User Interface Paradigms for VR-based Surgical Planning:
Lessons Learned over a Decade of Research

Kevin Montgomery PhD[1], Michael Stephanides MD[1], Stephen Schendel MD DDS[1], Muriel Ross PhD[2]
[1] National Biocomputation Center, Stanford University, Stanford CA 94305
[2] University of New Mexico, Albuquerque, NM 87131

Abstract
This paper covers our work in virtual reality-based surgical planning over the past decade. It aims to comprehensively examine the user interface paradigms and system designs during that period of time and to objectively analyze their effectiveness for the task. The goal is to provide useful feedback on these interface and implementation paradigms to aid other researchers in this field.

First, specialized systems for specific clinical use were produced with a limited set of visualization tools. Later, through a collaboration with NASA, an immersive virtual environment was created to produce high-fidelity images for surgical simulation, but it underestimated the importance of collaboration. The next system, a networked, distributed virtual environment, provided immersion and collaboration, but the immersive paradigm was found to be of a disadvantage and the uniqueness of the framework unwieldy. A virtual model, workbench-style display was then created using a commercial package, but limitations of each were soon apparent. Finally, a specialized display, with an integrated visualization and simulation system is described and evaluated.

Lessons learned include: surgical planning is an abstract process unlike surgical simulation; collaboration is important, as is stereo visualization; and that high-resolution preoperative images from standard viewpoints are desirable, but interaction is truly the key to planning.

Keywords: Surgical planning, virtual environments, surgical simulation
Key Links: http://biocomp.stanford.edu

Introduction
Computer-based surgical planning has been investigated by many researchers over the past decade. The promise of the technology is to provide better surgical results (with fewer procedures), decreased time in the operating room, lower resulting risk to the patient (increased precision of technique, decreased infection risk), and a lower cost.

As described below, a great deal of work has been attempted in this area by our labs and by excellent researchers throughout the world. Despite these efforts, routine use of surgical planning is still in its infancy. This paper seeks to detail our history in this field and to analyze the paradigms, assumptions, efforts, designs, and use of our surgical planning systems in order to provide the user the insights and hard lessons learned over this decade of research.

Related Work
While a great deal of technology development has been necessary to enable this goal, it has also been the case that a number of clinical studies have demonstrated the benefits of the application of these technologies in actual patient outcomes. Among the technology groups, a few bear special mention.

The Surgical Planning Laboratory at Harvard has a long history of applications of advanced visualization in the predicting of surgical outcomes [Harvard]. Similarly, the Mayo Clinic's Virtual Reality Assisted Surgery Program (VRASP) and Biomedical Imaging Resource (BIR) also merit special distinction [Mayo].

The Center for Medical Robotics and Computer Assisted Surgery (MRCAS) at Carnegie Mellon University has focused on technology development in surgical planning, robotics, and remote collaboration. Their collaboration with Shadyside hospital has also provided a conduit for real, clinical application of these technologies [CMU].

The Graphics, Visualization and Usability Center (GVU) at Georgia Tech develops technologies in biomedical imaging and visualization, visual computing, interaction and collaboration, communication and education. In particular, their virtual environment work in surgical simulation is especially noteworthy. Their collaboration with the Medical College of Georgia also affords them a mechanism for application and exchange of clinically relevant information [GaTech].

The University of Illinois at Chicago is home to one of the best visualization groups in the world, the Electronic Visualization Lab (EVL). The application of these visualization technologies in medicine is concentrated at the Health Sciences Center in the newly formed VRMed Lab. This group focuses on the application of advanced visualization technologies in medicine [UIC].

The Human Interface Technology (HIT Lab) of the University of Washington has done outstanding work in the interfaces between humans and visualization systems. Their aim to develop a new generation of human-machine interfaces has also been applied in a number of medical applications [HITL].

Military applications of these technologies have been developed over a number of years. The United States Air Force Institute of Technology's Computer Graphics Laboratory has focused on the design and implementation of interactive environments, 3D medical imaging, and advanced visualization with integrated simulation for medical uses in the US military [USAF].

In Europe, the Fraunhofer Institute's Center for Research in Computer Graphics (CRCG) has developed technologies in the areas of volume visualization, virtual environments, collaborative work tools, and user interface design. The Institut Graphische Datenverarbeitung (IGD) has focused on: anatomy education, surgical training, preoperative planning, intraoperative support and prototyping in projects for virtual anatomy, virtual endoscopy, VR-based training, preoperative planning, and intraoperative support [Fraunhofer].

Also in Germany, the Friedrich-Alexander-Universitat in Erlangen-Nurnberg, has developed some of the best and most clinically relevant work has taken place in the area of craniofacial reconstructive surgery planning and simulation [Erlangen]. Finally, the University of Mannheim's Virtual Reality in Medicine (VIRIM) project has also been very active in application of advanced visualization techniques in medicine [Mannheim]

In England, the University of Sheffield's Virtual Reality in Medicine and Biology Group (VRMBG) also focuses on medical applications of virtual reality, including biomolecular visualization, surgical simulation (including an arthroscopic knee surgery simulator), and blood flow modeling [Sheffield]

The Technical University of Denmark (DTU) 3D-Med project works on the design and analysis of a medical imaging application. Their Hybris system uses both volume rendering and surface to display isosurfaces generated from medical 3D volume scans [DTU]

In Asia, the Institute of Systems Science (ISS) in Singapore hosts the Center for Information-Enhanced Medicine (CieMed) group. This group also has done a great deal of impressive work in the areas of visualization, simulation, and artificial intelligence, as applied in medicine [ISS].

Also in Asia, the National Cancer Center of Japan investigates medical applications of virtual reality including surgical simulation and the virtual medical communication project, including medical VRML [NCCJ].

While this provides a very short list of the excellent, technically-focused groups in this field, the clinical side has seen an even greater number of studies and real application. Clearly, the application of computer-based cephalometric measurement and surgical planning is of great interest to surgeons.

Vannier, Marsh and others at the Mallinkrodt Institute in the Washington University in St Louis have done some of the most far-reaching and consistent work in this field. In addition, the orthopedic surgery group [Sutherland94] has also done planning in their area.

Surgeons at the University of Texas and Baylor University [Byrd92] have worked on rhinoplasty planning, Johns Hopkins has done orthopedic planning [Chao93], and the excellent orthopedic group at UCLA [Gautsch93] continue to work in this area as well.

Internationally, clinical and engineering groups at the University of Heidelburg [Haasfeld98, Pokrandt96] have done exceptional work in applications of computing and visualization in surgical planning. In addition, the University hospitals at Leuven Belgium [Lamoral90,Verstreken96] have done good work in quantification for maxillar implants.

Researchers at the University College in London [Moss88, Perry98, Ayoub96] have done work in maxillofacial surgery prediction and craniofacial reconstruction. Maxillofacial implant optimization has been researched by the University Hospital Groningen, Netherlands [Rozema92] and other craniofacial visualization has been done at Utrecht University Hospital [Zonneveld89].

While there are so many technical and clinical groups around the world working in this area, it is hoped that this number will pale in comparison to those beginning to enter this field. Toward that end, we wished to catalog, checkpoint, and disseminate a snapshot of our work over the past decade to hopefully aid future researchers in their endeavors. In particular, we wished to offer the often hard-learned lessons of surgical planning user interface design that we have accumulated throughout our research.

The work described in this paper covers nearly a decade of research at the Stanford University Division of Plastic and Reconstructive Surgery and the NASA Ames Research Center. Early work at Stanford in this area demonstrated the great clinical utility of advanced visualization. At this time, the Biocomputation Center at the NASA Ames Research Center was developing 3D reconstruction and visualization technologies for studying the effects of gravity upon biological systems. The demonstrated benefits and strong desire to accelerate the medical application of these technologies led to a strong and fruitful collaboration between the labs beginning in 1993. On the basis of this collaboration, NASA and Stanford formed the National Biocomputation Center in 1998 to focus the development and applications of these technologies in medicine and surgery.

The focus of the National Biocomputation Center has been to lead a national effort in developing and applying technologies in medicine and surgery. The idea was that these technologies would not only revolutionize medicine and surgery on Earth, but would enable health care and medicine for long duration human spaceflight. Specifically, our emphasis has been on surgical planning, intraoperative assistance, and training/education (including surgical simulators). The surgical planning work detailed in this paper reflects our focus on craniofacial reconstructive surgery, due to our clinical interest and the great technical challenges of this type of surgery. However, these technologies have been extended to other surgical areas and applied in surgical simulation for training and education as well.

Throughout our development, we continually applied our technologies on real patients. Our desire was to ensure that our engineering and research efforts were guided along the paths necessary for productive clinical use, but also to enjoy the benefits and feedback from early clinical application. At the time of this writing, the systems described have been applied on the data of over 40 patients during development. As our work has progressed, our adoption and use of these technologies has accelerated and now are using the system on at least one patient every two weeks. Further efforts to automate portions of the process are continuing and full deployment of the system throughout a number of medical departments should be accomplished this year.

