

## 44 Terminating Threads

Termination of a thread (usually `:-`) returns a value. There are two (regular) ways to terminate a thread:

1. The initial function call has terminated. Then the return value is the return value of the call.
2. The thread executes the statement `exit (e)`; Then the return value equals the value of `e`.

### Warning:

- We want to return the return value in the bottom stack cell.
- `exit` may occur arbitrarily deeply nested inside a recursion. Then we de-allocate all stack frames ...
- ... and jump to the terminal treatment of threads at address `f` .


Therefore, we translate:

```
code exit (e); ρ = codeR e ρ
                    exit
                    term
                    next
```

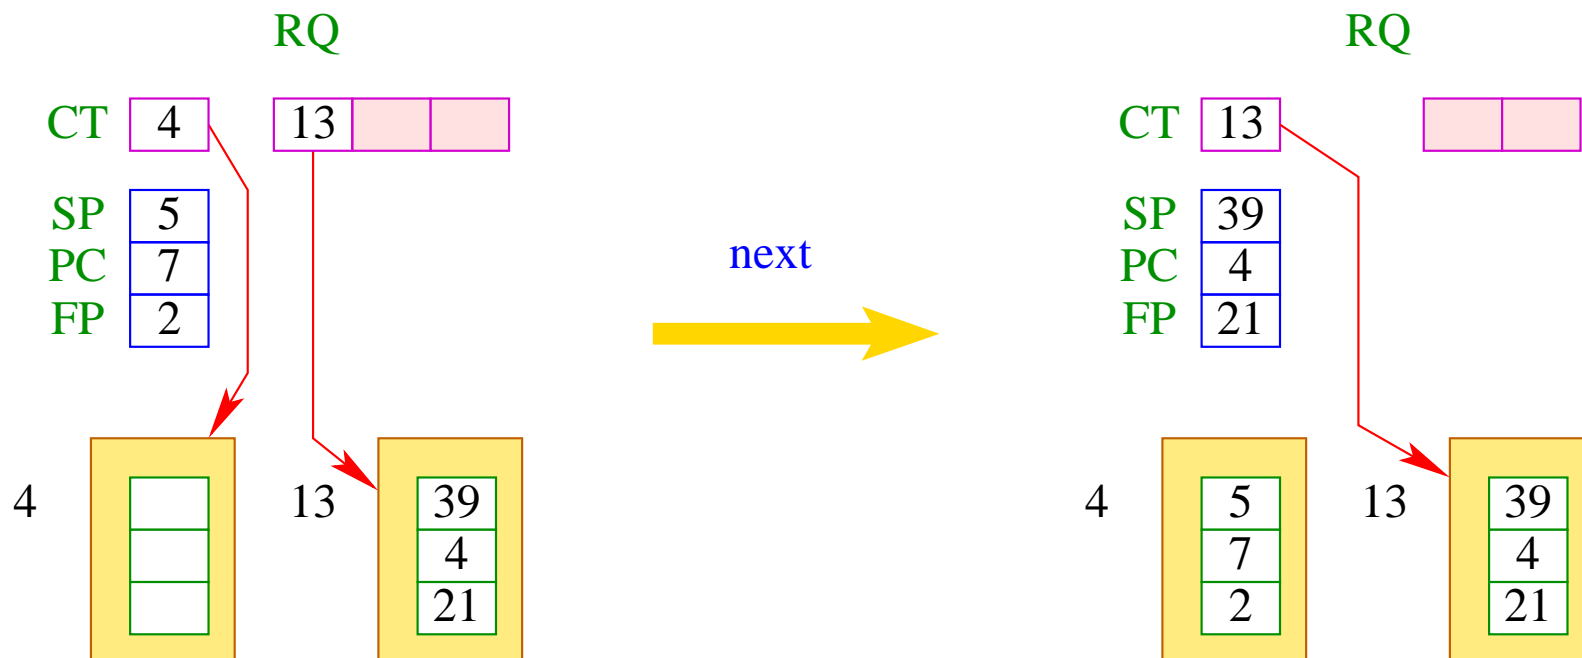
The instruction `term` is explained later :-)

The instruction `exit` successively pops all stack frames:

```
result = S[SP];
while (FP ≠ -1) {
    SP = FP-2;
    FP = S[FP-1];
}
S[SP] = result;
```



The instruction `next` activates the next executable thread:  
in contrast to `yield` the current thread is **not** inserted into `RQ`.



