INSTITUTE OF ELECTRICAL
& ELECTRONICS ENGINEERS

CENTENNIAL TECHNICAL
CONVOCATION

October 8-9, 1984

Tuesday, October 9, 1984

Finding the Technologies

Session II

Franklin Institute
Philadelphia, Pennsylvania

SecrePhone

RECORDING AND TRANSCRIPTION SERVICE

Administrative Headquarters
2874 Limekiln Pike, North Hills, PA 19038
(215) 885-3815)

New York: (212) 425-7520

MR. J. L. EVERETT, III: Good morning, ladies and gentlemen. It's my privilege to officiate at this second session which I think you'll find exceptionally stimulating.

I believe those of you who have attended the other sessions know the format. We'll have a presentation followed by a discussion by a discusser and then a chance for the audience to ask questions and make comments about each of the presentations. We'll take a break after the conclusion of the second presentation and all the discussion that follows. We'll move out quickly for a spot of coffee and a stretch and then come back for our last paper.

I know many of you were at the banquet last night. You know it ran a little late. Those of you who are here, my commendations for getting up and moving at an early hour on a dismal day. And for those who continue to filter in, I'm going to take a little time right now rather than at the end of the program to let as many get seated as possible before introducing this most fascinating program.

With respect to schedule, lunch is scheduled at noon. Whether you make it or not will be up to this audience. I'll do all I can to help you make it, and I'm sure the speakers and discussers will.

If you will look in your program, you'll notice that we have really an all-star cast. Collectively they represent a wealth of experience and scientific information and achievement.

I'm a little bit reminded of the story of the itinerant preacher in North Carolina where I was brought up. In a farming community around the turn of the century, most rural churches could not afford a full-time minister and so they employed a circuit rider as he was called in those days. Usually the circuit rider, the preacher, took care of four small country churches, and he would appear at each church once a month on the fourth or the third or the second Sunday and preach the sermon. Of course, that was the Sunday that all the people looked forward to.

One Sunday the minister was headed for this small, remote church and he found the sanctuary empty with one exception -- one elderly farmer. He said to the farmer, "I suppose there's been some mistake in the schedule."

"Well," he said, "preacher, I don't know. All I knows is how to tend to my critters. But this is Sunday."

"Well," the minister said, "without anyone here but you and I, I don't suppose we ought to have the service."

"Well," he said, "preacher, I don't know nothing about preaching. All I knows is critters. But if I was to take a load of hay out in the field to feed my critters and only one of them showed up, I think I'd feed him"

He said, "I got the message. We're going to have the service."

And they did. They had the full service. As a matter of fact, he preached his longest and firiest sermon and he prayed for a full half hour. As they were walking out together very late in the morning, he said to the farmer, "You know, I'm awfully glad you told me that story about feeding the critters or we wouldn't have had this service."

"Well," the old farmer said, "preacher, I told you I didn't know nothing about preaching and I don't. All I knows is critters. And I did say if I'd take a load of hay out in the field to feed my critters and only one of them showed up, I'd feed him. But I didn't say I'd drop the whole damn load."

(Laughter)

Well, you can see from the all-star cast that in just three hours this morning there is no way these gentlemen can drop the whole load. But nevertheless, I think they're going to try to condense it into a very, very interesting and informative presentation.

Our first speaker this morning is Dr. Robert N. Noyce who is Vice Chairman of the Board of Directors of Intel Corporation from Santa Clara, California. A co-founder of Intel in 1968, he was President until 1975 and Chairman from 1975 to 1979. Dr. Noyce is co-inventor of the integrated circuit with Jack Kilby. They have jointly received the Ballentine Medal of the Franklin Institute and the Cledo Bernetti Award (Phonetic) of the IEEE for this work. With Gordon Moore, he has received the AFIPS Harry Goode Award for leadership in computer science. Dr. Noyce was awarded the National Medal of

Science and the IEE Farraday Medal (Phonetic) in 1979 and the IEEE Medal of Honor in 1978. He's a member of the National Academy of Science, the National Academy of Engineering, the American Academy of Arts and Sciences, and a Fellow of the IEEE.

