

iClay: Digitizing Cuneiform

Jonathan Cohen, Donald Duncan, Dean Snyder, Jerrold Cooper, Subodh Kumar,
Daniel Hahn, Yuan Chen, Budirijanto Purnomo, and John Graettinger

Johns Hopkins University, Baltimore, Maryland, USA
Applied Physics Laboratory
Department of Computer Science
Department of Near Eastern Studies

Abstract

Advances in digital technology for the graphic and textual representation of manuscripts have not, until recently, been applied to the world's oldest manuscripts, cuneiform tablets. This is due in large part both to the three-dimensional nature of cuneiform tablets and to the complexity of the cuneiform script system. The Digital Hammurabi Project and the Initiative for Cuneiform Encoding announce success in encoding Sumerian-Akkadian cuneiform in Unicode while also demonstrating advances in 3D scanning and visualization of cuneiform tablets, showcased by iClay, a cross-platform, Internet-deployable, Java applet that allows for the viewing and manipulation of 2D+ images of cuneiform tablets.

Categories and Subject Descriptors: H.3.0 [Information Storage and Retrieval]: Character Encoding; H.3.7 [Digital Libraries]: Standards; I.3.3 [Picture/Image Generation]: Digitizing and scanning; I.4.1 [Digitization and Image Capture]: Scanning; I.7.2 [Document Preparation]: Standards

1. Introduction

Modern digital technologies greatly enhance the access to, reproduction of, and publication of cultural heritage materials, particularly manuscripts. Parchment, papyrus, and paper manuscripts lend themselves to effective two-dimensional visualization; and the scripts in which most are written have standard computer encodings in place. Neither, however, is the case for the world's oldest manuscripts, cuneiform tablets.

Cuneiform tablets exhibit several levels of three-dimensionality. (Figure 1) The media are three-dimensional - rounded clay tablets; the characters are three-dimensional - wedge impressions in damp clay; and the writing streams themselves are three-dimensional - sentences often running over the fronts and onto the sides and backs of the tablets. Two-dimensional reproduction simply falls short of capturing the three-dimensional reality of cuneiform tablets. As



Figure 1: 3D scan of tablet T24. Johns Hopkins Archaeological Collection

a result, most scholars to this very day still publish cuneiform tablets in hand-drawn copies. (Figure 2)

Another issue impacting cuneiform research and publication, is the fact that, unlike for most writing systems in the world, there is still no standard computer encoding for cuneiform text. The general practice among cuneiformists of working almost exclusively in Latin alphabetic transliteration, although useful for its intended purposes, is, in fact, a hindrance to the pervasive use of computers in cuneiform research and instruction. Latin transliteration represents the researcher's context-bound interpretation of the cuneiform text, but for effective automated cuneiform text processing, we also need the simple context-free description of the text provided by a native cuneiform computer encoding.

The Digital Hammurabi Project [DHP] is a multi-discipline research effort begun at

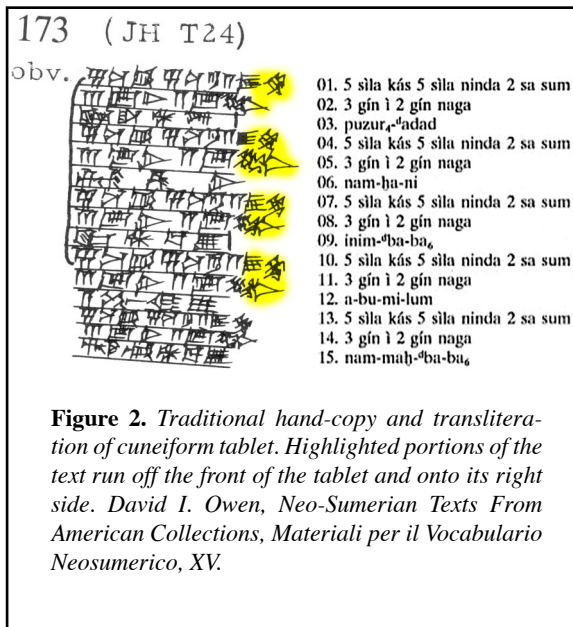


Figure 2. Traditional hand-copy and transliteration of cuneiform tablet. Highlighted portions of the text run off the front of the tablet and onto its right side. David I. Owen, *Neo-Sumerian Texts From American Collections, Materiali per il Vocabulario Neosumerico*, XV.

