphs

#### Kronecker [1]



#### Erdős-Rényi [2]

### Real-world SNAP graphs [3]

[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.

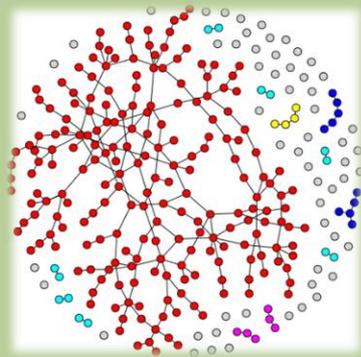
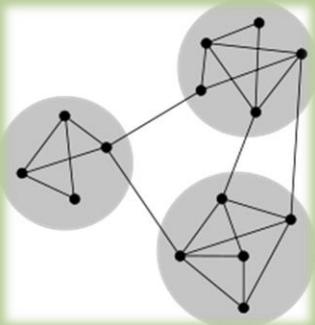
[3] <https://snap.stanford.edu>

# EVALUATION

## CONSIDERED TYPES OF GRAPHS

### Synthetic graphs

#### Kronecker [1]



Erdős-Rényi [2]

### Real-world SNAP graphs [3]

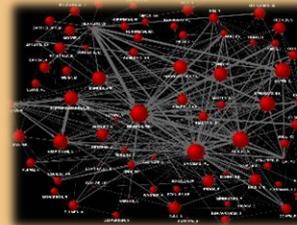


Road networks

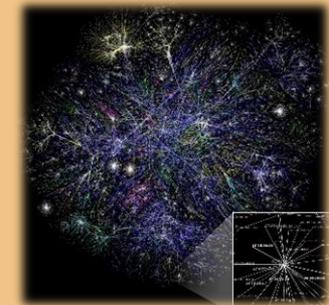


Comm. graphs

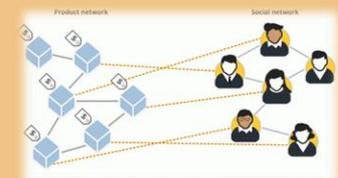
Citation graphs



### Social networks



Web graphs



Purchase networks

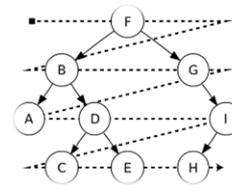
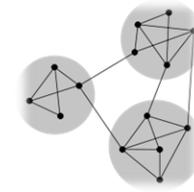
[1] J. Leskovec et al. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Research. 2010.

[2] P. Erdos and A. Renyi. On the evolution of random graphs. Pub. Math. Inst. Hun. A. Science. 1960.

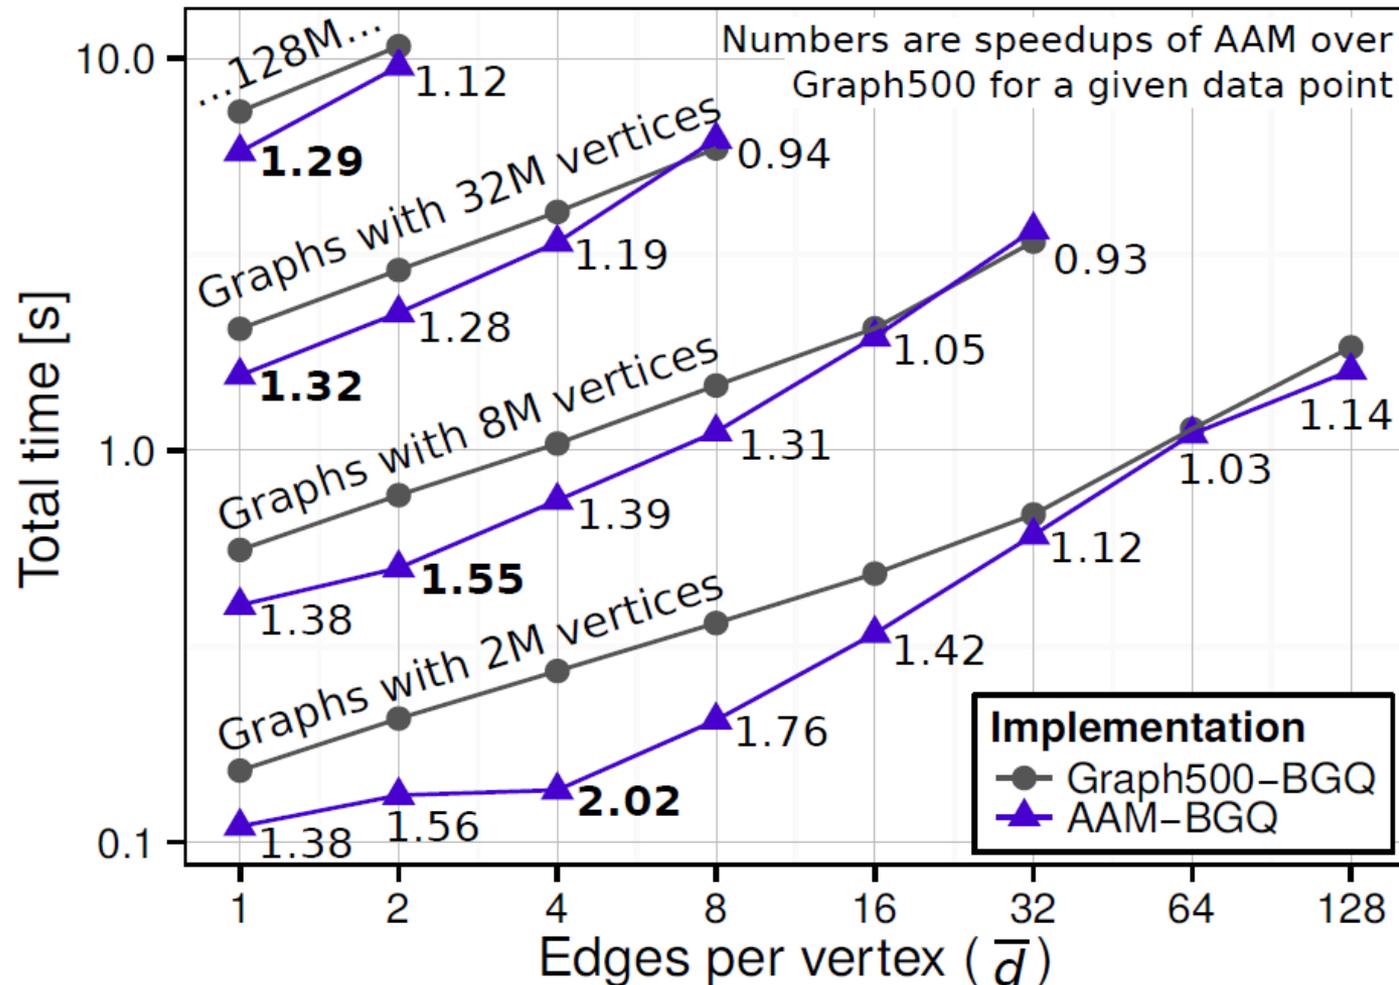
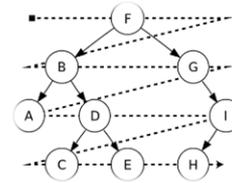
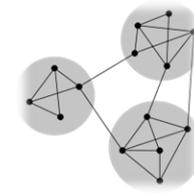
[3] <https://snap.stanford.edu>

# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)

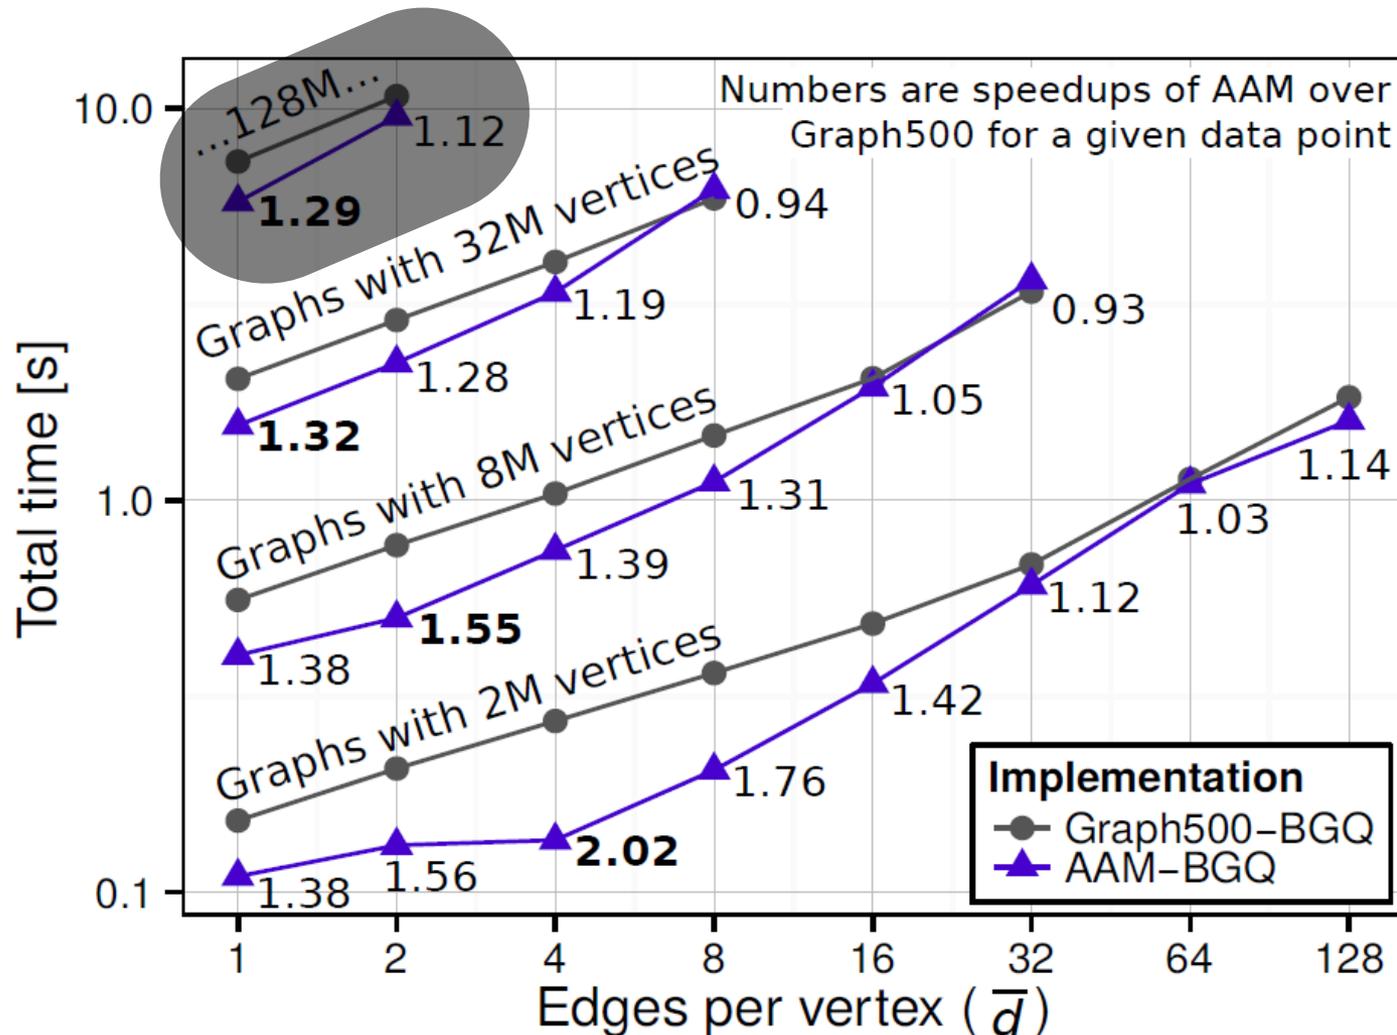
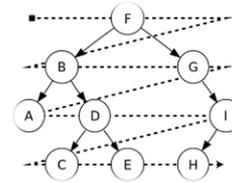
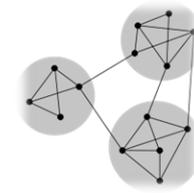
# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)



# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)

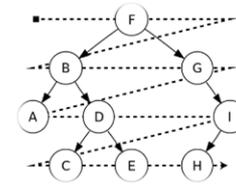
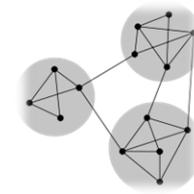


# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)

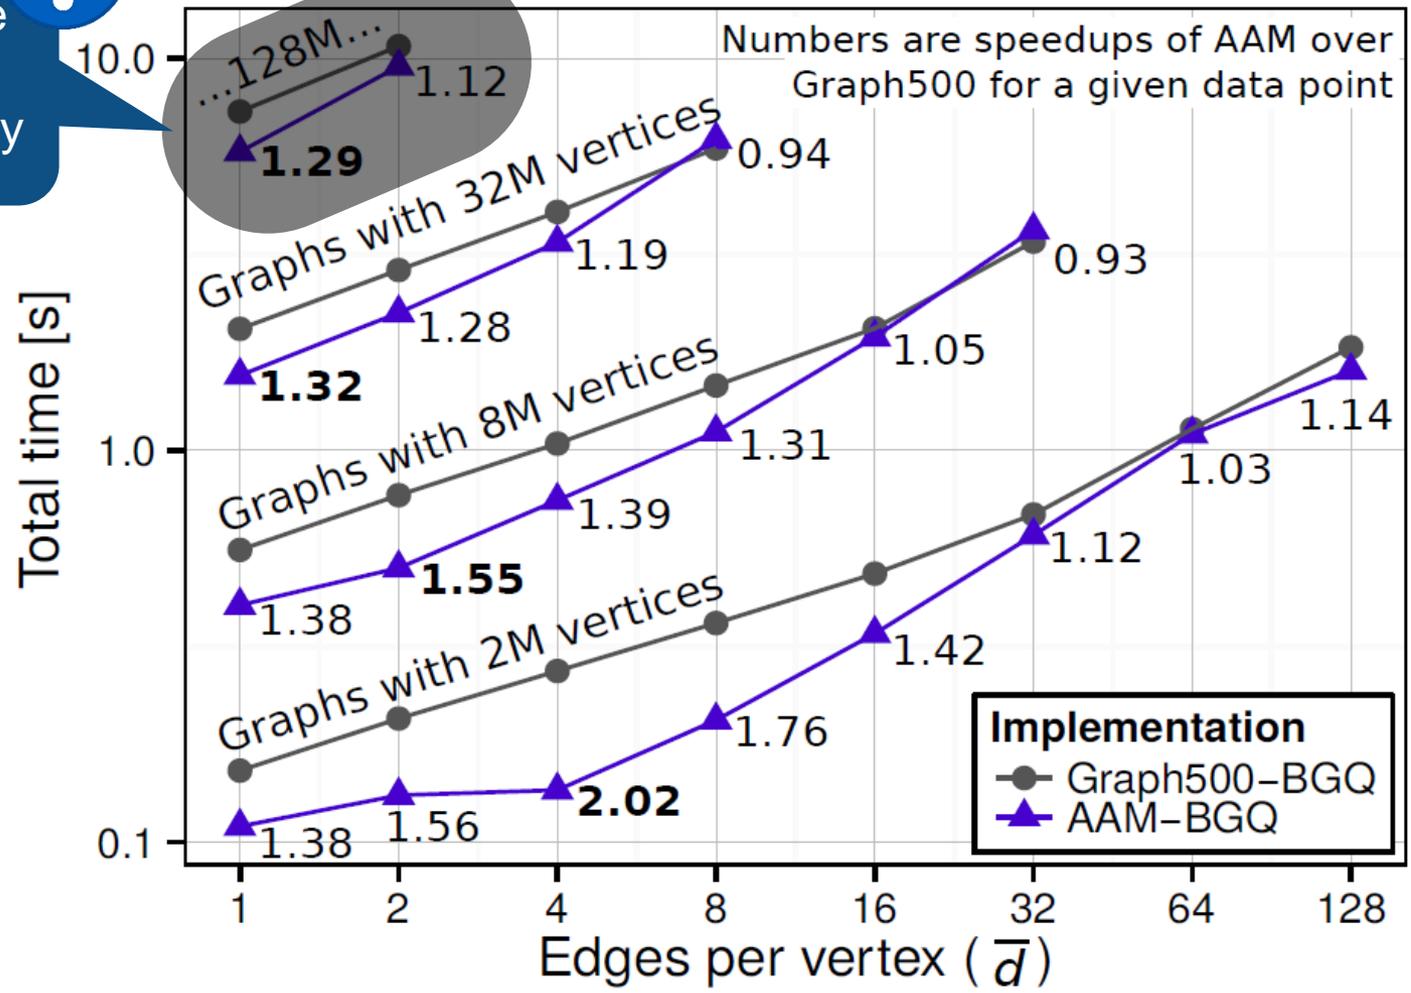




# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)



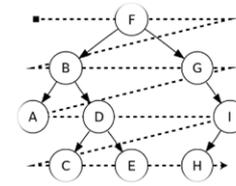
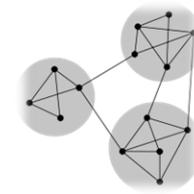
Fill the whole memory



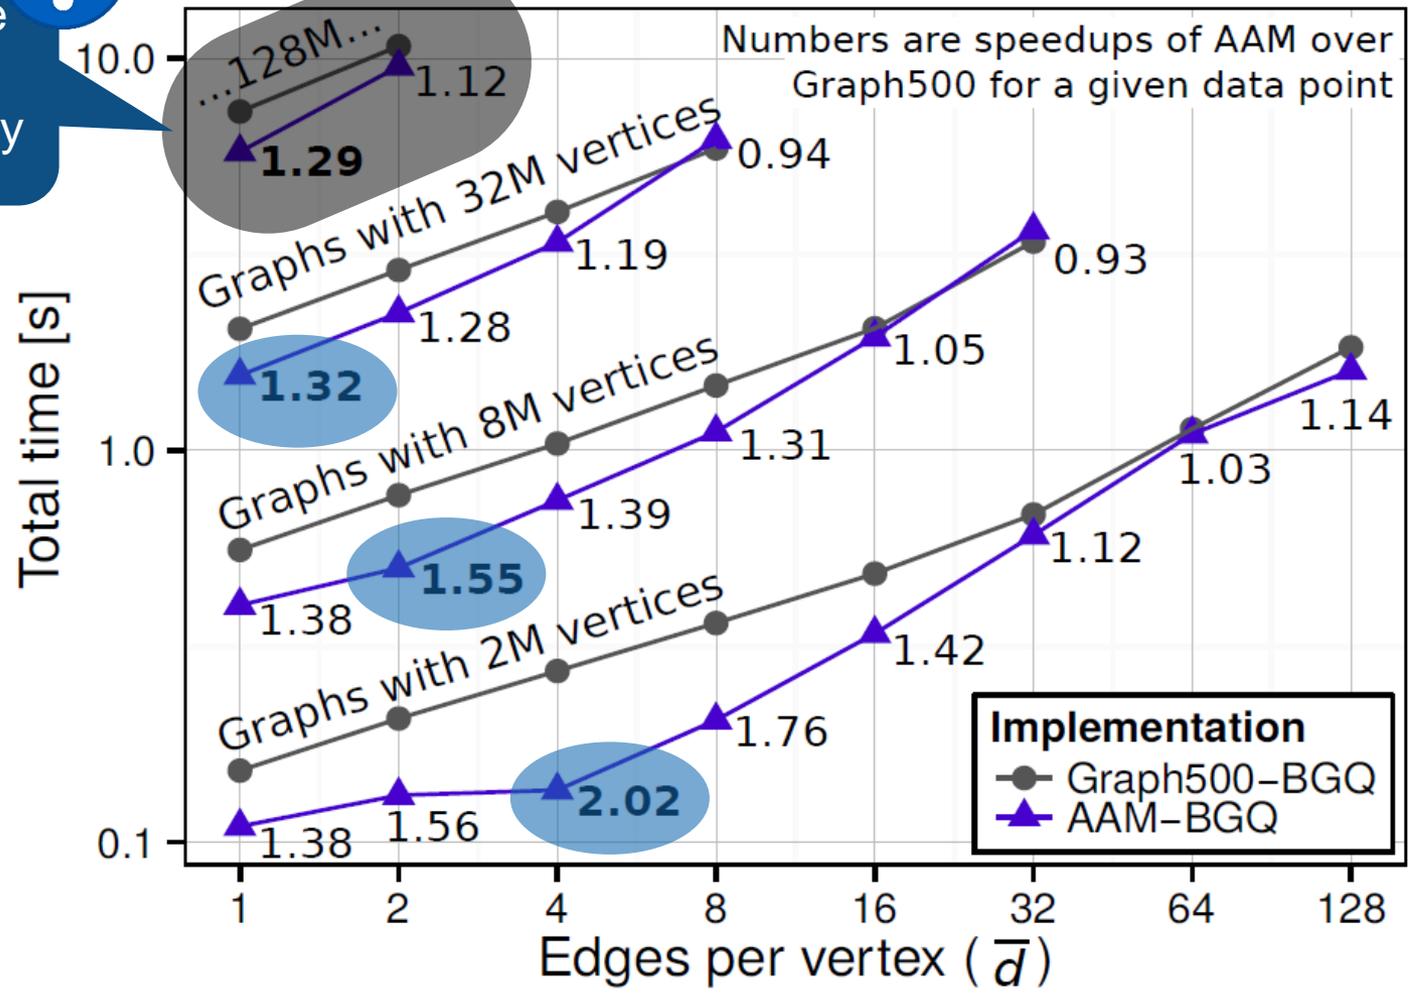
**Implementation**  
 ● Graph500-BGQ  
 ▲ AAM-BGQ



# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (BLUEGENE/Q)

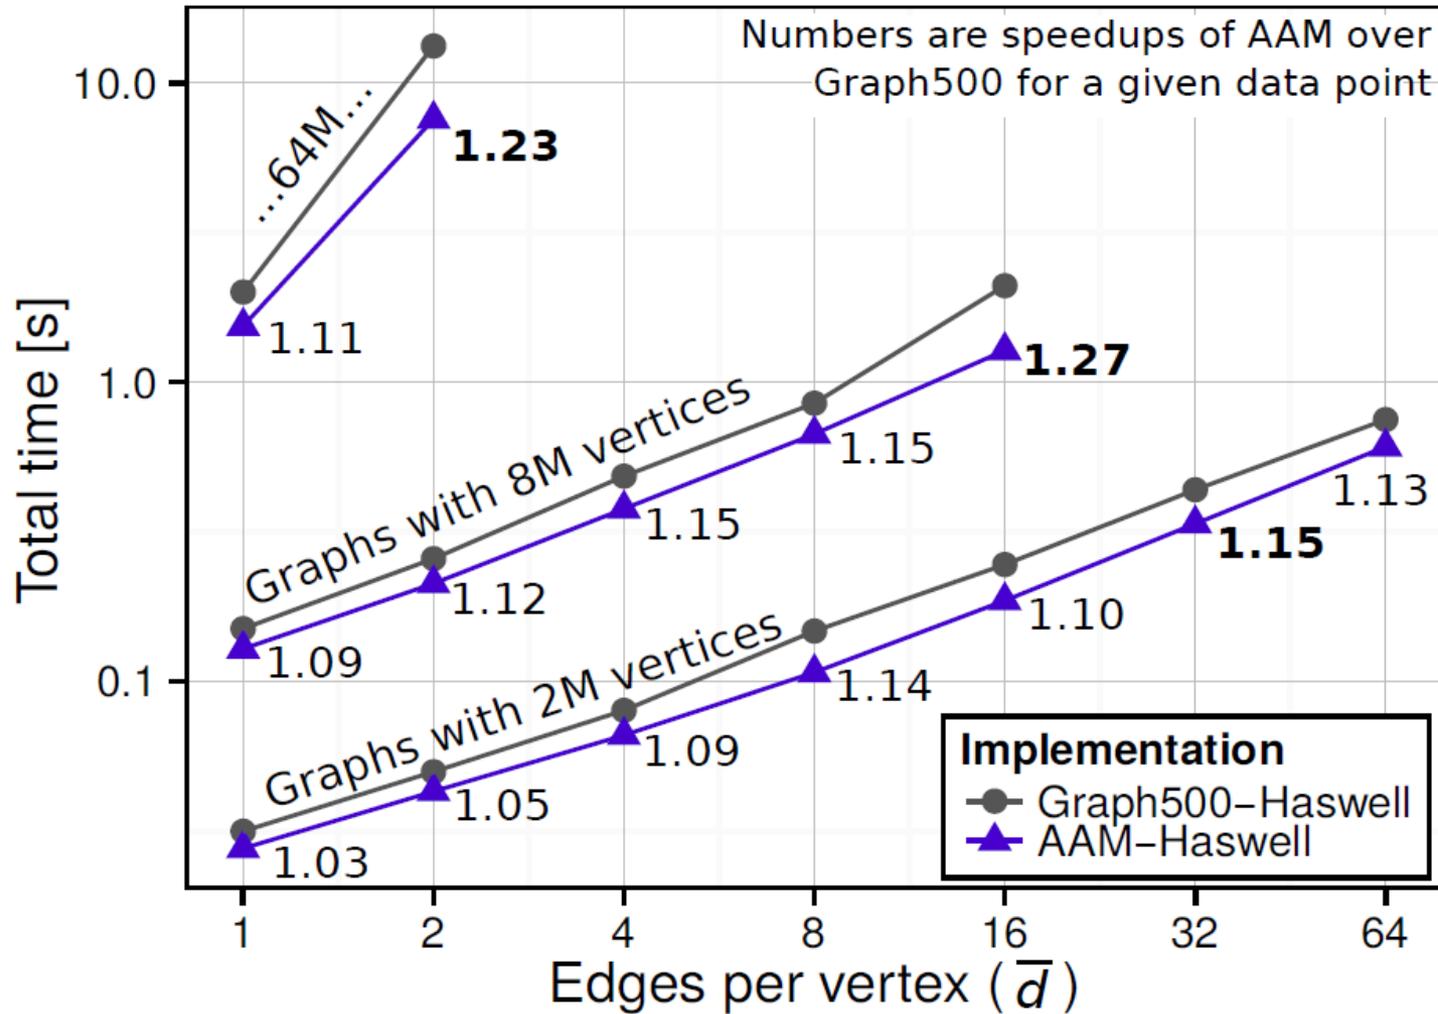
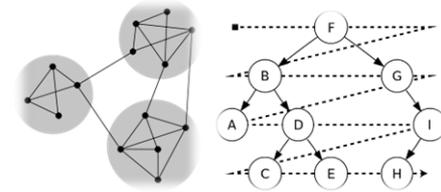


Fill the whole memory 

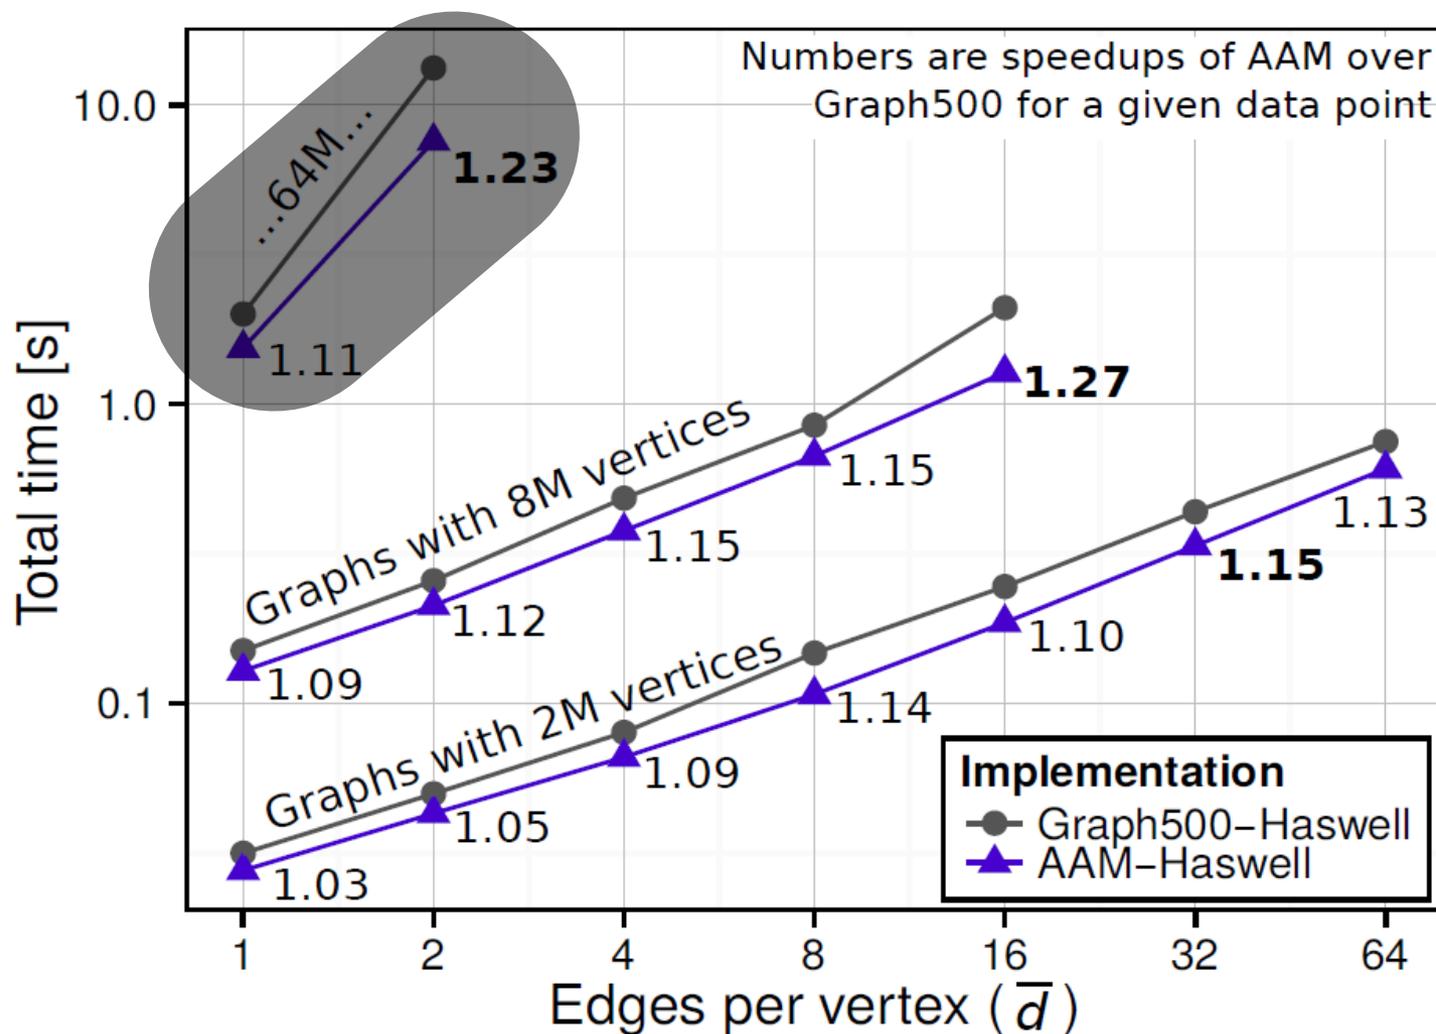
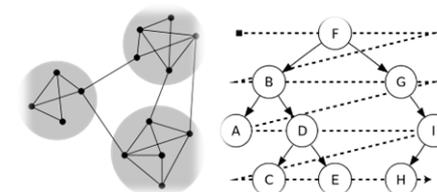


**Implementation**  
 ● Graph500-BGQ  
 ▲ AAM-BGQ

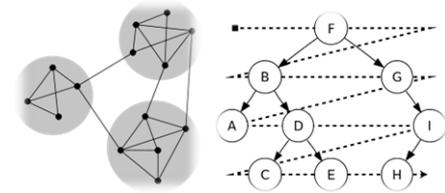
# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (HASWELL)

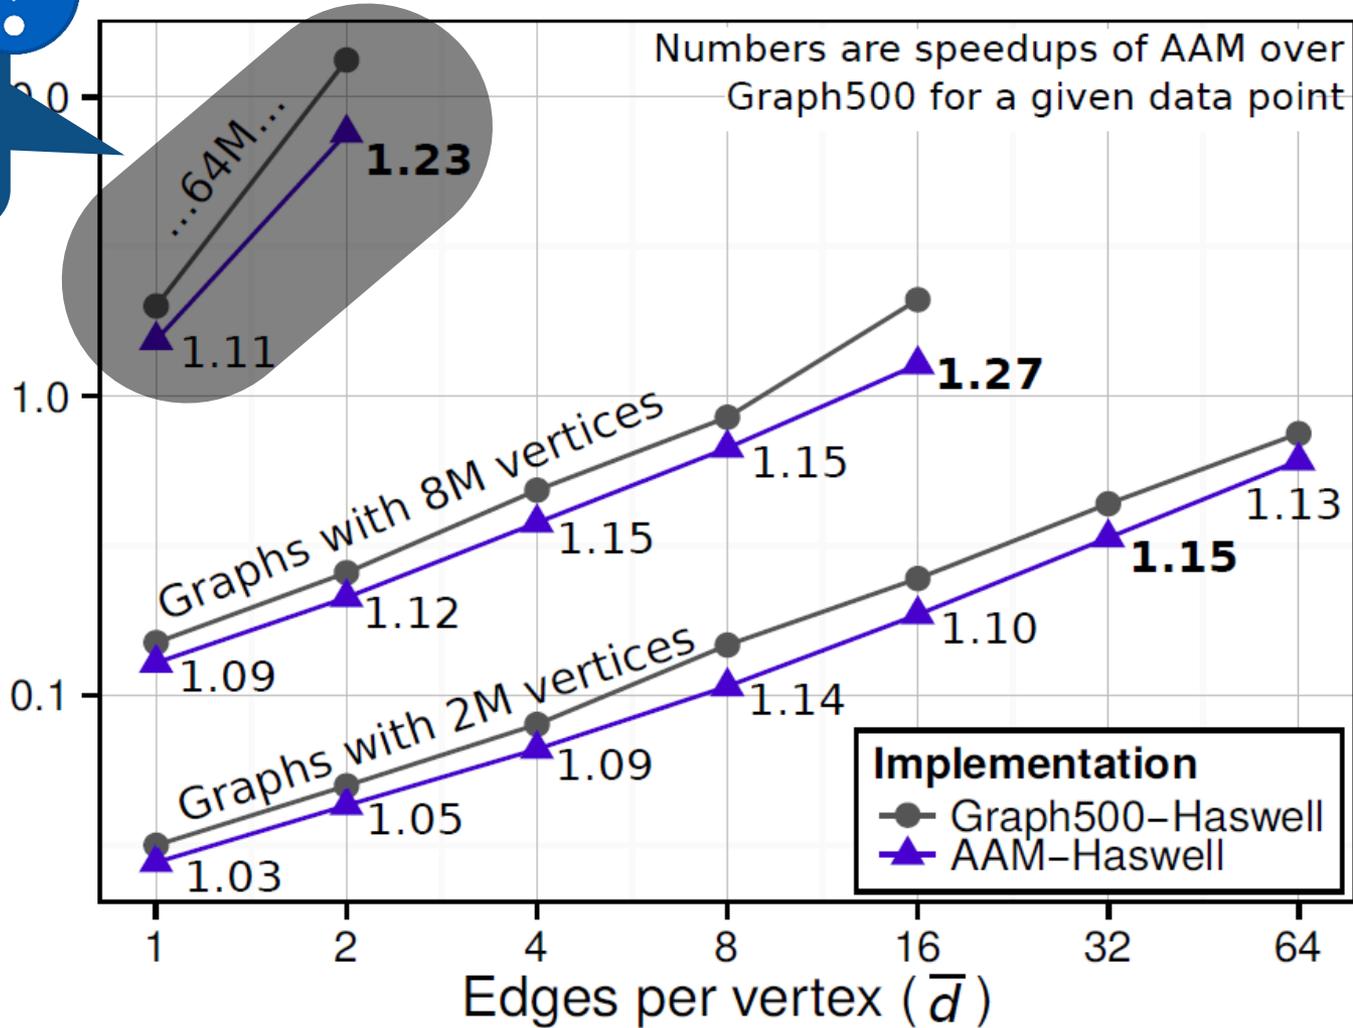
# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (HASWELL)



# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (HASWELL)



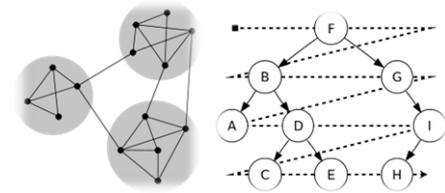
Fill the whole memory



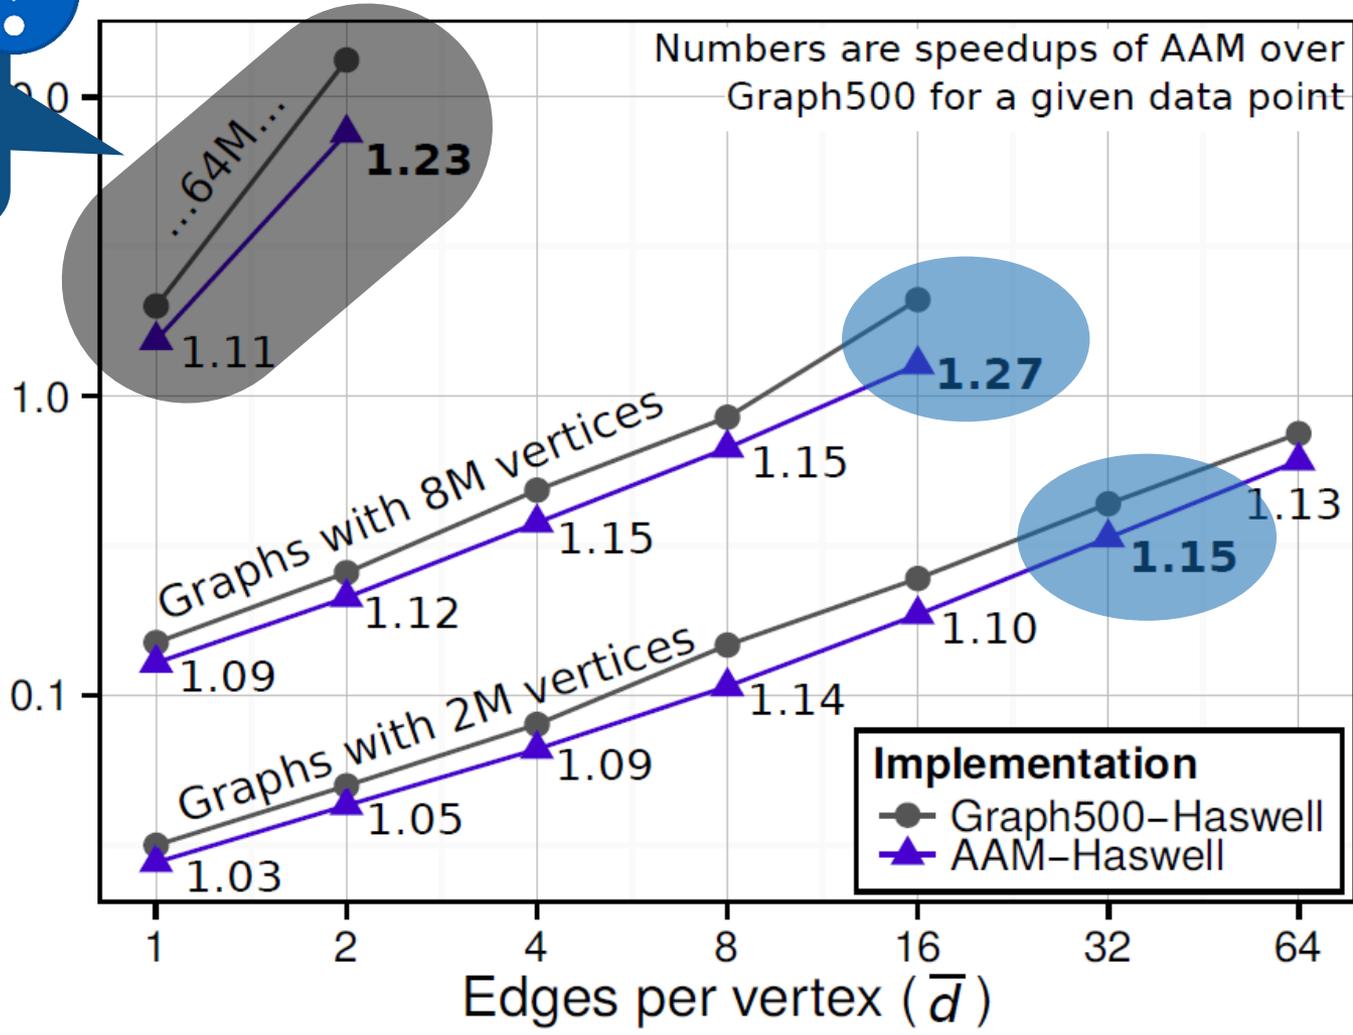
**Implementation**  
 ● Graph500-Haswell  
 ▲ AAM-Haswell



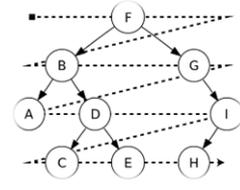
# ACCELERATING STATE-OF-THE-ART GRAPH500 + AAM (HASWELL)



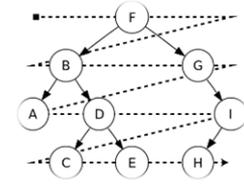
Fill the whole memory



# OUTPERFORMING STATE-OF-THE-ART

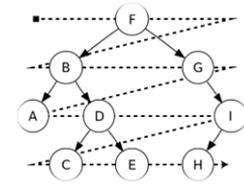


# OUTPERFORMING STATE-OF-THE-ART



Input graph properties					BG/Q analysis			Haswell analysis					
Type	ID	Name	$ V $	$ E $	$S$ over g500 ( $M = 24$ )	$M$	$S$ over g500	$S$ over g500 ( $M = 2$ )	$S$ over Galois ( $M = 2$ )	$M$	$S$ over g500	$S$ over Galois	$S$ over HAMA
Comm. networks (CNs)	cWT	wiki-Talk	2.4M	5M	2.82	48	3.35	0.91	1.22	6	0.96	1.28	344
	cEU	email-EuAll	265k	420k	3.67	32	4.36	0.76	0.88	4	0.97	1.12	1448
Social networks (SNs)	sLV	soc-LiveJ.	4.8M	69M	1.44	12	1.56	1.05	1.1	3	1.07	1.12	$> 10^4$
	sOR	com-orkut	3M	117M	1.22	20	1.27	1.06	0.69	4	1.13	0.74	$> 10^4$
	sLJ	com-lj	4M	34M	1.44	12	1.54	1.03	1.03	4	1.04	1.04	603
	sYT	com-youtube	1.1M	2.9M	1.67	8	1.84	0.96	1.1	5	0.98	1.11	670
	sDB	com-dblp	317k	1M	1.33	8	1.80	$\approx 1$	2.5	2	$\approx 1$	2.53	2160
	sAM	com-amazon	334k	925k	1.14	8	1.62	1.04	1.64	2	1.04	1.64	1426
Purchase network (PNs)	pAM	amazon0601	403k	3.3M	1.45	8	1.91	$\approx 1$	1.25	3	1.03	1.30	618
Road networks (RNs)	rCA	roadNet-CA	1.9M	5.5M	$\approx 1$	2	1.59	1.33	1.74	8	1.38	1.80	$> 10^4$
	rTX	roadNet-TX	1.3M	3.8M	$\approx 1$	2	1.53	1.29	1.89	6	1.42	2.08	$> 10^4$
	rPA	roadNet-PA	1M	3M	$\approx 1$	2	1.52	$\approx 1$	2.00	9	1.07	2.16	$> 10^4$
Citation graphs (CGs)	ciP	cit-Patents	3.7M	16.5M	1.16	8	1.57	1.01	1.26	2	1.01	1.26	1875
Web graphs (WGs)	wGL	web-Google	875k	5.1M	1.78	12	2.08	0.98	1.26	6	1.06	1.35	365
	wBS	web-BerkStan	685k	7.6M	1.91	24	1.91	0.93	1.31	5	1.07	1.40	755
	wSF	web-Stanford	281k	2.3M	1.89	24	1.89	0.98	1.54	5	1.07	1.58	1077

