electronic design

**DANIEL HARRIS** | Digital Technology Editor

# Architecting New Dimensions Of Medical Imaging

**Several technologies—like** 4D (3D over time) ultrasound imaging *(Fig. 1)*—have taken the medical-imaging market by storm. The medical field will continue to benefit from Moore's Law as speed and resolution continue to improve. Take for example the joint effort between engineers and scientists from IBM and the Mayo Clinic that seeks to exploit recent parallelism advances in processors such as the Cell.


**1. Philips' iU22 Ultrasound System captured this 3D ultrasound image of a 24-week old fetus.** ***(courtesy of Philips Medical Systems)***

The result is a dramatic acceleration in 3D medical-image processing, which significantly advances the image-fusion process. Also known as registration and overlay, this process creates 3D images by aligning two or more images captured by different devices (e.g., MRI and CT), or the same type of device on different dates. Using alignment algorithms, images are "fused" to provide more complete visual information for easier detection of tissue changes like tumor growth or shrinkage.

But University of Calgary students have taken a different approach in creating the most complete 4D model of a human yet *(Fig. 2)*. Using a joystick, the object-oriented hologram, dubbed CAVEman, can provide a view of up to 3000 distinct body parts. This technology will help physicians plan for complex surgeries and allow patients to see a map of their body before surgery.

**LET'S TALK ABOUT GOALS** • Instead of using the hospital's ICU equipment, patients can be monitored at home. The patient's quality of life improves and medical costs are reduced, achieving two key goals of these new technologies. Another goal is improved accuracy—for example, imaging the heart in a single beat or the lungs in a single breath. Researchers also hope to improve diagnostic capabilities via the least invasive procedure in as close to real time as possible.

Let's not forget about reducing or eliminating false positives and false negatives. Traditional mammograms have a high percentage of false positives, resulting in the unnecessary removal of tissue in far too many patients. Added costs for false positives include evaluation costs, treatment costs (of the observed breast cancer), and the immeasurable emotional cost associated with a false-positive result. Of course, a false negative can be much worse, possibly leading to death. And for the physicians responsible for interpreting these images, the ongoing goal is to increase the potential to find anomalies in organs, tissue, and cells via the least invasive means.

**ARCHITECTING MEDICAL IMAGING SYSTEMS** • These goals imply one ongoing theme for all new or redesigned medical imaging systems: the need for maximum computing power to provide the highest-resolution processed images in the least amount of time. Typically, that means maximizing the number of cores and threads for the target form factor, since many imaging algorithms are parallel-processing friendly.

But before deciding which brand of multicore processor to use, carefully consider the system's scalability and upgradability. Due to jumps in performance and data rates within the semiconductor and storage industries, it's important to be able to drop in the next-generation device or add more nodes to the system (when using clustering) without redesigning and retesting the entire system. If you can get away with only a recompile, you're ahead of the game.

"Scalability of solutions is key to enabling customers' reuse of software and algorithms across products," says Bob Ghaffari, manager of the Medical Segment for Intel. "Having a silicon architecture that can address a variety of performance and power bands ranging from high-end CT equipment down to a low-power portable ultrasound product requires an architecture that can scale."

Ghaffari said Intel is focused on meeting a variety of

## High-Speed Memory Options

| High-speed memories | DDR2 (x16) | GDDR3 (x32) | GDDR4 (x32) | XDR (x16) |
|---|---|---|---|---|
| Data rate (Gbits/s) | 1 | 2.4 | 3.2 | 3.2 |
| Bandwidth/device (Gbytes/s) | 2 | 9.6 | 12.8 | 6.4 |
| Total bandwidth on a 256-bit bus | 32 | 76.8 | 102.4 | 102.4 |
| Number of devices/system | 16 | 8 | 8 | 16 |
| Total memory size (Mbytes) | 1024 | 512 | 512 | 1024 |

**2. The University of Calgary's CAVEman project delivers the most complete 4D medical diagnostic model of a human to date. The hologram can provide a view of up to 3000 distinct body parts.** ***(courtesy of the University of Calgary)***

medical application requirements by providing highly functional and flexible system-level building blocks, thereby minimizing the cost of ownership and significantly accelerating time-to-market.

