PANEL C

Frequency Stability and Its Interpretation to the Users

Panel Members

Andrew Chi	Goddard Space Flight Center
David Allen	National Bureau of Standards
William Lindsey	University of Southern California
Robert Vessot	Smithsonian Astrophysical Observatory
Leonard Cutler	Hewlett-Packard Company
Gernot Winkler	U. S. Naval Observatory
James Mullen	Ratheon Co.
Warren Smith	Bell Telephone Labs

PANEL C DISCUSSION

MR. CHI: I am delighted to have a special panel discussion on frequency stability in this PTTI meeting. The reason for me to consider this special session is the fact that you probably will know many of the faces, especially those who participated about 12 years ago when we had a special symposium on short-term frequency stability in 1964 at Goddard Space Flight Center.

It was organized on the basis that users generally need stable frequency sources, and do not know how to communicate with another group of people whose specialties were either to make the stable frequency sources or to maintain the frequency sources.

The communication was so inadequate that it became obvious that there should be some kind of conference to open up the communication and establish a common language so that the users can communicate with those who actually make the frequency standards. However, after a dozen years or so, the two groups could not provide any kind of feedback. It was just like a phase lock loop. The time constant was extremely long.

It brings to my mind that it was almost as if we were talking Greek to each other except no one really understood what we were talking about and very few really understood Greek.

Then after some time, we thought we made some progress; enough for us to communicate with the users. It appeared to me that instead of talking in Greek, we were talking in Chinese. Well, I thought that is very good; then I shall be able to understand them. However, it turned out they did not really converse in my dialect.

So what I thought would be appropriate at this time is not to try to impress any particular group of people of how much we know or what kind of theory we can use to improve on the past difficulties but to describe the actual situation. If we try to do this, the panel discussion will be able to establish a link of communication which is the objective of this panel.

However, if you look at the faces, you will notice a few people who are not on the stage – absent, for instance, are Jim Barnes and David Leeson who were not able to attend the meeting. In essence we have essentially the original group of the frequency stability subcommittee of 1964.

Now, I would like further to bring out the fact that in some of the past papers and also in the publications, you will see that the Subcommittee of IEEE on Frequency Stability was comprised to do certain things.

The present chairman of that Committee on Frequency and Time is Dr. Robert Vessot, sitting next to Dr. Winkler, and the Subcommittee chairman on Frequency Stability is Warren Smith, who is to my left.

In the future, should you have any more questions with regard to frequency stability and time, those are the two persons to whom you should address your questions and ask for solutions.

I would like further to make one more statement in regard to frequency stability. The point I would like to say is that when we observe the frequency variations or phase variations, we are looking at the imperfection of a system. What we can measure is the imperfection of the behavior of an oscillator.

What we seek is not the imperfection of the oscillator but ways to improve it. What we need is the best possible performance of the oscillator under the conditions that it encounters. What we want to describe is how good the oscillator is.

It is subject to noise processes. What we require is to understand statistically, how it behaves.

I would like to first allow a few questions to be raised on those papers which were presented this morning. We did not have time to answer any questions.

Are there any questions in regard to the earlier papers? If so please address your question to the author. identify the author and the paper so that author can answer the questions.

MR. CHI: Are there any more questions?

DR. WALL: Fred Wall.

I, too, had at least a comment on Professor Lindsey's paper. I thought he was just a little bit unfair in describing us by saying people who use frequency standards or make them, only describe them in terms of, say, the pair variance.

In fact, is that the IEEE Committees have recommended a number of measures, including a spectral density of frequency or phase fluctuations. And these fully characterize a standard method of characterization. They have all the information there. Further, there is the systematic effects Dr. Winkler described, quite clearly.

The other thing, the reason one uses the pair variance or Allan variance is to condense the data so one has a small number of numbers to talk about rather than 10,000 data points. Furthermore, there are a number of different regions. You might have one that describes, say, the shot noise process or fundamental noise process in the frequency determining element.

Then, as one includes various phase lock or frequency lock loops, one has other regions which are described by perhaps a different power lock. So, typically, in a short-term, one has tau to the minus 1/2 or tau to the minus 1 then a flat region and then some long-term systematic variations.

I don't think any of us have tried to say one describes an oscillator by a single number whether one picks an Allan variance to describe it or something else. I think most of us use the frequency domain notation if he wants to look at the details of an oscillator. If one wants to know what the phase spectrum is one looks at frequency domain, not the time domain.

DR. LINDSEY: I would like to comment on that. I apologize if I was not fair in my comments with respect to the utility of the power spectral density of the frequency process.

I well recognize that in the Abstract of the 1971 or '66 special issue on frequency stability this comment was made. It was stated either can be used as a measure of frequency stability.

However, in the circles that I travel and in the papers that I read, no one surprises me with that data. And in particular, I can quote a couple or three jobs in which I have been concerned where we had to specifically measure the power spectral density of the process in order to make an assessment of performance of the system.

