HAL
open science

Synchronous Programming in Audio Processing: A Lookup Table Oscillator Case Study

Karim Barkati, Pierre Jouvelot

To cite this version:

Karim Barkati, Pierre Jouvelot. Synchronous Programming in Audio Processing: A Lookup Table Oscillator Case Study. International Workshop on Synchronous Programming (SYNCHRON 2011), Nov 2011, Dammarie-les-Lys, France. 10.1145/2543581.2543591 . hal-00751490v2

HAL Id: hal-00751490

https://minesparis-psl.hal.science/hal-00751490v2

Submitted on 31 Jan 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A

Synchronous Programming in Audio Processing: A Lookup Table Oscillator Case Study

KARIM BARKATI and PIERRE JOUVELOT, CRI, Mathématiques et systèmes, MINES ParisTech, France

The adequacy of a programming language to a given software project or application domain is often considered a key factor of success in software development and engineering, even though little theoretical or practical information is readily available to help make an informed decision. In this paper, we address a particular version of this issue by comparing the adequacy of general-purpose synchronous programming languages to more domain-specific languages (DSL) in the field of computer music. More precisely, we implemented and tested the same lookup table oscillator example program, one of the most classical algorithms for sound synthesis, using a selection of significant synchronous programming languages, half of which designed as specific music languages – Csound, Pure Data, SuperCollider, ChucK, Faust – and the other half being general synchronous formalisms – Signal, Lustre, Esterel, Lucid Synchrone and C with the OpenMP Stream Extension (Matlab/Octave is used for the initial specification). The advantages of these two approaches are discussed, providing insights to language designers and possibly software developers of both communities regarding programming languages design for the audio domain.

Categories and Subject Descriptors: A.1 [**General Literature**]: Introductory and Survey; C.3 [**Special-purpose and Application-based Systems**]: Real-time and Embedded Systems; Signal Processing Systems; D.1.3 [**Programming Techniques**]: Concurrent Programming—*Parallel programming*; D.1.m [**Programming Techniques**]: Miscellaneous; D.2.11 [**Software Engineering**]: Software Architectures; D.3.2 [**Language Classifications**]: Concurrent, Distributed, and Parallel Languages; Data-flow Languages; Specialized Application Languages; Very High-Level Languages; D.3.3 [**Programming Languages**]: Language Constructs and Features; E.1 [**Data Structures**]: Arrays; H.5.5 [**Information Interfaces and Presentation**]: Sound and Music Computing—*Signal analysis, synthesis, and processing*; *Systems*; J.5 [**Arts and Humanities**]: Performing Arts; J.7 [**Computers in Other Systems**]: Real Time; K.2 [**Computing Milieux**]: History of Computing

General Terms: Design, Languages

Additional Key Words and Phrases: Synchronous programming languages, Music programming languages, Computer music, Signal processing, Timing

ACM Reference Format:
Karim Barkati and Pierre Jouvelot, 2012. Synchronous programming in audio processing: a lookup table oscillator case study. *ACM Comput. Surv.* V, N, Article A (YYYY), 35 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

The understanding of the existence of a close relationship between music and mathematics has been mentioned since the ancient Greeks. This is therefore not surprising that programming language designers have considered the musical domain as a venue of choice for their investigations from the early days of the field of computing

This work was supported by the French National Research Agency, under grant ANR 2008 CORD 003 01.
Authors' addresses: Karim Barkati, IRCAM, France; Pierre Jouvelot, CRI – MINES ParisTech, France.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© YYYY ACM 0360-0300/YYYY/-ARTA $15.00
DOI:http://dx.doi.org/10.1145/0000000.0000000

[Van Roy 2009]: programming is indeed a constructive approach to mathematical reasoning. There are of course multiple ways to use programming languages in music applications, from low-level audio processing to more abstract music notation manipulation processes to the higher sphere of music composition frameworks.

Current digital audio applications are often based on the algorithmic real-time processing of streams of sound samples; these signals are subject to strong timing requirements since latencies and delays in music signals are easily detectable even by untrained human ears. Such constraints call for programming languages that are able to manage somewhat significant data throughputs while enforcing timing properties. Even though traditional programming languages offer means of addressing timing issues within programs, e.g., via the `sleep` system call in C or `clock` in C++, they do not consider the notions of time and data streams as the foundation for their design paradigm, syntax and semantics. The goal of synchronous and music languages is in fact to hide low-level implementation details to the programmer as much as possible, narrowing the semantic gap between specification and implementation. For instance, in all music languages, one only needs to specify what processing actions are requested for each audio sample. All the management of low-level issues such as system calls, buffers, callbacks or audio frames is performed via the compiler and run-time system.

Moreover, most audio processing applications happen to be modeled as large sets of concurrent interacting dataflow processes. The specification of the behaviors and synchronization patterns of such computations would be unwieldy at best using traditional all-purpose languages, in which difficult-to-track "temporal bugs" [Simon and Girault 2001], such as race conditions or deadlocks, often occur. Traditional languages seldom provide strong formal guarantees regarding timing constraints for concurrent processes operating over sampled input and output streams, as opposed to synchronous and music languages: correctness, ease of programming and other benefits are what such specific languages provide. Finally, the growing importance of the practice of "livecoding" [Nilson 2007], where computer musicians perform on stage by programming and running code snippets on the fly, calls also for lightweight and dedicated programming formalisms.

A family of programming languages happens to have been developed to specifically deal with timing issues, in particular for embedded and time-critical systems: synchronous languages [Benveniste and Berry 1991b; Benveniste et al. 2003]. This particular programming paradigm is based on the key idea of synchronizing concurrent computing processes on clocks, while ensuring the mathematical correctness of the resulting programs, obviously needed given the target domains. The use of clocks corresponds to a discretization of the time domain, and thus of the number of states a programmer or an automatic correctness prover might have to consider when trying to ensure the correctness of source code. Traditional programming models adopt what amounts to a continuous time domain vision, more prone to the introduction of subtle run-time errors.

Independently, the computer music community has been working on music-specific languages, initially centered on designing sound generators, seen as virtual instruments; the issues of time and processes were also addressed there, but in less formal ways and more as an afterthought in the language design process. The relationship of synchronous languages with the temporal structure of music and real-time audio computing appears thus to be manifest and justifies the current interest of the music community for the synchronous paradigm.

The purpose of this paper is to provide a comparative survey of the existing portfolio of major synchronous programming languages that can be used in the specific field of audio processing. We believe such an analysis is particularly pertinent today, since this domain has seen a recent significant growth both in the industrial and research

worlds, with the widespread use of gadgets such as MP3 players or the introduction of new programming languages such as Faust [Orlarey et al. 2009] or ChucK [Wang et al. 2003]. This survey adopts a two-pronged approach: we look both at key music programming languages, i.e., audio-specific languages designed in the computer music community, and general signal-processing programming languages to see how they relate to each other while adopting, with some variations, the synchronous trait.

Our comparison work focuses on design issues. Both synchronous and computer music programming languages strive to help lower the semantic gap between informal specifications and their formal program development. Even though other facets of programming languages such as their performance, implementation details, portability or availability might justify other comparison studies, we believe that design choices are *in fine* the key drivers in the acceptability, success and ultimately sustainability of programming languages; they deserve an in-depth analysis of their own. We adopt a pragmatic approach when looking at these design issues: we use the same running example throughout the article, namely the implementation of a lookup table oscillator. For each implementation, we stick as much as possible to the programming style of the corresponding language in order to highlight its idiosyncratic design aspects.

We believe our work can be of help to language designers interested in the handling of time and/or audio in programming languages, since, in particular for DSL[1] applications, cross-fertilization between languages via the borrowing of existing language features is a common way to improve language designs. Moreover, software developers and project managers who need to decide which language is best adapted to a given digital audio project may also find our survey of interest, in particular the audio language part. Indeed, the adequacy of a programming language to a project is often considered a key factor of success in software development, even though little theoretical or practical information is readily available to help make an informed decision. Even though all the programs in this paper have been tested and run, we do not expect our readers, whom we assume have already some working knowledge of current programming languages, to understand all the gory details of these implementations. We hope though that they will get from our work (1) a feel for the design principles of each language, (2) a new perspective on how wide the spectrum of design choices is when trying to deal with timing issues, and possibly (3) a keen interest in delving further into one or the other of these key representative languages.

The structure of the paper is the following. Section 2 provides a brief general overview of the most prominent music and synchronous programming languages. Section 3 contains a description of our running example, which we consider a typical use case for most current audio applications; we provide both an informal and a formal, in Matlab/Octave, specification of our target program, `osc`, a standard wave synthesis algorithm. The two following sections, namely Section 4 for music-specific languages and Section 5 for synchronous languages, present our attempts at implementing `osc` in a few selected languages. For each one of these languages, we use the same presentation format: (1) a brief overview introduces the language – using the own words of its author(s) –, (2) the `osc` example we coded in this particular formalism and (3) a set of notes, when we felt that some explanations were needed for what we assumed would be the most difficult programming details to understand. Section 6 puts our different implementations in perspective, highlighting the specificities of the various approaches and the opportunities for mutual interactions between synchronous and computer music language designs. We conclude and suggest future work in Section 7.

[1]Domain-Specific Language.

2. THE SYNCHRONOUS PARADIGM AND COMPUTER MUSIC

From theoretical time-complexity issues to user interaction management, the concept of time has a structuring effect on the way computer technology impacts the programming world. In languages specifically dedicated to music programing, music events are expressed along strict timing constraints. Synchronous general-purpose languages also consider (logical) time, along which control and computation are scheduled, as a key design ingredient; they adhere to the "synchronous hypothesis", which emphasizes time constraints and determinism. Even though developed within a totally different research community, music-specific languages also follow this synchronous hypothesis. We survey below these two approaches, and end this section with the list of key representative languages we use in the remaining of this paper.

2.1. Computer Music Languages

Even though digital music sequencing tools such as, nowadays, Steinberg Cubase or Apple Logic provide friendly ways to edit and synchronize music sequences on a fixed timeline, more refined timing constraints and higher-level abstraction mechanisms have been asked for by artists over the years, calling for programming mechanisms that go way beyond those offered in sequencers. These issues, among others, are addressed by the vibrant and dynamic research field of computer music, one of the oldest application domains of computers. Started in the late 1950s on mainframes, computer music has two main branches [Loy and Abbott 1985]: computer-aided composition (CAC) and digital audio processing. For this study, we are interested in the latter.

The CAC branch, which usually aims at producing scores, deals mostly with the *note* paradigm, using a symbolic approach where the characteristics of notes (e.g., their pitch and time location) and not their actual sound are encoded. This programming path started in 1956 with the MUSICOMP[2] language of Lejaren Hiller and Robert Baker (both chemists at that time) at the University of Illinois using an Illiac I [Baker and Hiller 1963; Hiller and Baker 1964]. Main CAC languages include PatchWork [Laurson and Duthen 1989], Common Music [Taube 1991], Haskore [Hudak et al. 1996], Elody [Orlarey et al. 1997], OpenMusic [Assayag et al. 1999] and PWGL [Laurson et al. 2009].

The digital audio processing branch is mostly focused on the *sound* paradigm, using signal processing and physical modeling approaches; the overall goal here is to synthesize and process files or streams of audio samples, while moving in the mid-1980s to a more "real-time" mode better fitted to live performances. This branch started around 1957 with the MUSIC I language of Max Mathews [Mathews et al. 1969], then an engineer at the Bell Telephone Laboratories (Murray Hill, New Jersey), on an IBM 704; at the time, hours of computation were necessary to get a few seconds of sound. The concept of *unit generator*, implemented in Mathews' Music-N languages, will prove itself a pervasive concept in audio signal processing, both in computer music languages and hardware synthesizers.

We classify below several significant audio synthesis languages and systems, hence providing a global picture of the domain – note though that some boundaries may be somehow artificial as some languages belong to several categories:

Textual Languages. Music-N family languages (I, II, III, IV, V) [Mathews et al. 1969], Csound [Vercoe 1992; Boulanger et al. 2000], SAOL [Scheirer and Vercoe 1999], Faust [Orlarey et al. 2009], Nyquist [Dannenberg 1997], SuperCollider [McCartney 1996], ChucK [Wang et al. 2003], Impromptu [Sorensen 2005];

[2]MUsic Simulator-Interpreter for COMpositional Procedures.

Visual Programming Environments. Max/MSP [Puckette 1991; Zicarelli 1998], Pure Data [Puckette 1996], jMax [Déchelle et al. 1999], Open Sound World [Chaudhary et al. 2000];
Physical Modeling Systems. Modalys [Eckel et al. 1995], Chant [Rodet et al. 1984], Genesis/Cordis-Anima [Castagné and Cadoz 2002; Cadoz et al. 1993];
Miscellaneous. Kyma [Scaletti 1987] (graphical sound design environment), STK [Cook and Scavone 1999] (C++ toolkit), CLAM [Amatriain et al. 2006], SndObj [Lazzarini 2000].

