UNITED STATES INTERNATIONAL TRADE COMMISSION

WASHINGTON, D.C.

Before the Honorable Thomas B. Pender

Administrative Law Judge

- x

In the Matter of

CERTAIN CONSUMER ELECTRONICS AND DISPLAY DEVICES WITH GRAPHICS PROCESSING AND GRAPHICS PROCESSING UNITS THEREIN	Investigation No. 337-TA-932

- x

CONTAINS ADVANCED MICRO DEVICES, INC. CONFIDENTIAL BUSINESS INFORMATION SUBJECT TO PROTECTIVE ORDER - VIDEOTAPED DEPOSITION OF ANDREW E. GRUBER, on behalf of Advanced Micro Devices, Inc. and Individually

Tuesday, March 17, 2015, 9:14 a.m.

Robins, Kaplan, Miller & Ciresi L.L.P.

800 Boylston Street

Boston, Massachusetts

Reported by:

MICHAEL O'CONNOR, RPR, CRR, CBC, CCP

Job No. 91486

Tuesday, March 17, 2015

9:14 a.m.

Videotaped Deposition of ANDREW E. GRUBER, individually, and as Corporate Designee on behalf of Advanced Micro Devices, Inc. before Michael D. O'Connor, a Registered Professional Reporter and Certified Realtime Reporter and Notary Public in the Commonwealth of Massachusetts.

APPEARANCES:

ROBINS KAPLAN MILLER & CIRESI
Attorney for Advanced Micro Devices:
800 LaSalle Avenue
Minneapolis, Minnesota 55402
BY: AARON FAHRENKROG, ESQ.

- AND -

STERNE KESSLER GOLDSTEIN & FOX
Attorney for Advanced Micro Devices:
1100 New York Avenue, N.W.
Washington, DC, 20005
BY: JONATHAN TUMINARO, ESQ.

ORRICK HERRINGTON & SUTCLIFFE
Attorney for NVIDIA Corporation:
1152 15th Street, N.W.
Washington, DC, 20005
BY: CHRISTOPHER HIGGINS, ESQ.
I. NEEL CHATTERJEE, ESQ.

APPEARANCES (Cont'd):

KEKER & VAN NEST
Attorney for Qualcomm Incorporated:
633 Battery Street
San Francisco, California 94111
BY: EUGENE PAIGE, ESQ.

Also Present: Leann McKeon, Videographer

Also Present via speakerphone:
Daniel Leventhal, Esq.,
Norton Rose Fulbright (US), LLP

P R O C E E D I N G S

(Gruber Exhibit 1, Document entitled "Complainant NVIDIA Corporation's Notice of Deposition of Andrew Gruber," marked for identification)

(Gruber Exhibit 2, Document entitled "Application For Issuance of Subpoena Duces Tecum and Subpoena Ad Testificandum to Advanced Micro Devices, Inc.," marked for identification)

VIDEOGRAPHER: This is the start of tape labeled number one of the videotaped deposition of Andrew Gruber in the Matter of Certain Consumer Electronics and Display Devices With Graphics Processing and Graphics Processing Units Therein, Investigation No. 337-TA-932 in the United States International Trade Commission, Washington, D.C., before the honorable Thomas B. Pender. 09:13 09:13 09:13 09:13 09:13 09:13 09:13 09:13 09:13 09:14

This deposition is being held at 800 Boylston Street, Boston, Massachusetts, on March 17, 2015, at approximately 9:14 a.m. My name is Leann McKeon. I'm the legal video specialist from TSG Reporting, Inc. headquartered at 747 Third Avenue, New York, New York. The court 09:14 09:14 09:14 09:14 09:14

reporter is Michael O'Connor in association with TSG Reporting. Will counsel please introduce yourselves. 09:14

MR. HIGGINS: Chris Higgins from Orrick, Herrington & Sutcliffe, on behalf of Complainant, NVIDIA Corporation. And with me is Neel Chatterjee, also from Orrick. 09:14

MR. FAHRENKROG: Aaron Fahrenkrog from Robins Kaplan, LLP on behalf of non-party Advanced Micro Devices, Inc. 09:14

MR. TUMINARO: Jonathan Tuminaro from Sterne, Kessler, Goldstein & Fox on behalf of non-party AMD. 09:14

MR. PAIGE: Eugene Paige of Keker & Van Nest on behalf of Respondent Qualcomm and the witness in his personal capacity. 09:14

A N D R E W E. G R U B E R,

called as a witness, having been duly sworn by a Notary Public, was examined and testified as follows:

EXAMINATION BY

BY MR. HIGGINS: 09:15

Q. Good morning, Mr. Gruber. 09:15

A. Good morning. 09:15

Q. You have been handed what's been 09:15

marked as Exhibits 1 and 2. Exhibit 1 is "Complainant NVIDIA Corporation's Notice of Deposition of Andrew Gruber." Exhibit 2 is "Application For Issuance of Subpoena Duces Tecum and Subpoena Ad Testificandum to Advanced Micro Devices, Inc." 09:15

Do you see those? 09:15

A. Yes. 09:15

Q. Are you here to testify on behalf of both the deposition notice and the subpoena to AMD? 09:15

A. Yes. 09:15

Q. When was the R400 built? 09:15

A. The R400 project was remained the R500 project, and eventually involved into the R600 project, which was the chip that finally made it to market. The same database also evolved into the internal code name Xenos chip, which was in the Microsoft Xbox 360 project. 09:15 09:16

That Xbox project shipped in late 2005. The R600 shipped after that. I believe it was in 2007. 09:16 09:16

Q. Was a chip ever built that was based solely on the R400 project? 09:16

A. No. By the time that the Xbox 09:16

shipped, there were certain evolution of the database. So there were changes. But the internals were largely based on the R400 unchanged. 09:16

Q. Was a chip ever built that was based solely on the R500 project? 09:16

A. No. The R500 project evolved into the R600 project. There were substantial changes from the R500 to the R600. But still a lot of the engineering in the R400 showed up in the R600. 09:17

Q. What were the substantial changes from the R500 to the R600? 09:17

A. The R600 still had the same unified shader idea as the R500. In the R600, the instruction set changed from a pure vector model, what was called a Vec 4 model, into four parallel scaler instructions. 09:17

The data flow in the R500 to R600 was very similar. The main change was in the, as I say, the instruction set and the approach to instruction execution. 09:18

Q. Did the R600 use the same sequencer as the R500? 09:18

MR. PAIGE: Vague and ambiguous. 09:18

A. The sequencer between the R500 and the R600 had substantial changes. The R500 team was primarily from the Marlboro office, and that was the same team that had just continued on as the R400 changed into the R500. 09:18

The R600 was more of a combined effort in between the East Coast office and the West Coast office. So the sequencer aspects changed substantially between the R500 and the R600. 09:18 09:19

Q. Why did the sequencer change substantially between the R500 and the R600? 09:19

A. There were just different people working on it. I was not involved as much in the R600 as in the R500. So I can't talk about a lot of the specific changes in the sequencer, but I do believe there were changes from the R500 to the R600. 09:19

Q. When was the R400 project started? 09:19

A. The R400 project started taking shape in late 2000. That's when the initial concepts were kind of first come up with and discussions started about it within ATI. 09:19 09:20

Q. Why was the R400 project started? 09:20

A. Well, we needed a follow-on chip. 09:20

The way that ATI worked at that time, we had a leapfrogging design system. So it was a group on the East Coast that was doing the odd sequence -- sorry -- the even sequence of chips. So the East Coast team did the R200 and then the West Coast team would do the R300, and the East Coast team would do the R400. 09:20

This is because the market needs were such that the need for a new chip was shorter than the design time needed for a new chip. So we had these leapfrogging teams to be able to meet the market windows. 09:20 09:21 09:21 09:21

So the R400 was intended as our new chip after the -- well, after the West Coast R300. But, from the East Coast point of view, after the R200. We needed both increased performance as well as we were looking at supporting next-generation APIs. 09:21 09:21 09:21 09:21 09:21

Q. Did the West Coast team work on the next odd numbered R500? 09:21 09:21

A. That was originally the goal. As schedules changed and the R400 took longer than anticipated to meet the development effort, what happened was that the -- we met that marketing window rather than with the R400 with a revision 09:21 09:22 09:22 09:22

of the R300, which we called the R420. 09:22

The R400 was renamed the R500, and 09:22
aimed at the next market window. As it turned 09:22
out, the R500 slipped as well. And so the West 09:22
Coast team did wind up producing a chip that went 09:22
into the R500 marketing window, and that was 09:22
called the R580. 09:22

Q. When was the 420 chip released? 09:22

A. I am not sure. I'd have to go back 09:22
into records and get back to you on that. 09:23

Q. Do you have a general idea of the 09:23
time frame when it may have been released? 09:23

A. The R420. It was probably released 09:23
in the 2004 time frame. I couldn't... 09:23

Q. Do you know when the 580 was 09:23
released? 09:23

A. It probably would have been released 09:23
about a year later, maybe 18 months later. 09:23

Q. Why did the R400 take longer than 09:23
anticipated to meet the development effort? 09:23

A. The R400 was a very aggressive chip 09:23
in a number of ways. The unified shader idea was 09:23
certainly new, and there was a lot of engineering 09:24
associated with that. But there were a lot of 09:24
other changes involved in the R400 as well. 09:24

There were compression changes in how we dealt with memory that took a lot longer than we thought. There was negotiations with Microsoft over how the API would shape up. 09:24

So there were back-and-forth discussions about that. So sometimes engineering just takes longer than you anticipate. 09:24

Q. Where was the work on the R400 done, in what location? 09:24

A. The primary work on the R400 was done in the Marlboro office of ATI. There was some work done as well in the Orlando office of ATI. The Orlando team came on board with the R400 a little later than the Marlboro team did, because the Orlando office was also involved in the R300 series of products. 09:25

So while we had a leapfrogging effort, the Orlando team was actually on both of those efforts. 09:25

Q. Which parts of the R400 did the Marlboro office work on? 09:25

A. As I said, the Marlboro office worked on virtually all of it. It worked on the sequencer, on the shader pipe, on what we called the render back-end, which NVIDIA called the ROP 09:26

unit, the texture system, on the command processor, although some of the command processor was done in Orlando. 09:26

Most of the vertex processing after shading was done in the Orlando office, that it's a clipping and viewport transform. 09:26

Q. What aspects of the engineering associated with the unified shader contributed to the delay in the R400 project? 09:26

MR. FAHRENKROG: Object to the form of the question. 09:26

MR. PAIGE: Join in the objection. 09:26

A. A lot of the data flow was challenging. How do you get the data from the vertex shader back into the pixel shader. Dealing with some of the shared resources for both vertices and pixels in a manner such that you didn't get deadlock and that you got adequate performance, that one didn't overwhelm the other was also challenging. 09:27

Q. In the R400, what were the shared resources between vertices and pixels? 09:27

A. The shader itself, that is the execution unit and the sequencer logic, were shared. The general purpose registers were 09:28

shared also on a slowly moving dynamic basis where there were two different pools, but the pool sizes could change. 09:28

Q. What were the issues with the general purpose register in the unified shader design? 09:28

MR. PAIGE: Object to the form of the question. 09:28

MR. FAHRENKROG: Objection. 09:28

A. Allowing the shader that had the most work to do to get more of the resources than the other shader, while at the same time having a reasonable allocation mechanism for those. We like to allocate in kind of a FIFO order, because that's easy to control, but it's difficult to do that when you have two different shader programs that take radically different amounts of time, that take unpredictable amounts of time. 09:29

Q. Were these issues with the general purpose registers ever resolved in the R400 project? 09:29

MR. PAIGE: Vague and ambiguous. 09:29

A. Yes. The R400 had a register allocation scheme that was functional and that worked well, and that was actually carried forward into the Xbox chip. I don't think there 09:30

were -- I can't think of a difference between the R400 design and what was produced in the Xbox, in that area. 09:30

Q. When was the design of the general purpose registers in the R400 finalized? 09:30

A. The general purpose registers were finalized fairly early on. I would say prior to the end of 2001, the overall scheme was established and really didn't change much after that. 09:30

Q. What was your role in the R400 project? 09:31

A. I was the head architect of the Marlboro site. Steve Ryan was under me, and he was nominally specifically in charge of the R400. As the project evolved, I took over more of that role from Steve. 09:31

Q. Who else worked on the R400 project? 09:31

A. Well, there was a very large team. I mean, by the end, it was certainly over 100 people. 09:31

Q. Who was on the initial team that started work on the R400 project? 09:31

A. Initially, there were just a small group of architects working. Steve Morein, 09:31

Laurent Lefebvre, Andy Skende were people who worked on it. 09:32

Q. When were additional people added to the R400 team? 09:32

A. Throughout 2001, as people came off of the final work on the R200, there were more and more people involved. I would say during the first half of 2001, it was a fairly small architectural team. 09:32

Q. When did the R400 project end? 09:32

A. I believe it ended -- well, it didn't end. We evolved it into the R500. That is, we realized we wouldn't make that market cycle in early 2003. 09:32

Q. When did work on the R500 project begin? 09:33

A. Immediately after. There was no work stoppage. It was just kind of a redirection. We were late on the R400, and we didn't want to be late on the R500 as well. So the nature of the work did not change at all between R400 and R500. 09:33

The API it was aimed at was the same API. 09:33

Q. When did work on the Xenos chip begin? 09:33

A. That also began in early 2003. I think that the Xenos work began shortly after the transition to R500. 09:33

Q. Did the same teamwork on the Xenos chip and the R500 project? 09:33

A. There was a lot of overlap, but there were also separate people working on the Xenos chip that were working on the R500. Most of the Xenos work wound up being focused in the Orlando site, while most of the Marlboro site continued work on the R500. 09:34

But the databases were so similar, that there was a lot of back and forth between the two. 09:34

Q. Which project did you work on? 09:34

A. I was primarily on the R500. 09:34

Q. Why was a different team assigned to work on the Xenos? 09:34

A. There were aspects of the Xenos that were unique to Xenos that weren't going to be part of the anticipated R500 chip. Xenos had a very different memory interface, for instance, there was a dedicated memory chip as opposed to a standard off-the-shelf memory chips in the Xenos. 09:35

There were also interfaces to the 09:35

host processor in the Xenos chip that were not anticipated to be part of the R500. And there were also just the general Microsoft interaction, as Microsoft was the project owner, that weren't a factor in the R500. 09:35

Q. Was Microsoft the project owner of the Xenos chip? 09:35

A. They were the overall project owner, yeah. 09:35

Q. What did they do as the overall project owner of the Xenos chip? 09:35

MR. FAHRENKROG: I'll insert an objection here. I caution you not to reveal any Microsoft confidential information. We've objected to that in response to your subpoena, and he's not here to talk about anything confidential to Microsoft here today. You're welcome to ask your questions about Xenos, and so forth. 09:35 09:36 09:36 09:36 09:36 09:36

Q. What did Microsoft do as the overall project owner of the Xenos? 09:36

A. They specified some of the interfaces and some of the scenarios that we had to support. 09:36

Q. Which interfaces did they specify? 09:36

A. There was an interface in between the 09:36

host and the Xenos chip that was Microsoft specific that would allow the host to directly write into Xenos memory. 09:36 09:36 09:36

Q. What were the specific scenarios that you had to support for Microsoft? 09:37 09:37

MR. FAHRENKROG: Object to the form of the question. 09:37 09:37

A. Well, for instance, there was -- I mentioned this dedicated memory associated with the part, and that had a certain size associated with it. 09:37 09:37 09:37 09:37

Microsoft was involved in specifying the size of that memory, and the size of that memory defined what screen resolutions would be supported in one pass versus multipass. 09:37 09:37 09:37

So, you know, whether given say 1280, or whatever standard high definition is, whether that would be a multipass. And that was primarily a cost issue. So that's why it fell to Microsoft. 09:38 09:38 09:38 09:38

The engineering of the chip was ATI, but a number of the specifications of the chip were Microsoft. 09:38 09:38 09:38

Q. So did Microsoft specify certain aspects of the chip design for the Xenos? 09:38 09:38

A. They specified certain capabilities of the chip. I wouldn't say Microsoft specified design aspects. 09:38

Q. What capabilities of the chip did Microsoft specify? 09:38

A. Well, as I mentioned, the ability for the host to directly write into the chip. The size of the external memory. They might have -- as I said, I wasn't on the Xenos project, but I expect that they had certain performance goals as well that had to be met. 09:38 09:39

Q. Who is the person or people at Microsoft that were involved with the Xenos project? 09:39

MR. FAHRENKROG: Object to the question as asking for confidential Microsoft information. 09:39

A. I don't recall offhand. 09:39

Q. Was the Xenos chip a single chip GPU implementation? 09:39

MR. FAHRENKROG: Object to the form of the question. 09:39

A. The GPU was a single chip. There was an additional dedicated memory part that held a tile buffer that was on the same -- was in the 09:40

same package, but it was not part of the same die.

Q. Were any rendering functions performed on a separate chip that had a tile buffer?

A. I believe it was purely a eDRAM part, an embedded DRAM part. I would have to go back to actual notes to check to see that there wasn't any logic functionality on that part.

I remember very early on there was some discussion about that, but I do not believe that was in the part.

Q. What notes would you have to go back and check?

A. Oh, I'd have to go through ATI proprietary information. It wouldn't be my personal notes, but I'd have to get proprietary information from ATI about Xenos, like the Xenos overall specification.

Q. Was the render back-end module on a separate chip?

A. Again, certainly the entire render back-end module would not have been on a separate chip. There may have been some aspect on that separate chip, but I can't testify one way or the

other. 09:42

Q. Why would certain aspects of the render back-end module be placed on a separate chip? 09:42

MR. FAHRENKROG: Object to the form of the question. Vague. 09:42

A. The reason why you would do that was to convert read-modify-write operations into pure write operations to preserve bandwidth. 09:42

Q. In general, what type of operations does the render back-end module perform? 09:42

MR. FAHRENKROG: Object to the form. Vague. 09:42

MR. PAIGE: Object to the form. 09:42

A. It generally performs blending operations and logic operations like and/or. 09:42

Q. What reason would blending operations be moved to a separate chip? 09:43

A. For the same reason to preserve memory bandwidth. If you were doing read-modify-write operations, could be converted to just a write-only operation. 09:43

Q. Did Microsoft request the use of a separate chip in the Xenos design? 09:43

A. I don't recall whether that aspect of 09:43

the design was originated with Microsoft or was suggested by ATI. 09:43

Q. Why didn't Microsoft use the R400 design? 09:43

MR. FAHRENKROG: Object to form. Foundation. 09:44

MR. PAIGE: Join in the objection. 09:44

A. We used the R400 as the basis for the Microsoft proposal. As we understood more about what Microsoft was looking for, there was, you know, design changes proposed to meet Microsoft-specific needs. 09:44

I don't know whether the specific change of the use of an embedded DRAM separate chip came from Microsoft or ATI. 09:44

Q. What parts of the R400 design did not meet Microsoft-specific needs? 09:44

MR. PAIGE: Object to the form. 09:45

A. Well, for one thing, Microsoft didn't want the graphics DRAM aspect of it, that is, as a separate memory pool. Microsoft wanted things to work out of a single system pool. 09:45

Once we lost the separate memory pool for graphics, that's when the eDRAM showed up as a way of at least providing some additional 09:45

memory bandwidth to meet GPU needs. 09:45
The actual shader system from R400 to 09:45
Xenos was pretty much identical. 09:46
Q. How did the R400 design use the 09:46
graphics DRAM? 09:46
A. The R400 generally rendered into 09:46
graphics DRAM as a separate memory pool, and it 09:46
had a high-speed bus into the graphics DRAM. It 09:46
used it purely as a memory. 09:46
There was a memory hub as well that 09:46
could direct rendering into a system memory at a 09:46
slower rate. 09:46
Q. Did the R400 have memory in addition 09:46
to graphics DRAM? 09:47
A. Again, the graphics DRAM was not part 09:47
of the R400. It was a separate chip, external. 09:47
It had internal memories and, as well, it had 09:47
access to system memory. 09:47
Q. What internal memory did the R400 09:47
use? 09:47
A. The R400 had various caches as well 09:47
as FIFOs and buffers that would join different 09:47
blocks. There were fairly small pieces of 09:47
memory. 09:47
Q. Were general purpose registers within 09:47

the R400 design? 09:47
A. Yes. 09:48
Q. Were those part of the various caches and -- 09:48
A. Yes. 09:48
Q. -- buffers? 09:48
A. Yes. 09:48
Q. Was the R400 ever taped out? 09:48
A. It was never taped out as the R400. The design evolved and was eventually taped out as either Xenos as well as the R600. 09:48
Q. Was the R500 ever taped out? 09:48
A. The R500 was never taped out. It evolved into the R600, which was taped out. 09:48
Q. Was the R600 based off the R500 or Xenos? 09:48
A. It was based off the R500. 09:49
Q. When did the R600 tape out? 09:49
A. Again, I'm not entirely sure. Tape out was probably 2005/2006 time frame, but it could have even been later than that. 09:49
Q. When was Xenos taped out? 09:49
A. Xenos shipped the end of 2005. Tapeout was probably late 2004, although it could have stretched into 2005. 09:49

Q. Was there ever any attempt to tape out the R400? 09:49

MR. PAIGE: Object to the form. 09:49

A. Not as the R400. Again, only in the evolved design that was called something else. 09:49

Q. Was the R600 a single-chip GPU implementation? 09:50

A. Yes. 09:50

Q. Was the R600 ATI's first single-chip GPU implementation with a unified shader? 09:50

A. Regardless of the capabilities of the memory in Xenos, I regard the Xenos chip as a single-chip GPU that had the unified shader. 09:50

Q. Did ATI tape out any single-chip GPUs with unified shaders prior to Xenos in late 2004 or early 2005? 09:50

MR. FAHRENKROG: Object to the form. 09:51

A. No. 09:51

(Gruber Exhibit 3, Article entitled "ATI R400 GPU Cancelled," marked for identification) 09:52

Q. You have been handed what's been marked as Exhibit 3. It's an article posted on neoseeker.com, entitled "ATI R400 GPU Cancelled." 09:52

A. Yes. 09:52

Q. About midway down on the page there's a box that says, "The R500 is the R400 which was canceled earlier this year." 09:52

Do you see that? 09:52

A. Yes. 09:52

Q. Is that an accurate statement? 09:52

MR. FAHRENKROG: Object to the form of the question. 09:52

A. The first sentence, yes, I believe that. It was just a name change. From an engineering point of view, it's pretty much the same. It took us a while to get all the documentation, even up to date. We continued on with a large mixture of things that said R400 and things that said R500. 09:52 09:53 09:53 09:53

Q. What do you mean by, It took a while to get all the documentation up -- 09:53

A. Well, there were various places in the documentation and in the code where you would have R400 in the same documentation and code that was being used in the R500 project. And it took a while before all of those name changes and comments were made up to date to reflect the fact that the name change occurred. 09:53 09:53 09:53 09:53

Q. Other than changing 400 to 500, were 09:53

there other changes made to bring the documents up to date? 09:53

MR. PAIGE: Object to the form. 09:54

A. As any design continues, there are small changes made in the design. Fewer as you get closer to tapeout, because you just have to go with what you have. 09:54

So there were probably small changes between, you know, in the R500 that weren't in the R400, but they were quite small. 09:54

Q. What were the changes between the R500 and the R400? 09:54

A. I'm not aware of them off the top of my head. It is possible that we added a dedicated vertex cache as opposed to using the texture cache, but I'm not entirely sure of that. It might have been aimed as a different process technology, for instance, but I don't think that would affect the design that much. 09:55

Q. Why would a separate dedicated vertex cache have been added? 09:55

A. It's part of the unified shading system, and we want to avoid, for instance, textures holding up -- that is pixel processing holding up vertex processing and vertex 09:55

processing holding up pixel processing. 09:55
So there was a concern that, you 09:55
know, we have a long latency texture, but 09:55
vertices would tend to hit more. So you might 09:55
have a situation where, you know, vertices are 09:55
waiting -- vertex processing is waiting on pixel 09:55
processing unnecessarily, because there were 09:56
certain order enforced in the data flow that was 09:56
not logically required. 09:56
Q. This concern with latency you just 09:56
mentioned, was that present in the R400? 09:56
A. The concern about latency? 09:56
Q. Yes. 09:56
A. Latency is always a concern in the 09:56
R400. One of the challenges of the unified 09:56
shader and one of the advantages of the unified 09:56
shader is it allowed us to do a more effective 09:56
job of latency hiding, especially for vertex 09:56
processing. 09:56
One feature that Microsoft was 09:56
pushing in their new API was the idea of being 09:56
able to fetch data from the vertex engine, like 09:56
being able to fetch displacements, if you were 09:57
doing a displacement map. 09:57
And normal vertex shaders have very 09:57

small number of threads and don't do a good job of latency hiding by unifying the vertex shader into the pixel shader. It could take advantage of the latency-hiding characteristics that pixel shaders naturally have. 09:57

Q. In order for a unified shader to take advantage of the latency-hiding characteristics, would it need a dedicated vertex cache? 09:57

A. No, that is not a requirement. 09:57

Q. Then why was a vertex cache added to the R500 to address latency? 09:57

MR. FAHRENKROG: Object to the form of the question. 09:57

A. I'm not testifying that it definitely was added. I have a memory of it being added, but between the R500 and the R600. I don't want to say for sure that it was in the R500. 09:58

I'm sorry, could you repeat the question? 09:58

Q. If a vertex cache is not required to hide latency in a unified shader, why was it added at some time after the R400 project ended? 09:58

A. There was a concern that the pixels were unnecessarily interfering with vertex fetching. It was just a fear and efficiency 09:58

concern at that point. And so it was more of a tweak in the design than something that was required. 09:58 09:59 09:59

It was just, you know, one of those engineering changes that you do because you have the opportunity to do so. 09:59 09:59 09:59

Q. In the R400, how were pixels unnecessarily interfering with vertex fetching? 09:59 09:59

A. The texture system was strictly ordered in the sense that you got data back in the order that you requested. When I say "you," you got data back, I mean the shader ALUs got data back in that order. 09:59 09:59 09:59 09:59

But vertices and pixels have no inherent order. However, once a vertex does a texture fetch, it will come back after the last texture fetch from the pixel shader. So that imposed an order that was otherwise not required. 09:59 10:00 10:00 10:00

So if you could have the vertices come back out of order, either with a separate cache or by changing the requirements of the texture system, you could break that dependency that was otherwise not required. 10:00 10:00 10:00 10:00 10:00 10:01

(Gruber Exhibit 4, Document entitled "Respondents' Invalidity Contentions," marked for identification)

Q. You have been handed what's been marked as Exhibit 4.

A. Yes.

Q. It's excerpts from a document called, "Respondents' Invalidity Contentions." Have you seen the Respondents' Invalidity Contentions?

