

**Script** generated by TTT

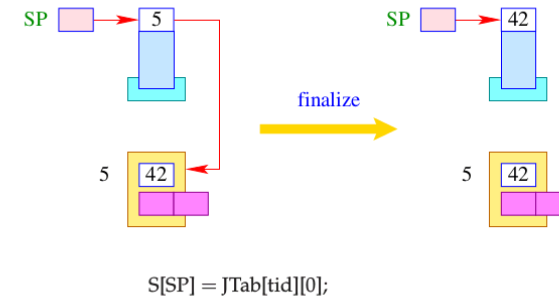
Title: Seidl: Virtual\_Machines (06.07.2015)

Date: Mon Jul 06 10:17:04 CEST 2015

Duration: 89:38 min

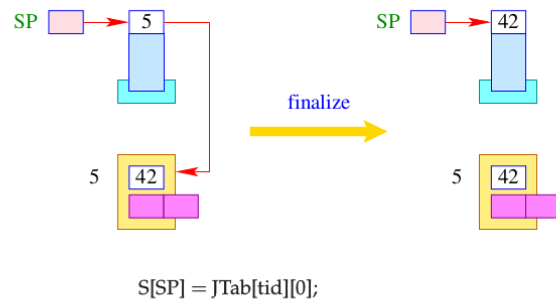
Pages: 62

... accordingly:



430

... accordingly:



430

Thus, we translate:

$code_R \text{ join } (e) \rho = code_R e \rho$   
join  
finalize

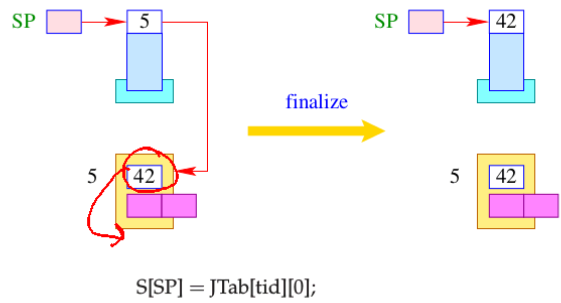
... where the instruction **join** is defined by:

```

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


429

... accordingly:



430

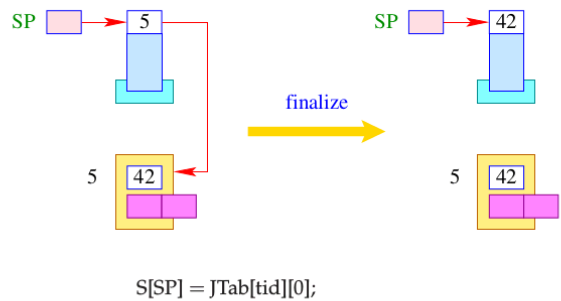
Example:



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.


428

... accordingly:



430

The instruction sequence:



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` at offset 0;
- ... 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:

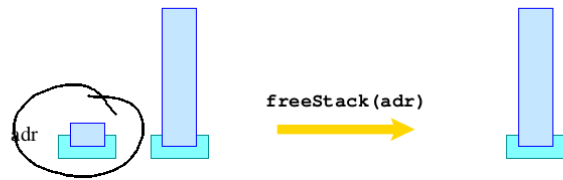
431

```

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

```

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



432

```

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

```

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



432

## 53 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:

```

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

```

**Warning:**

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

433

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.




434

Then we translate:

`codeR newMutex () ρ = newMutex`

where:




435

Then we translate:

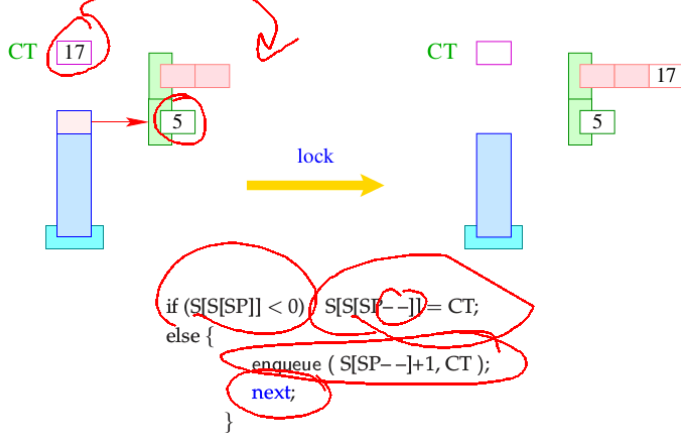
`code lock (e); ρ = codeR e ρ`  
`lock`

where:



436

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




437

Accordingly, we translate:

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

where:




438


Accordingly, we translate:

```
code unlock (e); p = code; e p
unlock
```

where:




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



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

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



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

## 54 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:

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

441

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

Therefore, we translate:

```

code wait (e0, e1); ρ = codeR e1 ρ
                        codeR e0 ρ
                        wait
                        dup
                        unlock
                        next
                        lock
    
```

where ...

443

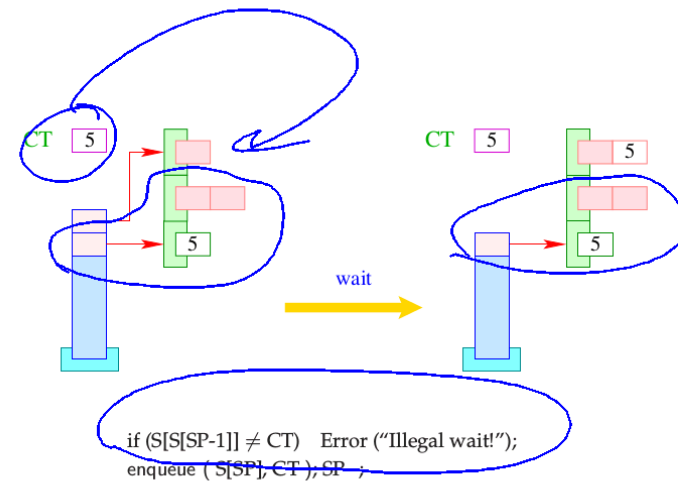
Then we translate:

```
codeR newCondVar () ρ = newCondVar
```

where:



442

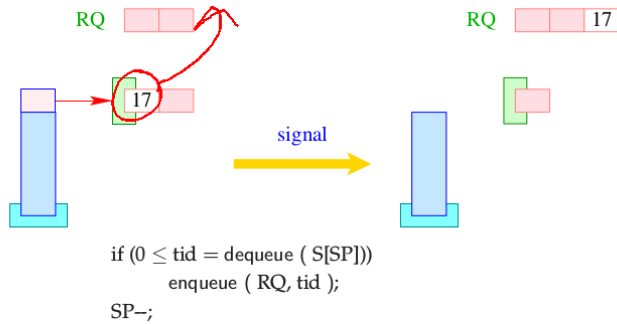


444

Accordingly, we translate:

$$\text{code signal } (e); \rho = \text{code}_{R} e \rho$$

signal



445

## 55 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.

447

Analogously:

$$\text{code broadcast } (e); \rho = \text{code}_{R} e \rho$$

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 :-)

446

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;
```

448

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 -2 | loadr 1   |
| loadc 3  | loadr 1  | loadr 1    | loadr 1  | storer -2 |
| new      | store    | loadc 1    | loadc 2  | return    |
| storer 1 | pop      | add        | add      |           |
| pop      |          | store      | store    |           |
|          |          | pop        | pop      |           |

*loadrc 1*  
*loadc 0*  
*add*

The translation of the body amounts to:

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



The translation of the body amounts to:

|          |          |          |            |
|----------|----------|----------|------------|
| alloc 1  | add      | loadc 0  | wait       |
| loadr -2 | load     | less     | dup        |
| load     | loadc 1  | jumpz A  | unlock     |
| storer 1 | sub      | loadr 1  | next       |
| lock     | loadr -2 | loadr -2 | lock       |
|          | loadc 2  | loadc 1  | A: loadr 1 |
| loadr -2 | add      | add      | unlock     |
| loadc 2  | store    | load     | return     |

