FEDERAL TRADE COMMISSION

I N D E X

EXHIBITS

UNITED STATES OF AMERICA

FEDERAL TRADE COMMISSION

In the Matter of:	)
Rambus, Inc.	) Docket No. 9302
----------	)

Friday, May 30, 2003

9:15 a.m.

TRIAL VOLUME 19

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: Sally Jo Bowling

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

APPEARANCES:

ON BEHALF OF THE FEDERAL TRADE COMMISSION:

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

ON BEHALF OF THE RESPONDENT:

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

P R O C E E D I N G S

- - - - - -

JUDGE McGUIRE: This hearing is now in order. Counsel, any housekeeping tasks we need to attend to this morning?

MR. ROYALL: I believe Mr. Oliver may have something.

JUDGE McGUIRE: Okay, Mr. Oliver.

MR. OLIVER: Your Honor, there is one exhibit I omitted to move into evidence yesterday afternoon that is CX-798, an email to Richard Barth from Richard Crisp.

JUDGE McGUIRE: Any objection?

MR. STONE: No, Your Honor.

(CX Exhibit Number 798 was admitted into evidence.)

JUDGE McGUIRE: Anything else from either side?

MR. ROYALL: One other thing, Your Honor, this may have been discussed yesterday, because we are trying to get two witnesses on and off today from the west coast, we were thinking if it would be possible to have a truncated lunch break.

JUDGE McGUIRE: Certainly. Certainly.

MR. ROYALL: We can decide later, 30 minutes, 45 minutes.

JUDGE McGUIRE: That's no problem with me,

whatever you all want to do, then that's what we'll do.

MR. ROYALL: Okay.

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

MR. ROYALL: At this time, complaint counsel calls as its next witness Mr. Richard Heye.

JUDGE McGUIRE: All right, Mr. Heye, would you please approach and be sworn in by the court reporter.

Whereupon--

RICHARD HEYE

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

JUDGE McGUIRE: Mr. Heye, have a seat right there if you would.

DIRECT EXAMINATION

BY MR. ROYALL:

Q. Good morning, Mr. Heye.

A. Good morning.

Q. Could I ask you to state your full name for the record, please.

A. My name is Richard Heye.

Q. And where are you currently employed, Mr. Heye?

A. Advanced Micro Devices in Sunnyvale, California.

Q. And is Advanced Micro Devices also known by the acronym AMD?

A. Yes, sir.

Q. Can you briefly explain the general nature of the business of AMD?

A. Yes, AMD's primary business is selling microprocessors into what is referred to as the x86 market which includes desktop PCs, mobile computers, which are normally laptops, and server workstations.

Q. Does AMD make products other than microprocessors?

A. Yes, it does, AMD is also involved in the flash memory business and it just recently acquired a company called Alchemy which makes microprocessors, that's a small business.

Q. Is the microprocessor part of AMD's business the largest aspect of AMD's business?

A. Yes, that's the predominant business.

Q. How long have you been employed by AMD?

A. A little over five years.

Q. And what is your current title or position with the company?

A. My current title is vice president and general manager of the microprocessor business unit.

Q. How long have you held that position?

A. About two months.

Q. And generally speaking, what are your job

responsibilities in that position?

A. The primary job responsibility is really managing the business of all three segments that we had mentioned earlier which were the desktop, mobile and server divisions. In addition to that, I'm responsible for all the platform engineering associated with delivering the infrastructure to support those three businesses and I'm actually responsible for the engineering infrastructure for the fulfillment as well.

Q. And roughly how many employees report to you in your current position?

A. A little over a thousand.

Q. Immediately before taking this position, and did you say it was only the last few months that you've had the title that you described?

A. Yes, sir.

Q. And immediately before that, what position did you hold within AMD?

A. Vice president of platform engineering and infrastructure.

Q. And how long did you hold that position?

A. Roughly two years.

Q. And how did your responsibilities in that position differ from your current responsibilities?

A. My responsibility at that time was focusing just

on the engineering aspect of delivering the infrastructure and what we call the fulfillment of the infrastructure which is making sure that all of the components that are required to sell microprocessors are available in the market.

Q. And did you have a smaller number of employees reporting to you in that position?

A. Yes, in that capacity, I had roughly 550 folks.

Q. Now, let me ask you a few questions about your educational background and your work experience before you joined AMD. First of all, where did you attend college?

A. Washington University in St. Louis, Missouri.

Q. And what did you study?

A. Engineering.

Q. Do you hold any degrees from Washington University?

A. Yes, I have three degrees, a Bachelor's in engineering, electrical engineering, Bachelor's in computer science, and a Master's in computer science.

Q. When did you get the Master's degree?

A. 1981.

Q. And when you completed your Master's degree, did you take a job?

A. Yes, sir.

Q. And with what company?
A. I joined Digital Equipment Corporation.
Q. And how long were you employed by Digital?
A. Approximately 11 years.
Q. Eleven years you said?
A. Eleven years, yes.
Q. Digital is no longer in existence today. Is that right?
A. That's correct, it's no longer in existence. It was acquired by Compaq, which was acquired by Hewlett Packard.
Q. At the time that you worked for Digital, how would you describe the basic nature of the company's business?
A. At the time, Digital was the largest manufacturer of minicomputers, in fact they developed the minicomputer market. So, they were pretty much number two after IBM in terms of size of computer makers.
Q. And generally speaking, what was the nature of your work for Digital?
A. When I joined Digital, they had just begun getting into the semiconductor business, and that was all in-house microprocessors, but I was on the original design team making microprocessors for Digital

Equipment's internal minicomputers.

Q. You said you were with Digital for 11 years?

A. Yes, sir.

Q. And you left in the early nineties some time?

A. '91.

Q. And where did you go from there?

A. Apple Computer.

Q. To Apple?

A. Yes, sir.

Q. How long were you employed by Apple?

A. Six years.

Q. And when you left Apple, was that '97?

A. '97, correct.

Q. And was it at that time that you joined AMD?

A. Correct, yes.

Q. In the time that you were employed by Apple during those six years, how would you describe the work of the company or the nature of the company's business?

A. Apple had multiple business units, the one I was most familiar with, and was involved with, was the Macintosh computer.

Q. And what type of work did you do for Apple?

A. When I first joined I was hired as what was called a hardware lead for developing Macintoshes. And Macintosh ironically got cancelled six months after I

joined the company, but after that I was in charge of all microdevelopment.

Q. Before going on to your work at AMD starting in the late nineties, let's take a moment to focus on what you did at Apple. You just mentioned in your last answer something called Macintosh and most of us probably know what that is, but just to be clear, can you describe what the Macintosh product is or was?

A. Sure. The Macintosh was and is a PC that competed with at the time it was euphemistically referred to as the IBM PC market. And the Macintosh was the first to have what is now known as the Windows interface and it's been a relatively niche player, I don't know that it ever surpassed 10 percent market share, but it's always had an avid following.

Q. And at the time that you arrived at Apple in the early 1990s, was the company already producing Macintosh computers?

A. Yes, sir.

Q. I think you said that the highest market share that you recall Apple having in the PC market was did you say 10 percent?

A. I believe that to be true.

Q. During the six years that you spent at Apple, was your work always somehow related to the Macintosh

product line?

A. Yes.

Q. And what responsibilities did you have with respect to the Macintosh product line?

A. When I joined the Macintosh group, they were focusing on -- they were using a microprocessor called the 68,000, and pretty shortly after my arrival, Apple made a strategic decision to stop using the 68,000-based microprocessor and begin to use what is now referred to as the PC -- Power PC microprocessor. And given my background at Digital being a microprocessor designer, I was one of the early members of the technical team that helped define the Power PC microprocessor for Apple.

Q. And so, this Power PC microprocessor was at the time that you arrived, it was the processor that was intended to be used in future generations of the Macintosh. Is that right?

A. That's correct, yes.

Q. And what precisely was your role with respect to the Power PC microprocessors?

A. Let me give you a real quick history on Power PC, to give you a little context in understanding what I did. The Power PC was based on a architecture that was specified by IBM called Power. IBM and Motorola combined to generate the Power PC. The first Power PC

processor used in a Macintosh was based off an existing IBM design of a 601. Then IBM and Motorola formed a joint design team which was referred to as Somerset. Somerset was physically in Austin, Texas.

The Somerset design team was tasked to generate actually three microprocessors, two for the Macintosh and one for a server. And the two processors that AMD was interested in were called the 603 and the 604. And the products that I would -- the Macintosh that I would be ultimately responsible for designing initially were using the 604 microprocessor, but because of my background, I was one of the initial interfacers from Somerset to Apple. So when they had a question about trade-offs from the 603 to the 604 initially I was one of the people that they turned to for resolution of these questions.

I should also mention parenthetically that over time Apple hired eight folks to physically reside in Austin, Texas, and once they built that team down there, I spent less and less time dealing with Somerset.

Q. You mentioned the two Power PC microprocessors that you were involved in, 603 and 604.

A. Correct.

Q. Very generally, how did those products differ?

A. Sure, the 604 was targeted for the high end of

the Macintosh, this would be used for what we call the power users, and the 603 was targeted for the value segment, and that would be more for educational market and low cost Macintoshes, which at the time was $2,000, which I hadn't mentioned today.

Q. You mentioned that you dealt with or interfaced with the IBM and Motorola people at the Somerset facility in Austin. Were you involved in negotiating the pricing or the terms on which Apple purchased the Power PC microprocessors that they manufactured?

A. I was not involved in the price negotiation, however, there are design considerations a person makes prior to the actual manufacturing of the microprocessor that has direct cost ramifications. So, on the design side, I was very influential. On the actual negotiating the price, I was not involved in that.

Q. Let me ask you to elaborate on that. You said that you were involved in design decisions that affected cost, the cost of these microprocessors. Can you be more specific about how your involvement in design issues related to cost?

A. Sure. First off, cost for a microprocessor is directly proportional to the die size, which is how physically large the chip is. The bigger the chip, the more expensive the chip, the smaller the chip, the

cheaper it is. So, what we -- discussions would always be around the trade-offs of performance versus die size. Because typically the more stuff you put in a microprocessor, the better it performs, at the same time, it would be bigger, therefore more costly. So, the trade-offs involved, should I put more stuff in there or should I make it cheaper and thereby take stuff out. And I was involved in the stuff discussions.

Q. And from your standpoint representing Apple, did you care about die size issues, was that important to you?

A. Absolutely. For example, the 604, we would err on the side of putting more stuff in there, at the expense of cost, while on the 603 we would err on the side of putting less stuff in there to make it cheaper.

Q. And is that because of the different market segments that those two microprocessors were targeted for?

A. Exactly.

Q. And the 603 was targeted more for the low end market. Is that what you said?

A. Correct.

Q. From Apple's standpoint, was there any down side to decreasing the die size of the Power PC microprocessors?

A. Well, the down side would be that the 603 did not have the same performance as the 604. You know, in the perfect world, you want the absolute fastest microprocessor for cheaper cost, although physically that's not allowed to happen. So, you're constantly making business decisions to differentiate do you want a low cost microprocessor with less performance or a high performing microprocessor with more performance. And the judgment of that is what I was involved with, so that's just the standard trade-offs we make.

Q. Are you familiar, Mr. Heye, with something called the PCI Bus?

A. Yes, I am.

Q. Can you explain to us what the PCI Bus is?

A. The PCI Bus is the standard bus that is used to interconnect components on a motherboard and it's also a bus used to plug in third party peripheral cards into PCs and laptops.

Q. Was the PCI Bus used in the Macintosh computers that you were involved in developing while at Apple?

A. Yes. The first PCI Bus used in a Macintosh was the Macintosh that I was designing.

Q. And were you personally involved in the decision to use the PCI Bus in that version of the Macintosh?

A. Yes, in fact, I drove that decision within

Apple.

Q. You drove that decision, meaning what?

A. Meaning that the team I worked with and myself determined that the PCI Bus was the right solution for the Macintosh. We then presented that to our senior management and got buy-in from our management to proceed down the PCI Bus. Which at the time, by the way, was a rapid departure for Apple.

Q. Now, let's talk about that just for a moment. Just to be clear, I think you made mention in your earlier answer, but what precisely, not too precisely, but what is the role that is played or the function of the PCI Bus as it was incorporated into the Macintosh computer?

A. So, think about a PC today. So, when you buy a PC today, a lot of times you may want to add, say, to the hard drive, Ethernet, a whole host of possible third party devices, and the question is, well, how do you plug those devices in, and a lot of times that requires plugging in some extra electronics, and those electronics reside in a card which then plugs into the PCI Bus. And so the PCI Bus is referred to as an expansion bus and also which allows a manufacturer PC to put in a bunch of these what are referred to as slots or connectors on the PCB, which then allows the customer to

plug in cards.

And it's really important to have that, because it's impossible to design a PC that can meet the needs of every individual. So, there has to be some flexibility in the way you configure a PC, both in production and in the after market. And the PCI Bus is the one we chose for future generation Macintoshes.

Q. And at the time that you and your team were making that choice, was there an alternative bus that was available to Apple to use in lieu of PCI Bus?

A. At the time of the PCI Bus, there was a basic transition occurring in the industry --

JUDGE McGUIRE: Okay, we keep talking about at the time of the PCI Bus, what time was this for the record, and what year, what point in time are we talking about again?

BY MR. ROYALL:

Q. Well, can you explain, Mr. Heye, during your six-year tenure at Apple, at what point in time were you involved in making judgments about whether or not to use the PCI Bus that you've described?

A. That was in the early '90s, so roughly '92-'93 time frame.

JUDGE McGUIRE: All right.

BY MR. ROYALL:

Q. And to go back to the earlier question, at that time that you were considering this decision, were there alternatives to the PCI Bus that you were also considering?

A. Yes, again as I was saying, in the early nineties, Apple Computer had an existing bus called the Nubus, spelled N U B U S. And the Nubus was the expansion bus for the Macintosh. The Nubus was running out of performance capabilities. In other words, the needs for the Nubus exceeded the capability of the Nubus. So, Apple at the time was in the process of redesigning a future generation bus to replace the Nubus, and that was an internal proprietary bus.

Q. And the Nubus that you mentioned, again, in case the court reporter didn't get it, I think it's N U B U S?

A. Correct.

Q. Is that right? Had a version of that proprietary bus been used in earlier models of the Macintosh?

A. Correct.

Q. So then the question is, you were working on future generations and the question was did you continue using a next generation of the Nubus in future generations of the Macintosh?

A. Correct, that's right.

Q. When you arrived at Apple in the early 1990s, was there already a plan in place to use this Nubus technology or a new generation of it in future generations of the Macintosh?

A. Yes, there was a design team actually working on the implementation and actually a test chip associated with the future generation Nubus.

Q. And once you joined Apple, did you at some point in time in the early 1990s develop your own views as to the quality of the next generation Nubus technology?

A. From a technical point of view, it was a great design. At that time, in the early nineties, it would have far and away been the most superior next generation PC interconnect bus on the market. It was better than the existing PC Solutions, it was better than the existing Apple solutions, and quite frankly, it was technically better than the PCI Bus.

Q. In your mind, as you were considering decisions, design decisions about future Macintoshes, did you see any down sides or drawbacks to Apple using the next generation Nubus technology?

A. The biggest drawback I saw was that, you know, Apple was not a majority player in the market, and what you find is that as -- if you're different from the

competitor and the competitor has a much, much larger market share, what happens is that it is much more difficult for you, the smaller share person, to get the -- in this case the peripheral cards that you want for your customers, because again, if I'm a card manufacturer, and I have a choice in designing and the manufacturing of a card, am I going to do that card for 90 percent of the market, or am I going to do that card for the 10 percent of the market?

Well, business would dictate, nine times out of ten you're going to do it for the 90 percent of the market. And so the concern I was having that if the PCI Bus was going to become the dominant industry standard bus, which at that time was not clear at all, by the way, but if it was going to become a dominant industry standard bus, Apple would always get a competitive disadvantage in our ability to get third party solutions for the Macintosh.

Q. Now, I believe you mentioned with respect to the PCI Bus that it was a standardized bus. Is that right?

A. Well, the PCI Bus was initially developed by Intel, and then Intel formed a consortium around the PCI Bus. So, Intel and nine other companies formed this consortium and they specified the PCI Bus. And then they made the PCI Bus open to anyone in the high-tech

community as long as you paid a certain fee, which I don't recall what it was, and you could join the -- they call it PCI SIG, which was for a special interest group.

Q. Was Apple a participant in the PCI Bus consortium that you just mentioned?

A. Not at the time when we were doing the analysis. And in fact, what had occurred was the following: The PCI Bus consortium that was already founded and they were pretty much closing the specification of the PCI Bus. At that moment in time, that's when Apple decided that we really wanted to get involved in the PCI Bus. Our concern was that we wanted to be involved in the inner sanctum, you know, those select eight companies plus Intel. And the way the bylaws were written was that you had to get voted into this inner sanctum and the votes happened on an annual basis and we were applying in the interim so we had to wait roughly an entire year. And that was unacceptable to Apple.

So, I personally went down and presented our case to the executive committee of the PCI group and the deal I wanted to strike was, okay, let Apple be a member of the executive committee. We can't vote, but we want to be a member. And the reason that was so important to Apple was that we wanted to make sure that any decisions, any early decisions that they were even

thinking about, we wanted to know about. That was very important to Apple that we understood every possible direction or nuance that PCI Bus may have in the future.

Q. And let me ask you, on that point, why was it important to Apple to have that role or that involvement in the PCI Bus consortium?

A. Because at the time Apple designed all of its chipsets for its Macintosh internally. So, even though you don't think of Apple as a silicon -- as a chip designer company, the reality was in the early nineties, every piece of custom silicon in that Macintosh were designed by Apple engineers. And so when you design silicon, you have got to have that information fairly early, as silicon has a long lead time.

Q. You mentioned earlier that you drove the decision ultimately to go with the PCI Bus over the proprietary Nubus technology in future Macintoshes.

A. Correct.

Q. Did you have to get management approval for that decision?

A. Yes, I did.

Q. And were you able to get management approval for that decision?

A. Yes, I was. It was actually contingent on me getting on the PCI executive committee, which I did.

Q. Was it at all a controversial decision within Apple?

A. Very.

Q. Why is that?

A. Because it was -- that was one of the first times that Apple conceded being the absolute best versus industry standard. Because they really did have a superior solution, and the engineering people felt, you know, this PCI Bus isn't as good as we're designing, so why would you ever take something that's not as good and bring it to market? And the counterargument is what I just said, which is at the end of the day, being different in a commodity market is a bad thing, and you will actually lose your competitive advantage over time. And so it was more important to be a standard and try to influence the standard as opposed to being different and better.

Q. So, would it be fair to say that it was your judgment at the time that it was in Apple's business interest to go with the industry standard over a proprietary technology, even though that proprietary technology was deemed to be superior?

A. Yes.

Q. And you mentioned commodities.

A. Yes.

Q. How does the -- does the issue of commodities come into play in your thinking or how did it come into play in your thinking in that time period?

A. Well, commodities in particular in this instance might have been a misnomer, but what I'm trying to say is when you look at the entire industry that develops third party cards for PCI Bus, they would take offense to me calling them commodities, but if you look at it from a global point of view, these folks are developing solutions for the mass market, and they're going to be very focused on cost. And market share. And to the extent that's what they're focusing on, they're always going to defer to the higher volume segment before they go to the low volume segment. And in which case, which is a similar characteristic of the commodity market, by the way. Which is why I use the word commodity in terms of defining the PCI cards.

Q. During your tenure at Apple, in the mid-1990s, again, I think you said was it '91 to '97?

A. Correct.

Q. That you were there. During that time period, to your knowledge, did Apple participate in any other standard-setting activities besides the PCI consortium?

A. Lots. You know, I couldn't even begin to remember what they all were. I mean, certainly we had

members in the JEDEC committee, there were groups -- there was a whole set of committees on mobile computing, which I know nothing about, but I know we were involved in the standards there. You know, in the high-tech world, there are standards, you can't turn around without there being a standard-setting body, but I think there were a number of people involved in the IEEE standard body, for example. So we had people in the IEEE, we had folks in JEDEC, and I'm sure there are dozens, quite literally, special interest groups in different subsegments of the PC mobile market that we were involved with.

Q. With respect to JEDEC, do you know what aspects of JEDEC's work Apple participated in during the time frame that you were involved with the company?

A. Well, certainly we had a member of our -- a member of Apple was involved in the memory JEDEC committee.

Q. The memory committee?

A. Yeah.

Q. Do you have an understanding as to why Apple chose to send a representative to the JEDEC memory committee?

A. Yes. And again, a quick background. It's hard to believe right now, but in the early nineties, Apple

was the largest consumer of semiconductors in the world outside of IBM. Because at that time, Macintosh was actually from an individual SKU point of view were some of the highest selling individuals in the world. So, when you're in that kind of volume in the '90s, you have got to ride the commodity curve. And by that I mean you have to ensure that your products, that is commodity parts, are using the parts that are the highest available lowest cost parts. And so, in memory, by the way, being a huge component from a pricing point of view, with respect to that memory wasn't all that cheap, you had to be sure you knew where the memory market was going. And so we had a person whose job was not only to be a member of JEDEC, but he would go literally around the world, I think at least twice a year, and talk to every memory vendor to understand the memory roadmaps.

Q. Let me ask you, to follow up on just your last point there, you mentioned memory roadmaps. Can you explain what that is or what you're referring to?

A. So, every memory in itself is a commodity. It's also a pretty high-tech device. And so each memory company has different times where they're going to introduce new memory technology, they're constantly cost reducing memory technology, so they have roadmaps on when new technology is going to be available, when cost-

reduced memory technology is going to become available. It's important for us to understand when the next generation of new memory is going to be available.

Memory comes, each one may have the same number of bits, sometimes they come by form factors, by 2s, by 4s, by 8s, and each memory vendor would have slightly different roadmaps. And again, what was important was we had a -- we being Apple, this wasn't so much my job, but we being Apple, from a supply-based management point of view, had to make sure that whatever was needed by Macintosh, that all those memory vendors could supply our needs in the time frame of importance.

Q. To your knowledge, was Apple's -- did Apple have a goal in participating in JEDEC to try to influence the technical content of the JEDEC memory standards?

A. Not really. I mean, our -- the primary goal for our participation was really to be aware of what was going to happen. Because the person who actually attended the JEDEC meeting, wasn't an engineer. Actually --

Q. Can I stop you there. Do you recall his name?

A. I think it was Mike Pierson. And Mike actually belonged to this -- to the group called actually component --

JUDGE McGUIRE: Okay, now what's the question

here? I'm losing track and I don't want him simply testifying when I'm not sure what the question is.

BY MR. ROYALL:

Q. The question, again, was to your knowledge, did Apple participate in JEDEC's memory committee with --

JUDGE McGUIRE: Okay, I think he's answered that. Let's go on to the next question. I mean, I'm not trying to cut you off, Mr. Royall, but I don't want to hear, you know, five-minute answers. I want you to ask him your next question and then perhaps you can still go where you're trying to go.

MR. ROYALL: I understand, Your Honor, I just interrupted him to just make clear, he referred to the person who was involved in JEDEC and I wanted to get the name out.

JUDGE McGUIRE: All right, that's fine. All right, go ahead.

BY MR. ROYALL:

Q. Do you want to complete that answer now that you've identified the name of Mr. Pierson?

A. So, Mike had two jobs. One was to work with the supply-based community to make sure that they could procure the parts at the best possible prices and he also would feed information to the engineer so when we designed the next generation chipsets, we would want to

know which memory technology we should design our parts for.

Q. Now, in your role at Apple, you personally, did you become involved in issues relating to the evaluation and selection of what type of computer memory to use in the Macintosh?

A. Yes.

Q. Can you explain how you became involved in those types of issues?

A. So that, again, my team was responsible for designing the future generation Macintoshes, so one thing you should do when you're designing the Macintosh is you would sit down with Mike and say, okay, Mike, we're issuing the Macintosh in 1994, what memory is going to be available in the 1994 time frame, Mike would tell us and we would design our Macintosh to work with the memory that would be available in 1994, for example.

Q. Who designed the chipsets that were used in the Macintosh computers that you were involved in developing?

A. Apple Engineering.

Q. And did Apple also manufacture those chipsets?

A. We worked with Texas Instruments and VLSI and those two companies would actually physical manufacture and test the parts.

Q. Before designing the chipsets for the Macintosh computers, was it necessary to make a decision of what type of memory to use in those Macintosh computers?

A. Yes, you had, in other words, memory design -- you had to know what memory design to design those chips for.

Q. And if you recall, what type of memory did Apple select for the Macintosh computers that you were involved in designing?

A. It was asynch DRAM, I don't remember the exact frequencies.

Q. To your knowledge, did Apple during the time period that you were with the company develop its own proprietary memory products for the Macintosh?

A. Never.

Q. Do you have an understanding as to why Apple did not consider that?

A. Even Apple, which would like to design everything themselves, knew that there were certain commodities you just couldn't influence or try to change. So, Apple had no interest in changing the memory standards. You know, our value -- the Apple value-added had nothing to do with memory, so we stuck to the value-added areas that we understood best, which were, you know, user interfaces and hardware.

Q. Were you involved while at Apple in decisions of which memory supplier or suppliers to do business with?

A. Well, Apple was so big at the time, we had to have a minimum of six suppliers for any given memory technology.

Q. And why was that?

A. Again, availability is very, very important and when you have a commodity like memory, you know, if you don't get the memory, you can't chip your Mac, you can't chip your Macintosh, you're out of business. And because the volume of memory is so great, Apple thought it was very, very important to have multiple suppliers. Again, the DRAM industry is very cyclical, so there are years when we have tremendous supply glut where you can't give it away and there's some years when there's a supply trough, and when there's a trough, you have to make sure you have enough suppliers.

Q. Let's return to your work at AMD, and I think you said you began working for AMD in 1997. Is that correct?

A. That's correct.

Q. And you also mentioned earlier that the primary business of AMD is the microprocessor business. Just to be clear, can you give a short explanation of what a microprocessor is?

A. The best way to think of a microprocessor is the brains of the computer, so when you want to think about how fast you can add, subtract, multiple, divide, that's done by the microprocessor. When people talk about graphic images and watching things spin around and move quickly, a lot of that is all done by a microprocessor. It's just the brains.

Q. What types of products are AMD's microprocessors used in?

A. Again, they're used in desktops, and in mobile computers, and within the last two or three years, Apple has been involved in the server market.

Q. Who are AMD's principal customers?

A. AMD sells to every major OEM with the exception of Dell, and we have a large what's called white box market, which is the -- these are guys who sell one, two, three, four, five, six computers and they buy all their microprocessors through distribution channels.

Q. Roughly speaking, what portion of your customer base is reflected by the PC, the OEM manufacturers as opposed to the white box manufacturers?

A. At this moment in time, about 30 percent of our business is in the white box market.

Q. Who does AMD compete with in the sale of microprocessors?

A. Intel.

Q. Other than AMD and Intel, are there any other firms that have a material presence or share in that marketplace?

A. Not a material, there are a few companies, but they're less than one percent share.

Q. Less than one percent?

A. Yeah.

Q. Do you have a rough estimate as of today of Intel's and AMD's respective shares in the microprocessor business?

A. We're roughly 20 percent of the desktop business, and about 10 percent of the mobile business. So our average share of all microprocessors sold is probably about 17 percent.

Q. In case you didn't make it clear earlier, when you refer to mobile, is that in part in reference to laptop computers?

A. Yes, everybody on that -- on these two tables have laptops, those are considered mobile computers, yes.

Q. Do you know how the market share splits between AMD and Intel have varied over, say, the past four or five years?

A. In the last -- since I've been at Apple, our

desktop market share has been pretty consistently increasing, although it's like the stock market, we have good quarters and bad quarters. For example, we're about 19 percent share right now, we've been as high as 23, but when I started we were like 14 percent. So, we've been on the desktop pretty much over time increasing share on Intel. And in the mobile work, pretty much flat, we go up or down plus or minus five percent in any given quarter.

Q. In the time that you've been with AMD for the past I guess five or six years.

A. Yes.

Q. Has the company developed or sold more than one generation of microprocessor products?

A. Yeah, in the time frame that I've worked at AMD, they've sold three.

Q. Three different generations?

A. Yes, sir.

Q. Let's take a moment and just make clear what generations of products -- if I can use the easel, Your Honor.

JUDGE McGUIRE: Yes.

BY MR. ROYALL:

Q. Forgive my handwriting. All right, you mentioned that there are three different generations of

AMD microprocessors that you've been involved in in the years that you've been at AMD. What was the first generation microprocessor that you had some involvement with?

A. The K-6.

Q. When you joined AMD in 1997, was the K-6 in development?

A. Actually it just started production.

Q. So, the development was already complete?

A. Yes, sir.

Q. And do you know when the development on the K-6 started?

A. The short answer is we acquired the K-6, so it really wasn't developed by AMD, there was a start-up in the Valley called NexGen and AMD acquired NexGen, and so it was done by acquisition. And I'm not sure quite frankly how long the K-6 was in development with NexGen.

Q. Have you heard the term "launch" in connection with microprocessors?

A. Yes.

Q. What does that term mean to you?

A. Launch is when you first start selling your microprocessors into the public.

Q. Into the marketplace?

A. Into the marketplace, yeah.

Q. Do you know when the K-6 product was launched? You may have already said this already.

A. April '97.

Q. '97. And are K-6 microprocessors still being sold today?

A. No.

Q. Do you know ballpark when they ceased to be in the marketplace?

A. 2000, say 2000, 2001.

Q. And did the K-6, AMD's K-6 microprocessor, when it was being marketed and sold, did it compete with any particular counterpart products from Intel?

