27th June 1975

After many a sobering experience.

EWD492 was written in a state of great excitement: it described --without proof-- for the on-the-fly garbage collection our first solution, which we had found on the previous Tuesday afternoon.(On Fool's Day! We should have known...) The next Tuesday afternoon, when we tried to give a proof of its correctness, we quickly found the bug. It was repaired by postulating the availability of the atomic action "make at least grey" instead of the original

```
if node is observed to be white → make it grey
 [] node is observed to be non-white → skip
fi
```

The bug was, that between the node being observed to be white and the subsequent action "make it grey", if done by the mutator, the collector could have made it grey and subsequently black. Such interleaving would violate the intended monotonicity of the colour history for each node; hence our introduction of "make at least grey". Our then correct solution went through a series of embellishments, and, eventually --more than two months later!-- EWD496 was written. And that was that.

A few weeks later the question was raised whether the dual-processor configuration with one mutator and one collector could be generalized into a multi-processor one, with many mutators and, more interesting, many sequential processors for the garbage collection as well. And we designed a system with as many mutators as we had LISP-programs, an arbitrary number of "markers" and one "collector proper".

The mutators would, as before, make the target node of each edge at least grey before placing the edge.

The markers would search for grey nodes: each grey node found they would make black after having made its successors at least grey. This would be a continuously ongoing activity, their searching for grey nodes would be such that eventually each grey node would be found by at least one of them.

The collector proper was proposed to consist of the repeated execution of

```
make the roots at least grey;
scan all the nodes in some order and make for each node observed to be black, its successors at least grey {now edges from a black node to a white node have disappeared};
wait until in a single scan in some order all nodes have been observed to be non-grey {all white nodes are garbage};
process all nodes in some order, where "to process a node" means that a white node is appended to the free list and a black node is made white.
```

In EWD496, the last activity caused all black nodes to disappear, so that edges from a black to a white node were then guaranteed to be absent. Because here our markers are supposed to continue their activity, the absence of edges from black to white cannot be guaranteed, hence the insertion of the second statement, which indeed establishes this absence (any edges from black to white that were present can no longer lead to a white node and neither mutators nor markers can have introduced new ones).

For nearly two weeks I thought that the above solution was correct, until I started to try to prove its correctness. It turned out to be wrong. Observing in a single scan in some order all nodes to be non-grey is no longer sufficient to conclude that all white nodes are garbage. The argument breaks down where in EWD496 - 6, Note 5, was written: "As the mutator leaves grey nodes grey, the collector must have encountered that grey node during the scan." This conclusion does not hold for our collector proper, because during that scan the ongoing marker activity may have made that grey node black! So there I was. Fooled again.....

So --rather reluctantly, I admit-- I introduced some more activity for the collector proper and a fourth colour, ultrablack say (which is regarded to be darker than black). The mutators remain as they are, the markers react as above on a node observed to be grey, but make it "at least black" (instead of just "black" as above) and the collector proper that I considered was the repeated execution of the following program:

```
{Q0}
make all the roots at least grey {Q1};
scan all the nodes in some order and make for each node observed to be
   black its successors at least grey {Q2};
repeat make each node, observed to be black, ultrablack
until in a single scan in some order all nodes have been observed to be
   white or ultrablack {Q3};
process all nodes in some order, where "to process a node" means that
   a white node is appended to the free list and an ultrablack node is
   made white.
```

The intended assertions are:

Q0: there are no ultrablack nodes

Q1: there are no ultrablack nodes and the roots are at least grey

Q2: there are no edges from black or ultrablack to white and the roots are at least grey. (As a result, the presence of a white reachable node implies the existence of at least one grey node.)

Q3: there are no grey nodes and Q2 still holds. (As a result, all white nodes are therefore garbage.)

Q4: there are no ultrablack nodes.

