Int. J. Radiation Oncology Biol. Phys., Vol. 25, pp. 123–129
Printed in the U.S.A. All rights reserved.

0360-3016/93 $6.00 + .00

Technical Innovations and Notes

VISTANET: INTERACTIVE REAL-TIME CALCULATION AND DISPLAY OF 3-DIMENSIONAL RADIATION DOSE: AN APPLICATION OF GIGABIT NETWORKING

JULIAN G. ROSENMAN, PH.D., M.D.,[1,2] EDWARD L. CHANEY, PH.D.,[1] TIM J. CULLIP, M.S.,[1]
JAMES R. SYMON, M.S.,[1] VERNON L. CHI, B.S.,[2] HENRY FUCHS, PH.D.[1,2]
AND DANIEL S. STEVENSON, M.S.[3]

[1] Departments of Radiation Oncology and [2]Computer Science, University of North Carolina at Chapel Hill, NC;
and [3]The MCNC Center for Communications

Three-dimensional treatment planning can allow the clinician to create plans that are highly individualized for each patient. However, in lifting the constraints traditionally imposed by 2-dimensional planning, the clinician is faced with the need to compare a much larger number of plans. Although methods to automate that process are being developed, it is not yet clear how well they will perform. VISTAnet is a 3 year collaborative effort between the Departments of Radiation Oncology and Computer Science at the University of North Carolina, the North Carolina Supercomputing Center, BellSouth, and GTE with the medical goal of providing real-time 3-dimensional radiation dose calculation and display. With VISTAnet technology and resources, the user can inspect 3-dimensional treatment plans in real-time along with the associated dose volume histograms and can fine tune these plans in real-time with regard to beam position, weighting, wedging, and shape. Thus VISTAnet provides an alternate and, possibly, complementary approach to computerized searches for optimal radiation treatment plans. Building this system has required the development of very fast radiation dose code, methods for simultaneously manipulating and modifying multiple radiation beams, and new visualizations of 3-dimensional dose distributions.

3-Dimensional treatment planning, 3-Dimensional displays, Computers.

INTRODUCTION

In 1989 a consortium of the Radiation Oncology and Computer Science Departments at the University of North Carolina, BellSouth Corporation, GTE, and the MCNC Center for Communications (MCNC) was formed to prepare a response to the high speed network initiative proposed by the National Science Foundation (NSF), and the Defense Advanced Research Projects Agency (DARPA) and coordinated through the Corporation for National Research Initiatives (CNRI). The initiative was aimed at developing the technology to boost computer networking speeds up to a gigabit per second (10^9 bits/second). This rate of data transfer is almost a thousand times faster than currently available and would provide the bandwidth necessary for near real-time (≥ 1 frame/second in this paper) interaction, from a remote site, with the most powerful computing equipment in the world. To accomplish this goal CNRI proposed to fund several "test bed" sites where gigabit networks would be developed using driving problems that required this rate of data transmission.

Our consortium, known as VISTAnet, proposed to use real-time radiation therapy treatment planning as the driving problem. The plan was to develop a system that could calculate and display a three-dimensional radiation dose distribution for any number of radiation beams in less than a second, with no restrictions on beam angles or modifiers. With this system the user could quickly search through potential treatment plans as fast as they could be specified, with immediate feedback as to whether a given modification made the plan better or worse. The overall medical goal of the project was to allow the user to rapidly arrive at a highly optimized treatment plan.

Reprint requests to: Dr. Rosenman, Department of Radiation Oncology, University of North Carolina School of Medicine, Chapel Hill, NC 27599.

Acknowledgements—We would like to acknowledge the National Science Foundation (NSF) and the Defense Advanced Research Projects Agency (DARPA)'s financial support through Research Grant NCR-8919038 "A very high bandwidth prototype network for interactive 3D medical imaging research." We also wish to acknowledge the financial and technical support of BellSouth and GTE. Finally, we wish to credit the entire consortium composed of the Corporation for National Research Initiatives (directed by Dr. Robert Kahn), the North Carolina Supercomputing Center, BellSouth, and GTE, for their time and effort on behalf of this project.

