


SCIENTIFIC AMERICAN



MAGIC CUBOLOGY

March 1981

METAMAGICAL THEMAS

*The Magic Cube's cubies are twiddled
by cubists and solved by cubemeisters*

by Douglas R. Hofstadter

Cubitis magikia, *n.* A severe mental disorder accompanied by itching of the fingertips that can be relieved only by prolonged contact with a multicolored cube originating in Hungary and Japan. Symptoms often last for months. Highly contagious.

What this stuffy medical-dictionary entry fails to mention is that contact with the multicolored cube not only cures the itchiness but also causes it. Furthermore, it fails to point out that the affliction can be highly pleasurable. I ought to know; I have suffered from it for the past year and still exhibit the symptoms.

Bűvös Kocka—the Magic Cube, also known as Rubik's Cube—has simultaneously taken the puzzle world, the mathematics world and the computing world by storm. Seldom has a puzzle so fired the imagination of so many people, perhaps not since Sam Loyd's famous 15 Puzzle, which caused mass insanity when it came out in the 19th century and is still one of the world's most popular puzzles. The 15 Puzzle and the Magic Cube are spiritual kin, the one being a two-dimensional problem of restoring the scrambled numbered pieces of a 4×4 square to their proper positions and the other being a three-dimensional problem of restoring the scrambled colored pieces of a $3 \times 3 \times 3$ cube to their proper positions. The solutions of both demand that the solver be willing to seemingly undo precious progress time and time again; there is no route to the goal that does not call for a partial but temporary destruction of the order achieved up to that point. If this is a difficult lesson to learn with the 15 Puzzle, it is much harder with the Magic Cube. And both puzzles have the fiendish property that well-meaning bumblebers or cunning rogues can take them apart and put them back together in innocent-looking positions from which the goal is absolutely unattainable, thereby causing the would-be solver indescribable anguish.

This Magic Cube is much more than just a puzzle. It is an ingenious mechani-

cal invention, a pastime, a learning tool, a source of metaphors, an inspiration. It now seems an inevitable object, but it took a long time to be discovered. Somehow, though, the time was ripe, because the idea germinated and developed nearly in parallel in Hungary and Japan and perhaps even elsewhere. A report surfaced recently of a French government official who remembers having encountered such a cube, made out of wood, in 1920 in Istanbul and then again in 1935 in Marseilles. Without confirmation the claim seems dubious. In any event Rubik's work was completed by 1975, and his Hungarian patent bears that date. Quite independently, however, Terutoshi Ishige, a self-taught engineer and the owner of a small ironworks near Tokyo, came up with much the same design within a year of Rubik and filed for a Japanese patent in 1976. Ishige also deserves credit for this wonderful insight.

Who is Rubik? Ernő Rubik is a teacher of architecture and design at the School for Commercial Artists in Budapest. Seeking to sharpen his students' ability to visualize three-dimensional objects, he came up with the idea of a $3 \times 3 \times 3$ cube any of whose six 3×3 faces could rotate about its center, yet in such a way that the cube as a whole would not fall apart. Each face would initially be colored uniformly, but repeated rotations of the various faces would scramble the colors horribly. Then his students had to figure out how to undo the scrambling.

When I first heard the cube described over the telephone, it sounded like a physical impossibility. By all logic it ought to fall apart into its constituent "cubies" (one of the many useful and amusing terms invented by "cubists" around the world). Take any corner cubie—what is it attached to? By imagining rotating each of the three faces to which it belongs you can see that the corner cubie in question is detachable from each of its three edge-cubie neighbors. How then is it held in place? Some people postulate magnets, rubber bands or elaborate systems of twisting wires in

the interior of the cube, yet the design is remarkably simple and involves no such items.

In fact, the Magic Cube can be disassembled in a few seconds [see bottom illustration on page 17], revealing an interior structure so simple that one has to ponder how it can do what it does. It actually does fall apart. To see what holds it together first observe that there are three types of cubie: six center cubies, 12 edge cubies and eight corner cubies. The center cubies have only one "facelet," the edge cubies have two facelets and the corner cubies have three. Moreover, the six center cubies are really not cubical at all—they are just façades attached to axles that issue from a sixfold spindle at the middle. The other cubies are nearly complete cubes, except that each one has a blunt little "foot" reaching toward the middle of the cube and some curved nicks facing inward.

The basic trick is that cubies mutually hold one another in by means of their feet, without any cubie's being attached to any other. Edge cubies hold corner cubies' feet, corner cubies hold edge cubies' feet. Center cubies are the keystones. As any layer, say the top one, rotates it holds itself together horizontally and is held in place vertically by its own center and by the equatorial layer below it. The equatorial layer has a sunken circular track (formed by the nicks in its cubies) that guides the motion of the upper layer's feet and helps to hold the upper layer together.

In his definitive treatise "Notes on Rubik's 'Magic Cube'" David Singmaster, professor of mathematical sciences and computing at the Polytechnic of the South Bank in London, defines the "basic mechanical problem" as that of figuring out how the cube is constructed. I sometimes wonder whether Rubik's intended visualization task for his students was to solve the unscrambling problem (Singmaster calls it the "basic mathematical problem") or to solve the mechanical problem. I suspect the latter is the harder of the two. I myself must have put in more than 50 hours of work, distributed over several months, before I solved the unscrambling problem, and I never did solve the mechanical problem until I saw the cube disassembled. Singmaster informally estimates that people who eventually solve the unscrambling problem (without hints) take on the average two weeks of concentrated effort. Of course, it is hard for anyone who has done it to say exactly how long it took (how can you tell play from work?), but it is safe to say that if you are destined to solve the unscrambling problem at all, it will take you somewhere between five hours and a year. I trust this is reassuring.


An important fact many people fail to appreciate at first is that restoring the cube to the Start position (the state

where each face is a solid color) is so hard that it is necessary to find a general algorithm for doing it from any scrambled state. No one can restore a messed-up Magic Cube to its pristine state by mere trial and error. Anyone who gets back to the Start position has built up a small science.


A word of warning: Proposed solutions to the mechanical problem are of-

ten lacking in clarity, having either too much detail or too little. It is certainly a challenge to come up with a mechanism that has the multifaceted twistability of the Magic Cube, but it is perhaps no less of a challenge to describe the mechanism in language and diagrams other people can readily comprehend. By the same token, to convey algorithms that restore the cube to Start calls for a good,


clear notation. Singmaster himself has an excellent notation that is now considered standard; I shall present it below. A second word of warning: I am not a "cubemeister" (defined as one who has contributed to the profound science of cubology). I am a mere cubist, an amateur who is amazed by the cube and by the virtuosos who have mastered it. Therefore I am not a suitable recipient




The Magic Cube in the Start position




A scrambled cube




The "pretty pattern" called Dots




The Giant Meson with Cherries position




The Pons Asinorum configuration




The Christman-cross configuration



The Plummer-cross configuration



An alternate coloring requiring 12 colors




An alternate coloring requiring eight colors



F stands for face, E for edge, C for corner

of novel solutions to the mechanical problem or the unscrambling problem, of which by now there are hundreds. I recommend that readers who believe they have some novel insight communicate it to Singmaster, who periodically updates his booklet. It will soon be distributed to bookstores in this country by Enslow Publishers of Hillside, N.J., at \$5.95. Singmaster's address is Department of Mathematical Sciences and Computing, Polytechnic of the South Bank, London SE1 0AA, England.