A. No.

Q. If you would turn to the page at the bottom labeled 252.

When was the R400 offered for sale to Microsoft?

MR. PAIGE: Objection. Foundation.

A. I do not know if the R400 per se was offered for sale to Microsoft.

Q. Do you know what the basis for the statement in Respondents' Invalidity Contentions stating that the R400 was offered for sale to Microsoft in the spring of 2002 was?

MR. PAIGE: Let me interpose an objection here. If there's any information that you've learned from communicating with attorneys, I would instruct you not to include that in your

answer. But, otherwise, you may answer the question. 10:02

A. What was the question again? 10:02

Q. Do you know what the basis for this statement in Respondents' Invalidity Contentions is? 10:02

A. Could you show me where on the page the statement is? 10:02

Q. Bottom of Page 252. 10:02

A. "Upon information and belief offered for sale in the spring of 2002." I don't know. I can't answer to that. 10:02

MR. HIGGINS: Let's take a quick break. 10:03

VIDEOGRAPHER: The time is 10:03 a.m. At this time we'll pause recording and go off the record. 10:03

(Recess) 10:13

VIDEOGRAPHER: The time is 10:13 a.m. We are on the record and recording. 10:13

MR. FAHRENKROG: If I could jump in real quick and just designate the transcript Confidential Business Information of AMD. 10:13

BY MR. HIGGINS: 10:13

Q. If you could turn to Exhibit 2, 10:13

Page 10. So as AMD's 30(b)(6) witness, you have an obligation to investigate -- perform an investigation as to Topic No. 4 as to whether there was a sale of the R400 to Microsoft, and I just want to confirm that it is AMD's position that it does not know if an offer for sale of the R400 to Microsoft occurred? 10:13 10:13 10:13 10:13 10:13 10:14 10:14

MR. FAHRENKROG: Let me jump in on that. That seems more of a question to me. It's AMD's position that we have objected to Topic No. 4 and the scope of it and him testifying about information related to Microsoft. So he's not here to testify about that. It's not AMD's position that AMD does not know when it made an offer for sale to Microsoft. 10:14 10:14 10:14 10:14 10:14 10:14 10:14

MR. HIGGINS: The objection to the depo topic didn't cover -- didn't state that you weren't going to produce the witness to state whether or not an offer for sale to Microsoft occurred. Microsoft information, whether it's confidential or not, is not the subject of that question. 10:14 10:14 10:14 10:14 10:14 10:14 10:14

MR. FAHRENKROG: Well, you can ask him questions around the knowledge he has. And I would let you know that it's not AMD's position 10:14 10:14 10:14

that we don't know whether an offer for sale was made. 10:14

MR. CHATTERJEE: He's a 30(b)(6) witness. Either they have a position or they don't. Are you saying that AMD has no position? 10:14 10:15

MR. FAHRENKROG: I am not taking a position one way or the other. I'm telling you that your statement of our position is incorrect. What I'm saying also is that we object to this topic -- 10:15 10:15 10:15 10:15 10:15

MR. CHATTERJEE: If that's the case, the witness needs to answer it, not you. 10:15 10:15

MR. FAHRENKROG: I'm stating our position around our objection. We've objected to this topic. We don't think it's an appropriate topic here. He's not here to testify on it. That's our objection. You can ask him the question. 10:15 10:15 10:15 10:15 10:15 10:15

MR. CHATTERJEE: Okay. Go ahead and ask the question. 10:15 10:15

BY MR. HIGGINS: 10:15

Q. As AMD's 30(b)(6) witness, was the R400 offered for sale to Microsoft? 10:15 10:15

A. I don't have personal knowledge of that, but I believe that within ATI there is such 10:15 10:15

knowledge. 10:15

Q. Prior to today's depo, did you perform any investigation as to whether that sale of the R400 occurred? 10:15

A. No, I do not. 10:15

Q. What information would you have to review in order to make the determination of whether an offer for sale of the R400 was made to Microsoft? 10:15

A. I think we'd have to go back into records, perhaps e-mails, both internal and between ATI and Microsoft to come up with the date and the details of the offer for sale. 10:16

Q. So there was an offer for sale to Microsoft? 10:16

A. Oh, there was definitely an offer for sale. I'm not sure if it was of the R400 in particular or some derivative of the R400. 10:16

I expect that there was an offer for sale of the R400, but I don't know when that occurred either. 10:16

Q. And when you say that you expect there was an offer for sale for the R400, do you mean to Microsoft or to another party? 10:17

A. No, to Microsoft. Again, I was not 10:17

-- I can't talk from personal knowledge here, but my experience in these type of things, because I was involved in an earlier attempt to sell IP to Microsoft, is that we tried to sell what we have. What typically happens is that there's a negotiation back and forth, and there are requests for differences. 10:17

Q. Did this back and forth occur between ATI and Microsoft? 10:17

A. You know, I can't really say. As I said, I think that I'd have to do some investigation into the matter that I haven't done at this point to provide a definitive answer. 10:17 10:18

Q. Was the R400 offered for sale to any other party? 10:18

A. Not that I know of. But again, I haven't done an investigation into that topic. 10:18

Q. Did Microsoft reject ATI's offer to sell the R400? 10:18

MR. FAHRENKROG: Object to the form. 10:18

A. I really can't answer. 10:18

Q. Do you not know whether Microsoft rejected an offer for sale of the R400? 10:18

A. I do not know at this time. I presume that there is documentation within ATI on 10:18

this negotiation, but I have not seen it. 10:19
Q. If there was a discussion with 10:19
Microsoft with respect to an offer for sale of 10:19
the R400, as the head architect, would you have 10:19
been involved in that discussion? 10:19
A. Not necessarily. I mean, it's 10:19
possible that I was, but it's also possible that 10:19
it was done more at a business level. 10:19
Q. When did negotiations with Microsoft 10:19
for the Xenos chip begin? 10:19
A. I don't have any recollection of 10:19
that. I may not have been directly involved with 10:19
it. Again, I'm sure there are documentation 10:19
within ATI, but I haven't seen it. 10:19
Q. Did the Xenos project begin after the 10:20
R400 was offered for sale to Microsoft? 10:20
MR. PAIGE: Object to form. 10:20
A. I can't answer that, because I do not 10:20
know the date of when or if the R400 was offered 10:20
for sale to Microsoft. 10:20
Q. Would the R400 have been offered for 10:20
sale to Microsoft after the Xenos project had 10:20
already started? 10:20
MR. PAIGE: Object to the form. 10:20
MR. FAHRENKROG: Object to the form. 10:20

A. No. I think that the Xenos chip would have been started in response to any offer or negotiation with the R400, because the Xenos chip would have simply been a response to the offer, and whatever modifications Microsoft wanted in the database. 10:20 10:20 10:20 10:21 10:21 10:21

Q. And when did work on the Xenos chip begin? 10:21 10:21

A. Early 2003. 10:21

Q. When did discussions with Microsoft regarding the Xenos chip begin? 10:21 10:21

MR. FAHRENKROG: Object to the form. 10:21

A. I don't know. I can't answer that. 10:21

Q. Would those discussions have happened prior to the Xenos project beginning? 10:21 10:21

MR. PAIGE: Object to the form. 10:21

A. Yes, they would. 10:21

Q. Do you have an idea how much earlier in time those negotiations would have started before ATI would have begun the project? 10:21 10:22 10:22

MR. FAHRENKROG: Object to the form. 10:22

A. No. It's purely a function of how well the negotiations would have gone. It would clearly have been measured in months, though, because nothing would have gone faster than that. 10:22 10:22 10:22

Q. Why would nothing have gone faster than that? 10:22

A. Well, to agree to, you know, make a special chip for Microsoft is a substantial decision. That simply takes a certain period of time. You also have to check the technical feasibility of doing whatever those changes were. 10:22

So that isn't a short time period. 10:22

(Gruber Exhibit 5, Document entitled "Non-Disclosure Agreement," Bates AMDITC-NVIDEA_0000157, marked for identification) 10:23

Q. You have been handed what's been marked as Exhibit 5, a document entitled "Non-Disclosure Agreement," with a Bates label AMDITC-NVIDIA_0000157. 10:23

Have you seen this document before? 10:23

A. Not to my knowledge. 10:23

Q. Do you see at the top where it says, "To protect confidential information, ATI Technologies, Inc., and its subsidiaries ('ATI'), and University of Waterloo"? 10:24

A. Yes. 10:24

Q. What was ATI working on with the University of Waterloo? 10:24

MR. FAHRENKROG: Object to the form. 10:24

Foundation. 10:24

A. My recollection is Dr. McCool had some technology involved with shading languages. 10:24

Q. What do you mean by "technology involved with shading languages"? 10:24

MR. FAHRENKROG: Objection. Outside the scope of the 30(b)(6) notice. 10:24

A. I don't recall much of the details of this. I just remember it had something to do with shading languages. 10:24

Q. Do you know what information about the R400 was disclosed to the University of Waterloo? 10:24

A. I just don't have any recollection after how many years it has been, 13 years. I remember the name Michael McCool. 10:25

Q. What types of shading languages were used with the R400? 10:25

A. Well, generally we were aimed at implementing the Microsoft DX9 assembly language, shader model 3.0. 10:25

Q. Did the R400 ever implement Microsoft DX9? 10:26

A. Under simulation, it did. I mean it was never produced, but we demonstrated working 10:26

shaders.

Q. Were working shaders demonstrated to the University of Waterloo?

A. I don't know.

Q. What was Michael Doggett's role on the R400?

A. Michael Doggett was on the architecture team. I think he did a number of things. I specifically remember him working in the texture system, but I think that he did other things as well.

Q. Did he do any work on the shading language?

A. I don't recall. Michael also acted as kind of a representative to academia. So even if it wasn't his particular area, he would have been involved in this.

Q. Did he have agreements with other academic institutions with respect to work being performed on the R400?

MR. FAHRENKROG: Object to form. Outside the scope.

MR. PAIGE: Lacks foundation.

A. I don't recall. It's possible.

Q. Did you do any work with academic

institutions while on the R400 project? 10:27
A. I remember a visit to Stanford, and 10:27
talking to a professor there who later went to 10:28
work for NVIDIA, but I forgot his name. 10:28
Professor Hanrahan and his team. But nothing of 10:28
substance came from that. 10:28
Q. What did you talk to the professor at 10:28
Stanford about? 10:28
MR. FAHRENKROG: Object to form. 10:28
Outside the scope. 10:28
A. I don't recall. It had something to 10:28
do with shaders and GPUs. But beyond that, I 10:28
don't recall. Stream processors. I remember 10:28
that. 10:28
Q. What was the date that you spoke with 10:28
the professor at Stanford? 10:28
MR. FAHRENKROG: Object to form. 10:28
Outside the scope. 10:28
A. I just don't remember. 10:28
Q. Do you know if you met with the 10:29
professor at Stanford before or after the R400 10:29
project started? 10:29
MR. FAHRENKROG: Object to form. 10:29
Outside the scope. 10:29
A. I can't even give you an answer on 10:29

that I'm afraid. 10:29
MR. FAHRENKROG: Please give me a 10:29
chance to make my objection. 10:29
THE WITNESS: I'm sorry. 10:29
MR. FAHRENKROG: Thank you. 10:29
Q. What about your discussions with the 10:29
professor at Stanford with respect to stream 10:29
processors, do you remember? 10:29
MR. FAHRENKROG: Object to form. 10:29
Outside the scope. 10:29
A. I just don't have a strong memory of 10:29
that at all. I remember the trip. I remember 10:29
talking to a couple of his grad students as well, 10:29
but I don't recall the nature of the 10:29
conversation. 10:29
Q. Do you remember the names of those 10:29
grad students that you spoke with? 10:29
A. I believe one's last name was 10:29
Proudfoot. 10:30
Q. Proudfoot? 10:30
A. (Witness nods). 10:30
Q. Do you remember any other names? 10:30
A. It shouldn't be hard to track that 10:30
person down. 10:30
No. Oh, I remember they were working 10:30

on a shading language called Brook. 10:30
Q. B-r-o-o-k? 10:30
A. Yes. 10:30
Q. What was the Brook shading language? 10:30
A. I don't recall anymore. It's amazing 10:30
I came up with that. 10:30
Q. Was any of the technology from the 10:30
Stanford stream processors incorporated in the 10:31
R400? 10:31
MR. FAHRENKROG: Object to the form. 10:31
MR. PAIGE: Join in the objection. 10:31
A. No. 10:31
Q. Did you say no? 10:31
A. Yes, no. 10:31
Q. Is the stream processor that Stanford 10:31
developed different than the unified shader in 10:31
the R400? 10:31
MR. FAHRENKROG: Object on form. 10:31
Lacks foundation. Outside the scope of the 10:31
notice. 10:31
A. I'm not familiar enough with the 10:31
Stanford design to say one way or the other. 10:31
Q. What is your understanding of the 10:31
stream processor? 10:32
MR. FAHRENKROG: Object to form. 10:32

Outside the scope. Lacks foundation. 10:32

A. A stream processor is generally something that works on a small amount of incoming data, processes -- does a lot of processing, and then outputs a small amount of output data. 10:32

Q. Was the Stanford stream processor a GPU? 10:32

MR. PAIGE: Object to the form. 10:32

MR. FAHRENKROG: Object to the form. 10:32

A. I believe that the goal was trying to see how well GPUs act as stream processors, and whether there are modifications that could make them into better stream processors. 10:32 10:33 10:33

Q. Did you begin working at ATI in 1994? 10:33

A. I did. 10:33

Q. What initial position did you hold? 10:33

A. I was the lead architect on the ATI Rage product. Lead 3D architect. 10:33 10:33

Q. How long did you work on the ATI Rage product? 10:34 10:34

A. The Rage were a series of products. It probably went on for five years. 10:34 10:34

Q. Do you recall what you worked on after the ATI Rage products? 10:34 10:34

A. Then we switched on to the next family, which were called Radeons, Radeon 100, Radeon 200, called the R100 and R200 as well internally. 10:34

Q. Did you work on the Radeon R100? 10:34

A. I did. 10:34

Q. Did you work on the R200? 10:34

A. I did. 10:34

Q. Did you work on the R300? 10:34

A. I did not. 10:34

Q. Why did you not work on the R300? 10:34

A. We acquired a group on the West Coast from a company called ARTX. I guess we acquired the company ARTX. And at that point we went to the leapfrogging design teams, where the R300 would be done by the West Coast team, and the R200 and R400 would be done by the East Coast team. 10:34 10:35 10:35 10:35

Again, with Orlando -- at some point in here we moved vertex processing into the GPU. I'm not sure whether that was in the R100 or the R200. The vertex processor was done by the Orlando group, and that was done for both the R200 and the R400 by the Orlando group. 10:35 10:35 10:35 10:35 10:35

Q. While you were at ATI, did they have 10:35

any formal system to track the status of projects? 10:35

A. Yeah. There was status tracking. I'm not sure if it was formal in, you know, a status tracker tool or via just, you know, Microsoft Project, Microsoft spreadsheets. But those certainly existed. 10:35 10:36 10:36 10:36

Q. Was there any status tracking done for the R400? 10:36 10:36

A. Yes. 10:36

Q. Do you know what form the status tracking of the R400 was in? 10:36 10:36

A. I don't recall. I believe it was Microsoft Project, but I could not say 100 percent. 10:36 10:36 10:36

Q. What types of things go into the status tracker that was done in Microsoft Project? 10:36 10:36 10:36

A. The various milestones, how far we were from the milestones, the resources associated with each of the blocks, the number of tests written and passed, both in the emulator and in the RTL code. 10:36 10:36 10:36 10:36 10:37 10:37

Q. If a milestone was not met, was it rescheduled to a later point in time in the 10:37 10:37

tracking system? 10:37
MR. PAIGE: Object to the form. 10:37
A. Yes. The milestone was adjusted. 10:37
Q. Was there a milestone for the tapeout of the R400? 10:37
A. I can't speak to direct knowledge, but I expect that there was. 10:37
Q. Do you know how many times the milestone of the tapeout for the R400 was changed? 10:37
A. I do not. 10:37
Q. What is the best source of information that would have been at ATI to determine what the status of a given project was at a point in time? 10:38
MR. FAHRENKROG: Object to the form. 10:38
A. It would have been this status tracker, as I say, I believe in Microsoft Project. I'm not sure if those were saved or under revision control. So I can't say whether that's available now. 10:38
Q. If someone viewed the status tracker of the R400, how would they confirm what stage of the development the R400 was in at a point in time? 10:38

MR. PAIGE: Object to the form. 10:38
A. I think they would be able to see 10:38
achievements against milestones. 10:38
Q. Would there be any kind of 10:38
documentation they could go to to confirm the 10:39
entry in the milestones? 10:39
MR. FAHRENKROG: Object to the form. 10:39
A. I don't know. 10:39
Q. What is the most accurate source of 10:39
information that describes the function of the 10:39
R400? 10:39
A. I would say the top level and 10:39
block-level specs are a good starting point for 10:39
understanding. The most accurate would be the 10:39
RTL code, which would be the definitive answer on 10:39
how something works. 10:39
Q. Are there differences between the top 10:39
level and block-level specs for the R400 and the 10:39
RTL code for the R400? 10:40
MR. PAIGE: Object to the form. 10:40
Lacks foundation. 10:40
A. Certainly not intentionally. But in 10:40
some cases, I think there are differences. 10:40
Q. Are you aware of any differences 10:40
between the top level and block-level specs for 10:40

the R400 and the RTL code for the R400? 10:40

A. I'm not aware of any -- well, it 10:40
depends on what version of the spec you are 10:40
looking at. I certainly looked at back-rev 10:40
versions of specs that I could see have like 10:40
known typos in them that are clearly not the same 10:40
as what's in the RTL. 10:40

But I don't, in the most up-to-date 10:40
specs, I don't know of any difference that was 10:40
present, but I suspect that there were still 10:41
differences. 10:41

Q. How can you tell what the most 10:41
up-to-date ATI specification is for the R400? 10:41

MR. FAHRENKROG: Object to the form. 10:41

A. I would just look at the last one in 10:41
revision control. 10:41

Q. Is the revision control included in 10:41
each document? 10:41

A. There is some revision control, but 10:41
you had to, you know, manually add it into the, 10:41
like, the front page of the document. A better 10:41
source is the edit date or what comes out of the 10:41
revision control when it was checked in, which 10:41
will automatically assign, you know, order 10:41
things. 10:42

MR. HIGGINS: We have to change tape. 10:42
So take a break. 10:42
VIDEOGRAPHER: The time is 10:42 a.m. 10:42
At this time we will end cassette number one. 10:42
We'll stop recording and go off the record. 10:42
(Recess) 10:49
VIDEOGRAPHER: The time is 10:50 a.m. 10:50
At this time we are beginning cassette number 10:50
two. We are on the record. We are recording. 10:50
BY MR. HIGGINS: 10:50
Q. Are the R400 specification documents 10:50
or the R400 RTL a better indication of the status 10:50
and development of the R400 project? 10:50
MR. FAHRENKROG: Object to the form. 10:50
Vague. Overly broad. 10:50
A. The RTL is more accurate in terms of 10:50
functionality, that is functionality at that 10:51
time. The documentation gives you information 10:51
beyond that, that in some ways goes to status as 10:51
well. 10:51
But I'd say that the accuracy of the 10:51
RTL is probably the best guide to where you are 10:51
in terms of full implementation. 10:51
Q. If a feature is not included in the 10:51
RTL code, could it be implemented in the GPU? 10:51

MR. FAHRENKROG: Object to the form. 10:51

A. Anything not in the RTL code is not in the GPU. 10:51 10:52

Q. Did the R400 include a primitive assembly block? 10:52 10:52

A. Yes, it did. 10:52

Q. What components were included in the primitive assembly block? 10:52 10:52

A. The primitive assembly block managed the parameter cache. It also performed setup operations for the triangle. It did clipping. It did viewport transform. 10:52 10:52 10:52 10:52

Q. Did it do anything else? 10:52

A. It did setup for the interpolators. So it created the barycentric values and the Z, that is the depth interpolation as well. 10:52 10:52 10:53

Q. What inputs was the primitive assembly block designed to have? 10:53 10:53

A. The primitive assembly block got indices to initiate, to kind of track whether the vertex was already in the parameter cache or not. 10:53 10:53 10:53

It also got input from positions that it would use as input to the whole setup function and the clipping and the viewport transform. 10:53 10:53 10:53

Q. From what units or modules would the 10:53

primitive assembly block receive these inputs from? 10:53 10:54

A. So the indices would be fetched from memory. The positions would come via the shader through some kind of position output buffer. 10:54 10:54 10:54

Q. What units or modules would the primitive assembly block output data? 10:54 10:54

A. It would output to the raster engine the clipped triangles, as well as the Z block information. It would also talk to the position cache to retrieve the particular positions as input it was interested in. Although, I'm not sure the position cache may actually have been inside the primitive assembly block, so it wouldn't have to talk to an external block. 10:54 10:54 10:55 10:55 10:55 10:55 10:55 10:55

