

Bill Burr: Well, here we are again. And it's been a fascinating and for me a somewhat exhausting week. I'll talk a little bit about what we've learned and will do from this.

First thing we've learned that Leuven is really a wonderful town. It's a great university, a great place. We have a great crew of people who helped us do all this from Bart and all of the people who carry the microphones around and run things. And people up there too, I believe, are doing the same thing. And the food was actually wonderful. And I can't think of a conference or a place I've been to where it was done better or run more smoothly, actually. So I really want to thank the people here in Leuven.

The second thing I've definitely learned is that we have a lot of good candidates. We have a very strong field here, almost strong enough it's going to be a problem I think. Just in terms of we're going to wind up not accepting for the second round some very good hash algorithms, algorithms that very probably if we had picked them to be the final winner, they wouldn't have disgraced us at all. But that's the problem I'd rather have than the opposite problem. So it's a good thing, really.

We've heard quite a bit this week that hashes are everywhere. And indeed, I suppose they are, the high end and the low end and everywhere in between. And we can't, I don't suppose, wish away the low end. There will always be uses for the smallest useful processor, and it will shrink down to some microscopic thing, and people will still want to use the smallest processor they can in many applications. And the only question that I really have there is, are these very often things at the very absolute dead low end of the scale, that involve collision-resistant hashes; or are they generally more simple authentication tasks that might be better handled by some simpler kind of a MAC?

So that's really, in some ways, the first of several kind of quandaries or dissonances that I think I've identified. So we've got then this extraordinary range, right? The computer on your desk now is by the standards of my youth a supercomputer of almost unimaginable power, something with two or four cores and running at not in megahertz but in gigahertz. And at the other end, the low end stays about the same. It's still something that's clocked relatively slowly and it doesn't have much memory, and it has a fairly simple instruction set. The PDP-12 conceptually is maybe just a little bit less than a typical 8-bit processor today, or the PDP-8, I should say. Getting my numbers confused here, it's been so long. When I had one, it was the size of a refrigerator. And now it would fit on a speck of sand, pretty much. But the machines are not too different. And so we have an enormous range to cover, whether it's practical or good sense or whatever to try and find a single algorithm that does reasonably well at either extreme of this is to me probably the most interesting question of this competition.

Another thought that I've heard here, and I may not have expressed it very well with the bullets, on the one hand, people are saying standards ought to have very few options, ideally none.

We started out with four, right? Which is the output sizes, the hash sizes. And we've also heard Adi say, and I think most people agree, that in practical, real terms, anything over two to the two hundredth or some number in that general field is likely to be overkill for almost any foreseeable future. So do we need some kind of parameters to let us trade off performance and security? But that's kind of in contradiction to the thought that we need standards with few options, and that if we give the user a tuning parameter, is he always going to tune it all the way down so that it runs as fast as possible? And so there's an interesting debate there that I don't have any very good idea on.

Another area that I always get in trouble on, NIST always gets in trouble on, is the importance and the necessity for proofs. And this is I think a generally contentious subject. I think we have to say that proofs are a good thing. Maybe at some point, we'll say they're a necessary thing. And then there's still always the question, what is the price you're willing to pay for a proof? And so in the range of options that we have here, we have some things that have really no particular proof but are claimed to be simple enough that they're a simple element to analyze. And at the other end, we have a surprising number, to me at least, of hash functions that offer rather strong proofs of one property or another that reduce to the elliptic curve, Diffie-Hellman problem or whatever, or the discrete log problem, anyhow. This is to me a very interesting tradeoff generally, but the proofs seem to have a cost in performance terms. I'm not sure that there's a law of nature that says they have to, but generally the more constraints you have on a design, the harder it is to make it fast, I think. And just my rough observation is that most of the algorithms with very strong proofs are among the slower of the algorithms. So that's something that I don't have an answer to at the moment.

And when we get into picking the finalists, I think we're going to have some very intense, or not the finalists, but the second round candidates at this point, we're going to have some very intense discussions at NIST. And probably we're going to get some fairly interesting inputs from people. And I had one or two additional thoughts just about what might be helpful as we go forward.

