HOW MUSIC WORKS

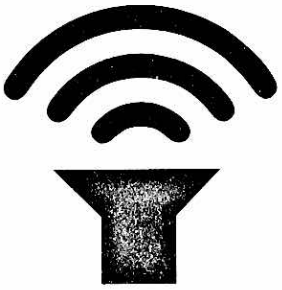

DAVID BYRNE

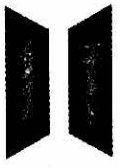

CANONGATE
Edinburgh · London

This paperback edition published by Canongate Books in 2013

First published in Great Britain in 2012 by Canongate Books Ltd, 14 High Street, Edinburgh EH1 1TE

www.canongate.tv

6

Cover design by Dave Eggers

First published in the USA in 2012 by McSweeney's, San Francisco

Some of these chapters first came to life in other forms. 'Creation in Reverse' began as a TED talk; 'How to Make a Scene' began as an introduction to a picture book about CBGB; 'Business and Finances' includes material from an article that originally ran in *Wired* magazine. Those pieces have been extensively revised and expanded. Thank you to the editors who first motivated me to write them. – DB

ISBN: 978-0-85786-252-5

Printed in China on acid-free paper by C&C Offset Printing Co., Ltd.

CONTENTS

CHAPTER FOUR

Technology Shapes Music

Part Two: Digital

I heard computer scientist Jaron Lanier speak at a symposium recently. After playing some pieces on a *shen*, a Chinese mouth organ, he said that it had a surprising and prodigious heritage. He claimed that this instrument was maybe the first in which the notes to be played were chosen by a mechanism, a mechanism that was a precursor to binary theory and thus to all computers.

That ancient bit of gear found its way to Rome via the Silk Road, and the Empire had a giant version built—as Empires are wont to do. This larger instrument required an assistant to pump the air—it was too big to play by mouth anymore—and, more significantly, a series of levers that selected the notes. This system was the inspiration for what we now know as the keyboard—the series of levers that are used to play the notes of the organ (which is also a large wind instrument) and piano. It was also inspirational to the Frenchman Joseph Marie Jacquard, who in 1801 made a weaving loom in which complicated patterns were guided by punch cards. One could control the design of a fabric by stringing the cards together.

Decades later, Jacquard's loom was inspirational to Charles Babbage, who owned one of Jacquard's self-portraits, in which he used these cards to make an image of himself woven in silk. Babbage designed his Analytical Engine—a computational

machine that, were it ever built, would also have been controlled by punch cards. In Babbage's version, the cards no longer controlled threads but made the leap to binary abstraction—making pure calculations. Babbage's young friend Ada Byron (daughter of the poet) was fascinated by the device, and many years later became celebrated as the first computer programmer. So, according to Lanier, our present computer-saturated world owes something of its lineage to a musical instrument. And computer technology, not too long after it came into being, affected music as well.

The technology that allowed sound information (and, soon thereafter, all other information) to be digitized was largely developed by the phone company. Bell Labs, the research division of the Bell Telephone Company, had a mandate to find more efficient and reliable ways of transmitting conversations. Prior to the sixties, all phone lines were analog and the number of conversations that could be handled at one time was limited. The only way to squeeze more calls through the lines was to roll the high and low frequencies off the sound of the voice, and then turn the resulting lo-fi sound into waves that could run in parallel without interfering with one another—much like what happens with terrestrial radio transmissions.

Bell Labs was huge and gave birth to a slew of new inventions. The transistor and semiconductors that form silicon-based integrated circuits (making today's tiny devices possible), the laser, microwave technology, solar panels—the list goes on and on. When you're a monopoly you can afford to spend on R&D, and they had the luxury and foresight to take the long view. Scientists and engineers could work on a project that might not show results for ten years.

In 1962, Bell Labs figured out how to digitize sound—to, in effect, sample a sound wave and slice it into tiny bits that could be broken down into ones and zeros. When they could do this in a way that was not prohibitively expensive and that still left the human voice recognizable, they immediately applied that technology to making their long-distance lines more efficient. More calls could now be made simultaneously, as the voice was now just a stream of ones and zeros that they could squeeze (via encoding and transposing), along with other calls, into their telephone cables. This was especially relevant considering the limitations imposed by long-distance underwater phone cables; you couldn't just go out and lay more lines down if suddenly it seemed more people wanted to talk to France.

A voice is, in the abstract sense, a kind of information, viewed from Bell's perspective. Therefore, much of their research regarding what made a transmission understandable, or how you could squeeze more transmissions in, involved applying the science of information in combination with insights gleaned from the science of psychoacoustics—the study of how the brain perceives sound in all its aspects. So understanding how we perceive sound became integrated with the quest for how to most efficiently transmit information of all kinds. It was even relevant to the meta question "What is information, anyway?"

Psychoacoustics has applications to the sound of ambulances (Why can we never tell where they're coming from?), the speaking voice, and, of course, music. The *psycho*-prefix is there because what and how we hear is not simply mechanical, it's mental (meaning the brain "hears" as much as the ear does—not mental as in insane, or insanely great).

Of course, much of what we hear is partly defined and limited by the mechanics of our ears. We know that we can't hear all the high-pitched sounds that bats emit or the full range of sounds that a dog can hear. There are low-pitched sounds that whales produce that we can't really hear either, though they are strong enough to do us physical harm if we are too close to the source.

But there are things we "hear" that have nothing to do with the physics of the eardrum and the auditory canal. We can, for example, isolate the voice of someone talking to us in a noisy environment. If you were to listen to a recording of a noisy restaurant it would sound like acoustic chaos, but sometimes we manage to make order out of it and carry on a rudimentary conversation. Repetitious sounds, the sound of waves or constant traffic, become somewhat inaudible to us after a while. We have the ability to selectively hear just the stuff we're interested in and make the rest recede into some distant acoustic background. We also have the ability to perceive patterns in sounds. This too has nothing to do with our ears. We can remember pitches, and some people with perfect pitch can accurately determine notes heard out of a musical context. We can tell if the sound of squealing subway brakes and the highest note on a clarinet are the same. We can remember sequences of sounds—a bird's song or a door creak followed by a slam—and the exact timbre of sounds—we sometimes recognize a friend's voice by hearing a single word.

