Integrated Optoelectronics

by AMNON YARIV

A new information era, which will cause a fundamental change in the way we go about our daily lives, is about to dawn as the result of a major technological development now under way in the United States and Japan. We may even call it a revolution in communication.

Three separate technological breakthroughs, almost simultaneous but independent of each other, are responsible for this revolution. The first is the incredible advance that has taken place in the last 10 years in the very large scale integration of semiconductor circuits. This technology, which in 15 years has condensed the calculating power of a room full of computers into a box that fits in a pocket, makes it possible to store and process vast amounts of information at high speeds, using small and inexpensive hardware. Major advances now in the works, aimed at making basic transistors and circuit elements with submicron dimensions, will cause a further increase within the next 10 years of several orders of magnitude in the number of electronic functions that can be packed onto a semiconductor chip one centimeter square.

This vast and rapidly increasing capability for data handling would be of limited usefulness if it were not possible to transmit the data from one location to another at very high rates; otherwise the transmission link would form a bottleneck. It is also important for the cost of transmission to be low enough to be afforded by the average household. This brings us to the second technological breakthrough — optical fiber communication.

Up to now, the mainstay of modern man's communication has been transmission systems based on copper. Coaxial cable and twisted wire pairs still carry the vast majority of telephone conversations and communication data within cities. These installations, however, are the last dinosaurs of the communication field, and they are already on their way to extinction. The new and more agile predators causing this demise are inauspicious in appearance. They are tiny (''hair-thin'' as the ads say) silica glass fibers with diameters between 5 and 50 microns (1 centimeter equals 10^4 microns).

Glass fibers have been used since the 1950s for specialized applications involving short-distance transmission of light (less than a few meters). Originally they were not considered as serious contenders for long-distance optical communication because of the huge losses, of the order of magnitude of 3000 decibels per kilometer (db/km). A loss of that amount causes the intensity of light propagating in a fiber to be halved during each meter of path length.

What changed this situation was the realization in 1968 by researchers at the Standard Telecommunication Laboratories in Great Britain that these high losses were not intrinsic to the fibers but were due to impurities incorporated during their fabrication. Two years later scientists at the Corning Glass Works succeeded in pulling a fiber with losses of only 20 db/km. Further advances have resulted in fibers with losses down to the intrinsic level of about 0.1 db/km, due to Rayleigh scattering from the frozen-in thermodynamic inhomogeneities of the glass. At this level we can consider the prospect of propagation links with lengths of a few hundred kilometers between repeater stations, where the attenuated signals are reamplified.

An important property of these fibers is their extremely large bandwidth, that is, the ability to transmit very short (about 10^{-12} seconds) optical pulses; this makes them capable of carrying vast amounts of information. A single glass fiber a few microns in diameter can carry the entire FM and television channel spectrum with plenty of room to spare, or the equivalent of one hundred thousand telephone conversations. Fibers can transmit, conservatively, at speeds of billions of bits (gigabits) per second (bits are the measure of the number of pulses carried by the wave); this is a tremendous increase in data-transmission rate over other forms of communication. At these rates it will take approximately 1/100 of a second to transmit the information contents of an average book, or about four hours for the entire collection of a major library of one million books.

So, someday this "hair-thin" glass fiber entering an average home will be able to put at the disposal of its occupants huge amounts of information, which at present may be available only at special, often remote, sites. It will be much cheaper then to transport data than to transport ourselves. Our grandchildren will find it difficult to believe that back in 1981 we were so primitive that an engineer in Los Angeles had to board an airplane (and the price of fuel at this future time will probably reach $100 per gallon) and spend three days on a trip to New York just to discuss some blueprints for a new project with a client.

But the fibers only *transmit* light at gigabits per second; first the light has to be launched into the fibers. This brings us to the third partner in this technological troika — the semiconductor diode laser — a device less than a millimeter in size, which, when energized with a small voltage (about 2 volts) emits a coherent infrared beam that can carry the information "piggyback" as it propagates in the fiber. These lasers, invented simultaneously in 1962 by researchers at IBM and General Electric Research Laboratories, are still undergoing intensive development, and some of their properties are still topics of current research.

The principles of operation of these semiconductor lasers involve semiconductor physics, quantum electronics, and optics. With later modifications, the laser consists basically of a number of layers (transparent thin films a few thousand molecules deep) of gallium arsenide (GaAs) and gallium aluminum arsenide (GaAlAs) and takes advantage of the fact that the former has a smaller energy gap than the latter. When voltage is applied, electrons from the n-type layer of GaAlAs combine in the central GaAs layer with electron vacancies (holes) from the p-type layer of GaAlAs, producing a photon, which becomes part of the laser beam.

