Turenas: the realization of a dream

John Chowning

[1] CCRMA, Department of Music, Stanford University, Stanford, California 94305
jc@ccrma.stanford.edu

Abstract. The culmination of two research paths, *Turenas* (1972), embodies both sound localization (begun in 1964) and Frequency Modulation (FM) synthesis (begun in 1967). Illusory motion of sound in space was a musical goal from the very beginning of my work with computers, a goal that led to a number of perceptual insights that not only set up my discovery of FM synthesis but insights that have enriched my thinking about music ever since.

In tracing the long compositional trajectory of *Turenas*, I will show the connections between the work of Mathews, Schroeder, and Risset, at Bell Telephone Laboratories, and the development of musical tools by David Poole, Leland Smith and myself within the research community at the Stanford Artificial Intelligence Laboratory (A. I. Lab). The presentation will include sound-synchronous animations.

1 Introduction

In 1957, more than fifty years ago, Max Mathews introduced a wholly new means of making music. Mathews (with the support of John Pierce, who was director of research) created out of numbers and code the first music to be produced by a digital computer using a program that was extensible, pointing to a future of which we are all now a part. In his famous 1963 article [4] Mathews described his work and this program, Music IV.

In January 1964 I was serendipitously given this article by a friend (who worked in the sciences), based upon her understanding of my interest in electroacoustic music. Before coming to Stanford as a graduate student in music composition in 1962, I had studied three years in Paris where electroacoustic music was often a part of new music concerts. I was fascinated by the idea of composing music for loudspeakers and, in particular, by the idea of composing music in spaces that were compositional constructions in which sound could be positioned, animated and even moved through space—as was suggested by some electroacoustic works, especially those of Stockhausen, e.g. Gesang der Jünglinge (1956). At the time I was completely naïve in regard to physical acoustics and the engineering sciences and had never seen a computer. But, I did understand Mathews' striking statement —"There are no theoretical limitations to the performance of the computer as a source of musical sounds, in contrast to the performance of ordinary instruments."

After taking a programming course that convinced me that I could learn to program a computer, the following summer I visited Mathews and he provided me with a box

of punched cards that represented the Music IV program and the Bell Telephone Laboratories (BTL) compiler. David Poole, an undergraduate math student and "hacker" at Stanford's newly established A. I. Lab, implemented the program using an IBM 1301 disc, the common storage unit between an IBM 7090 and DEC PDP-1 computer. Samples were generated by Music IV and written onto the disc drive. Poole wrote a double buffer program that read the samples at a constant rate to the x and y ladders of the PDP-1 Precision CRT display, having 1024 by 1024 addressable locations, plotting 20k points per second. The two analog signals were then recorded on a stereo tape recorder.

Always at Poole's side, I learned from him and others about computers, programming and acoustics. In 1966 the A. I. Lab acquired a PDP-6 computer (predecessor of the PDP-10). Poole rewrote the Music IV program in PDP-6/10 assembly language and I wrote the function generators in FORTRAN IV. By 1967 we called this Music 10 with the arrival of the PDP-10. A 4-channel DAC was designed and built by an electrical engineer who was interested in the project. Thus, I had the computer, programs and special DAC hardware that were required to begin my quest that would finally result in *Turenas* five years later.