I've come to the conclusion that the eBash page is a pretty good thing, or site. And that it would be helpful if folks were to get good code in the format. Dan will have a talk on that tomorrow, so I guess we'll know more about just what the requirements are. But it seems like the cheapest, easiest, fastest way to get a lot of performance data with a relatively small effort. And that's a good thing.

The next bullet there I guess is rally Adi's idea. And it seems to me a reflection, a pretty good one, is that if you're making your little change or your small tweak, don't tweak a bunch of different things, tweak one thing. Two small tweaks in different areas, the complexity of figuring out what goes on probably goes up as some power of the number of tweaks. So try and confine yourselves to just one thing and not a variety of things.

Also, in considering the performance and trying to decide if something can be adjusted so that it's approximately at least as fast as the SHA-2s. Probably more practical, if we decide we need to do that, to make a judgment, for us to contact the submitters and confer with them about what's a reasonable way to do that. Because we have some multi-dimensional algorithms, that is multi-tweak, or multi-adjustable, whatever it is, with several things that you can-- and it may not be obvious to us what the right way to adjust it is in terms of getting the most security for the least computational cost. So if we decide that to make a decision that we need to do any adjusting of parameters, we'll talk to the submitters about that first and try and work something out with them for what's reasonable.

Okay. So here's our rough plan. It's pretty simple at this point. We want to cut down to about 15 second round candidates. And the target date for that is before Crypto 2009. Okay, so we've got until this August to do that, okay? That's not a long time, but it's a reasonable length of time. We can do a lot more thinking and focusing than we have so far.

We intend to come up with a balanced group of 15. We have quite a diversity of proposals here, and we expect to preserve a good bit of that diversity in the final group. They are not going to be all proposals that use the AES round function. They're going to include some stream ciphers, base stuff. They're going to be a reasonably diverse selection. So if there are three or four very similar good proposals, it could be one or two of them wind up being cut just because we have a surplus of that sort of proposal.

The truth is that we expect to wind up with a group where most, if not all, hopefully all of them, would actually make a good final selection. Something with at least decent performance, and that's very unlikely to be broken for quite a while. And we'll wind up cutting some good hashes. And I guess I've already said that.

So we're then figuring that if we make this decision by the time of Crypto 2009, this coming Crypto, we'll have about a year for people to bang on things before we have our next conference. The current plan is to try and hold that in conjunction with Crypto 2010. And my understanding is that the plan is to hold CHES also in Southern California or in Santa Barbara, just before Crypto. So this will make one awful week for people who stay for the whole thing. But it's a nice place to have an awful week or whatever. And maybe the second nicest university to hold a meeting at in the world, I think.

So what we'd like is, to help us with the selection of the 15 now, if you can get us results, preliminary results, whatever it is you have, by June or so, that gives us some time to digest them before we have to make the final cut. We have our email list and we have our official comments, or whatever, and we pay attention to them. And we'll also look at what's posted on the Zoo and on the eBash site. But if you have results that really bear on the selection, we'd appreciate getting an official comment so that we make sure we have a record of it and we know about it.

So that basically is what I have to say. And I guess this is a chance now to open the floor to additional comments, discussion, suggestions, complaints, whatever you have. So does anybody have anything they want to say at this point?

Man 1: You ask us to post comments and analysis of hash functions by first of June. But it's a very limited time. And we have 50 or 40 proposals. What should we concentrate on?

Bill Burr: I'm sorry. I'm just having trouble understanding your words.

Man 1: There are plenty of proposals. Which should we take for analysis? Because there are few groups of cryptanalysts so far who do analysis. And we cannot analyze all the proposals by the first of June.

Bill Burr: No you can't. That's clear. And I'm sorry, but I think the length of time it would take to do a good job of analyzing all the proposals that we have is just too much. And so the idea here is, I've got to make, I feel, a relatively quick chop to get it down to a tractable number. Ideally, we'd take the next five or ten years to complete this competition, but my masters don't have that long a time horizon. And I think probably most people want a process that has a reasonable termination time. One of the interesting things to me in the discussions we had a couple of years ago at NIST when we had people come in was that the industry people are always saying, "I need it next year. I've got some deadline I need it by, and if you don't get it to me by then, everything slips a decade or something like that." Well, we haven't it for them next year, but they will lose interest if we don't make some progress and show them that we're actually going to give them something. And I want this to be used. I want it to be implemented.