What I would like to say is that I should like to see more emphasis on supplying the user with the power spectral density of the frequency.

If you give him only the Allan variance measurement, say between the interval .1 second up to 1000 seconds or whatever, he is given a curve. How does he use this in a certain particular application?

If he needs the power spectral density of the instability, then he has to make a transformation based upon the measurement and he only has a finite number of points so that means extrapolation of the Allan variance into regions where the measurements aren't made and which subject the model to error.

Consequently, in that transformation, in fact, for the most part, it has not been recognized how to transform from the Allan variance to the frequency domain. However, in cases where you have to select a performance measure, which contains a dead zone on the order of some minutes, an hour in certain cases, you need power spectral density in order to get to the zero point. The Allan variance does not provide you with that capability.

So I assert that, given the power spectral density of the frequency fluctuations, it contains all the information in accordance with the two sets of structure functions that I described. These two sets of structure functions contain generally the information that the user needs. These represent the performance measures for various systems.

MR. CHI: Dr. Reinhardt.

DR. REINHARDT: I would like to disagree with you strongly. The problem with spectral density measures is that the integrals are not doable.

DR. LINDSEY: Why aren't they doable?

DR. REINHARDT: They are doable numerically, but in many cases very difficult. The spectral density measurements frequently are very difficult in the zero to thousandths Hertz range. From the Allan variance charts from the spectral power law behavior you can usually refer to charts. And the National Bureau of Standards and others have published chart after chart to convert from one measure to another.

Now, I think that most of us will agree that some form of the Allan variance, whether the two, three or four sample variance, can be obtained in real time. You can write them in that form and use the variance weighting functions that are calculated and published to go from the two sample variance with zero dead time to sample variances with finite dead time if you know the spectral process.

DR. LINDSEY: I guess may comment to that would be, I disagree with you the integrals are impossible. If you look in tables of integrals, frequently, you will not find them. For instance, the flicker frequency effect produces a divergent interval if you extend it over all regions.

But in my opinion, integrations which are not found in tables shouldn't frighten the engineer these days. There are well known tabulations of numerical techniques which are available such as Gaussian quadrature techniques. These are quite sophisticated.

However, for experimental data generally engineers are not familiar with them. But we do have computers these days of tremendous power; there is no reason to be frightened by the fact you have an integral you cannot invert or do in closed form.

You go right to the machine and use sophisticated forms if need be. I have used these to get the inversion and get the number.

MR. CHI: I would like to exercise the prerogative of being the moderator. And what I would like to do is to set a time limit on the discussion of this paper to 12:00 o'clock by the clock on the back of the wall. So that we will keep that in mind so that we will be able to move in the discussion of other areas.

MR. CHI: I would like to recognize Bob Vessot.

DR. VESSOT: I would just like to remind people that, in fact, values of both the Allan variance and the power spectral density either of the phase or frequency are now being given almost as a matter of course by most manufacturers of equipment and that both which are complementary in a real sense give the information that I think is required.

The Allan variance, in fact, does not tell it all. We have many ambiguities at the tau to the minus 1 slopes which we can resolve either by tinkering with the bandwidth or adjusting the number of samples.

I think what Professor Lindsey has said is overstated. These transformations, in fact, are done. There is a way of going from the time to the frequency domain as Cutler and Searle, I think, have elegantly pointed out.

And at the risk of perhaps exposing my ignorance, there may have been transforms of this form in existence even before that time, but under another name. And I think this fact has since been identified. So I don't think we are that badly off. I think we are somewhat enlightened, but I think we could stand a lot more.

DR. LINDSEY: I hope I haven't been misinterpreted, maybe I have, but I don't think that I have claimed that anybody is off. I represent the user. And as a user and in the circles that I travel over the past two years, in the work that I have done and the applications of selecting an oscillator, the power spectral density has not been supplied to me. It has been something we have had to fight for.

I can give you a couple of examples and how I got involved in the subject, as a matter of fact. It was at TRW, I was doing a consulting job where we needed to specify an oscillator.

One of the problems that we had was that we had, of course, the Allan variance. They were able to make measurements of the variance. However, it was not the Allan variance I needed at that point in time. I was not able to transform it into the frequency domain.

It turns out that this spurred TRW to look into the problem of measuring in the frequency domain. As a matter of fact, I believe over the last two years, they have developed this technique to a certain extent in terms of measuring the phase noise power spectral density of the frequency which we have used. How accurate is it? I can't make a statement with respect to that.

But I am also aware of work in which other areas that I travel having to do with communication systems performance in these cases the performance of a system was not adequately achieved because of the fact the phase noise degraded the system performance and in an unknown way, not in accordance with the Allan variance.

In that context, the engineers required the power spectral density to evaluate performance. And it was not a function explicitly of the Allan variance.

