DYNAMIC SEMANTIC NOTATION: JAMMING TOGETHER MUSIC ENCODING AND LINKED DATA

David M. Weigl, Kevin R. Page
Oxford e-Research Centre, University of Oxford, UK
`first.last@oerc.ox.ac.uk`

ABSTRACT

The Music Encoding Initiative (MEI) [1] provides a framework for expressing musical notation that enables the identification (via XML identifiers), and thus addressing, of score elements at various levels of granularity (e.g. individual systems, measures, or notes) [2]. Verovio [3], an open-source MEI renderer that produces beautiful SVG renditions of the score, retains the MEI identifiers and element hierarchy in the produced output, enabling dynamic interactivity with score elements through a web browser. We present a demonstrator that combines these capabilities with semantic technologies including RDF, JSON-LD, SPARQL, and the Open Annotation data model, anchoring into the musical notation by using the MEI XML IDs as fragment identifiers to enable the fine-grained incorporation of musical notation within a web of Linked Data. This fusing of music and semantics affords the creation of rich Digital Music Objects supporting contemporary music consumption and performance.

1. MOTIVATION

In March 2015 the all-hands meeting of the FAST project [1] (*Fusing Audio and Semantic Technologies for Intelligent Music Production and Consumption*) was supplemented with a late-night "jam session", where an eclectic line up of participants performed on guitars, bass, keyboard, violin, and percussion. During the jam, a participant would signal to the other players a new piece that he or she wished the group to transition to; or reference a particular prior recording or artist they wished the group to emulate in style or interpretation; or call out direction to shape the structural elements of the performance, such as to repeat a chorus or verse, or move to a bridge section.

Each of these call-outs invokes concepts transcending the symbolic representation of the music being played (be that full parts or chord sequences), and reference significant context that can be captured or supplemented by metadata from related material (e.g. about the artist, or a particular style, or the music structure). But, once identified and retrieved, all these contexts imply modifications to the musical notation that could be dynamically rendered to the individual performers, e.g. to modify articulation in a passage; or to jump to a different section, or queue up the next song to transition to. The tool presented here provides an initial implementation of this idea, using semantic technologies to reference external contextual information and dynamically generate Linked Data to both drive the interactions with the score and between the performers, and capture provenance information providing insight into the temporal evolution of the performance in terms of these interactions.

2. DATA MODEL

By treating the XML IDs of elements within the MEI resource as fragment identifiers, URIs can be straightforwardly generated for each notation element of interest. We employ the Open Annotation data model [4], using these URIs as *annotation targets* of annotations representing each requested action (*call-out*). Corresponding *annotation bodies* are associated with *semantic tags* defined to encode the different supported call-out types, e.g., `meldterm:Jump` requesting performers to jump from one section of the score to another. Each call-out is itself an annotation body of a *top-level annotation* targetting the URI of the MEI file that is currently being performed.

3. IMPLEMENTATION

We have developed a dynamic semantic notation client using HTML/CSS and JavaScript, served by a simple web service implemented with Python Flask. The procedure driving the rendering and multi-user interaction is illustrated in Figure 1. The client processes a framed [2] JSON-LD [3] representation of the RDF [4] graph instantiating the data model. It then performs an HTTP GET call to acquire the MEI resource targeted by the top-level annotation, and renders the corresponding musical score to SVG using Verovio. Capture of user interactions and visual indication of call-out results is handled using HTML divs drawn as bounding boxes over portions of the SVG corresponding to MEI elements of interest; this is simplified

[1] http://www.semanticaudio.ac.uk

[2] http://json-ld.org/spec/latest/json-ld-framing/
[3] https://www.w3.org/TR/json-ld/
[4] http://www.w3.org/TR/rdf11-concepts/