Q. Did the primitive assembly block output to any other module other than the rasterizer? 10:55 10:55 10:55

A. Any other than the rasterizer? Oh, well, sure, because it would have to output basically the indices that would have to be transformed into the shader. So if it missed in the parameter cache, it would create a vector of indices that would be loaded through the interpolator block. 10:55 10:55 10:55 10:56 10:56 10:56

It was just a data path into the shader, and it would talk to the sequencer to tell it, okay, I have a vector of vertices ready to be transformed. 10:56

Q. What is a vector of vertices? 10:56

A. The shader works in a SIMD fashion. That's single instruction multiple data. So by a "vector," I mean the group of vertices that will be processed as a unit. Some people might call that a thread once it gets into the shader. 10:56

In this case, each vertex is represented by an index into a vertex buffer so the rest of the vertex can be fetched based on that index. 10:57

Q. Is there any difference between vector of vertices in a vertex thread? 10:57

A. Under some definitions, no. But "thread" can be used in a lot of different ways. So I just use the term "vector," because that's what most of the R400 documentation uses it. 10:57

I think "thread" is, at least for some meanings of a thread, that's reasonable. 10:57

Q. With respect to the R400, is a "vertex thread" the same thing as a "vector of vertices"? 10:58

MR. FAHRENKROG: Object to the form. Asked and answered. 10:58

A. The only distinction that I'd make is that something is generally called a thread once it acquires control flow information about it and is in the process of being executed. 10:58

In the case of a vector of vertices, these may be queued up for input into the shader system, and may not have acquired any kind of control flow information or temporary storage available, attached to them, that we normally think of when we say a thread, meaning like a thread of execution. 10:58

There is no, you know, these vertices, as they come out of the PA, they don't have any, like, program counter associated with them. 10:58 10:59 10:59

Q. Did the vector of vertices get a program counter when they enter the shader processor? 10:59

A. Yes, they do. 10:59

Q. When did the development of the primitive assembly block begin? 10:59

A. Well, that was part of the initial R400 plan. So I don't think RTL work on the 10:59

block began in late 2000, early 2001. The initial design of the block did begin in that time frame. 10:59 11:00 11:00

Q. Was the final design of the primitive assembly block completed as part of the R400 project or as part of the Xenos project? 11:00 11:00 11:00

MR. PAIGE: Object to the form. 11:00

A. Well, I'm sure the primitive assembly block evolved as any design evolves through time. But it was largely completed in the R400 time frame. I don't think that was one of the blocks that caused the R400 to be delayed. 11:00 11:00 11:00 11:00 11:00

Q. What were the blocks that caused the R400 to be delayed? 11:00 11:00

A. The primary block that had an issue was the RB, the render back-end. 11:00 11:00

Q. What was the issue with the render back-end? 11:01 11:01

A. The caching system for the decompressed logic just took a long time for us to actually get working. 11:01 11:01 11:01

Q. Why did it take a long time to get working? 11:01 11:01

A. Sometimes you have too ambitious a design. Sometimes you have less competent 11:01 11:01

engineers than you thought you had. 11:01

Q. Were there any other blocks in the R400 that caused it to be delayed? 11:01

A. The RB sticks out in my mind as the long pole block. I don't recall any issues with other blocks, whether there were, you know, other poles that also meant that we couldn't make the particular market window. 11:01

Q. Did the R400 have a rasterizer? 11:02

A. Yes, it did. 11:02

Q. Where was the rasterizer in the R400? 11:02

A. The rasterizer took the output of the primitive assembly. This is not the indices, but the triangles that are the result of the output of the clipper and the setup unit within the PA. And then it fed the sequencer with the resulting pixels. It would produce what we called barycentric coordinates per pixel. So what we called an I and a J were a representative pixel, as well as an X/Y location. 11:02

Q. What components were included within the rasterizer? 11:03

A. Well, there were edge functions that determined whether a pixel was inside or out. So some kind of triangle walker. 11:03

There was also an interface to the RB and back, which would do the Z processing, the depth processing logic, to determine whether a pixel should be culled or not rather than go to the shader. 11:03 11:03 11:03 11:03 11:04

Q. When did the development of the rasterizer in the R400 begin? 11:04 11:04

A. Again, the design took place or the initial architecture in 2000, in late 2000 and early 2001. My recollection is the rasterizer itself was largely based off of previous versions of rasterizers. So it was not -- that wasn't that challenging of a block. 11:04 11:04 11:04 11:04 11:04

Q. Which previous versions was the R400 rasterizer based on? 11:04 11:04

A. It would have come from the R200 rasterizer. 11:04 11:05

Q. Did the R400 include a scan converter? 11:05 11:05

A. Yeah, by "rasterizer" and "scan converter," we mean very similar things. 11:05 11:05

Q. Are there different blocks in the R400 for a rasterizer and scan converter? 11:05 11:05

A. No. 11:05

Q. Did the R400 include a sequencer? 11:05

A. Yes, it did. 11:05
Q. And where was the sequencer in the 11:05
R400? 11:05
A. So the sequencer was very tightly 11:05
tied to the shader system, that it took input 11:05
from the rasterizer and from the primitive 11:05
assembly and setup blocks. The primitive 11:05
assembly really. 11:06
Q. What inputs did the sequencer receive 11:06
from primitive assembly? 11:06
A. So it received this vector of indices 11:06
that represented triangles -- sorry, represented 11:06
vertices, that hadn't been transformed yet. 11:06
Q. What did it do with the vector of 11:06
vertices when it receives them? 11:06
A. So it would try to allocate a thread 11:06
for them. That is, it would try to allocate GPR 11:06
space and it would try to allocate a control 11:06
storage space associated with that vector of 11:06
vertices so it could be submitted for processing 11:06
to the GPU -- to the ALUs. Excuse me. 11:07
Q. How did the sequencer allocate 11:07
storage space to the vector of vertices? 11:07
A. So we had two different types of 11:07
storage space, at least two different types 11:07

within the sequencer it was put further. There was space for the control information, which are things like the program counter, predicates, a return stack. And we had space for temporaries in the GPRs. 11:07

In the control flow, we had a pool of control flow resources that were allocated on a FIFO basis dedicated to the vertices. 11:08

In the GPRs, there was a pool of registers, and over time the registers could go either to vertices or pixels. But at any given period of time, there was a circular buffer associated with vertices and a separate circular buffer associated with pixels, and we'd allocate space out of the circular buffer, again, in order, based on following this ring buffer for that particular vector of vertices based on state information that told us how many registers this particular program -- the particular vertex shader program that was loaded at the time required. 11:09

Q. Did the sequencer in the R400 assign a program counter to a vector of vertices? 11:09

A. Yes, it's based on state. Right. There was some state that said, Here is where the 11:09

starting program counter is for vertex shader, for the vertex shader. 11:09

So when a vector was loaded into the thread buffer, that would be the initial program counter associated with that vector. 11:09

Q. Is the sequencer part of the shader processor? 11:09

A. It's intimately tied to the shader processor. We, just in terms of organization in the R400, we broke it up into two blocks; a sequencer and a processor. Where the sequencer would have all of the control flow information and do more sophisticated things, while the shader processor was the -- the SP block was more of a pure data path. 11:10

But if you were talking in a more generic level, I would include the sequencer and the shader processor. 11:10

Q. Which unit, the sequencer or the shader processor, assigned a program counter to the vector of vertices? 11:10

A. The sequencer would assign the program counter, and would manage how that program counter got incremented over time. 11:11

Q. Where was the program counter stored 11:11

in the R400? 11:11

A. So the program counter was part of 11:11
this thread pool state that was inside of the 11:11
sequencer. 11:11

Q. And where was the thread pool state 11:11
stored? 11:11

A. It was in some local memories within 11:11
the sequencer. 11:11

Q. Was there a name for the local 11:11
memories within the sequencer where the thread 11:11
pool state was stored? 11:11

A. I don't recall whether we had a 11:11
particular name for it. We had two different 11:11
pools of memories in the sequencer associated 11:11
with thread state or I really should say that a 11:11
single thread was split over two different types 11:12
of memories, and one of the memories was only 11:12
single ported and one of the memories was 11:12
multiported, because the information associated 11:12
with the thread in the multiported memory was 11:12
needed more frequently than the information in 11:12
the single pool of memory. 11:12

So we might have had a different name 11:12
for them. Like we might have called one state 11:12
memory and one register memory, but I'd have to 11:12

go back to the actual documentation. I mean, I could look at -- I could probably -- I don't have it here -- to come up with the actual names that we referred each of those pools of memory by. 11:12

Q. When did development of the R400 sequencer begin? 11:12

A. Again, the initial development began with the initial R400 in late 2000, continuing into 2001. Because this was new, that is the unified shader was new and the sequencer was the key part of it, that was certainly a focus early on. 11:13

Q. When was the development of the sequencer finished? 11:13

A. Well, again, it evolves over time. But I think it reached something close to its final state in the spring of 2002. 11:13

Q. Were changes made to the sequencer in the R400 after the spring of 2002? 11:13

A. I'm sure changes were made, yeah. There was a continuing, you know, evolution of a design where, you know, things generally change, even, you know, up to a few months prior to tapeout. 11:13

The overall architecture and the 11:14

basic functionality was established by the spring of 2002. 11:14

Q. Were there changes made to the sequencer during the Xenos project? 11:14

A. I don't -- well, are you asking as part of the R400 to Xenos transition or just throughout the Xenos project? 11:14

Q. Both. 11:14

A. I don't think there was any change made as part of the transition, and I don't know of any changes made through the Xenos project. 11:14

But knowing that we run tests and sometimes bugs crop up, I wouldn't be surprised if there were changes made. The RTL check-in code under source control would be the best way of assessing when a given change was made. 11:15

Q. Were there any bugs in the sequencer code while the Xenos was in development? 11:15

MR. PAIGE: Object to the form. 11:15

A. I think my answer is similar to the last one. I don't know offhand of bugs in the sequencer code, but it wouldn't surprise me if they were. And the way to establish that would be to look at the history of check-in and the history of changes, which generally the check-in 11:15

comments would say, you know, oh, I'm fixing a bug at this point. 11:15

Q. Was there a particular person in charge of the development of the sequencer? 11:15

A. Laurent Lefebvre was the architect in charge of it. There was probably a block lead on it as well. I'm not sure whether Laurent did the RTL code for it or whether somebody else was doing the RTL code. So I'm not sure about that. I could tell, again, by looking at the check-in. 11:16

It's probable that Laurent did a lot of the RTL code as well. 11:16

Q. Was anyone else working on the sequencer with Laurent? 11:16

A. I don't recall, but I could tell by whatever the history of check-ins were. 11:16

Q. Did Laurent work on the R500 project? 11:16

A. Yes, he did. 11:17

Q. Did Laurent work on the Xenos project? 11:17

A. I don't think so. That is, everybody on the R400 did it at times, you know, help out on the Xenos project, but I don't think he was really on the Xenos project. But I can't testify definitively one way or the other. 11:17

(Gruber Exhibit 6, Document entitled "Xenos Sequencer Specification, Version 3.0," Bates AMD-ITC-0020788 through AMD-ITC-0020843, marked for identification) 11:18

Q. You have been handed what's been marked as Exhibit 6, a document entitled "Xenos Sequencer Specification, Version 3.0," Bates label AMD-ITC-0020788 through 20843. 11:18

Have you seen this document before? 11:18

A. I have seen versions of this document. I'm not sure if I've seen this particular version. 11:18

Q. When is the last time you viewed a version of this document? 11:18

A. Yesterday, as a matter of fact. 11:18

Q. Which version did you view yesterday? 11:18

A. I believe it was 2.02. 11:19

Q. This Version 3.0, what date does it reflect the development of the sequencer? 11:19

A. The date is June 27, 2005. 11:19

Q. Turn to Page 6 in the revision history. 11:19

A. Okay. 11:19

Q. Between May 1, 2003 and September 24, 2004, what work was done on the development of 11:19

the sequencer? 11:19

A. It doesn't show any additional 11:19

changes in the specification. That doesn't mean 11:20

work wasn't continuing to be done on the RTL 11:20

database. 11:20

Q. Was work done on the sequencer RTL 11:20

database between May 1, 2003 and September 4, 11:20

2004? 11:20

A. I can't say whether the database -- 11:20

whether the RTL database changed or not. I'd 11:20

have to go into the CVS. Certainly work was 11:20

continuing on the chip. Whether the sequencer 11:20

needed additional work or not, I couldn't say. 11:20

Q. What does it mean when it says 11:20

"Updated export address table to reflect hardware 11:21

implementation"? 11:21

A. Sometimes that the hardware will find 11:21

an easier way of doing things that doesn't 11:21

actually change the functionality. And so we'll 11:21

change the spec to allow the hardware to do the 11:21

easier path, because it really doesn't matter one 11:21

way or the other. 11:21

But as it turns out, once you do the 11:21

hardware implementation, you find out that one 11:21

way is easier than the other. So the idea was 11:21

the hardware guys went back and said, oh, can we change the spec, because this will be easier, so the spec was updated. 11:21

Q. Then in September 2004 did the hardware guys go back and change the spec? 11:21

A. The hardware guys asked Laurent to change the spec. 11:21

Q. Do you know what changes were made in 2004? 11:22

A. The detail of the changes, no. I'd have to get into the details of the RTL code to really understand this or to, you know, go into the spec if -- I'd have to compare the two revs of the spec to see what actually changed. 11:22

It looks like it's something fairly minor. 11:22

Q. What work was done on the development of the sequencer between September 24, 2004 and June 27, 2005? 11:22

A. Again, it appears that there wasn't a lot of specification changes, but there was continued development on the chip itself, and, if so, there might have been fixes in the RTL. 11:23

So, you know, it's not that the architecture changed, but part of designing a 11:23

chip is to implement the architecture. So that's what was going on. And the spec changes become fewer, because the changes associated with the development of the block are mainly implementation changes. 11:23

Q. What is meant by "Update of spec to reflect bug fixes and features"? 11:23

A. Sometimes you will fix a bug by changing a spec as opposed to actually making the hardware meet the spec. Other times, a fix in the hardware or an implementation in the hardware will result in a small change in the specification that, you know, nobody really cares much about. 11:24

I don't know what this particular refers to. Features, a lot of times, a late-adding thing is like a debug feature, where you'd add some additional visibility, and you'd want that updated in the spec, but you're not really sure at the time that you're doing the initial architecture what those debug features would be. 11:24

It also looks like here that, you know, a lot of the time was simply that the spec was fairly stable, and that Laurent added all of 11:24

the implementation-related changes at one time. 11:25
Q. At what point did the spec become 11:25
stable? 11:25
MR. PAIGE: Object to the form. 11:25
A. As I said, the last major change that 11:25
I know about was in the spring of 2002. 11:25
Q. April 8, 2003, what were the R500 11:25
modifications that were added? 11:25
A. I don't know what this refers to, how 11:25
much of it -- how many changes were associated 11:25
with the R500. I'd have to go into the details 11:25
of this. 11:25
Q. Are you aware of any changes made to 11:25
the sequencer for the R500? 11:25
A. As I said, I have a memory that we 11:25
added a vertex cache. So fetching from the 11:26
vertex cache would certainly be something that 11:26
the sequencer would be -- would have to be aware 11:26
of. 11:26
Q. How does the sequencer interact with 11:26
the vertex cache in the R500? 11:26
A. So the vertex cache is now a second 11:26
input into the shader, and the sequencer has to 11:26
be aware of whether the vertex data has returned 11:26
or not to be able to know whether a given thread 11:26

is ready to be launched if it's waiting for its vertex data to come back. 11:26

Q. When would the vertex cache have been added to the R500? 11:26

A. I am not sure. I'd have to go back and specifically look for the change list in various blocks to see the updated vertex cache. 11:27

Q. Does the sequencer receive any inputs from the vertex cache? 11:27

A. Yes. It would receive a bit saying that data is back to be able to know that that thread has received its data. 11:27

Q. What do you mean by the "thread has received its data"? 11:27

A. I mean, it will do a fetch instruction. This is the same mechanism that is used for fetching textures, and in the R400, the same signal was used for both textures and vertices. 11:28

In the R500, as we switched -- as we, you know, assuming that we did move the vertex cache functionality into a separate path, then you'd have a separate wire saying "vertex data is back" versus a wire saying "texture data is back." 11:28

In both cases, the shader program is capable of saying "I need to stall until data is back." So any shader program that did that would be marked as not ready for execution until the control bit in the sequencer winds up being set the same, based on the incoming data. 11:28 11:28 11:28 11:28 11:28 11:29

Q. Did the Xenos chip have a vertex cache? 11:29 11:29

A. I'd have to look. I'm not sure whether we added a vertex cache. I don't want to guess. 11:29 11:29 11:29

Q. Where would you have to look to determine whether the Xenos chip had a vertex cache? 11:29 11:29 11:29

A. We would have to look in the top-level Xenos document with a block diagram in it. I'm not sure that I have this available right here. It wouldn't -- well, I may be able -- it's possible that it is in the sequencer document, if it would be in the interface. 11:29 11:29 11:29 11:29 11:29 11:29

MR. FAHRENKROG: Is there a current question pending other than where would you have to look? I'm just trying to follow the track here. 11:30 11:30 11:30 11:30

MR. HIGGINS: I'm just waiting for 11:30

him to look. 11:30
MR. FAHRENKROG: I didn't hear you 11:30
ask him. Okay. 11:30
A. Yeah, there is a vertex cache, at 11:30
least in this spec, because it does have an SQ to 11:30
VC control bus. 11:31
Q. Does this document tell you whether 11:31
the vertex cache was implemented in the Xenos or 11:31
would you have to look to the RTL to determine 11:31
that? 11:31
A. No. I think we can rely on this that 11:31
there is a vertex -- well, this is the Xenos 11:31
spec. Let me just see the dates. No, that's a 11:31
large enough change that this can be relied upon. 11:31
Q. How much of a change is large enough 11:31
where you can rely on the documents and not have 11:31
to resort to the RTL? 11:31
MR. PAIGE: Object to the form. 11:31
A. At least in this case, where there's 11:31
a substantial interface consisting of multiple 11:32
signals, tells me that while I can't say that 11:32
every detail is correct, the whole interface with 11:32
multiple signals is more than a detail. So I can 11:32
rely on that. 11:32
There are multiple references to 11:32

vertex cache in this document. 11:32

Q. Could a feature be written in the 11:32

Xenos specification that is not ultimately 11:32

implemented in the RTL for the Xenos? 11:32

MR. FAHRENKROG: Object to the form. 11:32

MR. PAIGE: Join in the objection. 11:32

A. While it's theoretically possible, I 11:32

don't see how our documentation would be that far 11:33

wrong that that's a realistic possibility. 11:33

Q. What do you mean by "that far wrong"? 11:33

A. I mean at least in the case of the 11:33

vertex cache, if you're talking about a very 11:33

substantial feature. If you're talking about a 11:33

minor feature, it is possible that there's 11:33

something in the documentation that isn't 11:33

reflected in the chip. 11:33

But if you're talking about major 11:33

blocks, like a vertex cache, there is no way that 11:33

the documentation would have something that the 11:33

hardware would not have. 11:33

Q. When a new feature is being 11:33

implemented on a project, is it first put into a 11:33

specification document or in RTL? 11:33

MR. FAHRENKROG: Object to the form. 11:33

Overly broad. 11:33

MR. PAIGE: Join in the objection. 11:33

A. It's first put into the specification document generally. I will not say that that's an absolute rule. 11:33 11:34 11:34

Q. How would you confirm whether a feature that's put in the specification document is ultimately implemented in the RTL? 11:34 11:34 11:34

A. To get 100 percent confirmation, you'd have to validate that the RTL has this same interface as it has in the underlying logic. 11:34 11:34 11:34

Q. What does it mean at the top of the document where it says, "Originate date 9 July 2003"? 11:34 11:34 11:34

A. It's put in automatically by Word. So it means either that the document started from scratch then, or a copy of something was made and then renamed at that time. 11:34 11:34 11:34 11:34

Q. Would changes such as that be reflected in the revision history of the document? 11:34 11:35 11:35

A. Yeah, they certainly might. For instance, this, you know, has a revision date of Revision 01 of 2001, even though the origin date has 2003. 11:35 11:35 11:35 11:35

What that tells me is that even 11:35

though a new physical file was created, it started with the same data that was further back. 11:35

So I imagine that this has to do with some major change, maybe when the R500 -- when they name changed from R400 to R500, rather than simply renaming the document, maybe it was copied. 11:35 11:36

Q. If you look at revision history, the R500 modifications were added on April 8, 2003, correct? 11:36 11:36

A. Right. 11:36

Q. So why would the document be changed on July 9, 2003 to reflect something that was already changed on April 8, 2003? 11:36 11:36

MR. FAHRENKROG: Object to the form. 11:36

MR. PAIGE: Object to the form. 11:36

A. I would just be guessing. It's possible that that's when we -- it was -- a copy was made to reflect the Xenos database, you know, and sort of have a separate copy of the spec for Xenos versus R500. It's possible that that was just when the forking of the Xenos spec database occurred. 11:36 11:36 11:36 11:36 11:36 11:37

Q. Would a major change to the specification be adding a vertex cache? 11:37

MR. FAHRENKROG: Object to the form. 11:37

A. I don't think that simply adding a vertex cache to the specification would result in copying of the specification. I think it more has to do with program management type changes as opposed to technical changes that would have led to this. 11:37

Q. When was the vertex cache added to the sequencer specification? 11:37

MR. PAIGE: Object to the form. 11:37

A. Well, all I can do is go by this, where it says -- oh, it doesn't actually say here. I don't know. I expect it had to do with when we added a vertex cache to the design as part of the R500. 11:38

Q. Would adding a vertex cache generate a new revision of the specification? 11:38

A. No, not automatically. If we had chosen to add that at some other point in the design cycle, it would not necessarily have led -- oh, to a new revision or to a new origin date? 11:38

Q. A new revision in the document. 11:38

A. Oh, yeah, I expect that it would have. 11:38

Q. Was the vertex cache added on July 9, 11:38

2003 in this revision? 11:38

A. I don't know. There's really no way 11:38

for me to know here. 11:39

I mean, it would be more likely it 11:39

was added when Laurent says on April 8th, "adding 11:39

R500 modifications." But that's just a guess. 11:39

Q. Turn to Page 24 of the document, 11:39

Bates number 20811. Look at Section 12.1 it 11:39

says, "SP Stall Conditions." Do you see that? 11:39

A. Yes. 11:39

Q. What is meant by PS and PV? 11:39

A. Those are bypass registers. What 11:40

that means is that PS means previous scaler and 11:40

PV means previous vector, and it means the result 11:40

of the previous operation. 11:40

The bypass logic is done in the SQ 11:40

rather than in the SP. So the idea is that if an 11:40

instruction is dependent on the previous 11:40

operation, and the data has not made it to the 11:40

register file yet but is kind of in flight to the 11:40

register file -- this is only within the shader 11:40

-- and so you may do an add operation and then 11:40

another add operation, with the output of one add 11:40

operation is used as the input for the second add 11:41

operation, if those two instructions follow each 11:41

other, the data bypasses the register file and is, instead, sent directly from the output of one instruction to the input of the following instruction. 11:41 11:41 11:41 11:41

Q. What does it mean when it says "none" in both of those entries? 11:41 11:41

A. It means that we don't have to stall based on that use, because of the bypass logic. 11:41 11:41

Q. So on the R400 sequencer, there were no stall conditions within -- strike that. 11:41 11:41

Within the SP of the R400, there were no stall conditions; is that correct? 11:42 11:42

A. No. I think this is an improvement from what was -- what at least was in early versions of the R400. I can't say when exactly all the stalls were taken out. 11:42 11:42 11:42 11:42

But there was a stall condition, at least in early versions of the R400, where if you were using only a part of, let's say, the previous vector, you wanted a mixture of the GPR data and the previous result. 11:42 11:42 11:42 11:42 11:42

