

Fractal Music, Hypercards and More ...

Mathematical Recreations
from SCIENTIFIC AMERICAN
Magazine

Martin Gardner



Library of Congress Cataloging-in-Publication Data

Gardner, Martin, 1914-

Fractal music, hypercards and more : mathematical recreations from
Scientific American / by Martin Gardner.

p. cm.

Includes index.

ISBN 0-7167-2188-0. —ISBN 0-7167-2189-9(pbk.)

1. Mathematical recreations. I. Scientific American. II. Title.

QA95.G26 1991

793.8—dc20

91-17066

CIP

Copyright © 1992 by W. H. Freeman and Company

No part of this book may be reproduced by any mechanical, photographic or
electronic process, or in the form of a phonographic recording, nor may it be
stored in a retrieval system, transmitted or otherwise copied for public or
private use, without written permission from the publisher.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 VB 9 9 8 7 6 5 4 3 2 1

White, Brown, and Fractal Music

"For when there are no words [accompanying music] it is very difficult to recognize the meaning of the harmony and rhythm, or to see that any worthy object is imitated by them."

—PLATO, *Laws*, Book II

Plato and Aristotle agreed that in some fashion all the fine arts, including music, "imitate" nature, and from their day until the late 18th century "imitation" was a central concept in western aesthetics. It is obvious how representational painting and sculpture "represent," and how fiction and the stage copy life, but in what sense does music imitate?

By the mid-18th century philosophers and critics were still arguing over exactly how the arts imitate and whether the term is relevant to music. The rhythms of music may be said to imitate such natural rhythms as heartbeats, walking, running, flapping wings, waving fins, water waves, the periodic motions of heavenly bodies and so on, but this does not explain why we enjoy music more than, say, the sound of ci-

cadence or the ticking of clocks. Musical pleasure derives mainly from tone patterns, and nature, though noisy, is singularly devoid of tones. Occasionally wind blows over some object to produce a tone, cats howl, birds warble, bowstrings twang. A Greek legend tells how Hermes invented the lyre: he found a turtle shell with tendons attached to it that produced musical tones when they were plucked.

Above all, human beings sing. Musical instruments may be said to imitate song, but what does singing imitate? A sad, happy, angry or serene song somehow resembles sadness, joy, anger or serenity, but if a melody has no words and invokes no special mood, what does it copy? It is easy to understand Plato's mystification.

There is one exception: the kind of imitation that plays a role in "program music." A lyre is severely limited in the natural sounds it can copy, but such limitations do not apply to symphonic or electronic music. Program music has no difficulty featuring the sounds of thunder, wind, rain, fire, ocean waves and brook murmurings; bird calls (cuckoos and crowing cocks have been particularly popular), frog croaks, the gaits of animals (the thundering hoofbeats in Wagner's *Ride of the Valkyries*), the flights of bumblebees; the rolling of trains, the clang of hammers; the battle sounds of marching soldiers, clashing armies, roaring cannons and exploding bombs. *Slaughter on Tenth Avenue* includes a pistol shot and the wail of a police-car siren. In Bach's *Saint Matthew Passion* we hear the earthquake and the ripping of the temple veil. In the *Alpine Symphony* by Richard Strauss, cowbells are imitated by the shaking of cowbells. Strauss insisted he could tell that a certain female character in Felix Mottl's *Don Juan* had red hair, and he once said that someday music would be able to distinguish the clattering of spoons from that of forks.

Such imitative noises are surely a trivial aspect of music even when it accompanies opera, ballet or the cinema; besides, such sounds play no role whatsoever in "absolute music," music not intended to "mean" anything. A Platonist might argue that abstract music imitates emotions, or beauty, or the divine harmony of God or the gods, but on more mundane levels music is the least imitative of the arts. Even nonobjective paintings resemble certain patterns of nature, but nonobjective music resembles nothing except itself.

Since the turn of the century most critics have agreed that "imitation" has been given so many meanings (almost all

are found in Plato) that it has become a useless synonym for "resemblance." When it is made precise with reference to literature or the visual arts, its meaning is obvious and trivial. When it is applied to music, its meaning is too fuzzy to be helpful. In this chapter we take a look at a surprising discovery by Richard F. Voss, a physicist from Minnesota who joined the Thomas J. Watson Research Center of the International Business Machines Corporation after obtaining his Ph.D. at the University of California at Berkeley under the guidance of John Clarke. This work is not likely to restore "imitation" to the lexicon of musical criticism, but it does suggest a curious way in which good music may mirror a subtle statistical property of the world.

