

# Synchronization: Implementing Barriers Promises + Futures

Chris Rossbach

CS378H

# Today

- Questions?
- Administrivia
  - Lab 2 due sooner than you'd like
- Material for the day
  - Barrier implementation
  - Promises & Futures
- Acknowledgements
  - Thanks to Gadi Taubenfeld: I borrowed from some of his slides on barriers



# Faux Quiz (answer any N, 5 min)

- How are promises and futures related? Since there is disagreement on the nomenclature, don't worry about which is which—just describe what the different objects are and how they function.

# Barriers

# Barriers




# Prefix Sum

# Prefix Sum

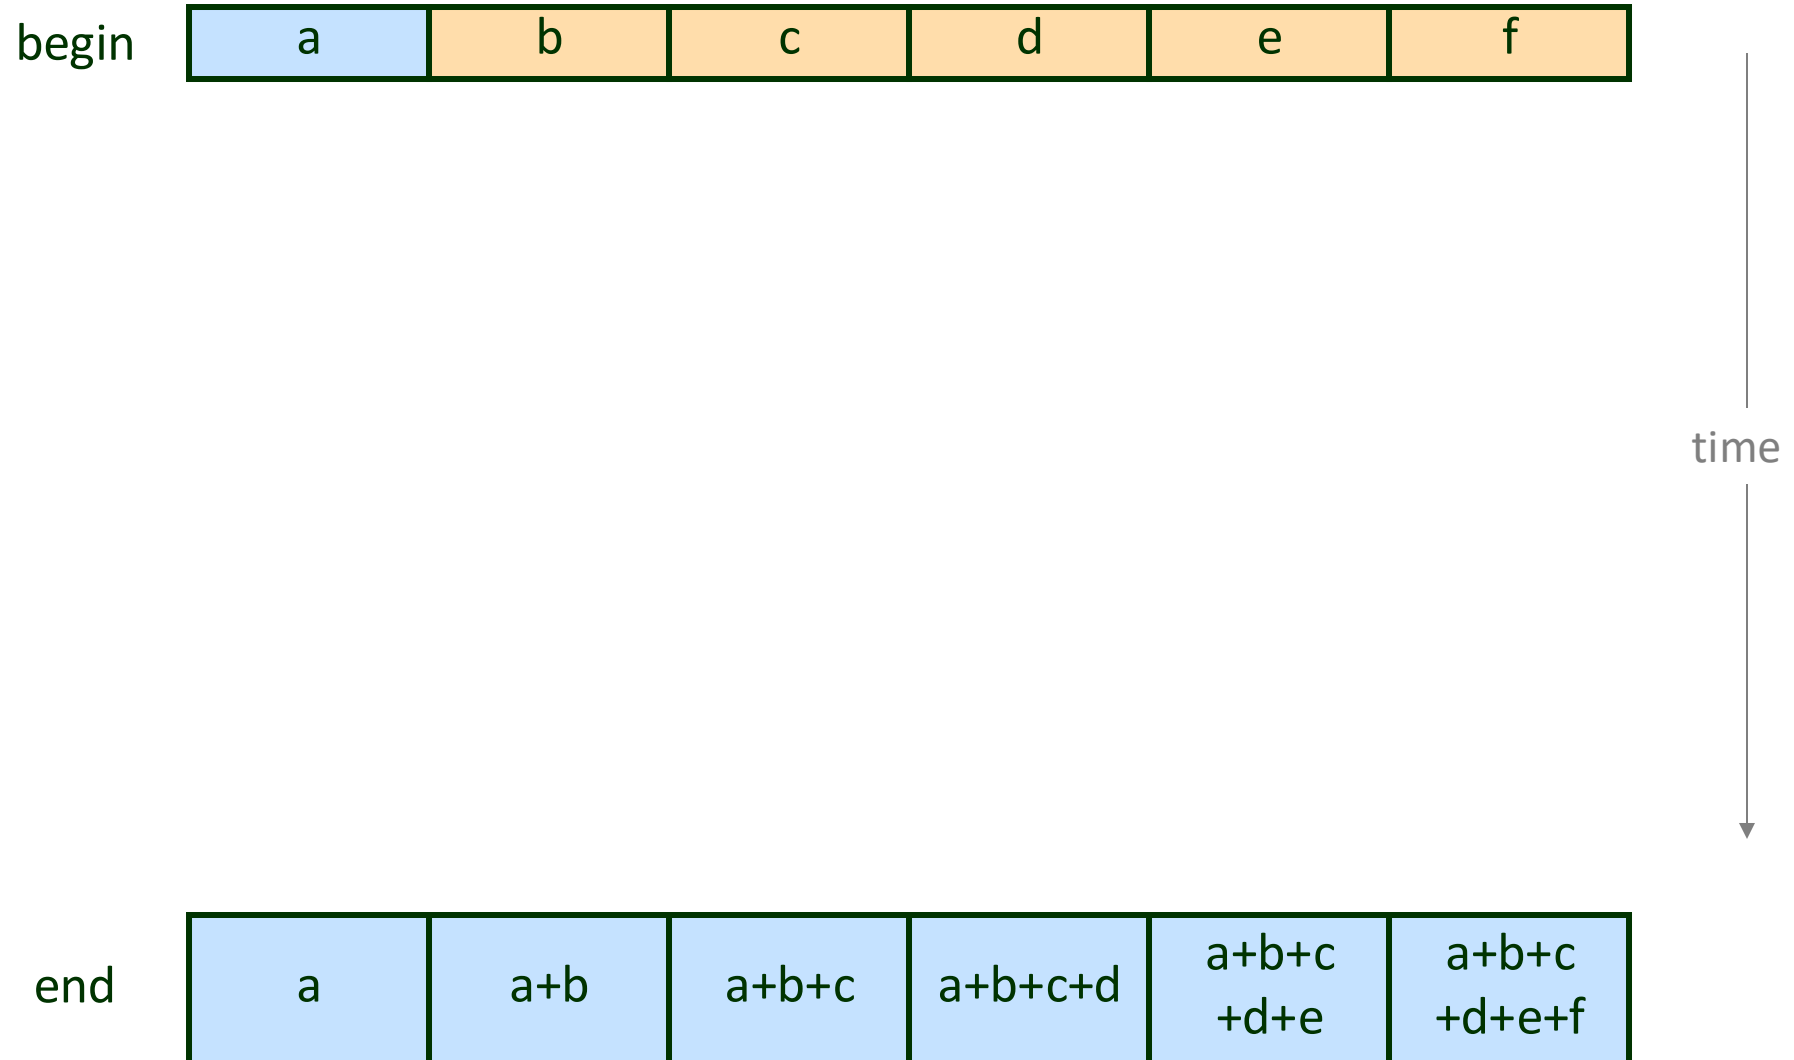


# Prefix Sum

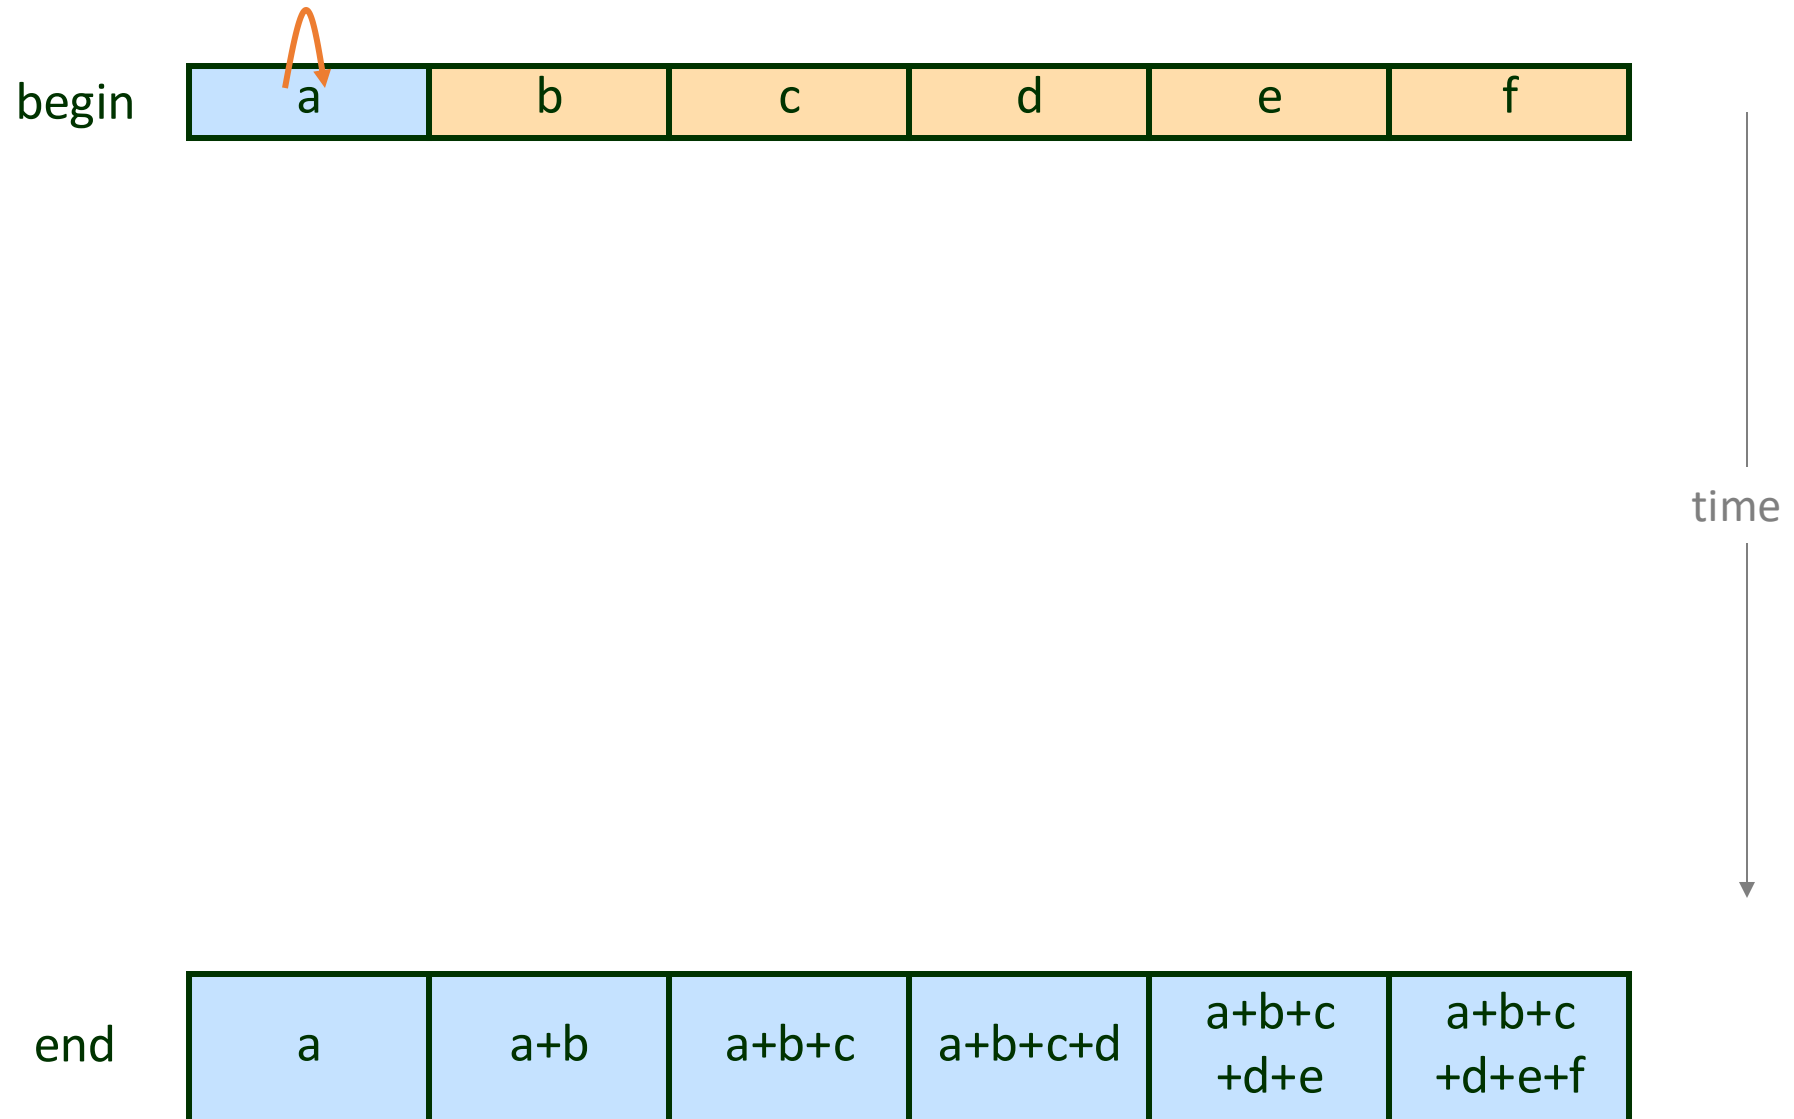





# Prefix Sum




# Prefix Sum




# Prefix Sum




# Prefix Sum




# Prefix Sum




# Prefix Sum



# Prefix Sum




# Prefix Sum






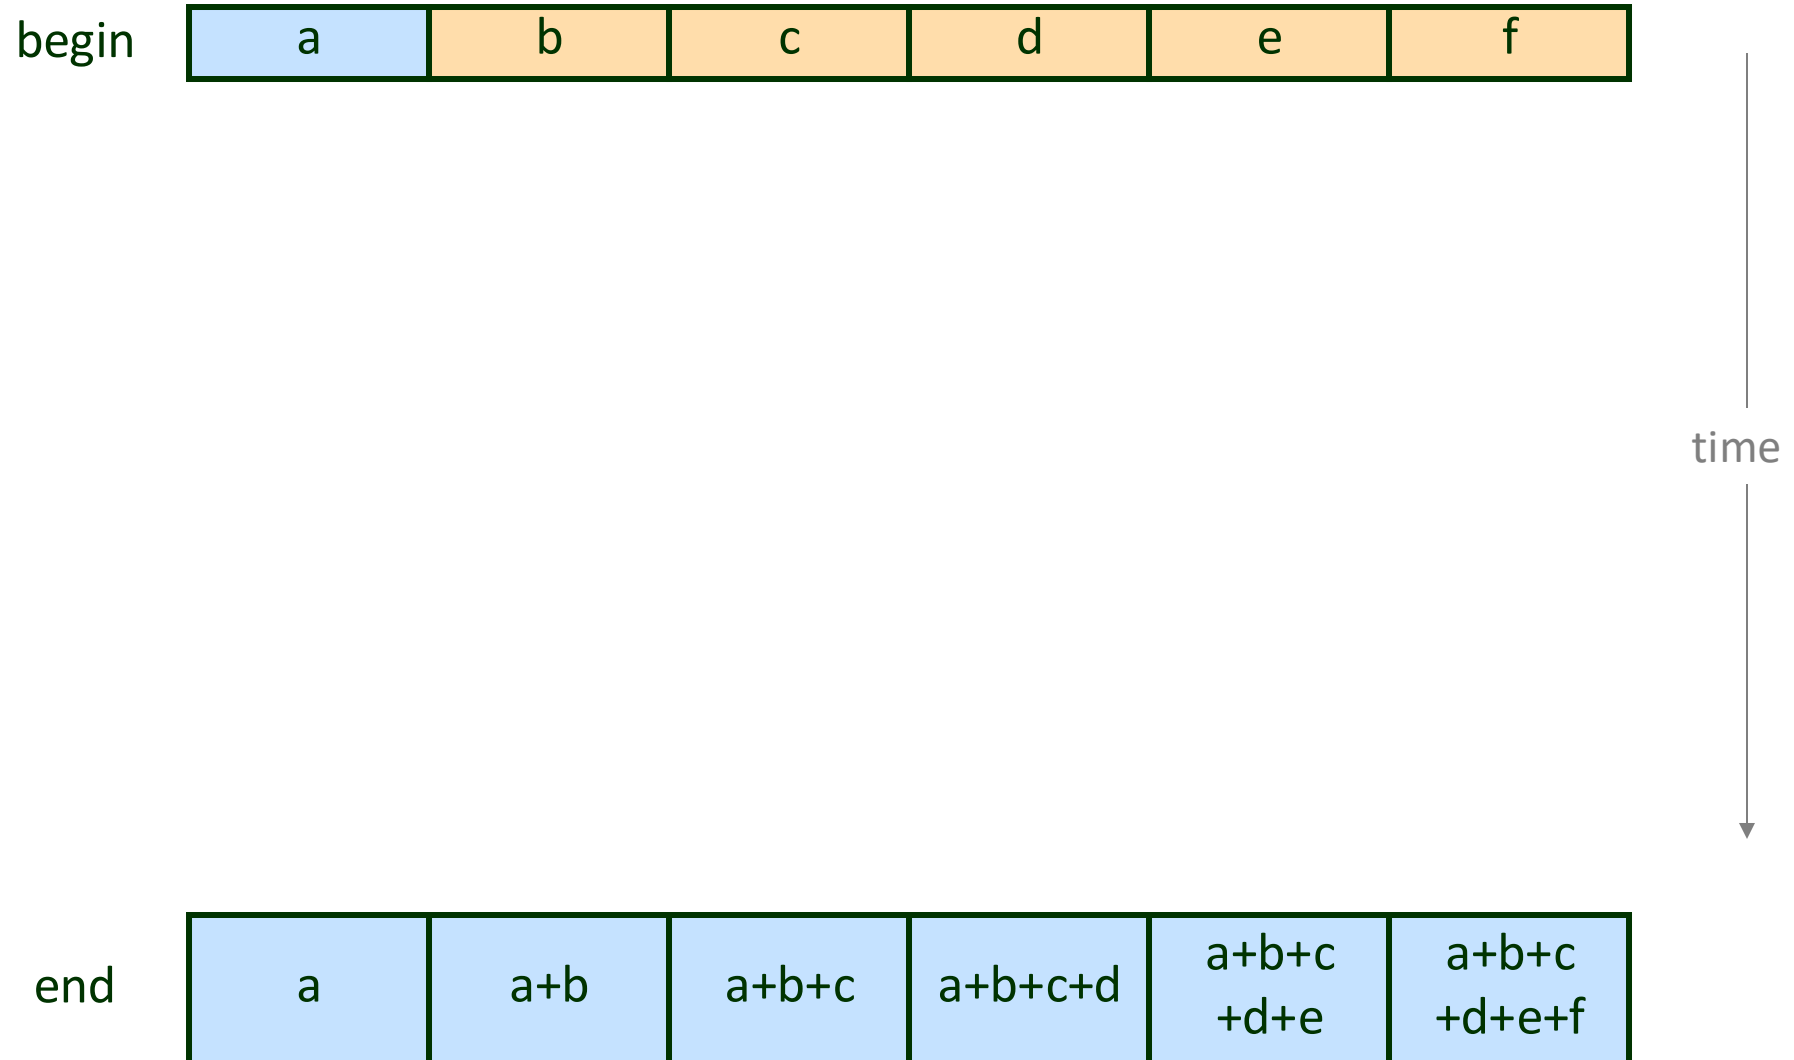
# Prefix Sum




# Prefix Sum




# Parallel Prefix Sum




# Parallel Prefix Sum



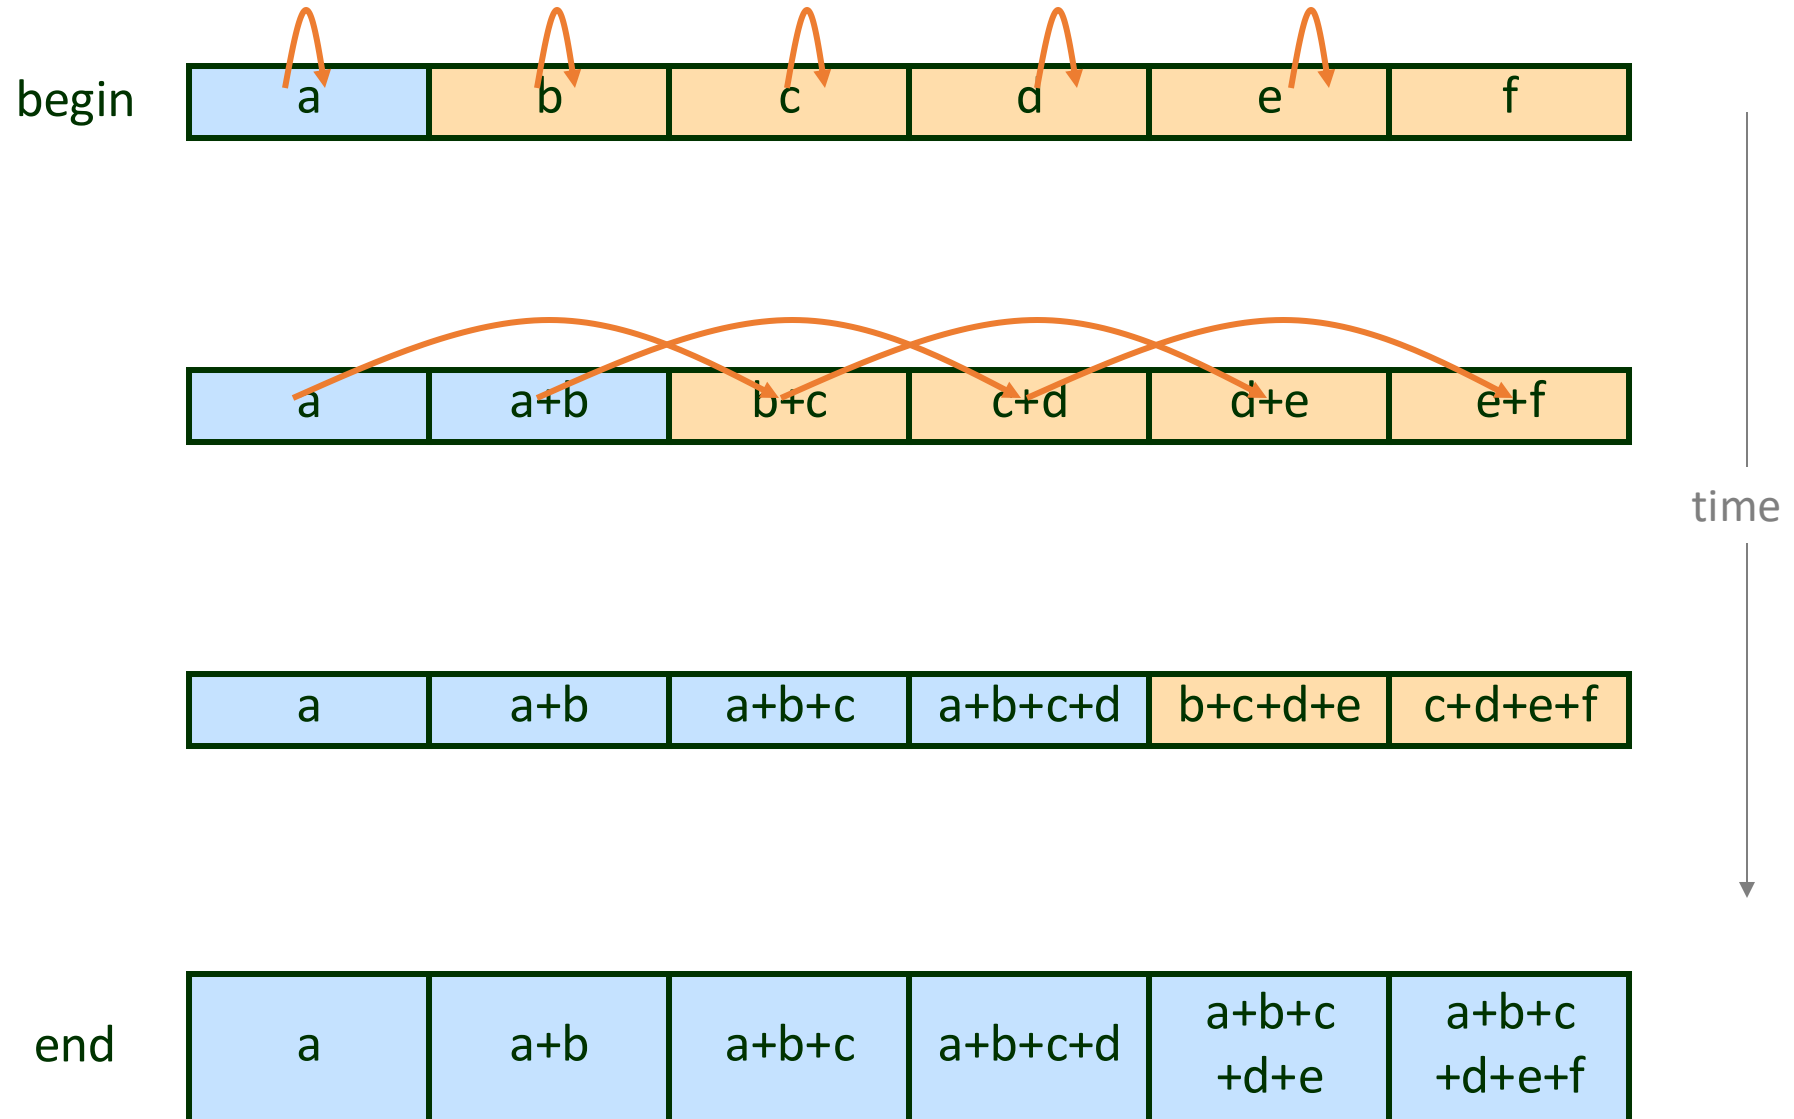
# Parallel Prefix Sum




# Parallel Prefix Sum



# Parallel Prefix Sum



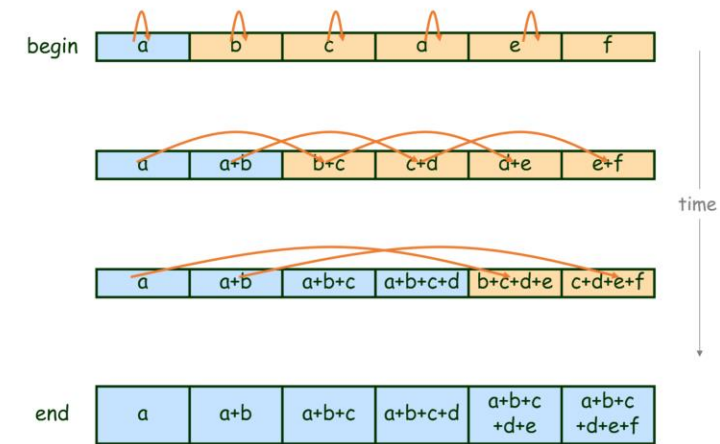
# Parallel Prefix Sum






# Pthreads Parallel Prefix Sum

```
int g_values[N] = { a, b, c, d, e, f };  
  
void prefix_sum_thread(void * param) {  
  
    int i;  
    int id = *((int*)param);  
    int stride = 0;  
  
    for(stride=1; stride<=N/2; stride<<1) {  
        g_values[id+stride] += g_values[id];  
    }  
  
}
```



# Pthreads Parallel Prefix Sum

```
int g_values[N] = { a, b, c, d, e, f };  
  
void prefix_sum_thread(void * param) {  
  
    int i;  
    int id = *((int*)param);  
    int stride = 0;  
  
    for(stride=1; stride<=N/2; stride<<1) {  
        g_values[id+stride] += g_values[id];  
    }  
  
}
```



Will this  
work?

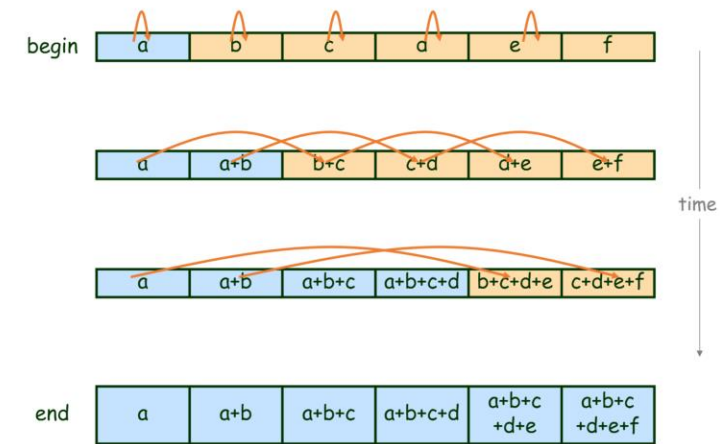
# Pthreads Parallel Prefix Sum

```
pthread_mutex_t g_locks[N] = { MUTEX_INITIALIZER, ...};
int g_values[N] = { a, b, c, d, e, f };

void prefix_sum_thread(void * param) {

    int i;
    int id = *((int*)param);
    int stride = 0;

    for(stride=1; stride<=N/2; stride<<1) {
        pthread_mutex_lock(&g_locks[id]);
        pthread_mutex_lock(&g_locks[id+stride]);
        g_values[id+stride] += g_values[id];
        pthread_mutex_unlock(&g_locks[id]);
        pthread_mutex_unlock(&g_locks[id+stride]);
    }
}
```



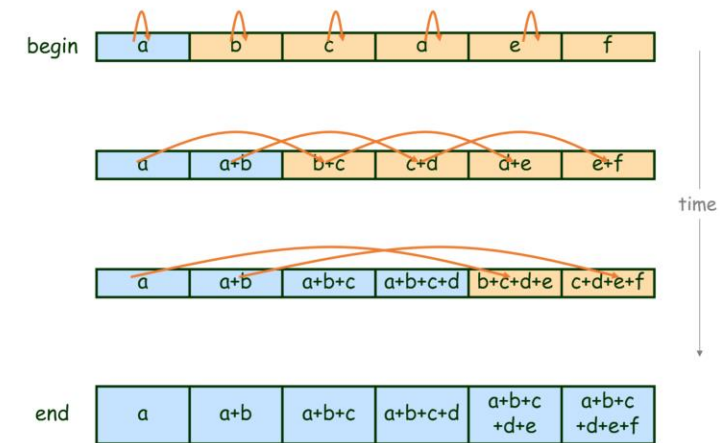
# Pthreads Parallel Prefix Sum

```
pthread_mutex_t g_locks[N] = { MUTEX_INITIALIZER, ...};
int g_values[N] = { a, b, c, d, e, f };

void prefix_sum_thread(void * param) {


    int i;
    int id = *((int*)param);
    int stride = 0;

    for(stride=1; stride<=N/2; stride<<1) {
        pthread_mutex_lock(&g_locks[id]);
        pthread_mutex_lock(&g_locks[id+stride]);
        g_values[id+stride] += g_values[id];
        pthread_mutex_unlock(&g_locks[id]);
        pthread_mutex_unlock(&g_locks[id+stride]);
    }
}
```



fixed?

# Parallel Prefix Sum



# Barrier Basics




# Barrier Basics

- *Coordination mechanism*
- *participants wait until all reach same point.*
- *Once all reach it, all can pass.*

# Barrier Basics

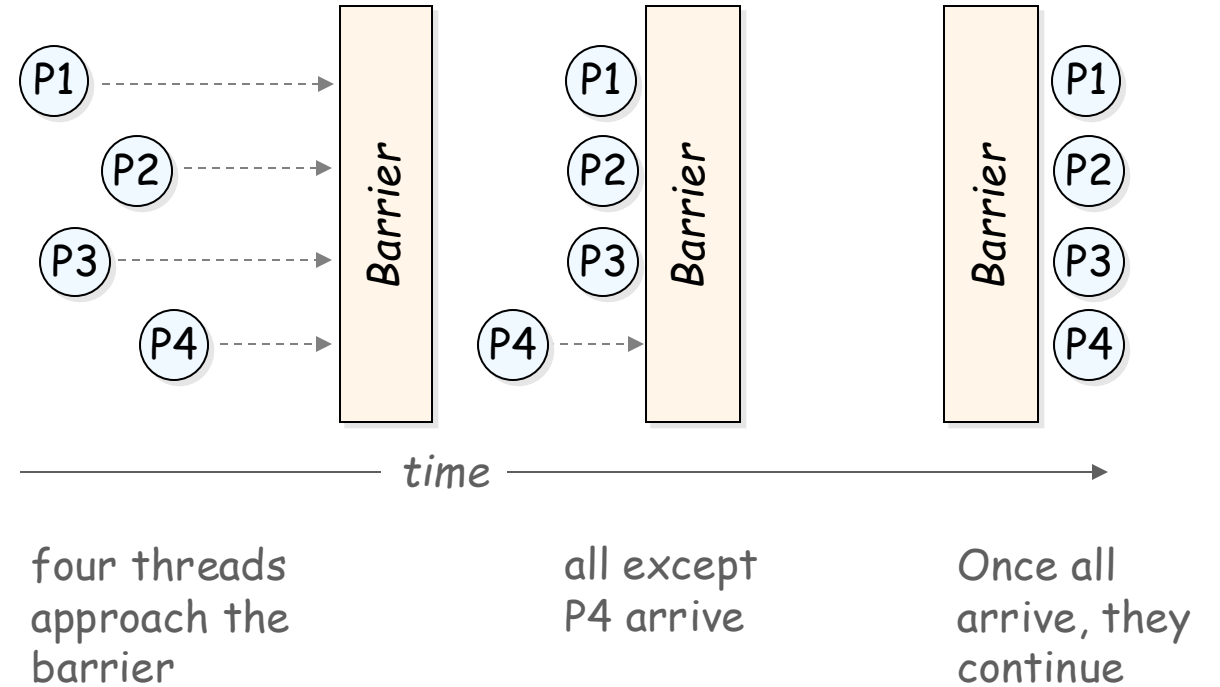
- Coordination mechanism
- participants wait until all reach same point.
- Once all reach it, all can pass.






# Barrier Basics

- Coordination mechanism
- participants wait until all reach same point.
- Once all reach it, all can pass.
- **Workhorse of BSP programming models**



# Barrier Basics


- Coordination mechanism
- participants wait until all reach same point
- Once all reach it, all can pass.
- Workhorse of BSP programming models



# Barrier Basics

- Coordination mechanism
- participants wait until all reach same point
- Once all reach it, all can pass.
- **Workhorse of BSP programming models**


*Fundamental primitive in many parallel models*



# Barrier Basics

Can you make a lock with a barrier?

