FEDERAL TRADE COMMISSION

I N D E X (PUBLIC RECORD)

UNITED STATES OF AMERICA

FEDERAL TRADE COMMISSION

In the Matter of:)

Rambus, Inc.) Docket No. 9302

---------)

Monday, June 16, 2003

9:30 a.m.

TRIAL VOLUME 28

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
CARY ZUK, Attorney
SUZANNE MICHEL, 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.

Counsel, good morning. We want to say hi to you back again, Mr. Detre. We heard the good news and we're very pleased to hear that.

MR. DETRE: Thank you very much, Your Honor.

JUDGE McGUIRE: Congratulations.

MR. DETRE: Thank you, Your Honor.

JUDGE McGUIRE: Any items we need to take up before we get started this morning?

MR. PERRY: Your Honor, we'd like to move in a few exhibits that we used on Friday. I have five exhibits that I discussed with Mr. Kellogg.

RX-236, an e-mail by Mr. Kellogg.

MR. OLIVER: No objection.

JUDGE McGUIRE: Entered.

(RX Exhibit Number 236 was admitted into evidence.)

MR. DETRE: RX-250, another e-mail by Mr. Kellogg.

MR. OLIVER: No objection.

JUDGE McGUIRE: Entered

(RX Exhibit Number 250 was admitted into evidence.)

MR. DETRE: RX-1443, a Toshiba document about a meeting Mr. Kellogg attended.

MR. OLIVER: No objection.

JUDGE McGUIRE: Entered

(RX Exhibit Number 1443 was admitted into evidence.)

MR. DETRE: RX-1695, another e-mail by Mr. Kellogg.

MR. OLIVER: No objection.

JUDGE McGUIRE: Entered.

(RX Exhibit Number 1695 was admitted into evidence.)

MR. DETRE: And RX-562, some JEDEC meeting minutes.

MR. OLIVER: No objection.

JUDGE McGUIRE: Entered.

(RX Exhibit Number 562 was admitted into evidence.)

MR. OLIVER: Your Honor, we also wish to move in one exhibit, CX-110, the meeting minutes from the 42.5 subcommittee of June 1996.

MR. DETRE: No objection.

JUDGE McGUIRE: So entered.

(CX Exhibit Number 110 was admitted into evidence.)

JUDGE McGUIRE: Thank you very much.

Anything else we need to take up this morning before we begin?

MR. OLIVER: No, Your Honor.

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

MR. OLIVER: Thank you, Your Honor.

Complaint counsel calls Professor Bruce Jacob.

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

- - - - -

Whereupon --

BRUCE LEDLEY JACOB

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

JUDGE McGUIRE: Go ahead, Mr. Oliver.

DIRECT EXAMINATION

BY MR. OLIVER:

Q. Good morning, Professor Jacob.

A. Good morning.

Q. How are you today?

A. I'm good.

Q. Could you please state your full name for the record.

A. Bruce Ledley Jacob.

Q. Professor Jacob, what is your current position?

A. I'm an associate professor at the University of Maryland in the electrical and computer engineering department.

Q. How long have you been a professor at the University of Maryland?

A. Six years.

Q. Are you tenured?

A. Yes, I am.

Q. When did you get tenure?

A. This past year.

Q. Congratulations.

A. Thank you very much.

Q. What does it mean to receive tenure at the University of Maryland?

A. It's a lifetime appointment and it means that you've done a very good job of doing research and advising students, employing students, that sort of thing, teaching.

Q. Could you please describe in general terms your field of research.

A. In general terms, I'm a computer architect and I study memory systems, meaning DRAM systems and cache

systems and that sort of thing.

Q. Now, why did you decide to specialize in memory systems?

A. Because it is perhaps the most important facet of computer design when you're trying to build faster systems, systems that perform better. The memory system is more important than speeding up the CPU at this point.

Q. Professor Jacob, let me take a step back and ask first about your educational background.

Did you receive an undergraduate degree?

A. Yes, I did.

Q. What was that degree?

A. The honors baccalaureate, a bachelor's, cum laude in mathematics from Harvard.

Q. Now, after graduating from Harvard, did you do any professional work relating to computer engineering?

A. Yes, I did.

Q. What was your first job in computer engineering?

A. I was a software engineer at Boston Technology.

Q. What did you do at Boston Technology?

A. I was -- I designed the software applications

for their embedded system.

Q. Can you explain in a bit more detail what your responsibilities were there?

A. I was building their voice mail system. It's a distributed telecommunications product and I was writing the software code that would implement things like foreign-language systems, Japanese, French, Spanish.

Q. Now, after Boston Technology, did you take another job in the computer engineering field?

A. Yes, I did.

Q. What was your second job?

A. I worked for a company called Priority Call Management. I was the sole engineer for the first nine months or so and I was the system architect. I implemented their product. I designed it and implemented it. It was also a distributed telecommunications product with a slightly different focus than Boston Technology's.

Q. Now, at some point did you decide to pursue graduate studies?

A. Yes, I did.

Q. And what graduate work did you pursue?

A. I obtained a master's and Ph.D. in computer science and engineering at the University of Michigan

in Ann Arbor.

Q. Can you please explain in a little bit more detail your field of study in your graduate studies?

A. I was trained to be a computer architect.

Q. What does that mean?

A. That means that I studied the design of complex chips and systems of chips, so, for example, CPU chips, memory controller chips and DRAM chips.

Q. Did you write a master's thesis?

A. Yes, I did.

Q. What was the subject of your master's thesis?

A. The subject was memory systems, specifically cache systems and DRAM systems and architectures.

Q. Can you describe in a bit more detail the subject of your master's thesis?

A. Yes. It was the optimization of cache systems and memory systems -- and DRAM systems, as well as disk subsystems. The optimization, for instance, in a given technology and cost, the performance parameters, what's the best arrangement of the system.

Q. Now, did your master's thesis work involve DRAMs?

A. Yes, it did.

Q. And what sorts of DRAMs were you studying as part of your master's thesis work?

A. The DRAMs of the day. EDO, fast page mode, that sort of thing.

Q. Now, did you write a Ph.D. dissertation?

A. Yes, I did.

Q. What was the subject of your Ph.D. dissertation?

A. It was also in memory systems, so specifically cache systems and DRAM systems, architectures as well as software means for controlling their operation.

Q. Now, when did you become a professor at the University of Maryland?

A. In the fall of 1997.

Q. And since becoming a professor at the University of Maryland have you done any research involving DRAM architectures?

A. Yes, I have.

Q. We'll talk in just a moment about the specific studies that you've conducted, but before we do, could you just mention what types of DRAM architectures you've evaluated in the course of your studies.

A. Yes. We evaluated fast page mode; EDO; synchronous DRAM, SDRAM; several different Rambus designs, including concurrent Rambus and direct Rambus; SyncLink DRAM, SLDRAM; DDR; DDR-II. So a very wide range.

Q. Now, I'd like to talk a bit about the specific studies that you've conducted, and perhaps we'll just simply try to take them in chronological order if we could.

What was the first work that you did at the University of Maryland involving DRAMs?

A. In the winter of '98, so this is the January-February time frame, I began a study of -- a comparative performance evaluation of different DRAM architectures that were commercially available at the time, including fast page mode, EDO, SDRAM, concurrent Rambus and direct Rambus, which had just been defined I believe, so we did a very accurate modeling of these architectures to determine what the performance differences would be and why.

Q. Now, when you talk about modeling the performance of DRAMs, what does that involve?

A. That involves -- that's a very accurate, very precise software representation of the behavior of the device or system under study, so you need to really thoroughly understand how something works and then represent that behavior in software.

Q. Now, what type of software do you use to do that?

A. You write the software yourself.

Q. Now, in the case of this study that you started in 1998, who wrote that software?

A. I wrote some of it and my graduate students wrote some of it.

Q. Now, can you please explain in a little more detail what your role in that study was?

A. Well, as I said, I wrote some of the software. I directed the work of my graduate students, told them what experiments to run, how to plot the results. I looked at the results, you know, gained the insights. I wrote the bulk of the paper.

Q. How long did it take to complete that study?

A. The study -- the research took around nine to -- somewhere between nine and twelve months, and the paper write-up took around three, four, five months.

Q. Why did it take so long to complete the study?

A. Because this is nontrivial work. This is what architecture design is about. That one study represents one architectural investigation.

Q. Now, did you publish the results of that study?

A. Yes, I did.

Q. And where was it published?

A. It was published in 1999 in ISCA, the International Symposium on Computer Architecture, which

is the premier forum for research in computer architecture.

Q. Now, did you receive any type of award in connection with that study?

A. Yes, I did. For this research I -- or based upon this research I received the Prestigious Career Award from the National Science Foundation for my work in DRAM systems and architectures.

Q. What is the Prestigious Career Award from the National Science Foundation?

A. It's an award for young professors who do outstanding work.

Q. Now, what was the next study that you conducted at the University of Maryland in connection with DRAMs?

A. Well, the first thing we did was to extend that study, for example, adding in DDR to the mix and looking at a handful of other parameters, and published that in the IEEE Transactions on Computers in 2001.

Q. Okay. Can you please describe that extension of the study in a little more detail?

A. Well, for example, we looked at more architectures, changed the model so that we could obtain more -- obtain more information.

Q. And was the type of modeling that you were doing similar to the first study?

A. Yes. It was very similar.

Q. And you say the conclusions of that study were also published?

A. Yes, they were.

Q. Where were they published?

A. In 2001 in the IEEE Transactions on Computers, a journal in the field.

Q. Now, what, if any, studies have you performed while at the University of Maryland?

A. We also did a study of a higher-level view. Rather than just DRAM architectures, we took a system-level perspective and did a very -- an in-depth modeling of a memory controller, a very high-performance memory control, investigated different DRAM interfaces as well as bus organizations.

So this was a very system-level approach to the design.

Q. And again, can you please explain your role in that study?

A. I wrote -- I wrote all of the software for that study and I directed my students to run the pertinent experiments, showed them how to graph things, told them

what parameters to graph and how, and I wrote pretty much the entire paper on that one.

Q. How long did it take to complete that study?

A. That study began in early 1999 and ran through late 2000, early 2001, and the study was published in 2001 also in ISCA, the International Symposium on Computer Architecture. That was about a year-and-a-half study.

Q. And again, why did it take so long to complete that study?

A. Because it's nontrivial. There's a lot of work to these types of studies.

Q. Now, since the completion of the extension study that was published in the IEEE and the higher-level study you just mentioned was published in the International Symposium on Computer Architecture, have you done any additional work at the University of Maryland involving DRAMs?

A. Yes. I continue to direct a fairly large group of Ph.D. students who are investigating advanced issues in the design of DRAMs and DRAM systems, and currently we are pulling together a number of studies to put into a large treatise on DRAM systems and architectures.

Q. Do you have plans to publish that treatise?

A. Yes, I do.

MR. OLIVER: Your Honor, at this point complaint counsel tenders Professor Bruce Jacob as an expert in the field of memory architectures and systems.

JUDGE McGUIRE: Any opposition by respondent?

MR. DETRE: No opposition to the extent that Professor Jacob can testify based on an academic understanding of those fields, but if he -- since he doesn't have any experience in actual DRAM design --

JUDGE McGUIRE: You can take that up on cross-exam. Okay? Otherwise, he's deemed qualified in that area as an expert.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Professor Jacob, before we discuss your conclusions in detail, could you please summarize the questions that you were asked to address in connection with this case.

A. Yes. I was asked to determine if alternative technologies to the technologies in dispute existed at the time that JEDEC was considering these things, so that in the early 1990s time frame.

I was asked to determine if engineers in that time frame would have suspected that Rambus could claim

intellectual property rights over the work that JEDEC was doing based on the '898 application.

I was asked to look at Rambus patents that were pending at the time and granted at the time, looking at the claims in those patents to see if those claims would cover work that was going on in JEDEC 42.3.

And I was asked to look at the implications of modern redesign, what if DRAM -- what if the DRAM industry had to design new DRAMs to remove the technologies under dispute and replace them with alternatives.

Q. Let's look at these questions one at a time if we could, and starting just with a general overview, could you please summarize briefly your conclusion as to whether in the 1991 to 1996 time period engineers had available to them technological alternatives to the four technologies at issue in this case.

MR. DETRE: Objection, Your Honor.

Professor Jacob has not done any kind of modeling or testing or study of the economic feasibility of his alternatives. Given that, it's speculative. There is case law that alternative design evidence should not be admitted in the absence of some sort of modeling or testing that those alternatives

would actually be feasible or work.

So we object on the grounds of speculation.

JUDGE McGUIRE: Overruled.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Professor Jacob, do you have the question in mind?

A. No, I do not.

Q. Let me restate it then.

Could you please summarize briefly your conclusion as to whether in the 1991 to 1996 time period engineers had available to them technological alternatives to the four technologies at issue in this case.

A. Yes, they did.

Q. And can you please state briefly your conclusions as to whether Rambus' '898 patent application or the patents listed in Rambus' withdrawal letter to JEDEC would have alerted reasonable engineers that Rambus could claim patent rights over the subject matter of JEDEC's SDRAM and DDR SDRAM work?

MR. DETRE: Objection, Your Honor.

Professor Jacob, based on his own definition of a person of ordinary skill in the art to which that

patent application relates, is not a person of ordinary skill because he does not have the DRAM design experience required, and he simply cannot testify about what somebody who is of ordinary skill in that art would have understood from looking at that patent application.

He's a very qualified gentleman in many respects, but he's not qualified to speak to that topic.

JUDGE McGUIRE: Mr. Oliver, any response?

MR. OLIVER: Yes, Your Honor. We will be going through the bases for his conclusion in quite a bit of detail. There's ample opportunity for Mr. Detre to go into that on cross-examination.

JUDGE McGUIRE: Overruled.

MR. OLIVER: Thank you, Your Honor.

THE WITNESS: I believe the question is related to whether they would have suspected anything based on the information, and the answer is no, they would not have.

BY MR. OLIVER:

Q. Can you please summarize briefly your conclusion as to whether a reasonable engineer would conclude that claims in certain of Rambus' pending patent applications in the 1991 to 1996 time period or

in the issued '327 patent could cover technologies that were the subject of JEDEC's SDRAM and DDR SDRAM standards?

A. Yes. Claims covered the work that was going on in JEDEC 42.3.

Q. And could you summarize briefly your conclusion as to what the technical implications would be today of redesigning SDRAM and DDR SDRAM products to incorporate alternative technologies in place of the four technologies at issue.

A. If one were to redesign DRAMs to get rid of all of the technologies in dispute and replace them with alternatives, it would result in a DRAM that is not compatible with any JEDEC-compliant system.

Q. Let's focus then on the first question that was posed to you, and before we walk through some of the specific alternatives, I'd like to explore just briefly the work that you did to arrive at your conclusions.

Could you please summarize first, again briefly, what types of textbooks, treatises, articles, publications, et cetera, or other engineering materials that you consulted in the context of doing your work.

A. Well, I consulted a large range of engineering

material, including, for example, treatises, technical articles, things that I found on the Web.

I also read through an enormous number of JEDEC minutes and read through the presentations made at JEDEC meetings. These are the attachments of the minutes.

I read through -- well, yeah, I read through those presentations.

And I also consulted with engineers in the DRAM industry to confirm my understanding of, for instance, what was done in those meetings, what was being presented in those meetings.

Q. Just to follow up from what you mentioned, the Web, did you conduct any Internet searches in connection with your work?

A. Yes. Absolutely.

Q. And did you find any useful materials in the course of --

A. Oh, yes. Yes.

Q. You mentioned JEDEC documents, particularly minutes.

Were you referring to minutes from the meetings --

A. Yes. I'm sorry. Minutes of the meetings of the 42.3 subcommittee.

Q. And what time --

A. And presentations made.

Q. What time period of the minutes did you consult?

A. 1991 to 1996.

Q. And you also mentioned that you interviewed a number of practitioners in the field.

Do you recall the names of any of the individuals that you interviewed?

A. Yes. Certainly. For example, I spoke with Howard Sussman, Desi Rhoden, Jacob Baker, Peisl -- Martin Peisl I believe his name is -- Kevin Ryan. There's a complete list in the back of my -- in the back of both reports. But between one and two dozen different engineers.

Q. And what was the purpose of interviewing these various engineers?

A. To confirm my understanding of the contents of the JEDEC minutes.

Q. Let's turn first if we could to programmable CAS latency, and if I could ask you to first just explain very briefly what programmable CAS latency does in the context of the JEDEC SDRAM and DDR SDRAM standards.

A. Yes. In the context of a JEDEC-styled bus

architecture, programmable CAS latency is a convenience that allows parts of different generations with potentially different performance characteristics to coexist in the same system and have the same performance in that system, have the same behavior.

Q. Now, focusing on the 1991 to 1996 time period, what alternatives, if any, existed to programmable CAS latency as used in the JEDEC SDRAM and DDR SDRAM standards?

A. Well, we have a nice demonstrative here.

For example, JEDEC could have chosen to use fixed CAS latency parts.

They could have decided to program CAS latency, essentially use a fixed CAS latency part, by doing so after the packaging steps or by blowing fuses on the DRAM.

They could have chosen to scale the CAS latency with the clock frequency.

They could have chosen to use dedicated pins to transmit the latency information to the DRAM, so between the memory controller and the DRAM.

They could have chosen to explicitly encode the latency information in the control packet.

And they could have decided, for example, to

stay with an asynchronous style DRAM.

Q. Okay. Let's look at these individually if we could, starting with fixed CAS latency.

What do you mean by "fixed CAS latency"?

A. That means that a part would only have the ability to perform one -- well, to perform with one latency, to respond with one latency.

Q. Now, how could a manufacturer fix the CAS latency?

A. Well, the manufacturer, for example, could decide to define a fixed latency in the design stage and it could decide to define the fixed latency in the processing stage or it could decide to define the fixed latency in the packaging phase.

Q. Okay. Starting with the design stage, how would a manufacturer fix the --

JUDGE McGUIRE: Mr. Oliver, before we proceed, do you intend to have the screens marked as DX exhibits?

MR. OLIVER: Yes, Your Honor, I do. Unfortunately, we don't have a color copy with us. We'll see if we can print it out during the lunch break and have it to the court by this afternoon.

JUDGE McGUIRE: And then at that time we'll go into each one and have it marked? Is that the

appropriate time or is it better to do that as we go? Only for the record.

MR. OLIVER: We could assign numbers as we go.

JUDGE McGUIRE: That might be easier so we'll understand it as we go back through it. I don't mean to interject so much at this point, but I thought it would be the proper time to perhaps do that.

MR. OLIVER: Actually I think that's a good suggestion and it would make sense.

JUDGE McGUIRE: Let's go back and just go ahead and get them marked as we go.

I think we're up to DX-62 at this point.

MR. OLIVER: So this one would be 62?

MR. STONE: Whatever your first one is would be DX-62.

MR. OLIVER: DX-62 would be a slide entitled Four Central Questions.

DX-63 would be a slide entitled First Central Question.

DX-64 would be a slide entitled CAS Latency.

DX-65 would be a slide entitled Alternatives to Programmable CAS Latency.

DX-66 would be a slide also entitled Alternatives to Programmable CAS Latency with the first item "use fixed CAS latency parts" highlighted.

DX-67 would be a slide entitled Fixed CAS Latency.

BY MR. OLIVER:

Q. Professor Jacob, if you could then explain with reference to DX-67 how a manufacturer could fix CAS latency at the design stage.

A. Yes. Yes. At the design stage, the DRAM manufacturer would design a part to only perform, you know, with one CAS latency.

For example, this picture shows that you would have some CAS latency circuitry, but it would only be told to use CAS latency 2, for example.

Q. Now, could you explain how a manufacturer could fix CAS latency in the processing stage.

A. Yes. I think there's another demonstrative for that.

Q. Well, if we could move up to DX-68, please.

A. This would be a metal mask option, and the idea is that you would have hard-wired onto the chip the value 2 and the value 3 and during the processing steps as one of the final steps within the processing you would just lay down a single wire connecting either the value 2 or the value 3 to the block labeled "CAS latency circuitry," so the part at that point would become a de facto fixed latency part.

Q. Now, you referred to a metal mask.

What is a metal mask?

A. A metal mask is one of the final steps in -- well, the metal mask is the actual -- the mask that helps you lay down a piece of metal onto the semiconductor chip and it's one of the final steps in the processing stages.

Q. Would it be fair to say that it helps to establish -- to lay down connections among different elements in the chip?

A. Yes. I'm sorry. Yes. I forget who I'm speaking with.

Yes, this is -- when you lay down pieces of metal on a chip, these establish different connections between the circuits. That's how circuits talk to each other, through metal wires.

So if there's no metal wire connecting two things, that means that there is no circuit between them, there's no connection between them. That means they can't interact.

Q. Now, DX-68 shows a connection between the box CAS 2 and the CAS latency circuitry, but I understand from what you said that it would be equally possible at the metal mask stage to establish a connection, instead of between CAS 2 and the CAS latency circuitry, between

CAS 3 and the CAS latency circuitry?

A. Yes. Exactly.

Q. Now, can you please explain how a manufacturer could fix CAS latency during the packaging phase?

A. Yes. I believe we have a demonstrative for that as well.

Q. That would be DX-69.

A. Yes.

In this option, the part through the processing phases would have the configuration shown here.

So for example, you have two hard-wired values of CAS 2 and CAS 3 both connected to the CAS latency circuitry, but they would be connected through a multiplexer, labeled "mux" in the figure, and the selection of whether it would be a CAS latency 2 or CAS latency 3 would be done by that mux and the mux would be hard-wired to either the power pins or the ground pins, and this is something that could be done with a bond wire during packaging.

So you would simply -- you would fabricate one chip, but during packaging you would connect that mux to either power or ground and then select either the value 2 or the value 3.

Q. Now, you referred to a bond wire.

What is a bond wire?

A. The bond wire is a packaging mechanism that connects the semiconductor chip or, rather, the semiconductor die to the pins of the chip.

Q. Now, in comparison with use of a mode register to program CAS latency, what advantages, if any, would have been realized by using fixed CAS latency in the 1991 to 1996 time period?

A. It would be potentially a simpler design. Certainly you don't have a mode register, so that's a simpler mechanism.

You potentially would have fewer testing stages, and again, that depends on where you decide to fix the CAS latency. For example, if you fix it earlier in the design stage, you don't actually have to test the fabricated part for multiple CAS latencies.

So the test costs and design costs can go down.

Q. You referred to a simpler design. Why do you include that as among the advantages?

A. Well, because you don't have to build and test a mode register.

Q. Are you familiar with the term "die size"?

A. Yes, I am.

Q. And what does "die size" mean?

A. Thank you.

It's the size of the semiconductor die. And the cost of manufacturing goes roughly with the area to the third power, the area of this semiconductor part, so if you have a part that is 1 percent larger, it's approximately 3 percent more expensive to manufacture.

So for example, if you eliminate a mode register, you eliminate some of the size of the part and it can make it smaller and therefore cheaper.

Q. Now, what, if anything, would be the impact on die size if JEDEC had chosen in the 1991 to 1996 time period to use fixed CAS latency rather than a mode register?

A. It would have been a simpler design and therefore a smaller design. In this instance you would have eliminated circuitry, you would have eliminated the mode register, and so it would have been a smaller design and therefore smaller die.

Q. Again compared to use of a mode register to program CAS latency, what, if any, would have been the disadvantages of using fixed CAS latency in the 1991 to 1996 time period?

A. The -- for instance, the manufacturer would

have to be a bit more clear about the labeling of these devices, because if you made no changes to the memory controller, for example, it would be possible to put -- well, if JEDEC had decided to standardize, for example, on two different CAS latencies and each part had a fixed latency, then it would be possible to put DIMMs with two different latencies in a system and that could potentially cause compatibility problems, but that could have been solved by building a more sophisticated memory controller.

So you would either have the scenario where the memory manufacturers would need to be and the memory module manufacturers would need to be more explicit about the behavior of these things and users would have to understand that or, again, you could put that onto the shoulders of the memory controller.