Johns Hopkins University aimed at providing solutions for both problems. We are developing 3D scanning and visualization technologies for cuneiform tablets along with a standard computer encoding for cuneiform text. iClay is the first publication of software showcasing our efforts to date and can be viewed over the Internet at this URL: <<http://www.jhu.edu/digitalhammurabi/iClay>>

2. Cuneiform as text

Cuneiform was used in the Ancient Near East for more than three millennia, from ca. 3200 BC to 100 AD, to write everything from ephemera to literary masterpieces, medical treatises to historical texts, mathematical and grammatical exercises, beer recipes, international treaties, musical scores, legal codes, religious rituals, sales receipts, maps, and astronomical tables. There are approximately 600,000 known cuneiform tablets and fragments in the world's museums, with perhaps 100,000 of them having been read by modern scholars. (The inventory of cuneiform tablets was dramatically increased just this last summer when, in Iraq, the birthplace of cuneiform, 200,000 previously unknown tablets were discovered in the basement of the Baghdad Museum. [CP04])

Sumerian-Akkadian cuneiform is a complex syllabic and ideographic writing system, with different signs for the various syllables, such as signs for “ba”, “bi”, and “sa” (but none for the phonemes “b”, “s”, “a”, and “i”). There are also signs for words, numbers, grammatical features (like “plural”), and certain unpronounced semantic classifiers (like “something made of stone” or “a god name”). The approximately 1000 different cuneiform signs form a complex script system in which practically every symbol is polyvalent, i.e., having

multiple phonemic and/or semantic realizations. Adding to the complexity of the script is the fact that the individual glyphs exhibit dramatic paleographic development over the three millennia of their use. In the early periods the signs were large, complex and curvilinear, resembling at times the very objects they signified; in later periods the signs become simpler and more abstract.

In 1999 members of the Digital Hammurabi team began efforts at Johns Hopkins University to encode Sumerian-Akkadian cuneiform in Unicode. This Initiative for Cuneiform Encoding (ICE) brought together an international team of leading cuneiformists, Unicode experts, software engineers, and font architects to work out the theoretical and practical issues surrounding this complex encoding.

ICE participants agreed on the following principles for the encoding of Sumerian-Akkadian cuneiform:

- Encode all characters occurring in UR III and later texts
- Include pre-UR III characters if well established
- Order cuneiform characters according to Unicode name
- Use UR III period for the representative glyphs
- Encode complex signs as single characters
- Encode compound signs as separate characters
- Use the standard, modern names for the characters
- Encode each unit number separately
- Encode all members of sign mergers and splits

After 4 1/2 years of effort, involving over 3,000 emails, two international conferences, and hosting a Unicode Technical Committee meeting at Johns Hopkins University, the Initiative for Cuneiform Encoding, under the auspices of the Digital Hammurabi Project, successfully completed the historic effort toward the international standard computer encoding of Sumerian-Akkadian cuneiform when both the Unicode Technical Committee and the ISO 10646 Working Group 2 unanimously approved the Initiative for Cuneiform Encoding proposal in their June 2004 meetings in Toronto, Canada.

The cuneiform sign repertoire embodied in the ICE proposal is a truly historic milestone in cuneiform studies, representing, as it does, the merger of the three most important cuneiform sign lists in the world, all of them previously unpublished, and with some of them in development for more than 40 years. The authors of these three lists, Dr. Miguel Civil, Professor of Sumerology, Oriental Institute, University of Chicago, Editor, Materials for the Sumerian Lexicon, Prof. Dr. Rykle Borger, Georg August Universitaet, Goettingen, author of Mesopotamisches Zeichenlexikon, and Dr. Robert England, Director of the Cuneiform Digital Library Initiative, UCLA, all gave ICE permission to use their unpublished materials to generate our master list of cuneiform signs for the Unicode encoding.

