

1 FEDERAL TRADE COMMISSION
2 I N D E X (PUBLIC RECORD)
3
4 WITNESS: DIRECT CROSS REDIRECT RECROSS
5 Polzin 3932 3997 4048
6
7 EXHIBITS FOR ID IN EVID
8 CX
9 Number 2153 3948
10 Number 2152 3987
11
12 RX
13 Number 2302 3931
14 Number 2100-13 4055
15
16 DX
17 Number 31 4048
18
19
20
21
22
23
24
25

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

UNITED STATES OF AMERICA
FEDERAL TRADE COMMISSION

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

Monday, June 2, 2003

1:00 p.m.

TRIAL VOLUME 20

PART 1

PUBLIC RECORD

BEFORE THE HONORABLE STEPHEN J. McGUIRE

Chief Administrative Law Judge

Federal Trade Commission

600 Pennsylvania Avenue, N.W.

Washington, D.C.

Reported by: Josett F. Hall, RMR-CRR

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

1 APPEARANCES:

2

3 ON BEHALF OF THE FEDERAL TRADE COMMISSION:

4 M. SEAN ROYALL, Attorney

5 GEOFFREY OLIVER, Attorney

6 JOHN C. WEBER, Attorney

7 MICHAEL FRANCHAK, Attorney

8 ROBERT DAVIS, Attorney

9 Federal Trade Commission

10 601 New Jersey Avenue, N.W.

11 Washington, D.C. 20580-0000

12 (202) 326-3663

13

14 ON BEHALF OF THE RESPONDENT:

15 GREGORY P. STONE, Attorney

16 STEVEN M. PERRY, Attorney

17 PETER A. DETRE, Attorney

18 SEAN GATES, Attorney

19 ANDREA WEISS JEFFRIES, Attorney

20 Munger, Tolles & Olson LLP

21 355 South Grand Avenue, 35th Floor

22 Los Angeles, California 90071-1560

23 (213) 683-9255

24

25

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

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

APPEARANCES:

ON BEHALF OF THE RESPONDENT:

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

P R O C E E D I N G S

- - - - -

1
2
3 JUDGE McGUIRE: This hearing is now in order.

4 Counsel, I want to take up a couple
5 housekeeping items before we start today.

6 First of all, regarding the opposition to the
7 proposed testimony of Dr. Oh, I just received the
8 response to that opposition about two hours ago. I am
9 currently involved in trying to draw up an order on
10 that and I anticipate being able to have that order
11 issued on Tuesday morning, so we'll anticipate having
12 that issued at that time.

13 The other item I want to take up is I've been
14 informed that I'm going to have to attend a
15 conference, I'm obligated to attend a conference on
16 Friday, June 13, and that is to run from 9:00 a.m.
17 through 12:00 noon, so I will be happy -- we can I
18 guess convene the hearing at any time on that date
19 after 12:00, or if it's going to cause too much
20 problem, we can go to trial as well on I guess this
21 coming Friday. I know we're otherwise off on our
22 schedule, but if that would help overcome some of
23 that, or we could start on that Friday at 12:30 and go
24 late if we have to.

25 So I'm just going to advise everyone of that

1 change in schedule.

2 Mr. Oliver, go ahead.

3 MR. OLIVER: Just on scheduling issues, I will
4 have to check our calendar to see how that affects us
5 and what we might be able to do. I did want to remind
6 you that we were planning to have next 3t8GrEy and
7 Wedn8GrEy off, and as a result of that we were planning
8 to be in schedule next FrirEy.

9 JUDGE McGUIRE: I'll tell you what this is.
10 This is not of my own choosing. As a supervisor of an
11 office here in this FTC agency I'm required to attend a
12 three-hour training session for my employees. I
13 understand that is being offered on other dates. I
14 will see if I can change to either what would be
15 3t8GrEy, the 10th --

16 MR. STONE: 10th and 11th we claim to be dark.

17 JUDGE McGUIRE: Let me see if I can't get that
18 changed. Otherwise, that was the only time they had.

1 with our schedule of witnesses to see.

2 JUDGE McGUIRE: I'll have an answer on that in
3 the morning.

4 MR. STONE: Okay. And so your plan is at the
5 moment at least we're scheduled to be dark this
6 Friday.

7 JUDGE McGUIRE: That was the plan, but I'm open
8 as well. I want to keep this thing on track, and I
9 know we've had a few off days here and -- but I thought
10 that was our understanding, but I'm free to do whatever
11 the parties can do, and I know that puts the onus on
12 the FTC in that regard.

13 MR. STONE: I think they might have been
14 planning to go forward on Friday.

15 JUDGE McGUIRE: Okay. Then let's go forward on
16 Friday.

17 MR. OLIVER: I should have been clear on our
18 schedule.

19 JUDGE McGUIRE: I didn't assume that. I
20 thought we wouldn't be going on Friday.

21 So we will be going this coming Friday, the
22 6th.

23 MR. OLIVER: If that's okay with you. We
24 thought that by having next Tuesday and Wednesday
25 off --

1 JUDGE McGUIRE: No. I didn't realize that, so
2 that's fine. So that would make it imperative then to
3 try to change the schedule if I could for the following
4 Friday if we're going to be off on Tuesday and
5 Wednesday.

6 MR. STONE: If you can, but if we lose a half a
7 day --

8 JUDGE McGUIRE: If we do, we do. I'll do what
9 I can do about it, but I'm not in charge, so we'll
10 see.

11 Mr. Stone, anything you want to add?

12 MR. STONE: I just want to introduce one of my
13 partners. Andrea Weiss Jeffries is going to be here.
14 Mr. Detre will be back. His wife is going to be due
15 any day now, so she's come to help us.

16 Then I want to move in RX-2302, which was the
17 press release that I used on Friday.

18 JUDGE McGUIRE: Mr. Oliver, any objection?

19 MR. OLIVER: No objection, Your Honor.

20 JUDGE McGUIRE: All right. So entered at this
21 time.

22 (RX Exhibit Number 2302 was admitted into
23 evidence.)

24 MR. STONE: Okay, Your Honor.

25 JUDGE McGUIRE: Okay. Very good.

1 are competitive components that surround the
2 microprocessors that AMD designs and sells. That
3 includes chipsets, motherboards, DRAM, voltage
4 controllers, clock chips, everything that surrounds the
5 microprocessor to make competitive computers for our
6 customers.

7 Q. How long have you been the chief platform
8 architect at AMD?

9 A. I think they gave me that title in 2001.

10 Q. And what was your position prior to that?

11 A. Before that, they referred to me as the Athlon
12 system architect where I was focusing on Athlon, our
13 previous-generation microprocessor.

14 Q. What were your responsibilities in that
15 position?

1 implementation of the Northbridge for K7, so I manage
2 the team, drive the schedule, drive the project.

3 Q. Were there any code names for that particular
4 chipset?

5 A. Yeah. The K7 or Athlon Northbridge chipset was
6 referred to internally as Irongate or IG.

7 Q. Did you have a position prior to that at AMD?

8 A. Before that, I managed pretty much the same
9 team working on a K6 Northbridge chipset.

10 Q. And how long have you been at AMD?

11 A. June 1996 is when I joined.

12 Q. And prior to that where were you?

13 A. I was at Apple Computer.

14 Q. And what did you do there?

15 A. At Apple Computer I was the -- they called me
16 the hardware lead or the system architect for power
17 Macintosh computers.

18 Q. And what did you do in that position?

19 A. Pretty much similar to what I described as the
20 system architect at AMD but more detailed. I managed
21 multiple designers working on chipsets in the memory
22 controller, the I/O controller. Details of the
23 motherboard design I worked on, system-level design
24 aspects, file cards in mechanicals, that sort of
25 thing.

1 Q. Did you have any responsibilities at Apple with
2 respect to DRAM?

3 A. Yes. My team was the team that designed the
4 memory controller chip.

5 Q. Did Apple use SDRAM while you were at Apple?

6 A. Not to my knowledge. Not in the product I was
7 working on.

8 Q. Were you aware of whether Apple was considering
9 using SDRAM while you were employed there?

10 A. We had considered that. Our procurement
11 manager, the folks that are in charge of buying
12 components for our -- for Apple products, kept us aware
13 of industry developments, and synchronous DRAM was on
14 the horizon. We chose not to implement it for the
15 machine that we introduced in 1995.

16 Q. And why did you choose not to implement it?

17 A. I'm sorry? Why?

18 Q. Why did you choose not to implement it?

19 A. The main factor was it was too new. It wasn't
20 in the price point that was established to be a
21 commodity. We didn't want to be disadvantaged when
22 customers bought a Macintosh that they would have to
23 buy more expensive memory than buying a PC, for

y ttry? lhe

1 synchronous DRAM. The commodity DRAM was fast page
 2 mode or EDO. I forget exactly which one was the
 3 commodity at the time, but it wasn't synchronous.
 4 Synchronous was too new.

5 Q. You've used the term "commodity." How have you
 6 been using that term?

7 A. "Commodity" means widely available at a low
 8 cost and7 T7lgeraly aeans wmultipe amanufctlurer wmak
 39 3he tsamepagr and7 he ycommpet orncost

11 A. "1992too"1996, almst afie yoears.

13 A. "Iwas temployd tt the tDigital Equipment

1 SCw(w cticosin Massachd ttt 1 SCw(w 9yd reeAsg9d Y

1 A. I have a BSEE from the University of
2 Connecticut.

3 Q. That's a bachelor of science in electrical
4 engineering?

5 A. Yes, sir.

6 Q. Since you've been at AMD, how many different
7 CPU generations have you worked on?

8 A. K6, K7, K8, and we're working on K9.

9 Q. Does AMD sell chipsets for those CPUs?

10 A. AMD sold chipsets for K7 and K8. We haven't
11 sold K9 yet, so we haven't sold the chipsets for it
12 yet.

13 Q. AMD didn't sell chipsets for the K6?

14 A. I don't think it did. I'm not quite certain
15 on whether we actually -- we worked on a design for
16 one, but I believe we chose not to take it to market.

17 Q. And why was AMD developing chipsets to work
18 with the K6?

19 A. The K6 microprocessor was pin compatible with
20 an existing Intel microprocessor, the Pentium, and
21 therefore we were able to sell our microprocessors to
22 folks that could get a motherboard that was compatible
23 with a Pentium. There was an existing infrastructure
24 of motherboards for Pentium that K6 fit into very well,
25 seamlessly.

1 We, AMD, were fearful that Intel was moving
2 away from that Pentium infrastructure and they were
3 going to a Pentium II infrastructure and they were
4 going to leave that infrastructure behind and we
5 wouldn't have anywhere to plug in our microprocessor,
6 so we thought we needed to continue that motherboard
7 infrastructure with our own chipset.

8 Q. Did AMD sell chipsets for the K7?

9 A. Yes.

10 Q. And what was your understanding of the business
11 purpose for AMD to sell chipsets for the K7?

12 A. With K7, we developed a unique interface to our
13 microprocessor, the so-called front-side bus. We did
14 not copy Intel or follow Intel at that time. We
15 licensed a design from Digital Equipment Corporation
16 that was used on the Alpha microprocessor.

17 We called that the S2K, which stood for
18 Socket 2000, just a geeky term, but that was a unique
19 thing to AMD and we needed to design a chipset that
20 interfaced to it to get our motherboard partners
21 developing motherboards that would be compatible with
22 K7 and also provide the design as a design example for
23 our chipset partners so that they could design chipsets
24 that supported K7.

25 Q. Now, with regard to the K8, does AMD plan on

1 selling chipsets for the K8?

2 A. Yes, and we are currently. We have introduced
3 that part. However, K8 has a radically different
4 system architecture, so while we do sell components
5 that can be referred to as chipsets, they don't
6 include the memory controller anymore. We've moved
7 the memory controller onto the K8 microprocessor
8 device.

9 Q. Do you have an understanding of why AMD chose
10 to move that onto the K8 microprocessor?

11 A. The reason we moved the memory controller onto
12 the microprocessor was to gain significantly more
13 performance out of the microprocessor by moving memory
14 logically closer in terms of latency to the
15 microprocessor.

16 Q. Did you have any role in the decision to
17 include that on the K8?

18 A. Yes.

19 Q. And what did you understand to be the benefits
20 to AMD for including the memory controller on the
21 microprocessor design?

22 A. Performance. We were able to deliver
23 significantly better performance delivered to the
24 customer by having the DRAM very close.

25 Q. That's what I'm wondering. What do you mean by

1 "performance" in this case?

2 A. Oh, okay. The way microprocessors work is
3 they're executing a program and they need to access
4 data that is not contained locally in the
5 microprocessor. It needs to go out to main memory to
6 get that data.

7 In the old, traditional way, that request would
8 have to go from the microprocessor across the
9 motherboard to the chipset. The chipset would have to
10 figure out what to do, Northbridge chipset. It would
11 then send that request to the memory. It would wait.
12 The memory would return the data. The data again would
13 have to go through that Northbridge chipset again
14 across the motherboard to the microprocessor. That's a
15 long time.

16 If the memory controller is integrated with the
17 processor, that time can be cut down significantly,
18 allowing programs to execute faster.

19 Q. Do you have an understanding of whether the K8
20 is designed to work with any particular kind of DRAM?

21 A. Yes. It's designed to work with DDR DRAM.

22 Q. I understand DDR has a number of different
23 speeds. Is it designed to work with any particular
24 speed of DDR DRAM?

