Recent Advancements in Roman Numismatics

Ethan Wooster Gruber
Kirby, West Virginia

BA, The Pennsylvania State University, 2006

A Thesis presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Master of Arts

McIntire Department of Art

University of Virginia
May, 2013

Table of Contents

Introduction

For more than five years, I have engaged in the conceptual and practical organization and development of digital numismatic projects—focusing mainly on coins produced in the Greco-Roman world. I did not intend to study numismatics as an academic or professional pursuit, but rather fell into the discipline by result of the intersection of circumstances: participation in a Roman numismatics graduate seminar taught by Professor John Dobbins at the University of Virginia in 2007 and my role as a Digital Humanities web applications developer at the University of Virginia Library. I oversaw the digitization of the University's numismatic collection during my tenure at the Library, which eventually grew into employment at the American Numismatic Society in January 2011, where I have since worked on numerous projects relating to Greek and Roman coinage.

This thesis, a hybrid of traditional humanities scholarship and cultural heritage informatics, reflects technical advances in digital numismatics made between 2011 and 2013. In this regard, this thesis is intended both as a potential guide to contemporary curators and information technology professionals who desire to make their collections broadly available on the World Wide Web and as documentation of the state of digital numismatics in the year 2013. This aspect of documentation is particularly important since many earlier digital numismatic projects have gone unpublished. It is paramount to record the thought processes which have informed the intellectual and technical decisions made during the development of these projects for the benefit of future generations of scholars, curators, and information technologists.

This paper introduces the reader to the history and evolution of numismatic description, from its inception during the thirteenth century to modern relational databases. It discusses conceptual methodologies of the semantic web applied to the description of coins, implemented in two American Numismatic Society projects (one a corpus of Roman Imperial coin types, the other a database of Roman Republican coin hoards).[1] Finally, it discusses the update of the University of Virginia Art Museum numismatic collection to the principles of linked data defined within this paper as a culmination of the technical and conceptual advances made in numismatic description. Arguably, the University of Virginia collection, together with its interaction with other services described in this paper, illustrates the future of numismatic collections on the web.

I would be remiss not to state that I work in a collaborative environment. While I have engaged in much of the data modeling and computer programming aspects of the projects detailed in this thesis, many others have lent feedback, insight, or have otherwise guided the development process. Especially important were scholars of numismatics with a far greater understanding of the discipline than I hold. Contributors include, but are not limited to, Andrew Meadows and Rick Witschonke of the American Numismatic Society, Gilles Bransbourg and Sebastian Heath of New York University, Kris Lockyear of University College London, and Sam Moorhead and Roger Bland of the British

1 A coin hoard is technically defined as a collection of three or more coins found within the same archaeological context.

Museum. A more complete list of contributors and their roles is included in the appendix.

The History and Evolution of Numismatic Description

From Medieval Inventories to 20th Century Type Corpora

Numismatics as a scholarly discipline can trace its roots at least to the private collections of Medieval lords of the thirteenth century. While numismatics was at first merely descriptive, it developed a more scientific methodology over time. By the Renaissance, connoisseurs of Greek and Roman coinage used the objects systematically to identify portrait busts in private collections.[2] Owing to their intrinsic metallic value, some ancient coins persisted in circulation into the modern era. Copper coins minted under Constantine (323-337) were found in circulation in remote parts of southern France during the reign of Napoleon III (1852-1870). In other cases, ancient coins, valued for their aesthetic beauty, were reused in decorative arts of the Medieval period: jewelry and reliquaries. The majority of coins that have survived into the present have come from the recovery of hoards, discovered underground by pure chance. In 1543, Transylvanian peasants discovered a treasure of the Dacian king Decebalus, which consisted of over 40,000 gold Greek staters.[3]

During the Renaissance, ancient coins became a source of learning rather than decorative or mystical objects. The earliest indicators of this Classical revival can be seen in the coinage of the Holy Roman Emperor, Frederick II of Hohenstaufen, who reigned from 1212 to 1250. Portrait motifs under Frederick II diverged from common

---

2 Clain-Stefanelli, *Numismatics*, 6.

3 Ibid. See IGCH 670, http://nomisma.org/id/igch0670.

Medieval designs and more closely approximated Roman models.[4] While no collection of Roman coins has been documented to have been within the possession of his court, Clain-Stefanelli postulates that return to Roman-styled portraiture must indicate direct influence from typologies in the emperor's possession.[5]

In the fourteenth century, Nicholas Oresme was the first scholar to approach the study of coins in a scientific manner. Throughout this century, Greek and Roman coins found their way into the collections of Italian humanists. The most significant collection of the time was held by Petrarch of Florence. A listing of collectors of the fourteenth century would include many great names, most of them Italian. With the upsurge in wealth brought in from maritime states such as Florence and Venice, coin collecting also saw significant growth. Cardinal Barbo, who would eventually become Pope Paul II (1464-71), assembled a collection of antiquities which would form the core of St. Mark's Museum in Venice. The collection of Antonius, Cardinal of St. Mark's and nephew of Pope Eugene IV (1431-1447), included 97 gold and more than 1000 silver coins. An inventory of Cosimo de' Medici's (1389-1464) collection taken the year after his death revealed more than 100 gold and 500 silver coins, to which 1844 coppers were later added.[6]

Beyond Italy, the Habsburgs possessed a collection of antiquities. An inventory taken in 1547 revealed a number of coins. The pre-Renaissance collector was interested

---

4 See 67.265.26, an "Augustalis" of Frederick II, in the possession of the Metropolitan Museum of Art: http://metmuseum.org/Collections/search-the-collections/170008718.

5 Clain-Stefanelli, *Numismatics*, 12.

6 Ibid., 13; Ogilvie, "Petrarch's Historicism and the Purpose of Philosophy," 11; Gnecchi, *Roman Coins: Elementary Manual*, 4.

in unusual specimens of aesthetic or mystical value. By the Renaissance, collectors were driven by a scholarly interest in science and history formed by empirical modes of thinking. With the rise of the book as a mass produced medium by the sixteenth century, iconographic and literary studies commenced, and illustrations of coins were published. Many numismatic monographs were published during this century, chief among them was Guillaume Budé's *De asse et partibus eius* in 1515, based upon his own collection of Greek and Roman coins. In 1553, Guillaume Rouille published *Promptuaire des médailles des plus renommées personnes qui ont esté depuis le commencement du monde* in Lyon, which contained drawings of coins, although many coins were misinterpreted or embellished with false legends. The important works of this period are too numerous to discuss in great detail within the scope of this thesis, so Clain-Stefanelli's historiography is recommended for providing a broader accounting of numismatic literature.[7]

With specific regard to the development of numismatic description comes *Commentatiorum vetustatorum numismatum*, published by Wolfgang Lazius (1514-1563) in 1558. Lazius, a Viennese doctor and humanist, conceived of the notion of a “corpus nummorum,” or a compendium of all of the coins of antiquity. His intellectual successors of the late sixteenth century, the German humanist, Adolph Occo of Augsburg and the Dutch scholar, Hubert Goltzius, were pioneers in the continual evolution of numismatics into a science.[8]

Occo pioneered a classification method based on chronology rather than by metal.

7 Clain-Stefanelli, *Numismatics*, 17–18.
8 Ibid., 19.

In 1579, he published *Imperatorum Romanorum Numismata a Pompeio Magno ad Heraclium*, a catalog of many of the known varieties of late Roman Republican and Imperial coin types. Figure 1 shows a portion of page 59 of the 1625 edition of this work. The work is organized chronologically by Republican personage and Roman emperor, with a listing of numismatic typologies within each of these sections. The description includes the Latin abbreviation of the metal of the coin, an “Orbis” number, a “V.C.” date referring to the year since the mythical foundation of Rome by Romulus (753 B.C.), the A.C./I.C. date (B.C./A.D.), transcription of legends, and the description (in Latin) of the iconography of the obverse and reverse of the coin.[9] While individual coin typologies are not given unique identifiers in Occo's catalog, Figure 2 shows that one of the “TVRPILIANVS IIIVIR” types contains a cross-reference to Aug. 38 in Goltzius' contemporary work, generally known by the title of its 1708 edition, *De re nummaria antiqua, opera quae extant universa*. Originally published in several parts from 1557 to 1579 in Antwerp and Brussels, *De re nummaria antiqua* served as a standard reference for Roman numismatics for two centuries.[10] Like Occo's catalog, coins were organized chronologically. The monograph consisted of hundreds of pages, including illustrations of coins (a feature which *Imperatorum Romanorum Numismata* lacked), glossaries of numismatic terminology, and indexes of people and places. Page 153 of Goltzius' work (Figure 3) illustrates coins issued during the Jugurthan war, with the year since the

9 Occo, *Imperatorum romanorum numismata, a Pompeio magno ad Heraclium*, 59. The role of the “Orbis” number is not clear. It is incremental, but not for every coin. It may refer to a physical tray number or a thematic grouping of coins, for example, Augustan coins which refer to the return of legionary standards by the Parthians.

10 Ibid., 57. These TVRPILIANVS IIIVIR types are identified as RIC 287 and 288. See numismatics.org/ocre/id/ric.1(2).aug.287 and http://numismatics.org/ocre/id/ric.1%282%29.aug.288.

foundation of Rome provided in the form of a Roman numeral.[11]

The increased attention upon numismatics which began in the sixteenth century continued into the seventeenth. France became one of the leading countries in this regard; royal collections bequeathed to Louis XIV were expanded as the king grew into an avid connoisseur of Greek and Roman coins. Jean Foy-Vaillant (1632-1706), Louis' chief agent responsible for acquiring coins for the royal cabinet, published numerous works about Seleucid, Ptolemaic, and Arsacid coins. One of his chief contributions to the discipline was *Numismata imperatorum Romanorum Praesantiora, A Julio Caesare Ad Postumum*, published in 1682. The catalog mainly consists of provincial coinage of the Roman Imperial period (including those of the eastern Mediterranean with Greek legends), following the format established by Occo. The region of production is listed if known, but years and materials are not explicitly denoted.[12] A Swiss contemporary of Foy-Vaillant, André Morrell (1646-1705), was hired to publish an inventory of the royal collection. He set out to establish a grand "corpus" of all European collections, but ultimately failed in his task. During this period, the English, Swedish, and Danish collections also grew significantly. James II built his cabinet upon the collections of his predecessors. In 1677, the antiquarian, Elias Ashmole (1617-1692), established what would become the Heberden Coin Room of the University of Oxford's museum.[13]

The eighteenth century saw much progress in the scientific evolution of numismatics. While the publication of catalogs of great private collections (of mainly

11 Goltzius, "Huberti Goltzii De Re Nummaria Antiqua, Opera Quae Extant Universa.," 153.
12 Foy-Vaillant, *Numismata Imperatorum Romanorum praestantiora a Julio Caesare ad Postumum et Tyrannos*.
13 Clain-Stefanelli, *Numismatics*, 22.

Classical coins) continued, Medieval and contemporary coins became subjects of interest among scholars. By the end of the 1600s, the discipline saw the publication of François Le Blanc's *Traité historique des monnoys de France* (1690) and Pierre Bizot's *Histoire metallique de la République de Holland* (1687). Shortly thereafter, Wilhelm Ernst Tentzel published four volumes of *Saxonia numismata* from 1705-1714. The body of coins to be classified expanded enormously, and alphabetical and chronological organization alone could no longer suffice. German numismatists of the period devised new organization methodologies based on denomination, which enabled the grouping of coins of multiple countries by similar denominations. This allowed scholars to rationalize more effectively the complex political and currency systems within Germany, which were, at the time, comprised of numerous bishoprics, cities, principalities, and even abbeys, all of which issued their own coins.[14]

By the end of the century, Classical numismatics saw the publication of one of its most important works—*Doctrina nummorum veterum*, an eight volume work published in Vienna between 1792 and 1798 by Joseph Eckhel (1737-1798). Eckhel was the director of the Imperial Coin Cabinet of Vienna and taught Classical archaeology at the University. In *Doctrina nummorum veterum*, he filtered out faulty interpretations and apocryphal data. The first 24 chapters of the first volume introduce the reader to the discipline. For the first time, Eckhel details the basic elements of ancient coins—metals, denominations, the organization of mints, significance of coin types, and the value of coins in relation to history and art. Eckhel abandoned the alphabetical arrangement of

14 Ibid., 27.

Greek coins in favor of one based on geography. For Roman numismatics, Eckhel established a chronological sequence of coin issues in a basic arrangement which has changed little in printed monographs in two centuries.[15] Volume VI of this series is focused on the early Principate. The catalog consists of legends, type descriptions, denominations, and, typically, several paragraphs which provide historical context for the issue. This format serves as a model for later catalogs, but there is still no numbering system by which to reference coins.[16]

By the early nineteenth century, catalogs of large private and public collections became more widespread once again. The year 1814 sees the publication of the British Museum collection by Taylor Combe with *Veterum populorum et regum numi qui in Museo Britannico adservantur* and, two years later, Christian Ramus' catalog of the coins in Copenhagen is published. Combe's work influenced Reginald S. Poole, later keeper of the cabinet at the British Museum, to initiate a great series of coin catalogs which today are counted among the standard reference works in numismatics. In 1883, Percy Gardner of the British Museum published *Types of Greek Coins*.[17] By the turn of the century, numerous advancements had been made in the classification and publication of Greek and Roman coins, too many to list within the confine of this brief introduction to numismatics. One should refer to Clain-Stefanelli's outstanding hisotry for a fuller picture of the evolution of the discipline, but I will summarize several publications which are directly relevant to this thesis.

15 Ibid., 28–29.
16 Eckhel et al., *Doctrina numorum veterum*, 91.
17 Clain-Stefanelli, *Numismatics*, 39.

By the late 19th century, Friedrick Imhoof-Blumer pioneered a die-linking methologogy that moved beyond simple stylistic considerations to establish relative chronologies in undated coin series "which opened new perspectives for the entire field of numismatics."[18] Inspired by Theodor Mommsen's *Corpus Inscriptionum Latinarum*, Imhoof-Blumer endeavored to create a grand corpus of all Greek types that would supersede all previous reference works. This laudable goal, dreamt initially by Wolfgang Lazius in the sixteenth century, was undertaken by German numismatists slowly until its abandonment in the 1930s. Contemporaneously, a series of important works emerged from the British Museum, including *Historia Numorum* (1887), written by Barclay V. Head, Reginald S. Poole's successor at the museum. Head codified the system of Greek geographic regions, first proposed by Eckhel, which remain the standard organizational means within Greek numismatics. The Greek catalogs of the British Museum are organized into volumes based on Head's regions, e. g., Ionia, with each volume sub-categorized into sections based on *poleis* or mints, e. g., Ephesus and Smyrna. Data are formatted into tables organized chronologically within these mint-based subsections. The tables contain weight standards, metals, size, and obverse and reverse legends and descriptions, including denominations. In order to conserve space on the page for the sake of printed publication, repetitive descriptive elements are replaced with "Similar" or other such abbreviations. Importantly, each type within the mint subsection is numbered, transforming the series into a reference work that other collections may use for attribution. For example, "BMC Greek Ephesus 11" refers to a tetradrachm of the late

---

18 Ibid., 42.

fifth century B.C. (Figure 4).[19] The reference numbering system employed by the BMC series is readily recognized throughout the discipline, and Greek numismatists and curators would make no mistake about the semantic meaning of the phrase. A computer, on the other hand, may find difficulty resolving the reference, a problem to be discussed in the next section of this thesis.

In the realm of Roman numismatics, Harold Mattingly and Edward Sydenham made progress where Imhoof-Blumer did not with the *Roman Imperial Coinage*, a ten volume type corpus encompassing all known (at the time of publication) imperial coin types from Augustus to Anastasius (27 B.C.-A.D. 518). Several of these volumes have been revised and republished since their original issue in the early decades of the twentieth century. Mattingly and Sydenham are also responsible for the publication of the catalogs of Roman coins in the British Museum in the early part of the twentieth century. With respect to Republican coinage, Sydenham later authored *The Coinage of the Roman Republic* (1952), which supplanted Ernest Babelon's obsolete *Monnaies consulaires* (1877).[20] Two decades later, in 1974, Michael Crawford published his corpus, *Roman Republican Coinage* (RRC), which continues to serve as the standard reference work for coins of that period.

The Advent of the Computer Database

In the 1980s, personal computers grew in prevalence. With this emerged the adoption of computer databases by various collections, including the American

19 Head and Poole, *Catalogue of the Greek Coins of Ionia*, 49.
20 Clain-Stefanelli, *Numismatics*, 43–44.

Numismatic Society, which was among the first organizations to migrate toward digital catalogs. Tracing the development of these databases from their inception to their current online rendition is extraordinarily difficult. Like many early projects in the "Digital Humanities," the conceptual and technical processes went unpublished for the majority of these databases. As a result, numismatic databases may only be discussed in broad generalities, with only occasional reference to concrete evidence.

What can be said, however, is that the earliest incarnations of databases typically followed the distinct curatorial practices of the organizations that produced them, i. e., database schemas and input workflows evolved from printed or handwritten cataloging methodologies. Gnecchi, in his early twentieth century handbook on coin collecting, dictates that every coin ought to have a number (e. g., a museum accession number) written on a round card, which was to be stored under the coin within the cabinet. The card should include provenance, value, reference to type corpus or other such scholarly publication, and descriptive attributes, such as metal, authority, or denomination.[21] Early databases inherited this model, especially those which predated the World Wide Web.