The key concepts behind Voss's discovery are what mathematicians and physicists call the spectral density (or power spectrum) of a fluctuating quantity, and its "autocorrelation." These deep notions are technical and hard to understand. Benoit Mandelbrot, who is also at the Watson Research Center, and whose work makes extensive use of spectral densities and autocorrelation functions, has suggested a way of avoiding them here. Let the tape of a sound be played faster or slower than normal. One expects the character of the sound to change considerably. A violin, for example, no longer sounds like a violin. There is a special class of sounds, however, that behave quite differently. If you play a recording of such a sound at a different speed, you have only to adjust the volume to make it sound exactly as before. Mandelbrot calls such sounds "scaling noises."

By far the simplest example of a scaling noise is what in electronics and information theory is called white noise (or "Johnson noise"). To be white is to be colorless. White noise is a colorless hiss that is just as dull whether you play it faster or slower. Its autocorrelation function, which measures how its fluctuations at any moment are related to previous fluctuations, is zero except at the origin, where of course it must be 1. The most commonly encountered white noise is the thermal noise produced by the random motions of electrons through an electrical resistance. It causes most of the static in a radio or amplifier and the "snow" on radar and television screens when there is no input.

With randomizers such as dice or spinners it is easy to generate white noise that can then be used for composing a random "white tune," one with no correlation between any

two notes. Our scale will be one octave of seven white keys on a piano: do, re, me, fa, so, la, ti. Fa is our middle frequency. Now construct a spinner such as the one shown at the left in Figure 1. Divide the circle into seven sectors and label them with the notes. It matters not at all what arc lengths are assigned to these sectors; they can be completely arbitrary. On the spinner shown, some order has been imposed by giving fa the longest arc (the highest probability of being chosen) and assigning decreasing probabilities to pairs of notes that are equal distances above and below fa. This has the effect of clustering the tones around fa.

To produce a "white melody" simply spin the spinner as often as you like, recording each chosen note. Since no tone is related in any way to the sequence of notes that precedes it, the result is a totally uncorrelated sequence. If you like, you can divide the circle into more parts and let the spinner select notes that range over the entire piano keyboard, black keys as well as white.

To make your white melody more sophisticated, use another spinner, its circle divided into four parts (with any proportions you like) and labeled 1, $1/2$, $1/4$ and $1/8$ so that you can assign a full, a half, a quarter or an eighth of a beat to each tone. After the composition is completed, tap it out on the piano. The music will sound just like what it is: random music of the dull kind that a two-year-old or a monkey might produce by hitting keys with one finger. Similar white music can be based on random number tables, or the digits in an irrational number.

A more complicated kind of scaling noise is one that is sometimes called Brownian noise because it is characteristic of Brownian motion, the random movements of small particles suspended in a liquid and buffeted by the thermal agitation of molecules. Each particle executes a three-dimensional "random walk," the positions in which form a highly correlated sequence. The particle, so to speak, always "remembers" where it has been.

When tones fluctuate in this fashion, let us follow Voss and call it Brownian music or brown music. We can produce it easily with a spinner and a circle divided into seven parts as before, but now we label the regions, as shown at the right in Figure 1, to represent intervals between successive tones. These step sizes and their probabilities can be whatever we like. On the spinner shown, plus means a step up the scale of

of seven white keys is our middle frequency—the one shown at the top of the spinner. The seven sectors and all what arc lengths are completely arbitrary. No probability has been imposed on the probability of being assigned to pairs of notes or to pairs of notes. This has the effect of

pin the spinner as a note. Since no tone precedes the tone that precedes the tone. If you like, and let the spinner on the keyboard, black

simplified, use an interval (with any probability 1/8 so that you get 1/8 of a beat to be used, tap it out on a keyboard, what it is: random or a monkey might play similar white music or the digits in an

noise is one that is not characteristic of small particles. Thermal agitation of a dimensional “randomly correlated members” where

let us follow Voss. We can produce a tune into seven parts shown at the right successive tones. It can be whatever we want to put up the scale of




FIGURE 1 Spinners for white music (left) and brown music (right)

one, two or three notes and minus means a step down of the same intervals.

