Interactive Mapping on Virtual Terrain Models

Tony Bernardin Eric Cowgill Ryan Gold Bernd Hamann Oliver Kreylos Alfred Schmitt

Abstract

We present an interactive, real-time mapping system for use with digital elevation models and remotely sensed multispectral imagery, which aids geoscientists in the creation and interpretation of geologic/neotectonic maps at length scales of 10 m to 1000 km. Our system provides a terrain visualization of the surface of the Earth or other terrestrial planets by displaying a virtual terrain model generated from a digital elevation model overlain by a color texture generated from orthophotos or satellite imagery. We use a quadtree-based, multi-resolution display method to render in real time high-resolution virtual terrain models that span large spatial regions. The system allows users to measure the orientations of geologic surfaces and record their observations by drawing lines directly on the virtual terrain model. In addition, interpretive surfaces can be generated from these drawings and displayed to facilitate understanding of the three-dimensional geometry of geologic surfaces. The main strength of our system is the combination of real-time rendering and interactive mapping performed directly on the virtual terrain model, with the ability to navigate the scene while changing viewpoints arbitrarily during mapping. User studies and comparisons with commercially available mapping software show that our system improves mapping accuracy and efficiency, and also yields observations that cannot be made with existing systems.

I. Introduction

The surface topography of a terrestrial planet such as Earth or Mars preserves a unique record of both the types of processes that have shaped the planet-atmosphere interface, and the temporal sequence in

Bernardin, Schmitt: Institut für Betriebs- und Dialogsysteme (IBDS), Fakultät für Informatik, Universität Karlsruhe, Kaiserstrae 12, 76131 Karlsruhe, Germany; ug1g@rz.uni-karlsruhe.de, aschmitt@ira.uka.de

Cowgill, Gold: Department of Geology, University of California, One Shields Avenue, Davis, CA 95616, U.S.A.; {cowgill,gold}@geology.ucdavis.edu

Hamann, Kreylos: Institute for Data Analysis and Visualization (IDAV), Department of Computer Science, University of California, One Shields Avenue, Davis, CA 95616, U.S.A.; {hamann,kreylos}@cs.ucdavis.edu

which those processes have occurred. Because erosion and deposition of sediment, volcanic eruption, and tectonic deformation all leave fingerprints in the landscape as they shape it, we can understand both the types of processes that have operated in the past, and the relative sequence in which they occurred, if we can identify and measure the distinctive sets of geometric patterns in the landscape that each process produces. To make use of this record, Earth and planetary scientists need both high-resolution (1-10 m/pixel) data that characterize the shape of the planetary surface, and methods for analyzing such data.

High-resolution terrain data (digital elevation models, orthophoto mosaics, or multispectral satellite imagery) with regional to planetary coverage are increasingly available, amplifying the need for tools that facilitate analysis of these data. Two different approaches have been used to design such tools. In the first the user specifies a set of criteria by which a given type of feature can be identified, and then uses these criteria to create a set of formal rules in a computational algorithm that is used to filter the data in an automated routine, permitting users to rapidly analyze large (GB to TB) data sets such as a $700\,\mathrm{km} \times 700\,\mathrm{km}$ mosaic of $15\,\mathrm{m}$ ASTER data draped over a $30\,\mathrm{m}$-resolution DEM (digital elevation model) from the Shuttle Radar Topography Mission. A second approach is to visualize the data and directly observe, measure, and record geometric patterns using the human brain to filter the data instead of an algorithm. This latter process of geomorphic and geologic mapping has the advantage of allowing users to adapt or evolve their selection criteria as the analysis proceeds, and to discover features or patterns in the data that were unrecognized prior to the analysis. But this approach also requires the data to be displayed, which typically exceeds current hardware capabilities.

Earth and planetary geologists have struggled for some time with this limitation and in particular with the difficult task of mapping their detailed observations across large areas directly on remotely sensed data. One compromise is to degrade the data by displaying it at reduced resolution. A second is to restrict high-resolution studies to small ($10\,\mathrm{km} \times 10\,\mathrm{km}$) areas at a few, widely spaced localities, and to then analyze those sites using stereo images which restrict geologists to plan views of the terrain. Unfortunately, both methods yield a strongly filtered view of the geologic and geomorphic record.

In response, we have developed RIMS, a *real-time, interactive, mapping system*. This brief aims to explain the design of RIMS, illustrate its utility using movies[1], and provide access to the application at http://keckcaves.ucdavis.edu/RIMSG3[2]. As part one of Movie 1 illustrates, the interactive 3D terrain visualization environment of RIMS operates much like Google Earth [1] in the way that it allows users to "fly" in real-time over virtual terrain models. These models are composed of texture data (e. g., a false-color satellite image) draped over a DEM. To facilitate real-time visualization of large virtual terrain models, RIMS employs level-of-detail display techniques, as is shown in the second part of Movie 1 and in Movie 2. Most importantly, RIMS also allows geologists to map by 1. drawing polylines directly on the 3D terrain visualization (Movie 3), 2. measuring the orientations of surfaces using a virtual geologic compass (Movie 4), and 3. reconstructing the 3D shape of geological surfaces by surface extrapolation above and below the surface of the virtual terrain model. For present purposes, we use *terrain data* to refer to remotely sensed data that characterizes the morphology and/or spectral response in the electromagnetic domain of a planetary surface. Examples of terrain data include digital elevation models, orthophoto mosaics, or multispectral satellite data. *Virtual terrain models* are built from terrain data, and consist of a texture image that is draped over a digital elevation model. A *terrain visualization* is a graphical display of terrain data, either in its native format, a stereopair, or a virtual terrain model.

II. Real-time Visualization of Terrain Data

A. Motivation

For geologic and geomorphic mapping, geologists need to study the detailed 3D structure of a planet's surface. Most display devices, however, only support 2D images, requiring perspective projection to represent 3D models. This technique is insufficient for conveying proper depth cues because 3D shapes

[1]Supporting materials (Movies 1-4) are available via Web browser or Anonymous FTP from ftp://ftp.agu.org/apend/ (Username = "anonymous", Password = "guest").