If the queue **RQ** is empty, we additionally terminate the whole program:

```
if (0 > ct = dequeue( RQ )) halt;  
else {  
    save ();  
    CT = ct;  
    restore ();  
}
```

## 45 Waiting for Termination


Occasionally, a thread may only continue with its execution, if some other thread has terminated. For that, we have the expression `join (e)` where we assume that  $e$  evaluates to a thread id `tid`.

- If the thread with the given `tid` is already terminated, we return its return value.
- If it is not yet terminated, we interrupt the current thread execution.
- We insert the current thread into the queue of threads already waiting for the termination.

We save the current registers and switch to the next executable thread.

- Thread waiting for termination are maintained in the table `JTab`.
- There, we also store the return values of threads `:-)`

## Example:



Thread 0 is running, thread 1 could run, threads 2 and 3 wait for the termination of 1, and thread 4 waits for the termination of 3.

Thus, we translate:

$$\text{code}_R \text{ join } (e) \rho = \text{code}_R e \rho$$


`join`  
`finalize`

... where the instruction `join` is defined by:

```
tid = S[SP];
if (TTab[tid][1] ≥ 0) {
    enqueue ( JTab[tid], CT );
    next
}
```



... accordingly:



$S[SP] = JTab[tid][1];$

The instruction sequence:

`term`

`next`

is executed before a thread is terminated.

Therefore, we store them at the location `f`.

The instruction `next` switches to the next executable thread. Before that, though,

- ... the last stack frame must be popped and the result be stored in the table `JTab` ;
- ... the thread must be marked as terminated, e.g., by additionally setting the `PC` to `-1`;
- ... all threads must be notified which have waited for the termination.

For the instruction `term` this means:

```
PC = -1;  
JTab[CT][1] = S[SP];  
freeStack(SP);  
while (0 ≤ tid = dequeue ( JTab[CT][0] ))  
    enqueue ( RQ, tid );
```

The run-time function `freeStack (int adr)` removes the (one-element) stack at the location `adr` :



## 46 Mutual Exclusion

A **mutex** is an (abstract) datatype (in the heap) which should allow the programmer to dedicate exclusive access to a shared resource (**mutual exclusion**).

The datatype supports the following operations:


<b>Mutex</b> * newMutex ();	—	creates a new mutex;
void lock ( <b>Mutex</b> *me);	—	tries to acquire the mutex;
void unlock ( <b>Mutex</b> *me);	—	releases the mutex;

### Warning:

A thread is only allowed to release a mutex if it has owned it beforehand :-)

A mutex me consists of:

