FEDERAL TRADE COMMISSION

I N D E X (PUBLIC RECORD)

EXHIBITS	FOR ID	IN EVID
RX		
Number 695		4982
Number 1532		4858
Number 1762		4981
DX		
Number 47	4994	
Number 48	4994	
Number 49	4995	
Number 50	4996	
Number 51	4996	
Number 52	4997	
Number 53	5001	
Number 54	5004	
Number 55	5007	
Number 56	5008	

UNITED STATES OF AMERICA

FEDERAL TRADE COMMISSION

In the Matter of:)

Rambus, Inc.) Docket No. 9302

----------)

Thursday, June 12, 2003

9:30 a.m.

TRIAL VOLUME 26

PART 1

PUBLIC RECORD

BEFORE THE HONORABLE STEPHEN J. McGUIRE

Chief Administrative Law Judge

Federal Trade Commission

600 Pennsylvania Avenue, N.W.

Washington, D.C.

Reported by: Josett F. Hall, RMR-CRR

For The Record, Inc.
Waldorf, Maryland
(301) 870-8025

APPEARANCES:

ON BEHALF OF THE FEDERAL TRADE COMMISSION:

M. SEAN ROYALL, Attorney
GEOFFREY OLIVER, Attorney
JOHN C. WEBER, Attorney
MICHAEL FRANCHAK, Attorney
ROBERT DAVIS, Attorney
Federal Trade Commission
601 New Jersey Avenue, N.W.
Washington, D.C. 20580-0000
(202) 326-3663

ON BEHALF OF THE RESPONDENT:

GREGORY P. STONE, Attorney
STEVEN M. PERRY, Attorney
PETER A. DETRE, Attorney
SEAN GATES, Attorney
Munger, Tolles & Olson LLP
355 South Grand Avenue, 35th Floor
Los Angeles, California 90071-1560
(213) 683-9255

APPEARANCES:

ON BEHALF OF THE RESPONDENT:

A. DOUGLAS MELAMED, Attorney
Wilmer, Cutler & Pickering
2445 M Street, N.W.
Washington, D.C. 20037-1420
(202) 663-6090

P R O C E E D I N G S

- - - - -

JUDGE McGUIRE: This hearing is now in order.

Any issues that need to come to the court before we start this morning?

MR. STONE: Mr. Detre had a son, Aaron Frederick, so everybody is well, so I wanted to share that.

JUDGE McGUIRE: Good news. The court is pleased with that information.

Anything from complaint counsel?

MR. ROYALL: Nothing here other than congratulations.

JUDGE McGUIRE: Very good. Then at this time complaint counsel may call its next witness.

MR. DAVIS: Thank you, Your Honor.

Complaint counsel calls Pete MacWilliams.

JUDGE McGUIRE: Sir, if you'll please approach the bench and you'll be sworn in by the court reporter.

- - - - -

Whereupon --

PETER D. MacWILLIAMS

a witness, called for examination, having been first duly sworn, was examined and testified as follows:

MR. DAVIS: Your Honor, before I start, could I approach the witness to give him documents?

JUDGE McGUIRE: Yes.

DIRECT EXAMINATION

BY MR. DAVIS:

Q. Good morning, Mr. MacWilliams.

A. Good morning.

Q. Please state your name for the record.

A. Peter D. MacWilliams.

Q. And where are you currently employed?

A. Intel Corporation.

Q. What's your current title?

A. I'm director of platform architecture.

Q. And where did you go to college?

A. I went to college at UC Berkeley and got a bachelor's degree in 1978 and a master's degree in 1979.

Q. Where did you start working after leaving Berkeley?

A. I went to Intel.

Q. And when did you start at Intel?

A. In August of 1979.

Q. What was your first position at Intel?

A. I started working on analog I/O boards for the multibus product line.

Q. And what --

A. As a design engineer.

Q. How long were you doing that?

A. Oh, roughly two years.

Q. What did you do after that?

A. I continued to work on analog I/O boards for the multibus product line as a design engineer and I started to spend some more time working with the component groups within Intel, trying to coordinate the components they were building to make sure that they would work better in our board-level products.

Q. When you say "component groups," what were you referring to?

A. Components would be microprocessors and peripherals that go around the microprocessors.

Q. And how long were you in that position?

A. Oh, in that specific position I was there roughly another two years or so. But the department where I worked with the component parts of Intel's continues to this day.

Q. And where did you go after that position?

A. I worked in the systems technology group. We were looking at a lot of the forward-looking technologies that we needed for the systems in the future in both the product-level technologies as well

as board, thermal, other aspects of systems.

Q. And about what year did you start at the systems technology group?

A. I actually started working there about the mid-'80s on a part-time basis while I was still working in the multibus board group and had transferred there full-time in 1990.

Q. Did you work with DRAM at all in that period?

A. We used DRAM, but I never worked with DRAM suppliers. We just used whatever they built.

Q. After that, what was your position?

A. I worked in the systems technology lab group when came to Intel. In 1990, at that point I transferred over there and we were looking at technologies to speed up our microprocessors. The primary focus then was other system components and high-speed caches.

Q. What's a high-speed cache?

A. A high-speed cache is a type of memory that we put behind the microprocessor. It's typically much faster than DRAM. It was SRAM initially. We looked at doing some integrated components and the idea was to supply instructions and data to the microprocessor that are most often used more quickly than we can supply by DRAM.

Q. And did you have a position after the position you just described?

A. I was in that position basically from that time until roughly 2000-2001, at which time I transitioned to the desktop platforms group, performing basically the same job. Throughout the last thirteen years, my job has been to look for issues with our microprocessors and platforms into the future and try to figure out what technology we need to deliver the full value.

Q. When you say "platforms," what were you referring to?

A. "Platforms" means the boards and systems that we build up around the microprocessors to make the computer.

Q. Do you have an understanding of the term "memory enabling" as it's used at Intel?

A. Yes.

Q. And what does that mean?

A. What it means is that as we started looking at what it takes to make platforms better in the future, around the middle of the '90s we concluded that we needed to have higher-speed memory technologies. It was no longer sufficient to simply work on CPUs and caches, and as such, we had to go out and work with

people in the industry who supplied the memory technologies, and "memory enabling" refers to programs that we put in place to help coordinate our product lines with the memory vendors' product lines and help them see what we wanted them to do in order to better complement what we were doing.

Q. Were you ever involved in a memory enabling project at Intel relating to RDRAM?

A. Yes.

Q. And what was your role in that project?

A. My initial role was to decide that's what we would do. I was one of the key people on the team that looked at memory technologies in the future and decided that the RDRAM technology would be the right one to pursue for the long term in terms of its technical merits.

Q. Was this the project that you were describing that began in the mid-1990s?

A. The specific RDRAM project started a little bit later than our memory-enabling efforts but roughly the same time frame, yes.

Q. So the memory-enabling effort began in the mid-1990s, in 1995 or so?

A. Yeah, about 1995. And we spent about a year looking at lots of options before we actually decided

to take the Rambus direction and then about a year or two to finalize all the details to confirm that was indeed the right way to go.

Q. Now, why was Intel interested in enabling RDRAM rather than some other DRAM?

A. When we did the analysis, we looked at several options. Back in the '95 time frame, we were using EDO, just starting to use SDRAM. We looked at speeding up SDRAM. We looked at what was referred to as DDR SDRAM. We looked at SyncLink. We looked at various versions of RDRAM. And we also asked the memory vendors to provide us with whatever ideas they would have to provide high-speed memories.

And as we looked through the options, the memory vendors didn't have any other good ideas and the scaling of SDRAM to DDR we judged to be somewhat limited in that we could put a lot of energy into this and only get another maybe step or two in terms of memory generations.

We looked at the RDRAM case, and while the current RDRAM implementations, which at the time were base RDRAM and concurrent RDRAM, had limitations, that made them not work well for system memory in PC. But we saw that because they had already implemented those, it was a fairly proven technology that we could

evolve forward and create a technology that would scale for the five-year time frame we were looking for.

Q. You said "five-year time frame."

Why would it be for a five-year time frame?

A. It was somewhat arbitrary, but it was important that we pick a technology that would allow some stability and longevity because we were going to ask the industry to go through a major transition in terms of the infrastructure. The connectors, the boards, the modules would all need to change, and that's not something that you can change on a yearly or every-other-year basis. You have to do it on a longer term and we thought five years was a reasonable time frame.

Q. Why couldn't you make those changes on a yearly basis?

A. Just the amount of investment to make the changes, to validate the changes are correct, to optimize the results based on the feedback you get from the first designs, will take longer than one year, so...

Q. When you're talking about investments, were you talking about -- take a drink of water. I didn't intend to stop you from drinking water.

Were you talking about solely Intel investments or investments by other firms as well?

A. Both. Intel would have to make investments, our OEM customers would have to make investments, and DRAM vendors would also have to make investments.

MR. DAVIS: Your Honor, I'd like to show the witness RX-904, and unfortunately, I didn't tell my legal assistant I needed RX-904, so I don't actually have a physical copy.

JUDGE McGUIRE: All right. That's fine. Go ahead.

MR. DAVIS: It will pop up on the screen.

BY MR. DAVIS:

Q. Can you identify what this is?

A. Yes.

JUDGE McGUIRE: Do you have of a copy of this, Mr. Stone? I'm sure you do somewhere.

MR. STONE: I'll find it, Your Honor. That's fine.

JUDGE McGUIRE: Do you want to take a break until you get that?

MR. STONE: No. And I have extra copies if counsel would like them for the witness and himself.

MR. DAVIS: Thank you. I appreciate that.

May I approach?

JUDGE McGUIRE: Yes.

BY MR. DAVIS:

Q. Could you describe what this is.

A. This is a Rambus program review. When we signed the agreement with Rambus in I believe it was November of '96, we decided we needed to make this a more formal program within Intel and we needed to take all aspects of it and do a review for the Intel executives, and this was the review.

Q. Is a program review something that's commonly done at Intel or is that something special to this case?

A. It's common but not universal. So we'll do it for lots of things but not everything.

Q. Do you recall if you gave a presentation at this meeting?

A. Yes.

Q. If you'd turn to the third page, there's a slide entitled Memory Enabling.

Was this part of your presentation?

A. Yes.

Q. Under the word "Goal" is the statement "Ensure memory subsystem does not limit processor business" and then below that there's a bullet that states "do not limit ramp" and then parentheses "price and

availability."

What did you mean by this subbullet?

A. Basically we were concerned that with any new memory technology, for it to achieve high volume we need to be price competitive with the previous technology that was already in high volume and that we also needed to guarantee that there would be enough availability for the Intel processor business, which is -- consumes a substantial percentage of the DRAMs worldwide.

Q. What was the importance of the availability?

A. If there wasn't enough DRAMs available, then we limit the amount of processors we might be able to sell.

Q. Below the goal is a list with a title Process.

Could you describe what that list pertains to.

A. I'm sorry. Can you be more clear?

Q. Sure.

Could you describe what that list is.

A. Basically this list is a set of items that we actually viewed as enabling activity.

Q. So this is what you were describing earlier?

A. Yeah.

Q. The first bullet under Process states, "Project future memory requirements (extrapolations,

discontinuities)."

What does that mean?

A. We basically want to look at what our platform products are. We take our CPUs and try to project out what the performance would be in the future -- what the performance of our CPUs will be in the future, try and project what the memory requirements for those CPUs would be. Initially we'd do that with just extrapolating what we have to -- what we have best and try and figure out what the future needs are.

We'd also look for discontinuities in the sense that if we come up with new processor technologies, for example, recent hyperthreading is discontinuity work where one CPU might demand more bandwidth than the previous one or graphics which demands a lot more bandwidth than previous chipsets.

Then we'd try and factor that in and we'd take the result and try and project what our needs are in the future.

Q. Now, is this a process that began in the mid-1990s, 1995-ish?

A. Yeah.

Q. When was the process finished?

A. Well, it's never really finished. We go through it over and over again. We went through one

process in '95, made the direction decision at the end of '95 to go with the RDRAM. We confirmed that late '96 and began that program, but during the whole period of the late '90s and even till today we continue to project future needs and try and figure how to meet those.

Q. If you could turn to page 7 of RX-904.

The title of that slide is Intel nDRAM Target Requirements.

What does nDRAM stand for?

A. It stands for the next-generation DRAM.

Q. Is that the result -- is this next-generation DRAM the result of the process you were describing a minute ago?

A. Yes.

Q. The second bullet on that slide is: Cost parity with 100 megahertz SDRAM.

What constituted cost parity in your analysis?

A. What we meant by this was that the next-generation DRAM once it was in volume production would have to equal the cost of 100 megahertz SDRAM when it was in production, cost measured by the DRAM cost per bit and the system costs associated with the module and connector, and so forth.

Q. Now, what did you mean by "cost per bit"?

A. By "cost per bit" we mean basically the die size of the DRAM.

Q. Why is cost per bit important as opposed to, say, cost per DRAM chip?

A. Cost per bit is important because we look at -- well, I'll step back.

Cost per bit was a way to look at it without specifying specific density. So if we'd said, you know, equivalent costs for 64-megabit DRAMs, that might not be fair because the next DRAM generation might not occur in the 64-megabit generation. If we said 256-megabit, the old generation of DRAM might not be around that long. So we use cost per bit to sort of equalize it to allow us to make the comparison at whatever density made sense.

Q. I see.

So cost per bit is a measurement that allows you to compare memory cost across generation?

A. Or allows us to compare memory cost at each generation as opposed to at one specific generation.

Q. I see.

Now, going back to page 3 of RX-904, the bullet -- the next bullet states: Assess memory technology options, availability, cost and risk.

And is that something you were describing

earlier as well?

A. Yes.

Q. What type of technology options are you referring to here?

A. What I'm talking about there is we'd spent time with the DRAM vendors asking them what they could do now and in the future, tried to look for their capabilities, tried to look at what process technologies would bring them in the way of benefits and tried to look at areas where our requirements might add cost to the devices, so it was a very important step to try and figure out what they could do to deliver our requirements without adding cost and which of our requirements would add cost, which allow us to backward think.

Q. Now, for the RDRAM-related memory-enabling effort, this process began you said in 1995?

A. The process to make the initial selection began in '95, yes.

Q. And when did it end?

A. The first phase ended at the end of '95. And the actual process ended in '96.

Q. Now, at the end of '95, was there some conclusion that was reached at Intel relating to this?

A. Yes.

Q. What was that conclusion?

A. So the conclusion we reached based on talking to the DRAM vendors and actually talking to Rambus as well is that from a technology point of view we would be much better off to define the next generation of Rambus technology to meet our goals.

Q. Now, would you turn to page 5 of RX-904.

Does that second bullet on RX-904 page 5 describe the process that you were describing earlier?

The bullet starting with "Considered several options for greater than one year"?

A. It describes the results of the process. It doesn't describe -- oh, actually it does, yeah, "with significant DRAM vendor involvement." That was a key part. Yes.

Q. Now I'd like to show you a document that's been marked for identification as RX-1546. If you look at the documents on the pile, the numbers are at the very bottom of the document. And the RXs are at the back.

A. Okay.

Q. Can you identify what this is?

A. Yes.

Q. And this is an e-mail from Kathy Garchow to you dated December 16, 1999; is that right?

A. Yes.

Q. And it's part of a string of e-mails, part of a string of e-mails between --

A. Yes.

Q. Who is Kathy Garchow?

A. Kathy Garchow was the technical assistant for Paul Ottelini at the time.

Q. And who is Paul Ottelini?

A. He was I think the chief operating officer of Intel.

Q. Did you have an understanding of why she wanted to know why Intel didn't choose faster DRAMs per the Intel initial investigation?

A. What she told me was that Andy Grove, who was the president at the time, asked her to go look at the RDRAM experience and try to take from the key learnings as to what we did right and what we did wrong.

Q. Looking at your response to Ms. Garchow, which is the middle e-mail, you describe why each technology other than RDRAM was not chosen; is that right?

A. That's correct.

Q. Now, looking at the first alternative, faster SDRAM, what were you referring to there?

A. We'd already made the decision to do

100-megahertz SDRAM, so we're looking at SDRAMs faster than 100 megahertz.

Q. And what speeds were you looking at at the time?

A. We looked at 133 and actually 200.

Q. Now, you say next it would not work in four-DIMM configuration past 100 megahertz without adding lots of buffers and then in parentheses you say "added cost to base platform and performance hit."

What did you mean by "four-DIMM configuration"?

A. At the time, motherboards typically had four DIMM slots or slots for memory modules, and the OEM customers that bought our products wanted chipsets with four DIMM sockets, so any memory technology that we chose in the '95 time frame had to support the four DIMM sockets or wasn't acceptable.

Q. Who was it that wanted the four DIMMs?

A. Our customers.

Q. And who were your customers?

A. People like Dell, IBM, Gateway, HP.

Q. Did you have an understanding of why they wanted four DIMMs -- slots? I'm sorry.

A. I think some of them wanted it because it could be done, so they didn't want to give up anything, but there was a good practical reason for at least three

slots.

At the time, a lot of these vendors would ship a system with a base memory in one slot, they wanted to have a slot for upgrading the memory at the time of sale, and they always wanted to have one more slot for upgrading the memory in the system after the sale. They needed a three-slot requirement.

Q. So this requirement by your customers drove part of your -- part of Intel's investigation of DRAM?

A. Yes, it did.

Q. Looking at the second alternative, DDR SDRAM, you state, "Specs did not work." Then in parentheses you have: "They have since improved but are still not showing positive margin."

What did you mean, the specs did not work?

A. When we looked at the first SDRAM proposals, the AC timings and the way that the signals were defined in the interfaces basically would not work in a system.

One of the key issues, for example, was that they had one strobe signal, which is a signal used to time the data, and it would generate by one DRAM device on behalf of another DRAM device and there was no way in the system to manage the skew between the two devices.

So there was issues like that where they just hadn't thought out the timings very thoroughly, and we were concerned that it just wouldn't work with a lot of effort we had to put in there to make them work.

Q. I see.

You said "AC timings." What did you mean by the term?

A. "AC timings" means the timings of the signals coming into and out of the device.

Q. Does AC stand for access time?

A. AC is -- no. It stands for -- AC. I don't know what it stands for. Actually AC as opposed to DC, which is the steady state. AC is the transient or non-steady-state timings.

Q. I understand. Okay.

Now, in your parenthetical you state that the specs have improved but are still not showing positive margin, and what did you mean by that?

A. We first looked at the DDR SDRAM back in the '95-96 time frame and it was -- it was not going to work then. If we fast-forward to the time frame of this memo, which is '99, they made many improvements to fix some of the issues we identified in the earlier years, but they still hadn't fixed enough to make the margin positive, meaning that it still wouldn't work.

It was better in the sense that the timing margins were less negative, meaning they were closer to working, but they weren't there yet.

Q. You said you considered DDR in the earlier years. When did you start hearing about DDR?

A. Oh, we first heard about DDR in '95 when we went out to ask for options, which was one of the options we considered, the higher-speed SDRAM. One of the options was DDR.

Q. And this is something you heard from DRAM manufacturers?

A. Yes. They claimed that they were working on this or concepts of this before that, which we weren't involved with.

Q. Now, below that you say, "Once done it would have carried us for one generation only."

Why would it have carried the industry for one generation only?

A. I believe at the time that the specs were very challenging to make work at the 200 megahertz speed, and if we spent a lot of effort and made them work at 200 megahertz, we'd basically double the performance of SDRAM from 100 to 200, but since they were so hard to make work at 100 -- actually I mean at 200, we thought scaling it beyond 200 would not work without going back

through and doing a major infrastructure change to the system, modules, connectors, so on.

Q. Now, in the last part of that sentence relating to DDR SDRAM you describe the cost to the chipset/platform to use wider interface.

What were you referring to there?

A. What we're referring to there is the same thing I mentioned for SDRAMs. If you run into a wall where the electrical performance will no longer work, you can always extend the number of devices by adding some buffers, so essentially you will take electrical signals and break them up into multiple electrical signals with less load.

And what I'm referring to here is that we could go faster potentially with DDR SDRAM by adding some buffers on the motherboard to break the signals up and make the timings work.

Q. I see.

And then in the very last part of the item related to DDR SDRAM you describe that using a wider interface with the chip -- with SDRAM or DDR SDRAM? Do you see that?

A. Yes.

Q. And what does that relate to?

A. Basically it was referring to the effect of the

buffers. What we'd wind up doing is using the buffers to create a wider memory. You know, instead of 64 bits it could be a 128-bit memory and use the buffers to bring them down to the 64 bits that's typical for our chipsets.

Q. So when you talk about 64 or 128-bit memory, you're talking about the width of the data bus?

A. Width of the data bus, yes.

Q. Now, describing RDRAM, you state at the end of the paragraph relating to RDRAM that it had a leader that had an interest in making the platform work rather than just a DRAM.

What were you referring to there?

A. When we worked on the DRAM vendors, we found across the board that they're most concerned about making the DRAM devices and they were used to providing the DRAM device and the data sheet for the DRAM device, and it was up to Intel to figure out how to use that DRAM device in the system.

So we not only had to design the chipset, but we had to design the interconnect between the chipset and the DRAMs, which would be the modules and the motherboards.

One of the attractive things of Rambus was that they worried about the whole solution. They

didn't just worry about the DRAM part. They saw the channel as a critical part of the solution, and the design trade-offs were not just made for the DRAM device but the DRAM device and the channel and the chipset.

Q. Now, going back to page 3 of RX-904, the next bullet says, "Choose a path (road map) with sufficient lead time and communicate (spec)."

What does that mean?

A. What we wanted to be able to do was to choose a path for us and the DRAM industry, and by "road map" we mean not just one step, we mean several steps that we can communicate so that the next step is a precursor to the following step, and so on.

And we wanted to do so with sufficient lead time to allow the DRAM industry and Intel products to react. It typically takes two to four years to do something new in the DRAM industry and something similar to do something new in the chipset, depending on the amount of change.

So for us to decide what technology to use with the two to four-year lead time was pretty important to both sides of being able to implement their parts.

Q. Why did it take two to four years to do this

kind of change in the DRAM industry?

A. Basically it's the latency for designs. The two-year number comes from trying to just take an existing design and trying to speed it up.

If you try and take, for example, an SDRAM, try and speed it up from 66 megahertz to 100 megahertz, you have to go back through your designs very carefully and look at all the speed paths, try to figure out how to make it go faster, and synchronize that with one of the shrinks they would normally do for a new-generation process technology, and that process typically takes a couple years to get right and to go not only to design but to validate and make ready for volume shipments.

If we're doing something quite different, as in the RDRAM case, it can be longer.

Q. So that's where the four-year number comes from?

A. Yeah. And that one would be a little bit faster than four years because they had a previous technology which we were extending, but it was more complicated.

Q. When you say a previous technology you're extending, what were you referring to?

A. Referring to the current and base RDRAMs.

Q. And what do you mean, you were extending concurrent RDRAMs?

A. What we mean is that the base electrical interface, that the concurrent RDRAMs was preserved in the move to direct RDRAM. We changed the voltage from 2.5 volts to 1.8 volts, but the basic structure of the I/O didn't change.

And the protocol for concurrent RDRAMs and the protocol for direct RDRAMs was actually quite different in that in the concurrent RDRAMs there was a multiplex protocol and used one set of signal lines to transfer addresses and data, and in direct we actually changed that to have a separate set of lines for address commands and a separate set for data.

So we made some changes, but we preserved the basic electrical infrastructure.

Q. Now, when you were involved in evaluating the DRAM technologies in late 1995, were you evaluating direct RDRAM or concurrent RDRAM?

A. We were actually evaluating concurrent RDRAM.

Q. And what was Intel's role in the development of direct RDRAM?

A. That was our development agreement or contract with Rambus, is our conclusion from concurrent RDRAM was that the technology had the ability to meet our

goals, but the concurrent RDRAM didn't. So we had to define the new technology, which turned out to be direct RDRAM, and we wanted to jointly develop it with Rambus, make sure that meets the needs of the PC.

Q. If you would turn to page 8 of RX-904.

The last bullet on that list says, "Proven ability to transfer compatible designs."

A. Yes.

Q. What does that mean?

A. That means that Rambus in the past with their base and concurrent RDRAM technologies had worked with several of the DRAM vendors in the industry and successfully taken some of the designs to production. There were production in Nintendo game machines and several other smaller applications.

So our confidence they were able to do a design, take it to the DRAM vendors, get through the design validation was higher.

Q. Now, does this describe the importance of the -- I'm sorry. Strike that.

Was Intel concerned about the ability of Rambus to relate well and work well with the DRAM manufacturers?

A. I'm trying to think back now. At this time maybe we were a bit concerned about it, but I don't

think it was a primary concern at all. Because the experience that we had talking to DRAM vendors that had used Rambus technology in the past, they spoke pretty well of it. There didn't seem to be issues.

Q. Okay. Now, do you have an understanding of the meaning of the term "revolutionary" as it pertains to the DRAM industry?

A. I believe so, yeah.

Q. And what's your understanding of that term?

A. "Revolutionary" is a term used to differentiate between evolutionary. Evolutionary typically means taking the existing standard and extending it in some way.

So SDRAM was perceived to be somewhat evolutionary in that it preserved the same pins as the old EDO memory but added the clock. DDR is perceived to be evolutionary in that it added some strobes for the data bus but preserved most of the paradigms of SDRAM.

Revolutionary talks about wholly different technologies. You know, the various versions of Rambus were perceived as revolutionary because they didn't look anything like previous ones. They had different electricals, different protocols, and so on.

Other revolutionary approaches were like

SyncLink, which never made it as a standard, but it was also proposed and important to JEDEC, and it didn't look like any of the previous SDRAM or EDO devices.

Q. So your understanding at the time was that RDRAM was a revolutionary device in the memory industry?

A. Yes.

Q. Did you have an understanding of whether RDRAM was chosen by Intel for nDRAM because it was revolutionary?

A. It's hard to answer that yes or no. We felt that we needed something revolutionary in the sense that to try to meet the performance goals we had was going to be hard, to try to meet the performance goals and get scalability beyond that was even harder, and the real chance that we had to do that was to pick a revolutionary technology.

The problem with revolutionary technologies is they're risky. They take a lot of work to get right and time. You have to go through multiple iterations typically.

So to kind of balance that, one of the attractive things of Rambus or something like the SyncLink was that they hadn't implemented parts of

their technology in base and concurrent RDRAM.

So we looked at this as a revolutionary technology from the point of view that we could get the benefits from it but evolutionary in the sense that we could rely on their past experience working with DRAM vendors to build the base and concurrent generations.

Q. If you would turn to page 9 of RX-904.

The title of this slide is Overview of Contract.

What contract does this refer to?

A. This is referring to the contract we signed with Rambus in November of '96, which started the joint development effort of direct RDRAM.

Q. If you go to the bottom of the slide, it states, "Overall system royalties is maintained at supportable level."

What does that mean?

A. It means basically one of the concerns we had in choosing the technology that had a third-party company chassis (phonetic) with it was if we were to ship a lot of chipsets that support the technology we'd put them in a very powerful position. We wanted to just make sure we had taken care of the aspect of that compared to royalties, so one of the clauses in our

contract basically said once this technology is in volume, the royalties beyond 2 percent would be provided back to Intel.

Q. Now, if you could turn to the next page of RX-904.

The second bullet on page 10 of RX-904 states, "Rambus to do the majority of industry enabling." And then it lists DRAM vendors, ASIC vendors, connectors, clocks and testers.

First of all, what do you mean by "industry enabling"?

A. With any technology, like Rambus, there's a whole bunch of supported technologies that are needed to make it applicable. Not only do we need DRAMs, but we need modules, connectors, channels, clocks to drive the DRAMs and testers to test the DRAMs and modules.

