UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN FRANCISCO DIVISION

ADVANCED MICRO DEVICES, INC., AND ATI TECHNOLOGIES ULC, Plaintiffs-Counter Defendants v. LG ELECTRONICS, INC., LG ELECTRONICS U.S.A., INC., AND LG ELECTRONICS MOBILECOMM U.S.A., INC., Defendants-Counter Claimants.	CASE NO. 3:14-CV-1012-SI

VIDEOTAPED DEPOSITION of ANDREW E. GRUBER

July 27, 2017

Boston, Massachusetts

Reporter: Michael D. O'Connor, RMR, CRR, CRC

UNITED STATES INTERNATIONAL TRADE COMMISSION

WASHINGTON, D.C. 20436

In the Matter of:	
CERTAIN GRAPHICS SYSTEMS, COMPONENTS THEREOF, AND CONSUMER PRODUCTS CONTAINING THE SAME	Investigation No. 337-TA-1044

Thursday, July 27, 2017

8:44 a.m.

VIDEOTAPED DEPOSITION of ANDREW E. GRUBER, held at Fish & Richardson, P.C., One Marina Park Drive, Boston, Massachusetts, pursuant to notice, before Michael D. O'Connor, Registered Merit Reporter, Certified Realtime Reporter, Certified Realtime Captioner, and Notary Public in and for the Commonwealth of Massachusetts.

APPEARANCES:

ATTORNEYS FOR DEFENDANTS - COUNTER CLAIMANTS:

FISH & RICHARDSON, P.C.
901 15th Street, N.W.
Washington, D.C. 20005
(202) 626-6355
BY: R. ANDREW SCHWENTKER, ESQ.
schwentker@fr.com
BY: MIN SUK HUH, ESQ.
huh@fr.com

ATTORNEYS FOR PLAINTIFFS - COUNTER DEFENDANTS:

ROBINS KAPLAN LLP
800 Lasalle Avenue, Suite 2800
Minneapolis, Minnesota 55402
(612) 349-0172
BY: BRYAN J. MECHELL, ESQ.
bmechell@robinskaplan.com

APPEARANCES (Continued):

ATTORNEYS FOR ADVANCED MICRO DEVICES, INC.:

MINTZ, LEVIN, COHN, FERRIS, GLOVSKY & POPEO, P.C.
One Financial Center
Boston, Massachusetts 02111
(617) 348-1884
BY: MICHAEL McNAMARA, ESQ.
mmcnamara@mintz.com

ATTORNEY FOR THE DEPONENT:

NORTON ROSE FULBRIGHT US, LLP
1301 McKinney, Suite 5100
Houston, Texas 77010
(949) 852-7705
BY: DANIEL LEVENTHAL, ESQ.
daniel.leventhal@nortonrosefulbright.com

Also Present: Alex Daunais, Videographer

I N D E X

*****

E X H I B I T S

E X H I B I T S (Continued)

E X H I B I T S (Continued)

*****

P R O C E E D I N G S

VIDEOGRAPHER: We are on the record. This is the videographer, Alex Daunais, speaking. Today's date is July 27, 2017, and the time is 8:44 a.m.

We are here in Boston, Massachusetts to take the deposition of Andrew E. Gruber in the matter of Certain Graphic Systems Components Thereof and Consumer Products Containing the Same, ITC investigation number 337-TA-1044, and also Advanced Micro Devices, Inc. and ATI Technologies ULC versus LG Electronics, Inc., LGE Electronics USA, Inc. and Electronics MobileComm USA, Inc., case number 3:14-CV-1012-S1.

Will counsel please state themselves.

MR. SCHWENTKER: Andrew Schwentker from Fish & Richardson on behalf of the LG Defendants in the Northern District of California case and on behalf of the LG Respondents in the ITC investigation.

Also with me is Andrew Huh from Fish & Richardson.

MR. MECHELL: Bryan Mechell from Robins Kaplan on behalf of Advanced Micro Devices, Inc. and ATI Technologies ULC in the Northern District of California matter for Plaintiffs.

MR. McNAMARA: Michael McNamara from Mintz Levin on behalf of AMD and ATI.

MR. LEVENTHAL: Daniel Leventhal, Norton Rose Fulbright US LLP on behalf of the witness.

ANDREW E. GRUBER

having been satisfactorily identified by the production of his driver's license, and duly sworn by the Notary Public, was examined and testified as follows:

EXAMINATION BY
MR. SCHWENTKER:

Q. Good morning, Mr. Gruber. Can you please state your full name for the record.

A. Andrew Evan Gruber.
Q.
A.

Q. Have you been deposed before?
A. Yes.
Q. How many times?
A. At least twice. It could be three times. Probably three.
Q. And what was the subject matter of those depositions?
A. It was an action between NVIDIA and Qualcomm over multiple patents. So I was deposed multiple times in that case.
Q. Multiple times in that same case?
A. In that same case, because there were multiple patents involved.
Q. What were the patents involved in those depositions?
A. I don't recall the numbers. They dealt in, in general, with 3D graphics.
Q. I assume you're aware that you're a named inventor on some patents with Steve Morein, Laurent Lefebvre and Andi Skende?
A. Yes, I am.

Q. Were the depositions about those patents or other patents or both?

MR. LEVENTHAL: Form.

A. They were about those patents. I'm not sure that they were exclusively about those patents.

Q. Okay. So you may have been deposed about other patents?

MR. LEVENTHAL: Form.

A. Yes.

Q. You have been through the deposition process two or three times at this point, but I'll just go over some quick background with you.

You understand that you are under oath as if you were in a court of law?

A. Yes.

Q. Is there any reason that you cannot give accurate and complete testimony today?

A. No.

Q. So I'll be asking a series of questions today, and unless you're instructed not to answer by your counsel, you'll be expected to answer my questions.

Do you understand that?

A. Yes.

Q. I will assume that you understand my questions unless you ask for me to clarify. Okay?

A. Okay.

Q. We can take a break any time you want. The only thing that I would ask is that you answer any pending question, and then we can take a break.

A. Okay.

(Document marked as Exhibit 1 for identification)

(Document marked as Exhibit 2 for identification)

(Document marked as Exhibit 3 for identification)

Q. Mr. Gruber, the court reporter has handed you three exhibits. The first exhibit, Exhibit 1, is a subpoena to testify at a deposition in a civil action.

Do you see that?

A. Yes.

Q. And Exhibit 2 is a subpoena to produce documents, information or objects or to

permit inspection of a premises in a civil action?

A. Yes.

Q. And Exhibits 1 and 2 are from the Northern District of California.

Exhibit 3 is a subpoena duces tecum and ad testificandum from the International Trade Commission. Do you see that?

A. Yes.

Q. Have you seen these documents before?

A. I received the subpoena to produce documents and to testify. I don't know if these are the exact same ones, but they look certainly similar.

I don't think that I've seen Exhibit 3 before.

Q. Do you understand that you are here today in response to a subpoena?

A. Yes.

Q. And do you understand that you're here to testify in connection with both a District Court litigation between AMD and ATI versus LG Electronics, as well as an

International Trade Commission investigation?

A. Yes.

Q. Could you turn to Exhibit 3. There are two attachments, Attachments A and B. If you could turn to Attachment A first. What I'm referring to is Attachment A.

A. Is that near the end?

Q. No. It's towards the beginning.

A. Attachment A. I see what you're saying, yes.

Q. If you could turn to Page 10 of Attachment A.

A. Okay.

Q. Do you see requests for production?

A. Yes.

Q. Have you seen these requests for production before?

A. I believe that I have. I certainly received a request for production, and this could have been it. I don't have a clear memory of whether this exactly was what the request was.

Q. Okay.

A. There was a request for production

associated with the subpoena that I received.

Q. Okay. So if we could turn to Exhibit 2, Page 4, there's a list of requests for production there. Is that the list of requests for production that you were referring to?

A. Yes.

Q. Did you search for documents in response to these requests for production?

A. Yes, I did.

Q. And did you find anything?

A. I found things related to my employment at AMD and ATI. Anything else that I had would have been transferred as part of Qualcomm's acquisition of ATI's mobile group, and was owned by Qualcomm, and on the advice of my counsel, I did not produce those.

Q. But you had -- you did find documents that you did not produce?

A. I didn't look through all of my electronic-related information that was on my Qualcomm-owned equipment. So I can't -- I can't say definitively that I found documents there.

I didn't look through those

electronic documents that were owned by Qualcomm based on the advice of my counsel that they weren't required as part of the production.

Q. Okay. You mentioned you found documents related to your employment at AMD and ATI. Did you provide those to your counsel?

A. I did.

Q. Do you recall what those documents were?

A. They were mostly related to employment agreements and stock option grants.

Q. Okay. Anything else?

A. I don't recall anything else.

Q. Then turning back to Exhibit 3, the request for production that we were looking at.

A. Yes.

Q. So I take it that you did not specifically search for documents responsive to these requests for production?

A. I did within my personally owned papers. I did not within information that I didn't have personal ownership of.

MR. LEVENTHAL: To interject, the

production that was made in ND Cal was reproduced in ITC with both confidentiality designations as responsive to Exhibit 3, and then the LinkedIn profile that were produced yesterday was produced responsive to request for production No. 1 in Exhibit 3.

MR. SCHWENTKER: Okay. Understood.

Q. If you could turn in Exhibit 3 to -- after Attachment A, there's an Attachment B --

A. Yes.

Q. -- which lists a number of deposition topics. Do you see those?

A. Yes.

Q. Have you seen those deposition topics before?

A. Yes.

Q. You understand you're here to testify with respect to these topics?

MR. LEVENTHAL: Object to the form.

A. Yes.

Q. What did you do to prepare for your deposition?

A. I produced my personal papers associated with my employment at AMD and ATI. Other than that -- and I had some depo prep as to form with my lawyer, but that was it.

MR. LEVENTHAL: I caution you not to reveal the contents of any communication.

Q. When did you meet with your attorney?

MR. LEVENTHAL: Form.

A. I met with him over the phone earlier in the week.

Q. For how long?

A. I'd say about 45 minutes.

Q. Did you review any documents to prepare for your deposition?

A. No.

Q. When did you find out that you would be deposed?

MR. LEVENTHAL: Form.

A. Well, I received a subpoena early in the year. I didn't find out the actual date of the deposition until a few weeks ago.

Q. Do you have an understanding of what -- strike that.

So you understand that there are two separate lawsuits at issue here today?

A. Yes.

Q. So one is the Northern District of California lawsuit and one is the ITC investigation; do you understand that?

A. Yes.

Q. Do you have an understanding of what the Northern District of California lawsuit is about?

A. No.

Q. No?

A. No.

Q. Do you have an understanding of what the ITC investigation is about?

A. No.

Q. Did you speak with anyone else in preparation for your deposition besides your attorney?

A. No.

Q. I'd like to turn to your background. Are you currently employed?

A. I am.

Q. Where?
A. At Qualcomm.
Q. What's your position?
A. I'm a vice-president of GPU architecture.
Q. GPU architecture?
A. Yes.
Q. How long have you been at Qualcomm?
A. Since 2009.
Q. And where did you work before that?
A. I worked for AMD.
Q. How long have you worked for AMD?
A. Since the acquisition of ATI.
Q. When was that?
A. I believe it was 2004, but I can't say with certainty.
Q. Before the acquisition, you were with ATI?
A. I was.
Q. How long had you been with ATI?
A. I started with ATI in 1994.
Q. Can you take me through your roles at ATI and AMD, starting in 1994?

A. Do you have a specific question about that or do you just want like an overall summary?

Q. An overall summary.

MR. McNAMARA: Objection.

A. I was involved in the architecture of the initial 3D graphics at ATI, their initial 3D graphics product, and I continued that role throughout my ATI employment.

At one point ATI had multiple design teams, and I was in charge of the GPU architecture for the East Coast design team.

Q. And when you say the East Coast design team, where is that based?

A. It was based in Marlboro, Mass. at the time. It's now based in Boxboro, Massachusetts. But I was at the Marlboro location.

Q. What was your position when you started at ATI in 1994?

A. I was a -- my role was to drive the 3D architecture. I might have started on a contracting basis and switched over to a full-time employment. I switched over at some point. I don't recall when that happened.

Q. From a contractor to a full-time employee?
A. Yes.
Q. And then did you have different titles during your time at ATI?
A. Yes. Most of the time I was a principal.
Q. And do you recall when you became a principal?
A. I do not.
Q. Do you recall any other titles you had?
A. I don't. I mean, it, you know, was some kind of engineering title, senior engineer or something like that, but I don't recall the exact title.
Q. Would that have been before or after principal?
A. That would be prior to principal. Principal was my last title at ATI and AMD.
Q. Okay. So you became a principal at ATI, and then that title carried over to AMD?
A. I cannot say for sure that I was a principal at ATI. I don't recall when I became

a principal.

Q. Okay. And did you have any other titles at AMD?

A. No. Well, as I say, if I started at AMD and I was not a principal, I could have been a senior engineer, and then made a principal during my time at AMD.

Q. But at the time of the acquisition --

A. I was a principal.

Q. -- by Qualcomm, you were a principal?

A. Yes.

Q. Okay. And what did Qualcomm acquire from ATI?

A. The mobile graphics --

MR. SCHWENTKER: Strike that.

Q. What did Qualcomm acquire from AMD?

A. The mobile graphics group, as well as some assets associated with that.

Q. How large is that group?

MR. LEVENTHAL: Form.

A. It was not very large as to the group that I was immediately associated with.

I would say under ten people, but they also acquired some display group assets that were based in Toronto that were larger, but I'm not sure how large.

Q. And you also mentioned some assets associated with the mobile graphics group. What are those assets?

A. Some of our like physical laptops, but I believe they also acquired some intellectual property assets as well.

Q. Do you know what those are?

A. I don't know in any detail. I just know that Qualcomm did acquire some rights to use AMD intellectual property, so that we did not have to worry about impinging on AMD intellectual property in future Qualcomm designs.

Q. What's your educational background?

A. I have a Bachelor's of Science in electrical engineering and computer science from MIT.

Q. When did you receive --

A. 1981.

Q. Do you have any postgraduate

education?

A. No.

Q. Did you work on specific products while you were at ATI and AMD?

A. Yes.

Q. What products did you work on?

A. As I mentioned, I worked on their first 3D product, which is called the Rage. We had a Rage II, a Rage III, after which I worked on a product called the R400, which -- I'm sorry, after the Rage series was the Radeon series. There was the Radeon 100 and the Radeon 200. The Radeon 300 or something called R300 was one that I did not work on, because that was done out of our California group. The R400 was done out of my group. The successor to that was the R500 and the R600, which I also worked on.

Q. Just to make sure I understand. You worked on the Radeon 100 and Radeon 200, you did not work on the R300, but then you worked on the R400, R500 and R600?

A. Correct. Some of those products were never produced, but there was still internal projects.

Q. Which ones did you not produce?
A. The R400 was never produced. The R500 was never produced either. The R400 and R500 -- the R400 was delayed and turned into the R500, and the R500 was delayed and turned into the R600, and the R600 was eventually produced.
Q. When was the R600 produced?
A. I don't recall.
Q. Do you recall an approximate time frame?
A. If you pressed me, I would say around 2006, but I could be wrong about that.
Q. You said the R400 was delayed. Why was it delayed?
A. There were two reasons for it. One was it was simply a difficult product. It was a lot of stuff that had to get right. There were changes throughout the design, and not all of those were ready in time for the market window.
Another reason for the delay was that a lot of the technology and the basic design was directed at the Xbox 360 contract that ATI had won from Microsoft. So that split

some of the resources that were needed to finish the design.

MR. MECHELL: I designate the transcript highly confidential and attorneys eyes only under the protective order of the Northern District of California at this time.

Q. What do you mean when you say that because it was directed at the Xbox 360 contract that split some of the resources that were needed to finish the design?

MR. LEVENTHAL: Form.

A. Well, the basics of the design were similar for both the R400 and the Xbox 360. There were differences as well. Microsoft wanted a different memory system, for instance, and we had to support all the requirements and customer interface of Microsoft from an engineering point of view.

So we split our engineering team. So some people were dedicated to dealing with the Microsoft-specific requirements while others of the engineering team pushed on with the desktop market product.

Q. And by "desktop market product,"

are you referring to the R400?

A. Yes.

Q. So the -- strike that.

Have you heard of the Xenos?

A. Yes. That was another code name for the Microsoft product.

Q. Okay. So the Xenos was the product for the Xbox 360?

A. Correct.

Q. And that was a different product than the R400?

MR. MECHELL: Objection to form.

A. Correct.

MR. MECHELL: For the record, when counsel for District Court objects, I'm not going to object as well. So you can keep the transcript as clear as possible. Is that fine?

MR. SCHWENTKER: That's fine.

MR. McNAMARA: And we'll also designate the entirety of the transcript as confidential business information in the ITC protective order.

MR. SCHWENTKER: Okay.

MR. MECHELL: And counsel, under a

similar approach, if counsel makes an objection, I'll try to keep the transcript clean as well, and please assume it's adopted as well.

MR. SCHWENTKER: Come on. You guys just are being lazy.

MR. McNAMARA: We can talk all over you. That's fine, too. Whatever way you want to go about it.

MR. SCHWENTKER: That's fine.

Q. Did you work on the Xenos?

A. I was on the R400 product. I, in that sense, interacted and had a role on the Xenos team as well, but my primary role was pushing forward the desktop effort.

Q. Where was the R400 team?

A. Most of the R400 team was in the Marlboro office. There were some people working in the Orlando office.

Q. Where was the Xenos team?

A. The Xenos team was primarily in the Orlando office, but there were some people in the Marlboro office working on it as well. It was not a clean split in between the two offices.

Q. When did you start working on the R400?

A. I haven't looked at that documentation recently. So I have -- I couldn't say off the top of my head when it was.

Q. What documentation are you referring to?

A. Design documentation that I recall seeing in the previous ITC case.

Q. You're not referring to the documentation that you have in your possession?

A. No. I don't have any of that.

Q. Earlier you said that one of the reasons -- one of the two reasons the R400 was delayed was because it was difficult. Do you recall saying that?

A. Yes.

Q. What was difficult about the R400?

A. The primary reason why it was late was that the memory interface associated with it was aggressive. It had various kinds of compression associated with it, and it just took us a long time to get that working correctly. That's why we were able to come out

with the Xenos product earlier than the R400, because it used a different memory interface that didn't have a lot of the complexity of the R400 interface.

Q. What do you mean by "memory interface"?

A. The part of the logic that talked to the frame buffer and did the final blending associated with the frame buffer.

Q. So you said that that was the primary reason. The primary reason why it was late was that the memory interface associated with it was aggressive?

A. Yes.

Q. Are there other reasons?

A. It was a difficult task on a lot of fronts. So I think that development didn't go as well as our initial schedule was. It was the memory interface that was the long haul.

I think a lot of the rest of the logic was working fairly well. That's why it gave us confidence to commit to Microsoft for the Xbox part, which aside from the memory interface, I mean, the rest of the Xbox part was very similar to what we were working on for

the R400 and used the same database.

Q. What do you mean by "the same database"?

A. There was a fork at one point, but they all started based on the same database. In other words, the R400 product was -- the R400 design was used and then forked to do the particular changes required by Microsoft.

Q. And what was in that database?

A. All of the graphics accelerator functions. So we had a command processor, we had a setup engine, we had a unified shader, we had a rasterizer, we had a texture fetch unit. Those are all elements that are typically part of a GPU.

Q. Are you referring to source code or RTL?

A. Yes, RTL source code.

Q. So that's what you were referring to by database?

A. Yes.

Q. What about documentation that went along with that?

A. Again, we had R400 documentation initially, and then we kind of split that off

and made Microsoft Xbox specific documentation. But they, you know, for a long time, they were the same, and then they got split off.