These are all vector instructions. So you use, say, call them the X, Y, Z and W components or call them, ARGB, and maybe the previous instruction updated the B component. 11:42 11:43 11:43 11:43

And then the next instruction uses all of them, the ARGB. 11:43 11:43

So the ARG needs to come from the register file. The old data of the register file is perfectly good data. It's sitting there. It wasn't modified by the previous instruction. The B data was modified by the previous instruction, and has to come from the previous vector register. 11:43 11:43 11:43 11:43 11:43 11:43 11:43

I believe the previous version of the -- at some point in the R400, there was a stall that was detected by the sequencer for that, and we would read from the GPR. So we would wait for the data to go into the GPRs, and then pull it out. 11:43 11:43 11:43 11:43 11:43 11:44

This says, no, you don't have to stall. You can mix the two. You can take some data from the GPRs and some data from the previous vector. 11:44 11:44 11:44 11:44

Q. About when was that change implemented? 11:44 11:44

A. I don't know. I'd have to go through it. That's a very specific change, and I don't know when it would have occurred. 11:44 11:44 11:44

Q. Did you have to look in the 11:44

specification documents or the RTL to figure out when that change was made? 11:44

A. The best answer would be from the RTL. 11:44

Q. Do you know what module of the R400 RTL you would look to? 11:44

A. It would be in the sequencer module. 11:44

Q. Do you know which module -- 11:44

A. Within the sequencer? No, I'd have to bury into that to really figure it out. 11:44

Q. Did you write any of the R400 RTL? 11:44

A. I wrote a little of the RTL in the shader pipe itself, in the shader pipe data path. 11:44

Q. You're able to read and write RTL? 11:45

A. Yes. 11:45

(Gruber Exhibit 7, Document entitled "R400 Architecture Proposal Version 0.1," Bates AMD-ITC-0005407 through AMD-ITC-0005422, marked for identification) 11:45

Q. You have been handed what's been marked as Exhibit 7, ATI document entitled "R400 Architecture Proposal Version 0.1," Bates label AMD-ITC-0005407 through 5422. 11:46

Have you seen this document before? 11:46

A. Yes, I have. I've seen this 11:46

document. 11:46
Q. When is the last time you've seen 11:47
this document? 11:47
A. I saw this document yesterday. 11:47
Q. Turn to Page 6 of the document. 11:47
Under Section 1.2, what does it mean by "Memory 11:47
is the most open issue on the R400"? 11:47
A. Well, I think what he refers to here 11:47
is that a lot of the design of a chip be built 11:47
around trying to get all of the memory bandwidth 11:47
that is available out of the memory. 11:47
So when he says "It's the most open 11:47
issue," I think that you can think about that not 11:47
as saying, you know, what is the technology or 11:48
what is the design, but merely how much memory 11:48
bandwidth are we going to have available to the 11:48
R400, because that tells you how big a chip you 11:48
have to build. 11:48
Q. Was that issue ever resolved for the 11:48
R400? 11:48
A. Yes. We had a target memory 11:48
technology and a target memory bandwidth. 11:48
Q. What was the target memory 11:48
technology? 11:48
A. I would have to look at some overall 11:48

specification of the chip that would tell me what our bandwidth was. 11:48

Q. What do you mean by "target memory technology"? 11:48

A. I mean, you know, there are various generations of DDR and various speeds. So maybe in this thing, maybe you wanted, you know, 1,600 megahertz PC DDR3. And especially -- that's more of a system memory. In graphics, we had specialized graphics memory, so there is like GDDR3. 11:49

And it was a matter of timing the release of the chip to the release of available memory, technologies and speeds. 11:49

Q. Turn to Page 8 of the document. Halfway down the page under heading 1.10, "Performance," there's a sentence -- 11:49

A. I'm sorry, I may be on the wrong page. What page am I looking for? 11:49

Q. Page 8. Under the heading 1.10 "Performance." 11:49

A. Yes. 11:50

Q. Where were the R400 team considering a dual chip solution for the very high end? 11:50

MR. FAHRENKROG: Object to the form 11:50

of the question. 11:50

A. Certainly the market looks for 11:50
scalability. You want to be able to win in the 11:50
high end, but you don't sell most of your chips 11:50
in the high end. So there's a business challenge 11:50
in terms of how you -- how do you win the high 11:50
end without doing a specialized chip for the high 11:50
end. 11:50

Q. How would the single-chip design be 11:50
split up into the dual-chip design? 11:51

MR. PAIGE: Object to the form. 11:51

A. It depends. Sometimes you might have 11:51
some dedicated bus connecting the two chips. In 11:51
other cases, you don't really do much special at 11:51
all. 11:51

Certainly when ATI finally came out 11:51
with multiple chip solutions, there was very 11:51
little -- I mean, I can't think of anything in 11:51
the chips that meant that the high end that was 11:51
present purely to support the dual chip solution. 11:51

So you just took one of your, you 11:51
know, medium high end chips and you put two on a 11:52
board for your very high end high-end chips. 11:52

There may be some on various register 11:52
settings that were added to the chip. And as I 11:52

say, sometimes there would be like a bus connecting kind of the memory, the back-end memory interfaces of the two chips in, like, some possible designs. 11:52

But there was very little done in the core of the chip to support dual chip. 11:52

Q. The dual chip solution they're referring to, is that two GPUs or just implementing a single GPU across two chips? 11:52

A. What I think we finally wound up with was the dual GPU solution. It wasn't like that one of the chips wasn't a GPU. Let me just read this so that I get the context of this particular statement correct. 11:52

Yeah, I think in this case it also meant dual GPUs. 11:53

Q. Did it also include taking a single GPU and implementing it on multiple chips? 11:53

MR. FAHRENKROG: Object to the form. 11:53

A. At this point in the design, that may have been a possibility as well. That's hard to say. 11:53

Q. Was the dual chip implementation pursued at all after November of 2000? 11:53

A. I don't recall any pursuance of this 11:53

for the R400, because any solution would also have to have a single-chip solution as well. So I can't recall anything like that. 11:54

There may have been some, you know, pursuit of specialized memories, but I think that will still be viewed as GPU memory rather than, you know, split up a GPU. 11:54

Q. What do you mean by "specialized memory"? Are you talking about placing some form of memory on a chip other than the GPU? 11:54

A. No. I'm talking about what was used as your memory pool. There was some talk about memory chips that have some logic built into them, you know, similar to what might have been done in Xenos. But there's talk about standard chips doing that as well. Whether they would ever occur or not, who knows. 11:54 11:55

I think the fact that we really had to go for standard chips focused the design fairly early on on not pursuing any of those paths. 11:55

Q. When you say "standard chips," are you referring to a single-chip design? 11:55

A. I'm talking about memories that are available from multiple vendors as opposed to 11:55

contracting with a memory vendor to produce a design specifically for your system. 11:55

Q. Was the Xenos a standard chip design? 11:56

MR. FAHRENKROG: Object to the form. 11:56

A. Yeah, the Xenos was a standard chip. I mean, the Xenos memories were specialized memories, the tile memory. 11:56

Q. Was the Xenos tile memory available from multiple vendors? 11:56

A. No. 11:56

Q. So was the Xenos tile memory a standard chip design? 11:56

MR. FAHRENKROG: Object to the form. 11:56

A. No. Well, it wasn't a standard memory design. It might have used standard chip design techniques to produce it. 11:56

Q. Was this Xenos tile memory on a separate chip than the GPU? 11:56

A. Yes. It's a single package, but a separate chip. 11:57

Q. Was this separate tile memory on another chip considered during the R400 development? 11:57

A. I think it was probably given some thought at some time. But the economics of it in 11:57

the R400 market meant that we couldn't really consider it. 11:57

Q. What was it about the economics of the R400 market that meant you couldn't consider it? 11:57

A. The cost of the memory was too high for a chip that had to compete against, you know, in the market with other GPUs, and it was especially being tied to one memory vendor meant that you couldn't take advantage of the competition for memories. 11:57 11:58

So you might be stuck, even if your GPU chip was competitive, if your memory chip cost was not competitive, then you'd have a business problem. 11:58

Q. Who designed the eDRAM in the Xenos chip? 11:58

A. I don't remember. I'd have to look that up. 11:58

Q. Was the design of the eDRAM in the Xenos chip specified by ATI or Microsoft? 11:58

A. I think that that was an ATI design. 11:58

Q. Did ATI choose to implement an eDRAM on a separate chip in the Xenos or was that a feature requested by Microsoft? 11:59

A. I have no specific knowledge of whether this was an ATI proposal or a Microsoft proposal. 11:59

Q. Turn back to Page 8, towards the bottom of the page, Section 3 that says "Schedule." Do you see that? 11:59

A. Yes. 11:59

Q. Was this schedule met? 11:59

A. No. 11:59

(Gruber Exhibit 8, Document entitled "R400 Sequencer Specification SQ, Version 2.02," Bates AMD-ITC-0015547 through AMD-ITC-0015599, marked for identification) 12:00

Q. You have been handed what's been marked as Exhibit 8, an ATI document entitled "R400 Sequencer Specification, Version 2.02," Bates number AMD-ITC-0015547 through 15599. 12:01

Have you seen this document before? 12:01

A. Yes. 12:01

Q. When is the last time you saw this document? 12:01

A. Yesterday. 12:01

Q. Where in the R400 sequencer specification document is there an algorithm or process described related to the creation of 12:01

threads? 12:02
MR. FAHRENKROG: Object to the form 12:02
of the question. 12:02
MR. PAIGE: Join in the objection. 12:02
A. 6.3 talks about the state and the 12:02
creation of threads. 6.3, "Implementation." 12:02
Q. What in Section 6.3 describes the 12:03
process of creating a thread in the R400? 12:03
A. Well, it talks about the buffers that 12:03
the threads would live in once they're created. 12:03
I can look elsewhere to see if there's any talk 12:03
specifically to the creation of a thread. 12:03
Q. Are threads created in two buffers? 12:03
A. Yes. I would say that when a thread 12:03
is put into that buffer, it is created; and when 12:03
it leaves the buffer, it no longer exists. 12:03
Q. How is a thread put into one of the 12:04
buffers? 12:04
A. So a thread comes from either a set 12:04
of vertices from the PA or pixels from the 12:04
rasterizer, and one of those vectors is selected 12:04
and assign a space in the thread buffer, either 12:04
in the vertex thread buffer or in the pixel 12:04
thread buffer, and all of its state bits are 12:04
initialized at that point. 12:04

It's assigned a GPR base, which means it has to allocate from the GPR register file that tells it where it can keep its temporary values. 12:04 12:04 12:05 12:05

Q. How is it determined which vector is assigned a space in the thread buffer? 12:05 12:05

A. I'm not sure what you're asking. You're saying if I have a pixel vector and a vertex vector, or if you're saying which of those do I choose, or are you asking how does it know where it goes in the thread buffer? 12:05 12:05 12:05 12:05 12:05

Q. How do you choose? There's multiple vectors to choose from. How do you choose which vector goes into the thread buffer? 12:05 12:05 12:05

A. Well, the inputs are FIFO'd. So I'm only going to choose in between a pixel vector or the vertex vector. In other words, the pixel vectors are ordered and the vertex vectors are separately ordered. 12:05 12:06 12:06 12:06 12:06

As far as selection in between vertex vectors and pixel vectors, there's an arbitration mechanism. It's typically giving vertex priority to handle which of them go in. But that's only -- that's only due to the fact that there's kind of a single-ported load into the thread buffer. 12:06 12:06 12:06 12:06 12:06

The thread buffer resources themselves are separate, such that you don't need to arbitrate for, say, the state bits for pixels versus the state bits for vertices, because they're not -- they're separate pools of resources. 12:06 12:06 12:06 12:07 12:07 12:07

The GPRs are shared over time. And so when somebody is selected, he has to first make sure that he has space in the GPR pool to be able to be allocated. 12:07 12:07 12:07 12:07

Q. In what way are vertex vectors given priority over the single-ported load into the thread buffers? 12:07 12:07 12:07

MR. FAHRENKROG: Object to the form. 12:07

A. You know, it's just if I have both of them -- as I said, if you have pixel vectors and vertex vectors currently wanting to load, and I have space for both of them, I'll just do the vertex vector first. 12:07 12:07 12:07 12:08 12:08

Q. What happens if you don't have space for both of them? 12:08 12:08

A. Then I can't select that. For instance, I can't start a vertex thread if I don't have any GPR space for that vertex vector thread, because part of the vertex vector -- 12:08 12:08 12:08 12:08

sorry -- part of the thread state is a GPR base that has to be initialized. And if I can't initialize that, then I can't start the thread. 12:08

MR. HIGGINS: We need to change tapes. So why don't we take a break. 12:08

VIDEOGRAPHER: The time is 12:08 p.m. At this time we'll end cassette number two. We'll stop recording and go off the record. 12:08

(Luncheon Recess) 12:31

VIDEOGRAPHER: The time is 1:11 p.m. At this time we are beginning cassette number three. We are on the record. We are recording. 01:11

BY MR. HIGGINS: 01:12

Q. I'm going to refer again to the R400 sequencer specification. I believe it's Exhibit -- 01:12

A. Eight. 01:12

Q. -- 8. Look on Page 7 of the document. Then at the first paragraph, it says, "There are two separate reservation stations, one for pixel vectors and one for vertices vectors. This way a pixel can pass a vertex and a vertex can pass a pixel." 01:12

Do you see that? 01:12

A. Yes. 01:12

Q. What does that mean? 01:12
A. It means that the vertices are ordered relative to vertices, and pixels are ordered relative to pixels. A vertex thread starts later than a pixel thread may finish earlier than a pixel thread. 01:12 01:12 01:12 01:13 01:13
Q. In what ways would a vertex thread that starts later than a pixel thread finish earlier than a pixel thread? 01:13 01:13 01:13
A. They are running two different shader programs. So how long it takes any one element to execute the shader program. Shader programs could be radically different. So there is no relationship in terms of how long it would take. 01:13 01:13 01:13 01:13 01:13
By keeping them in separate, these reservation stations are FIFO oriented, meaning first in, first out. By keeping them in separate reservation station pools, you're able to order them independently. 01:13 01:13 01:13 01:13 01:14
Q. So when it says "a pixel can pass a vertex and a vertex can pass a pixel," is this really just referring to the length of time it takes a thread to execute its instructions? 01:14 01:14 01:14 01:14
MR. FAHRENKROG: Object to the form. 01:14
A. Yes. This is not an individual 01:14

pixel. This is really talking about the threads, and that pixel -- the lifetime of a pixel thread may be shorter than the lifetime of a vertex thread or vice versa. 01:14 01:14 01:14 01:14

Q. How many entries are there in the vertex reservation station for vectors of vertices? 01:15 01:15 01:15

A. As I recall, there are 16. 01:15

Q. How many entries are there in the pixel reservation station for vectors of pixels? 01:15 01:15

A. As I recall, there's 48. 01:15

Q. Did that number ever change during the development of the R400? 01:15 01:15

MR. PAIGE: Object to the form. 01:15

A. I don't know off the top of my head. There's nothing magic about those numbers. So they certainly could have changed without affecting much in the design. 01:15 01:15 01:15 01:15

Q. What was the reasoning behind the 16 entries for vertices and 48 entries for pixels? 01:15 01:15

A. We were trying to make it such that the thread buffer capacity was not going to be the limiting factor. So we wanted to have enough such that we didn't think that we'd have idle ALU cycles. 01:15 01:16 01:16 01:16 01:16

The limiting factor in terms of the number of threads that are active at any given point we were hoping would be the general purpose register files, which are more expensive than these are. So we were hoping to optimize the use of the general purpose register files. 01:16

I explained earlier how you can't really allocate a thread unless you can also allocate space in the general purpose register file. 01:16 01:17 01:17

Vertices -- vertex shaders tend to use more general purpose registers than pixel shaders do. So you don't need as many active threads before you become limited by the general purpose register file. 01:17 01:17 01:17 01:17

Q. You said you were hoping the general purpose registers would be the unit or components that optimize the number of threads running. Is that what ultimately is reflected in the RTL of the R400? 01:17 01:17 01:17 01:17 01:17

MR. FAHRENKROG: Object to the form. 01:17

A. Whether you're limited because of the GPRs or because of the reservation stations is not really a function of the RTL. The RTL provides a certain number of reservation station 01:17 01:18 01:18 01:18

resources, a certain number of GPR resources. 01:18
But the shader program itself says 01:18
how many of those resources a given thread 01:18
actually needs. And so it's a function of the 01:18
programs that get run on the machine as well as 01:18
the compiler for those programs. 01:18
Q. Did the functions of the general 01:18
purpose registers at the end of the R400 project 01:19
function as you had hoped they did, as you 01:19
mentioned earlier? 01:19
MR. FAHRENKROG: Object to the form. 01:19
MR. PAIGE: Object to the form. 01:19
A. I think we typically were register 01:19
file limited as opposed to reservation station 01:19
limited. Again, there could be oddball cases 01:19
where that was not true. 01:19
Q. If you turn to the last page of 01:19
Exhibit 8. Under the heading "Open Issues," it 01:19
reads, "Need to do some testing on the size of 01:19
the register file as well as on the register file 01:19
allocation method (dynamic versus static)." 01:19
Do you see that? 01:19
A. Yes. 01:20
Q. What is that referring to? 01:20
A. That is referring to whether we have 01:20

enough general purpose registers to run the number of threads that we were hoping to run. 01:20

It's not so much as to whether you're limited by the static resources -- by the reservation station or not. It's that you might be limited at the GPR by having too few threads to hide your latency such that we have the reservation station room, but we don't have the GPR room. 01:20

And so the reason why we need to do testing here is that latency is one of those things that's very hard to test up front without a very good model. And it's the number of threads that you're using to hide latency. 01:21

Q. Was testing performed on the size of the register files? 01:21

A. We did performance testing and we, you know, saw, you know, how well we were doing against plan. 01:21

Q. When was that performance testing done? 01:21

A. I can't say for sure whether we were at that point where we could run a full performance test on the R400 database, that is a systems test. 01:21

We certainly ran some block-level testing, you know, throughout the design development stages in 2001 and 2002. But whether we were doing system testing at the time when we changed over to the R5000, I'd have to go back and look at the testing logs to see what tests we could actually run. 01:21 01:22 01:22 01:22 01:22 01:22 01:22

Q. In order to determine which register file size was required, would you have had to do the full system level performance test? 01:22 01:22 01:22

MR. PAIGE: Object to the form. 01:22

A. I think it would be best to do that, but it would have been possible to do some, at least, preliminary testing based on fixed latency models earlier in the design. 01:22 01:22 01:22

Q. During the R400 project, was the size of the register file ever determined? 01:23 01:23

A. Yes, it was. I mean, we had a stake in the ground for the size of the register file. 01:23 01:23

Q. When was the size of the register file in the R400 determined after the testing was performed? 01:23 01:23 01:23

MR. PAIGE: Object to the form. 01:23

A. Again, I'm not sure when the testing was performed. But the testing was mainly, you 01:23 01:23

know, to say, you know, do we have a serious problem here, and the lack of that we would have just gone with our current register file size. 01:23

Q. Did the size of the register files ever change during the R400 project? 01:23

A. Not that I recall. 01:23

Q. Did the size of the register files change during the R500 project? 01:23

A. Not that I recall. 01:24

Q. Is the size of the register file in the R500 the same as in the R400? 01:24

A. I do not recall the size of the register file in the R500 or whether it changed. So you would have to look at the RTL to figure out whether it changed or not. 01:24

Q. Is the size of the register file in the Xenos the same as the size of the register file in the R400? 01:24

A. This is one of those parameters that is very easy to change. So I do not know whether it was changed or not. 01:24

My guess is no. I think that they were built of 128-by-128 register pieces and did not change. But I cannot answer that for sure. 01:24

Q. What would you have to refer to in 01:25

order to answer that for sure? 01:25

A. For sure, it would have to be the 01:25

RTL. 01:25

Q. For the Xenos? 01:25

A. Well, if you're asking whether it 01:25

changed, then I'd have to look at both the Xenos 01:25

and the reference that you're asking whether it 01:25

changed from. 01:25

Q. During the R400 project, was there 01:25

any testing done on the register file allocation 01:25

method? 01:25

A. Yes. 01:25

Q. When was that testing performed? 01:25

A. I can't say for sure. I'd have to 01:25

look at the testing log to see. But, I mean, the 01:25

basic register file allocation mechanism was 01:25

basic to the operation of the part. 01:25

So I'm sure that it was -- that it 01:25

was tested in the design process, both at the 01:25

block -- certainly at the block level. 01:26

Q. As of May, 2002, what was the issue 01:26

with the register file allocation method? 01:26

MR. FAHRENKROG: Object to the form. 01:26

MR. PAIGE: Object to the form. 01:26

A. I think this refers to how well the 01:26

dynamic allocation method works. The static was certainly functional, and it was an issue of how much better performance do we get using the dynamic mode of the register file. 01:26

Q. Was the dynamic mode of the register file allocation method ever functional in the R400? 01:26

A. I believe it had some bugs associated with it. I mean, it was basically functional, but I'm not sure it was completely functional. 01:27

Q. What bugs were associated with the dynamic register file allocation method in the R400? 01:27

A. I'd have to look at the bug log to tell you. I recall that there were bugs, but I don't recall the details of the bugs. 01:27

Q. Were the issues related to these bugs ever resolved? 01:27

MR. FAHRENKROG: Object to the form. 01:27

A. Yes. The basic design was proven to work. 01:27

Q. When was the basic design proven to work? 01:27

A. Well, I don't -- the nature of these bugs were subtle enough that they would not come 01:27

up until you had silicon. So it wouldn't have been until the R600 or possibly -- oh, no, I guess Xenos would have had the same mechanism. 01:28

I don't know whether this worked in the Xenos silicon or not. There may have been an issue with the Xenos silicon with this. 01:28

Q. When you say "this," are you referring to the dynamic register file allocation? 01:28

A. I'm talking about the dynamic register file allocation, yes. 01:28

Q. What were the issues with the Xenos silicon -- 01:28

A. Again, I don't recall the details. I just know that things didn't work when -- again, I'm not 100 percent sure there was a problem with the Xenos, but I seem to have a recollection that there was a problem found in silicon, you know, the first silicon with this. 01:28

Q. So when the first silicon was made for the Xenos chip, the dynamic register file allocation method did not function properly, correct? 01:29

MR. PAIGE: Object to the form. 01:29

A. Correct. 01:29

Q. Was the problem with the dynamic register file allocation method, the Xenos chip, ever corrected? 01:29

A. I can't say without looking at the revision history whether there was a version of Xenos that did have corrected. 01:29

Q. When you say "looking at a version of the revision history," which document or RTL coder you're referring to? 01:29

A. ATI had some bug-tracking software, you know, system. I can't recall what the details are, but it would be in the bug tracking. It wouldn't be -- or possibly in the RTL as well. If I looked in the RTL revision history, that would probably have the fix associated with this in it. 01:30

Q. How did the fixed static that's referred to here register file allocation method work? 01:30

A. Register pools, the size of the vertex pool versus the pixel pool was set by a software register, and vertices allocated out of the vertex portion of the pool, and pixels allocated out of the pixel portion of the pool. 01:30

Q. Was there any register file 01:30

allocation set by the hardware? 01:31

A. Allocation was done by the hardware every time that a thread was launched. The source of the allocation depended on whether it was a vertex thread or a pixel thread. 01:31

Q. How did the software set the division of the register file? 01:31

A. It was purely heuristic based probably on experimentation. Ideally you would want to set it based on the individual shaders that were loaded. I'm not sure if the driver actually did that or just set it up once and left it alone. 01:31

Q. Was there any default hardware setting of the division of the registers between vertices and pixels in the R400? 01:32

A. There's a register that defines how much space is used for vertices and how much space is used for pixels. I don't know if that gets initialized if nobody touches it, because typically, you know, the driver, when it comes up, would set it to some setting. 01:32