2.2. Synchronous Languages

Synchronous programming languages appeared in the early 1980s in France, with Esterel (École des mines de Paris and INRIA, Sophia Antipolis), Lustre (Verimag/CNRS, Grenoble) and Signal (INRIA, Rennes), as an academic research field mixing control theory and computer science [Benveniste and Berry 1991a; Halbwachs 1993; 2005], before becoming of high industrial interest for critical systems [Benveniste et al. 2003] such as those present in avionics, trains and nuclear power plants. The idea of *synchrony* was arising also through Milner's work on communicating systems [Milner 1980], AFCET[3]'s Grafcet [Baker et al. 1987] and Harel's Statecharts formalism [Harel 1987].

Synchronous languages are high-level, engineer-friendly, robust, specification formalisms, rooted in the concepts of *discrete time* and *deterministic concurrency*. Time is usually not explicitly mentioned in the definition of traditional programming languages. Such a notion is however of paramount importance in the design and implementation of data-flow and control software for *reactive systems* [Harel and Pnueli 1985; Halbwachs 1993]; indeed interactions with external environment processes are there subject to time constraints, memory constraints, security constraints and determinism requirements. Synchronous languages, which strive to reach such demanding objectives, are often equipped with timing and concurrency mathematical models that are structured around automata theory and a typical core hypothesis of bounded calculus and communication between logical instants; this paradigm is called the "synchronous hypothesis" and is typically checked by calculating the *worst case execution time*. Since the semantics of synchronous languages assumes that computation and communication are performed within logical instants, these languages provide, contrarily to traditional ones, programmers and the running environment explicit access to time, via clocks. Computations are specified with respect to these explicit clocks, ensuring that timing constraints can be stated by programmers and verified by compilers later on.

More specifically, in synchronous languages, time is seen as a succession of shared logical instants generated by regular (hence synchronous) support clocks, which make programs actually move forward. Synchronous programming uses specific programming languages designed to offer a programming paradigm centered on clocks, and clocks only. The overall goal here is not to provide a general-purpose programming set of features, as can be found in traditional languages such as C/C++, that can be used to address any programming issue, but to help with the sole implementation of critical real-time systems that need to satisfy stringent, "hard real-time" timing constraints. The specification and execution of a synchronous program amount, at least in principle, to the sequencing of an infinite loop of tightly time-constrained sets of *atomic reactions*. All computations and communications must always be performed within one base clock period. This property is the crux of synchronous programming languages

[3]*Association française pour la cybernétique économique et technique*.

design: it ensures that both concurrency, inherent to reactive systems, and determinism, highly desirable for critical systems, are preserved. Traditional, general-purpose programming languages are usually unable to guarantee such a correctness feature, at neither the logical nor the temporal level, while the very model of synchronous languages implementation makes it easier to check both run-time timing load and memory footprint. Moreover, limiting the generality of the programming paradigm ensures that implementations will be closer to specifications, thus increasing readibility, debugging and maintainability.

To illustrate the richness and diversity of the synchronous programming research field, we provide below a list of forty synchronous languages of interest. This list is neither intended to be exhaustive nor limited to a particular paradigm; it is rather a large-spectrum overview, mixing together very different languages toward a taxonomy of a significant set of synchronous and synchronous-oriented languages. We loosely categorized them using the criteria of syntax (textual, graphic), language definition approach (full-fledged or language extension) and application domain specificities (generic, hardware, models):

Textual Languages. Esterel [Berry and Cosserat 1985], Lustre [Caspi et al. 1987], Signal [Gautier et al. 1987; Gamatié 2009], ConcurrentML [Reppy 1999], Larissa [Altisen et al. 2006], Lucid Synchrone [Caspi and Pouzet 1996], Prelude [Pagetti et al. 2011], Quartz [Schneider 2000], ReactiveML [Mandel and Pouzet 2005], RMPL [Ingham et al. 2001], SL [Boussinot and De Simone 1996], SOL [Bharadwaj 2002], StreamIt [Thies et al. 2002], 8 ½ [Giavitto 1991];
Visual Languages and Environments. Argos [Maraninchi 1991], Statecharts [Harel 1987], SyncCharts [André 1996], Argonaute [Maraninchi 1990], Polis [Balarin 1997], Polychrony [Le Guernic et al. 2003], Scade [Dormoy 2008], Simulink/Matlab [Caspi et al. 2003];
Language Extensions. ECL (C)[4] [Lavagno and Sentovich 1999], Jester (Java) [Antonotti et al. 2000], Reactive-C (C) [Boussinot 1991], Real-time Concurrent C (C) [Gehani and Ramamritham 1991], RTC++ (C++) [Ishikawa et al. 1992], Scoop (Eiffel) [Compton 2000], SugarCubes (Java) [Boussinot and Susini 1998];
Hardware Description Languages. Lava [Bjesse et al. 1998], SystemC [Initiative 2006], Verilog [Thomas and Moorby 2002], VHDL [IEEE standard 1988];
Models and Intermediate Formats. Averest [Schneider and Schuele 2005], DC+ [Pnueli et al. 1998], OC [Girault 2005], SC [Girault 2005], DC [Girault 2005], CP [Girault 2005], SDL [Ellsberger et al. 1997], ULM [Boudol 2004], UML Marte [Mallet and André 2009].

2.3. Key Language Representatives

The family of significant synchronous languages dedicated to music is clearly more limited than the one of general purpose languages, although music seems an obvious application field for the synchronous programming paradigm and its associated languages. Indeed, many concurrent processes (associated to instruments or artists) are intimately linked to the audio sampling frequency that drives the production of sound samples. One of our goals with this research work is to illustrate that a bridge can be made between the synchronous and music programming paradigms we just surveyed.

Even though the brief presentation, above, of existing music and synchronous languages is by no means exhaustive, it makes it obvious that these families of languages offer a very wide variety of possible candidates for our comparison survey. To make our use case analysis project realistic, we need to make a selection to end up with a

[4]The original language is put in parenthesis.

manageable small subset of these languages. We based our selection criteria on the availability of each language and its associated tools, the import of its design principles on the history of the synchronous and computer music paradigms and a rough assessment of the size of its user base. We thus end this section with the selection of the representative languages that are the basis for our use case study, listed in Table I with the name of the main institution where they were created and their authors.

Table I. Selection of music and general-purpose synchronous languages

Csound	MIT	B. Vercoe, J. Fitch *et al.*
SuperCollider	Univ. Texas	J. McCartney *et al.*
Pure Data	UCSD	M. Puckette
ChucK	Princeton Univ.	G. Wang, P. Cook
Faust	GRAME	Y. Orlarey *et al.*
Signal	IRISA/INRIA	A. Benveniste, P. Le Guernic
Lustre	CNRS/Verimag	P. Caspi, N. Halbwachs
Lucid Synchrone	Verimag & Paris 11	P. Caspi, G. Hamon, M. Pouzet
Esterel	MINES/INRIA	G. Berry *et al.*
OpenMP Stream	MINES ParisTech	Antoniu Pop

3. THE OSCILLATOR USE CASE

Our survey of significant technological tools for the synchronous programming of audio applications is grounded on practical terms. We decided to perform for such an analysis a use case study and picked `osc`, an implementation of a sound oscillator, as our test case. We selected this application since it is particularly simple and can be coded in a few lines, making it a program of choice for a comparison survey tackling multiple languages. This choice is also perfectly adequate on a more functional level, since this simple truncated lookup table oscillator algorithm is particularly significant for the audio domain, being one of the most classical algorithms of the sound synthesis field. Notably, it is also involved in other important computer music algorithms, such as wavetable synthesis, additive synthesis or FM synthesis (frequency modulation).

Of course, typical audio processing applications such as delays or reverbs are more complex than our simple `osc` use case, in particular since many of them can be seen as filters, taking input sample streams and producing modified ones. Dealing with such more involved applications might possibly provide some additional insights into the nature of the relationships linking DSLs and synchronous languages; yet, we believe that the ubiquitous nature of `osc` in audio processing makes it illuminating enough to expose the gist of each programming language addressed in our survey.

Although it would theoretically be possible to program the algorithm for `osc` at the same abstraction level for all languages, we felt that such an *a priori* fair approach would not lead to a convincing comparison. Indeed, by essence, the very concept of DSL has been introduced to provide idioms intending to ease programming. One of the goals of this survey is to compare programming paradigms, and possibly provide a case for the introduction of more *domain-specific* constructs in languages, including synchronous ones. Using languages, be they domain-specific or synchronous, at their top potential is thus a key ingredient to reaching that goal.

3.1. Presentation

The purpose of `osc` is to output, in a programming language-specific manner, the successive samples of a sinusoidal waveform stored in a vector[5]; the wave frequency is a parameter of this process, and can be changed at start-up time. The basic idea is to always loop over the same single sinusoidal vector for all frequencies, but to decide which sound samples to output according to what the requested frequency `freq` is; for instance, picking every other sample will provide a signal with a frequency twice that of the original if the sinusoidal vector is always looped over at the same rate. In more details (see also Figure 1):

— during the initialization stage, one period of the sin function is sampled and `tablesize` samples are stored in the vector `sinwaveform`;
— the main function `osc(freq)` loops over this vector indefinitely while outputting the successive sound samples of frequency-dependent phase, i.e., vector index `int(phase(freq))`, for each time tick;
— each `phase(freq)` is the product of `tablesize` and one of the steps $i(n)$ defined as $i(n) = \{i(n-1) + \texttt{freq}/\texttt{samplingfreq}\}$ and $i(0) = 0$, where $\{x\}$ denotes the fractional part of any number x and `samplingfreq` is the audio output sampling frequency — this kind of recursive equation is a typical tenet of both synchronous languages and digital signal processing (DSP) applications;
— the audio sample corresponding to each particular phase is provided by the `rdtable` function, returning `sinwaveform[int(phase(freq))]`; the integer typecast `int` computes the integer part of its argument by truncation of the fractional part[6], thus ensuring that table indexes are of integer type;
— each sampled data point is finally output, in a more or less platform-specific manner.

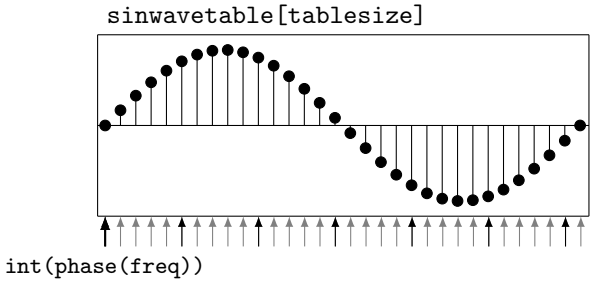

Fig. 1. Truncated lookup table oscillator

3.2. Interface

At the highest level, nine general constants and functions have to be provided to ensure the existence of a working implementation of `osc`; the corresponding signature is listed in Table II. Of course, depending on the particular programming language and its abstraction level, some of these functions will be indeed visible in the `osc` program text; for others, they will only be implicitly present in our implementation.

[5]Note that adapting this technique to generate a different waveform, and thus yield a different sound, amounts to simply changing the values stored in the vector of samples. This approach is, by the way, at the core of most current commercial wavetable-driven music synthesizers.
[6]`int` corresponds to the *floor* function if its argument is positive, which is the case for `phase`, and otherwise to *ceiling*.

Table II. Oscillator signature: general constants and functions

```
const tablesize = 65536              // number of sound samples
const sinwaveform[tablesize]         // sampled sinusoid (one period)
const samplingfreq = 44100           // audio sampling rate (Hz)
const freq = 440                     // 'A' diapason frequency (Hz)
const twopi = 6.28318530717958623
void osc(freq)                       // main function
float rdtable(table, index)          // dynamic table read access
float phase(freq)                    // phase for each tick
float fracpart(x)                    // fractional part of float x in [0,1)
```

Environment. In order for `osc` to be implemented in a particular language, the following features need to be available in the programming environment:

— the two mathematical functions `sin` and `floor` (used to compute the fractional part of a number $\{x\} = x - \lfloor x \rfloor$);
— the ability to perform dynamic reads of tables (vectors of samples);
— a looping construct for table initialization and `osc`.

Usage. The `osc` process is launched by calling the main function `osc` with the chosen frequency as argument, e.g., `osc(440)`. In this study, we use a constant frequency value, but ideally the `freq` argument should be an input signal, i.e, a stream of frequencies, for example [523.25, 587.33, 659.26, ...] – this particular succession of frequencies would in practice yield a sequence of notes, here [C, D, E, ...]. Such an extension would not have a significant impact on the structure of our `osc` algorithm, and thus our example, kept short due to space limitations, is enough to present the key ingredients of each programming language.

