

## Recent Developments at the Columbia University Computer Music Center

By Brad Garton

### Introduction

Columbia University has had a long involvement with music technology, establishing one of the first, if not the first, research/music centers devoted to electronic music in the United States. Officially recognized in the late 1950s as the Columbia-Princeton Electronic Music Center, the EMC was a hotbed of musico-technological work in the ensuing decades.<sup>1</sup>

A few years ago I became Director of the Center—its new advisory board comprising Fred Lerdahl, Tristan Murail, and myself. We managed to secure a sizable boost in funding from the Columbia University Administration and from several external sources, and with this influx of new support we decided to rebuild a number of our studios and to undertake a major overhaul and revamping of the Center's facilities. We also decided to rethink the operation of the Center, seeking to renew the status it enjoyed for decades as an advanced and progressive workplace for musicians and researchers who use new music technologies.

At that time we officially changed the name from the Electronic Music Center to the Computer Music Center (CMC) to better reflect the new organizational structure as well as the renewed research/music focus. We have since enjoyed a tremendous increase in activity at the CMC, with all of the attendant excitement and difficulties associated with explosive growth.

Several months ago Dan Thompson asked if I would write a description of some of the changes that have taken place at the CMC for *Current Musicology*, perhaps thinking that some of what we do might be of interest to *CM* readers. Rather than merely describing hardware and software projects, I thought it might be more interesting for me to try to articulate my version of the philosophy driving what we now do at the CMC.<sup>2</sup> What follows is an attempt to do just that. I feel I must apologize in advance for the decidedly personal tone of this article; however, the operation of the CMC is indeed a personal odyssey for everyone involved. One final caveat: what I describe is truly my own version of how the Center is, and it may or may not reflect the actual reality of the CMC. I like to pretend that it does.

## The CMC

I remember when I was in graduate school, my thesis advisor used to say to us: "Don't become the director of any sort of 'center.' It's the kiss of death!" And of course I now find myself, a little over a decade later, director of the Columbia University Computer Music Center. I guess I learned my graduate school lessons well. Indeed, many days do feel like a slow and painful act of final mortality, but I hope that I am not quite yet the corpse my advisor envisioned. I think he had a particular fatality in mind with his "kiss of death" statement—the death of creativity, of innovation, and of music (we are both composers). I also think that the "kiss of death" observation was motivated by a conception of what it meant to *be* a director of a "center" back then, especially considering the circumstances that provided a context for the definition of a "computer music center" more than a decade ago.

The point of this essay is not to outline those circumstances or to describe the context that existed for computer music work in the 1980s (see Georgina Born's (1995) fascinating description of IRCAM, a well-known music-technology research center in Paris, in the mid-80s for a detailed look at this world). To be sure, the environment for a contemporary computer music center probably hasn't changed much in ten years. However, we are attempting to build a different sort of computer music center at Columbia, and hopefully learning from past "kiss-of-death" types of mistakes. What I would like to do in this paper—at the risk of appearing massively self-delusional—is to highlight a few of the alternative organizational and philosophical approaches we are implementing as we move the CMC into the new millennium.

As far as centers go, the Columbia CMC is rather decentralized. This is partly a result of the recent history of computer music at Columbia. The relatively gradual growth of support for computer music within the older structure of the Columbia University Electronic Music Center, instead of a single "establishing moment," precluded the adoption of a strong central authority overseeing all computer music activities. The decentralization is also partly by design, for we have noticed that many practitioners of computer music work best in a rather loosely structured environment. We also consider one of our primary goals to be the creation of a center that exists to support the work done by students, researchers, and composers, regardless of the particular aesthetic or musical direction engendered by this work. In other words, the direction of the CMC is charted primarily through use. In place of formalized schemes or organized N-year plans we tend to go where users of the Center are taking us.

This self-organizing approach to defining the CMC's direction has several immediate consequences. The Center's hardware and software

foundation must necessarily be broad, because often a particular research project or musical composition requires specific software packages that run on a certain make and model of computer, possibly with unique peripherals and input devices (MIDI controllers, data gloves, distance sensors, etc.). To meet this need, we have attempted to purchase as wide a range of digital machinery as our budget will allow. At present, most of the major combinations of hardware, operating systems, and software currently used for computer music work are represented at the Center. We are committed to maintaining this array of equipment and software resources. Within budgetary constraints, our purchasing decisions are generally dictated by the needs of the CMC user community. We want to buy machines that will be used!