This paper details the progression of our work in computer-based surgical planning, with a special emphasis on user interface and design issues. It is hoped that the lessons learned over this time will aid other researchers in understanding the issues in practical surgical planning systems. Over time, we have developed essentially five different systems for surgical planning and visualization, which are described below. For each we will discuss the following:

- Data (how the data were obtained)
- Platform (what the system was built around and it's specifications)
- Displays (how data was presented to user)
- Interface (hardware to enable interaction)
- Interaction (how the user interacts with system)
- Software design (how the system fits together)
- Manipulation/Operation (how the user performs task of surgical planning)
- Examples of patients and outcomes
- Analysis (what worked, what didn't and why)
- Lessons (what we learned)

Initial Work: 1990

Our earliest work in this area began in 1990, with Dr. Stephanides' early attempts at producing a semiautomated 3D cephalometric system using SuperCard on an Apple MacIntosh.

In this monolithic (one program), specialized system, conventional radiographic cephalograms were digitized using a RasterOps scanner. The data for 40 patients were scanned and these high resolution images were presented to the user on a simple CRT. The scanned image data as presented showed much lower resolution than the original cephalograms, but provided adequate resolution for analysis.

The computer mouse was used as the primary method of input and the user would be instructed to pick the standard cephalometric landmarks on a given film. Then the software would compare the location of these landmarks with that contained in a normal anatomy database and analyze the deviation from the normal locations. In addition, it was used for both pre- and post-operative comparison studies in order to relate the long-term effects of surgical interventions.

While this early system was an effective (if limited) tool to study and provide care to a limited and specific patient population, future clinical research areas required more specialized means of data acquisition, presentation, and analysis.

A second system created in 1991 was to study the effects of corticotrophin releasing factor (CRF) on preventing postoperative swelling in eyelid surgery. A total of 40 blepharoplasty patients were scanned with a Cyberware 3030 3D laser-striping digitizer attached to a Silicon Graphics (SGI) Indigo Elan workstation. The Cyberware device creates a 3D mesh representation of the face and, together with a stereotactic head bar, can be used to scan a patient before and after an experimental condition. In this case, half of the patient population was provided CRF and the other half a placebo. Patients were scanned before and after surgery and each individual's pre- and post-operative scans were subtracted to obtain the volumetric changes around the eyes.

In interaction with this monolithic, specialized system, the user moved the mouse to visualize the patients' data, as well as the subtracted volumetric area. This allowed for user validation of the operation of the system and began to show the importance of 3D visualization (even if not stereo) in the understanding and analysis of patient anatomy.

The advantages of each of these systems were apparent- they each did one highly specialized task well and used appropriate (if primitive) forms of visualization to accomplish their goal. A further advantage was that these systems were created with a very small amount of funds and used rather inexpensive hardware.

The difficulties of these systems also start to become apparent. Clearly, writing a separate system for each form of cephalometric analysis would start to become unworkable at some point. In addition, more advanced forms of visualization could provide some benefits in understanding the data- the patient's anatomy is 3D, there may be benefit in seeing it in 3D. Finally, the ability to interact with the patient's data was also very limited by these early programs. These drawbacks, together with a desire to accomplish more work than would be possible with the limited resources at his disposal, led Dr Stephanides to seek a collaborator that could aid his work in surgical planning and analysis.

First Steps (VERS): 1994

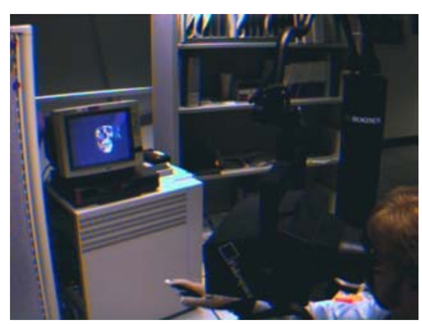

Our first collaborative work began in 1994 when Dr Stephanides approached the NASA group to inquire whether any tools existed to aid him in his task of planning complex craniofacial surgeries. Up to this point, the NASA group had been working on 3D reconstruction and visualization of biological tissue for space-related research. Since the techniques were judged to be similar and the benefit great, we began working together to apply our existing technologies to problems in medicine and surgery.

Data:
In this first system, we had envisioned using CT data to derive a computer model of the bone and soft tissues, then also acquiring a higher resolution model of the skin surface using a Cyberware 3030 laser striping digitizer. Our prediction was that the models that we could derive from the CT data would be of insufficient resolution for surgical planning purposes. Therefore, we would derive the low resolution bone and skin surface models from CT, then register the high resolution Cyberware image to the low-resolution CT skin image and thereby yield a dataset with bone at CT resolution and skin surface at the resolution of the Cyberware (approximately 0.5 mm).

As noted below, our first clinical application of this software was with a patient with limited compliance. We therefore found that working with the CT data alone was necessary. The first issue requiring a surprisingly enormous amount of effort was to obtain information on the proprietary encoding of the CT data. In the world before DICOM, each manufacturer had their own format for storing CT data and often guarded it as a trade secret. We were able to obtain a description of an older version of the format and used this, together with a great deal of trial and error to derive the format that the manufacturer had produced. Once the data were decoded, we converted the imaging data into TIFF image files for subsequent reconstruction.

We had developed a software package named Reconstruction Of Serial Sections (ROSS) to aid in our existing work in space-related neurobiological research. This package allowed the user to display an image and then trace out the contours (outlines) of the objects of interest in that image. Further, once this was accomplished for each slice of the data in question, the software could automatically register (align) the data, as well as produce a mesh for computer graphics-based visualization. This package was used on the TIFF images from the CT data of the patient. As would be expected, serial-section tracing of bone outlines is a time consuming, labor-intensive process. However, the process was completed and a mesh produced that, while somewhat limited in aesthetics, nonetheless was found to be clinically useful and promising as a clinical tool.

One issue we found during this process was that our system produced an extremely dense mesh of hundreds of thousands of polygons. While this level of detail was initially rewarding, the difficulty of visualizing such a model were soon apparent and work on mesh reduction (decimation) was begun. Initial work in this area focused on polygon merging by a chord-height (planarity metric). We therefore could reduce the size of the data to that which could be used interactively on existing hardware.

Finally, we found that our ability to extract useful and highly accurate CT-based models of bone and soft tissue removed the necessity for the separate Cyberware scanning step. Furthermore, the automated 3D registration of the CT and Cyberware-based skin surfaces would require a great deal of processing time and still yield a system with possibly additional registration errors. For these reasons, Cyberware scanning of the patients was discontinued.

Platform:
This work was initially begun on lower-end Silicon Graphics (SGI) computers, namely the Indigo and Indy personal workstations. As the work progressed toward immersive visualization, a SGI Onyx RealityEngine2 superworkstation was purchased to aid in this and other projects. The added graphics performance of the RE2 (2M polygons/second) allowed for interactive visualization even of models up to 100,000 polygons at acceptible frame rates.

Display:
The interactive display device used in this work was the BOOM3C from FakeSpace Inc. This device is CRT-based, counterweighted, armature-style immersive display. This device provided the highest resolution and best dynamic range available and supported full stereo display. Our thoughts were that image fidelity was of the utmost importance in surgical planning, so a high-end display was a critical need. In addition, the tracking method of the device (optical encoders on the joints of the counterweighted armature) provided very precise tracking of user position and orientation.

Interface:
The interface to the system consisted of a Virtual Technologies CyberGlove hand input device, with attached tracker. The glove was an 18-element electroresistive sensing glove that could report the flexion and abduction of each of the fingers, along with some wrist motion. The paradigm we were exploring was gesturally-based input.

The tracking device we employed was a Flock of Birds electromagnetic tracking device manufactured by Ascension Technologies. This device was selected after a literature and user survey that found that the pulsed DC method of tracking used by Ascension was more robust to metallic interference in the area. In addition, we used their Extended Range Transmitter (ERT), which is a 1 ft^3 transmitting coil also in hopes of overcoming (or overwhelming) any metallic interference induced by the metal of the BOOM display.

Interaction:
The user interacted with the system by producing hand gestures to indicate the desired operation. For example, the user could make the "drill" gesture and their virtual hand in the environment would turn into a virtual drill and could be applied to the patient data (skull) in the environment. Other gestures such as the "grab" gesture could be used to rotate and translate the patient data in the environment.

Software Architecture:
The software that brought the display and interface components together to allow for interaction, was a monolithic, home-grown system known as the Virtual Environment for Reconstructive Surgery (VERS). In this early program, we wrote all device driver/interface code and incorporated it into a large interface-render ("handle the sensors, then draw the scene") loop.

Example Patients:
This system was used to plan the surgery of a man in his late 50s who had been assaulted with a baseball bat to the forehead. The reconstruction, using the serial-section method was complicated by the presence of many small and irregularly shaped bone fragments. The resulting reconstruction, while primitive, showed that these technologies had promise for future work.

