

For the record: painting the squared plane

Thanks to Richard S. Bird, the following problem circulated at the 3rd International Conference on the Mathematics of Program Construction, Kloster Irsee, Germany, 1995.

Starting with a sufficiently large piece of (white) squared paper, we play the following one-person game, that consists of two phases. In the first phase, we may select 9 squares and colour them red. In the second phase, we may repeatedly colour a white square red provided it already has at least 2 red neighbours (where being neighbours means sharing an edge). The question to be answered is "Is it possible to colour the 10 by 10 square?"

* * *

When, in the second phase, a white square may be coloured on account of exactly two "enabling" neighbours, we can draw the topological conclusion that there are two possibilities for the relative position of those two enabling neighbours: either around a corner or opposite to each other. When we demonstrate a possibility by showing how to

do it, we can hardly avoid that each move in our witness is of some type. For instance: in order to show that the 9 by 9 square can be coloured, we can colour in the first phase the squares along a diagonal; in the second phase the enabling neighbours are then always situated around a corner. (For the 9 by 9 square there are also solutions in which opposite enabling neighbours occur as well.)

When demonstrating an impossibility, we consider no solutions because there aren't any, and in this type of argument it would be very nice to avoid the case analysis by not drawing the topological conclusion at all, i.e. by ignoring the cyclic arrangement of the four neighbours and of the four edges they share with the square they surround. In short, for an impossibility argument it is very tempting to try to get away with the simplicity of a counting argument. The only still open question is then what to count.

Counting red squares is too crude: their number increases by 1 at a time, which is only mildly interesting; worse is that the notion "red square" says nothing about

the notion of being neighbours. In order to capture the latter, we have to turn to the edges.

How do we relate edges to colours of squares? Each edge is identified by -and identifies- the pair of squares that share it. This enables us to define three "edge colours":

an edge is red, when shared by two red squares,

an edge is white, when shared by two white squares, and

an edge is pink, when shared by a red square and a white square.

Accordingly, a white square has no red edges, and a red square has no white edges. More precisely: colouring a white square red, turns its pink edges into red ones, and its white edges into pink ones. The rule in the second phase of "at least two red neighbours" means that at least 2 pink edges turn red, and at most $4-2$ white edges turn pink. In other words: in the second phase, the number of pink edges does not increase. Since the first phase does create at most 36 pink edges, we never end up with the 40 pink edges that would

occur in a red square of 10 by 10,
and thus the question is settled.

* * *

Faced with a problem, people usually react as follows. In their first explorations they try to colour the 10 by 10 square, and, finding they cannot do it, they change to trying to show the impossibility, but in that second effort they drag all the topological ballast into the argument, make pictures, talk about "occupied columns and rows" etc. In that second effort they are heading for a counting argument, but are insufficiently aware that avoiding case analysis is counting's major purpose.

* * *

The final argument nicely illustrates the linguistic demands that the doing of mathematics makes: a term for what we called "pink edges" is absolutely essential.

Austin, 30 August 1995

prof. dr. Edsger W. Dijkstra
Department of Computing Sciences
The University of Texas at Austin
Austin, TX 78712-1188
USA