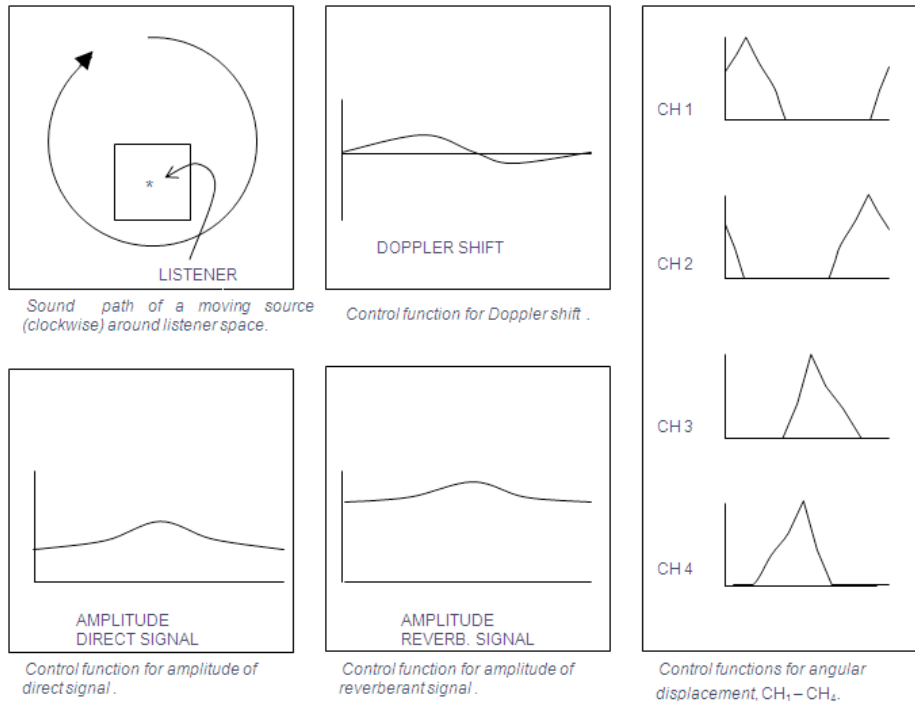## 2 Moving Sound Sources and Score

While much was understood in the early 1960s about stereo imaging in the azimuthal dimension, I was unable to find convincing information about distance or the radial dimension. Even nowadays there seems to be no agreement as to the perceptual cues for distance, as reflected in this statement from 2002 "Remembering that timbral detail is the primary determinant of distance location …"[5]. I knew from careful listening that reverberation must be an important component of distance perception, but I was unable to determine in what way. Music IV did not have any "reverb" unit generators, but it did have delay. Beginning on the IBM 7094/PDP-1 system and then moving to the PDP-6/10, I built reverberators using re-circulating delay units (comb filters). I then found Schroeder's article on natural sounding reverberation [8] and implemented a form of same in Music 10.

Leland Smith joined me in the development of the music system in 1966. Seeing a need to structure the input process to Music IV/Music 10 programs, Smith wrote the program Score that accepted lists of data by "instrument" as input that were correlated and turned into the sequenced "score" as output. Score was used to produce the first compositions from the music project at the A. I. Lab and I used Score to create stimuli in my development of the spatialization project.

2.1 Doppler shift, azimuth and distance

By 1967 I had completed the spatial processing system that would allow me to place a sound at a discrete location outside the perimeter of a square where there was a loudspeaker at each of its corners. In the following year I achieved a graphic solution to controlling the Doppler shift in frequency as a function of radial velocity. While I

had understood that to properly model the Doppler effect required variable delay lines, it was an option that was not practical because of the computational limits.

Sound path of a moving source (clockwise) around listener space.

Control function for Doppler shift.

Control function for amplitude of direct signal.

Control function for amplitude of reverberant signal.

Control functions for angular displacement, $CH_1 - CH_4$.

Fig. 1. From the sound trajectory drawn or computed around the listener space, as shown in the upper left, the program automatically computed the six control functions

The program to control the movement of sounds through an illusory space was largely complete by 1968 as seen in Figs. 1 & 2. I had generated a number of simple geometrical sound paths in order to evaluate the strength of the illusions. Of course, the azimuth (perceived angle) of the source was highly sensitive to the listener position within the listener space. But as long as the listener was not very close to one of the loudspeakers, the spatial distortion did not seem to destroy the illusion. Having four channels in a surround configuration seemed to produce effective illusions even for poor listener positions, as the direct signal could move through two or three of the channels even when the listener was just next to one of the four speakers.

2.2 Distance and localization of the reverberation

My conception of the spatial processing system evolved as I gradually learned about perception, the physics of sound and reverberant spaces. The perception of distance seemed not so sensitive to the listener's position relative to the loudspeakers. While subtle and easily covered, the distance cue seemed to me to be more robust than the azimuthal cue. The direct signal's amplitude decreases in inverse proportion to an