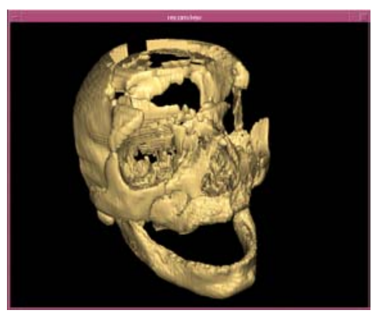

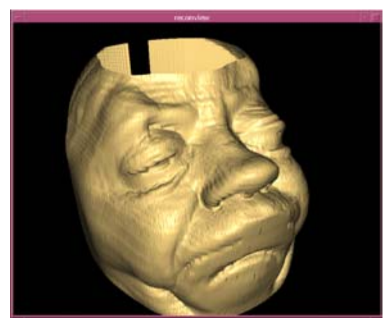

Analysis:
In analyzing the drawbacks of this system, a number of issues can be traced either to assumptions that were invalidated, or to naivety on the part of the prevailing interaction paradigm that we were operating under.

First, while we originally thought that the CT-based models would be insufficient to provide adequate resolution and detail, we found that the resulting reconstructions provided overwhelming detail (submillimeter), requiring reduction in the size of the mesh in order to provide reasonable visualization. However, our efforts in mesh reduction using a simple chord-height (planarity) metric found that the resulting reconstructions demonstrated large facets in the regions of the face where the skin surface is relatively planar. Clearly some tradeoff should be made to allow for decimation in areas of low curvature, but some additional metric to prevent over simplification was necessary to produce a visually acceptable result.

Second, when we applied the above method for mesh reduction to the point where the model started to lose visual acceptability, we found that the resulting model still contained nearly 400,000 polygons and was too large to enable real-time interaction on existing hardware. While the theoretical peak graphics performance of the RE2 was nearly 2 million polygons per second, we found that we could only obtain ~3 frames per second with a reduced model, for a total of 1.2 million polygons per second. This slow frame rate limited the system's utility. Either the model was of such low fidelity to be fast but useless, or the model had sufficient detail and was too slow to be usable.

In other areas of assessment, the BOOM was found to provide exceptional resolution, great dynamic range, high contrast due to the immersive aspect of the display (including a rubber face-coupling ring that eliminated stray light onto the display), and precise tracking. The disadvantage of the display is primarily its high cost ($95,000).

The electromagnetic trackers derived no noticeable interference from the metal of the BOOM, unless the tracker was near the back of the BOOM display where the CRT deflection coils produce a field. However, the Extended Range Transmitter (ERT) of the tracker did produce such a strong electromagnetic field that it effected not only the BOOM display, but also any other display within a 6 foot radius. While these effects could be minimized by orienting the user to have their head further from the transmitting cube and their hand (with tracker) closer to the cube (the Ascension trackers vary the strength of the electromagnetic field based on distance of the receiver to the transmitter).

As far as operation of the system, we learned a number of valuable lessons. First, the software architecture of a monolithic, single-threaded interface-render loop meant that the system was often blocking, waiting on either the BOOM or tracker to report the latest position information before continuing.

Second, we found that the gestural interface that was predominant in VR-based interaction at the time was counterintuitive to the surgeons. The fact that the surgeon was required to learn each gesture and its meaning was found to decrease enthusiasm and increase the artificial feel of the system. In short, gestural interfaces were found to be counterintuitive.

Third, while we initially felt that stereo visualization of the models was important but not critical, we found that stereo cues are essential in surgical planning. In the case presented above, the ability to visually separate and localize the separate bone fragments of the skull was extremely important for the surgeon to understand the complex anatomy of the defect.

Forth, we had originally envisioned that surgical planning is much as we envisioned an operating room- with one surgeon directing the operation, requesting information, making decisions, and directing actions. What we found was that surgical planning is not a 1-person task, but the ability to collaborate with other surgeons on a surgical plan is very important.

Fifth, we found that replicating the real world is a difficult task. We sought to produce an environment that was identical to a surgical suite, both visually, as well as in the methods of performing the task. Simulating the real world is a very high standard to satisfy.

Finally, we found that surgical planning is not surgical simulation. For example, our system allowed the surgeon to go through the act of cutting the bones of the skull (simulation). However, the surgeon really wanted to define where a cut would be made as simply as possible (planning). This distinction of the act of making a cut versus defining where a cut is made is subtle but important. If using a virtual environment is the same steps/difficulty/time as doing the real task, then the system is of limited use. A surgical planning system should instead allow the surgeon to specify the cognitive primitives of planning a surgery as quickly and efficiently as possible, without having to require them to be bound to the methods imposed by the real world.

For these reasons, we embarked upon a complete rewrite of the system that would allow for immersion, collaboration, faster rendering, and a more intuitive method of interaction.

Second Steps (Networked VERS): 1996

In this system [Montgomery96a-b] we sought to still provide an immersive environment, but to also provide a framework for collaboration. While we still felt that immersion was important in order to minimize distractions and allow fusion of the user with the data/task, we wished to determine how important full immersion was in the process. In addition, as our previous system had taught us, collaboration is important, so we wished to provide this capability via supporting multiple immersed users in a distributed, shared virtual environment.

Data:
In this complete redevelopment of the system, we produced fully volumetric methods of segmentation, mesh generation, and artifact removal. The desire was to get around the manual labor and imprecision inherent in serial-section reconstruction if this system was ever to gain wide acceptance.

An overview of the process follows. The surgeon orders a CT scan of the patient for preoperative evaluation. The patient's dataset is provided to our center by magnetic tape or over the Internet. Once obtained, the vendor-specific (proprietary) data format is first converted to a standard format (DICOM). Next, our software is used to visualize the raw data, segment the bone and soft tissues, and generate a triangular mesh over those surfaces.

Due to the drawbacks noted previously with the planarity metric for mesh reduction, this time we used triangle area as the metric for deletion. Our rationale was that the smallest triangles in the mesh are not contributing much to the visual display and, in some cases, are subpixel. Therefore, elimination based upon triangle area seemed a reasonable approach.

Platform:
As above, a SGI Onyx (2x150MHz R4400s, 512MB memory, 38GB disk, RE2 graphics) was used, but other devices (such as individual SGI workstations) were also employed as networked clients as well.

Display:
We again primarily used the FakeSpace BOOM 3C device, but due to the distributed nature of the design, we could also have other users on more traditional display devices, such as a workstation with CrystalEyes stereo eyeware (Stereographics Corp).

Interface
We also again used the Virtual Technologies CyberGlove with Ascension Flock-of-Birds trackers, but the flexibility of the system made it easy to add other devices- just a single device driver process was needed.

Interaction:
In this environment, we steered away from gestural interfaces. Instead, we decided that a simple grasp-release set of gestures was certainly intuitive, so we developed a “virtual toolbox” approach to providing virtual tools. The virtual toolbox existed within the environment (much like the tray of surgical instruments exists in an operating room) and the user can reach over, grasp a virtual tool, then apply it to the patient data. This also provided great flexibility and expandability in virtual tool design. Using the same grasp-release paradigm, the user could also grab the patient data (skull) to rotate it for visualization purposes.

Software Architecture:

The architecture consists of a number of modules, which send messages to graphical objects stored in shared memory. Each graphical object contains the mesh representing the object, position (location and rotation) of the object, and other attributes necessary for describing and rendering the object. All objects in the virtual environment are represented as graphical objects. Examples of these objects include the patient's skull, virtual tools, the user's hand, and even the user themselves. A directory of graphical objects also exists for indexing purposes.

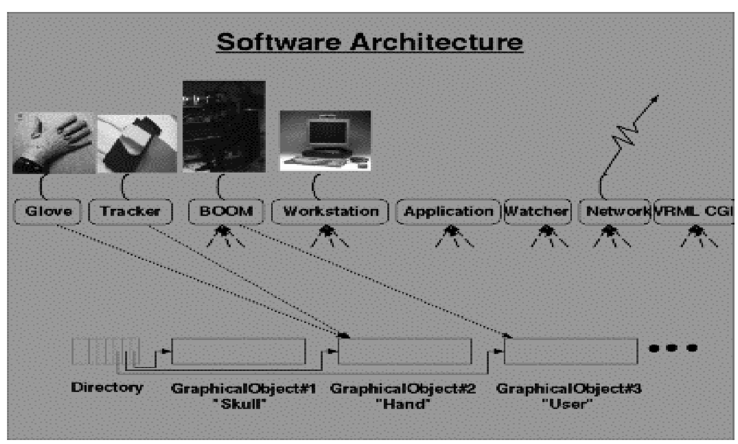

Each module operates as a separate Unix process. In this way, rendering modules can be decoupled from I/O modules for optimum performance and scalability. Modules include:

FakeSpace BOOM 3C Display and I/O module: This module communicates with the FakeSpace BOOM 3C to obtain the current user position, reads the graphical objects stored in shared memory, and renders them from the user's perspective. The BOOM buttons and joystick are used for modifying the nature of the display (such as displaying the objects in wireframe or solid objects, turning on and off objects, etc). Besides these duties, this module also updates the Graphical Object representing the BOOM user with the new position information.	*Ascension Technologies Flock-of-Birds tracker I/O module:* This module communicates with the Ascension Flock-of-Birds tracking device over an RS232 serial line (each bird has a separate serial line for minimum latency). When this module is started, it is given the name of a Graphical Object to which it should be associated. It then will continually update that object's position information. Typically the receiver is mounted on the CyberGlove and it is used for updating the user's hand location. However, by keeping the interface general, it could also be used with a head-mounted display or any other device without modification.
Virtual Technologies CyberGlove I/O module: This module communicates with the CyberGlove device over an RS232 serial line and updates the information in its associated Graphical Object.	*Workstation Module Display and I/O module:* Since expensive hardware is not always available to all users, a module was created that allows display to a standard SGI workstation. As in the other display modules, it reads the Graphical Objects in shared memory and renders them. The user's location and orientation may be updated by moving the mouse and/or keyboard controls.
Application Module: While the various hardware I/O modules update the information contained in a single Graphical Object, it is necessary to simulate interactions between objects. For example, when the hand Graphical Object comes in contact with a virtual tool, the virtual tool's position needs to be tied to that of the hand. A more complex example would be when a virtual drill comes in contact with the skull. These inter-object interactions are handled by the Application module.	*Network Module:* The network module keeps the shared memory area in one machine synchronized with that of another machine (cache coherency). The network module, running on the primary machine, communicates over the Internet (using TCP/IP) with a network module running on a remote machine. When communication is first established, the primary host sends all the information about Graphical Objects to the remote machine. Thereafter, it sends only those object attributes which have changed (as indicated by revision numbers tied to each attribute group) to the remote host. In this way, an extremely terse (and low-latency) protocol can be created (for example, the position information for 30 objects can fit into a single ethernet packet).
VRML CGI Module: A module was written to read the Graphical Objects contained in shared memory and create a Virtual Reality Markup Language (VRML) "snapshot" of the current state of the environment. This enables a user to view the current state of the environment using standard World-Wide Web browsing tools. While the later VRML standards support some object motion, real-time, server-controlled updates of object attributes are not yet supported to support a dynamic VRML environment.	*Web Image Module:* A process was written that performs a screen capture from the given screen device (such as the BOOM) and saves the image in a JPEG file on our World-Wide Web server. When combined with a script employing the "Refresh" Meta HTML command, a near-time, continuously updating image of the virtual environment can be viewed from a web page.
Watcher Module: Finally, a debugging module was created. This module reads all the Graphical Objects in shared memory and prints current object attributes on a standard text window.	

Patients:

In each case, the patient's CT data was transferred either by tape or over the Internet, converted from the manufacturer's proprietary format to DICOM, segmentation was performed, followed by mesh generation, artifact removal, and mesh reduction.

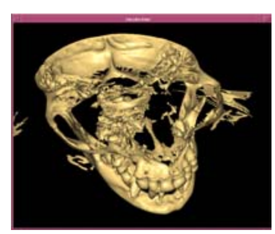

One patient was a 38 year-old male with a mandibular tumor. In this case, the surgeon needed to understand the extent of the defect in the mandible to make the clinical assessment of whether a complete mandible reconstruction was necessary, or if a partial reconstruction could be performed. Visualization for understanding the extent of damage due to infiltration of cancers into tissue demonstrates an important and useful feature of 3D reconstruction and visualization.

A more frequently occuring and notable case was that of a 6 month-old infant girl with craniosynostosis. In this condition, the sutures of the skull fuse prematurely, leading to deformation, increased cranial pressure and, if left uncorrected, can lead to mental retardation. In this case, we reconstructed the skull and soft tissue of the patient and allowed the surgeon to visualize the extent of the defect. Visualization using this software was found to be extremely useful in the understanding of such a three-dimensional defect and provided tremendous benefits over the existing paradigm of serial cross-section image presentation of CT data.

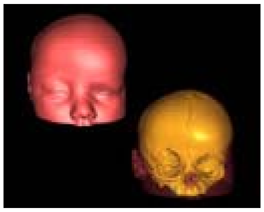
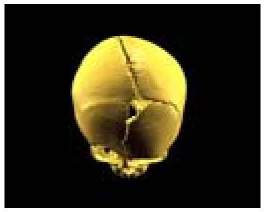
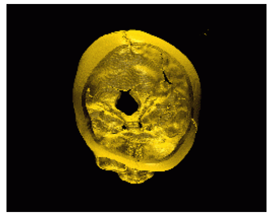

Analysis
As before, a number of lessons were learned in the application and use of this system.

First, in the data processing, we found that our segmentation software worked well and was largely automated, representing a vast improvement over the previous system. However, we again found we had a great deal of resolution (on the order of the CT data itself- 0.3mm in XY, 1mm in Z) and needed to reduce the mesh for visualization purposes.

The triangle area metric initially appeared to provide superior results but, in the case of the craniosynostosis patient, we found a disturbing artifact caused by this method of mesh reduction. Because this method eliminated very small triangles, it would tend to erode away the sutures in the skull, since the high curvature of these features is often captured by small triangles. So, while reduction to some point was beneficial, further reduction would degrade the fidelity of the reconstruction in ways that were clinically relevant. Therefore, we again found the same issue of graphical performance limitations (now 300,000 polygons, but still only 5 frames per second).

The software architecture provided several advantages. First, due to its modular nature, one could write a module specific to a particular piece of hardware without needing to learn a large software toolkit or worrying about how other devices are interfaced.

Second, the system was very flexible. One example of this was noted in the description of the tracker module above- a module can be bound to a particular object very easily. Modules in the system may come and go, may be running on any machine and provide access to many kinds of hardware. Even the behaviors of objects (as governed by the Application module) could be changed on the fly.

Third, the system yielded high graphics performance and low latency. Display modules may process graphics as fast as possible, without having to wait on any I/O. I/O devices can update their associated graphical objects as quickly as the devices allow. In addition, the terseness of the network module protocol also provided a very low-latency mechanism for inter-machine updates.

Forth, the system was collaborative. A researcher in one location may be in the same virtual environment space as one in another location, perhaps thousands of miles away or merely the next console.

Fifth, the system was interactive. It supported many types of interaction devices, all of which could be active simultaneously and in use by many users.

Finally, the system was distributed. Through the use of the network module, processing or I/O modules could exist on other machines to share the load when the tradeoff of network latency versus processing power is justified.

The shortcomings of this system should also be addressed. First, the system did not use a standard VR library (such as Sense8's WorldToolKit) and, as such, must perform all functions itself. While a library could be used in a particular module quite well, this has not yet been employed. Instead, the module developer had to interface to the proprietary code of each vendor, which is a difficult task. In short, while this home-grown system enjoyed many advantages, one spent a great deal of time writing device drivers, rather than focusing on the application.

Second, this system, with multiple processes all running as quickly as possible, was well tuned to a multiprocessor configuration, but may find difficulty in scaling down to a lower-end single processor machine.

Third, while the Application module concept provided a great deal of flexibility, it also represented a bottleneck of processing. Currently, if only one Application module is running and multiple remote users are on the system, there may be some lag in object interaction due to networking delays (e.g., a grasped object may not smoothly be moved by the hand, but would "jump" as the application module updates it's position more slowly over the networking link). Perhaps interactions between objects should be handled by the objects themselves, instilling each Graphical Object with its own innate behaviors. This is an area for further research.

As far as the interaction of the user with the system, we found that the virtual toolbox approach had some drawbacks. First, the requirement that the user had to incur a separate cognitive and motor task to change tools (i.e. dropping the first tool , finding the desired one, and picking it up), slowed progress. When the user wants a tool, they should have immediate, direct access to that tool, not to have to hunt for it.

Also, we found that a virtual tool-based paradigm was good for surgical simulation/training, but not necessarily for surgical planning. Sometimes, for the operation the surgeon wished to perform, there exists no real-world correlate, but merely an abstract concept. For example, if the surgeon wished to measure the volume of an orbit, what tool would that be? Again, issues of learning the environment arose.

Finally, we found that person-to-person collaboration in a shared virtual environment by way of computer graphics-based user representations (avatars) was not nearly as productive as face-to-face collaboration. The desire to have this kind of personal interaction, but working on a shared patient became more and more important.

Therefore, while the system was technically appealing, it was still limited in its applicability toward the domain of surgical planning. However, the insights gained led us forward to the next paradigm.

Third Instantiation (VMD VERS): 1997

In this third system [Montgomery97], we had learned that collaboration was more important than immersion. The surgeons wanted to work together on a shared dataset and to be able to interact with each other naturally.

Data:
In this system, we refined our previous segmentation and mesh generation software to improve speed, accuracy, and usability. In addition, we used a hybrid metric for mesh reduction or planarity/area. This metric afforded a nice tradeoff between the shortcomings of both other methods.

Platform:
For this project, we upgraded the previous SGI Onyx to an InfiniteReality graphics subsystem, capable of a theoretical peak of 10 million polygons per second. Our hope was that the additional graphical performance would enable us to finally have models of reasonable detail and complexity, while maintaining reasonable frame rates.

Display:
From our experiences in the previous system, it became clear that collaboration was most important, over immersion and perhaps even over image clarity. Said another way, it appeared more important for the surgeons to collaborate and discuss the case than for the image to be perfect. For these reasons we decided to move from the immersive BOOM display to a Virtual Model Display (VMD) system.