A. Yes, the Pentium II and then the Pentium III.

Q. Now, when you joined AMD in 1997, was there any development work being done on another generation of microprocessors beyond the K-6?

A. Yes, the K-7.

Q. The K-7. Is K-7 known by any other name?

A. Yes, in the marketplace it's known as Athlon and Duron.

Q. Duron, is that D U R O N?

A. Correct.

Q. Do you know when the development of the K-7 began?

A. '95-'96 time frame.

Q. And do you know when the development of the K-7 product was completed?

A. '99.

Q. When, if you know, was the K-7 -- AMD's K-7 microprocessor launched in terms of being sold in the marketplace?

A. It was sold in '99, in the fall.

Q. And are K-7 microprocessors still being produced and sold today?

A. Yes, sir.

Q. And in terms of competition from Intel, are there any particular Intel microprocessors that the K-7 competes with or has competed with during its life cycle?

A. Pentium III and then the Pentium IV.

Q. Now, you mentioned that there were three different generations of AMD microprocessors that you've had some involvement in. What is the third generation?

A. K-8.

Q. And is the K-8 microprocessor in development today?

A. Actually, the K-8 is shipping today. The version of K-8 for the server market is referred to as Opteron and that was launched in April of this year.

Q. Opteron is O P T E R O N?

A. Correct.

Q. And you said it was shipped in?

A. April of this year, 2003.

Q. And do you know during what time period the K-8 microprocessor was under development?

A. '98 to 2003.

Q. And understanding that this is just beginning to come onto the market, but do you have an understanding as to whether there's any particular Intel-based microprocessor that the K-8 will compete with?

A. Yes, it competes with Xeon, X E O N. And Itanium.

Q. Itanium?

A. Yes.

Q. I T A N I U M?

A. Yes. There's also a desktop version of K-8, but that has not been announced yet. We publicly stated that it will ship in September, and that will compete against Pentium IV.

Q. So, the desktop version of K-8 is something that will be launched some time later this year?

A. Correct.

Q. Does AMD, to your knowledge, have development work ongoing at the present time relating to any future generations of microprocessors?

A. Yes. K-9.

Q. The K-9. Can you say anything about the expectations in terms of when development will be complete or when that product will be launched?

A. My preference is not to since that is not public record at this time. That is all considered proprietary and confidential.

Q. That's fine. Now, generally speaking, how long does it take, or again, based on your experience at AMD over the past five to six years, how long does it take to develop a new generation of AMD microprocessors?

A. Well, as you can tell from the chart, it's roughly five years. And the reason I always kind of hesitate in telling you what the date was, you know, in the beginning there were two or three folks and then it builds up to a team of literally hundreds of people. So, but if you start the clock out when you have the first two or three really smart architects thinking about this, it's about five years.

Q. And again, generally speaking, based on your experience at AMD, what would you say is the typical product life cycle of an AMD microprocessor?

A. About four years. And again, the reason I keep on hesitating is the very last microprocessor sold at the end of life, you know, those things can last quite a

few extra years and in all actuality they're not out of the marketplace, although they may be selling a few hundred thousand, which is a pretty small number.

Q. Over the product life of an AMD microprocessor, does the company -- well, let's take the K-7, for instance.

A. Um-hmm.

Q. Does the company do anything to enhance performance or to upgrade the product over the period of years that it's out in the marketplace?

A. Absolutely. The way to think about it, we have actually two sort of independent efforts. The first effort is we continually improve the technology of our manufacturing process. And what that enables you to do is with the same design, you're able to increase the performance of your microprocessor. So, if you look at both AMD and Intel, for example, we sell Athlon, Intel sells Pentium IV, every quarter we see that they have higher and higher performance. One of the ways you get that is by improving your manufacturing technology.

The other way of doing that is we actually make modifications to the design which, again, allows us to tune the design to better enable higher performing parts.

Q. Are you familiar with the term "system" as it's

used in reference to microprocessors?

A. Yes.

Q. What does that term mean to you?

A. So, again, to me a system is what the end customer buys. So, for example, the PC you buy at the store or the laptops that, you know, folks have on their desk right now, those are end systems.

Q. And does AMD during the life cycle of a given microprocessor do anything to update or enhance the system that goes along with the microprocessor?

A. Absolutely. So, what is common practice in the industry for both AMD and Intel is that we are constantly working -- in AMD's case with our partners -- to develop new technologies that go into the system that enable the overall performance of the system to improve. And that can take its form in terms of different I/O devices, take the form of faster memory, faster front side buses. There's a whole bunch of different changes that you make to assist the microprocessor to perform better.

Q. And what is it that drives or motivates AMD to make those types of enhancements or improvements in a microprocessor system over time?

A. The best way to describe it is by analogy. Another way to look at the microprocessor, if that's

your high performance of the car, and as you continue to improve and improve your engine on the car, you've got to improve the rest of the car to take advantage of that engine. So, for example, you take a Porsche engine, drop it into a Ford pick-up truck, that Ford pick-up truck is just not going to perform like a Porsche that has the exact same engine. So, all the other stuff that Porsche does to make that car run really well, Ford pick-up trucks don't have that.

So, it's to the same extent, is that if the performance of my processor keeps getting better and better and better, that the components of the system which feed the microprocessor need to also get better and better and better to take advantage of the technology of the microprocessor.

Q. Let's go back, Mr. Heye, to the time when you were first hired by AMD, you said it was 1997, I've forgotten the month, if you recall.

A. June.

Q. June?

A. June, yes, sir.

Q. Did AMD hire you to work on any particular microprocessor product?

A. Yes. I was -- my primary reason for joining AMD was to generate an infrastructure for K-7.

Q. Did -- well, strike that.

And the K-7 microprocessor, from what you've told us earlier, was under development at the time you joined the company.

A. That's correct.

Q. Is that right?

A. Yes.

Q. Did the K-7 microprocessor differ in any material respect from the earlier generation, the K-6?

A. It differed in a lot of different areas. The one area that was of most concern to AMD, and the reason I was hired, was that the K-7 did not have the same front side bus as Intel. And the implication of that is as follows: When I take a K-6, I could have walked in any store and had an Intel motherboard or I should phrase it a motherboard that can be used by any Intel processors, I can plug my K-6 into any of those boards into anyplace in the world and it would just work.

On K-7, I couldn't do that. So, K-7 was the first microprocessor that was not pin compatible with an Intel-existing microprocessor. So, all that infrastructure and all those companies designing stuff for Intel-based motherboards, which AMD was leveraging in the K-6 days, went away. And we had to generate our own for K-7.

Q. And do you have an understanding as to why in the mid to late 1990s AMD made a decision to develop a new generation of microprocessors, namely the K-7, that did not have capability with the Intel infrastructure?

A. Yeah. It actually came from a legal document. AMD and Intel were negotiating a cross license patent agreement, and the implication was that AMD was no longer allowed to use Intel front side bus on its future generations of microprocessors. So, we were actually contractually obligated not to use an Intel bus.

Q. And I believe you said that when you were hired by AMD in 1997, you were hired to put in place an infrastructure to support the K-7 microprocessor. Is that right?

A. That's correct.

Q. And is that an infrastructure that did not already exist?

A. Did not exist.

Q. Just to make sure we're very clear, what we're talking about by the term infrastructure and what really that entails and also to gain a better understanding of your role at AMD, I would like to see if you could walk us through and explain to us what constitutes an infrastructure. Could you do that for us?

A. Sure.

Q. And would it be helpful to use the easel?

A. It would be great. If I may.

MR. ROYALL: Your Honor?

JUDGE McGUIRE: All right, go ahead, Mr. Heye. Do you want to mark that as DX-29, just for the record.

MR. ROYALL: Yes, Your Honor, that's where we are?

JUDGE McGUIRE: All right, thank you, Mr. Royall.

(DX Exhibit Number 29 was marked for identification.)

THE WITNESS: Let's start with the microprocessor. So, initially we have a microprocessor and a microprocessor talks in what's called the front side bus. And a bus, by the way, is just a collection of wires that's used to communicate to other components in the PC. So, if the microprocessor is the brains, the chipset is the central nervous system. And there are two -- and we say it's in chipsets because there are two chips associated with it. And these chips have always been called euphemistically the north bridge and the south bridge.

The north bridge has four -- is responsible for communicating to four different subsections of the microprocessor. The first thing is the north bridge

talks to the microprocessor on the front side bus. The north bridge talks to memory. The north bridge talks to graphics. And it turns out the way it talks to graphics is there's also another industry standard bus called AGP, it stands for advanced graphics port. And north bridge also talks to the south bridge, and the way it talks to the south bridge is by our friend the PCI Bus. So, the PCI Bus is actually an industry standard in every PC and every laptop you all have has the PCI Bus in it. So, the north bridge talks to the south bridge.

Now, the south bridge talks to all the components that everybody here is familiar with on their PC today. So, it talks to the keyboard, it talks to the mouse, most of you have what's called USB ports on your PCs now. If you have a digital camera, you plug it into you USB.

The other thing is it has this thing that's called the IDE bus, and I can't remember what that stands for. And the IDE bus is when you have your hard drive, it talks to the hard drive, it talks to your CD, it talks to your DVD. It also has audio, and it also talks to what's called the BIOS.

Q. The BIOS, B I O S?

A. B I O S, and that stands for binary input output software. Basically, the BIOS is a -- is software that

is stored on non-volatile memory. Real briefly, though, all that really means is the following: Here's what it's trying to solve. When the computer is turned off, all the memory is empty. There's nothing there. You turn on the computer, the microprocessor wakes up, and the poor guy goes, what am I supposed to do, who do I talk to.

Well, it turns out the microprocessor goes to north bridge, north bridge goes to the south bridge, the south bridge goes to the BIOS, the BIOS then sends information to the microprocessor, it gives the microprocessor enough information to now go like this to the hard drive, the hard drive -- it tells the hard drive, take the operating system, put it into memory and then you're running.

It's like when you turn on your PC you see all those weird sentences come up in the beginning, that's the BIOS and then that flashes away and you see the OS logo, and the transition is you're talking to BIOS which is talking to the hard drive.

Q. Is all of what you have drawn on this exhibit, which we will mark for identification as DX-30, does that sit on something?

A. That sits on a piece of fiberglass, and that fiberglass is called the printed circuitboard, in fact I

brought one with me, I can show you in a minute, if you like, but it's just on a printed circuitboard and that printed circuitboard is euphemistically called a motherboard. So, every PC has a motherboard. And let me draw out two little things to make you aware. I'll just make it two little blobs, there's two things that you have to feed the microprocessor with, one is voltage, and it turns out there's a special set of circuitry called VRMs, which are voltage regulator modules, and clocks. Believe it or not, microprocessors have to have a clock. And VRMs and clock chips are microprocessor-specific. So, again, when I'm doing Intel compatible, I can leverage these two, when I do my own microprocessor, I have to develop my own clock chip and VRM.

So, all this right here makes up the motherboard. So, when you -- when AMD said, I no longer can use the same front side bus, right here, as Intel, the first thing you have to do is say, I better find me a new chipset, because there's not a single chipset in the world that can talk to my microprocessor. So, the first thing you have to do is somehow get the chipsets.

Q. Let me, if you don't mind, just stop you there. You mentioned that you brought some equipment with you.

A. Yes, sir.

Q. Do you have an actual physical motherboard that corresponds with this?

A. Yes.

Q. Could you just quickly show us that?

A. Yes, sir.

MR. STONE: I thought I was supposed to get my copy.

THE WITNESS: I beg your pardon?

MR. STONE: Yeah, I can take it home, put a case around it and I have my computer.

MR. ROYALL: Well, this may be more antiquated technology, I'm not sure you want to do that.

JUDGE McGUIRE: What's your memory at home, by the way?

THE WITNESS: I'll tell you, I'll swap, you give me your laptop, I'll give you this beautiful thing.

BY MR. ROYALL:

Q. If you can just hold that up, just so we can see and you can explain to the judge how the physical motherboard you have in your hand corresponds with the diagram that you drew.

JUDGE McGUIRE: Can we come up here and Mr. Stone, feel free, if you want to, so I can get a better view of it.

THE WITNESS: Sure. So, this socket --

actually, could you hand me those two black boxes, please.

MR. ROYALL: Sure.

THE WITNESS: Bring me the cards, too. So, the first thing is, this socket right here, the microprocessor socket, and this is the microprocessor. So, again, in that diagram, it's where you plug in the microprocessor goes there. That's the north bridge, that's the south bridge. So, the north bridge has those four heads. So, first, you see those little wires?

JUDGE McGUIRE: Yes.

THE WITNESS: Those wires talk to the north bridge, that's the front side bus. These are memory DIMMs, so again, my memory is in here. So, that's my memory, so I have to plug my memory in here. This right here is the graphics port, this is the graphics card right here, so I plug my graphics card like so. Anyhow, you plug your graphics card in there. So, that's the AGP port. There's even more wires that go from here to these slots, that's the PCI Bus. And then right here is your south bridge and the south bridge talks all sorts of stuff, so that little thing right there is the BIOS, and believe it or not, there's a whole bunch of wires that snake itself around and talk to these connectors.

Just for clarity, this is your voltage regulator

and I believe that's the clock chip. So, that's the whole physical thing. And so, what at the end of the day, I have to generate one motherboard for every microprocessor we sell, but knowing that when you say a motherboard, you're really talking about a whole family of components that go into the motherboard.

JUDGE McGUIRE: All right, thank you.

MR. ROYALL: Now, if you can hold onto these, and you can take your seat, Mr. Heye. Your Honor, we can deal with this later, but Mr. Heye said that he's happy to leave these, they can be marked as demonstratives at a later point.

JUDGE McGUIRE: Okay, fine.

BY MR. ROYALL:

Q. Thank you for that explanation, Mr. Heye, and going back to my earlier questions, in your initial role when you joined AMD, you said it was to create an infrastructure for the K-7 microprocessor.

A. That's correct.

Q. And very briefly, how does what you have drawn on the easel relate to the infrastructure that you described earlier?

A. So, again, my job is to make sure that that whole easel accepts the microprocessor that's available, so when we started selling microprocessors, all that

other stuff was the marketplace. So, if one of those components wasn't in there, you can't sell the motherboard; if you can't sell the motherboard, you can't sell your microprocessor.

Q. Of these various components that you have described, motherboards and various other things, how much of that in terms of the physical equipment does AMD itself manufacture?

A. Only the microprocessor, and let me -- I should be a little clearer on that. AMD at times is involved in the chipset business. So, we have made our own north and south bridges. In fact, when we first launched Athlon, we actually, AMD manufactured what is designed and manufactured as north and south bridge. However, today, for example, every motherboard that is sold for an AMD microprocessor has a third party chipset on it. We are not in the chipset business for Athlon today.

Q. If AMD is not, other than the microprocessor, is not actually manufacturing these various components, does that mean that AMD goes out and purchases these things?

A. We don't buy any of those components. What we have to do is we establish an industry-wide business model with many, many partners, and those partners based on the business model would go off and design north

bridges, BIOSes, motherboards, clock chips, VRMs, and we call it the -- it's a virtual system.

Q. And what do you mean by the term "virtual system?"

A. Well, for example, Intel does its own north bridge, it does its own south bridge, it does its own motherboards. So, Intel is not virtual, they're physical.

Q. So, you're saying that unlike AMD, Intel not only makes microprocessors, but it makes the north bridge and south bridge which together make up the chipset?

A. Um-hmm.

Q. And it makes the motherboard itself?

A. That's right.

Q. Okay.

A. And we made a strategic decision that AMD would not do that. Primarily AMD's value-add in the industry, where we have all the smart engineers and we have tremendous, tremendous microprocessor design teams. And that's our value-added. You know, our value-added is not building or procuring or subcontracting out, you know, motherboard manufacturing factories. We made a determination that quite frankly Taiwan does a better job than Intel in manufacturing motherboards. We, in

fact, made a determination that third party chipset guys are better than AMD and in some areas better than Intel in doing chipsets. So, if we can work with them as partners, we believe that ultimately we have a lower cost, and the original infrastructure. The risk being, of course, that you really have to have good partnerships with all of these folks, because any one of them can cause you trouble if they don't support you.

Q. So, to make sure we understand, when you say that you developed, referring to the K-7 for the moment, the infrastructure for the K-7, by that do you mean that you developed through business relationships with other companies a virtual system that would allow you along with your partners to deliver a motherboard with all of these features to the companies that are your customers, the computer manufacturers?

A. Yes, sir.

Q. Well, can you walk us through what was involved in that process of putting together that virtual system or that inventory for the K-7?

A. Sure. Well, again, like I said, the first step is get the chipsets. So, AMD's chipset business was -- had the following strategy: AMD would develop chipsets for technology that did not currently exist in the PC marketplace. So, for example, the K-7 front side bus

did not exist.

Q. That's this here?

A. Yes, FSB.

Q. And what you said earlier the K-7 by contrast to the K-6 in earlier generations, this front side bus was no longer compatible with the Intel infrastructure?

A. That's correct.

Q. So, you had to develop a new front side bus for --

A. We developed a new front side bus for our microprocessor, and then we had to make sure that there was a family of chipsets to work with that front side bus.

Q. Okay, and can I ask you before -- was there anything other than developing the front-side bus that AMD had to do in terms of building this infrastructure, is there anything else it had to do before developing the chipset?

A. Well, I mean, what you have to do before you develop a chipset is you have to first -- you have to go around that circle of the north bridge, so then you go to graphics. Well, it turns out the AGP bus, there's actually, there's what they call AGP, 2X AGP, 4X AGP and 8X AGP, and the only difference in 2X, 4X and 8X, is the 2X gives you capacity, the 4X --

MADAM REPORTER: Please slow down.

THE WITNESS: I'm sorry, the 8X is really fast. So, what happens is there are different frequencies of the AGP bus, so you have to determine when you do the north bridge, which of these AGP buses to develop.

BY MR. ROYALL:

Q. Other than the front side bus and the AGP bus, is there any other piece of the equation that you have to figure out before you are able to develop the chipset portion of the infrastructure?

A. The next is memory.

Q. And can you explain why you have to develop issues relating to memory before you develop the chipsets?

A. You always want to make sure you're riding the commodity curve. You don't want to be different from what I call the Intel-based systems. So, for example, in the time frame that we did that north bridge, there was a transition going on. As I said, Apple was designing for asynch memory. There was a transition going that they called synchronous memory, synchronous DRAM, and in the industry it's what is known as PC-100. So, our original north bridge was designed to work with PC-100 memory.

Q. Is there a period of time before a new

microprocessor is launched that you need to resolve issues about what type of memory to use in the infrastructure supporting that microprocessor?

A. There are. It's what I would call second order of effect, reason being the microprocessor doesn't talk memory, but there are some characteristics of memory that you would kind of like to know about if you're the microprocessor, and at the time we designed K-7, the thinking was, quite frankly, that K-7 would go from asynchronous DRAM to actually DDR, and so the block size of the transfers between the north bridge and the microprocessor took into account what we thought the actual block size would be for asynchronous DRAM and DDR.

Q. Once you made decisions about these features of the infrastructure, or the front side bus, the graphics or AGP bus, and memory, what is the next step in terms of building out and completing a microprocessor infrastructure or system?

A. Well, so once you have the north bridge -- once you have your chipset worked out, you have to work with the BIOS vendors to make sure that they will -- that their software will work with that chipset. You have to work with the clock chip vendors and the voltage regulator vendors to make sure that their chips will

work with the microprocessors. But the next really big step is going to the motherboard vendors and saying, hey, why don't you guys build a motherboard that works with the AMD microprocessor, even though historically you've never done that before, we think it would be a great idea if you started. And so that was the next big hurdle to overcome was getting motherboards.

Q. And when you say the next hurdle to overcome, are you referring to your actual experience in developing the K-7 infrastructure?

A. Absolutely. Again, the difficulty is you go to Taiwan, these guys have been making motherboards that work with both AMD and Intel microprocessors, because at the time they were the same from their point of view, because their motherboards were really either vendors, and that represented 100 percent share of the PC market. And now here comes -- here I come and saying, hey, AMD has 10 percent share, or 14 percent share, some number, but something a lot less than Intel's, and we say we want you to design a motherboard just for that business. And there's still millions and millions there, there's still a business proposition there. That was my job to actually go to Taiwan and convince them to do that. And there was a lot of concerns in Taiwan. We've never done a chipset before. We had never had our own

infrastructure before, and it took a lot of control and work on AMD's part and actually a lot on the motherboard's part, because they were encouraged by our competitor not to support AMD and to make motherboards for us.

Q. So, were the motherboard manufacturers part of this group of business partners that you mentioned that along with AMD built the virtual system for the K-7?

A. Yes, sir.

Q. On the memory side, were there other companies that you dealt with that were business partners in building that infrastructure?

A. Well, actually, we did a similar thing with like Apple did, we talked to other memory vendors to find out what their roadmaps were, and again, you know, I have a person who reports to me and his job is to keep track of all the memory roadmaps, and so he would -- he was actually the person who told our design team, design for the PC-100 because that's going to be the dominant standard at the time of launch.

Q. How many different memory manufacturers did you deal with in building the infrastructure for the K-7, if you recall?

A. Roughly six.

Q. Did you deal separately with any memory module

manufacturers?

A. Oh, yeah. We -- I don't know how many, but there's -- there are dozens of memory module guys. If I show you, what you plug into your PC isn't chips, you plug in this little what's called a dim, and every memory vendor makes their own DIMMs, but not all DIMMs are made by memory vendors, so there are a lot of third parties who make DIMMs and we work with the three or four big companies to get DIMMs.

Q. So, you mentioned in building the infrastructure for the K-7, you had to work with chipset manufacturers, memory module manufacturers, motherboard manufacturers, what other types of third party companies did you need to partner with or deal with to build the infrastructure for the K-7?

A. Graphics vendors, you had to make sure the graphics drivers worked with the microprocessors, we had to work with BIOS vendors, moving away from the motherboard, you have to worry about heat safes and fans, so we dealt with Heat Sync vendors, we dealt with fan vendors and we had to ensure that the third party power supplies worked with our parts. These are the big several part supplies that's in your PCs, so when you plug your computer into a wall, you don't -- it doesn't just take that electricity and put it in that

motherboard, it actually goes through a box, that box changes the wall voltage into voltage that operates the electronics on the motherboard. So, you have to make sure all that stuff is available.

The good news is, you know, hard drives, CDs, DVDs, we didn't work with those guys because that's an industry standard up and down the line, so we didn't talk to them because that would just work with our boxes and Intel boxes.

Q. In terms of the third party companies that you did work with in building this infrastructure, do you have a ballpark as to how many different business partners AMD had in that process?

A. Yeah, currently right now, we have roughly about 100 partners that are involved in the hardware side. We've got another 100 partners on the software side, but that's a different discussion. But the topic that we're talking about right now, it's 100. But I would say that 25 are critical. You know, like I personally talk to folks at about 25 companies and my team picks up the balance of the other 75.

Q. Once the relationships were in place and you met with these various companies, and were headed in the direction of building this -- building out this infrastructure together with these companies, what was

the actual process in terms of design through completion of a physical motherboard of the sort that you talked about? Were there steps that were involved in designing through completion of motherboards that could then be used in the marketplace?

A. So, step one is you define your chipset. Once you define your chipset, then you in parallel talk to the motherboard vendors to get them to design the motherboard to use your chipset, and you start talking -- actually to the BIOS vendors to writing software to work with that chipset.

So, it's a lot of parallel efforts, but the real sequence of events, the critical path of events is to find the chipset, get the motherboard guy, develop the chipset, manufacture the chipset, put the chipset on the motherboard, debug the motherboard, get the motherboard to market and now you've got a product.

Q. And did you have a team of people working with you within AMD in the course of building that infrastructure in terms developing that infrastructure?

A. Yeah, I had a team and I had to essentially build up the team.

Q. How large was the team?

A. When I started it was about 100 and now it's about 500.

Q. And how long did it take you and your team to develop the infrastructure to support the K-7 microprocessor?

A. About two years.

Q. You mentioned in an earlier answer the word "debug."

A. Yes, sir.

Q. Can you explain what that term refers to?

A. It's pretty rare that the first time you design and then manufacture a complicated piece of silicon, like a microprocessor, like a north bridge or south bridge, that it works perfectly the first time. Odds are against it. So, what happens is you build your motherboard, you manufacture your silicon, you put it all together, you turn it on, and in all likelihood it won't work the way you expect it to work. So then you have to figure out what's wrong and fix it and then the jargon to be used for figuring out what is wrong and fixing it is called debug.

Q. And is the debugging process, is that the same as the testing process, or are those different?

A. Testing is different.

Q. Can you refer to what testing is in referring to the types of components on the motherboard that you've described?

A. Let me define two terms. There's a term called validation and a term called testing. Testing typically refers to ensuring that what you manufactured is correct. So, typically when you manufacture silicon, some of the parts will work, and some of the parts will not work. And you test the manufactured material to determine which parts work and do not work.

The term validation typically means I designed a component and the question is does that design work. So, validation validates the design, testing ensures that the manufacturing process was done properly.

Q. And do you have any incident of how long it took in connection with the K-7 microprocessor to complete these validation and testing steps that you've described?

A. So, again, from the time that the first K-7 showed up, to the time we shipped, was roughly a year. So, it took us a year to debug the microprocessor and the associated chipset and motherboards.

MR. ROYALL: Your Honor, I am at a convenient breaking point, if it would be all right.

JUDGE McGUIRE: Okay, let's take a five-minute break and then come back.

MR. ROYALL: Thank you.

(Whereupon, there was a recess in the

proceedings.)

JUDGE McGUIRE: Please continue, Mr. Royall.

BY MR. ROYALL:

Q. Thank you, Your Honor.

Mr. Heye, I asked you about your involvement in memory-related issues at Apple.

A. Yes.

Q. Now I would like to ask you a little bit about your involvement in memory-related issues since you've been at AMD.

A. Okay.

Q. Have you, in fact, in your work at AMD been involved in making decisions about which type of memory to use in AMD's systems?

A. Yes.

Q. And in your role, how have you become involved in those types of issues?

A. Given the nature of my job, I'm the focal point, as an executive, for bringing decisions regarding memory to our executive staff. So, and specifically in the time frame we're talking about, you know, there's a lot of transitions going on. They were going from PC-100, that's what we ship with, and then there was an evolutionary change to PC-133, then after that is when we had -- we made a decision initially to use Rambus,

and then we made the decision not to use Rambus, and I was involved both in the decision to pick Rambus and I was certainly the driving force in determining not to use Rambus.

Q. And I'll come back to a number of these things, but just to make the record clear, when you said that you were involved in the decision not to use Rambus, were you involved in the decision at that point to use something other than Rambus in the K-7 system?

A. Yes, using what's referred to as a DDR technology.

Q. And you said, I think, in your earlier answer that initially you shipped with PC-100, and when you say shipped, you're referring to the launch of the K-7 microprocessor?

A. That's correct.

Q. And in the initial version of the K-7 system, was it PC-100 memory that was used?

A. Yes.

Q. And to your knowledge, does the term "PC-100" have any relation to synchronous DRAM?

A. PC-100 refers to synchronous DRAM running at 100 megahertz.

Q. And so the initial versions of the K-7 system that were launched in 1999 used 100 megahertz SDRAM

memory. Is that correct?

A. Yes.

Q. And you said something about PC-133?

A. Yes.

Q. What is that?

A. That's synchronous DRAM running at 133 megahertz.

Q. So, after the initial launch of the K-7, at some point, did AMD develop another infrastructure for K-7 that was supported by the 133 megahertz version of synchronous DRAM?

A. Yes, but to be very clear, the chipset we used to do that was not AMD, it was based off the VIA chipset.

Q. And then at some point after that version of the K-7 infrastructure was released into the marketplace, was there another version of the K-7 infrastructure that was developed and released that used DDR SDRAM memory?

A. Yes.

Q. And do you know roughly when that version of the K-7 system was launched or released into the marketplace?

A. Roughly October of 2000.

Q. Now, going back to the initial decision of what type of memory to use in the K-7 system, you've told us

now that the choice was to go with PC-100 SDRAM memory, correct?

A. Yes.

Q. And were you involved in the decision to use that type of memory in the initial launch of the K-7?

A. Yes, again, my -- yes. The design team had to know what to use and we worked with the team to determine which memory would be available in the 1999 time frame, and that team worked for me, so yes, I was involved in the decision.

Q. Do you recall when that decision was made?

A. Probably two years, roughly two years before we shipped. It takes about, you know, from the time you start thinking about a chipset to implementing it, especially when it's brand new like the one for AMD, it's about two years prior to shipping.

Q. So, if I'm understanding you, you're saying that the decision of what type of memory to use in the initial version of the K-7 system that was launched in 1999, that decision was made roughly two years before the launch of that version of the system?

A. Yes.

Q. In making decisions about what type of memory to use in support of an AMD microprocessor, do you consult with people within AMD?

A. Well, yes, it's a collaborative effort.

Q. And what types of people do you talk with or does your team consult with in making the decision -- within AMD -- in making decisions of what type of memory to use?

A. Well, it's pretty much between the engineering community and again, I have a small team who is responsible for tracking the memory trends, and it would be those two folks, those two teams working together to decide what made the most amount of sense for implementing the memory, the chipset.

Q. And is it important to you in your capacity in making judgments about what type of memory to use to have information and input from those types of people within the company?

