PANEL DISCUSSION ON WORKSHOPS

Moderator: Raymond L. Filler
US Army Research Laboratory

RAY FILLER: Welcome to Part II of the audience moderator discussion which occurred yesterday. Today we're going to have our three session chairpersons (one is missing in action) give us a brief summary of what transpired at their session yesterday. Then for the rest of the time, we'll have audience questions. We're going to start with Joe White from the NRL whose session was entitled "Real Time Automated Systems."

JOE WHITE (NRL): We had a good crowd yesterday, we had about 30 or so people, pretty much a roomful. And we started off trying to define what a real-time automated system was, and basically came up with this kind of thing – that it was system that provided time or frequency, or both, to the user specification actually in real time; that it might include some sort of a historical calibration feature; but that basically what he wanted, he got out of the spigot right when he asked for it.

The other thing about the automated part, in particular, was there was not a frequent operator action required. In fact, in many cases, there wouldn't be an operator around it at all; we talked about fully-unattended and remotely-controlled type applications. The applications of these systems would typically include things like national time scales, remote time stations, and, as embedded pieces of equipment in military systems, telecommunication systems.

The class of performance that we were looking at for these systems, as far as time went, was on the order of 100 ns or better time accuracy; frequency accuracy to at least a part in 10^{11}; and again, this depended with some of them being as good as part in 10^{14}; and frequency stability, ranging from hydrogen maser systems, like a radio observatory system, to parts in 10^{13} at a second to other systems that might only be in parts in 10^{13} at a day. The other factor in this performance was that we required a synchronization to some national standard, or at least some network standard, and usually by a GPS or two-way time transfer measurements.

When we talked about the measurements, one of the things that came out that people thought was important there was that the measurements be accurately time-tagged when they're collected. Those of you that played with these systems, particularly things run by PCs, know that those time tags can often be in large error. And we talked about means of doing that, including having a hardware clock in the measurement system that provided very accurate time; or, alternatively, using one of the telephone or network time synch mechanisms for the control computer to keep it on time to the millisecond range.

Naturally, we all wanted nice quiet, unambiguous measurements, and we decided, in general, that meant making time measurements – or frequency measurements, I should say – at

5 MHz to get the smoother performance there. While one pps measurement was certainly necessary for things like GPS measurements, two–way time transfer measurements, in general, there were a lot of problems with those, as far as having a clean pulse to measure, establishing the right to triggering levels, the effects of long cables, those kinds of things.

We next talked about distribution systems, and we started off talking about the effects of the local environment on the distribution; that is, that the temperature, humidity, those kinds of things, often had an effect. The other thing that went with that is having a good way of connecting to it, that the connectors that were used and the types of cable were very important to achieving a good distribution, that just the distribution amplifier alone didn't really cover everything. We were typically looking for isolation of at least 100 dB between ports, and also 100 dB from output to input, which we have seen some systems not doing.

The other thing that was kind of interesting in distributions, we talked about widely–distributed systems, for instance, a communications network where the real–time automated system wasn't two racks sitting on one site, but a rack here, and a rack 100 miles away, and another that really is – in the terms of the way that system worked, really that was the system that they wanted to have as a real–time automated system. So sometimes the whole interconnection and distribution gets to be a pretty large problem.

From there, we went to software, or actually, robustness, which got us to software pretty quickly. Sam Stein gave what I thought was a nice definition of robustness; and that is that the small error in the system caused only small problems to the system operation. For instance, losing one device in the system shouldn't cause it all to die. That got us immediately to computers, and we decided there that you really need both stable user software, the specific software you wrote to make that system work, and stable underlying operating systems for the computer itself. A lot of times that's UNIX or OS–2, or something like that; that there often was great peril in changing versions of operating systems that ran the whole thing.

Also, in the robustness area, we talked about the trade–off between single point failures and the things that you do to try to avoid single point failures; there is a point of diminishing marginal returns as you add more and more redundancy and put in the switches to put the redundant sides together, that often you actually got to a system that was worse than what you started with; and that one of the solutions to that was to encourage your user of the system, the people that take the time and frequency outputs, to design their systems to be tolerant of small glitches; so that you really had a robust system in total, not just in the time and frequency part, but also in the piece that used the time and frequency.

We ended the robustness part with trying to define how you put robustness in the specification. And I think we came to the conclusion it was difficult to define that. There are really two problems. One was that you had to define what the users environment was, because what was robust for one environment may not be robust at all for another. And the other problem was that it's awfully hard to think of everything that can go wrong. You try to come up with very blanket–type statements that will cover everything; and when you field the system, you almost always find out there is something you left out. So I think we wound up agreeing that we had a difficult problem that we didn't quite know how to define.

