Solving All 164,604,041,664 Symmetric Positions of the Rubik's Cube in the Quarter Turn Metric

Tomas Rokicki

March 18, 2014

Abstract

A difficult problem in computer cubing is to find positions that are hard—positions that are as far from solved as possible. An effective way to find such positions is to examine positions exhibiting symmetry; empirically, we see that hard positions have a higher frequency in the subset of positions exhibiting symmetry than in general. While there are many symmetric positions, they are a small fraction of all positions, and they belong to a few large subgroups which enables us to solve them effectively with a coset solver.

1 Introduction

In 2006 Silviu Radu (with some help from Herbert Kociemba) achieved perhaps the most astounding feat of computer cubing when he computed optimal solutions in the face-turn metric to all 164,604,041,664 positions of the Rubik's Cube that exhibit symmetry. This paper, eight years later, confirms all of his results and extends them to the quarter-turn metric.

We begin with a discussion of metrics, and then define what it means for a position to be symmetric. We outline how the group structure of the symmetric positions can be exploited to quickly solve all such positions, and present some results from our explorations. We close with a discussion of the results and what the next steps might be.

2 Face-Turn Metric versus Quarter-Turn Metric

Ford versus Chevy. Emacs versus vi. IOS versus Android. In the mathematics of the cube, the division is face-turn metric versus quarter-turn metric. These are two distinct ways of counting moves on the cube; in the quarter-turn metric, every 90 degree twist counts as one move, but a 180 degree twist is two. The face-turn metric counts that 180 degree twist as only a single move.

From the very beginning of discussions on the cube, there were these two camps. Stanford cubers preferred the half-turn metric (lining up with David

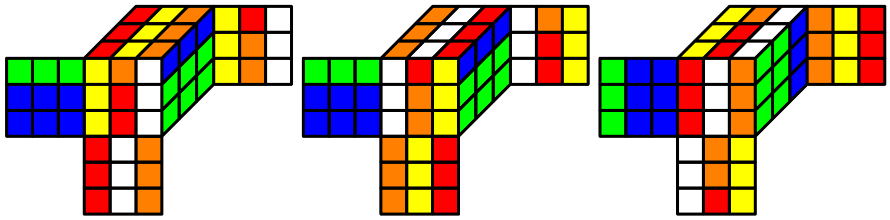

(a) (b) (c)

Figure 1: Positions (a) and (b) are recoloring-equivalent. Positions (a), (b), and (c) are all symmetry-equivalent.

Singmaster), while MIT cubers liked the quarter-turn metric. Most popular presentations of the cube implicitly used the half-turn metric; most mathematical analysis seemed to prefer the quarter-turn metric. Our recent result on God's Number for the cube[5] was for the face-turn metric; no equivalent result currently exists for the quarter-turn metric. Similarly, Radu's result that stimulated the creation of this paper is for the half-turn metric; this paper presents results in the quarter-turn metric.

Indeed there are other metrics, such as the slice-turn metric, which permits twists of a slice (holding the two outer faces in the same position) as a single move.

3 What is Symmetry?

We define what it means for a position to be symmetric in three steps, starting with two equivalence relations.

We define two positions of a cube to be recoloring-equivalent if there is a one-to-one mapping between the sticker colors of the two positions. One way to do this is to label all stickers that are the same color as the center up cubie 'U', all stickers that are the same color as the center right cubie 'R', etc, and then define the position by the these labels rather than by the actual colors. This is conventional when talking about the cube; we don't care what the specific coloring of a cube is, just how the stickers relate to each other, and since moves on the cube generally do not move the center cubies, we simply adopt a convention relative to the center cubie colors. For instance, the two positions (a) and (b) in Figure 1 are recoloring-equivalent, but position (c) is not.

We define two positions of a cube to be symmetry-equivalent if, when you reorient (rotate and/or reflect) the first in some manner, and it becomes recoloring-equivalent to the second, then the positions are symmetry-equivalent, or related by symmetry. For instance, when you rotate the up (U) face of position (c) in Figure 1 90 degrees to the front (F) face, the position becomes identical to that of position (a) and recoloring-equivalent to position (b), so it is symmetry-equivalent to them both.