25 A. We designed it to run with DDR 200, 266, 333

1 and 400.

2 Q. Are you aware of whether AMD has attempted to
3 develop products for use with RDRAM?

4 A. Yes.

5 Q. And what was AMD attempting to make work with
6 RDRAM?

7 A. We were planning to have the follow-on chipset
8 for K7 support RDRAM.

9 Q. And what was your role in this project?

10 A. I was again the design manager of the team.

11 Q. Was there a code name for this project?

12 A. Yes. The code name for that project was IGR4.

13 Q. And what did that stand for?

14 A. So IG is Irongate from the original chipset,
15 R was Rambus, and 4 was 4XAGP, which is the graphics
16 support that we also put on that second-generation
17 chip.

18 Q. When did that activity take place?

19 A. I believe we started up the design team working
20 on it in earnest in the fall of 1998.

21 Q. And were you aware of how the decision was made
22 at AMD to begin this project?

23 A. Yes.

24 Q. Did you have any role in this decision?

25 A. Yes.

1 Q. Now, what was your impression of DDR at this
2 point?

3 A. Well, DDR in I guess '97 or so was not in very
4 good shape. A lot of the DRAM folks were very
5 supportive of it and were pushing it and saying this is
6 a great technology. The problem was that the
7 technology was addressed at a component level. The
8 DRAM manufacturers were specifying a component, a DRAM
9 component.

10 There were no other chipset folks signed up
11 that we knew of. No one had looked at the
12 system-level aspects of implementing DDR. No one had
13 laid out a DIMM, a DIMM module, for example. No one
14 had figured out the placement for termination
15 resistors, enabled clock chip vendors, power supply
16 vendors, et cetera. There was a lot of work left to
17 be done.

18 Additionally, the DDR DRAMs being offered by
19 the DRAM companies differed slightly in between company
20 to company. One company would have a slightly
21 different spec than the other one. It wasn't a unified
22 spec at that point.

23 Q. What was the importance of there being
24 different specs between the different DDR
25 manufacturers?

1 A. That would have been very bad to design a
2 memory controller to work with different flavors of the
3 same device. It was crucial that we design systems
4 that could accept a DDR device from any manufacturer,
5 again, to get to that commodity state where many
6 manufacturers offer the same device at low cost.

7 If the devices were different in slightly
8 different ways, it would cause incompatibilities or
9 force the memory controller design A and B to implement
10 all sorts of weird hacks and it would just be a mess.
11 It was crucial that we had a common standard that would
12 allow interoperability.

13 Q. I'd like to show you a document that's been
14 marked for identification as CX-2153. I hope we have a
15 copy of it there.

16 A. Yep.

17 MR. GATES: Do you have a copy for me?

18 MR. DAVIS: Yes. I'm sorry.

19 BY MR. DAVIS:

20 Q. Do you recognize this document?

21 A. Yes.

22 Q. And what is it?

23 A. It is an e-mail that I wrote in response to a
24 note from Eric Hsu, who was an account representative
25 in our infrastructure enablement group. I believe he

1 was responsible for the Samsung relationship at this
2 point.

3 Q. Now, at the bottom of the e-mail -- at the
4 bottom of the document is an e-mail from Eric Hsu to
5 you.

6 What was your understanding of what he was
7 telling you in this e-mail?

8 A. At this point in time Samsung was trying to
9 have us adopt DDR for our K7 chipset, our follow-on K7
10 chipset, and they would come and present us information
11 and I would ask, well, we need, you know, A, B, C, D,
12 you know, a list of things we needed to help us make
13 the decision one way or another, and Eric was saying
14 that his contacts at Samsung were calling to make sure
15 I had all the information I needed.

16 Q. What sort of information was Samsung offering
17 you?

18 A. They were offering device specifications.

19 Q. Was that the information that you were
20 interested in receiving?

21 A. No. There were two main points. We needed
22 more of the system-level information. We needed the
23 layout design guidelines, a DIMM specification, a DIMM
24 layout, termination guidelines, power supply, clock
25 chips. The whole system solution is what we were

1 really looking for.

1 if AMD tried to do that on their own, it would be a
2 huge effort on our part.

3 Q. And 2 of the e-mail under the information you
4 wanted states, quote, "The Samsung person that I spoke
5 to also agreed that the system-level stuff was not very
6 well understood. I pointed out that the difficulties
7 that Intel had getting PC100 DIMMs correct and we
8 agreed this would be a big effort."

9 What did you mean by a system-level spec and
10 design guide? I think you referenced those.

11 A. That's exactly what I meant.

12 Q. And what is a system-level spec with a design
13 guide?

14 A. A design guide, a system-level planning
15 budget, termination resistor placement, termination
16 guidelines, voltage regulator specs, clock specs,
17 et cetera.

18 Q. And why are these important to AMD in whether
19 they're going to adopt DDR?

20 A. To get all of that, all those specifications,
21 correct required a lot of work, and AMD at this point
22 wasn't prepared to do all that work by ourselves. We
23 were hoping that an industry consortium, you know,
24 the -- everybody would work together to arrive at the
25 right answer.

1 MR. DAVIS: I'd like to move CX-2153 into
2 evidence.

3 MR. GATES: No objection, Your Honor.

4 JUDGE McGUIRE: Entered.

5 (CX Exhibit Number 2153 was admitted into
6 evidence.)

7 BY MR. DAVIS:

8 Q. Could you describe the work that was done at
9 AMD in order to allow the K7 to work with RDRAM.

10 A. That would be the design of the IGR4
11 Northbridge chipset, and the work that was done there
12 was broken into two parts.

13 The first part was the actual design of the
14 chipset itself. We started with the base Irongate
15 design -- that was the IG -- and added a Rambus memory
16 controller to it.

17 In parallel, part of the deal that was arrived
18 at with Rambus was that Rambus was going to design the
19 I/O cell. There's a complicated set of logic and
20 analog circuitry called a rack that needs to be
21 designed very, very carefully, that Rambus does that
22 design in the semiconductor process of interest for the
23 partner.

24 So Rambus did that design for us. We took that
25 information and taped it out on a test chip and into

1 the semiconductor fabrication facility that we were
2 going to use, which I believe was UMC at this point,
3 and we built a package for it and built the part and
4 got it back and tested it and qualified it, made sure
5 that that design was solid for later inclusion in the
6 main part.

7 Q. Do you know if AMD ended up implementing the
8 Irongate Northbridge for DRAM?

9 A. We brought it to the point of pretty much
10 completion of the design, but we decided late in the
11 design not to tape it out but rather to put the
12 project on the shelf for later, later production
13 possibly.

14 Q. What do you mean, put it on the shelf?

15 A. We took the design database and added the
16 libraries and the CAD tools and the system-level
17 elements in a big data backup so that we could, when
18 we chose to restart the project, we could just

1

1 reason was that we had earlier decided to start a
2 parallel effort designing a DDR chipset and we decided
3 that there was a lot of work to bring both of those to
4 market and we needed some resources freed up from the
5 IGR4 team to help out our microprocessor design effort
6 at the time.

7 Q. I'd like to show you a document that's been
8 marked for identification as CX-2158.

9 A. Okay.

10 Q. Do you recognize this document?

11 A. Yes.

12 Q. Did you write this document?

13 A. Yes.

14 Q. Why did you write CX-2158?

15 A. I was requested to give a quick history of our
16 engagement with Rambus and DDR to Dirk Meyer.

17 Q. And who is Dirk Meyer?

18 A. Dirk Meyer at that point was vice president of
19 microprocessor design in our group.

20 Q. Did you do anything to arrive at the
21 information that was presented on CX-2158?

22 A. Yeah. I most likely went through my e-mail
23 folders and probably through my notebooks to get the
24 dates right, put it in chronological order.

25 Q. I'd like to take a look at this document if we

1 can.

2 The first bullet says, "Initial direction was
3 DDR main memory for K7," and then it says, "Top-level
4 DDR attributes were very attractive to K7."

5 What did you mean by that?

6 A. K7, when we designed it, we aimed for clear
7 leadership performance in the marketplace, and a large
8 part of microprocessor performance is memory bandwidth,
9 and they designed -- one of the reasons that we
10 designed our own front-side bus, we adopted the S2K,
11 was that had offered much higher performance, much
12 higher bandwidth capabilities, and the initial look out
13 into the DRAM industry suggested that DDR DRAM had
14 bandwidth numbers that matched what we were looking at
15 for K7, so it seemed like a good match, just on a
16 performance bandwidth perspective.

17 Q. Okay. Below that the note that there were
18 discussions with Samsung, Micron and NEC it looks like
19 May of 1997?

20 A. Yes.

21 Q. And it quickly became apparent, and there's a
22 number of bullets below that.

23 First, it says "not one standard" and I think
24 you've already talked about that?

25 A. Yes.

1 Q. There are a couple of sort of technical things,
2 different I/O voltages, different protocols.

3 What were you referring to there?

4 A. Some manufacturers were attempting to start DDR
5 using 3.3-volt I/O signaling and others were going to
6 2.5 volt. There was no commonality between a few of
7 them.

8 The differing protocol, some manufacturers were
9 proposing bidirectional strobes, others were proposing
10 no strobes. It was just basic parts of the protocol
11 just were not nailed down in common between all the
12 manufacturers.

13 Q. You just mentioned the term "strobe." What do
14 you mean by that?

15 A. A strobe is a signal that is used to qualify
16 data transfers. When data is transferred between the
17 DRAM and the memory controller.

18 Q. What do you mean, "qualify data transfers"?

19 A. Consider it a clock. It's a clock that goes
20 along with the data to specify where a data bit starts
21 and ends.

22 Q. Below that you say "no DIMM infrastructure
23 planned" and then in parentheses you say "two DIMM
24 standards."

25 What did you mean by a DIMM infrastructure?

1 A. Back to some of what I said earlier, no one had
2 come up with a common DIMM. DIMM stands for dual
3 in-line memory module and it's the circuit board that
4 has the DRAM component soldered to it for eventual --
5 the DIMM gets plugged into the motherboard.

6 And there were -- there were no common DIMM
7 standards. Everybody had a slightly different version
8 of it, different pinouts, different form factors.
9 Again, back to the commodity statement, it's crucial
10 that we have a common standard for the DIMM definition
11 so that anybody's DIMM plugs into the motherboard and
12 works just fine.

13 Q. Now, below that you say, "No support
14 components, power supply, clock chips and registers."

15 What were the importance of -- first of all,
16 what's a clock chip and what's a register?

17 A. A clock chip is a chip, you know, an electronic
18 circuit that generates clocks to be fed out to DRAM
19 DIMMs.

20 Q. And I'm sorry. And what's a register?

21 A. A register was -- one flavor of DIMMs are
22 referred to as register DIMMs, and what they do is they
23 take some set of signals from the memory controller and
24 they register them. That's a small component that
25 consists of a series of flip-flops that register the

1 commands and then redistribute them on the DIMM to all
2 the DRAM components.

3 Q. I'm not sure if this is a good question to ask,
4 but what is a flip-flop?

5 A. Oh, boy. It's a basic state element that
6 captures data and holds it.

7 Q. Now, why was the availability of the system
8 components important to AMD's adoption of DDR?

1 Rambus in-line module I believe. Instead of DIMM,
2 it's RIMM.

3 But they had that standard. They had the
4 layout done. They had the termination scheme
5 specified. They had power supply components
6 specified. They had clock chips specified. All of
7 the support structure and components required to make
8 a computer using Rambus technology, they had figured
9 it all out.

10 Q. And you also say here "committed system logic
11 partner." What was the importance of that?

12 A. That was obviously Intel had committed to
13 including Rambus in their product line, and therefore,
14 Intel was designing a chipset that would interface to
15 Rambus, and given that there was a chipset manufacturer
16 working on how to talk to Rambus that was driving a lot
17 of the issue closure on the system-level solutions, so
18 issues around pinout and layout and guidelines,
19 et cetera, were obviously being worked with a committed
20 system partner.

21 Q. Below that you say: "We realize that creating
22 an S2K infrastructure was going to be tough enough. We
23 really could not afford to create the DDR
24 infrastructure at the same time."

25 I think you mentioned what S2K was earlier.

1 Could you remind me what that is.

2 A. So S2K was the term we used for the new K7
3 front-side bus, the new interface that we designed for
4 K7 to talk to the rest of the system.

5 And many of the things I've been talking about,
6 system-level components and design guidelines, a lot of
7 that work had to be done also to get our partners

1 work that AMD did with Rambus. The second bullet in
2 particular says, "We worked closely with Rambus to use
3 their RMC logic in the process, debugging numerous
4 issues with RMC."

5 First of all, tell me, what does RMC mean?

6 A. RMC I believe stands for Rambus memory
7 controller, and what it was was Rambus developed this
8 chunk of verilog RTL code that implemented an abstract
9 memory controller, the set of logic that you would want
10 to design to be a memory controller for their Rambus
11 device.

12 We quickly found out that it was not
13 implementable. It was a very high-level abstract
14 design, but when we tried to actually take this
15 abstract and make a chip out of it, it wasn't very
16 translatable, so we had to do a lot of work to make it
17 implementable in a real chip.

18 In addition, when we started doing extensive
19 simulation on the chip to verify its functionality, we
20 discovered some bugs, a couple big ones, a lot of
21 little ones, that were -- just had to be worked
22 through, and there were questions on the exact protocol
23 and the spec and a lot of interaction at the
24 engineer-to-engineer level to get that right.

