No. 18-

IN THE

Supreme Court of the United States

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Petitioner,

v.

APPLE INC.,

Respondent.

ON PETITION FOR A WRIT OF CERTIORARI
TO THE UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

PETITION FOR A WRIT OF CERTIORARI

MORGAN CHU
Counsel of Record
GARY N. FRISCHLING
JASON SHEASBY
ALAN J. HEINRICH
CHRISTOPHER ABERNETHY
AMY E. PROCTOR
IRELL & MANELLA LLP
1800 Avenue of the Stars, Suite 900
Los Angeles, CA 90067
(310) 277-1010
mchu@irell.com

Counsel for Petitioner

June 5, 2019

QUESTIONS PRESENTED

During the patent infringement jury trial in this case, the district court found defendant-respondent Apple, Inc. ("Apple") had "waived" any construction for a key claim limitation. The court thus instructed the jury to give the limitation its "plain and ordinary meaning as viewed from the perspective of a person of ordinary skill in the art [POSITA]." The jury heard expert testimony on that meaning in the relevant technical context. The jury then found for plaintiff-petitioner Wisconsin Alumni Research Foundation ("WARF") on infringement.

But on appeal, the Federal Circuit construed the limitation for the first time, while disregarding the jury's fact-finding role. Instead of assessing whether substantial evidence supported the jury's implicit fact findings regarding the limitation's ordinary meaning to a POSITA, the Federal Circuit construed the limitation *de novo*, stating "our view" of its meaning. Then, instead of remanding for further proceedings, the Federal Circuit applied its *new* claim construction to the *existing* trial record to grant JMOL of non-infringement.

The questions presented are as follows:

1. Where the district court properly instructed the jury to give a claim limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]," does *Teva Pharmaceuticals USA, Inc. v. Sandoz, Inc.*, 135 S.Ct. 831 (2015), allow the Federal Circuit to construe that limitation *de novo*, giving no deference to the jury's implicit fact findings regarding a POSITA's understanding?

2. May the Federal Circuit apply a new claim construction issued on appeal to grant JMOL based on a trial record developed without that construction, instead of remanding for proceedings consistent with the new claim construction?

PARTIES TO THE PROCEEDING

The petitioner here, and the plaintiff-appellee in the Federal Circuit, is the Wisconsin Alumni Research Foundation ("WARF"). The respondent here, and the defendant-appellant in the Federal Circuit, is Apple Inc. ("Apple").

RULE 29.6 STATEMENT

Pursuant to Rule 29.6 of the Rules of this Court, petitioner WARF states that it has no parent corporation and that no publicly held company owns 10 percent or more of its stock. WARF is a not-for-profit Wisconsin corporation and is the designated patent management organization of the University of Wisconsin – Madison.

TABLE OF CONTENTS

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

TABLE OF CONTENTS (Cont'd)

TABLE OF AUTHORITIES

TABLE OF AUTHORITIES (Cont'd)

TABLE OF AUTHORITIES (Cont'd)

TABLE OF AUTHORITIES (Cont'd)

TABLE OF AUTHORITIES (Cont'd)

TABLE OF AUTHORITIES (Cont'd)

PETITION FOR A WRIT OF CERTIORARI

Petitioner WARF respectfully petitions for a writ of certiorari to review the judgement of the United States Court of Appeals for the Federal Circuit.

OPINIONS BELOW

The Federal Circuit's opinion is reported at 905 F.3d 1341 and is reproduced at App. 1a-25a. Its order denying rehearing is unreported and is reproduced at App. 26a-27a.

The district court's Rule 50(a) opinion and order is unreported and is reproduced at App. 142a-150a. Its Rule 50(b) opinion and order is reported at 261 F. Supp. 3d 900 and is reproduced at App.151a-199a.

JURISDICTION

The Federal Circuit issued its judgment on September 28, 2018. App. 1a-25a. It denied rehearing on January 7, 2019. App. 26a-27a. This Court has jurisdiction under 28 U.S.C. § 1254(1).

STATUTORY AND CONSTITUTIONAL PROVISIONS INVOLVED

U.S. Constitution, Amendment V, provides in pertinent part:

> No person shall ... be deprived of life, liberty, or property, without due process of law;

U.S. Constitution, Amendment VII, provides:

> In suits at common law, where the value in controversy shall exceed twenty dollars, the right of trial by jury shall be preserved, and no fact tried by a jury, shall be otherwise reexamined in any court of the United States, than according to the rules of the common law.

Federal Rule of Civil Procedure 50 is reproduced at App. 243a-245a.

INTRODUCTION

The Federal Circuit has again forgotten that its "appellate function" requires "appropriate deference be applied to the review of fact findings." *Apple v. Samsung Elecs.*, 839 F.3d 1034, 1039 (Fed. Cir. 2016) (en banc). WARF requests this Court's guidance on the proper role of the appellate court.

WARF brought suit accusing Apple's processors of infringing U.S. Patent No. 5,781,752 ("the '752 patent"). During trial, mere hours before closing arguments, Apple asked the district court to construe a claim limitation that neither party had previously asked be construed—the phrase "associated with the particular" load instruction. The court refused, finding Apple had "waived any request" for a construction. The court accordingly instructed the jury to give the "particular" limitation its "plain and ordinary meaning as viewed from the perspective of a person of ordinary skill in the art [POSITA]." The jury heard expert testimony on a POSITA's understanding of this limitation in the relevant technical context. The jury found Apple's processors literally infringed WARF's patent, resulting in a $506 million judgment.

On appeal, despite the district court's waiver ruling, Apple asked the Federal Circuit to construe the "particular" limitation in the first instance. The Federal Circuit granted Apple's request and construed that limitation for the first time, without reviewing or even acknowledging the district court's waiver ruling. Further, the panel construed the "particular" limitation *de novo*, stating "our view" of its meaning, giving no deference to the jury's implicit fact findings regarding a POSITA's understanding.

In *Teva Pharmaceuticals USA, Inc. v. Sandoz, Inc.*, 135 S.Ct. 831 (2015), this Court held that (a) claim construction can be a mixed question of fact and law, (b) language's "meaning to a person of ordinary skill in the art [POSITA]" is a "factual finding," and (c) the Federal Circuit must review such fact findings under "the ordinary rule governing appellate review of factual matters." *Id.* at 838, 841. But *Teva* addressed a case where the district court, not jury, was the fact-finder regarding a POSITA's understanding.

This case thus presents an important question on the application of *Teva*: Where the district court properly instructed the jury to give a claim limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]," does *Teva* allow the Federal Circuit to construe that limitation *de novo*, giving no deference to the jury's implicit fact findings regarding a POSITA's understanding?

This case provides an ideal vehicle to resolve that question. Given Apple's waiver, the Federal Circuit did not find error in the district court having instructed the jury to give the "particular" limitation

its "plain and ordinary meaning." The jury reached its verdict after hearing competing expert testimony regarding a POSITA's understanding of that phrase. Nevertheless, the Federal Circuit disregarded the jury's fact-finding role, disregarded the relevant expert testimony, and construed the "particular" limitation *de novo*. App. 12a. The Federal Circuit adopted a construction that conflicts with the expert testimony the jury should be presumed to have credited had "substantial evidence" review been properly applied. The Federal Circuit's choice regarding the standard of review—*de novo*, rather than substantial evidence—was dispositive.

Teva should not permit "turning what is fundamentally a factual question for the jury regarding whether the accused systems and features infringe the patent claims into a legal one for the court." *NobelBiz, Inc. v. Global Connect, L.L.C.,* 876 F.3d 1326, 1326-27 (Fed. Cir. 2017) (O'Malley, J., dissenting from denial of rehearing en banc). "It is time [for] much-needed guidance" regarding the standard of review for "plain-and-ordinary-meaning construction[s]." *Id.*

The Federal Circuit did not stop after issuing its new construction. Despite construing the "particular" limitation for the first time on appeal, the Federal Circuit did not remand for proceedings consistent with its new construction. Instead, the panel assessed its new claim construction against the existing trial record, reversed the jury's infringement verdict, and granted Apple JMOL of no literal infringement.

This case accordingly presents another important question: May the Federal Circuit apply a new claim

construction issued on appeal to grant JMOL based on a trial record developed without that construction, instead of remanding for proceedings consistent with the new claim construction?

WARF submits that in failing to remand, the Federal Circuit disregarded due process, the Seventh Amendment, and fundamental fairness.

WARF formulated its trial strategy in reliance on Apple having sought no construction for "particular." This strategy included WARF narrowing its evidence presented, while proceeding to trial on only literal infringement, without trying infringement under the doctrine of equivalents ("DoE"). If this case were remanded, WARF would offer new evidence of literal infringement under the new claim construction, while also arguing DoE for the first time. Importantly, this would not be an exercise in futility; the Federal Circuit suggested WARF might have prevailed on literal infringement had the existing trial record contained just a modicum of additional evidence directed to the panel's new construction. App. 17a n.9.

The issue of whether to remand after the Federal Circuit issues a new claim construction is a recurring and widespread problem. The Federal Circuit tends to issue a new claim construction in over a third of patent appeals. Despite this, it has developed lines of divergent panel opinions applying different standards to remand decisions, resulting in a body of law that lacks any clear or consistent guidelines for when remand is appropriate after a new claim construction on appeal. Numerous appeals therefore risk panel-dependent outcomes. This Court's guidance is needed.

STATEMENT OF THE CASE

A. The Patented Technology

Microprocessors can run faster by executing instructions "out-of-order." But this is not always possible, as some instructions require data from prior instructions—a condition called "data dependence." For example, a "store" instruction writes data to memory, and a "load" reads data from memory. "Data dependence" exists when a load must read data previously written by a store. For correct results, such instructions must execute in-order.

At execution time, the processor may not yet know whether a given store and load are data dependent. But it may nevertheless execute those instructions out-of-order, "speculating" that they are probably not dependent. If speculation turns out correctly, the processor gains performance. But if a "mis-speculation" occurs, it causes an error and harms performance. Thus, the benefits of speculative out-of-order execution depend upon the accuracy of the processor's speculations.

Addressing this problem, WARF's '752 patent claims a circuit in a processor that makes "predictions" regarding whether data dependence likely exists between instructions. The predictions unlock performance gains by enabling the processor to engage in speculative out-of-order execution only when the predicted risk of harmful mis-speculations is acceptably low.

As described in the '752 patent, predictions (109) are maintained in entries of a prediction table (44):

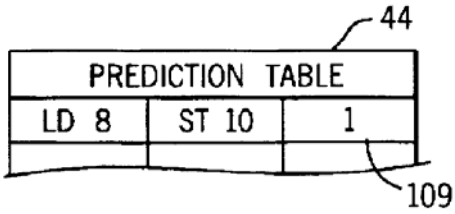

App. 34a. A table entry is created upon detecting that a given load-store pair has caused a mis-speculation. App. 66a. The circuit thereafter prevents speculation by later instances of those instructions if the prediction indicates a high risk of mis-speculation. App. 61a-62a.

As noted, a prediction is created upon detecting a mis-speculation. App. 66a. The '752 patent's claims recite a circuit for "detecting a mis-speculation where a [load] instruction" caused an error, and a circuit configured "to produce a prediction associated with the particular [load] instruction." App. 68a (Claim 1).[1] The phrase "prediction associated with *the particular* [load] instruction" is the sole limitation at issue. *Id.*[2] For shorthand, we call this the "particular" limitation.

B. The District Court Proceedings

As the Federal Circuit acknowledged, "neither party asked the district court to construe 'particular' before trial." App. 10a.

WARF filed this case in January 2014. App. 8a. The district court issued its claim construction and

---

[1] For clarity and convenience, the parties agreed to substitute "[load] instruction" for "data consuming instruction." The Federal Circuit did the same. App. 6a n.1.

[2] All emphases are added unless otherwise noted.

summary judgment order in August 2015. App. 72a-123a. The parties agreed to constructions for several claim limitations, and the district court adopted those constructions. App. 82a-83a. The parties disputed the meaning of one limitation ("prediction"), for which the court adopted WARF's proposal. App. 81a-82a, 87a-102a, 122a. Neither party sought construction of the "particular" limitation. *Id.* Moreover, neither party moved for summary judgment on infringement. App. 122a-123a.

Trial commenced in October 2015. Given the district court's claim construction and summary judgment order (App. 72a-123a), WARF declined to try DoE infringement and "proceeded only on a theory of literal infringement" at trial. App. 10a n.5.

On the second trial day, Apple indicated it planned to offer a non-infringement theory that assumed a construction of "particular." App. 128a. Apple's expert planned to testify that "a prediction associated with the particular [load] instruction" requires a prediction "associated with *one and only one* load instruction." App. 10a-11a; *see* App. 131a. Because Apple had never requested the "particular" limitation be construed, WARF moved during trial to exclude this non-infringement theory. App. 10a-11a, 128a-132a.

Opposing WARF's motion, Apple argued that the "particular" limitation "does not require construction" and that "plain and ordinary meaning should apply." App. 11a. The district court agreed and ruled in Apple's favor, concluding that "plain meaning" was sufficient and that "there is no need for instructing the jury on the meaning of this term." App. 132a.

Apple accordingly presented its non-infringement theory to the jury. App. 211a.

But on the fifth trial day, after both parties had rested their cases-in-chief, Apple for the first time requested that the district court instruct the jury on a claim construction for the "particular" limitation. Apple filed a motion stating that, while "the Court has likely ruled on this issue already," Apple "requests that the Court include the following language in the closing jury instructions:"

> "a prediction associated with *the particular* [load] instruction" means "a prediction associated with *a single* load instruction"

App. 133a.

The district court denied Apple's motion, stating:

> Consistent with Apple's position, the court denied WARF's request to exclude expert testimony and argument on this issue. The court also sided with Apple in its request that the term simply be given its plain and ordinary meaning. As such, Apple has *waived any request* to now insert a construction of the term into the closing jury instructions.

App. 136a-137a.

The court accordingly submitted the case to the jury without construing the "particular" limitation. The jury instructions provided eight constructions for other terms, before concluding with: "All other claim terms should be given their plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a.

Throughout the trial, the jury heard expert testimony from both sides regarding a POSITA's view on the ordinary meaning of "prediction associated with the particular [load] instruction" in the relevant technical context. App. 68a. As noted, Apple's expert Dr. August contended this requires a prediction "associated with *one and only one* load instruction." App. 10a-11a; *see* App. 131a.

But WARF's expert Dr. Conte offered competing testimony. He noted that the '752 patent's claims first recite a circuit for "detecting a mis-speculation where a [load] instruction" caused an error, then a circuit configured "to produce a prediction associated with *the particular* [load] instruction." App. 125a; *see* App. 68a (Claim 1). Dr. Conte testified that a POSITA would have understood "a prediction associated with the particular [load] instruction" means a prediction associated with "*the load that produced the mis-speculation*," referring back to "a [load] instruction" earlier in the claim. App. 125a. In other words, he explained, a POSITA would have understood that "particular just is identifying in this case an association"—*i.e.*, "the association with *the load instruction that you detected [produced] the mis-speculation*." App. 126a-127a. Further, Dr. Conte explained, a POSITA would have understood that the claimed association need not have "uniqueness" or be tied to "one and only [one]" load instruction. App. 125a-126a. In Dr. Conte's view, a POSITA would have understood that a prediction can be associated with *more than one* load instruction while still meeting this claim limitation, so long as that prediction is associated with *at least* "the load that produced the mis-speculation." App. 125a-127a.

The jury heard both parties' experts and agreed with WARF's expert Dr. Conte. It found the '752 patent literally infringed by Apple's processors, resulting in a judgment for WARF. App. 200a-202a.

In post-trial motions, Apple asked the district court to determine that no reasonable jury could find literal infringement of "a prediction associated with the particular [load] instruction." The court denied Apple's motions. App. 145a-146a, 156a-160a. It explained that, based on the evidence and testimony presented, "a reasonable jury could conclude that a prediction was associated with a particular load []instruction *even if that same prediction may be associated with other load instructions*." App. 157a; *see* App. 146a. The jury was free to credit Dr. Conte's testimony and to reject that of Apple's Dr. August.

C. The Federal Circuit's Decision

On appeal, Apple reprised the argument from its post-trial motions. App. 10a. However, despite the district court's express ruling that Apple had "waived any request" to construe the "particular" limitation (App. 136a-137a), Apple urged the Federal Circuit to construe that limitation for the first time. The panel granted Apple's request.

The Federal Circuit began by noting that "neither party asked the district court to construe 'particular' before trial." App. 10a. It further recognized that, when WARF filed its motion to exclude Apple's non-infringement theory for relying on a construction Apple had never sought, "Apple responded by arguing that the term 'particular' should carry its plain and ordinary meaning." App. 11a. The panel further recognized that "the district court had agreed with

Apple that the term 'particular' should be given its plain and ordinary meaning and thus ruled that no jury instruction was necessary to define that term." App. 11a.

Importantly, the Federal Circuit did not find the district court erred in instructing the jury to give the "particular" limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a; *see* App. 10a-12a. The Federal Circuit also did not find the district court erred in ruling that Apple had "waived any request" to construe the "particular" limitation. App. 136a-137a. A district court's waiver ruling is reviewable solely for "abuse of discretion." *Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1385 (Fed. Cir. 2016). But the Federal Circuit did not assess the waiver ruling for abuse of discretion. Notwithstanding thorough briefing on the issue (App. 208a-214a), the panel's opinion ignored the district court's waiver ruling entirely, omitting any reference to it. App. 1a-25a.

Despite not finding error in the jury having been instructed to give the "particular" limitation its "plain and ordinary meaning" (App. 11a), the Federal Circuit construed that limitation anyway, for the first time on appeal, reasoning:

> Giving a term its plain and ordinary meaning does not leave the term devoid of any meaning whatsoever. Instead, "the 'ordinary meaning' of a claim term is its meaning to the ordinary artisan after reading the entire patent." *Phillips v. AWH Corp.*, 415 F.3d 1303, 1321 (Fed. Cir. 2005) (en banc).

App. 12a (citing *Phillips*, 415 F.3d at 1321, which set forth the general standard for claim construction).

However, the Federal Circuit did not assess the factual record to determine what "plain and ordinary meaning" *the jury* should be presumed to have given the "particular" limitation. The Federal Circuit did not, for example, assess the testimony of Dr. Conte on that issue. App. 125a-127a. Instead, the panel ignored the record and construed the "particular" limitation *de novo*, announcing "our view" of its meaning:

> In our view, the plain meaning of "particular," as understood by a person of ordinary skill in the art after reading the '752 patent, requires the prediction to be associated with a *single* load instruction. A prediction that is associated with more than one load instruction does not meet this limitation.

App. 12a (original emphasis). This, of course, directly conflicts with the testimony of Dr. Conte, which the jury presumably credited in finding infringement. App. 125a-127a.

Despite construing the "particular" limitation for the first time on appeal, the Federal Circuit did not remand. Instead, the panel assessed its new claim construction against the existing trial record to grant Apple JMOL of no literal infringement, stating: "Applying the plain and ordinary meaning of the term 'particular,' and drawing all reasonable inferences from the evidence in favor of WARF, we hold that no reasonable juror could have found literal infringement in this case." App. 13a.

The Federal Circuit concluded this by finding WARF had presented "insufficient evidence" at trial to support the jury's literal infringement verdict under the panel's new construction, even though the trial record under review had been developed without that construction. App. 15a ("we find that there is insufficient evidence to support WARF's theory"); App. 18a ("there is insufficient evidence to support the jury's finding that Apple's products literally satisfy the 'particular' limitation"). Inexplicably, the Federal Circuit even suggested its JMOL decision on literal infringement may have differed had WARF only presented certain evidence relevant to the panel's new claim construction during the previously-held trial. App. 17a n.9.

The law should not require that litigants be prescient. Nor should litigants need to waste valuable trial time presenting evidence directed to claim constructions no party has sought and no court has issued. A remand was necessary to allow WARF to present evidence directed to the Federal Circuit's new claim construction.

The Federal Circuit also should have remanded for a trial on DoE infringement. Its opinion ruled solely on literal infringement, while expressly noting that DoE was not before the court on appeal. App. 10a n.5; 13a; 18a. When WARF decided not to try DoE, neither party had sought to construe the "particular" limitation. Had the Federal Circuit's new claim construction governed at trial, WARF would have tried DoE infringement under that construction.

REASONS FOR GRANTING THE WRIT

I. UNDER *TEVA*, THE FEDERAL CIRCUIT CANNOT IGNORE JURY FACT FINDINGS ON A POSITA'S UNDERSTANDING OF "PLAIN AND ORDINARY MEANING"

In patent cases, claim construction is often "the single most important event," as it "is often the difference between infringement and non-infringement, or validity and invalidity." *Retractable Techs., Inc. v. Becton, Dickinson & Co.*, 659 F.3d 1369, 1370 (Fed. Cir. 2011) (Moore, J., dissenting from denial of rehearing en banc). Litigants and district courts spend millions of dollars and countless hours trying cases, only to have judgments subjected to inconsistent law applied by divergent Federal Circuit panels, some overturning jury verdicts by issuing new claim constructions based on the panel's own view of the limited factual record presented on appeal. When "claim construction appeals are 'panel dependent,'" it "leads to frustrating and unpredictable results for both the litigants and the trial court." *Id.*

In *Teva*, this Court held that (a) claim construction can be a mixed question of fact and law, (b) language's "meaning to a person of ordinary skill in the art [POSITA]" is a "factual finding," and (c) the Federal Circuit must review such fact findings under "the ordinary rule governing appellate review of factual matters." *Teva*, 135 S.Ct. at 838, 841. But *Teva* addressed a case where the district court, not jury, was the fact-finder regarding a POSITA's understanding.

This case thus presents an important question on the application of *Teva*: Where the district court

properly instructed the jury to give a claim limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]," does *Teva* allow the Federal Circuit to construe that limitation *de novo*, giving no deference to the jury's implicit fact findings regarding a POSITA's understanding?

This case presents an ideal vehicle to decide that question, as the flaws inherent to *de novo* review of jury fact findings are both crystalized and dispositive. The Federal Circuit did not find error in the district court instructing the jury to give the "particular" limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a; *see* App. 10a-12a. Moreover, the jury reached its verdict after hearing competing expert testimony regarding a POSITA's understanding. App. 125a-127a. Nevertheless, the Federal Circuit disregarded the jury's fact-finding role, disregarded the relevant expert testimony, and substituted "our view" of the limitation's meaning without addressing whether the jury's view, implicit in the verdict, was supported by substantial evidence.

A. Finding That Apple Had "Waived Any Request" For a Claim Construction, the District Court Appropriately Instructed the Jury to Apply "Plain and Ordinary Meaning" as Understood By a POSITA

Prior to this Court's decision in *Teva*, the Federal Circuit held that, "[w]hen the parties raise an actual dispute regarding the proper scope of [patent] claims, the court, not the jury, must resolve that dispute" by construing the claim limitation and providing that construction to the jury. *O2 Micro Int'l Ltd. v. Beyond*

Innovation Tech. Co., 521 F.3d 1351, 1360 (Fed. Cir. 2008). However, the *O2 Micro* court "recognize[d] that district courts are not (and should not be) required to construe *every* limitation present in a patent's asserted claims." *Id.* at 1362 (original emphasis). For example, a district court need not construe a limitation if the requesting party "has waived its right to request a construction of [the limitation]." *Broadcom Corp. v. Qualcomm Inc.*, 543 F.3d 683, 694 (Fed. Cir. 2008).

That waiver rule applies here. Finding waiver (App. 136a-137a), the district court instructed the jury to apply "plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a. The Federal Circuit never addressed waiver on appeal, much less found the abuse of discretion necessary to overturn a waiver finding. App. 10a-12a; *see Wi-LAN*, 830 F.3d at 1384-85 (claim construction waiver findings reviewable solely for "abuse of discretion"); *Nuance Commc'ns, Inc. v. Abbyy USA Software House, Inc.*, 813 F.3d 1368, 1372-73 (Fed. Cir. 2016) ("[T]he district court found in Nuance's favor by adopting the plain and ordinary meaning," and "[t]he fact that shortly before trial Nuance became dissatisfied with its own proposed construction and sought a new one does not give rise to an *O2 Micro* violation."). Nor did the Federal Circuit otherwise find error in the district court instructing the jury determine and apply the "particular" limitation's "plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a; *see* App. 10a-12a.

Accordingly, the jury in this case was properly instructed to act as the fact-finder to determine and

apply a POSITA's understanding of the "particular" limitation's plain and ordinary meaning. Despite this, the Federal Circuit flatly ignored the jury's implied fact findings on that issue.

B. Following *Teva*, Implicit Jury Fact Findings On a POSITA's Understanding of "Plain and Ordinary Meaning" Should Be Reviewed for Substantial Evidence

In *Teva*, this Court ruled that claim construction can be a mixed question of fact and law. 135 S.Ct. at 841. One may need "to consult extrinsic evidence in order to understand, for example, the background science or the meaning of a term in the relevant art during the relevant time period." *Id.* If such "subsidiary facts" are disputed, the fact-finder must "make subsidiary factual findings about that extrinsic evidence." *Id.* This may include "resolv[ing] a dispute between experts and mak[ing] a *factual finding* that, in general, a certain term of art had a particular *meaning to a person of ordinary skill in the art* at the time of the invention." *Id.*

After that factual finding comes "a legal analysis: whether a skilled artisan would ascribe that same meaning to that term *in the context of the specific patent claim under review*." *Teva*, 135 S.Ct. at 841 (original emphasis). However, "a factual finding may be close to dispositive of the ultimate legal question of the proper meaning of the term in the context of the patent." *Id.* at 841-42. "Simply because a factual finding may be nearly dispositive does not render the subsidiary question a legal one." *Id.* at 842.

Teva also made clear that such fact findings must be reviewed under "the ordinary rule governing

appellate review of factual matters." *Teva*, 135 S.Ct. at 838. As this Court explained, "appellate courts must constantly have in mind that their function is not to decide factual issues *de novo*." *Id.* at 837 (citation omitted).

In *Teva*, the relevant fact findings were made by a district court, so the "ordinary rule governing appellate review of factual matters" was the "clearly erroneous" standard set forth in Federal Rule of Civil Procedure 52(a)(6). *Id.* at 836. But in the present case, the jury was instructed to make the factual findings. *See* Part I.A, *supra*. This presents an opportunity for the Court to make clear that the "ordinary rule governing appellate review of factual matters" applies concerning *all* fact-finders, and thus when the jury is properly instructed to apply "plain and ordinary meaning," the standard of review is for substantial evidence, not *de novo*. *See* Fed. R. Civ. P. 50(a); *Corsicana Nat. Bank of Corsicana v. Johnson*, 251 U.S. 68, 80 (1919) ("[T]here was substantial evidence tending to support … plaintiff's contentions. What weight should be given to it was for the jury, not the court, to determine."); *Apple*, 839 F.3d at 1039 ("[T]he appellate court is required to review jury fact findings when they are appealed for substantial evidence.").

As *Teva* made clear, claim construction may include (1) a "factual finding" regarding language's "meaning to a person of ordinary skill in the art," and (2) a "legal analysis" of whether a POSITA "would ascribe that same meaning to that term *in the context of the specific patent claim*." 135 S.Ct. at 841 (original emphasis). Apple waived the "legal analysis" prong by arguing that the "particular" limitation did "not require construction" and that "plain and ordinary

meaning should apply." App. 11a; *see Broadcom*, 543 F.3d at 694; *Nuance*, 813 F.3d at 1372-73.

Apple's waiver thus left only the "factual finding" prong regarding "meaning to a [POSITA]." *Teva*, 135 S.Ct. at 841. It also made that fact finding "close to dispositive of the ultimate legal question of the proper meaning of the term in the context of the patent." *Id.* at 841-42. Given Apple's waiver, the district court properly tasked the jury with deciding that fact issue.

To assist fact-finders, "[e]xperts may be examined to explain terms of art, and the state of the art, at any time." *Teva*, 135 S.Ct. at 841. The fact-finder is well-positioned to "make 'credibility judgments' about witnesses," particularly when presented "with competing fact-related claims by different experts." *Id.* at 838, 840. Here, the jury heard competing expert testimony regarding a POSITA's understanding.

The Federal Circuit should have applied "the ordinary rule governing appellate review of factual matters." *Teva*, 135 S.Ct. at 838. In this case the jury served as fact-finder on the appealed issue, so the substantial evidence standard of review governs. *See* Fed. R. Civ. P. 50(a). On "substantial evidence" review, the Federal Circuit must "presume the jury resolved all underlying factual disputes in favor of the verdict." *Apple*, 839 F.3d at 1040. The Federal Circuit accordingly should have presumed the jury credited Dr. Conte, who testified that a POSITA would have understood the "particular" limitation does not exclude predictions associated with more than one load. App. 125a-127a.

But the Federal Circuit did not assess the record to determine what "plain and ordinary meaning" the

jury presumptively gave the "particular" limitation. Nor did it address Dr. Conte's related testimony. App. 125a-127a. Instead, it construed "particular" *de novo*, announcing "our view" of its plain meaning. App. 12a.

In contrast, the district court applied the correct standard at trial. In a sidebar, the court explained it was allowing Dr. Conte's testimony on a POSITA's understanding of plain meaning because "*it is the jury who will decide* what the ordinary or plain meaning is of a term of those skilled in the art." App. 126a.

Moreover, in denying Apple's post-trial JMOL motions, the district court found that "a reasonable jury could conclude that a prediction was associated with a particular load []instruction *even if that same prediction may be associated with other load instructions*." App. 157a; *see* App. 146a. In other words, the jury was free to reject Dr. August's theory that the "particular" limitation requires a prediction "associated with one and only one load instruction." App. 10a-11a; *see* App. 131a. The jury was free to instead credit Dr. Conte, who testified that a POSITA would have understood the "particular" limitation does not exclude predictions associated with more than one load instruction. App. 125a-127a. The Federal Circuit "note[d]" this aspect of the district court's ruling in a cryptic footnote, then said nothing further about it. App. 12a n.6. The district court applied the proper standard, consistent with *Teva*. The Federal Circuit did not.

