

THE 8-MILLION CHANNEL NARROWBAND ANALYZER

Paul Horowitz
Harvard University

John Forster
Cambridge, Massachusetts

Ivan Linscott
Stanford University

ABSTRACT. An 8.4 million channel narrowband spectrum analyzer is nearing completion, and will be used to expand the frequency coverage of the ongoing search at Oak Ridge by a factor of 200. The new system -- project META -- will cover 420kHz at 0.05Hz resolution, utilizing a swept receiver to cancel the effect of the earth's rotation. The increased bandwidth will permit observation of CW beacons transmitted at magic frequencies in any of three preferred frames: the local standard of rest, the galactic barycenter, and the cosmic blackbody rest frame.

This is a literal transcription of the talk, slightly edited for readability.

Yesterday I described the system that we're now using out at the Oak Ridge Observatory in Harvard, Massachusetts, in which we are looking for directed CW radio frequency signals at "magic frequencies". It's a very special search in the sense that we wouldn't see signals that were not predopplered to arrive in our heliocenter at 1420MHz (or whatever special frequency we guess), mostly because we don't have enough bandwidth to cover the doppler offsets caused by various relative motions of order 10km/sec (50kHz at 1420MHz). What I will discuss today, with my colleague John Forster and absent colleague Ivan Linscott, is an expansion of this system to approximately 8 million channels in order to lift the restriction of predopplered signals. The new system is called META (these projects have to have names, I understand) -- the Mega-Channel Extra-Terrestrial Assay, (ET's my middle name) an 8.4 million channel enhancement of the current 128,000-channel project Sentinel. And here are the credits: The Planetary Society is supporting the upgrade. We are building on equipment we already had before, the receiver built with support from Planetary Society and NASA, some contributions from Hewlett-Packard corporation and the observatory itself owned by Harvard and Smithsonian. We are using a portion of NASA's multichannel

361

M. D. Papagiannis (ed.), The Search for Extraterrestrial Life: Recent Developments, 361-371.
© 1985 by the IAU.

analyzer -- the "bandpass filter #2" -- as a fast Fourier preprocessor; its creators are Allen Peterson, Ivan Linscott, and Kok Chen at Stanford University. And finally, those eager undergraduates at Harvard, Brian Matthews and Teddy Kim. It's one of the great benefits of doing science at a University. They're wonderful. They stay up all night and write programs and just get a great thrill out of it.

This is going to be a sort of Huntley-Brinkly talk (or maybe it should be Bert and Ernie): I will start out by giving the scientific rationale, John will describe the actual implementation of the processors, and I'll finish with a wrap-up. Why should we go to 8 million channels? What's wrong with what we have now? The current system has a pair of 64,000 channel analyzers, and the bin width is 0.03Hz per channel. You can calculate for yourself that gives a pair of 2kHz bands. Two kilohertz is OK if you're willing to accept a magic predoppler CW signal as I described yesterday, in other words a signal that's guaranteed to arrive in our heliocenter at the right frequency. That's not a totally outlandish thing for an extra terrestrial civilization to do, particularly if they have a good reason to target our star and really want us to get the message, and assume that we're primitive and can only receive that much bandwidth. But it makes an awful lot of assumptions. The assumption we really would like to lift is the necessity of pre-compensating for the target star's motion along the line of sight, because we'd like the extra terrestrials to be able to transmit an omnidirectional beacon, or perhaps a signal that's at a magic frequency in some frame we can agree upon, perhaps the galactic center, perhaps the cosmic black body rest frame.

But with the current system we don't have enough bandwidth for that. A typical stellar radial velocity for a nearby star is its peculiar velocity within the local neighborhood, or of the order of 20 kilometers per second: That's 100kHz at the 21cm line and of course scaling like frequency. And in addition, because of differential rotation in the galaxy you need to add an additional peak offset of another 75kHz per kiloparsec. So we're talking bandwidths, then, of order 100 or 200kHz, if we want to receive signals that have been transmitted at a magic frequency in the local standard of rest, for example.

