

# Exam Questions

## Proof Methods for Concurrent Programs

11 February 2010

*Instructions: only printed documents are authorised. You can admit the result of one question and move on. Leave optional questions until the end.*


Your goal is to implement and prove correct some support functions used by the *A-Steal* multiprocessor scheduling algorithm.

### Exercise 1. (Data structures)

A *double ended queue* (often abbreviated to *deque*) is a data structure that implements a queue for which elements can be added to or removed from the front (head) or back (tail). We follow the C++ interface (`std::deque`) to name the different operations:

<code>push_back</code>	insert element at back
<code>push_front</code>	insert element at front
<code>pop_back</code>	remove element from back
<code>pop_front</code>	remove element from front

We use *doubly-linked lists* to implement deques. A doubly-linked list can be graphically represented as



where `f` and `b` are the pointers to the front and back elements. When the doubly-linked list is empty, `f = b = null`. An implementation of the `push_front` procedure is given below:

```
push_front(a,f,b) {
  t := new (a,f,null);
  if (f = null) then b := t else [f+2] := t;
  f := t;
}
```

Remark that our notation for procedures assumes that the arguments (e.g. `a`, `f`, `b`) are passed by reference (in other terms, they can be considered global variables), while the other variables (e.g. `t`) are local.

1. Give a sequential implementation to the `pop_back(f,b,r)` and `pop_front(f,b,r)` procedures:
  - `f` and `b` are as above, while `r` is used to return either the popped value or `null` if the deque is empty;
  - your code should not leak memory.

*Answer.*

```
pop_front (f,b,r) {
  if (f = null)
  then r := null
  else {
    t := f;
    f := [f+1];
    if (f = null)
```

```

        then b := null
        else [f+2] := null;
    r := [t];
    dispose (t,t+1,t+2);
}
}

pop_back (f,b,r) {
    if (b = null)
    then r := null
    else {
        t := b;
        b := [b+2];
        if (b = null)
        then f := null
        else [b+1] := null;
        r := [t];
        dispose (t,t+1,t+2);
    }
}
}

```

Let  $\alpha$  range over lists of values ( $\epsilon$  denotes the empty list and  $\cdot$  is the concatenation operator). Recall the recursive specification of doubly-linked list segments:

$$\begin{aligned} \text{dlseg } \epsilon (f, f', b', b) &= \text{empty} \wedge f = b' \wedge f' = b \\ \text{dlseg } (a \cdot \alpha) (f, f', b', b) &= \exists j. f \mapsto a, j, f' * \text{dlseg } \alpha (j, f, b', b) \end{aligned}$$

and consider the definition of doubly-linked lists below:

$$\text{dls } \alpha (f, b) = \text{dlseg } \alpha (f, \text{null}, \text{null}, b)$$

Intuitively  $\text{dls } \alpha (f, b)$  should be read “ $f, b$  are the ends of a doubly-linked list representing the list of values  $\alpha$ ”.

2. Prove the following properties:

- (a.)  $\text{dls } (a \cdot \epsilon) (f, b) \Leftrightarrow f \mapsto a, \text{null}, \text{null} \wedge b = f$
- (b.)  $\text{dls } \alpha (f, b) \Rightarrow (f = \text{null} \Rightarrow (\alpha = \epsilon \wedge b = \text{null}))$
- (c.)  $\text{dls } \alpha (f, b) \Rightarrow (f \neq \text{null} \Rightarrow (\exists a, \alpha'. \alpha = a \cdot \alpha'))$

*Answer.*

(a.)

$$\begin{aligned} \text{dls } a (f, b) &\Leftrightarrow \text{dlseg } a (f, \text{null}, \text{null}, b) \\ &\Leftrightarrow \exists j. f \mapsto a, j, \text{null} * \text{dlseg } \epsilon (j, f, \text{null}, b) \\ &\Leftrightarrow \exists j. f \mapsto a, j, \text{null} * (\text{empty} \wedge j = \text{null} \wedge f = b) \\ &\Leftrightarrow f \mapsto a, \text{null}, \text{null} \wedge f = b \end{aligned}$$

(b.) Suppose that  $(s, h) \vdash \text{dls } \alpha (f, b)$  and  $f = \text{null}$ . Since  $f = \text{null}$ , it cannot exist a  $j$  and a  $h' \subseteq h$  such that  $(s, h') \vdash f \mapsto \_, j, \_$ . It must then hold that  $(s, h) \vdash \text{empty} \wedge f = \text{null} \wedge b = \text{null}$ . We conclude that  $(\text{dls } \alpha (f, b) \wedge f = \text{null}) \Rightarrow (\text{empty} \wedge b = \text{null})$  and the result follows by uncurrying.