In theory, the `osc` algorithm never terminates, as it is a synchronous program that, given an input frequency, generates sinusoidal sample values in an output stream. In practice, depending on the language at hand, the user will either have to interrupt the execution after enough samples have been output, in many cases by typing `'ctrl-c'`, or get a finite vector of `outputsize` samples. However, even though the output behaviors of our implementations vary according to the language at hand, this is a rather peripheral issue linked to the actual input/output environment used to run our code; this does not significantly alter the synchronous specification of the `osc` process, which is the focus of interest here.

For the validation of our tests, we computed a reference vector of 200 sound samples; it begins with the following rounded values, corresponding to the 440 Hz diapason at the 44,100 Hz sampling rate: [0.0000, 0.0626, 0.1250, 0.1869, 0.2481, 0.3083, 0.3673, 0.4249, 0.4807, 0.5347, 0.5866, 0.6362, 0.6833, 0.7277, ...]. The output of each implementation has been compared to this reference vector, any mismatch being an indication of something wrong with the corresponding implementation.

3.3. Specification

We provide here an implementation of `osc` using indifferently Octave[7] or Matlab[8], a well-known "high-level technical computing language and interactive environment for algorithm development, data visualization, data analysis and numeric computation". This straightforward imperative implementation should be easily understood by most readers, and serves here as a more formal specification of our use case. Here, the `osc` output is the array `waveform`.

[7]Octave is an open-source variant of Matlab: http://www.gnu.org/software/octave/.
[8]http://www.mathworks.com.

```
function [waveform] = osc(freq)
tablesize = bitshift(1, 16);
samplingfreq = 44100;
outputsize = 200;
twopi = 2 * pi;

indexes(1) = 0;
waveform(1) = 0;

for i = 1 : tablesize
  sinwaveform(i) = sin( (i * twopi) / tablesize);
end

for i = 2 : outputsize
  indexes(i) = fracpart( (freq / samplingfreq) + indexes(i-1) );
  phase = tablesize * indexes(i);
  waveform(i) = sinwaveform(uint16(phase));
end
end

function [y] = fracpart(x)
y = x - floor(x);
end
```

Design Notes

— The `bitshift` expression shifts here a `1` sixteen times on the left, yielding 2^{16} as required.

— The keyword `function` introduces the definition of a function. Here the `fracpart` function yields in the return value named `y` the fractional part of its argument `x`, and is used to perform a round-robin in the $[0, 1)$ interval (closed at 0 and open at 1).

— The intrinsic function `uint16` returns an unsigned 16-bit integer that approximates its floating point argument, here `phase`. This truncation of the floating point phase values provides the integer indexes needed to access the wave table in read mode.

4. MUSIC LANGUAGES

In some loose sense, all music-specific programming languages use, in one way or another, synchronous idioms, since they have to deal with temporal streams of audio samples. We decided to adopt here a somewhat historical order to present key music programming languages:

— Csound is, in a way, the father of modern audio synthesis languages [Vercoe 1992; Boulanger et al. 2000];
— SuperCollider adopts an object-oriented programming approach, inspired by the SmallTalk language [McCartney 1996];
— Pure Data is, like Max/MSP [Puckette 2002], a typical representative of the visual programming paradigm often adopted by the computer music community, thanks to its appeal to the contemporary music composers [Bresson et al. 2009];
— ChucK exemplifies the importance of on-the-fly programming that now can occur even during music performance through “live coding” practices [Wang et al. 2003];
— finally, Faust promotes the functional paradigm onto a block-diagram algebra, striving to balance expressivity and run-time performance [Orlarey et al. 2002; 2009].

4.1. Csound

Presentation. The following position statement is extracted from the Csound official site, http://www.csounds.com.

> Csound is a sound design, music synthesis and signal processing system, providing facilities for composition and performance over a wide range of platforms. It is not restricted to any style of music, having been used for many years in the creation of classical, pop, techno, ambient, experimental, and (of course) computer music, as well as music for film and television.

Oscillator. The Csound implementation `osc.csd` of the oscillator can be found below. We tested Csound version 5.11 (float samples) Sep 24 2009, with the graphical interface QuteCsound version 0.4.4. The `osc` execution result is an infinite sample stream, output on the default Csound audio port.

```
<CsoundSynthesizer>
<CsInstruments>
sr              =       44100
kr              =       441
ksmps           =       100
nchnls          =       1

                instr 1
aosc            oscil   p4, p5, 1
                out     aosc
                endin
</CsInstruments>
<CsScore>
; use GEN10 to compute a sine wave
f1      0       65536   10      1
;ins    strt    dur     amp     freq
i1      0       2       20000   440
e
</CsScore>
</CsoundSynthesizer>
```

Design Notes

— Csound is the oldest musical language of this study: it is a C-based audio DSL following Music11, also developed by Barry Vercoe at MIT in the 1970s, and the MUSIC-N languages initiated by Max Mathews at the Bell Labs in the 1960s. Several aspects present in the Csound language (possibly inherited from previous languages) still persist in later musical languages, so we detail here several aspects of Csound that will be of use for the understanding of most of the sections dedicated here to musical languages.

— In the header, the two assignments "`sr = 44100`" and "`kr = 441`" stand for *sample rate* and *control rate*, which are two fundamental concepts in computer music. The first one specifies the discrete audio sampling rate, set according to the Nyquist frequency, which is twice the 20,000 Hz upper human listening bound needed at sampling time to avoid aliasing. The second one specifies the "control" rate; to save computing resources, it is usually set to a value smaller than the audio rate, since it is mostly used to manage music control information, which do not require the very high temporal resolution requested by sound signals.

— The Csound code of `osc.csd` presented here gathers into one file both the so-called *orchestra* and *score* parts, using XML-like sections, although Csound code had originally been stored in two distinct '`.orc`' and '`.sco`' files. "An *orchestra* is really a computer program that can produce sound, while a *score* is a body of data which that

program can react to." [Vercoe 1992]. Beside the score, other means of launching the instruments exist, such as via MIDI[9] or real-time events (see Section 6.5).

— The Csound syntax makes intensive use of one-letter prefixes:

— in an instrument definition, "`a`" and "`k`" specify signal rates (so that "`aosc`" is an audio-rate oscillator signal);

— in an instrument definition, "`p`" followed by a number specifies a reference to the corresponding *parameter* field in the score part (here `p4` corresponds to `20000` and `p5` to `440` when used in `i1`, see below);

— in the score part, "`f`" followed by a list of numbers declares a *function* table (here, the table `f1` will be computed at `0` second, on `65536` points, using the *unit generator* `10`, with a relative energy of `1` for the fundamental frequency; see next notes);

— in the score part, "`i`" followed by a list declares a Csound note that references an instrument to be played (here `i1` requests the computation of `instr 1`, from time `0`, during `2` seconds, with an amplitude of `20000`, at a frequency of `440` Hz);

— finally, `e` asks for the execution of the score.

— The oscillator is synthesized by `instr 1` in the orchestra part, using the `oscil` generator, which is a simple direct synthesis oscillator without interpolation. This generator has three arguments: its amplitude `p4`, its frequency `p5` and its function table number `1`, which refers to `f1` in the score part; this latter function table relies on the tenth *unit generator* called "GEN10", specified as the fourth argument of `f1`.

— As GEN10 is also used in several musical languages in the next sections, we cite here its extensive description from the Canonical Csound Reference Manual [Vercoe et al. 2007]:

> **GEN10** – Generate composite waveforms made up of weighted sums of simple sinusoids.
> **Description**: These subroutines generate composite waveforms made up of weighted sums of simple sinusoids. The specification of each contributing partial requires 1 pfield using *GEN10*.
> **Syntax**: `f# time size 10 str1 str2 str3 str4 ...`
> **Initialization**: *size* – number of points in the table. Must be a power of 2 or power-of-2 plus 1. *str1, str2, str3, etc.* – relative strengths of the fixed harmonic partial numbers 1, 2, 3, etc., beginning in p5. Partials not required should be given a strength of zero.
> **Note**: These subroutines generate stored functions as sums of sinusoids of different frequencies. The two major restrictions on *GEN10*, namely that the partials have to be harmonic and in phase, do not apply to *GEN09* or *GEN19*. In each case the composite wave, once drawn, is then rescaled to unity if p4 was positive. A negative p4 will cause rescaling to be skipped.

4.2. SuperCollider

Presentation. The following position statement is extracted from the SourceForge http://supercollider.sourceforge.net site, since the SuperCollider official site, http://www.audiosynth.com, seems to be less well maintained.

> SuperCollider is an environment and programming language for real-time audio synthesis and algorithmic composition. It provides an interpreted object-oriented language which functions as a network client to a state of the art, real-time sound synthesis server.

[9] *Musical Instrument Digital Interface*, an industrial protocol defined in 1983.

> SuperCollider was written by James McCartney over a period of many years, and is now an open source (GPL) project maintained and developed by various people. It is used by musicians, scientists, and artists working with sound.

Oscillator. The SuperCollider implementation `osc.scd` of the oscillator can be found below. We tested SuperCollider version 3.4, rev 10205. The `osc` execution result is an infinite sample stream, output on the default SuperCollider audio port.

```
(
  var tablesize = 1 << 16;
  b = Buffer.alloc(s, tablesize, 1);  // allocate a Buffer
  b.sine1(1.0, true, false, true);    // fill the Buffer
  {OscN.ar(b, 440, 0, 1)}.play        // N: Non-interpolating
)
b.free;
```

Design Notes

— In SuperCollider [McCartney 1996], which has over 250 unit generators (cf. [Valle et al. 2007]), such an oscillator could have been achieved in at least four different ways:
(1) the one presented here, using a `Buffer` filled by a `sine1` pattern and played by a non-interpolating `OscN` wavetable oscillator;
(2) replacing `OscN` by an interpolating `Osc` wavetable oscillator;
(3) using `BufRd`, `BufWr` and `SinOsc` (since `BufRd` is to be filled by a *unit generator*);
(4) using directly `SinOsc`.

— Since Version 3, SuperCollider is build upon a client/server architecture (communicating via OSC[10]), with a synthesis application on the server side and a remote language application on the client side. So here the `free` method allows the client to tell the server to free the memory of the buffer previously used.

— The first level of parenthesis surrounding the main block of code is a syntactic trick that is used to ensure that the enclosed lines of code will be launched at the same time by the SuperCollider interpreter; this occurs, in practice, when one simply double-clicks inside one of the parentheses (cf. [Valle et al. 2007]).

4.3. Pure Data

Presentation. The following statement is extracted from the Pd-FlossManual, available at http://en.flossmanuals.net/PureData. The Pure Data official site is http://puredata.info.

> Pure Data (or Pd) is a real-time graphical programming environment for audio, video, and graphical processing. Pure Data is commonly used for live music performance, VeeJaying, sound effects, composition, audio analysis, interfacing with sensors, using cameras, controlling robots or even interacting with websites. Because all of these various media are handled as digital data within the program, many fascinating opportunities for cross-synthesis between them exist. Sound can be used to manipulate video, which could then be streamed over the internet to another computer which might analyze that video and use it to control a motor-driven installation.
>
> Programming with Pure Data is a unique interaction that is much closer to the experience of manipulating things in the physical world. The most

[10] *Open Sound Control*, a content format for digital device communication defined at UC Berkeley CNMAT [Wright 2005].

basic unit of functionality is a box, and the program is formed by connecting these boxes together into diagrams that both represent the flow of data while actually performing the operations mapped out in the diagram. The program itself is always running, there is no separation between writing the program and running the program, and each action takes effect the moment it is completed.

Oscillator. The Pd implementation `osc.pd` of the oscillator can be found in Figure 2. We tested Pure Data version 0.42.5-extended-20091222. The `osc` execution result is an infinite sample stream, output on the default Pure Data audio port.

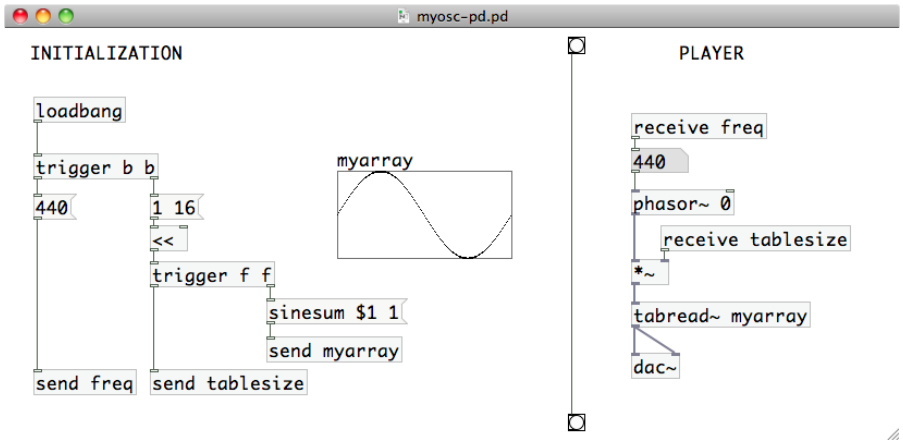

Fig. 2. The `osc.pd` implementation in Pure Data.

