CHAPTER 10

Challenges of Large Datasets: Commentary from the Vendors' Perspective

Robert Taylor, PhD

This chapter's title speaks to challenges and large volumes of information. And that is really the essence of being a medical imaging equipment vendor today. We are faced with huge technological obstacles to overcome if we are to truly realize the promise and potential that recent advances in digital imaging have uncovered, and that is exactly the challenge relished by various companies engaged in medical imaging equipment research and development. We are racing toward the answers, shoulder to shoulder in competition. The race has been building up speed over recent years, and it is now truly a thrilling time to be engaged in breaking down some of the walls that can unleash a true revolution in the medical imaging world.

Radiology today is turning a corner on the road that leads from the old analog world to the new digital era. I often wonder how these few years in radiology's history will be regarded from the distant future. This is a time when almost all imaging modalities are available in a digital form, yet film is still widely used for diagnosis and working paradigms based on film are still being applied, despite the fact there is a good chance that they are no longer really the best approach. The most advanced imaging systems, such as modern computed tomography (CT) and magnetic resonance (MR) scanners, are home to a locked treasure trove of new procedures, new diagnostic principles, and higher standards in patient comfort, prognosis, and care. Presumably, the radiologists of the future will think back to the days when medical imaging systems with all the capabilities that will become standards of care in the future were used at a fraction of their potential. I wonder if our ignorance of their true capability will appear to be a big disservice to the patient community. Perhaps it will, but of course there is a great luxury in 20–20 hindsight. At least for now, we vendors are fully occupied working to unlock those new procedures and ways of working, in concert with radiology professionals at institutions around the world.

A BRIEF HISTORY OF RADIOLOGY

It has been little more than a century since radiology was born with Röntgen's exposure of photosensitive media with an absorption image of the human hand. Since that time, a whole industry has evolved using this technology for its medical benefits to help with diagnosis of disease and planning of therapy. It provides a helpful perspective to understand that the commonplace radiography room in a modern hospital is really in essence not so different from its 1900 counterpart, although other parts of the hospital, of course, house far more advanced imaging machines, rapidly developed to take advantage of huge leaps forward in computer technology and medical imaging physics made in the 1970s, 1980s, and 1990s.

The headlines in the closing decades of the 20th century were clearly claimed by the modalities themselves, with CT and MR scanners and advanced x-ray and ultrasound devices leaping forward in sophistication, image quality, and throughput. Now, in the new millennium, there is sometimes a feeling that it is the morning after the big party and we have rather an imposing mess to clear up. The thousands of images and wide variety of acquisition possibilities that have been thrust upon the radiology community have produced a whirlwind of confusion where there are plenty of roads but very few accurate maps showing how to navigate them. This is where our greatest future challenge will lie—in shaping these new powerful modalities into elegant, tested, and truly beneficial diagnostic and clinical applications for radiologists, physicians, and the patients in their care.

THE ROLE OF MEDICAL IMAGING TECHNOLOGY COMPANIES

In order for a company to survive and grow in a highly competitive industry (as any industry worth participating in is bound to be), it is essential that the company offer something different and valuable over its competition. This is especially true if that company is a smaller, newer corporation, because the competition will be large, well-resourced corporations with billions in consolidated revenues and strategic marketing budgets to match. This has given rise to a crop of highly motivated, highly mobile, and innovative companies, seeking to capture the imagination of the medical imaging market with special technologies and with a well-defined focus of corporate attention on some of the specific key problems that radiologists face every day. In essence, one of the problems with being a big company is, in fact, not being small. Small comes with the benefit of quick, decisive decision making, the ability to U-turn on a dime when wrong moves are made, and the ability to make big technological leaps forward by focusing the best talent right at the point where the problems to be solved are at their most evident—in the hands of the medical imaging professional.

So the handful of medical imaging companies working on the challenges of these large datasets generated by modern scanning modalities are each building valuable resources, both in intellectual property and in teams of people who accumulate the collective knowledge and experience required to speak with authority about the most effective and economical way to apply these exciting new tools. Naturally, at the core of such a company will be a strong team of professionals who fully understand how to shape and mold the technology that is behind these new devices and solutions. That means creative scientists and engineers skilled and familiar with software development, hardware engineering, x-ray generation and detection, and networks—local, wide-area, wired, and wireless.

These professionals do not necessarily have the benefit of wide experience or understanding of the ways in which medical professionals will try to apply the technologies they are developing. To bridge this gap between engineering and radiology, companies build in a team of specialized professionals who speak both languages sufficiently well to enable the company to organize itself and prepare for growth. The company's management should include these specialists, because this group is key to ensuring that the needs of radiology are properly addressed by the products and services the company develops. The information they learn and the experience they accumulate in the field can move medical professionals toward the realization that fundamental shifts in direction are required. This must be communicated to the very top of the organization as soon as possible and not lost in the frustrations of field employees who would otherwise end up dealing with short-term issues on a daily basis as the company loses its way.