How does this work? Can we simulate that mental process with a mathematical formula or a computer program? As you can imagine, such questions—how little information do we need to recognize someone's voice, for example—were of prime importance to a phone company. If they could understand what exactly makes speech understandable and intelligible, and isolate just that aspect—refine it, control it—then they might increase the efficiency of their phone system by eliminating all the superfluous parts of the transmissions. The goal was to eventually communicate more and more using less, or the same amount, of the mechanical and physical electrical stuff. This possible increase in information flow would make them a lot more money. Psychoacoustics would eventually lead to an increased understanding of information transmission. This arcane science was suddenly hugely useful.

An unforeseen consequence of this phone-related research was the emergence of digital-based audio technology that was eventually used in, among other places, recording studios. In the seventies a new piece of equipment the size of a briefcase appeared in recording studios. It was called a harmonizer, and it could change the pitch of a sound without changing its speed or tempo, as would happen if you changed the pitch by speeding up a tape. It achieved this by slicing up the sound waves into digital slivers, mathematically transposing what were now merely numbers and then reconstructing those as sounds at a higher or lower pitch. The early versions of this machine sounded pretty glitchy, but the effect was cool, even when it didn't work.

Around that same time there appeared other devices called digital delays, which were in effect primitive samplers. The digital samples they created in order to mimic acoustic echoes were usually much less than a second long and they would be used to produce very short delay effects.

More devices followed: machines that could grab and hold longer sound samples with greater resolution, and some that could manipulate those "sounds" (they were really just numbers) more freely. All sorts of weirdness resulted. Bell Labs was involved in manufacturing a sound processor called a vocoder that could isolate certain aspects of talking (or singing) like speech formants (the shape of the sounds that we use to form words). This device could remove these aspects of our talking or singing from the pitch—like isolating just the percussive parts—the T's

and B's and the sibilants of S's and F's. This machine could transmit these formant aspects of a voice separate from the rest of a vocalization, and the resulting gibberish, when transmitted, was more or less unintelligible. But the components of intelligent speech were still there. The elements of the sound of speech or singing had been deconstructed and could make sense again when put back together. Wonderful, but what do you do with this technology? One use for it was a sort of cryptology for the voice: the garbled nonsense could be "decoded" at the other end if you knew what had been taken out and where. These machines were also adopted for music production. Below, the German band Kraftwerk's vocoder, made especially for them.[A]

A vocoder was typically used to apply those isolated and separated speech formants to the sound of a pitched instrument. The instrument then appeared to be talking or singing. Often the resulting "voice" was somewhat robotic sounding, an aspect that likely appealed to Kraftwerk. I once used a vocoder like this that I borrowed from Bernie Krause, a musician and early synthesizer pioneer, whom I met when Brian Eno and I did the *Bush of Ghosts* record. The vocoder was beautifully made, but rather complicated, and very expensive.

An early harmonizer (that digital pitch shifter) cost thousands of dollars. A good digital reverberation unit set a studio back maybe ten thousand dollars, and a full-fledged digital sampling device, like the Fairlight or a Synclavier that emerged soon after, cost much, much more. But soon the price of memory and processing dropped, and the technology became more affordable. Inexpensive Akai samplers became the backbone of hip-hop and DJ mixes, replacing the earlier use of vinyl, and sampled or digitally derived drum sounds took the place of live drummers in many recordings. We were off to the races, for better or worse.

With the digitization of sound, digital recording and consumer products like the CD became possible, and entire record albums were soon sliced into these tiny slivers of ones and zeros. Not long after that, the capacity and speed of home computers became sufficient to allow individuals to record, archive, and process music. All of this follows from Bell Labs's desire to improve the efficiency of their phone lines.

A

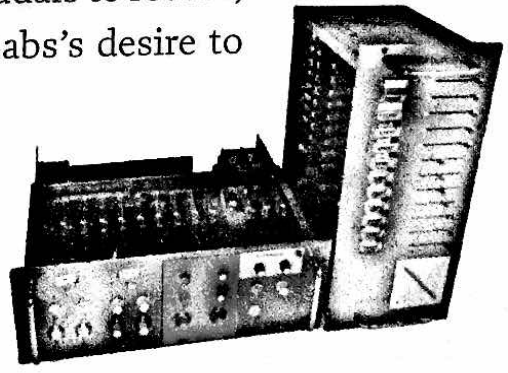

Bell Labs eventually became Lucent. I visited their labs in the mid-nineties and they showed me a

processor that could squeeze what sounded to the ear to be CD-quality music into a minuscule bandwidth. I believe encoding music as MP3s had already been invented in Germany by that time, so this extremely efficient compressing/encoding trick was not a complete surprise. And it was certainly no surprise that squeezing more sound information into smaller spaces continued to be a priority for a subsidy of a phone company. But like many people, I worried that the quality of music might somehow get sacrificed in this "rezzing down" process.

I was right. Those early, low-bandwidth digital files sounded slightly off, as if something ineffable was missing. It was hard to put your finger on why they sounded wrong, but they did. All the frequencies seemed to be there, but something seemed to have been sucked out in the process. Zombie music. MP3s have improved quite a bit since then, and now I listen to most of the music I own in that format. I believe what Lucent was working on ended up being used for satellite-radio transmission: getting "CD-quality" sound into smaller bandwidth transmissions, so that a satellite could send out lots of channels of sound that seem to be of high quality. Similar processing would be applied to photographs and video signals, which allows us to stream movies without them looking completely grainy or pixelated.