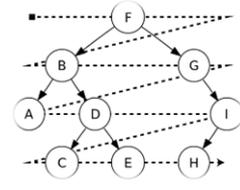
# OUTPERFORMING STATE-OF-THE-ART



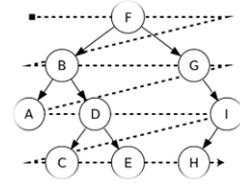
Input graph properties					BG/Q analysis				Haswell analysis				
Type	ID	Name	$V$	$E$	5 runs gG00 ( $M = 24$ )	$M$	5 runs gG00	5 runs gG00 ( $M = 2$ )	5 runs Galois ( $M = 2$ )	$M$	5 runs gG00	5 runs Galois	5 runs FLAMA
E-mail networks (09%)	wWT	wiki-Talk	2.4M	5M	2.52	45	1.35	0.91	1.22	8	1.06	1.25	364
	zEU	zoo-EDUAll	265k	420k	3.07	32	1.36	0.97	0.95	4	0.97	1.19	1445
	zLV	zoo-Lives	1.6M	95M	1.16	13	1.55	1.05	1.16	5	1.07	1.13	107
Road networks (09%)	rCA	roadNet-CA	1.9M	5.3M	>1	2	1.59	1.03	1.74	8	1.38	1.80	>10 <sup>5</sup>
	rTX	roadNet-TX	1.3M	3.8M	>1	2	1.53	1.20	1.80	6	1.42	2.05	>10 <sup>5</sup>
	rPA	roadNet-PA	1M	3M	>1	2	1.52	>1	2.00	9	1.07	2.10	>10 <sup>5</sup>
Citation graphs (05%)	citP	cit-Patents	3.7M	10.5M	1.10	8	1.67	1.01	1.26	2	1.01	1.26	1875
Web graphs (05%)	wGL	web-Google	875k	5.1M	1.75	12	2.08	0.98	1.26	6	1.06	1.35	365
	wBS	web-BerkStan	645k	7.6M	1.91	24	1.91	0.93	1.31	5	1.07	1.40	759
	wSP	web-Stanford	281k	2.3M	1.80	24	1.80	0.98	1.54	5	1.07	1.58	1077

😊 No, you don't have to read it. All details are in the paper. Here: just a summary.

# OUTPERFORMING STATE-OF-THE-ART HASWELL

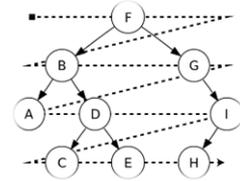


# OUTPERFORMING STATE-OF-THE-ART HASWELL



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000

# OUTPERFORMING STATE-OF-THE-ART HASWELL



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000

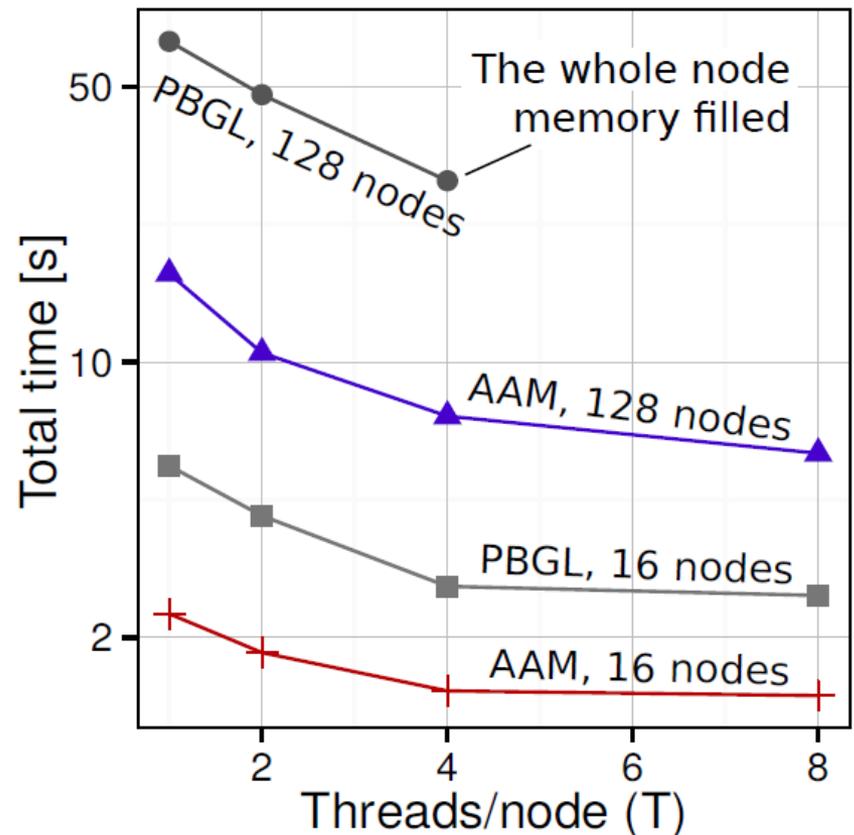
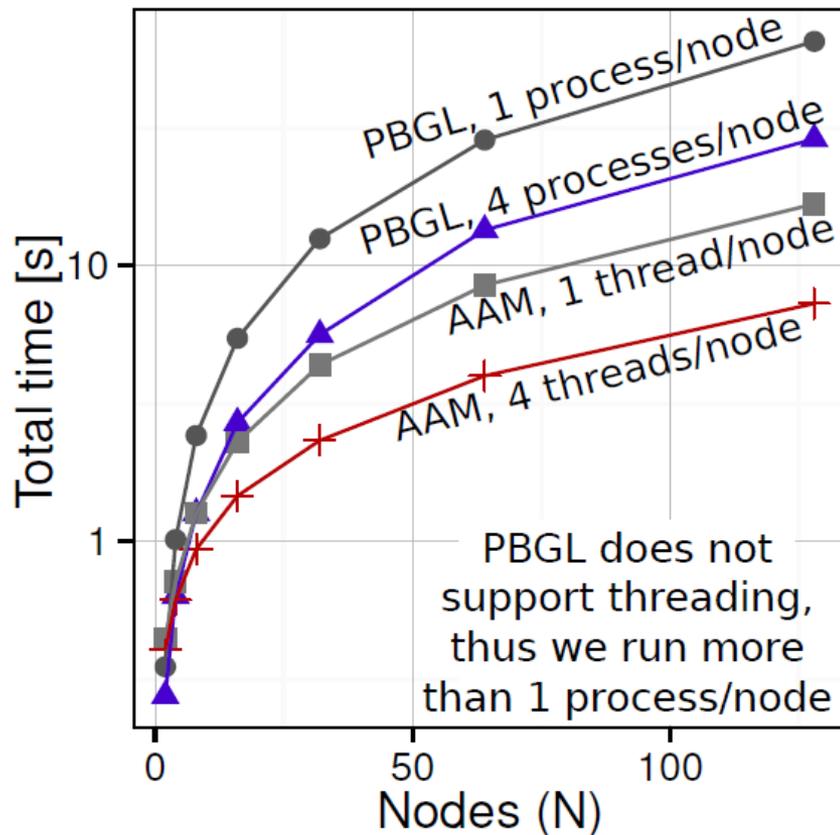
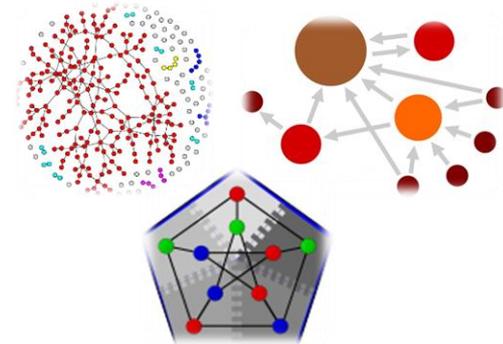


1.85x on average, up to 4.3x



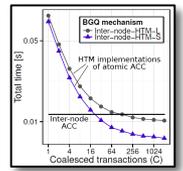
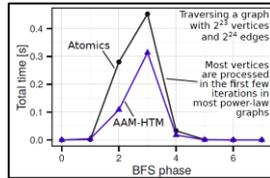
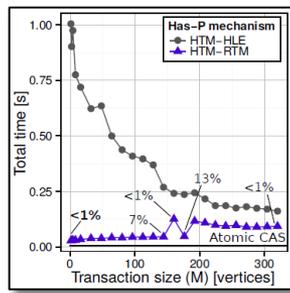
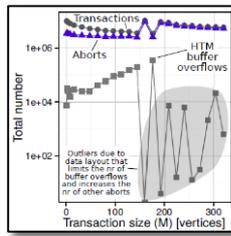
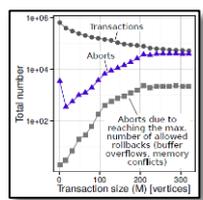
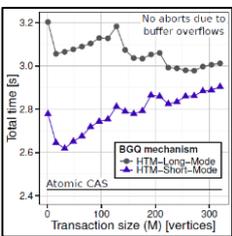
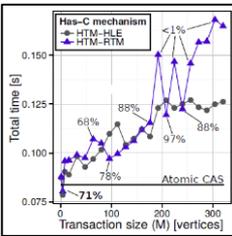
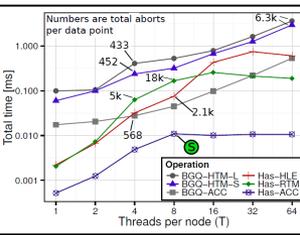
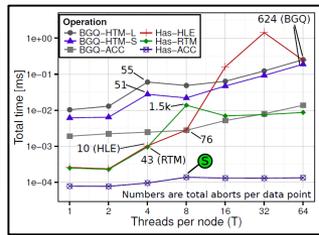
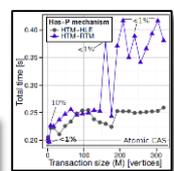
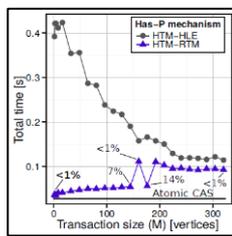
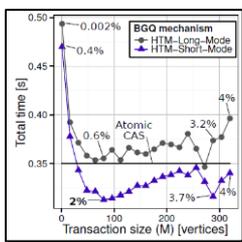
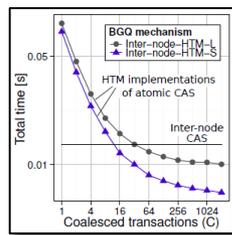
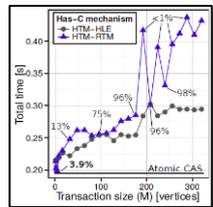
# OUTPERFORMING STATE-OF-THE-ART

## SCALABILITY ANALYSIS: DISTRIBUTED-MEMORY



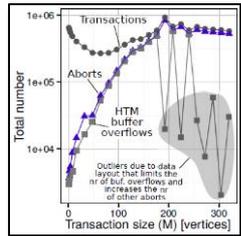
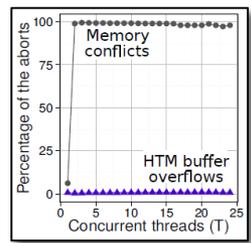
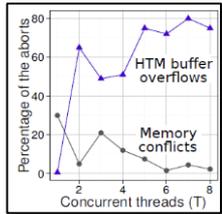
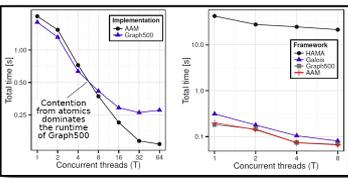
# OTHER ANALYSES

# OTHER ANALYSES



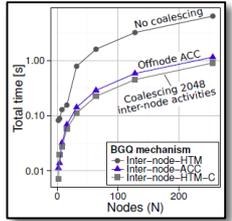
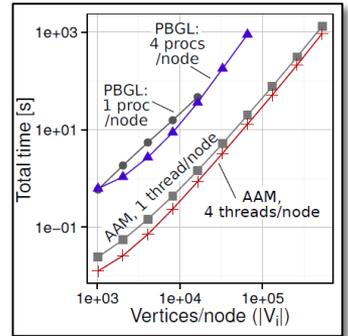
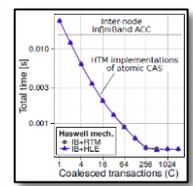
Aborts due to:

	Memory conflicts	Buffer overflows	Other reasons
10 ops	Has-RTM 1,520	1	0
	BGQ-HTM-L 624	62	614
	BGQ-HTM-S 623	62	613
100 ops	Has-RTM 18,952	33	0
	BGQ-HTM-L 6,374	637	6,360
	BGQ-HTM-S 6,392	639	6,380

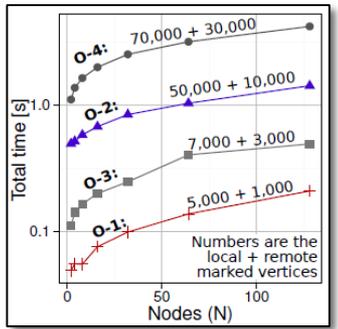


Aborts due to:

	Memory conflicts	Buffer overflows	Other reasons
10 ops	Has-RTM 2	2	0
	BGQ-HTM-L 802	3	1
	BGQ-HTM-S 1,118	46	180
100 ops	Has-RTM 2	2	0
	BGQ-HTM-L 1,539	5	1
	BGQ-HTM-S 2,242	13	2



Input graph properties		BG/Q analysis				Haswell analysis						
Type	ID Name	V	E	S over g500 (M = 24)	M	S over g900 (M = 2)	S over g500 (M = 2)	S over Galois (M = 2)	M	S over g500	S over Galois	S over HAMA
Comm. networks (Cn)	cWT wiki-Talk	2.4M	5M	2.82	32	3.35	0.91	1.22	4	0.96	1.28	344
	cEU email-EuAll	265k	420k	3.67	32	4.36	0.76	0.88	4	0.97	1.12	1448
Social networks (Sn)	sLV soc-LiveJ.	4.8M	69M	1.44	12	1.56	1.05	1.1	3	1.07	1.12	> 10 <sup>4</sup>
	sOR com-orkut	3M	117M	1.22	20	1.27	1.06	0.69	4	1.13	0.74	> 10 <sup>4</sup>
	sLJ com-lj	4M	34M	1.44	12	1.54	1.03	1.03	4	1.04	1.04	603
	sYT com-youtube	1.1M	2.9M	1.67	8	1.84	0.96	1.1	5	0.98	1.11	670
Purchase network (Pn)	sDB com-dblp	317k	1M	1.33	8	1.80	≈1	2.5	2	≈1	2.53	2160
	sAM com-amazon	334k	925k	1.14	8	1.62	1.04	1.64	2	1.04	1.64	1426
Road networks (Rn)	rCA roadNet-CA	1.9M	5.5M	≈1	2	1.59	1.33	1.74	8	1.38	1.80	> 10 <sup>4</sup>
	rTX roadNet-TX	1.3M	3.8M	≈1	2	1.53	1.29	1.89	6	1.42	2.08	> 10 <sup>4</sup>
	rPA roadNet-PA	1M	3M	≈1	2	1.52	≈1	2.00	9	1.07	2.16	> 10 <sup>4</sup>
Citation graphs (Cn)	cIP cit-Patents	3.7M	16.5M	1.16	8	1.57	1.01	1.26	2	1.01	1.26	1875
Web graphs (Wn)	wGL web-Google	875k	5.1M	1.78	12	2.08	0.98	1.26	6	1.06	1.35	365
	wBS web-BerkStan	685k	7.6M	1.91	24	1.91	0.93	1.31	5	1.07	1.40	735
	wSF web-Stanford	281k	2.3M	1.80	24	1.89	0.98	1.54	5	1.07	1.58	1077



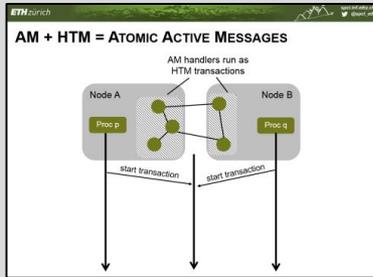
# CONCLUSIONS

# CONCLUSIONS

HTM for graphs in SM & DM environments

# CONCLUSIONS

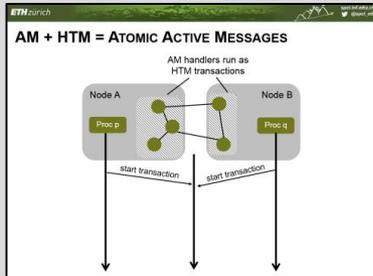
## HTM for graphs in SM & DM environments



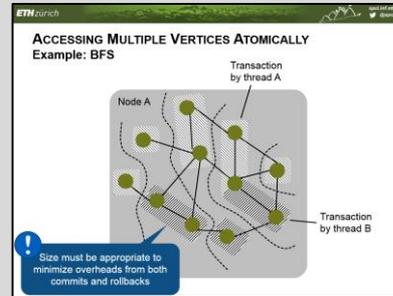
AAM: Combine the advantages of HTM and Active Messages

# CONCLUSIONS

## HTM for graphs in SM & DM environments



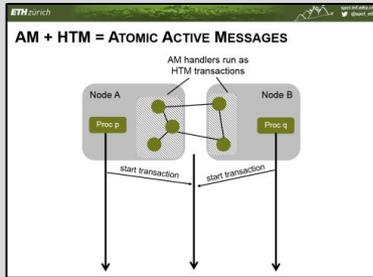
AAM: Combine the advantages of HTM and Active Messages



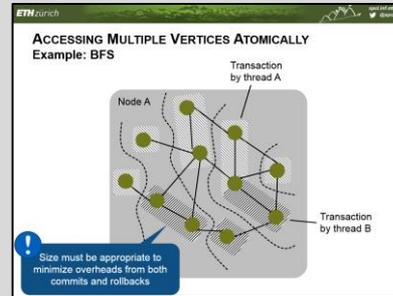
Illustrate HTM's advantages in performance, next to programmability

# CONCLUSIONS

## HTM for graphs in SM & DM environments



AAM: Combine the advantages of HTM and Active Messages



Illustrate HTM's advantages in performance, next to programmability

**Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages**

Mucelj Besta  
 Department of Computer Science  
 ETH Zurich  
 Universitätsstr. 8, 8092 Zurich, Switzerland  
 {mucelj,besta}@inf.ethz.ch
 

 Thorsten Hoeller  
 Department of Computer Science  
 ETH Zurich  
 Universitätsstr. 8, 8092 Zurich, Switzerland  
 hoeller@inf.ethz.ch

**ABSTRACT**

We propose Atomic Active Message (AAM), a mechanism that accelerates irregular graph computations on both shared and distributed-memory machines. The key idea behind AAM is the hardware transactional memory (HTM) used to avoid the complex and slow processing of regular updates due to concurrent and conflicting updates. Instead, hardware transactions to read-only vertices of graphs guarantee the regularity of graph processing. We present a detailed performance analysis of AAM on both shared and distributed-memory machines and we illustrate our implementation results. Across different HTM architectures we report the benefits of graph processing. AAM can be used to implement algorithms affected by regular graph processing overheads and to improve the performance of irregular graph processing codes such as Graph500 or Ginkgo.