25 Q. Now, you mentioned the term "verilog RTL."

1 Just so I understand, what is verilog RTL?

2 A. Yeah. I apologize. That's a common
3 programming language that chip designers use to design
4 chips. It's the base design language of most chips
5 these days.

6 Q. Now, next you say between the summer of
7 1997 and -- the summer of '97 and the fall of '98,
8 JEDEC DDR turned around due mostly to the efforts of
9 the folks at Micron who partnered with Micron PC folks
10 to address the DDR system issues.

11 What are you referring to there?

12 A. Well, as I mentioned, what was missing in the
13 whole DDR equation was a system partner that would
14 drive all the system issues to get the layout, design
15 guides, et cetera, put together. And it turned out
16 that Micron had a division of their company that was
17 developing PCs and chipsets, so they hooked up --
18 they, the DRAM Micron folks, hooked up with the
19 Micron PC folks and started working out all these
20 issues.

21 Apparently, the Micron PC folks designed a
22 chipset, a DDR chipset, talked to an Intel
23 microprocessor, but other than that, it was fine and it
24 worked great with DDR. They built a motherboard,
25 worked through the termination issues. They made a

1 DIMM layout that seemed to be pretty good, got some
2 power supply vendors on board to develop power supplies
3 that met some pretty reasonable specs, clock chips,
4 et cetera. It seemed like they had the whole package
5 sort of put together.

6 Q. Now, next you refer to a trip, AMD traveled to
7 Micron PC in Minnesota.

8 Were you involved in that trip?

9 A. Yes.

10 Q. And what happened in that trip? First of
11 all -- I'm sorry. I shouldn't have asked that.

12 Why did you go on that trip?

13 A. Well, we were invited by Micron PC -- by the
14 Micron folks, both the DRAM and PC folks, to come see
15 that DDR was real and that we should consider it for
16 our future products.

17 So myself and Jim Keller, who at the time was
18 the K8 processor architect, flew to Minneapolis and
19 went and saw what Micron had done. We saw the
20 motherboard. We looked at their layouts. We went into
21 the lab and saw it working and put a scope on some
22 signals and saw the pretty signal integrity and it was
23 real, things were working well.

24 Q. And who at Micron were you working with on
25 this?

1 A. The Micron DRAM person was Terry Lee and the
2 Micron PC person was Joe Jeddlow, J-E-D-D-L-O-W.

3 Q. Next you say, "In parallel, lots of
4 information from DRAM manufacturers on the real cost
5 of DRAM."

6 Before we talk about the different
7 information, what was the importance of receiving
8 information from the DRAM manufacturers on the cost of
9 DRAM manufacture?

10 A. The DRAM manufacturers are the final say in
11 what it costs to manufacture their part. If you want
12 to find out how good or bad a DRAM part is, you ask the
13 DRAM manufacturer. They kept us up-to-date on what
14 their costs were for various technologies moving
15 forward.

16 Q. And why was it important for AMD to talk to the
17 DDR manufacturers about the cost?

18 A. We needed to make sure that whatever memory we
19 chose in our systems for our microprocessors was a
20 commodity and met the performance requirements at the
21 lowest possible cost.

22 Q. Now, some of the items listed below, let's go
23 through some of those.

24 What was your understanding of the term "die
25 area"?

1 A. What we were being told by the DRAM
2 manufacturers was that comparing a like density DDR
3 chip to a like density Rambus chip, the Rambus chip had
4 larger die size. It was physically a bigger chip,
5 which implied it was more expensive.

6 Q. Why does the larger size mean it is more
7 expensive?

8 MR. GATES: Objection, Your Honor. It lacks
9 foundation.

10 JUDGE McGUIRE: I can't hear you, Mr. Gates.

11 MR. GATES: It lacks foundation, Your Honor.

12 JUDGE McGUIRE: Sustained.

13 BY MR. DAVIS:

14 Q. Have you been involved in chip manufacture?

15 A. Yes.

16 Q. And in your involvement as a chip manufacturer,
17 have you observed the importance of die size to the
18 cost?

19 A. Yes.

20 Q. Okay. So could you answer what was the
21 importance of die area.

22m0: ansity DDf.mportan02cnie miconductorce of die area.

1 kay.trynce doour involvedesignchipur ine of die area.

e miconductorcment as a chipr sce ty Dabsolutwer wh 3.033961

1 minimum die size possible to meet your product
2 requirements because there's only a certain number of
3 die that fit onto a wafer and only a certain number of
4 wafers that can go through a manufacturing facility in
5 a given time period and that equals your -- that's the
6 basis of your economic model.

7 Q. And test issues, what were you being told about
8 test issues?

9 MR. GATES: Objection, Your Honor. It calls
10 for hearsay.

11 MR. DAVIS: I'm not entering it for the truth
12 of the matter but his state of mind in the decisions
13 regarding the RDRAM chipset.

14 MR. GATES: If it's coming in under state of
15 mind and not for the truth of the matter, I don't have
16 any objection, Your Honor.

17 JUDGE McGUIRE: Noted.

18 BY MR. DAVIS:

19 Q. Do you have the question in mind? Do you
20 remember the question?

21 A. No.

22 MR. DAVIS: Would you read back the question,
23 please.

24 (The record was read as follows:)

25 "QUESTION: And test issues, what were you

1 being told about test issues?"

2 THE WITNESS: DRAM manufacturers were telling
3 us that the Rambus devices required new, expensive
4 testers and required longer test time than equivalent
5 DDR devices.

6 BY MR. DAVIS:

7 Q. And why was that important to you?

8 A. Similar to the die size issue, the longer your
9 test time, the higher your cost. You want to minimize
10 your test time when you manufacture semiconductors.
11 The longer you have to test a device, the less your
12 throughput through your manufacturing line.

13 Q. The next item you have there is yield. What
14 does that refer to?

15 A. That refers to after manufacture, whether a
16 device meets its specifications or not. Devices that
17 meet the specification are a yield, that's your yield,
18 your positive yield, and the ones that don't are thrown
19 away.

20 Q. And what were you being told about yield on
21 this topic?

22 MR. GATES: Your Honor, it's again calling for
23 hearsay. If it's coming in under state of mind,
24 that's fine. I would just like to have a standing
25 objection.

Ctquired longer test timhrown

1 MR. DAVIS: And these questions will all be for
2 the state of mind.

3 JUDGE McGUIRE: All right. So noted.

4 MR. GATES: Thank you.

5 THE WITNESS: So on yield we were being told
6 that the Rambus devices weren't coming in as fast as
7 they needed to be and therefore they were going to be
8 quite expensive initially.

9 BY MR. DAVIS:

10 Q. The next item below that is power. What did
11 that mean?

12 A. We were being told by the DRAM manufacturers
13 that the Rambus devices were quite a bit hotter in
14 operation than the comparable DDR device.

15 Q. And why is that important?

16 A. Higher power would imply higher cost in the
17 system. A larger power supply and perhaps a heat sink
18 to remove the power.

19 Q. And what were you being told about power by the
20 DRAM manufacturers?

21 A. We were being told that the Rambus devices
22 required heat sinks and the DDR devices didn't,
23 implying that the Rambus devices would be more
24 expensive.

25 Q. The next item below that is packaging. What

1 does that refer to?

2 A. We were being told by the DRAM manufacturers
3 that Rambus devices required a new, higher-cost
4 packaging mechanism -- I believe it was BGA -- whereas
5 the comparable DDR devices were still being packaged in
6 the current-generation package technology. I believe
7 it's referred to as TSOP, T-S-O-P.

8 Q. And why was that important?

9 A. Again, BGA was more expensive than TSOP.

10 Q. And I'm sorry. Did you say what you were
11 hearing from the DDR manufacturers?

12 A. Yes. This is all what the DDR manufacturers
13 were telling us.

14 Q. And the final item there is RIMMs, and it's
15 capital R-I-M-M-S, all caps up to the last M. Sorry.

16 What does that refer to?

17 A. A RIMM is the term that Rambus coined for
18 their memory module, their -- analogous to the DIMM.
19 Rambus in-line memory module I believe is what it
20 stands for.

21 And in the context of this statement, what the
22 DRAM folks were telling us is that the RIMMs were hard
23 to manufacture because of the yield issues and the test
24 issues. They had to do a lot of testing after the
25 devices were assembled onto the RIMM itself, and

1 therefore, if they found a bad component out of the
2 multiple components on a RIMM, they'd have to throw the
3 whole RIMM away or rework it, which is just adding
4 cost, large amounts of cost, to the actual RIMM
5 manufacture process.

6 Q. And why was that important to AMD?

7 A. Again, we needed to make sure that the memory
8 required for use with AMD microprocessors was
9 cost-effective and commodity, widely available at the
10 lowest possible cost.

11 Q. The next bullet refers to future K8
12 implementations. K8 again was -- what was K8?

13 A. K8 was our next-generation microprocessor,
14 sometimes referred to as Hammer.

15 Q. And then you say here "It became clear WRT,"
16 does that mean with respect to?

17 A. With respect to, yes.

18 Q. The difficulties of getting a Rambus controller
19 on die.

20 What are you talking about there?

21 A. Here, I think I referred to previously that our
22 memory controller for K8 we decided to integrate onto
23 the microprocessor die itself. Traditionally, memory
24 controllers are implemented in the Northbridge chipset,
25 a separate chip from the microprocessor, and there's

1 some fundamental differences between the semiconductor
2 processes used for the manufacture of chipsets and the
3 semiconductor processes used for the manufacture of
4 microprocessors.

5 With chipsets, the goal is lowest possible cost
6 with a constant control of the process technology so
7 that when you start manufacture of your Northbridge
8 chipset, you assume that the semiconductor process
9 technology parameters aren't going to move around over
10 the lifetime of the product. That product will remain
11 stable.

12 And that's particularly crucial in Rambus with
13 their rack design. As I referred to earlier, the rack
14 is their sophisticated I/O cell design that has a lot
15 of very detailed analog, careful analog design, that
16 depends upon, critically upon stable semiconductor
17 process technology parameters. That works well and
18 when the Northbridge -- the memory controller is
19 implemented in the stable process technology such as,
20 you know, the ones from UMC or TSMC.

21 On the microprocessor side, it's a completely
22 different equation. The way microprocessor
23 semiconductor process technology is targeted, you tape
24 out the microprocessor initially and then you're
25 constantly tweaking the process technology to ever

1 increase the speed of that microprocessor over its
2 lifetime.

3 If we had decided to implement the Rambus
4 memory controller on the microprocessor using the
5 microprocessor semiconductor technology, we would have
6 been constantly tweaking and constantly having to tape
7 out new versions of the microprocessor to adapt to the
8 changing process technology underneath it in order to
9 keep that rack cell working.

10 The rack cell is quite dependent upon --

11 JUDGE McGUIRE: Okay. I'm not going to let him
12 just go on and on. I understand what you're trying to
13 say, sir, but let's ask more tighter questions so we
14 don't have four or five pages of narrative, if you
15 would.

16 MR. DAVIS: That was my fault, Your Honor. I'm
17 sorry.

18 BY MR. DAVIS:

19 Q. Just so I understand it, AMD had made a
20 decision not to use RDRAM with the K8; is that
21 accurate?

22 A. That was the direction at that time, yes.

23 Q. Did that decision have any effect on the
24 decisions that were made regarding the K7?

25 A. Yes. We wanted to make sure that when we

1 had -- when we introduced K8 with DDR that we had an
2 established infrastructure for DDR, and a good way to
3 do that would be to introduce a DDR chipset for K7.

4 Q. The next item refers to a complete design team
5 freed up in AUS. Does that refer to Austin?

6 A. Yes.

7 Q. And what's that item in parentheses that says
8 "EPD"?

9 A. Embedded products division, which was a
10 division that more or less went out of business at
11 about that time.

12 Q. Now, what was the importance of the design team
13 being freed up in Austin?

14 A. That team could be applied to doing a parallel
15 chipset effort, and in fact we let -- asked them to
16 take on the task of designing a DDR version of the K7
17 chipset.

18 Q. Okay. Now, one item below that, you say, "Due
19 to resource constraints and continuing bad news about
20 RDRAM in April 1994, IGR4 was put on the shelf with
21 full intentions to restart later."

22 What was the bad news that you were referring
23 to here?

24 A. There started to be industry rumors, industry
25 articles about problems with Rambus. The Intel

1 A. We had to make sure that all of our partners
2 designed the motherboards properly, that all of the
3 DRAM DIMMs that were manufactured by all our partners
4 worked properly, a big qualification, making sure all
5 the DRAM devices and all the DIMMs worked interoperably
6 with the different motherboards that were all designed
7 to work with IGD4.

8 At the same time, we had a few bugs towards the
9 end of, you know, the end of the summer getting that
10 into production. It's the general ramp-up to
11 production.

12 MR. DAVIS: Now, I understand that this has
13 already been moved into evidence. If it hasn't, I
14 would move it in, but I understand it has been.