In 1988, I got an advance peek at this technology as it was applied to visual information when the designer Tibor Kalman and I visited a printing studio on Long Island. The studio had a machine that could digitize and then subtly manipulate images (we wanted to "improve" the image that was to be used on a Talking Heads record cover). Like the early computers and recording-studio gear, this machine was incredibly expensive and rare. We had to go to it (it couldn't be brought to the design studio), and we had to book time in advance. A Sytex machine, I think it was called. Impressed as we were, its cost and rarity meant we didn't think much about incorporating its talents into future projects.

After a while, as with sampling, the price of scanning images dropped, and manipulating images using Photoshop became common. There are some film holdouts, and I have no doubt that, as with MP3s, something has been lost with digital images, but, well, for most of us the trade-off seems acceptable—and inevitable. Needless to say, as images become digitized, they enter the river of networked data. Images for us are increasingly sequences of ones and

zeros—information, like everything else. The digitization of every form of media enabled the Web to be what it is, much more than a way of transmitting text-based documents. This slicing of content allowed a wide variety of media to flow into that river, and in a way we owe all the pictures, sounds, songs, games, and movies that are part of our Internet experience to the phone company, information science, and psychoacoustics.

CDS

CDs, which made their debut in 1982, were jointly developed by Sony in Japan and Philips in Holland. Previously, digitized movies had been stored on LaserDiscs, which were the size of LPs, and the prospect of encoding an entire record album's worth of sound therefore seemed within reach. If the discs could be made smaller, it could be lucrative. Philips had the laser aspect in development and Sony had the manufacturing prowess, so they agreed to work on this new format together. The arrangement was unusual; usually one company developed a format on its own and then tried to exert control over it so it could start charging others for using it. As a result, a lot of proprietary nonsense that could have burdened the acceptance and dissemination of CDs was avoided.

It was rumored that the length of the CD was determined by the duration of Beethoven's Ninth Symphony, because that was Norio Ohga's favorite piece of music, and he was the president of Sony at that time. Philips had designed a CD with an 11.5 cm diameter, but Ohga insisted that a disc must be able to hold the entire Beethoven recording. The longest recording of the symphony in Polygram's archive was seventy-four minutes, so the CD size was increased to a 12 cm diameter to accommodate the extra data.

Unlike LPs, whose grooves and bouncing needles limited the volume, the low frequencies were practically unlimited with the super-high-end CD technology. The music was no longer mirrored by physical grooves but was now encoded in a series of digital ones and zeroes. Though these discs spun around like LPs, they were technically nothing like the old records. Their extended audio range resulted from the fact that since there was no physical analogue of the sound, the coded

messages "told" the CD player what frequencies to play. The ones and zeroes could tell the stereo system to play anything audible to the human ear, at whatever frequency or volume was desired. This sonic range in digital music was really only limited by the playback and sampling mechanisms that allowed sounds outside the human hearing range to be recorded. The expanded and unlimited sound range was now, or would soon be, available to everyone.

Inevitably, this sonic freedom got abused quite a bit. Some records (the writer Greg Milner mentions most Oasis albums and *Californication* by the Red Hot Chili Peppers) were made so artificially loud that though the music seemed amazing on first listen (it was louder, and more consistently louder, than anything else), it rapidly wore on the ears. Milner claims that this "volume war" was spurred by radio DJs and technicians who wanted their stations to seem louder than the stations near them on the radio dial.[1] To achieve this, inventor Mike Dorrough developed a device in the sixties called a "discriminate audio processor," which caught on widely years later when every station was trying to be louder than every other. Milner speculates that musicians and record producers responded to this competition by figuring out how to make their records sound louder, and stay louder, for the whole length of the record.[2] Pretty soon there was ear fatigue all around. The listener never got a break; there was no dynamic range anymore. Milner suggests that even rabid music fans can't listen to these records over and over, or very much at all. The actual enjoyment of them is short-lived, and he proposes that this might have had something to do with pushing consumers away from purchasing recorded music. The technology that was supposed to make music more popular than ever instead made everyone run away from it.

CRAPPY SOUND FOREVER

Early CDs, like the MP3s that followed, didn't sound all that great. Dr. John Diamond treated psychotic patients with music, but by 1989 he sensed that it had all gone wrong. He claims that the natural healing and therapeutic properties of music were lost in the rush to digitize.[3] He believes that certain pieces of music can help soothe and heal, if they are the entirely analog versions, while the digital

versions actually have the reverse effect. When his test subjects are played digital recordings, they get agitated and twitchy.

Throughout the history of recorded music, we have tended to value convenience over quality every time. Edison cylinders didn't really sound as good as live performers, but you could carry them around and play them whenever you wanted. LPs, revolving slower, didn't sound as rich as 45s or 78s, but you didn't have to attend to them as much. And cassettes? Are you kidding? We were told that CDs would last forever and sound squeaky clean, but they really don't sound as good as LPs, and the jury is out regarding their durability. The spectrum of sound on analog mediums has an infinite number of gradations, whereas in the digital world everything is sliced into a finite number of slivers. Slivers and bits might fool the ear into believing that they represent a continuous audio spectrum (psychoacoustics at work), but by nature they are still ones and zeros; steps rather than a smooth slope. MP3s? They may be the most convenient medium so far, but I can't help thinking that the psychoacoustic trickery used to develop them—the ability to cause the mind to think and feel that all the musical information is there when in reality a huge percentage has been removed—is a continuation of this trend in which we are seduced by convenience. It's music in pill form, it delivers vitamins, it does the job, but something is missing. We are often offered, and gladly accept, convenient mediums that are "good enough" rather than ones that are actually better.