Design Notes

— Pure Data is a graphical programming environment [Puckette 1996], where a graphical window containing Pure Data code (i.e. boxes and connections) is called a *patch*. The implementation shown on Figure 2 is a screenshot of the patch `osc.pd`.

— This patch is graphically separated into two main parts labelled `PLAYER` and `INITIALIZATION`; the graphical separation is materialized by an idiosyncratic line made up of a connection between two dummy `bang` objects. A `bang` object (a circle in a box) is used to trigger the most primitive of all event messages, in which only the information that a constant *bang* value has to be sent is encoded.

— Two `trigger` objects are used in order to enforce a sequential order of message emission (*right-to-left* order in Pure Data), which is crucial for such an event-driven language. The syntax of the `trigger` object includes a list of types, where the number of elements determines the number of outputs and where the type names (or one-letter type aliases) determine the type for the corresponding output, allowing type conversion. Here, the letters `f` and `b` stand respectively for types `float` and `bang`.

— For instance, the `[trigger f f]` object receives as input the value of `1 << 16`, i.e., 2^{16}. Proceeding right to left, it first initializes the local array `myarray` with the values of the first harmonic (specified by `1`) of the `sinesum` function, predefined in Pure Data and having a semantics similar to *GEN10*'s; the number of samples for one period is the value of the first argument `$1`, here 2^{16}. Then, in a second step that exhausts its list

of types, the trigger launches the command `[send tablesize]`, which will be received by the `PLAYER` process[11].

— Object names ending with a "∼" (tilde) character denote *signal objects*, running at audio rate, and bold connections denote *signal connections*. "[Tilde objects] use continuous audio streams to intercommunicate, and also communicate with other ('control') Pd objects using messages." [Puckette 2007]

— To read our array `myarray` of sine samples, we use a `phasor∼` object, which outputs a sawtooth signal between 0. and 1., here multiplied by the table size for table lookup to yield the successive indices. The dummy argument `0` allows `phasor∼` to receive a non-signal message for the frequency (at control rate).

— The `dac∼` object, for *digital-to-analog converter*, transfers the real-time audio outputs of Pure Data patches to the audio driver of the underlying operating system.

4.4. ChucK

Presentation. The following position statement is a mix of texts from the official site http://chuck.cs.princeton.edu and the ChucK manual.

> ChucK is a new (and developing) audio programming language for real-time synthesis, composition, performance, and now, analysis. ChucK presents a new time-based, concurrent programming model that's highly precise and expressive (we call this strongly-timed), as well as dynamic control rates, and the ability to add and modify code on-the-fly. In addition, ChucK supports MIDI, OSC, HID, and multi-channel audio. It's fun and easy to learn, and offers composers, researchers, and performers a powerful programming tool for building and experimenting with complex audio synthesis/analysis programs, and real-time interactive control.

Oscillator. The ChucK implementation `osc.ck` of the oscillator can be found below. We tested ChucK version 1.2.1.3 (dracula). The `osc` execution result is an infinite sample stream, output on the default ChucK audio port.

```
Phasor drive => Gen10 g10 => dac; // gen10 sinusoidal lookup table

[1.] => g10.coefs; // load up the partials amplitude coeffs
440 => drive.freq; // set frequency for reading through table

while (true) // infinite time loop
{
    500::ms => now; // advance time
}
```

Design Notes

— The ChucK language is specifically designed to allow *on-the-fly* audio programming [Wang et al. 2003].

— The heart of ChucK's syntax is based around the massively overloaded *ChucK operator*, written as '=>'. "[This operator] originates from the slang term 'chuck', meaning to throw an entity into or at another entity. The language uses this notion to help express sequential operations and data flow" [Wang et al. 2003]. The ChucK operator's behavior relies on the strong typing system of this imperative language, depending on the type of both its left and right arguments.

[11] `send` commands are, in fact, spurious, since they could replaced by graphical connections; they have been introduced in Pure Data to help programmers better visually structure their programs.

— Several elements of ChucK, such as `Gen10` (cf. Csound section), `Phasor` or `dac` (cf. Pure Data section), are inspired by features existing in previous musical languages.

— Here, the `drive` phasor is declared and piped to the `g10` generator, itself connected to the digital-to-audio converter.

— Like in Csound, `Gen10` has to be fed with a list of relative coefficients specifying the harmonics of the spectra; here a single-element array `[1.]` yields a single sinusoid.

— The infinite time loop allows the computing and playing processes declared above it to run; the loop body merely advances time by the arbitrary duration of `500::ms`. Note that the timing model mandates the attachment of a time unit to each duration, such as milliseconds in "`500::ms`".

— Modifying the special variable `now` has the effect of advancing time, suspending the current process until the desired time is reached, and providing the other processes and audio synthesis engine with the computing resources needed to run in parallel.

— The value of `now`, which holds the current time, only changes when it is explicitly modified [Wang and Cook 2007]. "The amount of time advancement *is* the control rate in ChucK." [Wang et al. 2003]

4.5. Faust

Presentation. The following position statement is extracted from the Faust official site, http://faust.grame.fr.

> FAUST is a compiled language for real-time audio signal processing. The name FAUST stands for Functional AUdio STream. Its programming model combines two approaches: functional programming and block diagram composition. You can think of FAUST as a structured block diagram language with a textual syntax.
>
> FAUST is intended for developers who need to develop efficient C/C++ audio plugins for existing systems or full standalone audio applications. Thanks to some specific compilation techniques and powerful optimizations, the C++ code generated by the Faust compiler is usually very fast. It can generally compete with (and sometimes outperform) hand-written C code.
>
> Programming with FAUST is somehow like working with electronic circuits and signals. A FAUST program is a list of definitions that defines a signal processor block-diagram: a piece of code that produces output signals according to its input signals (and maybe some user interface parameters).

Oscillator. The Faust implementation `osc.dsp` of the oscillator can be found below. We tested Faust version 0.9.13. The `osc` execution result is an infinite sample stream, output as the `process` output signal, linked to the standard audio port via Jack[12].

```
import("math.lib");   // for SR and PI

tablesize       = 1 << 16;
samplingfreq    = SR;
twopi           = 2.0 * PI;

time            = (+(1) ~ _) , 1 : -;  // 0,1,2,3,...
sinwaveform     = twopi*float(time)/float(tablesize) : sin;

fracpart(x)      = x - floor(x);
phase(freq)     = freq/float(samplingfreq) :
                  (+ : fracpart) ~ _ : *(float(tablesize));
osc(freq)       = rdtable(tablesize, sinwaveform, int(phase(freq)));
```

[12] http://jackaudio.org

```
process = osc(440);
```

Design Notes

— The Faust language combines a block-diagram algebra [Orlarey et al. 2002] with a functional paradigm [Orlarey et al. 2004].

— The keyword `process` is analogous to `main` in C and has to be defined [Orlarey et al. 2004].

— The sample rate constant `SR` is defined in the imported library file `math.lib` as a *foreign constant*, which is linked to the actual sampling rate of the host application through the architecture compilation mechanism of Faust and determined at initialization time [Smith III 2010]. This typical DSL feature protects against incompatibilities.

— Faust uses five block-diagram composition operators [Orlarey et al. 2004]: sequential composition `A:B`, parallel composition `A,B`, recursive composition `A~B`, split composition `A<:B` and merge composition `A:>B` (the last two are not used here).

— The `time` processor definition "`time = (+(1) ~ _) , 1 : -;`" uses the three essential block-diagram composition operators plus two more key elements. From right to left, there are:

— the sequential composition operator "`:`", to connect the two inputs of the "`-`" processor with the output signals of the preceding processors to compute differences;
— the parallel composition operator "`,`", to combine the two parallel processors that feed the "`-`" processor (note that order matters for the subtraction);
— the recursive composition operator "~", to specify a one-sample feedback increment that generates a series of natural numbers, starting at 1 (Figure 3 shows the block-diagram schema of the `time` processor, where the small square on the output denotes the implicit sample delaying operation);
— the identity block "`_`", used here in the one-sample recursive loop to directly connect the output to the input of the increment processor, with no other processor than the *identity* one;
— the partial application "`+(1)`", using the curried form of the processor "+" to fix one of its arguments with the "1" value.

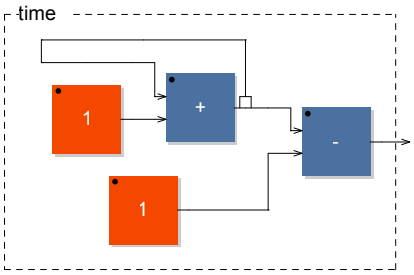

Fig. 3. Block-diagram schema of the `time` processor

— Integer operations are defined modulo the (implementation-dependent) machine integer size, so the increment operation above never overflows, but simply performs a round-robin on its domain.

— Infix notation is allowed in Faust as syntactic sugar, as in the `fracpart` definition where "`x - floor(x)`" is equivalent to "`x, floor(x) : -`".

— The read-only table `rdtable` stores values from a stream at initialization time. The data will then be accessed during execution. Its three parameters are the size of the table, the initialization stream, and the read-index signal [Gaudrain and Orlarey 2003].

— Faust default output is the audio system within which a Faust `process` is run.

5. SYNCHRONOUS LANGUAGES

There are various ways to introduce synchronicity in general-purpose programming languages. We focus, in this section, on our oscillator use case example to survey the usual approaches described in the literature:

— the direct approach, taken by Signal, Lustre and Esterel, is to make the notion of synchronous computation at the core of the language design *per se*;
— a more indirect route is to add the notion of streams to an existing language, and among the multiple existing proposals we decided to present two such integrations as illustration: Lucid Synchrone, for the functional paradigm, over OCaml, and OpenMP Stream Extension, for the imperative one, over C and OpenMP.

5.1. Signal

Presentation. The following position statement is extracted from the Signal official site, http://www.irisa.fr/espresso/Polychrony.

> Signal is based on synchronized data-flow (flows + synchronization): a process is a set of equations on elementary flows describing both data and control.
>
> The Signal formal model provides the capability to describe systems with several clocks (polychronous systems) as relational specifications. Relations are useful as partial specifications and as specifications of non-deterministic devices (for instance a non-deterministic bus) or external processes (for instance an unsafe car driver).
>
> Using Signal allows to specify an application, to design an architecture, to refine detailed components down to RTOS[13] or hardware description. The Signal model supports a design methodology which goes from specification to implementation, from abstraction to concretization, from synchrony to asynchrony.

Oscillator. The Signal implementation `OSC.SIG` of the oscillator can be found below. We tested Signal version V4.16. The `osc` execution result is an infinite sample stream, output on the `output` signal; additional helper files, in particular clock files specifying explicit time ticks (not shown here), are needed to generate a user-specified output file of 200 samples.

```
process osc =
     ( ? event inputClock;
       ! dreal output;
     )
   (| output ^= inputClock
    | output := rdtable(integer(phase(freq)))
    |)
   where
   constant dreal freq = 440.0;
   constant integer samplingfreq = 44100;
   constant integer tablesize = 2**16;
```

[13]Real-Time Operating System.

```
constant dreal twopi = 6.28318530717958623;
process rdtable =
      ( ? integer tableindex;
        ! dreal sample;
      )
    (| sample := sinwaveform[tableindex]
     |)
    where
    constant [tablesize]dreal sinwaveform =
      [{i to (tablesize-1)}:sin((dreal(i)*twopi)/dreal(tablesize))];
    end;
process phase =
      ( ? dreal freq;
        ! dreal phi;
      )
    (| index := fracpart((freq/dreal(samplingfreq))+index$)
     | phi := dreal(tablesize)*index
     |)
    where
    dreal index init 0.0;
    end;
function fracpart =
      ( ? dreal fracpartIn;
        ! dreal fracpartOut;
      )
    (| fracpartOut := fracpartIn-floor(fracpartIn) |);
end;
```

Design Notes

— Signal processes manipulate signals, i.e., named streams of typed data, either as input "?" or output "!", here of floating point numbers in double precision `dreal`. Local subprocesses are defined similarly.

