Optical Interconnects to Silicon

David A. B. Miller, *Fellow, IEEE*

Invited Paper

***Abstract*—This paper gives a brief historical summary of the development of the field of optical interconnects to silicon integrated circuits. It starts from roots in early optical switching phenomena, proceeds through novel semiconductor and quantum well optical and optoelectronic physics and devices, first proposals for optical interconnects, and optical computing and photonic switching demonstrators, to hybrid integrations of optoelectronic and silicon circuits that may solve basic scaling and other problems for interconnections in future information processing and switching machines.**

THE IDEA of using light beams to replace wires now dominates all long-distance communications and is progressively taking over in networks over shorter distances. But should we use light beams at the much shorter distances inside digital computers, possibly connecting directly to the silicon chips, or even for connections on chips? If so, why, and, ultimately, when?

In this short paper, I will try to summarize the history of how we got to the point we are at now of seriously contemplating the introduction of optics into mainstream computing and switching as a way of interconnecting electronic chips. This history cannot avoid being one person's perspective. I hope that I get the broad sweep of events correct, and that my colleagues will forgive my omissions and misperceptions. Hopefully also, by being one person's perspective, the paper may appear more coherent to the reader (even if that coherence is merely an illusion).

The history of optical interconnects[1] is strongly intertwined with the broader field of optics in digital computing and switching. With the benefit of hindsight, aspects of optical computing have been misguided or even naïve; that field, however, provided much of the impetus for key technical innovations that make optical interconnects a serious possibility today, and, incidentally, generated some important spinoffs along the way that may more than compensate for the errors of the field.

Modern thinking about the use of optics in digital computing dates from the 1960s. The laser had been invented. The semiconductor diode laser in its earliest form had been demonstrated. Nonlinear optics was an exciting emerging field. The idea of using lasers or other nonlinear optical devices as the logic devices in computing was examined. The conclusion of these first analyzes was that optical logic devices were not a good substitute for transistors in an general-purpose computing machine (see, e.g., [1]). Their biggest problem would be that they would take too much energy. The ideas of optics for communication were being developed also at that time. The notion that the best use of optics inside computers might be for making connections was not seriously considered, perhaps because wires were still capable of delivering the performance needed.

The optics research community continued to be interested in the idea of optical switching, driven by the idea that optics could make switches that would be much faster than any electrical transistor. This particular claim remains valid; nonlinear optics is capable of making logic devices much faster than any electrical device. Much of the work was in the field of optical bistability, which flourished from about the mid 1970s to the late 1980s [2].

The observation of optical bistability in semiconductors[2] [3], [4] increased the interest in the field, because semiconductors offered more nearly practical devices than some of the other systems being investigated (e.g., sodium vapor). Smith [5] published a more optimistic analysis of optical logic devices in which he argued that the use of resonator structures could reduce the operating energy below electronic devices. The nonlinear effects in semiconductors became relatively well understood [6]. These semiconductor bistable devices were very simple, with some (e.g., [4]) requiring only a semiconductor material with plane parallel surfaces. As a result, semiconductors offered the possibility that large parallel arrays of devices might be made; a simple slab of semiconductor might be sufficient. In the early 1980s, it was found that there were particular enhanced optical properties in semiconductor quantum well structures[3] [7]–[9]. In particular, quantum wells showed very strong peaks (excitonic peaks) in their absorption spectra at room temperature, and

Manuscript received October 9, 2000.

The author is with Ginzton Laboratory, Stanford University, Stanford, CA 94305-4085 USA.

Publisher Item Identifier S 1077-260X(00)11592-8.

---

[1] Here the author should make a grammatical apology for the field. The word "interconnects" errs twice. First, adding the prefix "inter" to the words "connect" or "connection" is surely redundant. How else can one make a connection other than between objects? Second, "connect" (and hence "interconnect") is a verb, not a noun. The term comes from the electrical community, however, where its use predates any use in "optical interconnects."

[2] By a remarkable coincidence, the two papers, [3] and [4], were received by the same journal on the same day, a fact of negligible direct scientific importance, but one that fueled competitive fires at the time, and led to increased activity in the field.

[3] Quantum wells are very thin semiconductor layers (e.g., 10 nm thick), sandwiched between other semiconductor layers. Shortly after the original observation of their optical properties [8], it was understood that they would improve the performance of semiconductor lasers [9], and they are the dominant method of making semiconductor lasers today.

relatively easy saturation of those peaks. This led to further interest in semiconductor optical switching devices.

One could now imagine planes of devices communicating with large arrays of light beams from one plane to another, with the information flowing perpendicular to the plane. The numbers that could be connected this way would be much larger than the number of wires of electrical busses. Here we see the emergence of an interconnection advantage as a reason for the use of optics in computing. At that time, however, there was no practical way of getting information in the form of light from electrical circuits at high densities and high speeds, so the discussion was still restricted to interconnecting optical logic devices. There was also early work, using liquid crystal spatial light modulators, to demonstrate elementary optical computers, that exploited the two-dimensional (2-D) parallelism of optics, though the devices available did not offer any serious possibility of making a practical system [10].