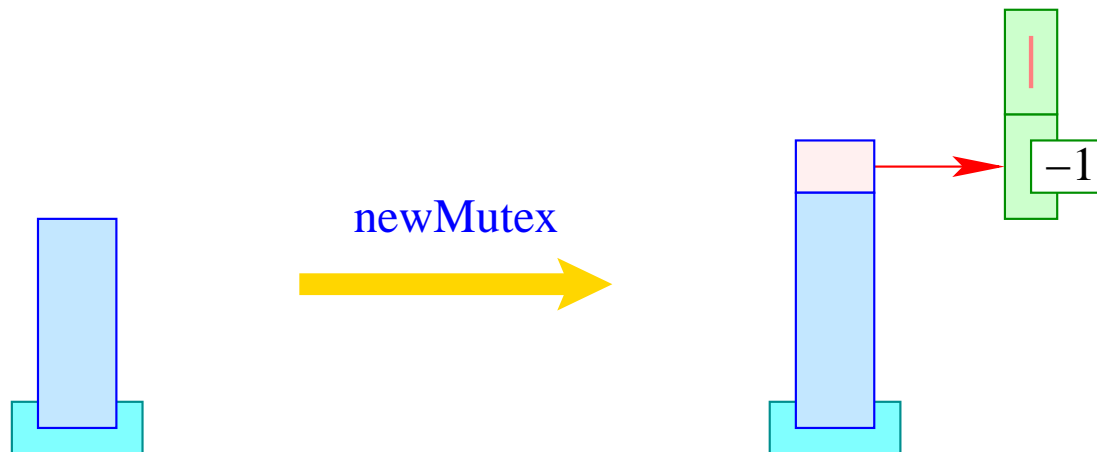
- the tid of the current owner (or  $-1$  if there is no one);
- the queue BQ of blocked threads which want to acquire the mutex.



Then we translate:

$$\text{code}_R \text{ newMutex } () \rho = \text{newMutex}$$

where:

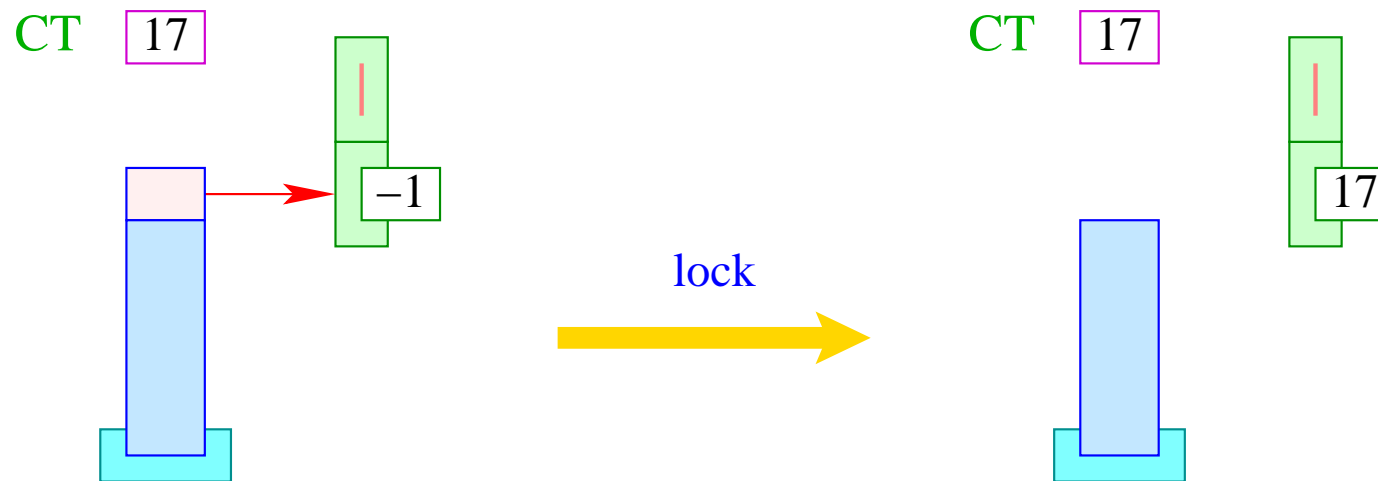


Then we translate:


$$\text{code lock}(e); \rho = \text{code}_R e \rho$$

lock

where:



If the mutex is already owned by someone, the current thread is interrupted:



```
if (S[S[SP]] < 0) S[S[SP--]] = CT;  
else {  
    enqueue ( S[SP--]+1, CT );  
    next;  
}
```