In particular, the American Numismatic Society's earliest MS-DOS-based database inherited peculiarities of the society's traditional curatorial practices, which varied across departments.[22] The ANS used Latin abbreviations to describe metal (AR for silver, AV for gold, etc.), but without authority control, the input of "Ar" or "Au" results in different computational interpretations, even if a numismatist understands the semantic

21 Gnecchi, *Roman Coins: Elementary Manual*, 15.

22 In numismatic collections, departments are divisions within the organization based on geographic, historic, or other thematic differences (e. g., Greek, Islamic, Modern, South Asian).

equivalence between "AV" and "Au." Furthermore, the database fields themselves had different semantic meanings between departments. The "person" field may denote the portrait which appears on the obverse of a coin within one department or may be used to denote the authority (which may or may not be the person depicted on the coin) in another. Compounding the confusion is the redundancy in names: Charles II may refer to the seventeenth century king of England or the seventeenth century king of Spain and Spanish Netherlands.

Even in a case where authority control is heavily enforced, such as the database of the Münzkabinett of the Staatliches Museen zu Berlin, aggregation of data across disparate collections described in varying native languages proves to be a complicated task. Berlin's database is bilingual—available in both German and English—but their database schema employs the full, unabbreviated word to describe the metal of the coin. *Silber* and AR are semantically equivalent intellectual concepts denoting the material composition of an object, but a computer database must be provided with an instruction that defines the two terms as identical concepts. Cross-collection querying is a fundamental requirement in pushing numismatics into the twenty-first century, and so it is imperative to provide a solution to this decades-old problem: the separation of important bodies of data into institution-specific silos.

Nomisma.org: Bridging the Past and Future

Applying standards of conformity to collections of library, archival, or museum materials is not a new idea. The Getty Research Institute maintains four thesauri which have become standards within the museum community: The Art & Architecture Thesaurus, The Cultural Objects Name Authority, The Getty Thesaurus of Geographic Names, and The Union List of Artist Names. The Getty's contribution to the standardization of terminology for art, architecture, and decorative arts cannot be overstated. A large museum may employ many curators and other staff responsible for inputting object metadata, and the Getty thesauri enable these institutions to encourage conformity in data entry. Even a minor difference such as an upper-case letter in a label (for example, “Denarius” instead of “denarius”) may result in two objects being classified by two denominations instead of the same one. The Getty maintains unique identifiers for the terminologies they define, but access to the underlying data that the institute produces requires subscription to its services, which is prohibitively expensive for smaller organizations. Even if one were to gain access to the Getty Research Institute's data, the information would not be accessible in a web standards-compliant format, but in a proprietary XML model. Due to these access restrictions, the Thesaurus of Geographic Names is being supplanted by Geonames.org for modern place names and the Pleiades Gazetteer for ancient ones, especially in open-source projects. Perhaps realizing this, the Getty plans to develop a system for providing their four thesauri as “linked open data”, which can be briefly summarized as a “recommended best practice for exposing, sharing, and connecting pieces of data, information, and knowledge on the Semantic Web using URIs [Uniform Resource Identifiers] and RDF [Resource Description Framework].”[23]

---

23 “Linked Data | Linked Data - Connect Distributed Data Across the Web”; Harpring, “Getty Vocabularies and Linked Data.”

In the library and archival fields, the Library of Congress developed a well-known system of subject headings which could be applied to bibliographic holdings to categorize subject-related materials, aptly named the Library of Congress Subject Headings (LCSH). Each subject heading may have broader or narrower terms, e. g., the term “Natural history museums” belongs to the broader subject of “Science museums” and may have subordinate categories of “Botanical” or “Zoological” museums. Each subject term maintained by the Library of Congress has a unique identifier, and recently these terms have been published in a system conforming to the tenets of linked open data. Tim Berners-Lee, the inventor of the World Wide Web during his tenure at the European Organization for Nuclear Research, has developed a simple list of specifications: use permanent URIs to define objects or concepts, provide standard-based machine-readable data at that URI (like RDF), and include links to other URIs so that users can discover more things.[24]

Before the introduction of the Library of Congress's new linked open data system for LCSH, each heading was identified by a unique number (e. g., sh85090241 for “Natural history museums”), but it was not possible to extract metadata about each heading and its relations. Now, this concept is represented by the URI, http://id.loc.gov/authorities/subjects/sh85090241, which directs the user to an HTML view of this information. Other machine-readable formats are available, such as RDF/XML at http://id.loc.gov/authorities/subjects/sh85090241.rdf. This information may be parsed by computer systems to extract broader or narrower terms, older variants of the

24 Berners-Lee, “Linked Data - Design Issues.”

heading, as well as the unique identifiers for exact or related headings maintained by other libraries, such as “Sciences naturelles – Musées,” defined by http://stitch.cs.vu.nl/vocabularies/rameau/ark:/12148/cb11967406z and maintained by the Bibliothèque nationale de France.

OCLC (Online Computer Library Center) has made similar progress with their project, the Virtual International Authority File (VIAF), in aggregating the unique identifiers and forms of personal and corporate names maintained by a growing number of national libraries. The person known as Augustus, born in 63 B.C. and who ruled Rome until his death in A.D. 14, is defined by http://viaf.org/viaf/18013086/. Despite libraries adopting different naming conventions, as result of both the native language of the collection and the libraries' own cataloging practices, adopters of VIAF IDs can make their materials available to cross-collection searching regardless of the language or script in which the name is written. VIAF provides RDF for computerized consumption, enabling the building of information systems which extract and maintain name forms, as well as provide more context about each person, such as Wikipedia articles.

The improvement in these linked open data systems over older, more rigid and closed systems of thesauri (like those of the Getty Museum) is significant. These systems are moving forward not only in the library and archive spheres, but also in museums. The British Museum disseminates their entire collection in RDF conforming to the CIDOC-CRM ontology for content of historical and cultural significance and makes their data available to the public through a query interface commonly referred to as a

"SPARQL endpoint," so-named for the recursive acronymn, "SPARQL Protocol and RDF Query Language."[25] If these few examples reflect the current trends of object description in cultural heritage, what steps should be taken to address the shortcomings of the plethora of numismatic databases operating in isolated silos in order to improve interaction between disparate datasets and move the field of numismatics into the twenty-first century?

In 2010, Sebastian Heath and Andrew Meadows established Nomisma.org, hosted by the American Numismatic Society, as "an effort to publish disciplinary-specific and stable http-based URIs for numismatic concepts so as to promote interoperability between numismatic collections and projects as well as links to and from other fields of study."[26] Nomisma uses URIs to define specific terminology used in the field and the concepts which belong to those categories of terminology. Terms include (among many more): reverse, represented by http://nomisma.org/id/reverse, and defined as "the reverse of a coin. Opposite of 'obverse'"; material (composition of an object, usually metallic, if it is a coin), represented by the URI http://nomisma.org/id/material; authority, defined as "generally, the personal, ethnic, or institutional authority under which a coin is issued. Frequently used to indicate the ruling monarch appearing on a coin," and identified with the URI, http://nomisma.org/id/authority. A visitor to one of these URIs is presented with a human-readable HTML page. This HTML is embedded with attributes conforming to

25 "British Museum - Collections Data." SPARQL is a W3C-defined RDF query language whose syntax is similar to Structured Query Language (SQL), commonly used for querying relational databases. The SPARQL Wikipedia article provides a concise definition. See http://en.wikipedia.org/wiki/SPARQL for more information.

26 Gruber et al., "Semantic Web Technologies Applied to Numismatic Collections."

the World Wide Web Consortium (W3C) RDFa 1.1 standard, chosen so that each concept defined on Nomisma could be both human and machine readable.[27] When the page is processed through a web service known as a distiller, machine-readable RDF is generated.

Below is an HTML fragment describing material:

```
<div xmlns="http://www.w3.org/1999/xhtml" typeof="numismatic_term"
about="material">
       <div property="skos:prefLabel" xml:lang="de">Material</div>
       <div property="skos:prefLabel" xml:lang="el">Υλικό</div>
       <div property="skos:prefLabel" xml:lang="en">Material</div>
       <div property="skos:prefLabel" xml:lang="es">Material</div>
       <div property="skos:prefLabel" xml:lang="fr">Matériel</div>
       <div property="skos:prefLabel" xml:lang="nl">Materiaal</div>
       <div property="skos:definition" xml:lang="en">The physical
material from which an object is made.</div>
</div>
```

The property attribute for each <div> is an RDF Property defined by the Simple Knowledge Organization System (SKOS). Represented as RDF/XML, below, the Nomisma namespace (http://nomisma.org/id/) defines an nm:numismatic_term for material at a particular URI, and presents SKOS Preferred Labels of the concept in various languages:

```
<rdf:RDF
  xmlns:nm="http://nomisma.org/id/"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:skos="http://www.w3.org/2004/02/skos/core#"
>
  <nm:numismatic_term rdf:about="http://nomisma.org/id/material">
    <skos:definition xml:lang="en">The physical material from which an
object is made.</skos:definition>
    <skos:prefLabel xml:lang="el">Υλικό</skos:prefLabel>
    <skos:prefLabel xml:lang="en">Material</skos:prefLabel>
    <skos:prefLabel xml:lang="de">Material</skos:prefLabel>
    <skos:prefLabel xml:lang="fr">Matériel</skos:prefLabel>
    <skos:prefLabel xml:lang="nl">Materiaal</skos:prefLabel>
    <skos:prefLabel xml:lang="es">Material</skos:prefLabel>
  </nm:numismatic_term>
</rdf:RDF>
```

27 Ibid., 3.

Many numismatic terms codified by Nomisma are related to concepts defined by those terms. For example, gold and silver are materials which are themselves represented by URIs, http://nomisma.org/id/av and http://nomisma.org/id/ar, respectively. Augustus, Alexander the Great, and the Aetolian League are all authorities with assigned URIs. The RDF metadata about these concepts also include labels in various languages, designated by skos:prefLabel and xml:lang.

A mint, or "generic term for a place of manufacture of a coin or for the issuing city," is also defined as a Nomisma concept.[28] The RDF structure is similar to that of material, above, but the RDF includes geographic coordinates encoded in the Geographic Markup Language (GML) namespace, as well as related resources on the web, defined by skos:related, such as a Wikipedia article or the URI identifying the mint in the Pleiades Gazetteer of Ancient Places. See an example of Rome, below:

```
<rdf:RDF
  xmlns:nm="http://nomisma.org/id/"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:skos="http://www.w3.org/2004/02/skos/core#"
>
  <nm:mint rdf:about="http://nomisma.org/id/rome">
    <skos:related
rdf:resource="http://pleiades.stoa.org/places/423025"/>
    <skos:related
rdf:resource="http://en.wikipedia.org/wiki/Ancient_Rome"/>
    <skos:definition xml:lang="en">The mint at the ancient site of Rome
in Latium.</skos:definition>
    <skos:prefLabel xml:lang="en">Rome</skos:prefLabel>
    <gml:pos xml:lang="en">41.9 12.5</gml:pos>
  </nm:mint>
</rdf:RDF>
```

A museum collection management system which interacts with Nomisma.org may import a label in whichever language is necessary, while creating a concordance between

28 See http://nomisma.org/id/mint for definition of "mint."

its own thesauri of controlled vocabulary and Nomisma URIs. In short, this makes it possible for the Berlin Münzkabinett to continue labeling their silver coins with the German “silber” and for the American Numismatic Society to continue their curatorial practice of using the Latin abbreviation, “AR,” but the two collections of coins may be linked together by the common language of Nomisma URIs. We are still in the early phase of establishing these interactions, but much progress has been made from 2012 to 2013 in this respect.

The Coin Type as an Intellectual Concept

Like materials, denominations, and mints, coin types are also intellectual concepts, often quantified with identifying numbers from printed corpora, and thus definable by Nomisma URIs. Various descriptors comprise a coin type, and it is this unique selection of descriptors that differentiates one typology from another. Currently, Nomisma has reserved URIs for those types from J. N. Svoronos' *Ta nomismata tou kratous ton Ptolemaion* (1908), Michael Crawford's *Roman Republican Coinage* (RRC), and part of *Roman Imperial Coinage* (RIC), a ten volume corpus of types from the Roman Empire from Augustus to Anastasius. Although Nomisma is focused, as of the publication date of this paper, primarily on the definition of concepts in ancient numismatics, type corpora listed on the site can be expanded beyond Greek and Roman both geographically and chronologically.

Technically speaking, a catalog, such as RRC, is defined by Nomisma as a type series (http://nomisma.org/id/type_series), and individual entries within the catalog are type series items (http://nomisma.org/id/type_series_item). Given that a coin type is composed of a combination of categorical attributes already defined as Nomisma IDs (denomination: denarius, mint: Rome, etc.), dates, and textual descriptions of the iconography and transcribed legends on both sides of the coin, the structure of a coin type requires a more complex model in RDF. Below is Nomisma's machine-readable representation of RRC 100/2, indicating that the coin is a bronze semis, minted 209-208 B.C. by an anonymous or unknown issuer:[29]

```
<rdf:RDF
  xmlns:nm="http://nomisma.org/id/"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:skos="http://www.w3.org/2004/02/skos/core#"
>
  <nm:type_series_item rdf:about="http://nomisma.org/id/rrc-100.2">
    <nm:start_date
rdf:datatype="http://www.w3.org/2001/XMLSchema#gYear">-
0209</nm:start_date>
    <nm:issuer xml:lang="en">Anonymous</nm:issuer>
    <skos:prefLabel xml:lang="en">RRC 100/2</skos:prefLabel>
    <nm:end_date rdf:datatype="http://www.w3.org/2001/XMLSchema#gYear">-
0208</nm:end_date>
    <nm:mint rdf:resource="http://nomisma.org/id/canusium"/>
    <nm:material rdf:resource="http://nomisma.org/id/ae"/>
    <skos:definition xml:lang="en">Roman Republican Coinage (Crawford
1974) no. 100/2</skos:definition>
    <nm:denomination rdf:resource="http://nomisma.org/id/semis"/>
    <dc:partOf rdf:resource="http://nomisma.org/id/roman_numismatics"/>
    <dc:partOf rdf:resource="http://nomisma.org/id/rrc"/>
    <nm:obverse>
      <rdf:Description>
        <nm:description xml:lang="en">Laureate head of Saturn, right;
below, mark; behind, denominational mark.</nm:description>
      </rdf:Description>
    </nm:obverse>
    <nm:reverse>
      <rdf:Description>
        <nm:legend xml:lang="en">ROMA</nm:legend>
        <nm:description xml:lang="en">Prow right; before, mark
.</nm:description>
      </rdf:Description>
    </nm:reverse>
    <dc:format rdf:resource="http://nomisma.org/id/coin"/>
  </nm:type_series_item>
</rdf:RDF>
```

29 Data from *Roman Republican Coinage* was provided to Nomisma.org courtesy of The British Museum.

There are several thousand types enumerated RRC, but RIC includes more than 40,000 distinct typologies, although this figure does not represent all known imperial coin types, as many new ones have surfaced through individual finds and hoards excavated over the last few decades.[30] RIC is the optimal corpus from which to generate references on Nomisma.org. It is one of the most complete catalogs of any particular area of ancient numismatics, organized more or less chronologically. Greek numismatics as a whole lacks any corpora similar to RIC. The body of Greek coinage is far more expansive than that of the Roman Empire, and the focus within the field has primarily been the documentation of coins within various public and private collections with the *Sylloge Nummorum Graecorum* (SNG), under the supervision of the International Numismatic Commission.[31] The SNG does not endeavor to develop a unique numbering system for types that can be applied across the entire spectrum of Greek numismatics.

Unlike the SNG, the volumes of RIC are organized by emperor. The types of Augustus occupy the first chapters of Volume I. The part for each emperor is divided into sections pertaining to some combination of geographic origin (mint or region) and typological theme (e. g., Nero's “Commemorative coins of Divus Claudius”).[32] RIC's numbering system generally applies on a per-emperor basis, which is to say, Augustus's types number from 1a to 535, Nero's from 1 to 622. This numbering system was adapted to generate web-compliant URIs from Nomisma. The identifier ric.1(2).aug.1a can be read as follows:

30 Moorhead, Booth, and Bland, *The Frome Hoard*, 31. The Frome Hoard, excavated in 2010, contained new Carausian types never previously documented.

31 http://www.sylloge-nummorum-graecorum.org/

32 Sutherland and Carson, *The Roman Imperial Coinage* Nero Vol. I (second edition), nos. 1-8.

ric: Catalog identifier, *Roman Imperial Coinage*
1(2): Volume I (second edition)
aug: Authority abbreviation, Augustus
1a: RIC-given coin type identifier

Http://nomisma.org/id/ric.1(2).aug.1a is a permanent, stable URI that organizations can use to express the intellectual concept of RIC Augustus Vol. I (second edition), no. 1a[33]. Nomisma seeks to capture identifiers within past, present, and future editions of RIC. The metadata associated with this identifier from the second edition will contain a cross reference to http://nomisma.org/id/ric.1.aug.1, i. e., its counterpart in the first edition of the volume. The linkage between equivalent types across different editions of RIC will facilitate museum database updates en masse. For example, if a third edition of Augustan types were to be published twenty years from now, the American Numismatic Society database records which link to Nomisma references within the second edition of RIC can be migrated to the third edition with an automatic process by observing changes in the source RDF. Changes to the data within the third edition coin type, ric.1(3).aug.1a, can be automatically ingested into the ANS database. These processes are theoretical now, but long-term data sustainability must always factor into the architecture of a web application.