The reader's appetite has by now, I trust, been whetted to the point where immediate possession of a Magic Cube is an urgent priority. Fortunately this can be arranged quite easily. Magic Cubes are being manufactured both by Logical Games, Inc., and by the Ideal Toy Corporation. The name Rubik's Cube applies only to the Ideal version, but in all intrinsic ways either version is equally "magic." Cubes are available by mail order from Logical Games (4509 Martinwood Drive, Haymarket, Va. 22069) for \$9 postpaid in the U.S. and at many toy and department stores for about \$10. It is likely that many people will buy them, little suspecting the profound difficulty of the "basic mathematical problem." They will innocently



Pattern on any face of the Slice Group

turn four or five faces and suddenly find themselves hopelessly lost. Then, perhaps frantically, they will begin turning face after face one way and then another, as it dawns on them that they have irretrievably lost something precious. They feel a little like a child watching a toy balloon drift into the sky.

It is a fact that the cube can be randomized with just a few turns. Let that be a warning to the beginner. Many beginners try to claw their way back to Start by first getting a single face done. Then, a bit stymied, they leave the partially solved cube lying around where a friend may spot it. The well-known "Don't touch it" syndrome sets in when the friend picks it up and says, "What's this?" The would-be solver, terrified that the hard-won progress will be destroyed, shrieks, "Don't touch it!" Ironically, victory can come only through a more flexible attitude allowing precisely that destruction.

For the beginner there is an awesome sense of irreversibility about destroying the Start position, a kind of fear of tumbling off the edge of a precipice. When my own first cube (I now have five) was first messed up (by a guest), I felt both relieved (because it was inevitable) and sad (because I feared the Start position was gone forever). The physicist in me was reminded of entropy. Once the Start position had become irretrievable each new twist of one face or another seemed irrelevant. To my naive eye there was no distinguishing one messed-up state from another, just as to the eye there is no distinguishing one plate of spaghetti from another, one pile of fall leaves from another and so on. The details meant nothing to me, hence they did not register. As I performed my "random walk" the vastness of the space of possible shufflings of the little cubies became vivid.

As with a deck of cards, one can calculate the exact number of possible rearrangements of the cube. An initial estimate would run this way. The first observation—a rather elementary one—is that on the rotation of any face each corner goes to another corner, each edge to another edge and the center of the face stays put (except for its invisible rotation). Therefore corners mix only with their own kind, and the same goes for edges. There are eight corner cubies and eight corner "cubicles" (the spatial niches, regardless of their content). Cubies and cubicles are to the cube as children and chairs are to the game of musical chairs. Each corner cubie can be maneuvered into any of the eight corner cubicles. This means that we have eight possible fillers for cubicle No. 1, seven for cubicle No. 2, six for cubicle No. 3 and so on. Hence the corners can be placed in their cubicles in $8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1 (= 8!)$ different ways. Each corner, however, can be in any one

of three orientations. Thus one would expect a further factor of 3^8 from the corners. One would expect the same for the 12 edge cubies: 12 objects can be permuted among themselves in $12!$ different ways, and then if each of them has two possible orientations, that gives another factor of 2^{12} . The center cubies never leave their Start positions (unless the cube is rotated as a whole), hence they do not contribute. If we multiply the numbers out, we get 519,024,039,293,878,272,000 possible positions, about 5×10^{20} .

But there is an assumption here: that any cubie can be got to any cubicle in any orientation, regardless of the other cubies' positions and orientations. As we shall see, this is not quite the case. It turns out there is a constraint on the orientation of the corner cubies: any seven can be oriented arbitrarily, but the last one is then forced, thus removing a factor of three. Similarly, there is a constraint on edge cubies: of the 12 any 11 can be oriented arbitrarily, but the last one is then determined, so that another factor of two is removed. There is one final constraint on the permutations of cubies (disregarding their orientations) that says you can place all but two of them wherever you want but the last two are forced. This removes a final factor of two, reducing the estimate above by a factor of 12, bringing the possibilities down to a mere 43,252,003,274,489,856,000, about 4×10^{19} . Still, it must be said, this does slightly exceed the assertion on Ideal's label: "More than three billion combinations."

Another way of thinking about this factor of 12 is that if you begin at Start, you are limited to a twelfth of the "obvious" states, but if you disassemble your cube and reassemble it with a single corner cubie twisted by 120 degrees, you are now in a formerly inaccessible state, from which an entire family of 43,252,003,274,489,856,000 new states is accessible. There are 12 such nonoverlapping families of states of the cube, usually called orbits by group theorists.


Speaking of impossible twists, I should like to mention a lovely parallel in particle physics that was pointed out by Solomon W. Golomb (a familiar name to many of Martin Gardner's readers). It is impossible to make any sequence of moves that will leave just one corner cubie twisted a third of a full turn and everything else the same. Now, recalling the famous hypothetical fundamental particle with a charge of $+1/3$ and its antiparticle with a charge of $-1/3$, Golomb calls a clockwise one-third twist a quark and a counterclockwise one-third twist an antiquark. Like their cubical namesakes, quark particles have proved to be tantalizingly elusive, and many theoretical physicists now believe in quark confinement: the principle

that it is impossible to have an isolated free quark (or antiquark). This correspondence between cubical quarks and particle quarks is a nice one.

