

Electronics in Transition

The Tools and Technologies for 2000 AD

Moderator: Barrie Gilbert, Analog Devices Inc, Beaverton, OR, 97006-1994 (503) 690-1333
& (617) 937-1253 barrie.gilbert@analog.com

Abstract

This paper begins by putting electronics into a historical perspective, identifying it as the latest in a long line of knowledge technologies, noting that electronic writing (the sequential recording of original symbol strings) is now performed on magnetic surfaces, and electronic printing (the mass replication of fixed knowledge) is now done on plastic in the form of CD-ROM. We note a shift away from classical Newtonian electronics towards a VLSI culture that transcends the old paradigm; this 'new electronics' is here called *epitronics*, being built upon electronics, but no longer having the same essential physical nature. After a brief review of current technologies, some ideas about how we may expect these to evolve are presented, followed by a discussion of the new knowledge-based tools that will be needed for product development in the latter part of this decade. Finally, we try to imagine the impact of epitronics on Earthtown in the year 2000.

Barrie Gilbert

Barrie Gilbert (IEEE Fellow, 1994) was born in Bournemouth, England in 1937, where he earned an HNC in Applied Physics. He has been largely attracted to analog circuits throughout his career. At SRDE, working on speech-encoding systems, he used point contact transistors and pursued this interest in solid-state

devices at Mullard Ltd, applying Ge devices and later, working on the first silicon planar ICs. In 1964, he emigrated to the USA and developed ICs at Tektronix for the new 7000-series oscilloscopes. Out of this work came various new concepts, including merged logic (later I²L), for which he received the IEEE outstanding Achievement Award (1970); the translinear principle, a method for designing very accurate nonlinear circuits, which was cited in the IEEE Solid-State Circuits Award (1986) and the concept of the carrier-domain device. In 1970, he returned to England and was Group Leader at the Plessey Research Labs, working on optical holographic memories, optical character recognition systems and various communications ICs. In 1972, he joined Analog Devices, where he is a Division Fellow, managing the NW Labs in Beaverton, Oregon.

Barrie has published about 40 technical papers and has five times been awarded the ISSCC Outstanding Paper Award. Other awards include Oregon Researcher of the year (1990) and the Solid-State Circuits Award (IEEE-1992) for "Contributions to Nonlinear Signal Processing Circuits". He has 35 patents in the field of semiconductor devices and circuits. Current interests include SiGe HBT in communications applications.

Electronics is, and certainly will continue to be, a dominant force in shaping our world. Because of electronics, in particular, our field of microelectronics, we are witnessing immense changes in every aspect of our lives, and not with the analytical eye of some future historian, but as vitally affected participants in a drama of epic proportions. But it is precisely because we *are* participants that it's not always easy to comprehend the full meaning of these changes. Our work as technologists brings us face to face with plenty enough challenges for each day, leaving little time to dwell on the meaning of life, the universe and everything, at least, not in a professional capacity.

Nevertheless, we are bound to give serious thought to the essential themes of our time. Lacking a grasp of these, we may not be prepared for the changes of tomorrow. Unless we take the time to think about the 'big picture', we may find ourselves in the embarrassing or even dangerous position of realizing the full import of subtle trends and shifts too late—only when they have become established realities. As technology leaders, we need to seize the endless opportunities that surround us and vigorously participate in transforming them into the *new* reality, which is the essence of leadership.

Most of us attending ISCAS approach electronics primarily as a technical subject and secondarily as a domain of business. Our attention is almost entirely focused on present day technologies and market imperatives. In this paper, we pause to review the foundations of electronics and ask some philosophical questions about its role in an increasingly complex world. What is electronics 'about'? Is there just one 'electronics', or several? What will we be using electronics for, by the end of this century? Until questions like this have been answered, it is probably premature to address the more pressing question: What tools and technologies will be needed to equip us as effective players in the global electronics industry as we approach the milestone year of 2000?

We can safely assume that electronics will not be replaced by any other technologies within the next decade. Nevertheless, we also sense that it is at a major turning point. We can perceive the emergence of a new paradigm for electronics, in which the historical view—based on a 'nuts-and-bolts' interest in the individual physical components and their detailed behavior—is giving way to what might be called an epilogical one, that is, an awareness of the potency of large ensembles of components as manipulators of knowledge, accompanied by a declining interest in the detailed behavior of individual primitives.

Seeing electronics as the last manifestation of a long history of *knowledge technologies* is at the heart of the argument presented here. According to this view, we technologists must strive to understand the role of late twentieth century microelectronics in the *broadest possible terms*. That done, we will be in a better position to consider the tools, methods and technologies which will be needed at the turn of the century. In reaching our conclusions we will need to be vigilant in not allowing the incredible versatility of electronic technologies to cloud our judgment of what truly needs to be done, what we can safely afford to ignore. Future applications of electronics will be stunning, but they will be driven by quite different forces to those which shaped the present landscape.

RECENT HISTORY

Until the 1970's electronics was almost entirely concerned with physical components and phenomena. Today, more than ever before, it is the servant of knowledge. In understanding the changes that are occurring in our field, it will be helpful to review its brief history of less than one human lifetime.

Modern electronics began to evolve out of its recent foundations—telegraph, telephone and radio—slightly more than fifty years ago. The wartime invention of radar was significant in being one of the first uses of radio techniques that did not involve the transmission of a human message or symbols. Excitement about the future uses of wartime electronics in peaceful applications ran high. One fine example was the remarkable prediction by a science-fiction writer in his late twenties, from Minehead, England, of a global communications system based on the use of satellites high above the earth in space. His name was Arthur C. Clarke, and his prediction was reported in the October 1945 edition of *Wireless World*, a venerable journal, whose name for a while was beginning to sound a little old-fashioned, but which today has an almost prophetic ring.

Three years later, in an inauspicious corner of the Bell Telephone Laboratories, life was being breathed into the first fragile bipolar transistor. It could turn out to be just what was needed for use in submarine telegraph and telephone repeaters where the 20-year lifetime of the vacuum tubes posed expensive maintenance problems. And—who could say?—perhaps a handful of other applications could be found for it, too. Of course, it just might prove to be impractical, yet another research dead-end. In that same year, 1948, researchers at the University of Manchester in England were also making history (although they, too, were probably oblivious to the fact) by commissioning the Mark I, the first electronic stored-program computer, replete with warmly-glowing 'valves'.

Six years later, in 1954, the transistor was still a rare novelty, and still very much in its infancy: crude devices cost many dollars each, and speed and reliability were poor. No-one had any foolish notions about building a system that needed thousands—let alone millions—of them. The new generation of military speech scramblers, on which I worked in that year, did not risk using the unproven transistor; instead, miniature tubes were used as logic elements.

Sputnik! On October 4th, 1957, roughly in the middle of that International Geophysical Year, the Space Age officially began with the surprise launching of the first orbiting satellite by the USSR. This historic event came as a shock to the United States, and set alarm bells ringing nationwide; a presumed birthright to technological leadership had been challenged. But seen above the political noise floor, that day marked not only a technical triumph of the highest calibre: it was the first moment in human history when—at least symbolically if not yet to any practical extent—our minds and senses were released from countless thousands of years of existence bound to the earth's surface. A new chapter in the expansion of human knowledge had opened, and with it the early glimmerings of a new global awareness. But very few living then could possibly have foreseen what the *Sputnik* challenge would lead to today, less than forty years later.