Accepted for publication 8 July 1992.

Table 1. Funded gigabet network test beds

Test bed	Universities	Industrial and other partners	Application
AURORA	University of Pennsylvania and MIT	Bell Atlantic, Bellcore, IBM, MCI, and Nynex	Virtual laboratory, advanced teleconferencing
BLANCA	Universities of California (Berkeley), Illinois, and Wisconsin, Lawrence Berkeley Laboratory	AT&T, Ameritech, Astronautics, Bell Atlantic, Cray Research, National Center for Supercomputing Applications, Norlight, and Pacific Telesis	Atmospheric storm modeling, radioastronomy imaging, multimedia digital library
CASA	Cal Tech, University of San Diego (Supercomputing Center)	JPL, Los Alamos National Laboratory, MCI, Pacific Telesis, and US West	Geophysical, global climate, and chemical reaction modeling
NECTAR	Carnegie-Mellon University and the University of Pittsburgh	Bell Atlantic, and Bellcore	Traveling salesman problem, chemical flowsheet modeling
VISTAnet	University of North Carolina	BellSouth, GTE, and MCNC	Real-time radiation treatment planning.

In the fall of 1990 it was announced that five test bed sites would be funded, including VISTAnet. The sites, listed in Table 1, are described in more detail in reference (8).

Each group in the consortium is working on a different task. The Radiation Oncology group is developing radiation dose code that could calculate a 3-dimensional dose grid (> 50,000 points) in less than a second on a CRAY YMP supercomputer owned by the North Carolina Supercomputing Center (a division of MCNC). Part of this task is to provide the appropriate user interface.

Computer Science is providing the hardware and software needed to visualize the vast amount of data that pours from the CRAY, and the hardware interface between the high speed network and Pixel-Planes, the high-speed multi-processing graphics computer under development in that department (3). MCNC is providing the CRAY support, and developing interface boards between the high speed network and ordinary workstations. MCNC is also providing administration for the project. Finally, BellSouth and GTE are responsible for providing the networking hardware and high speed cable to link the CRAY, Pixel Planes, and the computer workstation in Radiation Oncology (see Fig. 1).

METHODS AND MATERIALS

The goals of the Radiation Oncology arm of VISTAnet were threefold: (a) To provide a user interface that would allow a user to quickly and easily explore a wide variety of treatment options; (b) provide software that accurately calculates a 3-dimensional radiation dose distribution in a second or less; and (c) work with Computer Science to find a strategy to display the dose distribution in a manner that can provide immediate feedback to the user.

The user interface

Early in the VISTAnet project it was decided to allow the user the following treatment options:

1. An unlimited number of radiation beams.
2. The radiation beams could be placed at arbitrary angles around one or more isocenters.
3. The radiation beams could be shaped in any fashion desirable.
4. The beams could have arbitrary weighting or wedges.

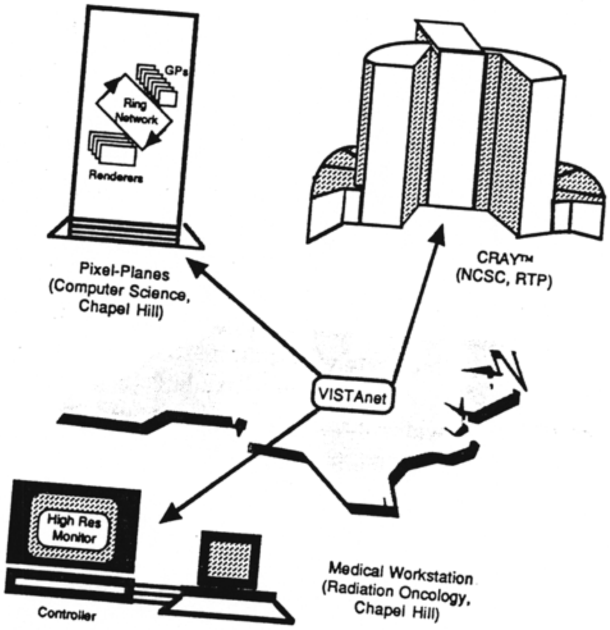

Fig. 1. The VISTAnet arrangement RTP (the Research Triangle Park) is located about 15 miles from Chapel Hill. Computer Science and Radiation Oncology are about 1/4 mile apart.

