

VORPAL: An Extensible and Flexible Middleware for Real-Time Soundtracks in Digital Games

Wilson Kazuo Mizutani and Fabio Kon

Department of Computer Science
University of Sao Paulo
{kazuo,kon}@ime.usp.br
<http://compmus.ime.usp.br/en/vorpal>

Abstract. Real-time soundtracks in games have always faced design restrictions due to technological limitations. The predominant solutions of hoarding prerecorded audio assets and then assigning a tweak or two each time their playback is triggered from game code leaves away the potential of real-time symbolic representation manipulation and DSP audio synthesis. In this paper, we take a first step towards a more robust, generic, and flexible approach to game audio and musical composition in the form of a generic middleware based on the Pure Data programming language. We describe here the middleware architecture and implementation and part of its validation via two game experiments.

Keywords: game audio, game music, real-time soundtrack, computer music, middleware

1 Introduction

Digital games, as a form of audiovisual entertainment, have specific challenges regarding soundtrack composition and design [10]. Since the player's experience is the game designer's main concern, as defended by Schell [13], a game soundtrack may compromise the final product quality as much as its graphical performance. In that regard, there is one game soundtrack aspect that is indeed commonly neglected or oversimplified: its potential as a *real-time*, procedurally manipulated media, as pointed out by Farnell [3].

Even though there is a lot in common between game soundtracks and other more "traditional" audiovisual entertainment media soundtracks (such as Theater and Cinema), Collins [2] argues that there are also unquestionable differences among them, of either technological, historical, or cultural origins. The one Collins points as the most important difference is the deeply nonlinear and interactive structure of games, which make them "*largely unpredictable in terms of the directions the player may take, and the timings involved*". Because of this, many game sounds and music tracks cannot be merely exposed through common playback (as happens with prerecorded media): they also need some form

of procedural control to be tightly knit together with the ongoing narrative of the game. However, this is seldom fully explored. Except in a few remarkable cases, most game musical soundtracks tend to have little real-time connections between what is happening in the game and what is going on with the music, for instance.

The ways with which real-time soundtracks are typically dealt with are poor and do not scale well. Basically, the development team needs to find a common ground for musicians, sound designers, and programmers where they can reach an agreement on how to introduce real-time behaviour into the game source code modules related to audio reproduction and sound assets management. Farnell [3] explains that the common approach is to produce as many assets as needed and then list all event hooks that must go into the game code to timely play the corresponding sounds and music (perhaps with one or two filters applied to the output). This is not only a waste of memory and a disproportional amount of effort, but also a very limited way of designing a game soundtrack. Farnell goes as far as to say that even advanced and automated proprietary middleware systems fail to provide a satisfactory solution, since they “*are presently audio delivery frameworks for prerecorded data rather than real ‘sound engines’ capable of computing sources from coherent models*”.

In our research, we provide an alternative to such excessively asset-driven solutions by empowering the musicians and sound designers with a tool able to express procedurally how a game soundtrack is to be executed, and by embedding it into a cross-platform programming library that can be easily integrated into the development workflow of digital games. As such, we present, in this paper, the *Vorpals Open Real-time Procedural Audio Layer (VORPAL)*¹, an open-source, extensible, and flexible middleware system for the development of real-time soundtracks in digital games as the result of our research. The middleware implementation is available at <https://github.com/vorpals-project/vorpals> under the MPL 2.0 open source license.

2 Soundtrack Implementations in Digital Games

Matos [9] states that soundtracks are essentially the union of all sound effects, voices, and music that are played along with a visual presentation. In the traditional asset-driven approach, each audio element from these soundtrack parts is implemented in the game code by specifying a playback trigger consisting mainly of [3]:

1. *Which* sample assets are to be played.
2. *How* they are to be played (that is, what filters should be applied).
3. *When* they are to be played.

