

CARTOGRAPHIC DISPLAY REQUIREMENTS

MR. DEAN EDSON: The last panel for today concerns the Cartographic Display Requirements. This is intended to be an overview of what people expect of the cartographic community in terms of a more meaningful presentation of information. We have asked Waldo Tobler, currently at the University of California at Santa Barbara, to gather a group together to discuss this very important subject.

Waldo, as I explained, is currently at the University of California at Santa Barbara. He received his Doctorate in Geography from the University of Washington in Seattle in 1961, and has been very active as a U.S. member of the International Geographical Union, the Commission on Geographical Data Sensing and Processing. This has been an extremely meaningful activity, servicing many spatial data activities throughout the world in the last few years. Waldo notes that he is currently teaching analytic cartography, and is also involved in teaching geographic information systems and regional analysis. He reminds me that he used a Benson-lehner plotter in 1957 to draw a U.S. outline map, a long time ago. So, this whole subject does date back. Waldo's hobbies include but are not limited to the invention of map projections, some of which are useful. (Laughter.) I think that is enough for me. If you are ready, I am. Waldo Tobler.

MR. TOBLER: Thank you very much, Dean. What I have tried to do here is get three people with contrasting backgrounds to show you and discuss some technologies. I will introduce them in more detail a little later. But in the order they will give their presentations, Jim Blinn in the middle, is an expert on computer graphics. The reason for asking Jim is to introduce to the cartographic community some of the techniques that they have been developing and which are not sufficiently familiar to cartographers. In the same vein, Harry Andrews is from the Image Processing Institute in Los Angeles, who will do the same thing for image processing techniques. Finally, Carl Youngmann, a cartographer is going to take the perspective of a cartographer on all of this.

Dean Edson mentioned that I had been in computer cartography for 20 years now. My recollection is that the first outline map on the Benson-lehner plotter consisted of 343 points. Since that time there have been much bigger and better data sets prepared. The history of the technology, however, has been that the first attempts were to mimic draftsmen--that is, get the machine to draw a map. We were not too much worried about data structure. I can foresee as an ultimate objective getting the information into the computer and never having to produce the map at all; solving the problems right in the computer. We are a long ways from that

objective, if it is really a valid one.

I think my first slide demonstrates, in fact, the importance of graphics. I will show just a few slides, and then we will get right into the substance with the speakers. This diagram is a set of profiles, but they are demographic profiles. On the Z axis is the number of people, and on the X axis--I may have the axes confused here--is the age group. Then, along this way is from 1801 to 1947. I find this is a terrific graphic because you can see the effects in the canyons as they progress. You can see the effects of wars and so on, and it is really a dramatic visual impact. It is quite clear that humans are very good at processing visual information. I do not expect that we will ever not want the computer to do illustrations, and to just solve the geographical problem in the computer. People are very good at solving problems.

The next slide is a map. It is Minard's map, again, very dramatic, of Napoleon's march and return into Russia. You see the broad life line starting with many, many troops, and it gets narrower and narrower as he gets to Moscow, and the black line shows us how he gets back, with hardly anybody left. Again, a very dramatic graphic. I think the point I am trying to make is fairly obvious. Can we go on to the next one? This, to cartographers, will be familiar: ways of showing relief. We have on the left a lady's face with contours and shading, and on the right the Crimean peninsula with relief shading.

The next slide shows this in more practical application. This is a Swiss topographical sheet with shaded relief. I think it is perhaps appropriate--it has been done once already today--but, to remind people in the computer field of the tremendous amount of ⁷ information on a topographic sheet. It is estimated at about 10^7 on a single sheet. You can calculate very quickly how many sheets it takes to cover the world surface.

The next slide shows a computer shaded relief. This one was done by Mr. Batson, et al., on the image processing facilities that they have in Flagstaff. I show this particularly because the people in computer-graphics have also been doing shading of objects for some years now. The first person in cartography of course was Pinhas Yoeli, who was doing it in 1961 on a line printer. This shows how far the technology has come. I think Jim Blinn will show some things that show how much further one can go with this in computer-graphics.

The next slide shows the kind of illustration you have seen now, quite frequently, probably; it demonstrates the increasing resolution when one changes the spot size. Some of you, I am sure, are

familiar with the digital terrain tapes--that is, the sampling using the multiplication with the Dirac brush of the continuous function at fairly coarse intervals. This demonstrates the question of resolution, which I think is going to become more important in cartography, and the next few slides illustrate that.