— Process behavior is defined via sets of functional equations on signals between the `(|` and `|)` enclosing symbols; these equations constrain either the values in a given signal, via the `:=` connector, or clocks, via the `^=` connector, used here to impose that signals `output` and `inputClock` share the same timing information.

— Arrays, such as `sinwaveform` of `tablesize` elements, are defined by intension, at initialisation time, using implicit quantification over indices such as `i` here.

— Data equations are functional, and the `$` postfix is used to reference the previous item in a stream, while `init` is used to specify the initial value.

— The execution of `osc` relies upon the local `rdtable` process, parametrized with the appropriate frequency, to define its `output` signal. The `rdtable` process, in turn, outputs, for each integer in its `tableindex` input signal, the corresponding value of the `sinwaveform` table, initialized once and for all with a sampled sine function.

— Signal allows the use of simple processes called `functions`, such as the straightforward `fracpart` definition here, to help modularize Signal specifications.

5.2. Lustre

Presentation. The site http://www-verimag.imag.fr/The-Lustre-Toolbox.html is Lustre official repository and the following position statement is extracted from Wikipedia[14].

> Lustre is a formally defined, declarative, and synchronous dataflow programming language, for programming reactive systems. It began as a research project in the early 1980s. In 1993, it progressed to practical, indus-

[14] http://en.wikipedia.org/wiki/Lustre_\%28programming_language\%29.

trial use, in a commercial product, as the core language of the industrial environment SCADE, developed by Esterel Technologies. It is now used for critical control software in aircraft, helicopters, and nuclear power plants.

Oscillator. The Lustre implementation `osc.lus` of the oscillator can be found below. We tested Lustre version V4. The `osc` expected result is an infinite sample stream, output on the `Y` signal of the `osc` node.

```
-- WARNING : Does *not* work, because of *dynamic* array accesses!

include "math.lus"

const samplingfreq = 44100;
const tablesize = 65536;
const timeTab = time(tablesize, 0);
const sinwaveform = sintable(timeTab);
const twopi = 6.28318530717958623;

node time( const n: int; start: int ) returns ( t: int^n );
let
  t[0] = start;
  t[1..n-1] = t[0..n-2] + 1^(n-1);
tel

node sintable ( X : int ) returns ( Y : real );
let
  Y = sin(((real X)*twopi) / (real tablesize));
tel

node fracpart ( X : real ) returns ( Y : real );
let
  Y = X - floor(X);
tel

node phase ( freq : real ) returns ( Y : real );
var index : real;
let
  index = 0.0 -> fracpart((freq/(real samplingfreq)) + pre(index));
  Y = (real tablesize) * index;
tel

node rdtable ( tableindex : int ) returns ( Y : real );
let
  Y = sinwaveform[tableindex]; -- Dynamic array access.
tel

node osc ( freq : real ) returns ( Y : real );
let
  Y = rdtable(int phase(freq));
tel
```

Design Notes

— Lustre sees computation as the processing of data exchanged between nodes. Within each node, functional definitions of data streams are expressed as possibly recursive equations.

— In a stream definition, `pre` is used to denote the previous value in the argument stream, while the arrow `->` operator ("followed-by") is used to distinguish the initial value from the recursive expression when defining a stream by induction.

— Streams are typed, and ^ is used to introduce aggregate vector types (its second argument is the vector size).

— The definition of the array `t` in `time` is by induction over array slices: `t[0]` is 0, while, for all i in `[1..n-1]`, `t[`i`]` is `t[`$i-1$`]`$+1$, since `1^(n-1)` is an array of `n-1` elements, all initialized to 1.

— Array `timeTab` is initialized via `time`, when the whole program is loaded; its elements are integers from 0 to `tablesize` - 1, numbering all samples in one sine period.

— The execution of `osc`, with a frequency argument as input, yields a stream based on the node `rdtable`. The `phase` node returns, for the same frequency, the recursively defined stream `index`, which is the succession of indices used to access, via `rdtable`, the `sinwaveform` array, also defined at load time.

— The `const` keyword flags identifiers denoting constant values. This is used here to compute the `sinwaveform` array.

— The Lustre version we used is in fact *unable* to manage dynamic access to arrays but only handles constant indexes, yielding a naming convenience for numbered constants like `a[1]`, `a[2]` and so on, while this dynamic access feature with variable index reading support is clearly required to implement variable frequencies in `osc` (see Sec. 6.4). Newer, non open-source versions of Lustre, such as Lustre V6, do provide dynamic arrays.

5.3. Esterel

Presentation. The following position statement is taken from the original Esterel site, http://www-sop.inria.fr/meije/esterel/esterel-eng.html.

> Esterel is both a programming language, dedicated to programming reactive systems, and a compiler which translates Esterel programs into finite-state machines. It is one of a family of synchronous languages, like SyncCharts, Lustre, Argos or Signal, which are particularly well-suited to programming reactive systems, including real-time systems and control automata.
>
> The Esterel v5 compiler can be used to generate a software or hardware implementation of a reactive program. It can generate C-code to be embedded as a reactive kernel in a larger program that handles the interface and data manipulations. It can also generate hardware in the form of netlists of gates, which can then be embedded in a larger system. Extensive optimization is available. We provide a graphical symbolic debugger for Esterel. We also provide support for explicit or BDD-based verification tools that perform either bisimulation reduction or safety property checking.
>
> Esterel is now experimentally used by several companies and taught in several universities.

Oscillator. The Esterel implementation `Osc.strl` of the oscillator can be found below (see Design Notes also). We tested version v5 (and GCC 4.2.1 for the C files handling array accesses). The `osc` execution result is an infinite sample stream, output on the default Esterel output port.

```
module Osc:
  function floor_int (double) : integer;
  constant tableSize_cte = 65536 : integer;
  input I : double;
  output O : double;

  signal index : integer, phase : double, sample : double in
    every I do
      run Phase [signal I / freq, phase / phi];
      ||
```

```
      loop
        emit index(floor_int(?phase));
        run RdTable [signal index / tableindex];
        emit O(?sample);
      each tick
    end every
  end signal
end module

module RdTable:
  function sinwaveform (integer) : double;
  input  tableindex : integer;
  output sample : double;
  emit sample(sinwaveform(?tableindex));
end module

module Phase:
  function floor_db(double) : double;
  constant samplingfreq_cte = 44100.0 : double;
  constant tableSize_cte_db = 65536.0 : double;
  input freq : double;
  output phi : double;

  signal step : double in
    emit step(?freq/samplingfreq_cte);
    var index := 0.0 : double, preindex := 0.0 : double in
      every immediate tick do
        emit phi(tableSize_cte_db*preindex);
        index := ?step+preindex;
        preindex := index - floor_db(index);
      end every
    end var
  end signal
end module
```

Design Notes

— The Esterel implementation of the oscillator relies on the Esterel `Osc.strl` module file provided above and C helper functions, not presented here, that handle arrays, a data structuring mechanism not provided by the public-domain version of Esterel we used. We wrote the following functions, declared as external C functions in Esterel modules via the `function` keyword:

— `void init_sinwaveform()`, that initializes a local C array with the double-formatted 65536 samples of a one-period sine function;

— `double sinwaveform(int i)`, that returns the `i`-th value in this local sampled sine C array;

— `double floor_db(double d)` and `int floor_int(double b)`, that return the floor value of `d` in either `double` or `int` format;

— and, finally, `int main()`, that calls `init_sinwaveform()`, provides the initial frequency value of 440 Hz to the Esterel input `I` in the `Osc` module via a call to `Osc_I_I(440.0)` and then calls `Osc()` to initiate Esterel processing. These last two functions are generated by the Esterel compiler.

— Esterel modules specify signals and signal computations that operate on the occurrence of events, including `ticks` issued by a logical clock.

— Esterel supports both internal or external signals. For instance, `Osc` defines internal signals `index`, `phase` (via Module `Phase`) and `sample` (via Module `RdTable`). As for external signals, frequency values appear on `I` while audio samples are output on `O`.

— Esterel signal computations `emit` values of interest on outputs from values read on inputs, using the `?I` notation. All computations are assumed to be performed at each tick in zero time.

— The keyword `every` is syntactic sugar for an event-controlled looping statement. For instance, each time a new frequency value appears on `I` in `Osc`, Module `Phase` is run in a separate thread, with appropriate bindings for input `freq` and output `phi`, to get the proper sequences of array indices in the `phase` signal while an infinite `loop` thread emits on `O` the samples obtained via the `RdTable` module, using the `index` and `phase` signals.

— `Phase` illustrates another aspect of Esterel, namely the concept of variables storing values. Here, the `index` variable takes, at each time tick, the successive values of the sample indices, while `preindex` keeps track of the previous value[15]. The successive indices of the samples in the sine wavetable appropriate for yielding a sine of the given frequency `freq` are emitted on the `phi` signal. The internal `step` signal is used to broadcast the sample index increment.

— The `immediate` keyword indicates that the corresponding code is evaluated immediately, even when the tick is already present, as is the case when the program starts.

5.4. Lucid Synchrone

Presentation . The following position statement is taken from the official Lucid Synchrone site, http://www.di.ens.fr/~pouzet/lucid-synchrone.

> Lucid Synchrone is an experimental language for the implementation of reactive systems. It is based on the synchronous model of time as provided by Lustre combined with some features from ML languages. The main characteristics of the language are the following:
>
> — It is a strongly typed, higher-order functional language managing infinite sequences or streams as primitive values. These streams are used for representing input and output signals of reactive systems and are combined through the use of synchronous data-flow primitives *à la* Lustre.
>
> — The language is founded on several type systems (e.g., type and clock inference, causality and initialization analysis) which statically guarantee safety properties on the generated code...
>
> — The language is built above Objective Caml used as the host language. Combinatorial values are imported from Objective Caml and programs are compiled into Objective Caml code. A simple module system is provided for importing values from the host language or from other synchronous modules.
>
> — It allows to combine data-flow equations with complex state machines (Mealy and Moore machines with various forms of transitions). This allows to describe mixed systems or Mode-automata as originally introduced by Maraninchi & Rémond.
>
> — Data-types (product types, record types and sum types) can be defined and accessed through pattern matching constructions.

Oscillator. The Lucid Synchrone implementation `osc.ls` of the oscillator can be found below. We tested Version 3.0b. The `osc` execution result is an infinite sample stream, output on the default Lucid Synchrone output port.

```
let static tablesize = 65536
```

[15]The `pre` operator on Esterel signals could have been used here too.

```
let static samplingfreq = 44100
let static twopi = 6.28318530717958623
let ftablesize = float_of_int tablesize

let static sinwaveform = Array.make tablesize 0.0
let static gen_sin () =
  let rec feed i =
    match i with
      | 0 -> ()
      | i ->
        (Array.set sinwaveform (i-1)
           (sin((float_of_int (i-1)) *. twopi /. ftablesize));
         feed (i-1))
    end
  in feed tablesize
let static sidefeeding = gen_sin ()

let fracpart x = x -. floor(x)

let node phase freq =
  let rec index = 0.0 ->
    fracpart((freq /. (float_of_int samplingfreq)) +. pre(index)) in
  int_of_float (ftablesize *. index)

let rdtable tableindex = Array.get sinwaveform tableindex

let node osc freq = rdtable(phase(freq))
```

Design Notes

— Lucid Synchrone imports most of its value and typing constructs from Objective Caml (OCaml), a mostly-functional object-oriented language in which (possibly recursive) functions are first-class values [Leroy et al. 2010]. It is also inspired by Lustre signal processing concepts.

— Following Lustre, Lucid Synchrone adds to OCaml the notion of a *node*, defined via `let node` declarations. These nodes are used to manipulate streams, which are infinite sequences of values linked to a particular clock.

— In a stream definition i -> s, i denotes the default, first value of the stream while s is the inductive definition of a stream element. A reference to the previous stream value is allowed using the `pre` operator.

— `Array` is an OCaml module, used here within Lucid Synchrone, that provides standard operations to define (`make`) or manipulate (`set` and `get`) array elements.

— In OCaml, integer operators use the traditional syntax (such as $*$ for multiplication), while floating-point constructs use a different notation, via the addition of a dot (`"."`) suffix to the integer notation.

— The execution of the `osc` node, given a frequency `freq` argument, yields a stream based on the implicit mapping of the `rdtable` function to the stream `phase(freq)`. As in Lustre, the `phase` node defines the local stream `index`, later used to iteratively access the `sinwaveform` array.

— Lucid Synchrone allows the definition of constants via the `let static` binding; it is used here to populate, in a simple recursive manner via `gen_sin`, the constant array `sinwaveform`.

5.5. OpenMP Stream Extension

Presentation. The OpenMP Stream Extension comes as a GCC CVS branch, at http://gcc.gnu.org/viewcvs/branches/omp-stream. Below follows its position statement.