As an example of this pattern, consider a gunshot sound effect. Using a prerecorded gunpowder explosion sound, one could produce different firearms

¹ A recursive acronym.

sounds by applying multiple filters to it and then mapping the corresponding configuration whenever a weapon is shot. This way, when the player’s character wields a specific type of pistol, its respective sound effect is triggered.

Being able to synchronize an audio element playback achieves an initial level of real-time behaviour on its own. It is often further expanded by allowing the other two parameters (the *which* and the *how*) to change according to the game state. In the game *Faster Than Light* (*Subset Games*, 2012), every musical piece in the soundtrack has two versions – one for exploration and one for combat – and they are cross-faded between each other whenever the game situation changes from exploration to combat and vice-versa. This consists of both a modification in the sample used (the *which*) and a real-time control over the filters, responsible for fading out the previous version of the music while fading in the next one (the *how*).

However, since this method specifies only whole samples to be played, it excludes from the sound design space other forms of audio construction, notably symbolic representation (e.g., MIDI) and Digital Signal Processing (DSP) audio synthesis. The *IMuse* system [8] was an example of how to use symbolic representation to allow music changes in real-time. Essentially, it provided if-then-jump commands among the other typical music score based messages, bringing symbolic representation closer to a programming language of its own. Regarding DSP audio synthesis, there are quite a few works on physically based real-time audio synthesis [1, 4, 6], which could contribute to many unexplored ways of dealing with sound effects in games using no sample files, but at a greater computational cost.

3 Real-Time Restrictions in Game Implementation

To be considered a real-time application, digital games rely on the **Game Loop** architectural pattern [5, 12]. It guarantees that the time difference between the continuous input processing and output generation is so short that the user experiences it as being instantaneous. This is accomplished by dividing the program execution into very short steps between each input handling and output rendering, then finely controlling the process rate of these cycles inside an endless loop.

Nystrom [12] shows how the Game Loop frequency is related to the Frames Per Second ratio (FPS) of the game, i.e., the ratio of how many graphical frame buffers are fully rendered per second. Ideally the FPS should be greater than or equal to the Game Loop frequency, which means that its visual output may change and adapt at least as often as changes are made to the game state. On the other hand, conventional asset-driven audio implementations in games provide a slower feedback mechanism, since they load as many audio samples as possible from the assets to the computer sound card (where they can no longer be promptly accessed) in each game cycle. The samples are transferred in chunks that are typically too big, causing effects applied to the game soundtrack to come

with a perceptible delay, thus not satisfying the desirable real-time requirements. Essentially, it incurs in too much *latency*.

For instance, in the *LÖVE* game framework², the default audio stream buffer contains 4096 samples, which, with an audio sample rate of 44100 Hz, leads to soundtrack changes being able to occur only every 93 milliseconds, approximately. The simple solution to this is to reduce the size of the audio buffers sent to the sound card, which actually means processing less audio in each game cycle. If a game is to be executed at 60 FPS and its audio is sampled at 44100 Hz, then each game cycle must provide only $44100/60 = 735$ audio samples. In the more general case, one cycle may actually have a variable time length. If we let f_{audio} be the audio sample rate (in Hertz) and Δt be the time difference (in seconds) between the current game cycle and the last, then the number N of maximum audio samples allowed for the current cycle would be:

$$N = f_{audio} \cdot \Delta t \tag{1}$$

As a matter of fact, the DSP graphical programming language *Pure Data*³ has a standard cycle buffer size of merely 64 samples. That would be precise enough for a game running at approximately 689 FPS. The drawback of reducing the audio buffer size is that if the computations needed to fill it actually last long enough to make a perceptible time gap between each buffer update, then the sound might come out chopped by the sudden lack of samples to play. There is also an increased overhead in having a larger number of smaller data copies sent to the sound card. Thus, even though reducing the buffer size is important to decrease latency, a point of equilibrium must be found or the audio quality may be compromised [7]. This depends on how much computation time the Game Loop has available for handling audio processing and on the technical capabilities of the sound hardware at our disposal. Weiner [15] proposes a possible audio pipeline that based on lazy evaluation, for instance.

