

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

MERCEDES-BENZ USA, LLC
Petitioner,

v.

DAEDALUS PRIME LLC
Patent Owner

IPR2023-01344
Patent 9,575,895 B2

Record of Oral Hearing
Held: December 9, 2024

Before JAMES P. CALVE, GREGG I. ANDERSON, and ARTHUR M.
PESLAK, *Administrative Patent Judges*.

IPR2023-01344
Patent 9,575,895 B2

APPEARANCES:

ON BEHALF OF THE PETITIONER:

SCOTT HUGHES, ESQUIRE
Hogan Lovells US, LLP
Columbia Square
555 Thirteenth Street NW
Washington, DC 20004
(202) 637-5600
scott.hughes@hoganlovells.com

KYLE XU, ESQUIRE
(415) 374-2442
Kyle.xu@hoganlovells.com

ON BEHALF OF THE PATENT OWNER:

TAREK FAHMI, ESQUIRE
Ascenda Law Group, PC
2150 N. First Street
Suite 420
San Jose, CA 95131
(866) 877-4883
tarek.fahmi@ascendalaw.com

The above-entitled matter came on for hearing on December 9, 2024, commencing at 1:00 p.m., via video teleconference.

P R O C E E D I N G S

- - - - -

JUDGE ANDERSON: This is a hearing on IPR2023-01344.

Petitioner is Mercedes-Benz, USA, LLC, and Patent Owner is Daedalus Prime, LLC, owner of U.S. Patent No. 9,575,895, which is the challenge patent here. I am Judge Anderson. I am joined by Judges Calve and Peslak. We are all participating via teleconference, so there are some things we need to keep in mind about that particular situation. Chief among this is demonstrative exhibits. When referring to the demonstrative, please describe any slide you are discussing by slide number so that everyone can follow your arguments. We have both parties' demonstratives, and we also have the complete records, so you can direct us to whatever you need to, to follow up on your arguments.

I will say, as far as my procedures, I do not like the presenting party to be interrupted. Therefore, no objections at the time of the presentation. If you have an argument you want to make about beyond the scope or some similar objection, you can make that during your own presentation. And then you can make your argument about what the objection is. The panel can address that later if it's needed. I also want you to be aware that this is a public hearing. There is a public line for this hearing, and be aware that absence -- well, that whatever you say will be potentially available to the public.

Per our hearing order, each party will have 60 minutes to present its argument, total of one hour. Petitioner is granted an additional 15 minutes during which LEAP practitioner Zhang, I hope I said that correctly,

1 also Kyle Xu, is permitted to present all or any part of Petitioner's argument
2 for a total of 75 minutes for Petitioner.

3 Because Petitioner has the burden to show unpatentability of the
4 challenge claims, Petitioner will proceed first, followed by Patent Owner.
5 Petitioner may reserve time to rebut Patent Owner's position and Patent
6 Owner may reserve time for a sur-rebuttal.

7 We do not have a timer available, though I am pulling out my
8 handy dandy smart phone and I will attempt to keep good time and warn all
9 of you who are presenting as the time gets close to being expired.

10 Let's see, anything else? I think that's it from my standpoint.
11 Absent any questions regarding this procedure, I'd like the counsel present to
12 introduce themselves, who they represent and who is present either virtually
13 or physically. And with that, Petitioner, please make your introduction.

14 MR. HUGHES: Good afternoon. Thank you, Your Honors. My
15 name is Scott Hughes from Hogan Lovells on behalf of Petitioner Mercedes-
16 Benz USA. And I have with me, Kyle Xu, who will be presenting on one of
17 the issues here as well. Also in the room with me, we have one of our
18 associates, Jake Falcetti, who is working our demonstrative slides today.

19 JUDGE ANDERSON: Very good, welcome to you all. For Patent
20 Owner who's presenting, please.

21 MR. FAHMI: Good afternoon, Your Honor. It's Tarek Fahmi on
22 behalf of the Patent Owner, and it's just myself here today.

23 JUDGE ANDERSON: Okay, Mr. Fahmi, welcome. We
24 understand the primary issue here. I'm not going to repeat it. I would expect
25 you to deal with it. It does relate to the single caching agent, and then there
26 is an issue relating to the admissibility of one exhibit. So with that,

1 Petitioner, you may proceed.

2 MR. HUGHES: Thank you, Your Honors. And I'd just like to
3 reserve 15 minutes for rebuttal before we get into our main argument here.

4 JUDGE ANDERSON: I have it, thank you.

5 MR. HUGHES: All right. Good afternoon, Your Honors. We are
6 here for IPR2023-01344 regarding U.S. Patent Number 9,575,895, which I
7 may refer to as the '895 patent. Going on to slide 2, we have here on the
8 table of contents the topics we'll be discussing today. First, we'll give an
9 introduction to the '895 patent and a discussion of that single caching agent
10 limitation. After that, we'll get into the disputes, the first of which is
11 whether the reference Sinharoy used in one ground of the petition discloses
12 the single caching agent limitation. The second dispute is whether the other
13 base reference in the petition, Tang, also discloses the single caching agent
14 limitation. Related to that, dispute three is specifically whether Tang
15 discloses that the caching agent has distributed portions, each associated
16 with a core. And then the last dispute is whether a secondary reference to
17 Xu was publicly available before December 13th, 2011.

18 Going on to slide 4, we have here an overview of the '895 patent.
19 The '895 patent relates to a multi-core processor computer system, having a
20 plurality of cores and an integrated input output, or IO module. And this IO
21 module interfaces between the multi-core processor and an IO device. The
22 multi-core processor further includes a caching agent that performs cache
23 coherency operations. And in particular, this caching agent performs cache
24 coherency operations for the plurality of cores. And it performs cache
25 coherency operations for the IIO module.

26 Going on to slide 5, you can see here on the left, we have claim 1

1 from the '895 patent. This is one of the independent claims. And there are
2 two other independent claims here, both of which are very similar to claim 1
3 and are going to have the same limitation and dispute of the single caching
4 agent. I note that one of the independent claims is directed to a computer
5 readable medium instead of a processor. And the other, claim 15, brings in
6 not just one processor, like we have in claim 1 here, but introduces
7 limitations that there is a second processor also having some of these
8 components.

9 So, looking here at claim 1, which will lay out the dispute between
10 the parties, claim 1 generally recites a processor, including a plurality of
11 cores, a shared cache memory, a memory controller, an integrated input
12 output IO module to interface between the processor and an IO device, a
13 caching agent to perform cache coherency operations for the plurality of
14 cores and the IIO module. The caching agent is a single caching agent for
15 the processor and includes a plurality of distributed portions, each associated
16 with the corresponding core. And we can see that at the bottom of claim 1
17 on the left, which is really the limitation in dispute here today.

18 Going on to slide 6, what is a single caching agent in the context of
19 the '895 patent? Well, the '895 patent broadly describes the single caching
20 agent for the processor. It does not discuss a particular structure or circuit
21 implementation from the single caching agent. It doesn't indicate that it's
22 one single component or one single structure. And in fact, it even indicates
23 that the caching agent can be formed from distributed portions and that it
24 may be logic. We can see at the top righthand side of slide 6, at column 2,
25 lines 41 to 45, the '895 patent talks about Figure 2, which is a block diagram
26 showing a caching agent. And this even states that that caching agent 200

1 may simply be logic that's interposed between the cores and an LLC.

2 Going down to column 3, lines 6 to 10, in addition, the '895 patent
3 again says, understand that a caching agent can be distributed such that each
4 of different portions of the caching agent can be associated with the
5 corresponding core and LLC bank or slice. So right here, this single caching
6 agent does not need to be one single structure or one single component. It
7 can be multiple distributed components that act together. We can see this if
8 we look at Figure 6 of the '895 patent, which we have here on the righthand
9 side of slide 7. And in describing Figure 6, the '895 patent says that the
10 processor 700 includes a distributed configuration having partitions or slices,
11 each including a core 710. We can see those on the left-hand side in the
12 middle, and a partition of a caching agent 715, those are highlighted in red,
13 and then LLC 720. And here, you can see there's two different sides here
14 with these distributed caching agents, and then they --

15 JUDGE PESLAK: Counsel?

16 MR. HUGHES: Yes.

17 JUDGE PESLAK: This is Judge Peslak. What are those sort of
18 dashed, I don't want to call them circles. They look like rounded rectangles,
19 but what does that represent in this figure?

20 MR. HUGHES: Yes, Your Honor. You're talking about the
21 dashed lines. For example, if we see that top arrow, the up and down double
22 arrow, it's connecting to that in between the two sets of caching agents and
23 cores.

24 JUDGE PESLAK: Yes, that's what I'm referring to.

25 MR. HUGHES: Yes, so there's no reference numeral on this.

26 Based on the description in the patent, it seems like this may be potentially a

1 ring interconnect or just something else that indicates that these caching
2 agents, for example, may be communicating with each other or otherwise
3 being able to share information. But there's no given reference numeral
4 here, and so it doesn't specifically call that out.

5 JUDGE PESLAK: I have another question though. If I look at
6 claim 1, is it fair to say that the caching agent is recited functionally and
7 there's no physical structural limitation to the caching agent?

8 MR. HUGHES: That's correct, Your Honor. If we see it says the
9 caching agent is a single caching agent for the processor and includes these
10 plurality of distributed portions, if we look up further in the claim, the first
11 time we see it mentioned in red there, this in fact calls out what this single
12 caching agent is. It's a caching agent that performs caching operations for
13 two different things. First, it performs cache coherency operations for the
14 plurality of cores, and it also performs cache coherency operations for the
15 IIO module. So, that's how the claim has defined this by what it does.
16 Instead of having one entirely separate and distinct caching agent to do
17 cache coherency operations for the cores, and another entirely separate and
18 distinct caching agent for the IIO module, it's saying that you can have one
19 that's doing the coherency operations for both of those things.

20 JUDGE PESLAK: Okay, thank you.

21 MR. HUGHES: And going back to slide --

22 JUDGE CALVE: This is Judge Calve. I have a follow-on
23 question to your very helpful answers. To Figure 6, so it's your contention
24 that each of the individual coaching, I'm sorry, caching agents that are
25 located with a respective core, those correspond to distributed portions. Is
26 that -- am I correct in understanding that's your position?

1 MR. HUGHES: That's correct, Your Honor. And you can see
2 here on the left-hand side, this is some of the text from the '895 that
3 describes this. And it calls Figure 6, this distributed configuration, having
4 the partitions or slices. Now, again, these are --

5 JUDGE CALVE: What makes them then they're distributed, what
6 makes them also a single caching agent? What aspect functional or
7 structural or both ties them together to satisfy that limitation?

8 MR. HUGHES: So, if we look on the left-hand side of the slide,
9 we have from column 6 of the patent, lines 35 to 40. And this is where the
10 patent indicates that while you may have these distributed caching agents,
11 they form a single caching agent. And again, it goes to what they're doing
12 functionally, what are they doing? They're handling the cache coherency
13 operations for the cores, as well as an IIO module. So, they may be
14 distributed. And in fact, like we saw in the earlier slide, these could even
15 just be logic. So, while they're shown as a block diagram here, these
16 different caching agents, by definition of the '895 patent, could in fact just be
17 logic that's associating some cache coherency operations being performed,
18 linking that to a specific core or LLC.

19 So overall, the patent is really describing this pretty broadly by
20 saying single caching agent, but then going back and saying, well, it could
21 be distributed portions, which all together form the single caching agent. It
22 just has to do these cache coherency operations for the cores in the IIO
23 module. And so, what we're really left with is when we see this in the claim,
24 it's something that is actually very broad.

25 JUDGE CALVE: So did the cache -- caching agent distributed
26 portions that function as a single larger entity have to be able to

1 communicate with one another. Is that how they perform cache coherence
2 one way?

3 MR. HUGHES: Well, now, I don't know if they need to
4 specifically communicate with each other, but for them to perform and form
5 this single caching agent, there does have to be some shared responsibility,
6 or if they're not communicating directly with each other, there's got to be
7 something or some information about what they're doing in the functionality
8 that they're performed is either being overseen or somehow controlled on
9 kind of that chip level. You know --

10 JUDGE CALVE: Is that aspect depicted in Figure 6? And this
11 goes back to Judge Peslak's question, is that what that double ring is trying
12 to depict logic or structure that enables these disparate distributed caching
13 agents to function as a single agent for cache coherence?

14 MR. HUGHES: It could be Your Honor, but like I said the patent
15 itself doesn't really talk about that and it doesn't give it a reference numeral,
16 so it's possible that it could be, but I don't want to speculate as it was not
17 specifically called out in the description of the patent itself.