Let us cast a more focused attention on Augustus 1a. This type is a silver quinarius issued under the authority of Augustus between 25 and 23 B.C. by the moneyer, P. Carisius. It was minted in Emerita, Lusitania. The obverse features the “head of Augustus, bare, left,” with a legend of “AVGVST;” the reverse indicates “P CARISI

33 http://numismatics.org/ocre/id/ric.1(2).aug.1a

LEG" with "Victory standing right, placing wreath on trophy with dagger and sword at base." The type is represented as a Nomisma ID, as previously discussed, but its metadata are partially composed of textual descriptors which oftentimes are distinct among types and also composed of other numismatic concepts represented by Nomisma IDs (http://nomisma.org/id/emerita, http://nomisma.org/id/augustus, http://nomisma.org/id/quinarius, etc.). A challenge emerged in developing a schema to effectively record this data. Looking at encoding standards within the library and archival realms as a model, the author commenced the development of a draft XML schema for describing and transmitting numismatic data online which meets the following specifications: 1) differentiate between two types of records: conceptual (coin type) and physical (an actual specimen within a physical collection, 2) ability to record all of the physical and typological attributes of coins while supporting linking to numismatic concepts defined by Nomisma.org, 3) record collections management data about an object (e.g., provenance, acknowledgments, and digital images), consistent with the practices of the American Numismatic Society and similar institutions, and 4) record metadata about the creation, revision, and publication of the XML record itself. In addition to developing a data model with which to encode numismatic data, software was and continues to be developed to create, manage, and publish these data on the web.

Although the draft of this new schema emerged in 2011, the seeds of this endeavor were planted in 2007 with the digitization and publication of the University of Virginia Art Museum Numismatic Collection, a project detailed in a paper presented at CAA 2009 in Williamsburg, Virginia.[34] The project, emerging from a Roman Numismatics graduate seminar taught by Professor John Dobbins in Fall 2007 and funded by a University of Virginia Library Innovation Grant, eventually became an open-source project called Numishare, which is freely available on the open-source code repository, Github.[35] At the time of the publication of "Encoded Archival Description for Numismatic Collections" in 2009, Numishare was built upon a foundation of several open-source applications which run in Apache Tomcat: Cocoon, a Java-based framework designed to effectively process XML (into HTML and KML, among other formats); Solr for faceted searching, eXist XML database, and Orbeon XForms for editing, managing, and publishing XML. These applications, including various XSLT stylesheets, Javascript files, and CSS stylesheets, comprise Numishare, which has since been made generalizable to all numismatic collections, regardless of era. Numishare has evolved considerably since 2009; virtually all code has been rewritten. First and foremost, the XML data model was completely reinvented.

Numismatic Description Standard

The data model of the U. Va. Art Museum Numismatic Collection was an adaptation of Encoded Archival Description (EAD) to coins. EAD is an XML metadata standard used throughout the library and archival communities for encoding electronic finding aids. Thus, EAD is focused primarily on the description of written documents, though photographs, maps, and other predominantly non-textual objects can also be

34 Gruber, "Encoded Archival Description for Numismatic Collections."
35 https://github.com/ewg118/numishare

described. At the time, no standard for numismatic metadata had been authored, and the adaptation of EAD to coins was satisfactory given the needs of that project. Over time EAD was found to lack descriptors needed distinctly by numismatists and museum curators, as well as attributes to facilitate linked data interactions.

Independent of the development of EAD for coins and the U. Va. digitization project, Sebastian Heath and Andrew Meadows of the American Numismatic Society developed, in a series of workshops funded by the UK's Arts and Humanities Research Council in 2006-2007, a list of recommended numismatic and curatorial fields.[36] Called the Numismatic Database Standard (NUDS), the list contained many fields which were not included in EAD. Some of these were related to provenance and collection management: auction history, sale lot and price, owner, etc. Others were related to description of coins themselves: artists and engravers, symbols and countermarks, edge descriptions, and findspots. Many of these fields, like artist and engraver, can be expressed generally about the coin as a whole or more specifically to the obverse or reverse of the coin. NUDS fields were proposed with a database management system in mind, thus hierarchical complexities and repetition cannot be easily or efficiently represented. Moreover, denoting uncertainty on any level within the record is also difficult. The EAD and NUDS approaches had distinct advantages and disadvantages, and the American Numismatic Society has endeavored to address the differences between both in order to come to a solution that meets the requirements for all implementations of the model.

36 Heath, “Numismatic Description Standard (NUDS).”

In summer 2011, work began on a data model which uses NUDS fields reconceived in an XML schema purpose-built for numismatics. By June 2012, the model, although still considered to be a draft, is stable, and several projects have been built on top of it. NUDS remains the title for the XML implementation of these fields, though the acronym now stands for "Numismatic Description Standard." The model is influenced by other common XML standards found in the Library, Archive, and Museum world, like EAD, Encoded Archival Context – Corporate, Personal, and Family (EAC-CPF), VRA Core, Metadata Encoding and Transmission Standard (METS), Metadata Object Description Schema (MODS), and Text Encoding Initiative (TEI). Like several of these standards, NUDS uses the World Wide Web Consortium xlink attributes for semantic linking, enabling links to URIs of Nomisma concepts, places on Geonames.org, and personal and corporate names defined by VIAF. A NUDS/XML record is fundamentally "linked data" because it is an electronic record which links to other resources on the web, although it is not linked data in the way that many of us may conceive of it or how Berners-Lee himself defines it—RDF in a triple store with a SPARQL endpoint for querying. Like traditional linked data architectures, querying for machine-readable data are supported in Numishare and will be discussed at a further point in this paper.

NUDS/XML

The document root of a NUDS/XML document is <nuds>, defined by the namespace http://nomisma.org/nuds. It requires one vital attribute for denoting the type of object represented by the electronic document. Objects are either physical or conceptual.

A record for a physical object is self-explanatory: the XML document describes coin, medal, token, paper note, or other numismatic object which physically exists in a collection. A conceptual object, however, is an intellectual construct. In numismatic terms, it is a coin type. The following outline illustrates the structure of the document:

- nudsHeader
- descMeta
  - Miscellaneous field types: title, department, subjects, abstract
  - typeDesc (type description)
  - physDesc (physical description)
  - undertypeDesc (undertype description for physical specimens which have been counterstamped or marked for recirculation)
  - findspotDesc (findspot description, including geographic location, date of discovery, disposition, etc.)
  - adminDesc (administrative description)
  - refDesc (references to catalogs or published work about the object)
- digRep (digital representations: links to physical specimens from coin type records or links to digital images of records describing real coins)

Regardless of the physical or conceptual nature of the object being described, all NUDS documents must contain a NUDS Header. Like headers in TEI and EAD, the NUDS Header contains metadata about the electronic record itself: unique identifier, publication information, rights statements, and revision history. Within the document root, below the NUDS Header is Descriptive Metadata about the object itself. This includes, title, subjects, and physical, typological, undertype, findspot, and reference descriptions, in addition to administrative history for encoding provenance, ownership and accession information, and other other fields necessary for collections management. A physical object may utilize any or all of these categories of descriptors, but a coin type record would not contain descriptors for collections management, physical state, or undertypes, which are specific only to physical objects. Thus, the Typological Description node is the only required Descriptive Metadata section.

```
<typeDesc>
	<objectType xlink:type="simple"
xlink:href="http://nomisma.org/id/coin">Coin</objectType>
	<manufacture xlink:type="simple"
xlink:href="http://nomisma.org/id/struck">Struck</manufacture>
	<dateRange>
		<fromDate standardDate=”-0025”>25 BC</fromDate>
		<toDate standardDate=”-0023”>23 BC</toDate>
	</dateRange>
	<denomination xlink:type="simple"
xlink:href="http://nomisma.org/id/quinarius">Quinarius</denomination>
	<material xlink:type="simple"
xlink:href="http://nomisma.org/id/ar">Silver</material>
	<obverse>
		<legend>AVGVST</legend>
		<type>
			<description xml:lang="en">Head of Augustus, bare,
left</description>
		</type>
		<persname xlink:type="simple"
xlink:role="portrait">Augustus</persname>
	</obverse>
	<reverse>
		<legend>P CARISI LEG</legend>
		<type>
			<description xml:lang="en">Victory standing right,
placing wreath on trophy with dagger and sword at base</description>
		</type>
		<persname xlink:type="simple"
xlink:role="deity">Victory</persname>
	</reverse>
	<geographic>
		<geogname xlink:role="mint" xlink:type="simple"
xlink:href="http://nomisma.org/id/emerita">Emerita</geogname>
		<geogname xlink:role="region" xlink:type="simple"
xlink:href="http://nomisma.org/id/lusitania">Lusitania</geogname>
	</geographic>
	<authority>
		<persname xlink:role="authority" xlink:type="simple"
xlink:href="http://nomisma.org/id/augustus">Augustus</persname>
		<persname xlink:type="simple" xlink:role="issuer">P.
Carisius</persname>
	</authority>
</typeDesc>
```

The Typological Description (<typeDesc>) of the NUDS document of Augustus 1a, above, contains a mix of elements which utilize the W3C xlink:role and xlink:href attributes for semantic linking and other elements which simply contain free text.[37] In the above fragment, the <typeDesc> contains an object type, denomination, manufacture method, and material, each linking to concepts defined by Nomisma URIs. Date element functionality is borrowed directly from EAC-CPF. A date, date range, and date set (which may contain any combination of other dates and date ranges) are available for use. The date set, for example, may be used to indicate an issue date of 45 B.C. *or* 40 B.C. to 35 B.C., explicitly excluding 44-41, using a combination of a date and a date range. The standardDate attribute allows a union of the XML Schema Datatypes of xs:date, xs:gYear and xs:gYearMonth. To express uncertainty, the notBefore and notAfter attributes may be used to denote *terminus post quem* and *terminus ante quem*, respectively. Below these elements are structured data: obverse, reverse, geographic, and authority nodes. The obverse and reverse may contain identical elements, including legend, type, symbol, and personal or corporate names defined by xlink:role. Roles may be authority, issuer, mint, region, deity, portrait, or others, as defined in Nomisma. Personal and corporate names may employ a title attribute for arbitrary, uncontrolled titles or positions (e. g., moneyer, emperor, etc.). The type element nested within obverse and reverse may contain repeatable description elements, with language defined by the xml:lang attribute. Numishare supports rendering documents in multiple languages, as long as labels denoted

37 "XML Linking Language (XLink) Version 1.0." Several elements in this XML fragment, such as that relating to the issuer and deity, describe concepts which have not yet been defined in Nomisma. Nomisma has defined Roman emperors, but magistrates from the Republic and early Principate have not yet been. Therefore, only textual descriptors can be used.

by xml:lang exist either within the NUDS/XML record itself or within the RDF representation of the Nomisma URI.

As a practical example, suppose that a NUDS record contains both German and English type descriptions. A language parameter for German (“de”), may be passed to Numishare to display the appropriate type description, and lookups can be performed upon all Nomisma URIs within the document, parsing the RDF for each Nomisma ID and extracting the German SKOS preferred label for serialization into HTML. This move toward internationalization, greatly aided by Nomisma, makes it easier to develop and maintain multilingual interfaces for a single collection, as well as making it possible to query across multiple collections, regardless of the native language of the XML documents. When a NUDS record is published in Numishare and indexed into Solr, the URIs for Nomisma, the Pleiades Gazetteer of Ancient Places, Geonames, and VIAF IDs are stored in the index, enabling queries of the URIs directly (e. g., query all silver coins by the ID http://nomisma.org/id/ar, regardless of local encoding practice for materials). Furthermore, indexing of URIs for Pleiades has led to the integration with the Pelagios project, allowing Roman Imperial coins from Ephesus to be cross-searchable with other archaeological finds from the ancient city or classical texts which reference it.[38]

Typological elements are typically explicitly declared within the <typeDesc> of a NUDS record. This is the case for conceptual records, but not necessarily a requirement for physical ones. The <typeDesc> element itself may include an xlink:href attribute which points to another resource on the web. The Typological Description node can then

38 See http://pelagios.dme.ait.ac.at/api/places/http%3A%2F%2Fpleiades.stoa.org%2Fplaces%2F599612

be left blank, as Numishare is designed to fetch the XML document from the remote resource for processing and display. For example, the object formally identified by 1989.19.5 in the University of Virginia Art Museum (which now belongs to the Virginia Museum of Fine Arts) is type RRC 286/1, referencing the Nomisma ID http://nomisma.org/id/rrc-286.1.[39] The typological metadata for this and many other objects referencing RRC are extracted directly from Nomisma, which greatly reduces the burden of data curation on those responsible for maintaining the U. Va. collection, deferring the long-term maintenance to the editors of Nomisma. Objects in the U. Va. collection that reference Nomisma are constantly up to date with the latest changes.

While the Typological Description section of Descriptive Metadata forms the core of NUDS records representing coin types, other sections comprise records which describe physical specimens. Foremost of these is the Physical Description, which contains a variety of fields for describing measurements, appearance, state of conservation, and other such attributes. The <physDesc> is outlined as follows:[40]

- axis
- chemicalComposition
- countermark
- testmark
- conservationState
  - condition
  - wear
  - completeness
  - grade
- dateOnObject
- measurementsSet

39 See VMFA 2012.88, available online at http://coins.lib.virginia.edu/id/2012.88.
40 For a more detailed look at the XML structure of <physDesc>, see http://wiki.numismatics.org/nuds:nuds#physical_description

- shape

The die axis is one of the most vital numismatic descriptors. The Chemical Composition node within the Physical Description encapsulates in a structured, machine-readable form the results of chemical analyses of coins, at least in theory. No real data from analyses have been used to form a cohesive data structure at this point. Countermarks and test marks can be described generally about the coin as a whole or on the level of the obverse and/or reverse. The <conservationState> includes information about the condition of coins. In numismatic terms, the phrase “secondary treatment” is commonly applied, i. e., modifications made to the coin post-manufacture. Modifications include cutting, clipping, piercing, among others. Physical wear may be described, as well as completeness (coins intentionally or unintentionally cut or damaged) and grade, the professional evaluation of its physical state. Condition, wear, completeness, and grade may also be described on a more granular level of obverse or reverse.

The <dateOnObject> may include a date or series of dates physically present on the object and, optionally, the accompanying era (B.C., A.D., or Hijri, for example). The <dateOnObject> should not be confused with the typological date. While the typological date may coincide with the <dateOnObject>, as is normally the case with modern coinage, dates on the die may be commemorative of past events or refer to the year of minting of an issue spanning multiple years. For example, the United States Washington Quarter type, defined by the head of Washington on the obverse and bald eagle on the reverse was issued from 1932 to 1998. Countless coins of this single type were issued, each with the current year on its obverse.[41]

The Measurements Set is influenced by the Measurements category of VRA Core, though with fields more specific to numismatics, as defined by Heath and Meadows.[42] The diameter field is employed for coins, while height and width may be used to define the dimensions of non-circular items, which are paper notes in most cases. The field for thickness of the object is also included, as is weight, one of the most important descriptors for any coin. A units attribute may be supplied to denote the units of measurement. The default is “mm.” Lastly are a few fields which are unlikely to be applied to the majority of coins: shape, serial number, watermark.

Like the Physical Description section, the Undertype Description may only be applied to records regarding physical specimens. The undertype is defined as the original type for a previously-circulated coin reused as a flan for a new issue. The contents of the <undertypeDesc> are identical to the Type Description, and also like the <typeDesc>, the xlink:href attribute may be employed to link the undertype to URIs expressing coin types.

The optional Findspot Description may be used to express the findspot of the coin. For individual finds, a findspot may be encoded with explicit geographic coordinates or by linking to URIs that define places. Currently, Numishare supports parsing and processing coordinates expressed in Geography Markup Language (GML) or Geonames URIs, and W3C Geospatial Ontologies, once formalized, will be supported in the future. For those coins which belong to a hoard which is expressed electronically as a

41 See U.S. Numismatics catalog.
42 “VRA Projects- VRA Core 3.0.”

stable URI, the NUDS/XML document may link to this URI, enabling geographic metadata to be pulled directly from that resource. A more detailed description of electronic encoding of coin hoards will be discussed at a later point in this paper.

The next metadata section of NUDS/XML, found exclusively in physical record types, is the Administrative Description. The <adminDesc> contains accession numbers or other such identifiers, provenance (including auction catalogs and prices), appraisals, acknowledgments, ownership, and information about the collection to which the object belongs.

The final large descriptive section which may be found within NUDS/XML documents is the Reference Description. This optional section may occur within the <descMeta> of physical and conceptual record types. The <refDesc> contains references about a coin or coin type, whether books or articles published in print media or electronic resources. There is a small handful of fields available within the Descriptive Metadata which are not large enough to fall within the larger descriptive sections: title (a requirement for all NUDS/XML documents), abstract, notes, subjects, and departments. Departments and subjects in particular are peculiar to the holding institution, and thus would not typically link to concepts defined by Nomisma.org.

While many of the NUDS elements have their own unique attributes (@clock for axis or @currency for <salePrice>, for example), there are several attributes which apply broadly to some or nearly all elements within the schema. A “certainty” attribute can be applied to all of the fields in the NUDS <descMeta>. There is a significant level of uncertainty in many areas of Roman numismatics. Insofar as coin types are concerned, anonymous moneyers and unknown mints are common occurrences across *Roman Imperial Coinage*. In other cases, coins may have originated from either of two mints.[43] With respect to records describing physical objects, the attribute may be used to express gradations of certitude where features are unclear, and thus unidentifiable. In theory, one may encounter this usage of the attribute with a badly preserved, archaeologically excavated coin where the legend and iconographic features may only be partially legible. The form of the obverse portrait might exhibit the general shape of Domitian while the reverse legend may read […] DES VIII. The identification of Domitian is likely, but not entirely beyond doubt, and so the certainty attribute is useful for communicating scholarly decisions regarding classification.

In some sense, these two examples from coin type and physical coin records use the attribute to communicate different ideas. On one hand, an unknown mint is the product of a lack of historical or contextual evidence. On the other, uncertainty of features is a product of the lack of clarity of the surface of the object. Adopting the approach of the Text Encoding Initiative schema regarding clarity with respect to written items (e. g., manuscripts) is under consideration. TEI employs an <unclear> element with a "reason" attribute. This concept can be adapted to NUDS in the form of a "clarity" attribute which will be available to all those elements within the schema in which the "certainty" attribute is also allowed. "Reason" can be employed optionally to cite the cause of the use of either "clarity" or "certainty."