Two positions that are related by symmetry have the same distance, since there is a one-to-one mapping of moves on one to moves on the other defined by the orientation. Since there are 48 different ways to orient a cube, the size of these equivalence sets is at most 48 (and it is usually 48). When talking about sets of positions (such as those at a given distance), it is conventional to say "the size of the set mod M" as the number of distinct equivalence classes (with respect to symmetry) contained in the original set; that is, the number of positions in that set that are not related to each other by symmetry. M is used to represent the group of 48 rotations and reflections of the full cube.

A position exhibits symmetry if and only if there is some non-identity orientation that is equivalent to itself. All the positions in Figure 1 exhibit symmetry, since a 180 degree rotation around one of the three major axis of the cube results in a recoloring-equivalent position; for position (a), that axis is the one passing through the center of the U and down (D) faces.

Another useful concept is that of antisymmetry; a position is antisymmetric if it is symmetry-equivalent to its own inverse. Symmetry and antisymmetry can be combined to generate equivalence classes of size up to 96. When talking about sets of positions we say "the size of the set mod M+inv" to indicate the count of distinct equivalence classes modulo symmetry and antisymmetry.

There are almost exactly 96 times fewer equivalence classes taking both symmetry and antisymmetry into account than there are positions of the cube. Since every position in one of these equivalence classes has the same distance, and has closely related solution sequences, most computer cubing takes advantage of these equivalences, and most reports of interesting positions give only a single member of one of these equivalence classes since the other positions can be trivially determined.

4 Hard Positions

A major interest in computer cubing has been determining the hard positions of the cube—the positions that take the largest number of moves. Indeed, God's Number is the distance of the position furthest from solved.

There is no known simple way to find hard positions. In smaller puzzles, full state space explorations, cataloging every position by its distance, finds all hard positions, but the state space of the Rubik's cube is too large for that approach to work (at the moment). Optimally solving many random positions can sometimes be lucky and find a hard position. For the half-turn metric, positions of length 20 are very rare (fewer than one in every 80 billion positions); in the quarter-turn metric, positions of length 24 are exceptionally rare (probably fewer than one in every trillion positions) so this approach will probably not work.

Radu's feat, solving more than 164 billion distinct positions in the face-turn metric, was an amazing tour-de-force, a combination of insight and talent, bringing together several ideas to accomplish something most would have thought impossible at that time. He tied together GAP, a program for group theory calculations, a modified version of Michael Reid's optimal cube solver, and some long runs of Herbert Kociemba's optimal solver from Cube Explorer to calculate distance distributions for all symmetric subgroups of the Rubik's Cube. Before his work, only a few thousand distance-20 positions were known; he increased this number to more than 1,000,000 (more than 32,000 mod $M + \text{inv}$), and showed that, if there were any distance-21 positions, they must be positions lacking in symmetry.

In the time since Radu's work, God's Number for the Rubik's Cube in the face-turn metric has been shown to be 20, so now we know there are no distance-21 positions at all. In addition, over ninety million (one million mod $M + \text{inv}$) distance-20 positions have been found so far. We have estimated there are a total of 490 million (five million mod $M + \text{inv}$) distance-20 positions.

God's Number in the quarter-turn metric is still unknown (although it is almost certainly 26). Exactly three positions at distance 26 are known, and they are all symmetry-equivalent. This position is that found composing the position known as superflip (all edges flipped) with the position known as four-spot (exchange two pairs of opposite centers). The only distance-25 positions known are the 36 immediate neighbors of the distance-26 position, representing only two symmetry-equivalence classes. Our work presented here more than triples the count of known distance-24 positions, from 12,602 (1,260 mod $M + \text{inv}$) to 39,818 (3,339 mod $M + \text{inv}$).