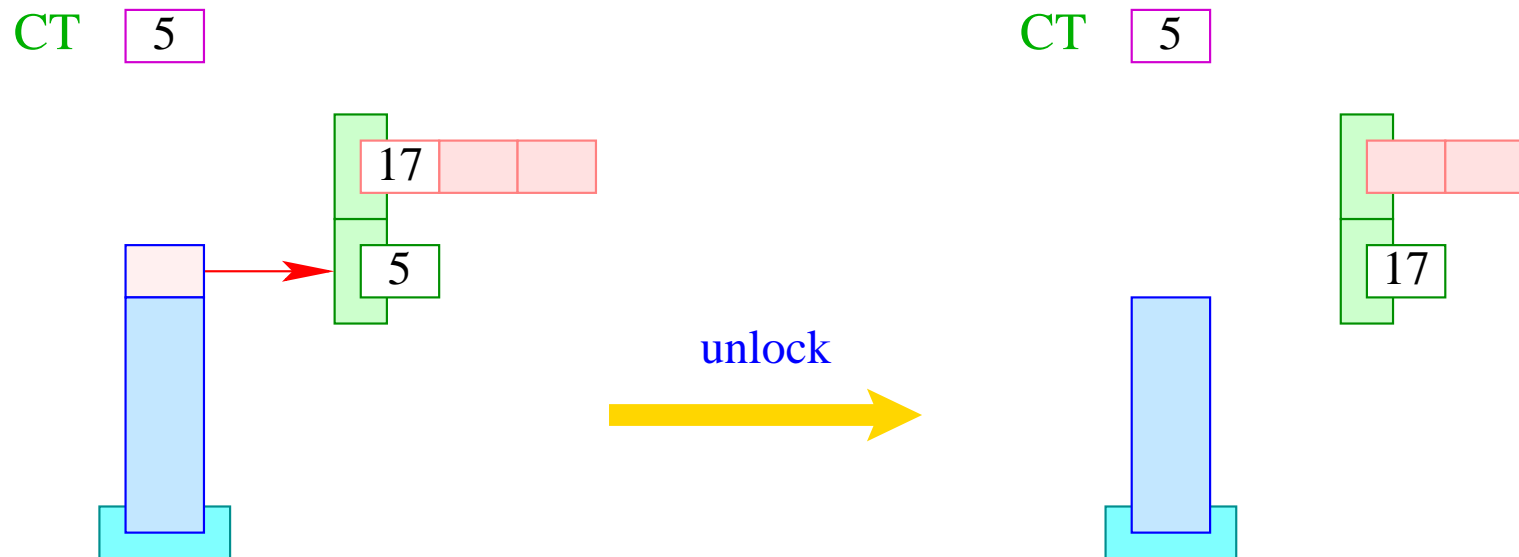


Accordingly, we translate:


$$\text{code unlock } (e); \rho = \text{code}_R e \rho$$

unlock

where:



If the queue **BQ** is empty, we release the mutex:



```

if (S[S[SP]]  $\neq$  CT) Error ("Illegal unlock!");
if (0 > tid = dequeue ( S[SP]+1)) S[S[SP--]] = -1;
else {
    S[S[SP--]] = tid;
    enqueue ( RQ, tid );
}

```

## 47 Waiting for Better Weather

It may happen that a thread owns a mutex but must wait until some extra condition is true.

Then we want the thread to remain in-active until it is told otherwise.

For that, we use **condition variables**. A condition variable consists of a queue **WQ** of waiting threads :-)



For condition variables, we introduce the functions:

<b>CondVar * newCondVar ();</b>	— creates a new condition variable;
<b>void wait (CondVar * cv, Mutex * me);</b>	— enqueues the current thread;
<b>void signal (CondVar * cv);</b>	— re-animates one waiting thread;
<b>void broadcast (CondVar * cv);</b>	— re-animates all waiting threads.

