Navigating Soundspace: Modelling the Sound Domain At Real World

Mark A. Simos
Organon Motives
45 Shea Road
Cambridge, MA 02140
simos@hpcesra.ce.hp.com

Real World (Studios) Ltd.
Box Mill, Box
Wiltshire UK SN14 9PL

Abstract

This paper described a project initiated by the author in collaboration with Real World Studios, that is exploring ways of applying software reuse concepts, methods and tools to the domain of synthesized and sampled sound artifacts. The paper provides background on the project, motivates the applicability of software reuse concepts and tools in multimedia domains, describes overall project aims from a reuse perspective, and summarizes current project status. Digital sound resources exemplify a growing class of multimedia artifacts (such as fonts, clip art, and video images) that challenge our conventional notions of software systems or modules, yet exhibit a number of "software-like" attributes. As collections of these artifacts grow larger, more complex, heterogenous and semantically rich we encounter problems of component classification, retrieval, configuration and generation strikingly similar to those addressed by software reuse research. Reuse techniques may therefore contribute unique approaches within the area of audio and other multimedia database technology. In addition, we expect to extend and generalize reuse methods and tools by applying them beyond traditional software domains; for example, widening our repertoire of domain modelling techniques and representations. Even initial inventory and modelling tasks have yielded important general lessons for domain analysis methods. The project will also explore the possible role of multimedia technology itself as a component of general-purpose library interfaces.

Keywords: Domain Analysis, Reusability, Multimedia, Domain Modelling

1 Background (Problem Under Investigation)

Real World Studios is a world-class recording studio in Wiltshire, UK, founded by Peter Gabriel in 1987, and well-known in the music industry. A successful commercial studio, Real World's design, architecture and organization reflect its original intent: to create an ideal environment for recording live traditional music performances and, more generally, to foster creative collaboration between artists and technologists from many cultures. Real World is in fact a "community of enterprises", with formal or informal relations between several businesses situated in close proximity to the studio proper, including: WOMAD, a non-profit group that organizes international world music festivals in Europe, Canada, and Japan; Peter Gabriel, Ltd., Real World Records, a joint venture record label/publishing company; Real World Design, an engineering company that designs and prototypes audio hardware designs; a video editing company; and a local computer music composer-in-residence.

For several years, studio engineers and musicians have struggled to organize Real World's extensive and ever-expanding holdings of sound material (including many unique recordings of musicians from African, Asian, and other non-Western musical traditions) into a library that can be more effectively accessed. The studio is a technologically sophisticated environment; while there is an immediate problem to be addressed, artists and engineers have given considerable thought to advanced technology issues involved in creating a library system that will intelligently augment (rather than supplant) the current reliance on informal, intuitive, memory-based retrieval strategies.

2 Project Aims and Approach

The current project is an attempt to create synergy between the studio's needs and vision and the tools, methods and open research questions in the reuse field. Initiating the project involved a lengthy "warming up" process, including extensive and mutual education on software reuse and domain analysis concepts and how they could be applied to the sound modelling problem, and on the nature of the work at Real World and the studio's concept of the potential sound library system. As a result, Real World has supported some initial inventory, tool development and domain analysis activity. In addition, under the auspices of the studio, and as affiliate research scientist with the Palo Alto-based Institute for Research on Learning, I am preparing a concepts document (or "Manifesto") integrating Real World's conception with my previous and parallel work in reuse ([SIMOS86, SIMOS88, SIMOS90b, SIMOS91]). Sun Microsystems' Multimedia Platform Products group is contributing additional support for the proposal effort. The Manifesto will serve as the foundation for a more extensive project proposal; the intent is to fund the overall project with a combination of grants from corporate, academic, and possibly governmental (U.S., U.K. and/or EC) sources. Extracts from the library and software tools developed may also be productized for the commercial music market or for educational purposes. As the following sections will make clear, we intend to pioneer innovative approaches to domain analysis, domain modelling, and library navigation that should contribute significantly, not only to the state of the art in the computer music industry, but to reuse technology in general.

2.1 Reuse in Multimedia Domains.

The domain of synthesized and sampled sound seems to fall outside the scope of the kind of software components addressed by reuse research. In fact, digital sound resources exemplify a broad class of multimedia artifacts such as font and clip art libraries, stackware, video libraries, and animation, emerging as a natural by-product of the computer as a digital medium. Since collections of these artifacts can be stored, copied and transformed as digital data, in managing these collections we confront classic reuse problems such as: proliferation of subtle and unmanaged variants; the need for semantically expressive classification schemes; and the interplay of static parts with a large repertoire of generative and transformational capabilities [SIMOS90a]. One objective of the project is therefore to stress techniques and tools developed for software reuse in a multimedia information domain. This research should help recast these techniques and tools in more general form, and refine our understanding of their range of application (i.e., "re-engineering for reuse" our reuse research).

