

Section 3

Practice: Spin Locks and Contention

HY486 - Panagiota Fatourou

1

Spin Locks and Contention

- In contrast to uniprocessor programming, in multiprocessor programming, it is crucial to understand the underlying machine architecture.
- What do you do if you cannot acquire the lock (i.e., if you cannot enter the critical section)?
 - Spinning (or busy waiting)
 - Sensible when you expect the lock delay to be short
 - Blocking
 - Sensible when you expect the lock delay to be long
- Many Operating Systems (OS) apply a combination of both techniques
- Here, we turn our attention to spin locks.

HY486 - Panagiota Fatourou

2

Welcome to Real World: Peterson's algorithm

Experiment

- In practice, once both threads have finished the execution of the 5000000 accesses to the CS, we may discover that the shared counter's final value may be slightly off from the expected value.

⇒ Even if this is a tiny error, why is there any error at all?

- Compilers re-order instructions to enhance performance!

- Program order is preserved for each individual variable but not always across multiple variables.

- Due to hardware, writes to multiprocessor memory do not necessarily take effect when they are issued!

- Writes to shared memory are buffered in a special write (or store) buffer, to be written to memory only when needed.

```
shared int count = 0;
```

Code for Process p_i

```
while (lcnt++ < 5,000,000) {
    i = ThreadId.get(); // either 0 or 1
    flag[i] = true
    turn = 1-i
    while (flag[1-i] and turn == 1-i)
        noop;

    count++; // critical section

    flag[i] = false

    remainder section;
}
```

HY486 - Panagiota Fatourou

3

Welcome to Real World

Memory Fences (or memory barriers)

- A *memory fence* forces outstanding operations to take effect!
- It is usually an expensive operation!
- Stronger primitives, like `Get&Set()` or `Compare&Swap()`, as well as reads and writes to *volatile* vars do not cause such errors (usually because they are implemented using memory fences).

```
ATOMIC boolean Compare&Swap(
    MemoryWord *pW, Value old, Value new) {
    Value tmp = *pW;
    if (*pW == old) {
        *pW = new;
        return TRUE;
    }
    return FALSE;
}
```

```
ATOMIC boolean
    Test&Set(MemoryByte *pB) {
    boolean tmp = *pB;
    *pB = 1;
    return tmp;
}
```

```
ATOMIC void Reset(MemoryByte *pB) {
    *pB = 0;
}
```

```
ATOMIC int Fetch&Inc(MemoryWord *pW) {
    int tmp = *pW;
    *pW = *pW + 1;
    return tmp;
}
```

```
ATOMIC Value
    Get&Set(MemoryWord *pW, Value nv) {
    Value tmp = *pW;
    *pW = nv;
    return tmp;
}
```

⇒ Given that memory fences cost as much as synchronization instructions, it may make sense to design ME algorithms directly from such synchronization primitives!

HY486 - Panagiota Fatourou

4

The TAS and TTAS Locks

THE TAS LOCK

```
void lock(MemoryByte *pB){
    while (Test&Set(pB)) noop;
    // wait until Test&Set(pB) returns 0
}
```

```
void unlock(MemoryByte *pB){
    reset(pB);
}
```

Theorem

The algorithm above is a correct ME algorithm.

- *Is it possible for some process to starve in this algorithm?*

✓ The TAS and TTAS Locks are equivalent in terms of correctness!

The TTAS Lock

```
void lock(MemoryByte *pB){
    while (TRUE){
        while (*pB == TRUE) noop;
        if (Test&Set(pB) == FALSE)
            return;
    }
}
```

```
void unlock(MemoryByte *pB){
    reset(pB);
}
```

Theorem

The algorithm above is a correct ME algorithm.

HY486 - Panagiota Fatourou

5

The TAS and TTAS Locks

- But are they equivalent in terms of performance?

- The TAS Lock performs very poorly!
- The TTAS Lock performs substantially better but still falls far short from the ideal!

The TAS Lock - Some remarks

- Each Test&Set() call causes:
 - a broadcast on the bus, and
 - all other processors to discard their own cached copies of the lock!
- Additionally, when the thread holding the lock tries to release it, it may be delayed due to bus traffic caused by the spinners!