18 JUDGE CALVE: Okay. Well, if we look at the description
19 please, this would be column 6 of the patent, granted there's no numeral 730
20 to denote this ring interconnect. But the ring interconnect is described
21 beginning around lines 58 to 67 of column 6 as being interconnected with a
22 number of other elements shown in Figure 6 that seemed to be in Figure 6
23 connected to that unlabeled loop. So, the ring interconnect 730 is connected
24 to the home agent 770 and also is connected to the QPI 780 and I think it
25 also communicates with the IIO which seems to be illustrating pretty much
26 what you've been describing here to for as the function of this single agent.

1 It ties together the integrated IO with the caching with some other elements.
2 So, would you say -- would you take a position that that element probably is
3 the ring interconnect?

4 MR. HUGHES: Yeah, you know, as I mentioned when I first
5 talked about this figure I said it's likely it is the ring interconnect, they just
6 didn't give it that reference numeral.

7 JUDGE CALVE: Well, if you look up at Figure 2 of the same
8 patent they do actually label a ring interconnect, and this is the caching agent
9 right, 200, and to the right in Figure 2 there is element 260 which is a ring
10 interconnect. So, that seems to be indicating that this ring or interconnect
11 form some function that interconnects these distributed caching agents, or is
12 that not a correct understanding?

13 MR. HUGHES: It's with Figure 2 we have single caching agent
14 200 shown and then 260 is off to the right of that. So, it's unclear from that
15 figure, I think, whether it's working within the single caching agent or
16 connecting different caching agents. For example, there's discussion in the
17 patent of having multiple processors, each with its own single caching agent
18 that the ring interconnect could also be connecting to those. But in general,
19 it seems like there is some kind of communication or connection, you know,
20 for these different distributed portions of the single caching agent for them
21 to work together.

22 JUDGE CALVE: Right. It seems to be the basis that allows this
23 caching agent to maintain coherence without the need for additional caching
24 agents. That's cited at column 2, say, lines 59 to 67 again. So, it's
25 elaborating a little bit on the functionality of the ring interconnect 260,
26 which is labeled in Figure 2. It seems to indicate, to me at least, that this

1 would be a means for each of these. Right. This is an individual agent, but
2 for them to intercommunicate for cache coherence, you're describing this
3 with the claims are looking for, apparently. This ring interconnect plays
4 some functional role. I don't know what structure it has.

5 MR. HUGHES: Yeah. And I will note, Your Honor, that the ring
6 interconnect is in fact, a limitation in certain dependent claims, but the
7 independent claims are broader than that and just recite the single caching
8 agent generally without requiring specifically a ring interconnect to be
9 performing anything for it to act as that single caching agent.

10 JUDGE CALVE: Okay. So, in the dependent claims, is the ring
11 interconnect denoted as one embodiment of something that provides cache
12 coherence? It's not the full scope, I guess, of the independent claims, but it's
13 indicating the independent claims include at least that the ring interconnect
14 could provide the cache coherence?

15 MR. HUGHES: Well, for what claim 2 recites, Your Honor, is
16 that the plurality of distributed portions of the caching agent are coupled via
17 a ring interconnect, that's the entirety of that claim. It doesn't go into any
18 more details than that. And then we have claim 10. Again, there's a ring
19 interconnect that couples the plurality of cores and the shared cache memory
20 via a caching agent. So, you know, right there, it's doing some kind of
21 connection, but it's not -- the claims, even those dependent claims are not
22 linking certain coherence operations with that ring interconnect.

23 JUDGE CALVE: Okay. Just within the scope of the claims, what
24 are some examples of cache coherence functions that are going to be
25 performed by this logic within the scope of the claims?

26 MR. HUGHES: So, within the scope of the claims, and like I can

1 get to if we go, like I can pause here, but, you know, you're going to have
2 things like read operations and write operations, fetches, and things like that
3 that are going to be going on and checking these different LLCs and the
4 cores. And those are going to be some of our, what we're going to need the
5 coherency for.

6 JUDGE CALVE: So, it's just so that the two different caching
7 agents don't try to grab the same program code. It's just one aspect of what
8 it does?

9 MR. HUGHES: That's part of it, Your Honor. And I think also
10 that if you have one caching agent modify a line, you don't want to have
11 another caching agent modify that same line in a different way. You want to
12 maintain the coherence so that a single cache line that may be used in
13 different places is going to be the same across them all.

14 JUDGE CALVE: Okay, thank you.

15 JUDGE ANDERSON: I had one question, counsel. With respect
16 to your position that the two prior art references each show a single caching
17 agent, you rely on a logical single caching agent. What do you mean by a
18 logical single caching agent?

19 MR. HUGHES: Well, what we have, and I'll go back, Your
20 Honors, to slide 6 here. And you can see at the top righthand side, the patent
21 itself even indicates that this caching agent 200 may simply be logic
22 interposed between the cores and the LLC. So, what we're talking about
23 when we have a logical single caching agent is these may not need to be
24 separate physical components that are linked together, but because this
25 caching agent can in fact be logic, you can have logic partitioned or
26 developed in such a way that certain portions or distributed portions of that

1 logic would correspond, for example, to a given core. So, within the logic of
2 how these controllers are operating, that would be what the single caching
3 agent is. You can have dedicated logic, you know, you don't have to, like I
4 said, have physically connected components, but you can have separate
5 components and then logic that may have portions that make it so that this
6 portion is going to go to this core, this portion is going to go to this core and
7 things like that, Your Honor.

8 JUDGE ANDERSON: You seem to say at one point in some of
9 the papers that where the caching agents do not work independently, they
10 work together or they're shared is another term that's used. Is that part of
11 what you're saying when you say a logical single caching agent?

12 MR. HUGHES: Yes, Your Honor, in that through their logic, they
13 have to know what each other is doing. They're interconnected, they're not
14 operating entirely independently of each other.

15 JUDGE ANDERSON: So, if they are operating independently,
16 they would not be a shared single caching agent?

17 MR. HUGHES: Correct. If you had one caching agent that all it
18 did was, for example, the cache coherency operations for a given core and
19 had no information or relation to what another caching agent is doing, let's
20 say with respect to a different core or an input output module, those could be
21 separate cache controllers that this patent is discussing. Maybe you can
22 reduce that by bringing in some of the functionality on both. But yes, if they
23 were entirely separate and distinct like that, with no relation to each other,
24 then they would not be a single caching agent.

25 Okay. And now, Your Honor, I know you mentioned we talked
26 about the two references disclosing this. We're going to get into that now.

1 First, going to slide 9, this is an overview of the first base reference used in
2 the petition, which is Sinharoy. And Sinharoy is a paper directed to the
3 IBM's POWER7 chip, which is a multi-core server processor. We see here
4 on the left-hand side of the slide, we have Figure 1, which is a die photo of
5 this IBM Power processor. And it includes the limitations we saw in the
6 claim. The plurality of cores, we can see that there's eight cores here, those
7 eight black boxes. There's a shared 32-megabyte L3 cache, which comprises
8 eight 4-megabyte L3 regions. There's two memory controllers that interface
9 with DDR memory chips, two integrated IO controllers that interface
10 between the processor and IO devices. And in Sinharoy, one of these IO
11 devices could be an IO hub chip, but then further connects on and is
12 interfaced with additional peripheral components. And then there's also
13 centralized controllers for the L2 and the L3 caches. And these together
14 form this single caching agent, which performs the overall cache coherency
15 operations.

16 JUDGE PESLAK: Counsel, this is Judge Peslak again. Can you
17 go back to your prior slide? Correct me if I'm wrong, but I don't see the
18 centralized controllers in this drawing.

19 MR. HUGHES: Yes, Your Honor, those are described and we
20 have on the slide here, pages 1-18 to 19. And I know all the pages in
21 Sinharoy have the 1:18-19. If Your Honors prefer, I can just refer to this as
22 pages 18 and 19 to be clear on the record. Would that be okay?

23 JUDGE PESLAK: I -- yeah, as long as we know what you're
24 referring to, if it's easier to say 18 and 19.

25 MR. HUGHES: So, I'll try to see if this makes sense to Your
26 Honors. While that's not shown in the figure, like we have here, pages 1:18-

1 19 is where Sinharoy discusses that these centralized controllers are present.
2 We can see, for example, the L2 cache and then we have the shared L3
3 cache in between. So, these controllers would be there within their, you
4 know, those regions in the figure, but they're not specifically called out or
5 identified in this figure from the reference.

6 JUDGE PESLAK: Do you know if they're in -- identified in any
7 other figure of Sinharoy?

8 MR. HUGHES: I can look for that specifically, Your Honor, and
9 come back to you on rebuttal if there's a particular figure. I believe that, you
10 know, they're discussed in general and there are some figures later on that
11 kind of talk about some of the functionality. But if there's, for example, a
12 block diagram that has these figures in there, I'll take a look and we'll let you
13 know on rebuttal time.

14 JUDGE PESLAK: Right. That might be helpful. I like to see
15 things sort of laid out on a drawing, you know, with the text, if you know
16 what I mean. That might clear it up as to, you know, what are the distributed
17 portions of this controller that form this single caching agents. Okay. That's
18 what I -- and when you're going through that, if you can focus on that, at
19 least for me, that would be very helpful.

20 MR. HUGHES: Yes, Your Honor. Thank you.

21 JUDGE CALVE: Okay. Counsel, Judge Calve, back to that
22 figure, please. So, if I understand this, you're saying the L2 caches, is that
23 physical real estate on the integrated circuit, that silicon we're looking at?
24 And the same with the L3, or is this logic?

25 MR. HUGHES: I think it could be either, Your Honor. I'm not
26 sure if the reference calls it out one way or the other. What this reference is

1 really talking about is some changes from the previous POWER6 chip, and
2 some of the additional efficiencies and benefits that they've gotten in this
3 new version, which includes bringing some things like certain controllers
4 and this shared, certain RAM, and some of the shared caches onto the chip,
5 instead of being external.

6 JUDGE CALVE: So, the cores are physical structures. They're
7 transistors. So, the L2 cache, perhaps one embodiment of the L2 cache is
8 physically situated next to the core for proximity for efficiencies.

9 MR. HUGHES: Yeah, it could be.

10 JUDGE CALVE: (CROSSTALK)--

11 MR. HUGHES: Yes.

12 JUDGE CALVE: For your contentions, do you need both L2 and
13 L3? I'm not sure this matters but it seems like either the L2 or the L3
14 separately if they have single controllers operating them would map to the
15 claim language. Is that right, or is there a reason you need L2 and L3 with
16 their separate independent controllers?

17 MR. HUGHES: So, they could read independently, Your Honor.
18 The reason we're pointing to them both is we want to make sure that we're
19 clear and the reference is clear how it's hitting that dual functionality that the
20 claimed single caching agent has to perform, that it's performing cache
21 coherency operations for the cores and cache coherency operations for the
22 IIO module for the IIO devices. And so, we're pointing to the L2 cache
23 controllers and the L3 cache controllers, because when you look at those
24 together, it's clear that they're hitting both of those functionalities.

25 JUDGE CALVE: Can you show that -- could you show that to us,
26 please, or show it to me at least please on that previous Figure 1 slide of

1 yours? The one previous. Yes.

2 MR. HUGHES: Yeah, so this is back to the slide 9, Your Honor.
3 And here we've got, you know, like I said the L3 cache controllers are not
4 shown on this figure, but we'll go when we get into the reference and it and
5 its particular descriptions, it will become clear what those are and where
6 they're operating within the chip.

7 JUDGE CALVE: Okay. I look here also at the disclosure. This is
8 the written description that goes with this reference, page 18. It says the L2
9 cache uses 128 byte lines and is eight-way set associative. That first
10 sentence, doesn't that by itself pretty much show that they're operating as a
11 single cache? I'm not sure if that's in the petition, but I think 128 byte lines
12 that sounds like physical structure, like a bus or something, an eight-way set
13 associative. Doesn't that mean they're interconnected some way? They're
14 performing cache coherence as a set, as a single element.

15 MR. HUGHES: Yes, and that is part of our argument in the
16 petition, Your Honor. And we focused a little bit on, for example, the L3
17 cache, where the same statement is made, and have it sent back to page 18 as
18 well. But again, like we explained for the L3 region, I think for that one,
19 you know, those act as a 32 set, or highly set associative of set. So yes, we
20 are pointing to this, that what the reference is saying, is that these are highly
21 set associative, they're working together.

22 JUDGE CALVE: Yeah, and if it's 32 megabyte, that connotes
23 some physical structure, as does, I guess, it needs a physical line in and out.
24 It's connected to the core to communicate, right? The bytes and bits?