For example, a distinguishing feature of the software reuse problem involves the ability to use both constructive and generative techniques in engineering a component. A desired program family might be implemented with a set of static variants, with a generator, or a hybrid solution, depending on features such as the binding-time characteristics of the implementation language. This duality of static vs. dynamic implementation is relevant in the digital sound domain as well: synthesizers and samplers have numerous capabilities for dynamically altering the envelope and other attributes of a sound resource, and to store this dynamically modified version as a separate entity. Domain modelling tools need to provide transparent access to variants obtained through static selection or dynamic operations on components.

2.2 Alternative Modelling Techniques

A second, related objective of the research is to explore ethnographic methods of data acquisition and alternative domain modelling techniques for reuse. Software reuse researchers have considered a wide spectrum of representation/retrieval schemes (e.g., free text searches, keyword-in-context, faceted schemes [Prieto87], entity-relationship models, semantic network, frame-based, and object-oriented systems). While these vary widely in terms of flexibility, formality and extensibility, most share the characteristic of being textual/verbal in nature (even when a graphic interface is provided to navigate through the taxonomy). Effective retrieval of multi-media information may need to rely on different kinds of models and navigation strategies—both using different kinds of linguistic cues and non-verbal cues as well. Certain "virtuoso" studio engineers and sound programmers are surprisingly skillful at accessing massive collections of sound samples with minimal tool support. By empirical study of their intuitive classification and retrieval techniques, we hope to gain important insights into effective navigation strategies that will prove relevant beyond multimedia domains.

For example, in a "spatial" model, navigation through the component collection is projected into movement (continuous or quantized) through a qualitative "space", each dimension or axis of which is mapped to some parameter or polarity between features. As a familiar example in the typeface domain, movement from "light", through "medium", "demi-bold", and "bold" faces to "extra-bold" could be modelled as movement along an axis of "font weight". Preliminary analysis of the kinds of categories used in the sound domain suggest that qualitative polarities like "fat-thin" and "warm-cool" may be important ways of structuring the sound palette. An attribute of "brightness" vs. "darkness" (or "fatness" vs. "thinness") might be associated with one axis of the sound (or timbre) space. [1]

2.3 Domain Modelling as Composition

Many researchers in music perception and cognition have explored spatial models for the perception of musical sound quality (timbre). Most have hypothesized a general timbre space model, attempting to correlate (and validate experimentally) specific parametric changes with intuitive, qualitative categories, and embodying these schemes in the interfaces to computer music-based software systems. In contrast, our approach will not rely on pre-determining a single classification scheme for the sound assets; rather, we see the sound modelling system itself as a testbed supporting research into these questions. Our intent is to create a software platform enabling performers, composers, engineers and listeners to iteratively evolve their own models of the sound domain through working (and playing) with the system. We assume that the sound model will be actively shaped to reflect individual, stylistic and cultural preferences, and will metamorphose over time, in part through use of the technology itself. We intend to study both the modelling and retrieval process among a selected set of expert sound librarians, and document the changing process resulting from use of various iterations of the toolset as it develops. This will yield insights on how domain models develop and evolve over time, both through the work of individuals and through joint use within "communities of practice". In the long term, we hope to create an environment that encourages musicians and engineers to treat model-building as a kind of "instrument-building", extending and interwoven with composition and performance.

2.4 Use of Multimedia Information in Library Navigation.

The software platform must support a variety of modalities and strategies for organizing, classifying, and retrieving sounds within the collection. The ability to audition sounds in real time as part of active library browsing will be crucial to the tool's long-term success as an intuitive aid. We will investigate where names, categories, and unorthodox linguistic strategies can help structure the sound model, and where auditory information itself can provide feedback to the navigation process.

In exploring this set of capabilities for the sound library, we will also be considering the role that auditory cues can play in the general task of navigating large and conceptually complex information bases. Since humans process information differently in auditory, visual, and other sensory modes, we may find that sound can be an important facet of any intuitive, instrument-like, "playable" interface to a large library. At the same time, musicians often use imagery drawn from other sensory modes—not only visual terms (e.g., color associations) but terms invoking taste, smell and temperature. It may be that intuitive, "instrumental" access to any collection structured with reference to a given sensory mode relies on synaesthetic "mappings" to other sensory modes. If this thesis is substantiated, we envision linkage to further domain analysis experiments in other multimedia domains, simultaneously exploring the organization of multimedia in libraries and its role in library interfaces.

[1] This work correlates in interesting ways with the notion of a "design space" in software design [Lane90, Bell71]. Open research questions involve exploration of what reasonable measures for distance or neighborhood there may be in different domain spaces that potentially exhibit different geometries or even topologies.

3 Current Status and Experience

The first phase of the project, currently underway, involves several distinct tasks. First, preparatory work has been done similar in many ways to the preparation one would do for ethnographic fieldwork. This preparatory fieldwork involved familiarization with studio personnel and the rhythms and culture of the studio environment through interviews, documentation of artifacts and physical characteristics of the workplace, both passive and participant observation, and education in the complex technical areas involved, such as comparative study of synthesizer and sampler architectures, digital audio processing and multi-track recording technology. It has been important to develop the project concept in collaboration with members of the Real World community as a whole. This collaborative design approach helps create a sense of ownership on the part of eventual users of the system that is vital to its adoption.

