

Little energy for perfect pixels

Getting the right light balance in pictures can be tricky, but a new chip could make expert photographers of us all.

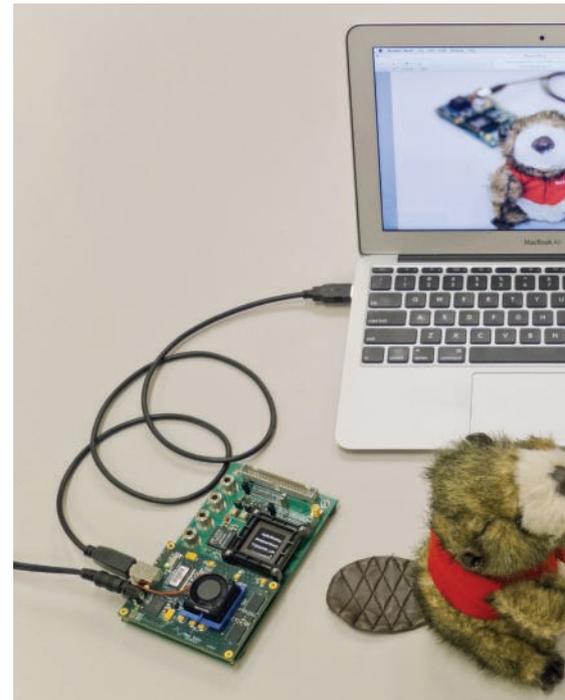
By **Tim Fryer**.

Whatever the problem, or quest for improvement, with image processing, the trend is to try and find solutions in software. The problem is that software complexity turns graphics processors into power hungry beasts.

Rahul Ridhe, a graduate student at MIT, is working towards his thesis looking at efficient systems for portable multimedia processing. He is one of the team that has developed a chip that enable huge strides in photograph and video quality.

“The processor that we have built is equally

applicable to any energy-constrained devices,” commented Ridhe. “Smartphones and cameras are great examples, but it could also apply to tablet computers or even laptops. What this chip offers is real time functionality while offering extremely low power compared to your smartphone or laptop processor. For example, if you are trying to do high dynamic range imaging on your laptop computer, it currently uses several Watts. With this new chip, you can do that with a few milliWatts – the energy reduction is more than 1000 times. So this chip can be used for all kinds of battery constrained devices.”



Such improvements in energy efficiency result from real-time image processing in hardware, rather than using software. Equally, the new chip adds functionality to the camera, allowing photographic applications such as lightfield photography, in which pictures can be ‘created’ in difficult lighting conditions that would not have been possible with a traditional camera (see left). The principal technique is bilateral filtering, which opens the door to a range of applications, including High Dynamic Range (HDR) imaging, Low Light Enhanced (LLE) imaging, tone management and video enhancement.

There are two techniques for creating pictures using HDR. The first involves taking a number of pictures – typically three – virtually instantaneously using a regulation camera. The second option – soon be available in some new digital cameras – is to use three sensors to take the three pictures instantaneously. In either case, one picture captures a normal shot, one

captures the brightest parts of the shot and the other the darkest. Each of these pictures has a low dynamic range, but a single HDR image can be obtained by combining them – and this is what the MIT chip, nicknamed Maxwell, can do.

Maxwell was named after Scottish scientist James Clerk Maxwell. Maxwell was prolific in his



scientific and engineering output – his work on electromagnetism was of huge importance to the modern world. But he also developed the first technique for creating colour photography. His trichromatic process split a picture into three images using red, green and blue filters, then recombined them in a single image. And this is the same fundamental concept used by Rihde for the image processing chip, hence the tribute.

Maxwell (shown right) is an ASIC built using TSMC's 40nm CMOS technology. Rihde commented: "The chip was developed by us completely, but TSMC has a University Shuttle Programme that allows universities to fabricate their ICs. Foxconn provided funding for the project. We had discussions during the project with them and had some feedback during the design process, but the development was done by MIT."