Maintaining this broadly based infrastructure poses two direct difficulties. The first is, well, simply *maintaining* the infrastructure. Hardware breaks, software configurations get trashed, wires come unplugged, disk drives fail . . . All of these ongoing (and very real) problems place an enormous load on the CMC staff. I have yet to visit a contemporary high-technology academic research center where the support staff wasn't overtaxed, overworked, and overburdened. The Columbia CMC is no exception. It is not unusual during peak times of the academic year for our technical staff to spend 12–15 hours a day putting out technological "brush fires" to keep the CMC running smoothly. This situation cannot continue indefinitely.

One way to lessen the burden of at least routine maintenance work is to involve Center users directly in our support procedures. I would hesitate to cast the CMC as a kind of post-60s technological commune, but we do try to nurture a communitarian spirit, of sorts. People working at the CMC generally recognize that a small investment of their time can help make the Center a more productive place. In general, we allow all of our users to take as much responsibility for configuring and maintaining our hardware as they wish—provided that this does not interfere with other users working on the system. We often find that individual students or researchers have detailed personal knowledge of a particular machine or software package. This "knowledge bank" among our users is invaluable to us as we confront the plethora of hardware/software possibilities, each with idiosyncratic configuration features that must be known for proper operation.

Another obvious way to lessen the maintenance burden is to hire more staff. This "solution" intersects with the second of the immediate difficulties encountered in trying to maintain a broadly based technological infrastructure: budget. In a world where hardware is nearly obsolete the day it is shipped, a solid foundation of financial support is a necessity. Even remaining barely "even" with new innovations in technology requires a con-

stant reinvestment in basic machinery at the Center. Compound this with the additional support and maintenance time and budget needed for incoming new hardware/software and the downward budgetary spiral begins to become apparent: new equipment and software needs more support, but additional support requires monetary commitments, leaving less for new equipment and software, which must be purchased to remain technologically current, but the new equipment and software needs more support . . . etc. Every center currently engaged in a fundamental way with new technology will probably *never* have a budget sufficient to meet demand.

At Columbia, we are fortunate in having an administration that recognizes the necessity of providing at least a modest amount of direct support for technology. It has become almost a cliché to say that most progressive universities and institutions of higher education are aware that a strong technological base will be essential for future survival in an increasingly competitive academic market. Columbia is no exception, and the CMC has been the beneficiary of this state of affairs. However, the amount of annual support we receive earmarked as operating budget for the Center does not begin to approach what is needed to maintain our technological viability. To make up for this difference, we have to seek—as many other centers do—outside sources of funding.

This is where our “open door” policy toward work done at the CMC has truly paid dividends. Nearly all of the projects that have generated external income for us in the past few years have originated in a use of our facilities that would not have been envisaged had we adopted a narrow, hierarchical definition of what the CMC should be doing. Certain individuals made specific uses of our facilities and capabilities—uses we often had not anticipated when setting up the Center—and these alternative uses grew into relatively lucrative income-generating projects.<sup>3</sup>

To be honest, this is probably how most other centers operate. To a greater or lesser extent, projects are generally driven by the individuals involved in them instead of by “official” institutional sanction. Our plan is that by *explicitly* articulating an “anything-sort-of-goes” attitude, together with the range of resources we provide, we will create an extraordinarily fertile environment for the gestation of new and innovative projects. We don’t want to eliminate *a priori* any possible avenues for fruitful musical investigations by adopting an artificial set of limits on what the CMC should be doing.

A side effect of this policy has to do with how the CMC relates to other divisions of the University. I recall that when I first became director of the Center, we had many long discussions on how to define our relationship to the music department, other departments, other university research

centers, etc. These discussions are ongoing, as very real issues surface concerning allocation of specific resources. To a large extent, however (and lengthy discussions notwithstanding), the CMC has already become tightly integrated into the workings of the music department. This integration has occurred not as a result of any planned effort on our part, but as a consequence of individual student and faculty projects that have become what the Center does.