When attempting to determine just how many cores and threads are needed, try to make the data path the bottleneck, because there's really no point in processing data faster than it can be stored. If a local hard drive will be used, then serial ATA (SATA) or serial attached SCSI may be the limiting factor. Otherwise, if you're writing data to a device on the network, the network connection (Ethernet or wireless) will have a known maximum bandwidth.

Intel suggests some guidelines for choosing the number and type of processors, as well as how to tweak them. First, determine board performance and form-factor criteria. Next, run the code. Continue to optimize code, and stop after a reasonable number of iterations. Then, adjust the performance. If this is adequate, you're done with architecture selection. If greater system performance is required, add external devices for acceleration offload.

Creating 4D images from a spattering of 3D images could require anywhere from 500 Mbytes to 5 Gbytes of data per patient. This is sure to grow as resolution and the number of image slices increase. Factor in the number of patients seen on a daily basis, and a thin-client network that stores all patient data on a fast central server and uses local PCs to display the images starts to look attractive.

But when portability is mandatory, a system based on the MicroTCA architecture may be the best bet. MicroTCA provides a rugged small form factor with lots of compute power, bandwidth, and built-in network connectivity. Meanwhile, the display of 4D images requires several gigabytes of temporary storage. Designers have to consider the amount, type, and speed of graphics memories like Graphics Double Data Rate (GDDR) *(see "High-Speed Memory Drives Visualization" at* www.electronicdesign.com*, Drill Deeper 15793)*.

If you segment out your architecture properly, with an overall goal of designing only what isn't readily available, chances are you're in good shape. The major building blocks include the analog front end, the digital back end, the graphics display renderer, and a system controller with optional networking *(Fig. 3)*.

Data acquisition and image pre-processing make up the analog front end. They rely heavily on the imaging modality, which may require one or more DSPs, FPGAs, or

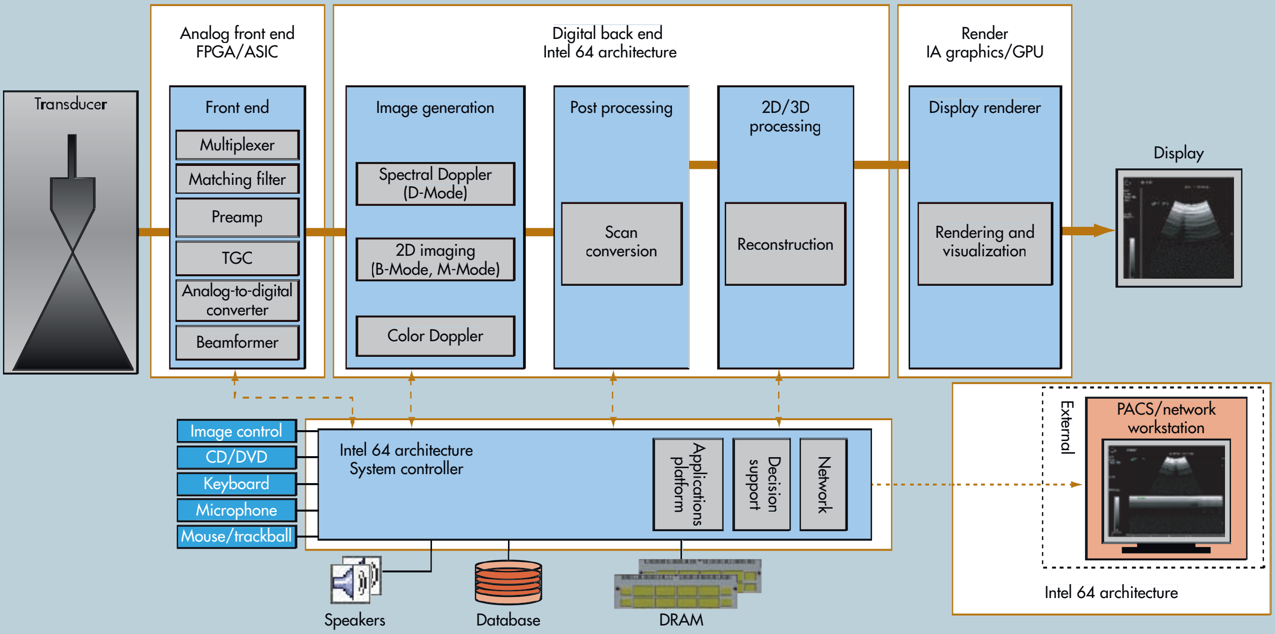


**3. A typical ultrasound system includes several major system blocks and associated sub-blocks, spanning both hardware and software.** ***(courtesy of Intel)***

