Ravi Athale and Demetri Psaltis

OPTICAL COMPUTING

Past and Future

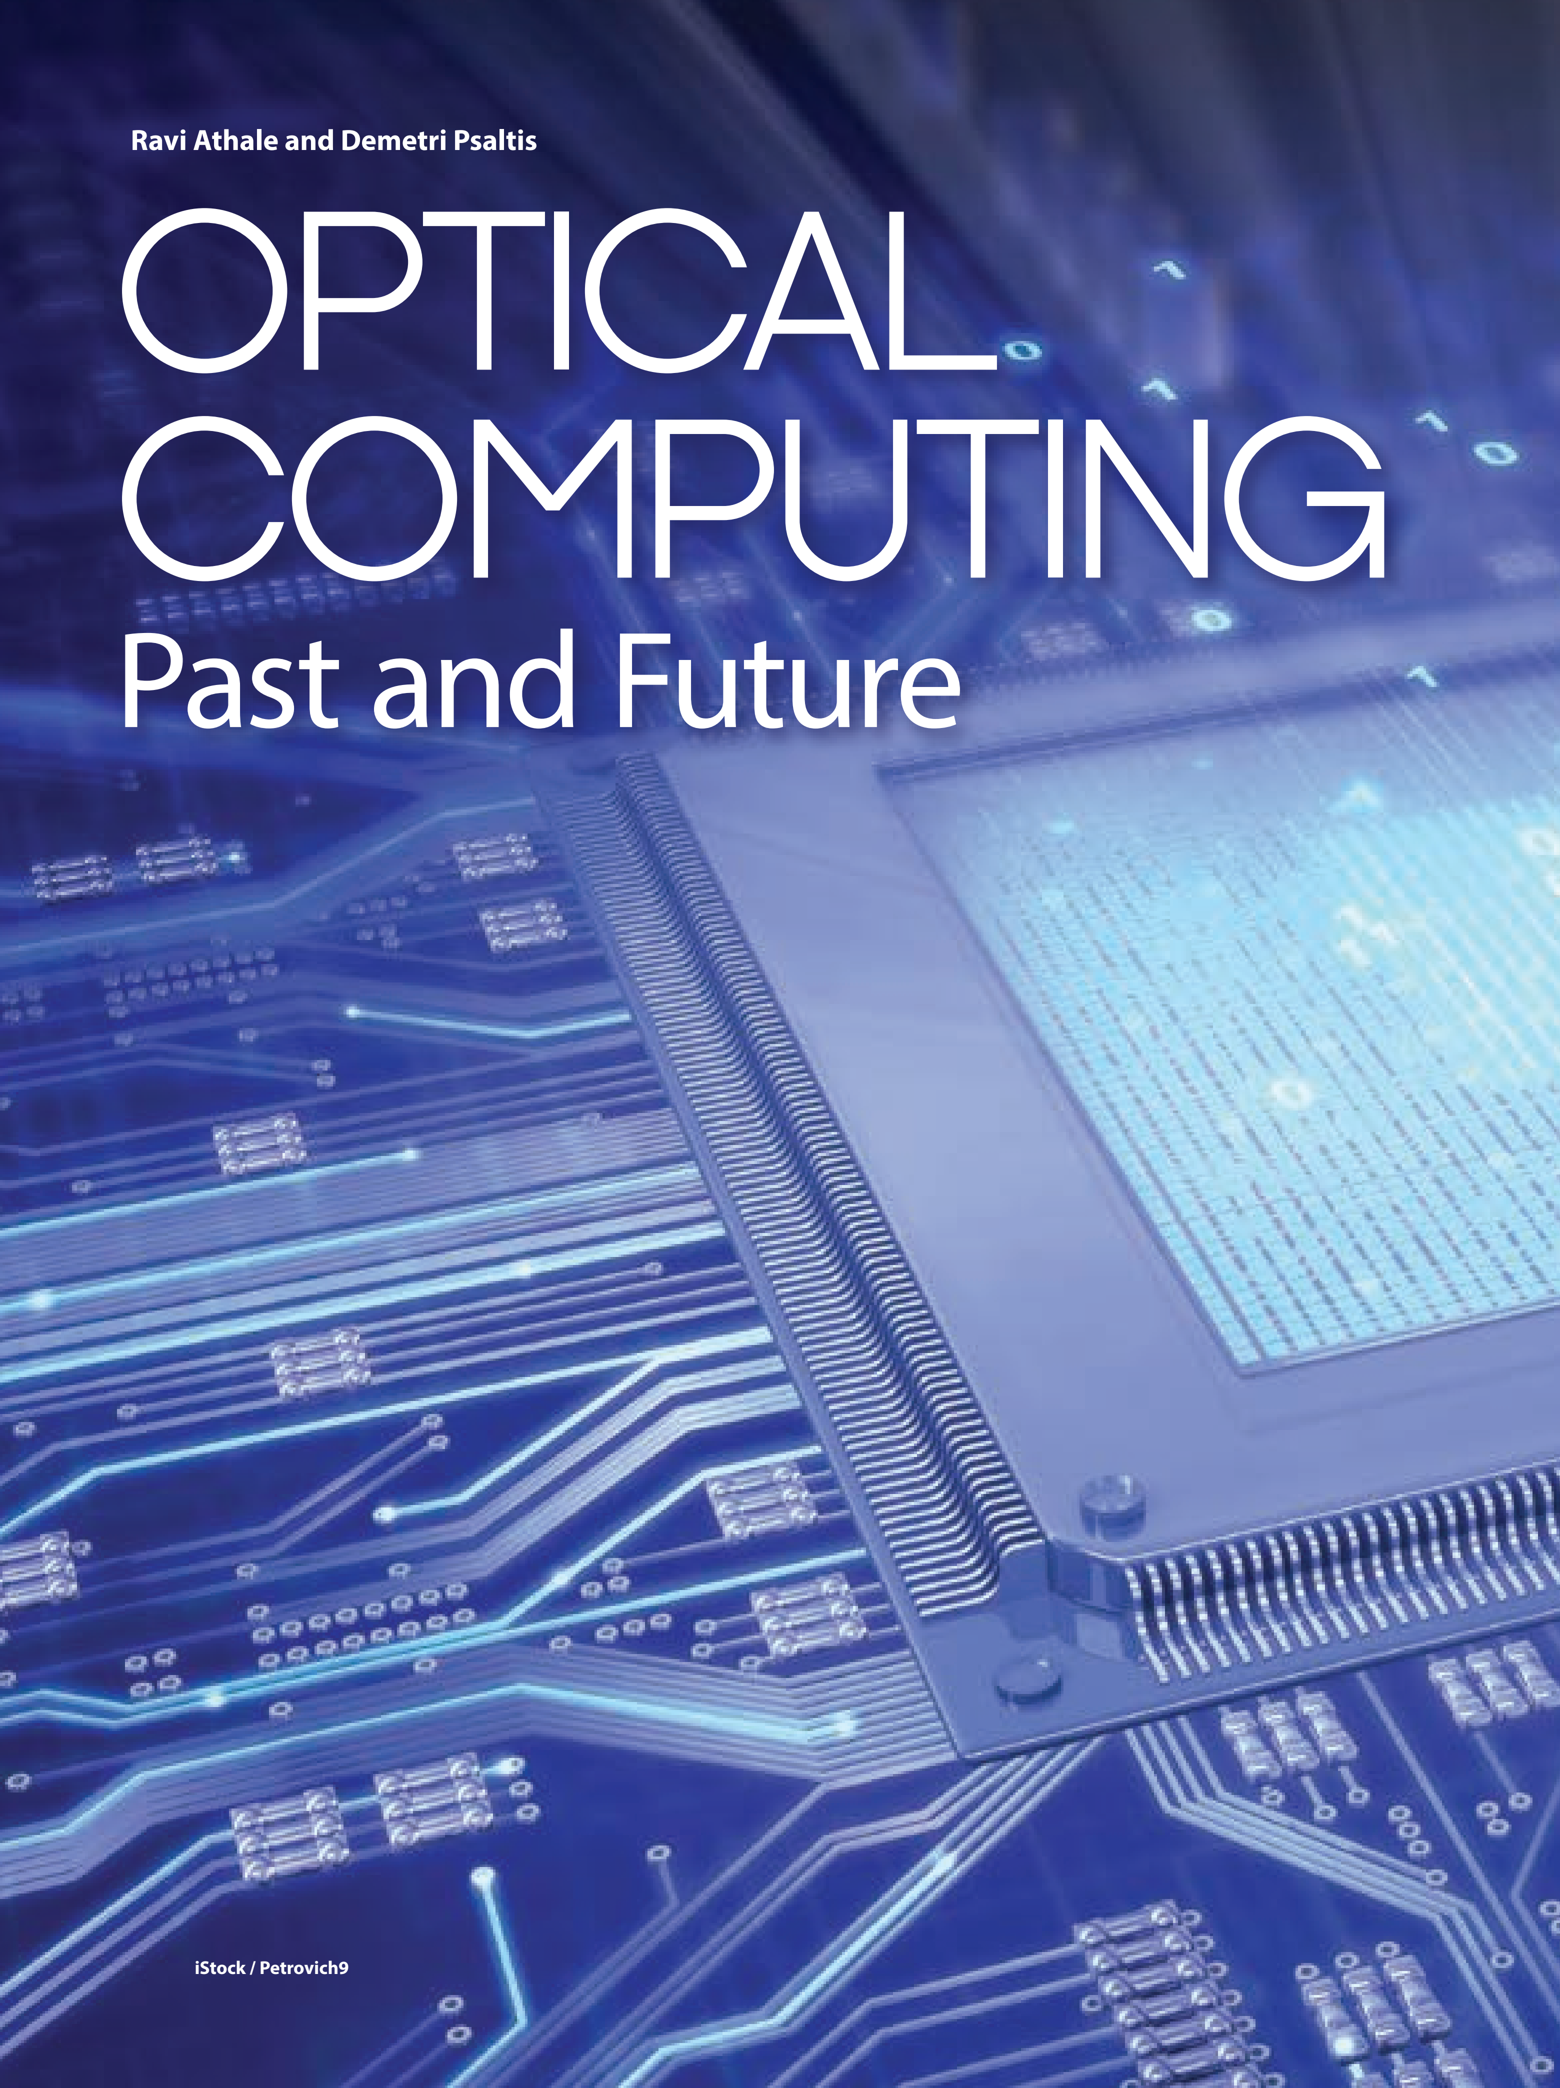

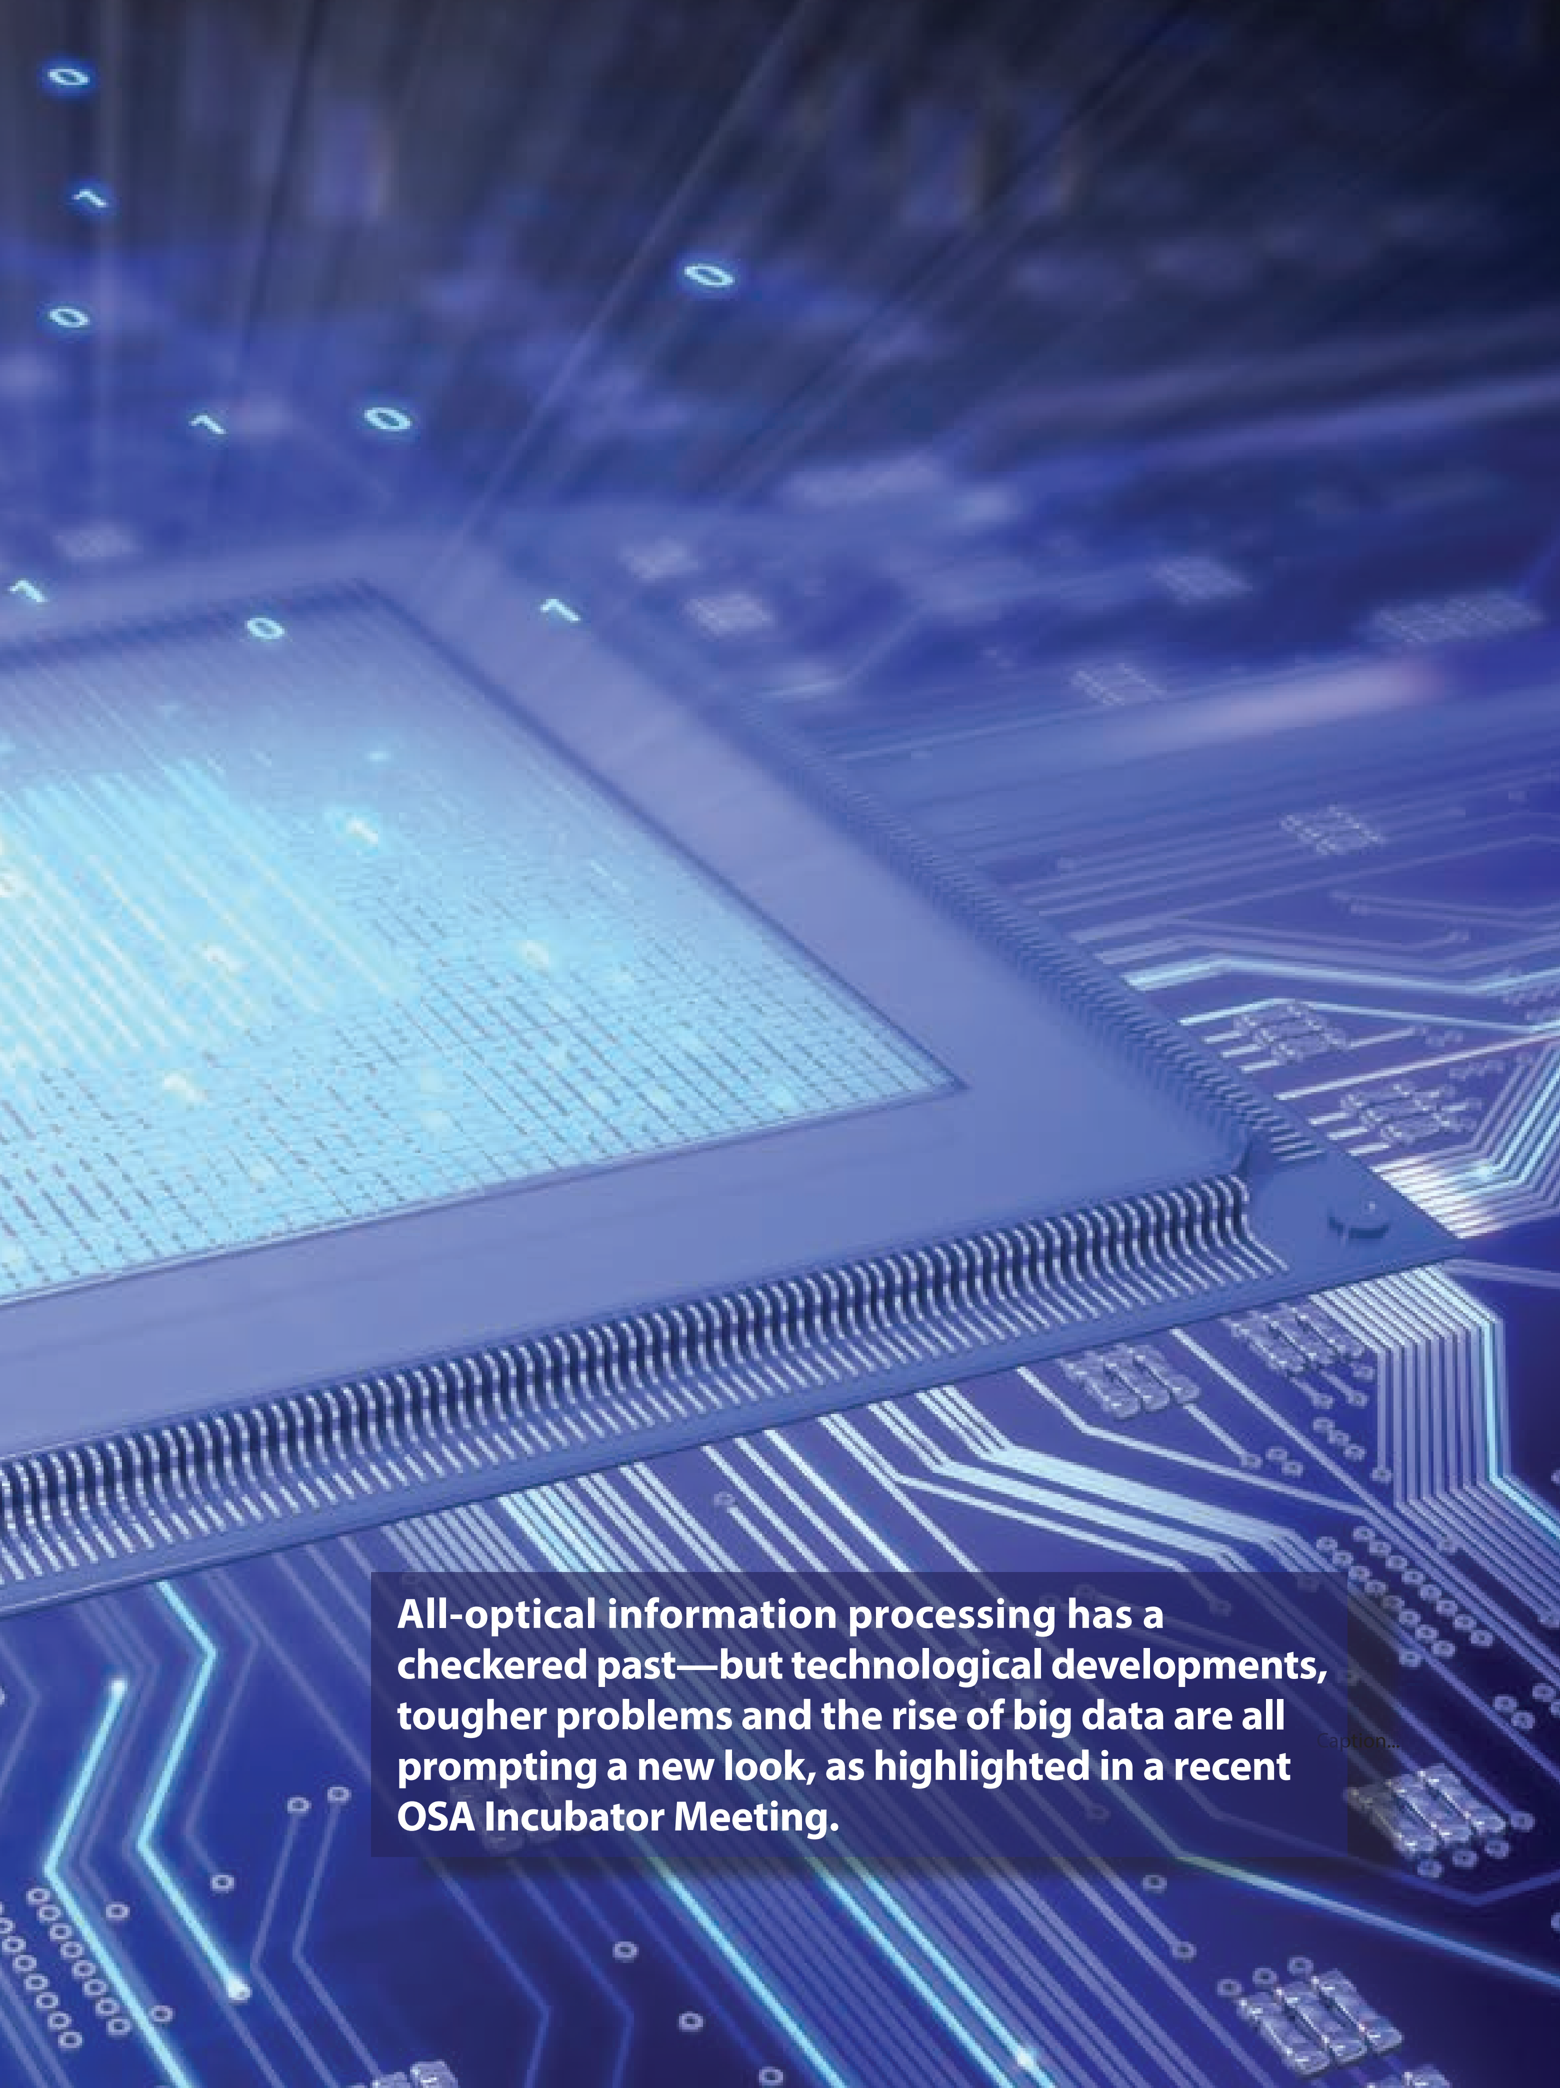

All-optical information processing has a checkered past—but technological developments, tougher problems and the rise of big data are all prompting a new look, as highlighted in a recent OSA Incubator Meeting.