> The stream-computing extension to OpenMP enables the expression of flow dependences between OpenMP tasks. This allows to statically specify the program's dynamic task graph, where tasks are connected through streams that transparently privatize the data. The programming model is conducive to making relevant data-flow explicit and to structuring programs in ways that allow simultaneously exploiting pipeline, data and task parallelism. Stream computations help reduce the severity of the memory wall in two complementary ways: (1) decoupled producer/consumer pipelines naturally hide memory latency; and (2) they favor local, on-chip communications, bypassing global memory.
>
> This extension provides dataflow semantics close to Kahn process networks and guarantees functional determinism, a major asset in the productivity race. In contrast with common streaming frameworks, the communication patterns can be dynamic, while preserving the determinism of arbitrarily merging and splitting data streams. The GCC prototype implementation of the OpenMP extension for stream-computing has been shown to be efficient to exploit mixed pipeline- and data-parallelism, even in dynamic task graphs [Pop and Cohen 2011]. It relies on compiler and runtime optimizations to improve cache locality and relies on a highly efficient lock-free and atomic operation-free synchronization algorithm for streams.

We need to emphasize here that the OpenMP Stream Extension is particularly interesting for our survey since, built on top of an imperative language (C) extended with asynchronous parallel constructs (OpenMP), this language extension is not strictly synchronous. Yet it offers to programmers the ability to perform parallel signal processing operations that loosely adhere to the synchronous hypothesis. Indeed, all stream operations are specified to be deterministic and to not require explicit synchronization actions. This illustrates how synchronous-like operations could be added to other existing traditional languages.

Oscillator. The OpenMP Stream Extension [Pop 2011] implementation `osc.c` of the oscillator can be found below. We tested a prototype directly with its author, Antoniu Pop (from MINES ParisTech's Computer Science Research Center). The `osc` execution result is the printing on the standard output of the first 200 values of the expected sinusoid.

```
#include <stdlib.h>
#include <stdio.h>
#include <math.h>

#define freq 440
#define outputsize 200
#define twopi 6.28318530717958623

static inline float fracpart (float x) { return x - floor (x); }

int main (int argc, char **argv) {
  int i;
  int tablesize = 1 << 16;
  int samplingfreq = 44100;
  float *sinwaveform = (float *) malloc (tablesize * sizeof (float));

  for (i = 0; i < tablesize; ++i)
    sinwaveform[i] = sin (((float) i) * twopi / ((float) tablesize));

  #pragma omp parallel num_threads (2) default (none)
                       shared (tablesize, sinwaveform, samplingfreq) {
```

```
    #pragma omp single {
      float f_sf_ratio, index, phase;
      float frac_add = 0.0; int i = 0;
      while (i++ < outputsize) {
        #pragma omp task shared (samplingfreq) output (f_sf_ratio)
                         num_threads (2) {
          f_sf_ratio = ((float) freq) / ((float) samplingfreq);
        }
        #pragma omp task input (f_sf_ratio) output (index)
                         shared (frac_add) {
          frac_add = fracpart (frac_add + f_sf_ratio);
          index = frac_add;
        }
        #pragma omp task shared (tablesize) input (index)
                         output (phase) {
          phase = index * tablesize;          // A stream processor
        }
        #pragma omp task shared (sinwaveform) input (phase)
                         shared (stdout) {
          fprintf (stdout, "%f \t %f\n",
                   sinwaveform[(int)phase], phase);
        }}}}
  return 0;
}
```

Design Notes

— OpenMP Stream Extension is an upward-compatible extension of the OpenMP standard [Dagum and Menon 1998], which extends sequential languages with options for parallel execution. OpenMP has multiple language bindings, and its C variant uses `#pragma omp` C preprocessor-like directives to describe thread-parallel tasking.

— The `parallel` pragma is an OpenMP-specific directive used to open a parallel section in which multiple tasks may be used on up to `num_threads` parallel threads. A separate task is launched for each `omp task` pragma, running the statement following it in parallel on one of these threads. The `single` pragma enforces its following sequence of code to be run by only one thread (which runs here the `while` loop that starts all required tasks).

— All variables declared within a parallel block are local to each task by default; global variables accessed by a particular parallel construct have to be listed in the `shared` parameter.

— OpenMP Stream Extension extends the `task` pragma, used to specify an OpenMP parallel task, with the `input` and `output` parameters that introduce stream processing into OpenMP. Variables such as `index` and `phase` are promoted to streams; memory is allocated and managed automatically by the OpenMP Stream Extension runtime in a pipeline fashion, hiding to the programmer the implementation details. The `input/output` arguments can be seen as queues on which the pipeline stages, implemented as separate parallel tasks, synchronize and exchange signal data.

— A stream parallel task such as the one defining `phase` (see commented line) processes its input signals to yield output data accordingly.

6. DISCUSSION

We look, in this section, at some of the issues raised by our implementations of `osc`. In particular, we discuss (1) the differences between DSL and general-purpose languages, (2) the subtle differences in the notions of time and signals these formalisms introduce, (3) how the synchronous hypothesis affects language design, (4) the way

these languages handle aggregate data structures such as arrays, and finally (5) the integration of asynchrony.

6.1. The DSL vs. General-Purpose Languages Debate

In this work, we surveyed 10 languages, 5 specific to computer music applications and 5 general synchronous languages. Each of these languages provides in one way or another answers to the same design questions, such as how to manage time or what are signals supposed to represent for the problems at hand. Yet, the final language design decisions, summarized in Table IV, vary widely, mostly along the line of whether the corresponding language is intended to be used in a somewhat limited application domain, i.e., strives to be a music-specific DSL, or able to tackle a wide range of time-constrained problems, i.e., is a general-purpose synchronous language.

Table IV. Design concepts comparison summary

	Music programming languages	Synchronous languages
Code size	much shorter	longer
Time	mostly a hardware notion	mostly a logical notion
Signals	simple tick mappings	abstract and complex clocks
Layers	low-level and interactive high-level	synchronous layer and GALS[a]

[a]See Section 6.5 for GALS.

Of course, the DSL vs. General-Purpose separation line is not enough to automatically imply which particular programming traits a given language should adopt. To get a feel for the spectrum of notions spanned by this survey, we summarize in Table V the main features of these programming tools: (1) their core computing paradigm, (2) the way they specify, in particular for recursive definitions, the memory values at previous time slots, via delays and initializations, and (3) their design approach of parallelism specification at the most abstract level, thus not taking in account any hardware constraints.

Table V. Salient design points of surveyed languages

	Paradigm	Delay	Initialization	Parallelism
Csound	orchestra and score	`delay`	2nd arg.	orchestra
SuperCollider	object-oriented	`DelayN`	3rd arg.	implicit
Pure Data	visual	`delread~`	2nd arg.	graph
ChucK	on-the-fly	`Delay`	`.delay`	implicit
Faust	functional	$\sim$	0	`,` or `par`
Signal	relational	*signal*`$`	`init`	implicit
Lustre	equational	`pre`	`->`	implicit
Esterel	imperative	`pre(?`*signal*`)`	`init`	`||`
Lucid Synchrone	functional	`pre`	`->`	implicit
OMP Stream	imperative	window	explicit	`omp task`

One obvious result of our work is that using a single simple yet significant application to illustrate the expressiveness power of various programming languages provides an interesting and practical point of view for software development tool selection. In particular, choosing an audio application as our main running test case yields a specific example of the intrinsic value of Domain Specific Languages as a general and practical approach to the software productivity wall [Mernik et al. 2005; Van Deursen et al. 2000]. Indeed and not surprisingly, we were able to use all these tools to get the work done, except for a technical limitation of the open-source version of Lustre we used. Yet, and as anyone could have expected, all programs using music-focused DSLs are

much shorter, and also readable, than the ones based on general-purpose synchronous languages.

A positive by-product of the conciseness of DSL-based programs is that they are consequently probably more often correct than those based on the other formalisms. Indeed, even though the idea that the number of bugs introduced in a particular program is mainly a function of the number of lines of code and is rather independent of the programming language used may be mostly folklore, we feel that our analysis illustrates in a very concrete manner that conciseness, and hopefully then lack of software defects, clearly lies within the field of DSLs.

6.2. Logical and Physical Times

Although time is, as we mentioned in Section 2, the core concept that structures the definition of all the languages used in this use case study, it is obvious that music-oriented and reactive systems have a somewhat different view of what this notion means. In the traditional synchronous programming world, time is mostly a logical notion, around which computations are scheduled; indeed, multiple clocks can even be defined, e.g., via the `^=` symbol in Signal. For music aficionados, time is a hardware, physical notion deeply linked to the speed at which sound is sampled by input and output converters; the key notion here is the "sampling rate", e.g., via the `SR` predefined identifier in Faust. In some sense, music languages, as DSLs, are more closely linked to the practical matters at hand than the more abstract, hence more general, traditional synchronous programming languages.

Consequently, the notion of what a signal is varies also in the two communities, even though both use the same foundation, i.e., the concept of time. Following a more pragmatic approach, music synchronous languages view signals as mappings from regularly-spaced, sampling rate-sequenced time ticks to values. Traditional languages have to deal with more complex clocks, for instance where time events might even be absent, which leads to more abstract notions of signals. In music applications, all values of a sampled signal are defined (barring computing errors such as a 1/0 division), while general signals may yield undefined values for some time events, and such undefined values are first-class in these languages [Benveniste et al. 2003].

The apparently limited approach of what time and signals must be in music applications is mitigated by the fact that there is a strong tendency in this community to address timing issues as a two-tiered problem. First, a low-level synchronous layer, synchronized to the audio rate, deals with concrete and predictive sampled signals. Then, a higher-level control layer, in fact mostly asynchronous (Csound being a notable exception, using the `ksmps` variable we discuss in Section 6.3), schedules the low-level activities in response to the user. These two strata are often embedded in the same language or environment, using one scheduler at audio rate (typically 44,100 Hz, with high priority and low latency based on buffering techniques) and another one at a much lower "control rate", usually managing MIDI events (medium priority and latency) and GUI objects (low priority and high latency). On the other hand, traditional synchronous languages tend to address only the issues relevant to the first class of problems, using either a *sample-driven* execution scheme as in Lustre or an *event-driven* one as in Signal, and rely on a different programming paradigm to link the synchronous modules together, possibly using a GALS[16] approach [Teehan et al. 2007]. Thus, they usually require more general and flexible notions for time and signals than the ones found in audio languages, since they do not have to follow the fundamental rhythm of the otherwise primary audio sampling rate.

[16] GALS stands for "Globally Asynchronous, Locally Synchronous".

6.3. The Synchronous Hypothesis in Computer Music

Synchronous languages, because of their more abstract and logical view of time and signals, are formally defined through complex mathematical semantic models [Manna and Pnueli 1995; Schneider 2004]. These formal specifications are moreover of key importance given the domains these languages target, i.e., within strongly reactive and very often mission-critical environments. Music, on the other hand, can, and has to, deal with more "soft" constraints: the notion of truth is more in the ear of the listener/composer than in the strict structure of a mathematical proof. Of course, the human ear is quite a subtle device, and professional listeners have been shown to be quite sensitive to even very small differences in two audio signals; even one missing sample may induce a dramatic sound artefact that any listener will hear. Moreover, even for music applications, a trend is appearing, which calls for more assurance in the fidelity of the audio processing methods, in particular when one wishes to address the issue of long-term and exact preservation of the world musical heritage [Guercio et al. 2007; Bachimont et al. 2003; Barkati et al. 2011].

These two approaches regarding this core notion of time in the synchronous layer lead to different approaches to the compilation process. Traditional synchronous languages are more specification-oriented than audio languages; the programmer provides equations defining clock and signal values, relying on the compiler to implement them in efficient sequences of computations. Audio/music frameworks have to deal less with the issue of reifying relationships between logically synchronized computations than with the efficient implementation of explicitly synchronized processes (see for instance Pure Data connected graphs or Faust functional expressions).

These two different ways of addressing the issue of correctness, i.e., the ability of performing a given set of computations under timing constraints, have had a significant impact on language design. Synchronous languages adopt the synchronous hypothesis, which mathematically ensures that all specifications will be met, both in the computation and time domains; no event will ever be lost. On the contrary, in the music realm, the synchronous hypothesis has never been a well-identified key design principle; audio samples may, in practice, be lost. To lower the number of such occurrences, music language designers have adopted a much more pragmatic approach: audio processing is not performed on a sample basis, but rather on buffers of multiple audio and control samples, thus lowering the I/O overhead associated to per-sample computations. This hands-on design decision is in fact visible in many music languages we surveyed above; for instance, our Csound example defines the variable `ksmps`, which specifies the number of elements to be used to size the control buffers used internally to process sound. Such low-level buffer management does not need to appear in synchronous languages. Note however that, as as consequence of this more "relaxed" approach to correctness handling, music languages have been sometimes explicitly equipped to deal with event loss; the outcomes of this research might, interestingly, be of use by the synchronous language community, having to handle hardware faults that may fall outside of the synchronous hypothesis protection.