In the interest of time, I will not read you his achievements. I would refer you to the program which has them in a very, very condensed form. I can only tell you that this gentleman is without a doubt one of the leaders in the development of large semiconductor integration -- and very large semiconductor integration. He's practically written the book in these and founded at least two companies that have had a great deal to do with the growth in that industry.

It is my pleasure at this point to present Dr. Robert N. Noyce.

(Applause)

DR. ROBERT N. NOYCE: In a centennial year for the IEEE, we have a tendency to look backward and to look forward. In that sense, it's living up to one of my tenets which is you can look forward about as far as you can look back. Now when we start talking about microelectronic technologies, it isn't quite possible to look back a century.

As a matter of fact, if we look at the seminole events, they were proably in the mid-40s. Harvard Mark I, the first of the electromechanical calculators was 1944. ENIAC (Phonetic) with vacuum tubes was 1946. The Noiman's computer concept came out that year. And then, of course, in 1947 was the invention of the transistor by R. Dean Bratman and Shockley which is the beginning of microelectronics technologies as we know them today.

Now if I follow my rule then and try to look forward about as far as we can look back, that suggests that we can look out to about 2020. 2020 was mentioned earlier in the conference. That was the time when we can put all books on one disk if I recall correctly.

But since the subject today is a look at the technology for the next century, I'll try to stretch a little bit. But in doing that, we'll find that we run into some very hard limits with the approaches that we're taking today and that those approaches are indeed going to become obsolete. Then I'll speculate as to how we might get around some of those limits to find a path such that microelectronics won't become an entirely extinct species. We are already beginning to look at biological models to suggest answers to the problems that we'll be encountering. And I suggest that inquiry into biological systems will increase in the future. Finally, I'll speculate about how the next century's technology will affect our society and vice versa.

Talking about limits. If we do extend our sidelines for another century and look at the limits that we'll encounter, then perhaps the song from "Oklahoma" gives us the fitting theme -- "We've Gone About as Far as We Can Go." Within this time frame, we'll have to change the current course drastically because of these walls. I see three main limits on the horizen -- the limits to our computing elements due to basic physical laws, the economic limits (that is how much money is our society willing to spend on these things), and finally the limits to the usefulness of the approaches that we're now taking.

Let me have the first slide and we'll take a look at the physical limits. Some 20 years ago Bob Keyes talked about three limits on semiconductor computing elements. Fundamental limits on the power/speed performance are indeed imposed by the uncertainty principle. They are imposed by the power required to propagate signals above the noise level at high speed, and they are limited by thermal energies. Let me concentrate on the thermal energy.

The limitation that he foresees there is KT. That implies that we have one electron crossing one thermal barrier -- that is the total energy involved is about 1/40 of an electron bolt at room temperature. It is indeed hard to see how we could get a computing element to dissipate less than that. You could say we could lower the temperature but then the other limits that Keyes mentioned would block us somewhat tighter than this limit implies.

If we look at the rate of progress toward that KT limit over the last 20 years and assume that that rate of progress will continue, we find that we'll run into that fundamental limit about the year 2020. Realistically, we'll reach that limit sooner because we will have to have more than one electron crossing that thermal barrier and that barrier will have to be higher than thermal energy in order to keep it there. So if we were to say that it would 10 electrons and 10 KT over E, then we are now a factor of some 5,000 away from that limit.

Let me have the lights again here.

As we get closer to that limit, random errors are going to creep in. Even at a quarter micron geometry for most demo S devices, any random event would destroy the charge that we're trying to read. Consequently, we'd no longer be able to rely on any discrete element to provide the value. That suggests that the approach that we are now taking to logic will become obsolete. Something there will have to change.

The second limit is an economic limit. Incidentally, I've put in the background of this slide an extinct microprocessor just to suggest that we're going to have a lot of those in the next century. But every simply, as this slide shows, if the rate of growth in the value of microelectronics were to continue as it has in the past, the sales of microelectronic devices in our country would equal the projected GNP for the United States in about the year 2020 also. Now I suggest that that will not happen.

(Laughter)

We will have some other things to do in our society besides building MOS devices. Again, something will have to give.