[2]Both the RIMS application and the supporting program needed to generate virtual terrain models from digital elevation and texture data are freely available via Web browser from http://keckcaves.ucdavis.edu/RIMSG3. This site also provides access to the movies and a demonstration version of the program with a preprocessed virtual terrain model.

can be difficult to perceive on static images. Allowing geologists to manipulate the viewing parameters so that they may look at an object from different angles, facilitates the perception of its 3D shape. Doing so interactively – providing immediate visual feedback through smooth transitions between views – provides a motion parallax effect that greatly enhances 3D perception. The third part of Movie 1 illustrates this point, and in particular, demonstrates how the depth of a river gorge is readily identified in RIMS using an interactive visualization but is difficult to observe in a static (2D) plan view. To support mapping effectively, RIMS must provide a highly detailed terrain visualization that allows a user to manipulate the virtual terrain model in real time. This means a user must be allowed to change viewing parameters at any time, with the display adjusting smoothly (at least 30 frames per second).

B. Challenge

The biggest hurdle in rendering detailed terrain data is the large data volume: in general, current hardware is not powerful enough to generate full-resolution 3D images in real-time from terrain data at its native resolution. We circumvent this problem by employing level-of-detail (LOD) display techniques, such that only a carefully chosen, triangle-mesh approximation of the virtual terrain model needs to be processed. Such approximations must balance the need for accurate representation of the data with the effective use of available processing resources. Generating them efficiently can be quite difficult and several different approaches have been developed in academia [2], [3], [4], [5], [6].

C. RIMS Approach

In designing the visualization component of RIMS, the specific needs of the application were carefully considered. During geologic or geomorphic mapping, users generally operate on a zoomed-in portion of the virtual terrain model (often down to the native resolution of the original terrain data), effectively defining a region of interest. Thus, the approach of our visualization is not to display a global approximation of the whole scene, but rather to show a limited focus area as precisely as possible (i. e., only detail that surpasses the limitations of display hardware can be safely hidden), while outlying areas

are shown at lower resolution to provide context. By moving a cursor, users chose what portion of the data is shown at maximum resolution, as shown in part 2 of Movie 1 and in Movie 2. A hierarchical LOD approach that manages the terrain data as tiles of progressively lower resolution, akin to that of Lindstrom et al. [7], is adequate for such a representation of the virtual terrain model. In the following, we outline this approach and its implementation in RIMS:

a) Preprocessing of the Tile Hierarchy: In our approach we use multi-resolution representations based on a quadtree subdivision. A quadtree's lowest-level nodes correspond to a tiling of a given data set at its native resolution. Higher-level nodes contain successively coarser representations, subsampled by a factor of two between levels. Since terrain data are typically provided only at their native resolution, we generate quadtree hierarchies of tiles for both the DEM and texture image in a preprocessing step. These hierarchies then provide the appropriate data for our approximation of the terrain (see Figure 1).

b) Defining Focus and Context: When using RIMS, users manipulate a cursor that is bound to the surface of the virtual terrain model. The cursor is the user's interface to the mapping capabilities of RIMS (see section III-C below). As shown by Movie 2, we define the high-detail focus area as a configurable radius around the cursor. The remaining terrain outside the focus is considered to be part of the context and is therefore approximated more coarsely so that it still provides spatial reference within the data. Movie 2 shows the virtual terrain model as a triangulated mesh to graphically illustrates how the high-resolution (i. e., higher density of triangles) portion of the display fades with distance from the cursor. It also illustrates the quadtree data structure.

c) Generating the Surface Approximation: To effectively visualize the multi-resolution virtual terrain model, we need a method for evaluating the *appropriateness* of each tile within the quadtree hierarchy, that is, whether or not the tile should be displayed as part of the visualization. A first criterion is given by

the focus: within the focus region the highest possible resolution is desired, with tiles gradually falling off in resolution with increasing distance from the focus point. A second criterion integrates view-dependent characteristics: tiles not currently within the view can be discarded; and tiles that present sub-pixel detail when projected onto the display are defined to have too high a resolution for the current approximation.

With such measures of tile appropriateness, the surface approximation can be generated by traversing the quadtree hierarchy from the root (the coarsest resolution tile) through the subnodes, evaluating the tiles encountered on the way and retaining only those that are neither too coarse nor too fine. Although the quadtree traversal must produce a tiled image without gaps or overlaps, cracks might appear between neighboring tiles of different resolutions due to hanging triangle vertices (see Figure 2, left). To address this problem, we modify the LOD criteria such that direct neighbors differ at most by one level of resolution (see Movie 2). Cracks can then be removed by simply *stitching* the edges of affected tiles, as illustrated in Figure 2.

d) Rendering: As a final step, the approximation computed for a given set of viewing and focus parameters must be rendered to produce the corresponding screen image. Since only a subset of the DEM and texture tile hierarchies needs to be drawn for a given frame, terrain data sets much larger than the workstation's main memory can be visualized. Data is swapped in and out of memory from a mass storage device as needed, with the memory serving as a cache for the data that are immediately consumable by the graphics hardware.

III. Geologic and Geomorphic Mapping

A. Motivation

The visualization technique described above provides perspective views of texture data draped over a DEM. A large number of widely available tools support interactive manipulation of such displays [8], [9], [10], [11], [12], [13] the most popular of which is arguably GoogleEarth [1]. Only Google Earth

and ArcGlobe [14] provide multi-resolution techniques. More importantly, it is generally not possible to easily conduct geologic and geomorphic mapping directly on the 3D scenes. To grasp how useful such a feature can be, it is necessary to understand that the mapping process is not one where geologic features are merely observed and then recorded using symbols (points, lines, etc.). Instead the process fundamentally relies upon iterative analysis: initially a feature is observed and tentatively mapped and classified, with final assessment coming only after further analysis, such as measuring the orientation of the feature, establishing the geometric relationship of the observed feature with other mapped features, or visually inspecting adjacent portions of the data to evaluate the continuity of the feature. A direct, interactive mapping capability would, thus, greatly empower geologists and geomorphologists.