Some of you may know that in the 1850's geographers were collecting data like the number of whales per five degrees square. If you see the movie, Moby Dick, for example, you will see that in there. Here is the world population by five degrees square, which I assembled. The next slide, please. You will see more of this, I am sure, in the future. Here is the U.S. population. I have now increased the resolution by a factor of 25. U.S. population by one degree squares. Well, quadrilateral, technically. The next slide shows the data by county resolution. If you define the resolution of a geographical data set as the number of pieces of data divided into the geographical area involved, to the K th root of the dimension--in this case, two dimensions, you find that the county resolution is--I have forgotten--3,200 divided into three million, the square root of that. From the sampling theorem we know that you can only see phenomena of a wave length which is twice the sampling interval, if it is not noisy data. So if you want to get real detailed information on the behavior of people, and people have a mean activity locus of about 20 miles a day, that means you need resolution of about a 20th of a county to detect anything interesting about individual behavior. The next slide, please. Also, geographical data are arranged hierarchically. Here we have a typical socioeconomic type of data where we have municipalities, economic areas, provinces, regions, and, finally, the whole country. This happens to be the Netherlands.

One question that I think will be developed further in the future is how you take these very interesting algorithmic manipulations that people do on raster pictures, as I am sure Dr. Andrews will talk about--how do you apply that to polygonal data?

I will introduce Jim Blinn, who is currently at the Jet Propulsion Laboratory, and also at the Computer Science Department at Cal Tech. I first met Jim at the University of Michigan, where he was doing work on a music synthesizer. He was at that time in the computing center. He received both his Master's and Bachelor's degree at the University of Michigan. He was also very influential in developing integrated graphics system, which runs under the Michigan thermal system, essentially a device-independent graphics system, where you sign on on the terminal, and the computer queries your terminal to find out what kind of terminal it is, and then the software appropriately modifies itself so that it will handle output for that terminal, so it does not matter whether you

are on a Tectronix or a Hewlett-Packard, or working with a Calcomp or what have you--you use one set of software for all of them. This obviously is the way technology is going to have to go.

He has also worked at the New York Institute of Technology on computer animation, and currently does consultant work for Information International, which does computer animation. Jim left the University of Michigan and went to Utah where, of course, Ivan Southerland and Dr. Evans were working, and got his Ph.D. there in, I guess it is technically, Electrical Engineering. But he is interested in computer representation of three dimensional objects. Jim, do you want to come up here?

(EDIT NOTE: The presentation by Dr. James Blinn has been omitted because his closely-related illustrations were not available for publication. Some other illustrations shown at the meeting were unavailable, or unsuitable for printing, without the need for omission of the speaker's remarks. Dr. Blinn's remarks are omitted, however, following his suggestion, and our concurrence, that they would have little meaning without accompanying illustrations.)

MR. TOBLER: Thank you, Jim. As a cartographer I would have a lot to say about these objects that do not exist, but I will leave that to Carl Youngmann to talk about. Jim has only recently joined JPL. As you probably know, JPL has had a lot of experience with image processing, and I look forward to seeing the two technologies merge.

The next speaker also has a lot of experience in image processing, Dr. Harry Andrews. He worked at Stanford and at Southern Cal. The Image Processing Institute is located there. He has some literature which describes that operation. Harry has published two books on this. I took a course from him one time at Purdue on image processing. He is going to tell us about the work they have been doing there. His own work, I know, has been in image transforms, and I found it very interesting. They are most recently, I think, getting into satellite picture processing, but I will let him tell you about that. Harry?

DR. HARRY ANDREWS: I would just like to say very briefly that, thank you very much, Waldo, for inviting me. I do not know anything about this group. I made the unfortunate mistake of spending two weeks at DMATC, if you know what that means, in Washington, D.C., during the summer. And after that I thought I was an expert. But it is true, I am not. I did leave a little brochure out there in the next room, if you are interested in finding out what we do. Some other material I will leave up here describing the Image Processing Institute, which is an educational facility at USC. Our goal in life is to train image processors. I will leave a lot of this up here. If you are interested, feel free to take some of that material.