43 Cite specific example here.

The second major attribute which may be associated with every element in NUDS is "localType," which, like date elements, is borrowed directly from EAC-CPF. The EAC-CPF tag library defines this attribute as follows:

> "Used to make the semantics of the element more specific or to provide semantic specificity to elements that are semantically weak. Value should be an absolute URI."[44]

Simply put, this attribute can be used to distinguish the semantic meaning of a NUDS element, particularly when the same element is employed by an organization to portray more than one intellectual concept. A current American Numismatic Society curatorial workflow illustrates the functionality of "localType." The curators of the Medals and Decorations department of the ANS define four distinct types of subjects which are used to categorize medals. The Filemaker database that the ANS curators use has four fields for these subject types, but NUDS employs only a singular subject element. The semantic meaning of this element is preserved by the use of the localType attribute after data is processed from Filemaker into NUDS and Numishare. The final aspect of NUDS/XML which requires discussion is the Digital Representations section.

This section is located below the Descriptive Metadata section within the NUDS document. In physical record types, one can embed a fragment of METS to link to digital images of a coin. NUDS documents defining coin types once supported aggregating links to URLs describing physical coins for the purpose of linking types to their associated physical specimens. This has since been replaced by an RDF database

44 "EAC-CPF Tag Library Draft."

architecture, detailed later in this thesis.

Numishare

Having introduced the NUDS/XML data model, it is now possible to discuss the applications which compose Numishare—how they function and how a user interacts with the public interface. As mentioned previously, there are four open-source, Java-based applications that run in Apache Tomcat. Essentially, Numishare is the code which links these separate applications together into one consolidated software suite. XSLT is the predominant scripting language within Numishare, with Javascript and CSS used to enhance the user experience. Apache Cocoon is the primary web publishing platform upon which the Numishare public interface is built. The Cocoon sitemap defines a myriad of pipelines which represent the Model-View-Controller relationship in XML applications development. Output serialization is formed by the combination of the data model (in some cases, NUDS/XML) and a controller and/or view (typically an XSLT stylesheet). The output serialization of a NUDS/XML document can be HTML rendered to a user of the Numishare interface in his or her browser, but XSLT stylesheets can also generate KML for rendering in maps or simply output the XML document itself to a user interested in taking advantage of the raw, unserialized data. The NUDS documents are stored in an eXist XML database, and therefore Cocoon reads the model directly through eXist's web service (or REST, respresentational state transfer) interface.[45]

Cocoon also handles interactions between the user and the Solr search index.

45 For more information about the REST design model, see https://en.wikipedia.org/wiki/Representational_state_transfer.

According to the official website for Solr, “its major features include powerful full-text search, hit highlighting, faceted search, dynamic clustering, database integration, rich document (e.g., Microsoft Word, PDF) handling, and geospatial search. Solr is highly scalable, providing distributed search and index replication, and it powers the search and navigation features of many of the world's largest internet sites.”[46] Netflix was among its earliest adopters in private industry, and it has since become the de facto standard search index throughout the library community in the United States. It serves as the index for Blacklight, VuFind, and various Fedora repository implementations, as well as countless other projects.[47] Solr has been a part of Numishare's application stack since 2007.

The following is a brief description of the functionality of the search and browse interfaces in Numishare: Solr accepts queries following the Lucene query syntax submitted by the user through the interface, and Solr responds through its REST API with query results in the form of XML (although it also supports JSON output). Cocoon intertwines this XML model with an XSLT stylesheet, which renders the results in the form of human-readable HTML. Query results can also serialized into several other formats. Results can be delivered in the form of an Atom feed, which allows users to subscribe to changes made to coins which fit their key area of interest (e. g., coins of Hadrian). More importantly, the Atom feed is the primary means by which data may be harvested from Numishare by programmatic means. The feed contains 100 hits per page, with links to next and previous pages, as well as links to alternates to the default HTML

46 “Apache Lucene - Apache Solr.”
47 Quote code4lib articles

serialization of a coin type, such as RDF, KML, and NUDS/XML. Therefore, it is possible for a script to page through the entire collection and extract the machine-readable NUDS metadata. In addition to the Solr-based Atom feeds, Numishare also supports serializing Solr search results as KML and CSV. While Cocoon handles output for the public user interface, XForms applications processed by Orbeon, which also runs in Tomcat, form the foundation of the private administrative interface.

XForms is a W3C specification for next-generation web forms which adhere to the Model-View-Controller philosophy. Orbeon is used in the administrative back-end to edit complex XML metadata and interact with REST services. XForms applications query various APIs on the web for enhancing controlled vocabulary and manage Create, Read, Update, and Delete (CRUD) operations between the eXist database and Solr index. Specific details about the architecture of Numishare's back-end are too complex to be detailed within the limited space of this paper, but advantages are briefly enumerated below:

1. XForms supports validation. For example, numeric fields such as diameter and weight must adhere to the xs:decimal datatype. In other cases, XForms validation warns the user that certain fields may not be left blank. Invalid data cannot be sent into eXist.
2. Authority control. XForms taps into auto-suggest for local controlled vocabulary delivered from the Solr index or interacts with REST APIs of Geonames, VIAF, and nomisma.org to embed URIs directly into the xlink:href attribute within the NUDS/XML document.
3. Link to images with Flickr's APIs.
4. Easily link to coins in other collections through automated processing.

A fuller description of XForms, as well as other use cases for similar applications are detailed in “XForms for Libraries, An Introduction,” published in Code4Lib 11.[48]

Together, this suite of applications known as Numishare forms the basis for a powerful research tool for those interested in the study of Roman Imperial Coins. Several features of the public interface have already been discussed (KML and Atom feeds), but this paper will now turn toward a more focused discussion of the public interface of Numishare as represented by Online Coins of the Roman Empire.

OCRE: A Powerful Research Tool for Roman Imperial Coinage

Online Coins of the Roman Empire (OCRE), http://numismatics.org/ocre/, is an open access, online equivalent of *Roman Imperial Coinage*, based on identifiers for RIC coin types established by Nomisma.org. Although the data provided by both Nomisma and OCRE for imperial coin types are virtually identical, the Numishare-based user interface provides the user with a more interactive experience than offered by Nomisma. Development on OCRE commenced in early 2011, and the project was officially launched to the public July, 2012.

Searching and Browsing

Despite numerous and significant advances in Numishare since its inception, especially with respect to collections management and publishing in the XForms-driven backend, the application in its core is designed to deliver a simple, yet sophisticated interface for navigating and interacting with coin collections. The search and browse interfaces are good places to begin a discussion about the features of Numishare, for

48 Gruber et al., “XForms for Libraries, An Introduction.”

these are what separates Numishare from most other collections. Solr supports faceted searching based on nine of typological categories: authority, deity, denomination, manufacture, issuer, material, mint, portrait, and region. Full-text search of the entire record is supported, as well as searches of other fields, like obverse or reverse legend or type, date ranges, among others. Search results can be sorted by a handful of fields, including all those listed above.

Importantly, Solr supports wild card searches, which are especially useful for legends that are only partially legible. Such searches can aid in the identification of coins acquired in archaeological excavations which may be in a poor state of preservation. For example, one may search for “P?CI AVG*” within the reverse legend to yield several dozen matches for the legend “PACI AVGUST” or “PACI AVGVSTAE.” A “?” wild card represents zero or one unknown characters while “*” represents one or more. Results may be filtered further by facets. Given this partial legend query, suppose the metal of bronze is readily identifiable, and the obverse portrait, however worn, appears to exhibit the features of the emperor, Domitian. In OCRE, this yields one result (through the reign of Antoninus Pius, ending in A.D. 160). Such queries can be represented by an Atom feed, and the metadata for coin types matching these criteria can be ingested into local databases by automated processes, reducing the data entry workload of archaeologists or other recorders.

With the adaptation of Numishare to the American Numismatic Society's collection, called Mantis, in early 2011, a new faceted geographic search interface was introduced, in addition to the traditional textual interface one typically associated with Solr searches. Javascript was written to connect faceted search with OpenLayers, an open-source Javascript library for mapping. These Javascript functions accept interactions from the user of the interface to query Solr, which replies with XML piped through a Cocoon+XSLT transformation into KML to rapidly update the OpenLayers map. Visualizing queries geographically is potentially immensely useful. One may map the distribution of coins of Pax over time and space (Figure 5) or map the distribution of a particular coin type over the whole of the Roman Empire, which would give an economic historians a glimpse of ancient trade networks.

The final interface constructed upon Solr is one that visualizes search results in the form of charts and graphs. A query of all coins which depict the deity, Victory, may be visualized as a column graph depicting the numeric count or percentage (of the total number of issues) per Roman emperor (Figure 6). In this, Vespasian leads all other emperors. Such visualizations may lead to research questions which may have otherwise never been asked.

Coin Type Records

Like most other aspects of Numishare, the coin or coin type record HTML serialization has evolved considerably since it was initially documented at CAA 2009, although some features of the HTML page have remained.[49] The design of the object

49 Gruber, "Encoded Archival Description for Numismatic Collections."

record page was conceived in 2007 to address the perceived shortcomings of record pages in other databases, namely the “dead-end” nature of those pages. Links from search results typically led to pages for coins which contained metadata, but no way to seamlessly traverse from those records to records of related coins.[50] The 2007 design created links for typological attributes—like denomination, material, and mint—which would direct the user to the search results page for that attribute. One could navigate from a particular denarius of Augustus to other denarii or other coins of Augustus. OCRE continues to provide this functionality, which had an unintended, nevertheless positive, consequence: Google and other robots are able to crawl from record to record through RESTful search results to make records available in search engines. The indexing of OCRE into these search engines broadens access to the collection.[51]

To enhance the user experience, when Cocoon builds the HTML representation of a NUDS/XML document for a coin type, links to Nomisma URIs are generated to the right of each label. Links to Nomisma.org are handled by the Ancient World Linked Data Javascript library (awld.js), developed by Sebastian Heath and Nick Rabinowitz. Awld.js creates a small pop-up window containing summary information about the concept when the user of the site hovers his or her mouse over the link. Additional features of the HTML page for a coin type include links to associated physical objects in other collections, an OpenLayers map which renders points for mints and known findspots

50 See, for example, object number 18207928 of SMK Berlin, http://www.smb.museum/ikmk/object.php?id=18207928. Search results pages are driven by session web architecture, not REST, so it is not possible to navigate from a coin record to a results page featuring coins of a similar typology.

51 From February 1, 2013 to April 15, 2013, Google Analytics indicates that the Google search engine alone accounts for nearly 72% of all traffic to OCRE. Bing is the second-leading search engine with 0.53%.

associated with the coin type, and a section for quantitative analysis of the type. While fairly rudimentary at the present stage, especially given the limited number of physical coins associated with each coin type, this section of the page lists the average weight for the coin type and allows the user to create graphs which compare the average weight of the specified coin type with the average weights of coins with similar attributes (e. g., the average weight of other bronze Roman coins compared to a bronze sestertius of Vespasian, no. 433; Figure 7). This interface for visualizing weights will be expanded to encompass other measurements, as well as standard deviations.

Lastly, the record page links to other serializations of the record: NUDS/XML, RDF, and KML. These links are established with icons represented visually in the browser, as well as in the HTML header to facilitate machine reading.

OCRE and the Role of Linked Open Data Systems

In the first version of OCRE, released July 2012, the NUDS records for imperial coin types contained a list of links to associated coins in the American Numismatic Society collection. Weights, diameters, and images were extracted directly from Mantis and stored in OCRE's Solr index. Updates to Mantis became difficult to maintain in OCRE, and a new approach was taken in associating physical coins with their conceptual counterparts. Moreover, the intense focus on relating physical specimens reflected by stable URIs on the web (e. g., from sophisticated collections management systems developed by the ANS, Berlin, or other museums) caused the author to overlook scenarios in which attribution by coin type is extremely important even if particular physical specifications of a coin are not a factor in recording, and disposition is not captured, for example, in a coin hoard record. Hoard information may not contain weights and diameters of individual to use for analysis, but they do offer geographic data necessary for displaying the extent of a coin's circulation. When OCRE was first launched, this geographic information could not be used unless a record for a physical coin explicitly contained it. None of the Roman imperial coins through the reign of Antoninus Pius in the ANS collection have attested findspots. The integration of coin hoards into OCRE will be discussed in the description of the Coin Hoards of the Roman Republic (CHRR) project, but here we will address the technical design.

On a theoretical level, we seek to associate coin types to other objects represented by those types. In most cases those other objects are physical coins, and those coins typically have a one-to-one relationship with the type, although countermarked coins are an exception. A hoard may have a one-to-one or one-to-many relationship with types. There could be several coins of a single type or many coins classified with hundreds of types. The most efficient method for maintaining the relationships between coin types in OCRE and their associated coins and hoards is with the use of an RDF database and SPARQL endpoint.

Testing for the new system, based on Apache Fuseki, commenced in late 2012, and it was launched into production in spring 2013.[52] The RDF requires three components of data: first, RDF representations of coin types, second RDF descriptions of physical coins or coin hoards, and finally, the RDF data from Nomisma.org, which

52 http://jena.apache.org/documentation/serving_data/index.html

enables links to be made from typological attributes (e.g., http://nomisma.org/id/ar for silver coins or http://nomisma.org/id/augustus for those of Augustus) between types and coins/hoards. The RDF model conforms to the numismatic ontology established by Sebastian Heath and Andrew Meadows on Nomisma.org. The structure for RIC Augustus 1a is similar to the model for RRC 100/2 in a previous section, “The Coin Type as an Intellectual Concept,” represented as nm:type_series_item defined by the URI http://numismatics.org/ocre/id/ric.1(2).aug.1a. A relationship is established by the nm:type_series_item in the RDF model, below, associating object number 18207296 in the Berlin Münzkabinett to Augustus 1a:

```
<rdf:Description rdf:about="http://www.smb.museum/ikmk/object.php?
id=18207926">
        <nm:type_series_item
rdf:resource="http://numismatics.org/ocre/id/ric.1(2).aug.1a"/>
        <dcterms:title>Augustus, ca. 25-23 v. Chr.</dcterms:title>
        <dcterms:publisher>Münzkabinett, Staatliche Museen zu
Berlin</dcterms:publisher>
        <nm:axis rdf:datatype="xs:integer">6</nm:axis>
        <nm:diameter rdf:datatype="xs:decimal">13</nm:diameter>
        <nm:weight rdf:datatype="xs:decimal">1.32</nm:weight>
        <nm:reverseReference
rdf:resource="http://www.smb.museum/mk_edit/images/8099/rs_opt.jpg"/>
        <nm:obverseReference
rdf:resource="http://www.smb.museum/mk_edit/images/8099/vs_opt.jpg"/>
    </rdf:Description>
```

The RDF contains a few metadata fields from Dublin Core (dcterms:title and dcterms:publisher) and a handful of Nomisma-defined concepts for encoding measurements (nm:axis, nm:diameter, and nm:weight). Image URLs are also recorded.

Note that the RDF description of this object in the Berlin collection does not explicitly denote its material (silver), denomination (denarius), or other typological attributes. Like a relational database, an RDF database enables queries by these attributes by the association of the nm:type_series_item between the coin and the coin type. Practically speaking, how does this affect OCRE? Rather than updating records in OCRE when the collections of new partners become available, the RDF database stands apart, and thus it is easier to update Fuseki with new collections and manage changes or deletions in collections already contained in the system. Using SPARQL, OCRE (through the Numishare application code) can query Fuseki to display images of coins associated with particular types in record or search results pages. SPARQL query results can be serialized into KML directly and displayed in maps, making it possible to make use of findspots from online coin hoard catalogs. Using mathematic functions inherent to SPARQL, average weights of coin types or specific typologies (e.g., denarii of Augustus) can be delivered to OCRE directly without complicated processing of numeric figures stored in the Solr search index. The updating of this RDF database is independent of progress made in adding new types into OCRE, and thus more findpots, images, and measurements become available immediately upon ingestion into Fuseki.

OCRE and the Future of Roman Numismatics

OCRE represents a fundamental shift in Roman numismatics. OCRE is freely available to everyone with Internet access, creating a low-cost and portable tool, accessible whether in the field or in a library. The public user interface serves as a powerful tool for searching, sorting, and visualizing Roman coinage in a manner that is simply impossible in the printed catalogs. Since the coin type metadata is open and considered to be canonical representations of entries in RIC, other collections may make use of OCRE's APIs for extracting this metadata for ingestion or reference in their own databases, saving time in the data entry workflow and deferring the burden of maintaining the data to the OCRE project editors. In theory, OCRE can be used by field archaeologists to quickly identify excavated coins, aiding in the classification and interpretation of stratigraphic contexts, even when the excavation lacks numismatic specialists.

OCRE will continue to evolve even after the publication of this thesis. More NUDS records will be added into OCRE, terminating with Anastasius in the early sixth century. New features for quantitative analyses and visualizations will be added into Numishare. Physical coins in other collections will continue to be linked to coin types. At this moment, the Portable Antiquities Scheme, British Museum, and the Münzkabinett of the Staatlichen Museen of Berlin have committed to sharing their data with the project, with some of Berlin's coinage of Augustus already available through OCRE as a proof of concept of the RDF database system. In some sense, OCRE, at the date of this publication, is merely a starting point in the development of one of the most comprehensive and sophisticated tools for the study of Roman coins on the web.