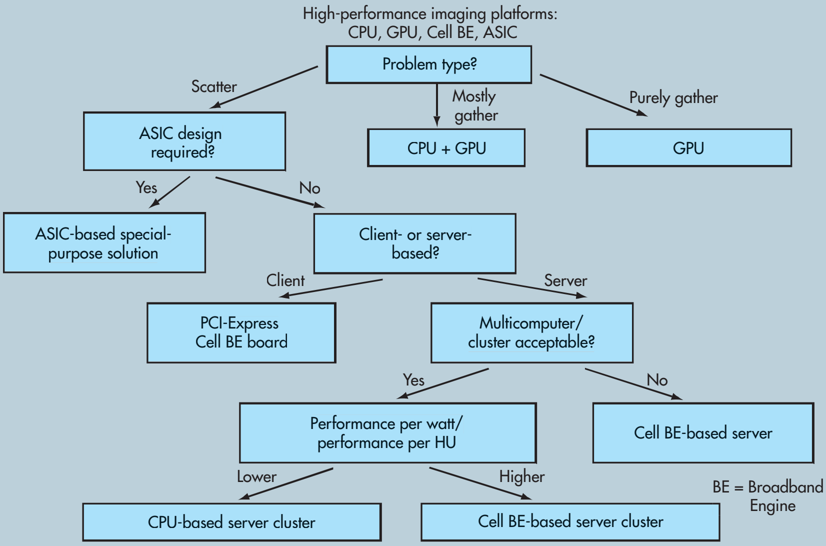


**4. Designers must make several key decisions in determining how much processing power to dedicate to their imaging system.** ***(courtesy of Mercury Computer Systems)***

ASIC ICs. The digital back end includes the image reconstruction and post-processing blocks. Depending on the modality's complexity, this block could be a simple processor (GPU) or one or more advanced processors (CPU and/or GPU) containing multithreading capabilities with multiple cores. For demanding tasks like image processing and reconstruction, when top performance is needed, processors like the Cell Broadband Engine may be more appropriate *(Fig. 4)*.

If your future involves multiple cores, seriously consider software-based decisions, such as the operating system, message passing interface, parallel programming language, and so on. Even otherwise trivial decisions like which type of file system to use become substantially more important and should be made carefully *(see "Parallel Programming And Multicore Environments" at ED Online 14691 and "Multicore My Way" at 14631, both at* www.electronicdesign.com).

TechniScan's UltraSound CT Imaging System produces fully digital breast images based on transmission ultrasound. This type of ultrasound can be used to produce two images of the breast based on both the speed and attenuation of sound *(Fig. 5)*.

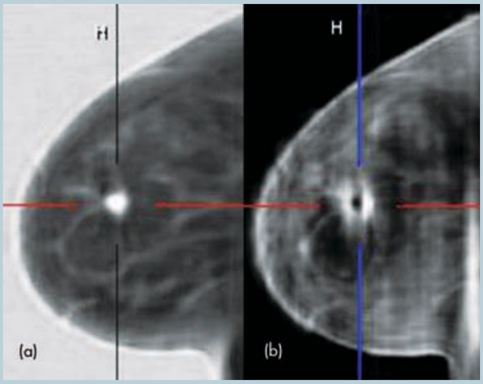


**5. The TechniScan UltraSound CT System's high-resolution whole-breast 3D image lets radiologists view abnormalities based on (a) the speed of sound and (b) attenuation of sound, which provide more accurate localization and characterization of problem areas compared to traditional methods.**

"When a vendor says that they can replace a major component in my system that doubles the performance of the original component and requires the same power and cooling as the original component, I get really interested," says Frank Setinsek, system architect for TechniScan *(see "Advances Trigger An Ultrasonic Boom," ED online 12682)*.

**IMAGING MODALITIES** • Except for X-rays, which are recorded directly on film, all medical imaging modalities use similar basic principles and rely on a similar data flow *(Fig. 6)*. The process starts with the imaging machine building an analog "image." It does so by applying one stimulus or more to the patient (subject) and then recording the response to the stimulus. Then the raw data is usually pre-processed and "scrubbed" to both suppress noise and enhance signal quality.