DR. MULLEN: At some times, it is quite hard to find one thing and much easier to find the other. As a matter of fact, as Bob Vessot pointed out, there are some ambiguities with some dependences on the Allan variance. And now that we have got your formula for getting back to the other domain, we can get the inverse of the table that goes from the frequency domain to the time domain, allowing for the cutoffs and get the table that goes back the other way with cutoffs on the Allan variances. So that will be very handy when we have it.

Then, that may save all of us the experimental measurement problem if, in fact, we have found it easy to get the one measurement and find it difficult to get the other. But what usually happens, is that in the first three to six months after you get the job, you budget everything. And if you made a mistake, you find out about it a couple of years later. And then it is several years before the system, any other system, comes around again. And then provided nobody has been promoted, then they remember what they did the first time. But there is definitely a problem in keeping everyone aware of the best or of standard ways to use things.

I think there has been an ongoing creeping tendency to use Allan variances and spectral densities of frequency. And we have more commonality and more understanding of the problem than we had before. But there is no doubt that we have got a problem in getting the user community tied together with the oscillator-producing community.

MR. CHI: Well, are there any questions from the audience on other papers? I do not see any hands. Therefore, I would like to start the panel discussion.

I would like to start by viewing from the user's point of view rather than by trying to use the viewpoint of the people who actually work in the frequency stability or frequency generation field. Let us start from the more fundamental level that, when you want to use a stable oscillator, what are the questions one should raise.

I would like to list perhaps three questions out of which I would like to ask each member who may wish to comment to provide an answer to any particular question he wishes.

The first one is, what is frequency stability and how is it characterized.

The second one, what is time domain measurement and frequency domain measurement. And under what condition is each selected to meet the user's need?

The third one, how should a user select oscillators from the specifications provided by the manufacturers?

These are the general type questions which a user would have to answer before he can select what particular oscillators he should buy and also a certain amount of trade-offs.

I would like to obtain answers or comments and clarifications addressed to these three questions and I would like to start from Len Cutler on my right and then we proceed around the table.

DR. CUTLER: These are pretty general questions, Andy, and they have been asked many times before. We have come a long ways, I think, in getting answers to these questions. I think there is probably a great deal more distance for us to go before we will be able to answer them completely. In a way, that is completely satisfactory both from the standpoint of scientists understanding the fundamental processes that are going on, the design of equipment and the utilization of some of the basic models to improve the designs of equipment. We need to be able to specify the equipment in the terms of users and to understand the specifications of the equipment and how to optimize systems to make use of the equipment that is being produced at a particular given time.

The first question that you gave is "what is frequency stability," and how is it characterized. Well, presently, I think that I would tend to agree with Dr. Mullen that a great deal of the characterization is given in terms of the Allan variance.

This is a characteristic which tends to deal with the very long times or the very low frequencies, because it is very easy to make such measurements in that domain. And indeed, you can transform between the domains as has been pointed out many times and as has been very elegantly presented by Professor Lindsey.

If one has a complete characterization of spectral density in terms of frequency or in terms of phase (by the usual relationship, which may have some mathematical problems) nevertheless it can be used in all practical cases, and one can get measures of stability in the time domain from the spectral density and vice versa.

Generally, but maybe not necessarily, the best way, most characterization is done for long times and very low frequencies in terms of the Allan variance. I don't think very many people would disagree with that.

For short times, frequency stability is generally characterized by the phase spectral density or L(f) been commonly used in the last few years, which is the single sideband noise spectral density using the carrier power in a one Hertz band width.

This is a characterization many people have found useful for very short-term stability measurements and also measurements or equipment which involve noise powers out in the sidebands. That is my answer to your question number one.

Question number two is on time domain and how they relate to the user.

I think that is pretty well covered in my answer to question number one so I won't dwell any further on it.

The third question is how does the user select an oscillator from the manufacturer's specification.

That is a good question. And if the specification of an oscillator were absolutely complete and all users were well educated, the answer to that question would be they just do it. They go through the necessary exercises to assure that an oscillator specification in a particular domain of interest meets their system requirements.

If not, they look for a better oscillator or they bend their system requirements.

I think that is about all I have to say on those three questions.

MR. CHI: Thank you.
David Allen.

DR. ALLEN: I would like to respond to the first question of what is frequency stability and how is it characterized by saying that I think that we have seen a very interesting thing historically in the way it was developed.

In '64, as Andy mentioned, the IEEE and NASA Goddard sponsored this Committee on Frequency Stability. And it was obvious at that time there was a critical need in the community to be able to communicate what is stability. And there were many interesting papers given at that symposium.

The IEEE Subcommittee was formed. And out of it came some recommendations. And these have been, I think, quite readily adopted by most of the community.

The interesting thing that has happened, as Dr. Mullen has indicated, that it has provided a high level of communication between laboratories, manufacturers, and users. I enjoyed very much the paper of Professor Lindsey and found it very insightful.

I think, there is still a need in the field of communications that is not being covered, and we need to address that need. And I think this is insightful and helpful.

The thing that has happened in developing time domain and frequency domain measures is that I think there have been some errors committed that have caused problems.