Q. When did ATI start working on the Xenos project?

A. I don't have any definite recollection of the dates.

Q. You mentioned a setup engine. What's a setup engine?

A. It's something that takes the vertices of the triangles and generates parameters that allow you to rasterize the individual pixels within the triangles, and essentially figures out slopes of the various parameters in the -- of the vertices so that you can interpolate what the pixel values are within the middle of the triangle.

Q. What provides the vertices that are input into the setup engine?

MR. McNAMARA: Objection.

A. The vertices are typically fetched from memory. In the case of the R400, those vertices were fetched by the shader.

Q. And that was a unified shader, I believe you said?

MR. McNAMARA: Objection.

A. Yes. I'm not sure that I said that, but yes, it was.

Q. What is a unified shader?

MR. McNAMARA: Objection.

MR. MECHELL: Objection.

A. A unified shader refers to the same hardware, the same instruction set processor executing both vertex commands and pixel commands.

Q. You said the R500 was never produced either. Why was that?

A. Well, we made changes as we went from the R400 to R500, to add features and get better performance needed, because it was going to be a year after the R400 was going into a different market window. And we ran into some of the same schedule issues, so it was not ready for production in time for that market window.

Q. What was your role on the R400?

A. I was the lead architect for the R400.

Q. Who did you work with?

MR. MECHELL: Form.

Q. On the R400?

A. I worked with Laurent Lefebvre, Andi Skende, Steve Morein, Michael Doggett, Larry Seiler. Those are the names who come to mind as leading various areas of the design. Jocelyn Houle as well.

Q. What did you do as the lead architect?

A. I kind of drove the overall direction of the design. Different people kind of headed up the architecture of different areas of the chip. We would have somebody who was in charge of, say, the texture system or somebody who was in charge of the shader system, somebody who was in charge of the memory system, which we called the render back end, somebody who was in charge of the setup engine; those kind of things.

Q. Had you been the lead architect on any prior products?

A. Yes.

Q. Which products?

A. All of ATI's 3D products, so that's all the ones that I mentioned, the Rage I, Rage II, and the Rage III, which was

productized as the Rage I, Radeon 100, Radeon 200, as I mentioned, not the Radeon 300, but on the 400.

Q. You were the lead architect on all of those products?

A. Yes. The lead 3D architect. There were other parts of the chip, the display system mainly, that I was not involved with.

Q. What was Laurent Lefebvre's role on the R400?

A. He was the sequencer architect. So that's the block that drives and controls the shader. You can think of it as processing instructions and ordering instructions for the shader and submitting work to the shader.

Q. Those are all functions of the sequencer?

A. Of the sequencer.

Q. Does the sequencer have any other functions?

A. I think that that captures the overall, you know, providing work, and then providing instructions to the shader for execution.

There are multiple threads and

scheduling those threads for the shader as well.

Q. What is a thread?

A. A thread --

MR. McNAMARA: Objection.

A. -- the way that I was using it. People use different terms for this. NVIDIA would call them warps. AMD sometimes called them waves. But it's a single group of items, which could be pixels or vertices, that share the same flow of execution. So they all execute the same instruction on the same cycle, even though they're using different data to execute that given instruction.

Q. How did the sequencer schedule the threads for the shader?

MR. LEVENTHAL: Object to form.

A. It had a heuristic to give priority to what got submitted next. Things would be split up into what were called clauses, which is a group of instructions that could not be interrupted.

So it would submit clause by clause for each thread or wave of execution, and it would have to arbitrate in between

whether it's going to submit a pixel thread or a vertex thread is the next clause to be executed.

Q. How would it do at arbitration?

MR. McNAMARA: Form. Vague.

A. I know it had some heuristic for doing it, but I don't recall the details of the heuristic that it used.

Q. By "heuristic," are you referring to an algorithm?

A. Yes, an algorithm. I say "heuristic," because there's no right answer. It's more of a strategy that it used.

Q. So you said Laurent Lefebvre worked on the sequencer. What was Andi Skende's role on the R400?

MR. McNAMARA: Objection.

A. Andi Skende was mainly involved in the implementation of the shader. So I would say that he was less involved with architecture than taking the architecture and implementing it and documenting the architecture.

He did have some role in the architecture of that system, but I would say it wasn't primary.

Q. When you say "implementation of the shader," are you referring to -- well, what are you referring to?

A. I'm saying that he wrote a good deal of the RTL code for the shader, and as well as documenting, you know, the details of the instruction set, what each instruction did on a detail basis so that the compiler could implement instructions for it.

Q. And what was Steve Morein's role on the R400?

A. Steve was the person who first came up with the idea and the impetus for the unified shader. So he was involved early on in the architecture of the unified shader. He and I worked closely on that, but I would say Steve was the initial driver of that.

Q. Is that something that was new in the R400?

A. Yes.

Q. So would you say that -- strike that.

Do you know when Mr. Morein came up with the idea for the unified shader?

A. I don't have any clear

recollection of when that was.

Q. Do you recall any conversations with him about the unified shader?

A. Well, I recall that I had numerous conversations with him about it, trying to understand how it would work, and both of us working out some of the details of how we would feed data from one shader stage to the next shader stage, while still utilizing a single piece of hardware.

So I don't recall any of the details of those conversations, but I recall the fact of having them.

Q. Do you recall him coming to you initially with the idea of the unified shader?

A. I know that he did, but I don't recall the details of when or how that happened.

Q. Okay. So you don't recall any initial conversation where he came to you and said I have this great idea for a unified shader?

A. No. But I know that I was -- well, I think I was the first one that he came to, and we spent a lot of time where, you know,

he convinced me of the basic idea, and we worked out the details. And then both of us would kind of act as proselytizers for the idea to some of the other engineers within the company.

I do remember a specific trip to Orlando, for instance, to present this idea to the engineers there and to convince them that it was a good idea.

Q. Do you recall when that was?

A. I don't.

Q. Do you recall any reactions that you received at that presentation?

A. I think there was openness, but some skepticism as well as to whether this idea would actually work.

Q. Why is that?

A. There was some concern over both the performance and whether it would have adequate precision for some of the operations that needed to happen.

Q. Do you recall what the concerns about performance were?

A. No. All I recall was a general performance concern.

Q. Do you recall what the concerns about precision were?

A. Surprisingly, I do. The concern was that this shader meant that we were changing the way that we did interpolation.

So rather than having the setup engine move the vertices -- sometimes a triangle can be large, larger than the screen, and to draw it, you only want to draw the portion that is on the screen. So you have to clip the triangle to the screen boundaries.

The way that we had previously done that was in the setup engine we would actually generate new vertices at the screen boundary. So that would be done in the setup engine.

The unified shader idea was that we would take the original vertices and interpolate them within the shader from their original positions. So that sequence of operations meant that there were different precision outputs, different precision results you would see, and the issue was whether that would cause a problem in terms of accuracy of the interpolation.

Q. Do you recall who was at that meeting in Orlando when you first presented the idea of the unified shader?

A. I know that Mike Mantor was there, Clay Taylor was there, Mike Mang was there, M-a-n-g, Scott Hartog was there as well.

Q. Do you recall anything else about their reactions at that meeting?

A. The only other thing I recall was that Steve was giving the presentation, but they wanted confirmation that I supported this, and was behind the idea.

Q. And what did you say?

A. I said that I was. I thought it was a good idea. I said I thought it was the future.

MR. LEVENTHAL: We're a little over an hour. We don't need to break now, but at your next point.

MR. SCHWENTKER: Okay. Let me just ask a couple more questions, and then we can break, if that's okay.

Q. So why did you think it was the future?

A. I could see the advantages of it,

because it's very difficult to predict vertex load versus pixel load on any given, I'll call it, changes on a dynamic basis, both within an application and between applications, and this was an efficient way of dealing with that dynamic loading problem.

Q. Were there any other advantages?

A. You could just engineer one shader system. I mean, as the APIs were moving toward very similar capabilities of the two, it was wasteful to devote separate engineering resources to being one guy working on a vertex shader and another guy working on a pixel shader.

So while it was going to take significant engineering resources to get the first iteration working, as we, you know, understood how the whole thing worked and the data flow worked, going forward it would mean an engineering savings, engineering resource savings.

MR. SCHWENTKER: Why don't we take a break now.

VIDEOGRAPHER: The time is 9:49 and we're off the record.

(Recess taken at 9:49 a.m. and reconvening at 10:05 a.m.)

VIDEOGRAPHER: We are back on the record. The time is 10:05.

BY MR. SCHWENTKER:

Q. Welcome back. Before the break, you mentioned you were involved with proselytizing for the unified shader idea, and you mentioned a meeting in Orlando.

A. Yes.

Q. Was there any other proselytizing you did for the unified shader idea?

A. That's the main effort that comes to mind. I'm sure within the Marlboro group, there were other meetings to explain what the unified shader was and how we're planning on implementing it, but I don't have specific recollection of those.

Q. That meeting in Orlando, was that an internal meeting?

A. It was an internal meeting, yes.

Q. There was no one else besides ATI employees at that meeting?

A. Correct.

Q. Did you ever talk to people

outside of ATI about the unified shader idea?

MR. MECHELL: Objection. Vague.

A. Well, we certainly did as part of the Microsoft effort to, you know, get them to use ATI as a supplier for the Xbox.

Q. Were you involved in those discussions with Microsoft?

A. I was probably involved with some of the early ones, you know, the initial technical sales efforts.

Q. What do you remember about those discussions with Microsoft?

A. I don't have any specific recollections of them.

Q. Do you remember if you traveled to Microsoft?

A. I believe that I did. But again, I don't have any specific recollection.

Q. Do you know when those discussions took place?

A. I don't have any specific recollection.

Q. You don't know when they started?

A. It would have been sometime during the R400 effort. The original idea was to get

them to use the R400 unchanged, and it would be just like another customer for the R400. That never happens when you're dealing with a large customer like Microsoft that has, you know, their own ideas about what their product should look like.

Q. So ATI originally went to Microsoft with the idea of the R400, pitching that to Microsoft?

A. Yes.

MR. MECHELL: Objection.

Misstates the testimony.

Q. Is that an accurate characterization?

A. Yes.

Q. Then is it fair to say that at some point the discussions with Microsoft turned from the R400 to the Xenos?

A. You know, I wasn't involved enough later in the process to really say how this happened.

Q. Do you recall any documentation that was shared with Microsoft in this process?

A. I know that at some point we shared the detailed shader documentation with

them. That included the shader instruction set, because they were able to generate their own compiler for the Xenos shader, but I don't know when that was shared with them.

I do remember being surprised that they were able to do that on their own, that the documentation was good enough that they could figure that out.

Q. Why were you surprised at that?

A. Well, because oftentimes engineering documentation is not thorough enough to make sure that there are no -- well, that it's thorough enough to write a compiler that has to work for all inputs, that there are no holes associated with the documentation that would cause the compiler to produce incorrect code.

Q. What do you mean by "compiler"?

A. The shader executes instructions, but at the API level, the application gives instructions in a higher level language. And the compiler takes that higher level language, which the application writes at higher level language and it can run on any hardware, not only ATI's hardware, but the compiler takes a

higher level language and compiles it into hardware specific instructions that only run on one vendor's hardware.

Q. So Microsoft was able to compile their higher level instructions so that they would run on the ATI product?

A. Correct.

Q. And when you say that, are you referring to the R400 or the Xenos?

A. The Xenos.

Q. Do you know if they ever were able -- strike that.

Do you know if they wrote a compiler for the R400?

A. I am unaware of them doing so, and I don't know why they would.

Q. Why do you say that?

A. Because we would supply such a compiler with our driver when we ship our, you know, the R400 was aimed at the desktop market. So that would be part of the driver that ATI would ship.

Microsoft doesn't produce compilers for desktop hardware. That's the desktop vendor's job.

Q. Why did Microsoft create a compiler for the Xenos?

A. I'm not sure. I know that they did, because I was surprised that they did. I had expected them to use ATI's compiler.

Q. Do you have any documents from your involvement with Microsoft on the R400 or the Xenos?

A. I don't know that I do. They may be in, you know, my electronic documents that are, you know, online somewhere in what was taken over -- taken with me from AMD to Qualcomm. But I don't have any personal possession of any documents associated with that.

Q. What about e-mails?

A. Again, the e-mails were a part of the electronic assets that were taken from ATI to Qualcomm.

Q. And are those part of the documents that you said earlier you did not search at the --

A. Correct.

Q. -- instruction of your counsel?

A. Correct.

Q. Do you recall when you traveled to Microsoft?
A. No, I don't.
Q. And do you recall how many times you traveled to Microsoft?
A. No. I know that I traveled to Microsoft many times, but they certainly weren't all specifically associated with trying to get Microsoft to use ATI products for the Xbox.
Q. What else were your trips to Microsoft for?
A. Well, the primary reason was that Microsoft was evolving APIs for applications to use for running graphics on the desktop. So we wanted to assure that the API capabilities matched our planned future hardware.
Q. Do you recall which products you discussed with Microsoft?
A. Well, we certainly wanted to make sure the R400 and DX9 were compatible. DX9 was their particular version of their API. Their APIs DX, and they have DX9, DX10, and they're up to DX12 now.
Q. So you at least had discussions

with Microsoft about the R400?

A. I would say about the capabilities of the R400. It was not from a product point of view, but for certain -- but from the point of view of being able to run DX9 applications on the R400.

Q. Did you have any discussions with Microsoft about other products that you worked on?

A. Yes. We certainly, you know, spoke to them in similar ways for like the DX10 and DX11, and those would apply -- I mean, the R400 was not a DX10 class machine, but some future product was. The R500 and R600 were DX10.

Q. When ATI was trying to get Microsoft to use the R400 in the Xbox 360, do you recall who else at ATI was involved in those discussions?

MR. MECHELL: Objection to form.

A. I know that Bob Feldstein was involved. I remember that.

Q. Bob Feldstein?

A. Yes.

Q. Anyone else?

A. I don't recall specifically.
Q. Do you recall who you dealt with at Microsoft in those discussions?
A. I don't. I have a vague memory of Andrew Goosen being involved from the technical side.
Q. Who is he?
A. He's a Microsoft -- he's high up on the technical ladder at Microsoft. That's all I can tell you. I don't know what his title is.
Q. Earlier you said that you worked on a series of products at ATI. I assume each one of those products was an advancement over the prior products?
A. Right.
Q. What were the differences between the R300 and the R400?
MR. MECHELL: Objection. Vague.
A. Well, the unified shader associated with the R400, I think, was a big difference. Along with that unified shader came 32-bit floating point capability in the fragment shader, which the R300 didn't support.
Q. What capability did the R300

support?

MR. MECHELL: Objection.

A. The R300 had separate vertex and pixel shaders. The fragment shader was limited to 24 bits.

Q. What's the fragment shader?

MR. McNAMARA: Objection. Vague.

A. So the fragment shader is the part of the GPU that runs instructions to determine the final color of the fragment or pixel that's output to the frame buffer.

Q. Is that a separate component of the GPU?

MR. McNAMARA: Objection.

A. It is. Although -- well, let me just make sure. Fragment shader and pixel shader are typically used interchangeably. So it is part of the unified shader system, at least in R400, that can do both vertex shading and pixel or fragment shading.

Q. Earlier when you said that the R300 had separate vertex and pixel shaders, pixel shader also refers to fragment shader?

A. Fragment shader. It's just a different term for the same thing. Fragment is

a little more correct in that you may, if you have what's called a multi-sampled frame buffer, there may be multiple samples for each pixel.

So the group of samples that are written with the output of a single shader invocation are called a fragment.

So on the edges of triangles where the edge may cut through a pixel, you may only write out a portion of that pixel, just the samples within that pixel that are hit, and those are called the fragments.

You may not color the entire pixel with the output from the shader. So that's why people sometimes use fragment rather than pixel.

In the typical case where you're not doing multi-sample rendering, where you just have one sample associated with every pixel, then a pixel and a fragment are identical.

Q. Okay. But we can refer to the fragment shader as a pixel shader?

A. Yes.

Q. So is it accurate to say that the

R300 had a vertex shader and a separate pixel shader?

A. Yes.

Q. Okay. So those were two separate components in the R300?

MR. McNAMARA: Objection.

A. Those were two separate components.

Q. And you said that the R400 was different from the R300 in that it had a unified shader?

A. Correct.

Q. So does that mean that in the R400 the pixels -- strike that.

So in the R400 where the -- strike that.

In the R400, were the vertex shader and the pixel shader combined into one component?

MR. McNAMARA: Objection to form.

MR. MECHELL: Objection.

A. Yes.

Q. And that's the unified shader?

A. Yes.

MR. McNAMARA: Objection.

Mischaracterizes.

MR. MECHELL: Objection.

Q. Going back to the distinction between fragments and pixels, I want to make sure I understand the distinction between those.

Is there a difference between fragment color and pixel color?

MR. McNAMARA: Objection to form. Vague.

A. In a multi-sample frame buffer, which is used for anti-aliasing, for making the jaggies along the triangle edge smoother, you can have multiple samples for each pixel in that frame buffer. So a pixel may have multiple colors in them.

When that happens, we say that there are multiple fragments within that pixel.

Q. So fragments can make up a pixel?

A. That's right. Or a pixel can be just one fragment. In the inside of a triangle, each pixel is a single fragment. On the edges of a triangle, where the edge bisects a pixel, and you have a multi-sample frame buffer, you can have multiple fragments within

that pixel.

Then when that frame buffer is finally displayed on the screen, you will blend the two fragment colors and produce the single color for that pixel that shows up on the panel.

Q. So what generates the fragment color?

A. The fragment color is the output of the pixel shader or the fragment shader, whatever you want to call it. That gets modified by the render back end, if there's any blending operations that happen.

If you're not doing blending, which is the common case, then the fragment -- then the output of the pixel shader is simply converted to the particular pixel format that you have in the frame buffer, and it's written to the frame buffer.

Q. Okay. And what generates the pixel color?

MR. McNAMARA: Objection.

A. As I said, it's a combination of the output from the pixel shader, which is the result of instructions being executed on input

data, which, you know, could be textures, it could be constants, it could be the result of vertex interpolation.

They are all combined in whatever way is specified by the shader program, and the shader program produces an output, color, for the pixel, which is then sent to the render back end, which may modify it further, and then it gets written to the frame buffer.

Q. Okay. So the ultimate pixel color is output from the render back end?

A. Correct. The render back end is sometimes called the ROp unit. AMD calls it the render back end, but AMD calls it the ROp unit.

Q. How do you spell that?

A. R-O-p.

Q. Thank you. So as I understand what you've said, the R300 had a separate vertex shader and pixel shader or also called the fragment shader, and the R400 had a unified shader that combined both the vertex shader and the pixel shader; is that accurate?

MR. McNAMARA: Objection.

A. Yes.

Q. Were there challenges in going from the two-shader model in the R300 to the unified shader in the R400?

MR. LEVENTHAL: Form.

A. Yes.

Q. What were the challenges?

A. Well, there were a lot of them. You had to figure out how the data would flow from one shader to the other, the buffering associated with that data.

The reason why that's a challenge is that, you know, you don't want to get into the situation where maybe you have all the -- the shader is running all vertices and it's using all of its resources to run vertex shader operations, and the only way that it can finish running those vertex shader operations is to export those vertices when it's done to some buffer.

If that buffer is all full up, because the only way that that buffer drains is by processing -- changing those vertices into pixels and processing the pixels, because you have a unified shader, you can't process pixels if the shader is full up processing vertices.

So you can run into a deadlock situation where we can't run any more vertices because the output buffer is full, but we can't run any pixels because the shader is full of vertices. So that's certainly one challenge.

Simply, you know, how to assemble work for the two and arbitrate in between the two tasks in an efficient manner, when you have to arbitrate in a functional manner, but then you also want to arbitrate in an efficient manner so that you don't wind up getting bubbles in your pipe.

I mean, maybe you didn't -- as an example, maybe you ran all of your pixel operations, you didn't run any vertex operations, you used all of your resources for running pixels, you're finally done with running pixels, and now it's time to run vertices, but the pixel pipe is now totally drained. All the logic in between the vertex pipe and the pixel pipe, some of that setup logic that we spoke about earlier, winds up emptying due to, you know, not enough interleaving in between vertex processing and pixel processing.