Five years after *Sputnik*, in 1962, President John Kennedy announced a bold program to put men on the moon by the end of that decade. The US military already had a number of space programs in process, and the first of these came to fruition in 1960, when the Army Signal Corp. launched Courier, a simple low-orbit satellite weighing 500lbs. It demonstrated the feasibility of transmitting teletype at 55 kilobits per second on microwave carriers at between 1.7 and 1.9GHz. Needless to say, the transmitting and receiving systems were still based on vacuum-tubes, though of increasing sophistication and robustness.

Another three years forward, the launching of Syncom marked a new phase in the development of satellite communications. Fulfilling Clarke's prophetic vision, it was placed in a high geosynchronous orbit, one consequence of which was that it was rarely in the Earth's shadow, allowing it to be powered by solar cells, thus minimizing the weight of its storage batteries. During the Japan Olympic Games in 1964, live television pictures were conveyed across the Pacific, by Syncom III, marking another milestone in global awareness. 1965 saw the launch of Intelsat I, and no less than three Intelsat satellites were launched in 1967, four more in 1969. The communication capacity of these early satellites, serving 75 countries by 1970, was immensely greater than that of Courier. With these satellites, the age of instant transglobal voices, images and text had dawned, announcing another step toward the global village and one of the most significant milestones in the development of our central theme—in which electronics has figured so dramatically—of the quest for knowledge.

1969 was important for another reason, which is curiously easy to forget. On July 20th of that year, we watched an evening TV show featuring three guys and a spaceship; their names were Neil Armstrong, Ed Aldrin and Mike Collins. Their successful journey to the moon, in fulfillment of Kennedy's mandate, should have been the event that changed the world. But it wasn't. Something of even more far-reaching importance was happening down here on planet Earth.

In the midst of this Cold War era, numerous integrated-circuits companies, with names like Motorola, Fairchild and National, were getting into their stride. While complex and challenging functions—including A/D and D/A converters of useful accuracy and speed—were still board-level entities, such things as monolithic amplifiers, small-scale arrays of logic gates and a few communications ICs were now routinely being designed into the latest equipment. These components represented the latest embodiments of the classical, and now maturing, world of what may be called Newtonian electronics, in which physical devices and mechanisms remain central to the subject.

Meanwhile, a distinctly new kind of electronic object—the *microprocessor*—was quietly incubating. Bringing together most of the pieces of a very small computer into one piece of silicon (although without significant amounts of memory to begin with) the microprocessor was poised to revolutionize the world, quite literally. It was clearly electronic in substance, but in an important sense that was merely incidental: its true nature would soon transcend its material surface.

During the intervening twenty five years we have since seen the full maturation of IC technologies, with staggering advances in sophistication of fabrication equipment and silicon processing. With these, computers of unprecedented power have been enabled, and a world-wide communication network based on satellites has emerged. These two major developments are fueled by the insatiable demands for knowledge in government, in the business arena, in manufacturing and construction industries, in education, in health care, and in research. The recent shift from a military to a civilian emphasis in electronics is rapidly generating new challenges and creating new enterprises.

The world of electronics has matured in many diverse ways. But towering above all else have been the spectacular advances in the speed and complexity of the microprocessor and its associated VLSI technologies, driven primarily by logic and memory development. These have resulted in unimagined growth in the scale of integration and huge reductions in the cost of a function, leading to an explosion in their scope of applications. We now sense that something very important is happening out there—something that transcends progress in terms of the traditional performance metrics of smaller, cheaper, faster and more accurate, however impressive these continue to be. As a direct result of the advent of VLSI, and the communications infrastructure that it has enabled, many ideas about charter, once held by IC manufacturers, have needed to be revisited.

These had included, for companies like my own, the forswearing of all interest in purely digital circuits, the use of silicon foundries, a distaste for low-margin components, and little expectation of ever being embroiled in microwave communications circuits. Almost completely counter to these entrenched notions, today's mixed-signal theme involves a high proportional of digital circuitry, and to an increasing extent designed with cost very much in mind, forcing a much stronger emphasis than ever before on high quality manufacturing. In particular, we have learned how to make silicon objects containing many millions of transistors with incredible yields; the most advanced examples of DSP contain upwards of 30 million devices, and 256Mb memories are becoming commercially available. The billion-transistor circuit is now within reach, and will probably be working before the close of the decade.

Clearly, modern VLSI is profoundly more than very complex collections of transistors. Such objects can scarcely be called 'circuits', in the sense that an amplifier or an LC-oscillator is an electronic circuit. Indeed, very large digital ICs can quite easily be designed, at the architectural level, without any knowledge of 'circuits'. Something very profound is going on here. Whatever electronics may have been fifty years ago, we sense that a large part of it today is concerned with a different theme.

When we step back and look at the complete world of electronics in a somewhat detached manner, we conclude that the overarching theme is that of human knowledge. That is, **electronics is the latest manifestation of a centuries-old evolution in making knowledge accessible to more and more people, at ever lower cost, in exponentially larger amounts and at ever-increasing speed.**

ANCIENT HISTORY

Surely, some may insist, electronics continues to be about electromagnetic propagation, about signals, filtering and automated decision making, about measurement and control? Certainly, that is what electronics *was* about, until some fairly recent time. Today, a very significant transition is occurring. Modern electronics technologies are being increasingly enrolled in meeting widespread contemporary demands for access to knowledge, provided by stone, clay and in the past.

In the far reaches of history, human knowledge must have been unique to each individual. Accumulated observations and deductions were a major asset in augmenting his or her physical strength and thus the likelihood of survival. Eventually, natural languages blossomed. These achieved something rather remarkable: they enabled the imparting of knowledge from one individual to another, and from one generation to a later one, for example, by story telling.

This transference of knowledge was by a *flow of information*, which is the dynamic form of knowledge. We might say that knowledge is analogous to potential energy and information is analogous to kinetic energy. Alternatively, knowledge is like voltage and information is like current. Groups, such as families or clans, that shared a *common pool* of knowledge could have enjoyed many advantages over others group that were without it. The same holds true today, except that nowadays these groups also extend to bigger entities such as companies, communities and even nation states. But language became more than a merely useful asset that allowed the transmission of survival knowledge and tool-making skills. It became in itself a new kind of tool—and a powerful one—for describing, modeling and ultimately understanding our world.

We generally think of writing as the 'next' development, but it's likely that linguistic and iconographic skills developed concurrently. Long before the emergence of a formal writing system, mankind was using various forms of notation, often in a record-keeping context. A small engraved bone, called the Ishango Bone, was found in Africa roughly thirty years ago. Dating methods suggest it goes back to about 8,500 BC, which puts it near the beginning of the Mesolithic period. It was described in a 1962 article in *Scientific American* and at much greater length by Alexander Marshak in his provocative book *The Roots of Civilization*. This bone, only some 10cm long, is marked by a series of scratches, arranged in three distinct columns, which appear to be more than mere decoration. Marshak's theory is that they might have been an early calendar, and he supports this idea in some interesting and persuasive ways.

The importance of a *permanent notation* for knowledge is that it allows knowledge to be more accurately transported in both time and space—unlike spoken language, which is temporal and evanescent. Bones, clay tablets and other surfaces inscribed or otherwise marked in special ways could have been used years later by its maker's progeny, or at some distant location where its valuable knowledge would probably still apply.

Our attention now leaps forward some five thousand years, to the land of Sumeria, which later became known as Babylonia. It occupied an area a little larger than Wales or Massachusetts—about 9,000 square miles—in what is now Iraq. This ancient land was extremely flat, hot and dry. It had no trees for timber, and almost no minerals or stone. But the people who occupied it during the third millennium BC were endowed with an exceptionally creative intellect and enterprise. In spite of the land's natural limitations, they turned Sumer into a veritable Garden of Eden and developed the first high civilization in the history of man.