Hence, medical imaging vendors have developed a spectrum of roles, from fundamental research and development into what will become the underlying technologies for future products, to consultants who help medical imaging professionals understand and overcome the obstacles between where they are today and where they want to be tomorrow, to clinical specialists who show the direct application of today's advanced technologies to the everyday problems faced by the vendor's customers, to "evangelists" and educators who cultivate understanding and stimulate the imagination of their future customers to generate the interest that turns into revenues, which become profits, which can be reinvested to keep pushing the envelope of what's possible.

TECHNOLOGICAL CHALLENGES OF MANAGING LARGE DATASETS

Many in the medical imaging world would agree that the single biggest problem facing radiology today is how to manage the huge volume of information generated by modern scanners. This is a key question, because it is in this deluge of information that the future lies: more accurate diagnosis, detection of more subtle disease, less invasive diagnosis and therapy for more patients—the list is long. However, the recently supercharged CT and MR scanners are being driven on roads that were designed for the horse and cart—in other words, radiographic film.

In recent years, vendors have offered picture archiving and communications systems (PACS) as the answer to these problems, touting the technology as a 12-lane freeway, but that is not a fair analogy. PACS has taken the dirt road of analog radiography and laid a two-lane blacktop, so that the earlier digital modalities can get from A to B. But the roads are still winding and narrow, the traffic is heavy, and PACS cannot claim to have really addressed the issues at the core of the problem.

This is because PACS takes the analog paradigm of film and transplants it directly into the digital world, assuming that computers and networks will just make things fast enough and sufficiently less labor intensive that all the problems will go away. Of course, we all now know that this is not the case at all.

The real power of a modern CT scanner is not realized by providing a system that lets the viewer look at slice images on a screen. When a single CT angiogram (CTA) can produce 2,000 slices of data representing more than 100 screens full of images at any useful resolution, the result is nothing but headache for the busy radiologist. The real power of a modern digital modality can only be realized by abandoning the slice-based paradigm entirely and taking the plunge into a volumetric world, where the scan is a volume of information that covers the anatomy in question, not just a set of arbitrary sections through it.

Unfortunately, the first foundation stones of PACS design were laid with the presumption that the most important task at hand was to emulate and streamline the analog radiographic workflow. As a result, volumetric processing and review are somewhat alien concepts to PACS vendors and fit poorly if at all. Moreover, a leap of faith into a new world of volumetric imaging forces the radiologist to abandon much of the experience and wealth of references that radiology has built up over more than a century in a world of two-dimensional (2D) imaging…or does it?

The truth is that volumetric rendering should be a navigation opportunity that allows the diagnosing radiologist or planning surgeon to choose the views they want—to select the correct 2D cut or curved plane that most efficiently and accurately represents the morphology or finding in question. Once used in this way, all the cross-sectional knowledge base of radiology today can be brought to focus on the phenomenally detailed volumetric information collected by modern scanners.

So what's stopping the PACS companies from just writing a little software to provide volumetric review and processing? There's a very good answer to that question.

WHAT'S UNDER THE HOOD?

The true challenge to managing volumetric datasets lies in the power of the engine. In the 1980s, demonstrations were shown of 3D volume rendering that now appear crude. Each of these images took a huge amount of time to calculate. Of course, computer power through the late 20th century increased at a phenomenal rate, and most of us assumed it was just going to be a matter of time before any imaginable task could be performed by a computer no larger than a pocket calculator.

The reality has turned out to be somewhat different. It is true that computer technology and processing power have surged forward tremendously, but what we have seen is the greatest concentration of research and intellectual prowess on the development of extremely high-performance, general-purpose processors. This means that if you are in the business of using spreadsheets and word processors or playing computer games, you have a multibillion dollar industry working day and night to satisfy your every need and remove every new frustration that their last software release brought into your life.

The problem for those of us in the world of medical imaging is that, unfortunately, our teenage children are not helping us out at all. If third-person 3D computer games required volumetric 3D rendering, then medical imaging would be in a different place today. But as it turned out, we could not use the commodity processors or specialized 3D graphics systems developed for the consumer market, and we have had to take on the challenge of developing our own advanced hardware and software solutions to make volumetric medical image review a reality.

It is a sobering thought that the world market for all medical imaging in 2003 was estimated around $16 billion, according to market analysts Frost and Sullivan. That includes all the CT, MR, positron emission tomography, ultrasound, all the radiography and fluoroscopy rooms, all the angiography suites… and it does not leave much room for investment in semiconductor technology and software to help us manipulate medical images in the volumetric form in which many are acquired. Compare that with annual revenues for Intel alone of $30 billion over the last 12 months, and suddenly medical imaging is thrown into perspective. If we want to drag medical imaging into the 21st century, we are going to have to do it ourselves. We cannot cut and paste ourselves out of this one.

SO HOW BIG AN ENGINE DO WE REALLY NEED?

The technological challenge of providing real-time interaction with volumetric datasets has attracted some of the best minds in computer graphics for several decades now, because it is such a demanding application. In the mid 1980s, some of the earliest 3D images were produced from medical imaging data. Back then, the datasets that were being manipulated were tiny in comparison with those being manipulated today, so it is a good thing that processing power has advanced so rapidly. Not only was there a dire need for more performance (several seconds or minutes of calculation per frame was never going to be acceptable), but the resolution of volumetric acquisitions like CT and MR was growing rapidly, meaning more information had to be processed to generate a diagnostically acceptable image.