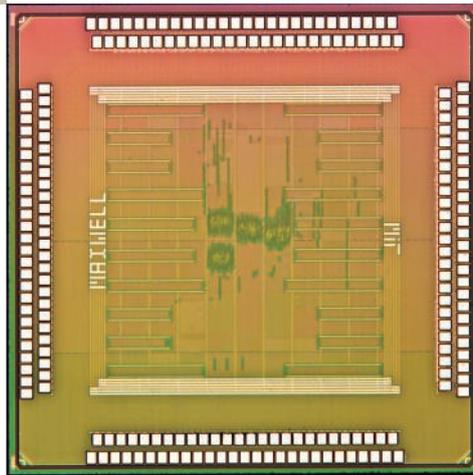
The test chip is verified to be operational from 25MHz at 0.5V to 98MHz at 0.9V. It is designed to function as an accelerator core as part of a larger microprocessor system, using the

system's existing DRAM resources. For standalone testing of this chip a 32bit wide 266MHz DDR2 memory controller was implemented using a Xilinx XC5VLX50 FPGA.

In tests that compare the runtime for a 10Mpixel image with GPU/CPU implementations of C++ code, that replicates the functionality of the testchip, the processor achieves an x15 reduction in run time compared to the CPU implementation, while consuming 17.8mW, a significant reduction compared to previous CPU or GPU implementations.

Bilateral filtering is a non-iterative process for smoothing images, while still preserving edge integrity. Rihde noted: "In this work, we implement bilateral filtering using a reconfigurable grid, which reduces the storage requirement to 21.5kbyte [compared to 65Mbyte for a 10Mpixel image using software filtering] by scheduling the filtering engine so that only two grid rows need to be stored at a time. The implementation is flexible to allow varying grid sizes for energy/resolution-scalable image processing.

"The reconfigurable filtering engine performs HDR imaging, LLE imaging and glare reduction. The filtering engine can also be accessed from off-chip and used by other applications. The implementation accelerates bilateral filtering



significantly and enables various edge-aware image processing applications in real time on HD images. The testchip can also process a 10Mpixel image in 771ms with 17.8mW power consumption while operating at 98MHz/0.9V."

The testchip contains two bilateral filter engines, each processing 4pixel/cycle.

Displaying HDR images on LDR media requires tone mapping that compresses image dynamic range by non-linear filtering. A tone-mapped HDR image is created by bilaterally filtering HDR intensity values in the log domain, followed by contrast reduction. In HDR mode, both bilateral grids are configured to perform filtering in an interleaved manner, where each grid processes alternate blocks in parallel. Glare reduction is similar to performing single image tone mapping and is integrated with the HDR architecture. LLE imaging is performed by merging two images captured in quick succession, one taken without flash and one with flash. The bilateral grid is used to decompose both images into base and detail layers. In this mode, one grid is configured to perform bilateral filtering on the non-flash image and the other to perform cross-bilateral filtering on the flash image using the non-flash image.

There is no critical lower limit to the picture – if the picture is very small, there won't be any detail in it and changes will not be detectable. On the higher side, the processor can handle up to 16Mpixels.

Although Maxwell is currently operated and tested through a laptop (see main image), its value will not be realised until it can be designed into the new breed of camera-bearing devices. While suited to the digital SLRs featuring the triple sensors mentioned above, a more important target could be smartphones, where battery life is emerging as a more important battleground parameter than functionality.

But would a camera featuring Maxwell appeal to everyone? Photographers may feel that this level of processing would remove the skill in taking good photographs and remove some of the effects that would have been deliberately introduced. Not so, according to Rihde: "The functionality can be enabled – it doesn't always have to be used. You don't want to be too prescriptive. You could enable the functionality and see the results in real-time as you would an ordinary picture. The real-time performance allows you to apply these techniques to video as well. So if you were shooting a video with a DSLR, you can do a high dynamic range video by using the chip as it will process the frames in real-time. That is another advantage of having a dedicated processor rather than doing it on a computer or doing it on a general-purpose processor."