A. Absolutely.

Q. And why is that?

A. Well, again, somebody in the company has to know what's going on in the memory industry, and again, those folks, it's not the engineers, so again, you have to have a group of folks who track what the memory community is going to be doing, and by community I mean they have to be following the roadmaps of all the large memory corporations, and that information has to then be worked with the engineering team to determine what is

the best solution at the time of launch for any given chipset.

Q. In addition to consulting with others within AMD, when you and your team make decisions about memory selection for AMD microprocessors, do you also consult with any third parties outside of AMD?

A. Again, obviously the memory vendors. We talk to our third party chipset partners, and quite frankly, we look at Intel. You know, Intel has these things called the Intel developer forums, and they also specify memory roadmaps and we look at what Intel is saying as well.

Q. Is it important to you in making memory selection decisions to consider the views or the input of memory manufacturers?

A. Absolutely.

Q. And why is that important?

A. Again, we have to track their roadmaps. At the end of the day, it's the memory manufacturers who actually deliver the DRAMs to market, and it's critical to understand what they believe is going to be the commodity part to the marketplace, because quite frankly, they're the ones delivering the parts.

So, you know, no matter what AMD and Intel says, at the end of the day, if the DRAM guys don't want to deliver it, they're not going to deliver it. So, you

really have to know what they're thinking.

Q. And do you have to know what they're thinking in regard to future production of memory?

A. Absolutely.

Q. And why is that?

A. Same argument, it's a two-year lead time. So, if I was designing a chipset today and I looked and said what's shipping today, odds are two years from now it's going to be different. You know it's going to be different. So, the question is, if I need to design something two years from now -- if I need to design something today that's going to ship two years from now, what type is that going to be, and you have to find that out. And the answer is you have to ask, and who do you ask? Memory guys.

Q. And in 1997 when you made the decision to go with 100 megahertz SDRAM in the initial launch of the K-7, were you convinced that that would be the dominant commodity product in the marketplace in 1999?

A. Yes.

Q. And was that based in part on input you were getting from memory manufacturers?

A. Yes. That was just very uncontroversial. That was conventional wisdom. Everybody was pretty confident that's the way it was going to go and Intel was making a

big push for PC-100, so it was pretty clear in the industry.

Q. To your knowledge, when the K-7 microprocessor was developed in the time period that you described earlier, 1995 to 1999, did thoughts about what memory would be used with K-7 influence in any way the development of that microprocessor?

A. Yes. Like I said earlier, we -- the design team determined what they called a block size where the memory transfers from the north bridge to the chipset -- the north bridge to the microprocessor, and they made some assumptions about the memory technology, and the assumptions they made at that time was it was going to go Sync DRAM to DDR.

Q. So, you're saying that when K-7 was developed, the design team on the microprocessor side was assuming that in the future after launching with an SDRAM memory in the future, K-7 would migrate to DDR memory?

A. Yes.

Q. And did you in the 1997 time frame agree that DDR memory would be used in future generations of K-7 microprocessor after the initial launch?

A. I'm sorry, what date did you ask me that question?

Q. In 1997, or some time shortly thereafter, that

is within the first year, let's say, after you joined AMD in 1997, did you agree with the assumption that the K-7 microprocessor would start with SDRAM memory and then that the system would be transitioned at some point in the future to DDR memory?

A. No. In the '97 time frame, we -- and when I say we, it included myself and my team which included folks in engineering and the folks who tracked infrastructure, the memory guys, we were thinking more and more that we should go Rambus. Intel had come out and said that Rambus was going to be the next generation high-speed memory as opposed to DDR. The memory community was saying that they were going to do Rambus.

At the same time, if you -- the DDR standard was not coming together as quickly as people would have thought. So, it was languishing in JEDEC. So, we made the decision that we should go Rambus. The other thing, too, I should point out, is that in dealing with Rambus, one nice thing was that when you purchased -- when you gave -- you had to pay Rambus a fair amount of money to get the license to use the Rambus technology. One of the things you got in return for all those dollars was they actually gave you a design that you could plug into your north bridge. So, we got -- we were able to offload some of our design work from internal design to

external design. So, it seemed like a pretty good deal at the time. And I personally negotiated the Rambus contract.

Q. Let's step back to make sure we're clear. First of all about time frame.

A. Okay.

Q. You said that you and your team made a decision at some point in time, if I'm understanding you correctly, that after the initial launch of the K-7 with SDRAM memory, that the next step in terms of memory supporting K-7 should be Rambus memory. Is that right?

A. That's correct.

Q. And what time frame was it that you and your team came to that judgment?

A. Mid-'97. I think I'm pretty sure I signed the contract -- well, I negotiated the contract, actually, I think Atiq Raza signed it, our president at the time, but I believe the contract was completed in December of '97.

Q. When you and your team made the judgment that starting with SDRAM that future versions of the K-7 system should be supported by RDRAM, was it your belief that in the future RDRAM would become the predominant commodity memory in the marketplace?

A. Correct, yes.

Q. And what was that belief based on?

A. Again, it gets back to the first principles, which is does the end user experience or perceive betterness, you know, improved quality, improved performance, by going to faster memory. And the answer was pretty demonstratively yes. So, it was very clear that both Intel and AMD needed memory that would perform better than PC-133. That was a given, and you could run all sorts of simulations and performance analysis that said, faster memory, better.

So, it was pretty clear the industry was going to go that way. The question is do you pick DDR or do you pick Rambus? And given that, you know, Intel, who owns 80 percent of the market, really put his wood behind the arrow, so to speak, on Rambus, you know, they had talked about the customers, well our customers were saying, hey, you ought to use Rambus, and we talked to the memory vendors. And the memory vendors were saying, you know what, Rambus, it's a revolutionary change, not evolutionary, but, you know, that's the way the industry is going, that's the way we're going to go, and Rambus is it. And furthermore, here's something that's another critical component --

JUDGE McGUIRE: I'm going to cut you off, Mr. Heye, because I don't want you to just keep going on, I

want you to be able to respond to, you know, counsel's question.

BY MR. ROYALL:

Q. Thank you, Your Honor, and I'll take this time to hand Mr. Heye some water, if I may approach.

JUDGE McGUIRE: Go ahead.

THE WITNESS: That would be great.

BY MR. ROYALL:

Q. Now, going back, Mr. Heye, to what you said in response to an earlier question. When you and your team in '97 made the judgment that you should migrate from SDRAM memory and then in the future to RDRAM memory in connection with the K-7 microprocessor, was that based in part on information you were getting from the memory manufacturers?

A. Yes.

Q. So, the memory manufacturers were supportive of the idea of moving to RDRAM memory in the future at that time frame. Is that correct?

A. Yes.

Q. You mentioned in your earlier answer, you mentioned, I believe the term "revolutionary" in reference to Rambus memory. Can you explain what you meant by that?

A. If you look at technology trends, let's use

PC -- let's use synchronous DRAM going to DDR and contrast that to Rambus. To go from synchronous DRAM to DDR did not require a brand new back end test methodology in the memory groups, in the memory companies. The basic design, the basic architecture, the basic understanding of the mechanism of how memory communicated to front side bus between Sync DRAM and DDR was pretty straightforward. Rambus really, really had a very different architecture. It was quite a rapid departure from the existing Sync DRAM and DDR communication conventions.

Q. And did that different Rambus architecture create any complications with respect to designing the rest of the infrastructure that would support the K-7 microprocessor?

A. Yeah, Rambus was a fairly high performing bus, and the way it was designed was you really, really had to be very careful how you designed both the north bridge, the memory chips, the RIMMs and the motherboard, and if you didn't get that all just tuned out just right, you're going to have some serious electrical problems. It was a really touchy design.

Q. Now, by comparison to Rambus memory, was DDR memory more evolutionary?

A. It was more evolutionary, yes.

Q. And did that mean that it involved less complication in terms of designing an infrastructure that would support that type of memory?

A. It was less complicated. It was still, you know, DDR was always faster than Sync DRAM, so that was certainly going to be some investment in terms of training the motherboard vendors and the dim manufacturers how to support the DDR infrastructure, but it was certainly less complicated than Rambus.

Q. But even despite the complications associated with Rambus, your initial decision in the 1997 time frame was that it made sense for AMD to use Rambus technology in future generations of the K-7 and not to use DDR. Is that right?

A. That's correct.

Q. And that was a decision that you made with input from memory suppliers?

A. That's correct.

Q. Was there anything -- any drawback or particular drawback that in your judgment at that time frame to using DDR as opposed to Rambus?

A. Well, the biggest drawback at the time was the standard wasn't a standard. They hadn't -- the memory guys -- the memory companies could not agree on what the DDR standard should be. And because of that, it was --

it was languishing in JEDEC, and furthermore, you need a chipset to help debug the memory. Intel was doing a Rambus-based north bridge. No one at the time that we knew of was doing a DDR-based north bridge. So, there was no way you're going to test DDR. So, we thought, well, heck, if Intel is going to throw its money at the problem, and help establish the infrastructure in Taiwan and establish the infrastructure in the rim manufacturers and do all that work, we'll leverage it and go with Rambus.

Q. And once the decision was made to go with Rambus, did you say that you were involved in negotiating a license with Rambus?

A. Yes.

Q. And what did that license cover, if you recall?

A. It covered NRE, which is nonrecoverable engineering costs, so basically we gave them, I think it was $2 million, and for that $2 million, like I said, you received a piece of technology you could literally plug into your north bridge design that could communicate to memory, and you also just got the right to use the Rambus technology.

Q. The piece of technology that you're referring to that you received from Rambus that connected with the north bridge part of the system, what was that Rambus

technology called?

A. I think it was a rack or a rim. I can't keep track, I can't remember right now. What it was was, it was an actual circuit and layout design in the technology of our north bridge that we could literally drop into our design, and that would communicate directly to the Rambus memory.

Q. So, when you signed the license with Rambus, you paid some amount of money up front?

A. Yes.

Q. Is that right?

A. I think it was $2 million.

Q. $2 million, and did you agree to pay royalties on the -- in the future on products that you -- that AMD produced?

MR. STONE: Your Honor, the specific royalty rates should be treated as in camera, that information is subject to the protective order as restricted confidential. I don't mind it being introduced in camera and would intend to bring it out on cross examination anyway.

JUDGE McGUIRE: Mr. Royall?

MR. ROYALL: I understand, Your Honor, and I'm not intending to bring out the actual rates at this point. I was just establishing that.

JUDGE McGUIRE: Okay, then that being established we'll leave it in the public record.

MR. STONE: I have no objection to it now.

JUDGE McGUIRE: I understand your objection. Go ahead, Mr. Royall.

BY MR. ROYALL:

Q. Mr. Heye, without getting into rates, if you even recall what royalty rates were, I'm simply asking to your recollection did the license that you signed or that AMD signed with Rambus involve payment of or agreement to pay royalties in the future based on production or sale of items by AMD?

A. Yes.

Q. And do you recall what types of AMD products were subject to that license provision?

A. Both north bridges and microprocessors.

Q. Now, once the decision was made within AMD to use Rambus technology in the future, and the license was signed, what did AMD do, if anything, at that point to begin to implement RDRAM memory in its designs?

A. Well, we mainly started a design team to do a north bridge that would encompass the Rambus technology. And actually the first thing we did was we wanted to build a test chip that would validate the Rambus design that they were giving us. So, we actually had a bunch

of design engineers working with a Rambus team to incorporate their design into a test chip which we would then manufacture and validate that the design that Rambus gave us was working well, and then we would continue to roll that into our north bridge and make a product out of it.

Q. Do you know or do you have an estimate of how many AMD engineers were working on the Rambus technology in the period during which that was the plan of the company to use Rambus in the future?

A. Probably around 10 to 15 engineers, and then there's probably another 10 folks that support the engineering community. So, 20, 25-ish.

Q. During that time frame, and just to be clear, we're talking about is it 1998 that we're talking about? Or do you recall?

A. Yes. Yeah, it had to be early '98.

Q. During that time frame, you said, I believe, that you were working with Rambus, AMD engineers were working with the Rambus engineers. Is that right?

A. That's right.

Q. Were you working with any other third parties in connection with Rambus such as memory manufacturers?

A. Not from an engineering point of view, no. We were talking to them, but no engineering work was being

done with it.

Q. And what was your purpose for talking to them about Rambus in that time period?

A. Just we still -- we were -- while the industry had decided that it was going to go Rambus, and by industry, again, it's the memory makers, the chipset guys, and certainly Intel, it was pretty clear that there were some serious challenges to getting Rambus to the high volume market, and so we just wanted to keep track of, you know, what the heck the memory guys were really doing and making sure that they were on track. Because if they started slipping, then we're going to have a problem.

Q. Were you personally involved in that time period in discussions with memory manufacturers relating to Rambus technology?

A. Yes.

Q. Did you make trips to visit those companies or did they come to see you in that time frame?

A. I made trips to visit them, I go to Japan, Korea, Germany.

MR. STONE: Your Honor, can we just clarify when he says "in that time frame," what time frame we're talking about here?

JUDGE McGUIRE: Mr. Royall?

MR. ROYALL: I think he just a few questions earlier had said 1998.

THE WITNESS: Early 1998.

MR. ROYALL: Early 1998.

MR. STONE: Thank you.

BY MR. ROYALL:

Q. And in your meetings -- strike that.

When you met with memory manufacturers in that early 1998 time period to -- and you discussed Rambus -- issues relating to Rambus technology, were you meeting with those companies collectively or individually?

A. Individually.

Q. Was there a reason why you met with them individually as opposed to collectively?

A. Well, they would go over their detailed roadmaps, which would include their technology roadmaps, and memory technology roadmaps are extremely confidential. That would be the equivalent of Intel and AMD in a joint session sharing confidential information in front of each other. It just wouldn't happen. I mean, the memory guys are incredibly competitive and there's no way that they would share their technology roadmaps with a competitor in the room.

The other thing, too, was they were very open to us in terms of cost. Because the other thing that I

would worry about --

MR. STONE: Your Honor, I would object that this goes well beyond the question, which is why did you meet with them individually.

JUDGE McGUIRE: Sustained, and try to confine your testimony to his questions. Let's try to stick to that and I am going to admonish you as well, Mr. Royall, on that point. So --

THE WITNESS: Yes, sir.

BY MR. ROYALL:

Q. Referring to these meetings that you recall in the early '98 time period with memory manufacturers in which you discussed issues relating to Rambus, do you have any present recollection of input or impact that you received from the memory manufacturers at that time relating to Rambus?

A. Yeah, the memory --

MR. STONE: Your Honor, this is beyond the answer yes, which I think is all he needed to answer the question that he has a present recollection, if he's going to go into what he was told by the memory manufacturers, that would be hearsay.

MR. ROYALL: Your Honor, I have laid a foundation, a very clear foundation that in decisions Mr. Heye made about memory selection, it was very

important to him to consider input from a variety of sources, including memory manufacturers. I'm not getting into that for the truth of the matter, but to understand his state of mind and the decisions that he made for his company in that time.

MR. STONE: Your Honor, I am quite happy for the evidence to come in, as long as we all understand it's not being offered for the truth of what the manufacturers said.

JUDGE McGUIRE: Mr. Royall, is that your understanding?

MR. ROYALL: Yes, Your Honor, as stated.

JUDGE McGUIRE: All right, you may continue.

MR. STONE: Thank you.

BY MR. ROYALL:

Q. Do you have the question in mind, Mr. Heye?

A. Could you repeat it, please.

Q. In the early 1998 time period when you met with memory manufacturers and you discussed Rambus technology, do you have any present recollection of input or feedback that you received from those companies during those meetings relating to Rambus?

A. Yes. They were concerned -- the cost of Rambus was starting to concern the memory vendors. This is as I understood it, as told to me by them, and it was

becoming pretty clear to me that in talking to memory vendors, that some of the initial cost projections of Rambus versus DDR were not coming in line to what had initially been expected, and that it was more expensive than people had thought.

In fact, one question I had asked them, I can't remember if it was early or late '98 when I was asking this question, but I had always asked the question as follows: If SDRAM costs one, how much does DDR cost? And the answer was 1.X and X would be different per memory vendor but would always be more expensive than synchronous DRAM. And I would always say, okay, what is Rambus, and the answer would be 1.Y, and again the Y was going to be greater than synchronous DRAM, but what I was also interested in seeing was that Y was always greater than X. In other words, every memory vendor that I spoke to would tell me that Rambus had a higher cost structure on a per part basis than DDR.

Q. Now, putting aside input that you were getting from memory manufacturers in this time period, were you receiving any input internally within AMD from the engineers that were working on Rambus memory?

A. Yeah, they were getting nervous. The --

MR. STONE: Again, Your Honor, this is hearsay.

MR. ROYALL: This is the -- it's the same issue,

Your Honor, I'm not offering it for the truth, it goes to his state of mind.

JUDGE McGUIRE: Let's make that clear at the onset of your inquiry, Mr. Royall, so we don't have to go through this every time.

MR. STONE: Thank you, Your Honor.

BY MR. ROYALL:

Q. Again, the question, Mr. Heye, and again I'm just asking you these questions from the standpoint of what you understood and what you heard and how that influenced your decisions.

A. Okay.

Q. But in the same time frame that we've been focusing on, early 19 -- early 1998, or thereabouts, were you and your team receiving any input internally within AMD relating to efforts to implement and design around or to design with the Rambus technology?

A. Yes, a couple of things were going on. One, the actual folks working on the silicon team were starting to become concerned on how difficult it was to implement Rambus in the technology under the current interpretation of the technology. Again, it was not impossible, but it was harder than they thought, and their concerns that it might not be as stable as they had hoped for. The second thing that was happening was

on a -- we were doing some analysis on the motherboard, and if I may, let me just show you this motherboard.

The motherboard is the piece of -- this is fiberglass, the board itself is made out of fiberglass, and this is nothing but fiberglass, a layer of copper, fiberglass, layer of copper. And in Taiwan, in the PC industry, all motherboards have four layers, that's it. You can manufacture boards with 100 layers, it's technically feasible, but in the high volume segment it's four layers.

The other thing to look at is the back side. If you look on that back side, and this is a motherboard that was made probably some time in the last six to nine months, there's nothing on the back side. Which means it's what they call single-sided manufacturing, which means the board goes on the top and they put it in a box and they ship it. Again, can you manufacture the back side? Absolutely. In fact, that's what Apple did. But it's more expensive.

So, one of the standard things you're always trying to do is single-sided -- single-sided four-layer motherboards. The concern that we were starting to have was that it looked like that you may have to go to your multiple layers, your out of bunch capacitors for the back side of the board, and that would add cost to

product, and we were concerned about that.

Q. And that was a concern that was raised to you internally within AMD. Is that correct?

A. Yes.

Q. In this same time period, roughly early 1998, did you have any interaction with memory module makers relating to Rambus?

A. Yes. The other thing that was occurring was --

Q. And let me just to make clear, I am not -- I am interested in understanding what, if anything, you may have learned from memory module makers relating to Rambus to understand your state of mind and how this may have influenced your thinking in that time period.

A. Right. You know, first clarifying, we keep saying early '98, I would like it to expand to like the first six months of '98, because it was kind of a continuum for this. There was a bunch of different data points coming in. But the other thing we were starting to hear was that the RIMMs, and again the RIMMs in accord with the DIMMs that Rambus --

Q. Can I stop you there just to make clear what you said. RIMMs is a reference to is it the memory module that goes with the Rambus memory?

A. Correct.

Q. As oppose to dim, which is the memory that's the

name given to the memory modules with synchronous or DDR memory. Is that right?

A. That's correct.

Q. Continue.

A. And so the RIMMs were also had some issues, one of which for example, in the heat syncs. If you looked at the dim that we looked at earlier on today, there are no heat syncs. If the rim has heat syncs, that's more expensive. The other thing again, the layout of the RIMMs, how you physically hook up to Rambus on those RIMMs was extremely expensive. Again, even in DDR, you have to be very -- there are strict rules on how you lay these things out and there are standards and all, but Rambus was just more difficult. And we were starting to, again, you know, hear, I admit that, that there were inneroperability problems with Rambus. We didn't have a Rambus design in-house, Intel was the only one who had their own chipset, AMD didn't have their own chipset inside.

JUDGE McGUIRE: Again he is going on and on again, I would ask you to confine yourself only to the question that's being asked, otherwise we'll be here all day and we cannot afford to do that.

THE WITNESS: Sorry.

BY MR. ROYALL:

Q. Now, the other issues that you heard described within AMD from memory makers and memory module makers concerning Rambus, did you or your team make any effort to work with Rambus to try to resolve or improve these issues?

A. Yes.

Q. What can you tell us regarding the efforts that AMD made to work with Rambus to resolve the sorts of issues and concerns that you've described?

A. Again, I was not part of the engineering team working with Rambus, but I know we had, you know --

MR. STONE: Your Honor, I'm going to object. If the prior testimony about his concerns was not being offered for the truth, then we don't have any foundation that there was any basis for these concerns. They were just things that were expressed to this witness. If he now says I didn't have any personal involvement in the engineering team, he's going to be again expressing his view based on hearsay and he lacks personal knowledge or foundation for this testimony and I object to that, Your Honor.

JUDGE McGUIRE: Mr. Royall, response?

MR. ROYALL: I can lay a better foundation for the question, Your Honor.

JUDGE McGUIRE: Okay, go ahead.

BY MR. ROYALL:

Q. During the 1998 time period and focusing on the first half of 1998, did you have any personal knowledge of efforts relating to the work of the AMD engineers that were working on Rambus, efforts on their part to work with Rambus employees relating to the design and implementation issues concerning Rambus?

A. I knew that we had -- that there were meetings taking place between the two companies on non-engineering issues.

Q. Now, you mentioned earlier that at some point in time, a decision was made to go with DDR memory in future versions of the K-7 infrastructure and not with RDRAM. Is that correct?

A. Yes.

Q. Do you recall when that decision was made within AMD?

A. Probably around late summer/fall time frame.

Q. Of 1998?

A. '98, yes.

Q. And were you involved in that decision?

A. Yes.

Q. And how were you involved in that decision?

A. I personally drove that decision through AMD.

Q. And what do you mean by saying that you drove

that decision through AMD?

A. Through all the information I was collecting throughout the industry, it was my personal belief that Rambus was going to fail as a commodity part, and that ultimately even Intel would have to go DDR, and that AMD should be the first -- should drive the DDR standard and not get tied up with the Rambus memory.

Q. What caused you personally to come to the conclusion that AMD should substitute DDR memory in its future infrastructure in place of Rambus memory?

A. Again, as I testified from information that I was hearing from memory vendors in terms of cost, from information that I was hearing from my own design team, saying that they were having concerns about the technology, the other thing that was happening in parallel on DDR was the JEDEC committee was actually getting more focused and was starting to get some closure on DDR. They hadn't closed on it yet, but they were getting there.

MR. STONE: Your Honor, I believe this witness, he's gone beyond the question, I think, but if he hasn't, he has no foundation to testify what JEDEC was doing. I don't know that he ever attended a JEDEC meeting and I don't think any foundation has been laid. I object on that basis.

MR. ROYALL: Your Honor, I don't think it does, in fact, go beyond the question, but I believe I can lay a better foundation.

JUDGE McGUIRE: Lay a better foundation.

BY MR. ROYALL:

Q. During this time period, I believe you said it was some time in the latter half of 1998, when you personally made the decision to substitute DDR memory in the K-7 infrastructure in place of Rambus memory.

A. Yes.

Q. Was that decision based in any part on knowledge or information that you personally had relating to the development of DDR standards?

A. I had members of my staff who did attend -- went to JEDEC meetings, come back and tell me that it was their belief --

MR. STONE: Your Honor, again, what he was told by members of his staff is hearsay.

MR. ROYALL: Your Honor, if I could get a yes or no answer to my question, I think I will have laid a foundation, again, all of this goes to his state of mind and his decisions. It's not hearsay, we're not offering it for the truth of the matter.

JUDGE McGUIRE: Okay, Mr. Stone, does that satisfy you?

MR. STONE: As long as it's not being offered for the truth. All we have is a witness who is testifying just --

JUDGE McGUIRE: And let's try to, if that's going to be the case, then let's try to put that up front, Mr. Royall, so we don't have to spend time going back through this same, you know, objection time and time again.

BY MR. ROYALL:

Q. Yes, Your Honor, thank you.

Now, first of all, Mr. Heye, let me just go back and ask you, and a yes or no answer would be fine, if you can give such an answer, was your decision to go with DDR memory in the K-7 architecture in place of RDRAM memory in future versions of K-7, was that decision based in any part on information or knowledge that you had relating to developments with the DDR standard-setting process?

A. Yes.

Q. Now, without asking you for to establish what, in fact, was happening with DDR, but simply asking for your understanding of that information and how you took that information into account, if you did, in your decisions, can you explain what information relating to the DDR standard-setting process you had available to

you and that you considered in that time frame.

A. It was my understanding that they were -- that the JEDEC community was closing in on a formal DDR standard.

Q. And again, in your decision-making process to substitute DDR memory for RDRAM memory in future versions of the K-7 infrastructure, did that information have some bearing on your decision?

A. Yes.

Q. And in what way? How did it bear on your decision?

A. Per my earlier testimony, you can't design a chipset to take advantage of a memory device if the memory device isn't specified. So, before I was going to go to our senior management and tell them to go DDR, I had to have a pretty good belief that ultimately there would be a DDR standard, otherwise we couldn't design a DDR chip and we would have failed.

Q. You mentioned going to your senior management relating to the choice of DDR, did you, in fact, go to senior management relating to that issue?

A. Yes, sir.

Q. Who within senior management did you raise that issue with?

A. Jerry Sanders, CEO.

Q. And what did you -- what did you say to him or what proposal did you make to him?

A. I basically proposed that AMD stop doing its work on the Rambus north bridge, immediately begin work on a north bridge that would work with DDR, and that AMD would be the first company to propose DDR as a industry standard in the PC industry.

Q. And did you receive approval for that decision at that point in time?

A. Yes.

Q. Was that decision to go to DDR from Rambus in that time period in terms of your future plans for K-7, was that a decision that involved any potential risk to AMD?

A. Huge risks. If I was wrong, we would have -- AMD would have been at a huge, huge competitive disadvantage to Intel.

Q. And why is that?

A. Because if I was wrong, the dominant commodity part for memory would have been Rambus, which was clearly higher performing than PC-133, the best AMD would have had would have been PC-133, and Intel would have just owned the performance space of the PC industry, and that means that's where your high priced processors are going, we would have been out of that

market and that would have absolutely killed us.

Q. Once the decision had been made to use DDR memory in future versions of the K-7 infrastructure as opposed to RDRAM, what, if anything, did you and your team at AMD do to prepare the company for that step in the future?

A. Well, there was actually multi-avenues we had to attack. We had to first attack the -- first you had to get the specification completed, you had to then work with the memory vendors to ensure that they actually had a supply of DDR parts out there. In parallel, we started this effort called Team DDR, because don't forget, at that time frame, Intel was spending lots and lots of money and energy telling the entire world that Rambus was the way to go. And we had to establish the feasibility in the mind of our partners and our customers and the analysts that DDR was real. And so we established this notion of Team DDR to help us do that.

Q. Let me stop you there. You mentioned something called Team DDR. What is or was Team DDR?

A. Team DDR was a group of partners which AMD led which represented a valid infrastructure around DDR so that Team DDR would consist of memory vendors, dim vendors, chipset vendors, and motherboard vendors, and the idea was that, again, we needed to establish the

feasibility of this infrastructure, because even our customers were telling us that Rambus was the way to go. And we had to say no, there's a viable alternative. And so Team DDR was one directed at analysts and the press saying, hey, DDR is real, and getting partners on board and ultimately telling our customers, because we had to tell our customers, hey, we have a live technology, we have a live alternative and it's going to be successful in the marketplace. So, that was Team DDR's function.

Q. In terms of your role in building a K-7 infrastructure, were there any complications created by your decision to change from the plan of going with RDRAM to the plan of going to DDR?

A. Well, it's just like what I just testified, I mean, we had to establish a whole brand new memory infrastructure, which was something that AMD had never attempted, and for all those cases that I had just rattled off, we had to address all that, and furthermore, we had a team that was designing the Rambus interface, we had invested quite a bit of dollars in that team, all that work was discarded and we had to start doing DDR.

Q. At the time that you made that judgment, were the various business partners that you worked with in developing an infrastructure, were they prepared to

develop products, their own products, in the way that would be compatible with a DDR-based infrastructure?

A. Yes, we -- yes. We got them to agree with us.

Q. I'm sorry, you got them?

A. Yeah, they ultimately agreed to us and were willing to work with us, yes.

Q. Did AMD need to do anything in terms of training or education to help its business partners be prepared to support DDR?

A. One thing we did is in addition to designing our own chipset, we designed what we call a reference board. So, actually, AMD designed and built a motherboard that had DDR on it, and we went around with the feasibility saying, look, this thing really works. We also went to Taiwan and said, okay, Taiwan, we want you folks -- you folks -- we want these large corporations to build motherboards that supported DDR, and again, DDR was faster than Sync DRAM, it's technically more complicated, and it works --

JUDGE McGUIRE: Okay, you have answered the question.

BY MR. ROYALL:

Q. And from the point at which you made the judgment to go with DDR in future versions of the K-7 infrastructure, how long after that decision was made

did it take you and your team to put in place the virtual system or infrastructure that would support DDR memory with a K-7 microprocessor?

A. We launched our first DDR product in October of '99 -- no, October of 2000.

JUDGE McGUIRE: So, then, how long did it take?

THE WITNESS: It -- well, we changed to it about a little over -- about 18 months. Yeah, that's 12 to 18 months, in that time frame.

