White Paper

Report ID: 111947

Application Number: HD-51836-14

Project Director: Raffaele Viglianti

Institution: University of Maryland, College Park

Reporting Period: 6/1/2014-10/31/2015

Report Due: 1/31/2016

Date Submitted: 2/25/2016

Enhancing Music Notation Addressability
White Paper

Grant number	HD-51836-14
Title of project	Enhancing Music Notation Addressability
Project director	Raffaele Viglianti
Grantee institution	Maryland Institute for Technology in the Humanities, University of Maryland
Date submitted	02/25/2016

Enhancing Music Notation Addressability White Paper

Introduction and background

Enhancing Music Notation Addressability (EMA), a research project directed by Dr. Raffaele Viglianti, was funded with a Level 2 Digital Humanities start up grant from the National Endowment for the Humanities between June 1, 2014 and October 31, 2015.

Prompted by the fact that addressing such units of music notation as measures, notes, and phrases has long been a powerful instrument in musicology for both analysis and historical narrative,[1] the project researched a method for addressing and extracting specific portions of music notation published in machine-readable formats on the web. Music notation, like text, can be "addressed" in new ways in a digital environment, allowing scholars to identify and name structures of various kinds, thus raising such questions as how can one virtually 'circle' some music notation? How can a machine interpret this 'circling' to select and retrieve the relevant music notation?

The project successfully investigated these questions by producing: 1) an Application Programming Interface (API) for addressing music notation on the web regardless of the format it is stored in; and 2) an implementation of the API for documents in the Music Encoding Initiative (MEI) format.

Both the API and the implementation were evaluated using a dataset of music analysis statements from the *Du Chemin: Lost Voices* project (Haverford College), which focused on reconstructing songs printed by Nicholas Du Chemin between 1549 and 1568 in Paris. This dataset has provided in particular the opportunity to evaluate how EMA methods can facilitate citation and attribution of credit, as is discussed in the Evaluation section below.

---

[1] When talking about music in general, it is important to say that addressing written music notation is not the only instrument of the musicologist. Music exists in several domains besides the written or "graphemic" one, each addressable in its own way (see Milton Babbitt, The use of computers in musicological research. Perspectives of New Music, 3(2): pp. 74–83, 1965). For the purpose of this paper, we focus on written Western music notation, because it both shares features with written language and plays a prominent role in musicological discourse.

The Music Addressability API

The API was created to provide a web-friendly mechanism for addressing specific portions of music notation in digital format. This is not unlike the APIs often provided by image servers for retrieving specific portions of an image. Such servers typically operate on a given large image file and are able to return different zoom levels and coordinate spaces. To use these servers effectively, one must know the image beforehand. One can ask for zoom level 3 coordinates 0 0 500 500, but cannot, for example, ask for all shapes looking like a vase. Though, by knowing beforehand that the vase is in a polygon of which the coordinates are known, one can ask the server to cut the image to size. The server will also do some processing, such as adjusting full-level coordinates to the requested zoom level. This basic service can then be used to build advanced applications, such as a map viewer that requests tiles through the server. The International Image Interoperability Framework (IIIF) has recently created an API to generalize this behavior, so that it can be implemented across multiple image servers and digital libraries. IIIF was used as a model for the Music Addressability API created for EMA.

At the early planning stage, EMA was focused on creating methods for addressing exclusively documents in the Music Encoding Initiative (MEI) format. This was because music notation requires substantial computational modeling even for the simplest musical text before any further operation is possible. There are many different ways of representing a single note; some aspects are common to all representation systems, such as information about pitch and duration, but some systems will prioritize certain aspects over others. Nonetheless, we have found that there are simple units that are typically represented by all music notation systems for common western music notation, such as measure, staff, and beat. Therefore it was decided to develop an API and a URI scheme to make it possible to target music notation resources on the web regardless of their format. Such a scheme may facilitate (and some cases enable) a number of activities around music notation documents published on the web. The following table gives a few basic examples of how an implementation of the URI scheme could be useful to musicological research:

Scholarly	**Visual**	**Procedural**
Analysis: being able to address components of music notation for analytical purposes. Example: precisely identify start and end of a pedal tone in Bach's Prelude no. 6 in D Minor, BWV 851.	*Rendering*: rendering music notation in an interactive environment such as a browser or a tablet requires the ability to cut up a large music document. For example to show only the number of measures that fit in a given space.	*Processing*: extracted portions of music notation can be passed on to another process. For example, given the MEI encoding of the Overture to Mozart's Don Giovanni, extract the string instrument parts and send them to another program that will return an harmonic analysis.
Citation: quote a passage from an encoded music notation file. For example the timpani in the opening bars of the Overture to Mozart's Don Giovanni.	*Highlighting*: address a segment of music notation to highlight it in a visual context (e.g. with color).	