Let me talk about a few of these when I talk about the second question. The first question, I think, has been already addressed by the subcommittee of which you were chairman at one time.

Later with Dr. Barnes as chairman, the committee put together a paper, actually two papers. One is a NBS technical note, and later a publication. I think that the committee many of the members of which are here, did an excellent job.

In regard to the precautions one needs to take, the time domain measure that is typically used, the sigma tau diagram is very powerful if you have power law equal densities. You can then readily translate from the time domain to the frequency domain and vice versa. This happens to be the happy situation of many precision oscillators.

Because it is the happy situation, the reality that models fit in practice, tends to make it a very powerful tool. This is true, I think, for three basic reasons. It is very simple to apply; it is insightful in how you use it; and you can understand in what you are doing in the process of measuring frequency stability in the time domain.

It has a sound theoretical basis for the power law spectra that are applicable. And, in fact, the subcommittee with Dr. Cutler really taking the lead, worked out the translations very neatly from the frequency domain to time domain.

So I think I would echo Fred Wall's concern that your presentation maybe was a little unfair in saying we didn't try to cover that ground in both the frequency domain and time domain.

Let me echo some concerns. If there is sideband structure in a spectral density, the time domain will lead you to problems. And this has been shown nicely by Mike Fischer in the previous presentation.

You get this funny looking sigma tau curve, and it is much more difficult to interpret that than if you do the analysis in the frequency domain. So really going to your third question and the difficulty that the user encounters, I think if one is looking basically (and this is a generalization that may have lots of flaws, generalizations always do) for sample times longer than one second, time domain measurements usually give you the information needed.

If you look for times shorter than one second or for Fourier frequencies greater than one Hertz the frequency domain is typically the best characterization.

MR. CHI: Would you repeat that just to make sure it got across?

DR. ALLEN: Okay. This is a generalization applying to precision oscillators and has some definite flaws in it, but typically for sample times longer than one second or for Fourier frequencies less than one Hertz, time domain measurements give you the necessary information. That is typically because usually there is no sideband structure there.

If you go higher than one Hertz in Fourier frequency or shorter than one second in sample time, very often you find structure in the spectra, and the time domain will be very misleading or hard to interpret.

And you will learn a lot more from the frequency domain analysis of the system whether it is phase of frequency spectral densities. I fully agree with Dr. Cutler, either one gives you the necessary information.

Another concern that I have is a point that Bob Vessot brought up earlier. And that is that we look at low frequency phenomena and try to classify them statistically. I think we can get into troubles there. If we can find the causal relationships, maybe it isn't the statistical phenomena we think it is. However, given such causal relationships it is much better to go and cure them, than to try to classify them and do predictions based on them.

So in the user sense, of course, if he sees a flicker flattening on the sigma tau curve, he has to live with it, assuming it has some environmental disturbances which causes that type of behavior. But for the manufacturer, that ought to be a real insight looking further into the heart of the problem. I think I have said enough.

MR. CHI: Bob Vessot.
Thank you, David.

DR. VESSOT: With reference to how the user should select an oscillator, I am reminded of a joke about porcupines' lovemaking and the answer is "very carefully" because one has to be, I think, not only wary of the manufacturer's intentions but really aware of our own requirements.

And I think too few people relate their requirements to the property of the oscillator and then go and get the proper oscillator. What they do is look at their requirements and then find somebody who is going to sell them an oscillator.

Naturally, the urge to conduct a transaction sometimes becomes more prominent than the urge to satisfy a need. I suggest that the buyer is really the man who has to do his homework. He has to understand what the properties of the oscillator are that he needs in order to accomplish the job he is doing or wants to do.

I suggest in order to do this, he makes models of how his system will work and that he applies measurements in the time domain if that is the way he wants to use it, if it is in the sense of a timing effort or in the frequency domain if it is in the sense of a spectrally purity as is so often required, for instance, in long baseline interferometry.

Then if he understands what he wants, he can go to the vendor and say, "I need an oscillator with an Allan variance like this and spectral behavior like that." And they better look decently relatable through the Cutler-Searle relationship.

And I think that is when the process of communication will have begun so that he will get what he wants. Personally, I would like to see a more complete representation in the frequency domain of the way data are taken in the frequency domain.

You realize the remarkable completeness of the way in which the time domain things were worked out. The papers that began in 1964 and even much earlier by Barnes, Allan and others, were the result of a very important requirement to understand what clocks were doing at the Bureau of Standards.

And naturally, I think they were led to making their analysis in the time domain. It led us to have this realization that what we saw depended on things like the number of samples, the dead time, the band width (although that came a bit later) and the averaging time; all these parameters have their analog in the case of a spectral analyzer.

What comes to mind to me is that if you run a spectral analyzer and are out to give a spectrum, I think you ought to say what the rate of the sweep, what is the band width of the sweep. After all, that tells you the resolution of what is going to happen, and the duration of the time segment that you actually performed this operation on.