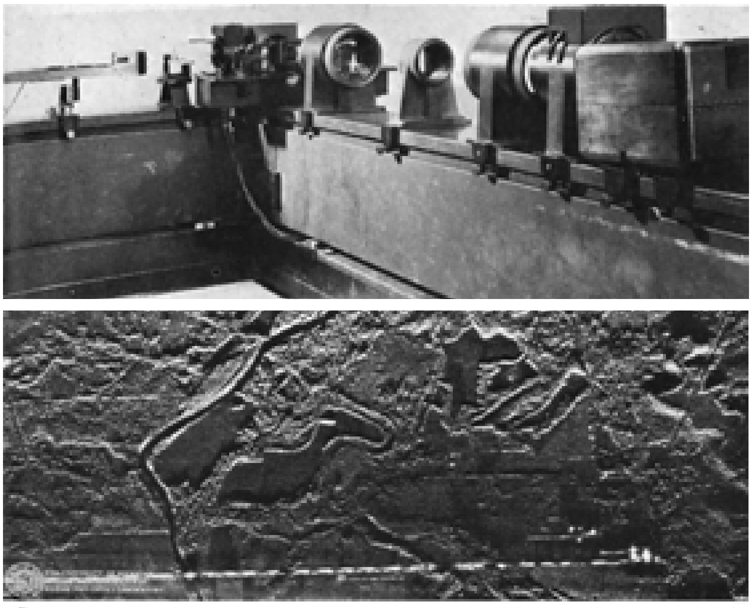

Early analog optical computing: In a classic 1972 paper in *Applied Optics*, Adam Kozma, Emmett N. Leith and Norman G. Massey described an all-optical setup built of "multilens, anamorphic telescopes including spherical and cylindrical lenses" (top) that could process synthetic-aperture radar data into photo-like images (bottom).

Kozma et al., Appl. Opt. **11**, 1766 (1972)

Optical techniques, used extensively to communicate, store, display and sense information, thus far have not found widespread acceptance in processing it—that is, in computing. Here, we look at some of the reasons that optical computing has yet to gain traction—and at possible future directions for optical-computing research. (We exclude discussion of recent research activities exploiting photons with quantum entanglement—optical quantum computing—and focus on systems that implement optical computation via more classical behaviors.)

Given the historical strengths of electronics, why did the notion of optical computing even arise? One answer lies in the highly localized encoding of information in optical systems.