I'm not sure what setting the driver used. It probably started out at 50/50, but then as a result of experimentation, it was set to 01:32

something else. 01:32

Q. When you say "result of 01:32

experimentation," are you referring to 01:32

experimentation by the application writer? 01:32

A. No. I'm referring to experimentation 01:32

by our driver team on typical or important 01:32

applications to see, you know, what is a setting 01:32

that would run well across a range of 01:32

applications. 01:32

Q. What was that setting for the 01:32

register file division in the R400? 01:33

A. Well, the R400 never got to the point 01:33

where that experimentation was done. I think in 01:33

our testing, we probably just used, you know, 01:33

half of the register file for the pixels and half 01:33

of the register files for the vertex. 01:33

Q. Do you know on what project this 01:33

experimentation was first done? 01:33

A. It would have to be Xenos. 01:33

Q. Do you know when this experimentation 01:33

was done on Xenos? 01:33

A. I do not. 01:33

(Gruber Exhibit 9, Declaration 01:34

under 37 C.F.R. 1.131, marked for identification) 01:34

01:34

(Gruber Exhibit 10, Declaration under 37 C.F.R. 1.131, marked for identification) 01:35

Q. Mr. Gruber, you have been handed what's been marked as Exhibits 9 and 10. They are both entitled "Declaration Under 37 C.F.R. Section 1.131." 01:35

A. Yes. 01:35

Q. Exhibit 9 was filed during the prosecution of Application No. 11/746,446, and Exhibit 10 was filed during the prosecution of Application No. 13/109,738. 01:35

Have you seen either of these documents before? 01:35

A. Yes, I have. I saw them when I signed them, and I saw them yesterday. 01:35

Q. Starting with Exhibit 9, do you know which patent this relates to? 01:35

A. No. I can read where it is entitled "Multithread Graphics Processing System." 01:36

(Gruber Exhibit 11, U.S. Patent No. 8,400,459, for identification) 01:36

(Gruber Exhibit 12, U.S. Patent No. 8,760,454, for identification) 01:37

Q. You have been handed what's been marked as Exhibits 11 and 12. Exhibit 11 is a 01:37

copy of U.S. Patent 8,400,459. Exhibit 12 is a copy of U.S. Patent No. 8,760,454. 01:37 01:37

You've seen these two documents before? 01:38 01:38

A. Yes. 01:38

Q. Are you listed as an inventor on both Exhibits 11 and 12? 01:38 01:38

A. I am. 01:38

Q. If you look at Exhibits 9 and 10, the two declarations you submitted with respect to these patents, attached to both declarations are portions of RTL code? 01:38 01:38 01:38 01:38

A. Yes. 01:38

Q. Is the RTL code that's attached to Declaration 9 RTL code from the R400? 01:38 01:38

A. I can't say for sure, but it does look like R400 code. 01:38 01:38

Q. Same question for Exhibit 10? 01:39

A. Let me look at Exhibit 10. Again, Exhibit 10 references structures and instructions that I'm familiar with from R400. So I think it's probably that it is -- that these do represent R400 codes. 01:39 01:39 01:39 01:39

Q. Do you believe that the declarations you submitted to the patent office in Exhibits 9 01:39 01:39

and 10 reflect a full disclosure of the R400 project? 01:39

MR. FAHRENKROG: Object to the form. Vague. 01:40

MR. PAIGE: Join in the objection. 01:40

A. What disclosure are you referring to? Where it says A, B, C, D here or are you talking about the disclosure of the patents themselves? 01:40

Q. Not the patents. Just what's in Exhibits 9 and 10, these two declarations. 01:40

A. Okay. 01:40

Q. There's the statements A through D, and then there's code attached to each declaration. Is what's attached to these declarations the full set of code for the R400? 01:40

A. No, this isn't the full set of code for the R400. This is the code for one portion of the R400. 01:40

Q. What portion of the R400 -- 01:40

A. This is a module from within the sequencing. 01:41

MR. FAHRENKROG: Which exhibit were you referring to? I apologize, just to make sure it's clear. 01:41

MR. HIGGINS: Both. 01:41

MR. FAHRENKROG: When you answer, please reference certain exhibits. 01:41

A. Okay. So I'm talking about Exhibit B looks to be a portion of the R400 sequencer code. 01:41

Q. Are you referring to Exhibit B in Exhibit 9? 01:41

A. Yes. Let me just make sure I'm looking at -- I'm not sure, actually, that this is Exhibit A or Exhibit B. It's not labeled. I'll look, figuring one is Exhibit A and one is Exhibit B. 01:41

Q. What is your understanding of reduction to practice? 01:42

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. 01:42

MR. PAIGE: Join in the objection. 01:42

A. Yeah, I can't really explain a legal term. I can tell you what my engineering understanding is. 01:42

Q. What was your understanding of what the purpose of the declaration in Exhibit 9 and Exhibit 10 was to do -- 01:42

MR. FAHRENKROG: Objection to the extent it calls for -- 01:42

MR. HIGGINS: Can I finish my 01:42

question? 01:42
MR. FAHRENKROG: I'm sorry, I thought 01:42
you were done. A couple pauses there. 01:42
Q. When you submitted Exhibits 9 and 10, 01:42
the declarations to the patent office, and you 01:42
signed them, what was your belief as to the 01:42
purpose of these declarations? 01:43
MR. FAHRENKROG: Is your question 01:43
finished? Can I object? 01:43
MR. HIGGINS: Go ahead. 01:43
MR. FAHRENKROG: Object to the extent 01:43
it calls for any attorney-client privileged 01:43
information, and to the extent your knowledge is 01:43
based on information communicated to you by your 01:43
attorneys, I instruct you not to answer. 01:43
A. I'm sorry, my attorney has instructed 01:43
me not to answer. 01:43
Q. Did you understand why you were 01:43
signing declarations reflected in Exhibits 9 and 01:43
10 when you signed them? 01:43
MR. FAHRENKROG: You can answer that 01:43
one yes or no. 01:43
A. Yes. 01:43
Q. What was your understanding of why 01:43
you signed declarations in Exhibits 9 and 10? 01:43

MR. FAHRENKROG: So again, to the extent that that understanding is based on what you were told or directed to do by an attorney at the time, I instruct you not to answer. 01:43

A. Sorry, my attorney instructs me not to answer. 01:43

Q. Did you have any understanding of what you were signing when you signed the declarations in Exhibits 9 and 10? 01:43

MR. PAIGE: Asked and answered. 01:44

A. Yes, I did. 01:44

Q. What was your understanding? 01:44

MR. FAHRENKROG: Same thing. Same instruction. Please do not communicate anything that was communicated to you by an attorney. 01:44

A. Based on my attorney's advice, I decline to answer. 01:44

Q. Did you read the declaration before you signed it? 01:44

A. Yes, I did. 01:44

Q. Did you understand what the declaration said before you signed it? 01:44

A. Yes, I did. 01:44

Q. What was your understanding of what was written in the declaration prior to you 01:44

signing it? 01:44

A. That I was attesting to the fact that these exhibits reflected the state of the R400 at a certain point in time. 01:44

Q. When you mean it reflected the state of the R400 at a certain point in time, are you referring to the entirety of the R400 project? 01:44

MR. FAHRENKROG: Object to the form. 01:45

A. I was referring to the particular pieces of the R400 that were embodied by these design codes. 01:45

Q. What particular pieces of the R400 were embodied by these design codes attached as exhibits to declarations in Exhibits 9 and 10? 01:45

A. In one case, this is the part of the sequencer. And in the other case, it appears to be a wider block that includes the shader execution unit as well. 01:45

Q. If you look at Exhibit 10 in Paragraph 2, what do you mean by "We conceived the invention prior to June 30, 2003"? Paragraph 2 on the bottom of the first page. 01:46

A. In this case, I believe the invention is the shader subsystem that was included in the R400. 01:46

Q. What is the shader subsystem in the R400? 01:46 01:47

A. The unified shader and sequencer. 01:47

Q. And what did you mean by you conceived of the invention of the unified shader and sequencer prior to June 30, 2003? 01:47 01:47 01:47

A. We had a design and a specification that we believed worked. We also had simulation and implementation at the RTL level. 01:47 01:47 01:47

Q. What date did you conceive of the invention you're referring to here? 01:47 01:47

MR. FAHRENKROG: Objection to the extent the term "conceived" is being used in a legal context outside of what's in the declaration. 01:47 01:47 01:47 01:47

A. Early aspects of it were conceived in late 2000, and the ideas were refined, I would say, over the next six to nine months. 01:47 01:48 01:48

Q. When was the entirety of the invention that you reference in Exhibit 2 in your declaration conceived? 01:48 01:48 01:48

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. 01:48 01:48 01:48

MR. PAIGE: Join in the objection. 01:48

A. There were certainly modifications as 01:48

the design evolved, but the fundamental aspects were designed prior to 2002 certainly. 01:48

Q. Do you know specifically what date you conceived of the entirety of the invention as you state in Paragraph 2 of your declaration in Exhibit 10? 01:48 01:49 01:49 01:49

MR. FAHRENKROG: Object to the form. 01:49

A. No, I cannot provide a date. 01:49

Q. What in Exhibit 10, in your declaration, shows that you conceived the invention prior to June 30, 2003? 01:49 01:49 01:49

MR. FAHRENKROG: Object to the form. 01:49

MR. PAIGE: Object to the form. 01:49

A. Well, if these code snippets adequately date from prior to -- this is not obviously the full code of the R400, but it's the shader code, as well as the RTL code here of a portion of the SQ. The fact that these date from prior to the date shown here shows that the invention must have been conceived prior to this date. 01:49 01:49 01:49 01:49 01:50 01:50 01:50

Q. In the RTL code that's attached to Exhibit 10, where are the dates shown? 01:50 01:50

A. I don't see any dates in this code snippet. Though this code snippet is certainly 01:50 01:50

compatible with my belief of the state of the development prior to 2003. 01:50 01:51

Q. What is your belief as to the state of the development prior to 2003? 01:51 01:51

A. That we had a fully functional shader system. "Functional," I mean that we were able to do RTL level simulation of the shader system successfully and pass tests showing a mixture of vertex and pixel processing. 01:51 01:51 01:51 01:51 01:51

Q. On what date did you have a fully functional shader system of the R400? 01:51 01:51

A. I can't give you a date off the top of my head. I'd have to go back and look at our status reports to see, you know, what tests were running and what tests weren't running in any given point. 01:51 01:51 01:51 01:51 01:51

Q. When you signed this declaration stating you had a fully functional shader system with the R400 prior to 2003, what was that determination based upon? 01:51 01:52 01:52 01:52

MR. FAHRENKROG: Object to the form. 01:52

A. It was based on my memory and my knowledge that we did have such a shader system running. 01:52 01:52 01:52

Q. Was your memory and knowledge in 2011 01:52

of the R400 shader system different than your memory and knowledge of the R400 shader system today? 01:52

MR. FAHRENKROG: Object to the form. 01:52

A. I suppose it's probably changed in some respects, but not in that one. 01:52

Q. So if this was based off of your memory of the R400, and your memory is not changed in that respect from the time of filing this declaration until today, how did you sign the declaration saying you had this fully functional shader system prior to June 30, 2003 if you can't give me a date today prior to June 30, 2003 when you had a fully functional shader system? 01:52

MR. FAHRENKROG: Object to the form. 01:53

MR. PAIGE: Object to the form. 01:53

MR. FAHRENKROG: Incomplete hypothetical. 01:53

A. I can't give you a specific date, but I can tell you that it was prior to June 30, 2003. 01:53

Q. How can you tell me it was prior to June 30, 2003? 01:53

A. Because I know when we were working 01:53

on this, and that the amount of time from the time that we conceived it until the time we had, you know, working tests through the shader system was much less than the time from conception to June 30, 2003. 01:53

Q. How much less than the time? 01:53

A. Again, I can't give you a specific date, but I can for sure say it was before the end of 2002. 01:54

Q. Why did the declaration not provide a specific date prior to June 30, 2003? 01:54

MR. PAIGE: Object to the form. 01:54

A. Are you asking me? 01:54

MR. FAHRENKROG: Objection. To the extent it was communicated by an attorney, don't answer. 01:54

Q. When you say you had a fully functional shader system prior to June 30, 2003, which of points A, B, C and D in Exhibit 10 are you referring to? Exhibit 10 is the larger one. 01:55

A. Okay. All of these are true. 01:55

Q. So is the fact that you say you had a fully functional shader system based on the model code that's attached to this declaration? 01:56

MR. FAHRENKROG: Object to the form. 01:56

A. Well, I know we had both model code and RTL code functional prior to this. 01:56

Q. When did you have both model code and RTL code functional prior to this declaration? 01:56

A. Again, I can't give a specific date. 01:56

Q. If you look at Exhibit A in your declaration that's Exhibit 10. The top says "Exhibit A - Model Code." Is this the computer simulation model code that's referenced in the declaration Part A? 01:57

A. Yes. 01:57

Q. Can you point to where threads are created in the model code in Exhibit A? 01:58

MR. PAIGE: Object to the form. 01:58

MR. FAHRENKROG: Object to the form. 01:58

A. This is very complicated code. I'm not sure that I'm going to be able to point to this as I sit here. 01:58

I see there's a lot of allocation being done here, but it's mostly allocation associated with execution. 01:59

Q. Is this computer emulation in Exhibit A, is it multithreaded? 02:00

MR. FAHRENKROG: Object to the form. 02:00

MR. PAIGE: Join in the objection. 02:00

A. I don't believe the code itself is multithreaded, that is the C++ code. I really haven't had enough exposure to the details of the C model code to give a definitive answer. 02:00

Q. Can C model emulation code be multithreaded? 02:00

MR. FAHRENKROG: Object to the form. 02:00

A. Yes, it can be multithreaded. 02:00

Q. Can a C model emulation execute things simultaneously? 02:00

MR. PAIGE: Object to the form. 02:01

MR. FAHRENKROG: Object to the form. 02:01

A. Yes, it can. 02:01

Q. How can it do that? 02:01

A. If you have -- I'm not necessarily saying this code. I thought you were asking me theoretically whether model code that is multithreaded could execute multiple threads simultaneously. 02:01

Yes, you can run a thread on two different processors that are executing simultaneously. 02:01

Q. Can the model code in Exhibit A perform two functions simultaneously? 02:01

A. I don't have enough experience with 02:01

this code to say whether it can or it can't. 02:01
Q. Is the model code in Exhibit A of 02:01
Exhibit 10 a cycle-accurate model? 02:01
MR. PAIGE: Object to the form. 02:02
A. No. Our C models were never intended 02:02
to be cycle accurate. 02:02
Q. Refer back to the beginning of the 02:02
declaration, Exhibit 10 on Page 2. The bullet 02:02
point labeled D, why was that part of your 02:02
declaration made, information and belief? 02:02
MR. FAHRENKROG: Objection to the 02:02
extent that your answer would reveal any 02:02
information communicated by an attorney, I 02:02
instruct you not to answer. 02:02
A. I'm sorry, I can't answer on the 02:02
advice of my attorney. 02:02
Q. When you signed this declaration, why 02:02
did you only state that what is contained in 02:02
Paragraph D only based upon your information and 02:02
belief? 02:02
MR. FAHRENKROG: Same instruction. 02:03
A. I can't answer on the advice of my 02:03
attorney. 02:03
Q. You had no understanding of what this 02:03
paragraph meant when you signed your declaration; 02:03

is that correct? 02:03
MR. FAHRENKROG: Objection. Same 02:03
instruction and asked and answered. 02:03
A. I'm sorry, I can't answer on the 02:03
advice of my attorney. 02:03
Q. When you signed the declaration, did 02:03
you understand what was written in the 02:03
declaration? 02:03
MR. PAIGE: Object to the form. 02:03
MR. FAHRENKROG: You can answer that 02:03
yes or no. 02:03
A. Yes. 02:03
Q. What did you understand Paragraph D 02:03
to mean when it says that what you're stating is 02:03
only on information and belief? 02:03
MR. FAHRENKROG: You can answer as to 02:03
your personal understanding, except to the extent 02:03
it would reveal communications from an attorney, 02:03
then I instruct you not to answer. 02:03
A. I'm sorry, my attorney instructs me 02:03
not to answer. 02:03
Q. You have no personal understanding of 02:03
what Paragraph D meant when you signed this 02:03
declaration? 02:03
MR. FAHRENKROG: Objection. Asked 02:03

and answered. 02:03
MR. HIGGINS: Actually, you 02:03
instructed him not to answer. 02:04
MR. FAHRENKROG: He gave his answer 02:04
pursuant to the instruction. 02:04
MR. CHATTERJEE: Counsel, your 02:04
instruction is actually wrong. He's not asking 02:04
about the communications. He should be answering 02:04
this whatever he understood when he signed 02:04
something under oath. He's not asking about 02:04
communications with lawyers. There's absolutely 02:04
no basis to say that his understanding of what he 02:04
said to the patent office is a privileged 02:04
communication. He's not asking about the 02:04
communication. 02:04
MR. FAHRENKROG: If you listen 02:04
closely to my instruction, it's not to reveal 02:04
anything being communicated to him by an 02:04
attorney. I said he can answer the question as 02:04
to his personal understanding. 02:04
MR. CHATTERJEE: Okay. 02:04
A. My personal understanding was that 02:04
the model code and the chip design code 02:04
demonstrated a functional embodiment of the 02:04
patent. 02:04

Q. When you signed the declaration, what facts did you have that supported your understanding that the model code and chip design demonstrated a functional embodiment of the patent? 02:04 02:05 02:05 02:05 02:05

A. That the model code and the RTL code were codes from the R400 design effort prior to June, 2003, and that -- and I knew that we had such codes running functionally at that time. 02:05 02:05 02:05 02:05

Q. Is the code that's attached to your declaration the same code that was incorporated into the Xenos chip? 02:05 02:05 02:05

A. I don't know whether it was the same code that was incorporated into the Xenos chip. I mean, it looks similar to a lot of R400 code that we had during the R400 design effort. 02:06 02:06 02:06 02:06

I don't know whether exactly this code was what was incorporated into the Xenos chip, but this code certainly performed the functions of whatever code was in the Xenos chip. There may have been slight differences. 02:06 02:06 02:06 02:06 02:06 02:06

MR. HIGGINS: Off the record for a second. 02:07 02:07

VIDEOGRAPHER: The time is 2:07. At this time we will stop recording and go off the 02:07 02:07

record. 02:07
(Recess) 02:07
(Gruber Exhibit 13, Source Code, 02:12
marked for identification) 02:18
VIDEOGRAPHER: The time is 2:18 p.m. 02:18
At this time we're beginning cassette number 02:18
four. We're on the record. We are recording. 02:18
BY MR. HIGGINS: 02:18
Q. What is an RTL simulator? 02:18
A. RTL was the chip design language. A 02:18
simulator does a cycle-by-cycle simulation of 02:18
what will finally be built in the chip. The chip 02:18
is produced from the RTL. 02:18
Q. Does an RTL simulation run on 02:18
hardware? 02:18
A. Well, an RTL simulation runs on a 02:18
computer. So in that sense it runs on hardware. 02:18
Q. If you look in Paragraph C of your 02:19
declaration, Exhibit 10. It says, "The chip 02:19
design code was used by an RTL simulator system 02:19
to validate the operation of an integrated 02:19
circuit version of the claimed graphics 02:19
processing system." 02:19
Do you see that? 02:19
A. Yes. 02:19

Q. When was an RTL simulation done on a complete integrated circuit version of the R400? 02:19

MR. FAHRENKROG: Object to the form. 02:19

A. We were doing simulation all through the design effort of the R400 as the blocks were built, they were simulated. As to when the complete GPU was simulated, I can't point to a specific date. 02:20

But, for instance, the shader subsystem, which this code encompasses, was certainly simulated fairly early in the design as it was built. 02:20

Q. What shader subsystem are you referring to? 02:20

A. I'm talking about the shader and sequencer blocks. Oftentimes we would simulate a block by putting it into an environment where it would be passed vectors into the block and out of the block. 02:21

Q. What RTL simulator program did you use when simulating the R400 RTL? 02:21

A. I don't recall. I seem to remember one from Mentor Graphics, but it has been too long, and I've used too many different ones for me to say for sure. 02:21

Q. Was an RTL simulator a software program? 02:22

MR. FAHRENKROG: Object to the form. 02:22

A. It generally is. There are hardware simulators as well. 02:22

Q. Which simulator was used on the R400 RTL? 02:22

A. Well, we definitely used a software program. We might have used a hardware program as well -- a hardware system as well. 02:22

Q. When using a software RTL simulator program, does it configure the hardware on the computer to perform graphics operations? 02:22

MR. FAHRENKROG: Object to the form. Vague. 02:22

A. I don't know what you mean by "configure the hardware"? It runs software, software code that simulates the RTL. 02:22

Q. Does an RTL simulation allocate actual register file space? 02:23

MR. FAHRENKROG: Object to the form. 02:23

MR. PAIGE: Join in the objection. 02:23

A. It allocates space in a structure managed by the software that represents register file space. 02:23

Q. Does it actually allocate any physical register space for a thread during the RTL simulation? 02:23

MR. PAIGE: Vague and ambiguous. 02:23

MR. FAHRENKROG: Object to the form. 02:23

A. It will allocate space in a memory structure. Whether you call that a register file, I mean, it's a simulated register file. It's simulated by having it just be in a structure understood by the software code. 02:23 02:24 02:24 02:24

Q. Is the code attached to Exhibit B in your declaration Exhibit 10 sufficient to perform an RTL simulation with the entire integrated circuit design of the R400? 02:24 02:24 02:24 02:24

MR. PAIGE: Object to the form. 02:24

MR. FAHRENKROG: Object to the form. 02:24

A. This is not a -- we're talking about Exhibit 10, correct? 02:24 02:24

Q. Exhibit B to Exhibit 10, yes. 02:24

A. This code, this software simulator code, does not simulate the RTL. It is -- it simulates the functions of the RTL. 02:24 02:25 02:25

Q. Go towards the back. It's part of the last 20 pages or so of Exhibit B, the chip design code. 02:25 02:25 02:25

A. I'm sorry. I'm sorry. 02:25
Q. It looks like it's after Page 93 of 02:25
Exhibit A. 02:25
A. Okay. So this is the chip design 02:25
code for a portion of the SQ. I'm sorry, what 02:25
was the question? 02:25
Q. The code in Exhibit B -- 02:25
A. Yes. 02:26
Q. -- that we're looking at -- 02:26
A. Yes. 02:26
Q. -- is this sufficient to simulate 02:26
the entire integrated circuit chip design of the 02:26
R400? 02:26
MR. FAHRENKROG: Object to the form. 02:26
MR. PAIGE: Object to the form. 02:26
A. It will not simulate the entire R400. 02:26
It would simulate important aspects of the 02:26
unified shader. 02:26
Q. What other code would be needed to 02:26
perform an RTL simulation of the entire 02:26
integrated circuit design of the R400? 02:26
MR. FAHRENKROG: Object to the form. 02:26
Overly broad. 02:26
MR. PAIGE: Join in the objection. 02:26
A. You would need the entire RTL code 02:26

for the R400. This is not all the code in the R400. 02:26

Q. If you move on to the small stack of documents to your right in yellow. 02:26

MR. FAHRENKROG: Objection to the characterization as small. I understand it was in gest. 02:27

Q. Exhibit 13 in front of you is AMD confidential source code, Bates label AMD932SC000001 through 2110. 02:28

Have you viewed these source code files before? 02:28

MR. CHATTERJEE: Let the record reflect a big, heavy sigh. 02:28

A. It is possible I've viewed some of them before. I'm quite sure I haven't viewed them in their entirety. 02:28

Q. Let me specify. Have you viewed this printed stack of source code before? 02:29

A. No. 02:29

Q. Do you know whether the source code printed in front of you between the Bates ranges 1 and 2110, represents the entirety of the R400 RTL code? 02:29