C. Courts and Litigants Need Guidance On the Standard of Review for "Plain and Ordinary Meaning" Constructions

Significant uncertainty will result if the Federal Circuit may construe "plain and ordinary meaning" *de novo* after a jury was properly instructed to make that factual determination. *Teva* does not permit "turning what is fundamentally a factual question for the jury regarding whether the accused systems and features infringe the patent claims into a legal one for the court." *NobelBiz,* 876 F.3d at 1326-27 (O'Malley, J., dissenting from denial of rehearing en banc).

Opportunities for gamesmanship will abound if litigants can waive claim construction below, as Apple did here, then request *de novo* claim construction for the first time on appeal. The "district court's case-management authority to set a schedule for claim construction that requires parties to take positions on various dates and holds the parties to these positions" will be a dead letter. *Wi-LAN*, 830 F.3d at 1385. Litigants will try to make alleged claim requirements express on appeal, shifting positions after it is too late for their opponents to respond with new evidence. *See* Part II.B, *infra*. Parties may seek no construction knowing that, if they lose before the jury on "plain and ordinary meaning," they can try again by asking the Federal Circuit to construe that meaning *de novo*.

The threat of "plain and ordinary meaning" constructions facing *de novo* review is widespread, because the vast majority of claim terms are not construed by district courts. *See* James R. Barney & Charles T. Collins-Chase, *An Empirical Analysis of District Court Claim Construction Decisions, January*

to December 2009, 2011 Stan. Tech. L. Rev. 2, 9 (2011), (analyzing "211 district court decisions" and finding that the "average number of disputed claim terms construed per decision was 8.8 (median = 7)"); *see also O2 Micro*, 521 F.3d at 1360 ("[T]his court has repeatedly held that a district court is not obligated to construe terms with ordinary meanings, lest trial courts be inundated with requests to parse the meaning of every word in the asserted claims."). Without required deference to jury fact findings on a POSITA's understanding of "plain and ordinary meaning," every unconstrued term will be a potential appeal issue lying in wait, ripe for appellants to invite the Federal Circuit to consider *de novo*.

Furthermore, whether the Federal Circuit will entertain first-instance claim construction arguments appears to be panel-dependent. Some panels reject attempts to raise a new "claim construction argument ... in the guise of a challenge to the sufficiency of the evidence." *Comcast IP Holdings I LLC v. Sprint Commc'ns Co.*, 850 F.3d 1302, 1311 (Fed. Cir. 2017) (quoting *ePlus, Inc. v. Lawson Software, Inc.*, 700 F.3d 509, 520 (Fed. Cir. 2012)); *see Abbott Labs. v. Syntron Bioresearch, Inc.*, 334 F.3d 1343, 1357 (Fed. Cir. 2003) ("Since Syntron did not urge a particular claim construction of the disputed language before the district court, it has waived the right to do so on appeal."). Other panels hold that, "[a]lthough the district court declined to construe the claims, that does not preclude us from making that legal determination on appeal." *Bancorp Servs., L.L.C. v. Sun Life Assur. Co. of Canada (U.S.)*, 687 F.3d 1266, 1274 (Fed. Cir. 2012). Yet other panels vaguely note that, while "hesitant to construe claim terms for the

first time on appeal," they might do so if "the record is sufficiently developed to enable us to construe the term." *Meyer Intellectual Properties Ltd. v. Bodum, Inc.*, 690 F.3d 1354, 1368-69 (Fed. Cir. 2012). Thus, whether the Federal Circuit may disrupt a jury's verdict by construing "plain and ordinary meaning" *de novo* appears dependent upon chance.

In sum, "[i]t is time [for] much-needed guidance" regarding the standard of review for "plain-and-ordinary-meaning construction[s]." *NobelBiz*, 876 F.3d at 1326-27 (O'Malley, J., dissenting from denial of rehearing en banc). This is particularly so following this Court's recent ruling that "meaning to a person of ordinary skill in the art" is a "factual finding." *Teva*, 135 S.Ct. at 841. WARF respectfully requests that the Court grant its petition for a writ of certiorari.

II. THE FEDERAL CIRCUIT DISREGARDS ITS APPELLATE FUNCTION WHEN IT FAILS TO REMAND AFTER ISSUING A NEW CLAIM CONSTRUCTION ON APPEAL

Despite construing the "particular" limitation for the first time, the Federal Circuit did not remand for proceedings consistent with its new construction. Instead, it assessed its new construction against the existing trial record, reversed the jury's verdict, and granted JMOL of no literal infringement. App. 13a.

This case accordingly presents another important question: May the Federal Circuit apply a new claim construction issued on appeal to grant JMOL based on a trial record developed without that construction, instead of remanding for proceedings consistent with the new claim construction?

WARF submits that, in failing to remand for proceedings consistent with its new construction, the Federal Circuit disregarded due process, the Seventh Amendment, and fundamental fairness.

This issue is a recurring and widespread problem. The Federal Circuit has developed divergent lines of panel opinions applying different standards to remand decisions, resulting in a body of law that lacks any clear or consistent guidelines for when remand is appropriate after a new claim construction on appeal. This Court's guidance is needed.

A. The Federal Circuit Has Failed to Articulate Any Clear Guidelines For When Remand Is Appropriate After a New Claim Construction on Appeal

The Federal Circuit tends to issue a new claim construction in over a third of appealed patent cases. *See* Kimberly A. Moore, *Markman Eight Years Later: Is Claim Construction More Predictable?*, 9 Lewis & Clark L. Rev. 231, 238 (2005) (finding "the number of cases in which one or more claim term was erroneously construed" to be "37.5%"); David L. Schwartz, *Practice Makes Perfect? An Empirical Study of Claim Construction Reversal Rates In Patent Cases*, 107 Mich. L. Rev. 223, 240 (2008) (finding the "[p]ercentage of cases with at least 1 wrongly construed term" to be "38.8%"). Yet the Federal Circuit has not articulated any clear guidelines for when remand is appropriate after such decisions.

Under a first line of Federal Circuit authority, after issuing a new claim construction on appeal, many panels remand if "the jury was not presented with the question of whether [the accused product]

infringes the asserted claims under [the new] construction." *Virnetx, Inc. v. Cisco Sys., Inc.*, 767 F.3d 1308, 1319 (Fed. Cir. 2014). Some "remand for a new trial," explaining that "a new trial is required so [the parties] can present evidence and argument that were not needed under the district court's original claim construction." *Cardiac Pacemakers, Inc. v. St. Jude Med., Inc.*, 381 F.3d 1371, 1374, 1382-83 (Fed. Cir. 2004); *see August Tech. Corp. v. Camtek, Ltd.*, 655 F.3d 1278, 1286 (Fed. Cir. 2011) ("Because the jury was given a flawed claim construction, … [we] remand to the district court for a limited trial on infringement."). Others "remand for further proceedings consistent with" the new construction, explaining that "the issue of whether JMOL or a new trial should be granted is an issue best addressed in the first instance by the district court." *Abbott Labs.*, 334 F.3d at 1351, 1358; *see Harris Corp. v. Ericsson Inc.*, 417 F.3d 1241, 1244 (Fed. Cir. 2005) ("We vacate the district court's denial of JMOL … and remand for reconsideration of that issue in light of our interpretation of the claims.").

Conversely, under a divergent line of authority, certain panels will issue a new claim construction on appeal, then apply it to the *existing* record to "reverse the district court's determination with respect to JMOL without remand." *SimpleAir v. Sony Ericsson Mobile Commc'ns*, 820 F.3d 419, 425, 431-32 (Fed. Cir. 2016) (finding "no reasonable jury could have found infringement under the proper construction" and directing "a judgment of no infringement") (quoting *Finisar Corp. v. DirecTV Grp., Inc.*, 523 F.3d 1323, 1333 (Fed. Cir. 2008)). Such panels believe "no remand is necessary" if "the record evidence does not

support an infringement verdict under the correct construction." *Eon Corp. IP Holdings v. Silver Spring Networks*, 815 F.3d 1314, 1320 & n.3 (Fed. Cir. 2016). Noticeably absent is any meaningful consideration of whether the parties might have presented different evidence or theories had the new construction governed below. Instead, finding a mere "failure of proof under the correct claim construction" based on the existing record purportedly suffices for the panel to enter "judgment of noninfringement as a matter of law." *CVI/Beta Ventures, Inc. v. Tura LP*, 112 F.3d 1146, 1161 (Fed. Cir. 1997) ("On the record before us, there is no evidence that the accused Turaflex frames meet the elasticity limitations … as those limitations are properly construed.").

Beyond the above lines of divergent authority, multiple Federal Circuit panels have also indicated remand decisions may hinge on amorphous factors. For example, some panels have vaguely suggested considering "the degree of difference between the incorrect construction and the correct construction." *Comcast*, 850 F.3d at 1310 (quoting *Finisar*, 523 F.3d at 1333).

As another example, some panels have asked whether it was sufficiently "clear from the record that the accused device does or does not infringe" under the new construction. *Praxair, Inc. v. ATMI, Inc.*, 543 F.3d 1306, 1324 (Fed. Cir. 2008); *see Lazare Kaplan Int'l, Inc. v. Photoscribe Techs., Inc.*, 628 F.3d 1359, 1370 (Fed. Cir. 2010) (remanding because "we cannot determine with any certainty that the accused machines infringe … under this new construction"); *Bowers v. Baystate Techs., Inc.*, 320 F.3d 1317, 1334 (Fed. Cir. 2017) (noting that "a change in the claim

construction at the appellate level generally necessitates a remand," yet granting JMOL because "the record contained undisputed evidence showing that the limitations [newly construed] are not met").

Simply put, the Federal Circuit has failed to articulate any clear guidelines for when remand is appropriate after a new claim construction on appeal. This case presents an ideal vehicle for the Court to provide much-needed direction.

B. The Federal Circuit Erred By Applying Its New Claim Construction to Grant JMOL Based On a Trial Record Developed Without That Construction

This case is not a close call on the remand issue. As the panel noted, "neither party asked the district court to construe 'particular' before trial." App. 10a. Then at trial, Apple contended "particular" did "not require construction" (App. 11a), and the district court ruled Apple had "waived any request" for a construction. App. 136a-137a. "Both parties tried the case and presented appellate arguments based on this understanding." *Apple*, 839 F.3d at 1043 n.4.

But on appeal, the Federal Circuit construed the "particular" limitation for the first time and then granted JMOL by finding "insufficient evidence" in the existing trial record to meet the panel's new construction. App. 15a, 18a. This was fundamentally unfair, because had the new construction governed below, the trial record would have been very different.

Importantly, this is not a case where attempting to establish infringement under the new construction would be futile. Indeed, the Federal Circuit's opinion

directly suggests that WARF likely could prevail if the case were remanded for proceedings consistent with the panel's new claim construction. If afforded that opportunity, WARF would offer evidence of literal infringement under the new construction, while also arguing DoE infringement for the first time.

1. Relevant Background

To recap, the '752 patent claims a circuit in a processor that makes "predictions" regarding whether load instructions are likely dependent on store instructions. Each prediction (109) is maintained in an entry of a prediction table (44) that associates the prediction with a load-store instruction pair:

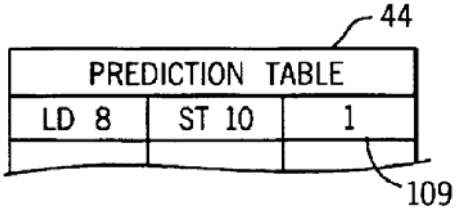

App. 34a. The '752 patent's claims recite "a prediction associated with the particular [load] instruction." App. 68a. On appeal, the panel construed this as requiring each "prediction to be associated with a *single* load instruction." App. 12a (original emphasis).

Apple's accused processors include a "Load-store Dependency Predictor ('LSD predictor')" circuit that "detects data dependencies between load and store instructions and uses a prediction table to make predictions based on those dependencies." App. 7a-8a. The load portion of each table entry holds a "load tag," which is a 12-bit number generated by a "hashing function." App. 8a. A particular "load instruction will always hash to the same 12-bit load tag." App. 14a.

But with "12 bits, only 4,096 load tags are available." App. 8a. This means "it is *possible* for multiple load instructions to hash to the same load tag," and thus "it is *possible* for multiple instructions to update the same prediction." App. 8a.

2. If Remanded, WARF Would Offer Evidence of Literal Infringement Under the New Claim Construction

The Federal Circuit construed the "particular" limitation as requiring each "prediction to be associated with a *single* load instruction." App. 12a (original emphasis). WARF had pointed to expert testimony indicating that Apple's predictions are associated with "a single load instruction at least sometimes (in fact, 99.9% of the time)." App. 15a; *see* App. 228a, 234a, 239a. However, the panel refused to credit this expert testimony, finding "there is not substantial evidence to support WARF's theory." App. 17a. The panel instead granted JMOL based on its own technical analysis, stating: "given that only 4,096 load tags are possible, and that Apple's operating system alone contains millions of load instructions, *the only reasonable inference to draw* is that load tags will *always* represent multiple load instructions." App. 17a.

Had the Federal Circuit remanded, WARF could have readily offered evidence directed to the new claim construction, showing that the panel's technical conclusion was not "the only reasonable inference." *Id.* The panel's opinion suggests this on its face.

For example, the panel granted JMOL by inferring that "load tags will always represent multiple load instructions." App. 17a. To infer this,

the panel reasoned that there are "only 4,096 load tags," yet some programs have "millions of load instructions." App. 17a. This was an appeal argument by Apple. App. 207a-208a. WARF countered by stating the obvious: "Programs can have fewer than 4,096 loads," *id.,* meaning the processors are capable of infringing. *See Finjan, Inc. v. Secure Computing Corp.*, 626 F.3d 1197, 1204 (Fed. Cir. 2010) ("[T]o infringe a claim that recites capability and not actual operation, an accused device 'need only be capable of operating' in the described mode."). The panel acknowledged WARF's argument in a footnote, but dismissed it on evidentiary grounds, stating: "Although WARF's brief states that programs can have fewer than 4,096 load instructions, *WARF has not pointed us to any evidence to support this assertion*." App. 17a n.9. But WARF would offer such evidence if this case were remanded for proceedings under the new claim construction. While not a focus at the first trial—because no claim construction necessitated it—the fact that programs can have fewer than 4,096 loads is indisputable. Apple never even attempted to deny this fact after WARF asserted it in its appeal brief. App. 207a-208a, 240a.

As another example, the focus on "millions of load instructions" (App. 17a) under the panel's new claim construction can be readily undercut by evidence on Apple's prediction "replacement policy." App. 235a-237a. As WARF explained in its rehearing petition, in Apple's processors the "predictions are short-lived, as they are constantly deleted and replaced." App. 236a. The panel's comparison of "4,096 load tags" to "millions of load instructions" was thus inaccurate (App. 17a), because "*only a small subset of*

instructions execute during a prediction's lifespan." App. 236a. The replacement policy was not a focus at trial—again because no construction necessitated it. If remanded, WARF would offer substantial evidence showing that the replacement policy supports literal infringement under the panel's new construction.

3. If Remanded, WARF Would Try Doctrine of Equivalents Infringement Under the New Claim Construction

At minimum, the Federal Circuit should have remanded for at least a trial on DoE infringement. The panel ruled solely on literal infringement, while noting that DoE was not before the court. App. 10a n.5. If remanded, WARF would try DoE under the new claim construction.

When WARF decided before trial not to try DoE, neither party had sought to construe the "particular" limitation, and the district court's claim construction order did not address it. App. 81a-82a, 87a-102a, 122a. WARF decided not to try DoE based on its reasonable beliefs (a) that no construction for the "particular" limitation would be given to the jury, and (b) that in the absence of any construction, the "particular" limitation would not entitle Apple to JMOL of no literal infringement. Notably, without a construction, Apple had not even moved for summary judgment of no literal infringement. App. 122a-123a.

The law should not require prescience. Nor should plaintiffs need to try alternative liability theories for each of numerous unconstrued claim limitations, fearful that the Federal Circuit might potentially construe one of them for the first time on appeal. *See* Part I.C, *supra*. Fundamental fairness required a

remand for a trial on DoE infringement under the Federal Circuit's new claim construction.

C. WARF's Due Process and Seventh Amendment Rights Required a Remand for Proceedings Under the Federal Circuit's New Claim Construction

WARF respectfully submits that both its Fifth Amendment right to procedural due process and its Seventh Amendment right to a trial by jury required a remand in this case.

1. WARF's Right to Procedural Due Process Required a Remand

When a new legal standard is adopted on appeal, "it [is] a violation of due process … to reverse a case and render judgment absolute, against a [party] who succeeded in the trial court" but has had "no occasion and no proper opportunity to introduce [] evidence" under the new legal standard. *Hamling v. United States*, 418 U.S. 87, 149-150 (1974); *see Saunders v. Shaw*, 244 U.S. 317, 319 (1917) (where state supreme court adopted new legal rule and applied it to reverse without remand, this Court found defendant "deprived of due process … because the case has been decided against him without his ever having had the proper opportunity to present his evidence" under the newly-adopted legal rule); *Carpet World, Inc. v. Riddles*, 737 P.2d 939, 942 (Okla. 1987) (where appeals court adopted new legal rule and then reversed without remand, Oklahoma Supreme Court found "Buyer's constitutional right to due process has been abridged in that Buyer has not been afforded the opportunity to present evidence" under the new rule).

The Federal Circuit's new claim construction is a new legal standard for which due process requires a remand. The trial record was not developed in expectation that the Federal Circuit would construe "particular" for the first time on appeal. This is especially so given Apple's contention at trial that this limitation "does not require construction" (App. 11a) and the district court's subsequent waiver ruling. App. 136a-137a. WARF did not have an incentive to fully develop and present evidence and theories directed to a nonexistent claim construction. Quite to the contrary, the first time Apple requested a construction for the "particular" limitation was hours before closing argument, after both parties had rested their cases-in-chief. App. 133a-135a. Simply put, this clearly was not the trial record "the parties would have presented had they known what [the Federal Circuit's] claim construction would be." *Furnace Brook LLC v. Aeropostale, Inc.*, 447 F. App'x 165, 170 (Fed. Cir. 2011) (O'Malley, J., dissenting). Failure to remand in such a case "deprives the parties of the ability to fully and fairly litigate the issues." *Id.*[3]

2. WARF's Right to a Trial By Jury Required a Remand

WARF's Seventh Amendment right to a jury trial also required a remand.

---

[3] It is perhaps ironic that Judge O'Malley herself was on panel in the present case. So too was Judge Bryson, who has likewise contended that in such cases "the remedy [sh]ould be, at most, a new trial," rather than JMOL. *Eon*, 815 F.3d at 1329 (Bryson, J., dissenting). This underscores the depth of the Federal Circuit's inconsistency—not just from panel-to-panel, but also across decisions involving the same judges.

A court construing claim language does "not deprive parties of their right to a jury trial," but only because the "right to a jury trial on the application of the properly construed claim to the accused device is preserved." *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 984 (Fed. Cir. 1995). Here, the Federal Circuit issued a new construction, but then went on to conduct "the application of the properly construed claim to the accused device" in the first instance. *Id.* The Federal Circuit, "by retroactively imposing new requirements for proving infringement, then re-finding the facts under the guise of determining whether these new requirements were complied with, has denied this litigant's historic right [to a jury trial]." *Malta v. Schulmerich Carillons, Inc.*, 952 F.2d 1320, 1345 (Fed. Cir. 1991) (Newman, J., dissenting). As this Court has admonished: "Maintenance of the jury as a fact-finding body is of such importance and occupies so firm a place in our history and jurisprudence that any seeming curtailment of the right to a jury trial should be scrutinized with the utmost care." *Id.* (quoting *Dimick v. Schiedt*, 293 U.S. 474, 486 (1935)).

The foregoing principle was applied correctly in *Odetics, Inc. v. Storage Tech. Corp.*, 185 F.3d 1259 (Fed. Cir. 1999). There, defendant STK argued on appeal that the jury had been instructed incorrectly on the term "first position" and that, under the correct construction, STK was entitled to JMOL of non-infringement. *Id.* at 1271. But the Federal Circuit rejected this argument due to "STK's failure to object to the jury instructions." *Id.* It further explained that an appellate "court cannot, consistent with the Seventh Amendment, evaluate a jury's verdict based

on ... a legal standard not given to the jury." *Id.* at 1272 (quoting *Mattison v. Dallas Carrier Corp.*, 947 F.2d 95, 108 (4th Cir. 1991)).

This case should have had a similar result. The district court found Apple "waived any request" to construe the "particular" limitation (App. 136a-137a), and the jury was instructed to give that limitation its "plain and ordinary meaning as viewed from the perspective of a [POSITA]." App. 140a-141a. As in *Odetics,* the Federal Circuit found no error in this. App. 10a-12a. Nevertheless, in contrast to its concern in *Odetics,* here the Federal Circuit evaluated the jury's verdict under the panel's new construction—"a legal standard not given to the jury." *Odetics*, 185 F.3d at 1272. *See Simmons v. City of Philadelphia*, 947 F.2d 1042, 1082 (3d Cir. 1991) ("[R]equesting a j.n.o.v. that depends on a jury instruction … never requested in essence asks the court to function as the jury in finding new facts and thus implicates equally serious seventh amendment concerns.").

* * *

Apple sought no construction before trial and was found by the district court to have "waived" any construction. At minimum, WARF deserves a fair opportunity to make its case under the new claim construction the Federal Circuit issued on appeal. The Federal Circuit's inconsistency on such remand questions—which potentially loom large in more than a third of patent appeals—warrants this Court's careful consideration.

CONCLUSION

"[T]he proper role of the court of appeals is not to reweigh the equities or reassess the facts but to make sure that the conclusions derived from those weighings and assessments are judicially sound and supported by the record." *Curtiss-Wright Corp. v. General Elec. Co.*, 446 U.S. 1, 10 (1980). That role is no different for the Federal Circuit, particularly where the jury was properly instructed to apply plain and ordinary meaning. A jury's verdict is owed no less deference merely because a patent is at issue.

WARF respectfully requests that the Court grant its petition for a writ of certiorari.

Respectfully submitted,

MORGAN CHU
Counsel of Record
GARY N. FRISCHLING
JASON SHEASBY
ALAN J. HEINRICH
CHRISTOPHER ABERNETHY
AMY E. PROCTOR
IRELL & MANELLA LLP
1800 Avenue of the Stars,
Suite 900
Los Angeles, CA 90067
(310) 277-1010
mchu@irell.com

Counsel for Petitioner

June 5, 2019

APPENDIX

APPENDIX A – Federal Circuit Opinion
(September 28, 2018)

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff-Appellee,

-v-

APPLE, INC.,

Defendant-Appellant.

2017-2265, 2017-2380

Appeals from the United States District Court for the Western District of Wisconsin in No. 3:14-cv-00062-wmc, Judge William M. Conley.

Decided: September 28, 2018

MORGAN CHU, Irell & Manella LLP, Los Angeles, CA, argued for plaintiff-appellee. Also represented by CHRISTOPHER ABERNETHY, GARY N. FRISCHLING, ALAN J. HEINRICH, AMY E. PROCTOR, JASON SHEASBY.

WILLIAM F. LEE, Wilmer Cutler Pickering Hale and Dorr LLP, Boston, MA, argued for defendant-appellant. Also represented by ANDREW J. DANFORD,

FELICIA H. ELLSWORTH, LAUREN B. FLETCHER, KEITH SYVERSON.

Before PROST, *Chief Judge,* BRYSON and O'MALLEY, *Circuit Judges.*

PROST, *Chief Judge*.

The Wisconsin Alumni Research Foundation ("WARF") sued Apple Inc. for infringement of U.S. Patent No. 5,781,752 ("the '752 patent"). After a two-week, bifurcated trial, a jury found Apple liable for infringement and awarded over $234 million in damages. The district court denied Apple's post-trial motions for judgment as a matter of law and for a new trial. Because no reasonable juror could have found infringement based on the evidence presented during the liability phase of trial, we reverse the district court's denial of Apple's motion for judgment as a matter of law. With respect to invalidity, we affirm the grant of summary judgment in favor of WARF.

I

The technology at issue relates to how computer processors execute a computer program's instructions. Computer programs are made up of lists of program instructions written in "program order." Although these instructions could be executed sequentially—i.e., in program order—some processors execute program instructions "out-of-order" to improve computer performance.

Of course, when executing instructions out-of-order, the processor must obtain the same result as if it had executed the instructions in program order. This can be complicated by the fact that "data

dependencies" may exist between individual program instructions. A data dependence exists between two instructions if one instruction relies upon data produced or modified by an earlier, or "older," instruction in the program order. To illustrate, the parties discuss "store" and "load" instructions that access the same location in memory. Memory can be thought of as a set of places to store data, where each place has an address by which the contents of that place can be accessed. *See* J.A. 1713 (testimony of Apple's expert, Dr. Colwell). A "store" instruction, or simply a "store," writes data to a given location in memory, overwriting any data that had previously been stored in that memory location. A "load instruction," or "load," reads data from a given memory location and then uses that data to perform some function. *Id.* at 1714. A data dependence exists between a store instruction and a load instruction if (1) the store instruction appears earlier than the load instruction in the program order; and (2) the store and load instructions will access the same memory location—i.e., the same address—if and when the store and load instructions are executed. In such a scenario, the load instruction depends on the store instruction having been executed first so that the data the load instruction reads from memory is current and correct.

At the time a processor decides whether to allow instructions to execute out-of-order, it may be unclear whether a data dependence exists between given store and load instructions. This is called an "ambiguous dependency." J.A. 1429 (testimony of WARF's expert, Dr. Conte), 1715 (testimony of Dr. Colwell). This ambiguity can occur, for example, if the

address where the store instruction will store data has not yet been determined, due to some independent calculation. Without knowing the ultimate storage location, the processor cannot determine whether the store and load instructions will access the same memory location and, thus, cannot determine whether a data dependence exists between those store and load instructions.

Even where an ambiguous dependency exists, the processor may nonetheless choose to execute the potentially dependent load instruction before the store instruction has finished executing. This is called "speculation" because the processor is effectively speculating that no data dependence exists between those store and load instructions. A "mis-speculation" occurs if a data dependence does exist between the two instructions, and the processor executes the dependent load instruction before the store instruction. If a processor correctly speculates—in other words, if the processor correctly guesses that a load instruction is *not* dependent on an earlier store instruction that has not yet executed—processor performance may be improved because the processor did not needlessly delay execution of that load instruction. J.A. 1431–32 (testimony of Dr. Conte). But, if a processor mis-speculates, the processor essentially has to discard work it has already performed and re-do the work in the correct order. J.A. 1433 (testimony of Dr. Conte), 1717–19 (testimony of Dr. Colwell). This recovery process is called "squashing" or "flushing." J.A. 1433 (testimony of Dr. Conte). As might be expected, mis-speculations do not help processor performance, and may in fact harm performance. In short, while out-of-order

execution of instructions with ambiguous dependencies may improve performance in cases where the processor speculates correctly, performance may be decreased by mis-speculation.

One method to minimize mis-speculation is for the processor to make an informed decision as to whether it should speculate. This is called "prediction." This case concerns a particular prediction method used to increase the accuracy of processor speculation such that mis-speculations are minimized.

A

The '752 patent, which expired on December 26, 2016, describes a specific prediction technique for an out-of-order processor. In this case, WARF asserted independent claims 1 and 9, as well as dependent claims 2, 3, 5, and 6. Claim 1 reads:

> 1. In a processor capable of executing program instructions in an execution order differing from their program order, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a [load] instruction dependent for its data on a [store] instruction of earlier program order, is in fact executed before the [store] instruction, a data speculation decision circuit comprising:
>
> a) a predictor receiving a mis-speculation indication from the data speculation circuit to produce a prediction associated with the particular [load] instruction and based on the mis-speculation indication; and

b) a prediction threshold detector preventing data speculation for instructions having a prediction within a predetermined range.

'752 patent claim 1.[1] Claim 9 reads:

> 9. In a processor capable of executing program instructions in an execution order differing from the program order of the instructions, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a [load] instruction dependent for its data on a [store] instruction of earlier program order, is in fact executed before the [store] instruction, a data speculation decision circuit comprising:
>
> a) a prediction table communicating with the data speculation circuit to create an entry listing a particular [load] instruction and [store] instruction each associated with a prediction when a mis-speculation indication is received; and
>
> b) an instruction synchronization circuit only instructing a processor to delay a later execution of the particular [load] instruction if the prediction table includes an entry.

---

[1] The modifications to the claim language reflect the parties' substitutions, for clarity, of the term "load" for "data consuming," and the term "store" for "data producing." *See* Appellant's Br. 12 n.1; Appellee's Br. 13. We adopt this helpful substitution in this opinion.

'752 patent claim 9.[2]

According to the claim language, when the data speculation circuit detects a mis-speculation, it sends a mis-speculation indication to a predictor. *Id.* at claim 1. The predictor then produces a prediction, based on the mis-speculation indication, as to whether a data dependence likely exists between the corresponding load and store instructions. A higher prediction value indicates a greater likelihood of data dependence and, therefore, a greater likelihood that a mis-speculation will occur if those instructions are executed out-of-order. *Id.* at col. 11 ll. 29–32. The prediction may be "updated based on historical mis-speculations detected by the data speculation circuit." *Id.* at col. 8 ll. 8–9. Going forward, the predictor "provides a dynamic indication to the data speculation circuit . . . as to whether data speculation should be performed." *Id.* at col. 8 ll. 1–3. And, if the prediction for a given load instruction exceeds a certain "predetermined range," speculation is prevented. *Id.* at claim 1.