I think what I will do today is represent the Department of Defense community. You see, what I do is size up an audience and make sure that nobody in the audience knows anything about what I am talking about, then I represent that other side. Probably you all know about the Department of Defense, but, in any event, we have been funded for -- I guess I should not admit this -- they may want to redirect their funds. We have been funded for a while by an organization known as ARPA. ARPA is a group in the Department of Defense that gathers its funds from the other services before they get their money, so they manage to offend everybody, and then they give it to the university. (Laughter.) We established the Image Processing Institute quite a few years ago, and have been actively involved in trying to manipulate imagery with digital computers for the benefit of mankind. If ARPA is paying for it, it may be to the detriment of the enemy but to the benefit of us.

In any event, what I would like to do today is briefly go through a somewhat historical synopsis fairly rapidly to show you some of the things that have been developed in the past, and to show you some of the direction that we are moving in right at the present. I will go as rapidly as possible, because I would rather entertain questions than entertain you with slides. May we have the first slide. The first slide is where I am from. The next slide is the block diagram which you also saw of our computer facilities. I will not spend much time describing it, simply pointing out that we have a couple of KL-10 computers, which we do our number crunching on interactively, and we pass that data to our exploitation facility station. One of the topics or highlights that I think this conference might be most interested in is in the area of exploitation facilities, the use of high speed digital interactive computers and displays to allow a human to interact with large digital data bases.

Next slide. Before we get into that, I want to show our typical USC girl. Actually, this is an SMPTE slide, but you have probably seen her around. Just to set the tone, let us look at the next slide. She is made up of bits. Those are the four most significant bits of the green component. The next slide shows the four least significant bits of green component. Essentially what we did was sliced her into 256 levels of brightness for each color, thereby resulting in 24 bits, or eight times three levels of brightness per pixel.

Next slide, please. That is what a digital girl looks like. Now, to just show you some of the things we do interactively on our display devices, we have developed what we call a little menu system, but it is an interactive or visible menu system. If you look very closely you see a little white dot that is approximately adjacent to a box to the right of the imagery. That essentially tells the computer what the operator wants to do next. In this case, the operator simply zooms down onto the airplane, which happens to be an aerial photograph of the Los Angeles Airport. The next slide. That sort of thing can be pursued in greater depth. You can magnify it all sorts of factors; two techniques showing the difference in magnification. All of this is done somewhat instantaneously on less than a 30th of a second type of interaction.

The next slide. Waldo mentioned the concept of image transforms. Really what that means in academia you had to do something mathematical or else you get fired. So what we did is we have tried applying a little mathematics to pictures. It turned out that we kept one step ahead of the reviewers for about three years and managed to publish a paper every time we invented a new transform. Actually, we didn't have to invent new transforms; we just looked in

some mathematical literature and then published in the engineering literature, and nobody reads both. (Laughter.)

The next slide. Essentially, one way of viewing an image, albeit somewhat artificial, is the fact that an image is nothing more than a matrix in a computer. It could be a very large matrix, maybe 2,000 by 3,000 pixels or picture elements, if you are talking about ERTS type of photography, or LANDSAT. You can break that image up into a sum of other images, as illustrated in the top row, or the bottom row, and you can add them all up and form the original. Now, why might you want to do this? Well, you may want to discover some underlying structure of the image that was not readily evident in the original photograph or original digital version. If you break it up into a set of basis functions known as frequencies, then you end up doing a two dimensional Fourier transform. If you break it up into a set of basis functions known as sequencies, then you are breaking the image up into a set of Walsh transforms.

As I mentioned earlier, we generated as many transforms as we could get away with. Let me show you what just a few of those look like on image. Next slide. Here is our little toy tank broken up into its frequencies. The frequencies are then colored: The low frequency, red; the intermediate, green; the high frequency, blue; and then put back together, so you can get a little feel for what spatial frequencies mean. That is the combination of the individual frequencies. Let us look at the next three slides in succession. There is the low frequency, red. Notice the center of the star, low frequency, and just the circle comes out. The next slide, green, shows the edges of the star, and the number twelve. The spatial frequency of the "12" was an intermediate band of frequencies. The next slide, the blues, show the edge effects and the tank treads.

Next slide. If you are breaking imagery up into square waves rather than sine waves, you could break them up into these functions known as Hadamard or Walsh basis functions. There was a tremendous interest in this sort of decomposition of imagery back when semiconductor technology was running rampant, I guess there still is. The idea here is that switching functions can very easily decompose an image into these sets of functions.

