

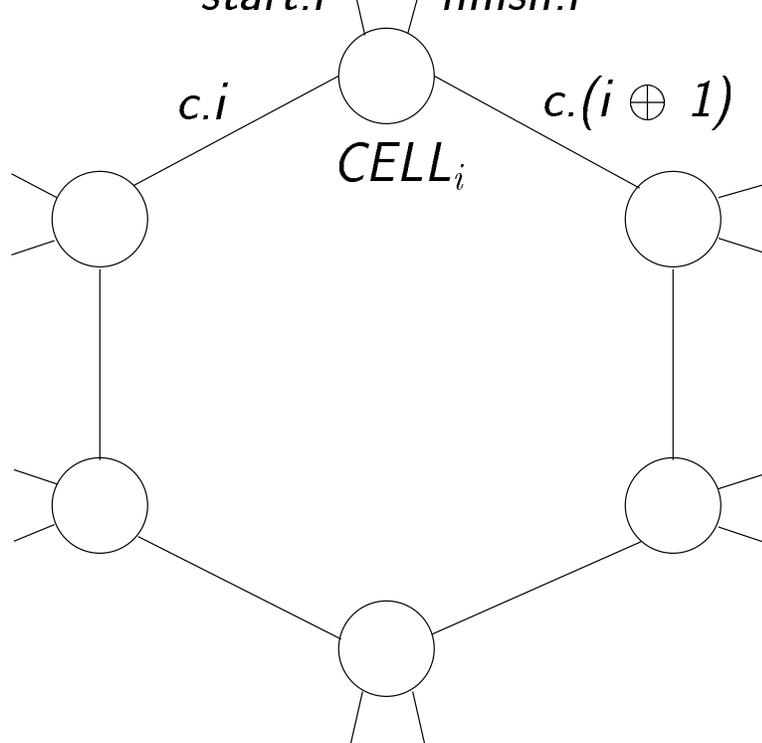
◇ The Cyclic Scheduler ◇

Suppose there are a number of processes which we need to control. Each process can be started by a *start* event and uses a *finish* event to indicate that it has finished. Suppose also that we want to start the processes in order, returning to the first when the last has been started; when a process has finished it can be started again, but only when its turn comes round.

We will define a *scheduler*, which uses *start* and *finish* events to control the processes. The scheduler will be implemented as a collection of cells (processes), each of which communicates with one of the processes being controlled, and also with other cells.

The next slide has a diagram of the case where there are 6 processes to control. \oplus denotes addition modulo the number of processes.

The idea is that each cell waits for a signal on the *c* channel to its left, which means that the process to its left has been started. Then the cell starts its own process, and sends a signal on the *c* channel to its right, to tell the next cell that it can start its process. Also, each cell has to wait for its process to do *finish* before starting it again.



In order to start everything off, one cell must begin by starting its process instead of waiting.

$$\begin{aligned}
 STARTCELL_i &= start.i \rightarrow c.(i \oplus 1) \rightarrow \\
 &\quad (finish.i \rightarrow c.i \rightarrow STARTCELL_i \\
 &\quad \square c.i \rightarrow finish.i \rightarrow STARTCELL_i)
 \end{aligned}$$

The other processes wait for $c.i$

$$WAITCELL_i = c.i \rightarrow STARTCELL_i$$

It is convenient to define

$$\begin{aligned}
 CELL_0 &= STARTCELL_0 \\
 CELL_i &= WAITCELL_i \quad (i > 0)
 \end{aligned}$$

and then

$$\begin{aligned}
 SCHED &= \left(\parallel_{0 \leq i < n} \{ start.i, finish.i, c.i, c.(i \oplus 1) \} CELL_i \right) \setminus \\
 &\quad \{ c.i \mid 0 \leq i < n \}
 \end{aligned}$$

There are three properties which we would like to verify for the scheduler. The first is that for each i , the events $start.i$ and $finish.i$ happen alternately, beginning with $start.i$. The second is that the events $start.0, \dots, start.(n-1)$ happen in the correct cyclic order. The third is deadlock-freedom.

For the first property, we can define a process specifying alternation of $start$ and $finish$ for each cell:

$$ALT_i = start.i \rightarrow finish.i \rightarrow ALT_i$$

and combine them in parallel to produce a specification for the scheduler as a whole.

$$ALTSPEC = \parallel_{\{start.i, finish.i\}}^{0 \leq i < n} ALT_i$$

In this parallel combination the alphabets are all disjoint, and no synchronisation is required. It is simply an independent parallel combination of the ALT processes.

The specification

$$ALTSPEC \sqsubseteq_t SCHED$$

can be checked with FDR.

For the second property, define

$$CYCLE_i = start.i \rightarrow CYCLE_{i \oplus 1} \quad (0 \leq i < n)$$

and specify

$$CYCLE_0 \sqsubseteq SCHED \setminus \{finish.i \mid 0 \leq i < n\}.$$