B

The products accused of infringement in this case are Apple's A7, A8, and A8X integrated circuit chips, which include one or more processors. These processors include a Load-Store Dependency Predictor ("LSD predictor"), which is the technology at issue in this case.

The LSD predictor detects data dependences between load and store instructions and uses a

---

[2] The parties have not distinguished these claims for purposes of the infringement issues on appeal.

prediction table to make predictions based on those dependences. Each entry in the prediction table includes (among other things) a load tag, a store tag, and a prediction (or "counter"). The load tag is generated by taking certain information about a load instruction, such as its address, and creating a 12-bit load tag using a hashing function. Because the load tags are limited to 12 bits, only 4,096 load tags are available. The hashing algorithm uses a one-way hash, meaning that a given load tag cannot be expanded back to the load instruction that generated that load tag. Moreover, based on this hashing algorithm, it is possible for multiple load instructions to hash to the same load tag.

When multiple load instructions hash to the same load tag, it is possible for multiple instructions to update the same prediction in the LSD predictor's prediction table. This is called "aliasing." *See* J.A. 2237 (testimony of Apple's expert, Dr. August), 2294 (same); *see also* Appellant's Br. 23–24; Appellee's Br. 14. This means that a given instruction's history may impact the behavior of all load instructions that share the same load tag. J.A. 2166 ll. 15–18, 2168 ll. 7–11 (testimony of Dr. August).

C

WARF filed this patent infringement suit against Apple in January 2014.[3] Apple answered and asserted

---

[3] WARF has since filed a second infringement suit against Apple with respect to additional products released by Apple that WARF also believes infringe the '752 patent. *See* Compl., *Wis. Alumni Research Found. v. Apple Inc.*, No. 3:15-cv-00621-WMC (W.D. Wis. Sept. 25, 2015), ECF No. 1. That case is currently stayed pending the outcome of this appeal. *See* J.A. 20346.

counterclaims for declaratory judgment of non-infringement and invalidity of the '752 patent.

Before trial, both Apple and WARF moved for summary judgment with respect to Apple's counterclaims and defenses of anticipation under 35 U.S.C. § 102, based on U.S. Patent No. 5,619,662 ("Steely"). Specifically, Apple asserted that claims 1–3, 5, 6, and 9 of the '752 patent were invalid as anticipated by Steely. The district court granted summary judgment of no anticipation in favor of WARF.

Once the case proceeded to trial, the district court bifurcated the trial into two phases: liability and damages. After the liability phase, the jury found the asserted claims infringed and not invalid.[4] After the damages phase of trial, the jury awarded WARF over $234 million in damages.

After trial, Apple moved for judgment as a matter of law ("JMOL") and, in the alternative, for a new trial. The district court denied Apple's post-trial motions in their entirety. Apple timely appealed. This court has jurisdiction under 28 U.S.C. § 1295(a)(1).

II

"We review a district court's denial of JMOL or a new trial under the law of the regional circuit." *LifeNet Health v. LifeCell Corp.*, 837 F.3d 1316, 1322 (Fed. Cir. 2016). The Seventh Circuit reviews a district court's denial of a motion for judgment as a matter of law de novo. *Clarett v. Roberts*, 657 F.3d

[4] We note that the invalidity issues presented to the jury are not before this court on appeal.

664, 674 (7th Cir. 2011). In doing so, the appellate court "review[s] the record as a whole to 'determine whether the evidence presented, combined with all reasonable inferences permissibly drawn therefrom, is sufficient to support the verdict when viewed in the light most favorable to the party against whom the motion is directed.'" *Id.* (quoting *Erickson v. Wis. Dep't of Corr.*, 469 F.3d 600, 601 (7th Cir. 2006)). A jury verdict will be overturned only if no reasonable juror could have found in the non-movant's favor. *Id.*

A

Apple contends that no reasonable juror could have found that Apple's processors literally infringe the asserted claims of the '752 patent.[5] Specifically, Apple argues that its processors satisfy neither the "particular" nor the "mis-speculation" limitations recited in each of the claims.

With respect to the "particular" limitation, independent claim 1 requires a predictor that "produce[s] a prediction associated with the particular [load] instruction." '752 patent claim 1. Likewise, independent claim 9 requires a prediction table that "create[s] an entry listing a particular [load] instruction and [store] instruction each associated with a prediction." *Id.* at claim 9.

Although neither party asked the district court to construe "particular" before trial, WARF moved during trial to preclude Apple's expert, Dr. August, from testifying that a prediction cannot be associated

---

[5] WARF abandoned its theory of infringement under the doctrine of equivalents before trial, and has proceeded only on a theory of literal infringement.

with a "particular" load instruction if each load tag represents multiple load instructions. J.A. 18646–62. Specifically, WARF argued that Apple's expert should have been forbidden from making any suggestion that each prediction must be associated with one and only one load instruction.

Apple responded by arguing that the term "particular" should carry its plain and ordinary meaning, and that its expert's theory of non-infringement was consistent with that meaning. J.A. 18728 ("Apple has always maintained that the phrase—which uses only an ordinary word—does not require construction; the plain and ordinary meaning should apply."); J.A. 18730 ("Apple believes that the word 'particular' does not require any construction, because the '752 patent uses the word in its ordinary sense."); J.A. 18728 ("With respect to the 'particular' limitation, Dr. August has consistently applied the plain and ordinary meaning of the claim language."). Apple explained that the plain and ordinary meaning of "particular" meant that the claimed "prediction" must be associated with a *single* load instruction (i.e., one and only one load instruction), rather than with a *group* of load instructions. *See* J.A. 18729–33; *see also* J.A. 144 (Dist. Ct. Op. (summarizing Apple's argument)).

The district court denied WARF's motion to exclude the testimony of Apple's expert. J.A. 142–46. In doing so, the district court agreed with Apple that the term "particular" should be given its plain and ordinary meaning and thus ruled that no jury instruction was necessary to define that term. J.A. 145. Consistent with Apple's understanding of the plain and ordinary meaning, the district court

reasoned that "[f]rom the court's reading of claim 1 as a whole, it contemplates a *single* load instruction." J.A. 144 (emphasis added). In the district court's view, this was "consistent with the plain meaning of the claim terms 'the' and 'the particular.'" J.A. 145. The court thus "conclude[d] that claim 1 discloses a prediction associated with a *single* load instruction." J.A. 145 (emphasis added).[6]

On appeal, WARF does not dispute the district court's decision to give the term "particular" its plain and ordinary meaning. *See* Appellee's Br. 11, 26. Instead, WARF appears to disagree with the district court's understanding of the plain meaning. Giving a term its plain and ordinary meaning does not leave the term devoid of any meaning whatsoever. Instead, "the 'ordinary meaning' of a claim term is its meaning to the ordinary artisan after reading the entire patent." *Phillips v. AWH Corp.*, 415 F.3d 1303, 1321 (Fed. Cir. 2005) (en banc). In our view, the plain meaning of "particular," as understood by a person of ordinary skill in the art after reading the '752 patent, requires the prediction to be associated with a *single* load instruction. A prediction that is associated with more than one load instruction does not meet this limitation.[7]

---

[6] We note that the district court, in its order denying Apple's JMOL, stated that "a reasonable jury could conclude that a prediction was associated with a particular load instruction even if that same prediction may be associated with other load instructions." J.A. 205 (alteration omitted).

[7] WARF contends that this view of the plain meaning reads out the preferred embodiment of the '752 patent, which purportedly uses partial instruction addresses to identify load instructions,

Applying the plain and ordinary meaning of the term "particular," and drawing all reasonable inferences from the evidence in favor of WARF, we hold that no reasonable juror could have found literal infringement in this case. As explained above, each entry in Apple's LSD prediction table includes, among other things, a load tag and a prediction. Each load tag is generated by hashing information about a load instruction, such as its address, down to a 12-bit load tag. Only 4,096 load tags are possible. And because of the way Apple's hashing algorithm is designed, multiple load instructions may hash to the same load tag. Each load tag can therefore be associated with a *group* of load instructions—namely, all of the load instructions that hash to the same load tag. The practical effect of this is that a given load instruction's history will impact the prediction associated with *all* load instructions that hash to that same load tag.

WARF first contends that the "prediction" corresponding to a load tag will necessarily remain associated with a "particular" load instruction that

---

similar to the load tags used in Apple's products. *See* Appellee's Br. 28–29. We are unpersuaded. Figure 1 of the patent shows a program stored in memory "at a plurality of physical addresses 19 here depicted as xx1–xx6 where the values xx indicate some higher ordered address bits that may be ignored in this example." '752 patent col. 6 ll. 62–67; *see also id.* at col. 5 ll. 44–49 (discussing Fig. 2). Although Figures 5–8, which show the prediction table, then refer to the instructions as "LD 8" and "ST 10," there is no indication in the specification that instruction addresses are hashed or truncated before being added to the prediction table. Instead, the specification explains that the prediction table of Figure 5 is reviewed to determine if a "particular" instruction "identified by its physical address" is in the prediction table. *Id.* at col. 11 ll. 3–7.

mis-speculates because that load instruction will always hash to the same 12-bit load tag. Appellee's Br. 13. But, even accepting that a load instruction will always generate the same 12-bit load tag, *see* J.A. 2518 ll. 19–24, this is insufficient to satisfy this claim limitation because this argument ignores the plain and ordinary meaning of the term "particular," as described above. Under that meaning, it is not enough that an instruction hash to the same tag every time; the dispositive issue is whether *other* instructions *also* hash to that tag, such that the prediction is associated with a group of instructions, rather than a *particular* instruction.[8]

WARF's second argument for upholding the jury verdict appears to be that, even if the prediction must be associated with a single load instruction, the products still infringe in at least *some* circumstances—i.e., those in which aliasing does not occur. Appellee's Br. 15–18. Certainly, a product that "sometimes, but not always, embodies a claim nonetheless infringes." *Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1333 (Fed. Cir. 2013) (alteration omitted) (quoting *Bell Commc'ns Research, Inc. v.*

---

[8] On this point, WARF also argues that even if multiple instructions hash to the same load tag, "the 'prediction' merely becomes associated with *two* loads, including 'the particular [load]' that mis-speculated." Appellee's Br. 16. WARF then contends that infringement still exists because the preamble of the claims uses the word "comprising," which allows for additional, unrecited elements. *Id.* at 15–16. But "'[c]omprising' is not a weasel word with which to abrogate claim limitations," *Spectrum Int'l, Inc. v. Sterilite Corp.*, 164 F.3d 1372, 1380 (Fed. Cir. 1998), and WARF's application of that term here would frustrate the plain meaning of "particular" as used in this patent.

Vitalink Commc'ns Corp., 55 F.3d 615, 622–23 (Fed. Cir. 1995)). But after reviewing the evidence and drawing all reasonable inferences in favor of WARF, we find that there is insufficient evidence to support WARF's theory that Apple's load tags are *sometimes* associated with a single load instruction.

The evidence WARF points us to in support of this theory is sparse. WARF contends that the frequency of "aliasing" in Apple's products is low (specifically, 0.1%), which WARF takes to mean that load tags represent a single load instruction at least sometimes (in fact, 99.9% of the time). This conclusion, however, does not follow from the evidence cited by WARF.

First, the inference WARF seeks to draw from the evidence cited is not reasonable. The 0.1% statistic comes from an email from Stephan Meier, an Apple engineer, that pertains to testing various hashing functions. J.A. 1499 (testimony of Dr. Conte). During trial, the parties disputed the meaning of the 0.1% statistic, with WARF arguing that it represents the *frequency* of aliasing, and Apple arguing that it represents the *performance impact* of aliasing, *see, e.g.*, J.A. 2238 ll. 10–18 (cross-examination testimony of Dr. August). Although there are a few isolated statements from Apple's fact and expert witnesses that WARF argues support its theory, *see* J.A. 2238–40, the most thorough explanation of this piece of evidence comes from WARF's expert, and his testimony undermines the inference WARF seeks to draw. According to WARF's expert, Mr. Meier was trying to determine the "performance impact" of using different hashing functions to determine which hashing function performs best, including an analysis of how many bits should be used for the load tag. J.A.

1499–500 (testimony of Dr. Conte). As explained by WARF's expert, Mr. Meier concluded that a 9-bit load tag would cause a loss of "0.9 percent in performance"; a 10-bit load tag would cause a 0.4% loss in performance; and a 12-bit load tag would cause less than a 0.1% loss in performance. J.A. 1500 (testimony of Dr. Conte). Despite this explanation, which indicates that Mr. Meier's statistics were indeed describing overall performance, WARF's expert jumped to the conclusion that "aliasing is very rare." J.A. 1501 (testimony of Dr. Conte). But, in light of Dr. Conte's testimony, it is unreasonable to infer that the 0.1% statistic was referring to the *frequency* of aliasing.

Second, even accepting WARF's unreasonable view of this evidence (that the *frequency* of aliasing is 0.1%), this does not support an inference that load tags sometimes represent a single load instruction. "Aliasing" does not simply refer to two load instructions hashing to the same load tag. Instead, "aliasing" occurs when two load instructions actually update the same prediction in operation because they share the same load tag. *See* J.A. 2294 (testimony of Dr. August) ("Q: First, what's the difference between the load tags' grouping of load instructions and the concept of aliasing? A: So the grouping is always present. Aliasing is when the program is running, what is the performance impact of that grouping."); *see also* Appellant's Br. 23–24; Appellee's Br. 14. It is therefore not reasonable to infer that load instructions rarely hash to the same load tag, merely because the frequency of load instructions actually updating the same prediction during operation is low.

Finally, WARF points to Apple's technical documentation, arguing that certain language in the documentation demonstrates that Apple's LSD predictor "uniquely" identifies load instructions. Appellee's Br. 14 (citing J.A. 10131); *see also* J.A. 1489 l. 8–1490 l. 8 (testimony of Dr. Conte, discussing J.A. 10131). Apple points out, however, that the documentation merely states that the LSD predictor "can be thought of" as uniquely identifying load instructions, Reply Br. 2–3 (quoting J.A. 10131), and that "in practice" the load tags are the result of applying the hashing algorithm. *Id.*; *see also* J.A. 2178 ll. 3–22 (testimony of Dr. August). Reading the quote in context, it is not reasonable to infer that the load tags, in practice, uniquely identify load instructions. And even if this inference were reasonable, it would not be enough to support a finding that Apple's processors actually practice the "particular" limitation.

In short, there is not substantial evidence to support WARF's theory that, in Apple's LSD predictor, a prediction (by way of a load tag) is at least sometimes associated with a single load instruction. And, given that only 4,096 load tags are possible, and that Apple's operating system alone contains millions of load instructions, the only reasonable inference to draw is that load tags will always represent multiple load instructions. *See* J.A. 1605–06 (testimony of Dr. Conte), 2296–97 (testimony of Dr. August).[9]

---

[9] Although WARF's brief states that programs can have fewer than 4,096 load instructions, WARF has not pointed us to any evidence to support this assertion. *See* Appellee's Br. 17.

In sum, drawing all reasonable inferences in favor of WARF, there is insufficient evidence to support the jury's finding that Apple's products literally satisfy the "particular" limitation. As this conclusion is sufficient to set aside the jury's infringement finding, we need not address Apple's arguments regarding the "mis-speculation" limitation.

B

Apple also contends that the district court erred in granting summary judgment of no anticipation based on the Steely prior art reference. The district court determined that Steely did not disclose the "prediction" claimed in the '752 patent. *Wis. Alumni Research Found. v. Apple, Inc.*, No. 14-cv-062-WMC, 2015 WL 4668247, at *13–16 (W.D. Wis. Aug. 6, 2015) ("*MSJ Order*"). In Apple's view, this determination was based on an incorrect construction of the term "prediction." Apple also contends that, even under the court's construction, a genuine dispute of material fact exists, making summary judgment improper.

1

Claim construction is ultimately a legal question reviewed de novo, with any subsidiary fact-findings regarding extrinsic evidence reviewed for clear error. *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 135 S. Ct. 831, 841 (2015).

The parties dispute the construction of the term "prediction." WARF contends that a "prediction" must be dynamic, meaning it is capable of receiving updates. Apple contends that while a "prediction" includes dynamic predictions, the term is also broad enough to include static predictions (i.e., those incapable of receiving updates). The district court

agreed with WARF, concluding that a prediction, as used in the patent, must be "capable of receiving updates." *MSJ Order* at *13.

On appeal, Apple argues that the plain and ordinary meaning of "prediction" encompasses both dynamic and static predictions. Apple further contends that the patent's specification does not limit a "prediction" to being dynamic, and that by requiring that the prediction be capable of receiving updates, the district court improperly imported a limitation from the preferred embodiment. *See* Appellant's Br. 38–39. Apple's arguments are unpersuasive, as explained below.

First, "the 'ordinary meaning' of a claim term is its meaning to the ordinary artisan after reading the entire patent." *Phillips*, 415 F.3d at 1321. Reading the patent as a whole, it is clear that the claimed prediction must be capable of receiving updates. The term "prediction" is used throughout the specification to describe a prediction value that updates based on a given load instruction's historical mis-speculation behavior. *See* '752 patent col. 11 ll. 33–35 ("Normally the prediction 109 starts at zero when an entry is first made in the prediction table 44 and is incremented and decremented as will be described below."); *see also id.* at col. 8 ll. 7–11 ("The prediction provided by the predictor circuit 33, as will be described, is updated based on historical mis-speculations detected by the data speculation circuit 30. For this reason, the data speculation circuit 30 must communicate with the predictor circuit 33 on an ongoing basis."). Specifically, the prediction is updated as new information is gathered regarding the likelihood of future mis-speculation. *Id.* at col. 12 l. 61–col. 13 l. 3

("[T]he predictor circuit 33 must also make adjustments in its prediction table 44 if there is a mis-speculation, ... [T]he prediction table 44 is checked to see whether the LOAD/STORE pair causing the mis-speculation is in the prediction table 44 already. If so then at process block 302, the prediction 109 is updated toward synchronize so that this mis-speculation may be avoided in the future."); *id.* at col. 12 ll. 14–17 ("In this case, the prediction that there was a need to synchronize was wrong and so at process block 120 the prediction 109 is decremented toward the do not synchronize state."); *id.* at col. 12 ll. 50–55 ("In this case, the prediction 109 is updated toward the synchronize condition indicating that the prediction that there was a need to synchronize was correct as there is in fact a LOAD instruction waiting to be synchronized."). Where, as here, "a patent 'repeatedly and consistently' characterizes a claim term in a particular way, it is proper to construe the claim term in accordance with that characterization." *GPNE Corp. v. Apple Inc.*, 830 F.3d 1365, 1370 (Fed. Cir. 2016) (quoting *VirnetX, Inc. v. Cisco Sys., Inc.*, 767 F.3d 1308, 1318 (Fed. Cir. 2014)).

Second, Apple has not pointed us to any portion of the specification that describes a static prediction. Although Apple directs our attention to "alternative embodiments" for obtaining the prediction—methods other than incrementing, such as "various weighting schemes" or "complex pattern matching techniques"—none of the passages concerning these embodiments describe a static prediction. *See* Appellant's Br. 39 (citing '752 patent col. 14 ll. 6–14); Reply Br. 13 (citing same). Instead, the embodiments merely illustrate methods other than "simply incrementing it in value

for each speculation" for calculating the value of the prediction. '752 patent col. 14 ll. 8–9. In short, by allowing the claimed "prediction" to also include static predictions, Apple's proposed construction would "expand the scope of the claims far beyond anything described in the specification." *Kinetic Concepts, Inc. v. Blue Sky Med. Grp., Inc.*, 554 F.3d 1010, 1019 (Fed. Cir. 2009); *see id.* (limiting the term "wound" to "skin wound," rather than allowing it to encompass "pus pockets," where all of the examples in the specification involved skin wounds).

In sum, rather than improperly reading a limitation from the preferred embodiment into the claims, the district court's construction, with which we agree, properly reads the claim term in the context of the entire patent.

2

Apple next contends that the district court erred in granting summary judgment of no anticipation, even under the district court's construction of "prediction," which requires that the prediction be "capable of receiving updates." Appellant's Br. 40–42. Specifically, it contends that a genuine factual dispute exists as to whether Steely discloses predictions capable of receiving updates.

Applying Seventh Circuit law, "[w]e review the grant of summary judgment *de novo*, construing all facts and drawing all inferences 'in the light most favorable to the non-moving party.'" *Austin v. Walgreen Co.*, 885 F.3d 1085, 1087 (7th Cir. 2018) (quoting *Zuppardi v. Wal-Mart Stores, Inc.*, 770 F.3d 644, 649 (7th Cir. 2014)).

Steely discloses out-of-order processors that use past mis-speculations (or "collisions") to predict whether load instructions should be allowed to execute out-of-order. *See* Steely col. 2 ll. 63–66. Steely uses "tags" to indicate whether instructions that were "previously reordered and executed had a collision" and to "ascertain whether the . . . instructions can be reordered." *See id.* at col. 2 l. 64–col. 3 l. 1. Specifically, Steely assigns "tags" to load and store instructions that mis-speculate (or "collide"). Although Steely discloses multiple techniques for generating tags, the parties focus on the first technique disclosed. According to this technique, "when a pair of load and store instructions cause a problem the first time"—i.e., when they are executed out-of-order and a mis-speculation occurs—"a portion of the address in memory which resulted in a load-store collision are [sic] saved." *Id.* at col. 48 ll. 1–4. The example in Steely uses five bits of the memory address as the tag. *Id.* at col. 48 ll. 26–28. That tag is associated with the load and store instructions that mis-speculated when reordered. *Id.* at col. 48 ll. 26–29, 37–50; *see also* J.A. 15069 ¶ 162 (invalidity report of Dr. Colwell). The next time that same pair of load and store instructions is called, the instructions' tags are compared. If the tags match, those instructions will not be executed out-of-order. Steely col. 48 ll. 33–36; J.A. 15069 ¶¶ 162–63, 15136–37 ¶¶ 301–02 (invalidity report of Dr. Colwell).

Apple contends that the outcome of this comparison is a "prediction," as it indicates the likelihood of mis-speculation if those instructions are executed out-of-order. We agree that this is a reasonable inference to draw at the summary

judgment stage. The only question, then, is whether the outcome of that comparison is also "capable of receiving updates," as is required under the proper construction of the term "prediction."

Apple's expert, Dr. Colwell, provided an example in his invalidity report explaining how the outcome of this comparison can change. He reasoned that the outcome may change because, "[a]s more mis-speculations occur, one or both of the tags of the same pair of load and store instructions may change to different values, resulting in different outcomes from a comparison of the tags." J.A. 15137–38 ¶ 303. In other words, if the store instruction from the first load-store pair is reordered with a *different* load instruction, and a mis-speculation occurs, both of those instructions would receive a tag based on the memory address that the instructions were accessing. According to Dr. Colwell, this necessarily causes the store instruction's tag to change. And because the store instruction's tag has changed, that tag will no longer match the tag of the original load instruction. Thus, Dr. Colwell concludes that the outcome of the comparison of the original load-store pair will be different. *Id.* Based on this example, Apple contends that Steely discloses a variable "prediction" that is "capable of receiving updates."

WARF responds that Dr. Colwell's example is mere speculation regarding how Steely *might* be implemented, and that Steely itself never discloses that tags, once generated, can change. Appellee's Br. 45.

We agree with WARF that no reasonable juror could find that Steely's specification discloses the

behavior described in Dr. Colwell's example regarding changing tags. Apple points us to just two statements from the specification as support for this disclosure: (1) a statement regarding the size of the tag storage table, *see* Steely col. 47 ll. 56–60; and (2) a statement that tags "will be stored" in that table, *see id.* at col. 48 ll. 29–30. Based on the size of the table, Apple contends that a fact-finder could infer that each instruction can be associated with only a *single* tag (as opposed to, for example, allowing an instruction to carry multiple tags). So, the argument goes, because tags "will be stored," when a given instruction receives a new tag, the new tag is stored in the table, and the former tag is necessarily changed. But the inference Dr. Colwell would have a fact-finder draw from the size of the table is not reasonable. And although Apple cites additional evidence in attempt to bolster this theory (namely, uncorroborated inventor testimony and another reference stating generally that Steely uses a "prediction"), such evidence is insufficient to create a genuine dispute of material fact.

We therefore agree with the district court that no reasonable juror could find that Steely discloses the "prediction" limitation of the '752 patent's claims. We therefore affirm the district court's grant of summary judgment on this issue.

III

For the reasons stated above, we reverse the district court's denial of Apple's JMOL motion with respect to non-infringement, but affirm its grant of summary judgment with respect to Apple's anticipation defense based on Steely.

AFFIRMED-IN-PART AND REVERSED-IN-PART

COSTS

The parties shall bear their own costs.

APPENDIX B – Federal Circuit Order Denying Rehearing (January 7, 2019)

NOTE: This order is nonprecedential.

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff-Appellee,

-v-

APPLE, INC.,

Defendant-Appellant.

2017-2265, 2017-2380

Appeals from the United States District Court for the Western District of Wisconsin in No. 3:14-cv-00062-wmc, Judge William M. Conley.

ON PETITION FOR PANEL REHEARING AND REHEARING EN BANC

Before PROST, *Chief Judge*, NEWMAN, LOURIE, BRYSON[1], DYK, MOORE, O'MALLEY, REYNA, WALLACH, TARANTO, CHEN, HUGHES, and STOLL, *Circuit Judges*.

PER CURIAM.

O R D E R

Appellee Wisconsin Alumni Research Foundation filed a combined petition for panel rehearing and rehearing en banc. The petition was referred to the panel that heard the appeal, and thereafter the petition for rehearing en banc was referred to the circuit judges who are in regular active service.

Upon consideration thereof,

IT IS ORDERED THAT:

The petition for panel rehearing is denied.

The petition for rehearing en banc is denied.

The mandate of the court will issue on January 14, 2019.

FOR THE COURT

January 7, 2019
Date

/*s*/ Peter R. Marksteiner
Peter R. Marksteiner
Clerk of Court

---

[1] Circuit Judge Bryson participated only in the decision on the petition for panel rehearing.

APPENDIX C – U.S. Patent No. 5,781,752

(19) **United States Patent**
Moshovos et al.

(11) **Patent No.: US 5,781,752**

(45) **Date of Patent: Jul. 14, 1998**

(54) **TABLE BASED DATA SPECULATION CIRCUIT FOR PARALLEL PROCESSING COMPUTER**

(75) Inventors: **Andreas I. Moshovos; Scott E. Breach; Terani N. Vijaykumar; Gurindar S. Sohi,** all of Madison, Wis.

(73) Assignee: **Wisconsin Alumni Research Foundation**, Madison, Wis.

(21) Appl. No.: **773,992**

(22) Filed: **Dec. 26, 1996**

(51) **Int. Cl.**[6] **G06F 9/38**
(52) **U.S. Cl.** **395/392**
(58) **Field of Search** 395/392

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,664,138 9/1997 Yoshida 395/395
5,666,506 9/1997 Hesson et al. 395/392

OTHER PUBLICATIONS

Gurinda Sohi et al., Instruction Issue Logic for High-Performance Interruptable Pipelined Processors; ACM 1987, pp. 27–34.

Primary Examiner—Krisna Lim
Attorney, Agent, or Firm—Quarles & Brady

(57) **ABSTRACT**

A predictor circuit permits advanced execution of instructions depending for their data on previous instructions by predicting such dependencies based on previous mis-speculations detected at the final stages of processing. Synchronization of dependent instructions is provided by a table creating entries for each instance of potential dependency. Table entries are created and deleted dynamically to limit total memory requirements.