6.4. Dynamic Array Management

Table lookup is about performance: computer music makes an intensive use of wavetables to avoid the expensive computation of trigonometric functions like sine functions for each sample at given audio rates, typically 44,100 times by second, for audio synthesis (wavetable synthesis, waveshaping, etc.). It is noticeable that most of the music programming languages studied here, except Faust, borrow the idea of the GEN routines introduced in Csound; they are used as data generators to fill so-called *function tables* [Boulanger et al. 2000]. For instance, our oscillator uses the *GEN10* routine to

fill the oscillator sample table at initialization time with a sum of sinusoids (only one here); this table is then read using a wrap-around lookup process. Faust provides support for load-time table initialization and reading operation via the `rdtable` ternary function.

Our study reveals that general-purpose synchronous languages are often poorly

Table VI. Table support

	Initialization	Dynamic Access
Csound	`f1 0 65536 10 1 ; GEN10`	`oscil p4, p5, 1`
SuperCollider	`b.sine1`	`{OscN.ar(b,440,0,1)}.play`
Pure Data	`sinesum $1 1`	`tabread~ myarray`
ChucK	`[1.] => g10.coefs`	`Gen10 g10 => dac`
Faust	`rdtable` 1st and 2nd arg.	`rdtable` 3rd arg.
Signal	`[i to (size-1):sin(...)]`	`sinwaveform[tableindex]`
Lustre v4	*external C code*	*imported function*
Lucid Sync.	OCaml `Array.make` and `.set`	OCaml `Array.get`
Esterel v5	*external C code*	*imported function*
OMP Stream	C code	C code

equipped to support tables. Of course, they cannot be expected to provide music-specific GEN-like routines, but, more surprisingly at first sight, most of them simply do not handle array data structures, as is the case with the Esterel and Lucid Synchrone languages in the versions we used. Furthermore, the Lustre version we tested do provide an array syntax, but only as syntactic sugar for variables numbering, not for dynamic array access. Of course, it is generally possible to handle array initialization and dynamic access by importing foreign functions, C functions in most languages, as we did in our Esterel implementation of the oscillator, or OCaml functions in Lucid Synchrone, as we showed. The underlying reason for not handling arrays in the languages themselves is often the difficulty of ensuring that the synchronous time and memory constraints are still enforced, which is crucial for critical synchronous applications.

Commercial and more recent versions of Lustre and Esterel do handle arrays. Nevertheless, our survey suggests that the designers of synchronous languages could look at the GEN-like mechanism inspired by music programming languages as a safe strategy for the introduction of array data structures in these formalisms. Another possible approach to ensure a mathematically-correct integration of arrays and synchronous constraints is to couple in a single analysis the rates of signals with the size of the elements they convey (see for instance [Jouvelot and Orlarey 2011]).

6.5. Event Management

Our survey focused on the audio signal processing part of the computer music domain, since audio DSP shares obvious features with synchronous applications, among which time and signal concepts – even though these are subtly different in both fields, as we have shown. The DSP part corresponds to the *sampled scheme* of evaluation of synchronous languages, where the main loop handles each sample, as opposed to their *control scheme*, where the main loop manages each interaction event. Synchronous languages often rely on a two-tier strategy to handle the integration of asynchrony, for instance via a GALS approach; most computer music languages are, them, inherently hybrid along the synchronous/asynchronous separation line.

Indeed, in addition to the DSP part, most music programming languages also embed event management for the inherently asynchronous occurences of musical notes and interactive remote control of parameters. Asynchronous event management is usually handled using message passing standards such as MIDI or OSC with no time tags. It is typically implemented via fixed-size event FIFOs or dedicated schedulers that

run at a lower priority than the DSP one: such implementations may lead to non-deterministic effects when overflows occur. Note that this approach can also be applied to user-induced asynchronous events such as I/O operations, as is the case for SuperCollider.

On the other hand, dealing with these asynchronous events can be performed by embedding them within a synchronous framework, as is done in Csound, Pure Data or ChucK. They are here synced on a clock running at a specific *control* rate, which is usually lower than the audio rate since these events occur less frequently than audio samples. This approach may lower the computing load but lead to a loss of precision or even data loss when events are issued at too fast a pace; such a behavior may occur, for instance, in SuperCollider, when events are synchronized on `AppClock` instead of `SystemClock`. Such a synchronous handling can also be used when dealing with more time-based events such as timeouts, which may also be synced to the audio or control rates.

A challenging idea suggested by our survey would be to study if and how musical programming languages could improve their event management processes by borrowing from the mathematically-well founded control handling of the sophisticated synchrony traits introduced by synchronous languages, and thus benefit from their formal consistency.

7. CONCLUSION

We performed a practical, use case-oriented survey of ten key music-specific and general-purpose synchronous programming languages, implementing in each of them a simple yet significant audio processing algorithm, namely a frequency-parameterized oscillator. We believe this survey provides the first bridge between two mature and widely successful computing fields, the more than 50-year old computer music domain and the 40-year old niche of synchronous programming languages. Our work showed that the wide variety of existing music and synchronous languages leads to a large spectrum of program sizes and styles, even for the simple case of the oscillator, which was our running example throughout the article.

Our application-oriented comparison work, and the discussion points it led to, can be of use to language designers interested in the relationships and opportunities for mutual interactions between synchronous and computer music language designs. Moreover, programmers will get from our survey a better feeling for what the two audio and synchronous families of programming languages have to offer, both in common and on their own, extending their views on how explicit timing issues can be dealt with at various levels of abstraction.

Our present work has focused on programming language design issues. It would be interesting to see whether our findings regarding DSLs' benefits can be leveraged to more complex use case applications. Future work needs also to address the implementation, performance, environment integration and event management aspects of such a comparison, since these factors are also key in the decisions leading to the choice of a particular language or language paradigm in software projects. Finally, the choice of our audio oscillator running example creates some bias in our comparison of general-purpose and music-specific languages; this calls for an other, symmetric study, which would use a typical synchronous application, such as for instance ABRO [Berry 2000], to provide a balanced assessment of both programming paradigms.

ACKNOWLEDGMENT

The authors thank Daniel Gaffé, Léonard Gérard, Yann Orlarey, Cédric Pasteur, Antoniu Pop, Marc Pouzet, Annie Ressouche, Xavier Rival and Valérie Roy for their help, Laure Gonnord and Jean-Pierre Talpin for

their kind feedback, and the reviewers for their detailed comments and suggestions. The motivation for this study comes from the French ASTREE[17] ANR 2008 CORD 003 01 research project, which addresses the preservation issues of synchronous programs in computer music, using the Faust real-time audio processing language as the core foundation for such preservation efforts.

References

K. Altisen, F. Maraninchi, and D. Stauch. 2006. Aspect-oriented programming for reactive systems: Larissa, a proposal in the synchronous framework. *Science of Computer Programming* 63, 3 (2006), 297–320.

X. Amatriain, P. Arumi, and D. Garcia. 2006. CLAM: A framework for efficient and rapid development of cross-platform audio applications. In *Proceedings of the 14th Annual ACM International Conference on Multimedia*. ACM, 951–954.

C. André. 1996. Representation and analysis of reactive behaviors: A synchronous approach. In *CESA'96 IMACS Multiconference: computational engineering in systems applications*. 19–29.

M. Antonotti, A. Ferrari, A. Flesca, and A. Sangiovanni-Vincentelli. 2000. JESTER: An Esterel based reactive Java extension for reactive embedded systems. In *Forum on Specification & Design Languages*.

G. Assayag, C. Rueda, M. Laurson, C. Agon, and O. Delerue. 1999. Computer-assisted composition at IRCAM: from PatchWork to OpenMusic. *Computer Music Journal* 23, 3 (1999), 59–72.

B. Bachimont, J.F. Blanchette, A. Gerzso, A. Swetland, O. Lescurieux, P. Morizet-Mahoudeaux, N. Donin, and J. Teasley. 2003. Preserving interactive digital music: a report on the MUSTICA research initiative. In *Proceedings of the Third International Conference on Web Delivering of Music*. IEEE, 109–112.

A.D. Baker, T.L. Johnson, D.I. Kerpelman, and H.A. Sutherland. 1987. GRAFCET and SFC as Factory Automation Standards Advantages and Limitations. In *American Control Conference*. IEEE, 1725–1730.

R. Baker and L.A. Hiller. 1963. MUSICOMP: MUsic Simulator-Interpreter for COMpositional Procedures for the IBM 7090. (1963).

F. Balarin. 1997. *Hardware-software co-design of embedded systems: the POLIS approach*. Vol. 404. Springer Netherlands.

K. Barkati, D. Fober, S. Letz, and Y. Orlarey. 2011. Two Recent Extensions to the FAUST Compiler. In *Proceedings of the Linux Audio Conference*.

A. Benveniste and G. Berry. 1991a. Another look at real-time programming. In *Special Section of the Proceedings of the IEEE*, Vol. 79.

A. Benveniste and G. Berry. 1991b. The synchronous approach to reactive and real-time systems. *Proc. IEEE* 79, 9 (1991), 1270–1282.

A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. Le Guernic, and R. De Simone. 2003. The synchronous languages twelve years later. *Proc. IEEE* 91, 1 (2003), 64–83.

G. Berry. 2000. *The Esterel v5 Language Primer: Version v5_91*. Centre de mathématiques appliquées, Ecole des mines and INRIA.

G. Berry and L. Cosserat. 1985. The ESTEREL synchronous programming language and its mathematical semantics. In *Seminar on Concurrency*. 389–448.

R. Bharadwaj. 2002. SOL: A verifiable synchronous language for reactive systems. *Electronic Notes in Theoretical Computer Science* 65, 5 (2002), 140–154.

P. Bjesse, K. Claessen, M. Sheeran, and S. Singh. 1998. Lava: hardware design in Haskell. In *Proceedings of the Third ACM SIGPLAN International Conference on Functional Programming*. ACM, 174–184.

G. Boudol. 2004. ULM: A core programming model for global computing. *Programming Languages and Systems* (2004), 234–248.

R.C. Boulanger and others. 2000. *The Csound Book*. Vol. 309. MIT Press.

F. Boussinot. 1991. Reactive C: An extension of C to program reactive systems. *Software: Practice and Experience* 21, 4 (1991), 401–428.

F. Boussinot and R. De Simone. 1996. The SL synchronous language. *Software Engineering, IEEE Transactions on* 22, 4 (1996), 256–266.

F. Boussinot and J.F. Susini. 1998. The SugarCubes tool box: a reactive Java framework. *Software: Practice and Experience* 28, 14 (1998), 1531–1550.

J. Bresson, C. Agon, and G. Assayag. 2009. Visual Lisp/CLOS programming in OpenMusic. *Higher-Order and Symbolic Computation* 22, 1 (2009), 81–111.

[17] ASTREE stands for "*Analyse et synthèse de traitements temps réel*", i.e., "Analysis and Synthesis of Real-Time Processes".

C. Cadoz, A. Luciani, and J.L. Florens. 1993. CORDIS-ANIMA: a Modeling and simulation system for sound and image synthesis: the general formalism. *Computer Music Journal* (1993), 19–29.

P. Caspi, A. Curic, A. Maignan, C. Sofronis, S. Tripakis, and P. Niebert. 2003. From Simulink to SCADE/Lustre to TTA: a layered approach for distributed embedded applications. *ACM SIGPLAN Notices* 38, 7 (2003), 153–162.

P. Caspi, D. Pilaud, N. Halbwachs, and J. A Plaice. 1987. LUSTRE: A declarative language for programming synchronous systems. In *Conference Record of the 14th Annual ACM Symp. on Principles of Programming Languages*.

P. Caspi and M. Pouzet. 1996. Synchronous Kahn networks. In *ACM SIGPLAN Notices*, Vol. 31. 226–238.

N. Castagné and C. Cadoz. 2002. GENESIS : a friendly musician-oriented environment for mass-interaction physical modeling. In *Proceedings of the International Computer Music Conference*. Goteborg, Suede, 330–337. http://hal.archives-ouvertes.fr/hal-00481717/en/

A. Chaudhary, A. Freed, and M. Wright. 2000. An Open Architecture for Real-time Music Software. *Proceedings of the International Computer Music Conference* (2000).

M. Compton. 2000. SCOOP: An investigation of concurrency in Eiffel. *Master's thesis, Department of Computer Science, The Australian National University* (2000).