Row 120: CUNEIFORM

| | 1200 | 1201 | 1202 | 1203 | 1204 | 1205 | 1206 | 1207 |
|---|------|------|------|------|------|------|------|------|
| 0 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 1 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 2 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 3 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 4 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 5 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 6 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 7 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 8 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| 9 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| A | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| B | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| C | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| D | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| E | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |
| F | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 | 𐎗 |

Figure 3: A sample glyph chart from the Unicode cuneiform proposal. Font by Dr. Steve Tinney, University of Pennsylvania.

Dr. Steve Tinney, editor of the Pennsylvania Sumerian Dictionary, University of Pennsylvania, edited the cuneiform proposal and created the cuneiform font (Cuneiform Classic) for the Unicode Consortium. [EFT04] (Figure 3)

We estimate that cuneiform will appear in version 5 of Unicode in 2005, with all the major computer operating systems and programming languages supporting it.

In the very near future cuneiform scholars will be able to work with native cuneiform text on their computers as easily as, for example, one does now with Japanese. And with an international standard encoding for cuneiform in place, we expect to see the long anticipated growth in software tools for cuneiform philology - automated character recognition, natural language processing, sign concordancing, etc.

3. Cuneiform as artifact

The multi-tiered three-dimensionality of cuneiform tablets has, of course, driven our decision to explore 3D techniques for data capture and visualization of cuneiform tablets.

3.1. 3D Surface scanning

Much research has been devoted over the last several decades to three-dimensional surface scanning technologies. One pioneer in this field, the Visual Information Technology group of the National Research Council of Canada, has developed a three-color laser surface scanner that can scan three-dimensional artifacts of widely varying sizes at approximately 250 dots per inch (100 dots per cm). The NRC technology has been licensed by several companies, including Arius3D, Kestral3D, and Neptec, and has been successfully used for artifact reproduction in museum settings. [RBB*00] [BBC*00] [Rio99]

The Digital Michaelangelo Project at Stanford University Computer Graphics Laboratory has done pioneering research with very large-scale 3D data sets resulting from their work on laser scanning and visualizing of massive statues. [LPC*00] In addition they have published some interesting work on flattening 3D models of cuneiform tablets. [AL02]

Hewlett-Packard has taken an image-based approach to representing and rendering cuneiform and other artifacts. They capture pictures of a tablet from a single viewpoint using a myriad of individual lighting directions. From this data, they derive Polynomial Texture Maps, which allow interactive relighting of the surface viewed from that single viewpoint. [MGW01] We favor a more geometric approach in our work, because it allows interactive viewpoint control, geometry-derived rendering enhancements, etc.

The Egyptian Culture Project, a joint venture between IBM Research and the Egyptian government, has recently used a combination of 2D and 3D imaging to populate the Eternal Egypt web site [FEE*03]. Like us, they use a combination of geometry scanning with photometric capture to provide the necessary detail for their models. Their system uses laser scanning rather than structured lighting techniques.

Though not using 3D images, we must mention here the impressive Cuneiform Digital Library Initiative at UCLA, which combines two-dimensional scans of cuneiform tablets with transliterations of the texts, providing access over the Internet to tens of thousands of hitherto unpublished cuneiform tablets. [CDLI]

In defining our scanning system concept, we considered the general requirements of resolution, accuracy, speed, complexity, size, and cost. Specific technical requirements called for acquiring color, texture, and shape of clay artifacts

at a resolution on the order of 25 micrometers (1000 dots per inch, or 400 dots per cm). Tablets to be scanned ranged from thumbnail size to briefcase size.

In approaching this problem, we reviewed a number of technologies that showed promise for reaching one or more of our goals. What we found was that a number of specific technologies satisfied one or more requirements, but fell substantially short on others. For example, structured lighting could provide the necessary shape resolution over small regions, but could not provide color or texture information. Another example is the confocal scanning technology from NRC, Arius 3D, Kestral3D, Neptec, and others. [RBB*00] [BBC*00] [Rio99] The color information is excellent, but the resolution, though the best in the industry for this technology, does not meet our requirements of 1000 dots per inch. Moreover the high cost, lack of portability, and data capture rates of existing technologies do not fit our need to scan hundreds of thousands of tablets in remote locations.