Then we translate:

$$\text{code}_R \text{ newCondVar } () \rho = \text{newCondVar}$$

where:

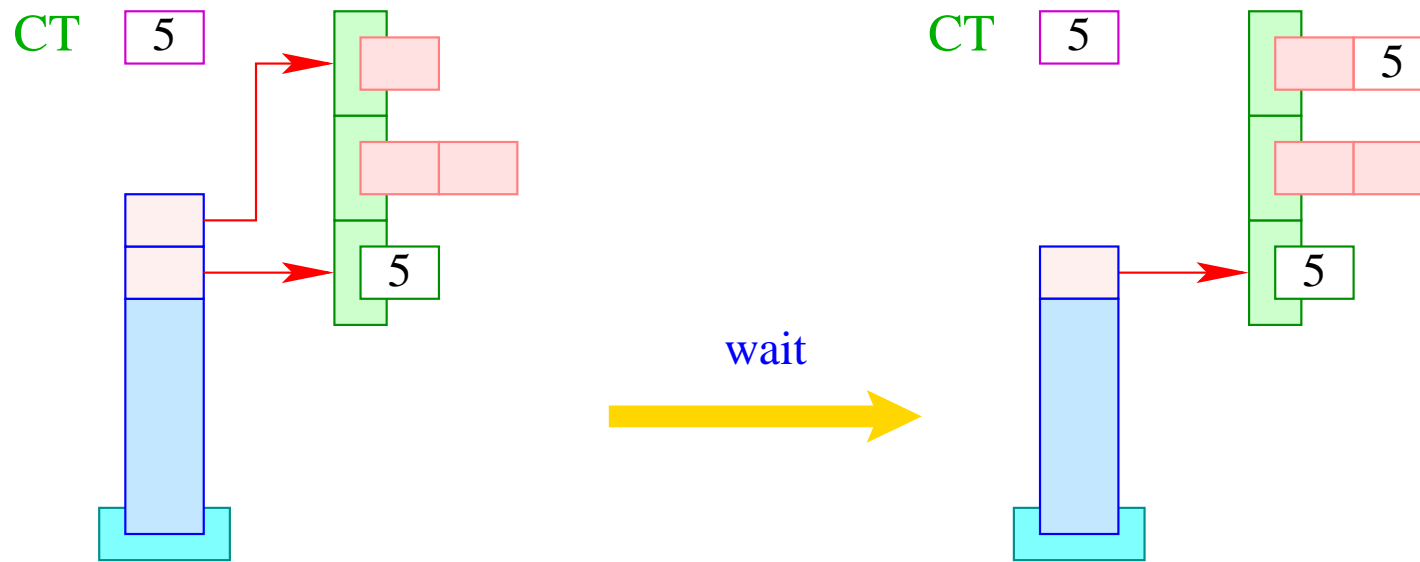


After enqueueing the current thread, we release the mutex. After re-animation, though, we must acquire the mutex again.

Therefore, we translate:

$$\text{code wait } (e_0, e_1); \rho = \begin{array}{l} \text{code}_R e_1 \rho \\ \text{code}_R e_0 \rho \\ \text{wait} \\ \text{dup} \\ \text{unlock} \\ \text{next} \\ \text{lock} \end{array}$$

where ...




```

if (S[S[SP-1]] ≠ CT) Error ("Illegal wait!");
enqueue ( S[SP], CT ); SP--;

```

Accordingly, we translate:

`code signal (e);  $\rho$  = codeR e  $\rho$`   
`signal`



```
if (0 ≤ tid = dequeue ( S[SP]))  
    enqueue ( RQ, tid );  
SP--;
```



Analogously:

`code broadcast (e); ρ = codeR e ρ`  
`broadcast`

where the instruction `broadcast` enqueues all threads from the queue `WQ` into the ready-queue `RQ` :

```
while (0 ≤ tid = dequeue ( S[SP]))
    enqueue ( RQ, tid );
SP--;
```

**Warning:**

The re-animated threads are not `blocked` !!!

When they become running, though, they first have to acquire their mutex :-)

## 48 Example: Semaphores

A semaphore is an abstract datatype which controls the access of a bounded number of (identical) resources.

