
Medical Imaging Implementation Using FPGAs

Introduction

Medical imaging equipment is taking on an increasingly critical role in healthcare, as the industry strives to lower patient costs and achieve earlier disease prediction using non-invasive means. To provide the functionality needed to meet these industry goals, equipment developers are turning to programmable logic devices.

Earlier prediction and treatment are driving the fusion of modalities such as positron emission tomography (PET)/computerized tomography (CT) and X-ray/CT equipment. The higher image resolutions that are needed require fine geometry micro-array detectors coupled with sophisticated software/hardware systems for the analysis of photonic and electronic signals. These systems must provide both highly accurate and extremely fast processing of large amounts of image data (up to 250 GMACS and 1 Gbps). Furthermore, to lower patient costs, each piece of equipment must be lower priced and possess a longer life utility. This calls for more flexible systems with the capability to continually update features and algorithms over the equipment's lifetime. Together, flexible algorithm deployment and modality fusion compel the use of programmable system electronic components, such as high-powered CPUs and FPGAs.

Several factors should be considered in the efficient development of flexible medical imaging equipment:

- Development of imaging algorithms requires high-level intuitive modeling tools for continual improvements in digital signal processing (DSP).
- The performance needs for near real-time analysis require system platforms that scale with both software (CPUs) and hardware (configurable logic). These processing platforms must meet various performance price points and be capable of bridging the fusion of multiple imaging modalities.
- System architects and design engineers need to quickly partition and debug algorithms onto these platforms, using the latest tools and intellectual property (IP) libraries, to speed their deployment and improve profitability.

With these factors in mind, Altera has introduced its new modular Video and Image Processing Suite, a blockset of key IP building blocks that can accelerate the development and implementation of sophisticated imaging algorithms into FPGAs. The new Video and Image Processing blockset, along with other Altera®/partner IP modules and reference designs (including IQ modems, JPEG2000 compression, fast Fourier transform (FFT)/inverse fast Fourier transform (IFFT), edge detection, etc.), provide a broad range of tools designers can use to speed FPGA implementations of computationally intensive tasks.

First, we will examine some of the trends in imaging algorithms for medical applications and the implementation use of imaging IP and FPGAs.

Algorithm Developments in Medical Imaging

Some of the most critical pieces of equipment in today's medical development environment include:

- X-ray, magnetic resonance imaging (MRI), CT scanner, ultrasound, and 3D imaging systems
- Measuring and analysis instruments
- Optical manipulation and analysis
- Surgical microscopes
- Telemedicine systems

Designers are looking for rapid imaging solutions with applications including:

- Image analysis and pattern recognition
- Image enhancement and restoration

- Image and data compression
- Wavelet transform capabilities
- Color space conversion

What types of algorithm developments are taking place within these applications and major modalities? The following sections will cover some of the key developments that are driving the integration of programmable logic into medical imaging equipment.

Image-Guided Therapy

Intraoperative image processing for surgical guidance uses the registration (correlation) of preoperative (CT or MR) images with real-time 3D (ultrasound and X-ray) images to guide the surgical treatment of disease using non-invasive therapies (ultrasound, MR interventional, and X-ray treatments). Various algorithms are being developed to provide the best registration results for the specific fusion of modalities and therapy types.

Molecular Imaging

Molecular imaging is the characterization and measurement of biological processes at the cellular and molecular level. Its purpose is to detect, capture, and monitor abnormalities that cause disease. To speed the study of symptoms and evaluation of therapies, small animals are used in molecular imaging applications. As a result, all medical imaging modalities are being downscaled to miniature equipment sizes. Areas of study include data acquisition, image reconstruction, image processing, and analysis. For example, X-ray, PET, and single photon emission computed tomography (SPECT) have been combined to map functional/cellular/molecular images at low resolution onto corresponding anatomy at high resolution, down to 0.5 mm. Miniaturization and algorithm exploration drive the use of FPGAs into these compact system platforms, which help accelerate performance beyond the capabilities of multicore CPUs.