One solution, of course, is to lower the resolution (that is, increase the binwidth of each channel) so that the same number of channels will cover enough bandwidth. You will be hearing about some searches that go in that direction later in the session. But the problem then, at least from our point of view, is that we don't maintain this nice match to the Drake-Helou minimum bandwidth of the galaxy, and that means we lose signal/noise ratio, which we can scarcely afford with our small antenna. We would lose a factor of at least the square root of the ratio of increased bandwidth. But secondly, and I think more important, is that we would lose the remarkable cancellation of terrestrial interference that I described yesterday, because at extremely narrow binwidths we have to sweep our receiver. That gives the extra terrestrial signal a different signature from terrestrial interference. From our experience of last year, this has worked well in rejecting interference. Clearly this is going to become even more important when we go from the **number of channels** that we now have (128,000) out to 5 or 10 million

channels, and therefore if we are going to increase channel numbers significantly, we would really like to keep the resolution high.

So the best solution is to keep high resolution, and simply expand the number of channels sufficiently to bracket the Doppler shift that we might expect to see in the galaxy. The kind of Doppler shifts we might expect are (1) the civilization may simply transmit a magic frequency in the local standard of rest, or perhaps in their own heliocentric frame, or (2) they might choose instead the reference frame of the galactic center. They can measure it better than we can. Or (3) probably best of all is Phil Morrison's suggestion the other night that we now have a preferred universal reference frame that we already know to reasonable accuracy and will soon know much better: The frame in which the remnant of the big bang is isotropic. (A related suggestion was made by Gott in the volume by Hart and Zuckerman, *Extra Terrestrials, Where Are They?*, Pergamon, 1982.)

Let me just summarize on the slide here (table 1) the kind of bandwidths we are talking about. I have put down here in order of increasing bandwidths or Doppler shifts, all the effects that contribute doppler shifts to interstellar radio beacons. The first -- the Drake-Helou spreading, is of course not a shift, but a spreading, and dictates the ideal resolution of a spectrometer that is a matched filter to the interstellar medium. The barycenter/heliocenter ambiguity refers to offset of the solar system center of gravity relative to the sun. If a transmitting society were good enough to know (and transmit in) our barycentric frame, it would be at most 60Hz from the heliocentric frequency.

<u>Effect</u>	<u>Shift</u>	<u>Comments</u>
Multipath FM ("Drake-Helou")	0.01-0.1Hz	A good <i>binwidth</i> ! Must chirp, gives RFI rejec.
Barycenter/heliocenter ambiguity	±60Hz	Jupiter
Beamwidth (84' → 0.5°)	1.5kHz	Untargetted only
Spin (uncorrected)	±2.5kHz	Spectral smear: <u>must correct</u>
Orbit (uncorrected)	±150kHz	
Stellar Radial Velocities		
Galactic Rotation	±75kHz/kpc	} Survey. No attempt to correct
Peculiar Motion	~±100kHz	
Big Bang Rest Frame Offset	1.5MHz	
Uncertainty	±150kHz	COBE should improve
Galactic "Center"	±125kHz	Dixon's value yesterday

Table 1. Doppler shifts relevant to acquisition of interstellar beacons at magic frequencies. All values are calculated at 1420MHz, at which 1km/s corresponds to 5kHz.

These frequencies are all calculated, by the way, at 1420MHz, the most magic of magic frequencies. The next entry is an additional bandwidth needed (and this is, in fact, very close to our current bandwidth) be-

cause with a finite beam size, we don't know exactly where in the beam the source might be. The different beam directions have different direction cosines with the Earth's orbit, and that gives rise to a minimum required bandwidth simply to make sure that a magic frequency source anywhere in your antenna beam will be within your spectrometer band. I add the comment here that this is relevant only if you are doing an untargeted search: If you are looking at a particular star, you know its right ascension and then there is no "beamwidth ambiguity" there whatsoever. Next in increasing bandwidth are the spin and orbit terms of the earth, which I have already discussed in yesterday's talk. The orbit is the big term, but the spin is the rapidly changing term. This latter term gives rise to a spectral smear unless you correct for it, and therefore bandwidth isn't what you need here, you need a frequency chirp correction. The current narrowband system does this, and we plan to continue that.

