The Dream and the Cross: a 3D scanning project to bring 3D content in a digital edition

Chiara Leoni, Marco Callieri, Matteo Dellepiane, ISTI-CNR, Pisa
Daniel Paul O'Donnell, University of Lethbridge
Roberto Rosselli Del Turco, University of Pisa
Roberto Scopigno, ISTI-CNR, Pisa

The Dream of the Rood is one of the earliest Christian poems in Old English and an example of the genre of dream poetry. While a complete text can be found in the 10th Century "Vercelli Book", the poem is considerably older, and its oldest occurrence is carved (in runes) on the 7-8th Century Ruthwell Stone Cross. In this paper, we present the work done in the framework of the "Visionary Cross" project, starting from the digitization of the Ruthwell Cross to the creation of a web-based digital edition of *The Dream of the Rood*, as it is carved on the Cross. The 3D data has been collected and processed with the explicit aim of creating a multimedia framework able to present the highly detailed digital model acquired with 3D Scanning technology, together with the transcription and translation of the runes that can be found on its surface. The textual and spatial information are linked through a system of bi-directional links called Spots, which allow the users to navigate freely over the multimedia content, keeping the 3D and textual data synchronized. The present work discusses the different issues that arose during the work, from digitization and processing to the design of a tool for the integration of three-dimensional content in the context of the presentation on the web platform of heterogeneous multimedia data. We end with the difficulties involved in the creation of an XML encoding that could account for the necessities of the visualization system but remain within the scholarly encoding standards of the relevant disciplinary community.

Categories and Subject Descriptors: I.4.1 [**Digitization and Image Capture**] Scanning; I.3.7 [**Three-Dimensional Graphics and Realism**]

Additional Key Words and Phrases: 3D scanning, digital edition, 3D and the web, multiresolution

1. INTRODUCTION

The Dream of the Rood is one of the first Christian poems in the corpus of Old English literature and one of the main examples of the genre of dream poetry in English [Swanton 1970]. It is also one of the earliest known Old English poems, due to a runic fragment carved on the late eighth-century *Ruthwell Standing Stone Cross* in Ruthwell Kirk, in Dumfries and Galloway, Scotland (see Figure 1). Due to its artistic qualities and overall relevance for Anglo-Saxon studies, the Dream is considered one of the most important achievements of Old English poetry. It is also one of the very few texts in this tradition that have been preserved in more than one witness [O'Donnell 1996]. In addition to the

Reference address: M. Callieri, ISTI-CNR, Via G. Moruzzi 1, 56124 Pisa, Italy; email: callieri@isti.cnr.it;

DOI:http://dx.doi.org/10.1145/0000000.0000000

fragment on the *Ruthwell Stone Cross*, an extended version of poem is found in the late tenth-century "Vercelli Book" (*Codex Vercellensis CXVII*, Archivio e Biblioteca Capitolare di Vercelli, Vercelli, Italy) and a small excerpt also appears on the eleventh-century reliquary, the *Brussels Cross*.

Fig. 1. The Ruthwell Cross: an incredibly rich artwork, with a very complex story to tell that links it to an Anglo-Saxon poetry milestone. The Cross presents several element types, like figurative scenes, runic and latin inscriptions, decorative elements.

The aim of the Visionary Cross project (http://www.visionarycross.org/) is to create a digital edition of a group of important objects related to *The Dream of the Rood* and to the Cult of the Cross in general in Anglo-Saxon England: in addition to the above mentioned artifacts containing the Dream poem, these objects include the Bewcastle Cross and the Old English poem *Elene* (also found in the Vercelli Book). The goal is to create a multimedia presentation that combines a "critical edition" of the poem (i.e. a specialized representation of the text intended for research purposes) with digitized images of the Vercelli Book and three-dimensional data related to the Crosses. Preserving and representing the strong connection between the poem and the objects on which it appears, from a spatial point of view as well as an artistic, is a primary desideratum.

In this paper, we present a project focused at presenting *The Dream of the Rood* and its close relationship with the Ruthwell Cross, through the design and implementation of a web-based integrated digital edition. The multimedia presentation system includes a highly detailed 3D model acquired with active 3D scanning technology, together with the transcription and translation of the runes that can be found on its surface, and a detailed description of the scenes carved of the cross. The textual and spatial information are linked through a system of bi-directional links called *Spots*, which allow the users to freely navigate any data source, supporting a *synchronized* browsing and analysis of the 3D model, the textual data and the digital editions of the poem texts.
Our approach provides a discussion and some solutions to several issues related to digital editions of poetic texts and multimedia presentations:

—*Tight integration of different media:* by adopting a web-based presentation platform, we are able to offer a full integration of 3D content with heterogeneous multimedia data. The web-based platform, based on WebGL and a system of XML configuration files, allows to easily encode and visualize interlinked modules, covering different types of data (i.e. text, 3D models, editions). This is the first time, to the best of our knowledge, that such a high level of integration of different media has been used in the presentation of medieval poetry.

—*Support to authoring and adaptation by non-expert users:* The framework designed proved to be flexible enough to be adapted to different types of media and to cope with the specific characteristics of this test case. One of the goals of our project was to develop a framework that could be used directly by non-experts during the authoring process and could be easily adapted to the presentation of similar poetic texts and objects.

—*Standards-compliant XML encoding:* our use of XML supports innovative functionalities for data organization and presentation; at the same time, however, we have followed encoding rules that are compatible with the standards of the Digital Humanities research community. This ensures interoperability and ease-of-use across disciplines, including within the humanities. The encoding of the runic text presented in this work is fully interoperable with the web-based visualization system and tools commonly used by Digital Humanities researchers.

—*Free and guided navigation:* the system also combines free navigation with some more structured access paths. We have included solutions to structure a ”guided visit” through a lecture that follows a number of pre-defined interesting points of the object (exploiting the classical slide-based approach).

—*Plugin-free operation:* finally, the multimedia presentation has been designed to run inside a standard web page, requiring only an up-to-date web browser (i.e. no plug-in installation required). This to allow easy access of interested users (from everywhere, with a standard software platform) and to reduce the ICT skills needed to use it.

2. RELATED WORK

In this section, we propose a short overview of the state of the art: the first subsection will be devoted to the problem of 3D content visualization on the web and integration with textual data; the second one will propose a short overview on the issue related to encoding of digital editions of ancient works.

2.1 Publishing 3D content on the Web, and its integration with textual descriptions

Text, images and videos have been ingredients of web pages from the very birth of the web. Three-dimensional content has been included only in a much later stage.
Initially, the visualization of 3D models in Web pages was devoted to embedded software components, such as Java applets or ActiveX controls [ACTIVEX]. There was no standardization, so that web-based rendering of 3D data was supported by a number of proprietary plug-ins, each one requiring

an explicit installation on the local computer. Providing standardization was the goal of the Virtual Reality Modeling Language (VRML) effort [Raggett 1995], started in 1995, and of the more recent X3D standard [Don Brutzmann 2007]. But 3D visualization was still delegated to external software components.
The introduction of the WebGL standard [Khronos Group 2009b], promoted by the Khronos Group [Khronos Group 2009], was a fundamental change. The WebGL API is a one-to-one mapping of the OpenGL|ES 2.0 specifications [Khronos Group 2009a] in JavaScript. This allows Web browsers to natively access the graphics hardware without needing additional plug-ins or extensions. A number of high-level libraries have been developed to help non-expert in using WebGL. They range from scene-graph- based interfaces, like Scene.js [Kay 2009] and GLGE [Brunt 2010], to more programmer-friendly paradigms, like like SpiderGL [Di Benedetto et al. 2010] and WebGLU [DeLillo 2009]. A declarative approach is implemented in X3DOM [Behr et al. 2009], whose use is simpler for people with expertise in web development, due to the similarity of the core principles.
Leaving aside integration issues, the presentation of complex datasets comprising images, videos and other media have been studied for a long time [Faraday and Sutcliffe 1997]. The lack of standards delayed a rich use of 3D content, but its integration with textual information has been thoroughly studied under all aspects, including not only navigation, but also the disposition and readability of the text [Sonnet et al. 2005; Jankowski et al. 2010], and the usability of interaction paradigms [Bowman et al. 2003; Polys et al. 2011]. All the above works are based on the idea of positioning the text in a spatial fashion, immersing the text tokens directly in the 3D environment. This approach does not fit applications cases in which different types of media are interconnected in a multi-directional fashion, as theorized by Halasz and Schwartz [Halasz and Schwartz 1994].
More recently, other approaches combined 3D models and textual information in an integrated fashion. The first of these [Götzelmann et al. 2007], presented a system to link textual information and 3D models where the 3D model is visualized from the best point of view to illustrate the corresponding portion of the text. Another recent work [Jankowski and Decker 2012] explored the possibility of the integration of text and 3D models on the web, by proposing two visualization modes: one is a typical hypertext, while the other one creates a virtual 3D environment, where text is attached to 3D hotspots. The Smithsonian X3D explorer (http://3d.si.edu) is an alternative example where 3D models are associated to additional content, but the structure and flexibility of the proposed system are not known.