A. I do not, no. 02:29

Q. What would you have to do to make that determination? 02:29

A. I would have to compare this code with all the modules that we use to produce the -- or to simulate the R400. 02:29

Q. If you go to the page Bates labeled 319 through 338. Let me know when you have it. 02:29

A. Okay. 02:30

Q. 319 to 338 is a directory of all the source code files that Respondents had printed off the AMD source code computers that's sitting in the small stack of papers in front of you. 02:30

A. Yes. 02:30

Q. If you looked through this directory listing of the files that were printed, would you be able to determine if it represents the complete R400 RTL code? 02:30

A. I'm not sure that I'd be able to attest to that myself. I think what I would probably need is the build script that is used for building R400 and check the reference of these against that build script. 02:31

I mean, I'm perfectly willing to believe you that these are all the ones. But if you're asking for my personal knowledge, I 02:31

certainly don't have all of those -- all the blocks involved committed to memory. 02:31

MR. PAIGE: Your representation is that this block is what Respondents had printed from the source code? 02:31

MR. HIGGINS: Yes. 02:31

MR. PAIGE: Okay. I believe up to Bates range 1450. 1451 through 2110 was a separate production. 02:31

Q. What is the build script for the R400 RTL? 02:32

A. So when you wanted to -- it's a mechanism for running simulation that would grab all the particular files involved in the simulation. 02:32

Q. Did a build script for the R400 exist? 02:32

A. I believe that's how we did it, yeah, that there was a script or an execution. 02:32

Q. When was the build script for the R400 created? 02:33

A. Well, it evolved with the R400 as modules were added. There would be various versions of it. So it would evolve with the R400. 02:33

Q. Was there a build script for the R400 that related to the complete integrated circuit chip design for the R400? 02:33

MR. FAHRENKROG: Object to the form. 02:33

MR. PAIGE: Object to the form. 02:33

A. I believe so. I cannot answer for sure. I don't see why there wouldn't be. 02:33

Q. When would the build script have been created? 02:33

A. Again, it would have -- I mean, you know, things would evolve. So as you add blocks, they get added to the build script. 02:33

You would also use it, for instance, you might use a version of one if you wanted some of the blocks from RTL, and then use other blocks modeled from the C model. 02:33

Q. When was the last block added to the R400 design? 02:34

MR. PAIGE: Object to the form. 02:34

A. I don't know. I'd have to look on the check-ins to see when that happened prior to the conversion to the R500. 02:34

Q. Were blocks added to R500? 02:34

A. I expect that there may -- that there were blocks added to R500. For instance, as we 02:34

added the vertex cache, that would be a block that would be added. 02:34

Q. If you look at the code in Bates range 1 through 64. 02:34

A. I'm sorry, did you say 1 through 64? 02:35

Q. 1 through 64. 02:35

A. Okay. This is the clause.arb. 02:35

Q. If you see at the top in the file directory, it has the date March 27, 2002? 02:35

A. Yes. 02:36

Q. It's the same date through all 64 pages. 02:36

A. Okay. 02:36

MR. PAIGE: Are you asking or representing that? 02:36

MR. HIGGINS: It's kind of a little of both. Let's go back. 02:36

Q. Is all of the code in Bates ranges 1 to 64 labeled with the date March 27, 2002? 02:36

MR. PAIGE: The document speaks for itself. He can look at the 64 pages if you want him to. 02:36

A. Yes. 02:37

Q. Does this collection of code, Bates ranges 1 to 64, represent the entirety of the RTL 02:38

code for the sequencer of the R400 as of March 27, 2002? 02:38

A. I can't be 100 percent certain, but it appears most of the functionality is here. 02:38

Q. Are you aware of any modules of RTL code for the R400 sequencer that are not in Bates ranges 1 to 64? 02:39

A. I'm not aware of any. 02:39

Q. If you go to Bates ranges 65 through 278. 02:39

A. All right. 02:40

Q. The second part, Bates ranges 339 through 537. 02:40

A. I'm sorry, what am I looking for? 02:41

Q. 339 through 537. Do these two ranges of code represent the entirety of R400 sequencer RTL code as of May 29, 2002? 02:41

A. I'm just not going to be able to give you a definitive answer. I would have to, again, look to see how the SQ was built in that time frame. I'm not familiar enough with the internals of the RTL code of the SQ, especially after 13 years, to tell you whether these are all of the codes. 02:41

They certainly do have a lot of the 02:42

modules that I'm familiar with. 02:42
Q. Did you do any investigation prior to 02:42
today's deposition to determine whether the RTL 02:42
code produced by AMD is the complete set of RTL 02:42
code for the R400? 02:42
A. I did not. 02:42
Q. Did you understand that you had an 02:42
obligation to be familiar with the documents 02:43
produced by AMD and to testify on those today? 02:43
MR. FAHRENKROG: Object to the form. 02:43
A. I understood that I had an obligation 02:43
to represent AMD. 02:43
Q. But you did not perform any 02:43
investigation as to the RTL code that AMD 02:43
produced; is that correct? 02:43
MR. FAHRENKROG: Object to the form. 02:43
MR. PAIGE: Misstates prior 02:43
testimony. 02:43
A. I did look over some ATI-produced RTL 02:43
code. I did not look over every piece of ATI RTL 02:43
code or assure that it represented the complete 02:43
SQ source. 02:43
Q. Which ATI-produced RTL code did you 02:43
review? 02:43
A. Well, I produced -- I certainly 02:43

looked at some of the code associated with those patents that we looked at earlier. 02:44

Q. Did you look at any code that was not included in the declarations of Exhibits 9 and 10? 02:44

A. No, I did not. 02:44

Q. So prior to today's deposition, did you review any source code produced by AMD in this investigation? 02:44

MR. PAIGE: Vague and ambiguous. 02:44

A. No, I did not produce -- I did not look at the ATI source code that was represented in these... 02:44

Q. Who wrote the RTL source code for the sequencer module, the R400? 02:44

A. I believe it was Laurent Lefebvre, but I couldn't say for sure that Laurent was the only author. 02:45

Q. Who else wrote RTL code for the R400? 02:45

MR. FAHRENKROG: Object to the form. 02:45

A. For the R400? There was a large team of people working on the source code for the R400. I could only, you know, come up with a handful of names. 02:45

Q. What are the names you can come up 02:45

with? 02:45

A. I remember Jay Wilkinson working on the render back end. Andi Skende working on the shader system. Jocelyn Houle working on the texture pipe. That's all that comes to mind now. 02:45 02:45 02:45 02:46

Q. Which portions of the RTL code did you write? 02:46 02:46

A. I rewrote some of the shader code. But other than that, I didn't write any RTL code in the R400. 02:46 02:46 02:46

Q. When did you rewrite some of the shader code? 02:46 02:46

A. I couldn't give you an exact date, but it was during some of the Xenos development effort. 02:46 02:46 02:46

Q. Why did you rewrite some of the shader code during the Xenos development? 02:46 02:46

A. We were trying to save area. It was not functionally different. 02:46 02:46

Q. Which modules of RTL code did you rewrite during the Xenos development? 02:46 02:47

A. I don't remember the names of the modules, but it would be the shader pipe itself, the data path. 02:47 02:47 02:47

Q. What about the data path of the 02:47

shader pipe did you rewrite? 02:47
A. The adder and the normalizer logic. 02:47
Q. Why did you rewrite those two 02:47
modules? 02:47
A. We were just -- we would compile them 02:47
and try to see where the timing was tight. The 02:47
idea is if you can reduce the timing, you can 02:47
reduce the area. 02:47
Q. So if we were to walk through the 02:47
rest of the RTL code that's sitting here as 02:48
grouped by date, would you be able to tell me if 02:48
any of those groups of RTL code represent the 02:48
complete code for the sequencer in the R400? 02:48
MR. FAHRENKROG: Object to the form. 02:48
Overly broad. 02:48
A. I wouldn't be able to say whether 02:48
there was a missing module or not without having 02:48
a list of modules, you know, based on the script 02:48
that was used to compile the simulation. 02:48
Q. Does AMD -- did AMD keep a record of 02:48
all the simulations performed on the RTL of the 02:48
R400? 02:49
A. I don't think we kept a record of all 02:49
the simulations performed. I know that we 02:49
certainly kept a record of tests, whether a test 02:49

was passing or not passing, at any given time. 02:49
Whether that still exists, I think we 02:49
would have to go into the -- go into AMD's files 02:49
to really know. 02:49
Q. Did an RTL simulation performed on 02:49
the entire integrated circuit design of the R400 02:49
ever pass the simulation test? 02:49
MR. FAHRENKROG: Object to the form. 02:49
MR. PAIGE: Vague and ambiguous. 02:49
A. I can't answer definitively. I 02:49
expect that it did, but, you know, we had lots of 02:49
tests, and I'm sure that some of them passed. 02:50
But I can't, sitting here without reference, say 02:50
that an entire chip silicon with all the modules 02:50
being RTL as opposed to the C simulator. 02:50
I would have to go back and really 02:50
see what we could dig up as to what the state of 02:50
the chip was at any given date. 02:50
Q. If an RTL simulation performed on an 02:50
entire chip silicon with all the modules as RTL 02:50
passes the test, what is the usual next step in 02:50
the development process? 02:50
MR. FAHRENKROG: Object to the form. 02:50
MR. PAIGE: Object to the form. 02:50
A. Well, you have a lot of tests, and 02:50

you have both directed tests and you have random tests. So a lot of times, you know, we would focus on getting the random test passing, and then you would try to get -- sorry, you would try to get the directed test passing, because it would be easier to debug, and then you would run the random tests. 02:51

And only when you got to a suitable level of time where, you know, you weren't seeing bugs or the bugs you were seeing were ones you really didn't care about would you say, okay, well I guess we can tape it out. 02:51

Q. Was that point ever reached during the development of the R400? 02:51

A. No. We never got enough of the test passing in all the blocks to be able to tape out the R400. 02:51

Q. Go to Bates range 1451 through 1694. 02:52

A. Wait a second. Let me just make sure that I'm -- this is 278. Okay. I'm sorry, I need the numbers again. 02:52

Q. 1451. 02:52

A. Okay. Through? 02:52

Q. 1694. Let me know when you have it in front of you. 02:52

A. I'm just trying to get a little more organized here. All right. 1451. 02:53

Q. At the top of the page in the directory it says "R400_sp_snapshots." Do you see that? 02:53

A. Yes. 02:54

Q. What does that represent? 02:54

A. The R400 snapshots within the source code? I don't know. I'd have to -- I'm not sure how I could find out. I presume this was -- represents times when a backup was made and dated and perhaps kept in a more secure space. 02:54

Q. What does snapshot refer to? 02:54

A. With a view of the database as it existed at the time. 02:54

Q. Is it a view of the complete SP database as it existed on March 27, 2002? 02:54

MR. FAHRENKROG: Object to the form. 02:54

MR. PAIGE: Lack of foundation. 02:54

A. I don't know, but that would be my assumption. 02:54

Q. What is SP? 02:54

A. Shader pipe. 02:54

Q. Similar to the sequencer code that we were just referring to, if we walked through this 02:55

code in groups by dates, would you be able to determine whether this is the complete set of RTL code for the R400 shader processor? 02:55

A. No, I could not tell you that it was the complete set without reference to the scripts that were used to build it. 02:55

Q. Was RTL code ever created for the R400 primitive assembly? 02:55

A. I believe that it was, but I cannot say for sure. 02:55

Q. Was any RTL code ever created for the R400 scan converter? 02:56

A. Yes, I'm sure that that was -- I'd like to look at what tests were running to say absolutely for sure, but I believe that, yes, there was scan converter code. 02:56

Q. Was the RTL code for sequencer module ever integrated with the RTL code for the SP module? 02:56

A. Again, I'd like to see what tests we run, but I'm pretty sure the answer is yes. 02:57

Q. When was the RTL code for the sequencing and shader processor modules integrated? 02:57

A. Again, it would have to have been 02:57

prior to the end of 2002 certainly, and, you know, perhaps even before the end of 2001. 02:57

Q. What does it mean for the sequencer code to be integrated with the shader processor? 02:57

MR. FAHRENKROG: Object to the form. 02:57

A. We were able to run a simulation where both of those are RTL simulators and they are talking to one another. 02:57 02:58 02:58

Q. In order for an RTL simulation to be run, have the sequencer and the SP talking to each other, would that mean that the inputs into the shader processor and the outputs of the sequencer to the shader processor would have to match up? 02:58 02:58 02:58 02:58 02:58 02:58

MR. FAHRENKROG: Object to the form. 02:58

MR. PAIGE: Join in the objection. 02:58

A. Yes. You would have to -- you would be driving the inputs of one with the outputs of the other. 02:58 02:58 02:58

Q. What would happen if an RTL simulation was run and the inputs and outputs of the sequencer and shader processor did not correctly match up? 02:58 02:58 02:58 02:58

MR. PAIGE: Objection. Hypothetical. 02:58

A. It simply wouldn't run correctly. 02:58

You'd have unknowns on certain of the inputs, which would propagate fairly quickly. 02:59

Q. Could you look at the bottom stack of code that starts at 1695 and runs through the entirety down to 2110. 02:59

A. Okay. 02:59

Q. What does this collection of code relate to? 02:59

MR. FAHRENKROG: Objection to the form. 03:00

MR. PAIGE: 500 pages of code, what it relates to is being asked? 03:00

A. So at least that the first page is related to the sequencer. It's the arbiter within the sequencer. 03:00

Q. Is this RTL code? 03:00

MR. PAIGE: When you say this, do you mean all 500 pages of RTL code or are you asking some part of it? 03:01

A. Well, it appears to be a simulator code. 03:01

Q. Does this code relate to an emulator for the R400? 03:01

A. Yes. 03:01

Q. Sorry, I just want to be clear. Is 03:01

this code for the emulator R400 -- 03:02
A. Yes. 03:02
Q. -- or for the simulator -- 03:02
A. Oh, okay. Well, we often call the 03:02
emulator and simulator the same thing. But if 03:02
you mean the code that runs the RTL code, by 03:02
simulator -- if you mean the code that runs the 03:02
RTL code, no. This is not the code that runs the 03:02
RTL code. This is the C model emulator. 03:02
Q. Did a fully functional C model 03:02
emulator exist for the entire chip design of the 03:03
R400? 03:03
MR. FAHRENKROG: Object to the form. 03:03
A. I believe that it did, but I cannot 03:03
answer you 100 percent without doing some 03:03
additional research. 03:03
Certainly the emulator existed for 03:03
the shader system. 03:03
Q. The R400 emulator, did it operate in 03:03
a multithreaded environment or single-threaded 03:03
environment? 03:04
MR. PAIGE: Vague and ambiguous. 03:04
A. I couldn't answer you without doing 03:04
further research. 03:04
Q. Could the R400 emulator process two 03:04

threads in parallel? 03:04
MR. FAHRENKROG: Object to the form. 03:04
MR. PAIGE: Join in the objection. 03:04
A. I can't answer you without doing 03:04
further research whether we did multithread the 03:04
emulator or not. 03:04
(Gruber Exhibit 14, Document 03:05
entitled "R400 Emulator Version 0.0 January 15, 03:05
2002," marked for identification) 03:05
Q. You have been handed what's been 03:05
marked as Exhibit 14, document entitled "R400 03:05
Emulator," dated January 15, 2002, Bates label 03:05
AMD-ITC-0019672 through 19678. 03:05
Have you seen this document before? 03:05
A. Not to my recollection. 03:05
Q. Is this the same R400 emulator that 03:05
we have just been discussing with respect to the 03:05
code? 03:05
MR. FAHRENKROG: Object to the form. 03:05
A. Yes. 03:05
Q. If you look on Page 4 of the 03:05
document, Bates number 19675, if you look at 03:05
point 3 under Section 2. 03:06
A. Okay. 03:06
Q. Does this refresh your recollection 03:06

as to whether the emulator operated in a single thread or multithread? 03:06

A. Well, my recollection was that it did run in a single thread. My uncertainty was based on whether in a later revision of the emulator we tried to multithread it to increase performance. 03:06

Q. Under point 1 in Section 2, what does it mean by "The emulator contains no knowledge about the parallelism used in the hardware"? 03:06

A. The emulator is a functional emulator. It does not match the hardware cycle by cycle. It does match the hardware as certain checkpoints, generally block boundaries. 03:07

Q. I'm going to have you turn back to Exhibit 8, and at the same time -- 03:08

A. Let me put things back together. 03:08

Q. Do you still have them in order or do we need to take a break. 03:08

A. They're in order. I kept them in order. 03:08

MR. FAHRENKROG: Very sophisticated. 03:08

MR. HIGGINS: That's impressive. 03:08

Q. So if you go to the section of code with the date June 27, 2003. It's going to be in Bates ranges 279 through 318. 03:08

A. They're not in order. 279 to 318? 03:09
Q. Yes. 03:09
A. It's the 279 part that's missing. 03:09
All right. 03:09
Q. 279 through 318; do you have that? 03:09
A. I do. 03:09
Q. The next sections -- why don't we go 03:09
off the record for a second. 03:10
VIDEOGRAPHER: The time is 3:10 p.m. 03:10
At this time we'll pause recording and go off the 03:10
record. 03:10
(Recess) 03:15
VIDEOGRAPHER: The time is 3:19 p.m. 03:19
At this time we're on the record. We are 03:19
recording. 03:19
BY MR. HIGGINS: 03:19
Q. If you pull out Exhibit 8, the R400 03:19
sequencer spec. 03:19
A. Okay. 03:19
Q. If you turn to Page 9 of the 03:19
document, Bates number 1555. 03:19
A. Yes. 03:20
Q. I'm going to be referring to this 03:20
figure -- 03:20
A. Okay. 03:20

Q. -- and if you turn to the source code Page 303. It should be the module sq_input_arb.v. 03:20

A. Did you say 303? 03:20

Q. Yes. 03:20

A. Okay. 03:20

Q. Does this module, the sq_input_arb, correspond to the block labeled input arbiter on the diagram Exhibit 8? 03:20

MR. FAHRENKROG: Object to the form. 03:20

MR. PAIGE: Object to the form. 03:20

A. No. No, it doesn't. The input arbiter shown here is talking about input arbitration to the reservation stations. This particular module is arbitrating access to the GPRs when there's data to be written. So these are not the same. 03:21

Q. What RTL module corresponds to the input arbiter shown on Page 9 of Exhibit 8? 03:21

MR. FAHRENKROG: Object to the form. Are you asking about the RTL you've given him or something else? 03:21

MR. PAIGE: Foundation. 03:21

MR. HIGGINS: RTL in general. 03:21

Q. What RTL module of the R400 03:21

corresponds to the input arbiter shown on Page 9 of Exhibit 8? 03:21 03:22

MR. FAHRENKROG: Object to the form. 03:22

MR. PAIGE: Assumes facts not in evidence. 03:22 03:22

A. It would take me a certain amount of time. I don't know off the top of my head. I could -- I could try and go through this, and see if I can find it, but I do not know. 03:22 03:22 03:22 03:22

Q. Do you know the name of the unit or component that performs the arbitration into reservation stations? 03:23 03:22 03:22

MR. FAHRENKROG: Vague and ambiguous. 03:22

A. You're talking about the code component; is that what you're saying? The name of the module; is that what you're asking for? 03:22 03:22 03:22

Q. Either one at this point. 03:23

A. No, I don't. I don't -- you know, I see it called input arbiter here, and that's -- and I'm hoping it will be called that, although it isn't this one; I can tell you that. Let me see what this says. 03:23 03:23 03:23 03:23 03:23

The pages that you've pulled out appear to be the sequencer -- let me know when you want to go on to something else. 03:24 03:24 03:24

Q. While we have these sections pulled that you were just referring to, the one that starts on 756, sq_pix_thread_buff.v. 03:24

A. You're talking about some other section. 756. Okay. 03:24

Q. Does this module correspond to the pixel reservation station? 03:24

MR. FAHRENKROG: Object to the form. 03:25

MR. PAIGE: Vague and ambiguous. 03:25

A. Pix thread buffer. Yes, I believe that this corresponds to the pixel reservation station. 03:25

Q. If you look at the last section of code that you pulled out that begins Page 908, sq_vtx_thread_buff, is this RTL module the vertex reservation station? 03:25

MR. FAHRENKROG: Object to the form. 03:25

MR. PAIGE: Vague and ambiguous. 03:25

A. Yes, I think it is. 03:26

Q. From where do these two reservation stations receive inputs from? 03:26

A. According to this, it's from the ISM, which is probably the input state machine. 03:26

Q. What is the input state machine? 03:26

A. I don't know. I'm just reading from 03:26

the code. I don't think I was ever familiar in that detail with the sequencer. I certainly haven't remembered anything after the extent of time we're talking about. 03:26 03:27 03:27 03:27

I'm just looking for it here. So if I can find it, I don't know if I can. I think I have to analyze this code in fairly high detail to really track it back. 03:27 03:27 03:27 03:27

Here is something called the input state memory. 03:27 03:27

Q. What is the input state memory? 03:28

A. I'm just looking on Page 934. I don't know. I'd be looking for the input data. Let me look at the top level -- I mean, I think the way I would do this is if I had this code on a -- in a computer as opposed to printed out, is that I would search, starting with the inputs, and trace them through. 03:28 03:28 03:28 03:28 03:28 03:28

It's a lot harder based on paper to really say how data gets into those buffers. 03:28 03:28

Q. Can you tell by looking at these two reservation stations where they receive inputs from? 03:28 03:29 03:29

MR. FAHRENKROG: Object to the form. Vague. 03:29 03:29

A. Again, all I see is event and control packet input, initial values for state and status. It's from something called the ISM. So it's hard to go backwards like that. 03:29

Q. Is the input from the ISM the only input that the vertex and pixel reservation stations receive? 03:30

A. No. It receives -- there are other inputs here, where it's receiving control from the vertex cache and from the texture cache, as well as from the ALUs. 03:30

It also gets input from the execution arbiters and the return data from the -- there are a lot of inputs here. But those are more once the thread is loaded. 03:31

Q. Do either of the pixel or vertex reservation stations receive an input from an input arbiter? 03:31

MR. FAHRENKROG: Vague and ambiguous. 03:31

A. I cannot tell from this source code without further analysis whether this is an arbiter -- whether this output comes from an arbiter or not. 03:31

Obviously at one point there was an input arbiter based on this picture of May, 2002. 03:31

So I expect the answer is, yes, that they do come from an input arbiter. 03:31 03:31

And my recollection, which I cannot attest to with 100 percent, is that vertices were simply given priority on the input, but that was just a temporary thing, and that there was certainly a check to see whether space was available in the GPRs. 03:31 03:32 03:32 03:32 03:32

Q. You say vertices were given priority on the input. What input are you referring to? 03:32 03:32

A. I'm talking about -- I'm looking at this picture here on that input arbiter. You know, I'd have to go through this code in great detail to find where that arbitration is done. 03:32 03:32 03:32 03:32

Q. The vertex reservation station, for example -- 03:33 03:33

A. Yes. 03:33

Q. -- Page 908. 03:33

A. Okay. 03:33

Q. How many inputs are there to the vertex reservation station? 03:33 03:33

A. Again, there are all of these signals listed as inputs. I mean, some of these are outputs as well. 03:33 03:33 03:33

Q. Are you referring to lines 241 03:33

through 419? 03:34

A. Here are some inputs. I was actually 03:34
just looking at the wire definitions in the 03:34
inputs here. 03:34

So there are the parameters. Yeah, 03:34
so if you look here, there's all of those test 03:34
inputs. Then these ISM inputs, those are the 03:34
input state machines that I'm talking about. 03:34

But there are a whole bunch of other 03:34
inputs as well that I think are not part of 03:34
loading the thread, but just part of having a 03:34
thread in the buffer as it executes, and it gets 03:34
updated. 03:35

Q. Other than the inputs on lines 262 03:35
through 268, that begin with ISM, are there any 03:35
other inputs that relate to loading a thread in 03:35
the buffers? 03:35