Part of the problem with working in 3D is that if the scan's dimensions are enlarged by just 25%, the amount of information that must be processed almost doubles, because that 25% is bigger in all three directions. As we moved through the 1990s, CT acquisitions grew from simple scans of the head or lumbar spine, with perhaps 15, 20, or 50 slices, to higher-resolution CTA examinations consisting of hundreds of images, each consisting of 512 × 512 pixels with a possible range of 4,000 values each. The early years of the 21st century have seen the advent of multiphase cardiac CTA examinations with as many as 3,000 images, and even more commonplace and routine, chest–abdomen–pelvis plus run-off CTA examinations with 1,500–2,000 images each (Fig. 10.1).

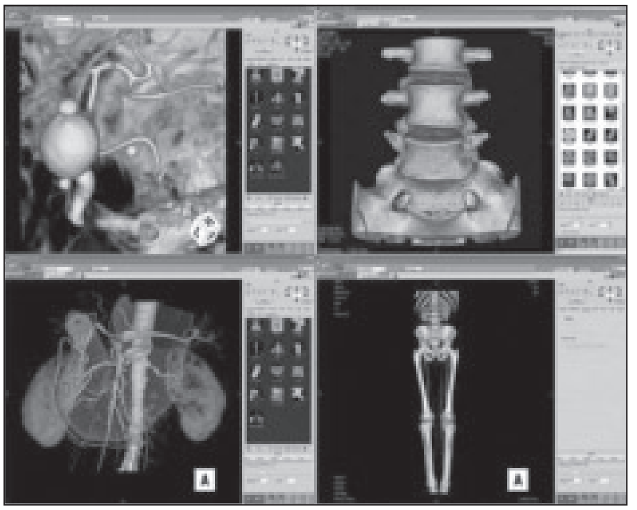

Figure 10.1 Four screenshots showing a cerebral aneurysm (top left), a lumbar spine with Parr's fracture (top right), a computed tomography angiogram of the upper abdominal vasculature and stomach (bottom left) and a computed tomography angiogram run-off examination from the thorax to the ankles. (AquariusNET Client software image courtesy of TeraRecon, Inc.)

Companies producing 3D review tools quickly realized that although pretty 3D pictures are nice, widespread acceptance of volumetric review was clearly going to be severely limited unless the equipment could keep up with the fast-moving mind of a highly specialized professional: the radiologist. In an environment in which images are read and processed in a fraction of a second, the image processing system that produces the 3D reconstructions or cross-sectional planes must operate extremely fast, and it has to be highly responsive to operator input. It became clear that if radiologists were to adopt volumetric review as a standard of care, the 3D vendors had to find a way to deliver fast, intuitive, and streamlined applications into the hands of reading physicians, so that they could effortlessly find the best view to illustrate the anatomy or pathology under review, without having to wait for progress bars to march across the screen or images to update at a rate far slower than the rate at which the specialized brain can process information or deliver instructions about the next view.

Upon analyzing this requirement, three distinct areas can be identified that need attention. First, the volumetric processing system must have a high-performance engine. Image rendering must be performed in real time from datasets as large as 3,000 CT images in a manner highly responsive to operator input. Second, the user interface must be carefully designed to deliver intuitive and responsive tools to allow effortless manipulation of images and oblique planes until the correct view is obtained. Finally, and perhaps most important, the volumetric review capability must be delivered directly into the hands of the reading radiologist. It will never be acceptable to ask a radiologist to work at two places, one of which they may have to share with other radiologists, to cover the range of studies on which they are working, some of which will and some of whch will not require 3D. Vendors have been waking up to the challenge of delivering volumetric review directly into the existing diagnostic reading workflow.

RENDERING

Rendering is one of the magic words of volumetric review. It is a computer graphics term used to describe the process of taking a collection of image data in one form (for example, a stack of CT or MR slices) and then applying some kind of image processing technique to generate a single output image from what may be a large collection of input image data. Rendering engines in software or hardware typically are used to provide image output in an interactive fashion. When the operator sees the first rendered image, he or she may determine that a view from a different direction would be helpful. By dragging a mouse, a sequence of commands is sent to move the viewpoint to the rendering engine. The rendering engine responds and calculates new images by applying the same image processing step with different parameters to the same input data, and an interactive volumetric review tool is achieved (Fig. 10.2).

For medical imaging data such as stacks of CT and MR slices, there are several ways to solve the raw 3D rendering problem. Many years ago, computer power was only sufficient for one "surface" to be selected from the volumetric dataset (for example, the surface of all bones in a CT scan, determined by a single threshold value). After some lengthy preprocessing, a 3D "model" could be constructed that could then be manipulated in something like real time (if the scan was not too large). This approach was referred to as shaded surface display and, in essence, was similar to the 3D graphics principles used for most computer games. However, so much information is discarded using this technique (that is, anything not related to the surface that was selected), that as soon as computer power was sufficiently potent to manage a higher level of rendering, it was rapidly adopted. The next level of processing is volume rendering. This technique does not discard any information from the original dataset and, if implemented well, does not require any preprocessing. Every single volume element (or "voxel") in the acquired image data is simply assigned a color and transparency, and then the computer calculates a projection through the entire dataset from a given point of view. The result is a far more detailed and life-like rendition of the scan data, and this rapidly became the only standard medical imaging professionals would accept.