The TTAS Lock - Some important points

- What happens the first time a thread B reads the lock?
- What happens each time B rereads the lock (finding it occupied)?
- Is a thread that releases the lock delayed by other threads?
- What happens when the lock is released by its holder thread A?

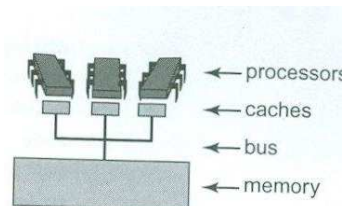
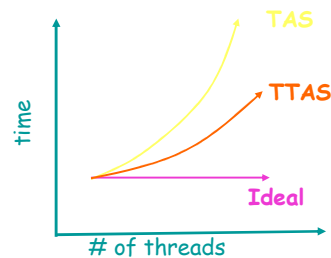


Figure B.4: M. Herlihy and N. Shavit, The Art of Multiprocessor Programming, Morgan Kaufman, 2008

HY486 - Panagiota Fatourou

6

Exponential Backoff

The idea

- If some other thread acquires the lock between the read step and the Test&Set step in the TTAS algorithm, there is probably high contention for the lock. Therefore, we'll back off for some time, giving to the competing threads a chance to finish.

How long should the thread backoff before it retries?

- The larger the number of unsuccessful tries, the higher the likely contention and the longer the thread should backoff.

Strategy

- The thread backoffs for a random duration.
- Each time the thread tries and fails to get the lock, it doubles the expected back-off time, up to a fixed maximum.

```
#define MIN_DELAY ...
#define MAX_DELAY ...

void BackOff(void) {
    static int limit = MIN_DELAY;
    int delay = rand_next(limit);
    // if 0 < limit < 32, that many low-order bits of the
    // returned value will be independently chosen bit
    //values
    limit = min(MAX_DELAY, 2*limit);
    usleep(delay);
}
```

The Exponential BackOff TTAS Lock

```
void lock(MemoryByte *pB) {
    while (TRUE) {
        while (*pB == TRUE) noop;
        if (Test&Set(pB) == FALSE)
            return;
        else BackOff();
    }
}

void unlock(MemoryByte *pB) {
    reset(pB);
}
```

HY486 - Panagiota Fatourou

7

Exponential Backoff

Advantages

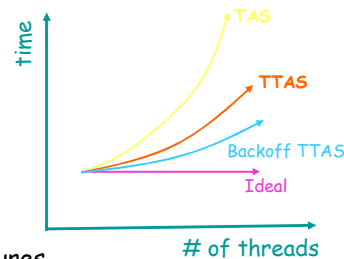
- Easy to implement
- Has better performance than TTAS Lock

Drawbacks

- Must choose parameters carefully
- Not portable across different architectures

Problems

- Cache-Coherence Traffic
 - All threads spin on the same shared location causing cache-coherence traffic on every successful lock access.
- Critical Section Underutilization
 - Threads might back off for too long causing the critical section to be underutilized.



HY486 - Panagiota Fatourou

8

Queue Locks

Idea

- A queue is formed.
 - In a queue, every thread can learn if his turn has arrived by checking whether his predecessor has finished.
 - A queue has better utilization of the critical section because there is no need a thread to guess when is its turn.

HY486 - Panagiota Fatourou

9

Anderson's Algorithm: The Array-Based Lock

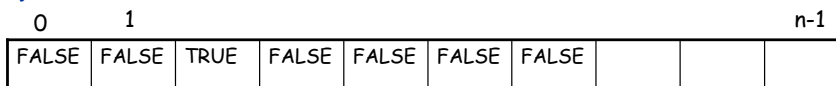
```
#define n <number-of-processes>

shared integer Tail = 0;
shared BOOLEAN flag[n] =
    {TRUE,FALSE,FALSE,... ,FALSE};

/* Code for process $p_i$ */
int slot = -1; /* global variable for process $p_i$; */

void lock(void) {
    slot = Fetch&Inc(&Tail);
    while (!flag[slot % n]) noop;
}

void unlock(void) {
    flag[slot % n] = FALSE;
    flag[(slot + 1) % n] = TRUE;
}
```