9 Claims, 7 Drawing Sheets

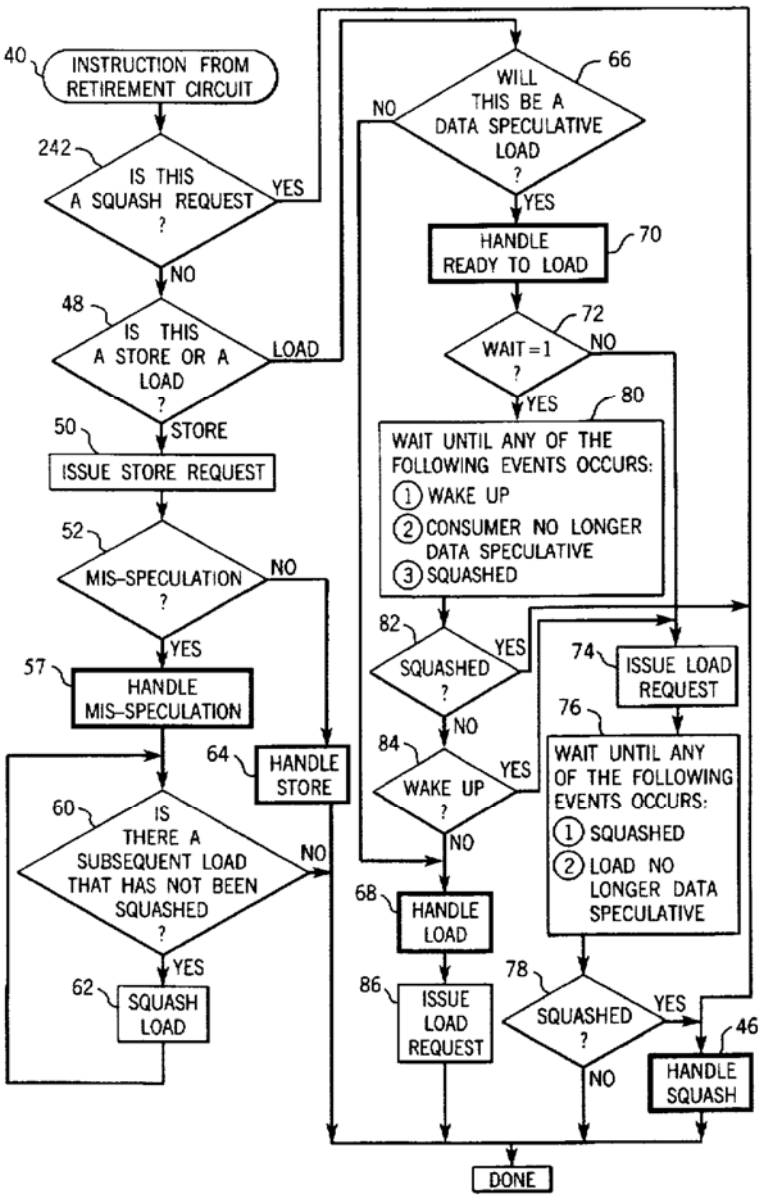
40
INSTRUCTION FROM RETIREMENT CIRCUIT
242
IS THIS A SQUASH REQUEST ?
YES
NO
48
IS THIS A STORE OR A LOAD ?
LOAD
STORE
50
ISSUE STORE REQUEST
52
MIS-SPECULATION ?
NO
YES
57
HANDLE MIS-SPECULATION
64
HANDLE STORE
60
IS THERE A SUBSEQUENT LOAD THAT HAS NOT BEEN SQUASHED ?
NO
YES
62
SQUASH LOAD
66
WILL THIS BE A DATA SPECULATIVE LOAD ?
NO
YES
70
HANDLE READY TO LOAD
72
WAIT = 1 ?
NO
YES
80
WAIT UNTIL ANY OF THE FOLLOWING EVENTS OCCURS:
1 WAKE UP
2 CONSUMER NO LONGER DATA SPECULATIVE
3 SQUASHED
82
SQUASHED ?
YES
NO
74
ISSUE LOAD REQUEST
76
WAIT UNTIL ANY OF THE FOLLOWING EVENTS OCCURS:
1 SQUASHED
2 LOAD NO LONGER DATA SPECULATIVE
84
WAKE UP ?
YES
NO
68
HANDLE LOAD
86
ISSUE LOAD REQUEST
78
SQUASHED ?
YES
NO
46
HANDLE SQUASH
DONE

Sheet 1 of 7

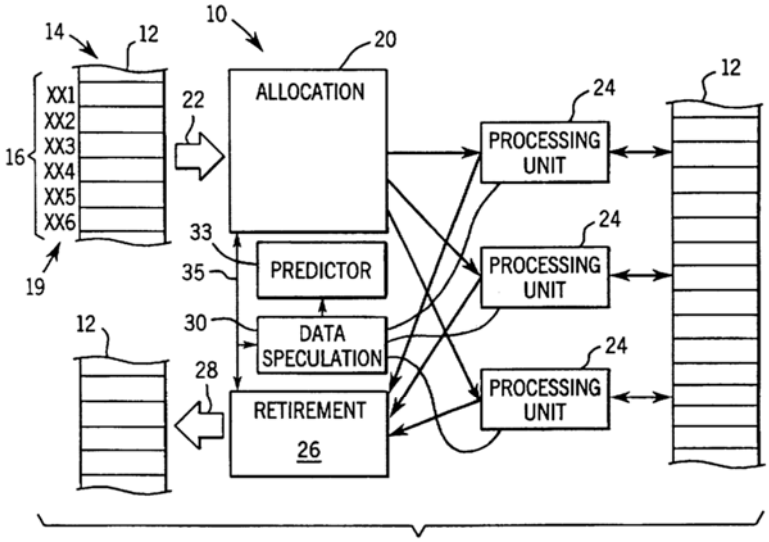

FIG. 1

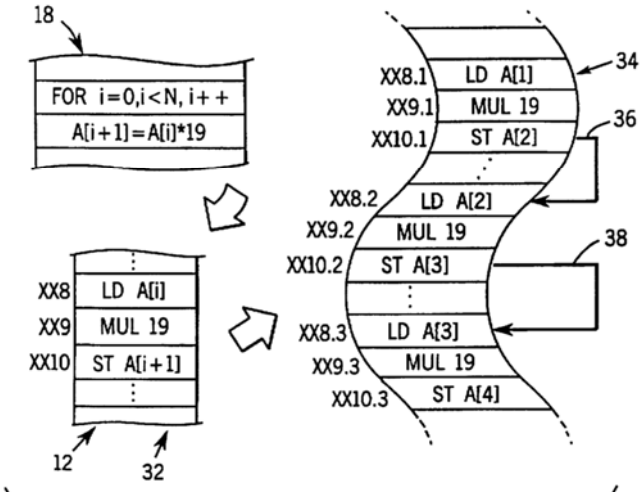

FIG. 2

Sheet 2 of 7

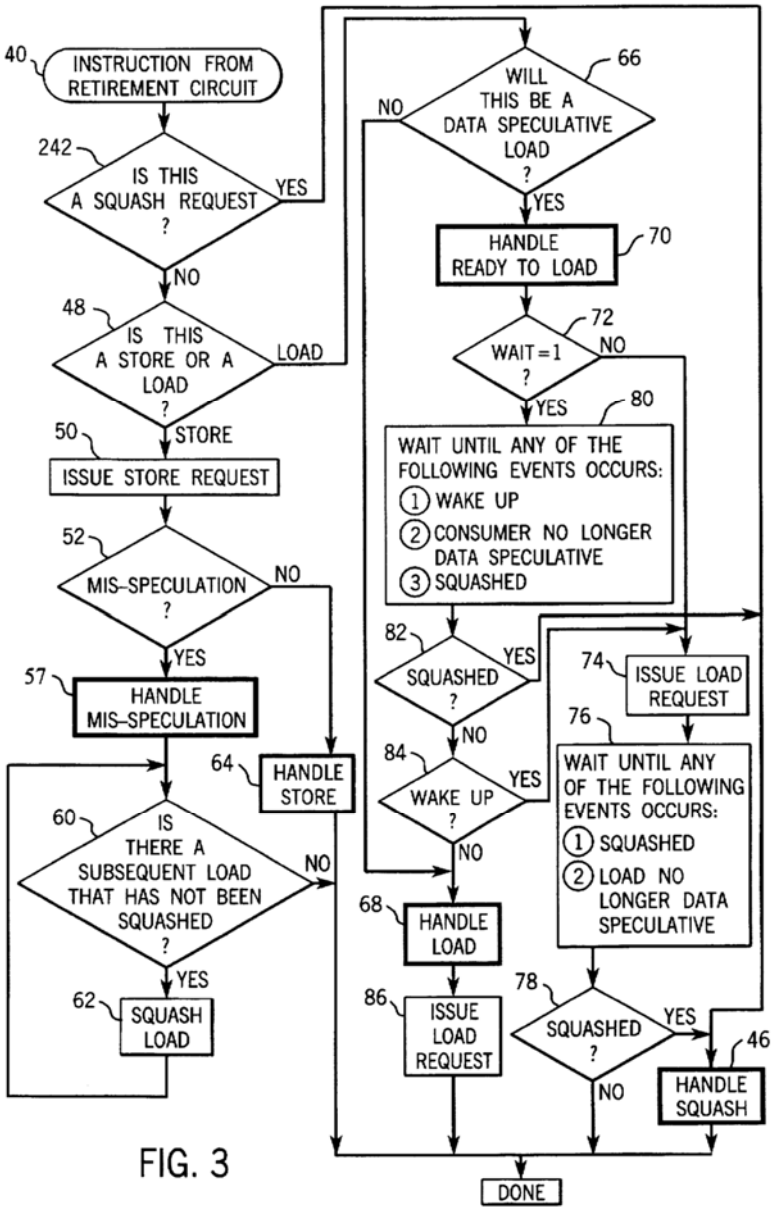

FIG. 3

Sheet 3 of 7

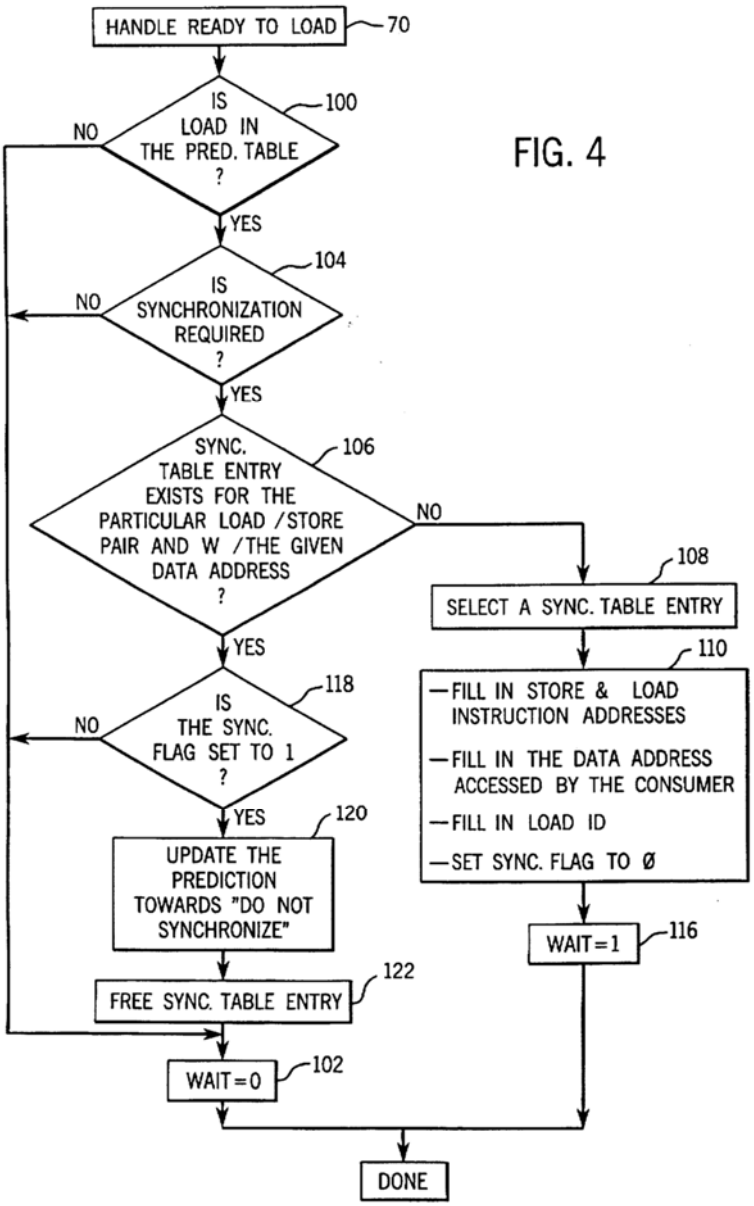

FIG. 4

Sheet 4 of 7

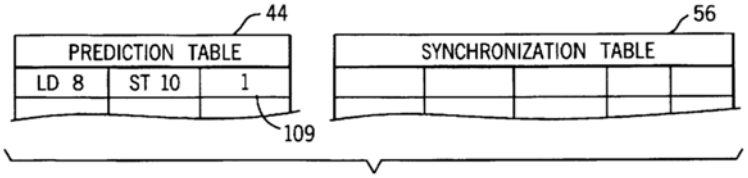

FIG. 5

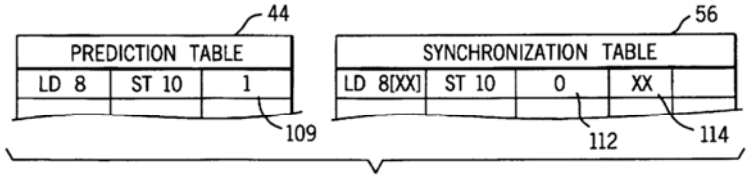

FIG. 6

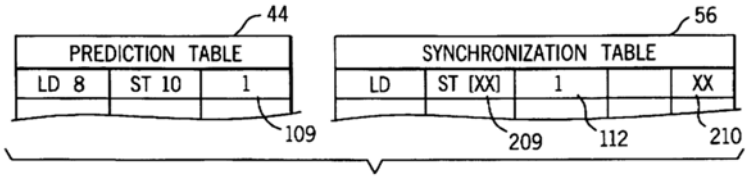

FIG. 8

Sheet 5 of 7

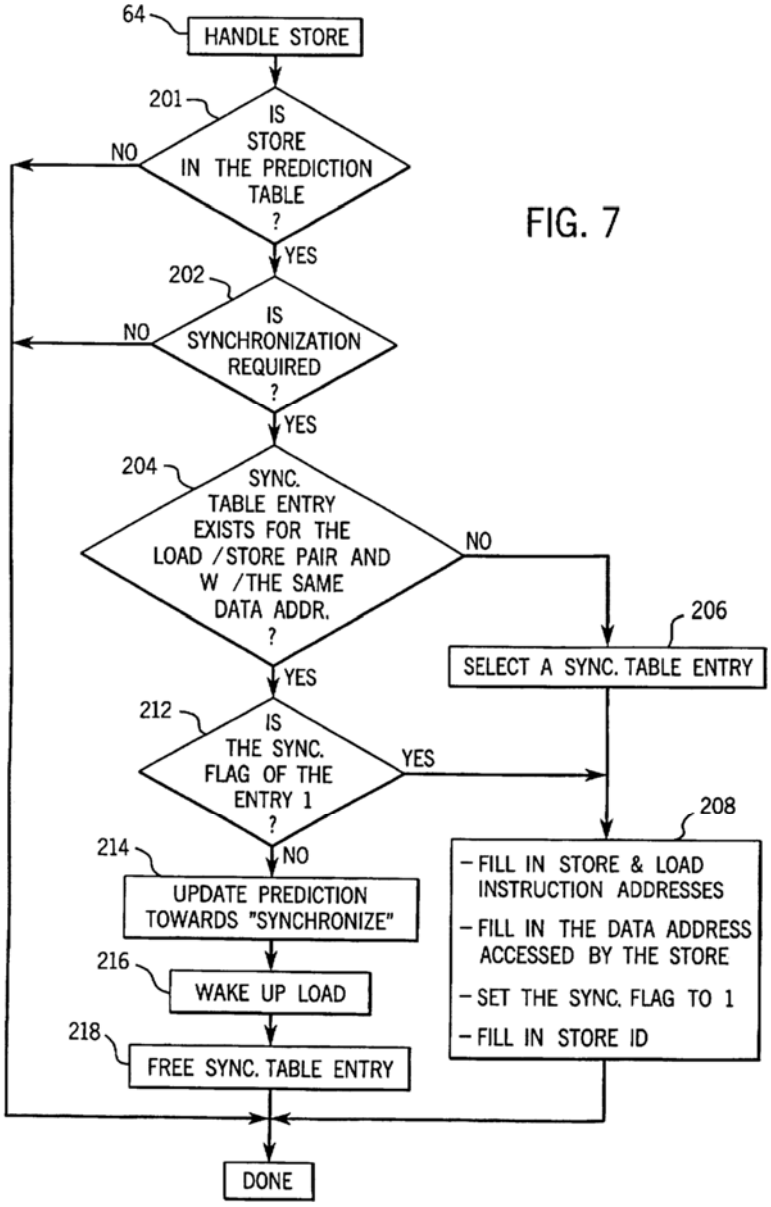

FIG. 7

Sheet 6 of 7

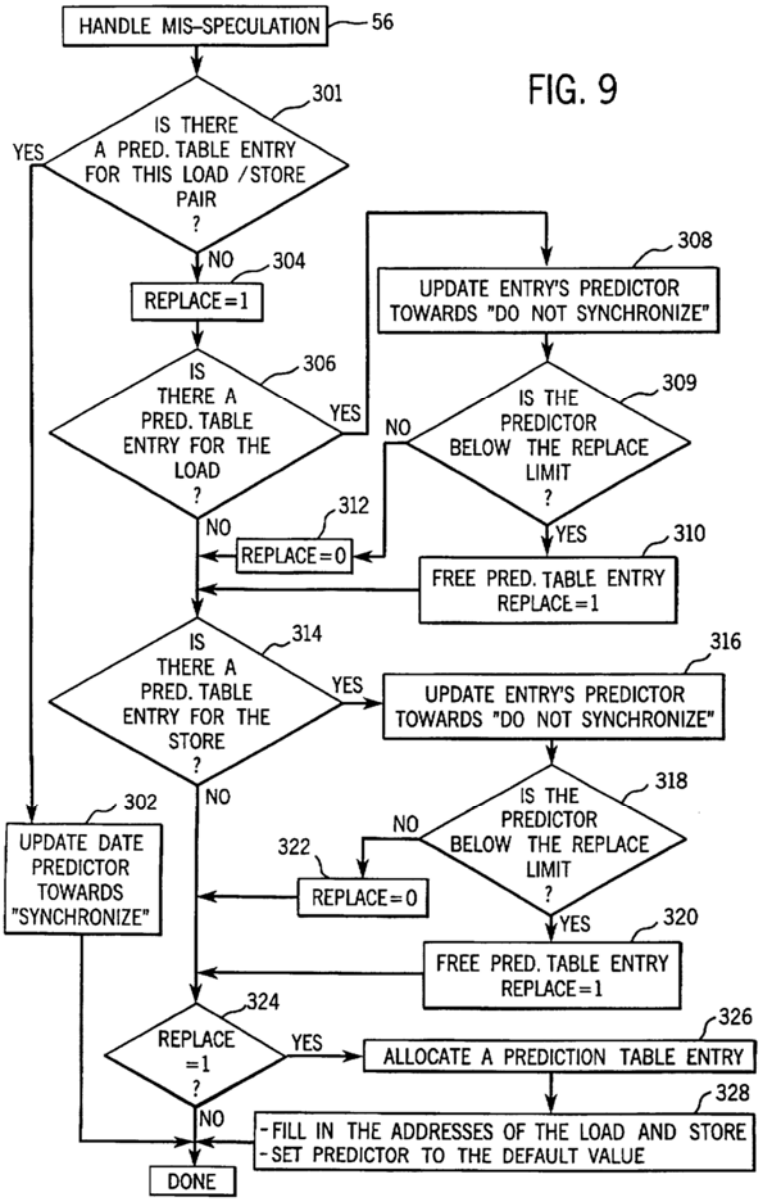

FIG. 9

Sheet 7 of 7

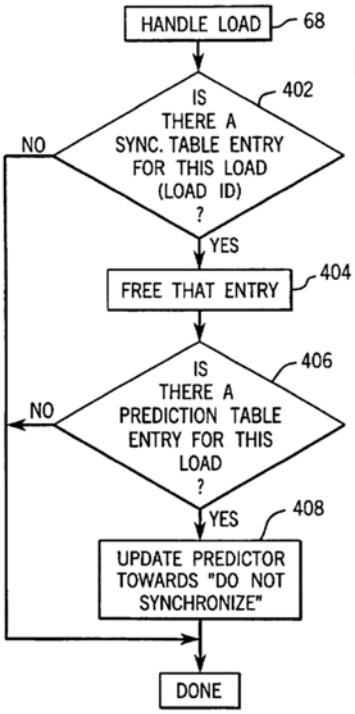

FIG. 10

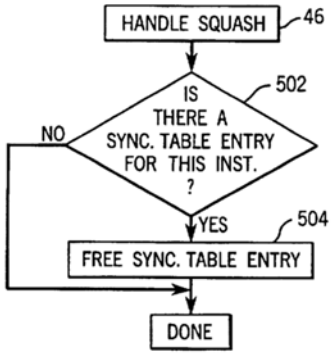

FIG. 11

TABLE BASED DATA SPECULATION CIRCUIT FOR PARALLEL PROCESSING COMPUTER

This invention was made with United States government support awarded by the following agencies:

ARPA Grant No. DABT63-95-C-0127;

ONR, Grant No. N00014-93-1-0465; and

NSF, Grant Nos.: CCR-9303030 and MIP-9505853.

The United States has certain rights in this invention.

FIELD OF THE INVENTION

The invention relates generally to architectures of electronic computers and specifically to electronic computers for parallel processing.

BACKGROUND OF THE INVENTION

General Computer Architecture and Instruction Level Parallel (ILP) Processing

In an electronic computer with a single processing unit, the processing unit may communicate with a memory holding the data and program instructions. The processing unit may also hold data in internal registers. The program instructions are executed by the processing unit which operates on the data according to the instructions.

Typically, the processing unit repeats a series of fetch/execute cycles in which each instruction is first fetched from memory and then executed. The order in which the instructions are executed is determined by

the value of a program counter within the processing unit. After each execution of an instruction, the program counter normally increases in value by one so that the next instruction memory is fetched. The order of the instructions in the stored program will be termed "memory order".

Some instructions, when executed, cause data to be loaded into the processing unit from memory or stored from the processing unit to memory. Other instructions may perform their operation on data that are stored in registers without loading or storing data from or to memory. Still other instructions change the value of the program counter permitting the processing unit to jump or branch through the program instructions according to a "program order" that normally differs from the memory order. The branch or jumps in a program may be conditional on values of data used by the program that may change as the program is executed.

One method of increasing the speed of electronic computers involves using multiple processing and/or functional units to execute multiple instructions at the same time or in an "execution order" differing from the program order. Computers using this technique are termed "parallel processing units". An instruction level parallel ("ILP processing unit") is one where individual instructions of a single program are separated to be run on different processing units in contrast to systems in which independent programs may be assigned different processing units, for example.

Control and Data Dependencies

There are two types of dependencies exhibited by instructions. "Control dependency" is a dependency of instructions after a conditional branch or jump instruction on whether the branch or jump was taken. For example, instructions immediately after a branch are properly executed only if the branch is not taken. "Data dependency" is a dependency of instructions that use data on earlier instructions that change the data. These latter instructions may correctly execute only if the earlier instructions using the same data do not change the common data or have completed the change of the common data. A dependency is "unambiguous" if it necessarily produces an error when the dependent instruction is executed before the instruction on which it is dependent. The dependence of an instruction may remain ambiguous until both instructions involved are executed. A dependency between two instructions is ambiguous if it cannot be determined whether a dependency really exists without executing the instructions.

Usually, instructions that are data dependent must be executed in the program order.

Control and Data Speculation

Since dependencies are often ambiguous, as the ILP processing unit prepares to execute an instruction, it cannot always determine if the instruction will in fact be dependent on earlier instructions that have not yet completed their execution. In the case of ambiguous dependencies and unless special circuitry is provided as will be explained in the next paragraph, the ILP processing unit is forced to assume dependencies exist.

However, it is quite often the case that an ambiguous dependency is resolved as no dependency at all. For this reason, some ILP processors may provide for "speculation", that is, execution of an instruction that has ambiguous dependency as if it had no dependency at all. One may speculate on control dependencies and on data dependencies. Control speculation, for example, might involve executing an instruction that follows a branch instruction without knowing the outcome of the branch (and thus whether the following instruction should have been executed or was branched around). Data speculation, for example, might involve reading from memory to obtain data for a later instruction, even though there are earlier STORES to that memory location that have not yet been completed and that may change the value of the memory location.

Squashing

Control and data dependencies are important in an ILP processing unit which, in the course of execution of instructions, may execute some dependent instructions before the instructions on which they are dependent. If the dependency is unambiguous, then the results of the prematurely executed dependent instructions must be discarded ("squashed") and the instructions of the correct branch executed.

Squashing instructions is a time consuming process that to some extent defeats the advantages to be gained from parallel processing. In order to avoid the problems associated with unambiguous

dependencies, it is necessary that when the ILP processing unit speculates, that it is mostly correct.

Speculation in an ILP Processor

In an ILP processor, the processor may fetch multiple instructions at a single time and an allocation circuit allocates those instructions to separate processing units. The separate processing units may read data from memory and perform arithmetic or logical operations on that data. A data speculation circuit may exist to detect data dependencies and report any mis-speculation to a retirement circuit. The retirement circuit collects results generated by the independent processing units and "retires" the instructions executed by those processing units by writing final results to memory. The retirement circuitry also resolves mis-speculation detected by the data speculation circuit, that is, instructions that were executed out of program order but were in fact dependent on an earlier instruction in the execution order and have produced an erroneous result. These instructions are squashed as is understood in the art. Data speculation circuitry currently does not decide when to do data speculation. Either all memory accesses are speculated or none at all.

For either type of speculation to be successful (control or data speculation), the performance cost associated with the speculation must be low. The performance cost is a function of the frequency that speculation is required, the time required to perform the speculation, the probability of mis-speculation and the time required to recover from a mis-speculation. A cumbersome or inaccurate speculation

system may hurt overall system performance, especially when many dependencies must be evaluated.

Because the ILP processing is relevant primarily for high speed processing units, the speculation process must be implemented in circuitry rather than software. The process of speculation must be capable of rapid execution using limited amounts of high speed memory.

Control speculation is well understood in the art. The points were control speculation is needed are clearly identified by conditional branch and jump instructions (we will refer to these instructions as "control transfer instructions"). Typically all control transfer instructions are speculated since: (1) control transfer instructions are relatively frequent, (2) relatively few instructions from those that appear between two consecutive control transfer instructions can be executed in parallel, (3) typically the performance of a processor that always mis-speculates on control is virtually the same as the performance of a processor that never speculates on control.

In contrast the points where data speculation is needed are not clear since any instruction loading data from memory can be data dependent on any previous instructions that writes to memory. Consequently, predicting and tracking data dependencies, "data dependence speculation" can easily become overwhelming. Furthermore, the cost of a data mis-speculation typically cannot be neglected, that is the performance of a processor that always

mis-speculates on data dependencies is far less than that of a processor that never speculates on data.

BRIEF SUMMARY OF THE INVENTION

The present inventors have recognized that most data dependence mis-speculations can be attributed to a few static STORE/LOAD instruction pairs. Furthermore, these instruction pairs exhibit "temporal locality" that is, if one LOAD/STORE pair causes a data mis-speculation at a given point in time, it is highly likely that a later instance of the same pair will soon cause another mis-speculation. For this reason, a table based approach may be used in which data dependent instructions likely to be a source of mis-speculation, are stored in a small, high speed memory. These particular instruction pairs can be identified based on previous mis-speculations.

Very low overhead in a table based data dependence speculation is obtained by a three-tiered approach. If there is no history of data mis-speculation, an instruction is executed without further inquiry. This will be the case for most data independent instructions. If there has been a mis-speculation with a given LOAD instruction, a predictor based on the past history of mis-speculations for that LOAD instruction is employed to determine whether the instruction should be executed or delayed. Thus, instructions that are typically not dependent may be executed immediately. If the instruction is to be delayed, a synchronization table is used to determine when the instruction is to be performed.

Specifically, the present invention provides a speculation decision circuit for use in a processor

capable of executing program instructions in an execution order differing from the program order of the instructions, the processing unit further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a data consuming instruction dependent for its data on a data producing instruction of earlier program order is, in fact, executed before the data consuming instruction. The speculation decision circuit includes a predictor circuit receiving the mis-speculation signal from the data speculation circuit to produce a prediction associated with the particular data producing/consuming instruction pair and based on the mis-speculation indication. A prediction threshold detector prevents data speculation for instructions having a prediction value within a predetermined range. This prediction threshold detector may include an instruction synchronizing circuit instructing a processing unit to delay a later execution of the particular data consuming instruction until after the execution of the particular data producing instruction when the prediction associated with the data producing/consuming instruction pair is within a predetermined range.

Thus, it is one object of the invention to provide a predictor circuit that may identify data dependencies on an on-going or dynamic basis. Recognizing that there are relatively few instructions which will cause data mis-speculations, these instructions are identified by reference to historical mis-speculation associated with the instructions as stored in a prediction.

The instruction synchronization circuit may include a prediction table listing certain data

consuming instructions and certain data producing instructions each associated with a prediction. The instruction synchronization circuit will delay the particular data consuming instruction only when the prediction associated with the data consuming instruction is within a predetermined range of predictions and when the particular data consuming instruction is in the prediction table.

Thus, it is another object of the invention to provide a prediction of data dependence which adds very little overhead to the execution of instructions that historically have not resulted in mis-speculation. If the particular data consuming instruction is not in the prediction table, no further inquiry into the prediction table is required.

The instruction synchronization circuit may also include a synchronization table associating certain data consuming instructions and certain data producing instructions, each with a flag indicating whether the respective data producing instruction has been executed. The instruction synchronization circuit delays the subsequent instances of the certain data consuming instruction only when the prediction associated with the data consuming instruction is within a predetermined range and when the particular data consuming instruction is in the prediction table and when the flag indicates that particular data producing instruction has not been executed.

Thus, it is another object of the invention to provide reduced instruction overhead even for instructions that have a history of mis-speculation when it is unlikely that, in the given instance, mis-

speculation will occur. If the prediction indicates that mis-speculation is unlikely, no further synchronization steps need be taken.

It is yet another object of the invention to allow data consuming instructions that are historically dependent on preceding data producing instructions to be executed rapidly if the preceding data producing instruction has already been executed. Thus, if the flag indicates that the data producing instruction has been executed, the data consuming instruction may be immediately executed without further waiting.

The predictor circuit may create an entry in this synchronization table only after mis-speculation has occurred for a previous instance of the particular data consuming instruction and particular data producing instruction of the entry. After synchronization has occurred, this entry may be removed.

Thus, it is another object of the invention to provide a predictor circuit that minimizes the need for storage. Synchronizing table entries, which may be more numerous than the prediction table entries as a result of possible multiple instances of each instruction, have entries in the synchronization table that are dynamically created and released to minimize storage requirements. This also minimizes search overhead in identifying instructions within this table.