However, volume rendering is a very demanding calculation. For example, a 2,000-slice CT dataset contains half a billion sample points. To fully render this image, tens of calculations must be performed per sample point. Assuming we could use some shortcuts to get this down to 10 calculations per point, that would still be 5 billion calculations to generate a single 3D image. Unfortunately for us, despite the billions of dollars Intel has invested in their chips, there is not much in the world of spreadsheets and word processors to generate a demand for that kind of raw processing performance, and commodity graphics processors are just not up to the task. In browsing the Internet, I found a Web site where some computer wizards were raving about how they had achieved 5 gigaflops by using 10 liters of liquid nitrogen to cool an Intel processor to nearly –200°C.

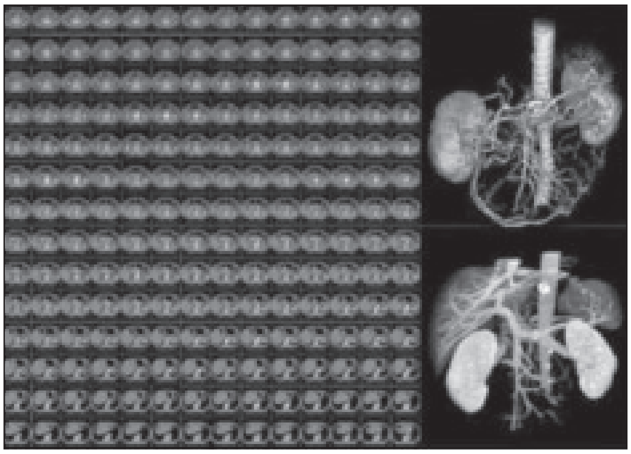

Figure 10.2 Rendering is the process by which a computer system takes a stack of slice images (left) generated by a volumetric acquisition modality, such as CT or MR, and converts them to a 3D volume-rendered representation (top right) using various look-up tables to convert source data values to colors and transparency or to a maximum intensity projection (bottom right) by projecting through the volumetric stack of slices in a given direction and recording only the brightest pixel encountered along each ray. (AquariusNET Client software image courtesy of TeraRecon, Inc.)

That superpowered processor would be able to manage only one image per second in our previous example, and, in reality, it's never as clean cut as that, so that it could be expected to be slower by another factor of 10.

Current commodity computer hardware is clearly underpowered for the task of volumetric image review. Because the semiconductor research dollars are driving development in the wrong direction (from our point of view), 3D medical imaging companies have all adopted their own approaches to working around this obstacle.

The first approach is to accept that there is at least a factor of 10 too little power in the processors available today and to try to streamline the problem to fit what is available. The classic example of this is the software approach that uses only a sparsely sampled subset of the input image data to calculate an image with enough performance for interactivity. Of course, it looks terrible, and no medical practitioner would wish to diagnose from it. But that's OK, proponents say, because if you let go of the mouse and wait, the fine-detail image will be calculated in the background and delivered on the screen. This is the "software only" approach, and it is not ideal. The nice thing is that it can run on almost any computer if the vendor designs the product carefully. The only drawback is its limited utility. And the description "any computer" should be qualified—any computer with two high-performance, expensive processors and lots of memory. Without this it will be difficult to manipulate large datasets.

The next approach is to try to leverage commodity graphics hardware. Millions of dollars have been invested in 3D graphics hardware, primarily for the mass consumer market of computer games. They have come up with some interesting square pegs for us to try to fit into the round hole that is medical imaging. Some companies have taken on the challenge of trying to adapt these commodity graphics boards, with some limited success to date. Because the graphics cards are being used for a task far outside their specification, they too are underpowered. Companies have invented all kinds of shortcuts and tricks to try to get every last drop of performance without sacrificing image quality. This approach also comes with an inherent drawback. Even more than with the "software only" approach, the company must specify (and probably provide) the hardware at which the radiologist will be sitting to perform the manipulation, which means that the critical obstacle of delivering volumetric review into the existing diagnostic workflow has not been overcome.

These two examples lead to two conclusions. First, available hardware is underpowered, and we must make some kind of dedicated system to deliver fully interactive 3D rendering for the typical medical imaging datasets of today and tomorrow. Second, we cannot place an unreasonable constraint on the hardware specification for the diagnosing radiologist. We need to be able to install a system at any typical hospital and take advantage of the existing computer equipment and infrastructure. Although it may seem at first that these two requirements are contradictory, in fact, they are not, when client–server networking is considered in the equation. This is where the most elegant innovation in recent medical image management comes in: the concept of a hugely powerful server, with sufficient processing resources for an entire enterprise's 3D needs, streaming interactive 3D images to multiple client computers across the hospital network, running a simple client application that can be installed on any regular personal computer (PC) in the diagnostic reading workflow, such as the PACS workstation itself or the dictation PC.