The Coin Hoard as an Intellectual Concept

Coin hoards, like types, are intellectual concepts. Hoards are named or numbered by scholars. Catalogs regarding a thematic subset (Greek, Roman Republic, Greco-Roman Egypt, etc.) have been published over the decades. Sydney Noe's pioneering 1925 publication *A Bibliography of Greek Coin Hoards* was among the first of these catalogs. A number of important hoards found after 1925 inspired Noe to issue a second edition of the work in 1937. In his introduction, he describes the format of the catalog: hoards would be arranged alphabetically by the modern Italianized or Anglicized place name under which the hoard had been published. Noe admits “this often brings difficulties in transliteration of Turkish, Balkan, and Asiatic names, and no hard and fast rule that is consistent has been found.”[53] This is a point to which this paper will return later in this section. Noe also includes deposit dates, when known, a brief summary of hoard contents, and disposition. Finally, Noe cites the article or monograph from which the catalog entry is derived.[54] In the second edition, the author introduces sequential numbering ameliorate problems which arose when there were numerous hoards for a single site (e.g., Taranto or Delos).[55]

Three decades later, the International Numismatic Commission sponsored a new edition of Noe's work, revised and supplemented with new references. Called *An Inventory of Greek Coin Hoards* (IGCH), it was published in 1973 by the American Numismatic Society and edited by Colin Kraay, Otto Mørkholm, and Margaret Thompson. The catalog increased from just under 1200 hoards in 1937 to 2387 with a chronological cut-off of 30 B.C. The hoards were re-arranged from alphabetical to geographical by region (Greece, Thrace, Asia Minor, Italy, etc.), and within these regions, hoards are organized chronologically by deposit date: Archaic (up to c. 480 B.C), Classical (480-330), and Hellenistic (330-30).[56] Importantly, the appendix of IGCH

53 Noe, *A Bibliography of Greek Coin Hoards*, 1.
54 Ibid., 2–3.
55 Ibid., 7.
56 Thompson, *An Inventory of Greek Coin Hoards*., vii.

includes a concordance list between Noe's numbering system and the new one.

For illustrative purposes, let us consider IGCH 195, which is the first hoard listed in Noe's bibliography. IGCH notes that this hoard, published in *JHS* 1896, p. 302, differs slightly from Noe's description owing to a reexamination of the hoard and based on J. N. Svoronos' papers.[57] The contents include 63 silver and 2 bronze coins from a variety of origins. These origins are not described in a standardized way. In some cases, the number of coins corresponding to a place of origin relates to a mint, but in other cases, the number corresponds to regions defined in Barclay Head's *Historia Numorum* or are identifiable regional confederations of city-states for which the particular mint is unknown.[58] With regard to Hellenistic royal coinage, the count corresponds to the authority irrespective of the mint. *An Inventory of Greek Coin Hoards* has been a standard reference work in Greek numismatics for decades, and like RIC, RRC, and other type corpora, hoard entires can be represented as URIs in a linked open data environment.

Nomisma.org maintains a URI for each hoard cataloged in IGCH; this hoard is represented on the web by http://nomisma.org/id/igch0195. The view of this web page is a nearly-verbatim facsimile of the content from IGCH, with the addition of links to Nomisma URIs for mints and geographic coordinates for the findspot. Nomisma.org is able to use these references to URIs defining mints in conjunction with findspot coordinates to generate a KML representation of a hoard for display on a Google Map. If one were to distill http://nomisma.org/id/igch0195 into RDF, nearly all semantic meaning

57 IGCH 195.

58 Noe, *A Bibliography of Greek Coin Hoards*, 5. IGCH 195 includes counts for the Achaean, Aetolian, and Arcadian Leagues, which issued coins within Achaea, Aetolia, and Arcadia, respectively, although the mints within these regions cannot be ascertained.

would be lost. Neither numeric coin counts nor denominations are encoded in a machine-readable format. Heath admits the copy/paste method was an expedient solution to generating the URIs; effort can be placed into making the data better in the future. With more than 2,000 hoards, it is impossible to publish in a single volume each and every permutation of coin counts per typological category. The lack of machine-readability of the IGCH entries on Nomisma reduces the potential for computational analyses of the data. With no standard approach to encoding coin hoards electronically, the American Numismatic Society invested in the development of a draft XML schema in order to address inadequacies in the printed format of hoard catalogs (and by extension, Nomisma's adaptation of IGCH's format to the web), and to serve as an interchange format between existing relational database systems designed for the recording of hoards and finds.

Coin Hoard Description: Print to Database

Development of a schema to fit the specifications of three concurrent projects began in 2011. While still considered to be preliminary, it is nonetheless a working draft on which Numishare has been adapted to display hoard metadata and perform analyses on records. Before the particulars of this XML schema can be discussed, it is imperative to identify the ways in which the contents of hoards have been and continue to be documented in print form.

IGCH represents one of the most simple formats for hoard description. A reference in the inventory rarely exceeds a few dozen lines, with the quantity of coins identified by merely mint or ruler. As hoard excavation methodology progressed considerably since Noe's seminal work in the 1920s and 30s, so too has the detail of data recording. The hoards of IGCH rarely exceeded a few dozen coins in total, but late Roman Imperial hoards which surface in Britain or continental Europe routinely contain hundreds of coins and, occasionally, number in the thousands or tens of thousands. The Normanby Hoard, discovered in 1985 one and a half miles north of a Roman settlement at Owmby, contained nearly 48,000 coins.[59] The catalog, published three years later by Bland and Burnett, consumes 45 printed pages, divided into four general sections encompassing the coinage of the Central, Gallic, and British Empires, as well as irregular coinage (e. g., imitations, forgeries, and uncertain issues). These divisions are further organized by ruler and mint, with each distinct coin type listed in conjunction with its concordance to RIC or Cunetio numbers, the number of coins per type, and weights per coin.[60] The prose description of the Normanby Hoard includes numerous tables of general statistics—the number of coins per emperor from each mint, percentages of coins from each emperor in Normanby compared to other contemporary British hoards, etc. While these tables are certainly useful to researchers, the salient point here is that the researcher may see and use the data only in the way that Bland and Burnett have presented them.

Needless to say, late Roman Imperial hoard publications of recent decades represent the opposite end of the spectrum from IGCH. Therefore, it is necessary from a

59 Bland and Burnett (1988), 114.

60 *The Cunetio Treasure* is the standard reference work for late third century Gallic and British coin types, revising some types found in RIC V and supplementing the catalog with new types revealed through archaeological excavation and exhaustive study of the hoard.

technical standpoint to design a schema capable of handling both descriptive practices, from quantities per general typology (mints, rulers) to granular description of coin types per physical specimen, including encoding of weights, diameters, and die axes.

As previously mentioned, this draft XML schema for encoding hoard information is designed to meet the specifications of three distinct projects. The cornerstone of this technical development was laid in the summer of 2011 with a database schema proposal authored by Sam Moorhead, Finds Advisor for Iron Age and Roman coins at the Portable Antiquities Scheme (PAS), British Museum. Moorhead provides a list of proposed database fields with a roughly one-page description outlining his “misty vision” for a database and its potential uses.[61] Moorhead and Bland intend in the long term to develop a database of Roman coins hoards found in Britain. The database would expand upon Anne Robertson's *An Inventory of Romano-British Coin Hoards* (2000), numbering just over 1,900 entries, to include several hundred additional hoards discovered since its publication. The project, in its earliest phase, would describe contents of hoards with the quantities of generalized typologies in a similar fashion to IGCH in order to expedite publication of the data, with the expectation that more complete typological and physical descriptions could be provided further in the future.

Where Moorhead's specifications differ dramatically from the format of IGCH is in the description of the findspot and its archaeological context. Moorhead desires descriptors not seen in earlier hoard publications—height above sea level, land-use, distance from nearest source of water, among others—which “would enable much more

61 Moorhead, “A Roman Hoard Database – Some Notes.”

effective consideration of the reasons for hoarding."[62] The aggregation of this data into a single database will allow not only visualization and analyses of the contents within a single hoard, but also comparisons across numerous hoards. Moorhead and Bland seek to undertake a major study of "radiate" hoards from Britain. A project of this magnitude requires both a sophisticated data model as well as sophisticated software, and thus the model was developed in conjunction with the expansion of the capabilities of Numishare.

Despite the immense influence of Moorhead's specifications on my development of a hoard description schema and associated software functions, development of the Romano-British coin hoard database has yet to begin in earnest. The British Museum and the University of Leicester received a grant of £645K from The Arts and Humanities Research Council in February, 2013 for this three year project, entitled, "Crisis or continuity? The deposition of metalwork in the Roman world: what do coin hoards tell us about Roman Britain in the 3rd century AD?"[63] More immediately, an opportunity to work with Republican hoard data emerged through partnership between the American Numismatic Society and Kris Lockyear, of the University College London Institute of Archaeology. Lockyear is a researcher specializing in Roman Republican coin hoards who has aggregated hoards into a Microsoft Access database schema that he himself defined. Building on the work of Michael Crawford's *Roman Republican Coin Hoards,* Lockyear's database, *Coin Hoards of the Roman Republic* (CHRR) consists of nearly 700 hoard entries, which in turn include more than 100,000 coins. The majority of these

62 Ibid.
63 "The British Museum and the University of Leicester Announce £645K to Study Roman Hoards Found in Britain."

coins are identified with reference numbers from RRC, with concordances between Crawford's catalog and other Republican corpora. A handful of these hoards include early imperial issues which are identified by RIC numbers.

Thus, after having studied traditional printed hoard catalogs (like IGCH and *Coin Hoards*), Moorhead's conceptual design document, and Lockyear's database, I designed a flexible XML schema that would accommodate these three methods of publication. The schema allows for general enumeration of typologies, seen in IGCH, or enumeration by coin type, including description of the physical attributes of individual coins within the larger hoard. This level of flexibility allows the schema to be applied to any digital collection of hoards. Accordingly, with the feedback of Lockyear, Meadows, and other colleagues, I extended Numishare's core functionality beyond coins and coin types to support the publication of hoards.

NUDS for Hoards

The heavy use of RRC and RIC to identify coins in CHRR makes it possible to extract metadata directly from Nomisma's RDF representation of these types. The use of data from Nomisma is especially important in this case because Lockyear had created a table for coin types, which on one hand was partially duplicative of Nomisma's data, but also rather incomplete for it lacked the mints, issuers, legends, and type descriptions that Nomisma's RDF contains. Much like those curators or archaeologists who may depend upon OCRE for delivering canonical metadata about RIC types, Lockyear can reduce his own workload by outsourcing long-term maintenance of RRC and RIC references to the editors of Nomisma. Other coins listed in Lockyear's database which are not identified by existing catalog numbers are described by a general typology (e. g., “Denarius”).

The schema for hoards is based largely on NUDS, and thus by extension influenced by those other metadata schemas and linked open data frameworks which influenced the development of NUDS. The basic structure of this NUDS-Hoard document is as follows:

- nudsHoard (root element)
  - nudsHeader (header)
  - descMeta (descriptive metadata)
    - title (titles in multiple languages accepted)
    - hoardDesc (hoard description)
      - findspot
      - discovery
      - deposit
      - disposition
    - contentsDesc (contents description)
      - contents (coins)
      - containers
      - otherObjects (objects other than coins—ingots, jewelry, etc.)

Like NUDS, NUDS-Hoard includes a <nudsHeader> which contains metadata about the electronic record itself. The content and structure of this element are identical to that of the <nudsHeader> in NUDS documents. Below the header is Descriptive Metadata which contains two subsections, in contrast to the same section in NUDS, which may contain six descriptive subsections in addition to several sets of descriptors.

The Hoard Description contains information about the findspot: the name and/or coordinates. Like NUDS records for coins, the findspot can link to URIs defined by

Nomisma or Pleiades for ancient places or Geonames for modern ones. Lockyear's CHRR, like IGCH, denotes the findspot by modern place name, and therefore Geonames URIs were adopted. Moorhead's specifications differed from Lockyear's in the level of detail. Recording the name of the nearest town is a requirement, but Moorhead's schema proposal includes parish, district, county, region, and nation.[64] Fortunately, Geonames can be relied upon for delivering this hierarchical geographic data without the need to explicitly input it into a NUDS-Hoard XML record. The XML datastream from Geonames' "get" API for Lincoln, UK (ID 2644487) includes administrative divisions, and therefore one can programmatically derive the hierarchy Lincoln-Lincoln District-Lincolnshire-England-United Kingdom and store them in the Solr search index by simply providing a link to http://www.geonames.org/2644487/ in the NUDS-Hoard record.[65] Furthermore, Geonames provides place names in a myriad of languages and scripts, ameliorating problems of ambiguity and transliteration noted by Noe in *A Bibliography of Greek Coin Hoards.*

Linking to URIs will likely be the most common method of defining the findspot for a hoard, but geographic coordinates can also be explicitly inputted. These coordinates may be public or private (withheld from the public display of the record on the web). This is particularly advantageous when used in conjunction with URI-defined findspots. The public may see that a hoard was discovered near Lincoln, but specific and highly accurate GPS coordinates of the findspot may be recorded for the benefit of

64 Cite schema?

65 http://api.geonames.org/get?geonameId=2644487&username=anscoins&style=full

archaeologists and administrators, but blocked from appearance on the web to prevent looting of the site. As per Moorhead's schema, the landowner can also be recorded, and this would likely be withheld from the public as well.

The Hoard Description contains information about the discovery of the hoard—the date (which may include the notBefore and notAfter attributes to express uncertainty), type of discovery (excavation, agricultural activity, metal detecting, etc.), the finder(s), or other notes about the circumstances of the discovery. After the discovery section is the deposit date. If this date is not explicitly expressed in the NUDS-Hoard record, the date can be derived programmatically by analyzing the contents of the hoard. The date should be explicit when archaeology provides more accurate context than can be derived from the dates of the coinage. Finally, the Hoard Description contains descriptors for the disposition of the hoard described generally as a whole. The disposition can also be defined on the coin level.

Below the Hoard Description section is the Contents Description, which comprises the bulk of the hoard record. The <contentsDesc> contains three subsections: contents (for listing coinage), containers (for defining the materials, dates, and typologies of the containers in which the hoards were placed), and a listing of other objects found in the hoard. It is possible to internally link containers defined in the hoard record to coins and other objects listed in the Contents Description in order to maintain the relationship between the container and its contents, which may be useful in large hoards which consist of more than one container. The latter two subsections are requested by Moorhead, but do not apply to Lockyear's Republican hoards.

With the structure of the NUDS-Hoard document model generally described, attention must be focused more directly on the structure of the numismatic contents section of the record. Two XML elements may appear in the <contents>: coin and coin group (<coinGrp>). The difference between these two elements is that the latter accommodates a “count” element to denote the quantity of coins which adhere to the descriptors listed within the group. Both elements allow any Descriptive Metadata elements found in NUDS—Physical Description, Typological Description, etc., defined by the NUDS XML namespace. A hoard very generally described by denomination may include one denarius and two quinarii in the structure outlined below:

```
<coin>
       <nuds:typeDesc>
              <nuds:denomination xlink:type="simple"
xlink:href="http://nomisma.org/id/denarius">Denarius</nuds:denomination>
       </nuds:typeDesc>
</coin>
<coinGrp count=”2”>
       <nuds:typeDesc>
              <nuds:denomination xlink:type="simple"
xlink:href="http://nomisma.org/id/quinarius">Quinarius</nuds:denominatio
n>
       </nuds:typeDesc>
</coinGrp>
```

Although this structure is simple, statistical analyses can be supported because the data is machine-readable by processors capable of interpreting XML, which is to say, Numishare's Apache Cocoon application layer. Denominations are explicitly captured by XML tags which also link Nomisma URIs. The quantity of quinarii is captured in a XML “count” attribute. This approach differs enormously from Heath's XHTML+RDFa structure of IGCH hoards in that neither quantities nor denominations are encoded in such a way as to be extracted by computational processes. Both the quantities and denominations of coins within IGCH entries on Nomisma are lost in the translation to RDF/XML by the W3C RDFa distiller.

One may expand the above XML fragment further by differentiating the two quinarii geographically by separating the <coinGrp> into two coins, each with a Typological Description including the denomination of quinarius, but one minted in Rome and the other Emerita, for example. Legends and type descriptions may also be inserted into the <typeDesc>. A <physDesc> element may be added for capturing the weight, diameter, and die axis of each coin. Thus, this approach to listing coins and coin groups within a hoard can accommodate the granular descriptive practices applied by Bland and Burnett in *The Normanby Hoard* and other such publications.

Like NUDS records that describe physical specimens which can be identified by types published by Nomisma, the <typeDesc> within a coin or coin group can also utilize the xlink:href attribute to link to Nomisma URIs.

```
<coinGrp count="5">
	<nuds:typeDesc xlink:href="http://nomisma.org/id/rrc-44.5"/>
</coinGrp>
<coin>
	<nuds:typeDesc xlink:href="http://nomisma.org/id/rrc-88.2a"/>
</coin>
```

The fragment above indicates that a hoard comprises of five instances of RRC 44/5 and one of RRC 88/2a. The majority of Lockyear's entries from his Access database were migrated into NUDS-Hoard XML records which contain little more than a reference to a Geonames URI for findspot and a list of Nomisma URIs for *Roman Republican Coinage* numbers and their associated quantities. The data model is relatively simple, especially compared to the organization of Lockyear's original database schema, but the analyses which may be conducted by users of the public interface are quite sophisticated. Like OCRE, Lockyear's CHRR is built upon Numishare, which was heavily modified to support the publication and dissemination of hoard data, a functionality that was never remotely considered when the application was conceived as part of the University of Virginia project in 2007.

Coin Hoards of the Roman Republic: Numishare for Hoards

Coin Hoards of the Roman Republic was officially launched in March, 2013 at the Computer Applications and Quantitative Methods in Archaeology conference, held in Perth Australia. It is available at http://numismatics.org/chrr/. From a presentational standpoint, CHRR does not differ significantly from OCRE. There is an HTML interface for searching and browsing the collection of hoards with keyword search and filtering by facets derived from typological attributes. Search results can be delivered in the form of Atom, KML, and CSV. It also contains a Solr facet-based mapping interface. CHRR's most significant departure from OCRE, or other Numishare collections for coins or coin types, lies in the HTML representation of the NUDS-Hoard record. While the HTML resource for a hoard is not wholly dissimilar to a coin record in that it contains a textual description of the record (including a list of hoard contents), a mapping component, and a quantitative analysis component, the processes by which these aspects of the HTML view are generated differ substantially.