Our conclusion was that no single technology would meet all our requirements. We were thus lead to combine a number of technologies, each of which complemented the shortcomings of the others. Further, we found that a flexible architecture provided the ability for the scanning system to evolve with improvements in the state of the various technologies.

In considering all these conflicting requirements, we arrived at a system concept that made use of a combination of photometric stereo and structured lighting. These separate technologies can be tightly coupled into a single system.

As illustrated in Figure 4, the system consists of a rotating platen on which we place the artifact to be scanned. Attached to this rotating platen is a high-resolution monochrome CMOS camera. Illumination is provided by a digital projector. This projector provides HDTV resolution and a contrast ratio of 2000:1.

Using these two basic components:

- photometric sensing provides local slope at the full resolution of the camera
- fringe projection provides shape at a somewhat reduced resolution, and
- color is provided at full resolution.

Recent advances in CMOS cameras have made this technology very attractive. For low cost, it provides small pixels, large dynamic ranges, and low noise. Further, it allows selective ROI readout and on-chip binning for selection of variable resolution. Square pixels 6.7 micrometers on a side along with a 4X telecentric lens yield 26.8 micrometer pixels in object space and an orthoscopic projection. Of course different pixel sizes and object fields can be effected by choice of proper lens magnification.




Figure 4: *Illustration of Scanning System*

The lighting makes use of the Texas Instruments Digital Light Processing (DLP) chip used in a number of consumer grade projectors. The novel feature of this projector technology is that it provides in one apparatus a means of fringe projection for the structured lighting, multiple colors for acquisition of color imagery with a monochrome camera, and illumination for photometric stereo. The existing projection optics are augmented with a custom lens to provide the corresponding orthoscopic fringe projection.

Photometric stereo is based on the assumption that the surface under inspection is Lambertian, i.e., that its reflectivity displays a cosine response. In such a case, the apparent brightness of the surface depends only on the position of the light source with respect to the local surface normal. Under this assumption, there are three degrees of freedom for the surface characterization, the two components of slope and the reflectivity. Thus, three light sources are required to infer the local surface slope and reflectivity. Generally, more than three sources are used and the reflectivity and slope are estimated based on a least-squares solution. For a typical scan we may acquire as many as 12 images, as a given point on the object may be shadowed in some of the images.

A further detail of the photometric stereo concept is the implicit assumption that the surface is illuminated by a point source at infinity. [RB99] [Woo80] The unique aspect of our implementation of photometric stereo is that instead of using multiple light sources, we use a single fixed light source and move the object/camera reference frame instead. This arrangement avoids the issue of inter-source illumination variation. The polar angle of the light source is fixed and the apparent azimuthal position of the light source can be varied

at will throughout 360 degrees. In any practical implementation of photometric stereo, there are departures from the ideal. To cope with such non-ideal features, we are implementing a singular value decomposition solution (rather than a simple least squares) that admits, for example, to some variation in the local illumination direction and small uncertainties in the azimuthal positioning.

It is well known that any real surface, however diffuse, departs somewhat from the idealized cosine behavior of the Lambertian model. As a result, photometric estimates of the surface normals are more accurate for small slopes and less so for large ones. This shortcoming leads us to our choice of the complementary technology of structured lighting.

