Self-stabilization with four-state machines.

We consider a sequence of N+1 finite state machines, numbered from 0 through N; two machines are "neighbours" if and only if their ordinal numbers differ by 1.

For each machine one or more "privileges" are defined, i.e. boolean functions of its own state and the state(s) of its neighbour(s). A "move" consists of selecting one of the privileges currently present and bringing the machine enjoying it in a new state that must be a function of its old state, the state(s) of its neighbour(s) and the privilege selected (the latter can only apply in the case of a machine that can enjoy more than one privilege simultaneously).

We call a state of the total system "legitimate" if and only if
1) exactly one privilege is present and, therefore, exactly one move is possible, and
2) the execution of the only possible move will again bring the total system in a legitimate state, and
3) successive moves will cause each privilige to be present with equal frequency.

We are interested in designing a set of machines, privileges and machines, such that, in addition, the system enjoys the property of "self-stabilization", i.e. regardless the initial state and regardless the privilege selected for the next move, it must always be possible to do a next move and after a finite number of moves the total system must arrive at a legitimate state. It is the requirement of self-stabilization that makes the problem non-trivial.

For each machine nr.i $(0 \leq i \leq N)$ we introduce two booleans, called "x[i]" and "up[i]" respectively. In the "bottom machine" up[0] = <u>true</u> holds permanently, in the "top machine" up[N] = <u>false</u> holds permanently, all other booleans are variables. In other words, the top machine and the bottom machine are two-state machines, the so-called "normal machines" (i.e. $0 < i < N$) are four-state machines.

For all machines nr.i with $0 < i \leq N$ the "privilege from below" is defined as

$$x[i] \neq x[i-1] \quad ;$$

for the normal machines the corresponding move is defined as

$$x[i]:= \underline{\text{non}}\ x[i];\ up[i]:= \underline{\text{true}}$$

but for the top machine --who has $up[N] = \underline{\text{false}}$ permanently-- the move is reduced to

$$x[N]:= \underline{\text{non}}\ x[N] \quad .$$

For all machines nr.i with $0 \leq i < N$ the "privilege from above" is defined as

$$x[i] = x[i+1]\ \underline{\text{and}}\ up[i]\ \underline{\text{and}}\ \underline{\text{non}}\ up[i+1] \quad ;$$

for the normal machine the corresponding move is defined as

$$up[i]:= \underline{\text{false}}$$

but for the bottom machine --who has $up[0] = \underline{\text{true}}$ permanently-- the move is given by

$$x[0]:= \underline{\text{non}}\ x[0] \quad .$$

If neither machine nr.0 nor machine nr.1 is to have a privilege, $x[1] = x[0]$ and $up[1] = \underline{\text{true}}$ should hold; if in addition machine nr.2 should also have no privilege, $x[2] = x[1]$ and $up[2] = \underline{\text{true}}$ should hold as well. Repeating the argument we see that the assumption of "no privilege at all" would lead to the conclusion $up[N] = \underline{\text{true}}$, but as this contradicts the truth we conclude

Theorem 1: There is at least one privilege.

Furthermore we can conclude --the proofs are left to the reader--

Theorem 2: Each move destroys the privilege that allowed it.

Theorem 3: If machine nr.0 has a privilege, machine nr.1 has not a privilege from below; a move by machine nr.0 will give a privilege from below to machine nr.1.

Theorem 4: If machine nr.N has a privilege, machine nr.N-1 has not a privilege from above; a move by machine nr.N may give a privilege from above to machine nr.N-1 (and will certainly do so in the legitimate state on account of Theorems 1 and 2).

Theorem 5: If a normal machine nr.i does a move on account of a privilege from above, it can only affect the privileges held by nr.i and by nr.i-1; machine nr.i is sure to lose its privilege from above, a new privilege from above may be created for machine nr.i-1 (and this will certainly happen in the legitimate state on account of Theorems 1 and 2).