Q. And why was Rambus expected to do the majority of the industry enabling?

A. Because they had told us at the time we were working out the deal that was their job, their role, was to try and make the technology available to multiple people in the industry, and they looked at themselves as doing the designs with DRAM vendors, providing them the means to test the designs. They wanted to take care of the connectors and clocks as

well, and they would also take the controller side of the design to any ASIC vendors that wanted to implement this as part of the deal.

Q. Now, what does ASIC refer to?

A. ASIC refers to the controller side, so Intel would do the chipset for our products, but if anyone else wanted to do a controller for a different product, they typically would go to the ASIC vendor to get the device abrogated, and Rambus said that they worked with ASIC vendors and they would enable them to develop RDRAM controllers that could be used by these other vendors.

Q. Do you have an understanding of what constitutes industry enabling with respect to the DRAM manufacturers?

A. Yeah. Actually I think I went through that before. But basically it means going out and working with them to understand the limits of their technology, the benefits of their technology, trying to figure out what you want them to do, and then working back and forth to try and figure out how to get them to use the technology to the best of their ability and to take back issues and try and fix them elsewhere in the system.

So it's a relationship where you try and take

your target, try to get them to build your target, where they have problems building your target come back and try to move your target around, so it's a give-and-take relationship.

Q. Do you have an understanding what constituted industry enabling with respect to connectors and clocks?

A. More or less the same thing. Although in those cases it was much easier in that they basically wanted to know what the specifications they needed to get were and that they'd just go do it.

So in those cases it came down to writing clear, crisp what the connector needed to be, the size, the pin count, the, you know, specifications for insertions, et cetera, electrical parameters for the connector, and then let the connector vendors go away and try and implement that.

Likewise for the clocks, give them the specification of the frequency, the drive strength, the jitter, all the parameters that were important for them, and then they'd go implement it.

Q. And what constituted industry enabling with respect to testers?

A. Actually I wasn't involved as much there, so I don't know all the intricacies and details, but for the

most part it's giving them a direction in terms of what the speeds are going to be in the future, when they're going to be important and what are the critical parameters they needed to test to make sure they got testers that are capable of doing that.

Q. And why is it important that there be testers that are capable of testing DRAM?

A. Because the DRAMs are going to go into high-volume production and achieve a quality metric, they have to test them, and one of the key things for a high-speed interface is testing, and what we're doing with Rambus was way beyond what they are currently doing with SDRAM, so we knew that the testing infrastructure would be different.

Q. Now, if you'd turn to page 13 of RX-904.

And that's a slide entitled Key Messages/Issues.

Is that like a summary slide?

A. Yes.

Q. The last major bullet says "technology development challenging." What does that refer to?

A. That means after we signed the contract and started to dig into some of the details beyond the basics of what direct RDRAM needed to be, we saw there were some issues that needed to be solved that had a

lot of work left to do.

Q. I'm sorry.

A. I was going to say it wasn't all done as we hoped.

Q. And what kind of work had you discovered you needed to do?

A. The biggest issues turned out to be the packaging of direct RDRAM. The original concept proposal by Rambus was something referred to as chip file in which individual DRAMs could be inserted to a set of rails that attached to just some pads on the motherboard. And after a more thorough mechanical evaluation of that approach, we didn't think it was producible in high-volume manufacturing, so we had to come up with a different way to package these high-speed devices for the system.

And the other key area we found was the voltage level. The original proposal was to try and do a 3.3-volt device with a 2.5-volt I/O, and after some analysis, we concluded that the power the device consumed would not be acceptable to our systems and the I/O we were concerned would be too high a voltage for some of our feature process technologies, so we had to drop both voltages to 2.5 volts at the core and 1.8 volts with the I/O, and that proved somewhat

challenging because it was the first time the DRAM vendors would have to implement a technology with the low voltages.

Q. Now, you're referring to two different types of voltage, voltage for the core and then voltage for the I/O.

What is the difference between those two voltages in terms of what they're used for?

A. The core voltage is what's used by the DRAM for the internal core device or the memory storage cells, and the I/O voltage is what's used on the interface for the signals that go out of the package.

Q. So the core voltage is sort of the power cord for the DRAM?

A. Yes.

Q. Now, the first major bullet says "Package POR not yet established."

What does POR stand for?

A. It's plan of record.

Q. And what does that mean?

A. That basically means what I said earlier about the chip file, is that the original thought was chip file would be the packaging used for direct RDRAM and we concluded that it would not work and so we needed to find something. We had no specific plan that we were

going to be able to bet the project on at that point. We had work to do.

Q. I see.

Now, two bullets -- actually three bullets from the bottom you state "Industry enabling less than expected from Rambus."

What did that mean?

A. I'm trying to think back now.

You know, it's hard to state exactly what this meant. What we learned over time kind of grew, but I think at this point what we're referring to is the design package Rambus gave to the vendors initially was perceived to be a good deal. They could take this. They could implement it. It was a pretty refined process.

After we signed the deal and spent some more time with the DRAM vendors and went into details, we found out that the design package that Rambus provided was actually pretty generic and most of the vendors had to take the data from there and reimplement it using their tools and technology, so it was quite a bit more work on their part than we initially thought.

We also understood Rambus provided a lot more support to the DRAM vendors and in the way it was played out is they were trying to support to one, maybe

two vendors, and the rest of the vendors would take their package and they're pretty much on their own to implement it.

Q. You say the package technology that was provided to the DRAM vendors -- did you say that was pretty generic?

A. Yes. In the sense that they took a DRAM technology that was a combination of several of the vendors' technologies, tried -- did their design on the more generic technology, and then provided a design package to the vendors, allowing them to take it, translate it and tailor it to their technology. Because all the DRAM vendors' technologies were not the same, similar but not the same.

Q. Okay. So when we're referring to technologies with respect to DRAM vendors, what do you mean?

A. Process technologies.

Q. So their ability to make the DRAMs?

A. Well, no. I mean the actual physical process technology that they use to lay out the transistors on. The characteristics of the process technology, the speeds of the transistors, the, you know, leaks of the transistors, the capacitance of the transistors, is not all identical.

Q. I understand.

When I read you that bullet relating to industry enabling, you said back in this time period you're trying to understand what that meant.

Did this term or this phrase mean something different over time?

A. Industry enabling?

Q. Yeah. With respect to the industry enabling being less than expected from Rambus.

A. Well, "industry enabling" is a more generic term with this interaction. I think what this refers to is, as we go back and we look at how much interaction was required to do something like an EDO memory, it was somewhat minimal. We just had to talk about some of the key specs and agree and everyone did their designs.

With SDRAM we did a lot more details. We had to not only do the key specs, but we had to look at some of the key functions and we had to actually provide standard layouts for the DIMM modules.

As we went into the next-generation RDRAM, we expected the standardization to be much more difficult and require us to not only provide the standard channel DIMM layouts but to provide standard modules for the DRAM vendors to implement in their DRAMs in the high-speed I/O, and so the perception here in what

we were expecting Rambus to do was provide DRAM vendors with designs to do the high-speed interfaces and to work with them in the same way we talked about enabling in the past, which meant a fairly tight relationship, to translate whatever their generic design would be, Rambus, to the vendors' specific requirements.

So the bottom line is we were expecting a lot more effort on Rambus' part to make those designs happen in DRAM vendors as opposed to a design package being given to the DRAM vendors and the DRAM vendors having to do a lot of the work on their own.

Q. I'd like to show you a document that's been marked for identification as RX-1532.

Can you identify this document?

A. Yeah.

Q. Is this the same Kathy Garchow who sent you an e-mail earlier?

A. Yes.

Q. Do you know how Ms. Garchow came to the information that she sent to you in this e-mail?

A. I'm sorry. I missed your question.

Q. Do you know how Ms. Garchow came to the information that she sent to you in this e-mail?

A. Well, it says here it's a summary of the

one-on-one with Pete MacWilliams, so basically I agreed to meet with her, and I think we met for a couple hours, and she had lots of questions about what happened in the RDRAM experience and what the timeline was, what the issues were, what the benefits were.

She took a bunch of notes. I believe this is her playing back to me what we talked about.

Q. I'd like to ask you a few questions about the timeline that starts in the middle of the first page.

Now, the first bullet is Q395, and about what calendar period is this referring to?

A. Q3 is calendar years, so it's July through September.

Q. I always have some trouble with that. Sorry.

The third and fourth bullets state: "Up to this point in time, memory vendors were strictly focusing on lowering costs and increasing density. Intel felt the memory vendors needed to get more focused on increasing access speed."

Did you have an understanding, based on your experience of working with the DRAM industry, of why the vendors were focused on lowering costs and increasing density rather than focusing on increasing the access speed?

A. I think over time that businesses tend to do

that. If we look back to the past thirty years, DRAM technologies advanced very quickly. The density has been the primary driver. As new process technologies come on-line, they're able to take the devices, shrink them, provide lower cost per bit, and the business model that they work on is pretty much driven by that.

They lived with fast page and EDO memory several years just going through successive strengths trying to improve the density.

Q. Would you go down to the bullet that starts January 1996. It states, "Met with Rambus and memory vendors to begin next level of engagement on RDRAM."

What constituted the next level of engagement with the memory vendors on RDRAM?

A. Well, basically at that point, January of '96, would be right after we made the internal decision as to the direction that we thought was most promising. We chose the RDRAM path. Now we had to go flesh out the details and figure out if we were to redefine a new standard direct RDRAM would this be something that we could get accepted based on the relationship with Rambus and is this something the memory vendors can actually build.

So we started this in January.

Q. So after choosing RDRAM as the direction you

wanted to pursue, you went back to the industry and talked to DRAM vendors about that choice?

A. Yes.

Q. And what were the DRAM vendors telling you about RDRAM at the time?

A. At the time they were -- there are actually multiple accounts. There were some that were fairly enthusiastic about that; some that were more lukewarm in the sense they didn't really care, they just wanted to know which way we were going; and there were some that didn't like it.

Q. Now, if you turn to the next -- I'm sorry -- just the very bottom of the first page.

The subbullet states: Toggle happened when volume hits, the memory vendors migrate to the highest volume -- I'm sorry -- migrate the highest-volume memory to the new processes first, therefore giving cost advantage.

Do you have an understanding of what's meant by "toggle" in that bullet?

A. Yeah.

Q. What does that mean?

A. What it means is the transition of volume from one standard to the next.

So what we're referring to is the fact the DRAM

vendors, for operating reasons, typically will shrink their highest-volume parts first, and therefore those parts have an advantage because they get the process technologies, the best cost structures first, it's in their economic interest to do so.

And if you go to the statement above, we'd actually changed our cost target based on the discussion with Rambus to say we want to hit a 5 percent delta over the current high-volume technology on an even technology basis, meaning the same-technology-to-same-technology comparison. And we believed that with the 5 percent difference, that once the volume of RDRAM was high enough that the vendors could start shrinking those parts first, that that would overcome the 5 percent and it would become a nonissue.

Q. Now, the second page of the document, at December 1996, the timeline states that Intel communicated to memory vendor that it had chosen RDRAM and signed a license with Rambus.

Do you see that?

A. Yes.

Q. The bullet below that states, "Intel made decision not" -- and "not" is all caps -- "to negotiate a contract for the memory vendors but did add

conditions into the contract to help the industry" and then in parentheses "limit royalties."

What sort of contract was Intel contemplating that it would negotiate for the DRAM vendors with Rambus?

A. I'm not sure I can go into specifics of that, but let me tell you what the business concern was and how we responded.

The concern was that Intel chose a technology where there was a lot of volume behind the technology and that we might be putting Rambus in a very powerful position with respect to their contracts in the DRAM industry.

And so at one point there was discussion of whether we needed some sort of a blanket contract or a boilerplate contract that everyone could sign up to at a minimum to make sure there was not a problem with people signing the Rambus contract to provide Rambus DRAMs. For various reasons. In the end, we chose not to do that. We just put the clause in our contract to effectively limit the royalties once the technology went to volume.

Q. How were royalties limited once the DRAM went into volume?

A. Our contract basically said that once the

technology was in volume, royalties in excess of 2 percent on the DRAMs would come back to Intel, meaning that Rambus would have no motivation to charge more and we had no motivation to collect them, so I think basically that effectively set them at 2 percent.

Q. Below -- right below that, the timeline states, "In retrospect, this was a mistake... Rambus took advantage of the memory vendors."

Was this your understanding of whether it was a mistake for Intel to have not negotiated the contract for memory vendors?

A. I don't recall exactly, but probably. It was a mistake in the sense that what we did didn't work. There was enough ways to get around the language in our contract to charge higher royalty rates before the technology hit volume.

We'd heard back from several DRAM vendors who were not very happy with Rambus that they charged them quite a bit more initially, and so basically that's a reflection of the fact that we heard back from DRAM vendors things that we had not intended when we put the original deal together.

Q. Now, what do you mean, they were charging quite a bit more for royalties than 2 percent?

A. I don't know exact numbers, but they were 2-1/2-3 percent maybe, so they were, you know, 50 percent more, somewhere in that ballpark. I'm guessing, but they were not specific. But when we said there was a 2 percent number, they were paying more.

Q. Is that what you meant by or what you understood that she meant by "Rambus took advantage of the memory vendors," that they charged more money?

A. That was -- that's one of the key areas. I think the -- yeah, in this context that's what it was.

Q. I'm sorry. Were there other areas that you were thinking of?

A. I think, yeah, in general we heard back from the DRAM vendors many times that Rambus was taking maximum advantage of their position. They were charging higher royalties for lower volumes.

They were charging for validation, which we had not assumed, and in fact they did stop that once we found that out.

They were charging to have Rambus engineers come and help them port the design, which was not part of our arrangement, and they did stop that as well.

But the idea that we didn't have some blanket agreement in terms of how Rambus would work with the

DRAM vendors to implement our DRAM standard per this joint development agreement we perceived as a mistake. We should have tried to treat this as much more of a level playing field instead of trying to set one cap and let the DRAM vendors individually deal with Rambus.

Q. In the -- next there is a bullet that says "1H97." That stands for first half of '97; is that right?

A. Yes.

Q. Below that is a bullet that states "worked with DRAM vendors... die size; package technology," and then it states the first major issue with Rambus was the chip file package.

Was this the package issue you described earlier?

A. Yes.

Q. Below that is a second major issue and it says: "Rambus said the product could be achieved with .3u technology, Intel said .25u technology, and today we are at .18u and still not hitting cost/die size targets."

First of all, what does .3u mean?

A. It's .3 micron, which is a critical dimension that's usually used to identify process technologies.

Q. What was Rambus saying when it said that the product could be achieved with a .3 micron technology?

A. Well, they believed that they could design an interface, a DRAM interface, that would run at 800 megahertz in a .3 micron technology, which was the technology back in the '96 time frame.

Q. So I understand it, as the number gets smaller, that's a more advanced technology?

A. That's correct.

Q. What was the problem with the technology being at .18 micron and the industry still not hitting its cost and die size targets?

A. Basically it was just -- well, one assumption first is that when we go from .3 to .25 to .18, the dies all shrink, but because of the smaller dies and better process technologies, the speed also is assumed to increase. Historically that's always been the case.

And we were concerned here because Rambus had made the statement to us during the first part of the deal which indicated they could implement an 800-meg transfer on .3 micron technology with a 5 percent die delta and we were looking at .18 micron technology in this time frame, which would have been '99, and they still weren't hitting the 5 percent mu yield and they

still weren't hitting high yields at the end.

Q. The next bullet states, "Realized at this time that the enablement would take much more effort from Intel than originally forecasted."

You were describing the enablement before. What effort was required from Intel?

A. At this time I think the major part of the effort was believed to be the packaging solution, the fact that we had to move from chip file to RIMMs.

Q. Now, if you look at the second bullet under the second half of 1997, the timeline states, "Tried to get Rambus to pay for the enablement areas which we believed they signed up to support and did not."

And these were the enabling -- enablement areas you described earlier on?

A. Yes.

Q. In the first half of 1998, the first bullet states that memory vendors said that they did not want to do PC133 and no lower voltages, but by the middle of '98 they become more positive on PC133.

What was the importance of the attitude of the memory vendors toward PC133 to Intel?

A. Well, at the time we actually looked at PC133 multiple times through the years, and what this is saying is that we went out in the first part of '98 and

asked again are they interested in doing a slightly faster SDRAM and/or doing a lower-voltage part based on mobile requirements.

At the time they told us no, don't do it. They were concerned that they had already optimized their die for PC100 or 100 megahertz speeds and the yields for 133 would not be very good and the effort to redesign the dies to get the yields up to 133 would take some effort and they wanted to spend their effort doing RDRAM designs and driving cost out of the PC100.

Q. Well, did that attitude have an impact on Intel's decision to support PC133?

A. Initially, yeah. In '98 when we asked them, we were actually looking at doing this possibly maybe for a mobile part or for a low volume -- excuse me -- a low-price-point part where we could restrict the four DIMM assumptions we talked about earlier to two or three.

So we were out there seriously asking them does this make sense, and they did not like the idea of lower voltage and/or higher speed because of the need to do a new design.

Q. I think we described -- we talked a little bit about the four DIMMs and what the importance of that was to your customers.

What was the impact of having multiple DIMMs or having four DIMMs rather than three or two on the DRAM designs?

A. On the DRAM designs it didn't have much of a direct impact. On the I/O interfaces on the DRAMs it would.

If you look at the I/O interfaces between the chipset and the DRAMs, you have a connection in the motherboard and some stubs that go off in the various modules, and all the interconnect board and the subsystem modules creates some timing loss, so as you increase the number of modules and the number of loads, you have more timing loss. If you want to run higher speeds, you have to have less modules for a given point in time or technology.

So the design of the DRAM to a 133 would not change dramatically, but the design of the system to use the 133 would have to.

Q. I see.

You used the word "stubs." Is that a specialized term?

A. Actually it's a pretty standard term in the electrical area. A stub refers to the fact that the line does not have a start point, the driver, and the endpoint, the receiver, but it actually has multiple

receivers, and the receivers have little visual stubs.

So you can envision a chipset has a line going all the way down through this array of DIMMs and each DIMM connector at the connection of that line goes up. Those connections are maybe one to several inches depending on the signal and they represent stubs (indicating). Electricals are very difficult to manage at high speeds.

Q. And the next bullet states, "Worked out packaging and I/O voltage issues, but cost target remained a major concern."

You described the I/O voltage issues previously I think; correct?

A. Yes.

Q. In the second to last bullet for the first half of '98, the timeline states, "Architecturally the costs were not comparable: four independent banks in SDRAM versus 16 dependent banks in RDRAM... overhead due to the redundancy cost and the interface."

What's the importance of the number of banks to the cost of the DRAM?

A. Well, in this case -- let me go back one step and just supply what happened to create 16 banks.

One of the key issues was to create with RDRAM was the amount of bandwidth we need into the core of

the DRAM device was much higher than what they were doing with SDRAMs. So we couldn't just take an SDRAM core and use it.

Rambus came up with a technique to use the same basic core, turn it on its side, and the result of that was we got 16 dependent banks, "dependent" meaning one bank would have actually shared some of the circuitry with adjacent banks.

So in the SDRAM the four banks are completely independent and can be accessed independently. In the dependent bank case, we'd actually have to restrict the accesses so we'd never access adjacent banks.

And at the time we made a decision to go with independent banks, the comparison that Rambus and the DRAM vendors provided us showed little to no overhead. And as they actually implemented it and they started adding redundancy, which is a technique in DRAMs to add extra storage cells and then in a test if they have bad cells they can actually swap out good ones for bad ones and they can save the die as opposed to throwing it away, when they started adding redundancy, the core sizes started growing.

The importance of banks in redundancy is each bank needs more redundancy structure, so because we have more banks, the amount of redundancy needed to get

the same amount of coverage for manufacturing was higher and the result of that was, you know, much larger core size in fact than when we did the original analysis back in '97.

Q. You referred to swapping redundancy and swapping out bad DRAMs. Is that used -- does that use fuses in any of the DRAM --

A. Yes, it was.

Q. Why was the cost target a major concern at this point in the first half of '98?

A. From our point of view, we couldn't really affect how DRAMs are priced because that's an issue between the customer and the DRAM vendors, but the one thing we could have a lot of effect on is what the basic cost structure was.

It was pretty important for us to try and keep the cost structure of the new technology with the original goal was the same as SDRAM, and I think we moved it to the 5 percent number after working with the vendors and Rambus, but we wanted to keep it very, very close so that the DRAM vendors had the flexibility to price the DRAMs in the same ballpark as the existing volume DRAMs without losing money.

We wanted them to basically make the same amount of margin on these devices or more than they had

in the previous devices.

Q. Why was that important to Intel?

A. I think it was important to Intel because for the -- for the DRAM vendor to be motivated to build a new technology, they had to see a way to make money at it. And if new technology costs a lot more money and they couldn't sell it for a large premium so they can make up the cost difference, they wouldn't make money. And if they couldn't sell the new technology at a cost pretty close to the old technology once it got in volume, we couldn't see how the market would transition.

Q. In the second half of '98 the timeline states, "First saw product in June."

What does that refer to?

A. It means the first samples, the first RDRAMs showed up in June.

Q. And the third bullet of the second half of '98 states: Set up good process to make technology work but did not focus enough on how to break it.

What does that mean?

A. That refers to our own development process.

If you look at the bullet above -- or actually the last bullet in the first half of '98, it says we built three test chips, two test boards, did thermal

and electrical analysis. And I think part of the problem of trying to do three test chips, two test boards is we had limited resources.

We got to the point in each one where we were able to make the key things work, but we didn't have the time or resources to go in and try to figure out how those things would break, so trying to apply all the boundary conditions, do all what we call margining, where we vary the voltage, vary the timings, try to create all the worst-case conditions that the technology is likely to see over its life and try to figure out where its weaknesses are and go fix those.

We didn't do enough of that.

Q. Two bullets below that it states, "Electrical side created several issues which were found late in the process."

What did you mean by that -- what was your understanding of what "electrical side" means?

A. It means that the -- you know, actually later in '99 we had several issues we uncovered with the channel itself, so the silicon of the controller and the DRAMs were okay, but we found some of those boundary cases on the channel which weren't okay. And we had errors that would show up like once every few

days or once a week that we'd have to capture, track down and fix.

Q. Now, you referred to a channel. What is a channel?

A. The channel is the connection between the chipset and the DRAMs. It's the physical traces on the motherboard and the modules.

Q. Sometimes is that called a bus as well sometimes?

A. Yes.

Q. And what did "late in the process" mean?

A. It means that we expected these things to have been sorted out way back in the time where we had test chips or even before we had our direct RDRAM and that Rambus had experience with this channel in previous generations and we were actually surprised to find issues popping up in '99 that were still related to the channel.

Q. What was the effect of the electrical issues on the adoption of RDRAM by the industry?

A. Actually, it was pretty significant. We had to delay the launch of our first product twice because of these electrical issues. The first target was I believe June '99 and we had a whole bunch of them that we found in the first part of '99 which caused us to

delay that until October. And then just before the launch in October we found one more which caused us to delay to November.

The DRAM vendors were preparing for launches and had built a lot of RDRAMs for the June launch before we found the problems and were left holding a lot of RDRAMs, and having built a bunch assuming the October launch wasn't as big of a problem then because we did launch it in November, so there was a month delay, but the first one was a pretty big deal to some of the vendors.

Q. You referred to the first product. What was the name of that product?

A. The first product was the 820.

Q. Was that also sometimes known as Camino?

A. Yes.

Q. There's a note below that bullet describing: What we believed the agreement between Intel and Rambus was regard bringing the product to market.

Do you understand who the "we" was in that sentence?

A. I'm sorry. Oh, okay.

The "we" here is Intel.

Q. I'm sorry?

A. Is Intel.

Q. Below that is a subbullet stating, quote, "Rambus focus: Modules, RACs (interface to chipset) and working with memory supplies/channel."

What was your understanding of Rambus' responsibilities regarding the implementation of RDRAM?

Well, actually -- first of all, what is a RAC?

A. A RAC is the interface in the chipset to the RDRAM channel.

Q. Below that is a statement: "The only thing that Rambus did well was the RAC."

What was your understanding as of November of '97 of Rambus' performance in the implementation of RDRAM?

A. As of November '97?

Q. Yes.

A. This is a comment back in '98.

Q. I'm sorry. Actually in November of '99.

A. Okay. Well, let me address this comment. I think, you know, we were very pleased with the way they worked with Intel in doing RAC designs for our chipsets. They did designs. Our engineers worked with their engineers and the designs were very good. We had no issues, no complaints. It was a very good experience.

The concern part here was the memory suppliers, at least many of them were coming back to us with issues of one form or another. I think we went through some of those earlier. The channel. We found a lot of issues in the channel that we thought were already wrung out but led to the delays in our product. We had not spent a lot of energy on that, actually largely because we thought that was already done.

And the other comment in this line was modules. We were assuming that -- Rambus had acknowledged at one time that they were going to implement the RIMM modules needed for the direct RDRAMs, and I'm not sure exactly when, I think it was late '98, we concluded that it wasn't going to happen, that the designs they did were coming too slow, they needed too much work, they weren't going to meet production and we were not confident on that, so we actually gathered up a team and did them ourselves.

Q. Did the --

A. Did the RIMM modules.

Q. RIMM is what?

A. I'm not sure if it means anything. RIMM is a term used to talk about the memory module the RDRAMs sit on and that unit plugs into the connector.

Q. So it's like a DIMM for SDRAM?

A. Like a DIMM for SDRAMs, yes.

MR. DAVIS: Could we take a break in the next 10-15 minutes?

JUDGE McGUIRE: Sure. We'll take off for ten minutes.

I want to inquire, are you going to go through each and every line on this document and ask him his understanding on it?

MR. DAVIS: No.

JUDGE McGUIRE: Because I'm a little concerned on where we're going with this.

All right. Let's take a break. We'll be back in ten minutes.

(Recess)

JUDGE McGUIRE: Mr. Davis, you may proceed with your inquiry of the witness.

MR. DAVIS: Thank you, Your Honor.

BY MR. DAVIS:

Q. Now, just turning back to the first page of RX-1532, do you see the statement "three key learning areas" that's right beneath the summary?

A. Yes.

Q. What was your understanding of the business aspect, that business aspect of the key learning areas?

A. Our perception when we started to deal with Rambus was that memory vendors were doing multiple DRAM technologies that they'd licensed from Rambus. There are NREs to do the technologies and then royalty agreements beyond the NREs.

And one of the items we discussed with Rambus up front and we believed they were motivated to do was to get out of the mode of turning the technologies charging NREs for each new type and to get into the mode of making their money based on royalties.

So our belief was that their motivation would be to try and see RDRAMs go into high-volume manufacturing and to make their money on royalties and to provide a valuable service to the DRAM industry to do so.

MR. DAVIS: And I actually don't know if my mike is working at this point.

JUDGE McGUIRE: Well, I can hear you.

Are you having any trouble over there hearing him from the opposing side?

MR. DAVIS: I'll just keep on moving closer and closer.

JUDGE McGUIRE: All right.

BY MR. DAVIS:

Q. What is key learning areas, what does that term

mean?

A. It means what we took away from the experience and we were wanting to apply to the next time we do something like this.

MR. DAVIS: I see.

I move to admit RX-1532.

JUDGE McGUIRE: Objection?

MR. STONE: No objection, Your Honor.

JUDGE McGUIRE: Entered.

(RX Exhibit Number 1532 was admitted into evidence.)

BY MR. DAVIS:

Q. We spoke earlier about PC133?

A. Yes.

Q. What is PC133 SDRAM?

A. It's basically SDRAM sped up to 133 megahertz.

Q. When the contract was signed between Intel and Rambus at the end of 1996, did you foresee that Intel chipsets for personal computers would use PC133 SDRAM?