So it would still functionally work, but you'd wind up with bubbles in your pipe, which means it would run slower because some of the intervening logic is not utilized as effectively as it otherwise could be.

Q. So the first challenge you mentioned related to buffering data, and the second challenge related to assembling work and arbitrating between the two tasks?

A. Right. I would also say that fetching data for the two in a manner that they don't wind up thrashing resources is a challenge as well.

You have to get vertices into the vertex engine. You have to get -- typically you'll have to get textures, which are also read from memory, into the pixel engine.

In the R400 we shared the same path for those, and you wanted to assure that one didn't hold up the other unnecessarily or that one didn't wind up -- that they didn't wind up fighting in whatever cache storage was in that data path.

Q. So that challenge relates to fetching data?

A. Yes.
Q. Were there any other challenges?
A. Those are the main ones that come to mind. I'm sure there were other challenges as well.

Allocating resources in between the two, what we call GPRs, which are the temporary variables used by the shader programs, they were a shared pool, but you wanted to make sure that that pool was used efficiently.

That became especially challenging, given that you could solve a lot of problems within a pixel shader or within a vertex shader by simply running those waves or those threads, whatever you want to call them, in the order that they were received, so that you knew that, you know, if a thread started -- if a thread one started before thread two, that thread one would finish before thread two, and therefore, give up its resources in a sequential order.

But that doesn't work well when those threads could be mixed in between vertices and pixels, because the shader

programs have nothing to do with one another. So making a short vertex shader wait for a long pixel shader or vice versa isn't a good idea from a performance point of view.

So you had to deal with the fact that these threads or waves would finish out of order.

Q. So that challenge related to allocating resources between the vertex and pixel shaders?

A. Yes.

Q. Any other challenges?

A. I'm sure that there were other challenges. Those are the ones that come immediately to mind.

Q. Going back -- I'd like to step back through those challenges that you identified. The first one related to buffering data.

How did you solve that challenge?

MR. McNAMARA: Objection.

A. Generally we allocated -- we made sure that we had room for output before we would allow a shader wave to start up and grab shader resources. So that you would eliminate

the deadlock possibility associated with that, with the linkage in between pixels and vertices. So I believe that's how we solved the functional problem.

The performance problem was solved via various heuristics about how we would arbitrate to decide whether we would run vertices or pixels next, and I don't recall the details of those heuristics.

But those heuristics were subject to these hard and fast rules that would prevent deadlock and keep functionality, even if while the heuristics were there, to assure performance.

Q. When you said that you made sure you had room for the output, is there -- I guess room in what?

A. So we had what we called a parameter cache, which was where the vertex shader would drain into. You had both a parameter cache and a position cache, but in some -- but they're similar.

It's just one holds a specific type of parameter called position that is needed by the setup engine, and the other one

holds parameters that are not needed by the setup engine, but are used later by the pixel shader.

But the vertex shader writes out both positions and parameters and the strategy for managing those two separate caches was similar.

Q. The second challenge that you mentioned related to assembling work and arbitrating between vertex and pixel shaders. How did you solve that challenge?

A. So again, there were some assembly buffers that we would, you know, group up the whole wave of pixel work or a whole wave of vertex work prior for that wave being eligible for arbitration, for entry into the shader.

And then the sequencer would decide which of those possible waves, we had a wave at the front of the queue from a vertices sees, a wave at the front of the queue for pixels, and it would look at available resources. It would look at how full we are with pixels versus vertices.

Again, these are the heuristics I mentioned. I don't recall the details of how

it actually made the decision, but it would look at these -- at this knowledge that it had about how the shader was operating to determine which of these eligible waves should grab available resources.

In some cases, for instance, there may not be enough resources to run a vertex shader, but there would be resources to run a pixel shader. So it would probably launch the pixel shader in that case, because the vertex shader maybe couldn't get access to enough GPRs so it could execute.

Q. What made that arbitration decision?

A. That would be the sequencer.

Q. So the sequencer would arbitrate between vertex and pixel shading?

MR. McNAMARA: Objection.

A. Yes. Again, there's the arbitration associated with launch as well as the arbitration associated with on any given cycle, because the shader has multiple waves loaded in it, some of which could be vertex, some of which could be pixel.

Once they have loaded, those waves

may execute in some arbitrary order, and the sequencer would have to choose that as well.

So I'm just trying to distinguish between execution arbitration and launch arbitration. Launch arbitration involves loading the wave into the shader and grabbing the resources that that wave will need for its execution.

Q. And what performed launch arbitration?

A. I want to say that the sequencer performed both of those tasks, but I'm not 100 percent sure the launch arbitration was in the actual sequencer block. There could have been a separate launch arbitrator block.

Q. Okay. How would you find out if that was the case?

A. You would have to go through the RTL code. Maybe going through the sequencer documentation would tell you that.

Q. But the execution arbitration was definitely performed by the sequencer?

A. Yes.

Q. The third challenge you mentioned was fetching data. How did you solve that

challenge?

MR. McNAMARA: Objection.

A. I think we looked at a number of ways of trying to separate vertex fetches, fetches from the vertex shader from fetches from the pixel shader.

I think ultimately in R400 we decided that we could live with the interaction of the two, and we didn't do much to actually solve the problem. We were not sure how serious the problem actually was, and the number of options that we looked at seemed to have their own problems.

In R500, I believe that we went to a separate path for the vertex fetching. I'm not sure what Xenos had, whether it had the R400 approach or the R500 approach.

Q. So in the R400, the same component fetched data for both vertex and pixel shading?

MR. McNAMARA: Objection.

A. Yes. To the best of my recollection, yes, that is true.

Q. And what components performed that data fetching?

A. That would be the texture unit or

the TP, the texture processor.

Q. Is that texture unit or texture processor separate from the unified shader?

MR. McNAMARA: Objection.

A. I would call it part of the shader system, but it was not -- it was a separate block from the unified shader.

Q. Okay. Then the last challenge, the fourth challenge you mentioned, related to allocating resources. How did you solve that challenge?

A. We had a number of schemes. One was kind of just a fixed pool for vertices versus pixels for the GPRs. The GPRs were the biggest challenge in terms of shader resources, because they vary as you go from application to application, and even within an application.

They're just how many GPRs you need is simply based on what the compiler has chosen to ask for for a given shader program.

So our base solution for that was to just use two separate pools, each of which was organized as a ring buffer organization. So the GPRs themselves were not shared.

On top of that, we had a mechanism

that would allow the resources to migrate from one pool to the other based on what the hardware perceived to be the loading. So it would try to move GPR resources from the pixel shader pool to the vertex shader pool when it thought that it was safe and it could manage to do so.

The GPRs were in the same physical memory. So the separation of the pools was just some very simple control logic that could be managed by the sequencer.

Ideally, if you had a very vertex-limited load, you would want to give as many of those GPRs as you could to vertex shading. Similarly, if you had a very pixel-dominant load, you would want to have as many GPRs available as you could for pixel shading.

Q. You said vertex-limited load. Did you mean vertex-dominant load?

A. Dominant load. If your bottleneck was vertices.

Q. I believe earlier you referred to a deadlock problem?

A. Yes.

Q. Can you describe that?

A. Yes. So as I mentioned, you need space to output from the vertex shader. If you can't output, then the vertex shader waves cannot release their resources and allow somebody else into the shader.

Ideally, you wouldn't want to preallocate space prior to launching the waves, because the problem is that that space is just wasted during the execution. You don't need that space until the very end of the execution.

The problem is if you wait until the end, and that space is not available, one reason why it may not be available is because you can't process pixels, because processing pixels is the way that that buffer drains.

So it fills by vertices -- by vertex shaders writing into it. It drains by those vertices being processed into pixels and those pixels making it through the pixel shader.

It may not drain if the vertex shader is trying to write into it. It can't write into it, because it's full, and it's never going to empty, because pixel shading is

not going to be able to execute, because it can't get into the shader, into the unified shader, because the unified shader is filled with vertices.

So ideally you would want to come up with a solution that doesn't have deadlock, but also doesn't have the preallocation. We didn't have a solution for that. We just preallocated, and that meant we needed a bigger buffer than we would otherwise need.

Q. Okay. One other thing you mentioned earlier was interleaving between pixel and vertex operations.

A. Yes. So you have to decide -- you have a lot of waves active, that is loaded, launched, in the shader. I think R400 could have up to 16 of them. Some of those are pixel waves, some of those are vertex waves.

You only have one shader that can execute. So on any given cycle, only one of those waves can be executing an instruction.

So you have to choose which of the launched waves you're going to run, and part of that choice is choosing whether you're going to run a vertex wave or a pixel wave.

Another part of that choice, even after you've made the choice between vertex and pixel waves, you may have multiple pixel waves loaded or multiple vertex waves loaded. You have to make a choice which of those waves is going to run.

MR. SCHWENTKER: I think the tape is almost out. Why don't we take a break now.

VIDEOGRAPHER: The time is 10:59. We are off the record.

(Recess taken at 10:59 a.m. and reconvening at 11:15 a.m.)

VIDEOGRAPHER: We are back on the record. The time is 11:15.

BY MR. SCHWENTKER:

Q. Before the break, we were talking about interleaving of vertex and pixel shading waves, I believe is the term you used?

A. Yes.

Q. I just want to make sure I understand what that looks like. I think you said that there's vertex waves and pixel waves, and then they have to be operated on by the same shader?

A. Right. Because it's a unified shader.

Q. Okay. So how does that work?

MR. MECHELL: Objection to form.

A. So one of the waves, based on some kind of priority heuristic, is going to be chosen and submitted to the shader. And as I mentioned, they get submitted a clause at a time, which is a unit of uninterruptible instructions. It's not the whole shader, it's just a small segment of instructions.

So that wave will be submitted, and it will run, and when it's done running its clause, another wave will be taken out of the shader and be put back in kind of a waiting mode by the sequencer, and the sequencer will submit another wave to be executed.

Again, it will execute a single clause, and that next wave may be a vertex wave, it may be a pixel wave, it may be the same wave over again.

Q. Okay.

A. I should add that I think, as I recall in the R400, actually the sequencer picked two waves to be executing on the shader

at the same time. When I say at the same time, they were not actually executed at the same time, but they were executing in a finally interleaved fashion.

So on the odd cycles you would run one wave, and on the even cycles you would run another wave. That has to do with just the details of how we handled some of the timing aspects of the unified shader.

Q. So in that last example that you -- your description of the R400, the sequencer would pick two waves to be executing at the same time, but only a clause from one of those waves would actually be executed by the unified shader at any particular time?

MR. MECHELL: Object to form.

A. I'm saying there would be two different waves active on the unified shader at any given time. When one of them ended, and they wouldn't necessarily end at the same time, they wouldn't necessarily start at the same time, but when one of them ended, the sequencer would pick another wave to submit and take its place in executing on the actual shader hardware.

So that fine grain interleave was just to note that there were actually two waves executing on the ALU hardware at any given period of time in that odd/even interleave fashion.

So if you think about a wave that is submitted to the shader, it will execute on cycle zero, on cycle two, on cycle four, on cycle six, maybe if it only had four instructions, and maybe right after cycle two, the other wave that was executing on the odd finished up.

So the HLSQ would find another wave to execute on the odd cycles, which then when the first wave executed cycle four, the new wave would execute cycle five, the old wave execute cycle six and finishes, but the new wave is going to execute cycle seven, and may continue on cycle nine and cycle eleven.

Q. Okay. You referred to the HLSQ. What did you mean by that?

A. I'm sorry, I shouldn't have said that, because that's not the acronym used by AMD. It's the sequencer. SEQ is what AMD called it.

I should also point out what I was describing is just the waves running on the ALU unit. There were other execution units associated with the shader.

For instance, there was the texture unit that we noted, and you could have a wave executing on the texture unit and a different wave executing on the ALU unit.

So these are kind of parallel execution units, and the sequencer's job is to find waves for all of these execution units.

Q. So what is the difference between the texture unit and the ALU unit?

MR. MECHELL: Objection to form.

A. The texture unit is associated with fetching data from memory, and it will execute things like fetch textures. It also fetched vertex information from memory. So it would execute and kind of load and store instructions.

The ALU is more associated with once the data has been loaded into the shader, executing math operations on it, so adds and multiplies, that kind of thing.

Q. I believe you said earlier today

that "thread" is another word for "wave"; is that correct?

A. In many cases it is used as another word for "wave." I really should try to stay away from using "thread," because while that is the nomenclature that Intel uses for their CPUs, in graphics some people use "thread" to mean one of the individual units within a wave that get executed as well. So it's somewhat ambiguous.

So I think it's cleaner to use "wave" rather than "thread." NVIDIA uses "warp" to mean "wave" and they use "thread" to mean an individual item within a wave. That's executing in parallel with all the other items in the wave.

Q. Was "wave" the term that you used for the R400 at ATI?

A. Yes. Though sometimes we would use "threads" as well, which became confusing. So we finally agreed to use "wave."

Q. When you used the word "thread" when you were working on the R400, what were you referring to?

MR. McNAMARA: Objection.

A. Well, the issue was sometimes we would mean wave and sometimes we would mean the individual item within a wave.

Q. And by "individual item," do you mean a clause?

A. No. A clause is a group of instructions. But a wave is a group of like pixels or vertices. So what I mean is the individual pixels, say, within a pixel wave, because the wave size that the R400 used, for instance, was 64.

So you'd have 64 pixels or 64 vertices within a wave. Sometimes I would call those 64 things a thread or sometimes I would call one of those 64 things a thread, which is why I say it can be ambiguous.

Q. Earlier when we were talking about the ALU unit, you said that it executes math operations?

A. Yes.

Q. How does it know what math operations to perform?

A. So those are contained in the instructions of the program. So the sequencer is the one that submits the instruction

associated with this particular point in the program that the ALU is supposed to execute.

Q. We have been talking about interleaving in the unified shader and the R400. Why did ATI decide to pursue interleaving in the R400?

MR. McNAMARA: Objection.

A. Well, once you have decided to have a unified shader, you have to have a mechanism for both types of waves to execute on that shader.

In some cases, you know, I suppose that you could have said, well, if I have a vertex wave executing, I'm only going to allow vertex waves to continue to execute, and I'm going to wait until all of those vertex waves are done and then I'm going to execute pixel operations. But that wouldn't be as efficient as interleaving, because a lot of times those vertex waves are not using up all the resources of the shader.

For instance, maybe you have a vertex wave operating, maybe all of your vertex waves operating, are waiting for data from memory, they're fetching vertex data in, and

there's room in the shader for more waves. You can't load any more vertex waves, because you don't have any output buffer. We already said we're not going to start up a vertex wave if we don't have any output buffer for it. But you do have resources for pixel waves.

It makes sense from an efficiency point of view to allow interleaving rather than have some kind of strict sequentiality.

Q. Is this idea of interleaving something that you all at ATI developed or is that something that was known beforehand?

MR. McNAMARA: Objection to form.

A. When we were working on the unified shader, we were not using anything else as a model. So, I mean, it's -- so we, you know, we worked on the unified shader and we said that this was -- the interleaving just seemed to us to be a natural thing to do as part of a unified shader.

So there was -- there was no prior reference that we were thinking of.

Q. What about outside of a unified shader, were there other instances of interleaving in other types of computer

processing or graphics processing?

MR. MECHELL: Objection to form.

MR. McNAMARA: Objection.

A. I've certainly seen and knew of at that point, you know, thread interleaving on CPUs, you know, being able to switch threads easily. There were CPUs that were called barrel processors that had that approach.

But I was unaware of any prior reference in the CPU world.

Q. And those barrel processors, were those made by any particular company?

A. I seem to recall a company called Hex, but I wouldn't swear to that, H-e-x.

Again, those were basically a latency hiding mechanism, and the interleaving of threads to hide latency, that aspect was not new in R400. Our previous parts used multiple, within the vertex shader or within the pixel shader, we had multiple threads or waves active, which would interleave execution, although they would start and finish in a sequential order. While they were executing, they would interleave. As a mechanism to hide latency, similar to how these CPUs had done it

in the past.

What was new for R400 was the interleaving of vertex and pixel threads within the processor, and that was new, because the unified shader was new. If you didn't have a unified shader, there would be no mechanism. There would be no way, even if you thought of the idea of interleaving, there would be no one mechanism for you to interleave on.

Q. I see. So in the R300, which had a separate vertex shader and a separate pixel shader, did the vertex shader use interleaving for vertex waves?

MR. MECHELL: Objection to form.

A. I can't say with 100 percent certainty the vertex shader used interleaving for the vertex waves. I suspect it did, but I don't have a clear memory of that.

One of the things that we enabled for the R400, and this is regardless of whether you have a unified shader or not, is the ability for the shader to fetch from memory.

In the R300 and prior, there was a dedicated fetch unit, and the vertex shader would only do ALU operations. So the need for

latency hiding is a lot less when you don't have to deal with a memory fetch. There is still some benefit associated with it. But I don't recall whether there was actual interleaving of threads within the vertex shader.

Q. What about for the pixel shader?

A. For the pixel shader, there definitely was.

Q. There was?

MR. MECHELL: Object to form.

A. Yeah, because the pixel shader dealing with memory was not a new thing. So latency hiding was something, was a problem, that had to be addressed, and we addressed it. The R300 wasn't the first to do thread interleaving to be able to address it.

Q. Was that something that you had used in the prior products as well?

MR. McNAMARA: Objection.

Q. Such as the Radeon 100 and Radeon 200?

MR. McNAMARA: Objection.

MR. MECHELL: Same objection.

A. I don't believe that we did,

because the R100 and R200 were not the -- the shader portion of the ALU was not a general purpose shader. So it was a hard-coded thing that you had some dedicated -- you had a fetch stage and then you had an execution stage and you had, you know, the render back end stage.

So there wasn't the kind of intermixing between ALU instructions and memory fetch instructions that would require the need for latency hiding and multi-threading to be able to deal with it as a mechanism.

So I can't say 100 percent, but I have no memory of that kind of interleaving in the R100 and R200.

Q. Just to make sure I'm clear, what was being interleaved in the pixel shader of the R300?

A. So you could have multiple waves executing, and they could be executing at different points in their programs, and they would interleave their execution.

So while one wave was fetching data from memory, it would then go to sleep until the memory showed up, and other waves would execute on the shader system. And you

were able to have waves at different points in their execution, executing on the shader system, while other waves were executing in the texture system.

Q. So if I understand correctly, the R300 did interleaving between color shading and texture shading?

MR. MECHELL: Objection to form.

A. Texture fetching, I would call it. All shading in the R300 was done in this one shader unit, in this pixel shader unit.

Q. So the R300 did interleaving between color shading and texture fetching?

MR. McNAMARA: Objection.

MR. MECHELL: Same objection.

A. Yes.

Q. Do you recall when the R300 was launched?

A. Not off the top of my head, I'm afraid.

Q. Was it before or after you started working on the R400?

A. It was launched after we started working on it.

Q. And I think you said that a team

in another office was working on the R300; is that correct?

A. Yes. So the Silicon Valley office was the driving force between -- on the R300. Although, the Orlando office that I mentioned earlier provided the vertex shader for the R300.

Q. Okay. So while those offices were finishing up the R300, the Marlboro office started work on the R400?

A. Yes.

Q. Was the R400 ever taped out?

A. No.

Q. Why not?

A. The design was not finished before we retargeted its market window and changed the name to the R500, as well as added features associated with that market window change.

Q. And do you recall when that took place?

A. I do not.

Q. How far did the R400 get?

A. Well, I mean, we were doing simulations of it, you know, it was passing tests. I remember that. We had enough -- we

were -- we had, you know, confidence in the shader system, and there was, you know, the Xenos effort going on along the same time that used a lot of the R400 as its base.

Q. And do you recall when you started modifying the R400 to the Xenos?

A. No.

Q. You said that the R400 was passing tests. What kind of tests were you --

A. Well, it was certainly passing block-level tests. I know that the shader was running. I don't recall whether we were, you know, passing entire system tests, including data going out to the memory.