Next slide. There is one transformation that is optimally matched to the image in terms of least squares approximation. That transform comes to us from a technique in numerical analysis known as singular value decomposition. What we will do is we will find the set of basis functions which are best suited for a given image, because that in turn will give us the best means of compressing an image into a few number of coefficients for image compression and

transmission over communication channels.

Let us look at the next slide. What I want you to do is tell me in the sequence of slides coming up when you recognize what the image is. This is the sum of the first image and its coefficient.

The next slide is the sum of the first two. The next slide is the sum of the first four. Does anybody wish to hazard a guess?

The next slide is the sum of the first eight. Are there any guesses?

MR. TOBLER: A girl?

DR. ANDREWS: It's a girl? Waldo. Will you be embarrassed. (Laughter.)

The next slide, the sum of the first 16. The next slide is the sum of the first 32. The next one, the sum of the first 64. The final one, the sum of the first 128. That is a building, Waldo. And I think we should send you back to the university. (Laughter.) That technique is utilized for image compression. The idea of image compression -- if we may have the next slide -- is to transmit imagery over communication channels with as few bits as possible.

I mentioned earlier the possibility of exploitation facilities where you had digital data bases and wanted to manipulate those data bases. One of the real problems with image exploitation and human interaction with such data bases is keeping the human entertained, or essentially keeping the bandwidth high enough so that the human is not sitting there waiting for an image to be brought up for display purposes. One way of observing this phenomena is simply to admit that we in the industrial, military and civilian complex will never control that industry out there known as television. So rather than trying to get higher resolution displays, why don't we live with the 512 by 512 color real time refresh monitor, and use that as a window to zoom around much, much larger data bases? What we are going to do is a scenario in the next sequence of slides in which we zoom in on this data base, and, if you can imagine, we would be doing this in real time.

Let us go through the sequence. The next slide we zoom down a little lower. The next slide, closer in. The next slide, even closer. The next one -- Does anybody know where this is? Yes, this is Gary, Indiana, one of the garden spots of the nation. (Laughter.) That is not smoke. That is effluent going out into the lake.

Let us go to the next slide and suggest that we may want to even go further than a one-to-one mapping. The previous slide was a

mapping. The previous slide was a mapping of one pixel in the original image domain to one pixel on the display. We may want to magnify and artificially introduce data that does not really exist. We might call that interpolation. The mathematicians tell us how to do interpolation under certain mathematical criteria. Here are four or five interpolating functions. The first is called a sample and hold, or a replicator; the second, a bi-linear interpolator; the third a quadratic; and the fourth, a cubic. These all have certain properties in two dimensions. The next slide shows what they look like.

These interpolating functions allow you to have continuity in the zeroth, first, second, third and higher order derivatives at the intersections of the true data points with the interpolated data points. The next slide shows a simple illustration of this where we have taken a 32 by 32 image and interpolated it up to 512 by 512 using the various different techniques. Now, I cannot say or you cannot prove it with this slide, but I am told that certain individuals with trained eyes can actually see the difference in the second derivative of an image and its interpolation function.

Next slide. I would like to now discuss very briefly the idea of spatial warping, the idea of automatically in a computer registering scenes or images from platforms that were never originally intended to be registered. As an illustration, we may consider taking an ERTS photograph or LANDSAT photograph and registering it with a U-2 photograph to see if there have been changes, to see if seasonal effects have been measured, et cetera. What one might consider doing then is taking the first image, number one, and the second image, number two, and finding control points or points of commonality between the two images. This might be done automatically or with the human interacting and then refined automatically, et cetera. Then a two dimensional polynomial might be computed and a "warping" function, as they are referred to, describing the mapping of Image 1 onto Image 2. Naturally, that warping function will not exist in terms of the discreet components or sampling rate of the second image, so we will have to do some interpolation.