452

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);  
}
```

451

The translation of the body amounts to:

|          |          |          |            |
|----------|----------|----------|------------|
| alloc 1  | add      | loadc 0  | wait       |
| loadr -2 | load     | less     | dup        |
| load     | loadc 1  | jumpz A  | unlock     |
| storer 1 | sub      | loadr 1  | next       |
| lock     | loadr -2 | loadr -2 | lock       |
|          | loadc 2  | loadc 1  | A: loadr 1 |
| loadr -2 | add      | add      | unlock     |
| loadc 2  | store    | load     | return     |

452

The function Down() decrements the counter.

If the counter becomes negative, wait is called:

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

451

The translation of the body amounts to:

|          |          |            |         |
|----------|----------|------------|---------|
| alloc 1  | add      | loadc 0    | wait    |
| loadr -2 | load     | less       | dup     |
| load     | loadc 1  | jumpz A    | unlock  |
| storer 1 | sub      | loadr 1    | next    |
| lock     | loadr -2 | loadr -2   | lock    |
|          | loadc 2  | loadc 1 A: | loadr 1 |
| loadr -2 | add      | add        | unlock  |
| loadc 2  | store    | load       | return  |

452

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);  
}
```

451

The translation of the body amounts to:

|          |          |            |         |
|----------|----------|------------|---------|
| alloc 1  | add      | loadc 0    | wait    |
| loadr -2 | load     | less       | dup     |
| load     | loadc 1  | jumpz A    | unlock  |
| storer 1 | sub      | loadr 1    | next    |
| lock     | loadr -2 | loadr -2   | lock    |
|          | loadc 2  | loadc 1 A: | loadr 1 |
| loadr -2 | add      | add        | unlock  |
| loadc 2  | store    | load       | return  |

452

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);  
}
```

451

The translation of the body amounts to:

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

454

## 56 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); }
```

455

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);
}
```

453

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);
}
```

453

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);
}
```

453

The translation of the body amounts to:

|                       |                       |                         |                      |
|-----------------------|-----------------------|-------------------------|----------------------|
| <code>alloc 1</code>  | <code>loadc 2</code>  | <code>add</code>        | <code>loadc 1</code> |
| <code>loadr -2</code> | <code>add</code>      | <code>store</code>      | <code>add</code>     |
| <code>load</code>     | <code>load</code>     | <code>loadc 0</code>    | <code>load</code>    |
| <code>storer 1</code> | <code>loadc 1</code>  | <code>leq</code>        | <code>signal</code>  |
| <code>lock</code>     | <code>add</code>      | <code>jumpz A A:</code> | <code>loadr 1</code> |
|                       | <code>loadr -2</code> |                         | <code>unlock</code>  |
| <code>loadr -2</code> | <code>loadc 2</code>  | <code>loadr -2</code>   | <code>return</code>  |

454

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);
}
```

453

The translation of the body amounts to:

|                       |                       |                         |                      |
|-----------------------|-----------------------|-------------------------|----------------------|
| <code>alloc 1</code>  | <code>loadc 2</code>  | <code>add</code>        | <code>loadc 1</code> |
| <code>loadr -2</code> | <code>add</code>      | <code>store</code>      | <code>add</code>     |
| <code>load</code>     | <code>load</code>     | <code>loadc 0</code>    | <code>load</code>    |
| <code>storer 1</code> | <code>loadc 1</code>  | <code>leq</code>        | <code>signal</code>  |
| <code>lock</code>     | <code>add</code>      | <code>jumpz A A:</code> | <code>loadr 1</code> |
|                       | <code>loadr -2</code> |                         | <code>unlock</code>  |
| <code>loadr -2</code> | <code>loadc 2</code>  | <code>loadr -2</code>   | <code>return</code>  |

454

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);
}
```

453

### 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 ...

456

## 56 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); }
```


455

### 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:


458

**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:



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:



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:

458

### 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!

462



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

464



Inparticular, the **local** variables reside in the new block ...

465

⇒ 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 ... :-)

466


⇒ 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 ... :-)

466



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

465