Q. If we could turn to the next alternative you mentioned, which is determining CAS latency by fuses.

A. Yes.

Q. Can you please explain what you mean by "program CAS latency by blowing fuses on the DRAM"?

A. Well, this would be similar to having the fixed CAS latency part because what you would have is a de facto fixed CAS latency part, but it would become that fixed CAS latency part after you've blown the

fuse, so this would be, for example, an option that you would do after packaging of the die so that the DRAM manufacturer could ship a part that was capable of performing as a CAS latency 2 part or a CAS latency 3 part, ship that part to the OEM and the OEM would blow a fuse and it would at that point become a fixed latency part, but it would have either 2 or 3.

Q. I believe we have a couple of demonstratives that help explain this, if we could bring up -- actually DX-70 I think will be the slide entitled Alternatives to Programmable CAS Latency with number 2 highlighted and DX-71 will be the first demonstrative relating to blowing fuses. It bears the caption Set CAS Latency With Fuses.

Would you please use DX-71 to help explain how fuses could be used to determine CAS latency.

A. Yes. Depicted in this figure is a -- in the box labeled "circuitry" we have a hard-wired value of 2 and a hard-wired value of 3, and these are both connected to the box labeled "CAS latency circuitry" through wires, and those wires have fuses on them, and if you blow one of those fuses, then that connection is no longer established so that after blowing one of those fuses, only the value 2 or the value 3 would be driving that CAS latency circuitry, so once you blow a

fuse, the part would behave as a de facto fixed latency part.

Q. Now, how would the fuse be blown?

A. It would be blown electrically after packaging.

Q. Is that the only way to blow a fuse?

A. No. No. There are laser-blown fuses as well, but those need to be blown before packaging.

Q. If we could pull up DX-72.

What does DX-72 show?

A. This shows a laser-blown fuse.

Q. And in DX-72, which would be the CAS latency that would operate on this DRAM?

A. In this instance CAS latency 2 would be the value that would be sent to the rest of the circuit; therefore, the circuit would behave as a fixed CAS latency 2 part.

Q. Now, if we could bring up the next demonstrative, which would be demonstrative DX-73.

This has the caption CAS Latency Hardware With Fuses - Electrical Pulse. Could you please explain what's shown in DX-73.

A. Yes. This is showing the blowing of the fuse, that the fuse would be blown using an electrical pulse rather than a laser, so this is something that could be

done after the part is packaged.

Q. And again on DX-73, which is the latency value that would operate on the DRAM?

A. It's the same as before. The value of CAS latency 2 would be driving the circuitry, so therefore the part would be a fixed CAS latency of 2 part.

Q. Now, with respect to the fuses that you've been discussing, do synchronous DRAMs today contain fuses of this sort?

A. Yes, they do.

Q. What's the purpose of using fuses of this sort in synchronous DRAMs today?

A. Most of the fuses are involved in enabling redundant storage elements so that if defects in the storage elements or the storage arrays are found, they can connect the redundant elements in place of the damaged elements.

Q. So in other words, it's a means of rerouting the circuitry to replace a portion of the circuitry?

A. Yes. Absolutely.

Q. Now, the fuses that are used in synchronous DRAMs today, are they laser blown or electrically blown or both or other?

A. They are both. They are -- some manufacturers use laser-blown fuses; other manufacturers use

electrically blown fuses.

Q. Now, compared to using a mode register to program CAS latency, what, if any, would have been the advantages in the 1991 to 1996 time period of using fuses to determine the CAS latency?

A. It would be potentially a simpler design. You would eliminate the mode register. It would be potentially a smaller design and therefore a cheaper design. After blowing the fuse, you would only need to test one CAS latency value instead of having to test all possible CAS latency values, so it would be a cheaper alternative potentially.

Q. Now, again compared with using a mode register to program CAS latency, what would have been the disadvantages to using fuses to determine CAS latency in the 1991 to 1996 time period?

A. Well, again, what you're dealing with is what is -- well, what is in effect a fixed CAS latency part once the fuse is blown, and it can't be used in a system until the fuse is blown, so it's a de facto fixed CAS latency part, so the manufacturers would have to be a little bit more explicit about these things and users would, you know, have to be more savvy if they were going to use these things in their systems. But again, this could be solved with the memory controller

redesign.

Q. Now, the next alternative that you mentioned I believe was scaling CAS latency with clock frequency. You have a demonstrative DX-74 which again reads "Alternatives to Programmable CAS Latency" with item number 3 highlighted.

Can you please explain what you mean by "scale CAS latency with clock frequency"?

A. Yes. Here the DRAM would not be told what latency to use but, rather, what clock frequency to use, and it would determine how many cycles that represented based upon its inherent latency.

Q. Well, I guess, what would actually determine what latency or what frequency to use?

A. Either the memory controller could tell the DRAM explicitly what frequency the bus would be running at or the DRAM could learn that information on its own by having an internal circuit that would sense the bus speed and determine, for example, if it's higher or lower than a reference and therefore it would choose between the two possible values or more.

Q. Could you please explain in a bit more detail how the DRAM could itself determine the operating speed of the bus.

A. Yes. For example, let's say that the part was

meant to work at either 100 megahertz or 150 megahertz. It could have an internal oscillator that would be somewhere in between, expected to run somewhere in between 100 and 150 megahertz, just to pick something, 133 megahertz.

And it would have a simple circuit that would look at the bus frequency, the existing bus frequency, and do an edge detect to see if the bus frequency is faster than or slower than the internal reference. And if it was faster than the internal reference, it would therefore be a 150 megahertz part. If it was slower -- or, rather, it would be a 150 megahertz bus. If the external bus clock would be slower than the internal reference, then it would be a 100 megahertz bus.

Q. And again, just to be clear, I think you said that was one of the two options. The other option would be the memory controller signaling to the DRAM what the bus speed was?

A. Yes. That would be the other -- another option.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to scale CAS latency with clock frequency rather than using a mode register to determine CAS latency in the 1991 to 1996 time period?

A. Well, for example, the part would always operate with the fastest latency possible rather than being told to wait longer than it would be possible.

So if it was capable of producing something with a latency of 2, it would produce something with a latency of 2.

So the systems would always be designed to have the best possible performance.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to scale CAS latency with clock frequency rather than using a mode register to determine CAS latency?

A. Well, again, the manufacturers would need to be more precise about the labeling of these parts and the labeling of their systems and labeling them so that everyone would understand what parts are compatible with what systems, and users would be -- would need to be a bit more savvy about plugging in a DIMM with these types of DRAMs on it into a system and know that the system would actually work with this DIMM.

Q. I believe the next alternative you identified is using an existing pin or a new, dedicated pin to identify the latency.

A. Yes.

Q. DX-75 is a slide also entitled Alternatives to Programmable CAS Latency with item number 4 highlighted.

Can you please explain the alternative of using a pin to identify latency?

A. Yes. The idea is that rather than placing the value in a mode register, you send that same value over a dedicated pin, and so that pin would contain only -- or it would transmit only that information during the lifetime of the DRAM.

While the system is operative, that value would never change, and so the information would be usable by the DRAM exactly as if the information came from a mode register, only it would be coming from the input from the pins instead, so it would be the same information. It would be used in an identical manner, only rather than holding it in the mode register you'd send it over a pin.

Q. Just to be clear with respect to this alternative here, you are talking about using a dedicated pin?

A. Yes. A dedicated pin that would, for instance, transmit only CAS latency information.

So you would have one pin that if it had a low voltage asserted on it, that would mean, for example,

CAS latency 2; if it had a high voltage, it would indicate use CAS latency 3.

Q. Now, would this option have required that additional pins be included in the DRAM?

A. Not in all cases. In many examples there are no-connect pins on DRAMs. There are pins left over after the specification is made that have no function assigned to them, and so these could have been used to transmit this information.

Q. Now, with respect to this alternative, would the signals on the pins change dynamically during the operation of the DRAM?

A. Not in a JEDEC-style system. In a JEDEC-style organization, you set that value at system initialization and it does not change throughout the lifetime of the system while the system is powered on. It's a constant value while the system is running, and so this alternative would be used in that same manner. While the system is running, the memory controller would assert the same value and it would not be changing dynamically.

Q. Now, what, if any, implications would that have for the type of current on the pin?

A. It would be a DC signal. It would not be an AC signal. It would not be changing. It would be

constant and therefore the receiver on the DRAM side would be much simpler.

The pin could be on a part of the package that's not desirable for the faster types of pins, so it would be a cheaper pin to implement. It would be a cheaper pad to locate on the DRAM die. It would be much simpler and much cheaper to add than, for example, adding another data pin or, you know, something that's expected to change rapidly.

Q. Now, in your opinion, what would the advantages of using a dedicated pin to determine CAS latency have been as opposed to using a mode register to determine CAS latency?

A. Well, it would be a simpler design because you would eliminate the mode register as well as the interface required to put information into the mode register, and it would be a smaller design and therefore a cheaper design to manufacture, so it would be simpler and cheaper.

Q. Now, in your opinion, what would have been the disadvantages, if any, had JEDEC chosen to use a dedicated pin to determine CAS latency as opposed to using a mode register?

A. If they had no-connect pins available, there would be no disadvantage. If there were no no-connect

pins available or not enough no-connect pins available, then you would have to add new pins to the package, and that would increase cost. But it would be relatively insignificant.

Q. Why do you say it would be relatively insignificant?

A. Because, again, as I said, these -- this type of interface, a DC-type interface, is much less expensive than adding, for instance, what they call a high-speed pin, a data-type pin.

Q. If we could turn then to your next alternative, explicitly identify CAS latency in the read command, and this will be marked then DX-76, the slide that's also entitled Alternatives to Programmable CAS Latency with number 5 highlighted.

Can you please explain what you mean by "explicitly identify CAS latency in the read command"?

A. By this I mean that you would have multiple commands for multiple latencies, so if you wanted to transmit or have the DRAM use two different latencies, you would have two different commands, one that says read with latency 2, one that indicates to read with latency 3, and so that information would be transmitted across the bus at the time of the command.

Q. Okay. So this command then would be

originating in the memory controller?

A. Yes, it would.

Q. And what pins would be used to transmit this command?

A. Well, for example, there are a number of existing pins that encode control information, for example, the RAS pin, the CAS pin, clock enable, and so forth, DQ mask and write enable, and these together form a de facto control bus with 32 possibilities for their combinations, and currently in the standard not all 32 possibilities are encoded, so you could use some of those unused combinations to encode this additional information.

Q. Can you please explain how five pins could send 32 different commands?

A. Well, it's basic mathematics. If you have five variables each of which can take on two values, you have to 25 combinations, which is 32.

Q. So in other words, one pin can send two possibilities?

A. And the next pin can send 2 different values, so the two pins together can send 4 values, three pins together can send 8 different values, four pins together can send 16 different values, and five pins together can send 32 different values, and so forth.

Q. And I believe you identified the RAS, CAS, clock enable, DQ and then write enable --

A. Yes.

Q. -- as the pins.

Do you have an understanding of how many different commands will currently send synchronous DRAMs across these pins?

A. Yes. Far less than 32. It's on the order of a dozen.

Q. So in other words, there would be sufficient remaining combinations to permit this alternative?

A. Yes, there would.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to explicitly identify CAS latency in the read command rather than using a mode register to program CAS latency?

A. The advantage would be that you would eliminate the mode register and the circuitry required to decode special commands and put that information into the mode register, so it would make the part potentially smaller and simpler.

Q. And would that have had any implication for cost?

A. Yes. That potentially would reduce the cost of the part.

Q. What, if any, would have been the disadvantages had JEDEC chose to explicitly identify CAS latency in the read command rather than using a mode register?

A. The disadvantage would be that it would make the decoding logic on the DRAM more complex because you would have these additional commands that would need to be decoded, so that would make the part more complex, so you'd have a trade-off there.

And if, for example, there were certain combinations that you had to support and you didn't want to redefine, for example, the DQ mask pins in the way I've described, it might require an additional pin.

Q. Focusing on the use of existing pins for the moment, you mentioned that it might require more complex decode circuitry?

A. Yes.

Q. How significant would that be?

A. Not very significant. It would be on the order of the complexity that you're removing by not having to decode the initialization commands.

Q. In other words, on the order of the complexity that would be removed by taking off the mode register?

A. Absolutely.

Q. Now, by the way, this alternative of explicitly identifying CAS latency in the read command, would that require a register on the DRAM?

A. No, it would not.

Q. Would it require some means of receiving information?

A. Yes, it would.

Q. And how would that be done?

A. You could latch the information.

Q. What is a latch?

A. A latch is a circuit that retains its state. It's a way of capturing data and holding on to it.

Q. Now, how, if at all, does a latch differ from a register?

A. A latch is a specific implementation. A register implies how a piece of storage is being used.

So for instance, a register might be built out of latches, but a register is not a latch. A register could be built out of latches or D flip-flops or any number of mechanisms.

Q. If we could turn to the final alternative that you've listed here, and we'll now look at DX-77, which is a slide also entitled Alternatives to Programmable CAS Latency with item number 6 highlighted, stay with asynchronous DRAM (e.g., burst EDO).

Can you please explain what you mean by the alternative of stay with asynchronous DRAM?

A. Yes. Since the latency on a synchronous DRAM, the time which it drives data out onto the bus, is determined by the memory controller and the point of programmable CAS latency is to make parts of different -- potentially different generations compatible with each other, an asynchronous DRAM gives you that compatibility inherently.

Q. Perhaps as just a start, if you could explain your understanding of what you mean by "asynchronous DRAM."

A. This is the term that is used to describe DRAMs prior to SDRAMs, those who are driven off the RAS and CAS signals where the RAS and CAS actually control the operation of the DRAM rather than a clock.

Q. Are there similarities between asynchronous DRAMs and synchronous DRAMs?

A. Oh, very much so. They're very similar. They have -- they do the same thing. They have the same internal circuits.

The primary difference between an asynchronous system and a synchronous system, for example, is in a synchronous system you have a system clock that is driving the memory controller and the DRAMs, in an

asynchronous system the system clock drives the memory controller directly. It does not drive the DRAMs directly but, rather, indirectly through the memory controller so that the clock drives the memory controller and the memory controller then drives the DRAMs through the RAS and CAS timing signals.

Q. Can you please explain how the memory controller would drive the DRAM through the RAS and CAS signals?

A. Well, those are the signals that cause the DRAM to do things. The RAS is the equivalent of a row activate system. The CAS is the equivalent of a, you know, read command or a write command.

Q. Now, focusing again on the early to mid-1990s time period, what, if any, would have been the advantages had JEDEC chosen to continue to develop asynchronous memory rather than using synchronous memory?

A. It would have been a simpler transition because the technology existed at the time. This was a technology that the engineers of the time were more familiar with. Asynchronous DRAM tended to have smaller die sizes like burst EDO at the time had a smaller die size than SDRAM and had better performance at the same speeds.

So asynchronous potentially had better performance and cheaper implementation.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to continue to develop future generations of asynchronous relative than moving to synchronous?

A. I don't see enormous disadvantages. The general -- the general view is that moving to synchronous allows you to scale to higher speeds more easily and so it's a faster upgrade path.

It's a simpler design mechanism for achieving higher rates of speed, but at some point it also tops out, because if you have a synchronous system with a global clock, at some point that global clock mechanism starts getting in the way of data transmission, so at some point you start moving back to mechanisms that look like asynchronous designs as well.

Q. If we could move on now to programmable burst length, and perhaps you could start by briefly explaining your understanding of how programmable burst length is used in the JEDEC SDRAM and DDR SDRAM standards.

A. Yes. Programmable burst length is the ability for a memory controller to tell all of the DRAMs in the

system to use the same -- to produce the same amount of data or to read the same amount of data per transaction, and usually this is chosen to correspond with the size of the cache block in the system.

So that if the cache block in the system is 64 bytes, the memory controller sets the DRAMs to read and write in granularities of 64 bytes, the DIMMs to -- the DIMMs read and write to with a granularity of 64 bytes.

Q. And there's a demonstrative DX-78 that consists of a slide labeled Burst Length.

Now, what, if any, were the alternatives to programmable burst length that existed and were available to JEDEC in the 1991 to 1996 time period?

A. I believe we have a demonstrative for that. Oh, yeah.

For example --

Q. Let me just mention that's DX-79, a slide entitled Alternatives to Programmable Burst Length.

A. So for example, the manufacturers could have used fixed burst length parts.

They could have programmed DRAMs to use a de facto fixed burst length by blowing fuses on the DRAM.

They could have used an existing pin or a new,

dedicated pin to transmit the information, identify the burst length.

They could have extended the command set so that different burst lengths would be transmitted with the read commands so that for two different burst lengths you would have two different read commands.

They could have used a burst terminate mechanism so that the part, for example, would use by default a burst length of eight, but systems that wanted to use a burst length of four would either implicitly or explicitly terminate the burst to a length of four instead of eight.

And lastly, they could have used a CAS pulse to control the data output, which means the DRAM would only return data when it saw a toggle on the CAS pin. So that if the memory controller wanted a length of eight, it would toggle CAS eight times; if it wanted a length of four, it would toggle CAS four times.

Q. Okay. Let's see if we can look at these one at a time.

Now, some of these I think will be fairly similar to the ones we looked at in CAS latency and we'll see if we can move through those a little bit more quickly.

The first one, use fixed burst length parts,

can you please explain briefly what that one entails?

A. Yes. So again for a fixed burst length part, you could either define the burst length to be fixed at the design phase, during the manufacturing phase or during the packaging phase.

And here, we are showing in this diagram the decision being made during the design phase so that the part would only have a value of 4, so this would be a hard-wired value of 4 driving your burst length circuitry, so this would be a design that would only be able to give you a burst length of 4.

Q. For the record, you're referring to a slide that has a caption of Fixed Burst Length and in the top box circuitry burst length 4. That will be marked as DX-80.

Now, is it also possible to determine the burst length in the manufacturing process?

A. Yes, it is.

Q. If we could perhaps pull up the next slide to be marked as DX-81.

A. This is the metal mask option, and the part as designed would have a hard-wired value of 4 and a hard-wired value of 8, but at the time of manufacturing, in one of the last steps of the processing stages you would put down a piece of metal

that either connects -- so you would have two masks and you would use either the one mask or the other mask to put down either one piece of metal connecting burst 4 to the burst length circuitry or you would use a different mask to put down a different wire that would connect the hard-coded value of 8 to the burst length circuitry.

So you would have the design, the design would give you the option of doing 4 or 8, but at the time of manufacturing you would choose only one of those to be connected to your burst length circuitry, so at that point the part would become a de facto burst length part.

In this instance, we're showing that the hard-wired value of 4 is connected to the burst length circuitry, so that's -- that would be the part's behavior. It would have a burst length of 4.

Q. And does it operate in a similar manner to the way you described the metal mask option in CAS latency?

A. Absolutely. The same mechanism.

Q. And would it also be possible to determine the CAS latency at the packaging phase of manufacturing?

A. Yes, it would.

Q. And perhaps we can bring up the next

demonstrative which will be marked as DX-82 and that reads "Fixed Burst Length Packaging Option."

Could you please explain what's shown in DX-82.

A. Yes. This is similar to the packaging option for determining fixed CAS latency, and here you would have a burst length of 4 or a burst length of 8 as options.

You would have a hard-wired value of 4 or a hard-wired value of 8, both being connected indirectly to the burst length circuitry through a multiplexer, which is labeled "mux" here, and the multiplexer would choose between either 4 or 8, would choose either the value 4 or the value 8, choose between the two based upon the control signal which would be connected to either power or ground.

It would be essentially hard-wired to either power or ground, and that decision whether it was hard-wired to power or ground would be made at the packaging time through the use of bond wires that would connect that wire to either the power pin or a ground pin.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use fixed burst length rather than use a mode register to program burst

length?

A. You would have a simpler design because you would have the -- you would eliminate the mode register and eliminate the circuitry used to initialize the mode register, so it would be a smaller part, a cheaper part to manufacture, and depending upon at what point you decide to fix the burst length, it would potentially be a cheaper part to test.

Q. And what, if any, would be the disadvantages had JEDEC chosen to use a fixed burst length rather than a mode register to program burst length?

A. Just as in fixed CAS latency, the manufacturers would need to be explicit about the behavior of these parts, be very clear about the labeling of these parts. Users who would plug these parts into their systems would have to be a little bit more savvy to ensure that they did not put parts with different behavior into the same system because they might not work together well.

Alternatively, you could have redesigned your memory controller to be more intelligent about handling the systems of -- that contain mixed-mode parts. Again -- but again, that depends upon what JEDEC chose. They could have standardized on a burst length of 8 and just stuck with that, in which case you would have no

compatibility issues at all.

Q. If we could turn to the next alternative that you identified, which I believe was use of fuses.

We'll pull up the next demonstrative, which will be marked as DX-83. This slide is captioned Set Burst Length With Fuses.

Can you please explain just very briefly what DX-83 shows?

A. Yes. This is similar to setting the CAS latency with fuses. And the idea is that you have the same burst length circuitry that you would have today and you would have two hard-wired values, for instance, a burst length of 4 or a burst length of 8.

So you would have the hard-wired value of 4 or the hard-wired value of 8, and these values would be connected indirectly to the burst length circuitry through fuses, one of which could be blown, thereby breaking the connection so that after the fuse is blown, only one of these hard-coded values would be connected to the burst length circuitry.

Q. As with the case of CAS latency, could these fuses be blown either with a laser or electrically?

A. Yes.

Q. If we can bring up DX-84, which does that show?

A. This shows blowing the fuse with a laser.

Q. And if we can bring up DX-85?

A. This shows the fuse being blown with an electrical pulse.

Q. Okay. Can you please explain briefly what, if any, would have been the advantages had JEDEC chosen to determine burst length with fuses rather than by programming it in a mode register?

A. The advantage would be the elimination of the mode register and the circuitry required to initialize it, which would make the DRAM potentially smaller, potentially cheaper to implement, cheaper to design, cheaper to manufacture, and potentially cheaper to test.

Q. And again, what would have been the disadvantages, if any, had JEDEC chosen to use fuses to determine burst length rather than programming it through the mode register?

A. Well, this is again similar to having a fixed latency part. Once the fuse is blown, the part becomes a de facto fixed burst length part and so the -- you could potentially have incompatibilities if you put parts with a burst length of 4 into a system with parts that had a burst length of 8, so the manufacturers would need to be more explicit about the

behavior and users might need to be more savvy about their use.

But again, this could be solved by making a memory controller a bit more sophisticated and adaptable and able to deal with mixed-mode parts.

Q. Now, I believe the next alternative that you identified was using pins to set the burst length.

And again, in this option I think that you're referring to using a dedicated pin; is that right?

A. Yes.

Q. Just for the record, we've brought DX-69 back up again.

Can you please explain, again very briefly, how use of a -- or how a dedicated pin could be used to determine burst length?

A. This would be similar to the mechanism used to define different CAS latencies. Rather than storing a value indicating burst length in a mode register, you would send the same value over a dedicated pin that would not change while the system is running or while the DRAM is operative.

So just as the value can be taken out of a mode register, a DRAM could take the same value off of a pin and use that to select between different burst length circuitry.

Q. Now, when you were discussing use of a pin, a dedicated pin, to determine CAS latency, I believe you described certain attributes, such as DC power, et cetera.

Would those attributes also apply to the pin that you'd have in mind to determine burst length?

A. Absolutely. The same -- the same conditions apply.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use a dedicated pin to determine burst length rather than using programming through a mode register?

A. Well, again you would eliminate the mode register and the circuitry required to initialize it, which would make the part potentially smaller, cheaper to manufacture, potentially cheaper to test, easier to design.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to use a dedicated pin to determine burst length rather than programming burst length through the mode register?

A. If there existed a no-connect pin available to transmit this information, there would be no disadvantage. If the -- if the specification did not have enough unused pins, then you would have to add a

new pin to the package, and that would increase cost.

Q. How significant would the cost increases have been had it been necessary to add an additional pin?

A. Not -- as I said before, not as significant as adding a data pin because this would be a signal that would not be changing over -- it would not be changing dynamically, so it would be a DC value, it would be a simpler receiver, the pin could be in an undesirable location on the package, the pad could be in an undesirable location on the DRAM die, and much simpler to add this.

Q. Now, if we can turn next to the fourth alternative you've identified, explicitly identify burst length in the read command, can you please explain briefly what that means?