The Flag Array

HY486 - Panagiota Fatourou

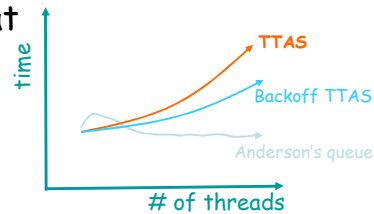
Tail

10

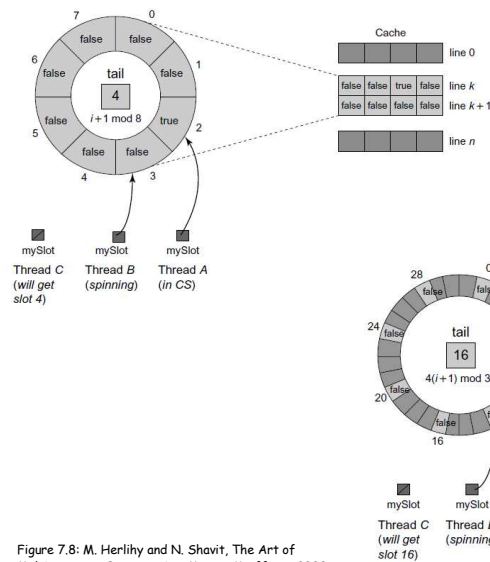
Anderson's Algorithm: The Array-Based Lock

Advantages

- ❑ At any given time, each thread spins on its locally cached copy of a single array location
 - Shorter handover than backoff
 - Curve is practically flat
 - Better Scalability
- ❑ FIFO Fairness



Anderson's Algorithm: The Array-Based Lock



Drawbacks

- ❑ Padding is required, so that distinct elements are mapped to distinct cache lines in order for false sharing to be avoided. ☹

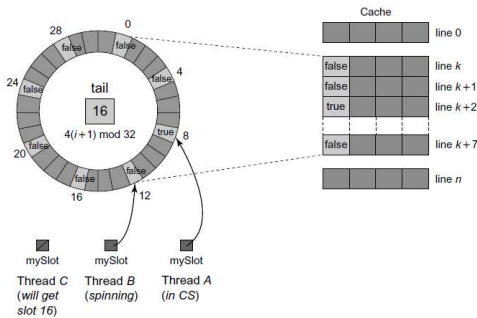
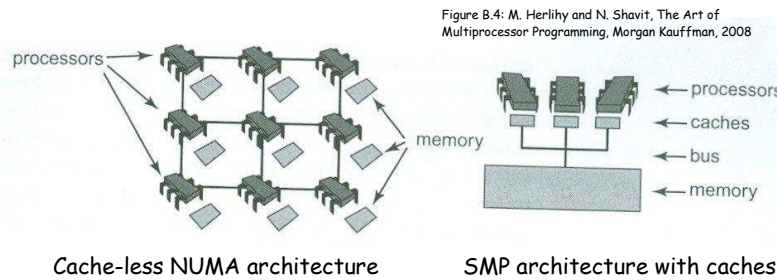


Figure 7.8: M. Herlihy and N. Shavit, The Art of Multiprocessor Programming, Morgan Kaufman, 2008

Anderson's Algorithm: The Array-Based Lock

Drawbacks

- Not very space-efficient. ☹
 - It allocates an array of size $O(n)$ per lock. Synchronizing L distinct objects requires $O(Ln)$ space, even if a thread accesses only one lock at a time ☹
- On cache-less NUMA architectures, the algorithm does not have a good performance since spinning during `lock()` might be performed on a remote variable ☹



HY486 - Panagiota Fatourou

13

The CLH Lock

A more space-efficient Algorithm

```
typedef struct node {
    BOOLEAN locked;
} NODE;

shared NODE *Tail;
/* initially, points to a NODE n with n.locked == FALSE */

/* Section of global variables for process $p_i$ */
NODE *MyNode, *MyPred = NULL;
/* Variable MyNode initially points to a struct NODE */

void lock(void) {
    MyNode->locked = TRUE;
    MyPred = Get&Set(&Tail, MyNode);
    while (MyPred->locked == TRUE) noop;
}

void unlock(void) {
    MyNode->locked = FALSE;
    MyNode = MyPred;
    /* recycling of nodes */
}
```