Electronic computing's historical strengths

Conventional wisdom says that electrons compute and photons communicate. That's because the strong Coulomb interactions of charged electrons can be leveraged to perform nonlinear computations (Boolean logic), whereas charge-free photons do not interact with each other at all in free space.

An electronic computer operates at baseband by manipulating the flow of charges in semiconductors, such as silicon, whereas most optical systems transfer information encoded on a carrier frequency of several hundred THz by the polarization of bound electrons in dielectric materials, such as glasses. Indeed, the lack of photon-photon interaction makes it possible to use a large number of spatial and spectral channels to increase the information-carrying capacity of optical communication systems.

Electronic telecommunication systems evolved from telegraphy, with baseband operation using simple metal wires for transmission, through telephony, whose increased demand for transmission bandwidth was initially met by higher carrier frequencies that require more complex guiding structures such as coaxial cables, and finally into fiber optic technology for long-haul communications beginning in the late 1980s. The starting point for computers was similar to that of early communications systems—that is, simple electrical circuits operating at baseband frequencies, with bandwidth of several MHz.

Unlike communications however, computers have continued to operate at baseband—and, thus far, with great success. Computers have acquired complexity and speed through improvements in the resolution of lithography, which allowed exponential gains under Moore's law. Transistors with ever-shrinking dimensions provide a highly localized interaction between electrical signals (typically the gate and source voltages), through the $1/r^2$ drop of the electric field established by the charge at the gate of the transistor. Such localization is essential for Boolean logic, in which only two bits typically interact at a time.

The miniaturization of transistors made possible by ever-finer-scale lithography leads to

1047-6938/16/06/32/8-$15.00 ©OSA

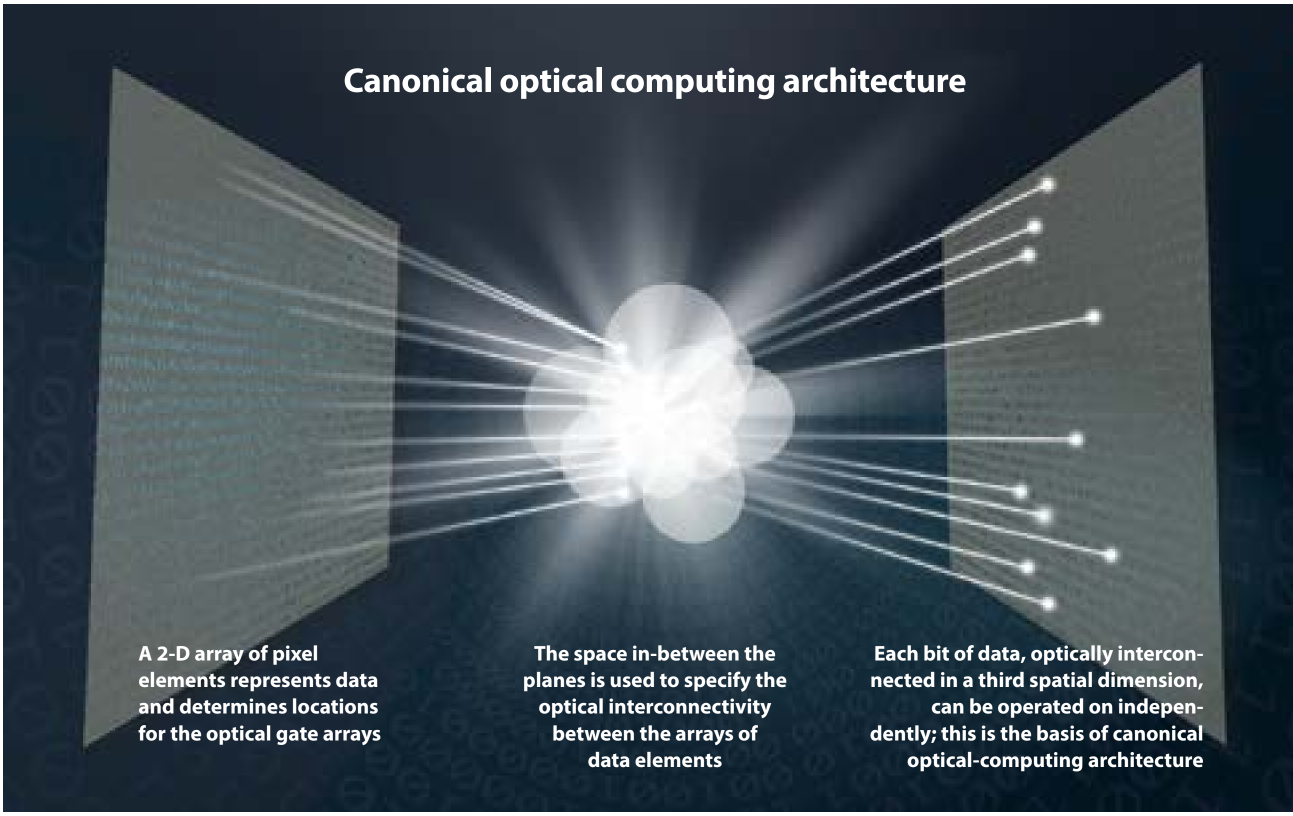

Illustration by Phil Saunders

increased speed, greater density, lower power and lower cost (through increased integration)—all at the same time—and has formed the basis of the microelectronic revolution. In contrast to communications, placing the data on a high carrier frequency before performing nonlinear logic computations has offered no obvious advantages.

Optical computing: Early successes and false starts

Given the historical strengths of electronics, why did the notion of optical computing even arise? One answer lies in the highly localized encoding of information in optical systems.

The transverse spatial resolution achievable in optical circuits is roughly equal to the wavelength of the light—smaller than 1 μm in the visible part of the spectrum. Thus a one-millimeter-square area conceivably might pack in some 10^6 optical data elements, each representing one bit or an analog value, that can be operated upon independently of all the other data elements in the plane. These data elements can then be optically interconnected in the third spatial dimension, an arrangement constituting the canonical optical-computing architecture.