Start the melody on the piano's middle C, then use the spinner to generate a linear random walk up and down the keyboard. The tune will wander here and there, and will eventually wander off the keyboard. If we treat the ends of the keyboard as “absorbing barriers,” the tune ends when we encounter one of them. We need not go into the ways in which we can treat the barriers as reflecting barriers, allowing the tune to bounce back, or as elastic barriers. To make the barriers elastic we must add rules so that the farther the tone gets from middle C, the greater is the likelihood it will step back toward C, like a marble wobbling from side to side as it rolls down a curved trough.

As before, we can make our brown music more sophisticated by varying the tone durations. If we like, we can do this in a brown way by using another spinner to give not the duration but the increase or decrease of the duration—another random walk but one along a different street. The result is a tune that sounds quite different from a white tune because it is strongly correlated, but a tune that still has little aesthetic appeal. It simply wanders up and down like a drunk weaving through an alley, never producing anything that resembles good music.

If we want to mediate between the extremes of white and brown, we can do it in two essentially different ways. The way chosen by previous composers of “stochastic music” is to adopt transition rules. These are rules that select each note

on the basis of the last three or four. For example, one can analyze Bach's music and determine how often a certain note follows, say, a certain triplet of preceding notes. The random selection of each note is then weighted with probabilities derived from a statistical analysis of all Bach quadruplets. If there are certain transitions that never appear in Bach's music, we add rejection rules to prevent the undesirable transitions. The result is stochastic music that resembles Bach but only superficially. It sounds Bachlike in the short run but random in the long run. Consider the melody over periods of four or five notes and the tones are strongly correlated. Compare a run of five notes with another five-note run later on and you are back to white noise. One run has no correlation with the other. Almost all stochastic music produced so far has been of this sort. It sounds musical if you listen to any small part but random and uninteresting when you try to grasp the pattern as a whole.

Voss's insight was to compromise between white and brown input by selecting a scaling noise exactly halfway between. In spectral terminology it is called $1/f$ noise. (White noise has a spectral density of $1/f^0$, brownian noise a spectral density of $1/f^2$. In "one-over- f " noise the exponent of f is 1 or very close to 1.) Tunes based on $1/f$ noise are moderately correlated, not just over short runs but throughout runs of any size. It turns out that almost every listener agrees that such music is much more pleasing than white or brown music.

In electronics $1/f$ noise is well known but poorly understood. It is sometimes called flicker noise. Mandelbrot, whose book *The Fractal Geometry of Nature* (W. H. Freeman and Company, 1982) has already become a modern classic, was the first to recognize how widespread $1/f$ noise is in nature, outside of physics, and how often one encounters other scaling fluctuations. For example, he discovered that the record of the annual flood levels of the Nile is a $1/f$ fluctuation. He also investigated how the curves that graph such fluctuations are related to "fractals," a term he invented. A scaling fractal can be defined roughly as any geometrical pattern (other than Euclidean lines, planes and surfaces) with the remarkable property that no matter how closely you inspect it, it always looks the same, just as a slowed or speeded scaling noise always sounds the same. Mandelbrot coined the term fractal because he assigns to each of the curves a fractional dimension greater than its topological dimension.

ple, one can
 certain note
 The random
 abilities de-
 adruplets. If
 n Bach's mu-
 irable transi-
 les Bach but
 t run but ran-
 eriods of four
 ed. Compare
 er on and you
 tion with the
 far has been
 ny small part
 grasp the pat-

n white and
 y halfway be-
 noise. (White
 ise a spectral
 nt of f is 1 or
 oderately cor-
 it runs of any
 ees that such
 vn music.
 poorly under-
 elbrot, whose
 Freeman and
 a classic, was
 e is in nature,
 ers other scal-
 at the record
 uctuation. He
 h fluctuations
 scaling fractal
 rn (other than
 e remarkable
 et it, it always
 aling noise al-
 e term fractal
 ctional dimen-

Among the fractals that exhibit strong regularity the best-known are the Peano curves that completely fill a finite region and the beautiful snowflake curve discovered by the Swedish mathematician Helge von Koch in 1904. The Koch snowflake appears in Figure 2 as the boundary of the dark "sea" that surrounds the central motif. (For details on the snowflake's construction, and a discussion of fractals in general, see Chapter 3 of my *Penrose Tiles to Trapdoor Ciphers* (W. H. Freeman, 1989).