2.2 Text encoding and digital editions

In recent years, a good number of digital editions of ancient and modern texts have been published, and more are under work. The first examples of such editions date to the second half of the 1990s. Many of these were based on simple HTML and/or Flash [McGillivray and Asgar-Deen 1999; Reed Kline 2001; Foys et al. 2009]. Others used SGML displayed in specialized and proprietary viewers [Robinson and Taylor 1998; Duggan ; Streuvels et al. 2003]. Slightly more rare is the use of SGML (later XML) for the master encoding, coupled with display texts in HTML (see [Prescott 1997; O'Donnell and Collins 2005] for early examples). All these early experiences adopted an optical support (CD/DVD-ROM) as distribution medium.

Current digital editions retain considerable diversities, since alternative approaches include WYSIWYG tools (such as the Classical Text Editor by Stefan Hagel http://cte.oeaw.ac.at/) or the use of LaTeX macros for typesetting purposes (e.g. adopting a customized form as in the Mauro-TeX project [Mascellani and Napoletani 2001] http://www.maurolico.unipi.it/mtex/mtexen.htm). Nevertheless, many digital editions have now settled on two standards: HTML and the World Wide Web for publication and dissemination, and the Text Encoding Initiative (TEI) XML schemas and related *Guidelines* [TEI Consortium 2007] for encoding of the edition text.

The TEI XML schemas are modular in nature, allowing for great flexibility. As a consequence, creating a digital edition requires only to pick and choose the modules needed for the task: besides the essential TEI modules (*core*, *tei*, *header* and *textstructure*) most users only need to add the *msdescription* (Manuscript Description), *transcr* (Transcription of Primary Sources) and *textcrit* (Critical Apparatus); if the edited text includes special or otherwise non-standard characters it is also possible to include the *gaiji* (Non-standard Characters and Glyphs) module. An *analysis* (Simple analytic mechanisms) module provides tags for linguistic, syntactical or other forms of analysis. A simple web interface (http://www.tei-c.org/Roma/) allows the creation of a new schema or to modify an existing one. It is also possible to select not only the XML elements really necessary for the text markup, but also to modify their characteristics and even create new elements and attributes using this interface.

Adding the transcr module allows to use the elements (<facsimile>, <surface>, <zone>) and attributes (@facs) necessary to create digital facsimiles: this is particularly useful in the case of diplomatic or semi-diplomatic editions, where it is possible to have the manuscript images and the corresponding edition text facing each other, with synchronized text-image linking down to the level of a single word or character. The TEI XML markup makes it possible to have different edition levels within the same XML file, which means that it is not necessary to have separate versions of the edition text (for instance, for diplomatic and interpretative levels), as it is very easy to extract these "views" of the edition from the XML document by means of XSLT style sheets. As it is often the case with complex XML markup, the editor may have to face one of the most grievous limitations of the XML language, that of conflicting hierarchies; fortunately there are several possible workarounds, and the TEI *Guidelines* suggest many useful solutions.

Fig. 2. The 3D scanning of the Ruthwell Cross (during the second day of scanning).

Thanks to their power, flexibility, comprehensiveness and excellent documentation (available in both on-line and off-line forms), the TEI schemas have been chosen by a large number of projects aiming at creating digital editions, encompassing texts of all historical periods and a variety of philological methods [Parker 2012; Emery and Toth 2009; Smith 2010].

The high flexibility and modularity of these XML schemas makes it possible not only to customize them for specific purposes, but also to expand them according to the most recent philological theories, possibly experimenting with a non-standard extension which, after a period of testing and examination, can be proposed and accepted within the existing collection of TEI modules. This was the case for the addition of facsimile to the transcription module, and is the case of the genetic criticism module (http://www.tei-c.org/SIG/Manuscripts/genetic.html), which has been accepted for inclusion and is now undergoing a testing period.

3. DIGITIZING THE RUTHWELL CROSS - GATHERING THE RAW 3D MODEL

A 3D digitization of the Cross was planned from the very early stages of the Visionary Cross project to produce digital data that could support the study, visualization and presentation of the artwork both as a whole and down to its fine carved details. For this reason, even though the object of interest is quite big (more than 5 meters in height), performing the acquisition using faster mid-long range scanners was not a valid option: the fine details would have been lost due to the poorer resolution and accuracy of those scanning devices (order of a few millimeters).

Therefore, we decided to use a laser triangulation scanner (Minolta Vivid 910) to ensure sufficient sampling density and accuracy in the acquisition of all the required data. The scanning campaign, which took place in April 2012, was planned carefully: the necessity to place the scanner at a short distance from the object (50-130 cm) forced us to use a scaffolding. Moreover, even though the cross is inside a church, its placement is not easily accessible: the cross resides in a pit behind the altar, and it almost reaches the top of the apse. The pit is around 1.5 meters deep, and has a metal rail all around it that, because it is of historical interest, could not be removed. Therefore, it was necessary to create a multi-layer scaffolding to work at three different heights. The Cross was acquired in three days of work, starting from the top layer and dismantling a portion of the scaffolding every day (see Figure 2).

The geometry of the cross was sampled with a total of 350 scans. The raw data were processed following the usual 3D scanning pipeline, using the MeshLab tool [Cignoni et al. 2008]. The final "master model", reconstructed from the sampled data with a cell resolution of 1 mm, is a heavy 112 million triangles model. We also acquired the surface color detail, to be able to produce a digital model encompassing both geometry and color. A photographic campaign of nearly 250 images was performed to acquire the apparent color. The photographic datasets were then projected on the model (using some recent MeshLab tools for color data processing [Ranzuglia et al. 2012]) in order to obtain a highly detailed representation of the surface (Figure 3 presents some snapshots of the final model).

The entire digitization activity took four days of work on site (shape and color sampling), while a few weeks were needed for post-processing the raw data, to produce several models (with different resolution and color encoding modalities) required for visualizing the Cross on different hardware platforms. Moreover, several higher-resolution models of portions of the Cross surface were also created, to account for more focused analysis.

The global model and the higher-resolution details were placed in the same reference system. This is required to switch from the global model (invaluable for the generic navigation) to the details (when a close-up inspection is needed). Since all the annotations, views and linked data will be referenced in this same space, this will also ensure interoperability between the different resolutions and 3D model subsets.

4. PREPARING THE DIGITAL EDITION

The Ruthwell Cross is an extremely interesting object of art and its relation with *The Dream of the Cross* goes beyond the presence of the carvings. There's a strong relation between the poem and the Cross, featuring at least three presentation levels:

—*Global* level, due to the shape of the object (the cross, intended as a monumental religious artifact).

—*Panel* level, where detailed textual descriptions of the single scenes depicted on the Cross are linked to the corresponding carved panels.

—*Rune* level, where the lines of the poem are linked to the object not only for their meaning, but also to display their spatial disposition on the stone and around the figured panels.

For this reason, the digital edition required not only the visualization of the poem's verses, but also to handle the other levels of presentation. Moreover, both the digital edition framework and the runes' encoding had to be structured in the most flexible and extensible fashion, so that the approach developed could later be applied to other objects that share similar features (e.g. artworks with carved inscriptions).

Defining the target audience is an essential step in the design and development of a presentation system. In our case, we wanted to build a presentation system oriented towards didactic purposes in the first instance (a research edition, aimed at the needs of humanities researchers, is planned for future work). The aim was to compile a sort of "interactive textbook", to be used by teachers and scholars when presenting this artifact. For this reason, we had the idea of presenting each level of exploration (the cross, the panels, the runes) in an independent way, but following a common presentation paradigm, to make the experience seamless. We chose to adopt the web platform to grant wide and easy access, and to provide a uniform access and visualization over different computing platforms. The web platform was chosen because of the wider and easier access, and to provide a uniform access and visualization of the data across different platforms.

5. VISUALIZING THE HIGH-RESOLUTION 3D MODEL

The first task in this project was to solve the issues related to the visualization on the web of a high-resolution 3D model. Displaying high-resolution models on a web browser is not just a matter of optimizing the rendering speed, but also involves considering the loading time and network traffic caused by transferring a considerable amount of data over the net. While WebGL gives direct access to the GPU resources, how data are transferred and managed is totally in the hands of the programmer. Loading a high-resolution model as a whole through the web requires to transfer data in the order of tens of megabytes: this is definitely unpractical, especially if the user has to wait the file transmission completion before seeing any visual result.