The last limit I'd like to talk about is the limit on the usefulness of today's approach. We have a great deal of capability with our highly integrated processors, dense memories and the like. And we will be able to make much more complex devices in the future -- but will they be useful? Here let me give you a very simple example of the limit of today's approach.

You can tell what that is. It's a matronly woman. Even as an infant, you were able to recognize your mother long before you were able to talk, but that's still something we cannot yet do with a computer. We could say then that the computer is indeed in its very early infancy. The simplest task that the child takes on cannot yet be accomplished by the computer, and yet we are increasingly asking the computer to do jobs that humans do. That's what artificial intelligence or the fifth generation project is all about.

In the next century we certainly will see even more exploration of new architectures to perform jobs for which the Noiman architecture is ill suited. These jobs take all shapes, including that of recognizing mother or any other pattern. And although the computer in my car calculates the right amount of fuel hundreds of times a second, I think very few of us would trust a computer to drive us to work and do all of the pattern recognition that is required for that job.

Again, looking at biological systems, they are in large part self-healing and self-repairing. We've not devised a way to make computer systems expert at those tasks. Perhaps one of the most astounding capabilities of biological systems is self-reproduction. And although computers are used to design computers or to make them, this kind of reproduction is nowhere nearly as complete as it is in biological systems. Thus I believe that the incapability of our present systems to do the task which we would deem desirable will exorably lead to new approaches to computer systems.

In spite of the limitations of the usefulness of our current machines, their range of applications is indeed very broad and will certainly continue to expand very rapidly over the next several decades. Semiconductor device engineers will continue to make more complex devices in ever shrinking dimensions. Our lithographic techniques will become more and more sophisticated. We'll use more layers of thin films to extend the usefulness of the underlying silicon or other material. And yet these approaches will run out of steam too and be their own limitation. Again, biological systems may offer us a way around these systems.

The basic characteristic of that biological system seems to be that it's self-organizing. In some senses, we've used self-organizing to aid our production of semiconductor devices already. Techniques such as self-aligned gates where one pattern....the existence of one pattern determines the geometry of the second pattern might be one example. Hopefully, we can extend those techniques even further by looking for affinities within materials which would allow the materials to organize themselves.

Another type of self-organization comes in redundancy. As we push the silicon limits, we find that the instance of errors will increase more and more and that redundancy will be required to assure that stored information has a high probability of being preserved correctly. The redundant rows and columns in memory circuits that exist today are an example and perhaps the precursors of a broader use of redundancy. Error correction on memories that we now use at the systems level is another example of self-healing and could also, will also probably be used at the chip level. In sum, those approaches can be seen as part of a growing tendency toward self-diagnosis and self-healing.

There's another limitation that the increasing complexity suggests to us. This is a slide that Gordon Moore presented in 1980 showing that the complexity of our design was on a steep rise and would eventually overpower us. If we do that study a bit more quantitatively, we find that the design time, the productivity of our designers is not keeping pace with the demands made upon them for new designs. And indeed, the design times are increasing exponentially. We know that anything that increases exponentially finally runs out of steam. And the answer to this, of course, has been the various attempts at design automation.

Design automation has many different levels. The design process itself has many different levels and many different representations of the final product -- from the architectural design and specification of the product to chip layout to a test program that the chip must pass. So far design automation has concentrated on the lower levels of this design process. These levels are relatively easy to automate and it is relatively easy to verify the correctness of the solution. The creative part of the work, the part that is much more difficult to automate, is the synthesis involved in both understanding what the market will require as the final product specification and in actually developing that new architectural specification. Yet, as time goes by, automation will move up this hierarchy and be applied to the more complex problems.

I can cite as an example the design of a 32 bit successor to a popular 16 bit microprocessors. One of the requirements was simply that the successor be software compatible with its predecessor. In achieving that compatibility, the design effort was between 15 and 20 person years. Yet that's a relatively straightforward job and it would have been, as far as I'm concerned, far more preferable to have had a primitive instruction in the CAD system which simply said make it compatible with the previous one that you did.

Everytime in the past that an extrapolation of existent trends has grown to ridiculously large numbers or created awkward design parameters, there has been either a change in methodology or we have suggested that it's a job for the government to do.