- **Coordination mechanism**
- participants wait until all reach same point
- Once all reach it, all can pass.
- **Workhorse of BSP programming models**



Fundamental primitive in many parallel models

# Barriers: Goals

Desirable barrier properties:

# Barriers: Goals

Desirable barrier properties:

- Low shared memory space complexity

# Barriers: Goals

Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process



# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization
- Symmetric: same amount of work for all

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization
- Symmetric: same amount of work for all
- Algorithm simplicity

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization
- Symmetric: same amount of work for all
- Algorithm simplicity
- Simple basic primitive

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization
- Symmetric: same amount of work for all
- Algorithm simplicity
- Simple basic primitive
- Minimal propagation time

# Barriers: Goals

## Desirable barrier properties:

- Low shared memory space complexity
- Low contention on shared objects
- Low shared memory references per process
- No need for shared memory initialization
- Symmetric: same amount of work for all
- Algorithm simplicity
- Simple basic primitive
- Minimal propagation time
- Reusability of the barrier (must!)

# Barrier Building Blocks

- Conditions
- Semaphores
- Atomic Bit
- Atomic Register
- Fetch-and-increment register
- Test and set bits
- Read-Modify-Write register

# Barrier with Semaphores





# Barrier using Semaphores

Algorithm for N threads

# Barrier using Semaphores

Algorithm for N threads



# Barrier using Semaphores

Algorithm for N threads

```
shared sem_t arrival = 1;      // sem_init(&arrival, NULL, 1)
sem_t departure = 0;         // sem_init(&departure, NULL, 0)
atomic int counter = 0;     // (gcc intrinsics are verbose)
```





# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1;      // sem_init(&arrival, NULL, 1)
sem_t departure = 0;         // sem_init(&departure, NULL, 0)
atomic int counter = 0;     // (gcc intrinsics are verbose)
```

```
type __sync_fetch_and_add (type *ptr, type value, ...)
type __sync_fetch_and_sub (type *ptr, type value, ...)
type __sync_fetch_and_or (type *ptr, type value, ...)
type __sync_fetch_and_and (type *ptr, type value, ...)
type __sync_fetch_and_xor (type *ptr, type value, ...)
type __sync_fetch_and_nand (type *ptr, type value, ...)
```

# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1;      // sem_init(&arrival, NULL, 1)
sem_t departure = 0;         // sem_init(&departure, NULL, 0)
atomic int counter = 0;     // (gcc intrinsics are verbose)
```





# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1;    // sem_init(&arrival, NULL, 1)
sem_t departure = 0;        // sem_init(&departure, NULL, 0)
atomic int counter = 0;     // (gcc intrinsics are verbose)
```

```
1 sem_wait(arrival);
2 if(++counter < N)
3   sem_post(arrival);
4 else
5   sem_post(departure);
6 sem_wait(departure);
7 if(--counter > 0)
8   sem_post(departure)
9 else
10  sem_post(arrival)
```



# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1;    // sem_init(&arrival, NULL, 1)
      sem_t departure = 0;   // sem_init(&departure, NULL, 0)
atomic int counter = 0;     // (gcc intrinsics are verbose)
```

Phase I

```
1 sem_wait(arrival);
2 if(++counter < N)
3   sem_post(arrival);
4 else
5   sem_post(departure);
```

Phase II

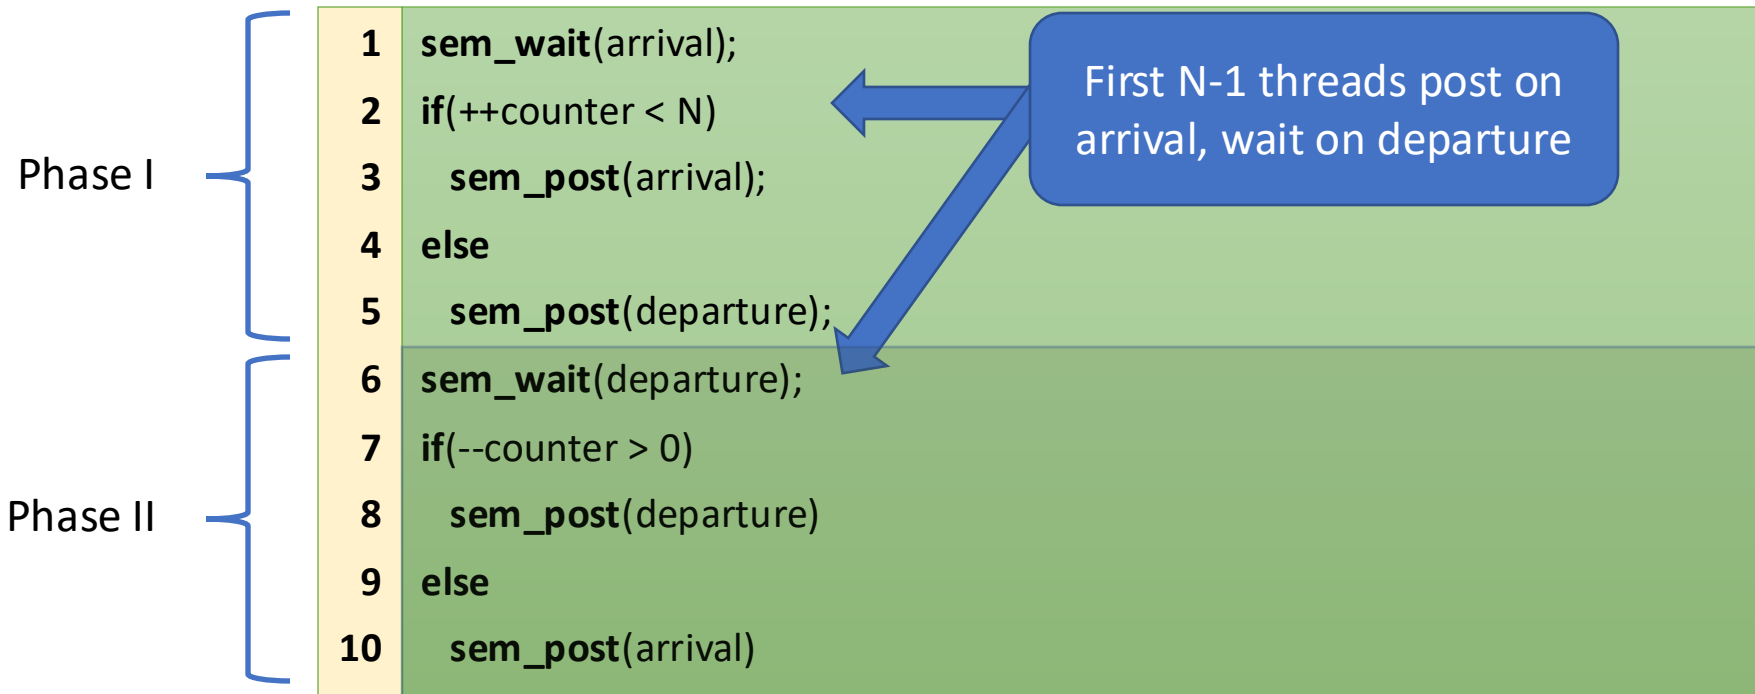
```
6 sem_wait(departure);
7 if(--counter > 0)
8   sem_post(departure)
9 else
10  sem_post(arrival)
```



# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1; // sem_init(&arrival, NULL, 1)
shared sem_t departure = 0; // sem_init(&departure, NULL, 0)
atomic int counter = 0; // (gcc intrinsics are verbose)
```



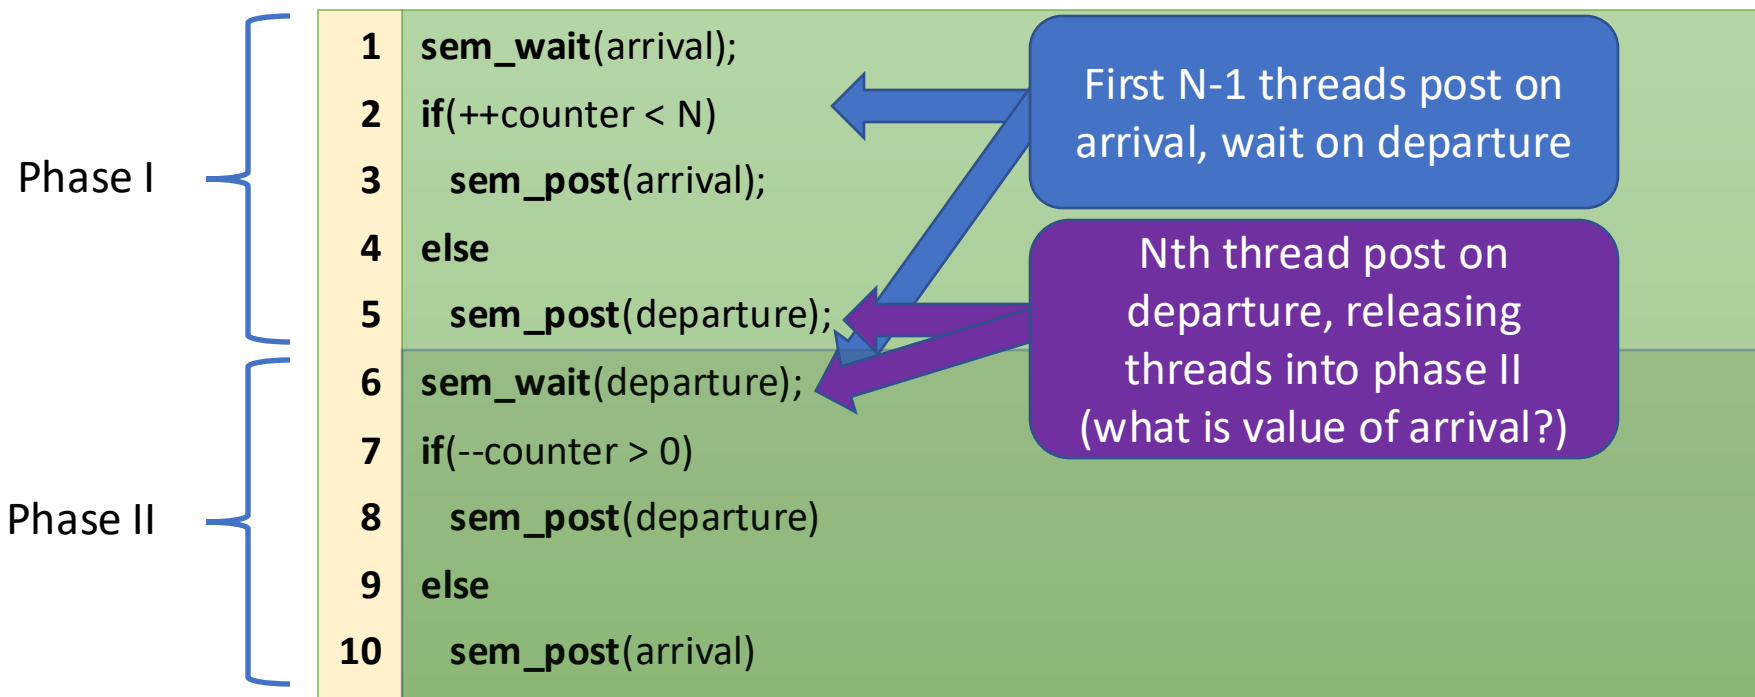




# Barrier using Semaphores

## Algorithm for N threads

```
shared sem_t arrival = 1; // sem_init(&arrival, NULL, 1)
shared sem_t departure = 0; // sem_init(&departure, NULL, 0)
atomic int counter = 0; // (gcc intrinsics are verbose)
```

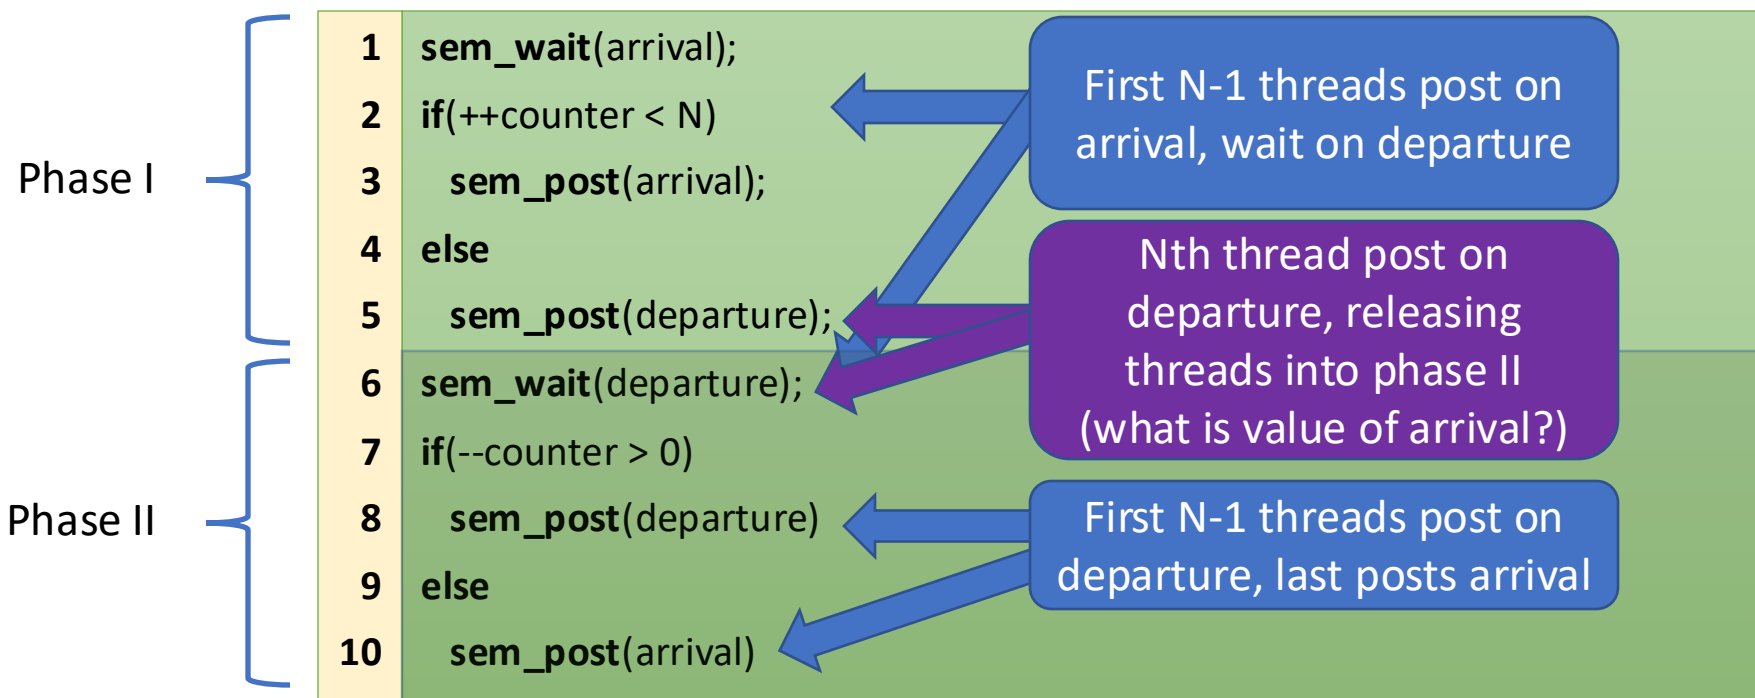




# Barrier using Semaphores

## Algorithm for N threads

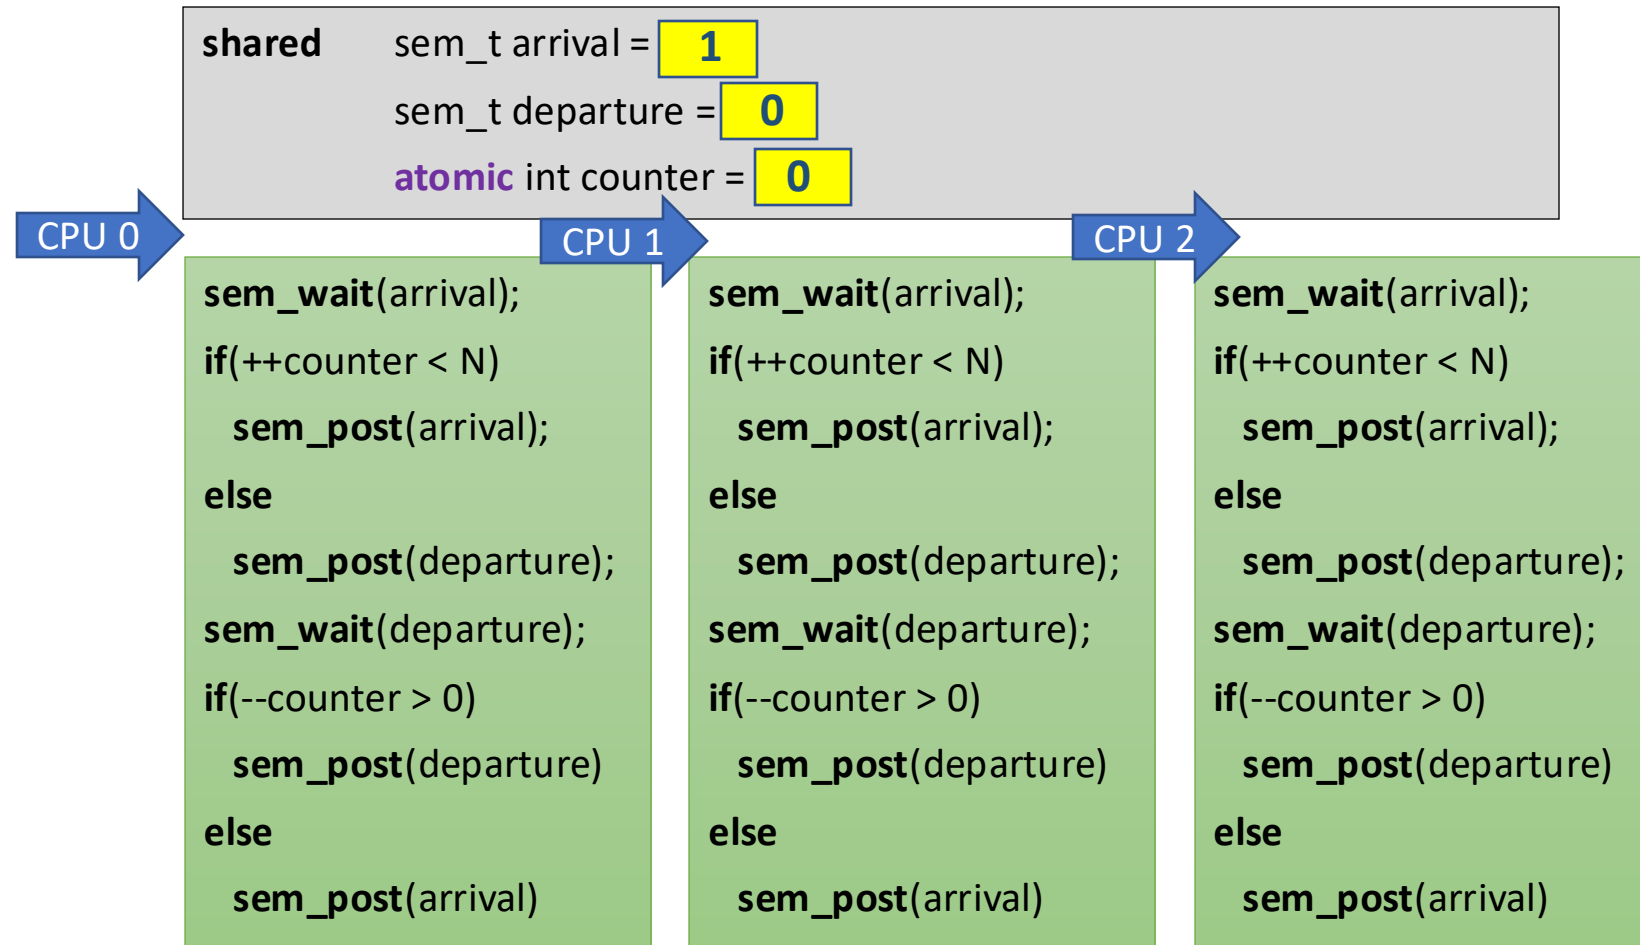
```
shared sem_t arrival = 1; // sem_init(&arrival, NULL, 1)
shared sem_t departure = 0; // sem_init(&departure, NULL, 0)
atomic int counter = 0; // (gcc intrinsics are verbose)
```





# Semaphore Barrier Action Zone

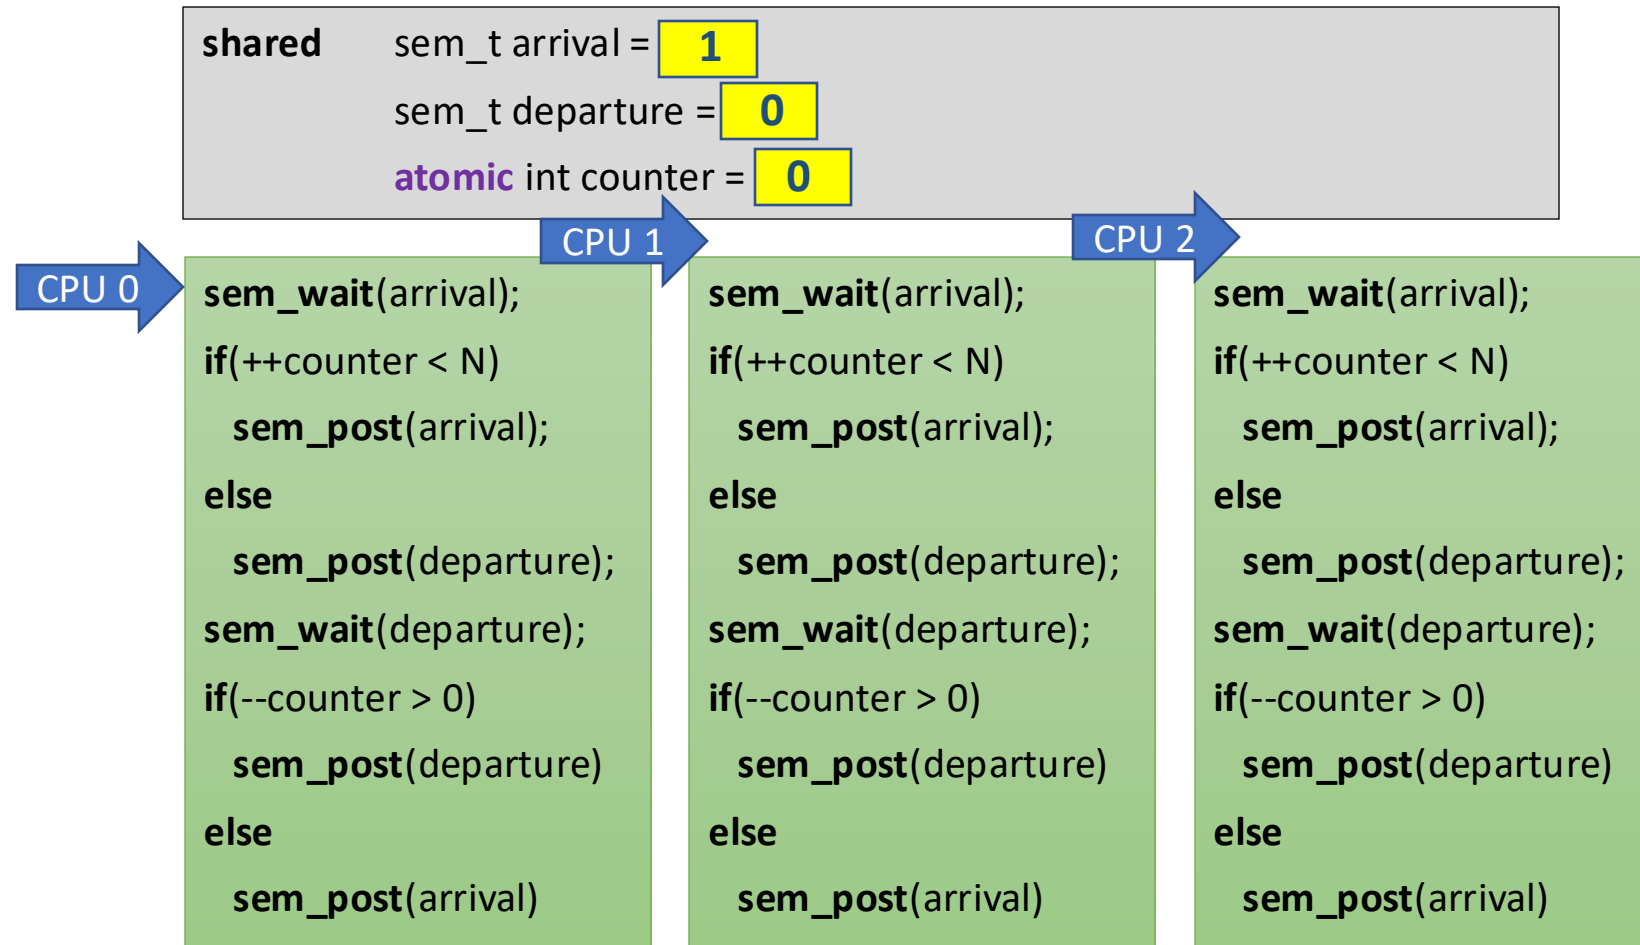
N == 3





# Semaphore Barrier Action Zone

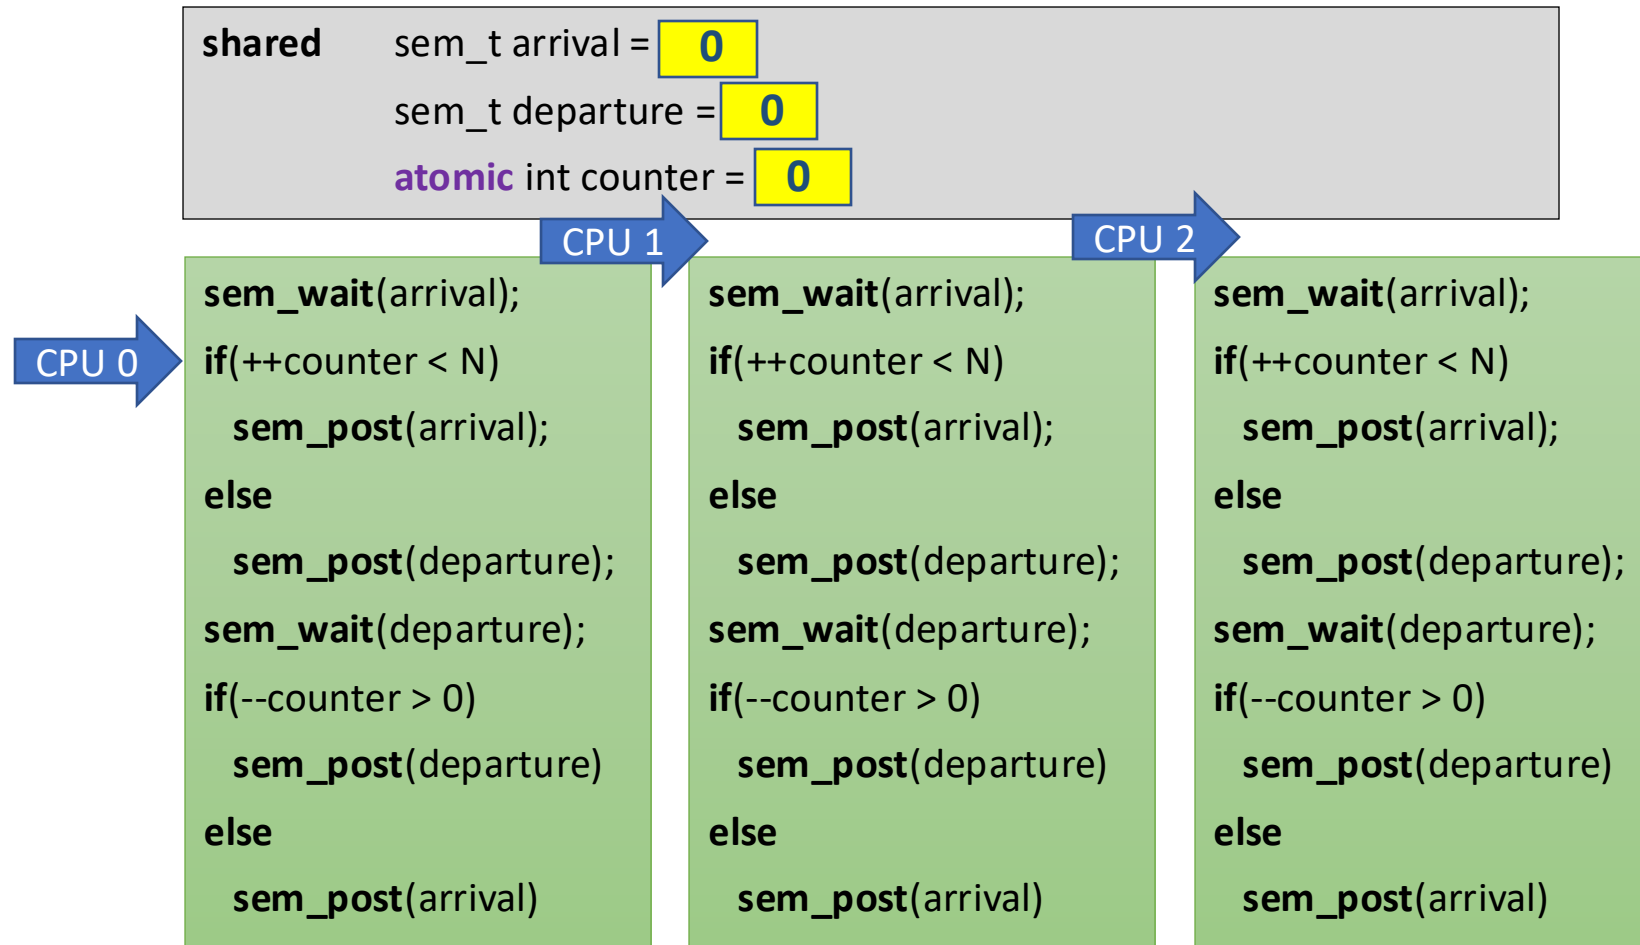
N == 3






# Semaphore Barrier Action Zone

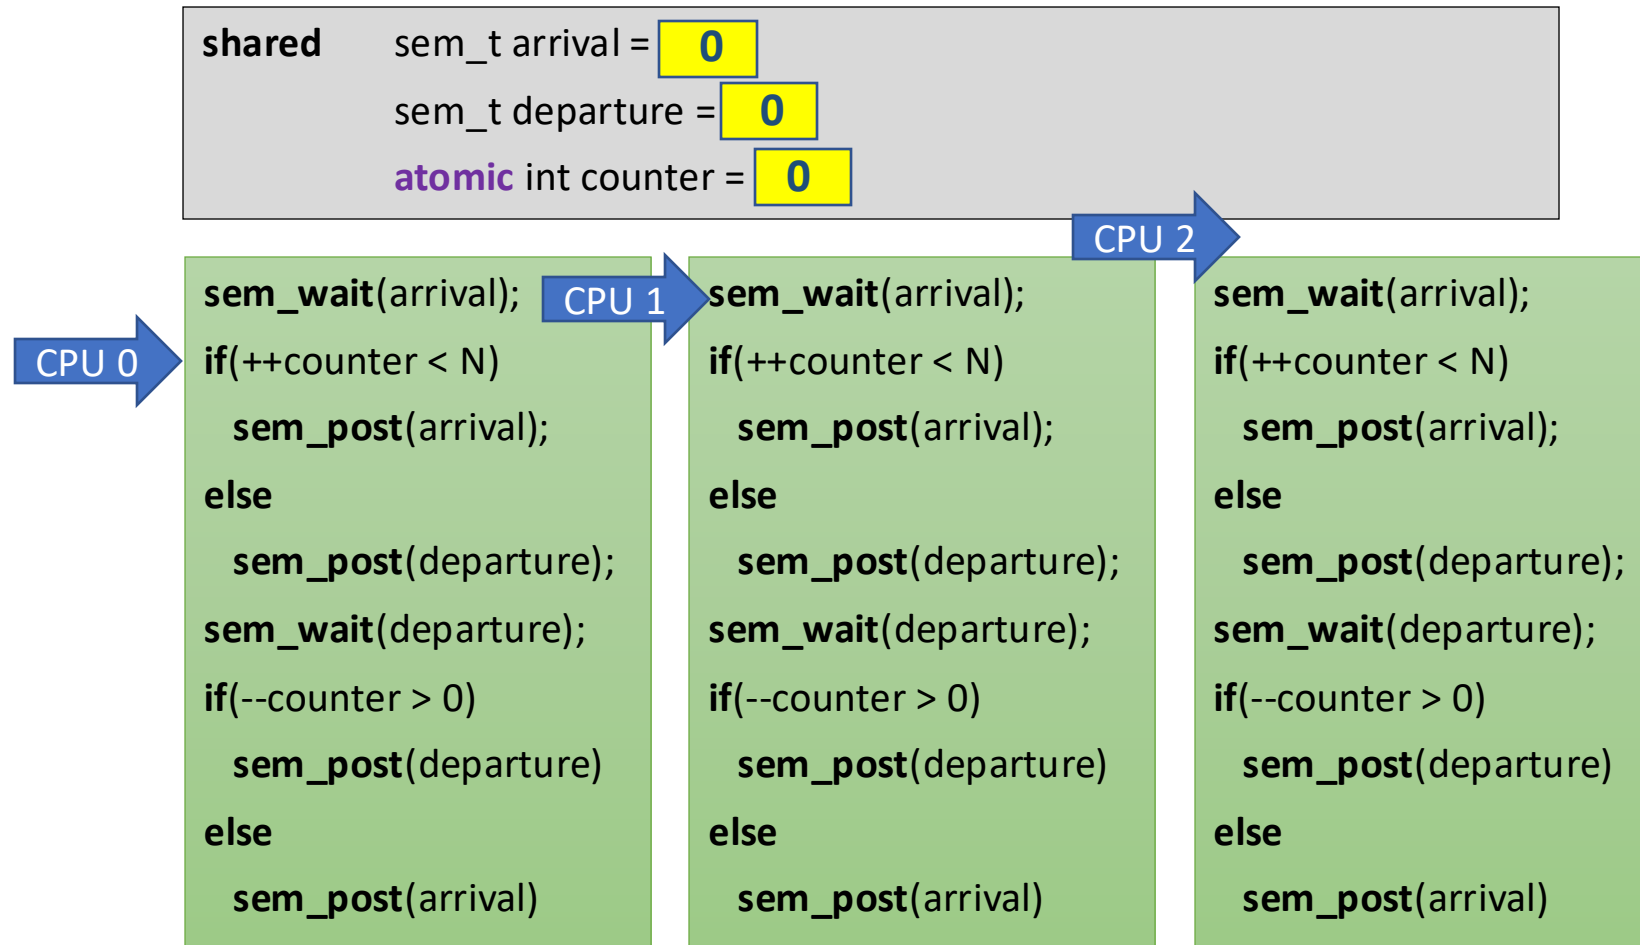
$N == 3$





# Semaphore Barrier Action Zone

N == 3

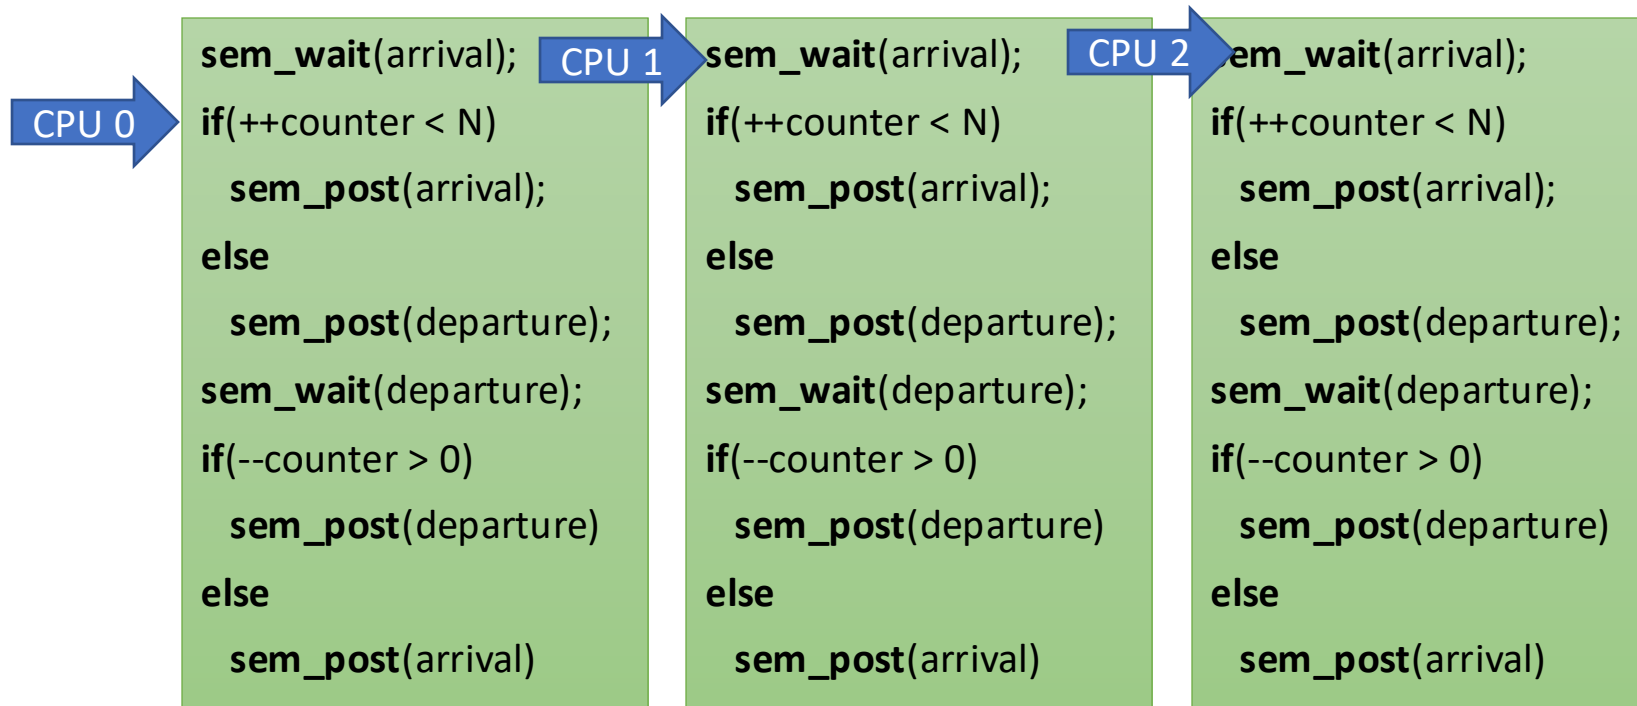




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
sem_t departure = 0
atomic int counter = 0
```

