Semaphores. Limits and Extensions CS 3SD3

Ryszard Janicki

Department of Computing and Software, McMaster University, Hamilton, Ontario, Canada

・ロン ・回 と ・ 回 と ・ 回 と

3

- Semaphore *s* is an integer variable that can take only non-negative values.
- The only semaphore operations are down(s) (wait(s), V(s)) and up(s) (signal(s), P(s)).

• Semaphores should be atomic and 'easy/fast/efficient' to implement, preferably at a low level.

イロト イヨト イヨト イヨト

Mutual Exclusion

```
var s: semaphore = 1;

P_1 : cycle begin P_1-instructions<sub>1</sub>;

down(s);

critical region;

up(s);

P_1-instructions<sub>2</sub>;

end
```

```
P_2: cycle begin P_2-instructions<sub>1</sub>;
down(s);
critical region;
up(s);
P_2-instructions<sub>2</sub>;
end
```

白 と く ヨ と く ヨ と …

æ

Theory of Semaphores

• Properties of Dijkstra's semaphore operations may be characterized in the following way. Let:

C(s) - the initial value of a semaphore variable sndown(s) - the number of times down(s) was executed nup(s) - the number of times up(s) was executed npdown(s) - the number of times down(s) was passed

• Using these notions we may define the results of actions *down* and *up* as follows:

 $\begin{array}{l} \textit{down}(s): \textit{ndown}(s) = \textit{ndown}(s) + 1:\\ \textit{if }\textit{ndown}(s) \leq \textit{nup}(s) + \textit{C}(s) \textit{ then }\textit{npdown}(s) = \textit{npdown}(s) + 1;\\ \textit{up}(s): \textit{if }\textit{ndown}(s) > \textit{nup}(s) + \textit{C}(s)\\ \textit{then }\textit{npdowns}(s) = \textit{npdowns}(s) + 1; \textit{nup}(s) = \textit{nup}(s) + 1; \end{array}$

Theorem/Definition (Semaphore Invariant)

npdown(s) = min(ndown(s), C(s) + nup(s))

• Anything that satisfies the equation above is a Dijkstra's semaphore!

(국물) 국물) 문

Three smokers are sitting at a table. One of them has **tobacco**, another **cigarette paper**, and the third one has **matches** - each one has a different ingredient required to make and smoke a cigarette but he may not give any ingredient to another. On the table in front of them, two of the three ingredients will be placed, and the smoker who has the necessary third ingredient should pick up the ingredients from the table, make a cigarette and smoke it.

Since a new set of ingredients will not be placed on the table until this action (i.e. smoking) is completed, the other smoker who cannot make and smoke a cigarette with the ingredients on the table, must not interfere with the fellow who can. Therefore, coordination is needed among smokers.

(4月) (1日) (日)

The actions of the smokers without coordination

- X the smoker with tobacco
- α_x : pick up the paper pick up the match roll the cigarette light the cigarette smoke the cigarette *goto* α_x
- Y the smoker with paper
- α_y : pick up the tobacco pick up the match roll the cigarette light the cigarette smoke the cigarette goto α_y
- Z the smoker with matches
- α_z : pick up the tobacco pick up the paper roll the cigarette light the cigarette smoke the cigarette *goto* α_z

回 と く ヨ と く ヨ と

'Obvious Solution' with Semaphores

| Rtobacco | Rpaper | Rmatch |] |
|---------------------|--------------------------|----------------------------|------------------------|
| rt: down(s); | rp: <i>down(s</i>); | rm: <i>down</i> (s); | |
| up(paper); | up(match); | up(tobacco); | $ $ \leftarrow agent |
| up(match); | up(tobacco); | up(paper); | |
| <i>goto</i> rt | <i>goto</i> rp | <i>goto</i> rm | |
| Smoker with Tobacco | Smoker with Paper | Smoker with Matches | s |
| at: down(paper); | ap: <i>down(match)</i> ; | am: <i>down(tobacco</i>); | m |
| down(match); | down(tobacco); | down(paper); | ←o |
| | | | k |
| up(smokert); | up(smokerp); | up(smokerm); | e |
| <i>goto</i> at | <i>goto</i> ap | <i>goto</i> am | rs |
| bt: down(smokert); | bp: down(smokerp); | bm: <i>down(smokerm</i>); | |
| up(s); | up(s); | up(s); | |
| <i>goto</i> bt | <i>goto</i> bp | <i>goto</i> bm | |
| | * | | - |

Rule: The set of new ingredients will not be placed on the table until an appropriate action of smoking is completed

< 注→ …

- Deadlocking sequence (not unique): *down(s)* in RTobacco → *up(s)* in RTobacco → *down(paper)* in Smoker with Tobacco → *up(matches)* in RTobacco → *down(matches)* in Smoker with Paper
- paper and matches have been put on the table and the process Smoker with Tobacco takes paper while the process Smoker with Paper takes matches!