25 MR. HUGHES: Yes, Your Honor.

26 JUDGE CALVE: Okay, thank you.

1 MR. HUGHES: So, going on to slide 10 here, it's clear from the
2 Sinharoy reference that it's performing cache coherency operations, as is
3 recited in the claim. I have some here, the fetches, the read-write core
4 prefetches, and then it also indicates that these same L2 and L3 cache
5 controllers support cache injection, and this is a coherent operation. And in
6 addition, just with respect to claim 15, like we noted, that requires not one
7 processor but two, Sinharoy discloses that you could have a system which
8 combines up to four of these POWER7 chips, meaning you would have the
9 first processor and the second processor.

10 Going on, this is the limitation in dispute from claim 1 on slide 11,
11 and again, it's really the single caching agent which is for the processor that's
12 performing the operations for the cores and the IIO module.

13 Going on to slide 12, Sinharoy has this caching agent. Again, the
14 L2 cache includes a single centralized controller, and each L3 cache region
15 also includes a single centralized controller. Now, the L2 and L3 caches
16 both have read, claim, or RC machines, and these manage the cache
17 coherence operations, including those core fetches, core prefetches, and re-
18 modify write operations. If we look at the righthand side of slide 12,
19 Sinharoy at page 18, indicates, like we were just talking about, that there's a
20 set of eight fully functional RC machines that are incorporated into the L2
21 cache controller, and then there's lightweight RC machines in the L3 cache
22 controller, and we can see these are managing things like core fetches and
23 stores and prefetches and stores.

24 Going on to slide 13, these L2 and L3 cache controllers work
25 together, and Dr. Horst explained this, including in his declaration, Exhibit
26 1017 at paragraph 18. There, he indicated that these RC machines are

1 shared across the L3 regions and hence are part of a logical single cache
2 controller. He says if the L3 regions had separate independent cache
3 controllers, they would each need their own RC machines, and there would
4 be no opportunity to globally reorder the storage operations. So, these RC
5 machines further support that the L2 and L3 cache controllers work with
6 each other and share information and logic across cores, making them part of
7 this logical single caching agent.

8 If we go on to slide 14, we can see here again, like I talked about,
9 Sinharoy is also discussing these L2 and L3 controllers performing cache
10 coherency operations for the IIO module. And in particular, Sinharoy
11 discloses that cache injection is one of the operations that's performed by the
12 L2 and L3 cache controllers. And these cache injections are something that
13 may come from or relate to the IIO module. And again, on the righthand
14 side, we can see here that Sinharoy is talking about these pool of read
15 machines, a pool of 24 lightweight RC machines. And clearly what they're
16 doing is, you know, these are what are operating on those cache coherency
17 operations as we see here.

18 Now, it's this knowledge of the other L3 caches that show that
19 these controllers are working together. If we see here specifically, Sinharoy
20 at pages 19 to 20 on the left-hand side towards the middle of that excerpt,
21 Sinharoy even states that, note, when the core and L2 associated with a
22 given L3 region are disabled, the remaining L3 regions enjoy full access to
23 the L3 region. When the core and the L2 are enabled, there are heuristics
24 that cooperate to give precedence to the associated L2 cache while allowing
25 load and capacity balancing among all eight L3 regions on the chip, thereby
26 causing them to behave as a highly set associative 64-way adaptive 32-

1 megabyte L3 cache. And Dr. Horst helped to further explain what this
2 would mean to a POSITA. As he said at paragraph 16 of Exhibit 1017, a
3 POSITA would have understood that the reason why the remaining L3
4 regions still enjoy full access to the disabled L3 region is that the other cache
5 controllers still have knowledge about the disabled cache region. They do
6 not operate independently from each other. And in addition, when the cores
7 are enabled, he indicates that this load and capacity balancing that Sinharoy
8 indicates is there, again, means that there's a single centralized controller for
9 all of the caches. If there wasn't, you could not have this load and capacity
10 balancing, and you couldn't have them behaving in this highly set associative
11 way.

12 Going on to slide 16, Patent Owner argues that because Sinharoy's
13 L2 and L3 caches have structurally individual cache controllers, there can't
14 be this single cache controller. But again, this is incorrect. We saw from the
15 petition and looking at the patent itself that these cache controllers don't
16 have to physically be one structure. Instead, they can be distributed portions
17 which act together or you have logic that act together to perform these
18 caching operations. And that's what we --

19 JUDGE ANDERSON: Wait a minute. And so, what's your
20 evidence that the controllers are working together? I think you went over
21 that a little bit before, but I missed it, and I want to hear more about it.

22 MR. HUGHES: Yes, Your Honor. And I'll go back here to slide,
23 for example, on slide 15 on the lefthand side. We do have here that, for
24 example, when you have a given -- the core associated with an L2 and a
25 given L3 region is disabled, the remaining L3 regions enjoy full access to
26 that region. And you can get this overall load and capacity balancing. As

1 Dr. Horst explained, the only way to have this is for them to be working
2 together. If they were entirely independent, and a core and an L2 and L3
3 region were disabled, and everything else is entirely independent of that,
4 there would be no way for those remaining L3 regions to continue to enjoy
5 that full access, for example.

6 JUDGE ANDERSON: So, is this related in any way to snooping
7 between caches or not?

8 MR. HUGHES: It could be, Your Honor, and that's one of the
9 next arguments coming up right here, so I'll go to slide 17 and address your
10 point here. You know, Patent Owner makes some arguments that there
11 might be these specific snooping operations that have to be performed, and
12 that this somehow indicates that there's not a single caching agent in
13 Sinharoy.

14 First, I'd like to note that this is very specific on certain snooping
15 operations, and none of that is recited in the claim. The claim doesn't get
16 down to the level of these very specific snooping operations in certain
17 instances. You know, it doesn't say that this single caching agent has to do
18 all snooping or, you know, no snooping anywhere else. It just says it does
19 the coherence operations for the cores and the IIO module. What Patent
20 Owner's argument is really getting into is a level of detail that's just not
21 required by the claim. But in any event, there are situations where we can
22 see from Sinharoy that even if there are some instances where there may be
23 snooping and other L3 caches, that doesn't mean that we don't have some
24 cooperation or knowledge in the claim single caching agent. We can see this
25 first, that there's some difference in the cache states for the L2 and L3
26 caches, including between the TEN state. Like we just talked about, and

1 very importantly, when you have that core and L2 associated with a given
2 region disabled, the remaining L3 regions still enjoy full access. And even
3 when it talks about this snooped, it says possibly by other L3 caches.

4 So first, I'd like to talk about the difference between these cache
5 states.

6 JUDGE CALVE: Counsel, I apologize. I'm basically asking you
7 to repeat yourself. That last slide if you put that up, please. And I think you
8 had answered Judge Anderson's questions. That second bullet.

9 MR. HUGHES: Yes.

10 JUDGE CALVE: You're saying that the logic, the controller logic,
11 whatever form it's in, is performing the logic for the L2 and the L3 caching
12 agents. Is that -- it seems what the clear import is of that. So, that's what
13 you're saying is evidence that connects them as a single agent?

14 MR. HUGHES: Yes, Your Honor, there's some logic or some
15 overall control that knows what the others are doing. Or what here, for
16 example, knows that this is disabled, but somehow the other L3 regions still
17 have full access to that disabled region. So, there is something that's
18 coordinating all of this together.

19 JUDGE CALVE: What's Patent Owner's position against your
20 contention there?

21 MR. HUGHES: So, this is where Patent Owner brings in their
22 argument that this can't be the case because they're snooping. And I'll go to
23 that and come back to the TEN states. If we go to slide 19, this is what
24 Patent Owner is pointing to on the left-hand side. Storage accesses that miss
25 both the L2 cache and the four megabyte local region are broadcast to the
26 coherence fabric and snooped by the memory controller, other L2 caches,

1 possibly other L3 caches, and by the seven remote four megabyte L3
2 regions.

3 JUDGE CALVE: Can I just interrupt you there? I'm sorry to
4 interrupt. This is complex technology for me. That first phrase or portion of
5 a sentence you wrote, that seems to support your position. It's showing once
6 again that snooping is done on a corporate basis. It's done with collaboration
7 and cooperation between these different caches. Or am I misunderstanding
8 how this aspect of snooping works?

9 MR. HUGHES: This does, Your Honor, go to that overall control
10 and coordination that we were talking about. You're right. This is operating
11 as one system. You're operating for this 32 megabyte shared L3 cache
12 region. And this goes into the aspect that that's shared. You do have this
13 overall control and operation for it.

14 JUDGE CALVE: Okay. Yeah, that's what I didn't understand.
15 Well, we'll ask Patent Owner when they get on, but their response at 17 is
16 citing this snooping and they're saying if something's missed, all the other
17 groups band together and perform some collective or cooperative snooping.
18 That sounds like a single controller, single caching entity functionality
19 within the scope of the claims. Is that -- am I missing something?

20 MR. HUGHES: No, I think that's fair, Your Honor.

21 JUDGE CALVE: Okay, thank you. I'm sorry to interrupt you. I
22 apologize. Just trying to keep up with this. Thank you.

23 MR. HUGHES: That's okay, Your Honor. I did want to mention
24 briefly there's some additional evidence of the sharing and table two has
25 these cache states. And I think this was explained well by Dr. Horst, but we
26 can see a T state is shares and scope. Yes, probably global, whereas a TEN

1 state is yes, local. And we have from Dr. Horst, he explained, you know,
2 this probably global means that when it's in this state and potentially shared
3 access to the local L3 regions would need to snoop those. But this yes, local
4 means that this snooping could in fact be avoided. So this is just some
5 additional evidence. It's in the papers from Dr. Horst that we may not need
6 this snooping.

7 JUDGE PESLAK: Counsel, this is Judge Peslak again. Just one
8 question. Go back on slide 18.

9 MR. HUGHES: Yes, Your Honor.

10 JUDGE PESLAK: We can infer from this slide that probably
11 global and local are not defined in Sinharoy. And so, what we'd be relying
12 on here would be just Dr. Horst's opinion of what those two things mean?

13 MR. HUGHES: That's right, Your Honor. They're not defined in
14 Sinharoy. But I'll give one example. If we look at the description of TEN at
15 the bottom, it says it was formerly ME, now shared. State ME was
16 exclusive. So, what this state is telling us is there used to be something that
17 was exclusive, now it is shared. And how is it shared? Yes, it's shared in
18 our shares in scope column. And how is it shared? Local. Like he says,
19 when we look at that, that changing state from exclusive to now shared, that
20 local means that the targeted cache line is shared, but it's known that no
21 other non-local sharers have modified it because we know that whatever's
22 local, that's what it is.

23 JUDGE ANDERSON: So, I am a little unclear on the state of the
24 caches. And so, look I'm still on, we're all on the same 18, slide 18. So,
25 we've got a list of states in that far left column. And I believe Dr. Horst said
26 at one point that some of these are now described differently. They are now,

1 and I think you used the word counsel, shared, or I think they can be
2 modified and there may be a third state. So, it sounds like this whole thing
3 has changed somewhat. I guess my first question is, is that right? Has it
4 changed from these TNIN, TEN to shared and modified? And there's a third
5 one. I know that, but I can't remember that one.

6 MR. HUGHES: So, I think again, if we can maybe look at the
7 TEN state at the bottom, this may clear it up, where it says formerly ME,
8 now shared, that doesn't mean that what state TEN used to be called state
9 ME. What it's saying, formerly ME, if we look up about midway through,
10 we see state ME is exclusive. So, the description for TEN says this used to
11 be ME, it used to be exclusive, now it's shared. So, Your Honor, the
12 description here is relating to the change in state from ME to something
13 that's shared, not that something that used to be called ME, that same thing is
14 now called TEN. Does that answer your question?

15 JUDGE ANDERSON: It does, but that really is just sort of a
16 preliminary to what I really want to know. What is shared then as it relates
17 to what is being shown here? Or what's testified to by Dr. Horst?

18 MR. HUGHES: Yeah, so shared would be that the other caches
19 could access to it or that it may be in other places as well. It's not -- so you
20 could have others that have the same line, it could be shared. It doesn't have
21 to just be local in this specific also.

22 JUDGE ANDERSON: So, if it's shared again to the question I
23 asked earlier, it's a logical single cache agent because it's shared. Is that
24 simplistic enough?