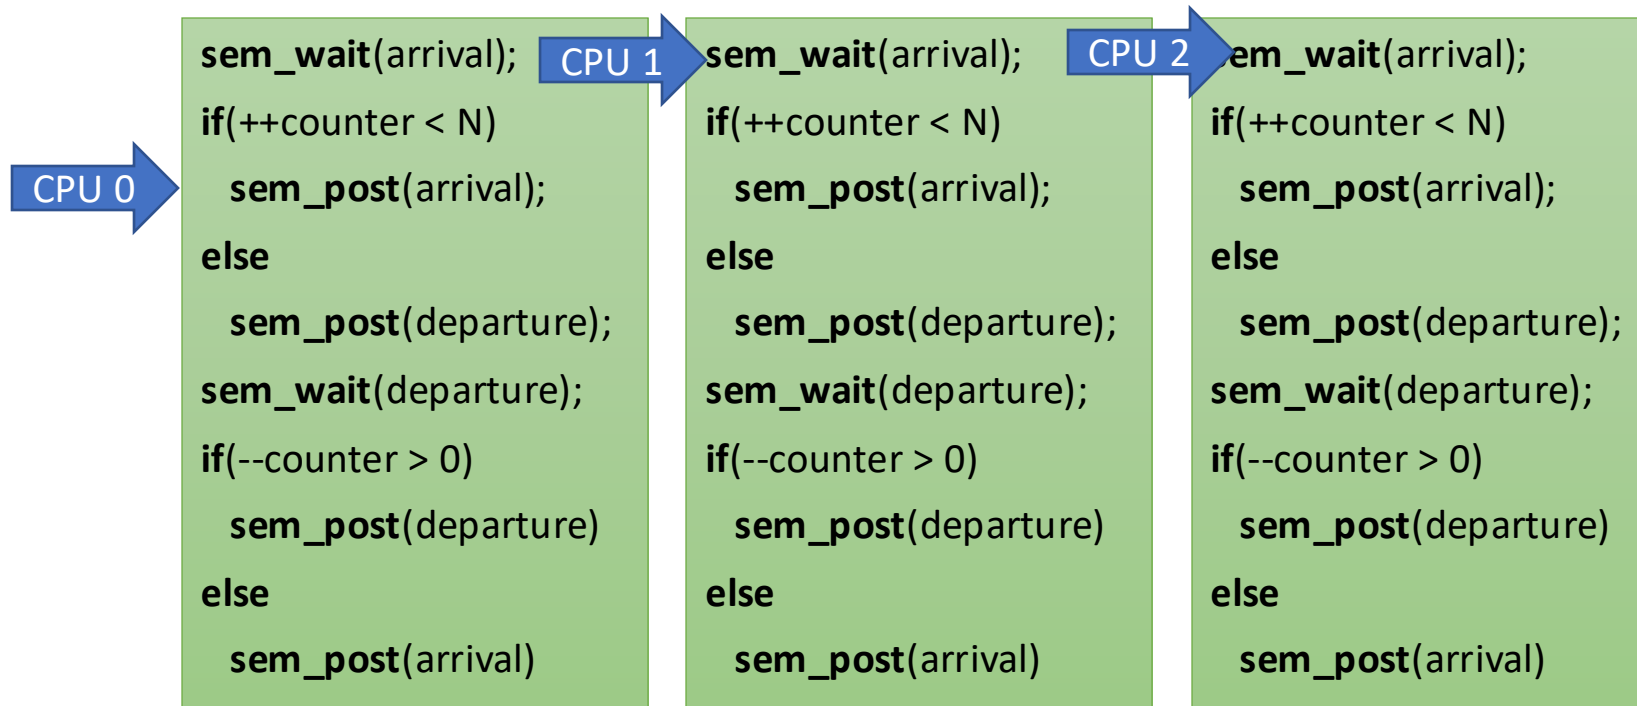




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 1
```



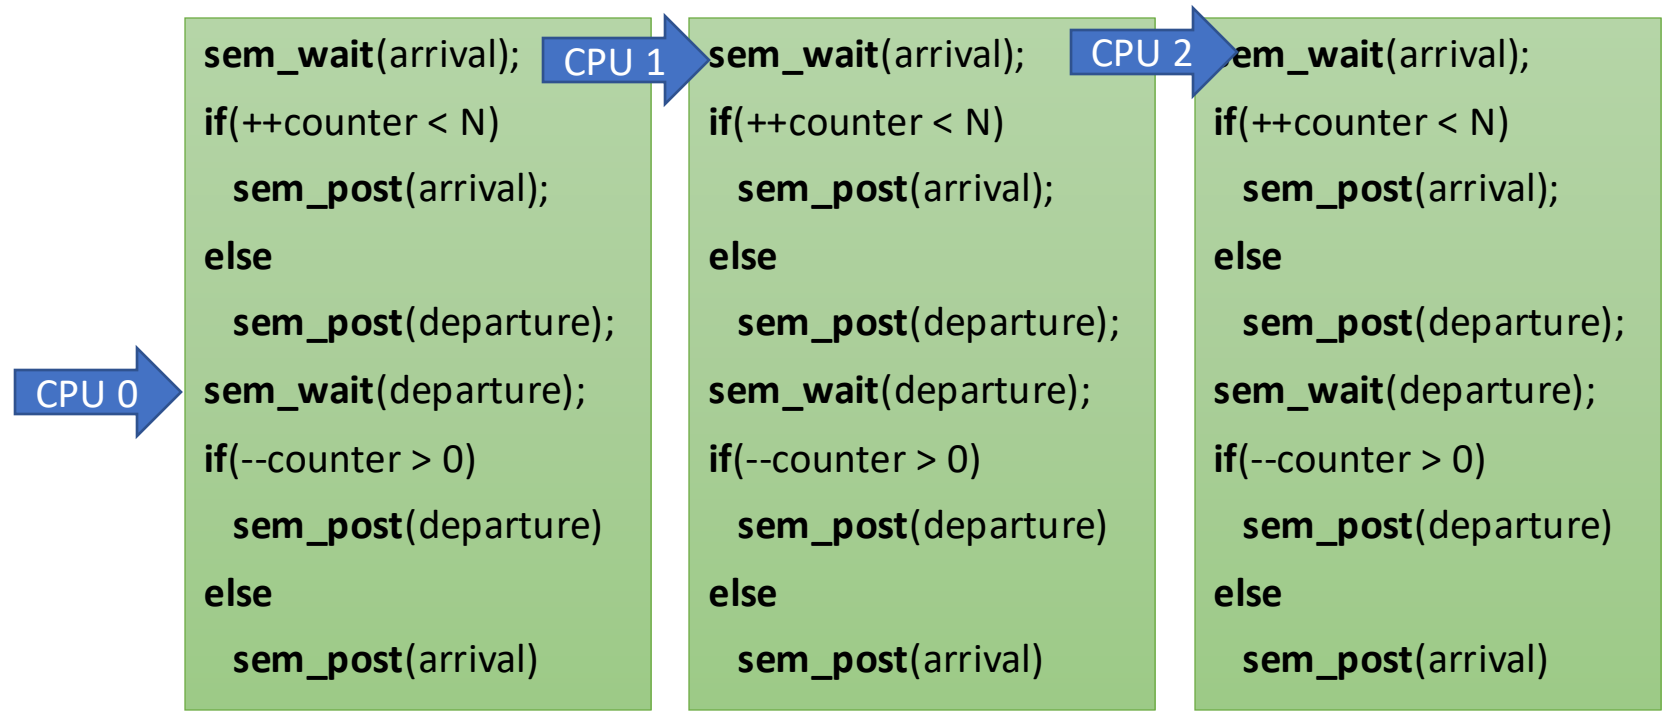





# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 1
sem_t departure = 0
atomic int counter = 1
```

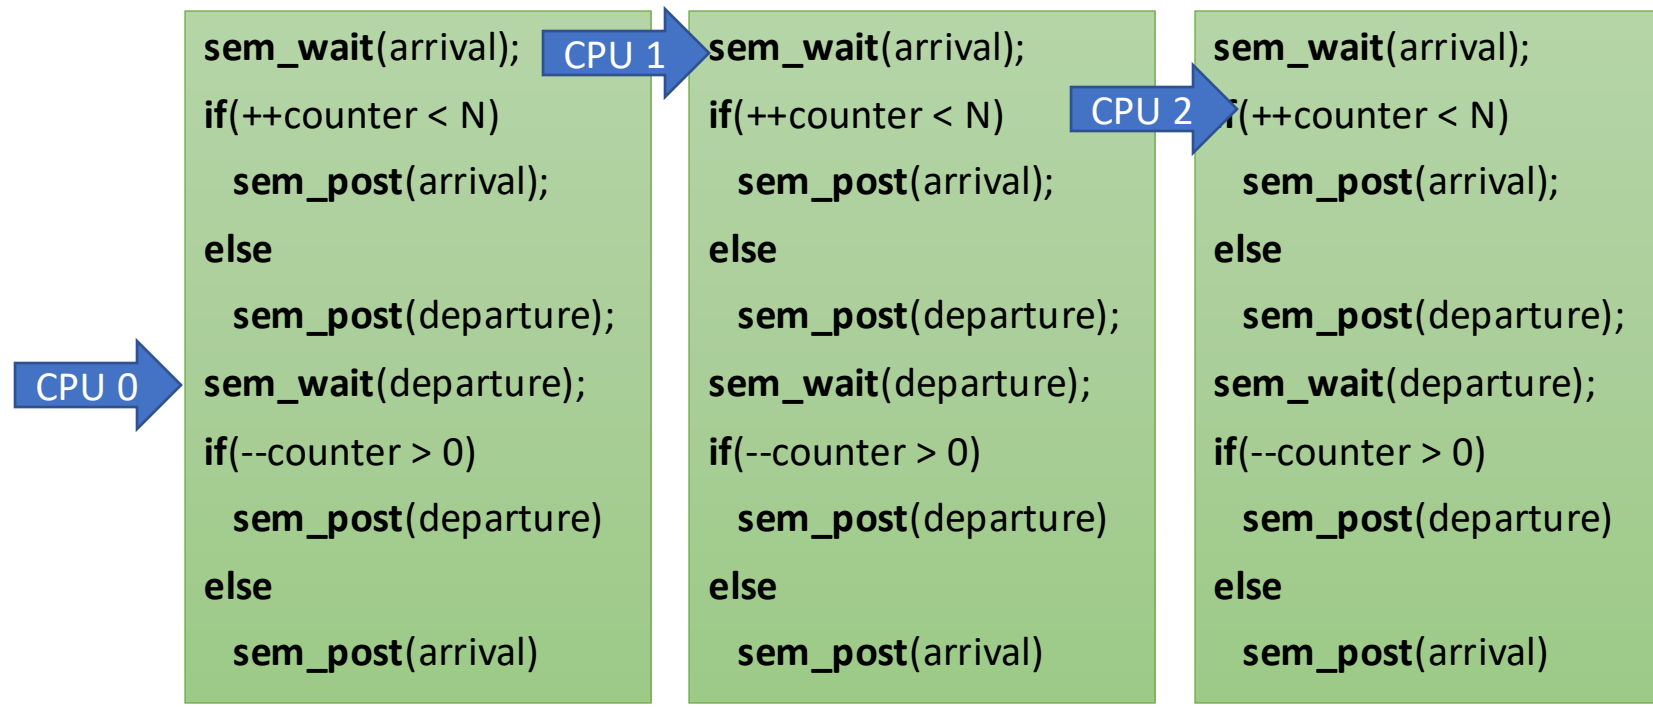




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 1
```

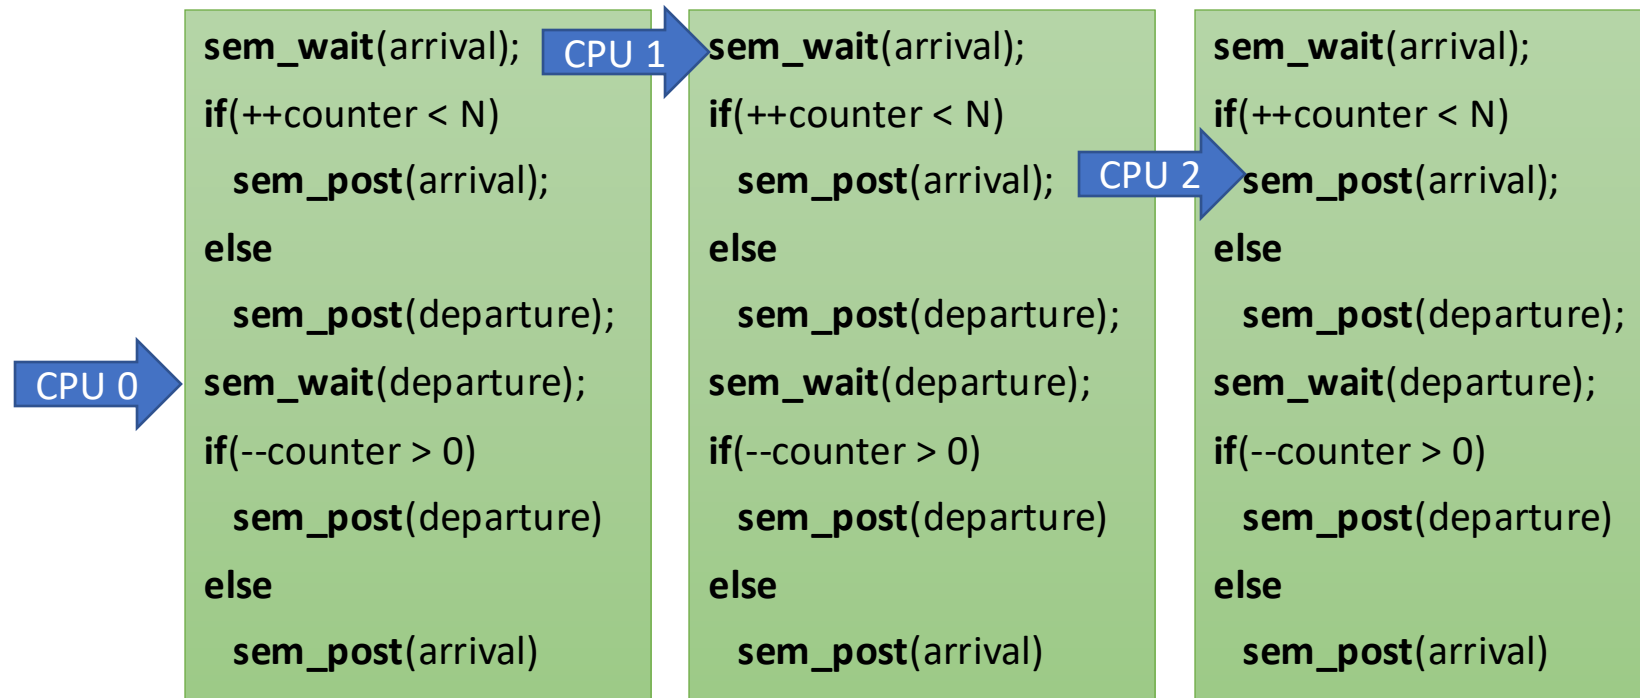




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 2
```

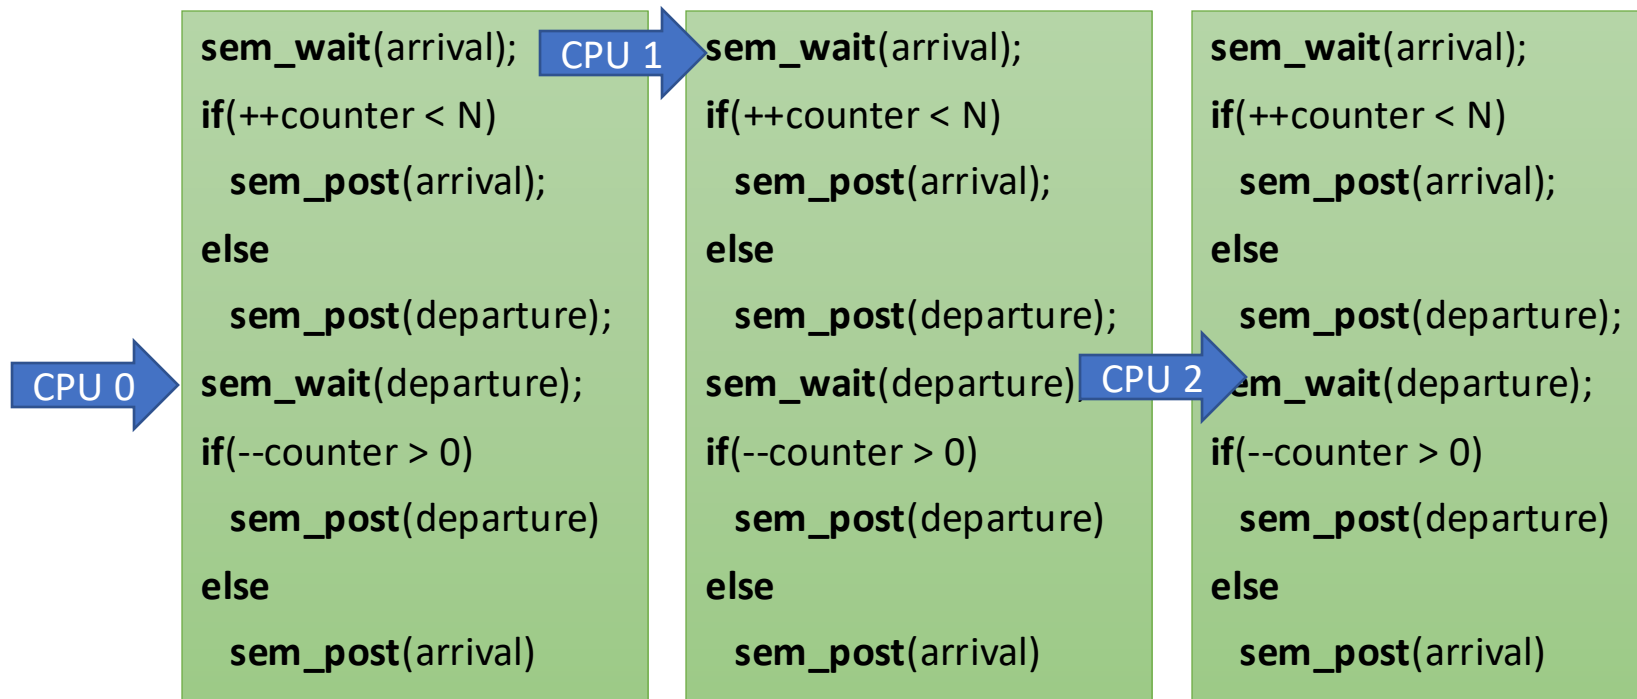




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 1
sem_t departure = 0
atomic int counter = 2
```

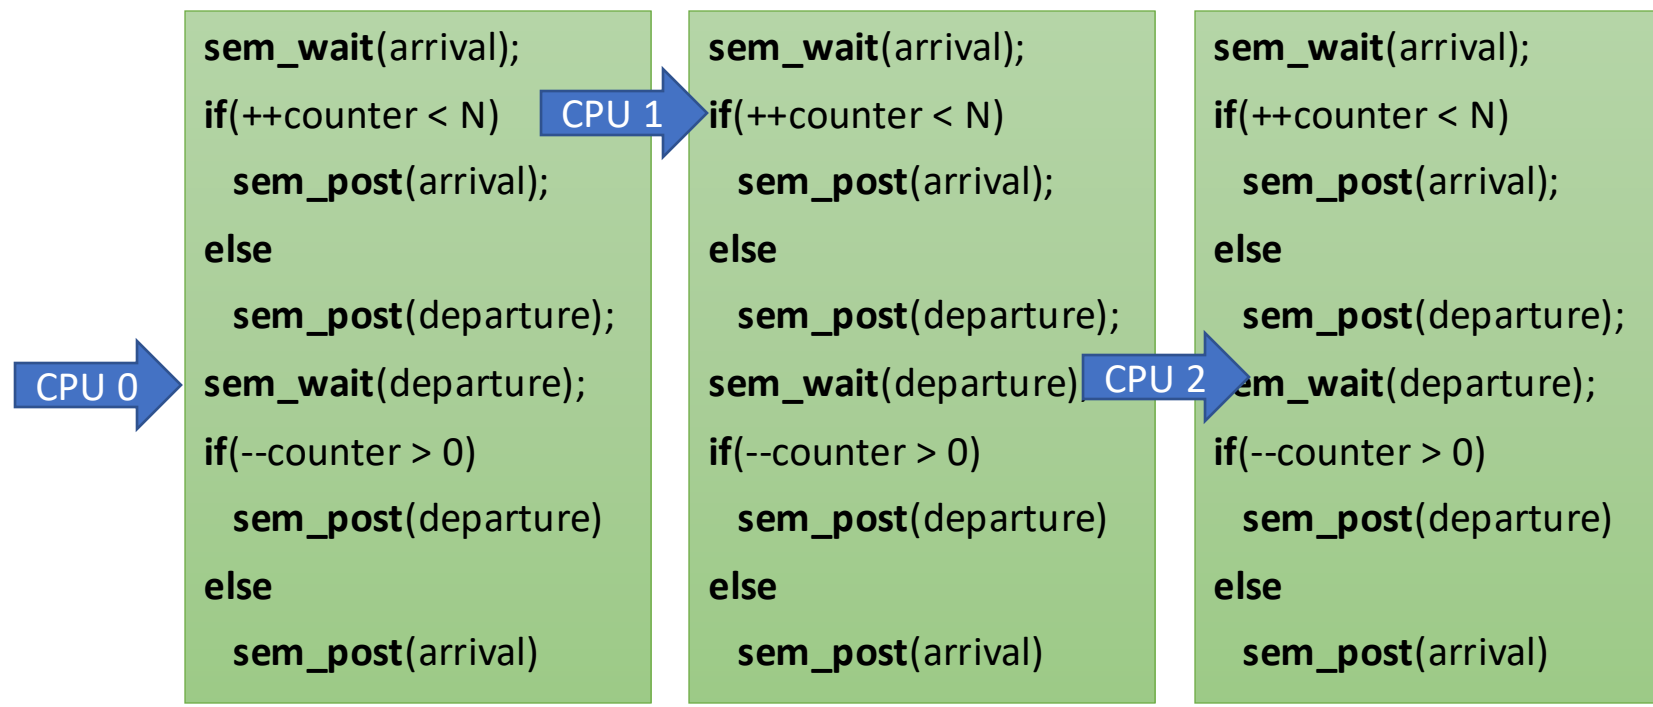





# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 2
```

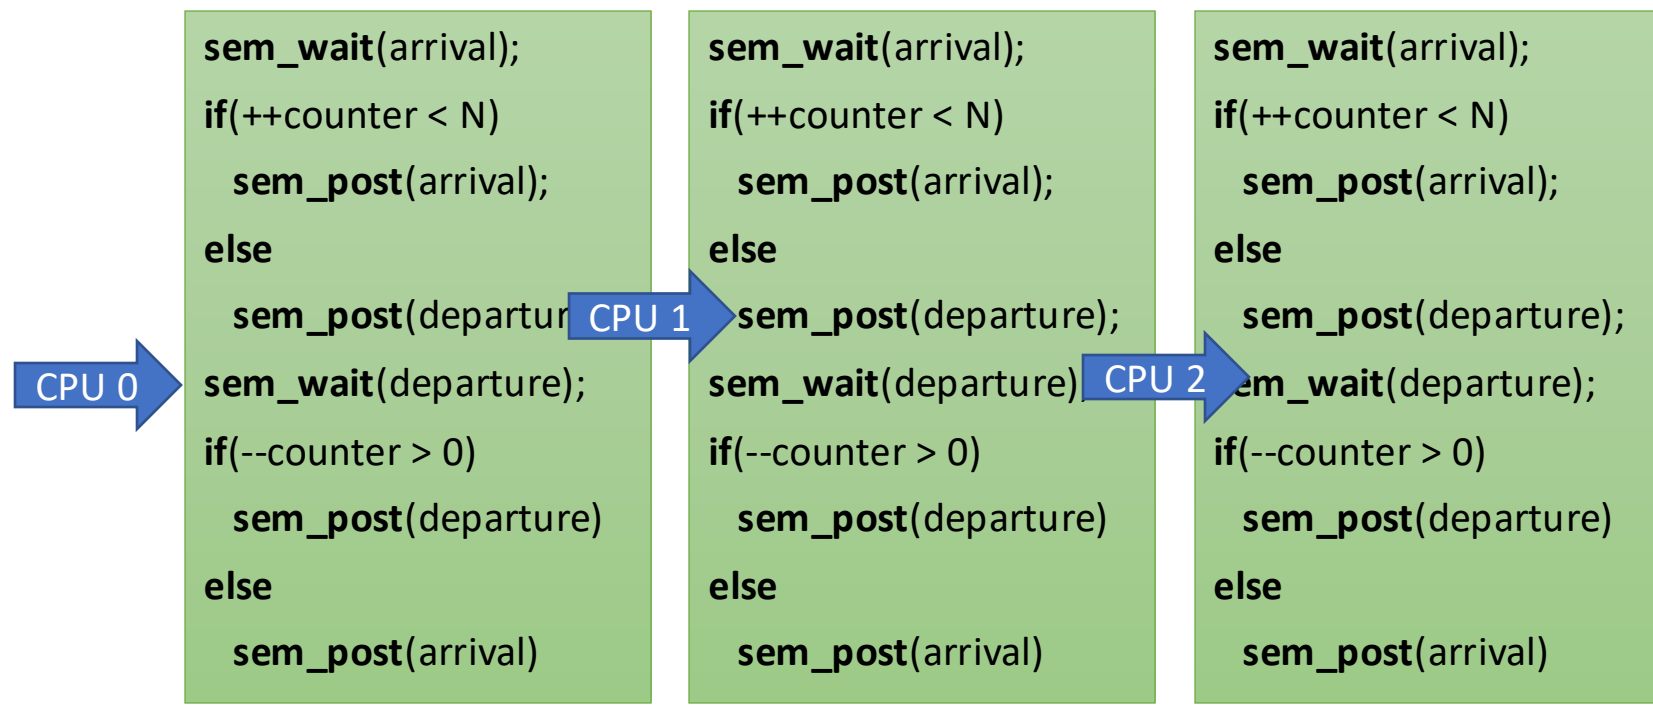




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 3
```

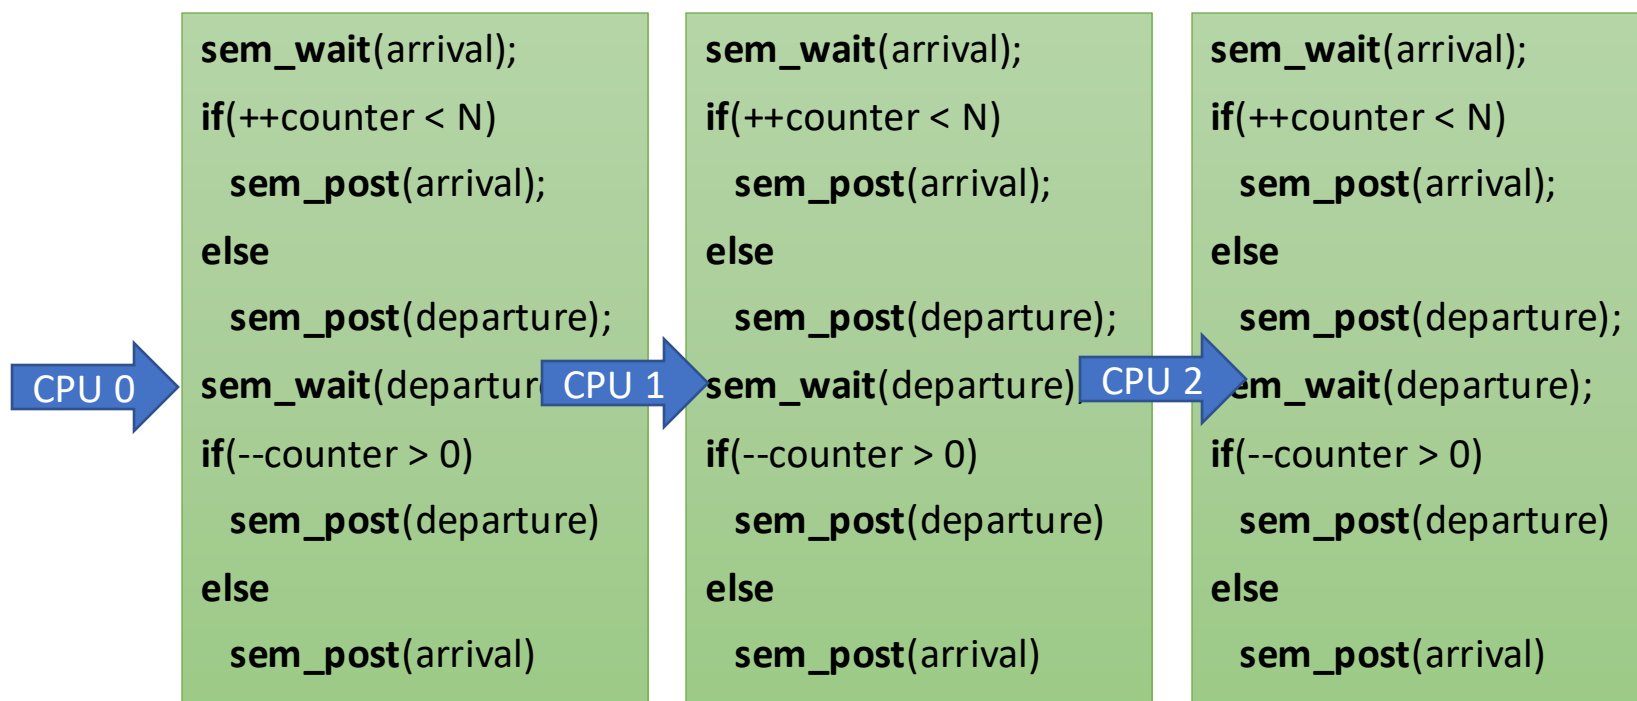





# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
sem_t departure = 1
atomic int counter = 3
```