These are exactly analogous to the behavior of the time domain analysis we have grown to understand as the sigma tau plots of all these parameters. I think the spectral domain is likely to be more and more important as the frequencies go higher and higher because, as Dave Allen says, one has a hard time doing a time domain representation of a laser signal.

I also would like to point out that the measurements that we are considering in terms of time and frequency domain are going to be replicated in our discussion of length measurements some day and that I can see a whole new committee coming up with the same problems of definition of length in view of the fact that they would like to relate it to something that is physically available.

In my opinion, the greatest precision now is available in the field of atomic standards. So, somewhere along the line, there is going to be an Allan variance with the determination of an object's length or distance.

MR. CHI: Thank you.
Dr. Winkler.

DR. WINKLER: I would like to read for the first question the definition of frequency stability from NBS Technical Note 679 by Dave Allen.

Frequency stability is the degree to which an oscillator signal produces the same value of frequency for any interval throughout a specified period of time.

That is the definition – degree out of which the same frequencies are produced. How is it characterized?

You characterize the deviations from a process by the same way as you characterize any random process. We define randomness.

What is randomness?

Randomness, you have a random process or random signal. If there is no correlation between the disturbance at one moment and the next one, this would be pure randomness.

As you leave that area, there is a completely continuous transition to complete determinus where you have no randomness at all.

In between, then, you have noise steps which are increasingly more internally correlated. That is the whole secret. And we do not have to reinvent the wheel.

I think statistics, the science of how to describe such a random process, has been in the forefront of modern science and technology. But what we are debating here is, insofar as time series are concerned, an attempt to reinvent the wheel. And I don't think it is necessary.

A correlation function, and a spectral density function, are essential characteristics for not completely random processes. In a completely random process, all you have to do is give the mean and the variance.

In view of that, I can only repeat that the only problem which we have here is how to place that cut between random description and deterministic description. That has to be somewhere in the middle. And we can debate that.

And I can only repeat what the other experts have already said. What is time domain and what is frequency domain measurement?

Time domain measurements are obtained through sample time. You sample a quantity, you generate a time series that you do by phase sampling, or you sample frequencies by measuring frequencies over a time interval. But that is only a distinction in the measurement process.

You measure in the frequency domain when you directly determine the side band power in respect to the carrier frequency power either of the variations of phase, where you establish a variance of the phase fluctuations, and you investigate its behavior versus frequency, you can do the same thing in frequency variations. But that is a distinction which is based upon the measurement process. And it is not identical to the distinction which is based on the language in which one states what one has measured.

And here, I must most emphatically remember one should use the same language in which one has the requirements. It would make no sense to insist, for instance, on a spectral density specification if your requirements are purely in longterm timekeeping.

You don't learn very much from that and vice versa. There is no point in using the Allan variance if your requirements are in applications of coherent Doppler radar, for instance.

So the language specification ought to be the same in which your requirement exist.

And so I come to the last thing, how does a user select equipment.

Amplifying what you just said, Bob. I think for an important project, for an expensive project, to select anything just on the basis of a piece of paper is sheer madness. The larger the project. the more important it is to set up a little pilot thing yourself. Get a test bench somewhere, spend a few hours, get a little cheap thing, and play around with it. Then you will understand what you need.

And we have seen in most cases the greatest problem is to bring the user really to define what he needs. I begin to use the term "user" in the same way as Congress uses the "taxpayer."

This is not accidental, incidentally, because that is where the money comes from.

But I cannot emphasize enough, that for systems design and for specifications and for putting everything together, if you try to save investment for a little initial experimental effort, the study effort (and now I am talking about bench work) you will pay dearly later on.

MR. CHI: Thank you, Dr. Winkler.
Jim Mullen.

DR. MULLEN: I would like to speak to the question of how to select oscillators, too. Certainly, the first thing is to decide what the system needs. And in that respect, some of the consolidated results that Professor Lindsey has shown to us will be very useful.

The fact of the matter is there are lots of oscillators that are needed and in many systems these turn out not to be critical. Often, the budget isn't big enough to do very much really to solve the problem of how to do the experiment, which would be much the safer thing to do.

And so somebody of junior level in the program goes out and buys the oscillator. And if anything is wrong, the money is already spent. So it is really important to be able to try to estimate what quality you have to have oscillator with some kind of consolidated or overall estimate of how much effect the oscillator performance will have on the system and then to be able to interpret specifications of available oscillators. The question is which standard descriptions are useful.

It often turns out to be the case that you can find the oscillator that meets all the environmental requirements. But unfortunately, often the measurements that have already been made, are made in the wrong domain, and you have to go back and forth. Either you can call up the oscillator manufacturer and tell him you want one, but want him to measure the whole thing all over again in some other way (in which case he doesn't have that much interest) or else you have got to be able to convert it yourself.

If the oscillator is critical to the whole system, then you generally get the opportunity to do a decent job. But otherwise, there is a fair amount of hard work that has to be done without completely fundamental digging at the problem.