Where does this road of compromise end, and does it really matter if we lose a little quality along the way? Isn't the quality or accuracy of a recording somewhat irrelevant to music's use and enjoyment? We laugh out loud at antics on fuzzy, grainy, and atrociously low-resolution YouTube postings, and we talk to our loved ones on mobile-phone networks with voice quality that would make Alexander Graham Bell roll over in his grave. Information theory tells us that the amount of bits needed to communicate certain kinds of content—what someone is saying, or the antics of a cat, for example—can really be much lower than we think. If we only need to understand the verbal content of someone on the telephone, then the quality can be surprisingly bad and we'll still know what our friends and family are saying. It doesn't seem to matter that so much is missing. Maybe "good enough" is okay.

Or maybe not. Reacting to this tendency, some musicians have decided to go back to analog recording, and some have perversely gone out of their way to make

their recordings sound as lo-fi as possible—as bad as they can get away with. They want to get as far from digital cleanliness as possible. Why would bad quality, fuzziness, and distortion imply that the music is more authentic? The idea is that if one accepts that crisp and clean recordings are inherently soulless, then the opposite, dirty and rough, must therefore be straight from the heart. That might not sound logical, but that's the way we think. It's all part of the recurring belief that conflates new technologies with being inauthentic. Bad—even fake bad, in this way of thinking—means good. It's confusing, because most digital music does not sound "bad." If anything, it sounds conventionally good—clean, spotless, with a full range of frequencies. Though it is actually less rich sounding than previous technologies, it fools the ear into believing that it sounds better. It's this shiny, glossy quality that is considered suspect by many music fans. In response, they overvalue the easily audible drawbacks of a previous era—the hiss, crackle, and distortion. In my opinion, realness and soul lie in the music itself, not in the scratches and pops of old records. So, while the cleanliness and "perfection" of much current music is not a guarantee of a moving musical experience, neither is their opposite.

If, following the lead of the phone company, we find ourselves talking about communication and information transmission when we talk about music, then maybe some of the sonic richness of LPs is indeed superfluous and can be eliminated with no serious loss. Could this work with speech as well? Yes and no. Music has more going on simultaneously than speech, for starters. Looking at a reproduction of a painting is certainly not the same as standing in front of the real thing, but an awful lot of the emotion, intent, ideas, and sensibility can indeed be communicated—even via a cheap reproduction. Similarly, I can be moved to tears by a truly awful recording or a bad copy of a good recording. Would I be moved even more if the quality were higher? I doubt it. So why bother?

There does come a time, however, when the richness of the retinal or aural experience is so diminished that the communication—in this case the enjoyment of the music—becomes unintelligible. But how can we define that? I first heard rock, pop, and soul songs on a crappy-sounding transistor radio, and they changed my life completely. The sound quality was atrocious, but that tinny sound was communicating a wealth of information. Though it was an audio transmission

that carried the news, it was the social and cultural message embedded in the music that electrified me as much as the sound did. Those extra-musical components that got carried along with the music didn't demand a high-resolution signal—good enough was good enough. I'm not saying that tinny sound should be considered satisfying or desirable, or that we should never strive for more than "good enough," but it's amazing how *much* lo-fi or lo-res information can communicate. Live concerts don't generally have perfect sound either, but they can move us deeply.

Now I begin to ask myself if the fuzziness and ambiguity inherent in low-quality signals and reproductions might actually be a factor that gives the viewer or listener a way in. I know from writing lyrics that some details—names, places, locations—are desirable; they anchor the piece in the real world. But so are ambiguities. By letting the listener or viewer fill in the blanks, complete the picture (or piece of music), the work becomes personalized and the audience can adapt it to their own lives and situations. They become more involved with the work, and an intimacy and involvement becomes possible that perfection might have kept at bay. Maybe the lo-fi music crowd has a point?

MUSIC SOFTWARE AND SAMPLE-BASED COMPOSITION

Music composition changed a lot with the advent of digital recording. As we've seen, the first digital samples were short and primarily used by the phone company. People used them for gimmicks and special effects, but these early developments did not have a wide musical influence. Soon enough, though, it became possible to grab or sample a whole bar of music, and though the samples were not of super-high quality, it was enough. Looping beats became ubiquitous, and rhythm tracks made of sampled measures (or shorter intervals) now function as the rhythm bed in many songs. You can "hear" the use of Akai, Pro Tools, Logic, and other digital recording and sample-based composition in most pop music written in the last twenty years. If you compare recent recordings to many made in previous eras, you might not know what makes them sound different, but you would certainly be able to hear it.

The software has affected not just the sonic quality but also the composition process. Maybe you can also hear the effects of the ones and zeros that make up digital recording, though that may be less true as time goes by and technology passes the limits of the differentiating aspects of our hearing. What you do hear, though, is a shift in musical structure that computer-aided composition has encouraged. Though software is promoted as being an unbiased tool that helps us do anything we want, all software has inherent biases that make working one way easier than another. With the Microsoft presentation software PowerPoint, for example, you have to simplify your presentations so much that subtle nuances in the subjects being discussed often get edited out. These nuances are not forbidden, they're not blocked, but including them tends to make for a less successful presentation. Likewise, that which is easy to bullet-point and simply visualize works *better*. That doesn't mean it actually *is* better; it means working in certain ways is simply easier than working in others. Music software is no different. Taking another avenue would make music composition somewhat more tedious and complex.