BY MR. ROYALL:

Q. Well, you said you launched the DDR version of K-8 in -- did you say October 2000?

A. Yes.

Q. And the decision, you said earlier, the decision to go with DDR as opposed to RDRAM was made in did you say late '98?

A. Yeah, mid to late -- I don't know exactly, but I'm assuming mid to late '98, so it would be a year, maybe 15 months.

Q. Now, we touched very briefly earlier on the K-8 microprocessor.

A. Yes.

Q. Which you mentioned is also known in the server marketplace as I believe Opteron.

A. Correct.

Q. Were you involved in any decisions relating to what memory to use in support of the K-8?

A. K-8 offered a radical departure --

JUDGE McGUIRE: Wait a minute, sir, that's not the question. He asked you if you were involved.

THE WITNESS: Yes. I'm sorry.

JUDGE McGUIRE: Okay, now, next question.

BY MR. ROYALL:

Q. Thank you, Your Honor.

And how were you involved in decisions relating to memory selection of K-8?

A. My team offered feedback to K-8 on what the dominant -- what we thought the dominant memory would be in the time frame of its launch.

Q. And what type of memory was selected for use in the initial launch of the K-8?

A. DDR.

Q. And has there been any decision about any future version of K-8, what type of memory will be used in the future versions of K-8?

A. The expectation is we will use DDR2.

Q. And can you very briefly explain what reasons caused AMD and your team to choose DDR and in the future DDR2 memory for the K-8?

A. Well, if you look at that chart, I mean, we

started developing K-8 in the '98-'99 time frame and that's just the time frame that we had made the determination that DDR was the memory of choice for the commodity infrastructure, so that was our decision.

Q. And when you say memory of choice for the commodity infrastructure, does that mean that the memory that you in your judgment believed will be the dominant commodity memory at the time of those versions of the K-8 infrastructure was introduced into the marketplace?

A. Yes.

Q. Now, relating to JEDEC, you mentioned a few things about JEDEC, to make it clear for the record, have you, sir, ever attended a JEDEC meeting?

A. No.

Q. Does anyone on your team at AMD attend JEDEC meetings to your knowledge?

A. Yes.

Q. And who is that?

A. Steve Polzin, and I know Steve has some people working for him who I think also attend JEDEC meetings.

Q. And do you know what portion or aspect of JEDEC your team members are involved in?

A. Specifically, I know they're involved in the memory aspect of it.

Q. Do you personally, I'm asking for your personal

views, if you have any, do you personally have any views regarding the value, if any, that AMD derives from participation in JEDEC?

A. AMD spends a lot of time -- AMD works collaboratively with the memory vendors through JEDEC to ensure that the memory standards going forward can be implemented both by the chipset vendors and the memory vendors.

Q. You mentioned earlier the various business partners that you work with in building an infrastructure to support AMD's microprocessors.

A. Yes.

Q. Are any of those business partners involved in JEDEC as well?

A. Well, I know for certain that the memory vendors are involved in JEDEC. I don't honestly know who all the -- I don't know who all the participants are in the JEDEC meetings.

Q. In addition to memory-related standards, are there other areas of AMD's business in which the company relies on industry standards?

A. Yes, when we have, again, there are IEEE standards we're involved with, AMD is involved in the PCI standard, similar comments when I was at Apple, there are dozens of standards in the high-tech industry

and AMD is involved in many of those standard-setting bodies.

Q. And are you familiar with the term "open standards?"

A. Yes.

Q. What does that term mean to you?

A. Open standards is a term that says that there's a standard that is available to be used in the industry that is -- that effectively is royalty free, that if you follow the standard, one, you will be able to interact with other folks using that standard, and that there's basically no royalties associated with it.

Q. And does AMD have any position or view to your knowledge as to -- well, strike that.

Do you personally have any view about the value, if any, that AMD derives from open standards relating to its business?

MR. STONE: Objection, relevance, Your Honor. I don't think this is a trial about whether open standards are good or bad or proprietary standards are good or bad and this witness's view on the value of open standards doesn't seem relevant to any issue in this case.

JUDGE McGUIRE: Overruled. He can answer the question.

MR. ROYALL: Thank you, Your Honor.

THE WITNESS: Yeah, AMD's use of open standards is absolutely critical for success in the marketplace.

BY MR. ROYALL:

Q. And why is that?

A. AMD has a minority share of the market, and what that means is, we really want to ensure equal playing field between us and Intel. And in order to do that, the best way to ensure that is that if Intel and AMD both adhere to open standard bodies so that the specification that I get, AMD gets, it's at the same time that Intel gets, you know, then may the best competitor win. That's totally fair.

With regard to proprietary standards that aren't open, then usually the dominant player, i.e. Intel, will get access to standards before AMD does, and I'm at a competitive disadvantage.

Q. Are you familiar with something called ADT?

A. Yes.

Q. What is ADT, or what do you understand ADT to be?

A. ADT stood for Advanced DRAM Technology, I believe, and it consisted of Intel and six memory vendors, the goal of which was to design the next generation memory interface.

Q. And do you recall when you first learned of ADT?

A. 2000, maybe 2001.

Q. Do you recall how you learned of ADT? Or the existence of ADT?

A. We heard about it through our memory partners, they told us they were doing that.

Q. Do you have any understanding concerning how or why the ADT consortium was formed?

A. Again, my understanding was that it was formed to address next generation memory interface, you know, somewhat in response to the fact that JEDEC is not -- does not close on issues quickly, so this was supposed to be a more expedient mechanism to get the standard to market.

Q. And when you refer to next generation memory interface, are you talking about types of memory designs beyond --

A. Beyond DDR.

Q. Beyond DDR?

A. Beyond DDR, yes.

Q. Do you know who formed the ADT consortium?

A. I honestly don't know which initiated it.

Q. Do you know whether the work of ADT had anything to do with Rambus?

A. I heard rumors, but I wouldn't say that it was formally -- no, to answer your question, so I won't get

objected on it.

Q. And when you heard of ADT, did you on behalf of AMD have a desire to participate in the ADT consortium?

A. Yes.

Q. And why is that?

A. It goes back to my previous testimony that, you know, Intel is a formidable competitor. The fact that Intel was getting information defining the next generation memory interface before AMD put AMD at a distinct disadvantage to Intel.

Q. And did you make an effort to have AMD join and participate in the ADT consortium?

A. Yes. The ADT consortium rules were that you had to get a unanimous vote with the six memory vendors plus Intel to allow us to get on what I'll call the executive committee. Actually, I personally called Pat Gelsinger at Intel to try to elicit his vote to try to get AMD on the executive committee. For the record, he was noncommittal. But yeah, we tried really hard, we talked to the memory vendors as well to try to get us on the committee, and there were votes taken and we lost the vote. And the votes were held in confidentiality, so I have no idea who voted for us and who voted against us.

Q. So, you were not -- in your efforts, you were not successful in your efforts to join ADT. Is that

what you're saying?

A. We were not successful in our efforts to join the executive committee of ADT.

Q. To your knowledge, does ADT still exist today?

A. No.

Q. Do you have any views as to whether AMD suffered any type of disadvantage as a result of not being permitted to participate in the executive committee of ADT?

A. I don't think we suffered any damages.

Q. And why not?

A. The span of the technology didn't go anywhere.

Q. Now, you mentioned earlier that AMD entered into a license with Rambus at the time that you and your team made the decision to pursue use of Rambus technology in future versions of the K-7 infrastructure. Is that right?

A. Yes.

Q. And you were personally involved, I believe you said earlier that you were personally involved in negotiating that license with Rambus. Is that correct?

A. Yes.

Q. Did you have meetings with Rambus representatives relating to that license or the negotiation of that license?

A. Yes.

Q. Do you recall how many meetings you attended with Rambus representatives relating to that license negotiation?

JUDGE McGUIRE: Mr. Royall, how is that pertinent at this point?

MR. ROYALL: I'm just laying a foundation, I think it will become clear, Your Honor.

JUDGE McGUIRE: Okay, proceed.

THE WITNESS: I have no idea exactly how many meetings they had, it was certainly greater than three and less than 100, I mean, we were negotiating.

BY MR. ROYALL:

Q. Well, no, I'm sorry, the question, maybe I misstated it, was not how many meetings there were, but do you recall how many meetings that you participated in? Or roughly speaking, I don't mean precisely. I mean, was there more than one?

A. Definitely more than one, yes.

Q. And during those meetings, whatever the number may be, but the meetings that you personally participated in with Rambus, do you recall Rambus making any presentations relating to its intellectual property?

A. Not specifically its intellectual property, no.

Q. Were there other meetings of AMD employees with

Rambus relating to this license negotiation that you did not attend?

A. Yes, many.

Q. And were these other AMD employees, were they -- that participated in those meetings, were they part of your team?

A. They didn't all report to me. You said my team, I'm not sure exactly what that means. There was a negotiating team from AMD which included lawyers, business development people and people on my staff. And that was the negotiating team for the contract. I was the lead executive on it.

Q. And before AMD signed a license with Rambus relating to RDRAM, do you recall ever hearing that Rambus had or might have intellectual property relating to JEDEC's SDRAM standard?

A. No.

Q. Did you at some point in time learn that Rambus had or might have or claimed to have intellectual property relating to JEDEC's SDRAM standards?

A. Yes.

Q. Do you recall when you first heard that?

A. It was in 2000, spring, early summertime period.

Q. Do you recall how you first learned that Rambus claimed to have intellectual property relating to

JEDEC's SDRAM standards?

A. Yes, they started suing memory vendors.

Q. And how did you learn of that, if you recall?

A. I read about it.

Q. Where did you read about it, if you recall?

A. On the web.

Q. And was that the first -- well, strike that.

At some point in time, did you have a meeting with or discussion with anyone from Rambus relating to intellectual property, Rambus intellectual property in its relation to the JEDEC standards?

A. Yes.

Q. Do you recall when that was?

A. Again, spring or early summer-ish of 2000.

Q. Do you recall one such meeting or was there more than one such meeting in which you interacted with Rambus representatives relating to JEDEC intellectual property?

A. Just one meeting.

Q. And do you recall where this meeting occurred?

A. At AMD.

Q. Do you recall what led to that meeting being scheduled?

A. Yes, the Rambus representative called me and said that she wanted to have a meeting with me and she

also said she was bringing her attorney. Which is their way of saying I should bring my attorney, so I did.

Q. And who was the Rambus representative who called you, if you recall?

A. Laura Fleming.

Q. And in the meeting that later took place, did Laura Fleming attend on behalf of Rambus?

A. Yes, Laura attended.

Q. Do you recall whether anyone else from Rambus attended?

A. Yes, there was an attorney there, and I think there were a few other folks from Rambus there as well.

Q. At this meeting, which I believe you said you recall being in the spring of 2000. Is that right?

A. Spring/early summer, yes.

Q. At this meeting with Laura Fleming and others from Rambus, did anyone else other than you attend on behalf of AMD?

A. An AMD attorney.

Q. And do you recall the name of that attorney?

A. No, I don't, I'm sorry.

Q. Anyone else other than the one AMD attorney and yourself from AMD?

A. No.

Q. Did Rambus at this meeting make any formal

presentation to yourself and the AMD attorney?

A. Yes. A PowerPoint presentation.

Q. Did they leave you with a copy of that PowerPoint presentation, or do you recall?

A. I believe they did, yes.

Q. Do you recall anything, without getting into any details or substance, but do you recall anything about the PowerPoint presentation?

A. Yes.

Q. You do?

A. Yes.

Q. Again, without getting into details, what do you recall?

A. Well, there was sort of two big topics they wanted to cover in that, one was our current status of our contract we have with Rambus, and the second one was they were -- they wanted to share with us the fact that they claimed that they had IP related to synchronous DRAM and DDR, they had patents in regard to those two technologies.

MR. ROYALL: Your Honor, at this point, I would like to show the witness a document, and it's a document that has been marked by Rambus as in camera and I gave Rambus notice that I planned to do this, but it is I believe an in camera document.

JUDGE McGUIRE: Then what we are going to do by all of you who are in the audience, by prior court order there is certain evidence in this case that cannot be offered to the public, so I am going to ask everyone from the courtroom who has not been cleared to hear this evidence to please go outside and I'm sure when we're done we will open the door or something and you will have a chance to come back in.

MR. ROYALL: One point of clarification, Your Honor, the one person other than Rambus attorneys and FTC representatives who is in the courtroom is the attorney representing AMD.

JUDGE McGUIRE: Well, has he been cleared to hear this, if he's not, he's going to have to leave. That's pretty clear now. It's not just how you two agree, it's they're either cleared to have access to in camera treatment evidence under the order or they're not. And I'm going to ask both of you to I guess certify that everyone at their table and everyone who is I guess behind them are certified under that protective order to have access to in camera treatment. It's going to be that simple.

MR. STONE: Everyone on our side is, Your Honor.

JUDGE McGUIRE: Thank you, Mr. Stone. How about you, Mr. Royall?

MR. ROYALL: Under the protective order, I'm not certain, so maybe as a caution.

JUDGE McGUIRE: If you're not certain, then he is going to have to leave. So, I'm sorry, sir, but that's just the way we're going to have to do this. So, Mr. Royall, does everyone else in this courtroom at your table and behind you have access to this evidence under that order, that protective order?

MR. ROYALL: One point of clarification, the protective -- the reason I wasn't sure is the protective order refers to outside counsel, but I believe that's outside counsel of the parties. There is one provision of the protective order that was just drawn to my attention that says and such other persons authorized by the producing party. So, if Rambus had no objection to Mr. Heye's attorney being here, he could be present, but that's up to them.

JUDGE McGUIRE: Well, it's not really their say. That's fine if he has no objection, Mr. Stone, but it's not his say as to whether it's okay. I'm going to ask you to certify to me that everyone behind you has access under the terms of that order to hear evidence. If you think he is, then he can come back in, but that's up to you.

MR. ROYALL: I understand, Your Honor, and

everyone in the courtroom on our side is authorized and the only question is I'm just asking whether Rambus would have any objection to Mr. Heye's company's attorney being present, and if they don't have an objection, I think under the protective order he can come in. If they do have a reason to object to that, then --

JUDGE McGUIRE: Mr. Stone, do you have an objection?

MR. STONE: The problem, Your Honor, is Mr. Heye's attorney is with a firm that represents Hynix, and I'm opposed to that firm in the private Hynix litigation, and so because of that, I hate to put him in a position where he has ethical constraints on what he and his firm can do.

JUDGE McGUIRE: All right, let's make it very clear.

MR. STONE: So I can't agree to it.

JUDGE McGUIRE: From here on out we're treating anything that's being in camera, if it's not clear, then the individuals who we're not sure about, they're going to be excused. Just so there's -- if we're going to err, it's going to be on the side of any caution.

MR. ROYALL: Yes, Your Honor, I perfectly understand that and I obviously have no dispute with

their position on this.

JUDGE McGUIRE: Okay, so are we ready to go then?

MR. ROYALL: Yes, Your Honor.

(The in camera testimony continued in Volume 19, Part 2, Pages 3883 through 3923, then resumed as follows:)

AFTERNOON SESSION

(1:27 p.m.)

JUDGE McGUIRE: This hearing is now in order. You may proceed at this time with your inquiry, Mr. Royall.

MR. ROYALL: Thank you, Your Honor. Actually, before I go any further, I think at this time I would like to offer into evidence the exhibit that we dealt with in camera, CX-1420.

JUDGE McGUIRE: Objection?

MR. STONE: No objection, it just needs to remain in camera, Your Honor.

JUDGE McGUIRE: Right, so entered on that basis.

MR. ROYALL: Thank you.

(CX Exhibit Number 1420 was admitted into evidence.)

FURTHER DIRECT EXAMINATION

BY MR. ROYALL:

Q. Mr. Heye, did you at some point in 2000 learn that Rambus had commenced asserting patents against memory manufacturers that were business partners of AMD?

A. Yes.

Q. And how did you learn of that?

A. Initially through reading it on the web.

Q. When you learned that Rambus was enforcing

patents relating to -- and just let's be clear about this, we're talking about Rambus asserting patents against memory manufacturers relating to SDRAM and DDR. Is that right?

A. Yes.

Q. And when you learned that Rambus was seeking to enforce or was enforcing patents against those memory makers relating to SDRAM and DDR SDRAM memory, did that cause you any concern from the standpoint of AMD's business?

A. Yes.

Q. I would like, if we could, if you could explain to me or identify for me what, if any, concerns you had from the standpoint of AMD's business relating to Rambus asserting patents against memory manufacturers, and if you could, identify whatever concerns there were, and I can come back and ask you about each separate concern.

A. Sure. So, again, to your point there are multiple concerns. The overarching concern was time to market. The second concern was a possible cost disadvantage we might incur in the infrastructure due to the incremental royalty fees.

Q. I'm sorry, you say cost disadvantage, what was it you said after that?

A. Due to incremental royalty fees.

Q. Okay. Any other concerns that you had?

A. There was some discussion of possibly changing the specification to work around some of the Rambus patent claims.

Q. And when you refer to specification, are you talking about the JEDEC DRAM standards?

A. Yes.

Q. SDRAM standards?

A. Yes.

Q. Okay.

A. And if that were to occur, that would introduce just a whole host of problems which would have been a major, major concern for AMD.

Q. And if I can follow up on that. Other than what you've mentioned, time to market, cost disadvantage, potential for changing the JEDEC specifications, and I understand there's some more you have to say on that, but are there other concerns that you had that you wanted to identify that we could come back and talk about?

A. No, because they would be covered as you go through each bullet.

Q. Why don't we start actually with the last of the issues that you identified. And again, what I'm asking you about is what concerns did you personally have, if

any, and you have identified some concerns, relating to the assertion by Rambus of patents against memory manufacturers relating to SDRAM and DDR SDRAM, and you mentioned that one of those concerns had to do with potential changes in the JEDEC SDRAM standards. Is that right?

A. Yes.

Q. And can you explain why that was a concern to you, and again, if there are multiple issues and you want to lay them out and identify them, I can come back and ask you about them separately.

A. Sure. So, the concerns around that would have been first, it would have taken time to establish the new standards; depending upon what they were, you would have had to change the memory component, the north bridge, possibly both, you would possibly have to change the motherboard. You may possibly have to change the dim, once you've made all those changes, you would have to implement them, of course, then you would have to test -- well, you have to debug them, because again, now you're talking about multiple vendors with multiple different components. That would take time.

You would have to revalidate all those modified DIMMs, and again, by DIMMs, I mean both the actual PCB and the memory chips against the north bridge, and by

the way, that is a really big deal. We spent a lot of time and effort trying to get that right in the first go-round. And you certainly had potential inventory issues. You had opportunity cost issues.

Q. Okay, if I can follow up on what you've said, and if you think of something else, and you identify it for me, I can ask you about that.

A. Sure.

Q. Now, before we go into these separate concerns, let me ask you this: Were your concerns relating to Rambus asserting patents against memory manufacturers in the 2000 time period?

A. Yes.

Q. Were those concerns in any way related to AMD's own plans in terms of marketing or launching microprocessors in that time frame?

A. Yes. I mean, to the extent that we -- our microprocessors are going to rely on DDR-based systems. Absolutely, the K-7, to take full advantage of the K-7 capabilities, you needed to have high performing memory, which DDR was the critical component for that.

Q. And I think you said earlier, but the K-7 microprocessor which was initially launched in '99 was later -- a later version of the K-7 infrastructure that is the system, virtual system supporting K-7 was

launched in did you say some time in 2000, or with DDR memory?

A. Yes, the first DDR-based systems came out in October of '99.

Q. Or October of --

A. Of 2000, I'm sorry.

Q. Of 2000.

A. I apologize.

Q. So, that was some time after you learned that Rambus was asserting patent rights over DDR and SDRAM memory?

A. Correct.

Q. Now, let's go through and talk about to the extent your concerns about Rambus's assertion of patent claims over those memory products related to potential changes in the memory standards of JEDEC. Can you first of all give me just a general explanation of why you were concerned about the assertion of patent rights resulting in changes in the JEDEC standards.

A. Well, I mean, I went through the list of the specific -- is the question how all those changes could actually -- in terms of just JEDEC itself, again, the knowledge that I was aware of was it took JEDEC a while to close on the original DDR specification. AMD had spent a lot of energy trying to get closure on that

specification, and I know that given the relationships the memory vendors have with one another, it's hard to get a consensus of change, and while they may all agree that changes are in the works, odds are against you that it's going to be the same change and you have to try to figure out which change they can all accept, and all of that takes time, and time is something that you don't have in this market.

Q. One of the things you mentioned in your answer, and I don't recall your exact words, but it had something to do with given how the memory manufacturers, did you say relate to one another?

A. That's a good word.

Q. And what do you mean by that?

A. Well, they're extremely competitive, and, you know, were -- for example, I'm making this up as I go along, but there could be a -- they could make a change -- a memory vendor could actually have a change that would take let's say two weeks to implement, that change might take ten weeks for the other guy, and vice versa. You know, they may have change B that has a reciprocal problem.

So, neither one of them will give ground, because if they give ground to the other and give them a competitive advantage, that's unacceptable to them. So,

the memory guys are very, very competitive. Time to the market is very critical to them as it is to AMD and Intel. They don't want to give up any potential competitive edge to one another. It's an extremely competitive market.

Q. Now, in terms of changing the JEDEC standards in response to Rambus's patent claims were to take time, and without specifying how much time that might take, but to the extent that were to take some period of time to do.

A. Yes.

Q. Would that in any way adversely affect your company AMD?

A. Absolutely. And again, this is all supposition, so the supposition path we're going down is, again, this was happening in the middle of 2000. We were planning a launch in the fall of 2000, October. By that time frame, the chipset was for all intents and purposes complete, we were in the validation testing, the DDR, the DIMMs, the memory was done, the DIMMs were being manufactured, the memory folks were actually starting production and waiting for it to start, you have this whole inertia on the path from being launched, because once you're about to launch, last minute design work is being done.

And one thing we thought of we took all the DDR memory from all the different vendors and we took all sorts of mixing and matching to make sure that all combinations were going to work with our north bridge, and that's a lot of work and a lot of dollars and we were in the process of doing that.

And as soon as you change that standard, and I don't know what, you know, because we use the word change, it has to be material, otherwise no one is going to know it's a change, that would impact some level of that whole inertia.

Q. Now, let me ask you this: You mentioned earlier that one of your concerns in terms of Rambus asserting patent rights against memory manufacturers had to do with something that you referred to as time to market.

A. Yes, sir.

Q. What did you mean by that?

A. Well, again, our plan was to launch DDR in the fall time frame. At that time frame, at that point, Intel was coming -- I don't know the exact time, but in that time frame, which would be, you know, late 2000, Intel was coming out with their Rambus part, and we knew that that Rambus part with an Intel processor --

MR. STONE: Objection, Your Honor, I think this goes beyond the question of what did you mean by that

when you said time to market. That was the question.

JUDGE McGUIRE: Sustained.

MR. ROYALL: Well, I understand, Your Honor --

JUDGE McGUIRE: If you want to restate another question, I might let you do that, but he is clearly beyond what the inquiry was.

THE WITNESS: I'm sorry.

BY MR. ROYALL:

Q. Let me ask, Mr. Heye, as best as you can, can you explain to us what you mean by the term "time to market" and just as precisely and clearly as you can explain what you mean by that term.

A. Time to market just simply states how long it takes you to get to market with a new product or feature.

Q. I'm sorry, how long it takes?

A. How long it takes you to bring to market a new product.

Q. Okay. Now, why is it that Rambus enforcing patents against memory manufacturers in the 2000 time period relating to DDR and SDRAM caused you to have concerns about AMD's time to market with its microprocessors?

A. In the time frame of interest, as I just stated, you know, AMD had -- and the industry had a lot of

momentum and were finalizing -- actually finalizing the manufacturing processes to bring us to production in four or five months. Anything that would cause a slip in any of those components would have thereby necessitated that the overall system was going to slip out in time. And thereby impacting our time to market.

Q. Okay. And so by that, are you saying that a change in the DDR standard, for instance, JEDEC's DDR standard could disrupt your efforts to release a new version of K-7 supported by DDR memory in October 2000 as actually occurred?

A. Yes.

Q. And how -- do you have any estimate of how long a change in the DDR standard might have delayed the release of a K-7, a new version of K-7 that was supported by a DDR-based infrastructure?

MR. STONE: Objection, lacks foundation, calls for speculation.

JUDGE McGUIRE: Overruled.

BY MR. ROYALL:

Q. Thank you.

A. A matter of months. Let me give you a quick data point. We typically state that if you change one transistor in let's say a chipset, it's going to take you four to six weeks to get that change to the market.

So, it doesn't matter what you're doing, because when you change just one transistor, again, not trying to figure out what the heck we're talking about in terms of actual content, but to change anything, not only does it require the design guys to change something, then you have to generate on the silicon side, you go from there, you develop a what they call a mass set, which is used for the manufacturer of that part. That mass set goes to the fab, it takes four to six weeks to fab that, that didn't include any design time to change that transistor.

There are human beings responsible for the design side, there are human beings designed for the mass set generation, there are human beings responsible for getting that product through the manufacturing process, it comes out of the fab, you have to put it into a package, you take that packaged part, you've then got to put it into the system, you have to have engineers test to make sure that one, the transistor change did what it expected, but equally important, you have to prove that that transistor change didn't provide an error in the part. Because I can tell you that --

JUDGE McGUIRE: All right, Mr. Royall, what's your next question?

BY MR. ROYALL:

Q. My next question, Mr. Heye, is focusing on the infrastructure that you laid out for us earlier in DX -- I don't know if we said this on the record, but I marked it as DX-30.

JUDGE McGUIRE: Right, DX-30, noted.

(DX Exhibit Number 30 was marked for identification.)

BY MR. ROYALL:

Q. If the DDR memory standard were to change.

A. Yes.

Q. What other parts of the infrastructure supporting the K-7 microprocessor might also have had to change in response to a change in the DDR memory standard?

A. Again, the north bridge may have to change, the physical dim may have to change.

Q. And the dim, again, you're referring to the module, memory modules?

A. The module that plugs into the motherboard. The motherboard itself may have to change.

Q. Anything else?

A. Let's see, memory, chipset, dim, motherboard, that's it.

Q. What about the BIOS?

A. It may have, it depends on the change, it may

have to change, that's a good point.

Q. I assume that all of this would depend on the precise change?

A. The precise change, correct.

Q. Now, once those changes were made, would there need to be any debugging process of the sort that you described earlier?

A. Absolutely. You just can't -- yes. You would have to debug it.

Q. What about the validation process you described earlier?

A. You would have to revalidate the whole matrix.

Q. Based on your experience and your work at AMD since 1997, do you have any understanding of what the cost of validating a new -- or a revised infrastructure are?

A. I would get out to the millions, because you in all likelihood would have to procure a bunch of new silicon, and then it really, it took us months to do the entire check-out of every possible combination of memory, and that's a huge concern.

Q. Now, you mentioned earlier when you listed or itemized for me the possibility of the JEDEC standards changing in response to Rambus's patent claims. One of the things which you mentioned was inventory.

A. Yes.

Q. Can you explain what you meant by that?

A. Well, again, the time frame we're talking about, again, if you recall it takes about three and a half to four months from the time a manufacturer starts the silicon to the time it hits production. So, if you want to ship in October, that means you have to start production wafers four months back, and four months from October is June. So, starting in the June time frame, those -- the memory folks have started production wafers.

Now, as they start going through the manufacturing line in terms of fabricating silicon, you change a standard, those wafers are scrapped. And that's more money. That's what they call WIP, work-in-progress in the fab, and if you wait longer, that stuff comes out the other end and is sitting in packaged parts and there's a whole bunch of inventory which you couldn't use because it violates the standard.

Q. Now, the next question is you mentioned earlier when you itemized the concerns relating to the potential change in the DDR standards, one of the things you mentioned was something called opportunity cost. Do you recall that?

A. Yes.

Q. Can you explain what you meant in saying that opportunity cost was one of the types of concerns that you had?

A. Sure. As I discussed earlier, the change would require a lot of engineers to get involved, and involved in the changing of the JEDEC standard, and that would bring no real value to the industry. They're changing the standard not to make it better, they're changing the standard to avoid litigation with Rambus. Those engineers would actually do what I would consider value-added work to actually improve performance and provide features to the customers.

Q. Are you familiar with the term "backward capability?"

A. Yes.

Q. Does the concept of backward capability, did that have anything to do with the concerns that you had about Rambus enforcing patents against DDR and SDRAM memory?

A. Well, absolutely, yes.

Q. And can you explain the nature of that concern?

A. Well, in terms of SDRAM, there are literally hundreds of millions of systems in the industry. I mean, I -- Intel and AMD-based platforms, over 100 million were sold per year for multiple years had

synchronous DRAM in there. The expectation of the user who buys that BC is that at some time in the future they want to upgrade their memory, they can go buy some PC-100 memory and plug it in. If you change that standard, and that memory no longer works for your existing PC, that's a problem.

Q. Now, are you finished?

A. I'm done, yes.

JUDGE McGUIRE: Yes.

BY MR. ROYALL:

Q. Now, one other thing that I wrote down that you said earlier, in terms of concerns that you had in how AMD could be impacted by a change in the DDR memory standard in 2000 had to do with cost disadvantage. And then you mentioned something about increased royalty fees. Do you recall that?

A. Yes.

Q. Can you explain what you meant by that?

A. Well, again, it's just you're adding incremental cost to your base -- to the commodity part. And typically system vendors have a certain cost they're going to pay for the hardware and if one component commodity starts going up, they try to take the cost of other components, like microprocessor pricing. So, you know, the goal is to always maintain the lowest possible

cost structure of the system.