A. At the time it wasn't our plan. We saw no reason. We wouldn't preclude it per se, but we didn't plan to do it, no.

Q. And why not?

A. Well, because it's most efficient to focus on one memory technology and put all the energy behind

making that one efficient, standard and at the lowest cost possible, and because of the four-DIMM constraint we had for a lot of our products, we thought, after PC100, moving to direct RDRAM would be the most efficient way to use our resources and that of the DRAM vendors.

Q. Did Intel eventually start to use PC133 SDRAM with its chipsets?

A. Yes.

Q. And do you have an understanding of why Intel began using PC133 with its chipsets?

A. It was a combination of things and one of the key ones was our customers were asking us to do it. And I think the motivations behind that were fundamental in twofold.

First, the four-DIMM constraint that I talked about earlier had gone away because the model for how PCs were shipped was no longer the same in that a lot of people were selling PCs directly to end users that didn't need multiple memory upgrades, two DIMMs was quite fine.

Coupled with the fact that the RDRAM prices were prohibitively high for it to be a waterfall throughout the entire product line, so they wanted to have something faster than the PC100, but they couldn't

afford to use the direct Rambus.

Q. I'd like to show you a document that's been marked for identification as CX-2527.

And this is an e-mail March 18, 1999 from Will Swope to a number of people, including you; is that correct?

A. Yes.

Q. Who is Will Swope?

A. Will Swope has worked at Intel for quite a few years. At this time I'm not exactly sure his official job, but I believe at this time he was asked to come in and take a look at the RDRAM program, try and figure out from a more strategic point of view if this is the right thing for Intel to be doing, should we be looking into other things, and basically as an outsider come into the program and give us a -- his assessment.

Q. So then would you turn to the second page.

A. Yeah.

Q. It says "Minutes/AR Assignments from RDRAM Review."

Is that the review you're referring to?

A. It was one of them. Yes. He had multiple meetings and a couple of reviews at that point. I think this is one of the meetings.

Q. And at the time that these meetings occurred, was Intel's chipset design geared to work solely with RDRAM?

A. Actually on multiple chipsets. In this time frame we had SDRAM chipsets that supported up to PC100 and we had RDRAM chipsets, both.

Q. And then going forward, were you planning for personal computers to continue to use PC100 as well as RDRAM?

A. We assumed we'd use PC100 for a few more years in the low-end segments, but we assumed for the most part that those applications would use existing chipsets in multiple operations, so our new chipset designs were focused on RDRAM.

Q. Now, did that turn out to be an issue for Intel, that the higher-end, newer chipsets worked solely with RDRAM?

A. Yeah. Yes, it did. In that the higher-end chipsets that work with RDRAM we assumed that within a year they would move down to lower price points in the same way the PC100 was going to take the chipsets of '98 and move down to lower price points.

So the fact that we saw a limited supply and a fairly high price for RDRAM made us very concerned that those higher-end chipsets could indeed come down to the

lower price points a year later.

Q. Now, in this -- in the summary of these minutes, Mr. Swope says that most aspects of the Rambus transition have been more difficult than we anticipated, and he described that.

Then he says, "To that end, we believe that a strategy that puts our chipset and value processor line dependent solely on Rambus is no longer viable."

Was Intel planning on putting its chipset and processor value line dependent solely on Rambus?

A. Over a period of time, yeah.

At the time, we had a product which we referred to here as Mobile Timna. And this product was an integration which included the chipset, the graphics and the CPU in one device. And due to a lot of benefits of RDRAM technology, low pin count and high performance, I think it was believed that it would be an ideal technology for creating that product.

Basically the substance of the cost that we had back in the '96 time frame, we thought that after RDRAM had been in production for a year or so that it had the ability to come down in price and be competitive with SDRAM, and therefore we chose the RDRAM technology to use on the Timna product which we believed would be the

vast majority of the low-end products and therefore we were putting our own neck at risk there.

Q. And Mobile Timna worked solely with RDRAM; is that accurate?

A. Say again?

Q. Mobile Timna worked only with RDRAM?

A. Yes. That's correct. The processor itself.

Now, we did design a translator MTH which took the -- we designed a memory translator or an MTH which connected the RDRAM channel from the Mobile Timna to SDRAMs, so it was a bridge to use the old technologies.

Q. And that's a mobile -- I'm sorry -- the MTH --

A. Yes.

Q. -- is described below here. Is that accurate? It says, "The MTH II is late."

A. Actually MTH II was a second generation of that product. The first MTH supported only PC100 and the MTH II was going to support PC133 as well. And that's the one they're referring to as being late.

Q. Now, was the MTH ever used with Intel processors and chipsets?

A. Yes.

Q. Were you involved in this RDRAM review? Do you recall?

A. Yes.

MR. DAVIS: I move to admit CX-2527.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2527 was admitted into evidence.)

BY MR. DAVIS:

Q. Did Intel begin considering at some point what RDRAMs -- or what DRAM technologies should come after RDRAM, direct RDRAM?

A. Yes.

Q. Were you involved in that consideration?

A. Yes.

Q. I'd like to show you what's been marked for identification as CX-2519.

At the top of the page it says "Memory MRC."

What does MRC stand for?

A. It's management review committee.

Q. And what is that?

A. MRC is typically a group of people that meet regularly to review a program and give it some basic direction or review direction, where we're headed.

Q. Were you involved in this memory MRC?

A. Yes.

Q. What was your role?

A. I was the coordinator of it.

Q. If you look at the agenda on the first page, it starts with Timna memory decision.

Where you describe Timna as mobile, is that the same thing as --

A. Yes.

Q. They're the same thing?

A. The same things.

Q. What was the memory decision?

A. I'm trying to -- there was actually a couple -- I'm trying to think of the time frame. I think in the '98 time frame the decision was whether or not to make the commitment to RDRAM based on Timna.

Q. If you'd turn to page 7 of CX-2519, it says, "Why renegotiate contract?"

Do you see that?

A. Yes.

Q. Do you know what contract this is referring to?

A. Yes.

Q. And what contract is it referring to?

A. I think it's the November '96 contract we had with Rambus.

Q. And then it says "environmentals."

What does "environmentals" mean?

A. It describes what the situation was at the time.

Q. The first bullet underneath that says, "Direct RDRAM ramp has high risk."

Did you believe that direct RDRAM ramp had high risk at the time?

A. Yeah.

Q. And what was the cause of that risk?

A. Well, yeah, it lists it pretty much here what we thought at the time, the first point being that we were hearing a lot back from DRAM vendors and others in the industry about Rambus and the business interface or dealings that happened with Rambus and that it was not very positive. It was pretty much anti-Rambus, as this statement says.

There was some concern about the Asian currency crisis. That came and went.

And there was a concern that we had planned a very aggressive ramp for our first chipset Camino and that the DRAM industry's ability to build enough of our DRAMs at a reasonable price to meet that ramp was not clear.

Q. Now, earlier you described some of the issues relating to the DRAM vendors' attitudes towards RDRAM relating to royalties.

As of March of '98, were there additional concerns being voiced to Intel about the relationship between the memory manufacturers and Rambus?

A. Yeah. I think royalties was one of the main issues they brought up, but there was other issues regarding IP and the flexibility that they had to use Rambus technologies in other parts of the product line, what they I guess would refer to as noncompatible products.

So actually to step back for a second, it was fairly rigid what was being asked Rambus willing to license their IP for use with direct RDRAM only that met the direct RDRAM specs at least.

So anything that was direct RDRAM that didn't quite meet the specs for any derivative products that weren't quite like our direct RDRAM is were precluded from that deal.

The memory vendors historically are used to doing that in that they could take out-of-spec parts and/or do derivative parts, and they work with their customers to figure out what they want, what they're willing to accept, and they can do business on that basis. And some of them were -- some of them came back to us and said that that was a real problem, that they couldn't do business that way.

Q. And what was the problem that they were describing?

A. The DRAM vendors?

Q. Yeah.

A. That they couldn't take the Rambus technology and either sell out-of-spec parts where OEMs had applications that could use them and/or do derivative parts that didn't exactly meet the direct RDRAM specifications.

Q. And what are out-of-spec parts?

A. Out-of-spec parts would be parts that do not meet the data sheet published by Rambus.

Q. And there were people who would want to buy those out-of-spec parts?

A. Yes.

Q. If you go to the next page, it has the title What Intel Wants. And the third bullet under the contract items is: Intel gets the rights to use IP in competing memory interface.

Do you see that?

A. Yes.

Q. Do you have an understanding of what IP is being referred to here?

A. Yes.

Q. What is that?

A. It's the Rambus IP.

Q. Did you have an understanding of what "competing interfaces" meant?

A. It meant interfaces other than the RDRAM interface, direct RDRAM interface.

Q. Did you have an understanding of why Intel wanted to get rights to Rambus' IP for competing interfaces?

A. I think that was part of a bigger concern we had, that in working closely with Rambus there was a lot of their IP that was becoming known to a lot of engineers at Intel and we were very concerned they were going to start using this IP probably inadvertently in many cases in other parts of the product line and we wanted to try and create a situation where we would not have IP issues in the future.

It wasn't an intent to use anything at that point, but we didn't want to have issues in the future. We wanted to deal with them right then and there.

Q. Do you have an understanding of when Intel started thinking about using alternative DRAMs to replace RDRAM?

A. To replace --

Q. To follow on?

A. "Follow on" is a little bit better word there in the sense that, you know, from day one we assumed that to be the case. We were looking in '95 for what the best technology is for the next step and we assumed that after direct RDRAM and its five-year window that we'd do the process again, we'd talk to Rambus I'm sure, we'd talk to vendors in the industry and we'd make a decision then. There was no guarantees that it would be a next-generation Rambus solution. It would be whatever the industry chose to make it.

Q. Do you have an understanding of when that understanding was passed to Rambus?

A. I thought it was pretty much up front. You know, they knew going into the deal that the way we executed the direct RDRAM program, the way the industry perceived the direct RDRAM program, would be key to them being part of the next generation.

Q. You're saying that the way the industry perceived Rambus would be key to Rambus being part of the next generation?

A. That's correct.

Q. What did you mean by the way the industry perceived Rambus?

A. Well, you know, from our point of view, the

way that the development was set up is we viewed Rambus as providing many of the enabling capabilities that we provided back in the PC100 days, and if done correctly, they would not only provide designs to the industry, they would provide validation support, they'd help vendors get to market more quickly with compatible parts, and as a result, vendors would actually do better and make more money and would appreciate having a third party take on a lot of the responsibility of solving the compatibility/interoperability issues that they had to struggle through in the past.

So what we had hoped for and what we had worked for actually in the first few years of this deal was to try and make Rambus a value-added part of this whole industry infrastructure, DRAM vendors to build DRAMs, Intel to build chipsets, and Rambus provides all of the glue to make the enabling pieces work and therefore they would be perceived as valuable.

Q. And you've been discussing the relationship between R -- between Rambus and the DRAM manufacturers.

Did you or any of your colleagues at Intel ever discuss with Rambus the relationship between Rambus and the DRAM manufacturers?

A. Many times.

Q. And what were those discussions about?

A. Basically we were relaying the messages we were hearing back from the DRAM industry. We were asking, you know, what Rambus was going to do to fix this, because it was our belief that unless they could fix the issues they have with other people in the industry, unless they could be perceived as a value-added part of the enabling process, that RDRAM would be very difficult and beyond RDRAM wouldn't be possible.

Q. Now, who at Rambus were these discussions held with?

A. They were held with Bill Davidow, Geoff Tate, Dave Mooring, Al Roberts.

Q. Who is Bill Davidow?

A. Chairman of the board.

Q. Who is Geoff Tate?

A. The CEO.

Q. And did you mention Dave Mooring?

A. Dave Mooring.

Q. And who is he?

A. He's a vice president. He's had several different jobs but generally been a vice president. I think he's now the president.

Q. And who is Al Roberts?

A. He was the vice president of engineering during the early part of the engagement.

Q. And what were the responses of these people from Rambus to the conversations that you or people at Intel had regarding the relationship between Rambus and the DRAM manufacturers?

A. In most cases they were very concerned about it and they acknowledged a lot of the issues and wanted to fix them.

I think some of the responses were also in the nature that they couldn't do anything to fix them, they didn't have -- they were concerned about how much resources and money they had to actually go out and do some of the things they needed to do in order to fix these.

So it was kind of they understood some of the issues, they seemed to be sensitive to the fact that they needed to fix some of them. We didn't see much action in many cases. What actions we did see didn't do much to fix the issues.

Q. What sort of actions did you see?

A. At one point they actually went out and -- well, let me step back a second.

Some of the nuisance issues like charging people to go provide design help, like charging for

validation services, they just -- once we pointed out there was an issue, they just stopped doing that.

Providing some incentives for the DRAM vendors to ramp, there was a time when they were talking about providing them some warrants and some value actually being first to market where they could actually get Rambus stock and make some money.

Q. Did you say warrants?

A. Warrants. Warrants in Rambus stock.

Q. I see.

A. They did do some of that.

Q. But you said there were instances when you also didn't see an effort relating to --

A. Yeah. Well, at the end, the DRAM vendors still were telling us that there was a problem. They somehow weren't able to take the step to get the DRAM vendors to perceive them as a key part of the value added to the industry.

Q. I'd like to show you what's been marked for identification as CX-2521.

This is a letter from -- to Pat Gelsinger from Yoon-woo Lee.

Who is Pat Gelsinger?

A. Pat Gelsinger is a vice president of Intel and during the early part of the RDRAM program he was

actually the sponsor. He's now the corporate technology officer.

Q. You said he was the sponsor?

A. Yes.

Q. Of what?

A. The Rambus sponsor. He's the executive sponsor at Intel.

Q. And what does that mean at Intel?

A. That means he's the vice president who's in charge of this program. It was initially signed by Carl Everett. He was Intel's executive sponsor, but he left Intel about a month or two after the program, after the contract was signed, and Pat picked up the responsibility.

Q. And do you know who Yoon-woo Lee was?

A. Yes.

Q. And who was that?

A. I believe he's the CEO of Samsung. He's the one responsible for the DRAM business and more, but our dealings with him have been in respect to the DRAM business.

Q. Now, if you look at this, the first paragraph of this letter -- the one below "How are you doing?" -- the last sentence reads: Unfortunately, we are seeing some difficulties relating to the Rambus program as

described below.

Do you recall Samsung being one of the DRAM manufacturers having difficulties with the RDRAM program in the fall of 1998?

A. Yep.

Q. Did the fact that Samsung was describing difficulties in the relationship with Rambus concern you?

A. Yes.

Q. Why?

A. Because Samsung was one of the most aggressive suppliers of new memory technologies and we worked with them quite extensively in the SDRAM generation and were working with them in the RDRAM generation.

Q. Do you recall the types of difficulties that Samsung was having with RDRAM in the fall of 1998?

MR. STONE: I would object that the question as framed calls for this witness to testify as to facts of what was going on at Samsung. I think all he can testify to is his understanding based on what Samsung told him.

JUDGE McGUIRE: Sustained as to that.

BY MR. DAVIS:

Q. Could you describe your understanding of the difficulties that Samsung was having with the RDRAM

program in the fall of 1998.

A. Yeah. Back in '98, I think they're one of the people that were telling us that they were concerned about the restrictions that they had. They wanted to be able to sell the out-of-spec devices and do derivative products.

They were concerned about the amount of royalties they'd be paying in the long haul as they're a very large volume supplier of memories and expect to be a very large supplier of RDRAMs and saw no real motivation to increase their volumes because the royalty structure was pretty flat and they wanted to see that somehow change.

They were also concerned that they were spending a lot of their resources and energy making the RDRAM implementations work. They weren't getting as much help from Rambus as we'd thought, and as such they thought, you know, they should be somehow compensated in the whole process or amount of effort they were putting in.

Q. The first paragraph -- sorry -- the last paragraph at the bottom of page 1 describes an issue with the Rambus validation program.

Is that what you were referring to relating to trying to sell out-of-spec parts?

A. Let me read this.
Q. Sure.
A. Which sentence are you referring to?
Q. It's really the second sentence of that paragraph.
A. Yes.
Q. Now, if you go to the second page of CX-2521, in the second to last paragraph of that page, it says, "Nowadays, there are some difficulties in investment of back-end equipment for RDRAM production to support stable supply of RDRAMs to our customers."
Do you understand what that's referring to?
A. Yes.
Q. What is that referring to?
A. Well, the belief at this point from Rambus -- excuse me -- from Samsung was that they needed to buy quite a few testers to test the RDRAM products in the volumes that we projected.
Q. Was there difficulty obtaining testers, to your understanding?
A. I think obtaining the testers wasn't as big an issue as paying for the testers.
Q. What do you mean by that?
A. Well, coming up with money to pay for the testers and timing their purchase of the testers to the

exact ramp. They didn't want to buy the testers too early because they'd have them sitting around depreciating. They didn't want to buy them too late because they wouldn't be able to ramp the volume. And they were also concerned with just the magnitude of the amount of money they'd have to spend on testers and the ROI associated with that and whether it made sense to build that much RDRAM initially.

Q. Do you have an understanding of how much money they would have to pay for testers at this time?

A. I probably did at the time, but I don't recall.

Q. Did you have an understanding of how much the testers cost themselves?

A. I don't recall the exact number, but it was -- it was in the $5 million range. You know, it wasn't hundreds of thousands and it wasn't hundreds of millions. It was in the $5 million range.

Q. And did you have an understanding of about how many testers Samsung thought would be required to produce the RDRAM?

A. Not exactly, no.

Q. Do you have an understanding of the general area?

A. Well, it was more than one, so maybe to expand a little bit on that, I think the understanding we went

into the project with was that the vast majority of the RDRAM tests could be done on the existing testers and a very short test for the high-speed interface only would be done on the new high-speed testers. Therefore, the amount of volume you could ship with a new tester would be very large.

In practice, the DRAM vendors didn't do that. That was okay for low-volume production, but for high-volume production they wanted a production tester that could do the whole test, high speed and the rest of the DRAM core. And that led to this need to buy more testers, so this was changed.

Q. Did you have an understanding of why the DRAM vendors wanted to use the tester to test the entire DRAM?

A. It was just -- it was how their manufacturing model went. They wanted to have one tester to do everything. You know, as to why in their model it didn't work, I don't have a lot of details.

MR. DAVIS: I'd like to move CX-2521 into evidence.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2521 was admitted into evidence.)

BY MR. DAVIS:

Q. One of the technologies that Intel considered for nDRAM was the DDR SDRAM in 1995 and '96; is that right?

A. That's right.

Q. But at the time you didn't think that DDR would work; is that accurate?

A. That's correct.

Q. Did there come a time when Intel began reconsidering DDR at least for servers?

A. Yes.

Q. And why was Intel starting to consider DDR for servers?

A. Well, two things happened. One, DDR itself was maturing, so when we looked at it back in the 1996 time frame, it didn't look like it would work with -- you know, you can make most anything work if you spend enough energy and time on it. So if we fast-forward in time to more of the 1998-1999 time frame, the industry had done a lot of work on DDR addressing the issues we saw back in 1996.

Coupled now with the server requirements were very different in desktop. Desktop, because they want to have very small minimum memory required high bandwidth out of memory devices. We wanted to get full

performance in the minimum size memory.

In the case of servers, they're different. They typically have very large memory subsystems and they can get bandwidth without having high bandwidth at the device, so they can use many devices in parallel, wide data paths. You know, memory size isn't an issue. Cost is less an issue.

So we perceived they'd keep using SDRAM for quite some time and just use wider interfaces.

They were looking for an evolution of SDRAM and DDR looks pretty good from that point of view.

Q. Do you have an understanding of why the server manufacturers were looking for an evolution from SDRAM?

A. They were just looking for a way to get more performance, and you know, they can grow the number of channels to some extent, but at some point they wanted faster memories.

Q. I'd like to show you what's been marked for identification as CX-2529.

The top e-mail is an e-mail from John Miner to, among other people, you on May 29 -- wait -- May 27, 1999?

A. Yes.

Q. And below that is an e-mail from Paul Close to,

among other people, you on May 26, 1999.

A. Yes.

Q. Who is John Miner?

A. At the time John Miner was the VP in charge of the server product group at Intel.

Q. And then who is Paul Close?

A. Paul Close was in the server product group and at the time he was looking at memory technologies for servers.

Q. Now, below the first paragraph there's a title that says "Main questions remain the same: Is server memory strategy POR competitive" and "Do we need to add DDR on Intel server memory road map."

Do you remember these being important questions related to Intel's server strategy?

A. Yes.

Q. And what does it mean for server strategy to be POR competitive?

A. Well, the current server strategy was to continue to use SDRAM until RDRAM achieved some volume in the desktop, and after it achieved some volume at the right cost points there was a way to use it in servers as well. And the question being asked was would there be -- would RDRAM achieve the cost points in time to be competitive for the server products or

would we need a technology after SDRAM that provided more performance.

Q. So below that where it says "Market Feedback," it says "FAE feedback on 870 memory."

What is FAE?

A. Field application engineer.

Q. So the bullets below that, did you understand those to be a summary of the market research that Mr. Close did?

A. Yep.

Q. The second bullet says: "OEMs skeptical that RDRAM issues will be resolved, some are waiting to see progress."

Do you remember hearing this?

A. Yep.

Q. And how did this affect Intel's strategy for memory for servers?

A. The "OEMs" means our customers. Our customers were skeptical of the RDRAM issues being resolved, their issues being availability and price, and because server design cycles are longer than desktop, they were going to make some decisions in terms of what they were going to build for a much longer time frame. It was a very pretty key factor in their decision process.

Q. You said server design cycles are longer than desktop?

A. Yes.

Q. How long are server design cycles?

A. When they design a platform, it usually takes a year or two to actually do the platform and it takes another six months to a year to actually -- to get it off and qualified into production, and then it typically would stay in production for a couple years.

In the case of desktops, the qualification can occur as fast as a quarter and the platform typically ships for only about a year.

Q. And when is the memory decision made with respect to server strategy?

A. It's typically made about a year after the platform.

MR. DAVIS: I'd like to move CX-2529 into evidence.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2529 was admitted into evidence.)

BY MR. DAVIS:

Q. I'd like to show you what's been marked for identification as CX-2535 and CX-2536. We'll start

with CX-2536.

Do you see this is an e-mail between yourself -- from Abid Ahmad to you and Patrick Gelsinger?

A. Yes.

Q. On September 1999?

A. Yes.

Q. Who is Abid Ahmad?

A. He's responsible for the memory enabling team at Intel.

Q. Does he work for you? Do you work for him?

A. Not directly. He works in the TMG or technology manufacturing group at Intel, which is responsible for all the purchasing, the fabs, the factories, et cetera. And he works in the procurement group there, which has the relationships with the DRAM vendors.

Q. The first e-mail is an e-mail that's on the top called Backwards Compat.

Do you have an understanding of what that means?

A. I think so. Yeah.

Q. And what does that mean?

A. Well, typically what we look for in memory technology is being able to take multiple steps while

maintaining backwards compatibility with the previous one and then every once in a while we have to make a change.

So for example, if we look at SDRAM, when we did the PC100 application, we made sure it was backwards compatible with the 66 megahertz, and therefore you could buy a 100 megahertz module and populate your PC100 designs as well as your old 66 megahertz designs, so backwards compatibility reduced the number of line items the vendor would have to carry to satisfy his whole product line.

Q. Now, right below this it says, in quotes, "Who makes the call on backward compatibility?"

Do you recall that being an issue between Intel and Rambus?

A. Yeah. It was somewhat. It wasn't one of the high ones, but we did have some discussion on that.

And in fact, it goes back to the original contract we signed and we were concerned that Rambus' track record had been to define new technologies, get license fees and not rely on royalties, and based on what we'd like to see in our business, we wanted to see direct RDRAM show up at one speed and have multiple speed creates over time that maintained backwards compatibility, which would be somewhat

counter to their current business practices, and therefore at the time of the contract we had the discussion that I talked about before regarding, you know, NREs versus royalties and their desire to ship the royalty route.

As such, we'd spent some time talking about backwards compatibility at that point saying it was really important to us that when vendors sign licenses for direct RDRAM that that license implied future devices that were backwards compatible.

So for example, we had 800 megahertz RDRAM. When we did 1066 megahertz RDRAM, they shouldn't be forced to go through the whole negotiation process again. They should just be able to use the contract they had signed, pay the royalties per the contract they'd signed and keep going.

Q. When you say "they," who are you referring to?

A. The DRAM vendors.

Q. So was there a fear that the DRAM vendors as they improved the speed of the DRAM would be forced to pay higher royalty fees to Rambus?

A. We didn't know. We just wanted to make sure that there wasn't a restriction or a bottleneck in the industry to going to higher speeds or better parts that were backwards compatible because of the negotiation

cycle.

Q. Now, there's a sentence here below the "Who makes call on backward compatibility" that says: "Intel's position is that we should make the call because if we make it a joint decision then Rambus will stall the whole process. Other concern is that Rambus will use this as an opportunity to potentially gouge more" and there's two dollar signs.

Do those two dollar signs mean money there?

A. I'm sorry. I missed that.

Q. The two dollar signs means money there?

A. Yes.

Q. And was this the issue you were just referring to?

A. Yes.

MR. DAVIS: I move to admit CX-2536.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

MR. STONE: I would note that I think the second page of the document is probably missing, but I don't think it's pertinent.

The second page of 2536 shows up when you look at the second page of 2535, you see that the e-mail was cut off --

JUDGE McGUIRE: Noted. Thank you, Mr. Stone.

(CX Exhibit Number 2536 was admitted into evidence.)

MR. DAVIS: I won't refer to that e-mail.

BY MR. DAVIS:

Q. Now, on CX-2535 -- do you have that?

A. Yes.

Q. That's also a series of e-mails between yourself and Abid Ahmad?

A. Yes.

Q. And there actually seems to be nested e-mails there?

A. Yes. Actually it's a string of e-mails I think.

Q. Okay. Well, let's then talk about the e-mail that's at the very top of the page from you to Mr. Ahmad. Do you see that?

A. Yes.

Q. In your first paragraph you say, "Does Dave really believe that it's Intel that solely determines backward compatibility?"

What was your understanding of what -- is that Dave Mooring?

A. Yes.

Q. What was your understanding of what Dave Mooring was saying about backwards compatibility?

A. I'd have to take a look there to recall.

(Pause in the proceedings.)

Ask me the question again.

Q. What was your understanding of what Dave Mooring's -- you state here in the first sentence, "Does Dave really believe that it is Intel that solely determines backward compatibility?"

What was your understanding that Dave was saying that bothered you there?

A. Well, I think he was raising the issue that Rambus had the really key role in trying to determine what backwards compatibility was or not. And he was taking the point that Intel was in their role. That wouldn't be good. And I was asking him does he really think this is the case, that we are the ones that determine backwards compatibility.

Q. And did the point of your paragraph -- what's your understanding of what determines backward compatibility?

A. Basically I think it's going to be the market. There's a lot of dynamics that go on that make the market try and drive solutions to be backwards compatible, and I think we help enable it, but we don't determine it.

I think the DRAM vendors prefer to have

backwards-compatible solutions so they can ship one device at multiple speeds. The OEMs prefer to have backwards-compatible solutions so they can have one device in inventory to satisfy multiple product forms. And you know, the point was that, you know, we don't determine it, they don't determine it, and it's a market-driven thing. We might help, but it's not going to happen unless they want it.

MR. DAVIS: I move to admit CX-2535.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2535 was admitted into evidence.)

BY MR. DAVIS:

Q. Now, earlier you discussed -- we were discussing Samsung and their relationship with Rambus. Do you remember that?

A. Yes.

Q. And at some point you said that Samsung itself did some enabling work with respect to the RDRAM launch?

A. Yes.

Q. What did you mean by that? What enabling work did Samsung do with respect to the RDRAM launch?

A. Well, Samsung took the basic design package

from Rambus. They did a lot of work to translate that to their technologies. They actually did multiple designs on different technologies and ported the design to their most advantageous technology.

