PV STOCH: A SPECTRAL STOCHASTIC SYNTHESIS GENERATOR

Luc Döbereiner
Institute of Sonology, The Hague
luc@doebereiner.org

ABSTRACT

PV Stoch is a phase vocoder (PV) unit generator (UGen) for SuperCollider. Its objective is the exploration of methods used in "non-standard synthesis", especially in Dynamic Stochastic Synthesis (Xenakis), in another domain. In contrast to their original conception, the methods are applied in the frequency domain. This paper discusses some of the compositional motivations and considerations behind the approach, it gives a description of the actual synthesis method and its implementation, as well as a summary of the results and conclusions drawn.

1 INTRODUCTION

PV Stoch is a generator for frequency domain stochastic synthesis. After having worked at the generalization of "non-standard" synthesis[4], the development of *PV Stoch* was driven by an interest in extending stochastic synthesis; an interest in testing the transferability of its principle workings and reapply them in another area, the frequency domain. In this paper, we will discuss the first result of this investigation.

1.1 Transferability

Iannis Xenakis used stochastic functions for the generation of sound after having used them on a higher-level before. They have been compositional tools to him. The step to synthesize the sounds themselves using probabilities, as well as the introduction of them in musical composition itself, follow the belief that a method which has successfully been employed on one level or one domain may successfully be transferred to another.

> Any theory or solution given on one level can be assigned to the solution of problems of another level. Thus the solutions in macrocomposition (programmed stochastic mechanisms) can engender simpler and more powerful new perspectives in the shaping of microsounds.

SMC 2009, July 23-25, Porto, Portugal

1.2 Overview

Firstly, the synthesis method itself is described. The individual parameters are presented, as well as brief description of their aural effects. It shall be noted that the descriptions are somewhat simplified and most of all, the control parametrization is not congruent with their multi-layered perception. Although, the perceptible effects of each of the parameters is briefly addressed, their inter-dependencies and trans-active nature is far too complex to be properly outlined here.

Subsequently, we will give attention to *PV Stoch*'s relation to the "non-standard" synthesis approaches and thereby place it in a historical and theoretical context. Although *PV Stoch* does not fulfill all the criteria to be classified as "non-standard", we are trying to demonstrate that it does indeed comply with and even extend some fundamental notions present in these approaches.

Furthermore, some of the challenges and features we have encountered in the practical work with the generator are discussed by means of the description of a 96-channel composition by the author which was realized exclusively with *PV Stoch*.

2 IMPLEMENTATION

PV Stoch is a phase vocoder UGen for SuperCollider (J. McCartney). SuperCollider features a robust and efficient framework for the design of frequency domain operators. As the development of *PV Stoch* has been a rather experimental investigation, SuperCollider's flexibility and real-time controllability proved to be crucial. The implementation framework is straight forward and the UGen can be combined with a variety of already existing UGens and control mechanisms.

PV Stoch takes the following parameters, which are explained below. Except of **nBps** and **lambda**, which only have effect during the initialization, all parameters are dynamically controllable:

2.1 Basic Functionality

PV Stoch is a frequency domain stochastic synthesis generator. Although, it operates on a FFT buffer, it does not process an analyzed sound, but rather synthesizes sound with-

```
PV_Stoch(buffer, nBps, lambda,
         phaseSwitch, specDec,
         interpBase, range,
         offset, deviation)
```

Figure 1. The parameters of *PV Stoch*

out input source. [1] The created spectra have an envelope, or spectral contour, which is constructed of interpolated breakpoints. The positions of these breakpoints deviates from frame to frame, as the time domain breakpoints deviate from cycle to cycle in Xenakis's Dynamic Stochastic Synthesis.

When the UGen is initialized, it generates an initial spectral envelope. The distribution of the breakpoints follows controllable probabilistic laws – an exponential random distribution – and the interpolation function may vary over time. Frame by frame the positions of the breakpoints, and thereby the spectral shape, deviate. The amount of deviation is dynamically controllable. The created spectrum can also be dynamically frequency shifted or stretched, which are familiar frequency domain techniques. Furthermore, the phase spectrum generation has three states and it can be interpolated between them.

