from *SIAM News*, Volume 36, Number 6, July/August 2003

Novel Imaging Systems Rely on Focus-Free Optics

By Dana Mackenzie

You wouldn't ask Yitzhak Perlman to play a few sour notes on his violin, and you probably wouldn't ask diver Greg Louganis to do a belly-flop. But a Colorado company recently managed to convince the legendary German lensmaker Carl Zeiss Inc. to manufacture an imperfect lens—a lens so bad that it doesn't just focus at the wrong place, it doesn't focus *anywhere*.

It's all part of a revolutionary idea in optics called "wavefront coding," in which optical engineers deliberately sacrifice image clarity in return for increased depth of field. The earliest models of the novel lenses seemed to be straight from a calculus textbook, with shapes based on graphs of cubic polynomials. The lenses are still often called "cubic phase masks," even though they now have more sophisticated shapes that are found by computer optimization.

When seen through a cubic phase mask, an object doesn't move into and out of focus; it stays uniformly out of focus over a wide range of distances. But the image is not meant to be seen in this raw form. Instead, the image will be processed by a mathematical transformation called deconvolution, which removes the blurriness. Because the amount of defocus is roughly constant—the unique property of the cubic phase masks—the same deconvolution process will work, and give a sharp image, no matter how far away the object is. In other words, the lens does not have to be focused.

"It's such a non-intuitive technology that most people flat can't believe it," says Merc Mercure, the CEO of CDM Optics in Boulder, Colorado. "They also don't believe it when we show them simulations. We have to show them a physical demonstration to convince them." However, Mercure and his colleagues, Thomas Cathey and Ed Dowski (the "C" and "D" of CDM Optics; Cathey is also an engineering professor at the University of Colorado), are starting to win some converts. Bob Plemmons, a mathematician at Wake Forest University, is working on adapting cubic phase masks to iris scanners—possibly an ideal application for our security-conscious times. Further in the future, biologists might use them to add a third dimension to microscope images. And some day, if cell-phone makers catch on to the idea of focus-free optics, the wacky lenses could become nearly ubiquitous.

According to Joseph Mait, an electrical engineer at the Army Research Laboratory in Adelphi, Maryland, wavefront coding isn't just about fancy lenses. "It's a change in the way we look at imaging systems," he says. "There are three components to an imaging system—the optics, the detector, and post-detector processing. If you look at the system as a whole, you can distribute the work among the three elements"—an approach that Mait sometimes refers to, lightheartedly, as "holistic" imaging. Mercure underscores this point as well: "It's very important to understand that wavefront coding is an end-to-end systems technology. You have to understand the complete characteristics of the system, from when light enters the lens to when it exits the detector."

Four(ier) Eyes Are Better Than Two

Engineers describe the image produced by any imaging system, whether in focus or not, as a combination of three ingredients: the actual scene, the optical qualities of the lens, and random noise from the atmosphere or from imperfections in the detector. For the simplest case of a two-dimensional, black and white image, this combination can be represented by the following formula:

$$I(x, y) = O(x, y) * h(x, y) + n(x, y),$$

where $I(x, y)$ represents the intensity of the light in the image at point (x,y); $O(x, y)$ represents the true object intensity; $n(x, y)$ represents the noise; and * represents the convolution operation. The most important ingredient in this formula is $h(x, y)$, the *point spread function* (PSF) of the imaging system, which expresses the way the system smears out a single point of light. If the system were perfectly in focus, $h(x, y)$ would be a delta function, concentrated at one point. In reality, no lens can ever be perfectly in focus because of diffraction effects, which make the PSF develop fine ridges. The point spread function of a typical out-of-focus lens would consist of broad circular rings, and the PSF of the lenses developed by CDM Optics have even more complicated shapes, such as triangles or wedges.

Mathematically, the effects of the lens can be disentangled from the image by applying a Fourier transform to the basic equation above, which converts the convolution into an ordinary product. If there were no noise, you could—in principle—perfectly compensate for the lens, no matter how in-focus or out-of-focus it was. You would take the Fourier transform of the image, divide by the Fourier transform of the point spread function (this Fourier transform is called the *optical transfer function*, or OTF), and then transform back again to the image space.