Theorem 6: If a normal machine nr.i does a move on account of a privilege from below, it can only affect the privileges held by nr.i and by nr.i+1. There are now two cases to consider:

a) Originally $x[i] \neq x[i+1]$: machine nr.i and machine nr.i+1 will both lose their privilege from below, for machine nr.i a new privilege from above may be created;

b) Originally $x[i] = x[i+1]$: machine nr.i will lose all its privileges and for machine nr.i+1 a new privilege from below will be created (and this will certainly happen in the legitimate state on account of Theorems 1 and 2).

Corrolary 1: If a normal machine nr.i does a move on account of a privilege from below (above) and this privilege is not transferred to machine nr.i+1 (nr.i-1), then the total number of privileges is decreased.

Corrolary 2: No move increases the total number of privileges.

Corrolary 3: A normal machine enjoying both privileges cannot do a move without reducing the total number of privileges.

From the above it follows that in the legitimate state a privilege from below will go upward until the top machine reflects it as a privilege from above, until that again is reflected by the bottom machine as a privilege from below, etc. All our requirements of the legitimate state have been met. Self-stabilization now follows from the corrolaries.

We can ask ourselves: what is the minimum value for $k(N)$, such that after $k(N)$ moves the system must be in a legitimate state? The answer is that $k(2) = 2$ and for $N > 2$, $k(N) = N^2 - N + 1$. My demonstration for this result is unelegant and the exact value of the bound is not important enough to waste printing space on it.

* * *

More important is the question, whether we can eliminate --or rather: distribute-- the tacitly assumed demon, who was supposed to see to it that at any moment in time only one move was made. For this purpose we assume that each machine state is equipped with a private switch, guaranteeing that while a machine changes its state, no neighbour can inspect it; the changing or inspecting of a state are regarded as undivided actions. The bottom machine and the top machine can now be regarded as cyclic processes, inspecting if their privilege is present and if so, doing their move, while each normal machine can be regarded as a cyclic process that investigates the presence of its two possible privileges alternatingly and upon finding a privilege will do the corresponding move before investigating the presence of its other privilege.

The analysis is simplified by the fact that inspection for the presence of a privilege implies only one neighbour and --if the privilege is present-- the move is then a completely private affair.

If a normal machine nr.i enjoys two privileges, simultaneous execution of the two possible moves would lead to conflicting assignments to up[i], but this conflict is prevented from arising by the postulated sequential nature of the normal machines.

Simultaneous moves by two different non-neighbouring machines cannot interact at all and we are therefore left with the analysis of simultaneous moves by two neighbours.

If machine nr.i observes a privilege from below and machine nr.i+1 a privilege from above, no common states are involved and the net result is as if the moves were done in some order. If machine nr.i observes a privilege from above, machine nr.i+1 can observe no privilege at all and the question of simultaneous moves by nr.i and nr.i+1 does not arise. We are left with the case that machine nr.i and machine nr.i+1 both observe a privilege from below; simultaneous moves would interfere in the sense that upon completion of the moves, machine nr.i+1 would not have lost its privilege from below. As a matter of fact a whole sequence of machines might simultaneously detect the privilege from below and act upon it simultaneously.

The net effect, however, is as if the moves of the sequence had been done in order of decreasing ordinal number. As a result, even with the distributed demon, self-stabilization has been achieved.

* * *

Finally we observe that any normal machine enjoys in the legitimate state twice as often a privilege as the two end-machines. We can close the ring, merging the two end-machines into a single exceptional machine: this then is a four-state machine like the others and will enjoy a privilege with equal frequency as all the others.

12th October 1973
BURROUGHS
Plataanstraat 5
NUENEN - 4565
The Netherlands

prof.dr.Edsger W.Dijkstra
Research Fellow

Errata to EWD391:

page EWD391 - 5, 6 lines from below: for "non-decreasing" read "non-increasing"
page EWD391 - 6, 10 lines from above: for "." read ",provided K > N."
page EWD391 - 6, 8 lines from below: for "Solutiong" read "Solution".

Note. The second erratum --courtesy drs.C.S.Scholten-- shows a serious flaw in my reasoning; I should have known better!

EWD.

Note to readers of EWD392: The reading of this report is made less difficult when you have a set of draughtsmen at hand; while writing it, I used the set of my oldest son.

EWD.