2.2 The Envelope

The initial envelope has a big effect on the resulting sound. Initially, its shape is determined by three parameters: the number of breakpoints (**nBps**), a random variable controlling the spread of an exponential random distribution which determines the horizontal (frequency) position of the breakpoints (**lambda**), and the base of the interpolation function (**interpBase**). If the base is 1, the interpolation is linear, if it is bigger or smaller than 1, the interpolation is exponential, resulting in concave and convex curves respectively.

A higher number of breakpoints (**nBps** > 20) results in more defined and more complex spectra, a lower number creates sounds similar to more simply filtered noise. If **lambda** is smaller (**lambda** $<= 1.0$), the resulting sounds are more distinguishable and the deviations are clearer, if the random variable is greater, the sound becomes more static, the changes less drastic. A more concave interpolation curve (**interpBase** < 1.0) articulates the attenuated frequency regions more clearly, whereas more convex shapes create blurrier noise regions.

The vertical (amplitude) positions of the breakpoints are determined by a beta random distribution. Additionally, the magnitudes of the whole spectrum are also scaled by an exponentially decreasing shape, whose steepness is variable (**specDec**).

[1] One exception is the phase, which is explained in 2.5.

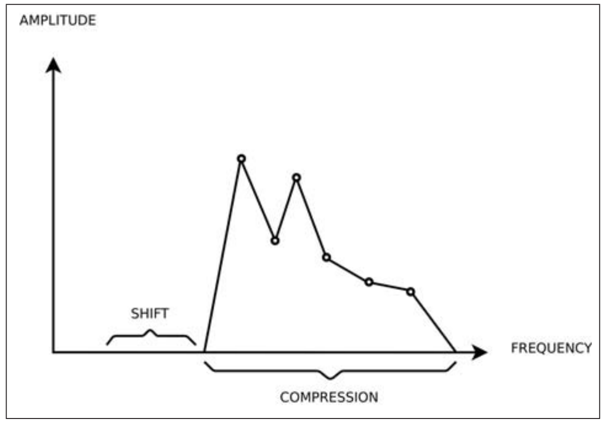

Figure 2. Shifting and compressing the created spectrum

2.3 Shifting and Stretching

Figure 2 illustrates two additional – and well known – operations which can have a drastic effect on the sound: frequency shifting and stretching/compressing. As can be seen, the entire spectral shape can be shifted (**offset**) along the frequency axis (in both directions) and stretched or compressed (**range**). Since the spectral shape is expressed by interpolated breakpoints whose position along the frequency axis does not need to coincide with the frequency grid imposed by the frame size, shifting and stretching or compressing occurs smoothly without making the frequency resolution audible. The shifting and stretching is similar to techniques presented by among others Trevor Wishart[8]. Although, Wishart's approach is regarded as "standard" synthesis, it is surely a compositionally motivated approach to sound synthesis.

2.4 Deviation

Figure 3 shows the deviation principle. The breakpoints deviate frame per frame from their previous position by a random amount, the maximum of which is controlled by the parameter **deviation**. Thus, similar to Dynamic Stochastic Synthesis', the breakpoints undergo random walks, however, only in their vertical position (amplitude).

2.5 Dealing with the Phase

There are three basic settings for the phase. It can be interpolated between them. The phases can be set zero, in which case the results are closer to additive synthesis using sine waves, the phase values can be generated randomly, which creates sounds closer to filtered noise, or they can be derived from an input source. Although *PV Stoch* has not been designed to process analyzed input sources, this phase setting was introduced in order to add "articulation" stemming