Of course, to have such a processing server with enough rendering power for the whole enterprise, we must create a 3D rendering engine that is up to the task. We have already established that the most powerful dual processor workstation today, equipped with several thousand dollars of commodity graphics hardware, can perhaps just about satisfy the demands of one radiologist, if that radiologist is prepared to go and sit at this heavy, dedicated workstation. How can we possibly conceive of a server that will have more than 10 times this 3D rendering power, all in one box?

ONE SOLUTION

At TeraRecon, Inc., of which I am chief operating officer and executive vice president, we have devised one set of solutions that I will describe here. These are by no means the only possible solutions, but the ones that we think address the demands outlined previously in this chapter. Our technology has been in development since the mid 1990s, when it was originally conceived with the intention of delivering a PC-based real-time volume-rendering system for medical imaging. Now in its second generation, our technology has realized in a single dedicated chip the power required to render modern volumetric datasets in real time. To learn more about how it works with our specialized server, readers are invited to our Web site at www.terarecon.com.

ENTERPRISE DISTRIBUTION

Having identified and addressed the formidable challenges that exist in getting a sufficiently powerful engine in place for the volumetric review task, it is clear that the next obstacle to be overcome is enterprise distribution. If volumetric review is to really take hold, it must be delivered directly into the diagnostic reading workflow, so that radiologists do not have to move from one location to another to access this new approach to managing huge datasets. Perhaps one day volumetric review will replace conventional diagnostic reading, but to reach that day, we must first live through an era in which traditional means are still available for reference, and, more important, all the other things a radiologist needs to do are properly addressed. That might mean reading chest x-rays from high-brightness displays with linearized grayscales (something volumetric review of CT or MR does not call for). For some time to come at least, we must operate within the world of PACS-based radiology as it is today.

So how best to integrate volumetric review into the diagnostic PACS reading station? For the "software only" and "specific hardware required" crowd, the options are limited. A sufficiently powerful PACS workstation must exist in the first place, possibly with the special commodity graphics card that the 3D company's product uses, and then this PACS station could double as a high-end 3D workstation. However, there are some considerable logistics to work out in terms of how the work will flow.

PACS have evolved to provide the reading radiologist with access that is as rapid as possible to the diagnostic images in a CT or MR scan. Because radiologists typically look at only one or a few images at a time, this means that PACS have been specially designed in recent years to fetch only the initial few images to the display station from the server, so that the transition from worklist to display is as rapid as possible. Then, in the background, other images are loaded as needed as efficiently as possible.

Unfortunately, this way of organizing information flow is not conducive to 3D processing locally at the workstation. To make one sagittal reconstruction from an axially acquired dataset with such "thick client" technology, every single image from the examination must be present on the display workstation running the 3D application (because a little piece of each slice image contributes to the sagittal reconstruction). The same is true for basically any other volumetric view, except for the pure axial image originally delivered. This means two things. First, the dataset must be pushed or pulled in its entirety to the display workstation before it can be used (a huge step backward in PACS technology), and second, additional delays may arise while the 3D application organizes the data in the special internal format required to process it.

However, with a client–server approach, a far more efficient workflow can be realized. Upon acquisition, images need only be routed to one location (the client–server) in order for the 3D rendering needs of the entire enterprise to be taken into account. At each display workstation, the client software is used to access the 3D resource, so that basically no minimum specification (by the standards of the last several years) is needed for the PACS workstation itself or whatever adjunct PC is used for displaying the images. When an examination is selected for review from the PACS worklist, while the PACS is organizing the first few images for display in their "native" format, the client will automatically notify the server which dataset will be worked on, and the server will immediately load the entire examination into its memory, ready for 3D rendering at a moment's notice. This loading can happen in seconds or less, even for very large datasets, because the transfer is from high-performance RAID disk arrays via the server's internal bus, which runs orders of magnitude faster than the hospital network that other technologies must use to summon a copy of the scan.

An initial 3D reconstruction, for example, is automatically performed and displayed on the PACS station. When the radiologist elects to manipulate this volumetric representation, the interactive streaming between client and server begins, with near-instantaneous calculation of new 3D images and reformats, for many different 3D sessions across the enterprise, concurrently. The radiologist experiences a real-time interactive volumetric review experience, without any need for the PACS to have prior knowledge of which radiologist will be reading at which PACS display station or which cases they will want to manipulate in a volumetric fashion.

CLINICAL RELEVANCE: THE IMPORTANCE OF APPLICATION DEVELOPMENT

Volumetric review is a new frontier, with its own set of challenges. In its most basic form, we can envision a 3D volume-rendered image of bones or blood vessels, but this does not get us very far in practice. Three-dimensional images are impressive, but so much more information is in the stack of slices that no radiologist would be comfortable abandoning the full volumetric information for a simple 3D reconstruction.

