

Script generated by TTT

Title: Seidl: Virtual_Machines (26.06.2013)

Date: Wed Jun 26 16:03:20 CEST 2013

Duration: 87:05 min

Pages: 48

Therefore, we translate:


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

The instruction `term` is explained later :-)

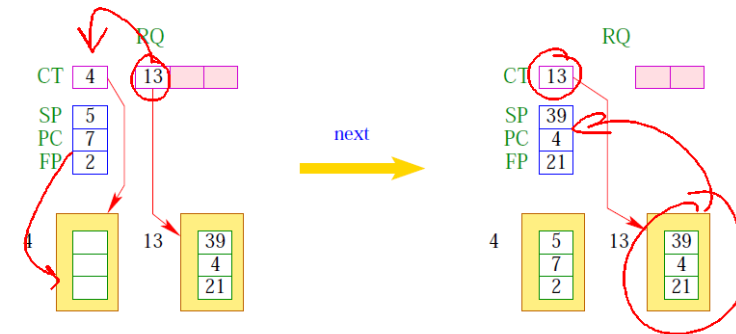
The instruction `exit` successively pops all stack frames:

```
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`.



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 ();
}


```

51 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:

```

codeR join (e) ρ = codeR e ρ
                    join
                    finalize

```

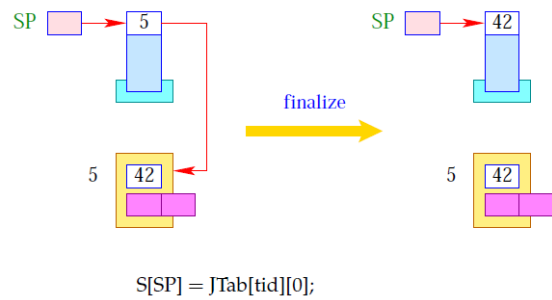
... where the instruction `join` is defined by:

```

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

```

... accordingly:



416

Thus, we translate:

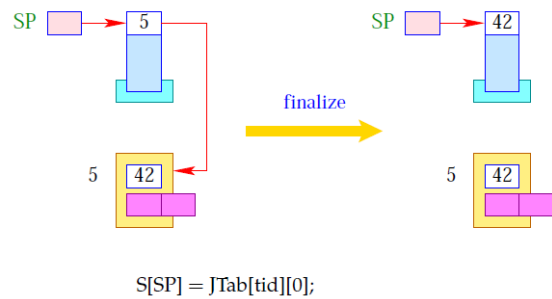
```
codeR join (e) ρ = codeR e ρ  
                    join  
                    finalize
```

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

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

415

... accordingly:



416

Thus, we translate:

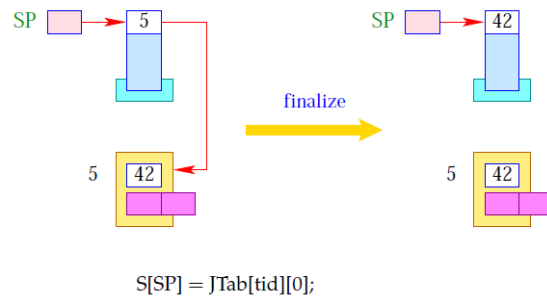
```
codeR join (e) ρ = codeR e ρ  
                    join  
                    finalize
```

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

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

415

... accordingly:



416

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` 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:

417

52 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 :-)

419




410

```

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` :



418

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` 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:


417

```

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` :



418

52 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 :-)

419

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.



420

Then we translate:

```
codeR newMutex () ρ = newMutex
```

where:




421

Then we translate:


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

where:



422

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[S[SP-]]+1, CT);
  next;
}
```

423

Accordingly, we translate:

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

where:




424

Accordingly, we translate:


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

where:



424


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 );
}
```

425


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 );
}
```

425

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 );
}

```

53 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:

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

Then we translate:

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

where:



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 ...

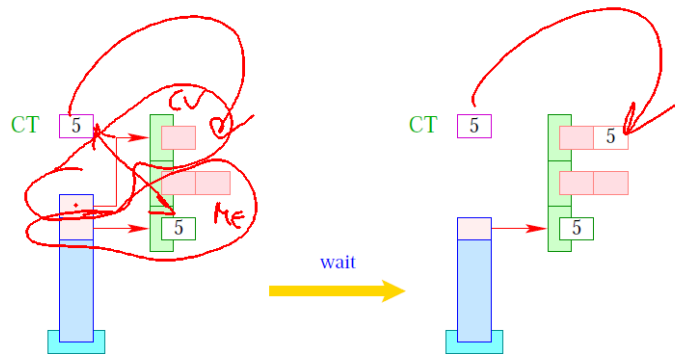
After enqueueing the current thread, we release the mutex. After re-animation, though, we must acquire the mutex again.

Therefore, we translate:

```

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

where ...



```

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


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 ...



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

430


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 ...

429




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

430

Accordingly, we translate:

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



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

431

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 :-)

432

54 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.
```

433

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;
```

434

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

435

```

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

```

435

The translation of the body amounts to:

```

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

```

436

The translation of the body amounts to:

```

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

```

436

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;

```

434

The translation of the body amounts to:

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