



© 1997–2004, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non–commercial use. For other uses, including electronic redistribution, please contact us.

05/06/2007

News

Rubik success in twenty–six steps



A Rubik's cube in a scrambled state. Image from [Wikipedia](#), reproduced under the [GNU Free Documentation License](#).

If you were deeply traumatised by the 80s Rubik's cube mania, you can now rest a little easier: it's just been proved that you can solve a Rubik's cube in only 26 moves. Computer science professor [Gene Cooperman](#) and graduate student Dan Kunkle of the [Northeastern University](#) in Boston used clever algebra and superfast parallel computing to show that no matter how scrambled the cube is, you can always get back to the *home state* in 26 steps. This is a small but significant improvement on the previously known bound of 27.

The new result is not just of interest to those who like to get their brain in a twist: since its conception in the late 1970s, the Rubik's cube has served as a guinea pig for testing techniques to solve large–scale combinatorial problems. "The Rubik's cube is a testing ground for problems of search and enumeration," says Cooperman. "Search and enumeration is a large research area encompassing many researchers working in different disciplines from artificial intelligence to operations. The Rubik's cube allows researchers from different disciplines to compare their methods on a single, well–known problem."

Whether or not you've ever been obsessed by Rubik's cube, you'll probably know that each of the cube's six faces consists of 9 smaller *facelets*, and that you can change their configuration by rotating any of the cube's faces through 90° , 180° or 270° . When you unwrap a pristine cube from its packaging, facelets belonging to the same face all have the same colour. Your task is to quickly scramble the cube up into one of its many possible *states*, and then painstakingly get it back to its original position.

Rubik success in twenty–six steps



Making a move. Image from [Wikipedia](#), reproduced under the [GNU Free Documentation License](#).

Mathematically speaking, any state of the cube corresponds to a *permutation* of its facelets – a rule that assigns to each facelet in the home state a new location on the cube in the scrambled state. Not every permutation of facelets can be achieved by a sequence of moves – for example no move can shift the centre facelet of a face so when trying to solve Rubik's cube you restrict your attention to a collection of *legal* permutations. These form a self–contained system. If you follow one legal permutation by another, the end result is also a legal permutation. If you've done one legal permutation, you can always find another one – namely the reverse of what you've just done – that will get you back to the position you started with. Doing absolutely nothing to the cube is also a legal move and corresponds to the permutation that leaves every facelet where it is, known to mathematicians as the *identity permutation*.

Collections of objects exhibiting these three characteristics – two objects combining to give you a third of the same type, being able to reverse operations, and having an identity object – crop up all over the place. You find them, for example, in the whole numbers combined by addition, the real numbers combined by multiplication, or the symmetries of geometrical objects like polygons. If your structure also satisfies a fourth rule called associativity, as the permutations of Rubik's cube and all the above examples do, then it is what mathematicians call a *group*.

When thinking about a group, you don't have to explicitly refer to the things it's made up of – permutations, numbers, symmetries, or whatever they may be. You can think of it in purely algebraic terms as a collection of *elements* together with a *binary operation*. This is what makes group theory so powerful: you can solve an explicit problem on a purely abstract level, and your solution will apply to all other objects that have the same group structure.

The problem of solving Rubik's cube can be visualised using what is called the group's Cayley graph – a network whose nodes are the legal states of the cube (corresponding to legal permutations) and which has two nodes linked up if you can get from one to the other by a legal move, which, incidentally, is itself a permutation. The home state of the cube corresponds to one of the nodes (and the identity permutation) and solving the cube corresponds to finding a path from one of the nodes to the home state along a sequence of linked–up nodes. A brute–force approach to showing that you can always solve the cube in 26 moves would be to show that no such path involves more than 26 steps. This sounds good in theory, but there is a huge problem: the Cayley graph of the group of legal permutations has 43,252,003,274,489,856,000 nodes, a number that defies even the fastest of supercomputers.

Rubik success in twenty–six steps



The home state. Image from [Wikipedia](#), reproduced under the [GNU Free Documentation License](#).

Cooperman and Kunkle proved their result by cleverly exploiting group theory to reduce the number of nodes. They started by only considering half–turns (ie 180° rotations) of faces as legal moves, resulting in a smaller number of legal permutations and hence a smaller group. Taking into account the symmetries of the cube that once you know how to solve the cube from one state you also know how to solve it from, say, its mirror image they came up with a Cayley graph with just 15,752 nodes. With this comparatively small number it took less than a day of computing time to show that the number of moves needed to solve the cube in this case is at most 13.

The permutations generated by half–turns form a smaller group a *subgroup* sitting within the full group of legal permutations. The permutations you're left with when you take this subgroup away don't themselves form a group, but you can divide them up into collections called *cosets*, which do form a group. Applying the Cayley graph technique to the group of cosets, Cooperman and Kunkle were able to prove that the "distance" between a quarter–turn solution and a half–turn solution is at most 16, giving them an overall bound of $16 + 13 = 29$. A refined analysis of the subgroup and its cosets knocked off the final three steps to give them their final answer of 26.

Despite these algebraic tricks, and although the efficiency of the researchers' computer program was boosted by a new method for calculating the behaviour of group elements, a huge amount of computing power was still needed to perform the calculations. Cooperman and Kunkle used 7 terabytes of distributed disc space from computers at [Teragrid](#) and Northeastern University to achieve their result in 8000 CPU hours. And although the computer now knows the 26 step solutions for all possible states of the cube, you still don't. It might take you a lot longer 8000 hours to find them.

Further reading

- These [Lecture notes on the maths of Rubik's cube](#) give a gentle and comprehensive introduction;
- If you're familiar with group theory, you will be able to follow Kunkle and Cooperman's paper [Twenty–six moves suffice for Rubik's cube](#);
- Find out more about group theory on [Plus](#).

Marianne Freiberger

Rubik success in twenty-six steps



Plus is part of the family of activities in the Millennium Mathematics Project, which also includes the NRICH and MOTIVATE sites.

—>