5 Solving Large Subgroups

Optimally solving 164 billion cube positions is a huge challenge. Using the best position-at-a-time optimal solvers, which can do about 4 positions a second per physical CPU core would take 1300 core years. Luckily these 164 billion positions reflect only 3.4 billion positions mod $M + \text{inv}$, so we can solve these in about 27 core years, a fairly large investment. But we can do better. Since all the symmetric positions belong to one of six large subgroups of the cube, we can use a faster technique, the same one used to find God's number, called a coset solver. (In this case we only solved the trivial coset of each of the subgroups).

A typical optimal solver for Rubik's cube[3] uses a function $f()$ on a cube position that gives a lower bound on the number of moves that position is from solved. This function is implemented with a large table, called a pruning table or pattern database. Iterated depth-first search is then used to find an optimal solution. The pattern database can be derived from a subgroup H of the cube, where the pruning table is indexed by the quotient coset space $G \setminus H$. Such pruning tables can have high average distance and thus yield effective search pruning, but since every position in H has distance 0, they can generate many sequences that do not solve the original position and thus must be individually rejected. A coset solver takes advantage of these normally rejected sequences by recasting the original problem: instead of solving a single position p, it solves all positions Hp (a coset), marking each solved position in a bitmap of size H as each sequence is found. With a few additional enhancements, this technique

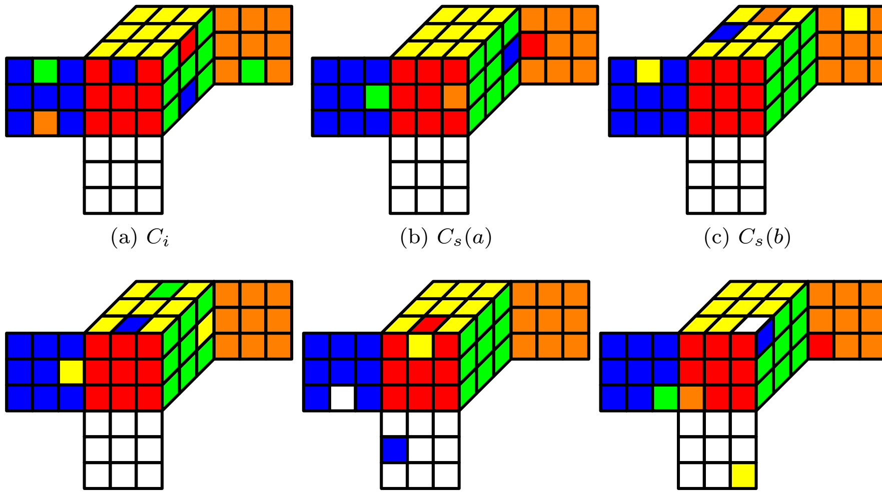

(a) C_i (b) $C_s(a)$ (c) $C_s(b)$

(d) $C_2(a)$ (e) $C_2(b)$ (f) C_3

Figure 2: Example positions from the six symmetry subgroups we solve.

can be used to find optimal solutions thousands of times faster than considering a position at a time.

Luckily, all symmetric positions fall nicely into a set of six subgroups, so the coset solving technique works well. For a discussion of symmetry and subgroups, see Kociemba's excellent website[2].

A particular position has a "degree of symmetry" which is the number of distinct orientations of a position that are recoloring-equivalent to the original position. The degree of symmetry of a position is always a factor of 48.

There are distinct symmetry subgroups with the same degree of symmetry. For instance, $C_s(a)$ is the subgroup where the orientation that preserves recoloring equivalence is a reflection through a plane midway between the up and down faces; it has a degree of symmetry of 2. C_i is the subgroup where the orientation that preserves recoloring equivalence is reflection through the centeral point of the cube; it also has degree 2.

There are actually three orientations of the subgroup $C_s(a)$—one where the reflection plane is between the U and D faces, one where it is between the L and R faces, and one where it is between the F and B faces; these subgroups are symmetry-equivalent, so we only consider one.