B. Challenge

One of the more significant technical challenges facing the addition of geologic mapping to a real-time, interactive terrain visualization is developing a method for annotating the 3D surface model with symbols. These annotations must closely follow topography to accurately reflect geospatial locality. Given the multi-resolution approach used, the topography of the terrain visualization is a continuously adapted approximation of the true terrain data (see section II-C.0.c). Thus, as Figure 3 illustrates, appropriately adapted annotations must also be generated each time the terrain visualization updates. Techniques have been proposed to address this problem for annotations rendered using geometric primitives [15] as well as texture images [16]. The changes in the surface approximation can produce analogous problems for other virtual mapping tools. Thus, when implementing these tools, care must also be taken as to how the data are best considered for the desired manipulation.

C. Implementation in RIMS

The support RIMS offers geoscientists for conducting geologic and geomorphic mapping can be classified into three activities: (a) “drawing” allows the direct annotation of the 3D surface model with lines; (b) “measuring” provides the ability to quantify the orientations of 3D features; and (c)

"extrapolation" permits geological surfaces to be reconstructed, thereby giving form to estimated 3D geologic models that are superimposed on the virtual terrain model for immediate evaluation and modification. It is important to note that users can continue to interactively manipulate the virtual terrain model while using these mapping tools, thereby facilitating precise and accurate geologic and neotectonic mapping.

a) Drawing: Most commonly, geologists record the locations of geologic/geomorphic features using a connected sequences of line segments, i. e., polylines. Thus, our system supports mapping with this primitive. By controlling a cursor bound to the surface of the virtual terrain model, users can perform actions such as creating, selecting, moving, deleting, and adding control points along a polyline, as Movie 3 illustrates.

Unlike the preprocessed terrain data, the vector mapping data are actively generated and modified via user interaction. Thus, the geometric representations of the mapped polylines must be generated in real-time with respect to the current surface approximation. To achieve this goal, RIMS exploits the locality of polyline manipulations (moving an inner control point, for example, modifies only the two adjacent line segments) and the strong frame-to-frame coherence (the terrain visualization changes only slightly from one image to the next). We store 3D representations for all polylines, and then *tweak* previously valid representations when polyline segments are edited or the level of detail of the terrain visualization changes.

b) Measuring: Not only must geologists be able to map the spatial distribution and 3D morphology of deformed topographic markers, they must also be able to measure the orientations of these features, particularly planar elements such as fault surfaces or dipping layers of rock on opposite sides of a fold. To accommodate this need, our system also provides a virtual geologic compass (VGC), as demonstrated in Movie 4. The VGC is a reference plane that intersects the virtual terrain model and that can be interactively manipulated by the user. Both the placement and orientation of the VGC plane

can be adjusted, allowing the user to either fit the plane to features visible within the texture or set the VGC plane to a known orientation for use as a reference plane. Multiple planes can be placed within the display, and each can be turned on or off independently. While a VGC plane is active or being manipulated, its orientation is displayed on screen as strike and dip angles to provide quantitative feedback.

c) Extrapolation: Many of the features geologists map and measure are undulatory lines formed by the intersection between a planetary surface and a geologic surface such as a fault or folded layer of rock (see Figure 4). RIMS facilitates the geometric interpretation and analysis of such features by allowing users to project and manipulate curviplanar surfaces through their mapping, thus displaying surfaces that can otherwise only be imagined or created through a time-consuming process of generating closely spaced serial sections.

In detail, geologic surfaces are reconstructed by considering the 3D positions defined by the control points along the polylines and locally fitting them to planes. These planes can be displayed directly to visualize the angle at which the mapped rock layer intersects the current topography. We can also consider a set of mapped polylines as defining the line of intersection between the topography and a geological surface that protrudes into (or dips beneath) the ground and under the topography. The desired surface is regenerated in alternating arcs above and below the surface of the virtual terrain model using pairs of polylines.

IV. Results

A. 3D Mapping

To evaluate the mapping capabilities of RIMS, we conducted two comparison tests between RIMS and an alternative mapping approach previously developed by one of the authors (Cowgill). This approach is analogous to viewing a pair of stereo images on a stereoscope and uses three components, a StereoGraphics Z-Screen, a stereo-capable graphics card, and software to drive them both. The Z-screen is an optically

polarized LCD panel that fits over a standard CRT monitor. The graphics card modulates the panel polarization at ≥ 70 Hz in synchronization with the image shown on the CRT. As a result, left- and right-eye images are displayed separately, permitting 3D visualization of stereo images by a user wearing passively polarized glasses. StereoAnalyst (SA) software [17] generates the terrain visualization by driving both the graphics card and the Z-screen, but requires users to first create a georeferenced input file called a digital stereo model (DSM). DSMs are built from a pair of stereo images by first picking matching pixels on both scenes to link them together. A set of ground control points is then used to link the images to a geographic coordinate system and thus georeference the DSM. The DSMs are used by SA to generate a terrain visualization by displaying the image pair in stereo using the Z-Screen. Polylines can then be mapped on this terrain visualization. In the alternative digital mapping method developed by Cowgill, DSMs are constructed from 15 m-pixel resolution ASTER stereo images using only the ground control point information provided with the Level 1A data. Individual DSMs comprise 4200×4200 pixel ($\approx 60\,\text{km} \times 60\,\text{km}$) images that can be mapped in 3D-space using (SA) at map scales up to 1:20,000. Although useful, the SA method is limited to a plan view, and thus does not permit a user to view the terrain in profile, a perspective that geologists rely upon heavily for analysis. In addition, the DSMs typically have lateral variations in the horizontal and vertical separation between matching points in the left- and right-eye images (i. e., X and Y parallax are not uniform). Users unconsciously compensate for these variations by changing how they look at the DSM, resulting in eye-strain after a few hours of visualization with SA. Finally, the lack of independently determined ground control information typically restricts the DSM to a single scene, thus only a $60\,\text{km} \times 60\,\text{km}$ area can be mapped at a time using SA.