At this time, VMD projection displays were just beginning to gain interest and we had the good fortune to obtain the first unit of a new projection-based VMD from FakeSpace called the Immersive WorkBench. This system employs a large, high quality Electrohome CRT-based video projector, reflecting onto a mirror and onto a table-top display. The resulting display is 6 feet across by 8 feet wide. The table top can be moved to as much as 30 degrees above horizontal using electric actuators built into the table. The projector supports scan rates high enough to provide stereoscopic visualization using CrystalEyes stereo eyeware, so the system essentially provides a large, multiperson, stereo, collaborative display in which an object appears to project upwards from the table.

Interface:
For this system, we wanted to use a stylus-based, tracked device as a multitool- a tool that can become any tool directly. We found that Polhemus had several patents covering stylus design and such devices were therefore not available from other vendors. Therefore, we selected the Polhemus stylus as the interaction device for this system.

The one problem of having a multitool however, is how to indicate modality. In this case, since the users were no longer immersed, we felt even having a keyboard on the workbech upon which the user could select the modality of the tool could be a reasonable interface.

Interaction:
In this environment, the user interacted with the patient data by using the keyboard to indicate which tool they wished to use, then use the stylus to apply that tool to the object in question. The button on the stylus was used to indicate activation of the tool.

For example, the movement tool could be selected by pressing ‘M’ on the keyboard. Then the user could move the tool near an object in the environment and press the button on the tool to activate the tool, which would grab the object. Similarly, a cutting tool was selected by hitting ‘C’, then moving the tool to the object and pressing the stylus button to activate the tool to being cutting.

These tools were 3D objects that were linked to the position of the stylus and included even abstract tools such as the measuring tool or marking tool.

Software Architecture:
In order to accelerate development, we chose to use the WorldToolKit product from Sense8. This toolkit is widely used within the VR community, provided a wealth of device drivers, and provided a framework for rapid application prototyping and creation.

Operation:
While previous systems were largely used for preoperative visualization, this was our first system where interaction with the patient’s data was used to directly affect patient outcome. The system allowed for the previous functionality of producing rough cuts, measurement of features and other forms of interaction, but one new feature was driven by clinical need- reflection. In this case, when a patient has a defect on one side of the body, the other, normal side can be reflected over to derive the piece of tissue (bone or soft tissue) that is missing. We found this abstract tool to be very useful and widely applicable in the reconstructive surgeries predominant in our research.

Patients:
This system was used on a number of patients, three of particular interest are detailed below.

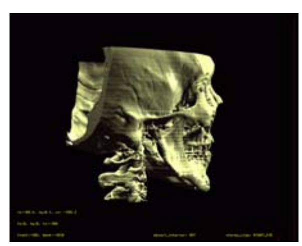

First, a collaborator from the Santa Clara Valley Medical Center, Dr. Leslie Hovey, was interested in examining the data of a mandible reconstruction patient. His individual had been in an automobile accident nearly a year earlier and had fractured his mandible in the front. Now, he presented with temporomandiblular (TMJ) joint pain and Dr. Hovey’s hypothesis was that perhaps when the mandible healed, it had changed the angle of the condyles with respect to the TMJ and thereby created uneven wear upon this joint. Therefore, he wished to determine whether the angle of the condyles matched that of a normal patient. In order to perform this task, we reconstructed the data of his patient and brought it up within our environment with the skull facing away and a near clipping plane eliminating the skull and allowing us to view the angle of the condyles as they project up to the TMJ. We similarly brought up the data of another patient that had an unrelated defect from our database. In comparing the angles of the patient with that of our “normal” (or at least not similarly affected) patient, we found a difference of nearly 10 degrees in the projection of the condyle into the TMJ. This visualization and quantification result led Dr

Hovey to make the clinical assessment that his hypothesis was correct and correction of the angle necessary.

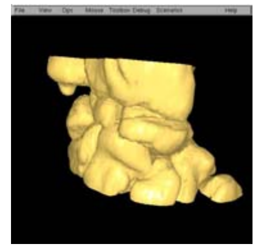

Another example case involving clinical diagnosis using our system was not a craniofacial case at all, but involved the bones of the wrist. Drs Hentz and Bowen from the Stanford Division of Hand Surgery wished to examine the data of a patient with a suspected scapholunate fracture. While the normal method of diagnosis involves measuring angles derived from X-ray films of the wrist, the diagnosis is often complex because the bones of the wrist are in such close proximity and occlude each other from the parallel projection offered by X-rays. Instead, a CT scan was performed of the patient's wrist and their data brought up in our environment. From the computer model, it was immediately obvious where the fracture was located, and the clinicians were very impressed with the clarity of the model and how simple it made diagnosis. As far as interaction with the system, the clinicians were operating the system by themselves as we observed their interaction. For nearly 10 minutes, we watched as they examined, moved and rotated, and discussed the data of their patient and the subsequent plan for the procedure. What was most impressive was that during this time, they appeared completely unaware that they were using a computer- they were truly integrated with the environment and able to concentrate on and discuss their patient without the medium (the virtual environment) getting in the way of the message (the patient's diagnosis). This level of interaction with the system is clearly where virtual environment surgical planning aspires to achieve.

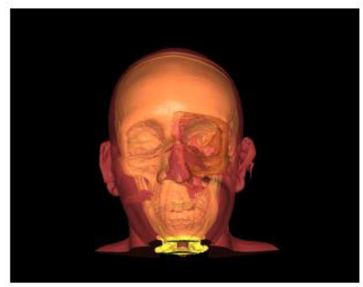

While the previous two cases presented involved diagnosis and advanced visualization of patient data, the final case demonstrates the power of real manipulation and interaction with patient data. A patient of Dr Stephanides was a 17 year-old boy that had a tumor resection at the age of 9 that involved the removal of most of the bone and soft tissue on the left side of his face. Now Dr Stephanides wished to reconstruct the boy's face using donor bone from the iliac crest of the hip. We first obtained the CT data of the patient and performed our 3D reconstruction process in a little over an hour. The resulting model of the boy's skull was 2.5 million polygons and his face was 1.6 million polygons. The surgeon could then visualize the full resolution version of his patient. With such a complex defect, stereo viewing of the data was found to be immensely useful for understanding the spatial relationships of the defect. After visualizing and printing full resolution preoperative images of the full, unreduced data, we used our mesh reduction software to reduce the size of the skull and face to approximately 400,000 polygons each. The surgeon could then rotate around the data and view it in a more interactive fashion. We then used our tools to reflect over the normal, unaffected side of the skull to the affected side to derive the piece of bone that would be needed to repair the defect. This surgical template was then examined, and images of it printed, creating a "blueprint" for repairing the defect. These template images were then traced onto radiological film so they could be sterilized and taken into the operating room, laid on the donor site, and used to extract the exact size and shape of bone needed to repair the defect.

Analysis

First, as far as the data, we found that most often surgeons wanted static, full resolution images of the patient printed for preoperative planning, but then were perfectly happy to use reduced models for interaction, as long as the details of the defect were not compromised.

The Immersive WorkBench display was also very well received by the users of the system. It enabled easy collaboration between users and fit the "many people working on a patient" (virtual model display) paradigm of real surgery well. In addition, the sheer size of the display (8 x 6 feet) was perceived as impressive the the users. The disadvantages of the display were that it was quite expensive ($100,000, fully configured) and other technical issues related to the fact that it was created as a general visualization device, not specialized toward our particular application. These issues included that the design employed an open area between the projector and the table which, due to stray background light and the lumen output of the projection, required us to turn out the lights in the room where the display is used in order to see the display clearly. Second, the maximum tilt of the screen was 30 degrees which required a remapping of our perspective when in stereo mode, since the top of the screen (i.e., top of the table) was roughly 5 feet away from the user, while the bottom of the screen was next to the user. (it is also interesting to note that many users of workbench-style displays still do not take this issue into account when developing applications).

As far as interaction, the selection of tool modality (moving tool, measuring tool, reflecting tool, etc) by pressing keys on a keyboard laying at the edge of the workbench table was clearly an inferior interface. In practice, Dr. Montgomery would be assisting the users of the system and handling these issues, but for a general tool this would clearly be a drawback.

In examining the software architecture, we found that the Sense8 WorldToolKit did, in fact, save us a great deal of time in that it eliminated our need to write any device drivers ourselves. In general, we found it to be a very well conceived and easy system to prototype and build applications rapidly. The company was also responsive to the performance issues we were faced with (our models were the largest they had dealt with) which are described in more detail below.

In short, performance on large meshes was a problem and we tried a number of techniques to provide reasonable tradeoffs. First, the WTK system has a preprocessing function that optimizes the geometry, creates display lists, and does other operations to increase the subsequent graphics performance. However, the optimization algorithm employed is O(n^2), which meant that it required nearly 20 minutes of processing time to optimize a mesh of over a million polygons. This delay before using the system was clearly unacceptable and greatly slowed the “implement, compile, test, refine” routine of application prototyping.

Another issue we found was that when the user wanted to interact with the model (by using the measuring tool, for example), the amount of data again overwhelmed the system’s abilities. When using the measuring tool, we would cast a ray from the tool into the dataset to determine the location that the user wanted to mark by searching through the the list of triangles and calculating the intersection. When processing on even smaller geometries, this test was found to be computationally expensive and incur great delays.