**Categories and Subject Descriptors:** Parallel Programming; Distributed Programming

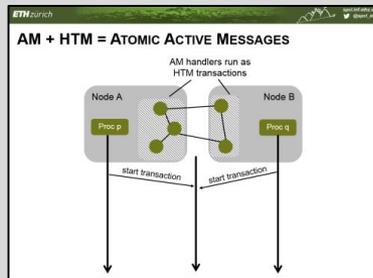
**General Terms:** Performance, Design

**1. INTRODUCTION**

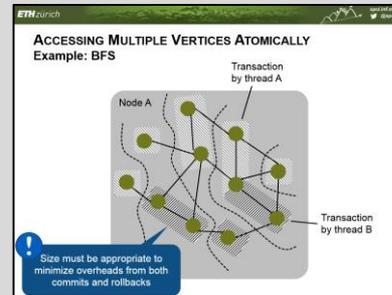
Deliver the of hierarchy of atomic messages that covers various graph algorithms

# CONCLUSIONS

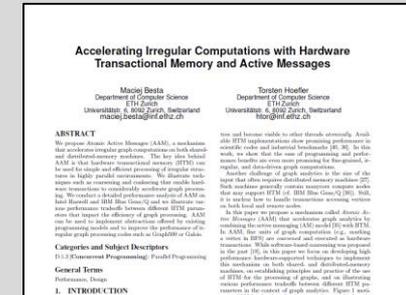
## HTM for graphs in SM & DM environments



AAM: Combine the advantages of HTM and Active Messages



Illustrate HTM's advantages in performance, next to programmability

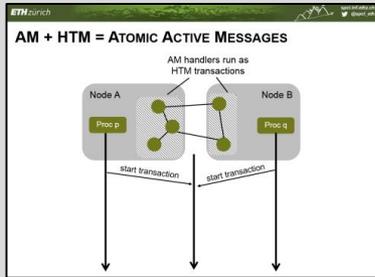


Deliver the of hierarchy of atomic messages that covers various graph algorithms

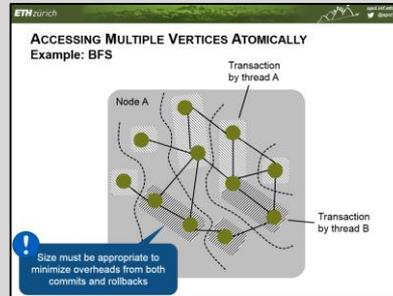
## Detailed performance analysis

# CONCLUSIONS

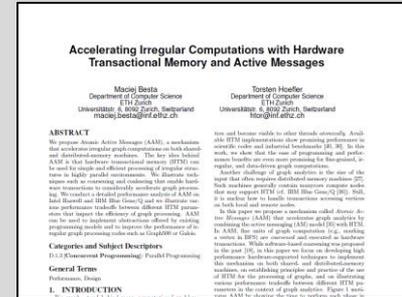
## HTM for graphs in SM & DM environments



AAM: Combine the advantages of HTM and Active Messages

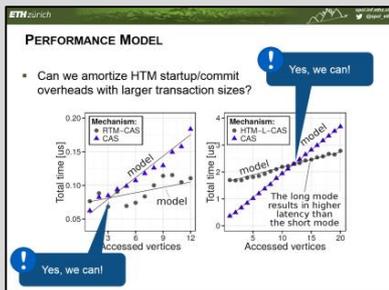


Illustrate HTM's advantages in performance, next to programmability



Deliver the of hierarchy of atomic messages that covers various graph algorithms

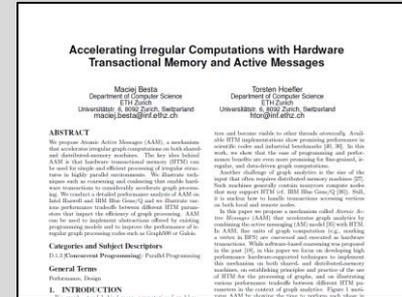
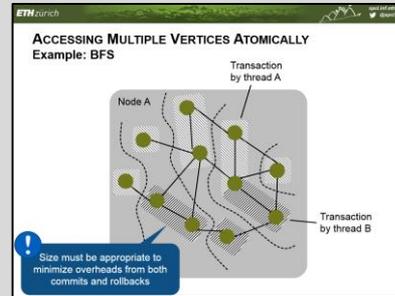
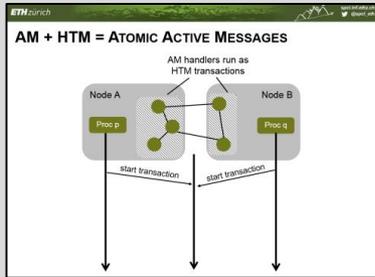
## Detailed performance analysis



Model & analyze performance tradeoffs

# CONCLUSIONS

## HTM for graphs in SM & DM environments

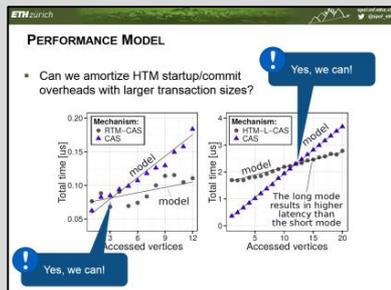


AAM: Combine the advantages of HTM and Active Messages

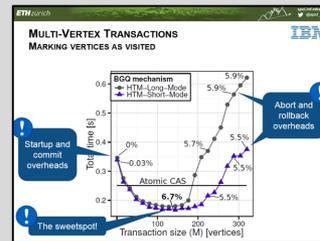
Illustrate HTM's advantages in performance, next to programmability

Deliver the of hierarchy of atomic messages that covers various graph algorithms

## Detailed performance analysis



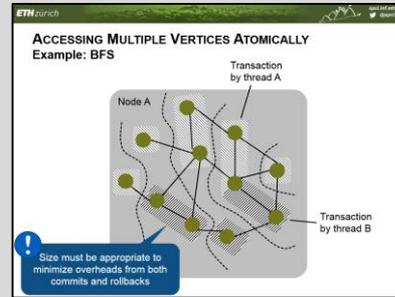
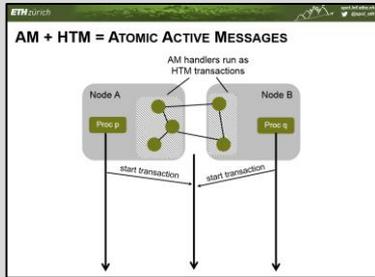
Derive close-to-optimal transaction sizes for Haswell & BG/Q



Model & analyze performance tradeoffs

# CONCLUSIONS

## HTM for graphs in SM & DM environments



**Accelerating Irregular Computations with Hardware Transactional Memory and Active Messages**

Muqoj Bebis, Department of Computer Science, ETH Zurich, Switzerland, muqoj.bebis@inf.ethz.ch  
 Torsten Hofer, Department of Computer Science, ETH Zurich, Switzerland, hofer@inf.ethz.ch

**ABSTRACT**  
 We propose Atomic Active Messages (AAM), a mechanism that enables regular graph computations on both shared and distributed memories. The key idea behind AAM is to leverage transactional memory (TM) mechanisms to enable different processors to operate on the same highly parallel computation. By allowing vertices to be accessed and updated by multiple hardware transactions to coarsely-grained graph processing, we realize a shared performance model of AAM on both shared and IBM Blue Gene/Q and we illustrate our new performance model. Active Message (AM) message sizes need to be implemented adaptively based on existing graph processing methods and to improve the performance of irregular graph processing such as Graph500 or Ginkgo.

**Categories and Subject Descriptors**  
 D.1.1 Operating Systems: Parallel Programming  
 D.1.2 Languages: Programming Languages

**General Terms**  
 Performance, Design

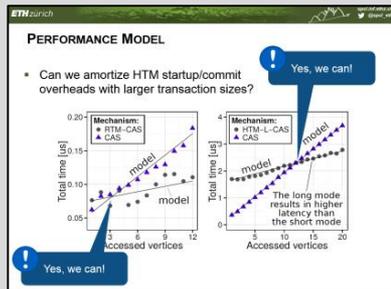
**1. INTRODUCTION**

AAM: Combine the advantages of HTM and Active Messages

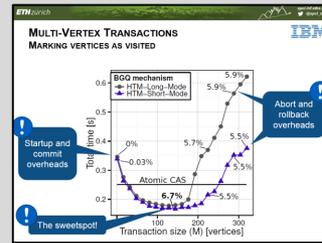
Illustrate HTM's advantages in performance, next to programmability

Deliver the of hierarchy of atomic messages that covers various graph algorithms

## Detailed performance analysis

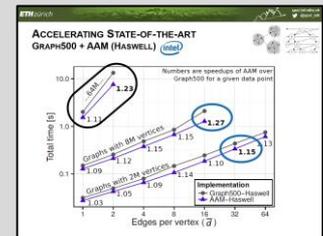
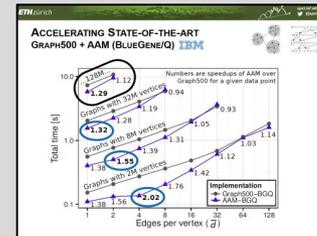


Derive close-to-optimal transaction sizes for Haswell & BG/Q

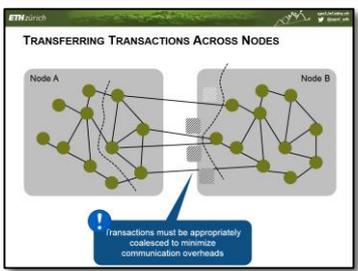
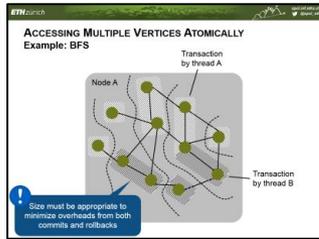
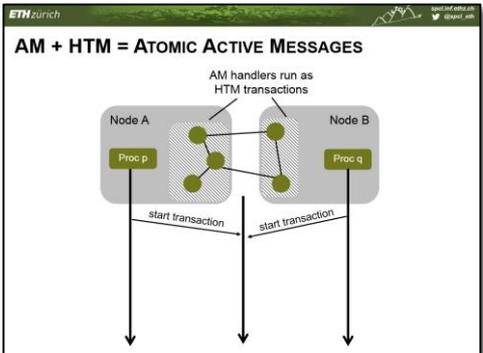


Model & analyze performance tradeoffs

## Accelerating state-of-the-art



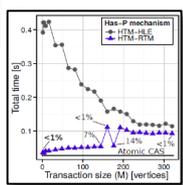
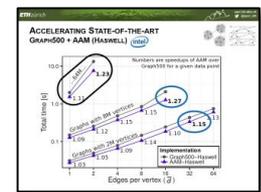
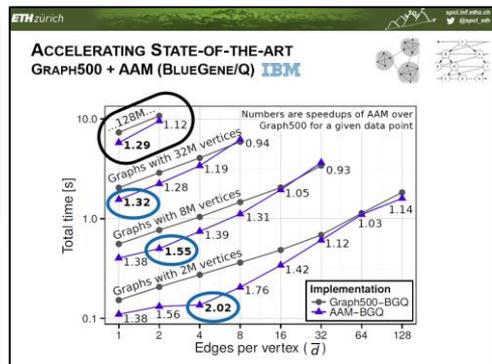
Average speedup 1.85x  
Up to 4x



### EVALUATION CONSIDERED ENGINES

- GRAPH 500 [1] Hand-tuned algorithm-specific codes
- PBGL [4] Distributed GPU libraries
- Galios [2] Runtimes that exploit architectural data-parallelism
- AAM + GRAPH 500 Improving Graph500 design
- HAMA [3] Hardware-based HTM engines

101. Hwang et al. Computing for Supercomputers, 2012  
 102. Hwang et al. Scalable Parallel Graph Algorithms, 2013  
 103. Wang et al. Scalable Graph Algorithms for High-End Parallel Systems, 2012  
 104. Wang et al. Scalable Graph Algorithms for High-End Parallel Systems, 2012  
 105. Wang et al. Scalable Graph Algorithms for High-End Parallel Systems, 2012

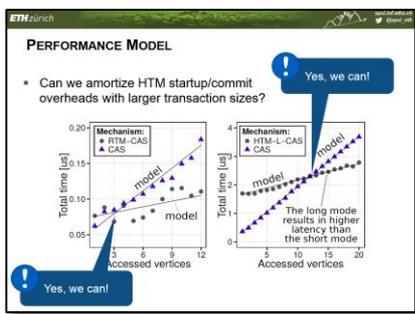
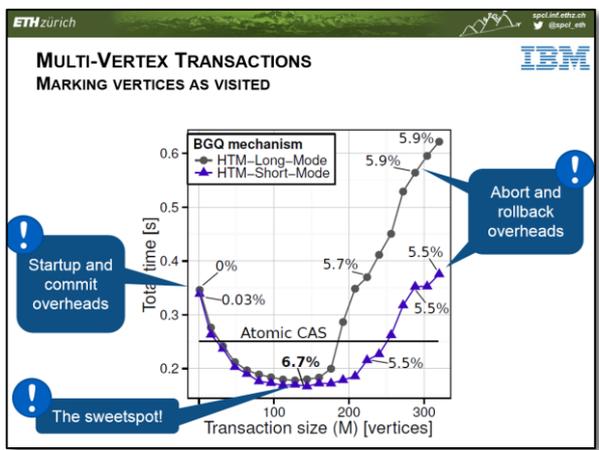


### EVALUATION CONSIDERED TYPES OF GRAPHS

- Synthetic graphs: Kronecker [1], Erdos-Renyi [2]
- Real-world SNAP graphs [3]: Social networks, Road networks, Comm. graphs, Citation graphs, Web graphs, Purchase networks

11. P. Kronecker and R. Melnikoff, Graphs as Approximations to Reality, in Graphs, Springer, 2005.  
 12. J. Lesk and A. Sany, Graphs and their Applications, in Graphs, Springer, 2005.  
 13. Stanford University, SNAP Datasets, 2015.

# Thank you for your attention



Total number of aborts vs Transaction size [M] [vertices]

Percentage of the aborts vs Concurrent threads (T)

Memory conflicts vs Concurrent threads (T)

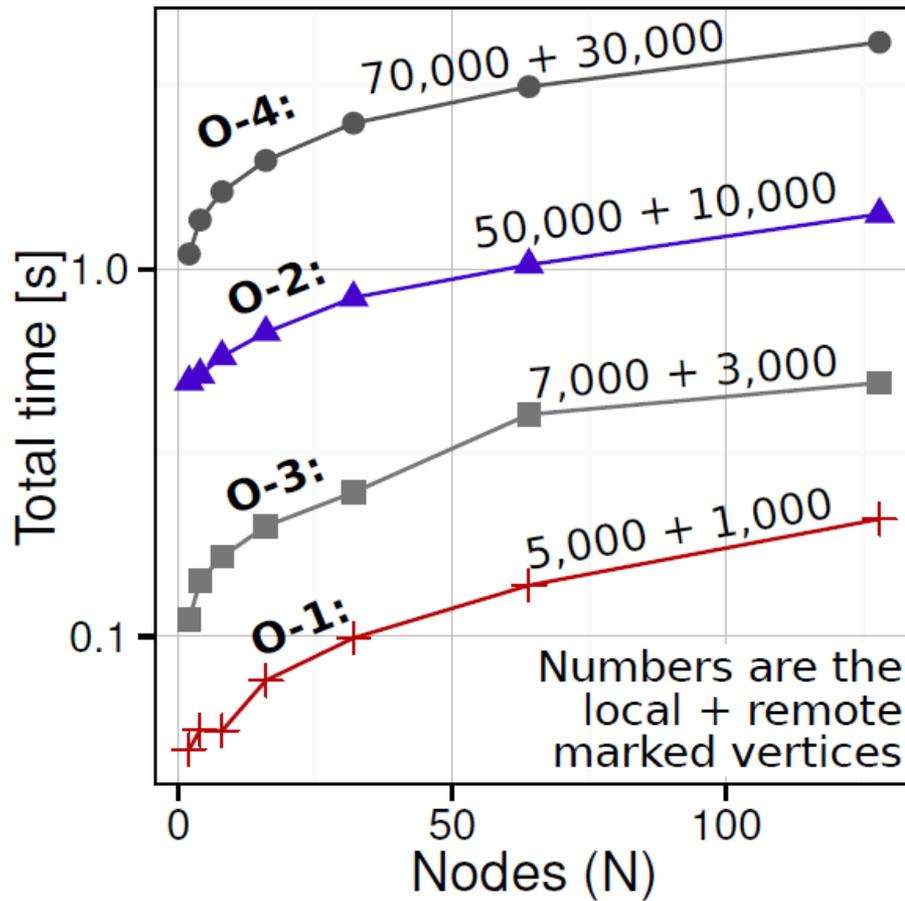
HTM buffer overflows vs Concurrent threads (T)

Scalability analysis: Distributed-memory

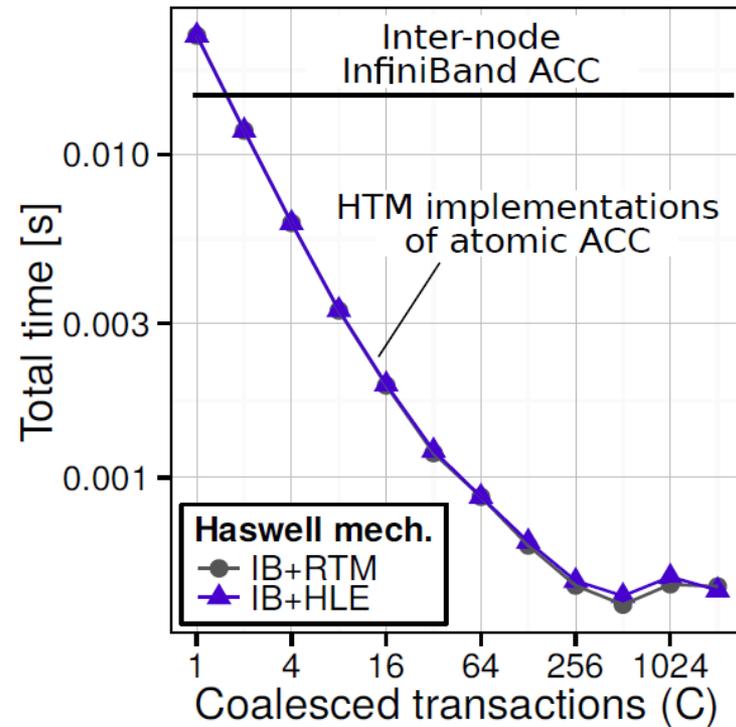
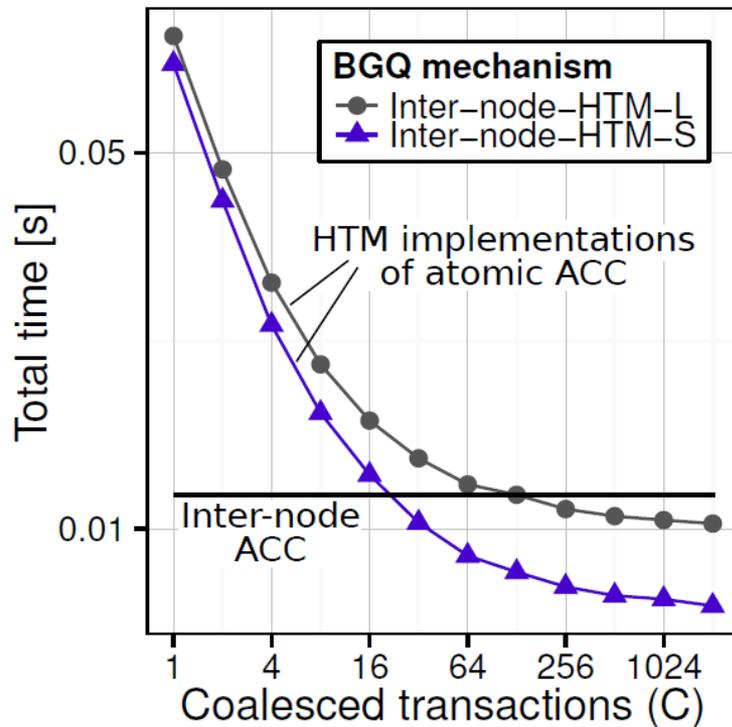
Nodes (N) vs Total time [s]

Threads/node (T) vs Total time [s]

# DISTRIBUTED HTM TRANSACTIONS

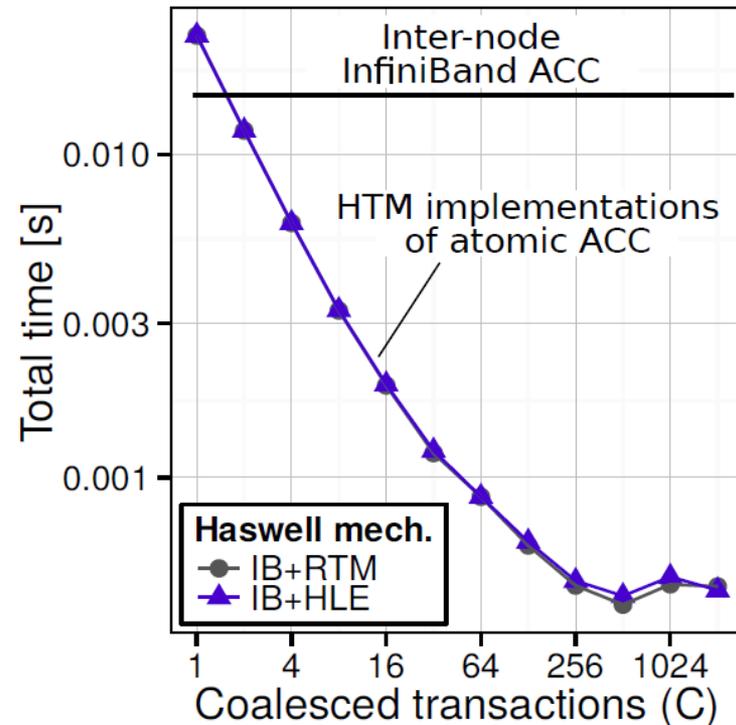
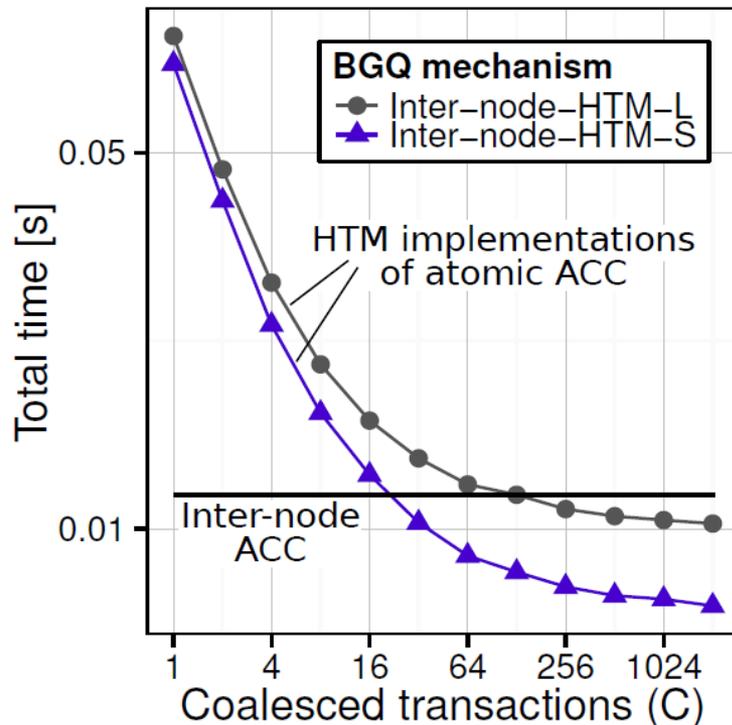