Imaging Algorithms

Image enhancement is commonly done with convolution (linear) filtering. High-pass filtering enhances the detail in an image, but also makes the noise more visible. Low-pass filtering suppresses the noise at the expense of blurring the detail. Most images contain some areas with detail and other areas without much detail. Linear combination filtering is a technique that enhances detail in the former and reduces noise in the latter, by producing both high-pass and low-pass filtered images and combining them according to a mask. This technique works because the eye is less sensitive to noise in areas that contain detail. The mask is the smoothed output of a Sobel edge-detection filter. It takes on values near one in areas with detail and values near zero in areas without detail. The linear combination of the high-pass and low-pass filtered images, weighted by the mask, produces an image in which detail has been enhanced, while noise has been reduced.

Video Image Stabilization and Registration (VISAR) is an algorithm for real-time video image stabilization. VISAR was developed to improve the quality of video images beyond simpler horizontal and vertical image registration techniques alone, by accounting for rotational and zooming effects in video data sequences. It co-aligns video image fields by removing the effects of translation, magnification, and rotation. Because VISAR allows the user to combine several video images together, noise can be averaged out among frames. VISAR also smoothes jagged edges found in still images extracted from video. VISAR can correct image jitter to about 1/10th of a pixel. This algorithm could potentially be applied to:

- Clarify cell images viewed through a microscope
- Stabilize eye images for retinal study
- Stabilize thermal infrared imaging
- Stabilize camera and body movement during endoscopic surgery
- Stabilize images transmitted for telemedicine
- Improve ultrasounds to correct for body movement when viewing MRI videos

Wavelet transforms is an analytical algorithm that overcomes some of the limitations of Fourier analysis. While Fourier analysis transforms signals from time into the frequency domain, it loses time information. This is why, when looking at a Fourier transform of a signal, it is impossible to tell when a particular event took place. Many imaging signals contain important nonstationary or transitory characteristics: drift, trends, abrupt changes, and beginnings and/or ends of events. To help acquire event information in signals, Dennis Gabor (1946) adapted the Fourier transform to analyze only a small section of the signal at a time—a technique called windowing the signal. More recently, wavelet analysis was developed using a windowing technique with variable sized regions. Wavelet analysis allows the use of long time intervals for more precise low-frequency information and shorter regions for high-frequency information. Wavelet applications include detecting discontinuities and breakdown points, detecting self-similarity, suppressing signals, denoising signals, denoising images, compressing images, and fast multiplication of large matrices. Altera's Video and Image Processing and DSP libraries provide key building blocks for wavelet manipulations including scaling, shifting, high-pass/low-pass filtering, I/Q decomposition, and reconstruction.

Distributed vector processing is an algorithm that enables faster computations. The recently developed S-transform (ST) combines features of the FFT and wavelet transforms, revealing frequency variation over both space and time. Applications include texture analysis and noise filtering. However ST is computationally intensive, making conventional CPU implementations too slow. This problem is addressed by combining vector and parallel computations for a 25-fold reduction in processing time. Such computationally intense algorithms can be significantly accelerated using a vector processor in conjunction with parallel computations implemented in FPGAs. (See the Development Tools section to learn how Altera products can speed algorithm implementations.)

CT Scanning

PET/CT fusion provides an alternative to software-based image fusion (registration), which is routinely used for the alignment of functional and anatomical images of the brain. For other parts of the body, image registration is more problematic because of differences in patient positioning, scanner bed profiles, and involuntary movement of internal organs. In PET/CT fusion, a scanner acquires both functional and anatomical images during a single imaging session, rather than fusing the images post hoc. The CT images provide essentially noiseless attenuation correction factors for the PET data.

MRI

MRI reconstruction creates cross-sectional images of the human body. Two distinct steps are necessary to reconstruct a 3D volume:

1. 2D reconstruction of each slice via FFT produces slices in gray-level, typically matrices, from the frequency domain data.
2. 3D volume reconstruction creates an interslice distance close to the interpixel distance via interpolation of slices, so that images can be viewed from any 2D plane.