15 JUDGE McGUIRE: And if it has, then it's not
16 pertinent that it be now moved.

17 MR. DAVIS: Okay. I'll move it in later.

18 JUDGE McGUIRE: I'm trying to be -- it has been
19 moved into evidence already; is that correct?

20 MR. GATES: Your Honor, I believe it was moved
21 in on Friday. I just don't have my list here.

22 JUDGE McGUIRE: But we aren't sure?

23 MR. DAVIS: No. I understand that it was.

24 JUDGE McGUIRE: It was or was not?

25 MR. DAVIS: I believe it was.

1 JUDGE McGUIRE: Well, in case it hasn't been,
2 is there any objection?

3 MR. GATES: There's no objection.

4 JUDGE McGUIRE: Then if not, it will be entered
5 if it hasn't already been entered.

6 MR. DAVIS: Thank you, Your Honor.

7 JUDGE McGUIRE: Thank you, Mr. Davis.

8 BY MR. DAVIS:

9 Q. Are you aware of an organization known as
10 JEDEC?

11 A. Yes.

12 Q. When did you first become aware of JEDEC?

13 A. When I was at Apple, the procurement folks
14 updated us periodically about the ongoing industry
15 situation and JEDEC was mentioned.

16 Q. Do your current job duties involve JEDEC in any
17 way?

18 A. Yes. The AMD representative to JEDEC works on
19 my team.

20 Q. Who is that?

21 A. That's Sam Patel.

22 Q. Could you tell me what JEDEC is.

23 A. It's an industry consortium that defines
24 standards that multiple manufacturers can design to to
25 have interoperable parts.

1 Q. Have you ever gone to a JEDEC meeting?

2 A. I've been to a few. Yes.

3 Q. How many, approximately?

4 A. Maybe three or four of them.

5 Q. Why did you go to the JEDEC meetings?

6 A. I go to back up Sam when he needs extra help.

7 There's usually a lot of issues that occur during a
8 JEDEC meeting. They take a week, but they try to pack
9 two weeks' worth of work in and there's a lot of
10 sidebar conversations, a lot of issues to address, and
11 when those issues become overwhelming for one person, I
12 go to support Sam.

13 Q. What is your understanding of the importance,
14 if any, of the JEDEC standards to AMD's business?

15 A. AMD views the JEDEC standards process as
16 crucial to its business. JEDEC allows manufacturers to
17 all design to a common standard and basically enables
18 the commodity marketplace. Everybody is designing
19 compatible parts at the lowest possible cost competing
20 on manufacturing cost.

21 Q. I'd like to show you a document that's been
22 marked for identification as RX-1839.

23 A. Yes.

24 Q. Do you recognize this document?

25 A. Yes.

1 Q. What is it?

2 A. It's an e-mail trail, set of exchanges, where
3 I was preparing for a participation in a panel
4 discussion at a VLSI symposium in Japan scheduled for
5 June 2001.

6 Q. What was the VLSI rump session that's
7 referenced in the subject line?

8 A. The VLSI symposium is this big conference of
9 all the VLSI designers, VLSI semiconductor designers,
10 in the world and they have day-long sessions of very
11 boring, highly technical presentations, and in the
12 evening they had some panel discussions, and the
13 request was to have a lively panel discussion about
14 DRAM and they asked various people to represent
15 different technologies and encouraged a lively
16 discussion.

17 Q. What do you mean, "a lively discussion"?

18 A. Controversy. Conflict. They wanted to put a
19 little juice in the proceedings from a day-long,
20 boring semiconductor discussion to a lively
21 presentation in the evening to keep the engineers
22 entertained I guess.

23 Q. And the presentation that you have in the back
24 is the presentation of that?

25 A. Yes.

1 Q. In the last paragraph, I guess the second to
2 last full paragraph, it's the one starting "Are DDR
3 memory controllers," there's a sentence that says,
4 "Alternate memory technology interfaces use roughly
5 65 percent of the memory controller resources that a
6 typical 64-bit DDR interface requires."

7 What is that referring to?

8 A. There were questions -- I don't honestly
9 remember the source of the questions, but there were
10 questions related to, gee, Rambus only requires, you
11 know, some number of pins and DDR requires a larger
12 number of pins, doesn't that mean that Rambus is less
13 expensive to implement. And this paragraph explains my
14 response that it's -- you just can't count pins to get
15 your cost. There's other issues to think about.

16 Q. In fact, the next sentence says, "However,
17 saving a few pennies in the memory controller while
18 spending quite a few dollars in other parts of the
19 system" -- and then there's a parenthetical there --
20 "is not the right trade-off."

21 And what were you talking about there?

22 A. At that point in time the -- our understanding
23 was that Rambus required -- implementing Rambus
24 technology in the motherboard required six-layer
25 motherboards. There were heat sinks attached to the

1 Rambus. There's -- the sockets for Rambus were quite
2 expensive compared to DDR at that point. The same
3 thing goes for the regulators and clock chips. They
4 hadn't achieved the commodity price point yet. In
5 addition to the basic cost issues of the device itself
6 that we talked about earlier.

7 Q. The next sentence starts, "AMD believes that
8 the sweet spot for mainstream memory controller
9 technology," and it goes on from there.

10 What do you mean by "sweet spot"?

11 A. Our Chinese partners who develop chipsets and
12 motherboards had -- were very familiar with given size
13 packages and layout rules, four-layer motherboards,
14 et cetera, that we believed that the DDR requirements
15 matched that very well, whereas the RDRAM requirements
16 did not.

17 Q. Okay. At the bottom of the next page, there's
18 an e-mail that I believe is from you and if you could
19 confirm this.

20 A. Yes.

21 Q. So starting with the -- it says "original
22 message"?

23 A. Yes.

24 Q. Is that -- so is the text underneath that, the
25 last original message, is that something that you

1 wrote?

2 A. Yes.

3 Q. The very last line says, "JEDEC is a pain."

4 What did you mean by that?

5 A. JEDEC is open to any and all parties, so any
6 and all parties have an opinion and can contribute or
7 delay, or everybody has a vote, so it's not always the
8 most straightforward thing to get a technical
9 specification through. It's sometimes long, laborious,
10 and you have to argue your points endlessly, probably

1 Q. I'd like to point you to the last -- I believe
2 it's the last three pages of the document.

3 A. Is that the actual presentation?

4 Q. Yes, the very last -- the last three pages
5 starting with the third from the end.

6 A. Okay.

7 Q. Which is actually page 7 of the document.

8 A. Okay.

9 Q. What are these slides?

10 A. These are the slides that I presented to start
11 out the panel rump session. They asked us to have a
12 few slides, speak for five minutes to give a base
13 opening presentation of what our position was, whatever
14 technology we were advocating.

15 Q. And who put these slides together?

16 A. I did.

17 Q. Were these the slides that were used at the
18 session?

19 A. Yes.

20 Q. A few bullets down there's an item that says
21 "die and controller flexibility."

22 What did you mean by that?

23 A. DDR technology allowed both DRAM components and
24 memory controller components to have -- to implement
25 both the old single data rate or a PC100/PC133 spec and

1 with the same component also implement the DDR spec,
 2 which would allow backwards compatibility, and in fact
 3 I believe most of the first DDR DRAM devices were in
 4 fact dual mode. They could work in single data rate
 5 mode and become a PC100 product or they could work in
 6 DDR mode and be a DDR product.

7 Similarly, a number of our chipset partners did
 8 the same thing with their chipsets. They made memory
 9 controllers that could work with either a single data
 10 rate or double data rate.

11 Q. The title of this slide is System Advantages of
 12 DDR.

n ov inight re's ramp or double data raty

fjTnew(or double data ra23)Tjtechnology(12 DD24)TjTf By 2

1 market demanded more single data rate in lower-cost
2 markets, they could supply that. If the market
3 demanded the DDR versions, they could supply it with
4 the same device.

5 Q. Now, earlier you had stated that AMD designs
6 and sells controllers sometimes?

7 A. Yes.

8 Q. Does it generally make them work with both
9 DRAMs and a transition, technology transition?

10 A. We close not to for our IGD4 device. When we
11 started the design of that, we decided to focus on
12 getting the DDR infrastructure going with our IGD4
13 device, get it out there as soon as possible with the
14 lowest risk. Adding an SDR certainly would be
15 possible, maybe even straightforward, but it added
16 risk, but it added not trivial schedule. We decided to
17 go out and get DDR going, get the infrastructure going,
18 get the industry going.

19 Q. Okay. Let's go to the next document.

20 Do you know whether it matters to AMD whether
21 or not the DRAM standard or a DRAM standard is
22 generated inside of JEDEC or inside -- instead of
23 inside a smaller consortium or industry group?

24 A. As long as it's an open, free consortium where
25 all interested parties are able to contribute, it's

1 fine with us. Generally, that means JEDEC.

2 Q. Has there been a situation in which a DRAM
3 standard was developed or was in the process of being
4 developed outside of JEDEC?

5 A. Yes.

6 Q. And what examples are you referring to?

7 A. Well, there was ADT and Rambus.

8 Q. What is ADT?

9 A. ADT I believe stands for Advanced DRAM
10 Technology. I believe it was a consortium of a handful
11 of the DRAM manufacturers and Intel who were defining
12 some next-generation DRAM technology.

13 Q. Were you involved in AMD's business plans
14 regarding DRAM when the ADT standard was being
15 developed?

16 A. Yes.

17 Q. Did you have any understanding of the effects
18 such a standard would have on AMD's business?

19 A. Yes. We were very, very afraid of what it
20 would do to our business. Our number one competitor
21 was essentially defining the new DRAM standard, getting
22 a huge head start in its implementation and details,
23 designing for its requirements without any input from
24 AMD.

25 Q. I'd like to show you a document that's been

1 marked for identification as RX-1746.

1 whether that would affect us or not and how it would
2 affect us.

3 Q. Okay. Do you know what DDR 333 is?

4 A. Yes.

5 Q. What is it?

6 A. It is the follow-on to the original DDR 200 and
7 266 speed specifications, sometimes referred to as
8 DDR 1.5, but it's an evolutionary speed increase in
9 DDR.

10 Q. Are AMD processors compatible with DDR 333?

11 A. Yes.

12 Q. Which processors are compatible with DDR 333?

13 A. Our K8 microprocessor. It's commonly referred
14 to as Opteron now that we've launched it.

15 Q. I'd like to show you a document that's been
16 marked for identification as CX-2152.

17 Do you recognize this document?

18 A. Yes.

19 Q. What is this document?

20 A. This is a -- I believe this is a meeting
21 minutes write-up that I wrote after a meeting.

22 Q. I'm sorry. You said you wrote this?

23 A. Yes.

24 Q. The first paragraph -- well, first of all, it
25 says "IGD4 Based DDR 1.5 Summary," so what was being

1 discussed in this meeting?

2 A. We were discussing the feasibility of adding
3 DDR 333 support to our existing IGD4 part that was at
4 that point I believe starting production.

5 Q. And the first paragraph describes a PPP
6 meeting. What does that stand for?

7 A. Yeah. PPP stands for propeller head platform
8 planning. It's a weekly meeting that is held at AMD
9 that I chair that's a cross-divisional lead
10 engineer/senior engineering discussion group.

11 Q. Okay. Now, what was the time frame in which --
12 who was requesting, first of all, that IGD4 be able to
13 support DDR 1.5?

14 A. Our marketing team at AMD asked for it.

15 Q. In what time frame were they hoping to see this
16 product?

17 A. I believe they were hoping to have a product in
18 the market in mid-2001.

19 Q. What did you do to provide the analysis that's
20 described in this memo?

21 A. We asked the design team, the IGD4 design team,
22 to go off and figure out what it would take to modify
23 the design to support DDR 333.

24 Q. I'd like to direct your attention to the item
25 in the memo entitled What Are the Constraints,

1 Obstacles, Risks, et cetera, to this Project? Do you
2 see that?

3 A. Yes.

4 Q. Could you identify or -- sorry -- could you
5 describe what you identified as the constraints related
6 to the project?

7 A. So typical engineering practice is when you're
8 asked to do something challenging, you want to
9 constrain the problem down to the minimum set of things
10 to solve, so we put some constraints and said, well,
11 we're going to only do this for registered only with
12 two DIMMs and we're not going to entertain other
13 feature additions to the part. We want to narrow the
14 focus of the problem we're trying to solve.

15 Q. Can you describe what you identified as the
16 risks related to the project?

17 A. The biggest risk was that we had other
18 projects in the pipeline that were slated to use the
19 resources that would be needed to make this change to
20 the chip.

21 Q. It says "imposes risks to tunnel projects"?

22 A. Yes. That's what that refers to. The tunnel
23 is what we refer to as the chipsets for K8.

24 Q. Is that similar to bridge?

25 A. Yes. For K8 we changed the system

1 architecture radically, as I discussed, to pull the
2 memory controller on the processor, and to make sure
3 everybody wasn't confused, instead of calling things
4 Northbridge and Southbridge, we called them tunnels.
5 A tunnel does the same thing as a bridge, but it's
6 constructed radically different, so it got the point
7 across.

8 More geeky engineer stuff here. Sorry.

9 Q. Could you identify what you described here as
10 the obstacles related to the project.

11 A. The biggest obstacle -- I think they're
12 detailed a little bit later, but the biggest obstacle
13 was we needed to make some changes inside the IGD4 part
14 itself. The current part had the frequency of the DRAM
15 in locked step with the frequency of the front-side
16 bus. If the front-side bus was going at 266 megahertz,
17 the DRAM would go at 266 megahertz.

18 By asking us to go to 333 megahertz memory, we
19 would have to put some sort of gearbox structure inside
20 the chip to adopt the different clock speeds.