Rich Schroepfel: Could you define a low-end platform or two?

Bill Burr: A low-end platform?

Rich Schroepfel: No, I don't mean tell me a word definition. I mean, tell me smartcard XYZ 43-22 is what you're going to be measuring on.

Man 3: It'll be _____.

Rich Schroepfel: Yeah, but...

Bill Burr: What you want me to do is to give you a very specific low-end platform to measure on?

Rich Schroepfel: Yeah.

Bill Burr: The way we've given you a dual core--

Rich Schroepfel: Right.

Bill Burr: Okay, I know a lot less about that. But I guess my instincts are that it'll-- I'll try to come up with some reasonable test platform that's available. I think the problem with the low-end stuff in general is that you buy this development kit or something and not everybody-- you know, whatever you pick, only a limited number of people are going to have it. But my notion is that it's not the dead low-end, it's far from the dead low-end. But something with an ARM processor on it, and maybe even a coprocessor, I don't know. Maybe it's easy to get stuff to develop for say an iPhone, I don't know.

Niels Ferguson: That's not low end.

Bill Burr: Pardon?

Niels Ferguson: iPhone is not low end.

Bill Burr: It's too high, you think? Okay. Give me some suggestions, people who understand this, for something that is important and common and available without too much effort and expense.

Donna Dodson: There were some people from my session who did offer to provide three or four-- During my session, where we talked about different platforms, there were a couple comments on that. And some people have said that they would suggest some low-end platforms and post them to some of the lists for people to look at, as well as what we come up with.

Bill Burr: I don't know what the smallest thing you could conveniently plug into eBash is. But I imagine that would still be well above what Niels would consider the low end. Adi?

Adi Shamir: Answering the question earlier, what we can do in order to concentrate the effort of cryptanalysts, I think that by now NIST can announce between 10 and 20 algorithms as being excluded from the next stage either because they are way too slow or because serious problems have been found in them. So if you want to help the community concentrate a little bit, you will not be able to go from 60-plus submissions to 15, but you will be able to go from 60 to maybe 45 algorithms. So you can indicate within the next two weeks a short list of algorithms that will definitely not be considered further. That's an interim solution. Instead of waiting until the summer with the whole range of algorithms, you can specify which algorithms you think have no chance.

Shu-jen Chang: At this point, we only have 41 left. So it's not 64, it's not 51. We have only 41 left.

Adi Shamir: The identity of the 41 have been announced, right?

Souradyuti Paul: There were 64 submissions initially. And we have selected 51 for the first round, 10 have been withdrawn. So we have 41 submissions.

Adi Shamir: Okay. But you depend on the voluntary retraction, that some algorithms were withdrawn by the authors. I think that you can also indicate, even though the algorithms were not withdrawn voluntarily, that some algorithms remaining in the field will not be considered.

Souradyuti Paul: On some of the algorithms, there are attacks reported. We don't say officially what we think about that, but people would be able to find out, people would be able to cut it down to say around 30, I think watching the results that are around.

Bill Burr: In any event though, your suggestion is that, to the extent that we can drop a few sooner rather than later, it would help.

Adi Shamir: <inaudible>

Man 7: On this question of moving to the second round, I understand the comments of NIST not wanting the process to drag on. But that process really does seem to be a very aggressive timetable to me. A lot of the hash functions that have been presented today, over the last few days, have basically been seen for the first time by many of the people in this room, I suspect. So you're allowing three months for analysis to take place. It just feels like a very tight deadline. And I'm wondering if aiming for Crypto 2009 isn't a little bit artificial. It's a nice deadline on the cryptographic calendar, but maybe just allowing even a month or two months more would really give the opportunity for the community to give much better information so that you can make a decision into the second round. That would make it much easier to make that decision, I think.

Bill Burr: I'm sorry. I'm a little hard of hearing, and sometimes I just have trouble understanding what people are saying. And I just really didn't get the substance of this.

Man 7: Okay. I'll speak slower, maybe that's the issue.

Donna Dodson: Two more months-- <inaudible>

Bill Burr: Two more months.

Donna Dodson: He's saying that while Crypto is a nice target, maybe it's a slightly artificial target because of the meeting it. And if the community had even two more months, it would be substantial.