4 Proposed Architecture

The main purpose of the *VORPAL* middleware is to bridge soundtracks and game engines. As such, we understand that its main user is the sound designer, although the programmers that bind the game code to the middleware must also be taken into consideration. This is commonplace for game audio middleware, requiring the division of the tool into two separate but complementary interfaces⁴. The first is a “soundtrack editor” – a visual application through which sound designers author audio content that can be later exported to the game.

² love2d.org

³ puredata.info

⁴ See *Firelight's FMOD Studio* (www.fmod.org), *Audiokinetic's Wwise* (www.audiokinetic.com/products/wwise), and *Elias Software's Elias* (www.eliassoftware.com)

The second is a programming library exposed through an Application Programming Interface (API) that is capable of loading the media exported by the editor and playing it during the game execution (see Figure 1).

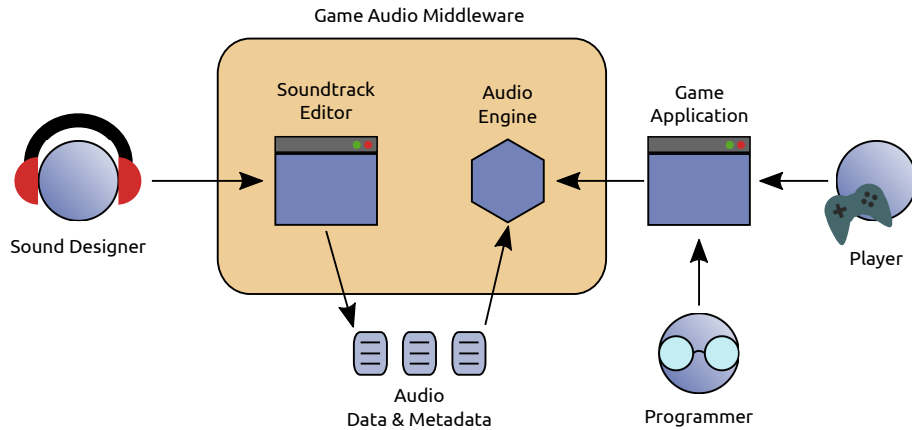


Fig. 1. Standard overview architecture for game audio middleware.

VORPAL follows this general architecture. However, instead of focusing the soundtrack editor in audio assets management, we chose to develop a procedure-oriented tool. We designed a collection of high-level abstractions in the Pure Data programming language, called the **Soundtrack Creation Kit**, that sound designers can use to produce Pure Data patches as part of a game soundtrack specification. The game engine can then link to *VORPAL*'s **Audio Engine** (a programming library) to load and *execute* the soundtrack patches authored this way. Our intention is to focus on giving sound designers full control over the sonic behaviour instead of just rigid sonic scheduling, since they can *program the soundtrack themselves* with an accessible, community-approved language such as Pure Data.

4.1 Pure Data Integration

Pure Data originally comes as a stand-alone application capable of creating, editing, and executing patches by itself, promptly serving as our soundtrack editor. However, even though it is capable of communicating with other applications through sockets or MIDI channels, ideally one would not want to have multiple applications launched when playing a digital game. The solution would be to run Pure Data's DSP *from inside the game*. There is a community developed programming library called `libpd`⁵ that provides access to Pure Data's core implementation, allowing its host program to load and execute Pure Data patches without the Pure Data application.

⁵ `libpd.cc`

However, when using `libpd`, there is no audio output handling. The host application is only given the processed signal and is responsible for sending it to the sound card or doing whatever it wants with it. Additionally, the processed signal is provided in blocks of 64 stereo samples, as mentioned before. Our audio engine must synchronize the retrieval of these blocks with the game run-time flow as in Equation (1). For that, the engine API routines demand precise timing information from the game execution *and* properly synchronized evocation, since time management is controlled by the game code.