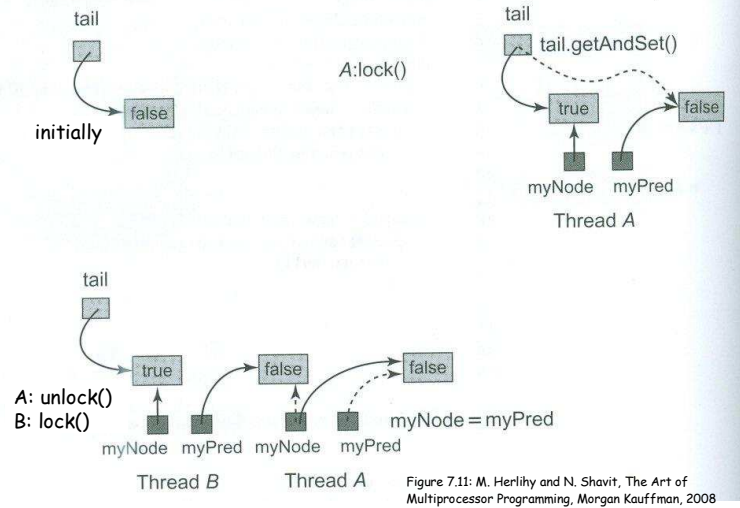
Brief Description

- ❑ An implicit list of NODES is created. Each thread allocates just one NODE.
- ❑ NODES are not connected to each other. Rather, each thread refers to its predecessor thread through a (non-shared) global variable, called `MyPred`.
- ❑ NODES are re-cycled by having each thread using the NODE pointed to by its `MyPred` Variable as its current NODE the next time it requires the lock.
- ❑ An initial NODE is placed in the queue implementing each lock.

HY486 - Panagiota Fatourou

14

The CLH Lock



HY486 - Panagiota Fatourou

15

The CLH Lock

Advantages

- ❑ Requires much less space than the Anderson's algorithm
 - Synchronizing L distinct objects requires $O(L+n)$ space (much better than the Array-based Lock).
- ❑ Does not require knowledge of the number of processes that might access the lock.
- ❑ Has all the performance advantages of the Anderson's algorithm.

Drawbacks

- ❑ Like Anderson's algorithm, on cache-less NUMA architectures, the CLH algorithm does not have good performance since it causes a lot of Remote Memory References (RMRs); i.e., spinning is performed on a remote variable.

HY486 - Panagiota Fatourou

16

The MCS Lock

```

typedef struct node {
    BOOLEAN locked;
    struct node *next;
} NODE;

/* shared variables section */
shared NODE *Tail = NULL;

/* Section of global variables for process $p_!$ */
NODE *MyNode, *MyPred = NULL;
/* Variable MyNode initially points to a NODE */

void lock {
    MyPred = Get&Set(&Tail, MyNode);
    if (MyPred != NULL) {
        MyNode->locked = TRUE;
        MyPred->next = MyNode;
        while (MyNode->locked) noop;
    }
}

void unlock {
    if (MyNode->next == NULL) {
        if (Compare&Swap(&Tail, MyNode, NULL) == TRUE) return;
        while (MyNode->next == NULL) noop;
    }
    MyNode->next->locked = FALSE;
    MyNode->next = NULL;
}
    
```

Major Ideas

- The created list of NODEs is now explicit.
- Each thread does not spin on the NODE pointed to by MyPred but on that pointed to by MyNode (which is a variable in the thread's local memory)