An obvious example is quantizing. Since the mid-nineties, most popular music recorded on computers has had tempos and rhythms that have been quantized. This means that the tempo never varies, not even a little bit, and the rhythmic parts tend toward metronomic perfection. In the past, the tempos of recordings would always vary slightly, imperceptibly speeding up or maybe slowing down just a little, or a drum fill might hesitate in order to signal the beginning of a new section. You'd feel a slight push and pull, a tug and then a release, as ensembles of whatever type responded to each other and lurched, ever so slightly, ahead of and behind an imaginary metronomic beat. No more. Now almost all pop recordings are played to a strict tempo, which makes these compositions fit more easily into the confines of the recording and editing software. An eight-bar section recorded on a "grid" of this type is exactly twice as long as a four-bar section, and every eight-bar section is always exactly the same length. This makes for a nice visual array on the computer screen, and facilitates easy editing, arranging, and repairing as well. Music has come to accommodate software, and I have to admit a lot has been gained as a result. I can sketch out an idea for a song very quickly, for example, and I can cut and paste sections to create an arrangement almost instantly. Severe or "amateurish" unsteadiness or poorly played tempo changes

can be avoided. My own playing isn't always rock steady, so I like that those distracting flubs and rhythmic hiccups can be edited out. The software facilitates all of this. But, admittedly, something gets lost in the process. I'm just now learning how to listen, value, and accommodate some of my musical instincts that don't always adhere to the grid. It makes things a little more complicated, as I still use the software, but I sense that the music breathes a little more as a result of me not always bending to what the software makes easiest.[B]

Sometimes, after having begun a songwriting process using computer software, I find that I have to sing or play it apart from that recording to free myself from its straightening tendencies. In singing a song freely, for example, I might find that the note at the high point or at the end of a melodic arc "wants" to hang on just a little bit longer than the grid of measures on the computer might indicate as normal. The result is that a verse might end up being nine measures long, rather than the traditional eight. Alternately, a half measure might feel like a nice emotionally led extension, as well as giving a short, natural-feeling breathing space, so then I'll add that half measure to the grid on the software. In a recent collaboration with Fatboy Slim, I discovered that he often added "extra" measures to accommodate drum fills. It felt natural, like what a band would do. Shifting off the grid is sometimes beneficial in other ways, too. If a listener can predict where a piece of music