The full API is published on GitHub at this address under a Creative Common Attribution 2.0 license: https://github.com/umd-mith/ema/blob/master/docs/api.md. We give here a brief overview of how the URI scheme operates.

Consider the following example,[2] and the notation highlighted in red:

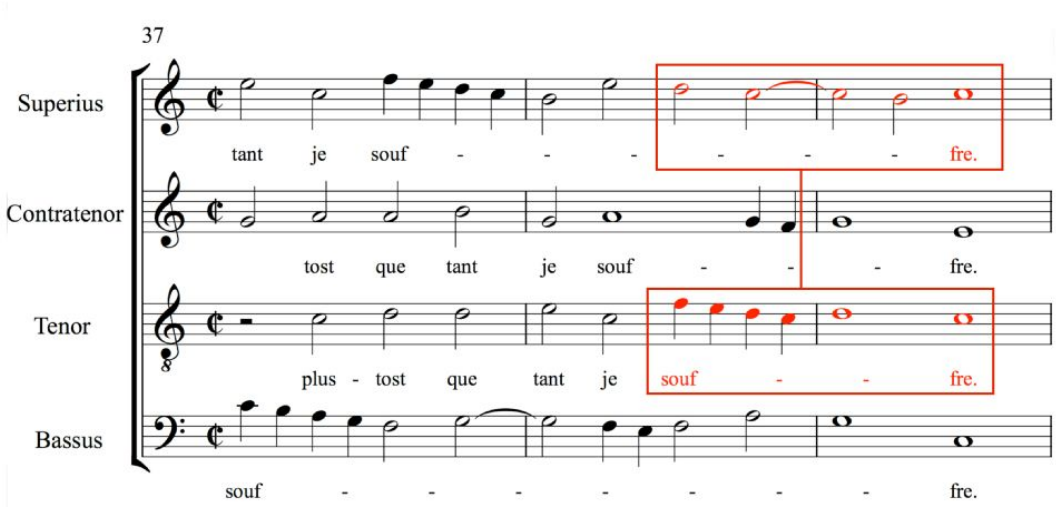

DC0519 L'huillier, Si je te voy

[2] Taken from *Du Chemin: The Lost Voices* project, at http://digitalduchemin.org.

The highlighted notation occurs between measure 38 and 39, on the first and third staves (labelled Superius and Tenor — this is a renaissance choral piece). Measure 38, however, is not considered in full, but only starting from the third beat. This selection can be expressed according to a URI syntax:

```
/{identifier}/{measures}/{staves}/{beats}

/dc0519.mei/38-39/1,3/@3-3
```

The measure is expressed as a range (38-39), staves can be selected through a range or separately with a comma (1,3), and the beats are always relative to their measure, so @3-3 means the third beat of the starting measure to the third beat of the ending measure. In this specification the beat is the primary driver of the selection: it allows for precise addressability of contiguous as well as non-contiguous areas. While this offers a sufficient level of granularity for most music notation, it has some limitations discussed in Conclusions and future work.

The API also includes a number of options to configure the “completeness” of the addressing act. These parameters help determine, for example, what to do when a beat only affects part of a note. We defer to the API specification for further detail, but we include here an overview of the URI scheme in Backus-Naur form as it is defined according to the version 1.0.0 of the API.

```
start ::= "start"
startOrEnd ::= "start" | "end"
all ::= "all"
measure ::= integer
measureRanges ::=
    {measure | startOrEnd | all / ","} | {measure | start, "-", measure | end / ","}
staff ::= integer
staffRange ::= {staff | startOrEnd | all / "+"} | {staff | start, "-", staff | end / "+"}
stavesToMeasures ::= {staffRange / ","}
beat ::= float
beatRange ::=
    {"@", beat | startOrEnd | all / "+"} | {"@", beat | start, "-", beat | end / "+"}
beatstoMeasures ::= {beatRange / ","}

selectionParameters ::= measureRanges, "/", stavesToMeasures, "/", beatstoMeasures
```

An implementation for MEI (Omas)

Unlike the API itself, its implementation cannot be format-independent. At the very least, an implementation of the API needs to know how a specific format models measures, staves, and beats to be able to retrieve the requested portion of music notation.