So whether the whole thing was passing tests -- I believe that it was prior to the R400, R500 change, but I couldn't say with absolute certainty.

Q. When you said that it was passing block-level tests, what kind of tests are you referring to?

A. So we had a C model, which modeled how the chip works, and it modeled it in enough detail so that at each of the block interfaces it would match the sequence of outputs.

So you could run a given test, and a test here is maybe a piece of an application, maybe it's a dedicated test to check a particular feature, maybe it's a random test where we just set up the registers randomly and make sure that it matches the RTL, the source code for the hardware matches the output of the C level model.

So you could check at the whole chip level, meaning that make sure they both get the same input and that they match at the same output, and we just check at the frame buffer level, or you could replace pieces of the C level model with their equivalent hardware model, and then check to see that it matched at the interface, as well as produced the same frame buffer image.

So we could run those tests on various blocks by replacing the block in the C model with the RTL, just checking at the boundaries, and using the data, feeding the RTL with the input from the C level model, and then taking the output from that block and pushing it onto the next blocks in the high level model.

So you could isolate each of the individual blocks and understand when you had a difference at which block were you seeing the difference. When I say "difference," I mean different from the C level model.

Q. The C model was developed before the RTL?

A. Yes.

Q. Was there ever a complete set of RTL for the R400?

MR. MECHELL: Objection to form.

A. I don't know. I don't know whether we actually finished it, but it just missed the deadline for the rest of the process to produce a part or, you know, there was still pieces of it that were not entirely finished prior to the R400 to R500 switch.

Q. Was there ever a complete set of RTL for the sequencer?

MR. McNAMARA: Objection.

MR. MECHELL: Objection to form.

A. I believe so. I think that the sequencer was in a producible -- I don't know if it was like passing every single test that we had or plan to write, but I think it was not

what stopped us from producing the part.

Q. Do you recall when a complete set of the RTL for the sequencer was completed?

A. No.

Q. How would you figure that out?

A. I think somehow you would have to compare the RTL versions. So in ATI's database, it's all versioned by the source control system against to see what level of tests you're passing that you would say this is good enough to be deemed complete.

Hardware is never really complete. If you give us extra time, we'll continue to fool around with it. But, you know, I think that you can say, you know, passing all directed tests, for instance, would be a level of completion.

So you could check by whatever test logs ATI or AMD has from that era, and, you know, see whether the sequencer was still failing any directed tests.

Q. Was there ever a complete set of RTL for the unified shader?

MR. MECHELL: Objection to form.

Q. In the R400?

A. Again, I can't say with 100 percent certainty, but I believe the answer is yes.

Q. Do you recall when?

A. No.

Q. When you started working on the R400, did you have discussions with the team working on the R300?

A. Yes.

Q. What discussions did you have?

MR. MECHELL: Objection.

A. I think a lot of them revolved around the differences in how it would look to the outside world. I mentioned the difference in precision earlier, and at the same time we were discussing with Microsoft what DX9 would look like. So we wanted to make it clear to Microsoft what we needed to fully support the R300, as well as what was needed to fully support the R400.

So what happened there was Microsoft actually came out with two different versions of DX9, one of which kind of mirrored the support of R300, and another one which mirrored the support of R400.

Q. And remind me what DX9 is?
A. DX9 is an application interface. So it's the way that applications, like games, talk to the hardware. So it's a set of calls that the application makes to get the hardware to do a particular task. It evolves with each generation of hardware to be able to surface the applications the new capabilities of hardware.
So DX9 is a specification really that comes from Microsoft.
Q. And do you recall when Microsoft came out with the version of DX9 that supported the R400?
A. I don't.
Q. Is DX an abbreviation for something?
A. DirectX.
Q. Are you familiar with OpenGL?
A. Yes. That's another API. It's a competing API.
Q. Have you done any work with OpenGL?
A. I've been involved with understanding what the specification requires,

and I've been involved with, you know, trying to craft extensions to it to support additional hardware features.

I haven't actually written code that interfaces with OpenGL.

Q. Okay. Did the R400 support OpenGL?

A. The R400 was certainly capable of supporting OpenGL. I do not know -- I don't recall whether we wrote an OpenGL driver for it. I suspect that we did, but I don't recall specifically for it or whether there was an effort to support OpenGL. I think that we probably did, because we typically did.

Q. When you say you typically did, do you mean for each of the ATI products you would --

A. Yeah. I know there were OpenGL drivers available. I don't recall whether they came from us or whether that was some kind of third party or independent effort to have OpenGL running on the parts.

Q. Why did you say that the R400 was certainly capable of supporting the OpenGL?

A. Well, because the requirements for

running DX9 were similar to the requirements for OpenGL, the underlying hardware requirements, and we did -- we were aware of the OpenGL requirements as well. So we made sure that we had capabilities for some of the OpenGL corners.

I can particularly think of like line stenciling, which is a GL thing and not a DX thing, and we made sure we had a method for doing that.

(Document marked as Exhibit 4 for identification - withdrawn)

MR. McNAMARA: Counsel, do you have authorization to show this to the witness?

MR. SCHWENTKER: Mr. Mechell can, I believe, confirm this, but my understanding is that this was made public by the PTAB in connection with the IPR of the '871 patent.

MR. MECHELL: I know that there are a number of documents that the PTAB denied AMD's request to maintain confidentiality on, and in other words are in the public record. However, I'm

not sure if this exact document is in the public record without looking at the corresponding public document.

MR. McNAMARA: Do you want to skip over this and just ask at lunch?

MR. SCHWENTKER: Sure.

MR. McNAMARA: Mark it as Exhibit 4 and then we'll figure out whether or not...

MR. SCHWENTKER: Sure. If you could set that to the side.

MR. McNAMARA: Counsel, we'll confirm on our end whether or not there's any issue with it at lunch, if that's all right?

MR. SCHWENTKER: Sure.

MR. McNAMARA: Could we actually have two minutes?

MR. SCHWENTKER: Yes, sure.

VIDEOGRAPHER: The time is 11:59 and we're off the record.

(Recess taken at 11:59 a.m. and reconvening at 12:10 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 12:10.

BY MR. SCHWENTKER:

Q. Mr. Gruber, when you were working on the R400, did you write any documents, any specifications?

A. I did.

Q. Do you recall what specification?

A. I believe I was involved with the shader spec. I don't recall what it was called specifically, but the spec that spoke about what instructions the ALU could operate and the description of the shader execution mechanism.

Q. Would it have been called shader processor?

A. Yes, that sounds right.

(Document marked as Exhibit 5 for identification)

MR. SCHWENTKER: Before you hand that to the witness, just to make sure there are no objections to handing this to the witness.

MR. MECHELL: Do you mind if we go off the record for a second?

MR. SCHWENTKER: Sure.

VIDEOGRAPHER: The time is 12:12 and we are off the record.

(Luncheon recess taken at 12:12 p.m. and reconvening at 1:04 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 1:04.

BY MR. SCHWENTKER:

Q. Welcome back, Mr. Gruber. So during the break, I confirmed, and I will let counsel confirm this as well, but I believe we've confirmed that Gruber Exhibit 5, which is a shader processor specification, with Andrew Gruber listed as an author, that there's no objection to showing that as an exhibit?

MR. McNAMARA: No objection.

MR. MECHELL: That's correct, no objection.

Q. Okay. Mr. Gruber, the court reporter has handed you Exhibit 5. For the recorder, it's "Shader Processor Revision 1.2," with Bates numbers AMDLG0147144 through 147185.

A. Yes.

Q. And before the break, you said that you did write a document called shader processor?

A. Yes, I did.

Q. Is this the document that you

wrote?

A. Yes, it is.

Q. So what is this document?

A. This document describes the interfaces and the operation of the shader processor. It also lists all of the op codes as well that the shader processor could execute.

Q. Which page are you referring to?

A. When I said the op codes, I meant 3.3, which is Page 11 of the document. AMDLG0147154.

Q. What was the purpose of writing this document?

A. To describe what the shader processor does, as well as to act as an implementation spec for the implementers of the shader processor block.

Q. If you could turn to Page 7 of the document. Do you see the Section 1.2.1 titled "Vertex Shader" and Section 1.2.2 entitled "Pixel Shader"?

A. Yes.

Q. What do those sections refer to?

MR. McNAMARA: Objection.

A. They refer to the shader when it's operating on vertices versus the shader when it's operating on pixels. These two paragraphs just describe the overall operation that gets on the starting conditions when the shader starts up.

Q. When you refer to the shader, you're referring to the unified shader?

A. The unified shader, yes.

Q. So depending on whether the unified shader is working on vertex operations or pixel operations, it is, I guess, referred to as the vertex shader or pixel shader?

MR. McNAMARA: Objection. Mischaracterizes.

A. Correct.

Q. Could you turn to Page 15, please. Do you see a Section 4.1 titled "Shader as an SIMD Architecture"?

A. Yes.

Q. What does that mean?

A. By "SIMD," we mean single instruction multiple data. So the idea is that when we were talking about waves earlier, that is the multiple data that is, in this case,

there were 64 elements of the wave that all are executing a single instruction at any point.

So the idea is you have a single instruction, but it executes in parallel over 64 elements.

Q. What do you mean by "elements"?

A. Pixels or vertices.

Q. Do you see the first sentence in that section, "As shown in the diagram below, four identical processing units comprise a shader unit"?

A. Yes.

Q. What does "processing unit" refer to?

MR. McNAMARA: Objection.

A. In this case it refers to an ALU data path that is capable of processing 32-bit floating point operations in it.

Q. What does "shader pipeline" refer to? Strike that.

Do you see in the next sentence it says, "There are four shader units in one shader pipeline"?

Do you see that?

A. Yes.

Q. What does "shader pipeline" refer to?

MR. McNAMARA: Objection.

A. In this case, the structure of the shader has four individual units that execute the same instruction, but one cycle offset from each other.

So the four units that execute across that pipeline -- across that single instruction across these four cycles are just grouped into a single unit that it refers to a pipeline.

Q. Earlier you were talking about interleaving vertex and pixel shading operations. How does that correlate to the shader units and shader pipelines discussed here?

A. It's somewhat orthogonal to this. So we launch -- in this case, you know, we launch an instruction, and what happens is that the instruction winds up getting launched over four clock cycles into each of these shader pipeline units, but the interleaving at a vertex or pixel level happens kind of higher than that, where it's based on what instruction

is picked as the next to launch over the next four cycles.

So, I mean, once I've chosen to launch an instruction, the fact that it gets actually launched over these four cycles is just an implementation detail.

Q. But once an instruction is chosen, that will be --

A. Yeah, so we'll have 64 elements in that instruction, and, you know, as this shows, you know, so we'll launch these first four quads, a quad is four, so you can think of like 16 elements out of that 64, get launched on the first cycle and 16 on the next cycle and 16 on the next cycle and 16 on the next cycle.

Q. Could you turn to Page 16 of the document. Do you see a diagram on Page 16?

A. Yes.

Q. Could you walk me through this diagram?

MR. McNAMARA: Objection.

MR. MECHELL: Objection to form.

A. So what this is showing is the path to load up either the initial vertex information or the initial pixel information as

a part of kind of loading a wave into the set of GPRs and MACs, MACs meaning multiplier accumulator. GPR is the storage and MAC is the data path.

So what you see at the top, you see a selection whether you're going to take in indices, which would be for the vertex, or you would take in the interpolated what we call barycentrics; that's the IJ data. So that shows how we initialize the data for a wave launch.

On the side, this TFM and the TAM, I believe those are the texture unit. So this is showing how we issue texture addresses to the TAM, and then the TFM brings the data back in, again, loading the GPR in these individual lanes.

So what this is trying to show is the data path for getting data into and out of the GPRs.

What this is showing at the bottom where it says "PC," that's how the output of the vertex shader works. So the output comes directly out of the data path, that is the MAC, and gets written into these PCs, which stand

for parameter cache, which -- I don't know whether it stands for parameter cache or position cache, but it is used to hold the output of the vertex shader. That purple data is then recycled as part of rasterization, and it goes into the interpolators to generate the barycentric information associated with pixel waves.

Do you see at the bottom of that PC block there's also something called OB, which is the output buffer, which is where the output of pixels goes when they're coming out of the shader, and those, rather than being recycled, at 250 bus, then goes back to the render back end.

Q. You can put that exhibit aside. Before the lunch break, we had some discussion about exhibits and whether to use exhibits from the IPR proceedings that the PTAB released to the public as opposed to the versions of the same documents that were produced with confidentiality designations in this case, and we've decided to use the public versions of those documents that were released by the PTAB.

So Exhibit 4, which was the R400

architecture proposal, I would propose swapping out this PTAB version for the previous version, if there's no objection to that.

MR. MECHELL: That's fine from our perspective. Just as a note, I'll note that AMD did produce the PTAB record in the Northern District of California action without a confidentiality designation. So at some point we'll just need to go back and correlate those, but we can do that later.

MR. SCHWENTKER: Okay. But you would agree that Exhibit 2040 from the IPR proceeding was released to the public?

MR. MECHELL: Yes. It was made publicly available over AMD's objections.

MR. SCHWENTKER: Any objections?

MR. McNAMARA: No objections.

(Document marked as Exhibit 4 for identification)

Q. Mr. Gruber, this is a document titled "R400 Architecture Proposal."

Do you see that?

A. Yes.
Q. Just for the record, on its face is labeled at the bottom right ATI 2040 LG v. ATI IPR2015-00325. Do you see that?
A. Yes.
Q. Have you seen this document before?
A. Yes, I have. I don't know if I've seen this particular version, but I've seen at least some version of this document.
Q. So you've seen the architecture proposal for the R400?
A. Yes.
Q. Did you work with -- strike that.
On the first page it lists Steve Morein as the author?
A. Yes.
Q. Did you work with him on this document?
A. I did not work with him on the writing of the document. I was more of a reviewer of the document.
Q. If you could turn to Page 6. Do you see Section 1.4 titled "Unified Processing Pipe"?

A. Yes.
Q. In that first sentence it says, "The most ambitious feature in this design is the 'truly unified pipe'"?
A. Yes.
Q. "A single programmable pipeline is used for 2D, video, 3D vertex, and 3D pixel operations."
Do you see that?
MR. McNAMARA: Objection.
A. Yes.
Q. Do you know what "truly unified pipe" means?
A. It was referring to merging the vertex processing and the pixel processing.
Q. Is there a distinction between a unified pipe and a truly unified pipe?
MR. McNAMARA: Objection.
A. I think the reason why Steve used the term "truly unified pipe," is that the R300 had something that they called a unified pipe.
What they were referring to is, you know, I mentioned how in the R200 you had kind of a texture-fetched stage and you had an ALU stage, and they were purely a pipelined

thing, where they didn't interact.

While in the R300, those two stages were merged so that the texture fetching was done by the same unit that did the ALU operations. This is all within the context of pixel shading.

But rather than having a separate texture addressing device that would fetch at textures and feed them to the ALU, they had a single device that could fetch textures and do ALU operations in an interleaved fashion.

The R300 documentation called that a unified shader. It wasn't unified in the sense of vertex and pixel being unified, and I think this "truly unified" was meant to distinguish the two.

Q. Could you turn to Page 8 of Exhibit 4. Do you see Section 3 titled "Schedule"?

A. Yes.

Q. And it was three dates there. "Tapeout April 2, 2002, Samples May 2002, Production November 2002."

Is it fair to say that that schedule was not met?

A. Yes.
Q. Were any of those deadlines met for the R400?
A. No. I mean, you know, I don't know what the R400 actually -- well, I know the R400 wasn't actually taped out. So, none of those were met.
Q. If you could go back to Page 6, that Section 1.4 that we were talking about earlier.
A. Yes.
Q. The first sentence that we read previously there talks about 2D as well as 3D vertex and 3D pixel operations.
A. Yes.
Q. What are the differences between 2D operations and 3D vertex and 3D pixel operations?
MR. McNAMARA: Objection.
MR. MECHELL: Objection to form.
A. 2D operations are usually much simpler operations than 3D operations, either 3D or vertex -- 3D pixel or vertex.
It is possible to map most 2D operations into simple 3D operations or you

could have a separate simpler engine dedicated to 2D operations.

Video is typically similar, but it usually involves special formats that are used for video, other than the RGB type formats that are typically used for graphics.

So if you want to support video, you typically have to understand these different formats, but you can sometimes use the shader to convert from the typical video color space into a RGB color space.

Q. In the next paragraph it says, "There is an area cost to the unified pipeline since we are forced to go to 32-bit precision for color, when application requirements may need less (22 to 24 bits)."

Do you see that?

A. Yes.

Q. What does that mean?

A. Here he's specifically talking about the pixel processing where pixel processing in general doesn't need as high precision as vertex processing needs.

So he's saying here at least on the pixel processing, things that we do, we're

going to spend more for that.

He goes on to say in this paragraph that by having a single structure, you know, where you're kind of sharing it among vertices and pixels, that you can optimize it.

Q. The last paragraph of that section, the last sentence says, "The unified pipeline presented here dynamically allocates its processing power between transform and raster."

Do you see that?

A. Yes.

Q. What do "transform and raster" refer to?

MR. MECHELL: Objection.

A. The transform refer to vertices. The typical job of the vertex shader is to do a matrix transform of the vertices from object space to -- well, not really to screen space, but what's called homogeneous space, which is oriented toward where the camera happens to be pointing. By raster, he means pixels, because the result of triangle rasterization are pixels.

Q. Is pixel shading part of the

rasterization process?

MR. McNAMARA: Objection.

A. Rasterization is sometimes used in the very narrow sense, meaning converting a triangle to a sequence of XY locations of the pixel within the triangle, and in that sense the shader is not part of the rasterization. But it's often used in a wider sense, meaning once you have a triangle, drawing that triangle onto the screen, and in that case the shader is used in that wider sense.

Q. You can set that Exhibit 4 to the side.

MR. SCHWENTKER: I hate to do this, but I think it might benefit us all if we go off the record for a few minutes. Since we just printed out these documents, I want to try to winnow them down to the extent possible. So if we could go off the record for a couple minutes.

VIDEOGRAPHER: The time is 1:40 and we're off the record.

(Recess taken at 1:40 p.m. and reconvening at 1:57 p.m.)

VIDEOGRAPHER: We're back on the record. The time is 1:57.

(Document marked as Exhibit 6 for identification)

BY MR. SCHWENTKER:

Q. Mr. Gruber, the court reporter has handed you Exhibit 6. For the record, it's labeled on the bottom right-hand corner ATI 2041 LG v. ATI IPR2015-00235.

Do you see that?

A. Yes.

Q. This is another document that was -- released to the public by PTAB over AMD's objections in the IPR.

MR. SCHWENTKER: So I understand there's no objections to showing this to the witness?

MR. MECHELL: Correct.

Q. Have you seen this document before?

A. I think I have.

Q. On the first page it's titled "R400 Top Level Specification."

Do you see that?

A. Yes.

Q. Would you have seen this as part of your work on the R400?
A. Yes.
Q. On the first page, do you see that the author is listed as Steve Morein?
A. Yes.
Q. Did you review his work on this document?
A. Yes.
Q. Did you provide input into it?
A. I don't recall. I'm sure that I did, though. I'm sure that I gave him feedback. I don't know if I actually gave him, you know, actual text or sections.
Q. Could you turn to Page 15. Do you see the block diagram in the middle of the page there?
A. Yes.
Q. Is this a block diagram of the R400?
A. Yes, it is.
Q. The date on this document is March 11, 2001. Do you see that?
A. Yes.
Q. So this is the -- this was a top

level block diagram of the R400 as of that date?

A. Yes.

Q. Do you recall if the block diagram of the R400 changed after this?

MR. McNAMARA: Objection.

A. It might have. The one thing that I'm looking at here is the way that the texture pipe is independent for each of the groups of four shader pipes. That doesn't match my recollection of the final organization of the R400, where there was more of the single texture pipe shared across all of them.