Assertions Q0 and Q1 present no problems, because neither mutators nor markers introduce ultrablack nodes. For Q2, the absence of edges from ultrablack to white follows from the absence of ultrablack nodes, the absence of edges from black to white follows from the fact that neither mutators nor markers introduce them and those originally present have disappeared or point no longer to a white node.

During the introduction of ultrablack nodes, Q2 remains valid: the only nodes made ultrablack are already black, therefore they don't point to a white node and edges from ultrablack to white are not introduced. (Remember that mutators make each new target at least grey.) But now the critical step: is conclusion Q3 justified? In a single scan in some order all nodes have been observed to be either white or ultrablack: does this fact justify the conclusion that there are no grey nodes? This must be hard to prove. Consider the following scenario:

Let there be an edge from a grey node A to a white node B; let marker 1 have observed node A to be grey and have decided to make node B at least grey; hereupon, marker 1 goes to sleep; let the edge from A to B be removed by a mutator; let marker 2 observe node A to be grey and make it black; let the collector proper make node A ultrablack. As soon as marker 1 wakes up again it will make the still white node B at least grey!

There seem two ways out: one of them could be to divide the memory into as many sections as we have markers, such that each marker only searches for grey nodes in its own section. Alternatively, we can try to weaken Q3.

I tried weakening of Q3 first (such as "There are no grey, reachable nodes.") and although that would be sufficient to conclude that all white nodes are garbage, I could not exclude the occurrence of grey garbage at stage Q3. In its last phase the collector makes dark garbage white, so that it can be collected the next time; the presence of grey garbage, however, can trigger a counter activity of a marker and as a result, existing garbage may forever escape being collected.

So: can we conclude the absence of grey nodes when each node can only be made black by a single marker?

When after a single scan all nodes have been observed to be either white or ultrablack, Q3 "there are no grey nodes" can only be violated provided since the start of the scan, a grey node has been introduced.

The introduction of a grey node by a mutator would have implied the existence of a white reachable node, and therefore of a grey node, when the scan started. This latter grey node would have been detected during the scan as either a grey or a black one, contrary to the hypothesis.

The introduction of a grey node by a marker implies at that moment a grey node in its section, which at the end of the scan will be grey or black. The assumption that the scan has observed it to be white implies the earlier introduction of a grey node since the scan started, and as a mutator cannot have caused the earlier introduction, also that earlier introduction must have been caused by a marker. The argument can be repeated and, the store being finite, leads to a contradiction.

All through the scan there have been no grey nodes; as a result, no black nodes have been introduced, no black nodes have been observed either, and we can strengthen Q3 "There are no grey nodes and no black nodes, only white and ultrablack ones, and all the white nodes are garbage."

During the last phase of the collector proper, the mutators and the markers will in general introduce grey and black nodes, but each of the nodes encountered by the collector will be either white or ultrablack: the ones that were originally ultrablack will remain so until the collector proper has made them white, the originally white ones will remain so until the collector proper has appended them to the free list, because before that moment, the white node was garbage, and garbage remains untouched by the mutators (by definition) and by the markers, because there were no grey nodes to start with. (Note that nothing prevents during the last phase of the collector proper the creation of --new!-- grey garbage! But none of the grey nodes introduced during that last phase can having an outgoing edge with one of the originally white garbage nodes as target.)

Salvo errore et omissione, this completes the argument.

* * *

The first moral of this story is never to believe the correctness of a solution without a proof for it. The second moral of this story is never to forget the first one. The third moral is that for the design of multiprocessor installations we cannot rely on the traditional approach of the optimistic engineer, who, when the design looks reasonable, puts it together to see if it works.

With LSI-technology existing it seems unavoidable that multiprocessor installations will be built, and, people being as they are, it seems equally unavoidable that many of them will be put together by aforementioned optimistic engineer. I shudder at the thought of all the new bugs: they will only delight the Devil. Am I too pessimistic? Nobody knows the trouble I have seen....

Plataanstraat 5
NL-4565 NUENEN
The Netherlands

prof.dr.Edsger W.Dijkstra
Burroughs Research Fellow