As 3D processing tools were initially developed and found their way into hospitals, they originally arrived more or less packaged with CT and MR scanners, as these modality vendors sought to take advantage of the ever-increasing performance of the scanners by tapping into new applications. Some of the earliest and most obvious applications were in orthopedic CT and MR examinations, replacing or complementing conventional radiography. More recently, the use of CT and MR in angiography has expanded rapidly. It is easy to understand why, because a modern CT scanner can perform a CTA in 20 seconds, which provides as much information if not more than an invasive catheter angiography, and all with a simple peripheral vein injection of dye.

As these new applications for CT and MR developed, the "heavy" stand-alone workstations provided by the modality vendors developed a suite of tools to facilitate their management. Three-dimensional volume-rendering tools were used to render orthopedic trauma or degenerative orthopedic conditions, and special segmentation algorithms were developed to remove the bones from CTA studies, allowing maximum intensity projection (MIP) and volume-rendered views more similar to the catheter angiogram that CTA seeks to displace. The list did not stop there, with a whole suite of new applications vying for the limelight and helping CT and MR vendors to extol the virtues of their ever improving systems. Time–density analysis for CT brain perfusion, segmentation, analysis and tracking of pulmonary nodules, coronary calcium scoring, virtual endoscopy, coronary CTA, and MR angio-graphic (MRA) cardiac functional analysis are only a few of these.

However, a common theme running through these applications is that they reside on a dedicated stand-alone workstation and, as such, have tended to be designed either for a dedicated radiologist who has the time and inclination to learn complex computer software or, more commonly, for a dedicated 3D technologist who stays at the workstation and performs these 3D workups for radiologists to review later from their PACS workstation. The latter solution has been relatively successful and workable for some time, with some institutions even setting up dedicated 3D laboratories with multiple workstations and technologists. In the end, however, it requires additional staff, additional workspace, and does not bring to the diagnosing radiologist's fingertips full control of the volumetric acquisition.

If diagnostic radiology and the greater health care enterprise are to truly enjoy the benefits of advanced volumetric review, new tools must provide fast, intuitive review procedures that take no more than 1 or 2 minutes as part of the normal diagnostic workflow. This is what radiologists have been demanding, and this is now what the vendors are racing to provide.

WORKSTATION VS. WORKFLOW

The difference between the workflow-centric tool as opposed to the workstation-centric tool is becoming clear as applications develop. When reviewing a CTA of the abdomen and thorax, for example, the radiologist needs to review all the acquired information to check for incidental findings, but the main point of the examination is review of the blood vessels, enhanced with dye to make them stand out as bright areas in the CT acquisition. Although removal of the bones and rendering of the examination in a manner similar to an angiogram is possible with today's technology, when automated it is an imperfect process that tends to work best on patients who do not have much wrong with them—a fact that rules out this approach, at least for now. However, the same abdomen and thorax CT examination, which may contain hundreds of images, can be reviewed swiftly with the use of coronal, sagittal, or oblique MIPs that capture extended sections of blood vessels following curved paths (because of the limited thickness of the slab deployed) and which can be directly controlled in real time by the interpreting radiologist, who can steer the volumetric review process directly to areas that attract attention and warrant more detailed workup (something that preprocessing by a radiologic technologist precludes). In addition, a volumetric tool should recognize the fourth dimension—time—as well as the three spatial dimensions.

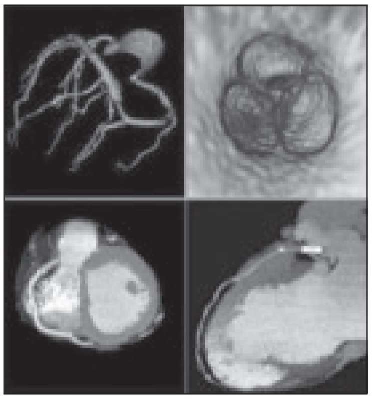

Figure 10.3 The advent of faster and more powerful multidetector CT scanners has made possible CT angiography (CTA) examination of the beating heart. However, capturing a dynamic organ in motion generates not just a single high-resolution scan but several. A typical coronary CTA consists of 300 slices per cardiac phase over 10 cardiac phases, for a total examination size of 3,000 slices, which must be reviewed volumetrically with time-dependent tools. Segmented coronary arteries (top left), aortic valve (top right), maximum intensity projection views of the right coronary artery (bottom left), and left anterior descending artery, with stent (bottom right). (AquariusNET Client software image courtesy of TeraRecon, Inc.)

In the diagnostic workflow, once an appropriate volumetric representation has been generated to illustrate a finding, images should be captured and automatically filed back into the PACS as a new series under the original patient study, to act as a record that volumetric processing was performed (potentially for support of a reimbursement claim) and also to be accessible by clinicians in the wider enterprise. With this few-minute procedure complete, the radiologist is free to move on the next examination without having to wait for reconstructions to be performed by a technologist, a delay and inconvenience that often precludes the use of volumetric processing if workstation-based tools are the only ones available.

