Edited Remarks of Speakers: Towards the Next Generation of Networks: The Last Links of the All-Fiber Networks

Columbia Institute for Tele-Information
Graduate School of Business
Columbia University
809 Uris Hall
New York, NY 10027
(212)854-4222

TELECOMMUNICATIONS AND INFORMATION SEMINAR

"Towards the Next Generation of Networks: The Last Links of the All-fiber Network."

Center for Telecommunications and Information Studies

809 Uris Hall, Columbia University, New York, NY 10027

For those of you who have been here before to our events, this is a bit of a different type of a meeting, and we're very excited about it. Usually we deal more with regulatory, policy, and debates, disputes, discussions, etc. This time we're going to discuss the underlying technology of a new network generation, the integrated broadband network that is coming after the more narrow band ISDN that's already been considerably discussed in its regulatory implications. The broadband networks are emerging technologically and business opportunities are being explored. Its regulatory implications and policy implications are not even yet understood, the networks are not understood. For any intelligent discussion to take place requires first an understanding where the technology is and where it is moving. For us, it is something at the Center, here in the Business School to continue doing. One of our research projects for the future is to understand the business, economic, legal societal implications of these new networks as they emerge. For this particular evening's events, we're very fortunate to have been considerably helped by our friends and colleagues across campus in the engineering school, the Center for Telecommunications Research. We'd like to thank them and the first speaker is indeed a professor from that school, and Center, Paul Prucnal, who we'd like to introduce. Paul is an associate professor at the Engineering School. He has advanced engineering degrees from Columbia, an MS in electrical engineering, a PhD and an M.Phil in optical communications. He's been at Columbia as a student since

1974, and as a professor since 1979. He's also had consulting experience with AT&T, GTE, Phillips, and others. He has many awards and honors, and a list of publications that is too long to repeat here, except to be impressed by it. And so, we're very glad to have you, Paul. Paul is under several instructions... one is to try to speak in a way that does not require advanced degrees in optical engineering to follow. Paul—

Thank you, Eli. I'm going to speak with the aid of view graphs today. If you can't hear me, or it sounds weak, just raise your hand, and I'll try to speak up. I have a rather difficult task tonight, and that is, in fifteen or twenty minutes, introducting you to the subject of photonic switching, and to attempt to identify some of the key issues in photonic switching, and not only that, but to somehow compare it in terms of feasibility and economics to the state of electronic switching, to give you some idea where the future might lie. I don't even hope to accomplish that in fifteen minutes. No, twenty minutes. OK, in that case I will also discuss optical computing today. Can you read this from the back row? I gave a short course for the Engineering Research Center on telecommunications last week at Arden House. I guess a lot of the business audience is familiar with Arden House, the Columbia conference facility. I attempted to introduce photonic switching at an elementary level. It was condensed into six hours. And Tingye(?) and I went to a conference last week on photonic switching which was three days. It generated a lot of excitement, and for me more confusion than answers, in some ways. So, as I said, it's a difficult task. I would like to introduce to you first what switching is, talk a little bit about what can be done with electronic switching, then about some of the virtues, as I see them, of photonic switching. The future is by no means rosey for photonic switching, but maybe we can identify what its limitations are; what it promises as well. Basically, photonic switching somehow means the switching of light and that can be interpreted in many ways. One has typically light beams travelling in optical fibers in telecommunciations systems, and these light beams carry extraordinary amounts of information, usually digital. This information is destined to travel at various places, and arriving at switching points needs to be re-routed. So, photonic switching is the process of re-routing that signal. Now, obviously, it would be desirable to allow that signal to remain in optical form for the switchin process, and ultimately we would like photonic switching to be switching in which light remains in optical form, and indeed switching in which switching is controlled by other light signals so that one could carry out the processing and control using light as well as the transmission using light. And this, for the future may lead to very high speed switching. For the present the situation isn't quite that. Photonic switching can also be interpreted to mean the conversion of the light signal coming down the fibers to an electrical signal which is then switched by conventional electronic means, then reconverted into an optical signal, and then transmitted wherever it's going to go. So that type of definition is often used as true photonic switching, but really, those of us who are interested in the field don't consider it to be true photonic switching. There are a couple of points that I want to make about conventional electronic switching. Prior to 1980, the field of telecommunications, as you all probably know was really separated into several categories, in which the transmission of information and switching of information could be easily distinguished. With the advent of data communications, in which there are all kinds of signals being transmitted digitally, and digital telephone transmissions, the distinctions between categories of transmission and switching became less distinct; that is one had bits of information in terms of on-off signals that could not only be transmitted, but remain in that form throughout the switching process. So distinction between the transmission and switching has become less clear and this is illustrated in part by this picture which shows a couple of different types of communications networks, one which is primarily a communications system with a number of users that are ultimately multiplexed into switching points in the network, and these signals are now digital signals. In addition to that it's more general purpose network communications, in which the signals are not only in the form of voice, but there may also be video signals, computer signals, as well, in digital form, going in the network. Now, again, the distinction between networks on this level, and networks on this level, that is telecommunciations and data networks or local area networks is not that .... Basically, these two fields are merging so that we now have networks that carry voice, video, and data in digital form and these can be local area networks in one room carrying information between computers or in a building or in a campus, or in the metropolitan area, and the geographic size can almost be arbitrary. So, what we're interested in then, is in primarily switching the telephone network, at central nodes, as well as the switching that might go on in central points in a data communications network.

Now just a bit more introduction on this subject. With that we can really identify two types of digital switching networks.. one which involves a physical switch, in which there are many incoming lines or trunks in the tel-communications system where a dedicated circuit, a switch has been thrown to make a connection from input to output.. generally called a circuit switch. This might be used for other purposes as well. The other kind really is the network application in which users are sharing the channel, in which there are many terminals, perhaps at very low rates, in which the switching does not need to be instantaneous, that is, data or packets can be queued for transmission, and there can be a delay, and sometimes packets can be lost. In this case, one does not need a switch per say, in the sense that we are discussion today. So the two types of switching are really for packet switching, and circuit switching, and just as telecommunications is evolving toward integrated networks, carrying voice, video and data, the distinction between these two types of switching, in which one is useful and the other is useful is not as clear either.