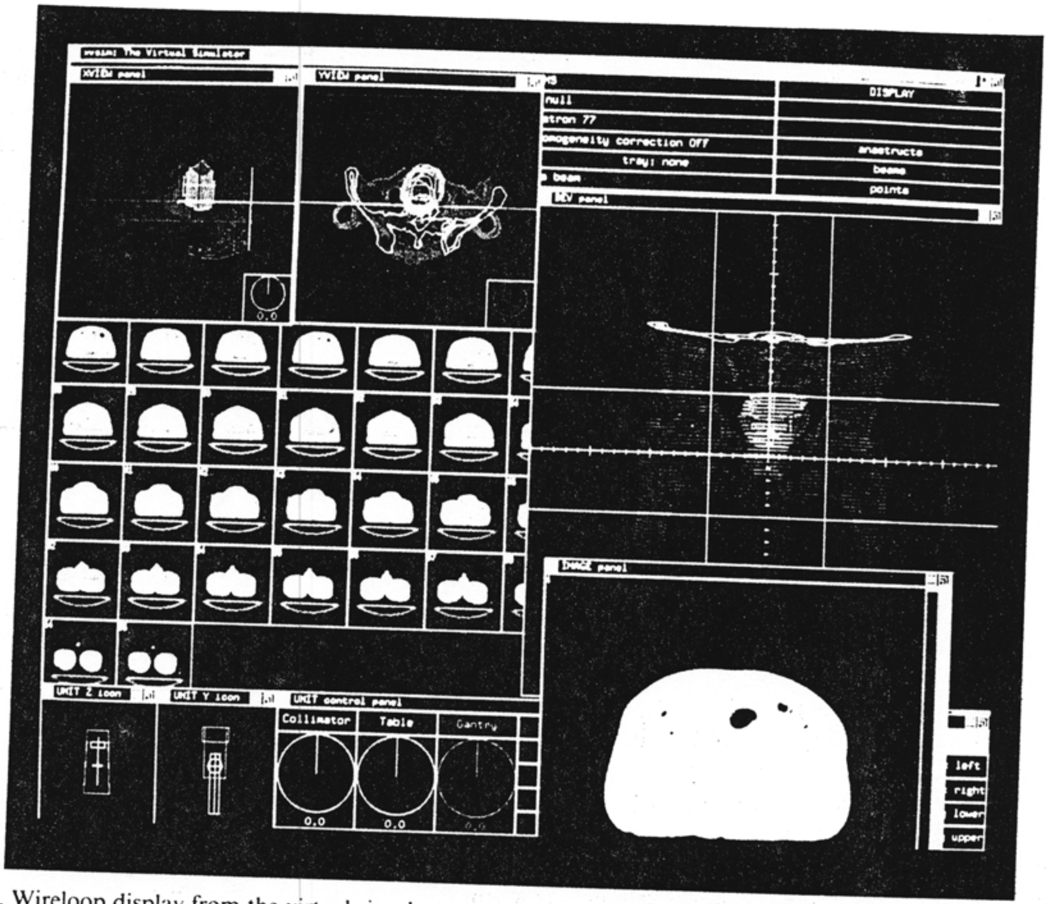

Fig. 2. Wireloop display from the virtual simulator. An anterior field is being placed over the pelvis as seen in the right upper window (the yellow lines mark the field edges as in a physical simulator). The beam angle can be rapidly changed using the "control panel" seen on the bottom left. This display (with modifications to allow for beam weighting, wedging, and shaping) serves as the interface for VISTAnet.