(Laughter)

For example, in the mid-50s, we learned that larger scale computers couldn't really be made because of the tyranny of numbers of the reliability of the wires and sockets and soldered joints that wouldn't let the computer operate long enough to be useful. In the mid-60s by extrapolating the design trends in integrated circuits from MSI to LSI and to VLSI, we determined that we would have had to use more than all of the engineers in the world designing integrated circuits in order to use large scale integration properly -- very similar to the telephone operator problem. And you cut off that extrapolation by finding a new way to do things.

In this case, we used programmable circuits, whether they be proms or programmable logic arrays or microprocessors wich could be programmed to many different applications. And that eliminated the tyranny of having so many designers making a unique design for each application.

If history is any example, then we are rapidly approaching the time when we'll have to create a new design approach in order to take full advantage of the capabilities that we've built for ourselves. What new approaches might we think of?

In this talk I've made references here and there to self-organizing and self-healing systems -- ideas that come straight from biology. I think there are other aspects of biology that suggest a path for the future. Biology also gives us a way of building structures and is a model of the system that has been extremely successful over time. It functions with less than perfect devices, and I might indicate that our industry is very good at making less than perfect devices.

(Laughter)

It seems to work on a basis of statistical accuracy rather than absolute accuracy. Input/output channels have their own processing. There is self-diagnosis. There's redundancy. There is adaptive behavior. And there is trainability. These are all hallmarks of biology's levels of compatibility that are either unattainable or poorly implemented in our current architectures.

In a recent article in Science by Waldrop, he noted that the brain is beating out computers with neurons that operate a million times slower than silicon and, of course, the secret is in the wiring.

(Laughter)

That, incidentally, is what microelectronics is all about is the wiring. Any yield model will say that the individual transistor is cheaper to make than the transistor in the integrated circuit because you don't have the multiplication of yield factors. So integrated circuits are wiring.

Comparisons like these imply that the von Noiman machine may go the way of the brontosaurus or the woolly mammals and be a dead-end response to the requirements and the dynamics of the world around it.

This next slide shows that we indeed have a long way to go. The U.S. market for semiconductor devices is about 10 to the 10th dollars per year. If all of those devices were made up of memory devices, memory disks, the cheapest element that the semiconductor industry produced, the bit costs us about 10 to the 4th, so consequently the total market would be of the order of 10 to the 14th bits per year. That probably is obsolete by now. That was true two weeks ago. I might note that that is about the number of synapses in the human brain.

Turning then to biological systems, we have a long way to go if we are going to begin to match the human brain in complexity. If we take all of the production of last year, we've produced about the complexity of the human brain. That suggests that we are far from exhausting the possibilities for there seems to be some reason to have on the order of 4½ billion human brains on this planet. I might also parenthetically note that if we were to continue the same rate of increase in the number of devices that we make, we would have produced about the number of elements that exist in all human brains, again, by -- you guess it -- 2020. It would appear then by the early part of the next century, we're going to run out of steam on all of our projections.

In contrast to the amount of brain power that we have produced, in the production of artificial muscle power we've gone far beyond human muscle power. If we look at the energy output for world oil consumption alone, it's equal to about five times the total human energy output today. That suggests that we still have a long way to go with brain enhancers as well as muscle enhancers.

Another telling facet of biological systems is that they seem to use constructive methods to put things together rather than subtractive methods such as our photolithographic processes. In the biological system there is an environment where raw materials exist to build cells, a pattern is introduced, and the molecules are formed upon that template. This molecular process conceivably could provide a way of achieving our new supercomplex, superfying brain structures that we would like to have and allow nature to help us build the complex structures of tomorrow.

As we further examine biological systems, we find that input and output channels are quasi-analogue, quasi-digital. The signals along the nerves appear to be digital signals without analogue's amplitudes or timing. And indeed, our sensors do a great deal of the processing at the sensor site itself as studies of the eye and ear have shown.