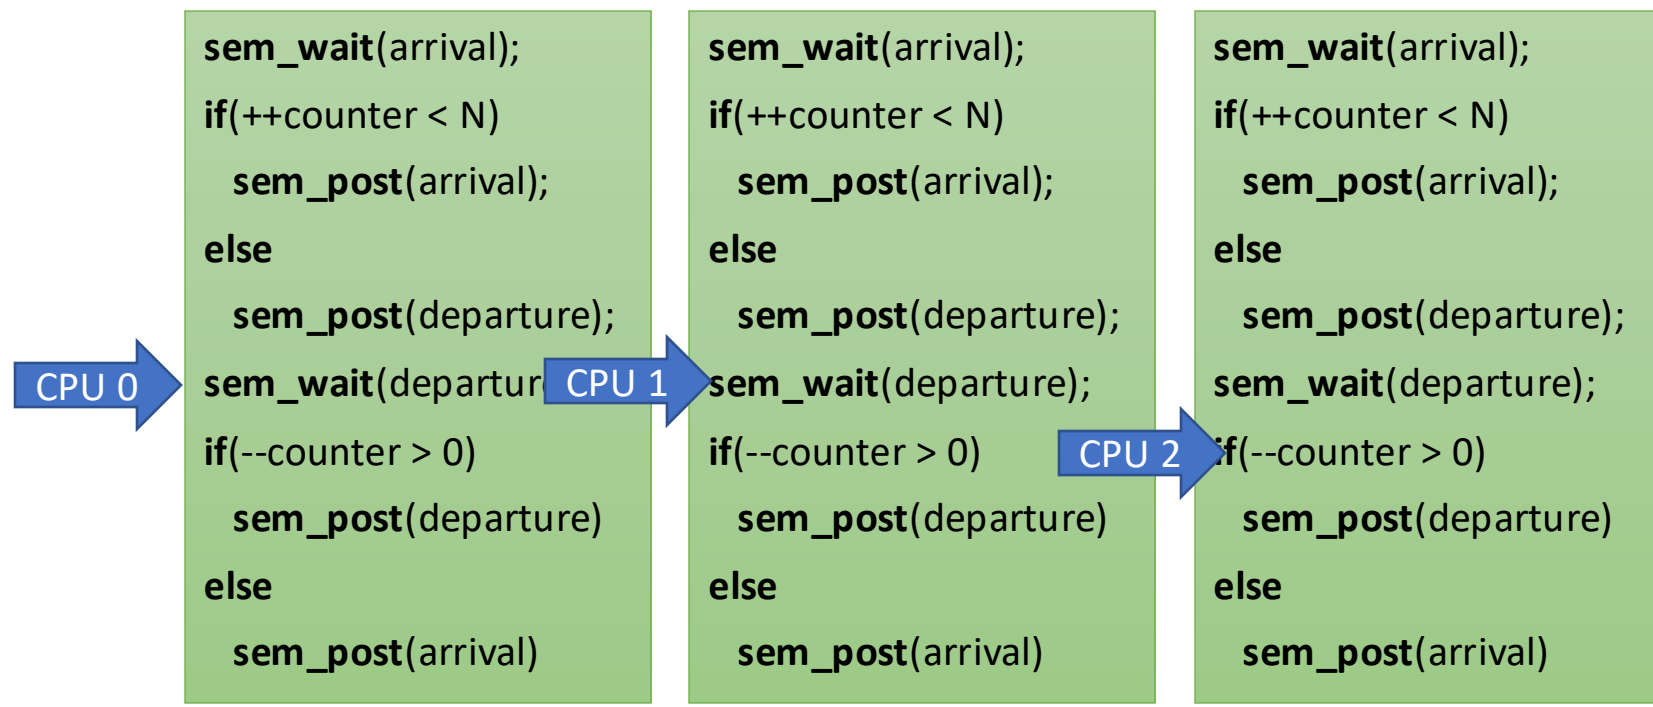




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 3
```



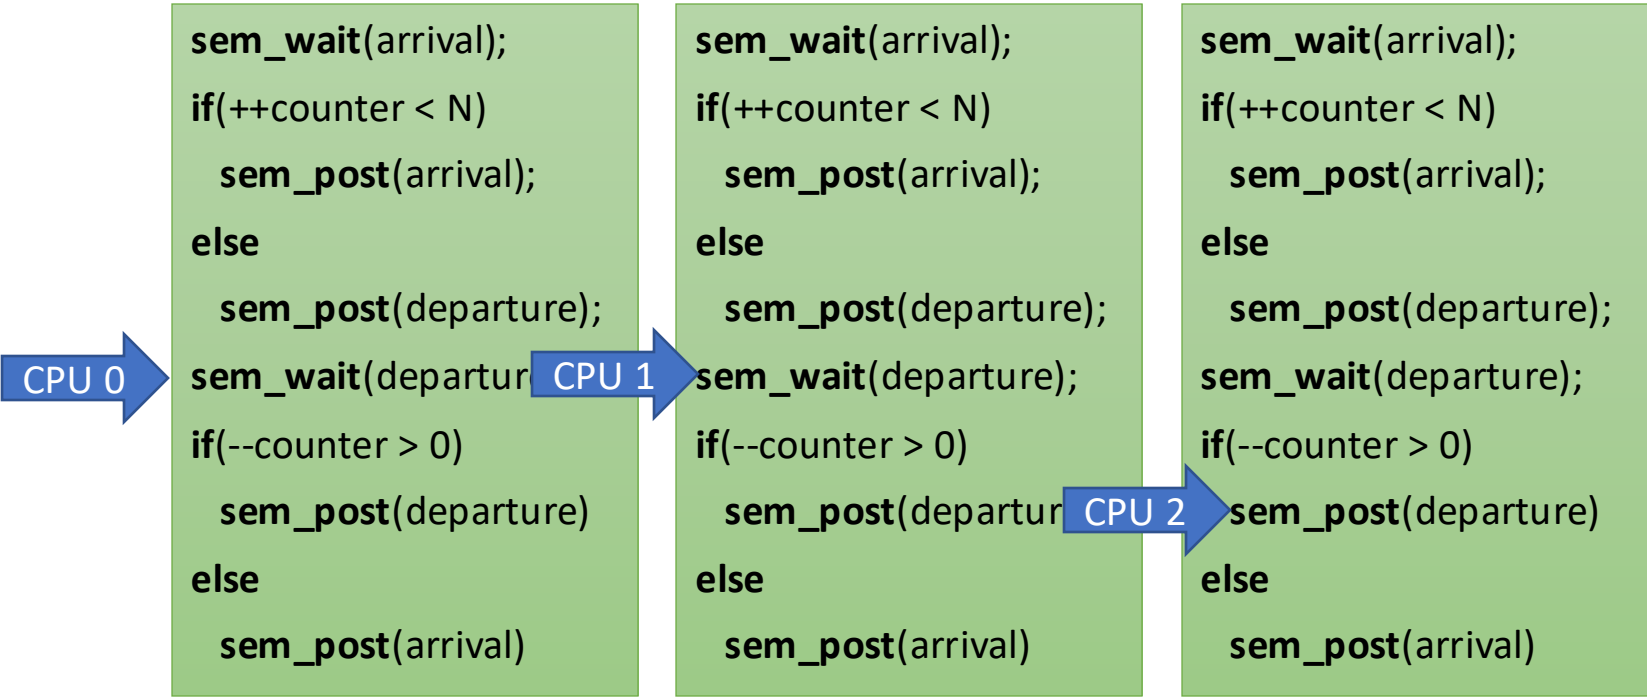





# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 2
```

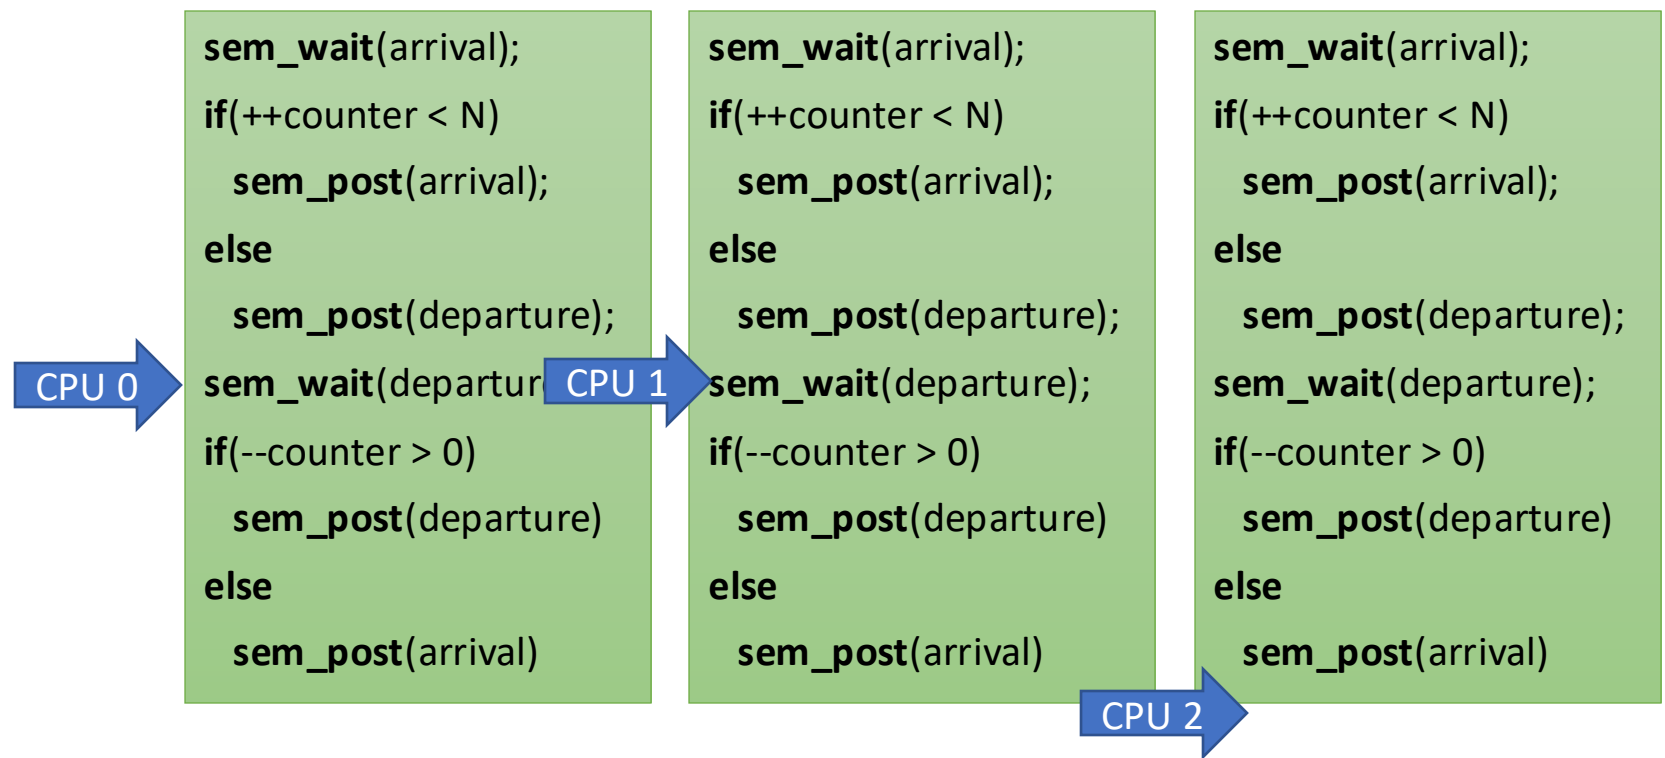




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 1
atomic int counter = 2
```





# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 2
```

CPU 0

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 2

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```



# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 1
```

CPU 0

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure);
else
    sem_post(arrival);
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure);
else
    sem_post(arrival);
```

CPU 2

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure);
else
    sem_post(arrival);
```



# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
sem_t departure = 1
atomic int counter = 1
```

CPU 0


```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 2

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```



# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 1
```

CPU 0

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 2

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```



# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 0
shared sem_t departure = 0
atomic int counter = 0
```

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 0

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 2



# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 1
sem_t departure = 0
atomic int counter = 0
```

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 0

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 1

```
sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)
```

CPU 2

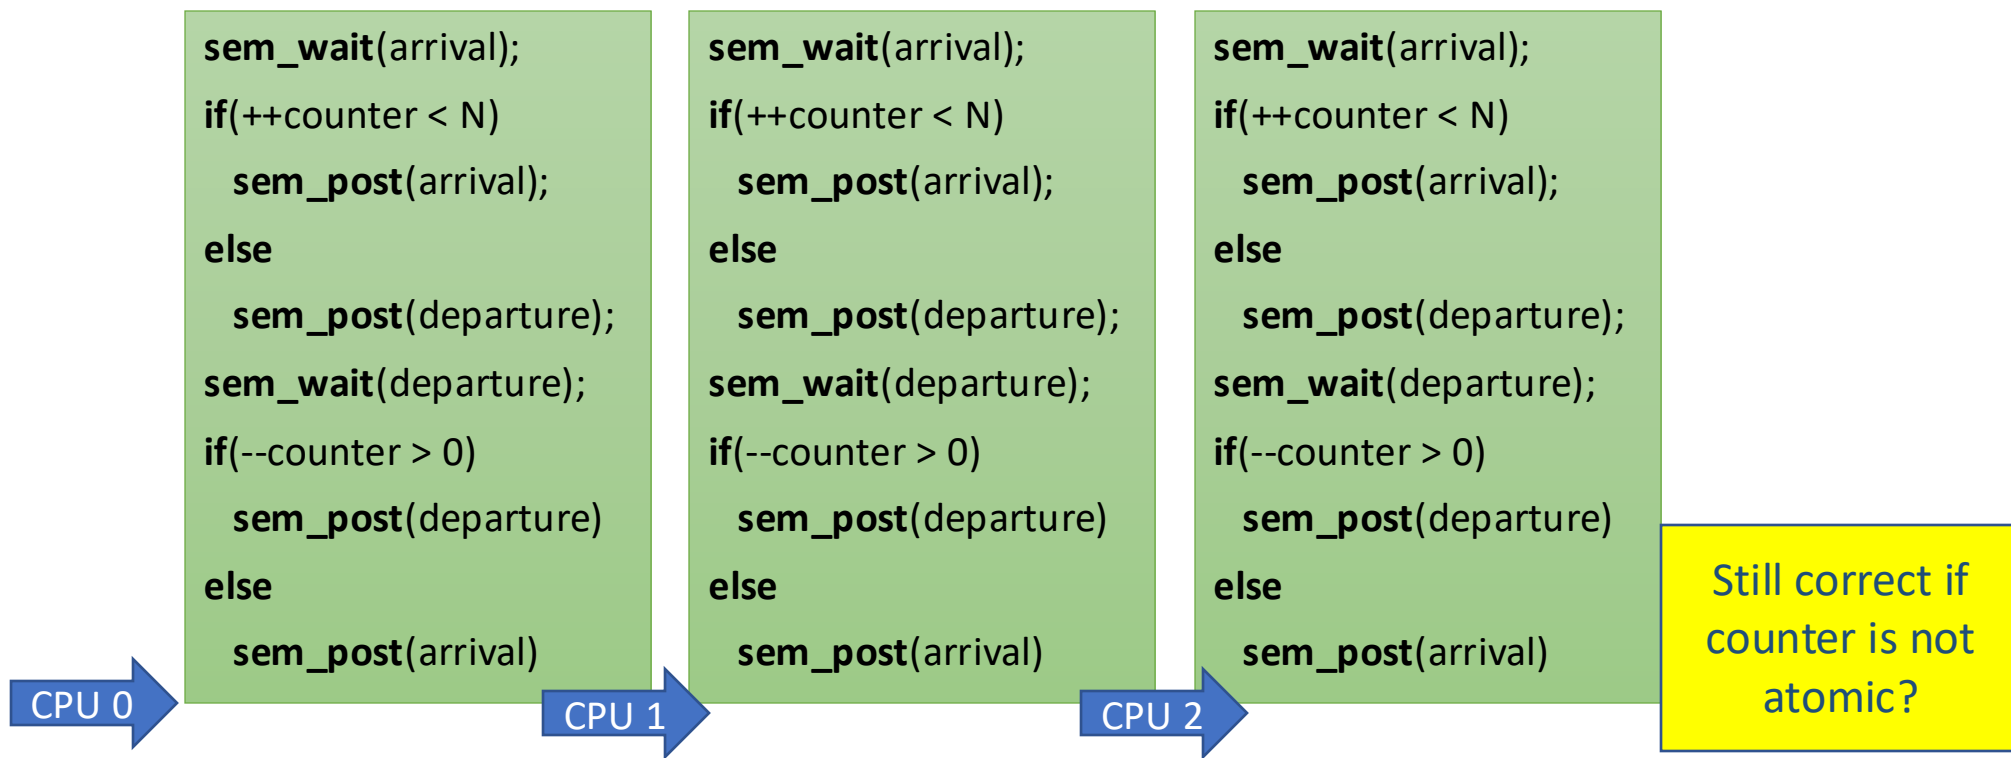




# Semaphore Barrier Action Zone

N == 3

```
shared sem_t arrival = 1
shared sem_t departure = 0
atomic int counter = 0
```





# Semaphore Barrier Action Zone

N == 3

```

shared  sem_t arrival = 1
        sem_t departure = 0
        atomic int counter = 0

```

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

CPU 0

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

CPU 1

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

CPU 2

Do we need two phases?

Still correct if counter is not atomic?



# Semaphore Barrier Action Zone

N == 3

```

shared sem_t arrival = 1
sem_t departure = 0
atomic int counter = 0

```

```

// why two phases:
for(...) {
    do_something();
    wait();
}

```

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

```

sem_wait(arrival);
if(++counter < N)
    sem_post(arrival);
else
    sem_post(departure);
sem_wait(departure);
if(--counter > 0)
    sem_post(departure)
else
    sem_post(arrival)

```

CPU 0

CPU 1

CPU 2

Do we need two phases?

Still correct if counter is not atomic?

# Barrier using Semaphores

## Properties

- Pros:

- Cons:

# Barrier using Semaphores

## Properties

- **Pros:**

- Very Simple
- Space complexity  $O(1)$
- Symmetric

- **Cons:**

# Barrier using Semaphores

## Properties

- **Pros:**

- Very Simple
- Space complexity  $O(1)$
- Symmetric

- **Cons:**

- Required a strong object
  - Requires some central manager
  - High contention on the semaphores
- Propagation delay  $O(n)$



Barriers based on counters



# Counter Barrier Ingredients

## Fetch-and-Increment register

- A shared register that supports a F&I operation:
- Input: register  $r$
- Atomic operation:
  - $r$  is incremented by 1
  - the old value of  $r$  is returned

```
function fetch-and-increment (r : register)
  orig_r := r;
  r := r + 1;
  return (orig_r);
end-function
```

## Await

- For brevity, we use the **await** macro
- Not an operation of an object
- This is just “spinning”

```
macro await (condition : boolean condition)
  repeat
    cond = eval(condition);
  until (cond)
end-macro
```



# Simple Barrier Using an Atomic Counter

<b>shared</b>	counter: fetch and increment reg. – {0,..n}, initially = 0
	go: atomic bit, initial value does not matter
<b>local</b>	local.go: a bit, initial value does not matter
	local.counter: register




# Simple Barrier Using an Atomic Counter

```
shared    counter: fetch and increment reg. – {0,..n}, initially = 0  
           go: atomic bit, initial value does not matter  
local    local.go: a bit, initial value does not matter  
           local.counter: register
```

```
1  local.go := go  
2  local.counter := fetch-and-increment (counter)  
3  if local.counter + 1 = n then  
4      counter := 0  
5      go := 1 - go  
6  else await(local.go ≠ go)
```

# Simple Barrier Using an Atomic Counter


Run for n=2 Threads



```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

# Simple Barrier Using an Atomic Counter


Run for n=2 Threads



```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```


# Simple Barrier Using an Atomic Counter

Run for n=2 Threads





# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






P1 →

```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






P1 →

```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```



# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






P1 →

```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






P1 →

```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

# Simple Barrier Using an Atomic Counter

Run for n=2 Threads





```
1 local.go := go
2 local.counter := fetch-and-increment
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

P1 →

0+1≠2

# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

P1 →

# Simple Barrier Using an Atomic Counter

Run for n=2 Threads




```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

P1 →

P1 Busy wait

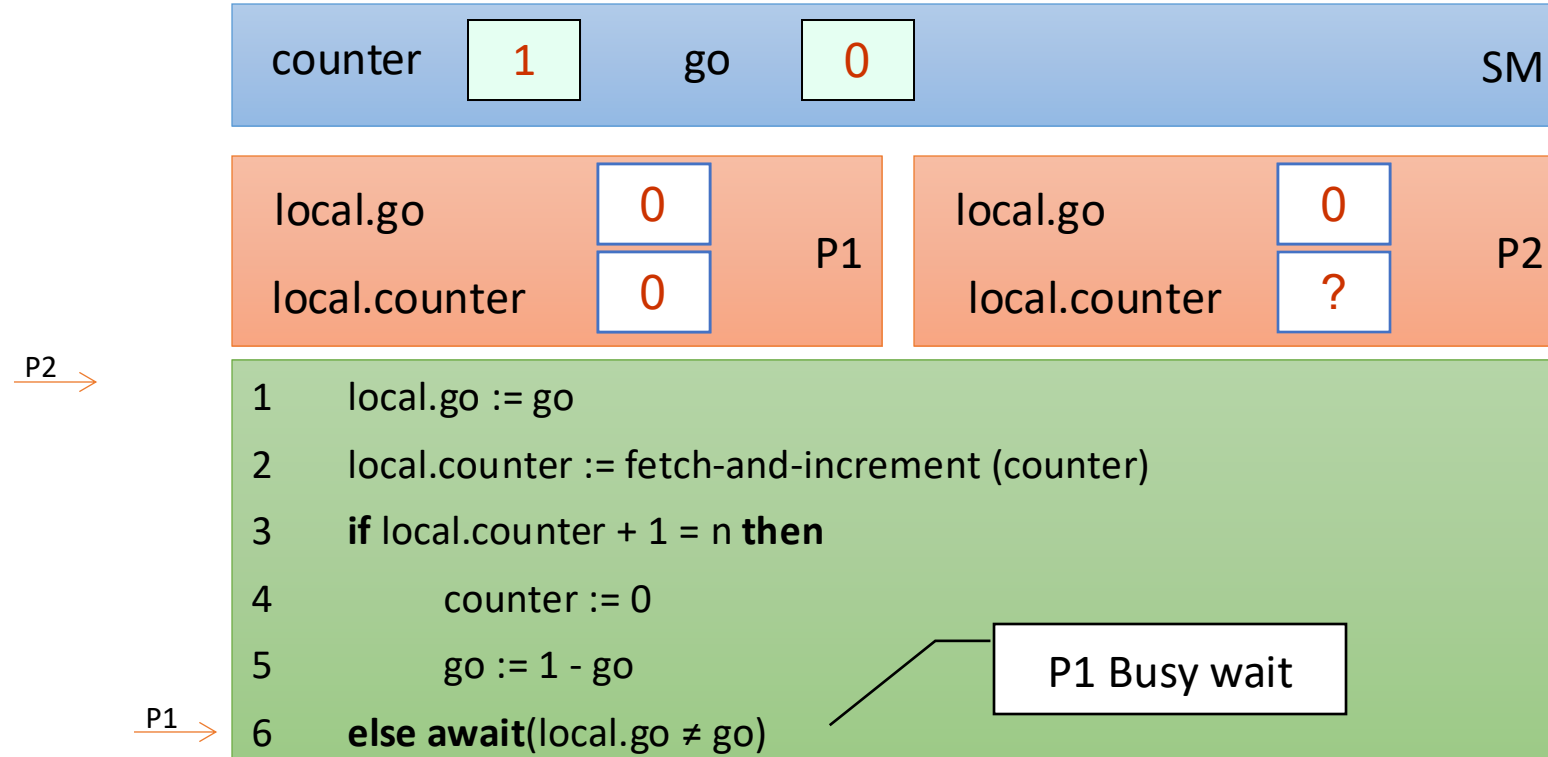
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



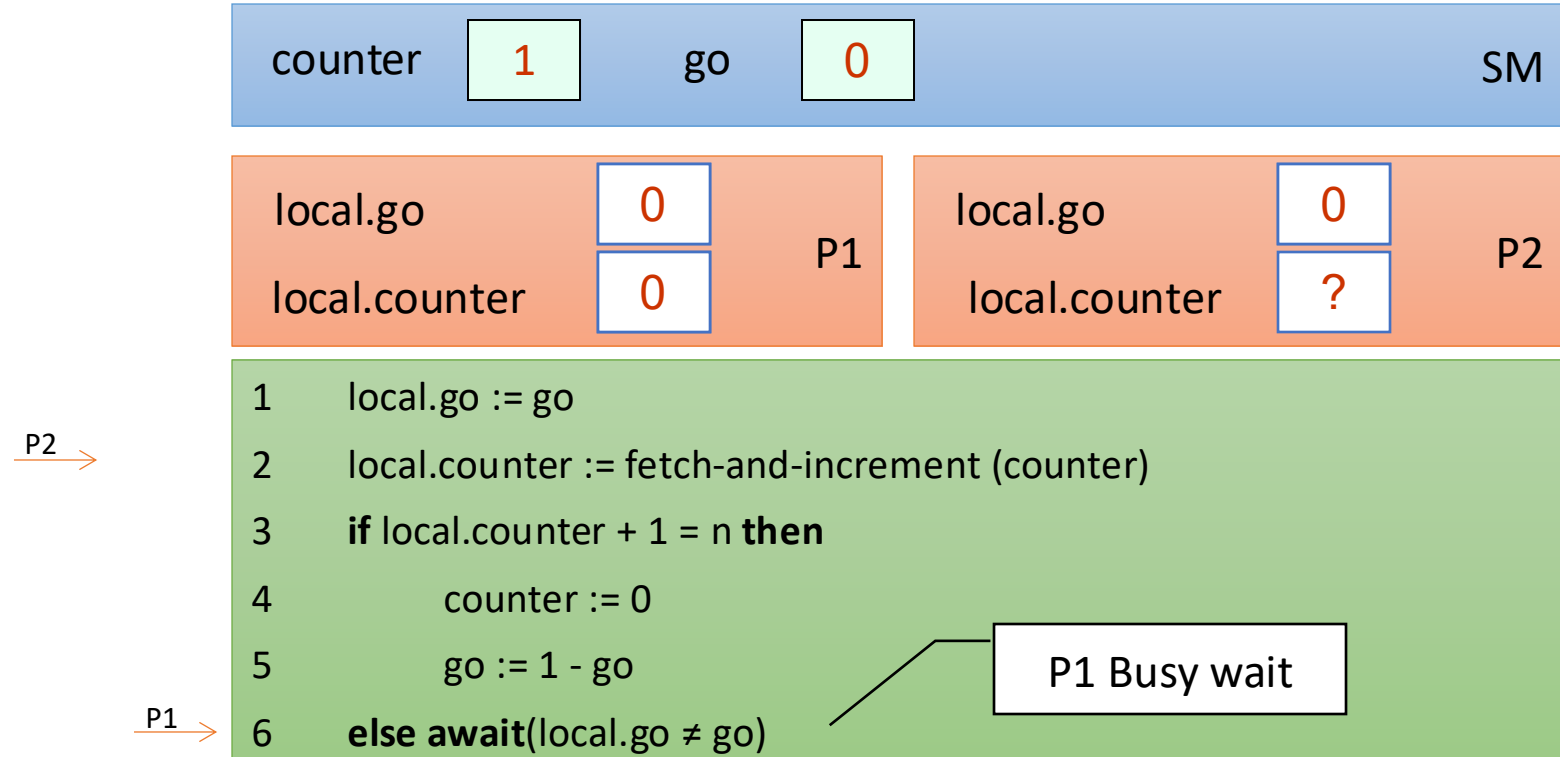
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



# Simple Barrier Using an Atomic Counter

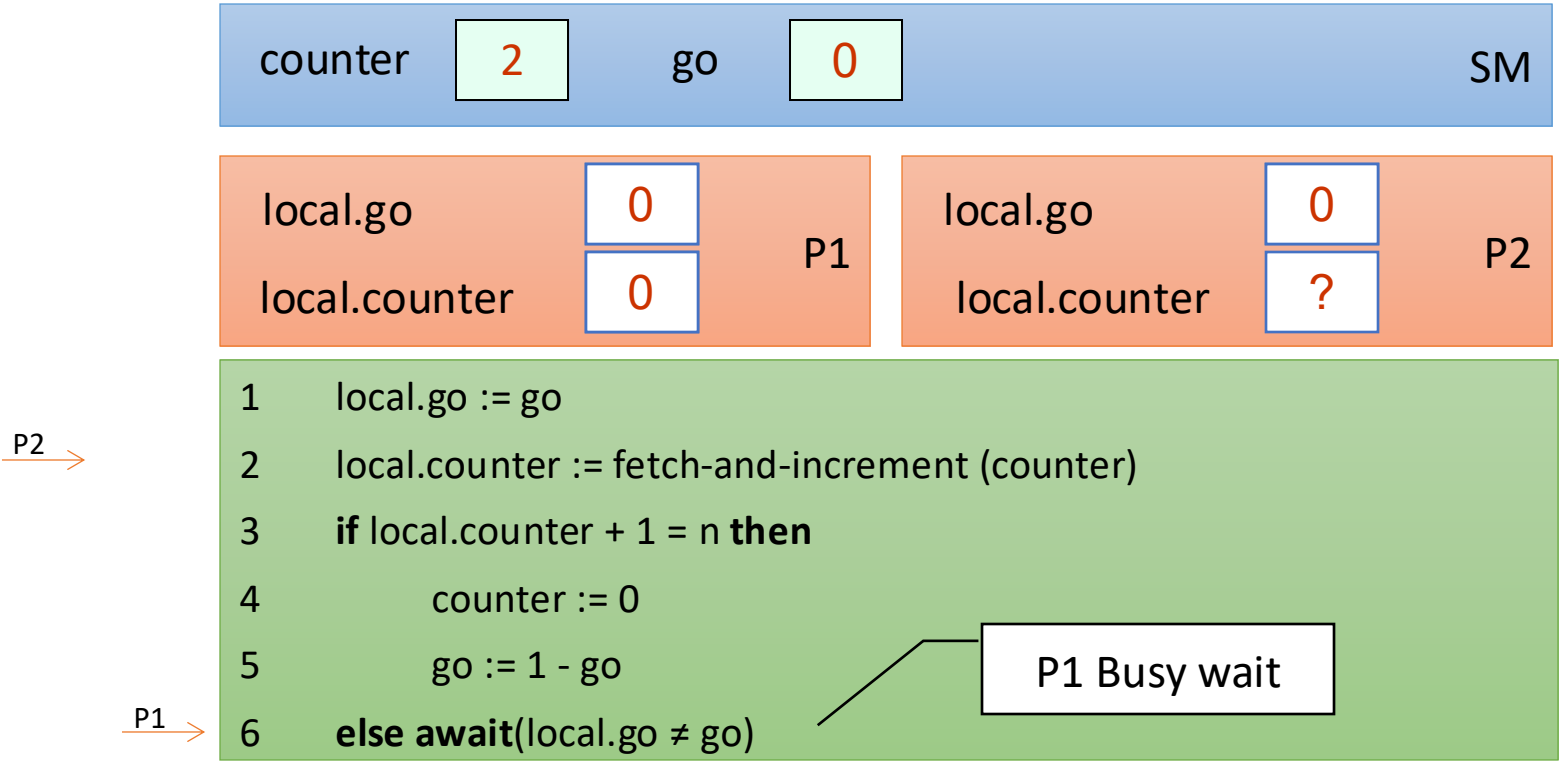
Run for n=2 Threads





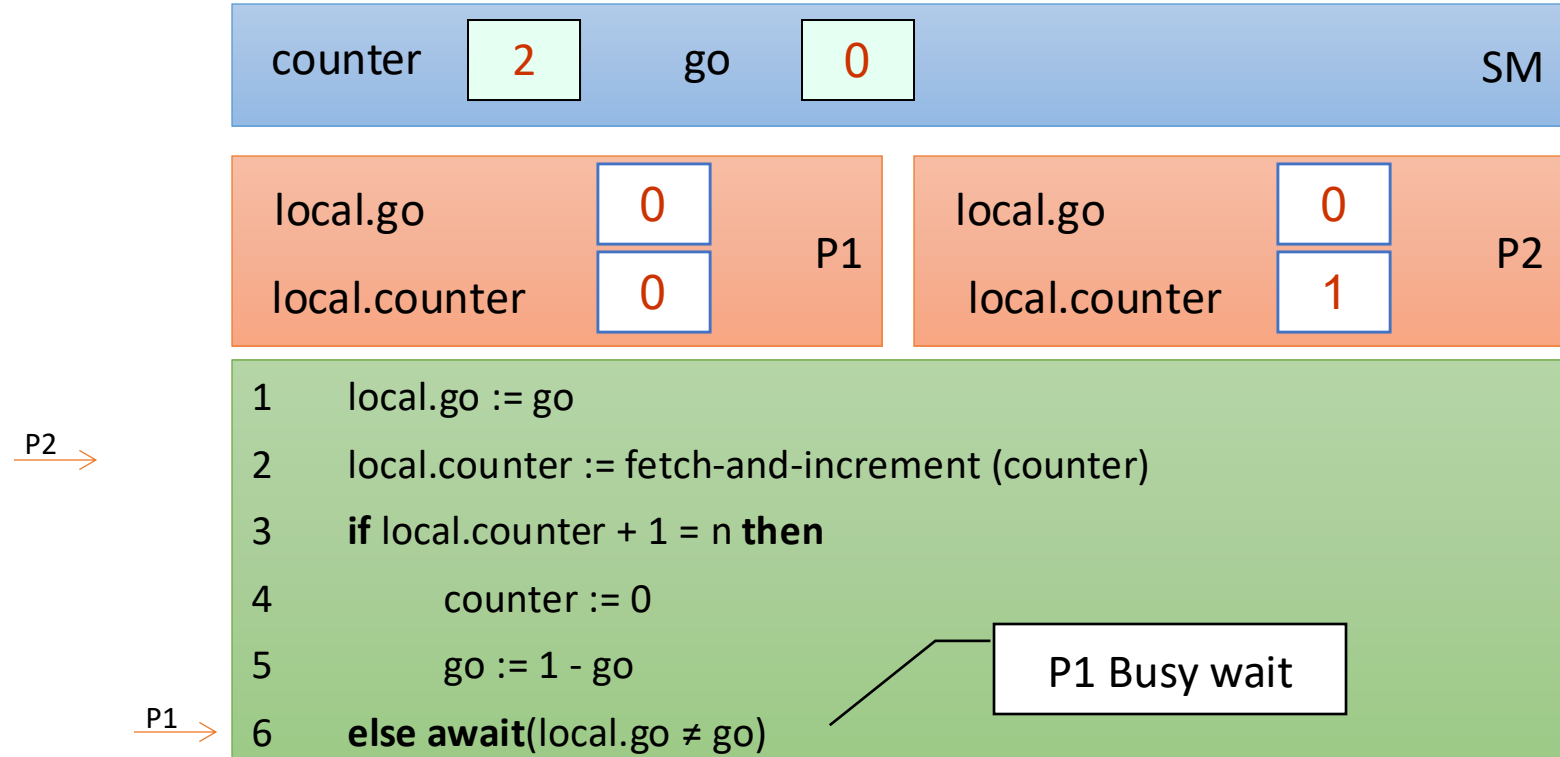
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