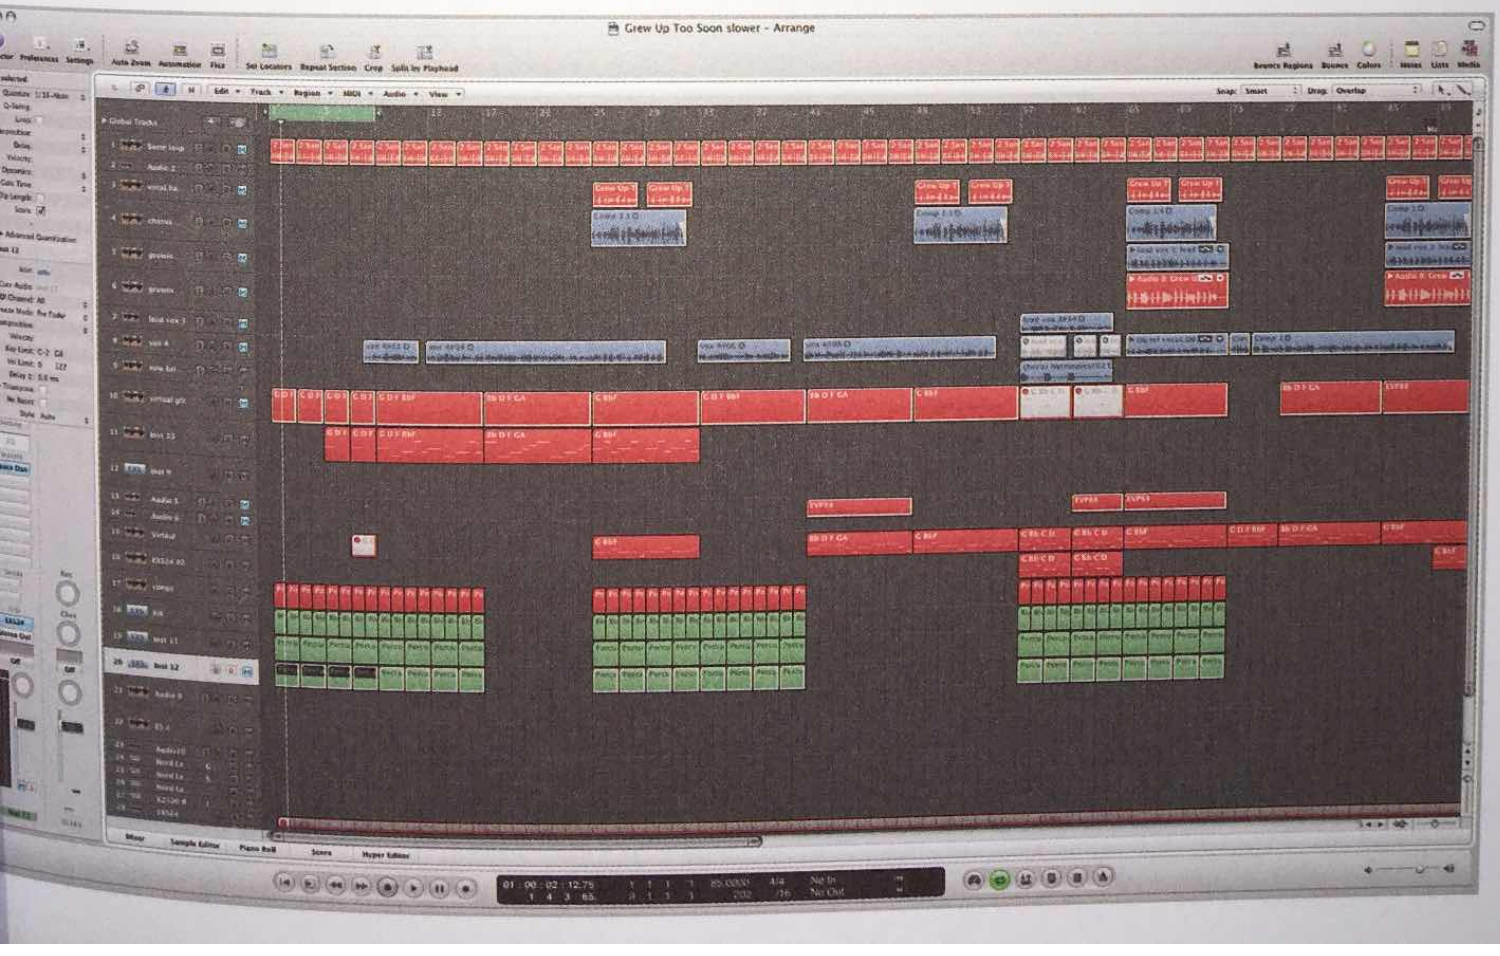

is going, he begins to tune out. Shifting off an established pattern keeps things interesting and engaging for everyone, though it sometimes means you have to avoid the path of least resistance that digital recording software often offers.

Quantizing, composing, and recording on a grid are just some of the effects of software on music. Other effects are created by the use of MIDI, which stands for musical instrument digital interface. MIDI is a software/hardware interface by which notes (usually played on a keyboard) are encoded as a series of instructions rather than recorded as sounds. If you strike a middle A on a keyboard, then the MIDI code remembers when that note was played in the sequence of notes in the composition, how hard or quickly it was struck, and how long it was sustained. What's recorded is this information and these parameters (a bit like a piano roll or the tines on a music box), but not the actual sound—so if that sequence of instructions is played back, it will tell the keyboard instrument to play that note and all the rest exactly when and how you played it previously. This method of "recording" takes up much less computer memory, and it also means the instructional information recorded is independent of the instrument played. Another MIDI-equipped instrument with a completely different sound can be told to play the same note, at the same time. What began as a piano sound might be changed later to synthesizer strings or a marimba. With MIDI recording, you can make arrangement decisions easily and at any time.

The way MIDI remembers how hard or fast you hit a note is by dividing the speed of the note strike into 127 increments. The speed of your hit will be rounded off so that it will fall somewhere in that predetermined range. Naturally, if you strike a key faster, slower, or with more subtlety than the MIDI software and its associated sensors can measure, then your "expression" will not be accurately captured or encoded; it will be assigned the nearest value. As with digital recording, music gets rounded off to the nearest whole number, under the assumption that finer detail would not be discernible to the ear and brain.

There are instruments that can be used to trigger MIDI fairly well: keyboards, some percussion pads, and anything that can be easily turned into switches and triggers. But some instruments elude capture. Guitars aren't easily quantifiable in this way, nor are wind, brass, or most bowed string instruments. So far, the nuances of those instruments have been just too tricky to capture. Using MIDI

therefore tends to entice people away from using those instruments and the kinds of expression they are uniquely good at. A lot of MIDI-based recordings tend to use arrays of sounds generated or at least triggered by keyboards, so, for example, it's easiest to play chord inversions that are keyboard friendly. The same chords on guitars tend to have a different order of the same notes. Those keyboard chords, then, in turn, incline composers to vocal melodies and harmonies that fit nicely with those specific versions of the chords, so the whole shape, melody, and arc of the song are being influenced, not just the MIDI parts and instruments. As soon as technology makes one thing easier, it leaves a host of alternatives in the dust.

The uncanny perfection that these recording and compositional technologies make possible can be pleasing. But metronomic accuracy can also be too easy to achieve this way, and the facile perfection is often obvious, ubiquitous, and ultimately boring. Making repetitive tracks used to be laborious and time-consuming, and the slight human variations that inevitably snuck in as a musician vainly attempted to be a machine were subliminally perceptible, if not consciously audible. A James Brown or Serge Gainsbourg track often consisted of a riff played over and over, but it didn't sound like a loop. Somehow you could sense that it had been played over and over, not cloned. Imagine a row of dancers moving in unison—something that has a huge visceral impact. Besides implying hard work, skill, and precision, it also works as a powerful metaphor. Now imagine that same row created by a series of mirror reflections or CGI technology. Not so powerful.

For many years, DJs, mixers, and hip-hop artists constructed tracks from digital samples of riffs and beats taken from existing recordings. Some artists lifted entire hooks and choruses from pop songs and used them like a knowing reference or quote (P. Diddy does this a lot, as does Kanye West), the way you might quote a familiar refrain to a friend or lover to express your feelings. How many times has "put a ring on it" been used in conversation? Song references are like emotional shortcuts and social acronyms.

In much contemporary pop, if you think you're hearing a guitar or piano, most likely you're hearing a sample of those instruments from someone else's record. What you hear in such compositions are lots of musical quotations piled on top of one another. Like a painting by Robert Rauschenberg, Richard Prince, or Kurt Schwitters made of appropriated images, ticket stubs, and bits of newspaper, this

music is a species of sonic collage. In some ways it's meta-music; music about other music.

However, many artists found out relatively quickly that when songs were created in this way, they had to share the rights and profits with the original record companies and the original composers of those fragments and hooks. Half of the money from a song often goes to the source of a hook or chorus. A few of my songs have been sampled in this way, and it's flattering and fun to hear someone add a completely new narrative to something you wrote years ago. The singer Crystal Waters sampled a Talking Heads song ("Seen and Not Seen"—whoa, what a weird choice!) in her hit "Gypsy Woman (She's Homeless)," and that wonderful song of hers (which I even covered at one point) bears no relation to the original. As far as I know, she or her producers were not intentionally referencing the original song; they simply found something about it, the way the groove felt or its sonic texture, useful.

