

Courtesy A.J.M. van Gasteren, C.S. Scholten, and J.G. Wiltink

Show that, for any set of grid points in the plane, it is possible to colour each of these points either red or blue such that on each grid line the number of red points and the number of blue points differ by at most 1.

* * *

Since constructive existence proofs are nice, we shall design a colouring algorithm; initially all grid points of the set are blank, finally each grid point of the set is either red or blue. Since the final requirement P_0 :

P_0 : on each grid line the number of red points and the number of blue points differ by at most 1

holds initially as well, we propose that P_0 be invariant.

We propose the simplest colouring algorithm we can think of, viz. one that colours blank grid points one at a time. Colouring a blank point can violate P_0 only in the presence of what we call "a surplus", defined by:

a surplus is a grid line that contains at least one blank point and on which the number of red points and the number of blue points differ by 1.

The invariance of P_0 would be guaranteed by

the permanent absence of surpluses, but since grid points are coloured one at a time, this constraint is stronger than we can live with. We must admit the occurrences of surpluses, but let us investigate the extent in which we can constrain their existence. But what are the relevant concepts in terms of which we can express such a constraint?

To begin with we observe that permuting a number of parallel grid lines would transform a solution into a solution: a reference to, say, "the left-most vertical surplus" would have introduced a concept foreign to the problem statement.

Next we observe that the problem statement is symmetric in vertical and horizontal grid lines. Hence, as far as the direction of surpluses is concerned, the only notion that can be of significance is whether two surpluses are "parallel" or "orthogonal."

Finally we observe that the problem statement is symmetric in red and blue. Hence, as far as the dominant colour of surpluses is concerned, the only notion that can be of significance is whether two surpluses are "equal" or "opposite".

Armed with these insights we investigate which surplus configurations we have to allow.

We have to admit state A, given by

A: there are no surpluses ,

since this holds initially and finally. Colouring in state A a blank grid point leads to state A, state B or state C, the latter ones being given by

B: there is one surplus, and

C: there are two surpluses that are equal and orthogonal.

Colouring in state B a blank grid point so as to remove the given surplus leads to state A or state B.

Colouring in state C a blank grid point so as to remove one of the given surpluses leads to state A, to state B, or to state D given by

D: there are two surpluses that are opposite and parallel

Colouring in state D a blank grid point so as to remove one of the given surpluses leads to state B or state C.

Thus we have found the closure and shown that it is possible to maintain $P \wedge (A \vee B \vee C \vee D)$, which concludes our argument.

*

*

*

Wiltink and Scholten solved this problem. (Wiltink did it with nested loops, $P \wedge A$ being the invariant of the outer one; he fully exploited the symmetries.) Van Gasteren taught us how to look for the concepts in which to conduct the argument.

Rutger M. Dijkstra, who happened to be at home, draw my attention to a shortcoming in my explanation.

Nuenen, 27 December 1986

prof. dr. Edsger W. Dijkstra
Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712-1188
United States of America