increase in distance. The amount of reverberation must remain more or less constant, however, providing for a distinction between a sound whose amplitude decreases because of an increase in distance as opposed to a decrease in
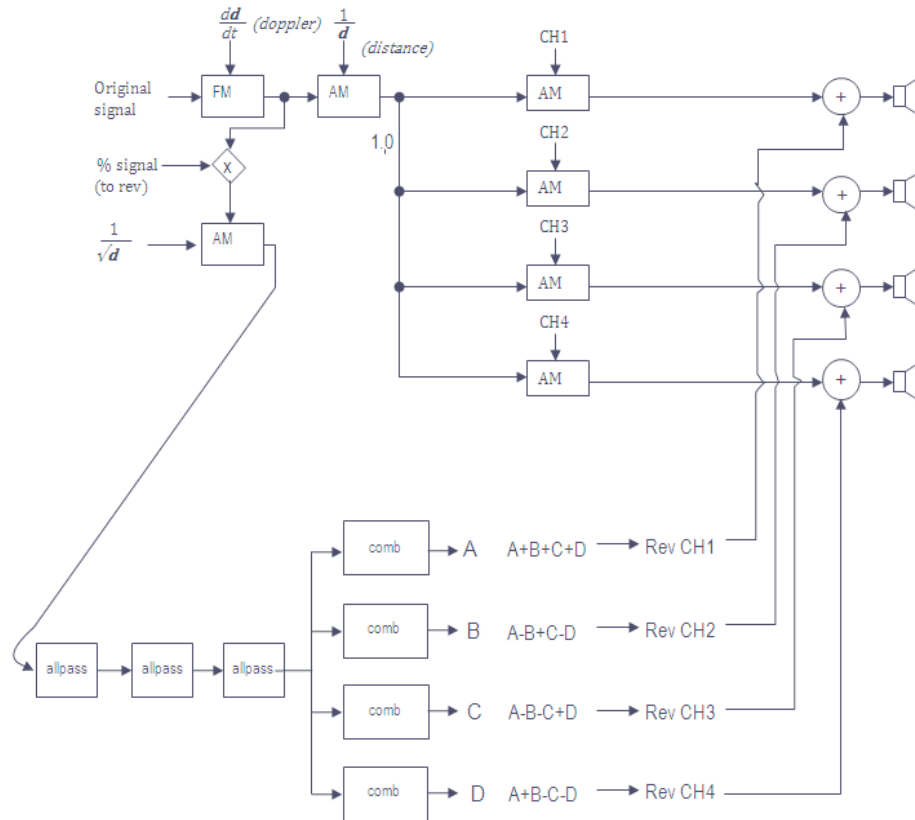

Fig.2. The sound spatialization system used in *Sabelithe* (1971) and *Turenas* (1972). The "allpass" are colorless and in series, increase the echo density. The "comb" generators summed to produce moderate (and desirable) coloring of the spectrum. All had prime number delays in samples according to Schroeder's recommendation

musical dynamic, in which case the reverberation decreases in the same proportion as does the direct signal. There is a point where the overall intensity of the reverberation is equal to the intensity of the direct signal. Beyond this point, the echo radius,[9] the direct signal is increasingly masked and thus unable to provide cues for either distance or azimuth.

In order to lessen this effect I conceived a system that localizes the reverberant field with the increase in distance of the direct signal, such that with the direct signal at infinity the reverberation is localized with the direct signal and with the direct signal maximally proximate the reverberation is global, i.e. equally distributed to all channels. The intention of this additional processing, seen in Fig. 3, was to increase the ability of the listener to perceive the distance and azimuth of the source when it is projected beyond the echo radius. Because this system required an independently

controlled reverberation unit for each of the four channels, it was never implemented in software. I was unwilling to trade high quality reverberation in all channels for sparse (computationally affordable), controllable reverberation in each.
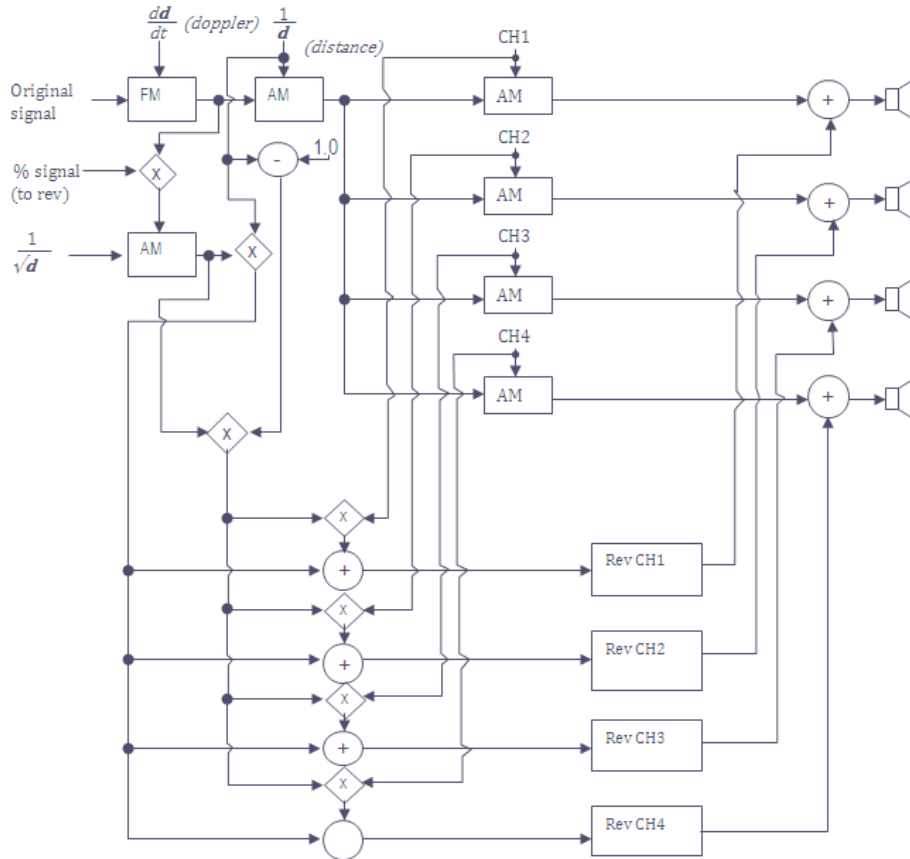
Fig. 3. Shown here is the processing scheme that provides for localization of the reverberation with increasing distance of the direct signal. Each RevCHn is comprised of allpass and combs as shown in Fig. 2. It is this system that was implemented by GRT Corp. although with spring reverb units