(c.) Suppose that  $(s, h) \vdash \text{dls } \alpha (f, b)$  and  $f \neq \text{null}$ . Since  $\text{dls } \alpha (f, b) = \text{dlseg } \alpha (f, \text{null}, \text{null}, b)$  and  $f \neq \text{null}$ , it must exist  $j, a, \alpha'$  such that  $(s, h) \vdash f \mapsto a, j, f' * \text{dlseg } \alpha' (j, f, \text{null}, b)$ . By the definition of  $\text{dlseg}$ , we deduce  $(s, h) \vdash \text{dlseg } a \cdot \alpha' (f, \text{null}, \text{null}, b)$ , and in turn  $(s, h) \vdash \text{dls } a \cdot \alpha' (f, b)$ . Thus,  $\alpha = a \cdot \alpha'$  for some  $a, \alpha'$ .

3. Prove that the operations on deques satisfy the specifications below:

- |               |  |                                  |  |
|---------------|--|----------------------------------|--|
| (a.)          | $\{\text{dls } \alpha \text{ (f, b)}\}$                  | <code>push_front(a, f, b)</code> | $\{\text{dls (a} \cdot \alpha \text{) (f, b)}\}$                   |
| (b.)          | $\{\text{dls } \epsilon \text{ (f, b)}\}$                | <code>pop_front(f, b, r)</code>  | $\{\text{dls } \epsilon \text{ (f, b) } \wedge \text{ r = null}\}$ |
| (c.)          | $\{\text{dls (a} \cdot \alpha \text{) (f, b)}\}$         | <code>pop_front(f, b, r)</code>  | $\{\text{dls } \alpha \text{ (f, b) } \wedge \text{ r = a}\}$      |
| (d. optional) | $\{\text{dls } \epsilon \text{ (f, b)}\}$                | <code>pop_back(f, b, r)</code>   | $\{\text{dls } \epsilon \text{ (f, b) } \wedge \text{ r = null}\}$ |
| (e. optional) | $\{\text{dls } (\alpha \cdot \text{a}) \text{ (f, b)}\}$ | <code>pop_back(f, b, r)</code>   | $\{\text{dls } \alpha \text{ (f, b) } \wedge \text{ r = a}\}$      |

Answer.

- (a.)  $\{\text{dls } \alpha \text{ (f, b)}\} \text{ push\_front(a, f, b) } \{\text{dls } \text{a} \cdot \alpha \text{ (f, b)}\}.$

```

{dls } \alpha \text{ (f, b)}
  t := new (a, f, null);
{t \mapsto a, f, null * dls } \alpha \text{ (f, b)}
  if (f = null) then
    {t \mapsto a, f, null * dls } \alpha \text{ (f, b) } \wedge \text{ f = null}
    {t \mapsto a, f, null \wedge b = null \wedge f = null \wedge } \alpha = \epsilon\}
    {t \mapsto a, null, null \wedge b = null \wedge f = null \wedge } \alpha = \epsilon\}
    b := t
    {t \mapsto a, null, null \wedge b = t \wedge f = null \wedge } \alpha = \epsilon\}
  {dls } \text{a} \cdot \alpha \text{ (t, b)}
  else
    {t \mapsto a, f, null * dls } \alpha \text{ (f, b) } \wedge \text{ f } \neq \text{null}
    {t \mapsto a, f, null * dls } \alpha \text{ (f, b) } \wedge \text{ f } \neq \text{null} \wedge } \alpha = \text{a}' \cdot \alpha'\}
    {t \mapsto a, f, null * \exists j. f \mapsto \text{a}', j, null * dlseg } \alpha' \text{ (j, f, null, b)}
    [f+2] := t;
    {t \mapsto a, f, null * \exists j. f \mapsto \text{a}', j, t * dlseg } \alpha' \text{ (j, f, null, b)}
  {dls } \text{a} \cdot \alpha \text{ (t, b)}
  f := t;
  {dls } \text{a} \cdot \alpha \text{ (f, b)}