Man 7: Yes, that's what I'm saying.

Bill Burr: Okay, okay.

Donna Dodson: Is that a correct paraphrase?

Man 7: Yeah it is. I think another couple of months of time would really benefit the process substantially in going to the second round.

Bill Burr: Do the people think another two months are a big deal?

Man 8: Well personally, as an academic, I find that I have most time to do research during the summer. I'm imagining that's true for many other academics as well. So if you make a deadline here of June first, that's not giving us a lot of summer time to do research.

Bill Burr: Stefan, you have something?

Stefan Lucks: Now, if the consequence isn't to extend later deadlines also by two months, that would mean that there are two months less to focus on the 15, I suppose semi-finalists. So I'm not sure if that would be really a good thing.

Bill Burr: Well that's the trade off, isn't it? Assuming I want to hold the next workshop in conjunction with Crypto, and CHES and Crypto together seems like a strong attraction.

Man 7: Can I just answer Stefan's point? Because actually the whole point of having those extra two months means that the information we go forward with into the second round is much better quality information when we've actually made a better choice of the algorithms and moved forward. We haven't chosen some that all of sudden we find some weakness in. So in fact, I think by having a little more time for analysis in the first round doesn't necessarily take away from having good analysis in the second round. I think it would strengthen the whole process.

Man 10: <inaudible> When I look up here, I see what you have planned as it is. Analyze until June, and then in August, I'll give you 15. And then for a year later, second conference. What will happen between that Crypto 2009 and 2010? That's not very clear either.

Bill Burr: What will happen then, my hope at least, is that people will have a year to do analysis on the 15 remaining candidates. And so it maybe a little bit arbitrary, but I figured that to get a start on analyzing 15 candidates, it would take people about a year to do it.

Man 10: That's what I wanted to hear from you. Now what you're saying is analyze 41 candidates in four months, then I'll give you one year to analyzing 15 candidates.

Bill Burr: Well, fair enough.

Man 10: So let's balance those things.

Bill Burr: What I guess I'm saying is that I'm going to make-- that I think the next six months are a long enough period to make a reasonable choice based on the features and the apparent properties of the hash submissions that we've got and that if I have 40-odd things, I might want two years then to do really good cryptanalysis of 40 of them. And so I'm trying to say that I think for performance reasons and for other reasons, we can probably pick a reasonable set of 15 in about six months, okay? But I don't imagine that we can claim in about six months that those 15 have been well vetted through cryptanalysis. Given the large field, we have to do something pragmatic to get it down to a group that can be cryptanalyzed in a reasonable length of time. Adi?

Adi Shamir: Let me try to explain why I think that your timetable makes a lot of sense. I support it. There are two stages. If in the second stage, when we have to reduce the number of candidates from 15 to 1, we make a mistake and we choose a wrong bad algorithm, this could be a total disaster. On the other hand, if we are looking at the problem that we are facing, the immediate problem we are facing, reducing the field to 15, it's very hard for me to imagine that all 15 that will be chosen are going to be bad algorithms. So for the designers, it could be a disaster if their algorithm is not chosen. But we will still have a field of very good choices to look at. So maybe NIST will not make the absolutely optimal choice at the end, but it will make a very good choice. So it is much more important to leave enough time for analyzing the small number of final candidates than to do a perfect job in vetting the large number of initial submissions.

Bill Burr: Very nicely put, actually. Rich?

Rich Schroepfel: I don't think there's any reason that you would need to synchronize your announcement with the Crypto meeting. So I think in that sense, the Crypto 2009/CHES meetings are artificial deadlines for announcing the cut down. The other thing is if you're asking people to submit analysis by June 1, but you're not going to put out your answers until end of August, then they still have 41 selections to work on between June 1 and August whatever. So your cut down is of no value until it's announced.

Bill Burr: Fair enough.

Rich Schroepfel: So you're putting your two-month gap in October might be a better decision.