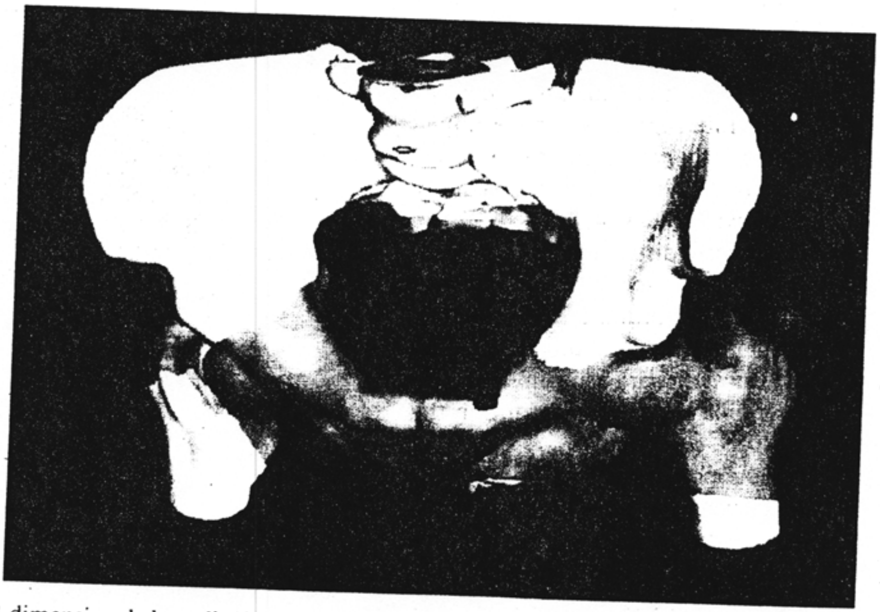

Fig. 3. The 3-dimensional dose distribution resulting from a four field beam arrangement shown in Figure 2. Shown are the 30% isodose surface (light red) and 90% isodose surface (dark red). The image is fully interactive in that one can rotate and "slice" it in real-time. It is also possible to put back skin and muscles and other organs, although this tends to obscure the dose distribution.

Furthermore, the interface had to allow the user to rapidly change all of these values in a smooth manner.

Another early decision in VISTAnet was to use the virtual simulator, developed at the University of North Carolina by Sherouse (6), as the starting point for the interface. The virtual simulator is a software implementation of a physical radiation simulator which allows the user to place an arbitrary number of beams around an arbitrary isocenter. Unmodified, it could satisfy requirements 1) and 2). The general plan of VISTAnet was thus set. The user would create and place radiation beams on a wireloop display of the patient (Fig. 2). The beam position and shape would be sent to the CRAY-YMP where the radiation dose would be rapidly calculated. The dose matrix would then be shipped via the gigabit network to Pixel-Planes, which would convert the dose matrix, and patient anatomy into a 3-dimensional image and send it to a monitor sitting next to the workstation running the virtual simulator (Fig. 3).

Consider the "four field box" of a prostate cancer shown in Figs. 2 and 3. In this patient, prostate tissue extends posteriorly along the right side of the rectum. To cover the entire prostate, but still spare some rectum, one might try obliquing all four fields slightly but keeping the relative angle between the fields fixed. VISTAnet can then be used to determine the preferred angle. However, it may be convenient to rotate two or more beams together, as moving the beams one at a time could lead to uninteresting dose distributions that are distracting. For this reason we have introduced the concept of "beam grouping," wherein the user can specify which groups of beams are moved together. More than one group can be specified. Similarly, beams that are "grouped" together for one purpose can be "ungrouped" for another.