The method we developed to address both these issues was to instead decompose the geometry into a number of rectangular regions named “slabs”. By breaking the model apart into a number of these subobjects, the preprocessing step could occur much faster (2 minutes) because the number of triangles in each slab was much less for the O(N^2) algorithm. In addition, we could also greatly speed up interaction. For example, now when using the measuring tool, we would cast a ray and determine which slab the ray hit first, then make the call to iterate over only the polygons in that slab. The drawbacks of this approach however, related to operations that spanned more than one slab.

Even with these interaction speedups, we still found that some computationally intense operations, such as cutting, would create unacceptably long delays. Originally, we tried to allow the user to use a cutting tool and actually cut the mesh on the fly. Because of the computational limits imposed by the large geometry, we found that e could not maintain interactive rates. Therefore, we changed the paradigm slightly such that the surgeon would lay down markers interactively (which could be done quickly with the enhancements listed above), but then we would defer the expensive mesh subdivision operation of actually cutting between the markers until the end of the operation (when the user indicated that they had laid down all markers). By separating these operations, the system was much more usable and this idea of obtaining user input quickly and deferring computationally intense operations (defined as anything taking over 1 second) until the user indicates the end of the operation and is more accepting of delays was a particularly useful insight.

As our work progressed however, we found that we required more and more direct control over the rendering and internal representations of the data. While using commercial, generalized packages for rapid prototyping was found extremely useful, in the end we found that it limited the continued specialization of our application.

One issue where this was particularly noticeable was that we started to realize that just visualizing and even interacting with static models did not provide all the information that a surgeon really requires. The ability to determine how the soft tissues will lay over a new bone structure, or to extract other simulation-type information from the patient data was found to be of increasing importance. In short, surgical planning with static patient data was good and useful, but having dynamic models where the soft tissues behaved like real soft tissues and other physical properties were imparted upon the model would be great. Trying to build in a full soft-tissue modeling system into the visualization framework imposed by WorldToolKit would simply be unworkable as we moved forward.

Forth Instantiation (VESPA): 1998

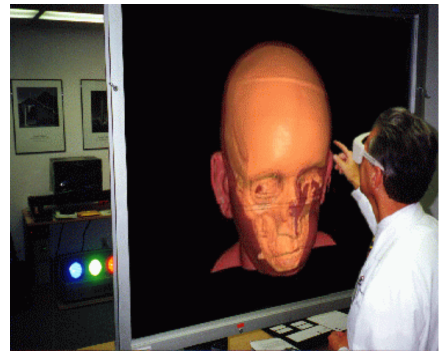

In the spring of 1998, NASA and Stanford formed the National Biocomputation Center in order to take these technologies and get them closer to the doctors and patients and to provide a focus that would not be possible with an organization physically and culturally located at some distance. Also, the goal was to accelerate the production of technologies that could not only dramatically improve medicine and surgery today on Earth, but to also make medical care for long-duration space flight (to Mars, for example) a reality. Finally, we were to lead the charge in bringing together all the groups across the country to organize and focus on these issues and to foster collaboration and communication.

During this first year, a tremendous amount of work was accomplished in setting up the new Center, the administration and logistical issues of starting such a new endeavor can not be underestimated. However, we also accomplished a surprisingly large amount of technical work during this time as well [Montgomery99a-d, Stephanides99a-b]. We rewrote all the software which we had previously prototyped at NASA Ames. We also switched platforms from Silicon Graphics computers to Sun Microsystems machines due to Sun's sufficient and increasing graphics performance, their superior computational performance, their stability as a company, and their enthusiasm and overwhelming support of the work we were attempting to accomplish.

Perhaps most significant was our realization that for a system to be truly useful, then it must not only provide for visualization and manipulation of static datasets, but that we needed to integrate in the ability to simulate how soft tissues behave and interact. This allows us to make predictions of surgical outcomes that were impossible previously.

Data:
After having spent considerable time over the past few years handling data media (optical disk, 4mm tape, 8mm tape, etc) and formatting (proprietary, DICOM flavors, etc) issues, we finally implemented our own DICOM server software, using the Mallinkrodt code base. By creating our own DICOM server, the surgeon could just request that the CT technician "push" the data to the Biocomp server and the data would be transferred directly to our lab over the medical center intranet. Once there, the DICOM server software would save each individual image slice as a DICOM file, which could then be read by our software.

In processing the data for a patient, our software would create an intermediate, internal, block-file format of the data. The DICOM image files would be examined, and split into separate directories based on image series and other parameters. Once there, the script would extract the image data from each DICOM image file and build a single, volumetric block file of the imaging data, as well as a separate file comprising all relevant DICOM header information. In this way, rather than parsing each DICOM image file when reading the data into our reconstruction software (~5 minutes of processing), we could instead read the header, then do a very fast block read of the data into memory (~2 seconds).

To perform segmentation and mesh generation, we recreated a similar (but improved) package to what we'd had at NASA and named it Reconstruction Of Volumetric Elements (ROVE). This package applies advanced segmentation techniques such as region growing, morphological operations, blob coloring, etc to classify voxels in the dataset. Then, it can generate a mesh using a Marching Cubes [Lorensen87] algorithm and apply many heuristics to clean up the data and produce a regularized mesh suitable for viewing or simulation.

For mesh reduction, we used the Qslim package [Garland97] that performed quadratic slimming of the given surface by creating and operating upon an intermediate surface description. We found this method to be vastly superior than the vertex elimination-based methods we had previously been using and it allowed us to create models that were as low as 100,000 polygons but retained sufficient detail for surgical planning.

Platform:
As mentioned above, we had moved from the Silicon Graphics (SGI) platform to a Sun Ultra60 with Elite3Dm6 graphics. This machine compared very favorably with the SGI machines in both computation and graphics performance and was only 30% the cost of the SGIs. This, together with an extremely supportive environment provided by Sun, made the production of this work greatly enhanced.

Display:
From our experiences with the FakeSpace Immersive WorkBench at NASA, we knew the drawbacks inherent in the design of general purpose projection display systems. We therefore sought to create the first workbench-style display specifically designed for surgical planning, the Surgical WorkBench (SWB). This system would be low cost, but employ the best projector and screen material available.

This visualization system employed a vertical screen, which we had found more conductive to the tasks in this domain. Because the screen was fixed in position, we could optimize the screen material for the system and use a material that gave a high gain (roughly brightness) and provided high contrast. After evaluating all screens in production, we selected the DiamondScreen from Draper Shade and Screen. This screen material has two layers of lenticular lenses embedded in the screen- one to produce a 180 degree field of view horizontally, the other to concentrate the light into a 30 degree field of view vertically. By concentrating the light, we felt that the system could be used with the room lights on, and would have a very bright image. We also felt that the 30 degree vertical field-of-view was not a limitation since most viewers fall within that height requirement. In addition, we build the system to be an enclosed rear-projection system to eliminate stray light from the back surface of the screen. For the projector, we selected the Sony VPH1292Q-3D (fast phosphor) projector. This projector had the greatest light output and stability and provided scan rates high enough to do 1600x1200 resolution in full stereo (120Hz). By building our own specialized system, we produced a superior device capable of exceptional visualization with limited expense (1/3 the cost of general purpose displays).

For stereo viewing, we continued to use the CrystalEyes stereo glasses that we had used while at NASA. These infrared glasses provided superior comfort and field of view than newer stereo eyeware that had emerged on the market.

Interface:
As for the interface, we realized that in order to support surgical simulation, we needed to support real surgical tools in our environment. Toward this end, we have instrumented a number of surgical tools with tracking devices and closure sensors to allow us to use them in the environment.

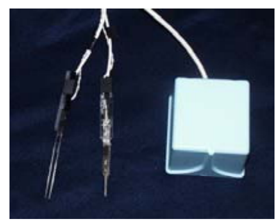

In addition, we have created a number of abstract tools which can be selected by using a stylus. All devices (real tools detecting closure) or stylus are wired into the system easily by connecting them to a connector wired into the system's mouse circuit board. Therefore, closing the forcep for example, is registered by the system as the left mouse button closing, etc. This very simple interface allows us to quickly create many different configurations of tools.

For tracking devices, we use Ascension Technologies electromagnetic trackers: Flock-of-Birds for tasks were 2mm accuracy is sufficient, a set of specially produced miniBird sensors for accuracy to .75mm, and a pcBird for low latency head-mounted display applications. For more limited tracking, we also support Intersense InterTrax inertial trackers. Finally, for high resolution, restricted workspace tracking (such as tracking drills or other large tools), we support the Microscribe armature-based tracking device (Immersion Corp).

As for haptic (force-feedback) devices, we currently use the Impulse Engine and 3GM devices from Immersion and the SensAble Devices PHANToM. These devices are becoming increasingly important to allow the user to perceive the data of their patient using other, perhaps more intuitive means.

While we have licensed voice recognition software from Dragon Systems, we have not yet integrated it with our core set of software. The ability to do voice I/O is very important toward providing a natural, hands-free method of indicating tool modality. Similarly, we have licensed speech synthesis technology from the A.I. DuPont Institute to allow us to provide speech output from the system as well.

Interaction:
In this system, we found that collaboration was key and interaction with computer models of vital importance. Besides the ability to rotate and translate models for visualization purposes, we also can select regions of objects, split objects into pieces, perform connected components analysis, and, because the system is integrated with our mass-spring dynamic simulation package, we can grab and pull on tissues, cut them, and perform other manipulations.