The idea of optical interconnection of very large scale integration (VLSI) electronics was proposed and analyzed in a seminal paper by Goodman, *et al.*in 1984 [11]. This paper pointed out several basic reasons why optical interconnection might be interesting and important, successfully anticipated many subsequent developments, and was arguably the start of the field of optical interconnects proper.

At that time, there still was no good way of getting large numbers of optical outputs from a silicon integrated circuit. Silicon itself, because of its indirect band gap, is fundamentally a very poor light emitter and not a very good light modulator. There has been much subsequent research on making silicon light emitters. Much recent interest (see, e.g., [12]) was stimulated by relatively efficient incoherent emission from porous silicon [13]. Incoherent emission is, however, likely not sufficient for dense, high-speed interconnects, mostly because of the basic optical inefficiencies in focusing incoherent light[4] [14]. The use of silicon for efficient optical output devices has remained a very stubborn problem. III–V materials (e.g., GaAs, InGaAs) remain the only viable ones for semiconductor light emission or high-performance modulation. In the early 1980s, the optoelectronic output devices available [III–V (incoherent) light emitting diodes, and III–V edge-emitting lasers] still had relatively high power dissipations, and there was no demonstrated technology for integrating them in large numbers with silicon integrated circuits. The edge-emitting geometry of the lasers was also not attractive for use of any free-space imaging optics; to this day, the packaging of any optics with edge-emitting lasers is a difficult and expensive process.

In 1984, a new optical modulation mechanism, the quantum-confined Stark effect [15], was discovered in III–V semiconductor quantum wells. This effect allowed efficient high-speed optical modulators. The first such modulator (and also the structure used to measure the physical effects) is shown in Fig. 1. Integrated waveguide versions of such modulators are extensively used today for high-speed low-chirp modulation in telecommunication systems (see, e.g., [16]). This effect was important for both optical computing and optical interconnects for three reasons: 1) it allowed low energy devices; 2) the changes in absorption were large enough to make a modulator or switch that would work for light beams propagating perpendicular to the surface, allowing 2-D arrays; 3) it had high enough yields to allow large arrays of devices to be made (e.g., thousands). When used with optical inputs controlling optical outputs, the resulting devices were called self-electrooptic effect devices (SEEDs) [17], [18]. Initial work on using such modulator arrays concentrated on their use as optical logic devices, partly because it was not clear how to integrate them with electronics and also because there was the hope that the parallelism of optical arrays would make them interesting in practical applications even if they were logically simpler than electronic systems.

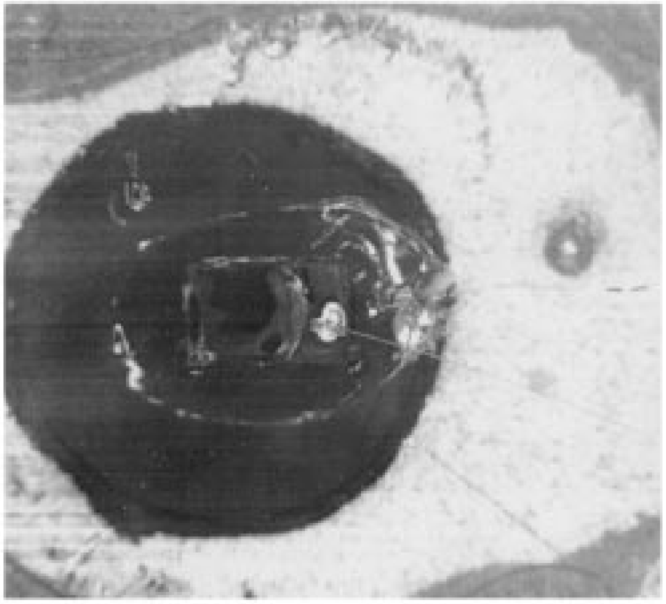

Fig. 1. The first quantum well electroabsorption modulator.

Simple optical bistability was found to be impractical for making optical logic systems of any substantial size. Simple optically bistable devices are examples of “two-terminal” logic devices, known to the electronic community to be troublesome for large systems [19]. As a result, work on optical computers based on such simple bistability largely died out in the late 1980s. The semiconductor nonlinear absorption effects that were researched for many of these earlier bistable devices did, however, find use in modelockers for short pulse lasers (see, e.g., [20]–[22]).

A variant of the SEED, the symmetric SEED (S-SEED), though bistable, was technically a “three-terminal” device by virtue of being bistable in the ratio of two beam powers[5] [23].

---

[4] If there is no predictable phase relationship between light emission from different points on a source, there is no way of persuading the light from these different parts to provide the constructive and destructive interference required to focus to a very small spot. Small spots are important because then we can use small photodetectors. Small photodetectors have small capacitance, which means larger voltage swings for a given optical input power (and/or higher speed). Higher voltage swings mean less amplification. Less amplification means less power dissipation, more noise immunity, and less latency. This difficulty of focusing incoherent sources is not merely a technological one; it is protected by the second law of thermodynamics (known in optics as the constant brightness theorem). It is not possible to take the light from two cooler black bodies and combine it with some linear optics so as to heat up a warmer black body without violating the second law of thermodynamics. Attempts to get round the focusing difficulties of incoherent sources tend also to permit this violation of the second law.