Iterative resolution sharpening uses a spatial deblurring technique based on an iterative inverse filtering procedure that reduces noise while the image structure is refocused. Thus, the overall visual diagnostic resolution of the cross section is significantly improved.

Ultrasound

Despeckling of ultrasound images is possible using lossy compression. Ultrasound speckle is generated by the interaction of various independent scatterers (similar to multipath RF reflections in the wireless domain), and manifests itself as multiplicative noise having a granular appearance. By taking the logarithm of the image, speckle noise becomes additive to the desired signal and can be minimized via lossy wavelet compression using a JPEG2000 encoder. Altera offers partners' IP for JPEG2000 encoding/decoding, as well as a median filter and reference design in the Video and Image Processing Suite for denoising speckle.

Cardiac motion estimation constitutes an important aid in quantifying the elasticity and contractibility of the heart muscle. Localized areas exhibiting abnormal movements are indicative of ischemic heart regions, where insufficient circulation exists. A developing algorithm includes the quantitative evaluation of elasticity from a series of ultrasound images using spatial-temporal registration techniques to find the deformation field with respect to a reference frame. Key Video and Image Processing and DSP building block functions applied here include 2D filtering, despeckling, correlation, and smoothing.

Video Imaging

In vascular imaging, the gold standard for many years has been radiopaque contrast angiography from X-ray images of blood vessels, using salt-based contrast agents. Today, CT angiography, time-of-flight/phase contrast MR angiography, and duplex/intravascular ultra-sonography are more commonly used. These techniques involve simultaneous acquisition and registration of optoacoustic and ultrasound images, segmentation of vascular and skeletal images, and the use of correlation-based enhancement filters to reduce false positives when diagnosing lung disorders.

For early detection of cancer using the disease's ability to recruit a new blood supply, the BioScan System™ can be implemented. The BioScan System is based on a sensor called the Quantum Well Infrared Photodetector (QWIP), which is sensitive to temperature changes of less than 0.027° Fahrenheit (0.015° Celsius) and has a speed of more than 200 frames per second. The digital sensor detects the infrared energy emitted from the body, thus "seeing" the minute differences associated with increased blood flow due to cancer. A typical implementation is based on a programmable systolic array implemented with a general-purpose workstation and a special-purpose hardware engine built from FPGAs. The engine can accelerate the core algorithm nearly 1,000 times more than the rate achieved by a state-of-the-art workstation.

X-Ray Imaging

Reverse Geometry X-radiography Lamography (RGL) was developed by Digiray Corp. to improve image clarity. Unlike conventional x-ray systems, the RGL system places the radiographic object close to the x-ray source. A point detector then captures the primary radiation without the image-degrading secondary radiation inherent in standard x-ray systems. By using scintillating crystal imaging detectors with a dynamic range at least 10 times superior to film, the digital image can be enhanced by a wide variety of standard image-processing tools such as averaging, filtering, image subtraction, and edge detection/enhancement. The clarity of RGL may provide x-ray imaging for medical applications including mammography, cardiac imaging, brain surgery, and orthopedics.

Motion correction of coronary X-ray images can help sharpen image results. For example, to minimize the effects of breathing and heart pumping (cardiac-respiratory cycle) during imaging, the motion of a '3D plus time' coronary model is projected onto 2D images and used to compute a dewarping function (translate and zoom) to correct for the motion, resulting in clearer images.

Critical Building Block Functions

What are some of the key building block functions required for these sophisticated imaging algorithms?

In CT reconstruction, interpolation, FFT, and convolution functions are required. In ultrasound, important processing methods include color flow processing, convolution, beam forming, compounding, and elasticity estimation. General imaging algorithms include functions like color space conversion, graphic overlays, 2D/median/temporal filtering, scaling, frame/field conversions, contrast enhancement, sharpening, edge detection, thresholding, translation, polar/Cartesian conversion, nonuniformity corrections, and pixel replacement.

The Video and Image Processing Suite, along with additional IP and reference designs from Altera and its partners, can accelerate integration of these algorithms onto FPGAs, including those systems with the highest performance and smallest footprint. Before describing the Video and Image Processing MegaCores®, it is important to consider the development methodology of algorithms and corresponding tools.