25 MR. HUGHES: Yes, Your Honor.

26 JUDGE ANDERSON: And more importantly, is it correct?

1 MR. HUGHES: I think that's right, Your Honor. Because these
2 are shared amongst these different regions and we know that they're
3 operating together, that is what is showing us that we do have a single
4 logical caching agent here. These are kind of working together and
5 operating together, not operating entirely independently from each other.

6 JUDGE ANDERSON: Thank you.

7 MR. HUGHES: About how much time? I'm sorry, Your Honor, I
8 just want to check how much time we have. Right there. Okay.

9 JUDGE ANDERSON: I've got you on that one. You have 16
10 minutes remaining of your hour. Now, if Mr. Xu is going to do some
11 presentation, you, of course, have another 15 minutes.

12 MR. HUGHES: Okay. I'd like to turn then, Your Honors, if it's
13 okay, to the second ground, which is Tang, and discuss how Tang also
14 discloses this single caching agent limitation.

15 JUDGE ANDERSON: Yeah, it's your rodeo.

16 MR. HUGHES: If we go on to slide 21, we have an overview of
17 Tang, and this is also to a processor system. The design we'll focus on is
18 this PBDC. That includes a dual processor system, so two processors, and a
19 dedicated DMMA cache.

20 Going on to slide 22, we see Figure 10 of Tang here at the left-
21 hand side of the slide. And remember, this is now two distinct processors.
22 We have the first processor on the left of the certain components, and a
23 second processor making it the dual processor system on the right, and they
24 both include the same components. Just like we saw in the claims, there's a
25 plurality of cores, this I/O bus, and then importantly, we've got the cache
26 controllers here. There's a cache coherence controller in each processor, and

1 that works through global control logic with what's called an LLC controller,
2 and that's in that longer dotted line box in the middle, and that LLC
3 controller includes two prefetchers, I/O data controller, CPU data controller,
4 and that global control logic.

5 Now, the first processor on the left, this processor here, is
6 performing the cache coherence for what we saw in the claim. It's
7 performing cache coherence operations for the cores, and it's performing
8 cache coherence operations for an I/O module. We, in fact, see the I/O data
9 controller and the prefetcher and the I/O bus. Those are cache coherence
10 operations for that input-output module. And we also see what we've got
11 here is the CPU data controller working with the cache coherence controller,
12 and again, that global control logic, which links in the I/O data controller,
13 that's performing the cache coherency operations for these plurality of cores,
14 core zero through core N that we see at the top.

15 JUDGE CALVE: Counsel, if you could go back to that previous
16 slide, please, I guess it's a 20. Yes. Just to help my understanding. What is
17 the purpose of the cache coherence controller? And my question really is,
18 why doesn't the LLC last level controller with its many components, why
19 doesn't that all by itself address the claim language? It works for the IO and
20 it works for the cores. The cache coherence controller beneath that appears
21 to be providing a coherence leak with the other CPU which has its own
22 course and things. Do you really need that cache coherence controller and
23 why?

24 MR. HUGHES: So, I don't think we really necessarily need it for
25 the claim. Like you said, the logic or what's performing these coherence
26 operations, including that global control logic is within the LLC controller.

1 We include it because Tang indicates that this cache coherence controller is
2 something that was left over from the DDC earlier design, and indicated that
3 it does have something to do with some of the coherence operations for the
4 cores. And we can see that with the double line coming out and then going
5 up to the cores there. So, that's the reason we included it here, Your Honor,
6 is that the reference indicated it is involved at least somewhat in some
7 certain of those operations with respect to the cores. But --

8 JUDGE CALVE: Thank you. Yeah, that was very helpful, thank
9 you.

10 MR. HUGHES: So, I'll go on to this. You know, I think going on
11 to slide 24, we can really see it at the bottom left here. These are the
12 components in that LLC controller and it tells us exactly what they're doing.
13 The IO data controller and the CPU data controller, and I will mention right
14 above this, it says that this is all organized by that global control logic, is
15 performing, and maintaining the cache coherence for the IO data ways and
16 the CPU data ways, which would be the core operations and the IIO module
17 operations.

18 So, Patent Owner argues, as we see on slide 25, that first, because
19 Tang included this pair of cache controllers which communicate via this
20 coherence link, it can't have a single caching agent. But as was previously
21 explained, Figure 10 is two distinct processors and the claim relates to a
22 processor claim 1. So, we're really looking within that one processor. And
23 as we saw within that one processor, Tang does in fact have this single
24 caching agent that does the cache coherence operations for both the cores
25 and the controller. The coherence link is really irrelevant to whether that
26 one processor, we'll say the link between the two different processors is

1 irrelevant to whether we have that single caching agent within each
2 processor. And if anything, that just goes to claim 15 where we do in fact
3 have two connected processors, we see that in Figure 10 as well. Now
4 Patent Owner raised this in its response, but it appears to have dropped it in
5 sur-reply. And there, like we see on slide 26, they're now arguing whether
6 that single caching agent has distributed portions. And in fact, it does.

7 If we can go on to slide 28 here, we see that this cache controller
8 and the LLC, there's certain arrows coming out of these, correct? And this is
9 coming from the global control logic. And we see certain of these arrows
10 are going up specifically to Core 0 or Core n. Now, as Dr. Horst explained,
11 what that means is that the LLC controller and the cache coherence
12 controller include signals and logic that are each connected to Core 0 or
13 Core n. The separate arrows to each of the cores show that the shared logic
14 must have signals to individual cores for handling, e.g., L2 cache misses and
15 signals are sent and received by the logic associated with these individual
16 cores. So, what we have is because this cache coherence controller and the
17 LLC cache coherence controller can send these different signals specifically
18 to different cores, there must be distributed portions of logic within them
19 that allow them to do that and direct certain of this control to Core 0 as
20 opposed to Core n.

21 JUDGE PESLAK: Counsel, this is Judge Peslak again. What we
22 seem to be looking at here is like a high-level schematic, and your expert
23 seems to be relying on these arrows to say that there's -- I know it's a one-to-
24 one correspondence to specific cores. Is there text in Tang that confirms
25 that? Or is he just relying on the arrows to say that, you know, there's this
26 correspondence with the cores? And that -- let me just get the language for a

1 second here, I'm sorry. It says, the single caching agent for the processor
2 included plurality of distributed portions, each associated with a
3 corresponding one of the plurality of cores. Where in Tang do we get that
4 correspondence that each distributed portion corresponds to one of the
5 plurality cores? Or are you just relying on Dr. Horst's interpretation of
6 Tang?

7 MR. HUGHES: So, we've got -- it's not just Dr. Horst's
8 interpretation. If you look at section 3.2.2 in the reference, that's the design
9 of this last level cache controller, which also brings in some mentions of the
10 CC controller or the cache coherence controller. And this states in here that
11 global control logic is responsible for routing the read-write requests and the
12 probe-write requests to the right destinations. That's the first thing it says.
13 Those write destinations are either Core 0 or Core N. So, the text is
14 explaining that there's this global control logic that is going to route these
15 coherence operations to make sure they get to the right spot. So, I think that
16 right there supports it and if we read further in that section there it gives
17 some details as to how this global control logic in fact works with all these
18 different components to do like it says routing those requests to the right
19 destinations.

20 JUDGE PESLAK: Okay, I understand.

21 MR. HUGHES: And then in its sur-reply this goes to, you know,
22 Patent Owner now argues that there will not be these distributed portions.
23 They all share this same cache coherence controller, so we can't have a
24 single logic, or sorry single caching agent. But as was explained, we do in
25 fact have this in Tang. There's that the global control logic that works with
26 the IO controller works with the CPU data controller and together make sure

1 that the different signals are sent to the different respective places. So, it has
2 these distributed portions of its logic that connect to those signals, including
3 each one going to the different cores.

4 And Your Honors, I will stop there. I believe we're getting close
5 on time for me. I was going to turn it over to my colleague Kyle. Okay.

6 JUDGE ANDERSON: Yeah, that's fine. That's fine, go ahead.
7 Turn it over and you have -- Mr. Xu has six and a half minutes, plus his 15.

8 MR. HUGHES: Okay, Your Honors, and we may reserve that six
9 and a half for rebuttal and I'd be happy to answer any further questions on
10 what we've discussed already during that time as well.

11 JUDGE ANDERSON: That sounds great. Thank you. Mr.
12 Zhang. Mr. Xu, I'm sorry. I got all confused here. Mr. Xu, am I saying that
13 correctly?

14 MR. XU: Yes, thank you, Your Honor.

15 JUDGE ANDERSON: Okay. Good, thank you. Proceed.

16 MR. XU: Thank you, Scott. Good afternoon, Your Honors. Now,
17 we have slide 31, which relates to dispute four, whether secondary reference
18 Xu was publicly available before December 13th, 2011, which is effective
19 filing date of '895 patent's contended by the Patent Owner. I will note that
20 this dispute only relates to ground 1B, whether Sinharoy in view of
21 secondary reference Xu renders obvious Dependent claims 2 and 10 of the
22 '895 patent, which actually relates to the rate interconnect dependent claims,
23 as Your Honor noted before.

24 Now, if we can turn to slide 32, the reference Xu was publicly
25 available before December 13, the effective filing date based on the
26 evidence in the record. First, Mr. Gordon McPherson, who is Director of

1 Board Governance and Policy Development at IEEE, through his role, he is
2 familiar with the paper dissemination and conference proceeding
3 publications at IEEE. And the Xu paper is actually an IEEE paper presented
4 according to Mr. McPherson's declaration in 2011 NORCHIP, which is a
5 conference, and the date of conference is November 14 to 15, 2011. He
6 declared, quote, copies of conference proceedings were made available no
7 later than the last date of the conference. What this means is a person of
8 ordinary skill in the art, and if they're interested in the subject matter and
9 relevant art of Xu through exercising reasonable diligence, can locate this
10 paper by the last date of the conference, which is November 15, 2011. This
11 predates the '895 patent.

12 Second, on the face of the Xu paper on the bottom of this and
13 reflected on the left hand side of the slide, it has a 2011 IEEE copyright
14 notice.

15 Third, according to the record of the copyright catalog from the
16 copyright office, this is a web capture and reflected in Exhibit 2019 in the
17 record. The conference proceeding 2011 NORCHIP, which is a collection
18 of the papers presented at the conference, including the Xu paper, was
19 published at least by December 22nd. What this means is a POSITA, I'm
20 sorry, a person of ordinary skill in the art and interested in the Xu paper
21 through exercising reasonable diligence would have been able to locate the
22 paper before then.

23 Now, if we can move to slide 33, on the left-hand side, it's an
24 excerpt from paragraph 11 from Mr. MacPherson's declaration, quote, he
25 states, copies of the conference proceedings, which is 2011 NORCHIP were
26 made available no later than the last date of the conference. And on the right

1 side is Exhibit 2019, which is a copyright record of the 2011 NORCHIP,
2 which is the collection of papers of the conference proceedings. Now, if we
3 can --

4 JUDGE CALVE: Counsel, let me ask you a question. Can you go
5 to your preceding slide 33?

6 MR. XU: Yes.

7 JUDGE CALVE: So, copies of the proceeding were available no
8 later than the last day of the conference. Where were they available? Paper
9 copies to people who attended the conference or were they already posted on
10 the IEEE website?

11 MR. XU: That is a great question Your Honor. So, based on Mr.
12 MacPherson's declarations, they would have been available at least to those
13 who attended the conference. It is unclear directly from the declaration itself
14 whether it was posted on IEEE explorer by that date.

15 JUDGE CALVE: Okay. So, public accessibility to the interested
16 public is usually the keystone of these sorts of issues. So what's your case
17 here? Where was it publicly accessible and who are the interested parties?
18 It's just within the confines of this conference that we have something before
19 December 13th that's publicly available?

20 MR. XU: No, Your Honor. So, like, the person of ordinary skill
21 in the art, if they are interested in the multi-core processors and, you know,
22 interested in the conference itself, the paper would have been available to
23 those people before the last day of the conference. So, it's not only confined
24 to the people who attended the conference.

25 JUDGE CALVE: I'm sorry, could you explain that again, please?

26 MR. XU: Yes, Your Honor. So, like, if I can actually go to the

1 next slide. So, like, to clarify on that, so, like, this also ties back to the
2 Patent Owner's argument. The Patent Owner has not put forth any
3 contradictory evidence regarding whether this paper was available or when
4 this paper was available, and they did not seek to depose Mr. MacPherson.
5 So, like Mr. MacPherson's role, his role at IEEE, he's familiar with how
6 these conference papers and proceedings were presented. And he already
7 said that the conference proceeding, including the Xu paper would have
8 been available at least by November 15th, 2011.