When you begin to think a little more about this problem you realize you do not even have to map Image 1 onto Image 2. You could map Image 1 onto Image Imaginary and Image 2 on Image Imaginary, and do all sorts of intermediate types of maps. So let us do that with our favorite building. The next slide will show an original. This is the original building that we will do some mapping with. The next slide is an original. These are the only two originals. Every image hereafter will be phony, phony in the sense that they were generated from those two,

The first sequence will be -- Let us look at the next slide. We will simply map the original, and rather than taking a photograph of the front of the building with the helicopter, we will simply let the computer rotate it. Now you see some difficulties with this process. The fellow who put in the control points forgot a little bit about perspective, and the building gets wider as it goes away from you rather than narrower. So the roof does not look too appropriate. But every pixel in there and every bit of information is in the original.

Next slide. Here we have taken the end view of the original building and simply magnified it by a factor of three and put it in the upper right-hand corner of that imaginary scene. Now, the next slide shows the warping of the front end of the building to exactly register on the end of the building. Now, this is difficult for you to envision without seeing both of them simultaneously, but, take my word for it, if we now take the original end view and this warped front end view and look at the two images through stereo we will see stereo information on the end of the building, because, naturally, you cannot perfectly register two image which were sensed from platforms at different positions. The resulting information, of course, must be stereo if you have done it right.

The next slide shows what happens if you introduce a bit of control point error. You see the curvature that you may not believe in. The next slide shows that on the second sequence.

The following slide is my favorite. This is the one that I always claim is the original, and the first one you saw is the result of undoing the rubber sheet. Imagine what you could do in a court of law with this sort of image process. That is why we keep getting thrown out. (Laughter.)

Next slide. This is a calibration slide. It serves no purpose other than to prove to you that you all have normal vision. You probably see scalloping. By scalloping, I mean you do not see steps of brightness change, you see differentiation or brighter edges at the left than at the right, when in fact with a densitometer or light measuring device there is nothing but a staircase. This is known as M \ddot{a} ch banding phenomena, and typically is the beginning of what we call the model of the psychophysics of vision. If indeed you are going to have a computer display imagery at an exploitation station, you might as well realize what your eye is going to do with the image so that you in the computer can precompensate and undo or take advantage of your own visual processes. So this is a technique now being utilized in image compression laboratories to try to get as much information out of the image ahead of the transmitter,

knowing that you will never see that information because of your own visual process.

Next slide. This introduces and displays the common spatial frequency response of the eye. You should see, hopefully, a ramp going up and a tapering off of the spatial frequencies over to the right. That is the MTF of your eye. In other words, you cannot see certain spatial frequencies beyond the sampling rate of the retina, and at low frequencies your eye tends to be a differentiator. That is why, in fact, you enjoy crisp images, the edge effects the information content and image enhanced in the high frequencies is very appealing to viewing. Next slide. Using that and some other models of human vision, we put together an image compression system which is compressing imagery in color. The eye has a very, very natural response to color in the following sense: That we do not see sharp edges in chromaticities, we only see sharp edges in brightnesses. You probably have already known that because of your television system in fact takes advantage of that; the NTSC color transmission system does not send high frequency edge information in color; it only sends high frequency brightness information in monochrome. Using this and other models of the human visual process, you will find that we can undo and remove a tremendous amount of information that the sensor gathers, but which you will never see, so we remove it prior to transmission through communication channels. That is what this block diagram is supposed to illustrate.

The next slide shows some intermediate stages. If you could in fact find a volunteer either in this audience or anywhere else who would be willing to let us remove the eyeball and see what an image looks like on the retina after the nonlinear devices, we would expect to see the images on the right. Now, there is no way you can disprove that, right? (Laughter.) Unfortunately, there is no way I can prove it, either. But we think that is what the retina sees after the nonlinear photochemical processes.

Next slide. This shows our original girl in four quadrants, and shows various different color combinations. We have not yet removed the bandwidth. These are 24 bit per pixel images. This is just to prove that we can take the mathematical models we have described, go to the coding domain, and come back. Yet we have not yet compressed. The next slide shows the effect of compression. This is from 24 down to two bits per pixel using what is called a block cosine transform with the various mean square errors there.

The next slide shows the result of a one bit per pixel communication system. Degradation is now starting to be introduced.

Next slide. We have considered developing what we call smart sensors. Actually, we haven't; ARPA has, and therefore we do. ARPA would like to be able to do the following: They would like to be able to put you people all out of business. But, don't worry, don't worry. With us on the project -- (Laughter.) -- you can feel confident.