So I wouldn't swear that that aspect didn't change, but the rest of it looks rather similar to my recollection of the final.

Q. I'd like to walk through some of these blocks on the diagram.

A. Okay.

Q. In the top left-hand corner there's a block labeled "HI"?

A. Yes.

Q. What is that?

A. That stands for host interface. That's when the host was directly talking to

the chip, sending in perhaps command register information.

Q. The host being --

A. The host being the main CPU, like an Intel processor in a PC.

Q. To the right of that there's a block labeled "CP/RBBM"?

A. Yes.

Q. What is that?

A. It stands for command processor, and the slash RBBM is registered backbone bus manager. The command processor is something that interprets a command stream.

So the command stream is in kind of an abstract language, as opposed to particular registered to be written in the chip, and the CPX is kind of a translator and changes that abstract command stream into a sequence of onboard register rights.

The CP itself was a processor. It was a homegrown processor, and would execute instructions. The RBBM is just the way that CP could write registers to the individual blocks within the chip.

Q. Below those there's a "Primitive

Assembly." Do you see that?

A. Yes.

Q. What did the primitive assembly do?

MR. McNAMARA: Objection.

A. The primitive assembly was in charge of kind of managing the execution of a draw call as the -- from the vertex point of view.

So it would -- it would fetch the indices associated with the vertices, and it would, using those indices, it would know whether that index was already transformed through the vertex shader or whether it had to be submitted to the vertex shader for vertex processing.

And it would also manage the parameter cache and the position cache that is the output of the vertices, and as that data came out -- I mentioned a setup engine, and that setup engine was part of the primitive assembly as well -- it would take the data out of the position caches and the parameter caches, and it would do a final hard-coded transform on them, what's called a viewport

transform, and it would do the setup in the format that the rasterizer needed it to produce pixels.

Q. We've talked about the unified shader today?

A. Yes.

Q. What corresponds to the unified shader in this diagram?

A. In this case --

MR. MECHELL: Object to form.

A. -- all the blocks labeled "Shader Pipe" would be the unified shader, as well as, you know, I suppose you would include part of the block labeled "Rasterizer Sequencer."

Q. Why do you say you would include part of the block labeled "Rasterizer Sequencer"?

A. Well, I think the sequencer would be considered part of the shader pipe in that it controls the execution of it. The rasterizer, not so much.

I think in later block diagrams, they were probably separated out into two different blocks.

Q. So you said the unified shader

would be the -- strike that.

You said that the shader pipe, plus the --

A. Plus the sequencer.

Q. -- would be the unified shader?

A. Yes.

MR. MECHELL: Objection to form.

Q. There are a series of four shader pipes -- well, actually, strike that.

There are 16 shader pipes --

A. Right.

Q. -- four per rasterizer/sequencer block?

A. Right.

Q. How many shader pipes would be in a unified shader?

MR. MECHELL: Objection to form.

MR. McNAMARA: Objection.

A. This is kind of an arbitrary division based on the implementation. Each of those represents a quad, but they're all executing the same instruction. So they're all really part of one shader pipe.

Q. Part of one shader pipe or --

A. I'm sorry, one unified shader. If

you go back to the other diagram we were looking at in the -- the shader processor, these correspond to each of the individual quads, quad pipes that were listed in that diagram as like just the next subdivision of a single shader pipe.

Q. Just so I'm clear, the group of four shader pipes would be all within one unified shader?

MR. MECHELL: Objection to form.

MR. LEVENTHAL: Objection to form.

MR. McNAMARA: Objection.

A. No. I'm saying that all of these pipes were in a single unified shader.

Q. All 16 of them?

A. Yes. At least the way that the chip was finally implemented, you know, this is a fairly early block diagram, and I think that some of the functions have moved around a little bit since this diagram.

Q. Okay. Earlier today you talked about two different types of arbitration; launch arbitration and I think execution arbitration. Is that accurate?

A. I said launch arbitration and

execution arbitration, yeah.

Q. Okay. What in this block diagram performed launch arbitration?

A. The sequencer would have.

Q. And what about execution arbitration?

A. The sequencer would have as well.

Q. And did that change during the development of the R400?

A. I don't recall whether the launch arbiter was part of the sequencer in what we finally implemented or not.

My recollection is that it was, but I'm not certain about that.

Q. But at least as of March 2001, it was part of the sequencer?

MR. MECHELL: Objection to form.

A. Yes.

(Document marked as Exhibit 7 for identification)

Q. The court reporter has handed you Exhibit 7. It has Bates numbers AMDLG0206392 through 206400.

A. Okay.

MR. SCHWENTKER: I understand from

counsel from AMD there's no objections to showing this document to Mr. Gruber?

MR. MECHELL: That's correct.

Q. Mr. Gruber, do you recognize this document?

A. I have no reason to doubt its authenticity, but I don't recall it.

Q. On Page 2, do you see that it says "Prepared by Andrew Gruber"?

A. Yes.

Q. Is that you?

A. Yes.

Q. So you have no reason to doubt that you prepared this?

A. Right.

Q. Do you know what this document is?

A. I'm just looking at it, it looks like a quick overview of the project.

Q. At the bottom in the right-hand corner it says, "December 13th, 2001"?

A. Right.

Q. Do you have any reason to think it's not from that date?

A. No.

Q. Was this -- do you know if this is

a presentation you gave?

A. I think it was. I wouldn't have prepared this except for a presentation. This is probably at the time when we were trying to present the R400 to a wider audience.

We had done enough work on it with a small group of people to have confidence in the overall plan.

Q. Do you know who the audience would have been?

A. It certainly would have been within ATI. I imagine it would have been an engineering audience.

Q. Why do you say that?

A. Because that was the primary audience that I would address within ATI. ATI is a very engineering-oriented company, and, you know, this, I think, would have been to explain where we were going with the next generation to people who were focused on the current generation. There might have been a marketing audience as well, a product marketing audience.

Q. On slide 3, do you see where it says "Plans/Deliverables"?

A. Yes.
Q. Do you know what "Deliverables" means there?
A. Well, I think it refers to the "Architecture Deliverables," on the next pages. So what we as an architectural group were planning to produce to prepare for the implementation.
Q. Do you know if at this time the deliverables listed on Pages 4 and 5 had been completed?
A. I don't know. I don't know whether this was saying this is what we have or this is what our plan to produce is.
Q. How would you figure that out?
A. Well, I would see the dates on some of this documentation in ATI's version control mechanism and see what state they were in at the time of this.
It's possible that even the initial versions were not in the version control at this time. So that would imply that this is a promise to the future rather than something that was currently present.
Q. If you could turn to slide 6, it's

titled "Current Status."

A. Right. So this implies to me that these deliverables are not available at the time of this documentation.

Q. Why do you say that?

A. Well, because the status here, you know, most -- first pass instruction set done versus something like a programming guide or a shader guide, I mean, you have to have an established instruction set to be able to write those other things, and if you're only on the first pass of the instruction set or the first pass of the register spec, you know, you're not going to have a detailed register specification or you're not going to have a synchronization or coherency documentation.

Q. Could you turn to the next slide.

A. Yes.

Q. The second bullet point -- strike that.

Slide 7 is titled "Open Issues/Concerns"?

A. Yes.

Q. And the second bullet point says "Power consumption is a concern especially for

Mobile Parts."

A. Yes.

Q. Why was that a concern?

A. Well, this is talking specifically about the overlay functionality. By the way, this was something that was not carried forward in the R400 implementation.

But the idea was that we would remove logic from the display unit, and instead implement it via the GPU via a memory-to-memory operation.

So what you do for overlays, overlays being like, let's say that you have -- let's say you're watching a video in a window where the video would overlay the rest of your desktop.

The way that that's typically handled and was handled up until the R400 was that you would have one frame buffer sitting in memory, which would be the entire desktop, including the area that's underneath the video, and you'd also have an area in memory that represented the video, and the display processor would read both of those surfaces and mix them on the way out to the display.

And even if you went to a full screen display, you would still only be reading the same size video surface in memory. So it was the display's job to scale and position the video overlay to the screen window that the user had selected.

What is being proposed here is to remove that overlay functionality, and instead, have the graphics unit read both of those surfaces, mix them, write them out to a single surface in memory, and then have the display unit pick up that single surface, and send that single surface to the display already premixed by the GPU.

So the reason why power is a concern is that if you think about what's happening to that video surface, rather than being picked up once out of memory, now it's being picked up by the graphics unit, scaled up, written back into memory, possibly larger, and then picked up by the display again.

So you have more memory traffic and you're also using the graphics to do the scaling operation as opposed to the video processor, the display processor, to scale up.

Both of those things, both using the graphics instead of the video, instead of the display, and going through memory more times, are going to add to the power consumption of the part, although they would save area, because you would not need the scalar unit -- the scaling unit in the display or this overlay unit.

Q. That ultimately was not implemented in the R400?

A. It was not. I mean, it's really -- from the graphics point of view, whether you do this or not, there may be a scheduling issue, but the hardware is pretty much identical.

It's the issue of about whether you can pull out the logic in the display or not, and I don't think ATI ever chose to pull that logic out of the display. So I don't think graphics was used for this typically.

In some cases, you may have more than one overlay, right? You may be watching four different YouTube videos at the same time. You only have a limited number of display overlays. So any time when you run out of

those, the job is given to the GPU to do the rest of them.

But I'm saying that ATI never pulled out all the overlays, as is being suggested here.

Q. The fourth bullet on Page 7 says "Worst Case Shader is still unresolved."

Do you see that?

A. Yes. So the issue here is that if you give the shader, meaning the GPU, the job of like video scaling, you have this issue that you have to match the refresh rate of whatever the source you're doing is; usually it's 24 frames a second, 30 frames a second, 60 frames a second.

And if you're using the graphics for that job, then this task that does this scaling of the video overlay has to get in those 60 -- every new frame, you have to scale it up to whatever the user said is the frame rate.

If you're doing graphics operations, the graphics unit may be busy, and especially if there's no way to -- if a single shader is running a long time, there's no way

to get in there and use the graphics unit for this realtime video need.

So what he's referring to here is that we don't have a way of guaranteeing any kind of minimum latency needed to assure you're not going to drop frames as you're displaying this user video.

People are used to videos just working, and the concern here is that there's a chance that you'd get some dropped frames or a loss of video quality due to trying to share the graphics system with using it for these video displays.

That may have been one reason why we never actually implemented that particular aspect.

Q. Could you turn to the next page, Page 8. The top of the page says "Target Schedule."

Do you see that?

A. Yes. So again, these are some of the plans for the previous deliverables that were...

Q. Are you finished?

A. Yes.

Q. Okay. Do you see where it says "Block Diagram"?
A. Yes.
Q. And beside that it says "Done"?
A. Yes.
Q. Does that indicate the block diagram was finished at that point?
A. I think it means a version of the block diagram was finished. I think, you know, the block diagram is going to be subject to the input from the implementation team, and we may make changes.

I think what this is saying there's a block diagram done. It's not a guarantee that that isn't going to change in the future.

Q. So not necessarily the final block diagram?
A. Right.
Q. Further down it says "Shader Guide"?
A. Yes.
Q. What is the shader guide?
A. The shader guide was the instruction set architecture definition,

similar to what I pointed out in the R400 documentation, but with more information on each instruction and what it did and the limitations of that instruction.

It may also have been more documentation on the shader itself, similar to a lot of the other stuff that's in this shader processor.

I think the shader guide eventually morphed into this shader processor spec.

Q. Besides "Shader Guide" it says "Plan, 12/19"?

A. Yes.

Q. Does that mean the shader guide was not complete at this time?

A. Yes.

Q. Do you know when it was completed?

A. I do not.

Q. Do you know if it was shortly after this presentation?

A. I do not. I just don't have a memory of this.

(Document marked as Exhibit 8 for identification)

(Document marked as Exhibit 9
for identification)
(Document marked as Exhibit 10
for identification)

Q. Mr. Gruber, the court reporter has handed you three documents. The first is Exhibit 8. Do you see it's U.S. Patent No. 6,897,871?

A. Yes.

Q. And for the record, it has Bates numbers AMDLG0002111 through 2121?

A. Yes.

Q. I'm going to go through the other two real quickly. Exhibit 9 has Bates numbers AMD1044_0000165 through 176.

A. Yes.

Q. And this is U.S. Patent No. 8,760,454?

A. Yes.

Q. Exhibit 10 has Bates numbered AMD1044_0000177 through 195, and this is U.S. Patent No. 9,582,846?

A. Yes.

Q. So if I refer to these as the '871 patent, the '454 patent and the '846 patent,

will you know what I'm referring to?
A. Yes.
Q. Do you understand that you are a named inventor on these three patents?
A. Yes.
Q. Do you recognize these patents?
A. Yes.
Q. All three of them?
A. Yes.
Q. What's your understanding of these patents?
MR. McNAMARA: Objection.
MR. MECHELL: Objection to form.
Q. Let me rephrase that. Turning first to the '871 patent.
A. Okay.
Q. Exhibit 8. What's your understanding of this patent?
MR. MECHELL: Objection to form.
MR. McNAMARA: Objection to form.
Calls for a legal conclusion.
A. It covers various aspects of the unified shader.
Q. And by "unified shader," are you referring to the unified shader we have been

talking about today?
A. Yes.
Q. So your understanding is that it covers various aspects of what you were working on with the R400?
A. Yes.
Q. What do you base that understanding on?
A. It's based on the abstract description. It talks about a shader that can perform one of vertex operations or pixel operations.
Q. You said your understanding is that the '871 patent covers various aspects of the unified shader?
MR. LEVENTHAL: Objection. Form.
MR. MECHELL: Objection to form.
MR. LEVENTHAL: Asked and answered.
A. Yes.
Q. What aspects of the unified shader does it cover?
MR. MECHELL: Objection to form. Calls for a legal conclusion.
A. It appears to cover aspects of

arbitration between the vertex and pixels.

Q. Why do you say that?

A. I'm just reading Claim 1, which talks about the arbiter circuit for performing vertex and pixel operations.

Q. What does the arbiter circuit referred to there?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. The arbiter circuit arbitrates between selecting either vertex operations or pixel operations.

Q. Now, earlier today we talked about both the launch arbitration and execution arbitration?

A. Right.

Q. Does this refer to one or the other?

MR. MECHELL: Objection to the form. Calls for a legal conclusion.

A. I haven't studied this well enough to say whether it covers one or both of those.

Q. Could you turn to Figure 4, please, or more precisely, Figure 4A?

A. Okay.

Q. Do you see that?
A. I do.
Q. At the -- well, strike that.
What's your understanding of Figure 4A?
MR. MECHELL: Objection to form. Calls for a legal conclusion.
A. It looks like a block diagram for a unified shader architecture.
Q. Do you see 62, block 62?
A. I do.
Q. Labeled "Unified Shader"?
A. Yes.
Q. So is that the unified shader?
MR. MECHELL: Objection to form.
MR. McNAMARA: Objection to form.
MR. MECHELL: Calls for a legal conclusion.
Q. So is this Figure 4A depicting the block diagram that includes the unified shader as well as blocks around surrounding the unified shader?
A. Yes.
MR. MECHELL: Objection to form. Calls for a legal conclusion.

MR. McNAMARA: Could you pause a minute before you answer so we can make sure the transcript is clear, please?

THE WITNESS: Okay.

MR. McNAMARA: Thank you.

Q. Do you see block 64 at the top left?

A. I do.

Q. Labeled "Arbiter"?

A. Yes.

Q. So is that the arbiter circuit?

MR. MECHELL: Same objection.

A. The arbiter circuit? What arbiter circuit?

Q. I think in connection with Claim 1, you were referring to the arbiter circuit?

A. I think that that's very plausible that it is.

Q. Is there anything else in this diagram that could refer to the arbiter circuit, in your opinion?

MR. MECHELL: Objection.

MR. McNAMARA: Objection to form. Calls for a legal conclusion. Counsel, are you confining his analysis to the

figure or the figure and the description?

A. Would you still like me to answer?

MR. SCHWENTKER: Well, I was referring specifically to the figure in this question.

MR. McNAMARA: So you don't want him to look at the spec. You want him just to look at the figure?

MR. SCHWENTKER: I didn't say that. Can you read back the question.

(Reporter read back pending question)

MR. LEVENTHAL: Form.

MR. McNAMARA: Same objection.

MR. MECHELL: Same objection.

A. No, I don't see anything else.

Q. Okay. Looking at the rest of the patent, do you see anything else that could refer to the arbiter circuit?

MR. McNAMARA: Objection to form. Calls for a legal conclusion.

A. That's hard for me to answer without, you know, a detailed study of the patent. The answer is do I see anything now?

The answer is no. Had I had more time to study the patent, perhaps.

Q. What about in the figures in particular?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. I do not.

Q. Could you turn to Figure 8 of the '871 patent -- sorry, Figure 5 of the '871 patent.

A. Yes.

Q. Do you have an understanding of what Figure 5 shows?

MR. McNAMARA: Objection.

MR. MECHELL: Objection. Calls for a legal conclusion.

A. I think Figure 5 shows the feeding of inputs into the unified ALU.

Q. When you say "unified ALU," what are you referring to?

A. I mean the unified shader that is the subject of this patent. Well, specifically the ALU portion of it. I suppose that the whole thing is a unified shader, but what this is showing is the sources going into the math

operants of the ALU.

Q. You said, "I suppose that the whole thing is a unified shader." When you said "the whole thing," what are you referring to?

MR. McNAMARA: Objection.

A. Everything in this diagram, except for -- yeah, everything in Figure 5. But what I was referring to, the inputs into the ALU, I was specifically referring to the box labeled "CPU."

Q. Okay. So you're referring to the CPU as an ALU?

A. Yes.

Q. Why are you referring to it as an ALU?

A. Because it's a purely math data path that's taking in three inputs and doing some mathematical operation on it. That's how I interpret these sources as inputs into a math unit.

Q. And you're referring to Source A, Source B and Source C?

A. Yes.

Q. And what makes you say it's a math

unit?

MR. McNAMARA: Objection.

A. Well, because it matches up with the data flow that I know that we implemented in R400. If you see, also there's a dotted box labeled "Scaler."

So I think what this is intended to show is that the math unit has both a vector portion in that solid box and a scalar portion in the dotted box. So it's really two parallel ALU units this is showing.

Q. And you referred to -- strike that.

The word "Scaler" there, is that -- it's spelled s-c-a-l-e-r?

A. Yes, I believe that's a typo in the patent. It should be a-r.

Q. Previously you said the box labeled "CPU," which you're referring to ALU, does some mathematical operation?

A. Yes.

Q. What are you referring to when you say "mathematical operation"?

A. I mean something like a multiply accumulate or a multiply add.

Q. Could you walk me through Figure 5?

MR. MECHELL: Objection to form. Calls for a legal conclusion. Overly broad.

A. I think this is a diagram showing the sources of data as each instruction is being issued. So I think the instruction store is meant to be driving both the selection of the constant that is possibly used or selecting the -- again, I haven't read the patent recently enough to know what box 65 is -- I'm sorry, box 92 is, but I believe it's intended to be the GPR store.

I think what this is showing is that the -- depending on the wave selected, these are the individual -- well, I don't know actually. I'd have to look at the patent to know whether the 1 to 63 represent individual waves or whether they represent individual locations within the GPR store that could be occupied by a single wave. It's very possible these are the possible GPRs associated with a single wave.

What it's showing is the inputs

into what's called the CPU here, but what I've been referring to as the ALU come from a mixture of the GPRs and the constants as directed by the instruction store.

Q. If you could turn to Column 4 of the patent --

A. Okay.

Q. -- of the '871 patent, starting at line either 8 or 9, it says, "Referring briefly to Figure 5, the unified shader 62 will now be described. As illustrated, the unified shader 62 includes a general purpose register block 92."

A. Okay.

Q. Is that the block 92 that you were just referring to?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

MR. McNAMARA: Objection.