9 With regard to the general public, so like it would be someone who
10 are interested in this conference and then if they exercise reasonable
11 diligence, they would have been able to locate this paper by the last day of
12 the conference based on the declaration.

13 JUDGE CALVE: Okay. So, the November 15th, 2011, date, that
14 date precedes December 13th. So again, what was available and where was
15 it accessible as of November 15th? Just paper copies, people's binders who
16 attended the conference, or had they already put it up electronically on a
17 publicly accessible site?

18 MR. XU: So, they would be available to those who attended the
19 conference, and they will be disseminated by the conference proceedings. The
20 declaration did not say whether the conference proceedings were directly
21 published by then or, you know, like where else it will be available, but
22 Your Honor, it is correct that the paper would have been available to those
23 who attended the conference by November 15th, 2011.

24 JUDGE CALVE: Okay. Did you have a case that says that's
25 sufficient for public availability? That's a group of interested people and it
26 was accessible before December 13th. Is there some authority that we

1 would be interested in that would fit these facts?

2 MR. XU: Your Honor, I can take a look at a case and get back to
3 you at the rebuttal.

4 JUDGE CALVE: Okay. And then the question, so the fact that
5 they're available at the Copyright Office by December 22, that seems to
6 postdate the effective date. So, does that help us at all for showing it was
7 available? That seems to be too late, right? Or did I misunderstand
8 something?

9 MR. XU: So, we included that date because the conference
10 proceedings, including that paper, a collective of the papers were published
11 by that date. The other reason why we had that in the record was the Patent
12 Owner appears to make a mistake in their reply and notice that if contended
13 that the effective filing date of the '895 patent is December 31 instead of
14 December 13. And that is why we included, but Your Honor, that is correct.
15 You know, like December 22nd, it is later than December 13, which is on
16 the face of the '895 patent.

17 JUDGE CALVE: Okay. So, we have some copies of this are
18 available by November 15th, but this is in paper form or on USB, some
19 form, it's available to conference attendees. We don't have evidence of a
20 practice of this conference of when they posted all these materials to the
21 IEEE website, so we don't have the IEEE website as a publication format.
22 But we do have the conference and the attendees. Is that what we're with --
23 what we're working with here?

24 MR. XU: Correct, Your Honor.

25 JUDGE CALVE: Okay. Thank you. That was good. You're
26 pretty clear. These are always difficult issues for me these factual

1 underpinnings. Thank you.

2 MR. XU: Thank you, Your Honor.

3 MR. HUGHES: And we'll pass on to Patent Owner for now. I
4 think we'll reserve the rest of our time for rebuttal.

5 JUDGE ANDERSON: Very good. You have six minutes 30
6 seconds. And with that we will turn it over to Patent Owner. Please
7 proceed.

8 MR. FAHMI: Thank you, Your Honor. I wonder if we could
9 share my screen for you. May it please the Board. Again, my name is Tarek
10 Fahmi here representing the Patent Owner. Your Honors --

11 JUDGE ANDERSON: I'm sorry, but before you get going Mr.
12 Fahmi, would you like to reserve some rebuttal time? If so, how much?

13 MR. FAHMI: Oh, thank you, Your Honor. Yes, let me reserve 10
14 minutes, please.

15 JUDGE ANDERSON: Okay, I will let you know when you have
16 five minutes of your opening remaining. Go ahead.

17 MR. FAHMI: Thank you very much. Your Honor, neither of the
18 primary references Sinharoy or Tang teach or suggests the limitation that's
19 been discussed this morning that is the single caching agent that includes the
20 plurality of distributed portions each associated with a corresponding one of
21 a plurality of cores. That limitation appears in each of the independent
22 claims of the patent and is not talked or suggested by either of the primary
23 references.

24 Just briefly, as to Sinharoy, and I'll get into this in more detail
25 later, but during deposition, Dr. Horace, the Petitioner's declarant, was
26 forced to admit during deposition that his alleged support for finding this

1 limitation in Sinharoy, which had to do with the caching states and whatnot
2 of the data, in fact, doesn't appear in Sinharoy. He confirmed that on a cache
3 miss, for example, a given controller would not have any knowledge about
4 the state of the cache line in question. And for that reason, all of the other
5 cache controllers would be forced to snoop that miss in order to determine
6 whether or not the cache line was present in any of their respective caches.
7 The fact that those snoops are required indicates that we're not dealing with
8 a single caching controller that would have knowledge of all the stored
9 cache lines. In fact, these cache controllers operate independently, and so
10 each is required to look to see whether that cache line was available in its
11 own cache.

12 And as to Tang, there's nothing since the filing of a petition that
13 should cause you to change or rethink your original assessment that Tang
14 shows in Figure 10 all of the ways of the last level cache being shared across
15 all of the cores of an individual processor. And so as a result, we don't find
16 the distributed portions of the caching agent, each associated with a
17 corresponding core, as recited in claims. So with that --

18 JUDGE CALVE: Counsel, this is Judge Calve. If I could
19 interrupt, please, because this is very helpful with the snooping issue. On
20 page 17 of your response, you're talking about the snooping, and a person of
21 ordinary skill in the art would understand that when a cache accessed by a
22 particular core misses that core's private L2 cache, and it's allocated portion
23 of the shared L3 cache, then this is the portion I have a question about, that
24 miss is broadcast and must be and is snooped by the cache controllers for
25 each of the cores. Isn't the common meaning of broadcast is it goes out to
26 everybody. So, isn't that evidence of the cache coherence? It's the control of

1 coordinating all of these activities and it's at least coordinating a cache miss,
2 among all the other cache elements. That seems to me to prove the opposite
3 point unless I'm misunderstanding the technology or your contention.

4 MR. FAHMI: Yeah, so this is I think the same quote that's
5 reproduced on our slide number 7, Your Honor. So, what's happening here
6 is you have one of the cores making an access trying to do a read out of its
7 L2 cache. And that L2 cache is individual to whichever core is trying to do
8 the read, and there's a miss. In other words, the information being sought is
9 not present. It also misses in the L3 for that cache. And so, now the L3
10 associated with the core that's trying to make the read broadcast, in other
11 words, puts out on an interconnect of some kind, an indication that, hey, I'm
12 about to do a read to memory. That's what's happening here. I need to get
13 the data. I don't have it in my L3. I'm going to do a read to memory, and
14 I'm informing the other caches that I'm about to take in a copy of this data.
15 That's necessary because the other L3 caches need to understand that another
16 core is about to make an access to data that they may have in their own
17 caches. They may need to change the state of that cache line in their caches.
18 If we were talking about one controller knowing everything about what's
19 contained in all the caches, that wouldn't be necessary. It would already --
20 that one controller would already understand the state of that cache line.

21 JUDGE CALVE: Well, let me interrupt there. Your claims don't
22 require that level. I mean, this language you're quoting says it broadcasts to
23 a coherence fabric. Isn't the coherence fabric evidence that you have a
24 single functional cache, the coherence fabric forms this? This is I think that
25 Petitioner's position is a coherence fabric, and it's going to go ahead and
26 coordinate its actions with all the of these other caching agents in their

1 respective caches. That's how I would understand that.

2 MR. FAHMI: No, Your Honor, you're mistaken. It's evidence of
3 their separate natures. This is what came up. There was discussion about
4 Sinharoy's table number two that showed the cache states. And there was
5 particular discussion about the TEN state. And this relates directly to your
6 question. So, the TEN state, if we look at Sinharoy page 17, table number
7 two, the TEN state relates to a cache line that was formally exclusive, i.e.,
8 formally ME. And it's now being shared. So what this means, and when a
9 cache line is shared, it's not, as was suggested during Mr. Hughes'
10 presentation, that each core is somehow accessing the same line in one of the
11 L2 or L3 caches. That's not what shared means in a cache state like this.
12 What shared means is that copies of a cache line, copies of data, exist in
13 more than one cache. So, they might exist in an L3 cache associated with
14 Core 0, and also existed in an L3 cache associated with Core 1 and so on.
15 That's what shared means in this context. So a line that was --

16 JUDGE PESLAK: Counsel, I have to -- I mean, hold your thought
17 for Judge Calve for a second. But I have a very basic question about this
18 quote. It starts with storage accesses that miss both the L2 cache and the
19 four megabyte local region or broadcast to the coherence fabric. What in
20 Sinharoy is doing the storage access and broadcasting?

21 MR. FAHMI: It's the individual L3 centralized controller. So,
22 we're told on page 19 of Sinharoy that each four megabyte L3 region has its
23 own single centralized controller. So, it's the individual controller of the L3
24 cache that's making this broadcast. And it's the controllers of the other L3
25 region caches that are doing the respective snooping for their L3 regions.

26 JUDGE PESLAK: So, but wouldn't that controller who sends

1 out -- who finds the miss in broadcasts, wouldn't that be the single caching
2 agent?

3 MR. FAHMI: It's one. There is no one caching agent in Sinharoy,
4 Your Honor. Each L3 region and each L2 has its own associated controller.
5 And those controllers are managing the cached data within their own
6 respective regions. So, they're keeping track of what one another are doing
7 in order to properly set the state of the cache line that they are managing in
8 their region to make sure that everything is coherent. That's what the
9 interplay of this state table and the snooping is all telling you, is that these
10 things are acting individually, but in such a way that they maintain
11 knowledge of the information that they are responsible within their own
12 caches.

13 JUDGE ANDERSON: But doesn't that require that each controller
14 has knowledge of the other controller?

15 MR. FAHMI: Not of the other controllers, Your Honor. No, all
16 they need is knowledge of the data which they have in their associated cache
17 regions. That's what this table two is about. So, if a particular data, it's
18 called a cache line, if particular data is stored in a particular region, it has a
19 state associated with it. It's either invalid, meaning that particular data is no
20 longer to be used by the core, or it's shared, meaning I have a copy and
21 another core's cache might also or also has a copy. It's modified. In other
22 words, I've received something from my core that has changed the data. It's
23 exclusive, meaning I'm the only cache that has a copy of this and it's clean.
24 My core hasn't yet modified it.

25 JUDGE ANDERSON: Yeah, let me ask it a little differently.
26 Maybe my question wasn't very good. Do you have any position with

1 respect to what Petitioner really is relying on, as I understand it, is that in
2 Sinharoy, and I believe this also applies to Tang, there was a single logical
3 cache agent. What's your position? Is that not the case? What is a logical
4 cache agent? Why is it not the case that there is not a single logical caching
5 agent with respect to these two references?

6 MR. FAHMI: Yeah. Two things about that, Your Honor. First of
7 all, the claim doesn't say logical. It says one. Second, even if you read it as
8 logical, their position on logical is unclear. I heard Mr. Hughes say, well,
9 that only means that there has to be some kind of logic associated with a
10 core.

11 JUDGE ANDERSON: Well, that is where your -- the patent says
12 that, right? It says logic. It says the whole cache can be logically connected,
13 I believe.

14 MR. FAHMI: Well, it says the cache can be formed from logic,
15 but it doesn't say that the cache controller -- the single cache controller is
16 somehow logical. It says that with respect to Figure 2, which shows the
17 components of a caching agent, that those components can be formed as
18 logic, which I don't --

19 JUDGE ANDERSON: Well, maybe I'm getting us off the track
20 here. If so, I apologize. What I'm trying to understand is why is it that these
21 references do not show a single logical caching agent?

22 MR. FAHMI: Yeah, so in the context of Sinharoy, Your Honor,
23 the caching agents, cache controllers are operating independently. They are
24 not acting as one single controller for the reasons that I've tried to describe.
25 In Tang, it's clear that we have the caching controller that's shown in Figure
26 10 is being shared across all the cores. We don't have a situation where the

1 controller's made up of distributed portions and each portion is associated
2 with an individual core. In fact, it's one controller that is shared across all of
3 the caches. So, it's a different kind of distinction.

4 JUDGE ANDERSON: Well, if we found that, notwithstanding, I
5 get your point about what the claim actually says, but if we found that a
6 single logical caching agent met the claim, why would you argue that is not
7 the case? I know you said it's -- they're independent, but why are the
8 controllers independent? And I think Mr. Hughes already said that if they're
9 independent, then it wouldn't be a single logical caching agent. Why are
10 they independent?

11 MR. FAHMI: They're independent -- in Sinharoy, Your Honor,
12 they're independent because we're told they're independent. They operate in
13 a fashion where each is autonomous. And we know this, for example,
14 because we're told that in a cast out, that is when data is being ejected from a
15 particular L3 region, and this is on page 20 of Sinharoy, there's two options,
16 right? Data could either be ejected to memory or data could be ejected or
17 cast out to another L3 region. And this is shown in the diagram in Figure 10
18 at the top of page 20.