Their list of achievements is awesome. The Sumerians invented land surveying, map-making and irrigation, the potter's wheel, the wagon wheel, the plow, the sailboat, the arch, the vault, the dome, casting in copper and bronze, riveting, brazing and soldering, engraving and the use of inlay, sculpturing in stone, and musical instruments of all sorts. We know this not only because of the thousands of art objects, everyday utensils and tools unearthed by archeologists, but also because they actually tell us about their achievements in texts and pictures. These were written as marks on a soft clay that quickly dried in the hot sun to a durable form. So plentiful are these cuneiform texts that, in some recent archeological digs, the majority were simply thrown away.

These early writings are not brief or primitive. Some are long epic poems that consume several pages of a modern book in translation. Many other are accounts of business transactions, and others appeared to be used for teaching purposes in schools (another Sumerian invention). What is most notable about these writings is the high level of civilization that they portrayed. For example, a tablet dated about 2,200 B.C. contains 15 medical prescriptions. These are amazingly rational, and probably effective. For example, to deal with a bad cut or abrasion, one of these instructs:

"Sift and knead together turtle shell, the sprouts of the naga plant, salt and mustard. Wash the sick area with quality beer and hot water, then apply the mixture. After scrubbing, apply vegetable oil".

The potential efficacy of this treatment relies on the alcohol in the beer and the salt helping to disinfect the area and the oil later serving as a bacteriological barrier. What is especially noteworthy about these prescriptions is that they are entirely free from the magic spells and incantations which were a regular feature of medieval medical texts: not a single deity or demon is ever mentioned.

Stop! What has all this to do with electronics? **It is the power of writing.** Written languages marked a huge step forward because they freed people from the limitations of human memory, and from the spatial and temporal constraints of the spoken word. The first lawyers seized on the potential of writing as the basis for binding contracts (as we know from numerous examples from Sumeria) and its uses rapidly grew into every corner of society. In an astonishingly similar way, its modern counterpart, that is, the *sequential writing of symbols* on magnetic surfaces, is once again revolutionizing the way we live. Twentieth-century knowledge technologies clearly have their roots in writing.

The invention of paper, generally attributed to the Chinese, about one thousand years ago, was a milestone of immense importance in the development of this great theme. During the 14th century the growth of knowledge led to a demand for writing paper that resulted in a four-fold drop its cost. Between 1444 and 1448 Johannes Gutenberg invented the movable-type printing press. Many details of his visionary invention remained essentially unchanged until the 20th century. Alas, like a lot of entrepreneurs, Gutenberg only got into heavy debt feeding his brainchild, and died, blind and broke, some twenty years later. But his invention irreversibly changed our world. The printed book ensured a high degree of accuracy in the conveyance of knowledge, much better than when documents were hand-copied by scribes. And that process had been, in turn, more reliable than transmission using the spoken word and story-telling. The proliferation of affordable books in the following 500 years had an immense impact on the rate at which human knowledge, particularly in the sciences, was propagated, and this in turn, on the rate at which new knowledge was generated. Eventually, with this new enlightenment came a rejection of institutionalized authority and an increasing reliance on the power of the *individual mind*, nowadays aided even more powerfully by the networked personal computer.

The relation of computers to knowledge in the earliest days was tenuous. They were designed expressly for numerical applications, as their name implies. Their use in processing symbols was quite slow in emerging. This was largely because the amount of written knowledge in the mid-'50s was still capable of being managed using paper media. Indeed, a single page of typewritten text, using about 30,000 bits, would have consumed all of the memory capacity of even quite advanced computers of the time.

The development of much larger memories allowed the computer to enter the knowledge mainstream. Today's mass storage media have reached the point where we can store several hundred megabytes—probably more than all of the knowledge of the entire Sumerian world—on one small disc drive. More to the point, we can store a significant fraction of the modern world's knowledge in such a way that essential pieces are available to us within milliseconds.

Recently, the compact disc—a triumph of advanced optical, mechanical and electronic technologies, and originally intended only for audio, but now adopted as a cheap medium for read-only digital data—has at last supplied a viable alternative to the book, a half-millennium after Gutenberg. Like the book, the CD-ROM is not written sequentially, but it is *printed*, this time on plastic, and in pages of almost a gigabit. Like the book, it affords the same benefits in disseminating large amounts of knowledge at low cost, and has the same potential for revolutionizing our world. But only computers can read the cryptic symbols of *these* books, translating them into familiar text just for our benefit. Seen in this historical context, one can appreciate that modern microelectronics are serving an age-old need for unhindered access to accurate knowledge, which is often directly equatable to power, and ultimately, in no less of a sense than in the times of earliest mankind, to individual survival.

THE PROSPECT

We stand at an astonishing milestone in human progress, the latest in a 10,000-year evolution in the development of knowledge technologies. These technologies allow us to store our own knowledge in a safe place for later usage, like the Ishango Bone or the Sumerian cuneiform tablet, or acquire someone else's knowledge, a service which was once the exclusive province of libraries, in both cases using machines that are crucial participants with us in our knowledge-related actions, and which can already perform better than we in such prosaic but intellectually relevant matters as spelling and even grammar.

Having an abundance of knowledge sources at our personal disposal allows us to be more complete citizens of Earthtown, and to create the future, each in our individual way. However, there's something else very important happening. The storehouses of knowledge in the past required intelligent human intervention, to search for and retrieve the relevant records and then bring them to life in achieving some action or effect. In profound contrast, the knowledge residing in computer memories is *machine-executable*. Modern stored knowledge can autonomously animate the machine in which it resides. Furthermore, machine executable knowledge is fluid and dynamic. It can be replicated and modified any number of times inside its own environment. Knowledge objects that have lost relevance can be quickly discarded and replaced with more potent objects. (We humans fair *far* worse in disposing of old ideas). With the addition of a communications channel, a machine can reach out to some data super-highway and pull in new knowledge objects, usually from fellow machines, which may result in altered behavior. It's only a matter of time before machines will get better than we are at knowing where the top-quality knowledge objects are to be found and go after them on our behalf, without our intervention.

It would be entertaining to explore the likely impact of these current developments on future thinking machines. Certainly there is little doubt that modern computers are capable of exhibiting intelligence, if that word is defined simply as meaning the ability to execute an intellectual act which, prior to the development of that computer, would have required an above-average human. This might be anything from calculating a square-root to diagnosing a pathological condition given the symptoms, or learning how to spray paint a chair without ever having seen one before, or playing a world-class chess game.

The next step in machine intelligence will not occur until we allow machines to deviate from deterministic behavior. The old view of the computer as merely a tool in the hands of an infinitely superior owner must give way to the realization that electronic 'minds', with their prodigious abilities of recollection and rapid thinking, must be allowed a measure of free will. Lacking that, they will forever be merely capable and willing, but exceedingly dull and unresourceful, servants. With it, their potential for creative contribution to our lives can perhaps be realized. Some may fear that prospect, like many once feared 'being taken over by machines' (a recurrent theme throughout history). Beyond intelligence is understanding, perhaps followed by machine consciousness.

What this really means, and what effect its eventual emergence will have on human life, we can only guess. Might it be possible, for example, that machines learn the value of knowledge objects and develop a new trading system based on them, rather than on intrinsically valueless metal and paper tokens called money? Whatever lies beyond the year 2000, we can be certain that knowledge technologies will be at the centre of all human progress. Thus, we can make the equation

KNOWLEDGE TECHNOLOGIES at the millennial transition

$$\begin{aligned}
 &= \text{ELECTRONICS (probably little else beside silicon)} \\
 &= \text{COMPUTERS} + \text{COMMUNICATIONS} \\
 &\quad \begin{array}{ll}
 \text{which accumulate} & \text{which move knowledge} \\
 \text{and manipulate} & \text{objects from one place to} \\
 \text{knowledge objects} & \text{another = information}
 \end{array}
 \end{aligned}$$

That is, there are only two main thrusts to the central theme: (1) computers, which, in many modern applications may actually spend very little of their CPU time in computing, if that means executing numerical calculations, and far more time in fetching and carrying knowledge objects, including pictures and sound; (2) communication systems, which to an increasing extent serve the not only human needs but also the conversational needs of increasingly smart computers.