What about these stellar radial velocities that we've ignored up to now? We've said we'll let the extraterrestrials worry about that, targeting our star with a pre-dopplered signal that is delivered at our heliocentric frame at the right frequency. The table indicates the size of doppler shifts involved: Something like 75kHz per kiloparsec due to galactic rotation, superposed on peculiar motions of the order of 100kHz. The big-bang rest frame has an offset with respect to the heliocenter of the order of $1\frac{1}{2}$ MHz. (If you want to convert all of these, by the way, to velocities, a km/sec is 5kHz at 1420MHz.) More important than the velocity offset of the big bang rest frame (which we can allow for), is the uncertainty in that offset, for which we must provide spectrometer coverage in the form of bandwidth. That uncertainty in the measured dipole moment of the 3°K background corresponds to something like ± 150 kHz of required bandwidth, but new measurements should improve the situation. The cosmic background explorer, COBE, I'm told by Sam Gulbis here should improve that by an order of magnitude or two. Then finally, the galactic center as a preferred frame. This is something I would like to get some comments from here, but I gather that you can make reasonably accurate measurements of that object at the center, in spite of its asymmetries and rapid motions. According to Bob Dixon yesterday, the uncertainties in velocity of that frame amount to something like 125kHz of doppler at 1420MHz.

So all three of these numbers here -- stellar radial velocities, uncertainty in big bang rest frame, and uncertainty in galactic center -- suggest that if our spectrometer could receive about 300 or 400kHz of instantaneous bandwidth, let's say ± 200 kHz might be a nice number, we could make sure that we wouldn't miss the signal in, for instance, the big-bang rest frame, by simply offsetting to our nominal velocity relative to it and then looking with enough bandwidth.

Let me show a slide (figure 1) that presents these bandwidths graphically. We begin with a transmitted carrier at 1420MHz of zero bandwidth. It propagates through the galaxy and, as prescribed by Drake and Helou, it spreads in bandwidth something like 0.03Hz. That's still very narrow -- it takes half a minute of observation just to see that -- but that signal chirps 10Hz per minute at our receiver because the earth is spinning. That's 300 channels in our current system of 64,000




Figure 1. Hierarchy of doppler shifts associated with spin, orbit, radial velocities, and choice of rest frame. All shifts are at 1420MHz.

channels whose total bandwidth is 2kHz. That system ("Sentinel"), which I described yesterday, chirps its local oscillator (LO) to compensate, and in addition offsets its LO as needed to compensate for orbital and spin velocities. However, Sentinel does not attempt to compensate for stellar radial velocities, whose magnitude is typically much larger than the current system bandwidth of 2kHz.

Our new system, dubbed "META" (Mega Channel Extraterrestrial Assay), has 420kHz instantaneous bandwidth, enough to bracket these stellar velocities, at least out to several kiloparsecs. As the figure shows, that spectrometer bandwidth is also adequate to bracket the uncertainty in the big bang rest frame, which however is offset by 1.5MHz from the heliocentric frame. The galactic center has a comparable velocity offset (though of opposite sign) and comparable uncertainty in velocity; therefore the META analyzer is well matched to observations of magic frequencies in that frame, also.

In actual observations we will switch between these favorite rest frames in successive integrations, since any source remains in our beam at least $2\frac{1}{2}$ minutes, long enough for eight integrations. That also gives us a chance to economize by alternating antenna polarizations, rather than duplicating analysis hardware (as we do in the present system). Let me now turn it over to John Forster, who will describe the overall implementation of the 8 million channel analyzer, then I

will return with some hardware details and a wrap-up.

