

Signal Processing Leads a Photographic and Imaging Revolution

Photography and imaging have been radically transformed over the past couple of decades in ways that 19th-century pioneers such as Louis-Jacques-Mandé Daguerre (Figure 1) and Henry Fox Talbot could have scarcely imagined. Traditional photography and imaging, rooted in chemical processes, have now largely given way to digital methodologies and technologies. The result has been faster, less expensive, and more convenient ways of acquiring and presenting images, and in many cases the creation of clearer, more detailed, and less distorted pictures on many different types of media.

Signal processing plays an important role in virtually all types of digital photography and imaging. In consumer, professional, industrial, and scientific still cameras, sophisticated integrated algorithms help determine how images are collected, interpreted, and stored. Algorithms, for example, ensure that captured raw sensor data are efficiently translated into color-corrected image data that can then be stored either in raw pixels or as compressed images. Image processing algorithms are also involved in image capture and compression, focus and exposure control, managing white balance, demosaicing, image storage, preview display rendering and scaling, and various post-processing tasks.

FREEZING STREAKS

As researchers work to extend the capabilities of existing imaging systems, as well as blaze new technologies, signal processing provides ways of adding new capabilities and improving the performance of

existing features and functions. Researchers at Laguna Hills, California-based MetroLaser, for instance, used signal processing in the creation of a camera that captures full-color images of projectiles traveling at speeds of up to 3,350 m/s, approximately ten times the speed of sound. The digital galvo mirror streak camera, designed to replace now-obsolete film-based streak cameras, records the motion of a projectile as it passes in front of its lens, creating a long, continuous composite image of the object.

Now that digital technology has completed its sweep across the photography industry, the specialized film required for analog streak photography cameras is no longer being manufactured. In 2007, the U.S. Air Force asked MetroLaser to design a modern digital system that could produce high-quality ballistic images. “The Air Force wanted a highly rugged digital camera system that would allow them to get full-color, high-resolution photos of rocket sleds moving up to Mach 10 with schlieren effects (optical inhomogeneities in transparent material that aren’t necessarily visible to the human eye) included,” says Ben Buckner, MetroLaser’s chief scientist.

The imaging system Buckner developed with coresearcher Drew L’Esperance, utilizes a precisely controlled mirror galvanometer to follow a rapidly moving object and freeze its image. Buckner explains that the mirror tracks the ballistic object as it moves past the camera lens and directs appropriate portions of the image onto specified areas of the image sensor to form a complete, undistorted picture. “It enables full-color 15-plus megapixel photography of objects moving at high speeds with standard photographic flashes, or even strong natural light, with schlieren photography of disturbances in the surrounding air,” Buckner says.

Since the mirror is synchronized to the ballistic object, the biggest challenge the researchers faced when developing the system was finding a way of accurately measuring the object’s speed and calculating the swing of the mirror to precisely match that object’s trajectory. “The main challenge is that the software has to control a galvanometer mirror very precisely and very quickly,” Buckner says. “The rocket sled velocity is variable, so you have to measure [the sled] as it’s coming down the track, and in a few milliseconds, the software has to calculate the required trajectory for the mirror to match the speed and generate the control signals,” he continues. At such high speeds, the mirror response tends to be nonlinear. “So in addition to working out the basic kinematic equations for the motion, you have to put in some corrections and then generate the required control waveform on the fly,” he says.

“I put a fair bit of time into optimizing the code for fast execution, since it’s on a



[FIG1] Photography pioneer Louis-Jacques-Mandé Daguerre. (Photo courtesy of www.wikipedia.com.)

fairly low-end embedded processor,” Buckner says. Other challenges were optimizing the acceleration curve and deriving the nonlinear response corrections. “We’re pushing the galvo controller to its limits, so the way you accelerate it is important,” Buckner says.