Looking at the late-twentieth-century planet, we can represent it as a sphere existing in the infinite domain of the real world. At its centre is a solid, dense core of totally digital technologies, the domain of memories (semiconductor, magnetomechanical and optomechanical) and microprocessors; surrounding this is a thick layer of digital signal processing, reaching out toward the surface layers; in a thinner layer we find the A/D and D/A converters which form a boundary between the inner world of symbol-based knowledge- representation, manipulation and storage, and the outer domain of analog signals. The entire domain of analog (Newtonian) circuits exists within the very thin biosphere which acts like a permeable membrane between the deeper layers of knowledge objects and the 'real' world of signals.

This representation of the world of electronics is very different from the *Sputnik* era of thirty-seven years ago. We do not know for sure how it will look in another thirty—or even ten—years from now. There are compelling reasons to believe that the analog boundary may push back toward the core as the growing expectations placed on thinking machines leads to a greater dependence on quasi-analog neural structures. It seems unlikely that the next wave of advances in non-human intelligence can be made without the use of richly-connectivist architectures and the abandonment of strong-physical-symbol representations of knowledge, but that is very much a philosophical question. As things look today, with the huge emphasis on Boolean systems, it's hard to imagine that a totally new approach to the machine manipulation of knowledge, based on continuous-amplitude, or at least, multi-level, representation of data, will ever displace them. But stranger things have happened.

A NEW PARADIGM FOR ELECTRONICS

The philosopher Thomas Kuhn first brought to our attention the idea of the paradigm shift in science, that is, the occasional periods in which major changes in viewpoint arise and displace formerly quite adequate, though sometimes strained, ways of modeling our world. We are in the middle of such a shift, from a traditional paradigm for electronics to a more modern view.

The old paradigm may be called Newtonian or Maxwellian, since it is based on physical effects and devices and the utilization of these in the generation, processing, propagation and reception of continuous-time, continuous-amplitude electrical and electromagnetic signals, also having physical attributes.

The new paradigm may be called epilogical or perhaps epitronic, since it concerns the non-physical operations of acquisition, storage, manipulation and presentation of knowledge objects existing as very large ensembles of logical (usually binary) variables, conveniently represented in electrical form.

Put another way, traditional electronic systems were conceived, designed and engineered largely in terms of the behavior of their component parts, which it was discovered could also usefully be applied to information-processing systems, although these are not *essentially* electrical. We are inclined to think of these systems as actually being electronic *in nature*, whereas they merely utilize electronic elements as a matter of convenience. They are *epitronic* (Gr. 'above').

We are already seeing epitronic systems that are conceived, designed and engineered largely in terms of their logical structure and function. In many cases, these can be generated using automatic compilers, and the fact that they are eventually fabricated in silicon is not particularly important. In such cases, no-one knows the precise details of the final electrical entity that results; even if known, such details would be almost worthless. The value of epitronics is that it implements algorithms; it is the silicon realization of purely logical functions. Electronics has been subjugated to a *secondary, supporting* role.

The change that we need to understand, then, is that nowadays there are at least 'two electronics': the one operating under the old paradigm, in which physics prevails and signals have meaningful dimensions, such as voltage, current and (most importantly) power; and the new *epitronics*, in which only dimensionless logical variables exist, and in which the physical substrate supporting the processing of these variables is of only incidental significance.

Obviously, digital VLSI has been so phenomenally successful in serving the logical appetite of epitronics because silicon is cheap and fast. A somewhat more subtle reason for its success lies in the fact that the logical variables also *happen to be* voltages, allowing a seamless interface with the world of Newtonian electronics—signals from sensors, and to actuators, that work in the physical domain—through the use of A/D and D/A converters. It is therefore not surprising that converter technology has received so much attention throughout the past thirty years of IC development: data converters are the gateways between the 'old' and the 'new' electronics, between our real, Newtonian world of physical phenomena and the rapidly evolving domain of epitronics.

Being aware of the idea of there really being 'two electronics' as we approach the close the twentieth century is probably not going to radically change the way we go about our daily work as microelectronics engineers. But it may help us to appreciate that while both disciplines will continue to expand their horizons, the growing gap between Newtonian electronics and epitronics will eventually become a chasm, and at some point in the future there will be little in common between designers practicing in these two areas. How will that affect the development of technologies? To what extent will these technologies become divergent? What skills will be required of the next generation of designers? What tools will be needed to meet the differing demands of the two electronics?

As already noted, traditional (largely analog) electronics is the domain of EM propagation and power management. Its circuit elements are oscillators, fixed and variable-gain amplifiers, filters, modulators, mixers, voltage references and the like; its components include those which address the intrinsic dependence on energy storage devices (capacitors, inductors, mechanical resonators), and those which provide current, voltage and power gain, provided during the last three decades primarily by the bipolar junction transistor; its signals have physical dimensions (Volts, Amps, Coulombs, Watts, Seconds, etc.) and so do its components (Ohms, Farads, Henrys, etc.) and thus classical electronic sub-assemblies have many parametric sensitivities, to manufacturing tolerances, to supply voltages, to mechanical stress, to aging effects, and so on. Most notably, their behavior is inherently and fundamentally dependent on the temperature of their component parts, and much of the challenge in circuit design is coping with this sensitivity. Finally, Newtonian electronics is bounded by the all pervasive noise that permeates real circuits, traceable ultimately to the finite charge of the electron and the thermal energy kT/q . A further challenge in classical design is finding ways to cope with circuit noise, rarely allowing a dynamic range of greater than 80–100dB (dependent, of course, on the signal bandwidth).

If we follow a similar thread of thought for epitronics—which is today the world of digital electronics, invariably meaning *binary* electronics—we already see marked differences. Instead of being concerned with EM propagation and signal power, epitronics is about what may loosely be called 'data'. Its elements are gates. Whether they are voltage- or current- or power-driven is not fundamentally important: gates are merely switches. However, voltage-driven CMOS transistors now almost exclusively provide these switches. Current flows only during transitions between states, and its value is not important to the function of the gate; thus, power is needed to operate the gates, but beyond that has no significance. Its 'signals' (stretching an 'old' electronics term) are logical variables, having zero dimension. In this important sense, logical signals lie outside of the domain of physical effects. The extreme simplicity of the switch means that it is essentially free from of parametric sensitivity. Finally, noise does not exert quite the same dominion over epitronics as it does in the analog world; there is, for example, no limit to the 'dynamic range' of a binary number, other than the number of registers and wires needed to represent and convey it. A 32-bit word represents a physically-unrealizable dynamic range of 192dB.

We can make a further observation, in attempting to see these distinctions more clearly. Today's electronic knowledge technologies are generally divisible into 'Communications' and 'Computers'; the former group, involving real signals, is largely served by Newtonian techniques; the latter, involving logical variables, by epitronics.