# TRANSFERRING TRANSACTIONS INCREMENTING RANKS OF VERTICES



# TRANSFERRING TRANSACTIONS INCREMENTING RANKS OF VERTICES

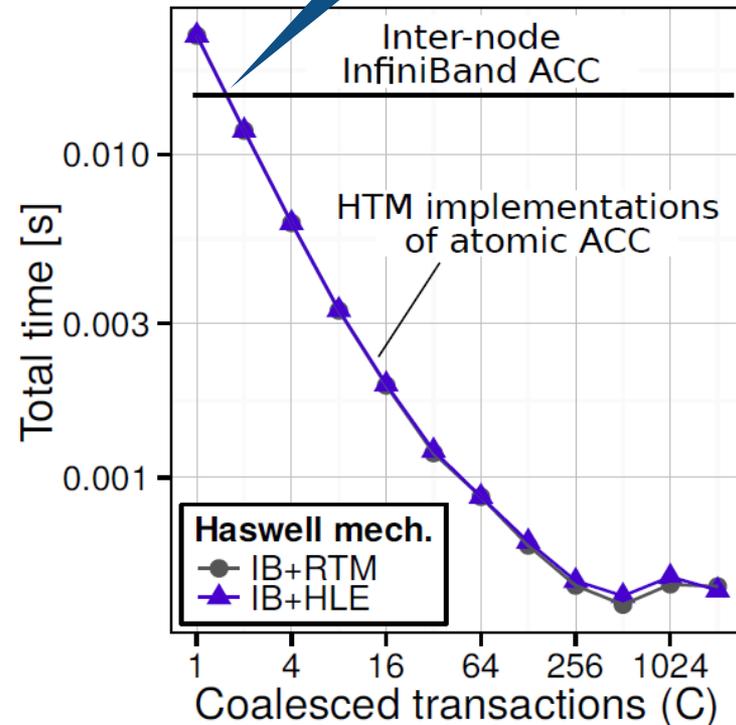
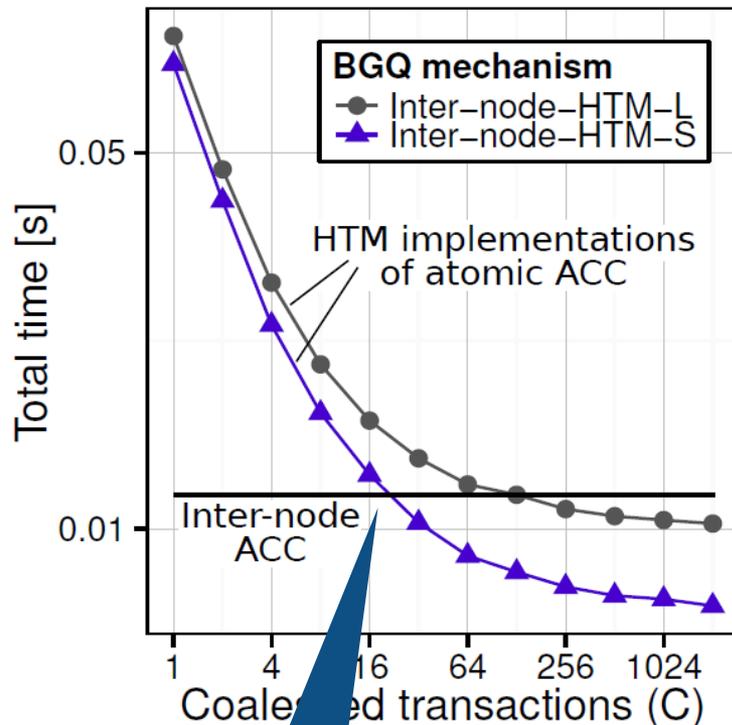
- Can we amortize HTM transactions' transfer overheads with coalescing?



# TRANSFERRING TRANSACTIONS INCREMENTING RANKS OF VERTICES

- Can we amortize HTM transactions' transfer overheads with coalescing?

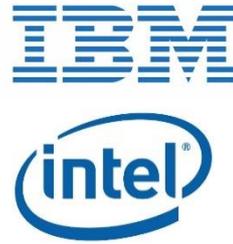
! Yes, we can!



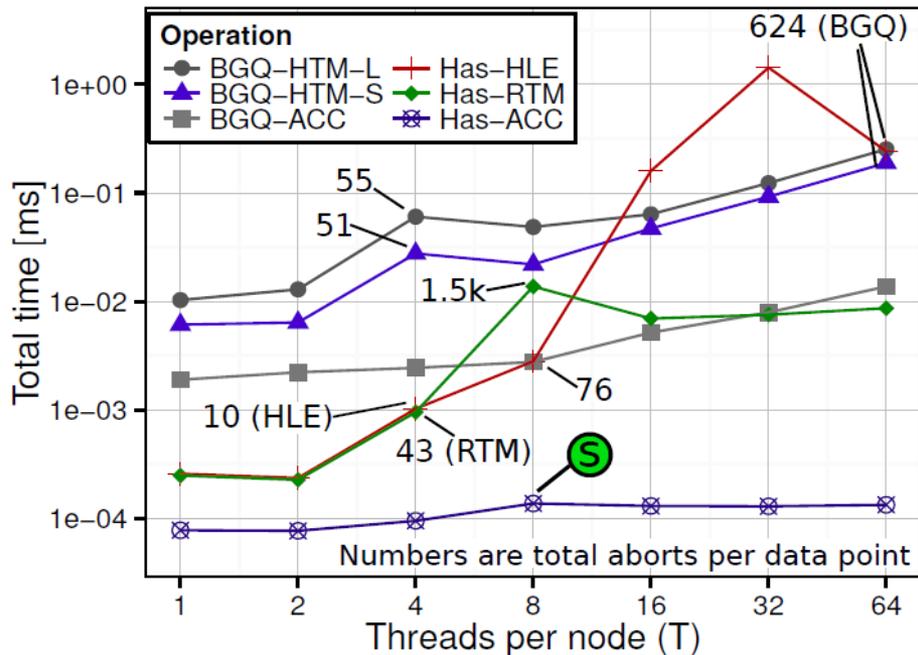
! Yes, we can!

# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

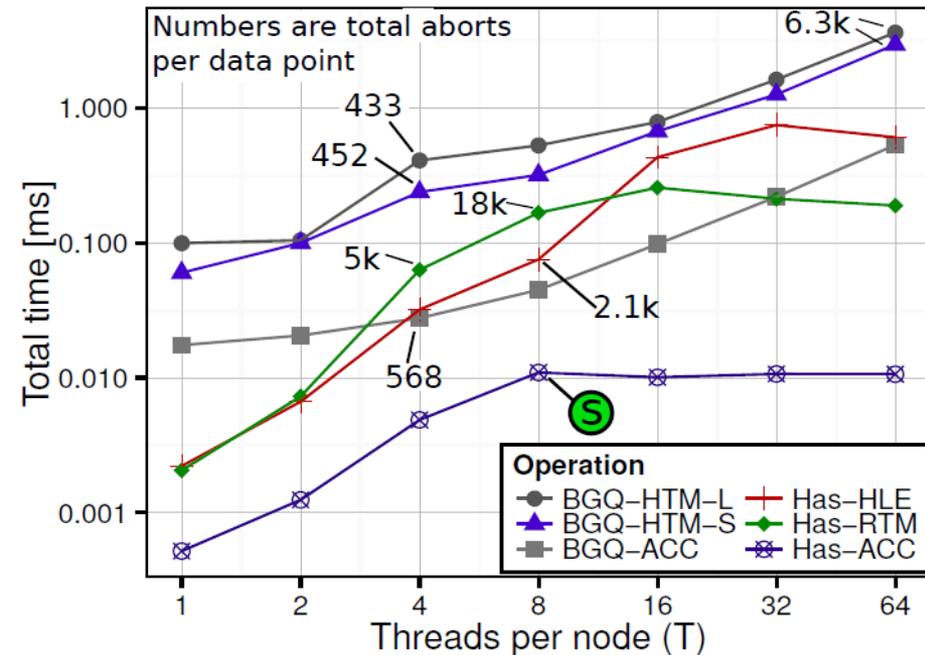
Used in  
PageRank



Lower contention  
(10 accesses/vertex)



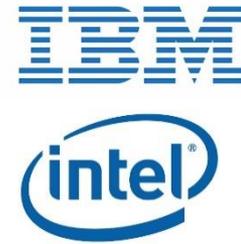
Higher contention  
(100 accesses/vertex)



# SINGLE-VERTEX TRANSACTIONS

## INCREMENTING VERTEX RANK

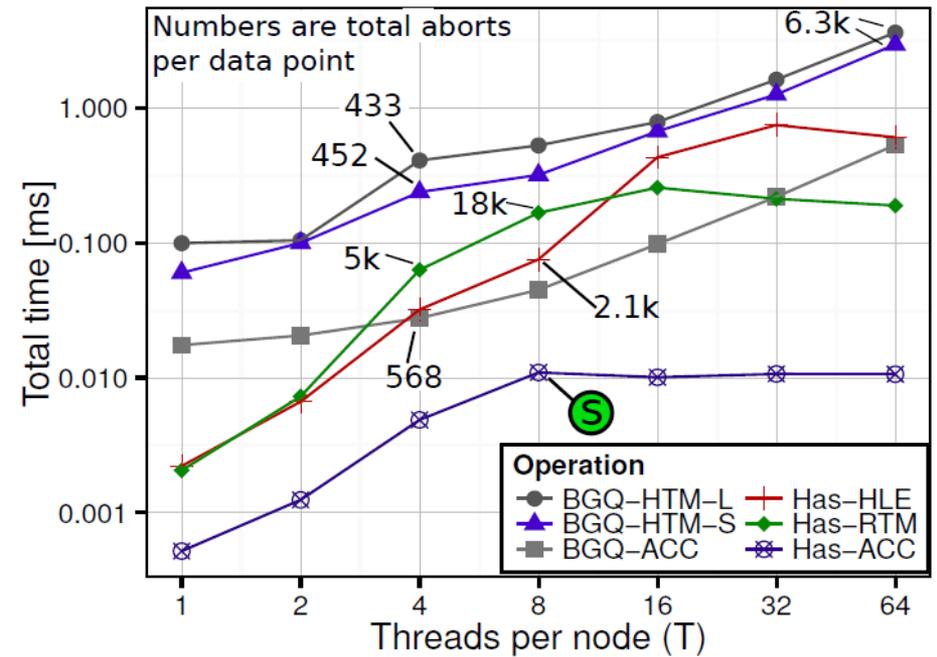
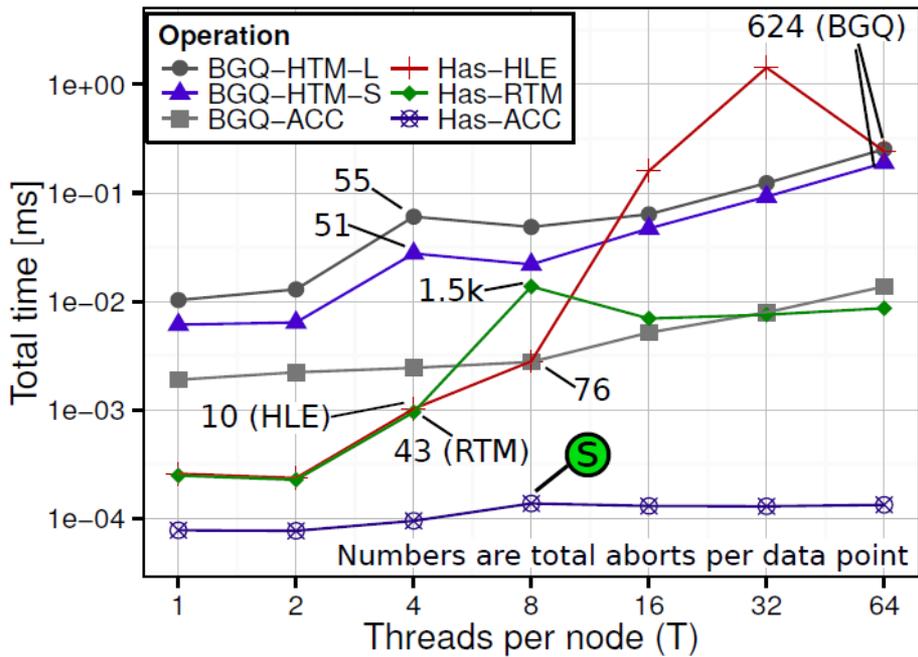
Used in  
PageRank



**!** Atomics always outperform HTM

Lower contention  
(10 accesses/vertex)

Higher contention  
(100 accesses/vertex)



# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

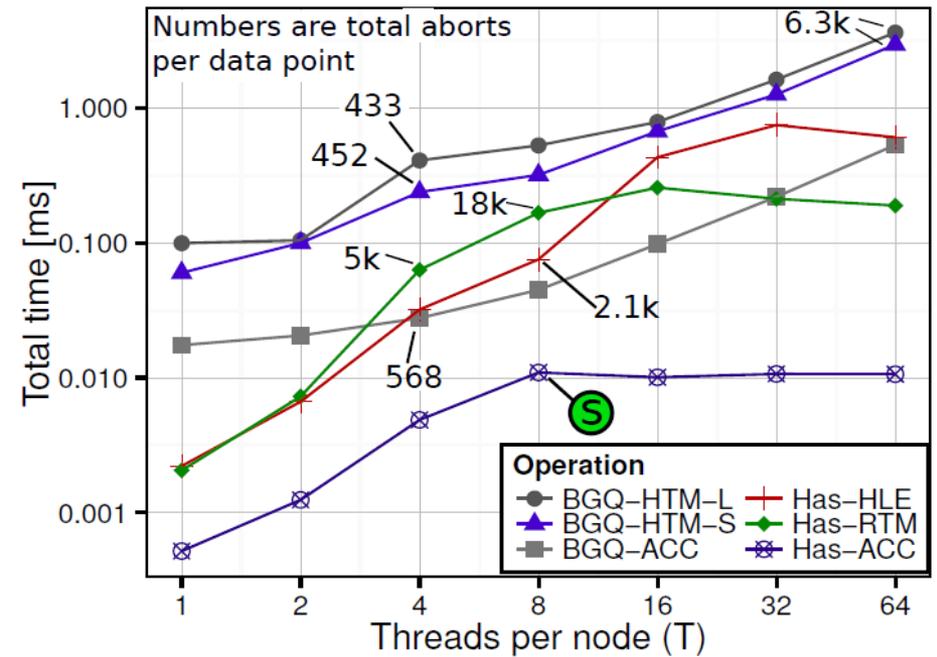
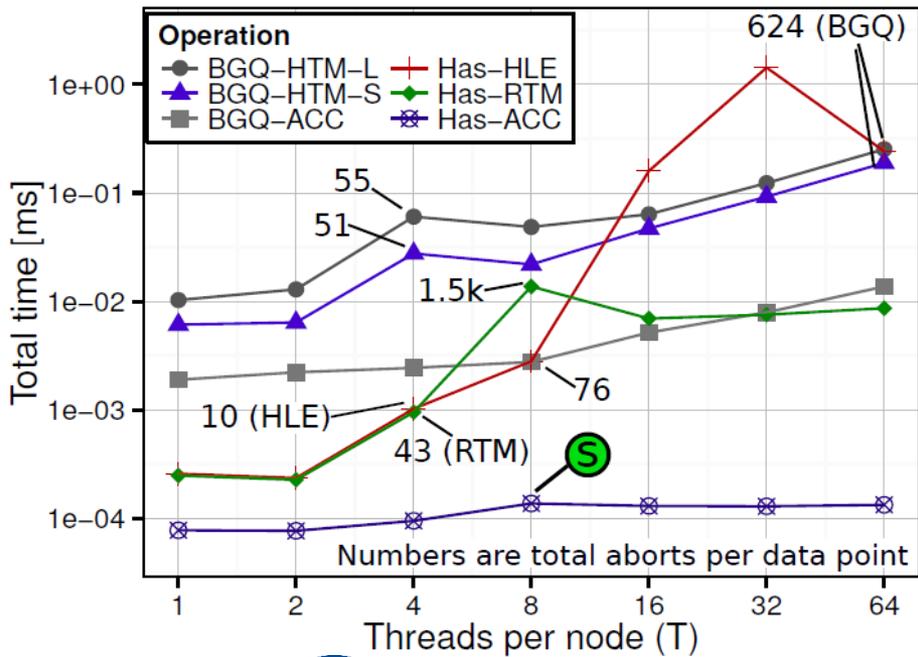
Used in  
PageRank



! Atomics always outperform HTM

Lower contention  
(10 accesses/vertex)

Higher contention  
(100 accesses/vertex)



! The reason: each transaction always modifies some memory cell, increasing the number of conflicts

# SINGLE-VERTEX TRANSACTIONS INCREMENTING VERTEX RANK

Used in  
PageRank

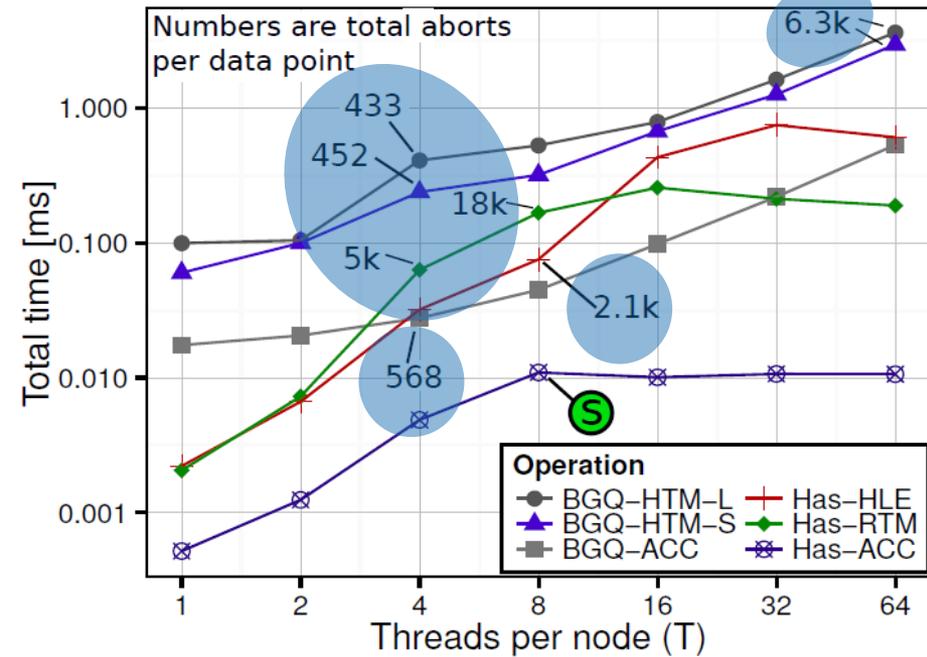
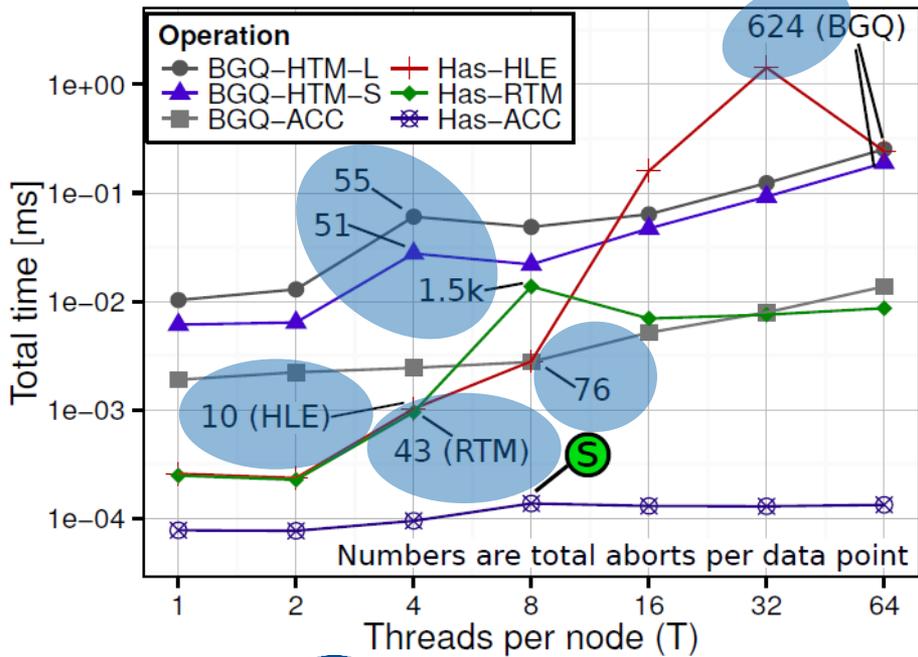