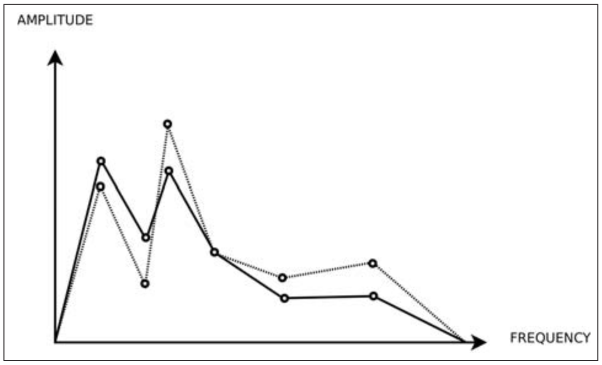

Figure 3. The deviation of breakpoints in two successive frames

from another source. It was primarily used with frequency modulated impulses.

3 *PV STOCH* AS "NON-STANDARD" SYNTHESIS

The starting points for the development of *PV Stoch* have been the so-called "non-standard" sound synthesis[6] approaches, especially Iannis Xenakis's Dynamic Stochastic Synthesis. The systems subsumed under the term "non-standard" have in common that they do not adhere to any superordinate acoustic models. [2] Instead, the models of sound are derived from compositional models. Sound synthesis is understood as the development of processes organizing the low-level units, as "microtemporal *compositional* processes."[3] *PV Stoch* takes up this idea of deriving higher level structural properties from the description of lower level processes. Here, the distinction between sound and music is blurred.

For different reasons, however, the "non-standard" approaches rejected the frequency domain. Xenakis heavily criticized the use of harmonic analysis for the synthesis of sound. The results he deemed uninteresting, the approach "inadequate". He ascribed the problems to the "synthesis by finite juxtaposed elements"-principle. "It is as though we wanted to express a sinuous mountain silhouette by using portions of circles"[9], he writes. Curiously, Xenakis's UPIC system is based on the very principle he had been criticizing so vehemently, it is a form of additive synthesis.

Perhaps, due to its mathematical nature and popularity among the more simulating sound synthesis methods, the frequency domain was considered inappropriate for a uniquely digital music. It seemed to be a concept which was not very well suited for answering the question, "what means of expression are idiomatic to computers?"[7] For Xenakis, the reason for his rejection may rather have been his association of additive synthesis with the electronic music of the Cologne studio.

In contrast to the lion share of the research done in sound synthesis, the "non-standard" approaches are truly experimental. The interest does not lie in "trying to reconstruct a sound based on analytic data", but in "composing sound using musical procedures."[1] They can be seen as explorations of compositional representations of sound. In Koenig's SSP, amplitude and time values are elevated to the level of musical unit elements. Surely, problems arise from that, because their treatment "would require parameters to have a recognizable identity"[2] *PV Stoch* continues to ask the question "what is the minimum of logical constraints necessary for the construction of a musical process"[9], but it changes the underlying form of representation, also in hope of creating elements with a more recognizable identity.

In fact, the so-called "non-standard" sound synthesis approaches are all characterized by the use of concepts which are initially alien to the description of sound. With SSP, for example, G.M. Koenig uses methods which he had developed for instrumental composition for the structuring of audio sample values. Similarly, Paul Berg's programs ASP and PILE derive musical and sonic relationships from instructions present in programs for numerical computation.

There is, thus, an element of transfer, of reapplication, in "non-standard" synthesis. The sound organizing principles arise from a compositional interest, the compositional idea is embodied in the 'sound material', it is not imposed on it.

In this line of thought, *PV Stoch* can be seen as an attempt at creating frequency domain "non-standard" synthesis. Although, this may stand very much in contrast to the rejection of superordinate acoustic models, it follows Xenakis's idea of transferability of concepts.

Instead of aiming at the (re)creation of specific sounds, it is rather a search for the remains of an organization principle, for the traces the prinicple may leave in the sound and through another representation.

4 AN APPLICATION: *SPACE STUDY 1*

Space Study 1: Order From Noise is a fixed medium (tape) piece for 96 independent channels which was composed by the author in 2009. For the sound production *PV Stoch* was used exclusively. Due to the immense amount of data and coordination necessary for the independent composition of 96 tracks, it became unavoidable to automate many processes in the production of the piece. A consequence of the automation was the necessity of clear distinctions, of parametric configurations on the one hand and strategies of transitions and transformations of the other hand.