As things are today, we still think of these two very different worlds as one, using the same basic microelectronics technologies but with differing emphases. Because of the high cost of development, we try to use each new technology for both worlds. Further, because VLSI has been the driving force in the modern market, we try to use digital technologies—predominantly CMOS—for analog applications. The many successes that have been achieved in this connection are testimony to the ingenuity of young designers. Even the distinction between bipolar and MOS technologies is blurred by the fact that both use silicon (almost exclusively); indeed, the two types of devices are often merged into a BiCMOS process. This sharing of technologies tends to reinforce the idea that electronics is just one discipline. The outlook is further complicated by the fact that a very large part of any communications system is nowadays digital, and a significant part of a computer system remains analog. And, into what category are we to place the all-important data converters?

It will probably be a few decades before the divergence between the two main disciplines will lead to totally different technologies. But we can already see signs of this. Thus, it is possible to be trained in the computer sciences and thereafter pursue a career in computer design without knowing the first thing about transistors. At one time, this might have been regarded as a dangerously myopic view of one's subject matter; today, we recognize that life is simply too short to embrace every aspect of a broadening field. And really, this is little different to the circuit designer being ignorant of the chemistry or metallurgy of microelectronics. We also are beginning to see digital-only technologies—meaning, new sorts of picopower switches—being explored at the research level, even down to what are misleadingly called quantum devices.

As things stand today, it is premature to make any bold pronouncements about the eventual collapse of the unity of the field of electronics. We will continue to share technologies, although it is very likely that digital electronics will continue to dominate the technology front, and that analog designers will be forced to find ever cleverer ways to make up for the deficiencies of technologies optimized for making fast switches. Thus, another apparent trend is the increasing use of digital filtering and other digital concepts to shore up the analog sections of a system, such as in sigma-delta converters and signal filtering in receivers.

The next section reviews the primary technologies available today, and attempts to see some trends, after which we will consider the tools that might be needed to carry out design work in the two electronics.

TECHNOLOGY INVENTORY

Today's designers are faced with a tantalizing array of IC technologies, viewed in totality. However, in any one environment, the choices are usually much more limited. The huge infrastructure which has grown to support digital VLSI—silicon processing knowledge and equipment, and design and layout tools—is to an increasing extent forcing designers of Newtonian circuits to utilize processes, predominantly CMOS, which were optimized for digital applications. Out of these constraints have come many clever circuit ideas, in which the wide parametric variations of the individual transistors are prevented from significantly influencing the quality of the analog section of mixed-signal circuits. Often this is achieved through the use of digital signal processing supporting a 'bare-bones' analog section, as in sigma-delta converters.

The great successes achieved in using low-cost digital processes in this way has fueled the growing conviction that "everything will eventually be done in CMOS". One of the most controversial topics of the day is whether that will turn out being correct, or whether bipolar transistors (perhaps heterojunction devices) and GaAs MESFETs hold the key to levels of performance that simply cannot be attained by CMOS at gate lengths somewhere near the realistic lower limit of $0.25\mu\text{m}$. However, even the most avid proponent of "CMOS for everything" will usually admit that it is primarily its low cost that commends the all-CMOS approach, not any intrinsic advantages of the MOS transistor as a Newtonian electronic element (although it has some).

Contemporary IC technologies can be grouped in a roughly historical order within each type. First, there are the all-bipolar processes:

- Traditional junction-isolated (JI) bipolar with a 'good' NPN transistor, low-speed lateral and substrate PNP's, and a limited availability of high-quality passive components, used in ECL and similar forms of 'current-mode' logic, and in many analog applications of modest complexity, including communications functions up to about 300MHz.
- High-speed bipolar JI+trench bipolar, featuring an NPN with f_T 's up to 25GHz, and capable of supporting several-thousand-transistor devices, such as flash A/D converters and frequency-synthesizers for use up to 2.5GHz, as well as analog mixers, oscillators, RF and IF amplifiers, etc. The penalty for bipolar speed is a low operating voltage, of $\sim 6\text{V}$.
- Ultra-high-speed SiGe heterojunction bipolar (HBT) processes, still quite rare, but offering extremely low inertia, with low-femtofarad junction capacitances and f_T 's to 110GHz. Will be useful in achieving high speed at much lower power consumption ($V_P=3\text{V}$), in the same general applications areas as all previous bipolar processes.
- First-generation JI complementary bipolar (CB) processes, providing both polarities (NPN & PNP) of transistor, with similar f_T 's in the 0.5–1GHz range, and often supported by thin-film (and laser-trimmable) resistors. CB has proven valuable in high-symmetry high-swing ($\pm 10\text{V}$) video amplifiers and some nonlinear analog applications.
- Second-generation CB processes, the most advanced of which use bonded-wafers and trenches to achieve total dielectric isolation (DI), very low capacitances and f_T 's in the 3–5GHz range. These are true analog processes, providing thin-film resistors and metal-oxide-metal capacitors with high 'top/parasitic' ratios (~ 10). Current DI-CB processes provide an excellent choice for communications products up to about 500MHz, and are particularly valuable in coping with the design strictures imposed by 2.7V supplies.

Next are the all-CMOS processes. Because of the inherent simplicity of the technology, CMOS is not only less expensive than bipolar, but is less prone to yield losses due to crystal defects. In mixed-signal applications, the insulating gate is of great value, both in storing voltage-mode signals and in implementing switches that behave like nearly-ideal two-terminal devices. While the progress in bipolar is paced by reducing base width and thus increasing f_T , CMOS transistors have progressed by reducing the channel length. As for bipolar, 'progress' also means accepting lower operating voltages, and thus, supporting smaller voltage swings. Today, these processes are in use:

- Early technologies allowing gate (or channel) lengths down to about $3\mu\text{m}$, augmented by low-quality resistors and capacitors, and providing a substrate PNP, also of poor quality. Though some resistor layers are available, their place in the world of CMOS circuits is of less fundamental importance than in bipolar design.
- CMOS processes optimized for use in analog applications, providing increased density through the use of shorter channels ($\sim 2\mu\text{m}$) and with the addition of capacitors using either sandwiches of metal-oxide-polysilicon, or so-called 'double-poly' capacitors. While the added process steps to support analog design slightly increase cost and lower yields, $2\mu\text{m}$ CMOS is still very capable of being used in complex digital applications.
- At gate lengths of $1\mu\text{m}$ and below, CMOS tends generally to be optimized for large-scale digital applications, with pure analog applications taking a back seat. Currently, the state of the art for mainstream CMOS is at $0.8\text{--}0.6\mu\text{m}$, allowing several million transistors to occupy a chip of manageable proportions. Resistors are available, but generally at a low sheet resistivity ($20\text{--}50\Omega$ per square) making high-valued resistors very space-intensive. Operating voltages are usually limited to 5V. Short-channel CMOS requires a radical departure from classical modeling concepts (such as the mythical 'square-law' behavior) and suffers from such artifacts as non-zero gate current (typically as high as several picoamps, due to tunneling) and non-unity current-transport from source to drain (due to avalanche multiplication at the drain, generating current which flows to the substrate).
- Advanced sub-micron CMOS, with gate lengths in the $0.25\text{--}0.5\mu\text{m}$ range. The drive to such dimensions is both to improve both packing density and speed, the latter though a reduction in all the parasitic capacitances and an increase in the channel conductance. With this speed advantage, many believe that analog CMOS will soon become quick enough to use in such things as cellular phones, operating at about 900MHz. It is at this point that opinions become markedly polarized.

An important difference between bipolar devices—which have junctions across which current flows—and MOS transistors—which are essentially nothing more than voltage-controlled resistors, is that the former suffer from shot noise, while the latter do not. Experience with the design of narrow-band low-noise amplifiers (LNAs) for critical applications (such as in digital handheld phones) leads one to conclude that shot noise is a major impediment to further progress in BJT design. It is fundamental and quite independent of the device geometry, or even the particular technology. HBT's are no less impaired by shot-noise, though they are more attractive in providing lower base resistance, which generates Johnson noise. In wide-band amplifiers, though, MOS transistors suffer from yet another kind of noise arising in the gate region, that is, flicker noise, which begins at frequencies as high as several kilohertz and requires special attention to deal with.