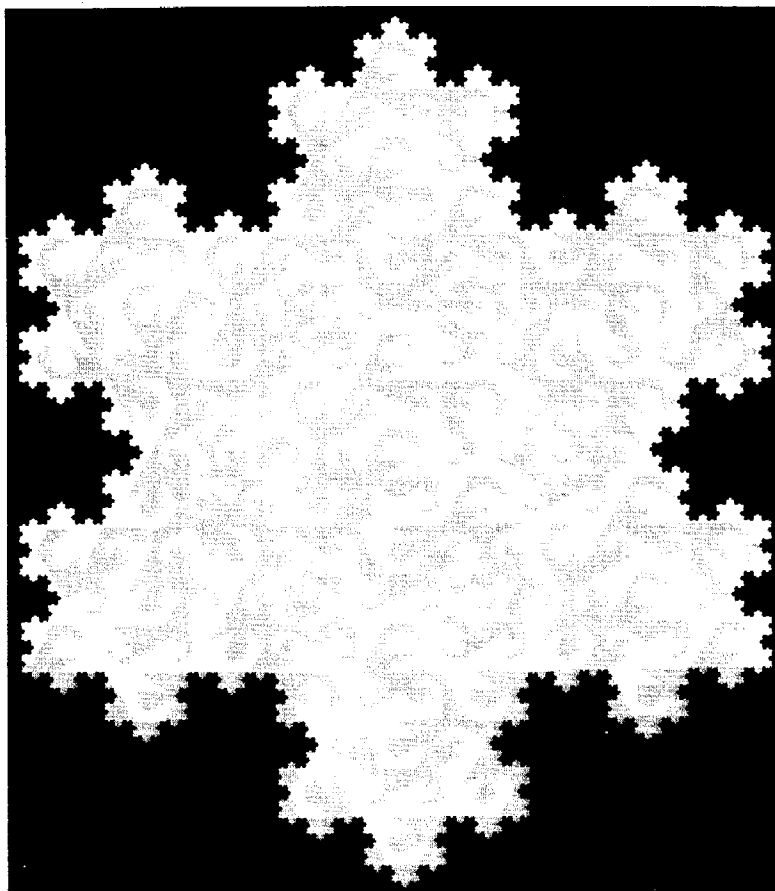


FIGURE 2 Mandelbrot's Peano-snowflake as it appeared on the cover of *Scientific American* (April, 1978). The curve was drawn by a program written by Sigmund Handelman and Mark Laff.

The most interesting part of Figure 2 is the fractal curve that forms the central design. It was discovered by Mandelbrot and published for the first time as the cover of *Scientific American's* April 1978 issue. If you trace the boundary between the black and white regions from the tip of the point of the star at the lower left to the tip of the point of the star at the lower right, you will find this boundary to be a single curve. It is the third stage in the construction of a new Peano curve. At the limit this lovely curve will completely fill a region bounded by the traditional snowflake! Thus Mandelbrot's curve brings together two pathbreaking fractals: the oldest of them all, Giuseppe Peano's 1890 curve, and Koch's later snowflake!

The secret of the curve's construction is the use of line segments of two unequal lengths and oriented in 12 different directions. The curve is much less regular than previous Peano curves and therefore closer to the modeling of natural phenomena, the central theme of Mandelbrot's book. Such natural forms as the gnarled branches of a tree or the shapes of flickering flames can be seen in the pattern.

At the left in Figure 3 is the first step of the construction. A crooked line of nine segments is drawn on and within an equilateral triangle. Four of the segments are then divided




FIGURE 3 The first two steps in constructing Benoit Mandelbrot's Peano-snowflake curve

into
of 13
of th
line.
as is
thirc
ure
Gos
Chaj
repe
a re
has
of in
fract


occu
clou
(sca
the
brot
ated
vert
mod
or si
to ei

fract
it w
see.
fore
cour
ture
from
that
a pa

can
rang
tuat
skyl
froze
can

the fractal curve
 ured by Mandel-
 ver of *Scientific*
 boundary be-
 of the point of
 of the star at
 a single curve.
 Peano curve.
 fill a region
 elbrot's curve
 best of them
 later snow-

use of line
 12 different
 vious Peano
 natural phe-
 Such natu-
 the shapes of
 construction.
 within an
 divided