More  
aborts

Lower contention  
(10 accesses/vertex)

Atomics always  
outperform HTM

Higher contention  
(100 accesses/vertex)



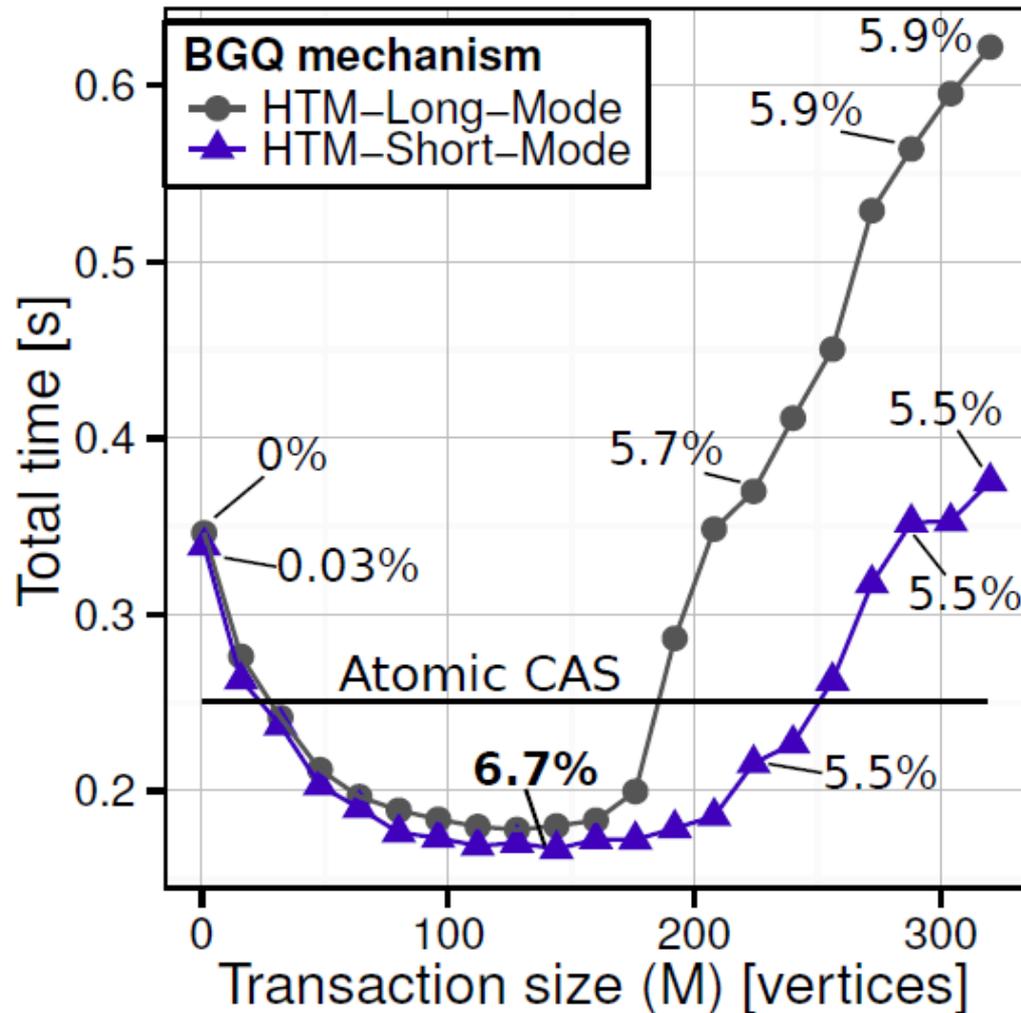
The reason: each transaction always modifies some memory cell, increasing the number of conflicts

# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED

# MULTI-VERTEX TRANSACTIONS

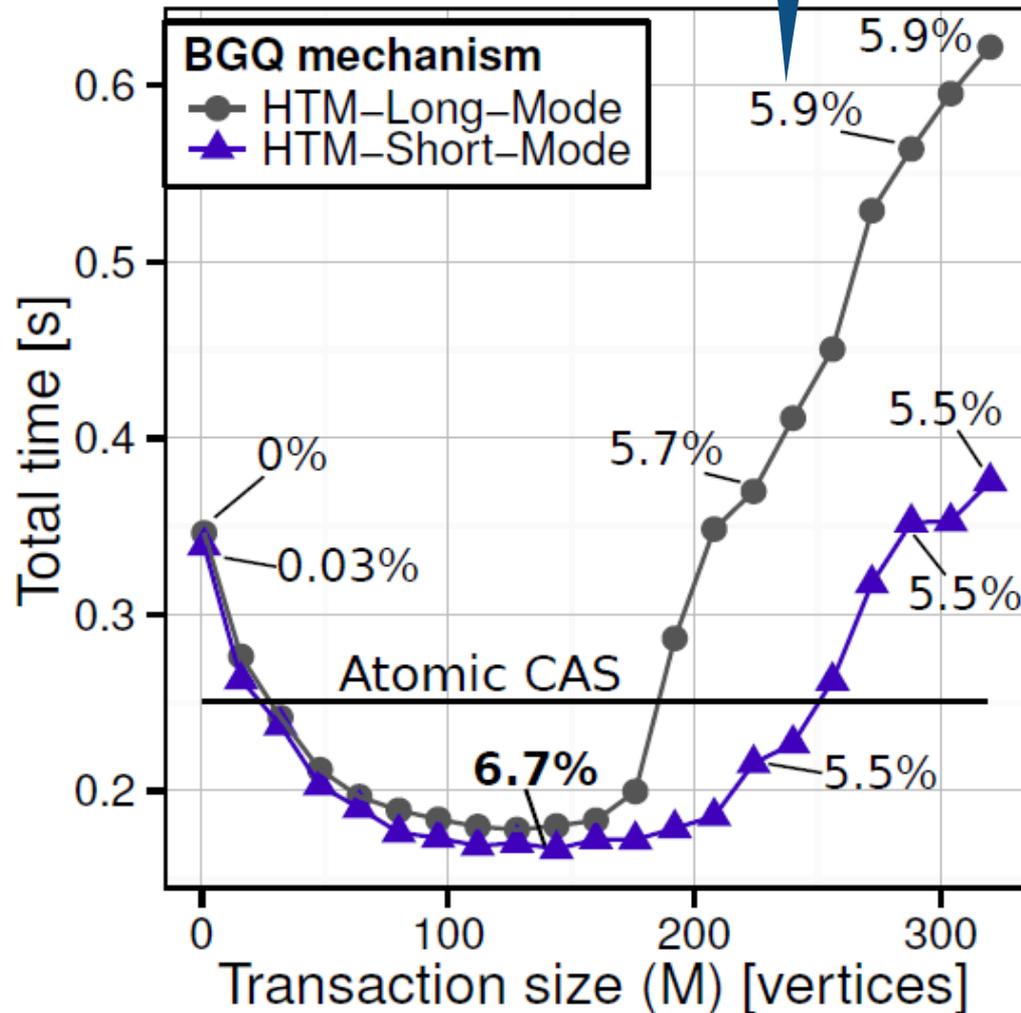
## MARKING VERTICES AS VISITED



# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED

Numbers: % of aborts due to the lack of HTM resources + memory conflicts

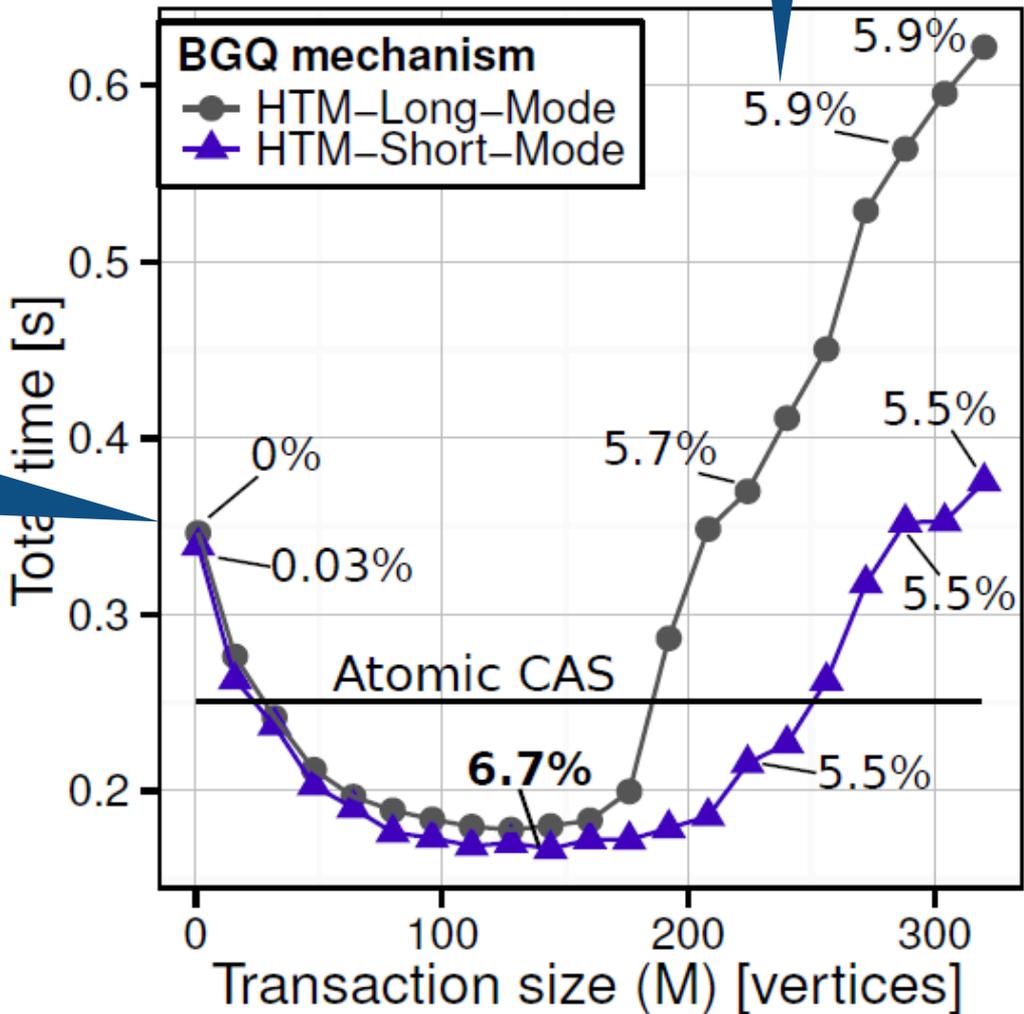


# MULTI-VERTEX TRANSACTIONS

## MARKING VERTICES AS VISITED

Numbers: % of aborts due to the lack of HTM resources + memory conflicts

Startup and commit overheads

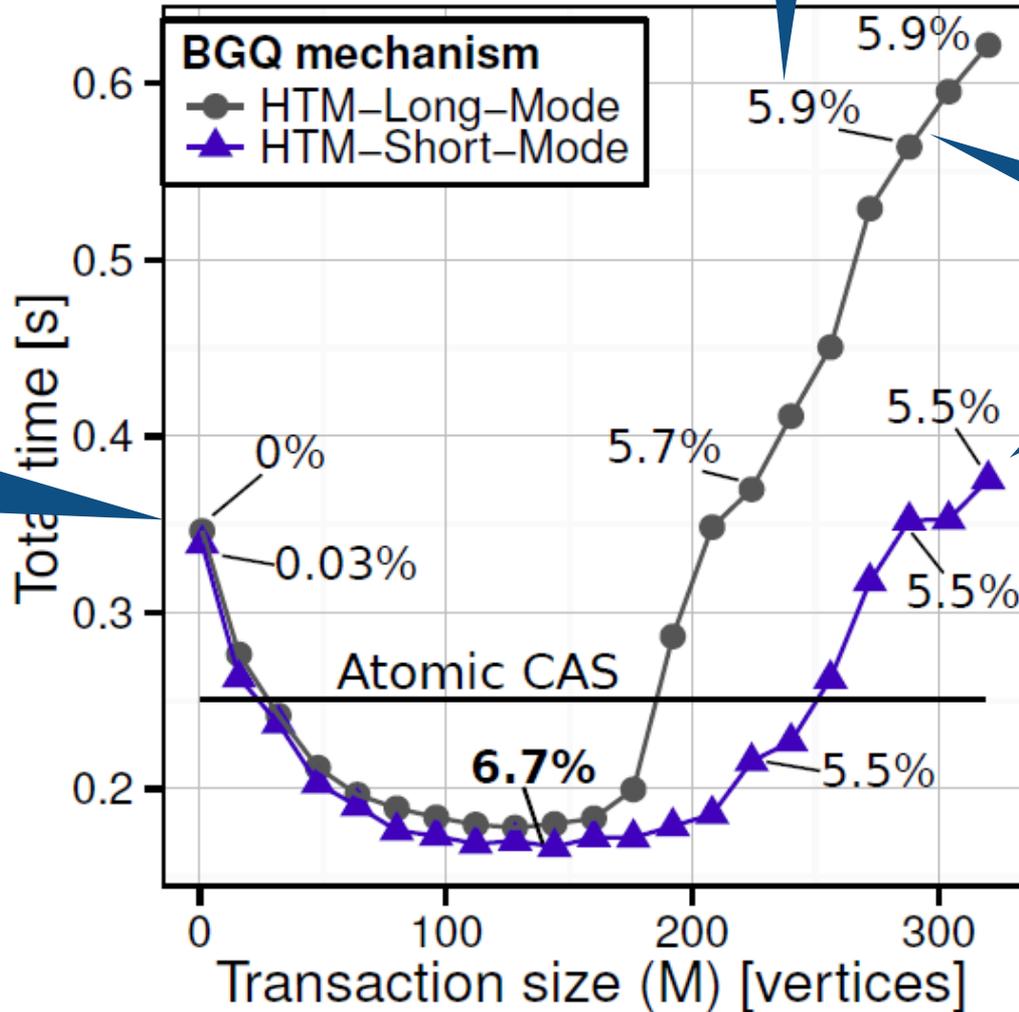


# MULTI-VERTEX TRANSACTIONS MARKING VERTICES AS VISITED

Numbers: % of aborts due to the lack of HTM resources + memory conflicts

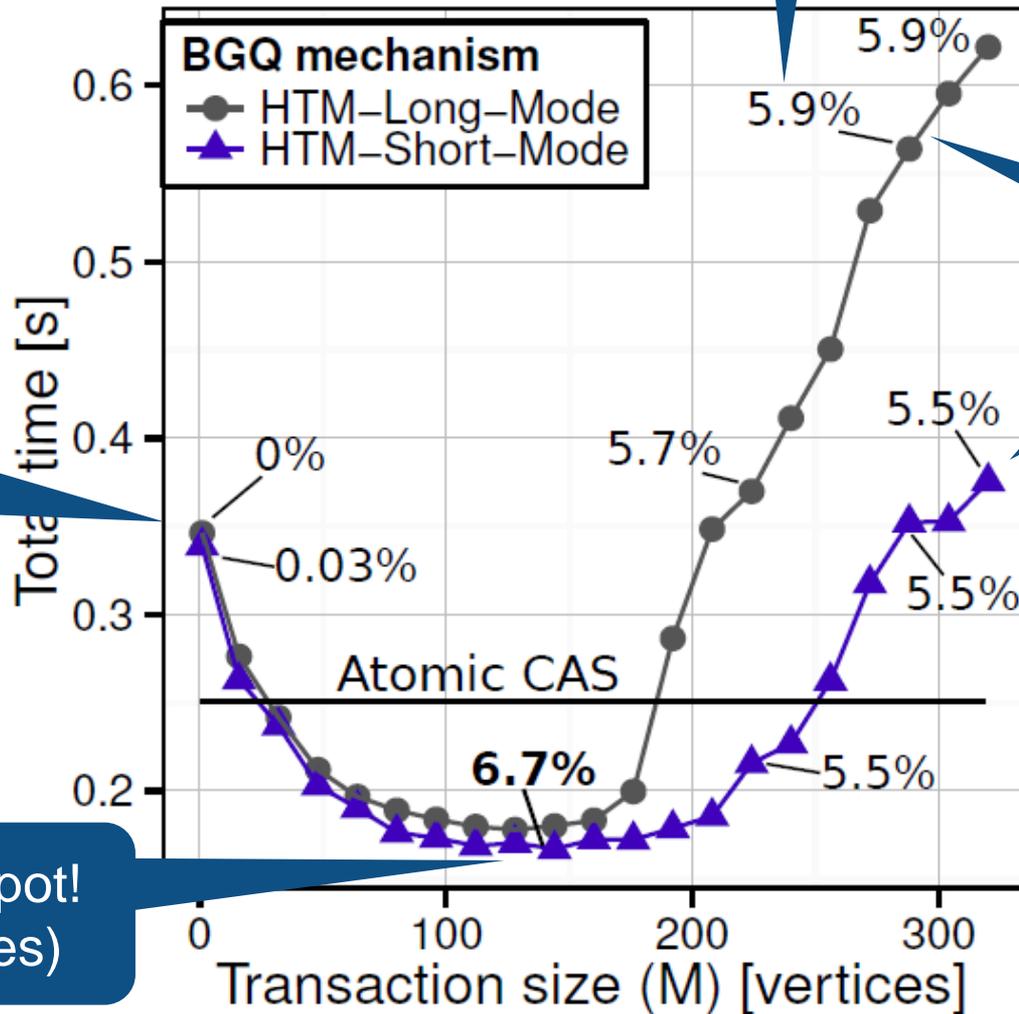
Startup and commit overheads

Abort and rollback overheads



# MULTI-VERTEX TRANSACTIONS MARKING VERTICES AS VISITED

Numbers: % of aborts due to the lack of HTM resources + memory conflicts



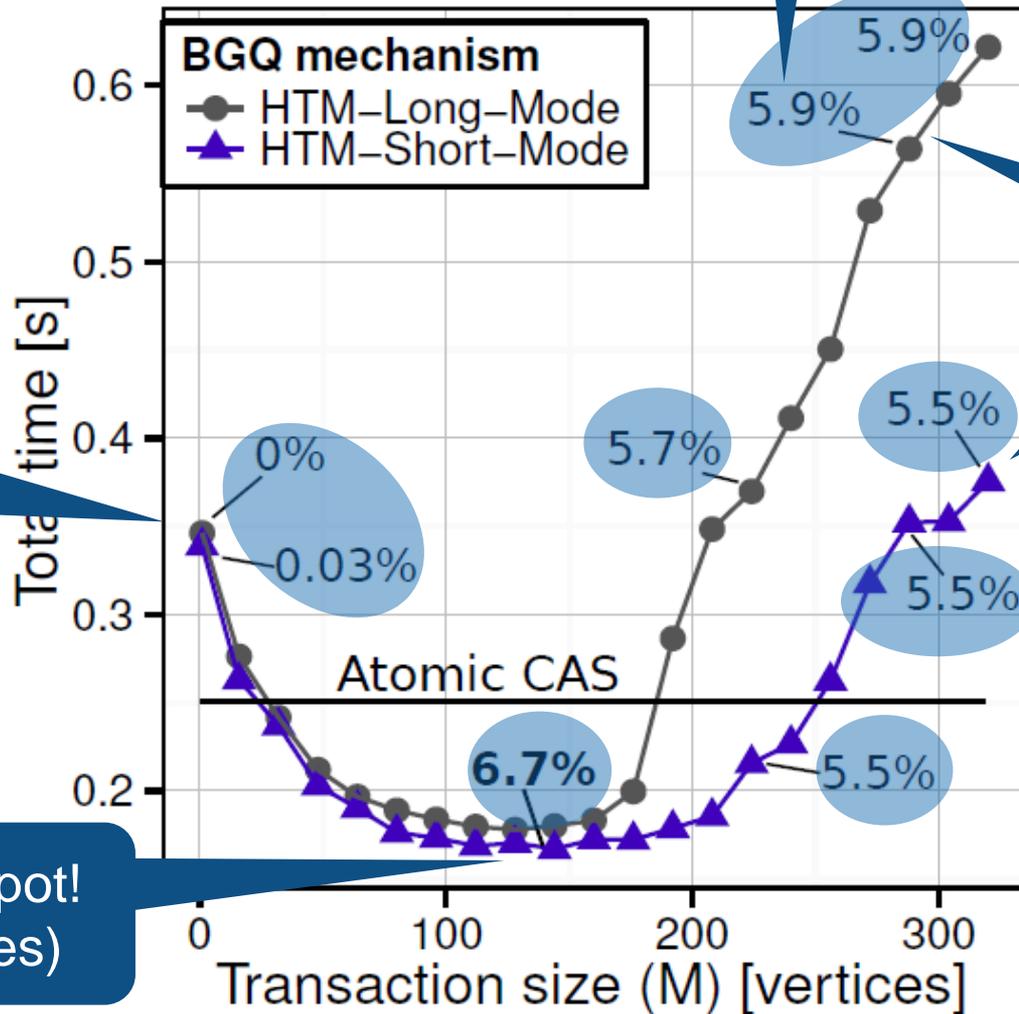
Startup and commit overheads

The sweetspot! (144 vertices)

Abort and rollback overheads

# MULTI-VERTEX TRANSACTIONS MARKING VERTICES AS VISITED

Numbers: % of aborts due to the lack of HTM resources + memory conflicts



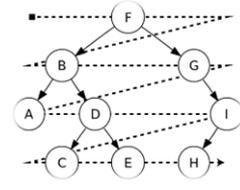
Startup and commit overheads

The sweetspot!  
(144 vertices)