    In fact, there are listening conditions in large reverberant spaces where the directionality of reverberation may be strong. For a listener near the center of a large space, the direct signal of a sound source that is close to a distant wall may have a sound path only slightly longer than some of the early reflections.

    I demonstrated my work to Stanford's Office of Technology Licensing (OTL), which applied for and was granted a patent.[1] A real-time prototype of this system was built by General Recorded Tape Corporation (GRT), which licensed the technology from Stanford. The prototype had a joystick, the radial motion of which con-

---

[1] METHOD AND APPARATUS FOR SIMULATING LOCATION AND MOVEMENT OF SOUND, US Pat. 3665105 - Filed Mar 9, 1970 Patented May 23, 1972

trolled a variable speed single channel tape player. A recorded sound source was processed based on the motion of the joystick, with angular motion positioning the source between loudspeaker pairs and radial motion controlling the Doppler shift and distance. Four spring reverberators were used and while not high quality, they were nonetheless effective within the system. The real-time system was a spectacular demonstration, but there were no applications, even on the horizon — cinema theater surround-sound was still decades in the future. GRT lost interest, but not before paying Stanford enough money so that we could purchase a professional quality 4-channel, 1/2in tape recorder. I wrote a paper that describes the PDP-10 based computer system that was published in 1971. [2]

3 Frequency Modulation (FM) Synthesis: a fortuitous diversion

It was nearly forty-five years ago that I 'stumbled" upon FM synthesis while searching for sounds having dynamic attributes that would be suitable sources in my sound localization experiments. The actual date is not known. Not having a scientific or engineering background, I did not have the habit of keeping dated lab notes, but I did keep notes. There is a record of my having visited BTL on December 18, 1967 when I showed the data that I used in my first trials to Max Mathews, Jean-Claude Risset and Pierre Ruiz and played for them the examples. Ruiz copied the processing diagram and Risset carefully copied *and dated* the data from my notebook. It was a few weeks before, almost certainly late at night, while experimenting with extreme vibrato frequencies and depths that I realized that I was not hearing changing frequency (pitch) in the time domain, but rather complex tones, both harmonic and inharmonic, in the frequency domain. Furthermore, the spectrum could be modulated from a pure sinusoid to a complex wave by means of a simple linear function. Again, a few years later Stanford's OTL applied for and was granted a patent.[2]

The discovery of FM synthesis was not a purposeful search—that is, stemming from a realization, from looking at the equation, that there might be some interesting experiments to try—rather, it was altogether a discovery of the "ear."

One must remember that while the theoretical potential for the production of rich dynamic sounds with the computer was great, the knowledge required for realizing this potential was meagre and the means (computer time) expensive. During the exchange of FM data mentioned above, Risset explained to me his analysis and resynthesis of trumpet tones and what he had learned as he applied this knowledge to the synthesis of other timbres. I was struck by the elegance of his synthesis. However, the cost in computer time was large, placing practical limits on the complexity of synthesis algorithms. Deep into the details of digital reverberation at the time, I was keenly aware of the computer time issue.