Before the transformation from XML into HTML can begin, the subject referenced by each unique xlink:href attribute associated with Typological Descriptions are imported into an XSLT variable. By default, the script will look for a NUDS datastream by appending '.xml' to the URI. If the <typeDesc> links to http://nomisma.org, '.nuds' is appended to extract the NUDS/XML model provided by Nomisma. A table row for each coin or coin group in the hoard contents is created. The denomination and date are displayed by default. For those coins or groups which link to coin types on Nomisma or OCRE, the <typeDesc> stored in the XSLT variable associated with that particular URI is passed through an XSLT template to generate a block of HTML which is hidden by default, but may be shown by Javascript with a mouse click. The XSLT template for rendering the Typological Description in hoard records is shared with XSLT for generating HTML views for coins and coin types, and therefore the HTML view of typologies in CHRR is identical to OCRE, complete with links to Nomisma or predefined search results (Figure 8).

Visualization and Analysis

Like the view for a coin or coin type, the hoard view also contains a map, but in this case the map is more sophisticated in that it is combined with SIMILE Timeline through the intermediary Javascript library, Timemap.[66] The map for a hoard in CHRR not only shows a point for the findspot and all of the mints associated with it, but the

66 SIMILE Timeline is an open source chronological Javascript library developed by MIT http://www.simile-widgets.org/timeline/). Timemap is developed by Nick Rabinowitz, which is a library that interlinks the timeline function with Google or OpenLayers maps. It is open source and available through Google Code at http://code.google.com/p/timemap/.

timeline shows every datable coin, which makes it easier for users to get a sense of the temporal, in addition to geographic, distribution of coins found within the hoard (Figure 9).

One of the most important features of this project is its analytical functionality. Quantitative analysis of hoards provides insight to scholars on the social and economic conditions in the time and place where the hoard was buried. While IGCH is merely a catalog of hoards containing very little consideration of the their contexts, modern hoard statistics are quantified into charts. Further historical context can be ascertained about a particular hoard by comparing its contents to hoards in neighboring geographic areas. A table in *The Normanby Hoard* shows a listing of the percentage of coins per mid-late third century emperor across fourteen different hoards.[67] The format of the data in *The Normanby Hoard* is purely tabular, but more recent publications, such as *Coin Hoards* X (2010) include computer-generated graphs. Wright includes a column chart showing the percentages of coins in two hoards originating from various mints. The following page in this article includes a representation of these data in a table.[68] In an article about the Gaziantep Hoard in the same volume, Meadows and Houghton include pie charts showing the geographic origins and general categories (Seleucid, Alexanders, etc.) of its contents.[69] According to Hoover, many charts published in this volume were generated in Microsoft Excel and submitted as images exported directly from that program or post-processed in Adobe Illustrator to improve the layout for publication (Figure 10).[70] Thus,

67 Bland, *The Normanby Hoard and Other Roman Coin Hoards*, 116.
68 Wright, “A Late Seleukid Bronze Hoard, C. 1988 (CH 10, 349),” 258–259.
69 Meadows and Houghton, “The Gaziantep Hoard, 1994 (CH 9.527; 10.308),” 147.
70 Hoover, “Coin Hoards X.”

each chart required manual entry into a spreadsheet the categories and values desired for visualization. The scholar is limited by the time available to create new spreadsheets in the categories desired for statistical analysis. A comprehensive database solution like CHRR dramatically improves upon the spreadsheet-based method of chart generation used even in current hoard publications.[71]

With these charts and tables as a model for the arrangement and display of data to which scholars are accustomed, it was immediately apparent that Numishare should be adapted to replicate these statistical analyses and visualizations. With a broader set of typological fields available through Nomisma-defined coin types, it is possible to quantify hoards not only by ruler or denomination, but by mint, region, material, dynasty, date, or even deity. The final category, deity, is rarely, if ever, analyzed in printed hoard catalogs, which is striking considering the distribution of deities may lend insight into a hoard's possible religious context (e. g., votive deposits of coins within a sanctuary).

Hoard analyses are available through two avenues in Numishare (and CHRR, more specifically): on the record page for a particular hoard, under the “Quantitative Analysis” tab, and through the “Analyze Hoards” page, which is accessible from the navigation menu. The user may select from a variety of options to generate results which may be visualized in the form of charts using the Javascript library, Highcharts (Figure 12), or downloaded in a CSV file. For chart visualizations, the user may select the numeric response type between the options of count and percentage (of total coins in a

71 See figures on page 258 and 259 of Coin Hoards X. Figure 11 shows the tabular format of the table represented graphically by Figure 10. Similar quantitative results can be downloaded in the form of CSV in CHRR.

hoard), as well as cumulative percentage when analyzing dates. Next, the user selects the chart format: bar and column for most typologies, but a variety of linear graph options are available for date visualizations. Next, the user may select one or more typological categories listed above, such as denomination or mint. Finally, the user may select hoards for comparison. When visualizing data in chart form, the user may select up to 8 hoards for comparison or up to 30 for generating a CSV file. It may take up to 10 or 15 seconds to process 30 hoards for comparison, which is a significant improvement over manual methods employed before the availability of computers for calculating the data. The scholar is no longer restricted to the presentation of data provided by Bland and Burnett in *The Normanby Hoard*, for example, but now has greater control over setting the parameters of his or her own research questions.

With respect to CHRR, Numishare introduces an option to generate linear graphs based on the cumulative percentage of issue dates within a hoard. This type of graph shows the temporal extent of the hoard, with particular focus on the most vigorous periods of hoarding activity. Perhaps no Republican hoard illustrates this more than the one found at Actium, where nearly 80% of its contents are denarii issued in 31 B.C. by Mark Antony (all from uncertain mints) and buried shortly thereafter (Figure 13).

Integrating CHRR Data into OCRE

After the release of OCRE in the summer of 2012, CHRR entered into a consistent phase of development in preparation for its public launch before the Computer Applications and Quantitative Methods in Archaeology conference held in March, 2013.

Not long into this phase of development, it was realized that CHRR's data would be very useful to OCRE and researchers of Roman Imperial coinage, but there were no means of incorporating it into OCRE given that application's current design. At the time, OCRE was only capable of linking to web resources describing physical coins, and no accommodation could be made for incorporating findspot data from hoard catalogs. Of the 694 hoards cataloged in CHRR, 72 contain references to Augustan types defined in RIC Volume I. These are 72 findspots which can be incorporated in to the maps of OCRE for perhaps one-quarter of all Augustan coin types. These data are too valuable to ignore, and this issue became the impetus for launching the Fuseki RDF database and SPARQL endpoint and redesigning the means by which OCRE interacts with objects associated with imperial coin types.

The premise behind OCRE-Fuseki integration has already been discussed in this thesis, but here we will briefly note CHRR's ingestion into the RDF database—a path other hoard catalogs might follow to become available in OCRE or other Numishare-based coin type corpora. Like RDF describing physical coins, the RDF describing a hoard associates its own URI to URIs of coin types by means of the Nomisma-defined nm:type_series_item. There may be many nm:type_series_item references in the RDF for each hoard. In the case of CHRR, most types will be be attributed to RRC numbers in Nomisma, but a significant minority link to types in OCRE. See below, an abbreviated RDF fragment (numerous RRC type references have been omitted):

```
<rdf:RDF>
   <rdf:Description rdf:about="http://numismatics.org/chrr/id/CAP">
```

```
      <dcterms:title xml:lang="en">Capilna (Romania;
CAP)</dcterms:title>
      <dcterms:publisher>American Numismatic Society</dcterms:publisher>
      <nm:findspot rdf:resource="http://www.geonames.org/682812/"/>
      <nm:closing_date rdf:datatype="xs:gYear">-0004</nm:closing_date>
      <nm:type_series_item rdf:resource="http://nomisma.org/id/rrc-
44.5"/>
      <nm:type_series_item
rdf:resource="http://numismatics.org/ocre/id/ric.1(2).aug.287"/>
   </rdf:Description>
</rdf:RDF>
```

This RDF associated Lockyear's Capilna (CAP) hoard in CHRR contains two coin types listed amongst its contents: RRC 44/5 and RIC Augustus 287. Additionally, the findspot is defined by the Geonames place, http://www.geonames.org/682812/, which defines Căpâlna, Romania. Once this RDF has been inserted into Fuseki, the OCRE record for Augustus 287 will immediately show a point on its map for this small Romanian town. The XSLT stylesheet which generates the KML file for this coin type (http://numismatics.org/ocre/id/ric.1(2).aug.287.kml) uses SPARQL to query its findspots from the database and process them into KML Placemarks, performing lookups to Geonames.org to extract latitude and longitude for each geographic identifier. Since Geonames is an authoritative resource for modern place names, storing coordinates explicitly in Fuseki is unnecessary. The closing date, defined by Nomisma as http://nomisma.org/id/closing_date and shown in the fragment above as nm:closing_date, can be used to insert dates into the KML, making it possible to render the temporal extent of a coin type's circulation in OCRE using the TimeMap library (Figure 14).

CHRR and the Future of Roman Numismatics

Coin Hoards of the Roman Republic was a logical starting point in the reconceptualization of coin hoards from the well-curated relational database developed by Kris Lockyear into a linked open data architecture. Lockyear's database structure evolved over the course of twenty-five years, inspired by the standard format established by print catalogs, such as IGCH and Crawford's *Roman Republican Coin Hoards*, but differed from these by its electronic nature. Lockyear meticulously created an entry for every unique coin type referenced in RRC, although he included only date and denomination. While Lockyear was able to query his own database within a Microsoft Windows environment, his relational database model would not have translated well into a web platform. Since the types referenced in CHRR are almost exclusively defined by RRC and RIC, his project served as an initial proof-of-concept for a web application based on an XML-based schema which made extensive use of linked open data principles.

Lockyear's CHRR makes use of data provided by the American Numismatic Society through Nomisma and OCRE while his own data are contributed back to OCRE. These projects are reciprocal. As a result, his efforts benefit both scholars of Republican coinage who use CHRR directly and scholars of Augustan coinage who make use of the mapping features of OCRE.

Where do coin hoards go from here? The logical answer is backward and forward, but there is much work that remains in order to exploit this new framework. While IGCH can be converted from the XHTML+RDFa of Nomisma IDs to NUDS-Hoard with some combination of programmatic processing and manual labor (a task which could be completed in a few short weeks or months), the resulting web project would not be as analytically powerful as CHRR. IGCH contents are described in a haphazard mishmash of denominations, mints, and authorities. Only by establishing Nomisma-defined type corpora for Greek coinage can IGCH realize its full potential, which itself is a gargantuan task which dwarfs the development of a standalone Numishare-based IGCH project.

On the other end of the spectrum are Roman Imperial coin hoards. Databases of finds already exist, united through the European Coin Find Network. Bland and Moorhead plan to implement a database of Romano-British coin hoards. Perhaps NUDS-Hoard will be its native format, or perhaps it will be a transmission standard. OCRE remains 30,000 coin types (roughly 75%) short of being a comprehensive representation of *Roman Imperial Coinage*. In order for existing databases of coin finds and hoards to integrate with OCRE, these later types must be defined by URIs, and databases must be willing to create concordances in their database between their system of references and Nomisma-styled URIs. These advancements in Greek and Roman numismatics will undoubtedly be made, but it may take many years for them to be realized.

The U.Va. Art Museum Numismatic Collection

All of the progress made in the dissemination and analysis of coins, described in the preceding sections, would not have been possible if not for this author's early work on digitizing the University of Virginia Art Museum Numismatic Collection. At that time, there were few numismatic collections available on the web, and the majority of those that were had cumbersome and unintuitive user interfaces. The U.Va. project was developed with only a meager understanding of the terminology of numismatists and their methods of classification, but it was impossible to anticipate at the time the role the project would have in shaping the current direction of semantic web applications in numismatics, or to co-opt a familiar phrase in the humanities, “digital numismatics.” XML, Cocoon, and Solr continue to be cornerstones of this framework, but many of the other aspects of this architecture have evolved considerably in the five years since the U.Va. project was established.

Until early 2013, the U.Va. numismatic collection existed in a silo. It resided on a server maintained by the Scholars' Lab and ITC, and although its coins are indexed into Google and also available through the Library's search interface, Virgo, the project did not interact with other related data portals. It could not: it predated the Pleiades Gazetteer of Ancient Places, Pelagios, and Nomisma.

Reflecting the culmination of advancements made in digital numismatics (including online coin type corpora and coin hoards), the U.Va. collection was re-launched in the spring of 2013 with data models migrated from EAD to NUDS in the newest version of Numishare. Although it is a small collection of slightly fewer than 600 coins, it reflects nearly every category of technical advancement made. It is essentially a microcosm of larger public collections of ancient coins, such as those at the British Museum or Bibliothèque nationale de France. Nearly all of the coins of the University of Virginia are ancient, and therefore the typological attributes of most of the coins are represented by Nomisma IDs. New URIs were created to accommodate numerous authorities (third century Roman emperors, Hellenistic kings, and Greek civic confederations, like the Aetolian League) in the U.Va. collection that Nomisma lacked. The Mosberg collection draws data from Nomisma-defined URIs for *Roman Republican Coinage*.[72] Roman Imperial coins comprise more than half of the collection, but only a few of these date to the reign of Antoninus Pius or before, and therefore link to OCRE. The rest of the imperial collection will link to types defined by OCRE in due course. Additionally, more than half of the collection comes from two coin hoards excavated in the 1980s: the Normanby and Oliver's Orchards Hoards. While these hoards may eventually be expanded upon in Bland and Moorhead's database (which will be able to link back to U.Va.'s collection), Nomisma IDs were created for these two hoards, enabling the display of these findspots in Numishare's mapping interfaces.

To summarize the migration workflow, the Encoded Archival Description XML fragment for each coin was transformed into NUDS/XML, with extensive linking to Nomisma.org for most coins and OCRE for the few which belong to the periods currently represented in that project. Linking to mints defined by Nomisma has opened the door to participating in Pelagios by means of Pleiades URIs extracted from Nomisma's RDF. Those coins with types defined by URIs (either Nomsima for RRC or OCRE for RIC) have been submitted into the Fuseki RDF database. OCRE now displays these coins and makes use of their measurements for quantitative analysis. Finally, the latest features of Numishare are available to the University of Virginia collection: mapping, rendering

72 These coins have since been sold to the Virginia Museum of Fine Arts, but the records will remain in the U.Va. Art Museum Numismatic Collection website until the VMFA publishes their own database online. The VMFA is noted as the repository for these coins.

quantitative analyses in the form of charts and graphs, downloading CSV, and extracting data from the collection by use of its Atom/OpenSearch-based REST APIs.

While the University of Virginia Art Museum Numismatic Collection continues to be among the most technologically advanced coin collections on the web, the most important feature of the updated project is not the mapping or analytical interface built into Numishare, but the sharing of its valuable data to other portals. Scholars are able to make use of these data, aggregated together with the contributions of other collections, to consider patterns which may never have become apparent without these recent technical advances in the field. Involvement in OCRE and Pelagios undoubtedly increases exposure and recognition of the collection, despite its humble size.

Conclusions

This paper reflects the advancements made over the last three years in publishing and interacting with numismatic collections on the web. The trends and standards detailed in this paper are relevant at the time of its publication, but will likely have shifted further in the coming years. The goal of this paper, then, is to document the thought processes which contribute to the current phase of the discipline's evolution, from 2012 to 2013.

Nomisma.org and Numishare, with the OCRE and CHRR projects, represent a shift from traditional relational databases (very much influenced by printed catalogs of the last century) to a new system of linked open data, which affects not only numismatics in the present, but is steadily being applied to other objects of cultural heritage, regardless of the type of object, its origin, or its time period. Projects like OCRE, CHRR, and the University of Virginia Art Museum Numismatic Collection perhaps signify the future of numismatics, where data and images from many collections—public or private, museums or auction houses—become available to everyone with access to the Internet; coins and hoards will be associated together with concepts defined by Nomisma.org, regardless of language or peculiarities in the curatorial practices of individual organizations. The models and frameworks discussed in this thesis will make it feasible, at last, to accomplish the Greek corpus first dreamt by Lazius five centuries ago and left unrealized by Imhoof-Blumer and his successors in the early twentieth century. While we are still many years from this possibility, the projects detailed above are an important first step in the larger aim of broadening access to coins, enhancing their appreciation among the general public, and improving the efficiency of academic research, that is to say, to enable scholars to perform analyses much faster than traditional analog methods and pose new questions in ancient studies which may have otherwise never have been considered.

Bibliography

"Apache Lucene - Apache Solr." Accessed March 3, 2013. http://lucene.apache.org/solr/.

Berners-Lee, Tim. "Linked Data - Design Issues." Accessed February 28, 2013. http://www.w3.org/DesignIssues/LinkedData.html.

Bland, Roger, ed. *The Normanby Hoard and Other Roman Coin Hoards*. London: British Museum Publications, 1988.

"British Museum - Collections Data." Accessed February 28, 2013. http://collection.britishmuseum.org/.

Clain-Stefanelli, Elvira E. *Numismatics: An Ancient Science*. Washington: U.S. Govt. Print. Off., 1965.

"EAC-CPF Tag Library Draft." Accessed March 3, 2013. http://www3.iath.virginia.edu/eac/cpf/tagLibrary/cpfTagLibrary.html.

Eckhel, Joseph, Anton Steinbüchel von Rheinwall, A. L. (Aubin Louis) Millin, and Emerich Thomas Hohler. *Doctrina numorum veterum*. Vindobonae : sumptibus Josephi Vincentii Degen, 1792. http://archive.org/details/doctrinanumorum03hohlgoog.

Foy-Vaillant, Jean. *Numismata Imperatorum Romanorum praestantiora a Julio Caesare ad Postumum et Tyrannos*, 1699.