The MCS Lock

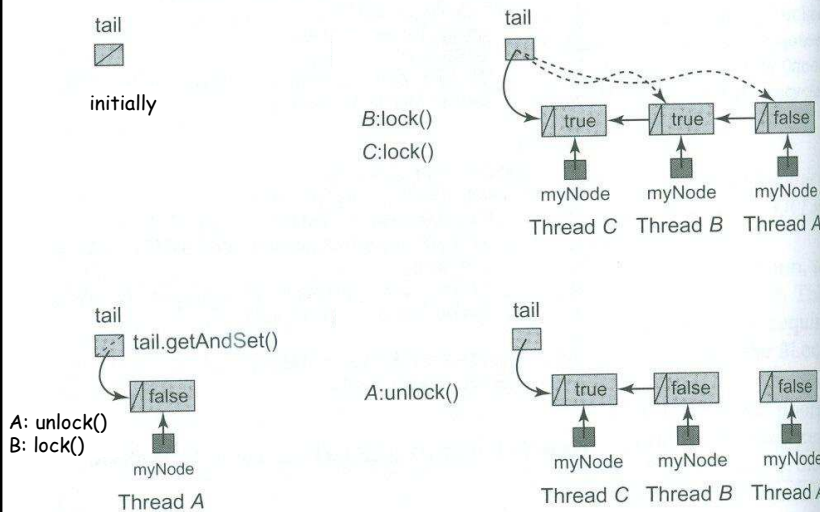


Figure 7.14: M. Herlihy and N. Shavit, The Art of Multiprocessor Programming, Morgan Kaufman, 2008

The MCS Lock

Advantages

- ❑ Each unlock causes just one invalidation (as with CLH)
- ❑ Lock/Unlock causes just $O(1)$ RMRs on a NUMA architecture
- ❑ Recycling can be applied (as in CLH) to obtain space overhead $O(n)$

Drawbacks

- ❑ Releasing a lock requires spinning.
- ❑ The algorithm executes more reads and writes to execute lock/unlock, and it requires `Compare&Swap()` for unlock.

HY486 - Panagiota Fatourou

19

J. Mellor-Crummey and M. Scott, "Algorithms for Scalable Synchronization on shared-memory multiprocessors", *ACM Transactions on Computer Systems*, 9(1):21-65, 1991.

Drawbacks of Queue Locks

- They perform well only in cases of no oversubscribing
- If oversubscribing occurs, their performance deteriorates significantly
- Binding may also affect their performance

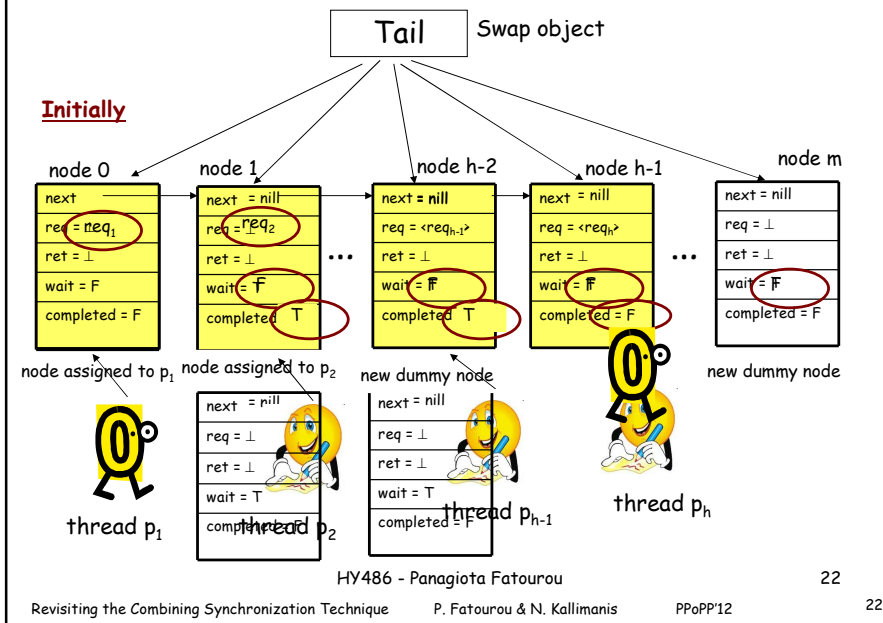
HY486 - Panagiota Fatourou

20

The Basics of the Combining Technique

- A thread attempts to become a **combiner** and serve in addition to its own request, active requests by other threads
- After announcing their requests, other threads may:
 - either perform the same actions (although not always "successfully")
 - synchronization by employing CAS or other similar primitives
 - or perform local spinning until the combiner performs their requests
 - synchronization by employing a coarse-grain lock

Blocking Combining: CC-Synch



Properties of CC-Synch

Fairness

- CC-Synch provides a strong notion of fairness: it serves requests in the order they enter the announced list.