In practice, it isn't quite that easy to "divide" by the optical transfer function, which for many frequencies has a value equal to or near zero. The practical effect of dividing by near-zero is that any noise in the system is amplified, and the reconstructed image looks like what you'd get from a TV station beyond your antenna's range. You might see the ghost of an image, but it would swamped by the static. In recent years, however, mathematicians have developed "regularization" techniques that work very well at separating the noise from the true image, often by incorporating prior knowledge of what the image is supposed to look like (a building rather than a snowstorm, say). Thus, to a large extent, engineers can now compensate for imperfections in a camera's optics by clever post-processing. This was demonstrated by the Hubble Space Telescope, which functioned serviceably if not optimally for three years, in spite of a gross optical defect, before astronauts repaired it.

A Warped Idea

Compensating for imperfections makes perfect sense. But deliberately building them in? That was an unheard-of concept until the mid-1990s, when Ed Dowski began his graduate research with Cathey. As is so often the case, Dowski was inspired by an idea from a seemingly unrelated area of science.

"I had worked for Raytheon on over-the-horizon radar before that," Dowski says. He knew, therefore, about a function called the ambiguity function, which radar engineers use to express the tradeoff between knowing the range (distance) to an object and knowing its Doppler shift (the speed at which it is moving). A few optical researchers had noticed that this function was mathematically identical—after a simple change of coordinates—to the optical transfer function. Some change in terminology was required: substitution of "depth of field" for "range" and of "misfocus" for "Doppler shift." But the implication, Dowski realized, was the same: Depth of field could be achieved at the expense of focus. He and Cathey were the first to come up with a practical way to do this. They derived the shapes of the earliest cubic phase masks directly from the mathematical formulas for optimal radar "chirps."

Unlike most people, Dowski didn't need a physical prototype to be convinced that his idea would work. "I was convinced by the mathematics," he says. "Just by visual inspection of the ambiguity function, [you can see] that an optical system with a cubic phase at the aperture will have an increased depth of field. It's that simple and powerful."

But with the help of Joseph van der Gracht at the Army Research Laboratory, physical proof was not long in coming. Van der Gracht made a prototype phase plate, in the shape of the cubic surface $f(x, y) = x^3 + y^3$, that Dowski and Cathey tested against more standard lenses. In conventional cameras, depth of field is usually increased by decreasing the aperture size or increasing the "f-stop." (Think of squinting in order to see more clearly.) This is not a perfect solution, however, because it decreases the amount of light coming into the imaging system, thus making the image dimmer or requiring a longer exposure time. With their unconventional lens and post-processing, Dowski and Cathey can increase the depth of field of many kinds of imaging systems by a factor of six or more. This would be comparable to increasing the f-stop by the same factor—and increasing the exposure time by a factor of 36. For some applications, such as imaging a living (and moving) cell under a microscope, getting the same image in 1/36 of the time is a huge step toward practicality.

A box of crayons, photographed through a conventional lens with the f-stop set at 8. The crayons in the midrange are in focus, but the close and distant crayons are not—the conventional lens does not have enough depth of field to get them all in focus.

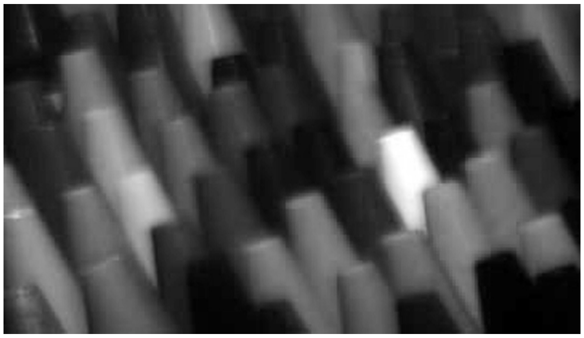

Crayons seen through a cubic phase plate, unprocessed and uniformly out of focus.

After processing, crayons seen through a cubic phase plate are all in focus. Images courtesy of CDM Optics.

From Tanks to Cell Phones

Mait points out that the Army is interested in wavefront coding for a variety of reasons. "In the M-1 tank, someone stands up in the turret and navigates, while someone else drives. In the future, they don't want that," he says. "They want to solve the problem of being able to navigate and detect long-range threats at the same time, keeping near and far objects in focus with one system." Because the crew would be inside the tank, the Army wanted to consider optoelectronic systems, in which a camera image would be transmitted to the crew electronically. Because of the depth-of-field requirements, wavefront coding was an obvious choice. In this case, the electronics would not simply transmit the image, but would also perform the mathematical calculations to transform it from the raw form into a sharp image.