In the process, they did a lot of work trying to tune up the circuits to work better. They also did some work in the platform area where they identified some issues with the channel that we talked about earlier and fed that back to both Intel and Rambus.

Q. I'd like to show you what's been marked for identification as CX-2537. This is an e-mail between yourself and Abid Ahmad from September of 1999?

A. Yes.

Q. This in fact is an e-mail from Abid Ahmad to Pat Gelsinger, but you're on the e-mail as well. I'm sorry.

The top e-mail is from Abid Ahmad to Pat Gelsinger, and the second item is: Samsung wants to be compensated for all the enabling work they've done. Samsung has done more work in debugging and ramping this technology than Rambus in most cases and there should be some benefit for that.

Do you recall there being an issue where Samsung wanted to be compensated for the enabling work they did on the RDRAM?

A. Yes.

Q. Do you remember how much they were asking for?

A. I don't think they were asking any specific amount. I think they were asking for some benefit probably in the form of what's talked about in item 1 of this same exhibit where they wanted to see the royalty rate they paid dropped.

Q. Oh, I see.

So in order to be compensated for the enabling work, they wanted to see their royalty rate dropped?

A. Yeah. Or they were quite open to other considerations as well, but you know, somehow they saw the situation as they put in a lot of effort and they were basically being treated like other vendors who had not put in the same amount of effort.

MR. DAVIS: I move to admit CX-2537.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2537 was admitted into evidence.)

BY MR. DAVIS:

Q. Now, I'd like to show you what's been marked for identification as CX-2540.

I'd like to point you to the second e-mail from the top that says original message from

Pete MacWilliams.

A. Yes.

Q. And it's an e-mail from you to Martin Rausch -- and is that the correct pronunciation, "Rausch"?

A. Yes.

Q. -- and Randy Bonella?

A. "Bonella."

Q. Who is Martin Rausch?

A. Martin Rausch is a manager of a group within the lab, Intel labs, that worried about system -- system technology issues. Thermal, mechanical, EMI, electromagnetic interference, interconnects.

Q. And who was Randy Bonella?

A. Randy worked for Martin, and he was specifically focused on memory technology enabling in all the ingredients that we need to like connectors and thermal solutions and EMI solutions, and so on.

Q. The first issue -- in the first sentence of the second paragraph you state, "Other issue we have is that making RDRAM work does not solve the business issues Rambus is causing in the industry."

What were you referring to there?

A. I was referring to basically the issues we already talked about, the fact that vendors were giving us feedback that there was a lot of problems in

negotiating with Rambus. There was some issues in terms of how much work Rambus was doing. In general, the vendors weren't really happy with the role Rambus was playing.

Q. And what was the effect of that on Intel's DRAM strategy?

A. I think there was two effects.

One, the Rambus technology, the RDRAM technology, was not moving forward fast enough to get the cost down to achieve the volume ramp we wanted to see.

And the second was the industry was getting a little bit upset with Rambus and the business model that they had, and some of them were more skeptical of whether they wanted to continue doing business in this way and have Rambus be a substantial part of the product line.

Q. Now, were you aware of any meetings in the fall of 1999 between Patrick Gelsinger, Geoff Tate and William Davidow regarding the RDRAM ramp?

A. I can't remember back then.

Q. Well, let me show you -- I'd like to show you what's been marked for identification as CX-2541.

I'd also like to move in CX-2540.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2540 was admitted into evidence.)

BY MR. DAVIS:

Q. Have you seen this letter before?

A. Yes.

Q. When did you see it before?

A. I believe I saw it when it was first written or sent.

Q. What was your role with respect to the materials in the letter?

A. I probably provided a lot of the input along with Abid Ahmad.

Q. If you look under item 1, where it says "Industry Acceptance of RDRAM," the last two sentences state: "Intel has, on several occasions, attempted to accelerate adoption/acceptance of Rambus technology, but on each occasion Rambus has failed to support our efforts. Two recent examples of this phenomenon are the RDRAM validation program and efforts to establish an RDRAM Implementers Forum."

First of all, what was the RDRAM Implementers Forum?

A. That was intended to be an industry group that would be -- that would include Rambus, Intel and DRAM

manufacturers to help define, you know, what would be the evolution of RDRAM technology.

Q. And what happened to the RDRAM Implementers Forum?

A. We decided it was a good idea, announced that we were going to do it, and then Rambus wouldn't agree to the terms that the vendors wanted to set, so it kind of fell apart before we got started.

Q. And then the -- we discussed the RDRAM validation program before; correct?

A. Yes. I guess the one thing we didn't discuss there is that we actually suggested this and formalized it and had them talk about it at an Intel Developer Forum, and the goal was to set them up in a position in the industry where they could indeed add value to the vendors, so we were trying to get them to a point where they would have this value-added role because we thought that was very important for their long-term success and our long-term success in ramping toward our target.

Q. So the validation program would have provided a role for Rambus in the industry?

A. Yes.

Q. And the next paragraph talks about availability. One issue that the first sentence --

the second sentence -- I'm sorry -- says, "Specific issues plaguing the DRAM suppliers are low RDRAM yields and low 400 megahertz and 356 megahertz yields."

Do you see that?

A. Yes.

Q. What is the difference between RDRAM yields and 400 megahertz yields and 356 megahertz yields?

A. Well, at the time we had two speeds actually 300 and 400 and the yields of the 400 were not very good, so we created another speed, which was 356, and basically we thought that 356/400 would be satisfactory yields to ramp our products with, satisfactory speeds to ramp our product with, but the feedback from DRAM vendors was that even at 356 the yields were not good. They needed to have more, more engineering work done on the RAC or the interface logic to get the yields up.

Q. Now, the RAC is something that was on the DRAMs?

A. There is a RAC on the DRAMs and then there's one in the controller as well.

Q. Now, the last sentence says, "We do not see any comprehensive and adequate plans from Rambus to work the DRAM industry to fix he's issues and bring up

additional DRAM suppliers who can provide the required volume for desktop platforms."

What sort of work were you expecting to see from Rambus relating to this effort?

A. We were expecting them to go to the vendors, find out what their speed-limiting paths and designs were, and then come up with ideas to fix those. They were the most knowledgeable about the interface in the world basically and we thought that, you know, they would be also the most knowledgeable to fix speed problems.

We were also hoping that, in addition to the suppliers that were building devices, they'd be willing to spend the effort on other DRAM vendors who are less capable and help get them up to speed and get the device in production as well.

Q. If you turn to the second page, item 5, it states "Robust RDRAM/RIMM and Channel Design," and then after that it talks about two major issues and in parentheses it says "single bit errors and lack of robust RDRAM/RIMM testing."

What is a single bit error?

A. A single bit error refers to a symptom we see in the system where we run memory tests for a long period of time and at some point we see a single bit in

the memory array fail. And it turns out that it was a single bit on the channel that would fail under certain conditions.

Q. Was this an issue with RDRAM or an issue with SDRAM?

A. No. It was an issue with the channel itself. The silicon, the controller and the DRAMs were okay. It was actually the layout of the channel and how we actually put the signals down the motherboards.

Q. So this is the channel bits being used for RDRAMs?

A. Yes.

Q. Now, in that same paragraph there's a sentence that says, "Our customers are rapidly losing confidence in us and in the technology, largely due to the lack of total, prioritized support from Rambus."

Was it your understanding that the customers were losing support for Intel?

A. Yeah, to some extent they were because we were saying this technology is solid and it's ready to ship and then we'd find problems.

Q. And what was the effect of that loss of confidence in Intel and Intel's business?

A. It's pretty substantial. We didn't ship our newest product on time and it didn't ship in the volume

we projected, so we had to ship older products for a longer period of time. I'm not sure what the dollar impact was, but it was pretty substantial. And at that time there was a lot of competitors who took some of our market share away, so it was -- we were selling less product at a lower price.

MR. DAVIS: I move to admit CX-2541.

MR. STONE: No objection, Your Honor.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2541 was admitted into evidence.)

BY MR. DAVIS:

Q. I'd like to show you what's been marked for identification as CX-2559.

If you look on page 3 of 2559, there's an e-mail from you to Louis Burns.

A. Yes.

Q. Do you see that?

And it's dated June 28, 2000.

A. Yes.

Q. The subject of this is that Rambus and Hitachi settle legal dispute.

Do you remember this discussion?

A. Yes.

Q. The last two sentences of the e-mail are: "NEC

was pretty bent out of shape" -- I'll skip the parenthetical -- "since they helped Rambus develop the technology. Bottom line is that Rambus appears to have taken yet another step in poisoning the industry. The extent of this was not expected."

What did you mean by "yet another step in poisoning the industry"?

A. What was meant there was that over the last few years, you know, things we already talked about here, in many cases Rambus had done things which the DRAM industry would later come back and talk to us about as these guys are hard to do business with. The way they negotiated contracts, the way that they provided service, it was just not a very positive experience for most vendors.

And in this case, because they were going to the extent of taking legal action against some of the same vendors, it was very concerning to us that these vendors might not even want to do business with Rambus in the future.

The position we'd hoped to achieve for Rambus of this key value-added part to the industry where people wanted them to be a part of the loop, they saw value in having them design and do validation of the parts, wasn't going to occur.

Q. Now, the last sentence says, "The extent of this was not expected."

What did you mean by that?

A. Up to this point they'd been difficult, but I think actually one of the items up -- up further maybe is a better illustration of this.

NEC was a Rambus performer from day one. They built the base Rambus part. They sold them in volume for the Nintendo game machine. They had a pretty good yield going. They had a pretty good working relationship going compared to most.

They merged with Hitachi and formed this new company Elpida. Because of the way they merged it, the Rambus contract with NEC was no longer valid. They wanted to negotiate a new contract.

MR. STONE: Your Honor, I don't think this witness has foundation as to what the contract said or what its terms were. I think he's reciting hearsay from what others told him, and my objection to this part of the answer is on the grounds of hearsay.

JUDGE McGUIRE: Sustained.

BY MR. DAVIS:

Q. Let me ask you this.

Do you recall anyone from Rambus ever telling you that Rambus owned patents or patent applications

relating to DDR SDRAM?

A. I'm sorry. I missed the first part of the question.

Q. Sure.

Do you recall anybody from Rambus ever telling you that Rambus owned patents or patent applications relating to DDR SDRAM?

A. Yeah. Dave Mooring either directly or through Dennis Lenehan had mentioned that they thought they might have some patent applications that related to DDR technology.

Q. Do you remember about when this was?

A. It was early '98, late '97. It was sometime back in the middle of the development.

Q. And what were you able to do with this information?

A. Not much. Because there was no specifics. We were related this notion that they might have something, but we were given nothing concrete. We had no specifics on what they had patents applications on. We had no patent applications in-house to look at. So it was kind of a piece of information with no immediate relevance. Coupled with the fact that we weren't even interested in doing DDR at the time, we didn't think much of it.

Q. But at some point after that, Intel started becoming interested in using DDR?

A. Yes.

Q. And -- but that information that you received from Rambus through Mooring or Lenehan wasn't useful for you?

A. Actually the first part still stands. We had nothing concrete to work with. And basically why we were starting to work on DDR was because our customers said they wanted DDR.

So we started the technical work on figuring out how to make DDR work, how to make server chipsets that supported DDR work, and we let the legal people worry about the other stuff.

MR. DAVIS: I move to admit CX-2559.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2559 was admitted into evidence.)

BY MR. DAVIS:

Q. I'd like to change the topic a little bit and ask you, do you know what PC100 is?

A. Yes.

Q. Would you describe what PC100 is.

A. PC100 was the effort by Intel to try and make

the 100 megahertz SDRAM that met a spec referred to as PC100.

Q. I'd like to show you what's been -- a document that's been marked for identification as RX-2103-14.

Can you identify what this is?

A. Yes. This is the PC SDRAM specification. This is --

Q. I'm sorry.

Did you have any role in developing the specification?

A. Yes.

Q. And what was your role in the development of the specification?

A. I was heavily involved in the memory-enabling activities at the time. The specification was actually written by an engineer who worked for Dennis Lenehan in the enabling group, but Dennis and I were kind of partners in trying to make this thing happen, he from the materials vendor relationship side and I was more focused on the Intel product lead side.

Q. And when was this specification put together or the first version of this specification together?

A. The very first version was probably put together through '95. My guess is it was probably published in the first half of or mid-'96.

Q. And this is PC100. Is there also something known as PC66?

A. Yes. In hindsight we created such a thing, but it was really PC100. In fact, the original draft of this specification was PC100 and over time it evolved to be a series of speeds, which included PC100, PC66 and PC133.

Q. I see.

And what was the purpose of this document, of RX-2103-14?

A. Well, at the time we started working on a 100 megahertz SDRAM, the environment was such that there was no industry spec to work from. JEDEC had worked very extensively on standardizing SDRAM, but they'd done so with a series of ballots. They had no comprehensive spec. You had to go back and search through the list of ballots to figure out what was in the standard and what was not.

Each vendor had their own data sheets for their parts, and when we looked at those, they didn't all match.

Q. What was the importance of those vendors' specs not matching?

A. The fact they didn't match meant they were going to build incompatible parts.

Q. But were these parts SDRAM-compliant? Were they compliant with the SDRAM standards?

A. They could all say they were compliant because they followed the JEDEC ballots, but the fact there was always an ongoing stream of ballots made it somewhat problematic.

And most of the differences were very subtle. The basic devices were all the same, but there were a few timing differences. There were basically two functions in the devices that were a little bit different between vendors, and so it wasn't an issue of being able to make the devices work, but if you're trying to do the device in high-volume manufacturing, you want all vendors to supply a compatible device. We thought there would be problems.

Q. If you would turn to page 15 of RX-2103-14, and focus on table 2.

Do you see that?

A. I'm on a different page.

Q. Oh, I'm sorry. I am focusing on the exhibit page at the very bottom of the document. It's also page --

A. I'm look at your thing. Okay. I see it.

Q. What does that table represent?

A. The table represents some bits in the mode

register and specifically to set the CAS latency.

Q. Is this sometimes referred to as programmable CAS latency?

A. Yes.

Q. Why is programmable CAS latency specified as part of the PC100 specification?

A. Because it was already there in the SDRAM parts. That was a previous feature that was balloted in JEDEC and approved.

Q. Other than the fact that it was already in the parts and it was already balloted by the group, did you have any preference for including that particular feature to set CAS latency over other ways of setting CAS latency?

A. No. It seemed like a good way to do it, they had already done it that way, the parts are already implemented that way, and so we just stuck with it.

Q. Now I'd like you to focus on table 4 of the same page. And that page is page 7 of the document and page 15 on the exhibit.

A. Okay.

Q. What does table 4 represent?

A. It represents the burst length, the same mode register, logic program, different burst lengths with the device.

Q. And is this sometimes referred to as programmable burst length?

A. Yes.

Q. Why is programmable burst length specified as part of this specification?

A. Again, because it was already something we did in the devices so we just carried it forward, and in a lot of applications, having different burst lengths is useful to kind of fine-tune it, and because people already used it, because the devices were already implemented, we didn't want to change anything.

Q. Other than the fact that the devices were already using it, did you have any preference for including the use of the mode register this way to set the burst length other than over other ways of setting the burst length?

A. No. Because it was already done, we just took it this way and used it.

Q. I'd like to show you a document that's been marked for identification as CX-2560.

This is an e-mail between yourself and Louis Burns from July of 2000?

A. Yep.

Q. It's actually an e-mail stream between yourself and Mr. Burns?

A. Yes.

Q. Who is Louis Burns?

A. Louis Burns is my boss. He's the VP in charge of the desktop platforms group.

Q. Looking at your e-mail, which is at the bottom of the first page -- do you see that, your e-mail at the bottom of the first page?

A. Yes.

Q. You state: "I do not understand the comment regarding PC133 nightmare. There is nothing wrong with PC133."

What was that referring to?

A. That was referring to the fact there was a lot of issues in the initial PC133 platforms that people were scared about, and when we looked at it in more detail, PC133 at current was basically the same device we did for PC100 with a few specs tightened up. PC133 DIMMs were basically the same modules we'd already done for PC100 and they're pretty robust.

So the issue wasn't with the DIMMs, it wasn't with the devices, it was with how those platforms were actually laid out.

Q. If you look at the top, your top e-mail, your second paragraph below where you say you're not sure if your boss is wet, it says, "PC133 specs were driven

without us." And what does that mean?

A. Well, it means there was other people in the industry that proposed the spec changes from PC100 to make PC133 spec 133 specs and those were adopted and we weren't part of the loop initially.

Q. So the initial -- the PC133 spec was driven without Intel?

A. That's correct.

Q. Then you say: "However, they would have been no different if we were in the loop. We evaluated it extensively when we made the decision last summer and found no improvements that were possible without impacting DRAM yields."

Do you have an understanding of how the PC133 specification was created even though Intel was not in the loop?

A. Vaguely. I wasn't part of the loop, so I don't know for a fact.

But based on what we heard from the DRAM vendors, a couple OEMs approached them with a proposed set of specs, and they provided feedback to those OEMs on what they could actually build, and between them they negotiated a set of specs that would work in the system and that they could build.

Q. Now, a few paragraphs below that, which is the

second to last paragraph in your e-mail, you state that "For DDR we are in a similar situation. DDR component specs are done. We were involved under the server umbrella and I do not think there is a lot more we could get out of the component."

What were the DDR component specs you were referring to?

A. It would be similar to this SDRAM spec. There was a set of specifications for the DDR component -- in this case JEDEC actually wrote the spec -- similar to this spec. And we were involved to some extent with the server group.

And you know, the comment that there's not much more we could get means that we tried to push everywhere we could to get better specs on the device to make the system design easier and we -- I figured that we were at the limit of what the DRAM vendors could actually deliver without impacting their yields across and therefore not a lot of improvement.

Q. I see in the last sentence of the paragraph you state that "At the end of the day we might be able to fine-tune the spec if we have a platform work done."

And what -- is that what you're referring to when you said "fine-tune"?

A. Yes.

Q. And what was platform work?

A. It means doing the equivalent of the channel, so we do the chipset, the motherboard, DIMMs, and all the associated simulations for the electrical timings between the devices.

MR. DAVIS: I'd like to move CX-2560.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2560 was admitted into evidence.)

BY MR. DAVIS:

Q. I'd like to show you a document that's been marked for identification as RX-1761. It's a big document, but I'm only going to refer you to page 18 -- I'm sorry -- page 16 of the document.

First of all, what is the Intel Developer Forum?

A. The Intel Developer Forum is an event that Intel holds twice a year, and it's an opportunity for Intel to talk about all the upcoming technologies that we have in the works, and we invite other people to talk about complementary technologies, and it's intended for the developers of products that use Intel products.

Q. So this was a presentation that was given to customers and other people in the industry?

A. Yes.

Q. And the first presentation was made by you?

A. Yes.

Q. If you'd turn to page 16 of RX-1761. It says "DDR Program Status" and "Intel is working to deliver robust DDR platforms for all IA CPUs." And the first item under Process below that states "Documenting our issues in a spec addendum."

Do you see that?

A. Yes.

Q. What is a spec addendum?

A. It's a process we evolved to after PC100 spec. Since JEDEC started writing the specs for these new technologies and we didn't think the specs captured all the issues that we were concerned about for our system designs, we started creating spec addendums.

Spec addendums document areas of the spec that we would hope would be a JEDEC spec eventually but are conditionally. We used the spec addendum as a tool to work with the DRAM industry to get agreement on how to build the devices to those specs. Once we have agreement, you know, we take those items to JEDEC and see if they want to incorporate them in the spec.

But in parallel with them developing the spec, it's our tool for working with the industry to make sure we have a robust device.

Q. If you turn to the next page, there's a page dedicated to the DDR spec addendum.

Now, were you involved in what ended up being the DDR spec addendum?

A. Yes.

Q. And what was your involvement?

A. I basically would look at it and review it and make sure that we'd done a good job of fleshing out our issues and documenting the ones we've added.

Q. Is the DDR spec addendum -- I'm sorry.

Is DLL in the DRAM technology -- is the DLL on the DRAM something that's in the DDR spec addendum?

A. No.

Q. I'm sorry. It's not?

A. No. It's in the DDR spec.

Q. And by "the DDR spec" what are you referring to?

A. The one that JEDEC produced.

So it's a base part of the DDR and we're not proposing to add it, we're not proposing to take it away or change it, so we wouldn't include it in the addendum.

Q. And what about dual-edge clocking? Is that something that's in the DDR spec addendum?

A. No. It's in the original JEDEC DDR spec.

Q. Now, I'd like to focus actually on a different topic. We were just talking a minute ago about the DLL on the DDR SDRAM. Do you recall that a minute ago?

A. I'm sorry. Restate it.

Q. We were just talking about using the DLL on the DDR SDRAM?

A. Yes.

Q. Do you have an understanding of what the DLL on the DDR SDRAM does?

A. Yes.

Q. And what's that understanding based on?

A. That's based on what the external specs of the device are.

Q. And what does the DLL on the DDR SDRAM do?

A. Basically the DLL aligns the internal clock to the external clock DIMM which allows all the external I/O timings to be aligned to the external clock more precisely.

Q. And now, have you ever been involved in analyzing the question of whether a DLL was required for DDR to operate properly?

A. Yes. We actually took a look at that back in the -- at the time we were looking at DDR for servers. And back in the earlier days of DDR we were very concerned that the DLL was a very high-risk item for the DRAM vendors to put on the device.

Q. And so was there an analysis done by Intel of whether a DLL was necessary for DDR device?

A. A brief one, yes.

Q. And what was found in that analysis?

A. We basically concluded that for the speeds that we're looking at that, which was DDR 200, we didn't need the DLL for the systems we were going to build.

You know, we had a pretty good experienced group at the time that had built high-speed caches for our Pentium II and Pentium III processors up to 400 megahertz without DLLs in the memory devices, so we thought we can do the same for system memory if this was going to be a risky item for the DRAM vendors.

Q. And how was it that you were able to use -- how were you able to accomplish -- deal with the data skew or the clock skew between the CPU and the SRAM caches?

A. Well, the key was the inclusion of strobes, so the fact that the clock didn't have a DLL meant that

there was more variability in the timings, but because we had strobes that were timed to the data and/or addresses, we could use the more tight timing constraints on the strobes to actually collect the data or send the data and we didn't need the precise timing to the clock.

Q. Are current Intel chipsets designed to deal with this kind of technology?

A. No.

Q. And how would Intel chipsets need to be changed in order to be able to do that?

A. Actually now it would be a pretty serious thing because back when we did the analysis, the feedback we quickly got from JEDEC was this was a done deal, they wanted the DLL there. The DRAM vendors did not see this as a high-risk item, so we kind of dropped it.

But if we were to go back and try to readdress this, I think what we'd have to do is add to the chipsets some FIFOs, F-I-F-O. And basically what a FIFO is, it's a series of charge elements.

So if we have to deal with more variability from the DRAM, data can come back quickly or it can come back slow, and we needed to make sure the FIFOs were big enough to take the data that comes back as

fast as possible and can store it in the chipset while at the same time be able to take the data that comes in as slow as possible and use that as a timing for the internal consumption of the data. And the difference between the fastest and slowest times dictates the size of the FIFO.

Q. FIFO, does that stand for first in first out?

A. Yes.

MR. DAVIS: Your Honor, I'm on my last topic, it's about two pages, but it's an in camera -- it uses a couple of in camera documents.

JUDGE McGUIRE: Let me ask you this, Mr. Davis. How much more time do you intend to take this morning before we go to lunch?

MR. DAVIS: I think I can be done in half an hour.

JUDGE McGUIRE: Why don't we go ahead and go into the in camera session and then we'll proceed until that point and then we'll break for lunch.

So to those of you in the audience, the court has issued an order in this proceeding such that the information that we're about to hear is confidential. For all of you then in the public and in the audience, I'll ask you to please excuse yourselves from this proceeding as it is closed to the public.

(The in camera testimony continued in Volume 26, Part 2, Pages 5069 through 5081, then resumed as follows.)

JUDGE McGUIRE: Okay. Very good. It's a quarter to one. Let's take a break for lunch and we'll reconvene here at 2:00 p.m.

MR. STONE: Thank you, Your Honor.

JUDGE McGUIRE: Hearing in recess.

(Whereupon, at 12:45 p.m., a lunch recess was taken.)

A F T E R N O O N S E S S I O N

(2:02 p.m.)

JUDGE McGUIRE: This hearing is now in order.

Mr. Stone, do you have a comment you want to make or do you just want to start your cross?

MR. STONE: No.

JUDGE McGUIRE: You seemed very anxious, so perhaps you are.

MR. STONE: Not at all.

CROSS-EXAMINATION

BY MR. STONE:

Q. Good afternoon, Mr. MacWilliams.

Earlier today Mr. Davis asked you about an implementer forum that Rambus was not agreeable to participating in?

A. That's correct.

Q. Was one of the terms of that implementer forum and one of the goals to revise the specifications for RDRAM?

A. All along, yes.

Q. And was it proposed by the DRAM manufacturers that the way those specifications would be revised was that each of the manufacturers would have a vote on revising the specifications?

A. Yes.

Q. And Rambus would not have a vote; correct?
A. I don't recall specifically, but I think so.
Q. And so Rambus was going to be compelled to revise its specs if the manufacturers voted in that direction and would not have any control themselves over changes; correct?
A. To the extent the manufacturers wanted to do something, yes, I think that's correct.
Q. And were you aware that that was the reason that Rambus was unwilling to participate in that implementer forum?
A. I don't think that was the only reason, but I was aware they were concerned about that, yes.
Q. They'd expressed that concern to you, hadn't they?
A. Yes.
Q. Could you look -- it's the first document in the stack that I set in front of you during the lunch break -- at CX-2535 that we looked at earlier.
If you'd look at the top e-mail of CX-2535, the one that starts -- well, you'll notice that the header at the top is from Abid Ahmad to you and then that's followed by original message from you to Mr. Ahmad.
Do you see that?

A. Yes.

Q. And is it correct that part of Mr. Ahmad's response to your e-mail is inserted in the text that begins right after it says "Abid" and ends right before it says "Pete"?

A. Yes.

Q. He inserted his responses in there?

A. Yes.

Q. And which of the paragraphs are his responses and which are part of your original e-mail?

A. My original e-mail was the first and third paragraph.

Q. And his is the part that looks bold?

A. And his response is the second and fourth, yes.

Q. And one of the subjects here was backward compatibility; is that right?

A. Yes.

Q. Why was that important at this point in time, September of 1999?

A. Because we envisioned that the direct RDRAM standard would last through several generations of speed upgrades and maybe a future upgrade or two and wanted to be able to move quickly with the industry to develop new versions of this basic technology and

wanted to be able to define backwards compatibility in the sense of what we defined in the original contract to incorporate most of those transitions, if not all of them, to make sure we can move quickly and we didn't have to go through another negotiation step with Rambus every step of the way.

Q. And there had been a negotiation over the definition of backward compatibility in the original agreement, hadn't there?

A. Yes.

Q. And was there an effort on your part to change that definition?

A. Not to my knowledge. I think we were just trying to enforce the same standard -- understanding that we had at that time.

Q. So you were just trying to make sure that whatever the provisions were in the contract everybody understood and followed them?

A. Yes.

Q. Okay. With regard to questions of compatibility or backward compatibility, it's correct, is it not, for example, that when you go from SDRAM to DDR SDRAM a new motherboard is required?

A. Yes.

Q. And are new sockets required?

A. Yes.

Q. Do you also need new modules?

A. Yes.

Q. And do you need new controllers?

A. Sometimes. We can actually do a controller that would support both SDRAM and DDR in the same device, so you use the same controller in both motherboards, but you required a different motherboard and connector and modules.