Abort and rollback overheads

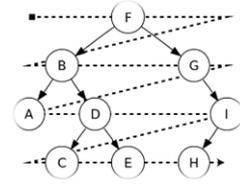
Not too many aborts due to the lack of HW resources (large cache size & associativity)

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



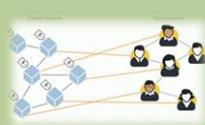
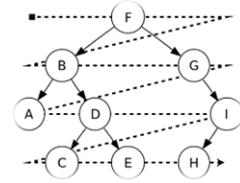
# OUTPERFORMING STATE-OF-THE-ART

## BLUEGENE/Q



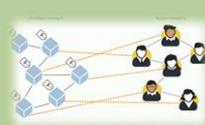
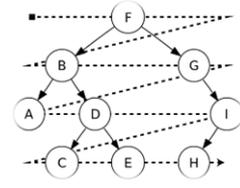
Average overall speedup over Graph500  
(geometric mean): 1.51 (1.85)

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



Average overall speedup over Graph500  
(geometric mean): 1.51 (1.85)

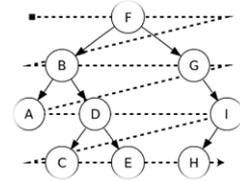
# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



Average overall speedup over Graph500  
(geometric mean): 1.51 (1.85)

The same  
transaction  
size for all  
graphs

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



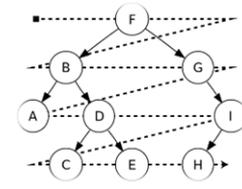
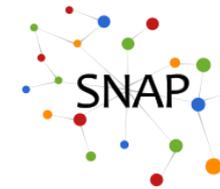
Average overall speedup over Graph500  
(geometric mean): 1.51 (1.85)



The same transaction size for all graphs

The same transaction sizes for each graph separately

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q




Average speedup: 1

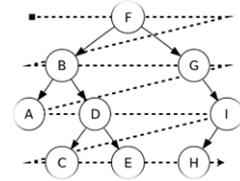


Average overall speedup over Graph500 (geometric mean): 1.51 (1.85)

The same transaction size for all graphs

The same transaction sizes for each graph separately

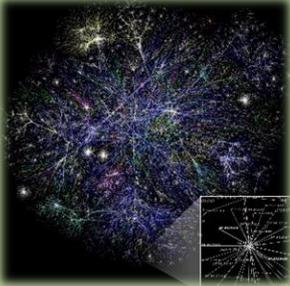
# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q




Average speedup: 1



Average overall speedup over Graph500 (geometric mean): 1.51 (1.85)

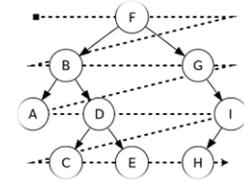


Average speedup: 1.85

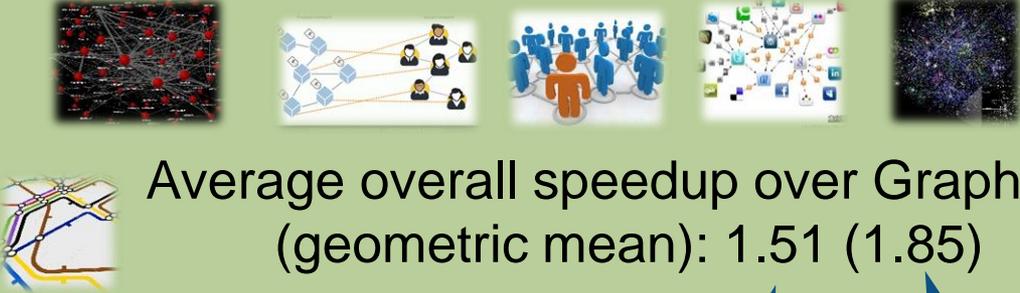
The same transaction size for all graphs

The same transaction sizes for each graph separately

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



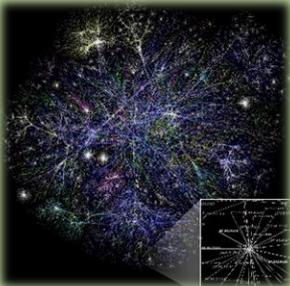

Average speedup: 1



Average overall speedup over Graph500 (geometric mean): 1.51 (1.85)



Average speedup: 3.20

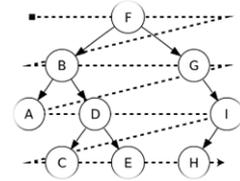


Average speedup: 1.85

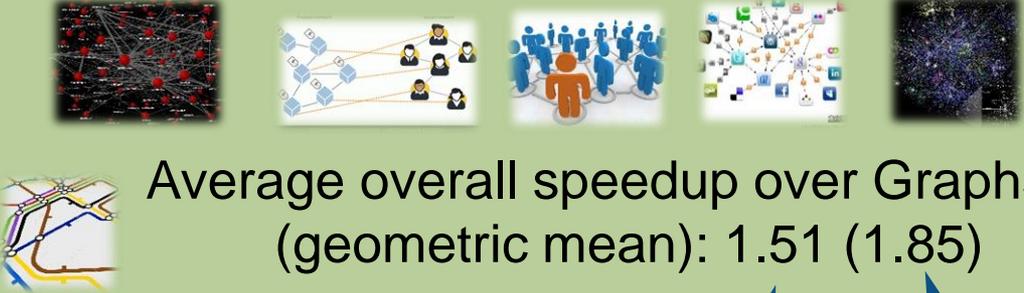
The same transaction size for all graphs

The same transaction sizes for each graph separately

# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



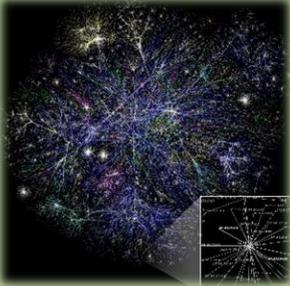

Average speedup: 1



Average overall speedup over Graph500 (geometric mean): 1.51 (1.85)



Average speedup: 3.20



Average speedup: 1.85

The same transaction size for all graphs

The same transaction sizes for each graph separately

 Best transaction size: ~24-100 vertices accessed

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

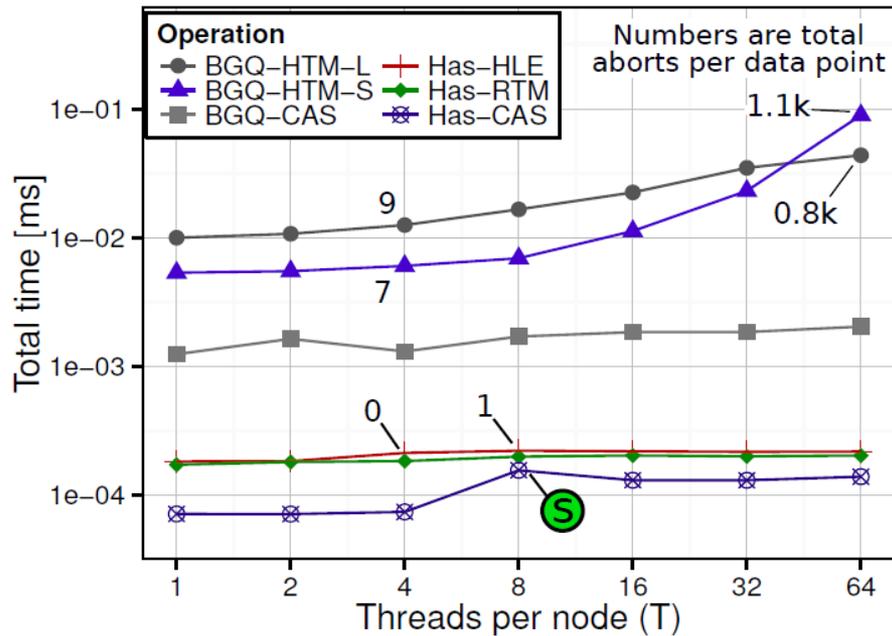
Used in BFS,  
SSSP, ...

# SINGLE-VERTEX TRANSACTIONS

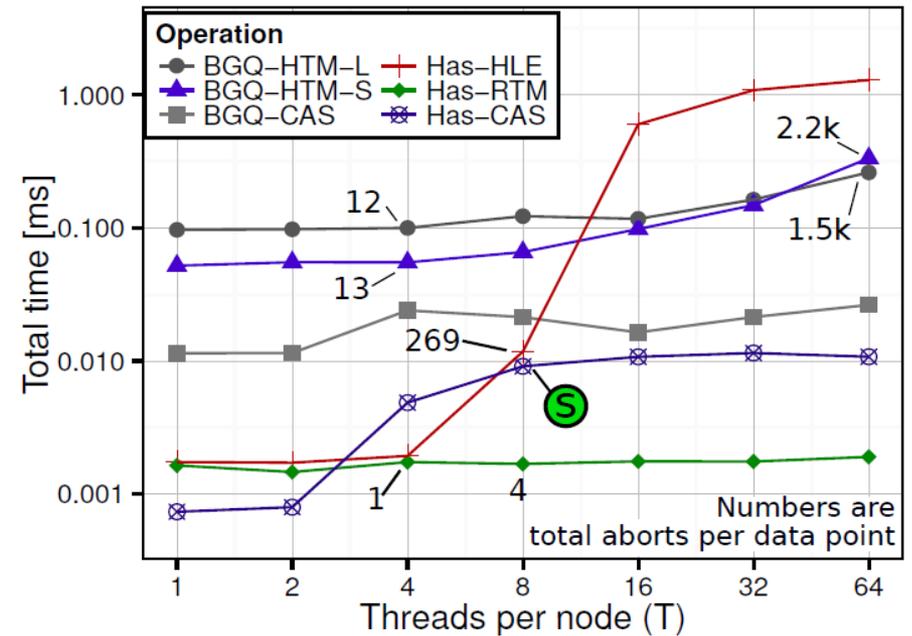
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 accesses/vertex)



Higher contention  
(100 accesses/vertex)

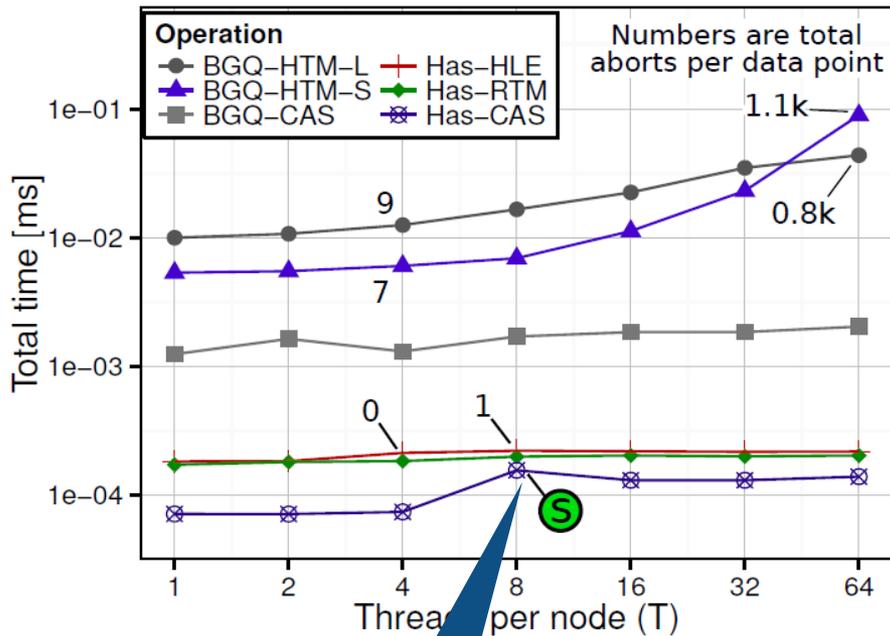


# SINGLE-VERTEX TRANSACTIONS

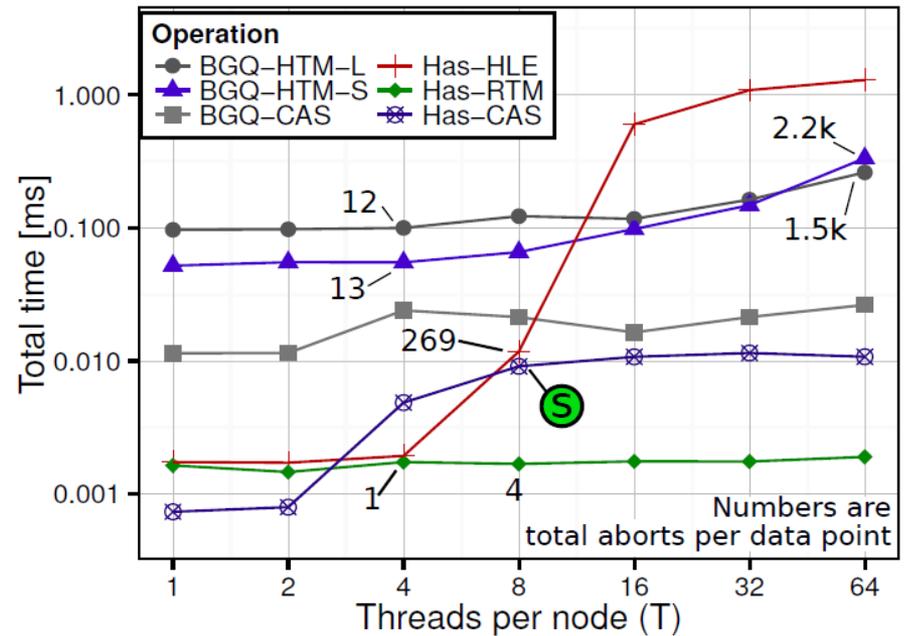
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 accesses/vertex)



Higher contention  
(100 accesses/vertex)



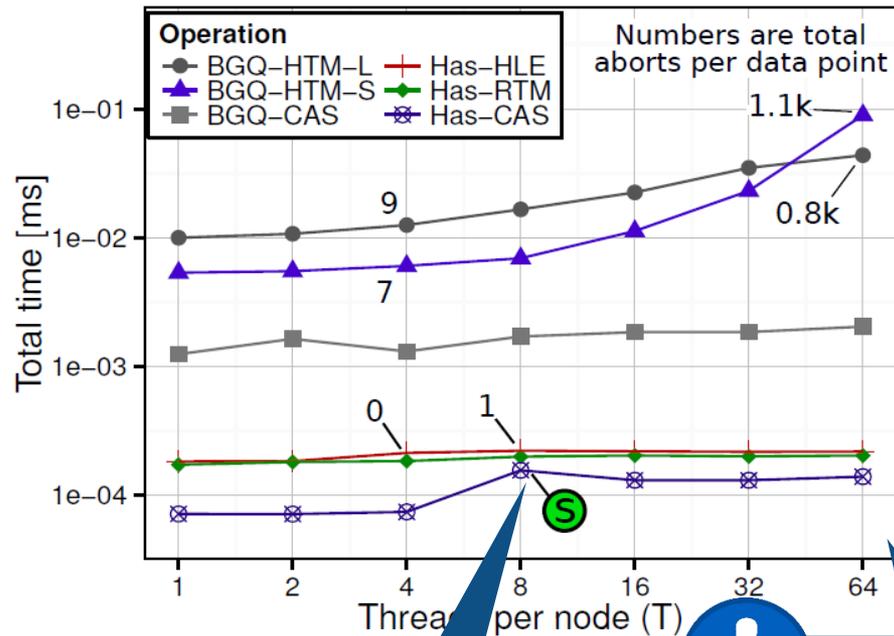
! Atomics (CAS) slightly faster than HTM

# SINGLE-VERTEX TRANSACTIONS

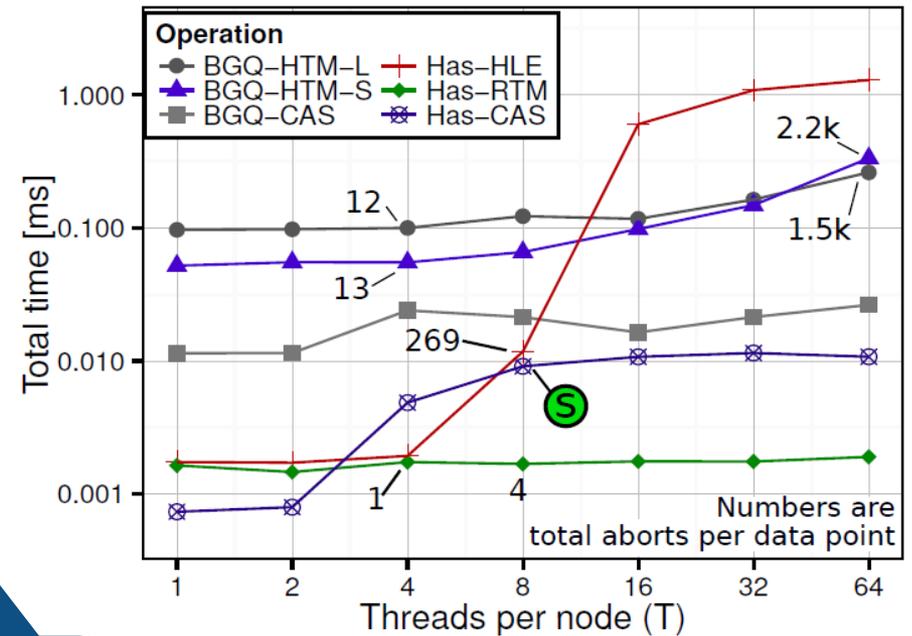
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 accesses/vertex)



Higher contention  
(100 accesses/vertex)



Atomics (CAS) slightly faster than HTM

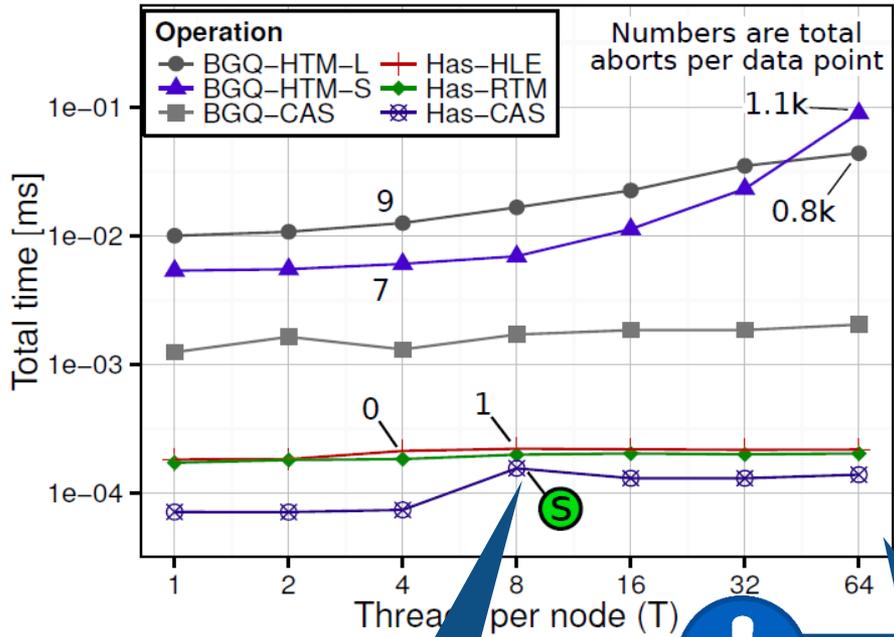
Commit overheads dominate

# SINGLE-VERTEX TRANSACTIONS

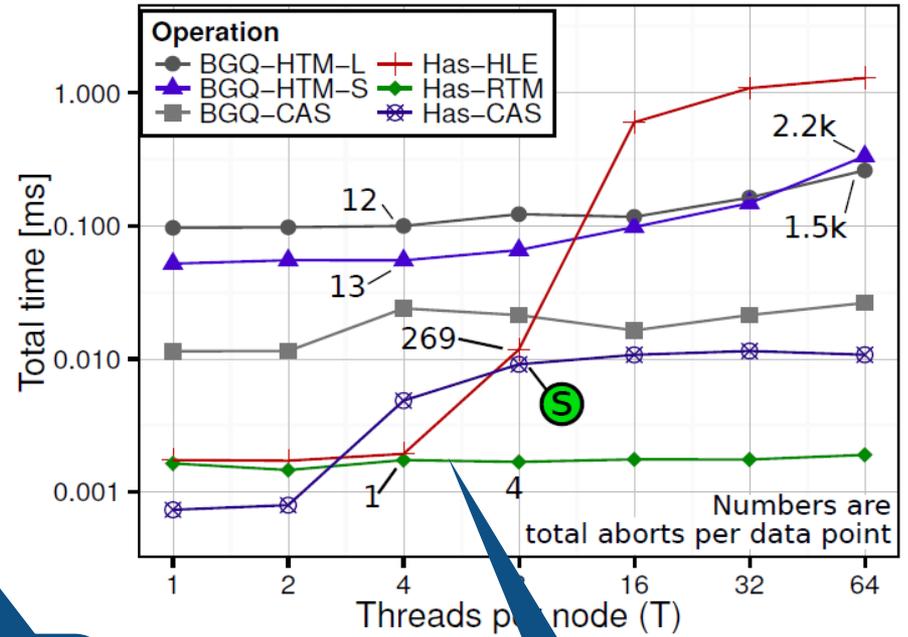
## MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 accesses/vertex)



Higher contention  
(100 accesses/vertex)