The communication between the sound designer patches and our Audio Engine consists of two main types of data transfer. The first, which we just described, results from `libpd`'s audio signal computation happening every frame. The other transfer occurs when our middleware must inform the designer's patch of a relevant change within the game state. Based on this communication, the patch may do whatever it deems necessary for the soundtrack to follow up the game narrative. The overall architecture of our engine can be seen in Figure 2.

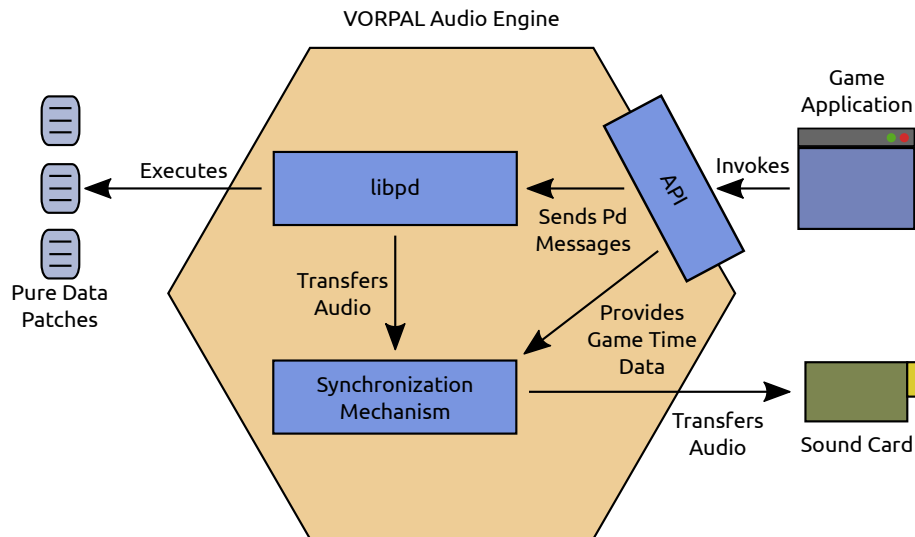


Fig. 2. *VORPAL* Audio Engine's architecture.

4.2 Middleware Abstractions

Finding a good model that links soundtrack elements to game elements depends mostly on how the latter are designed. But, aside from the Game Loop pattern, there is not much one can assume about the software architecture of a game without losing compatibility or discouraging users by imposing unwanted conventions. Thus, in this work we choose a rather weak assumption that within the

code base of a game there is the concept of *game objects*, ideally implemented through Object-Oriented Programming. Game objects range from obvious candidates such as the player’s avatar, the items lying on the floor, or a blazing fireball falling from the skies, to implicit or merely bureaucratic elements like the camera, a monster spawning spot, or an inventory interface button. The job assigned to the *VORPAL* middleware is, then, to map real-time soundtrack events to each game object – even a music track, which could be traced to an invisible auxiliary game object called “Player’s Headphone”.

This is done through the **Soundtrack Event** abstraction we introduce with the *VORPAL* middleware. Each Pure Data Patch the sound designer creates using the Soundtrack Creation Kit describes a Soundtrack Event type (like a class in Object-Oriented programming languages). When the programmer loads these patches through the Audio Engine API, he or she can create Soundtrack Event instances. Soundtrack Events are supposed to roughly present a one-to-one mapping with game objects. A character avatar could have a Soundtrack Event assigned to it to play its voice and dialogue, just as a door could carry a Soundtrack Event for when it creaks due to a player opening it.