21 Q. What do you mean by the term "gearbox"?

22 A. We would have to make sure that as data went
23 across the chip from the K7 front-side bus side up to
24 the memory it had to change time demands from a
25 266-megahertz time demand up to a 333-megahertz time

1 demand. And that required a bunch of logic that wasn't
2 in there. It wasn't in the chip.

3 Q. You also mention here that there were DDR

1 Q. And what was the nature of the analysis that
2 they asked you to perform on those patents?

3 A. They wanted to see, A, if we were in violation
4 or did we, you know, did we infringe on the patents, I
5 guess is the right word, and B, were there any
6 reasonable work-arounds.

7 JUDGE McGUIRE: Now, let me interject. When
8 you say your boss handed you some patents, whose
9 patents were these? Rambus patents?

10 THE WITNESS: Yes, they were Rambus patents.

11 JUDGE McGUIRE: All right.

12 MR. DAVIS: Thank you, Your Honor.

13 BY MR. DAVIS:

14 uess claim soAMDhe nausn t

1 and the future K8 microprocessor.

2 Q. At what stage was that future microprocessor --
3 this is the K8 you said?

4 A. Yes.

5 Q. And what stage was the K8 at at the time you
6 learned of the Rambus claims against AMD?

7 A. It was about midway through its design cycle.
8 Its base architecture was defined. It was well into
9 the design phase. It hadn't quite reached the phase of
10 physical design, but it was well along the path to
11 completion.

12 Q. Now, you said you reviewed the Rambus patents
13 or three of the Rambus patents?

14 A. Yes.

15 Q. Now, what did you conclude after reviewing the
16 patents?

17 A. They were pretty simple things to work around
18 if we had known about them a long time ago, but we
19 were in the middle of ramping up an infrastructure.
20 This was just when we were trying to get our IGD4 and
21 the first DDR motherboards out the door. They were
22 pretty trivial to work around, but we were in the
23 middle of ramping and they're pretty tough to change
24 things.

25 Q. What was the importance of being in the middle

1 of ramping to the issue of working around the Rambus
2 patents?

3 A. The work-arounds that were obvious required
4 some big changes to the device, to the chipsets, to the
5 motherboards, et cetera.

6 Q. What work-arounds were you referring to?

7 Start with programmable CAS latency. What
8 work-arounds did you have in mind in this analysis?

9 A. There were any number of options. The bottom
10 line is any change when you're trying to do a
11 production ramp is extremely difficult and hard, so you
12 asked about programmable CAS latency; correct?

13 The few work-arounds we talked about were --

14 MR. GATES: Your Honor, can I interject here
15 an objection? I think there's a lack of foundation
16 and we're just getting opinion testimony at this
17 point.

18 MR. DAVIS: This is not opinion testimony.
19 This is his state of mind in 2000 when he learned of
20 the patents, but I can establish a foundation if you
21 like.

22 JUDGE McGUIRE: Overruled. I'll entertain the
23 question.

24 BY MR. DAVIS:

25 Q. Do you remember the question?

1 A. What work-arounds were we contemplating.

2 Q. For programmable CAS latency.

3 A. For programmable CAS latency.

4 So the three on the table were pretty obviously
5 how to fix CAS latency, have the parts manufactured to
6 be a fixed CAS latency. The second one was to have a
7 pin-programmable CAS latency where the value of CAS
8 latency would be set by pins tied high or low on the
9 DIMM. And the third one was some sort of serial
10 programming using existing single-bit serial load
11 technology that was ubiquitous in the industry at the
12 time.

13 Q. So let's talk about pin strapping for a
14 second.

15 What did you mean by pin strapping?

16 A. With pin strapping one could allocate either
17 dedicated or multiplexed pins on the DRAM device that
18 depending on their state either pulled high or low
19 would tell a device what the CAS latency should be.

20 Q. Now, when you say "pulled high or low," what
21 are you referring to (indicating)?

22 A. Pulled to a high voltage to indicate a one or
23 pulled to a low voltage to indicate a zero.

24 Q. Did you think that pin strapping was a more
25 costly or less useful future than the current JEDEC

1 standard of setting CAS latency?

2 A. Certainly no more costly. Maybe a little bit
3 more inconvenient, but in the end it probably could
4 have been made to work just fine. The problem was,
5 we'd have to change everything in the middle of this
6 production ramp.

7 Q. What about fixing CAS latency?

8 A. Fixed CAS latency would have been pretty
9 onerous for the DRAM manufacturers.

10 Q. Let's -- first of all, did you think that
11 fixing CAS latency was more costly or less useful than
12 the current JEDEC standard to AMD?

13 A. Probably it had -- would have a significant
14 cost impact for the DRAM manufacturers.

15 One of the advantages of programmable CAS
16 latency is that DRAM manufacturers can bin their
17 devices. They can have fast devices with a short CAS
18 latency and sell them for more money, and parts that
19 were perhaps yielding slower, they could be programmed
20 with a longer CAS latency and sold for less cost.

21 Q. Now, I forget -- did you remember the name of
22 the person who gave you the patents and the
23 assignment?

24 A. It was either Rich Heye or Dirk Meyer. I
25 forget exactly which one.

1 Q. Did you recommend to either of them that AMD
2 change its products to change the CAS latency that it
3 used?

4 A. No. I recommended we didn't, given that we

1 it relates to AMD's business?

2 A. Fixed burst length would have been very, very
3 bad for AMD. AMD designed its microprocessors to have
4 its natural burst length to be 64 bytes, which is eight
5 cycles of data. Knowing that the DRAMs had that
6 capability, we decided to take advantage of that
7 capability for performance reasons.

8 If the work-around was to fix the burst length,
9 the most likely burst length chosen would have been an
10 Intel-compatible burst length or a burst length of four
11 cycles or 32 bytes. That would have been very bad for
12 us. A, it would have required lots of redesign in the
13 memory controllers and also caused us a performance
14 hit.

15 Q. Now, why did you say that you thought that the
16 most likely burst length would have been the -- you
17 said burst length of four?

18 A. Yes.

19 Q. And why did you think that would be the most
20 likely burst length that would have been used?

21 A. That is the burst length that Intel uses for
22 their microprocessors and their systems and they hold
23 the majority of the market and the DRAM manufacturers
24 would have manufactured the part to hit the majority.

25 Q. I understand.

1 The last technology you would have added, is
2 that dual-edged clocking?

3 A. Yes.

4 Q. And what is dual-edged clocking?

5 A. Dual-edged clocking allows data to be captured
6 on both a rising edge of a clock and the falling edge
7 of the clock. Traditionally, prior to -- I don't
8 know -- the mid-'80s, data was always transferred on a
9 single edge, on the rising edge of the clock. Data was
10 captured when the clock rose from a zero to a one.

11 DDR techniques allowed you to capture the data
12 on the falling edge and the rising edge to effectively
13 double the data rate, hence the word "double data rate"
14 or "DDR."

15 Q. What were the alternatives that you had in mind
16 in 2000 when you reviewed the Rambus patents regarding
17 dual-edged clocking?

18 A. Well, the placement of your clock edges and
19 your data is more or less arbitrary. We could have
20 slowed the clock down by half the rate or doubled the
21 rate of the clock itself. Either way would have been
22 reasonable to implement to capture the data.

23 Q. Did you propose to Mr. Heye or Mr. Meyer or
24 whoever gave you the patents that AMD change its
25 products to accommodate changes to the dual-edged

1 clocking?

2 A. I recommended that we wouldn't make any changes
3 for similar reasons as before. We were in the middle
4 of a production ramp. It would be impossible for us to
5 stop and change.

6 Q. Now, earlier you were describing an analysis in
7 CX-2152, which is -- (indicating).

8 A. Yep.

9 Q. What was the date this analysis occurred? What
10 was your understanding of that?

11 A. The date on the document says June 29, 2000.
12 It probably occurred the previous couple of weeks to
13 that date.

14 Q. And how does that relate in time when you
15 learned about the Rambus patents, the Rambus -- I'm
16 sorry -- the Rambus claims against AMD?

17 A. This preceded that, that knowledge.

18 Q. Is there any difference between the analysis
19 you conducted here and the analysis that you conducted
20 relating to the patents that Rambus showed you -- I'm
21 sorry -- that Rambus claimed against AMD?

22 A. Yes. Here was a very straightforward
23 incremental speed increase and I had the design team
24 look at what it would take.

25 For the work-arounds that we were previously

1 discussing, I communicated with a few of our DRAM
2 partners and a couple of senior folks inside AMD. I
3 didn't go to any team and ask them to do a detailed
4 analysis. It was a very top-level lead work, very
5 simple changes that -- you know, the work-arounds that
6 we talked about were very simple in concept and the
7 reason -- the reasons around doing them are not --
8 didn't surround the technical feasibility. It
9 surrounded the logistical nightmare of trying to change
10 something in mid-production.

11 MR. DAVIS: Thank you. No more questions.

12 JUDGE McGUIRE: Okay. Why don't we take a
13 ten-minute break and when we return we'll start with
14 cross-examination.

15 We're in recess. Off the record.

16 (Recess)

17 JUDGE McGUIRE: Mr. Gates, at this time I'll
18 entertain your cross-examination of the witness.

19 MR. GATES: Thank you, Your Honor.

20 CROSS-EXAMINATION

21 BY MR. GATES:

22 Q. Good afternoon, Mr. Polzin. How are you?

23 A. Hi, Sean.

24 Q. You said earlier that you helped design the
25 front-side bus for the Athlon processor; is that

1 right?

2 JUDGE McGUIRE: Mr. Gates, I'm going to ask you
3 to stand a little closer to the mike. I'm having a
4 little trouble hearing you from that far away.

5 MR. GATES: Okay. Thank you.

6 JUDGE McGUIRE: Thank you.

7 8 THBYR. GATES: jT* 7 9 7 Q ThWell, let mes

51 liinterfaced thit

1 Q. It went from 200 and later on you got a faster
2 bus?

3 A. Yes.

4 Q. And I think at one time you had a 333-megahertz
5 bus; is that right?

6 A. Yes.

7 Q. Well, let me try to get some dates on these
8 changes.

9 Okay. So the AMD Athlon, that's the K7; is
10 that right?

11 A. Yes.

12 Q. And you said the front-side bus speed when it
13 first came out was 200 megahertz; is that right?

14 A. Yes.

15 Q. Okay. And you said you worked on the chipset
16 for that initial version of the K7?

17 A. Yes.

18 Q. And the initial chipset that came out with K7
19 was the AMD 750?

20 A. That's correct.

21 Q. Okay.

22 A. I referred to that previously as Irongate. If
23 that helps.

24 Q. Okay. Or Irongate.

25 And when was that -- when did that chipset come

1 out, what year was that?

2 A. In June 1999 when we launched the first K7.

3 Q. And what type of memory did that chipset
4 operate with?

5 A. PC100.

6 Q. Okay. So it's PC100 SDRAM?

7 A. That's correct.

8 Q. And we'll say that's 6-99; is that right?

9 A. Yes.

10 Q. Later on, you developed at AMD a faster
11 front-side bus, 266 megahertz, for the Athlon chip; is
12 that right?

13 A. Yes.

14 Q. And you worked on the controller that would go
15 with that new front-side bus; is that right?

16 A. No. At that time a different design team was
17 working on that.

18 Q. Okay.

19 A. I was the system guy at that point.

20 Q. You were the system guy at that point.

21 Did it have a different controller than the
22 AMD 750?

23 A. Yes.

24 Q. And that was the AMD 760?

25 A. Yes.

1 Q. And when did this front-side bus come out?

2 A. September 19 -- or September 2000. That's
3 the -- it was coincident with the introduction of the
4 AMD 760 which was our IGD4 chipset.

5 Q. And that AMD 760 chipset was compatible with
6 what memory technologies?

7 A. DDR 200 and 266.

8 Q. Now, did you ever develop a chipset that was
9 compatible with PC133 SDRAM?

10 A. No.

11 Q. Third-party vendors created chipsets that were
12 compatible with PC133 and the AMD processor; is that
13 right?

14 A. Yes.

15 Q. Were they compatible with the 200-megahertz
16 front-side bus or the 266-megahertz front-side bus?

17 A. I believe they were compatible with the 200,
18 but I don't know for certain if they ever were not
19 compatible with the 266. They may have been; they may
20 not have been.

21 Q. Do you know whether or not those chipsets came
22 out in between the AMD 750 and the AMD 7 -- the AMD 750
23 and the AMD 760?

24 A. I believe they do. I don't have any direct
25 knowledge. I can't give you a date.

1 Q. Okay. So at some point third-party vendors
2 came out with chipsets compatible with the Athlon that
3 operated with PC133?

4 A. Yes.

5 Q. And we just don't know the date?

6 A. Yeah.

7 Q. Now, is the architecture of the AMD 760 chipset
8 different from the AMD 750?

1 That wasn't their goal. That's correct.

2 Q. Later on, AMD developed a 333-megahertz
3 front-side bus for the Athlon processor?

4 A. That's correct.

5 Q. And do you know when that came out?

6 A. I really have no idea.

7 Q. Were you aware of it when it came out?

8 A. Yes.

9 Q. Okay. So at some point you knew what the date
10 was?

11 A. Yeah. I just don't recall when it was.

12 Q. Would it refresh your recollection if I showed
13 to you a press release of when that chipset was
14 released?

15 A. Sure.

16 Q. Then why don't we bring up on the screen
17 AMD 04. I'll go ahead and give you a paper copy as
18 well so that you can see that and it's probably easier
19 to read.