Multiresolution techniques provide solutions for both rendering and data transfer phases. Multiresolution schemes generally split the geometry in smaller chunks. For each chunk, multiple levels of details are available. Transmission is *on demand*, requiring only to load and render the portions of the model strictly needed for the generation of the current view (thus, reducing transmission times by using several heuristics: view-frustum culling, view-dependent resolution, visibility culling, etc.). While this approach is key to be able to render very large models at an interactive frame rate, it is also highly helpful with respect to the data transfer over a possibly slow network, since the data transferred will be divided in small chunks and only transferred when needed. The advantages of using this type of methods are the fast startup time and the reduced network load. The model is immediately available for the user to browse it, even though at a low resolution, and it is constantly improving its appearance as new data are progressively loaded. On the other hand, since refinement is driven by view-dependent criteria (observer position and distance from the 3D model sections), only the data really needed for the required navigation are transferred to the remote user. Finally, multiresolution allows also some (limited) degree of data protection. Most institutions do not want their 3D data be easily copied and reused without permission. When we use a multiresolution encoding, the entire 3D model is never transmitted to the remote user in a single file. In most cases, multiresolution encoding schemes adopts a fragmented encoding and a proprietary file encoding structure, making malicious copying of the 3D data more complicated and requiring the design of ad-hoc procedures for downloading the data in its entirety and for the recombination of the geometry description.

Multiresolution techniques provide solutions for both rendering and data transfer phases. Multiresolution schemes generally split the geometry in smaller chunks. For each chunk, multiple levels of details are available. Transmission is , requiring only to load and render the portions of the model strictly needed for the generation of the current view (thus, reducing transmission times by using several heuristics: view-frustum culling, view-dependent resolution, visibility culling, etc). While this approach is key to be able to render very large models at an interactive frame rate, it is also highly helpful with respect to the data transfer over a possibly slow network, since the data transferred will be divided in small chunks and only transferred when needed. The advantages of using this type of methods are the fast startup time and the reduced network load. The model is immediately available for the user to browse it, even though at a low resolution, and it is constantly improving its appearance as new data are progressively loaded. On the other hand, since refinement is driven by view-dependent criteria (observer position and distance from the 3D surface), only the data really needed for the required navigation are transferred to the remote user. Finally, multiresolution allows also some (limited) degree of data protection. Most institutions do not want their 3D data be easily copied and reused (maybe for commercial purposes). When we use a multiresolution encoding, the entire 3D model is never transmitted to the remote user in a single file. In most cases, multi resolution encoding schemes adopts a fragmented encoding and a proprietary file encoding structure; thus, malicious copy of the 3D data is more complicated and requires the design of ad-hoc procedures for downloading the data in its entirety and for the recombination of the geometry description.

We implemented one of those multiresolution schemes, Nexus (http://vcg.isti.cnr.it/nexus/), on top of the SpiderGL library [Di Benedetto et al. 2010], obtaining excellent performances. Nexus is a multiresolution visualization library belonging to the family of cluster based, view-dependent visualization algorithms (see for example the Adaptive Tetrapuzzles [Cignoni et al. 2004] and the Batched Multi Triangulation [Cignoni et al. 2005] algorithms). The Nexus scheme employs a patch-based approach: the granularity of the primitive elements is moved from triangles to small contiguous portions of a mesh (patches composed by a few thousand triangles), to reduce the number of per-element CPU operations. Moreover, a batched structure allows for aggressive GPU optimization of the triangle patches; the latter are usually encoded with triangle strips, boosting GPU rendering performances. To hold and manage the large number of alternative patches that compose the multiresolution encoding, Nexus adopts a spatial partitioning strategy based on KD-trees, which combines fast streaming construction of the multiresolution model with efficient adaptive spatial partitioning of the mesh.

The Nexus library has been ported on SpiderGL/WebGL; therefore, it now works directly on all recent browsers including Chrome and Firefox and Internet Explorer, without the need of additional components or plug-ins. The entire visualization system is based on HTML5, JavaScript and XML.

The Ruthwell Cross multiresolution model currently used in our presentation system (see Figure 3) is based on a base model of 15 Millions of triangles (obtained simplifying the original 122 Mtr master model). The color channel was reconstructed by integrating and blending digital photographs over the surface and mapping this information onto the triangle mesh by using a color-per-vertex encoding. Some tests are ongoing to exactly evaluate the network load, to define optimal parameters for tuning the data exchange, and to find the best compromise with regard to the model resolution (total number of triangles). While theoretically possible, we discarded the idea to use the master full-res model (112 Million triangles) for the on-line system, because such a very high resolution would be mostly not used, given the resolution of the visualization web page and the level of zoom normally used in 3D navigation and rendering. Moreover, since we have used a high-quality surface simplification tool for producing the 15 Mtr model from the master 112 Mtr model, the accuracy of the reduced model is still extremely detailed. This decision, while appropriate for use in a pedagogical context (and with a remote access), may require reconsideration for future work on a research edition. In such a context, however, alternate delivery methods may need to be considered.

6. FIRST LEVEL OF PRESENTATION - THE CROSS AS A WHOLE

The first objective was to present the cross as a whole, giving the user the possibility to freely navigate and to explore its shape and minutes geometric details.

Navigating a 3D model is not a simple task, especially if the target audience is not experienced with 3D graphics and virtual navigation. In this specific case, we tried to help users by taking into account the shape of the cross and the more common navigation paths. Our goal was to define some constrained interaction modes to offer users more handy instruments for supporting the virtual exploration of the artwork.

Given the vertical structure of the cross, a cylindrical trackball (where the user can rotate around the object, and pan on its height) seemed the most sensible choice at this level of navigation, because it is simple to use, can reach every part of the carved surface, and is easy to *animate*. The trackball can be controlled by using the mouse, keyboard, or on-screen buttons: both the trackball and the interface have also been designed to be compatible with a touch-based interface, such as might be found on a multimedia kiosks (i.e. for use for in situ cultural heritage interpretation) or on mobile devices (when WebGL will be fully available on that hardware).

Free exploration is certainly helpful for study and classroom presentation of the artifact; but saving interesting points of view for future reference could even be more useful. For this reason, we implemented a simple bookmarking mechanism. The user can store its current position in one of the available slot, and return to it later on (thanks to the animated trackball). Each bookmark is simply a trackball position and a name, these data are saved currently as cookies on the local machine and exportable as XML files for sharing; this is because, at the moment, the system is entirely client-side. However, the bookmarking mechanism could be easily integrated into a client-server architecture, for easier sharing among users. The other parameters of the 3D visualization GUI are the usual zooming factor, the light settings and the rendering modes. In order to improve the insight in the carved details, we have added to the browser the capability of changing interactively the light direction, thus supporting the simulation of the "raking light" often used by experts. The user can switch between two rendering modes: mesh rendered with mapped photographic color or by the pure geometry; this is provided because some details may be easier to perceive in one mode but not in the other (Figure 3).

7. SECOND LEVEL OF PRESENTATION - THE CARVED PANELS NARRATING THE GOSPELS

The second level of presentation was the exploration and explanation of the various panels showing biblical scenes on the North and South faces of the Cross. Our idea was to make it possible for the user to explore the textual description of these panels and have these descriptions linked to the relevant portion of the 3D model. While working on this specific problem, we realized that it was just a particular instance of a much wider need of presenting a 3D object strictly interconnected to some other media (in this case, a textual description, but also images, videos or other media). Our work was then steered towards the design of an extensible tool for this purpose. The presentation system was

Fig. 3. Images grabbed from the web-based multiresolution 3D model of the Cross. Top left: the patches used by the multiresolution rendering engine in this specific frame. Top middle: the geometry rendered with photographic color. Top right: the same geometry rendered using a simple Lambertian shading (no color). Bottom row: changing the light direction, to simulate raking light, helps to improve insight of the small carvings, especially the carved runes.

organized following a modular and extensible paradigm, where all the elements of the system could be easily reconfigured, arranged and assembled in a visualization page, and completely parametric with respect to a given input dataset.

8. AN INTEGRATED FRAMEWORK TO VISUALIZE 3D AND TEXT ON THE WEB

Our approach is to allow the user to explore each media or data source (in this case the 3D model, the descriptive text and the digital editions/transcriptions) independently ,but also to allow the user to jump from one data source to another by exploiting predefined connection points. Most multimedia presentation systems simply give priority to a single media, which becomes the predominant one in navigation, having the rest of the information "float aroud" it. Conversely, we wanted to present to the user the different media on a true peer level. A diagram presenting the architecture of our system is shown in Figure 4. Multiple *Viewer* modules handle different types of data; in the current system version, only 3D models and structured textual information are taken into account, but we are planning to add at least a Viewer for 2D images in the forthcoming extended version. Each viewer component is able to display its specific type of data, supporting user navigation and inspection.