The foregoing and other objects and advantages of the invention will appear from the following description. In this description, references made to the accompanying drawings which form a part hereof and in which there is shown by way of illustration a preferred embodiment of the invention. Such

embodiment does not necessarily represent the full scope of the invention, however, and reference must be made therefore to the claims for interpreting the scope of the invention.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

FIG. 1 is a block diagram showing the architecture of an instruction level parallel processor having multiple processing units, an allocation unit allocating instructions of a program in memory to the processing units, and a retirement unit retiring the parallel processed instructions, the latter unit incorporating the predictor circuit of the present invention;

FIG. 2 is a fragmentary listing of a source code program and its corresponding object code instructions stored in memory order at physical addresses in memory and as unwound in program order producing multiple instances of each object code instruction in execution;

FIG. 3 is a flow chart of the operation of a typical data speculation circuit of the processor of FIG. 1 as modified to work with the predictor circuit of the present invention to provide READY TO LOAD, HANDLE STORE, HANDLE MIS-SPECULATION, HANDLE LOAD, and HANDLE SQUASH messages to the predictor circuit of the present invention;

FIG. 4 is a flow chart showing the steps performed by the predictor circuit of FIG. 1 upon receiving a READY TO LOAD message from the data speculation circuit;

FIG. 5 is fragmentary view of prediction and synchronization tables used by the predictor circuit of FIG. 1;

FIG. 6 is a figure similar to that of FIG. 5 showing the prediction and synchronization tables at a later point in the operation of the predictor circuit;

FIG. 7 is a flow chart showing the steps performed by the predictor circuit of FIG. 1 upon receiving a HANDLE STORE message from the data speculation circuit;

FIG. 8 is a figure similar to that of FIGS. 5 and 6 showing the prediction and synchronization tables at yet a later point;

FIG. 9 is a flow chart showing the steps performed by the predictor circuit of FIG. 1 upon receiving a HANDLE MIS-SPECULATION message from the data speculation circuit;

FIG. 10 is a flow chart showing the steps performed by the predictor circuit of FIG. 1 upon receiving a HANDLE LOAD message from the data speculation circuit; and

FIG. 11 is a flow chart showing the steps performed by the predictor circuit of FIG. 1 upon receiving a HANDLE SQUASH message from the data speculation circuit.

DETAILED DESCRIPTION OF THE INVENTION

An Example Program with Data Dependency Referring now to the following Table I, an example object code program may have three instructions: I1, I2 and I3.

TABLE I

I1	st $M(R_1)$
I2	ld $M(R_2)$
I3	ld $M(R_3)$

These instructions provide for a storage of data from a processor to a memory location and the loading of data from the memory location to the processor.

The first instruction I1 stores data from a processor register (implied in the instruction) to memory at a specific address determined by the contents of register R_1 of the processing unit. Generally, the contents of register R_1 will have been set by an earlier instruction and cannot be deduced simply by looking at instruction I1.

The following instructions I2 and I3 each load data from a specific physical address of memory to the processor register. Again, the registers R_2 and R_3 have been set by earlier instructions and the relevant addresses cannot be determined from these instructions alone.

If a parallel processor were to load the three instructions I1–I3 for parallel processing, there would be two possible data dependencies. I2 may be dependent on I1 insofar that the contents of register R_1 and R_2 may be the same and therefore the information stored in $M(R_1)$ may be the same that is loaded from memory location $M(R_2)$. For instruction I2 to be executed correctly, it may be necessary that instruction I2 wait for instruction I1. Likewise, instruction I3 may be dependent on instruction I1. That is, the address $M(R_3)$ may be the same as the physical address $M(R_1)$.

Each of these dependencies is ambiguous. That is, it cannot be determined whether there is in fact a dependency without knowing the contents of registers R1, R2 and R3 which cannot be deduced from the instructions alone.

In a parallel processor that seeks to load blocks of instructions and execute them in parallel, possibly out of their memory or program order, it is important to identify and address ambiguous data dependencies to eliminate mis-speculation and time consuming instruction squashing.

General Processor Architecture

Referring now to FIG. 1, an ILP processor 10 suitable for use with the present invention includes a memory 12 having a portion 14 holding a stored program 16 at a plurality of physical addresses 19 here depicted as xx1–xx6 where the values xx indicate some higher ordered address bits that may be ignored in this example.

The data speculation circuit 30 receives signals from the allocation circuit 20 that notify it of the program order of any instructions that are allocated to the processing units 24 and that will access memory. The data speculation circuit is responsible of keeping track of order of the memory operations as they are performed by the processing units so that it can detect any mis-speculations.

The ILP processor 10 includes an allocation circuit 20 which may read memory 12 and in particular program memory portion 14 to fetch a subset of the program 16 encompassing multiple instructions of an instruction window 22. Generally, as is understood in the art, the allocation circuit 20

sends different ones of these instructions to different independent processing units 24 for execution.

The processing units 24 may communicate with memory 12 for the purpose of obtaining data, but generally do not modify portions of the memory 12 that may be also read by other processing units 24. Thus, for example, an instruction which requires data to be obtained from memory 12 and added to a constant may be completely executed. However, an instruction, which stores a register to memory 12, may read the register but stop before the store operation which may modify memory 12 used by other processing units 24. An instruction that has been executed as far as possible is considered "ready to commit the operation". Prior to reading data from memory or requesting a store operation the processing units 24 notify the data speculation unit 30 of the operation so that the latter can, in conjunction with the allocation unit, keep track of the program and execution order of the operations.

A retirement circuit 26 receives signals from the processing units 24 indicating that their instructions are ready to perform the operation. The retirement circuit 26 then retires the instructions by writing any computed values to memory 12.

Prior to the retirement circuit 26 writing values to memory, a data speculation circuit 30 communicating with the allocation circuit 20 and the retirement circuit 26 detects mis-speculation. As described above, mis-speculation occurs in an instruction that has been executed prematurely and erroneously. Whenever a store instruction is ready to commit and write its data to a memory address, the

data speculation circuit 30, checks to see if any subsequent in the instruction window load instructions have accessed the same memory address, and if so instructs the allocation circuit 20 and retirement circuit 26 that these load instructions are to be squashed and re-allocated by the allocation circuit 20 at a later time. Thus, for example, in the program Table I, if instructions I1 through I3 represent the instruction window 22 and instruction I3 has accessed memory prior to I1 writing to memory, the data speculation circuit 30 and at the time I1 is ready to commit checks if I3 has accessed the same memory address as I1, and if so it instructions the allocation circuit 20 and the retirement circuit 26 that I3 is to be squashed and reallocated at a later time. If so, the data speculation circuit 30 instructs the allocation circuit and retirement circuit 26 that instruction I3 is to be squashed and must be reallocated by the allocation circuit 20 at a later time. The writing of results from instructions that have not been squashed is then done by the retirement circuit 26 as indicated by arrow 28.

The retirement circuit 26, the data speculation circuit 30, and the allocation circuit 20 communicate by a number of control signals 35. Each of these elements described above are well known in the art.

The data speculation circuit 30 also communicates with the predictor circuit 33 of the present invention. The predictor circuit 33 provides a dynamic indication to the data speculation circuit 30 as to whether data speculation should be performed. The data speculation circuit 30 may then, based on the indication of the predictor circuit 33 stall the

execution of a memory operation at the processing units 24 in order to avoid mis-speculation.

The prediction provided by the predictor circuit 33, as will be described, is updated based on historical mis-speculations detected by the data speculation circuit 30. For this reason, the data speculation circuit 30 must communicate with the predictor circuit 33 on an ongoing basis. Generally, as will be described below, the data speculation circuit 30 provides five signals to the predictor circuit 33 as listed in the following Table II.

TABLE II

Signal Name	Description
HANDLE MIS-SPECULATION	The data speculation circuit has detected a mis-speculation.
HANDLE STORE	The data speculation circuit has decided to issue a STORE operation.
READY TO LOAD	The data speculation circuit is about to perform a speculative LOAD operation and needs information from the predictor as to whether the LOAD should wait.
HANDLE LOAD	The data speculation circuit has decided to perform a non-speculative LOAD operation without data dependency.
HANDLE SQUASH	The data speculation circuit has issued a

squash for a particular instruction either as the result of data or control mis-speculation.

Generally, prior to each instruction being retired, the instruction is provided to the data speculation circuit 30 which detects mis-speculations. The retirement circuit 26 also instructs the data speculation circuit 30 when a squash instruction is required.

Referring to FIG. 2, an example source code program 18 includes two lines as follows:

for i=0, i<N, i++

A[i+1]=A[i]*19

These instructions represent typical source code instructions 18 which may create a data dependence speculation problem. The source code instructions 18 describe a loop executed N times in which each element of an array variable A(i) receives a value of 19 times that of its preceding array value.

The source code instructions 18, when reduced to machine language instructions 32, involve repeated LOAD and STORE operations. That is, to realize the second line of the source code instructions 18, the memory value of the array element A[i] must be: (1) loaded, (2) multiplied by 19, and (3) stored in memory location for array element A[i+1]. The LOAD and STORE instructions of the machine language instructions 32 have logical addresses within the memory 12 of xx8 and xx10, respectively, which may be used to identify the instructions uniquely.

As the machine language instructions 32 execute in a loop, they create an execution thread 34 of successive LOAD and STORE operations in which the variable 'i' is incremented from 0 to N. This execution thread 34 produces multiple instances of the LOAD and STORE instructions. For illustration purposes, each instance may be uniquely identified and is designated in FIG. 2 by an integer to the right of a decimal point forming a suffix to the physical address of the operation. Thus, the first instance of the LOAD instruction of the execution thread 34 is designated 8.1, the physical address of the LOAD instruction being 8 and the instance being 1. Likewise, the first STORE instruction is designated 10.1 and so forth.

The LOAD instruction 8.1 loads the contents of memory location A[1]. The STORE instruction 10.1 then stores a value in memory location A[2]. The LOAD instruction 8.2 then loads a value from the same memory location A[2]. This LOAD instruction 8.2 is thus unambiguously dependent on instruction 10.1, the data dependency 36 as indicated by an arrow. If instruction 8.2 were to be executed prior to instruction 10.1, it would operating on erroneous data.

Likewise, the LOAD instruction 8.3 which loads data from memory location A[3] is dependent on the STORE instruction 10.2 which stores data in that same memory location. This unambiguous dependence 38 is indicated by an arrow.

If the ILP processor 10 is to process the machine code instructions out of execution order, the data dependencies 36 and 38 must be observed to ensure that the out of order processing instruction 8.2 does

not execute before instruction 10.1 and instruction 8.3 does not execute prior to instruction 10.2.

The circuitry necessary to construct an ILP processor will typically involve many millions of transistors, however, its operation can be described by means of a description of a flow chart and the data structures modified during the steps described. Implementing the device described below in discrete circuitry, although laborious, is a routine engineering matter whose practice will be well understood to those of ordinary skill in the art. Generally the data structures described below will be implemented in solid state memories and the steps of the flow chart will be incorporated into a state machine of a type well understood in the art.

Operation of the Data Speculation Circuit

Referring now to FIG. 3, the normal operation of the data speculation circuit 30 such as is known in the prior art, must be modified slightly to accommodate the present invention. These modifications provide the necessary signals referred to in Table 2 to the predictor circuit 33.

Starting at process block 40 of FIG. 3, a program instruction may be received from the retirement circuit 26 by the data speculation circuit 30 together with an indication that the instruction should be either squashed or is about to execute. At decision block 42, if the instruction is to be squashed, the data speculation circuit 30 provides a HANDLE SQUASH signal to the predictor circuit 33 as indicated by process block 46. Otherwise, the program proceeds to process block 48 which determines whether the

instruction, which must be then assumed to be ready to issue is a LOAD or STORE instruction.

If the instruction is a STORE instruction, the program proceeds to process block 50 and a STORE request is issued. This STORE request may, for example, authorize the retirement circuit 26 to perform the STORE operation for the data.

At decision block 52, the data speculation circuit 30 checks other concurrent LOAD instructions to see if they have been prematurely executed and thus whether there has been a mis-speculation. At process block 57 if there has been a mis-speculation, a HANDLE MIS-SPECULATION signal is sent to the predictor circuit 33. This signal is used by the predictor circuit 33 in adjusting its prediction 109 as will be described. The dependent instructions are then squashed by the loop of decision block 60 and process block 62.

If at decision block 52 there was no mis-speculation, then the data speculation circuit 30 sends a HANDLE STORE signal to the predictor circuit 33 as indicated by process block 64 as will be described below.

If at decision block 48 the instruction received by the data speculation circuit 30 is a LOAD instruction, then at decision block 66 it is determined whether this is a data speculative LOAD, that is whether there are prior STORE instructions on which it might depend. If the answer is no, then the data speculation circuit 30 provides a HANDLE LOAD signal to the predictor circuit 33 as indicated by process block 68 as will be described below. Otherwise, the data speculation circuit 30 provides a HANDLE READY TO LOAD

signal to the predictor circuit 33 as indicated by process block 70 as will also be described below.

The predictor circuit 33 will address the READY TO LOAD request from the data speculation circuit by making a prediction as to whether the LOAD should take place through the use of a wait flag. Thus, at subsequent decision block 72 if the wait flag equals 1 indicating that speculation should occur, the program proceeds immediately to process block 74 and a LOAD request is generated.

Next at process block 76, the data speculation circuit 30 waits for that particular instruction, either to be squashed indicating that it had been erroneously speculated or for an indication that it is no longer data speculative, that is any previous STORE instructions were for different memory addresses. At decision block 78, if the condition of case block 76 was that the instruction was squashed, the data speculation circuit 30 proceeds to the handle squash block 46 as is previously described.

At decision block 72, the predictor circuit 33 may have indicated that data speculation was not appropriate by providing that the wait flag equal 1. In this case, a LOAD request is not issued but instead, at case block 80, the data speculation circuit 30 waits for a wakeup signal indicating that the dependent STORE instruction has been executed or a signal indicating that the LOAD instruction is no longer speculative because the earlier STORE instructions did not write to its memory location or for a squash signal indicating that the instruction should be squashed as a result of a later occurring control or data dependency mis-speculation. At decision block

82, if the condition of process block 80 was that of being squashed, the program branches to the handle squash process block 46. If not, then at decision block 84 if the event was a wakeup signal, the program branches to previously described process block 74 where the LOAD request is issued. If the triggering event was that the LOAD instruction is no longer data speculative, then at process block 68, the HANDLE LOAD signal is provided by the data speculation circuit 30 to the predictor circuit 33 and the LOAD request is issued as indicated by process block 86.

Operation of the Predictor

1. Handle Ready to Load

Referring also to FIG. 1, after instructions of instruction window 22 have been received by the allocation circuit 20 and allocated to the processing units 24, the processing units 24 will begin execution of the instructions and at various times certain instructions will become ready for their operations to be performed and will be received by the retirement circuit 26 and the data speculation circuit 30.

As an example, assume that instruction 8.2 in FIG. 2 becomes ready for its operation to be performed. Referring now to FIG. 3, instruction 8.2 is a LOAD instruction and thus will cause a HANDLE READY TO LOAD signal to be passed to the predictor circuit 33. Referring to FIG. 4, at decision block 100, the predictor circuit 33 reviews a prediction table 44 shown generally in FIG. 5 to see if the particular instruction 8.2 identified by its physical address is in the prediction table 44.

The prediction table 44 includes three elements for each entry, the elements depicted as columns, the

entries as rows. The first column identifies, by physical address, an instruction that is ready for its operation to be performed; the second column identifies the instruction on which the instruction in the first column may be data dependent; and the third column holds a prediction 109 as will be described.

For the purpose of this example, it will be assumed that the prediction table 44 is initially empty. In this case at decision block 100, the operation of the predictor circuit 33 is to set the wait flag equal to zero as indicated by process block 102 and return that flag value to the data speculation circuit 30 as has been described. Generally, the prediction table 44 has an entry if there has been a historical pattern of mis-speculation and thus at least one mis-speculation. Accordingly, if no entry is found in the prediction table 44, the reasonable assumption is that speculation can proceed. How the prediction table 44 is loaded will be described below.

If there is an entry in the prediction table 44 at decision block 104, then the prediction 109 of the prediction table 44 is examined to see if it indicates that it is likely there is data dependence of this instruction. The higher the prediction 109, the more likelihood of mis-speculation if the instruction of the first column is executed before the instruction of the second column. Normally the prediction 109 starts at zero when an entry is first made in the prediction table 44 and is incremented and decremented as will be described below. At decision block 104, if synchronization is not required as indicated by the prediction 109, the program proceeds to process block 102. On the other hand, if the prediction 109 indicates that there is a likelihood of mis-speculation, the

program proceeds to process block 106 and a synchronization table 56 is examined.

The synchronization table 56 is generally similar to the prediction table 44 but whereas the prediction table 44 compiles a prediction statistic as to whether data dependence exists between a LOAD/STORE pair, the synchronization table 56 indicates whether there is in fact a pending LOAD instruction awaiting its dependent STORE instruction.

Assuming that there is no entry in the synchronization table 56 as shown in FIG. 5, then an entry must be added to reflect the fact that there is a pending LOAD instruction that must wait for its preceding STORE instruction. At process block 108, a new row of the table is selected and at process block 110 an entry is created. In this creation of an entry, the STORE instruction address is filled in as indicated in FIG. 6. The address of the LOAD instruction is also stored. A synchronization flag 112 is set equal to zero indicating that the STORE instruction with which this LOAD instruction must be synchronized has not occurred and a LOAD identifier 114 is entered providing a unique number for this instance of the LOAD instruction used for tracking the instruction and generated arbitrarily with the only requirement that it be unique to this particular instruction.

At process block 116, the wait flag is set. This is the same wait flag that is used by the data speculation circuit 30 as shown in process block 72 of FIG. 1.

Referring again to decision block 106, if there is an entry in the synchronization table 56 then the program proceeds to decision block 118 where the

synchronization table 56 is checked to see if the synchronization flag 112 is equal to 1. If so, the program branches to process block 102 and sets the wait flag equal to zero which causes the data speculation circuit 30 to go ahead and issue a LOAD request. This is a situation where another instance of the LOAD instruction is in fact waiting for the same predicate STORE instruction and the decision is made simply to release the current LOAD instruction.

On the other hand, at decision block 118, if the synchronization flag 112 is set to 1, that indicates that the predicate STORE instruction has already occurred and the LOAD instruction may be released as no longer being data dependent. In this case, the prediction that there was a need to synchronize was wrong and so at process block 120 the prediction 109 is decremented toward the do not synchronize state. Finally, at process block 122 the entry of the synchronization table 56 is erased as no longer needed so as to preserve room in the synchronization table 56.

2. Handle Store

Referring now to FIG. 7, the initial stages of the HANDLE STORE routine of process block 64 are similar to that of the HANDLE READY TO LOAD. That is at a process block 201, the prediction table 44 is checked to see if the STORE instruction is in the prediction table 44. If not, the program exits, but if so at decision block 202, the prediction 109 is checked to see if synchronization is required between this instruction and another LOAD instruction. If not, again the program exits but if so, at decision block 204 the synchronization table 56 is checked to see if a previous dependent LOAD instruction is awaiting

execution. If not, at process block 206, a new entry is added to the synchronization table 56 and at process block 208 that new entry is loaded with the STORE and LOAD instruction addresses, the data address 209 of the STORE instruction (as indicated in FIG. 6), the synchronization flag 112 is set to 1, and a STORE ID 210 identifying uniquely that STORE instruction is added to the table. The program then exits again, the STORE instruction having been executed.

At decision block 204, if a synchronization table entry is present, then at decision block 212 the synchronization flag 112 is checked. If the synchronization flag is 1, indicating that entry exists already indicating that a predicate STORE instruction has been enrolled (but never claimed by a dependent LOAD instruction), the program proceeds to process block 208 and that entry is replaced with the data from the present STORE instruction.

More typically, the synchronization flag 112 will be zero indicating that there is a pending LOAD instruction. In this case, the prediction 109 is updated toward the synchronize condition indicating that the prediction that there was a need to synchronize was correct as there is in fact a LOAD instruction waiting to be synchronized.

The LOAD instruction is released at process block 216 and the entry in the synchronization table 56 is erased at process block 218.

3. Handle Mis-Speculation

Referring now to FIG. 9, the predictor circuit 33 must also make adjustments in its prediction table 44 if there is a mis-speculation, an occurrence that provides considerable information as to whether

synchronization is required. At an initial decision block 301, the prediction table 44 is checked to see whether the LOAD/STORE pair causing the mis-speculation is in the prediction table 44 already. If so then at process block 302, the prediction 109 is updated toward synchronize so that this mis-speculation may be avoided in the future. If not, a 'replace' flag is set equal to 1 at process block 304 and the program proceeds to decision block 306 and the prediction table 44 is again examined but this time for an entry having only the LOAD instruction.

Such an entry indicates a possible lack of temporal locality, that is, indicates that there are different instances of the instructions having different data dependencies that can make the prediction wrong. It is desired in this case to neutralize these table entries, but slowly, as it cannot be determined which particular instance represents the better prediction. Accordingly at process block 308, the prediction 109 is moved toward the do not synchronize direction or decremented and if it is below a limit as determined at process block 309, the table entry is deleted at process block 310 and the value of the flag replace is set equal to 1.

In cases where the prediction 109 is still above the limit, then the program proceeds to process block 312 and the replace flag is set equal to zero. In all cases, the program next proceeds to decision block 306 and a similar inquiry is made for the STORE instruction, that is whether a prediction table entry exists having only the STORE instruction. As with process block 308 to 312, the same steps are taken to eliminate possibly erroneous predictions. That is at process block 316, the prediction 109 is moved toward the do

not synchronize value at decision block 318 if the prediction 109 is below a certain limit, the program proceeds to process block 320 and the prediction table entry is erased and the replace flag set equal to 1.

If the prediction 109 is not below the limit, then replace is set equal to zero at process block 322 and the program proceeds to decision block 324, the same destination being arrived at from decision block 314 if there is no entry for either the STORE or LOAD instructions at decision blocks 306 and 314.

At decision block 324, if the replace flag equals 1 indicating that there are absolutely no prediction table entries left that match either one of the instructions involved in this mis-speculation, then at process block 326 a prediction table entry is allocated and at process block 328 the addresses of the LOAD and STORE instructions are inserted in the prediction table 44 and the prediction 109 is set to the default value, typically zero. If the value of the replace flag did not equal the 1 indicating that there was an entry, no further processing is needed.

4. Handle Load

Referring now to FIG. 10, the HANDLE LOAD routine of process block 68 is relatively straightforward and first examines, at decision block 402, whether there is a synchronization table entry for this particular LOAD instruction which has now been released for execution. If so at process block 404 that entry is erased. Next at decision block 406, prediction table 44 is examined for this particular LOAD instruction. If there is an entry in the prediction table 44, then the prediction 109 is

decremented toward a do not synchronize condition at process block 408.

5. Handle Squash

Referring now to FIG. 11, the predictor circuit 33 must also receive the handle squash message of block 46 with regard to updating the synchronization table 56. Specifically, at a decision block 502, if there is a synchronization entry, that synchronization entry is eliminated at process block 504.

As will be understood from this description, the prediction 109 is used to determine the likelihood of a dependency between two instructions in the future. The higher the prediction 109 the more likelihood of mis-speculation if the instruction in the first column is executed before the instruction in the second column. It will be understood that the prediction 109 may be obtained by methods other than simply incrementing it in value for each speculation as is described herein. For example, various weighting schemes can be provided to cause the predictor circuit 33, for example, to be less sensitive to the earliest mis-speculations. More complex pattern matching techniques may be also used, for example, to catch situations where mis-speculations occur in groups or regular patterns.

The present inventors believe that a relatively limited number of LOAD/STORE pairs will create mis-speculation and so the operation described above prevents the majority of the LOAD/STORE pairs from being slowed in execution. The list of critical LOAD/STORE pairs is prepared dynamically in a synchronization method for those LOAD/STORE pairs is created that can expand or shrink depending

on the operation of the program. It should be noted that the present invention may be used in any processor where execution of instructions deviates from the program order, for example, processing units that during execution of time consuming instructions may move to other instructions out of order to begin their execution.

The above description has been that of a preferred embodiment of the present invention. It will occur to those that practice the art that many modifications may be made without departing from the spirit and scope of the invention. In order to apprise the public of the various embodiments that may fall within the scope of the invention, the following claims are made:

We claim:

1. In a processor capable of executing program instructions in an execution order differing from their program order, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a data consuming instruction dependent for its data on a data producing instruction of earlier program order, is in fact executed before the data producing instruction, a data speculation decision circuit comprising:

a) a predictor receiving a mis-speculation indication from the data speculation circuit to produce a prediction associated with the particular data consuming instruction and based on the mis-speculation indication; and

b) a prediction threshold detector preventing data speculation for instructions having a prediction within a predetermined range.

2. The data speculation decision circuit of claim 1 wherein the prediction threshold detector includes an instruction synchronization circuit instructing the processor to delay a later execution of the particular data consuming instruction until after the particular data producing instruction when the prediction associated with the data consuming instruction is within the predetermined range.

3. The data speculation decision circuit of claim 2 wherein the instruction synchronization circuit includes a prediction table listing certain data consuming instructions and certain data producing instructions each associated with a prediction and wherein the instruction synchronization circuit delays the particular data consuming instruction only:

i) when the prediction associated with the data consuming instruction is within a predetermined range; and

ii) when the particular data consuming instruction is in the prediction table.

4. The data speculation decision circuit of claim 3 wherein the certain data consuming and data producing instructions are identified in the prediction table only by the address of the instructions in a program memory.

5. The data speculation decision circuit of claim 2 wherein the instruction synchronization circuit includes a synchronization table associating the certain data consuming instructions and the certain data producing instructions each with a flag value indicating whether the respective certain data producing instruction has been executed and wherein

the instruction synchronization circuit delays the particular data consuming instruction only:

i) when the prediction associated with the data consuming instruction is within a predetermined range; and

ii) when the particular data consuming instruction is in the prediction table; and

iii) when the flag indicates the particular data producing instruction has not been executed.

6. The data speculation decision circuit of claim 2 wherein the instruction synchronization circuit creates an entry in the synchronization table including the particular data consuming instructions and data producing instructions and the flag value only after a mis-speculation indication is received for the particular data consuming instruction and the particular data producing instruction.

7. The data speculation decision circuit of claim 5 wherein when the flag indicates the particular data producing instruction has been executed, the instruction synchronization circuit removes the entry from the synchronization table.

8. The data speculation decision circuit of claim 1 wherein the prediction produces the mis-speculation indication by tallying the mis-speculation indications for a data consuming instruction.

9. In a processor capable of executing program instructions in an execution order differing from the program order of the instructions, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a data consuming

instruction dependent for its data on a data producing instruction of earlier program order, is in fact executed before the data producing instruction, a data speculation decision circuit comprising:

a) a prediction table communicating with the data speculation circuit to create an entry listing a particular data consuming instruction and data producing instruction each associated with a prediction when a mis-speculation indication is received; and

b) an instruction synchronization circuit only instructing a processor to delay a later execution of the particular data consuming instruction if the prediction table includes an entry.

* * * * *

APPENDIX D – District Court Opinion and Order on Claim Construction and Summary Judgment (August 5, 2015)

UNITED STATES DISTRICT COURT

FOR THE WESTERN DISTRICT OF WISCONSIN

Case No. 14-CV-062-WMC

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff,

-vs-

APPLE, INC.,

Defendant.

Opinion and Order

In this lawsuit, plaintiff Wisconsin Alumni Research Foundation ("WARF") alleges that defendant Apple, Inc. infringes U.S. Patent No. 5,781,752 (the "'752 patent"), which concerns a "table based data speculation circuit for parallel processing computer." Before the court are the parties' cross-motions for summary judgment and claim construction. (Dkt. ##116, 117.) For the reasons that follow, the court will adopt WARF's proposed construction of the term "prediction" and grant summary judgment to WARF on Apple's

counterclaims and defenses based on anticipation under 35 U.S.C. § 102 with respect to U.S. Patent No. 5,619,662 ("Steely" or the "Steely patent"), as well as indefiniteness under 35 U.S.C. § 112 ¶ 2 with respect to claims 5 and 6. In turn, the court will deny Apple's motion for summary judgment based on those same defenses and counterclaims. As for Apple's motion for summary judgment on WARF's claim of willful infringement, the court will deny Apple's motion with respect to any defenses premised on (1) Apple's claim construction, (2) anticipation by Steely, and (3) indefiniteness of claims 5 and 6, but will reserve on Apple's motion in all other respects.

UNDISPUTED FACTS[1]

A. The Parties and Overview of This Lawsuit

Plaintiff Wisconsin Alumni Research Foundation ("WARF") is a Wisconsin corporation, with its principal place of business in Madison, Wisconsin. WARF is the owner of the '752 patent. Defendant Apple, Inc. is a California corporation, with its principal place of business in Cupertino, California.

On January 31, 2014, WARF filed suit against Apple alleging infringement of the '752 patent. Apple answered and asserted counterclaims for declaratory judgment of non-infringement and invalidity of the '752 patent. Material to the present motions, Apple contends that claims 1-3, 5, 6, and 9 of the '752 patent are invalid as anticipated by the "Steely patent. Apple also alleges that claims 5 and 6 of the '752 patent are invalid as indefinite.

---

[1] Except as otherwise noted, for purposes of summary judgment, the court finds the following facts to be material and undisputed.

B. Technology Overview

A modern computer device includes both hardware and software. Hardware typically includes memory, a microprocessor and peripherals, while software typically consists of sequences of instructions or "programs" that run on the hardware. At a general level, the microprocessor is responsible for fetching instructions and data, executing those instructions to modify the data, and then saving the results.[2] Typically, individual instructions call for the performance of a relatively simple task, such as reading a value from or writing a value to a memory location, or adding, subtracting or comparing two numbers. There are generally three types of software instructions: (1) memory instructions; (2) computing instructions; and (3) control instructions.