Our example implementation targets the Music Encoding Initiative (MEI) format not only because the dataset from *Du Chemin: Lost Voices* project (henceforth Digital Du Chemin, for brevity) is encoded in MEI, but also because MEI maintains a standard for representing music notation that is inspired by the same text encoding principles that motivated the Text Encoding Initiative, the *de facto* standard for encoding text in scholarly archives and editions.[3] As such, MEI has proven to be useful in the academic community, particularly for its ability in dealing with ambiguity and variance; this is useful, for example, for optical music recognition applications as well as digital scholarly editions.

The implementation is structured as a web service and is called Open MEI Addressability Service, or Omas. The code is available on GitHub under an Apache 2.0 license: https://github.com/umd-mith/ema/tree/master/Omas. Omas interprets a conformant URI, retrieves the specified MEI resource, applies the selection, and returns it. An additional parameter on the URI can be used to determine how "complete" the retrieved selection should be (whether it should, for example, include time and key signatures, etc.).

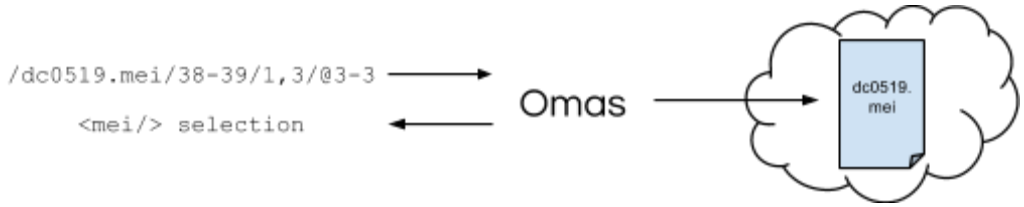

Similarly to an image server, Omas assumes that the information specified by the URL can be retrieved in the target MEI file. If requested, the web service can return metadata information about an MEI file, such as number of measures, staves, beats and their changes throughout the

[3] The Music Encoding Initiative has been jointly funded by NEH and the Deutsche Forschungsgemeinschaft (DFG) under the Bilateral Digital Humanities Program. The group has been awarded two grants for a total duration of four years. http://www.music-encoding.org.

document. This can be used to facilitate the creation of URL requests able to return the selection required.

A demo is published online for testing Omas: http://mith.us/ema/omas/. Any MEI file publicly accessible on the web can be addressed through this demo. A form is provided to enter the parameters for the URI scheme; after submission, Omas retrieves the MEI file and returns the addressed portion of music notation and the corresponding URL conformant to the Music Addressability API. It is also possible to attempt to read the music notation using Verovio, an MEI rendering engine.[4] Since Verovio is still experimental software, however, results will vary substantially.

Evaluation

In order to evaluate both the API and Omas, we took music analysis statements from the Digital Du Chemin project. Each musical piece in the corpus is annotated with analytical statements concerning various aspects of the notation, such as voice role and type of cadence.[5] All these statements address the notation at measure level and are stored in a relational database. We exported the data and re-structured it as Linked Open Data (LOD), using URLs conformant to the Music Addressability API to address the measures in the MEI files targeted by each analytical statement.

Each analytical statement was modelled as its own graph, according to the Nanopublication guidelines.[6] This specification is currently being used in the sciences to publish datasets; each scientific assertion is modelled independently so that it can be cited with a unique identifier (a "trusty" URI). A nanopublication is structured in three parts (the following list is quoted from the Nanopublication website):

1. The Assertion: An assertion is a minimal unit of thought, expressing a relationship between two concepts (called the Subject and the Object) using a third concept (called the Predicate).

[4] http://www.verovio.org/

[5] See for example *Qui souhaitez avoir tout le plaisir* by Pierre Certon at http://digitalduchemin.org/piece/DC0101/. The page shows the texts of the song, a rendering of the MEI data, and a table of analytical statements.

[6] http://nanopub.org/guidelines/working_draft/

2. The Provenance: This is metadata providing some context about the assertion. Provenance means, 'how this came to be' and includes the methods that were used to generate the assertion and attribution metadata such as authors, institutions, time-stamps, grants, links to DOIs, URLs about the assertion.
3. The Publication Information: This is metadata about the nanopublication as a whole, and pertains to both the assertion and provenance. Similar to the provenance graph, the Publication Information contains "citation-like" metadata but pertains to the nanopublication and not just the assertion.[7]

Each of these sub-graphs can be populated using any LOD vocabularies. We modelled the Assertion graph using the Open Annotation specification,[8] which provided helpful building blocks for associating each analytical statement with a specific portion of the music notation documents. Fundamentally, each Open Annotation is made of one or more *targets* and *bodies*. In our case, the body is the analysis, and the target is a URL formulated according to the Music Addressability API defined above.