The first of our two tests was designed to identify the most sensitive mapping system (see Figure 5). This test compared the maximum amount of geological information that could be extracted from a single data set in an unlimited period of time using each method. The second test was designed to identify the most efficient mapping system, and compared the number and quality of geologic observations collected from a single data set in a two-hour period of time (see Figure 7). The tests were performed by the

authors (Test 1: Cowgill, Test 2: Gold). In both cases, the study areas were mapped first with SA, then with RIMS. Any benefit derived from having already mapped the scene at the start of the RIMS session was likely offset by our lack of familiarity with the RIMS navigation/mapping environment at the time, as both Cowgill and Gold had been using the SA method for at least 6 months before using RIMS. The tests were the first time Cowgill and Gold used RIMS.

The two comparison tests indicated that RIMS reduced eye strain and provided faster zoom, pan, and rotation speeds. More importantly, the tests also revealed five key differences that make RIMS particularly useful for geologic/geomorphic applications. Relative to SA, RIMS provided: (a) better understanding of mapped structural geometries; (b) better confidence in feature identification and location; (c) the ability to identify and map a larger number of features; (d) more accurate mapping (i. e., a larger number of vertices per polyline); and (e) better ability to locate and identify small features. Specific examples of each difference are provided in the following sections, highlighting the utility of RIMS.

a) Understanding: The most important difference revealed by the tests is that RIMS allowed both users to obtain a more sophisticated understanding of the structural geometry of their areas. For example, in Figure 5, arrow A' on the right-hand side of the figure points to a structure that was obvious in the RIMS environment. The lack of a structure at arrow A on the left-hand side of the figure indicates that the user was not able to see and interpret this feature using SA.

b) Confidence: RIMS provided both users with higher confidence in their vector mapping, as indicated by the type of lines selected to represent mapped features. For this test the geologists expressed their confidence in their ability to accurately locate a mapped feature by using solid, dashed, or dotted lines (in order of decreasing confidence). Figure 7 shows that the RIMS project contains 20 boundaries mapped using solid lines, two using dashed, and two with dotted. In contrast, the SA project has only two boundaries defined with solid lines, 21 with dashed lines, and one with a dotted line.

c) Features: Both users were able to identify a larger number of features using RIMS than SA. The RIMS output shown in Figure 5 has 289 mapped features whereas only 172 features were extracted using

SA. Likewise, Figure 7 indicates that 14 major structures were defined using RIMS, in contrast to eight structures on the SA map.

d) Accuracy: RIMS allows users to more accurately locate features and map them using more vertices per feature because it does not demand constant manual parallax adjustments. Because the polylines have more vertices in the RIMS output, they better track short wavelength variations in the feature geometry and thus more accurately follow subtle changes in the boundaries between geologic units. In contrast, the maps generated from SA show a prevalence of long straight line segments. Differences in detail are especially evident in Figure 7 at comparison points A–A', B–B' and C–C' in SA and RIMS, respectively.

e) Small Features: Finally, RIMS is more effective for locating small geologic features. For example, a series of river terraces located at point D' in the RIMS output were not located at point A using SA (see Figure 7). Likewise, points E–E' indicate a small outcrop that was not seen in SA at E but that was mappable using RIMS at E'. Although these features are small, their identification has important implications regarding the geometry of active deformation in the mapped area.

B. 3D Measuring

The SA mapping method [17] does not contain a utility analogous to the virtual geologic compass tool in RIMS. Thus, to evaluate the performance of the VGC tool, we designed two tests using RIMS alone. The first was designed to determine how the accuracy of the VGC measurements varies as a function of dip angle and topographic relief. When geologic surfaces dip steeply or intersect areas of smooth topography, the mapped intersection between the surface and the terrain is highly linear, and a wide range of dipping planes appear to fit the observed line. Therefore, the first test was designed to quantify the extent to which these factors inhibit a user's ability to reproduce a plane of known orientation. The test design and results are described in detail in the appendix and show that the VGC is able to reproduce known orientations to within a few degrees.

In many cases, intersections between geologic surfaces and the terrain are not expressed as features that are clearly visible within the texture, making it difficult for users to fit the VGC plane to the image.

Therefore, a second test was designed to evaluate the extent to which the VGC tool can reproduce field-based measurements.

As is explained in the Appendix, the second VGC test indicates that the mean absolute difference between a set of field-based data and those measured with the VGC was 12.5 ± 8.5 degrees (1 SD) for strike and 21.1 ± 12.7 degrees (1 SD) for dip. These results indicate the VGC tool is useful as a semiquantitative utility that indicates where geologic surfaces are gently, moderately, or steeply dipping. The test also indicates that users can have more confidence in strike determinations than in dip measurements generated with the VGC plane.

C. Surface Extrapolation

Figure 6 illustrates the interpretive value of the surface projection utility in RIMS. Here, two lines mapped along opposite sides of a ridge delineate exposures of the same distinctive layer of sediment (left). The opposite dip directions of this layer and their positions along both flanks of the ridge indicate that ridge is likely a partially eroded fold. Disruptions of drainage networks that once crossed this fold suggest it is presently active. To gain an appreciation for the amplitude of the fold, and the magnitude of horizontal shortening it records, we use these mapped lines to generate an extrapolated surface (center) that shows how the marker bed might project across the fold. It becomes immediately evident that the initial interpretation of the fold geometry causes the marker surface to locally intersect the crest of the ridge, and thus predicts that the marker layer should also be present in this area. However, further inspection of the 3D dataset in this key area indicates that such exposures are not present. Thus, the amplitude of the fold must be increased to that shown on the right in the figure, so that the marker bed can clear the ridge crest and all observations are satisfied. Although the fold geometry is not uniquely constrained, this analysis does yield the minimum amplitude, and thus the minimum amount of shortening for this area.

V. Conclusions and Future Work

While the tests described above show that the maps generated using both utilities capture many of the same major geologic features, it is clear that RIMS provides both a more sensitive and a more efficient mapping method, and greatly advances geologists' ability to remotely map patterns of active deformation in fine detail while spanning continental collision zones that are thousands of kilometers wide and often inaccessible for field study. The advantages of RIMS over the previously used system are due to RIMS' ability to map directly on the virtual terrain model in real time and its interactive visualization of large terrain data sets.