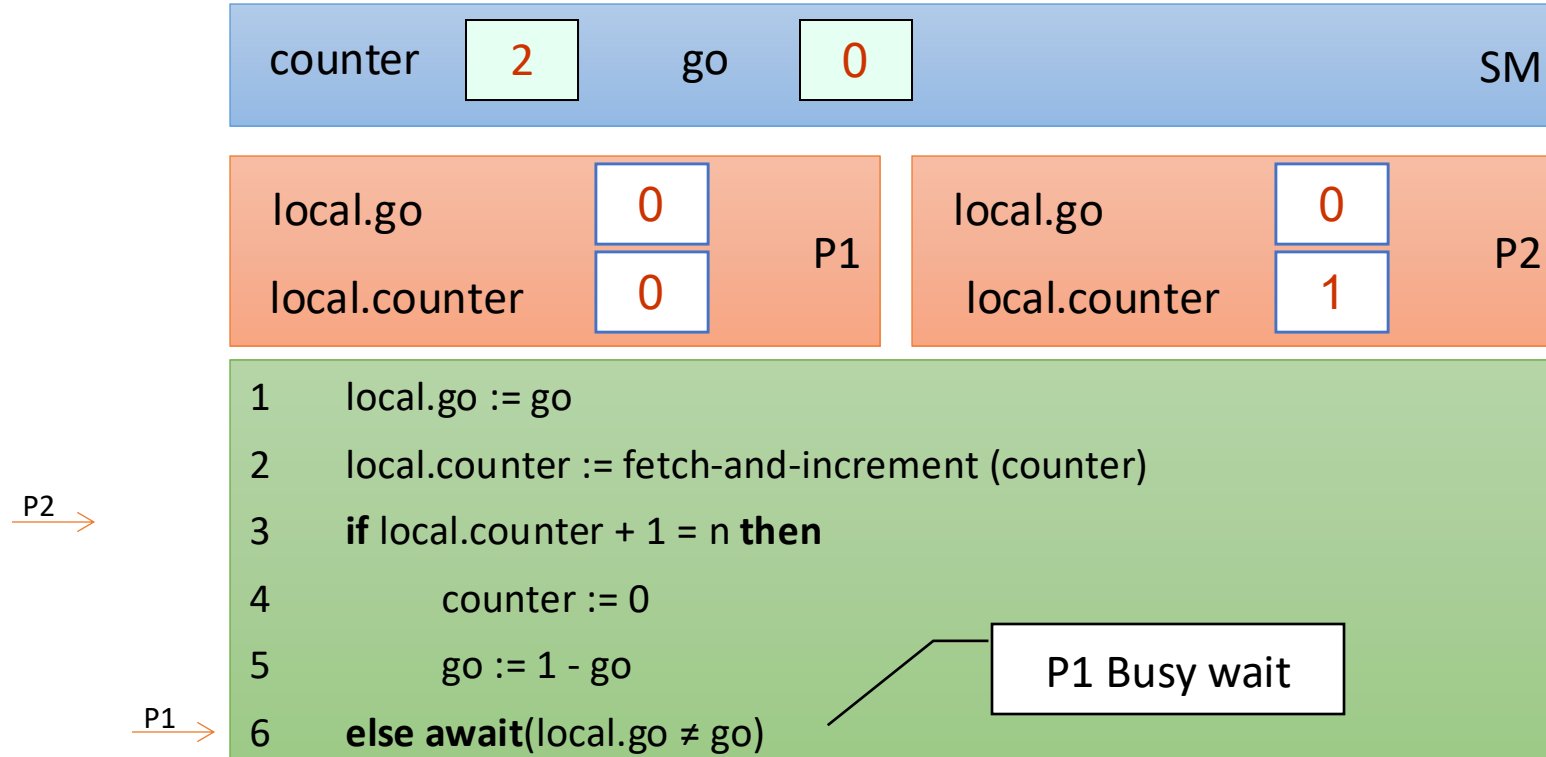
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



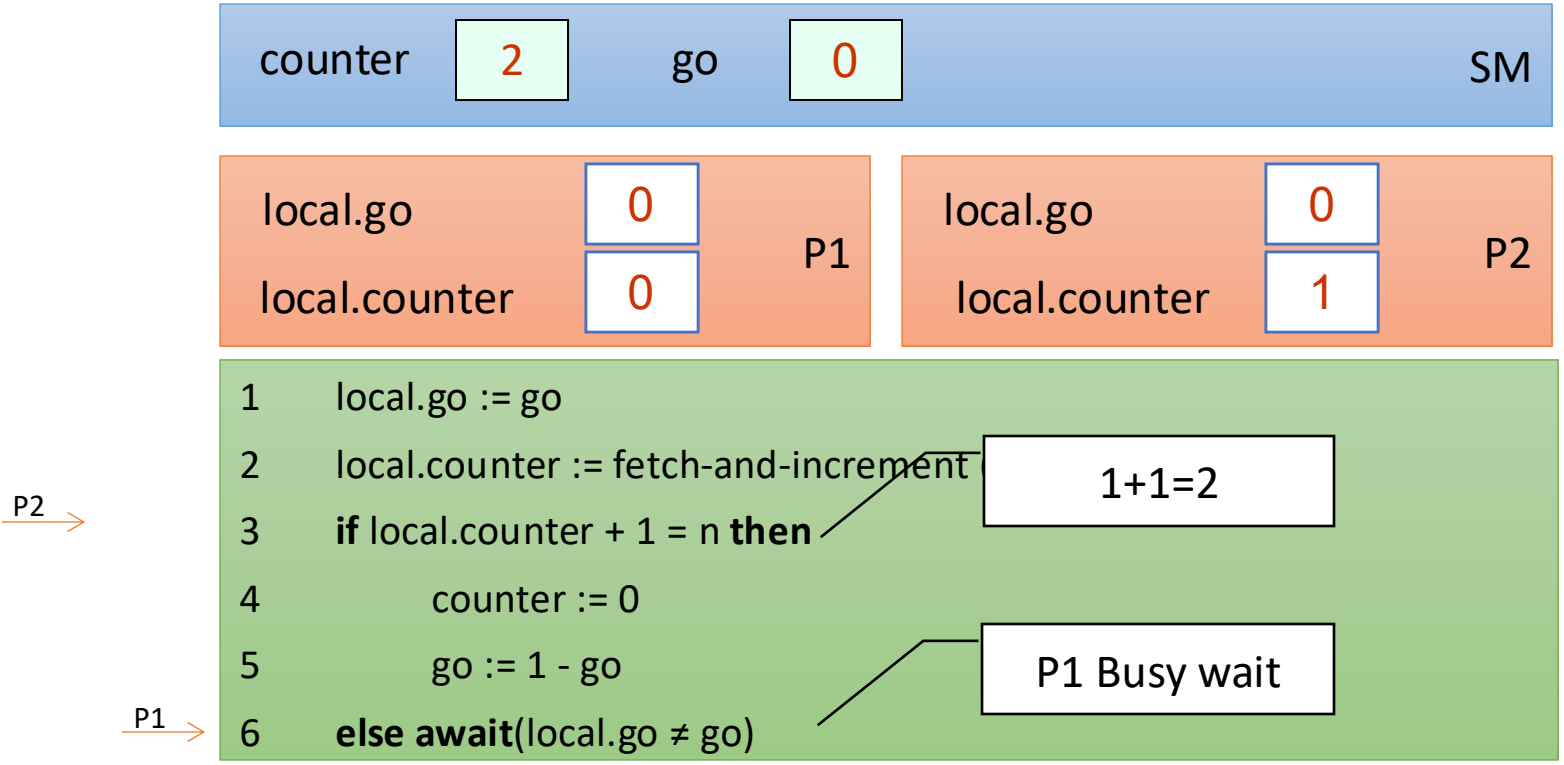
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads






# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

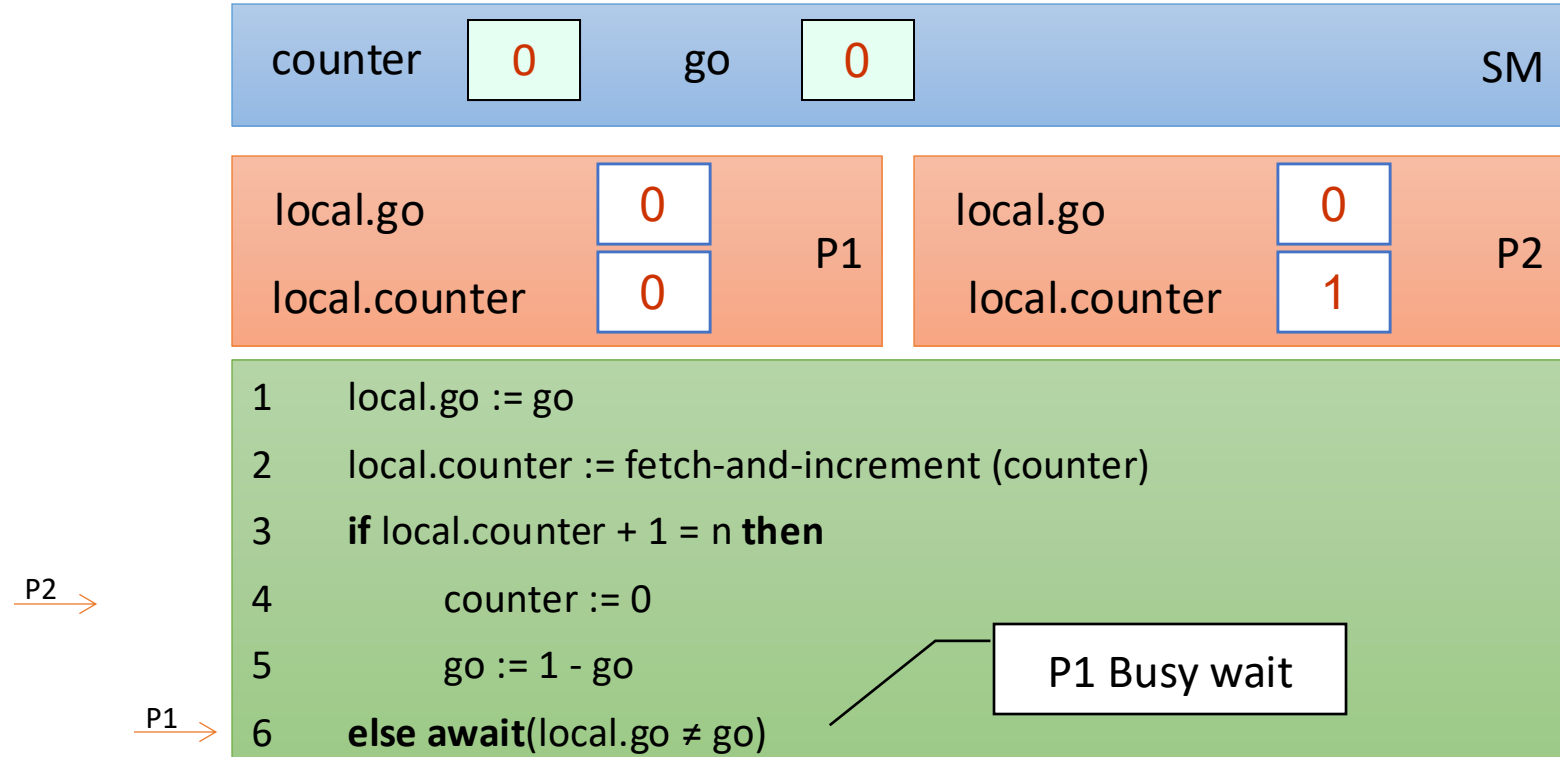
P2 → (points to line 4)

P1 → (points to line 6)

P1 Busy wait (points to the else branch of line 6)

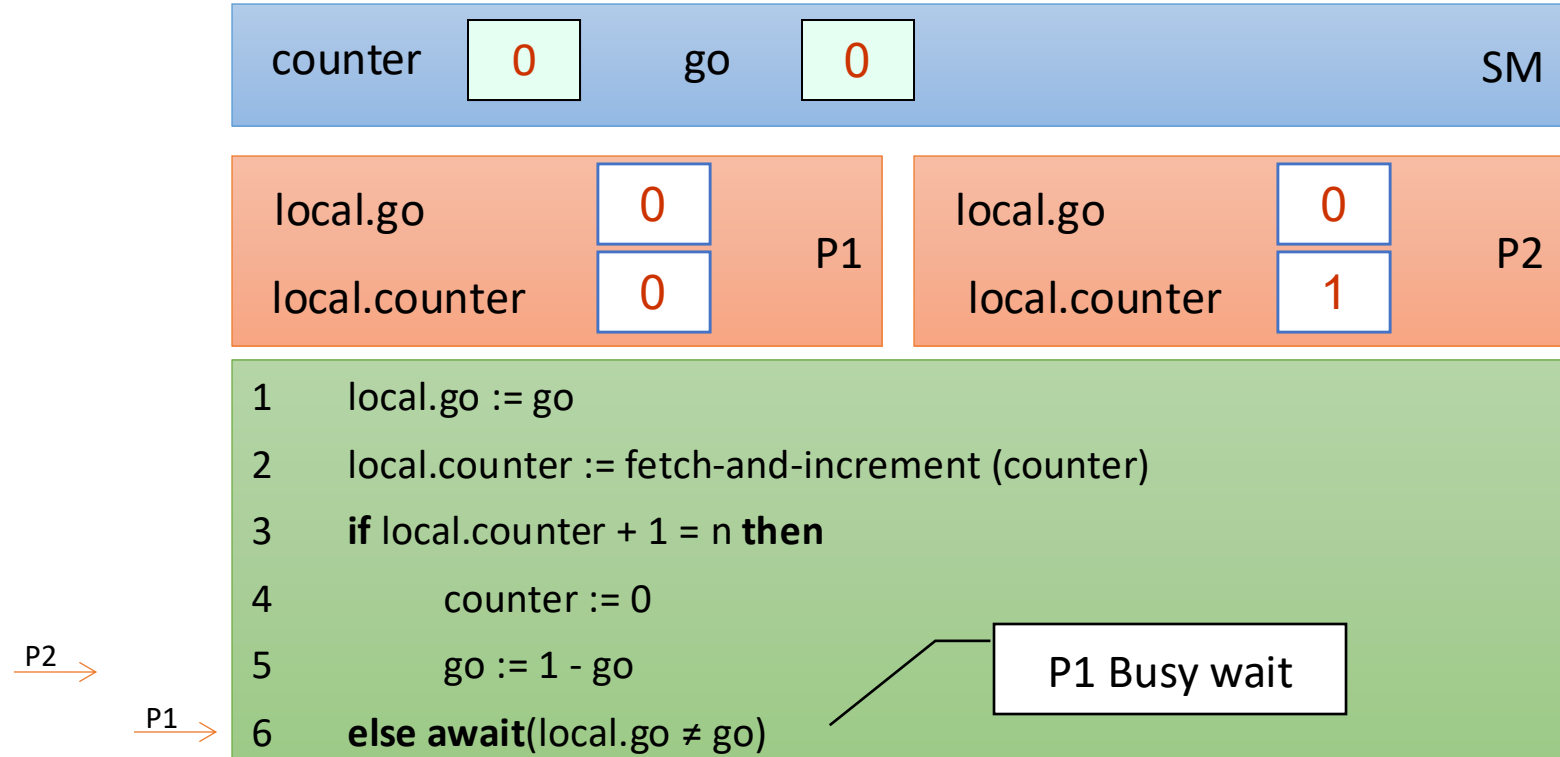
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



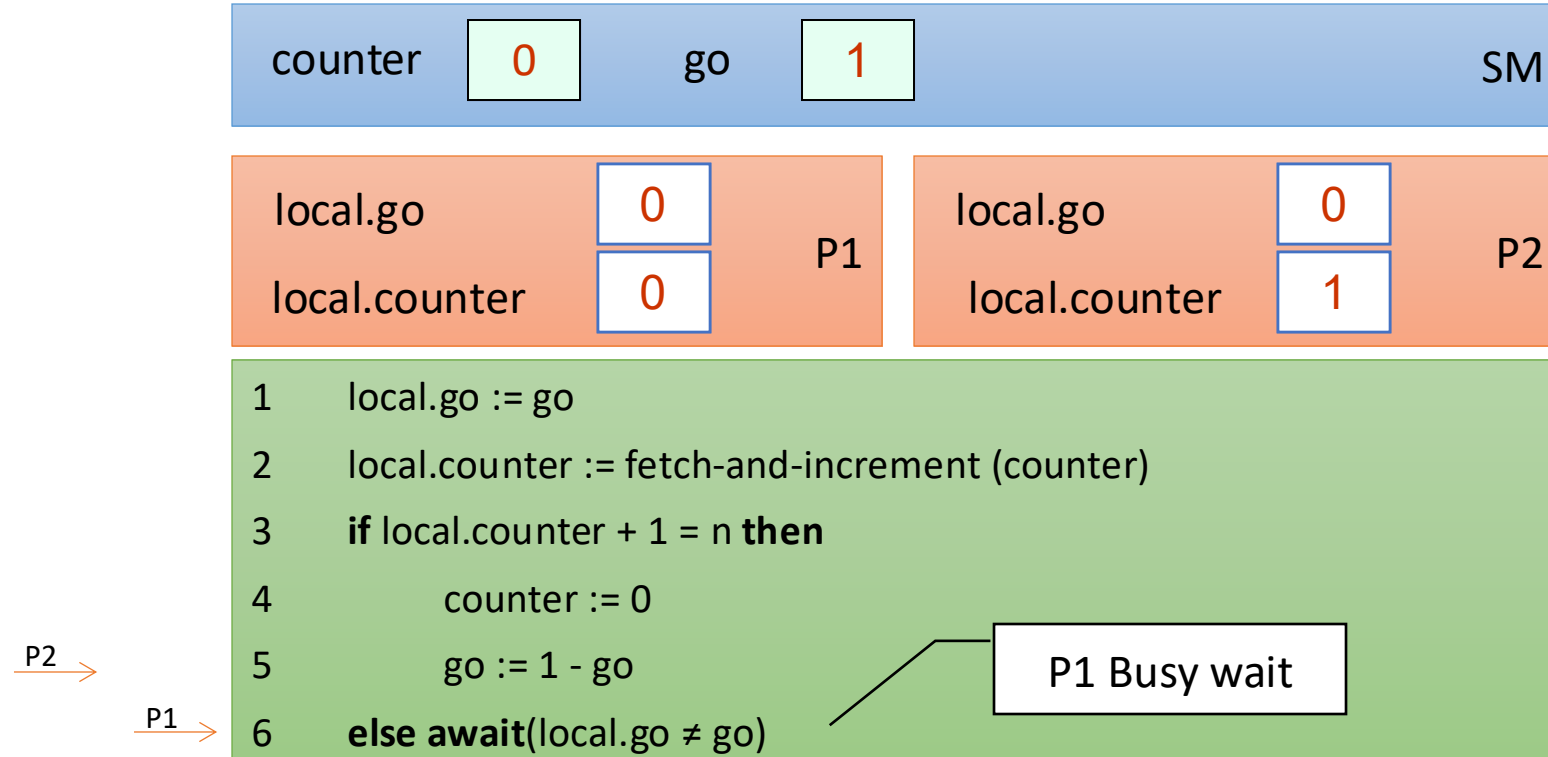
# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



# Simple Barrier Using an Atomic Counter



Run for n=2 Threads





# Simple Barrier Using an Atomic Counter

Run for n=2 Threads





```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```



# Simple Barrier Using an Atomic Counter

Run for n=2 Threads





```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

Pros/Cons?




# Simple Barrier Using an Atomic Counter

Run for n=2 Threads





```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

Pros/Cons?



# Simple Barrier Using an Atomic Counter

Run for n=2 Threads



```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```



Pros/Cons?



- There is high memory contention on go bit

# Simple Barrier Using an Atomic Counter

Run for  $n=2$  Threads



```
1 local.go := go
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     go := 1 - go
6 else await(local.go ≠ go)
```

Pros/Cons?



- There is high memory contention on *go* bit
- Reducing the contention:
  - Replace the *go* bit with  $n$  bits:  
 $go[1], \dots, go[n]$
  - Process  $p_i$  may spin only on the bit  $go[i]$

# A Local Spinning Counter Barrier

Program of a Thread  $i$

<b>shared</b>	counter: fetch and increment reg. – $\{0,..n\}$ , initially = 0
	go[1..n]: array of atomic bits, initial values are immaterial
<b>local</b>	local.go: a bit, initial value is immaterial
	local.counter: register

# A Local Spinning Counter Barrier



Program of a Thread  $i$

```
shared    counter: fetch and increment reg. – {0,..n}, initially = 0  
           go[1..n]: array of atomic bits, initial values are immaterial  
local    local.go: a bit, initial value is immaterial  
           local.counter: register
```

```
1  local.go := go[i]  
2  local.counter := fetch-and-increment (counter)  
3  if local.counter + 1 = n then  
4      counter := 0  
5      for j=1 to n { go[j] := 1 – go[j] }  
6  else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



# A Local Spinning Counter Barrier



Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier



Example Run for n=3 Threads



```
P1 → 1 local.go := go[i]
      2 local.counter := fetch-and-increment (counter)
      3 if local.counter + 1 = n then
      4     counter := 0
      5     for j=1 to n { go[j] := 1 - go[j] }
      6 else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier



Example Run for n=3 Threads



```
P1 → 1 local.go := go[i]
      2 local.counter := fetch-and-increment (counter)
      3 if local.counter + 1 = n then
      4     counter := 0
      5     for j=1 to n { go[j] := 1 - go[j] }
      6 else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1 →

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1 →

# A Local Spinning Counter Barrier



Example Run for n=3 Threads



```
1 local.go := go[i]
P1 → 2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier



Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
P1 → 3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P1 →

0+1≠3



# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1 →

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1 →

P1 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
P2 → 1 local.go := go[i]
      2 local.counter := fetch-and-increment (counter)
      3 if local.counter + 1 = n then
      4     counter := 0
      5     for j=1 to n { go[j] := 1 - go[j] }
P1 → 6 else await(local.go ≠ go[i])
```

P1 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
P2 → 1 local.go := go[i]
      2 local.counter := fetch-and-increment (counter)
      3 if local.counter + 1 = n then
      4     counter := 0
      5     for j=1 to n { go[j] := 1 - go[j] }
P1 → 6 else await(local.go ≠ go[i])
```

P1 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P2 → (points to line 2)

P1 → (points to line 6)

P1 Busy wait (points to the await statement in line 6)

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P2 →

P1 →

P1 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P2 → (points to line 2)

P1 → (points to line 6)

P1 Busy wait (points to the await statement in line 6)

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P2 →




P1 →

P1 Busy wait





# A Local Spinning Counter Barrier

Example Run for n=3 Threads



# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1,P2 Busy wait



# A Local Spinning Counter Barrier

Example Run for n=3 Threads



P3 →



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P1,P2 Busy wait

P2 → P1 →

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



P3 →

```



1  local.go := go[i]
2  local.counter := fetch-and-increment (counter)
3  if local.counter + 1 = n then
4      counter := 0
5      for j=1 to n { go[j] := 1 - go[j] }
6  else await(local.go ≠ go[i])
    
```

P1,P2 Busy wait

P2 → P1 →

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P3 →

P2 → P1 →

P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



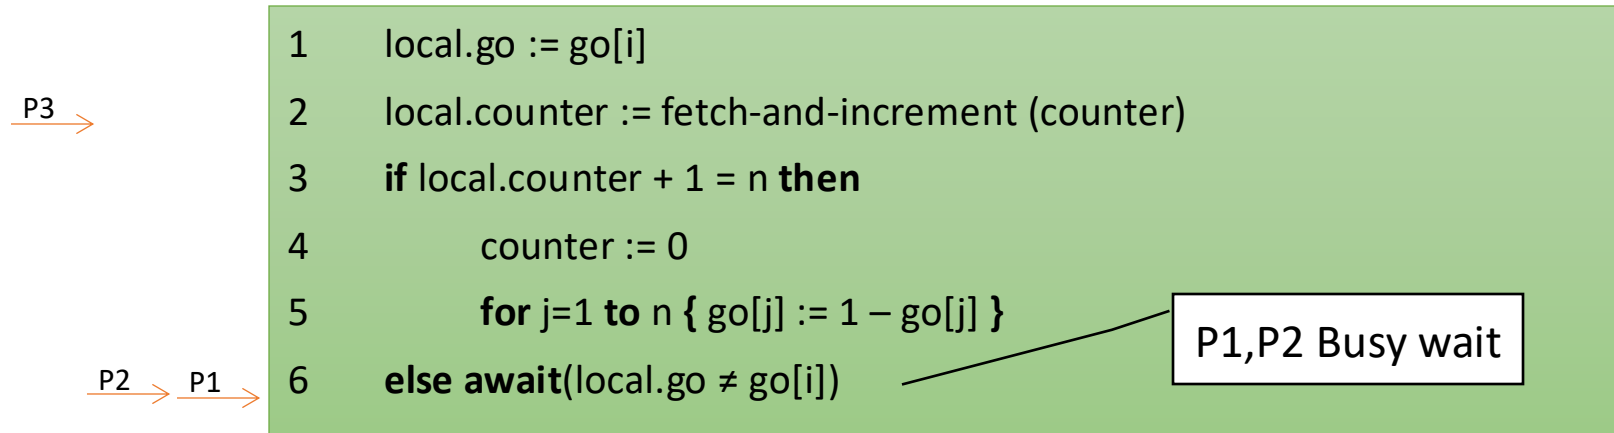


```

1  local.go := go[i]
2  local.counter := fetch-and-increment (counter)
3  if local.counter + 1 = n then
4      counter := 0
5      for j=1 to n { go[j] := 1 - go[j] }
6  else await(local.go ≠ go[i])
    
```

Annotations: P3 → (line 2), P2 → P1 → (line 6), P1,P2 Busy wait (pointing to line 6)



# A Local Spinning Counter Barrier

Example Run for n=3 Threads



# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```

P3 →

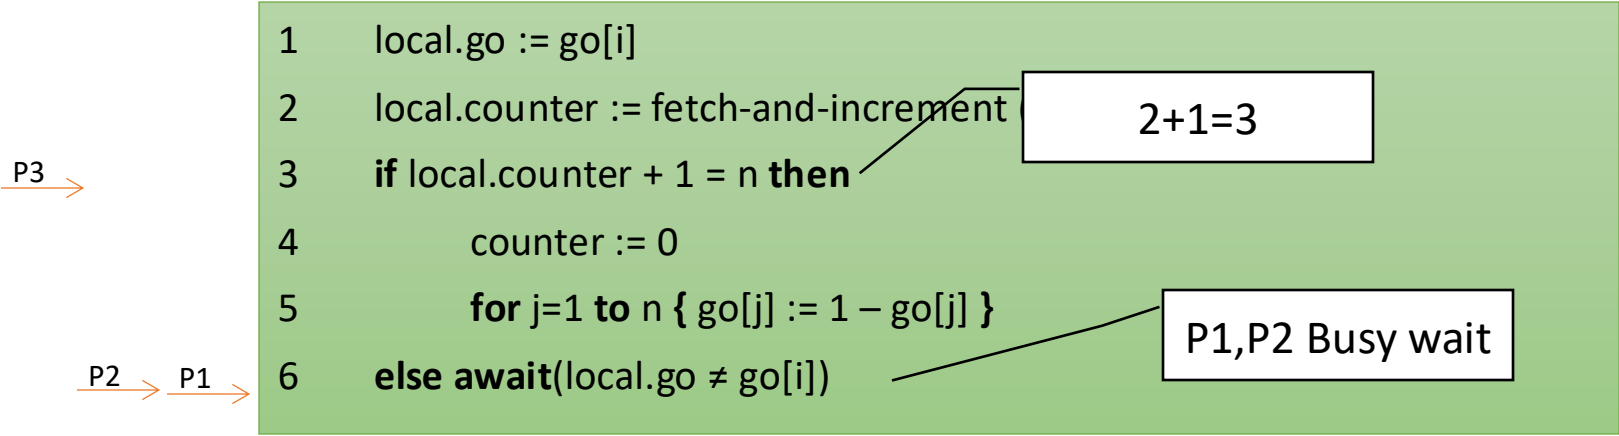


P2 → P1 →

P1,P2 Busy wait





# A Local Spinning Counter Barrier

Example Run for n=3 Threads



# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P3 →

P2 → P1 →

P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P3 →

P2 → P1 →

P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



P3 →

P2 → P1 →


P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```




P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```




P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```




P1,P2 Busy wait

# A Local Spinning Counter Barrier

Example Run for n=3 Threads





```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```





# A Local Spinning Counter Barrier

Example Run for n=3 Threads