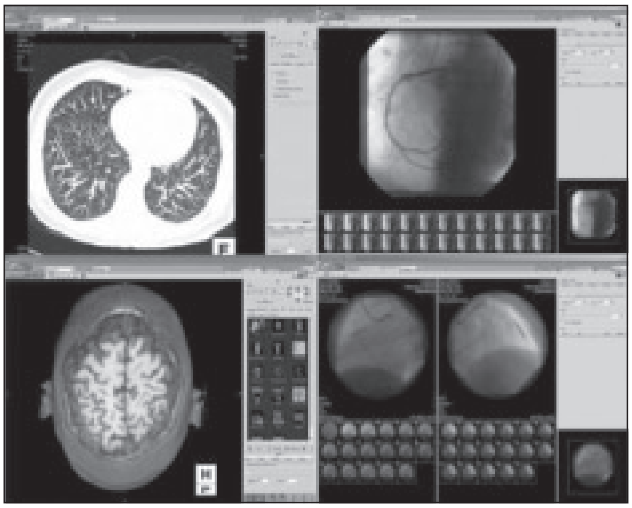

Figure 10.4 Images from various medical imaging modalities displayed on a thin-client software application. Angiography and biplane angiography (top and bottom right), pulmonary CT angiogram in axial maximum intensity projection (top left), and a volume-rendered representation of an MR image of the head (bottom left). (AquariusNET Client software image courtesy of TeraRecon, Inc.)

Diagnostic reading of MR examinations also benefits greatly from deployment of sufficiently powerful and flexible "in-the-workflow" volumetric review tools. A great advantage of MR is its ability to acquire entire volumes in rapid succession, allowing for a detailed time evolution to be obtained. This technology has been used to study the time-dependent behavior of contrast uptake (for example, to detect tumors that rapidly enhance as a result of high vascularization) and also to study the beating heart in motion (Fig. 10.3). Through application of these techniques, time-dependent MRA examinations can rival multidetector CT examinations for size, and similarly powerful and 4D volumetric review tools are required.

The high-powered thin-client approach is well-equipped to bring review of these examinations (that were once time consuming at best and requiring dedicated workstations at worst) directly into the reading workflow (Fig. 10.4, 10.5). Time-dependent cardiac examinations can be reviewed in any oblique plane, with constant cycling through all phases so that the heart "beats" on the screen during reading. Multiphase MR angiography examinations for detecting dye uptake by a tumor can be reviewed with a tool that allows the precontrast image to be subtracted from the postenhancement images in real time, at the press of a button, to make up for the trade-off in MR that must be made between high acquisition repetition rate and fat suppression techniques.

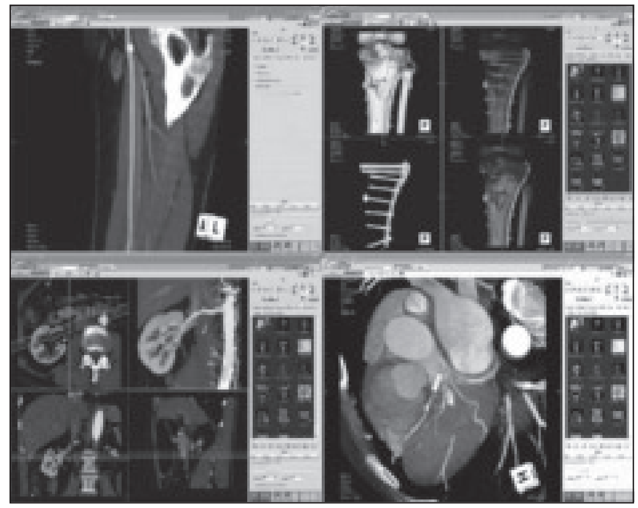

Figure 10.5 Sample screenshot images from thin-client software running on a regular PC without remarkable memory resources, processing power, or graphics capability. The technology allows real-time manipulation of volumetric images residing on the remote server. Three-dimensional, maximum intensity projection (MIP), and MPR images are calculated server-side, responsive to operator input on the client side, and streamed across the network for display on the client. Images shown: femoral artery (top left) in MIP mode, isolated from a run-off CT angiography acquisition with 1,500 slices; various volume-rendered views of a fractured tibial plateau (top right), postsurgery; MIP visualization of the renal artery (bottom left); and a volume rendered image of the coronary arteries (bottom right), revealing a stent and calcified plaque in the proximal left anterior descending artery. (AquariusNET Client software image courtesy of TeraRecon, Inc.)

In the orthopedic and trauma context, joints can be rapidly disarticulated. If so desired, interactive models of the disarticulated components can be delivered directly to the orthopedic surgeon's computer, because the same client technology used by radiologists to diagnose can be used by clinicians for review and consultation with patients, giving them access to multigigabyte studies out in the enterprise at large, even if available network bandwidth is limited.

Workstation-based processing is not rendered obsolete by the presence of advanced processing in the reading workflow; instead, it takes on an enhanced and complementary role. As much as the workflow-based tool removes the burden of "1–2-minute" review tasks from the dedicated workstation, it frees that workstation to concentrate on more time-consuming and specialized review and postprocessing tasks, which maximizes the value of time spent at the workstation and frees regular diagnostic reading stations from more time-consuming applications so that they can be dedicated to primary diagnostic review. Such applications currently include virtual colonoscopy, advanced vascular measurements for endograft planning, detailed volumetric and distance calculations for tumors and other pathology, and detailed semiautomatic segmentation of bones from complex CTA cases, before rendering in an angiogram-like form for both diagnosing radiologists and referring physicians (Fig. 10.6).