In 1960s and 1970s, when optical computing was first proposed, such densities were far greater than those that could be achieved with electronic gates. And the 3-D interconnection capability of free-space optical propagation offered a unique, clear competitive advantage over electronic computers of that time. Those early advantages initially played out in research in both the analog and digital domains.

Analog processors. The earliest optical computers, in the 1960s, were analog processors performing image and signal processing tasks—the most successful example being synthetic aperture radar (SAR). SAR required a substantial amount of computation to form a high-resolution image from a series of radar returns collected by an aircraft. At the time, photographic film provided the only viable mean for storing such vast quantities of data compactly, and in real time. Since the data were already in an optically accessible form, elegant optical systems were developed to create the high-resolution radar images by exploiting mathematical transformations that described propagation of light in free space as well in lenses.

Analog optical processors based on the mathematical equivalence between optical diffraction and the Fourier

Despite the mixed record of early efforts, today's data-intensive information environment raises some new possibilities for optics in computing.

transform and other useful linear operations were also developed for image power spectrum measurement, 2-D image correlation, RF frequency analysis and high-bandwidth signal correlation. These processors achieved varying degrees of success in specialized military and commercial settings—but, by the mid-1980s, rapid progress in digital signal processing technology made such linear analog optical computers seem largely obsolete.

Optical logic gates. The early 1960s also saw the first research in optical logic, driven by the advent of semiconductor lasers. Optical logic devices seemed to have the potential for much faster performance than their electronic counterparts, as the optical devices don't suffer from RC time constants but are instead governed by excited-state lifetimes. But by the early 1970s, the thermal limits on injection laser logic and the Moore's-law progression of conventional electronics caused those efforts to be abandoned.

Research in optical logic regained some momentum in the 1980s, however, owing to an increasing focus on the parallelism and unique topologies that 3-D, free-space optical interconnects could provide. The canonical configuration of analog optical processors (planes of data optically interconnected in the third dimension) was used, with one important modification: a nonlinear optical gate was introduced at each pixel in the 2-D planes of the optical computer. Two (or more) optical beams converged at each pixel and interacted in the nonlinear optical material with which the gate is fabricated.

Several optical gate technologies were developed in the late 1980s and early 1990s to implement optical logic, most notably the self-electro-optic effect device (SEED), which formed the foundation of system demonstration platforms at AT&T's Bell Laboratories (USA). A SEED generates free electrical charges in response to illumination, which in turn modify the electric field in the material and change its absorption. In this way the presence of one light beam modifies a second, and logic operations can be synthesized.

SEEDs require the movement free charge carriers to achieve the desired nonlinear optical interaction. Their performance, therefore, is limited by the same charge-transport restrictions encountered by electronic gates. In the late 1980s, the speed, size and switching energy of optical switches were comparable to those of electronic transistors, and the idea of combining optical gates, with 3-D optical interconnects replacing metal wires, seemed like a winning proposition. But electronic transistors continued to become smaller, faster and more power efficient, whereas optical gates had to remain larger than the wavelength of light. As a result, by the mid-1990s, active research on digital optical computing, as with analog computing, had been largely abandoned.

Taking another look at optical computing

Despite the mixed record of these early efforts, fast-forwarding to today's data-intensive information environment raises some new possibilities for optics in computing. One, of course, is optical interconnects.

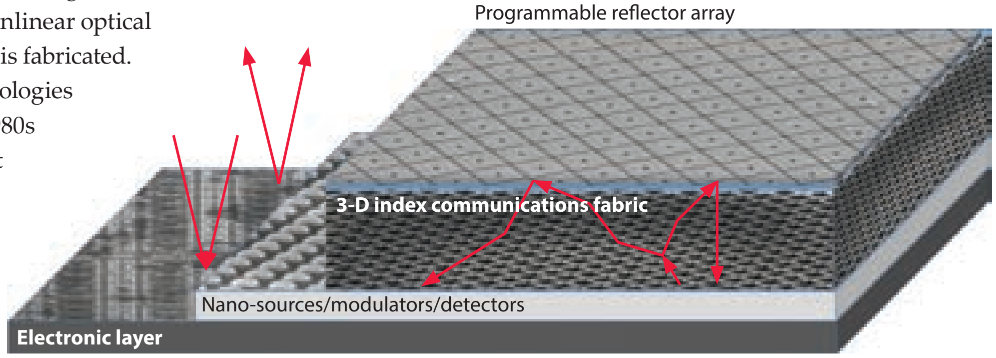

Optical-electronic integration

Schematic 3-D optical interconnect fabric, integrated with electronic processing elements on a printed circuit board. Using 3-D propagation of light in engineered media could allow connection topologies and densities not possible with 2-D electrical technology.

Mark Neifeld, University of Arizona

OSA's Optical Computing Incubator

From 9 to 11 December 2015, OSA held an Incubator Meeting on optical computing, hosted by Ravi Athale, Mark Neifeld and Demetri Psaltis. The Incubator brought together some 40 scientists and engineers from academia, industry and government to look at the past and future of optical computing. OPN talked with Neifeld about the meeting's agenda and outcomes.

Q. Mark, why this Incubator—what makes now a good time for this kind of meeting, and what were you trying to achieve?

During the 1980s, a community of researchers investigated the prospects for optical computing, and eventually found that optics was uncompetitive with electronics. Ravi, Demetri and I felt that a lot had changed during the intervening years. For example, the new field of nanophotonics can provide compact low-power light sources and other nonlinear devices; photonic integration technologies have matured since the 1980s; the end of Moore's-law gains is approaching, and will require new technologies; there is already a lot of data communicated in the optical domain; and new computational problems are arising that continue to challenge existing computing paradigms.