An important future enhancement to the system will be to move it to a virtual reality environment for even greater usability and interaction. We also plan to incorporate a texture-based mapping representation to support a higher quantity of mapped polyines with a more varied appearance to facilitate understanding while mapping.

Acknowledgments

This work was supported by the National Science Foundation under contract ACI 9624034 (CAREER Award), through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251, through the National Partnership for Advanced Computational Infrastructure (NPACI), and a large Information Technology Research (ITR) grant. This work was also supported by the W. M. Keck Center for Active Visualization in the Earth Sciences (KeckCAVES) and NASA grant EOS/03–0663–0306. We thank Magali Billen, Louise Kellogg, and Nickolas Raterman for providing valuable input, and we thank the members of the Visualization and Computer Graphics Research Group at the Institute for Data Analysis and Visualization (IDAV) at the University of California, Davis. We thank two anonymous reviewers for their constructive comments on an earlier submission of this work to G-Cubed, one of whom provided a very helpful second review for the final work.

REFERENCES

[1] GoogleEarth. http://earth.google.com/.

[2] Hugues Hoppe. Smooth view-dependent level-of-detail control and its application to terrain rendering. In *Proceedings of the conference on Visualization '98*, pages 35–42. IEEE Computer Society Press, 1998.

[3] Junho Kim and Seungyong Lee. Truly selective refinement of progressive meshes. In *No description on Graphics interface 2001*, pages 101–110. Canadian Information Processing Society, 2001.

[4] P. Cignoni, F. Ganovelli, E. Gobbetti, F. Marton, F. Ponchio, and R. Scopigno. BDAM – batched dynamic adaptive meshes for high performance terrain visualization. *Computer Graphics Forum*, 22(3), 2003.

[5] Peter Lindstrom and Valerio Pascucci. Terrain simplification simplified: A general framework for view-dependent out-of-core visualization. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):239–254, 2002.

[6] Mark Duchaineau, Murray Wolinsky, David E. Sigeti, Mark C. Miller, Charles Aldrich, and Mark B. Mineev-Weinstein. ROAMing terrain: Real-time optimally adapting meshes. In *Proceedings of the 8th conference on Visualization '97*, pages 81–88. IEEE Computer Society Press, 1997.

[7] Peter Lindstrom, David Koller, Larry F. Hodges, William Ribarsky, Nickolas Faust, and Gregory Turner. Level-of-detail management for real-time rendering of phototextured terrain. Technical Report 6, 1995.

[8] J. M. Lees. Geotouch: Software for three- and four-dimensional GIS in the Earth sciences. *Computers and Geosciences*, 26:751–761, 2000.

[9] GRASS GIS. http://grass.baylor.edu.

[10] ArcScene utility in 3D Analyst extension of ArcGIS. http://www.esri.com/software/arcgis/arcgisxtensions/3danalyst/.

[11] 3D SurfaceView utility in ENVI. http://www.rsinc.com/envi/.

[12] VirtualGIS module in ERDAS IMAGINE. http://gis.leica-geosystems.com/Products/Imagine/.

[13] FLY! http://www.pcigeomatics.com/product_ind/fly.html.

[14] ArcGlobe, to be released in Spring 2004. http://www.esri.com/news/arcnews/summer03articles/introducing-arcglobe.html.

[15] Zachary Wartell, Eunjung Kang, Tony Wasilewski, William Ribarsky, and Nickolas Faust. Rendering vector data over global, multi-resolution 3D terrain. In *Proceedings of the symposium on Data visualisation 2003*, pages 213–222. Eurographics Association, 2003.

[16] Oliver Kersting and Jürgen Döllner. Interactive 3D visualization of vector data in GIS. In *Proceedings of the tenth ACM international symposium on Advances in geographic information systems*, pages 107–112. ACM Press, 2002.

[17] Stereo Analyst for ArcGIS.

http://gis.leica-geosystems.com/Products/StereoAnalyst/.

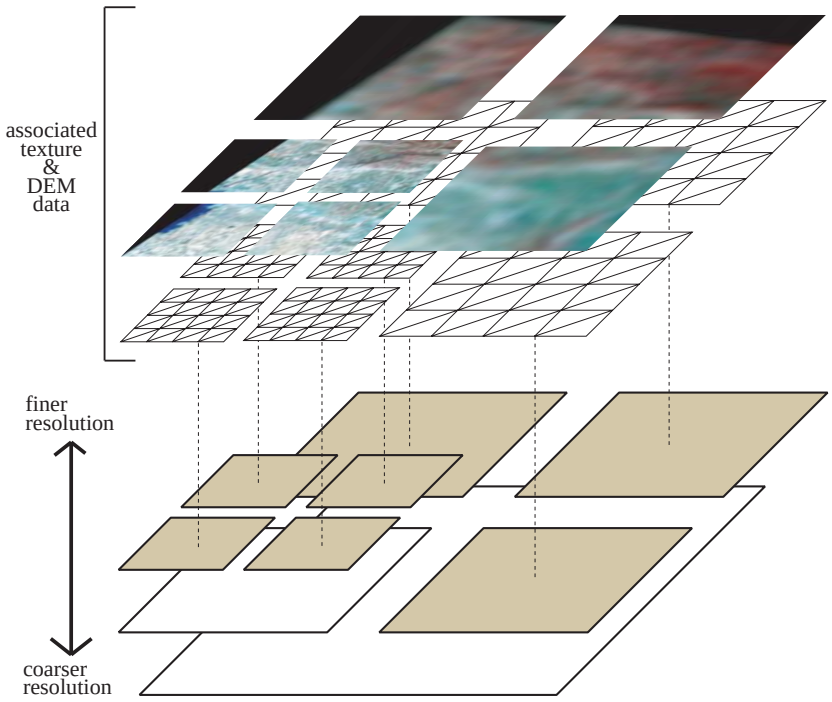

Fig. 1. A terrain tree illustrated with the root at the bottom and children on top. The bottom-most white box represents the whole scene at its coarsest resolution (the root). The colored boxes show the *active* tiles, which are those that are actually displayed on screen. For each active tile the associated terrain data are shown, although the triangulated DEM meshes have been flattened to illustrate the differences in geometric density with respect to area, rather than the actual 3D positions.