Since the original fields are shaped to conform to the tumor they must be reshaped as they are moved. Furthermore the reshaping must be done quickly so that the continual feedback is not interrupted. We have implemented an elementary "autoshaper" using the "convex hull" method. Although the convex hull method cannot follow a concavity (as shown in Figs. 4a and 4b) it can shape a radiation field a specified distance around a tumor target volume sufficiently rapidly so as to introduce no noticeable delay, even when many beams are being moved simultaneously. More sophisticated methods of beam shaping than can follow concavities are being investigated, but at present all are too time consuming for use in the VISTAnet project.

Beam weighting is also supported in VISTAnet. Beams can be given a weight factor of zero to 10 (in increments of 0.1). Because beam weighting is often done in pairs, beam grouping can be used with beam weighting. Thus, for a four field box, one could gradually "turn up" the weighting for the lateral fields as a pair. Continual feedback could be lost if one had to weight each field individually.

Wedges are another form of beam modulation. In VISTAnet we will allow wedging for each beam between 0° and 60°. Wedge angles other than the standard 0°, 30°, 45°, and 60° are achieved by using two appropriately weighted beams with wedges bracketing the desired value. For example, a 40° wedge would be approximated with a pair of 30° and 45° wedged beams, weighted 1:2. Beam grouping will be allowed for only two fields at a time, and the wedges will automatically be mirror images unless otherwise specified.

Rapid beam editing is also supported. In VISTAnet a beam outline is a polygon and the current implementation allows one to "grab" a vertex and move it. In essence one can "move a block" under VISTAnet and watch the doses change.

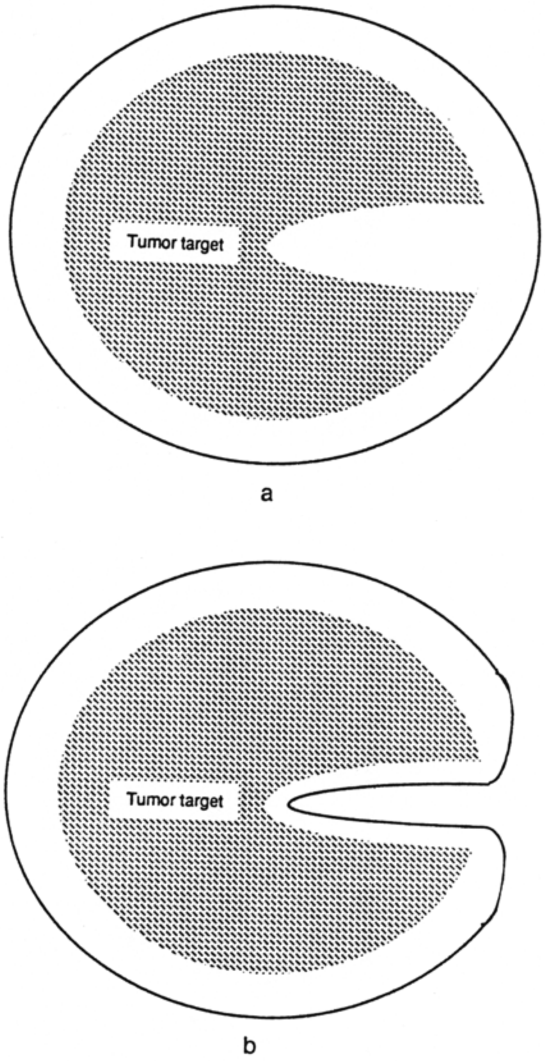

Fig. 4. (a) The convex hull method puts a convex shape around the tumor target. This is not as geometrically accurate as the field shape in Figure 4b, but it is much simpler to implement. (b) This tumor target outline is more geometrically accurate than the convex hull method in Figure 4a. It is more complex to implement and there is a real question as to whether radiation fields ought to be shaped like this.

Dose code

The calculation of a full 3-dimensional radiation dose distribution rapidly enough to allow for continual feedback was a formidable task, as a decision was made not to sacrifice accuracy. The VISTAnet team chose to develop high speed 3-dimensional dose code by modifying a differential SAR algorithm already in use in our department (1).