So it made sense, in 2016, to reconsider the potential for optics to serve as a viable computing substrate. By bringing together experts in both optical and electronic computing in this Incubator, we hoped to encourage collaboration toward solving some important computational problems—and to identify opportunities for optics in those solutions.

In particular, we wanted to look at where traditional digital computing approaches are failing, what characteristics are causing those problems, whether alternative problem formulations can make these issues more amenable to optical implementations, and what advances in optical materials and device technologies might be necessary to implement those alternative formulations.

Q. Can you tell us something about the meeting's format?

The meeting format was intended to encourage discussion. We invited a few highly distinguished speakers to make presentations, and those provided some unique insights, but the panel discussions were the main focus. We wanted these discussions to seed the interdisciplinary collaborations that will underpin future optical-computing research.

The only structure that we imposed on the process was a distinction between two kinds of optical computing: algorithmic, which involves procedural methods that combine computing primitives to produce a desired output; and metaphoric, which involves nonlinear dynamical systems for which the final state represents the desired output.

Q. What do you think was the meeting's most important takeaway?

My own top four takeaway messages were, first, that contrary to what we learned in the 1980s, there may be a role for optical digital computing—especially in very high speed applications for which the data already resides in the optical domain, such as in all-optical router functionality. Second, the high cost of analog-to-digital conversion appears to make many high-speed signal processing applications, such as radar processing and broadband signals intelligence, amenable to analog optical solutions.

Third, several recent examples of optical metaphoric computation suggest that this approach offers promise, but the costs associated with those solutions (for example, the required fabrication precision) could have as-yet-unknown impacts on the quality of the eventual solution. And, fourth, power dissipation—for example, flops per watt—may be the most important metric for evaluating future computing systems. There are clear arguments for why optics will provide a lower-power communication alternative inside the box, but the jury may still be out on the question of all-optical switching.

Q. You direct the University of Arizona's Optical Computing and Processing Laboratory (OCPL)—can you tell us a bit about what's going on at that lab?

OCPL was created in 1990, at the very end of the previous era of optical-computing research. Since that time we have pursued quite a lot of non-computing research at the intersection of optical physics and information/communication theory. There was definitely resonance between some of our ongoing OCPL activities and the Incubator discussions.

For example, my group is currently part of the U.S. Office of Naval Research's Multidisciplinary University Research Initiative program on optical computing, which involves other researchers at the University of Arizona and at several campuses of the University of California. Together we are studying optical Ising machines and other nonlinear dynamical optoelectronic systems for solving NP-hard problems; optically interconnected digital and analog electronics to address power and area challenges in current all-electronic computers; all-optical graphical inference machines for solving problems in big data; and tabletop electromagnetic "wind tunnels" to accelerate large-scale FDTD simulations.

Q. How would you sum up the future prospects for optical computing?

I am cautiously optimistic. We have a pretty good idea of where conventional computing breaks and why. Optical interconnects might be viewed as an existence proof—an example of where a good match between technology capabilities of and application are making technology insertion cost-effective. The length over which optical communications is superior to electrical communications continues to decrease. When might we begin to call this 'optical computing'? There is a lot of work to do in order to understand how recent progress in optical materials and devices can provide competitive solutions to new problems in optimization and big data. Hopefully this Incubator was a small step toward building the community required to answer these important questions.

Beyond interconnects, optical computing could have significant potential in extremely large-capacity channels, such as those found in routing, switching and security applications.

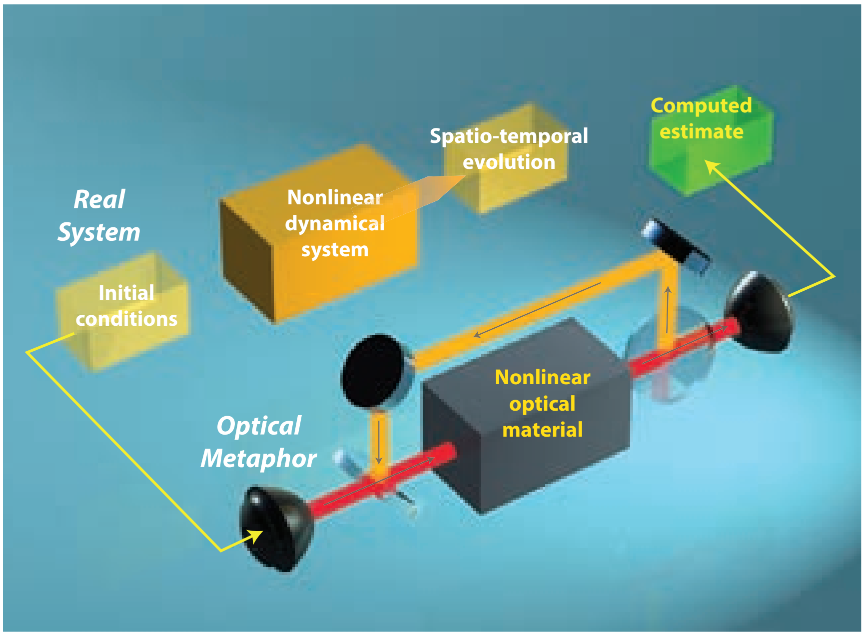

Metaphoric optical computing

In metaphoric computing, complex, nonlinear problems—such as weather—for which behavior is difficult to simulate with digital computers are analyzed and mapped onto a nonlinear optical system for computation. Illustration by Phil Saunders

Although Moore's law made it impossible for optical gates to compete as individual transistor components, the same scaling argument does not apply to wires. High-performance computing systems are increasingly performance-limited by interconnects, at levels ranging from between cabinets to all the way down to within the chip itself. As optoelectronic technologies for electrical-to-optical transduction (lasers and modulators) and optical-to-electrical transduction (detectors) have gotten simpler and cheaper, optical communications have begun penetrating the markets for shorter links in local area networks. Optical techniques are now being employed for connecting multiple cabinets in a server or router, as well as in high-performance computing environments (see "Optical Interconnects and Extreme Computing," OPN, April 2016, p. 32).