20 May I approach, Your Ew18 e3' to read.

1 front-side bus for the Athlon chip?

2 A. Yes.

3 Q. 400 megahertz?

4 A. Yes.

5 Q. And do you remember when that was released?

6 A. Very recently. I know that. I don't have the
7 precise date, but it was within the last couple
8 months.

9 Q. So in about May of this year; is that --

10 A. Yeah. Sounds about right.

11 Q. And what type of memory is that 400-megahertz
12 front-side bus compatible with?

13 A. I believe it's 200, 266, 333 and 400.

14 Q. So it can go with any of the DDR family up to
15 DDR 400?

16 A. I believe so, yes.

17 Q. And did AMD develop a chipset that was
18 compatible with DDR 400?

19 A. Not for K7, no.

20 Q. Not for the K7?

21 A. Yeah.

22 Q. So that was developed by third parties?

23 A. Yes.

24 Q. Now, we looked earlier at a document that you
25 should still have. I think it's CX-2152. It's the one

1 with the copy of this slide in the binder and it has a
2 discussion of DDR 1.5?

3 A. Yes.

4 Q. And that was a discussion of building a chipset
5 that would be compatible with DDR 333; right?

6 A. Yes.

7 Q. In order to do that, at least with the slower
8 bus speed, you had to rearchitecture the Northbridge;
9 right?

10 A. Correct.

11 Q. And wasn't the case that there were different
12 DIMM specifications for the DDR 333 than there were for
13 the previous generations of DDR?

14 A. Different DIMM specifications.

15 Q. Well, let's look at this document on page 3.
16 If you'll look at the third page, and if you look
17 under -- there's a heading DDR Infrastructure Issues
18 and Tasks?

19 A. Yes.

20 Q. If you look at the fourth bullet point, the
21 third and fourth bullet points, it says that there's no
22 device specifications available and later on it says
23 "no DIMM layout available." Doint, clable ,1

12 Q. WS do tIunder stad that wcrrect.lyto rmen thea

1 there is a different DIMM specification for DDR 333
2 products?

3 A. Well, at this point in time it was not known
4 whether the existing one would work fine or whether a
5 new one was required.

6 Q. Well, isn't it -- later on isn't it in fact the
7 case that there was developed a different specification
8 for the DDR?

9 A. I believe you're correct.

10 Q. Okay. Are you familiar with a -- something
11 called a Hot Chip Symposium?

12 A. Yes. Isn't that the one that happens in
13 Stanford once a year? Okay. Yes.

14 Q. And didn't you give a presentation at the
15 Hot Chip Symposium in 1999?

16 A. Me personally or me, AMD?

17 Q. You as part of a group from AMD.

18 A. I don't recall. I don't think I attended.

19 Q. Did you ever give a presentation on the IGR4
20 that you were developing at AMD to a symposium?

21 A. I don't believe I did, no, not personally.

22 Q. But do you know whether or not anybody from AMD
23 did that or not?

24 A. I have no knowledge one way or another. They
25 might have.

1 Q. Now, you spoke earlier about the importance
2 to -- of JEDEC to AMD?

3 A. Yes.

4 Q. And I think you said it was important that it
5 was an open process, everyone could be involved;
6 right?

7 A. Yes.

8 Q. And the reason why it's important for AMD to
9 be participating in an open process is so that you
10 don't get a competitive disadvantage to Intel; is that
11 right?

12 A. That's one of the reasons, yes.

13 Q. And that's one of the reasons why you were very
14 concerned about the ADT consortium?

15 A. Yes.

16 Q. Because AMD was not participating, but Intel
17 was?

18 A. Yes.

19 Q. We looked earlier at an e-mail that you sent.
20 It's RX-1746. It's just a one-pager.

21 A. Okay.

22 Q. Okay. And this was an e-mail that you wrote in
23 response to something that Intel was doing; right?

24 A. Yes.

25 Q. And they were specifying in an addendum

1 A. Correct.

2 Q. -- didn't you have to look at some of Intel's
3 addendums on the DDR 200 specification, for example?

4 A. No. Intel wasn't participating to our
5 knowledge at that point.

6 Q. So you're not aware of whether or not they have
7 addendums to the DDR specifications?

8 A. No.

9 Q. Didn't they write the PC100 standard?

10 A. I have no knowledge. I didn't think so, but I
11 have no knowledge one way or another.

12 Q. Okay. Why don't I show you -- let's bring up
13 RX-2103-14.

14 May I approach, Your Honor?

15 JUDGE MCGUIRE: Yes.

16 BY MR. GATES:

17 Q. Okay. Mr. Polzin, do you recognize that as the
18 Intel PC SDRAM standard?

19 A. That's what it says on the front, yes, sir.

20 Q. And did you use that standard when you were
21 designing the AMD 750 chipset?

22 A. Probably among others we used this, yeah, among
23 other specifications this is probably one we looked at,
24 yeah.

25 Q. So when you were designing the AMD 750 chipset,

1 you were looking at the Intel standard?

2 A. Among others, yes.

3 Q. You also talked earlier about your efforts with
4 the -- to develop the Rambus controller. Do you
5 remember that?

6 A. Yes.

7 Q. Okay. And you had looked at a document where
8 you laid out kind of the chronology of your efforts?

9 A. Yes.

10 Q. Okay. Let's look at that, CX-2158. Do you
11 have that?

12 A. Yep.

13 Q. Okay. Now, if you look down at something
14 Mr. Davis pointed to where it says, on the first page,
15 Rambus offered to complete -- offered a complete system
16 picture?

17 A. Uh-huh.

18 Q. And you explained to us that Rambus had already
19 specified a number of items in the infrastructure;
20 right?

21 A. Yes.

22 Q. And it was important to you at the time
23 because the DDR specification hadn't been settled
24 upon; right?

25 A. Yes.

1 Q. And the DDR infrastructure had not been
2 specified?

3 A. Yes.

4 Q. But the Rambus infrastructure at least had been
5 specified at this point; right?

6 A. Yes.

7 Q. And I think you referred to something -- you
8 referred to a term, you said it's important that the
9 specifications be vendor neutral; right?

10 A. Correct.

11 Q. Were the Rambus -- those Rambus specifications
12 were vendor neutral, weren't they?

13 A. Yes.

14 Q. Okay. And you also said that, well, you might
15 have a DRAM, but just having a DRAM is not enough, you
16 have to have all this infrastructure; right?

17 A. Yes.

18 Q. So without the system infrastructure, you can't
19 have a memory system even if you have that DRAM?

20 A. Correct.

21 Q. Okay. And the reverse is true as well; right?
22 If you don't have the DRAM, all that infrastructure is
23 useless; is that right?

24 A. Yeah.

25 Q. So if the DRAM manufacturers aren't going to

1 be producing a particular type of DRAM, no matter what
2 you do with all the infrastructure, it's not going to
3 work?

4 A. Correct.

5 Q. And at the time when you were designing this
6 chipset for the Rambus RDRAM, it was later on when
7 you'd made the decision to switch to DDR, it was your
8 understanding that RDRAM was not going to become the
9 commodity product; right?

10 A. No. Our understanding was it was -- the reason
11 that we shelved it -- we were very careful to make sure
12 we could restart the project -- it was a timing thing.
13 It was clear it wasn't going to be a commodity product
14 in the time frame of interest.

15 Q. So it was your understanding at least that the
16 reason why you switched over and put all your efforts
17 into a DDR controller was because DDR was going to be
18 the volume product?

19 A. Yes.

20 Q. And you understood that from what you were
21 hearing from the memory manufacturers?

22 A. Yes.

23 Q. At the very bottom of that page, of that
24 e-mail, the first page, you refer to future K8
25 implementations became clear with respect to

1 difficulties of getting the Rambus controller on the
2 die; right?

3 A. Yes.

4 Q. And that you were talking about putting the
5 memory controller actually on the CPU or on the same
6 die as the CPU?

7 A. Yes.

8 Q. And you're referring -- I guess you said
9 earlier that there were some difficulties in doing that
10 with the K8?

11 A. Yes.

12 Q. Now, you're familiar with the Alpha processor;
13 right?

14 A. The old ones.

15 Q. Okay. When you were at DEC, you worked on the
16 Alpha processors?

17 A. Correct. Yes.

18 Q. And DEC later changed and finally became part
19 of Compaq; is that right?

20 A. That's correct, yes.

21 Q. And Compaq is using the Alpha processors in
22 their high-end servers; right?

23 A. Yes.

24 Q. And are you aware of the fact that the Alpha
25 processors have incorporated the Rambus controller onto

1 the die with the CPU?

2 A. I wasn't aware of that, but it sounds
3 reasonable.

4 Q. Why does it sound reasonable to you?

5 A. I know that was their direction they told us

1 it -- keeping that processor competitive as you
2 constantly tweak the process to get faster and faster
3 processors, which you require in the desktop
4 microprocessor marketplace.

5 Q. Do you consider the Alpha processor
6 competitive?

7 A. No.

8 Q. No. Okay. So just because Compaq uses it in
9 its high-end servers it's not competitive?

10 A. It's not very competitive anymore in the
11 marketplace. That's why it's dying.

12 Q. At one time it was; right?

13 A. At one time it was, yes, sir.

14 Q. And do you consider Intel's network processors
15 to be competitive?

16 A. I'm not very aware of how the network processor
17 space, you know, stacks up, so I don't have any
18 knowledge of that.

19 Q. Further down on that second page of the e-mail

2 13 ce. Thcorrectt was, sir.

Waldorf, Marylstet in

1 A. Correct.

1 A. The motherboard companies?

2 Q. The motherboard companies.

3 A. No. The motherboard companies -- we enabled
4 the motherboard companies at about the same time in
5 March of 2000 for our chipset.

6 Q. So you enabled the motherboard companies in
7 March 2000 and so it was your samples --

8 A. Yes.

9 Q. -- motherboard samples?

10 Okay. So by September of 2000 the motherboard
11 manufacturers were able to mass-produce motherboards
12 compatible with the IGD4 chipset?

13 A. Yes.

14 Q. We talked about earlier something -- the VLSI
15 rump session. Do you recall that?

16 A. Yes.

17 Q. Where you gave a presentation?

18 A. Yes.

19 Q. And when you gave that presentation, it was
20 your understanding that each of the participants had a
21 different role; right?

22 A. Yeah. Yeah.

23 Q. So you were the DDR guy --

24 A. Correct.

25 Q. -- right?

1 And you were there to advocate the DDR
2 position; right?

3 A. Yes.

4 Q. And at that time AMD had decided to go with
5 DDR?

6 A. Oh, yes. Yes. Yes. We had been in production
7 I believe at the time.

8 Q. And so you were there to present all the
9 positive aspects of DDR; right?

10 A. Yes.

11 Q. And you weren't intending to present any of the
12 negative aspects, if there were any, of DDR?

13 A. No, I was not intending to present any negative
14 aspects of DDR.

15 Q. And there were other people there who were
16 advocating other memory technology; right?

17 A. That's correct.

18 Q. There was someone there advocating SDRAM?

19 A. Yes.

20 Q. And there was someone there advocating RDRAM?

21 A. Yes.

22 Q. Was that Mr. Kim from Samsung?

23 A. No.

24 Q. Who was advocating RDRAM?

25 A. It was a gentleman from Intel. His first name

1 Q. And so in order to put forth the total DDR
2 position, you were trying to spit back some of the
3 things that the DRAM manufacturers had told you about
4 DDR?

5 A. Yes.

6 Q. So some of the things that are in your
7 presentation at least were just coming directly to you
8 from the DRAM manufacturers; right?

9 A. Yes.

10 Q. And you didn't go out to the DRAM
11 manufacturers to verify everything they were telling
12 you, did you?

13 A. No.

14 Q. So you were just putting forward what they were
15 telling you in your presentation?

16 A. Yes.

17 Q. And if you look at the first page of your
18 presentation -- I guess that was RX-1839, page 7 -- in
19 the first page of your presentation your number one
20 point, the first point that you made was about pricing;
21 is that right?

22 A. Yes.

23 Q. And your point there was that the price of DDR
24 had dropped tremendously because there was volume of
25 DDR shipping; right?

1 A. Yes.

2 Q. And you also -- your second point was that it
3 was available. That means there were -- the DRAM
4 manufacturers were making it; right?

5 A. Yes.

6 Q. And that was in contrast to RDRAM, for example?
7 RDRAM had less availability?

8 A. I didn't make any specific statements to that
9 effect, but the intent was to show that this was
10 available.

11 Q. You were trying to advocate the kind of
12 advantages of DDR over other types of memory; right?

13 A. Yes.

14 Q. And one of the advantages was that it was
15 available while others were not?

16 A. Yes.

17 Q. And Mr. Davis asked you about some different
18 alternatives for Rambus technologies.

19 A. Yes.

20 Q. And the first time you thought of any of these
21 alternatives was in 2000; right?

22 A. Yes.

23 Q. And let me just understand some of your
24 background.

25 Had you ever designed a DRAM chip before?

1 A. No.

2 Q. You've designed microprocessors, but those are
3 different from DRAMs; right?

4 A. Yes.

5 Q. And are you familiar with DRAM manufacturing
6 processes?

7 A. No.

8 Q. Are you familiar with DRAM manufacturing
9 costs?

10 A. No.

11 Q. And so when you told us that some of the
12 alternatives that you had thought of wouldn't have an
13 impact on costs, did you know whether or not it would
14 have an impact on the manufacturing costs of the DRAM?