When the user encounters one of the predefined *points of interest* while browsing a dataset (presence of points of interests is highlighted graphically to the user during navigation), he/she may select it. This selection action forces a synchronization of all active viewers, which will focus their visualization on that specific point of interest. The Viewer modules communicate with the *Manager* module through synchronization signals, which encode the selected *Spot* entity. The Spot is a semantic element, which

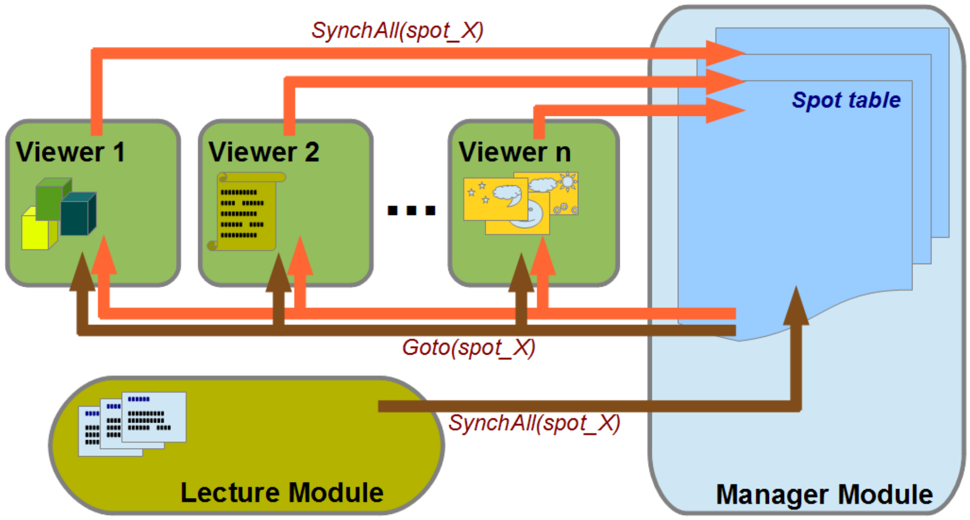

Fig. 4. The diagram presents the architecture of the presentation system: the Manager Module communicates with the Viewer modules using synchronization signals. The Manager module receives messages from the Viewers and broadcasts commands to focus on each activated Spot. The Lecture module is able to use the same signals to provide a path through pre-defined spots to the Manager module

defines a specific location in a dataset and is fully described in a data structure (Spot Table) stored by the Manager module. Each Spot connects different locations or data portions of the multimedia dataset presented by the viewer modules, e.g. in the current system it associates a specific viewpoint and region of the 3D model with a portion of the text presented by the Text Viewer module. The Spot concept is also used by the Lecture module, which allows to define predefined thematic exploration paths through the different media and Viewers, suggesting a linear navigation between spots that are not physically adjacent.

All the mentioned modules are configured using an XML structure. This drastic separation between the presentation code and the data makes the system extremely flexible and reusable. XML was chosen due to both its flexibility and the fact that potential content producers involved in the project (mainly Cultural Heritage operators who are assumed to have relatively low ICT expertise) were more familiar with this type of encoding, and in many cases will either already know how to access relevant tools for editing this XML or have access to workshops on the topic. The XML files are parsed using JavaScript, which also gives the possibility of easily handling the messaging procedure. The independence of the viewers and the simple message mechanism enable us to introduce new types of data by simply introducing a new viewer. At the same time, this also means that the system can be easily reconfigured to deal with different Cultural Heritage objects and information, just by changing the XML configuration files. The visualization web page starts almost empty: no data is present, beside the inclusion of JavaScript files and the ¡div¿ placeholders for the various elements. At startup, the page loads the XML configuration files specified and the DOM (Document Object Model) is populated with the elements. In this way, the appearance of the web page (i.e. the spatial arrangement of the elements in the HTML and the content) can be defined, as usual, with a standard CSS, making it much easier to integrate the system in an existing website. Please refer to [Callieri et al. 2013] for a more detailed description of the XML structure that we designed to implement the system.

All the mentioned modules are configured using an XML structure. This drastic separation between the presentation code and the data makes the system extremely flexible and reusable. XML was chosen for its flexibility and because potential content producers (mainly Cultural Heritage scholars and students) are assumed to be familiar with this type of encoding, and in many cases will either already know how to access relevant tools for editing this XML or have access to workshops on the topic.
The XML files are parsed using JavaScript, which also gives the possibility of easily handling the messaging procedure. The independence of the viewers and the simple message mechanism enable us to introduce new types of data by simply introducing a new viewer type.
At the same time, this also means that the system can be easily reconfigured to deal with different CH artworks and information, just by changing the XML configuration files. The visualization web page starts almost empty: no data is present, beside the inclusion of JavaScript files and the *<div>* placeholders for the various elements. At startup, the page loads the XML configuration files specified and the DOM (Document Object Model) is populated with the elements specified in the XML configuration files. In this way, the appearance of the web page (i.e. the spatial arrangement of the elements in HTML and the content) can be defined, as usual, with a standard CSS, making it much easier to integrate the system in an existing website. Please refer to [Callieri et al. 2013] for a more detailed description of the XML structure that we designed to implement the system.

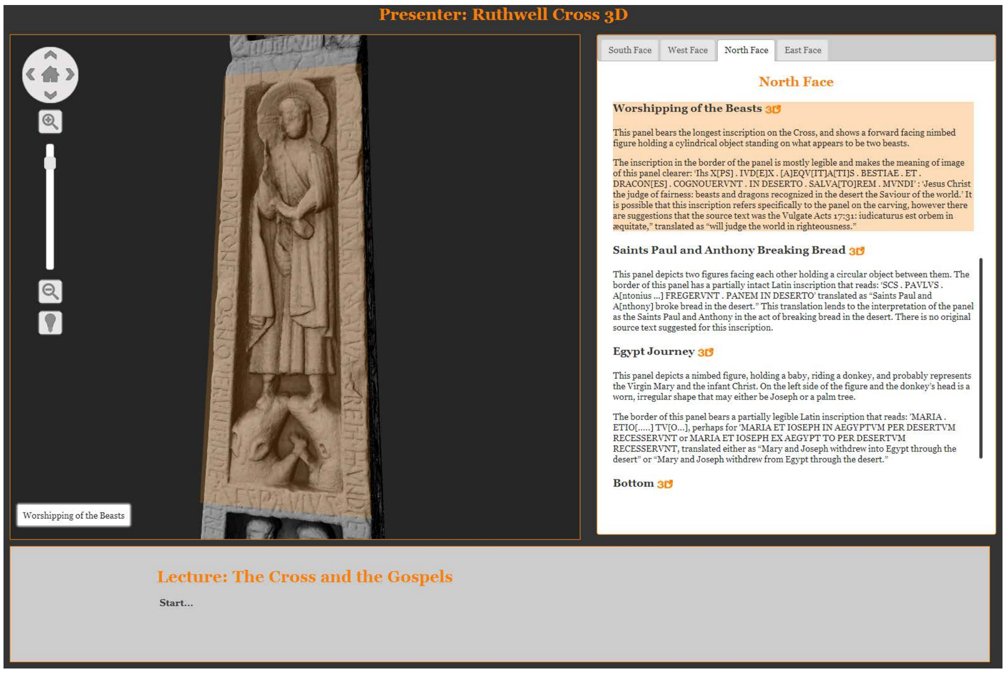

Fig. 5. A snapshot of the presentation framework: during the navigation, the user has detected and activated a Spot and consequently the link between the carved panel on the Ruthwell Cross and the corresponding portion of text describing it is visually enhanced.

The subdivision of the Ruthwell Cross in panels (each one containing a carved scene and some inscriptions) fits our paradigm perfectly, based on the definition of Spots that interlink the diverse representations and description media. Therefore, it was simple to configure the system to present the Cross at *panel* level: Figure 5 shows a snapshot of the system presenting one of the panels of the Cross together with the corresponding textual description. On the left side, the 3D viewer allows the user to navigate the 3D model (usig the same cylindrical trackball described in Subsection 6). The right side of the presentation window is dedicated to the Structured Text Viewer, which presents the descriptive text of the panels, divided in tabs (one for each side of the Ruthwell Cross).

During the navigation of the 3D model, every time the mouse passes over a region of the 3D model that has a Spot associated, the system shows the presence of an active areas by overlaying a colored transparent polygon. The presence of these colored regions warns the user of the presence of Spots that might disclose associated information in the Viewer modules corresponding to the other media. Once the spot is selected (mouse click), the Manager Module send communication to all active Vierwers: the 3D Viewer zooms in to better frame the selected area (according to the specific view defined in the Spot definition); the Structured Text Viewer moves to the tab defined in the Spot definition and scrolls the text to present the appropriate descriptive text to the user.
On the other side of the page, the Structured Text Viewer presents the descriptive text of the panels, divided in tabs (one for each side of the Ruthwell cross). The user may navigate the tabs (when a tab is selected, the Ruthwell cross rotates to show the corresponding side), browse the text and, by selecting the different links, force the 3D viewer to move the point of view to frame the appropriate carved panel.

Given our system architecture, it is easy to see how the *digital edition* might be simply considered as a new type of media: following this scheme, we have added a new Viewer, to display digital editions, and establish connections between the elements of the digital edition and the 3D model. The definition of the XML encoding at the *runic text* level proved to be more complex, due to the more fragmented nature of the content, and to the need to comply as much as possible with disciplinary standards for the encoding of text in digital editions. Section 9 will describe the implementation choices we made and our first results with the runic verses of *The Dream of the Rood* as it appears on the Ruthwell Cross

8.1 The Lecture module

The visualization scheme outlined in the previous sections, even considering the spot synchronization, can still be considered a free exploration. We stated at the beginning, however, that we wanted also to add a tool to support teachers in the organization of a lesson (and consequently in the setup of pre-defined navigations). While the synchronized exploration mimics a (non-linear) textbook, and may be used to present the material on the Ruthwell cross in an interactive session with students, it is also true that it lacks a specific presentation-oriented feature. We introduced the *Lecture* module specifically for this purpose. This module provides a way to explore the dataset which is orthogonal to the standard exploration described above; it is intended to be used to guide the user though the dataset by following a thematic path.