My "ear" was continually scanning for any sound having internal dynamism, coupled oscillators, vibrato, etc. That I found it within such a computationally efficient

---

[2] METHOD OF SYNTHESIZING A MUSICAL SOUND, US Pat. 4018121- Filed May2, 1975 Patented April 19, 1977

algorithm was certainly partly chance, but then I was also certainly prepared for that chance.

The first experiments were each only a few seconds duration, because of the tens of minutes of compute time on a time-shared system. But they do show that from the outset, all of the essential features were noted that would eventually be developed and used in musical contexts:

- both harmonic and inharmonic spectra could be produced
- a change in frequency deviation (Δf) produced a change in bandwidth of the spectrum
- the spectrum is conserved through the pitch space with a constant ratio of FM frequencies

As it turned out, these parameters of FM synthesis have a remarkable perceptual relevance.

As mentioned above, Risset's study of trumpet tones [7] had a major influence on my own development of FM synthesis. It was not until 1970, however, that I fully appreciated the importance of his discoveries about trumpet tones.

While working on the FM synthesis of percussive sounds, I noted that in nearly all tones of this class the amplitude envelope and the envelope controlling the modulation index were very similar if not identical. I also noted that there was as strong a correlation of the perception of 'strike force' to the modulation index as there was to intensity. I considered other classes of tones where this might be the case, and I remembered Risset's explanation of the "signature" of trumpet tones, some three years previous. With only a few attempts I was able to create credible brass-like tones by simply coupling a single function to the amplitude and modulation index envelopes with appropriate scaling. I realized that this correlation of force or effort (strike force, breath and bow pressure velocity, etc.) to the bandwidth and/or high-frequency emphasis of partials can be generalized to all natural sound and that the parameters of FM synthesis provided a straightforward implementation of this important correlation[3].

Then began a rapid development of FM synthesis[4], and the eventual licensing of the technology by Stanford University to Yamaha. I published a paper describing this work in 1973 [3].

4 Sabelithe: Caccia chasing Escher

While developing the spatialization and FM synthesis systems, the compositional possibilities and limitation of their use were evolving as well. It was clear to me, for example, that projected spatial trajectories could be quite complex but the auditory system could not easily track simultaneous trajectories — simple of complex. It was

---

[3] The ease with which spectral change could be coupled to effort (key velocity) is one of the reasons for the YAMAHA DX7's remarkable success.

[4] The first real-time FM synthesis was programmed on a DEC PDP-15 computer by Barry Truax in 1973, while studying in Utrecht. At Stanford, Bill Schottstaedt developed a particularly powerful form of the algorithm that was used in many compositions for many years.

also clear that sudden angular changes in direction within a trajectory were awkward sounding while curvilinear paths were agreeable and with the Doppler shift, natural sounding within a magical un-natural imagined context. I often thought of Brancusi's *Bird in Space*, which I had seen several times while studying in Paris, perhaps at the Musée d'Art Moderne in those years. I imagined the curves that might be drawn if the elegant lines of the object itself were extended into space according to their implied curves.

FM synthesis was capable of producing a large number of differentiated timbres with only small changes in a few numbers or control functions. This attribute suggested continuous changes in timbre or metamorphoses, for which there were compelling graphic analogs especially in the designs by M. C. Escher.

In 1971 I wrote my fist composition using computers, *Sabelithe*, that I thought of as test case for some of these ideas. In addition to spiral trajectories I included a section where short percussive sounds followed a path —left rear-to-front, front-across, right front-to-rear — that was then followed by a second and then a third instance of the same. It was a spatial canon, in fact, for which there is a musical analog, the 14th century Caccia (chase or hunt), a precursor of the fugue. There is some irony in the fact that a 20th century technology would produce a form for which there is a 600-year-old metaphor. The brilliant Stanford music historian, Imogene Horsley, found this fact amusing and somewhere mentions *Sabelithe* in her scholarly writings.

At various points in the composition there are timbral transformations, but the most dramatic is near the end where a short, noisy, discontinuous percussive sound is gradually transformed into a Risset inspired brass tone. The carrier/modulation frequency ratio begins at 1:1.414, producing an inharmonic spectrum and a large index of modulation causes aliasing and thus noise. With each repetition of the element, the ratio converges toward 1:1 and the index reduces to one appropriate for a brass-like tone. It is a compelling example and demonstrates the value of having a large timbral space controlled within a small parameter space. *Sabelithe* was first presented at Dinklespiel auditorium in May 1971.