Besides the fact that being sampled pays well, I don't feel that it compromises the original. Anyone can tell it's a quote, a sample, right? Well, that depends. Trick Daddy, Cee Lo, and Ludacris used another Talking Heads song, the demo "Sugar on My Tongue," in their hit song "Sugar (Gimme Some)." In that case the reference to the original was obvious, to me at least. They used our chorus hook for their own. But since that was never a popular Talking Heads song (it was a demo included in a boxed set), that "quotation" aspect was probably lost on most listeners. (For those keeping track, I did get paid and credited as a songwriter on that track, not a bad second life for what was originally a demo!) Since the song was already a bit of silly sexual innuendo, those guys taking it one step further was no big deal. But, if someone hypothetically proposed repurposing the hook of a song I'd written as a new song about killing Mexicans, blowing up Arabs, or slashing women, I would say no.

Many artists soon began to find that the ubiquitous use of samples could severely limit the income from "their" songs, which led many of them to eventually stop or curtail their use of sampling technology. Sometimes even what is easiest (grabbing a beat from a drum break on a CD, which only takes minutes) becomes a thing to be avoided. Bands like the Beastie Boys picked up instruments they'd put down years earlier, and hip-hop artists either disguised their samples better, found more obscure ones, or, more often, began to create or buy tracks made from scratch (often on synthesizers and drum machines). Technology, or

rather the aspect of technology that enabled the use of copyrighted material, had, thanks to the efforts of rights enforcers and guardians, sent some musicians back to the drawing board. A group of hot young programmers soon emerged whose skills were in constructing, for use by others, these grooves made from scratch. With a relatively inexpensive piece of gear or software, you could make contributions to major songs from your bedroom. In contemporary hip-hop, there is now often no relationship between a composition's backing track and a simulation of a live performance by musicians in the traditional sense. In the early days, there were live DJs using vinyl to loop drum breaks, but now everything—every instrument—is sampled, processed, or in some way shamelessly and boldly artificial.

This music floats free of all worldly reference. Most other pop genres retain some link to simulated live performance, or at least to the instruments used in one, but a song put together with finger snaps, super-compressed or auto-tuned vocals, squiggly synths, and an impossibly fat and unidentifiable bass sound doesn't resemble any existing live band at all. In my opinion, this isn't a bad thing. A new musical avenue has opened up, one that maybe had its roots with Kraftwerk and other electronic acts, but that has now morphed into something very different. It's music that, by design, affects the body. It's very sensuous and physical, even though the sounds themselves don't relate to any music that has ever been physically produced. You can't play air guitar or mimic playing an instrument to a contemporary hip-hop record; even the sounds that signify "drums" don't sound like a drum kit.

These artists therefore have a difficult decision to make when they're expected to perform. Nothing on their record was played by anything that sounds like a "real" instrument, so the performance becomes a kind of karaoke spectacle with most of the sounds pre-recorded—or sometimes artist and band deconstruct the recording in order to "play" its component parts in a way that resembles, visually at least, a traditional band. This sounds like I'm being critical, but actually playing and performing to samples and pre-recorded tracks can free the artist to create something more theatrical. The R. Kelly *Light It Up* tour pushed a traditional R&B concert into the realm of surreal theater. The music in these shows can be viewed as a soundtrack for a spectacle; a gathering; a show of camaraderie, visual bombast, and effects. Grace Jones and Pet Shop Boys did this years ago. They were mainly dance-club bands and visual icons, not live acts, and they were both very

arty about it. Now this kind of karaoke spectacle has gone mainstream. Even if the lead vocal is live, the audience often doesn't care whether there is a band or not—the band is often just there to dress the set. I will admit that there is an inevitable excitement ceiling when it comes to these karaoke spectaculars, at least as far as the music goes, since there is never the possibility of the music rising beyond its preexisting programmed level (and I don't mean volume). But if the social and the visual elements are good enough—and using this technology they are freer to become so—then it can be a reasonable trade-off.

Hip-hop artists have also revived the mixtape (or maybe it never went away). Sampled beats and digital recording have allowed mixtape artists to create seamless sequences of songs, dialogue, beats, and assorted weirdness in ways that were never really possible using cassette recorders. Most of these are "released" as CDs, though the name harks back to the cassette era. But I have a feeling that the era of the CD mixtape is also coming to an end. I have digital mixes that are one continuous hour-long piece of music, but now, in a technical sense, there's really no limit. The artificial length that Sony imposed on CDs no longer applies. You could make a downloadable album that's ten hours long, or one that's only ten minutes. Jem Finer recently created a piece of music that runs for a thousand years. Marketing and promotion aren't very cost effective for one song at a time, so temporarily at least we still market and acquire songs in clumps (an hour or less is typical), as we have for decades. But there's no reason that will remain the standard indefinitely.

Looking at the short end of the musical-length spectrum, ringtones haven't yet been recognized as a valid form of stand-alone musical creativity, and they may never be. But shortness shouldn't matter anymore. One could argue that the Mac start-up sound is a musical composition, as is the short, mysterious, ascending five-second bumper that signaled the end of each scene in the TV series *Lost*.

C

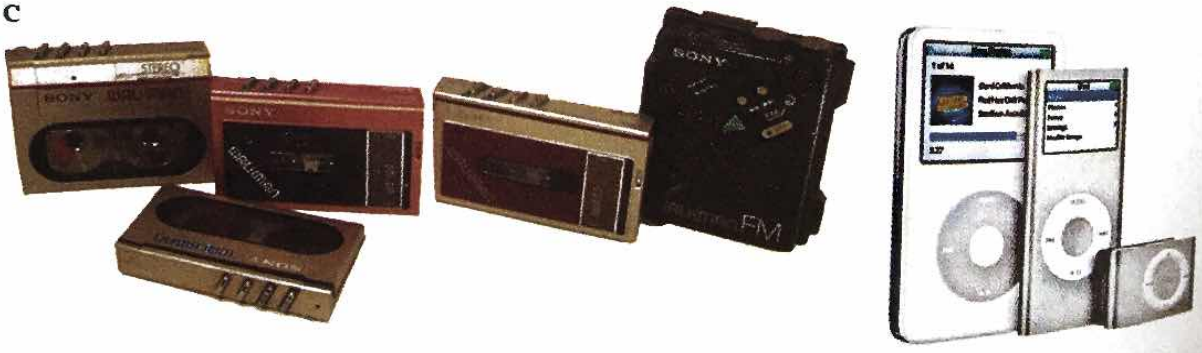

Doorbells, whooshes, email alerts, and car horns are all valid forms of composition. Our musical landscape is indeed broadening, as length doesn't matter anymore: short, long, and in between all coexist.