```

- (b.)

```

{dls } \epsilon \text{ (f, b)}
  if (f = null) then
    {dls } \epsilon \text{ (f, b) } \wedge \text{ f = null}
    r := null
    {dls } \epsilon \text{ (f, b) } \wedge \text{ f = null} \wedge \text{ r = null}
  else {
    {dls } \epsilon \text{ (f, b) } \wedge \text{ f } \neq \text{null}
    {\exists \text{a}, \alpha. } \epsilon = \text{a} \cdot \alpha\}
    {false}
    t := f;
    f := [f+1];
    if (f = null)
      then b := null
    else [f+2] := null;
    r := [t];
    dispose (t, t+1, t+2);
  }
  {dls } \epsilon \text{ (f, b) } \wedge \text{ f = null} \wedge \text{ r = null}
  {dls } \epsilon \text{ (f, b) } \wedge \text{ r = null}

```

(c.) We prove separately the case where  $\alpha = \epsilon$  and the case where  $\alpha = \mathbf{a}' \cdot \alpha'$  for some  $\mathbf{a}', \alpha'$ . If  $\alpha = \epsilon$  we can derive:

```

{dls a · ε (f, b)}
  if (f = null) then
{false}
  r := null
  else {
{dls a · ε (f, b) ∧ f ≠ null}
{f ↦ a, null, null ∧ b = f}
  t := f;
{t ↦ a, null, null ∧ b = f = t}
  f := [f+1];
{t ↦ a, null, null ∧ b = t ∧ f = null}
  if (f = null) then
{t ↦ a, null, null ∧ b = t ∧ f = null}
  b := null
{t ↦ a, null, null ∧ b = null ∧ f = null}
  else
{false}
  [f+2] := null;
{t ↦ a, null, null ∧ b = null ∧ f = null}
  r := [t];
{t ↦ a, null, null ∧ b = null ∧ f = null ∧ r = a}
  dispose (t, t+1, t+2);
{empty ∧ b = null ∧ f = null ∧ r = a}
{dls ε (f, b) ∧ r = a}
}
{dls ε (f, b) ∧ r = a}

```

If  $\alpha = a' \cdot \alpha'$  for some  $a', \alpha'$ , we can derive:

```

{dls a · a' · α' (f, b)}
  if (f = null) then
{false}
  r := null
  else {
{dls a · a' · α' (f, b) ∧ f ≠ null}
{∃j. f ↦ a, j, null * dlseg a' · α' (j, f, null, b)}
  t := f;
{∃j. t ↦ a, j, null * dlseg a' · α' (j, t, null, b) ∧ f = t}
  f := [f+1];
{t ↦ a, j, null * dlseg a' · α' (j, t, null, b) ∧ f = j}
  if (f = null) then
{false}
  b := null
  else
{t ↦ a, j, null * dlseg a' · α' (j, t, null, b) ∧ f = j}
{t ↦ a, f, null * dlseg a' · α' (f, t, null, b)}
  [f+2] := null;
{t ↦ a, f, null * dlseg a' · α' (f, null, null, b)}
{t ↦ a, f, null * dls a' · α' (f, b)}
  r := [t];
{t ↦ a, f, null * dls a' · α' (f, b) ∧ r = a}
  dispose (t, t+1, t+2);
{dls a' · α' (f, b) ∧ r = a}
}
{dls a' · α' (f, b) ∧ r = a}

```

4. Recall the definition of *precise predicate*. Is  $\text{dlseg } \alpha (f, f', b', b)$  a precise predicate?

*Answer.* If  $f = b$  then the doubly-linked-list segment might be empty (thus denoting the empty heap), or circular (thus denoting a non-empty heap). This implies that  $\text{dlseg } \alpha (f, f', b', b)$  is not precise. It can be shown that if  $f' = b' = \text{null}$  then the segment is non-touching:  $\text{dls } \alpha (f, b)$  is precise (however the proof is far from easy - see Sec. 4.3 of Reynold's lectures).

In what follows you can assume that  $\text{dls } \alpha (f, b)$  is precise.

### Exercise 2. (The *A-Steal* scheduler)

The *A-Steal* algorithm implements task scheduling for several processors. A separate deque with the thread-identifiers (**tid**s) of the threads to be executed is maintained for each processor. To execute the next thread, the processor gets the first element from the deque (using the **pop\_front** deque operation). If the current thread forks, it is put back to the front of the deque (**push\_front**) and a new thread is executed. When one of the processors finishes execution of its own threads (i.e. its deque is empty), it can *steal* a thread from another processor: it gets the last element from the deque of another processor (**pop\_back**) and executes it.

We introduce a handy C-like notation for arrays: if **a** points to a series of  $n$  contiguous memory locations, we write  $\mathbf{a}[i]$  for  $\mathbf{a}+i$ .