Development Tools

Imaging architects use high-level software tools to model various algorithms and results. Image modeling tools, such as the Interactive Data Language, are available from companies like RSI, BIR, and Advanced Digital Imaging Research (ADIR). Another trend is the use of open-source toolkits for image registration, segmentation, and image-guided surgery. These tools are optimized for imaging applications and algorithm development via software, but not for implementation into FPGAs.

Algorithmic development in MATLAB software and system-level design in Simulink software are very common DSP design approaches. For example, ADIR, a research and development organization specializing in digital imaging software and algorithm development, needs flexible tools to create fast and accurate image-processing algorithms. These algorithms define and implement various techniques, manipulate 3D images and statistical data, solve equation sets, and display/document algorithms. ADIR has been using MATLAB as a development tool for almost 15 years in its various specialties, including digital image processing, quantitative image analysis, pattern recognition, digital image coding and compression, automated microscopy, forensic image processing, and 2D wavelet transforms. In addition to algorithm development, MATLAB can also simulate the use of fixed-point arithmetic commonly used in FPGAs.

Using MATLAB coupled with Altera's tools—DSP Builder, SOPC Builder, Nios[®] II CPU development kit, Nios II C2H Compiler acceleration, and Quartus[®] development suite—designers can accelerate their implementations of designs onto FPGAs. These toolsets enable companies like ADIR to accelerate implementation of algorithms in FPGAs, when CPUs alone are not sufficient.

The DSP Builder tool provides a design flow based on IP libraries (such as Altera's Video and Image Processing, DSP, and connectivity IP MegaCore libraries) for bridging between high-level algorithms developed in Simulink and FPGA implementation using the Altera Quartus toolset. SOPC Builder is a system integration tool that automatically generates the interconnect code (Verilog or VHDL) between IP blocks, Altera MegaCores (including Nios II CPUs), partner IP, and user-defined functions. The Nios II development kit enables C programs to be moved or debugged onto the FPGA into one or more Nios II CPUs. The Nios II C2H Compiler is a new tool that significantly increases performance of software running on Nios II CPUs by profiling the inner loops of C-coded algorithms and generating coprocessor logic within the FPGA, thereby achieving acceleration factors of 10 to 100 with only modest increases in logic utilization. The Quartus development suite is the primary tool for implementing programmable hardware and software functions onto FPGAs. It provides IP-importing capability, simulation, placement and routing of all functions, and FPGA programming in situ on Altera development boards or user target systems.

Video and Image Processing Suite

Altera's Video and Image Processing Suite contains nine configurable MegaCore functions, which perform common imaging functions on Altera's FPGAs. Table 1 summarizes the functions provided.

Table 1. Functions Available With the Video and Image Processing Suite

Function	Description
Deinterlacer	Converts interlaced video formats to progressive video format
Color Space Converter	Converts image data between a variety of different color spaces
Scaler	Resizes and clips image frames
Alpha Blending Mixer	Mixes and blends multiple image streams
Gamma Corrector	Performs gamma correction on a color plane/space
Chroma Resampler	Changes the sampling rate of the chroma data for image frames
2D Filter	Implements 3x3, 5x5, or 7x7 finite impulse response (FIR) filter operation on an image-data stream to smooth or sharpen images
2D Median Filter	Implements a 3x3, 5x5, or 7x7 filter that removes noise in an image by replacing each pixel value with the median of neighboring pixel values
Line Buffer Compiler	Efficiently maps image line buffers to Altera on-chip memories

Each function is implemented with a simple streaming interface for transmitting image data through the function. It enables:

- Easy connection of any Video and Image Processing function, in any order
- An efficient means to support designs with widely different data rates

The Video and Image Processing IP Suite functions use common image data format parameters to provide compile time parameters. A MegaWizard® GUI is used to configure the core. The functions are available from within Simulink from the Altera DSP Builder blockset in the Video and Image Processing Suite. For medical imaging, some of the cores are more applicable to video and others to imaging functions. One core particularly applicable to both is the line buffer compiler.