The idea is to build a predefined path, by using the same Spot concept. A “lecture” is a series of elements that resemble the classical slides used in presentations, with a title and a short text; each of these slides is connected with a specific Spot (consequently, at each context and spot change, the system will activate the Viewers to present the associated content). The user can go back and forth on this path; when a given slide is selected, the associated title and short text is presented in the lower part of the system window; at the same time, the Manager module is triggered to synchronize all the viewers to the associated Spot.

Analogously to the other modules, the Lecture module component is rendered in a <*div*> of the web page, and its content is stored in an XML file that is loaded and parsed at startup. Multiple lectures may be stored in the same XML file. A simple interface allows the user to select the lecture and to follow the slides. Each page in the visualization system (the 3D cross, the panels and the runes) may have a different set of lectures, related to the level of detail described in that page. Figure 6 shows an example of a lecture in use.

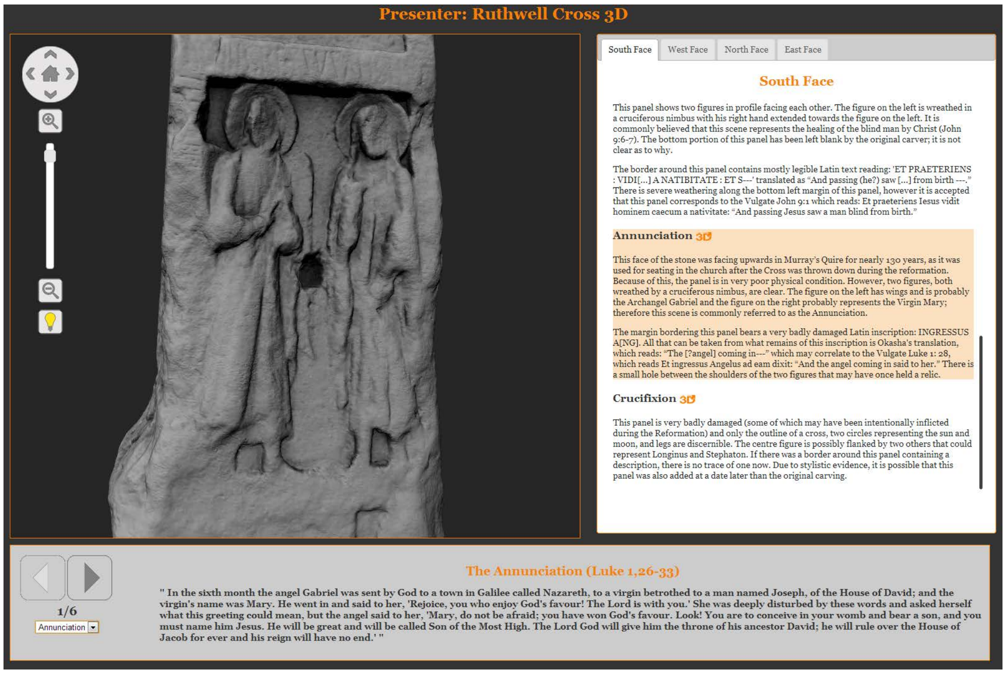

Fig. 6. A snapshot of the lecture model: a spot is associated with the current step of the lecture, allowing the user to follow a pre-defined path in the exploration of the system.

In the present status of development, a lecture has to be created manually by editing an XML file, in a manner similar to what required for the editing of the XML configuration files. Thanks to the fact that the configuration files are based on a very simple XML scheme, this is not a cumbersome task. Nevertheless, as a future development, we are planning to design an authoring tool to let users create their own lectures (and share them as XML files) without having to code in XML. Additionally, to make this tool more flexible, we will make possible to use not only the pre-existing spots but also user-defined locations, like the bookmarks described in Subsection 6.

9. THIRD LEVEL OF PRESENTATION - THE RUNIC TEXT OF *THE DREAM OF THE ROOD*

Using our presentation system it would have been easy to just take the Structured Text viewer and, with minimal changes, change it to display a digital edition of the text. However, this would have made both the encoding and the visualization completely unusable outside this specific system.

Our aim, on the other hand, was to have this part of the presentation system completely interoperable with the tools and workflow developed for the rest of the Visionary Cross project. For this reason, we decided not only to build a specific viewer for the digital editions, able to correctly display the complex formatting of the different editions (which is much more complex than the formatting used in the Structured Text Viewer described in the previous sections), but also to make the XML dataset used at this level completely interoperable with the standard tools used in the Digital Humanities and Digital Textual Studies. The first issue to take into account, then, was the encoding of the text.

9.1 XML encoding of the text using the TEI standard and XLST

Encoding a text in XML, easy as it may seem, is not a straightforward task. It is necessary to follow existing standards and conventions in order to obtain an encoding that is meaningful and complete, but also reusable and comprehensible to an external user (e.g. an unnegotiated exchange). We adopted the guidelines defined by the Text Encoding Initiative consortium (TEI) [TEI Consortium 2007], ensuring the encoding to be fully interoperable with the other digital tools used in the Visionary Cross project. To this purpose, in addition to the basic set of modules prescribed by TEI, we selected those modules that addressed the peculiarities of the poem and the researchers' interests: *transcr*, used to represent primary sources such as manuscripts and epigraphy; *gaiji*, used to encode non-standard characters, such as runes; and *linking*, used to aggregate and add structure to the entities.

To transform the XML specifications into HTML, instead of developing homebrew JavaScript code, we used XSLT (eXtensible Stylesheet Language Transformation), another widely used standard in the field of digital philology and digital editions. XSLT, a standard upheld by W3C, is a declarative language used to transform an XML document in another type of document. It is interesting that the XSL document that defines the transformation is itself an XML document, where a series of substitution/recursive/branching templates are encoded. A processor, which may be integrated into a web browser or as a stand-alone tool, reads the XSL templates, and applies them to the XML document to be transformed. Note that the XSL Transformation was used not only in the on-line system, but also in the other activities of the project.

An important point for the integration of this new media element in the system is to preserve its dynamic nature. This can be achieved, since XSL Transformation may be invoked dynamically, making it possible to follow the same strategy used for the structured text viewer: loading the XML at the page display time, transforming the XML into HTML (this time, by using XSL) and injecting the HTML into the appropriate empty element in the page.

9.2 Encoding the *Type-Facsimile*

The encoding and presentation of historical texts is a research topic in its own right. Most importantly, there is no single "correct" presentation: the precise form (and even content) depends on the purpose to which the text is being transcribed or presented, the nature and inter-relationships of the media on which they are found, and the extent to which the goal of the transcription is to describe or interpret the meaning of meta- and para-textual cues in the source documents. For this reason, when presenting critical editions of historical texts, it is often necessary to provide different versions or, more precisely, different editions of the text.

In our project, we wanted to begin by focusing on three different editions of *The Dream of the Rood* poem:

—The *Type-Facsimile* edition, a transcription focused on the shape and characteristics of the physical appearance of the runes as they are found currently on the cross with minimal interpretation. In this view, the rhetorical structure of the poem is not represented in the encoding.

—A *Diplomatic* edition presenting a more structured (and interpreted) representation of what the shape and characteristics of the poem (as it appears on the cross) indicate about its composition, carving, and meaning. This is a compromise format, which records aspects of the physical appearance of the text, while formatting the text according to its rhetorical structure;

—A *Critical* edition focused on the rune's literary and rhetorical interpretation. In this version, the poem is formatted according to its rhetorical structure, the runes are transliterated, and reconstructions (based on philological research and comparisons with other source-texts) are allowed. Information about the physical appearance of individual runes is ignored.

Our work started from the encoding and visualization of the *Type-Facsimile* edition; this choice was made because the *Type-Facsimile* edition offered three unique challenges:

—This edition should preserve the runic characters; this had to be considered in the XML encoding, in the parsing, and in the rendering of the web page.

—The *Type-Facsimile* has to mimic as closely as possible the actual spatial arrangement of the runes on the cross (an inverted U), which is more complex to manage with respect to a formatting more akin to poetry (which is easier to encode and display).