So I would like to talk about digital switching, and I won't really have time to go into all of the categories of switches, but those of you in telecommunications will know about time division switching and space division switching, and historically, space division switching has been used in the old days to make analog connections in a telephone channel, more recently switches are digital time division switches, in which essentially information is interleaved in time; many telephone conversations are interleaved in effect, and in order to change the destination one does a shuffling of that order; so that's what we call in rough terms, time division switching. Now an example of what I believe is a state of the art electronic switch, this is a space division switch, and what you can imagine here is that there are eight electronic inputs to the switch, here are digital signals, and eight electronic outputs, the digital switch is silicon right here., and is called a cross fire switch, because there are, in this case, eight by eight, and there are 64 cross points. At any input one is able to choose the appropriate output switching point, and switch it that place. This, as I said, is pretty much the state of the art in electronic switching, which is something on the order of the speed of tens of megabits or hundreds of megabits. This particular switch, I believe, was made at Rockwell. I believe it's either Rockwell or Honeywell, and I heard about this about three weeks ago at a ..Darpa(??).. meeting in Washington. So it's this kind of switch that has cascaded to form larger switches in telephone networks, and they're very useful and I think quite reliable, and very satisfactory for high bit-rate digital signals. Now, these switches can be configured in a variety of forms, and the simple one that I showed you is called the cross fire switch and you see that if you have eight inputs and eight outputs, there's sixty-four cross points. What happens if you have ten thousand inputs and ten thousand outputs? There are a hundred million cross points, ok? So very quickly it becomes not feasible to build such a switch. So the alternative, which is actually used in telecommunications systems is to use stages of these switch which are cascaded, are interconnected in a complicated way that allows the reduction of complexity of the switch, but also a corresponding reduction in performance, that is occasionally, when you pick up the phone you may get a busy signal because you just can't make a connection from input to output. But even if that's not the case, the connection of these switches is more complex, but the number of cross points is less. OK without going into that any further, that is more or less, what electronic switching is. I haven't done it justice by any means. An example of digital switching system is the AT&T 4ESS switch. This is a switch which has about 110,000 input trunks, and 110,000 output trunks corresponding to voice service. So with the state of technology the ability to make these kinds of switching nodes is very high. One can awke them in silicon, or high speed gallium arsenide or other technologies, and I think the future is very promising. Now the question is, why the concern in the photonic switching? Well, one reason is that fiber optics have proliferated in the last five or six years, and we have a lot of optical signals to carry the networks, and it seems that if we can use light to carry signals, we should also be able to use it to switch signals with the current devices that we have. This is a problem. Because from a physics point of view it's almost a contradiction, and Tirgy may disagree with me, or he may have something else to say about this, but photons are a kind of physical enitity that do not interact with one another, and therefore they are very good at travelling along optical fibers and going long distances without interferences, and without being influenced by other kinds of signals, essentially. In that sense, it's also very difficult to remove them from the fibers without physically taking sume of the signal that's travelling in the fibers. Electrons, on the other hand, are another kind of physical entity, which interact very strongly; that is as you know, like charges repel and unlike charges attract. Therefore, one can use electrons to control other electrons, very naturally, and that's not the case with photons which are basically the unit of light energy, if you will. the packet of light energy. So, it's not all clear from a fundamental point of view why one might want to do photonic switching, because it's almost a contradiction, but there are many practical reasons for wanting to do it, and as I mentioned at the beginning, we can eliminate the conversion from optical signals to electronic signals for switching in a fiber optic network. Also these switches have the potential for very high speed. Although the interaction between physical effect from light is generally not very strong, so it may be hard to build these circuits. And I think that's primarily the reason, ...there are two reasons that we're looking toward in the future for the use of light for switching and that is that light is inherently parallel, in the sense that the visual information you receive for example through your eye is an image that's coming to you in two dimension, you're seeing a two-dimensional image coming to you, and with lenses and things, one can process light that is in two dimensions, that is in parallel. And you may have heard about the super computers that are being worked on these days in which instead of processing that goes on step by step, many things are going on in parallel. So it seems that the natural, quote-unquote, form that light takes is in a parallel. And we could do switching and processing in parallel as well. This has not really been achieved, although very little significant steps have been taken towards doing this. In addition to that, handling of light is generally not in two dimensions on a chip, but in three dimensions between lenses and other kinds of optical elements, and so this leads to a more general type of processing in switching than one could do simply on a chip where light is travelling around on a surface. So, as

I said, these advantages have yet to be exploited, but let me at least give you a flavor for what is being done with photonic switching in the few minutes that I have.