Q. Are you saying that you had the concern with the potential for Rambus to charge royalties and thereby raise the cost of memory manufacturers?

A. Memory manufacturers and also the chipset manufacturers, because they were -- there's a double taxation.

Q. What do you mean by double taxation?

A. It means you had to pay for the rights to the Rambus on both ends of the bus, which means that the chipset vendors had to pay more, had to pay the royalty, and the microprocessor -- and the memory manufacturers had to pay royalty. So, you had to pay royalties on both sides.

Q. And did your concerns about Rambus asserting patents against DDR and SDRAM have anything to do with AMD's competitive position compared to Intel?

A. Yes.

Q. And in what way was that a factor in your thinking relating to this issue?

A. Well, it gets back to time to market, anything that impacts time to market would put it -- would mean that Intel would come out before we did with a competitive platform, which again would put us at a competitive disadvantage vis-a-vis Intel and its

platforms.

Q. Now, after you learned that Rambus was asserting patents against memory manufacturers relating to SDRAM and DDR SDRAM.

A. Yes.

Q. Did you at any point express the sorts of concerns that you've testified about this afternoon directly to the memory manufacturers that AMD worked with?

A. Yes.

Q. When was that?

A. Summer or fall time frame.

Q. Of 2000?

A. 2000, I don't remember the exact time.

Q. And did you meet collectively with memory manufacturers to talk about this issue?

A. No.

Q. You met individually with memory manufacturers?

A. Yes.

Q. Which ones?

A. Elpida, Samsung, Infineon, those were the three that I personally talked to.

Q. And you expressed concerns about the way in which changes to the DDR or SDRAM standards could impact AMD. Is that your testimony?

A. Yes.

Q. Did you ask these memory manufacturers to do anything or did you encourage them to do anything relating to this issue?

A. Well, my number one concern was to make sure that they maintained -- I'm trying to form it -- I wanted to make sure that in October of 2000 there would be a supply of memories in the market, and I went around to find out from an AMD perspective that we were going to do the AMD chipset that worked with DDR, we were still working with motherboard vendors to do DDR-based motherboards and it was my expectation that there would be a viable infrastructure for DDR in the fall of 2000. And I wanted them to participate in that infrastructure.

Q. Now, as we've been discussing these various concerns, we've been primarily focused on DDR, what about SDRAM, I don't mean to go back over the whole territory, but did you have concerns also about the potential for Rambus's patent claims to result in changes to the SDRAM standard as opposed to DDR?

A. Yes.

Q. And did those concerns differ in any way from the concerns you had about potential changes in the DDR standard?

A. Yes.

Q. How did they differ?

A. PC-100 was in existence, it was -- it was a huge installed base. There was a huge -- the infrastructure was alive, well, and established. The DDR infrastructure was in its infancy, it was just beginning, it hadn't been established in the marketplace, and the concern I had was that this could be a defocusing activity to actually establish DDR in the marketplace.

Q. How many microprocessors does AMD, in terms of current numbers, manufacture and sell today, if you know?

A. Roughly seven and a half to eight million microprocessors per quarter.

Q. Seven and a half to eight million per quarter?

A. Yes.

Q. Do you know what percentage of those seven and a half to eight million AMD microprocessors today use the DDR standard established by JEDEC?

A. 100 percent.

Q. If the JEDEC standard for DDR were to change today, would that have any impact on AMD, or do you have any concerns that it might have an impact on AMD?

A. It would have a very big impact on the whole PC industry, including AMD, yes.

Q. And how would that compare, if you have any views on this, to the nature of the impact that there would have been on AMD had the DDR standard been changed in 2000?

A. Now it's more the conversation of the PC-100 spec, because Intel right now has also adopted DDR as its preferable memory choice, so again, you have the vast majority of all PCs sold in the world today use DDR as its memory, and now you do have a huge installed base, now discussions of backwards capability come into play, and it would be a major problem for the industry.

Q. In terms of the impact that a change in the DDR standard would have on AMD today, would the impact be any different with respect to the K-7 microprocessor as compared to the K-8 microprocessor?

A. Actually, it would have worse ramifications for the K-8 microprocessor.

Q. And why is that?

A. As I've testified, the memory controller resides physically on the K-8 microprocessor, and that far and away has the most complicated IP in the system. If you make any changes to the memory controller, you have to really retest and revalidate the entire microprocessor. If there's any change to that silicon, you must treat as a -- you just have to be incredibly thorough and

rigorous and test the whole microprocessor in case you made a mistake. So, it would take a lot of time and effort. It would cause a disruption for K-8.

MR. ROYALL: Your Honor, I don't have any further questions at this time.

JUDGE McGUIRE: All right, thank you, Mr. Royall. You may cross examine, Mr. Stone.

MR. STONE: Thank you, Your Honor.

CROSS EXAMINATION

BY MR. STONE:

Q. Good afternoon, Mr. Heye.

A. Good afternoon.

Q. Earlier when we were in the in camera session, I asked a couple of foundation questions that I said that I would come back to and put on the record when we weren't in camera.

A. Yes, sir.

Q. The first one was the license agreement between Rambus and AMD, you signed that document, didn't you?

A. Yes, sir.

Q. And you had a meeting, I think you testified before we went into the in camera session, with Rambus representatives in April of 2000?

A. Yes.

Q. And it was shortly prior to that meeting that

you notified Rambus that AMD had made the decision to use DDR, correct?

A. It was --

MR. ROYALL: Your Honor, I don't know that there's foundation for that question.

MR. STONE: Your Honor, I --

JUDGE McGUIRE: Overruled. Go ahead.

BY MR. STONE:

Q. Do you have the question in mind, Mr. Heye?

A. Yeah. I'm not sure about the word shortly, but we had informed Rambus, I don't know exactly the time frame relative to when.

Q. Within a few months prior to that meeting?

A. Months, yes.

Q. And as soon as you -- within a few months, then, after you told Rambus that you were going to use DDR instead of RDRAM, Rambus came to you and said, well, we think that would infringe on some of our patents, didn't they?

A. Yes.

Q. When you made the decision to go to DDR, that was some time in the fall or summer of 1998, correct?

A. Yes.

Q. And you met with Mr. Sanders at that time?

A. Yes.

Q. And he was the CEO?

A. Yes.

Q. And the two of you made a decision in the summer or fall of 1998 to go to DDR and not use RDRAM, correct?

A. Correct.

Q. And you waited a year and a half or so to tell Rambus of that decision, correct?

A. I -- I'm not sure of the question, I did say months, but here's what I do not remember. I do not remember precisely when I told Rambus of our decision to go to DDR. It had to be -- they would have known about that decision as soon as we transitioned from doing the north bridge for supporting Rambus to DDR, so they would have to know within months of me making that decision that we were changing our minds, because we would have -- so, after going through the time line with you right now, you're correct, I told Rambus within a couple of months, it had to be within a couple of months of that decision that we were going DDR, because of the engineer effort.

Q. I'm not sure that now I understand. I want to make sure we're clear on this.

A. Yeah.

Q. You told me earlier that it was some time within a few months of your meeting in April of 2000 that you

told Rambus that you were going to go to DDR and Rambus then came to you and said that the use of DDR infringes on our patents.

A. I misspoke.

Q. So, you're now telling me that you told Rambus much earlier than that?

A. I'm telling you that -- and I'm hesitant to give you months, because I'm not clear. I can tell you in terms of sequence of events. Upon making the decision to go to DDR, we transitioned the team from doing a Rambus-based north bridge to a DDR-based north bridge. That event would have triggered us telling Rambus of that decision, because at that point, you know, they're meeting with the engineering team quite a bit. Those meetings terminated.

Q. Well, didn't you have two separate teams? Didn't you have one team working on the DDR north bridge and one team working on the Rambus north bridge?

A. We had -- we had one team working on -- we had -- we kept on going with the test chip for Rambus to keep on -- we kept on going on the technology thing just quite frankly the test chip -- I'm hesitant -- what I'm hesitating on is the term "design team." We had a full-fledged design team for the north bridge for DDR. We had some design engineers working -- if you recall I

said the first thing we were doing was a test chip for Rambus. That test chip team didn't disappear when we started focusing on the north bridge activity on DDR. We spent some time on that test chip. Now, I have to be honest with you, I don't remember exactly what we terminated that test chip activity.

Q. But it was not until after you terminated the test chip activity that you told Rambus of your decision, right?

A. I don't remember. I honestly don't. I'm sorry, I don't remember.

Q. Let me make sure if I can get one date accurate.

A. Sure.

Q. You did make the decision to go from Rambus to DDR in the summer of 1998?

A. Yes.

Q. And Rambus came to you some of time after you notifying them of your decision in April of 2000 and said that we think you're infringing, correct?

A. Yes.

Q. And at that time, you had a couple of options, didn't you?

A. Yes.

Q. You could go back to the Rambus chip, correct?

A. Yes.

Q. You could negotiate a license with Rambus, correct?

A. Yes.

Q. And you could go to the memory manufacturers and see if they thought it was going to be a problem? Right?

A. I could go and see if what was a problem, I'm sorry?

Q. To see if the fact of the use of DDR might infringe on Rambus patents was going to be a problem.

A. Yes.

Q. And so the first thing you did was you went and met with three memory manufacturers yourself, correct?

A. The first thing I did, as opposed to the other two?

Q. Yeah, you didn't try to negotiate a license agreement with Rambus, did you?

A. For DDR, no, I did not.

Q. And you didn't switch back to the RDRAM.

A. So then we did a -- right, so we -- correct, I went around to the memory vendors, that's right.

Q. And you went and met with Elpida.

A. Yes.

Q. And Samsung?

A. Yes.

Q. And Infineon?

A. Yes.

Q. And in the course of those meetings, did Elpida tell you that they had a license from Rambus for DDR so there weren't any problems, they would manufacture DDR as long as you wanted it?

A. No, my -- my recollection of the timing of that meeting was that Elpida was considering, although had not decided, at that moment in time, if they were going to sign a license agreement for getting the DDR IP.

Q. And they later did, correct?

A. They later did, yes.

Q. And you were pleased when they signed the license agreement?

A. I was ambivalent.

Q. Okay. And what did Samsung tell you? Did they tell you that they had signed a license agreement?

A. No, at the time, they hadn't signed it either.

Q. But they told you they expected to?

A. They didn't tell me that either. They said they were considering it.

Q. And they ultimately did sign?

A. Correct.

Q. And Infineon, to your knowledge, did not sign, correct?

A. At the time we talked to Infineon, they said they disagreed with some of the IP claims and they were going to challenge them in court.

Q. And Infineon told you that they thought the patents were invalid, didn't they?

A. I'm not sure they used those words. Again, I don't recall the specifics of the conversation. The specifics I take away was they had some issues with their patents and they were challenging them in court.

Q. And one of the things that you could have done that you didn't do was ask Rambus for a license, correct?

A. That's correct.

Q. Correct?

A. Yes, we did not ask for a -- we did not ask for a license from Rambus, that's right.

Q. And isn't it true that it is AMD's corporate policy when notified that you may infringe the intellectual property rights from others to seek to obtain a license?

A. I'm not aware of that as a formal policy.

Q. Well, I'm going to show you the AMD annual report for the year 2001.

A. All right.

Q. Marked as CX-2164.

MR. ROYALL: Do you have a copy?

MR. STONE: I do not, because it was not scanned by complaint counsel when they gave it to us, so all I can do is show it on the ELMO.

MR. ROYALL: This is, I take it, not on the exhibit list?

MR. STONE: It's on complaint counsel's exhibit list, CX-2164.

BY MR. STONE:

Q. Do you recognize the cover of this as the AMD annual report for 2001?

A. Actually, the picture I'm seeing it looks like an annual report cover, I can't see, you're on it, so I can't really say that it's on 2001.

Q. Let me zoom down on the year and show you.

JUDGE McGUIRE: Be careful there, Mr. Stone.

THE WITNESS: Yes, it says Annual Report 2000, so I'm sure that's it, yes.

BY MR. STONE:

Q. Let me go back, I want to turn you to one provision of this, the one I've labeled with patents, and you'll see, I'm going to give you my hard copy.

A. Thank you.

Q. But I just want to show you where it is on this, so we all have it, it's page 23 of the document.

A. Right.

Q. And it starts in the lower, left-hand column, "We may not be able to adequately protect our technology or other intellectual property in the United States and abroad through patents, copyrights, trade secrets, trademarks and other measures." Do you see that?

A. Yes, sir.

Q. And then on the carry-on up here, it says, what I want to direct your attention to, is the paragraph that begins, "From time to time." That reads, "From time to time we have been notified that we may be infringing intellectual property rights of others. If any such claims are asserted against us, we may seek to obtain a license under the third party's intellectual property rights. We cannot assure you that all necessary licenses can be obtained on satisfactory terms, if at all. We could decide in the alternative to resort to litigation to challenge such claims. Such challenges could be extremely expensive and time consuming, and could have a material adverse effect on us. We cannot assure you that litigation related to the intellectual property rights of us and others will always be avoided or successfully concluded."

Do you see that language?

A. Yes.

Q. Now, did you in 2000 --

JUDGE McGUIRE: You may approach, Mr. Stone.

MR. STONE: I'm sorry, Your Honor.

I just wanted to let you see in case you wanted to see the language.

THE WITNESS: Thank you.

BY MR. STONE:

Q. If it helps you. Did you in April of 2000 or thereafter seek to obtain a license from Rambus for any of its intellectual property in the event that anything AMD was doing might infringe on Rambus's patents?

MR. ROYALL: Could I ask for clarification whether this is a question as to what AMD did as opposed to what Mr. Heye individually did?

MR. STONE: Yes, it is about what AMD did.

MR. ROYALL: What AMD did, then I object for a lack of foundation.

JUDGE McGUIRE: Sustained.

BY MR. STONE:

Q. Let me ask it this way: To your knowledge, to your knowledge, Mr. Heye, did AMD seek to take a license from Rambus?

A. Not to my knowledge, no.

Q. Okay. To your knowledge, did AMD make a decision that if Rambus chose to assert those patents

against AMD, AMD was going to litigate?

A. I don't know.

Q. Did you make a decision one way or the other to either seek a license or choose to litigate?

A. No.

Q. What decision, if any, did you make after Rambus notified you about their patents?

A. I mean, our counsel, you know, our lawyers were given those patent numbers and were looking into it, and then after that, that was on the part of our general counsel office.

MR. ROYALL: Your Honor, I would just caution the witness to be careful about privileged issues.

MR. STONE: I am glad that he is represented by his own lawyer and Mr. Royall to caution him.

MR. ROYALL: Your Honor, I would ask that that statement be stricken from the record, and I do not appreciate those type of statements.

MR. STONE: Your Honor, I am not seeking anything that's attorney/client from this witness.

JUDGE McGUIRE: I realize that you're not doing that, but let's keep our tone here, Mr. Stone.

BY MR. STONE:

Q. I appreciate that, Your Honor.

Mr. Heye, after you were notified by Rambus that

AMD might be infringing on Rambus's patents.

A. Yes, sir.

Q. Did AMD suffer any adverse economic consequences?

MR. ROYALL: I object as vague and ambiguous as to whether these are consequences relating to that or something else.

MR. STONE: Relating to that.

THE WITNESS: To the best of my knowledge, there were no financial consequences.

BY MR. STONE:

Q. Have you suffered any delay in time to market for any products as a result of Rambus advising AMD that AMD might be infringing on Rambus's patents?

A. Not to my knowledge, no.

Q. Has any memory supplier told you that they would be unable to supply you with DDR DRAM?

A. Based on --

Q. On Rambus's patents.

A. Me personally, no. And to AMD, not to my knowledge.

Q. Okay. Earlier today, one of the things you were asked by Mr. Royall was backward capability. Do you remember that?

A. Yes, sir.

Q. When the K-7 was introduced, it was introducing SDRAM, am I right?

A. The initial chipset that supported Athlon was SDRAM-based, yes, sir.

Q. And it was a PC-100?

A. PC-100, that's correct.

Q. And then later you switched from the use of a PC-100 to DDR, correct?

A. That's correct.

MR. ROYALL: Your Honor, I believe that misstates the record, I believe the witness testified earlier today that there was PC-133.

MR. STONE: Whatever the witness thinks is the truth is what I would like.

MR. ROYALL: Well, I object, it misstates the record.

JUDGE McGUIRE: Well, if it misstates the record, restate.

MR. STONE: I don't know whether it does.

JUDGE McGUIRE: Well, we can ask Mr. Heye whether he can so indicate and whether that was your true and accurate testimony on that point.

BY MR. STONE:

Q. Let me do it this way, Mr. Heye.

A. Sure.

Q. When the K-7 was introduced, it used SDRAM?

A. Let me refresh my memory, when you're referring to it as opposed to the K-7, that's a misnomer, if you go back to the chart, the very first systems that supported K-7 were based off PC-100 memory, yes, that's correct.

Q. And who designed the memory system that was used with the K-7?

A. The first north bridge for K-7 was by AMD and that was PC-100 based, that's correct.

Q. And who designed the second north bridge?

A. The second north bridge came from a company called VIA and that supported both PC-100 and PC-133.

Q. And who designed the third north bridge?

A. Either ALI or SiS, I don't know which one. Capital S, small i, capital S.

Q. Am I correct that the only north bridge that was designed by AMD for use with the K-7 was designed to use PC-100?

A. Not the only, the first.

Q. Okay, the first, and then you later designed another one?

A. Yes.

Q. And when was that?

A. That was the -- we launched that product in

October of 2000 and that was the north bridge that supported DDR memory.

Q. Okay. So, the first north bridge designed by AMD was introduced into the marketplace in '99?

A. Yes, sir.

Q. And the development occurred as best you understand from some time prior to your joining the company, and you estimated '95, up until '99, correct?

A. That was the K-7 microprocessor, right, not the K-7 north bridge.

Q. When did they start work on the K-7 north bridge?

A. The K-7 north bridge is more -- that was right around when I joined, so that would have been '97.

Q. And what month was that in '97 when you joined?

A. I joined in June.

Q. So, the middle of '97 they started designing the north bridge that used the PC-100 and that was introduced in the marketplace in '99?

A. Yes, sir.

Q. And is it consistent with your recollection that the design of the north bridge took no more than 15 to 18 months?

A. The design took -- yeah, 15 to 18 months, yes.

Q. And that covers starting the designing up

through the time it's introduced into the market?

A. Yes, through the design, yes.

Q. And the next one that was designed by AMD for use on the K-7 with the DDR north bridge was started when?

A. In the '98 time frame.

Q. Okay. And it was introduced in the market in 2000?

A. Correct.

Q. And again, that next north bridge took in the range of 15 to 18 months from design to first sale?

A. Yes.

Q. And the second north bridge was intended to be used with the DDR?

A. That's correct.

Q. And the DDR-compatible north bridge that you designed could not operate with an SDRAM, could it?

A. That's correct.

Q. And you couldn't put a DDR product into your north bridge that was designed to use the PC-100, could you?

A. That's correct. That's right.

Q. So, you would agree that as to the two north bridges that were designed by AMD for use with the K-7, they are not backward compatible?

A. That is correct.

Q. And when you switched from the design for a PC-100 compatible north bridge to the design for a DDR-compatible north bridge, you had to switch some things other than just the north bridge, didn't you?

A. From a system point of view?

Q. Yes.

A. Oh, absolutely, yes.

Q. So, you undertook to change the motherboard?

A. Yes.

Q. And what else did you change?

A. Again, I just want to caution, when you say you changed, I would like to say, the changes that were required to implement the infrastructure, which is different than Rich Heye or AMD changing, but the memory, obviously the DDR memory is different from PC-100 memory, so the memory vendors had to have different memory. The dim standard is different for synchronous DRAM versus DDR, the motherboard is different and the north bridge is different. So, those all had to change and we facilitated that as opposed to actually doing.

Q. Now, when the change was made in the north bridge used with the K-7 to go from a PC-100 compatible product to a PC-133 compatible product, which you said

was done by VIA.

A. Yes.

Q. What changes were necessitated by that?

A. The VIA chipset required a different motherboard. That is the same -- the dim structure and the sockets were the same for PC-100 and PC-133.

Q. So, a new north bridge and a new motherboard?

A. Correct.

Q. One of the things you told us -- may I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. One of the things you talked about with Mr. Royall was whether the BIOS had to change. Do you recall that?

A. Yes, I did.

Q. Did the BIOS change when VIA designed the PC-133-compatible north bridge?

A. The BIOS had to change with the VIA chipset, that's correct.

Q. And did the BIOS change again when AMD designed the DDR-compatible north bridge?

A. Yes.

Q. And you told us earlier today, I think, that BIOS stands for binary input output software, correct?

A. Yes, sir.

Q. And that's wrong, isn't it?

A. I'm not sure. I thought it was right, but if you tell me it's wrong, I'll take your word for it. What do you think it stands for?

Q. I think it stands for basic input output system.

A. All right, I'll take your word for it. I'm not going to argue one way or the other. I've been calling it BIOS all my life and if that's what you say it is, I'll talk to my engineers and confirm that.

Q. Now, you told us earlier when you were at Apple -- if I may approach.

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. You told us that one of the things you worked on was the connection between the north bridge and the south bridge.

A. No, I said I worked on the PCI Bus.

Q. Isn't the PCI Bus what goes between the north bridge and the south bridge?

A. That is very Macintosh-centric, not AMD-centric.

Q. So, you used a PCI Bus at Apple for some purpose of connecting the north bridge and south bridge for a PCI-oriented chipset?

A. We used it to enable cards, just like I showed

you. It's to enable these connectors, that's right.

Q. And the PCI Bus that you persuaded Apple to use, you told us you had to get management approval at Apple to do that.

A. Yes, sir.

Q. Because Apple was in the process of making a big change while you were there from doing everything on an Apple proprietary system to trying to become more PC compatible, correct?

A. No. No, what I said was, I didn't say that at all.

Q. No, I didn't say you said that, I said while you were at Apple, Apple was in the process of making the change from relying on Apple proprietary systems to more PC compatible systems.

A. No, I didn't say PC compatible, what I would say is that Apple was on its peripheral card strategy, was adopting the same interconnect as the PC industry so it could leverage the PC -- the PC infrastructure for its I/O cards. The Mac is not very PC-centric.

Q. And it was trying to make the peripherals a part of that product more PC compatible. Is that a fair statement with which you could agree?

A. Yeah, the way I would phrase it is that, like I said before, the way I would phrase it, I guess we can

keep on restating our rephrasing, is that Apple was leveraging the PC interconnect and leveraging the PCI Bus in the PC industry for connector cards. I mean, the connector cards are not just for Macintosh, they're for the PCI Bus in a PC.

Q. And the PCI Bus that you had to get approval from management to use.

A. Yes, sir.

Q. Was a bus that was originally designed by Intel, correct?

A. It was specified to Intel, yes.

Q. So, they came up with the original specifications?

A. That's my understanding, that's right.

Q. And then a consortium of other people developed to continue to refine those specifications and they became part of the PCI consortium?

A. That's correct, yes.

Q. And was it a JEDEC specification?

A. No.

Q. Was it ever a JEDEC specification?

A. Not to my knowledge.

Q. And did it matter to you when you were at Apple that the PCI specification had been developed by Intel and then refined by a consortium of private companies?

A. It was important to me at the time that the PCI Bus was part of an open standards body.

Q. And it was important to you at the time that you would also have inneroperability for a lot of products with that bus, correct?

A. That's correct.

Q. And when you went to get management approval at Apple, who did you go?

A. A gentleman by the name of Eric Haslem [phonetic].

Q. And was that as high as you had to go? Did you have to go to -- was Mr. Skully in charge of Apple at the time?

A. Yes, Mr. Skully was in charge of the company, and I never personally talked to Skully about that decision. Whether Eric Haslem did or not, I have no knowledge.

Q. You described earlier the PCI consortium as something that you wanted to be part of the inner sanctum. Do you recall that?

A. Yes, sir.

Q. And your concern during the time you were at Apple was that if you weren't part of that consortium that the group might do things that were contrary to your competitive interests, correct?

A. Yes.

Q. One of the things that you were afraid was that they would know things about the design or specification before you would?

A. Yes.

Q. And when you agreed with Rambus to license the Rambus technology, did you get access to all of that technology?

A. Define "all that technology," please.

Q. If I can approach.

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. You told us earlier that Rambus gave you what you needed to plug into the north bridge design so that the interface with the Rambus DRAM would work, correct?

A. Yes.

Q. And they provided you with that technology, correct?

A. Yes.

Q. In the form of essentially drawings or a computer data that could be used in the design, right?

A. Yes.

Q. And you also understood that that information had been made available to Intel, correct?

A. Yes.

Q. And between AMD and Intel, that accounts for, what, 90 percent of the microprocessor market?

A. Yes.

Q. So, did you understand at the time that the two players in the microprocessor market both had access to the Rambus technology that they needed to design north bridge portions of the chipset to work with the RDRAM product?

A. I understood that Intel -- actually, I didn't know that Intel had, I made that -- Rambus would not tell me precisely what was given to Intel, so it's hard for me to tell you what Intel had, and the concern that I had was always that Intel had more than I did.

Q. You were concerned that Intel might have more than you and you were concerned that Intel might somehow be paying less for it than you were?

A. Yes, sir.

Q. And the concern that Intel might be getting a competitive advantage over you was a big worry?

A. It was a worry, yes.

Q. And one of the reasons you wanted to go out and support what you called Team DDR was that you wanted to have a product where you were part of the inner sanctum? Correct?

A. No, the real reason I went with DDR, quite

frankly, was I really believed that Rambus was going to fail in the market.

Q. You believed they were going to fail in the market?

A. Yes.

Q. And that was based on the information that you got that we heard about earlier from the memory manufacturers, the chipset manufacturers, and your people internally?

A. Yes, sir.

Q. And the memory manufacturers shared with you their projections for volume, correct?

A. Um-hmm.

JUDGE McGUIRE: I didn't hear that answer, Mr. Heye.

THE WITNESS: Yes, sir.

JUDGE McGUIRE: Okay.

BY MR. STONE:

Q. And their roadmaps?

A. Yes, sir.

Q. And you told us earlier that you didn't think that the memory manufacturers would ever share their roadmaps with each other, correct?

A. I would -- I had I said that the memory roadmaps were not the -- I do not believe that the memory

manufacturers would show the level of detail and specificity of their roadmaps that they shared with me to the general public. All had roadmaps on the web pages so of course the roadmap is out there in the public domain, but the level of detail that they were sharing with AMD was greater than the public statements.

Q. And you didn't think that the DRAM manufacturers would get together and share cost information either, did you?

A. Again, it's all the level and degree. I knew that there were manufacturers -- it's all degree. First of all, I have no idea what DRAM manufacturers' state of mind are, so that's the short answer, I just don't know.

Q. And earlier when you told us that you understood that the memory manufacturers were competing with each other, you don't know whether or not the memory manufacturers, in fact, have agreed on certain things, do you?

MR. ROYALL: Your Honor, I object to this on relevance grounds. The witness had mentioned something about why he chose to meet with memory manufacturers on an individual basis, he did not get into the sorts of issues that Mr. Stone is now attempting to probe in his cross examination.

JUDGE McGUIRE: Mr. Stone, response?

MR. STONE: Yes, Your Honor. The witness has said that he views the market as competitive. He has no foundation or basis for saying that. I want to bring out that he doesn't know whether they're competing or whether they're conspiring. I don't expect that he knows either way, but his testimony shouldn't stand that he uses them --

JUDGE McGUIRE: Overruled, Mr. Royall, I will hear the question.

THE WITNESS: I'm sorry.

BY MR. STONE:

Q. Do you have the question in mind?

A. Please rephrase it, please.

Q. Let me see if I can ask it this way: You don't know, do you, Mr. Heye, whether memory manufacturers have gotten together and agreed on anything that might impact the sale of memory products one way or the other?

A. That's correct, I don't know first hand that to be true.

Q. Was there a K-7 product or an Athlon product known as the XP?

A. Yes.

Q. And what was distinguished that from other Athlon products?

A. Let's see, the XP had -- I believe that's when

we first came out with 333 megahertz front side bus, I think that's the most distinguishing factor. It had a 333 front side bus -- I mean, one of the biggest differences on the XP was that it is a first-time we changed our nomenclature from megahertz to model numbers, so that was the way we -- that was the way we measured our performance change in the XP.

Q. And what memory would the Athlon Xp port with?

A. DDR.

Q. And was that a chipset designed by AMD?

A. When XP came out, there were chipsets from AMD, VIA, I think SiS and I believe nVidia at that time, all had DDR memory. I'm not sure about nVidia, they have memory today, I'm not sure when that came out.

Q. The RDRAM product that you signed the license agreement with Rambus to license their technology, you knew at the time that that was not standardized by JEDEC, correct?

A. Rambus, yes, I knew that.

Q. And you were -- you made some effort to become involved in ADT. Is that right?

A. Yes.

Q. And you knew that at the time you were trying to become involved in ADT that they were working outside of JEDEC?

A. Yes, sir.

Q. And you knew at the time that you were at Apple that much of the Apple system was proprietary to Apple and not standardized by JEDEC, right?

A. Absolutely.

Q. And the microprocessor that is designed by AMD today, it's proprietary to AMD, isn't it?

A. Yes.

Q. It's not standardized by anyone?

A. It's proprietary.

Q. And it's covered by AMD patents?

A. And a cross patent license, yes, and a cross patent license.

Q. And if Intel, for example, tried to come out with an identical version of the microprocessor that AMD is selling today, that is something that you would think they are not entitled to?