When the GaAs semiconductor laser appeared in 1962, I was working at Bell Laboratories, where it was *not* invented, so we were, of course, eager to see what our competition had done. As we were studying the physics of it, to try to discover the reason for its low light loss, we realized that dielectric waveguiding was going on at the p-n junction of the laser — that the layered structure trapped light as well as electrons. Light guiding, or waveguiding, confining and channeling a propagating beam of light (even a laser beam eventually spreads out), was first demonstrated in 1870. It involves the optical principle of total internal reflection, propagating a beam of light without any loss when an inner layer of the waveguide has a higher index of refraction (the degree of retardation with respect to the velocity of light in air) than the outer layer. In the original GaAs lasers, the waveguiding resulted fortuitously from the p and n doping of the different layers, causing a

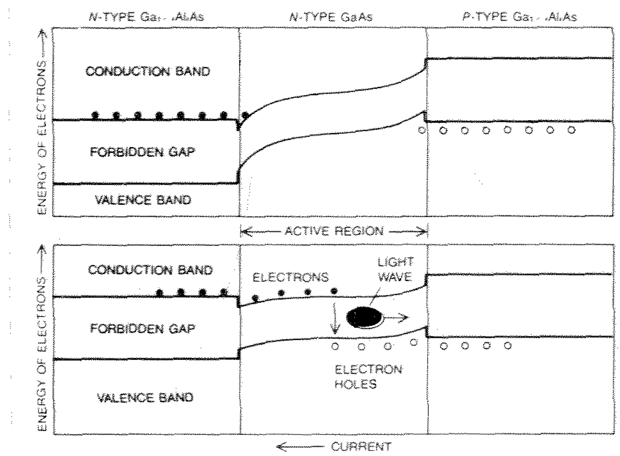

In the layered structure of the semiconductor diode laser, the electrons (black dots) of the n-type GaAlAs and the holes (open circles) of the p-type GaAlAs have the same energy but are separated by the central GaAs layer (in the top diagram). When voltage is applied (lower diagram), the energy of the electrons is raised, allowing both the electrons and the holes to flow into the central well where they can jump the energy gap and annihilate each other. The resulting photons are confined in the central trap, which also acts as a waveguide.

single light ray to zigzag back and forth across the inner layer by total reflection.

Even before coming to Caltech, I became convinced on theoretical grounds that it would also be possible to guide light at the surface of semiconductor crystals as well as at the p-n junction. This would make the light more accessible to launching and removal as well as to manipulation by surface techniques, and thus much more adaptable to practical use. I asked David Hall, one of the first graduate students to join my group at Caltech, to undertake this problem as his doctoral thesis project. He also attempted to use the electrooptic effect in the same crystal to make a light valve (modulator), in which an applied electric field turns the light on or off. At the time we did not have the facility for growing the crystals, so it was a rather frustrating experience for David, who had to depend on the generosity of a few industrial laboratories to supply us with gift crystals (a generosity that later stopped altogether as the donors became our competitors). Fortunately, David persevered and succeeded in 1969 in observing waveguiding and switching in a dielectric waveguide in a GaAs crystal.

Up to this point, research on the semiconductor lasers and waveguides had been proceeding without any particular direction — it was a nice thing to play with. But shortly after David's experiment, the first low-loss glass fibers were announced, and it became obvious that the era of optical communication via fibers was almost upon us. All of a sudden there was a need for a whole new technology to generate and manipulate the light that the fibers could transmit. You couldn't use conventional lenses, prisms, and the like because they are just too big and clumsy compared to the micron-size fibers (picture a watch assembly line where the workers are elephants). We needed a new approach to making optical systems that include miniature lasers, the modulators to impress the information on the light beam, the detectors to recover the information, and electronic amplifiers — a whole system to couple to the microscopic fibers and take advantage of their high data rates. Thinking about all this back in 1970, it suddenly dawned on me that the GaAs and GaAlAs crystals that we were already working with were a material system capable of performing all the functions needed to feed and retrieve information from the fibers. It's a very versatile material in both its electrical and its optical properties. GaAs/GaAlAs was already the material of choice for making lasers; field-effect transistor amplifiers in GaAs had just been demonstrated, and we had just shown how to make waveguides and modulators in these crystals. The variety of problems to be solved was large, but the technology definitely feasible.