The computational speed of the primary dose computation has been dramatically improved by replacing ray casting algorithms with scan conversion algorithms for patient contours, filters, and beam outlines. The method consists of a two step process:

1. Replace ray casting with polygon scan conversion from the beam's eye view, projected onto a plane at SAD before dose computation. This produces a 2-dimensional table of linked lists of tissue entry/exit pairs, a 2-dimensional table of filter thickness, and a beam map.
2. Use table lookup/interpolation during dose computation.

For scatter dose, the differential SAR algorithm speed has been improved by implementing a modification of a sampling method first proposed by Niemierko and Goitein (4). Finally, whenever possible the dose code was "vectorized" to take advantage of the pipeline architecture of the CRAY YMP.

Display

Several major problems had to be overcome to satisfy the display requirements for VISTAnet. First, high quality interactive 3-dimensional anatomic images, created for the CT data set, were needed as context for the radiation dose distribution. Second, certain critical organs, such as spinal cord or kidney, had to be visible within the 3-dimensional image. Third, 3-dimensional display of multiple dose levels was needed. The latter is a particularly difficult task because simply adopting different colors for different dose levels can be confusing in 3-dimensions where one has to view one dose level through another. Finally, the display had to combine the 3-dimensional anatomy, critical organs, and dose distributions in a meaningful way for the viewer.

In addition to the display of radiation dose other real-time measures of the quality of a candidate treatment plan were needed. These include the dose volume histogram and models of tumor control probability (TCP) and normal tissue complication probability (NTCP).

RESULTS

An intermediate version of the planning system is maintained at all times during development. As progress and improvements are made they are incorporated into the current version. This approach allows the demonstration system to gradually evolve into a final version. All of the results described below were obtained from this system which is still under development.

Interface

The modified virtual simulator has been integrated into VISTAnet and appears to serve adequately as an interface. Beam position and shape data flows from the virtual simulator over an interim 15 megabit/second network to the CRAY YMP. Several modifications of the virtual simulator had to be made. It is now possible to group and ungroup two or more radiation beams together. Whatever operation is carried out on one beam in a group is carried out on all beams in that group. Such operations include moving the beam, "autoshaping" the beam, and beam weighting. For wedging, no more than two beams at a time can be grouped, and shape editing can take place only on one beam at a time.

Autoshaping using the convex-hull method proceeds transparently to the user. That is, the beams move around at essentially the same speed regardless of whether or not autoshaping is taking place.

In the four-field pelvic setup the four beams were grouped together and then rotated (and individually autoshaped) as a unit. In a matter of minutes an angle was found for which the prostate was totally covered by the 95% isodose surface with a minimum of rectum within that dose level. Thus an "optimal" treatment plan was easily found for this highly constrained problem.

Dose code

The speedup in dose calculation (on a 20 mips DEC 5000 workstation) using the above described methods was about 100 in the primary dose calculation and about 10 in the scatter dose calculation. The overall speedup is by approximately a factor of 50. A $128 \times 128 \times 64$ dose grid can be computed in about 45 seconds/beam with these improvements. Porting the code to the CRAY brought the calculation time to ~ 1 second/beam/CRAY processor. Since our CRAY YMP has four processors it seems likely that we can achieve our goal of ~ 1 sec for the entire calculation.

Display

In the earliest stage of VISTAnet development some kind of display was needed so that other components of the system could be evaluated. It was decided to use a digitally reconstructed radiograph (DRR) (7) as the anatomical model and display a singe user defined dose level on it. This meant, in effect, producing DRR's sufficiently rapidly so as to have, in effect, computed fluoroscopy. Using a modification of our DRR code and adapting it to Pixel-Planes we have achieved a speed of 15 low quality images/second, and about 1 high quality image per second (2).