[5] These two-beam devices rely on “time-sequential” gain. Basically, the bias power is turned down to a low value, the device is able to be switched by small absolute powers in this condition, and the bias powers are then turned up to allow the device to be read out at high power, giving an effective signal gain in the switching.

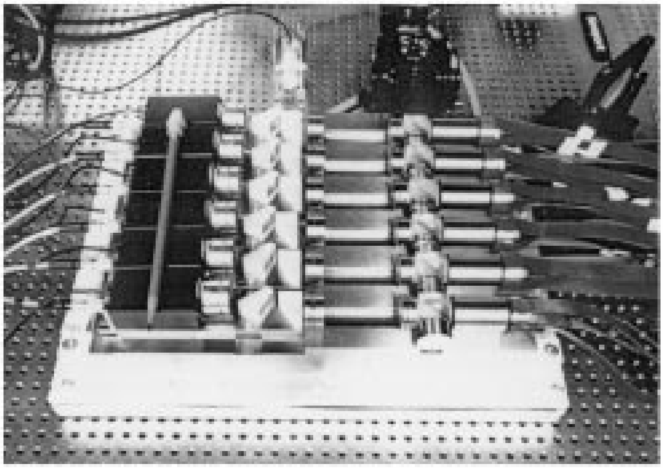

Fig. 2. A six-stage photonic switching system. The devices are mounted on slugs at the right. The units at the left are the laser power sources. The system as shown has more than 60 000 light beams, with arrays of light beams generated from each laser using diffractive optical elements.

This device allowed relatively substantial processors to be constructed. A team led by Huang [24] at Bell Laboratories demonstrated a simple but complete computer [25] with four arrays of 32 S-SEED devices. This work also led to substantial efforts on architectures and on optics for handling arrays of light beams (see, e.g., [26]). A related team led by Hinton, also at Bell Laboratories, constructed experimental switching machines for telecommunications routing. This work, which overlapped in time with the optical computing work, led to several generations of experimental systems [27], including likely the largest free-space digital optical system built so far [28], with over 60 000 light beams. This system is shown in Fig. 2. This working system demonstrated the feasibility of dense free-space optical array connections. It should be noted that, during this same time period of the late 1980s to early 1990s, there were several other important efforts and devices in such digital optics. These included the group at Heriot-Watt University (see, e.g., [29]) that had remained active from the earliest semiconductor bistability observations [4], and work by others on other optoelectronic logic devices based on light emission during this time and beyond (see, e.g., [30]–[32]).

The work on the computing and switching systems experiments exposed two issues with digital optical systems. 1) To be efficient in overall performance, the computing and switching systems both appeared to want greater logical complexity than a very simple logic function between the optical inputs and outputs. This lead to the notion of "smart pixels"—devices or units with optical inputs and outputs, but with significant functional complexity between those. 2) Complex optical systems tended to have loss that was not negligibly small, so the available powers to switch the next device in the system were less than originally hoped. This meant that either more optical power or even more sensitive devices were required. The SEEDs were already much more sensitive than other optoelectronic devices available at the time. Hence, some additional gain mechanism was required. A solution to this was to incorporate electronic gain. Both of these issues could be solved by integrating electronic devices between the optical inputs and outputs. One of the first devices to explore this solution was a monolithic integration of GaAs field effect transistors and quantum well modulators and detectors (FET-SEEDs) [33]. This approach was also successfully incorporated into a substantial system demonstration [34]. One of the "smart pixel" units in this system is shown in Fig. 3. These are likely the most complex monolithic optically interconnected systems demonstrated to date.

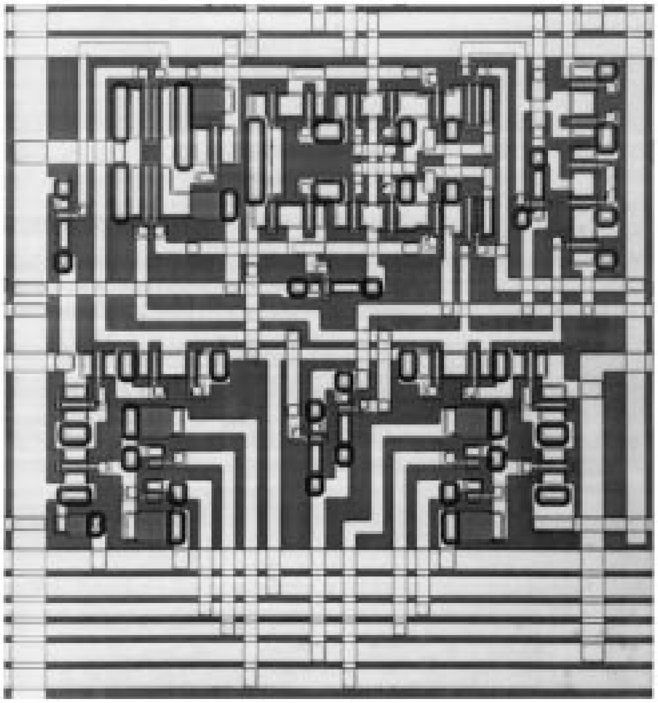

Fig. 3. An FET-SEED smart pixel. The area shown is approximately 200 μm in size.