The piece consists of four sections which undergo a similar macro-level development, there can be seen as variants of a common higher-level description. For the most part,

[2] Oddly, the time domain is usually not considered an acoustic model in the descriptions of "non-standard" synthesis.

the synthesis settings are the variable element among the sections. The four sections are briefly described:

1. Impulses whose frequencies follow exponential curves, and ranging from 1 to 100 Hz thus creating rhythm and pitch, serve as the input for the phase values. *PV Stoch* initially creates ”resonances” and gradually the phases become more random, the impulses are thus replaced by noise and the ”resonances” become the central sound itself.

2. The phases alternate between noise, impulses, and zero, thus creating clearly distinguishable types of events. Instead of gradual change and slow transitions, the different timbres are clearly opposed to each other.

3. Blocks of quickly deviating bursts form gestural units. The deviation is high, lambda is low.

4. Finally, the phases are set to zero. The section is rather soft in volume and the spectra act as clusters, slowly shifted in frequency and space.

Each of the sections creates its own timbre space. When the phases are derived from impulses, the generator creates ”resonances”, when it is random, the deviation is high, and lambda is low, the random walks are most audible and the output strongly resembles time-domain stochastic synthesis. Since, all the timbre states are outcomes of the same process, they can easily be related to each other.

5 FURTHER WORK

Peter Hoffmann writes about Xenakis’s GENDYN:

> The key idea of stochastic synthesis is its non-linear waveshaping, where the waveshaping function changes stochastically from period to period. Consequently, it is not the waveform as such that defines the aural result [...] but rather the dynamic behavior of its deformation over time. [5]

PV Stoch behaves similarly. It is not the specific spectrum created but rather the way it changes from frame to frame that determines the aural quality of the result. The behavior is also what is most controllable. Since the initial envelope has a big effect on the resulting sound and since it is not completely predictable from the parameter settings, several instantiations of the same parametric configurations can result in a great variety of different sounds. The generator is thus not very well suited for the purposeful creation (simulation) of pre-conceived sounds. By controlling deviations, “spectral definition”, pitch and noisiness, types of sounds and types of sonic behaviors can be created.

Several improvements and additions suggest themselves and need to be tested regarding their musical effectiveness. The deviation may be further refined. Since the dynamic behavior of the system is the perceptibly most significant element, it should be further developed.Similarly to Xenakis’s models, second order random walks could be included and the breakpoints could move on the frequency axis as well. Furthermore, the number of breakpoints should be dynamically variable. The impact of different random distributions on the various stochastic processes should be investigated.

6 REFERENCES

[1] Berg, P., Rowe, R., Theriault, D. “SSP and Sound Description”, *Computer Music Journal*, Vol. 4, No. 1, 1979.

[2] Berg, P. “Composing Sound Structures with Rules”, *Contemporary Music Review*, Vol. 28, No. 1, 2009.

[3] Di Scipio, A., Prignano, I. “Synthesis by Functional Iteration. A Revitalization of Nonstandard Synthesis”, *Journal of New Music Research*, Vol. 25, 1996.

[4] Döbereiner, L. “CompScheme: A Language for Composition and Stochastic Synthesis”, *Proceedings of the 5th Sound and Music Computing Conference*, Berlin, Germany, 2008.

[5] Hoffmann, P. “The New GENDYN Program”, *Computer Music Journal*, Vol. 24, No. 2, 2000.

[6] Holtzman, S. R. “An Automated Digital Sound Synthesis Instrument”, *Computer Music Journal*, Vol. 3, No. 2, 1979.

[7] Holtzman, S. R. *Digital Mantras*, MIT Press, Cambridge, 1994.

[8] Wishart, T. *Audible Design*. OTP Ltd, 1994.

[9] Xenakis, I. *Formalized Music*. (Rev. ed). Pendragon Press, New York, 1992.