Progress

- In CC-Synch, no thread ever starves.

Performance

- The lock is implemented using a highly-efficient queue-like lock.
- CC-Synch is the first that provides bounds on the number of RMRs it executes:
 - A combiner thread performs $O(h+t)$ RMRs, where h is an upper bound on the number of requests that a combiner may serve, and t is the size of the shared data that should be accessed in order to execute the h requests.
 - Other threads perform just a constant number of RMRs.
 - The amortized number of RMRs performed is $O(d)$, where d is the average number of RMRs required to serve a single request.

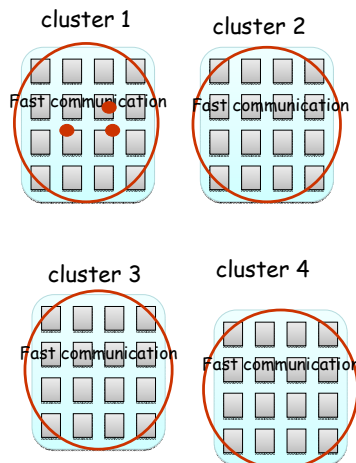
Required Primitives

- Only a swap object is used (no CAS) and r/w registers.

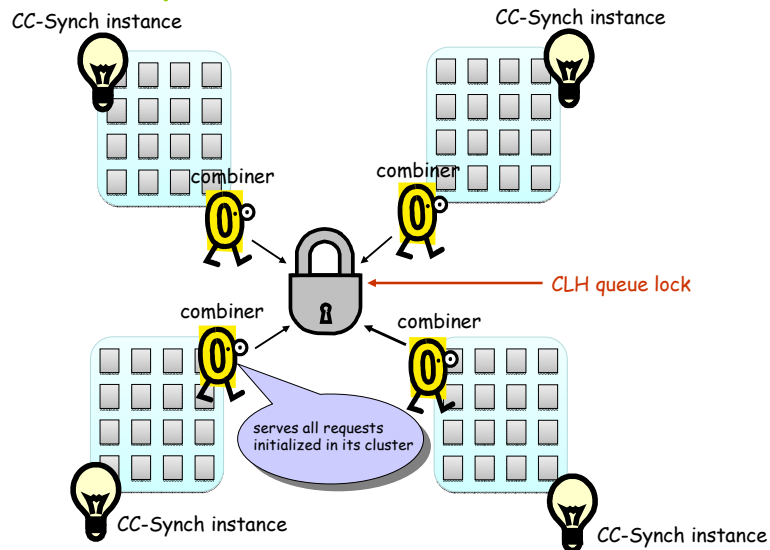
Hierarchically Structured systems The H-Synch Algorithm

- **H-Synch** is an hierarchical version of CC-Synch

⇒ Exploits hierarchical communication nature to achieve better performance



H-Synch - More Details



HY486 - Panagiota Fatourou

25

Revisiting the Combining Synchronization Technique

P. Fatourou & N. Kallimanis

PPoPP'12

25

Experimental Analysis

Machines on which the experiments were performed

- 32-core machine consisting of 4 AMD Opteron 6134 processors (Magny Cours)
- 128-way Sun consisting of 2 UltraSPARC-T2 processors (Niagara 2)

Synchronization approaches that were compared

- CC-synch and DSM-Synch
- P-Sim (Fatourou and Kallimanis, SPAA 2011)
- Flat-combining
- CLH-spin locks
- OyamaAlg
- Simple lock-free implementation
- H-Synch and Hierarchical NUMA lock (Marathe et al., SPAA 2011) in Niagara 2 machines

Experiment

A simple Fetch&Multiply object is simulated:

We measure the average throughput (number of Fetch&Multiply executed per second) that each approach exhibits when it executes 10^7 Fetch&Multiply requests for different values of n .