Gnecchi, Francesco. *Roman Coins: Elementary Manual*. 2d ed.rev. London: Spink, 1903.

Goltzius, Hubertus. "Huberti Goltzii De Re Nummaria Antiqua, Opera Quae Extant Universa.," 1708. http://search.lib.virginia.edu/catalog/001963195.

Gruber, Ethan. "Encoded Archival Description for Numismatic Collections." In *Proceedings of the CAA Conference*, 98–105. Williamsburg, Virginia: Archaeopress, 2009. http://www.academia.edu/1099828/Encoded_Archival_Description_for_Numismatic_Collections.

Gruber, Ethan, Chris Fitzpatrick, Bill Parod, and Scott Prater. "XForms for Libraries, An Introduction." *Code4Lib Journal* 11, no. 11 (2010): 1–11.

Gruber, Ethan, Sebastian Heath, Andrew Meadows, Daniel Pett, Karsten Tolle, and David Wigg-Wolf. "Semantic Web Technologies Applied to Numismatic Collections." In

Proceedings of the CAA Conference. Southampton, UK, 2012.

Harpring, Patricia. "Getty Vocabularies and Linked Data." August 2012. http://www.getty.edu/research/tools/vocabularies/Linked_Data_Getty_Vocabularies.pdf.

Head, Barclay Vincent, and Reginald Stuart Poole. *Catalogue of the Greek Coins of Ionia*. The Trustees, 1892.

Heath, Sebastian. "Numismatic Description Standard (NUDS)." *Nomisma.org*. Accessed February 27, 2013. http://nomisma.org/nuds/numismatic_database_standard.

Hoover, Oliver. "Coin Hoards X," February 8, 2013.

"Linked Data | Linked Data - Connect Distributed Data Across the Web." Accessed February 28, 2013. http://linkeddata.org/.

Meadows, A. R., and Arthur Houghton. "The Gaziantep Hoard, 1994 (CH 9.527; 10.308)." In *Coin Hoards*, edited by Oliver D Hoover, Andrew Meadows, and Ute Wartenberg. Vol. X. New York; Oxford: American Numismatic Society, 2010.

Moorhead, Sam. "A Roman Hoard Database – Some Notes," June 15, 2011.

Moorhead, Sam, Anna Booth, and Roger Bland. *The Frome Hoard*. London: British Museum Press, 2010.

Noe, Sidney P. *A Bibliography of Greek Coin Hoards*. New York: American Numismatic Society, 1937.

Occo, Adolf. *Imperatorum romanorum numismata, a Pompeio magno ad Heraclium*, 1625.

Ogilvie, Brian W. "Petrarch's Historicism and the Purpose of Philosophy." Western Michigan University, Kalamazoo, 2000.

Sutherland, C.H.V., and R.A.G. Carson, eds. *The Roman Imperial Coinage*. Revised ed. London: Spink, 1984.

"The British Museum and the University of Leicester Announce £645K to Study Roman Hoards Found in Britain." *Portable Antiquities Scheme*. Accessed April 11, 2013. http://finds.org.uk/news/stories/article/id/244.

Thompson, Margaret. *An Inventory of Greek Coin Hoards*. Edited by Otto Kraay Mørkholm. New York: American Numismatic Society, 1973.

"VRA Projects- VRA Core 3.0." Accessed February 27, 2013. http://www.vraweb.org/projects/vracore3/categories.html.

Wright, Nicholas L. "A Late Seleukid Bronze Hoard, C. 1988 (CH 10, 349)." In *Coin Hoards*, edited by Oliver D Hoover, Andrew Meadows, and Ute Wartenberg. Vol. X. New York; Oxford: American Numismatic Society, 2010.

"XML Linking Language (XLink) Version 1.0." Accessed April 17, 2013. http://www.w3.org/TR/xlink/.

Appendix

I. Figures

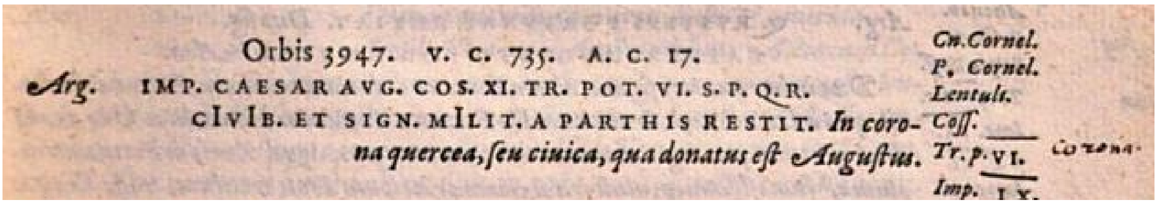

Orbis 3947. V. C. 735. A. C. 17.

Arg. IMP. CAESAR AVG. COS. XI. TR. POT. VI. S. P. Q. R.

CIVIB. ET SIGN. MILIT. A PARTHIS RESTIT. *In corona quercea, ſeu ciuica, qua donatus eſt Auguſtus.*

Cn. Cornel.
P. Cornel.
Lentuli.
Coſſ.
Tr. p. VI.
Imp. IX.

corona.

Figure 1: Occo - Page 159

Arg. CAESAR AVGVSTVS SIGNIS RECEPTIS. *Parthus reſtituens ſigna militaria, de quibus Sueton. Parthi ſigna militaria quæ M. Craſſo & M. Antonio ademerant, repoſcenti reddiderunt.*

MARS VLTOR. *Caput Martis galeatum.*

L. AQVILLIVS FLORVS IIIVIR. *Caput radiatum. Item Parthus, ut ſuperior Galli.*

CIVIB. ET SIGN. MILIT. A PARTH. RECVP. *Arcus triũph. cui inſiſtunt quadrigæ triumph. cum duabus figuris, quæ milites ſunt ſigna militaria in altum exporrigentes, nempe aquilam Legionariam quæ argentea erat, quo magis eſſet conſpicua & ſigna cohortium, cui de more affixæ erant Impp. imagines.*

TVRPILIANVS IIIVIR. *Caput hedera coronatũ. Goltz. Aug. 38.*

Arg. TVRPILIANVS IIIVIR FERON. *Caput Feroniæ.*

CAESAR AVGVSTVS SIGN. RECEP. *Parthus genibus flexis reſtituens ſigna militaria.*

Cap martis

Cap radiatum.

Cap. hedera cor

parthus genibus flexis

Figure 2: Occo - Page 157

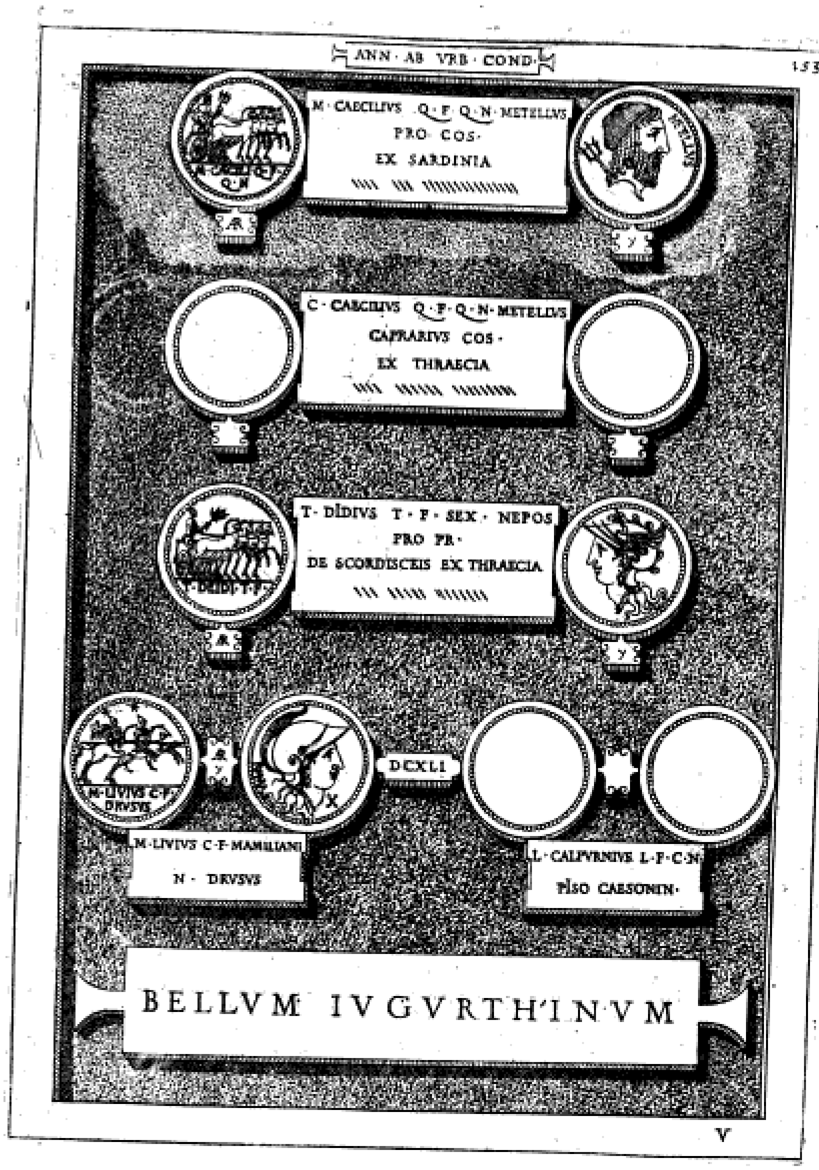

Figure 3: Goltzius - Page 153

Circ. B.C. 450—415.

Tetradrachm.

11	205·	Æ 1·	ΕΦ ΕΞΙ Ʌ Ο Bee with curved wings: border of dots.	Incuse square quartered. [Pl. IX. 4.]

Drachms.

12	49·	Æ ·65	Similar.	Similar.
13	47·7	Æ ·55		

Figure 4: BMC Greek Ephesus 11-13

Figure 5: Distribution of Pax, 27 B.C. - A.D. 160

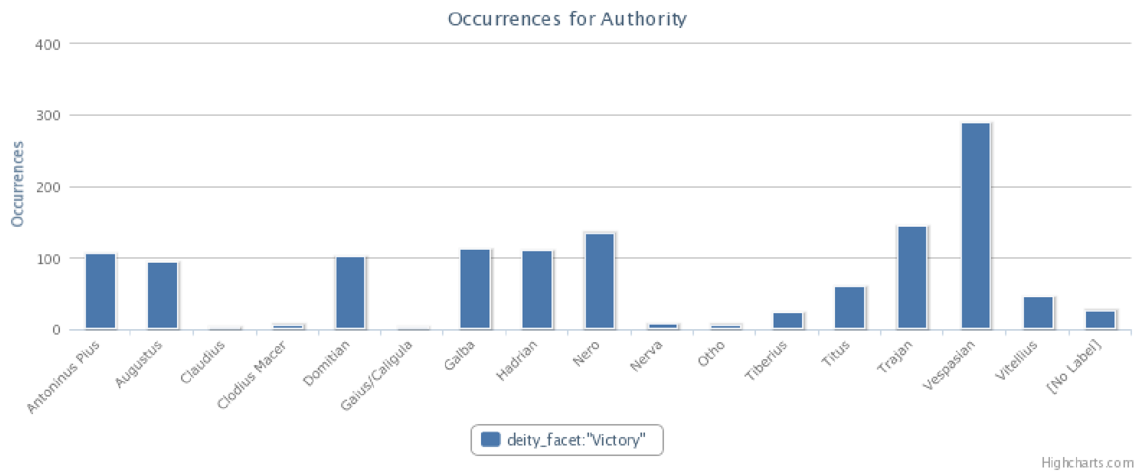

Figure 6: Distribution of Victory by Emperor

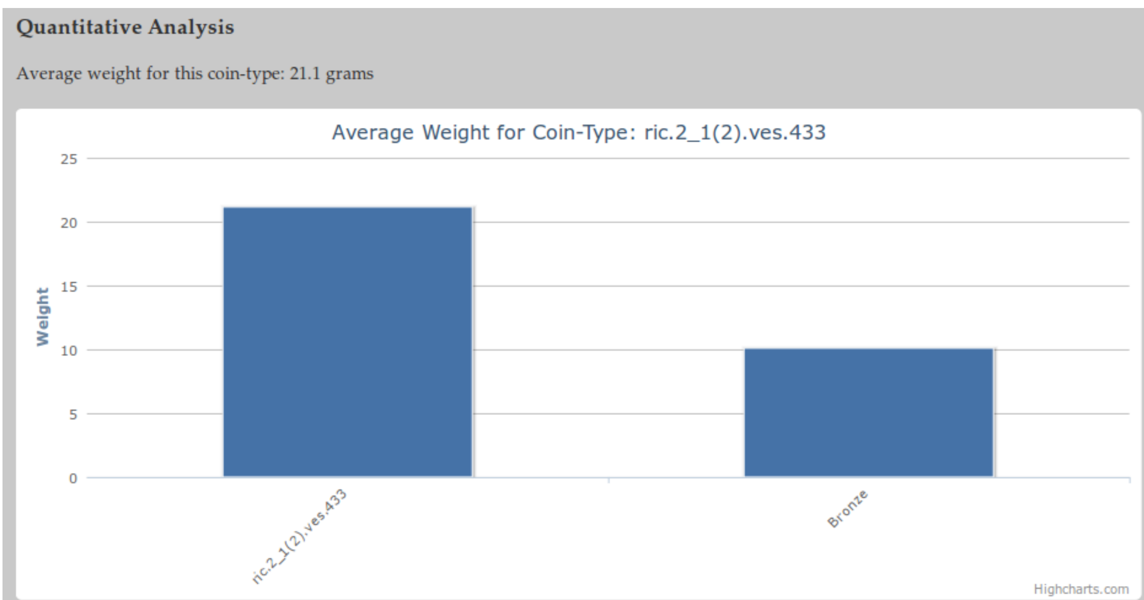

Figure 7: Comparison of Vespasian 433 with Bronze coins

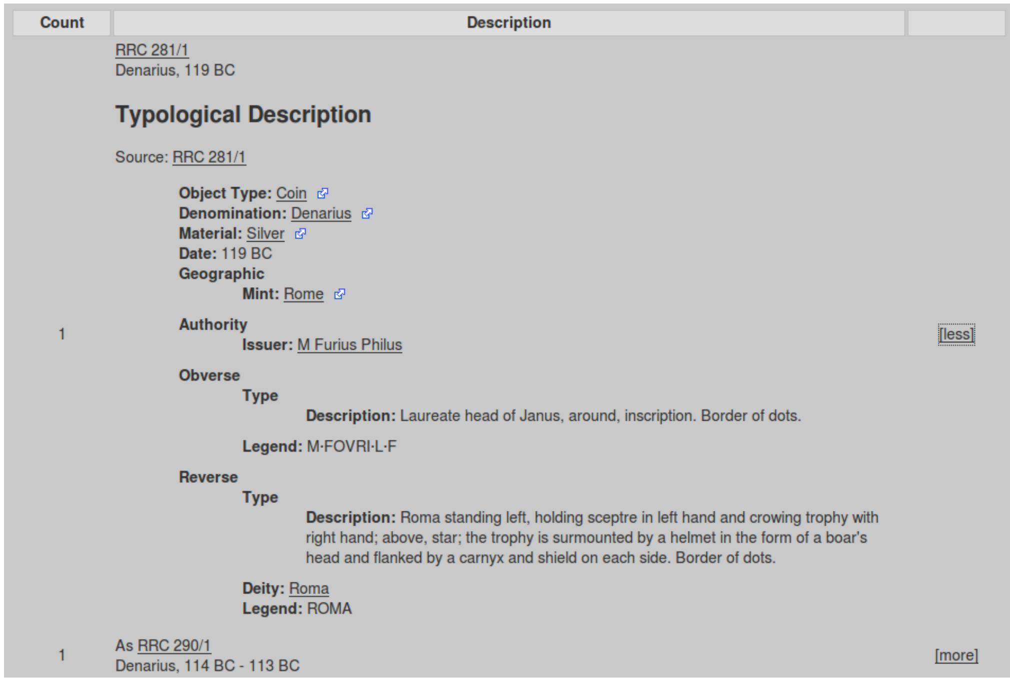

Figure 8: Screenshot of Hoard Record Page

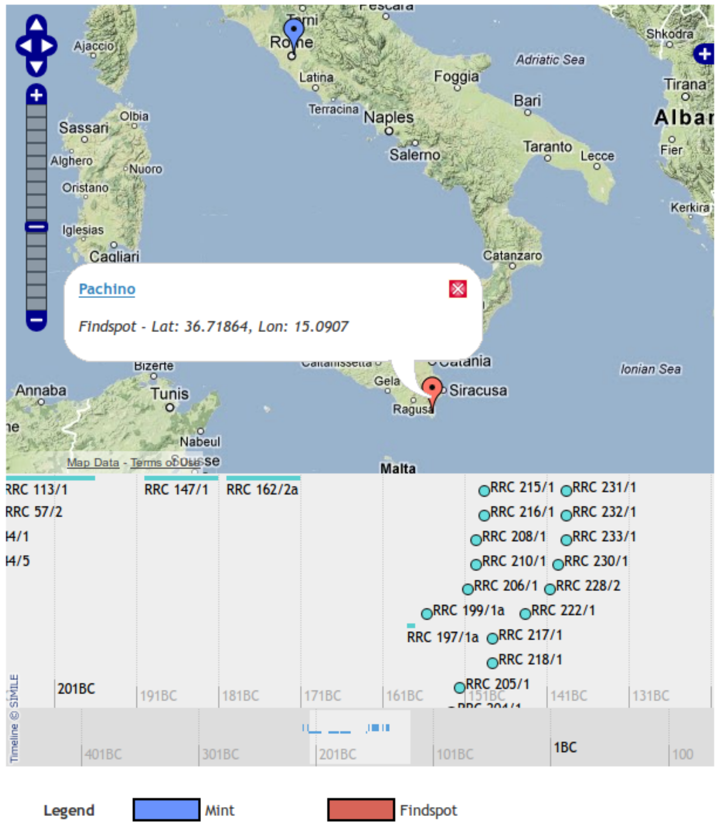

Figure 9: Temporal and Geographic Distribution of the Pachino Hoard

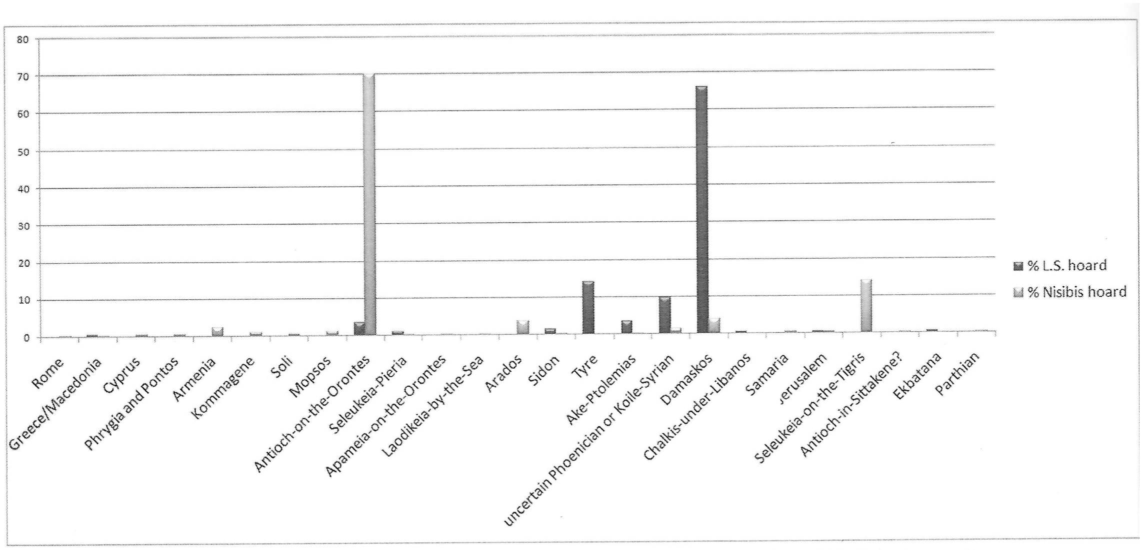

Figure 2. Mints represented in comparison with the Nisibis hoard.