into two equal parts, creating a line from A to B that consists of 13 long and short segments. The second step replaces each of these 13 segments with a smaller replica of the crooked line. These replicas (necessarily of unequal size) are oriented as is shown inside the star at the right in the illustration. A third repetition of the procedure generates the curve in Figure 2. (It belongs to a family of curves arising from William Gosper's discovery of the "flow-snake," a fractal pictured in Chapter 3 of my above cited book.) When the construction is repeated to infinity, the limit is a Peano curve that totally fills a region bordered by the Koch snowflake. The Peano curve has the usual dimension of 2, but its border, a scaling fractal of infinite length, has (as is explained in Mandelbrot's book) a fractal dimension of $\log 4/\log 3$, or 1.2618. . . .

Unlike these striking artificial curves, the fractals that occur in nature—coastlines, rivers, trees, star clustering, clouds and so on—are so irregular that their self-similarity (scaling) must be treated statistically. Consider the profile of the mountain range in Figure 4, reproduced from Mandelbrot's book. This is not a photograph, but a computer-generated mountain scene based on a modified Brownian noise. Any vertical cross section of the topography has a profile that models a random walk. The white patches, representing water or snow in the hollows below a certain altitude, were added to enhance the relief.

The profile at the top of the mountain range is a scaling fractal. This means that if you enlarge any small portion of it, it will have the same statistical character as the line you now see. If it were a true fractal, this property would continue forever as smaller and smaller segments are enlarged, but of course such a curve can neither be drawn nor appear in nature. A coastline, for example, may be self-similar when viewed from a height of several miles down to several feet, but below that the fractal property is lost. Even the Brownian motion of a particle is limited by the size of its microsteps.

Since mountain ranges approximate random walks, one can create "mountain music" by photographing a mountain range and translating its fluctuating heights to tones that fluctuate in time. Villa Lobos actually did this using mountain skylines around Rio de Janeiro. If we view nature statically, frozen in time, we can find thousands of natural curves that can be used in this way to produce stochastic music. Such




FIGURE 4 A modified Brownian landscape generated by a computer program

music is usually too brown, too correlated, however, to be interesting. Like natural white noise, natural brown noise may do well enough, perhaps, for the patterns of abstract art but not so well as a basis for music.


When we analyze the dynamic world, made up of quantities constantly changing in time, we find a wealth of fractal-like fluctuations that have $1/f$ spectral densities. In his book Mandelbrot cites a few: variations in sunspots, the wobbling of the earth's axis, undersea currents, membrane currents in the nervous system of animals, the fluctuating levels of rivers and so on. Uncertainties in time measured by an atomic clock are $1/f$: the error is 10^{-12} regardless of whether one is measuring an error on a second, minute or hour. Scientists tend to overlook $1/f$ noises because there are no good theories to account for them, but there is scarcely an aspect of nature in which they cannot be found.

og
ar
tli
a
th
in
a
su
re

in
pu
en
en
co
br
its
ne
de
ho
Th
vic
the

ref
br
ob
nat
a s
anc
cor
ph
tur
are

of
but
tha
ove
the
f s
qua



e generated by a com-

ated, however, to be natural brown noise may ns of abstract art but

ld, made up of quan- id a wealth of fractal- lensities. In his book inspots, the wobbling membrane currents in uating levels of rivers ed by an atomic clock whether one is mea- hour. Scientists tend e no good theories to an aspect of nature in

T. Musha, a physicist at the Tokyo Institute of Technology, discovered that traffic flow past a certain spot on a Japanese expressway exhibited $1/f$ fluctuation. In a more startling experiment, Musha rotated a radar beam emanating from a coastal location to get a maximum variety of landscape on the radar screen. When he rotated the beam once, variations in the distances of all objects scanned by the beam produced a Brownian spectrum. But when he rotated it twice and then subtracted one curve from the other the resulting curve—representing all the changes of the scene—was close to $1/f$.

We are now approaching an understanding of Voss's daring conjecture. The changing landscape of the world (or, to put it another way, the changing content of our total experience) seems to cluster around $1/f$ noise. It is certainly not entirely uncorrelated, like white noise, nor is it as strongly correlated as brown noise. From the cradle to the grave our brain is processing the fluctuating data that comes to it from its sensors. If we measure this noise at the peripheries of the nervous system (under the skin of the fingers), it tends, Mandelbrot says, to be white. The closer one gets to the brain, however, the closer the electrical fluctuations approach $1/f$. The nervous system seems to act like a complex filtering device, screening out irrelevant elements and processing only the patterns of change that are useful for intelligent behavior.