During this period when much of digital optical computing was moving toward optical interconnects, there was continued analysis of optical interconnects. For example, engineering analysis showed specific energy dissipation benefits of optical interconnects [35]. Other physical analysis [36] showed that one reason for potential energy saving in optical interconnects was that optical devices, being fundamentally quantum devices, could act as efficient impedance transformers between the high impedance of small electronic devices and the low impedance of classical wave propagation. The first experiments on the use of optics for clock distribution were also performed [37].

Another important development during this time was the demonstration of viable vertical-cavity surface-emitting lasers (VCSELs). These had originally been pioneered in the group of Iga [38]. Jewell (in Huang's optical computing group at Bell Laboratories) and collaborators had been working to make resonant-cavity semiconductor optical switching devices. In an extension of this work, they were able to demonstrate the first electrically pumped VCSEL at room temperature [39]. There has since been much work in turning these into practical devices, especially for low-cost optical fiber connections, where they now dominate such systems. They are also candidate devices for optical interconnects to silicon chips, though they lag behind the quantum well modulators in demonstrations of dense interconnects. (See, e.g., [40] for a critical comparison of optical output devices for optical interconnects to silicon chips.) In dense interconnects, low power dissipation will likely be critical. An important development in low-power VCSELs was the demonstration of oxide-confined devices for potentially much lower thresholds [41].

With digital optics and optical interconnects converging, a logical next step was to proceed to interconnecting silicon electronics. The FET-SEED work did manage to demonstrate high-performance integrated systems, but it required a custom GaAs electronic technology. Silicon VLSI technology continued to advance, and the idea of being able to join with the mainstream technology was appealing. It also meant that the optics would then no longer be competing against silicon technology itself (a battle optics had lost many times with general-purpose silicon replacing special purpose analog optical processors), but rather against the metal interconnect technologies. With optics joined to silicon VLSI, as silicon transistors became better, so also would the optical connections.

The monolithic integration of optical emitters with silicon has been, and continues to be, a substantial basic challenge; III–V emitters monolithically integrated to silicon during growth typically have problems with lifetimes because of progressive propagation of the crystal defects normally inherent in such lattice-mismatched growth. Quantum well modulators were quite successfully grown on silicon substrates, however, with apparently good lifetimes [42] (a fact that is perhaps less well known than it ought to be). Despite this success, and subsequent success integrating with actual circuits [43], these workers soon realized that the problem of integration with silicon VLSI was not really solved in a practical sense. The issue is that in such a practical monolithic integration the VLSI process is changed in some ways. Additional high-temperature cleaning steps are needed. The circuits must be made on substrates that are cut off-axis. Considerable care must be taken with extra protection steps to prevent any of the silicon oxide in the circuit from being exposed to gallium metal because gallium makes silicon oxide conducting. Most important from a practical point of view is that the silicon production line will likely not accept the wafer back into the line, e.g., for deposition of final metal levels, if there is gallium on the wafer.

The practical solution to integration, at least in the short to medium term, is therefore some hybrid technique that can be performed on the finished silicon VLSI wafers. Various techniques have been proposed and demonstrated (see, e.g., [40]). Solder bonding has been particularly successful. Large arrays (e.g., thousands) of quantum well modulator/photodiode devices have been successfully bonded to functioning silicon complementary metal–oxide–semiconductor (CMOS) VLSI circuits [44], [45]. Fig. 4 shows one of the first of these bonded circuits as used in photonic switching experiments. More recently, smaller arrays of VCSELs have also been successfully bonded to silicon circuits [46].

During the mid to late 1990s, the scaling arguments for why optical interconnects would be interesting were further clarified. One analysis [47], for example, showed that a likely limit on future dense optical interconnects to silicon CMOS would be the dissipation of receiver circuits, but that, because of the continuing improvement of the silicon transistors themselves, the optical interconnects would likely be able to scale almost to match the growth of the ability of silicon circuits to perform logic. Also during this time period, it was becoming increasingly clear to the silicon semiconductor industry that electrical interconnects were beginning to run into serious scaling limitations. A basic scaling problem had been understood for some time, which is that, as an electrical line is scaled down in all three dimensions, its resistive–capacitive time constant does not change at all (see, e.g., [48] for early electrical discussions). This is, to say the least, undesirable, because it means that, since the transistors get faster as they are scaled down, the wires do not scale to keep up with the transistors.