Buckner foresees other potential applications of the galvo mirror streak camera, such as determining the finishing order in high-speed races. “There are other ways of doing high-speed imaging, but this one is particularly effective at doing very high-resolution images, color images, schlieren images, and large-scale subjects where it’s very difficult to get a submillisecond flash of sufficient energy,” he explains.

Most other types of high-speed imaging techniques are hampered by resolution and color reproduction limitations, Buckner says. The galvo mirror streak camera, however, has many of the same benefits as conventional professional-level cameras. “We can work with the same flash rigs that any commercial photographer uses, and

largely [with] the same camera back,” Buckner explains. “Our system really just replaces the camera lens, and all the rest of it is the same equipment you would find being used to take school pictures or magazine glamour shots.” The system’s modular design also allows it to be easily upgraded. “Commercial camera backs are always improving,” Buckner says. “We could easily get [the system] up to 80 megapixels now just by putting one of the newer camera backs on it.”

UPGRADING MICROSCOPES

Researchers at the California Institute of Technology (Caltech) relied on signal processing techniques to help develop a method of converting relatively inexpensive conventional microscopes into high-end billion-pixel imaging systems. The new approach, called Fourier Ptychographic Microscopy (FPM), promises to significantly enhance the efficiency of digital pathology, particularly in situations where specialists need to review large


numbers of tissue samples (Figure 2). The researchers also hope that the technology will bring high-performance microscopy capabilities to medical clinics that can’t currently afford high-density imaging systems.

“A microscope’s pixel count is fundamentally limited by the physical nature of the optical lenses—all physical lenses have aberrations that ultimately degrades the imaging process,” says Changhui Yang, a Caltech professor of electrical engineering, bioengineering, and medical engineering. A standard digital microscope typically creates images with approximately ten megapixels of resolvable pixels. “You can choose between a large field of view and a poor resolution, or small field of view and a high resolution,” Yang says. “If we are simply examining microscope slides with our eyes, this pixel count is quite sufficient, but this pixel count is woefully low to address digital pathology needs.”

FPM provides a computational-oriented approach that aims to free microscope developers from the physical limitations of optical lenses. The microscopy industry’s current approach for creating high-quality microscopes is to use very complicated—and expensive—stacks of exotic glass lenses to cross-compensate for aberration. FPM makes this type of development model unnecessary. “To FPM, the distortions in optical elements are simply mathematical functions that it can manipulate computationally and zero out of the final processed image,” Yang says. “We can take a [poor] microscope, make some cheap modification to its lighting scheme, and use it to collect a sequence of poor-quality images. “The algorithm will then take the data and render a high-quality and high pixel count image.”

FPM stitches the low-resolution images together to create high-resolution intensity and phase information, providing a more complete picture of a particular clinical sample’s entire light field. To create a complete image of a particular sample, the system acquires approximately 150 low-resolution images with each image corresponding to a single element in the light-emission diode (LED) array.

When their work began several years ago, the researchers struggled to reach



[FIG2] (a) An FPM installation that converts a relatively inexpensive conventional microscope into a billion-pixel imaging system. The inset in (a) shows a magnified image of the LED chip that contains a red-green-blue LED. A raw image of human blood smear taken with a 2X objective lens is shown in (b) along with (c) the reconstructed image produced by the FPM system. (Photo courtesy of the California Institute of Technology.)

their goal by completely eliminating lenses from microscopes. The approach found them trying a variety of chip-scale microscope systems, none of which met their performance or cost targets. "In the past couple of years, we started asking ourselves whether we can tackle optical aberrations head-on rather than side-stepping the problem," Yang says. "If we could, we would not have to throw out the compelling advantages of using lenses in microscopes." Benefits associated with lenses include the ability to concentrate light and easier color handling. By following a computational-oriented approach and developing FPM, the researchers were able to bring the resolution of a conventional 2X objective lens to the level of a 20X objective lens.