! Atomics (CAS) slightly faster than HTM

! Commit overheads dominate

! RTM outperforms other (overcontended) targets

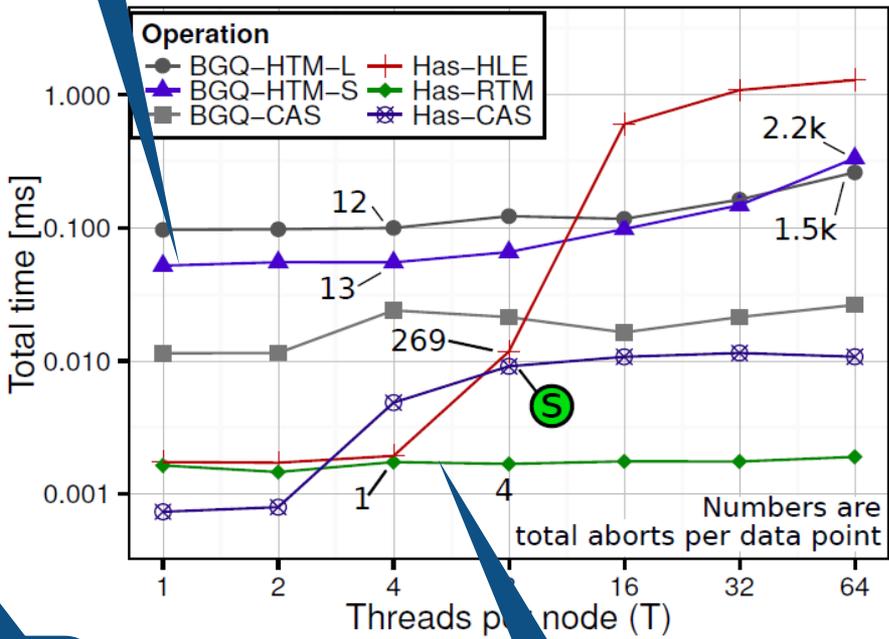
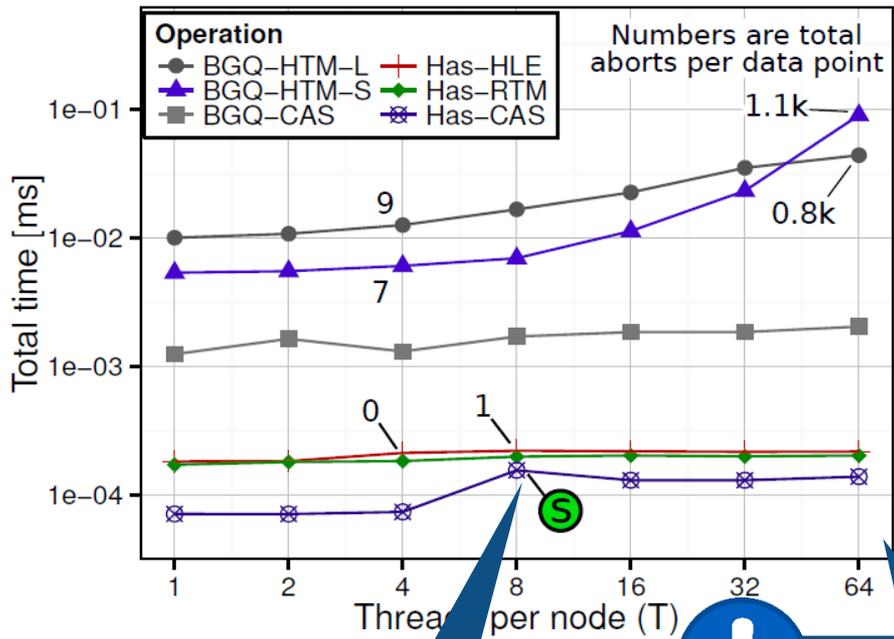
# SINGLE-VERTEX TRANSACTIONS MARKING A VERTEX AS VISITED

Used in BFS,  
SSSP, ...

Lower contention  
(10 accesses/vertex)

! BG/Q HTM still worse (L1 vs L2 matters!)

Higher contention  
(100 accesses/vertex)



! Atomics (CAS) slightly faster than HTM

! Commit overheads dominate

! RTM outperforms other (overcontended) targets

# SINGLE-VERTEX TRANSACTIONS

## MARKING A VERTEX AS VISITED

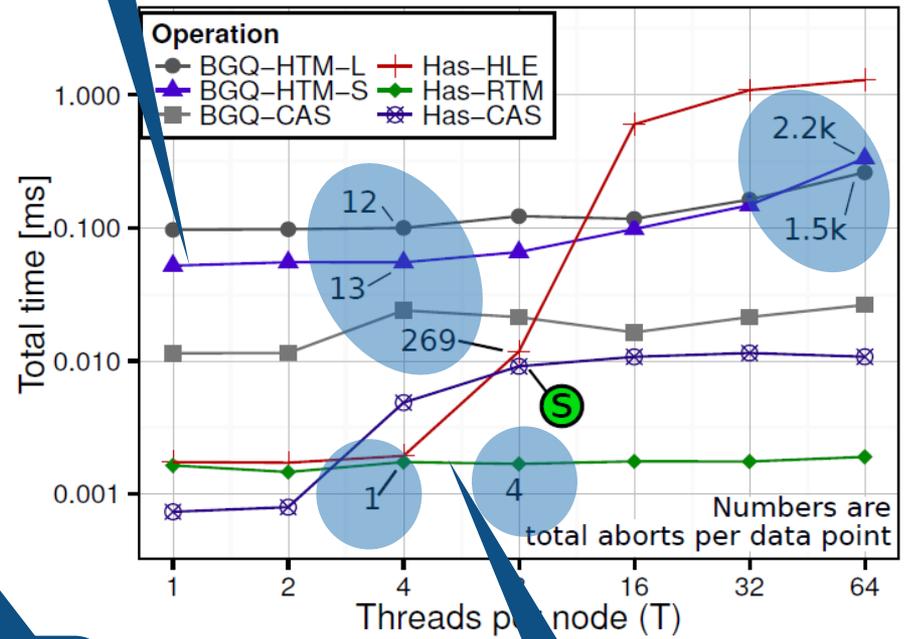
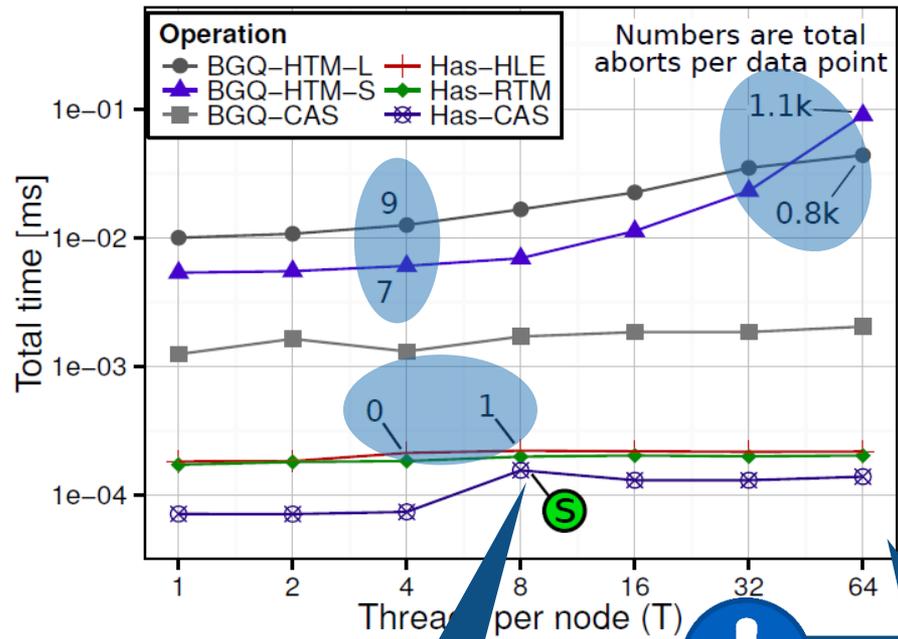
Used in BFS, SSSP, ...

Very few aborts

Lower contention (10 accesses/vertex)

BG/Q HTM still worse (L1 vs L2 matters!)

Higher contention (100 accesses/vertex)

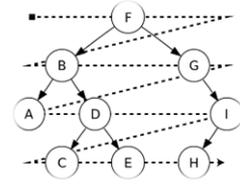


Atomics (CAS) slightly faster than HTM

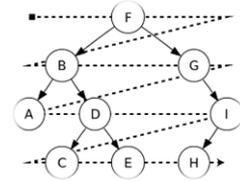
Commit overheads dominate

RTM outperforms other (overcontended) targets

# OUTPERFORMING STATE-OF-THE-ART HASWELL



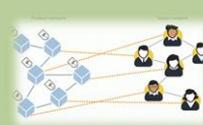
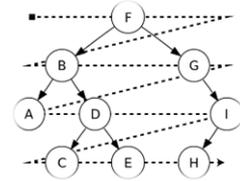
# OUTPERFORMING STATE-OF-THE-ART HASWELL



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



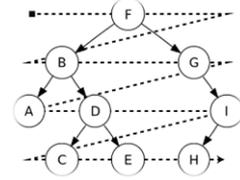
# OUTPERFORMING STATE-OF-THE-ART HASWELL



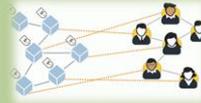
Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



# OUTPERFORMING STATE-OF-THE-ART HASWELL



Average  
speedup: 1

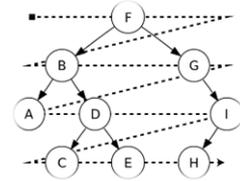


Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



# OUTPERFORMING STATE-OF-THE-ART

## HASWELL



Average speedup: 1



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000

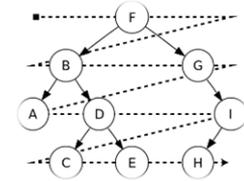


Average speedup: 1.85

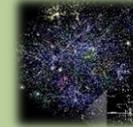
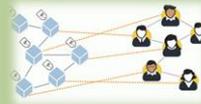


# OUTPERFORMING STATE-OF-THE-ART

## HASWELL



Average speedup: 1



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



Average speedup: 3.20

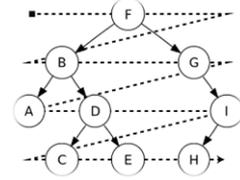


Average speedup: 1.85



# OUTPERFORMING STATE-OF-THE-ART

## HASWELL



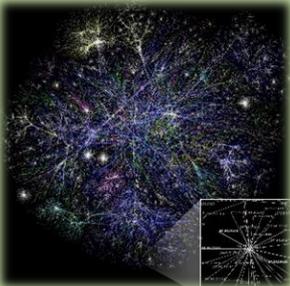

Average speedup: 1



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



Average speedup: 3.20



Average speedup: 1.85

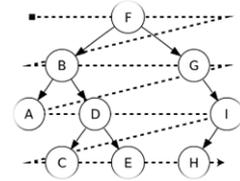
1.85x on average, up to 4.3x




Best transaction size: ~2-9 vertices accessed

# OUTPERFORMING STATE-OF-THE-ART

## HASWELL



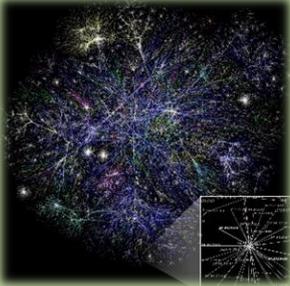

Average speedup: 1



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



Average speedup: 3.20



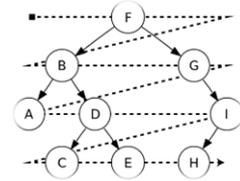
Average speedup: 1.85

! 1.85x on average, up to 4.3x



! Best transaction size: ~2-9 vertices accessed

# OUTPERFORMING STATE-OF-THE-ART HASWELL




Average speedup: 1



Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



Average speedup: 3.20



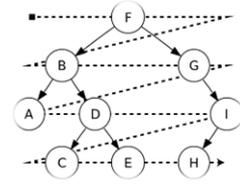
Average speedup: 1.85

! 1.85x on average, up to 4.3x



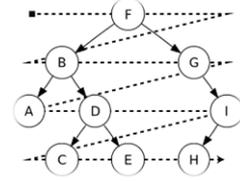

! Best transaction size: ~2-9 vertices accessed

# OUTPERFORMING STATE-OF-THE-ART HASWELL



Best transaction size:  
~4 vertices   
~14 vertices

# OUTPERFORMING STATE-OF-THE-ART HASWELL

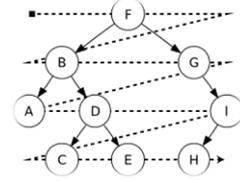


Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



Best transaction size:  
~4 vertices   
~14 vertices

# OUTPERFORMING STATE-OF-THE-ART HASWELL




Average overall speedup (geometric mean) over Graph500: 1.07, Galois: 1.40, HAMA ~1000



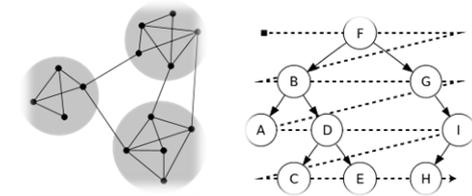

1.85x on average, up to 4.3x

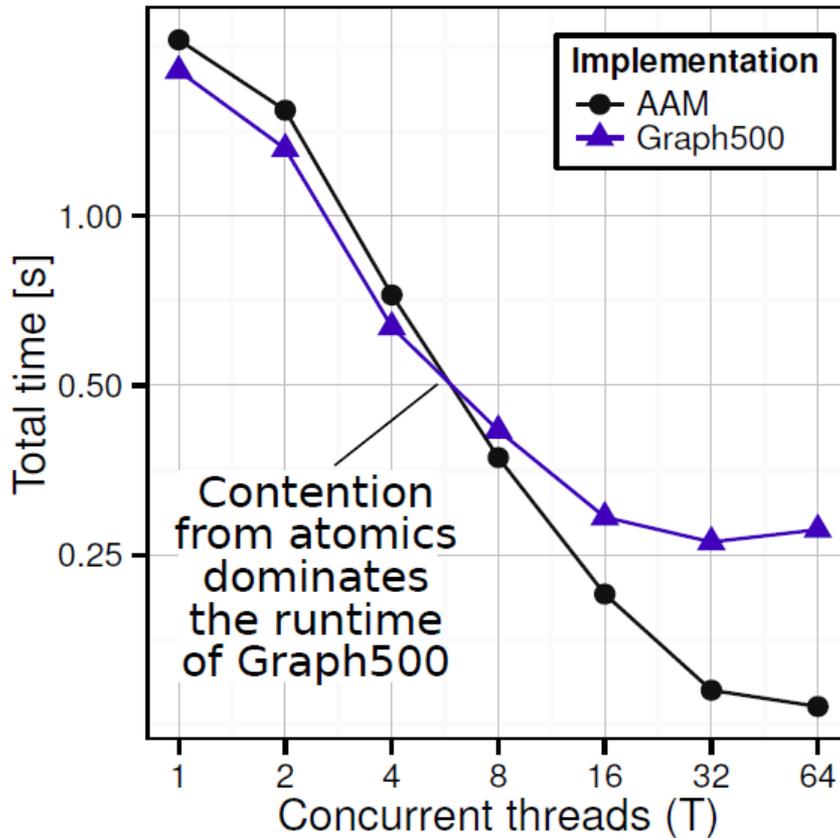


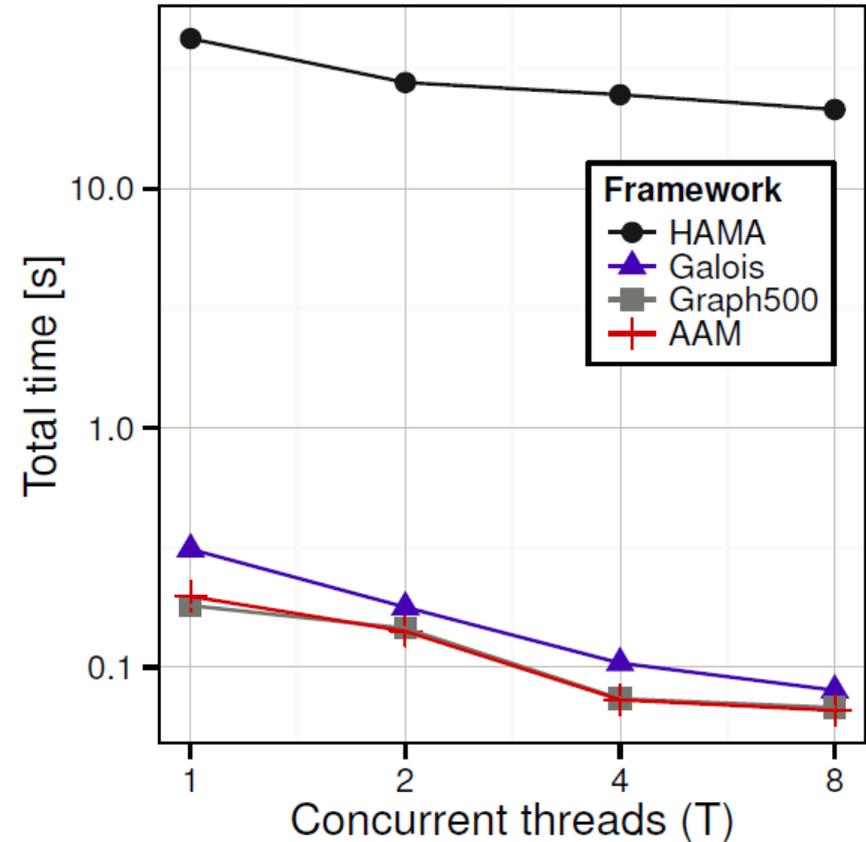


Best transaction size:  
~4 vertices   
~14 vertices 

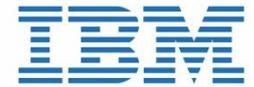
# OUTPERFORMING STATE-OF-THE-ART SCALABILITY ANALYSIS: SHARED-MEMORY







# OUTPERFORMING STATE-OF-THE-ART BLUEGENE/Q



Input graph properties					BG/Q analysis		
Type	ID	Name	$ V $	$ E $	$S$ over g500 ( $M = 24$ )	$M$	$S$ over g500
Comm. networks (CNs)	cWT	wiki-Talk	2.4M	5M	2.82	48	3.35
	cEU	email-EuAll	265k	420k	3.67	32	4.36
Social networks (SNs)	sLV	soc-LiveJ.	4.8M	69M	1.44	12	1.56
	sOR	com-orkut	3M	117M	1.22	20	1.27
	sLJ	com-lj	4M	34M	1.44	12	1.54
	sYT	com-youtube	1.1M	2.9M	1.67	8	1.84
	sDB	com-dblp	317k	1M	1.33	8	1.80
	sAM	com-amazon	334k	925k	1.14	8	1.62
Purchase network (PNs)	pAM	amazon0601	403k	3.3M	1.45	8	1.91
Road networks (RNs)	rCA	roadNet-CA	1.9M	5.5M	$\approx 1$	2	1.59
	rTX	roadNet-TX	1.3M	3.8M	$\approx 1$	2	1.53
	rPA	roadNet-PA	1M	3M	$\approx 1$	2	1.52
Citation graphs (CGs)	ciP	cit-Patents	3.7M	16.5M	1.16	8	1.57
Web graphs (WGs)	wGL	web-Google	875k	5.1M	1.78	12	2.08
	wBS	web-BerkStan	685k	7.6M	1.91	24	1.91
	wSF	web-Stanford	281k	2.3M	1.89	24	1.89

# OUTPERFORMING STATE-OF-THE-ART HASWELL



Input graph properties					Haswell analysis					
Type	ID	Name	$ V $	$ E $	$S$ over g500 ( $M = 2$ )	$S$ over Galois ( $M = 2$ )	$M$	$S$ over g500	$S$ over Galois	$S$ over HAMA
Comm. networks (CNs)	cWT	wiki-Talk	2.4M	5M	0.91	1.22	6	0.96	1.28	344
	cEU	email-EuAll	265k	420k	0.76	0.88	4	0.97	1.12	1448
Social networks (SNs)	sLV	soc-LiveJ.	4.8M	69M	1.05	1.1	3	1.07	1.12	$> 10^4$
	sOR	com-orkut	3M	117M	1.06	0.69	4	1.13	0.74	$> 10^4$
	sLJ	com-lj	4M	34M	1.03	1.03	4	1.04	1.04	603
	sYT	com-youtube	1.1M	2.9M	0.96	1.1	5	0.98	1.11	670
	sDB	com-dblp	317k	1M	$\approx 1$	2.5	2	$\approx 1$	2.53	2160
	sAM	com-amazon	334k	925k	1.04	1.64	2	1.04	1.64	1426
Purchase network (PNs)	pAM	amazon0601	403k	3.3M	$\approx 1$	1.25	3	1.03	1.30	618
Road networks (RNs)	rCA	roadNet-CA	1.9M	5.5M	1.33	1.74	8	1.38	1.80	$> 10^4$
	rTX	roadNet-TX	1.3M	3.8M	1.29	1.89	6	1.42	2.08	$> 10^4$
	rPA	roadNet-PA	1M	3M	$\approx 1$	2.00	9	1.07	2.16	$> 10^4$
Citation graphs (CGs)	ciP	cit-Patents	3.7M	16.5M	1.01	1.26	2	1.01	1.26	1875
Web graphs (WGs)	wGL	web-Google	875k	5.1M	0.98	1.26	6	1.06	1.35	365
	wBS	web-BerkStan	685k	7.6M	0.93	1.31	5	1.07	1.40	755
	wSF	web-Stanford	281k	2.3M	0.98	1.54	5	1.07	1.58	1077