A. This is similar to the method of identifying the CAS latency. You would encode the information for burst length in the read command using the control bus, redefining the definitions of those pins I mentioned earlier so that you would transmit several different read commands to identify several different burst lengths.

Q. And this would use the same five pins that you identified earlier in connection with CAS latency?

A. Yes, it would.

Q. Are there sufficient commands available to permit explicitly identifying both the CAS latency and the burst length in the read command?

A. Yes, there would be.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to identify burst length in the read command rather than programming it through a mode register?

A. Well, again, you would get rid of the mode register and therefore the circuitry required to initialize it, which would make the part simpler to design and test and potentially cheaper to manufacture.

Q. What, if any, would have been the disadvantages had JEDEC chosen to identify burst length in the read command rather than using a programmable mode register?

A. Well, you would increase the complexity of the decoding logic so that the part would have to recognize more, more commands.

Q. Now, the next alternative that you have identified is use a burst terminate command. Perhaps we could start simply by explaining briefly what is a burst terminate command.

A. The way that SDRAMs are defined currently, if

you have a long burst, let's say the part is defined to work with a burst length of 8, it's been programmed to work with a burst length of 8, and you send a read command at one instance of time and then follow that up with another read command four cycles later, the DRAM is going to implicitly interrupt itself.

The first command will only transmit four beats of data back and then the DRAM will start handling the second command, so the DRAM implicitly terminates the first burst, interrupts it with the second burst, so you -- what you would get is a de facto burst length of 4 in that instance.

Q. And then can you please explain what you had in mind with this alternative of using a burst terminate command as an alternative to programmable burst length?

A. Yes. For example, if JEDEC wanted to support two burst lengths, such as 4 and 8, it could choose to define all parts to use a fixed burst length with 8 which would, as I've mentioned, make the part simpler to design and it would use the burst terminate feature, which was already in the spec, to terminate those bursts to a length of 4 for any OEM or any system designer who desired a burst length of 4 rather than a burst length of 8.

So that if a memory controller only wanted four pieces of data, four beats of data to return on a read, it would either send a following read command four units of time later or it would explicitly terminate the burst if it had no other pending requests.

Q. Now, the burst terminate command would originate at the memory controller?

A. Yes, it would.

Q. And would it be sent to the DRAM via a pin?

A. Yes, it would.

Q. Now, would this alternative require use of any additional pins?

A. No, it would not because this is already in the specification.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use a long fixed burst length together with a burst terminate command rather than programmable burst length through the mode register?

A. You would have had a simpler part because you'd have no mode register, no -- the part would not need to have to determine -- behave with several different burst lengths, and so it would be a simpler part to design and test and manufacture.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to use a long burst length with a burst terminate command rather than programming burst length through the mode register?

A. You potentially could run into inefficiencies on the bus depending upon how you -- depending upon how the memory controller handles those situations where you want to terminate the burst down to 4 from 8.

Q. How significant a disadvantage would that have been?

A. I don't believe it would have been very significant.

Q. And then finally you identified use CAS pulse to control data output. Can you please explain what you have in mind with that alternative?

A. Yes. The way that burst length is defined currently in the standard, you -- all you need to do is toggle CAS once to get, for example, four beats of data back or eight beats of data back, and while data is coming back on the DRAM bus, the CAS is not toggling. It's sitting there idle.

Rather than having that scenario, they could have used CAS to explicitly control the driving of data onto the bus or the receiving of data off the bus at the DRAM side. So that if the memory controller wanted

four beats of data, then it would toggle CAS four times; if the memory controller wanted to send or receive eight beats of data, then it would toggle CAS eight times.

Q. So in other words, the DRAM would return data only in response to a toggle of the CAS signal?

A. Yes. Correct.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to control burst length through the CAS pulse as opposed to programming it in the mode register?

A. It would have been a simpler design. You would have eliminated the mode register and the circuitry required to initialize it, which would make the part simpler, smaller, easier to test, for example.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to control burst length through the CAS pulse rather than programming it in the mode register?

A. I don't see any significant disadvantages.

MR. OLIVER: Your Honor, I'm about to turn to the alternatives to dual-edge clocking. We're probably about halfway through the alternatives. Would this be an appropriate place to --

JUDGE McGUIRE: I think it would be. Let's take a ten-minute break.

Off the record.

(Recess)

JUDGE McGUIRE: Mr. Oliver, you may proceed.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Professor Jacob, if we could turn next to the alternatives for dual-edge clocking, and first if I could ask you to please explain your understanding of how dual-edge clocking works.

A. Yes. Dual-edge clocking is the transmission and receiving of data on both edges of the clock, meaning both edges of the clock cause data to be latched or sent via DRAM to the memory controller and the same thing at the memory controller side, so rather than sending or receiving one piece of data per clock cycle, you're sending two pieces of data per clock cycle.

Q. Let me state for the record that we're showing DX-86, which is a slide that simply reads "Dual-Edge Clocking."

And Professor Jacob, I believe you have another demonstrative that we'll call DX-87 that helps explain dual-edge clocking?

A. Yes, I do.

Q. DX-87 will be a slide marked Single-Edged Clocking versus Dual-Edged Clocking?

A. Yes.

Q. Can you please explain what this slide indicates?

A. Well, this is showing you the difference between using a single-edged clocking scheme versus a dual-edged clocking scheme. And in both systems you have data transmission that's at about the same rate, but what's shown is that in the top we have a single-edged clocking scheme, which means that the clock transitions from low to high at twice the rate of the data.

The data only goes from low to high, say, once per cycle, where the cycle is the clock going from high to low and then back again, so the data can only transition from zero to one or one to zero once per clock cycle, whereas the clock goes from zero to one -- makes two transitions, goes from zero to one and one to zero twice per clock cycle, so the clock is transitioning at twice the speed of the data.

The advantage here is that you have, as shown in the middle there, you have two clock edges per data, per beat of data. You have a clock edge that coincides

with the driving of the data and you have a clock edge that corresponds to the receiving of the data, the vertical lines showing that one edge of the clock is synchronous with the edge of the data and one is center aligned with the data.

On the bottom, you've got a dual-edged clocking scheme so that the clock transitions at the same rate as the data, so if the data transitions, say, 100 times per second, let's say a hundred megabits per second, the clock is going -- is essentially a 50 megahertz clock, so the difference here is that you only have one clock edge per beat of data.

So instead of having an edge at the beginning of the data and an edge at the middle of the data, you only have a clock edge at the beginning of the data in this example.

Q. Now, in your opinion, did JEDEC have available to it alternatives to dual-edged clocking as that concept was incorporated in the DDR SDRAM standard?

A. Yes, I do.

Q. Could you please explain what alternatives to dual-edged clocking were available to JEDEC.

A. Yes. Again it would help me if we have a demonstrative.

Q. Let me just note for the record that this will

be DX-88. It's a slide entitled Alternatives to Dual-Edged Clocking.

Sorry to interrupt. Could you please explain the alternatives that you identified.

A. Yes. So the idea of dual-edged clocking is to increase the data width without having to, for example -- doubling the data, the data bandwidth, without having commensurate increase in the clock speed and/or the clock energy, the power dissipated by the clock.

So alternatives would be, one, use interleaving, so you would interleave memory banks, two or more memory banks, on the DRAM chip itself, so use the concept of interleaving on the DRAM chip.

The second is to use the concept of interleaving at the module level, so you would have separate chips and interleave between the separate chips on the module and keep a single-edged clocking scheme in both of these scenarios. You'd have a single-edged clocking scheme.

And number three, they could have achieved higher bandwidth. Rather than by moving to higher speeds, they could have achieved higher bandwidth by using more pins.

So for example, this would give you the same

bandwidth, this doubling the clock rate, but you wouldn't have to double the clock rate. It's a slower interface in terms of toggling, but you know, it gives you the same increase in bandwidth.

The fourth alternative is to apply that same thinking at the module level, so don't change the width, the number of pins per DRAM, but, rather, change the number of connectors on the module, have a wider memory bus on the motherboard, so you would get your bandwidth that way.

The fifth alternative is to have the data rate increase by, you know, whatever factor and also have the clock increase at that same level, so use a single-edged clock instead of a dual-edged clocking scheme.

The sixth alternative would be to use simultaneous bidirectional I/O, which is a signaling scheme that allows you to send data in both directions at the same time, between a memory controller and a DRAM so that, for example, you can be processing reads simultaneously with writes as opposed to the present scheme where one has to wait for the other to finish.

And last but not least, they could have used the toggle mode DRAM, for example, IBM's toggle mode DRAM.

Q. Do you regard toggle mode as an alternative to dual-edged clocking?

A. I believe -- I personally believe toggle mode is a dual-edged clocking scheme, but it is listed here sort of for completeness because Rambus apparently believes that toggle mode is distinct from their implementation of dual-edged clocking, it's not the same thing as dual-edged clocking, so I've listed it here as an alternative.

Q. If we could focus first on use two or more interleaved memory banks on-chip, could you please explain in a little more detail how that would work.

A. Yes. Interleaving is a mechanism that has been around for quite a while, since the '60s, and the idea is to go back and forth between two banks.

Here we go. We have a demonstrative showing this.

The idea is that you go back and forth between two banks and so you delay the commands sent to the second bank so that you -- if bank 1 can give you data every ten nanoseconds, for example, and bank 2 can give you data every ten nanoseconds, but you delay the operation of bank 2 by five nanoseconds and you go back and forth between bank 1, you say read, read, read, read back and forth, so read to bank 1, read to bank 2,

read to bank 1, read to bank 2, every five nanoseconds, then in effect you're getting data every five nanoseconds, whereas each bank is only giving you data every ten nanoseconds, so you are doubling the bandwidth of the part without having to double the bandwidth of a single bank.

And as this -- the illustration on the left shows, here's one implementation where on the DRAM chip you would have two banks and you would send them both the same clock signal, you would send them both the same command signal, only to the second bank on the chip you would insert the delay of half a cycle.

Q. Let the record reflect that the witness is referring to a slide entitled Interleave On-Chip Banks that will be marked as DX-89.

Now, Professor Jacob, there are two separate diagrams on DX-89. What's represented by the two separate diagrams?

A. Well, the second diagram is another implementation of essentially the same thing where your interleaving the banks is also on the DRAM chip, but the way that you get the delayed clock signal would be, rather than having the DRAM being responsible for delaying the clock signal to the second bank, the memory controller would be

responsible for delaying the clock to the second bank, but what this would require would be two clock signals sent to the DRAM chip, so this would double the number of clock signals to the DRAM chip between the memory controller and the DRAM, so it would be doubled from one to two. So each bank would have its own separate clock.

Q. Apart from that difference, is there any fundamental distinction between the way that the diagram on the left side of DX-89 would operate as opposed to the way the diagram on the right side of DX-89 would operate?

A. No. They would be the same.

Q. Now, can you please explain what, if any, would have been the advantages had JEDEC chosen to interleave on-chip banks rather than using dual-edged clocking?

A. Well, you would have -- rather than having a dual-edged clocking scheme, you'd have a single-edged clocking scheme, which means that your design of the clock is significantly simpler.

In a dual-edged clocking scheme you have data transitioning on both edges of the clock, and so to have data being sent across the bus at a regular rate, then your clock needs to be very symmetric. Your duty

cycle needs to be very close to 50 percent and your rise time and fall time of the clock needs to be -- those two need to be very symmetric. They need to be equal. So you have to have a very, very symmetric clock signal, a very regular shape, a very even shape.

Whereas for a single-edged clocking scheme, which is what you could use in both of these examples, for a single-edged clocking scheme, your edge rates, meaning the rise time and the fall time, need not be the same and your duty cycle need not be 50 percent. And you still get, you know, very good behavior.

So it's simpler to design a single-edged clock than it is to design a clock with a dual-edged clocking scheme.

Q. I believe you have another demonstrative that may help to explain that.

A. Oh, okay.

Q. If we could bring up the slide entitled Duty Cycle and Slew Rates.

A. Thank you.

Q. And we'll label this as DX-90.

Now, you referred in your answer to duty cycle?

A. Yes.

Q. Can you please explain with reference to DX-90

if it helps what you mean by "duty cycle"?

A. Yes. The duty cycle represents the proportion of time that the clock is high versus the proportion of time that the clock signal is low, so if you have a duty cycle that is 50 percent, that means the clock is -- 50 percent of the time it's high, 50 percent of the time the clock is low.

And for a dual-edged clocking scheme, you need to have something that's very close to 50 percent because otherwise that means that every other piece of data is going to have a shorter window and that makes it harder to design the system.

So that's what's shown in the top figure there. For a dual-edged clocking scheme you have to have the symmetric duty cycle, you have to have something that's close to 50 percent, and the slopes that go up and the slopes that go down need to be rising and falling at the same time for your system to behave well.

Whereas in the bottom figure we're showing the type of clock that could be used to implement a single-edged clocking scheme, and notice that the clock is not nearly as symmetric as the one on top.

For example, we're showing here that the rise time of the clock is relatively fast and the fall time of the clock is relatively slow. That's perfectly

acceptable for a single-edged clocking scheme.

And note that it also shows an asymmetric duty cycle where it's -- the clock period or the clock signal is high for longer than it is low. That's also acceptable in a single-edged clocking scheme, whereas, again, it's -- that type of behavior is much less desirable if you've got a dual-edged clocking scheme.

So the top signal is harder to produce; the bottom signal is easy to produce.

Q. Professor Jacob, you included an awful lot in that answer.

A. Yeah. I'm sorry.

Q. I'd like to be certain that we keep the record clear on this.

Your demonstrative DX-90 refers to slew rates.

Could you please explain what is meant by "slew rates."

A. Slew rate corresponds to the rise time and the fall time. That's the slopes of those wires that go up and down.

Q. So in other words, the duty cycle is the amount of time above 50 percent -- the amount of time high versus the amount of time low?

A. Yes.

Q. And the slew rates are the rise slope --

A. The rise time and the fall time.

Q. The rise time versus the fall time?

A. Yes.

Q. Now, when you were explaining the potential advantages of interleaving on-chip banks through use of a single-edged clock, you made reference to duty cycle and slew rates.

Could you please explain how the concept of duty cycle and slew rates that you just explained would translate into an advantage for interleaving banks on-chip rather than using a dual-edged clock.

A. Yes. In an interleaved scheme you can use the type of clock that is depicted in the bottom of this figure. You can have a clock that is -- that has asymmetric duty cycle and slew rates. It need not be as clean, pristine, as the clock in the top. You can have -- you can have asymmetric duty cycles and slew rates, whereas that type of behavior is not acceptable in a dual-edged clocking scheme.

Q. And why is it an advantage to be able to use the type of clock illustrated in the bottom portion of DX-90?

A. Because it is much easier to build a clock generation scheme. It's a far simpler circuit, easier to build, easier to test, than it is to build a circuit

that produces a perfectly symmetric scheme.

Q. Now, returning to the alternative of interleaving on-chip banks, what, if any, would have been the disadvantages had JEDEC chosen to interleave banks on a chip rather than using dual-edged clocking?

A. I don't think there would have been much of a disadvantage. You would have to have the delayed clock and potentially an extra wire if you were going to transmit two clocks, but -- so perhaps an increase in the cost.

Q. So in other words, with reference to DX-89, if you were to do the implementation depicted on the right-hand side, would that might require the extra clock wire?

A. Yes. It would require an extra pin on the DRAM and an extra wire.

And on the left-hand side it would require an extra delay element on the DRAM.

Q. And how significant would those disadvantages have been?

A. Not very.

Q. Okay. If we can turn to your next alternative, if we can go back to the slide DX-88, which was the slide entitled Alternatives to Dual-Edged Clocking.

Number 2 here is interleave banks on the module. Can you please explain what you mean by that alternative?

A. Yes. This is similar to interleaving the banks on-chip except that rather than interleaving banks that are on-chip, you would have individual banks on the chip, one bank on-chip, for example, but you would interleave different chips on the memory module so that on your memory module you would have one logical bank -- we'll call it a rank -- and you would have another logical bank and you would interleave back and forth between, say, bank 1 or, rather, rank 1 and rank 2 on that module in a similar fashion with, you know, two clock signals.

Q. Do you have a demonstrative that helps illustrate this?

A. Yes, I do.

Q. We can now pull up the demonstrative with the title Interleave On-Module Ranks and we'll label this as DX-91.

Now, can you please explain what is depicted in DX-91?

A. Yes. This depicts the scenario that I just described where you have two logical ranks of memory on your module, so each of those white boxes that are

pointed to by rank 1 and rank 2, four of those white boxes, each of those white boxes represent a DRAM chip, so this would be a module that has four DRAM chips on it logically partitioned into rank 1 and rank 2, and you would interleave between rank 1 and rank 2 so that the same command would be sent to each of those ranks, but the clock signal would be delayed to rank 2 with respect to rank 1 so that rank 2 would operate just a fraction of a second behind rank 1, therefore toggling back and forth between rank 1 and rank 2 when you were handling data.

And as shown in this picture, the delay mechanism could be on the module, or as mentioned in the bottom of the figure, you could have two clock signals where the delay mechanism is on the memory controller. That would require an extra connector on the module and an extra wire connecting the memory controller to the module.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to interleave on-module ranks rather than use dual-edged clocking?

A. It simplifies the design of the DRAM. You don't have to move to higher bandwidths coming off the DRAM itself and you achieve that bandwidth at the module level instead.

Q. And what, if any, would have been the disadvantages had JEDEC chosen to use interleaving of on-module ranks rather than dual-edged clocking?

A. Well, it pushes the complexity out of the DRAM onto the shoulders of the module designer so that now your module, for instance, would have to have an extra clock line or would have to have that delay element that's pictured in the figure, so it would complicate the module slightly.

Q. By the way, going back to the advantages of using interleaving on-module ranks, would this alternative also permit use of a clock signal that does not have a close to perfect duty cycle or slew rate?

A. Yes. Yes. This is similar to the interleaving on-chip banks where you would be able to use one of those -- a single-edged clocking scheme where you would not have to have a perfect 50 percent duty cycle and you would not have to have matching rise times and fall times and you would not have to have as perfect a clock signal as you would have to have if you did a dual-edged clocking scheme.

Q. If we could pull back up DX-88, which was the slide entitled Alternatives to Dual-Edged Clocking.

And if we could turn next to the third of the

alternatives that you've identified, increase the number of pins per DRAM.

Can you please explain how increasing the number of pins per DRAM could increase the speed of transferred data?

A. Yes. The idea of dual-edged clocking is to, for example, double the bandwidth of the part by doubling the pin bandwidth, meaning having the data pin transition at twice the rate previously without having to increase the speed of the clock.

Alternatively, you could get the same bandwidth out of the part by simply doubling the number of data pins and going with a single-edged clocking scheme and you would not have to increase the clock speed at all. You would also not increase the data pin speed at all. You would simply increase the number of data pins. You would achieve the same bandwidth without having to increase the speed.

Q. In order to double the rate of data transfer, would it be necessary to double the total number of pins in the DRAM?

A. No, it would not.

Q. What pins or what number of pins would have to be increased?

A. You would have to -- for example, if you wanted

to double the bandwidth of the part, you would have to double the number of data pins and then you would have to increase the number of power and ground pins to support the new data pins that were added.

Q. What about command pins?

A. No. The number of command pins and address pins would remain the same.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to increase the number of pins per DRAM rather than using a dual-edged clock?

A. Again, you could retain the use of the single-edged clocking scheme, which means that you could use a far simpler clock circuit design. It would mean that your signals are transitioning at a slower rate than, for instance, a DDR-type interface, so rather than having a 200-megabit-per-second data pin, now you stick with a 100-megabit-per-second data pin, so the power of the DRAM actually goes down comparatively.

Let's see. That's it.

Q. And what, if any, would have been the disadvantages had JEDEC chosen to increase the number of pins rather than using a dual-edged clock?

A. Well, it would be the cost of increasing the number of pins, the number of pads and receivers on the

DRAM, so it would increase the size of the DRAM die. It would increase the size of the package and therefore the cost of the die and the cost of the package. And it would potentially increase the noise in the DRAM, but you would have more pins to spread that noise around, so it would be a net balance.

Q. Now, the next alternative that you identify is increase the number of pins per module. Can you please explain what that concept is?

A. That is the same general idea as the previous one, except now you push the bandwidth issue out to the module level so that you don't increase the DRAM's bandwidth at all, you don't increase the number of pins on the DRAM, you don't increase the speed of the DRAM, so you can use the same width of the DRAM as before in a single-edged clocking scheme, but you have a wider memory bus between the memory controller and the module.

Q. Now, would this require more DRAMs to be mounted on the module?

A. No, it would not.

Q. Now, would this alternative have required any differences in the design of the DRAMs?

A. No, it would not.

Q. What, if any, would have been the advantages

had JEDEC chosen to increase the number of pins per module rather than using a dual-edged clock?

A. It's a very simple mechanism as far as the DRAM is concerned because you don't have to use a dual-edged clock, you don't have to have the DRAM pins transitioning at twice the rate, so it's a far simpler design, it's a far cheaper design to build and test. And -- so yes, the noise and the energy levels are reduced compared to a DDR part.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to increase the number of pins per module rather than using a dual-edged clock?

A. The -- well, now, the onus of producing bandwidth is pushed out to the module level which has implications for your motherboard design. Your module now has to have a larger number of data pins connected to your memory controller.

Currently, you've got a lot of connectors on the module already. You've got -- modules already have two ranks of DIMM -- two ranks of memory on them, so you've got, you know, typically several hundred pins, 288 pins, connectors on it, so you would just use -- basically, instead of having a two-sided DIMM, you would create a big fat SIMM and gang together this

thing and have 128 bits of data coming off of the module, and this would require 128 wires on your motherboard and 128 data pins on your memory controller.

So it would increase the cost of your memory controller and your motherboard but not necessarily the cost of the memory module.

Q. If we can turn to the next alternative, double the clock frequency.

Now, what do you mean by "double the clock frequency"?

A. That means using a single-edged clocking scheme and simply doubling that clock, at least for read commands, and doubling the data bandwidth as well.

Q. Can you explain in a bit more detail how this alternative would work?

A. Yes. It would be similar to a DDR system of today except that the data strobe, for example, would run at twice the frequency that it runs at now.

Q. This example, would this require any increase in the speed or frequency of command signals?

A. No, it would not. You would still send your command and addresses at the same rates that you send them now. The only thing that would change would be

the data rate and it would be similar to or, rather -- I'm sorry.

The only thing that would change would be the clock frequency that accompanies your data transmission. Your data rates would be the same as in DDR parts of today, but your clock frequency would be twice what it is.

Q. Now, what, if any, would have been the advantages of running a single-edged clock at twice the frequency rather than using a dual-edged clock?

A. The advantages include the fact that you have the single-edged clock versus a dual-edged clock, meaning that the edge rates need not be symmetric, the duty cycle need not be 50 percent, and it gives you those extra edges per data packet that are not present in a dual-edged clocking scheme, you have an edge to transmit data as well as another edge to receive data, whereas in a dual-edged clocking scheme you only have an edge to drive data or you have an edge to receive data, but you don't have both.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to double the frequency of a single-edged clock rather than using a dual-edged clock?

A. You have a clock signal that is transitioning

at twice the rate of present, of present systems, which means it would be burning twice as much power as present systems.

Q. Turning to the next alternative, you've listed use simultaneous bidirectional I/O drivers.

Can you explain, first of all, what is an I/O driver?

A. This is the circuit on the DRAM that is responsible for reading and writing data on the data bus.

Q. What is a simultaneous bidirectional I/O driver?

A. This is a signaling scheme that has been around for a while. It's a scheme that allows a reader and a writer to exist on the bus at the same time where the reader and writer are communicating, for example, a memory controller and an individual DRAM, they can be sending data to each other simultaneously and that would not represent a bus conflict, whereas in the signaling protocols that we use today in DRAM systems, if two -- if the memory controller and the DRAM tried to write to the bus at the same time, that would represent a bus conflict.

Q. Can you please explain how using a simultaneous bidirectional I/O driver could increase the effective

rate of data transfer.

A. Well, there are several reasons why it would be beneficial.

One, it allows you to perform reads and writes at the same time, simultaneously. In traditional systems, writes have to wait for reads to finish and reads have to wait for writes to finish, but here you can be processing both simultaneously on the bus.