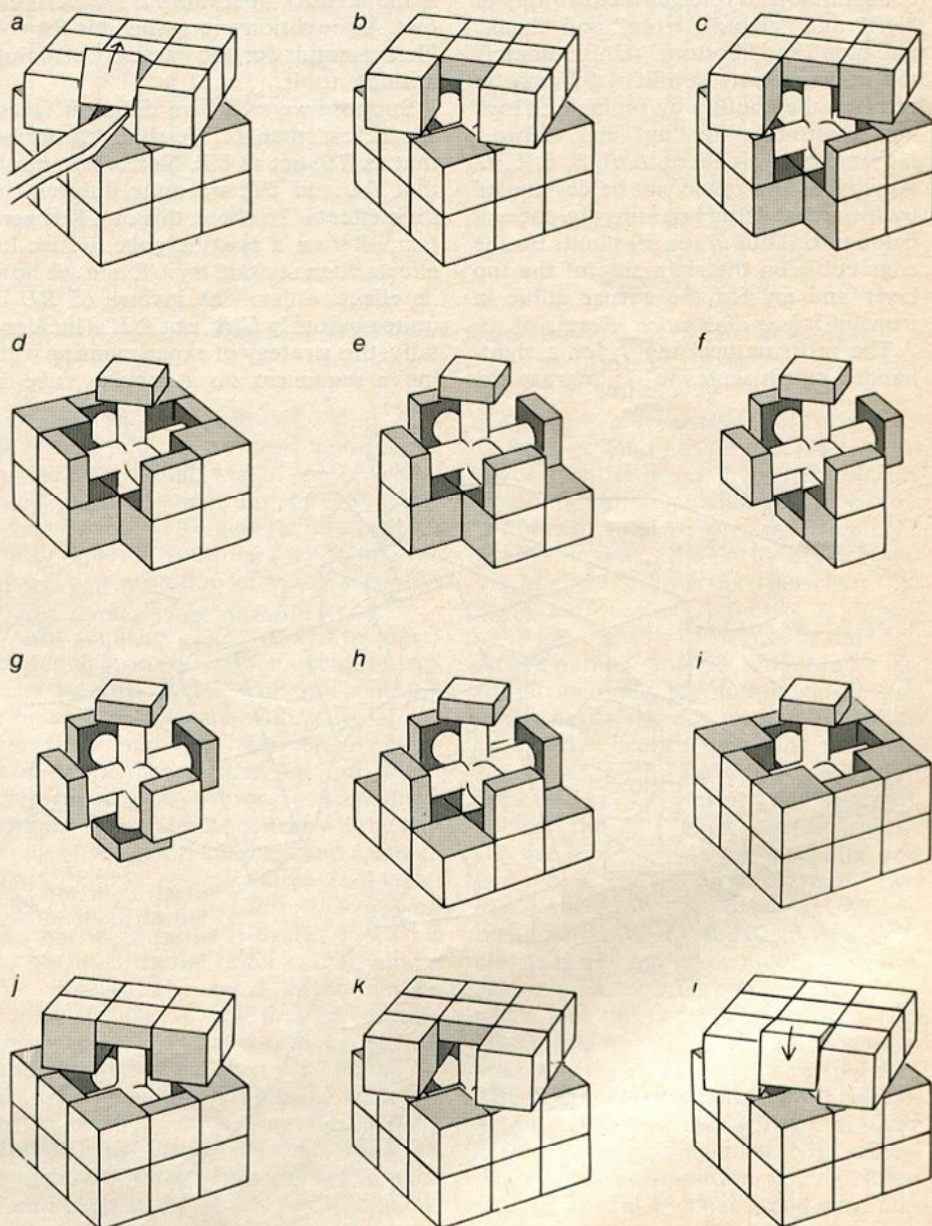
Actually the connection runs even deeper. Quark particles cannot exist free, but they can exist bound together in groups: a quark-antiquark pair is a meson, and a quark trio with integral charge is a baryon. (An example is the proton, with a charge of +1.) Now, in the Magic Cube, amazingly enough, it is possible to give two corner cubies one-third twists, provided they are in *opposite* directions (one clockwise, the other counterclockwise). It is also possible to give three corner cubies one-third twists, provided they are all in the *same* direction. Thus Golomb calls a state with two oppositely twisted corners a meson and one with three corners twisted in the same direction a baryon. In the particle world only quark combinations with an integral amount of charge can exist. In the cubical world only quark combinations with an integral amount of twist are allowed. That is just another way of saying the orientation of the eighth corner cubie is always forced by the first seven. In the cubical world the underlying reason for quark confinement lies in group theory. There may be a closely related group-theoretical explanation for the confinement of quark particles. That remains to be seen, but in any event the parallel is provocative and pleasing.

If we have a pristine cube (one in the Start position), what kind of move sequence will create a meson or a baryon? Here we have an example of the most powerful idea in cubology: the idea of "canned" move sequences that accomplish some specific reordering of a few cubies, leaving everything else untouched (invariant, as group theorists say). There are many different terms for such canned move sequences. I have heard them called operators, transforms, words, tools, processes, maneuvers, routines, subroutines and macros, the first three being group-theoretical terms and the last three being adapted from computer-science argot. Each term has its own flavor, and I like to use them all.

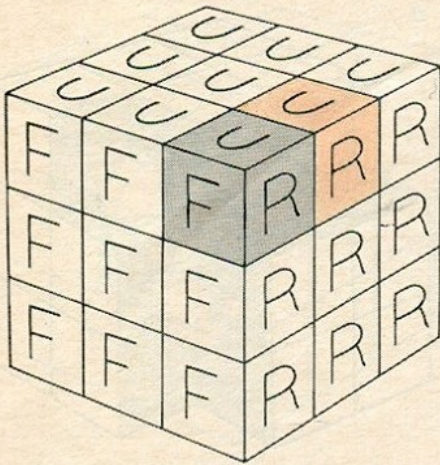
In order to talk about processes we need precision, and that means a good technical notation. I shall therefore present Singmaster's notation now. First we need a way of referring to any particular face of the cube. One possibility is to use the names of colors as the names of the faces, even after the colors of the cubies have become mixed up. It might seem that calling a face white, say, would be meaningless if white is scattered all over the place. Remember, however, that the white center cubie never moves with respect to the five other center cubies, and so it defines the "home face" for white. Then why not use color names for faces? The problem



The spindle (a), an edge cubie (b), a corner cubie (c) and a view of the internal mechanism (d)



Disassembly (a-f) and assembly (g-i). Always reestablish original Start position!



The urf cubie (gray) and the ur cubie (color)

is that different cubes come with their colors arranged differently. Even two cubes from one manufacturer may have different Start positions. A more general convention is to refer to faces simply as "left" and "right," "front" and "back" and "top" and "bottom." Unfortunately the initials create conflicts. Singmaster resolves the conflict by replacing "top" and "bottom" with "up" and "down." Now we have six faces: *L, R, F, B, U, D*. Any particular cubie can be designated by lowercase letters naming the faces it belongs to. Thus *ur* (or *ru*) stands for the edge cubie on the right side of the top layer and *urf* for the corner cubie in front of it [see illustration above].

The most natural move for a right-handed cubist seems to be to grasp the


right face with the thumb pointing up along the front face and to move the thumb forward. Seen from the right side this maneuver causes a clockwise quarter twist of the *R* face. This move will be designated *R* [see illustration below]. The mirror-image move, where the left hand turns the *L* side counterclockwise (as seen from the left), is L^{-1} , or for short *L'*. A clockwise twist of the *L* side is called, naturally, *L*. A 90-degree clockwise turn of any face (from the point of view of an observer looking at the center of that face) is named by the letter for that face, and its inverse—the counterclockwise quarter turn—has a prime mark (') or a superscript -1 following the face's initial. Quarter turns will henceforward be called *q* turns.

With this nomenclature we can now write any move sequence, no matter how complex. A trivial example is four successive *R*'s, which we shall write as R^4 . In the language of group theory this is the identity operation: it has zero effect. An equation expressing it is $R^4 = I$. Here *I* stands for the "action" of doing nothing at all.

Suppose we twist two different faces, say *R* first, then *U*. We shall transcribe that as *RU*, not as *UR*. Notice first of all that *RU* and *UR* are quite different in their effects. To check this out, first perform *RU* on a pristine cube, notice its effects, then undo it, try *UR* and see how its effects differ. The inverse of *RU* is quite obviously $U'R'$, not $R'U'$. (Incidentally, this strategy of experimenting with move sequences on a pristine cube is


most helpful. Very early I found it useful to buy a second cube so that I could work on solving one while experimenting with the other, never letting the second one get far from Start.)