Memory instructions are instructions that "when executed, cause data to be loaded into the processing unit from memory or stored from the processing unit to memory." ('752 Patent (dkt. #1-1) 1:38-40.) So-called "LOAD" instructions copy or read a value stored at a memory location specified by an address and return a value. LOAD instructions are also called "data consuming instructions," because they consume data by obtaining data from memory, though as Apple cautions, other types of software instructions also "consume" data. "STORE" instructions, on the other hand, copy or write a value to a memory location specified by an address. For that reason, STORE

---

[2] While the court uses the term "executing," the court acknowledges that the parties have agreed on a construction for the term "in fact executed" described below in subheading "E" of this Facts section.

instructions are also called "data producing instructions," as they produce data by providing data to memory. (Apple similarly points out that other types of instructions "produce" data.) When a STORE instruction executes, it overwrites any value previously stored at that memory location. Both LOAD and STORE instructions are memory instructions.

Generally speaking, software instructions in a program have a predefined "program order," where the processor performs the instructions sequentially. Instructions, however, need not always be executed in the listed order. Instead, they may be executed "out of order." In out-of-order executions, instructions are typically executed when ready -- in other words, based on the availability of their input data, or "operands," rather than a specified program order.[3] There are some obvious benefits to permitting instructions to execute out of order. For instance, because some instructions in a program take longer to execute than others, performing instructions in program order may slow processor performance since it requires waiting for earlier instructions to execute before performing later instructions in the program. Out-of-order execution, therefore, may result in increased efficiency since it allows the processor to use free time to execute other instructions that are ready to be processed. On the other hand, out-of-order execution may have a detrimental effect on performance if it leads to errors that require the processor to expend resources to correct.

---

[3] Apple clarifies that out-of-order executions also depend on the availability of necessary hardware.

A key requirement of efficient out-of-order execution, therefore, is that it must yield the same results as would the execution of instructions in program order. This requirement touches on the concept of "instruction dependency." A dependent instruction is one that must wait for the result of an earlier-in-order execution before it can safely execute.[4] For example, data dependency exists when an earlier-in-order STORE instruction writes data to the same address that is accessed by a later-in-order LOAD instruction. In that situation, the STORE and LOAD must execute in program order for the LOAD to read the correct data from the shared memory address that both instructions access.

In some situations, whether a given LOAD instruction depends on a STORE instruction from an earlier earlier-in-order program step cannot be known until after one or both of the instructions are executed. In other words, the processor lacks sufficient information to resolve whether or not a dependency actually exists. This uncertainty is known as "ambiguous dependency." Ambiguous dependencies may occur, for example, when the memory addresses that must be accessed by a given LOAD or STORE instruction are computed "on the fly" as the program executes. In those circumstances, the processor may have to perform additional

---

[4] While "[a] processor may permit dependent instructions to execute out-of-order and then invoke a recovery process to return to a correct machine state," as Apple describes, Apple fails to dispute WARF's point that to execute "safely," the dependent instruction must wait to execute until after the instruction on which it depends has been executed. (Pl.'s Reply to Pl.'s PFOFs (dkt. #157) ¶ 21.)

computations with data that are not currently available in order to resolve whether one instruction is dependent on another.

To maximize processing speed, however, the processor may elect to execute a LOAD instruction before an earlier STORE instruction. The out-of-order execution of instructions without knowing if there is an actual dependency between them is known as "speculation" or "speculative execution," because the processor is speculating that there is no actual dependency. Speculation can be advantageous if it turns out to be correct (i.e., the LOAD instruction in fact was not dependent on the STORE instruction); then the out-of-order execution will yield the correct result and the performance will improve.[5] In contrast, if a LOAD instruction is speculatively executed ahead of a STORE instruction of earlier program order and it turns out that the speculation was incorrect (i.e., the LOAD instruction was in fact dependent on the earlier STORE instruction), then the instructions will

---

[5] The parties dispute whether "it is quite often the case that an ambiguous dependency is resolved as no dependency at all," as the '752 patent represents. ('752 patent (dkt. #1-1) 2:26-27.) As Apple contends, "[t]he degree to which ambiguous dependencies will turn out to be resolved as no dependency depends upon the workload." (Def.'s Resp. to Pl.'s PFOFs (dkt. #141) ¶ 30.) At the same time, Apple also proposes on summary judgment that "[m]any program instructions are 'independent' of each other and can safely execute-out-of-order with respect to each other." (Def.'s PFOFs (dkt. #119) ¶ 24.) Even if this could be construed as a dispute, it is not material to the issues before the court on summary judgment.

cause an error -- the prematurely executed LOAD instruction having obtained incorrect or stale data.[6]

In the patent-in-suit, this error is referred to as "mis-speculation," although the Steely patent -- as described below -- refers to it as a "collision." As discussed generally already, and as the patent-in-suit explains specifically, a mis-speculation can be detrimental to processor performance because it requires "the results of the prematurely executed dependent instructions [to] be discarded" or "squashed," and the instruction will need to be re-executed in program order. ('752 patent (dkt. #1-1) 2:46-49; Def.'s PFOFs (dkt. #119) ¶ 33.)[7]

C. The '752 Patent

i. Overview and Prosecution History

The '752 patent, entitled "Table Based Data Speculation Circuit for Parallel Processing Computer," was filed on December 26, 1996, and issued on July 14, 1998. The listed inventors are Drs. Andreas I. Moshovos, Scott E. Breach, Terani N. Vijaykumar, Gurindar S. Sohi. Plaintiff WARF is listed as the original assignee. WARF maintains that the named inventors conceived of the claimed invention no later than December 11, 1995.

---

[6] Apple maintains that there will be no error, at least technically, if the yet-to-execute STORE instruction would not change the value already written to the memory address, because the LOAD instruction still obtains the correct data. (Def.'s Resp. to Pl.'s PFOFs (dkt. #141) ¶ 31.)

[7] Apple contends that there may be times that the performance cost of mis-speculation does not outweigh the performance benefit of speculation. (Def.'s Resp. to Pl.'s PFOFs (dkt. #141) ¶ 34.)

During prosecution of the '752 patent, the named inventors provided no prior art to the Patent Office. On October 8, 1997, the patent examiner issued a Notice of References Cited, which listed four pieces of prior art. The patent examiner rejected pending claims 1-2, 6 and 8-11 as anticipated in light of U.S. Patent No. 5,555,432 ("Hinton"). On January 5, 1998, WARF filed a response cancelling pending claims 9 and 10, but arguing that claims 1, 2, 6, 8 and 11 were allowable over Hinton. On February 3, 1998, the examiner allowed those claims.

ii. Objectives and Specification

The '752 patent recognizes that "[t]he performance cost is a function of the frequency that speculation is required, the probability of mis-speculation and the time required to recover from a mis-speculation." ('752 patent at 3:14-18.) The '752 patent also observes that "most data dependent mis-speculations can be attributed to a few static STORE/LOAD instruction pairs," and that mis-speculations typically "exhibit 'temporal locality,'" such that "if one LOAD/STORE pair causes a data mis-speculation at a given point in time, it is highly likely that a later instance of the same pair will soon cause another mis-speculation." (*Id.* at 3:51-57.) The patent further observes that

> The present inventors believe that a relatively limited number of LOAD/STORE pairs will create mis-speculation and so the operation described above prevents the majority of the LOAD/STORE pairs from being slowed in execution. The list of critical LOAD/STORE pairs is prepared dynamically

> in a synchronization method for those LOAD/STORE pairs . . . .

(*Id.* at 14:15-22.) Based on these observations, the inventors concluded that load-based memory dependencies may be amenable to history-based prediction.[8] As such, the '752 patent associates predictions with particular LOAD instructions that have mis-speculated in the past.

The specification of the '752 patent describes a processor containing a "data speculation circuit" that detects dependence between LOAD and STORE instructions. The data speculation circuit also detects mis-speculations where a LOAD instruction that is dependent for its data on a STORE instruction appearing earlier in program order is in fact executed before the STORE instruction. According to the preferred embodiment, the data speculation circuit sends a "mis-speculation indication" to a "predictor circuit" if a mis-speculation is detected, which uses the indication to then produce a prediction. The greater the "prediction," the greater the likelihood that the speculative execution of its associated LOAD instruction will cause a mis-speculation; the lower a prediction at a given time, the lower the likelihood of mis-speculation. The processor uses each prediction to decide whether its associated LOAD instruction should be allowed to execute speculatively.

The patent discloses a "three-tiered approach" to dealing with ambiguous dependency. The first tier

---

[8] Apple disputes that the named inventors were the first to develop history-based techniques for load-based memory dependencies. (Def.'s Resp. to Pl.'s PFOFs (dkt. #141) ¶ 42.)

considers whether a LOAD instruction has a history of mis-speculation. "If there is no history of data mis-speculation, [the instruction] is executed without further inquiry." ('752 patent (dkt. #1-1) 3:64-66.) At this tier, the '752 patent describes a "prediction table," in which entries are created when the processor detects a mis-speculation by a LOAD instruction. "[I]f no entry is found in the prediction table," then "the reasonable assumption is that speculation can proceed." (*Id.* at 11:22-24.) The second tier becomes relevant with a LOAD instruction has mis-speculated in the past. In this tier, "a predictor based on the past history of mis-speculations for that LOAD instruction is employed to determine whether the instruction should be executed or delayed." (*Id.* at 4:1-4.) With respect to the second tier, the patent explains that "it is an object of the invention to provide a predictor circuit that may identify data dependencies on an on-going and dynamic basis." (*Id.* at 4:31-33.) Finally, in cases where the prediction indicates that the LOAD instruction should not be executed speculatively, the third tier may be employed to decide when the LOAD instructions should be allowed to execute. This part of the patent describes a "synchronization table," which "indicates whether there is in fact a pending LOAD instruction awaiting its dependent STORE instruction." (*Id.* at 11:45-47.)

iii. Claim Construction

a) "Prediction"

The heart of the parties' dispute turns on the meaning of the term "prediction" as used in claim 1 and all other independent claims. As context, WARF

contends that the term "prediction" should be construed to mean "a variable that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation," where "a 'prediction' must be capable of receiving ongoing updates." (Pl.'s PFOFs (dkt. #122) ¶ 67.) In contrast, Apple contends that "prediction" need not be capable of receiving updates, and it, therefore, proposes a construction of "a value that indicates that likelihood that the data speculative execution of a load instruction will result in a mis-speculation," but does not necessarily contemplate a revision to that value based on regular updates. (*Id.* at ¶ 68.)

b) Other Agreed-Upon Terms

The parties agree to the following constructions of claim terms:

- "data speculation circuit" (claims 1 and 9): "a circuit that detects data dependence between load and store instructions and that detects mis-speculation by load instructions"
- "mis-speculation" (claims 1, 6, and 9): "when a load instruction that is dependent for its data on a store instruction appearing earlier in the program order is in fact executed before the store instruction wrote its data to a memory address shared with the load instruction"
- "in fact executed" (claims 1 and 9): "when a load instruction has actually accessed a memory address that has not yet been updated by a store instruction appearing earlier in the program order"

- "predictor" (claim 1): "a circuit that receives a mis-speculation indication from the data speculation circuit to produce a prediction"

D. State of the Prior Art

i. Overview

By 1995, out-of-order processing was well-known in the field of computer architecture design. Also by 1995, techniques for detecting data dependence were well-known in the art. On this much, the parties are in agreement.

Apple further maintains that by 1995, data speculation was well-known in the art, as were techniques for detecting and recovering from mis-speculations. WARF disputes this, asserting that the prior art techniques do not resemble the solutions proposed by the '752 patent inventors. Apple also contends that by 1995, prediction techniques to improve the accuracy of speculation in an out-of-order processor were well-known in the art. WARF also disputes this, and in particular contends that the techniques disclosed in the prior art did not satisfy the "prediction" claimed in the '752 patent -- the heart of the parties' dispute addressed in the opinion below. Finally, Apple also maintains that by 1995, data speculation involving LOAD and STORE instructions was well-known in the art. WARF disputes that too, arguing the prior art techniques bore no resemblance to the solutions proposed in the '752 patent.[9]

---

[9] Apple proposes finding of facts about other prior art references, including: a technique developed by Digital Equipment Corporation ("DEC"); U.S. Patent No. 5,666,506 ("Hesson"); and

ii. The Steely Patent

The "Steely patent" is titled "Memory Reference Tagging" and names Simon C. Steely, Jr., David J. Sager and David B. Fite, Jr. as inventors. The application was filed on August 12, 1995, and claims priority to an earlier application filed on November 12, 1992. The Steely patent issued as U.S. Patent No. 5,619,662 on April 8, 1997, and was assigned to DEC. As such, it is prior art to the '752 patent. Apple contends that the Steely patent anticipates claims 1-3, 5, 6, and 9 of the '752 patent.

Pertinent to this anticipation defense, all claims of the '752 patent require a "prediction" associated with a LOAD instruction or with a LOAD/STORE pair. Apple maintains that the Steely patent describes a processor that executes instructions out of order and uses a prediction to determine whether to allow speculation for LOAD and STORE instructions. WARF asserts that Steely fails to disclose any "prediction" capable of receiving ongoing updates -- or even a "prediction" under Apple's proposed construction of that claim term.

E. Person of Ordinary Skill in the Art

The parties dispute what characteristics a person of ordinary skill in the art would possess, though this

---

a commercial processor known as the Alpha 21264 or "EV6." As best as the court can discern, however, these prior art references are only material to Apple's motion for summary judgment on the objective prong of WARF's willful infringement claim. As discussed below, the court reserves on that based on any arguments not developed fully at summary judgment, waiting instead to hear the evidence of infringement and invalidity to be introduced during the first phase of the trial.

dispute is not material to the parties' respective motions for summary judgment, or at least the reasons for this court's disposition of those motions. Apple maintains that for purposes of the '752 patent, a person of ordinary skill in the art would have a Ph.D. in electrical engineering, computer engineering or computer science with a focus on computer architecture or microprocessor design; or an M.S. or B.S. degree in electrical engineering, computer engineering or computer science with significant work experience relating to computer architecture or microprocessor design.

WARF maintains that a person of ordinary skill in the art would have at least a bachelor's degree in electrical engineering or computer science, and at least three to five years of experience in computer design and computer architecture. Alternately, WARF asserts a person of ordinary skill in the art would have a master's degree in electrical engineering or computer science, and at least two to three years of experience in computer design and computer architecture. The experience could be derived from either industry or academia.

F. IPR Decision

Finally, Apple filed a petition with the United States Patent and Trademark Office's Patent Trial and Appeal Board ("PTAB") seeking *inter partes* review ("IPR") of all claims of the '752 patent. In the petition, Apple argued that claims 1-9 are invalid as obvious in view of Hesson and Steely, relying on a declaration of its expert Dr. Colwell. On April 15, 2015, after briefing by Apple and WARF, the PTAB denied Apple's petition "as to all challenged claims,"

finding that Apple "has not shown . . . that there is a reasonable likelihood that it will prevail" on its obviousness theory for any claim of the '752 patent. (4/17/2015 Declaration of Christopher Abernathy ("4/17/15 Abernathy Decl."), Ex. A (dkt. #151-1) p.3.)

In particular, the PTAB construed "prediction" as "a variable that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation." (*Id.* at p.10.) In so finding, the PTAB reasoned:

> We agree that in the '752 patent, the mis-speculation prediction at any point in time is a function of the mis-speculation history of load-store instruction pairs. Thus, the prediction is a variable. The fact that the prediction has a particular value at each point in time is merely an indication of its functional relationship and does not change the nature of the prediction from a variable to a constant value.

(*Id.*)[10]

---

[10] In *Wisconsin Alumni Research Foundation v. Intel Corp.*, No. 08-cv-78-bbc (W.D. Wis. filed Feb. 5, 2008), Judge Crabb similarly construed "prediction" in the same patent to mean "a variable that indicates the likelihood that the data speculative execution of a load instruction will result in mis-speculation," and later clarified that "a 'prediction' must be capable of receiving ongoing updates." *Wis. Alumni Research Found. v. Intel Corp.*, No. 08-cv-78-bbc, 2008 WL 4279975, at *7 (W.D. Wis. Sept. 18, 2008); *Wis. Alumni Research Found. v. Intel Corp.*, 656 F. Supp. 2d 898, 922 (W.D. Wis. 2009). While the court agrees with Apple that this decision has no binding effect on this court (Def.'s Opening Br. (dkt. #118) 23), any more than the PTAB's decision does, the court will obviously consider Judge Crabb's and the PTAB's reasoning and analysis in the opinion below.

OPINION

I. Claim Construction

"It is a 'bedrock principle' of patent law that 'the claims of a patent define the invention to which the patentee is entitled the right to exclude.'" *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005) (en banc) (quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 1111, 1115 (Fed. Cir. 2004)). The court exclusively determines claim construction as a matter of law. *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 372 (1996). The words of the claims are always the "appropriate starting point" for proper construction, *Comark Commc'ns, Inc. v. Harris Corp.*, 156 F.3d 1182, 1186 (Fed. Cir. 1998), with the court asking "how a person of ordinary skill in the art understands a claim term" as an "objective baseline from which to begin claim interpretation," *Phillips*, 415 F.3d at 1313.

"Importantly, the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification." *Id.* In fact, "[t]he best source for understanding a technical term is the specification from which it arose, informed, as needed, by the prosecution history." *Multiform Desiccants, Inc. v. Medzam, Ltd.*, 133 F.3d 1473, 1478 (Fed. Cir. 1998). As the Federal Circuit has recognized, however, "there is sometimes a fine line between reading a claim in light of the specification, and reading a limitation into the claim from the specification." *Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 904 (Fed. Cir. 2004) (quoting *Comark*

Commc'ns, 156 F.3d at 1186-87). "[A]n inherent tension exists as to whether a statement is a clear lexicographic definition or a description of a preferred embodiment. The problem is to interpret claims 'in view of the specification' without unnecessarily importing limitations from the specification into the claims." *E-Pass Techs., Inc. v. 3Com Corp.*, 343 F.3d 1364, 1369 (Fed. Cir. 2003).

In addition to intrinsic evidence like the specification and prosecution history, the Federal Circuit has "authorized district courts to rely on extrinsic evidence, which 'consists of all evidence external to the patent and prosecution history, including expert and inventor testimony, dictionaries, and learned treatises.'" *Phillips*, 415 F.3d at 1317 (quoting *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 980 (Fed. Cir. 1995)). "However, while extrinsic evidence 'can shed useful light on the relevant art,' [the Federal Circuit has] explained that it is 'less significant than the intrinsic record in determining 'the legally operative meaning of claim language.'" *Id.* (quoting *C.R. Bard, Inc. v. U.S. Surgical Corp.*, 388 F.3d 858, 862 (Fed. Cir. 2004)). Accordingly, the court can consider extrinsic evidence in construing patent claims, but it must do so in the context of the intrinsic evidence and while keeping in mind the flaws inherent in each type of extrinsic evidence. *Id.* at 1318.

As previously mentioned, the parties dispute the proper construction of only one claim term, "prediction," which appears in claims 1, 2, 3, 5 and 9

of the '752 patent.[11] The parties propose the following constructions for that term:

"Prediction"	
Plaintiff WARF's Proposed Construction	**Defendant Apple's Proposed Construction**
"A variable that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation" A "prediction" must be capable of receiving ongoing updates.	"A value that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation" A "prediction" need not be capable of receiving ongoing updates.

The obvious and sole substantive difference between the parties' competing constructions is whether the prediction *must* be capable of change (while implicit in the "value"-"variable" dispute, the second sentence of each definition makes that disagreement explicit). According to WARF, a "prediction" must be able to receive updates -- in other words, it must be dynamic. Apple, on the other hand, argues that a "prediction" *may* be dynamic, but it may *also* be static -- that is, incapable of receiving ongoing updates and changing to reflect those updates.

---

[11] The parties also mention a possible dispute about the meaning of "table," but neither party sought construction at summary judgment.

The language of claim 1 reads as follows:

> In a processor capable of executing program instructions in an execution order differing from their program order, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where a data consuming instruction dependent for its data on a data producing instruction of earlier program order, is in fact executed before the data producing instruction, a data speculation decision circuit comprising:
>
> a) a predictor receiving a mis-speculation indication from the data speculation circuit to produce a *prediction* associated with the particular data consuming instruction and based on the mis-speculation indication; and
>
> b) a prediction threshold detector preventing data speculation for instructions having a *prediction* within a pre-determined range.

('752 patent, 14:36-52 (emphasis added).)

The court can dispense with one of Apple's arguments at the outset. Apple points out that none of the claims at issue contain an *express* limitation requiring the prediction to be updated on an ongoing basis, suggesting that this means predictions need not be capable of update. (Def.'s Br. Support Summ. J. (dkt. #118) 19-20.) As appealing as that simple construction might be, since it would alleviate the need for further analysis, the lack of an express

limitation actually requires further inquiry: if the claims contained an additional limitation requiring a *dynamic* prediction, construing the term "prediction" to be intrinsically dynamic would render that limitation superfluous. *See Digital-Vending Servs. Int'l, LLC v. Univ. of Phoenix, Inc.*, 672 F.3d 1270, 1275 (Fed. Cir. 2012) (discussing the "well-established rule that claims are interpreted with an eye toward giving effect to all terms in the claim") (internal quotation marks and citation omitted); *cf. LSI Indus., Inc. v. ImagePoint, Inc.*, 279 F. App'x 964, 972 (Fed. Cir. 2008) ("Some claims specifically recite 'an illuminated display device,' while others recite only 'a display device.' . . . Thus, the language of the claims counsels against imposing an illumination limitation on the display device term because it would make the limitation superfluous where it explicitly appears."); *Phillips*, 415 F.3d at 1314 ("To take a simple example, the claim in this case refers to 'steel baffles,' which strongly implies that the term 'baffles' does not inherently mean objects made of steel."). Thus, although the claims include no limitations explicitly requiring predictions to be dynamic, the word "prediction" itself still might (or might not) include that requirement depending on the claim language, specification, prosecution history and extrinsic evidence.

WARF relies heavily on the fact that the claimed "data speculation decision circuit" prevents data speculation for instructions having a prediction "within a pre-determined range." According to WARF, the claimed function of determining whether a prediction falls within a given range makes sense only if, "at any given time *after* the 'prediction' is produced,

it *might* be 'within a predetermined range' or it might *not* be." (Pl.'s Br. Opp'n Summ. J. (dkt. #148) 8 (emphasis in original).) Apple argues in response that this interpretation narrows the claims in a way not supported by the text or the state of the prior art. In particular, Apple points out that prior art in the field, including an article entitled "A Study of Branch Prediction Strategies" by James E. Smith (Decl. of Bryan S. Conley, Ex. 12 (dkt. #124-12) [hereinafter "Smith" or the "Smith article"]), used the word "prediction" in the context of speculation strategies tracking single past events, rather than a dynamic history of such events.[12] According to Apple, the claims certainly *permit* a dynamic prediction but are also broad enough to encompass a prediction incapable of receiving updates.

In *Wisconsin Alumni Research Foundation v. Intel Corp.*, 656 F. Supp. 2d 898 (W.D. Wis. 2009), this court relied in part on the same language WARF cites, finding that:

> Claim 1 describes "producing" a "prediction" from a "mis-speculation indication" generated in a data speculation circuit and determining whether that "prediction" is "within a predetermined range" to decide whether to prevent data speculation. Thus, the claim language itself establishes that a

---

[12] The Smith article deals with control speculation, rather than data speculation. Control speculation involves "branch prediction." In the words of the '752 patent, it "might involve executing an instruction that follows a branch instruction without knowing the outcome of the branch (and thus whether the following instruction should have been executed or was branched around)." ('752 patent, 2:32-36.)

> "prediction" is something other than a stored "indication" and is capable of having a "range" of values . . . On its face, this language suggests that a prediction must be capable of changing over time.

Id. at 922.

Revisiting this same claim language here, the court again finds the contemplated use of a predetermined range of values to assess whether instructions should be permitted to speculate favors WARF's narrower interpretation. By way of example, imagine a pair of instructions that mis-speculates for the first time. The parties agree that a single mis-speculation is enough to produce a prediction; thus, in this instance, the predictor of the invention would receive that mis-speculation indication from the data speculation circuit and use it to produce a prediction of "1," representing the single mis-speculation. Under Apple's construction, the development of the prediction can end here, because it need not be capable of further updates. Thus, the prediction would be set permanently at its initial value of "1."

Under this approach, the next time the instructions execute, there is no need for the prediction threshold detector to assess whether the prediction falls within a "pre-determined range." In a static situation, there are only two possibilities: either there *is* a prediction with a value of "1," because the instructions have mis-speculated a single time; or there is *no* prediction, because the instructions have not yet mis-speculated and, therefore, the prediction has not yet been created. Thus, under Apple's construction, the question for the prediction threshold

detector is a binary determination of whether a prediction *exists at all*, rather than whether a prediction "falls within a given range." Indeed, there would be no need for "a data speculation *decision* circuit" in claim 1 at all, feeding ongoing mis-speculation outcomes, since the "data speculation circuit" itself would provide the single piece of information required for a static prediction.

Effectively, the prediction threshold detector would prevent data speculation.[13] Said another way, Apple's construction would read out the words "within a pre-determined range" from subsection (b) of claim 1, or at least render them superfluous in the context of "predictions" incapable of receiving updates; in those cases, the prediction threshold detector would prevent data speculation "for instructions having a prediction," full stop.

Of course, as Apple argues, a "range" can consist of a single value, which could technically allow for a "predetermined range" including *only* the value 1. Superficially, this provides some support for Apple's construction, but it still does not explain *why* it would ever be necessary to compare an existing prediction to a range of 1 for so-called static "predictions." The choice is still binary -- either there is no prediction or the prediction is set to 1 -- and so the notion of

---

[13] Theoretically, it is possible that the prediction of "1" would *not* fall within the predetermined range and the instruction would be allowed to execute regardless of the previous mis-speculation. But if that were so, the invention would appear to serve no purpose, since the prediction would not prevent speculation and could not change, much less improve, the processor's performance over time. Likely for this reason, Apple does not advance this argument, so the court does not consider it further.

"comparison" remains a poor fit for the kind of theoretical static "predictions" Apple posits, regardless of whether the predetermined range is set to encompass multiple values or a single value.

The remainder of the '752 patent further supports WARF's construction. The brief summary of the invention describes a three-tiered approach for determining when an instruction should execute. The first tier encompasses instructions with no history of mis-speculation; they may execute "without further inquiry." ('752 patent, 3:66.) The second tier implicates instructions that have previously mis-speculated. At that point, according to the description, the invention employs a predictor "to determine whether the instruction should be executed or delayed." (*Id.* at 4:1-4.) If the prediction were static, however, the mere fact of its existence would be enough to prevent execution. In contrast, the predictor as described in the '752 patent instead uses "the past history of mis-speculations" to determine whether the instruction may execute, allowing those that are "*typically* not dependent" to execute immediately. (*Id.* at 4:1-5 (emphasis added).) This language, too, suggests a prediction capable of update; it makes little sense to speak of instructions that are "typically not dependent" when a single instance of mis-speculation could, under Apple's construction, foreclose future speculative execution without the possibility of updates to reflect what *typically* occurs. If the predictor ultimately delays the instruction, the third tier then employs a synchronization table to determine when the instruction should execute, delaying it "until after the

execution of the particular data producing instruction" on which it depends. (*Id.* at 4:5-7, 27-28.)

Furthermore, this three-tiered approach appears in the brief summary of the invention, rather than as a description of a single embodiment, making it more persuasive as a source of support for WARF's narrower construction. *See C.R. Bard, Inc. v. U.S. Surgical Corp.*, 388 F.3d 858, 864 (Fed. Cir. 2004) ("Statements that describe the invention as a whole, rather than statements that describe only preferred embodiments, are more likely to support a limiting definition of a claim term. . . . Statements that describe the invention as a whole are more likely to be found in certain sections of the specification, such as the Summary of the Invention.").

While less persuasive given its location in the patent, the detailed description of the invention provides further context suggesting that a "prediction" must be dynamic. As this court recognized in describing the preferred embodiment of the invention in *Intel*, "the specification explains in unequivocal terms that '[t]he prediction provided by the predictor circuit 33 . . . is updated based on historical mis-speculations detected by the data speculation circuit 30. For this reason, the data speculation circuit 30 must communicate with the predictor circuit 33 on an *ongoing* basis.'" *Intel*, 656 F. Supp. 2d at 922 (quoting '752 patent, 8:7-11) (emphasis added). WARF also points to other examples supporting its position in the description of the preferred embodiment, including the description of the way the prediction normally "is *incremented and decremented*" such that "the higher the prediction **109**, the more likelihood of mis-speculation[.]" ('752

patent, 11:29-35 (emphasis added).) Indeed, throughout the description of the preferred embodiment, the specification consistently refers to the prediction as dynamic. (*See, e.g.*, *id.* at 12:14-17 ("[T]he prediction that there was a need to synchronize was wrong and so at process block **120** the prediction **109** is *decremented* toward the do not synchronize state.") (emphasis added); 12:52-54 ("[T]he prediction **109** is *updated* toward the synchronize condition indicating that the prediction that there was a need to synchronize was correct[.]") (emphasis added); 12:67-13:3 ("If [a mis-speculation occurs and the pair is already in the prediction table] then at process block **302**, the prediction **109** is *updated* toward synchronize so that this mis-speculation may be avoided in the future.") (emphasis added).)

Acknowledging, as it must, that the preferred embodiment describes a dynamic prediction that receives ongoing updates (Def.'s Br. Support Summ. J. (dkt. #118) 21; Def.'s Br. Opp'n Summ. J. (dkt. #140) 16), Apple relies on the general principle that "it is improper to read limitations from a preferred embodiment described in the specification -- even if it is the only embodiment -- into the claims absent a clear indication that the patentee intended the claims to be so limited." *Liebel-Flarsheim Co.*, 358 F.3d at 913. But this is not a case in which the "claim language is sufficiently broad that it can be read to encompass features not described in the written description, either by general characterization or by example in any of the illustrative embodiments." *Id.* at 905. Rather, as described above, the claims *themselves* suggest that the contemplated "prediction"

must be capable of change; the preferred embodiment merely provides further support for that conclusion.