As I said, what is presently being used for photonic switching is really not what we call a photonic switch. It is an electronic switch that has an optical to electronic conversion, that is, a photodetector, and then switching going on here, and then electrical to optical conversion, that is LASERS, at its output. And, really, this is very nice, because everything is doing what it does best. Switching is done best by electronics in the causes that I described and we have light sources and light detectors that have been refined to a great degree for fiber optic point to point communications. This is an example by the same company that I described before, of an eight by eight silicon switch with the necessary optical converters, optical to electronic converters on it. So you may hear about this kind of device as a photonic switch. And really, for the near term, this is what I recommend for use in switching systems, because it's a reliable device, it can be built, and I think that it can be used in more complex systems than the one that I showed you. Really, what we're interested in is switches in which the ingoing and outgoing light signals remains in optical form, never converted. And basically the idea is that one puts light through some kind of a material that is transparent to the light. It might be a crystal of some kind, and as the light traverses that material, one applies some kind of a physical effect to that material that changes structure, and the structure in that material moves the light beam, directs the light beam in a different direction. Here are some of the possible effects that could be used. One could have a thermal optical effect which changes the temperature of the material, and that distorts it in such a way as to route the light beam. Thermal things tend to be slow, however. Now magneto optics is another effect. One can apply a magnetic field to a material that is susceptible to that magnetic field and that can accomplish diffraction of the light, and diffraction of the light is really deflecting it according to wavelength. I don't know if I can express it a better way. But diffraction can be accomplished by the magneto optic effect, and also by the acoustal optic effect, meaning, an acoustic field is applied across a crystal. It distorts the crystal in such a way as to set up a series of reflectors in the form of a diffraction grating, and that can deflect the light in the desired direction, from input to output. Something a bit more promising in these, right now, is the electro optic effect. As the name implies, one applies voltage to a crystal and that voltage results in a change in the optical pathlink in that material, it's not very important what that is, but again, it's the electric field producing a physical effect in the material that can indirectly produces an effect on the light beam. OK. I would say two of the most important effects for photonic switching right now are magneto optic and electro optic. There's one conspicuous category left off in this list, and this is really present technology, and ther is opto optic effects, and perhaps Tang will say more about this, but opto optic effects is one can use light to change the material and cause a redirecting of the light beam. I hope the questions that you ask will help you to clarify this because I'm not sure how much of this you are following. This is an example of a popular device which is electro optics and I'm going to talk a little bit about the research at Columbia in which we use these devices. Basically, this is called an integrated optic device. Many of you may have heard of about the field of integrated optics. It is the beginnings of fabrication of circuitry, optical circuitry on a chip, by the same processing techniques as electronic circuitry is produced on a chip. Basically what you have is rather than an semiconductor on the surface, you have a crystal of some kind and one makes waveguides, that is paths that light can follow in this crystal material, and one can then either just have the light come in and go out, or one can produce changes on the light; one could do processing of various kinds. One of the simplest processing steps that you might want to take, signal processing steps, is to have two waveguides that are close together, and one can have light coming in here, and either go out this way or go out this way, depending on the voltage that's applied to these materials. See, this is a very common form of what is called a photonic switch these days. It's an electro optic waveguide switch. These kinds of devices are still rather expensive to make but they switch relatively easily in the range of gigahertz switching rates. Now when you look at a photonic switch there are two bandwidth parameters that you're going to be interested in. First of all, is the switching speed, that is how fast can you direct the light to either this output or this output. OK that's the switching speed. Now the other bandwidth that you might be concerned with concerns transmission capacity of this channel, and essentially one does not reduce the bandwidth in fiber obtic systems by putting a signal through this waveguide. That is, if one has an optical fiber carrying the signal in and an optical fiber carrying the signal out, then essentially the light remains in optical form and its bandwidth is not limited. Now that's not strictly true, but this channel, suffice it to say, has a very high transmission bandwidth. That's one of the advantages of this type of switch. One could have very high speed video signals or analog signals which remain in that form and take advantage of the high bandwidth of the optical channel. Is that clear? That's one distinction.. the high bandwidth of the optical channel. The other is the potentially high switching speed of this type of device.

Q. bandwidth in terms of the frequency or the ... of the light?

Frequencies... in terms of the useful communications bandwidth. OK that's one category of photonic switch. Another, I just picked up a sampling of some things that might be interesting, that you might hear about in the future.. is something called a spatial light modulator. Can't do justice to this of course, but here you see four input fibers carrying different signal, different information channels, and four output fibers. And I don't even know if I can explain easily how this works. If you see, the input fiber, let's just take this one, is split into four paths. And each of these little squares can be made transparent or opaque, by one of the effects I described before. In very rough terms, this isn't strictly accurate. But, by making it transparent or opaque one can choose the output channel that it goes through. The input channel is in vertical dimension, the output is in horizontal dimension, and say for example, I make the bottom three opaque, and the top one transparent. That means that the light from this is going to get blocked in the bottom from passing the top one, which takes us to the first output. Well, all possible combinations are possible. Without straining your eyes, just take my word for the fact that you can switch any input channel to any output channel by using this kind of switch. Now this switch has certain communications advantages over the ones that I described. You can see that by choosing combinations of these patterns, you can broadcast, say from the input to various output channels. And one can also do concentrations, that is, input channels can be concentrated into a single output channel. It has the potential for very large number of elements in these spatial arrays, and that's another reason that this is an interesting proposition for photonic switching. Another reason is the fact that we're all thinking about using parallel switching, parallel processing of light in the future in two dimensions, so this can lend itself to two and three dimensional switching architecture.

I was asked to tell you a little bit about the work going on at Columbia in this area, and naturally, I'm biased towards this work and I have to do a little advertizing. I'll just simplify this whole thing because of time, and I'll just say that at Columbia, we're interested in photonic switching, but we're more interested in the control of the switch, that is, in all of the switches that I've drawn pictures of, I haven't said who or what or how we determine the state of the switch at any given time. That is, you have input signals and output signals, but we don't have little elves throwing the switches. We have to have some signals to throw the switches, and those are the electrical, magnetic, or acoustical signals that are driving it. And basically we are working on ways to avoid bottlenecks in that processing to control the switch. We are atempting to design optical processing techniques to control the switch as well, and without going into a lot of detail, I think I can just show you a block diagram. Here's an incoming optical fiber and here's a switch. There's a black box here. We're replacing the conventional electronic element that decides how this switch is controlled by an optical processing element and we've demonstrated such processing techniques. I will spare you the details of the results of that. The work here has been involved in encoding the transmitted signal in such a way that it carries its own address in optical form and the optical form lends itself to the kinds of optical processing you can do. This kind of idea extends to the cross fire switch that I described earlier and this schematically shows a cross fire switch which has eight inputs and eight outputs and sixty four cross points, like the one I showed you earlier. And what these are, optical wave guides that carry these signals in, and each of these little boxes represents a switching point. When we apply a voltage we can decide whether the signal goes this way or this way, and its just traversing a wave guide like I showed you before. So we can extend our optical processing arguments to tapping off the input signals doing optical processing thereby controlling the switch. And the reason that I showed you that is that we're working with one of these switches, and it's made in an electro optic crystal, that is a crystal that could be controlled by a voltage, and we have one on loan from Sweden, from Ericsson Corporation, and this is an awful picture of it, but it gets some points across. The actual photonic switch is a little flat piece of crystal that looks like a microscope slide. It's here. It's about this size; a few centimeters long by a centimeter wide. It has those sixty four cross points. What we do is bring the eight fibers from this end and we have eight fibers coming out that end. Then you have to control it, and that's what this whole mess is. OK? The mess is sixty four fibers, in fact each cross point requires two voltages, so you need, in fact, I guess, 128 voltages. You have all these wires coming in, and signals are concentrated down to the switch. So this is a photograph of the device that we have set up in our lab. But I don't want to make it look pessimistic because this is an extraordinary device. It's truly a triumph in technology in making optical wave guides on a chip and it will have capabilities of switching at, can you believe, hundreds of megahertz, each of the elements that we've seen have a switching speed of gigahertz or more, so we thing that the optical device itself has very promising switching speed and it also clearly will be able to transmit at very high bandwidth signals. So that's one of the future technologies, but it's not clear to me how far it will go, and I think we can have a discussion on that. I don't want to stick my neck out too far right now.