It also makes your bus usage more efficient by allowing you to more naturally pipeline the data. So if you don't interleave reads and writes on the bus, you get a very efficient pipelining effect, whereas if you have to interleave reads and writes, this introduces dead cycles on the bus where the bus can't be used for anything. So this would eliminate all those inefficiencies.

And this would also -- this scheme thirdly would eliminate instances where reads would have to stall waiting for previous writes, which currently slows down your reads.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use simultaneous bidirectional I/O drivers to increase data transfer rates rather than dual-edged clocking?

A. Well, it would allow you to increase the

bandwidth of the system, as I just mentioned, without having to increase the speed of the system, so this would not increase the power consumption of the system. It would not increase the power consumption of the clock or power dissipation of the clock.

Q. What, if any, would have been the disadvantages had JEDEC chosen to use simultaneous bidirectional I/O drivers rather than dual-edged clocking?

A. Well, for example, it's a more complex driver design, so they would have had to use a more complex mechanism.

The other disadvantage is a potential disadvantage because the benefit that you would receive from using this is dependent upon your application, so for instance, if an application never performed any writes whatsoever and all it ever did was read data, it would receive no benefit from this at all because, well, you can't perform writes at the same time as reads, so you would receive no benefit.

Q. So would it be fair to say then that the potential advantages of simultaneous bidirectional I/O drivers depend upon the usage to which the system is put?

A. Yes. Yes. Yes. I should have mentioned that.

The advantages are limited by your applications, but most applications have a healthy balance of reads and writes.

Q. And then again, the seventh item is toggle mode, and just so the record is clear, did you regard that as an alternative to dual-edged clocking?

A. I believe that toggle mode is a dual-edged clocking scheme and it is included here more for completeness.

Q. If we could turn next then to alternatives to on-chip PLL and DLL, and if -- actually let me just mention for the record that we have pulled up the next demonstrative, which will be DX-92, which is simply a slide reading "PLL/DLL."

Perhaps if I could start by asking you to please explain the problem of aligning data to the system clock on a read operation.

A. I believe we have a demonstrative.

Q. Okay. We've pulled up a slide reading "The Problem" which we've marked as DX-93.

Could you please, if it helps you to do so, use DX-93 to explain the problem of aligning data to the system clock in read operations.

A. Yes. The problem is that when you are dealing with high rates of speed, the signal propagation time

in the system becomes significant, whereas perhaps at lower speeds that system we were pressing using -- at lower speeds these propagation times appeared to be instantaneous and therefore negligible. As you move to higher speeds, suddenly propagation times can no longer be ignored. And this illustrates what's going on (indicating).

For example, what we have illustrated here is a memory controller talking to two modules in the system and you have several different components of delay through the system and these correspond to elements of uncertainty in the timing of the system.

So what we have here is an illustration of read operations where the memory controller is sending, for example, a read to module A and also sending perhaps at a different time a read to module B, and this shows that if you send the request to module A, let's say, at time twelve o'clock, the -- so this is shown by that green clock signal at the top left-hand corner of the memory controller. This shows that you send the request at time twelve o'clock. It arrives at module A, let's say, five minutes later. That same request would arrive at module B ten minutes later, so five minutes after module A.

Then there's some internal delays that are due

to the fact that signals have to propagate through the DRAM chip and be amplified within the DRAM chip, and that's represented by the black circuits on the DRAMs (indicating).

And so by the time that module A or the DRAM on module A would be sending the corresponding result back to the memory controller, it would be -- actually those clocks are meant to be different. I'm sorry. Those two black clocks, one labeled "a" and one labeled "b", should be swapped. I'm sorry. This would be my fault.

So there's a delay through the chip so that at time quarter after, module A would be sending a result back out onto the bus. At time twenty after, module B, the DRAM on module B, would be sending the result back out onto the data bus.

And so the time that the data would arrive at the memory controller would be five -- let's say five minutes later from module A and five minutes later from module B so that the data would finally arrive at the memory controller from module A -- the data arriving from module A would be somewhere in the neighborhood of quarter after and it would arrive at the memory controller somewhere in the order of twenty-five after for module B.

So all of these delays introduce timing uncertainties, and what this means is that you have to have a relatively wide window of time that you're looking for the arrival of the data at the memory controller.

And corresponding to the older, slower rates of speed, this would be, for example, this would correspond to, let's say, a window of time that would represent an hour or two hours, and so the fact that the data would arrive back at the memory controller, you know, in a fifteen-minute window or ten-minute window, depending upon which module you're talking about, that would be insignificant relative to this two-hour time period or even this one-hour time period.

But as you go to higher rates of speed, now you're talking about thirty-minute windows and twenty-minute windows and fifteen-minute windows and smaller windows of time, and now the difference at which that window of time that differs depending upon which module you're talking about, suddenly that window of time becomes very significant.

Q. Now, the various delays that you've described, are they static or do they vary?

A. They vary, depending upon -- some depend upon

the process variations. Some -- and all of them depend at least to some extent on the current temperature of the system and voltage fluctuations in the system.

Q. And then the idea that they vary, is that true of the delays in the return path as well as the internal clock?

A. Yes. This corresponds to the wire delays, the propagation delays, as well as the internal delays to the DRAM.

Q. Now, is it necessary for operation of a system that 100 percent of the delay be corrected for?

A. No. You only need to correct enough of this timing problem so that all of the data arrives at the memory controller within whatever window you care about.

So if you have a fifteen-minute window that you care about, these signals need to be synchronized to within fifteen minutes of each other; if you care about a five-minute window, then you need to synchronize these to within five minutes, and so forth.

Q. Now, within the context of the problem that you've just explained, can you explain what the function of an on-chip DLL in a JEDEC-compliant

DDR SDRAM is?

A. What that does is it cancels out the internal -- or it rectifies the internal delays so that the DRAMs appear to have less of an on-chip delay.

Well, less timing differential. Not less of a delay. Less of a -- less skew, less differences in timing.

Q. By the way, are you familiar with a phase-locked loop or a PLL?

A. Yes, I am.

Q. Can you briefly compare a phase-locked loop or PLL with a delay-locked loop or DLL?

A. They are very similar. The primary difference is that a PLL contains an oscillator and a DLL doesn't.

Q. Now, in your opinion, is an on-chip PLL or DLL necessary for high-speed DRAMs?

A. No, it is not.

Q. Why not?

A. Because all that is necessary is that you cancel out some of these timing uncertainties and there are numerous mechanisms that can do that just as effectively as an on-chip PLL or DLL.

Q. Now, in your opinion, did JEDEC have available

to it other options for canceling out portions of the delay that you've described at the time that it was working on what became the DDR SDRAM standard?

A. Yes, they did.

Q. Can you please identify what some of those alternatives were?

A. Yes. I believe we have a demonstrative that would list them.

Q. If we can pull up the demonstrative entitled Alternatives to On-Chip PLL/DLL. We'll mark this as DX-94.

A. For example, they -- JEDEC could have decided to put a DLL on the memory controller.

JEDEC could have decided to put a DLL on the module.

They could have used a vernier method to account for skew, which is a static timing mechanism. It's not a dynamic one the way a DLL is.

They could have achieved a higher bandwidth using more pins, either at the DRAM level or at the module level, rather than trying to increase clock frequency, because as I just showed, the reason that these problems occurred is because you were trying to increase your clock frequency and your data frequency, and so if you can achieve higher bandwidth without

increasing your clock and data frequency, then the DLL becomes less of an issue.

Lastly, you could dispense with the DLL and rely upon source synchronous data strobes such as the DQS data strobe to provide timing.

Q. Okay. If we could look at these one at a time, starting with the first one, what do you mean by putting the DLL on the memory controller?

A. Well, initially, the idea is that you're trying to cancel out uncertainties in the timing of things and eliminate some of the perceived timing delays on the DRAM by having the DRAM itself decide how in sync it is with the rest of the system by looking at the global clock, but if what you're trying to do is make sure that all of the data arrives back at the memory controller at the same time regardless of which DIMM is producing the data or which DRAM in the system is producing the data, then who better than the memory controller to decide how much to delay that data, rather, delay the production of that data.

So the memory controller here would make sure that all of the DRAMs are in sync with each other rather than each DRAM doing it on its own.

Q. Well, how would a memory controller use a DLL to adjust for the arrival of the data?

A. For example, you could have two clocks, one that drives the control, the control information and the latching of the control data to the DRAMs, and another that drives the data portion of the interface, so one that's latching control on addresses and the other that tells the DRAM to begin driving data out onto the bus. And the memory controller would maintain the amount of phase difference between these two clocks so that for the receipt of data all of it would arrive in sync.

Q. Now, I believe in the previous demonstrative you had identified outbound delay, internal delay and return delay.

Which, if any, of those delay elements would be addressed by placing the DLL on the memory controller?

A. Every single one of them.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to place a DLL on the memory controller rather than using on-chip DLL?

A. Well, for one, you would eliminate the on-chip DLL, which would reduce the power consumption of the DRAM. It would reduce the die size of the DRAM, which would reduce the manufacturing cost of the DRAM. You would reduce the testing costs of the DRAM because you

don't have this PLL or, rather, this DLL that would be part of the DRAM.

It would be a simpler design because it would not include a DLL and therefore cheaper, take less time. And it would cancel out more timing uncertainty than simply putting the DLL out on the DRAM itself, so you could potentially reach higher rates of speed than just using an on-chip DLL alone.

Q. Now, what, if any, would have been the disadvantages had JEDEC chosen to place DLL on the memory controller rather than on the DRAM?

A. Well, for example, it could require extra clocks, as I described, one for the command and one for the data, which would increase -- potentially increase the number of pins on the DRAM or the number of pins on the module and -- or connectors on the module and the number of wires on the bus.

It would increase the design complexity of the memory controller because now the memory controller would have a DLL on board or as part of its design, so the memory controller design would be more complex and therefore more costly, and you would have potentially more pads on the memory controller, more pins on the memory controller package.

Q. Now, taking it in totality, how significant

would those disadvantages have been?

A. Some of them would have been relatively significant. Well, the advantages would have been significant, the reduction in power, but then you had the corresponding increase in power on the memory controller level. You simplify the design of the DRAMs, but you increase the design complexity of the memory controller. I think it would probably balance out.

Q. Now, the second alternative you have listed here is use off-chip or on-module DLLs. Can you please explain what you mean by that?

A. Yes. The idea is that rather than having the DRAM decide for itself whether or not it's in sync with the global system clock, the module would decide whether or not each DRAM is in sync with the global system clock.

And so you would have a DLL, a single DLL chip on the module with perhaps one DLL on that chip or perhaps more than one DLL on that chip, depending upon what rates of speed you want to reach, but the module would then decide how in sync with the global system clock each DLL would be and delay its local concept of clock.

Q. Just to follow up on that one point you just

mentioned, you said that you'd have one DLL chip on the module?

A. Yes.

Q. But then you said you might have one DLL or you might have more than one DLL?

A. Yes.

Q. Can you please explain what you mean by that?

A. Yes. Well, in looking at the rates of speed that they were considering, you would probably need simply just one DLL to ensure that all DRAMs are in sync with the system, with the system clock, but if you envision going to much higher rates of speed, you could require a separate DLL per DRAM on the module, but these could be put into the same chip.

Q. So in other words, you'd have one chip on the module that would contain multiple DLLs, one for each DRAM?

A. Yes.

Q. Now, with respect to the three sources of delay that you had mentioned when explaining the problem, the outbound delay, the internal delay and the return delay, which, if any, of those sources of delay would be accounted for by a DLL on the module?

A. This would be -- this would account for the internal delay of the DRAM.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to place DLL -- one or more DLLs in a single DLL chip on the module as opposed to placing DLLs in the DRAMs?

A. You eliminate the on-chip DLL from the DRAM, thereby reducing its power consumption, reducing its cost, reducing the design time.

Q. And what, if any, would have been the disadvantages had JEDEC chosen to place a single DLL chip containing one or more DLLs on the module as opposed to placing DLLs on the DRAMs?

A. Well, you then move that design complexity onto a special DLL chip that goes onto the module, so you would be trading one for the other.

Q. Now, the next alternative you have listed is using a vernier method to account for skew.

Can you please explain first what is a vernier method?

A. A vernier is a delay component that's a very accurate variable-delay circuit that provides a static delay to a circuit, but that static delay can be changed by modifying the vernier circuit dynamically.

Q. And then you say "to account for skew." What did you mean by "skew"?

A. Oh, this is what I'm talking about, the timing

uncertainty between parts. This is what we were talking about with the demonstrative showing the problem. The problem is technically skew. That's the problem that you're trying to eliminate. The differences in timing between different things.

Q. So in other words, in this context "skew" might refer to data being out of sync or out of alignment with the system clock?

A. Yes.

Q. Can you please explain how a vernier method might be used to account for skew?

A. Yes. For example, you could put a vernier on each DRAM instead of a DLL on each DRAM and the memory controller would initialize the system to determine what the timing of each DRAM was in the system and set its local vernier so that all of the DRAMs, as far as the memory controller is concerned, all the DRAMs would produce data that would arrive at the memory controller in sync with each other, so it would cancel out the skew between the DRAMs.

Q. And again with reference to the three components of delay that you explained in your -- when explaining the problem, which, if any, of those components would be accounted for by a vernier system?

A. All three.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use a vernier method to account for clock skew rather than using on-chip DLLs?

A. It's simpler to design than a DLL and it would cancel out potentially more skew than a DLL so you could potentially achieve higher data rates using it. And burn less power.

Q. Would that also include the other advantages you've described of not having the DLL on the chip?

A. Yes. Yes. I'm sorry.

Q. In other words, the reduction of power, simplicity, et cetera?

A. Yes.

Q. Now, what, if any, would have been the disadvantages of using a vernier method to account for clock skew?

A. Well, by itself, because the mechanism is not dynamic, as is a DLL, a DLL is continuously updating its concept of how much to delay the local clock signal, a vernier mechanism would not be a dynamic mechanism, and so it would not account for dynamic changes in skew, so these are, for instance, fluctuations in temperature of the system and fluctuations in voltage levels.

So those types of fluctuations would cause

skew, but this would not account for by itself.

Q. Is there a means to compensate for that using vernier methods?

A. Yes, there is.

Q. And what would that be?

A. The memory controller could periodically recalibrate the system to account for these changes.

Q. Could you explain briefly how that would work.

A. Well, just as the memory controller would initialize the system by testing each individual DRAM and setting its vernier to give the appropriate delay to keep all of the DRAMs in sync, here, rather than simply doing it at initialization time, the memory controller would do it periodically.

So for example, voltage fluctuations tend to occur 60 times a second due to the AC power supply and temperature tends to fluctuate about a millisecond -- that's the time granularity that you care about, so if the memory controller recalibrated at every millisecond, that would account for both voltage and temperature fluctuations.

Q. If we could turn then to the next alternative that you've identified, achieving high bandwidth using more DRAM pins and not clock frequency, can you please explain what you mean by that alternative?

A. Yes. The use of the DLL in DDR is there because you're attempting to achieve high bandwidths by scaling the frequency of the data bus. And this alternative is to say let's achieve high bandwidths, rather than by scaling the speed of the data bus, let's scale the width of the data bus, either -- for instance, by having more pins, data pins on the DRAM.

Q. Now, is this the same alternative that you had proposed as an alternative to dual-edged clocking?

A. Yes, it is.

Q. So in other words, it might have been possible for JEDEC to use more DRAM pins in order to avoid both a dual-edged clock and an on-chip DLL?

A. Yes.

Q. Now, what, if any, would have been the advantages had JEDEC chosen to use more DRAM pins rather than placing a DLL on the DRAMs?

A. Well, it would be your data rates would be slower, so you would consume less power on your -- in your data I/O. It's easier to achieve slower rates of speed than it is to achieve higher rates of speed, so it would make your system simpler. You would eliminate your DLL, so you would eliminate the costs associated with that as I've described before.

Q. If I could follow up on something you just

said, you said the data rates would be slower?

A. Yes.

Q. You're referring to data rates across any particular wire and any particular pin?

A. Yes. Yes. I'm sorry. Across the data pins, so rather than toggling your data pin at 200 million times a second, you would toggle the data across any particular pin at 100 million times a second, but you would have twice as many pins so that you would have total bandwidth that's the same, but each pin would toggle at half the rate.

Q. Again, so the record is clear, in terms of system-wide performance, would the data rate be any slower?

A. The total bandwidth would be the same.

Q. And by "bandwidth" you're referring to the --

A. The amount of data received per unit of time.

Q. Now, you've already described the disadvantages of using more DRAM pins when you explained this option in connection with dual-edged clocking.

Are there any other disadvantages that would arise here other than what you've already described previously?

A. No.

Q. And then finally you list relying on DQS data strobe to provide timing.

First of all, can you please explain what is a DQS data strobe?

A. This is a timing signal that accompanies your data that indicates -- for example, on DRAM writes, when the memory controller sends the data to the DRAM, it indicates to the DRAM here is when you should expect to see data on the data bus.

And so presently the way the spec works is that you do not need to look at the timing strobe on the return path. When the DRAM sends the data to the memory controller, the memory controller can ignore the data strobe because the data -- because of the DLL, the data is in sync with the system clock, so the memory controller can simply look at the global system clock to capture the data.

This alternative suggests that you would have to look at the DQS data strobe to determine the timing of the data and ignore the global clock because there would be no guarantee that the data would be in sync with the global clock.

Q. Would it be fair to say then that this alternative doesn't so much attempt to align the data with the system clock as opposed to providing a

different means of measurement?

A. Exactly. This provides a different means of alignment, so the data is aligned with the strobe and so the recipient of the data would look at the strobe rather than the system clock.

Q. Now, with respect to the three components of delay that you identified when you explained the problem, which components, in other words, outbound, internal and return, which of those components would be accounted for by relying upon a DQS data strobe to provide timing?

A. Potentially all.

Q. Now, what, if any, would be the advantages of relying on a DQS data strobe to provide timing rather than using on-chip DLLs?

A. Well, you would eliminate your DLL, which would make your design simpler. It would consume less power. The design would be smaller, cheaper to manufacture, and so forth.

Q. And what, if any, would be the disadvantages of relying on the data strobe to provide timing rather than using an on-chip DLL?

A. Well, what this suggests is that you would not -- that the memory controller would not be relying upon the global clock for the return data, which would

mean that the memory controller would have to have two concepts of time, so the concept of the arrival time of the data would not necessarily be in sync with the global clock and so the memory controller would have to deal with multiple time demands, which means your memory controller design would be more complex.

Q. Is that a significant disadvantage?

A. It's not a significant disadvantage, no.

MR. OLIVER: Your Honor, we're about to move to a new topic. I don't know if this would be an appropriate place to break for lunch.

JUDGE McGUIRE: I have no objection.

How about respondent? Do you have any preference one way or another?

MR. DETRE: We have no objection, Your Honor.

JUDGE McGUIRE: Then why don't we take a break at this time. It's quarter after twelve. How about if we return at 1:30?

MR. OLIVER: That will be fine, Your Honor.

JUDGE McGUIRE: All right. Hearing in recess.

(Whereupon, at 12:17 p.m., a lunch recess was taken.)

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

(1:30 p.m.)

(DX Exhibit Numbers 62 through 94 were marked for identification.)

JUDGE McGUIRE: This hearing is now in order.

At this time you may proceed, Mr. Oliver.

MR. OLIVER: Could we have just one moment, Your Honor?

JUDGE McGUIRE: Sure.

(Pause in the proceedings.)

MR. OLIVER: I apologize, Your Honor. We were short a document and our computer operator went to look for the document. She's not back yet.

JUDGE McGUIRE: Do you want to take a break here, just a short two minutes?

MR. OLIVER: I think if we had just a couple minutes, we could figure it out. I apologize for that.

JUDGE McGUIRE: All right. Go ahead.

(Pause in the proceedings.)

MR. OLIVER: Your Honor, respondent has come to our rescue with copies of the document.

JUDGE McGUIRE: Thanks very much, Mr. Detre.

BY MR. OLIVER:

Q. Good afternoon, Professor Jacob.

A. Good afternoon.

Q. If we could turn then to the second question that you were asked, whether reasonable engineers in the early to mid-1990s would have understood from Rambus' '898 patent application that Rambus might be able to obtain patents with claims covering the technologies in question as proposed for use in JEDEC's SDRAM and DDR SDRAM standards.

And first, can you please summarize the conclusion that you reached to this question?

MR. DETRE: Objection, Your Honor.

Just renewing the objection from earlier, there's been no foundation laid for this witness to testify about the state of mind of engineers in the early to mid-'90s. The witness was a graduate student at the time, he has no personal experience, and there's nothing in his expertise that would give him insight into the state of mind of those engineers in that time frame.

So this is not reliable testimony under Daubert and its progeny.

JUDGE McGUIRE: Mr. Oliver, any response?

MR. OLIVER: Your Honor, he was working in memory at the time as a graduate student.

Furthermore, as he has said, his studies

through time up to the present have used many different forms of memory, including asynchronous memory, such as EDO and fast page mode memory, which were in -- which were developed in the late 1980s and in use in the early 1990s.

In addition, as I think the testimony will make clear, what we are doing is looking at the Rambus '898 application in order to determine the understanding that can be drawn out of that application.

JUDGE McGUIRE: Overruled. I'll hear the question.

MR. OLIVER: Thank you, Your Honor.

BY MR. OLIVER:

Q. Professor Jacob, can you please summarize your conclusion to that second question?

A. No, engineers would not have suspected infringement.

Q. Why not?

A. Because everything described -- the technologies described in the '898 application are either different implementations than the -- we'll compare these to the technologies in dispute.

They're either different implementations or they were mechanisms that are clearly there in the

Rambus application or the Rambus specification to solve problems that are inherent in Rambus' narrow, multiplexed, packetized bus structure.

MR. OLIVER: Your Honor, may I approach?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a copy of a document marked CX-1451. This is a copy of Rambus' original '898 patent application.

Is this a document that you reviewed in the course of your work in this matter?

A. Yes, it is.

Q. And if we could start, please, with the so-called description of the invention -- or summary of the invention -- excuse me -- that appears on page 9 of CX-1451, internal page 7. That runs through page 12 and followed by a brief description of the drawings.

Did you review that portion of the application at the time that you reviewed CX-1451?

A. Yes, I did.

Q. Can you please give a brief overview of your understanding of the Rambus invention as described in CX-1451?

A. It is a narrow bus, multiplexed bus, meaning

all of the information required to carry on a transaction is transmitted over the same wires rather than having separate, dedicated wires to carry the information. And there are no point-to-point connections in it. It's -- for instance, there's no chip-select network.

Q. If on page 9 I can direct your attention underneath Summary of the Invention to the line, which is apparently I believe in line 16, "and the bus has substantially fewer bus lines than the number of bits in a single address."

Do you see that?

A. Yes, I do.

Q. Can you please explain your understanding of what that means?

A. Well the addresses of the day were typically several dozen. There are several dozen bits in an address. So for instance, if you had a 32-bit address, that's a four-gigabyte space.

So this is saying that your bus width in this invention is significantly narrower than a single address. And if a single address has 24 to 32 bits in it, then their width of the bus is significantly or, rather, substantially smaller than that, so this is what they're talking about, for instance, an eight-bit

bus.

Q. And how does that compare with the bus used with a JEDEC-compliant synchronous DRAM?

A. These days buses are about 144 bits plus chip-select information -- well, there's a lot of wires generally.

Q. Now, Professor Jacob, do you have an understanding of the term "multiplexed" as it would be applied to a bus architecture?

A. Yes, I do.

Q. And what does the term "multiplexed" refer to?

A. It means at different times different information, different classes of information, are transmitted over the same wire.

So for example, the main classes of information would be control information, address information, and data, and in a multiplexed bus, for example, if you transmit data and/or address and/or control over the same wire at different points in time, then that would be a multiplexed wire, so a bus is one that transmits the different pieces of information at different times rather than having a separate, dedicated set of wires for each class of information.

Q. Now, how, if at all, does the term or does the

concept of multiplexing relate to the Rambus invention as described in this '898 application?

A. Could you repeat the question.

Q. Yes.

How, if at all, is the concept of multiplexing related to the Rambus invention as it's described in the '898 application?