—This edition is the closest to the 3D shape of the cross; thus, this made it the ideal candidate to experiment with the concept of a digital edition strongly interlinked with the 3D representation.

```
29    <charDecl>
30       <char xml:id="feoh">
31          <charName>RUNIC LETTER FEHU FEOH FE F</charName>
32          <mapping type="transliterated">f</mapping>
33          <mapping type="runic">&#5792;</mapping>
34       </char>
35       <char xml:id="ur">
36          <charName>RUNIC LETTER URUZ UR U</charName>
37          <mapping type="transliterated">u</mapping>
38          <mapping type="runic">&#5794;</mapping>
39       </char>
40       <char xml:id="thurs">
41          <charName>RUNIC LETTER THURISAZ THURS THORN</charName>
42          <mapping type="transliterated">&#254;</mapping>
43          <mapping type="runic">&#5798;</mapping>
44       </char>
[.....]
212         <ab type="strip_left" xml:id="RC_ES_ab2">
213            <lb n="1"/><w xml:id="RC_W_016"><g xml:id="RC_G_069" ref="#ac"/><g xml:id="RC_G_070" ref="#haegl"/>
214            <lb n="2"/><g xml:id="RC_G_071" ref="#os"/><g xml:id="RC_G_072" ref="#feoh"/></w>
215            <lb n="3"/><w xml:id="RC_W_017"><g xml:id="RC_G_073" ref="#is"/><g xml:id="RC_G_074" ref="#cen"/></w>
216            <w xml:id="RC_W_018"><g xml:id="RC_G_075" ref="#rad"/><lb n="4"/><g xml:id="RC_G_076" ref="#is"/>
217            <g xml:id="RC_G_077" ref="#is"/><g xml:id="RC_G_078" ref="#cen"/><g xml:id="RC_G_079" ref="#nyd"/>
218            <lb n="5"/><g xml:id="RC_G_080" ref="#aesc"/></w>
[.....]
340         <join xml:id="RC_L_039" target="#RC_G_101 #RC_G_102 #RC_G_103"/>
341         <join xml:id="RC_L_040" target="#RC_G_104 #RC_G_105 #RC_G_106 #RC_G_107"/>
342         <join xml:id="RC_L_041" target="#RC_G_108 #RC_G_109 #RC_G_110 #RC_G_111"/>
```

Fig. 7.
Snippets of the XML encoding for the Type-Facsimile edition. The first chunk (lines 29-44) encodes the character "dictionary", where each rune type is mapped to its modern transcription and Unicode char code. The second chunk (lines 212-218) shows a portion of the encoding of one of the vertical strips of runes; all the elements, runes and words, have their own unique ID in order to be individually addressable (e.g. RC_G_71 is the $71st$ rune, RC_W_018 is the $18th$ word). The third chunk (lines 340-342) shows a different grouping of runes; in this case, each line of the carving has been expressed as a <join>. Each of these IDREFs is individually indexable.

The encoding copes with the mentioned peculiarities by using XML tags and properties (see Figure 7).
The inverted U arrangement was achieved by dividing the runes among three containers (upper horizontal line, plus two vertical arrangement of lines).
The runes were managed using a sort of *dictionary*: the actual characters were encoded as *glyph* elements, each one with a unique ID (to be able to individually address each rune) and a type, which specifies the kind of rune. This dictionary stores, for each kind of rune, the Unicode character used to display it and its transliteration in Latin characters. Finally, the different indexed entities (runes,

words, lines) were grouped using basic XML tags, or more complex elements such as <join> elements (for an earlier use of this technique, using SGML entities, see [O'Donnell and Collins 2005]). Defining the XSL Transformation used to transcode the described XML was not easy, for the same reasons mentioned earlier (the use of runes, and their spatial arrangement). Each container was transformed in a <*div*> element, shaped and formatted using CSS.

A level of indirection (from the text to the dictionary) was required to transform the runes. Thanks to this double encoding, it was possible to print the runes using both the appropriate runic symbol and its transliteration in Latin characters, obtaining a side-by-side visualization of the *Type-Facsimile*. Besides the required formatting, it was also necessary to add the HTML elements that are used by the system to manage spots (selection and viewer focusing)

As a minor issue, all existing browsers do support XSLT 1.0 but not the newer version 2.0. While it is possible to do the transcoding in both formats, XSLT 2.0 is more stable and it is easier to use in order to manage the more complex XML structures used in our encoding (the same transformation, in XSLT 1.0, would require tricks and a substantial use of recursion). This resulted in having to provide two different versions of the transformation: a more elegant and optimized XSLT 2.0 version to be used in the off-line tools, and an XSLT 1.0 version to be used in the web-based presentation system.

9.3 Multiple levels of linking

In the first part of the system (see Section 8), the 3D model and the textual data were connected by using, as connecting spots, the carved panels on the cross surface. This choice is definitely sound, since this data granularity was easy to understand but, at the same time, complete enough to properly describe the object.

However, when working with the digital editions, we realized very soon how the definition of the level of linking in the dataset was an issue. Since the *Type-Facsimile* edition is in runic, should we allow linking each single rune? Each word? Each line (as it appears on the cross)?

We had to face two problems:

(A) The granularity of the linkage is much smaller.

(B) There may be more than one level of linking for each edition.

While the point A) only requires more spots to be created, the point B) needs some additional work to make possible to change the level of linking on the fly, according to the choices of the user.

In the case of the *Type-Facsimile* edition, we opted to maintain the linking at two different levels:

—**Word**: each word of the poem is a synchronization spot; see Figure 8, top.

—**Line**: each line of the carving is a synchronization spot; see Figure 8, bottom.

We chose these these two levels because one was much closer to the physical shape of the carving, while the other was closer to the semantic structure of the poem. The user may change the level of linking at any time (using a button), depending on his interest. While a linking at the level of each single rune is entirely possible with this XML encoding (each rune has a unique ID), it seemed uninteresting to reach this level of granularity in the viewer, given the intended (primarily pedagogical) purpose and audience.

9.4 Other digital editions

As discussed above, we started from the *Type-Facsimile* edition because of its availability and its peculiarities (the runic alphabet, the shape, and its close relationship with the 3D shape). The work on the other digital editions, the *Diplomatic* and *Critical* ones, is still ongoing.

Thanks to the work done in the *Type-Facsimile*, in terms of XML encoding and XSL Transformation,

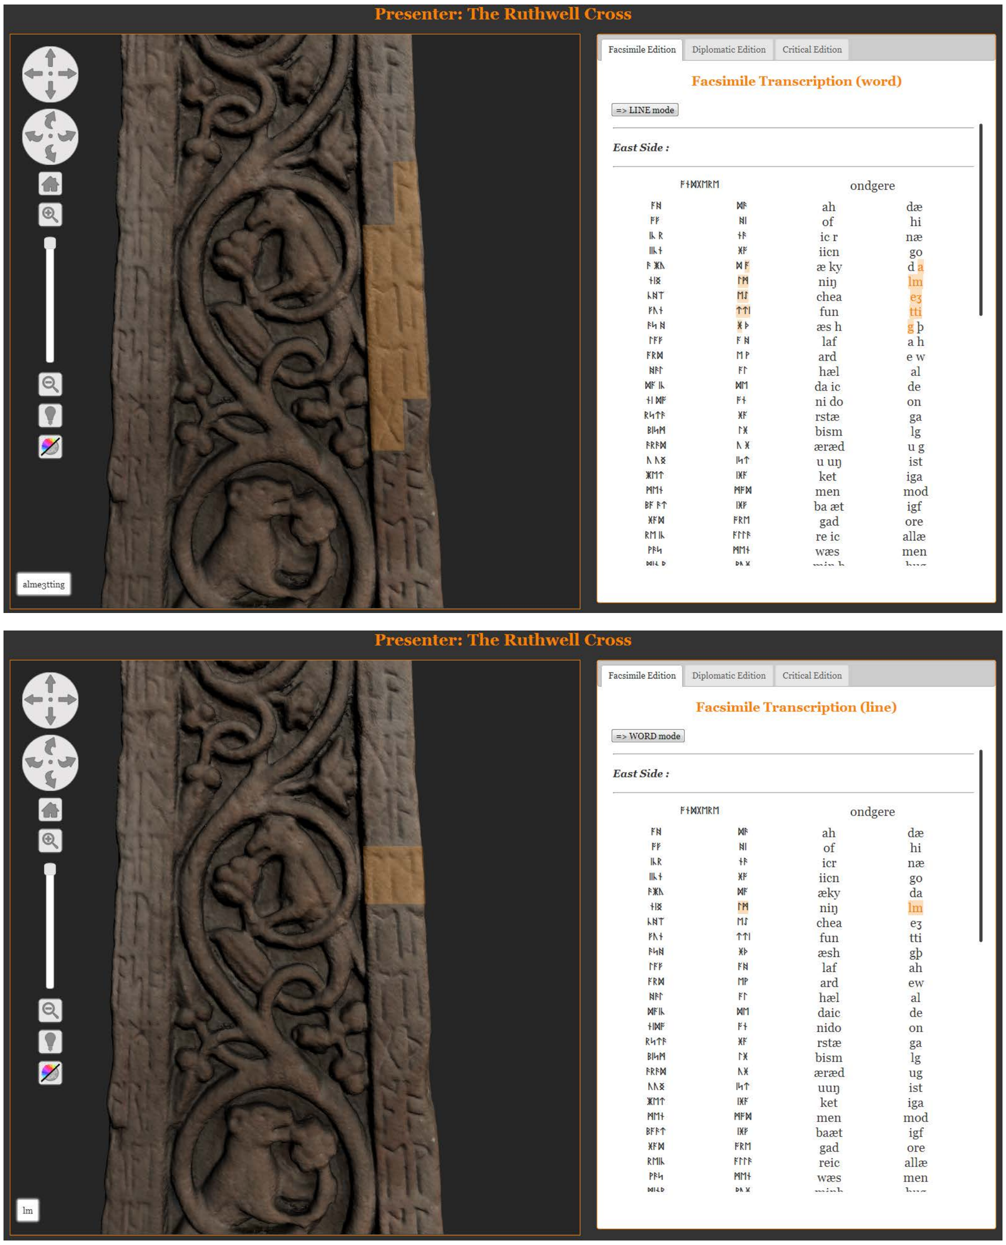

Fig. 8. The *Facsimile* edition with Spots linking at word and line level. Top, linking at word level. Bottom: linking at line level.

the addition of these two new editions to the system is easier. The kind of formatting required, closer to the one used in poetry, is easy to obtain in HTML and displaying the properties of the characters (missing, damaged, supplied) is also feasible. It is clear that multiple levels of linking will also be required in the other two editions.