Essentially, they would like to design what are known as smart sensors. A smart sensor is a device that tries to do things as close to the focal point as is possible, such things as possibly taking advantage of the fact that images are highly non-stationary. By that I mean that the information content, the energy, the edge information, the things we as humans like to see, are very different as you move around the field of view. So we have tried to design a technique that captures automatically this information content and essentially resamples or adaptively scales the imagery such that we can only send over the communication link that which is of most interest to the viewer. So, here we have a little APC, armored personnel carrier. You can see on the right the density of what we call sampling or knots, which is a mathematical term. It explains how you would have to transmit and re-sample the image to get the various parameter reduction ratios you see on the left. The next slide is a down-looking slide of our airport. You can see the system automatically keys on the airplanes and the high edge effect.

The next slide shows the result of an artificial ERTS type of photograph. You can see the airport, et cetera. Next slide. We have done quite a bit of modeling of the process of imaging and what mathematically you might expect an object to suffer when it indeed is imaged onto some sort of film or CCD or other digital sensor.

If we can go through these slides. Typically, today with the technology of discrete sensors, we use a model which is known as a continuous-to-discrete model. We hypothesize that the object we are looking at is continuous, but one that object is in digital form is discrete. Then we ask the question, how can we mathematically manipulate our discrete information to better appropriate the original object, knowing all about the lenses, knowing about the atmosphere, knowing all about the sensor and its nonlinearity, knowing as much as we can about the statistics of the image and of the noise, et cetera. So that is the typical mathematical model that we would play with. Incidentally, if anyone is interested in a lot of the gory details of this and other research, you are certainly welcome to contact me or come up and look at the publications list. We put out a report every six months.

Next slide. Here is a technique developed by the Aerospace Corporation in which we observe a real world imaging system. This slide

does not really do it justice. On axis this system is space-variant, but off axis the point spread function is non-symmetric. It really tells us that the image that you would expect to obtain from such a set of optics would have blurs that would vary as you moved around the field of view. The next slide shows what you might do with this sort of information. Here is a simulation in which the imagery is better in focus in the center of the squares and more out of focus at the edges. The question is, what type of mathematical techniques can we utilize to undo this type of blurring distortion? We call this a space-variant distortion because indeed the blur varies as you move around the field of view or the space of the image. The three dimensional perspective shows that effect.

Keep in mind the idea of moderate distortion and severe distortion. The next slide will show what we can do to try to undo this type of blur. Now we are talking about restoration; we are trying to undo the degradations of the previous slide model. Here we have our favorite subject. She is known as Tiffany. Tiffany is available in digital form on magnetic tape, for anybody who might be interested. Up there we see the blurred version of Tiffany with a moderate blur, a blur that was in the middle slide, the previous slide. Here we have an iterative restoration of Tiffany up to about K equals 100; that is our index of information. We call this coarse tuning.

The next slide shows the fine tuning and also illustrates what happens if we try to put too much back into the image that was not allowed to pass through the optical system. The system blows up on us. This is typical of most image restoration processes. They are known as ill-conditioned, a mathematical expression that says "there ain't no free lunch"; you don't get more out than went into this system. So, in fact, we are getting or approximating singularity. You can see in the very last slide a very faint version of Tiffany.

We will look at the next slide, the same simulation with coarse distortion, where she falls apart much earlier at a lower index, K equal to about 58, compared to 100 previously. This was the severe situation. Let us go on. This final sequence is just to illustrate what we are doing with an industrial group that is tied to USC.

A couple of years ago ARPA decided that they wanted to speed up what is known as technology transfer -- that is, getting the ideas out of the universities into the industry without waiting for publication and review cycles, and then somebody in industry picking up the ideas, et cetera, taking two or three more years before they became product lines. So ARPA decided that they would let the tail

wag the dog, and they gave the monies to the universities for them to distribute to the industry. You can imagine how far \$100,000 goes at a university where graduate students still work dirt cheap, and you can imagine what pain it does when I have to give that same amount of money to the Hughes Aircraft Corporation with an overhead of whatever it is, and they tell me, "Well, that will buy ten minutes of somebody's time." In any event, we twisted and begged and borrowed, and we have teamed up with the Hughes Research Labs to build a smart sensor. This sensor is essentially a CCD device that is designed to fit on the focal plane of an imaging camera that does some non-linear processes known as the Sobel operation; this is an operation that essentially looks for edges, and is very insensitive to noise.