A. Yes.

Q. Is that consistent with what you were just saying about Figure 5?

MR. McNAMARA: Object to form. Calls for a legal conclusion. The document speaks for itself.

MR. MECHELL: Same objection.
A. Yes.
Q. So --
A. But that doesn't tell me what 0 to 63 represent.
MR. LEVENTHAL: We have been going about an hour. When you reach a breaking point, and just before I forget, I want to put on the record that Mr. Gruber will review and sign.
MR. SCHWENTKER: Yeah, why don't we take a break now.
VIDEOGRAPHER: The time is 2:55 and we're off the record.
(Recess taken at 2:55 p.m. and reconvening at 3:11 p.m.)
VIDEOGRAPHER: We are back on the record. The time is 3:11.
A. Could I modify something I said earlier regarding the R400 top level specification block?
Q. Sure.
A. I think four shader pipes here constitute a single unified shader. I think I said 16.

Q. Which exhibit are you referring to?
A. I'm referring to the R400.
Q. What's the exhibit?
A. Exhibit 6.
Q. So the unified shader for those shader pipes, are they shader pipelines?
A. Shader pipes.
Q. Four of those shader pipes make up a unified shader?
A. Yes.
Q. Plus the sequencer?
A. Yes.
Q. Anything else?
A. No.
Q. Okay. And what did you base that on?
A. Thinking about back to the R400 as we implemented it, and we had multiple unified shaders each executing a different instruction simultaneous in them. I had forgotten about that. I thought maybe we only had one.
(Document marked as Exhibit 11 for identification)
Q. The court reporter has handed you

Exhibit 11.

A. Yes.

Q. It's a document titled "R400 Sequencer Specification SQ"?

A. Yes.

Q. In the bottom right-hand corner it says ATI 2028 LG versus ATI IPR2015-00325.

Do you see that?

A. Yes.

MR. SCHWENTKER: I understand from counsel for AMD that there's no objection to showing this document to the witness?

MR. MECHELL: No objection to showing the document.

MR. SCHWENTKER: Thank you.

Q. Mr. Gruber, do you recognize this document?

A. Yes.

Q. What is it?

A. It's the sequencer block specification.

Q. Did you review this document as part of your work at ATI?

A. Yes.

Q. At the top of the first page, do you see the --

A. When you say first page --

Q. Page 1 of Exhibit 11.

A. Okay.

Q. Do you see at the top left-hand corner it says "Originate Date"?

A. Yes.

Q. "24 September 2001"?

A. Yes.

Q. I'd like to take you to Page 5 of Exhibit 11.

A. Yes.

Q. Page 5 lists a number of revision changes. Do you see those?

A. Yes.

Q. What do these revision changes indicate?

A. Well, they obviously indicate new releases of the specification. But I believe that these were all maintained manually, that is when Laurent made a suitable enough change to the documents, that he thought it merited a bump in the revision number, he would bump the revision number and enter it here.

Q. At the bottom, the last entry in the list says, "Rev 2.0 (Laurent Lefebvre), Date: April 19, 2002"?

A. Right.

Q. Is it your understanding that that's the date of this document?

A. It probably is. As I say, he could have made minor changes and thought that that didn't merit a bump in the rev number. This isn't necessarily accurate, but I have no reason to doubt that it is accurate.

Q. If you go to the top of the page under "Originate Date," again it says "24 September 2001," of those -- that date is obviously earlier than the April 19, 2002 date at the bottom?

A. Right.

Q. Of those two dates, is there one that you think would be more accurate?

A. Well, I would trust the date at the bottom. The date at the top, when it says "Originate Date," that could be the date that the document was first created under a given name under this name, and the fact that he has revs prior to that may mean that he had it

under a different name and he copied it somewhere under a new name.

So I would believe the actual date that was put in by Laurent as being probably more accurate.

Q. Okay. And that Rev 2.0, if you turn to the first page, it says "Version 2.0"?

A. Yes.

Q. So does that match up?

A. Right.

Q. Could you turn to Page 7 of Exhibit 11. There's a "Figure 1: General Sequencer Overview."

Do you see that?

A. Yes.

Q. What is this diagram showing?

MR. MECHELL: Objection to the form.

A. It shows the sequencer and it interfaces to the shader data path.

Q. Which block is the sequencer?

A. The SQ block.

Q. On the right-hand side do you see a series of blocks labeled "SP"?

A. Yes.

Q. Four blocks labeled "SP"?
A. Yes.
Q. What do those represent?
A. Those are -- SP, I think we called them pipelines in the previous diagram. They're all part of one unified shader, but they're offset in time, and that's why, I think, they're showed as four separate units.
Q. So in connection with Exhibit 6, you --
A. Yes.
Q. -- said that four shader pipes and a sequencer constituted a unified shader?
A. Yes.
Q. So these, turning back to Exhibit 11 on Page 7, the four shader pipes there would all be part of the same unified shader?
A. Yes.
Q. So looking at this diagram, which blocks would constitute the unified shader?
MR. McNAMARA: Objection.
A. I would include -- I mean, I think the sequence, the SQ block and all of the -- you could view, you know, the entire block as a shader system. But if I had to pick out what I

would call a unified shader, I would include the thing labeled "CF Constants," the "Instruction Store," the "SQ," the "CStore," the four "SP" units and the "PC/OB" units.

Q. Anything else that you would include in the unified shader?

A. The "IJ" crossbar as well.

Q. Anything else?

A. I think it's a matter of opinion whether you want to include the texture pipe as part of the unified shader or just an auxiliary to it. I would probably keep it separate.

Q. What does the CF Constants block do?

A. It includes controls for control flow. So one of the instructions that we had would allow you to loop, and the loop instruction had some constants associated with it that would tell you the size of the loop, the increment of the loop, and the end condition of the loop, and that's what those CF constants are holding.

Q. What does the instruction store do?

A. The instruction store stores the

actual instructions that are being executed. They're the compiled result of the high level shader.

Q. Are the instructions passed from the instruction store to another block?

A. Yes. So they're passed to the SQ block. It doesn't actually -- yes, it does. If you see, there's a line on top of the instruction store that's labeled "INST" that goes from the instruction store to the SQ.

Q. What happens to the instruction after they're passed from the SQ to the sequence block?

A. Well, the SQ has to understand for what unit they're intended. As I mentioned, you know, they could be for ALUs, they could be for the texture unit, they could be control flow instructions.

So the SQ has to decode the instruction and use that to determine what unit it's going to drive the instruction to. Essentially the SQ has decided that, you know, let's say the SP is executing a given wave, it has to fetch the instruction for that wave, and then it has to send the instruction to the SPs.

-- I think it decodes it into a more understandable format. It could be kind of compressed in the instruction, but then it's expanded and it goes to the ALU.

Do you see there's a line that says "ALU Instruction"? So all instructions come from the instruction store to the SQ, and then the SQ sends the instruction to the various units, depending on what the instruction is intended to do.

Q. So after the sequencer sends the instruction to the shader pipe in the R400, what happens?

A. So the instruction is applied to the various execution units within the shader pipe. So, you know, part of this instruction bus goes to the GPRs to select the source inputs. You saw that previous patent that had Source A, Source B and Source C.

So the instruction bus controls where those sources come from, and what address of the GPR that we're reading.

In some cases you saw that it came from the Constant Store. That's this box labeled "CStore." That doesn't come from each

of the individual SPs because it's a constant. It's the same for all of them. So instead the SQ kind of reads it directly, and then kind of broadcasts it to all the SPs.

It also controls -- it has kind of an op here that says am I doing an add, am I doing a multiply, am I doing a multiply add, am I doing some logical function, shift or something.

So that's all contained in this ALU instruction, and it's all timed. The ALU instruction just kind of flows out of there and controls the data path, and it's the SQ's job to make sure that the right control is applied to the particular section of the data path at the correct time to implement the intended instruction.

Q. The instruction has, I think you said, an op code; is that right?

A. Yes.

Q. Does that specify the operation that's performed?

A. Yes; you know, multiply, shift type of thing, add. In this R400 -- no, in the shader processor spec there's a list of the op

codes I pointed out in the instructions.

Q. So those op codes would be specified -- strike that.

So a particular op code would be specified by an instruction?

A. Yes. It's part of the instruction.

Q. And that part of the instruction would determine what operation is performed?

A. Correct.

Q. I think earlier you talked about how the sequencer decodes instructions. When it does that, what -- I guess, does it tell or determine what the type of instruction is?

MR. McNAMARA: Objection.

MR. MECHELL: Objection. Vague.

A. The SQ needs to know it's an ALU instruction, for instance. It doesn't really care beyond that whether it's an add or a multiply, but it needs to know that this next instruction is for the ALU, and therefore, I have to -- I'm going to issue it to the ALU as opposed to it's a texture instruction or it's a control flow instruction.

Q. Okay. I think I understand. So

the sequencer, when it decodes the instruction, it determines whether it's an ALU instruction or a memory fetch instruction, for example --

A. Right.

Q. -- because it needs to know where, which block to send the instruction to?

A. Yes.

MR. McNAMARA: Objection.

Q. Earlier we also talked about, going back to the arbitration schemes we talked about earlier, we talked about both launch arbitration and execution arbitration.

A. Right.

Q. So what on this diagram in the R400 performed launch arbitration?

MR. McNAMARA: Objection.

A. Well, if you see these inter-blocks, do you see going into them are these IJs, which are the barycentric coordinates associated with pixels. Also going into them, which is vertex indices, those are the -- that's the work associated with vertices.

So I think the control is in the -- again, I am just going by this diagram,

because I really don't have recall of exactly how this worked. But it appears from this diagram that the SQ is telling the inter-blocks which of them to give priority to load into the SPs.

Q. So if the SQ block is telling the inter-blocks which of them to give priority to, does that mean the SQ block is performing the launch arbitration?

A. Again, I have no recollection from the details of this design back in 2000, but that would be my interpretation and my guess based on this.

Q. What about the execution arbitration?

A. That would definitely be done in the SQ, because the SQ is the one who has to issue the instructions. It had to pick a wave in order to know from which instruction point or two are used to go into the instruction store to fetch the instruction associated with that wave.

Q. Okay. You can set Exhibit 11 aside.

When we were looking at Exhibit 8,

the '871 patent, we were talking about an arbiter circuit. Do you recall ever implementing an arbiter circuit in the R400?

MR. McNAMARA: Objection.

A. I did not personally implement the arbiter. I recall discussing arbitration with the people who did implement it.

Q. Do you recall if an arbiter circuit was implemented in the R400?

MR. McNAMARA: Objection.

A. I'm sure an arbiter circuit was implemented, because you have to -- because a unified shader needs an arbiter circuit. You have to choose whether you're going to execute vertices or pixels, because you have the choice of both.

Q. Was that -- to your recollection, was that arbiter circuit within the sequencer?

MR. McNAMARA: Objection to form.

A. Again, I don't have a recollection -- I do know that the, what I called execution arbitration, was within the sequencer. I don't have a recollection of where the other arbiter was, the launch arbiter.

Q. Do you have any recollection of

the launch arbiter being located outside the sequencer?

MR. MECHELL: Objection. Lacks foundation.

A. I don't have a recollection one way or the other.

Q. But there's not a separate arbiter circuit shown in the block diagram we were just looking at?

MR. McNAMARA: Objection.

A. Correct. If I go by the block diagram, I would say the arbitration is within the sequencer.

Q. Do you remember when ATI first implemented an arbiter circuit in the R400?

A. No.

Q. I'd like you to turn back to Exhibit 8, please.

A. Okay.

Q. That's the '871 patent.

A. I have it operate here.

Q. Okay. If you could turn to Column 1 of the '871 patent, starting around Line 60. It says, "Conventional graphics processors require the use of both a vertex shader and a

pixel shader in order to generate an object."

A. I see.

Q. Do you know what the pixel shader there refers to?

MR. MECHELL: Objection. Calls for a legal conclusion.

A. A pixel shader, such as the one in the R300.

Q. So your understanding is that that's describing the -- strike that.

So the '871 patent was meant to be an improvement over the R300?

MR. McNAMARA: Objection. Mischaracterizes. Calls for a legal conclusion.

MR. MECHELL: Objection to form. Calls for a legal conclusion. Overly broad.

A. Yeah, I wouldn't say that it was specifically related to an improvement over the R300. This was relative to really all the graphics, the current state of the art graphics chips at the time.

Q. Including the R300?

A. Including the R300, yes.

Q. Okay. So what were the problems with the prior art graphics chips at the time?

MR. McNAMARA: Objection to form. Calls for a legal conclusion.

MR. MECHELL: Objection to form. Overly broad.

A. The specific issue that was being addressed by the unified shader was the fact that it was difficult to size the vertex and pixel shaders in a dynamic environment. You had to size both of them to be bigger than needed, because during a vertex-dominant case, the pixel shader would be idle, and during a pixel shader-dominant case the vertex shader would be idle.

It was rare that you had a balanced case where you could pick a reasonable size and both would be reasonably efficient.

Q. A little while ago you said that when it says -- when it refers to "pixel shader" there on Line 60 or 61, that that's a pixel shader such as the one in the R300; is that correct?

A. Yes.

Q. So that, I think, if I remember

correctly, the pixel shader in the R300 was able to perform color shading and texture shading?

A. And texture fetches I would characterize it as. Shading involving texture fetches.

Q. So the pixel shader referred to here on Line 61 is talking about a pixel shader that was able to perform color shading and texture fetching?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. I'm not sure it was limited to that. Even if you had a shader that, you know, the texture fetching was outside of the shader, I still think a unified shader makes a certain amount of sense, though perhaps not as much as one that included texture fetching as well.

Q. Could you turn to Column 4, please. Starting around Line 36 it says, "The instruction store 98 contains the necessary instructions that are executed by the processor 96 in order to perform the respective arithmetic and logic operations on the data maintained in the general purpose register

block 92 as provided by the source registers 93 to 95."

Do you see that?

A. Yes.

Q. What's the distinction between arithmetic and logic operations?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. Arithmetic generally refers to multiply or adds, and logical operations refer to things like ands, or ors, or shifts, things that, you know, you don't really care what the meaning of the -- where you're not treating the piece of data as a number.

Q. Could you turn to Column 5 of the '871 patent. Starting at line 14 or so, it says, "The unified shader 62 has 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 98, very quickly."

Do you see that?

A. Yes.

Q. What's your understanding of "manipulation operations"?

MR. McNAMARA: Objection.

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. My interpretation is that it's able to execute their respective shader programs. So each individual instruction in a shader program is a manipulation.

Q. Then further down in that paragraph, at line 27 or so, it refers to "pixel calculation operations." Do you see that?

A. Yes.

Q. What's your understanding of pixel calculation operations?

MR. MECHELL: Same objection.

A. Again, it's executing one or more pixel shader instructions.

Q. Is there any distinction in your mind between pixel manipulation operations and pixel calculation operations?

MR. MECHELL: Objection to form.

A. No.

Q. Can a single instruction cause a

manipulation operation?

A. Sure.

Q. And can a single instruction cause a calculation operation?

A. Yes.

Q. Going back to the top of that paragraph, that first sentence that I previously read, that discusses simultaneously performing vertex manipulation operations and pixel manipulation operations?

A. Right.

Q. What's your understanding of that?

MR. MECHELL: Objection to form.

MR. McNAMARA: Objection to form.

Calls for a legal conclusion. Vague.

A. My understanding is that the shader for a given pixel and a given vertex can be -- can be active simultaneously, meaning that you do not run a pixel to completion, and then run a vertex to completion or a pixel wave to completion and then a vertex wave to completion. But instead, you're able to execute both, I'll say it, at the same time in the same sense that this paragraph uses at the same time.

That doesn't mean that each instruction is being executed at the same time. You launch a pixel instruction and then you follow it by a vertex instruction, but the pixel shader, which is still in the middle of operating, even when no instructions are executing, if it hasn't finished yet, if it has further instructions to execute, and the vertex shader is still executing at the same time.

Q. Is this like the interleaving that we were discussing earlier today?

A. Yes, exactly.

MR. MECHELL: Objection to form. Calls for a legal conclusion. Overly broad.

Q. With an "L"?

A. Yes. I believe that's what's referred to, when it says "freely switched," between some programs instructions.

MR. SCHWENTKER: How long have we been going?

VIDEOGRAPHER: Five more minutes left for the disk.

MR. SCHWENTKER: Why don't we break, then.

VIDEOGRAPHER: The time is 3:58 and we're off the record.

(Recess taken at 3:58 p.m. and reconvening at 4:15 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 4:15.

BY MR. SCHWENTKER:

Q. Welcome back, Mr. Gruber. Before the break, we were talking about the '871 patent, Exhibit 8. I'd like for you to turn to Column 8. At the top of Column 8, do you see Claim 15?

A. Yes.

Q. Are you familiar with the concept of patent claims?

A. Yes.

Q. What's your understanding of patent claims?

MR. McNAMARA: Object to form. Calls for a legal conclusion.

A. Well, there are independent and dependent patent claims. This appears to be an independent patent claim.

Q. Is it your understanding that a patent claim kind of defines the metes and

bounds of an invention?

MR. MECHELL: Objection to form. Calls for a legal conclusion.

A. Yes.

Q. So Claim 15 says "A unified shader," and then -- let me stop there.

Do you know when you and your co-inventors first came up with the idea of a unified shader?

MR. McNAMARA: Objection.

MR. MECHELL: Objection to form.

A. I don't have any clear recollection.

Q. Do you recall a general time frame?

MR. MECHELL: Same objection.

MR. McNAMARA: Objection.

A. It was around 2000, but nothing better than that.

Q. It could have been later than that?

MR. McNAMARA: Objection.

MR. MECHELL: Objection.

A. I'm not sure. I suppose it's possible.

Q. Do you recall when ATI first implemented a unified shader?

MR. MECHELL: Objection to form. Vague.

A. I don't have a clear idea.

Q. The next line says "A general purpose register block for maintaining data."

Do you recall when you and your co-inventors first had the idea of a unified shader with a general purpose register block for maintaining data?

MR. McNAMARA: Objection to form.

MR. MECHELL: Objection to form.

A. No, but the general purpose register block would have been part of that from the beginning. It's not like that would be something added on.

Q. Do you recall when a unified shader with a general purpose register block for maintaining data was first implemented?

MR. MECHELL: Objection to form. Vague.

A. No.

Q. Below that it says "A processor unit." Do you see that?

A. Yes.
Q. Do you recall when you and your co-inventors first had the idea for a unified shader with a processor unit?
MR. MECHELL: Objection to form. Calls for a legal conclusion.
MR. McNAMARA: Objection.
A. No. But again, a processor unit would have been part of the whole unified shader from the beginning as opposed to a sequential thing.
Q. Do you recall when a unified shader with a processor unit was first implemented?
MR. MECHELL: Objection to form. Vague.
A. No.
Q. Then finally below that it says, "A sequencer, coupled to the general purpose register block and the processor unit, the sequencer maintaining instructions operative to cause the processor unit to execute vertex calculation and pixel calculation operations on selected data maintained in the general purpose register block."

Do you see that?

A. Yes.

Q. Do you recall when you and your co-inventors first had the idea for a unified shader with that sequencer?

MR. McNAMARA: Objection.

A. No.

MR. MECHELL: Objection to form. Calls for a legal conclusion. Overly broad.

Q. Do you recall when a unified shader with that sequencer was first implemented?

MR. MECHELL: Objection to form. Vague.

A. No.

Q. Further down in Column 8 there's Claim 20. Do you see that?

A. Yes.

Q. And it says, "The shader of Claim 15, wherein the processor unit executes vertex calculations while the pixel calculations are still in progress."

Do you see that?

A. Yes.

Q. Do you understand this is a dependent claim from Claim 15?

A. Yes.

Q. What's your understanding of when it says "wherein the processor unit executes vertex calculations while the pixel calculations are still in progress"?

MR. MECHELL: Objection to the form. Calls for a legal conclusion.