Shape data are acquired using a phase-stepped fringe projection concept. [GLT96] [HH99] [JB94] A series of various spatial frequencies is used to resolve any range ambiguities associated with a single frequency, and phase-stepping is used to obviate any effect due to non-uniform illumination or object reflectivity. At the moment we are using a four-phase technique with seven orders of gray-scale sinusoids. Binary periodic patterns are also easily implemented and for such grating patterns, a greater number of orders is possible. The optimum modulation scheme and number of fringe orders depends upon a number of parameters such as camera noise level and projector resolution and grayscale quantization step size. We are currently studying this aspect of the problem. Of course the resulting shape information is at a lower spatial resolution than that of the slope information provided by the photometric measurement. The utility of this shape information is twofold. For rendering images, it provides a hull for draping of the higher resolution slope information from photometric stereo. Further, it provides a means of correcting the long-range order of the shape description derived from integration of the slope information. In other words, it allows correction for some of the known distortions inherent in photometric techniques. This correction step can be implemented in several ways. A global spatial frequency domain integration technique is possible as well as a local spatial domain integration technique using an over-determined set of equations [BFL*03]. This latter technique is preferred as it allows matching of the local slope integration domain to the resolution of the structured light data. Implementation is via a simple convolution of the pseudo-inverse gradient matrix with the slope "image."

Color information is provided by sequential illumination using the projection of uniform light fields in red, green, and blue. This approach avoids the reduced resolution of a color camera and retains perfect pixel registration with the photometric stereo and structured lighting data.

For acquisition of photometric stereo and color data, careful characterization of the camera and light source are required. Flat-fielding of the camera removes the fixed pattern

noise that is typically seen with CMOS technology. Linearity is established through use of a series of differing integration times. Camera noise level is also measured. This is an important parameter in establishing thresholding levels to designate shadowed pixels prior to calculation of local slopes. Local illumination variations throughout the sample volume are also measured prior to data acquisition. Color balance is established by acquisition of red, green, and blue images of a white diffuse reference surface.

Our fundamental approach to this problem is the use of multiple technologies and an architecture that can accommodate evolution in the state of the art in each component. For example, we previously made use of CCD cameras, but the fundamental architecture allowed easy substitution of CMOS cameras. Likewise, we previously used a laser stripe projection scheme for structured lighting and separate red, green, and blue LED's for acquisition of color and photometric data. We had explored liquid crystal light projectors, but rejected them due to their low contrast and resolution. Nevertheless, rapid changes in MEMs technology driven by the consumer market lead to economical high performance digital light projectors.

If one considers the whole of the data acquisition process as a system, it quickly becomes apparent that it consists of hardware and software components. One approach is to concentrate on the hardware and relegate the software to simple data management. We have chosen the somewhat different approach of trading processing complexity for hardware performance. This philosophy too, provides a straightforward means of exploiting the rapid evolution of computing capability.

3.2. 3D and 2D+ visualization

The scanner prototype provides us with geometry, normal, and color data, as shown in Figure 5. Given this initial data, we have developed a novel geometry-normal filter that removes additional noise while moving the geometry to more closely match the normal data. We successively place the filter kernel at each geometry location. Each sample within the kernel approximates the local surface as a plane using its geometry and normal information. Each of these planes is used to predict the surface position at the center of the kernel, and we take a weighted average of these predictions to replace the geometry at the kernel center. The weighting may be based on a simple Gaussian scheme, or weighted by confidence. Although the resulting geometry does not exactly match the normal data, it is much closer, and the wedges become more easily distinguished (even when we render in a "flat-shaded" style).

In our computer graphics laboratory, we regularly visualize these tablet models using 3D visualization. Such visualizations involve interactively manipulating a virtual camera




Figure 5. Data from the Johns Hopkins Applied Physics Laboratory scanner prototype. (a) original geometry mesh acquired using structured light (flat shaded) (b) normal map acquired using photometric stereo (c) color imagery with illumination factored out (d) geometry mesh with geometry-normal filter kernel applied (flat shaded) (e) filtered mesh rendered with normal map applied (f) filtered mesh rendered with normal and color maps applied

in 3D space to achieve the desired view while simultaneously rendering an image using some combination of local illumination and shadows as well as more non-photorealistic enhancements such as accessibility [Mil94], curvature [AL02], and depth-based shading effects [KSD03].

However, pure 3D visualization has a significant learning curve for unpracticed users (e.g. many cuneiformists). The most challenging aspect of typical 3D visualization software is the manipulation of the virtual camera. Such a camera actually has six degrees of freedom -- three degrees of freedom for translation and three more for rotation. Although we can work to make the interface as simple to use as possible, for example adding constraints to restrict motion to one dimension of translation or rotation at a time, the inherent dimensionality of the problem remains.