Let me briefly go through these slides, because I really do not know what I am talking about -- You have to be some sort of solid-state physicist to understand. Let us look at the next one. What we want to do is a two dimensional signal processing process in which the charge passes under the various CCD devices that has been sensed with an electronic optical-to-charge transducer. We then take the outputs of these devices, and we want to do a non-linear process. In image processing, non-linear processes are probably the most useful ones. We take the outputs of these four lines, and pass them through an absolute value detector. The unique thing about this smart sensor -- next slide -- is the fact that here is what they call a spill and fill network. In addition to being a solid-state physicist you have to know about plumbing, because they do all sorts of things with charges and wells and liquids which really do not exist.

In any event, this does take the absolute value of the difference of the outputs of the previous pixels, and the output of this then becomes essentially the Sobel operator. The next slide shows essentially the chip that exists. This is a rather large magnification. This device actually is 190 mils by 190 mils, which is very, very tiny.

We are now in the process of building a much smarter device that measures texture off of sensors, and hopefully someday we hope to be able to automatically on-board the platform or the sensor, segment images so we can just transmit segments of interest rather than the entire imagery. Thank you very much. (Applause.)

MR. TOBLER: Thank you very much. I certainly found that very interesting, even if I can't tell a girl from a building in the truncated transform. (Laughter.) Image processing people always pick pictures of interesting things like baboons and girls and buildings.

The next speaker, Carl Youngmann, from the University of Washington, trained in cartography at the University of Kansas. He then went to Columbus, Ohio, and taught cartography there at Ohio State University, did some work with lasers and worked on geographical information systems. At the University of Washington he teaches computer cartography, and has also been involved in what might be called semi-automatic generation of the Coastal Zone Atlas of Washington. Carl?

MR. CARL YOUNGMANN: When Waldo asked me to provide the cartographic perspective on the use of computer graphics, image processing and other advanced technologies, I felt like the housewife or house-husband who had won an unlimited shopping trip in a grocery store promotion. I could buy anything I wanted in the technological store, but I didn't know how many guests there would be and whether there would be any vegetarians. To try to sort out what the present menu of technologies can do for us on a practical day-to-day basis, my feeling is that after we have accepted the existence of the things we have talked about and have seen here in the past few minutes, the real problem will be being able to know what we want and how we are going to utilize these devices and processes to represent what we want: What do we require in a computer-generated display? How are we going to utilize image processing systems such as that we have seen? Or a smart sensor? How would we utilize a device such as the one Jim Blinn talked about with which we could create an unnatural world as we would like to see it, and move through it? Where does such a device fit into the day-to-day problems we have to solve as cartographers? We must communicate spatial information to people who are not necessarily going to believe that these things really exist and that the manipulations are, as we have seen them, quite fascinating, but more importantly, reliable.

From the cartographic point-of-view the concerns are many and they are difficult to sort out--indeed, it is hard to get on top of a high enough hill to look over the variety of things we feel are important about cartographic displays. Immediacy of response, accuracy of communication, volume of data to be handled, scope of data to be represented, level of detail, scale, updating and extending images, inversions and analyses that can be made of the information, static-dynamic aspects of information--all seem to be

very important. Can we reduce these concerns so that we can at least start to give some structure to the cartographic shopping list?

One thing that we could say first is that accuracy is not really a representation problem. We have been shown that we can take something at one level of accuracy and really give the impression that it is much more accurate than it is, up to a point. Graphic devices are accurate. We will have to accept that.

Secondly, we can assume that the ultimate needs will generate the ultimate response; that technology does provide us the ability as cartographers to do many things--really the capability of doing more things than we actually want to do or need to do at any given point. Basically, we can take the approach that cartography is not wholly, from the computer scientist's point-of-view, an "NP Complete" problem--that is, although there are minor problems that we would say defy computation by sequential digital analysis, cartography is basically a computable problem. It can be addressed by sequential digital processes and probably can be handled in most contexts.

Successful methods, thirdly, are or will be the ones that are low cost, high quality, widely distributed and technologically simple. It is hard to make a technological forecast for 25 years from now. Various advances in solid state physics may end up radically changing the cartographic displays that we have to deal with. We may never ever have a machine such as the one that Harry Andrews works with day to day, but maybe some day we will.