Figure 10: Comparison of mints between two hoards. Coin Hoards X, page 258.

Mints represented in hoard	% L.S. hoard	% Nisibis hoard
Rome	0	0.19
Greece/Macedonia	0.41	0.19
Cyprus	0	0.39
Phrygia and Pontos	0	0.39
Armenia	0	2.14
Kommagene	0	0.78
Kilikia	0	1.36
Antioch on the Orontes	2.90	69.98
Seleukeia in Pieria	0.4	0.19
Apameia on the Orontes	0	0.19
Laodikeia by the Sea	0	0.19
Arados	0	3.51
Sidon	1.24	0.19
Tyre	14.05	0
Ptolemaïs (Ake)	3.72	0.19
Uncertain mint in Phoenicia or Koile Syria	9.50	1.17
Damascus	66.51	3.70
Chalkis under Libanos	0.41	0
Samaria	0	0.39
Jerusalem	0.41	0.39
Seleukeia on the Tigris	0	14.04
Antioch in Sittakene	0	0.19
Ekbatana	0.41	0
Parthia	0	0.19

Table 1. Attributable mints.

Figure 11: Tabular results for mints in Figure 10.

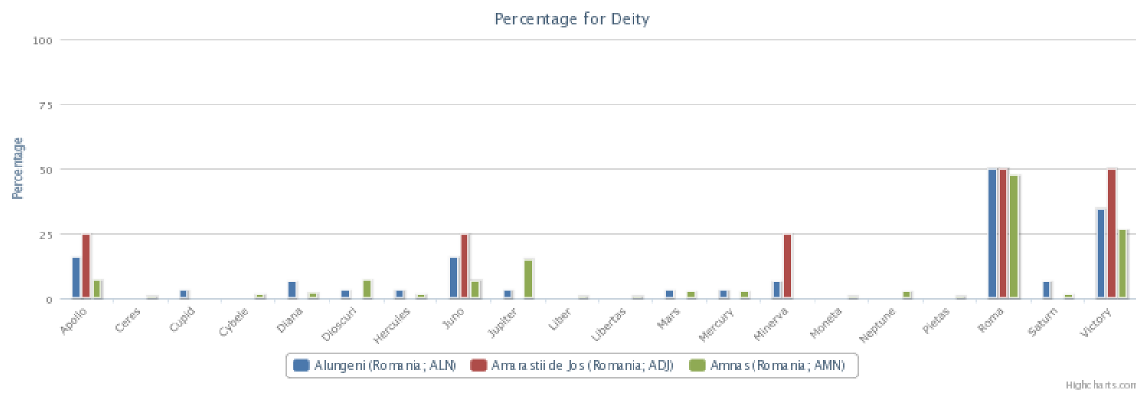

Figure 12: Visualization of Deities in a Selection of Romanian Hoards

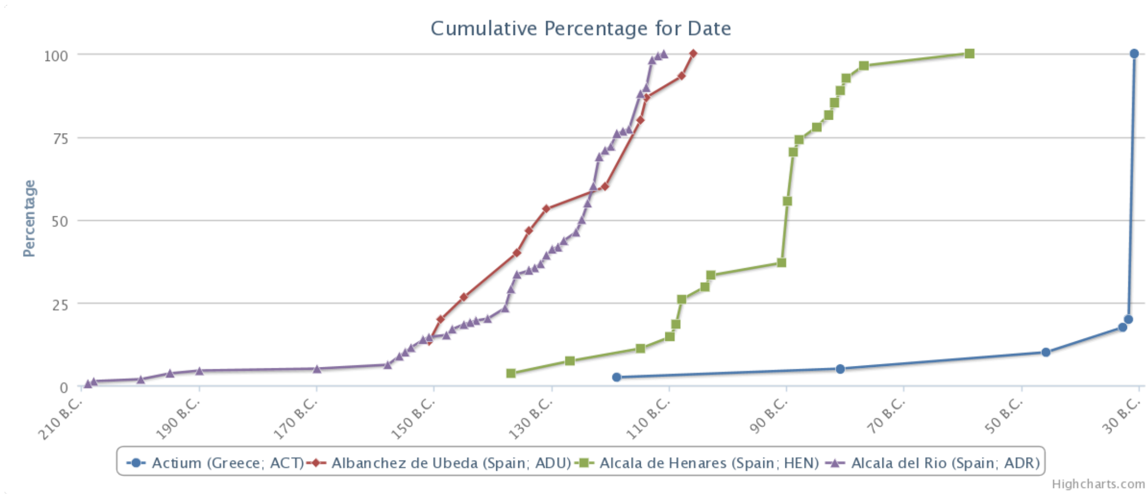

Figure 13: Cumulative Date Percentage Comparison

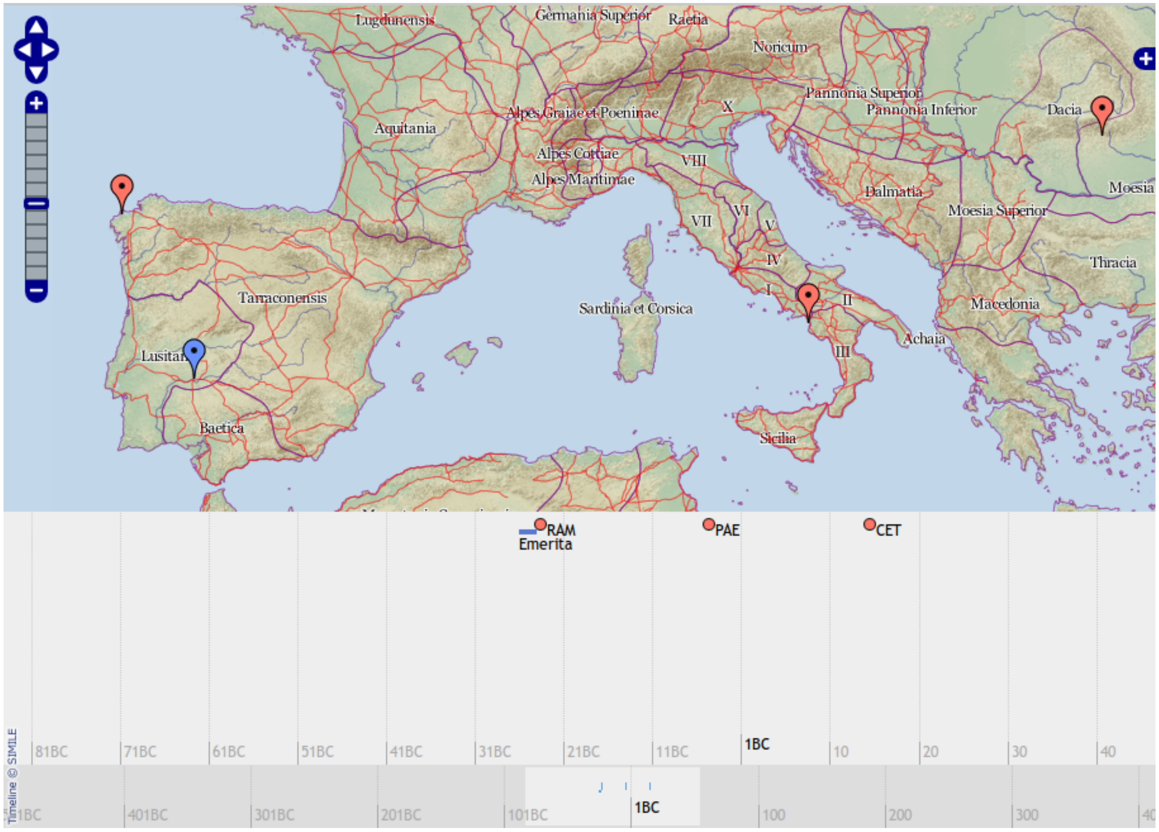

Figure 14: Distribution of Augustus 2b

II. Glossary

Technical Terminology

API – Application Programming Interface, which in web terms, refers to HTTP request and response methods. Also referred to as a "web service."

CIDOC-CRM – A conceptual reference model developed by the International Council of Museums for describing concepts and relationships in cultural heritage documentation. It is a recognized ISO standard (http://cidoc-crm.org/).

Cocoon – An open-source Java-based web application maintained by the Apache Foundation which drives the Numishare public interface (http://cocoon.apache.org/).

CSV – Comma separated values, tabular data (e. g., spreadsheet data) represented in a plain text file. Numismatists are accustomed to using spreadsheet files to generate graphs in Microsoft Excel or similar programs.

EAD – Encoded Archival Description, an XML markup schema for describing archival structure and contents.

eXist – The open-source, Java-based XML database component to Numishare (http://exist-db.org).

Fuseki – A Java-based SPARQL server maintained by the Apache Foundation (https://jena.apache.org/documentation/serving_data/). It powers the RDF/Linked Open Data aspects of OCRE and is planned to be implemented in the next version of Nomisma.org.

GML – Geographic Markup Language, an ontology for describing geographic attributes: points, polygons, etc.

HTML – Hypertext Markup Language, the cornerstone of "web pages" on the World Wide Web.

JSON – Javascript Object Notation, a text-based model for transmitting data on the web. The syntax varies from XML, but the structure is similar.

KML – Keyhole Markup Language, an XML notation for geographic annotation, developed for use in Google Earth.

LCSH – Library of Congress Subject Headings (http://id.loc.gov/authorities/subjects.html).

Linked Open Data (LOD) - “recommended best practice for exposing, sharing, and connecting pieces of data, information, and knowledge on the Semantic Web using URIs [Uniform Resource Identifiers] and RDF [Resource Description Framework].” See http://en.wikipedia.org/wiki/Linked_data for more information.

NUDS – Numismatic Description Standard, an XML schema following linked data principles and numismatic ontologies for describing coins and coin types.

NUDS-Hoard – an XML schema, partially derived from NUDS, for encoding coin hoard data.

Numishare – The open-source collection of scripts and stylesheets which bind various Tomcat applications together to form a publication framework for coins, coin types, and coin hoards (https://github.com/ewg118/numishare).

Orbeon – The XForms processor component to Numishare. Orbeon drives the administrative interface used for creating, editing, deleting, and publishing objects in Numishare (http://www.orbeon.com/).

RDF – A conceptual model used to describe resources and their relationships to one another.

RDFa – Resource Description Framework – in – attributes. W3C standard for embedding semantic web content within HTML using attributes.

RDFa Distiller – A web service for converting HTML conforming to RDFa into an RDF serialization, e.g., XML.

RDF/XML – A data serialization of RDF into XML.

SKOS – Simple Knowledge Organization System, an ontology for describing classification systems and thesauri.

Solr – An open-source Java-based search server maintained by the Apache Foundation used to power the search features in Numishare (http://lucene.apache.org/solr/).

SPARQL - SPARQL Protocol and RDF Query Language (*http://en.wikipedia.org/wiki/SPARQL*).

Tomcat – an open source web server and servlet container developed by the Apache Foundation.

TypeDesc – Typological Description, a wrapper element in a NUDS document which contains typological attributes.

URI – A Uniform Resource Identifier is string of characters used to identify a web resource.

VIAF – The Virtual International Authority File. An OCLC-sponsored project to create concordances between name authority files of various national libraries (http://viaf.org/).

VRA Core – XML standard developed by the Visual Resource Association for describing art historical content.

World Wide Web Consortium (W3C) – International standards organization for the World Wide Web.

XForms – A W3C XML-based specification defining next-generation web forms, designed to process XML models. See http://en.wikipedia.org/wiki/XForms for more information.

Xlink – XML Linking Language, a W3C specification for linking resources, internally or externally.

XML – Extensible Markup Language. A World Wide Web Consortium-defined markup language for encoding documents in a format that is both human and machine readable.

XSLT – Extensive Stylesheet Language Transformations, a functional language for transforming XML documents into other schemes, usually other XML formats.

Bibliographic and Digital Numismatics Projects Abbreviations

CAA – Computer Applications & Quantitative Methods in Archaeology conference.

CHRR – Coin Hoards of the Roman Republic (http://numismatics.org/chrr/), joint project between University College London researcher Kris Lockyear and the American Numismatic Society.

ECFN – European Coin Find Network, an online aggregation engine for the myriad of national Roman coin find databases of Europe.

IGCH – *Inventory of Greek Coin Hoards* (1973).

Nomisma – Collaborative effort to provide stable identifiers for numismatic concepts (http://nomisma.org).

OCRE – Online Coins of the Roman Empire (http://numismatics.org/ocre/), join project between the American Numismatic Society and the Institute for the Study of the Ancient

World at New York University.

RIC – *Roman Imperial Coinage.*

RRC – *Roman Republican Coinage*.

RRCH – *Roman Republican Coin Hoards* (1968).

SNG – Sylloge Nummorum Graecorum (http://www.sylloge-nummorum-graecorum.org/).

III. Contributors

As discussed in the introduction to this thesis, while I am responsible for much of the data modeling and the programming aspects of the projects detailed above, input from colleagues on conceptual or technical matters was an invaluable contribution to the thought processes involved in my development of these projects. This appendix provides a list of contributors and their roles in the projects.

Benton, Jared – Responsible for the identification of the coins of the Mosberg Collection in the University of Virginia Art Museum Numismatic Collection.

Bland, Roger – Keeper of the Portable Antiquities Scheme, British Museum: although Bland has not directly provided conceptual schematics or feedback on CHRR, his publications have informed the modeling processes of NUDS-Hoard.

Bransbourg, Gilles – Faculty member at the Institute for the Study of the Ancient World, New York University: OCRE project leader and sets design specifications and feedback on features.

Heath, Sebastian – Faculty member at the Institute for the Study of the Ancient World, New York University: the primary technical driving force behind Nomisma.org, created in partnership with Andrew Meadows of the American Numismatic Society.

Lockyear, Kris – Faculty, UCL Institution of Archaeology: primary intellectual contributor to CHRR; has had much influence on the design of CHRR, and the NUDS-Hoard data model was designed to fully accommodate his database.

Meadows, Andrew – Deputy Director of the American Numismatic Society: Nomisma.org Steering Committee and has had a significant role in setting specifications for OCRE and CHRR, as well as influenced the conceptual development of NUDS and NUDS-Hoard by means of his high-level knowledge of numismatic description.

Moorhead, Sam - Finds Advisor for Iron Age and Roman coins at the Portable Antiquities Scheme (PAS), British Museum: his database schema guidelines were instrumental in shaping the development of the NUDS-Hoard XML model.

Pett, Dan - ICT Adviser, Department of Portable Antiquities and Treasure, British Museum: developer of PAS, technical user of Nomisma, and member of Nomisma steering committee with whom I have collaborated on writing a paper for CAA 2012.

Sulosky-Weaver, Carrie – Responsible for the identification of many of the coins of the University of Virginia Art Museum Numismatic Collection (all those except the Mosberg Collection, the coins of Postumus, the bronze-diseased Greek civic coins, and the approximately 80 coins later found in a cigar box in the museum's vault) and creation of EAD files describing the coins.

Tolle, Karsten – Database and Information Systems, Goethe University Frankfurt am Main: developer of the Antike Fundmünzen in Europa (AFE), hosted by the Römisch-Germanische Kommission (RGK), one of the first technical users of Nomisma.org.

Witschonke, Rick – Curatoral Associate at the American Numismatic Society specializing in Roman Republican coins: initiated partnership between Lockyear and the ANS for the CHRR project; has tested and provided feedback during the development of the project.