15 A. I had some discussions with a few DRAM
16 partners, and that's where that data came from. That's
17 where my opinion came from on that matter.

18 Q. Now, let me ask, when you at AMD are confronted
19 with a new DRAM design, something that's been changed,
20 do you just take the manufacturer's word that the
21 change is going to work?

22 A. So it depends on what you mean by "change," so
23 if it's a minor spec change or a minor manufacturing
24 change, we will do a retest to make sure everything is
25 fine.

1 Q. Okay.

2 A. If that's what you're referring to.

3 Q. Let's take a change on the magnitude of some of
4 the things you suggested where you would -- might
5 change a function so that instead of doing it in a
6 register, you're doing it through pins.

7 A. Uh-huh.

8 Q. If a manufacturer came to you and said, well,
9 we're going to do it differently now, you would want to
10 see some kind of simulation to make sure that works; is
11 that right?

12 A. Well, there were -- yeah. Yes. We would want
13 to see some sort of verification that it worked.

14 Q. And if you don't have an actual part, you might
15 simulate it on a computer; right?

16 A. Yes.

17 Q. And if you do have an actual part, you would
18 test that; right?

19 A. Assuming you had a memory controller that
20 interfaced to it, yes.

21 Q. Okay. So if you had a memory controller that
22 interfaced to it, you would test it to make sure it
23 works; right?

24 A. Yes.

25 Q. Did you do any kind of simulations for any of

1 Q. And the way to implement that would be to
2 actually add pins to the DRAM and use those, post
3 those?

4 A. That's one way of doing it, yes.

5 Q. And if there were no available pins on the
6 DRAM, you'd have to add pins to the package; right?

7 A. Or multiplex existing pins.

8 Q. Well, one way to do it would be to add
9 dedicated pins?

10 A. One way to do it.

11 Q. And if there are no connected pins, then you're
12 going to have to add pins to the package?

13 A. You don't necessarily have to. You could use
14 existing pins and multiplex them.

15 Q. And I'm asking you in the situation where
16 you're not multiplexing them.

17 A. Yes. You would have to add them in, yes.

18 Q. For CAS latency, how many values are specified
19 by the JEDEC DRAM/SDRAM standard?

20 A. Oh, boy.

21 Q. Is it three?

22 A. It's either one, two or three. It's -- it's
23 either one or two bits, which is up to four states.

24 Q. So --

25 A. I don't know for certain. I don't have the

1 spec memorized.

2 Q. If it's two bits, then you would need two?

3 A. Two pins, that's correct.

4 Q. Two pins. Okay.

5 Let me call up RX-2100-13.

6 Your Honor, may I approach?

7 JUDGE McGUIRE: Yes.

8 BY MR. GATES:

9 Q. Now, Mr. 5nstmSzin, u wo've seyouMa

1 MR. GATES: Your Honor, he's just told us that
2 he's familiar with these types of data sheets, that
3 he's familiar with these types of diagrams and that he
4 would be able to read the pinout diagram that's there
5 on the page.

6 JUDGE McGUIRE: Overruled. If he can answer
7 that, I mean, he's gone through quite a bit here, so if
8 he can answer the question, he can go ahead.

9 MR. GATES: Thank you, Your Honor.

10 THE WITNESS: The question is how many
11 no-connect pins are available on the x16 version of
12 this part. And I believe the answer is one. If I can
13 read this properly.

14 BY MR. GATES:

15 Q. So under the alternative that we were just
16 discussing, in order to implement that alternative on
17 this particular DRAM, you would have to add a pin;
18 right?

19 A. Yes.

20 Q. Now, you also talked about a multiplexing
21 option.

22 So we understand that, you would send data
23 over at reset that would specify the CAS latency;
24 right?

25 A. Not necessarily send. Pins could be allocated

1 at reset to be again pulled up to a high voltage or
2 pulled down to a low voltage through a high value
3 resistor so that they wouldn't affect normal operation,
4 but that value could still be read by some internal
5 logic at reset time, as you suggested.

6 Q. And so that information about what the CAS
7 latency is would be received by the DRAM at reset?

8 A. Correct.

9 Q. And in order for that to operate, that
10 information would have to be stored somewhere in the
11 DRAM; is that right?

12 A. Correct.

13 Q. And you would store that in the register?

14 A. Correct.

15 Q. You also talked about various options,
16 alternatives for programmable burst length; right?

17 A. Yes.

18 Q. And one of the options that you talked about
19 was pin strapping again; right?

20 A. Yeah.

21 Q. And in order to implement that without
22 multiplexing you would have to have a dedicated pin, at
23 least one?

24 A. Yes.

25 Q. So if you wanted to have two burst length

1 options you'd have to have one pin; is that right?

2 A. Correct.

3 Q. If you wanted to have three burst length
4 options you'd have to have two pins; is that right?

5 A. Yes.

6 Q. So if you implemented that alternative with --
7 the pin strapping alternative that we discussed with
8 CAS latency, you would have to add two or three pins;
9 right?

10 A. I think you're combining both work-arounds
11 together.

12 Q. I am.

13 A. And yes, so you would have to add pins or use
14 more multiplexed pins as you added more features.

15 Q. And not only would you have to add pins on the
16 DRAM, but you might have to add pins on the DIMM as
17 well if there were not one -- DIMM pins available; is
18 that right?

19 A. That -- in one case, yes.

20 Q. And then you'd have to have pins at either the
21 memory controller or some other controller; is that
22 right?

23 A. Yes.

24 Q. And pins add cost?

25 A. Yes.

1 Q. The more pins, the more cost?

2 A. Correct.

3 Q. I think you talked earlier about when you were
4 at Apple -- well, I'm sorry. Let me back up. I've
5 missed something on programmable burst length. I
6 apologize.

7 You also talked about multiplexing with -- in
8 order to program the burst length; right?

9 A. Yes.

10 Q. And the same as we discussed with CAS latency,
11 the DRAM would receive that information at reset?

12 A. Yes.

13 Q. And that information would have to be stored in
14 the DRAM as to what the burst length is; right?

15 A. Yes.

16 Q. And that would have to be stored in a
17 register?

18 A. Sure. Yes.

19 Q. I'm sorry. Now we'll go to Apple.

20 You discussed earlier with Mr. Davis when you
21 were at Apple that Apple was using EDO technology;
22 right?

23 A. I believe so, yes.

24 Q. And EDO is an asynchronous technology?

25 A. That's correct.

1 1996-1997 time frame, wasn't it your opinion in that
2 time frame that synchronous technology had more
3 headroom than asynchronous technology?

4 A. Yes.

5 Q. And why was that?

6 A. The -- well, the industry was pushing that way
7 to start with. DRAM manufacturers were telling us
8 that's where they were going. They were obsoleting
9 their synchronous technology. In other areas of
10 computer systems, not necessarily DRAMs, but other
11 areas, synchronous technology had overtaken long ago.
12 Semi clocks with data was an established way for many,
13 many years and getting DRAMs onto that technology was
14 clearly the right thing to do.

1 higher, those sorts of things will dominate and cause
2 errors.

3 Q. Now, you're also familiar, aren't you, with a
4 technology developed by Kentron called QBM?

5 A. Top-level familiarity, yes.

6 Q. You've seen presentations?

7 A. Yes, I've seen presentations from Kentron.

8 Q. Did you do a preliminary evaluation of that
9 technology?

10 A. Very preliminary. Basically I saw their
11 presentation.

12 Q. Okay. And you understood that their technology
13 is basically interleaving banks on a DIMM, interleaving
14 memory banks on a DIMM?

15 A. Yes.

16 Q. And they use a FET switch, F-E-T switch, to do
17 that?

18 A. Yes. That's my understanding.

19 Q. And your preliminary evaluation of that
20 technology was that it would have signal integrity
21 problems?

22 A. Yes.

23 Q. Can you explain to me why you came to that
24 conclusion?

25 A. My first exposure to the Kentron stuff was in

1 the context of using signal data rate DRAMs using the
2 FET switch that you described to multiplex two devices
3 at twice the data rate on one wire. And our immediate
4 concern was that trying to use the signaling technology
5 or I/O specification of the synchronous data --
6 synchronous DRAM to drive twice as fast wouldn't work
7 very well. Given that it was having problems working
8 very well at its low speed, doubling it didn't seem to
9 be the right thing to do.

10 Q. You're familiar with a term called a burst
11 terminate command?

12 A. Yes.

13 Q. Can you explain what that is, what is a burst
14 terminate command?

15 A. This is a command used in synchronous DRAM
16 technologies, including DDR, where a burst command is a
17 command that says I want to do an operation with a long
18 sequence of data, four or eight cycles worth of data,
19 and the burst terminate command is I believe used after
20 a burst command has started to stop that command before
21 it completes.

22 Q. And is another name for that a burst interrupt
23 command?

24 A. I believe so, yes.

25 Q. And so, for example, when we're talking about

1 programmable burst length, that's a burst of data, is
2 that the same type of data that -- burst of data that
3 we're talking about with the burst terminate command?

4 A. Yes.

5 Q. Let me show you a document that's RX-1388.

6 May I approach, Your Honor?

7 JUDGE MCGUIRE: You may.

8 BY MR. GATES:

9 Q. Mr. Polzin, I've handed you a document that is
10 an e-mail from Jim Keller to you?

11 A. Yep.

12 Q. Attaching something from a future DRAM task
13 force. Do you see that?

14 A. Okay.

15 Q. Okay. I just want to ask you about something
16 on a particular page, and that is page 6.

17 If you look down near the bottom of the page,
18 there is a bolded text that says "No read or write
19 burst interrupt commands."

20 Do you see that?

21 A. Yes.

22 Q. And then it says -- I think it should be "at,"
23 but it says, "A high data rates burst interrupt
24 commands are of less value and are more difficult to
25 engineer."

1 Do you see that?

2 A. Yes.

3 Q. Okay. Do you agree with that statement, based
4 on your experience?

5 A. Yeah, this is referring to the details of a
6 DRAM design of which I'm not expert in. But it sounds
7 plausible to me.

8 Q. Well, let me ask you it this way.

9 Isn't your understanding that the use of a
10 burst interrupt command can waste command bandwidth?

11 A. It could be used in a manner that would waste
12 bandwidth, yes.

13 Q. And isn't it your understanding that the use of
14 a burst terminate command can lead to a less efficient
15 system overall?

16 A. Yes.

17 Q. And why is that?

18 A. The -- my understanding of the issue is that to
19 use the burst terminate command you cannot start an
20 operation until all of the data has been received in
21 case there was a terminate command halfway in between.
22 You have to wait for the entire thing to happen to make
23 sure no terminate commands happened while you were
24 receiving the data and then actually perform the
25 operation inside the chip.

1 Q. Have you heard of the term "pipelining"?

2 A. Yes.

3 Q. And are you familiar with what that term
4 means?

5 A. Yes.

6 Q. What is your understanding of pipelining?

7 A. Pipelining means that small parts of -- let me
8 back up.

9 Pipelining is that an operation can be broken
10 into small parts and executed on a series of data in
11 successive stages, so in cycle one the first data will
12 get operated on the first part of the operation, cycle
13 two the second, but immediately following the next data
14 can be working on the first operation.

15 It's better if I could draw it if you really
16 want me to, but --

17 Q. Just a high-level explanation is fine.

18 A. Yes.

19 Q. If I understand the concept, pipelining is used
20 to increase the efficiency of microprocessors?

21 A. Yes.

22 Q. And for synchronous DRAMs, those are pipelined;
23 is that right?

24 A. Yes.

25 Q. And is it your understanding that the use of a

1 burst terminate command can mess up your pipelining?

2 A. Yes.

3 Q. You were talking about earlier with Mr. Davis
4 about some alternatives for dual-edge clocking. Do you
5 remember that?

6 A. Yes.

7 Q. And you were concerned when you were looking at
8 these alternatives in 2000 that to implement them it
9 would mess up things because you were in the middle of
10 a product launch; right?

11 A. Yes.

12 Q. Okay. And that was in the middle of the
13 product launch for DDR; right?

14 A. Correct.

15 Q. And you had been designing things for DDR in
16 1999 prior to 2000; right?

17 A. Yes.

18 Q. And so in fact the decision to go to DDR was
19 sometime in early 1999?

20 A. Correct.

21 Q. So if you had known prior to that time that
22 Rambus' patents would be infringed by the use of
23 dual-edged clocking and one of your alternatives were
24 adopted in the JEDEC standard, would you have had those
25 same concerns?

1 MR. DAVIS: Objection. Calls for speculation.

2 JUDGE McGUIRE: Sustained.

3 BY MR. GATES:

4 Q. You were concerned because you were in the
5 middle of a product launch; right?

6 A. I was concerned about what?

7 Q. You were concerned about implementing these
8 alternatives because you were in the middle of a
9 product launch?

10 A. Correct. Yes.

11 Q. Would you have had those same concerns if you
12 weren't in the middle of a product launch?

13 MR. DAVIS: Objection. It's the same
14 question.

15 JUDGE McGUIRE: Sustained.

16 MR. GATES: Thank you, Your Honor.

17 BY MR. GATES:

18 Q. What was it about the fact that you were in
19 the middle of a product launch that raised these
20 concerns?

21 A. We had enabled a number of DRAM manufacturers,
22 motherboard manufacturers, our own manufacturing of our
23 chipset assuming one standard, one spec. To change
24 that, we'd have to stop all that production,
25 reengineer, redeploy, start production again. It

1 would be detriment -- seriously detrimental to our
2 business.