I think the standardizations that we make are not quite capable yet; but not too far from complete and what we have is going to make the problem a lot easier in the future than it has in the past.

MR. CHI: Thank you.
Warren.

DR. SMITH: Well, in the interest of time and also due reference to the fact that almost everything that can be said in answer to your questions has been said, I will keep my comments very brief.

It has been very gratifying over the past 10 years or so to see a great increase in the ability to communicate about clocks and about oscillators and frequency standards. I think things are infinitely better than they were in 1964. And in that respect, I think we have accomplished something. It has been very interesting to see the presentation of data here today.

I was particularly interested in some of the slides shown by Dr. Winkler. I would like to make a brief comment.

If one takes the pathological data presented for the cesium standard this morning by merely changing the time scale and, of course, the frequency deviation scale, one sees a perfect picture of a pathological quartz crystal oscillator with which I am much more familiar, having worked in that area a great deal.

But the point of interest that I would like to make here and I think will be of general interest, is that we find in quartz crystals on almost continuous variation in the magnitude of this kind of pathological behavior, which is usually characterized in a plot of raw data as frequency deviations with time that are quasi cyclic.

If you look for any given day, it looks almost like it is sinusoidal perturbation. And it has been possible to take those very bad offenders, open them up, and physically see defects usually in the form of badly adhering plating or small contaminants, semiattached to the surface of the quartz. This kind of behavior is pathological. The point is well taken, that when you see this kind of behavior you should be aware you're not really dealing with statistics, you're dealing with something that is out of the ordinary and abnormal, you should really back off and think.

I might say one word about your last question, the selection of an oscillator, the particular specification. I have been in the unique position of being bitten on this subject many times, and I would only caution that it has been my experience that most catastrophies in this area arise from a lack of knowledge at the beginning of a project as to what the specifications and requirements really are.

And you usually find out what you really need in a system when it is just too late to do anything about it. And it has little to do with the problem of communications between user and manufacturer. The great difficulty, particularly in sophisticated systems, is of assessing the whole question of requirements of your signal source in clocks at the beginning.

MR. CHI: Yes, please.

MR. TURLINGTON: Tom Turlington, Westinghouse Electric in Baltimore.

I have heard a lot about Allan variance and spectral density measurements this morning. There is also another, Hadamard variance, I used in the past for getting in very close to the carrier, less than 20, 30 cycles (all the way into tenths of a cycle) on oscillators that have clean spectral densities. That is, no discrete sidebands.

I find that to be a very useful technique. I haven't heard much about it here today. Do you care to comment?

MR. CHI: Anyone who would like to comment.
Dr. Winkler.

DR. WINKLER: You're absolutely correct. But it is particularly useful if you operate at a time constants for integration time where you have a high power spectral density. If that happens, you must use filtering. The transfer function of the Allan variance is very broad. So you have more dangers in that aspect than if you use the Allan variance.

I consider these may be variations of one and the same thing. Your point is absolutely correct.

Also, I think the so-called curvature variance may find some useful applications because it is insensitive to frequency drifts. And it is not more difficult to compute than the standard version.

For a calculator today, this is not a problem. You're absolutely right, but, of course, we cannot go into every detail here.

I think the least we could do, and I hope we have accomplished that, is to give some idea of the complexity of the subject. And I can only remind you that in order to really go into details, you have to get into details. There is no king's way to success.

Of course, there is an enviable way how to go round if you have no knowledge – simply buy the most expensive. And that's being done very effectively. But I don't think it is engineering.

MR. CHI: I would like to hear some questions from the users' viewpoint rather than trying to go into the theory of modeling or frequency stability. Question, please.

MR. KAHAN: Kahan, IDC.

This will be semi-user point of view. I am still worried about of characterizing oscillators. Assuming I have a transient, for example, a "Burster" variation and gather transient recovery and frequency as a function of real time, is there any way to present this data aside from reams and reams of data as a function of time in terms of a few parameters in a time domain or frequency domain? Or does it make any sense to characterize a transient response by Allan variance or whatever you want to call it in that sense?

MR. CHI: David.

DR. ALLEN: As a part of the work of CCIR, right now, they are trying to talk about questions of that nature. And I think that a point Harry Peters made and has been alluded to by Dr. Winkler at this conference is an exceedingly important one here.

That is, when you have nonstatistical phenomena, you should characterize that as a coefficient function. So much frequency change with certain radiations. So much frequency change per degree C in a certain range, that you should establish these coefficients that you might model these deterministic phenomena rather than trying to do it statistically.

I think it would be a mistake to use the Allan variance in such situations.

DR. CUTLER: I would certainly agree with that. What I would think would best do for that sort of thing would be to try and model it. In other words, you would assume that you have a ideal oscillator and this ideal oscillator under the burst of radiation undergoes a rapid frequency change.