So, almost exactly 10 years ago, in March 1971, we proposed that it should be possible to construct complete single-crystal monolithically integrated optical circuits, consisting of lasers, waveguides, modulators, detectors, and amplifiers, that would function as transmitters, receivers, and repeaters in fiber communication systems.

What is amazing is that we were then left alone for eight years to work on the problem by ourselves — despite frequent proselytizing on my part. The giant communication and computer companies in the United States were not convinced by our initial arguments and didn't pick up the idea (although the Japanese started watching us very closely). This afforded us the opportunity to pursue our goal at a fairly leisurely pace with the tolerant support of the Office of Naval Research and the National Science Foundation. Thanks in large part to the ingenuity of our graduate students and postdocs, Caltech succeeded where American industry, with its great resources, balked.

Before attempting the integration of all the functions in the "mother crystal," GaAs, we wanted to demonstrate that they could all be performed individually — as separate building blocks. A string of talented graduate students and postdoctoral fellows attacked many of the individual problems. One of our first tasks was to become self-sufficient in growing our own crystals. We do this with a liquid epitaxy (crystal-growing) system, starting with a semiconductor substrate, such as GaAs, and then bringing it in contact with successive wells of molten gallium saturated with the desired compound, for instance GaAlAs. When the temperature is lowered, a thin layer of that substance is grown on the crystal. Since GaAs and GaAlAs have the same crystal structure, the many-layered sandwich will act essentially as a single crystal. Impurities to create the p-type and n-type layers can also be added to the melt.

Many researchers contributed to the overall effort. Harold Stoll, PhD '74, and Sasson Somekh, PhD '74, demonstrated waveguiding in ion-implanted GaAs and developed the basic analytical and experimental tools to describe propagation in periodic waveguides. Shlomo Margalit, on leave as professor of electrical engineering at the Technion in Israel, added his enormous understanding of semiconductor device physics.

A great deal of our work has involved making lasers, and our lasers at Caltech are superior to those from any U.S. firm. Graduate students Huan-Wun Yen, PhD '76, and Willie Ng, PhD '79, and postdocs Michiharu Nakamura and Abraham Katzir, developed the distributed feedback laser, in which a fine surface corrugation, or grating, with a period of a fraction of a micron, replaces the end

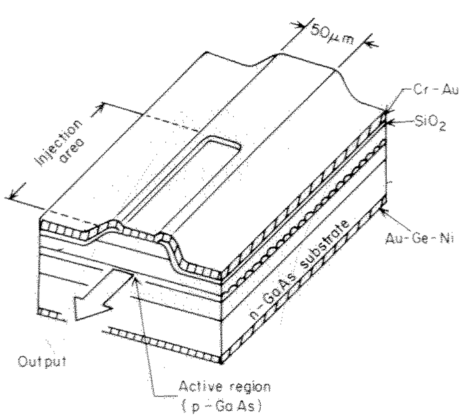

The distributed feedback laser uses a built-in corrugation for reflection (feedback) instead of the conventional end mirrors. The light generated in the active region by electron-hole combination is continuously Bragg-reflected back and forth by the corrugation, whose period determines the oscillation wavelength with great accuracy and stability.

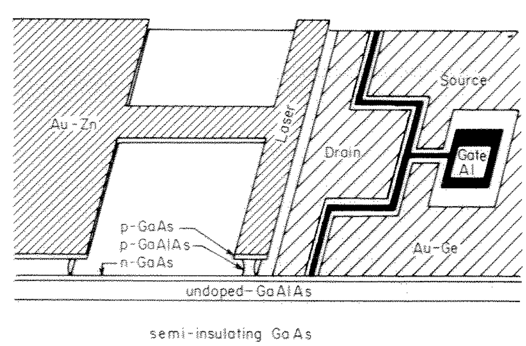

In this integration of the main electronic and optical elements, a field-effect transistor (FET), consisting of a source, gate, and drain, is grown monolithically along with a semiconductor laser on a single crystal of GaAs. The FET acts as a gate, which is capable of interrupting the laser current, turning it on and off at rates of a few times 10^9 per second.

mirrors in creating laser feedback by Bragg reflection of light (like x-rays in crystals). The gratings, produced by lithographic techniques, can also act as optical filters and multiplexers, since they discriminate among wavelengths. Along the way they also demonstrate field controlled coupling of light between adjacent waveguides.