Indeed, a truly enterprise-aware volumetric processing solution will allow advanced workups to be performed on dedicated high-end workstations and then allow for these models to be delivered to the workflow-based volumetric review tools for use in diagnostic and clinical review.

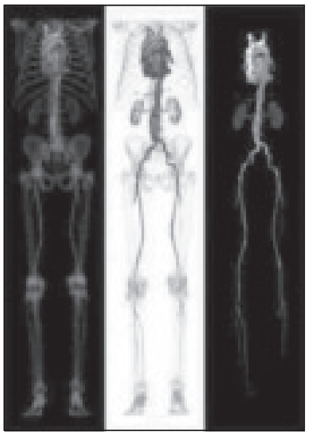

Figure 10.6 Modern multidetector CT scanners can acquire as many as 1,500 slices in 20 seconds or less. This rapid acquisition makes possible CT angiography (CTA) of the entire vascular system from the neck to the toes, as in this CTA chest–abdomen–pelvis examination, with a run-off to the feet. All this information is obtained from a single peripheral vein injection of dye, and 3D technology is then used to render the volumetric data with volume rendering (left and center images) and maximum intensity projection (right image). The bones occupy the same signal range as the blood vessels, so advanced segmentation tools have been used to extract the bone and reintroduce it as a semi-transparent anatomical landmark. In many cases, this can be achieved with just a few clicks of the mouse. (AquariusNET Client software image courtesy of TeraRecon, Inc.)

THE FUTURE

From where we stand today, some clear milestones must be passed along the road to the future, when today's world of simply emulating the analog, film-based workflow will be a thing of the past remembered by radiologists using a suite of intuitive, effortless volumetric review tools that truly allow them to focus all their attention and experience on the task of detecting disease.

Early steps along that road are being taken today, with volumetric tools such as 3D volume rendering, multiplanar reconstruction, and MIP becoming available as a part of PACS reading stations. These implementations today are generally partnerships between specialized 3D companies and the PACS vendors themselves, and as discussed here, different approaches are being pursued, determined by the varied technologies available to the industry participants. Certainly in terms of the power and simplicity required for volumetric image processing, the thin-client approach, driven by extremely powerful dedicated rendering hardware appears to have the upper hand technologically. As the novelty of these relationships matures into deeper integration, complexities and inefficiencies in workflow will be ironed out and protocols streamlined as radiology and the health care enterprise at large gain confidence with new reading and review paradigms.

Today, we are in an era of collaboration and partnership between commercially driven companies seeking to deploy the perfect solution and the radiologists, surgeons, and other health care professionals who ultimately are the ones who will vote for the winning solution with their purchase orders and publications. The key areas of research and development that we should expect to see emerge over coming months and years are those of better human–computer interfaces and computer-aided detection, where software will provide a useful second review of diagnostic images, to ensure against the momentary distractions and "bad days" to which all humans—including diagnostic radiologists—are susceptible.

Technology can also assist by bringing modalities together and presenting the radiologist with a complete set of examinations relevant to a patient and not merely their presenting condition. Each modality will be presented in the most appropriate form, selected automatically by context-aware software algorithms that intuitively understand the likely scenario being investigated, saving the radiologist all the effort and wasted time of collecting different examinations and reports from a variety of modalities and sources, allowing the radiologist or clinician to concentrate on the patient, his or her condition, and the appropriate remedy.

Meanwhile, we should never be complacent about the power of our engines. Scanners are collecting ever more information, and before long it will be easy to imagine a world in which the question is no longer "how many slices." Instead, we will be empowered with a true continuum of diagnostic information, and vendors will be faced with the continual challenge of coming up with ever faster and more powerful technologies for rendering this information in real time. New paradigms of thinking about image data are already taking form, with trailblazing companies seeking to obviate the need for slice-based representations at all and going instead directly from the raw data acquired by CT or MR scanners to volumetric representations (of which “the slice” is just one subset). For this reason, a clear criterion for success must be a company’s commitment to the most powerful and innovative volumetric processing technology and continued investment to ensure that this underlying engine remains sufficiently powerful to address the performance needs of modern health care.

Meanwhile, vendors must continually refine applications to make the task of the medical imaging professional as frustration free as possible. Indeed, it would serve the designers of this technology well to remember that the radiologist’s brain is a highly tuned and efficient instrument for detecting disease. If it is to operate in its most effective and efficient form, all that needs to be delivered by medical imaging equipment is the removal of distractions, obstacles, and inefficiencies that take away attention and break up concentration. The more a radiologist has to struggle to get to the desired representation of the information the examination contains, the less hope there is of getting into the “zone” of total concentration, where all experience and knowledge are brought to bear on the problem at hand in one, parallel, almost unconscious process. When medical imaging equipment has reached that level of efficiency and performance, this most exciting and challenging field of volumetric diagnostic imaging review will truly have reached a maturity from which all physicians and the patients in their care can benefit.