Then it may decay with an exponential or some other law to some other new frequency. And if these things were completely characterized in a deterministic sense, then you would put back the other things, the random variations and so forth. And indeed, some of those random variations may be modulated by this transient effect.

Those modulations can again be a deterministic things that are applied to the random characteristics that are associated with the oscillator under normal conditions.

MR. CHI: Dr. Winkler.

DR. WINKLER: I have not gone into the details, but you will find them in my paper. There are three classes of functions which one can just empirically apply to modeling. And, according to which one you select, they behave differently.

The three classes are, number one, polynomials of degree N that we discussed this morning. Number two, Fourier series. If you have any periodic phenomena, the Fourier series, of course, is the thing to use. Number three, exponential functions. Dr. Cutler mentioned that. In general, it would be a sum of exponentials. If you have several different phenomena interacting over time contents, you could end up with a sum of exponentials. All these transform into another function of the same kind under a transformation when you shift the time axis, but only the first one, the polynomials, have the additional property that they will remain polynomials if you transform time with a different scale.

So these are details which are completely covered in books like Digital Analysis. There is an excellent book by Hamming, 1962, which has an excellent discussion on that.

But here, again, you have two things to consider. You consider the black box approach where you just phenomenologically describe what you see, but you must remember that many of these are parameter-dependent. So one really has to consider whatever function one uses as a function of the excitation, radiation, temperature variations, and what have you. Then you end up in the simplest situation with coefficients which you have determined previously - temperature coefficients or magnetic field coefficient.

MR. CHI: One more question from the audience.
Tom Healy.

MR. HEALY: There is I think a larger number of semi precision oscillators in the world rather than clocks. And there is a big market for them.

Usually, the user, when he is faced with an Allan variance or so forth, it can't really transform these into its system. He knows what his system requirements are. Usually, in many systems, communications and radar and so forth, the script L of F is a much more suitable description of an oscillator for the user.

And the other fact, is tied in with his environment. Usually, when there are thousands of oscillators we are talking about systems where the environment isn't as nice and benign as in the Naval Observatory Laboratory. Temperature, vibration, shock, all these causal phenomena, are very important. So his linear coefficient should be specified. And there should be some specification of the spectral density of frequency for offsets greater than one Hertz. And the thing is whether it should be $S\dot{\phi}$, $S\phi$ or L(f), that is a question.

Most people I have dealt with, are system designers who can more easily determine the script L of F because they have to convert all these other things in order to be able to incorporate the data into their system.

Another thing I would like to point out is the Hadamard variance isn't a cure-all of evils anyway because the function does have side lobes. And it can land right in a bright line and louse up the observation. So you have to be very careful.

MR. CHI: In the remaining few minutes which we have, I propose to assume that the user now has his oscillator. While he has it, presumably, he can separate all the systematic errors out and find the performance of the specification or whatever information he obtained.

The question that should be raised, in my view, is how much confidence he should have in the oscillator performance. Now, if he wishes to make tests, should be completely redo the specification test or should he make nominal tests.

And then, once he has established a certain amount of confidence in the oscillator which he owns, the problem is that when he uses it, it will be in the real world of systematic variations which consists of transient for voltages, temperature.

Can he actually predict the performance such that he can make sure that the frequency which he obtains will fall within the band which he is allowed to stay? In most activities, this is a real world situation.

I will not go through every panel member to answer all these questions, but whoever wishes to answer any or all or one question, please raise your hand. And we will use some of the lunch hour, if we may, for about 10 minutes or so. We will terminate the panel discussion at 1:00 o'clock.

Dr. Winkler.

DR. WINKLER: Confidence in your data. That is the reason why one must give the number of measurements or one must give the confidence interval. Remember the slide I showed on the probability distribution function?

There were two lines. These were the 95 percent confidence. That is a phenomenological side. When it comes how much confidence can you place in an actual application, you must derive that from the range of environmental conditions and environmental sensitivities which you expect. That's all.

MR. CHI: Warren.

DR. SMITH: Just a short statement. It seems to me that we have spent most of our time here talking about the difficulties in assessing the random behavior of signal sources. In the real world, I agree with Mr. Healy, that the problems are primarily those of the causal or deterministic effects.

Temperature coefficients, voltage coefficients, shock and vibration, all the rest of it, are the things that make real world specifications extremely difficult and which all of you and all users need to keep well in mind. These are the things that one really has to be careful in tying down for any particular application.

The discussion of random variables is much more interesting and can be treated in much greater depth. And in particular, as it applies to precision sources and clocks and timekeeping it is probably the fundamental one.

But in the field of communications, and the things that go in the field, I would caution you that the toughest thing to do is meet your deterministic and causal specification.

MR. CHI: Len.

DR. CUTLER: I agree. Yes, I agree. And I would like to add one other note of caution.

Very often, one must be concerned with rates of change of temperature, durations of shock, and things like this. Very often, these things are extremely difficult to specify or put limits on.

I caution the users of such situations; that in many cases, they may have to make their own very specific tests or request the manufacturers to make very specific tests involving rates of change of things.