There's one other device which, when we talk about photonic switching, this is really the only true photonic switch. It's one which an optical signal's controlled by another optical signal. When you think of a signal controlling another signal, what's the kind of electronic device you think of? It's an electronic transistor. You have a three terminal device, input, output and control. And, oh, I don't know how many years ago, sometime in the last ten years or so, something analagous to an electronic transistor was developed called an optically bi-stable device. It's really like an optical transistor. And this just shows the characteristic of the device. This dimension is light in, and this dimension in light out, and what we do (we don't do this in my lab) is have an optical signal that's called the bias level; the bias level is an existing level of light that is right here, and basically, when the logical pulse comes in, it pushes this input signal over the threshold that turns the device from this output state (this is the output light level) to this output state. And it sits there. Because the bias level is here, it sits at this level. And in order to bring it back down to low state, one has to take the control signal and remove it so that it comes back down here. So this is an optical device that is bi-stable, that is, at any input level, at least any input level in this range, it is capable of having two states, bi-stable; the stable in the low state or the high state. It's also an optical storage element. Now it seems that this kind of device will be very useful for photonic switching. There's work going on at Bell Labs which Tingy will tell you about, perhaps, in which one can use optical beams to control other optical beams using this bi-stable device. We're working at Columbia on a hybrid version of this which involves a little bit of electronics but it's possible for us to build in our laboratory and learn from. So, I guess I should wrap this up at this point and go to the next speaker. Basically, I think what I want you to understand is that electronic switching is probably here to stay; it has extraordinary capability. It seems like optical switching or photonic switching will be very useful in certain applications, certainly applications where you want very very high band width signals going through a switch. Certainly in applications where you want the switch to change at extremely high speeds. That is not necessarily always the case in every application but we think that the future is promising, and I think that most of us are interested in the future going toward devices that can use light to switch to switch other light and devices which can carry optical signals in parallel form and switch them and process them in parallel form as well. So again, I don't think that I've done justice to the field at this point, maybe I've stimulated your interest enough; certainly I'll be happy to answer questions that you might have at this point. So I'll turn it over to Tingy at this point. Thank you.

Thank you very much. I will introduce the next speaker. I would hope, although it's not exactly been your assignment, but perhaps if you could also enlighten us what does exist in various laboratories, Ericsson was mentioned here with various other telecommunications or component manufacturers where they are at this point in the development of operationalization of this technology. Well, let me introduce Tingye Li who comes certainly with distinguished credentials, with engineering degrees from Universities that run from Johannesberg, South Africa, Masters and PhD degrees at Northwestern. He has been at Bell Labs since 1957 where he has been involved in research in microwave, lasers, optical communications, with 48 pamphlets and 60 papers currently he is head of the Light Waves Systems Research Department where he is in charge of research and optical fiber communications systems. He was head of the Light Wave Medix Research Department and other parts of the organization dealing with this subject at hand. He is a fellow of the Optical Society of America, and various other professional associations in optics and engineering and electrical engineering in general. He was guest editor of a special issue of the IEEE journal of quantum electronics, of devices for optical fiber communications. Well, time does not permit me to list all of his accomplishments. He is elected to the National Academy of Engineering. He has many other honors which he received for his work and we're very honored for your attendance here.

Thank you, Professor Noam. I've been asked to speak on the integration of broad band technology into the network. By that, I interpret it to mean the application of light wave or photonic technology for broadband services in the future network. And therefore, I would like to talk about, first, our present network. The network of the future must evolve, presumably in an orderly fashion, from the present network, especially if the lawyers don't do too much too it. I'm sure there are some lawyers in the audience; I have great respect for lawyers, but I do believe they create chaos as well as order. Now, the second thing I'd like to talk about is the research work that is presently going on in preparation for the introduction of the future wide band services and thereby requiring very high speed transmission rates, and also special switches, be it both photonic or electronic. So, let us now first look at the motivation; this is just a little bit of fun.. What drives us on to do the high speed work.. the broadband technology. And it's the broadband capability of singemod fibers. This is loss, or attenuation signal, attenuation loss of the fiber versus the wave length. And you can see that the loss of the singemod fiber is extremely low in the wavelength region around 1.5 microns, and in fact, if we look at the loss near .2 dB per kilometer, which is something like 5% of loss per kilometer, or light loses at the rate of, in fact, 2 every 10 miles. Now, for the windows, a very clean window pane will lose light at the rate of 2 for every 10 feet or so, and this is for 10 miles! Now, if we can center the band associated here we see that for the attenuation to be low, something about 25,000 jigaherz, towards 25,000 billion cycles. Now, what does that mean? It's such a staggering amount of bandwidth. So, let's just have some fun and look at the transmission window in this range of wavelength. We say, well 25,000 jigaherz enables you to transmit 25 times 12, or 75,000 jigabits of information. Let's see, what does that mean. So we can look and see what the total accumulated human knowledge residing in books not pictures, words. We can go to the Library of Congress, possibly the biggest depository of books in the world, and count up the number of books, and we look in each book and count up the number of bits in them, do a little bit of multiplication, we see that an average 500 page book contains 10x10 power 6 bits, and the Library of Congress contains something like 15x10 power 6 volumes.

So we have a capacity of that much. So we multiply this by this. We get 150 times 10-power 12. We can throw a few Chinese books in, and get up to 500, perhaps. Divide that by this, and you can see that in 20 seconds we can transmit the whole accumulated human knowledge residing in books. That is really amazing and astounding, and later on when we talk about bandwidth we have to evaluate this in terms of what the human mind can work with or appreciate; it's really at a pretty slow rate. So let's just keep that in your mind for the time being. And before I talk about high speed digital transmission I'd like to remind you of the digital hierarchies that are in use in the U.S. I'm sure you're familiar with this so called TS0 for single telephone circuits or telebits per second. Multiplexed up to 1.5 megahertz per second, so called TS1 or T1, there's T2, 96 circuits. This is TS3 which can accomodate one TV channel and much of the future lightweight systems will be working at this level. We have now in the field for long distance transmission, as well as for metropolitan area applications, systems that operate 417 megahertz per second. And since three weeks ago, a system operating at 1.7 gigahertz per second, permitting 24,000 telephone conversations to be transmitted. Now, this is the growth. And historically, in 1977, the first experimental system for metropolitan area application was installed in Chicago as a field trial and this one ran at 45 megahertz per second, or 672 telephone conversations. Then the so-called Northeast corridor began operating at twice that bit rate, 90 megahertz per second, linking