A. That gets down to the fact that we have a cross patent license with Intel and I am not an expert on that license, so I am the wrong person to have on the stand with respect to that question.

Q. I'm not going to push that if you don't know the answer. You told us earlier that the reason why you switched from producing Intel compatible microprocessors, correct?

A. We have to be real clear on the English on that one.

Q. If I can go back and if I can approach and just point you to this.

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. I think when we were talking on Mr. Royall's examination of you about the K-6 as a product that was or was not one that could be just put into a motherboard in place of an Intel microprocessor.

A. That's correct.

Q. It could be?

A. The K-6 could be.

Q. Okay. And so what I'm saying, I'm probably using the words compatible the wrong way.

A. Yes, you're making that -- again, the K-7 front side bus was not compatible with the Intel front side bus at the time, whereas the K-6 was, correct.

Q. And the reason AMD moved to the use of a front side bus that was not compatible at Intel, was because of patent litigation and other intellectual property litigation with Intel, correct?

MR. ROYALL: Objection, Your Honor, I believe this misstates the witness' prior testimony and there's no foundation for it.

BY MR. STONE:

Q. Sure, let me -- I don't mean to misstate, let me just ask. Why was it, to your knowledge, if you know, that when AMD went from the K-6 to the K-7, they went to a front side bus which was not compatible with Intel's?

A. It had to do with our cross patent license. I'm not sure there was any litigation involved. I mean, we had a -- it was a license negotiation and part of that negotiation was that the requirement was that we could no longer use the Intel front side bus in our microprocessors.

Q. Okay. Let me show you -- well, let me ask this: So, when you went to the front side bus that was no longer Intel compatible, you had to go around and persuade vendors to manufacture products that would work with your microprocessor.

A. Yes, sir.

Q. And that was the first big challenge you faced when you got to AMD?

A. Yes.

Q. And you had to get one of the chipset manufacturers, in fact, to manufacture for you by agreeing you would keep their identity secret and wouldn't let anybody know that they were manufacturing for you, correct?

A. Chipset, no.

Q. Motherboard?

A. The motherboard, that's true.

Q. Let me show you, if I can, a couple of documents on the ADT issue that you were asked about earlier.

May I approach, Your Honor?

JUDGE McGUIRE: Yes.

BY MR. STONE:

Q. Let me show you RX-1879.

A. Thank you.

Q. And if you can bring up whatever you can of the header part of this. Do you see at the top that this is a email that bears the date of August 17th, 2001 and it's addressed to Levi Murray and Richard Heye?

A. Yes, sir.

Q. And Levi Murray is someone who works with you, correct?

A. Yes, sir.

Q. And who is it from?

A. Quddus Mian.

Q. I was hoping that you could pronounce it and I could avoid it.

A. No.

Q. Mian Quddus?

A. I'll take your word for it.

Q. What I want to direct your attention to is it refers down below, if we can go back to the document, there's a heading, Attendees.

A. Yes.

Q. And it shows, Attendees, AMD, Richard Heye, Levi Murray and Samsung, that's what SSI stands for, right, Samsung?

A. Yes.

Q. Dieter Mackiewicz and Mian Quddus. Anyway, we see the names.

A. Sure, I see the names.

Q. And right under the summary, if you could go to the summary and just pull up that first paragraph under summary. It reports in this email that "Richard Heye stated that AMD has about 25 percent of the CPU market and their target is to be approximately 30 percent soon." Is that a correct statement of what you would have said or did say about AMD's market share in August of 2001?

A. Yes.

Q. And then help me just reconcile that with the numbers you gave us earlier that AMD's market share I think you told us earlier was over desktops and mobile about 17 percent.

A. Right. So, my comments would have been in

reference to desktop, although they didn't capture the desktop in their summary. And the only thing, too, was I was being a bit bullish, AMD can get about as high as 20 percent desktop and it has gotten down to about 15 percent now. So, I was being somewhat rounded to 25 percent. I wasn't sworn to -- I was being a bit aggressive in my salesmanship.

Q. Okay.

A. And our goal was to get to 30 percent, although we didn't hit that target, our goal was 30 percent, but we're not there yet.

Q. So, what has happened since 2001 is your share of the desktop market has dropped from where it was then?

A. It's gone from about, like I said, I haven't memorized the quarterly numbers. We were about -- the highest we've been is 23 percent, and this quarter was 19 percent, and that's a little on the low side. We're about 20 percent, maybe plus or minus a percent for the last couple of quarters.

Q. And what I want to focus on is it also reports that since Intel is part of ADT, they have a competitive advantage over AMD."

A. Yes.

Q. And that was your view in 2001, correct?

A. Absolutely.

Q. And that's why you made a strong effort to get into ADT as a developer, right?

A. Yes.

Q. And you wanted to -- they finally let you on the board, did they, or not?

A. No. Never did.

Q. Never let you in?

A. No.

Q. One other document I want to show you with respect to ADT, if I can, and this is RX-2106.

May I approach, Your Honor?

JUDGE McGUIRE: You may.

BY MR. STONE:

Q. And you'll notice this document talks about an ADT meeting with Intel, and it shows certain attendees and not you. Do you see that?

A. Yes.

MR. ROYALL: Well, Your Honor, if I can interject, I would ask that Mr. Stone, before he asks any questions about this, establish a foundation that this witness has seen this document before. Otherwise, I don't think it's appropriate to be asked about it along the lines of what I understand the questions have been recently.

JUDGE McGUIRE: Lay a foundation, Mr. Stone.

BY MR. STONE:

Q. I will, Your Honor.

And you'll notice that it's addressed to distribution.

A. Yes.

Q. But there's no distribution list attached to this document.

A. Yes.

Q. At least I couldn't find one. Is this a document that in the ordinary course at AMD you would have seen?

A. Well, I've got to read it a little more. I mean, in general, Levi had gone to dozens and dozens of meetings that I didn't always get the minutes. Depending upon the comment and what they wanted me to see, they could have sent it to me. Let's put it this way: I don't have a recollection of reading this document. I'm not saying I never got it, I just don't remember reading it. So, I would have to make a judgment of whether they sent it to me based on the content.

Q. Based on its content, can you tell whether it's something that your best recollection is that you saw?

A. You've got to give me a few minutes to read the

thing.

Q. Sure, whatever time you need.

A. Can I ask who sent this memo? There's no from either.

Q. I don't know. It was produced to us by AMD. You'll notice the AMD-0137 number.

A. Yes.

Q. But I can't tell you beyond that.

A. Okay.

Q. Is that jogging your recollection at all, Mr. Heye?

A. Well, let me get back, to answer your question is this something that they would have sent me, this is a -- I'll give it a 50/50, something they could have sent me. Something out of the -- it's not so mundane where if they sent it to me -- it's not something so mundane that would be sent to a vice president. At this level ADT stuff, they may have sent it to me, absolutely.

Q. But you don't have any recollection one way or the other?

A. I certainly have read documents about ADT, so it certainly may be a document that I have seen.

Q. Independent of this document, let me ask you a couple of questions.

A. Sure.

Q. Was it your understanding that if -- that if AMD was not permitted to join ADT.

A. Correct.

Q. And the ADT product came to fruition, that AMD might be required to pay royalties in order to use that product?

A. No, we were -- I don't have any recollection at all about worrying about royalty as an issue on ADT to be honest. That was not a concern.

Q. And was that because at that time you viewed whatever product ADT came out with as a product that would be purchased to work with one of your chipsets rather than you incorporating that technology into what you actually manufactured, you meaning AMD?

A. No. It's more complicated than you're making it out to be. The memory -- I don't know if you -- if you read, if I may reference the document you gave me, the structure of ADT had three structures, and I don't remember the names, board of directors, developers and participants. And if you read how those three are structured, if let's say, for example, ADT had taken off AMD could have joined a developer and participant and gotten actual IP and actually could have developed parts that were ADT compliant and furthermore, to the best of

my knowledge, and in fact he mentions here, that, you know, royalties was not part of ADT, that ADT was explicitly trying to get around the royalty of the Rambus patent. So, I can read that if you give me just one second. It said, i.e. did not like the Rambus model.

Q. Right.

A. So, royalty wasn't an issue, it was time to market.

Q. Let me ask you if you would to just look to see if it jogs your recollection, at the bottom of page 3, question number 10, and the top of page 4, answer number 10. You'll see where it says, "Royalty bearing, license to AMD versus non-royalty bearing (AMD's license) uneven treatment of IP." And then the answer, on the top of page 4, "Participants agree to license their technology, blocking IP at no charge. Developers won't pay royalties, participants could pay royalties if not cross-licensed with other developers (DRAM vendors)."

Does that at all refresh your recollection or your understanding that one of the things that you were trying to avoid that just being a participant in ADT could pay royalties?

A. Yes, they wanted us to be a developer is my recollection, so participants, you know, for those who

don't have the document in front of them, board of directors is the highest level, developers is the second ring down and participants was the third ring down. You know, we had actually been asked to join the developer level, so my expectation was that in the absolute worst case scenario, should ADT take place, and we couldn't get into the -- I forget what it's called, excuse me a moment, if we couldn't get into the board of directors, we would probably have joined the developers to not participate in any royalty arrangements, gotten actually the IP, done the -- you know, done the ADT development on our own.

Q. And when you said you voted on the membership, you didn't know what it was, and it had to be unanimous for you to be voted in. What was the vote for which status were you applying?

A. Board of directors.

Q. One last line of questions, Mr. Heye, and then we'll be finished here. You told us earlier today that it was the summer or fall of '98 when you made the decision to go with DDR, correct?

A. Yes. I believe I told you that.

Q. And pursuant to your license agreement with Rambus, you issued a press release announcing that AMD had chosen Rambus, correct?

A. Yes.

Q. And that press release announcing to the world that you had chosen Rambus came after or at the same time that you had already decided to go with DDR, didn't it?

A. I'm sure it was before, because we would have said -- we would have done the press release when we signed the contract, which was in December of '97, so I'm guessing, although I don't have the press release in front of me, although you may, but I'm guessing that the press release would have come out within weeks of us signing the agreement. You're shaking your head.

Q. October of 1998 is when it came out, Mr. Heye, isn't it?

A. If you say so. You may have the documents to back it up, I'm just explaining to you that I don't recall.

Q. This is a document that I pulled off the website last night from AMD. Do you have it up on the screen? Do you see the date on this press release, Mr. Heye?

A. Yes, I do.

Q. October 8, 1998.

A. Yes, I see it.

Q. You announced to the world in October of '98 that you had chosen Rambus, correct?

A. That's what the press announcement says, yes.

Q. And you have told us time and time again today under oath that in the summer or fall of '98 you had already decided to go with DDR, haven't you?

A. Yes, I did.

MR. STONE: I have no further questions of this witness, Your Honor, thank you.

JUDGE McGUIRE: Mr. Royall, redirect?

MR. ROYALL: If I could just have a moment, I need to locate an exhibit.

JUDGE McGUIRE: Go ahead.

MR. ROYALL: I'm sorry, Your Honor, something that his very last question raised prompted the need for an exhibit.

JUDGE McGUIRE: Go ahead.

MR. ROYALL: That I had a hard time finding.

REDIRECT EXAMINATION

BY MR. ROYALL:

Q. Mr. Heye, in his questions of you just this moment, Mr. Stone drew attention to the fact that you have said today that from your recollection, your recollection is that it was some time in the summer or fall of 1998 or the fall of 1998 that you recall making the decision to go with DDR as opposed to RDRAM. Do you recall that question of you just now?

A. Yes.

Q. Now, do you have a precise recollection as to when it was that you made the judgment and obtained management approval within AMD to go with DDR as opposed to RDRAM in future versions of the K-7 infrastructure?

A. No, I don't have the exact dates. I'm bad at recognizing dates. I remember sequence of events better than I do dates.

Q. Now Mr. Stone has shown you a press release dated October of 1998, do you know whether Rambus -- I'm sorry, did you know whether AMD made the decision to go with DDR as opposed to Rambus before or after this press release?

A. I would be pretty -- I believe that if we made this press release in October 1998, that at that time we believed that we were going to be Rambus-based memory subsystem and that my recollection would be faulty. Because we would not go out and mislead our customers or partners.

Q. So, it's your recollection that it was some time after this press release that you went to the CEO of AMD and got approval to go with DDR? Is that your testimony?

A. AMD would not deliberately -- let me be real clear, AMD would never deliberately mislead the press

analysts or the public of our intentions. So, to the extent this press release came out in October of '98, and I am sure that at this time frame we did not -- we did not make the DDR decision, I must be incorrect in my time reference in terms of when we made the DDR decision.

Q. Now, I would like to approach, Your Honor. I don't have multiple copies of this exhibit, because it just came up, but we'll flash it up on the screen.

Could I approach?

JUDGE McGUIRE: Go ahead.

BY MR. ROYALL:

Q. Mr. Heye, I have just handed you what's been marked for identification as CX-2158. I'll give you a moment, but I guess my first question is do you recognize this document?

A. I remember seeing this document in the documentation prior to coming here, but I didn't remember this document prior to this other than looking at it a few days ago.

Q. Now, let me ask you if it refreshes your recollection, do you see the subject line of this email, CX-2158, which was written by Steve Polzin, or it's from Steve Polzin, the subject line is "DDR and Rambus History." Do you see that?

A. Yes.

Q. And do you see, then, the discussion below on the first page of CX-2158, where it says, about roughly halfway down, "decided to stick with SDRAM PC-100." Do you see that language?

A. Yes.

Q. And do you see, then, that there's a reference below that to Rambus license signed and then question mark.

A. Right.

Q. And then below that, do you see IGR4 started in earnest in September 98. Do you see that?

A. Yes.

Q. And do you have an understanding of what IGR4 refers to?

MR. STONE: Your Honor, if I can object. Mr. Royall objected earlier to my questions on the grounds that I hadn't laid a foundation that the witness had seen the document before. Now, the witness just said I never saw this document at the time, that it was shown to me this past weekend before I came here. That's his testimony just now. That's not a foundation for showing him the document, if the rules are going to be applied the way Mr. Royall asked they be applied to me.

Now, I know we've used documents that witnesses

haven't seen before in efforts to refresh their recollection, but Mr. Royall is asking what words in the document mean, and that's not to refresh the witness' recollection, that's to get him to testify about a document for which he has not laid a proper foundation.

JUDGE McGUIRE: Mr. Royall, response?

MR. ROYALL: Your Honor, I am happy to move on, I am seeking simply to refresh his recollection with the document and I will focus him now on specific language.

JUDGE McGUIRE: Don't ask him what that word means.

MR. ROYALL: I will withdraw that question.

JUDGE McGUIRE: Noted.

BY MR. ROYALL:

Q. Now, going further down in the document it says, "Between summer '97 and fall '98 JEDEC-DDR turned around.

A. Yes.

Q. And then there's more language and below that it says, "In late 1998 AMD traveled to Micron, PC in Minn to see their DDR chipset." Do you see that?

A. Yes.

Q. Now, let me ask you, does that language refresh your recollection at all as to when it was that AMD made the final decision to go with DDR as opposed to Rambus?

A. Yeah, yes, it does. I just -- I misspoke. I'm frustrated with my memory right now. The reality is as I'm looking at this, again, the sequence of events which I remember but I got the dates wrong was we started working on the -- so, we started working on the Rambus-based chipset in October of '98, which was exactly when that press announcement was in October, so that presence announcement is consistent with what AMD did, which is inconsistent with my memory. So, I apologize for my bad memory and I apologize to the court.

Q. So, once you started on that work in October of '98, how long did you continue down the path on developing an infrastructure for RDRAM before you made the decision to go with DDR instead?

A. It was months, but I don't know how many months.

Q. So, you mean months after October '98?

A. Yes.

Q. Okay. Thanks. Now, when you first started dealing with Rambus, with the possibility in mind of licensing Rambus's technology for use in conjunction with AMD microprocessors.

A. Yes.

Q. Do you know whether it was known to Rambus through your communications with Rambus or whether it

was stated to Rambus that AMD had considered or even was considering the possibility of using DDR memory as opposed to Rambus?

A. I honestly don't know who told me that or not.

Q. You don't know?

A. I don't know.

Q. Now, Mr. Stone asked you whether the -- if I recall this correctly, whether the assertion of patent rights by Rambus over the JEDEC standards DDR and SDRAM standards had had any actual adverse consequences to AMD. Do you recall that?

A. Yes.

Q. In his questioning. And I believe you said that it did not have any actual adverse consequences. Is that right?

A. Yes.

Q. Would it have had adverse consequences on AMD if in response to Rambus's patent claims the JEDEC standards had been changed?

A. Yes.

MR. STONE: Objection. Calls for speculation. Lacks foundation.

JUDGE McGUIRE: Sustained.

BY MR. ROYALL:

Q. Were you concerned, sir, in terms of your own

concerns about AMD's business and the aspects of the business that you are responsible for that changes in the JEDEC standards might adversely impact your company?

A. Yes, I was concerned.

Q. Now, Mr. Stone also asked you about changes that had been made in various chipsets supporting the K-7 microprocessor. Do you recall that?

A. Yes.

Q. And he asked you about the change going from PC-100 to PC -- was it 133?

A. Yes.

Q. And he asked you about potential changes to the BIOS as well as the chipset. Do you recall that?

A. Yes.

Q. Now, in any of these changes that have actually been made in the life cycle of the K-7 infrastructure, from one chipset to another or one memory to another, have any of those changes been detrimental to AMD in terms of time to market?

A. Could I ask you to restate that question slightly different to make sure I understand the question that you're asking me.

Q. Let me withdraw the question, it was a little unclear. In the changes that have been made to the chipsets that have supported K-7 infrastructure.

A. Yes.

Q. During its product life cycle, has AMD been forced to make any changes that were outside of what it was planning to do in terms of its own roadmaps for the development of that chipset program or, I'm sorry, that microprocessor program?

A. No, I mean, we weren't forced to make any changes. I mean, we continuously made improvements to the platform which would cause us to change the platforms, but --

Q. And was that consistent with your preplanning for future versions of the infrastructure?

A. Yes.

Q. And had you ever had to make changes to the K-7 infrastructure in a way that was at odds with what you were planning to do in a material way?

A. Give me a moment to think. No, we never had any changes that were at odds. The reason I'm hesitating is at times we would have goals of having infrastructure available at a certain time and for engineering reasons it would slip out. So, we would sometimes not be able to accomplish our goals, so to the extent we would have technologies in the market longer than we had liked, that did occur, so it wasn't like we executed falsely to our roadmap, but to the extent that we had a roadmap and

to the extent that in the absence of bugs or nonplanned slips, we did okay. We weren't forced to make changes.

Q. Do you recall that you testified earlier about a number of concerns that you personally had in terms of how changes in the JEDEC memory standards might adversely impact AMD's business. Do you recall that?

A. Yes.

Q. And you talked about things like competitive disadvantages and opportunity costs and inventory issues. Do you recall that?

A. Yes.

Q. Have any of those concerns, those same types of adverse effects on AMD been -- have any of those types of concerns been caused as a result of the actual changes in the chipset that you have made on K-7 moving from PC-100 to PC-133 to DDR?

A. I hate to do this to you, but could you restate that question?

Q. Sure. In your earlier testimony in response to my questions, you had laid out the concerns that you had that caused you to believe that changes in the JEDEC SDRAM and DDR standards would adversely impact AMD's business. Do you recall that?

A. Yes. That's correct.

Q. And you talked about the disruption that would

be caused by having to redesign chipsets and delays in validation and time to market and competitive disadvantages and other things. Do you recall that?

A. That's correct.

Q. Now, when the actual changes were made in the chipset supporting K-7 to go from PC-100 to PC-133 and to DDR.

A. Yes.

Q. Did any of those adverse consequences that you talked about in your earlier testimony, did they -- did -- was AMD forced to suffer any of those adverse consequences when it made the change, the evolutionary change from PC-100 to 133 to DDR?

A. No, we never had a -- we did not suffer any consequences. If I may, the key here is that there's no question that you have to change things when you go from let's say PC-100, PC-133 to DDR. The thing is you want to manage it, in other words you want to work with all your partners because, you know, they're right, you know, DDR is different from synchronous DRAM, you did need a different chipset, you did need a different memory solution. The trick is to work with all of your partners in lining it all up such that you put a lot of engineering effort and validation effort such that when you do make that transition, it is not a major

disruption to the supply chains to your customer, i.e. the ones manufacturing PCs. That's your goal. And if you manage it and properly communicate with all your partners, the plans, you can manage change, because change does occur. The trick is all of the sudden when you're almost ready to launch and then you change something unexpectedly, then you have to go back and redo those plans and that's the adversity that I was referring to would occur.

Q. And that is a major disruption to the supply chains in response to an unexpected change, is that the type of concern that you had with respect to the potential of the DDR and SDRAM standards being changed in response to Rambus patent claims?

A. Yes.

MR. ROYALL: No further questions, Your Honor.

JUDGE McGUIRE: Mr. Stone, recross?

MR. STONE: Yes, thank you.

RECROSS EXAMINATION

BY MR. STONE:

Q. You still have Exhibit 2158 in front of you? The one that Mr. Royall showed you just a moment ago.

A. From --

Q. The Steve Polzin email.

A. Yes, I do.

Q. And that's addressed to Dirk Meyer as well as Levi Murray, correct?

A. Levi, yes.

Q. Levi Murray. Who is Dirk Meyer?

A. Dirk Meyer is my boss.

Q. If we can use this now, if we have a new set of dates to work with, I just want to go back.

A. Sure.

Q. And see if we can get this down once and for all.

A. Okay.

Q. You made the decision to switch from Rambus to DDR after you made the visit to Micron, correct?

A. Yes.

Q. After you visited the PC -- the other DRAM manufacturers, correct?

A. Yes. Well, say that again once more, I'm sorry.

Q. After you visited the other DRAM manufacturers.

A. Yes.

Q. And after a complete design team freed up in early '99. That reference is on the second page of the document.

A. Unfortunately I don't --

Q. Mr. Royall only gave you the first page?

A. Yes, sir, I apologize, I only have the first

page.

Q. Let me give you the second page, if I can approach.

A. Thank you, sir.

Q. Take a look at the second page of the document.

A. Yes.

Q. Can we bring it up? Yeah. Do you see where it says, "A complete design team freed up in AUS (EPD) in early '99?"

A. Yes.

Q. And so you didn't make the decision, then, to switch to DDR based on this document, at least, and let me ask if that refreshes your recollection, until some time in '99?

A. Early '99, yes.

Q. Okay. And you were able to design a new chipset, a new set of DIMMs, a new memory controller, and a new motherboard after you made the switch beginning in early '99 in order to power up by December '99, correct?

A. Your question had a lot of -- just take it back a step.

Q. Let me back up. One of the things that you had to do when you switched from the Rambus plan to the DDR plan was you had to design DIMMs to replace the RIMMs.

A. Again, this is where we have to be real clear. It depends who we is. We would include me, which is not me. In other words, the we in the context you're talking about are the memory vendors and the dim manufacturers. So, when you ask the question did -- AMD doesn't design DIMMs, and there was clearly work, according to Steve's memo, that Micron had DDR DIMMs up and running in late '98. So, that doesn't surprise me that any modifications that they would have had to make those DIMMs between late '98 and late '99 could occur, yes.

Q. Okay, so let me --

A. So, on the dim side.

Q. I understand the problem with my question, and let me ask it this way: Certain changes needed to be made by someone.

A. Yes.

Q. When AMD made the decision to switch from RDRAM to DDR?

A. Yes, sir.

Q. One of those changes was to make sure that you had DIMMs to replace the RIMMs?

A. Yes.

Q. One change was to redesign north bridge.

A. Yes.

Q. One change was to redesign the motherboard.

A. Yes.

Q. And one change was to possibly modify the BIOS.

A. Yes.

Q. And Mr. Royall asked you about all those changes earlier.

A. Yes.

Q. And those changes were accomplished so that you could power up a K-7 DDR design by December of '99, correct?

A. The first instance of those changes were made available -- the first prototypes were in '99, don't forget it took about ten months, so it went from '99 to power-up to October of '2000 to ship. So, this was confirmed by conversation, and I also said that the odds are when you turn on everything it all works at the same time, is pretty slim, but I don't know the details, but I know, for example, that our chipset when powered up was good enough to run a lot of stuff, but not good enough to ship. We found some bugs in it, I couldn't tell you what they are.

Q. And the decision that you made to go from RDRAM to DDR was one that you were able to plan for and accomplish in a time frame consistent with AMD's business goals, right?

MR. ROYALL: Objection, vague and ambiguous.

BY MR. STONE:

Q. Let me rephrase. When you made the decision to go from RDRAM to DDR, you did it voluntarily?

A. Yes.

Q. Nobody said, we're going to sue you if you use RDRAM?

A. That is correct, no one said those words.

Q. And it cost you money to make the change.

A. Yes, sir, it did.

Q. And you would have spent less money if you had continued to use RDRAM?

A. No. I don't believe that.

Q. Okay, because you believe in the long-term, you would have had issues?

A. That's correct.

Q. And when you made the change to go from RDRAM to DDR, you started it some time in early '99, correct?

A. Yes.

Q. And you were powered up by December of '99, correct?

A. Yes.

Q. And you were shipping product in 2000, late 2000?

A. Late 2000, yes.

Q. And if there were any of the issues about opportunity costs and inventory and the other issues that Mr. Royall asked you about, those costs weren't ones that AMD incurred as the result of switching from RDRAM to DDR, correct?

A. The context of that question versus the context of his question was different.

Q. Okay.

A. So, that's not a yes/no question, from my perspective.

Q. Explain it, please.

A. So, again, you know, we kept on saying, if you look at from the time we were talking about it, which was early '99, to the time we ship, which was late 2000, again, that's about an 18-month window from let's do this to the time we're actually doing it. So, that's the 18-month cycle. And it's also consistent that it takes about, in this case, December to October, so about ten months from the power-on to actually getting the system out there.

And, you know, there's no lost opportunity cost here because this is real -- this is bringing real value to the customer. You know, no one's -- I don't think anyone would state technically that it's not good for the consumer to have higher speed memory. And the

question in that time frame was, with Rambus versus DDR and for the last five hours my testimony has been that DDR was better. So, when we spent all those 18 months, that wasn't to me lost opportunity cost, that was good old-fashioned business acumen and engineering.

Opportunity costs would have said, we've done this for let's say 15 months and then three months before you ship, now I have to -- again, we're going down this whole supposition path, three months before I ship now you've made a change which has no real end user benefits but to get around these patents we had to make these changes and that slips everything out. Now I'm using the engineers not to finalize and get all this stuff to market, I'm spending it and moving on to next generation projects.

The opportunity cost now is that any time after October these engineers are working on the problem, that's the opportunity cost question which is different. So, that's why opportunity cost versus doing the actual business doesn't make sense.

The same with inventory, again, when you're three months prior to shipping a product and you're a semiconductor manufacturer, not a motherboard, but a semiconductor manufacturer, you have already started wafers, production wafers. So, again, like I said

earlier, the whole thing with doing infrastructure change, it's hard, I'm not going to deny it was a change to the industry, it was a change, it was gosh darn hard, quite frankly, a lot easier than Rambus, but it was hard and you have to plan it and manage it and it takes 18 months for everything from Team DDR to engineering stuff to make that happen.

When you do all that work and then you change last minute, that's where all the problems are and that's where your inventory hits. And so that's why we asked for the context, pre-18 month loss is a very different discussion of just engineering 101 and getting everything ready for planning versus now I'm about to ship and changing everything on you. I'm sorry.

Q. I appreciate that. And for a DRAM manufacturer before they go to silicon to avoid inventory costs, they're going to need at least three months lead time, right?

A. Three months prior to shipping they must start production of wafers.

Q. Just one more set of questions. Did you make the trip to Micron in Minnesota that's referred to in this document?

A. No.

Q. Do you know if the people who visited Micron

that's referred to in this document were told by Micron, well, you know, one issue is, Rambus has told us that there may be some patents that Rambus has that relate to SDRAM DDR?

A. I don't know if they were told that my Micron and if they were, they certainly didn't tell me.

Q. And did your JEDEC representative at any time after you joined AMD tell you that he learned at JEDEC meetings about Rambus patents?

A. Not to my recollection, no.

Q. When you looked at this document earlier today when Mr. Royall showed it to you and it allows you to conclude right away that you had not told Rambus -- you had not made the decision to go to DDR in the summer or fall of '98, correct?

A. Yes.

Q. All you had to do was pick up this document and say I can see now that summer and fall of '98 is not the right date.

A. Yes.

Q. And you saw this document most recently when?

A. A couple of days ago, which I probably didn't read it that carefully because I had figured out the dates already.

Q. And who did you see it with?

A. As is common practice, you get a set of documents that might be shown, so you look at them.

Q. Who put the set of documents together for you?

MR. ROYALL: Your Honor, I would caution the witness if there are any attorney client issues, his lawyer can speak if there are, but I just would do that.

JUDGE McGUIRE: Well, his attorney is here, he can speak for his client.

MR. BEDDO: Well, I can make the same caution not to disclose attorney/client privilege.

JUDGE McGUIRE: Sir, what is your name again for the record?

MR. BEDDO: B E D D O, David Beddo.

JUDGE McGUIRE: Okay, thank you.

BY MR. STONE:

Q. Let me ask it this way, Mr. Heye, let me withdraw that.

JUDGE McGUIRE: He is going to withdraw the question.