We ended up talking about maintenance and testing. The general consensus, as far as maintenance went, was that we thought that systems should be maintained generally at the box level in the field; that the modern hardware is simply too complex to deal with in the field; that no matter how well you train your technicians, it's very difficult, it's very expensive; that, in general, you ought to have a lot of spares and rotate them around and let the manufacturer or at least some highly-trained depot deal with most of those issues. To support determining when we had problems, we talked about built-in tests; and also, about a remote diagnostics capability.

That's pretty much it.

RAY FILLER: Thank you. Next, we'll have Dick Sydnor from JPL. His session was entitled "Real World User Requirements."

RICHARD SYDNOR: None of us seemed to know exactly what that title meant, so it took a little bit to get the thing going and we sort of wandered over a large area.

The first part of the discussion was sort of a déjà vu; we have talked about this many times in the past, and it's the problem of communication between the supplier and the user. We had a number of examples of a user having incomplete specifications. He forgets that he's going to take the spacecraft oscillator and launch it. So it has to have a shock and vibration specification, and he's left that out. Then he comes and says "Gee, it broke." That kind of thing happens more often than you might think.

Also, on the other hand, sometimes the oscillator or frequency standard supplier doesn't have a really complete set of specifications in his catalog. He doesn't say what effect vibration has on phase noise, for example; so sometimes it's difficult to figure out exactly what this particular item is going to do in your environment.

It was suggested that the supplier whongets a set of specifications from a user should question those requirements. He knows more about his oscillators than the user does probably. And if something looks a little bit awry, then he should question that and find out if the user means what he says, or if he has left something out. Many times the user is not very familiar with the oscillator and how it works, and its problems. And so there is a misunderstanding of what some of the specifications need. So there is a need for user education.

But who is responsible for that? That was kicked around for quite awhile. And John Vig had some comments about availability of literature that would outline tests and give information to the user. Some users say there is no information out there. And it just means that they haven't really looked very much.

I think the best suggestion, but probably the hardest to implement in that area, was that the supplier should be involved in the procurement from the very beginning. And that's a little hard to do with the present legal situation where you have competitive bids, how you get all these suppliers involved in it. But still, it looks like the most logical way to handle some of those problems. Those problems have been discussed many times in the past, and no solution has been forthcoming as yet.

Then we sort of wandered away from that area, and we started talking about problems, various specific problems in terms of, say, distribution systems, time delay variations in cables, fiberoptics, how you stabilize fiberoptic systems, good connectors, that sort of thing; how you make sure that if you have a large network and you distribute it in time to, say, a bunch of people that are all various distances away from your main control clock, how they all have the same time, rather than varying all over the place due to the length of the cables. We had quite a bit of discussion on that.

Somebody asked what do the margins mean in a specification; and there is 90 percent probability that it will do such–and–such. Do people really understand that? I think the answer on that one was that nobody really knows exactly what is meant by that margin statement, and most people would rather have a specification that says it's guaranteed to do no worse than such–and–such.

There were some comments about various problems with crystal oscillators. It was brought to our attention that crystal oscillators stored at a very low temperature sometimes comes back out of that as a completely different crystal oscillator than the one you put in. There are aging rate changes and everything else.

That pretty much handles it. We had a large group in here. I would say the room was half full. But we had only five or six people that really contributed. Thank you.

RAY FILLER: I'm sorry that our third session chairman is not here. But if anybody who was there wants to make some comments, that's fine.

We're going to open the floor now to anybody for questions, comments, discussion of any sort, on this topic or maybe any other.

GERNOT M. WINKLER (USNO): It may be useful to elaborate a little bit more on your comments about margins and specifications. It's a problem which comes up over and over again; and that is that a system, whatever kind, has certain system performances; and then you have accidents. The two come from different distributions. And I think they should be separated.

It makes no sense to include accidents in a system specification; if you separate them, you can put a limit on how many you will tolerate per year, or per month, or whatever. But the system should be characterized after these accidents have been separated; because otherwise, you characterize two different processes with one number.

RICHARD SYDNOR: The margin discussion would have more to do with things like radiation exposure; after a certain number of rads of radiation, the probability is ninety peercent that it will be within a certain range. That sort of thing is typically what you get with radiation exposure, for example. The specs you see in manufacturers' catalogs on something says, for example, at a second, a part in 10^{13}. To me, that means that it's no worse than that, under any condition. A benign environment, obviously.