Boston and eventually Richmond. And 1983 the 417 megabit system was in production, and two years ago one was in service that operated at 417 megahertz per second. Today we have the so-called FT Series G system that operated at 1.7 jigahertz per second which was first now installed and in operation three weeks ago at 1.7 jigahertz per second. So much for the various systems. We might want to look at the areas of application. The first one is the local loop. There's a great deal of fibers in the loop now, although none of them go to your home. The area of application here is from the central office to a remote terminal which is perhaps half a kilometer to a couple of kilometers from your home. And this is the so-called subscriber loop feeder area and the system is called SLC, which stands for subscriber loop carrier system. These systems operate at 6 megabits per second, 45 megabits per second, and in the future 90 megabits carrier signal. At the remote terminal multiplexing takes place and the various services are homes and small businesses, while other multiplexing can take place at big business centers. In future, of course, we will see the movement of fibers into the final leg of the subscriber loop. In the metropolitan area, this is for central offices to communicate with central offices, there is a system called metrolux being introduced. The fibers that run from central office to central office run at 146 578 and very soon, at 1.7 jigahertz per second. In a system like this, the signals that run around are synchronous so you can do all sorts of fancy things with all these inputs at 1.5, 3, 6, and 45 megabits; you can multiplex them, you can add and drop chunks; this type of thing, and this enables the Bell operating company to rearrange their networks efficiently and economically. Not only that, in the future it will allow the customers to manage their own network in a way which will serve them best. Now, for the long haul, which is the system that I mentioned, FT series G, operates at either 417 or 1.7 jigaherz per second, the signal format is binary, non-return to zero. The interface of this is at 15 megabits, and the wavelength of operation is 1.7 microns of fiber with least dispersion. The feeder spacing is somewhat between 28 and 50 kilometers. Since this view graph was made, we've had improvement. So if you remember, some years back there was a coaxial system running between Newark and New York, running at the speed of 300 megabits per second. For that coaxial system the repeater spacing was one mile. Now, we are seeing repeater spacing between 25 and 30 miles, and running at much higher bit rates. At the moment the country is criss crossed with various microwave systems, except for this part. Everything is here. So we can call from Boston to Miami to San Francisco by fibers. Not only GTE line, AT&T. Now, in the very near future, namely, next year, an undersea cable system will be in operation, and this system has been developed and installed, and it operates at 200 megaherz per second at a wavelength of 1.3 microns. It's an international effort, apart from Tukerton N.J. to a branching point. It is supplied by AT&T at the cost of 250 million dollars. 5,600 kilometers. This little branch is being built by STC in

England and this one by Samarcom in France. Insurants are England, France, in two year warrantes and we have given a ten year warrentee. I'm not too sure what will happen after two years, but you can see the international politics of the whole thing. Pretty much at the same time, the Pacific system will be installed. California to Hawaii to a branch in Japan. Here the Japanese KDD will supply the branching here and then this leg will be operating at the same time. Now let's look into the future. See what will really bring the fiber to the homes and the broadband services. Perhaps in, I don't know, three, four five years from now, there'll be fibers going from the remote terminals in the homes and this fiber will be a singomol fiber carrying perhaps 600 megabits per second, to a jigahertz per second. Several TV channels and many voice channels, data, as you wish, and the trunks and feeders will also have to go to higher bit rates. What will they be specifically? This is the singomol fiber in the home, optical to electrical conversion takes place, and demultiplexed down to the various services. This is 64 kilobits per line, run at a few megabits if you wish. Right now the data channels run at a few tens of kilobits per second at most. The video will be running at 45 megabits at first, but in future when high definition TV comes in, it might be 150 megabits per second. At the central office each one will be multiplexed out and the TV will be sent to a video switch and the rest might go to a much lower bandwitch. And the thing is this: at present you could implement this with the cheapest possible electronic technology known to us now, the same type of electronic technology that you build your electronic watches with now. This technology can switch, as Paul has shown, 150 megabits, way off into the future. So the question is, if you want to build photonic switch you have to compete with the electronic switch. Here we're trying to provide a service, and it will be a matter of which one of the technologies, photonic or electronic, will provide this type of switching services most economically. Now, as far as switching data, things like this can be packet switches. Again, as long as you have your systems architecture this way, it looks like electronic implementation, at least at this time, is the way to go, and perhaps, this way to go for a long time to come. The types of things that Paul had talked about I regard as way off into the future and it's going to come, and I'm glad that he's doing it, and we're also doing it. In the future, perhaps we'll change our signalling format and network configuration such that much of the traffic is carried by packets going at very high bit rates. When that day comes we will need very high speed optical interaction switches like the one that he has shown. And that also has to be competitive with any of this technology on a cost basis. So that's my little remark about switching. Now, back to high speed transmission, preparing the network to route high speed data. Now we are working on photodetectors, lasers, and electronics on a device level as well as sub-systems level to enable us to transmit and receive at the highest bit rate possible. We're trying to push the frontiers of technology toward higher and higher speeds. Once we have the devices we'd like to test them in a system. We have a pseudo-random word bit screen generator at very high speed driving the laser and we put a fiber in there. At the other end we have an avalanche photo diode which converts the photons to electrons, and we filter them, and we can see whether it has made any errors or not. So we would have done these experiments at various bit rates over the past four five or six years; at first 420 megabits, 2 jigabits, and this is a part of the bit error rate versus the received power. It's a thing that for us to evaluate whether we have achieved the goal of 10-power minus 9, or 10-power minus 10 error probability or not, and you can see for 4 jigabits per second modulating the laser directly, turning the laser on and off, we were able to send this signal over a distance of over 103 kilometers, yet achieving our goal of 10-power minus 9 error rate. Clearly the sensitivity, sensitivity means that if you are able to couple, say one milliwatt into a fiber, then one hundredth of one milliwatt or something like 500 microwatts or 500 nanowatts is about what you need to enable you to transmit 103 kilometers. Later on we were able to push on to higher states than this, that no one has done. Only recently the Japanese were able to accomplish 4 jigabits per second, whereas we have done that a year ago, transmitting 5 jigabits over 70 kilometers progress. Discussing error rates, it's due to the random noise in the receiver. We just lower the power down until the thermal noise in the receiver began to hurt you and that's the error. Now you can go to higher and higher bit rates by going to a thing called wave length multiplexing; that is, you use lasers of different wave lengths or colors, you have three lasers and you can multiplex them, combining them up one notch. We have done an experiment like that three years ago at Bell Labs; ten lasers is multiplexed by means of a grating of a single fiber over 70 kilometers, separated out ten different beams, one of which was picked out and looked at. And each one of these operated at 2 gigahertz per second. Now very recently there is a great deal of interest in doing coherent microwave systems. That is, operating the lightwave system very much like a radio or TV model. So far, what we do is we have a detector. We just convert the light pulses directly to an electrical pulse. This is called electrical detection. But a more sophisticated way of doing this is to have a local oscillator or laser at the receiver. This wave is combined with incoming waves to generate another wave which has a frequency in the microwave range. Then you can apply all the microwave technology this way to filter it, to amplify it, to do all sorts of processing, so we're at the state of the devices such that we're able to do that. You have a laser, modulated either in frequency or phase, and after the fiber we have another laser here which combines with it and generates a different frequency in the microwave range. Doing that, what are the advantages? We can improve the receiver's sensitivity by something like a factor of ten to a hundred and best of all, increase the selectivity, that is, we can pack a bunch of channels, a thousand times closer than you could do otherwise by direct detection. We've done experiments with them and here's sort of a summary which says what the situation is throughout the world. It's sort of interesting as you can see. This is the bit rate, up to 1 gigabit per second, and you can see that the message is received globally. This is a measure of the sensitivity of the receivers. This is how well you're doing technically. The smaller the number the better it is. So far, the record to date is that you only need 45 photons per bit to achieve an error probability of 10-power minus 9. This was done at 400 megabits per second by somebody in my department. And at AT&T they have high powered lasers so you're able to get longer distances but their sensitivity is not as good. And there's that other measure of bit rate times light which gives a measure of overall performance. There's a little competition going on. And this sort of sums up the state of affairs. Again this bit rate times distance problem here: this is the .9 micron technology, 1.3, 1.5, and over here... And it's interesting to note that this increases by a factor of 2 every year and that they're also kind of popping out here because the physical limits of the dispersion of a fiber, the noise of the receiver, and what have you. And this thing is coming up, but it will overtake it a little bit again. And I look at this and I don't know if its true in Columbia or not, but this is certainly reminiscent of the salary curve of people here. Here's an older guy, here's a younger guy. But you know, it sort of merges, and then the younger guy gets old. And the physical limits set in.