Furthermore, with the biological system's very high degree of connectiveness, truth can exist in several places at the same time and can exist statistically, so that the failure of any single element will not cause the failure of the entire organism. Even in cases of major failure the body can in many instances adapt to that failure and continue to function. That ability to adapt will likely become another important part of the architectures of tomorrow.

Adaption is part of the quality that we call intelligence -- the ability to learn from the environment and to adapt to it, to react to and anticipate change. We are beginning to recognize that the ability to deal with fuzzily defined problems is essential to the development of an adaptive capability and the development of artificial intelligence. There are those who feel that artificial intelligence is an abstract problem and that it is not related to the physical structures embodying that intelligence. Increasingly though there is a school of thought which believes that the basic structures will have to be changed in order to realize that quality which we call intelligence. For this to happen, the pathways of knowledge and communications will have to be increasingly connected.

Borrowing again from Rodgers and Hammerstein who wrote that the cowboy and farmer should be friends, for maximum advancement I believe the engineer and the biologist are going to have to be friends. Progress will come from the cross-pollinization of various fields of knowledge. This cross-pollinization will offer new approaches to the solution of evermore complex problems.

As we learn how the brain functions, our increasing understanding for new ideas for organizing electronics systems. Of course, we don't need to know exactly how it works. We can build airplanes that are powered by different mechanisms than those by which birds fly. The point is that if we want to generate approximate solutions to fuzzily defined problems, if we want to provide superprocessing input/output devices like the eye or the ear, it will probably require an interdisciplinary approach.

Until now, we've been going the other way. In order to understand the brain, we have tried to use the computer as a model of it. Perhaps it is time to reverse that reasoning. To understand where we should go with the computer, we should probably use the brain as a model. If we do so, we may indeed find that today's computer is just a fossile of our age.

With increases of knowledge, each generation has had a more exciting challenge than the previous one. The future would appear to be no different than the past in this respect. This means we have both opportunities and problems. The new technology will solve some of those peoblems as it expands our capabilities and create other problems as it makes us more god-like in our ability to mimic nature. At the same time, information is becoming a strategic resource as physical wealth has been in the past.

The change in our capabilities will mean a change in our society. Just as the age of mechanization moved most people in the world from subsistence agriculture to manufacturing, the increasing level of automation will mean that human activity will concentrate more and more on those things which machines are unable to do. Consequently, as the machines become more capable the definition of work will change.

There's certainly going to be opposition to that change as there has been opposition to change in the past. But there is little doubt that the more powerful information machines will someday be like telephones -- we'll all have them, we won't think twice about using them, and we would not want to live without them.

In conclusion, the next century will bring many challenges and changes. We'll reach the limits of our traditional approaches to devising new computer elements. We will reach the limits of the utility of our present architectures. Approaching those limits, we'll look for new ways to progress. Biological models hint at a coming mutation in methodology. There are certainly problems and opportunities enough to keep us occupied for a century to come.

Thank you.

(Applause)

MR. EVERETT: Thank you, Dr. Noyce.

To discuss Dr. Noyce's paper, it's a pleasure to present to you Dr. Carver A. Mead who is Gordon and Betty Moore Professor of computer science at the California Institute of Technology. He has taught there for over 20 years and his current research focus and teachings are in the area of VLSI design, ultraconcurrent systems, and physics of computation. He has worked in the fields of solid state electronics and the management of complexity in the design of very large scale integrated circuits. In addition to his wide range of interests in solid state physics, microelectronics and biophysics, he's written with Lynne Conway a book, the standard text for VSLI design. The name of the book is "Introduction to VSLI Systems."

Professor Mead is a fellow of the American Physics Society, a member of the National Academy of Engineering, a recipient of the T. D. Callinan Award in 1971, and the Electronics Achievement Award in 1981 and was recently presented the Harold Pender Award.

It's a pleasure to present to you Professor Mead.

(Applause)

PROFESSOR CARVER A. MEAD: It is really interesting to look at an industry -- the industry we have come to call the semiconductor industry -- which has evolved in its basic technology over a factor of about a million in the last 20 or so years and is still making products that are based on the single idea presented by Allen Turey in the mid-1930s.

What I want to talk about today is an extension of what Bob said. We have not begun yet to use the semiconductor technology for anything except cost reduction of old ideas. It is just now that we are beginning to see it as a medium for innovation of new kinds of systems.