Partly in keeping with this collaborative notion, and partly to get something practical done as a first step, short-term development tasks were identified to provide some immediate benefits to the studio while examining in more depth some of the larger issues involved. This short-term work has had three main facets: 1) definition and implementation of a relational database system for tracking synthesizer patches and samples from the various pieces of hardware in the studio; 2) inventory and data collection, logging the actual sound materials themselves; and 3) initial domain modelling: documenting the existing, ad hoc classification system; analyzing, re-clustering and consolidating existing categories; and finding intuitive representations for these revised categories and their inter-relationships.

4 Issues

These preliminary stages of the project—beginning an inventory of the current, ad hoc sound collection, assessing patterns of usage, and design of a simple relational database—have yielded valuable and practical insights into domain analysis and library retrieval issues. It's quite useful to start with a simple approach, even to see where it proves inadequate. This approach is in keeping with Prieto-Diaz' sound recommendations for an "Initiation" phase in getting a reuse program established within an organization [Prieto90].

Perhaps most illuminating have been the unexpected complexities of even the "simple" inventory task. "Software", understood loosely as data stored in a digital medium, breaks down our notions of physical commodities, because it can be replicated ad infinitum with a simple copy command. On a practical note, this means that inventorying existing assets in a software environment may present formidable problems in sorting out true variant versions from different copies of the "same" piece of software. For example, in cataloguing sounds on the KORG WaveStation, I found that approximately one third of the inventory consisted of duplicate or multiple copies. There are many and varied reasons for copying sounds in the studio environment: touring disks, temporary "palettes" created for composition, or "fly-in" recordings for sessions. The problem is aggravated by copies that have been "tweaked" at a level of granularity far below their overall size (like two giants differing by a toenail). Duplication and redundancy can also be the result of automated system activity: some locations were filled with automatic loads from other banks within the machine; in addition, some seeming inventory was in fact empty (though misleadingly titled "Unnamed"!)—

an "artifact" of the hardware. All these variants must be catalogued and sifted through in order to isolate the library assets proper. These sorts of problems apply to more than the artifacts themselves. One machine that provided an on-line database with user-definable categories sabotaged most advantages of the category field, simply by neglecting to provide a way of viewing the list of categories themselves; this resulted in many spurious variants of category names, making the database extremely error-prone to use. Similar issues are likely to arise in any software inventory activity.

5 Summary

The Real World Sound Modelling project is a novel application of reuse methods and techniques, exemplifying a broad class of multimedia database/library applications which may well become important new software "media" in their own right. By bringing together the perspectives of multiple disciplines, we hope to bring a more methodical approach to the structuring of sound libraries. At the same time, by viewing domain modelling as art form as well as engineering discipline [ARANGO89], we hope to bring a different approach to conventional software domains as well—taking one semi-tone step towards the transformation of merely "workable" into "playable" systems.

References

[ARANGO89] Arango, G., "Domain Analysis- From Art Form to Engineering Discipline", in *Proceedings*, 5th International Workshop on Software Specifications and Design, 1989, pp. 152–159.

[BELL71] Bell, C.G., and Allan Newell, *Computer Structures: Readings and Examples*, McGraw-Hill, New York, 1971.

[LANE90] Lane, T. *User Interface Software Structures*, CMU-CS-90-101, doctoral thesis, Carnegie Mellon University, 1990.

[PRIETO87] Ruben Prieto-Diaz, "Domain Analysis For Reusability", *Proceedings*, COMP-SAC'87, 1987, pp. 23–29.

[PRIETO91] Ruben Prieto-Diaz, "Making Software Reuse Work: An Implementation Model", position paper, First International Workshop on Software Reusability, Dortmund Germany, 1991.

[SIMOS86] Simos, M., *Alternative Technologies for Software Reusability*, Proceedings, 3rd STARS Workshop on Applications Systems and Reusability, March 1986.

[SIMOS88] ibid., "The Domain Life Cycle: Steps Toward A Unified Paradigm for Software Reusability", *Proceedings*, RMISE Workshop on Software Reuse, Boulder, CO., October 1987; reprinted in *Software Reuse: Emerging Technology*, IEEE Tutorial, ed. Will Tracz, 1988.

[SIMOS90a] ibid., "Methods and Tools for Reuse: Domain Analysis Issues", position paper, 3rd Annual Workshop on Methods and Tools for Reuse, Syracuse, NY, 1990.

[SIMOS90b] ibid., *Alternative Approaches To Domain Analysis: Steps Toward A Job Description*, course notes, 1990.

[SIMOS91] ibid., "The Growing of an Organon: A Hybrid Knowledge-Based Technology and Methodology for Software Reuse", in *Domain Analysis and Software System Modeling*, ed. Prieto-Diaz & Arango, IEEE Computer Society Press Tutorial, 1991, pp. 204-221.