Software Architecture:
As stated briefly above, we rewrote all software and created a much more general and usable system that included the ability to do soft-tissue modeling within the system. Since our scope of work at the center included surgical planning, but also intraoperative assistance (augmented reality in the operating room), and surgical simulation/training, we needed a platform that supported fast and efficient simulation of soft tissues, as well as all the visualization features offered with our previous system.

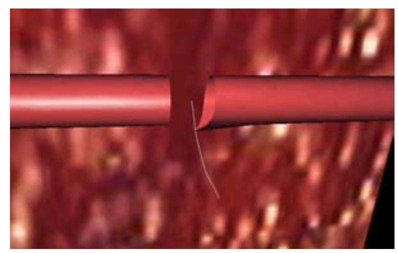

The mass-spring soft-tissue simulation software employs a point-mass and linear spring/damper model to rapidly approximate the deformation of soft-tissues. We early on investigated and abandoned using true finite-element methods for calculation of soft tissue dynamics, due to concerns over processing performance, large deformations, internal collisions, and non-linear operations. FEM systems, while precise, require enormous computation and were found inappropriate for real-time simulation. Instead, our mass-spring model has employed a number of different numerical techniques (Euler, Runge-Kutta, quasi-static, etc) for fast approximation of soft tissue dynamics.

In addition, our software simulation framework supports integrated networking. Through this facility, access to other, remote tracking or input devices (such as a haptic server) over the Internet is possible. In this case, the haptic server is a dedicated “black box” computer that maintains a high update rate with the haptic device and streams the current location of the physical tool to the simulation. The simulation then uses this location to render the virtual tool and calculate collisions and subsequent forces on the virtual tool. These forces are then streamed back to the haptic device for rendering. The ability to collaborate over the Internet and have multiple users in the same environment is a feature we wished to preserve for training and possible long-distance collaboration purposes. Further, the area of distributed haptics is one we wished to pursue.

Operation:
In this system, we not only used the system for visualization and interaction, but this time we have used the system for producing surgical templates which have been taken to the operating room and used to affect patient outcomes. The ability to produce templates of one’s surgical plan has shown itself to be a very important method for interfacing preoperative planning with the real surgery.

We now routinely use the system for preoperative visualization, 3D cephalometric analysis (distances, angles, areas, volumes), perform abstract operations (reflecting, etc) to plan surgeries.

Patients:
A woman in her late 40s presented with squamus cell carcinoma (T4N0M0) that required excision of the tumor and reconstruction of the floor of the mouth. Since the tumor had infiltrated a great deal of the bone and soft tissue, the entire jaw required removal and reconstruction. The fibula is a natural candidate for donor bone, but presents the difficulty of shaping a straight, stick-like bone into the complex shape of the jaw. We therefore performed a 3D

reconstruction of the woman's bone and soft tissue from her preoperative CT data. The surgeon, Dr Stephanides, then used the resulting model to visualize and understand the extent of the defect and to measure distances and angles of her existing mandible. Then, a fibula was extracted from an existing dataset of a woman of similar age and height and the same distances and angles were marked on this virtual fibula. From this information, small wedge-like templates were created that, when applied to the fibula, would produce the shape of the woman's original jaw. These templates were traced onto film, taken to the OR, and used to perform the reconstruction.

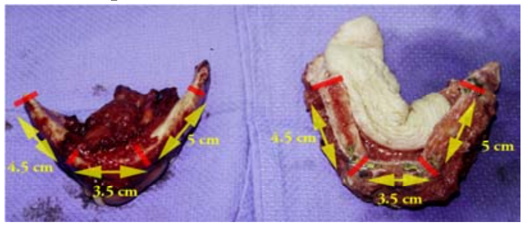

Analysis:
The complete recreation of all the software and inclusion of soft-tissue modeling has been a very large undertaking. The resulting software base, while complex in its functionality, can be applied in many areas of surgical planning and simulation and we are only now beginning to create end-user applications exploiting this framework.

We have learned a few important lessons in the application of these technologies. First, while surgical planning templates are a very good method for instantiating and using a surgical plan, it is sometimes the case that, during an operation, one must modify one's plan on the fly because of situations that could now have been foreseen. One that we have encountered is that a donor site may not be compatible or appropriate for use and this sometimes is only determined during the procedure. Clearly it would be useful to have access to all surgical planning software during the procedure itself, in order to modify the surgical plan when one encounters such a difficulty.

Second, we are beginning to deploy this system by automating the reconstruction process and providing a web-based (Java 3D) surgical planning tool on our web site to selected surgeons. We wish to limit the number of users at this point in order to gain knowledge about their use and plan for a higher capacity. But it is our goal to make surgical planning software available to every surgeon's office over the Internet. The benefits that we have seen, when replicated a thousand-fold, would provide a great advantage in producing superior surgical outcomes, with less risk to the patient, and at less cost.

Conclusion:

To summarize all the work we have done in this field over the past decade, a number of common threads appear.

First, in display technology, we started with a nonimmersive, noncollaborative display (computer screen); progressed to an immersive, non-collaborative display (BOOM); continued to an immersive, collaborative display (distributed, networked BOOM); and ended with a non-immersive, collaborative display (Surgical WorkBench).

Along this path, we have also seen a number of assumptions be cast aside. The idea that surgical planning equates to surgical simulation was a major assumption invalidated over this time. This perception of a lone surgeon performing their surgery ahead of time was one that seemed simple enough to be true, but did not bear resemblance to the reality. The reality is that surgical planning is an abstract process and is very different from the needs imposed for surgical simulation and training. The goal here is planning, not realism.

Further, the idea that surgical planning required the application of real tools was one that further became obsolete over time. We found that surgical planning was an abstract process that some surgeons can't even explain well themselves. It truly appears that we have progressed to the point where the surgeon has opened the door to a new place where there need be no correlates to the existing world. We are free to produce tools and method of interaction, which bear little resemblance to the real world, but are found to be of enormous benefit to the surgeon's practice and the patient's outcome.

We hope that we have finally produced a system that "thinks like a surgeon", although we realistically believe that user interface work in this area will certainly continue. To derive the cognitive primitives of a task is a very hard thing to do.

Ideally, the system should be intuitive to the surgeon (they should be unaware they are using a computer), or certainly unobtrusive to the process.

We found that collaboration is more important than we certainly imagined, and that stereo visualization is truly useful and provides real benefits to the surgeon's understanding of their patient's anatomy. We also found that surgeons often like to be provided with a few, static, very high resolution images of their patient's data, taken at a few well defined locations. Sometimes these locations are standard (lateral, frontal, ¾ lateral, etc) and sometimes completely new (cutting plane through the back of the head to examine condyle angles). We also found that surgeons are very tolerant of reduced data, if it allows them to interact with the data to accomplish their goal.

As far as adoption of the system, we have never found a surgeon that wasn't very ready to adopt any technology that could help them perhaps obtain an improved outcome. While there have been some surgeons that have been less tolerant of the system's shortcomings (in some ways rightly so), by and large they have been extremely supportive of this work.

There are some barriers to full adoption and deployment however. The costs of such systems are still prohibitive to most users. It is our hope that through centralization and use of the world-wide web as an access medium, we may have one shared resource able to enable any number of remote surgeons to have access to these visualization and interaction tools. Further, we must do outcomes analysis to prove and validate whether computer-based surgical planning is a cost-effective tool in patient care before widespread adoption can take place. Toward this end, we must continue to develop our system, obtain feedback, and refine the system to produce a usable tool whose benefits and ease of use outweigh any disadvantages.

We sincerely hope that by exposing our own biases and openly and objectively discussing the design, paradigms, and results of our work in this field, that others will find this information useful, derive benefit from our work, and avoid the pitfalls that we have encountered. After nearly 40 cases of virtual environment surgical planning, we are still learning with each new case and our paradigms continually evolving. We sense that the pace of this work has been accelerating, especially as we begin to open it up to more surgeons of different specialties and patients of different needs.

Finally, after all this work, we are unshaken in our belief that computer-baser visualization, interaction, and simulation is a useful tool in surgery today, and will be a standard of care in surgery tomorrow, providing better outcomes, with less time in the OR, lower cost, and decreased risk to the patient.