The concept of FM timbral transformations would be exploited in several pieces of the future, including Turenas and Pnoné, the latter of which is based upon Risset's remarkable example from *Mutations* (1969) where timbre, harmony and melody all derive from common pitch material, what I have referred to as structured spectra [1].

But to realize my dream of sound in space, I would have to first discover how to produce complex elegant curves.


5 Turenas: Doppler and Lissajous

Even before finishing *Sabelithe*, I had begun experimenting with a drafting arm that had sensors on the articulation points such that the cursor on a display would track the position of a mouse-like device at the end of the arm on a table. I wrote a program that plotted points at a constant rate as I moved the cursor through the two-dimensional space. Measuring the length of successive radii relative to a center point (ideal listener position) gave me the Doppler shift in their differences. The azimuth was easily derived from the angular position of the points. Using this mouse-like

control was much like the previously described real-time prototype built by GRT. But there was no sound. I quickly tired of this method of input, but not before D. Poole remarked that one of my drawn patterns looked like a Lissajous figure. My curiosity was sparked by his comment leading me to learn about and program Lissajous figures.

I quickly advanced through the well-known looping patterns and discovered that interesting figures could be generated, as shown in Fig. 4 and Eq.1, whose sound manifestation possessed a graceful motion that seemed to me natural — as the sound followed the path of a Lissajous figure it decelerated as it approached a change in direction.

Fig. 4. The iconic Lisajous figure of *Turenas* showing the amplitude and frequency control functions for distance and Doppler shift respectively

$$x = \sin 2\pi t + \sin 6\pi t$$
$$y = \cos 3\pi t + \cos 7\pi t$$

(1)

The coordination of the Doppler shift and correlated change in amplitude as a function of distance produces a convincing auditory illusion that may draw up images of accelerating and decelerating mass in motion. Perhaps in the real world the closest analog might be a figure skater on ice.

6 Conclusion

It has been nearly forty years since I finished Turenas[5] and eight years before that I began the search for a means of liberating sound from its loudspeaker source. The composition has been carefully analyzed [6, 10] and even reconstructed [11].

The dream was to be able to compose sound in space that was free of physical constraints and realities, yet would evoke images that were believable. However, had I been given the most advanced computer system in 1964—even by the standards of today, I could not have realized my dream. I had to pass through the years of discovery and learning, first mentored and guided by Max Mathews and David Poole in matters technical, then enlightened and inspired by Jean-Claude Risset. FM synthesis, which is integral to Turenas, was a gift from *nature* along the way—and hence the title from its possessive form.

References

1. Chowning, J.: Composing the Sound Itself, *John Chowning Portraits Polychrome*,. Nouvelle edition augmentée. Institute national de l'audiovisuel, Paris (2007)
2. Chowning, J.: The Simulation of Moving Sound Sources, J. Audio Eng. Soc. 19, 2-6, (1971)
3. Chowning, J.: The Synthesis of Complex Audio Spectra by Means of Frequency Modulation. J. Audio Eng. Soc. 21, (1973) Reprinted in Computer Music Journal 1 (1973)(2)
4. Mathews, M. V.: The Digital Computer as a Musical Instrument, *Science*, Vol. 142 No. 3592, (1963) 553-557
5. Moylan, W.: The Art of Recording: Understanding and Crafting the Mix**,** Elsevier Science, (2002) 52
6. Pottier, L.: Analyse de Turenas, *John Chowning Portraits Polychrome*,. Nouvelle edition augmentée. Institute national de l'audiovisuel, Paris (2007)
7. Risset, J-C.:Computer Study of Trumpet Tones, J. Acoust. Soc. Am. Volume 38, Issue 5, (1965) 912-912
8. Schroeder, M. R.: Natural-sounding artificial reverberation, *Journal of the Audio Engineering Society*, vol. 10, no. 3, (1962) 219-223
9. Wendt, K.. "The Transmission of Room Information," J. Audio Eng. Soc. 9, (1961)

URLs
10. Jure, L.: Escuchando Turenas de John Chowning, Músicas al Sur, n°1, Escuela Universitaria de Musica, Montevi-deo, Uruguay, janvier 2004 (http://www.eumus.edu.uy/revista/nro1/index.htm).
11. Pottier, L.: Turenas- vers une version interactive http://www.musimediane.com/numero6/POTTIER/index.html

---

[5] *Turenas* was first performed in Dinklespiel auditorium in April 1972. Present in the audience was Martin Bresnick, Gyorgy Ligeti, Loren Rush and Ivan Tcherepnin.