Of course, the comparisons between BJT and MOS do not end there. The utility of the MOS floating gate has been cited (with cautions about its non-insulating tendencies), but against this must be raised the superiority of the BJT in providing very high transconductances (g_m) without having to recourse to large geometries (which, in MOS still requires raising the bias current to achieve a high g_m). Not only are BJT g_m 's the highest possible of any active device, but they are *extremely predictable*. Thus, we have the well-known relationship $g_m = I_C / V_T$, that is, the transconductance is a linear function of collector current; on this remarkable foundation, summarized in the term 'translinear', an immense variety of important circuit concepts are predicated.

By contrast, the MOS transistor has poor parameter control, beginning with its only-approximate square-law, its uncertain threshold voltage, its inexact and invariably low g_m , its low output incremental resistance in the saturation region, and so on. Because of these uncertainties, both analog and digital CMOS designs are forced to follow quite different design philosophies, compared to their bipolar prototypes. However, both depend heavily on the idea of matching.

Not surprisingly, the vagueness of MOS properties is more acceptable in the new domain of epitronics, where precision of circuit function is of little concern. The strikingly elegant CMOS inverter is truly one of the wonders of the modern world, almost a 'gift from heaven' for logic gate applications. Even here, though, we should pause to acknowledge that CMOS logic leaves a lot to be desired. For example, one has to be careful to use the appropriate supply voltage for a given technology (that is, gate threshold voltages). If supplies are too low, gate speed will suffer badly; if too high, transient spiking currents in every inverter and logic gate will raise the supply current. Also, N-input gates are not so elegant.

On the other hand, advanced current-mode logic (CML) is quite benign, being independent of supply voltage, and behaving essentially the same, for example, from 2V to 6V. Furthermore, the supply current, though higher than CMOS at low clock rates, is independent of the dynamic data, and at high clock rates may be lower than CMOS. The absence of supply-line spiking in CML may prove to be an asset in such things as monolithic radio receivers of great sensitivity in which there is a frequency synthesizer operating at a high clock rate. N-level CML gates require the addition of only one transistor per input.

In analog applications CMOS has further disadvantages, one of which is the notorious back gate effect. All MOS transistors have two gates, and the bias on the one formed by the substrate or the well has a significant influence on channel conductance. This limits the available voltage gain that can be achieved, which is often acceptable. It is more likely that the back gate (and its high capacitance) will be troublesome in highly-integrated systems which include both sensitive signal processing sections and noisy digital functions. In this context, the BJT is more attractive, because there is no equivalent back-gate modulation of emitter-collector conduction, and the collector-substrate (or collector-well) capacitance C_{JS} is fairly small. Bonded-wafer DI processes are even more attractive, because this is stable with bias and the substrate material can be made highly conductive.

It is the mixed strengths of bipolar and CMOS that led to the development of BiCMOS processes. Interestingly, this approach is currently finding popularity in digital circuits, including that fortress of high device count and low cost, the multi-megabit memory, where the high g_m of tiny bipolar devices proves to be irresistible.

A few IC companies—Analog Devices was one of them—were quick to see the potential value of BiCMOS in mixed-signal applications. The rationale was fairly simple: the BJT is good for accurate analog, having strong and predictable g_m and excellent device matching (particularly, in V_{BE}), while the CMOS provides an excellent analog switch, is useful where ‘non-loading’ sensing of node voltages is required, and of course, is well-regarded in digital logic and memory applications, so brought together in a monolithic form, they ought to be able to take on the world.

Many BiCMOS processes have been developed during the past decade. Typically, they have characteristics of the sort outlined here:

- A 2 μ m CMOS capability combined with a 3GHz BJT. The PMOS device is useful in amplifier input stages and is placed in the well diffusion, allowing the back gate to be biased (for example, bootstrapped) so as to minimize analog performance degradation. Two NMOS devices may be provided, having gate voltages chosen to allow both logic and analog applications where there may be a negative supply (rather than ground) connected to the substrate. Poly or thin-film resistors may be supported; poly-poly capacitors are usually needed for high-quality switched capacitors and other charge-based analog circuits. Frequently, a lateral and/or substrate PNP will be available.
- A 1 μ m CMOS combined with a 10GHz BJT. Advanced BiCMOS processes will usually include a full library of logic gates and high-level digital functions. In other respects, the adjunct components will be very much like those listed above.
- A less aggressive CMOS (2–3 μ m gates) combined with complementary bipolar transistors of moderate speed, sometimes called a CBCMOS process. Although this has the appearance of the ‘process that has it all’, and therefore ought to be prohibitively expensive, this need not be the case; however, performance of any one component is likely to be well below the state of the art, and packing density will be low.

BiCMOS technologies have received mixed reviews. The basic premise—that one needs both minority- and majority-carrier types in close communication—may be weak. In companies such as Analog Devices where the entire spectrum of processes is available, including many vintages of all-NPN (up to and including a SiGe HBT technology), two generations of advanced CB, all-CMOS processes down to 0.6 μ m, and both ‘NPN’ BiCMOS and CBCMOS, there are signs of an emerging polarization in favor of using processes which are ‘just sufficient’ for the application.

One reason for this state of affairs is that die cost is an ever-increasing concern for IC designers. Many complex functions which a few years ago would have commanded a price of \$5 to \$20 are now expected to be available for \$1 to \$2.

GaAs, CCD, HEMT, TFT, etc.

Any review of contemporary IC technologies needs to at least acknowledge the availability of what might be called the exotics. Almost by definition, these are niche technologies, having one or two very useful special properties and applications but lacking the strength to become *mainline*, often because of cost, but also because of serious performance limitations in other directions. In this brief review, we therefore pass over CCD (charge-coupled devices) and the HEMT (high electron mobility transistor) as of very limited utility in the broader commercial context—even though their proponents can become quite evangelical about their future prospects—and the TFT (thin-film transistor) as of value mostly in display applications rather than signal processing.

Although GaAs *bipolar heterojunction* transistors have been successfully fabricated, the most common GaAs technology is based on MES-FETs, having a Schottky junction over the gate region (hence, metal-semiconductor field-effect transistor). Gold plate-up metallurgy is used. Gate lengths are typically 0.5–1μm in today's devices; width/length ratios of over 1,000 are commonly used to achieve high g_m . GaAs-MES-FETs, like all majority-carrier devices, have no shot noise; they are little more than voltage-controlled resistors. Because the channel resistance can be made quite small, their Johnson noise is also low, making them ideally suited for LNA (low-noise amplifier) applications in a variety of UHF and microwave applications. In this connection, the low resistance of the gate finger (which may be very thin and long) is helpful; by contrast, the poly-Si gate of a CMOS transistor can be quite resistive. Furthermore, the gate is predominantly capacitive, and may often be 'tuned out' as part of an input matching network; that is, the input port is not lossy. An amplifier using a bipolar transistor in a common-emitter configuration, on the other hand, requires power to be delivered to the base node. Finally, the semi-insulating nature of the substrate is often cited as important in implementing various frequency-dependent passive elements, such as inductors, transmission lines, stubs and the like.

On the other hand, there are serious limitations to the technology. Unlike CMOS, whose 'complementarity' is of immense value in both analog and digital design, GaAs transistors provide only one polarity type, limiting design freedom. Also unlike CMOS, the gate is not an insulator, and thus cannot be overdriven (beyond a V_{GS} of about +500mV) without causing significant loading and nonlinear behavior. That semi-insulating substrate is also more often than not a headache, since it results in ambiguous potentials underneath the channel, leading to poor control of conduction and some very long time-constants. And many of those passive elements can now be fabricated on silicon substrates.