A. Oh, it's inherent in the design.

Q. The concept of multiplexing is inherent in the Rambus design?

A. Yes. Yes. They are using the same bus to transmit the request to the DRAM that they use to transmit the data from the DRAM back to the memory controller, so you're transmitting request information and data information all in the same set of wires, so that, by definition, is multiplexed.

Q. Now, how, if at all, does the concept of multiplexing apply to bus architecture in JEDEC-compliant synchronous DRAMs?

A. There's very little multiplexing going on. You have a separate control bus, you have a separate address bus, you have a separate data bus, and you have a separate chip-select bus, so those classes of information are being passed down their own separate wires.

Q. And Professor Jacob, are you familiar with the concept of packets?

A. Yes, I am.

Q. Can you please explain your understanding of the concept of packets?

A. A packet would be when you send a -- well, it's the term that is used to mean a bundle of information and the implication is that that bundle is sent over multiple cycles of time rather than transmitted all at once, and so if you had a narrow bus, if you have, say, more information than can be put onto the bus in one cycle, then that means it takes more than one cycle to transmit that information. Therefore, what you end up with is a packetized structure.

Q. Now, how, if at all, does the concept of packets or packetized system relate to the Rambus invention described in the '898 application?

A. Again, it's inherent in the design.

Q. In other words, a packetized system is inherent in the Rambus design as described in the '898 application?

A. Yes, it is.

Q. And how, if at all, does the concept of packets or packetized system relate to JEDEC-compliant SDRAMs?

A. They're not packetized in the same way.

Q. Now, based on your review of the '898 application, CX-1451, did you see a description of any particular clock system in CX-1451?

A. Yes, I did.

Q. Can you please explain the clock system that you saw described in that application?

A. Yes. The clock is what has been usually referred to as a U-shaped clock, meaning that what you do is you send out on a piece of wire a clock signal, and it traverses through the system and sends its signal to every component in the system, and then at the far end it turns around and comes back, and every component in the system sees the return signal as well, so you have an early version of the signal and a late version of the same signal that are then propagated to each component in the system, meaning each DRAM in the system as well as the memory controller.

And so what they're doing is they're sending on one wire two logical signals, an early version of a clock signal and a late version of a clock signal.

Q. I believe that you have a demonstrative that helps to explain this?

A. Yes, I do.

Q. And while we're pulling that up, let me state for the record that we had shown a demonstrative entitled Second Central Question which would be DX-95, and the next demonstrative that we'll bring up will be DX-96.

DX-96 has a caption Rambus Clock Synchronization, and can you please explain what is shown in DX-96?

A. Well, this is the picture of that U-shaped -- this is a picture of that U-shaped clock where you have an outgoing version of the signal and then it turns around at the far left-hand side -- I'm sorry -- far right-hand side, and it turns around and comes back.

And so what's happening is you send the clock pulse -- for example, I've got little clock faces here to show the arrival of the signal, the time at which a signal arrives at a given point.

The signal is sent out at noon. It arrives at the first DRAM several minutes later. It arrives at the second DRAM several minutes later after that. It arrives at the third DRAM a little later, and so forth. Finally it arrives at the turnaround point of the U-shaped wire at quarter past and it makes its return path and connects to each of those DRAMs as a

separate signal, so that's why it's called clock 1 and clock 2.

In Rambus' terminology they use both early and late clocks and clock 1 and clock 2.

So this shows the time at which it arrives at each of those locations as well. Finally, the signal arrives back at the bus master at half past, and so if the bus master and each of the DRAMs has a local circuit that can take the early version of the signal and the late version of the signal and find a time average between these two, then they can all synthesize an internal clock signal that is essentially a representation of what time it is out at the turnaround point, meaning quarter past.

Q. Now, again just so the record is clear, the clock signal would be leaving the bus master where it reads "outbound clock" at the top of the bus master box?

A. Yes.

Q. And then it would traverse out to the right, turn around at the far right and come back to the point where it reads "inbound clock"; is that right?

A. Yes.

Q. Now, Professor Jacob, I've directed your attention specifically to pages 9 through 13 of

CX-1451, which is the '898 application.

Looking at page 13, there's a caption that reads "Detailed Description" and then behind that of course there are many pages that continue.

Just in very general terms, what is described starting at page 13 through the rest of CX-1451?

A. This describes the Rambus mechanism, the details of the system.

Q. Let's look if we could, please, at figure 2 of the application. This appears at page 130 of CX-1451.

Now, did you review figure 2 when you were reviewing the '898 application?

A. Yes, I did.

Q. Can you please explain what is depicted here in figure 2?

A. This is that bus. This is the Rambus bus organization.

Q. What do the four boxes across the top indicate?

A. The four boxes across the top are the different components that are in this example set-up, including a processor CPU, a ROM and two DRAMs.

Q. And then what do the horizontal lines indicate?

A. Those are the wires that comprise the bus.

Q. And what are the sets of vertical lines?

A. Those are each component's connection to the bus.

Q. Now, how, if at all, does the bus illustrated in figure 2 differ from the bus structure of a JEDEC-compliant SDRAM?

A. It's very different. For one thing, all of the major components of the bus are bused -- bus interfaces. They're multidrop interfaces so they're shared wires, they're shared between all the components using the system as opposed to having point-to-point connections.

And there's only a small number of bus data lines. There's eight them, bus data 7, 6, 5, 4, 3, 2, 1 shown along the right-hand side. You have eight bus data lines, and all the necessary information passes across those and added valid -- including control information, address information and data, everything propagates those few wires as opposed to a JEDEC system which has lots, lots of wires, an order of magnitude more.

Q. Professor Jacob, are you familiar with the concept of a chip-select line?

A. Yes, I am.

Q. What is a chip-select line?

A. A chip-select line is an -- it's a wire in JEDEC architecture, JEDEC memory architecture, that enables a rank of DRAMs and that identifies a rank of DRAMs.

Q. Now, on figure 2, does that contain a chip-select line?

A. No, it does not.

Q. And I believe you mentioned that JEDEC-compliant SDRAMs do?

A. Yes. Yes, they do.

Q. In the Rambus architecture, how does that architecture designate a particular DRAM or a rank of DRAMs?

A. It does it by encoding information in the packet itself.

So for example, when a packet is sent out by the CPU or the bus master, every DRAM in the system needs to decode that packet and determine if the packet is designated for that particular DRAM or some other DRAM in the system.

Q. If I could ask you to turn within CX-1451 to pages 20 through 25, please, and again, these are pages designated in the very lower right-hand corner.

Beginning on page 20, towards the bottom the caption reads "Bus." Do you see that?

A. Yes, I do.

Q. And then underneath that, continuing on the next page, the description of bus followed by a caption Protocol and Bus Operation?

A. Yes.

Q. Can you explain just in general terms what is described in the pages 20 through 25 of CX-1451?

A. Yes. This describes the physical organization of the bus as well as the operation of the bus, its protocol, the format of the information that's put out onto that bus.

Q. If I could direct your attention in particular to page 23.

And towards the bottom of page 23 there's a paragraph that begins "Any preferred implementation of this invention shown in figure 4, a request packet 22," et cetera, and then it continues.

Do you see that?

A. Yes, I do.

Q. Now, what is the request packet that is being referred to there?

A. That is the information that's required to handle particular requests, the information that's sent to a DRAM to get it to perform a request or a read request or a write request or something along that.

Q. That paragraph references figure 4, so why don't we turn to figure 4. That appears on page 131 of CX-1451.

Can you please explain what is depicted in figure 4 of CX-1451?

A. Yes. This is the format of that request packet.

Q. Let me start, if I could, by directing your attention to the top box and then within that box in the top there are a number of vertical lines that go partway down the box just above "access type master 03."

Do you see those vertical lines?

A. Yes, I do.

Q. What do those vertical lines indicate?

A. Those are dividing the -- that information into separate bits, each of which will travel along a separate wire.

So the top box, starting with the adder valid bit, continuing on with the bus data, gives you a total of nine bits of data that would travel along the nine wires, the bus data wires and the adder valid wire.

Q. What do the rows in figure 4 indicate?

A. The rows indicate the passage of time, so the

top row would be transmitted on the first bus cycle, the next row would be transmitted on the next bus cycle, and so forth, so that the packet takes six bus cycles to transmit.

Q. And that was the particular packet illustrated in figure 4 would take six --

A. Yes. That packet would take six cycles, yes.

Q. And then could you give us a brief description of the information that's being transmitted in the six various clock cycles.

A. Yes. Along the top row we have access type and master. The access type is a four-bit field. It shows access type 0 through 3. That's what that means, 0 bit, 1, 2, 3. That is information identifying what type of access it is, for instance, whether it is a read or a write. And it also identifies whether the read or write should perform a row activate before a column read or not. And it ultimately indicates a latency to use.

The other four bits, the master bits, indicate something about the entity that's driving the request out onto the bus.

Q. And then the rows below that reads address?

A. Yes.

Q. What does that indicate?

A. This indicates that those bits are the -- together several different cycles worth of information make up a single address, so first -- the first row of address, address 0 through 8, that indicates that that's the first nine bits of the address. The next row is the next nine bits of the address, and so forth. And finally, in the last row, you have the last four bits of the address and four bits worth of block size information.

Q. Now, within the address information, that would contain both the row address and the column address information?

A. Yes, it would.

Q. We'll return to block size in a moment.

Now, how, if at all, does the system depicted in figure 4 differ from the operation of a JEDEC-compliant SDRAM?

A. In a JEDEC-compliant SDRAM, much of the same information is transmitted, but instead of being transmitted over multiple cycles over the same sets of wires, the information in the JEDEC system is transmitted in one bus cycle over a large number of dedicated wires.

Q. Now, with reference to the '898 application as a whole, that is, CX-1451 as a whole, does that

application describe any bus structure other than the narrow, packetized bus structure that you've described?

A. No, it does not.

Q. If I could ask you to turn next, please, to pages 29 and 30 of CX-1451.

And were these pages among the pages that you viewed when you reviewed the '898 application?

A. Yes, they are.

Q. There's a discussion on these pages of block size.

Can you please explain the concept of block size as it relates to the '898 application?

A. Block size is a -- as I was saying, it's -- it is a four-bit field in the packet, and in those four bits it encodes an amount of data to be transferred either for a read or for a write.

And on page 30, at the top there's a table that indicates for which values of block size that you find in the packet how much data that corresponds to, so the part is able to transmit or receive anywhere from zero bytes to one kilobyte, 1024 bytes worth of data.

Q. Let's take a look at that table. Could you please explain in a bit more detail how that table

delineates the block size.

A. Okay. So for example, if you have four bits of data, that can represent any number from 0 to 15, so if a DRAM sees the number 0 in that bit field in the block size bit field of the packets, it would take that to mean that zero bytes were to be transferred.

If it saw the value 1, it would think that there should be one byte transferred, and the same goes up to 7 and 8.

If it sees the value 9 in the block size field, it takes that to mean that 16 bytes should be transferred.

If it sees the value 10 in the block size field, it should take that to mean that 32 bytes should be transferred, and on, so on, up to if it sees the value 15, it should take that to mean 1024 bytes are transferred.

Q. Now, anywhere within Rambus' '898 application, that is, anywhere within CX-1451, does the '898 application anywhere describe use of a variable block size feature in the context of DRAM generally, as opposed to in the context of a narrow, multiplexed bus?

A. No, it does not.

Q. Now, how, if at all, does the implementation of

the variable block size feature that is depicted here in pages 29 and 30 of CX-1451 differ from the implementation of programmable burst length in the JEDEC-compliant SDRAM?

A. In the Rambus mechanism it is a dynamic mechanism, meaning that this block size information is carried in the request packet, so it can change from request to request very easily, very efficiently.

In a JEDEC mechanism, the programmable burst length feature is something that must be programmed by a special initialization command, so -- and it is typically used in a JEDEC system such that it's set once at the system start-up and never changed again, so the JEDEC mechanism is not dynamic. This mechanism, the Rambus mechanism, is very dynamic.

Q. Well, is the use of the block size feature as described in the '898 application related in any way to the narrow bus, multiplexed bus structure?

A. Yes. The block size feature as described here is really -- it's clearly there to solve a scheduling problem that is inherent in the narrow, packetized bus structure that Rambus has described.

If you have a narrow, multiplexed bus architecture where the requests are being -- are using the same set of wires that the data is using, then

without a feature such as this where you can very efficiently change the block size on a request-by-request basis, it becomes very inefficient to try to use this in a system where you expect your data size needs to change.

Q. Now, how, if at all, does that differ from the way in which programmable burst length is used in a JEDEC-compliant SDRAM?

A. In a JEDEC-compliant SDRAM, programmable burst length is there as a convenience. It is something that is set once and not changed again, because it's not necessary for bus scheduling in the JEDEC system because you have separate buses. Because the control information uses a totally separate bus from the data information, you naturally get this -- a very efficient pipelined nature of data transfer without having to resort to something like this.

Q. Now, in your opinion, would an engineer reading the '898 application, CX-1451, in the early to mid-1990s have thought that Rambus could obtain patent rights over the programmable burst length feature as it was used in the JEDEC SDRAM standard?

A. No, they would not.

Q. Why not?

A. Because this mechanism is clearly there to

solve the problems that are inherent in the narrow, multiplexed bus architecture. It's inherently tied to that narrow, multiplexed bus architecture.

Q. Let me ask you to turn next, please, to pages 16 and 17 of CX-1451.

And if I could direct your attention in particular on page 16 to the bottom paragraph appearing on that page.

You'll see that there's a reference to access-time registers on that page?

A. Yes, I do.

Q. Can you please explain the concept of access time as used in Rambus' '898 application?

A. Access time refers to the transpiring of time between a -- the receiving of a request and the response to that request, so it's a delay between receiving a request and, for instance, driving the corresponding data out onto the bus.

Q. If I could then ask you to turn, please, to page 29 of CX-1451.

And I'd like to direct your attention to the table appearing at the top of page 29, please.

And one of the columns of the table in the right-hand side reads "Access Time."

Do you see that?

A. Yes, I do.

Q. Now, could you please explain what is being depicted in this table with particular reference to access time.

A. Well, this is showing the relationship between those access-time registers and the concept of access time as well as the bits in the request packet that are called access type. Remember, the first four bits of each packet are called this access type value.

So depending upon what value you see in that access type field in the request packet, this causes the -- and so that's corresponding to the first column in the table. Depending upon the value that you see in that field, the DRAM -- or depending upon what value the DRAM sees in that field, the DRAM would then use an access time corresponding to the value held in one of these access registers, and that's shown in the far right column.

So for example, if the DRAM sees the value 6 or 7 in the access type field, it would use access register 3 to determine its latency. If it sees the value 4 or 5, it would use access register 2 to determine its access time. And in the middle column it tells you what type of DRAM access would be then performed, whether it's a column read only or a row

activate and a column read.

Q. Now, how, if at all, does the implementation of the access-time register described in the '898 application differ from programmable CAS latency as implemented in JEDEC-compliant SDRAMs?

A. It differs in that it's a very dynamic mechanism. The fact that you can choose between multiple access times by setting different values in your request packet enables you to dynamically change the latency of each request on a request-by-request basis and do so very efficiently, whereas in the JEDEC mechanism you have to change the programmable CAS latency value through a special initialization or a special command that takes extra time and you can't change it on a request-by-request basis without going through that special command.

So in JEDEC-compliant systems you program that feature once at start-up and never change it again, whereas in the Rambus mechanism the fact that you can select different latencies based on bits in the request packet enables a far more dynamic system.

Q. Now, based on your review of the '898 application, how, if at all, is the implementation of the access-time register that you've described related to the narrow bus or packetized

system?

A. It is inherently tied to that narrow, packetized, multiplexed bus structure, because when you have a narrow, packetized, multiplexed bus structure where requests use the same wires as the data, you get -- sort of by definition you get inefficiencies in your scheduling that can only be resolved by having a mechanism that allows you to efficiently change the latency on a request-by-request basis.

Q. Now, how, if at all, does the access time feature as described in the '898 application differ from the concept of programmable CAS latency as it's used in the JEDEC standard?

A. In JEDEC, programmable CAS latency is used merely as a convenience to allow parts of different generations that have potentially different performance characteristics to coexist in the same system and yet behave identically.

So for instance, you can tell a faster part to run slower so that it would match the behavior of other slower parts in the system.

Q. Now, in your opinion, would an engineer reading the '898 application in the early to mid-1990s have understood that Rambus could obtain patent rights over

programmable CAS latency as that feature was used in the JEDEC SDRAM standard?

A. No, they would not.

Q. Why not?

A. Because this is clearly there to solve -- it's clearly in the Rambus mechanism to solve scheduling problems that are inherent in the narrow, multiplexed, packetized bus structure.

Q. Next I'd like to turn to the clocking scheme described in Rambus' '898 application, and if I could ask you to turn, please, to page 145 of CX-1451.

And specifically I'd like to direct your attention to figure 8a appearing at the top of page 145.

Again, was this part of the -- or did you review this page at the time that you reviewed the entire '898 application?

A. Yes, I did.

Q. Can you please explain what is depicted in figure 8a at page 145 of CX-1451?

A. This shows the U-shaped clock organization and two chips attached to that clock organization, so the clock chip labeled "CLK," that block on the upper right-hand corner, sends a clock signal out on the wire that then traverses to the left and it's called

clock 1. And when it gets to the end of the wire on the left-hand side, it turns around and comes back.

And note that each chip has two connections to that one wire, so each chip gets two logical clock signals. It gets clock 1, which is the early version of the clock, and it gets clock 2, which is the late version of the clock.

Q. Again, just to be certain the description is clear, how many actual clock wires are there in the bus?

A. There is one clock wire and that one clock wire carries two logical signals.

Q. And again, to be certain the record is clear, the outbound signal would be on the top horizontal wire in figure 8a?

A. Yes.

Q. And that would loop around the left-hand side and come back in the lower horizontal line in figure 8a?

A. Yes. Labeled "clock 2." Yes.

Q. Now, does this correspond to the demonstrative that you explained a moment ago with the outbound clock and the inbound clock?

A. Yes. Absolutely.

Q. If we could next turn to page 148 of CX-1451.

And on page 148, the figure 12, again, was that -- did you review this page at the time that you reviewed the entire '898 application?

A. Yes, I did.

Q. Can you please explain what's depicted in figure 12?

A. Well, as I said, in that U-shaped clock organization, if every chip is able to look at the early clock and the late clock and determine the time average between the two, then every chip can synthesize an internal version or an internal clock signal that would represent the midpoint between those two.

This is the circuit pictured in figure 12 that performs that time average. It shows the pad, the pin, clock 1, early clock. It takes the early clock signal and it shows underneath that clock 2, the late clock input, and this averages those two.

Q. Now, within the Rambus architecture, where would the circuit depicted in figure 12 be located?

A. A version of this circuit would appear in every chip that is attached to the bus. The memory controller as well as all of the DRAMs would have a copy of this circuit in them.

Q. Now, does figure 12 show a PLL?

A. No, its does not.

Q. Does figure 12 show a DLL?

A. No, it does not.

Q. How does the circuit depicted in figure 12 differ from a DLL circuit as it would appear in a DDR SDRAM?

A. Actually I have a demonstrative that would illustrate that.

Q. If we could pull up the next demonstrative, please, I believe that will be DX-97.

This is a slide with the title Rambus' Delay Circuit versus a DLL.

Professor Jacob, perhaps you could explain using DX-97 how the circuit depicted in figure 12 of Rambus' '898 application differs from a delay-locked loop circuit.

A. Well, as I said, Rambus' delay circuit takes an early clock and a late clock and finds the time average between the two, and a delay-locked loop takes two signals as input and delays one so that it becomes in sync with the other. They perform very different functions.

And as this demonstrative shows, the circuits are very different. One is not equal to the other.

The one on the right is a typical delay-locked

loop implementation. I pulled that out of Dally and Poulton's book on high-speed circuits.

And the figure on the left is Rambus' figure 12 from their '898 application.

Q. Now, how, if at all, does the function of the circuit depicted in figure 12 of Rambus' '898 application differ from the function of a DLL in a DDR SDRAM?

A. It -- they don't have the same function at all.

The circuit depicted in the figure on the left takes two clocks and finds the time average -- takes two clock signals and finds the time average between them.

The figure on the right delays the top signal, which is labeled "ref. clock input," delays that signal so that it becomes in sync with the bottom input signal, which is called "ref. data/CK."

So that the final delayed clock output signal which would be in the upper right-hand corner of that little diagram, that delayed version of the clock would be in sync with the reference data or the reference clock.

Q. Now, considering the '898 application in its entirety, that is, CX-1451 in its entirety, does the

'898 application ever refer to a PLL or a phase-locked loop?

A. No, it does not.

Q. Does the '898 application ever refer to a DLL or a delay-locked loop?

A. No, it does not.

Q. In your opinion, would an engineer reading the '898 application during the mid to late 1990s have understood that Rambus might claim patent rights to on-chip DLL as it was used in the JEDEC DDR SDRAM standard?

A. No, they would not.

Q. Why not?

A. Thoroughly different implementation, different circuit, different function. They are nothing alike.

Q. If I could ask you to turn, please, to page 149 of CX-1451. And I'd like to direct your attention to figure 13 appearing at the top of the page.

And again, did you review figure 13 at the time you reviewed the entire '898 application?

A. Yes, I did.

Q. Can you please explain what is depicted in figure 13 at page 149 of CX-1451?

A. Yes. This is a timing diagram showing the time at which certain events happen.

So for example, in this time going to the right, for example, the top lines labeled "bus clock 1," that shows the arrival of the early clock at the DRAM side, and bus clock 2, the next line down, that shows the arrival of the late version of the clock, and you can see that the late version of the clock arrives at the DRAM a little bit later, so the two are out of sync with each other.

A little further down where it says "internal clock" you see that that now has -- that is a signal that has an edge that is in between the edges of bus clock 1 and bus clock 2, so the internal clock is that synthesized signal that represents the time average between the early clock and the late clock.

So you have this internal clock and you also have the complement of that clock shown underneath it called internal clock complement 74, so you've got an internal clock signal and a complement of it.

Also underneath that you have vertical arrows labeled "input sample" that show the arrival of data at the DRAM or perhaps command, but information arrives at the DRAM and it is latched by those internal clock signals. It shows the negative edge of the internal clock and the negative edge of the internal clock complement latching those signals.

Q. The data signals you referred to, that would be 127 and 125?

A. Yes. Yes.

Q. Now, in figure 13, is the data latched in time with either of the external clock signals?

A. No, it is not.

Q. With what clock signal is the data latched with?

A. It is latched in sync with no external signal in the system. It's latched out of phase with respect to the external clock. And moreover, each DRAM in the system will bear a different phase relationship between that clock signal and when the data is latched, the latched --

Q. So in other words, it's latched only in phase with the internal clock?

A. Yes.

Q. And the internal clock is out of phase with each of the external clock signals?

A. Yes, it is.

Q. Now, how, if at all, does the implementation of the clocking scheme as described in the '898 application differ from the implementation of the dual-edged clocking scheme in JEDEC's DDR SDRAM standard?

A. Well, for one thing, the JEDEC DDR clocking scheme, as opposed to the Rambus clocking scheme which uses one wire to transmit two logical clock signals, the JEDEC DDR standard uses two wires to transmit one logical clock signal. That's the definition of the differential signal.

In addition, the data in DDR is latched in sync with the external clock and for each DRAM in the system, each DRAM bears the same phase relationship between the external clock and when the data is latched, as opposed to Rambus where it's latched out of phase with the external bus clock and each DRAM bears a different phase relationship between when the data is latched and the external signal.

Q. Now, in your opinion, would an engineer reading the '898 patent application during the 1990s have thought that Rambus could obtain patent rights over the dual-edged clocking feature as it was proposed for use or used in the JEDEC DDR SDRAM standard?

A. No.

Q. Can you please explain why?

A. Because they're different implementations.

Q. Next, Professor Jacob, I'd like to consider the claims contained in Rambus' '898 application.

Have you reviewed the original 150 claims of the '898 application?

A. Yes, I have.

Q. In your opinion, would any of those claims have alerted an engineer in the 1990s that Rambus might seek to obtain patent rights over features proposed for use or used in the JEDEC SDRAM or DDR SDRAM standards?