The META system, shown in block diagram form in figure 2, uses the front end from the Sentinel system described in our previous paper, together with a new data processor made up of 128 Sentinel-like Fast Fourier Transform (FFT) computers operating in parallel.

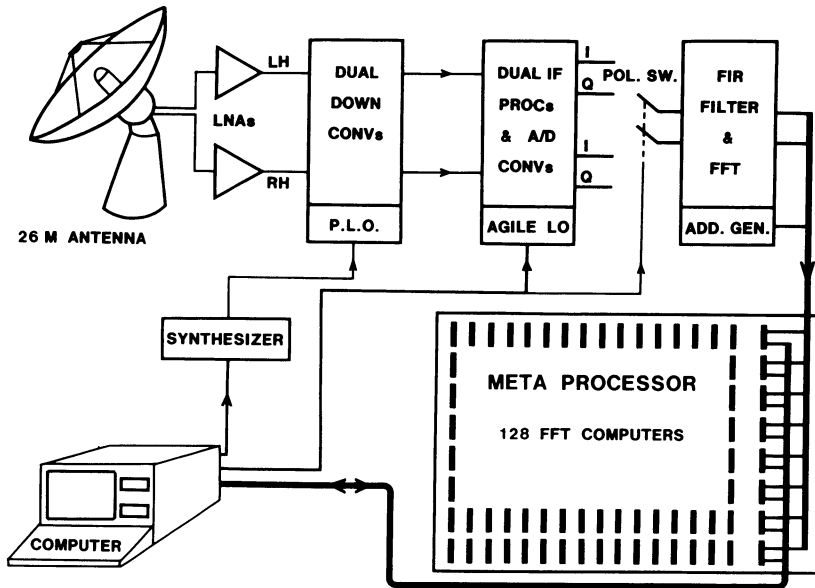


Figure 2. Block diagram of the 8-million channel narrowband analyzer. Both the control computer and the META processor array use the 68000 CPU.

Candidate signals from space are gathered by the fully steerable 26 meter Cassegrain antenna, located at Oak Ridge Observatory in Harvard, Massachusetts. At the feed horn, the incoming signal is separated into Left and Right circular polarizations and each polarization is amplified by a Low Noise Amplifier (LNA). These GaAs FET amplifiers were built with the help of Tap Lum at Berkeley, and have a noise temperature of about 45 degrees Kelvin when operated at ambient temperature, as in the present system. This amplified signal is fed to the Down Converters which translate the signals from a 420kHz wide band centered on the current favorite "Magic Frequency" (F_m) to a similar band centered on the Intermediate Frequency (IF) of 30.0MHz (figure 3).

The frequency translation is accomplished by mixing the incoming RF with the first Local Oscillator (1st LO) in an image-rejecting mixer. The use of an image-rejecting mixer permits the system to be frequency-agile, since very narrow bandpass filters are not required before the first frequency conversion. The 1st LO signal is produced by phase-locking the output of a microwave oscillator to a multiple of the synthesizer output frequency. The Phase Locked Oscillator (PLO) is necessary since we are searching for very narrow band signals, and phase




Figure 3. The META down-converter. Quadrature mixers convert the 30MHz IF to 420kHz baseband, which the NASA Bandpass Filter #2 then splits into 128 adjacent passbands, each centered about DC.

noise in the local oscillators directly contaminates any narrow band signals that we might receive. The operating frequency of the synthesizer is selected by the control computer such that the net Doppler shifts in the received frequency due to the component of the Earth's orbital and rotational motion relative to the chosen reference frame (heliocentric, galactic barycenter or big bang rest frame) along the current line-of-sight is cancelled out. This frequency is updated by the computer at times when data is not being taken for spectral analysis. The outputs of this first frequency conversion are amplified, filtered, and sent to the control room for further processing.