Bill Burr: There's no reason that this announcement has to be exactly synchronized with Crypto. And the sooner we make it, the better, I think. So I was trying to say what's the shortest length of time that I think people have a chance to do much cryptanalysis on things? And I said to myself, probably they could do it by June. I want a little time after I think most of the cryptanalysis is in before I have to announce the list, okay? And so it just seemed like that was a reasonable schedule given the amount of arguing I think that we'll wind up doing ourselves. But it makes sense, I think, for us to release the list as soon as we're comfortable with it. But I have to give us a deadline, and it seems to me I want the deadline to be not later than Crypto. If I want to hold a

meeting a year later, let's give people a year. If we can make the announcement on the 15th of June, we will. And that will give people more time. I don't want to let it drag past Crypto. And I'm hoping that if I give a reasonable length of time and say you have until June, some people will be motivated to try and do some serious cryptanalysis on something. The more cryptanalysis we can get, the better, right? Everything that's broken between now and then is a godsend to me.

Niels Ferguson: I'd like to support the early selection to 15 candidates. When we get down to 15 candidates, we will be where we started the AES process. And even then in AES, we still had a bunch of very weak candidates left, so we eliminated a bunch very quickly. So I think that's good. I think NIST could help us by early down selecting those candidates which you already feel will not make it to round two. And I guess from the 41 left, there's probably a dozen or so where it's extremely unlikely they would ever make it. If you could do that in the next two weeks, it would help the community in not spending time on things that <inaudible>

Bill Burr: Okay, well that's Richard's comment earlier really. And to tell you the truth, my reaction to that is mainly that if I had it to do over again, I think I would have accepted all 60-odd submissions I got, even the ones that were woefully incomplete, just because people are really offended when they're dropped conspicuously early. And they don't like it and they complain. So the bureaucrat in me says that the safe bureaucratic thing to do is to wait.

Niels Ferguson: <inaudible>

Bill Burr: If we can clearly identify ones that we think are--

Man 11: Okay. I also have a question related to something that Paul said. Do I understand correctly that if a designer decides to withdraw his submission, it's removed from the competition automatically?

Bill Burr: You're saying if people withdraw their submission, are they removed? Yes.

Man 11: That's the question.

Man 12: <inaudible>

Bill Burr: I think I'll have enough who want to be selected.

Man 11: <inaudible>

Bill Burr: There are a number of things here that I think we could say. Well, they've got a black mark against them, we can just drop them, okay? But some of them are actually pretty interesting. And so I'd like to take some time thinking about that. There are probably a few others that are so slow we can just say, "Well, you know, you're just too slow." But I think that we have a number of, several at least I can think of, really interesting decisions to make here in terms of

yeah, they have blemishes, but they also have pretty exceptional strength. The easy way to say it is “Well, take anything with a blemish and throw it away.” Okay? And we’d still have 30 or more things left, right? But actually, there’s a couple of blemished algorithms that are pretty interesting, and I want to give them some thought surely. But we’ll go home and start sorting things into piles. This is a little bit I think, at this stage, like selecting papers for a conference, right? You know fairly early on that some aren’t going to make it, right? And you know fairly early on that if you are and you spend most of your time with the ones in the middle. And let us go home and actually start sorting before we decide whether we can do reasonable early cuts.

Man 11: So even if a designer withdraws his proposal, it can still be selected for the next phase. Is that what you’re saying?

Donna Dodson: <inaudible>

Man 11: <inaudible>

Donna Dodson: So your question is, if we have somebody who says, “I would like to withdraw my submission,” then is it out of the competition?

Man 11: Yes, that’s my question.

Donna Dodson: Is that the question?

Bill Burr: Yes.

Donna Dodson: And the answer is?

Bill Burr: I thought I answered that earlier.

Donna Dodson: Yeah, I think you did.

Bill Burr: Yes, we have enough people who want to stay in the competition that we don’t have to drag people who are happy to withdraw along. If the designers want to withdraw, they can withdraw.

Danilo Gligoroski: When you make a decision about 15 out of 41, is it possible to get the ordering from the 16 to the 41, but not from number 1 to 15?

Bill Burr: Yeah, I think not.

Man 14: I don’t think there’s an ordering.

Bill Burr: I don't--

Danilo Gligoroski: What will happen if during the first year three of the 15 candidates are broken and the 16th, 17th, and 18th place are there, not broken, what will happen?