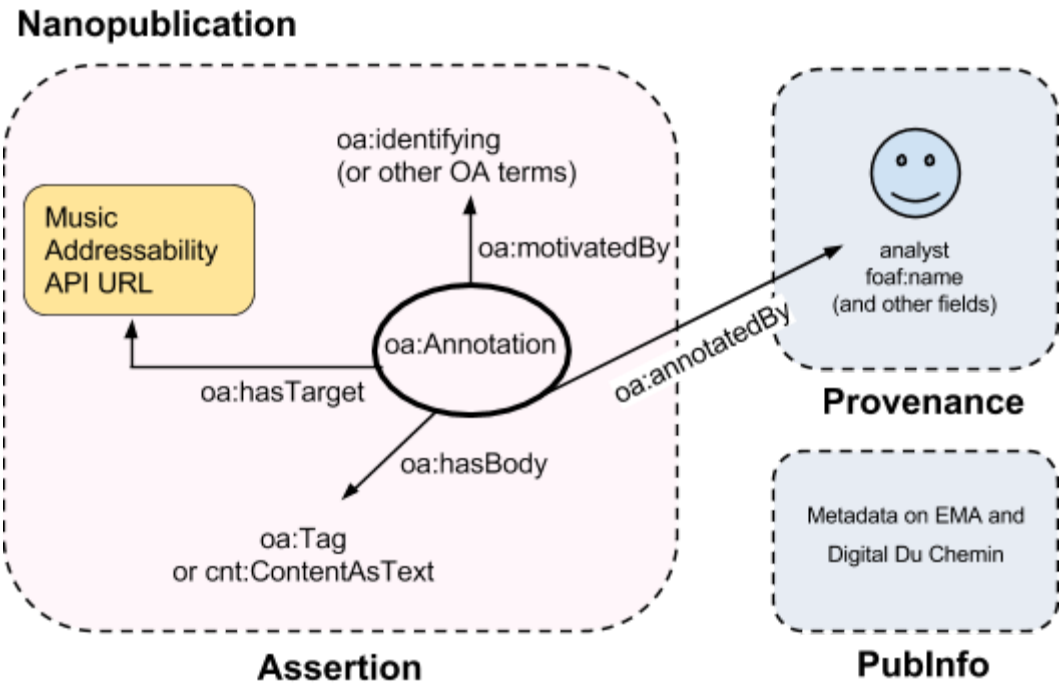

A simplified depiction of the structure of a nanopublication for the Digital Du Chemin dataset.

[7] Quoted from: http://nanopub.org/wordpress/?page_id=65.
[8] http://www.openannotation.org/spec/core/.

We wrote a Python script to generate the nanopublications and made it available on GitHub under an Apache 2.0 license: https://github.com/umd-mith/ema/tree/master/nanopub/csv2np.[9] Besides converting a CSV export of the analysis database, the script also parses the data and the references to measures to attempt and narrow them down to staff level. For example, some annotations will refer to specific parts (e.g. tenor), but will only be associated with full measures (this is how the database data is structured). With the Music Addressability URL scheme, however, it is possible to address the music at a greater level of granularity.

The generated nanopublications have been loaded to a Nanopublication Server available on Digital Du Chemin's site: http://digitalduchemin.org/np/. Each analysis can now be cited via a unique URI, which can return the nanopublication data in a variety of LOD formats. We also created a human-readable version that is displayed when the URI is visited from a browser, for example:
http://digitalduchemin.org/np/RAy6WBhrJFLDt0zSbZ75Ivcp902wS5u3mBluf6yqamSVY.

Nanopublication is more than a LOD specification, it is also a distributed network of servers, each mirroring published datasets. This means that once nanopublications are published to the network, they are impossible to delete or change, which is a powerful capacity that can critically improve the reliability of data published on-line. We find that the decentralization of nanopublications offers a novel paradigm for publishing humanities facts, both by providing an apparatus for responsibility metadata, and by potentially enabling discussion and revision in a space larger than the single humanities project. However, we decided not to publish Digital Du Chemin nanopublications to the wider Nanopublication network because this project's primary goal was evaluating methods developed by EMA, not producing reliable musicological data. Nonetheless, we believe that the Nanopublication approach is a valid and useful one for the humanities in general, and we will consider publishing to the network data from future projects.

The Digital Du Chemin analyses only target measures and staves, but the Music Addressability API can deal with a much greater level of granularity. In order to test this functionality, Richard Freedman (PI of Digital Du Chemin) has created a handful of more complex analyses.

[9] Naturally, this script will only work with a CSV export from Digital Du Chemin's database.