Consider a system with  $n$  processors, indexed from 0 to  $n - 1$ . To implement the A-Steal algorithm we use two arrays, named **a** and **l**, and  $n$  deques. The array **a** has  $2n$  entries, and  $\mathbf{a}[2*i]$  and  $\mathbf{a}[2*i+1]$  are the front and back addresses of the deque associated to the processor  $i$ . The elements of the deques are **tid**s (you can assume that a **tid** is just an integer). The array **l** has  $n$  entries, and each  $\mathbf{l}[i]$  is a resource that protects the locations  $\mathbf{a}[2*i]$  and  $\mathbf{a}[2*i+1]$  and the associated deque.

1. Implement the **fork** and **schedule** procedures of the A-Steal algorithm, according to their informal specification below:

**fork**(**proc**,**tid**) stores **tid** at the front of the deque of processor **proc**;

**schedule**(**proc**):**tid** pops (and returns) the **tid** at the front of the deque of processor **proc**; if the deque of processor **proc** is empty, it pops (and returns) the **tid** at the back of the deque of another processor which has a non-empty deque.

The procedures **fork** and **schedule** should not fail (**schedule** might loop), and, needless to say, should be robust to concurrent invocations.

*Answer.*

```
fork (proc,tid) {
  with l[proc] do {
    push_front (tid, a[2*proc], a[2*proc+1]);
  }
}

schedule (proc) {
  with l[proc] do {
    tid := pop_front (a[2*proc], a[2*proc+1]);
  };
  while (tid = null) do {
    proc = (proc+1) mod n;
    with l[proc] do {
      tid := pop_back (a[2*proc], a[2*proc+1]);
    }
  }
}
```

2. Define the resource invariant associated to each resource  $l[i]$ .

*Answer.* The resource invariant of the resource  $l[i]$  is

$$\exists \alpha. \text{dls } \alpha (a[2 * i], a[2 * i + 1])$$

3. Prove that

$$\begin{aligned} \text{(a.)} \quad & \{0 \leq \text{proc} \leq n - 1\} \text{fork}(\text{proc}, \text{tid}) \{true\} \\ \text{(b.)} \quad & \{0 \leq \text{proc} \leq n - 1\} \text{schedule}(\text{proc}) \{\text{tid} \neq \text{null}\} \end{aligned}$$

*Answer.*

(a.)

$$\begin{aligned} & \{0 \leq \text{proc} \leq n - 1\} \\ & \quad \text{with } l[\text{proc}] \text{ do } \{ \\ & \{0 \leq \text{proc} \leq n - 1 \wedge \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1])\} \\ & \quad \text{push\_front } (\text{tid}, a[2 * \text{proc}], a[2 * \text{proc} + 1]); \\ & \{ \text{dls } \text{tid} \cdot \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1]) \} \\ & \{ \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1]) \} \\ & \quad \} \\ & \{true\} \end{aligned}$$

(b.)

$$\begin{aligned} & \{0 \leq \text{proc} \leq n - 1\} \\ & \quad \text{with } l[\text{proc}] \text{ do } \{ \\ & \{0 \leq \text{proc} \leq n - 1 \wedge \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1])\} \\ & \quad \text{tid} := \text{pop\_front } (a[2 * \text{proc}], a[2 * \text{proc} + 1]); \\ & \{0 \leq \text{proc} \leq n - 1 \wedge \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1]) \wedge (\text{tid} = a \vee \text{tid} = \text{null})\} \\ & \quad \}; \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null})\} \\ & \quad \text{while } (\text{tid} = \text{null}) \text{ do } \{ \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null})\} \\ & \quad \text{proc} = (\text{proc} + 1) \bmod n; \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null})\} \\ & \quad \text{with } l[\text{proc}] \text{ do } \{ \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null} \wedge \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1])\} \\ & \quad \text{tid} := \text{pop\_back } (a[2 * \text{proc}], a[2 * \text{proc} + 1]); \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null} \wedge \exists \alpha. \text{dls } \alpha (a[2 * \text{proc}], a[2 * \text{proc} + 1])\} \\ & \quad \} \\ & \{0 \leq \text{proc} \leq n - 1 \wedge (\exists t. \text{tid} = t \vee \text{tid} = \text{null})\} \\ & \quad \} \\ & \{0 \leq \text{proc} \leq n - 1 \wedge \exists t. \text{tid} = t\} \end{aligned}$$