

## The knight's tour

Early September, after 2 weeks in Austin, I was for a few days back in the Netherlands, where even Netty van Gasteren's excellent hospitality did not prevent me from suffering from jet lag. Hopelessly awake during the 2nd night in Waalre, I was reminded of the problem of the knight's tour, a problem at which I had looked a number of times, but each time only to turn away from it because I found it too ugly. This time, I was in a very different mood: with so much experience now in the conscious avoidance of complications, I felt that now I should be able to solve the problem. Here is how I arrived at a solution.

\* \* \*

The move of the knight changes one of the coordinates of its position by 2, the other coordinate by 1; thus, from a square sufficiently far from the edge of the board, 8 different moves are possible. Show that there exists a knight's path that visits each square of the  $8 \times 8$  chessboard exactly once.

My first decision was to show the existence by constructing a witness, my second decision was to aim (for a start) at a cyclic path, the cyclic path being simpler because it has


no endpoints.

Next I tried to "simplify" the knight's move by looking for an environment in which I could concentrate on the coordinate being changed by 2 and in which the change by 1 of the other coordinate was a (temporarily) ignorable perturbation. I obtained that environment by


(i) covering the chessboard by 16 "super-squares", each formed by 4 adjacent original squares

(ii) by restricting myself to moves between supersquares that share a side - these are called the "straight" moves, the others are called the "skew" moves.


ad (i)



ad (ii)



the 4 straight moves



the 4 skew moves

For the straight moves the set of reachable super-


squares is independent of the internal position within the starting supersquare: they are (in the usual wind nomenclature) N, W, S and E. This independence is not found with the skew moves. This observation justifies a (temporary?) restriction to straight moves.

When we name the subsquares of a supersquare and give reachable subsquares in W and S the corresponding names, we observe

a	b	c	d
c	d	a	b
b	a		
d	c		


that is, a horizontal step - to W - corresponds to an interchange of the 2 rows, and a vertical step - to S - corresponds to an interchange of the 2 columns. These two interchanges are commuting involutions, and hence we can thus name the subsquares of the whole board. To give an alternative phrasing: reachability via straight moves generates 4 equivalence classes, which we labelled a, b, c, and d.

As the adjoining maze - a so-called "Gomory maze" - shows, the 16 supersquares can be arranged in cyclic paths so that neighbours




in the path share an edge on the board. In other words: using straight moves only the original problem of all 64 squares on a single cycle cannot be solved, but the  $4 \times 16$  squares divided over an a-cycle, a b-cycle, a c-cycle and a d-cycle is easily solvable (in many ways even: the different cycles don't need to use the same Gomory maze).

The remaining problem is to combine these 4 cycles into a single one. From each of the 4 cycles at least 1 straight move has to be taken away and an equal number of skew moves have to be added. I tried what seemed the simplest: remove from each of the 4 cycles precisely 1 straight move. We get a single cycle if and only if the 4 straight moves taken away and the 4 skew moves added form a cyclic path (on which, by necessity, the two types of move alternate):



This was the moment I left my bed for now I needed pencil to construct the following cycle of length 8:

drawn moves are straight, dotted ones are skew.



In particular

the analysis that the 4 straight moves were from the 4 different cycles was much more than I could do in my head. The above configuration of straight moves can for instance be found when the 4 cycles are formed with the same Gomory maze, and this observation concludes my proof.

The above solution is nice because it is so simple that people can remember it, and as such it is cause of some satisfaction, if not pride. That it was constructed following the discipline that is now so well-known, tempers the pride, but strengthens the satisfaction.

Austin, 26 September 1992

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