PRIVATE MUSIC

The iPod, like the Walkman cassette player before it,[C] allows us to listen to our music wherever we want. Previously, recording technology had unlinked music from the concert hall, the café, and the saloon, but now music can always be carried with us. Michael Bull, who has written frequently about the impact of the Walkman and the iPod, points out that we often use these devices to "aestheticize urban space."[4] We carry our own soundtrack with us wherever we go, and the world around us is overlaid with our music. Our whole life becomes a movie, and we can alter the score for it over and over again: one minute it's a tragedy and the next it's an action film. Energetic, dreamy, or ominous and dark: everyone has their own private movie going on in their heads, and no two are the same. That said, the twentieth-century philosopher Theodor Adorno, ever the complainer, called this situation "accompanied solitude," a situation where we might be alone, but we have the ability via music to create the illusion that we are not.[5] In his somewhat Marxist way, he viewed music as an opiate, especially popular music. (I've met some serious Wagner fans, and I'd be wary of limiting the accusation that music is an addictive palliative solely to pop.) Adorno saw the jukebox as a machine that drew "suckers" into pubs with the promise of joy and happiness. But, like a drug, instead of bringing real happiness, the music heard on jukeboxes only creates more desire for itself. He might be right, but he might also have been someone who never had a good time in a honky-tonk.

Private listening could be viewed as the height of narcissism—these devices usually exclude everyone else from the experience of enjoying music. In *Brave New World*, Aldous Huxley imagined a drug called soma that blissed everyone out. It was like taking a holiday, and you could regulate the length of the holiday by the dosage. Has technology turned music into a soma-like drug? Is it like a pill you take that is guaranteed to generate a desired emotion—bliss, anger, tranquility?

ACOUSTIC ASSOCIATIONS

Some of us sing to ourselves or even just whistle when we think no one is listening, and often there is music "playing" in our heads. A region of the brain seems to be devoted primarily to sonic memory, and that includes not just ringtones, dog growls, and ambulance sirens, but also snippets of songs, mainly recordings, that we've heard as well. These sonic fragments function as nodes in a network of related memories that stretch beyond their acoustic triggers. Everyone has had a song "transport" them to a vivid memory of an early romance or some other formative experience. Songs are like smells that way; they dredge up worlds, very specific places and moments. Other sounds do this, too: intense rain, the voice of a well-known actor, a knife on a cutting board, a distant train.

Are mobile music devices and the musically cluttered world we inhabit starting to substitute for our interior voices? Do we, little by little, stop singing and whistling because professionals are now singing and playing for us, right into our ears? A slew of musical associations bounce around in our heads, linking to recurring memories and feelings that, after a while, facilitate the creation and reinforcement of specific neural pathways. These pathways help us make sense of those experiences. They make us who we are. Is that space now inhibited by the inundation of the music and sounds of others? Are the voices in our heads, the inaudible chatter that we use to sort out who and where we are, being replaced by the voices of professionals? Well, I haven't stopped singing to myself, so maybe not.

THAT WHICH CANNOT BE PRESERVED

I listen to music only at very specific times. When I go out to hear it live, most obviously. When I'm cooking or doing the dishes I put on music, and sometimes other people are present. When I'm jogging or cycling to and from work down New York's West Side bike path, or if I'm in a rented car on the rare occasions I have to drive somewhere, I listen alone. And when I'm writing and recording music, I listen to what I'm working on. But that's it. That relatively short list defines, to a large extent, where and how I hear music. I find music somewhat intrusive in restaurants

or bars. Maybe due to my involvement with it, I feel I have to either listen intently or tune it out. Mostly I tune it out; I often don't even notice if a Talking Heads song is playing in most public places. Sadly, most music then becomes (for me) an annoying sonic layer that just adds to the background noise. It might sound like I'm a picky eater, but I actually do listen to a lot of music.

As music becomes less of a thing—a cylinder, a cassette, a disc—and more ephemeral, perhaps we will start to assign an increasing value to live performances again. After years of hoarding LPs and CDs, I have to admit I'm now getting rid of them. I occasionally pop a CD into a player, but I've pretty much completely converted to listening to MP3s either on my computer or, gulp, my phone! For me, music is becoming dematerialized, a state that is more truthful to its nature, I suspect. Technology has brought us full circle.

I go to at least one live performance a week, sometimes with friends, sometimes alone. There are other people there. Often there is beer, too. After more than a hundred years, we're heading back to where we started. A century of technological innovation and the digitization of music has inadvertently had the effect of emphasizing its social function. Not only do we still give friends copies of music that excites us, but increasingly we have come to value the social aspect of a live performance more than we used to. Music technology in some ways appears to have been on a trajectory in which the end result is that it will destroy and devalue itself. It will succeed completely when it self-destructs. The technology is useful and convenient, but it has, in the end, reduced its own value and increased the value of the things it has never been able to capture or reproduce.

Technology has altered the way music sounds, how it's composed, and how we experience it. It has also flooded the world with music. The world is awash with (mostly) recorded sounds. We used to have to pay for music or make it ourselves; playing, hearing, and experiencing it was exceptional, a rare and special experience. Now hearing it is ubiquitous, and silence is the rarity that we pay for and savor.