We have also established an excellent working relationship with the Columbia music library (and in fact with the greater Columbia library system), and we are becoming involved in collaborations with the film division, the engineering school, the medical center, and the chemistry department, as well as with many other units within the university. Again, these connections have all occurred through specific projects and initiatives arising from the CMC user community—not from the implementation of some pre-planned and agreed-upon “CMC Objective.” Rather than forcing a conception of how the CMC should relate to other university entities, the collaborative projects have already produced the best possible definition of Center policy regarding interdepartmental relations: relationships based upon a mutual pursuit of common goals as embodied in actual, ongoing work.

We would like to expand this approach to collaborative ventures beyond the walls of Columbia University. In the past, there has been little substantive cooperation between different centers for computer/contemporary music. In truth, it was probably necessary for centers to establish their independent, autonomous existence before any intercenter collaboration was possible. Recently, however, several music technology centers have begun to work together on joint projects. We are among those centers, for we believe that diverse perspectives can greatly assist the development of these projects. As with our other work, the approach we are attempting to take is to aid the creation of self-generating projects, rather than dictate from the top which “collaborations” (even if in reality they may be quite empty of real content) we will undertake. So far, this approach appears to be functioning quite well. We are currently engaged in fledging projects with Princeton University, the University of Virginia, the University of Thessaloniki, the Tokyo College of Engineering, the National Center for Supercomputing Applications, the National University of Uruguay, IRCAM, and a number of commercial enterprises.

Stepping back a bit from the local circumstances of the CMC, this notion of decentralized planning seems part of a larger phenomenon—call it postmodern management if you will. The fragmentation and lack of central authority that have been cited as salient features of postmodern philosophies are generally seen in a negative light: hierarchical high-

modernist edifices are *destroyed*, are *demolished*, are *deconstructively reduced* by postmodern thinking. I suggest that the activities and organization of the CMC represent a more positive postmodernism. To borrow a concept from artificial life or neural network research in computer science, the CMC functions as a self-organizing, “bottom-up” system, where aspects of centralized control and coordination are emergent features of a *confederation of autonomous users*. The Center operates almost as a logical or virtual construction, an entity that comes explicitly into existence to meet the demands of a particular situation.<sup>4</sup>

The decentralized nature of the CMC has a pronounced effect upon some of our products. A culture of shared information is nurtured by a self-organizing approach; thus, the software we develop is—at the discretion of the individual responsible for the work, of course—nearly always public domain, and source code is generally freely available to all. I contrast this with a suite of distribution-protected software packages we recently purchased from a notoriously hierarchical, epitome-of-high-modernism organization. The software was a nightmare to install, mainly because it assumed very specific machine and network configurations. We were required to duplicate large and relatively tangled parts of the selling organization’s hardware/software structure just to get the programs running. Had we access to the protected source code (or had the seller been a bit more accommodating of diversity in machine configuration), several days of painful, frustrating installation work would have been reduced to a matter of minutes. I’d like to imagine that software developed at the CMC—perhaps because it is developed in a heterogeneous and constantly shifting environment—is a little more tolerant of different computer configurations. We are actually developing some of our larger software applications on several machine/operating-system architectures simultaneously because the range of machines at the CMC makes the expediency of doing this quite obvious. And of course, most source code for our programs is available for the taking, making any reconfiguration (or enhancements!) easy for knowledgeable individuals to do.

As fabulous as this decentralized, bottom-up approach to Center organization may seem, it does dodge important decisions that still must be made from a top-level perspective. Although it is theoretically wonderful to speak of the CMC as this fantastic logical/virtual construct that forms and reforms for various user-initiated projects, many of our larger undertakings do require a serious allocation of limited resources. In actual practice, we administer the Center by making decisions about the level of support we can provide for individual projects. These decisions are becoming more and more difficult as the activity of the Center increases. Although we must act in a decidedly hierarchical fashion when choosing what we

can and cannot afford to support, a crucial distinction to be made is that the projects forcing the decision originate in the user community at large; they do not begin as officially sanctioned projects arising from a top-level, central CMC authority.