P.R. Cook and G. Scavone. 1999. The synthesis toolkit (STK). In *Proceedings of the International Computer Music Conference*. 164–166.

L. Dagum and R. Menon. 1998. OpenMP: an industry standard API for shared-memory programming. *Computational Science & Engineering, IEEE* 5, 1 (1998), 46–55.

R.B. Dannenberg. 1997. Machine tongues XIX: Nyquist, a language for composition and sound synthesis. *Computer Music Journal* 21, 3 (1997), 50–60.

F. Déchelle, R. Borghesi, M. De Cecco, E. Maggi, B. Rovan, and N. Schnell. 1999. jMax: an environment for real-time musical applications. *Computer Music Journal* 23, 3 (1999), 50–58.

F.X. Dormoy. 2008. Scade 6: a model based solution for safety critical software development. In *Proceedings of the 4th European Congress on Embedded Real Time Software (ERTS'08)*. 1–9.

G. Eckel, F. Iovino, and R. Caussé. 1995. Sound synthesis by physical modelling with Modalys. In *Proc. International Symposium on Musical Acoustics*. 479–482.

J. Ellsberger, D. Hogrefe, and A. Sarma. 1997. *SDL: formal object-oriented language for communicating systems*. Prentice Hall.

A. Gamatié. 2009. *Designing embedded systems with the Signal programming language: synchronous, reactive specification*. Springer Verlag.

E. Gaudrain and Y. Orlarey. 2003. *A Faust Tutorial*. Technical Report. Grame, Lyon.

T. Gautier, P. Le Guernic, and L. Besnard. 1987. Signal: A declarative language for synchronous programming of real-time systems. In *Functional Programming Languages and Computer Architecture*. Springer, 257–277.

N. Gehani and K. Ramamritham. 1991. Real-time concurrent C: A language for programming dynamic real-time systems. *Real-Time Systems* 3, 4 (1991), 377–405.

J.-L. Giavitto. 1991. A Synchronous Data-Flow Language for Massively Parallel Computers. In *Proceedings of the International Conference on Parallel Computing'91*, D. J. Evans, G. R. Joubert, and H. Liddell (Eds.). London, UK, 391.

A. Girault. 2005. A survey of automatic distribution method for synchronous programs. In *International Workshop on Synchronous Languages, Applications and Programs, SLAP*, Vol. 5.

M. Guercio, J. Barthélemy, and A. Bonardi. 2007. Authenticity issue in performing arts using live electronics. In *Sound and Music Computing Conference Proceedings*.

N. Halbwachs. 1993. *Synchronous Programming of Reactive Systems*. Kluwer Academic Pub.

N. Halbwachs. 2005. A synchronous language at work: the story of Lustre. In *Proceedings of the 2nd ACM/IEEE International Conference on Formal Methods and Models for Co-Design*. IEEE Computer Society, 3–11.

D. Harel. 1987. Statecharts: A visual formalism for complex systems. *Science of computer programming* 8, 3 (1987), 231–274.

D. Harel and A. Pnueli. 1985. *On the development of reactive systems*. Weizmann Institute of Science, Dept. of Computer Science.

L.A. Hiller and R.A. Baker. 1964. Computer Cantata: A study in compositional method. *Perspectives of New Music* 3, 1 (1964), 62–90.

P. Hudak, T. Makucevich, S. Gadde, and B. Whong. 1996. Haskore music notation: An algebra of music. *Journal of Functional Programming* 6, 3 (1996), 465–483.

IEEE standard. 1988. VHDL language reference manual. *IEEE Std* (1988), 1076–1987. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=26487

M. Ingham, R. Ragno, and B. Williams. 2001. A Reactive Model-based Programming Language for Robotic Space Explorers. In *Proceedings of ISAIRAS-01*.

OSC Initiative. 2006. IEEE Standard SystemC Language Reference Manual. *IEEE Computer Society* 2002, March (2006).

Y. Ishikawa, H. Tokuda, and C.W. Mercer. 1992. An object-oriented real-time programming language. *Computer* 25, 10 (1992), 66–73.

Pierre Jouvelot and Yann Orlarey. 2011. Dependent vector types for data structuring in multirate Faust. *Comput. Lang. Syst. Struct.* 37 (July 2011), 113–131. Issue 3. DOI:http://dx.doi.org/10.1016/j.cl.2011.03.001

M. Laurson and J. Duthen. 1989. PatchWork, a graphical language in PreForm. In *Proceedings of the International Computer Music Conference*. San Francisco, CA, 172–173.

M. Laurson, M. Kuuskankare, and V. Norilo. 2009. An overview of PWGL, a visual programming environment for music. *Computer Music Journal* 33, 1 (2009), 19–31.

L. Lavagno and E. Sentovich. 1999. ECL: a specification environment for system-level design. In *Proceedings of the 36th Annual ACM/IEEE Design Automation Conference*. ACM, 511–516.

V.E.P. Lazzarini. 2000. The SndObj Sound Object Library. *Organised Sound* 5, 1 (2000), 35–49.

P. Le Guernic, J. P Talpin, and J. C Le Lann. 2003. Polychrony for system design. *Journal of circuits systems and computers* 12, 3 (2003), 261–304.

X. Leroy, D. Doligez, A. Frisch, J. Garrigue, D. Rémy, and J. Vouillon. 2010. *The Objective Caml system release 3.12 Documentation and user's manual*. Technical Report. INRIA.

G. Loy and C. Abbott. 1985. Programming languages for computer music synthesis, performance, and composition. *ACM Computing Surveys (CSUR)* 17, 2 (1985), 235–265.

F. Mallet and C. André. 2009. On the semantics of UML/MARTE clock constraints. In *2009 IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing*. IEEE, 305–312.

L. Mandel and M. Pouzet. 2005. ReactiveML: a reactive extension to ML. In *Proceedings of the 7th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*. ACM, 82–93.

Z. Manna and A. Pnueli. 1995. *Temporal verification of reactive systems: safety*. Vol. 2. Springer Verlag.

F. Maraninchi. 1990. Argonaute: graphical description, semantics and verification of reactive systems by using a process algebra. In *Automatic Verification Methods for Finite State Systems*. 38–53.

F. Maraninchi. 1991. The Argos language: Graphical representation of automata and description of reactive systems. In *IEEE Workshop on Visual Languages*.

M.V. Mathews, J.E. Miller, F.R. Moore, J.R. Pierce, and J.C. Risset. 1969. *The technology of computer music*. The MIT Press, Boston.

J. McCartney. 1996. SuperCollider, a new real time synthesis language. In *Proceedings of the International Computer Music Conference*. International Computer Music Association, 257–258.

M. Mernik, J. Heering, and A.M. Sloane. 2005. When and how to develop domain-specific languages. *ACM Computing Surveys (CSUR)* 37, 4 (2005), 316–344.

R. Milner. 1980. *A calculus of communicating systems*. Vol. 92. Springer-Verlag.

C. Nilson. 2007. Live coding practice. *Proceedings of New Interfaces for Musical Expression (NIME)* (2007).

Y. Orlarey, D. Fober, and S. Letz. 1997. Elody: A Java + MidiShare based music composition environment. In *Proceedings of the International Computer Music Conference*. Thessaloniki, Greece.

Y. Orlarey, D. Fober, and S. Letz. 2002. An Algebra for Block Diagram Languages. In *Proceedings of International Computer Music Conference*. 542–547.

Y. Orlarey, D. Fober, and S. Letz. 2004. Syntactical and Semantical Aspects of Faust. *Soft Computing-A Fusion of Foundations, Methodologies and Applications* 8, 9 (2004), 623–632.

Y. Orlarey, D. Fober, and S. Letz. 2009. FAUST: an Efficient Functional Approach to DSP Programming. In *New Computational Paradigms for Computer Music*, Assayag G. and Gerzso A. (Eds.). IRCAM/Delatour France.

Claire Pagetti, Julien Forget, Frédéric Boniol, Mikel Cordovilla, and David Lesens. 2011. Multi-task implementation of multi-periodic synchronous programs. *Discrete Event Dynamic Systems* 21, 3 (2011), 307–338. http://hal.inria.fr/inria-00638936

A. Pnueli, O. Shtriehman, and M. Siegel. 1998. Translation validation for synchronous languages. *Automata, Languages and Programming* (1998), 235–246.

A. Pop. 2011. *Leveraging Streaming for Deterministic Parallelization : an Integrated Language, Compiler and Runtime Approach.* Ph.D. Dissertation. MINES ParisTech.

A. Pop and A. Cohen. 2011. A Stream-Computing Extension to OpenMP. In *Proc. of the 6th Intl. Conf. on High Performance and Embedded Architectures and Compilers (HiPEAC'11)*.

M. Puckette. 1991. Combining event and signal processing in the MAX graphical programming environment. *Computer Music Journal* (1991), 68–77.

M. Puckette. 1996. Pure Data: Another Integrated Computer Music Environment. *Proceedings of the Second Intercollege Computer Music Concerts* (1996), 37–41.

M. Puckette. 2002. Max at seventeen. *Computer Music Journal* 26, 4 (2002), 31–43.

M. Puckette. 2007. *The Theory and Technique of Electronic Music*. World Scientific Pub Co Inc.

J.H. Reppy. 1999. *Concurrent Programming in ML*. Cambridge University Press, Cambridge, England.

X. Rodet, Y. Potard, and J.B. Barriere. 1984. The CHANT project: from the synthesis of the singing voice to synthesis in general. *Computer Music Journal* 8, 3 (1984), 15–31.

C. Scaletti. 1987. Kyma: An object-oriented language for music composition. In *Proceedings of the International Computer Music Conference*. 49–56.

E.D. Scheirer and B.L. Vercoe. 1999. SAOL: The MPEG-4 structured audio orchestra language. *Computer Music Journal* 23, 2 (1999), 31–51.

K. Schneider. 2000. A Verified Hardware Synthesis of Esterel Programs. In *Proceedings of the IFIP WG10*. Kluwer, BV, 205–214.

K. Schneider. 2004. *Verification of reactive systems: formal methods and algorithms*. Springer Verlag.

K. Schneider and T. Schuele. 2005. Averest: Specification, verification, and implementation of reactive systems. In *Conference on Application of Concurrency to System Design (ACSD)*.

D. Simon and A. Girault. 2001. Synchronous programming of automatic control applications using OrCAD and Esterel. In *Proceedings of the 40th IEEE Conference on Decision and Control*, Vol. 4. IEEE, 3290–3295.

J.O. Smith III. 2010. *Audio Signal Processing in Faust*. Technical Report. CCRMA.

A. Sorensen. 2005. Impromptu: An interactive programming environment for composition and performance. In *Proceedings of the Australasian Computer Music Conference 2009*.

H. Taube. 1991. Common Music: A music composition language in Common Lisp and CLOS. *Computer Music Journal* (1991), 21–32.

P. Teehan, M. Greenstreet, and G. Lemieux. 2007. A survey and taxonomy of GALS design styles. *Design & Test of Computers, IEEE* 24, 5 (2007), 418–428.

W. Thies, M. Karczmarek, and S. Amarasinghe. 2002. StreamIt: A Language for Streaming Applications. In *International Conference on Compiler Construction*. Grenoble, France. http://groups.csail.mit.edu/commit/papers/02/streamit-cc.pdf

D.E. Thomas and P.R. Moorby. 2002. *The Verilog hardware description language*. Vol. 1. Springer Netherlands.

A. Valle and others. 2007. The SuperCollider Help Book. (2007).

A. Van Deursen, P. Klint, and J. Visser. 2000. Domain-specific languages: An annotated bibliography. *ACM Sigplan Notices* 35, 6 (2000), 26–36.

P. Van Roy. 2009. Programming Paradigms for Dummies: What Every Programmer Should Know. *New Computational Paradigms for Computer Music* (2009).

B. Vercoe. 1992. *Csound: a manual for the audio processing system and supporting programs with tutorials*. Massachusetts Institute of Technology.

B. Vercoe and others. 2007. The Canonical Csound Reference Manual. (2007).

G. Wang and P.R. Cook. 2007. *The ChucK Manual 1.2.1.3*. Princeton University.

G. Wang, P.R. Cook, and others. 2003. ChucK: A Concurrent, On-the-fly Audio Programming Language. In *Proceedings of the International Computer Music Conference*. 219–226.

M. Wright. 2005. Open sound control: an enabling technology for musical networking. *Organised Sound* 10, 03 (2005), 193–200.

D. Zicarelli. 1998. An extensible real-time signal processing environment for Max. In *Proceedings of the International Computer Music Conference*. 463–466.

Received Month Year; revised Month Year; accepted Month Year