19 But in the event a cast out is being done to a different L3 region,
20 Sinharoy clearly states that only works if the cast out is accepted by the
21 targeted L3 region. This is in the righthand column about the middle of that
22 column on page 20. If it were the case that we had a single logical or
23 otherwise, single cast controller, there wouldn't be a need for the controller
24 of an individual L3 region to accept the cast out or not. That operation
25 wouldn't make sense because the single controller would be managing
26 everything and would have already accounted for that cast out. But that's

1 not what Sinharoy says happens. What Sinharoy says happens is the
2 controller seeking to make the cast out broadcasts this intent, but it's up to
3 the targeted region to accept or not that transfer. And that only happens
4 when you're dealing with independent or autonomous cache controllers.

5 JUDGE CALVE: Well, why does that have to be autonomous?
6 They could have been internal rules of this coherence fabric -- this says
7 coherence fabric. A fabric means a cooperative structure. That alone shows
8 it's single, it's unitary in the understanding in the art. And you seem to be
9 poking at Tang because it's a single, and then you flip it around and you
10 poke at Sinharoy right because they're distributed and you don't like either of
11 those and they both seem to do exactly when you argue one, you argue
12 against the other one.

13 With Sinharoy, it says in Sinharoy, as I think I asked Petitioner's
14 counsel, each L3 region is an eight-way set associative, and this is the same
15 thing for L2. If that isn't singularity of function what else would -- how else
16 should we understand that? And the third issue I'm getting to, even if you're
17 snooping, we've been, I think, wrong on the snooping, but even if we said
18 you were right, to have a single caching agent, you don't have to have the
19 snooping work that way. You're going to have to address what Petitioner is
20 relying on for this single agent, and they're not relying on snooping,
21 although this snooping seems to prove their point, because it says that it's
22 going to a coherence factor, and it's operating together. And it might
23 operate -- it might have a coherence rule that says each individual one has
24 right of first refusal.

25 That still doesn't mean that there's not an overall logic here that's
26 coordinating the operations, although your claims don't get anywhere near

1 that, and I'll just conclude, you can address any of those if you wish. Your
2 own patent just describes this connecting element 260, and all it does is it
3 maintains coherence. It connects, it maintains coherence. So, that's a pretty
4 broad concept. But that's the specification, not even the claims. So, I don't
5 think they need to show much more. What is wrong? They're showing
6 coherence here as a coherence fabric. Why isn't that coherence that
7 underpins as an embodiment of a single agent? That seems to me to be a
8 slam dunk.

9 MR. FAHMI: So, you're basically wrong about everything, Your
10 Honor, and I don't mean to be rude, but it seems like you're not familiar with
11 the way these caches are operating. First of all, set associative has nothing
12 to do with the issues that we're discussing here. Set associative has to do
13 with the way data is organized internal within each region of these caches.
14 So the LP cache, which is private, only associated with one core, is set
15 associative in the way that it organizes data into blocks. That's to allow for
16 accesses from the core. The same is true of the individual L3 region. Set
17 associative when it's used in this context is distinguishing from direct map
18 caching. It has to do with the organization of the data within blocks. It does
19 not have to do with the way in which the controllers are coordinated or not
20 with one another.

21 JUDGE CALVE: So what is the set? I'm sorry. What is the set
22 that's the eight-way set?

23 MR. FAHMI: It's a set of eight blocks in this case, blocks of
24 memory being a particular size. I'm not sure.

25 JUDGE CALVE: Then what is associative? If these are all
26 individual blocks that never talk to each other according to you, what's the

1 meaning of associative? How is this set -- what are they associating with,
2 each other? What?

3 MR. FAHMI: It's associative in the sense of how the data is
4 mapped for memory, memory addresses to cache addresses, as opposed to
5 being direct mapped.

6 JUDGE CALVE: Where does it say that? Is there some evidence
7 of that in Sinharoy?

8 MR. FAHMI: No, Sinharoy doesn't explain set associative versus
9 direct mapped. It's written for an audience that is familiar with these types
10 of terms of art in caching.

11 JUDGE CALVE: Mm-hmm. Well, the only thing that seems to be
12 eight-way associative are these L2 caches with the cores. And then I guess
13 there's an L3 cache that's associated with these cores as well. That's
14 associative that we actually see in Figure 1. And so, why wouldn't that be
15 the set associative? Did you have expert testimony on this, by the way?

16 MR. FAHMI: No, Your Honor. This was never an issue that came
17 up in either parties' case.

18 JUDGE CALVE: Okay. Great. Thank you.

19 MR. FAHMI: Now, you asked about the snooping. The reason we
20 brought up the snooping in our papers, Your Honor, was to explain or help
21 to explain how these cache controllers are autonomous. And I think I
22 covered that, but let me just reiterate that because of the need for each L2
23 and each L3 controller to snoop accesses by other controllers, it shows the
24 autonomous nature. They're not -- there is no single guiding intelligence or
25 single controller that is operating here.

26 JUDGE CALVE: Well, let me interrupt you there. What is

1 broadcasting the miss then? Isn't there a single entity? Or is this peer-to-
2 peer?

3 MR. FAHMI: Peer-to-peer is an interesting way to describe it. I
4 don't think any of the papers, whether ours or Petitioners have described it in
5 that way. But I suppose it is --

6 JUDGE CALVE: So what does the broadcast mean? Who's
7 broadcasting? Doesn't broadcast mean from one point to many? So what's
8 doing the broadcasting?

9 MR. FAHMI: So, each individual L3 controller is broadcasting its
10 storage access to the fabric. The fabric is the way in which these things are
11 interconnected, the coherence fabric. They certainly talk to each other.
12 That's the purpose of the coherence fabric. But that doesn't mean they're
13 acting as a single entity.

14 JUDGE CALVE: So, to answer my question again, what is
15 broadcasting the miss? So, L3 cache number one does this miss? What
16 element broadcast did a miss occur? How does that occur?

17 MR. FAHMI: The L3 region's controller broadcasts the miss. It's
18 an indication that that element --

19 JUDGE CALVE: Let me interrupt. So, does the L3 controller that
20 broadcast the miss, is that the one with each core? Or is that -- that's a
21 centralized controller? Is that what --

22 MR. FAHMI: Your Honor, let's look on page 19 of Sinharoy in
23 the lefthand column under the heading L3 cache. The third sentence, a given
24 L3 region comprises a single centralized controller, et cetera. It's telling you
25 that each L3 region of which there are eight has its own controller. It's that
26 controller that broadcasts the miss for a particular L3 region.

1 JUDGE CALVE: Sure. It's like Bluetooth, one to all, I
2 understand. So, it's still broadcasting over some sort of interconnected ring,
3 right? It's able to broadcast to these other L3 controllers. So, isn't that a
4 coherence fabric? Isn't that a higher consciousness, a connectivity that
5 makes it work as a single entity? It's the best of both worlds. It's a synergy.
6 Each one operates individually for its circumstances that it can pull up a
7 level and work corporately. I think that's within the scope of your claims.

8 MR. FAHMI: That's not the way it's -- that's not correct, Your
9 Honor. The coherence fabric is simply a -- it's a set of wires. It's an
10 interconnect.

11 JUDGE CALVE: Well, isn't that what you're showing in your
12 Figure 2 and in your Figure 6, sort of a physical interconnect that connects
13 your various singular caching agents, the element 260 and Figure 2 of the
14 patent?

15 MR. FAHMI: Figure 2 of the patent is showing a ring interconnect
16 260, yes.

17 JUDGE CALVE: Okay. And isn't that what enables some of the
18 coherency actions within the scope of the claims? I mean there may be logic
19 behind it but that connectivity is evidence that there's -- it's a highway for
20 singularity of action of coherence right, and that's also what's shown in
21 Figure 6. And so, if we have evidence of some coherence fabric, some sort
22 of structural element that tends to suggest to me at least a unitary agent that's
23 all. That that's my question.

24 MR. FAHMI: I think you're mistaken, Your Honor. I think that
25 what the ring interconnect is just that. It's a set of wires that are connecting
26 different controllers.

1 JUDGE CALVE: And what's the purpose of that?

2 MR. FAHMI: To allow signaling between these different
3 controllers.

4 JUDGE CALVE: And isn't that how coherence then would be
5 handled in your patent? Between that and the IO, it connects all these
6 different elements together so they can coordinate their operations. And I
7 think that's what a single action agent is encompassing in the claims. Or is
8 that not correct?

9 MR. FAHMI: I think you're trying to imbue the interconnect with
10 some kind of intelligence. And I don't think there's anything in the patent
11 that indicates that's the case.

12 JUDGE CALVE: No, no, I think, no, maybe I said that perhaps
13 the interconnect is evidence that there is an intelligence and it's working
14 through that highway. So, if we see evidence that there's connectivity in
15 Sinharoy, for example, there's a broadcast element. I'm saying the broadcast
16 element is analogous to your 260 element. It's a communication medium
17 through which the singularity of action occurs. Otherwise, it wouldn't be
18 there. So, we're just we're saying that that seems to be in the parlance of the
19 claim language.

20 MR. FAHMI: Well, I don't think we've ever disputed that
21 somehow there's interconnection in Sinharoy, right?

22 JUDGE CALVE: Okay. You've made it clear. We're at a
23 misunderstanding here, and I'll take a look at the record. You've got your
24 points clear, but we'll take a look at the transcript. I don't want to consume
25 more time if you have some other points you'd like to make. I understand
26 your position fully. Thank you.

1 MR. FAHMI: Well, Your Honor, I don't want to just leave it
2 because it seems like we're not on the same page. But perhaps let me say
3 some other things about Sinharoy, and perhaps it will shed some light on this
4 point.

5 So, on slide number 9, I think I made the point that the fact that
6 each L3 region has to decide whether or not to accept or reject a cache line
7 that another L3 controller is trying to eject to it, indicates their autonomous
8 nature. In other words, there's no single cache controller that's operating.

9 On slide number 10, excuse me, before we get to Tang. There are
10 other things that the Petitioner has tried to say about Sinharoy
11 (CROSSTALK) --

12 JUDGE CALVE: Well, let -- I think this is the confusion here, Mr.
13 Fahmi. Let me ask you, we keep going back and forth like a toggle here.
14 The claim says the caching agent is a single agent, but it has a plurality of
15 distributed portions. So, you've not put anything in the record, really, about
16 what the single caching agent is. And I understand the burden is not yours
17 to explain anything. You can sit. But your description of the patent and
18 your response talked about Figure 2. So, that's where I started with the most
19 of my review. The single caching agent has a plurality of distributed
20 portions. Isn't that what Sinharoy is?

21 So, if we wrote an opinion that said Sinharoy has these individual
22 controllers, which you agree they're down at each floor with an L2 cache.
23 And if there's a miss, they act as a single caching agent because they
24 broadcast misses as a caching fabric and that's a single caching agent. What
25 are you going to show in your claim language or your spec that makes that
26 wrong?

1 MR. FAHMI: There's nothing in Sinharoy that says they act as a
2 single agent, Your Honor. I think that's where our fundamental
3 disagreement is. Each caching agent is deciding for itself. There's nothing
4 indicating in Sinharoy that all of these individual agents are acting as one.
5 That is where I think the dispute lies. Our contention would be that the
6 Petitioner hasn't shown how they are acting as one.

7 JUDGE CALVE: Well, isn't that acting as one? Your very
8 snooping example shows they act as one. They broadcast a miss and then
9 they act as a group, however they choose, to solve that miss. That's a single
10 caching agent. What does the claimant argue on that?

11 MR. FAHMI: But they don't act as one, Your Honor. I think that's
12 our point. Each one operates individually. Each agent, each individual L2
13 agent, each individual L3 agent has to examine that miss.

14 JUDGE CALVE: What is wrong with that? They have to examine
15 the miss to what, to see if they have a similar miss to see how their state of
16 registers compares? I mean, why is that -- why is that not a good corporate
17 way to handle a miss?

18 MR. FAHMI: Well, you wouldn't need to do it if you had a single
19 controller. If you had a single controller, then that single controller would
20 already know the state of the cache line, regardless of whether it was stored.

21 JUDGE CALVE: What you're positing something that maybe in
22 your world that's more efficient and maybe Sinharoy picked the way that
23 was less efficient. So, in other words, you're saying there's a miss and
24 there's a quarterback up there that goes down and gets the miss and it
25 broadcasts it. Sinharoy says we don't want to do it that way. When there's a
26 miss we're going to have each of you guys check it. Maybe that's too much

1 bandwidth, too much bus line, but that's how they've chosen to do that. And
2 that's their coherence fabric and that's their single caching agent. And that
3 meets the scope of the claims, doesn't it?