3/5/2015 15:43	Richard Freedman	DC0320	18-23	3,1-3,1-2,all,3-4,3-4 @4.5,@2-4+@3-4+@all,@all,@all+@1+@all+@3-4,@all,@1	Imitative Duo: S>Ct @5 below, T>B @5 below [Comment: Tenor motive *anticipates* main subject. It is not part of the ID itself.]	ID

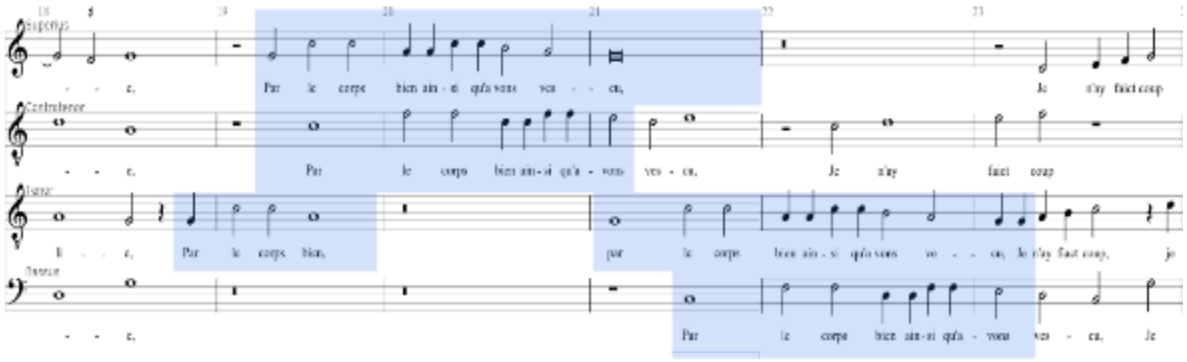

Example of an analytical statement with a complex selection.

A PDF listing the examples has been added to the GitHub repository: https://github.com/umd-mith/ema/blob/master/nanopub/examples/EMA_complex_selection_examples.pdf. The resulting nanopublications have been loaded to the Nanopublication server together with the other analyses, but they have also been added to the GitHub repository in JSONLD format to serve as examples: https://github.com/umd-mith/ema/tree/master/nanopub/examples.

In conclusion, building this corpus has successfully demonstrated that the Music Addressability API is capable of modelling complex analytical statements containing references to music notation. It also shows that it can be used together with other LOD vocabularies, such as Nanopublication, to enable and facilitate citation, attribution of credit, and distribution.

Conclusions and future work

The Enhancing Music Notation Addressability project was successful in creating a system for addressing specific portions of machine-readable music notation on the web. The Music Addressability API is arguably its most valuable contribution, while the API's implementation for MEI (Omas) can provide a working example for future implementation of the API for MEI and other formats. The evaluation completed using Digital Du Chemin data has provided a practical

demonstration of the applicability of the API to a musicological project. The API, however, has some known limitations.

- The API is fundamentally based on beat, therefore it is not possible to address music notation with relative or no beat. Cadenzas, for example, are ornamental passages of an improvisational nature that can be written out with notation that disregards a measure's beat, making it impossible to address subsets of the cadenza with the syntax discussed above. This, however, can be resolved in a future iteration of the API, though it is likely to complicate the URL scheme syntax.
- The API will not work well with notation such as neume and mensural notation; while it would be possible to address such notation at beat level as if it were contained in one sole measure, this is not ideal for specialists working with these notation systems. The best solution for this problem may be creating a separate URL scheme that does not rely on measure indexes.
- The API may or may not work with non-western music notation. The project has focused exclusively on western music notation for the time being.

While the project is now concluded and the promised goals have been achieved, there is certainly more work to be done to improve and stabilize the API with future releases. In particular the API could benefit from a large test suite for future implementations, documentation for edge and complex cases,[10] and reference implementations for other formats beyond MEI.

The work around the Music Addressability API and its implementation is going to be continued with the Project Director's personal research efforts, and with *Citations: The Renaissance Imitation Mass* (CRIM), a project funded by Mellon and the Fondation Maison des Sciences de l'Homme until 2018. For this project, Dr. Viglianti will develop a system for generating expressions conformant to the Music Addressability API by selecting portions of music notation from a human-readable rendering in the browser. The analysts will be identifying citations and other imitative cross-references across a corpus of Renaissance masses. These will be modelled as Linked Open Data with a dedicated ontology of relationships (currently being developed by CRIM participants), will be referring to MEI-encoded music documents with Music Addressability API expressions, and will be published as Nanopublications after the findings of the Enhancing Music Notation Addressability project.

[10] See for example the discussion about incomplete measures and *anacruses* here: https://github.com/umd-mith/ema/issues/14.