What is the effect of a particular "word"? That is to say, which cubies move where? To answer the question we need a notation for the motions of individual cubies. The effect of *R* on edges is to carry the *ur* cubie around to the back face to occupy the *br* cubicle. At the same time the *br* cubie swings around underneath, landing in the *dr* position, the *dr* cubie moves up like a car on a Ferris wheel to fill the *fr* cubicle and the *fr* cubie comes to the top at *ur*. Symbolically this becomes either of the following:



This is called a 4-cycle, and we shall write it in a more compact way: (*ur, br, dr, fr*). Of course, it does not matter where we start writing; we could equally well write (*br, dr, fr, ur*). On the other hand, the order of the letters in cubie names does matter. We can reverse all of them or none of them but not just some of them. If you think of the letters as designating facelets, it will become clear. For example, if we wrote (*ur, rb, dr, rf*), it would represent a 4-cycle involving the same four cubicles, but one in which each cubie flipped before moving from one cubicle to the next. Of course, such a cycle cannot be accomplished by a single *q* turn, but it may be the result of a sequence of *q* turns of different faces (an operator). Or consider the following 8-cycle: (*ur, uf, ul, ub, ru, fu, lu, bu*). This has length eight but involves only four cubies. Each cubie, after making a full swing around the top face, comes back flipped. After two full swings it is back as it started. Each facelet has made a "Möbius trip." We can designate this "flipped 4-cycle" as (*ur, uf, ul, ub*)₊, where the plus sign designates the flipping. The designation (*ru, fu, lu, bu*)₊ and numerous others would do as well. Thus the cycle notation tells you not only where a cubie moves but also its orientation with respect to the other cubies in its cycle.

To complete our description of the effect of *R* we must transcribe the 4-cycle of the corners. As with edges we have the freedom to start at any corner we want, and once again we must be careful to keep track of the facelets so that we get the orientations right. Still, *R* has a rather trivial effect on corners: (*urf, bru, drb, frd*), which could also be written (*rub, rbd, rdf, rfu*) and many other ways. Summing up, we can write $R = (ur, br, dr, fr)(urf, bru, drb, frd)$. This says that *R* consists of two disjoint 4-cycles. (If we



R, a natural move for a right-handed cubist


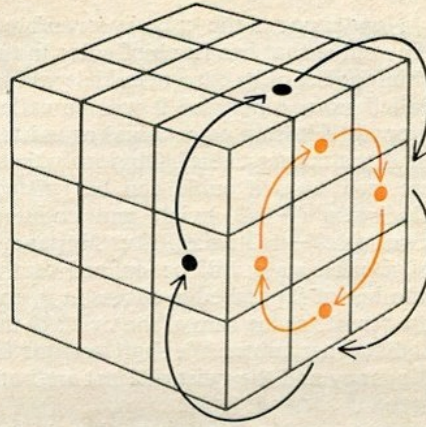
wanted to, we could throw in a term standing for the 90-degree rotation of the R face's center, but since such rotation is invisible, we shall not do so.)

What about transcribing a move sequence such as RU ? Well, take a pristine cube and perform RU . Then start with some arbitrary cubie that has moved and describe its trajectory. For example, ur has moved to br . Therefore br has been displaced. Where has it gone? Find the new location of that cubie (it is dr) and continue chasing cubies around and around the cube until you find the one that moved into the original position of ur . You will find the following 7-cycle: $(ur, br, dr, fr, uf, ul, ub)$ [see illustration on next page].


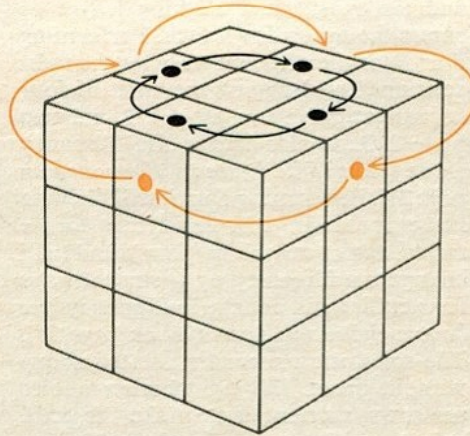
What about corners? Well, suppose we trace the cubie that originated in urf . Where did RU carry it? The answer is: nowhere; it took a round trip but got twisted along the way. It changed into rfu . We can designate this clockwise twist—this quark— $(urf)_+$. This "twisted unicycle" is shorthand for the following 3-cycle: (urf, rfu, fur) . You can even see this as cycling the three letters u, r and f inside the cubie's name. If the cycle had been an antiquark, we would have written $(urf)_-$ and the letters would cycle the other way.

What about the other seven corners? Two of them— dbl and dlf —stay put and the other five almost form a 5-cycle: $(ubr, bdr, dfr, luf, bul)$. It is unfortunate that the cycle does not quite close, because bul , although it gets carried into the original ubr cubicle, does so in a twisted manner. It gets carried to rub , which is a counterclockwise twist away from ubr . This means we are dealing with a 15-cycle. It is so close to the 5-cycle above, however, that we shall just add a minus sign to represent the counterclockwise twist. Our twisted 5-cycle is then $(ubr, bdr, dfr, luf, bul)_-$, and the entire effect of RU , expressed in cycle notation, is $(ur, br, dr, fr, uf, ul, ub)(urf)_+(ubr, bdr, dfr, luf, bul)_-$.

Now that we have RU in cycle notation we can perform rotations mentally by sheer calculation. What would be the effect, for instance, of $(RU)^5$? Edge cubie ur would be carried five steps forward along its cycle, which would bring it to ul . (This can also be seen as moving two steps backward.) Then ul would go to fr and so on. The 7-cycle is replaced by a new 7-cycle: $(ur, ul, fr, br, ub, uf, dr)$. Let us now look at the twisted 5-cycle. Corner cubie ubr would be carried five steps forward along its cycle, which brings it back to itself negatively twisted, namely rub . Similarly, all the corner cubies in the 5-cycle would return to their starting points, but negatively twisted; thus on being raised to the fifth power a negatively twisted 5-cycle becomes five antiquarks. But if that is so, how is the requirement of integral twist satisfied? Do we not have one quark—



The 4-cycle (ur, br, dr, fr) is at the left, the 4-cycle (ur, rb, dr, rf) at the right



The 4-cycle (ur, uf, ul, ub) is at the left, the flipped 4-cycle $(ur, uf, ul, ub)_+$ at the right

$(urf)_+$ —and five antiquarks, and does that not add up to four antiquarks, which have a total twist of $-1\frac{1}{3}$? Well, I have slipped something by you here. Can you spot it? To gain facility with the cycle notation you might try to find the cycle representation of various powers of RU and UR and their inverses.


Any sequence of moves can be represented in terms of disjoint cycles of various lengths: cycles with no common elements. If you are willing to let cycles share members, however, any cycle can be further broken up into 2-cycles (called transpositions, or sometimes swaps). For example, consider three animals: an alligator, a bobcat and a camel. They initially occupy three ecological niches: A, B and C [see illustration on page 22]. The effect of the 3-cycle (A, B, C) is to put them in the order camel, alligator, bobcat. The same effect can be achieved, however, by first performing the swap (A, B) (what was in A goes to B and vice versa) and then performing (A, C) . Of course, this can also be achieved by the two successive swaps $(A, C)(B, C)$ and $(B, C)(A, B)$. On the other hand, no sequence of three swaps will achieve the same effect as (A, B, C) . Try it yourself and see. (Note that a niche is like a cubicle and an animal is like a cubie.)