BY MR. STONE:

Q. Let me ask you, was it your understanding that you were given a set of documents to review over the weekend that you understand might be shown to you by complaint counsel when they examined you during the course of this proceeding?

A. Those words were never said. I mean, I think --

Q. I don't want to know the words, I don't want to know the words that were said, I just want to know your understanding. Wasn't it your understanding that these were documents that complaint counsel might show you during their examination of you?

A. Actually, it was more --

MR. BEDDO: Your Honor, I would like to caution the witness again not to disclose attorney/client communications.

JUDGE McGUIRE: All right, he has been cautioned. So, we don't need to stand up and say I want to caution him again. He's been cautioned. So --

THE WITNESS: I've been cautioned. I don't know who I ask this question to, but when you have a conversation with your lawyer, that's attorney/client privileged issues and I'm not supposed to --

BY MR. STONE:

Q. I am not asking you what was said to you by your lawyers at all. I just want to know was it your understanding that the set of documents that you were given to look at were documents that you understood complaint counsel might show you and ask you about in this proceeding?

A. Actually, to be quite honest with you, it was

more my understanding to help jog my memory, because I have the history of a bad memory, which has been documented for the last two hours.

MR. STONE: Thank you.

JUDGE McGUIRE: Thank you very much. That concludes your testimony from today and, sir, thank you very much, you are excused from this proceeding.

THE WITNESS: Thank you.

JUDGE McGUIRE: Counsel, I am going to suggest that we take a very short, five-minute break, and when we return, you can call your next witness. Before we do that, though, let me just inquire how much time do we anticipate that we are going to spend with the next witness?

MR. ROYALL: I will let Mr. Davis comment on that.

MR. DAVIS: I think his testimony should be an hour and a half to two hours.

JUDGE McGUIRE: That is going to put us at 5:30, and then some idea on the cross?

MR. STONE: I'm not sure until I hear the direct, but maybe an hour.

JUDGE McGUIRE: So, we're going to be here until probably close to 7:00. Is that correct? I'm trying to get everyone on the same page and make sure that's

doable. I am going to urge both sides for the sake of everyone concerned that we attempt to expedite each side's inquiry to the extent possible. All right, let's take a quick five-minute break. Off the record.

(Whereupon, there was a brief recess in the proceedings.)

JUDGE McGUIRE: On the record. Mr. Oliver, did you have a comment you wanted to make?

MR. OLIVER: I simply wanted to introduce to the court Mr. Robert Davis who will be handling the next witness on behalf of complaint counsel.

JUDGE McGUIRE: All right, Mr. Davis, thank you.

MR. STONE: Could I just move in some exhibits, Your Honor, from the last examination. I would offer the AMD annual report, CX-2164, the presentation of October '98, RX-2302 and the Polzin email, CX-2158.

JUDGE McGUIRE: Objection?

MR. OLIVER: We have no objection to CX documents, which was the third one, please?

MR. STONE: RX-2302.

MR. OLIVER: If I could perhaps just see if I could pull that up at some point during the afternoon.

JUDGE McGUIRE: The other two are entered exit.

(CX Exhibit Numbers 2164 and 2158 were admitted into evidence.)

JUDGE McGUIRE: Oh, are you still looking for the RX?

MR. OLIVER: I was going to try to pull it up while the proceedings are going on.

JUDGE McGUIRE: Then you may proceed at this time, Mr. Davis.

MR. STONE: I should clarify, that was the press release, Your Honor, that we showed him. RX-2302 would be the hard copy press release that I used.

MR. OLIVER: Perhaps if we could look at that.

JUDGE McGUIRE: Let's go off the record for a moment so all this chatter doesn't get on.

(Discussion off the record.)

JUDGE McGUIRE: On the record, you may call your next witness, Mr. Davis.

MR. DAVIS: Thank you, Your Honor. Complaint counsel calls Mr. Barry Wagner.

JUDGE McGUIRE: Mr. Wagner, would you approach and you will be sworn in by the court reporter.

Whereupon--

BARRY WAGNER

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

JUDGE McGUIRE: If you could have a seat over there, Mr. Wagner.

All right, Mr. Davis, proceed.

DIRECT EXAMINATION

BY MR. DAVIS:

Q. Thank you, Your Honor.

Please state your name for the record.

A. Barry Wagner.

Q. And where are you currently employed?

A. NVidia.

Q. And what is your current position?

A. Manager of technical marketing.

Q. And how long have you been at nVidia?

A. About six years.

Q. Could you describe your educational background, please.

A. Bachelor of science in electrical engineering and a minor in computer science.

Q. Thank you. We are going to talk a little bit about nVidia's background now, for the next set of questions. What is nVidia's line of business?

A. We make reference processors and chipsets for the PC workstation in game console markets.

Q. What is a graphics processor?

A. A graphics processor is the piece of silicon that's generally put on a graphics card and it's what ultimately you see on the displays of the computer.

Q. Does nVidia fabricate the graphics processors they design?

A. No, we do not.

Q. And who fabricates those?

A. We have a source of graphics, TMPC and we have relationships with IBM.

Q. And who validates the graphics?

A. A graphics card, primarily.

Q. What is a graphics card?

A. A graphics card is it's a PC board, it plugs typically into a graphics slot in the motherboard, and then ultimately your monitor hooks up to this graphics card.

Q. Does nVidia manufacture its graphics card?

A. We manufacture some, we mostly sell just the chip to other people who manufacture the card.

Q. And who are some of the firms that build nVidia graphics cards?

A. When we build it, we typically outsource our manufacturing to somebody like a Flextronics, for example, other companies, primarily in the far east, will build the card and sell it into a lot of OEM markets, customers like Dell or HP or others will buy the graphics cards from them.

Q. I would like you to look at what's been marked

for identification as CX-2833.

Your Honor, that's a document that's not in the database, it's a glossy brochure. May I approach?

JUDGE McGUIRE: You may.

BY MR. DAVIS:

Q. Do you recognize this document?

A. Yes.

Q. And what is CX-2833?

A. This is a kind of a product guide that the marketing group put together. It was used around the time of the launch of one of our products and it talks about our partners, some of our customers, some of the industry groups that we work with to try to enable better graphics in the industry.

Q. Thank you. If you could turn to the third page, that page with the table entitled nVidia Workstation Graphics Solutions.

A. Okay.

Q. And what is being described in this table?

A. This is -- it's a table of our workstation graphics product line, it shows several chips that are in different segments of the market, some different features about the graphics card, just some of the functionality that those chips have.

Q. If you look at the second column of the table,

the column with the title Memory, what does that refer to?

A. That's the memory that is on the graphics card.

Q. And what kind of memory is used in these graphics cards?

A. Most of them are DDR, one of them says SDR.

Q. And who buys these products?

A. Typically customers for our workstation graphics cards are engineers that need to produce some sort of hardware to model it on a computer. That's the kind of market that we sell the workstation products to.

Q. And what is a workstation?

A. A workstation is basically a PC, it's usually a higher performing PC. It often times will have multiple processors in it, things to improve its performance, and other than that it's very much like a PC.

Q. If you can turn three pages later to the page with the table titled nVidia Consumer Graphic Solutions?

JUDGE McGUIRE: What page is that, Mr. Davis?

MR. DAVIS: Unfortunately, the document doesn't have any page numbers, that's why I'm sort of flashing it. It's about six pages into the document.

JUDGE McGUIRE: I've got it.

BY MR. DAVIS:

Q. What's being described in this table?

A. It's another list of products that we have available, in this case in the consumer graphics market, again, it's broken out by segment that we sell these chips into.

Q. And what do you mean by segment?

A. nVidia produces a lot of different graphics chips, each one of them is a different level of complexity, different cost structure, we sell them into different segments of the market. Somebody in the value end of the market, that's typically something like a $79 to $99 graphics board. That's all those users willing to pay for that kind of product so it's lower performing. The Enthusia is the higher end typically, and they are more expensive boards, they can be as much as $300, $400.

Q. In this column is the header Recommended Memory Config on the second page of the table. Do you see that?

A. Yes.

Q. And what does that reflect?

A. I'm sorry, what does it?

Q. What does that column reflect?

A. I'm sorry, it's recommended memory config. We don't actually sell the graphics board in this case, so what you have here is a list that we recommend for each

one of these segments for how much memory should go onto the graphics card, and that's typically what the user in that market segment would expect to find if they bought a graphics card in that class of product.

Q. Well, in the lower right-hand corner of the document is a -- there's a picture in the lower right-hand corner of the page, can you describe what that is?

A. That's a picture of a graphics card.

Q. And I think everyone has probably seen a number of graphics cards and other cards, we'll speed up. Could you describe what some of the blocks on those -- on that video card are?

A. The chip in the middle of the card with our logo on it, that's our graphics processor. To the right there's four memory devices that are all on the graphics card.

Q. Is there a memory controller on the video card?

A. Part of our graphics processor has the integrated memory controller, that hooks up to the memory interface.

Q. Could you turn three more pages to the table entitled nVidia Platform Processor solutions.

A. Okay.

Q. What's being described on this table?

A. It's another list of products that are available, in this case this is for our chipset business. We have essentially a north bridge/south bridge kind of function, and again, it's some features and functionality for each of those platforms.

Q. Do you need me to indicate what north bridge and south bridge are?

JUDGE McGUIRE: No, I'm sorry, I've had enough.

BY MR. DAVIS:

Q. In the second to the last column of the table is a column titled Memory Controller. Do you see that?

A. Yes.

Q. In the first two rows of that column are entries that include the parenthetical, "Dual 64 bit DDR." Do you see that?

A. Yes.

Q. Could you describe what that means?

A. These product lines have 128 bit memory bus, and the data bus with it. We've architected it as two 64-bit memory controllers, each one of those 64 bits hooks up to a JEDEC standard DDR dim module.

Q. And who would buy the products described in this table?

A. One of our customers today is HP, they buy this product and put it into a motherboard. These particular

chipsets work with AMD processor base.

Q. Now, if you could -- Your Honor, I would like to -- I'm sorry, could you turn to the page entitled nVidia Mobile Graphics Solutions.

A. Okay.

Q. What's being described in that table?

A. It's a list of graphics products that we sell into the notebook market. Again, it's broken out by segment and lists a variety of features for each of those, the segments.

Q. And how are these processors different from the processors that you described earlier?

A. The primary difference is they're focused on lower power. We try to reduce power to maintain battery life in the notebook, it's a little different from a desktop in that sense.

Q. And who would buy the products described in this table?

A. Our biggest customers in mobile today are I believe Toshiba and Dell.

MR. DAVIS: Your Honor, I would like to move CX-2833 into evidence.

JUDGE McGUIRE: Any objection?

MR. STONE: No objection.

JUDGE McGUIRE: So entered.

(CX Exhibit Number 2833 was admitted into evidence.)

BY MR. DAVIS:

Q. I would like to return to your personal background and your work at nVidia. What do you do as the manager of technical marketing?

A. I have a group that works with various technology segments of our business to try to get the technology advances that we want to see into our products sooner. So, we will work with in my case the memory industry a lot to drive our requirements for future products into their roadmaps and hopefully keep our companies aligned with each other in terms of when products come on the market.

Q. And why do you work with the memory manufacturers?

A. The memory industry is -- it puts most of its emphasis into the activities in JEDEC to drive the dim sockets for the motherboard and what those definitions look like. In the graphics space, they don't really put a lot of effort in JEDEC focused on our specific needs. So, we work directly with the memory vendors to define our requirements for them and ultimately get hopefully better products for the graphics market out as well.

Q. Do you work with JEDEC at all?

A. I do work with JEDEC as well. It's one of my responsibilities.

Q. And why, do you have an understanding of why that's part of your responsibilities to work with JEDEC?

A. Yeah, so we originally got involved with JEDEC, partially on request from memory vendors. We had -- we were making requests of them for certain features that JEDEC was trying to eliminate, they encouraged us to participate in JEDEC to ensure that those features didn't get dropped from the standards. So, we got involved at that point. We got much more actively involved in --

JUDGE McGUIRE: Now, what point in time are you talking about, Mr. Wagner?

THE WITNESS: We joined -- I'm not exactly sure when we joined, but this would have been -- I forget the time exactly, but I remember the feature set that they were trying to drop. There was a function, a data mass function that was going to go away possibly in the standard, we wanted it, so at that time we got back in. I don't recall if it was --

JUDGE McGUIRE: So, you can't give me some time frame?

THE WITNESS: It was probably in the '99, 2000, it might have been sooner, I'm not sure.

JUDGE McGUIRE: All right, go ahead.

BY MR. DAVIS:

Q. Why was it important to you that you have an understanding of what memory technologies were in that -- in the JEDEC standard?

A. At this time we got involved more directly with JEDEC is when we started working on the chipset side of our business. The PC main memory platforms, ultimately if we're going to have to produce a chipset, we need to be using what that standard is, otherwise we're just not cost effective in the market. So, as the main memory standards were evolving, we wanted to be sure that we knew exactly what they were going to be. So, we built the chipset business, the standard product line that matched up with that technology requirement.

Q. Was there any relevance to the video card set of the business to your membership at JEDEC?

A. The video card side was -- it was one getting involved in JEDEC, you know, helped make it much more common that we wanted this commodity device as well. We worked with many vendors offline, the graphics piece of it is unique in JEDEC in that although most of the time we buy JEDEC standard device, we often buy it at a higher frequency than what JEDEC actually defines. So, an example is when the PC-100 was coming out, DRAM

vendors were offering 150 and 153 megahertz kind of die bins. We were actually buying say 166 megahertz. The 166 megahertz standard wasn't in JEDEC as something that they were working on at the time, but it was a speed bit of the memory that we could buy from our point-to-point application.

Q. When you use the term bin twice in your answer, what does that mean in your answer?

A. So, when the memory is built, there's a natural distribution that happens, some parts are faster than others. They will bin out these faster parts, cull them at faster speed grades and sell them for typically for more money. So, the faster are typically at a premium versus the slower parts.

Q. Prior to being the manager at nVidia, what was your position?

A. I was manager of systems engineering.

Q. And what did you do in that position?

A. I also worked with the memory vendors. I spent the majority of my time, however, on the board design and chip bring-up and debug kinds of activities. I had less time to spend on memory in that role.

Q. And did you have any positions within nVidia prior to being manager of systems engineering?

A. Prior to that I was a group member just to the

systems engineering team and prior to that I was a technical marketing engineer. During all these times I worked with memory industries and worked on more bring-up and chip debug activities.

Q. What do you mean by chip debug?

A. When the silicon comes back from fab it's not useful to put down on the graphics board. The teams that I have been involved with do the board design and then when the chip arrives we go through the process of getting the board working into something that you can actually sell and make a product out of.

Q. Just so I understand when you say the word "board," what are you referring to?

A. The PCB, the graphics board that the chip sits on.

Q. Is that the video card?

A. The video card.

Q. And you said the board bring-up, what does that mean?

A. I said what?

Q. I think you used the term board bring-up?

A. Board bring-up, yeah. The process that we go through in the lab to get something up and running in a PC is what we refer to as bring-up, it's the development process.

Q. Now, what DRAM manufacturers have you worked with while with nVidia?

A. It's changed a lot over the years. There's a lot less suppliers now than there used to be. Most of our time is spent with Samsung and Micron, Infineon, Hynix, a little bit with Elpida. There's ESMT, out of Taiwan, Nanya and Winbond. There are others.

Q. And Nanya is N A N Y A and Winbond is W I N B O N D?

A. Yes.

Q. And what work do you do with these DRAM manufacturers?

A. We do a variety of things, we talk to them about roadmaps, when we're going to have new products coming out, we get a sense of when their products are going to be coming out and we go through and have meetings with their engineers about future sets that we want, and ultimately work to define what a particular memory is going to have in the future, and then work towards getting that into the market.

Q. Does nVidia use internal code names to identify its individual graphics chips?

A. Yeah. The graphics products are all NV something and the something is usually a number, typically one or two digits.

Q. And which nVidia graphics chips have you worked on?

A. Since I've been there, there have been a lot. NV-3 was the first product I joined with, there's NV-4, NV-5, NV-10, NV-11, NV-15, NV-17, NV-20, NV-25, NV-28, NV-30, now NV-31, 34, NV-35.

Q. Now, are these products products that you work on one and you stop and then you work on the other and you stop?

A. Some -- most of the time it's just one of these devices in the lab. As we've become a bigger company now it's not uncommon for it to be multiple in the lab at the same time. So, some of them do go through the lab.

Q. And what's the nature of the work that you have done with respect to these graphics chips?

A. On the memory side of it I'm involved in the definition choices of what frequencies and what kinds of memory technology we're going to pick for each product. Once it actually comes back and the silicon is ready, I'm involved a lot of times in the chip bring-up activity that's actually in the lab trying to get the memory interfaces to run faster.

Q. Now, does nVidia design or manufacture DRAM?

A. No.

Q. Does nVidia buy DRAM?

A. We buy some. We -- for the Quadro product line, which is our workstation graphics cards, we buy the memory for those cards since we're actually having the card manufactured. We also buy memory for prototypes for testing out our product and sampling boards with customers. Then other than that, the only memory we buy is typically on the high end of our product line, and in that case, we'll buy memory that's hard to get, essentially the highest speed memory in the market, bundle it with our graphics policy serve processors to sell that to the market.

Q. What are the graphics processors that nVidia is currently selling?

A. There's quite a few that are still being sold. The most recent products are NV-30, NV-35, NV-31 and NV-34. There's some older products, NV-28 and NV-18, NV-11. There may be a few others that are still shipping.

Q. And what DRAM do these graphics processors work with?

A. They all have at least support for standard SDR SDRAM, many of them have support for DDR, a few of them have support for DDR2.

Q. Do you have an understanding of whether the DRAM

used with nVidia's graphics processors are always JEDEC compliant?

A. The graphics processors are not necessarily compliant in a true JEDEC sense. He many times we're using the same device that JEDEC would have defined, but because it's a different speed bin, there isn't a JEDEC standard for it per se. Generally what we try to do is use devices that are compatible with whatever that JEDEC standard was, so that we can leverage off that inventory that's in the market.

Q. Does nVidia ever attempt to get DRAM manufacturers to design DRAMs specifically for nVidia?

A. Yes, we frequently work with them on the high end of our product line to try to get a faster memory in the market than something that's what JEDEC might be working on.

Q. And how would these DRAMs differ from the JEDEC standard DRAMs that are being produced at the time?

A. We try to make them as close as possible to whatever the next standard in JEDEC is going to be, or we predict it's going to be. The main differences relate to in our applications it's a point-to-point application, there's usually only one or two memory loads on the data bus. It's a little bit different than a dim and we can typically run at a higher speed. So,

the differences usually relate to something about being a point-to-point graphics system or needing to run at a higher frequency.

Q. And why is it important that you make a similar to JEDEC standard DRAMs?

A. As we've talked about in several different product segments before, the very high end segments that we can sell for a higher price we can afford sometimes the premiums associated with the high speed memory. The low end of the market, you can't afford that premium and most of the low end products use a very JEDEC compatible device. So, we intermix those kinds of technologies depending on the price points in the market and the volume requirements.

Q. Is there a current DRAM being produced, a specialized DRAM being produced for nVidia?

A. The most recent one that went into production was what's now being called a GDDR2 DRAM, it was first sold with our NV-30 product and it's in one of our workstations products right now.

Q. And how does that differ from DDR2 SDRAM?

A. It's pretty similar, the DDR2 SDRAM standard focuses its definition on clock speeds that were basically 200 megahertz or 266 megahertz. We were focused on a 500 megahertz memory standard. So, the

main things that are part of a DDR2 standard in JEDEC are essentially a burst length 4 functionality, on-die termination and some differential clock strobes for the data bus. Those three things are in this GDDR2 part, which kind of separate from the DDR parts that are out there.

Q. I would like you to look at what's been marked for identification as CX-2832. Your Honor, this is another one.

JUDGE McGUIRE: You may approach.

BY MR. DAVIS:

Q. Do you recognize this document?

A. Yes.

Q. And what is CX-2832?

A. It's a product guide, it's for our Quadro Fx product. This marketing mainly is what was used for the workstation version of our NV-30 product line.

Q. And how long has nVidia been selling the Quadro Fx?

A. Since early this year.

Q. Do you know when work began on that nVidia Quadro Fx?

A. I don't know exactly when it began. The NV-30 product line was a major new architecture for us. It's fairly typical for a new architecture to take on the

order of two years to produce. So, it's somewhere around two years ago probably.

Q. Did nVidia have an internal code name for the Quadro Fx while it was being developed?

A. Yes, that was NV-30.

Q. And did you work on that product?

A. Yes.

Q. And what did you do on that project?

A. I was involved in the memory definition working with the memory suppliers to get the 500 megahertz memory defined, and ultimately when it came back in silicon form, I was in the lab working on it, getting it up to speed.

Q. Could you turn to the second page of CX-2832.

A. Yes.

Q. There's a table on that page, the lower portion of the page. Do you see that?

A. Yes.

Q. Can you describe what that table is.

A. It shows several graphics cards that are part of our workstation product family.

Q. Could I focus your attention on the row involving the Quadro Fx, I believe that's the top row.

A. Um-hmm.

Q. The fourth bullet down, and also, I'm sorry, the

second column of that table, the fourth bullet down states "128 megabytes DDR2 SDRAM." Could you describe what that means in this table?

A. So, that column is architecture, it talks a little bit about the basic architecture of the graphics board. The 128 megabytes DDR2 SDRAM is the memory that is on the graphics card, it's actually since this was produced, the name has effectively changed to GDDR2 SGRAM.

MR. DAVIS: Your Honor, I move CX-2832 into evidence.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2832 was admitted into evidence.)

BY MR. DAVIS:

Q. Earlier you described some of your graphics chips as being able to operate with both SDRAM and DDR SDRAM, does nVidia produce processor chips that are capable of operating with both types of DRAMs?

A. That was our NV-10 product line.

Q. And when did the NV-10 -- I'm sorry, when did nVidia video cards using DDR SDRAM begin to ship?

A. I'm sorry, could you say the question again?

Q. I'll restate it. When did the nVidia video

cards using DDR SDRAM begin to ship?

A. In the fall of 1999.

Q. And when did nVidia begin working on a graphics processors that was used to support those cards?

A. NV-10 was also a major new architecture, so roughly two years earlier is a reasonable estimate.

Q. Is this amount of time to ship a graphics card typical in your experience at nVidia?

A. It's typical for a major new architecture. It's typically not that long for minor changes to architecture for each generation of a product family.

Q. And how long does a minor change usually take?

A. Very minor things that are maybe almost kind of like bug fixes can be a few months. A more typical product is probably like a year, somewhere around that range.

Q. Do you know why the NV-10 was designed to support both SDRAM and DDR SDRAM?

A. Because DDR was new in that time, we had concerns whether or not it would be available in time for when we wanted to ship. So, the decision was made to support both SDR SDRAM and DDR SDRAM so that if DDR didn't show up, we still had a fallback plan and could still ship our product on the market, but at a slower speed.

Q. Were you involved in that decision?

A. Yes.

Q. Now, were there any other DRAMs considered for use with the NV-10 chip?

A. There are a number of technologies that were considered. In the end, it basically came down to DDR or the direct RDRAM. We also looked at I think Mitsubishi at the time had a -- something they called a 3DRAM and Fujitsu I think was doing an FC RAM back then.

Q. And why was DDR chosen instead of these other DRAMs?

A. The Mitsubishi and Fujitsu technologies were basically sole sourced kind of technologies. They were viewed as being kind of risky for us to produce, or risky to go with because we didn't know if there would be a wide supply base if there were manufacturing issues.

Q. I'm sorry, did you mention RDRAM as well?

A. RDRAM was -- I guess our decision came down to basically do we do direct RDRAM or do we do DDR. We came close to using direct RDRAM, but ultimately decided the alignment of our schedule in the market, it looked very risky to go with RDRAM. We wanted to come out in the fall of '99, at that same time period is when Intel was supposed to launch their chipset that was going to

use RDRAM, and we had a lot of concerns that supply was going to be limited, and if there was a supply problem, we probably weren't going to be the company to be able to get the product or we would have to pay a substantial premium to get the product. So, we chose to go a more evolutionary approach and go with DDR and maintain the backward capability of SDR in case there was a supply problem.

Q. And what was your understanding of what the effect would have been had there been a supply problem?

A. Well, for us it would have -- if we had chosen a direct RDRAM, we would have had no fallback plan, unless we had -- and technically Rambus also had a previous generation of technology, we could have designed support for that. It wasn't particularly compelling to our application, so it wasn't a good fallback plan, it was also a very high premium fallback plan.

The way things kind of worked out for us was although we chose direct RDRAM, ultimately to even -- we ultimately chose DDR SDRAM. We actually shipped first with SDR, we did end up having trouble getting the DDR as well.

Q. Do you know what other nVidia graphics chips are capable of operating with both SDRAM and DDR SDRAM?

A. Everything that we've built since NV-10 has had

support for both.

Q. And why are the more recent nVidia graphics chips capable of operating with both SDRAM and DDR SDRAM?

A. Because we serve a lot of different markets with these products, the cost associated with DDR versus SDR throughout much of the life of our -- of these products, it was still cheaper in many cases to buy an SDR SDRAM, so for the low end of the market that didn't really care about the performance, car manufacturers would buy and use the SDR SDRAM, and only recently has the cost of DDR essentially crossed over and it's now typically cheaper or more cost effective or a better performance trade-off to build with the DDR SDRAM than SDR, so almost all products now are shipping with DDR SDRAM.

Q. I would like to show you what's been marked for identification as CX-2828. You should have that in front of you.

A. Okay.

Q. Do you recognize this document?

A. Yes.

Q. Now, there's handwriting on the front page of the document, do you recognize that handwriting?

A. Yes. That's a highlighter.

Q. Could you tell me what this document is?

A. This first page looks like a photocopy of a file folder that I would have used to keep track of in this case next generation memory development efforts that we were working on.

Q. And what is the document itself?

A. The document is sort of a combination of an email thread and a report of the meeting minutes that we were having with Infineon to discuss a 500 megahertz DDR style device. Ultimately, this was the product that became the GDR2 SDRAM, or the SDRAM that we shipped with NV-30.

Q. And were you involved in these discussions?

A. Yes.

Q. What was your role in these discussions?

A. For nVidia's side, I was the primary technical contributor to this effort to get the 500 megahertz spec defined within Infineon. There were a number of other members of our team that depending on the meeting may or may not have been present.

Q. Now, were there any differences between this DRAM and the existing DDR SDRAM?

A. I'm sorry, say that again.

Q. Were there any technical differences between this DRAM and the existing DDR SDRAM?

A. So, this DRAM was targeting 500 megahertz. The

DDR standards that we were using is for other graphics memory solutions, at the time I don't recall exactly how fast they were, but up until recently, we had only gotten to about 325 megahertz. We just launched a product that's gotten up to 450 with a DDR kind of technology. This was targeting 500 megahertz and at the time was roughly twice the frequency that the JEDEC kinds of DDR devices were.

Q. Were there any other differences between this DRAM and the JEDEC SDRAM?

A. The feature sets that we settled on ended up being more in line with what was being developed for DDR2, we went with the burst 4 architecture, we wanted on-die termination and we wanted the differential clock strobes.

Q. If you turn to page 7, which is CX-2828-7.

A. Okay.

Q. And I refer you to the top entitled Technical Features. Do you see that?

A. Yes.

Q. Do you see the second paragraph in this section where the document states, "Infineon proposed to take into the definition some of the JEDEC DDR2 features."

A. Yes.

Q. And I think you just mentioned some of the DDR2

features that they were planning to take in. Do you know if there was work going on at JEDEC relating to the DDR2 standard at the time?

A. Yes, JEDEC was working on DDR2, and the frequencies that we were after, we agreed with Infineon could take advantage of the same basic technologies that the JEDEC was working on for their next generation. We wanted to be compatible in general with what JEDEC was doing, so it was convenient that things aligned and we were able to go in the same direction.

Q. If you focus on the topic at the top General Targets, which is right above that.

A. Yes.

Q. Under the topic General Targets is a two-sentence passage that I'll read out. "For nVidia, high-speed memories will not only be used in graphics products, but also in other applications like e.g. in sealed boxes as UMA memory. Accordingly, nVidia is focused on a new high speed technology and wants to drive new standard for high speed memory, which may differ partially from the future main memory."

Do you see that?

A. Yes.

Q. What does the term UMA mean?

A. That stands for unified memory architecture.

Q. And what is that?

A. For example, in our Enforce product line which is our chipset business, we integrate the graphics core into the north bridge product, in that case, we have a unified memory architecture. We use the same memory subsystem as the processor in that case, which in the platform it's in the DIMMs from JEDEC rather than have a separate memory, say either on a motherboard or on a graphics card, if we do the stand-alone cards.

Q. And what is your understanding of the term "main memory?"

A. Main memory is the -- whatever the JEDEC standard is at any given time being used by the PC market.

Q. Was it your understanding of nVidia's strategy here that nVidia was trying to drive the new standard for high speed memory?

A. Yes.

Q. And why was it necessary for nVidia to drive that type of a standard?

A. In the graphics space, one of the big differentiators for everybody's product is how fast is the memory subsystem, how much memory bandwidth do we have available to us. So, we put a lot of effort into making the memory subsystem faster, and as a company if

we can get access to faster memories ahead of our competitors, it gives us an advantage in the market. So, we work with the memory companies to try to ensure that we're defining what that standard is and that we're first to market with it.

Q. In the next paragraph under general -- I'm sorry, in that same paragraph, under general targets, is the statement that "JEDEC is a nice-to-have, but more important is to have a second source which provides compatible products." And I believe I misstated, it's in the next paragraph. Do you see that?