The IF processor takes the 420kHz band centered about 30.0MHz from the down converter and mixes it with a pair of 2nd Local Oscillator (2nd LO) signals, one a Sine and the other a Cosine. The 2nd LO signal is changed by the control computer to compensate for the changes in the Doppler shifts between updates to the 1st LO. The 2nd LO has a limited range (300Hz) but can be changed rapidly and without phase discontinuity, while the 1st LO has a much larger range (~1MHz) but requires considerable settling time. Thus this combination has the advantages of both. The voltage outputs of the mixers, the In-phase (I) and Quadrature (Q) channels, are then band-limited and converted to 8-bit signed integers.

Up until this point, there have been separate sets of hardware for

each polarization, the existing hardware scheme from Sentinel. Since the META hardware is much more costly, it becomes necessary to share it between the polarizations. The control computer switches the appropriate polarization to the preprocessor in synchronism with the data acquisition times of the META processors.

The preprocessor is the Number 2 Bandpass Filter from the NASA SETI program, which Dr. Alan Peterson described in the previous paper. In the META system it is used to divide the 420kHz band of frequencies centered about 0 frequency (DC) into 128 separate, equal bands, and to translate the frequency of each individual band by the appropriate amount so that all the 128 bands are centered about DC. The Finite Impulse Response (FIR) filter is used to preprocess the data to improve the adjacent channel rejection and to square off the bandpass shape, while the subsequent Fast Fourier Transform (FFT) algorithm is used to do the frequency translation. The data output consists of the first pair of data points (a complex quantity) for channel 1, then channels 2,3,4, and so on up to channel 128, and then the second pair of data points for channel 1, etc. The Address Generator takes the synchronizing signals from the NASA Bandpass Filter #2 and produces the correct control signals to route the data bytes to the appropriate META processor.

The META array consists of 128 separate computers, each with its own CPU, memory, and data bus interfaces, organized as eight crates of 16 processors each. Paul will describe the processors in detail later. Each crate also contains a crate controller card that recognizes addresses and buffers the data buses. One set of two data buses and an address bus are used to transfer data from the preprocessor to the META processors, and the other set is used to communicate with the control computer. Each of the META processors accepts data from the preprocessor corresponding to one of the 3.2kHz wide channels centered about DC, and uses an FFT algorithm to produce a 64K point complex spectrum. The processor then examines this spectrum and reports any noteworthy peaks to the control computer upon request. In the Sentinel system, the peak searching was done by the control computer, but in the META system that task has been reassigned to the individual processors to reduce the control computer's workload.

Because of the modest size of our SETI development effort we have tried to use proven designs to minimize the technological risks as follows:

- The Sentinel LNA's, Down Converters, IF Processors, and Control Computer have been used for the META system. In fact, minimizing changes to existing hardware and software has been an important design consideration.

- The NASA Bandpass Filter #2 will be used as a preprocessor, which relieves us of having to design a very sophisticated piece of hardware and software.

- The META processor is made up of microcomputers whose design has already been proven to work reliably in the Sentinel system.

Another important consideration in the design of the META system has been ease of manufacture and maintenance, since the system has over 20,000 integrated circuits and over 350,000 connections. To reduce

these problems to manageable proportions, the system has been designed to:

- Use identical processors throughout the META system. This results in economies in design, parts procurement, assembly, and testing, as well as the ability to repair the system rapidly by just swapping out a failed card with a known good one. The only other card with active components in the processor is the crate controller, identical in all eight crates.

- Use printed circuit backplanes and mass-terminated cabling throughout the processor to reduce the number of wiring errors and possibilities for cable failures.

- Permit test of the system components at the card level, before the system is assembled. In fact, both the printed circuit cards and the components for them will be tested before assembly. To reduce the likelihood of premature failure of the ICs they will all be stressed by operating them at elevated temperature prior to acceptance testing. Once in the system, the parts will be operated conservatively, because the system has been designed to minimize power dissipation.