4 MR. FAHMI: It's not a single caching agent at that point, Your
5 Honor. It's a number of individual caching agents. And it's not just me
6 saying it. It's Dr. Brogioli only saying it. And he's a guy who used to build
7 these things, these caching protocols.

8 So, it's not correct that, oh, it's just the way for a single caching
9 agent to operate. It's not a single caching agent. And everything else in
10 Sinharoy is telling you it's not a single caching agent. Sir --

11 JUDGE CALVE: Let me ask you have one question here. This
12 single caching agent for the processor. Are you talking about the processor
13 is the CPU and the CPU has cores, right? So, you have one agent for a CPU
14 and then you have distributed portions with each core. Do we at least agree
15 on that scope of the claim?

16 MR. FAHMI: Yeah, I think that's right, Your Honor.

17 JUDGE CALVE: Okay, that shouldn't be too controversial. So,
18 Sinharoy has the distributed portions with each of the cores. Do you agree
19 with that?

20 MR. FAHMI: Yes, I'm still -- yeah. I'm sorry. Did you send the
21 claim or Sinharoy, Your Honor?

22 JUDGE CALVE: Then Sinharoy has distributed portions with
23 each core, like the L2 portions are distributed. And I think he describes the
24 L3 portions.

25 MR. FAHMI: No, no, no. I wouldn't agree with that at all, Your
26 Honor.

1 JUDGE CALVE: Okay. Why is that?

2 MR. FAHMI: Sinharoy says that each L2 has its own controller.
3 It doesn't say that they're distributed parts of a single controller.

4 JUDGE PESLAK: So counsel, this is Judge Peslak --

5 MR. FAHMI: The answer is L2 cache. There are eight L2 caches,
6 and each of those eight has its own controller.

7 JUDGE PESLAK: So, Counsel, are you saying that in the thing
8 that there's a, for lack of a better term, there's something missing in
9 Sinharoy, namely some controller who would control each of the individual
10 cache controllers. That's what's missing in Sinharoy?

11 MR. FAHMI: We don't see that in Sinharoy, Your Honor.
12 Correct.

13 JUDGE PESLAK: So, when -- I think what Petitioner said was
14 that there are centralized controllers for the L2 and L3 caches which form a
15 single caching agent performing task coherency operations. And he was
16 referring to pages 18, 19 in Sinharoy. So, that's the part that the Patent
17 Owner is disputing that these centralized controllers exist in Sinharoy. Is
18 that a fair statement?

19 MR. FAHMI: Well, we don't dispute there are controllers,
20 centralized controllers for each individual L2 and each individual L3. What
21 we dispute is that there is a single controller for the entire processor.

22 JUDGE PESLAK: I understand what you're saying. But okay.

23 MR. FAHMI: Let me go on and just talk about Tang briefly, Your
24 Honors, so that I make sure to get it in as part of our presentation. Nothing
25 has changed in our view since your institution order with respect to any
26 evidence regarding Tang. In your institution order, you agreed that Tang did

1 not show the single caching agent that includes the plurality of distributed
2 portions that are each associated with a corresponding one of a plurality of
3 cores. And nothing, as I say, has changed since that time.

4 Dr. Horst, in his testimony, confirmed that in the last level cache
5 shown in Tang, this is on our slide 10, it's also Figure 10 from Tang. Dr.
6 Horst confirmed that the last level cache in Tang is shared by all the cores of
7 the processors, and so, too, are all of the components of the controller, the
8 CPU data controller, the pre-fetcher, the global control logic, the cache
9 coherence controller, all of those are also shared by all of the cores of the
10 processor. This is in Dr. Horst's testimony. The transcript of his deposition,
11 reply deposition at page 47, lines 5 through 22. We don't see a situation in
12 Tang where there's this correspondence required by the claim between the
13 distributed portions of the caching agent and each of the individual cores.
14 That simply doesn't exist.

15 Finally, as to the Xu reference that we have filed our motion to
16 exclude with respect to the Exhibit 1019, the basis for excluding that exhibit,
17 Your Honor, is that it hasn't been properly authenticated. None of the
18 hallmarks associated with the printout of a webpage that might be self-
19 authenticating, but none of those hallmarks are present with respect to the
20 exhibit that the Petitioner has provided. For example, there's no web
21 addresses, there's no dates indicating when the printing is done, et cetera.
22 There's simply nothing that would indicate -- there's not even a declaration
23 from counsel who, you know, downloaded, and printed the pages to where it
24 came from. So, there's nothing that indicates the authenticity of the exhibit.

25 And as to the Xu reference itself, we've argued that the Petitioner
26 hasn't shown that it's properly prior art. The declaration they rely on for Mr.

1 McPherson is deficient in a number of respects, and Mr. McPherson testifies
2 as to his current position, not the position he would have had at the time of
3 the purported publication of the Xu reference for him to have personal
4 knowledge of when it was published. Mr. McPherson doesn't state that he
5 was actually at the conference to know what was or wasn't distributed at the
6 conference. And, in fact, Mr. McPherson said he obtained a copy of the
7 reference from the IEEE archive website, and it's undisputed that the
8 reference itself was not posted to that website until after the critical date.

9 So, we don't think there's been sufficient showing that the Xu
10 reference is publicly available to make it prior art to the '895 patent.

11 JUDGE ANDERSON: You don't really dispute that the exhibit is,
12 in fact, what it says it is, that it's authentic, or do you?

13 MR. FAHMI: Which exhibit are you referring to, Your Honor?

14 JUDGE ANDERSON: The one you're trying to exclude, 1019.

15 MR. FAHMI: It hasn't been authenticated, Your Honor.

16 JUDGE ANDERSON: Well --

17 MR. FAHMI: We don't know what it is.

18 JUDGE ANDERSON: Well, Gordon McPherson says it's IEEE.
19 What does he say? He says what it is. He says it's an IEEE.

20 MR. FAHMI: I'm sorry, Your Honor, to interrupt. His declaration
21 refers to the IEEE archive page, not the purported Copyright Office
22 webpage. It's the copyright office webpage that we filed a motion with
23 respect to.

24 JUDGE ANDERSON: Yeah, I still don't quite understand why it
25 isn't what it says it is. Because that's not what authentication is. Usually in a
26 trial, everybody agrees, yeah, it is what it is. You're saying it's not what it

1 says it is. Why?

2 MR. FAHMI: Your Honor, it's not our burden to say what it is.
3 All we're saying is Petitioner hasn't shown that it is what the Petitioner says
4 it is.

5 JUDGE CALVE: What is wrong with the declaration, Counsel,
6 where his declaration attached that exact exhibit to it and said this is it? Isn't
7 that authentication? We're not talking about apples and oranges. He said
8 Exhibit A is what I'm authenticating. I think that's authentication, isn't it?
9 There's no confusion or ambiguity there.

10 MR. FAHMI: Your Honor, the declaration doesn't relate to this
11 exhibit. The declaration related to a different exhibit.

12 JUDGE CALVE: Okay.

13 MR. FAHMI: There is no authenticating declaration with respect
14 to 1019.

15 JUDGE ANDERSON: Well, maybe I'm confused, but it attaches
16 as Exhibit A, "Explorations of Optimal Core and Cache Placements for Chip
17 Multiprocessor" by, among others, Thomas, I can't pronounce the name, but
18 it's Xu at the end. Isn't that Exhibit 1019?

19 MR. FAHMI: No, it's not, Your Honor.

20 JUDGE ANDERSON: You're right. So, what is the problem
21 here? Exhibit 1000, the Xu is -- I thought Xu was the reference that
22 Petitioner is relying on. So --

23 MR. FAHMI: Xu is what they're relying on. Our contention is
24 they haven't shown Xu is prior art because they haven't properly indicated
25 the date of its public availability.

26 JUDGE CALVE: Well, isn't the appendix to the declaration the

1 same reference as Exhibit 1005, which is relied on in the petition?

2 MR. FAHMI: I believe it is, Your Honor. Yes.

3 JUDGE ANDERSON: It's Xu, right?

4 MR. FAHMI: Correct.

5 JUDGE ANDERSON: I'm confused. Maybe it's me.

6 MR. FAHMI: Let me see if I can clarify, Your Honor, that there
7 are two issues. Issue one, which we raised in our response, is Xu is not prior
8 art. The evidence that Petitioner has brought forward to try to establish Xu
9 as prior art concerns an alleged publication date, and a publication date that's
10 alleged is at least November 15, 2011, and they rely upon the McPherson
11 Declaration to establish that date. The problem with the McPherson
12 Declaration is it's not based on personal knowledge. There's nothing that
13 indicates he was at the conference. There's nothing that indicates he was the
14 custodian of records at the time that the alleged publication in November of
15 2011 would have occurred. So, those are the two issues with respect to the
16 reference itself.

17 Now, in the reply, the Petitioner brought another exhibit in, Exhibit
18 1019, which they allege is a printout of copyright office records showing,
19 again, this November 2011 publication date. The copyright office record has
20 not been properly authenticated.

21 JUDGE ANDERSON: So, you don't have a problem with Xu.
22 Your problem is that the Copyright Office effort to show public availability
23 is not authenticated.

24 MR. FAHMI: It's -- correct. The Copyright Office record is not
25 authenticated. Correct.

26 JUDGE ANDERSON: Okay. What about the position that it's a

1 government record?

2 MR. FAHMI: Well --

3 JUDGE ANDERSON: I guess that just goes to hearsay would be
4 your position.

5 MR. FAHMI: Yeah, in cases where Copyright Office printouts
6 have been deemed to be self-authenticating, those records have had indicia
7 that would suggest they are reliable. Like, for example, a copy of the web
8 address from which they were printed and the date of printing. But none of
9 those indicia are present here. And there's not even a declaration from
10 counsel saying when this web page was obtained or where it was obtained
11 from. You don't have any of those hallmarks that would lend credibility that
12 this is an actual printout of the web page.

13 JUDGE ANDERSON: Okay. I kind of let you go beyond. So, if
14 you need to have an additional couple of minutes, I would allow it, but
15 you're at the end of your one hour, Mr. Fahmi.

16 MR. FAHMI: Thank you.

17 JUDGE CALVE: If I could ask one question, Judge Anderson.
18 Counsel, you had some great arguments and great points to make, and I just
19 want to give you a chance to comment on. So in the reply, at page 8 of the
20 reply, if you can pull that up, please, I just want to give you a chance to
21 comment. This goes on this single caching agent issue we've been debating
22 back and forth. And I think this is just Sinharoy. But as you know, it's
23 Petitioner's position that even though they have these individual controllers
24 they nonetheless coordinate with each other and act as a single controller
25 and agent. Okay. And you can read that. What is your rebuttal response?
26 What is the error in their contentions here that you'd like us to be aware of?

1 If you have something you'd like to put on the record for that.

2 For example, a person -- a POSA would understand cache
3 controller for L3 region has knowledge of what the other lines in the L3
4 region are held in L2 and L3. Anything, and they specifically address
5 snooping, but they're talking about the larger context. What's your position
6 that differs from that, that we should be aware of?

7 MR. FAHMI: Yeah, I think what we said in our sur-reply with
8 respect to that, Your Honor, is that this point is irrelevant. The state of a
9 cache line doesn't give you any indication of single controller or not. It's
10 simply telling you that the L3 controller associated with a given core either
11 has a copy of that cache line or not. And if it has a copy of the cache line,
12 there's an indication of whether that same cache line is present in other L3
13 caches, or in the case of an L2 in other parts of the caches.

14 JUDGE CALVE: Okay. So, you're saying that's the evidence that
15 they're citing as evidence of coordination between all these individual
16 controllers. Is that your understanding? You're saying that is not evidence
17 of a single caching agent?

18 MR. FAHMI: Correct.

19 JUDGE CALVE: Okay. Thank you. If there's anything else you
20 wanted to add, that's what I wanted to clarify.

21 MR. FAHMI: No, I think that's the point, Your Honor.

22 JUDGE ANDERSON: Mr. Fahmi, you have 10 minutes for
23 rebuttal time left if you want to go into it to finish anything, you're welcome
24 to do that.

25 MR. FAHMI: Thank you, Your Honor. We'll reserve that for our
26 sur-rebuttal.

1 JUDGE ANDERSON: Very good. Petitioner, you are up, and you
2 have six minutes remaining.

3 MR. HUGHES: Okay. Thank you, Your Honor. And as we
4 indicated, we did go back and look at Sinharoy. Unfortunately, there is not a
5 figure in there that specifically shows the L2 and L3 controllers in it. But
6 like I said, and as we set forth in the petition, we think it's clear from the
7 reference and those specific descriptions of those controllers, kind of what
8 they are and how they operate. So, just to let Your Honor know, we did go
9 and look at that.