When the user first begins a session with the planning system the computer fluoroscopic display is "locked" to the active beam in the virtual simulator. That is, the DRR on display at any time is the beam's eye view of the active beam. This is called the "mode A" display. The display can be decoupled from the virtual simulator ("mode B") so that the user can study a plan in greater detail by rotation, plane-clipping and other manipulations. At any time the display can be re-locked to the virtual simulator (brought back to mode A).

Based on 4 years of experience we now believe that some form of volume rendering produces displays that most accurately portray human anatomy for CT and MR. Until very recently such high quality renderings could not be produced at interactive speed. We have now developed a volume renderer that, when implemented on Pixel-Planes 5, can produce high quality images at the rate of one per second. Effective displays of radiation dose have been limited so far to two dose levels (Fig. 3), although experimentation with meshes and textures suggest that a clear 3-dimensional representation of three nested, discrete dose levels may be possible. We believe that because these dose levels are dynamically adjustable, the need for more than three dose levels (as in traditional static displays) will not be necessary.

Dose volume histograms, and TCP and NTCP models have been implemented and run in near real-time.

Need for a gigabit network

Taking into consideration the results so far it is possible to estimate the bandwidth necessary for the smooth functioning of the planning system. We intend to generate a 128 × 128 × 64 dose grid at the CRAY and then to interpolate it up to 256 × 256 × 128 points. (The third number will depend somewhat on the number of CT slices in the study, but 64 is an "typical" number in our experience). The interpolation is done because our experience with the existing slower (provisional) network has indicated that 128 × 128 displays look jagged when rendered in 3-dimensions. The CRAY can perform this interpolation in just a fraction of a second and therefore only a tiny time lag will be introduced. If the goal is two display updates per second, the CRAY must send

$$\frac{256 \times 256 \times 128 \text{ dose points}}{\text{frame}} \times \frac{16 \text{ bits}}{\text{dose point}} \times \frac{2 \text{ frames}}{\text{second}} = \frac{2^{28}\text{bits}}{\text{second}}$$

or $\frac{256 \text{ megabits}}{\text{second}}$ to Pixel-Planes for rendering. To fine tune the beam weights or wedges, an update frame rate of 10 per second is possible and desirable. Possible, because re-summing the pre-calculated contribution from each beam can be completely vectorized on the CRAY, and desirable because "fine tuning" requires a very responsive display. Studies with flight simulators have indicated that comprehension improves with frame rates up to 30/second or more. A frame rate of 10/second would completely saturate our network. The demands on the network are somewhat less for the link between Pixel-Planes and the medical work station. For our anticipated display we will need

$$\frac{1200 \times 100 \text{ pixels}}{\text{frame}} \times \frac{24 \text{ bits}}{\text{pixel}} \times \frac{2 \text{ frames}}{\text{second}} \sim \frac{60 \text{ megabits}}{\text{second}}$$

DISCUSSION

True three-dimensional treatment planning (as opposed to multiplaner 2-dimensional planning) removes many of the traditional restraints of 2-dimensional planning. Radiation beams need not lie in the transverse plane, indeed their central rays need not even be coplanar (for example, three beams could be mutually perpendicular). The advantage going from even multiplaner 2-dimensions to 3-dimensions is that it makes highly conformal radiation therapy planning possible. However, comparing the multitude of possible treatment plans made possible with 3-dimensional treatment planning puts a tremendous burden on the user. Because of the overwhelming number of planning choices implicit in 3-dimensional treatment planning the approach at most centers (including our own) is to consider plans that are only a minor modification of tried-and-true 2-dimensional plans. We are, in effect, reimposing the constraints of 2-dimensional planning so as to avoid being overwhelmed with treatment choices.

Algorithmic methods to optimize treatment plans are being studied at a number of centers. Several different approaches are under investigation, but all of them involve having a computer calculate or find a treatment plan that meets pre-specified criteria (5). The planning system under VISTAnet is different in that a human being is doing the search; therefore, we feel that VISTAnet offers another, possibly complementary approach to the optimization problem. Although it may be the case that an algorithmic approach will prove more effective than a human using the technology and resources provided by VISTAnet we feel it likely that elements of both approaches will be present in future treatment planning systems.