```
1 local.go := go[i]
2 local.counter := fetch-and-increment (counter)
3 if local.counter + 1 = n then
4     counter := 0
5     for j=1 to n { go[j] := 1 - go[j] }
6 else await(local.go ≠ go[i])
```



Pros/Cons?  
*Does this actually reduce contention?*

# Comparison of counter-based Barriers

## Simple Barrier

- Pros:

- Cons:

## Simple Barrier with go array

- Pros:

- Cons:

# Comparison of counter-based Barriers

## Simple Barrier

- **Pros:**
  - Very Simple
  - Shared memory:  $O(\log n)$  *bits*
  - Takes  $O(1)$  until last waiting  $p$  is awoken
- **Cons:**
  - High contention on the go bit
  - Contention on the counter register (\*)


## Simple Barrier with go array

- **Pros:**
  - Low contention on the go array
  - In some models:
    - spinning is done on local memory
    - remote mem. ref.:  $O(1)$
- **Cons:**
  - Shared memory:  $O(n)$
  - Still contention on the counter register (\*)
  - Takes  $O(n)$  until last waiting  $p$  is awoken

# Tree Barriers




# A Tree-based Barrier




# A Tree-based Barrier

- Threads are organized in a binary tree
- Each node is owned by a predetermined thread




# A Tree-based Barrier

- Threads are organized in a binary tree
- Each node is owned by a predetermined thread
- Each thread waits until its 2 children arrive
  - combines results
  - passes them on to its parent




# A Tree-based Barrier

- Threads are organized in a binary tree
- Each node is owned by a predetermined thread
- Each thread waits until its 2 children arrive
  - combines results
  - passes them on to its parent
- Root learns that its 2 children have arrived → tells children they can go
- The signal propagates down the tree until all the threads get the message






# A Tree-based Barrier: indexing




# A Tree-based Barrier: indexing




Step 1: label numerically  
with traversal

# A Tree-based Barrier: indexing

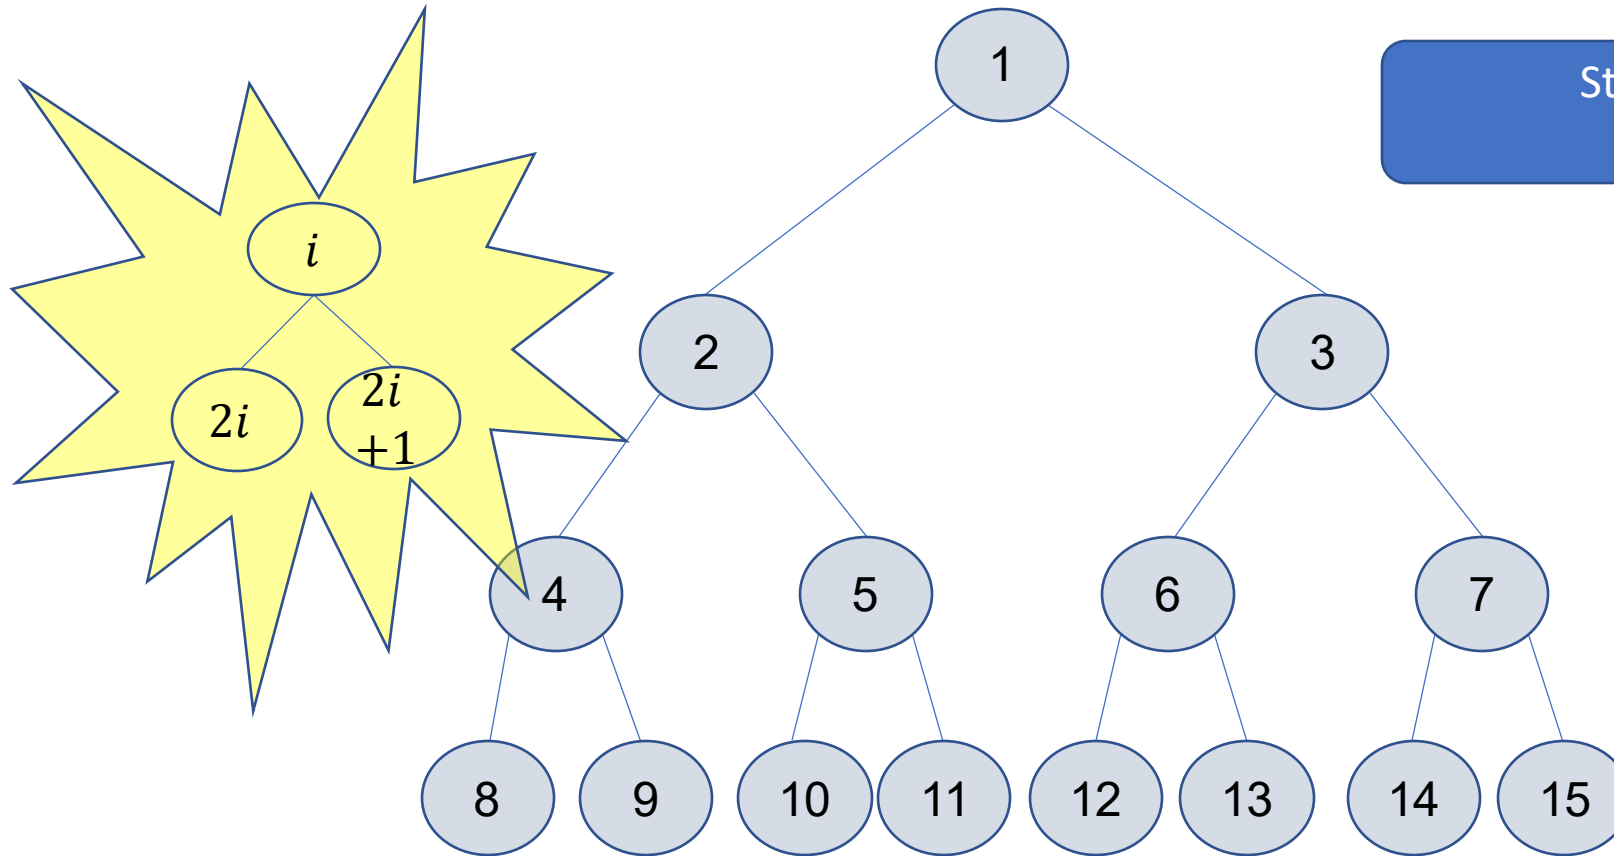


Step 1: label numerically with traversal

# A Tree-based Barrier: indexing

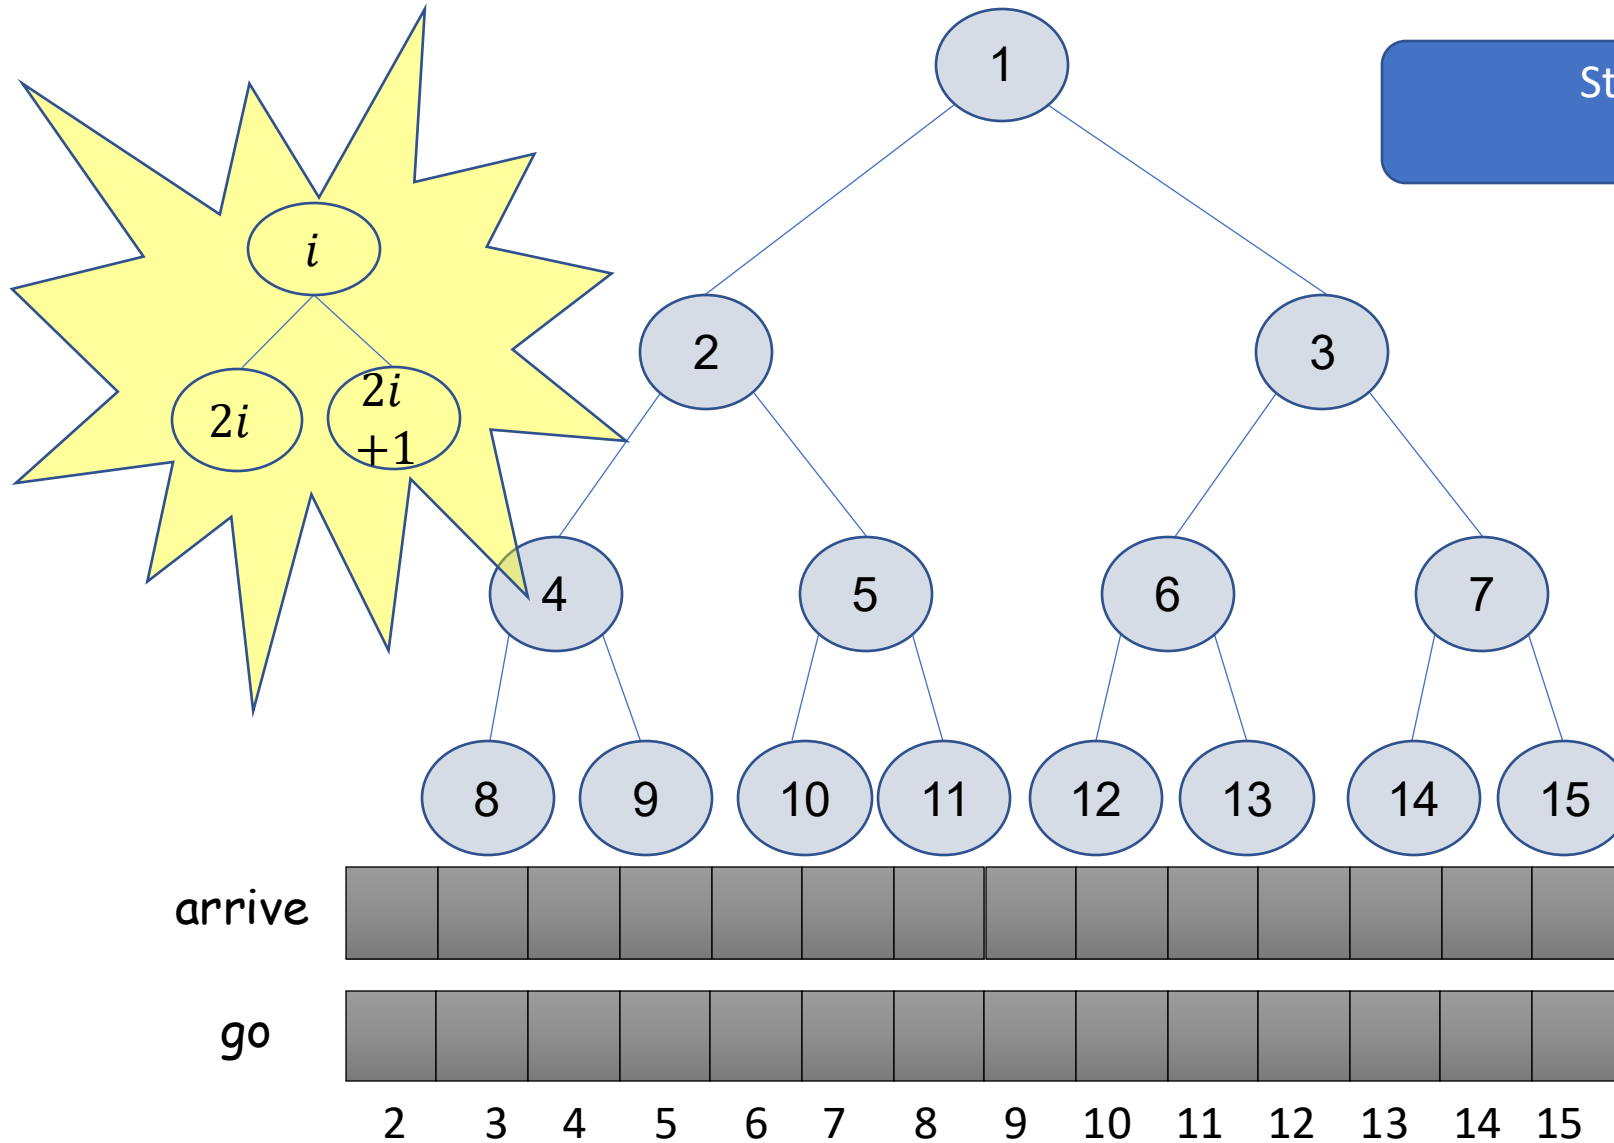


# A Tree-based Barrier: indexing



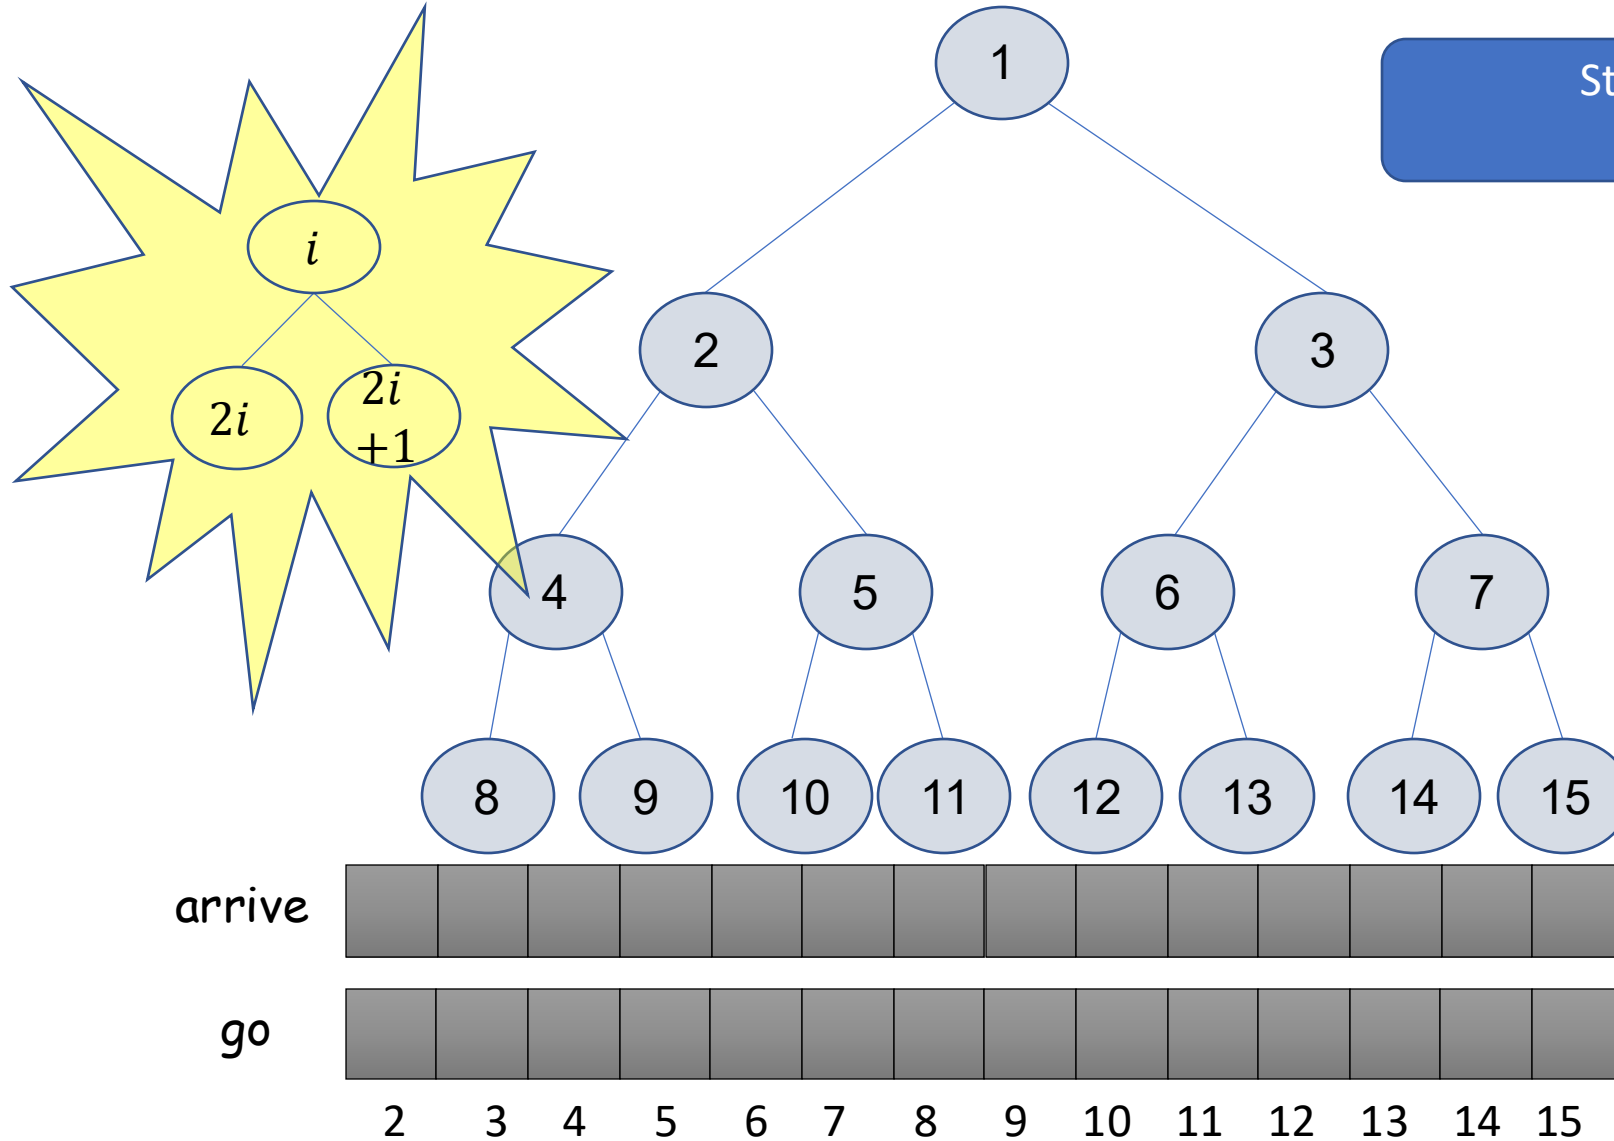
Step 1: label numerically with traversal

# A Tree-based Barrier: indexing



Step 1: label numerically with traversal

# A Tree-based Barrier: indexing



Indexing starts from 2  
Root  $\rightarrow$  1, doesn't need wait objects

# A Tree-based Barrier program of thread i

```
shared   arrive[2..n]: array of atomic bits, initial values = 0  
          go[2..n]: array of atomic bits, initial values = 0
```

```
1  if i=1 then                                     // root  
2      await(arrive[2] = 1); arrive[2] := 0  
3      await(arrive[3] = 1); arrive[3] := 0  
4      go[2] = 1; go[3] = 1  
5  else if i ≤ (n-1)/2 then                         // internal node  
6      await(arrive[2i] = 1); arrive[2i] := 0  
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0  
8      arrive[i] := 1  
9      await(go[i] = 1); go[i] := 0  
10     go[2i] = 1; go[2i+1] := 1  
11  else                                             // leaf  
12     arrive[i] := 1  
13     await(go[i] = 1); go[i] := 0 fi  
14  fi
```



# A Tree-based Barrier program of thread i

```
shared   arrive[2..n]: array of atomic bits, initial values = 0  
          go[2..n]: array of atomic bits, initial values = 0
```

```
1  if i=1 then                                     // root  
2      await(arrive[2] = 1); arrive[2] := 0  
3      await(arrive[3] = 1); arrive[3] := 0  
4      go[2] = 1; go[3] = 1  
5  else if i ≤ (n-1)/2 then                         // internal node  
6      await(arrive[2i] = 1); arrive[2i] := 0  
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0  
8      arrive[i] := 1  
9      await(go[i] = 1); go[i] := 0  
10     go[2i] = 1; go[2i+1] := 1  
11  else                                             // leaf  
12     arrive[i] := 1  
13     await(go[i] = 1); go[i] := 0 fi  
14  fi
```

Root

Internal

Leaf

# A Tree-based Barrier program of thread i

```
shared arrive[2..n]: array of atomic bits, initial values = 0  
go[2..n]: array of atomic bits, initial values = 0
```

```
1  if i=1 then                                     // root  
2      await(arrive[2] = 1); arrive[2] := 0  
3      await(arrive[3] = 1); arrive[3] := 0  
4      go[2] = 1; go[3] = 1  
5  else if i ≤ (n-1)/2 then                          // internal node  
6      await(arrive[2i] = 1); arrive[2i] := 0  
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0  
8      arrive[i] := 1  
9      await(go[i] = 1); go[i] := 0  
10     go[2i] = 1; go[2i+1] := 1  
11  else                                             // leaf  
12     arrive[i] := 1  
13     await(go[i] = 1); go[i] := 0 fi  
14  fi
```

Root

Internal

Leaf

Root:

- Wait for arriving children
- Tell children to go

# A Tree-based Barrier program of thread i

```
shared arrive[2..n]: array of atomic bits, initial values = 0  
go[2..n]: array of atomic bits, initial values = 0
```

```
1  if i=1 then                                     // root  
2      await(arrive[2] = 1); arrive[2] := 0  
3      await(arrive[3] = 1); arrive[3] := 0  
4      go[2] = 1; go[3] = 1  
5  else if i ≤ (n-1)/2 then                         // internal node  
6      await(arrive[2i] = 1); arrive[2i] := 0  
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0  
8      arrive[i] := 1  
9      await(go[i] = 1); go[i] := 0  
10     go[2i] = 1; go[2i+1] := 1  
11  else                                             // leaf  
12     arrive[i] := 1  
13     await(go[i] = 1); go[i] := 0 fi  
14  fi
```

Root

Internal

Leaf

Root:

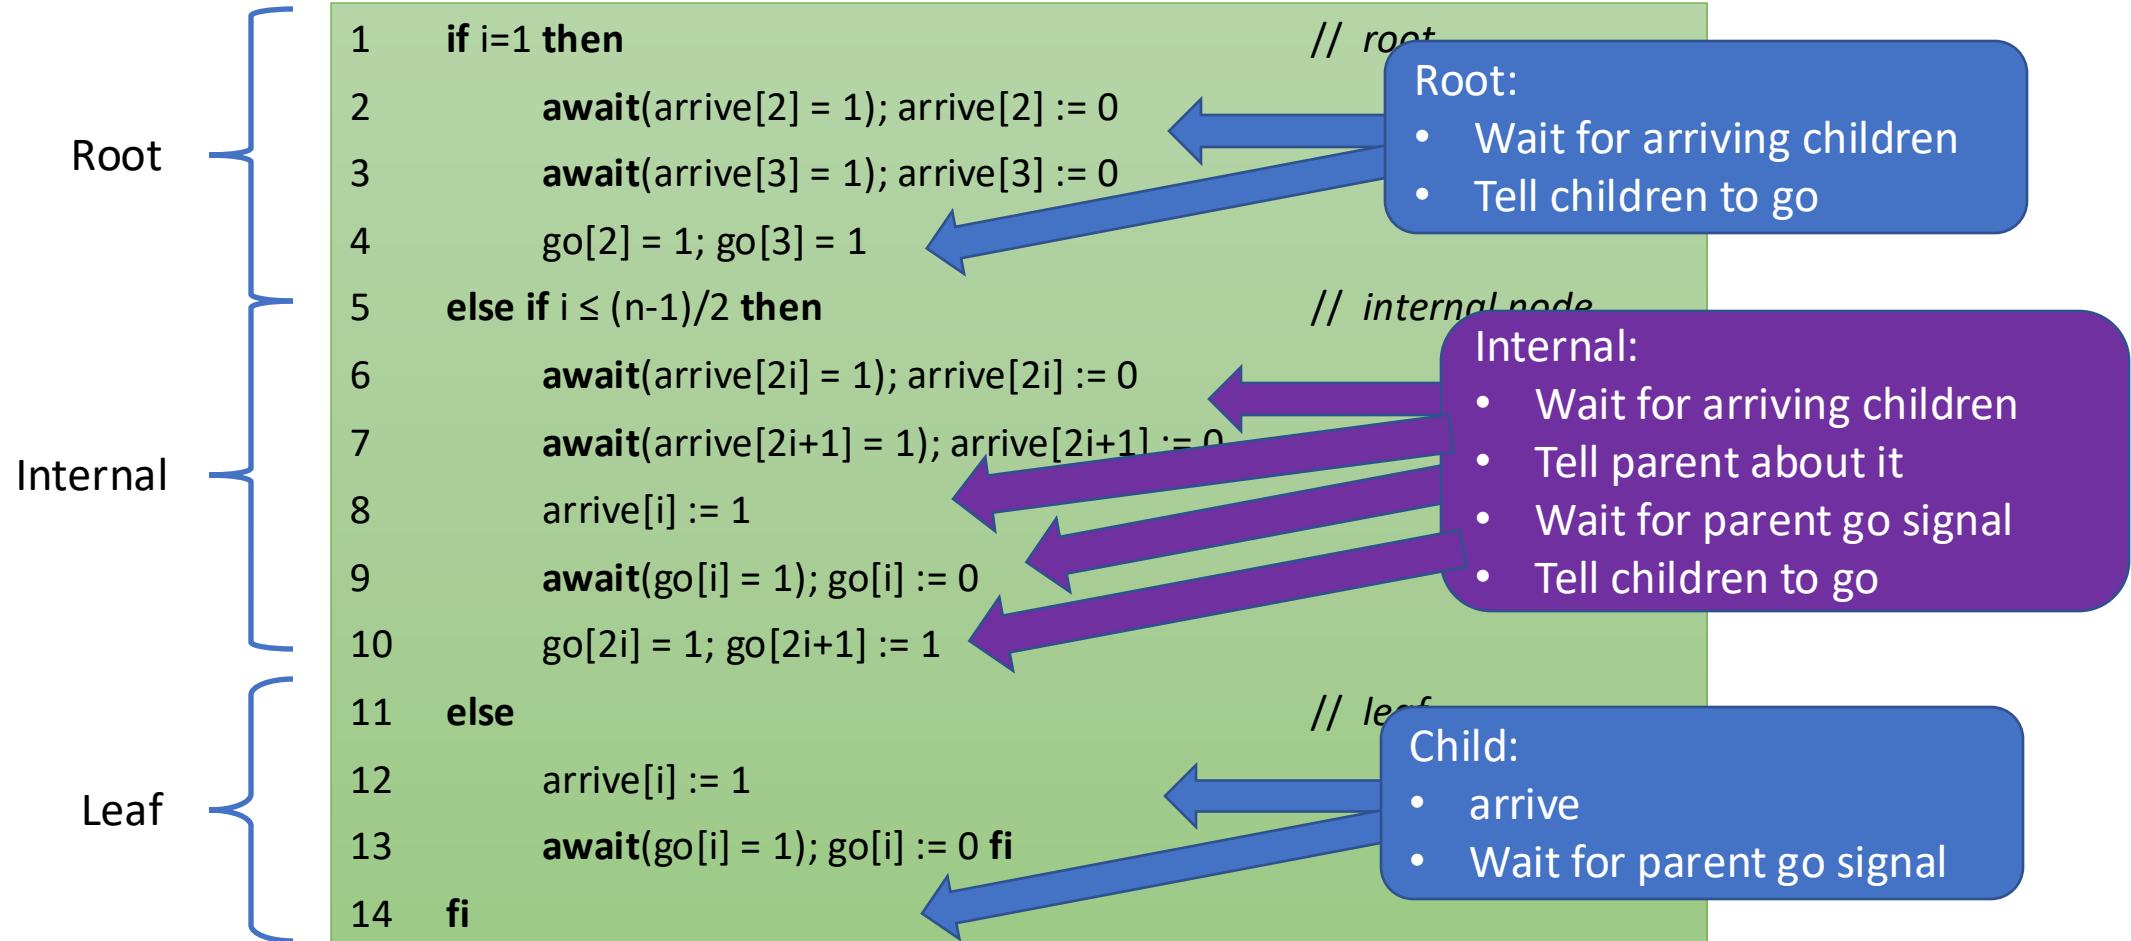
- Wait for arriving children
- Tell children to go

Internal:

- Wait for arriving children
- Tell parent about it
- Wait for parent go signal
- Tell children to go


# A Tree-based Barrier program of thread i

```
shared arrive[2..n]: array of atomic bits, initial values = 0  
go[2..n]: array of atomic bits, initial values = 0
```




# A Tree-based Barrier

## Example Run for n=7 threads





```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
```



# A Tree-based Barrier

## Example Run for n=7 threads






```

shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
```

# A Tree-based Barrier


## Example Run for n=7 threads



```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0


1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```





# A Tree-based Barrier

## Example Run for n=7 threads




Waiting for  
p<sub>4</sub> to arrive

```



shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads





```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




Waiting for  
go[5]

```



shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads





```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier


## Example Run for n=7 threads



```



shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads



```




shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads





```

shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads





```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```



```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads





```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads





```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```

```


1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```





# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```




```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
```

# A Tree-based Barrier


## Example Run for n=7 threads



```


shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```




# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```


```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```




# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```



```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads



```


shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier


## Example Run for n=7 threads



```


shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```


```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```




```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2      await(arrive[2] = 1); arrive[2] := 0
3      await(arrive[3] = 1); arrive[3] := 0
4      go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6      await(arrive[2i] = 1); arrive[2i] := 0
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0
8      arrive[i] := 1
9      await(go[i] = 1); go[i] := 0
10     go[2i] = 1; go[2i+1] := 1
11 else // leaf
12     arrive[i] := 1
13     await(go[i] = 1); go[i] := 0 fi
14 fi
    
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```


```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```


```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```


```

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
  
```




# A Tree-based Barrier

## Example Run for n=7 threads



```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```




```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
```





# A Tree-based Barrier

## Example Run for n=7 threads




```

shared  arrive[2..n]: array of atomic bits, initial values = 0
        go[2..n]: array of atomic bits, initial values = 0

1  if i=1 then // root
2      await(arrive[2] = 1); arrive[2] := 0
3      await(arrive[3] = 1); arrive[3] := 0
4      go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6      await(arrive[2i] = 1); arrive[2i] := 0
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0
8      arrive[i] := 1
9      await(go[i] = 1); go[i] := 0
10     go[2i] = 1; go[2i+1] := 1
11 else // leaf
12     arrive[i] := 1
13     await(go[i] = 1); go[i] := 0 fi
14 fi
  
```


# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```
1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
```



# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
  
```

```



1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



At this point all non-root threads in some await(go) case

# A Tree-based Barrier

## Example Run for n=7 threads




```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0

1 if i=1 then // root
2   await(arrive[2] = 1); arrive[2] := 0
3   await(arrive[3] = 1); arrive[3] := 0
4   go[2] = 1; go[3] = 1
5 else if i ≤ (n-1)/2 then // internal node
6   await(arrive[2i] = 1); arrive[2i] := 0
7   await(arrive[2i+1] = 1); arrive[2i+1] := 0
8   arrive[i] := 1
9   await(go[i] = 1); go[i] := 0
10  go[2i] = 1; go[2i+1] := 1
11 else // leaf
12  arrive[i] := 1
13  await(go[i] = 1); go[i] := 0 fi
14 fi
```

# A Tree-based Barrier

## Example Run for n=7 threads




```

shared arrive[2..n]: array of atomic bits, initial values = 0
      go[2..n]: array of atomic bits, initial values = 0
  
```


```

1  if i=1 then // root
2      await(arrive[2] = 1); arrive[2] := 0
3      await(arrive[3] = 1); arrive[3] := 0
4      go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6      await(arrive[2i] = 1); arrive[2i] := 0
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0
8      arrive[i] := 1
9      await(go[i] = 1); go[i] := 0
10     go[2i] = 1; go[2i+1] := 1
11 else // leaf
12     arrive[i] := 1
13     await(go[i] = 1); go[i] := 0 fi
14 fi
  
```



# A Tree-based Barrier

## Example Run for n=7 threads



shared arrive[2..n]: array of atomic bits, initial values = 0  
 go[2..n]: array of atomic bits, initial values = 0


```

1  if i=1 then // root
2    await(arrive[2] = 1); arrive[2] := 0
3    await(arrive[3] = 1); arrive[3] := 0
4    go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6    await(arrive[2i] = 1); arrive[2i] := 0
7    await(arrive[2i+1] = 1); arrive[2i+1] := 0
8    arrive[i] := 1
9    await(go[i] = 1); go[i] := 0
10   go[2i] = 1; go[2i+1] := 1
11 else // leaf
12   arrive[i] := 1
13   await(go[i] = 1); go[i] := 0 fi
14 fi
  
```

arrive	0	0	0	0	0	0
go	1	1	1	1	1	1
	2	3	4	5	6	7

# A Tree-based Barrier

## Example Run for n=7 threads



```
shared arrive[2..n]: array of atomic bits, initial values = 0
shared go[2..n]: array of atomic bits, initial values = 0
```

```

1  if i=1 then // root
2      await(arrive[2] = 1); arrive[2] := 0
3      await(arrive[3] = 1); arrive[3] := 0
4      go[2] = 1; go[3] = 1
5  else if i ≤ (n-1)/2 then // internal node
6      await(arrive[2i] = 1); arrive[2i] := 0
7      await(arrive[2i+1] = 1); arrive[2i+1] := 0
8      arrive[i] := 1
9      await(go[i] = 1); go[i] := 0
10     go[2i] = 1; go[2i+1] := 1
11 else // leaf
12     arrive[i] := 1
13     await(go[i] = 1); go[i] := 0 fi
14 fi

```

arrive	0	0	0	0	0	0
go	1	1	1	1	1	1
	2	3	4	5	6	7

# Tree Barrier Tradeoffs

- Pros:

- Cons:



# Tree Barrier Tradeoffs

- **Pros:**


- Low shared memory contention
  - No wait object is shared by more than 2 processes
  - Good for larger  $n$
- Fast – information from the root propagates after  $\log(n)$  steps
- Can use only atomic primitives (no special objects)
- On some models:
  - each process spins on a locally accessible bit
  - # (remote memory ref.) =  $O(1)$  per process

- **Cons:**


- Shared memory space complexity –  $O(n)$
- Asymmetric – all the processes don't do the same amount of work
- Corner cases for  $n \neq 2^k - 1$

# Butterfly Barrier

# Butterfly Barrier

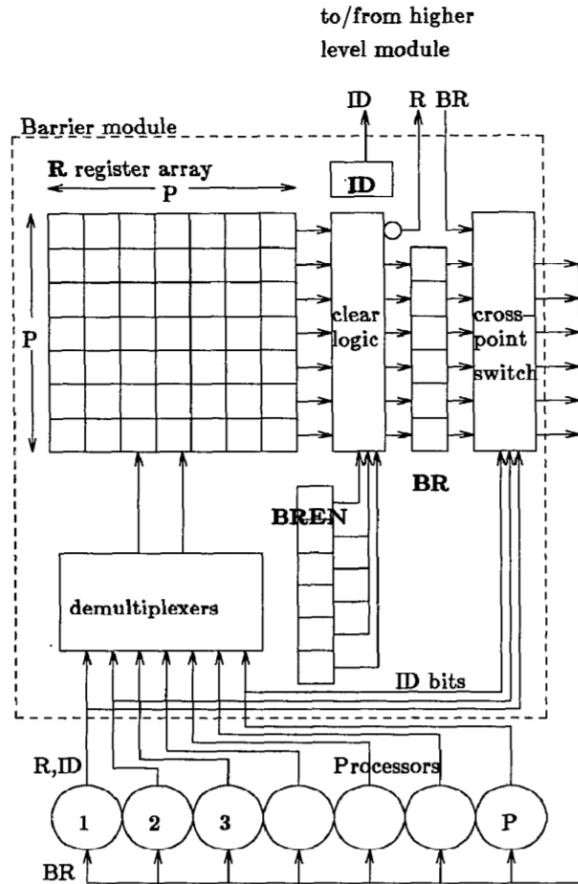


# Butterfly Barrier



- When would this be preferable?

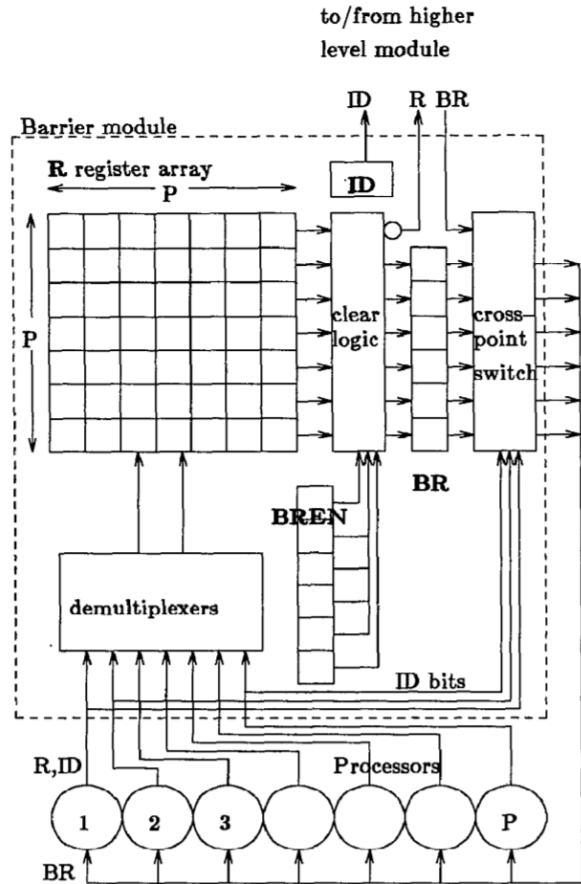
# Hardware Supported Barriers



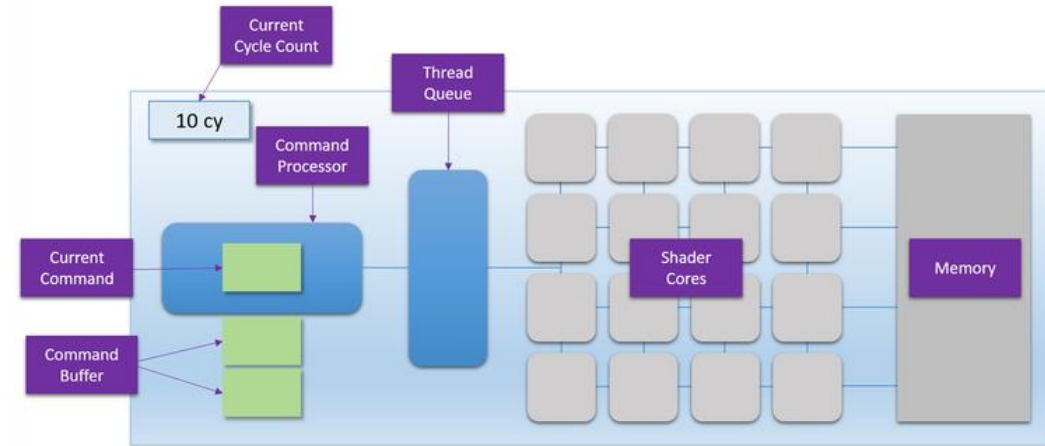
CPU



# Hardware Supported Barriers



CPU



GPU

- When would this be useful?

# Barriers Summary

## Seen:

- Semaphore-based barrier
- Simple barrier
  - Based on atomic fetch-and-increment counter
- Local spinning barrier
  - Based on atomic fetch-and-increment counter and go array


- Tree-based barrier

## Not seen:

- Test-and-Set barriers
  - Based on test-and-test-and-set objects
  - One version without memory initialization
- See-Saw barrier
- Book has condition barriers



# Asynchronous Programming Events, Promises, and Futures




# Programming Models for Concurrency

# Programming Models for Concurrency

- Hardware execution model:



# Programming Models for Concurrency

- Hardware execution model:
  - CPU(s) execute instructions sequentially





# Programming Models for Concurrency

- Hardware execution model:
  - CPU(s) execute instructions sequentially
- Programming model dimensions:
  - How to specify computation
  - How to specify communication
  - How to specify coordination/control transfer





# Programming Models for Concurrency

- Hardware execution model:
  - CPU(s) execute instructions sequentially
- Programming model dimensions:
  - How to specify computation
  - How to specify communication
  - How to specify coordination/control transfer
- Techniques/primitives
  - Message passing vs shared memory
  - Preemption vs Non-preemption



# Programming Models for Concurrency



- Hardware execution model:
  - CPU(s) execute instructions sequentially
- Programming model dimensions:
  - How to specify computation
  - How to specify communication
  - How to specify coordination/control transfer
- Techniques/primitives
  - Message passing vs shared memory
  - Preemption vs Non-preemption
- Dimensions/techniques not always orthogonal



# Programming Models for Concurrency

- Hardware execution model:
  - CPU(s) execute instructions sequentially
- Programming model dimensions:
  - How to specify computation
  - How to specify communication
  - How to specify coordination/control transfer
- Techniques/primitives
  - Message passing vs shared memory
  - Preemption vs Non-preemption
- Dimensions/techniques not always orthogonal

*Futures & Promises touch all three dimension*





# Futures & Promises

# Futures & Promises

- *Values that will eventually become available*

# Futures & Promises

- Values *that will eventually become available*
- Time-dependent states:
  - **Completed/determined**
    - Computation complete, value concrete
  - **Incomplete/undetermined**
    - Computation not complete yet

# Futures & Promises


- *Values that will eventually become available*
- Time-dependent states:
  - **Completed/determined**
    - Computation complete, value concrete
  - **Incomplete/undetermined**
    - Computation not complete yet
- Construct ( future X )
  - immediately returns value
  - concurrently executes X

# Java Example

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

# Java Example