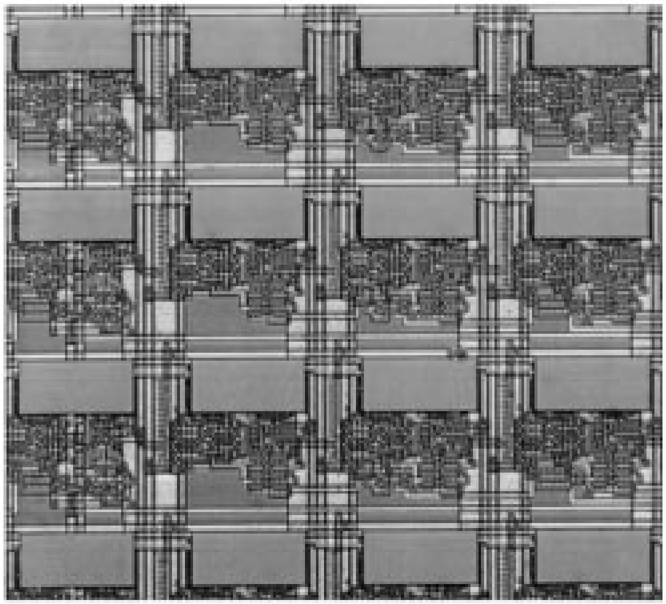

Fig. 4. Quantum-well diodes bonded onto active CMOS circuitry. The horizontal rectangles are the diodes, which can be used as reflection modulators or as photodetectors.

The analysis of the scaling of electrical connections was extended to include other kinds of electrical lines, with the relatively general conclusion that there is an underlying scaling to simple digital connections at any length scale [49]. This scaling concludes essentially that once an electrically connected architecture runs into speed limitations from the interconnect, and once all the available space is filled with wiring, the performance of the system cannot be improved either by miniaturizing the whole system or making larger; it is necessary to change the interconnect technology. This limit is known as the “architectural aspect ratio” limit [49]. Optical interconnects avoid this problem altogether because they do not have the resistive loss physics that gives rise to this behavior.

In the 1990s, the Semiconductor Industry Association had begun to publish “roadmaps” for performance and technologies for future generations of silicon technology [50]. These roadmaps show substantial problems for interconnects on silicon chips, with no known solution after approximately 2006 for interconnects to keep pace with the desired performance. Interconnects off of the chip, with their greater lengths, are likely a much larger problem. The issue of where optics might make sense as a solution to these problems has recently been analyzed in some detail [51], with the conclusion that optics is a very attractive solution on physical grounds to solve off-chip interconnect problems, and with some possibilities for use on the chips also. There is apparently no physical solution other than optics to the problem of dense off-chip interconnect.

The idea of changing the technology used for interconnects to silicon chips in computing and switching systems is a very radical one. It is important to understand exactly what the physical benefits are of doing that. The physical benefits have been codified in [52], and we will not repeat that classification here. Some of the resulting possible practical benefits of optical interconnects are as follows [51].

1) *Design Simplification:*
   a) Absence of electromagnetic wave phenomena (impedance matching, crosstalk, and inductance difficulties such as inductive voltage drops on pins and wires).
   b) Distance independence of performance of optical interconnects.
   c) Frequency independence of optical interconnects.
2) *Architectural Advantages:*
   a) Larger synchronous zones.
   b) Architectures with large numbers of long high-speed connections (i.e., avoiding the "architectural aspect ratio" scaling limit of electrical interconnects).
   c) Regular interconnections of large numbers of crossing "wires" (useful, for example, in some switching and signal processing architectures).
   d) Ability to have 2-D interconnects directly out of the area of the chip rather than from the edge.
   e) Avoidance of the necessity of an interconnect hierarchy (once in the form of light, the signal can be sent very long distances without changing its form or amplifying or reshaping it, even on physically thin connections).
3) *Timing:*
   a) Predictability of the timing of signals.
   b) Precision of the timing of the clock signal.
   c) Removal of timing skew in signals.
   d) Reduction in power and area for clock distribution.
4) *Other Physical Benefits:*
   a) Reduction of power dissipation in interconnects.
   b) Voltage isolation.
   c) High interconnect density, especially for longer off-chip interconnects.
   d) Possible noncontact, parallel testing of chip operation.
   e) Option of use of short optical pulses for synchronization and improved circuit performance.
   f) Possibility of wavelength division multiplexed interconnects without the use of any electrical multiplexing circuitry.

This short paper has tried to lay out a history of the development of ideas for optical interconnects to silicon chips. It is not a comprehensive review by any means of that field. One recent review is given in [52]. There has been much work in the 1990s on various demonstration systems that we have not reviewed here. The reader is also referred to the recent Special Issue of the PROCEEDINGS OF THE IEEE [53], which contains a much more comprehensive discussion of optical interconnects in digital systems, including optical and device approaches not discussed in detail here (e.g., waveguided systems). A review of proceedings from *Optical Computing and Optics in Computing* conferences over the last decade will show an increasing enthusiasm for optical interconnects, with this topic now arguably dominating the field.

With the idea of optical interconnects, the prognosis for the use of optics in digital computing and switching is more optimistic and realistic now than it likely has been at any point in the past. There are many challenges, including difficult issues such as packaging, cost, and reliability, and simply getting used to joining silicon circuits with optical links (e.g., receiver circuit issues). There is, however, a growing acceptance and even hope in the mainstream electronic community that optics may be successfully incorporated at some point in the not too distant future. That hope is doubtless shared by the optics community, which has shown an unflagging optimism that the days of optics as a major part of information processing would finally arrive.

