About This Issue

The year 2006 marks the 30th anniversary of the founding of *Computer Music Journal.* To honor the occasion, we invited the Journal's founder and first Editor, John Snell, to write the opening article of Volume 30. A brief history of *Computer Music Journal* had been published in Volume 10, Number 1 (Spring 1986). That article described the main events through the decade following the Journal's first issue. In the present issue, John Snell goes further back, reminiscing about the circumstances and thinking leading up to the founding in 1976. His article also discusses in greater depth the Journal's first few years, when it was published by the People's Computer Company in Menlo Park, California. Mr. Snell's writing brings to life a certain ethos of those early years, when a fondness for accessibility, shared knowledge, and creativity eclipsed any notion of profit, as evidenced by the remarkably long hours he himself dedicated to the Journal. (Not coincidentally, these were the years that also nurtured what came to be known as the open-source movement.) We are indebted to John Snell for contributing this glimpse into the past, and—need we mention?—even more so for having laid the foundations of this publication so diligently and deftly.

The next three articles of this issue deal with expressive control of music performance, with an emphasis on the mapping of high-level control parameters to lower-level rendering parameters. The three articles' approaches differ, particularly in terms of the nature of the target data, i.e., what is being controlled.

For example, one can control the playback of a stored audio (and video) recording, mapping the user's gestures into functions for time-varying tempo and amplitude. Such is the focus of the article by Eric Lee, Thorsten Karrer, and Jan Borchers. Their work involves interactive systems in which a user "conducts" digital video and audio streams that contain a pre-recorded orchestral performance. The systems the authors describe—Personal Orchestra, You're the Conductor, and a recent hybrid of the two—permit gestural control of tempo, dynamics, and instrumental balance. Their present article, which focuses on tempo, introduces the authors' improved audio time-stretching algorithm as well as a general model of time that is applicable to a wider range of multimedia applications.

Instead of controlling an audio (and video) signal, one can control a higher-level musical representation, such as a MIDI sequence. In this case, the conversion of gestural parameters to musical output can benefit from research on the ways in which performing musicians expressively deviate from the nominal values of a musical score. The article by Anders Friberg explains how the pDM sequencer from the Royal Institute of Technology (KTH) in Stockholm maps real-time, simple user controls onto an intricate matrix of rules for expressive performance. (As mentioned in this issue's News section, the KTH system won a recent competition of expression-rendering software.)

Finally, one can control the actual generation of musical material in real time, as opposed to simply applying expressive alterations to a previously stored audio recording or MIDI sequence. Ali Momeni and Cyrille Henry's article presents a model for mapping high-level gestural parameters into lower-level controls for synthesis of both audio and video. The authors stress the importance of an independent layer that sits between the gestural data and the audio and video synthesis parameters. This layer, which they call a dynamic visual mapping layer, is independent in that it has no direct knowledge of either the input or the output values, being insulated from them by additional layers that do have such

Front cover. An abstract, music-inspired image by Kyong Mee Choi.

Back cover. This sequence of four graphics, borrowed from the article by Ali Momeni and Cyrille Henry, shows a simulated mass-spring network changing over time. The network serves as a mapping layer between the user's gestures and real-time audiovisual synthesis.

knowledge. It is dynamic, in that the mapping algorithm can be time-varying, and visual, in that it can somehow be reflected in the user interface. Two specific examples of such mapping layers are described: one using systems of masses and springs, and another using multidimensional control spaces.

The articles by Mr. Lee and his coauthors and by Messrs. Momeni and Henry both deal with simultaneous control of audio and video signals (prerecorded signals, in the former article, and signals synthesized in real time, in the latter). This overlap of topics is not coincidental: these two articles, along with the final one in this issue, were submitted for last year's special issue on computer music/video, which was organized by associate editor Brett Terry. His call for submissions was so fruitful that not all the chosen manuscripts (which were only a fraction of what we received) could fit into the Winter 2005 issue. These three were held over for the current issue.

The final article, by composer Rodrigo Cádiz, presents a technique for mapping sounds to images and vice versa; but unlike the previously mentioned systems, his method is employed out of real time. Mr. Cádiz's approach employs fuzzy logic to relate acoustic parameters to visual ones. Two multimedia artworks are discussed: one in which the musical composition served as the input to an algorithm that generated the corresponding graphics, and another with the reverse directionality.

The Reviews section of this issue covers, among other things, books on the composers Herbert Brün and John Chowning, a textbook on audio encoding (such as various MPEG and Dolby formats use), a conference, a festival, and two discs. It may be worth reminding readers that some reviews written for the Journal are published only on the Web. (See the URL given at the beginning of the Reviews section.) The issue's concluding section, Products of Interest, which is edited by Margaret Cahill and James Harley, describes a variety of new audio- and music-related software and hardware.