A. The pixel shader for a given pixel has not finished; that is, there are still instructions to be issued for that pixel shader. Yet we have issued instructions for a vertex wave as well on the same unified shader.

Q. Is that like the interleaving we have been discussing today?

MR. MECHELL: Objection to form. Calls for a legal conclusion. Vague.

A. Yeah, I think it's -- I think it is -- interleaving is one aspect of it.

I will also point out that given that there are multiple execution engines, like we mentioned the ALU, the texture, the control flow, all of which are capable of executing an instruction in a given cycle.

So, you know, I think that this covers both interleaving on a given unit as well as simultaneous execution across two different units.

Well, maybe this claim doesn't, because it specifically talks about the processor unit.

Q. And by "processor unit," you understand that to mean what?

MR. MECHELL: Objection. Calls for a legal conclusion.

A. I'm not sure actually whether processor unit is restricted to the ALU or could be interpreted to mean to include the texture unit or the control flow unit. I'd have to look at a lot more detail in this patent to really give an opinion.

Q. Could you turn to Exhibit 9, please.

A. Yes.

Q. Actually, before we do that, going back to Exhibit 8, the '871 patent.

A. Yes.

Q. The first page, page AMDLG0002112.

A. Yes.

Q. Do you see you're listed as one of the inventors?
A. Yes.
Q. So this is your patent?
A. It is.
Q. Were you involved in drafting this patent?
A. Yes.
Q. Do you recall what you drafted?
MR. MECHELL: Objection to form.
A. I have looked at the patent since issue, but I haven't looked at this patent recently. So I don't have any immediate recall.
Q. When was the last time you looked at it?
A. It was associated with the ITC court case that must have been two years ago.
Q. Okay. So you looked at it two years ago?
A. Yes.
Q. Have you looked at it since then?
A. No.
Q. Okay. Now turning to Exhibit 9.
A. Okay. Let me say that I

previously said that I recognized this patent based on the title. But looking at the date of issue and the filing date, I'm no longer certain that I've seen this patent.

Q. Which one?

A. Exhibit 9.

Q. Oh, Exhibit 9. Okay.

A. You said turning to Exhibit 9. That's why I brought it up.

Q. Is that because the filing date and the issue date were after you left AMD?

A. Correct.

Q. Were you involved in the filing of the patent application for this patent?

A. Not to my knowledge.

MR. MECHELL: Objection. Vague.

Q. But you would agree that you're listed as one of the inventors of the patent?

A. Yes.

Q. And is it your understanding that this patent is related to the '871 patent?

A. Yes.

Q. And is it your understanding that the figures in the text of the patent are the same as the '871 patent?

MR. McNAMARA: Objection.
MR. MECHELL: Objection. Lack of foundation.
A. I haven't examined the patent close enough to really say one way or the other.
MR. LEVENTHAL: Could you just represent that to him, if it's necessary?
Q. I will represent to you that the figures in the text are the same as the '871 patent, with perhaps minor differences where it claims priority to related applications.
A. Okay.
MR. McNAMARA: Objection. Lacks foundation. Mischaracterizes the document.
Q. And, of course, the patent numbers and the filing date and application date or the issue date are also different.
MR. McNAMARA: Same objections.
Q. Could you turn to the last page of Exhibit 9, please.
A. Yes.
Q. In Column 8 there's a Claim 11.

Do you see that?

A. Yes.

Q. I'm going to read the first part of Claim 11. It says, "A unified shader comprising: A processer unit flexibly controlled to perform vertex manipulation operations and pixel manipulation operations based on vertex or pixel workload."

Do you see that?

A. Yes.

Q. Do you have an understanding of vertex or pixel workload?

MR. McNAMARA: Object to form. Calls for a legal conclusion.

A. Yes.

Q. What does that mean to you?

MR. McNAMARA: Same objections.

A. I think it means the number of vertices or the number of pixels submitted to the processor. For instance, you may, you know, if you're dealing with large triangles, you have very few vertices and you have a lot of pixels, and if you're dealing with very small triangles, you have a lot of vertices and very few pixels.

So the ratio of vertex-to-pixel workload may change based on that ratio, and that ratio is simply a function of if it's a game, you know, how far you are from interesting objects.

Q. Is workload a commonly used term in this context?

MR. McNAMARA: Object to form. Calls for a legal conclusion.

A. Yes. I would say yes.

Q. The understanding that you just provided, is that based on your experience in the industry?

MR. McNAMARA: Objection to form. Calls for a legal conclusion. Lacks foundation.

A. Yes.

Q. Below that, Claim 11 goes on to say "an instruction store and wherein the processor unit of the unified shader performs the vertex manipulation operations and pixel manipulation operations at various degrees of completion based on switching between instructions in the instruction store."

What's your understanding of what

it means when it says "based on switching between instructions in the instruction store"?

MR. McNAMARA: Objection to form. Calls for a legal conclusion. Lacks foundation. Vague.

A. My understanding is that is that you have waves in various stages of completion of their respective shader.

For instance, if a shader has 10 instructions for pixel and 20 instructions for vertex, you may be in a situation where the pixel shader has completed five instructions and the vertex shader has -- the vertex wave has completed seven instructions of the vertex shader, and that's because before completing the pixel shader, you switch to execution of the vertex shader; of the vertex shader wave, pixel shader wave and vertex shader wave.

Q. Could you turn to Exhibit 10, please. This is the '846 patent. Do you see that?

A. Yes. I'm going to make the same comment here. I mistakenly said I recognize this patent based on the title and the subject, but based on the filing dates and the issue

dates, I have not seen this patent.

Q. And that's because it was filed and issued after you left AMD?

A. Correct.

Q. Turning to the first page of the patent, you would agree, though, that you are listed as a named inventor on the patent?

A. Yes.

Q. You don't have any reason to doubt that this is your patent?

A. No.

Q. So of these three patents, it sounds like the '871 patent is the only one that you were personally involved in seeking patent protection?

MR. MECHELL: Objection to form.

A. Correct.

MR. MECHELL: Vague.

Q. Did you receive any compensation for the '871 patent?

A. I believe that I did. I believe that ATI had a patent compensation program.

Q. Do you recall what you received?

A. No. It would have been somewhere between $500 and $1,000, but I don't remember

what. Oh, and a balloon.

Q. What kind of balloon?

A. It was one of these ones, you know, the polyethylene ones, so they lasted for a while. I remember that my kid loved those when I would bring them home.

Q. Did you receive any compensation in connection with the '454 patent or the '846 patent?

A. No. Not even a balloon.

Q. Are you being compensated for your time here today?

A. No -- well, when I was served the patent, I got a $50 check associated -- when I was served the subpoena. So I suppose I am.

Q. But other than that, you haven't received any compensation for appearing here today?

A. Right.

Q. Do you have any plans to attend trial in the Northern District of California case?

A. No.

Q. Do you have any plans to attend a trial or hearing in the ITC investigation?

A. No.
Q. Has anyone asked you to attend?
A. No.
Q. If you were asked to attend, would you?
A. Without a subpoena, do you mean?
Q. Correct.
MR. MECHELL: Objection. Incomplete hypothetical.
A. No.
Q. Let me restate that. If you were asked to attend trial in the Northern District of California case, would you, without a subpoena?
A. No.
MR. MECHELL: Again, incomplete hypothetical objection.
Q. If you were asked to attend trial in the ITC investigation without a subpoena, would you?
A. No.
Q. Now, earlier today you said that you testified at least, I believe you said, two or three depositions --
A. Right.

Q. -- in connection with a different ITC investigation?

A. Correct.

Q. And I think you said that testimony was about two years ago?

A. Yes.

Q. Okay.

MR. SCHWENTKER: Counsel, this is -- Bryan, I'm handing you a copy of a transcript that was produced in the Northern District of California case. It says on the front that it's the videotaped deposition of Andrew E. Gruber.

I request your confirmation I can show this to Mr. Gruber?

MR. MECHELL: I can only speak to the extent it contains ATI Technologies ULC confidential business information, in that matter, in which case there would be no objection for showing it to the witness, because it is the witness's testimony from that matter.

To the extent there's any other third-party or Respondent's confidential

business information, I can't speak to that. Just to be clear, I don't know whether there's any in there.

MR. McNAMARA: What deposition is that, the date?

MR. SCHWENTKER: This is a March 17, 2015 deposition.

MR. McNAMARA: No objection.

MR. SCHWENTKER: No objection from you?

MR. McNAMARA: No.

MR. LEVENTHAL: On behalf of Qualcomm, who is the Respondent in this case, no objection either.

MR. SCHWENTKER: Okay. So where does that leave us?

MR. LEVENTHAL: Can we go off the record for a second?

MR. SCHWENTKER: Sure. Why don't we go off the record.

VIDEOGRAPHER: The time is 4:41 and we are off the record.

(Discussion off the record at 4:41 p.m. and reconvening at 4:45 p.m.)

VIDEOGRAPHER: We are back on the

record. The time is 4:45.

BY MR. SCHWENTKER:

Q. Mr. Gruber, do you recall being deposed in the ITC investigation on or around March 17, 2015?

A. Yes, I do.

Q. Okay. Do you recall you were designated as a corporate witness on behalf of Advanced Micro Devices and -- yeah, I believe just Advanced Micro Devices?

A. Yes.

Q. Have you reviewed your transcript from that deposition?

A. Yes.

Q. When did you review it?

A. It was soon after the transcript was available. It was probably about a month after the testimony.

Q. Did you find any errors in your testimony from that deposition?

A. As I recall, there were one or two errors or things that I wasn't sure whether they needed to be corrected or not, and I informed my lawyers of that, when I returned the signed copy with whatever markup I did, I

think they might have said those are not material and don't have to be changed. I don't recall.

Q. So as far as you're aware, there were no material errors in your transcript?

A. As far as I'm aware, yeah.

Q. Did you fill out an errata sheet where you specified changes to be made?

A. Yeah -- I can't be 100 percent sure I filled out an errata sheet as opposed to an e-mail communication.

Q. Okay.

A. It was definitely written, but I just don't remember whether it was actually on paper.

Q. Okay. Other than those non-material errors that you noted in your transcript, do you stand by your testimony from that deposition?

A. Yes.

Q. Why did you agree to appear as AMD's corporate witness in that deposition?

A. Because it was in the interest of my current employer.

Q. Your current employer being

Qualcomm?

A. Yes.

Q. Why was it in Qualcomm's interest?

A. Qualcomm was a party to the lawsuit. It was one of the Defendants.

Q. And how was your testimony in Qualcomm's interests?

MR. LEVENTHAL: Let me caution you, you can answer, but as you do so, do not reveal any attorney-client communication.

A. I'm sorry, could you repeat the question?

Q. How was your testimony in Qualcomm's interests?

A. My understanding was that one of the defenses to the NVIDIA infringement claim was that their patents were invalid because of prior art demonstrated by the AMD patents.

Q. Is it your understanding that the AMD patents are prior art to the NVIDIA infringement claims?

MR. LEVENTHAL: Form.

A. That is my current understanding.

Q. What do you base that

understanding on?

MR. McNAMARA: Objection.

MR. LEVENTHAL: Again, you can answer that to the extent you have a basis, other than attorney-client communications. Otherwise, I instruct you not to answer.

A. Okay. Simply, my recollection of how early we were in the unified shader development. I mean, I know when NVIDIA came out with theirs, and I know that we had done work significantly prior to that.

Q. When did NVIDIA come out with theirs?

A. I don't recall now, but I remember at the time when we were looking at the ITC case and the validity of the NVIDIA patents that the time frame was such that I was pretty certain that we were prior to that.

Q. So you said that your testimony in that prior deposition in 2015 was in the interests of Qualcomm, your current employer?

A. Correct.

Q. Do you know whether your testimony in this deposition is in the interests of

Qualcomm?

MR. McNAMARA: Objection.

A. I do not know.

Q. You have no reason to think it is or is not?

MR. McNAMARA: Objection.

A. I would say I have reasons to think on both sides.

Q. What reasons do you think it might be in Qualcomm's interests?

MR. MECHELL: Objection. Calls for speculation.

MR. LEVENTHAL: Form.

A. Well, because Qualcomm has a license to the IP from AMD. This could hurt Qualcomm's competitors and not hurt Qualcomm. So I suppose in that sense, it could help.

Q. Any other reasons you think it would help?

MR. MECHELL: Same objection.

A. No.

Q. Any reasons you think it would hurt Qualcomm's interests?

MR. MECHELL: Same objection.

A. Yeah, because, you know, LG is a

Qualcomm customer, and so we don't want to see our customers hurt. AMD is not.

Q. Going back to your work on the unified shader in the R400, do you know whether others copied the unified shader idea that you, Mr. Morein, Mr. Lefebvre and Mr. Skende had?

A. Well, I know the unified shader has become very popular, almost standard. I don't think you can find a GPU out there that doesn't use unified shaders. It doesn't mean that they copied as opposed to independently invent it.

Q. Do you have any opinions on whether -- strike that.

So you're not aware of anyone copying the unified shader idea that you all had?

A. No.

Q. From your prior response, I understand or I take it that you understand that graphics processors in the market today have unified shaders?

A. I think that's almost universally true, yes.

Q. Do you have any opinions on

whether any unified shaders on the market today infringe any of your patents?

MR. MECHELL: Objection. Calls for a legal conclusion.

A. I don't really have an opinion on it. I can say I have suspicions because of the breadth of our patent. But it could be there are ways around it.

I certainly haven't examined in detail the implementation of others to really draw an informed opinion.

Q. Okay. Do you know if others came up with the idea of a unified shader around the same time as you, Mr. Morein, Mr. Lefebvre and Mr. Skende?

MR. MECHELL: Objection. Calls for a legal conclusion.

A. I don't know, other than, you know, by looking at when NVIDIA introduced their product and estimating a reasonable development schedule.

Q. Other than the patent compensation of $500 to $1,000 that you mentioned earlier and the balloon, have you received any awards or praise for your work that led to the unified

shader patents?

A. At AMD I got this thing, like Plexiglas, it's like a trophy type of thing. I think that was specifically associated with the unified shader patent.

Q. Okay. Do you know why you received that?

A. I just think because they thought it was an important patent, so it was additional recognition.

Q. Any other awards?

A. Not that I can think of.

Q. Do you own any AMD stock?

A. No. I wish I did. It has gone up a lot lately.

MR. SCHWENTKER: I have no further questions. I did want to note for the record that I don't believe we have a copy of the errata, if there was an errata, from the 2015 deposition transcript.

With that, I will pass the witness.

VIDEOGRAPHER: The time is 4:59 and we're off the record.

(Recess taken at 4:59 p.m. and reconvening at 5:03 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 5:03.

EXAMINATION BY

MR. MECHELL:

Q. Good afternoon, Mr. Gruber.

A. Good afternoon.

Q. Earlier counsel for LG asked you about whether your testimony and Qualcomm's interests are related in any way. Do you remember that?

A. Yes.

Q. Does the question of whether your testimony is in Qualcomm's interests factor in any way into the content of your testimony today?

A. No, it does not.

Q. And prior to this deposition, have you had any communications with counsel for LG?

A. No.

Q. Have you had any communications with LG employees regarding this case?

A. No.

Q. And you also may recall counsel

asking a number of questions about three patents today; the '871 patent, '454 patent and the '846 patent. Do you remember that?

A. Yes.

Q. Are you an attorney?

A. No.

Q. Are you a registered patent attorney?

A. No.

Q. I believe you testified that you had not read the '871 patent in two years; is that about right?

A. Correct.

Q. Did you study the '871 patent two years ago?

A. I would say that I looked it over. I would not claim a deep study of it.

Q. Did you read the entire specification in the '871 patent at that time?

A. Yes, I read the entire patent.

Q. Did you review the R400 documentation that we discussed today and do a claim element-by-element comparison against the claims in the '871 patent?

A. No, I did not.

Q. Did you review any of the R400 RTL code and do a claim-by-claim element comparison for the '871 patent?

A. No.

Q. Would your answers be the same for the '454 and '846 patents as well?

A. Yes, they would, except the thing about looking them over. I did not look them over two years ago.

Q. And you had not looked at the '454 or the '846 patents prior to today?

A. Correct.

Q. And today, how would you characterize your review of those two patents?

A. Cursory at best.

MR. MECHELL: Thank you. No further questions.

VIDEOGRAPHER: The time is 5:07 and we're off the record.

(Off the record at 5:07 p.m. and reconvening at 5:08 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 5:08.

EXAMINATION BY

MR. McNAMARA:

Q. Good afternoon, Mr. Gruber. My name is Mike McNamara. You're here in response to a subpoena from AMD and ATI as well, correct?

A. Correct.

Q. Now, if you could turn your attention to the architecture proposal that you testified about earlier today.

A. Yes.

Q. It's Exhibit 4. Just let me know when you're there.

A. I'm there.

Q. Now, you had a hand in drafting this architecture proposal, correct?

A. I gave some feedback on it, I would say.

Q. And when this architecture proposal was drafted, you were working on the R400 project, right?

A. Correct.

Q. Now, you were involved with some of the initial meetings to convince people within your own company that going with a unified shader was the right thing to do for the R400 project, right?

A. Correct.
Q. Now, when you testified earlier today, you said their main issues were about performance and precision of a GPU with a unified shader, right?
MR. SCHWENTKER: Objection.
A. Correct.
Q. Did the GPU with the unified shader ultimately deliver with sufficient precision?
MR. SCHWENTKER: Objection. Form.
A. Yes. I mean, there may have been corner cases that showed precision issues, but they weren't significant enough to mean that there was any kind of real problem with the product.
Q. Did the GPU with the unified shader deliver on the performance needed?
MR. SCHWENTKER: Objection. Form.
A. Eventually in the produced products, yes.
Q. Is that one of the reasons why the rest of the industry has adopted your unified shader approach to GPUs?
MR. SCHWENTKER: Objection to

form. Calls for speculation.
A. Yes.
Q. When this architecture proposal, which is Gruber Exhibit 4, was drafted, were you working on the R400 at that time?
A. I was in the early stages of it, yes.
Q. Were other people working on the R400 project as well?
A. Yes.
Q. And you continued working on that R400 project up until that R400 team became the R500 team; is that correct?
A. That's correct.
Q. And that development was consistent from when you started out the R400 through when it transitioned to the R500 project, correct?
MR. SCHWENTKER: Objection. Vague.
A. It was a continual effort.
Q. And that R400/R500 team included 100 engineers? How many engineers were involved in the R400 project?
A. Probably close to 100, I would

say.
Q. And how many engineers were involved in the R500 project?
A. The same.
Q. And that team worked every workday to continue to develop that GPU, right?
MR. SCHWENTKER: Objection. Vague. Calls for speculation.
A. Yes.
Q. Now, this architecture proposal bears a date of November 13, 2000, correct?
A. Correct.
Q. And if you take a look at Page 19 of the document, the unified shader is described at a high level on that page, correct?
A. 19? There's only 16 pages.
Q. I apologize. Page 9. I merged 9 of 16 together to get 19.
A. Yes. It describes the unified shader.
Q. So the unified shader was conceived of by the time this document was drafted, correct?
A. Correct.

MR. SCHWENTKER: Objection. Calls for a legal conclusion.

Q. Take a look at Page 15 of the document, there's also a reference to a sequencer, Section 9.3.3. Do you see that?

A. Yes.

Q. So by this time, you had understood that a sequencer was necessary, right?

A. Correct.

MR. SCHWENTKER: Objection. Vague.

Q. Now, when you go through and work on the design effort at AMD, you write down the description of the R400 as it evolves, right?

A. Correct.

Q. So these design documents are a record of what the R400 project consisted of, right?

MR. SCHWENTKER: Objection. Vague.

A. Yes.

Q. And so, if I went to a particular document that was written in 2001 about the R400 project and read it, that would be a good

approximation of what the R400 project consisted of at that time?

A. At that time, correct.

Q. And an even better record of what the R400 project consisted of would be the RTL, that's the actual embodiment of the GPU, right?