```
1 static void runAsyncExample() {  
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {  
3         assertTrue(Thread.currentThread().isDaemon());  
4         randomSleep();  
5     });  
6     assertFalse(cf.isDone());  
7     sleepEnough();  
8     assertTrue(cf.isDone());  
9 }
```



# Java Example

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

- CompletableFuture is a container for Future object type

# Java Example

```
1 static void runAsyncExample() {  
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {  
3         assertTrue(Thread.currentThread().isDaemon());  
4         randomSleep();  
5     });  
6     assertFalse(cf.isDone());  
7     sleepEnough();  
8     assertTrue(cf.isDone());  
9 }
```

- CompletableFuture is a container for Future object type
- cf is an instance



# Java Example

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

- CompletableFuture is a container for Future object type
- cf is an instance
- runAsync() accepts
  - Lambda expression
  - Anonymous function
  - Functor

# Java Example

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

- CompletableFuture is a container for Future object type
- cf is an instance
- runAsync() accepts
  - Lambda expression
  - Anonymous function
  - Functor
- runAsync() immediately returns a waitable object (cf)

# Java Example

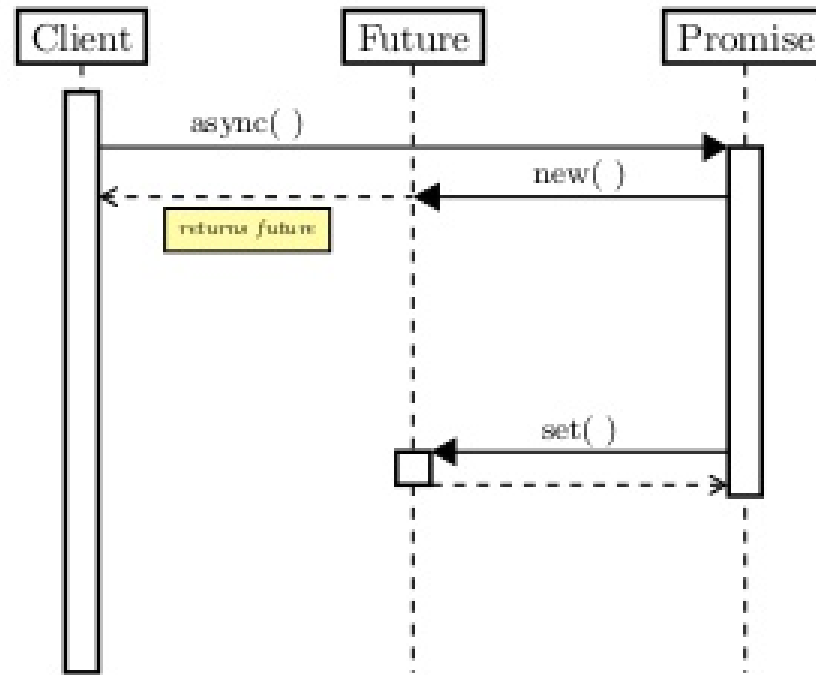
```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

- CompletableFuture is a container for Future object type
- cf is an instance
- runAsync() accepts
  - Lambda expression
  - Anonymous function
  - Functor
- runAsync() immediately returns a waitable object (cf)
- Where (on what thread) does the lambda expression run?

# Futures and Promises:

Why two kinds of objects?

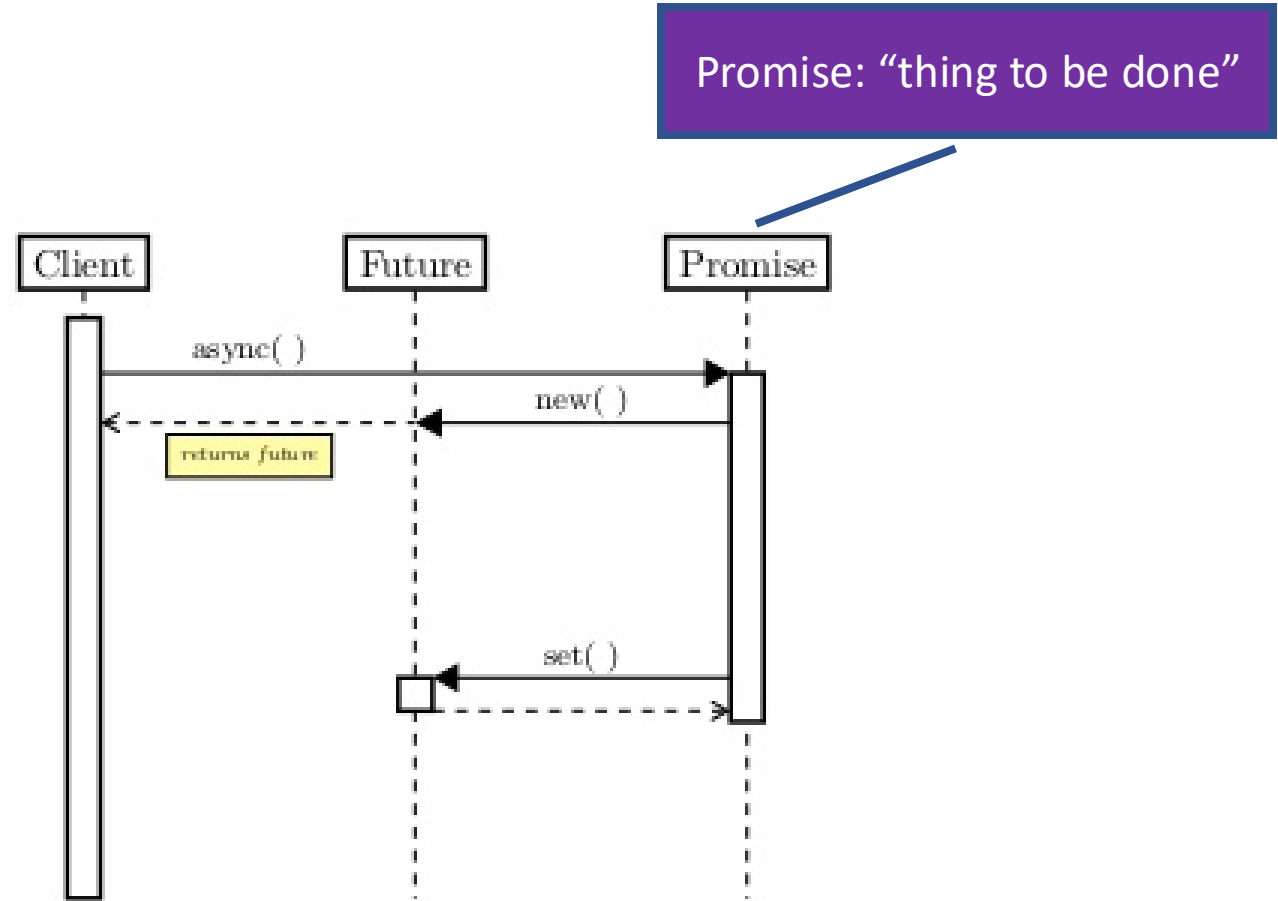
```
future<int> f1 = async(foo1);  
...  
int result = f1.get();
```



# Futures and Promises:

Why two kinds of objects?

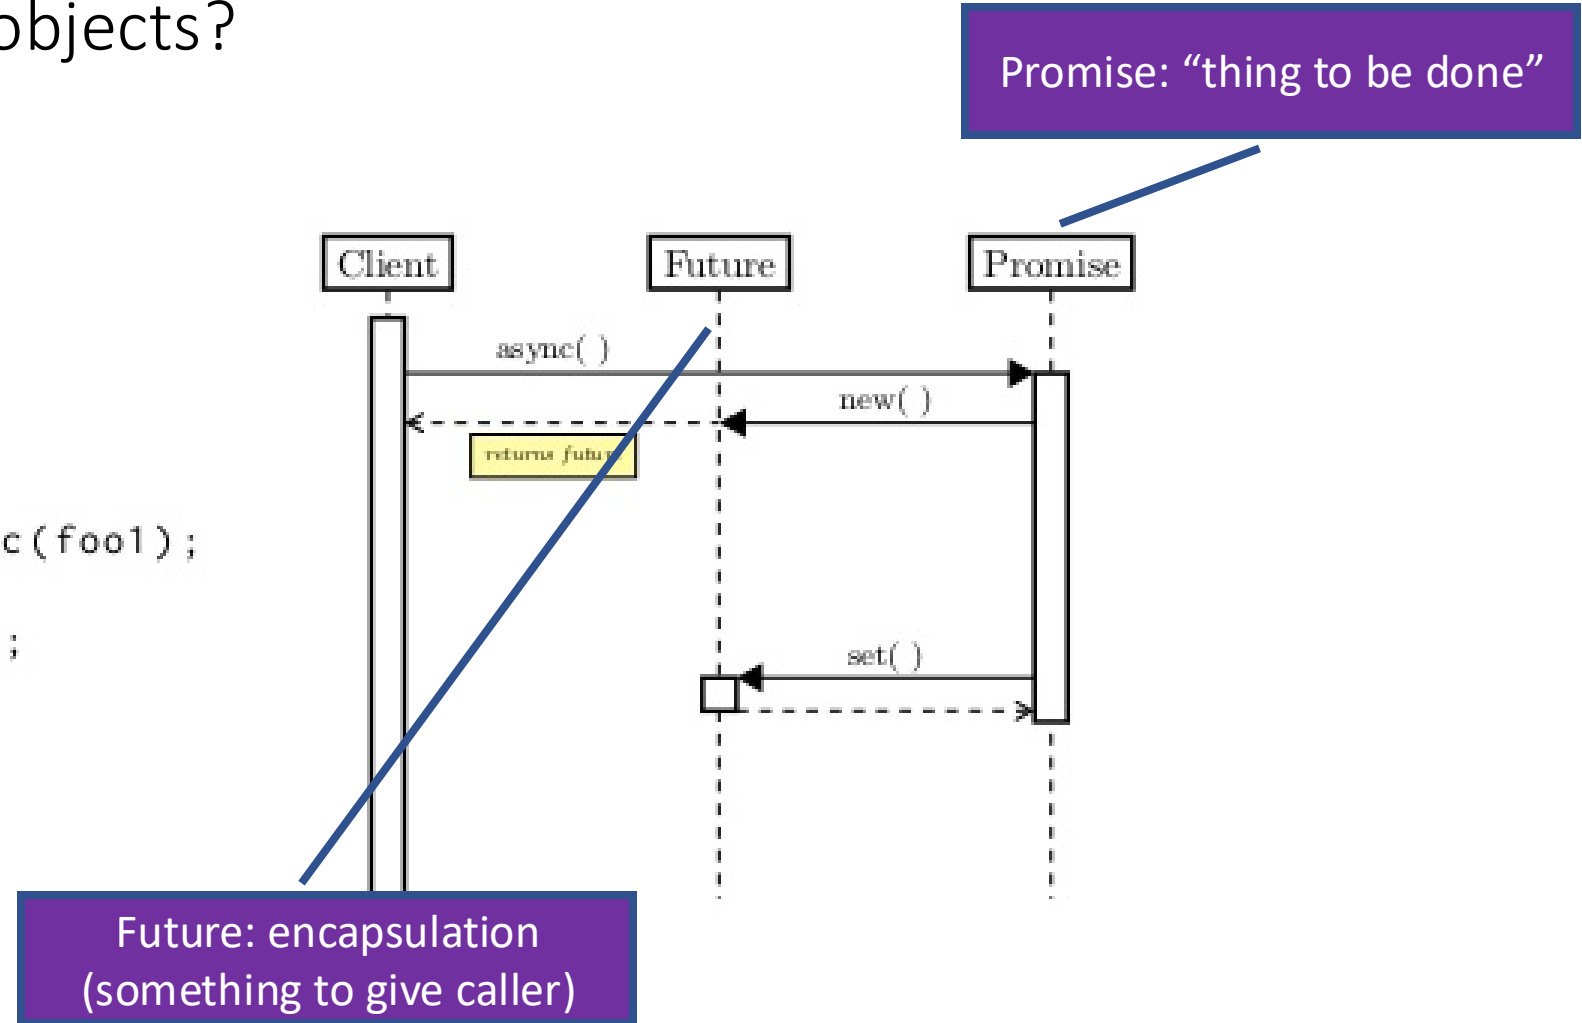
```
future<int> f1 = async(foo1);  
...  
int result = f1.get();
```



# Futures and Promises:

Why two kinds of objects?

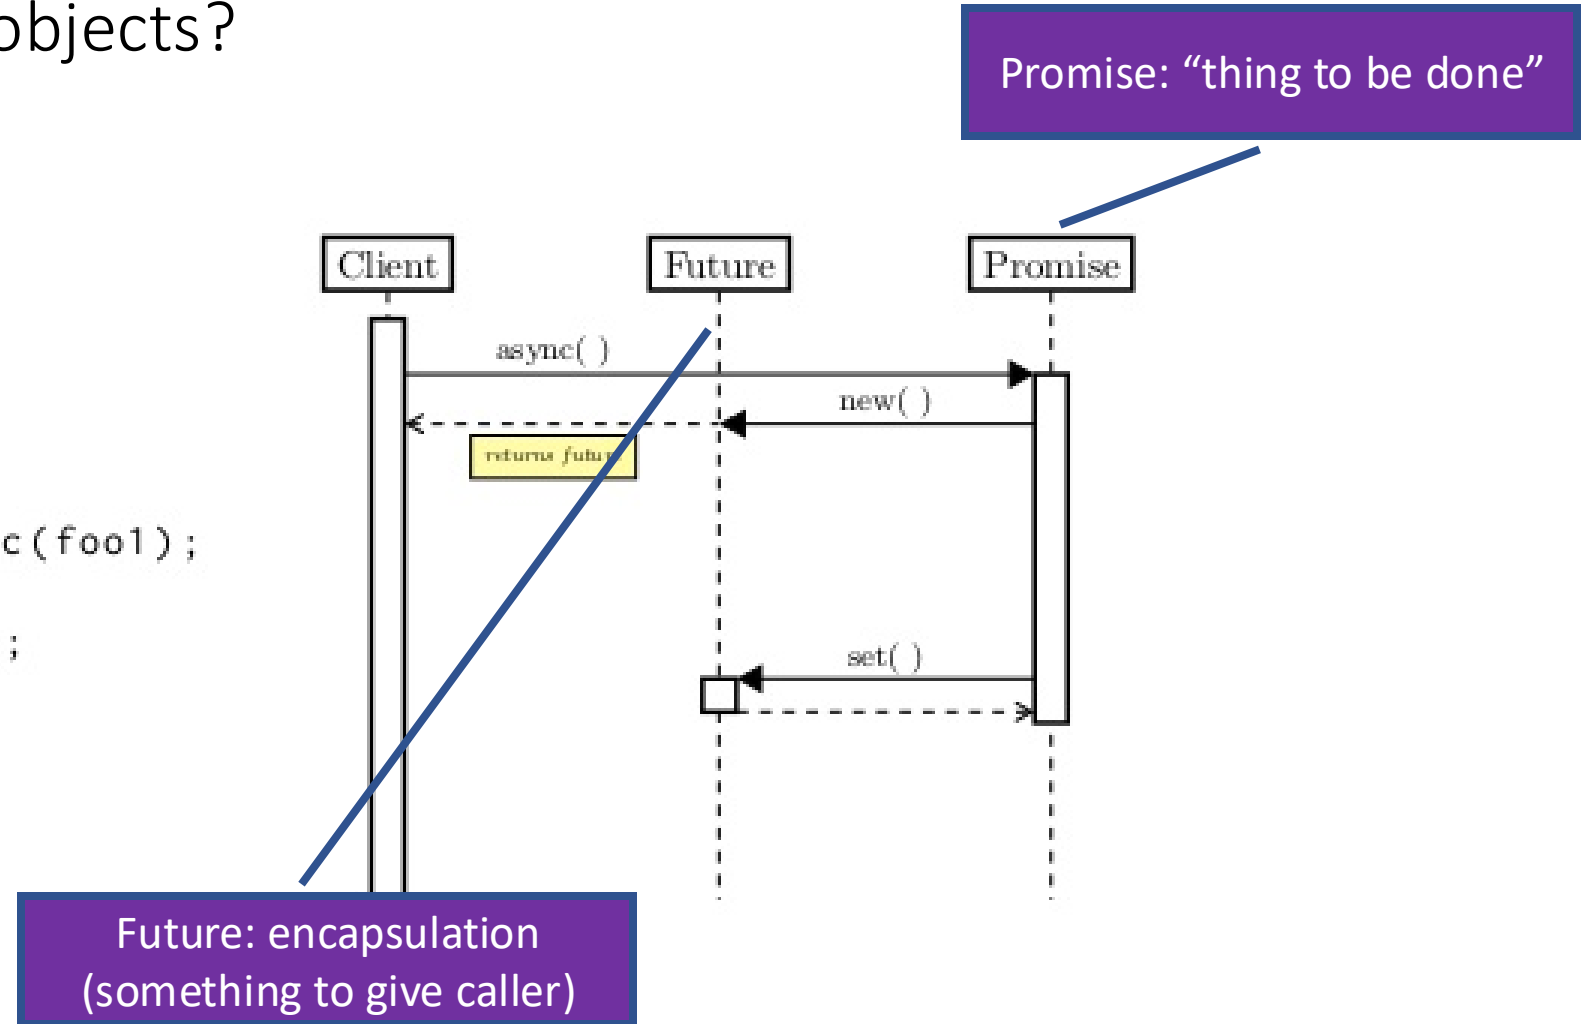
```
future<int> f1 = async(foo1);  
...  
int result = f1.get();
```



# Futures and Promises:

Why two kinds of objects?


```
future<int> f1 = async(foo1);  
...  
int result = f1.get();
```



**Promise to do** something in the future

# Futures vs Promises

- **Future:** read-only reference to uncompleted value
- **Promise:** single-assignment variable that the future refers to
- Promises *complete* the future with:
  - Result with success/failure
  - Exception






# Futures vs Promises

- **Future:** read-only reference to uncompleted value
- **Promise:** single-assignment variable that the future refers to
- Promises *complete* the future with:
  - Result with success/failure
  - Exception


Language	Promise	Future
Algol	Thunk	Address of async result
Java	Future<T>	CompletableFuture<T>
C#/.NET	TaskCompletionSource<T>	Task<T>
JavaScript	Deferred	Promise
C++	std::promise	std::future



# Futures vs Promises

- **Future:** read-only reference to uncompleted value
- **Promise:** single-assignment variable that the future refers to
- Promises *complete* the future with:
  - Result with success/failure
  - Exception

Language	Promise	Future
Algol	Thunk	Address of async result
Java	Future<T>	CompletableFuture<T>
C#/.NET	TaskCompletionSource<T>	Task<T>
JavaScript	Deferred	Promise
C++	std::promise	std::future




# Futures vs Promises

Mnemonic:  
Promise to *do* something  
Make a promise *for* the future

- **Future:** read-only reference to uncompleted value
- **Promise:** single-assignment variable that the future refers to
- Promises *complete* the future with:
  - Result with success/failure
  - Exception

Language	Promise	Future
Algol	Thunk	Address of async result
Java	Future<T>	CompletableFuture<T>
C#/.NET	TaskCompletionSource<T>	Task<T>
JavaScript	Deferred	Promise
C++	std::promise	std::future



# Putting Futures in Context

My unvarnished opinion

# Putting Futures in Context

My unvarnished opinion

Futures:

# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*

# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

# Putting Futures in Context

## My unvarnished opinion

### Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```



# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

Compromise Programming Model between:

# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

Compromise Programming Model between:

- Event-based programming

# Putting Futures in Context

My unvarnished opinion

Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

Compromise Programming Model between:

- Event-based programming
- Thread-based programming

# Putting Futures in Context

My unvarnished opinion

Futures:


- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative/sequential
  - Threads of control peppered with asynchronous/concurrent tasks

Compromise Programming Model between:

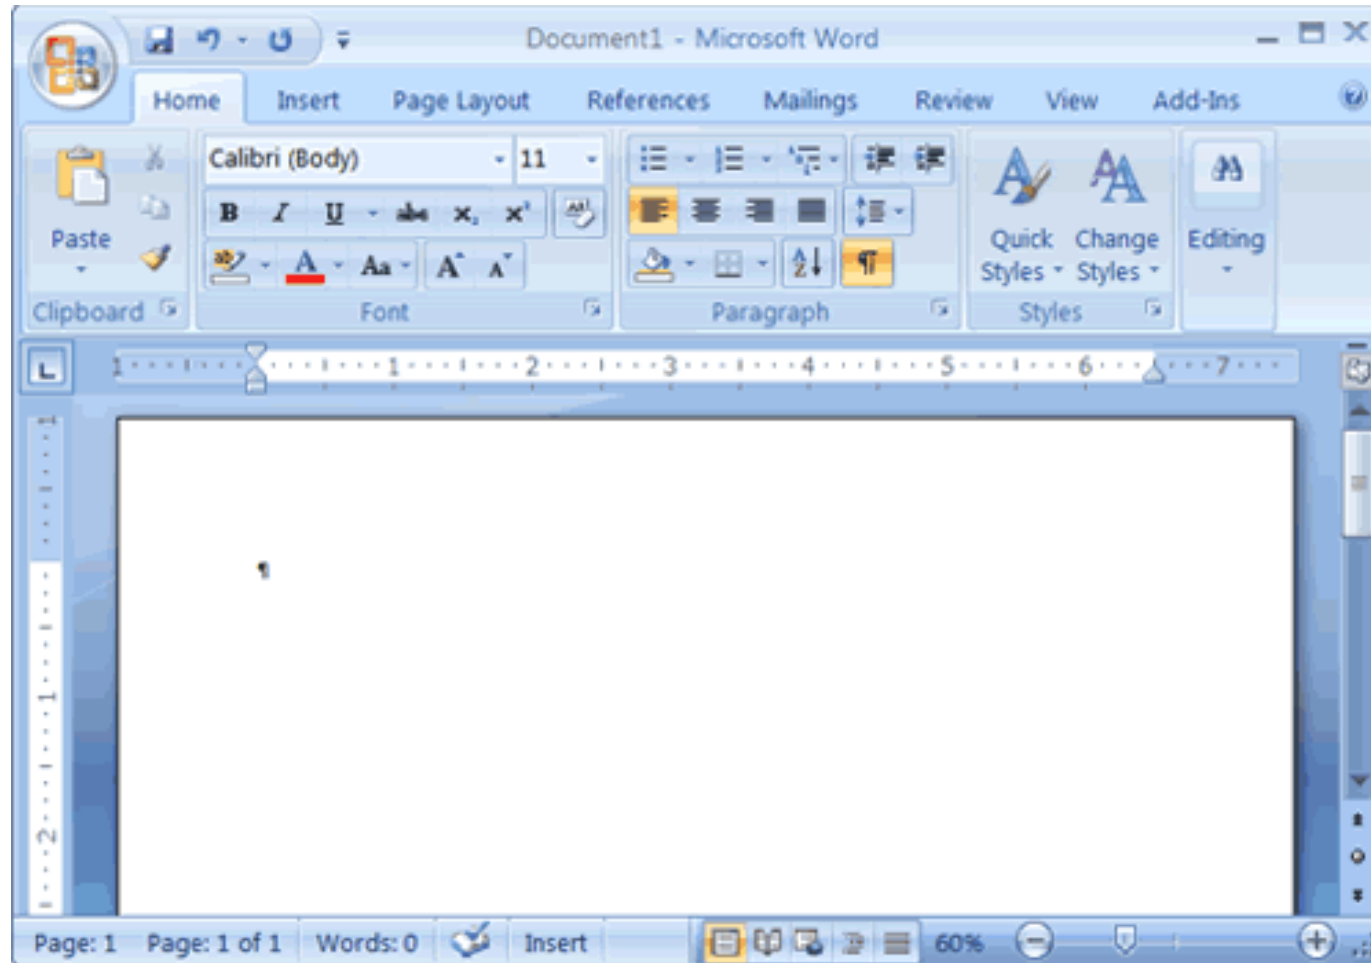
- Event-based programming
- Thread-based programming

Events vs. Threads!

# GUI Programming



# GUI Programming



```
do {  
    WaitForSomething();  
    RespondToThing();  
} until (forever);
```

# GUI Programming

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

    //Step 1: Registering the Window Class
    wc.cbSize      = sizeof(WNDCLASSEX);
    wc.style       = 0;
    wc.lpfnWndProc = WndProc;
    wc.cbClsExtra  = 0;
    wc.cbWndExtra  = 0;
    wc.hInstance   = hInstance;
    wc.hIcon       = LoadIcon(NULL, IDI_APPLICATION);
    wc.hCursor     = LoadCursor(NULL, IDC_ARROW);
    wc.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wc.lpszMenuName = NULL;
    wc.lpszClassName = g_szClassName;
    wc.hIconSm     = LoadIcon(NULL, IDI_APPLICATION);


    if(!RegisterClassEx(&wc))
    {
        MessageBox(NULL, "Window Registration Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    // Step 2: Creating the Window
    hwnd = CreateWindowEx(
        WS_EX_CLIENTEDGE,
        g_szClassName,
        "The title of my window",
        WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
        NULL, NULL, hInstance, NULL);

    if(hwnd == NULL)
    {
        MessageBox(NULL, "Window Creation Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    ShowWindow(hwnd, nCmdShow);
    UpdateWindow(hwnd);

    // Step 3: The Message Loop
    while(GetMessage(&Msg, NULL, 0, 0) > 0)
    {
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```





# GUI Programming

```
// Step 2: Creating the Window
hwnd = CreateWindowEx(
    WS_EX_CLIENTEDGE,
    g_szClassName,
    "The title of my window",
    WS_OVERLAPPEDWINDOW,
    CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
    NULL, NULL, hInstance, NULL);
```

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

    //Step 1: Registering the Window Class
    wc.cbSize      = sizeof(WNDCLASSEX);
    wc.style       = 0;
    wc.lpfnWndProc = WndProc;
    wc.cbClsExtra  = 0;
    wc.cbWndExtra  = 0;
    wc.hInstance   = hInstance;
    wc.hIcon       = LoadIcon(NULL, IDI_APPLICATION);
    wc.hCursor     = LoadCursor(NULL, IDC_ARROW);
    wc.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wc.lpszMenuName = NULL;
    wc.lpszClassName = g_szClassName;
    wc.hIconSm     = LoadIcon(NULL, IDI_APPLICATION);

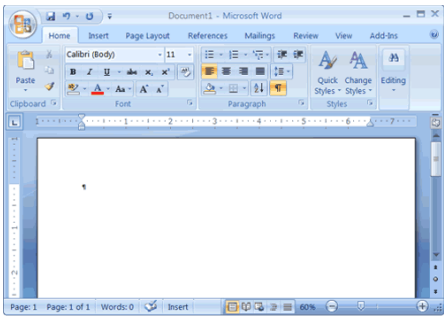

    if(!RegisterClassEx(&wc))
    {
        MessageBox(NULL, "Window Registration Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    // Step 2: Creating the Window
    hwnd = CreateWindowEx(
        WS_EX_CLIENTEDGE,
        g_szClassName,
        "The title of my window",
        WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
        NULL, NULL, hInstance, NULL);

    if(hwnd == NULL)
    {
        MessageBox(NULL, "Window Creation Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    ShowWindow(hwnd, nCmdShow);
    UpdateWindow(hwnd);

    // Step 3: The Message Loop
    while(GetMessage(&Msg, NULL, 0, 0) > 0)
    {
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```



# GUI Programming

```
// Step 2: Creating the Window
hwnd = CreateWindowEx(
    WS_EX_CLIENTEDGE,
    g_szClassName,
    "The title of my window",
    WS_OVERLAPPEDWINDOW,
    CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
    NULL, NULL, hInstance, NULL);
```

```
// Step 3: The Message Loop
while(GetMessage(&Msg, NULL, 0, 0) > 0)
{
    TranslateMessage(&Msg);
    DispatchMessage(&Msg);
}
```

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

    //Step 1: Registering the Window Class
    wc.cbSize      = sizeof(WNDCLASSEX);
    wc.style       = 0;
    wc.lpfnWndProc = WndProc;
    wc.cbClsExtra  = 0;
    wc.cbWndExtra  = 0;
    wc.hInstance   = hInstance;
    wc.hIcon       = LoadIcon(NULL, IDI_APPLICATION);
    wc.hCursor     = LoadCursor(NULL, IDC_ARROW);
    wc.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wc.lpszMenuName = NULL;
    wc.lpszClassName = g_szClassName;
    wc.hIconSm     = LoadIcon(NULL, IDI_APPLICATION);


    if(!RegisterClassEx(&wc))
    {
        MessageBox(NULL, "Window Registration Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    // Step 2: Creating the Window
    hwnd = CreateWindowEx(
        WS_EX_CLIENTEDGE,
        g_szClassName,
        "The title of my window",
        WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
        NULL, NULL, hInstance, NULL);

    if(hwnd == NULL)
    {
        MessageBox(NULL, "Window Creation Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    ShowWindow(hwnd, nCmdShow);
    UpdateWindow(hwnd);

    // Step 3: The Message Loop
    while(GetMessage(&Msg, NULL, 0, 0) > 0)
    {
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```



# GUI Programming

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

    //Step 1: Registering the Window Class
    wc.cbSize      = sizeof(WNDCLASSEX);
    wc.style       = 0;
    wc.lpfnWndProc = WndProc;
    wc.cbClsExtra  = 0;
    wc.cbWndExtra  = 0;
    wc.hInstance   = hInstance;
    wc.hIcon       = LoadIcon(NULL, IDI_APPLICATION);
    wc.hCursor     = LoadCursor(NULL, IDC_ARROW);
    wc.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wc.lpszMenuName = NULL;
    wc.lpszClassName = g_szClassName;
    wc.hIconSm     = LoadIcon(NULL, IDI_APPLICATION);


    if(!RegisterClassEx(&wc))
    {
        MessageBox(NULL, "Window Registration Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    // Step 2: Creating the Window
    hwnd = CreateWindowEx(
        WS_EX_CLIENTEDGE,
        g_szClassName,
        "The title of my window",
        WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
        NULL, NULL, hInstance, NULL);

    if(hwnd == NULL)
    {
        MessageBox(NULL, "Window Creation Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    ShowWindow(hwnd, nCmdShow);
    UpdateWindow(hwnd);

    // Step 3: The Message Loop
    while(GetMessage(&Msg, NULL, 0, 0) > 0)
    {
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```



# GUI Programming

```
switch (message)
{
    //case WM_COMMAND:
    // handle menu selections etc.
    //break;
    //case WM_PAINT:
    // draw our window - note: you must paint something here or not trap it!
    //break;
    case WM_DESTROY:
        PostQuitMessage(0);
    break;
    default:
        // We do not want to handle this message so pass back to Windows
        // to handle it in a default way
        return DefWindowProc(hWnd, message, wParam, lParam);
}
```