Even with these precautions, it is certain that parts of the system will fail from time to time. The failure of a single processor should not cause the whole system to fail, and in most cases will just cause a 3.2kHz hole in the spectrum, since the control computer has the capability to ignore a processor that doesn't respond properly.

I should mention that John is a real Jack of all trades, and these beautiful slides, multi-color objects that you just saw were prepared in his home laboratory. In the remaining 4.8 minutes I'm going to describe the actual processor itself and how we are going to use it in the finished system. Figure 4 is a block diagram of the processor itself. It's really very close to what we are using now in the Sentinel Fourier processor, and all processors look the same these days anyway. You just hang on a bus all the things you want, in this case a 68000 running at 10MHz without wait states, which is equivalent to the computational power of a VAX 780 running in integer arithmetic, according to the Berkeley benchmark published in ACM. The communication with the control computer and also with the pre-processor, the NASA bandpass filter #2, is done through four byte-wide bi-directional parallel ports, individually addressable via a pair of backplane address buses. There is enough RAM here (384K) so that we can both buffer interrupt-driven incoming data and do FFT's simultaneously, so no time is lost. The program itself is stored in static RAM on these processors, initially downloaded via a monitor ROM. That's the way NASA loads their processors, and it makes it easy to re-configure software, download diagnostics, and generally do lots of good stuff. Here is the CMOS hardware 16x16 integer multiplier, which speeds up the FFT computation by almost a factor of three. We also put in a set of asynchronous flags, a dual serial asynchronous/synchronous port, and various lights and switches. The "bit-reverser" is a cute little port in there that lets you turn around an address and do that famous FFT shuffle without having to do laborious bit-shifts or table lookups in software.

I'm holding in my hand an actual processor board, 23.3cm x 40.0cm (9"x16"), a double-sided design with 140 chips. It seems to be quite

clean electrically, because we gridded power and ground quite extensively. This thing can do a 64K complex FFT in 14 seconds, and, as John explained, with a parallel processing array of 128 of these horses we do an 8.4 million point transform in 14 seconds. We use the remaining 6 seconds to compute power spectra and look for peaks. These things cost a little under \$500 apiece and are equivalent, as I have said before, to a VAX 11/780 when doing integer computation. The processor board, by the way, consumes just 5 watts, and can do other tasks besides FFT's, since it's really a general purpose computer.

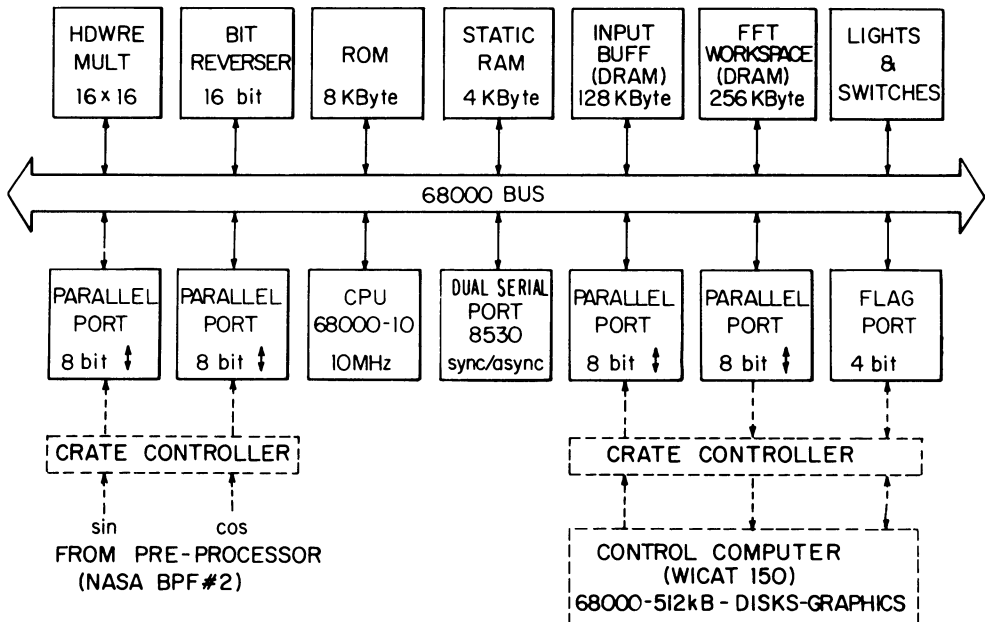


Figure 4. META-system Fourier processor block diagram.