ACKNOWLEDGMENT

The author is very pleased to acknowledge the collaboration and support of all of his colleagues at Bell Labs, where much of this work was performed, especially P. J. Anthony, W. Brinkman, D. S. Chemla, A. M. Glass, H. S. Hinton, A. Huang, H. Kogelnik, C. V. Shank, the extraordinarily talented groups of people that they led, and all of the people whom the author had the honor to lead.

REFERENCES

[1] R. W. Keyes, "Power dissipation in information processing," *Science*, vol. 168, pp. 796–801, 1970.

[2] H. M. Gibbs, *Optical Bistability: Controlling Light with Light*. Orlando, FL: Academic, 1985.

[3] H. M. Gibbs, S. L. McCall, T. N. C. Venkatesan, A. C. Gossard, A. Passner, and W. Wiegmann, "Optical bistability in semiconductors," *Appl. Phys. Lett.*, vol. 35, pp. 451–453, 1979.

[4] D. A. B. Miller, S. D. Smith, and A. Johnston, "Optical bistability and signal amplification in a semiconductor crystal: Application of new low-power nonlinear effects in InSb," *Appl. Phys. Lett.*, vol. 35, pp. 658–660, 1979.

[5] P. W. Smith, "On the physical limits of digital optical switching and logic elements," *Bell Syst. Tech. J.*, vol. 61, pp. 1975–1993, 1982.

[6] A. Miller, D. A. B. Miller, and S. D. Smith, "Dynamic nonlinear optical processes in semiconductors," *Adv. Phys.*, vol. 30, pp. 697–800, 1981.

[7] D. A. B. Miller, D. S. Chemla, D. J. Eilenberger, P. W. Smith, A. C. Gossard, and W. T. Tsang, "Large room-temperature optical nonlinearity in GaAs/$Ga_{1-x}Al_xAs$ multiple quantum well structures," *Appl. Phys. Lett.*, vol. 41, pp. 679–681, 1982.

[8] R. Dingle, W. Wiegmann, and C. H. Henry, "Quantized states of confined carriers in very thin $Al_xGa_{1-x}AS$–GaAs–$Al_xGa_{1-x}As$ heterostructures," *Phys. Rev. Lett.*, vol. 33, pp. 827–830, 1974.

[9] R. Dingle and C. H. Henry, "Quantum effects in heterostructure lasers," U.S. Patent 3 982 207.

[10] B. K. Jenkins, A. A. Sawchuk, T. C. Strand, R. Forchheimer, and B. H. Soffer, "Sequential optical logic implementation," *Appl. Opt.*, vol. 23, pp. 3455–3464, 1984.

[11] J. W. Goodman, F. I. Leonberger, S.-Y. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE*, vol. 72, pp. 850–866, 1984.

[12] P. M. Fauchet, "Progress toward nanoscale silicon light emitters," *IEEE J. Select. Topics Quantum Electron.*, vol. 4, 1998.

[13] L. T. Canham, "Silicon quantum wire array fabrication by electrochemical and chemical dissolution of wafers," *Appl. Phys. Lett.*, vol. 57, pp. 1046–1048, 1990.

[14] R. Bockstaele, J. Derluyn, C. Sys, S. Verstuyft, I. Moerman, P. Van-Daele, and R. Baets, "Realization of highly efficient 850 nm top emitting resonant cavity light emitting diodes," *Electron. Lett.*, vol. 35, pp. 1564–1565, 1999.

[15] D. A. B. Miller, D. S. Chemla, T. C. Damen, A. C. Gossard, W. Wiegmann, T. H. Wood, and C. A. Burrus, "Bandedge electro-absorption in quantum well structures: The quantum confined stark effect," *Phys. Rev. Lett.*, vol. 53, pp. 2173–2177, 1984.

[16] M. Aoki, M. Takahashi, M. Suzuki, H. Sano, K. Uomi, T. Kawano, and A. Takai, "High-extinction-ratio MQW electroabsorption-modulator integrated DFB laser fabricated by inplane bandgap energy control technique," *IEEE Photon. Technol. Lett.*, vol. 4, pp. 580–582, 1992.

[17] D. A. B. Miller, D. S. Chemla, T. C. Damen, T. H. Wood, C. A. Burrus, A. C. Gossard, and W. Wiegmann, "The quantum well self-electrooptic effect device: Optoelectronic bistability and oscillation, and self linearized modulation," *IEEE J. Quantum Electron*, vol. QE-21, pp. 1462–1476, 1985.

[18] D. A. B. Miller, "Quantum-well self-electro-optic effect devices," *Opt. Quantum Electron.*, vol. 22, pp. S61–S98, 1990.

[19] R. W. Keyes, "Optical logic—In the light of computer technology," , vol. 32, pp. 525–535, 1985.

[20] P. W. Smith, Y. Silberberg, and D. A. B. Miller, "Mode locking of semiconductor diode lasers using saturable excitonic nonlinearities," *J. Opt. Soc. Amer.*, vol. B2, pp. 1228–1236, 1985.

[21] U. Keller, D. A. B. Miller, G. D. Boyd, T. H. Chiu, J. F. Ferguson, and M. T. Asom, "Solid-state low-loss intracavity saturable absorber for Nd : YLF lasers: An antiresonant semiconductor Fabry–Perot saturable absorber," *Opt. Lett.*, vol. 17, pp. 505–507, 1992.