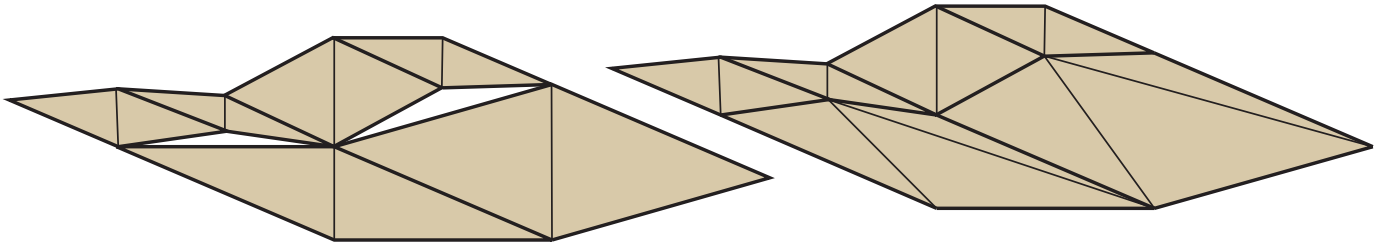

Fig. 2. Stitching between tiles. Left: Cracks can appear at the edge of neighboring tiles of different resolutions. Right: Stitching adapts tiles edges for continuous transitions.

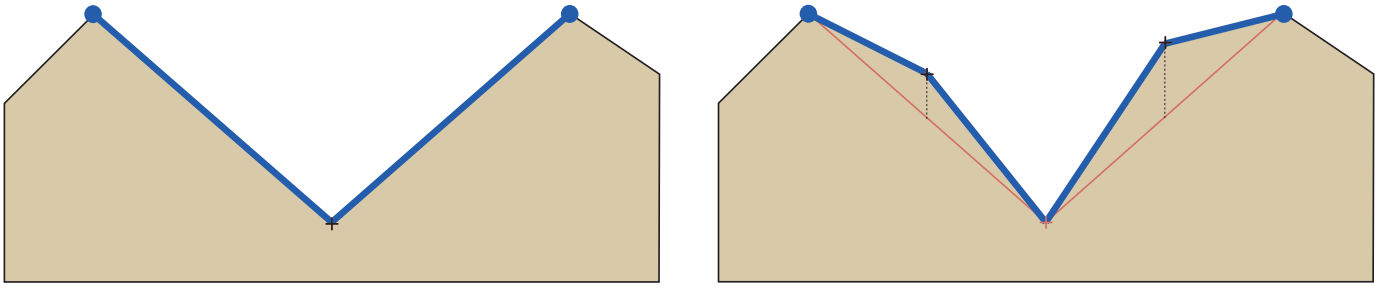

Fig. 3. Polyline representation (shown in a cross-section view). Left: Line strips connecting the segment control points follow the surface of the virtual terrain model. Right: Mesh representation has changed to show more detail. Because the old line strips (red) are hidden in the updated view, a new line strip (blue) must be generated to connect the same control points.

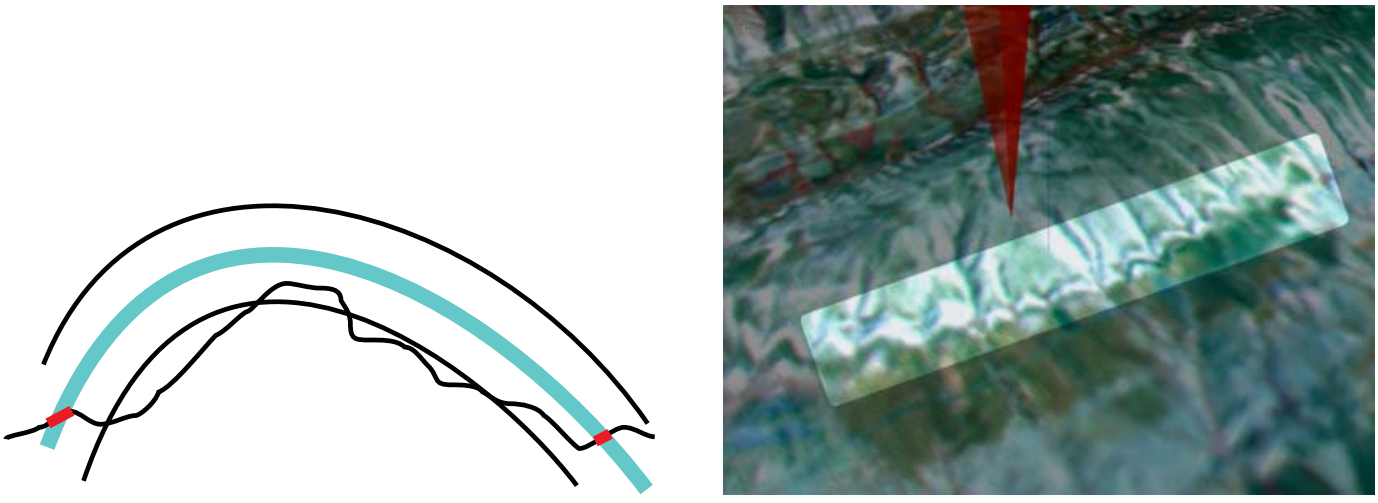

Fig. 4. Left: The folded rock layers have been eroded and their specific rock structure is apparent on the current landscape at the intersection (red highlights). Right: Dented intersection as seen in RIMS when using false-color satellite imagery.