FPM's main design strategy is similar to that of interferometric synthetic aperture microscopy: expanding summation









by parts (SBP) in Fourier space through multi-image fusion. However, because no measured phase information is needed for FPM, the researchers' approach eliminates the design challenges associated with interferometric detection. Yet FPM's image recovery procedure follows a strategy common with ptychography scanning diffraction microscopy, iteratively solving for a sample estimate that is consistent with many intensity measurements. Unlike ptychography, however, FPM's object support constraints are imposed in the Fourier domain, offering several unique advantages and opportunities.

FPM's data collection procedure is straightforward, according to Yang. The process involves placing a two-dimensional (2-D) sample at the focal plane of a low-numerical aperture microscope objective and collecting a sequence of images, with the sample successively illuminated

by plane waves at different angles. Unlike other synthetic aperture techniques, the procedure acquires intensity images of the sample, so no interferometric measurements are required. The use of a low-numerical aperture objective lens allows a large field of view to be captured at the expense of a low spatial resolution.

A major advantage of the new approach is relatively pain-free hardware compatibility. Manufacturers only need to add an LED array to an existing microscope—no other hardware modifications are necessary. A computer then handles the rest of the work. The researchers say that their method could have wide applications not only in digital pathology but also in everything from hematology to wafer inspection to forensic photography. "A broad swath of imaging modalities can benefit from this computational approach of tackling imaging," Yang says. Satellite

Engineering from CRC Press



CRC Press will be Exhibiting at this year's ICASSP, please visit our booth!

SAVE 20% when you order online and enter Promo Code EKN02.

FREE standard shipping worldwide!

www.crcpress.com



imaging is a particular area of interest. “The features we can resolve in satellite images is tied to the size of the camera you send up to space,” Yang says. “We think FPM can ... allow satellite imaging at unprecedented resolution.” X-ray imaging is another potential application. “X-ray imaging is confounded by the lack of high-quality lenses,” Yang says. FPM makes this a nonissue since it simply treats the distortions as a mathematical function.

Yang is optimistic that FPM microscopy will soon become a scientific mainstay. “Because the hardware is so simple, we hope it will be commercially available in a couple of years,” he says.

FIXING PHOTOS

Sophisticated computational processing also promises to benefit the everyday users of smartphones and various other types of consumer-level cameras. Targeting such individuals, researchers at the Massachusetts Institute of Technology (MIT) have developed a chip-based processor that’s dedicated to helping almost any camera user—amateur or pro—create high-quality photographs.

At MIT’s Microsystems Technology Laboratory, Rahul Rithe, a graduate student in the school’s Department of Electrical Engineering and Computer Science, recently worked on the team that developed the “Maxwell” processor (named after James Clerk Maxwell, who in 1855 first proposed creating color photographs by using red, green, and blue filters to merge together three captured images). The chip (Figure 3) aims to help shutterbugs by

almost instantaneously creating more realistic or enhanced lighting in a shot without destroying the scene’s ambience. “This energy-efficient and scalable implementation is ideal for integration with mobile devices such as smartphones, tablets, digital cameras, and even laptops, to enable live computational photography on these energy-constrained devices,” says Rithe, who was lead author of a paper on the project.

RITHE NOTES THAT SIGNAL PROCESSING TECHNIQUES LIKE NONLINEAR FILTERING ARE ESSENTIAL TO THE PROCESSOR’S OPERATION.

Most current computational photography applications are software based. “Performing [image optimization] tasks on general purpose CPUs and GPUs consumes a significant amount of power and is typically not fast enough to support real-time performance,” Rithe says. He states that the Maxwell processor can perform optimization operations in real-time while consuming dramatically less power. “While software-based systems typically take several seconds to perform an operation like high dynamic range (HDR) imaging, the chip can do it in a few hundred milliseconds on a ten-megapixel image,” says Rithe, who notes that the high-performance chip can also enhance video output.