(Eli Noam) Thank you very much. We have now two speakers who will take somewhat less time. They're listed as discussants but in fact they're not really discussants; they're giving additonal details. The first one largely on local loops, and we've already seen some of these graphs on trans-Pacific and trans-Atlantic and transmission and switching. Now the details of the local applications which will be presented by Howard Bruhnkee who is at Teleport Communications, the company that's a joint effort by Merrill Lynch and the Port Authority of New York and New Jersey. His primary responsibilities include building, implementing, managing the teleport on Staten Island, which is the center of the Company, constructing a fiber optic network interconnecting key cities in New Jersey, New York, Brooklyn, and Queens with the teleport, and maintaining 27 antennas at the Staten Island facility. Before being with the teleport, Howard was at NY Telephone and NYNEX for 36 years. He had responsibilities in marketing, tariff implementation, new technologies, and particularly over fiber optics throughout New York and New England companies etc. He, in fact, introduced the fiber optics systems into New York Telephone system. He was in the army, discharged as a master sergeant and he was at the Brooklyn Polytech where he received a degree in Engineering.

(H.B) I don't have slides or anything to present to you, but I'd just like to tell you where we think the technology is going, and that's what I believe we were called on to make the presentation on. You saw all kind of things of the future, and I have to tell you, that's not the future, it's today. A lot of that is in application today, and I'll try to describe some of the things we're doing. The main subject of all these things is digital technology and bandwidth. Let's go back to 1964 for the introduction of digital technology, you had T1 circuits. The unreliable part was the ability to provide those because of the wide band circuits. No one would put that density on. That didn't change until 1978 when fiber was introduced. Now, remember one thing, in my position I have to look at two things. Reliability and cost. If you're going to be in business, that's a major consideration. Otherwise, you won't be in business too long. With the introduction of fiber you brought 45 megabit, and it accelerated rapidly until 1982 when single mode came in and single mode now opened up a new horizon of band width. Now all these technologies evolved around improvements of fiber, laser technology, and the application of these systems. The thing that has lagged is the development of the end product that the customer uses. The only reason for that is you can't develop a product unless it's economical. So we're back to the figure of cost. The use is out there. There are computers that would love to be connected together, but the cost of those connections has been prohibitive. With the introduction of the fiber, the lower cost of the technology now has brought it down to a point where we put in a mile of fiber and the cost for a T1 circuit to me is ten cents. So therefore, it is now distance insensitive. If we're going to charge those kinds of rates of somewhere between 500 and a thousand dollars for a T1 circuit with excellent reliability, I guarantee 10 power minus 9 on all of our systems, we take them out of service and switch, at 10-power minus 8 and I'll buy anyone lunch who can pick a circuit I have that doesn't run at 10-power minus 11 or better, or I'll take a dollar from you, and that's a pretty fair exchange. So the reliability of these circuits is excellent, and they're here today. We have about 200 miles of fiber cable around and through New York City ranging down to Princeton N.J. And the reason we're down there, is this project, Teleport, was founded by Merrill Lynch, to fulfill communications systems that would provide the bandwidth and lower their cost. On the same token, we built a satellite farm on Staten Island, in conjunction with the Port Authority and New York City, where we can up- and down-link signals and bring them back into the city. Now, those systems were originally defined at one-and-a-half megabit levels. It is only in the past year that we're now expanding those. And to give you an idea of the requirements, within the last two months, we now have 11 circuits going in at 10 megabits, very common ethernet levels. There is a requirement in the market that various local area networks have to be interconnected. Can we do them? Yes, we're doing them today at very costly rates because the technology isn't there to merge the ten megabits into the common denominator of 45 megabit signal. And incidentally, I want every network system at 560 megabits, and we will trial 1.2 gigabits very shortly, multiplexed. We could put wave division multiplexing in tomorrow and run it at 1.2, but that's not what our objective is. We don't need that much yet. The other portion of it is the application of the high bandwidth to interconnect the customers into the various carriers in the telephone network. 1.5 will be the minimum penetration into a digital switch network. It'll be a matter then of how the customer controls his network, switches it, internal to his carrier, to get the best usage of that. An example of that right now is that we run 1.5 megabit out to Staten Island to the base antenna of Comsat, where Comsat switches those lines between London, Zurich, and out to the West Coast. So we use the same common circuit as the time domain changes. With Zorich being ten hours behind the west coast the data lines for the financial market can now be utilized as that time zone changes, and we do that via the switching on the ground in Staten Island to various different satellites, down links, around the country. The other piece that we just brought in, in the Times, about a week and a half or so ago, we just did a digital dubbing and reporting utilizing Stevie Wonder and Quincy Jones in California who dubbed the record over Niles Rogers and a choir group in New York City in the Astoria studios. And only through Digital transmission, could we achieve the laser transmission that was necessary in the high quality of the audio. And that went off excellently. The quality was incredible.