The real usefulness of this abstraction lies in the concept of *messages* from Object-Oriented Programming. When a piece of code invokes a method on an object, we say that it is sending that object a message. Analogously, since our abstraction is based off on Object-Oriented Programming, Soundtrack Events can be sent **Commands**, which trigger real-time effects in them. For instance, the Soundtrack Event we previously assigned to a character avatar could respond to the Commands “shout”, “speak dialogue line 322”, “step in mud”, etc. Each of these would then be sent to the Pure Data patch that implements the Soundtrack Event type, where it could be used to perform the corresponding real-time DSP effect the sound designer defined. An example for music themes would be a Soundtrack Event that responds to Commands like “play”, “increase tension”, “roll dramatic transtion”, “finish up in at most 8 seconds”, etc.

5 Implementation Decisions

For the sake of game engine compatibility and performance [5, 12], we chose to develop our middleware in C++, except for the parts that must be made in Pure Data. The Audio Engine uses *OpenAL*⁶ for cross-platform open-source-friendly access to the sound card, enabling us to produce proper playback of the desired audio. It also comes with convenient spatialization features that are explored by the system.

5.1 Audio Engine

To satisfy the real-time restrictions described in Section 3, the *VORPAL* Audio Engine strongly relies on *OpenAL*’s buffer queueing mechanism. It enables the

⁶ See www.openal.org and kcat.strangesoft.net/openal.html

allocation of multiple buffers whose purpose is to send audio data to the sound card in First-In-First-Out order. Each buffer can have arbitrary sizes, but we fit them to Pure Data’s cycle block size (64 samples). Then, *OpenAL* automatically switches the current buffer to the next one when it has finished playing. That way, even when the game cycles do not match Pure Data’s cycles, we can schedule the next block. Doing so increases latency, but since the block is very small the difference is minimal and allows us to reduce the previously discussed overhead.

The actual audio transfer from a patch to the engine uses the Pure Data array API instead of its standard output, since this simplifies the use of multiple simultaneous audio buses. The engine recognizes the audio buses thanks to a naming convention. The Soundtrack Creation Kit then wraps it, hiding this and other implementation details needed to properly synchronize the filling of the arrays with the engine cycles.

Still on the engine side, we implemented the Command abstraction as follows. It relies on the messaging mechanism from Pure Data, accessed via `libpd`. With a single routine call, a message is sent to the sound designer’s patch containing as much information as desired, so long as it can be represented by a list of numbers and character strings (Pure Data symbols). Besides, each patch is represented by a C++ class designed after the Soundtrack Event abstraction. That is, each instance of that class is a Soundtrack Event instance, and it provides the API for sending Commands to the corresponding Pure Data patches.

5.2 Soundtrack Creation Kit

For a Pure Data patch made by the sound designer to actually work properly with the *VORPAL* middleware, it must contain at least one specific object from the Soundtrack Creation Kit. We developed an abstraction patch called **Audio Bus**, which captures a signal produced inside Pure Data and sends it to the corresponding Soundtrack Event instance in the Audio Engine. Each patch is then responsible for sending to its Audio Buses the corresponding sound signal. The captured signals are managed and mixed by the Audio Engine, then played when the game application demands so. Besides the Audio Bus, there other auxiliary abstraction patches the Soundtrack Creation Kit provides. Their intention is to guide the sound designer’s workflow when using the middleware, but they can be ignored if one wishes to work directly with the low level mechanisms of the Pure Data language.

6 Results

To validate our proposal, we used our middleware to create soundtracks for two very different games. First, we forked the open-source game *Mario0*⁷ and replaced its default soundtrack for one entirely produced by our middleware. The game is a parody of two other famous games: *Super Mario Bros.* (*Nintendo*, 1985)

⁷ stabyourself.net/mario0

and *Portal* (Valve, 2007). The focus was the music track, and we experimented only with Koji Kondo’s “Overworld Main Theme”, the music track for the first stage of the game. By dividing the music in its three voices – melody, bass and percussion – and the score bars into a few sections, we were able to add the following real-time behaviours to the game soundtrack:

1. The music advances through the bars as the player progresses through the stage, so that each of its parts has a particular music segment associated.
2. Mario’s size directly influences the bass voice. The stronger he gets, the louder the bass line becomes.
3. The quantity of enemies nearby also increases the percussion volume, in an attempt to suggest that the situation became more action-intensive.