Problem (Smokers' Problem)

The Smokers' Problem is to define additional semaphores and processes and to introduce appropriate down and up statements so as to make them deadlock-free. No alternation can, however, be made to the processes defining the agent. No conditional statement and assignment statement instructions may be used.

Theorem (Patil 1970)

The Smokers' Problem has no solution.

Parnas (1974) solution to Smokers' Problem!?

initially: s = mutes = 1, t = tobacco = paper = match = smokert = smokerp = smokerm = 0, S[1] = S[2] = S[3] = S[4] = S[5] = S[6] = 0

| 3110 Kerm = 0, 5[1] = 5[2] | = 3[3] = 3[4] = 3[3] = | | |
|-----------------------------|----------------------------------|----------------------------------|------------------------|
| Rtobacco | Rpaper | Rmatch | |
| rt: down(s); | rp: <i>down(s</i>); | rm: <i>down</i> (<i>s</i>); | |
| up(paper); | up(match); | up(tobacco); | \leftarrow agent |
| up(match); | up(tobacco); | up(paper); | |
| <i>goto</i> rt | <i>goto</i> rp | <i>goto</i> rm | |
| bt: down(smokert); | <pre>bt: down(smokerp);</pre> | bt: down(smokerm); | const- |
| up(s); | up(s); | up(s); | \leftarrow raints |
| <i>goto</i> bt | <i>goto</i> bp | <i>goto</i> bm | |
| Smoker with Tobacco | Smoker with Paper | Smoker with Matches | s |
| at: down(S[6]); | ap: <i>down</i> (S [5]); | am: <i>down</i> (<i>S</i> [3]); | m |
| t = 0; | t = 0; | t = 0; | $\leftarrow o$ |
| | | | k |
| up(smokert); | up(smokerp); | up(smokerm); | e |
| goto at | <i>goto</i> ap | goto am | rs |
| dt: <i>down(tobacco</i>); | dp: <i>down(paper)</i> ; | dm: <i>down(match</i>); | |
| down(mutex); | down(mutex); | down(mutex); | push- |
| t = t + 1; | t = t + 2; | t=t+4; | ←ers |
| up(S[t]); | up(S[t]); | up(S[t]); | |
| up(mutex); | up(mutex); | up(mutex); | |
| <i>goto</i> dt | <i>goto</i> dp | <i>goto</i> dm | |
| d1: down(S[1]); | d2: <i>down(S</i> [2]); | d3: <i>down</i> (<i>S</i> [4]); | no over- |
| goto d1 | <i>goto</i> d2 | goto d3⊐ → ∢ @ → ∢ ≣ → | ≪≣ flow ≘ 🧠 |
| | Pyczard Janicki Som | anhouse Limits and Extensions | 0./20 |

- *t* is just and integer variable, not a semaphore variable.
- The array *S*[...] is *an array of semaphores*, a non-standard construction, very seldom implemented, however formally OK.

• 3 > 1

Both!

- Parnas, if the *letter of law* is more important.
- Patil, if the *spirit of law* is more important.
- Unfortunately, Patil's paper was not well written and has many *implicit* assumptions that were not spelled out.
- It was implicitly assumed: no semaphore extension to arrays, no extension to many variables, no assignment statements, etc.

- 4 同 6 4 日 6 4 日 6

Multidimentional Semaphores of Agerwala

• The extended primitives *edown* and *eup* are atomic (indivisible) and each operates on a set of semaphore variables which must be initiated with non-negative integer values.

$$edown(s_1,\ldots,s_n,s_{n+1},\ldots,s_{n+m})$$
:

inhibitor values if for all $i, 1 \le i \le n, s_i > 0$ and for all $j, 1 \le j \le m, S_{n+j} = 0$ then for all $i, 1 \le i \le n, s_i = s_i - 1$ else block execution of calling processes.

 $eup(s_1, s_2, ..., s_n)$: if processes blocked on $(s_1, ..., s_n)$ then awaken on of them else for all $i, 1 \le i \le n, s_i = s_i + 1$

Theorem

Agerwala's semaphores can simulate the action of an arbitrary Turing machine.

Inhibitor Nets

- A transition *t* can only be fired if all places connected by inhibitor arcs are empty.
- The transition *t* below can be fired as it has all input placed filled and all inhibitor places empty.



• The transition *t* below **cannot** be fires since its inhibitor places are not all empty (some or all contain tokens).





Ryszard Janicki Semaphores. Limits and Extensions 13/20

3.0

Nets with Inhibitor Arcs and Turing Machines

Theorem

Nets with Inhibitor Arcs are equivalent to Turing Machines.