Although the Army Research Laboratory constructed a successful prototype for the Tank and Automotive Command, it has not yet gone into production. It's not a question of whether the system works, Mait says. "It has more to do with whether the system solves the military problem as well as the technical problems."

Another technology in which depth of field is important, and the fuzziness of the raw (pre-processing) image less so, is iris recognition. Although discussed for years by science fiction writers (moviegoers may remember the ubiquitous iris scanners in last year's sci-fi thriller *Minority Report*), iris recognition is just beginning to take off as a real security measure. The first iris-recognition system in an airport was installed in Amster-dam's Schiphols Airport in 2001. A frequent flyer inserts a "smart card" into the machine and then looks into a mirror. A person whose iris pattern matches the data on the card can skip the passport line, avoiding a 10- to 15-minute wait. At present, about 500 passengers a day use the service. Boston's Logan Airport and New York's John F. Ken-nedy Airport also have iris-recognition systems, but only as an optional screening device for employees.

Iris-recognition systems have unparalleled sensitivity—far better than fingerprints or any other "biometric" you can think of.

Researchers have never found a single false match. False negatives, however, do happen with some regularity, according to Harald Bresser of Dartagnan Biometrics, which developed the software for the Schiphols Airport installation. “The system at Schiphols has a false rejection rate of 0 to 1%,” Bresser says. “For people who use it a lot, false rejections are very rare.” But people who are unfamiliar with the system have at least a 1% chance of being turned back. The problem is that the least expensive systems, which run around $150, are simply cameras with a fixed focal distance and a shallow (typically 4-inch) depth of field. New or infrequent users have more difficulty positioning themselves at the right distance from the camera, even with the recorded voice that coaches them. The problem can be solved by equipping the camera with auto-focus—but that drives the price up to about $1200.

“Less than $100 is what the government wants,” Bob Plemmons says. With a price tag like that, security systems based on iris recognition would become practical not only for buildings but also for computers, for which they could replace or supplement passwords. Plemmons believes that equipping the camera with a cubic mask might eliminate the need for auto-focus without driving up the price tag, because cubic masks, if mass-produced, would cost about the same as ordinary lenses. He and his group are designing a prototype with funding from the Army Research Office, and expect to finish this aspect of the work in 2004. (They are also working on more general kinds of “pupil phase engineering” that can correct aberrations other than misfocus.)

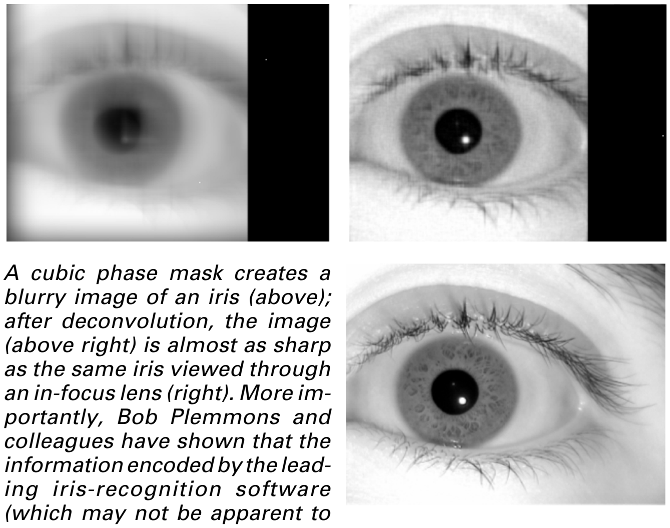

A cubic phase mask creates a blurry image of an iris (above); after deconvolution, the image (above right) is almost as sharp as the same iris viewed through an in-focus lens (right). More importantly, Bob Plemmons and colleagues have shown that the information encoded by the leading iris-recognition software (which may not be apparent to the naked eye) is also preserved after wavefront coding and deconvolution.

Mercure considers the application of wavefront coding to iris recognition a “done deal.” Unfortunately, iris-recognition systems themselves are an unproven niche market. The market he dreams about is video cell phones, which sell in the millions per year. Even more than iris scanners, cell-phone cameras need to capture images at a variety of distances without elaborate zooming equipment. “This technology is very well suited to them,” Mercure says. “In our view, cell phones could be a holy grail.” Soon you, too, may see how the world looks through a warped lens.

Dana Mackenzie writes from Santa Cruz, California.