Q. And you can do a controller that would work with RDRAM and other products as well if you wanted?

A. In theory, although at the time we didn't believe that was possible because the RDRAM interface is quite different than the DDR interface. The drivers were a different configuration and the number of pins in the interface was different, making it very difficult.

Q. And are the number of pins different between DDR SDRAM and SDRAM?

A. They're different, but they're close. DDR adds strobes for each data line, data lane, so for a 64-bit interface it would be an incremental 16 signals.

Q. Okay.

A. So small difference.

Q. And when you go from SDRAM to DDR SDRAM, does

that also mean there will be a change in the chipset?

A. There will be a change in the chipset to support DDR, but the same chipset would support SDRAM.

Q. Okay. As you used the term "backward compatibility" then, would you consider DDR SDRAM to be backward compatible with SDRAM?

A. In this -- in the sense that we're trying to talk to here, it would be. We worked at backwards compatibility at some degrees in the way we did our product forms. We could do a common chipset that was one degree of backwards compatibility and in the strongest sense we wanted to maintain that so we could do one chipset that spans the old technology and the new.

I believe this time we were also trying to make sure that additional backwards compatibility in trying to preserve the motherboard, socket and module was also doable as long as we could.

Q. And I wasn't quite sure I heard the very beginning of your answer.

Would you consider, in the context that we're talking about backwards compatibility with respect to CX-2535, would you consider DDR SDRAM to be backward compatible with SDRAM?

A. Okay. Let me think.

In terms of this context, yes, in the sense that one chipset can support both technologies.

Q. Okay. As long as the chipset is designed with both technologies in mind?

A. Yes.

Q. And the chipsets that were designed to be used with SDRAM originally wouldn't support the DDR SDRAM?

A. No.

Q. But you could design a chipset for DDR SDRAM that would support both?

A. Yes.

Q. Okay. Was there a period of time when you felt there was a memory bottleneck such that the speeds of the CPUs were getting much faster and the memory devices were not keeping up in terms of their speed or their data transfer rates?

A. Yes.

Q. When did you first come to recognize the existence of such a bottleneck?

A. We saw the bottleneck coming about '95, at which time we decided we needed to do something to make memories speed up more quickly.

Q. And was the recognition of that bottleneck the incentive or the instigating event for Intel to begin its investigation of various memory technologies?

A. Yes.

Q. Okay. Today we looked at many letters and documents from the '98-99 time frame about concerns that were expressed, among other things, about concerns that were expressed by manufacturers to Intel; correct?

A. Yes.

Q. Did there come a time then later that you had a presentation -- I think we looked at it earlier -- in 2001 where you presented your road map after you had heard about the various issues in the '98-99 time frame?

A. We did many presentations. I'm not sure which one you're referring to.

Q. Well, take a look -- I've put in front of you a document -- it should be the next one in your stack -- which is a document marked RX-1792. It's maybe identical or much like the RX-1761 you saw earlier.

MR. DAVIS: Did we use that?

BY MR. STONE:

Q. 1761. I'm sorry. I wrote the wrong number down.

It's much like 1761, but I'm going to show you 1762 which is printed in landscape as opposed to up and down, and the reason is some of the charts come out a

little easier to read in this version.

A. Okay.

Q. But I otherwise think it's the same.

This is a document, am I correct, RX-1762, that was prepared for a developer forum sponsored by Intel?

A. Correct.

Q. And presentations were given at this forum by yourself and others from Intel?

A. At the forum, yes.

Q. And then presentations were also given by manufacturers, among others?

A. Yes.

Q. And the document we have, RX-1762, does this consist of some slides or foils that you used as part of your presentation as well as slides used by others?

A. Yes.

Q. Okay. And was it your effort at the time you gave this presentation in February of 2001 to give accurate and correct information to the people in attendance as to Intel's current plans?

A. Yes.

Q. Okay. And was it your understanding that the information the manufacturers were giving was information that they had shared previously with

Intel?

A. Yes.

Q. Okay. So at the point in time that they were giving their presentations in February of 2001, did you hear any of the manufacturers provide information to the audience that was inconsistent with anything they had been telling you in the same time frame?

A. No.

Q. Okay. Look if you would at page 5 of this document.

You'll see, under the first big bullet point, you'll see there's three smaller bullet points, and I want to draw your attention to the third one that begins "Memory pricing that limited ramp" -- in the 2000 I think -- "is coming down. Cost look competitive with alternatives."

Do you see that bullet point?

A. Yes.

Q. Was this your view in February of 2001, that the memory and pricing that had limited the ramp under the year 2000 was coming down?

A. Yes, that was.

Q. And was that based on what manufacturers had told you they expected to have happen with the prices?

A. Yes. And based on some programs that we'd

talked with Rambus about working with the DRAM vendors to try and drive the costs of the interface down.

Q. And when you say "cost look competitive with alternatives," were the alternatives you had in mind then SDRAM?

A. Actually alternatives were faster versions of SDRAM and DDR both.

Q. Okay. And how important was the price of RDRAM to its success during the time period of 1998 through 2001?

A. It was very important. The price of RDRAM needed to come down to be very close to what the volume memory technology was in order for it to transition to be the mainstream technology, and it can be higher priced at introduction, it can be higher priced for the first part of the ramp, but OEMs had to have confidence that it was going to come down in order to make a larger commitment to it.

So it became the determining factor as to how big the volume could be.

Q. And one of the factors going into the price of RDRAM was the number of manufacturers who would produce it; is that right?

A. Yes.

Q. And one of the factors also would be whether

the volume of the RDRAM being produced was sufficient to meet all of the demand?

A. Yes.

Q. And if you'd look at page 10 of Exhibit RX-1762, you'd see the second bullet point says "Memory pricing eventually determined by market."

Do you see that bullet point?

A. Yep.

Q. Was it your view in February of 2001 that the pricing of RDRAM had not yet been something that was determined by market?

A. No. It was actually determined by market. Typically what happens in a new technology, in the ramp, the market will decide that they need a premium. Because it costs more, they need a higher price in order to justify building these parts.

Longer term, they determine the price based on what the cost is, what the number of vendors that build it is and what the supply-demand situation is.

So I think it was true then and it is true in the future.

Q. So you didn't mean to suggest that it hadn't been determined by the market till then; you just meant to suggest that the market you felt in 2001 was going to continue to drive the price lower?

A. We thought so, yeah.

Q. Okay.

A. And very specifically what I want to make clear is that we had nothing to do with the pricing, you know, we couldn't. We built chipsets. We projected volumes. We give the information to the DRAM vendors and then the pricing was between them and their customers.

Q. And I don't mean by my questions to suggest otherwise, so let me ask you one more question to make sure we're clear on that.

From your perspective, the price at which RDRAM sold was determined by those companies who manufactured and sold it, not by Intel?

A. That's correct.

Q. Let me ask you then to turn if you would to page 24, still on RX-1762.

Does this page summarize Intel's views regarding RDRAM as of February of 2001 where it states that RDRAM provides the best Pentium 4 processor platform now and in the future?

A. Yes.

Q. Okay. And was it also perceived in this time frame that SDRAM was the lowest-cost memory?

A. Yes.

Q. And the price of SDRAM just like the price of RDRAM was set not by Intel but by the memory manufacturers who produced it?

A. And the OEMs that bought it, so it's a market supply-demand, yes.

Q. And then I want you to turn if you would -- you'll see there is an Elpida presentation which follows yours, and I'd like if you would to turn to page 53 of that Elpida presentation, still RX-1762.

Now, you mentioned earlier that there had been some concerns expressed by NEC with respect to contract negotiations with Rambus?

A. Yes.

Q. And was -- did NEC ultimately merge its DRAM business into a company that came to be known as Elpida?

A. Yes.

Q. And so by February of 2001, was Elpida a manufacturer of RDRAM?

A. I'm not exactly sure what the legal relationship was. Elpida existed, but I'm not sure if it was Elpida or NEC who was building the RDRAM officially.

Q. Okay. And I may have jumped to the Samsung presentation in any event, so let me go back for a

minute.

I wanted to go to page -- I took you to the wrong page. I want to take you to page 39 if I could and keep you in the Elpida one, so I apologize for that.

In any event, was Elpida a presenter at the February 2001 developers conference?

A. Yes.

Q. And did they present their plans and projections for their production of RDRAM?

A. Yes.

Q. And if I direct your attention to page 39 of RX-1762, do we see that they were making particular projections here about the die overhead?

A. Yes.

Q. And can you tell us what was meant by "die overhead"?

A. That was the amount of die size on top of what the standard SDRAM part that was in volume production would be.

Q. And were they showing on this chart that the die overhead for the 256-megabit design was about 15 percent over SDRAM?

A. That's correct.

Q. And then they have a box on this chart which

says "4i"?

A. Yes.

Q. And you can -- can you tell us what that stands for?

A. That's a different architecture. The initial RDRAMs as we talked about this morning had a 16d bank structure and the 4i bank structure was believed to be a lower-cost structure because of the redundancy overhead we talked about this morning.

Q. So in 2001, February of 2001, was Elpida projecting that with the 4i architecture the die overhead compared to SDRAM would drop from it looks like 8 percent to 2 or 3 percent?

A. Yeah. We've got to be careful with a little bit of this in that the dropping die was not just 4i. There are other factors that make the die size go down.

Q. Okay. And I didn't mean to suggest that was the only one.

But with that architecture and other changes that they were going to introduce, they expected the die overhead to drop to --

A. Below 10 percent.

Q. Below 10 percent?

A. Yeah.

Q. And then look if you would at page 42 of RX-1762.

Does this reflect cost projections presented by Elpida in February of 2001?

A. Yes.

Q. And what are these cost projections compared to? In other words, is this price or is this cost of manufacture or is this something else?

A. These were cost projections, assumed to be cost of manufacturing, including dies, package and test.

Q. And the cost overhead again, is that compared to SDRAM?

A. I believe so, yes.

Q. And does this chart reflect that at least in February of 2001 Elpida felt that the cost of manufacturing, including package, die and testing, for the RDRAM was going to be less than for the DDR?

A. Yes.

Q. And then let me ask you if you would turn to the next page.

And does page 43 of RX-1762 reflect in summary form what Elpida was telling the attendees at the Intel Developer Forum in February of 2001 with respect to their plans and projections for RDRAM?

A. Yes.

Q. Now, we also talked a bit about Samsung today, and Samsung was also a presenter at this forum, weren't they?

A. Yes.

Q. And their presentation I believe begins on page 44. That appears to be the first page. And I want to direct you back to page 53 if I could, still in RX-1762.

In February of 2001, did Samsung at the Intel Developers Forum present to the audience that they were achieving yields for RDRAM that were approaching the yields that they were achieving for SDRAM?

A. I'm not sure if they said those exact words. They said the yields for PC800 were good.

Q. And you'll see on this chart it says under the big bullet Improving Cost Structure it says "higher yields" and then it says, colon, "approaching SDRAM"?

A. That's what it said then.

Q. And when we use the word "yields" in this context, does that mean the number of good products that comes out of any particular batch?

A. It can actually mean two different things. One would be the number of good products that comes out of a wafer and the other would be the number of die that

are good as a percentage of the total number of die in the wafer. It's not clear from this statement which they were talking about.

Q. And both are sometimes used in --

A. Yes.

Q. -- within the framework of what we mean by "yield"?

A. Yes.

Q. Okay. And on this same subject, the Samsung presentation, if you would, turn to page 60.

And at least as of February of 2001, was it correct that what Samsung was telling the attendees at this Intel developer conference that the cost that they were projecting for RDRAM was going to be lower than the cost they were projecting for DDR at least until the end of 2002?

A. Yes.

Q. Okay. And then if you would turn to page 63, based on what Samsung was presenting at the February 2001 Intel Developer Forum, did you understand their position to be that RDRAM was the best cost/performance solution for the Pentium 4 system as indicated on this page?

A. That was their position, yeah.

Q. And is it correct that Samsung has continued to

be a producer of RDRAM up till today?

A. Yes. And including today, they're still doing it.

Q. And do you have any understanding one way or the other as to whether that has proven to be a particularly profitable business for them?

A. I don't know firsthand.

Q. Have they ever shared any of that information with you?

A. They've told us that it's been very profitable, yes.

Q. Okay. And then there also was a presentation -- I believe if you turn to the next page, page 64, you'll see the first page of a presentation by Toshiba, and I have just a couple of questions about the Toshiba presentation as well.

First, if you'll just turn to page 70, still in RX-1762, and does this diagram reflect that Toshiba at the February 2001 developers conference was telling the audience that the die size overhead as they were manufacturing RDRAM had been reduced to under 10 percent?

A. Yep.

Q. Okay. And if you look at page 79 -- not surprising, now you'll see I'm now going to go to the

cost chart -- does this chart, as you understand it, reflect that Toshiba was showing a cost overhead due to die size at least -- it's hard for me to read this. I'm not sure -- but at least a die size overhead of a little under 10 percent and dropping to something under 5 percent by the end of the year 2002?

A. That's correct.

Q. And if you would look at the next page, page 80, does this reflect that as of February of 2001 Toshiba was telling the audience that the cost of RDRAM was less than what they then were projecting to be the cost of DDR?

A. Actually in this picture it does show the cost curve for DDR is lowest cost, the RDRAM with the next two curves and then the x32 DDR is much higher.

Q. That's another -- so we have two DDRs on this chart?

A. Yes. So this perception was that the standard system version of DDR was a little less expensive than RDRAM even in this time frame.

Q. So they had a view somewhat at odds with one of the charts we looked at earlier?

A. And I think one of the key differences in the different points of view was the point that each vendor was on their learning curve. Vendors who started RDRAM

much earlier and DDR later would have a higher cost for DDR compared to RDRAM because they hadn't built DDR long enough to get down the learning curve. And in the case of Toshiba, they had been building RDRAM for quite some time and they'd got down the learning curve -- they had been building DDR for quite some time and got down the learning curve. Some of the other vendors had not.

Q. And had Toshiba been building RDRAM as long as Samsung?

A. Yes, they had.

Q. Okay. Now, did Micron ever produce RDRAM in commercial quantities?

A. No.

Q. Did Intel make an investment in Micron in an effort to encourage their production of RDRAM?

A. Actually we made an investment in Micron and the primary reason for the investment was to make sure that the industry capacity for DRAMs didn't start contracting. We were very concerned at the time that the current pricing was not allowing vendors to invest in new capacity as needed.

While we were doing the investment, we also put some requirements in the deal to incentivize them to do more RDRAM work, but that wasn't the primary driver of

the deal.

Q. When the press release announcing that investment was made, did the press release and the other public statements indicate that the investment was part of Intel's strategy to help drive PC industry growth by accelerating the adoption of direct RDRAM?

A. I don't know.

Q. Let me show you a document from -- it's CX-2522. It's not there in front of you, and let me just bring you a copy, if I may approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. Let me show you.

Directing your attention to the first page of CX-2522, do you recognize this to be an October 16, 1998 e-mail from Mr. Calder to various people, including yourself?

A. Yes.

Q. And this was about the time that -- let me just ask you, is this consistent with your recollection that this is about the time that Intel made an investment in Micron?

A. Yes.

Q. And was the amount of that investment $500 million?

A. Yes.

Q. And look if you would at the second page of Exhibit CX-2522 under Background. And if you would, just take a look at the first paragraph there.

A. Yes.

Q. And was it, in the first paragraph at least, was it indicated that Intel at least was announcing to the public that its investment was to support the development and supply of next-generation memory products and help drive PC industry growth by accelerating the adoption of direct RDRAM?

A. That's what it says.

Q. Okay. And you don't have any reason to think that the public announcements Intel made were in any way inaccurate with respect to the purpose of the investment, do you?

A. No. I think that's a key focus of the investment as it came together. The motivation for the investment as it started was not this.

Q. Well, look if you would at page 7 of this document.

You recognize this to be a copy of a press release that was actually issued by Intel, don't you?

A. Yes.

Q. And if you'd look at the next to the last

paragraph on this page where there's statements attributed to Mr. Barrett -- he at that time was Intel's president and CEO, wasn't he?

A. Yes.

Q. And you'll notice the quote that follows the identification of Mr. Barrett in that paragraph says, "Our goal in making this equity investment is ensuring an adequate supply of memory components, particularly direct RDRAM."

A. Yes.

Q. Do you see that statement?

Now, is that a correct statement at the time the investment was actually made of its purpose?

A. I believe so.

But if you read it careful, it says "ensuring an adequate supply of memory components." That was the primary thing we started out with from day one, and then particularly the RDRAM was where our road map was going at that point in time, so we wanted to make sure that Micron was strategically aligned with us to build those technologies as well.

Q. And after the investment was made in October of '98, didn't there come a time when Intel became quite disturbed with Micron because they weren't manufacturing and producing RDRAM?

A. Yes.

Q. Okay. Do you have familiarity with whether or not Intel was required under the terms of its agreement with Rambus to primarily or exclusively use RDRAM in the Pentium 4 product?

A. There is no constraint that said we needed to exclusively use RDRAM in any product.

Q. Was there a best-efforts provision or other obligation on how hard Intel needed to work to promote RDRAM? To your knowledge.

A. To my knowledge, the contract required us to make RDRAM the primary memory for our product line for a period of time, which I believe was '99 through 2002. And "primary" in my mind means that we drive it as the high-performance memory, the showcase memory, all the things that we did.

Q. And when did Intel begin to consider using other memories in its product line?

A. In the server product line we started looking at that probably in the '99 time frame, maybe a little sooner. In the desktop product line it was probably late '99.

Q. And if you would, I think the next document on the stack that I left for you earlier is the letter we looked at earlier, CX-2541, that was sent to Mr. Tate

and Mr. Davidow?

A. 2542?

Q. I might have pulled the wrong one. I have -- yours says 2542?

A. Yes.

Q. Give me one second and I'll get you another copy.

What's the date on yours? Let me just make sure they're different.

A. November 2.

MR. STONE: Your Honor, may I approach?

JUDGE McGUIRE: Go ahead.

MR. STONE: I pulled the wrong one out of the stack earlier, so...

(Pause in the proceedings.)

BY MR. STONE:

Q. I think we have the right one now, CX-2541.

Did you testify earlier that you were one of the people who sort of gathered some of the information that was included in this letter?

A. Yes.

Q. And at the time this letter was written, Intel was contemplating the possibility of some dispute with Rambus, were they not?

A. Actually '99 I'm not sure.

Q. There resulted a conference call and a meeting that you participated in following this letter, wasn't there, at which these and other issues were discussed?

A. These issues were discussed, but I'm not -- the issue of a dispute with Rambus I'm not sure was this time frame or later.

Q. At this point in time, was Intel concerned that their efforts to begin to think about other memories that might be used with their product line might put them at odds with Rambus in regard to making the RDRAM their primary memory?

A. There was some concern that Rambus might interpret it different, yeah.

Q. And ultimately the issues set forth in this letter as the result of meetings and conference calls, if not resolved, at least went no further than what was set forth in this letter, did they?

A. I think so, yeah. At this time we were still extremely motivated to make RDRAM successful and we saw a lot of technical advantages to using that in our product line. We saw no reason we couldn't do so, assuming these issues were addressed.

Q. Okay. Let me ask you to look at the next document on your stack that you were shown earlier -- hopefully I got the number right -- CX-2559.

A. Yes.

Q. And if you would, turn to page 3.

You were asked earlier about the NEC-Hitachi joint venture that they would need -- they were told that the NEC-Hitachi joint venture would need a new license to build RDRAM. Do you see that?

A. Yes.

Q. And is that what came to be known as Elpida?

A. Yes.

Q. And Elpida ultimately took a license sometime, from the date of this e-mail, June of 2000 till February 2001 based on the Intel Developer Forum materials we just looked at; correct?

A. I'm not sure exactly how they resolved the issue.

Q. And was the source of information that you set forth in this e-mail on page 3, the one dated June 28th of 2000, was this based on information you received from Rambus or information you received from NEC or Hitachi or someone else?

A. It's based on information I received from Abid Ahmad who had visited DRAM suppliers and he visited either NEC or Hitachi or both, probably both.

Q. So it was not based on information from Rambus; correct?

A. That's correct.

Q. And did you ever have a conversation with Rambus to see how they had resolved any issue that might exist in connection with the Elpida license agreement with Rambus?

A. No, I didn't.

Q. This e-mail chain that we have here, Exhibit CX-2559, was written shortly after there was an announcement of a settlement of the Rambus and Hitachi patent dispute; correct?

A. Correct.

Q. And by that time of the Hitachi patent dispute, you were well aware, weren't you, that Rambus believed that its patents covered various features that were incorporated in products that were then being manufactured?

A. Yes.

MR. DAVIS: Objection. Calls for a legal conclusion.

JUDGE McGUIRE: I can't hear you.

Right. That's sustained.

BY MR. STONE:

Q. You were aware -- and I'm asking just for your understanding, not for a statement as to what in fact the law was -- you were aware that Rambus was

contending at this point in time that products then being manufactured, SDRAM and potentially DDR SDRAM, infringed upon patents that Rambus owned; correct?

A. That was my understanding based on these press releases, yes.

Q. Yes.

And at this point in time, June of 2000, were you aware that Intel had an observer who attended various of the Rambus board meetings?

A. Yes.

Q. And initially that observer was Pat Gelsinger, wasn't it?

A. I was only aware of Pat. I didn't know he had anyone else.

Q. And did you ever talk with him about his role in discussions with respect to whether the patent lawsuit should be filed against Hitachi?

A. No.

Q. Do you know whether Mr. Gelsinger was given an opportunity to object and did not object to the bringing of the lawsuit against Hitachi?

A. I have no idea.

Q. Did you ever talk with Mr. Gelsinger about whether he encouraged Rambus to bring litigation against some of the other DRAM manufacturers who were

infringing on Rambus patents, according to what Rambus contended?

A. No.

Q. Let me ask you to look at the next document if you would, CX-2519. And if you would, turn to page 8 of that document. This is again one that you were shown this morning I believe, isn't it?

A. Yes.

Q. Under Contract Items -- and this document is as of 1998; correct?

A. Yes.

Q. And this is either after or around the time frame that Mr. Mooring had told you, either directly or indirectly, that Rambus had patent applications which they thought might cover DDR; correct?

A. Yes.

Q. And he also told you that they had patent applications that they thought might cover SyncLink; correct?

A. Yes. Actually I'm not sure if he told me or he told Dennis Lenehan who told me. Somehow I got the word that they thought they had some coverage, but they had nothing to show us that was specific.

Q. Okay. And so at the time of these -- this document was prepared, you were aware that Rambus

thought that it might obtain patent coverage for DDR and it might obtain patent coverage for SLDRAM or the SyncLink-designed product; correct?

A. That's what they told us, yes.

Q. Okay. And here where you talk about the contract items under What Intel Wants, where it says "Intel gets rights to use IP in competing memory interface," you were using the phrase "IP" here to refer to patent, weren't you?

A. Correct.

Q. And competing memory interfaces that were thought of at the time in 1998 included DDR and SLDRAM, did they not?

A. I'm not sure actually if SLDRAM is still around. It would have been DDR or follow-on DDR, follow-on to memory technology, basically anything.

And actually to put this in more context, I don't think the fact that they had mentioned that they might have patent coverage affected us at all. The most important thing we were thinking about here was the amount of exposure we'd put Intel to if we committed to more products in our product line to support RDRAM, and before we did that, we wanted to take a look at what our contract said, how much exposure we had and whether we should try and reduce

that exposure. That was the motivation.

Q. Earlier today you said that with respect to the information you received from Mr. Mooring that one of the -- I think your words were you let the legal people deal with that issue; is that right?

A. That's correct.

Q. And you don't know or at least you're not in a position to testify today as to then what the legal people did after you left the patent issues or patent application issues to them, are you?

A. No.

Q. Okay.

A. Actually one other thing to put that in context, this was in March of '98, this document and this discussion in the contract.

Q. Yes.

A. And the whole issue of Rambus patents and the discussion I talked about this morning was in 2000 after they filed the lawsuits against Hitachi. So there's a couple years difference in time.

Q. Well, the first time you heard from Mr. Mooring was in late '97 or early '98; correct?

A. Correct.

Q. And that was not the only notification that Intel received from Rambus about Rambus patent

coverage, was it?

A. I'm not sure. That's the one I can recall.

Q. Well, let me -- do you know one way or the other whether in April of 1997 it was the view at Intel that Rambus had expressed its view that it had patent coverage or would have patent coverage over the use of DDR in any memory?

MR. DAVIS: Objection. Foundation.

MR. STONE: I'm asking for foundation, Your Honor, I think.

JUDGE McGUIRE: Mr. Stone, let me read it.

MR. STONE: Certainly.

JUDGE McGUIRE: Overruled. I think he is asking for foundation.

THE WITNESS: Could you repeat the question, please.

BY MR. STONE:

Q. Sure. Let me see if I can make it simpler and try to keep it as direct.

Were you aware whether or not in April of 1997 Intel's view was that Rambus had expressed the view that it had or would have or might have patent coverage over the use of DDR in any memory?

A. I'm not sure actually. I think it was later than '97, but I couldn't be sure. And it didn't

really matter to us that much because we were so focused on trying to get the RDRAM technology to work and at the time we weren't even thinking about using DDR, so it probably wouldn't have been something that registered.

Q. Was it ever brought to your attention that Mr. Davidow had sent an e-mail to Mr. Parker in about July of '97 in which one of the things he talked about was Rambus' hoped-for or expected patent coverage?

A. Yeah, at the time I didn't know anything about that.

Q. At the time you did not?

A. No.

Q. And at the time, July of '97, what was Jerry Parker's position?

A. He was the senior VP or executive VP in charge of the technology manufacturing group.

Q. Did that put him in -- were you in his chain of reporting or not?

A. No, I wasn't. But the memory enabling team, which Dennis Lenehan was in charge of at the time, would have been.

Q. And did you later become aware of that e-mail?

A. Yes.

Q. Okay. But that was in connection with

depositions or something?

A. That's correct.

Q. Back in 1996 at the time of the original agreement between Intel and Rambus, were you given any of the Rambus patents or patent applications to look at then?

A. Yes.

Q. And did you look at them at that time?

A. Yes.

Q. And did you form any views on -- I want to be careful because I know this is an area in which there have been privilege objections in the past sought, so let me go bit by bit and I'll try and steer clear of the areas I think have drawn these objections.

Did you -- and this is just a yes-or-no question -- form any views as to the scope of Rambus patent coverage based on your review of patents or applications in 1996? And I'm not asking you what views you formed, just whether you did or didn't.

A. Yes.

Q. And do you recall whether you looked at patents or applications?

A. Only patents. There was no access to the applications.

Q. And did you at that time, 1996, look at the

WIPO or the European patent application?

A. I don't believe so.

Q. And any views that you formed as to the scope of patent coverage that Rambus had or might in the future get, those views, my understanding is, you shared with your lawyers and consider to be privileged?

A. Yes.

Q. Okay. I think the next one in your stack, I hope, is RX-1432?

A. Yes.

Q. Okay. And is this a document that you prepared in April of 1999?

A. Yes.

Q. And does this set forth some of your concerns about Micron's performance with respect to the manufacture of RDRAM?

A. I'd just like a second to read through it.

Q. Certainly.

(Pause in the proceedings.)

A. Yes, although the primary motivation of the memo was some issues between Micron and Rambus due to the contract which Micron claimed to be an impediment to make the forward progress.

So we were trying to make sure we facilitated a

discussion between Micron and Rambus, resolve their issues, so we could get going.

Q. And was this before or after the investment of money in Micron we looked at earlier?

A. I'm not a hundred percent sure, but I believe that it would have been after. We can go back and look at the press release to make sure.