A. I don't believe so. I think those 03:35
are the inputs. 03:35

Q. Do the ISM -- let me start over. 03:35
Lines 262 to 268, within those ISM modules, do 03:35
any of those contain a priority for vertices over 03:35
pixels? 03:35

MR. FAHRENKROG: Object to the form. 03:35

MR. PAIGE: Objection. Vague and 03:35

ambiguous. 03:35

A. These things would be after any priority in code, priority selection or priority arbitration would be done. So this would be the output of an arbiter. 03:35 03:35 03:35 03:36

Q. Is the ISM the arbiter? 03:36

A. The ISM might be the arbiter. I'd have to look at the ISM block, if there is a block labeled ISM. 03:36 03:36 03:36

Q. Do you know if there's a module labeled or titled ISM -- 03:36 03:36

A. I don't know. That's why I would have to analyze the code to figure out. 03:36 03:36

I mean, what I would do is I would take these ISM wires, these inputs, and trace them back to see what module produces them. 03:36 03:36 03:36

Q. If you look at pages 319 through 338. 03:36

A. Hang on. 03:36

Q. It's the directory of code we were looking at earlier. 03:36 03:36

A. Did you say 309? 03:37

Q. 319 through 338. 03:37

A. Okay. 319 through 338. 03:37

Q. Let me when you get it. 03:37

A. I'm there. 03:37

MR. PAIGE: Just to be clear, your representation is what Respondents printed off on the computer and not everything found on the computer, correct? 03:37

MR. HIGGINS: Yes, this represents the prints that Respondents made up through Bates number 1450. 03:37

MR. PAIGE: So it's not on the computer; you're not saying that? 03:37

MR. HIGGINS: No. This is just the -- actually, this is everything that's on the computer. This is not what was printed. This is everything that's on the computer. This is the directory listing of what's on the computer. 03:37

MR. PAIGE: Okay. So what you said before about Respondents printed is incorrect? 03:38

MR. HIGGINS: Respondents printed from this collection of code in the computer. This is a collection of code for the sequencer that is on the computer. 03:38

MR. PAIGE: Okay. So what you're saying is this is everything on the computer. The stuff we have been looking at here and here and here, this stuff that comes from the computer, but is not necessarily everything on 03:38

the computer? 03:38

MR. HIGGINS: Yes. The stuff we have been looking at is what Respondents chose to print. 03:38

MR. PAIGE: Okay. 03:38

Q. Mr. Gruber, in this directory listing of RTL code for the sequencer, where is the IMS? 03:38

A. ISM. 03:38

Q. ISM. 03:38

MR. FAHRENKROG: Object to the form. 03:38

MR. PAIGE: Join in the objection. 03:38

A. I'm looking. 03:38

MR. FAHRENKROG: I'd also like to state a general objection with respect to the questions that ask him to search through the source code. I think that's outside the scope of the subpoena on AMD to know every line of the source code set which has been set in front of him, which is over 2,000 pages or however many modules listed in this directory. 03:39

So I think you've had a lot of latitude with that questioning, but I do think it's outside the scope. 03:39

MR. HIGGINS: Well, we would disagree. 03:39

A. I do see one module called sq_pism.v. and another module called sq_vism.v. That would be my guess is where those things are. 03:39 03:39 03:40

Q. Where are you referring to? 03:40

A. I'm on Page 323 and 324. 03:40

Q. You said 324? Do you know what sq_pism refers to? 03:40 03:40

A. My guess is pixel input state machine, and "v" is vertex input state machine. As I said, I really never had strong familiarity with the details of this code as opposed to the functionality of it. 03:40 03:40 03:41 03:41

Q. Is sq_pism and sq_vism two separate input arbiters? 03:41 03:41

MR. FAHRENKROG: Object to the form. 03:41

A. I don't know. 03:41

Q. If you go to Page 439. Pull out 439 through 537. In the module sq_pism, is there an output labeled ism_rts? 03:41 03:41 03:42

A. I'm sorry, what are you looking for? 03:42

Q. The ISM outputs we were talking about. Does the pism module receive -- 03:42 03:42

A. I'm looking. I don't actually see them here. 03:43 03:43

Q. Does the pism include any arbitration 03:43

method between pixels and vertices? 03:44
MR. FAHRENKROG: Object to the form. 03:44
MR. PAIGE: Vague and ambiguous. 03:44
A. All I can say is I don't see any 03:44
arbitration between pixels and vertices in this. 03:44
Q. Would the pixel input state machine 03:44
receive vertices as an input? 03:44
A. I don't see any vertices received as 03:44
an input. 03:44
Q. Would the vertex input state machine 03:44
receive pixels as an input? 03:44
MR. FAHRENKROG: Objection. 03:44
Foundation. 03:44
A. I'd have to see that code. 03:44
Q. It's on Page 530. 03:44
A. I don't see any pixel-related input 03:45
into the vism. 03:45
Q. If the vism doesn't receive pixel 03:45
input, can it arbitrate between vertices and 03:46
pixels? 03:46
MR. FAHRENKROG: Object to the form. 03:46
MR. PAIGE: Join in the objection. 03:46
A. No. I don't see how. 03:46
Q. Does either the vertex input state 03:46
machine or the pixel input state machine perform 03:46

any arbitration between vertices and pixels? 03:46
MR. FAHRENKROG: Object to the form. 03:46
A. I don't believe so, based on my 03:46
current knowledge of this code, which is five 03:46
minutes old. 03:46
VIDEOGRAPHER: The time is 3:46 p.m. 03:46
At this time we'll end cassette number four. 03:46
We'll stop recording and go off the record. 03:46
(Recess) 03:56
VIDEOGRAPHER: The time is 3:57. At 03:57
this time we'll begin cassette number five. We 03:57
are on the record. We are recording. 03:57
BY MR. HIGGINS: 03:57
Q. Mr. Gruber, if you go back to the 03:57
page that begins on 319. It's the directory 03:57
listing. 03:57
A. Okay. 03:57
Q. In this directory listing of RTL 03:57
modules in the sequencer, can you identify any 03:57
module that performs arbitration between vertex 03:57
and pixel threads to the reservation stations? 03:57
MR. FAHRENKROG: Object to the form. 03:58
And again, outside the scope of the topics, 03:58
asking him to search through source code for an 03:58
answer to your question. 03:58

MR. PAIGE: Actually, searching the directory listings, not even the source code itself. 03:58

A. No. I'd have to analyze the code, probably backwards, from the thread buffer or the pixel buffer to say whether such arbitration exists or not, and if so, what module it appears in. 03:58

Q. Sitting here today, can you identify any RTL module in the R400 that arbitrates between vertex and pixel threads as an input to the reservation stations? 03:58

MR. FAHRENKROG: Same objection. 03:59

A. Sitting here today, I cannot. I can tell you how I would go about finding whether such a module does exist, and it would involve tracing the inputs and outputs back. But it's very hard to do with a lot of paper. 03:59

Q. If you go to Page 286. It's going to be in that big stack to your right most likely. 03:59

A. Why is it not in one of these? 286 is gone. Where did it go? 04:00

Q. You may have pulled a stack earlier that started at 279. 04:00

MR. FAHRENKROG: I'll also state an 04:00

objection that requiring the witness to navigate over 2,000 pages of code in random places, nine or more places, is very burdensome for the witness. 04:00

A. This is 900, 756, 319, 430 -- I'm sorry, what page? 04:00

Q. 286. 04:01

A. Let me see if I for some reason put it way down here. 04:01

MR. HIGGINS: Do you want to go off the record to help him find it? 04:01

MR. FAHRENKROG: Sure. 04:01

VIDEOGRAPHER: The time is 4:01. At this time we'll pause recording and go off the record. 04:01

(Discussion off the record) 04:02

VIDEOGRAPHER: The time is 4:02. We are on the record. We are recording. 04:02

BY MR. HIGGINS: 04:02

Q. Mr. Gruber, do you have the source code pages 286 through 302? 04:02

A. Yes, I do. 04:02

Q. What is the function of sq_gpr_alloc module? 04:02

MR. PAIGE: Object to the form. 04:02

MR. FAHRENKROG: Object to the form. 04:02

A. This allocates space in the general purpose register file as threads are allocated. 04:03

Q. Does the GPR alloc module include any priority? 04:03

MR. FAHRENKROG: Object to the form. 04:03

MR. PAIGE: Vague and ambiguous. 04:03

A. There is some priority associated with this dynamic mode in the sense that based on the dynamic mode, it chooses how to move the head and tail pointer boundary. 04:03

If you look on Page 293, there's where that code starts. 04:03

Q. How does the dynamic mode utilize a priority? 04:04

A. Well, it's determining whether to allocate more space for pixels or space for vertices via this pointer control state machine. So it does that. 04:04

I believe the inputs are based on how busy the output buffers are associated with either vertices or pixels. 04:04

Q. Which inputs are based on how busy the output buffers are? 04:04

A. The inputs into the boundary control 04:04

mechanism, that is, the boundary between pixels and vertices in the GPR register file. 04:05

Q. How is it determined whether to allocate more space for pixels or vertices? 04:05

A. I'm just trying to find it in the codes. Since I have the code here, I don't want to -- well, it uses the amount of space in the output buffers to decide whether we're more pixel limited or more vertex limited. 04:05

I know that from the high-level documentation. I'd have to analyze this code to really point out where in the code that is done. 04:06

Q. If the GPR alloc module is not operating in the dynamic mode, is it operating in static mode? 04:06

A. Yes. 04:06

Q. When it operates in static mode, does it utilize any priority? 04:06

MR. PAIGE: Object to the form. Vague and ambiguous. 04:06

A. No, it does not. 04:06

MR. FAHRENKROG: Counsel, is it okay for the witness to move the yellow paper out of the way to clear space for the next exhibit? 04:07

MR. HIGGINS: Yes. Your choice. 04:07

(Gruber Exhibit 15, U.S. Patent No. 6,897,871, marked for identification)

(Gruber Exhibit 16, U.S. Patent No. 7,239,322, marked for identification)

Q. You have been handed what's been marked as Exhibits 15 and 16. Exhibit 15 is U.S. Patent No. 6,897,871. Exhibit 16 is 7,239,322.

Have you seen these exhibits before?

A. Yes.

Q. Are you listed as an inventor on both of these patents?

A. Yes, I am.

Q. The register file allocation we were just discussing in the GPR and alloc module, is that functionality disclosed in these patents?

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. Also overly broad.

A. Which functionality were you talking about?

Q. The register allocation. In particular, the dynamic register allocation.

A. No, these don't cover the dynamic register allocation.

Q. Do either of these patents describe how threads are created?

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. Overly broad. 04:10

MR. PAIGE: Vague and ambiguous. 04:10

A. Well, this patent does talk about thread creation. I'm looking at arbiter here 101 for loading threads. 04:10 04:11 04:11

Q. Do either patent -- does either patent discuss a priority related to thread creation? 04:11 04:11 04:11

MR. FAHRENKROG: Same objection as to form. Legal conclusion. Overly broad. 04:11 04:11

MR. PAIGE: Join in the objections. 04:11

A. I don't see anything in here about priority on thread creation. Although, I do see lots about priority on thread execution. No, I take that back. 04:13 04:13 04:13

If I look at Line 65, Column 3, on the '322 patent, "one embodiment an input arbiter provides to the command threads to each of the first reservation station, and the second reservation station 304 based on whether the command thread is a pixel command thread or vertex command thread." 04:13 04:13 04:14 04:14 04:14

Q. Is that a priority or is it just sending vertex threads to the vertex reservation 04:14 04:14

station and pixel threads to the pixel reservation station? 04:14

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. 04:14

MR. PAIGE: Argumentative. 04:14

A. I'm not sure what the legal interpretation of this would be. 04:14

Q. What is your interpretation of this? 04:14

A. Well, the fact that it's an arbiter suggests to me that there is some arbitration. I expect that the arbitration is such that even though a pixel thread has to go to a pixel station and a vertex thread has to go to a vertex station, there's an arbitration simply due to time when that happens. 04:14 04:15

In other words, the reason why you would call it an arbiter here is because there's only one set of wires that come out, and that's either loaded into the pixel station or the vertex station. 04:15

So you may have done that either to save wires or to save control logic where you can just -- you can just load one thread in a given cycle, and that's why you would do an input arbiter as opposed to just having two entirely 04:15

separate paths. 04:15

Q. Is there any disclosure in the patent of how this arbiter receives the command threads? 04:16

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. Also this line of questioning about the disclosure of these patents is outside the scope of the deposition topics. 04:16

MR. PAIGE: Join in the objections. 04:17

A. Let me just look at the picture that this refers to. It says "not illustrated in Figure 4." Well, that's helpful. 04:17

I don't see anything further referenced in this paragraph. 04:17

Q. The arbiter that you're referring to in Figure 4, is that arbiter 306? 04:18

A. No, it's not, because Line 66, I guess, says "although not illustrated in Figure 4." 04:18

Q. So there's no input arbiter described in this specification? 04:18

MR. FAHRENKROG: Object to the form. Asked and answered. Legal conclusion and outside the scope of the topics. 04:18

MR. PAIGE: And misstates prior testimony. 04:18

A. I pointed to the language on page -- on Column 3, Line 67 that does talk about an input arbiter. 04:18 04:18 04:19

Q. Does that input arbiter that you're referring to utilize any priority with respect to vertex and pixel threads? 04:19 04:19 04:19

MR. FAHRENKROG: Same objection. 04:19

A. Well, it's an arbiter, and it has to arbitrate on some priority basis. As I say, this patent does not seem to expand on how that arbiter works. 04:19 04:19 04:19 04:19

Q. Does an arbiter have to use a priority? 04:19 04:19

MR. FAHRENKROG: Object to the form. 04:19

A. Well, it doesn't have to use a fixed priority, but it has to arbitrate. So it has to select one, and the one who he selects has higher priority. At least that's how I interpret the word "arbiter." 04:19 04:19 04:19 04:19 04:19

Q. What do you mean by "the one who he selects has higher priority"? 04:19 04:19

A. I mean, an arbiter selects in between one or more inputs, and I would interpret the one that he selects as being the one who had higher priority, at least for that cycle, even if he's 04:19 04:20 04:20 04:20

selecting on a random basis, then on that cycle, whoever won that spin of the dice, would have higher priority. 04:20

Q. If the arbiter just goes back and forth between two inputs, is that a priority? 04:20

MR. FAHRENKROG: Object to the form. Vague. Incomplete hypothetical. 04:20

A. I would say on any given cycle, someone has priority, and if you didn't have priority in the last cycle, then you'll have priority on this cycle. 04:20

Often, like a round robin arbiter, you would talk about who has priority on this cycle. 04:20

Q. Exhibit 15, the '871 patent. Does the '871 patent describe a multithreaded graphics processor or a single-threaded graphics processor? 04:20

MR. FAHRENKROG: Object to the form. Vague. Calls for a legal conclusion. 04:21

MR. PAIGE: Join in the objections. 04:21

A. It's a multithreaded processor. 04:21

Q. What leads you to believe that it is a multithreaded processor? 04:22

A. Well, I'm looking at Column 5, Line 04:22

14. It says, "The unified shader 62 has the ability to simultaneously perform vertex manipulation operations and pixel manipulation operations at various degrees of completion by being able to freely switch between such programs or instructions, maintained in the instruction store very quickly." 04:22

Q. Can it perform those operations without operating on threads? 04:22

MR. FAHRENKROG: Object to the form. Vague. Also outside the scope of this deposition. 04:22

MR. PAIGE: Join in the objections. 04:23

A. Can it do it without threads? Well, no, because the operations are in various degrees of completion. So you have to track where you are. 04:23

So that implies that you have -- that you're multithreading; at least it does to my mind. 04:23

Q. Does the '871 patent mention threads anywhere? 04:23

MR. FAHRENKROG: Object to the form. 04:23

A. I can't give you a definitive answer without spending time looking at every single 04:23

point. I don't see the word "thread" anywhere in the -- off the top of my head or in my quick glance. 04:23

Q. Can a unified shader operate without threads? 04:24

MR. FAHRENKROG: Object to the form. Overly broad. 04:24

MR. PAIGE: Vague and ambiguous. 04:24

A. I think you could design such a unified shader. It wouldn't be a good unified shader. 04:24

Q. What did you do to prepare for today's deposition? 04:24

A. I went over numerous versions of top-level specs and block-level specs. I clarified in my mind the sequence of chips that ATI produced. 04:25

I looked over the patent -- the patent applications to the follow-on patent application that I had signed. I looked over the code attached to those, as well as obviously meeting with my lawyers. 04:25

Q. Which top-level and block-level specs did you review? 04:25

A. I looked at the sequencer 04:25

specification, I looked at the shader specification. I looked at the R400 -- I don't know whether it was called the overview spec or it was the initial spec produced by Steve Morein. I looked at various versions of that. 04:25 04:26 04:26 04:26 04:26

Q. You said you met with your lawyers. Which lawyers are you referring to? 04:26 04:26

A. Both Qualcomm lawyers and AMD lawyers. 04:26 04:26

Q. Did you meet with Qualcomm and AMD lawyers at the same time? 04:26 04:26

A. No, I did not. 04:26

Q. Did you talk to anyone at AMD to prepare for today's deposition? 04:26 04:27

A. No. 04:27

Q. Did you talk to anyone at Qualcomm to prepare for today's deposition? 04:27 04:27

A. No one beyond the legal department. I mean, I spoke to both in-house Qualcomm counsel as well as legal support. But that was more in just knowing where to go. 04:27 04:27 04:27 04:27

Q. Since you left AMD, do you still communicate with any employees of AMD? 04:27 04:27

MR. FAHRENKROG: Objection. Outside the scope of the AMD topics. 04:27 04:27

A. Occasionally I will have lunch or an occasional dinner with AMD and ex-AMD employees. 04:27 04:28

Q. Which ex-AMD employees do you meet with? 04:28 04:28

A. I had dinner with Bob Feldstein a couple months ago. I occasionally have lunch with Paul Mitchell who's at NVIDIA now. 04:28 04:28 04:28

Q. Which current AMD employees have you met with? 04:28 04:28

A. I occasionally have lunch with Mark Fowler. 04:28 04:28

Q. Anyone else? 04:28

A. It has been over a year, but I had lunch with Phil Rogers about a year ago. 04:28 04:28

Q. Who is Phil Rogers? 04:29

A. Phil is CTO, I believe, at AMD software CTO. 04:29 04:29

Q. Did you speak with any of these people about the topics of your deposition today? 04:29 04:29

A. The last time I spoke to Bob Feldstein, I asked him whether he thought -- I think I said something like, Can you guys come up with better patents to advocate than this, and he said, Yeah, maybe they're not that good. That was the extent. 04:29 04:29 04:29 04:29 04:29

Q. Why did you ask that? 04:29
MR. FAHRENKROG: Objection. Outside 04:29
the scope of the AMD-related topics on this line 04:29
of questioning. 04:29
A. I was surprised that NVIDIA didn't 04:30
have stronger patents. 04:30
Q. Why is it your view that these are 04:30
not strong patents? 04:30
MR. FAHRENKROG: Same objection. The 04:30
witness is testifying as an individual, not on 04:30
behalf of AMD. 04:30
MR. PAIGE: And in giving an answer, 04:30
please do not reveal the content of any 04:30
communications you've had with Qualcomm counsel. 04:30
A. I just didn't think they were very 04:30
fundamental. 04:30
Q. Why did you not think they were very 04:30
fundamental? 04:30
MR. FAHRENKROG: Same objection. 04:30
MR. PAIGE: Same caution. 04:30
A. The particular -- the particular 04:30
aspects of the patents were having to do with 04:30
lighting and, you know, special lighting 04:30
processors, which I thought nobody used any more. 04:31
That was the main issue. 04:31

Q. Which patents are you referring to? 04:31
A. I don't remember. I just glanced at 04:31
these patents to see what the general subject 04:31
was, and there was one patent involved in 04:31
prosecution that had to do with a specialized 04:31
lighting processor. 04:31
Q. What do you mean there was one patent 04:31
involved in the prosecution? 04:31
A. I believe there were five or six 04:31
patents asserted against Qualcomm and Qualcomm's 04:31
customers, and one of them had to do with the 04:31
lighting. 04:31
Q. When you said involved in 04:31
prosecution, were you referring to the patents 04:31
asserted in this case? 04:31
A. Yes. 04:31
Q. When you made that comment to Bob 04:31
Feldstein, was that just about this one patent or 04:31
was it about all seven asserted patents? 04:31
A. I said it that I didn't see any of 04:31
them that were particularly strong. 04:32
Q. Why, in your view, were they not 04:32
particularly strong? 04:32
MR. FAHRENKROG: Restate the 04:32
objection that the witness is testifying as an 04:32

individual, not as a corporate representative of AMD on this topic. 04:32

A. Just didn't -- I just didn't see anything that would really apply to most existing graphics processors, most modern ones. 04:32

Q. So who, in your view, has stronger graphics patents than NVIDIA? 04:32

MR. FAHRENKROG: Same objection. 04:33

A. I don't know. I mean, you know, I don't know that anyone does, but I thought that NVIDIA's would be stronger. That's all I can say. Maybe no one has a strong patent. 04:33

Although, I like this one, the one with my name on it. 04:33

Q. Which one are you referring to? 04:33

A. The unified shader. 04:33

Q. Which one is that? 04:33

A. '871. 04:33

Q. Why do you like that one? 04:33

A. Because I think it covers a fundamental aspect, because I think unified shading is important, and I think this covers it. 04:33

Q. What fundamental aspect of unified shading does the '871 -- 04:33

A. I said it covers unified shading, 04:33

which I think is fundamental to modern processors. 04:33

Q. When did you meet with your AMD lawyers prior to today's deposition? 04:33 04:34

A. Yesterday. 04:34

Q. How long did you meet with AMD lawyers? 04:34 04:34

A. Six hours. 04:34

Q. What did you discuss during those six hours? 04:34 04:34

MR. FAHRENKROG: Objection. Calls for a privileged communication between counsel for the non-party AMD and former employee, which is privileged. So I instruct you not to answer that question. 04:34 04:34 04:34 04:34 04:34

A. I'm not going to answer that question. I could have figured that one out by myself, though. 04:34 04:35 04:35

Q. When did you leave AMD and go to Qualcomm? 04:35 04:35

A. Early 2009. 04:35

Q. Why did you leave AMD and go to Qualcomm? 04:35 04:35

A. I went to Qualcomm as part of the mobile team, an effort that I had been involved 04:35 04:35

with. 04:35
MR. HIGGINS: Let's take a quick 04:35
break. 04:35
VIDEOGRAPHER: The time is 4:35 p.m. 04:35
At this time we're stop recording and go off the 04:35
record. 04:35
(Recess) 04:36
VIDEOGRAPHER: The time is 4:40 p.m. 04:39
At this time we are on the record. We are 04:40
recording. 04:40
EXAMINATION BY 04:40
MR. CHATTERJEE: 04:40
Q. Hello, Mr. Gruber. I introduced 04:40
myself earlier. I'm Neel Chatterjee, and I 04:40
represent NVIDIA in this case. 04:40
I want to ask you some follow-up 04:40
questions from my colleague's questioning, and I 04:40
want to start with where he ended which was 04:40
talking about your review of NVIDIA's patents. 04:40
When did you do that? 04:40
A. I'm sorry, when did I look over 04:40
NVIDIA's patents? 04:40
Q. Yes. 04:40
A. The day that they became public about 04:40
this lawsuit. So I don't know when that was, but 04:40

it's probably a couple months ago. 04:40
Q. So you looked at it the day that 04:40
Qualcomm was sued? 04:40
A. I don't know if it was the day that 04:40
Qualcomm was sued. The day it was publicly 04:40
announced. 04:41
Q. Okay. Other than your counsel, have 04:41
you talked with anyone else, other than Mr. 04:41
Feldstein, about these patents? 04:41
A. No. As I say, I really didn't talk 04:41
to Bob about the patent. I just kind of gave him 04:41
my opinion of the value of the patent. We didn't 04:41
talk about any technical aspect of the patent. 04:41
Q. Did you study the claims in the 04:41
patents? 04:41
A. I just -- I wouldn't say that I 04:41
studied them. I just looked at them. I read 04:41
through the patents. 04:41
Q. And did you read the claims? 04:41
A. Yes. 04:41
Q. And formed opinions about them? 04:41
MR. PAIGE: Object to the form. 04:41
A. Did I form any opinions about them? 04:41
I generally judge them on like, you know, whether 04:41
I can think of areas where there's infringement 04:41

and stuff that I've worked on. I didn't see anything that I was overly worried about. 04:41 04:42