A. That's correct.

Q. Were you involved at all in the testing of the RTL?

A. I was aware of what stage we were in testing and how far along we were in terms of pass/fails. We would generally keep track in terms of number of tests written, number of tests passed on the C level model, and then number of tests passed on the RTL.

Q. And throughout the process of this project, you would synthesize the RTL into net lists, right?

A. Yes.

Q. And those net lists were tested on a FPGA prototype, right?

A. I believe that that's true, that we had one of those.

Q. So the design process would be to develop the RTL first, correct?

A. Yes.
Q. Then you would test it using kind of cadence type environment to test the RTL?
A. Yes.
Q. Then synthesize it into net lists, right? You've got to give me a yes or no.
A. Well, I mean, I think to synthesize into net lists, that would have to wait until we had enough of the chip implemented in RTL to have a full chip simulation.
But we would do a lot of work just purely in the RTL doing testing of individual RTL blocks just using the RTL and the C level model as I described earlier.
So I'm saying that the path to the FPGA implementation netlist prototype happened later in the process, after there was already a lot of initial testing done at the block level on all the blocks.
Q. And so when you got to the point where you generated a netlist, that would be tested on those IKOS FPGAs, right?
A. Right.
Q. Just closing one thing out, was a

unified shader used in the AMD products prior to the R400?

A. No.

Q. Was a unified shader used by others, to the best of your knowledge, prior to the R400 project?

MR. SCHWENTKER: Objection. Form.

A. No.

Q. To the best of your understanding, was the unified shader a new concept for use in GPUs as of the R400 development?

A. Yes.

Q. Was a sequencer used in AMD products prior to the R400?

A. I think so, yes.

Q. Was a sequencer used in GPUs prior to R400?

MR. SCHWENTKER: Objection.

A. I believe the R300 had a sequence.

Q. But a sequencer wasn't used in connection with a unified shader prior to the R400, right?

A. Correct.

MR. SCHWENTKER: Objection. Calls for speculation.

Q. Now, during your earlier testimony, you talked a little bit about NVIDIA. Do you recall that?

A. Yes.

Q. And you were saying that NVIDIA didn't develop its unified shader until after AMD developed its unified shader, correct?

A. That is my belief. But I cannot say for certain that that's true.

Q. And as an engineer, you keep up with what's going on in the industry and read industry articles, right?

A. Right.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 12 a document bearing the Bates numbers AMD1044_0256586, please. It goes through 588.

(Document marked as Exhibit 12 for identification)

Q. Sir, do you see that Exhibit 12 bears a title that reads "ATI and NVIDIA Proclaim Different Graphics Processors Architecture Goals"?

A. Yes.

Q. Then it goes on and says, "ATI Says Unified Rendering Engine - the Way to Go, NVIDIA Disagrees."

Do you see that?

A. Yes.

Q. Is it your recollection that NVIDIA, as of 2004, was going in a different direction than using a unified shader?

MR. SCHWENTKER: Objection. Form.

A. Yes.

Q. And that article is dated December 23, 2004, correct?

A. Correct.

MR. SCHWENTKER: Objection. Calls for speculation.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 13 a document bearing Bates numbers AMD 1044_0256589 through 6590, please.

(Document marked as Exhibit 13 for identification)

Q. This document has a title that says, "NVIDIA Chief Architect: Unified Pixel and Vertex Pipelines - The Way to Go. NVIDIA

says it would make a chip with unified pipes 'When it Makes Sense'."

Do you see that?

A. Yes.

Q. Is it your understanding that NVIDIA was not yet making a unified shader in 2005?

MR. SCHWENTKER: Objection. Calls for speculation and a legal conclusion.

A. Yeah, based on this article, yes.

Q. Based on your understanding of what was going on in the industry in 2005, is it your understanding that NVIDIA didn't have a product with a unified shader on the market?

MR. SCHWENTKER: Same objections.

A. I remember that we beat them to a market with a unified shader, but I don't recall specifically in 2005 whether NVIDIA had a unified shader or not.

Q. And when did AMD first come out with a product with a unified shader?

A. I know that the product was the R600, but I don't recall when that was.

Q. And that beat NVIDIA to the market with a GPU with a unified shader, correct?

MR. SCHWENTKER: Objection. Calls for speculation and a legal conclusion.

A. Yes.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 14 a document bearing the Bates numbers AMD1044_0256582 through 585, please.

(Document marked as Exhibit 14 for identification)

Q. Sir, what is this document?

MR. SCHWENTKER: Objection. Calls for speculation.

A. This is an announcement that ATI and Microsoft -- that ATI is going to supply the graphics for the next Microsoft Xbox.

Q. Is it your understanding that the Xenos project within AMD started before August 14, 2003?

A. Yes.

Q. That Xenos product -- project was an extension of the R400 project, right?

MR. SCHWENTKER: Objection. Form.

A. Yes.

Q. The development process extended

seamlessly from the R400 through into the Xenos process, right?

MR. SCHWENTKER: Objection. Form.

A. Yes.

Q. And the design process also extended seamlessly from the R400 through to the R500, correct?

MR. SCHWENTKER: Same objection.

A. Yes. Both of the teams continued to work on their products.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 15 a document bearing Bates numbers AMD1044_0016660 through 6679, please.

(Document marked as Exhibit 15 for identification)

Q. Sir, do you recognize this document?

A. Yes. This is another version of the sequencer spec.

Q. Now, this version of the sequencer spec has revision changes to it, correct?

A. Correct. Let me just make sure --

Q. The third page.

A. Yes.
Q. And the most recent revision to this sequencer specification was in August 24, 2001, right?
A. Yes.
Q. So this document would be current as of August 24, 2001, right?
A. Correct.
Q. And it continued to evolve through the R400 design cycle, right?
A. Yes.
MR. SCHWENTKER: Objection to form.
Q. But as of August 24, 2001, this is the state that the sequencer specification existed in, right?
MR. SCHWENTKER: Objection. Calls for speculation.
A. Yes.
Q. Now, AMD keeps track of changes to its design specifications of a Perforce system, right?
MR. SCHWENTKER: Objection to form.
A. Yes.

Q. It also keeps -- let me rephrase that. When you worked at AMD, AMD kept track of changes to its design specifications, right?

A. I don't recall how much of this used the Perforce database or not. We certainly used it for the RTL. I don't recall actually whether we used it for the specifications or not.

Q. So outside of just Perforce, when you worked at AMD, AMD kept track of changes to its technical documents, right?

A. It did, which is why we had version numbers and we had Word keep track of changes.

Q. And for some documents, the Perforce system was used, right?

A. Yes.

Q. Were there any other document revision governing software used at AMD, to your recollection?

A. I don't recall.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 16 a document bearing Bates numbers AMD1044_0253421 through

53436, please.

(Document marked as Exhibit 16 for identification)

Q. Sir, do you recognize this document?

A. I don't really recognize it, no. I don't recall it.

Q. Do you recall whether you would have been part of an executive review of the R400 project?

A. I think that I would have.

Q. Could you take a look at page ending in Bates numbers 53427, that talks about the general status?

A. Yes.

Q. Now, this general status, the first bullet point talks about hardware emulation?

A. Yes.

Q. And that's the IKOS FPGA we talked about a little bit earlier?

MR. SCHWENTKER: Objection. Calls for speculation.

A. Correct.

Q. Here it reports that the netlist

delivery was completed on October 4, 2002. Do you see that?

A. Yes.

Q. Is that consistent with your recollection?

A. It is consistent, but my recollection is not that good of the actual dates.

Q. Do you have any reason to doubt that the netlist was actually delivered on October 4, 2002?

A. No.

Q. Do you know what the PD team was in connection with the R400?

A. "PD" refers to physical design. So it's the actual layout of the gates into silicon.

Q. And were you part of the PD team?

A. No. That was a separate team.

Q. Would you have attended PD team meetings? Not internal PD meetings. I was probably present when a PD representative reported in a wider meeting.

MR. McNAMARA: I would like to have marked for identification purposes

as Gruber Exhibit 17 a document bearing the Bates numbered AMD1044_0169851 through 9863, please.

(Document marked as Exhibit 17 for identification)

Q. Sir, do you recognize this document?

A. I don't recognize it, no.

Q. Do you know whether you would have been part of this PD team meeting?

A. I can't say for sure whether I would have or I wouldn't. I mean, I would imagine I probably wouldn't have been actually.

Q. If I could direct your attention to Page 11 of the PowerPoint presentation. Just let me know when you're there.

A. I'm there.

Q. Page 11 of the PowerPoint presentation lists out three different releases of a netlist?

A. Right.

Q. Was it commonplace to have three different releases, so here it's 3.0B, 3.0C, 3.0D of a netlist?

MR. SCHWENTKER: Objection to

form.

A. Yeah. I mean, you would do it based either on PD feedbacks saying there was congestion in this area and maybe there could be something done at the design stage to fix it or sometimes just, you know, the design implementation is moving in parallel with the PD placement, and as you get updated design, you then feed it back into the PD team.

Q. So the PD team would actually test the hardware prototype, right?

A. They were not concerned with testing the hardware prototype. They were concerned with the layout. So the placing of gates and making it fit in the smallest area and the routing of the metal from one gate to the next.

Q. And it was another team that did kind of hardware implementation?

A. Yeah. That was the design team that I was associated with. I was mainly associated with the architecture team, but the architecture team and the implementation team were not totally separate.

MR. McNAMARA: I'd like to have

marked for identification purposes as Gruber Exhibit 18 a document bearing the Bates numbers AMD1044_0242863, please.

(Document marked as Exhibit 18 for identification)

Q. Sir, do you recognize this document?

A. No.

Q. Do you know whether this is a screen capture of the R400 hardware prototype?

MR. SCHWENTKER: Objection. Form. Calls for speculation.

A. I don't know. I don't know if this is from IKOS, whether it's from the C level model or whether it's from the RTL running on cadence mentor type tools.

Q. I need your help with pronunciation. Lefebvre?

A. Lefebvre.

Q. Do you know whether Mr. Lefebvre would be closer to determining whether this was an output from the IKOS system or C++ synthesis models?

MR. SCHWENTKER: Same objections.

A. Yeah, he certainly might know

better than I would, but he isn't the ideal person to say either.

Q. Who would be the ideal person?

A. It has been a long --

MR. SCHWENTKER: Same objections.

A. -- time.

Q. If you know?

A. I don't know.

MR. McNAMARA: Could we take a short break.

VIDEOGRAPHER: The time is 5:37 and we're off the record.

(Recess taken at 5:37 p.m. and reconvening at 5:44 p.m.)

VIDEOGRAPHER: We are back on the record. The time is 5:44.

BY MR. McNAMARA:

Q. Sir, if you recall before we took a break, we were talking about what components of the R400 were part of the design when it was kicked off in November of 2000. Do you recall that?

A. Ask the question again?

Q. Sure. Before we took a break, early in that session, we were talking about

what components were part of the R400 design when it was kicked off right around November of 2000. Do you recall that?

A. Yes.

Q. Was a rasterizer part of the conceptual design at that time?

A. It was part of the conceptual design, but we weren't anticipating major changes in the rasterizer. The unified shader was pretty much agnostic to -- or the rasterizer was pretty much agnostic to a unified shader.

Q. And the render back end was also part of that conception?

A. Yes.

Q. One of the things that the unified shader uses are general purpose registers, right?

A. Yes.

Q. The design of the general purpose registers was finalized fairly early on in the design process, right?

MR. SCHWENTKER: Objection to form.

A. Yes.

Q. And that was finalized prior to the end of 2001; is that right?
A. I can't really say.
MR. McNAMARA: I would like to introduce for identification purposes as Gruber Exhibit 19 a document bearing the Bates numbers AMD1044_0175251 through 75463, which is the March 17, 2005 deposition of Andy Gruber.
MR. SCHWENTKER: Counsel, just to correct the record. You said 2005. This is 2015.
MR. McNAMARA: I did, and I apologize. Thank you for the correction.
(Document marked as Exhibit 19 for identification)
Q. Sir, do you recognize this document?
A. Yes, I do. This is my deposition.
Q. If I could direct your attention to Page 15 of the document, particularly lines 4 through 10.
A. Okay.
Q. Could you please just read those

through.

A. Okay.

Q. In this particular portion of the transcript, you were testifying regarding the general purpose register of the R400, right?

MR. SCHWENTKER: Objection.

A. Yes.

Q. Here you testify that the general purpose registers of the R400 were finalized prior to the end of 2001. Do you see that?

A. Yes.

Q. Is that consistent with your recollection of when the general purpose registers were finalized for the R400?

A. It is consistent, and back then, you know, I had familiarized myself with the timeline of the R400. So I was much more aware of the timeline than I am now, but I don't have any reason to doubt or contradict the testimony that I gave earlier.

Q. Sir, do you know why Qualcomm purchased the mobile business unit from AMD?

MR. LEVENTHAL: Form.

MR. SCHWENTKER: Objection. Same objection.

A. I understand at a high level what their motivations were.

Q. Was Qualcomm's motivation at all to purchase the know-how related to the GPUs developed by AMD?

MR. SCHWENTKER: Same objection.

A. That was certainly one of their motivations, yes.

MR. LEVENTHAL: Form.

Q. Qualcomm didn't develop on its own a GPU with a unified shader, correct?

MR. LEVENTHAL: Form.

MR. SCHWENTKER: Same objection.

A. That is correct.

Q. And in your own experience developing that GPU with the unified shader was a very time-consuming and laborious process, right?

A. Yes.

MR. McNAMARA: I'd like to have marked for identification purposes as Gruber exhibit -- let me withdraw that.

Q. Qualcomm also purchased an earlier GPU from AMD with a fixed function pipeline; isn't that right?

MR. LEVENTHAL: Form.
MR. SCHWENTKER: Same objection.
A. Yes.
Q. And then instead of developing a GPU with a unified shader, Qualcomm purchased that know-how from AMD, right?
MR. SCHWENTKER: Same objection.
MR. LEVENTHAL: Form.
A. Well, Qualcomm had also purchased IP from AMD that included a unified shader as well.
Q. And by purchasing that IP from AMD that included a unified shader within a GPU, that made it so Qualcomm didn't have to develop a GPU with a unified shader on its own, right?
MR. SCHWENTKER: Same objection.
MR. LEVENTHAL: Form.
A. That is true.
MR. McNAMARA: I'd like to have marked for identification purposes as Gruber Exhibit 20 a document bearing the Bates numbered AMD1044_0256597 through 599.
(Document marked as Exhibit 20 for identification)

MR. LEVENTHAL: While that's being marked, to the extent the ITC protective order requires separately identifying different third parties, I'd also like this to be indicated Qualcomm confidential business information.

Q. Now, sir, to your understanding, did NVIDIA eventually adopt the use of a unified shader in its GPUs?

MR. SCHWENTKER: Objection. Calls for speculation and a legal conclusion.

A. Yes.

Q. Now, the document, Gruber Exhibit 20, is an FAQ from the NVIDIA website. Do you see that?

A. Yes.

Q. Now, this FAQ discusses the NVIDIA GeForce product, right?

A. Yes.

Q. NVIDIA GeForce product is a GPU, right?

MR. SCHWENTKER: Objection. Calls for speculation.

A. Yes.

Q. Now, if you take a look at the

second page, there's a question here about the unified shader, the sixth question down.

A. I see that.

Q. Now, is it your understanding that the first time NVIDIA incorporated a unified shader into one of its products, it was incorporated into the GeForce 8 series of GPUs?

MR. SCHWENTKER: Same objection.

A. I don't remember whether their previous ones had a unified shader or not. So I would just be answering based on what's in this particular document, which it appears that this is a new introduction.

MR. McNAMARA: That's all I have.

MR. SCHWENTKER: Off the record for just a minute.

MR. McNAMARA: No. Sorry, that's fine.

VIDEOGRAPHER: The time is 5:55, and we're off the record.

(Recess taken at 5:55 p.m. and reconvening at 5:57 p.m.)

VIDEOGRAPHER: We're back on the record. The time is 5:57.

RE-EXAMINATION BY
MR. SCHWENTKER:
Q. Mr. Gruber, counsel for AMD asked you some questions about Exhibit 19.
A. Yes.
Q. And so that Exhibit 19 is your deposition transcript from your 2015 deposition that we've talked about earlier today; is that correct?
A. Correct.
Q. And that when you testified in that deposition, you were appearing on behalf of AMD as a corporate witness; is that correct?
A. Correct.
Q. And I believe you said earlier today that your testimony or your appearance in that deposition was in the interests of your current employer, Qualcomm, as well?
MR. LEVENTHAL: Form.
A. That is correct.
Q. If you could turn to Exhibit 12, please. This is the document showing a title of "ATI and NVIDIA Proclaim Different Graphics Processors Architecture Goals."
Have you ever seen this before?

A. I might have. I was certainly keeping track with some of these types of articles back in the day, but I don't have any specific recollection of this.

Q. So you have no personal knowledge of this document, that you're aware of?

A. Correct.

Q. Could you turn to Exhibit 13, please. This is the document with the title "NVIDIA Chief Architect: Unified Pixel and Vertex Pipelines - The Way to Go."

A. Yes.

Q. Have you ever seen this before?

A. Again, I have no specific recollection of seeing this prior to today.

Q. So you have no personal knowledge of this document that you're aware of?

A. Correct.

Q. Exhibit 14, please. This is the document with the title "Microsoft and ATI Technologies Announce Technology Development Agreement."

Have you ever seen this document before today?

A. I have no specific recollection of

this, seeing it before today.

Q. So you have no personal knowledge of this document?

A. Correct.

Q. Exhibit 16 is the document titled "Executive Review R400" -- it says "October 15, 2002" on it?

A. Right.

Q. Do you have any personal knowledge of this document?

A. No.

Q. If you could turn to Exhibit 17, please. This is the document titled "PD Team Meeting," with a date of January 9, 2003 on it?

A. Correct.

Q. Do you have any personal knowledge of this document?

A. No.

Q. If you could turn to Exhibit 18, please. This is the page showing a series of four pictures, it looks like.

A. Right.

Q. Do you have any personal knowledge of this document?

A. No.

Q. And Page 20 -- strike that.

Exhibit 20, this is the -- it appears to be a FAQ. Down at the bottom it says -- it looks like this is from the Web Archive?

A. Yes.

Q. Do you have any personal knowledge of this document?

A. No.

Q. You've never seen this before?

A. I'm not saying I never saw it. I just don't have any recollection of seeing it.

Q. So as far as you know, sitting here today, you've never seen it before?

A. Right.

MR. SCHWENTKER: I have no further questions.

MR. McNAMARA: No questions.

MR. MECHELL: I would state again for the record that this is marked highly confidential outside counsel only pursuant to the protective order in the Northern District of California matter.

MR. LEVENTHAL: I don't have any questions. As I said before, Mr. Gruber

will review and sign. Off the record.

VIDEOGRAPHER: The time is 6:02 and we are off the record.

(Whereupon the deposition concluded at 6:02 p.m.)

C E R T I F I C A T E

I, ANDREW E. GRUBER, do hereby certify that I have read the foregoing transcript of my testimony, and further certify that it is a true and accurate record of my testimony (with the exception of the corrections listed below):

Page	Line	Correction/Reason
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____
____	____	_____

Signed under the pains and penalties of perjury this day of , 2017.

_____
ANDREW E. GRUBER

C E R T I F I C A T E

COMMONWEALTH OF MASSACHUSETTS

SUFFOLK, SS.

I, Michael O'Connor, Registered Merit Reporter/Certified Realtime Reporter, and Notary Public in and for the Commonwealth of Massachusetts, do hereby certify:

That ANDREW E. GRUBER, the witness whose testimony is hereinbefore set forth, was duly sworn by me and that such testimony is a true and accurate record of my stenotype notes taken in the foregoing matter to the best of my knowledge, skill and ability.

IN WITNESS WHEREOF, I have hereunto set my hand and Notarial Seal this 27th day of July 2017.

_____

MICHAEL O'CONNOR, RMR, CRR, CRC
Notary Public

My Commission expires:
November 22, 2022