Q. And let me direct your attention to the two numbered paragraphs at the bottom.

Was it your understanding in April of 1999, based on your visits and conversations with people at Micron, that Micron was not providing any visible support to the RDRAM ramp?

A. That's correct. In fact, now that I read that, this was after the investment.

Q. And had they committed to volumes in the contract with Intel that you had felt in April of 1999 they were not going to meet?

A. I'm not sure actually. They did commit to tape-out dates and to build the device and they had missed the dates they were supposed to supply prototypes, and I think that's what I'm referring to here.

Q. And did it in April of 1999 appear to you -- and I'm directing your attention to point number 2 --

appear to you that Micron was doing what they could to stall the RDRAM ramp?

A. I'm not sure of the exact time frame, but at some point they did do that very aggressively, yes.

Q. Let me ask if you'd look at the next document I've placed in the stack in front of you, which is RX-1445.

And as with many of the documents we've seen today, this is a series of e-mails, an e-mail chain, and I just want to draw your attention if I can to the second page of Exhibit RX-1445.

And the bottom half of the page is an e-mail, is it not, from you to Mr. Ahmad and Mr. Gelsinger with copies to a few other people?

A. Yes.

Q. Is the content of this e-mail based upon information that you obtained through the course of meetings and conversations with people at Micron?

A. Let me just read it.

Q. Certainly.

(Pause in the proceedings.)

A. Actually, no. This is a summary of several issues I see which involve conversations with Micron as well as information we got back from our OEM customers, information we got back from DRAM suppliers

I believe, and information we were actually reading in the press.

Q. Okay. Let me see if I can ask you about a couple of the items then and we can just break it down.

If you'd look at the very first bullet point where you say "Running ads for PC133 now like the 'wall' and 'sheep' ads they ran for SyncLink and DDR."

Do you see that point?

A. Yes.

Q. Can you explain what that refers to?

A. They are running full-page ads in technical publications such as EETimes.

JUDGE McGUIRE: Now, let me intervene so I'm clear. Who are you talking about? When you say "they," who are you --

THE WITNESS: They -- Micron is running ads.

JUDGE McGUIRE: Okay.

THE WITNESS: In technical publications like EETimes, which were full-page. You opened the magazine up and it was across, you know, the page.

And a couple of the ads actually showed a herd of sheep looking for a leader and another one showed a wall similar to the Berlin Wall that was knocked down. And the context of the ad and the words that were on

the ad were that they were concerned that the DRAM industry is just following along and not driving their own technology for themselves, and they were trying to advocate some of these other technologies through these ads.

BY MR. STONE:

Q. And the ads were ads that were negative vis-a-vis RDRAM and Intel's expressed preference for RDRAM at that point in time; correct?

A. Not directly. But there was some implication about following a lead. It wasn't clear whether it was Intel's lead or whether it was Rambus' effect on the industry that they were after.

What they were trying to do was to galvanize the DRAM industry and the OEMs to try and keep the business model as is with the DRAM vendors pretty much defining the technologies.

Q. And if you would look at the fourth bullet point, which says: "Told a DRAM supplier that they are intentionally driving excess volume into the market to make it painful for marginal players. Want people to drop out or consolidate."

Do you see that?

A. Yes.

Q. Now, was this based on information you received

from somebody other than Micron?

A. Yes.

Q. Do you recall who it was?

A. No.

Q. And on the basis of this, were you concerned yourself in April of 1999 that Micron was trying to drive people out of the DRAM manufacturing business?

A. Actually -- yeah. We were concerned that the goal of our investment was to improve the supply of memory products for the PC industry and that the investment was actually being used to possibly constrain the supply by driving some of the smaller players out, thereby not meeting the goal of our investment.

Q. One of the documents we looked at earlier today and you testified to talked about the price differential between RDRAM and SDRAM.

Do you recall that and the price differential growing? The price differential got greater over time?

A. Yes.

Q. SDRAM got cheaper by a greater percentage?

A. Yeah, I don't recall the price graph. I remember a cost.

Q. Oh, cost. I'm sorry.

Was there a price differential between SDRAM and RDRAM in the 1999 time frame?

A. Yes.

Q. And was the price differential a factor in Intel's ability to persuade OEMs to use RDRAM?

A. Yes.

Q. Was it your impression at the time that Micron was driving down the price of the SDRAM?

A. It wasn't clear. I don't know.

Q. In this document when you said that they are intentionally driving excess volume into the market, was it your understanding that driving excess volume into the market would have the effect of lowering the price?

A. It could. Although the concern expressed here was for the overall supply and the vendors that might have to go out of business, so it wasn't specifically price-related.

Q. Let me ask you if you'd look at the next one, which is a month or so later. It's RX-1453.

Is RX-1453 a copy of an e-mail that Mr. Gelsinger sent to Mr. Barrett and copied you as well as some other individuals on?

A. Yes.

Q. And is this a document that reflects the status

of Micron's progress in producing RDRAM?

A. Give me a second to look over it.

Q. Certainly.

(Pause in the proceedings.)

A. I'm finished reading. Could you please repeat the question, please. Sorry.

Q. I will.

Does this document reflect Intel's views as to Micron's status in producing RDRAM in the May of 1999 time frame?

A. Yes. I think it was specifically focused on the concerns. If there's anything good happening, it wouldn't be in here. These are the concerns that we needed to go fix.

Q. And you'll notice in the middle section where there's the headings Technically, Marketing, Relationship and Motivation -- let me just bring those up on the screen so they're legible.

Let me ask you about under the section Marketing, the last sentence says, "Their advertising implies that the rest of the industry is blindly following the Intel road map (sheep, communism, et cetera)."

Is that a reference to the sheep ads and the Berlin Wall ads?

A. Right.

Q. Okay. So at this time, May of 1999, the understanding at Intel was that the ads were intended to suggest that the industry was blindly following Intel?

A. It says the Intel road map.

Q. Yes.

A. Which may have been Intel or may have been Rambus. It's not clear at the time.

Q. Okay. Under Motivation, you'll see that it says, "Generally, their business plan currently appears to be: Drive the most aggressive pricing on PC100 forcing other players to lose money/share."

And let me just stop right there and ask you, PC100 would be the SDRAM product; correct?

A. Correct.

Q. And you were asked earlier some questions by Mr. Davis about the Intel specifications for PC100?

A. Yes.

Q. And was it your view that those specifications that Intel developed were necessary in order to ensure interoperability of SDRAM 100 megahertz that was manufactured by different of the DRAM manufacturers?

A. Yes.

Q. Okay. The JEDEC specifications alone were not

sufficient to ensure interoperability in your view, were they?

A. Well, between -- there was no JEDEC specification. There were a collection of JEDEC ballots, and if correctly interpreted, that may have worked, but the fact that they were interpreted different by different vendors and there were slight differences between the parts, it took something to kind of collect the whole industry into one common standard. It was very small, subtle differences in most cases.

Q. To your knowledge, was the first specification ever written as such for SDRAM that was used by multiple manufacturers the one that Intel put together?

A. To my knowledge, that's true.

Q. Let me direct you back to RX-1453 if I might.

When it talks here about the most aggressive pricing, do you understand that to be referring to lowering the pricing?

A. Yeah.

Q. I mean, "aggressive" in this context means low as opposed to high? Correct?

A. Yes.

Q. And then if you would look at the very end of

that paragraph on motivation where it says "create as much turmoil to prevent RDRAM as possible," was it Intel's view based on the information that you were able to gather at the time that as of May of 1999 Micron was trying to create as much turmoil to prevent RDRAM as possible?

A. It appeared so, yes.

Q. Let me ask if you would -- I think I can skip a document and speed this along.

Look if you would at RX-1295. I don't believe this is a document that you've ever seen. You may have been shown this. I'm not sure. Parts of it you may have seen.

You'll notice on the first page of RX-1295 down at the bottom it says "Ahmad, Abid" or Abid Ahmad, if I reorder it, is down there maybe as the author. Do you see that?

A. Yes.

Q. And he was at Intel in October of 1998, was he not?

A. Yes.

Q. And were you aware that in October of 1998 Intel or Mr. Ahmad were trying to collect information from DRAM suppliers about their expected production of RDRAM?

A. Yes.

Q. And was that something that was a normal part of what Intel normally did?

A. Yes. In fact, it was not just RDRAM. We collected data on SDRAM, PC100 and RDRAM, all technologies.

Q. And the purpose of that was what?

A. The purpose was to try and give the industry some more insight into what the supply-demand mix was going to look like.

Because Intel's chipsets drove a substantial portion of the demand, we had a unique position to be able to see where the future was going to go, to which technologies and speeds, and we wanted to work with the DRAM vendors to understand what they were going to build, understand what our plans were, and then we could feed back the collection of all the data to various vendors so they could see how the balance would work.

Q. And let me just see if I understand this correctly, and tell me if I have this wrong.

If for every chipset it took one memory, let's just assume a one-to-one relationship, and Intel wanted to sell a million chipsets, you wanted to make sure that there were a million memories that would work with

that chipset available in the market or else your OEMs would find themselves with unusable chipsets; is that fair?

A. That's correct.

Q. Okay.

A. And the DRAM suppliers saw value, too, because if we were going to build a million chipsets and the aggregate of the vendors were going to build five million DRAMs, they'd want to know so they could back off, and if they were going to build a hundred thousand DRAMs, they'd want to know so that they could take advantage of the opportunity.

Q. So you were sharing with the DRAM manufacturers sort of your projection for the number of units of each particular chipset that you expected Intel to build?

A. That's correct. Actually we shared the aggregate of what the demand would be.

Q. Okay. And it was important ultimately to match up the demand with the supply of memory and consistent with the demand and supply of the chipsets; correct?

A. That's correct.

Q. Do you know Mr. Tabrizi?

A. Yes.

Q. And did you understand him in 1998 to be an employee of Hyundai?

A. Yes.

Q. Were you aware -- I want you to look at the top part of this e-mail which I don't know that you've ever seen before -- and if you bring up the first couple of paragraphs -- were you aware that there were statements made at Hyundai in the October of '98 time frame that "From HEA's perspective, we can overstate our direct Rambus production so Intel can feel we are more aggressive on our ramp-up"?

A. I never heard anything like that before.

Q. If you were given -- if Intel was given false information about a manufacturer's production plans, that would ultimately have an effect on the supply and demand for memory because you might have more chipsets than memory available; correct?

A. It could. It depends on how much of an error it was. Supply-demand is not exact, but if it made more than 10 percent or so difference, it could impact that, yes.

Q. And that could have an impact on the price of the memory that was being sold in the market; correct?

A. Yes.

Q. Let me ask you to look if you would at one

other document, which is the bottom one on your stack I think, RX-2192. And I don't believe you've seen this before, but I want to ask you whether some of the information in this document was ever communicated to you.

But let me just start at the top with the names of the persons if I can and see if you know any of these persons.

So under Present Parties.

You've already told us about Mr. Tabrizi.

Do you know any of the Micron individuals listed: Terry Lee, Kevin Ryan, Brent Keeth, Jeff Mailloux?

A. Yes.

Q. Which of those do you know?

A. Terry Lee, Kevin Ryan and Jeff Mailloux.

Q. And next to NEC it says Jeffrey Lee. Do you know him?

A. Yes.

Q. And do you know the person from MOSAID?

A. No.

Q. And how about either of the persons listed from Siemens?

A. No.

Q. Okay. And I want to draw your attention to the

second page of RX-2192 down at the bottom, sort of the fourth and third bullet points up from the bottom, if we can highlight those.

There's one that starts "According to Farhad Tabrizi" -- are you on the second page?

A. Yes.

Q. It says, "According to Farhad Tabrizi, Hyundai has given Rambus ASP projections for end of next year of two to three times of today's SDRAM prices," and then it says: "They also gave to Intel a production projection of three times their actual plans. They encourage every DRAM manufacturer to do the same in order to let Intel not generate a Rambus oversupply."

Do you see that?

A. Yes.

Q. What I want to ask you, Mr. MacWilliams, is whether any of the OEM manufacturers ever told you that they had heard that Hyundai had given Intel a production forecast that was three times too high.

A. No.

Q. Did any of the DRAM manufacturers ever tell you that they'd been encouraged also to inflate their production forecasts for Intel?

A. No.

Q. If in fact there had been inflation on the

order of one or two or three times actual production of what people were forecasting, would that in your view lead to a market imbalance in the pricing because of a market imbalance in supply and demand for RDRAM?

MR. DAVIS: Objection, Your Honor. Hypothetical. Calls for speculation.

JUDGE McGUIRE: Overruled. I think the question is probative of the issue and I'll hear it.

THE WITNESS: Yeah, if there was a major imbalance, it would affect either the price and/or the amount of our chipsets that we could ship.

BY MR. STONE:

Q. Switching subjects slightly, Mr. MacWilliams, was it your -- you're familiar with an entity called ADT?

A. Yes.

Q. And was there a period of time when Intel was active in ADT?

A. Yes.

Q. Was one of the purposes of ADT to try to find technologies that would avoid Rambus patents?

A. I don't know if that was the purpose, but I think there were some discussions amongst DRAM vendors to that end, yes.

Q. Because they were aware that at least Rambus

claimed that certain of the features in SDRAM and DDR SDRAM infringed on Rambus patents, and one of the things they talked about at least at ADT meetings was to try to develop designs that would avoid those patents?

MR. DAVIS: Objection. Foundation.

MR. STONE: Let me back up.

JUDGE McGUIRE: Sustained.

BY MR. STONE:

Q. Was Intel involved in ADT?

A. Yes.

Q. And in the course of your duties at Intel, did you become aware of the extent of Intel's involvement in ADT?

A. Yes.

Q. Tell us if you would --

A. Let's just say I was aware of certain involvements but not necessarily all involvements.

Q. Tell us if you would what you were aware of with respect to Intel's involvement with ADT.

A. Well, the main focus from my point of view was we were working with several DRAM vendors to implement the enabling model we talked about this morning. We wanted to understand what their technology was capable of. We wanted to be more specific about what our needs

were in the future and try and come up with some options of how we can move DRAM technology forward.

We had worked with individual DRAM vendors and they'd asked that we try and form a group to allow us to talk amongst ourselves, so that was the primary reason and most of my involvement was to deal with that.

Q. And Intel was a member of ADT for a while?

A. Yes.

Q. And based on your involvement with ADT, was it your understanding that one of the things ADT was trying to do was develop or design products that would avoid what they understood might be the scope of Rambus' patent coverage of certain features in DDR SDRAM and SDRAM?

A. Yeah, my understanding was there was a group looking at what the implications of that might be and was not attached to the technical people that I had discussion with, so I don't know exactly where they went or what they did.

Q. And did a product -- was a product ever developed out of the ADT efforts?

A. It was not intended to develop a product out of the ADT efforts. It was a technology effort and a lot of the work that was done there was taken to JEDEC and

became part of the DDR-II standard.

Q. And was it your understanding some of that work was folded into the work that led to DDR-II?

A. Yeah. That's correct.

Q. Earlier today when you were being examined by Mr. Davis you I think used the phrase that the prices were prohibitively high for RDRAM and is it -- let me ask you whether there was a point in time when you did think that the prices for RDRAM were prohibitively high.

A. Yes. Or let me clarify that. This is in the context of making a technology that would ramp from top into the volume segments. They were prohibitively high.

Q. And the prices, just so we're clear, the prices at which the RDRAM were sold were not set by Intel?

A. No.

Q. And you had no -- Intel had no control over those prices; correct?

A. No. Other than for the memories we bought ourselves, which we negotiated.

Q. And one of the things you did do that had an effect on those prices was you did perform some function of trying to balance the number of chipsets

being manufactured with what you hoped would be the number of RDRAM products that were being manufactured; correct?

A. Yes.

Q. And to some extent in that regard you relied on the manufacturers to provide you with accurate information about their production plans?

A. Yes.

MR. STONE: I have no further questions, Your Honor. Thank you.

JUDGE McGUIRE: All right. Thank you.

Complaint counsel, redirect?

MR. DAVIS: Could I have one minute, please.

JUDGE McGUIRE: Sure. Go ahead, Mr. Davis.

(Pause in the proceedings.)

MR. STONE: Your Honor, I have two exhibits, while they're conferring if I could just ask --

JUDGE McGUIRE: Why don't you give him just a second. Then when he's set, you can do that.

JUDGE McGUIRE: Mr. Davis, are you ready?

MR. DAVIS: No questions, Your Honor.

JUDGE McGUIRE: He wants to enter in a couple of exhibits.

MR. STONE: RX-1762, which is the press release regarding the Micron investment, and CX-2522.

MR. DAVIS: I have no objection to the press release and I have no objection to CX-2522.

JUDGE McGUIRE: So entered.

(RX Exhibit Number 1762 was admitted into evidence.)

(CX Exhibit Number 2522 was admitted into evidence.)

JUDGE McGUIRE: And complaint counsel has no further inquiry of this witness, Mr. Davis?

MR. DAVIS: No, Your Honor.

JUDGE McGUIRE: All right, sir. Thank you very much for your testimony. You're excused from these proceedings, and thank you.

You can call your next witness, but before we do that, why don't we take just a short break, five or ten minutes.

(Recess)

JUDGE McGUIRE: At this time complaint counsel may call its next witness.

MR. PERRY: Your Honor, before we start, if I could, one housekeeping matter.

JUDGE McGUIRE: Go ahead.

MR. PERRY: There was one document that I used with Mr. Crisp a long time ago that I'd like to move into evidence now. It's RX-695.

MR. OLIVER: No objection.

JUDGE McGUIRE: So entered.

Thank you, Mr. Oliver.

(RX Exhibit Number 695 was admitted into evidence.)

MR. OLIVER: Your Honor, complaint counsel calls Mr. Mark Kellogg.

JUDGE McGUIRE: All right. Mr. Kellogg, would you please approach the bench and then you'll be sworn in by the court reporter.

- - - - -

Whereupon --

MARK WILLIAM KELLOGG

a witness, called for examination, having been first duly sworn, was examined and testified as follows:

JUDGE McGUIRE: Have a seat if you would, Mr. Kellogg.

DIRECT EXAMINATION

BY MR. OLIVER:

Q. Good afternoon, Mr. Kellogg.

A. Good afternoon.

Q. Could you please state your full name for the record.

A. My name is Mark William Kellogg.

Q. Mr. Kellogg, are you currently employed?

A. Yes, I am.
Q. With whom are you employed?
A. IBM Corporation.
Q. What is your position with IBM?
A. My position title is that of distinguished engineer.
Q. What is a distinguished engineer?
A. A distinguished engineer is an engineer that IBM has identified as an IBM executive and has broad responsibilities in regard to IBM product definition and industry activities.
Q. Approximately how many distinguished engineers are there at IBM?
A. My recollection would be two to three hundred.
Q. How many employees does IBM have, approximately, with engineering or technical background?
A. I believe it's in the range of 120,000.
Q. Would it be fair to say then that distinguished engineer puts you the somewhere in the top 1 percent of the engineers at IBM?
A. Certainly. Yes.
Q. Could you describe your responsibilities at IBM, please.
A. My responsibilities are associated with the

definition of memory technology that would be used in IBM systems to influence the development of technology for IBM system memory usage in the industry and to help define the entire memory subsystem itself in those systems.

Q. Mr. Kellogg, what is your educational background?

A. My formal education consists of a two-year technical degree in engineering technology, a broad range of technical, business, liberal arts courses, and of course extensive IBM education.

Q. When did you start work at IBM?

A. Just after my 19th birthday.

Q. Wow. And how long have you been at IBM?

A. 29 years.

Q. And can you briefly describe the different positions that you've held at IBM?

A. I was hired in as an engineering technician, working in a laboratory doing card design, characterization, predominantly logic-related.

After two years activity in the logic card development, I relocated from Fishkill, New York to Burlington, Vermont and joined the memory card development organization. And for the last 27 years I've been in the memory card development organization,

initially doing memory card design through the '70s, mid-'80s for low-end, midrange and high-end systems.

I became a manager at IBM in the mid-'80s, managing a card development and memory applications team. Two of my employees were JEDEC representatives from IBM to or -- to JEDEC. Sorry. For clarification.

In the late '80s, I was technical staff to a director in the subsystems area.

In the '90s to present I've been doing memory card development as a lead engineer.

Q. Now, at some point did you become a senior engineer?

A. Yes, sir, I did.

Q. And when was that? Approximately.

A. Yeah, not having worked that back recently, I would say that was in the early to mid '90s.

Q. Now, once you became a senior engineer and continuing to the present, have your responsibilities changed at all?

A. The responsibilities of a senior engineer, senior technical staff, distinguished engineer tend to broaden responsibilities to include far more than just the product design and definition. In my case, it led me to working to a much greater extent with end

customers, doing product architecture, participating in industry standard bodies, things of that nature.

Q. And can you please explain your understanding of why IBM is concerned with memory?

A. IBM is very, very concerned with memory. To put things into perspective, IBM in the early 2000 was purchasing somewhere in the range of 200 -- or no -- excuse me -- 300 to 400 million memory devices a year, so first of all, we certainly were purchasing great volumes of memory.

The memory itself is part of what we call the central electronics complex or CEC -- and I think later on I'll be describing a little more about that -- but it's so closely coupled to the processor and the overall system performance that it's a critical element in our systems; therefore, it's critical to IBM.

Q. Now, can you please now give a brief indication of the range of IBM products that use memory?

A. Most of IBM's systems utilize memory in some form.

In the case of DRAM, we would use DRAM memory in systems ranging from something as small as a printer through clearly any personal computers,

midrange systems, high-end computers. We would use them in our disk drives. We would tend to use them in retail terminals. Just about, again, all of the hardware we're producing would use memory in some form.

Q. During the course of our proceedings, we've heard quite a bit about use of memory in personal computers, but could you give at least a very brief explanation of how memory is used in printers.

A. In printers, the memory would be used in one of several forms. A most common form would be a printer buffer such that a document could be transferred to a printer for later printing or to allow the system to move off and work on other tasks.

So a printer buffer would probably be the most common.

Other cases that use memory sometimes are font buffers, for example, identifying if the printer should print the document in Japanese, English or some other language, things of that nature.

Q. And would use -- would printers use standard JEDEC-compliant SDRAMs or DDR SDRAMs?

A. In general, that is the case.

Q. I believe you also mentioned disk drives; is that right?

A. Yes, I did.

Q. Again, could you give just a very brief explanation of how memory is used in disk drives.

Perhaps I should ask you to explain briefly first what a disk drive is.

A. A disk drive is another form of memory. It generally consists of a rotating medium, a disk is what we refer to it as, magnetic medium, that stores information.

The memory associated with that disk drive would tend to be a location where information would be stored on a temporary basis such that, again, the system can apply data to the device at a very high speed and the device could store the information at the speed of the disk drive, which is relatively slow as compared to the total performance of the system, but the disk drive could perform the actual storage of the information at a much slower rate while the system is off doing other tasks.

Q. And again with respect to a disk drive, would that typically use standard JEDEC-compliant SDRAMs or DDR SDRAMs?

A. Yes, it would.

Q. Now, I believe you also mentioned servers.

Can you please describe what a server is?

A. A server could be described as a computer that is optimized for running businesses. It's a system that a business is going to rely on to operate both internal and external activities.

For example, for internal activities, it may be processing payroll, it may manage mail, calendars, scheduling, other activities inside the business, printer interface, managing large DASD farms (phonetic) for retention of data. That server would generally also or a similar server would serve as a portal to the rest of the world.

For example, a server would very often support, let's say, in retail a series of stores and manage the inventory control. It might manage your credit card processing. It may manage your actual purchases real-time. It would tend to handle invoicing. I may have mentioned inventory control and other things.

So the server is a critical element of most current businesses and the performance and reliability of the server are critical to the success of businesses in general.

Q. Can you briefly describe the different types of servers that IBM sells?

A. Now, that may be best accomplished with some

charts I brought. Would you like me to do that?

Q. Certainly.

Can we bring up the demonstrative, please.

Is this the chart you had in mind?

A. Yes, it is. I thought it might be useful just to put things in perspective a little bit.

This is a very simplistic chart and it's intended just to show scale. On the left we have what we call our xSeries product line, which is our current name convention. Our system has evolved over time.

Our xSeries describes a set of systems that are based on Intel processors or Intel-compatible processors and on operating systems consistent with those processors and they tend to be relatively low cost. Now, by "relatively low cost," they may run anywhere from four to five hundred dollars up to clearly tens of thousands of dollars.

The system just to the right, iSeries, is a series of servers that are optimized for small and midsize businesses. We've actually been producing this family of products for many, many years under different names, such as AS-400, and this family of systems is noted for performing, again, business tasks, very high reliability and for a great deal of compatibility to code that was written twenty years

ago, for example.

The pSeries servers are our UNIX-based servers. These are the systems that in addition to running UNIX code are more technologically focused or technology-focused, high-performance servers.

The servers to the far right are the zSeries. The zSeries servers, we call these our enterprise servers. These are the servers that are most likely to be found with very, very large retail customers, banks, investment firms, and other very large companies that need extreme high reliability and fault-tolerant features.

Q. Looking in particular at the zSeries, what's the typical life cycle or lifespan of a zSeries server?

A. We have a minor problem with zSeries in that quite often we can't convince our customers to stop using them, which can become very painful.

JUDGE McGUIRE: That's a real problem.

THE WITNESS: After a long time.

So the zSeries, I would expect that a customer would be using that system actively for at least five years but quite often ten and maybe beyond that. It's strictly a question of cost/performance.

But these systems tend to be what we call

full-time operation. They are always on. They are intended never to fail.

BY MR. OLIVER:

Q. You mentioned reliability I believe in a previous answer.

Why is reliability a concern with respect to servers?

A. One aspect of the server that differentiates it greatly from, let's say, the laptops that we're all getting so used to is that the laptops or desktop systems do periodically go down. It's an inconvenience. We get frustrated. We reboot. We get that blue screen of death, whatever the case may be. We're actually somewhat used to the fact that systems fail.

When a server goes down, anyone associated with that application of that system is no longer able to function in some manner. He may not be able to access the Internet. He may not be able to access his mail, which may be good or bad. He can't schedule. He can't ship his trucks. Retailers can't process credit cards.

It's unimaginable, but some applications when you scan your credit card through that terminal, that request is processed thousands of miles away and a

response comes back in less than one second confirming that purchase. It also tracks everything you just bought so that the next order is on the next truck to backfill all the purchases for that day.

That system is extremely important, and if it goes down, the business is down in some manner.

Q. Can you explain in general terms how a large server such as the zSeries would use DRAMs?

A. A system such as zSeries will predominantly use DRAM for what we call main memory.

Q. And what do you mean by that?

A. A main memory is generally the memory that's used in the system to retain programs that are currently being executed as well as information that is being processed, so it is all activity underway at a given point in time.

Q. Then I believe that you earlier used the term "central electronic complex." Can you please explain in a little more detail what that refers to?

A. May I do that via a picture?

Q. Certainly.

A. So if we can bring up the next picture.

Actually -- okay. Can I step back to this picture?

Q. I'll tell you what. Why don't I simply ask you

to explain what the second picture in your demonstrative illustrates.

A. Okay. The purpose of this picture was simply to show you the picture from the previous page for the zSeries product in a scale, because the next series of pictures won't show scale. They're two-dimensional.

So this is strictly the person we stood next to a zSeries processor to give you a feel for the size. Okay?

MR. OLIVER: Actually before we proceed any further, I think it might make sense if we were to mark these series of slides with a DX number.

JUDGE McGUIRE: Yeah. Where are we? Does anyone remember?

So let's mark the one we just discussed as DX-47 and we'll mark this one as DX-48.

(DX Exhibit Numbers 47 and 48 were marked for identification.)