Q. Had you seen any of them before? 04:42

A. Any of those patents or any of those claims? 04:42 04:42

Q. Yes, any of those patents or any of those claims. 04:42 04:42

A. I don't think so. 04:42

Q. Were you aware that Exhibit 9 and 10, the 131 declarations that you submitted, were actually in response to an assertion by the examiner that the patents you got were invalid, and this was an attempt to swear behind the dates of those patents? 04:42 04:42 04:42 04:42 04:42

MR. FAHRENKROG: Object to the form. Calls for a legal conclusion. 04:42 04:42

A. No, I wasn't fully aware that was the reason for this. 04:42 04:42

Q. Would it surprise you to learn that? 04:42

MR. FAHRENKROG: Object to the form. 04:42

A. Surprise me? Maybe a little bit. 04:42

Q. Why would it surprise you a little bit? 04:42 04:42

MR. FAHRENKROG: Object to the form of the question. 04:42 04:42

A. Simply because I know -- my feeling was that ATI was significantly in advance on the development of a unified shader. 04:42 04:43 04:43

Q. Why is that your feeling? 04:43

A. Because I know that we had been working on one for a long time, and then when Xenos came out, you know, we were ahead of -- that was the first unified shader on the market, and at the time the Xenos came out, NVIDIA had been poopooing the whole idea of a unified shader. 04:43 04:43 04:43 04:43 04:43 04:43

Q. Who at in NVIDIA was poopooing the idea of a unified shader? 04:43 04:43

A. I just know what I read in things like AnandTech. 04:43 04:43

Q. So it was kind of by reading kind of the marketing -- 04:43 04:43

A. Yeah, the marketing. I'm not talking about talking to any individual. 04:43 04:44

Q. Kind of like the industry pundits? 04:44

A. Right. 04:44

Q. Is that right? 04:44

A. Yes. 04:44

Q. Are you proud of the unified shader technology you helped create? 04:44 04:44

A. Yes. 04:44

Q. And how many patents do you have? 04:44

A. I don't know the exact number. It's probably 40, 50. 04:44

Q. How many of them relate to this unified shader architecture? 04:44

A. A handful, I believe. Not more than that. 04:44

[redacted]

[redacted]

Q. Do you have any interest in the outcome of this case? 04:45

MR. FAHRENKROG: Vague and ambiguous? 04:45

A. No direct interest. Kind of rooting for Qualcomm, but I don't have any monetary interest. 04:45

Q. You're rooting for Qualcomm because you work there? 04:45

A. Yes. 04:45

Q. You made a comment when Mr. Higgins was talking to you, you told Mr. Feldstein you were surprised that NVIDIA did not have stronger patents. Why was that surprising to you? 04:45

A. NVIDIA has been in the graphics field a long time, and they were first in many things in the graphics field. So I just would have expected that, you know, maybe they had some killer patents. 04:46

Q. Can you give me some examples of things where they were first? 04:46

A. I think they were first in probably a vertex shader. Other areas they were first? I don't know. It's a long time. I'm trying to think back. 04:46

There's a long history of graphics there. They would have been first at least in the PC world in a lot of the floating point aspects of graphics. 04:46 04:46 04:47 04:47

Q. Anything else you can think of? 04:47

A. Not off top of my head. 04:47

Q. Would you say NVIDIA is an innovative company? 04:47 04:47

A. Yes. 04:47

Q. I want to go to the very beginning of your deposition with Mr. Higgins. You had talked about -- I think the term you used was evolution, that you talked about going from the R400 to the R600, and you said there was an evolution. Do you recall that? 04:47 04:47 04:47 04:47 04:47

A. Yes. 04:47

Q. Why were the product numbers changed? 04:47

A. Well, because we generally decide on product numbers aimed at market windows, and we like to have a sequence of product numbers. 04:47 04:47 04:47

So, you know, when you come out with a brand new architecture, you want to have a brand new product number associated with it, and it would have been awkward, for instance, to have the West Coast team do an R300 and then do an 04:47 04:47 04:48 04:48 04:48

R500 with a missing R400. 04:48

I'm just speaking for the marketing 04:48
group. It's a guess on my part. But that was 04:48
kind of how we did both marketing and 04:48
engineering. 04:48

Q. Now, was the R420 part of this 04:48
evolution? 04:48

A. It was part of an evolution of the 04:48
R300 series. 04:48

Q. Was it part of the evolution from the 04:48
R400 to the R600? 04:48

A. No, it was not. 04:48

Q. So do you know why it was released as 04:48
the R420? 04:48

A. Well, because they wanted something 04:48
slotted into something called the R400 series. 04:48
They didn't want to use the word R400, because 04:48
that would lead to internal confusion. So they 04:48
picked something that was R4, but not R400. 04:49

Q. Who is the "they"? 04:49

A. I don't know. 04:49

Q. It was lore within the company? 04:49

A. Hmm? 04:49

Q. It was lore within the company? 04:49

A. Yes. I don't know where the name 04:49

came from, but this is the reason. 04:49

Q. Now, this term "evolution," what do you mean by the term "evolution"? 04:49

A. Well, any design as you work on it, you learn things about the design. You learn things during the simulation process, what you find out, where there are bottlenecks that you didn't know about otherwise. 04:49

You also, as the time goes on, the requirements for the part have changed. Something that was fast enough if released at time end is not fast enough if released at time end plus six months. 04:49

The technology changes such that, you know, trade-off between how much memory costs and logic costs and the overall size of the part that you're trying to build. So those types of things. 04:50

Q. And with those requirements changing, does that cause alteration of the designs? 04:50

A. Yes. 04:50

Q. Now, how long a time period was this for the evolution of the R400 to the R600? 04:50

A. I don't recall off the top of my head when the R600 shipped, because I didn't review 04:50

that as part of this -- my preparation here. But the R200 -- sorry, the R400, you know, changed to R500 in early 2003. So I think the R600 was about three years later in terms of when it finally shipped. 04:50 04:50 04:50 04:50 04:50

Q. So when was the R400, when did the development of it start? 04:50 04:51

A. The development -- the engineering -- the architectural effort started in late 2000, and then over 2001, it moved from a pure architecture effort to an architecture and implementation effort. 04:51 04:51 04:51 04:51 04:51

Q. So I just want to make sure I have the timing right. So the R400, the engineering effort, started around late 2000? 04:51 04:51 04:51

A. No -- well, if you include architecture effort, yeah, as opposed to, you know, people actually working coding RTL. People working coding RTL started later than that. I would say, you know, probably in mid-2001. 04:51 04:51 04:51 04:51

There's an architecture section of a schedule, and then there's, you know, implementation. It may overlap, but architecture leads. 04:51 04:51 04:51 04:51

Q. When was the R400 moved into the R500 04:51

again? 04:51

A. Early 2003. 04:51

Q. Early 2003. Then did you say two or three years after that the R600 development started? 04:51 04:52 04:52

A. No, I didn't mean that's when it started. I was thinking about when it was done. 04:52 04:52

Q. Okay. So it was done -- the R600 was done two or three years later? 04:52 04:52

A. Yes. I would say probably about three years later, because there was an R500 in between. 04:52 04:52 04:52

Q. And was there an R500 that was actually commercially released? 04:52 04:52

A. No -- well, again, similar to the R420, there was an R500 part in that time frame that was unrelated to this R500 effort. 04:52 04:52 04:52

Q. So was the first product that had the unified shader architecture that was commercially released the R600? 04:52 04:52 04:52

MR. PAIGE: Objection. Misstates prior testimony. 04:52 04:52

A. No, because the Microsoft Xbox had the unified shader architecture very closely related to what we were working on in the R500. 04:52 04:53 04:53

Q. And that was the Xenos chip? 04:53
A. The Xenos chip. 04:53
Q. When about was that? 04:53
A. That was like for the holiday season in 2005. 04:53
Q. So it would have been commercially released when, six months earlier? 04:53
A. Well, I don't know what you mean by -- it was, you know, a part for one customer. So it was never really offered for sale, except to that one customer. 04:53
Q. I understand that. But when you say for the holiday season, obviously AMD or ATI would have had to sell it before -- 04:53
A. Oh, you're saying when it went on sale? 04:53
Q. To Microsoft. 04:53
A. Oh, to Microsoft. Yeah. I wish I knew. It had to be at least six months prior to that. 04:53
Q. You're not certain exactly how long? 04:53
A. I'm not certain. 04:53
Q. Do you know why the R400 did not meet its anticipated release date? 04:54
A. The schedule just slipped. It was a 04:54

more complicated part than we thought. We had some -- we had some issues. I don't recall the shader being the real issue, because, you know, we started that early, and we did a lot of engineering on that. 04:54

I know, as I said earlier, that the long pull was the RB, the render back end. That just was not coming together. 04:54

Q. So the render -- it's your recollection that for the R400 the render back end was the big obstacle to the release? I need you to say with a yes or no. 04:54

A. Yes. I'm sorry. 04:54

Q. I always say there's two important people in the room, the witness and the court reporter. If you don't say it, the court reporter can't get it down. And that's just not because I'm married to a court reporter. 04:55

I want to talk -- let's see. Would the RTL code for the R400 be the best indicator of what had been done on the chip? 04:55

MR. FAHRENKROG: Object to the form. 04:55

A. Yes. The RTL code defines what was done on the chip. 04:55

Q. A couple small questions. How long 04:55

did you meet with AMD counsel? I don't want to know what was said, just how long was it? 04:55

MR. FAHRENKROG: Objection. Asked and answered. 04:55

A. About six hours. 04:55

Q. How about Qualcomm counsel? 04:56

A. Yesterday, because I had a deposition a couple weeks ago as well? 04:56

Q. Right, yesterday. 04:56

A. Maybe two hours. 04:56

Q. Do you know someone named Michael Doggett? 04:56

A. I do. 04:56

Q. Was Mr. Doggett involved in the R400? 04:56

A. Yes. 04:56

Q. What was his involvement? 04:56

MR. FAHRENKROG: Objection. Asked and answered. 04:56

A. I remember him -- Michael did a number of things, but I specifically remember him working on the texture pipe, as well as I mentioned that he was involved in some academic relations. 04:56

Q. You mentioned someone named Laurent earlier today, saying that he wrote much of the 04:56

RTL code. Do you remember talking about that? 04:56
A. Yes. 04:56
Q. And what was his involvement in the 04:56
development of the R400? 04:56
A. He was involved in the architecture 04:56
fairly early on. He was hired in to work on the 04:56
architecture of the R400. 04:56
One of the things that happens in a 04:56
development cycle is that you have people who are 04:57
busy, because they're getting the current chip 04:57
out. So you have this kind of overlapping 04:57
development. And in the early part, it's hard 04:57
getting a lot of people working on it. So you 04:57
are working with a small group of architects. 04:57
Laurent had the advantage, he didn't 04:57
have a previous project to work on. 04:57
Q. Did you talk to him at all to prepare 04:57
for your deposition today? 04:57
A. I did not. 04:57
Q. Why not? 04:57
MR. FAHRENKROG: Object to the form. 04:57
A. He just simply wasn't offered to me 04:57
to talk to. 04:57
Q. You were talking earlier today about 04:57
a register file versus a preservation station. 04:57

Do you remember talking about that -- 04:57
A. Yes. It was a reservation station. 04:57
Q. Reservation station. I'm sorry. So 04:57
I'm just going to ask you a really basic station. 04:57
What is a register file? 04:58
A. So what we meant by the register file 04:58
-- a register file is generically just a bunch of 04:58
storage. 04:58
But what we meant by the general 04:58
purpose register file was the storage associated 04:58
with a thread where the thread would keep all of 04:58
the temporary data that it calculated while it 04:58
was producing its final output. 04:58
Q. And a reservation station, what is 04:58
that? 04:58
A. So a reservation station is also 04:58
associated with a thread, but rather than holding 04:58
data that's produced by the shader, it holds all 04:58
the control information associated with the 04:58
thread; for instance, the current program counter 04:58
associated with the thread. 04:58
Q. Mr. Higgins asked you about a vertex 04:58
cache. Do you remember talking about that? 04:59
A. Yes. 04:59
Q. Do you remember the first time AMD or 04:59

ATI came up with the idea of adding a vertex cache? 04:59

MR. PAIGE: Object to the form. Vague and ambiguous. 04:59

A. No. We had vertex caches prior to the R400. The issue was whether to add a vertex cache to the R400. 04:59

I know that that came up somewhere in the R400 design phase, and we tried a number of -- well, one, we didn't know whether we really needed it. It wasn't needed for functionality. 04:59

And as a matter of fact, we later produced parts -- I say "we," in this case, Qualcomm has produced parts based on the R400 database. It doesn't have a vertex cache. It's more of a performance issue, and whether we needed it for performance. 05:00

So we kind of fought it in the R400, and finally gave in at some point. 05:00

Q. Do you know if when you gave in, it was at the R400 or the R500 stage of things? 05:00

A. My recollection, which may not be perfect, was it was part of the R500 changes, between the R400 and the R500. 05:00

Q. Now, you referred in the early part 05:00

of the deposition to a C model emulator. What is a C model emulator? 05:00

A. So what I mean by that is a model of the chip that's in -- that's written in C or C++ that is functionally the same as the chip; that is, it produces the same outputs, and it matches at certain checkpoints within the chip. 05:00 05:01 05:01 05:01

Q. And the C model is something that operates on the simulation software, right? 05:01 05:01

A. It is simulation software. It just runs. You compile it and it's a program like anything else. It's not built on top of anything. 05:01 05:01 05:01 05:01

Q. Do you recall what simulation program you used? 05:01 05:01

A. Well, this is what I'm saying. The C model doesn't need a simulation program to run. It is a program, and it runs just like Microsoft Word runs. 05:01 05:01 05:01 05:01

Q. So you don't need something like VCS or anything like that? 05:01 05:02

A. There is something else like VCS that takes in the RTL as the input, and that simulates that. 05:02 05:02 05:02

MR. PAIGE: Could I just register an 05:02

objection here, Neel. I understand we had an agreement in the case that we would have different questioners asking questions based on different patents, different areas, things like that. 05:02

It appears what's happening now is you're asking a bunch of follow-up questions to questions Mr. Higgins asked. I'm not sure why we need two questioners for that. I object to it, but you can continue. 05:02

Q. In Exhibit 9 and 10 of the deposition, you had these declarations that you submitted and you referred to RTL simulation in the declaration. Do you remember that? 05:02

A. Yes. 05:02

Q. Did you use a particular RTL simulator? 05:02

MR. FAHRENKROG: Objection. Asked and answered. 05:02

MR. PAIGE: Objection. Asked and answered. 05:02

A. We did. I don't recall the ones that we used, and we might have used more than one during the course of the project. 05:02

Q. Did you use a hardware emulator? 05:02

A. We might have. I know that we were doing things with those. I don't recall whether the R400 was ever ported to a hardware emulator.

Q. Do you know if there was ever a physical representation of a chip for the R400 ever made?

MR. FAHRENKROG: Objection. Vague. Asked and answered.

A. There was no chip made on the database called R400. There was chips made on databases that were based on the R400, the R500, and the Xenos came pretty soon after the change from R400 to R500.

Q. Was there a way to assign vertex data and pixel data to threads based on allocation priority between the two?

MR. PAIGE: Objection. Vague and ambiguous.

MR. FAHRENKROG: Object to the form.

A. The thread pools themselves, that is the control mechanism, was separate. There was a dynamic mechanism for sharing the GPRs that was based on a priority that was kind of heuristically developed in that part of the RTL, based on looking at how busy the output buffers,

be it the vertex and pixel buffers, was. 05:04

So if there weren't enough vertices 05:04

being generated, it would allocate more room for 05:04

the vertices. And at least in some version of 05:05

the code, there was an input thread arbiter, but 05:05

I'd have to go through the code to say whether 05:05

that was in -- whether that remained in the 05:05

design. 05:05

But that input thread arbiter would 05:05

obviously make a decision, and I think it made a 05:05

decision based on pixel versus vertex, but I'd 05:05

have to -- I can't say that for sure without 05:05

finding the code associated with that input 05:05

arbiter. 05:05

Q. And you don't know whether that was 05:05

for the 400, the 500 or the 600 version, right? 05:05

MR. FAHRENKROG: Object to the form. 05:05

A. Well, I am talking specifically about 05:05

the 400 version. 05:05

Q. Okay. Was that arbiter programmable? 05:05

MR. FAHRENKROG: Object to the form. 05:06

MR. PAIGE: Vague and ambiguous. 05:06

A. The dynamic mechanism of the GPRs was 05:06

a programmable mechanism. The input arbiter, I 05:06

can't say for sure, because I have to find the 05:06

actual code. But my belief is that it was not programmable. 05:06

Q. And for the one that was programmable, do you know when that was first conceived of? 05:06

A. That was present in the design very early on. I see mention of it before 2002. 05:06

Q. Did the R400 have a loop instruction? 05:06

A. Yes. 05:06

Q. When was that first conceived? 05:06

A. I'd have to go back and look. The loop instruction was very heavily based on Microsoft DX assembly language, which also have a loop instruction. 05:06

So I don't think that AMD can even claim invention credit for that in that we were simply implementing a requirement from Microsoft. 05:07

Q. Okay. And do you know when that was done? 05:07

A. I don't. I'd have to -- I think you could look it up in terms of Microsoft when they -- when they released DX specs. 05:07

Q. Did R400 have a call instruction? 05:07

A. It did. 05:07

Q. And when was that first conceived? 05:07

A. Again, that was mostly based on the Microsoft DX assembly language. So I wouldn't claim any great invention on the part of AMD. 05:07

Q. Did the R400 have a return instruction? 05:07

A. Yes, it did. 05:07

Q. And when was that first -- 05:07

A. Again, it's simply based on the API requirements. There was some innovation in terms of how to implement these things. But the implementation, you know, followed the requirements, and the requirements were set by Microsoft. 05:08

Q. Do you have any sense of what time frame that work was done? 05:08

A. I don't as I'm sitting here. I could find out. 05:08

Q. Did R400 have a branch instruction? 05:08

A. R400, it did have a form of a branch instruction. It was not a typical branch instruction. 05:08

Q. And why was it atypical? 05:08

A. Well, because R400 had really two types of instructions. There were control flow instructions, and the control flow instructions 05:09

kind of kicked off execution or ALU instructions. 05:09
And so the ALU instructions were like 05:09
add, multiply, subtract, and the control flow 05:09
instructions said, okay, execute, you know, eight 05:09
ALU instructions, and then the following eight 05:09
instructions would be the add, multiply, 05:09
subtract. 05:09
So the branch instructions were part 05:09
of those control flow instructions. So you 05:09
couldn't just put a -- you know, it wasn't put in 05:09
the middle there. And the branch instructions 05:09
only worked -- what do I want to say here? It's 05:09
a SIMD machine. So that means there's one 05:10
instruction flow for multiple pieces of data. 05:10
There was not a branch instruction 05:10
that you could do like a test on an individual 05:10
piece of data and say if this thing is a true 05:10
branch. The branch only functioned on all of the 05:10
pieces of data. 05:10
So that's what I mean when I say it 05:10
wasn't a fully general branch instruction. 05:10
Q. When was that done, do you know? 05:10
A. That was also done early if you look 05:10
at the early SQ documentation. There's 05:10
documentation about how the flow instructions 05:10

work. 05:10

Q. SQ, what is that? 05:10

A. That's the sequencer block. Most of this pile. 05:10

Q. Did the R400 rasterizer convert vertex data output into pixel data? 05:10

A. It would take in triangles, which were based on vertex data output. But vertices don't produce pixels. It's the grouping of them into 2D primitive that produces pixels. Obviously the triangles are a function of the output vertices. 05:11

Q. So is that a yes or a no or kinda yes, kinda no? 05:11

A. It's a kinda yes. It's not a real yes, because we don't -- the rasterizer did not directly take in vertices. It took in -- vertices first had to go through a fair amount of processing prior to getting to the rasterizer. 05:11

Q. Now, did your unified -- I'm forgetting the word, because it's getting late -- 05:12

A. Shader. 05:12

Q. The unified shader, could it also process primitives? Let me ask the question again so the record is clear. 05:12

Did the unified shader process primitives? 05:12

A. Not in the R400 or R500. It would process either vertices or it would process pixels. 05:12

Q. What about the 600? 05:12

A. The 600 did support processing primitives. 05:12

Q. So could the R400 assign threads to vertex, primitive and pixels? 05:12

A. No. 05:12

Q. It could only do it for vertex and pixels? 05:12

A. Correct. 05:13

Q. But the R600 could do it for all three? 05:13

A. Yes. 05:13

Q. Give me just one second. This was a question I had. Could pixels and vertex data proceed through threads out of order in the R400? 05:13

MR. PAIGE: Object to the form. 05:13

A. Could pixels and vertex data proceed through threads out of order? Not within a thread, because it was an SIMD machine. So all of the pixels in a given thread were pretty much 05:14

lockstep.

But pixels could be processed out of order with pixels, although they had to finish in order. Similarly, vertices could be processed out of order with other -- vertex threads could be processed out of order with other vertex threads, but they had to finish in order.

There was no ordering requirement on pixels versus vertices.

Q. Did you perform any of the simulations on the R400?

A. In terms of actually running them and seeing how they did?

Q. Yes.

A. I don't think so.

MR. CHATTERJEE: I think that's all I got. Thank you very much.

MR. FAHRENKROG: Okay. I just want to restate that we'll designate the transcript AMD confidential business information and the witness will read and sign the transcript.

(Continued on next page to include jurat)

VIDEOGRAPHER: The time is 5:15. At this time we'll end today's deposition of Andrew Gruber. We'll stop recording and go off the record. 05:15

(Time Noted: 5:15 p.m.)

ANDREW E. GRUBER

Subscribed and sworn to before me
this day of 2015.

C E R T I F I C A T E

Commonwealth of Massachusetts)
) ss:
County of Suffolk)

I, Michael D. O'Connor, a Notary Public within and for the Commonwealth of Massachusetts, do hereby certify:

That ANDREW E. GRUBER, the witness whose deposition is hereinbefore set forth, was duly sworn before me and that such deposition is a true record of the testimony given by such witness.

I certify that I am not related to any of the parties to this action by blood or marriage; and that I am in no way interested in the outcome of this matter.

IN WITNESS WHEREOF, I have hereunto set my hand this 18th day of March, 2015.

_____

Michael D. O'Connor, RPR, CRR, CBC, CCP

I N D E X

---------- E X H I B I T S ----------

Realtek Ex. 1017
Case No. IPR2023-00922
Page 211 of 213

---------- E X H I B I T S (Cont'd)----------

ERRATA SHEET FOR THE TRANSCRIPT OF:

Case Name: In the Matter of Certain Consumer Electronics Investigation and Display Devices with Graphics Processing and Graphics Processing Units Therein

Dep. Date: March 17, 2015

Deponent: Andrew E. Gruber

CORRECTIONS:

Pg.	Ln.	Now Reads	Should Read	Reason
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____
___	___	_____	_____	_____

_____

Signature of Deponent

SUBSCRIBED AND SWORN BEFORE ME

THIS ___ DAY OF _____, 2015

_____

(Notary Public) MY COMMISSION EXPIRES: _____