Use of the preferred embodiment as context, rather than as a source of limitations that do not otherwise appear in the claims, is permissible. *Compare, e.g.*, *Teleflex Inc. v. Ficosa N. Am. Corp.*, 299 F.3d 1313, 1327-28 (Fed. Cir. 2002) (district court erred in holding that "clip" was limited to a "single pair of legs," even where that was the only embodiment described, where claim language did not support that limitation, specification and prosecution history included no statements of restriction and the ordinary meaning of "clip" was not so restricted), *with Toro Co. v. White Consolidated Indus., Inc.*, 199 F.3d 1295, 1301-02 (Fed. Cir. 1999) (holding that construction of "including" required attachment between structures where that was the only embodiment disclosed *and* where nothing in the remainder of the specification supported an unattached embodiment; "[T]he specification describes the advantages of the unitary structure as important to the invention. . . . No other, broader concept was described as embodying the applicant's invention, or shown in any of the drawings, or presented for examination.").

Still, Apple argues that the '752 patent *does* expressly contemplate alternative embodiments of the invention, pointing out that the detailed description of the invention states:

> It will be understood that the prediction **109** may be obtained by methods other than simply incrementing it in value for each speculation as described herein. For example,

> various weighting schemes can be provided to cause the predictor circuit **33**, for example, to be less sensitive to the earliest mis-speculations. More complex pattern matching techniques may also be used, for example, to catch situations where mis-speculations occur in groups or regular patterns.

('752 patent, 14:6-14.) Apple contends that a person of ordinary skill in the art would read this discussion to allow for alternative embodiments in which a prediction is *not* updated on an ongoing basis and urges the court not to "improperly exclude a disclosed embodiment" by adopting WARF's construction. *See Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1333 (Fed. Cir. 2013). The problem with this argument is that the patent simply does *not* disclose the embodiment Apple advocates.

As a beginning point, *neither* of the alternative embodiments disclosed in the specification contemplate a static "prediction." To the contrary, both proposed alternatives implicitly contemplate arrangements involving dynamic predictions. Both schemes that assign a different weight to later mis-speculations and techniques that identify mis-speculations occurring in groups or regular patterns assume a *developing history* of mis-speculations that the predictor circuit can use to obtain its prediction. There is no need to weight different instances of mis-speculation if the prediction is static and will never be updated to reflect those weights.[14] There is likewise

---

[14] Apple uses the mention of "weighting schemes" to propose its own take on an alternative embodiment as well -- a "weighting

no need to develop complex matching techniques to identify patterns in mis-speculation if the prediction can never take that information into account in determining how likely an instruction is to mis-speculate. Adopting WARF's construction, therefore, does not exclude a "disclosed embodiment" from the scope of the claims.

Nor is there support for Apple's proposed construction in the intrinsic evidence, Apple's arguments to the contrary notwithstanding. For instance, Apple argues that the specification makes clear that a single mis-speculation is enough to produce a prediction. (Def.'s Br. Support Summ. J. (dkt. #118) 18.) This is true enough, but it has no bearing on whether a prediction must be capable of update *after* its creation -- nor does the fact that the claim language "does not specify any minimum number of times that the instruction must mis-speculate before the 'prediction' is above the threshold required to prevent speculation." (*Id.*) Apple also argues that a "prediction" must be construed as a "value" (i.e., static) rather than a "variable" (i.e., dynamic) because the specification "explicitly describes the prediction as a 'value.'" (*Id.* at 15-16.) But the examples it cites speak of the prediction being

---

scheme that always prevents speculation by a load instruction for which mis-speculation recovery would be especially costly." (Def.'s Br. Support Summ. J. (dkt. #118) 23.) Such a weighting scheme would *not*, however, be a means of producing a "prediction." Both parties agree that a prediction indicates the *likelihood* that a pair of instructions will mis-speculate; Apple's embodiment has nothing to do with the likelihood of mis-speculation, but rather assesses whether the costs associated with a single mis-speculation are prohibitive, regardless of how likely or unlikely that mis-speculation is.

set to a "default value" or being "incremented in value." All this confirms is that the prediction is some number that *has* a value; it does not suggest the value of that prediction cannot change. To the contrary, the portions of the specification Apple cites refer to "incrementing" the value of the prediction, suggesting that it can and does change. Thus, the court again adopts Judge Crabb's conclusion in *Intel* that "[n]either the claim language nor the specification supports defendant's proposed construction that a 'prediction' may include values that are fixed once to indicate a single incident of mis-speculation." 656 F. Supp. 2d at 922.

Finally, Apple contends that its construction finds support in *extrinsic* evidence, citing to the Smith article discussed above, as well as the reports of its two experts, Dr. August and Dr. Colwell. Both reports, however, primarily rehash Apple's legal arguments by purporting to analyze the language in the specification and claims. (*See* August Report (dkt. #103) ¶¶ 136-47; Colwell Report (dkt. #104) ¶¶ 131-41.) The court rejects Apple's positions on those issues, and so, too, expert reports that echo those same arguments.

Apple is, therefore, left with the Smith article and another paper, "Memory Dependence Prediction [U]sing Store Sets," by George Z. Chrysos and Joel S. Emer (the "Chrysos paper") (Conley Decl., Ex. 1 (dkt. #143-1)), both of which Apple contends describe techniques that produce static predictions. Even if Apple's characterization were accurate, these two extrinsic references are wholly underwhelming compared to the language of the patent itself and contrary intrinsic evidence. Moreover, extrinsic

evidence "can be used only to help the court come to the proper understanding of the claims"; it cannot be used to vary or contradict the claim language or specification. *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1584 (Fed. Cir. 1996).

For all these reasons, the court finds WARF's proposed construction of the term "prediction" compelling and will construe that term as requiring a prediction that is capable of receiving updates.

II. Invalidity

A. Anticipation by Steely

On summary judgment, both parties devote most of their invalidity briefing to the question of whether the '752 patent is invalid as anticipated by the Steely patent, U.S. Patent No. 5,619,662. Evaluating a claim of anticipation involves a two-step inquiry. The first step requires proper construction of the meaning and scope of the claims. *Power Mosfet Techs., L.L.C. v. Siemens AG*, 378 F.3d 1396, 1406 (Fed. Cir. 2004). "The second step in the analysis requires a comparison of the properly construed term to the prior art[.]" *Id.* To demonstrate anticipation, "the proponent must show 'that the four corners of a single, prior art document describe every element of the claimed invention.'" *Net MoneyIN, Inc. v. VeriSign, Inc.*, 545 F.3d 1359, 1369 (Fed. Cir. 2008) (quoting *Xerox Corp. v. 3Com Corp.*, 458 F.3d 1310, 1322 (Fed. Cir. 2006)). Although anticipation is ultimately a question of fact, "it may be decided on summary judgment if the record reveals no genuine dispute of material fact." *Leggett & Platt, Inc. v. VUTEk, Inc.*, 537 F.3d 1349, 1352 (Fed. Cir. 2008)

(quoting *Golden Bridge Tech., Inc. v. Nokia, Inc.*, 527 F.3d 1318, 1321 (Fed. Cir. 2008)).

Both parties have moved for summary judgment of anticipation in light of Steely, and each relies on its own, preferred construction of the disputed term "prediction." (Def.'s Br. Support Summ. J. (dkt. #118) 25-37; Pl.'s Br. Support Summ. J. (dkt. #120) 38-64.) Having just rejected Apple's construction of that term, Apple's motion for summary judgment will be denied. Even if the court adopts WARF's construction, however, Apple maintains numerous disputed issues of fact preclude entry of summary judgment against it on grounds of anticipation. (Def.'s Br. Opp'n Summ. J. (dkt. #140) 27.) It is to that question the court now turns.

i. Background of the Steely Patent

The Steely patent is entitled "Memory Reference Tagging" and describes a processor that "includes a memory reference tagging store associated with the instruction scheduler so that the scheduler can reorder memory reference instructions without knowing the actual memory location addressed by the memory reference instruction." (U.S. Patent No. 5,619,662 (dkt. #131-4) Abstract.) Most relevant to the issue of anticipation, Steely discloses four different techniques in which a "write buffer" assigns "memory reference tags" involving a mis-speculation to load and store instructions. Each of those techniques appears in the section of the patent entitled "Memory Reference Tagging." (*See id.* at 47:35-49:8.)

In the first technique, a mis-speculation generates a memory reference tag from a portion of

the address in memory that resulted in the LOAD-STORE collision. (*Id.* at 48:2-4.) Once that portion is placed in the memory tag store, every time an instruction is retrieved from memory to be executed, the memory reference tag circuit "will provide the tag bits to be used by the instruction scheduler." (*Id.* at 48:30-33.) If the instructions appear with identical tag bits (indicating a previous mis-speculation), the instruction scheduler will not reorder them. (*Id.* at 48:33-36.)

In the second technique, the pair of instructions after a mis-speculation is tagged not with a portion of the memory address, but instead with "a problem number which could be a number provided from a counter." (*Id.* at 48:55-57.) As a result, "[t]wo memory reference instructions with the same address and number will not reorder." (*Id.* at 48:57-59.) "However, if the two memory reference instructions have a different number, the instructions will reorder." (*Id.* at 48:59-61.) The counter does not appear to increment with respect to that for the same pair of instructions once it has assigned the "problem number"; rather, it increments when a mis-speculation occurs with respect to a *different* pair of instructions. For instance, if a pair of instructions mis-speculates and is assigned the problem number 0, the next pair to mis-speculate might be assigned the problem number 1.

The third technique is to assign an instruction a "bit to indicate that an instruction should not be reordered." (*Id.* at 48:62-63.) Thus, using this technique, "for a store that previously caused a problem in the write buffer, the instruction is tagged with a bit indicating that the ISCHED **38** [instruction

scheduler] cannot reorder memory reference instructions around the Instruction tagged with the bit." (*Id.* at 48:63-67.)

The final technique is to "turn off reordering" entirely in response to a mis-application under certain circumstances. (*Id.* at 49:4-5.) For example, the patent suggests turning off reordering when entering a subroutine, based on the general observation that "during a subroutine call, there are some initial stores and some exiting loads" and "[i]t would not be desirable to reorder the exiting loads before the initial stores." (*Id.* at 49:4-9.)

ii. Analysis

Construing "prediction" as a dynamic (updating) "variable that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation," the remaining question for deciding the anticipation issue before the court is whether Steely discloses a prediction that can change over time. In *Intel*, this court found that it did not, holding that the requirement of a dynamic prediction was "fatal to defendant's contention that the four techniques disclosed in the '662 patent anticipate the '752 patent." 656 F. Supp. 2d at 922. WARF urges the court to adopt the same result here, arguing that Steely's four techniques do not disclose predictions that update on an ongoing basis. According to WARF, those techniques simply involve tagging instructions to reflect a single mis-speculation event, without providing a mechanism to update those tags. *See also Intel*, 656 F. Supp. 2d at 922 ("For each [technique], the tag is designed to indicate only that a mis-

speculation has occurred, not keep track of mis-speculations on an ongoing basis.").

Unsurprisingly, Apple objects to this characterization. Apple instead contends that the outcome of Steely's tag comparison "*can* change over time for the same pair of load and store instructions." (Def.'s Br. Opp'n Summ. J. (dkt. #140) 29.) According to Apple's expert, Dr. Colwell, that can occur if, after the write buffer assigns the tags, additional mis-speculations involving one of the pair of tagged instructions occur. As an example, Dr. Colwell presumes a situation in which a load instruction, "Inst 1011," and store instruction, "Inst 1007," have been tagged with the same memory address of "10010," such that Steely would prevent speculation. (Colwell Report (dkt. #104) ¶ 303.) Colwell then posits another situation in which a *different* load instruction, Inst 1012, *also* mis-speculates with store instruction Inst 1007:

> Another load instruction, for instance Inst 1012, may later be reordered ahead of store instruction Inst 1007, both Inst 1012 and Inst 1007 accessing the same memory address, this address ending in a different set of 5 bits, for instance "00110." Inst 1007 would then be associated with the tag "00110," which would no longer be identical to the tag "10010" associated with the load instruction Inst 1011. Because the tags for Inst 1007 and Inst 1011 are no longer identical, Steely predicts they are not dependent and may reorder them. Thus, the "prediction" disclosed by Steely is a "variable" that is "capable of receiving ongoing updates," as required by

> WARF's proposed construction of the term "prediction."

(*Id.*)

Whether Steely actually discloses this means of "updating" its tags within the four corners of the patent is certainly open to debate. Apple asks the court to infer as much, based on the fact that: (1) the memory reference tag store is large enough to store just one tag per instruction; and (2) the patent describes how tags for each mis-speculation "will be stored" regardless of other tags that may already exist for those instructions. According to Apple, these two facts demonstrate that Steely overwrites previously stored tags, thereby "updating" the result of any comparison that Steely performs between the two.

The flaws in this argument are multiple. Essentially, Apple and its experts assume what amounts to a defect in Steely, which prevents it from assigning more than one tag preventing mis-speculation to a single instruction (in Dr. Colwell's example, Inst. 1007), even though the example posits tags with different memory addresses (Inst. 10010 and Inst. 00110), depending on the store instruction with which the load instruction 1007 is paired (here, Inst. 1011 and Inst. 1012). Not only is this assumption contradicted by the language in Steely, *see* U.S. Patent no. 5,619,662, at 47:37-43 ("The memory reference tag store . . . provides *at least* one bit associated with said instruction . . ."') (emphasis added), but it would undermine the whole purpose of Steely, which is to prevent future mis-speculations, since it would result in a never-ending loop for load instructions causing multiple mis-speculations each

time the 10010 and 00110 tags overwrite one another.[15]

More importantly for summary judgment purposes, even assuming one might infer that such overwriting occurs and constitutes "updating," the above-described "tag replacement system" would hardly constitute a "prediction" as this court has construed the term. Properly construed, a "prediction" must communicate the likelihood of mis-speculation *and* must be capable of update. Using the example offered by Dr. Colwell for the sake of simplicity, Inst 1011 and Inst 1007 proved to be dependent and were, accordingly, tagged with the same memory address. Thereafter, *another* load instruction, Inst 1012, also proves to be dependent on store instruction Inst 1007. Accordingly, Steely overwrites the first tag on Inst 1007, tagging it to match Inst 1012, but that is not so much an "update" of the comparison between Inst 1007 and Inst 1011 as it is the wholesale *elimination* of that comparison. By Apple's and Dr. Colwell's own description, *no* record of the previous mis-speculation remains; the next time the tags are compared under

---

[15] The inventor of Steely does appear to have admitted in his deposition that in his view, this is how his invention would function, though he was asked the question out of context and without being asked about the obvious defect this would appear to create in his patent. (*See* Def.'s Br. Opp'n Summ. J. (dkt. #140) 31.) As WARF points out, this is why after-the-fact testimony of the inventor is of limited relevance when unsupported by the patent itself. *See Howmedica Osteonics Corp. v. Wright Med. Tech., Inc.*, 540 F.3d 1337, 1346 (Fed. Cir. 2008) ("The testimony of an inventor 'cannot be relied on to change the meaning of the claims.'") (quoting *Markman*, 52 F.3d at 983). What matters for purposes of anticipation is what the patent *actually* discloses, not what the inventor says it would do in a situation the patent does not clearly address.

Steely, they fail to reflect that any mis-speculation has occurred in the past and, therefore, fail to communicate the likelihood that the data speculative execution of the load instruction 1011 and store instruction 1007 will result in a mis-speculation. In contrast, the invention of the '752 patent incrementally increases the prediction for each mis-speculation associated with an instruction pair, while it decrements the prediction associated with a pair of instructions when they *do not mis-speculate*, thereby updating its assessment of the likelihood of mis-speculations in the future.

Steely's tag replacement system, even as explained by Dr. Colwell, *discards* the prediction associated with a pair of instructions when a different pair of instructions mis-speculates. While this data elimination admittedly yields a *change* in the result of the tag comparison, that change has nothing to do with updating the likelihood that the first pair of instructions will mis-speculate again in the future. Accordingly, no reasonable jury could find that the Steely patent discloses each and every limitation of the '752 patent as properly construed. The court will, therefore, grant summary judgment to WARF on defendant's Steely anticipation defense and counterclaim.

B. Indefiniteness

Finally, Apple contends that claims 5 and 6 of the '752 patent are invalid as indefinite. "[A] patent is invalid for indefiniteness if its claims, read in light of the specification delineating the patent, and the prosecution history, fail to inform, with reasonable certainty, those skilled in the art about the scope of

the invention." *Nautilus, Inc. v. Biosig Instruments, Inc.*, 134 S. Ct. 2120, 2124 (2014). A party raising an indefiniteness challenge, like other invalidity challenges, bears the burden of proving that invalidity by clear and convincing evidence. *Microsoft Corp. v. i4i Ltd. P'ship*, 131 S. Ct. 2238, 2242 (2011); *see also* 35 U.S.C. § 282.

Here, Apple contends that claims 5 and 6 of the '752 patent should be held indefinite under *Nautilus* solely because certain terms in those claims lack an antecedent basis. Claim 5 is a dependent claim and reads:

> The data speculation decision circuit of claim **2** wherein the instruction synchronization circuit includes a synchronization table associating *the certain data consuming instructions* and *the certain data producing instructions* each with a flag value indicating whether the respective certain data producing instruction has been executed and wherein the instruction synchronization circuit delays the particular data consuming instruction only:
>
> i) when the prediction associated with the data consuming instruction is within a predetermined range; and
>
> ii) when the particular data consuming instruction is in *the prediction table*; and
>
> iii) when the flag indicates the particular data producing instruction has not been executed.

('752 patent, 15:7-20 (emphasis added).) Similarly, Claim 6 likewise depends from claim 2 and reads:

> The data speculation decision circuit of claim **2** wherein the instruction synchronization circuit creates an entry in *the synchronization table* including the particular data consuming instructions and data producing instructions and *the flag value* only after a mis-speculation indicating is received for the particular data consuming instruction and the particular data producing instruction.

('752 patent, 15:21-27 (emphasis added).) Apple focuses on the italicized portions of each of the above claims in making its § 112 argument.

According to Apple, the use of the definite article "the" in each of the above italicized instances suggests that the terms that article introduces must refer to specific claim elements already previously discussed. (Def.'s Br. Support Summ. J. (dkt. #118) 39.) *See also, e.g.*, *Warner-Lambert Co. v. Apotex Corp.*, 316 F.3d 1348, 1356 (Fed. Cir. 2003) ("[I]t is a rule of law well established that the definite article 'the' particularizes the subject which it precedes. It is a word of limitation as opposed to the indefinite or generalizing force of 'a' or 'an.'") (quoting *Am. Bus Ass'n v. Slater*, 231 F.3d 1, 4-5 (D.C. Cir. 2000)). As Apple points out, the italicized terms above do *not* appear elsewhere in claims 5 and 6 themselves, or in claims 1 and 2, on which both claims 5 and 6 ultimately depend. In Apple's view, this makes it impossible for a person of skill in the art to determine the scope of claims 5 and 6, rendering them indefinite.

In *Halliburton Energy Services, Inc. v. M-I LLC*, 514 F.3d 1244 (Fed. Cir. 2008), the Federal Circuit held that "a claim could be indefinite if a term does not have proper antecedent basis where such basis is not otherwise present by implication or the meaning is not reasonably ascertainable." *Id.* at 1249; *see also Energizer Holdings, Inc. v. Int'l Trade Comm'n*, 435 F.3d 1366, 1370 (Fed. Cir. 2006) (citing *Slimfold Mfg. Co. v. Kinkead Indus., Inc.*, 810 F.2d 1113, 1116 (Fed. Cir. 1987)). The specification can, however, provide sufficient context for a person skilled in the field of the art to understand the claim to render it definite. *See, e.g.*, *In re Skvorecz*, 580 F.3d 1262, 1268 (Fed. Cir. 2009) ("We agree with Mr. Skvorecz that the clause 'welded to said wire legs at the separation' does not require further antecedent basis in claim 1, for a person skilled in the field of the invention would understand the claim when viewed in the context of the specification."). Here, the terms in question are "reasonably ascertainable" in light of the patent's specification.

Taking first the terms "the certain data consuming instructions" and "the certain data producing instructions" in claim 5, the patent's specification summarizes the invention and notes that the invention's instruction synchronization circuit:

> may also include a synchronization table associating *certain data consuming instructions and certain data producing instructions*, each with a flag indicating whether the respective data producing instruction has been executed. The instruction synchronization circuit delays the

> subsequent instances of the certain data consuming instruction only when the prediction associated with the data consuming instruction is within a predetermined range and when the particular data consuming instruction is in the prediction table and when the flag indicates that particular data producing instruction has not been executed.

('752 patent, 4:54-65 (emphasis added).) As WARF points out, this portion of the specification tracks the language of claim 5 almost exactly. There is no reason why a person of ordinary skill in the art would not read "the certain data consuming instructions" and "the certain data producing instructions" to be those included in the synchronization table in light of the specification. At the very least, the brief summary of the invention allows one skilled in the art to proceed with "reasonable certainty," as *Nautilus* requires.

The term "the prediction table" in subsection (ii) of claim 5 would likewise inform a person of ordinary skill in the art that the "prediction table" is contained in the instruction synchronization circuit. As the brief summary of the invention states, "[t]he instruction synchronization circuit may include a *prediction table* listing certain data consuming instructions and certain data producing instructions each associated with a prediction." ('752 patent, 4:39-42 (emphasis added).) The instruction synchronization circuit then employs the entries in that prediction table in determining whether to delay subsequent instances of the data consuming instruction -- the instruction must be in the prediction table for delay to take place. (*Id.* at 4:48-53.)

As for claim 6, the "synchronization table" is the one that "may" be included in the instruction synchronization circuit (which is explicitly claimed in independent claim 2) and "associate[s] certain data consuming instructions and certain data producing instructions, each with a flag indicating whether the respective data producing instruction has been delayed." ('752 patent, 4:54-58.) The "flag value" likewise takes its meaning from this portion of the specification, which indicates that each pair of instructions in the synchronization table has "a flag indicating whether the respective data producing instruction has been executed." (*See id.*) The invention then uses the flag to determine when to delay execution of subsequent instances of the data consuming instruction. (*Id.* at 4:58-65.) A person of ordinary skill in the art would understand the scope of "the flag value" in claim 6 in light of this relatively clear context. (*Id.*)

Importantly, because Apple does not dispute that the specification offers context for the claim terms it identifies, that argument is waived. *See Fresenius USA, Inc. v. Baxter Int'l, Inc.*, 582 F.3d 1288, 1296 (Fed. Cir. 2009) ("If a party fails to raise an argument before the trial court, or presents only a skeletal or undeveloped argument to the trial court, we may deem that argument waived on appeal."); *Jordan v. Binns*, 712 F.3d 1123, 1134 (7th Cir. 2013) (undeveloped arguments considered waived); *Ultratec, Inc. v. Sorenson Commc'ns, Inc.*, No. 13-cv-346-bbc, 2014 WL 3565409, at *1 (W.D. Wis. July 17, 2014). Regardless, Apple takes an entirely different tack, one which requires a bit of explanation. According to Apple, in light of the antecedent basis

problems in claims 5 and 6, a person of ordinary skill in the art *might* simply look to the specification to understand the scope of the invention. However, Apple argues, she might also assume that claims 5 and 6 do not, in fact, depend from claim 2 but instead were intended to depend from claims 3 and 5, respectively, which would provide the requisite antecedent basis for the identified terms, but would also include additional limitations by virtue of depending from different claims. (Def.'s Br. Opp'n Summ. J. (dkt. #140) 41.)

The court does not find Apple's argument persuasive. Apple cites no cases in which courts found indefiniteness due solely to a lack of antecedent basis, at least where the specification so clearly delineates the structure of what the patent intended to claim. Instead, Apple cites *Novo Industries, L.P. v. Micro Molds Corp.*, 350 F.3d 1348 (Fed. Cir. 2003), for the proposition that claims are indefinite where "in light of the mistakes in the claims there is no clear choice as to how to interpret their scope." (Def.'s Br. Opp'n Summ. J. (dkt. #140) 46.) But *Novo* involved an obvious typographical error amenable to no fewer than four possible interpretations (at least one of which would have significant substantive implications for the scope of the claims).[16] *Novo* does

---

[16] In *Novo,* the claim included a "stop means formed on a rotatable with said support finger." 350 F.3d at 1352 (emphasis removed). Novo suggested correcting the claim either by deleting the words "a rotatable with" or by deleting the words "with said." *Id.* at 1357. The district court raised another possibility by changing the word "a" to "and." *Id.* And Micro Molds proposed as a fourth possibility that a word, such as "skirt" or "disk," might have been erroneously omitted, which would add an additional

not support this court *reading in* a typographical error to *create* ambiguity where the specification otherwise indisputably provides context to delineate the scope of the invention "with reasonable certainty." *Nautilus*, 134 S. Ct. at 2124.

The other case Apple cites, *Automed Technologies, Inc. v. Microfil, LLC*, 244 F. App'x 354 (Fed. Cir. 2007), is similarly unhelpful to its indefiniteness argument. In *Automed*, the Federal Circuit vacated and remanded a grant of summary judgment of non-infringement because the district court had based its ruling on a finding that the accused systems lacked a "controller" -- a limitation that was actually *absent* from the asserted claims. *Id.* at 359. In the midst of that discussion, the Federal Circuit observed:

> We also note that claim 27 of the '671 patent, which recites "the controller," appears to be mistakenly dependent on claim 20, in which this term finds no antecedent basis. . . . Because claim 21 - and not claim 20 - recites a "controller" limitation, perhaps claim 27 was intended to depend from claim 21.

Id. Even so, the Federal Circuit said nothing about that potential error rendering claim 27 indefinite. Rather, it "[left] to AutoMed any corrective action it deem[ed] necessary." *Id.* The Federal Circuit's observation that claim 27 might have been intended to depend from claim 21, not claim 20, certainly does

---

substantive limitation to the claims. *Id.* Because the Federal Circuit "[could not] know what correction [was] necessarily appropriate or how the claim should be interpreted," it concluded that the claim was necessarily indefinite "in its present form." *Id.* No comparable indefiniteness is even arguable in this case.

not compel, or even do much to support, a finding of indefiniteness in this case.

Accordingly, the court finds that the specification provides ample guidance as to what elements the claims are referencing when they refer to "the certain data consuming instructions," "the certain data producing instructions" and "the prediction table" (claim 5), as well as "the synchronization table" and "the flag value" (claim 6). Even the authority upon which defendants rely indicates that a lack of antecedent basis renders a claim indefinite only if "it would be unclear as to what element the limitation was making reference." Manual of Patent Examining Procedure § 2173.05(e) (9th ed. 2014); *see also Halliburton*, 514 F.3d at 1249. That is simply not the case here.

III. Willful Infringement

WARF has alleged a claim that Apple's infringement was willful, thereby permitting (but not requiring) the court to award enhanced damages. 35 U.S.C. § 284 ("[T]he court may increase the damages up to three times the amount found or assessed."); *Beatrice Foods Co. v. New Eng. Printing & Lithographing Co.*, 923 F.2d 1576, 1578 (Fed. Cir. 1991) ("It is well-settled that enhancement of damages must be premised on willful infringement or bad faith.") (citations omitted). Apple seeks summary judgment on this claim on the basis that WARF cannot as a matter of law meet the threshold for proving willfulness on an objective basis.

To establish willful infringement, WARF "must show by clear and convincing evidence" (1) that "the infringer acted despite an objectively high likelihood

that its actions constituted infringement of a valid patent," and (2) that "this objectively-defined risk . . . was either known or so obvious that it should have been known to the accused infringer." *In re Seagate Tech.*, 497 F.3d at 1371. The former "objective determination of recklessness" is a question for the court, not the jury. *Bard Peripheral Vascular, Inc. v. W.L. Gore & Assocs., Inc.*, 682 F.3d 1003, 1006-07 (Fed. Cir. 2012).

"[T]he 'objective' prong of *Seagate* tends not to be met where an accused infringer relies on a reasonable defense to a charge of infringement." *Id.* at 1005-06 (internal citation and quotation marks omitted); *see also Spine Solutions, Inc. v. Medtronic Sofamor Danek USA, Inc.*, 620 F.3d 1305, 1319 (Fed. Cir. 2010) (overturning jury's finding of willful infringement, finding that defendant raised a "substantial question as to the obviousness" of the patent in suit); *Douglas Dynamics,* 747 F. Supp. 2d at 1112 (granting summary judgment on willful infringement claim where there was "reasonable difference of opinion" and a "close question").

In cursory fashion, Apple's opening brief advances a wide range of arguments for seeking summary judgment on this objective prong. Some of the bases were fully briefed for review on the merits – namely, Apple's claim construction of "prediction," its related argument on anticipation by Steely and its indefiniteness defense and counterclaim as to claims 5 and 6. The court will take up Apple's motion on these bases in the discussion below.

Other bases, including ones on which Apple bears the burden of proof like obviousness, were not,

however, the subject of the parties' motions for summary judgment. While the court appreciates that it is WARF's burden to demonstrate that Apple's defenses to infringement or claims of invalidity are not objectively reasonable, Apple's scattershot approach in its motion renders the task near impossible to resolve on summary judgment. Perhaps if Apple had identified two or three of its strongest arguments, this may have been a manageable task. Instead, Apple's treatment of each basis is limited to a paragraph or two in its opening brief and reflects ships passing in the night in reply to WARF's responses.[17] In any event, WARF did come forward with evidence and law that, despite Apple's attempt to refute it in reply, could lead to a finding that Apple's belief that it either did not infringe the '752 patent or that the patent was invalid was not objectively reasonable.[18] As such, the court will await a more robust demonstration of the merits of Apple's defenses and WARF's infringement claims at trial.[19]

---

[17] Perhaps most telling, the few defenses that Apple moved on the merits do not offer grounds for the court to find for Apple on the objective prong of WARF's willful infringement claim.