10 And then with that, I'm going to give the rest of the time to Kyle to
11 talk about the Xu reference.

12 JUDGE PESLAK: I think just before you run off, --

13 MR. HUGHES: Yes.

14 JUDGE PESLAK: -- it seemed to me that the crux of what Patent
15 Owner is saying is you need to have something besides the controllers for
16 the L2 and L3 caches. You need an overall centralized controller I'm
17 paraphrasing what he said. But I believe that's Patent Owner's position.
18 What do you say to that?

19 MR. HUGHES: So I would say, one, that's not what's in the '895
20 patent itself. It's not what's in the claim. The claim even says that this single
21 caching agent could be made of distributed portions. And if we look at the
22 figures, including Figure 6 of the '895 patent, even there, we're not seeing
23 this alleged one overall controller. Instead, what we have is like we have in
24 the references, there's these individual portions or individual controllers that
25 are just shown as being somehow connected. So, the fact that there must be
26 some higher level or overall controller, one, it is not something that's in the

1 claim. And we think even with the description in the '895 patent of what this
2 single caching agent can be, you know, it doesn't there require this overall
3 higher level single controller otherwise.

4 JUDGE PESLAK: Okay, I understand.

5 JUDGE CALVE: Okay. Counsel, let me -- that got me thinking.
6 Judge Peslak always has great ideas here. So, it's a similar response to their
7 attack on Tang. So in Tang, they take the opposite attack and they say Tang,
8 you have this LLC, which is a single one, but you don't have any distributed
9 ones on the Core 0 Core n. So ergo, you don't have distributed portions, you
10 got a single one. What's your response to that?

11 MR. HUGHES: So, I think like you pointed out, and I think this
12 goes to show that the claim and how single caching agent is used is very
13 broad. They are arguing one way, you have multiple things, it can't be the
14 single. Here you have one thing, it can't be the single. But like I talked
15 about in that section of the Tang reference, the first thing it says is that there
16 is that global control logic that is able to direct things to specific recipients.
17 So in Tang, we do have this logic and it's like the patent says it can be logic.
18 In Tang, we very clearly do have this logic that is going to from those
19 controllers, direct certain things to different cores or different destinations.
20 So, if there's --

21 JUDGE PESLAK: I think we've lost Mr. Fahmi here. I don't see
22 his box on the screen here.

23 JUDGE CALVE: Well, hopefully they can pull him in on the dial-
24 in line if he's able to do that, communicate.

25 MR. HUGHES: Would Your Honors like me to wait and answer
26 that question until we have counsel on the line so everyone's here?

1 JUDGE PESLAK: Right, I think we've got to figure this out,
2 sorry.

3 MR. HUGHES: No problem, I can wait.

4 JUDGE CALVE: I was just speaking out loud to Charle. Charle,
5 you don't have to respond to me, but I assume you're checking the lines and
6 you can reconnect and pull him back in.

7 CHARLE: Yes, I have a phone number that disconnected.

8 JUDGE CALVE: Okay.

9 MR. FAHMI: Yes, this is Tarek Fahmi, Patent Owner's counsel.

10 JUDGE CALVE: Are you on the dial-in line, Mr. Fahmi? That's
11 good if you are. Can you hear us okay?

12 MR. FAHMI: Yes, I can, Your Honor. I just dialed in. We lost
13 audio and we're not seeing anything other than the individual's names on the
14 video.

15 JUDGE CALVE: Okay, can you hear us? Are you able to proceed
16 satisfactorily in this fashion? Are you able to understand what the
17 discussions are?

18 MR. FAHMI: Yes, absolutely.

19 JUDGE CALVE: Okay, we're going to have to back up. Counsel
20 for Petitioner, I asked him the follow-on question for Tang. Whereas you
21 were concerned with Sinharoy because they were distributed, but no central,
22 you're concerned with Tang because there's a central LLC, but no
23 distributed. So, I threw that out to Petitioner and we'll hear, he will repeat
24 his answer for you to hear why Tang has a single caching agent with this
25 point distributed portions with each core.

26 MR. HUGHES: Yes, and Your Honor, like I said, what we have in

1 Tang is we have, you know, like you said, they're arguing well, Sinharoy,
2 you've got multiple that can't be single, but then Tang you have single that
3 can't be allegedly distributed portions or multiple.

4 Now, again I think that goes to the breadth of the claim including
5 how broad the '895 patent describes what its single caching agent can be.
6 But with respect to Tang we do have that global control logic which is this
7 overall higher level control logic for all of these caching operations that are
8 going on. And in particular, Tang indicates that that's responsible for
9 routing these requests and to the right destinations. So, what Tang is telling
10 us there is its global control logic has these portions that are telling it okay,
11 for this, I can route it to specific places. Those are the partitioning, the
12 slicing, the logic, that's letting it coordinate with a specific core and specific
13 LLC from another core.

14 JUDGE CALVE: And that's an example of what we would call
15 coherence. And that coherence is an example that falls within single caching
16 agent. Is that your position?

17 MR. HUGHES: That's correct, Your Honor.

18 JUDGE CALVE: Okay, thank you. Oh, I was going to ask you
19 something here. I know, I don't want to go through the looking glass, where
20 if you use a word, it can mean whatever you choose it to mean, which we
21 have called lexicography. Is there any lexicographic definition or any
22 disavow, or anything in the prosecution history that narrows the scope of
23 single caching agent beyond what we've been discussing today that you're
24 aware of?

25 MR. HUGHES: Not that I'm aware of, Your Honor. And I think,
26 like I say, if you look at just the different examples in the '895 patent, they're

1 not limiting. They further, every time it talks about it, it could additionally
2 be this. It could also include this. It can be logic. So, there's not one
3 singular definition, but what we do have is lots of examples that just broaden
4 out what it can be. It doesn't have to be single. It can be distributed, but
5 working together. It can be logic. Those are just some examples, Your
6 Honor.

7 JUDGE CALVE: Okay. You heard over my colloquy with Mr.
8 Fahmi. Do you have anything in the record or thoughts on coherence fabric
9 and the eight-way set associative? He disagreed with that. Do you have a
10 different position on that or is there anything in the record to guide us?

11 MR. HUGHES: Yeah, I think if you go and we look at Dr. Horst's
12 declaration on that, where it's talking about those things, that's where Dr.
13 Horst explains what's going on there is this is all of these components
14 working together, sharing their knowledge, that's performing this overall
15 coherence for the chip.

16 JUDGE CALVE: All right. Thank you.

17 MR. HUGHES: Thank you.

18 JUDGE ANDERSON: Mr. Xu, you have, I'm going to give six
19 minutes. If you feel like you need more, let me know. I've kind of lost a
20 little bit of track of time here, but go ahead. I'll give you six minutes to start.
21 And if you aren't done, we'll go ahead and give you a little bit more. Go
22 ahead.

23 MR. XU: Thank you, Your Honor. And Judge Calve, I have the
24 cases for why the secondary reference Xu was publicly available. So first, a
25 *Medtronic v. Barry* is a Federal Circuit case in 2018. To determine whether
26 a paper presented at a conference qualifies as prior art when the conference

1 date predates the priority date, the Federal Circuit considers a few factors.
2 First, the size and the nature of the conference. Second, whether they were
3 open to the people who are interested in the subject matter. And third, a
4 little bit more importantly, whether this conference was confidential or
5 there's no confidentiality requirement attached.

6 And then in the District Court case, in Eastern District of Virginia,
7 *Netscape Communications v. Valueclick*. The District Court of Virginia
8 found that a paper -- in addition to the paper author's deposition testimony,
9 the paper public was presented at an international conference, which is an
10 international conference on World Wide Web in 1994, May. The District
11 Court concluded that that date was the public accessibility date for that paper
12 and credited that to be prior art.

13 Similar here, the Xu paper was presented at 2011 NORCHIP,
14 which is, first, IEEE is a well-known organization related to processors and
15 computer systems. And second, this is an international conference. I
16 believe in Mr. McPherson's declaration happened in Lund, Sweden. So, and
17 also like Patent Owner has not put forth any evidence that there is any
18 confidentiality requirement attached to this conference. So, based on the
19 case law so far, a person of ordinary skill in the art and interest in the subject
20 matter would have attended this conference and this paper and the
21 conference proceedings would be public to whoever is interested in this
22 conference. And the last date of the conference is November 15, 2011. This
23 predates the priority date of the '895 patent.

24 And just in addition to that, the Board has found in the post-grant
25 proceeding that a poster posted at an international medical conference, its
26 conference date is sufficient to be the public accessibility date to be the prior

1 art.

2 JUDGE CALVE: Right, right. You might take a look, not
3 immediately, but *Nobel Biocare v. Instradent*, it's a 2018 case, talks about a
4 trade show catalog publicly accessed before the critical date was a printed
5 publication(INDISCERNIBLE). And then probably more relevant, *Quest v.*
6 *Hirshfeld*, 2021, abstract on a CD was distributed to specific people most
7 interested in subject and was tech searchable to find the article. So, those
8 are some more I was kind of alluding to. They get in --they drill down to the
9 level of a trade show functionality without having to be in a library for the
10 whole world. And that might be also pertinent to your facts here if you can
11 show that you have evidence that actually was made available at that
12 conference.

13 MR. XU: Thank you, Your Honor.

14 JUDGE ANDERSON: Thank you, Mr. Xu. Anything else Mr.
15 Hughes?

16 MR. HUGHES: That's it. I think I'll just add, you know, we do
17 have the declaration from McPherson indicating the public availability. But
18 Patent Owner has nothing else unless there's any further comments from the
19 Judges.

20 JUDGE ANDERSON: Oh, we will turn to Mr. Fahmi. Mr. Fahmi
21 go ahead and you have 10 minutes to -- for sur-rebuttal. So, go ahead.

22 MR. FAHMI: Thank you, Your Honors. I think the comments
23 from Mr. Hughes really pertain to Tang. And again, with respect to Tang,
24 we have Dr. Horst's admission that all of the components of the last-level
25 cache and the controller associated with the last-level cache are shared by
26 each of the cores. And so, what's missing then from Tang is the requirement

1 that the caching agent have distributed portions each associated with a
2 corresponding one of the plurality of cores, it's that individual association of
3 a distributed portion with a corresponding core that is not found in Tang,
4 again, by their own expert's admission. And so, we don't see any reason for
5 you to deviate from your conclusion in that regard that was set forth in the
6 institutional order.

7 With respect to the Xu reference, there's no evidence or otherwise
8 that that paper was actually presented at the conference. Certainly, Mr.
9 McPherson in his declaration doesn't allude to that and there's no other
10 evidence in the record regarding a presentation of the paper. In fact, what
11 Mr. McPherson says is he obtained a copy from the IEEE Explorer website
12 and it's uncontested that the paper was not made publicly available via that
13 website until after the critical date of the subject patent.

14 Unless there's other questions, Your Honor, I think we'll leave it at
15 that and rest on the arguments today and those presented in our papers.

16 JUDGE ANDERSON: Okay, thank you, Mr. Fahmi. Well, that
17 concludes the hearing today. Let me just say that I found all of counsel and
18 their presentations to be very good. There was good interaction, and I
19 always appreciate that. I think that helps us get to the ultimate decision,
20 which will be coming in due course.

21 With that, we are -- oh, I will say one thing. For the purposes of
22 helping out the court reporter, would counsel stay on the video here, and I
23 know you're not videoing here, Mr. Fahmi, but hang in and answer any
24 questions the court reporter has.

25 And with that, we are adjourned. The final written decision will be
26 forthcoming within the deadline. Thank you all very much and have a good

1 afternoon.

2 (Whereupon the above-entitled matter went off the record at 3:12

3 p.m.)

IPR2023-01344
Patent 9,575,895 B2

PETITIONER:

Celine Crowson
Joe Raffetto
Scott Hughes
Nicholas Rotz
Ryan Stephenson
Helen Trac
HOGAN LOVELLS US LLP
celine.crowson@hoganlovells.com
joseph.raffetto@hoganlovells.com
scott.hughes@hoganlovells.com
nicholas.rotz@hoganlovells.com
ryan.stephenson@hoganlovells.com
helen.trac@hoganlovells.com

PATENT OWNER:

Tarek Fahmi
Jonathan Tsao
ASCENDA LAW GROUP, PC
tarek.fahmi@ascendalaw.com
jonathan.tsao@ascendalaw.com