The five symmetry subgroups of degree 2 are C_i, $C_s(a)$, $C_s(b)$, $C_2(a)$, and $C_2(b)$. There is only one symmetry subgroup of degree 3: C_3. The 24 other symmetry subgroups that positions of the Rubik's cube fall into are all subgroups of these six main subgroups, so if we solve these six subgroups, we will have solved all symmetric positions.

The size of these six subgroups are shown in Table 1. The sizes do not sum to the total of 164 billion because some of the subgroups occur in multiple

Subgroup	Size
C_i	45,864,714,240
$C_s(a)$	18,345,885,696
$C_s(b)$	424,673,280
$C_2(a)$	15,288,238,080
$C_2(b)$	2,548,039,680
C_3	3,779,136

Table 1: The sizes of the six subgroups we solve.

equivalent orientations, and because some positions occur in multiple subgroups.

For the degree-3 symmetry subgroup C_3, we enumerated the positions, selected a single position from each symmetry/antisymmetry class, resulting in only 500,704 positions, and solved each with an optimal solver; this took less than a day for each of the quarter-turn and half-turn metrics.

We solved each of the five degree-2 symmetry subgroups with a coset solver in both the half-turn and the quarter-turn metric. In each case, we ran each coset solver until the number of remaining positions was sufficiently small (a few thousand) to quickly finish things off with a general optimal solver. (Coset solvers can solve many positions very quickly, but it is usually most effective to use a fast optimal solver to "finish up" the last positions.) This took a couple of core days per coset, easily dominated by the time required to write the program and check the results. No effort was made to multithread these coset solvers.

As an optimal solution to each new position was found by the coset solver, the overall symmetry of the position was evaluated and tabulated, so we were able to find a distance distribution not only for the large degree-2 subgroup but also all higher-degree subgroups. Since many symmetry subgroups were subgroups of more than one of the subgroups we wrote coset solvers for, this provided some cross-checking among the runs.

6 Results

Table 2 summarizes the results for the half-turn metric; Table 3 for the quarter-turn metric. More information is available at `http://cube20.org/symmetry/`. All of our half-turn metric numbers were checked against Radu and Kociemba's results as published on `http://kociemba.org/symmetric2.htm` and they were found to be in complete agreement.

As a result of this effort we were able to extend the count of known distance-24 positions in the quarter-turn metric from 12,602 to 39,818. This is the first such extension of this set for many years. If any additional positions of distance 25 or 26, or any positions of distance greater than 27 exist, they lack any symmetry; this increases our confidence (but does not prove) that God's number in the quarter-turn metric is 26.

d	positions	mod M	mod $M + \text{inv}$	mod 48
0	1	1	1	1
1	18	2	2	18
2	51	5	5	3
3	312	14	8	24
4	1,335	61	42	39
5	4,380	191	115	12
6	17,782	784	461	22
7	70,188	2,974	1,631	12
8	229,336	9,729	5,294	40
9	851,139	35,615	18,604	3
10	2,989,204	125,103	65,125	4
11	9,732,164	406,615	208,603	20
12	35,024,904	1,461,865	746,313	24
13	122,054,340	5,089,766	2,581,241	36
14	436,197,214	18,183,156	9,195,333	46
15	1,763,452,505	73,491,944	37,037,332	41
16	8,035,307,127	334,845,106	168,373,488	39
17	37,542,012,922	1,564,371,587	785,249,161	10
18	95,387,902,305	3,974,821,633	1,994,938,045	33
19	21,267,102,443	886,299,819	446,607,275	11
20	1,091,994	46,514	32,625	42
total	164,604,041,664	6,859,192,484	3,445,060,704	0

Table 2: Depth distribution in the half-turn metric.