Active research is being carried out to explore optical backplanes for board-to-board and chip-to-chip interconnects, with proposals for investigating on-chip optical networks. The recent establishment of the American Institute for Manufacturing Integrated Photonics (AIM Photonics), a US$600 million U.S. public-private partnership announced last July, highlights the drive to make this kind of photonic integration a practical reality.

Beyond interconnects, optical computing could have significant potential in extremely large-capacity channels, such as those found in routing, switching and security applications, in which the information is already in optical form, and all-optical processing could avoid the potentially high performance and cost penalty of optical-to-electrical and electrical-to-optical conversions. Optical microscopy represents another important area where information is naturally in the optical domain—Zernicke's dark-field technique might be thought of as an analog optical computer for spatial filtering of the image to perform linear mapping from phase to intensity before detection.

Modern optical imaging and sensing instruments have incorporated more advanced ideas whose origin can often be traced to early analog optical computing based on diffraction. Concepts such as adaptive optics, structured illumination or wavefront shaping (for imaging through scattering media) manipulate information in the optical domain to facilitate and empower post-detection digital processing, and thus to realize the ultimate system performance as efficiently as possible. The exploding field of "computational imaging" relies on the parallelism of optics to perform complex analog (linear and potentially nonlinear) signal-processing tasks while still in the optical domain, reducing the computational load on the post-detection digital processor.

Another approach: "Metaphoric" computing

A different driver for using optics in computing could arise not from resolving bottlenecks in existing computing systems, but from the increasing complexity of the scientific and mathematical problems those systems are addressing. Early optical computers implemented desirable linear transformations by mapping them to the physics of the optical system. We refer to this approach as metaphoric computing, to distinguish it from the more classical term analog computing, which entails an algorithmic (step-by-step implementation of simple primitives) approach to computation.

A broad class of computationally challenging problems, including combustion modeling, economic forecasting, and chemical and biological reactions, involves coupled nonlinear differential equations in a high-dimensional space. Mapping such differential equations directly onto nonlinear optical propagation, in engineered materials and structures, could dramatically enhance the ability to solve these difficult computational problems. In this approach, the research team designs an optical system whose spatiotemporal dynamics mimic the nonlinear physical system the team wishes to analyze, and then uses measurements on the real physical system to set the initial conditions of the optical system.

Optical systems provide a compact, high-bandwidth platform with considerable flexibility for designing nonlinear dynamic behavior, setting initial and boundary conditions, and monitoring temporal evolution through detectors. The Navier-Stokes equations that govern fluid dynamics (including weather systems), for example, can be reduced with a transformation of variables to Maxwell's equations—the propagation of light in media with negative third-order nonlinearity becomes equivalent to fluid flow in incompressible media. Granted, the challenges involved in building optical systems that actually reveal something useful about difficult fluid dynamics problems are enormous—but so is the potential payoff.

Big data and physical limits

In recent years, it's been widely recognized that conventional scaling in CMOS processors is reaching its physical limits, and can't provide the same exponential improvement in computational capabilities as in the past. The computational challenges now posed by so-called big-data analytics are also necessitating a rethinking at a fundamental level.

All of this has driven increased interest in alternatives to silicon-CMOS-based hardware for digital computation, a trend captured by a variety of campaigns by industry groups to "reboot" information technology. These initiatives envision tight integration among specific applications, alternative models of computation, and new, potentially unconventional hardware platforms. Proposals for building optical systems implementing a "reservoir" model of computation—a variation on neural-net models—constitute one recent example.

In those efforts, nanophotonics could play a key role. The same advances in lithography and manufacturing that have driven Moore's law have also, in the past decade, brought a veritable revolution in photonics technology, making it possible to precisely create features far smaller than wavelength of light. As a result, photonic-crystal structures, metamaterials, plasmonics and highly resonant nanostructures are now enabling unprecedented control over light propagation, modulation, generation and detection. Novel ideas in bottom-up self-assembly of materials are also opening new vistas in light-matter interactions through tools such as quantum dots.

These integrated-photonics developments are leading to exploration of ever-smaller, ever-higher-performance devices for electrical-to-optical and optical-to-electrical conversion. They have also rekindled interest in nonlinear optical switches as logic devices for special-purpose digital optical circuits—if not as CPUs in general-purpose digital computers. The rapidly evolving landscape of information processing—and the increasing limits faced by Moore's law—makes now an opportune time to explore such advanced optical-computing techniques. OPN

Ravi Athale (ravindra.athale@navy.mil) is with the Office of Naval Research, Arlington, Va., USA. Demetri Psaltis (demetri.psaltis@epfl.ch) is with the École Polytechnique Fédérale de Lausanne, Switzerland.

References and Resources

- J.T. Tippett et al. *Optical and Electro-optical Information Processing* (conference proceedings), MIT Press, Cambridge, Mass., USA (1965).
- L.J. Cutrona et al. "On application of coherent optical processing techniques to synthetic aperture radar," Proc. IEEE **54**, 1026 (1966).
- J.W. Goodman et al. "Optical interconnections for VLSI systems," Proc. IEEE **72**, 850 (1984).
- D.A.B. Miller et al. "Electric-field dependence of optical absorption near the band gap of quantum well structures," Phys. Rev. B **32**, 1043 (1985).
- M. Tsang and D. Psaltis. "Metaphoric optical computing for fluid dynamics," Proc. SPIE **5735**, doi:10.1117/12.601656 (11 April 2005).
- F. Duport et al. "All-optical reservoir computing," Opt. Express **20**, 22783 (2012).
- Semiconductor Industry Association and Semiconductor Research Corporation, "Rebooting the IT Revolution: A Call to Action" (September 2015), www.src.org/newsroom/rebooting-the-it-revolution.pdf
- IEEE, Rebooting Computing, http://rebootingcomputing.ieee.org/