But if you are talking about systems, then you have to know not only, say, an upper limit, you have to know what the spread, what the distribution of the things are. And that's not in the manufacturers' catalogs. And many of them probably don't even know what it is. Some manufacturers will supply that information, if it's available, and they give it in terms of a histogram or something like that, a performance of the different ones that were produced. And that's essential if you're doing a system design. But that wasn't discussed during our meeting.

DICK KLEIN (LOCKHEED AT KENNEDY SPACE CENTER): One of the things we've noted with more than one vendor, they'll take the specification, particularly a short–term specification of an oscillator, and publish it as the short–term specification of the GPS receiver, ignoring the pertubation of the circuitry within the receiver itself. And we found that to be a problem in more than one vendor. Particularly one problem, you could almost see a IRIG A on the 1 MHz output. And it turned out that they were able to correct it. But apparently, it wasn't tested at the factory, only the specification that the oscillator manufacturer gave.

JOE WHITE: I think that happens.

FRED WALLS (NIST): One of the limitations and specifications for almost all oscillators and synthesizers, and things of that sort, is a lack of specification for AM noise. And in many system applications, it is the AM noise that limits noise floor for residual measurements on amplifiers and other things; you have AM to PM conversion in your amplifiers and on mixers and on non–linear things. You can have two oscillators with the same phase noise, and yet different AM; and one will work and one won't work. And so, we need to raise the consciousness of both manufacturers and users to insist on AM noise specifications.

RICHARD SYDNOR: That's a good point. Many manufacturers don't even know what the AM noise performance of the oscillators are, because they measure just the phase component and not the AM component.

JOHN VIG: In our experience in the Army, many of the problems that come to us originate from the fact that people who are assigned the job of writing a specification, and this often involves major systems – people just sit down and write specifications in isolation, without regard to what's been written before; and they invent their own definitions, invent their own way of measuring certain parameters for which others have already worked out the details. For example, Ray came back from a meeting recently on a major radar system. He was asked to review the specification for the oscillator, and he found several things that were just basically wrong with the specification; one, of which, was that a frequency of zero –

RAY FILLER: Yeah, a frequency of zero. The frequency aging specification was plus or minus F zero, I think, or something.

JOHN VIG: Yes, totally nonsensical specifications are being written by people who don't know what they're doing. And this is for multi–billion dollar systems. So I think the manufacturers probably could perform a service by including in their literature a list of existing specifications that people could at least start with. There are IEEE specifications, there are military specifications, there are IEC specifications; we have a set of definitions in a CCIR[1] glossary. That means they are all internationally recognized and accepted documents.

If somebody has a job of writing a specification, it's so much easier to go to the existing document and just call out a paragraph of an existing document rather than to sit down and scratch your head, 'How should I define 'aging,' how should I define 'phase noise?' " and

[1]International Radio Consultative Committee, now named the ITU–R.

invent things when there is no need for that.

JIM DeYOUNG (USNO): I think you said that Dr. Hellwig wasn't here. I took some notes, and so maybe I could give a short synopsis of what happened in our group, "User Environmental Effects."

Dr. Hellwig introduced a document that is going to be published, I believe, in the spring of '95, discussing user environmental effects, including radiation, acceleration, temperature, humidity, et cetera. It's going to be IEEE Standard 1193–1994.

Our group – after Dr. Hellwig gave this little bit of introduction to get us going, he also introduced three areas he thought were important, which is fitness of use. Does your device or system really meet your requirements that you originally had formed? He had another consideration: "How do I characterize this?" or, optimize the design is the bottom line on that. And then he discussed liability and survival of systems that are important in your timing or frequency.

We talked about complex systems, as that's getting to be a problem. We have specifications on individual devices, but then how do you merge those specifications on those devices and get a global picture of how the system is going to perform? We decided communication; in my few years in PTTI, that's always been one of the things we discussed in most of these forums, is communication as one of the most important things that can happen.

There were a few specifics that we discussed, and that happens to be related GPS clocks on board the satellites. At least one gentleman – I'm not sure of his name – mentioned something about the Block II–R clocks where, in the early incarnations of the GPS clocks, they were doing frequency stability measurements; I believe it was temperature variation in a vacuum. Those tests were done and they found some problems with specific clocks. But those tests aren't even being done now in the Block II–R clocks. So that was pointed out as possibly a problem.

Then one final thing we discussed was that the design materials and the components are very important; therefore, you want the highest quality of those things. That's pretty much everything I have in my notes from that group.

RAY FILLER: Anybody have anything else to add to that or to any other topic of discussion? Thank you.