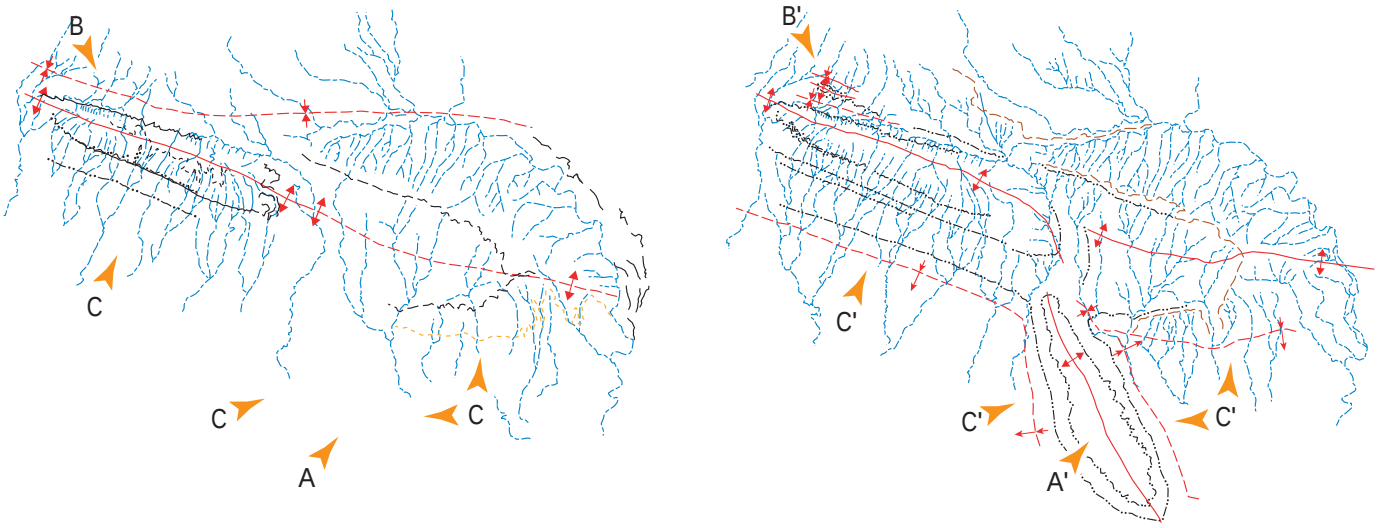

Fig. 5. Results of sensitivity test. Left: Map generated by a geologist using StereoAnalyst. Right: Map generated using RIMS. Gold arrows highlight points where the maps differ significantly. Red lines are fold hinges and are solid where confidently located and dashed where their position is less clear. Broken black lines indicate contacts between two different geologic units, dotted black lines are marker beds. Dashed yellow line denotes the edge of a geomorphic surface. Brown lines indicate drainage divides. Blue lines denote drainages. All lines were mapped by the geologist.

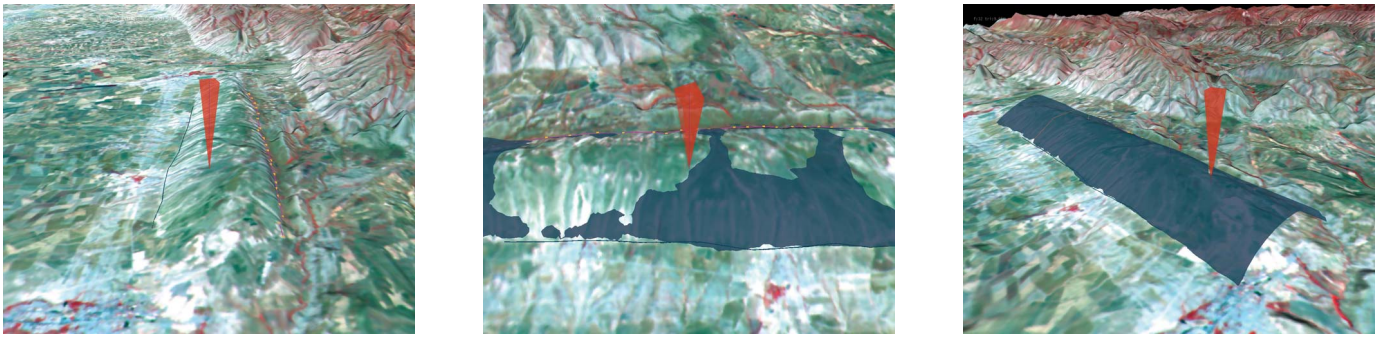

Fig. 6. Reconstructing mapped structures. Left: A distinct layer of rock intersects the surface of the earth on opposite sides of a ridge. Extrapolation of this layer indicates that prior to its removal by erosion it arched over the top of the ridge. Using RIMS, the location of the layer has been mapped on either side of the ridge, using a blue line on the left and a red line with yellow squares on the right. Center: Automatically reconstructed surface intersects topography. This solution is not acceptable because inspection of the virtual terrain model indicates no exposures of the bed on the top of ridge where the reconstructed surface intersects the surface of the earth. Right: Surface is tweaked manually until the minimum fold amplitude is achieved such that the bed projects over the crest of the ridge, consistent with geological observations.

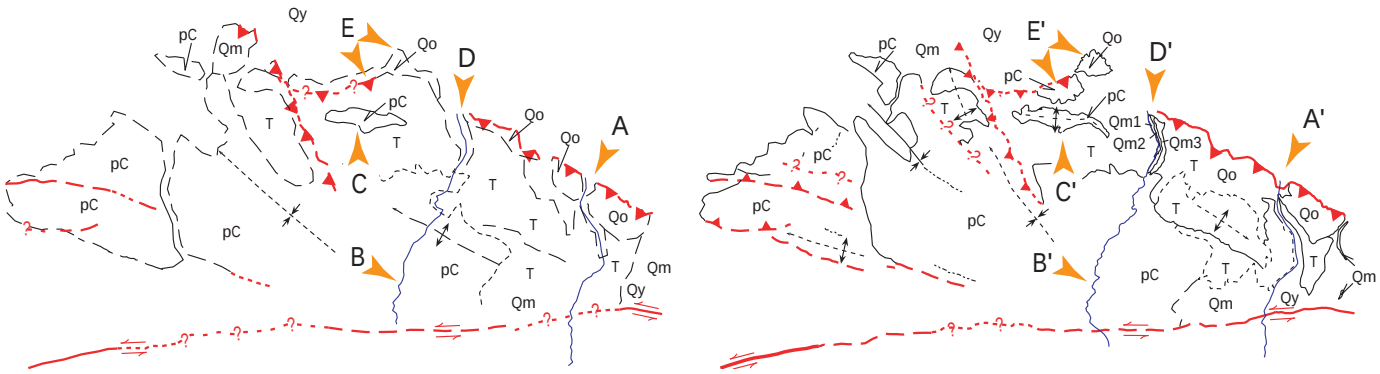

Fig. 7. Results of efficiency test. Left: Map generated by a geologist using StereoAnalyst. Right: Map generated using RIMS. Gold arrows highlight points where the maps differ significantly. Decorated red lines are various types of active faults. Decorated black lines represent folds, undecorated black lines are contacts between different geologic units. Red and black lines are solid where features are confidently located and dashed, dotted, or queried where position is increasingly less clear. Solid blue lines are drainages. Text labels (pC, T, Qo, Qm#, Qy) denote units of different apparent ages. Ages were inferred from weathering pattern and reference to existing low-resolution geologic maps. All mapping was done by a geologist.