3 Q. And when you decided to go to DDR instead of
4 Rambus, did you have to reenable your vendors and
5 redesign your chipset at that point?

6 A. We, based on -- well, from what basis?

7 Q. Well, in early 1999, you decided that instead
8 of going with Rambus to go with DDR; right?

9 A. Yes.

10 Q. Okay. And I'm wondering whether or not you had
11 to go through the same types of things that you were
12 concerned about when you looked at the Rambus patents
13 in 2000 in the middle of a product launch.

14 A. Well, given that we were starting essentially
15 from, you know, a certain base and starting forward,
16 you know, we could have done pretty much anything, if
17 you want to call it that.

18 Q. So if I understand what you're saying, if
19 you're starting at a zero base you can implement these
20 alternatives?

21 A. Yes.

22 Q. Costlessly?

23 A. Yes.

24 Q. So for example, you're working on chipsets
25 for -- you worked on chipsets for DDR-II; right?

1 A. No.

2 Q. You worked on a processor that's going to be
3 compatible with DDR-II?

4 A. Personally, no, I'm not working on that, that
5 product.

6 Q. You're not working on the K8?

7 A. No. I'm the system guy. I'm not the
8 microprocessor guy. I'm sorry. I'm not following your
9 question here.

10 Q. So you're the system guy?

11 A. Yes.

12 Q. For K8?

13 A. Yes.

14 Q. So your job is to make sure that there's an
15 infrastructure for the K8?

16 A. Correct. Yes.

17 Q. And you based your determination of using
18 DDR-II on what you knew was going to be or what you
19 thought was going to be available in the marketplace;
20 right?

21 A. So DDR-II is a future technology. We don't
22 have any products out for DDR-II. I'm --

23 Q. You're developing products for DDR-II; right?

24 A. Yes.

25 Q. And you're developing -- you're working on

1 developing an infrastructure for DDR-II; right?

2 A. We're starting, yes.

3 Q. When did that start?

4 A. It technically hasn't started. We haven't
5 really started that yet.

6 Q. It hasn't started?

7 A. Yeah.

8 Q. So if the DDR-II specification were to adopt
9 some of your alternatives, that would be starting at
10 ground zero as far as AMD is concerned?

11 JUDGE MCGUIRE: All right, Mr. Gates. I'm not
12 sure where you're going at this point. I'm not sure
13 you're still within the scope and I'm going to ask you
14 to change your line.

15 MR. GATES: Okay. Thank you, Your Honor.

16 BY MR. GATES:

17 Q. You said earlier that you had discussions with
18 some DRAM manufacturers about the alternatives that you
19 had come up with?

20 A. Yes.

21 Q. And those discussions happened in 2000; is that
22 right?

23 A. Yes.

24 Q. And you're aware of the DDR-II specification;
25 right?

1 A. I believe it was winding its way through JEDEC
2 at the time, yes.

3 Q. So the DDR-II specification at the time in
4 2000 was winding its way through the -- through JEDEC;
5 right?

6 A. Yes.

7 Q. Okay. The DDR-II specification uses
8 programmable CAS latency; right?

9 A. That's my understanding, yes.

10 Q. And it uses --

11 MR. DAVIS: Objection, Your Honor. This is
12 beyond the scope. We didn't talk about DDR-II at all
13 in direct.

14 JUDGE MCGUIRE: Sustained.

15 MR. GATES: Your Honor, I'd like a little
16 leeway here only because I don't want to have to call
17 this witness back on direct in our case. I have one
18 small line of questioning on DDR-II and that will be
19 it. I thought that was our understanding with

1 MR. GATES: All right. Thank you, Your Honor.

2 BY MR. GATES:

3 Q. Okay. So it's your understanding that the
4 DDR-II specification specifies programmable CAS
5 latency; right?

6 A. Yes.

7 Q. And it specifies using a mode register to set
8 the CAS latency?

9 A. I believe so, yes.

10 Q. And the DDR-II specification is not completed
11 yet; right?

12 A. The JEDEC specification is pretty darn close to
13 completion, but yeah, I guess it's never quite at the
14 end.

15 Q. And it's now 2003, it's not quite at the end;
16 right?

17 A. Yeah. Yeah.

18 Q. And the DDR-II specification specifies
19 programmable burst length; is that right?

20 A. I believe so.

21 Q. And were you aware of the fact that at one
22 point it was considered to just use a burst length of
23 four in the specification?

24 A. I'm not aware of that. But I have no knowledge
25 one way or another.

1 A. Yes.

2 MR. GATES: No further questions, Your Honor.

3 JUDGE McGUIRE: All right. Thank you,

4 Mr. Gates.

5 Mr. Davis?

6 You may want to mark this sheet as DX-31,

7 Mr. Gates, if you would.

8 MR. GATES: Thank you, Your Honor.

9 (DX Exhibit Number 31 was marked for
10 identification.)

11 JUDGE McGUIRE: All right. Mr. Davis?

12 MR. DAVIS: Thank you, Your Honor.

13 JUDGE McGUIRE: Redirect.

14 REDIRECT EXAMINATION

15 BY MR. DAVIS:

16 Q. Mr. Polzin, would you describe what "front-side
17 bus" means on that (indicating).

18 A. "Front-side bus" is an industry term that

1 Q. Were you involved in the decision to do that?

2 A. Yes.

3 Q. Why was AMD moving from 200 megahertz to
4 266 megahertz on the front-side bus?

5 A. To increase performance of the system.

6 Q. Why was that important?

7 A. To remain competitive in the marketplace.

8 Q. And why were you moving -- were you involved in
9 the decision to move from 266 megahertz to
10 333 megahertz?

11 A. Yes.

12 Q. And do you know why that was done?

13 A. Similar reasons, always to get more and more
14 performance for our customers.

15 Q. And so it's the same reason -- were you
16 involved in the decision to move from the 333 to the
17 400?

18 A. Yes.

19 Q. Was that also the same --

20 A. Same, same reason.

21 Q. Did that improvement in performance from the
22 200 to 266 and the 266 to 333 and the 333 to 400, did
23 that have any implication for your chipset or
24 motherboard suppliers and partners?

25 A. Yes. Our chipset partners needed to design

1 faster circuitry in their chipsets and our motherboard
2 partners needed to adhere to stricter design rules in
3 their manufacture of their motherboards.

4 Q. Was there any benefit to the chipset
5 manufacturers or the motherboard manufacturers if they
6 did design to the higher speeds?

7 MR. GATES: Object, Your Honor, on foundation
8 grounds.

9 JUDGE McGUIRE: Sustained.

10 BY MR. DAVIS:

11 Q. Do you have any understanding of why your
12 chipset partners manufactured your chipsets for your
13 front-side buses?

14 A. Why they manufacture our chipsets for the
15 faster front-side buses? That's the question? They
16 want to keep up with the latest technology. They can
17 get higher prices for more advanced chipsets. A
18 chipset that supports DDR 333, for example, is worth
19 more than a chipset that supports DDR 200.

20 Q. What's your basis for saying that?

21 A. Well, when we were in -- when AMD was in the
22 chipset business and selling our DDR chipset, our DDR
23 chipset we could sell for more money than our single
24 data rate chipset. And we closely track our
25 third-party chipset prices and motherboard prices as

1 part of our ongoing business to make sure our
2 microprocessors match up in the right marketplace.

3 Q. Now, do you have an understanding then of
4 why -- of the benefit to the chipset manufacturers and
5 the motherboard manufacturers of increasing -- of
6 working with the increased front-side bus?

7 A. Yes. The benefit is they can then use
8 faster -- better-performing microprocessors and deliver
9 a better-performing system to customers and therefore
10 be able to keep up with the competition and retain
11 their prices, retain their market share in the given
12 price points that they want to participate.

13 Q. Thank you.

14 Now, earlier you were discussing synchronous
15 DRAM with Mr. Gates. Do you remember that?

16 A. Yes.

17 Q. And the difference between synchronous DRAM and
18 asynchronous DRAM?

19 A. Yes.

20 Q. When you were referring to the term
21 "synchronous DRAM," were you -- sorry.

22 When you were referring to the term
23 "asynchronous DRAM," did that include the DRAM where
24 there was a data strobe or some other strobe?

25 A. No. Asynchronous DRAM -- I'll just use the

1 industry terms. Asynchronous DRAM I refer to as the
2 fast page mode and the EDO versions of DRAMs.
3 Synchronous DRAM -- synchronous technology started in
4 my mind with SDRAMs.

5 Q. So when you were talking about synchronous
6 DRAM and asynchronous DRAM, a DRAM that was
7 asynchronous except that it had a data strobe, that
8 would come under your definition of synchronous or
9 asynchronous?

10 A. I believe that that would have to fall under
11 asynchronous. My definition is synchronous starts out
12 with the synchronous DRAM specs. The PC66, the PC100,
13 that's the start of synchronous technology in my -- in
14 what I was referring to.

15 Q. Okay. When you were referring to Kentron and
16 the signal integrity problems that you observed when
17 you saw that presentation --

18 A. Yes.

19 Q. -- what was the cause of those signal
20 integrity problems?

21 A. The basic cause is that they were trying to use
22 a driver technology that was designed to run at a
23 hundred megahertz speeds, for example, using PC100 DRAM
24 parts and run that twice as fast using the same I/O
25 signaling technology.

1 Q. Did you have any understanding at the time
2 whether that I/O signaling technology could have been
3 improved so that this Kentron solution could work?

4 MR. GATES: Objection, Your Honor. It lacks
5 foundation and it calls for opinion.

6 JUDGE MCGUIRE: It calls for opinion?

7 MR. DAVIS: This is exactly the topic that he
8 was referring to earlier.

9 MR. GATES: Well, you're --

10 JUDGE MCGUIRE: Well, I'm going to -- that's
11 overruled. I'll let him, to the extent he had an
12 understanding at the time, I'll let him answer.

13 THE WITNESS: I'm sorry. Could you repeat the
14 question.

15 BY MR. DAVIS:

16 Q. Yes.

17 Did you have an understanding at the time of
18 whether an improved signaling technology would have
19 allowed the Kentron system to work whereas the
20 signaling technology then on SDRAMs wouldn't allow it
21 to work?

22 A. I believe in the time frame that we were
23 discussing, this particular Kentron technology, they
24 probably could have invented a new signaling
25 technology, certainly the one used for DDR, for

1 example, but it would have required a brand-new DRAM,
2 manufactured DRAM I/O cell design, and I don't believe
3 that's the basis of their technology.

4 The basis of their technology from my
5 understanding is they use, quote, old DRAMs and you run
6 them twice as fast through your magic FET switches, and
7 it just doesn't work that way.

8 MR. DAVIS: Okay. Thank you, Your Honor.

9 JUDGE McGUIRE: Mr. Gates, any recross?

10 MR. GATES: No, Your Honor.

11 JUDGE McGUIRE: All right, sir. Thank you very
12 much for your testimony today. You're excused from
13 these proceedings.

14 Let me ask from I guess complaint counsel,
15 what's on tap for Tuesday? I know we talked the other
16 day that that was still somewhat up in the air. Maybe
17 we can get an update.

18 MR. OLIVER: Yes. Excuse me. We had had a
19 scheduling problem tomorrow. We don't have a witness
20 for tomorrow. We're hoping to be able to read in some
21 depositions, and I think, based on what you said at the
22 beginning of this afternoon's session, I think the
23 question I have for you is, after you issue your order
24 tomorrow, would we then be in a position to play and
25 read in Dr. Oh's testimony?

1 JUDGE McGUIRE: Yes, you would.

2 MR. OLIVER: Then I think I guess I would
3 suggest if we could await your order and then do that
4 tomorrow.

5 JUDGE McGUIRE: Okay. You mean -- I'll plan on
6 having it out in the morning. Do you want to postpone
7 then the hearing for an hour or so after I issue the
8 order? What are you suggesting that we do?

9 MR. OLIVER: I think that would work very well
10 for our side.

11 (Pause in the proceedings.)

12 We were discussing perhaps starting at 11:00 in
13 the morning.

14 JUDGE McGUIRE: That's fine. That's fine.
15 We'll start at that time then.

16 Anything else?

17 MR. GATES: No, Your Honor. I just had one
18 exhibit that I needed to move in.

19 JUDGE McGUIRE: Go ahead, Mr. Gates.

20 MR. GATES: I'd like to move in RX-2100-13.
21 That's the spec sheet.

22 JUDGE McGUIRE: Any objection?

23 MR. DAVIS: No objection, Your Honor.

24 JUDGE McGUIRE: So entered.

25 (RX Exhibit Number 2100-13 was admitted into

1 evidence.)

2 JUDGE McGUIRE: Okay. Very good. We'll
3 convene tomorrow morning.

4 And Mr. Stone, I'm hopeful you'll keep us
5 informed on the pending birth of your colleague's child
6 immediately.

7 MR. STONE: I certainly will, Your Honor.
8 Thank you for inquiring.

9 JUDGE McGUIRE: Thank you very much. Hearing
10 in recess.

11 (Time noted: 3:44 p.m.)

12

13

14

15

16

17

18

19

20

21

22

23

24

25