Since that, in the past week and a half, we've done three more programs utilizing this feature. Where an artist is at one city and he will broadcast on the high digital signal, bring it into New York, and then we dub it back onto the tapes in New York. So it now eliminates the travelling of all of these different people to get them all together. And I don't know how long it would take to get four artists like that together, but it certainly would be very costly. The other thing that we now propose and offer is the video signals. We have links in New York directly into CBS and NBC, ABC, DSN, CTN, all the different broadcasters, and we deliver these signals from New York City out to Staten Island approximately 60 kilometers. And we deliver it on an alternate route that if one system fails it runs on the other. That's an electronic switching system built into all of our systems. In every system we provide, if there's a failure rate at 10-power minus 6, or below it, I should say, it automatically switches to the protection channel, the internal system does its analysis, and literally prints out the fault in the network, so that we can take corrective action. On the video signals, we're running at 45 megabits from the city over the digital signals, but the latest application that we are now using is 140 megabit system. And that appears to be the quality that's necessary to provide the video signal required by the broadcasters. Now that runs, and again, we have to look at why this wasn't done in the past, and the very reason is the lower cost of the fiber, and the higher bandwidth that's available. Combine the two, and the end product is that the cost of the service comes down low enough that we can now offer it economically. So consequently from our concept, we see bandwidth as being the provider to the business customer, where the customer will control it and select his option of services over this service. Our company, at the present time, has service into approximately 80 buildings in New York City. We serve all the way out through Brooklyn, Queens, all of Manhattan, and we go down through Jersey, down as far as Princeton. And the services are available, and they're open today and we are a tariff company registered with the Public Service Commission, and I guess we're really the only alternative providing high band width capabilities throughout New York City other than New York Telephone. Thank you.

(Eli Noam) Could I just ask you a question concerning the fiber in the local loop? Could you perhaps enlighten us as to the technical and economic obstacles or hurdles in wiring up, so to speak, the last mile of residential or office buildings. When you say that is a kind of expensive proposition, what leads to the expense, what are the elements that make it expensive?

(H.B.) Well, let's take two different areas, which is what you're looking for. In New York City, where you have a high business concentration, the cabling of buildings becomes the most expensive part because you've got the vertical cable to place in the building. The streets of Manhattan are highly congested;

however, with a franchise license you can get conduits through the street. You try not to splice by any of the cable in any of the underground conditions. Consequently we pull all continuous length fibers right from our office to a customer's building. We have pulled better than two kilometers or better than 7 thousand feet of cable in one shot. The only reason we haven't gone further is because there's been no requirement. That is a high cost operation in the city. The cable is pulled at night and access to the street and the buildings is very difficult. Once you come into the building you now have to distribute. Unlike copper wire where you could terminate and make tap offs at any point you can't do that with fiber. So consequently you have to vertically cable the building and provide access points every five floors. Now you get into the situation of the type of cable, fire protection, and all kinds of other things that come into this and drive the cost considerably higher in the buildings. The bandwidth, though, is available, because fiber still becomes the cheapest part of that provisioning of service. So we concentrate the electronics back at our central points and deliver low speed electronics out to the customer. Now, when you get out into the rural area, you get into the reverse condition where linear length becomes your problem of providing the service. Consequently there is the item that Dr. Li said of using the SLC 96 where they're concentrating telephone circuits and to a point utilizing fiber. This is, again, in lieu of the copper facilities, or digital transmission. From that point to the home, the fiber cost would not be that expensive. We use very simple two path fiber cable and we use this in buildings throughout the city that we place through the air plenums and drop off, and our service runs right to the customer's termination. We are fiber, end to end, all the way. And when you get into a home, then, you can deliver the fiber into the home, but the demultiplexing becomes the expensive piece of equipment. Now, there are two thoughts on the technology of the demultiplexing. One is, you would run in a maximum of let's say, three video channels and that's assuming that you don't have more than three kids in a home so that each one has his own TV and you can still watch one of your own; then you get into troubles; you have to make a break off point somewhere. Thoughts are, three video channels would be available from a video switch. Now, a single fiber could deliver all of the video channels at CATV grade to that switch and therefore the local distribution to the home can be via just one pair of fibers to put in the telephone service and three video channels. That cost of that technology has not come down to the point where it is economical today. That's the only thing prohibiting it. Will it? I think inevitably, it will. Also through some of the subject areas that Prof. Frucht talked about with wave division. I belive if you bring in multiple light waves on a fiber, the splitting of those waves may lower that cost in the light wave area rather than in the electronic area. So that's a second thought that could come in on this. It'll be a case of which one is developed first that'll bring fiber directly into the home. For the home user, a copper line, the average home can't exceed the use of that. The average home today has a coaxial cable system and therefore video is supplied with 34 72 channels over the 400 megabit band.

from a central point, it then feeds in by the distribution system. I think that'll be the controlling point, and it'll be a matter of industry over which one will be done. And I'm sure legislature and a few other things will come into it also.

Q. You say you ran all the fibers directly from one office to customer premises?

A. Yes.

Q. There's no tapping off the sort of longer runs uptown?

A. Absolutely not. Now there may be more than one customer on a cable sheath. That's a minor point which I'll talk to you about afterwards.