A random number of dummy loop iterations are executed between the execution of two consecutive Fetch&Multiply by the same thread.

HY486 - Panagiota Fatourou

26

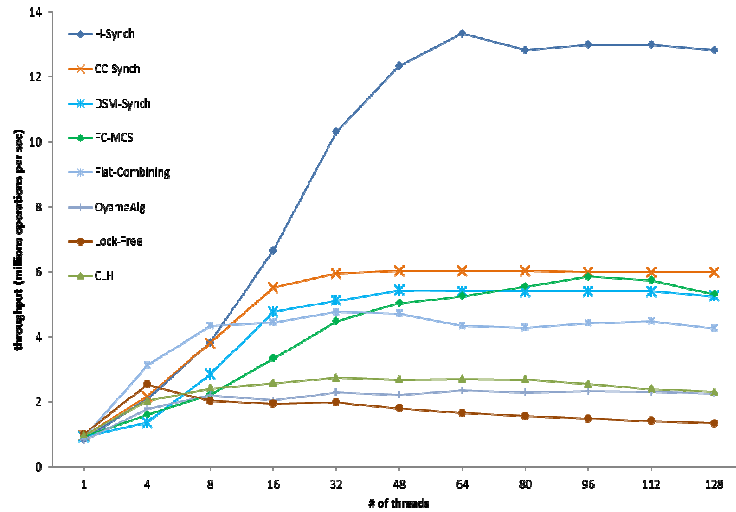
Revisiting the Combining Synchronization Technique

P. Fatourou & N. Kallimanis

PPoPP - Feb 2012

26

Experimental Analysis Niagara 2



HY486 - Panagiota Fatourou

27

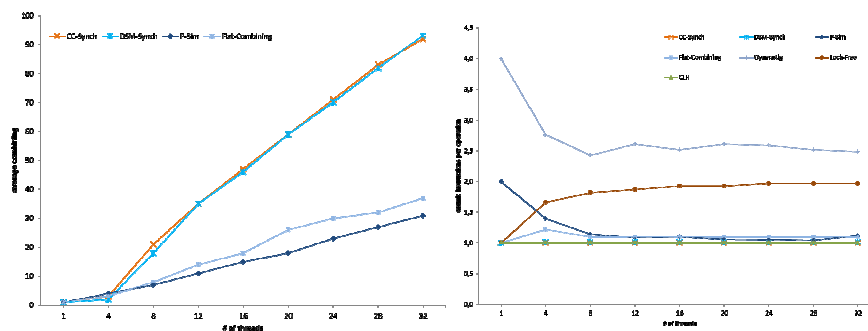
Revisiting the Combining Synchronization Technique

P. Fatourou & N. Kallimanis

PPoPP'12

27

Why is CC-Synch Efficient?



Factors that significantly impact performance

- Combining degree
- Number of primitives performed

HY486 - Panagiota Fatourou

28

Concurrent Computing

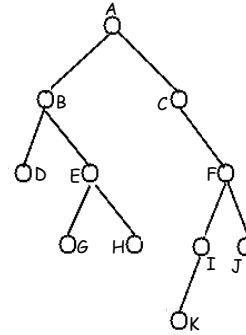
1st Seminar

Panagiota Fatourou

28

PART III: Limitations of the Combining Technique

- What is the cost to apply k concurrent search operations on a tree?
- A "sophisticated" concurrent implementation could allow multiple searches proceeding in parallel and being executed in $O(h)$ time in total, where h is the height of the tree.



Bibliography

These slides are based on material that appears in the following books and papers:

- M. Herlihy and N. Shavit, *The Art of Multiprocessor Programming*, Morgan Kaufman, 2008 (Chapter 7)
- Panagiota Fatourou, [Nikolaos D. Kallimanis](#): Revisiting the combining synchronization technique. [PPOPP 2012](#): 257-266
- Panagiota Fatourou, [Nikolaos D. Kallimanis](#): A highly-efficient wait-free universal construction. [SPAA 2011](#): 325-334