Furthermore, if we want to describe some portion of a cuneiform text, we do not always want to require the reader to navigate to the appropriate portion of the tablet model to follow the discussion. We wish to, in some sense, pre-digest the tablet, presenting appropriate parts to the reader in conjunction with descriptive text in modern fonts as well as transcription in cuneiform font.

To solve both these problems, we consider the cuneiform visualization problem in the context of document creation and document viewing.

We have developed an extended 2D image format (2D+) and an associated Java viewing applet which may be embedded along with standard Unicode text in the context of an HTML document. (Figure 6) The main text of the document is created in any Unicode-compliant text or html editor. The cuneiform images are created using our 3D tablet viewer. The author navigates using the virtual camera to get the desired view, sets the desired lighting direction and other shading parameters, chooses a rectangular crop area, and saves an associated image file. The author then embeds the viewing applet in the html document at the desired location, with the newly created image file as the source.

But the image file does not merely store the colors resulting from the 3D visualization program. Our custom image format stores a number of additional data attributes at each pixel: surface color, surface normal, accessibility, curvature, and depth (or some subset of these). We can think of these layers as a color image, normal image, accessibility image, curvature image, and depth image. Each of these values is taken from the 3D surface data and essentially resampled onto the rectangular grid of this 2D image.


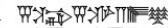

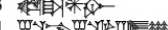


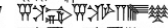

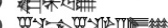

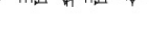

| 2D+ IMAGE | ENCODED CUNEIFORM | TRANSLITERATION | TRANSLATION |
|---|--|--|---|
|  | 1  | 5 sila kás 5 sila ninda 2 sa sum | 5 liters of beer, 5 liters of gruel, 2 strings of garlic, |
| | 2  | 3 gín ì 2 gín naga | 1/20 liter of oil, 1/30 liter of soap |
| | 3  | puzur ₄ - ^a adad | Puzur-Adad |
| | 4  | 5 sila kás 5 sila ninda 2 sa sum | 5 liters of beer, 5 liters of gruel, 2 strings of garlic, |
| | 5  | 3 gín ì 2 gín naga | 1/20 liter of oil, 1/30 liter of soap |
| | 6  | nam-ha-ni | Namhani |
| | 7  | 5 sila kás 5 sila ninda 2 sa sum | 5 liters of beer, 5 liters of gruel, 2 strings of garlic, |
| | 8  | 3 gín ì 2 gín naga | 1/20 liter of oil, 1/30 liter of soap |
| | 9  | inim- ^a ba-ba | Inim-Baba |
| | 10  | 5 sila kás 5 sila ninda 2 sa sum | 5 liters of beer, 5 liters of gruel, 2 strings of garlic, |
| | 11  | 3 gín ì 2 gín naga | 1/20 liter of oil, 1/30 liter of soap |

Figure 6. This screen shot illustrates just one way the publication of cuneiform corpora can be enhanced by the use of 2D+ images of cuneiform tablets combined with encoded and transliterated cuneiform text.

Advantages of our 2D+ technique include the fact that, since we are not doing any triangle rasterization, the rendering is more portable across platforms and the technical requirements for computers, operating systems, and web browsers are less than they would be for full 3D rendering. Future work on adding lossy compression to our custom 2D+ image format will further decrease file sizes and, therefore, download times.

The iClay 2D+ viewer applet initially uses the lighting direction as well as parameters for depth or accessibility shading as specified by the author as default values, so the image initially appears exactly as the author intended. But the applet also allows the reader to interactively manipulate the lighting direction and other parameters. These controls are considerably easier to operate than virtual camera manipulations. For example, the control of lighting direction has only two degrees of freedom, and the remaining parameters

each require only a slider for a single degree of freedom. Thus these enhanced images allow useful interactive manipulation of the tablet visualization without requiring the reader to do complex camera manipulations to find the area of interest.