For the final portion of this talk, I'll briefly describe the way the whole META system plays together. The software in the individual processors takes the preprocessed data (as generated by the NASA front-end preprocessor) and computes an FFT, then a power spectrum. We use the distributed processors to do peak searching also, since it's awkward to move 8.4 million 16-bit numbers from 128 individual processors into a central computer (that is itself no more powerful than one of the satellite processors), look for peaks, do archiving, and get everything out of there in time for the next 8.4 million numbers 20 seconds later. The central computer will simply poll these guys and find out whether they have found anything interesting. If they have, move some more data and turn on a quick-start instrumentation recorder as per Bob Dixon's instructions yesterday, namely, when you think you've got something, don't let it get away.

We are going to do polarization switching between runs rather than analyzing simultaneously two polarizations: Since we get at least 7 or

8 shots at any target going through the beam, it seems foolish to waste the processing power looking at two polarizations simultaneously, when frequency seems to be the space that is more important to explore. Likewise, we will do frequency hopping, to cover the best rest frames in which one might hope to find narrow spectral features at magic frequencies. We will also add to the software Woody Sullivan's very good suggestion that a telescope of this size doing a regular survey of the galaxy (and in fact everything else in the sky) would be foolish not to compare maps from run to run to see if there is a new supernova in the galaxy, or perhaps some other kind of flare source not yet expected.

In summary, let me reiterate the limitations of the META strategy: This is not a do-everything system. In fact, this is a do-only-a-little-bit system, but it does more than we now do. In particular, we are still limited to narrow band, i.e. CW, transmissions at magic frequencies only, because of our coverage of 420kHz. Again I would just like to point out by comparison that the developing NASA strategy will detect not only CW signals but also "chirp CW" and pulses, and also will not confine itself to magic frequencies but will cover the whole waterhole and additional frequency bands. The sky survey, in particular, will cover quite a large range of the microwave window. Our search is insensitive to signals with source chirps, or for pulses or binary phase shift keying, for instance. That is our primary limitation.

On the other hand, the META strategy removes the most serious limitation of our current system, namely that signals must be beamed at our star, precompensated for the relative velocity along the line of sight. The new META analyzer has enough bandwidth to detect signals transmitted without velocity compensation, or compensated to make them magic in the galactic barycenter or big bang rest frame. We can't do all these scenarios at once, but we can do them sequentially in successive integrations during a single pass of the sky. Our sensitivity stays the same but since we are doing 100 times more searching, perhaps we are entitled to increase by a factor of 100 our insult to Frank Drake here, and say we are doing one Mega Ozma per minute (100,000 years of Ozma per minute). Or rather, perhaps we should view the comparison with Drake's pioneering work as a celebration of the power of advancing technology.

We still cover 80% of the sky in a half year, but now at 420kHz coverage with somewhat broadened resolution (0.05Hz compared with 0.03Hz) in order to give us the bandwidth we need with 8.4 million channels. Finally, with a system of this magnitude it probably pays to cool the front end. We haven't done it at this point, because it is a somewhat of a pain and gets you something like 3dB, although as Frank pointed out 3dB is really better than it seems. As a result the system noise temperature will drop to 50°K. For those fiscal planners in the audience, the budget for the META upgrade is \$95,000, with an operating budget of \$20,000 per year, both provided by the Planetary Society. I thank you for your patience.