10. RESULTS: AN INTEGRATED DIGITAL EDITION OF THE DREAM OF THE CROSS

The prototype of the integrated edition is built entirely using the described presentation system. As discussed in previous sections, the idea was to provide a tool for didactic purposes, to be used as an active textbook. In this sense, we wanted to mimic the way books divide different levels of information in different chapters. The three different level of significance of the cross (the cross, the panels and the poem) are then arranged in three different *chapters*, each one implemented as a single web page built using our presentation tool. When moving from one page to another, the user moves through the level of presentation:

—**The Cross**: in this web page, the cross is presented as a whole, using only the 3D viewer. The Lecture component is used to provide a guided walkthrough of the cross, presenting the signs left on the cross by its conservation history.

—**The Panels**: as described in Section 8, by using the 3D viewer and the Structured Text viewer, this page presents the different panels and their depicted stories (see Figure 9, top). The Lecture component provides thematic walkthroughs (stories from the gospels, mythological figures).

—**The Poem**: in this page, the poem is the main focus. The 3D viewer is sided by the Edition viewer, where the different editions are presented (see Figure 9, bottom).

A video showing the navigation features and the interactions is available at http://www.youtube.com/watch?v=Wov-2ik4ibY

11. CONCLUSIONS

In this paper we presented the results of a project aimed at designing the tools for an innovative visual presentation of a complex, unmovable object (the Ruthwell Cross).

The aim of the project was to design technology able to deploy interlinked and interactive visualization of an high-resolution 3D model (obtained by active 3D scanning) with a complex story-telling case study, by means of new 3D web technology. The goal has been therefore not just to provide a web-based visualization of a digital clone, but also to create a platform for sharing the knowledge on this complex artwork, including an innovative exploration of *The Dream of the Rood* poem by linking it to the Ruthwell Cross, showing their deep interconnection.

In order to be able to represent the different levels of connection (global, panels, runes) between the different elements, we created an integrated framework to visualize 3D and textual content. The entire system is based on HTML5, JavaScript and XML, works on most modern browsers, and does not require additional components or plug-ins to work. Moreover, the configuration of the web page is obtained using a set of XML files, which are easy to modify even by non-expert users. The system proved to be easy to use and flexible enough to encode the different levels of presentation. An extremely important point is that the tool is extensible and configurable; thanks to the modular structure and the complete separation between visualization code and data, it allows to add new types of media, and to reconfigure it to display a different set of artifacts and media.

Additional effort was put in the encoding of the editions of the poem verses that were carved on the Cross. The XML encoding was implemented following the standards defined by the TEI. Hence, it

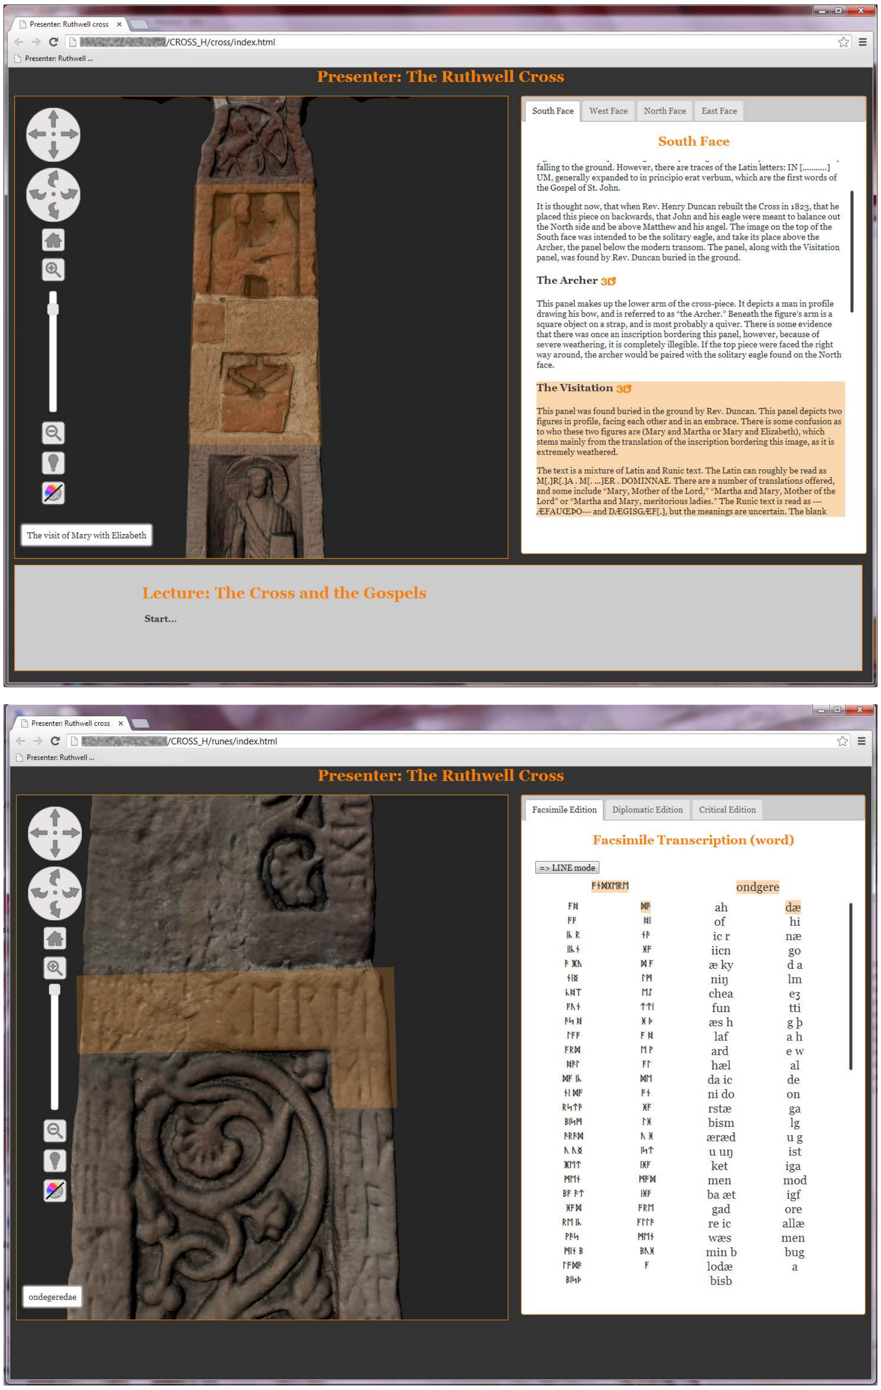

Fig. 9. Screenshots from the Integrated Digital Edition: on top, the cross explored at the level of the carved panels; on bottom, the digital editions of the poem, connected with the geometry.

is fully interoperable with both the web-based visualization system and the tools commonly used in the field of digital philology. Moreover, different levels of linking for the runes (lines or words) were implemented to enrich the modalities of navigation.

As far it concerns future work on the project, we have plans concerning extensions of content, usability assessment and technological development/extensions. In terms of content, our goal is to encode all the necessary editions of the poem (i.e. *Diplomatic* and *Critical*). Other possible extensions of the current system include: a new mechanism to create content (designing an authoring tool that should simplify the setup of a similar system describing a different artwork), and extending the presentation Modules, adding the possibility to visualize and integrate other types of media, like images and videos, that could add other type of visual content to be interconnected to the other presentation media.

Acknowledgment

The different authors of this work have been funded by the following projects: EU Community FP7 Project 3D-ICONS (FP7–ICT–PSP–297194); EU FP7 Network of Excellence V-MUST.NET (Grant Agreement 270404); Social Sciences and Humanities Research Council of Canada (SSHRC) Standard Research Grant and as part of the Digital Humanities Subproject of the GRAND National Centre of Excellence. The authors wish to thank the team of the Visionary Cross Project (http://www.visionarycross.org/), who contributed considerably to the functional specification phase of this system.