BY MR. OLIVER:

Q. I'm sorry. You were saying that the slide that has been marked as DX-48 simply illustrates the scale?

A. Yes, this is simply to illustrate the scale.

Q. Perhaps we could then turn to the next slide. Would this then be DX-49, Your Honor?

JUDGE McGUIRE: Yes.

MR. OLIVER: Okay.

(DX Exhibit Number 49 was marked for identification.)

BY MR. OLIVER:

Q. And what does this slide show?

A. This picture is a cutaway of a z900. It is not the exact system from the previous page. That's actually a functional system.

This system has been cut away to permit the inside to be viewed, and what we're going to do in a series of slides is just zoom in so you can see what the contents look like, and I'll explain the central electronics complex as part of that.

Fundamentally, without dwelling a great deal, I'm going to ask that we look into the center vertically of this picture, and you're going to see some pipes running in there, and that's for cooling. We have a lot of heat being dissipated by the processor complex, so the center area.

And then if we go to the next page, we'll have a zoom-in.

Now, this is a system that's been shipping since December of 2000, and in fact we have our new system replacing that system which is starting to ship

this month called the z990, just for reference.

The z900 uses SDRAM memory. We're going to talk about that in a second. But what we have to the center and slightly to the right you're going to see an orangeish area with a number of silicon chips with the backs exposed.

I don't know if it's possible to zoom in. How do we -- there's a way to zoom in.

MR. OLIVER: If I could reflect for the record the slide we just described is DX-50 and the slide now on the screen will be DX-51.

(DX Exhibit Numbers 50 and 51 were marked for identification.)

THE WITNESS: We are now looking at the processor portion of the central electronics complex, and I'll get back to the memory in a second. This is a processor multichip module, and I'm going to in the next slide give you a little description of what it's composed of. This is actual hardware, real hardware, cut away.

To just describe what you're seeing, facing towards you is the back sides of a series of silicon chips. These chips are all designed by IBM.

The area around and kind of protruding out towards you is part of the heat sink and cooling, is

the heat exchange system. That's what the pipes were hooking up into.

So if we're okay with that chart, then I would propose we go to the next page.

MR. OLIVER: And the next one will be DX-52.

(DX Exhibit Number 52 was marked for identification.)

THE WITNESS: Now, what we've done with this chart is and what we're leaning to is to kind of show you how memory interfaces to processors, things like that.

This is a picture drawn of the multichip module itself with each of the chips identified. And I will just very quickly describe the types of chips on this multichip module complex and describe a few attributes of the complex itself.

So if you look to the right, you see some text. There are 35 chips on this carrier. The package itself, just for relative size, is about five inches by five inches.

It has about two and a half billion transistors, and those two and a half billion transistors are logic devices, memory devices, things like that. They all pretty much work, which is cool. If you're an engineer, it's great when you actually get

one of these to power up and show you a log-on screen or something.

But it has 20 processors, so the PUs are actually processors.

Now, of the 20, we generally use about, well, maybe 8, 10, 12, 14, 16, depending on the customer. The customer can actually use more or less of the processors depending on what they purchase, what they would like to buy at any given time.

It actually has spare processors, too. If one fails, we can actually switch a workload over to another processor to keep operation again. Again, fault tolerance is key here. The system cannot go down.

So the processors each have about 47 million transistors.

Beneath that I describe we have eight cache chips. Those are actually SRAM caches. Those are the SD devices in gray. Those are local storage. There's actually storage on the processor for very fast execution.

The next level of storage is sitting next to the processor called the cache, and that cache stores local information and ultimately will connect back to the main memory, which we'll get to in a minute.

The two storage control chips right in the center are the chips that interface between the processor, the local cache and ultimately off to the main memory.

The four bus adapters in the corner are actually to communicate to I/O devices just in case you want to print something or store something to an external device.

And there's a clock chip down at the bottom center which distributes clocks or timing signals to all the chips on the board, on the card.

The package it's attached to -- and we have one, don't we?

Q. Yes. Actually, let me bring it up.

May I approach, Your Honor?

JUDGE McGUIRE: Yes.

THE WITNESS: I forgot we brought this.

MR. OLIVER: Your Honor, may I be permitted to approach as well as opposing counsel to point?

JUDGE McGUIRE: Yes. We'll look at that up at the bench.

Is this then what DX- --

MR. OLIVER: Actually they've requested that they have this back at the end, so this would just be for illustrative purposes.

JUDGE McGUIRE: All right.

THE WITNESS: I'm clearly folding the package in the direction as the image on the screen, so just for reference, for example, here's the clock chip (indicating).

The purpose of showing you this is this particular assembly not only has some several billion transistors, the package that it's attached to, since it has to attach all these devices together, has a hundred layers of wiring.

Now, those hundred layers of wiring are actually on a ceramic substrate, so we actually have a hundred layers of ceramic, each layer having wiring on that.

There are also several thin film layers in the very surface which are a higher-speed communication channel for transferring data at very high speeds between the processors and the caches.

So -- and on the back side you can see a series of pins, many of which are bent since this is actually handled by people, so if you'd actually like to touch it, you're welcome to.

JUDGE McGUIRE: Now, that's not like the one we saw the other day that costs eight or nine hundred thousand, is it? Because I was afraid to touch that

one.

THE WITNESS: This does cost more, yes. Anyone familiar with the technology.

So this is the central electronics complex processor, and sitting next to it we'll show you the memory.

JUDGE McGUIRE: Thank you.

THE WITNESS: So if there's no questions, the next page?

MR. OLIVER: And this will be DX- -- is it 52?

JUDGE McGUIRE: 52.

(DX Exhibit Number 53 was marked for identification.)

THE WITNESS: Now, the central electronics complex is, again, the central region of this mainframe, this large computer.

We're looking just now to the left of the processor and we're looking at the memory card. We call that a card. It's also in what we call a book package. The memory card is slightly pulled back, and you can see there's a connector.

Just to the left of the processor there's a bunch of gold pins. Those are the pins into which the card installs.

And we brought one of those cards as well?

MR. OLIVER: Actually, Your Honor, I believe that the demonstrative he's discussing is actually DX-53.

JUDGE McGUIRE: All right.

THE WITNESS: Now, the memory card that we have --

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

THE WITNESS: So this memory card is actually the memory card in that picture. We thought you'd want to see the real thing.

And this memory card installs just to the left of the processor. If you'd gone back several pictures, and don't, you would notice that there was one of these cards to the left of the processor, one to the right of the processor, there's one on the other side of this board which you can't see, and then there's to the left and to the right, so there's actually positions for four of these memory cards.

Now, do you want to --

MR. OLIVER: Would it make sense for Mr. Kellogg to approach again?

JUDGE McGUIRE: Sure.

THE WITNESS: Now, I will in a subsequent chart actually show you the key elements. I'll just point it

out here so you can see it.

The memory devices themselves, there aren't that many on this card because these are fairly costly, so we didn't want to damage a very expensive one. These are the memory devices. These are actually synchronous DRAMs, not double data rate. The current system is double data rate.

And I'm going to talk a little bit about these assemblies with the heat sinks. These are heat sinks with the black. And I'm going to talk about this assembly but also to the heat sink.

This package, by the way, is a metal-encased package that has cooling that passes through it again. Okay?

So I'm going to show a picture and describe this.

BY MR. OLIVER:

Q. Mr. Kellogg, before you proceed, could I ask you to explain how many memory book packages of that sort there would be in a typical z900 server.

A. Between two and four.

That's actually on a chart I'm going to show in a few pages.

MR. OLIVER: If we could bring up the next slide, which I believe is DX-54.

(DX Exhibit Number 54 was marked for identification.)

BY MR. OLIVER:

Q. Mr. Kellogg, can you please explain what is shown in DX-54?

A. So this picture is actually a picture of the card I brought, and we've submitted a picture because we would like the opportunity to take this hardware back with us at some point in the near future.

So what we've done on this picture is pointed to the key elements, and on the next page we'll actually describe them further.

But fundamentally, if you start from the left -- and we'll go around it clockwise -- we point to an SDRAM or synchronous dynamic random access memory, which I think you're familiar with.

So that's the memory device and that's on the memory carrier, which we'll describe in a minute.

Up in the top left is a key store module. This is a carrier that's in this assembly that tells us what information is actually stored on this card, because there are many processors all storing information in this card and we have to keep track of it.

So this is where we store the keys. This is

the information telling us what memory is on it.

By the way, that key storage is triple redundant and it works, so if any memory device fails, chances are we'll continue to operate without fault. That key store uses synchronous dynamic random access memory, so it's another industry standard memory.

Directly to the right of that and in the center is the memory controller. This is the logic chip that actually passes information from this assembly back to the MCM, the processor module itself.

To the right is a FRU gate array module. Effectively what this is is this is attribute data. This is what the system goes out and reads to find out what's installed on this particular card.

Down to the right you'll see a term "memory module." The memory module is referring to -- and I'll hold this back up again -- the memory modules are named for these carriers, these vertical assemblies that have the memory devices attached (indicating). Those memory modules are simplistically very similar to what you see in a PC.

Directly below that you'll see a pointer to what we call the memory interface chip with the black heat sinks or the little pins sticking up for cooling. This is the device that actually directly communicates

to the memory device.

Now, just to put it in perspective, for the processor to talk to the memory device, the processor has to pass information from the processor chip itself to the memory controller on the MCM itself, from that chip to the memory controller on this card, from the memory controller to the memory interface chip, and from the memory interface chip to the memory device in the memory module. So the path is relatively lengthy.

The last item pointed out on there is the bulk decoupling module, which is simply a whole bunch of capacitors.

Q. Now, Mr. Kellogg, if I noted correctly, it appears that there are nine SDRAMs in a typical module. Is that right?

A. On the carrier here there are nine DRAMs, yes.

Q. And I think that compares to, you know, four or eight in a typical module for a personal computer?

A. A personal computer would typically have on a module 4, 8 or 16.

Q. Can you explain why this has nine memory devices on the module?

A. I can. The difference is this particular memory module is only what I would describe as four bytes wide or 36 bits and I actually access many of

these memory modules in parallel to get my, what I call data word.

Can we go to the last page?

Q. Yes. We can turn to the next demonstrative, please. This is DX-

JUDGE McGUIRE: 54 is it?

MR. OLIVER: 55 I believe, Your Honor.

JUDGE McGUIRE: 55. Thank you.

(DX Exhibit Number 55 was marked for identification.)

BY MR. OLIVER:

Q. And what is illustrated on DX-55?

A. So this is actually the next to the last page, and what this page describes is the fact that a memory card such as the one I brought would have 16 or 32 memory modules. It would have one memory controller, as I pointed out. It would have eight memory interface chips -- there's actually nine, the ninth one being on the key store. It would have the key store module itself with, as I mentioned, triple redundancy. It would have this FRU gate array module, which is what I described as vital product data, the bulk decoupling module, and there's a connector off the back, the VHDM or very high density module connector.

MR. OLIVER: If we could then turn to the next demonstrative, which would be DX-56.

(DX Exhibit Number 56 was marked for identification.)

BY MR. OLIVER:

Q. Could you please explain what is illustrated on DX-56.

A. This last chart is intended then to take what we've just shown you and put it in perspective in regard to the number of devices and other complexity.

We have then on the back side of this assembly -- I'll hold it up for a second -- there's a memory card. The memory card is the carrier for all these components we're discussing. That memory card has 26 layers.

The purpose of the 26 layers are to distribute power as well as signals to all components on that card.

The memory card then supports anywhere from 144 to 1,152 synchronous DRAMs. This is one of the special elements about servers. Servers tend to have a lot of memory and now they have significantly more than even what we had in 2000.

Those chips allow us to put anywhere from 2 gigabytes to 32 gigabytes of memory on that

particular card.

The memory DIMM or the memory module, those vertical cards I described before, there are 16 or 32 of those, which I mentioned in the previous slide. The memory module can have 9, 18 or 36 DRAMs, actually SDRAMs for clarification. The memory module has 12 wiring layers and it supports 128 or 256-million-bit SDRAMs, allowing us 128 to one gigabyte per DIMM.

As I mentioned before, you have two or four of these cards per z900 frame and at time of introduction in 2000 that allowed us 64 gigabytes of system storage.

Q. Mr. Kellogg, just to help summarize what we've just seen, can you please list again the elements of what you would consider to be the central electronic complex?

A. So the central electronics complex that I showed in the original slides then consists of in the center of this system an MCM which carries the processors, the local cache, the interface chips. Just to either side of that it includes the memory, the main memory which is on these memory cards. And then just to the right of that it has the primary interface to the I/O.

The rest of that whole system is technology to

support that central electronics complex. This is really the heart of that mainframe.

Q. And then within the memory book package there, would you consider the memory controller, the memory interface, the memory card and the modules with chips themselves also to be part of that central electronic complex?

A. Yes, I would.

Q. Now, of that central electronic complex in the z900, how much of that was designed by IBM?

A. All of it.

Q. Including the DRAM chips?

A. Just for clarification -- thank you -- the DRAM chips themselves are industry-standard memory devices. They're designed by a number of manufacturers and standard with JEDEC.

Q. So in other words, IBM did not design the DRAM chips?

A. IBM did not design the chips. IBM participated in the definition of those chips.

Q. And those chips again would be the JEDEC-compliant chips?

A. Yes, they are.

Q. Now, putting the chips aside, at least with respect to all the remaining elements of the central

electronic complex, how many of those did IBM design?

A. To my estimation, all circuitry on that central electronic complex was designed by IBM in regard to the silicon chips other than the memory devices. There may be one or two discrete, small devices such as E-square EPROM, but I would say that 95-98 percent of the circuitry on that central electronics complex was designed by IBM.

Q. So in other words, IBM designed almost everything except for the DRAM chips themselves?

A. That is correct.

Q. What's the typical time frame for development of an IBM server?

A. The design cycles for IBM servers are relatively long, typically in the two to three-year range, depending on the complexity of the server.

Now, I could also provide some further information on that if the court is interested.

Q. Okay. Go ahead.

JUDGE McGUIRE: Well, go ahead. I mean, I'm going to have you ask the questions, and if I need to interject, I will.

BY MR. OLIVER:

Q. Thanks, Your Honor.

Again, we'll keep it relatively brief, but if

you could provide a little more information about the time frame of development of IBM servers.

A. If I state so verbally instead of drawing that out, the -- if I was to look at a midrange or high-end server, in other words, a server that's developed using IBM technology, that server, as I indicated, would be in the range of two to three years, typically in the range of two and a half years, and that's really related to products that we have recently developed or are developing.

If I was to break that schedule into the key elements, of the key elements, about one and a half years of that initial development cycle is the actual architecture and design of all -- of the system itself as well as the semiconductors that we showed in the central electronics complex and on this memory card other than the DRAMs.

Once that design is complete, it takes about three months to get those chips manufactured, first parts returned and a system powered up.

Assuming the system powers up and diagnostics can be run and testing can be initiated, in then about six months we would be prepared to finalize what we'd call a final pass design where we make logic changes required to correct any errors in the initial design.

That design activity might take three to six months depending on the complexity of the change.

Again, the chips would be manufactured in a semiconductor facility. Over a two to three-month period, qualification testing could be initiated based on whatever parts are fully functional.

We would continue to merge devices out of the facility, semiconductor facility, to continue testing, and within about nine months we would be prepared to actually start selling the system.

If you were to add those numbers up with a little bit of overlap, you would find a schedule, again, in the range of two to three years, typically two and a half years.

Q. Now, with respect to the z900 series server, when did IBM start selling the z900?

A. This box began to be sold in December of 2000.

Q. Do you recall when IBM started development of the z900 servers?

A. It was in that particular case almost three and a half years before from an architecture standpoint and three years before from a design standpoint.

Q. So it first became -- began architectural work in, say, mid-1997?

A. Yes.

Q. Now, does IBM have plans to sell a follow-on model to the z900?

A. Yes. The follow-on model was announced in May and begins sale this month.

Q. What is that model?

A. The z990.

Q. When did IBM begin development of the z990?

A. Approximately three years prior to this date.

Q. What was the first IBM server that incorporated DDR SDRAMs?

A. I can't say that I recall the name under which it was sold in the market. I do know the date. Would that be adequate?

Q. Okay. What's the date?

A. The first introduction of double data rate by IBM in a server product was in 5 of '01. May of '01.

Q. And when did IBM begin development work on its server to incorporate the DDR SDRAMs?

A. I didn't have a great deal of involvement in that particular system, but I did in one that started to ship shortly after that.

Q. Do you have an understanding of when development work started in the system you are familiar with?

A. Right. The system I'm more familiar with is a

family of pSeries systems which began to ship in 12 of '01. That was a broad-based family which we continue to ship today. That product set, again, was on about a three-year development cycle from initial design to production.

Q. Now, at some point did IBM begin development of memory interface units for its servers that would work with both SDRAM and DDR SDRAM?

A. Yes, we did.

Q. And do you recall when IBM began to work on memory interface units that would support both SDRAM and DDR SDRAM?

A. It turns out that the memory interface chips on the memory card that I brought today, those are the eight, as it turns out, nine chips on this carrier with the black heat sinks. Those memory interface chips are designed to communicate to both SDR and DDR memory. Those chips had a design initiation in the third quarter of 1997.

Q. And in what products today does IBM use that memory interface unit?

A. This particular chip is used predominantly in the z900. We have two other versions of this chip. We have a low-cost version with a different package. We have a slightly modified version also used in another

form factor.

This chip is used in one of three form factors, in xSeries, pSeries, iSeries and zSeries systems.

Q. Mr. Kellogg, I'd like to turn now to your experiences with JEDEC.

First, let me ask generally, are you a member or a participant in any industry standard organizations?

A. Yes, I am.

Q. And what organizations?

A. I'm involved in JEDEC.

Q. Any others?

A. No others.

Q. Okay. When did you first become involved in JEDEC?

A. 1989.

Q. And do you still participate in JEDEC today?

A. Yes, I do.

Q. In what JEDEC committees have you participated?

A. I've participated in JC-42.5, JC-42.3, JC-40, JC-16, and to a limited extent in JC-11.

Q. Let's focus if we could on JC-42.3 and JC-42.5.

Do you recall when you began attending meetings of the JC-42.3 subcommittee?

A. Yes, I do. It was approximately -- it was around 9-89, possibly 12-89.

Q. And do you recall when you began attending meetings of the JC-42.5 subcommittee?

A. The same time.

Q. Now, did you attend any JEDEC committee meetings as IBM's official representative?

A. Yes, I did.

Q. Which committees or subcommittees?

A. I was IBM's official representative to JC-42.5 pretty much from the start of that committee, again, in 1989.

I also have been the official representative to JC-42.3 as the primary since later in the '90s, as an alternate since early 1990s.

Q. Who was the primary IBM representative to the JC-42.3 subcommittee in the early 1990s?

A. That was Gordon Kelley.

Q. Now, would you say that during the early to mid-1990s you attended JC-42.3 and JC-42.5 subcommittee meetings regularly?

A. Yes, I did.

Q. Did you attend JC-42.3 and JC-42.5 subcommittee meetings in the capacity of a memory manufacturer or a memory user or both or some other capacity?

A. That was the slightly special thing about my IBM participation in that I was representing in most cases both a supplier as well as a customer.

Q. And by "a customer" you mean a user of memory?

A. That is correct.

Q. Now, does JEDEC have a policy requiring disclosure of patents or patent applications?

A. Yes, they do.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document that's been marked as CX-2375 for identification.

Let me ask first, do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document is a set of minutes that I wrote based on a 3-94 JC-42.3 meeting.

Q. Now, is this a document that you wrote at the meeting itself or at some point thereafter?

A. This document was prepared at the meeting.

Q. Let me ask you to turn if you would, please, to page 2.

And there's a reference about a quarter of the

way down the page or so to IBM. Do you see that?

A. Yes, I do.

Q. The paragraph there, if I read your handwriting correctly, I believe it reads, "IBM: Not full IBM position, IBM agrees to warn of potential applicable patents, not to determine applicability. Companies must warn of potential applicability."

Do you see that?

A. Yes, I do.

Q. Do you recall a discussion of IBM position along these lines at this meeting?

A. Yes, I do.

Q. Let me ask first, do you recall who made the statement that you reflected here in these minutes?

A. Gordon Kelley.

Q. Now, can you please explain a little bit of the context in which Mr. Kelley made this statement? If you need to take a moment to look at your notes, please feel free to do so.

A. I recall this discussion as being associated with a TI quad CAS patent disclosure which I believe had occurred at this -- during this time frame, and there's a discussion about patents underway. And Gordon Kelley has taken a position -- now, at the time, Gordon had a leadership role in the committee and

Gordon was reflecting the fact that it was -- it was necessary for companies to warn of patents, and he was confirming that IBM agrees to warn of patent or patent -- potential applicable patents.

Q. Let me take a step back first and ask if you can give a brief description of your recollection of the TI dispute that arose at this meeting.

A. I'm going to try to put this in the context of time, so I'm double-checking.

So my recollection is that prior to this situation, TI had filed suit -- I think against Micron -- associated with the production of what we described as quad CAS devices, a device that had four column address select lines.

So that was kind of the background of the discussion underway.

Now, would it be possible to have the rest of the question repeated?

Q. Let me ask another question then.

I see on your notes starting on page 1 and turning over to page 2 a number of company references, a reference to TI, a reference to Sharp, a reference again to TI, a reference to IBM, a reference to VLSI, et cetera.

Do you see those?

A. Yes, I do.

Q. And it appears that your notes are reflecting some type of ongoing discussion; is that right?

A. That is correct.

Q. Can you please explain your recollection of the discussion that you were summarizing here in your notes?

A. There was some concern in the committee, again associated with this TI patent on quad CAS devices because there was JEDEC activity associated with quad CAS. TI had brought some information to the meeting which they were presenting and that led to a discussion amongst some of the people in the room.

I attributed certain comments, I'm sure far from all of the comments, but some comments that I was able to write down as they were being discussed and I attributed them to the company that was making the comment.

Q. By the way, based on your recollection, was this, you know, a fairly slow discussion or was this a rapid discussion?

In other words, I'm trying to figure out how many -- how accurately you would have summarized -- excuse me -- how many of the comments you would have captured in here.

MR. PERRY: Objection. Compound. There's five questions in there.

MR. OLIVER: I'll withdraw that. I'll withdraw the question, Your Honor.

JUDGE McGUIRE: Okay.

BY MR. OLIVER:

Q. I'm trying to get some sense of the nature of the discussion at this meeting at this time.

A. From the notes, it's clear to see that several companies were concerned about the quad CAS patent, mainly because I believe it was disclosed at a time when validating activity was underway. At least in my recollection that was the situation.

So there was some emotion, and that is reflected in my comments.

I don't recall the length of the discussion. I don't know if I captured all of the comments or a portion, so I can really only speak to what's on this page.

Q. Okay. Well, let me come back then to the specific IBM comment that you've recorded here. Again, it reads, "IBM: Not full IBM position, IBM agrees to warn of potential applicable patents."

Now, what do you recall Mr. Gordon Kelley stating with respect to that, the sentence you've

written there in your notes?

A. Gordon Kelley was stating that IBM was willing and in fact we were already disclosing patent activity, so he was simply stating that we would warn of potential applicable patents.

Now, beyond this we were disclosing applications and other preapps, so he was simply making a statement, and what I mean by "not full IBM position" is I have here an excerpt of the IBM position. I don't know if I attributed that portion or Gordon did, but he is addressing at least the patent side of IBM's position, which is we'll disclose.

Q. I believe you referred in your answer to preapps. What did you mean by "preapps"?

A. That's a term I use for disclosing plans to apply for a patent. In other words, the application hasn't yet been submitted to the patent office.

Q. Now, do you recall any other instances in which Mr. Kelley described the IBM position with respect to patent disclosure to the JEDEC 42.3 subcommittee?

A. I can't say that I recall by meeting what Gordon stated. I know that Gordon stated several times, likely many times, what IBM's patent policy was, when it was appropriate to do so. And I believe in

some ways he was pulling that up just so people understood that it was important and that IBM was following that policy.

So I recall the discussion. I can't attribute specific meetings. I'm sure my notes would show several.

Q. Without worrying about which specific meeting any statement occurred at, can you please summarize your recollection of what Mr. Kelley told the JC-42.3 committee with respect to IBM's position?

A. I believe that Gordon made it clear that IBM and in fact that any company was obligated to disclose patent activity, and sometimes we used the word "patents," sometimes "patent and patent applications," but the term I'm most familiar with was "patent activity," and Gordon was indicating that IBM would disclose patent activity that the participating members were aware of.

Q. Do you recall Mr. Kelley ever telling the JC-42.3 subcommittee anything that IBM would not do?

A. Yes.

Q. What do you recall?

A. Yes, I do. We're admittedly a big company, and memory is a small piece of this very big company, so one of the things that we felt impossible for us to do

would be to study/evaluate the entire patent portfolio from IBM.

So Gordon more than once disclosed the fact that we would not research patents, in other words, attempt to study via keywords whether or not IBM had patents associated with activities that we were not directly familiar with.

So he did state that.

Q. Did Mr. Kelley ever say anything about what IBM representatives would or would not do if there was a specific request made to them?

A. I believe Gordon stated at least once and actually followed up -- that's why I'm comfortable with this -- but Gordon stated that if we were requested to evaluate the possibility of patents and if we were able to do so, we would investigate.

And I believe the one case that I recall is associated with what we described as bulk data arrays or what IBM used to describe as C4 interconnect packages.

Q. Now, do you recall Mr. Kelley ever telling JEDEC that IBM would not comply with the JEDEC disclosure policy?

A. No, I do not.

Q. Do you ever recall Mr. Kelley telling JEDEC

that IBM representatives at the meetings would not disclose patents or patent applications of which they were aware?

A. No, I do not.

Q. Do you recall anybody from IBM ever telling a JEDEC meeting that an IBM representative at the meeting would not disclose patents or patent applications of which they're aware?

A. No, I do not.

Q. If I could ask you to look back again at CX-2375, please, page 2, and if I could direct your attention in particular to about three-quarters of the way down the page.

A. Yes.

Q. There's a statement about Rambus that reads "Rambus: Policy seems unworkable."

Do you see that?

A. Yes.

Q. I'd like to ask what your recollection is of the statement that is reflected in your notes there.

A. My recollection is I heard something like this, and my recollection doesn't go beyond that.

Q. After the date of this meeting, did you have any belief that Rambus would not follow the JEDEC disclosure policy?

A. I think the earliest that I became aware that that might be the case -- and this was indirect -- would have been upon Rambus' withdrawal from JEDEC.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kelley, I've handed you a document that's been marked as JX-19.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document is a package of what I would describe as the official JEDEC minutes for the JC-42.3 meeting on March 9, 1994.

Q. If I could ask you to turn, please, to page 4.

And on the page underneath the caption Patent Policy, if I could ask you just to read that to yourself, the remainder of page 4 and onto the first paragraph of page 5. I'd like to ask you specifically a couple of questions about the top of page 5, but I would like for you to get the context first.

(Pause in the proceedings.)

A. Okay.

Q. Now, the discussion in JX-19 beginning under the caption Patent Policy on page 4 and turning to the

top of page 5, does that appear to reflect the same events that were recorded in your notes, CX-2375, that we looked at a moment ago?

A. Yes.

Q. If I could ask you to turn, please, to page 5 in JX-19, and specifically the statement at the top of the page: "The committee was asked if the patent policy is clear. The committee felt it was clear."

Do you see that?

A. Yes, I do.

Q. Do you have any recollection of the events reflected in those two sentences?

A. Yes, I do.

Q. What is your recollection of those events?

A. My recollection is that after the discussion that was described in my notes, the discussion had led some to question whether or not there really was any confusion with the patent policy, and it was necessary to determine whether or not there was confusion with the basic patent policy, the need and obligation to disclose patent activity.