Next, the pre-processed image is typically reconstructed by converting (e.g., using a Fourier transform) thousands of transmission measurements into a pixel map that makes up a physically meaningful image or volume. The image or volume then is post-processed to improve its appearance and usefulness. The image display may be standalone or a composite built using overlaying images captured with different technologies, like MRI and PET. If slicing techniques were used, the slices may be viewed one at a time or combined for a 3D view.

Finally, computer-aided diagnosis (CAD) may be employed to aid in analysis and interpretation of images. CAD works by using the post-processed data and applying segmentation, followed by feature selection for the regions of interest and feature classification using pattern-recognition algorithms. The physician or radiologist then enters the equation as the final interpreter. After analyzing the images and optionally using historical data as a base for comparison, the physician delivers the diagnosis or update to the patient *(see "Video Processing Brings New Meaning To Motion," ED Online 13291)*.

**ADDITIONAL WEB RESOURCES** • One Web site, *www.rtstudents.com*, was designed with radiology students in mind. This portal to other useful sites also contains a plethora of great links for research, discussion, and resources to aid learning.

GE's Medcyclopaedia includes a medical-imaging encyclopedia, a glossary, and an outstanding interactive e-learning section with a complete anatomy breakdown *(www.medcyclopaedia.com)*. With this site, you'll never get the cerebellum confused with the temporal lobe again; the e-learning module also includes a virtual index-card-by-picture or -by-name learning system for medical-imaging terms.

If you're looking for information on high-performance computing (HPC) using clusters, some helpful Web sites include IEEE's Computer Society Task Force on Cluster Computing *(www.ieeetfcc.org)*, the Linux HPC site *(www.linuxhpc.org)*, the Windows HPC site *(www.winhpc.org)*, the Sun HPC site *(sun.com/hpc)*, and IBM's deep computing site *(www.ibm.com/servers/deepcomputing/)*.

Also, be sure to read *GPU Cluster for High Performance Computing* by Zhe Fan, et al. Other written resources include white papers such as Intel's *Optimizing Software for Multi-core Processors* and *How Much Performance Do You Need for 3D Medical Imaging?*, Toshiba's *The Next Revolution: 256-Slice CT* by Richard Mather, PhD, and Altera's *Medical Imaging Implementation Using FPGAs*.

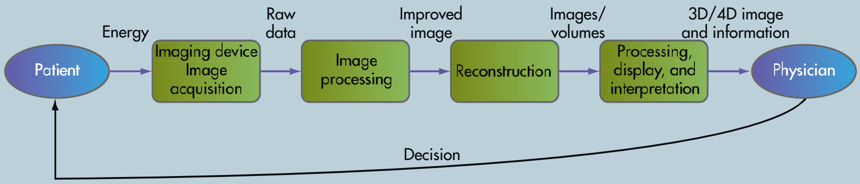


**6. Data in most modern medical-imaging modalities follows a specific flow from patient to physician.**

**DIAGNOSTIC FOOD FOR THOUGHT** • Before designing your next medical imaging system, there's one last thing to consider. With the correct image analysis and diagnostic programming, is it possible for a computer to "out-diagnose" a physician or radiologist? It certainly seems feasible for some ailments even now, and this possibility grows stronger with each generation of processor power and knowledge.

**ED ONLINE 15795**

## MEDICAL-IMAGING MODALITIES

**Medical imaging includes a broad spectrum of devices and** technologies, using a variety of modalities. For example, magnetic resonance imaging (MRI) uses three kinds of electromagnetic radiation to generate the desired images.

First, a main powerful magnet (normally 0.5 to 3.0 tesla) can be used to polarize hydrogen atoms in the tissues of interest (brain, muscle, etc.) because hydrogen has a large magnetic moment. Second, gradient magnets located within the main magnet are switched on and off rapidly to alter the main magnetic field for the area targeted to create image slices.

Third, radio-frequency (RF) pulses that are specific to hydrogen are directed at the target tissue perpendicular to the main magnetic field, causing the protons within the hydrogen atoms to spin in a different direction (resonate) at a particular frequency (i.e., they become "excited").

When the pulse is removed, the hydrogen atoms "recover" and return to their natural alignment. In doing so, they emit pent-up energy in the form of radio waves, producing a signal received with an RF antenna. The signal is then converted using a discrete Fourier transform (DFT) into a photo slice. **ED ONLINE 15794**