To create an HDR image, Maxwell tells the camera to take three individual low dynamic range photos: a normally exposed image, an overexposed image capturing details in the dark areas of the scene, and an underexposed image capturing details in the bright areas. The processor then merges the photos to create a single image that captures the scene’s full color and brightness range.


The processor uses bilateral filtering, Rithe says, a nonlinear filtering technique that effectively reduces noise and smooths out an image’s defects without blurring sharp edges, thereby preserving important details. “Nonlinear filtering techniques like bilateral filtering are used in a wide

range of computational photography applications,” he notes. Unfortunately, due to its high computational complexity, bilateral filtering is generally inefficient and slow. “We leveraged the bilateral grid structure ... and developed an optimized hardware implementation that represents a 2-D image using a three-dimensional (3-D) data structure and performs the processing in the 3-D domain,” Rithe says. “This significantly reduces both the computational complexity and the amount of memory required to process large images.”

Rithe notes that signal processing techniques like nonlinear filtering are essential to the processor’s operation. “Signal processing is vital to our research in the form of image processing techniques that enable us to manipulate and create images that could have only come from a handful of prolific artists, like Ansel Adams, in the past,” Rithe says. The algorithms implemented on the chip were inspired by the computational photography work of Fredo Durand and Bill Freeman (an associate professor and professor, respectively), at MIT’s Computer Science and Artificial Intelligence Laboratory.

Multimedia processing applications such as computational photography have very high computational complexity and memory requirements. “The major challenge was to come up with a combination of algorithmic-, architectural-, and circuit-level innovations that significantly brought down the computational complexity, memory requirement, and bandwidth,” Rithe says. “To enable real-time processing while being extremely power efficient, we developed a highly parallel architecture that is able to support real-time processing of high-definition (HD) images while operating at less than 100 MHz frequency, as opposed to CPUs and GPUs that operate at several GHz.” One of the key components in maximizing the processor’s energy-efficiency is voltage/frequency scaling. “Careful circuit design for low-voltage operation ensured reliable performance from 0.9 V down to 0.5 V,” Rithe explains. “This enables voltage/frequency scaling to maximize the energy-efficiency for a required performance level.”

Rithe developed the processor on a team that included Anantha Chandrakasan, MIT’s Joseph F. and Nancy P.




[FIG3] A printed circuit board containing MIT’s “Maxwell” computational photography processor. (Photo courtesy of MIT.)

Keithley Professor of Electrical Engineering. Other members included fellow graduate student Priyanka Raina, research scientist Nathan Ickes, and undergraduate student Srikanth Tenneti.

Work on the processor began in January 2011 when Rithe and his coresearchers started exploring different types of computational photography algorithms. After the team completed algorithmic optimizations, developed a highly parallel architecture to enable real-time processing, and finalized circuit implementations, the chip was sent for fabrication in April 2012 through Taiwan Semiconductor Manufacturing Company's University Shuttle Program.

The researchers presented their work at the IEEE International Solid-State Circuits Conference in February 2013.



[FIG4] A demonstration system that integrates the processor with DDR2 memory and connects with a camera and a display through the USB interface. The system provides a platform for live computational photography. (Photo courtesy of Nathan Ickes/MIT.)

The live demonstration system prototype combined the processor with external memory, camera, and display (Figure 4). “We received significant interest from

the leading mobile processor and device makers,” Rithe says.

Future processors designed along the lines of Maxwell will permit more complex computational photography applications, Rithe says. He notes that Raina is currently leading an effort to develop a processor capable of sharpening images that are blurred due to camera shake during image capture. “We are also exploring ways of extending computational photography and computer vision techniques to enable portable smartphone-based medical imaging applications,” he says.

AUTHOR

John Edwards (jedwards@johnedwardsmedia.com) is a technology writer based in the Phoenix area.



Innovation doesn't just happen.
Read first-person accounts of
IEEE members who were there.

IEEE Global History Network
www.ieeeahn.org

Photo: NASA