• Agerwala's semaphores can model Nets with Inhibitor Arcs, so they can model Turing Machines as well.

・ 回 と ・ ヨ と ・ ヨ と

2

'Not Later Than'

The nets below allows the sequence a → b and the simultaneous step {a, b}, but the sequence b → a is disallowed.



• The above net models 'a is not later than b'.

回 と く ヨ と く ヨ と

'Only Simultaneously'

• The net below allows only the simultaneous step $\{a, b\}$, neither $a \rightarrow b$ nor $b \rightarrow a$ are allowed.



- The above net models 'only simultaneous execution of a and b'
- This net with this interpretation is a little bit controversial as the step $\{a, b\}$ does not have any sequential interpretation, so cannot be simulated by any sequential system!

白 ト く ヨ ト く ヨ ト

Comments on Generalized Semaphores

- Both inhibitor values and $eup(s_1, s_2, ..., s_n)$ are rarely used in practical applications.
- Inhibitor values are needed to simulate Turing Machines.
- Releasing resources seldom needs to be done in a specific order, so eup(s₁, s₂,..., s_n) are not so often used.
- *Problem*: Any kind of semaphores except the classical Dijkstra's semaphores are not so easy to implement, especially on a very low level, and implementations are usually expensive and not very efficient.
- Hence, very often we only have standard Dijkstra's semaphores to use.

소리가 소문가 소문가 소문가

| Rtobacco | Rpaper | Rmatch | |
|-------------------------|---------------------------|--|---------------------|
| rt: down(s); | rp: <i>down(s</i>); | rm: <i>down(s</i>); | |
| up(paper); | up(match); | up(tobacco); | \leftarrow agent |
| up(match); | up(tobacco); | up(paper); | |
| <i>goto</i> rt | <i>goto</i> rp | <i>goto</i> rm | |
| bt: down(smokert); | bt: down(smokerp); | bt: down(smokerm); | const- |
| up(s); | up(s); | up(s); | \leftarrow raints |
| <i>goto</i> bt | <i>goto</i> bp | <i>goto</i> bm | |
| Smoker with Tobacco | Smoker with Paper | Smoker with Matches | s |
| at: down(paper, match); | ap: down(tobacco, match); | am: <i>down</i> (<i>tobacco</i> , <i>paper</i>); | mo |
| | | | ⊢ k |
| up(smokert); | up(smokerp); | up(smokerm); | e |
| <i>goto</i> at | <i>goto</i> ap | <i>goto</i> am | rs |

• Intuition: smoker can pick only two ingredients *simultaneously*, i.e. as *atomic* one action.

白 と く ヨ と く ヨ と …

2

Binary Semaphores. Why they are admired?

- Simplicity, simplicity,
- In FSP formalism: 'Normal' semaphores:

const Max = M (must be a concrete number) range int = 0..MaxSEMAPHORE(N = K) = SEMA[N] (K is the initial value) $SEMA[v:int] = (when(v < Max)up \rightarrow SEMA[v+1])$ when (v > 0) down $\rightarrow SEMA[v - 1])$ ****** *********************** For M = 3 it expands to: $SEMA[0] = (up \rightarrow SEMA[1])$ $SEMA[1] = (up \rightarrow SEMA[2] \mid down \rightarrow SEMA[0])$ $SEMA[2] = (up \rightarrow SEMA[3] \mid down \rightarrow SEMA[1])$ $SEMA[3] = (down \rightarrow SEMA[2])$ Binary semaphore: M = 1, so $SEMA[0] = (up \rightarrow SEMA[1])$ $SEMA[1] = (down \rightarrow SEMA[0])$, we can use substitution: ********* $SEMA[0] = (up \rightarrow down \rightarrow SEMA[0])$ $SEMA[1] = (down \rightarrow up \rightarrow SEMA[1])$ SEMAPHORE(N = 0) = SEMA[0]SEMAPHORE(N = 1) = SEMA[1](日) (同) (E) (E) (E)

Binary Semaphores. Why they are admired?

- Hence, in FSP formalism, the binary semaphores are very simple!
- If the initial value is 0 (*False*), then: $SEMAPHORE = (up \rightarrow down \rightarrow SEMAPHORE)$
- If the initial value is 1 (*True*), then: $SEMAPHORE = (down \rightarrow up \rightarrow SEMAPHORE)$
- In reality, binary semaphores are much easier to implement, especially on low level, then 'normal' semaphores.
- Usually, the implementation of binary semaphores is conceptually different, i.e. it is not just a special case of 'normal' semaphores.
- Binary semaphores are much simpler than the normal ones in virtually any high level formalism!

・同下 ・ヨト ・ヨト