The function of cartographic communications really determines the need for cartographic media. May we have the first Vu-graph (Figure 1), please. To this point, in many cases we have had the tail wagging the dog in that the media available to us many times has determined how we were going to communicate. Perhaps we should now look in the other direction and define what we want to communicate and then select a medium.

Cartographic communication is used in five general functions. Those functions include first what we might call the dispatching function, in which we provide information for decision making in a network--routing of trucks--giving delivery people something on a day-to-day basis so that they can perform their job more efficiently; routing of emergency vehicles; and similar questions of spatial dispatching. The second area is what we might call status reporting, very quickly delivered information about the location and distribution of objects interacting in space. Aircraft control and vessel traffic control systems fit within such a scheme. Item

three deals with research, an all-encompassing topic which involves summarizing and analyzing the distribution of geographic phenomena. Such communication is not so much directed to somebody else as it is to oneself, giving oneself the information about the location and distribution of phenomena that are of interest to him or her and being able to change the representation in some way or another so that one might learn something. In that way, researchers come to generalizations, so that they can then proceed to the fourth function: reporting. For reporting, information must be presented in a manner that meets special criteria, like the high quality graphics used for planning documents, census reports, and other similar items. The fifth area is recording, the actual placing of information in a highly retrievable spatial context, useful for reference--surveying data, cadastral information, general purpose atlas information, etc.

These five functions dictate certain kinds of needs. May I have the next Vu-graph (Figure 2), please. I think the graphic environment in which these functions fall really can be boiled down to three major axes. These axes may be independent orthogonal, or they may be interrelated. I show them here as orthogonal. The axes relate the requirements of a cartographic communication function to a display system's capability for immediacy of response, the data capacity, and graphic quality. Each axis is divided into a set of relative classes ordered by increasing capability.

In regard to immediacy of response, we have available cartographic display systems which run from instantaneous response--a second or less, to those systems which provide a timely response--within a few seconds to one minute, and systems that provide what I would call an "assured" response--"If I submit the job, I know I am going to get it back sometime." The second axis, data capacity, starts at the lowest end with one element of information and runs through low data capacity systems--up to say a hundred map elements; then, moderate capacity--a hundred to a thousand elements; next high capacity--a thousand to fifty-thousand elements; and finally, very high capacity systems--over 100,000 individual elements in an image. The third dimension is graphic quality of the final product. This axis extends from extremely crude resolution, the one-tenth by one-sixth of an inch SYMAP style graphic, to what we might call a coarse graphic--the matrix printer, a tenth of an inch by a tenth of an inch output; to moderate graphic quality, on the order of a hundredth of an inch resolution and finally, to fine, high quality graphics of one-five-hundredth of an inch resolution.

The five communication functions are then placed within a framework to show the different levels of graphic quality, immediacy of response, and data capacity required for each.

For dispatching the delivery truck driver can get by with a coarse, SYMAP style output. The output has to be ready once a day when the trucks go out, and the data capacity is of a moderate level. Status reporting has to have a moderate capacity. The information must be clearly presented, although it does not require extremely fine resolution. However, the system must respond almost instantaneously. In research we can deal with coarse quality, but we must have assured response and a high capacity. Reporting and researching require higher quality and higher capacity but only assured response.

Each of us, of course, will see our own specific communication needs in a different light. Indeed, my presentation of the dimensions of our requirements oversimplifies the complexities encountered in matching a graphic medium to a communication function. Furthermore, I have ignored two other significant dimensions: mode of use, which ranges from static representation to dynamic, interactive graphics and in the extreme to intelligent, reactive representation; and dimensionality--representations of variation in multiple phenomena, multiple spaces, or time.

Now, with this understanding of our requirements we can approach image processing and computer graphics. I think this shows us a place to begin our search. I do not know how interested designers are in creating systems that will match these requirements, but I hope that at least now we have some idea of what kind of dinner we can begin to serve, having won the cartographic grocery shopping contest. Thank you. (Applause.)

MR. TOBLER: Thank you very much, Carl. As an academic, I would have a question for Harry. How do you convince the government to give you money for academic research? Are there any questions?

Thank you all very much, and have a good evening. (Applause.)

Figure 1

Cartographic Communication Functions

- ☆ DISPATCHING
- ☆ STATUS REPORTING
- ☆ RESEARCH
- ☆ REPORTING
- ☆ RECORDING

Figure 2

Cartographic Communication Needs