REFERENCES

[USAF] Air Force Institute of Technology's Graduate School of Engineering, Virtual Environments, 3D Medical Imaging, and Computer Graphics Laboratory http://www.afit.af.mil/ENGgraphics/GRAPHICS.html

[CMU] Carnegie Mellon University, Center for Medical Robotics and Computer Assisted Surgery (MRCAS), http://www.mrcas.ri.cmu.edu/

[Fraunhofer] Fraunhofer Institute, Center for Research in Computer Graphics (CRCG), http://www.CRCG.edu/, Institut Graphische Datenverarbeitung (IGD), http://www.igd.fhg.de/www/igd-a4/flyers/medicine/

[GaTech] Georgia Tech, Graphics, Visualization and Usability Center (GVU), http://www.cc.gatech.edu/gvu/gvutop.html, Medical College of Georgia, Modeling, Imaging and Visualization, http://www.cc.gatech.edu/gvu/visualization/surgsim/

[Harvard] Harvard, Surgical Planning Laboratory (SPL), http://splweb.harvard.edu

[Erlangen] Friedrich-Alexander-Universitat Erlangen-Nurnberg, http:// www-nt.e-technik.uni-erlangen.de

[UIC] University of Illinois at Chicago, VRMed Lab, http://www.bvis.uic.edu

[ISS] Singapore Institute of Systems Science (ISS), http://www.iss.nus.sg/, CIeMed: Center for Information-Enhanced Medicine, http://www.ciemed.nus.edu.sg

[Mannheim] Mannheim, University of, Virtual Reality in Medicine (VIRIM), http://www-mp.informatik.uni-mannheim.de/groups/virim/virim.html

[MIT] Massachusetts Institute of Technology, Biorobotics Group, http://biorobotics.mit.edu/

[Mayo] Mayo Clinic, VRASP: Virtual Reality Assisted Surgery Program, The Mayo Biomedical Imaging Resource (BIR), http://www.mayo.edu/bir/vrasp/vr_intro.html

[NCCJ] National Cancer Center of Japan, http://medvr.res.ncc.go.jp/

[Sheffield] Sheffield, University of, Virtual Reality in Medicine and Biology Group (VRMBG), http://www.shef.ac.uk/~vrmbg

[DTU] Technical University of Denmark (DTU), 3D-Med project, http://www.it.dtu.dk/~hahl/3dmed/index_old.html

[HITL] Washington, University of, Human Interface Technology Laboratory (HIT Lab), http://www.hitl.washington.edu/

[Ayoub96] Ayoub, A.F., Wray, D., Moos, K.F., Siebert, P., Jin, J., Niblett, T.B., Urquhart, C. and Mowforth, P. Three-dimensional modeling for modern diagnosis and planning in maxillofacial surgery. Int. J. Adult Orthod. Orthognath. Surg. 11:225-233 (1996).

[Bechtold93] Bechtold, J.E. and Powless, S.H. The application of computer graphics in foot and ankle surgical planning and reconstruction. Clin. Podiatr. Med. Surg.10: 551-562 (1993).

[Byrd93] Byrd, H. S. and Hobar, P.C. Rhinoplasty: a practical guide for surgical planning. Plast. Reconstr. Surg. 91:642-654 (1993).

[Chao95] Chao, E.Y.S. and Sim, F.H. Computer-aided preoperative planning in knee osteotomy. Iowa Orthop. J. 15:4-18 (1995).

[Garland97] Garland, M, Heckbert, P "Surface simplification using quadratic error metrics", Proc SIGGRAPH 97.

[Gautsch94] Gautsch, T.L., Johnson, E.E. and Seeger, L.L. True three dimensional stereographic display of 3D reconstructed CT scans of the pelvis and acetabulum. Clin. Orthop. 305:138-151 (1994).

[Gerbo97] Gerbo, L.R., Poulton, D.R., Covell, D.A. and Russell, C.A. A comparison of a computer-based orthognathic surgery prediction system to postsurgical results. Int. J. Adult Orthod Orthognath. Surg. 12:55-63 (1997).

[Haasfeld98] Haasfeld, S., Zöller, J., Albert, F.K., Wirtz, C.R., Knauth, M. and Mühling, J. Preoperative planning and intraoperative navigation in skull base surgery. J. Cranio-Maxillofac. Surg. 26:220-225 (1998).

[Lamoral90] Lamoral, Y., Quirynen, M., Peene, P., Vanneste, F., Lemahieu, S.F., Baert, A.L. and van Steenberghe, D. Computer tomography in the preoperative planning of oral endo-osseous implant surgery. Rofo Fortschr Geb Rontgenstr Neuen Bildgeb Verfahr 153:505-509 (1990).

[Linney93] Linney, A.D., Tan, A.C., Richards, R., Gardener, J., Grindrod, S. and Moss, J.P. Three-dimensional visualization of data on human anatomy: diagnosis and surgical planning. J. Audiov. Media Med. 16:4-10 (1993).

[Lorensen87] Lorensen, W. E. and H. E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," Computer Graphics, vol. 21, no. 3, pp. 163-169, July 1987.

[Lotwin93] Lotwin, R.S. and Baerg, R.H. Orthographics. A two-dimensional surgical planning system. Clin. Podiatr. Med. Surg. 10: 539-549 (1993).

[Montgomery99a] Montgomery, K; Stephanides, M; Schendel, S; "Craniofacial Surgery Planning Using a Virtual Environment: Virtual Tools and Their Interactive Application"; Computer Assisted Radiology and Surgery (CARS) `99, Paris, France.

[Montgomery99b] Kevin Montgomery PhD, Michael Stephanides MD, Stephen Schendel MD, Muriel Ross PhD; "Preoperative Planning of a Craniofacial Defect Using a Virtual Environment", Medicine Meets Virtual Reality, January 1999.

[Montgomery99c] Kevin Montgomery PhD, Michael Stephanides MD, Stephen Schendel MD, Muriel Ross PhD "Virtual Surgical Tools for Planning Craniofacial Surgery", Medicine Meets Virtual Reality, January 1999.

[Montgomery99d] Kevin Montgomery Ph.D., Michael Stephanides MD, Joel Brown BS, Jean-Claude Latombe Ph.D., Stephen Schendel MD, DDS; "Virtual Reality in Microsurgery Training", Medicine Meets Virtual Reality, January 1999.

[Montgomery98a] Montgomery K, Stephanides M, Schendel S, Ross M; "A Case Study Using the Virtual Environment for Reconstructive Surgery", IEEE Visualization `98, Research Triangle Park, NC, October, 1998.

[Montgomery98b] Montgomery, K; Stephanides, M; Cheng, R; Schendel, S; Ross, M; "Clinical Application of VR for Surgery", SPIE Electronic Imaging- Engineering Reality of Virtual Reality, San Jose, CA, January 1998.

[Montgomery97] Montgomery, K, "VERS- A virtual environment for reconstructive surgery planning", Engineering Reality of Virtual Reality, SPIE Electronic Imaging, San Jose, CA, February 1997.

[Montgomery96a] Montgomery K. "VERS: Virtual Environment for Reconstructive Surgery", Presentation: University of California, Santa Cruz, April 10, 1996.

[Montgomery96b] Montgomery K. "VERS: Virtual Environment for Reconstructive Surgery", Presentation: Sterling Noontime Seminar, NASA Ames Research Center, February 22, 1996.

[Moss88] Moss, J.P., Grindrod, S.R., Linney, A.D., Arridge, S.R. and James, D. A computer system for the interactive planning and prediction of maxillofacial surgery. Am. J. Orthod. Dentofac. Orthop. 94:469-475 (1988).

[Perry98] Perry, M., Banks, P., Richards, R., Friedman, E.P. and Shaw, P. The use of computer-generated three-dimensional models in orbital reconstruction. Br. J. Oral Maxillofac. Surg. 36: 275-284 (1998).

[Pokrandt96] Pokrandt, P., Bohner, P., Knolle, J.P., Hassfeld, S. and Raczkowsky, J. A communication system supporting simultaneous planning and execution in cranio-maxillo-facial surgery. J. Telemed. Telecare 2:3-7 (1996).

[Rozema92] Rozema, F.R., Otten, E., Bos, R.R.M., Boering, G. and van Willigen, J.D. Computer-aided optimization of choice and positioning of bone plates and screws used for internal fixation of mandibular fractures. Int. J. Oral Maxillofac. Surg. 21: 373-377 (1992).

[Stephanides99a] Stephanides, M; Montgomery, K; Schendel, S; "Microsurgical Reconstruction of the Mandible using a new 3D Surgical Planning System"; Computer Assisted Radiology and Surgery (CARS) `99, Paris, France.

[Stephanides99b] Stephanides M, Montgomery K, Bowen V, Hentz V, Schendel S; "Computerized Diagnostic and Surgical Planning Tools for the Surgery of the Wrist", Medicine Meets Virtual Reality, January 1999.

[Sutherland94] Sutherland, C.J., Bresina, S.J. and Gayou, D.E. Use of general purpose mechanical computer assisted engineering sofware in orthopaedic surgical planning: advantages and limitations. Comput. Med. Imaging Graph. 18:435-442 (1994).

[Vannier96] Vannier, M.W. and Marsh, J.L. Three-dimensional imaging, surgical planning, and image-guided therapy. Radiol. Clin. North Am. 34: 545-563 (1996).

[Verstreken96] Verstreken, K., Van Cleynenbreugel, J., Marchal, G., Naert, I., Suetens, P. and van Steenberghe, D. Computer-assisted planning of oral implant surgery: A three-dimensional approach. Int. J. Oral Maxillofac. Implants 11:806-810 (1996).

[Yune93] Yune, H.Y. Two-dimensional-three-dimensional reconstruction computed tomography techniques. Dent. Clin. North Am. 37:613-626 (1993).

[Zonneveld89] Zonneveld, F.W., Lobregt, S., van der Meulen, J.C.H., and Vaandrager, J.M. Three-dimensional imaging in craniofacial surgery. World J. Surg. 13: 328-342 (1989).