If I can have the first slide. As Bob said, current computers are really good at doing extremely well defined problems in a very precise way and getting very precise answers. One of the things you notice when you try to work real system problems is the sort of exponential explosion of the band widths of computation as you get right to the inputs or you get right to the outputs. The problems that come to mind are image processing, seeing and hearing if you like, on the input end, and the generation of graphic images or high quality sound on the output end. I'd like to show you just a couple of examples out of my own lab because I'm the most familiar with that work.

The first example is when you try to actually generate really high quality music. Music has been an important part of human culture for a long time, and it's natural that we try to use our latest technology to do a really good job of making musical instruments. When you look at that problem hard, as we've done, one of the really fruitful approaches to synthesizing music is basically to build a model of a physical instrument. That doesn't have to be a real physical instrument, but it can be. It can also be an imagined physical instrument. There are continuity problems with sound that allow this kind of synthesis to do a really, really good job.

If you try to run that kind of a program on an ordinary computer, you find that it takes about ten minutes to generate one second of sound for one voice, like one string on a guitar or one bar of a marimba or that sort of thing. With the kind of approach that we've been taking, modeling these equations with finite different schemes, you can do that amount of synthesis with a dedicated architecture on one chip. That gives you a factor of about 600. In other words, one chip is in this application because of its dedicated architecture, a tremendous amount of concurrency in its computation, is doing as much computation as about 600 medium sized computers.

This is an experimental chip we've been using for this job. Each of those horizontal slices there is one small finite difference engine. The new chip we have under design right now will have about 60 of those finite difference engines together with the interconnect scheme which allows them to synthesize musical sounds.

The second example I'd like to talk about is that in vision. You all know, and Bob has said, we can do with millisecond logic with our eye a task that our most powerful computers take hours to do. There's a small factor of 10 to 10th or so there. It suggests that maybe we should think about new ways of doing that job.

This is a reconstruction of the wiring -- just the bottom level -- in the retina of the rabbit. The red cells are the....these are all the ganglion cells that directly feed the optic nerve. They don't show the other levels of the retina. The red ones are those that detect a dark object coming into a light background. The black ones show a light object coming into a dark background. And I should point out to you that it's very important if you're a rabbit to be able to detect the motion of objects in real time.

(Laughter)

You can't wait around for a supercomputer to do that several hours later.

This is a picture of an experimental one dimensional retina that we have just designed and are experimenting with at the time. The individual circuits in this retina have been tested and are known to work. I only got this chip back a couple of weeks ago so I haven't been able to test the whole thing yet.

It starts out at the top with photo sensors. This entire design was motivated by our intention to understand how the mammalian retinas work and as near as we've been able to reconstruct it is sort of cell for cell an embodiment of our understanding of how the biological system works. Down at the bottom, those cells are our version of the ganglion cells that feed our very, very poor version of the optic nerve which comes off the chip at the bottom. And up at the top are the row of photo sensors.

One of the things that strikes you immediately with this technology is that you can understand a great deal more about what biological systems do after you've tried to do one yourself. The interaction of biology and engineering is not a one-way path. Biologists tend to view themselves as strictly observational in their interaction with the biological systems. Engineers tend to review themselves as synthetic, and actually there's a lot of feedback between the two approaches. We will learn a lot more about biological vision systems as we try to build our own vision systems. We're already learning a lot.

Now this is a one dimensional retina and it's a very small fragment. We don't have the many levels of nerves that the living system does so we've had to map what would be different levels in the retina down into two dimensions. Nonetheless, within the next decade we will certainly have built very, very authentic real time vision systems on silicon. There's no question about it.

One thing you notice if you do this game a while -- that is the game of mapping real applications down onto silicon -- is that they're all different. I have tried myself and many people have tried to find the general purpose highly concurrent system. At this point I don't believe such a system exists.

When you get to highly concurrent systems, the algorithms are inseparable from the architecture. That tells me that there's going to be more and more mapping directly the applications down onto silicon or what today's newspapers call custom chips.