A. No.

Q. Why not?

A. Because all of the claims deal with -- or each of the claims either deals with Rambus' clocking mechanism, which was different from JEDEC's clocking mechanism, or the claim deals with Rambus' packaging techniques, which are different from the packaging techniques that JEDEC was dealing with, or the claim is explicitly limited to the narrow, multiplexed, packetized bus structure, which was unlike the bus organization considered by JEDEC.

Q. I'd like to take a look at three claims and to help illustrate what you just said.

If I could ask you to turn, please, first to page 64, looking at claim number 1.

Do you have that in front of you?

A. Yes, I do.

Q. Now, what, if anything, in claim number 1 indicates to you that this claim would not apply outside of the specific Rambus architecture?

A. Well, for instance, the last two clauses said "bus containing substantially fewer bus lines than the number of bits in a single address" and "said bus carrying device-select information without the need for separate device-select lines connected directly to individual memory devices."

Both of these are very different from the memory bus architectures of the day, the DRAM bus architectures of the day.

Q. Looking at the first clause you identified, said bus containing substantially fewer bus lines than the number of bits in a single address, what about that clause indicates it will be different from the way that JEDEC implemented its SDRAM standard?

A. The way JEDEC implemented its DRAM standard was similar to the way DRAMs had been -- DRAM interfaces had been orchestrated, used, had substantially more bus lines than the number of bits in a single address. In a typical memory bus architecture.

Q. Now, looking at the second of the two clauses that you identified, said bus carrying device-select information without the need for separate

device-select lines connected directly to the individual memory devices, again, what is it about that clause that would serve to distinguish this claim from the way in which JEDEC implemented its SDRAM standard?

A. Well, JEDEC-compliant DRAM organizations did have a need for a separate device-select line connected directly to each individual memory device. That was the way things had been built and this is the way JEDEC continues to build their memory systems, so this is very different from the way things have been done.

Q. Would that device-select line be the chip-select line you referred to earlier?

A. Exactly. That would be the chip-select line that I referred to.

Q. Let me ask if you could, please, to turn to claim number 73. I believe it appears at page 89 of CX-1451.

It actually begins on page 89 and carries over to the top three lines of page 90.

Now, what, if anything, about claim 73 indicates to you that this claim would not relate to the JEDEC SDRAM standard?

A. This describes the Rambus clocking scheme that

generates an early version of the clock signal as well as a late version of the clock signal.

Q. I assume from what you're saying that is then different from the JEDEC --

A. I'm sorry. Yes. This is very different from the clocking scheme that was considered by JEDEC, where you don't consider early and late versions of the clock signal; you just have one, one logical signal.

Q. And I asked you specifically about JEDEC's SDRAM standard.

A. Yes.

Q. Is there anything about claim 73 that would serve to indicate that it would not apply to the JEDEC DDR SDRAM standard?

A. The same, the same is true. The SDRAM and DDR SDRAM both use a single logical clock signal, not this dual local clock signal with early and late signals that are described, for instance, in the second, third and fourth elements of this claim. This is nothing like what was used in SDRAM or DDR SDRAM.

Q. If I could then ask you to turn, please, to claim number 91, which I believe appears at page 99.

Now, what, if anything, in claim 91 indicates to you that this claim would not apply outside of the

particular Rambus architecture?

A. So this describes Rambus' packaging scheme which was to place the connectors of a DRAM only on one edge of the DRAM chip as shown, for example, in the last two lines or the last three lines of the claim, talking about the connections are placed along a single side of the package, and this is very unlike the DRAM packages of the day which used both sides of the package. They used, you know, more than one edge of the package to connect to the rest of the system.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked as CX-1460. I'll give you just a moment to look at it.

(Pause in the proceedings.)

A. All right.

Q. CX-1460 is a copy of the Rambus patent number 5,243,703 issued September 7, 1993.

Professor Jacob, did you review Rambus' '703 patent in connection with your work in this case?

A. Yes, I did.

Q. In your opinion, would Rambus' '703 patent have alerted an engineer during the 1990s that Rambus might

seek to obtain patent rights over features contained in the JEDEC SDRAM or DDR SDRAM standards?

A. No, it would not.

Q. Why not?

A. Because this patent deals with the U-shaped clock or the reflected clock. It's a -- it deals with schemes using two logical clock signals, an early version of the signal and a late version of the signal, finding the time average between the two.

Q. If I could ask you to turn to page number 24, please, in CX-1460.

And I'd like to direct your attention to claim number 1, starting a third of the way down the left-hand column on page 24 and continuing onto the right-hand column.

And can you please explain what it is about this claim number 1 that indicates to you that the claim was limited to the particular Rambus clocking scheme?

A. Well, the claim deals with what I was talking about, a clock system that delivers an early version of a signal and a late version of a signal, and the reason that you can determine this is because it talks about path length matching between clock 1 and clock 2 and a turnaround point.

For example, lines -- somewhere around line 38 or -- we'll start at line 35 in the first column, wherein the first clock signal generation means is coupled at, one, a first point of the transmission line means for receiving the clock signal, the global clock signal, and two, a second point of the transmission line means for receiving the global clock signal wherein the first point is between the first end and the midpoint, wherein the second point is between the midpoint and the second, they're talking about the path length matching things so that the reflection time for the first -- for the first clock and the second clock, they're equal, their transmission to the midpoint are equal so that you can have a nice midpoint between the two so that every chip in the system has the same time average.

And you know, you get the same sort of language at the end of the claim that talks about synchronizing between these two.

Q. And that describes the U-shaped, the loop clock that you had described earlier?

A. Yes. Exactly.

Q. Now, within the '703 patent there's the so-called specification, which is the language appearing before you reach the actual claims.

Is there any -- based on your review of that part of the patent, is there anything in that portion of the patent that differed in any substantial way from the description you mentioned provided in the '898 application?

A. Nothing that I noted. It seemed to be substantially the same as what was in the '898 application.

Q. Was there then anything in the specification portion of the '703 patent that would have alerted an engineer during the 1990s that Rambus might be able to obtain patent rights over features contained in JEDEC SDRAM or DDR SDRAM standards?

A. No.

MR. OLIVER: May I approach?

JUDGE McGUIRE: You may.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked CX-887. It's a two-page document. The first page is a letter to Mr. Ken McGhee of Electronic Industries Association, dated June 17, 1996, from Mr. Richard Crisp. And the second page is also dated June 17, 1996, with a caption Rambus U.S. and Foreign Patents, followed by a list of numbers.

Professor Jacob, did you review CX-887 in

connection with your work on this case?

A. Yes, I did.

Q. And did you review the various patents that are listed on page 2 of CX-887?

A. Yes, I did.

Q. Now, in your opinion, would the patents listed in the attachment to the letter that is on page 2 of CX-887 have alerted an engineer in the 1990s that Rambus might be able to obtain patent rights over features incorporated in the JEDEC SDRAM or DDR SDRAM standards?

A. No.

Q. Why not?

A. Because in each of the cases the patents fall under one of several different categories. They either are restricted to the -- explicitly restricted to the narrow, packetized, multiplexed bus organization, or they contain -- or deal with the topic that lies outside of the scope of JEDEC 42.3, or they are describing material that could have related to or could have been within the scope of JEDEC 42.3 but cover minor implementation details that JEDEC did not consider in the definition of the standard.

Q. Professor Jacob, I'd like to turn now to the third of the four questions that you addressed, and

that's the question of whether Rambus had issued patents or pending patent applications before June of 1996 that contained claims that a reasonable engineer would believe covered the four technologies at issue in this case.

Now, Professor Jacob, have you reviewed the claims of any patent applications that Rambus submitted to the Patent and Trademark Office before June of 1996?

A. Yes, I have.

Q. And have you reached a conclusion as to whether any of those claims would or might be considered to cover the technologies at issue in this case?

A. Yes.

Q. Could you summarize very briefly what your conclusion is.

A. That there do exist patent claims that cover each of the technologies in dispute.

Q. By the way, I should note for the record that we have pulled up another demonstrative slide, entitled Third Central Question. I believe that would be DX-98.

JUDGE McGUIRE: Mr. Oliver, can I get some hard copies of these screens, these demonstratives, at some

point?

MR. OLIVER: Yes. Absolutely, Your Honor. I don't know if we'll have a set for you by today but certainly by tomorrow morning.

JUDGE McGUIRE: That's fine.

BY MR. OLIVER:

Q. Professor Jacob, let's start with programmable CAS latency.

Have you formed any opinions as to whether an engineer could reasonably construe any of the claims that you reviewed to cover programmable CAS latency as that technology is used in the JEDEC SDRAM and DDR SDRAM standards?

A. Yes.

Q. And what is your opinion?

A. That they did. The claims did cover and you could reasonably infer that.

Q. And looking also at programmable burst length, have you formed any opinions as to whether an engineer could reasonably construe any of the claims that were pending at the time to cover programmable burst length as that technology is used in the JEDEC SDRAM and DDR SDRAM standards?

A. Yes.

Q. And briefly, what is your conclusion?

A. That one could -- yes, an engineer could construe that claims did cover programmable burst length as was discussed in JEDEC 42.3.

MR. OLIVER: Your Honor, if we could have just a moment to set up an easel in order to -- in a moment or two we want to set up a couple of demonstratives on the easel.

JUDGE McGUIRE: Take a few moments.

We'll just go off the record for two minutes. If you want to get up and stretch, go ahead.

(Recess)

JUDGE McGUIRE: Mr. Detre.

MR. DETRE: Your Honor, before Mr. Oliver begins this line of questioning, it looks to me like he's about to get into programmable CAS latency, purported programmable CAS latency claims which are, according to the demonstratives they shared with us, are the same ones that Mr. Nussbaum testified about.

In other cases, the FTC has always taken a position that there should be one expert per topic. This is cumulative, Your Honor, and I don't think we should hear it again.

JUDGE McGUIRE: Mr. Oliver?

MR. OLIVER: We're simply presenting two points. First of all, we're merely presenting it from

a different point of view, from a technical expert as opposed to a patent expert, depending on whether either you or the commission wishes to attach more weight to a patent expert or more weight to a technical expert. It's not cumulative because it's coming from an additional point of view.

In addition, I believe that Professor Jacob may have additional information and understanding with respect to certain aspects of the technology that may bear upon his testimony.

JUDGE McGUIRE: Overruled.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked CX-1504, which is a prosecution history of a Rambus patent, and I'd like to ask you to turn in particular to page 216 of CX-1504.

A. Okay.

Q. And beginning on page 216 is a portion of this prosecution history that consists of an amendment at the top captioned In The United States Patent and Trademark Office. In the box on the right-hand side is a date, January 6, 1995. The left-hand side indicates the serial number 07/847,961.

Professor Jacob, did you read this amendment to the '961 patent application in the course of your work in this matter?

A. Yes, I did.

Q. If I could ask you to turn to pages 220 to 221 -- I apologize -- 221 through 222. And I'd like you to focus in particular on claim number 160 beginning at the top -- excuse me -- at the bottom of the page 221, carrying over to the top of page 222.

Now, Professor Jacob, in your opinion, could a reasonable engineer construe claim 160 of this amendment to cover programmable CAS latency as that feature was proposed for use in the JEDEC SDRAM and DDR SDRAM standards?

A. Yes.

Q. Could you please explain the basis for your conclusion.

A. If we take this element by element --

Q. Actually if I could interrupt for just one moment.

A. Oh, there we go.

MR. OLIVER: Your Honor, we have an additional set of demonstratives that consist of claim charts that I would propose to hand out, and again we'll mark these also as demonstratives.

JUDGE McGUIRE: Okay. Go ahead.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. I apologize for interrupting you, Professor Jacob. You were starting to analyze claim 160 element by element.

Let me ask first whether there's a page in the demonstratives that I've just handed out that relates to that.

A. Yes. Page 1 of this demonstrative, claim 160 of application number 07/847,961, contains the wording of that claim in the first column or sort of the second column after the number, and in the far right column we have descriptions in JEDEC work that was going on. So --

Q. Why don't we indicate for the record that this will be marked as DX-99.

Your Honor, this demonstrative consists of eight numbered pages. Should we simply just mark the entire demonstrative as just --

JUDGE McGUIRE: Yes. Let's just mark it as the same demonstrative.

BY MR. OLIVER:

Q. Okay. Professor Jacob, I believe then you're

referring to page 1 of DX-99?

A. Yes.

Q. And Professor Jacob, we have also placed a copy of a page from release 4 of the JEDEC SDRAM standard on the easel, and if it is helpful for you to do so, please feel free to approach that, and if you wish to mark up or make any marks on the blowup, please feel free to do so. I have a marker here that I believe will work. The ones next to you probably don't work, but I think this one here should work.

A. Okay. Great. Thanks.

Q. And with those materials, please make use of any of those materials you find helpful, but can you please explain your conclusion as to why you believe that a reasonable engineer could construe claim 160 of the amendment to the '961 application as potentially covering a CAS latency feature as proposed for use in the JEDEC SDRAM standard?

A. All right. Well, let's go through this element by element.

So element 1 describes a memory storage system including a bus and a semiconductor device, and this is clearly the work of JEDEC 42.3 that describes or defines specifications for semiconductor devices and the buses used to interface to them, so buses, for

example, up here on page 134 of JX-56.

Element 2 reads "having that is configurable by a device external to the semiconductor device," and I'm reading that as if the word "that" is a typo. "That" should be eliminated so that the phrase would read a semiconductor device that is configurable by a device.

So element 2 describes a -- the fact that the semiconductor device is configurable by an external device and in comparison to JEDEC work is the configuration information that's provided to the SDRAM by the bus controller, and the type of information that's provided to initialize the DRAM to configure it is described in JX-56 at, for example, pages 114, 115 and 116.

Element 3 describes a -- the fact that the semiconductor must have at least one pin to couple the semiconductor to the bus, and for example, pinouts can be found within the JEDEC specification, for example, page 106 of JX-56.

Element 4 describes a register operative to store information within the semiconductor device.

And for example, up in the diagram on the easel or the illustration up on the easel we have pictured the SDRAM mode register. That is a register within the

semiconductor device that -- the SDRAM is a semiconductor device. That's a register that's on the DRAM and it's operative to store information. And the diagram, the illustration up on the board, shows what types of information are stored there within that register.

For example, a burst length value is stored, a burst type value is stored, and a latency mode is stored within that register.

Element 5 describes the fact that the register -- that the information that's held in the register should specify a manner in which the semiconductor device is to respond to transaction requests.

So for example, if you consider burst length, burst length -- depending upon the value that's held in that burst length field at the mode register in SDRAM, for example, if the value -- if the -- if the value held within the mode register is 010, the burst length should be 4. If the value is 011, the burst length should be 8.

So for example, depending upon the value stored in that portion of the mode register, the DRAM will drive either four bits of data out onto the bus per pin or eight bits of data out onto the bus per

pin.

So that clearly specifies a manner in which the DRAM is to respond to transaction requests.

Element 6 describes the fact that the -- describes how the information is received by the semiconductor device, so it says that the information is received when the semiconductor device is configured and the semiconductor device stores that information that is received from the bus lines into the register.

And for instance, in the JEDEC spec, JX-56, pages 14 -- 114, 115 and 116, the power-on configuration sequence is described wherein the memory controller puts information out onto the bus. That information is grabbed directly off of the bus by the DRAM device and that information is put directly into that SDRAM mode register. And that's all described within the spec.

So there's clearly a relationship there.

And element 7 is sort of a restatement of element 5. It says thereafter the semiconductor -- or thereafter responds to transaction requests in the manner specified by the information in a mode register.

So for example, we'll return to the burst

length example. As shown in the SDRAM mode register, you have a burst length value that can change the way that the DRAM part responds to a read request or a write request.

JUDGE McGUIRE: All right. Just so I'm clear, I want to inquire here. What is meant in this column 2 when it says "JEDEC Work"? What exactly does that term imply?

THE WITNESS: In column 2?

JUDGE McGUIRE: DX-99, at the top of the page there. It seems sort of broad. I'm trying to get an understanding of what that entails.

THE WITNESS: Oh, what this is meant to show is a comparison between the claim language in claims that were outstanding at the time that Rambus was a member of JEDEC to work that was going on within the JEDEC 42.3 subcommittee while Rambus was there, so this points to specific instances of things that were happening at the time that Rambus was a member of 42.3.

JUDGE McGUIRE: Okay. Go ahead, Mr. Oliver.

MR. OLIVER: Thank you, Your Honor.

Your Honor, could I invite Professor Jacob to mark on the demonstrative any particular elements that he was referring to?

JUDGE McGUIRE: He may approach.

BY MR. OLIVER:

Q. Professor Jacob, could I ask you to use this pen to mark on the demonstrative the mode register that you've referred to.

A. (Witness complies.)

Q. And Professor Jacob, can you mark the particular portion of the mode register that you referred to that would determine the burst length?

A. (Witness complies.)

These are simply two different representations of the same register, whereas this simply says that this is reserved to test mode. This bit is not one that's for test mode (indicating).

These are both the mode register. The only difference is that when this bit is a 1, it's in test mode.

So the bottom three bits correspond to the burst length information. This box right here indicates for the eight different values that these three bits can represent, these three bits can represent numbers between and including 0 to 7. For the eight different possibilities that can be held here, how the DRAM should respond (indicating).

So there are eight different things that the

DRAM could possibly do.

Q. Would you please label the two circles you've just drawn as relating to burst length.

A. Yes.

(Witness complies.)

Q. And Professor Jacob, while you're there, could you also mark on that demonstrative which portion, if any, of the mode register would relate to CAS latency.

A. (Witness complies.)

The three bits that are labeled "LTMODE" within the mode register, those bits hold CAS latency information (indicating).

And this is tied to the bottom-most box down here that shows for the eight different values that could be held in LTMODE how the DRAM should respond to requests.

So clearly different values of LTMODE correspond to different behaviors caused by the -- caused -- different behaviors that the DRAM would have.

Q. Let the record reflect that the witness is pointing to a box labeled "latency mode" towards the bottom of the demonstrative.

Your Honor, could we mark this demonstrative as DX-100?

JUDGE McGUIRE: Yes.

THE WITNESS: So -- (witness indicating).

And I have labeled the diagram further to show that the burst length subtable -- or I don't know what you want to call this, but this little table that's call Burst Length and this little table that's called Latency Mode, both of these tables specify manners in which the DRAM is to respond to requests received by the DRAM from the memory controller.

BY MR. OLIVER:

Q. Thank you, Professor Jacob.

A. You're welcome.

JUDGE McGUIRE: Do you want to go ahead and just have that marked as DX-100 while you're up there.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

(DX Exhibit Number 100 was marked for identification.)

BY MR. OLIVER:

Q. Now, Professor Jacob, just so the record is clear, I believe my question was framed specifically with respect to CAS latency, but I believe in your discussion you referred both to CAS latency and burst length.

A. Oh. I'm sorry.

Q. That's quite all right.

I just wanted to be clear, though, that first of all with respect to the explanation that you've given, the claim table on page 1 of DX-99, does that set forth the basis for your conclusion that claim 160 of the amendment to the '961 application could be reasonably interpreted by an engineer as covering the CAS latency feature as incorporated in the JEDEC SDRAM standard?

A. Yes, it does.

Q. And does the claim chart set forth on page 1 of DX-99 also set forth the basis for your conclusion --

A. Yes, it does.

Q. Let me make sure the record is clear.

It also sets forth the basis of your conclusion that claim 160 of the amendment to the '961 application could also be construed by a reasonable engineer to cover the burst length feature as incorporated in the JEDEC SDRAM standard?

A. Yes.

Q. Now, Professor Jacob, in interpreting the claim 160 of the amendment to the '961 application, did you interpret the word "bus" in that claim to be a narrow, multiplexed bus of the type that you described earlier?

A. No, I did not.

Q. Why not?

A. I interpreted "bus" to -- using the normal and common definition of the term, which is simply a collection of wires.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document that's been marked as CX-1892. It's a report of William Huber regarding claim construction and infringement of U.S. patent numbers in connection with the Micron versus Rambus litigation.

Is this a document that you considered in the course of your work on this matter?

A. Yes, it is.

Q. And generally, what is the subject matter of CX-1892?

A. It describes claim construction and infringement of patents.

Q. Could you turn, please, to page 25 of CX-1892.

And if I could direct your attention in particular to the bottom box on page 25, the left-hand side that reads "a bus or external bus." Do you see that?

A. Yes, I do.

Q. And then in the right-hand side there's a caption towards the top that says "Meaning and Basis" and then in the box next to "bus" under Meaning and Basis it reads, "A bus is a set of signal lines used by an interface system to which a number of devices are connected and over which information is transferred between devices."

Do you see that?

A. Yes, I do.

Q. Did you consider that definition of a bus in connection with your claim interpretation?

A. Yes, I did.

Q. And is that definition of a bus appearing at page 25 of CX-1892 consistent with the way that you defined "bus" for purposes of this analysis?

A. Yes, it is.

Q. And by the way, Professor Jacob, did you understand Mr. Huber to be working on behalf of Micron or on behalf of Rambus in connection with this litigation?

A. I believe he was working on the behalf of Rambus.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked as CX-1875. It's also an expert report of William Huber and in this case in connection with the litigation of Rambus versus Infineon.

Now, did you consult CX-1875 in connection with your analysis of the patent claims?

A. Yes, I did.

Q. If I could ask you to turn, please, to page 77 of CX-1875.

I'd like to direct your attention here to the line starting on the left-hand side with "bus" that appears in the second row on page 77.

A. Uh-huh.

Q. And again --

A. Yes.

Q. -- in that row on the right-hand side under Meaning and Basis, "A bus is a set of signal lines used by an interface system to which a number of devices are connected and over which information is transferred between devices."

Again, is that definition of the term "bus" consistent with the way that you defined "bus" in your interpretation of claim 160 in the amendment to the '961 application?

A. Yes, it is.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document that's been marked as CX-1893. This is an expert report of Robert Murphy in connection with the Micron versus Rambus litigation.

Did you consult CX-1893 in connection with your claims analysis?

A. Yes, I did.

Q. If I could ask you to turn, please, to page 27.

Professor Jacob, on page 27, the term -- focusing particularly on the paragraph in the middle of the page, stating "Bus: A plurality of conductors capable of being connected to at least two communicating entities," is that statement of bus consistent with your interpretation of the term "bus" as you used it interpreting claim 160 of the amendment to the '961 application?

A. Yes, it is.

Q. Now, Professor Jacob, returning for a moment to the language of the claim 160 as set forth on page 1 of DX-99, there's a term "transaction request."

Now, what is a transaction request?

A. That's a general term meaning to be inclusive of read requests and write requests and perhaps control information as well.

Q. Is that the way that you interpreted that term in connection with your interpretation of claim 160 of the amendment to the '961 application?

A. Yes, it is.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked as CX-1881. This is the supplemental report of William Huber in the Rambus versus Infineon litigation.

And if I could ask you to turn, please, to page 18 of CX-1881.

And I'd like to focus on the box in the left-hand side that reads "transaction request" and then on the right-hand side under Meaning and Basis it reads: "Transaction request is an instruction to perform one of a set of possible memory operations, such as writing data to or reading data from the specified memory cells of the memory. These operations are specified by binary logic levels provided to the memory device during a single clock cycle and received

by the memory device in response to a clock transition."

Do you see that?

A. Yes.

Q. Now, is that description of a transaction request consistent with the way that you interpreted "transaction request" when you were interpreting claim 160 of the amendment to the '961 application?

A. Yes, it is.

Q. Now, returning again to the language of claim 160 as set forth in page 1 of DX-99, does claim 160 contain any limitation to a device having a device identifier feature?

A. No, it does not.

Q. And if we could turn next to claim 164 of the amendment to the '961 application.

I'll pause for just a moment until we catch up.

(Pause in the proceedings.)

We've pulled up on the computer screen claim 164 of the amendment to the '961 application that appears at page 223 of CX-1504.

Professor Jacob, I believe that you've also summarized the language on page 2 of DX-99.

Let me start first, though, just by asking your

opinion as to whether a reasonable engineer could conclude that claim 164 of the January 1995 amendment could cover the programmable CAS latency feature as contained in the JEDEC SDRAM standard.