An elementary theorem of zoo theory (a field we shall not go into here) states that no matter how a given permutation of animals among niches is reduced to a product of successive swaps (which can always be done) the parity of the number of such swaps is invariant, that is to say, a permutation cannot be expressed as an even number of swaps one time and an odd number another time. Moreover, the parity of any permutation is the sum of the parities of any permutations into which it is broken up (using the rules for the addition of even and odd numbers: odd plus even is odd and so forth).

Now, this theorem has repercussions for the Magic Cube. In particular you can see that any q turn consists of two disjoint 4-cycles (one on edges and one on corners). What is the parity of a 4-cycle? It is odd, as you can work out for yourself. Hence after one q turn both the edges and the corners have been permuted oddly, after two q turns evenly, after three q turns oddly and so forth. The edges and corners stay in phase, in the sense that the parities of their permutations are identical. Now, clearly the null permutation is even (zero swaps). Thus if we have a null permutation on corners, the permutation on edges must also be even. Conversely, a

null permutation on edges implies an even permutation on corners. Imagine a state identical with Start except for two interchanged edges (that is, one swap). The state is even in corners but odd in edges, hence it is impossible. The best we could do would be to have *two pairs* of interchanged edges. The same argument holds for corners. In short, single swaps are impossible; swaps must always come in pairs. (This is the origin of one of our factors of two in the calculation above of the number of states of the cube.) There are processes for exchanging two pairs of edges, two pairs of corners and even for exchanging one pair of edges along with one pair of corners. (This last process necessarily involves an odd number of q turns.)

To round out the subject of constraints let us ponder the origin of the constraints on corner twisting and edge flipping. Here is a clever explanation provided by the mathematicians John Horton Conway, Elwyn R. Berlekamp and Richard K. Guy, building on ideas of Anne Scott. The basic concept is that we want to show that the number of flipped cubies is always even and that the twist is always integral. In order to determine what is flipped and what is twisted, however, we need a frame of reference. To supply it we shall define two things called the "chief facetlet of a cubicle" and the "chief color of a cubie." (Remember that a cubicle is a niche and a cubie is a solid object.) The chief facetlet of a cubicle will be the one on the up or down surface of the cube, if there is one; otherwise it will be the one on the left or right wall [see top figure in illustration on page 24]. There are nine chief facetlets on U , nine on D and four on the equator. We can forget about centers, however, because they never can be flipped or twisted. The chief color of a cubie is defined as the color that should be on the cubie's chief facetlet when the cubie "comes home" to its proper cubicle in the Start position.



The 7-cycle (ur, br, dr, fr, uf, ul, ub)

Now suppose the cube is scrambled. Any cubie that has its chief color in the chief facetlet of its current cubicle will be called sane; otherwise it will be called flipped (if it is an edge cubie) or twisted (if it is a corner cubie). Obviously there are two ways a cubie can be twisted: clockwise (+1/3 twist) and counter-clockwise (-1/3 twist). The "flippancy" of a cube state will be defined as the number of flipped edge cubies in it, and the "twist" as the sum of the twists of the eight corner cubies. We shall say that the flippancy and the twist of Start are both zero.

Next consider the 12 possible q turns out of which everything else is compounded. Performing U or D (or their inverses) preserves both the flippancy and the twist, since nothing leaves or enters the up or down face. Performing F or B (or their inverses) leaves the twist constant, by changing the twist of four corners at once: two by +1/3 and two by -1/3. It also leaves the flippancy alone [see middle figure in illustration on page 24]. Performing L or R will likewise leave the twist constant (four corner twists again cancel in pairs) and will change the flippancy by 4, since always four cubies will change in flippancy [see bottom figure in illustration on page 24]. The conclusion is what we stated above without proof: the eight corner cubies are always oriented to make the total twist a whole number, and the 12 edge cubies must always be oriented to make the total flippancy even.

After this discussion of constraints you should be convinced that no matter how you twist and turn your Magic Cube you cannot reach more than a twelfth of the conceivable "universe," beginning at Start. It is another matter, however, to show that every state within that one-twelfth universe is accessible from Start (or what amounts to the same thing, only backward: that Start is accessible from every state in the one-twelfth universe). For this we need to show how to achieve all even permutations of cubies and how to achieve all orientations that do not violate the two constraints described above. What it comes down to is that we have to show there are operators that will perform seven classes of operations: (1) an arbitrary double edge-pair swap, (2) an arbitrary double corner-pair swap, (3) an arbitrary two-edge flip, (4) an arbitrary meson, (5) an arbitrary 3-cycle of edges, (6) an arbitrary 3-cycle of corners and (7) an arbitrary baryon.

Of course, each of these operators should work without causing side effects on any other parts of the cube. With these powerful tools in our kit we would be able to cover the one-twelfth universe without any trouble. In the case of the overlapping swaps of animals, I showed that a 3-cycle is really two overlapping 2-cycles. This means that classes 5 and 6

can be made out of the first four classes. Similarly, a baryon can be made from two overlapping mesons. Hence all we really need is the first four classes.

To show that all the operators belonging to these four classes are available we use another of the most crucial and lovely ideas of cubology: that of conjugate elements. It turns out that all we need is one example in each class; given one example, we can construct all the other operators of its class from it. How does this work? The idea is very simple.

Suppose we had found one operator in class 1 that swapped, say, uf with ub and ul with ur , leaving the rest of the cube undisturbed [see colored arrows in illustration on page 25]. Let us call it H . Now suppose we wanted to swap two entirely different pairs of edge cubies, say rf with df and rb with dr [see black arrows in same illustration]. One daydreams: "If only those cubies were in the 'four magical swapping spots' on the top surface. . . ." Well, why not just put them up there? It would be fairly simple to get four cubies into four specific cubicles. The obvious objection is: "Yes, but that would have an awful side effect—it would totally mess up the rest of the cube." There is, however, a clever retort. Let the destructive maneuver that gets those four cubies into the magical swapping spots be called A . Suppose we were smart enough to transcribe the move sequence of A . Then right after performing A we perform our double swap H . Now comes the clever part. Reading our transcript in reverse order and inverting each q turn, we perform the exact inverse of A . This will not only unmaneuver the four cubies back into their old cubicles but also undo the side effects A created. Does that restore the cube intact? Not quite. Remarkably, since we sandwiched H between A and A' , the four edge cubies go home permuted, that is, each one winds up in the home of its swapping partner! Otherwise the cube is restored, and so we have accomplished precisely the double swap we set out to accomplish.

When you think this through, you see that it is flawless in conception. The inverse maneuver, A' , does not "know" we have exchanged two pairs of edges. As far as it is concerned, it is merely putting everything back where it was before A was executed. Hence we have "snuck" our swaps in under A' 's nose, which is to say we have "fooled the cube." Symbolically, we have carried out the sequence of moves AHA' , which is called a "conjugate" of H .