The user interface for this viewer applet has been purposely designed to keep it out of sight during ordinary perusal of the document. It is fully encapsulated in a pop-up menu which remains hidden when the image is not being manipulated. This makes the document easy to read, while retaining the flexibility for the user to adjust parameters. Even with the expert selection of default parameters by the document author, we find that the ability to adjust the lighting and other parameters dynamically gives the reader a better understanding of the 3D shape of the tablet and its wedges.




Figure 7. Screen shots of the iClay Java applet illustrating a few of the various graphic operations that can be applied to a 2D+ image of a cuneiform tablet.

We have also developed a Java version of our full 3D tablet viewer for the purpose of simple remote manipulation of geometrically simplified versions of the 3D tablet data. In the future, we plan to develop a 3D viewer that incorporates view-dependent, progressive level of detail that can selectively refine the detail of tablet model in those areas the reader wishes to examine more closely.

3.3. Visualization techniques

This version of iClay demonstrates two visualization techniques useful for reading virtual cuneiform tablets - accessibility shading and depth shading. (Figure 7)

Accessibility shading computes the maximum radius of a sphere that can fit in each depression thereby approximating how much light reaches inside it. It finds narrow areas and darkens them. [Mil94]

Depth shading computes the local depth of each point from a statistically constructed surface in its neighborhood and uses that to modify the amount of light reaching that point for shading. This assumes that on average deeper recesses receive less light.

4. Future work

In this paper we have described our ongoing efforts to bring cuneiform tablets and cuneiform text into the digital age through standards initiatives, 3D scanning, visualization, and mixed-mode interactive document creation. There is still a great deal to be done on these fronts and others. [KSD03]

In 3D scanning, we are working to more rigorously characterize and improve the feature size resolution of our scanner, as well as to better employ the data gathered through photometric stereo and structured light techniques to mutually correct each other.

In 3D visualization, we are investigating feature detection and feature-based registration techniques to improve automation of the registration and reconstruction processes.

For end-user software, we continue developing better, more powerful cross-platform software tools to aid in the research and visualization of cuneiform tablets.

In the long term, we plan on creating a 3D digital archive of the world's cuneiform tablets and provide easy and convenient access to them as widely as possible.

5. Acknowledgements

We would like to thank the National Research Council of Canada for providing us with some initial tablet scans.

Thanks to Marc Levoy and the Stanford Computer Graphics Laboratory for making a number of useful suggestions and for the use of their freely available software tools and data. A special thanks goes to all the participants in the Initiative for Cuneiform Encoding and especially to Miguel Civil, Robert England, and Rykle Borger for providing access to their unpublished cuneiform sign list materials, and to Steve Tinney for doing most of the work on bringing it all together into the formal Unicode proposal. Apple Computer provided an Xserve G4 server and RAID storage device in support of the project. This project has also benefitted from the help of the other members of the Johns Hopkins Computer Graphics Lab, the Johns Hopkins Applied Physics Laboratory, and the Department of Near Eastern Studies. This work has been supported by NSF Medium ITR IIS-0205586.