On the canvas of a painting, colors and shapes are static, reflecting the world's static patterns. Is it possible, Mandelbrot asked himself many years ago, that even completely non-objective art, when it is pleasing, reflects fractal patterns of nature? He is fond of abstract art, and maintains that there is a sharp distinction between such art that has a fractal base and such art that does not, and that the former type is widely considered the more beautiful. Perhaps this is why photographers with a keen sense of aesthetics find it easy to take pictures, particularly photomicrographs, of natural patterns that are almost indistinguishable from abstract expressionist art.

Motion can be added to visual art, of course, in the form of the motion picture, the stage, kinetic art and the dance, but in music we have meaningless, nonrepresentational tones that fluctuate to create a pattern that can be appreciated only over a period of time. Is it possible, Voss asked himself, that the pleasures of music are partly related to scaling noise of $1/f$ spectral density? That is, is this music "imitating" the $1/f$ quality of our flickering experience?

That may or may not be true, but there is no doubt that music of almost every variety does exhibit $1/f$ fluctuations in its changes of pitch as well as in the changing loudness of its tones. Voss found this to be true of classical music, jazz and rock. He suspects it is true of all music. He was therefore not surprised that when he used a $1/f$ flicker noise from a transistor to generate a random tune, it turned out to be more pleasing than tunes based on white and brown noise sources.

Figure 5, supplied by Voss, shows typical patterns of white, $1/f$ and brown when noise values (vertical) are plotted against time (horizontal). These patterns were obtained by a computer program that simulates the generation of the three kinds of sequences by tossing dice. The white noise is based on the sum obtained by repeated tosses of 10 dice. These sums range from 10 to 60, but the probabilities naturally force a clustering around the median. The Brownian noise was generated by tossing a single die and going up one step on the scale if the number was even and down a step if the number was odd.

The $1/f$ noise was also generated by simulating the tossing of 10 dice. Although $1/f$ noise is extremely common in nature, it was assumed until recently that it is unusually cumbersome to simulate $1/f$ noise by randomizers or computers. Previous composers of stochastic music probably did not even know about $1/f$ noise, but if they did, they would have had considerable difficulty generating it. As this article was being prepared Voss was asked if he could devise a simple procedure by which readers could produce their own $1/f$ tunes. He gave some thought to the problem and to his surprise hit on a clever way of simplifying existing $1/f$ computer algorithms that does the trick beautifully.

The method is best explained by considering a sequence of eight notes chosen from a scale of 16 tones. We use three dice of three colors: red, green and blue. Their possible sums range from 3 to 18. Select 16 adjacent notes on a piano, black keys as well as white if you like, and number them 3 through 18.

Write down the first eight numbers, 0 through 7, in binary notation, and assign a die color to each column as is shown in Figure 6. The first note of our tune is obtained by tossing all three dice and picking the tone that corresponds to the

50

40

30

20

50

40

30

20

60

40

20

0

0

FI

here is no doubt that
 ibit $1/f$ fluctuations in
 anging loudness of its
 ssical music, jazz and

He was therefore not
 r noise from a transis-
 out to be more pleas-
 n noise sources.

s typical patterns of
 (vertical) are plotted
 is were obtained by a
 eneration of the three
 e white noise is based
 f 10 dice. These sums
 ties naturally force a
 wnian noise was gen-
 g up one step on the
 a step if the number

y simulating the toss-
 remely common in nat-
 t it is unusually cum-
 mizers or computers.
 probably did not even
 they would have had
 this article was being
 evised a simple proce-
 neur own $1/f$ tunes. He
 o his surprise hit on a
 nputer algorithms that

onsidering a sequence
 5 tones. We use three
 e. Their possible sums
 otes on a piano, black
 mber them 3 through

0 through 7, in binary
 h column as is shown
 is obtained by tossing
 at corresponds to the




FIGURE 5 Typical patterns of white, $1/f$ and Brownian noise

	Blue ↓	Green ↓	Red ↓
0	0	0	0
1	0	0	1
2	0	1	0
3	0	1	1
4	1	0	0
5	1	0	1
6	1	1	0
7	1	1	1

FIGURE 6 Binary chart for Voss's $1/f$ dice algorithm

sum. Note that in going from 000 to 001 only the red digit changes. Leave the green and blue dice undisturbed, still showing the numbers of the previous toss. Pick up only the red die and toss it. The new sum of all three dice gives the second note of your tune. In the next transition, from 001 to 010, both the red and green digits change. Pick up the red and green dice, leaving the blue one undisturbed, and toss the pair. The sum of all three dice gives the third tone. The fourth note is found by shaking only the red die, the fifth by shaking all three. The procedure, in short, is to shake only those dice that correspond to digit changes.