It is this kind of marvelously concrete illustration of an abstract notion of group theory that makes the Magic Cube one of the most amazing things ever invented for teaching mathematical ideas. Normally in group-theory courses the examples of conjugate elements are either too trivial or too ab-

stract to be enlightening or exciting. The Magic Cube provides a vivid illustration of conjugate elements and of many other important concepts of group theory.

Suppose you want to get a quark-antiquark pair on opposite corners but know how to do so only on adjacent corners. How do you do it? Here is a hint: There are two nice solutions, but the shorter and prettier one involves using a conjugate. Incidentally, any maneuver that creates a quark on one corner (with other effects, of course) might be called a quarkscrew.

What we have shown for edges goes also for corners: the ability to swap two *specific* corners enables you to swap *any* two corners. Conjugation allows you to build up an entire class of operators from any single member of that class. Of course, the question still remains: How do you find some sample operator in each of the four classes? For example, how do you find an operator that creates a meson on two adjacent corners (a combination of a quarkscrew and an antiquarkscrew)? How do you find an

operator that exchanges two edge pairs both of which are on the top surface? I shall not give the answer here but shall follow Singmaster, who points the way by suggesting quasi-systematic exploration of some small "subuniverses" within the totality of all cube states, that is, he suggests you look at *subgroups*. This means limiting your set of moves deliberately to some special types of move. Here are five examples of interesting subgroups created by various kinds of restriction:

1. The Slice Group. In this subgroup every turn of one face must be accompanied by the parallel move on the opposing face. Thus R must be accompanied by L' , U by D' and F by B' . The name comes from the fact that any such double move is equivalent to rotating one of the three central slices of the cube. Singmaster abbreviates the slice move RL' by R_s , $R'L$ by R'_s and so forth. With this restriction faces cannot get arbitrarily scrambled. Each face will have a pattern in which all four corners share one color. A special case is the pattern called

Dots, wherein each face is all one color except for its center [see illustration on page 15]. Can you figure out how to achieve Dots from Start? How many different ways are there of arranging the dots? How does the Dots pattern resemble a meson? (Note: The reader will find answers to all these questions, along with much else, in Singmaster's book. Please do not send me answers. Send novel ideas to Singmaster.)

2. The Slice-squared Group. Here we restrict the Slice Group further, allowing only squares of slice moves, such as R_s^2 (which is the same as R^2L^2) or F_s^2 (which is the same as F^2B^2).

3. The Antislice Group. Here instead of rotating opposing faces always in parallel we rotate them always in antiparallel, so that R is accompanied by L , F by B and U by D . An antislice move has a subscript a , as in R_a , which equals RL . (Of course, the Antislice-squared Group is no different from the Slice-squared Group.)


4. The Two-Faces Group. Allow yourself to rotate only two faces, say F and R .

5. The Two-Squares Group. As in the Two-Faces Group, you may rotate only two faces, using only 180-degree turns at that. This is a very simple subgroup.

If you limit your attention to just the Two-Faces and Two-Squares groups, you will be able to find processes that achieve double swaps, some of edges and others of corners. It is a remarkable fact that these processes alone, together with the notion of conjugation, will allow us—in a theoretical sense—to solve the entire unscrambling puzzle.

Why do we not also need a meson maker and a double edge flipper? Well, consider how we might make a double edge flipper from the two classes of tools one may assume will be found. In order to flip two edges without creating any side effects we shall perform two successive double edge-pair swaps, and both times they will involve the same pairs! For example, we might swap uf with ub and df with db and then re-swap them. This seems to be an absolute "nothing process," but that need not be the case. After all, just as before, we can sandwich the second swap between a process X and its inverse X' , where we carefully choose the process X so as to... (Oh dear, I totally lost my train of thought there. I am sure you can finish it up, though. I do remember that it was not too tricky, and that I thought the idea was elegant. I am sure you will too.)

The same kind of thinking will show how you can build up a meson maker out of mere corner-swapping processes and conjugation. Given mesons you can build up baryons. And with mesons and baryons, double edge flippers and double edge-pair swappers and double corner-pair swappers you have a full set of



An alligator, a bobcat and a camel (a, b, c) are permuted in their ecological niches (A, B, C)

tools with which to restore any scrambled cube to Start, as long as it belongs to the same orbit as Start. This is, needless to say, a highly theoretical existence proof, and any practical set of routines would be organized quite differently. The type of solution we have described has the advantage of being compact in description, but it is enormously inefficient. In practice a cube solver must develop a fairly large and versatile set of routines that are short, easy to memorize and highly redundant. There is an advantage to being able to carry out required transformations in a variety of

ways: one can choose whichever tool seems best adapted to the situation at hand, instead of, for instance, using some theoretical tool that can make a baryon in several hundred q turns.

The typical solver evolves a set of transforms partly by intuition, partly by luck, sometimes with the aid of diagrams and occasionally with abstract principles of group theory. One principle nearly everyone formulates quite early is that of "getting things out of the way." This is once again the idea of conjugates, only in a simpler guise. The typical patter that goes along with it is


something like this (I have included sound effects of a sort): "Let's see, I'll swing *this* out of the way [flip, flip] so that I can move *that* [flap, flap, flap], and now I can swing *this* back again [unflip, unflip]. There—now I've got *that* where I wanted it to be." You can hear the conjugate structure inside the patter.

The only problem with being conscious of why it all works as you carry it out is that it may be too taxing. My impression is that most cubemeisters do not think in such detail about how their tools are achieving their goals, at least not while they are in the midst of restoring some scrambled cube. Rather, expert cube solvers are like piano virtuosos who have memorized difficult pieces. As one M.I.T. cubemeister said to me, "I have forgotten how to solve the cube, but my *fingers* remember."

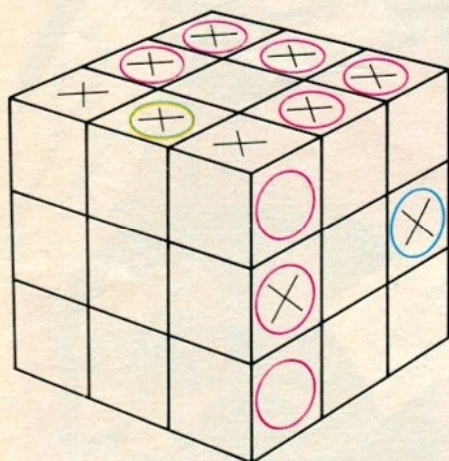
The average operator seems to be about 10 to 20 q turns long. You do not ever want to get lost in mid-operator, because if you do, you will have a totally scrambled cube on your hands, even if you were carrying out your final transform. As cubemeister Bernie Greenberg, who with Dave Christman is responsible for the beautiful computer graphics of the cube on the cover of this issue, said to me once, "If I were solving a cube and somebody yelled 'Fire!' I would finish my transform before clearing out."