REFERENCES

ACTIVEX. Microsoft ActiveX Controls. http://msdn.microsoft.com/en-us/library/aa751968(VS.85).aspx. (????).

Johannes Behr, Peter Eschler, Yvonne Jung, and Michael Zöllner. 2009. X3DOM: a DOM-based HTML5/X3D integration model. In *Proceedings of the 14th International Conference on 3D Web Technology (Web3D '09)*. ACM, New York, NY, USA, 127–135. DOI:http://dx.doi.org/10.1145/1559764.1559784

Doug A. Bowman, Chris North, Jian Chen, Nicholas F. Polys, Pardha S. Pyla, and Umur Yilmaz. 2003. Information-rich virtual environments: theory, tools, and research agenda. In *Proceedings of the ACM symposium on Virtual reality software and technology (VRST '03)*. ACM, New York, NY, USA, 81–90. DOI:http://dx.doi.org/10.1145/1008653.1008669

Paul Brunt. 2010. GLGE: WebGL for the lazy. http://www.glge.org/. (2010).

Marco Callieri, Chiara Leoni, Matteo Dellepiane, and Roberto Scopigno. 2013. Artworks narrating a story: a modular framework for the integrated presentation of three-dimensional and textual contents. In *ACM WEB3D - 18th International Conference on 3D Web Technology*. ACM, ACM, 167–175. http://vcg.isti.cnr.it/Publications/2013/CLDS13

P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and Guido Ranzuglia. 2008. MeshLab: an Open-Source Mesh Processing Tool. In *Sixth Eurographics Italian Chapter Conference*. Eurographics, 129–136.

P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponchio, and R. Scopigno. 2004. Adaptive TetraPuzzles: Efficient Out-of-Core Construction and Visualization of Gigantic Multiresolution Polygonal Models. *ACM Trans. on Graphics (SIGGRAPH 2004)* 23, 3 (2004), 796–803.

Paolo Cignoni, Fabio Ganovelli, Enrico Gobbetti, Fabio Marton, Federico Ponchio, and Roberto Scopigno. 2005. Batched Multi Triangulation. In *Proceedings IEEE Visualization*. IEEE Computer Society Press, Conference held in Minneapolis, MI, USA, 207–214. http://vcg.isti.cnr.it/Publications/2005/CGGMPS05

Benjamin DeLillo. 2009. WebGLU: A utility library for working with WebGL . http://webglu.sourceforge.org/. (2009).

Marco Di Benedetto, Federico Ponchio, Fabio Ganovelli, and Roberto Scopigno. 2010. SpiderGL: a JavaScript 3D graphics library for next-generation WWW. In *Proceedings of the 15th International Conference on Web 3D Technology (Web3D '10)*. ACM, New York, NY, USA, 165–174. DOI:http://dx.doi.org/10.1145/1836049.1836075

Leonard Daly Don Brutzmann. 2007. *X3D: Extensible 3D Graphics for Web Authors*. Morgan Kaufmann.

H.N. Duggan. *The Piers Plowman Electronic Archive*. University of Virginia Institute for Advanced Technology in the Humanities. http://books.google.it/books?id=BYJnnQEACAAJ

Doug Emery and Michael B. Toth. 2009. Integrating Images and Text with Common Data and Metadata Standards in the Archimedes Palimpsest. *Digital Humanities Abstracts 2009* (June 2009), 281–283.

Peter Faraday and Alistair Sutcliffe. 1997. Designing effective multimedia presentations. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems (CHI '97)*. ACM, New York, NY, USA, 272–278. DOI:http://dx.doi.org/10.1145/258549.258753

M.K. Foys, K.E. Overbey, and D. Terkla. 2009. *The Bayeux Tapestry: New Interpretations*. Boydell Press.

Timo Götzelmann, Pere-Pau Vázquez, Knut Hartmann, Andreas Nürnberger, and Thomas Strothotte. 2007. Correlating Text and Images: Concept and Evaluation. In *Proceedings of the 8th international symposium on Smart Graphics (SG '07)*. Springer-Verlag, Berlin, Heidelberg, 97–109. DOI:http://dx.doi.org/10.1007/978-3-540-73214-3_9

Frank Halasz and Mayer Schwartz. 1994. The Dexter hypertext reference model. *Commun. ACM* 37, 2 (Feb. 1994), 30–39. DOI:http://dx.doi.org/10.1145/175235.175237

Jacek Jankowski and Stefan Decker. 2012. A dual-mode user interface for accessing 3D content on the world wide web. In *Proceedings of the 21st international conference on World Wide Web (WWW '12)*. ACM, New York, NY, USA, 1047–1056. DOI:http://dx.doi.org/10.1145/2187836.2187977

Jacek Jankowski, Krystian Samp, Izabela Irzynska, Marek Jozwowicz, and Stefan Decker. 2010. Integrating Text with Video and 3D Graphics: The Effects of Text Drawing Styles on Text Readability. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1321–1330. DOI:http://dx.doi.org/10.1145/1753326.1753524

Lindsay Kay. 2009. SceneJS. http://www.scenejs.com. (2009).

Khronos Group. 2009. Khronos: Open Standards for Media Authoring and Acceleration. (2009).

Khronos Group. 2009a. OpenGL ES - The Standard for Embedded Accelerated 3D Graphics. (2009).

Khronos Group. 2009b. WebGL - OpenGL ES 2.0 for the Web. (2009).

Paolo Mascellani and Pier Daniele Napoletani. 2001. MauroTeX - A Language for Electronic Critical Editions.. In *ICHIM (2)* (2002-01-03). 223–241. http://dblp.uni-trier.de/db/conf/ichim/ichim2001-2.html#MascellaniN01

M. McGillivray and T Asgar-Deen. 1999. *Geoffrey Chaucer's Book of the Duchess a hypertext edition*. University of Calgary Press.

D.P. O'Donnell. 1996. *Manuscript Variation in Multiple Recension Old English Poetic Texts*. Yale University.

D.P. O'Donnell and D. Collins. 2005. *Cædmon's Hymn: A Multimedia Study, Archive and Edition*. D.S. Brewer. http://books.google.it/books?id=8iEAcwwTMoIC

D.C. Parker. 2012. A Transcription of Codex Sinaiticus. (January 2012). http://epapers.bham.ac.uk/1690/

Nicholas F. Polys, Doug A. Bowman, and Chris North. 2011. The role of Depth and Gestalt cues in information-rich virtual environments. *Int. J. Hum.-Comput. Stud.* 69, 1-2 (Jan. 2011), 30–51. DOI:http://dx.doi.org/10.1016/j.ijhcs.2010.05.007

A. Prescott. 1997. The electronic Beowulf and digital restoration. *Literary and Linguistic Computing* 197 (1997), 185–195.

D. Raggett. 1995. Extending WWW to support Platform Independent Virtual Reality. *Technical Report* (1995).

Guido Ranzuglia, Marco Callieri, Matteo Dellepiane, Paolo Cignoni, and Roberto Scopigno. 2012. MeshLab as a complete tool for the integration of photos and color with high resolution 3D geometry data. In *CAA 2012 Conference Proceeding*. Pallas Publications - Amsterdam University Press (AUP). http://vcg.isti.cnr.it/Publications/2012/RCDCS12

Naomi Reed Kline. 2001. *A Wheel of Memory: The Hereford Mappamundi*. University of Michigan Press.

Peter Robinson and Kevin Taylor. 1998. Publishing an Electronic Textual Edition: The Case of The Wife of Bath's Prologue on CD-ROM. *Computers and the Humanities* 32, 4 (1998), 271–284. DOI:http://dx.doi.org/10.1023/A:1000943530396

Neel Smith. 2010. Digital Infrastructure and the Homer Multitext Project. In *Digital Research in the Study of Classical Antiquity*, Gabriel Bodard and Simon Mahony (Eds.). Ashgate Publishing, Burlington, VT, 121–137.

Henry Sonnet, Sheelagh Carpendale, and Thomas Strothotte. 2005. Integration of 3d data and text: the effects of text positioning, connectivity, and visual hints on comprehension. In *Proceedings of the 2005 IFIP TC13 international conference on Human-Computer Interaction (INTERACT'05)*. Springer-Verlag, Berlin, Heidelberg, 615–628. DOI:http://dx.doi.org/10.1007/11555261_50

S. Streuvels, M. de Smedt, and E. Vanhoutte. 2003. *De Teleurgang van den Waterhoek*. Amsterdam University Press. http://books.google.it/books?id=R77WLwEACAAJ

Michael Swanton (Ed.). 1970. *The Dream of the Rood (Old and Middle English texts)*. Manchester University Press - Barnes & Noble, Manchester - New York. http://books.google.it/books?id=hxsNAQAAIAAJ

TEI Consortium (Ed.). 2007. *TEI P5: Guidelines for Electronic Text Encoding and Interchange*. TEI Consortium, Charlottesville, VA, USA. http://www.tei-c.org/Guidelines/P5/

Received February 2009; revised July 2009; accepted October 2009