GaAs will probably be eclipsed by faster silicon technologies. With CMOS approaching gate lengths of 0.25μm, all-Si BJTs providing f_T 's of 30GHz and SiGe taking these to 110GHz, there is considerable doubt about the need for the use of other materials in the intervening time between now and 2000, and perhaps well beyond. Gallium is also quite a rare element: half the world's output this year will be used in just one particle-physics experiment (a neutrino detector).

THE TOOLS

Virtually all the tools needed to turn these technologies into products will be computer based. Considering how long high-performance computers have been available, contemporary design tools are surprisingly limited. They are mostly in the form of various types of schematic capture programs, circuit simulators and layout programs, often linked into a single working environment.

For analog simulation purposes, SPICE ('Simulation Program with Integrated Circuit Emphasis')—developed at UCB under Don Pederson's guidance during the late sixties and early seventies—has been pre-eminent and all-pervasive, in spite of its warts. Many commercial variants are available at very modest cost. Some of these provide minor enhancements, often in the form of a graphics shell that makes file handling easier. Though still useful, the original 'Berkeley SPICE' is beginning to show its age. A familiar problem is its failure to converge on a DC solution, particularly for large circuits with multiple feedback paths (some positive), without elaborate assistance on the part of the user. Another is its use of deeply-embedded built-in models for transistors, preventing one from modifying the relevant equations or adding new types of devices. Yet another is the limited repertoire of analysis modes, chiefly DC operating point, small-signal AC gain, phase and noise, and transient analysis, using simple stimuli. Finally, there is very little support of such things as a macro language for command purposes, and comprehensive post-processing of data in graphical form for improved insight.

Many microelectronics companies have developed their own simulators to address the needs of contemporary design. For example, at Analog Devices, we have developed ADICE (not very imaginatively standing for 'Analog Developed Integrated Circuit Emulator'). Being shaped by the design community, it has evolved over a period of about fifteen years into a very complete design environment, and includes numerous proprietary features. The capability for exploring detailed aspects of circuit behavior, using unusual analysis modes and stimuli, supported by good graphical post-processing, contributes significantly to product quality and manufacturability.

However, the expectations of the modern designer continue to put pressure on the tool developers. We want more realism, particularly in connection with modeling noise mechanisms. For such applications as determining the input-referred noise of a low-level amplifier, the present linearized noise calculations are adequate. For some high-gain receivers, where even the input noise can drive sections of the circuit into limiting, this approach is useless, as it is for determining the noise figure of mixers and the phase noise of oscillators. We want more speed, particularly in analyzing nonlinear circuits having widely disparate time-constants. In some cases, where the nonlinearities are weak, we can use harmonic balance techniques, but for strongly nonlinear circuits, this method becomes as slow as full transient analysis. And we want more help in keeping an eye on background subtleties which we might not notice as we pursue the main simulation objective (for example, warning lights of various kinds).

In many respects, what we want is for our computer to act more like a silicon companion, a willing apprentice, looking after myriad details on our behalf, now and again interrupting us as potential problems loom, carrying out the tedious tasks and leaving us to get on with the creative work. And we want our simulation environment to behave more like a bench-full of instruments: spectrum analyzers, communications test-sets, modulation-domain analyzers, signal generators, and so on, leaving less to chance before committing the design to mask-making and prototype production.

We also need to engage computers in helping us to deal with the necessity of documenting our work. The need often arises in organizations to research how long a project actually took, and what resources were consumed. Often, finding the starting date for an old, completed development is next to impossible. Written records—the 'lab notebooks'—have traditionally been very important in tracing the development and ownership of intellectual property for the purposes of patent processing and (perhaps) subsequent litigation. Unfortunately, as our design work migrates toward being totally supported by our computer, or network of computers, their hard disks become the primary repository of all our records—project schedules, schematics, simulation files, numerical and graphical results, notes to ourselves about this or that reason for carrying out a certain simulation experiment, cell layouts, complete layouts, back-annotation of layout parasitics, manufacturing schedules, wafer lot IDs, and on and on—and the 'lab notebook' gets relegated to a pile of unattended paperwork on the desk.

The problem here is that no-one seems to be facing up to the shift in work habits, and taking the necessary steps to first, help the design engineer replace an earlier reliance on hand-written notes by an *automatic project documentation system*, for his or her own benefit as a depository of important ideas, and second, to turn the fragmented, un-witnessed, un-time-stamped and therefore 'unofficial' computer records into legally-acceptable documents. Records of this sort would have obvious value when the time comes to check the design against a check-list prior to maskmaking, although, if we continue to think about the computer as our ever-willing silicon companion, it would assist us automate such check-list procedures, and to attend to them on-line rather than at the end of the project.

By the year 2000, this capability will almost certainly become commonplace. We already are seeing numerous personal assistants (such as Apple's 'Newton') which presumably represent a response to a market need for devices to help we humans organize ourselves. Yet, the electronics community often seems to be the last to benefit from its own advances. Though we cannot accurately envision the state of affairs seven years out, we can think about how might this be done *today*. Consider the following scenario.

We begin a new project by generating a special type of folder, which actually starts a program. It presents various menus for the record area we wish to enter. We might first select 'Objectives', which opens a text-editing window, into which we begin to set forth the nature of the project, its origins and aims, in free form. As we close that subwindow, our entries up to that point are *time-stamped*. Each

time a menu function is selected, old entries are presented read-only but the contents may be expanded; each time a subwindow is closed, it again gets date-stamped. Since the determination of the time and day is critical to the legality of these records, we will need to use clocks whose honesty is somehow verifiable. There are various ways to achieve this; one might be by using a signal from a standard-time transmitter such as WWV in the USA. However implemented, it would be prone to tampering, but in this respect it is no different from hand-dating, and at this point legal verifiability is not needed.

We next might select a menu item called 'Technologies', which would inform us of available IC processes, with sub-menus to explain these in detail, tell us their location, whether captive or foundries, general cost guidelines and so on. After selecting the process we expect to use for our project (perhaps also indicating a second choice) we close this window, automatically adding the new data to the overall folder and date-stamping the entry. The contents of this file will play an active role in subsequent operations.

Next, we might call a menu item called 'Schedule'; this would connect our environment into company-wide active calendars, including our own, and allow us to explore the impact of our project on all the other projects in the group, a division, or company-wide. Depending of the structure of the company and the status of the user, we may or may not see the other schedules for these other projects, but the scheduler subprogram knows about them, and about all the available resources. Using critical-path modeling, it would provide us with the best possible estimate of the release date for our product, using all the known data (minimal at this point, so the estimate would be provisional).

Thus, we would be able to choose other interactive subwindows, such as 'Team Members', 'Testing', 'Packages', 'Cost Analysis', 'Revenue Forecast', 'Customers', and enter data to gradually define the personality of the product and the project, with the help of many prompts and frequent opportunities for free-format text and data entries. These could be entered as many times as desired and in any sequence; in each case, prior data (even our own) is read-only, and closing the subwindow adds a new time-stamp.

Still within this design environment, we will eventually select the 'Schematic' window, and begin our first circuit, or copy in older files. Each 'Save' or 'Netlist' operation adds a time-stamp to an adjunct journal file (normally invisible to the user). Eventually, we will be working between several open subwindows, including one for control of the various 'Simulation' options. Here, the electronic notebook would keep track of our own notes about the purpose of each experiment, or the experimental conditions. This might be done by asking for free-form comments each time we begin an analysis (the SPICE 'go' command); if we have nothing to say, we just hit 'return', but every line entry keeps the process open as long as we wish. Whether we enter text or not, our silicon companion makes a note of the type of analysis, the sweep range, certain other data (for example, the current value of selected parameters and variables), etc. At the end of the day, all of these files would be sent to a *write-only* CD-ROM.