Line Buffer Compiler

FPGA memory is a valuable resource for many video and imaging applications, particularly when developing systems that require high-definition resolutions and maximum data throughput. The line buffer compiler provides an efficient means to map video line buffers to Altera's on-chip memories.

Imaging Cores

The 2D FIR MegaCore function provides a flexible and efficient means to perform 2D finite impulse response (FIR) filtering operations using matrices of 3x3, 5x5, or 7x7 constant coefficients with 8-, 10-, and 12-bit pixel input data.

The 2D Median Filter function provides a means to perform 2D median filtering operations using matrices of 3x3, 5x5, or 7x7 kernels. Each output pixel is the median of the input pixels found in a 3x3, 5x5, or 7x7 kernel centered on the corresponding input pixel. Zeros are filled in where this kernel runs over the edge of the input image.

The Scaler MegaCore function provides an efficient means to resize image frames using either bilinear filtering or nearest neighbor scaling, including support for 8- and 10-bit pixel input data and image clipping. The initial version includes conversion support for a fixed number of standard and high-definition resolutions with later releases supporting arbitrary horizontal/vertical scaling ratios.

2D Interpolation Methods

The interpolation method uses the real-valued point from where a sample is to be taken, called (i', j') , and its nearest (rounded-down) integer point called (i, j) , where $(i, j) = (\lfloor i' \rfloor, \lfloor j' \rfloor)$. d_i and d_j are defined as $d_i = i' - i$ and $d_j = j' - j$, respectively, while F is the function that returns an intensity value for a given point.

Nearest neighbor method: $F(i', j') = F(i, j)$

Linear interpolation: $F(i', j') = \sum_{m=0}^1 \sum_{n=0}^1 [(m - (1 - d_i)) \cdot (m - (1 - d_j)) \cdot F(i + m, j + n)]$

Nearest neighbor and linear interpolation typically provide high quality results when scaling or downscaling by ratios up to 2X. Support for bicubic interpolation is planned for later releases.

Video Cores

The color space converter function provides a flexible and efficient means to convert image data from one color space to another. Conversions are achieved by providing an array of 12 coefficients that relate various color spaces, including computer RGB, studio RGB, YIQ/YUV, and YCbCr.

Chroma resampling allows the user to change between 4:4:4, 4:2:2, and 4:2:0 sampling rates. Independent control of horizontal and vertical format conversion methods is provided, with quality choices between pixel replication/decimation of nearest neighbor and bilinear interpolation. The user can specify how the resampling

behaves with the occurrence of overflow (i.e., when output values exceed 2^{bps}) and underflow (i.e., negative output values). The available options are IGNORE and SATURATE.

The alpha blending mixer function provides an efficient multiple color plane mixing up to eight image layers. The function provides support for both picture-in-picture mixing and image blending with per pixel alpha support.

The de-interlacer function provides a flexible and efficient means to convert interlaced video to progressive video using bob and weave algorithms. The de-interlacer produces an $N/2$ -Hz output frame rate for an N -Hz input field rate (e.g., 1080i @ 60 Hz to 1080p @ 30 Hz).

The weave algorithm creates an output frame by filling in all of the missing lines in the current input field with lines from the previous input field. This produces good results for still parts of an image but unpleasant artifacts in moving parts. Bob de-interlacing scales input fields up by a factor of 2 vertically. This function supports two types of scaling for bob de-interlacing: scanline duplication and scanline interpolation. Scanline duplication simply scales by repeating each scanline in input field F0 twice to make the output frame. Input field F1 is discarded. Scanline interpolation recreates the line missing from input field F0 by performing an unweighted mean of the lines above and below them. Input field F1 is discarded. At the bottom of field F0 there is only one line available, so this line is just duplicated as per scanline duplication.

Example Video Design

A typical video system using the Video and Image Processing Suite is shown in Figure 1.