A. Yes.

Q. Does this passage accurately state your understanding of nVidia's interests in the importance of JEDEC for this particular DRAM chip?

A. For this particular DRAM chip, yes. We participate in JEDEC to try to drive our requirements in, but in the graphics space, for the highest performance devices, we know JEDEC is not going to be the one defining that requirement, they focus on the DIMMs at a much lower speed, therefore if we can get it out of JEDEC, that would be great, but the reality is we have to do it ourself.

MR. DAVIS: I would like to move CX-2828 into evidence.

JUDGE McGUIRE: Objection?

MR. STONE: No, I'm sorry, Your Honor, no objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2828 was admitted into evidence.)

JUDGE McGUIRE: It's getting late in the afternoon, Mr. Stone.

MR. STONE: I'm just trying to be easy.

JUDGE McGUIRE: That's fine.

BY MR. DAVIS:

Q. I'm going to show you what has been marked for identification as CX-2828.

Your Honor, this is a document that there is some information on it that was apparently confidential information, but it hadn't been marked in camera previously. I have spoken with Mr. Detra and we just blocked out the -- it's really just a number.

JUDGE McGUIRE: All right, as long as it's been --

MR. STONE: Yeah, I don't think it's 2828, though.

JUDGE McGUIRE: All right, as long as it's been -- as the issue has been addressed and there's no concern about it being disclosed proper improperly.

MR. STONE: No, they redacted the information. It was a third party's concern, I have no problem with that.

JUDGE McGUIRE: All right, proceed.

MADAM REPORTER: What was that CX number?

MR. DAVIS: CX-2829.

THE WITNESS: I have the document.

BY MR. DAVIS:

Q. Do you recognize this document?

A. Yes.

Q. And what is this document?

A. It's an email thread, it started from one of our architects asking some questions about the burst length that we expected to see for future memories and whether or not we expected any changes coming soon.

Q. If you turn to the third page of this document, the bottom email I suppose is the email you were just referring to?

A. Yes, it started with Matthew Pappakippos.

Q. And do you have an understanding of why Mr. Robinson forwarded that question to you?

A. John Robinson is one of the engineering managers for one of our chips, he's responsible for NV-35. He knows that I'm the primary person in the company that works with the memory vendors about the technology

direction. So, if he's got any questions, he forwards them on to me.

Q. That Robinson email on that page refers to a GS1000 and a Micron part. Do you have any idea what those terms refer to?

A. The GF-1000 I believe is Samsung's code name for the GDR DDR2 DRAM that we sold with NV-30.

Q. And the Micron part?

A. The Micron part is the product that is not yet on the market.

Q. Now, what was nVidia's interest in these two DRAM chips, the Micron chip and the GF-1000?

A. Those two products were the leading candidates for high-speed memory in a time frame that we wanted to bring out NV-30. They're the two technologies we were watching very closely.

Q. And if you could go to the first page of CX-2829, in the bottom email, which is an email from Matthew Pappakippos. Do you see that?

A. Yes.

Q. There's a sentence there that reads, "These changes are hard for us to work around in arch, and our lead time to fix them is 1.8 years." Do you see that?

A. Yes.

Q. Do you have an understanding of what the term

"arch" means there?

A. Arch is short for architecture.

Q. And what changes is he referring to?

A. This discussion thread was about the burst length. Our products at the time were being architected for a burst length 4, and if the next high speed device that we wanted was going to be a burst length 8, we needed about 1.8 years to change the architecture design to be efficient for burst length of 8 implementation.

MR. DAVIS: Your Honor, I would like to move CX-2829 into evidence.

MR. STONE: No objection.

JUDGE McGUIRE: Entered.

(CX Exhibit Number 2829 was admitted into evidence.)

BY MR. DAVIS:

Q. I think earlier you stated that nVidia is a member of JEDEC?

A. Correct.

Q. And why did nVidia join JEDEC?

A. For a variety of reasons, but in part to keep up with what technology was coming into the main side of our business with the chipsets, and also because we were requesting things from memory vendors that was different from the direction that JEDEC was going, and they had

asked us to get involved in JEDEC to drive our requirements. There weren't many members in JEDEC that were from the graphics industry, so most of the feedback that the JEDEC team was driving towards was based on system memory, and with what the server applications wanted.

Q. And who is nVidia's JEDEC representative?

A. I am.

Q. How many meetings have you attended since nVidia joined JEDEC?

A. On the order of a dozen.

Q. And --

JUDGE McGUIRE: All right, let's get some time frames here. It's not clear to me the time that he has joined JEDEC, so that I understand the context of the testimony. Can you tell me when you joined?

THE WITNESS: Sure. We became very active in JEDEC around 2001. Prior to that, it was -- we would go to JEDEC meetings when memory vendors would kind of give us an alert that it looks like something is going to get voted on that was important to us. So, we would go then.

JUDGE McGUIRE: Prior to that meaning in the preceding, what, two years or so?

THE WITNESS: Yeah. My belief is somewhere

around '98, '99 we probably joined or at least got involved at some of them.

JUDGE McGUIRE: All right, go ahead, Mr. Davis.

BY MR. DAVIS:

Q. And now, which committees of JEDEC have you attended?

A. Typically I go to JC-42, sometimes JC-16, they often have meetings that kind of overlap with each other in certain locations.

Q. And what is JC-16?

A. JC-16 focuses on the I/O technology, standards like SSTL and LLVTL come out of JEDEC's JC-16 committee. It's typically the I/O subsystem definition for whatever JEDEC 42.3 memory definition is agreed on.

Q. So when you say I/O subsystem definition, could you give me an explanation of that. I'm sorry, let me finish the question, could you give us a lower level definition of that?

A. Sure. So, the I/O interface is basically from the memory component pin out to the graphics board and then back to the memory controller there's a signaling protocol that is used or some sort of voltage levels that are transmitted across that wire. Whatever technology is used for that transmission, that is generally referred to as the I/O technology.

Q. Are you aware of what the term "programmable CAS latency" is?

A. Yes.

Q. Could you define that?

A. The CAS latency function is basically when you do a read to a memory, some number of clocks later, in the case of SDRAM, you would get that memory out of the data bus of -- you get the data out of the memory. So, a CAS latency of 3 as an example, three clocks after you do your read request the data comes out and that's CAS latency 3. Being programmable indicates of SDRAM you write a mode register and you pick whether you want it to be CAS latency 2 or CAS latency 3 or something else and then the DRAM operates in that mode.

Q. Are you aware of what -- I'm sorry, do you know whether the current JEDEC SDRAM/DDR SDRAM standards specify programmable CAS latency?

A. Yes, they do.

Q. Are you aware of what the term "programmable burst length" is?

A. Yes.

Q. Could you define that?

A. Again, being programmable that there fields in the mode register that represent the burst length, the programmable nature of it is, you know, typically a

burst length 2, burst length 4, burst length 8 are common values. Burst length represents how many bits of data come out of the memory every time you make a request. In the case of a read or a write you are going to send them that many bits of data.

Q. Do you know what the current SDRAM and DDR SDRAM standards specify programmable burst length?

A. Yes, they do.

Q. Do you understand what the term "dual edge clocking" means with respect to JEDEC compliant DRAM?

A. Yes.

Q. Could you define that for me?

A. The dual edge clock refers to the data interface, so the data bus for every time you set a clock there's a data strobe being used where every edge of the data strobe clocks out the data. So, a burst of 4 you have two clocks, there's four edges to that clock, those four edges, the data come out.

Q. Are you aware of Rambus's lawsuits against the DRAM manufacturers?

A. Yes.

Q. And when did you learn of those lawsuits?

A. When it became very public in the Internet forums for technology sites that DRAM companies were being sued, I was aware of it.

Q. And did you have an understanding of what those lawsuits are about?

A. I understand it relates to the burst length/CAS latency kinds of patents.

Q. At the time you learned of the Rambus lawsuits against the DRAM manufacturers, do you recall whether nVidia had any graphics processors then in the market that were designed to work with DRAMs that used JEDEC standard programmable CAS latency?

A. All of the devices that we have on market use either SDR or DDR. I believe that the time DDR wasn't quite available in the market yet, so we were using SDR, which does have programmable functionality for burst lengths and CAS latency.

Q. At the time you learned of the Rambus lawsuits against the DRAM manufacturers, do you recall whether any of the nVidia graphics processors then on the market were designed to work with a DRAM that used JEDEC standard dual edge clocking?

A. At some point we began using DDR with maybe ten. We ended up shipping in the fall of '99, so somewhere around the year of '99 we were working on that memory controller and so we would have had the dual edge clock in that device.

Q. Now, between the time you learned of the Rambus

lawsuits against the DRAM manufacturers and today, did you have any understanding of whether any of the features of SDRAM and DDR SDRAM that you've described, that's programmable CAS latency, programmable burst length and dual edge clock, as those technologies are specified by JEDEC, could have been replaced in the JEDEC standard by alternatives?

MR. STONE: Objection, Your Honor, lacks foundation, calls for opinion testimony or speculation that this witness has not been designated as an expert.

MR. DAVIS: I am not asking for an opinion, I am asking for his state of mind as a JEDEC representative.

JUDGE McGUIRE: Still lack of foundation. So, you can restate. Sustained.

BY MR. DAVIS:

Q. Well, did you have any understanding about whether there were alternative technologies for these -- for programmable CAS latency, programmable burst length and dual edge clocking?

MR. STONE: Your Honor, that's the same question that we just had, and whether he had an understanding or not, it's calling for opinion testimony. This is not his business. He doesn't design DRAMs. He doesn't design these features in or out of them. Whether there's alternatives or not is proper testimony for

expert testimony or at least a designer.

MR. DAVIS: He designs and works with the graphics processor which has the memory controller on it, his job is to work with DRAM and make it work with that memory controller.

JUDGE McGUIRE: It's a close call. I'm going to hear the question. Overruled.

THE WITNESS: Yeah, I believe there are solutions to those issues. You have to treat them on a case-by-case basis, to which what solution is right for each of those options.

BY MR. DAVIS:

Q. And what have you understood to be the alternatives to JEDEC standard programmable CAS latency in the mode register?

MR. STONE: Your Honor, same objection, if I can just have a standing objection. There's no foundation that he knows the alternatives, he is not an expert.

JUDGE McGUIRE: Okay, I said it was a close call. So, that doesn't mean that you are going to get it in on the same line of questioning on this, Mr. Davis. Now, I allowed you, you know, the one episode to allow that question, I'm not going to allow this same inquiry when I don't think there's a whole lot of foundation here that's going to give him the opportunity

to state what he's stated. So, you either have to change your inquiry or I'm going to uphold every objection on this issue. So --

BY MR. DAVIS:

Q. Okay, Your Honor.

Did you ever propose to change the JEDEC SDRAM standard at any time after you learned of the Rambus lawsuits against the DRAM manufacturers to eliminate programmable CAS latency or burst length?

A. No, I didn't propose for them to go off and change the technology. It's painful for us as a company to go and change our direction.

MR. STONE: Objection, Your Honor, the question was did you do something, he said, no, I didn't do it, and that should be the end of his answer and we should now have another question so that I have an opportunity to object. So, I move to strike everything after the words "and change the technology -- no, I didn't propose for them to go off and change the technology." I move to strike everything beyond that as beyond the grounds of the question.

JUDGE McGUIRE: Well, you didn't ask the question and you don't have any standing on that ground to propose that objection.

MR. STONE: Well, Your Honor, if I didn't have

standing to raise that objection, then you could put a witness on the stand and say, what's your name, and the witness could talk for four hours, even if he was no longer being responsive, if I don't have the right to object when he stops being responsive. I think I do have a right to object when he goes beyond the scope of the question. If not --

JUDGE McGUIRE: Well, beyond the scope is one thing, on that grounds, you know, I will sustain that objection as beyond the scope. But, you know, unresponsive, that's something else all together.

MR. STONE: Fine, Your Honor, I don't mean to phrase it incorrectly, I appreciate that.

JUDGE McGUIRE: We'll sustain in terms of scope.

BY MR. DAVIS:

Q. And why didn't you propose to change the SDRAM or DDR SDRAM standards to eliminate programmable CAS latency and programmable burst length?

A. We were trying to launch products into the market and if the standard was going to change, that meant we had to change our development plan and go change to something new that was yet undefined. For us it's a painful process to go through and not be able to release a product that's basically ready to be released and have to go start over again. So, we weren't trying

to get the industry to promote a change when we were working within the standard in the industry that was proposed by JEDEC and was perfectly viable.

Q. Did you ever propose to change the JEDEC DDR SDRAM standard at any time after you learned about the Rambus lawsuits against the DRAM manufacturers?

A. No, I did not.

Q. And why didn't you?

A. For the same reason, that we felt that it was a lot of work for the industry to have to go through and make those changes. As a graphics company, we try to create demand for our products primarily through the game industry. So, we have feature sets that are new to each architecture generation that enable a very new set of type of game technologies to come out. The time lines of those games get aligned to the time lines of our chips. If we can't release the chip because we have to go redesign for some new technology, then, you know, we miss the opportunity to align with this new game or whatever is coming out. So, we try not to change things late in the design cycles.

MR. DAVIS: Thank you. No more questions, Your Honor.

JUDGE McGUIRE: Thank you, Mr. Davis. You may cross examine.

MR. STONE: Thank you, Your Honor.

CROSS EXAMINATION

BY MR. STONE:

Q. Mr. Wagner, how are you?

A. Good.

Q. Are you a gamer?

A. I'm not.

Q. Cool games featured in your catalog 2833, aren't there?

A. There's a lot of neat stuff coming out.

Q. Are the gamers sort of the high-end group concerning what they're willing to pay for performance?

A. Yeah, the real enthusiast gamers, they'll save up their money and go out and buy the latest and greatest things every few months if they can afford to.

Q. And so one of the things that you find is true is that the highest speed, the fastest chip will sell for more money?

A. Yes.

Q. And in order to achieve that one of the things you look for is memory that runs at the fastest speed?

A. Yes.

Q. And to do that, you've pushed manufacturers to supply you with memory which goes faster than anything JEDEC has specified?

A. Yes.

Q. And you've pushed Infineon to do that, you've pushed Micron to do that, you've pushed Hynix to do that, you've pushed other companies as well, right?

A. I wouldn't say I push. We tell them what we want and if they want to serve our market, they go off and produce high speed memories.

Q. You encourage them to do it?

A. Absolutely, and our selling factor on it quite frankly is that they get to be the leading edge performance guy, by serving our market they can go out and say they're the fastest in the world as well.

Q. And often times what they provide you is something that is not JEDEC -- is not within the JEDEC specifications?

A. Correct.

Q. Okay. And you talked to us earlier today about DDR2.

A. Yes.

Q. The DDR2 specification has not been published by JEDEC, has it?

A. I actually think there is a version of it that's been published now.

Q. You think it's been published as opposed to just circulated within JEDEC?

A. I believe it is on the JEDEC website now, but I'm not positive.

Q. And you started, in any event, you started making use of something that you thought would meet the JEDEC specification before that specification was finalized?

A. Correct.

Q. And the first product, am I right that the first product you started to work on that made use of what you thought would ultimately be a DDR2 product was the one featured in 2832, the nVidia Quadro Fx?

A. Not exactly. It wasn't a -- it was the first device that we targeted to be compatible with hopefully what would become a JEDEC DDR2 standard. The device that we were developing was not being developed in the context of JEDEC, so we knew that it was not true JEDEC DDR2, we hoped they would be compatible so that we could ultimately have a lower cost version of it also further down the product line.

Q. Let me try and just make sure. Was the first -- let me just ask you, what was the first product that you designed that you thought would make use of something which might become DDR2?

A. That was our NV-30 product.

Q. That was your which one?

A. NV-30.

Q. And the NV-30, was that then released with a product name?

A. Yes, that bulletin that you are holding, Quadro Fx is one of the product names associated with the NV-30 product.

Q. And the NV-30 product, when did you start work on that?

A. I had said previously it was on the order of a two-year development cycle. I don't know exactly when it started.

Q. And when was it first released?

A. We released that product into the market early this year, maybe February time frame, something like that.

Q. So, you would have started on it some time in 2001? Late 2000, early 2001?

A. Yeah, something like that.

Q. And in any event, you started work on that product after you knew about the Rambus lawsuits, correct?

A. Yes.

Q. You started work on the product after you knew that the use of what might become DDR2 might infringe on the Rambus patents?

A. The development of that technology, there really wasn't any discussion whether it did or didn't infringe on patents. I don't know whether I knew much about the patents at that point in time. I knew there was lawsuits going on, I don't think I knew the specifics.

Q. Well, I thought you -- maybe I misheard earlier, I thought you told us that the lawsuits involved programmable CAS latency.

A. I do understand that now, in 2000, I don't know if I knew specifically it was a CAS latency they were concerned about or something else.

Q. When did you first learn it involved dual edge clocking?

A. I don't know exactly when I learned the specifics of the patents in question, it was probably some time around getting asked to participate in this trial what were the specific items involved.

Q. So, when Mr. Davis asked you if you ever went to JEDEC and proposed that they changed the standard with respect to programmable CAS latency or programmable burst length or dual edge clock, when is the first time that you even knew those issues might be involved in the lawsuit?

A. At some point, I don't recall exactly what year, memory vendors came to us and said would you like to

change this, we're finding some IP issues over it, and our feedback was, well, if the industry is going to change, we need to change with it and we basically left it up to them to go off and decide to change or not.

Q. And when was that that the memory vendors came to you and said should we change to avoid these IP problems?

A. It was probably some time in the 2000 time frame. I don't recall exactly.

Q. Okay. And so, and you gave your feedback to the memory vendors when they asked you if they should change?

A. The way discussions would have happened, they wouldn't have asked me should I change, they would have, you know, brought in suggestions to change the technology and we would have said, we already have a standard, we don't really want to change, or we're on a development cycle that cannot tolerate the schedule hit.

Q. And do you remember what manufacturers came to you in 2000 and asked you what you thought about changing to avoid IP problems?

A. I don't know who all of them would have been, I believe Micron would have come, we probably would have talked to Samsung in that time frame as well, I don't know for sure.

Q. And do you recall that one of the changes proposed was to simply go to fixed burst length?

A. Yes, and fixed burst length we've never had an issue with. Any one of our products has only ever been fixed burst length.

Q. And they also proposed going to fixed CAS latency, didn't they?

A. Yeah. There were proposals about fixed CAS latency at the time.

Q. Now, with respect to the dual edge clocking, one of the things you've done in the lab is see how fast you could run a single clock, an SDRAM, correct?

A. Correct.

Q. And you've got SDRAMs running in the lab at between 300 and 350 megahertz. Is that right?

A. It depends on what you call the SDRAM, I don't know if you're talking about standard SDRAM or the DDR device. We run the DDR2 SGRAM device in the lab at a 500 megahertz clock, which is a gigabit data rate. To me that's an SDRAM still, but because the basic clocking structure that's going across the bus is 500 megahertz.

Q. And you're still -- and in those instances you're using just one edge of the clock?

A. The way the memory subsystem works, the address and command portion of the bus is at the 500 megahertz

clock, the data portion of the bus is at the gigabit frequency, so a thousand megahertz kind of a clock.

Q. Have you run tests to see how fast you could run an SDRAM, which was not using any dual edge clocking?

A. We have, and at the time of the standard SDRAMs, I think they stopped making them at about the 250 megahertz for a standard SDRAM.

Q. But you've pushed those past the 250 megahertz, right?

A. No, we didn't really have any need to spend the energy in that time frame. It was a -- the devices that could go to 250 megahertz were in a density that wasn't interesting to the new crop of devices we were coming out with, so we just didn't put the engineering effort into it.

Q. And do you put your memory in a module or does it fit directly next to the memory -- I'm sorry, the -- what's it called, the graphics controller?

A. It depends on the market segment and what you refer to as a module. We sell a graphics card, so the graphics card is not a module in most people's sense, but in many ways the bulk of a graphics card is a graphics chip and memory plus power supply components. So, that's kind of an example for a graphics card. We have in a mobile market, we have a device we call MAP

which is literally just a little tiny package that we kind of refer to as a little memory graphics module. It's for the mobile market, takes up a small amount of space. But we don't do DIMMs that are those kind of memory modules.

Q. Let me ask you just about the picture that you were shown earlier by Mr. Davis, which is in exhibit CX-2833, on the page entitled nVidia Consumer Graphics Solutions. I think that's one he showed you that you identified the square box in the center with the nVidia written on it as being the graphics chip.

A. Yes.

Q. And then over to the right-hand side you showed four memory chips.

A. Correct.

Q. And are those individual DRAMs?

A. Yes.

Q. Okay. Not in a dim in this configuration?

A. Correct.

Q. Okay.

A. Mind you, there are times in our market where there have been memory DIMMs that plug into the graphics card, also. It's not cost effective these days.

Q. Okay. The NV-10 that you told us about earlier that could use both the SDRAM and the DDR, because you

had designed the controller to use either one. Do you recall that?

A. Correct.

Q. And am I correct that that's the NV-10?

A. Correct.

Q. Is that what you would call a bimodal controller?

A. I don't generally use the term bimodal in my talk of it. It's a memory controller that's compatible for either SDR or DDR SDRAM.

Q. And does it detect what type of DRAM you've put in?

A. It doesn't detect, the BIOS we have on our card gets programmed to tell the chip what type of memory is on the card. Because we aren't using a dim, we don't have to detect, we know what we do manufacturing time, so we just program the BIOS to represent what is manufactured.

Q. And so there's different circuitry that comes into operation depending whether it's an SDRAM or DDR?

A. Yeah, there are some differences.

Q. And in your operations, the use of SDRAM and DDR are not something you would say are backward compatible, are they?

A. The use of which?

Q. SDRAM and DDR, the DDR is not backward compatible to your products which were designed for SDRAM?

A. I disagree with that, it is backward compatible.

Q. Backward compatible beginning with the NV-10?

A. So --

Q. Am I right, beginning with the NV-10?

A. As we design a memory controller, NV-10 is the first one we did both DDR and SDR, the memory controller is backward compatible in supporting SDR, as well as DDR. So, we developed for the new technology, we maintained backward capability to the old.

Q. Beginning with the NV-10?

A. Yes.

Q. And that's because you designed the controller specifically to do that?

A. Yes.

Q. And could you plug in a DDR2 product into the NV-10?

A. Not into the NV-10. Actually --

Q. You have to go up to -- go ahead.

A. Actually, NV-10 could have worked with the GDR-2 SGRAM. If I could elaborate on that, if you like, it is possible. The reason I know this is the first chip that we actually tested the GDR SGRAM was actually an NV-20

chip even though we shipped it with NV-30, actually NV-25. In order to support -- well, the way we specified the GDR-2 device is we knew it would come out before our chip that supported DDR2, so we built in capability into the memory to be -- to kind of act like a DDR-1 device in the sense that as long as the controller could operate in burst 4 mode, you could still operate the memory because you could turn off the offline termination and you could tie the differential strobe one way to be directional. So, because it was a fixed burst 4 DDR in essence, so NV-10 had support for burst 4 and probably could have used it.

Q. The number of different product launches you've worked on since you joined nVidia, if I counted them correctly, would be 14?

A. If you counted them, I haven't tried to count them. I can walk through them all in my mind again if you like.

Q. You did 3, 4, 5, 10, 11, 15, 17, 20, 25, 28, 30, 31 and 34 and 35?

A. There was an 18 in there also.

Q. And I thought earlier when you gave the list, I thought you left 18 out of the list. You worked on that one as well, right?

A. I didn't do as much on that one, but yes.

Q. And you've been with the company how long?
A. About six years.
Q. Since you've -- and your first JEDEC meeting that you attended was March of 2000, correct?
A. I don't believe that's the first, but that's at least kind of in that time frame.
Q. Well, I could be wrong. Because I certainly wasn't there, but let me show you your deposition transcript for one second, if I may.
May I, Your Honor?
JUDGE McGUIRE: Yes.
BY MR. STONE:
Q. Mr. Wagner, you recall giving a deposition in this case?
A. Yes.
Q. And you were deposed I believe it was January of this year. Is that right?
A. Sounds correct. The document says that, so yes.
Q. I hand you the transcription of that deposition, I would ask you if you could to turn to page 44 of this transcript and look at the testimony that starts at the bottom of page 44, line 24, and continue on to page 45, line 9.
A. I'm sorry, how far down?
Q. Just line 9. Did you have a chance to read?

A. Yes, I did.

Q. I'm just trying to jog your recollection. If you remember a meeting before March 4 of 2000, could you tell us when it was and where it was?

A. Like I said, the very first meeting that I ever went was to when I was invited to attend, at the time we weren't members and we were talking about a data mass function. I know that's before Atlanta. I don't know exactly when that was. It was in the Bay Area, though, so it was in California somewhere.

Q. And that was the first one that you can remember attending as a member, March 2000 in Atlanta?

A. I think it is.

MR. STONE: That's all. Thank you, no further questions.

JUDGE McGUIRE: Mr. Davis, any further redirect?

MR. DAVIS: No, Your Honor.

JUDGE McGUIRE: Then thank you, sir, for your testimony. You are excused from this proceeding. Let me ask complaint counsel again, get me up to date for what's on tap again for Monday.

MR. OLIVER: Actually, if I could have just a moment, please, Your Honor.

JUDGE McGUIRE: Okay.

MR. OLIVER: Your Honor, we were expecting to

call Mr. Steve Polzin of AMD, but he will not be available until 1:00 in the afternoon. We had actually been hoping to resolve the deposition issue with Mr. Reese Brown, but --

JUDGE McGUIRE: I can do that on Monday. I can do that Monday morning if you want.

MR. OLIVER: We have -- we have filed a response to --

JUDGE McGUIRE: Right, I just got it today and I haven't really had a chance to go through it, but if we wanted to go -- I mean, it's only -- your opposition only pertains to pages, what, 89 through 115, something like that?

MR. STONE: Yeah, I don't have the exact page numbers in mind but it's a limited number of pages that we had objected to.

JUDGE McGUIRE: That's what I'm asking. Is it complaint counsel's intent to go ahead and put that on and then at the time we get to the area that's in dispute then I'll rule? I mean is that what you plan on doing?

MR. OLIVER: With three caveats, Your Honor, first I haven't had a chance to consult with the other side to see if they are willing to agree to that. The second, I agree that there are some additional

objections that may have to be ruled on after we go that I think could be handled as we go. The third is that to my understanding that we still have a few technical issues to work out with the video, which I think we can do over the weekend, but I have not confirmed that.

JUDGE McGUIRE: That's why I'm asking, because I didn't know what you had planned. So, on that grounds, then, you know, how do you want to proceed?

MR. OLIVER: Would it be possible to try to give you an update either later today or first thing Monday morning?

JUDGE McGUIRE: Sure. I mean, you know, I'm agreeable to whatever you have.

MR. STONE: That's fine. I mean, I think --

JUDGE McGUIRE: I'm not sure, I mean the question is whether I'm going to have time to consider, I mean, you say there's other concerns that you have in opposition to any respondent's objections?

MR. OLIVER: My understanding is there's certain specific objections to certain questions.

JUDGE McGUIRE: Pages 89 through 115 I believe is the crux of your objection.

MR. STONE: I believe that's right, Mr. Perry is handling that.

JUDGE McGUIRE: It's close. But you said

however there's other objections.

MR. OLIVER: That was my understanding, again, Mr. Swindell has been coming up on our behalf and he is not here. My understanding is that there were a few additional specific objections that would have to be ruled on.

JUDGE McGUIRE: Well, I'll tell you right now having just gone through your response today, I tend to agree with the respondent that there was no -- there's certainly an improper foundation, if any, you know, foundation laid for the proposed testimony on pages 89 through 115. So, that's how I intend to hold, you know, on that issue. So, I'll just advise you at this time of that. Now, if there's other objections, we could perhaps consider them as we go on Monday, if that's agreeable. But, you know, I don't know what you have in store at this point.

MR. STONE: That's fine with us, Your Honor, if you rule on them as we go. We have no problem with that. I don't think the other ones are large in volume and I think you can rule on them. That's fine with us.

JUDGE McGUIRE: Now, I've had time to go through each of the pleadings, and like I say, I don't see where there's a proper foundation, at least on the pages in question.

MR. OLIVER: Okay, thank you, Your Honor. I will try to find out as quickly as I can whether we can resolve the technical issues with the video over the weekend.

JUDGE McGUIRE: So then you want to plan on I guess convening on Monday in the morning at the same time, 9:30, or do you want to advise me, you know, on Monday morning and then we'll play it by ear on Monday morning?

MR. OLIVER: I would recommend, Your Honor, that we advise you first thing Monday morning and play it by ear.

MR. STONE: We probably need to have the reporter here, though.

JUDGE McGUIRE: Let's plan on having a court reporter here in the morning, early on Monday morning at 9:30. So, with the anticipation that we might well start at 9:30.

MR. STONE: And we'll work, if we can help resolve the technical issues, we will.

JUDGE McGUIRE: Okay. Anything else?

MR. STONE: Nothing else, Your Honor, thank you.

JUDGE McGUIRE: This hearing is now in recess. See you Monday.

(Whereupon, at 4:47 p.m., hearing adjourned.)

C E R T I F I C A T E O F R E P O R T E R

DOCKET/FILE NUMBER: 9302

CASE TITLE: RAMBUS, INC.

HEARING DATE: MAY 30, 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:

Sally Jo Bowling

C E R T I F I C A T E 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.

Sara J. Vance