Bill Burr: It's not a case if we have to have 12 jurors. We don't need-- in the U.S. in jury trials, if it's an important case, we commonly have two or three alternates sitting there listening to the case, in case a juror gets disqualified or gets sick or something, so that they have a full quorum of 12 so they can carry out the death sentence or whatever is required. And in this case, I think if something gets broken in the 15, then we're down to 14, or we're down to-- no. We're not going to order-- and that's a problem I have actually with early cuts, is it's a kind of ordering. And people are going to take it negatively.

Man 14: <inaudible>

Bill Burr: Well, fair enough. But I still think not.

Man 14: Well then also, number-- <inaudible>

Bill Burr: Well, but then you have even more incentive to breaking another one of the 15.

Man 15: I ask if you cannot tell us what algorithms not to consider, maybe you can tell us on which we should draw more attention.

<laughter>

Man 15: Seriously.

Bill Burr: I'm sorry. Can you repeat the question, first of all?

Man 15: I'm repeating the question.

Woman 3: So the question now is can you tell us which ones to look at? <inaudible> --don't look at, then which ones to look at.

Man 15: It's not the end of the question.

Bill Burr: Which ones to look at, okay. That's another way of asking the same questions.

Man 15: No. No, it's not the same question.

Bill Burr: That I just dodged in the first place.

Donna Dodson: I have to say that I think there are a lot of fairly clever people in the room. And given some of the discussions that we've had today and yesterday about some of our considerations, some of your considerations, it's probably fairly reasonable for people to start thinking about their own piles and what it is they want to take a look at and that sort of thing. So I kind of would like to give the challenge a little bit back to the audience. Because again, this is a community process, it's not just NIST.

Bill Burr: Well, if you're doing this tactically, pick the one that's closest to yours that runs a little faster.

Man 15: No. The point is that there are many algorithms, there are many proposals that are partly broken, or some reduced versions are broken, or they have another flaw. And people who will do cryptanalysis the next three months, they will not start from scratch. They will likely extend their current work. And many algorithms, many proposals, will remain untouched. And if you are satisfied with this situation, okay. If you want for some algorithms to be considered more, then perhaps you should state about it. Because so far, there are many algorithms with no attacks. If this situation will continue until first of June, what will you do?

Man 16: May I speak? I have the feeling that not all algorithms are attractive to everybody. And certainly, not all of them will receive an equal share of attention, especially since some candidates here are offering prizes in money and food or beer or whatever. Certainly, it's no joke, really. Some people will be attracted to do that and neglect many of the others. So some candidates will attract all effort while the others will be neglected. What will happen in the end? You are going to reject candidates that have no analysis at all? So I wonder what could be a good strategy to allow each candidate to receive an equal share of attention, so we have an equal share to be overviewed, analyzed by all candidates?

Bill Burr: An equal share of the attention.

Man 16: Yeah, because some of them are becoming more attractive to analyze than the others because if you successfully break them, you get some prizes, but the other candidates have no prize.

Bill Burr: Okay, so Bruce wants to answer that, I think.

Stefan Lucks: Two quick points. To first the issue of telling which candidates to look at or which candidates not to look at, if you are able to do proper cryptanalysis, then you should have a very decent idea which ones are likely to be chosen and which ones are borderline, and which ones will very, very likely not make it. A second point to, you know, ordering and giving the 15th and 16th and 17th candidate. I am sure the 16th candidate still has good motivation to break one of the first 15, and even if it's just for revenge and for the fun.

Bruce Schneier: I think inherently we have to trust the random process here. This is not a corporation; we can't allocate equal hours for cryptanalysis. We have to believe that there are enough of us here looking at the things that we notice, we want to look at, we think are breakable, we think are good, we think will get us a good publication, we think will get the most kudos for breaking, will improve someone else's attack, that there will be enough random people doing enough random things that the good things will tend to bubble up. The perfect thing won't be at the top, it won't be a perfect ordering, but for the next 15, it'll be a good ordering. That is the fundamental conceit of this process. If we don't buy that, then we don't believe in the process. The process seemed to work with AES, there's a good chance it'll work here. I think it's been working great so far. So I think rather than sort of min/max this, let's trust the randomness.

Rich Schroepel: I propose a sort of secret Santa scheme.

Bill Burr: What?

Rich Schroepel: We put the 41 names in a hat and everyone draws out a name, and if it's theirs they put it back and rescrumble. Their job is to analyze the name they receive.