My own style is probably overly blind. Not only do I not think about why my operators work as I am carrying them out; I must also admit that with some of them I have not the foggiest idea why they work at all. I found these "magic operators" through a long and arduous trial-and-error procedure. I used some heuristic approaches: "Explore various powers of simple sequences," "Use conjugates a lot" and so on. One thing I hardly used at all—alas, poor Rubik—was three-dimensional visualization. I do, however, know one Stanford cubemeister, Jim McDonald, who can give the reason for every q turn he makes. His operators do not seem magic to him because he can see what they are doing at every moment along the way. In fact, he does not have them memorized as I do mine; he seems to reconstruct them as he unscrambles cubes, relying on his "cube sense." He is like an expert musician who can improvise where a novice must memorize. For interested readers the central idea of Jim's method is first to solve the top layer, minus one corner, and then to let a vertical column of three cubicles act in the manner of a driveway one uses as an aid in turning a car around. The other two layers are cleaned up by shunting cubies in and out of the "driveway."


Perhaps not coincidentally, the abstract approach has been carried to its extreme by Singmaster's officemate Morwen B. Thistlethwaite (I am not jok-



The cube is in the Start position. The chief facelets of the cubicles are shown by the black X's and the chief colors of cubies by the colored O's. (The centers can be ignored, since they are stationary.) Think of the X's as floating in space and the O's as being attached to the cube, so that when the faces turn, the X's stay where they are but the O's move.



The q turn F has been executed. The two empty colored circles indicate that those two facelets have lost their "sanity." For them to be returned to sanity one would have to be twisted one-third clockwise and the other one-third counterclockwise. The same thing has happened on the invisible left-hand face.



The q turn R has been executed from the Start position. Empty circles occur in pairs. The top and bottom corner yellows have canceled twists. The blue edges from a pair, and the yellow edge is paired with a flipped edge on the invisible back face.

Proof that flippancy is even and twist is integral

ing), who currently holds the world record for the shortest unscrambling algorithm. It requires at most 52 "turns." (A turn is either a q turn or a half turn, that is, a 180-degree turn of one face.) Thistlethwaite has used ideas of group theory to guide a computer search for special kinds of transforms. His algorithm has the property of not seeming to converge toward the solved state at all—until the last few turns.

This must be contrasted with the more conventional style. Most algorithms begin by getting one layer, say the top layer, entirely correct. (In saying "top layer" rather than "top surface" I mean that the "fringe" has to be right, too: the cubies on top must be correct as seen from the side as well as from above.) This represents the first in a series of "plateau states." Although further progress requires any plateau state's destruction, that state will later be restored, and each time this happens more order will have been introduced. These are the successive plateau states.

After getting the top layer the solver typically works on corners on the bottom layer, or perhaps on getting the horizontal equator slice all fixed up. Most algorithms can, in fact, be broken up into about five natural stages, corresponding to natural classes of cubies that get returned to their home cubicles. My personal algorithm, for instance, goes through the following five stages: top edges, top corners, bottom corners, equator edges and bottom edges. In the first two of my stages placement and orientation are achieved simultaneously. Each of the last three stages breaks up into a placement phase and then an orientation phase. Naturally the operators of any stage must respect all the accomplishments of preceding stages. This means they can damage the order built up as long as they then repair it. They are welcome, however, to indiscriminately jumble up cubies scheduled to be dealt with in later stages. I find that other people's algorithms are usually based on the same classes of cubies, but the order of the stages can be wildly different.

Virtually all algorithms have the property that if one were to take a series of snapshots of a cube at the plateau states, one would see entire groups of cubies falling into place in patterns. This is called monotonicity at the operator level, that is, a steady, visible approach toward Start, with no backtracking. Of course, there is backtracking at the *turn* level, but that is another matter.


Very different in spirit is Thistlethwaite's algorithm. Instead of trying to put particular classes of cubies into their cubicles he makes a "descent through nested subgroups." This means that starting with total freedom he makes a few moves, then clamps down on the types of move that will thenceforward be allowed, makes a few more moves,

clamps down a bit more and so on until the constraints become so heavy that nothing can move any more. Just at this point, however, the Start position has been achieved! Each time, the clamping down amounts to forbidding q turns on two opposite faces, allowing only half turns thenceforward in their stead. The first faces to be thus "clamped" are U and D , then come F and B and finally L and R . The strange thing about this approach is that you cannot see Start getting nearer, even if you take a series of snapshots at carefully chosen moments. All of a sudden there it is.

This raises a natural question. Is there any easy way to tell how far you are from Start? It could obviously be quite useful. For example, it is rather embarrassing if one has to resort to the full power of a general unscrambling algorithm to undo what some friend has done with four or five casual twists. For that reason alone it would be nice to be able to assess quickly if some state is "really random" or is close to Start. But what does "close" mean? Distance between two states in this vast space can be measured in two fairly natural ways. You can count either the number of q turns or the number of turns needed to get from one state to the other (where "turn," as above, means either a q turn or a half turn). But how can one figure out how many turns are needed to get to Start without doing an exhaustive search? A reliable and at least fairly accurate estimate would be preferable, one that could be carried out quickly during a cursory inspection of the cube state. A naive suggestion is to count the number of cubies that are not in their home cubicle. This estimator, however, can be totally fooled by the Dots position, in which nearly all cubies are on the "wrong" side. That position is only eight q turns away from Start. Perhaps the flippancy and the number of quarks could also be taken into account by a better estimator, but I do not know of one.

There are sophisticated group-theoretical arguments estimating that the farthest one can get from Start is 22 or 23 turns. This is quite amazing, considering that most solvers' early algorithms take several hundred turns, and highly polished algorithms take a number somewhere in the 80's or 90's. Indeed, many mere operators take considerably more turns than Thistlethwaite's entire algorithm does. (My first double edge flipper was nearly 60 turns long.)

One thing we know, and it can be demonstrated easily, is that there exist states at least 17 turns away from Start. The argument goes as follows. At the outset there are 18 possible turns we might make: L, L', L^2, R, R', R^2 and so on. After that there are 15 reasonable turns to make. (One would not move the same face again.) The number of dis-



How to swap arbitrary edge pairs

tinct turn sequences of length 2 is therefore 18×15 , or 270. Another turn will add another factor of 15, and so on. How long does it take before we have reached the number of accessible states? It turns out that 17 turns will create more turn sequences than there are distinct states, and that 16 turns are too few. Now, not every turn sequence leads to a distinct state, not by a long shot, and so we have not shown that 17 turns will reach every accessible state. We have simply shown that *at least* 17 turns are needed if you want to reach every state from Start. Hence conceivably no two states are much more than 17 turns away from each other. But which 17 turns? That is the question.

So far only God knows how to get from one state of the Magic Cube to another in the fewest turns. "God's algorithm" is the one that would tell you how to do it. A burning question of cubology is: Is God's algorithm just a gigantic table without any pattern in it, or is there a significant amount of pattern to it, so that an elegant and short algorithm based on it could be mastered by a mere mortal?


If God were to enter a cube-solving contest, he might encounter some rather stiff competition from a few prodigious mortals, even if they do not know his algorithm. I am told there is a young Englishman from Nottingham named Nicholas Hammond who has got his average solving time down to close to 30 seconds! Such a phenomenal performance calls for several skills. The first is a deep understanding of the cube. The second is an extremely polished set of operators. The third is to have the operators down so cold that you could do them in your sleep. The fourth is sheer speed at executing twisty hand motions. The fifth is having a well-oiled "racing cube": one that turns at the merest twitch of a finger, eagerly anticipating every operator before it is needed. In short, the racing cube is a cube that *wants* to win.