### Operations:

- `Sema * newSema (int n)` — creates a new semaphore;
- `void Up (Sema * s)` — increases the number of free resources;
- `void Down (Sema * s)` — decreases the number of available resources.

Therefore, a semaphore consists of:

- a **counter** of type **int**;
- a mutex for synchronizing the semaphore operations;
- a condition variable.

```
typedef struct {  
    Mutex * me;  
    CondVar * cv;  
    int count;  
} Sema;
```

```
Sema * newSema (int n) {  
    Sema * s;  
    s = (Sema *) malloc (sizeof (Sema));  
    s→me = newMutex ();  
    s→cv = newCondVar ();  
    s→count = n;  
    return (s);  
}
```

The translation of the body amounts to:

alloc 1	newMutex	newCondVar	loadr 1	loadr 2
loadc 3	loadr 2	loadr 2	loadr 2	storer -2
new	store	loadc 1	loadc 2	return
storer 2	pop	add	add	
pop		store	store	
		pop	pop	

The function `Down()` **decrements** the counter.

If the counter becomes negative, `wait` is called:

```
void Down (Sema * s) {  
    Mutex *me;  
    me = s→me;  
    lock (me);  
    s→count--;  
    if (s→count < 0) wait (s→cv,me);  
    unlock (me);  
}
```

The translation of the body amounts to:

alloc 1	loadc 2	add		loadc 1
loadr 1	add	store		add
load	load	loadc 0		load
storer 2	loadc 1	le		wait
lock	sub	jumpz A	A:	loadr 2
	loadr 1	loadr 2		unlock
loadr 1	loadc 2	loadr 1		return

The function `Up()` increments the counter again.

If it is afterwards **not yet positive**, there still must exist waiting threads. One of these is sent a signal:

```
void Up (Sema * s) {  
    Mutex *me;  
    me = s→me;  
    lock (me);  
    s→count++;  
    if (s→count ≤ 0) signal (s→cv);  
    unlock (me);  
}
```



The translation of the body amounts to:

alloc 1	loadc 2	add	loadc 1
loadr 1	add	store	add
load	load	loadc 0	load
storer 2	loadc 1	le	signal
lock	add	jumpz A	A: loadr 2
	loadr 1		unlock
loadr 1	loadc 2	loadr 1	return

## 49 Stack-Management

### Problem:

- All threads live within the same storage.
- Every thread requires its own stack (at least conceptually).

### 1. Idea:

Allocate for each new thread a **fixed amount** of storage space.



Then we implement:


```
void *newStack() { return malloc(M); }  
void freeStack(void *adr) { free(adr); }
```

## Problem:

- Some threads consume much, some only little stack space.
- The necessary space is statically typically unknown :-)

## 2. Idea:

- Maintain all stacks in one joint **Frame-Heap** **FH** :-)
- Take care that the space inside the stack frame is sufficient at least for the current function call.
- A global stack-pointer **GSP** points to the overall topmost stack cell ...



Allocation and de-allocation of a stack frame makes use of the run-time functions:

```
int newFrame(int size) {  
    int result = GSP;  
    GSP = GSP+size;  
    return result;  
}
```


```
void freeFrame(int sp, int size);
```

## Warning:

The de-allocated block may reside inside the stack :-)



We maintain a list of freed stack blocks :-)





This list supports a function


```
void insertBlock(int max, int min)
```

which allows to free single blocks.

- If the block is on top of the stack, we pop the stack immediately;
- ... together with the blocks below – given that these have already been marked as de-allocated.
- If the block is inside the stack, we merge it with neighbored free blocks:









## Approach:


We allocate a fresh block for every function call ...