The second validation is a soundtrack for an original title we created in partnership with a professional sound designer. The game developed was called *Sound Wanderer*, and has been accepted as a demonstration in an international conference on computer music [11]. We also evaluated the tool’s usage by potential users during that process. The source code of both games is available under an open source license at <http://compmus.ime.usp.br/en/vorpal/games>.

7 Conclusion

There are many details regarding our research that we had to skip over here. A detailed account of our findings and a throughout description of the *VORPAL* middleware can be found in the first author Masters Thesis available at the *VORPAL* web site. But we can say that there are many interesting challenges in this research area. First, as widespread as Pure Data is, it lacks a simple user interface that would be appealing to sound designers and composers with no programming experience. Second, with a new soundtrack composition workflow comes a whole new skill set for sound designers and composers to reach out to – an unavoidable cost of breaking well established paradigms [14]. Our expectations are that, with this work, we will at least open the way for new methods of producing soundtracks for digital games that actually try to exploit the dynamic and interactive nature of the medium, avoiding further waste of this game design space.

References

1. Bonneel, N., Drettakis, G., Tsingos, N., Viaud-Delmon, I., James, D.: Fast modal sounds with scalable frequency-domain synthesis. *ACM Trans. Graph.* 27(3), 24:1–24:9 (Aug 2008), <http://doi.acm.org/10.1145/1360612.1360623>
2. Collins, K.: *Game Sound: An Introduction to the History, Theory, and Practice of Video Game Music and Sound Design*. The MIT Press (2008)
3. Farnell, A.: An introduction to procedural audio and its application in computer games (2007), <http://cs.au.dk/~dsound/DigitalAudio.dir/Papers/proceduralAudio.pdf>

4. Farnell, A.: *Designing Sound*. The MIT Press (2010)
5. Gregory, J.: *Game Engine Architecture*. A. K. Peters/CRC Press (2014)
6. James, D.L., Barbič, J., Pai, D.K.: Precomputed acoustic transfer: Output-sensitive, accurate sound generation for geometrically complex vibration sources. In: *ACM SIGGRAPH 2006 Papers*. pp. 987–995. SIGGRAPH '06, ACM, New York, NY, USA (2006), <http://doi.acm.org/10.1145/1179352.1141983>
7. Lago, N.P., Kon, F.: The quest for low latency. In: *Proceedings of the International Computer Music Conference*. pp. 33–36 (2004)
8. LucasArts: Method and apparatus for dynamically composing music and sound effects using a computer entertainment system. United States Patent 5315057 (1994)
9. Matos, E.: *A Arte de Compor Música para o Cinema*. Senac, Brasília, DF, Brasil (2014)
10. Meneguette, L.C.: Situações Sonoras e Jogos Digitais. *Simpório Brasileiro de Games* pp. 30–33 (2013)
11. Mizutani, W.K., Vicente, D., Kon, F.: Sound wanderer: An experimental game exploring real-time soundtrack with openda. *Demonstration at the 12th International Symposium on Computer Music Multidisciplinary Research, 2016, São Paulo* (2016)
12. Nystrom, R.: *Game Programming Patterns*. Genever Benning (2014)
13. Schell, J.: *The Art of Game Design: A Book of Lenses, Second Edition*. A. K. Peters/CRC Press (2014)
14. Scott, N.: Music to middleware: The growing challenges of the game music composer. In: *Proceedings of the 2014 Conference on Interactive Entertainment*. pp. 34:1–34:3. IE2014, ACM, New York, NY, USA (2014), <http://doi.acm.org/10.1145/2677758.2677792>
15. Weiner, K., Ltd., D.: Interactive processing pipeline for digital audio. In: DeLoura, M. (ed.) *Game Programming Gems II*, pp. 529–538. Charles River Media (2001)