It is not hard to see how this algorithm produces a sequence halfway between white and brown. The least significant digits, those to the right, change often. The more significant digits, those to the left, are more stable. As a result, dice corresponding to them make a constant contribution to the sum over long periods of time. The resulting sequence is not precisely $1/f$ but is so close to it that it is impossible to distinguish melodies formed in this way from tunes generated by natural $1/f$ noise. Four dice can be used the same way for a $1/f$ sequence of 16 notes chosen from a scale of 21 tones. With 10 dice you can generate a melody of 2^{10} , or 1,024, notes from a scale of 55 tones. Similar algorithms can of course be

impleme
drons an

With
three ty
The com
7, 8, and
tone dur
are show

Ove
played a
many hu
sic too r
right." I
how the
Voss's e
ally elec
rives fro
He has
he appli
varied th
bled Ori
by addin
stochast
the unde
white no

Note
in a frac
sition ru
sic rever
But $1/f$ r
mountain
note to r
tune into
forgets
with its

It is
joy good
prise. H
prise if
what is
listening
the keyf
Good m

implemented with generalized dice (octahedrons, dodecahedrons and so on), spinners or even tossed coins.

With the same dice simulation program Voss has supplied three typical melodies based on white, brown, and $1/f$ noise. The computer printouts of the melodies are shown in Figures 7, 8, and 9. In each case Voss varied both the melody and the tone duration with the same kind of noise. Above each tune are shown the noise patterns that were used.

Over a period of two years, tunes of the three kinds were played at various universities and research laboratories, for many hundreds of people. Most listeners found the white music too random, the brown too correlated and the $1/f$ "just about right." Indeed, it takes only a glance at the music itself to see how the $1/f$ property mediates between the two extremes. Voss's earlier $1/f$ music was based on natural $1/f$ noise, usually electronic, even though one of his best compositions derives from the record of the annual flood levels of the Nile. He has made no attempt to impose constant rhythms. When he applied $1/f$ noise to a pentatonic (five-tone) scale and also varied the rhythm with $1/f$ noise, the music strongly resembled Oriental music. He has not tried to improve his $1/f$ music by adding transition or rejection rules. It is his belief that stochastic music with such rules will be greatly improved if the underlying choices are based on $1/f$ noise rather than the white noise so far used.

Note that $1/f$ music is halfway between white and brown in a fractal sense, not in the manner of music that has transition rules added to white music. As we have seen, such music reverts to white when we compare widely separated parts. But $1/f$ music has the fractal self-similarity of a coastline or a mountain range. Analyze the fluctuations on a small scale, from note to note, and it is $1/f$. The same is true if you break a long tune into 10-note sections and compare them. The tune never forgets where it has been. There is always some correlation with its entire past.

It is commonplace in musical criticism to say that we enjoy good music because it offers a mixture of order and surprise. How could it be otherwise? Surprise would not be surprise if there were not sufficient order for us to anticipate what is likely to come next. If we guess too accurately, say in listening to a tune that is no more than walking up and down the keyboard in one-step intervals, there is no surprise at all. Good music, like a person's life or the pageant of history, is a

$1/f$ dice algorithm

001 only the red digit
dice undisturbed, still
toss. Pick up only the
ll three dice gives the
transition, from 001 to
ange. Pick up the red
undisturbed, and toss
es the third tone. The
he red die, the fifth by
short, is to shake only
nges.
orithm produces a se-
own. The least signifi-
often. The more signif-
re stable. As a result,
onstant contribution to
e resulting sequence is
that it is impossible to
y from tunes generated
used the same way for
m a scale of 21 tones.
y of 2^{10} , or 1,024, notes
ithms can of course be






FIGURE 7 White music





FIGURE 8 Brown music