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

    //Step 1: Registering the Window Class
    wc.cbSize      = sizeof(WNDCLASSEX);
    wc.style       = 0;
    wc.lpfnWndProc = WndProc;
    wc.cbClsExtra  = 0;
    wc.cbWndExtra  = 0;
    wc.hInstance   = hInstance;
    wc.hIcon       = LoadIcon(NULL, IDI_APPLICATION);
    wc.hCursor     = LoadCursor(NULL, IDC_ARROW);
    wc.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wc.lpszMenuName = NULL;
    wc.lpszClassName = g_szClassName;
    wc.hIconSm     = LoadIcon(NULL, IDI_APPLICATION);

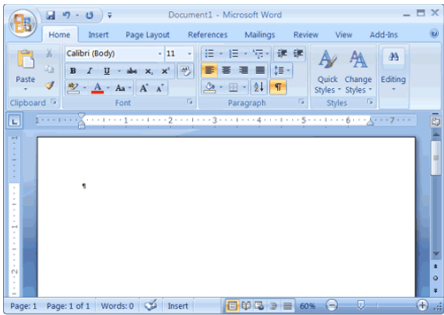
    if(!RegisterClassEx(&wc))
    {
        MessageBox(NULL, "Window Registration Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    // Step 2: Creating the Window
    hwnd = CreateWindowEx(
        WS_EX_CLIENTEDGE,
        g_szClassName,
        "The title of my window",
        WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, CW_USEDEFAULT, 240, 120,
        NULL, NULL, hInstance, NULL);

    if(hwnd == NULL)
    {
        MessageBox(NULL, "Window Creation Failed!", "Error!",
            MB_ICONEXCLAMATION | MB_OK);
        return 0;
    }

    ShowWindow(hwnd, nCmdShow);
    UpdateWindow(hwnd);

    // Step 3: The Message Loop
    while(GetMessage(&Msg, NULL, 0, 0) > 0)
    {
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```




# GUI programming

```
switch (message)
{
    //case WM_COMMAND:
    //    // handle menu selections etc.
    //break;
    //case WM_PAINT:
    //    // draw our window - note: you must paint something here or not
    //break;
    case WM_DESTROY:
        PostQuitMessage(0);
    break;
    default:
        // We do not want to handle this message so pass back to Winc
        // to handle it in a default way
        return DefWindowProc(hWnd, message, wParam, lParam);
}
```

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;
```

Hex	Decimal	Symbolic
0000	0	WM_NULL
0001	1	WM_CREATE
0002	2	WM_DESTROY
0003	3	WM_MOVE
0005	5	WM_SIZE
0006	6	WM_ACTIVATE
0007	7	WM_SETFOCUS
0008	8	WM_KILLFOCUS
000a	10	WM_ENABLE
000b	11	WM_SETREDRAW
000c	12	WM_SETTEXT
000d	13	WM_GETTEXT
000e	14	WM_GETTEXTLENGTH
000f	15	WM_PAINT
0010	16	WM_CLOSE
0011	17	WM_QUERYENDSESSION
0012	18	WM_QUIT
0013	19	WM_QUERYOPEN
0014	20	WM_ERASEBKGDND

```
        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}
```



# GUI programming

```

switch (message)
{
    //case WM_COMMAND:
    //    // handle menu selections etc.
    //break;
    //case WM_PAINT:
    //    // draw our window - note: you must paint something here or not
    //break;
    case WM_DESTROY:
        PostQuitMessage(0);
    break;
    default:
        // We do not want to handle this message so pass back to Winc
        // to handle it in a default way
        return DefWindowProc(hWnd, message, wParam, lParam);
}

```

```

int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance,
LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

```


Over 1000 last time I checked!

Hex	Decimal	Symbolic
0000	0	WM_NULL
0001	1	WM_CREATE
0002	2	WM_DESTROY
0003	3	WM_MOVE
0005	5	WM_SIZE
0006	6	WM_ACTIVATE
0007	7	WM_SETFOCUS
0008	8	WM_KILLFOCUS
000a	10	WM_ENABLE
000b	11	WM_SETREDRAW
000c	12	WM_SETTEXT
000d	13	WM_GETTEXT
000e	14	WM_GETTEXTLENGTH
000f	15	WM_PAINT
0010	16	WM_CLOSE
0011	17	WM_QUERYENDSESSION
0012	18	WM_QUIT
0013	19	WM_QUERYOPEN
0014	20	WM_ERASEBKGDND

```

        TranslateMessage(&Msg);
        DispatchMessage(&Msg);
    }
    return Msg.wParam;
}

```



# GUI programming

```
switch (message)
{
    //case WM_COMMAND:
    //    // handle menu select
    //break;
    //case WM_PAINT:
    //    // draw our window -
    //break;
    case WM_DESTROY:
        PostQuitMessage(0);
    break;
    default:
        // We do not want to handle this message so pass back to Winc
        // to handle it in a default way
        return DefWindowProc(hWnd, message, wParam, lParam);
}
```

```
void OnMove() { ... }
void OnSize() { ... }
void OnPaint() { ... }
```


Hex	Decimal	Symbolic
0000	0	WM_NULL
0001	1	WM_CREATE
0002	2	WM_DESTROY
0003	3	WM_MOVE
0005	5	WM_SIZE
0006	6	WM_ACTIVATE
0007	7	WM_SETFOCUS
0008	8	WM_KILLFOCUS
000a	10	WM_ENABLE
000b	11	WM_SETREDRAW
000c	12	WM_SETTEXT
000d	13	WM_GETTEXT
000e	14	WM_GETTEXTLENGTH
000f	15	WM_PAINT
0010	16	WM_CLOSE
0011	17	WM_QUERYENDSESSION
0012	18	WM_QUIT
0013	19	WM_QUERYOPEN
0014	20	WM_ERASEBKGDND

Over 1000 last time I checked!

```
int WINAPI WinMain(HINSTANCE hInstance, HINSTANCE hPrevInstance, LPSTR lpCmdLine, int nCmdShow)
{
    WNDCLASSEX wc;
    HWND hwnd;
    MSG Msg;

```

```
TranslateMessage(&Msg);
DispatchMessage(&Msg);
}
return Msg.wParam;
}
```



# GUI Programming Distilled

```
1  winmain(...) {  
2      while(true) {  
3          message = GetMessage();  
4          switch(message) {  
5              case WM_THIS: DoThis(); break;  
6              case WM_THAT: DoThat(); break;  
7              case WM_OTHERTHING: DoOtherThing(); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```



# GUI Programming Distilled

```
1  winmain(...) {  
2      while(true) {  
3          message = GetMessage();  
4          switch(message) {  
5              case WM_THIS: DoThis(); break;  
6              case WM_THAT: DoThat(); break;  
7              case WM_OTHERTHING: DoOtherThing(); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

Pros

# GUI Programming Distilled

```
1  winmain (...) {  
2      while (true) {  
3          message = GetMessage ();  
4          switch (message) {  
5              case WM_THIS: DoThis (); break;  
6              case WM_THAT: DoThat (); break;  
7              case WM_OTHERTHING: DoOtherThing (); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

## Pros

- Simple imperative programming

# GUI Programming Distilled

```
1  winmain (...) {  
2      while (true) {  
3          message = GetMessage ();  
4          switch (message) {  
5              case WM_THIS: DoThis (); break;  
6              case WM_THAT: DoThat (); break;  
7              case WM_OTHERTHING: DoOtherThing (); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

## Pros

- Simple imperative programming
- Good fit for uni-processor

# GUI Programming Distilled

```
1  winmain (...) {  
2      while (true) {  
3          message = GetMessage ();  
4          switch (message) {  
5              case WM_THIS: DoThis (); break;  
6              case WM_THAT: DoThat (); break;  
7              case WM_OTHERTHING: DoOtherThing (); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

## Pros

- Simple imperative programming
- Good fit for uni-processor

## Cons

# GUI Programming Distilled

```
1  winmain(...) {  
2      while(true) {  
3          message = GetMessage();  
4          switch(message) {  
5              case WM_THIS: DoThis(); break;  
6              case WM_THAT: DoThat(); break;  
7              case WM_OTHERTHING: DoOtherThing(); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

## Pros

- Simple imperative programming
- Good fit for uni-processor

## Cons

- Awkward/verbose

# GUI Programming Distilled

```
1  winmain(...) {  
2      while(true) {  
3          message = GetMessage();  
4          switch(message) {  
5              case WM_THIS: DoThis(); break;  
6              case WM_THAT: DoThat(); break;  
7              case WM_OTHERTHING: DoOtherThing(); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```

## Pros

- Simple imperative programming
- Good fit for uni-processor

## Cons

- Awkward/verbose
- **Obscures available parallelism**

# GUI Programming Distilled

```
1 winmain(...) {  
2   while(true) {  
3     message = GetMessage();  
4     switch(message) {  
5       case WM_LONGRUNNING_CPU_HOG: HogCPU(); break;  
6       case WM_HIGH_LATENCY_IO: BlockForALongTime(); break;  
7       case WM_DO_QUICK_IMPORTANT_THING: HopeForTheBest(); break;  
8     }  
9   }  
10 }  
11 }
```

## Pros

- Simple imperative programming
- Good fit for uni-processor

## Cons

- Awkward/verbose
- **Obscures available parallelism**

# GUI Programming Distilled


```
1  winmain(...) {  
2      while(true) {  
3          message = GetMessage();  
4          switch(message) {  
5              case WM_THIS: DoThis(); break;  
6              case WM_THAT: DoThat(); break;  
7              case WM_OTHERTHING: DoOtherThing(); break;  
8              case WM_DONE: return;  
9          }  
10     }  
11 }
```



# GUI Programming Distilled

```
1 winmain(...) {  
2     while(true) {  
3         message = GetMessage();  
4         switch(message) {  
5             case WM_THIS: DoThis(); break;  
6             case WM_THAT: DoThat(); break;  
7             case WM_OTHERTHING: DoOtherThing(); break;  
8             case WM_DONE: return;  
9         }  
10    }  
11 }
```

How can we  
parallelize  
this?




# Parallel GUI Implementation 1

```
1 winmain(...) {  
2     while(true) {  
3         message = GetMessage();  
4         switch(message) {  
5             case WM_THIS: DoThis(); break;  
6             case WM_THAT: DoThat(); break;  
7             case WM_OTHERTHING: DoOtherThing(); break;  
8             case WM_DONE: return;  
9         }  
10    }  
11 }
```


# Parallel GUI Implementation 1

```
1 winmain(...) {  
2   while(true) {  
3     message = GetMessage();  
4     switch(message) {  
5       case WM_THIS: DoThis(); break;  
6       case WM_THAT: DoThat(); break;  
7       case WM_OTHERTHING: DoOtherThing(); break;  
8       case WM_DONE: return;  
9     }  
10  }  
11 }
```




# Parallel GUI Implementation 1

```
1 winmain(...) {  
2   while(true) {  
3     message = GetMessage();  
4     switch(message) {  
5       case WM_THIS: DoThis(); break;  
6       case WM_THAT: DoThat(); break;  
7       case WM_OTHERTHING: DoOtherThing(); break;  
8       case WM_DONE: return;  
9     }  
10  }  
11 }
```




# Parallel GUI Implementation 1

```
1 winmain(...) {  
2   while(true) {  
3     message = GetMessage();  
4     switch(message) {  
5       case WM_THIS: DoThis(); break;  
6       case WM_THAT: DoThat(); break;  
7       case WM_OTHERTHING: DoOtherThing(); break;  
8       case WM_DONE: return;  
9     }  
10  }  
11 }
```



# Parallel GUI Implementation 1

```
winmain() {  
    pthread_create(&tids[i++], DoThisProc);  
    pthread_create(&tids[i++], DoThatProc);  
    pthread_create(&tids[i++], DoOtherThingProc);  
    for(j=0; j<i; j++)  
        pthread_join(&tids[j]);  
}  
  
DoThisProc() {  
    while(true) {  
        if(ThisHasHappened)  
            DoThis();  
    }  
}
```




# Parallel GUI Implementation 1

```
winmain() {  
    pthread_create(&tids[i++], DoThisProc);  
    pthread_create(&tids[i++], DoThatProc);  
    pthread_create(&tids[i++], DoOtherThingProc);  
    for(j=0; j<i; j++)  
        pthread_join(&tids[j]);  
}
```

Pros/cons?

```
DoThisProc() {  
    while(true) {  
        if(ThisHasHappened)  
            DoThis();  
    }  
}
```



# Parallel GUI Implementation 1

```
winmain() {  
    pthread_create(&tids[i++], DoThisProc);  
    pthread_create(&tids[i++], DoThatProc);  
    pthread_create(&tids[i++], DoOtherThingProc);  
    for(j=0; j<i; j++)  
        pthread_join(&tids[j]);  
}  
  
DoThisProc() {  
    while(true) {  
        if(ThisHasHappened) {  
            DoThis();  
        }  
    }  
}
```


Pros/cons?

Pros:

- Encapsulates parallel work

Cons:

- Obliterates original code structure
- How to assign handlers → CPUs?
- Load balance?!?
- Utilization





# Parallel GUI Implementation 2




```
winmain() {  
    for(i=0; i<NUMPROCS; i++)  
        pthread_create(&tids[i], HandlerProc);  
    for(i=0; i<NUMPROCS; i++)  
        pthread_join(&tids[i]);  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```



# Parallel GUI Implementation 2

Pros/cons?




```
winmain() {  
    for(i=0; i<NUMPROCS; i++)  
        pthread_create(&tids[i], HandlerProc);  
    for(i=0; i<NUMPROCS; i++)  
        pthread_join(&tids[i]);  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```



# Parallel GUI Implementation 2

Pros/cons?

```
winmain() {  
    for(i=0; i<NUMPROCS; i++)  
        pthread_create(&tids[i], H  
    for(i=0; i<NUMPROCS; i++)  
        pthread_join(&tids[i]);  
}
```

Pros:

- Preserves programming model
- Can recover some parallelism

Cons:


- Workers still have same problem
- How to load balance?
- Shared mutable state a problem

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```



# Parallel GUI Implementation 2

Pros/cons?

Pros:

- Preserves programming model
- Can recover some parallelism

Cons:

- Workers still have same problem
- How to load balance?
- Shared mutable state a problem

```
winmain() {  
    for(i=0; i<NUMPROCS; i++)  
        pthread_create(&tids[i], H  
    for(i=0; i<NUMPROCS; i++)  
        pthread_join(&tids[i]);  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```

```
threadproc(...) {  
    while(true) {  
        message = GetMessage();  
        switch(message) {  
            case WM_THIS: DoThis();  
            case WM_THAT: DoThat();  
        }  
    }  
}
```



*Extremely difficult to solve  
without changing the whole  
programming model...so*

***change it***

# Event-based Programming: Motivation

# Event-based Programming: Motivation

- Threads have a *\*lot\** of down-sides:
  - Tuning parallelism for different environments
  - Load balancing/assignment brittle
  - Shared state requires locks →
    - Priority inversion
    - Deadlock
    - Incorrect synchronization
  - ...

# Event-based Programming: Motivation

- Threads have a *\*lot\** of down-sides:
  - Tuning parallelism for different environments
  - Load balancing/assignment brittle
  - Shared state requires locks →
    - Priority inversion
    - Deadlock
    - Incorrect synchronization
  - ...
- Events: *restructure programming model to have no threads!*

# Event Programming Model Basics



# Event Programming Model Basics

- Programmer *only writes events*

# Event Programming Model Basics

- Programmer *only writes events*
- Event: an object queued for a module (think future/promise)

# Event Programming Model Basics

- Programmer *only writes events*
- Event: an object queued for a module (think future/promise)
- Basic primitives
  - `create_event_queue(handler) → event_q`
  - `enqueue_event(event_q, event-object)`
    - Invokes handler (eventually)

# Event Programming Model Basics

- Programmer *only writes events*
- Event: an object queued for a module (think future/promise)
- Basic primitives
  - `create_event_queue(handler) → event_q`
  - `enqueue_event(event_q, event-object)`
    - Invokes handler (eventually)
- Scheduler decides which event to execute next
  - E.g. based on priority, CPU usage, etc.

# Event-based programming

# Event-based programming

```
switch (message)
{
    //case WM_COMMAND:
    // handle menu selections etc.
    //break;
    //case WM_PAINT:
    // draw our window - note: you must paint something here or not trap it!
    //break;
    case WM_DESTROY:
        PostQuitMessage(0);
    break;
    default:
        // We do not want to handle this message so pass back to Windows
        // to handle it in a default way
        return DefWindowProc(hWnd, message, wParam, lParam);
}
```

# Event-based programming



# Event-based programming

```
PROGRAM MyProgram {  
    OnSize () {}  
    OnMove () {}  
    OnClick () {}  
    OnPaint () {}  
}
```






# Event-based programming

```
PROGRAM MyProgram {  
    OnSize () {}  
    OnMove () {}  
    OnClick () {}  
    OnPaint () {}  
}
```

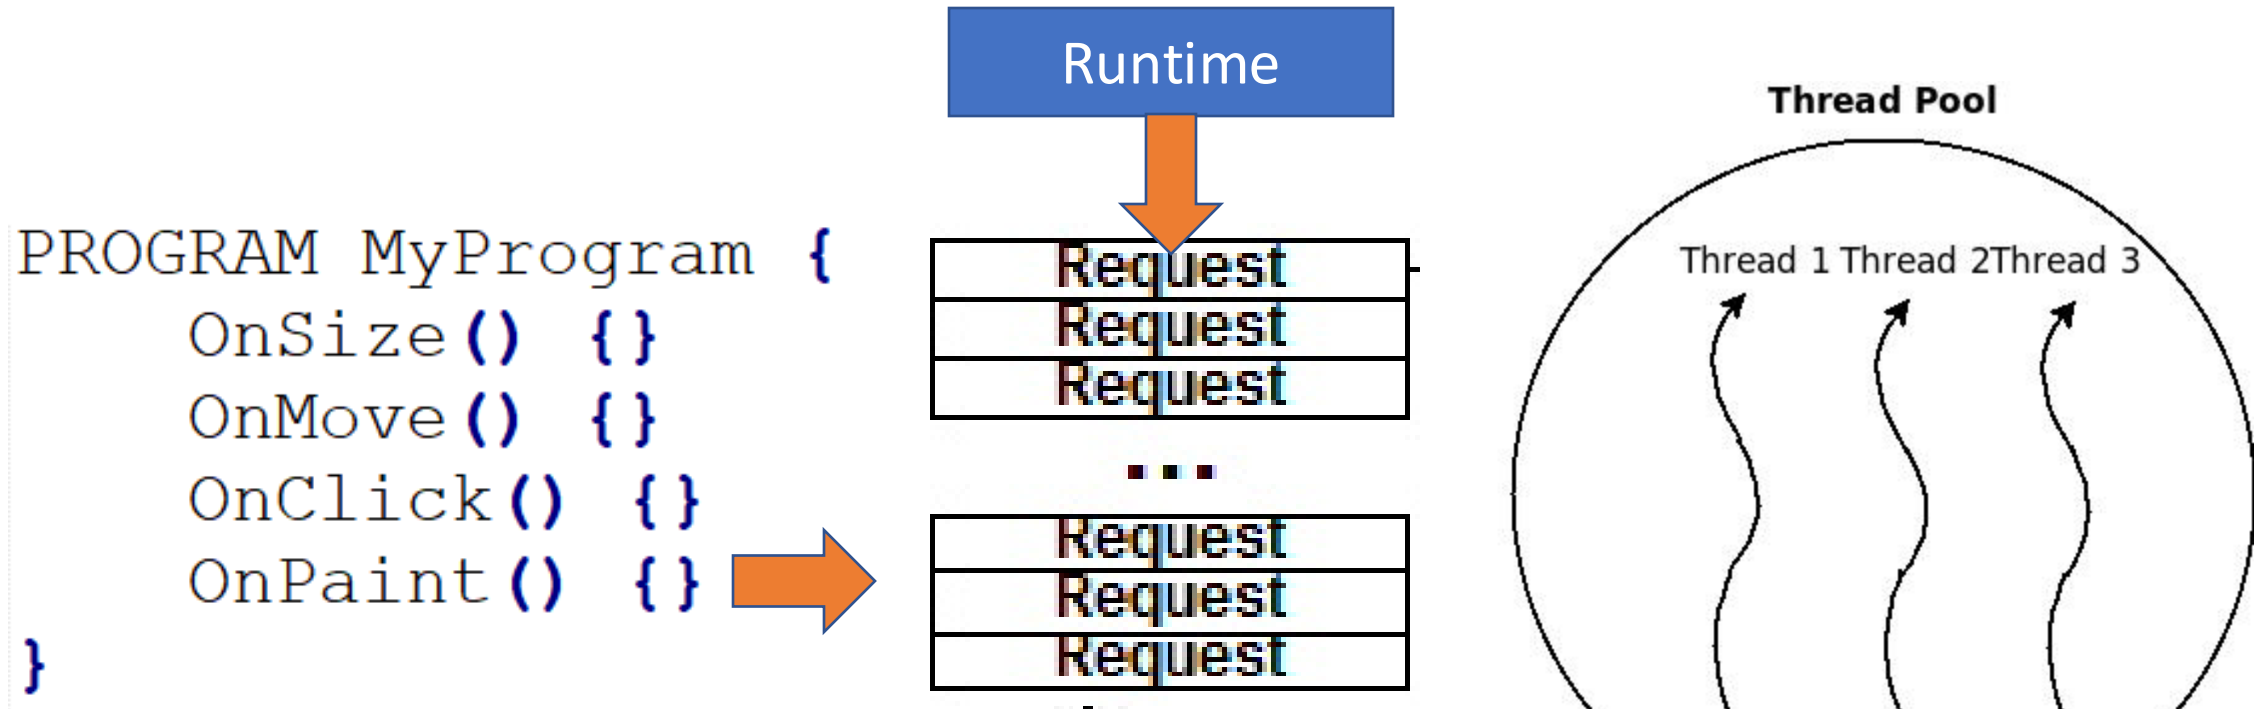


# Event-based programming

```
PROGRAM MyProgram {  
    OnSize () {}  
    OnMove () {}  
    OnClick () {}  
    OnPaint () {}  
}
```



# Event-based programming



Is the problem solved?

# Another Event-based Program

# Another Event-based Program

```
1 PROGRAM MyProgram {
2     OnOpenFile () {
3         char szFileName [BUFSIZE]
4         InitFileName (szFileName);
5         FILE file = ReadFileEx (szFileName);
6         LoadFile (file);
7         RedrawScreen ();
8     }
9     OnPaint ();
10 }
```

# Another Event-based Program

```
1 PROGRAM MyProgram {
2     OnOpenFile () {
3         char szFileName [BUFSIZE]
4         InitFileName (szFileName);
5         FILE file = ReadFileEx (szFileName);
6         LoadFile (file);
7         RedrawScreen ();
8     }
9     OnPaint ();
10 }
```



Blocks!

# Another Event-based Program

```
1 PROGRAM MyProgram {
2     OnOpenFile () {
3         char szFileName [BUFSIZE]
4         InitFileName (szFileName);
5         FILE file = ReadFileEx (szFileName);
6         LoadFile (file);
7         RedrawScreen ();
8     }
9     OnPaint ();
10 }
```

Burns CPU!

Blocks!

# Another Event-based Program

```
1 PROGRAM MyProgram {
2   OnOpenFile () {
3     char szFileName [BUFSIZE]
4     InitFileName (szFileName);
5     FILE file = ReadFileEx (szFileName);
6     LoadFile (file);
7     RedrawScreen ();
8   }
9   OnPaint ();
10 }
```

Uses Other Handlers!  
(call OnPaint?)

Burns CPU!

Blocks!



# No problem!

## Just use more events/handlers, right?

```
1 PROGRAM MyProgram {
2     TASK ReadFileAsync(name, callback) {
3         ReadFileSync(name);
4         Call(callback);
5     }
6     CALLBACK FinishOpeningFile() {
7         LoadFile(file);
8         RedrawScreen();
9     }
10    OnOpenFile() {
11        FILE file;
12        char szName[BUFSIZE]
13        InitFileName(szName);
14        EnqueueTask(ReadFileAsync(szName, FinishOpeningFile));
15    }
16    OnPaint();
17 }
```

# Continuations, BTW

```
1 PROGRAM MyProgram {
2     OnOpenFile () {
3         ReadFile (file, FinishOpeningFile);
4     }
5     OnFinishOpeningFile () {
6         LoadFile (file, OnFinishLoadingFile);
7     }
8     OnFinishLoadingFile () {
9         RedrawScreen ();
10    }
11    OnPaint ();
12 }
```

# Stack-Ripping

```
1 PROGRAM MyProgram {
2     TASK ReadFileAsync(name, callback) {
3         ReadFileSync(name);
4         Call(callback);
5     }
6     CALLBACK FinishOpeningFile() {
7         LoadFile(file);
8         RedrawScreen();
9     }
10    OnOpenFile() {
11        FILE file;
12        char szName[BUFSIZE]
13        InitFileName(szName);
14        EnqueueTask(ReadFileAsync(szName, FinishOpeningFile));
15    }
16    OnPaint();
17 }
```

# Stack-Ripping

```
1 PROGRAM MyProgram {
2     TASK ReadFileAsync(name, callback) {
3         ReadFileSync(name);
4         Call(callback);
5     }
6     CALLBACK FinishOpeningFile() {
7         LoadFile(file);
8         RedrawScreen();
9     }
10    OnOpenFile() {
11        FILE file;
12        char szName[BUFSIZE];
13        InitFileName(szName);
14        EnqueueTask(ReadFileAsync(szName, FinishOpeningFile));
15    }
16    OnPaint();
17 }
```

# Stack-Ripping

```
1 PROGRAM MyProgram {
2     TASK ReadFileAsync(name, callback) {
3         ReadFileSync(name);
4         Call(callback);
5     }
6     CALLBACK FinishOpeningFile() {
7         LoadFile(file);
8         RedrawScreen();
9     }
10    OnOpenFile() {
11        FILE file;
12        char szName[BUFSIZE];
13        InitFileName(szName);
14        EnqueueTask(ReadFileAsync(szName, FinishOpeningFile));
15    }
16    OnPaint();
17 }
```

# Stack-Ripping

```
1 PROGRAM MyProgram {
2     TASK ReadFileAsync(name, callback) {
3         ReadFileSync(name);
4         Call(callback);
5     }
6     CALLBACK FinishOpeningFile() {
7         LoadFile(file);
8         RedrawScreen();
9     }
10    OnOpenFile() {
11        FILE file;
12        char szName[BUFSIZE];
13        InitFileName(szName);
14        EnqueueTask(ReadFileAsync(szName, FinishOpeningFile));
15    }
16    OnPaint();
17 }
```

Stack-based state out-of-scope!  
Requests must carry state

# Threads vs Events

- Thread Pros

- Event Pros

- Thread Cons

- Event Cons

# Threads vs Events

- Thread Pros

- Overlap I/O and computation
  - While looking sequential
- Intermediate state on stack
- Control flow naturally expressed

- Thread Cons

- Synchronization required
- Overflowable stack
- Stack memory pressure

- Event Pros

- Easier to create well-conditioned system
- Easier to express dynamic change in level of parallelism

- Event Cons

- Difficult to program
- Control flow between callbacks obscure
- When to deallocate memory
- Incomplete language/tool/debugger support
- Difficult to exploit concurrent hardware



# Threads vs Events

- Thread Pros

- Overlap I/O and computation
  - While looking sequential
- Intermediate state on stack
- Control flow naturally

- Thread Cons

- Synchronization required
- Overflowable stack
- Stack memory pressure

- Event Pros

- Easier to create well-conditioned system
- Easier to express dynamic change in level of parallelism

- Event Cons

- Difficult to program
- Control flow between callbacks obscure
- When to deallocate memory
- Incomplete language/tool/debugger support
- Difficult to exploit concurrent hardware

Language-level  
Futures: the  
sweet spot?

# Threads vs Events



- Thread Pros

- Overlap I/O and computation
  - While looking sequential
- Intermediate state on stack
- Control flow naturally

- Thread Cons

- Synchronization required
- Overflowable stack
- Stack memory pressure

- Event Pros

- Easier to create well-conditioned system
- Easier to express dynamic change in level of parallelism

- Event Cons

- Difficult to program
- Control flow between callbacks obscure
- When to deallocate memory
- Incomplete language/tool/debugger support
- Difficult to exploit concurrent hardware

Language-level  
Futures: the  
sweet spot?

# Thread Pool Implementation

```
///-----  
/// <summary> Starts the threads. </summary>  
///  
/// <remarks> crossbac, 8/22/2013. </remarks>  
///  
/// <param name="uiThreads"> The threads. </param>  
/// <param name="bWaitAllThreadsAlive"> The wait all threads alive. </param>  
///-----  
  
void  
ThreadPool::StartThreads(  
    __in UINT uiThreads,  
    __in BOOL bWaitAllThreadsAlive  
)  
{  
    Lock();  
    if(uiThreads != 0 && m_vhThreadDescs.size() < m_uiTargetSize)  
        ResetEvent(m_hAllThreadsAlive);  
    while(m_vhThreadDescs.size() < m_uiTargetSize) {  
        for(UINT i=0; i<uiThreads; i++) {  
            THREADDESC* pDesc = new THREADDESC(this);  
            HANDLE * phThread = &pDesc->hThread;  
            *phThread = CreateThread(NULL, 0, _ThreadProc, pDesc, 0, NULL);  
            m_vhAvailable.push_back(*phThread);  
            m_vhThreadDescs[*phThread] = pDesc;  
        }  
    }  
    m_uiThreads = (UINT)m_vhThreadDescs.size();  
    Unlock();  
    if(bWaitAllThreadsAlive)  
        WaitThreadsAlive();  
}
```

# Thread Pool Implementation

```
///-----  
/// <summary> Starts the threads. </summary>  
///  
/// <remarks> crossbac, 8/22/2013. </remarks>  
///  
/// <param name="uiThreads"> The threads. </param>  
/// <param name="bWaitAllThreadsAlive"> The wait all threads alive. </param>  
///-----  
  
void  
ThreadPool::StartThreads(  
    __in UINT uiThreads,  
    __in BOOL bWaitAllThreadsAlive  
)  
{  
    Lock();  
    if(uiThreads != 0 && m_vhThreadDescs.size() < m_uiTargetSize)  
        ResetEvent(m_hAllThreadsAlive);  
    while(m_vhThreadDescs.size() < m_uiTargetSize) {  
        for(UINT i=0; i<uiThreads; i++) {  
            THREADDESC* pDesc = new THREADDESC(this);  
            HANDLE * phThread = &pDesc->hThread;  
            *phThread = CreateThread(NULL, 0, _ThreadProc, pDesc, 0, NULL);  
            m_vhAvailable.push_back(*phThread);  
            m_vhThreadDescs[*phThread] = pDesc;  
        }  
    }  
    m_uiThreads = (UINT)m_vhThreadDescs.size();  
    Unlock();  
    if(bWaitAllThreadsAlive)  
        WaitThreadsAlive();  
}
```

Cool project  
idea: build a  
thread pool!

# Thread Pool Implementation

```
DWORD
ThreadPool::ThreadPoolProc (
    _In_ THREADDESC * pDesc
)
{
    HANDLE hThread = pDesc->hThread;
    HANDLE hStartEvent = pDesc->hStartEvent;
    HANDLE hRuntimeTerminate = PTask::Runtime::GetRuntimeTerminateEvent();
    HANDLE vEvents[] = { hStartEvent, hRuntimeTerminate };

    NotifyThreadAlive(hThread);
    while(!pDesc->bTerminate) {

        DWORD dwWait = WaitForMultipleObjects(dwEvents, vEvents, FALSE, INFINITE);
        pDesc->Lock();
        pDesc->bTerminate |= bTerminate;
        if(pDesc->bRoutineValid && !pDesc->bTerminate) {
            LPTHREAD_START_ROUTINE lpRoutine = pDesc->lpRoutine;
            LPVOID lpParameter = pDesc->lpParameter;
            pDesc->bActive = TRUE;
            pDesc->Unlock();
            dwResult = (*lpRoutine)(lpParameter);
            pDesc->Lock();
            pDesc->bActive = FALSE;
            pDesc->bRoutineValid = FALSE;
        }
        pDesc->Unlock();
        Lock();
        m_vhInFlight.erase(pDesc->hThread);
        if(!pDesc->bTerminate)
            m_vhAvailable.push_back(pDesc->hThread);
        Unlock();
    }
    NotifyThreadExit(hThread);
    return dwResult;
}
```

# ThreadPool Implementation

```
///-----  
/// <summary> Starts a thread: if a previous call to RequestThread was made with  
/// the bStartThread parameter set to false, this API signals the thread  
/// to begin. Otherwise, the call has no effect (returns FALSE). </summary>  
///  
/// <remarks> crossbac, 8/29/2013. </remarks>  
///  
/// <param name="hThread"> The thread. </param>  
///  
/// <returns> true if it succeeds, false if it fails. </returns>  
///-----
```

```
BOOL  
ThreadPool::SignalThread(  
    _In HANDLE hThread  
)  
{  
    Lock();  
    BOOL bResult = FALSE;  
    std::set<HANDLE>::iterator si = m_vhWaitingStartSignal.find(hThread);  
    if(si != m_vhWaitingStartSignal.end()) {  
        m_vhWaitingStartSignal.erase(hThread);  
        THREADDESC * pDesc = m_vhThreadDescs[hThread];  
        HANDLE hEvent = pDesc->hStartEvent;  
        SetEvent(hEvent);  
        bResult = TRUE;  
    }  
    Unlock();  
    return bResult;  
}
```

# Redux: Futures in Context

# Redux: Futures in Context

Futures:



# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

## Compromise Model:

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

## Compromise Model:

- Event-based programming

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

## Compromise Model:

- Event-based programming
- Thread-based programming

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

# Redux: Futures in Context

## Futures:

- *abstraction* for concurrent work supported by
  - Compiler: abstractions are *language-level objects*
  - Runtime: scheduler, task queues, thread-pools are *transparent*
- Programming remains **mostly** imperative
  - Threads of control peppered with asynchronous/concurrent tasks

## Compromise Model:

- Event-based programming
- Thread-based programming

Currently: 2nd renaissance IMHO

```
1 static void runAsyncExample() {
2     CompletableFuture cf = CompletableFuture.runAsync(() -> {
3         assertTrue(Thread.currentThread().isDaemon());
4         randomSleep();
5     });
6     assertFalse(cf.isDone());
7     sleepEnough();
8     assertTrue(cf.isDone());
9 }
```

Questions?