FEIL Noam? Let's continue this discussion after our last speaker. Our second discussant, Mr. Joseph Nadan, has an interesting career which includes a Ph.D. in Electrical Engineering from the City University of New York, 1969, was an associate professor with tenure at City College from '69 to '74, became assistant dean of the School of Engineering, 1973, this no doubt brought him into despair, and he left to the private sector where things are much quieter and became the scientific consultant to the president of Phillips Laboratories, part of North American Phillips, where he worked on computer aided designs, electronic devices and systems. He became group director for electronic and optical systems research, was in charge of consumer electronic products, digital television, CATV, etc. Now, while doing all other things he also managed to be right here in this building to get an Executive MBA, in 1982, specializing in marketing and consumer behavior, etc. In 1985 he left Phillips and joined Rich, which is a division of Rogers. That company designs, installs, and services financial trading systems, combining video switching and digital local area networks that provide market information for trading rooms on Wall Street and the rest of the world. He's now Senior Vice President for Corporate Planning and development there and his responsiblities include evaluation of technologies.

[illegible]: Thank you. When Eli asked the last question I thought he was giving away my entire talk. But, basically, what I'd like to discuss is the cost aspects of installing fiber in a trading environment. We're not talking about going between cities, we're not talking about going between buildings; we're really talking about going between floors of a building or between several locations on the same floor of that building. A typical trading environment consists of a lot of video switching and a lot of digital switching. Both of these are essential in the trading environment because people who are traders want to look at live video. There are contractual obligations that require the conversion of digital information to video information before it goes into the trading floor so that people can't steal the information, and suffice it to say that the overwhelming majority of trading floors today combine both video and digital information. We are probably the world's largest installer of these trading floors with over 250, and only two of them have anything related to fiber optics on them. Tingye and I were talking before and he gave me a phrase which I will be grateful forever for. These two systems, I am absolutely certain, are fiber optically oriented because of "photonic fever." Thank you Tingye. Why is this so? People keep talking about the band width that fiber can deliver. Band width isn't a beneficial aspect of financial trading systems at the present time. The largest system that we have designed to date involves something like 1500 desks with 5000 monitors on it, with large amounts of digital information and telephone systems going down to each desk. And when you get all through, and you do a traffic analysis of what is necessary to support that, it turns out that you're using less than 2% of the backbone capacity of an Ethernet network. That's remarkably low. And therefore, the need to go up to data rates that are substantially more than that are probably five to ten years away. The second consideration: fiber costs. People talk about the cost of fiber. Each fiber today will roughly cost you twice what an RG59U coax would cost you per foot. Putting on connectors is a whole other story. At the present time in New York, the labor rate that we pay to our union installers of electronic things is $50 an hour. They can, in one hour, put on one fiber optic connector, and this fiber optic connector will have a wide variation in the insertion loss characteristic. In the same amount of time, one hour, they can easily put on five BNC connectors, plus take a coffee break. When you get all through with this the cost of actually putting on the fiber into the physical plant would more than double. In a trading floor that costs forty million dollars total, the cost of coaxial cable and connectorization of that cable is probably typically [illegible] million dollars. So we're talking about, by going over to fiber, adding an additional cost of at least 4 million dollars on that type of facility. Next characteristic that really precludes fiber getting onto the trading floor, or in the equipment room immediately behind it, is that there's a tremendous cost involved in constantly going back and forth between electrons and photons and electrons and photons. Today there is no optical switching that we know of that is viable in that type of environment. People constantly talk about the cost per cross point of electronic switching and I'd like to give a new number, and that is the price per cross point electronic switching, and may my professors remember that I remember the difference. In this type of environment the price per cross point that we can sustain to a client is on the order of ten dollars, which is far and above what the cost is. And even with that price we are not able to economically justify that with fiber. Another characteristic that is important to us with switching is the physical volume of the switch and the amount of heat that it has to dissipate. The footprint that a switch occupies in an equipment room is a very large cost factor to the operator of a trading floor and if we had optical switching this would be a major benefit, but we don't. Further, in this conversion process there's always extra electronics all over the place. Can I go out today and buy a CRT monitor having a fiber input connector? I could almost do that, but you can't go to 47th Street Photo and get one. Lastly, there are some benefits that people talk when they talk about fiber. And these are real benefits. One of them is commonly handied about and that is the security of the information. It's a lot harder to steal an optical signal than it is an electrical signal. And people are very concerned about this, and this does have a certain incremental value and people are willing to pay a minimal amount for that. By far, the major benefit of fiber to us would be that it provides electrical isolation between the subsystems that go into a financial trading room. When you are establishing an equipment room, which has on the order of 700 to 1000 racks of electronics, and you are distributing this over four floors of the World Financial Center, the grounding problems that you have become very very difficult to manage. If we could use fiber in that system in a cost effective way, it would solve many of these grounding problems, and we constantly follow the cost patterns to us for fiber. But in our opinion, at this time, we are very very far away from justifying its use in that type of an environment

[Eli Noam] Thank you very much. We unfortunately are kind of being pressed for time because Faculty House will have dinner for those of you who would like to go. It expects us at eight o'clock at the latest, and it's a five minute walk, so we have I would say, at most three minutes time for questions and answers if such exist. Is there a question from the floor?

Q. Would you say that a major factor of the research is to generate computer applications which aren't really valid for this problem?

A. I think that's absolutely correct. I also have a great deal of experience in high definition television areas; as a matter of fact, one of Paul Druenell's former Ph.D. students used to work for me at Phillips and he was the fellow who developed a lot of this 525 to 1050 line conversion circuitry and algorithms and I just believe that's one of the biggest boondoggles around. I don't believe that you'll, in the next ten to fifteen years, see high definition television. There are many alternate technologies that most consumer electronic companies are working on that utilize far less band width. There are two-channel compatible methods of transmitting television signals that will look subjectively exactly the same as a high definition television signal, and that's really what's keeping that from happening. So I don't see anyone beating down the doors of fiber except for long distance super trunking applications.

Q. Do you think that if you redesign satellites and transponders for high definition television it will be easier to market it?

A. Anybody want to take this from the panel?

E.J.R.: I don't know. I haven't studied that, so I don't want to give an uninformed comment. But I can tell you my intuition is no. I don't see the market for it.

Q. You said ten cents a mile for a t1 circuit for 560 megahertz?

A. Yes.

Q. And that works out to roughly forty dollars a mile when you install a cable -- is that right?

A. We're talking about a high density fiber cable? Isn't it more like 400 or 500?

Q. That sounds awfully low.

A. 4,500!

ED.N.: I hope you will join me thanking our panel for a certainly interesting discussion and review of the technology of the future and the development of these technologies of fiber and switching that will reach us very shortly -- or have reached us as we've already heard. Thank you very much.