Bill Burr: Greg.

Philip Hawkes: Sorry, okay. You were saying before that the best people to attack are the ones that are similar to yours because you're most likely to choose one of similar ones. So if you could provide a possible grouping. This is maybe not the best idea, but so at least all the people, all those teams within that group know who their competition is, and then they will make sure that they focus attention on whoever else is within their group, to make sure they have the best chance of going through. That's not a perfect scheme either. I thought an interesting mathematical thing is this process ends up sort of being orthogonal then, because the first round is looking at things similar to you, so sort of like looking at diffusion to closer bits. And then the next round will be sort of trying to look at things orthogonal to you, because we'll have eliminated everything that's similar to you.

Bill Burr: Okay, in the back. Dan, I guess.

Dan Bernstein: Yeah. I'm wondering a little bit whether NIST could engage in one of those famous U.S. government anonymous leaks of information.

<laughter>

Bill Burr: Okay, Dan.

Dan Bernstein: In all seriousness, from a designer's perspective, if my hash function disappears, I'm sure we'll end up with something just fine in the end. But from the cryptanalysts' perspective, if they're working on some function and then NIST says all of a sudden, "Sorry, we're dropping it because it's too big or it has this previous attack," or whatever, then suddenly the cryptanalyst has been wasting quite a bit of time. And if NIST would give some hints in advance, for instance, NIST could just mention on the mailing list that, you know, "These are some interesting proposals, and these are some very interesting proposals, and these are some very, very interesting proposals, then cryptanalysts will have less chance of feeling like they've been wasting their time.

Bill Burr: No, no. There may not be a good analogy, but you don't want the judge in a case to just be sort of hinting how he's going to rule. I want to reserve the right to change my mind up until the last moment, and not gratuitously tick people off at me either.

Ron Rivest: So proposals that don't get any analysis at all probably would end up being rejected in the end. So I think maybe if somebody's looking at a scheme, it'd be courteous to inform the developers that you're doing so. And if you're a developer and nobody's looking at your scheme, you've got an opportunity then to hawk it on a hash forum list or something like that and try to generate some interest in it.

Bill Burr: Yeah.

Man 21: This is a very short comment. If you want to see every one of the hash functions touched, the only way to solve that problem is to put a couple thousand dollars on each one of them. Therefore, the monies will be won, the ones that's not won is what we go after. That will cost you \$82,000, Mr. Burr.

Man 22: <inaudible>

<laughter>

Bill Burr: Okay. I think maybe we've reached a point of silliness here.

Man 23: <inaudible>

Bill Burr: I think there are-- what we will try and do is if it's clear we can cut some early, we'll give that a lot of thought and maybe do that. But we're not going to release any more information about relative rankings than is required to make the process work. We're going to try very hard not to hint who we're going to pick. We're perfectly happy to discuss what's important and why and what some of the dilemmas are and so on. But we're going to not try and signal everything in advance. And so there may be a-- we can get rid of a few of them early. And that maybe we'll decide also that even that is-- you know, you people are supposed to actually know what's good.

And what I'd like you to do is pick something good and try and hack it now. That's the way it seems to me you really earn your bones is break something good. And certainly, I know John Kelsey doesn't actually go looking for easy things to break. I mean, he gets really excited if he thinks he's got something good to break.

Donna Dodson: <inaudible>

Bill Burr: Oh yeah, fair enough. I should announce that, or should if I haven't. I thought I had. John Kelsey, who many of you know, is not here because he's a new father. His third child, his first daughter, who I believe is names Janetta, was born. I guess it's now been two weeks since the birth. And otherwise, we've actually been running around trying to figure who was going to substitute for John where. That's why you haven't seen John Kelsey, not because somehow he's out of favor or whatever. Okay. I think we have run out of steam at this point. I want to thank you all. I want to also again thank the people at Leuven and Bart for the really wonderful job you've done here hosting us. And for the hardcore folks, we actually have quite a number of Rump session talks scheduled tomorrow morning and some very interesting talks on some of the websites that have grown up around the competition, and which we actually use quite a lot. So for those of you who'll be here tomorrow, I'll look forward to seeing you again in the morning at nine a.m. And for the rest of you, thank you very much for coming.

<applause>

End of Session12.mp3