MR. CHI: Anyone else? Anybody in the audience who may want to make any comment?

DR. KLEPCZYNSKI: I would like to ask any member of the panel if they feel that development of crystal oscillators is far behind in terms of its ability to function in a deterministic way, i.e., in an environment that is somewhat hostile. And compare that with the development of the oscillator in terms of low noise levels.

In other words, has the crystal oscillator been developed to a degree that is much, much better in terms of its noise floors and various other noise properties in proportion to its performance in an environment?

MR. CHI: Warren Smith.

DR. SMITH: I will make a brief statement on that point. Crystal oscillators have been around a long time. A lot of work has been done to achieve low noise floors. The name of that game as in any other game that we talk about is power.

The higher the power that you can dissipate in the crystal unit, the better the signal to noise that I can come up with. Unfortunately, quartz crystals are still the same. The other performance factors degrade as you increase the power level in the crystal. So you're back to the same trade-offs.

I would say there has probably been some improvement over the years. I don't know of any, I personally have not been in contact with any great breakthroughs. However, the crystal oscillators are still very widely used, especially in the communications of radar-type of application.

MR. CHI: David.

DR. ALLEN: I would like to add to those comments. We have one of our people from NBS here who is involved with crystal studies that we are doing. I see this as an extremely exciting field right now even though crystals have been around for a long time.

We see already some breakthroughs that have occurred with some commercial products. And in fact, some of the work we are doing is directed toward some significant breakthroughs, both as to environmental insensitivity and low noise.

Just to throw out some numbers, we have hope that even in long term, one might have a crystal oscillator that would exceed in performance the stability of rubidium in atomic devices. Whether that is true in a harsh environment is another thing. Those are separate problems.

MR. CHI: Lauren Rueger.

MR. RUEGER: I would like to come a little to the defense of the system engineers who overdesign their oscillators in systems. Part of that is because we find these beautiful characteristics you get when you deliver the oscillator, but what do you have one, three, five, 10 years later if you put it in a place where you can't get to it like in orbit?

A little extra margin is a pretty important factor in the system. It has been our practice in general to try to design or ask for specifications far enough over the margins needed so that as time goes on and you get radiation from the natural effects, for example, reduce the gain of the transistors or it may cause some other drift or other aging characteristics, you're still in business.

So I would like to defend in the design initially, you would like to overdesign so that you will retain all these margins.

MR. CHI: Dr. Winkler.

DR. WINKLER: I cannot agree more with you. My comment should not be misunderstood. I think there is a very important difference between specifying something out of ignorance versus application of a conservative safety margin.

I would be on the side of the conservative approach.

MR. CHI: You have a question?

MR. BRESHON: John Breshon, University of Maryland.

I was concerned about your statement about crystals. And are there some other types of crystals? Are quartz crystals considered for these clocks?

DR. ALLEN: I think some of the very fine work that is being done in new crystals is being done outside of this country even, and it is still quartz crystal as far as I know.

DR. VESSOT: I am not in the quartz crystal game, but I have heard of sapphire crystals of considerable size that have enormously high Q from the mechanical point of view, far, far higher than we would expect from quartz. They have the difficulty that they require a somewhat more eloquent approach in order to communicate with some electrical circuit.

However, I, too, wonder whether or not there isn't some opportunity to look for better mechanical oscillators. The sapphire crystals I know of are made in Salem, Mass. Dr. Fred Schmidt of Crystal Systems, Inc., makes single crystals that are of the size of a basketball. It is really extraordinary. These crystals are being used by Prof. D. Douglass of Rochester University for gravity wave detectors. I expect we will be hearing more about these in the future.

MR. CHI: Dr. Winkler.

DR. WINKLER: I believe (and I am really speaking with very little definite knowledge about this) that other than quartz is a very unique material. And it is my feeling that it will be hard to find something much better to make further improvements in the quartz crystals which are under development right now. This will likely have to do with improvements in the circuitry, and so on. There are, however, two further aspects. One. Either one could put the quartz back into the cryogenic environment (which was done at NBS 20 or 25 years ago the first time) and immediately one get much better mechanical performance, or we could consider the super conductive cavity again (which is also an oscillator which is not based on the quartz resonator, but nevertheless based on mechanical stability), or one could refer to the NBS development of ammonia; a relatively inexpensive ammonia-controlled oscillator which I think is a very interesting development.

So there are a variety of things cooking at the moment. And I think what one bets on will depend to a great degree on what one has available. If one can tolerate a cryogenic complication, I think that may be something to look forward to.

MR. CHI: Unless you really want to look at some earlier activities in this area, there is a book by Warren Mason who was formerly with Bell Telephone Lab, presently at Columbia University, I think called Sonics and Ultrasonics. You will see a whole list of materials which were examined as possible material in this field.

I would like to use the remaining minute to express, in behalf of the Executive Committee, my thanks to all the authors for this session who have given such excellent papers and the members of the panel.