[22] U. Keller, K. J. Weingarten, F. X. Kartner, D. Kopf, B. Braun, I. D. Jung, R. Fluck, C. Honninger, N. Matuschek, and J. A. der Au, "Semiconductor saturable absorber mirrors (SESAM's) for femtosecond to nanosecond pulse generation in solid-state lasers," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 435–453, 1996.

[23] A. L. Lentine, H. S. Hinton, D. A. B. Miller, J. E. Henry, J. E. Cunningham, and L. M. F. Chirovsky, "Symmetric self-electro-optic effect device: Optical set-reset latch," *Appl. Phys. Lett.*, vol. 52, pp. 1419–1421, 1988.

[24] A. Huang, "Architectural considerations involved in the design of an optical digital-computer," *Proc. IEEE*, vol. 72, pp. 780–786, 1984.

[25] M. E. Prise, N. C. Craft, M. M. Downs, R. E. LaMarche, L. A. D'Asaro, L. M. F. Chirovsky, and M. J. Murdocca, "Optical digital processor using arrays of symmetrical self-electrooptic effect devices," *Appl. Optics*, vol. 30, pp. 2287–2296, 1991.

[26] N. Streibl, K. H. Brenner, A. Huang, J. Jahns, J. Jewell, A. W. Lohmann, D. A. B. Miller, M. Murdocca, M. E. Prise, and T. Sizer, "Digital optics," *Proc. IEEE*, vol. 77, pp. 1954–1969, 1989.

[27] H. S. Hinton, T. J. Cloonan, F. B. McCormick, A. L. Lentine, and F. A. P. Tooley, "Free-space digital optical systems," *Proc. IEEE*, vol. 82, pp. 1632–1649, 1994.

[28] F. B. McCormick, T. J. Cloonan, F. A. P. Tooley, A. L. Lentine, J. M. Sasian, J. L. Brubaker, R. L. Morrison, S. L. Walker, R. J. Crisci, R. A. Novotny, S. J. Hinterlong, H. S. Hinton, and E. Kerbis, "Six-stage digital free-space optical switching network using symmetrical self-electro-optic-effect devices," *Appl. Opt.*, vol. 32, pp. 5153–5171, 1993.

[29] R. G. A. Craig, B. S. Wherrett, A. C. Walker, F. A. P. Tooley, and S. D. Smith, "Optical cellular logic image-processor: Implementation and programming of a single channel digital optical circuit," *Appl. Opt.*, vol. 30, pp. 2297–2308, 1991.

[30] G. W. Taylor, P. A. Evaldsson, P. A. Kiely, T. Vang, P. R. Claisse, S. L. Daryanani, D. P. Docter, S. K. Sargood, and P. W. Cooke, "Integrated inversion channel optoelectronic devices and circuit elements for multifunctional array applications," *IEEE J. Quantum Electron.*, vol. 29, no. 2, pp. 785–800, 1993.

[31] K. Kasahara, "VSTEP-based smart pixels," *IEEE J. Quantum Electron.*, vol. 29, pp. 757–768, 1993.

[32] A. G. Kirk, H. Thienpont, A. Goulet, P. Heremans, G. Borghs, R. Vounckx, M. Kuijk, and I. Veretennicoff, "Demonstration of optoelectronic logic operations with differential pairs of optical thyristors," *IEEE Photon. Technol. Lett.*, vol. 8, no. 3, pp. 467–469, 1996.

[33] D. A. B. Miller, M. D. Feuer, T. Y. Chang, S. C. Shunk, J. E. Henry, D. J. Burrows, and D. S. Chemla, "Field-effect transistor self-electrooptic effect device: integrated photodiode, quantum well modulator and transistor," *IEEE Photon. Technol. Lett.*, vol. 1, pp. 61–64, 1989.

[34] F. B. McCormick, T. J. Cloonan, A. L. Lentine, J. M. Sasian, R. L. Morrison, M. G. Beckman, S. L. Walker, M. J. Wojcik, S. J. Hinterlong, and R. J. Crisci, "5-stage free-space optical switching network with field-effect transistor self-electro-optic-effect-device smart-pixel arrays," *Appl. Opt.*, vol. 33, pp. 1601–1618, 1994.

[35] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee, "Comparison between optical and electrical interconnects based on power and speed considerations," *Appl. Opt.*, vol. 27, pp. 1742–1751, 1988.

[36] D. A. B. Miller, "Optics for low-energy communication inside digital processors: Quantum detectors, sources, and modulators as efficient impedance converters," *Opt. Lett.*, vol. 14, pp. 146–148, 1989.

[37] P. J. Delfyett, D. H. Hartman, and S. Z. Ahmad, "Optical clock distribution using a mode-locked semiconductor-laser diode system," *J. Lightwave Technol.*, vol. 9, pp. 1646–1649, 1991.

[38] H. Soda, K. Iga, C. Kitahara, and Y. Suematsu, "GaInAsP–InP surface emitting injection-lasers," *Jpn. J. Appl. Phys.*, vol. 18, pp. 2329–2330, 1979.