## Problem:


When ordering the block **before** the call, we do not yet know the space consumption of the called function :-)



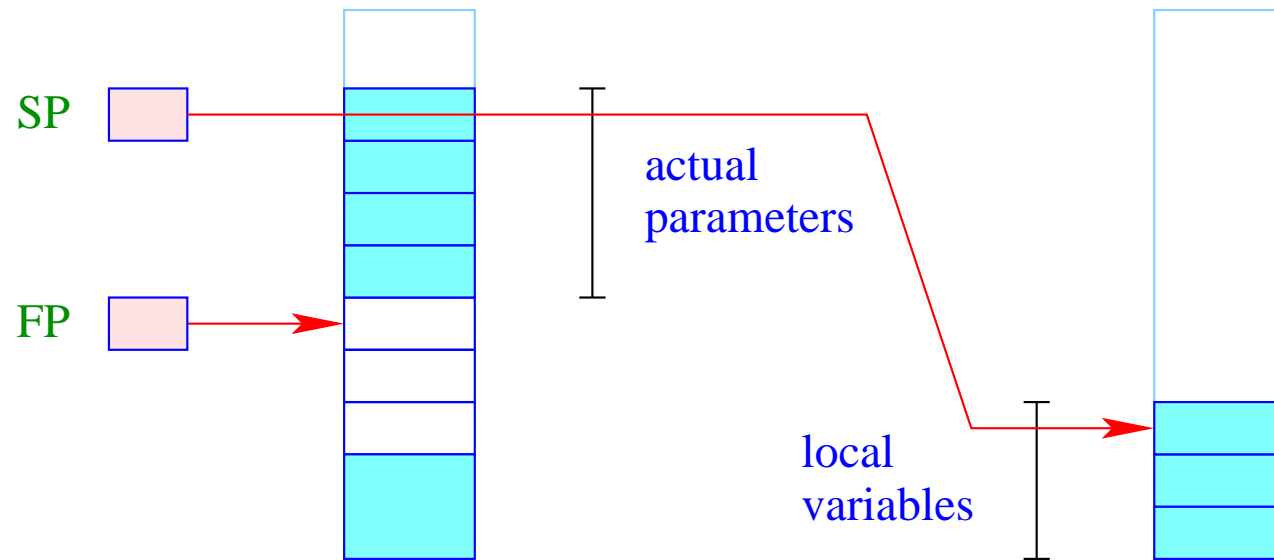
We order the new block **after** entering the function body!



Organisational cells as well as actual parameters must be allocated inside the old block ...



When entering the new function, we now allocate the new block ...



In particular, the **local** variables reside in the new block ...




We address ...

- the formal parameters **relatively** to the frame-pointer;
- the local variables **relatively** to the stack-pointer :-)