We noticed an important thing when we got personal computers. Probably most of you remember when computer companies thought they should write all the software there was in the world, and we didn't have much software. And what we did have didn't work very well. When we got the personal computer, it disconnected the creation of the computing engine from the creation of the software and the applications, and we got an enormous wave of innovation in software. The same thing is starting to happen with silicon.

The mapping of applications onto silicon is not the province of a semiconductor company. The advent of silicon foundries which will fabricate designs made by people who are good at mapping applications onto silicon is really rejuvenating the whole industry. Instead of a single monolithic standard components business as we've had in the past, we're seeing an industry which is growing a new structure. And along side of that standard components business imbedded in all the applications, we see growing up a semiconductor foundry services, people in the business of fabricating wafers for those that have applications to put on them, and we see the evolution of a new design tool industry which provides the bridge from the system application to the silicon.

If we look back -- as Bob said, we can see forward about as far as we can see back -- I cnn't see much further back than the '40s. Sorry about that. There have been major revolutions in the business -- the transistor, the integrated circuit, and the microprocessor. We're fortunate to have Bob Noyce here who was personally responsible for the middle one and in his company he nurtured the microprocessor.

If we look forward, we notice that this has happened on about a 12 year cycle. Things happen very fast in this industry. 1984 where we're standing now I think will be remembered as the year of the silicon compiler. There are a number of products appearing on the market this year which allow system designers to directly experiment with architectures on the silicon without the intermediary of the 100 people years to get it translated down into little boxes and rectangles.

As we look into the future, I can only see.... Bob Noyce may have 20/20 foresight. I don't have foresight that far. I can see about another 12 years. I think 12 years from now we'll have real collective computation systems, those modeled in some sense after biological systems, that use this peculiar mixture of analogue and digital processing that you find in biological systems. I think we'll find some real applications running 12 years from now.

Thank you.

(Applause)

MR. EVERETT: Thank you, Dr. Mead. This presentation is open for questions or discussion from the audience. If you'll raise your hand, I'll try to call on as many of you as possible. Who would like to ask the first question of either one of our speakers this morning on this topic?

MR. OSCAR GARCIA: You have both indicated that we should perhaps look at biological systems as the model to follow in investigating both artificial intelligence, vision, and many of our programs. I submit to you that there is so little knowledge about biological systems that just as the rat cannot wait for the computation of the facts to take place to run away from the cat, we cannot wait for the biologists to unravel the way in which we think to be able to make progress in artificial intelligence.

That's a personal opinion. I'd like to hear your counterarguments on that.

MR. EVERETT: Did you all hear the question? I'll ask first perhaps Dr. Noyce to comment on it. And if he could repeat the essence of the question that he's going to answer, it would be helpful.

(Laughter)

DR. NOYCE: I think the essence of the question was that since so little is known about biological systems that if we are going to wait until biological systems are understood, it will be a long time before we make any progress on artificial intelligence.

Artificial intelligence is really posing problems that can't be solved by what we call computers. And as soon as we can solve them, then they're just another computer problem and they really aren't artificial intelligence. So perhaps under that definition, we'll never make any progress whatsoever in artificial intelligence. It will just be another computer program once we've gotten the job done.

I suggest that there's certainly going to be many, many things that can be done within today's technology and today's approaches that are going to be very, very useful to our society. No question about that. The issue I think is whether there will be a revolutionary approach sometime in the future. What I'm suggesting is that we're going to run out of steam on the approach that we're taking now and we're going to have to have that revolution in approach or we're not going to make substantial progress in the future.

MR. EVERETT: Dr. Mead.

DR. MEAD: I would say there's two sides to how much is known about biology. When you start asking very specific questions about very specific things like the retina, for example, you'll find there's an enormous amount known. And you start looking at it, and most of it doesn't help very much in trying to synthesize an engineering solution.

Fortunately we have good vision people at Cal Tech and I have a captive biology student that I can tap for access to that literature. And after, you know, going through several feet of literature on this subject, you can get answers to important questions about the functions of some of the cells. My own experience has been everytime we find one of those....you know, I had started out thinking, what these funny nerve systems really are is a concession to the lousy hardware they have to work with and we can do a lot better. Let me tell you, I have developed an enormous respect for the engineering that was put in the vision system of mammals.