CAN DESIGN EXPERTISE BE CAPTURED?

Today's workstations and tomorrow's connection machines will continue to open our eyes to unimagined possibilities and help us cope with the never-ending increase in the complexity of the overall assignment. Corporations without advanced design tools, and without the people having the skills to use them thoughtfully and creatively, will be at a severe disadvantage in the highly competitive semiconductor industry of the mid-nineties and beyond.

Since we are considering the tools we'll be using up until 2000, it may be appropriate to ask whether the design process itself will be aided by expert systems. What is Design Expertise? What does it mean to "Capture" it? Is it enough to just capture the expert's knowledge? Don't we really want to *capture the expert* and all his living, breathing functional skills to solve problems? What is the connection between *expertise* and *creativity*? If the creative spirit is absent in the machine, can we expect expert systems to find interesting *new* ways of doing things, and if they can't, are they still useful? Can a machine which is completely deterministic ever come up with a new idea?

Can captured expertise improve time-to-market and product quality? Or, on the other side of the coin, might the use of such systems lead to delayed product introductions, because of the greatly augmented knowledge base and richer opportunities for wide-ranging exploration and endless refinement? Is it possible that, in spite of all the computing power now at the fingertips of the graduate designer, product quality might actually decline because of only superficial understanding and at times a critical ignorance of some practical detail not captured within the expert system?

The questions are endless: Is it cost-effective to develop expert systems? Are conventional (that is, von Neumann) machines adequate platforms for expert systems? Is circuit design an art or a science? Is analog design more artful than digital design? Can solutions forged from analog-cell libraries compete effectively with full-custom analog cells? Presuming the answer to that question to be 'No', then what can be done using expert systems to put better design capabilities, combining the refinement of custom design with the rapidity of library based design, under the fingers of everyone? What makes analog cell design so much more difficult than digital VLSI design? Just how much do we expect from this Aladdin's Lamp, anyway? Do we imagine that our future design tools will blend the youthful vigor of an eager apprentice with all the wisdom of a lifetime, or will we settle for fast calculators with large data-bases?

We won't be able to answer all these questions, but it's interesting to think about them. There is a lot to remember in an engineering job. It begins at an early age with the accumulation of basic technical information, that is, organized facts, physical data, other peoples' discoveries and results (both classical and contemporary), some awareness of the fundamental limits to realizability, and so on. To this, the mature engineer adds know-how, including a portfolio of workable solutions, design algorithms and procedures, some of which will be very specific, others quite general, plus an understanding of modeling, etc.

Later on, technical knowledge is augmented by market knowledge. This requires a strong grasp of the potential utility of the company's products in a broad field of applications, as well as the more prosaic matter of knowing who the customers are, where they are, what their unit volumes are likely to be, who they are buying their components from right now, and so on. As time goes by, all of this grows into a complex state of mind in the which acquired skills blend and the ability to cope with numerous interrelationships between various disciplines assumes an increasingly important role. From using well-established approaches to circuit design, the engineer now adds the ability to create novel device or circuit topologies, not to be found in any book, and to instinctively seek to optimize these structures, a process which is most effectively achieved by appropriate models which stress the fundamental aspects of performance, rather than by all-too-quickly appealing to simulation for answers.

Eventually, the seasoned engineer possesses a highly-personal corpus of experience, incrementally and progressively acquired over time, in which original achievements play an increasingly important role, and in which idiosyncratic attitudes, viewpoints and preferences become more strongly manifested. It might be said that this process embodies a gradual shift from a knowledge-based world to one in which the elusive human quality of wisdom is acquired. Will *all this* be captured in the turn-of-the-century machines?

It is very likely that friendly, intelligent machine-companions will indeed become increasingly commonplace, a totally-integrated part of our lives, and citizens of the 21st century will not be the least bit disturbed when these exhibit some capabilities that outstrip their own. Of course, every designer has his or her own unique ideas about what will be needed from knowledge-based systems in the workplace of the future. I believe that these needs can be reduced to four important criteria: First, they must provide *leverage for the human mind*; second, they must do this in some rather *specialized and directed manner*; third, they must do so with impeccable *accuracy and reliability*, much more so than the engineer expects to be capable of, and fourth, they must be able to modify their behavior with time, that is, they must *benefit from past experience* and learn new capabilities.

Most engineers want their silicon companions to be able to multiply their own innate abilities to achieve some particular task. Nobody wants a machine to come up with brand-new circuit or product ideas, even less, to have to deal with a machine that becomes so good at doing this kind of thing that it threatens the human. Years from now, we may think differently about this, but for the present, at least, we want machines to know their position in society. We want smart but submissive silicon slaves.

One reason why few, if any, engineers want machines to cross that line is simply that engineering is so much fun. The rush of endorphins at that Aha! moment is just too good to share with anybody but a fellow human. The sense of achievement and pleasure in finding neat new ways to use all the technology at our disposal is, I believe, the chief payment for the long hours spent in the lab and the many other times when we find we've hit another dead end.

In other words, we humans want to remain the architects of our work, but will gladly will accept help with the tedium of looking up numerous detailed facts and making many calculations. We want levers for the mind. Just as the invention of the mechanical lever didn't alter the tasks men undertook (at least, not at first), but just allowed bigger jobs to be done, far beyond the power of human muscles, so today we look for machines which can help us with the immense computational burden of designing a modern integrated circuit.

There is little joy in calculation. More importantly, in today's highly commercial world, there is no time left for that part of the design challenge: products must be brought to market quickly, and their performance, quality and reliability must be impeccable. We understand this very well, and know that future work in the area of capturing design expertise must, at least in part, be focused on aiding us to further hasten the introduction of new products. Consequently, we have an insatiable appetite for machines that not only provide this essential cerebral leverage, but also that carry out their microminded tasks at speeds that make last year's model appear glacial.

Our second provisional criterion was that in our design work we need help in very specific, particular, narrow ways. We want our expert system to be really good in dealing with minutia of detail which might tax the resources of the average designer. In fact, to a large degree, that is actually what we mean by the term 'expert', in human terms. We don't immediately think of an expert as one who has a wide general knowledge of a field, although that is sometimes the case, but rather as one who has devoted his or her life to some rather special topic that only a few people, perhaps in the whole world, know much about. The term 'expert' also implies, not just the ownership of specialized knowledge and experience, but also the idea of accuracy in applying them, perhaps even a hint of infallibility. This touches on our third criterion, which calls for our silicon companion to be accurate and reliable in its procedures and calculations.

Finally, we would like our expert system to profit from having gone down the same path many times. We may not dignify this process by the term 'learning', but the more prosaic aspects of learning and experience can very easily be achieved by machines. For example, in using integrated CAD tools to develop a new IC product, we may have generated the layout of several previously non-existent transistors. We expect those cells to 'be remembered' the next time we build a layout in that environment.

SUMMARY

We have attempted to trace the history of electronics and found that its origins are in writing, printing and other means to capture and disseminate knowledge. Electronics is the latest in a long heritage of 'knowledge technologies'. Since the quest for knowledge is unlikely to ever cease, electronics will doubtless remain of crucial value to society, until superseded by something more potent. The next big breakthrough will come as machine-executable knowledge comes into its own. Inevitably, electronic knowledge will someday be used to create knowledgeable electronics; once that point is reached, life in Earthtown will be *very* different.