At about that stage, in 1975, we felt ready to begin integrating a number of optical devices on a single-crystal chip. One of the first examples involved, naturally, bringing together the two main actors of optics and electronics — the laser and the transistor. Taking advantage of the basic similarity in the epitaxial layer structure of these two building blocks, Israel Ury, PhD '80, grew them together on a single crystal.

Another interesting example involved the monolithic integration of a microwave oscillator (Gunn oscillator) with a diode laser. The oscillator current, which flows through a crystalline channel directly into the laser, oscillated in our first experiment at a rate of 10^9 Hz, causing the light output to turn on and off at this rate.

The world's first completely integrated optoelectronic circuit was an optical repeater consisting of a light detector, electronic field-effect transistor, current amplifier, and semiconductor injection laser fabricated on a single chip of GaAs. The attenuated light pulses (as a result of long-distance propagation) enter the repeater from the incoming fiber at the detector. The current pulses are translated by the detector into voltage pulses, which are applied to the current field-effect amplifier. The output current pulses are fed into the laser and converted into large output light pulses that make their exit via the output fiber on their way to their ultimate destination. The technology is now available for long-distance communication over tiny optical fibers.

What happens next? At this point the biggies in industry have seen the light and are entering the game. At home, giant Ma Bell has finally picked up her skirts and waded in. In Japan, Hitachi already has a big effort, headed by our ex-postdoc Nakamura. The Japanese government has also established the Optoelectronic Collaboration Laboratory, which, as its name indicates, involves the cooperation of the large electrical companies there (Nippon Electric, Hitachi, Toshiba, Mitsubishi, and Fujitsu) with a charter to do research and basic development in optoelectronic integration. Further advances in this field will involve large expenditures of manpower and money — both more available in industry than in academia.

Our own effort has already changed direction. When we were trying to convince industry of the advantages of optoelectronic integration, we had in some cases to become more practical and technological than is our custom. Our research group is now moving back into a more basic pursuit of the material and electromagnetic aspects of the field. We are looking at the effects of the properties of the layer interfaces on the carrier lifetimes. Graduate students Tom Koch and Liew-Chuang Chiu are doing this with the aid of picosecond lasers. We are planning on growing synthetic crystals by molecular beam epitaxy, which makes it possible to control the thickness of the layers within a few angstroms. Chris Harder and Kam Lau are considering the ultimate frequency response of semiconductor lasers, while Dan Wilt has developed what is probably the first realistic numerical model of a semiconductor laser that allows for the interdependence of the optical and electronic properties. Hank Blauvelt and Joseph Katz are cooking up lasers with radically new layer structures and properties, and P.C. Chen is now working with a new "mother crystal" — indium phosphide — which holds a great deal of interest. This should keep us happy and busy for some time to come. □

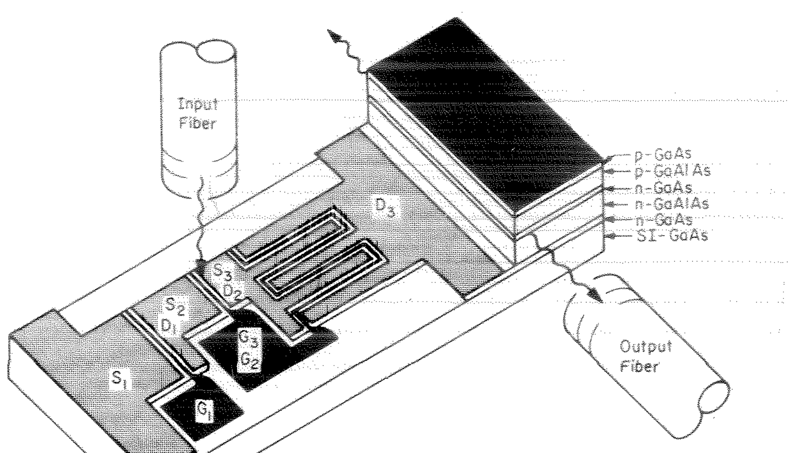

This complete monolithic optical repeater station regenerates the weak pulses from the input fiber. Current generated by FET 1 (source S_1, gate G_1, and drain D_1) flows through FET 2, whose source-to-drain resistance is modulated by the incoming light pulses. The resulting small voltage pulses are then fed to the gate of FET 3. The resulting large drain (D_3) current pulses enter the laser (raised mesa at right) through the common n-GaAs layer, causing it to turn on and off, feeding the regenerated optical pulses into the outgoing fiber.