(Laughter)

...however that came about. Once you get a hint as to what something's doing there, you very often say, "Oh, I never would have thought of that, but that's a good idea." And that has happened repeatedly on our little project, and I'm sure we're going to find more and more of that.

We have also found that we're asking a lot of questions that get the biologists interested in going and poking more electrodes into more places. So there really is going to be a synergy there I think. No question about that.

MR. EVERETT: Yes, sir.

MR. STEPHEN KAHNE: You've been talking about what probably we should call a revolution in technology. Should there also be a revolution in engineering education to prepare people for this revolution in technology?

MR. EVERETT: Dr. Noyce.

DR. NOYCE: The question is whether there should be a major change in our educational systems to accomplish the revolution in technology that we're foreseeing here.

I guess my feeling is that our educational system is always organized around solved problems and that the really interesting part of new discovery is where the Teutonic plates are grinding together where various fields are intersecting. So yes, I will say there will be a change. I can't predict where that would be, but I think we will find that the major advances come from the cross-pollinization as I suggested, whether it be between biology and engineering or any other two fields.

MR. EVERETT: Dr. Mead.

DR. MEAD: There are several aspects of that. As you know, we university types tend to divide education into undergraduate education and graduate education. Graduate education is really sharing with your students the ongoing research program, and as such they get tangled up in these cross-pollinization things. In fact, they're the carriers of the pollen.

(Laughter)

Rubbing students against each other is the really effective way of getting knowledge across between fields.

Undergraduate education is a much bigger problem. We have I think at all universities much too rigidly defined fields for this thing to be easy to do. Much more of the interdisciplinary stuff needs to happen, and we're all struggling with how to do it.

MR. EVERETT: Yes, sir.

MR. GEORGE H. HEILMEIER: The IC industry has pretty much been built on scaling things down, and there's a lot of analysis that indicates that the fundamental limits of MOS technology will occur somewhere around 2/10 of a micron. So projecting from where we are today, the rate of progress suggests that sometime in the early '90s perhaps we reach a fundamental limit on what's been driving the IC industry for the last two decades or more.

I guess I'd be interested in both your comments on whether you any new device structures on the horizen that are going to enable the IC industry to continue this progress below 2/10 of a micron which is primarily the limit on MOS technology.

DR. NOYCE: Did you get that question? The question generally was directed toward the limits in MOS devices being in geometries of the order of 2/10 of a micron and are there any new devices which might be on the orizen which could displace MOS devices then in the future.

My reaction is, no, I don't know of any. I think that some of the things that have been suggested can be eliminated as contenders. That is I guess why I think that other approaches are going to become important because that 2/10 of a micron limit on the MOS devices is not very far from these other fundamental limits that we're running into anyway. So even if you could devise something else, you wouldn't have very far to go beyond what that 2/10 micron MOS device could do.

And then there's the question of whether there's enough motivation to do that. With the literally tens of billions of dollars that have gone in....will have gone into the developing the technology for a silicon device at that time, is it worthwhile to make a similar investment in another technology for marginal benefit? I think the answer is no. So it's very dangerous to say that there won't be fantastic new inventions here, but I don't see any yet. I wish I did. I'd invest in it.

MR. EVERETT: We'll take one more question.

Yes, sir.

MR. GARCIA: I haven't heard anybody say anything about integrated optical devices. Isn't that on the horizon too? I consider that a major breakthrough for communications so you don't have to go back and forth between electrical signals and optical signals.

DR. MEAD: Integrated optical devices exist and have for ten years or ten years or so, as you know, and they're a major component of the fiber communications system of course. In terms of the bulk of computation, I would merely point out that a light wave is a lot bigger as an electron so they'll never be as dense.

MR. GARCIA: But for communications, you still have to make the conversion back. And that's the place where we may have a difference.

MR. EVERETT: Well, ladies and gentlemen, in order to keep us on time, I'm going to have to sadly call this stimulating discussion to a halt. But I know you'll want to join me in thanking these two speakers for a most interesting presentation.

(Applause)