I have not heard of people naming

their racing cubes, although it is sure to come. It would seem, however, that there is an association between colorful names and major contributors to cubology. Apart from Singmaster and Thistlethwaite there is Dame Kathleen Ollerenshaw (recently Lord Mayor of Manchester), who has discovered many streamlined processes, has written an article on the Magic Cube and has the distinction of being the first to report an attack of cubist's thumb, a grave form of the disease mentioned at the beginning of this column. Then there is Oliver Pretzel, the discoverer of a delicious twisted 3-cycle and the creator of a lovely "pretty pattern" called the 6-*U* state, which can be reached from the start position by way of the word $L'R^2F'L'B'UBLFRURRLR_sF_sU_sR_s$.

Pretty patterns are of interest to many cube lovers, but I cannot do them justice here. I shall mention only a few of the best. My favorite is the Worm, whose "genotype," or turn sequence, is $RUF^2D'R_sF_sD'F'R'F^2RU^2FR^2F'R'U'F'U^2FR$. Then there is the Snake, a similar sinuous pattern that winds around the cube: $BR_sD'R^2DR_sB'R^2UB^2U'DR^2D'$. If you cut off the Snake's tail R^2D' and stick on instead $B^2R_aU^2R_a'B^2D'$, you will create a curious bi-ringed pattern. All of these are from pretty-pattern-meister Richard Walker.

A beautiful pattern is the Giant Meson, made from a giant quark (a $2 \times 2 \times 2$ corner subcube rotated 120 degrees) and a giant antiquark. To top it off you can use quarkscrews to twist a standard-size quark and an antiquark onto the corners of the giant quark and antiquark, like cherries on top of sundaes [see illustration on page 15]. I shall let you figure this one out. A good warm-up exercise is to figure out how to make the Pons Asinorum (Bridge of Asses) state [see same illustration], so called because, as one M.I.T. cubemeister remarked to me, "If you can't hack this one, forget about cubing." Then there are two kinds of cross, known to the M.I.T. cube-hacking community as the Christman cross and the Plummer cross [see same illustration]. The Christman cross involves three pairs of colors (*U-D*, *F-R* and *L-B* in the illustration); the Plummer cross



The Magic Domino in a scrambled state

involves two triples in the quark-anti-quark style.

I should like to leave the reader with a set of hints and some things to think about. A difficult challenge, good for cubists at all levels of cubistry, is for someone to do a handful of turns on a pristine cube, for him to give it to you in this mildly scrambled state and for you to try to get it back to the Start position by finding the exact inverse word. Cube-meisters will be able to invert a bigger handful of turns than novices. Apparently Kate Fried can invert seven turns regularly, and once after a full day of staring at the cube she undid 10. (I can undo about four.)

My royal road to discovering an algorithm is based on two challenging exercises involving corner cubies only. The preliminary exercise is as follows. Maneuver the four corner cubies with white on them to the top face with their white facelets pointing upward. Do not worry about which cubie is in which cubicle. Simultaneously do the same thing on the bottom face (of course with *its* color pointing downward). The advanced exercise is to do the preceding one but in addition to make sure that all the corner cubies are in their proper cubicles. This amounts to solving the $2 \times 2 \times 2$ Magic Cube puzzle, and it will take you a long way toward mastery of the Magic Cube.

To help you with your edge processes, here is a wonderful trick discovered by David Seal, based on a type of operator called a monoflip. I shall give it to you as a puzzle. How can you make a double edge flipper out of a process that messes up the lower two layers but leaves the top layer invariant except for flipping a single edge cubie? I shall give you a hint: The answer involves the important group-theoretical idea of a commutator, a word of the form $PQP'Q'$. I shall also leave it to you to find your own monoflip operator. After I found out about it I incorporated the trick into my method.

Here is a small riddle: Why do 5- and 7-cycles crop up so often in an object whose symmetries all have to do with numbers such as 3, 4, 6 and 8? Where do cycle lengths such as 5 and 7 come from? A somewhat related question is: What is the maximum order a word can have? The order of a word is the power you have to raise it to in order to get the identity. (For example, the order of *R* is 4.) You can show that the order of *RU*, for instance, is 105 by inspecting its cycle structure.

Where do we go from here? I must mention that I have only scratched the surface of cubology in this column. Rubik and others are working on generalizations of various types. There already is a Magic Domino, which is like two-thirds of a magic cube: two 3×3 layers. You can rotate it only by half turns about two of its three axes and by *q* turns about the third axis. In its Start posi-

tion one face is black, the other is white and both faces have numbers from 1 through 9 in order. It resembles the 15 Puzzle even more strongly than the cube does. Various people have made $2 \times 2 \times 2$ cubes, and such cubes may go on sale one day. You can make your own by gluing little three-cornered hats over each of the eight corners of a $3 \times 3 \times 3$ cube. Readers will naturally wonder about such enticing possibilities as a $4 \times 4 \times 4$ cube. Rest assured, it is being developed in the Netherlands, and it may be ready soon. Inevitably there is the question of both higher and lower dimensionalities. Cube theorists are beginning to discuss the properties of higher-dimensional cubes.

The potential of the $3 \times 3 \times 3$ cube is not close to being exhausted. One rich area of unexplored terrain is that of alternate colorings. This idea was mentioned to me by various M.I.T. cube hackers. One can color the cubies in a variety of ways. Each new coloring presents a different kind of unscrambling problem. In one variant coloring, edge-cubie orientations become irrelevant and center-cubie orientations take on a vital importance. In another variant, corner-cubie orientations are irrelevant and centers matter. Then, moving toward simplicity, one can color two faces the same, thereby reducing the number of distinct colors by one. Or one can paint the faces with just three colors. An extreme would be to have three blue faces meet at one corner and three white ones meet at the corner diagonally opposite. The French government official I mentioned above says that on the cubes he saw, five faces had one color and the sixth face had another color!

Who knows where it will all end? As Bernie Greenberg has pointed out: "Cubism requires the would-be cubist to literally invent a science. Each solver must suggest areas of research to himself or herself, design experiments, find principles, build theories, reject them and so forth. It is the only puzzle that requires its solver to build a whole science." Could Rubik and Ishige have dreamed that their invention would lead to a model and a metaphor for all that is profound and beautiful in science? It is an amazing thing, this Magic Cube.

I should like to thank the many cubemeisters and cubists who have contributed directly or indirectly to my knowledge and enthusiasm. Among them are Clark Baker, Alan Bawden, Jim Boyce, Larry Breed, Charles Brenner, Bob Filman, Carl Hoffman, Scott Kim, Bill McKeeman, Jeannine Moseley, Richard Pavelle, Dave Plummer, David Policansky, David Singmaster, Ann Trail, Allen Wechsler, Dan Weise and John Woodcock. My warmest thanks, however, are reserved for Greenberg, who set up a cube-in for my benefit at M.I.T. and who infected me with his enthusiasm for the beauties of cubology.