APPENDIX

A. Rims performance

To evaluate the performance of RIMS, we simulated mapping usage on several test data sets of different sizes. Preprocessing times were generally less than a few minutes: the largest data set used for this study was an 8193×8193 pixel ($163\,\text{km} \times 163\,\text{km}$) lidar DEM and color hillshade of the Puget Sound area, which took 750 seconds to preprocess. Running on a hyper-threaded Pentium4 $3.06\,\text{GHz}$ machine with $1\,\text{GB}$ RAM and a GeForce 6800 GT ($128\,\text{MB}$ dedicated memory) frame rates never dropped below $25\,\text{fps}$ averaging at $82\,\text{fps}$ for the Puget Sound data set.

B. VGC – Test one

The first VGC test was designed to determine how the accuracy of the VGC measurements varies as a function of dip angle and topographic relief, and thus to quantify the extent to which these factors inhibit a user's ability to reproduce a plane of known orientation.

For this test, a user (Gold) first assembled a benchmark dataset by dispersing a set of 18, variably oriented VGC planes of known strike and dip across the Mosul dataset. He then precisely mapped the intersections between these VGC planes and the surface of the virtual terrain model. Six planes were placed in each of three types of topography: mountainous, moderate, and undulatory. In addition, the VGC plane dips varied within each type of topography, with two planes in each of three ranges: shallow (0-30 degrees), moderate (31-60 degrees), and steep (61-90 degrees) dips. Once the intersections between the VGC plane and the surface of the virtual terrain model were mapped, the mapped polylines were checked by extrapolating planar surfaces from them using the surface interpolation tool described above. The projected planes matched the true VGC planes to within a degree in both strike and dip. After these benchmark data were generated, the 18 mapped lines were given to a second user (Cowgill), who was not told the orientations of the VGC planes from which the lines had been derived. The second user then fit VGC planes to the mapped lines, recorded the strike and dip values reported by RIMS at onefold and

twofold vertical exaggeration, and then calculated the measurement errors. Errors were determined by finding the absolute value of the difference between the measured strike and dip values and those used to generate the test data.

Table I reports the average absolute errors and their standard deviations for the first test. Only results from the measurements made at twofold vertical exaggeration are reported because they were found to be slightly smaller. Standard geological compasses can be used to measure strike and dips to within about 0.5 degree. However, natural roughness along geological surfaces means that a given strike and dip measurement can differ from the average surface orientation by up to five degrees. Thus, it is clear from the values in Table I that the VGC tool is quite accurate, even where steeply dipping surfaces cross areas of low relief, as long as the mapped features are coplanar and clearly defined. Although the errors do show small variations as a function of the amount of topographic relief and the surface dip angle, it is clear that all values are well within the five degree variation typical of natural data.

		Strike		Dip	
Group	n	a.e.	s.d.	a.e.	s.d.
All Data	18	0.06	0.05	1.23	2.37
High Relief	6	0.06	0.05	0.27	0.26
Medium Relief	6	0.07	0.06	0.65	0.69
Low Relief	6	0.05	0.07	2.76	3.77
Shallow Dip (0-30)	6	0.02	0.02	0.98	1.00
Medium Dip (31-60)	6	0.08	0.06	2.45	3.87
Steep Dip (61-90)	6	0.08	0.06	0.26	0.30

TABLE I

AVERAGE ERROR (A.E.) AND STANDARD DEVIATION (S.D.) OF MEASUREMENTS MADE AT TWOFOLD VERTICAL EXAGGERATION.

C. VGC – Test two

The second VGC test was designed to evaluate the extent to which the VGC tool can reproduce field-based measurements. For this test, a user (Gold) scanned and georeferenced a published map, draped this

texture over a DEM with $10\,m$ horizontal postings, and then mapped the locations of 25 points where the published map reported field-based strike and dip data.

Measurements were only chosen from areas where the corresponding 5 m-resolution air photo revealed exposed layering that would likely be visible in RIMS, because a user is unlikely to use the VGC tool in areas without such exposure. The measurement sites, DEM and texture, were then passed to a second user (Cowgill) without providing the published strike and dip data. This second user then set the vertical exaggeration to fivefold and used the VGC tool to determine the orientation of the dipping layers of sediment at each site. Unlike the first test, in this case the absolute errors are significant: the mean absolute difference between the published and measured values was 12.5 ± 8.5 degrees (1 SD) for strike and 21.1 ± 12.7 degrees (1 SD) for dip. A number of explanations for this discrepancy are possible. First, errors in the registration of the texture image with the DEM will propagate into errors in the measured orientations. Second, a number of the measurement sites were in areas where the trace of bedding across the topography was poorly defined in the texture. Thus because these are sites where a RIMS user would not typically collect data, the RIMS data set is likely to be anomalously poor. Third, we are comparing data measured at two very different scales: while the published values report the orientation of a surface observed over a few meters, the VGC measurements cover much larger areas.

Although the source of the discrepancy remains unclear, this second test indicates that the VGC tool is useful as a semiquantitative utility that indicates where geologic surfaces are gently, moderately, or steeply dipping. The test also indicates that users can have more confidence in strike determinations than in dip measurements generated with the VGC plane. This result is not surprising, because the intersection lines between the topography and a geologic surface generally show more along-strike continuity than they due variability due to the dip, making it easier to for users to constrain the strike than the dip.