A. Yes. Absolutely.

Q. Could you please explain your conclusion.

A. Well, claim 164 is an extension of claim 160, further narrowing the scope of the 160, and as we described earlier, claim 160 describes the mode register. This further limits that to a mode register or a register holding the value indicative of an access time.

So for example, if we go through this claim by -- element by element, element 2 is that restriction, an access-time register, and it's described -- that LTMODE portion of the SDRAM mode register contains CAS latency information, which is an access time. It is -- it defines the amount of time that should transpire between the DRAM receiving a read request and the moment that the DRAM begins placing the corresponding data out onto the DRAM, so that is a latency, that is an access time, that is a variable amount of delay. So that is a de facto access-time register in the SDRAM specification.

Element 3 states that the DRAM should use this

access time information in response to a transaction request specifying the semiconductor. And for example, that -- such a transaction request could specify the semiconductor using a JEDEC or, rather, in the JEDEC organization using the chip-select bus.

Q. Focusing --

A. I'm sorry. As shown on pages 21 and 121 of JX-56.

Q. Focusing on the language you just mentioned, the transaction request specifying a semiconductor device, does claim 164 indicate how the transaction request would specify the semiconductor device?

A. No, it does not.

Q. And could a chip-select line specify the semiconductor device?

A. Yes, it could.

Q. In your opinion, would a chip-select signal be part of a transaction request?

A. Yes, it would.

Q. In your opinion, is there anything in claim 164 that would limit that claim to use of Rambus' device identifier feature?

A. No.

Q. If I could ask you to turn next to claim 151 of the January 1995 amendment to the '961 application.

This can be found I believe at pages 218 and 219 of CX-1504.

Now, did you form any conclusion as to whether a reasonable engineer could conclude that claim 151 would cover any of the features contained in the JEDEC SDRAM standard?

A. Yes.

Q. What was your conclusion?

A. This would also cover work within 42.3.

Q. And can you please explain to us briefly your conclusion?

A. This is very similar to claim 160 that we just looked at, the primary difference being that the focus of this particular claim being a computer system comprising a bus, a bus master and a -- it's basically the focus is the computer system whereas in claim 160 the focus is the semiconductor device within that memory storage system.

So claim 151 is the same thing just in a more broad scope.

Q. In other words, claim 151 would cover a computer system that incorporated the memory storage system covered in claim 160?

A. Yes, it would.

Q. If I could ask you to turn next, please, to

claim 165. This appears I believe at page 223.

Now, did you form any conclusion as to the scope of coverage of claim 165 of the January 1995 amendment to the '961 application?

A. Yes, I did.

Q. And what was your conclusion?

A. This again is similar. This covers the same thing as -- well, not the same thing, but this is very similar in its intent as claim 160 in that it's describing a -- this type of behavior where the DRAM was configurable, only the focus here is the method for configuring the semiconductor device rather than the mode register.

Q. Just to be clear then, did you reach any conclusion as to whether claim 165 would cover the method of configuring operation of CAS latency as described in the JEDEC SDRAM standard?

A. Yes. Absolutely.

Q. And what was your conclusion?

A. It would cover the method of programming CAS latency.

Q. And did you also reach a conclusion as to whether a reasonable engineer could conclude that claim 165 would also apply to the method of determining burst length --

A. Yes.

Q. -- as set forth in the JEDEC SDRAM standard?

A. Yes.

Q. If I could ask you to turn next, please, to page 258 in CX-1504.

Do you have that page?

A. Yes, I do.

Q. Page 258 is another document with the caption at the top reading United States Patent and Trademark Office, in the box on the right-hand side is the date June 23, 1995, and in the middle of the page about three-quarters of the way down Preliminary Amendment.

Professor Jacob, did you review this document in the course of your work on this case?

A. Yes, I did.

Q. If I could ask you to turn, please, to claim 183. This appears at pages 264 and 265.

Professor Jacob, did you reach any conclusion as to whether claim 183 -- excuse me -- did you reach any conclusion as to whether an engineer might reasonably believe that claim 183 of this amendment would cover a computer system used in a JEDEC-compliant SDRAM?

A. Yes. Yes, I did.

Q. And could you please explain how you arrived at

that conclusion.

A. Well, for this I would like to look at page 3 of DX-99, the claims analysis for claim 183 of application 08/469,490.

So this shows that -- how this particular claim relates to, for instance, CAS latency -- the CAS latency of JEDEC SDRAM parts.

So we'll take it element by element.

The first element describes a computer system with a semiconductor device and a bus, and the JEDEC work was clearly focusing on SDRAMs and which operate in computer systems. And buses are shown within the standard, for example, on page 164 of JX-56.

Element 2 describes the access-time register of the semiconductor device that contains an access time for the semiconductor device, and as described earlier and is shown on DX-100, there is a component within the SDRAM mode register of JEDEC-compliant DRAMs that contains CAS latency information and that that CAS latency information corresponds to an access time. It tells the DRAM how long to wait before driving data out onto the bus.

So element 3 describes the -- how the memory controller, a bus master, programs that access-time register. It says that the bus master should transmit

the value -- transmit a value to the semiconductor via the bus and that the semiconductor device should read that value off of the bus and store it locally in its access-time register.

And in JX-56, the SDRAM specification, it shows the power-on configuration sequence where the memory controller transmits configuration information over the bus that is read by the DRAM off of the bus and put directly into this SDRAM mode register, and that's shown on pages 114, 115 and 116.

Element 4 states that the value in the access-time register should indicate to the DRAM or, rather, the semiconductor device how long to wait in response to a request before driving data out onto the bus, for example, how long to wait before satisfying that request.

And as described, this is exactly what the programmable CAS latency feature of JEDEC SDRAMs does, that the programmable CAS latency feature tells the DRAM how long to wait before driving data out onto the bus in response to a read request.

Q. Now, with respect to the term "transaction request specifying a semiconductor device," did you interpret that in the same way that you interpreted the phrase in connection with the January 1995

amendment?

A. Yes, I did.

Q. If I could ask you to look, please, at claim 184 of the same amendment. This appears on page 265.

Did you reach any conclusion as to whether a reasonable engineer might understand that claim 184 applied to a semiconductor device manufactured in compliance with the JEDEC SDRAM standard?

A. Yes, I did.

Q. And what was your conclusion?

A. That it does relate to JEDEC-compliant SDRAMs with the programmable CAS latency feature.

Q. Could you please explain just very briefly what the basis of that conclusion is.

A. This claim is very similar to claim 183, whereas claim 183's focus is on a computer system that includes a bus and a semiconductor device, claim 184 is focusing on the semiconductor device.

Q. If I could direct your attention to claim 185 appearing at pages 265 and 266 of CX-1504.

Did you reach any conclusion with respect to the potential scope of coverage of claim 185?

A. Yes, I did.

Q. And what was your conclusion?

A. That it does cover the JEDEC SDRAMs.

Q. Again, can you explain just very briefly the basis of that conclusion?

A. This is similar to claims 183 and 184, except that the focus of this claim is on the method for programming the access-time register rather than the computer system or the semiconductor device.

MR. OLIVER: Your Honor, I'd like to move into evidence CX-1892. This is the report of William Huber in the Micron versus Rambus matter.

JUDGE McGUIRE: Any objection?

MR. PERRY: Yes, Your Honor.

We had numerous conversations with complaint counsel on the subject of expert reports from other cases, and we've always said they shouldn't be in. They've always been excluded from the stipulations. We didn't know he was going to offer them in today. We certainly do object. We'd be happy to brief that or have a further discussion with you.

JUDGE McGUIRE: Mr. Oliver, any response?

MR. OLIVER: I apologize. I didn't mean to take you by surprise.

Why don't I withdraw it for the time being. We'll consult and figure out how to proceed.

JUDGE McGUIRE: All right. Very good.

BY MR. OLIVER:

Q. Professor Jacob, I'd like to turn now to the topic of on-chip PLL and on-chip DLL.

Your Honor, may I approach?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked as JX-21. This consists of the meeting minutes of the September 1994 meeting of the JC-42.3 subcommittee in JEDEC.

Professor Jacob, did you review JX-21 in the course of your work on this matter?

A. Yes, I did.

Q. If I could ask you to turn, please, to page 86 and take a quick look at pages 86 through 92.

(Pause in the proceedings.)

I'd just like to ask whether you recognize the presentation that appears at these pages.

A. Yes, I do.

Q. And if I could ask you to focus in particular on page 91.

And if I could ask you to explain this very briefly, please, what is depicted on page 91 of JX-21.

A. This demonstrates a proposal for the use of a PLL on a synchronous DRAM and it -- they're showing how

PLL can be used to synchronize an external clock and an internal clock, the external clock being a CLK and the internal clock being ICLK.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

MR. OLIVER: By the way, Your Honor, I should just mention for the sake of updating my earlier time estimate, unfortunately this is taking longer than I expected with the demonstratives and documents and the like, so I do expect that we'll finish today, but I think it's going to take me longer than originally estimated.

JUDGE McGUIRE: Okay. Go ahead.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked as CX-1502. This is the file wrapper for U.S. Patent Number 5,657,481.

And let me ask you to turn in particular to page 205 of CX-1502.

At page 205 is a document with the caption In the United States Patent and Trademark Office. There's a box in the right-hand side with the handwritten date June 28, 1993. About halfway down the page is the caption Preliminary Amendment, in the upper left-hand side in the box Serial Number 07/847,692.

Professor Jacob, have you -- or in the course of your work in this matter, did you review the preliminary amendment appearing at page 205 of CX-1502?

A. Yes, I did.

Q. I'd like to direct your attention in particular to claim 151, which appears on page 208.

Looking at claim 151 appearing at page 208 of CX-1502, did you reach any conclusion as to whether an engineer reasonably could construe claim 151 of this amendment to cover a JEDEC-compliant SDRAM with the addition of a PLL circuit as set forth in the September 1994 NEC proposal?

A. Yes, I did.

Q. And what was your conclusion?

A. That it would -- it would cover.

Q. I'd like to ask you to explain that conclusion, please. Before you do, I think that we have another demonstrative that I'd like to place up on the easel.

A. Yes.

MR. OLIVER: Could I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. I believe you also have a claim chart appearing

at page 4 of DX-99?

A. Yes, I do.

Q. Could you please explain your conclusion with respect to the scope of coverage of claim 151 of the amendment to the '692 application using the claim chart as you wish and also if -- Your Honor, if he could be permitted to approach the easel?

JUDGE McGUIRE: Yes. Go ahead.

THE WITNESS: I'm sorry.

JUDGE McGUIRE: And while you're up there, let's mark that as DX-101.

MR. DETRE: Your Honor, is it all right if I just move over here to observe (indicating)?

JUDGE McGUIRE: Sure.

(DX Exhibit Number 101 was marked for identification.)

THE WITNESS: So we can go through the claim chart element by element, and again this is the claim chart on page 4 of DX-99, which gives the text of claim 151 of application 07/847,692.

So element 1 describes a memory device on a single substrate. That's what DRAMs are, single-chip memory devices.

Element 2 describes a memory array that stores data at addresses. Well, that would be this. This is

your memory array (indicating). This is the box that is in the figure labeled "memory array." That is for storing data at addresses. That's sort of the definition of arrays and memory arrays.

Element 3 describes a clock signal receiving circuit coupled to receive an external clock signal, so this right here, this sort of right triangle on the upper left-hand side of the right half of the diagram, that is -- the triangle is labeled "receiver." This is the clock signal receiving circuit (indicating).

And it is coupled to receive an external clock signal. This is the external clock signal that is labeled "CLK" (indicating).

And in these diagrams, note that the little tiny circles are representing pins, so that's something outside of the little circuit, is meant to be an external signal, so this CLK is an external signal and over here, DQ, this represents an external connection to the data bus, so this is a pin, that is a pin, these two little circles represent pins (indicating).

So we've got a clock signal receiving circuit coupled to receive an external clock signal for generating a local clock signal. That would be this I-clock signal. This is my local clock signal. And it's local because it's internal to the DRAM rather

than being outside of the DRAM.

So this is my local clock circuit -- my local clock signal.

And element 3 says that the local clock signal should be performing memory operations with respect to the memory array. Well, the figure shows the local clock signal coupled to the memory array through this output driver that drives data out onto the data bus, so you know, so this is, you know -- performs read operations (indicating).

As shown in this diagram, the clock is driving data out onto the bus. It's orchestrating that timing.

So that's the operation with respect to the memory array that the clock is performing.

Element 4 of the claim describes a phase-locked loop -- I don't need to label that again. That's that PLL box that's in the middle.

The phase-locked loop is coupled to the clock signal receiving circuit, and as you can see, it's coupled to the clock signal receiving circuit (indicating).

The element says that it should also be coupled to the memory array. And we see that it is coupled to the memory array through this driver circuit

(indicating).

And the element states that it should provide a variable delay to the local clock signal such that the delayed local clock signal is synchronized with the external clock signal. And that's what is shown in the bottom half of both of these two diagrams. On the left we have a system that has no PLL and the corresponding timing diagram. On the right-hand side we show the inclusion of the PLL and the resulting timing diagram.

And as you can see, in the left-hand timing diagram, the external clock signal, this CLK, and the internal clock signal, ICLK, these two are not in sync with each other. They're out of sync. But with the inclusion of the PLL, it now synchronizes the CLK and ICLK signals so that the timing diagram on the right-hand side of this illustration, the two are now in sync.

And that's what the PLL is doing. Its job is to delay this clock so that -- the internal clock, so that the internal clock becomes in sync with the external clock.

Q. Thank you.

A. No problem.

Q. Now, if I could ask you to look once again at

CX-1502, and if I could ask you to turn, please, to page 233.

This is an amendment of -- that was filed with the typewritten notation in the box in the upper right-hand side October 23, 1995. The box in the upper left-hand side again indicates Serial Number 07/847,692.

Professor Jacob, did you review this document in connection with the work you did in this case?

A. Yes, I did.

Q. And if I could ask you to turn, please, to claim 151, which begins at the bottom of page 233 and continues over to the top of 234.

A. Okay.

Q. And did you reach a conclusion as to whether a reasonable engineer could conclude that claim 151 in this October 1995 amendment could cover a JEDEC-compliant SDRAM plus the inclusion of a PLL circuit as proposed in the NEC proposal of September 1994?

A. Yes, I did.

Q. What was your conclusion?

A. That an engineer reading this would conclude that it covers that scenario presented by NEC.

Q. Can you please explain briefly the basis for

that conclusion?

A. Well, this is a rewording of the previous claim and the wording changes are cosmetic; they don't actually change the meaning of the claim.

Q. So would it be fair to say that the substance of your prior analysis as set forth on page 4 of DX-99 would also apply to claim 151 in the October 1995 amendment?

A. Yes.

Q. Professor Jacob, let's turn next to, if we could, to dual-edged clocking.

And have you formed any opinion as to whether an engineer could reasonably construe any of the claims that you've reviewed to cover dual-edged clocking as that technology was proposed for use in the JEDEC DDR SDRAM standard?

A. Yes, I have.

Q. And what is your conclusion?

A. My conclusion is that an engineer would get that understanding by reading some claims.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Go ahead.

BY MR. OLIVER:

Q. Professor Jacob, I have handed you a document marked CX-1494, which is the U.S. Patent

Number 5,513,327 issued to Rambus with the issue date of April 30, 1996.

Professor Jacob, did you review this patent in connection with your work on this case?

A. Yes, I did.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked CX-34. These are the minutes of the JC-42.3 subcommittee meeting of May of 1992.

Is this a document that you reviewed in connection with your work in this case?

A. Yes, it is.

Q. I'd like to ask you to turn, please, if you could to page 32.

And I'd like to direct your attention to item I.D appearing about three-quarters of the way down that page.

It's the section that begins "I.D IBM William Hardell (Austin)."

Do you see that?

A. Yes, I do.

Q. And the third line reads "asynchronous RAS/CAS with synchronous DQ" and the line underneath that reads

"dual clock edge"?

A. Yes.

Q. Now, what is your understanding of how, if at all, the proposal reflected here relates to dual-edge clocking?

A. This is dual-edged clocking. This proposal right here proposes dual-edged clocking.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked JX-28. These are the meeting minutes of the 42.3 subcommittee of December 6, 1995.

Did you review this document in connection with your work in this case?

A. Yes, I did.

Q. And if I could ask you to turn, please, to page 34. It's the section that begins "Future SDRAM Features Survey Ballot." Do you see that?

A. Yes, I do.

Q. And then continuing on to the next page and the next to last bullet point on that page 35, using both edges of the clock for sampling inputs. Do you see that?

A. Yes, I do.

Q. Now, how, if at all, did that reference relate to dual-edged clocking?

A. That is the definition of -- well, a definition of dual-edged clocking, using both edges of the clock to sample inputs.

Q. And sampling inputs there would be a write operation?

A. Yes. At the DRAM side, yes.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Go ahead.

BY MR. OLIVER:

Q. Professor Jacob, I have handed you a document marked JX-31. These are the minutes of the JC-42.3 committee meeting of March 1996.

Is this a document that you reviewed in connection with your work in this case?

A. Yes, it is.

Q. If I could ask you to turn, please, to page 6 of JX-31.

There is a presentation there that reads "Future SDRAM" of Samsung, and within that I'd like to direct your specific attention to page 71.

And underneath the caption Proposed Clocking Scheme there are -- the fourth bullet point, data in sampled at both edge of clock into memory, and the last

bullet point, use both edge of the strobe clock to sample the memory data into controller?

A. Uh-huh.

Q. How, if at all, does this proposal relate to dual-edged clocking?

A. This describes using both edges of the clock to both read and write data to and from the DRAM, so this is dual-edged clocking.

Q. Now, in your opinion, could an engineer reasonably construe claim 1 of the '327 patent to cover a JEDEC-compliant SDRAM that also incorporated the dual-edged clocking proposals that we've just discussed?

A. Yes, an engineer could.

Q. Could you please explain your opinion.

A. Yes. Well, we can go element by element through the claim. I'm going to refer here to page 5 of DX-99 that gives the wording of claim 1 of the '327 patent.

And element 1 describes the DRAM, dynamic random access memory, with a first circuit for providing a clock signal. And as you can see, for example, in the Samsung presentation, specific -- specific mentions of clocks, so -- an SDRAM uses an external clock signal generating an internal clock

signal, so for instance, in the Samsung presentation in JX-56, page 124.

Element 2 describes a conductor pin that couples the DRAM to a bus and a receiver coupled to the conductor and the first circuit, and that simply means the DRAM has a clock pin with a receiver circuit, so, for example, see pinout diagrams within the spec at page 106 of JX-56.

Element 3 describes that the receiver circuit should latch -- the receiver circuit should latch information on a rising edge of the clock and a falling edge of the clock, and I believe we've got -- well, this just -- this is just a description of a dual-edged clock. We've latched data in response to the rising edge of the clock and the falling edge of the clock, which is basically what's being said in the Samsung presentation, data sampled at both edges of the clock, meaning the rising edge and the falling edge of the clock.

Element 4 describes a specific implementation of that where you pingpong back and forth between two input receivers, a first input receiver and a second input receiver, where one latches information corresponding to the rising edge of the clock and the other one latches information in response to the

falling edge of the clock, so this is sort of an implementation detail that's how you would implement dual-edged clocking because any given latch cannot operate on both edges of a clock, a latch can only respond to either of the rising edge or the falling edge, so the only way to perform this is really to have two latches, one in response to the rising edge, one in response to the falling edge.

Q. I notice a reference to interleaving on page 5 of DX-99.

Now, is that the same interleaving technique that you referred to this morning when you were discussing alternatives?

A. Oh, no. No. No, it isn't. That is what I'm describing by saying you would pingpong back and forth between two latches. That is just a form of interleaving, but it's very different from what I was describing this morning where what you interleave between are two complete banks, which are large structures on the DRAM. This would be a very small input receiver that you toggle between an input latch.

Q. Now, in your opinion, could an engineer reasonably construe claim 7 of the '327 patent to cover a JEDEC-compliant SDRAM that also incorporated the dual-edged clocking proposals that we discussed a few

moments ago?

A. Yes.

Q. And what is your opinion?

A. That it would -- that an engineer would construe it could cover.

Q. Can you please explain your conclusion?

A. Well, if we look at the following page of DX-99, this is page 6 that has the wording for claim 7 of the '327 patent. It has a claim chart for that claim.

If we look at element 1, again, this is a -- it describes a DRAM with a clock signal, and again, that's the -- for instance, the Samsung proposal describes clocks explicitly and that's the point of synchronous DRAM. You have a clock signal, so for example, see page 124 in JX-56.

Element 2 describes a conductor pin that couples the DRAM to the bus, so for example, look at page 106 of JX-56 for a pinout diagram to show pins.

Element 3 describes an implementation of dual-edged clocking for driving data out onto a bus and which you toggle between two output drivers through a multiplexer.

And this is the most reasonable form of memory interleaving in -- if you were going to have a

dual-edged clocking scheme, this is the most reasonable implementation.

Q. By the way, just to be clear, in claim 7 of the '327 patent, would that apply to a read operation or to a write operation?

A. I'm sorry. I should have been more precise.

Element 3 describes the DRAM driving data out on both edges of the clock, so that would be a read operation.

Q. Now, with respect to the comparison between the claim 7 of the '327 application and certain of the earlier proposals we looked at, specifically the April 1992 IBM proposal and the survey ballot, would your analysis differ at all comparing the claim 7 of the '327 patent to --

A. No. No. This describes dual-edged clocking.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a document marked CX-1493, which is the copy of the file wrapper for U.S. Patent 5,513,327. I'd like you to turn in particular to page 183 in CX-1493.

Page 183 is a document marked Preliminary Amendment. The box in the right-hand side has a

handwritten date of September 6, 1994. The left-hand side of the box has a caption Serial Number 08/222,646.

Professor Jacob, did you review this preliminary amendment in the course of your work on this matter?

A. Yes, I did.

Q. I'd like to direct your attention to claim number 151 of the '646 amendment. It begins on the bottom of page 184 and carries over to the top of page 185.

Did you reach any conclusion as to whether claim 151 in the September 1994 amendment to the '646 application could be reasonably construed to cover a JEDEC-compliant SDRAM that also incorporated the dual-edged clocking proposals that we discussed a few moments ago?

A. Yes, I did.

Q. And what was your conclusion?

A. That it would.

Q. Can you please briefly explain your conclusion?

A. This describes the general concept of dual-edged clocking, so if you walk through this element by element, we have a DRAM -- of course we've been talking about DRAMs that are coupled to a bus -- a

first circuit for providing a clock signal -- and we've been talking about synchronous DRAMs that had clock signals.

We have a conductor that couples the DRAM to the bus. That's simply a pin and we've -- or I've pointed you to pinout diagrams in the spec.

And then the last component describes the receiver circuit that latches information in response to both a rising edge of the clock and a falling edge of the clock. That is -- that's the general definition of dual-edged clocking.

MR. OLIVER: May I approach, Your Honor?

JUDGE McGUIRE: Go ahead.

BY MR. OLIVER:

Q. Professor Jacob, I've handed you a copy of a document marked JX-57. This is the JESD79 JEDEC double data rate DDR SDRAM specification dated June 2000.

Did you review JX-57 in the course of your work in this matter?

A. Yes, I did.

Q. Did you form any opinion as to whether an engineer could reasonably construe claim 1 of the '327 patent to cover dual-edged clocking as described in JX-57, the DDR SDRAM specification?

A. Yes, I did.

Q. And what was your conclusion?

A. That claim 1 does cover it.

Q. Could you please explain how you arrived at that conclusion.

A. I'd like to --

MR. OLIVER: I believe we have another demonstrative, Your Honor.

JUDGE McGUIRE: Go ahead.

THE WITNESS: May I?

JUDGE McGUIRE: Yes.

THE WITNESS: Thank you.

So I will walk through claim 1 of '327, and if you want to follow, this is on page 7 of DX-99.

So if we go through it element by element, element 1 describes a dynamic random access memory, a DRAM, with a first circuit for providing clock signal. Well, what we're talking about is DRAMs here.