[39] J. L. Jewell, A. Scherer, S. L. McCall, Y. H. Lee, S. Walker, J. P. Harbison, and L. T. Florez, "Low-threshold electrically pumped vertical-cavity surface-emitting microlasers," *Electron. Lett.*, vol. 25, pp. 1123–1124, 1989.

[40] D. A. B. Miller, "Dense two-dimensional integration of optoelectronics and electronics for interconnections," in *Heterogeneous Integration: Systems on a Chip*, A. Husain and M. Fallahi, Eds. Bellingham: SPIE, 1998, pp. 80–109.

[41] D. L. Huffaker, D. G. Deppe, K. Kumar, and T. J. Rogers, "Native-oxide defined ring contact for low-threshold vertical-cavity lasers," *Appl. Phys. Lett.*, vol. 65, pp. 97–99, 1994.

[42] K. W. Goossen, G. D. Boyd, J. E. Cunningham, W. Y. Jan, D. A. B. Miller, D. S. Chemla, and R. M. Lum, "GaAs–AlGaAs multiquantum well reflection modulators grown on GaAs and silicon substrates," *IEEE Photon. Tech. Lett.*, vol. 1, pp. 304–306, 1989.

[43] K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, and D. A. B. Miller, "Monolithic integration of GaAs/AlGaAs multiple quantum well modulators and silicon metal–oxide–semiconductor transistors," in *OSA Proc. Photonics Switching*, vol. 16, 1993, pp. 94–98.

[44] K. W. Goossen, J. A. Walker, L. A. D'Asaro, B. Tseng, R. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaS MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 360–362, 1995.

[45] A. L. Lentine, K. W. Goossen, J. A. Walker, L. M. F. Chirovsky, L. A. D'Asaro, S. P. Hui, B. J. Tseng, R. E. Leibenguth, J. E. Cunningham, W. Y. Jan, J. M. Kuo, D. W. Dahringer, D. P. Kossives, D. Bacon, G. Livescu, R. L. Morrison, R. A. Novotny, and D. B. Buchholz, "High-speed optoelectronic VLSI switching chip with >4000 optical I/O based on flip-chip bonding of MQW modulators and detectors to silicon CMOS," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 77–84, 1996.

[46] A. V. Krishnamoorthy, L. M. F. Chirovsky, W. S. Hobson, R. E. Leibenguth, S. P. Hui, C. J. Zydzik, K. W. Goossen, J. D. Wynn, B. J. Tseng, J. Lopata, J. A. Walker, J. E. Cunningham, and L. A. D'Asaro, "Vertical-cavity surface-emitting lasers flip-chip bonded to gigabit-per-second CMOS circuits," *IEEE Photon. Technol. Lett.*, vol. 11, pp. 128–130, 1999.

[47] A. V. Krishnamoorthy and D. A. B. Miller, "Scaling optoelectronic-VLSI circuits into the 21st century: A technology roadmap," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, no. 1, pp. 55–76, 1996.

[48] K. C. Saraswat and F. Mohammadi, "Effect of scaling of interconnections on the time delay of VLSI circuits," *IEEE Trans. Electron Devices*, vol. ED-29, pp. 645–650, 1982.

[49] D. A. B. Miller and H. M. Ozaktas, "Limit to the bit-rate capacity of electrical interconnects from the aspect ratio of the system architecture," *J. Parallel Distrib. Comput.*, vol. 41, p. 4252, 1997.

[50] "International technology roadmap for semiconductors," Semiconductor Industry Association, 1999.

[51] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, pp. 728–749, 2000.

[52] D. A. B. Miller, "Physical reasons for optical interconnection," *Int. J. Optoelectron.*, vol. 11, no. 3, pp. 155–168, 1997.

[53] Y. Li, E. Towe, and M. W. Haney, "Optical interconnections for digital systems," *Proc. IEEE*, vol. 88, no. 6, pp. 723–863, 2000.

David A. B. Miller (M'84–SM'89–F'95) received the B.Sc. degree from St. Andrews University and the Ph.D. degree from Heriot-Watt University in 1979.

He was with Bell Laboratories from 1981 to 1996, as a Department Head from 1987, latterly of the Advanced Photonics Research Department. He is currently the W. M. Keck Foundation Professor of Electrical Engineering at Stanford University, Stanford, CA, and Director of the Ginzton and Solid State and Photonics Laboratories. His research interests include quantum-well optoelectronic physics and devices, and fundamentals and applications of optics in information, sensing, switching and processing. He has published more than 200 scientific papers, and holds over 40 patents.

Dr. Miller has served as a Board member for both OSA and IEEE Lasers and Electro-Optics Society (LEOS), and in various other Society and conference committees. He was President of the IEEE Lasers and Electro-Optics Society in 1995. He was awarded the Adolph Lomb Medal and R. W. Wood Prize from the OSA, the International Prize in Optics from the International Commission for Optics, and an IEEE Third Millennium Medal. He is a Fellow of the Royal Society of London, OSA, and APS, and holds an honorary degree from the Vrije Universiteit Brussel.