The question of what kind of work the CMC should be doing reappears in resource-allocation choices. My tidy picture of a happy group of users cooperatively determining the direction of the CMC begins to crack and crumble for many people in the face of criticisms about the legitimacy of the kind of music composition and research done at the Center. Every user has a different concept of what she thinks should be of central concern for the CMC, a few people going so far as to dismiss composition and research not in line with a particular viewpoint or aesthetic as somehow not *really* doing music composition or research. An example of this dynamic in action can be seen in the CMC's engagement (or non-engagement) with "multimedia" (film/video/etc.). We have not invested heavily in video production or graphic design systems at the Center, although I have seen many wonderful and exciting multimedia computer music works in the past few years. This is partly because of the high costs of these systems, but it is also partly due to a prevailing attitude at the Center (and in the music department) that this sort of work isn't "pure" music. This situation is changing, however, as several students have recently become heavily involved in multimedia projects. I predict that it will be through these collaborative, cross-disciplinary projects that we eventually move to include film/video equipment as part of the CMC facilities. This is a real case where specific individual projects are driving the direction of the Center, but in this particular instance there exists a strong budgetary counter-pressure, partially fueled by notions of musical legitimacy, that works against the realization of these individual projects. The challenge that we face at the CMC is to meld disparate conceptions of what we should be doing into a manageable, affordable environment that will not discourage innovative work. The boundaries that we set must be semi-permeable.

Concerns about what is "central" to the CMC's direction are related to another issue: the pedagogical role played by the CMC. What should we be teaching? How should we teach it? Terry Pender (the technical director and a composer who works at the Center) recently participated in a panel discussion that took place during a computer music conference in Tokyo. The panel was charged with addressing the question, "What does a computer musician need to know?" This pernicious issue seems to surface repeatedly as people strive to codify and make sense of a rapidly changing technological/musical world. Despite many rather heavy-handed pronouncements made about what musicians must know to create True Art, Terry's basic response was that we should—to the best of our abilities—teach what people need to know to accomplish their own personal goals.

The emphasis here is placed again on the individual instead of an external doctrine or tradition that must be absorbed in order to produce real music. Our students and researchers at the CMC are now coming from a wide variety of cultural backgrounds. Enforcing a unitary view of musical knowledge could produce a cultural schizophrenia that might easily destroy individual creativity.

While I certainly endorse Terry's attitude, the truth is that the very act of selecting the issues to address pedagogically in a field as broad as contemporary computer music does carry heavy aesthetic presuppositions. My hope is that we can maintain a relatively "open" approach by relying on a philosophy of individual engagement coupled with a commitment to maintaining a diverse population of students and researchers working at the CMC. I would also want to replace the "need to know" question with an alternative: "What sort of musical community do we want to build?" For me, ephemeral questions of musical knowledge often reduce to issues of social relations and the reinforcement of particular social hierarchies and structures—I guess I subscribe to the Foucauldian view of how knowledge/power works in the world. Knowledge of computer music is still relatively young, however, and my personal desire is that we may build a community of knowledge at the CMC that is reflective of the liberal, egalitarian values I cherish. As is obvious by now, I also believe that this is the best possible way to organize a center dedicated to nurturing creativity and promoting musical diversity (diversity and creativity being inextricably interrelated in my view).<sup>5</sup> Even if we fail, the attempt to create a light and shallow controlling structure at the Center will be an interesting experiment.

I do think that circumstances in the world have changed to allow space for the kind of infrastructure and management philosophy manifested by the CMC. The growth of international travel is allowing personal contacts to break down deep-seated cultural prejudices and foster a heightened intersocial awareness, and the Internet (although probably not the liberator of mankind that techno-idealists proclaim) has surely created a climate of individual potential. Even my former advisor, with his dire "kiss of death" warnings about administration, is now chair of a prestigious music department. Since becoming chairman, his music-compositional output has nearly tripled. Maybe there is hope for us yet.

#### *Notes*

1. A detailed history of the Center may be found on our web site: <http://www.music.columbia.edu/cmc>

2. For an up-to-date description of the CMC facilities, along with a description of current projects (and software available for downloading), see our web site: <http://www.music.columbia.edu/cmc>



3. I realize I don't give specific examples of these projects—this paper is intended to provide a philosophical overview of the CMC. I refer again to our web site for a good listing of work being done at the CMC.

4. Kelly (1994) presents an interesting set of essays about this sort of organizing methodology in a variety of areas.

5. Many "postmodern management" theorists agree with this approach. See Kao (1997) for a popular, albeit somewhat breathless, example.

*References*

Born, Georgina. 1995. *Rationalizing Culture*. Berkeley: University of California Press.

Kao, John. 1997. *Jamming*. New York: HarperCollins Publishers.

Kelly, Kevin. 1994. *Out of Control*. Reading, MA: Addison-Wesley Publishing Co.