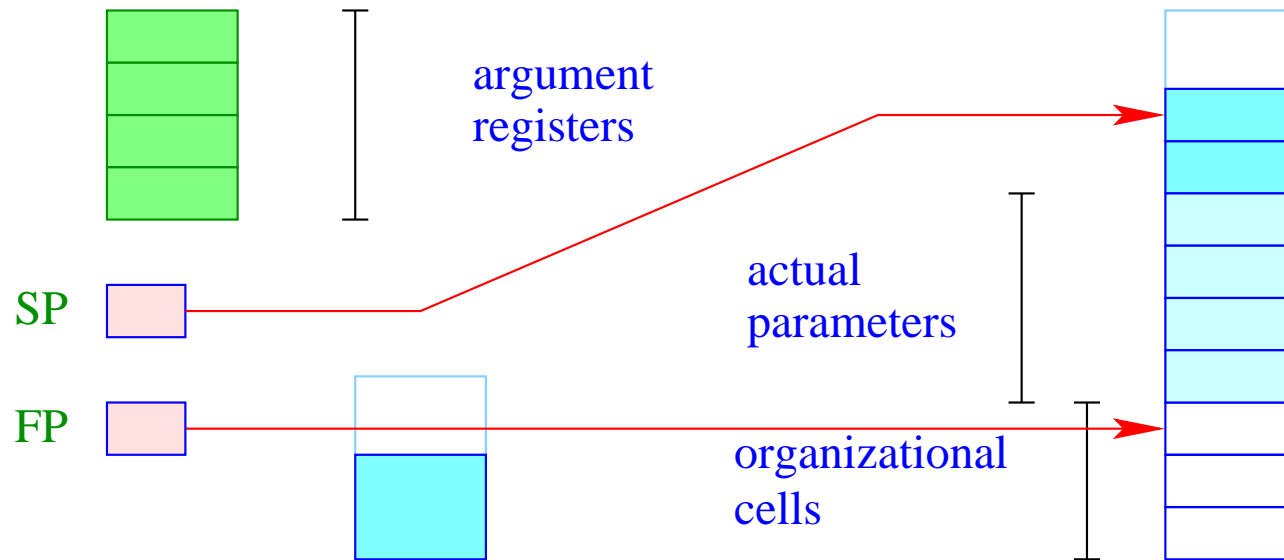


We must re-organize the complete code generation ... :-)

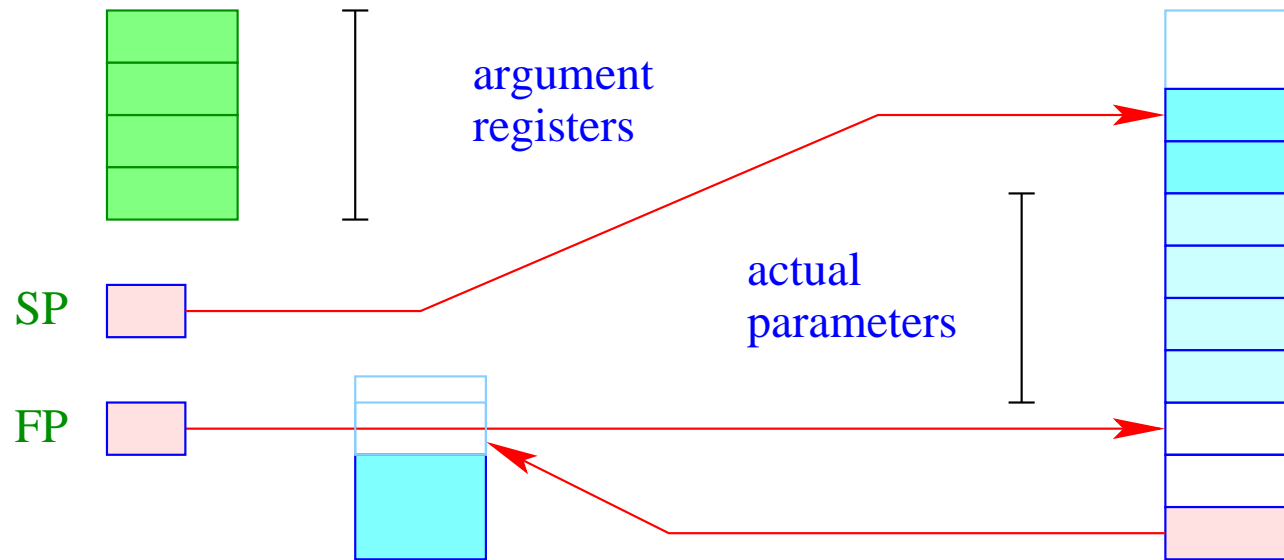
**Alternative:** Passing of parameters in registers ... :-)



The values of the actual parameters are determined **before** allocation of the new stack frame.



The **complete** frame is allocated inside the new block – plus the space for the current parameters.



Inside the new block, though, we must store the old **SP** (possibly +1) in order to correctly return the result ... :-)



### 3. Idea: Hybrid Solution

- For the first  $k$  threads, we allocate a separate stack area.
- For all further threads, we successively use one of the existing ones !!!



- For few threads extremely **simple** and **efficient**;
- For many threads **amortized** storage usage :-))