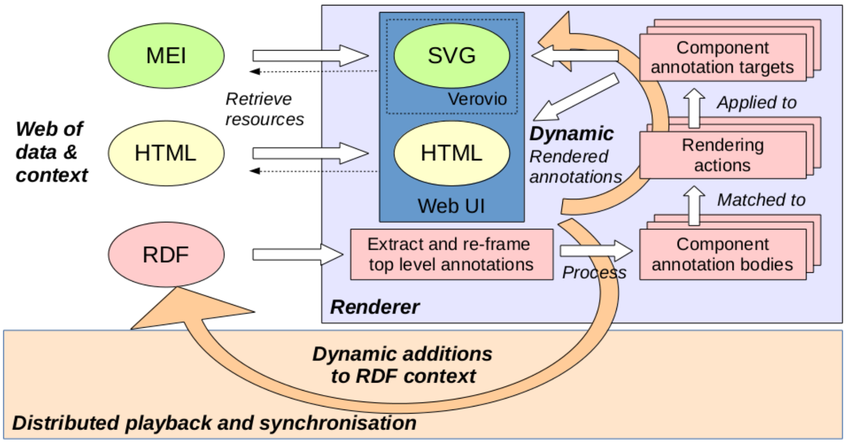

Figure 1. Dynamic rendering of musical notation and semantic markup corresponding to user call-outs

by Verovio's retention of MEI identifiers in the produced SVG output. Consequences of user call-outs are pushed to the server (HTTP POST), where they are incorporated into the RDF graph. The client then repeats this sequence in an iterative polling procedure, enabling multiple users (i.e., performers in the jam session) to interact dynamically with a shared representation of the score, presenting call-out outcomes to each user in near-real-time.

A screenshot of the system in action is provided in Figure 2. The client is capable of filtering multi-voiced MEI files to only show portions of the score relevant to the respective performer, via a simple XML transform using XPath; the MEI file displayed in the screenshot encodes a score for a string quartet, filtered to show only the viola part. Within the MEI hierarchy, voice-specific layers are situated underneath the level of measures. Thus, annotations targeting individual measures remain visible to each performer, even if those measures are located on different pages of the score according to the layout requirements of the respective voices being rendered.

Figure 2. Screenshot illustrating a call-out requesting performers to jump from the end of the 13th measure back to the introduction. In jumps across pages, users can click on the jump source (red highlight) to switch pages appropriately in order to land at the jump target (green highlight).

4. FUTURE WORK

Future work will focus on extending the supported actions and semantic capabilities of the system; handling the interface constraints of real-time user interaction within a live performance context; and expanding upon the applicability of our implementation and data model to scholarly, musicological use cases by enabling matching and linking of notation examples as they are described in scholarly text, via semantic hyperlinks, to and from the score, including exact matches and variants, illustrating interpretations, and situating the examples back in context.

5. ACKNOWLEDGEMENTS

This work is undertaken as part of the Fusing Semantic and Audio Technologies for Intelligent Music Production and Consumption project, funded by the UK Engineering and Physical Sciences Research Council under grant number EP/L019981/1, a collaboration between Queen Mary University of London, the University of Nottingham and the University of Oxford.

6. REFERENCES

[1] Andrew Hankinson, Perry Roland, and Ichiro Fujinaga. The music encoding initiative as a document-encoding framework. In *Proceedings of the 12th International Society for Music Information Retrieval Conference*, pages 293–298, 2011.

[2] Kevin R. Page. Music encoding in context: MEI and Linked Data. In *Music Encoding Conference*, 2015.

[3] Laurent Pugin, Rodolfo Zitellini, and Perry Roland. Verovio: A library for engraving MEI music notation into SVG. In *Proceedings of the 15th International Society for Music Information Retrieval Conference*, pages 107–112, 2014.

[4] Robert Sanderson, Paolo Ciccarese, Herbert Van de Sompel, Shannon Bradshaw, Dan Brickley, Leyla Jael Garc a Castro, Timothy Clark, Timothy Cole, Phil Desenne, Anna Gerber, et al. Open annotation data model. *W3C community draft*, 2013.