Up to this point we have used the term "optimal plan" without defining it. Clinically proven measures of plan quality (figures of merit) do not yet exist, in part because of the scarcity of data on which to build these models. For this reason VISTAnet is display oriented so as to allow a subjective evaluation of treatment plans. At the same time, real-time dose volume histograms, and TCP and NTCP models have been included to provide the user with the widest possible range of quantitative feedback.

The relative merits of the real-time, display oriented approach of VISTAnet has not yet been explored. Once the gigabit network is in place (expected by September, 1992), several treatment sites will be investigated exhaustively, with nasopharynx and lung being the most likely candidates. The plan is to set up a number of generically placed tumors and use VISTAnet to study a large number of reasonable combinations of radiation beams and modifiers. If successful, we may be able to define "classes" of 3-dimensional setups for these areas. Future treatment planning for patients with these tumors could then be reduced to consideration of a few minor variations of these plans.

The hardware and software making up VISTAnet is not clinically practical at this time. Reproducing the hardware alone would cost more than $20 million. VISTAnet is a state-of-the-art research program. However, given the pace of improvement in computer hardware, it is not unreasonable to expect that the cost of real-time radiation treatment planning could fall to reasonable levels in this decade.

The VISTAnet project has produced a number of useful techniques. In the first 18 months of the project we have developed accurate radiation dose code that can provide a complete 3-dimensional plan in a few minutes on a low cost workstation. We have developed the concept of computed fluoroscopy and believe that a version of it will one day be available in all 3-dimensional treatment planning software. We are well on our way to having adequate 3-dimensional dose displays via volume rendering. These are visualizations that have pushed the state-of-the-art in computer graphics. By the end of the project we will have demonstrated how multiple computers, linked with a gigabit network can solve a problem (real-time radiation treatment planning) that no single computer can. This effort, then contributes to the concept of "meta-computing" wherein different, specialized types of computers can be linked together by high speed networks. And finally, we have demonstrated that medical imaging in general, and radiation therapy treatment planning in particular is a worthy application of gigabit networks, supercomputers, and advanced computer science concepts.

REFERENCES

1. Cullip, T. J.; Sherouse, G. W.; Symon, J. R.; Rosenman, J. G.; Chaney, E. L. Methods for increasing the speed of computation of 3D dose distributions using a differential SAR algorithm. Med. Phys. 18:624;1991.
2. Cullip, T. J.; Symon, J. R.; Rosenman, J. G.; Chaney, E. L. Computed fluoroscopy in virtual simulation. Med. Phys. 18:622;1991.
3. Fuchs, H.; Poulton, J.; Eyles, J.; Greer, T.; Goldfeather, J.; Ellsworth, D.; Molnar, S.; Turk, G.; Tebbs, B.; Israel, L. A heterogeneous multiprocessor graphics system using processor-enhanced memories. Comp. Graph. 23:79–88;1989.
4. Niemierko, A.; Goitein, M. The use of variable grid spacing to accelerate dose calculations. Med. Phys. 16:357–365;1989.
5. Niemierko, A.; Urie, M.; Goitein, M. Optimization of 3D radiation therapy with both physical and biological end points and constraints. Int. J. Radiat. Oncol. Biol. Phys. 23:99–108;1992.
6. Sherouse, G. W.; Chaney, E. L. The portable virtual simulator. Int. J. Radiat. Oncol. Biol. Phys. 21:475–482;1991.
7. Sherouse, G. W.; Novins, K.; Chaney, E. L. Computation of digitally reconstructed radiographs for use in radiotherapy treatment design. Int. J. Radiat. Oncol. Biol. Phys. 18:651–658;1990.
8. Stix, G. Gigabit connection. Scientific Am. 263:118–120;1990.