References

- [AL02] Anderson S., Levoy M.: Unwrapping and Visualizing Cuneiform Tablets. IEEE Computer Graphics and Applications Vol. 22, No. 6, November/December, 2002, pp. 82-88. <<http://graphics.stanford.edu/papers/cuneiform/>>.
- [BBC*00] Beraldin J.-A., Blais F., Boulanger P., Cournoyer L., Domey J., El-Hakim S.F., Godin G., Rioux M., Taylor J.: Real world modelling through high resolution digital 3d imaging of objects and structures. ISPRS Journal of Photogrammetry & Remote Sensing 55:230 P250,2000.
- [BFL*03] Barchers J. D., Fried D. L., Link D. J., Tyler G.A., Moretti W, Bremann T. J., Fugate R Q., Adaptive performance of wavefront sensors in strong turbulence, in Adaptive Optical Systems Technologies II, Proceedings of the SPIE, V. 4839, pp. 217-227. 2003.
- [CDLI] Englund R., Damerow P.: Cuneiform Digital Library Initiative. UCLA <<http://cdli.ucla.edu/>>.
- [CP04] Colonnelli, L., Pettinato, G.: Così abbiamo ritrovato 300 mila tavolette inedite. Corriere Della Sera. Roma May 5, 2004. <<http://www.unipi.it/rassegna/archivio/2004/05/24817077.pdf>>
- [DHP] Digital Hammurabi Project. Johns Hopkins University <<http://www.jhu.edu/digitalhammurabi/>>.
- [EFT04] Everson M., Feuerherm K., Tinney S.: N2786: Final proposal to encode the Cuneiform script in the SMP of the UCS <<http://www.dkuug.dk/jtc1/sc2/wg2/docs/n2786.pdf>>
- [FEE*03] Farouk M., El-Rifai I., El-Tayar S., El-Shishiny H., Hosny M., El-Rayes M., Gomes J., Giordano F., Rushmeier H., Bernardini F., Magerlein K.: Scanning and Processing 3D Objects for Web Display. 4th International Conference on 3D Digital Imaging and Modeling (3DIM '03) Banff, Alberta, October 2003 pp 310-317. <<http://www.eternalegypt.org>>
- [GLT96] Gartner H., Lehle P., Tiziani H.: New, Highly Efficient, Binary Codes for Structured Light Methods. SPIE, Vol. 2599, 1996.
- [HH99] Horn E., Hiriyati N.: Toward Optimal Structured Light Patterns. Image and Vision Computing, Vol. 17, 1999.
- [iClay] iClay, Digital Hammurabi Project, Johns Hopkins University <<http://www.jhu.edu/digital-hammurabi/iClay>>.
- [JB94] Jiang X., Bunke H.: Range Data Acquisition by Coded Structured Light: Error Characteristic of Binary and Gray Projection Code. SPIE, Vol. 2252, 1994.
- [KSD03] Kumar S., Snyder D. A., Duncan D., Cohen J. D., Cooper J.: Digital Preservation of Ancient Cuneiform Tablets Using 3D Scanning. Proceedings of IEEE Fourth International Conference on 3-D Digital Imaging and Modeling, pp. 326-333, 2003.
- [LPC*00] Levoy M., Pulli K., Curless B., Rusinkiewicz S., Koller D., Pereira L., Ginzton M., Anderson S., Davis J., Ginsberg J., Shade J., Fulk D.: The digital Michelangelo project: 3D scanning of large statues. Proceedings of the 27th annual conference on Computer graphics and interactive techniques, p.131-144, July 2000 <<http://graphics.stanford.edu/projects/mich/>>.
- [MGW01] Malzbender T., Gelb D., Wolters H.: Polynomial Texture Maps. Computer Graphics, Proceedings of ACM Siggraph 2001 <http://www.hpl.hp.com/personal/Tom_Malzbender/research/research.htm>.
- [Mil94] Miller G.S.P.: Efficient Algorithms for Local and Global Accessibility Shading. Computer Graphics (Proc.Siggraph 94) ACM Press, New York, 1994, pp.319-326.
- [RBB*00] Rioux M., Blais F., Beraldin J.-A., Godin G., Boulanger P., Greenspan M.: Beyond Range Sensing: XYZ-RGB. Digitizing and Modeling. ICRA 2000: 111-115 and Visual Information Technology group, National Research Council of Canada <http://iit-iti.nrc-cnrc.gc.ca/about-sujet/vit-tiv_e.html>.
- [Rio99] Rioux M.: Three dimensional colour imaging. Canadian patent no.2,043,336. issued 1999/07/27.
- [Woo80] Woodham R. J.: Photometric method for determining surface orientation from multiple images. Optical Engineering 19, 1, 139P144. 1980.




Figure 6. Data from the Johns Hopkins Applied Physics Laboratory scanner prototype. (a) original geometry mesh acquired using structured light (flat shaded) (b) normal map acquired using photometric stereo (c) color imagery with illumination factored out (d) geometry mesh with geometry-normal filter kernel applied (flat shaded) (e) filtered mesh rendered with normal map applied (f) filtered mesh rendered with normal and color maps applied