This vote was taken to see if everyone in the room agreed to the patent policy. My recollection is that this was a unanimous vote. What I don't know is if we followed the normal procedure, which is asking

for companies that disagreed to raise their hands first. That's the easiest way to ensure unanimity. However, this was a unanimous vote.

Q. Now, in terms of a vote, how did this come about?

A. Typically when a question like this is asked, and we have a history of asking questions to this group, the question is asked again in a form and generally we look for the negatives to see if anyone disagrees and a show of hands is used to determine what the vote is.

Q. So it would be a show of hands of everyone present in the room?

A. That is correct.

Q. And your recollection is that that show of hands resulted in unanimity?

A. My recollection is it was unanimous.

One of the reasons I feel comfortable with that is this was a critical question. I do remember getting discussion because the quad CAS activity affected IBM, for example. It was a serious question.

And if a company had indicated they were not familiar with it, I would have logged a comment. I tended to log comments, especially comments that I believed were important or potentially important.

Q. Now, Mr. Kellogg, have you yourself made any disclosures of patents or patent applications during your time at JEDEC?

A. Yes, I have.

Q. On how many occasions have you disclosed patents or patent applications at JEDEC?

A. I believe at least five.

Q. We won't go into all of those, but I would like to look at one at least.

May I approach, Your Honor?

JUDGE McGUIRE: Go ahead.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document marked as CX-21 for identification.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document represents the official JEDEC minutes for a JC-42.5 meeting held on September 16, 1991.

Q. Did you make a presentation at this meeting?

A. Yes, I did.

Q. If I could ask you to turn, please, to page 43 in CX-21.

A. Okay.

Q. Is this the presentation that you made?
A. Yes, it is.
Q. I see the letters MK in the lower left-hand corner. Are those your initials?
A. Yes, they are.
Q. And I see the caption at the top of JC-42.5 8 Byte SIMM Proposal. Do you see that?
A. Yes, I do.
Q. Could you just explain very briefly what an eight-byte SIMM is.
A. An eight-byte SIMM is a memory module that is capable of providing 64, 72 or 80 bits of information to a -- some other source.
Q. Now, at the time that you made this presentation in September of 1991, what, if any, was your understanding as to whether IBM had any relevant patent activity in the area?
A. This proposal was actually out of our PC division in Boca Raton and they had asked me to propose an eight-byte memory module that they were considering for use in some of their future products.

This proposal was compiled by me in an effort to describe that proposal to the JEDEC committee.

Q. At the time that you made this presentation, did you understand one way or another as to whether IBM

had any patents or patent applications that were relevant to this proposal?

A. When asked to make a showing, I asked the engineers if they had any patent activity planned, and they indicated they did.

So I was aware of patent activity; hence, that affected my showing and my disclosures within JEDEC.

Q. Now, when you refer to patent activity, what do you mean?

A. I described that before and I'll try to be very consistent here. Patent activity to me is intent to file, file, the actual filing itself or the issuance of a patent, so I use that general term.

Q. Now, as of September 1991, had IBM actually filed a patent application relevant to the eight-byte SIMM?

A. No, we had not.

Q. If I could ask you to turn, please, to page 45 of CX-21.

A. Yes.

Q. And at the top of that page under the bullet Proposal Status it reads, "IBM wishes to disclose full details of this proposal in 12-91 product definition remains volatile resolution of any potential patent issues prior to showing."

Do you see that?

A. Yes, I do.

Q. Now, what did you mean by the term "resolution of any potential patent issues prior to showing"?

A. My objective in making this statement was to respond to the fact that I believed that patents would be filed and I needed to get resolution on that filing activity prior to actually making a showing to the committee.

Q. But this was a statement that you did make to the JEDEC committee; is that right?

A. Yes. It's in the chart.

Q. So in other words, you were putting them on some type of notice; is that right?

A. I was warning the committee of expected or planned patent activity within IBM Corporation associated with the material that I was showing.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document marked as JX-9.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document represents the official JEDEC meeting minutes for a JC-42.5 meeting dated December 2, 1991.

Q. So in other words, this is the next meeting after the set of minutes we just looked at?

A. Yes.

Q. Now, did you make another presentation with respect to the eight-byte SIMM at the December 1991 meeting?

A. Yes, I did.

Q. If I could ask you to turn, please, to page 5.

If we could pull up item 9.9. It's a little bit difficult to read on the hard copy, but I think we can pull it up on the screen.

If you'd take a look at item 9.9, is that the presentation that you made?

A. Yes, it is.

Q. If I could ask you then to turn next, please, to page 21 in JX-9.

And at page 21, is that the beginning of the actual presentation that you made?

A. Yes, it is.

Q. And if I can ask you to turn to the next page, on page 22, please.

If you could please summarize just in very,

very general terms what's reflected on page 22.

A. This page reflects the basic attributes of a memory module being proposed, including the operating voltage, the pin count, the rationale behind the form factor that we were proposing, a few other attributes.

Q. Now, is this essentially a repeat of the presentation that you made in the September meeting or was this in some way different?

A. It was a repeat in that it was the same fundamental concept. Some of the words are the same. I edited the previous file. I simply added the material.

Q. So you added some additional detail in this presentation?

A. Yes, I did.

Q. If I could ask you to turn to the next page, please, page 23.

There it reads "patent position." Do you see that?

A. Yes, I do.

Q. Under that it reads: "IBM currently has a patent application on this product proposal for which no final determination has been made regarding the breadth of claims nor the final disposition of the application. IBM has a standard policy regarding the

licensing of patents, which is attached for reference."

Do you see that?

A. Yes, I do.

Q. Now, first of all, with respect to the reference to a patent application, had an application actually been filed with the Patent and Trademark Office as of December of 1991?

A. No, it had not.

Q. So this is still before the application was actually filed?

A. Yes, it is.

Q. Now, the second line reads "for which no final determination has been made regarding the breadth of claims." Do you see that?

A. Yes, I do.

Q. And what did you intend to convey in that phrase?

A. This is clarification that in fact the patent application had not been finalized and therefore we weren't really sure what we'd be claiming yet in the patent application itself.

Q. And in the next line there's a reference to IBM has a standard policy regarding the licensing of patents. Do you see that?

A. Yes, I do.

Q. If I could ask you to turn, please, to page 24.

This is a page with the caption IBM Worldwide Patent Licensing Practice. Do you see that?

A. Yes, I do.

Q. Is that the standard policy regarding license of patents that you referred to in page 23?

A. Yes.

Q. Now, the statement that is included here on page 24 of JX-9, was that the final IBM position with respect to licensing policy with respect to JEDEC members on this particular product?

MR. PERRY: Objection. Lacks foundation and vague as to what he means by "final."

MR. OLIVER: Your Honor, this is the IBM representative who was presenting this particular document to the committee. My question is whether he or others at IBM had any intention of ever presenting any other document that might reflect different terms.

MR. PERRY: Well, that's a different question about the intention of others. I just think the question is unclear at this point.

JUDGE McGUIRE: Sustained.

Just restate the question, Mr. Stone.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Based on your understanding, did you understand the document reflected at page 24 of JX-9 to be IBM's final position with respect to its licensing policy vis-a-vis JEDEC members?

A. No. This was not the final IBM position.

Q. What was your understanding?

A. My understanding is this is a standard document. It does not reflect the results of any negotiations on a specific patent activity.

Q. Did IBM ever come back to JEDEC with any other statement of its licensing position with respect to the eight-byte SIMM product?

A. Not reflecting this proposal, no.

Q. Why not?

A. Because this proposal was not well-received by the committee and ultimately dropped by IBM.

Q. What do you mean by "not well-received by the committee"?

A. They didn't really like it.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document marked as CX-30 for identification.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document is a copy of the official JC-42.5 committee meeting dated February 26, 1992.

Q. If I could ask you to turn, please, to page 6.

A. Yes.

Q. And if I could direct your attention in particular to item 10.1 towards the bottom of the page. And it's very difficult to read on the hard copy, but it might be easier on the computer screen.

Item 10.1 reads: Item 370 IBM eight-byte SIMM 162 pins. IBM gave a second showing on this proposal (see Attachment R).

Do you see that?

A. Yes, I do.

Q. Now, was this the same item that you had discussed at the September 1991 and presented at the December 1991 meetings?

A. It's fundamentally the same item and it shared the same item number I believe.

Q. Were there any differences between this presentation and the one at the earlier one?

A. This presentation had differences in it as compared to the prior showing and those differences

were adequate that the patent activity no longer applied.

Q. Let me direct your attention to the following sentence. It reads: "A patent has been applied for on the 180-pin version by IBM. The patent does not apply to the 162-pin proposal IBM noted."

Do you see that?

A. Yes, I do.

Q. Now, the presentation during the February 1992 minutes, that was the 162-pin version?

A. Yes. I believe it was.

Q. And the presentation that you had made in September and December of 1991, was that the 162-pin version or was that the 180-pin version or was that a different version?

A. That was what became the 180-pin version, yes.

Q. So in other words, the patent application that you had disclosed in December related to the 180-pin version, that was the application you're referring to here with this statement?

A. Yes.

Q. Now, what led you to make the statement that a patent had been applied for on the 180-pin version but it did not apply to the 162-pin version?

A. The changes associated with the proposal we are

making here were sufficient that we had effectively designed around our own patent application.

Q. I guess the question I'm asking is, as I understand it, the JEDEC policy would not require you to disclose if something did not apply, and so I was wondering why it is that you made a statement that the IBM patent application did not apply.

A. Just to clarify for the committee that this proposal was sufficiently different that the discussion from the December meeting in regard to patent activity didn't apply to this showing. I wanted to differentiate the showings.

Q. And by the way, had the patent application actually been filed with the Patent and Trademark Office by this time?

A. To my recollection, it had, but I can't confirm that without checking.

Q. I'd like to ask a couple questions about JEDEC process and JEDEC goals if I could and then if I could ask based on your understanding, based on your attendance at JEDEC meetings.

Did you understand it to be a goal of JEDEC to standardize the best technologies?

A. I have difficulty using the term "the best technologies."

The way I describe the JEDEC process is a lot of good ideas are brought into a room and good companies and good people debate attributes associated with those proposals, and what ultimately emerges from the JEDEC process is ideally a workable solution, one that's acceptable to a large majority of the participants in JEDEC, and ideally one that we can classify as truly an open standard without a lot of patent concerns or other impediments to adoption.

Q. And is it, based on your understanding, is it usually the best technology that prevails?

A. I think I've used the term before that I lose sleep over a lot of these standards. Do I perceive JEDEC standards as the best? I'm sure some are, some aren't, and then there's a whole bunch in the middle. So I have difficulty with the word "best."

Q. By the way, we referred earlier today to the quad CAS technology. Do you recall that?

A. Yes, I do.

Q. In the 1994 time period, when your notes reflect a discussion of that, did you have an understanding of what the quad CAS technology was?

A. Yes, I did.

Q. And can you please explain just very briefly what your understanding of that technology was?

A. The quad CAS device was a concept which permitted a parity memory module, in this case typically a four-byte parity memory module, to apply all of the parity that's associated with that parity memory module into a single device.

Now, just to clarify that, in effect, one device could replace four discrete devices that would otherwise be required to provide that function.

Q. Now, based on your understanding in the early 1994 time period, did you view that as a valuable technology?

A. Yes, I did.

Q. Is that a technology that you recommended that IBM use?

A. Yes. And IBM did use that technology.

Q. Is that a technology that in the early 1994 time period that you recommended that JEDEC standardize?

A. I believe I did.

Q. I believe your notes that we looked at earlier reflected a certain discussion that took place at a JEDEC meeting. If you want, I think the notes are still in front of you at CX-2375, if you would like to look at them.

A. Yes.

Q. If I could ask you to summarize generally what the issue was that was being discussed as reflected on pages 1 and 2 of your notes, CX-2375.

MR. PERRY: Your Honor, could I just be clear that we're talking about whether or not he has a recollection or not. If he's just summarizing these notes, I don't think it's a proper question.

MR. OLIVER: I'll withdraw the question, Your Honor.

JUDGE McGUIRE: Restate.

BY MR. OLIVER:

Q. If you could please look at your notes, and then I'll ask you a question to see if that has refreshed your recollection.

A. Okay.

Q. If I could just ask you to set your notes down.

If I could just ask your recollection of what was the subject of the discussion at that meeting with respect to quad CAS technology.

A. The subject discussion at this point in time -- and I must state, when questions are asked in the context of time when I don't have the material in front of me, it is often difficult to place specific events to specific times without specific reminders -- this

March '94 meeting was associated with our reaction to the existence of known patents and ballots that were in process within JEDEC.

Q. Perhaps my questions were not sufficiently clear.

If I could ask specifically to an earlier time period, before JEDEC became aware of any TI patents relating to quad CAS, at that point in time did you recommend that JEDEC standardize the quad CAS technology?

A. Yes. That was the clarification I was hoping to achieve. To go beyond that, not only was I in support of the concept, we were designing that function onto our own DRAM devices.

Q. I see.

And why were you in support of the concept at that time?

A. Because it was a good solution to allow us to replace four discrete devices with a single device providing the same function.

Q. Now, at some point in time did you come to learn that Texas Instruments had one or more patents relating to that technology?

A. Yes, I did.

Q. And do you recall approximately when you

learned that?

A. Only within the context of the notes. It appears as though it was late '93, but I don't recall other than via the notes.

Q. Now, what, if any, effect did your understanding that Texas Instruments had one or more patents relating to quad CAS have on your decision to support a JEDEC standardization of the quad CAS technology?

A. Once IBM became aware of the patents or of the patent that read on quad CAS, from a JEDEC perspective, we were obligated to request all activity be stopped. We were obligated to consider work-arounds.

By "obligated," I'm simply describing the fact that the JEDEC guidelines discuss the fact that it's important for the committee to consider alternatives, and this was certainly a case where we'd have to consider alternatives.

Q. And did JEDEC in fact do that?

A. By "that" you mean we did terminate activity, the ballots were stopped, there was lots of discussion as evidenced by this particular meeting and I believe that continued for some series of meetings.

Q. Okay. Mr. Kellogg, are you aware of any circumstances when a company advised JEDEC that it had

relevant patent rights and agreed to the RAND costs, in other words, assured JEDEC that it would make licenses available on reasonable and nondiscriminatory terms, yet JEDEC nevertheless chose to investigate alternative technologies?

A. Yes, I do.

Q. Can you please describe the examples that you can think of?

A. One very good example I remember was associated with Cypress. Cypress disclosed a patent associated with a PLL power-down mode. This is a device that we're using on memory modules for our synchronous memory standard.

And in that case, Cypress disclosed that the method by which we were powering or reducing the power dissipation on the device was covered by one of their patents.

The committee did consider the alternative of continuing to use the method that Cypress was claiming and that we had standardized, but we also investigated alternatives, and ultimately we did adopt an alternative which -- which was somewhat painful but not significantly so, fortunately, in that case, but we did adopt an alternative.

Q. Are you aware of any other circumstances in

which a company advised JEDEC that it had relevant patent rights and agreed to RAND policy, but JEDEC nevertheless investigated alternative technologies?

A. Well, there are several. A couple of them are associated with a company called Kentron.

Kentron was regularly coming in with innovative -- what I perceived as relatively innovative concepts, generally associated with memory modules. And one of them was a card-on-card assembly which would allow us to build high-density modules without this stacking technology that we were currently using.

Another was a way of doubling the data rate off a module without the need of using a technology that doubled the data rate on the memory device itself, so it was a way of doubling data rate on a card.

They also had some others.

And I think we had another -- for example, I think Hyundai came forward with a TRIMM, a triple in-line memory module, for which they disclosed patent activity and we simply decided there wasn't enough value there, we're going to do something different, so we did something different in that case.

And there are others, but those are immediate ones that come to mind.

Q. Did IBM ever have a situation in which it disclosed a relevant patent application and promised RAND terms and yet JEDEC considered alternative technologies?

A. Yes. There were several IBM proposals for which we disclosed patent activity and the committee showed no interest. I'm not exactly sure why they didn't. We think they were good ideas. But alternatives existed.

MR. OLIVER: Your Honor, I see it's approaching five o'clock. This witness does have I think quite a bit more to cover. I think I would have a logical breaking place in about 20 to 25 minutes. Would it be acceptable to perhaps continue to that?

JUDGE McGUIRE: Yes, that's fine.

BY MR. OLIVER:

Q. Mr. Kellogg, let me turn now to any involvement that you might have had with Rambus outside of JEDEC, and let me ask first, did you have any involvement with Rambus outside of JEDEC?

A. Yes, I did.

Q. Can you please describe what that involvement was?

A. The initial involvement was associated with some meetings held at IBM where Rambus had

representatives present their basic structure associated with their first-generation Rambus device.

Q. And did you attend those meetings?

A. I attended several meetings. I don't recall what percentage of those meetings I attended.

Q. Now, at that time what did you learn about the Rambus products or Rambus technology?

JUDGE McGUIRE: What time is it we're talking about? That's not clear.

BY MR. OLIVER:

Q. Okay. Mr. Kellogg, can you please explain what time period you're referring to?

A. The time frame I was referring to in my response is associated with the early '90s.

Q. And can you please explain at that time what you learned about the Rambus technology?

A. I became familiar with the Rambus technology or proposal as being a combination of a memory device, a memory form factor, a memory bus structure, a clocking structure, which I would describe as a high-speed, narrow-I/O, packetized interface.

Q. I believe you in the beginning part of your answer referred to a memory device, a form factor and a bus structure.

Can you please explain first what you mean by

"a form factor"?

A. Well, I'm almost using "memory device" and "form factor" in a similar manner. The package that the memory chip was placed in was also unusual in that it was a -- I believe at the time it was a vertical package, which was ideally suited for some of the proposals they were making to us.

Q. Now, at the time that Rambus was making these presentations to IBM, did you have any understanding as to whether the Rambus technology would be suitable for general use in the IBM product line?

A. I spent some portion of time on my own as well as talking to others considering this Rambus concept, this initial proposal, in pretty much the full range of IBM systems, PCs up through servers, just to see if there was something here that we could take advantage of to have some value to our systems, density, performance, reliability, something of that nature.

Q. And what was your conclusion?

A. My conclusion for the products that we were considering that this was not an optimal solution.

Q. Why not?

A. It was somewhat dependent on the application.

In the case of the servers, it didn't provide us a means of achieving the reliability or density, the

amount of memory chips per unit area that we needed to achieve our performance objectives. Performance to some extent is related to density.

In our PC products and some of the others, it just didn't seem to be the right answer in regard to things such as latency or some of the commands, some of the operations, the structure itself, the pluggable nature. It just didn't seem to be well-suited for the PC environment. It seemed cool for a game. And --

Q. You may have anticipated my next question. My next question is whether there were any particular uses that you understood Rambus technology might be well-suited for.

A. Well, I believe even the Rambus team that we visited with in this time frame was using as a demonstration vehicle an application that did not require a large number of memory devices and could take advantage of the bandwidth that the device was offering. And those applications were very predominantly game or in some cases more video-intensive than what we saw in our systems.

Q. Now, at the time that Rambus was making these presentations to IBM in the early 1990s, did you have an understanding as to whether Rambus had or had applied for any patents relating to its technologies?

A. In the early '90s time frame, Rambus approached us through an IBM executive to disclose information. My recollection is in that disclosure -- I think we had to sign a nondisclosure agreement of sorts -- we were made aware of the fact that there was patent activity underway.

So I was aware and I have recollection of awareness of patent activity on Rambus' part associated with the products they were showing to us at the time.

Q. Now, at that time what, if any, was your understanding of the scope of coverage of Rambus' patents or patent applications?

A. The only understanding I would have had would have been my own expectations, and my expectations would have been that patent activity would be associated with the product they were showing us, the specification, which to me was a narrow-I/O, high-bandwidth, packetized memory device or a card with a loop-back clock structure and a few other elements.

Q. Now, with respect to the meetings that you attended, the meetings involving Rambus representatives at IBM, did anyone from Rambus ever suggest to you or to others in those meetings that you attended that its

proprietary technology extended to products outside its own RDRAM architecture?

A. Not to my recollection.

Q. Now, at any time did anyone from Rambus ever tell you or suggest to you that its proprietary technology extended to products outside its RDRAM architecture?

MR. PERRY: Objection. Compound.

JUDGE McGUIRE: Sustained.

BY MR. OLIVER:

Q. At any time did anyone from Rambus ever suggest to you that its proprietary technology extended outside the RDRAM architecture?

A. I don't believe they did. I have no recollection of that.

Q. Now, do you have an understanding as to whether IBM ever entered into a license agreement with Rambus?

A. Yes.

Q. Do you know for what technology IBM entered into a license agreement with Rambus?

A. I recall that IBM signed at least two licensing agreements, one associated with the design and manufacturing of a memory device or devices of varying densities and the second being a logic or

license of some form -- I never saw it; I was just familiar with it -- that permitted us to design an interface set of circuits that would allow us to communicate between our ASICs or our application-specific integrated circuits and Rambus devices on the channel.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document that's been marked as CX-2370.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document is a set of notes that I took during a JEDEC JC-42.3 meeting in New Orleans on May 7, 1992.

Q. If I could ask you to turn, please, to page 3 of these notes.

And I'd like to direct your attention to the fourth heading if you will. It reads "Siemens." Do you see that?

A. Yes, I do.

Q. And it says "Siemens: Kernel of chip. Similar to Rambus. Patent concerns? (No Rambus comments)."

Do you see that?

A. Yes, I do.

Q. Now, following this meeting, what, if anything, did you do with respect to the information that's reflected in those notes?

A. I don't recall taking any action following the meeting on that comment.

Q. And why not?

MR. PERRY: Your Honor, I just want to make sure the question is do you recall or based on his recollection as to why not. Otherwise --

JUDGE McGUIRE: Sustained.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Following the May 1992 meeting, did you have an understanding that Rambus had intellectual property applicable to SDRAMs?

A. No, I did not.

Q. And why not?

A. Why would I?

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Mr. Kellogg, I've handed you a document marked as CX-2374.

Do you recognize this document?

A. Yes, I do.

Q. What is this document?

A. This document is a set of notes that I compiled at a JC-42.3 meeting in 9-22-93.

Q. If I could direct your attention on the first page, the third main bullet "patent discussion." Do you see that?

A. Yes.

Q. Under that it reads "added Motorola BGA patent" and there's a number and then under that "added Rambus patent number 5,243,703 9-7-93."

Do you see that?

A. Yes, I do.

Q. Now, did that indicate to you that Rambus might have intellectual property relating to JEDEC's ongoing work?

A. No, it did not.

Q. Why not?

A. This is a disclosure of a patent with no reference to any JEDEC activity, so it's strictly a patent.

Q. Did you have any understanding of the scope of coverage of this patent?

A. I don't think I read this patent, so I don't

think I had a view of the coverage of the patent itself.

Q. Now, just so the record is clear, did anyone from Rambus ever make any disclosures at JEDEC of any pending patent applications, to the best of your recollection?

A. I don't think so. I don't recall one.

Q. During the time that you were at -- or you were participating in JEDEC, do you recall anyone from Rambus ever doing anything to put you on notice that Rambus had or might have intellectual property relating to ongoing JEDEC work?

A. No.

And I think it's important to clarify something. Coming to JEDEC and saying you have a patent is nothing unusual. Lots of companies have patents. It's extremely important to have a reference for the patent.

In general, patents are disclosed during activity, during a discussion on a product. And for a patent to be thrown on the table in front of a room full of engineers or just described doesn't allow us to comprehend the implications of the patent.

So unless there's a relationship, unless we understand how the patent affects an activity, then I

don't know what we'd do with that.

Q. Is there any relationship between your last answer and the Rambus disclosure of its '703 patent?

A. My intention was to describe this and other disclosures of this nature in the context that I believe this disclosure and any other disclosure that is unrelated and not coupled to a specific activity doesn't allow us to comprehend the implications of the disclosure.

Q. At the time that Rambus disclosed its '703 patent, did you have any understanding of the context in which Rambus was disclosing that patent?

A. "Understanding" might be the wrong word. Did I have an expectation, a belief? My belief was that this patent was a patent that was associated with the devices that they were showing us in the early '90s and that the patent was associated with a high-bandwidth, narrow-I/O, packetized interface with loop-back clock and other similar attributes.

Q. At the time that Rambus disclosed its '703 patent, did Rambus provide any explanation of how, if at all, that patent might relate to ongoing JEDEC work?

A. I don't recall Rambus disclosing. If they had, I feel I would have written something down.

JUDGE McGUIRE: Well, then let me ask a question.

Under your understanding of the patent policy, when one discloses a patent, are you saying then that if they haven't also disclosed the implications of the patent, have they I guess adequately then disclosed the patent under the patent policy?

THE WITNESS: No. That's kind of my point, and I appreciate the clarification. Because I'm not a patent attorney and these patents are pretty complex.

JUDGE McGUIRE: We're just talking about -- the question I'm asking you just deals with your understanding of the patent policy at JEDEC.

THE WITNESS: Yes.

JUDGE McGUIRE: Okay.

THE WITNESS: Within the context of the patent policy at JEDEC, disclosure of a number I don't believe meets the patent policy. If the number is disclosed not in any context of anything else.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Now, when, if at all, did you first learn that Rambus contended that it had patent rights relating to JEDEC-compliant products?

A. I believe my first awareness of that would have

been with the first suits that were filed and somewhere in the 2000 time frame.

MR. OLIVER: Your Honor, this will probably be a convenient place to break for the day.

JUDGE McGUIRE: Okay. Very good. Now, I understand he's going to be coming back on Friday.

Do we have some time frame as to how much more time will involve his examination, for certain at least complaint counsel, on direct?

MR. OLIVER: On direct, Your Honor, I expect it could take another three hours or so.

JUDGE McGUIRE: And then Mr. Perry, how about cross? Are you able to estimate?

MR. PERRY: Well, Your Honor, it's hard to know. I didn't think it would be another three hours, so there may be subject matters that I was unaware of, which means I would have to go longer than I thought I had to go. At the moment I think I probably have an hour to an hour and a half.

JUDGE McGUIRE: So that will put us tomorrow to about three hours to 12:30 and then convene and then be done by 3:30 or so. Okay. I just want to get some feel.

If we adjourn, if we complete that, you know, before, say, 3:00, is there any alternative plan as to

what else we want to do, if anything? Involving reading of depositions or anything like that?

MR. OLIVER: Yes, Your Honor. We still have transcripts of Mr. Karp and --

JUDGE McGUIRE: I know. And I was hoping to get an order out on that on Friday. I'm not sure I'm going to be able to do that.

Can we go into that in areas that won't be impacted by the pending motion on that deposition?

MR. OLIVER: I'm not the best person to answer that, but I believe we probably can, Your Honor.

JUDGE McGUIRE: I'm just asking this in case we get done at 2:30 or so, I don't want to adjourn and we get still -- we've only got about another two weeks for you to -- well, we have the time it takes for you to complete your case in chief, but at least the understanding is you're trying to have that done by the end of June, so I want to take full advantage of any time we have.

So I'm just asking these questions so I'll have some idea as to what we're going to be doing on Friday.

Any comments?

MR. STONE: I just was going to say I think there probably is a way. We can confer with complaint

counsel about the Karp deposition and find sections that don't involve the issues that are --

JUDGE McGUIRE: Okay. Very good.

Then we will convene tomorrow morning at 9:30.

This hearing is in recess.

(Time noted: 5:19 p.m.)

C E R T I F I C A T I O N O F R E P O R T E R

DOCKET NUMBER: 9302

CASE TITLE: RAMBUS, INC.

DATE: June 12, 2003

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the notes taken by me at the hearing on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: June 13, 2003

JOSETT F. HALL, RMR-CRR

C E R T I F I C A T I O N O F P R O O F R E A D E R

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

DIANE QUADE