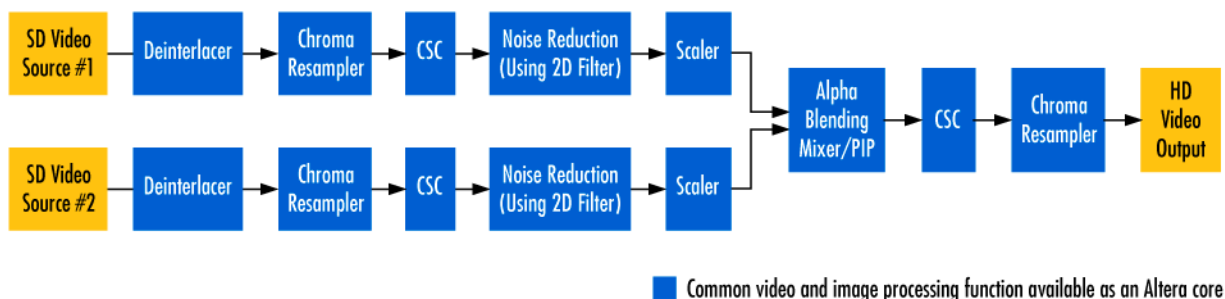


Figure 1. Example Design Block Diagram

Leveraging FPGAs for Medical Imaging

An example of an OEM leveraging Altera devices and tool suites highlights the benefits of FPGA and imaging IP use for today's algorithms and systems.

3D-Computing Inc. developed a breakthrough medical technology called 3D Complete Body Screening (3D-CBS), based on Altera FPGAs. 3D-CBS is expected to revolutionize the use of PET technology by making it safe enough to use for routine preventative-measure patient examinations. 3D-Computing uses Altera devices to accelerate the processing of the algorithms that allow this new technology to more accurately examine a patient's body, using only four percent of the radiation they would receive in a traditional PET scan.

Preventive-measure full-body screening of a healthy person using current PET technology is inadvisable because of the dangers of radiation overexposure. Current PET technology is also very slow and costly, making its routine use prohibitive in a time of rising medical costs. Altera FPGAs were used in a variety of ways inside the 3D-CBS system, the most important of which used a matrix of FPGAs to perform the high-speed processing required to more accurately and effectively capture a far greater percentage of the photons emitted from a patient's body during an examination. As a result, the patient's level of radiation exposure is reduced, and image quality is improved.

Key Third-Party Video and Imaging Partners

In addition to the Video and Image Processing Suite and DSP library, Altera has partnered with multiple suppliers of video and imaging IP to help design solutions that meet the requirements of medical imaging equipment developers.

- ATEME provides H.264 (MPEG-4 AVC) main-profile SD/HD encoding
- Barco provides JPEG/JPEG2000 encoders/decoders and MPEG-4 advanced simple profile SD encoding
- BroadMotion provides an efficient JPEG2000 encoder/decoder
- CAST provides JPEG/JPEG2000 encoders/decoders, MPEG-4 advanced simple profile, and H.264 baseline profiles (SD/HD encoding)

Conclusion

Baby boomers are seeking new and more accessible therapies to treat the most common diseases (especially heart-related ailments and cancer), including earlier detection and non-invasive surgical treatments. Advances in the fusion of diagnostic imaging modalities and their associated algorithm developments are the primary drivers in advancing new equipment to meet these patient needs. Advanced algorithms require scalable system platforms with significant increases in image processing performance, yet in smaller, more accessible portable equipment.

Integrated into multicore CPU platforms, FPGAs provide the DSP horsepower for the most flexible, highest performance systems. To help accelerate the implementation of sophisticated imaging algorithms onto these platforms, high-level development tools and IP implementation libraries are required. Altera developed the Video and Image Processing Suite with these concerns in mind. The suite contains key building block functions for imaging and is integrated into Altera's complete toolset for rapid development.

Further Information

- Altera's Video and Image Processing Solutions website:
www.altera.com/video_imaging
- Altera's Video and Image Processing Reference Design:
www.altera.com/end-markets/refdesigns/sys-sol/broadcast/ref-post-processing.html



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>

Copyright © 2006 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.