Here is the first circuit for providing a clock signal (indicating). It is this wire in the bottom right-hand side of this figure, the wire that has a number of arrows coming out of it, and for instance, the end of the wire goes into the -- this box labeled "right FIFO and drivers" and where the wire goes into that box is labeled "CKN," meaning that this is the --

we have a clock input to this component here called "right FIFO and drivers," so that means that this wire right here is a clock wire. So this is the first circuit that provides the clock signal right there (indicating).

Element 2 describes a conductor for coupling the DRAM to a bus, so we will consider these pins right here that are labeled DQ 0 through DQ 3, this would be with the conductor -- conductors for coupling the DRAM to the bus (indicating).

Element 2 continues "a receiver circuit coupled to the conductor and the first circuit," so that would be this component that I have circled, which includes the -- well, actually not those so much but mostly this box labeled "receivers" and then the bottom two boxes in this stack that has been labeled "input registers" in the diagram, so this whole circuit right here is the receiver circuit (indicating). That is described in element 2.

And as you can see, it is coupled to the conductor, this conductor, the DQ conductors through that wire. And it is also coupled to the first circuit through that wire coming up from underneath.

Element 3 describes the general behavior of this system. It says that the receiver circuit should

latch information received from this conductor on both the rising edge of the clock signal, of this clock signal which is DQS, and the falling edge of that clock signal, and as shown on page 32 of JX-57, we have a timing diagram that illustrates that. I think we'll pull that up on the screen.

So what we have here in the second to bottom line down here that's labeled "DQ," we have a bunch of stuff happening there over time. What that shows is the little sort of -- I don't know -- hexagons that are filled with white, that's considered valid data.

So these are valid data windows over time. DQ, that means data. And several lines above that is the timing diagram for DQS, and as we see, the edges of DQS, where DQS goes from low to high and where DQS goes from high to low, that corresponds to these valid data windows on DQ. That means that the data on DQ is latched in sync with the rising edge and the falling edge of DQS. That illustrates this behavior. That says this DQS signal right here that's connected up to these latches is causing the data coming off of this pin to be latched on both the rising edge and the falling edge of that DQS clock signal (indicating).

So that's what element 3 is saying.

Element 4 says "wherein the receiver circuit

comprises a first input receiver coupled to the conductor and the first circuit," so I'll label this top box the first input receiver and the bottom box will be the second input receiver (indicating).

So you see two of them stacked on top. One is the first input receiver; the other is the second input receiver. And one is going to latch on the rising edge of the clock and the other is going to latch corresponding to the falling edge of the clock.

And I believe that's it.

JUDGE McGUIRE: Then while you're up there, let's mark that as DX-102.

THE WITNESS: Thank you.

JUDGE McGUIRE: Thank you.

(DX Exhibit Number 102 was marked for identification.)

BY MR. OLIVER:

Q. Professor Jacob, at the bottom of what's been marked as DX-102, there's a notation in note 1 that this is a functional block diagram.

Do you see that?

A. Uh-huh. Yes.

Q. Now, what's your understanding of that statement?

A. I read that as a disclaimer to indicate to the

reader that this is not a fully functional implementation. This is a simplified version of a real implementation and it just happens to be missing some signals to -- basically for clarity. There are places where, for instance, there should be a clock signal, but there isn't.

So it's just a -- it's meant to convey a general idea of how the system works but not necessarily be a fully functional circuit implementation. This is the standard, but you can't -- it doesn't have enough information to completely build the system based on that.

Q. Are the features depicted in this functional block diagram contained in an actual JEDEC-compliant DDR SDRAM?

A. Yes, they are.

Q. Now, Professor Jacob, have you formed an opinion as to whether claim 7 of the '327 patent can reasonably be construed to cover dual-edged clocking as used in the JEDEC DDR SDRAM specification?

A. Yes, I have.

Q. What is your conclusion?

A. My conclusion is that it can be construed to cover it.

Q. I'd like to ask you to explain that conclusion,

and I think we have one more copy of the same page in which I'd like to ask you to again point out the specific elements.

Your Honor, may I approach?

JUDGE McGUIRE: Yes.

THE WITNESS: May I?

JUDGE McGUIRE: Go ahead.

THE WITNESS: Thank you.

And shall I label this?

JUDGE McGUIRE: DX-103.

(DX Exhibit Number 103 was marked for identification.)

THE WITNESS: So I will be reading through the claim 7 as part of the claims analysis, the claims flowchart on page 8 of DX-99.

So element 1 describes a DRAM, a DRAM with a first circuit for providing a clock signal. So for example, here is a first circuit for providing a clock signal. This CLK is a clock. And you see several clocks around the diagram, so there is a clock on the chip. This happens to be one instance of that clock (indicating).

Number 2 or, rather, element 2 describes a conductor for coupling the DRAM to the bus. That would be this little circle off to the right edge that's

labeled DQ 0 through DQ 3. This would be our conductor (indicating) -- conductors coupling to the bus, coupling the DRAM to bus. This is our data bus connection.

And element 3 describes the implementation of driving the data out onto this bus. It describes a multiplexer coupled to the first circuit. So this is our multiplexer that is coupled to the first circuit (indicating). It is the box within the figure that is labeled "MUX." Mux is just a shorthand notation for multiplexer. And it is coupled to col. 0 and we note that col. 0 is the bottommost -- the least significant bit coming out of the column address counter/ -- the column address counter latch box kind of at the bottom of the figure.

Because this is a counter, there is an implicit clock going into this circuit even though it's not showed. That's how counters work.

So we have a clock going into this counter and the bottom -- the least significant bit of that counter, which of all counters the least significant bit goes 010101, which starts to look a lot like a clock, so that col. 0 ultimately goes into this mux, so here our multiplexer that is coupled to the first circuit which is the clock.

The multiplexer having an output, okay, so this is the wire that exits the box called mux to the right is the multiplexer's output. The multiplexer has a first input, and due to space reasons, I'm just going to label that 1. That will be where the top wire enters into that box labeled "mux." The wire with the "/4" through it, where that enters the mux we'll call that the first input at the bottom of the mux. Entering from the left-hand side, we'll call that the mux's second input.

So this is how mux have been drawn, two inputs and one output and a select function.

Let's see, where are we?

Element 3 further describes a first output line, so that would be this one, first output line (indicating). That would be the line with the "/4" through it that's connecting the read latch to the mux. The one on top and the line on bottom would be the second output line.

So we have the element describes a first output line connected to the first input of the multiplexer, and then further down in the element we have a second output line connected to the second input of the multiplexer, and then the second half of each of those paragraphs or subsections is wherein the multiplexer

couples the first output, the first output line to the output of the multiplexer in response to a rising edge of the clock and the multiplexer connects the second output line to the multiplexer output on the falling edge of the clock.

That's what's going on, this select line as it together goes up and down between zero and one, causes this multiplexer to alternate between driving this first output through this circuit, the output label and then selecting this wire, so it goes back and forth between selecting these two wires (indicating).

The first one selects that one and then that one and that one. It goes back and forth between the first output line and the second output line which corresponds to this read latch, has eight bits of data in it, so what this multiplexer does is it first grabs the top four bits and then the bottom four bits and then the top four bits and then the bottom four bits, and so forth, so that's what this element states and that's what's happening in the block diagram.

Q. Thank you, Professor Jacob.

Your Honor, I'm about to turn now to the last of the four topics, and first, I apologize. My time estimate was way off this afternoon. I think we can finish this in probably 40 to 45 minutes.

JUDGE McGUIRE: Okay. I would like to be done this afternoon if we could by quarter after five at the latest, you know, if we could be, which should give you 45 minutes.

MR. OLIVER: That would be fine, Your Honor.

JUDGE McGUIRE: Okay. Go ahead.

BY MR. OLIVER:

Q. Professor Jacob, we're working you pretty hard. Do you need a short break?

A. No. I'm good.

Q. Okay. Professor Jacob, let's turn to the fourth question that you were asked to address today, namely, the implications of a redesign today.

I've brought up on the screen another demonstrative. This simply reads "Fourth Central Question." We'll label this demonstrative as DX-104.

Professor Jacob, can you first please explain how you went about answering this question?

A. Well, I read numerous engineering documents, treatises and technical articles and I read the specification and consulted with a number of engineers in the field to confirm my understanding of what the implications would be.

Q. Now, is there a demonstrative that would help you explain what you mean by the redesign

implications?

A. Yes. That would be the following slide I believe.

Q. We've pulled up a slide entitled DRAM Basics Graphics Card and we'll label this as DX-105.

And Professor Jacob, if you could please explain briefly what DX-105 shows.

A. This shows a typical desktop system, the connections that are existent in your desktop system, and it shows the place of DRAM within that system.

So for example, you have a number of DRAM chips on the module off to the right and you have another group of DRAM chips that are on the graphics card that are talking to the CPU, and everywhere that you have DRAM you have a similar style bus connecting that to your CPU or controller.

Q. I assume that this is a simplified --

A. Yes. Yes. This is rather simplified.

Q. Now, if we could again just very briefly take the various elements one at a time, starting with the memory controller. Could you please remind us just very briefly of what the memory controller does.

A. The memory controller is the centralized access point for the DRAM subsystem, so all of the requests for the DRAM system, whether they come from

the CPU or multiple CPUs or the graphics card, all requests for the DRAM system go through the memory controller.

Q. Now, in general terms, if the specification for the DRAM were to change, what, if any, would be the potential redesign implications for the memory controller?

A. Well, it would depend upon the change, but in general, it could require anything from logic changes, simple logic changes to the memory controller to pinout changes, signaling changes. It could require any number of different changes.

Q. Focusing next on the memory module, can you again remind us briefly of what the memory module does in a memory system?

A. The memory module provides you a convenient package for DRAM. It's a way to collect a number of DRAM into a small chunk that can be used rather than having to deal with individual DRAMs at a time.

Q. Again, in general terms, if the specification for the DRAM were to change, what, if any, would be the potential redesign implications with respect to the memory module?

A. It could require -- depending upon the types of changes you make to the DRAM, it could require trivial

changes to nontrivial changes like changes of -- add a number of connectors to the rest of the system or inclusion of extra chips on the module.

Q. Now, Professor Jacob, are you familiar with the so-called serial presence detect or SPD?

A. Yes, I am.

Q. What is that?

A. That's a chip that's on the module that identifies to the memory controller the configuration of the DRAMs that are on that module.

Q. Now, in general terms, if the specification for the DRAM were to change, what, if any, would be the potential redesign implications with respect to the serial presence detect?

A. It could change trivially or nontrivially.

For instance, if you had fixed latency parts, the serial presence detect would have to identify that information, that sort of thing.

Q. Focusing then on the motherboard, which is not specifically depicted here but which would be in the area in blue on which the various bus lines appear, in general terms, if the specification for DRAM were to change, what, if any, would be the potential redesign implications with respect to the motherboard?

A. Again, it depends upon the types of changes

that would go into the DRAM, but the types of changes could range from trivial or no changes whatsoever to nontrivial changes like the addition of more wires, a lot more wires, termination, that sort of thing.

Q. Looking next at the central processing unit or the CPU, and again in general terms, if the specification for the DRAM were to change, what, if any, would be the potential redesign implications for the CPU?

A. Well, the CPU is the originator of most of the requests in the system. It's the thing that's generating requests for data that are ultimately turned into reads and writes. And in most general-purpose systems like this, the CPU requests data at the granularity of a cache block size, so if, for example, you change the block size available from the DRAM, if you change the burst length to a different value than the one that the CPU expected, you might want to change your cache organization.

Q. Now, we've also heard discussion of a BIOS. BIOS is not reflected in this particular diagram here on DX-105.

Can you please briefly what the BIOS is?

A. It is a set instructions and information that helps the memory controller to configure the memory

system at start-up, and depending upon how -- depending upon what types of changes you make to your DRAM, you would have to rewrite the BIOS.

Q. Again -- okay.

Let's look for a moment at the graphics card.

Can you please explain, again very briefly, what SGRAM is?

A. That is an optimized form of DRAM that's used in graphics applications.

Q. And can you please explain what the CPU on the graphics card is?

A. That would be a graphics coprocessor that would integrate to processing functions as well as memory controller, and it would connect directly to the DRAMs instead of connecting to the DRAMs through a memory module.

Q. Now, in general terms, if the specification for the SGRAM were to change, what, if any, would be the potential redesign implications with respect to the graphics card CPU or the graphics card itself?

A. Depending upon the types of changes, it could be anything from trivial modifications to serious modifications, nontrivial modifications.

Q. Let's reflect back if we could on the various alternatives that we discussed this morning. Perhaps

we can start with the alternatives to programmable CAS latency.

And perhaps we could pull up -- oh, we have it -- the demonstrative.

It's the demonstrative you used this morning entitled Alternatives to Programmable CAS Latency, DX-65.

Now, with respect to the various alternatives to programmable CAS latency listed on DX-65, did you consider the degree of disruption, if any, that would be caused by these various alternatives?

A. Yes, I did.

Q. And can you explain in general terms, if the JEDEC SDRAM or DDR SDRAM standard were to be changed today, which of the alternatives listed on DX-65 would be the more disruptive alternatives?

A. Well, in general, some alternatives would be more disruptive than others. Some would cause the DRAM to be incompatible with existing JEDEC-compliant systems and others might be incompatible, and I think we have a demonstrative to show that.

The --

Q. We'll label this as DX-106.

Sorry to interrupt. If you could please explain what DX-106 shows.

A. This shows highlighted in red those alternatives that would create DRAMs that would be incompatible with JEDEC-compliant systems. The other two alternatives 1 and 2 would produce DRAMs that may or may not be compatible with any given JEDEC-compliant system, and that would depend upon the configuration of the system in question.

Q. Let's focus more if we could on the least disruptive of the alternatives, and can you please explain what the redesign implications might be if the JEDEC standard would be changed to use fixed CAS latency parts?

A. Yes. The disruption would be that if you were to put a fixed latency part into a system that expects a different latency and is optimized to use a different latency and would be expecting to use a certain latency, it would expect that the part is able to be programmed to use a certain latency, but the part would have a fixed latency that's something other than what the memory controller would try to program it to be. In that instance, you would have a part that would fail to work.

Q. Is there any redesign that could be done to any other component to ensure that parts with fixed CAS latency would work?

A. Yes. If you changed your memory controller to accommodate systems with different CAS latency and you changed your serial presence detect chip to notify the memory controller that this particular part can't support programmable latency, then you would be able to support it.

Q. Now, with respect to the second alternative, program CAS latency by blowing fuses on the DRAM, I believe that you testified this morning that once the fuses are blown, there's a similarity between the part with the blown fuse and a part with a fixed latency?

A. Yes, I did.

Q. Now, your testimony with respect to the redesign implications of fixed CAS latency, would that then also apply to parts that were -- in which the CAS latency was set by blowing a fuse?

A. Absolutely. Yes.

Q. Now, just in very general terms, again, if a change in the JEDEC standard would be made today, can you compare in general terms the redesign implications for the four alternatives in the red box with the redesign implications that you've just described for fixed CAS latency?

A. Well, if -- actually, could you repeat the question.

Q. Yes.

I was wondering if you could explain in general terms how a change today in the JEDEC standard to use one of the alternatives in the red box, namely alternatives 3, 4, 5 or 6, how the redesign implications of that type of change would compare with the redesign implications you have just described with respect to fixed CAS latency.

A. I see. Yes. The redesign implications would be at least as extensive. You would have to change at least the memory controller in the scenarios as well as other things such as the module, potentially the motherboard, and so forth.

Q. Now, Professor Jacob, what, if any, is the difference between making a change today in how CAS latency is determined as opposed to having selected a different method of determining CAS latency in the early to mid-1990s?

A. These would have been very straightforward designs had they been -- had this been decided the first time around. If you tried to redesign them today, that impacts -- this means that for one thing you throw away design work that has already been done. It would cause you to slip schedules for existing designs. Future designs would not be able to use

already generated designs or, rather, existing designs.

So this would throw off schedules. It would cause, you know, wastage of engineering effort. And if this had been designed in at the beginning, it would have been as straightforward as the present-day design.

Q. What about a comparison of the time, of the timing of the design change today versus design change -- or the selection of alternative in the early to mid-1990s?

A. I'm not sure I understand the question.

Q. What, if any, would be the implication with respect to timing of completion of design of products?

Obviously my question is not clear. I'll just move on then.

A. Okay.

Q. Thank you.

A. Sorry.

Q. Let's turn then to the alternatives that you had identified for programmable burst length.

If we could pull up DX-79, which lists the alternatives that you identified for programmable burst length.

Now, again, did you consider the degree of

disruption that the various alternatives listed in DX-79 would cause if a change were to be made today?

A. Yes, I did.

Q. If a change were to be made today, which of the alternatives listed on DX-79 would be among the more disruptive?

A. In general, some would be -- some would cause more disruption than others, and I believe we have a demonstrative to indicate which.

So for example, the alternatives highlighted in red would be more disruptive than those that are not highlighted in red.

If one were to build a DRAM using one of the alternatives highlighted in red, you would produce a DRAM that's incompatible with present JEDEC-compliant systems.

If one were to use one of the other alternatives that are not highlighted in red, you would produce a part that may or may not be compatible with existing JEDEC-compliant systems, and it would depend upon the system in question.

Q. Let the record reflect that the demonstrative to which Professor Jacob was referring is DX-107 entitled Alternatives to Programmable Burst Length with some items placed in red boxes.

Looking for a moment at the first alternative, for example, which is not one of the ones you identified as more disruptive, using fixed burst length parts, again in order to try to save time, is there any significant difference with respect to the redesign implications today of using fixed burst length parts as opposed to using fixed CAS latency?

A. No. No.

Q. So the discussion you provided a moment ago with respect to design implications of fixed CAS latency would apply here as well?

A. Correct.

Q. Now, what, if any, is the difference between making a change today in how burst length is determined as opposed to JEDEC having selected a different method of determining burst length in the early to mid-1990s?

A. Well, again, if one were to create these designs or introduce these designs into present-day technology, it means that you would throw away existing designs. You would have to retest the new parts with the new designs. This would cause slippage of schedules. This would cause projected arrival times of parts to change.

Whereas if this were known -- if this were --

if one of these alternatives was chosen in the early '90s, it would have been a very straightforward design to accomplish, and you know, it would be just as easy as the present-day system had they designed it in from the beginning.

Q. Let's turn back now to the alternatives that you discussed this morning with respect to dual-edged clocking.

If we could perhaps again pull up the demonstrative, DX-88, that listed the alternatives you identified to dual-edged clocking.

And again with respect to the alternatives listed -- actually I should specify that I think you said that number 7, in your opinion, is really not an alternative, so we'll focus on items 1 through 6.

And with respect to items 1 through 6 on DX-88, did you consider whether some of these alternatives might be more disruptive than others?

A. Yes, I did.

Q. And which alternatives did you identify as being more disruptive?

A. And in general, some would be more disruptive than others, and I believe we have a demonstrative to show.

And in this instance, all of the alternatives

would produce parts that would be incompatible with JEDEC-compliant systems of today. But perhaps the least disruptive of these would be the number 5, doubling the clock frequency.

Q. Why do you say that that would be the least disruptive?

A. Because it wouldn't require additional number of pins. It wouldn't require additional hardware or additional -- or changes to the I/O driver definitions.

Q. What would be the -- or what, if any, would be the redesign implications of switching to alternative number 5, doubling the clock frequency, today in place of use of the dual-edged clock?

A. One would need to change the memory controller.

Q. Why is that?

A. To use a different clock frequency, to use a higher clock frequency, as well as using a different clock chip on the motherboard to provide a higher rate.

Q. In other words, you're saying that the memory controller and the clock chip today would not function with a faster single-edged clock?

A. Correct.

Q. Now, in general terms, how would the redesign implications of alternatives 1, 2, 3, 4 or 6 of DX-108 compare with the redesign implications of number 5, doubling the clock frequency?

A. They would be at least as extensive changes. Some would require more pins. It would require changes to the module. They would require changes to the motherboard, changes to the driver designs. So the changes would be at least as extensive.

Q. Now, what, if any, would be the difference between making a change today to replace dual-edged clocking with one of the alternatives listed in DX-108 as opposed to JEDEC having selected a different method of clocking before the DDR SDRAM standard was finalized?

A. Again, it would have been straightforward to add any of these alternatives or to implement any of these alternatives from day one, and trying to retrofit them into an existing design flow causes disruption in schedules and so you have to throw away existing designs. You can't reuse those designs. And so it's more disruptive if you try to put this into an existing infrastructure.

Q. Now, let's turn back to the alternatives to on-chip DLL that you also testified about this

morning.

We've pulled up DX-94, listing the five alternatives that you identified to on-chip DLL.

Now, with respect to the five alternatives listed on DX-94, again, did you consider whether some of these alternatives would be more disruptive than others?

A. Yes, I did.

Q. And could you please identify which alternatives would be more disruptive.

A. And again, some would produce parts that would be incompatible with JEDEC-compliant systems and some would produce parts that may or may not be incompatible with JEDEC-compliant systems, and I believe we have a demonstrative to show.

So those highlighted in red, alternatives 1, 2, 3 and 4, would produce parts that are incompatible with existing systems, and alternative 5 would produce a part that may or may not be incompatible with existing JEDEC-compliant systems, and it would depend upon the system in question.

Q. Let the record reflect that Professor Jacob was referring to a demonstrative that will be marked as DX-109 with the caption Alternatives to On-Chip PLL/DLL --

JUDGE McGUIRE: Is it 109 or is it 108?

MR. OLIVER: Your Honor, I believe 108 was the alternatives to dual-edged clocking.

JUDGE McGUIRE: Okay. Very good.

BY MR. OLIVER:

Q. Now, Professor Jacob, perhaps you could explain in a bit more detail what, if any, would be the potential redesign implications today of the least disruptive of the alternatives you have identified, relying on the DQS data strobe to provide timing.

A. What this means is one would disable the DLL or eliminate the DLL and expect the memory controller to use DQS to latch the data upon a DRAM read rather than being able to use the global clock to latch the data.

So in a system that already uses the DQS data strobe and disables DLL -- for example, many graphics applications do this -- then this would require no changes. This would be perfectly compatible. But in any system that expects the DRAM to be in sync with the global clock, this would fail to be compatible.

Q. Now, in general terms, how would the redesign implications of the alternatives 1 through 4 highlighted in the red box compare with the redesign implications of relying on the DQS data strobe to

provide timing?

A. In general, they would be at least as -- the changes would be at least as extensive.

Q. Now, what, if any, would be the difference between changing today to one of the alternatives that you have identified to replace on-chip DLL as opposed to JEDEC having selected a different method to capture data at the memory controller before the DDR SDRAM standard were finalized?

A. Again, these are straightforward engineering techniques, and had they been designed in from the beginning, it would have been a very straightforward thing to do as opposed to trying to retrofit them into existing infrastructure today.

Q. Now, we've looked at the redesign implications with respect to each of the technologies at issue individually.

How would you describe the redesign implications, if any, if one were to try to replace all four of these technologies at once?

A. Well, as shown, each of the technologies has a number of alternatives and some of the alternatives would produce parts that, you know, may or may not be compatible and some alternatives produce parts that are clearly incompatible. If one were to replace all of

the technologies in dispute with one of the alternatives, you would produce a DRAM part that would fail to be compatible with any existing JEDEC-compliant system.

Q. And again, what, if any, would be the difference between making a change today to replace all four of the technologies at issue as opposed to JEDEC having selected different technologies at the time that JEDEC was forming the SDRAM and DDR SDRAM standards?

A. Again, today the changes would be more disruptive because it would cause engineers to have to throw away designs. You wouldn't be able to reuse existing designs. Future designs would need be to redesigned from scratch. Designs that are in current process might have to be restarted and which would cause slippage of schedules. As opposed to if any of these had been -- or if all of these technologies had been replaced with alternatives in the beginning, none of this would happen.

MR. OLIVER: Your Honor, I have no further questions of the witness at this time.

JUDGE McGUIRE: Okay. Thank you very much, Mr. Oliver.

Counsel, it's 5:00, so I suggest we go ahead

for the day and break and we'll reconvene here on Tuesday morning at 9:30. Okay?

This hearing is in recess.

(DX Exhibit Numbers 95 through 99 and 104 through 109 were marked for identification.)

(Time noted: 5:02 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 16, 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 16, 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