d	positions	mod M	mod $M + \text{inv}$	mod 48
0	1	1	1	1
1	12	1	1	12
2	18	3	3	18
3	108	5	3	12
4	411	19	14	27
5	1,104	46	25	0
6	3,744	163	102	0
7	11,760	490	258	0
8	36,731	1,582	902	11
9	111,144	4,632	2,370	24
10	358,138	15,054	7,908	10
11	1,028,848	42,895	21,647	16
12	3,266,949	136,691	70,284	21
13	9,443,588	393,602	198,149	20
14	29,201,318	1,218,544	618,945	38
15	83,765,676	3,490,523	1,752,632	12
16	268,136,523	11,177,948	5,644,771	27
17	819,440,112	34,146,994	17,140,681	0
18	3,083,699,868	128,516,587	64,687,074	12
19	11,628,867,276	484,563,973	243,091,188	12
20	41,538,350,563	1,730,948,894	869,224,590	19
21	67,617,360,740	2,817,563,628	1,413,565,363	20
22	37,373,063,137	1,557,455,774	783,700,221	1
23	2,147,815,036	89,510,797	45,330,245	28
24	78,820	3,635	3,324	4
25	36	2	2	36
26	3	1	1	3
tot	164,604,041,664	6,859,192,484	3,445,060,704	0

Table 3: Depth distribution in the quarter-turn metric.

When we started this work, we hoped to solve the quarter-turn, half-turn, and also the slice-turn metric; our code was general across all three of these metrics. Unfortunately, the coset solvers in the slice-turn metric did not perform as effectively, so we terminated the programs after a few days of runtime. Further investigation will be required to determine why our technique was so effective in the half-turn and quarter-turn metric but not in the slice-turn metric.

7 Discussion

Solving all symmetric positions in the quarter-turn metric accomplishes several goals. First, it enhances our confidence that God's number in the quarter-turn metric is 26, since symmetric positions have a higher incidence of high distances than non-symmetric positions. In addition, it shows that if there are positions of distance more than 26, there must be at least a multiple of 48 of them, so our extensive searches through trillions of positions would likely have encountered one of them, one of their neighbors, or another position at distance two from one of them; yet, we have not found any additional distance-25 positions beyond the neighbors of the known distance-26 positions.

Since unsymmetrical positions occur in multiples of 48, and since we have solved all symmetrical positions in both the half-turn metric and the quarter-turn metric, we know the value of the count of positions at each distance in the quarter-turn and half-turn metric, modulo 48. These are given in the rightmost columns of Tables 2 and 3, and they help confirm the known values of positions through distance 15 in the half-turn metric and through distance 17 in the quarter-turn metric.

About 90% of the known distance-24 positions (and all known distance-25 and distance-26 positions) exhibit antisymmetry. The logical next step is to determine a method to solve antisymmetric positions in some efficient way; there are many more antisymmetric positions than symmetric positions, but this is still only a tiny fraction of overall cube space so it may be a fertile ground to search. Current results are available at `http://cube20.org/distance20s/`.

Of course at some point we will want to prove God's number in the quarter-turn metric. At this point, we estimate it will take about the same amount of CPU time that the face-turn metric did, so perhaps we will wait a few years for a better technique or faster machines.

We continue to search for distance-20's in the half-turn metric and distance-24's in the quarter-turn metric by solving large cosets. We can typically find 200,000 distance-20's in the half-turn metric in a day of computation (with four desktop machines); we typically find less than fifty distance-24's in the quarter-turn metric on average in an equivalent day.

References

[1] Joyner, David. *Adventures in Group Theory: Rubik's Cube, Merlin's Magic & Other Mathematical Toys.* Baltimore: The John Hopkins University Press, 2008.

[2] Kociemba, Herbert. Symmetric Patterns.
http://kociemba.org/math/symmetric.htm.

[3] Korf, Richard E. "Finding Optimal Solutions to Rubik's Cube Using Pattern Databases." *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97)*, Providence, RI, July, 1997, pp. 700–705.

[4] Radu, Silviu. "How to Compute Optimal Solutions for All 164,604,041,664 Symmetric Positions of Rubik's Cube." 27 July 2006.
http://cubezzz.dyndns.org/drupal/?q=node/view/63

[5] Rokicki, Tomas; Kociemba, Herbert; Davidson, Morley; and Dethridge, John. "The Diameter of the Rubik's Cube Group is Twenty." *SIAM J. Discrete Math.* (2013) vol 27 n. 2 pp. 1082–1105.