[18] Certainly, Apple seems to make an objectively reasonable argument as to claims 1 and 2 being obvious, and perhaps as to claims 3 and 9, but the court cannot say on this record whether the supposed links drawn between Steely, Hesson, Chen and EV6 are obvious or pure sophistry. Similarly, while Apple raises a number of arguments that appear to objectively establish non-infringement on a literal basis, it has left the court unconvinced as to infringement under the doctrine of equivalents.

[19] To clarify, while the jury is deliberating on liability, the court can take up any additional evidence and argument relevant to the objective prong of WARF's willful infringement claim and likely will render a decision on the objective prong before the

Returning to those bases which *were* fully briefed for review on the merits, Apple's claim construction is arguably "objectively reasonable" if viewed purely in a vacuum. Apple presented some evidence that "prediction" *can* describe a static value in the context of computer circuits and speculation, for example, in the form of the Smith article and Chrysos paper; they are also correct that the patent does not *explicitly* define "prediction," ostensibly leaving at least some room for debate. The problem is that *nothing* in the patent -- not the claim language, not the specification, not the purpose of the invention -- supports Apple's construction. As discussed above, the claim language from the outset suggests that a prediction must be dynamic in the context of this particular invention. The specification, including both the brief summary of the invention and the detailed description of the embodiments, further supports this construction. And Apple's resort to extrinsic evidence fails to render its arguments to the contrary any more reasonable, given that extrinsic evidence cannot be used to vary the intrinsic evidence under settled principles of claim construction.

While superficially appealing, not unlike a siren's song, Apple's construction crashes against the rocks of the patent language itself and intrinsic evidence. Given how strongly the patent itself supports WARF's narrower construction, and how little Apple has to offer in support of its broader one, Apple's position is

---

parties present any evidence on the subjective prong during the second phase of the trial (assuming the jury finds infringement and does not find invalidity).

not objectively reasonable. *Compare Cohesive Techs., Inc. v. Waters Corp.*, 543 F.3d 1351, 1374 (Fed. Cir. 2008) (no willfulness where disputed term "was susceptible to a reasonable construction under which [the] products did not infringe"), *with SSL Services, LLC v. Citrix Systems, Inc.*, 769 F.3d 1073, 1091 (Fed. Cir. 2014) (affirming district court's finding of willful infringement, in part, because defendant's non-infringement defense based on an unwarranted limitation of a claim term was not objectively reasonable); *cf. Raylon, LLC v. Complus Data Innovations, Inc.*, 700 F.3d 1361, 1369 (Fed. Cir. 2012) (finding position on claim construction frivolous under Rule 11 where proffered construction was "contrary to all the intrinsic evidence and does not conform to the standard canons of claim construction"). Finally, while Judge Crabb granted summary judgment to the defendant in *Intel* on WARF's willful infringement claim, she did so on a basis unrelated to claim construction and one not before this court. *Intel*, 656 F. Supp. 2d at 924 (finding defendant's licensing defense objectively reasonable). Accordingly, the court will deny defendant's motion for summary judgment on plaintiff's willful infringement claim that depends on Apple's claim construction, finding this defense objectively unreasonable.

As for Apple's anticipation challenge to the validity of the '752 patent based solely on Steely, the court finds this defense not objectively reasonable as well, though it will reserve on any obviousness defense involving Steely. Much of Apple's anticipation argument depended upon its claim construction, which was not objectively reasonable as discussed

above. Admittedly, Apple attempted to maintain an anticipation defense even under WARF's construction, but its dependence on Steely's purported, defective "tag overwriting" scheme is likewise unreasonable, given that this dubious overwriting defect would simply *dispose* of previous predictions, rather than "updating" them to reflect an increased likelihood of future mis-speculation.

The court will also deny Apple's motion with respect to its indefiniteness defense. Apple points to *no* case in which a lack of antecedent basis led to a finding of indefiniteness *despite* clear context providing that basis in the specification. Even the case law Apple cites explain that there is no invalidity for indefiniteness so long as the antecedent basis is present by implication, and Apple waived any contention that the specification did *not* serve to provide such context. As a whole then, this defense was not objectively reasonable, and Apple cannot use it to escape the possibility of enhanced damages.

ORDER

IT IS ORDERED that:

1) Plaintiff Wisconsin Alumni Research Foundation's construction of the disputed term "prediction" is ADOPTED as described in this opinion.

2) Defendant and counter claimant Apple, Inc.'s motion for summary judgment (dkt. #116) is DENIED as to its counterclaims and defenses of anticipation by Steely and indefiniteness, and DENIED as to plaintiff's willful

infringement claim premised on (1) Apple's claim construction, (2) anticipation by Steely, and (3) indefiniteness of claims 5 and 6. The court RESERVES on the motion in all other respects.

3) Plaintiff's motion for summary judgment (dkt. #117) is GRANTED. Entered this 5th day of August, 2015.

BY THE COURT:

/s/
WILLIAM M. CONLEY
District Judge

APPENDIX E – Excerpts from Trial Transcript: Testimony of Dr. Thomas Conte (October 6, 2015)

UNITED STATES DISTRICT COURT

FOR THE WESTERN DISTRICT OF WISCONSIN

Case No. 14-CR-62-WMC

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff,

-vs-

APPLE, INC.,

Defendant.

Stenographic Transcript of
Second Day of Jury Trial Held Before
Chief Judge William M. Conley, and a Jury

Madison, Wisconsin

October 6, 2015

Dkt. 665

Excerpts from Pages 147-148 & 157-158

* * *

BY MR. SHEASBY:

Q Applying the plain meaning of a particular data consuming instruction, how did you determine whether this limitation was met in Apple's LSD predictor?

A Sure. Well --

MR. SHEASBY: Your Honor, may I publish the slide?

THE COURT: Yes.

THE WITNESS: Sure. Well, remember what I was saying about you read a patent claim like a book. So if you're reading the patent like a book, you'll see here it talks about detecting a mis-speculation where a data consuming instruction. That's the load involved in the mis-speculation.

Now, down here it's talking about predictor receiving a mis-speculation indication from the data speculation circuit, that's what's up here, to produce a prediction associated with the particular data consuming instruction. So that's referring back to the load, the load that produced the mis-speculation.

* * *

Q Now, did you analyze the meaning of particular as part of your report?

A Yes, I believe I did.

Q And to a person of ordinary skill in the art does particular require uniqueness or one and only?

MR. LEE: Your Honor --

THE COURT: We'll have a brief sidebar.

(Discussion at sidebar at 12:25 p.m.)

THE COURT: Did his report indicate that those skilled in the art would understand the meaning of particular?

MR. SHEASBY: Yes.

THE COURT: All right. And ultimately it is the jury who will decide what the ordinary or plain meaning is of a term of those skilled in the art. I understand that ultimately it's my obligation to construe something which does not have its plain meaning. But I'm not -- I don't understand the basis for your objection.

MR. LEE: My objection was not that. My objection was leading because we're at a very important place.

THE COURT: I'm fine with that. Thank you. And if you would just state your objection, that would be fine.

(End of sidebar at 12:26 p.m.)

THE COURT: I'll sustain the objection to leading. And you may rephrase.

BY MR. SHEASBY:

Q To a person of ordinary skill in the art, how does the concept of uniqueness or one and only relate to particular?

A I don't think that they're equivalent to one of ordinary skill in the art and I think particular just is identifying in this case an association.

Q Particular is referring to the first chapter; is that correct?

A That's correct. That was the association with the load instruction that you detected the misspeculation.

APPENDIX F – District Court Order Denying WARF Motion to Exclude Evidence and Argument (October 8, 2015)

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

WISCONSIN ALUMNI
RESEARCH FOUNDATION,

Plaintiff,

v.

APPLE, INC.

Defendant.

OPINION AND ORDER

14-cv-062-wmc

On the second day of trial, Apple described a theory of noninfringement based on its understanding of the meaning of "the particular" in claim 1(a) of the '752 patent. WARF objected during the hearing to Apple presenting expert testimony as to the meaning of this term, arguing that the issue was one of claim construction and Apple waived any construction of that term by failing to raise it timely. While this issue was only presented to the court during trial, and indeed was raised for the first time the night before Apple intended to develop this theory through an expert's testimony, the dispute has been known to the parties for some time. For that reason, it is unclear who is more at fault. Apple's expert Dr. August articulated its reading of claim 1a in his March 2015 report (August Noninfringement Rept. (dkt. #103) ¶¶

242-61) and reiterated the theory in his August 2015 deposition (August Dep. (dkt. #238) 51-52). As such, WARF knew -- or at least should have known -- that its own interpretation of the term "the particular" differed from Apple's. On the other hand, the meaning of this term now appears central to Apple's theory of noninfrigement, and therefore Apple should have raised the need for a construction once the dispute presented itself.

Regardless of who is to blame, unless the court were to find waiver, some construction is required for the same reasons previously articulated by the court in its ruling on WARF's motion *in limine* 12. (9/28/15 Op. & Order (dkt. #464) 26.) "When the parties present a fundamental dispute regarding the scope of a claim term, it is the court's duty to resolve it." *O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351, 1362-63 (Fed. Cir. 2008)) (*citing Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 976 (Fed. Cir. 1995) (*en banc*), *aff'd* 517 U.S. 370 (1996) (purpose of claim construction is to "determin[e] the meaning and scope of the patent claims asserted to be infringed")). This is true even if one or both of the parties insist, as in this case, that the claim terms should be given their "plain" or "ordinary" meanings. *O2 Micro Int'l Ltd.*, 521 F.3d at 1361-62 ("plain meaning" construction not useful if "reliance on a term's 'ordinary' meaning does not resolve the parties' dispute").

After this issue was raised at a hearing outside of the presence of the jury, both sides filed briefs explaining their respective positions. (Dkt. ##550, 552.) The dispute concerns the following language of claim 1:

> In a processor capable of executing program instructions in an execution order differing from their program order, the processor further having a data speculation circuit for detecting data dependence between instructions and detecting a mis-speculation where ***a data consuming [load] instruction*** dependent for its data on a data producing instruction of earlier program order, is in fact executed before the data producing instruction, a data speculation decision circuit comprising:
>
> a) a predictor receiving a mis-speculation indication from the data speculation circuit to produce a prediction associated with ***the particular data consuming [load] instruction*** and based on the mis-speculation indication; and
>
> b) a prediction threshold detector preventing data speculation for instructions having a prediction within a pre-determined range.

'752 patent, 14:36-52 (emphasis added).

Both parties agree that "the particular data consuming instruction" in subpart (a) refers back to "a data consuming instruction" in the preamble of claim 1. WARF's interpretation would stop there. Apple, on the other hand, argues that the word "the" is sufficient alone to refer back to the prior reference to "a data consuming instruction" and, therefore, "[t]he word 'particular' must have some additional meaning." (Def's Br. (dkt. #552) 3.) Apple then contends that "particular" has its plain and ordinary meaning, directing the court to dictionary definitions of "particular" as "of, relating to, or being a single person

or thing distinguished from some or all others," or "relating to or considered as one thing or person as distinct from others." (*Id.*, Exs. 1, 2 (dkt. ##552-1, 552-2).) From this, Apple concludes that subpart (a) discloses a predictor associated with one and only one load instruction. This interpretation of the claim may be material because Apple contends that the LSD Predictor in the accused products "treats load instructions collectively and does not associate a prediction with 'the particular' load instruction." (Def.'s Br. (dkt. #552) 4 (summarizing August's likely testimony).)

In a way, this dispute may be a red herring, since the parties do *not* dispute that the language in subpart (a) refers back to the preamble description of "*a* data consuming instruction." From the court's reading of claim 1 as a whole, it contemplates a single load instruction. Perhaps recognizing this, WARF alternatively argues that the word "comprising" at the end of the preamble to claim 1 supports its position:

> [B]ecause claim 1 uses the transition "**comprising**," it merely requires the prediction to be associated with ***at least*** the particular Load that mis-speculated. It ***does not exclude*** predictions from ***also*** being associated with one or more ***additional*** limitations. In other words, if the prediction is associated with *at least* the Load that mis-speculated, then the presence or absence of any further association is irrelevant to the analysis.

(Pl.'s Br. (dkt. #550) 4 (emphasis in original).) While the court agrees with WARF -- and will instruct the jury that the word "comprising" means that "the

invention includes the listed requirements, but is not limited to those requirements" (Draft Jury Instructions at 5) -- "comprising" in claim 1 precedes the two subparts, requiring, in other words, that the invention contain both subparts (a) *and* (b). If an accused product has (a), (b) and (c), that would still constitute infringement. WARF, however, cannot rely on "comprising" to expand what is claimed in subpart (a).

From all of this, the court concludes that claim 1 discloses a prediction associated with a single load instruction, albeit one that is "dynamic." Because this language is consistent with the plain meaning of the claim terms "the" and "the particular," the court concludes there is no need for instructing the jury on the meaning of this term. Dr. August is free to rest his analysis on this plain, general reading of claim 1 without over emphasizing the importance of the terms "the particular" in subpart 1a.

ORDER

IT IS ORDERED that WARF's motion to exclude evidence and argument on Apple's hash "aliasing" non-infringement theory (dkt. #550) is DENIED.

Entered this 8th day of October, 2015.

BY THE COURT:

/s/

_____

WILLIAM M. CONLEY
District Judge

APPENDIX G – Apple Motion for a Jury Instruction on Claim Construction (October 9, 2015)

UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

WISCONSIN ALUMNI RESEARCH FOUNDATION, Plaintiff, v. APPLE INC., Defendant.	Case No. 3:14-cv-00062-WMC

APPLE'S REQUEST FOR A CLOSING INSTRUCTION CONSTRUING "A PREDICTION ASSOCIATED WITH THE PARTICULAR DATA CONSUMING INSTRUCTION"

Apple recognizes that the Court has likely ruled on this issue already. But for purposes of the record and in light of the Court's construction of the term "particular" (Dkt. 559), Apple respectfully requests that the Court include the following language in the closing jury instructions:

"a prediction associated with the particular data consuming instruction" means "a prediction associated with a single load instruction"

See O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co., 521 F.3d 1351 (Fed. Cir. 2008).

Dated: October 9, 2015 Respectfully submitted,

/s/ Derek Gosma

William F. Lee (*pro hac vice*)
Lauren B. Fletcher (*pro hac vice*)
Bryan S. Conley (*pro hac vice*)
WILMER CUTLER PICKERING
HALE AND DORR LLP
60 State Street
Boston, MA 02109
Tel: (617) 526-6000
william.lee@wilmerhale.com
lauren.fletcher@wilmerhale.com
bryan.conley@wilmerhale.com

David C. Marcus (*pro hac vice*)
Andrea W. Jeffries (*pro hac vice*)
Derek Gosma (*pro hac vice*)
WILMER CUTLER PICKERING
HALE AND DORR LLP
350 South Grand Avenue, Suite 2100
Los Angeles, CA 90071
Tel: (213) 443-5300
david.marcus@wilmerhale.com
andrea.jeffries@wilmerhale.com
derek.gosma@wilmerhale.com

Mark D. Selwyn (*pro hac vice*)
WILMER CUTLER PICKERING
HALE AND DORR LLP
950 Page Mill Road
Palo Alto, CA 94304
Tel: (650) 858-6000
mark.selwyn@wilmerhale.com

Catherine Cetrangolo
CETRA LAW FIRM LLC
20 North Carroll Street, 2d Floor
Madison, WI 53703
Tel: 608-535-9220
Email: cetrangolo@cetralaw.com

Attorneys for Defendant Apple Inc.

APPENDIX H – District Court Order Denying Apple Motion for a Jury Instruction on Claim Construction (Oct. 9, 2015)

UNITED STATES DISTRICT COURT

FOR THE WESTERN DISTRICT OF WISCONSIN

Case No. 14-CR-62-WMC

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff,

-vs-

APPLE, INC.,

Defendant.

Text Only Order

Madison, Wisconsin

October 9, 2015

Dkt. 575

** TEXT ONLY ORDER **

The court is in receipt of Apple's "request for a closing instruction construing 'a prediction associated with the particular data consuming instruction.'"

(Dkt. # 572 .) Consistent with Apple's position, the court denied WARF's request to exclude expert testimony and argument on this issue. (See 10/8/15 Op. & Order (dkt. # 559).) The court also sided with Apple in its request that the term simply be given its plain and ordinary meaning. (*See Id.*; *see also* Apple's Opp'n (dkt. # 552) 2 (asking the court to give the term its plain meaning).) As such, Apple has waived any request to now insert a construction of the term into the closing jury instructions. In any event, for the reasons already explained in its prior opinion, plain meaning is sufficient. Apple's request is therefore DENIED. Signed by District Judge William M. Conley on 10/9/15. (jat) (Entered: 10/09/2015)

APPENDIX I – Excerpts from Trial Transcript: Closing Jury Instructions (October 9, 2015)

UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

Case No. 14-CR-62-WMC

WISCONSIN ALUMNI RESEARCH FOUNDATION,
Plaintiff,

-vs-

APPLE, INC.,
Defendant.

Stenographic Transcript of
Fifth Day of Jury Trial Held Before
Chief Judge William M. Conley, and a Jury

Madison, Wisconsin
October 9, 2015
Dkt. 668

Excerpts from Pages 236-243

* * *

THE COURT: Please be seated. Members of the Jury, I am about to read to you the closing instructions before closing arguments and as you'll see they're actually going to be displayed for you.

You're about to hear closing arguments of the parties. Before these arguments, I will instruct you on the law. After the closing arguments, I will provide very brief instructions governing your deliberations. After that, the case will be in your hands.

* * *

In patent law, the requirements of a claim are often referred to as "claim elements" or "claim limitations." For example, a claim that covers the invitation -- invention of a table may describe the tabletop, four legs, and glue that holds the legs and tabletop together. The tabletop, legs, and glue are each a separate limitation or requirement of the claim.

The claims are the main focus when a patent's validity is challenged. In this case, we are concerned with claims 1, 2, 3, 5, 6, and 9 of the '752 patent as reflected in your special verdict form.

WARF contends that Apple infringed these claims. Apple denies this and contends that these claims are invalid. When a product or process is accused of infringing a patent, it is the claims of the patent that must be compared to the accused product or process to determine whether or not there is infringement. It is the claims of the patent that are infringed when patent infringement occurs. The claims are also the main focus when a patent's validity is challenged. For example, in deciding whether a patent is invalid because it is anticipated

or obvious, you must compare the claims to the asserted prior art. In reaching your determinations with respect to infringement and invalidity, you must consider each claim of the patent separately.

Therefore, the first step in any patent case is to understand the meaning of the words used in the patent claims. It is my job as the judge to determine what the patent claims mean and instruct you about the meaning. You must accept the meanings I give you and use them when you decide whether or not the asserted claims of the asserted patents are infringed and whether or not the asserted claims of the asserted patents are invalid.

These are the specific claims I've construed. I have provided you with a copy of the '752 patent. Actually you'll have that when you get in deliberations. The following eight terms found in the claims are defined as follows:

Data consuming instruction, which you'll find in claims 1, 2, 3, 5, 6, and 9, means instruction that consumes data by obtaining data from memory, such as a load instruction.

Data producing instruction, also found in the same claims, means instruction that produces data by providing data to memory, such as a store instruction.

Data speculation circuit, which appears in claims 1 and 9, means a circuit that detects data dependence between load-and-store instructions and that detects mis-speculation by load instructions.

Mis-speculation in claims 1, 6, and 9 means when a load instruction that is dependent for its data on a store instruction appearing earlier in the program

order is in fact executed before the store instruction wrote its data to a memory address shared with the load instruction.

In fact executed, as you previously heard, appears in claims 1 and 9. It means when a load instruction has actually accessed a memory address that has not yet been updated by a store instruction appearing earlier in the program order.

Predictor in claim 1 means a circuit that receives a mis-speculation indication from the data speculation circuit to produce a prediction.

Prediction, claims 1, 2, 3, 5, and 9, means a variable that indicates the likelihood that the data speculative execution of a load instruction will result in a mis-speculation where a prediction must be capable of receiving ongoing updates.

Prediction table and synchronization table in claim 5 can be located in the same structure, including in a single table, and can share data and circuitry.

The asserted claims of the '752 patent all use or include or incorporate the term comprising. When a patent term uses the term comprising, it means that the invention includes the listed requirements, but is not limited to those requirements. All other claim terms should be given their plain and ordinary meaning as viewed from the perspective of a person of ordinary skill in the art or field of the invention.

APPENDIX J – District Court Rule 50(a) Opinion and Order (October 26, 2015)

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

WISCONSIN ALUMNI
RESEARCH FOUNDATION,

Plaintiff,

v.

APPLE, INC.,

Defendant.

OPINION AND ORDER

14-cv-062-wmc

Before the court are motions for judgment as a matter of law, pursuant to Federal Rule of Civil Procedure 50(a). (Dkt. ##555, 586, 634, 637.) Both parties filed two motions -- one for the liability phase of trial and one for the damages phase. The court implicitly denied or reserved on all of the motions -- save one of Apple's noninfringement theories, for which the court directed verdict in WARF's favor. (Dkt. #639.) Since the jury found for WARF on all liability and damages questions, its motions are effectively moot. In this opinion and order, the court will briefly set forth its reasons for denying Apple's motions for directed verdict on all claims asserted by WARF at trial, as well as touch briefly on part of WARF's motion as to Apple's damages theories. In its motion on the liability phase of trial, Apple also seeks

entry of judgment on claims WARF opted not to pursue at trial. For the reasons that follow, the court will deny that request. Finally, to the extent necessary, the court finds that all arguments raised timely at trial are preserved for purposes of pressing the same arguments in a Rule 50(b) motion.

OPINION

Under Federal Rule of Civil Procedure 50(a), a court may "enter judgment against a party who has been fully heard on an issue during a jury trial if 'a reasonable jury would not have a legally sufficient evidentiary basis to find for the party on that issue.'" *Schandelmeier-Bartels v. Chi. Park Dist.*, 634 F.3d 372, 376 (7th Cir. 2011) (quoting Fed. R. Civ. P. 50(a)). In considering a Rule 50 motion, the court is to "construe the facts strictly in favor of the party that prevailed at trial," including drawing "[a]ll reasonable inferences in that party's favor and disregarding all evidence favorable to the moving party that the jury is not required to believe." *May v. Chrysler Group, LLC*, 692 F.3d 734, 742 (7th Cir. 2012) (internal citations and quotation marks omitted). In particular, the court is not to make credibility determinations or weigh the evidence; it need only determine whether "more than 'a mere scintilla of evidence' supports the verdict." *Id.* (quoting *Hossack v. Floor Covering Assoc. of Joliet, Inc.*, 492 F.3d 853, 859 (7th Cir. 2007)). Said another way, the court's "job is to decide whether a highly charitable assessment of the evidence supports the jury's verdict or if, instead, the jury was irrational to reach its conclusion." *May*, 692 F.3d at 742. Under this highly deferential standard, Apple's 50(a) motions both fall short in all respects, as do WARF's for that matter.

I. Defendant's Motion on Liability (dkt. #555)

A. Abandoned Infringement Claims

Apple seeks judgment in its favor on claims that WARF opted not to pursue at trial. Specifically, Apple seeks judgment in its favor on WARF's infringement by equivalence claim and multiple indirect infringement claims. (Def.'s Mot. (dkt. #555) 2-4.) Apple fails to cite any caselaw in support of its motion, which was a failing already pointed out by the district court in *Apple, Inc. v. Samsung Electronics Co.*, 67 F. Supp. 3d 1100, 1117 (N.D. Cal. 2014).

In response, WARF directs the court to *Alcon Research Ltd. v. Barr Laboratories, Inc.*, 745 F.3d 1180, 1193 (Fed. Cir. 2014), which affirmed a district court's refusal to enter judgment on claims that were "not litigated or fairly placed in issue, during the trial." Instead, the Federal Circuit explained that the scope of any judgment as a matter of law "should conform to the issues that were actually litigated . . . during the trial." *Alcon* might be distinguished in that, here, Apple asserted a counterclaim for a declaration of non-infringement. (Def.'s Answ. & Countercl. (dkt. #19) ¶¶ 9-11). *See also Alcon*, 745 F.3d at 1193 (noting that defendant "never filed a counterclaim for declaratory judgment of noninfringement," and if it had, "the district court might have exercised its discretion differently"). At the same time, Apple never formally withdrew this counterclaim, it also never pressed a declaration of noninfringement on abandoned claims at trial, and for

good reason.[1] Accordingly, the court will deny Apple's request for a further noninfringement finding, while agreeing that any attempt by WARF to resurrect abandoned claims would be futile in light of preclusion rules. *See, e.g., Acumed LLC v. Stryker Corp.*, 525 F.3d 1319, 1326 (Fed. Cir. 2008) (holding that a claim for patent infringement is "barred by claim preclusion if that claim arises from the same transactional facts as a prior action").

B. Literal Infringement Claims

Apple also seeks a directed verdict of noninfringement under the literal infringement claim WARF pursued at trial. Apple argued that the accused products do not meet three elements of the asserted claims of the patent-in-suit: (1) "detecting a mis-speculation"; (2) "the particular data consuming instruction"; and (3) "flag value." (Def.'s Mot. (dkt. #555) 4-9.) As for the "detecting a mis-speculation" element, the court denied Apple's motion, among other reasons, because WARF had submitted sufficient evidence through Dr. Conte's testimony and the admissions of Apple's own engineers for a reasonable jury to conclude that Apple's Load-Store Device ("LSD") triggers a Store-Hit-Younger-Load redirect. (Pl.'s Opp'n (dkt. #586) 5-6.) Although there was conflicting factual and opinion testimony as to

---

[1] The court in the *Apple* case in the Northern District of California granted Apple's request in part, entering judgment of noninfringement for Apple's accused iPad products with respect to one claim of one of the asserted patents. Unlike here, however, this infringement claim was presented to the jury in the preliminary instructions, and as such was at issue during the trial, even though the jury did not consider it as part of its verdict. *Apple*, 67 F. Supp. 3d at 1125-27.

whether the LSD detected only the possibility of a "speculation" or also "mis-speculation," there was more than sufficient evidence for the jury to find the latter, and thus to find the first element satisfied.

As for the second element, the court rejected Apple's argument that the hashing function foreclosed a finding that "the particular data consuming instruction" was present in the LSD predictor. More specifically, WARF submitted evidence using the Cyclone specification from which a reasonable jury could conclude that a prediction was associated with a particular load instruction (even if that same prediction may be associated with other load instructions), or at least the jury could find that element regularly, if not almost always, met. (Pl.'s Opp'n (dkt. #586) 11.)

Third and finally, the court also rejected Apple's motion on the "flag value" element found in claim 5 of the patent-in-suit because WARF submitted sufficient, credible evidence through Dr. Conte's testimony and Apple's own engineers that the armed bit functions as and carries a flag value as called for by that claim. (Pl.'s Opp'n (dkt. #586) 15-16.)

II. Plaintiff's Motion on Damages (dkt. #634)

For its part, WARF sought a directed verdict on damages generally, asserting two core arguments in support. First, WARF contends that the court should have directed a verdict on the issue of whether WARF's license agreement with Intel represents an "established royalty" for the '752 patent. The precise nature of this request is unclear at best since the court instructed the jury, consistent with Federal Circuit

caselaw, on how to assess the proper weight, if any, it should place on WARF's license with Intel in determining a reasonable royalty in this case. (Closing Damages Instructions (dkt. #649) p.3.) Not only was the jury *not* instructed that the WARF-Intel license represented an "established royalty," but both sides were allowed to introduce substantial evidence and arguments as to its relative importance or unimportance in a hypothetical negotiation. The court finds no basis to grant plaintiff's motion for a directed damages verdict on this basis.

Second, WARF argues that it is somehow entitled to a monetary judgment as a matter of law because no reasonable jury could adopt any of Apple's damages theories. In particular, WARF criticizes the reliance of Apple's principal damages expert on her: (1) Intel analysis; (2) ARM analysis; and (3) a Base X Rate analysis. As an initial matter, it would appear that the jury ultimately agreed with WARF, or at least placed little weight on any of Apple's damages theories. Even if the jury credited Apple's damages theories in part, however, the court rejects WARF's motion to direct verdict on a damages amount for the same reasons it rejected much of WARF's motion to strike Apple's expert Julie Davis's testimony. (*See* 9/28/15 Op. & Order (dkt. #464) § I.O.) Consistent with the *Georgia Pacific* factors themselves, Davis was free to emphasize her own view as to which considerations were most likely to drive the parties' hypothetical negotiation, just as did WARF's principal damages expert, Catharine Lawton.

III. Defendant's Motion on Damages (dkt. #637)

Lastly, Apple moves for directed verdict on the damages phase of trial, arguing that it is entitled to judgment as a matter of law on (1) both questions concerning Samsung products, and (2) the damages award. The latter argument is apparently based on the assertion that WARF failed to submit sufficient evidence to justify its request for a 50/50 split of profits attributable to the patented invention. The court ultimately rejected both bases for the reasons explained below.

First, the court concluded that there was sufficient evidence from which a reasonable jury could find that Apple's products manufactured by Samsung in the United States and sent overseas infringe at the time they leave the United States. Specifically, the jury reasonably could credit WARF's evidence and argument that the processors were complete at that stage of the manufacturing process, rejecting Apple's evidence and essential argument that the LSD predictor must be fully capable of running software for the processors to be capable of infringing.

The court also found overwhelming evidence that WARF had proven that Apple exercised sufficient control over Samsung's manufacturing process to be liable for its infringement. If anything, the court considered directing a verdict in *WARF's* favor on this argument given the manufacturing terms of the contract and undisputed evidence as to how it functioned in practice. While the law on vicarious liability may be unsettled, the Federal Circuit is certainly moving in the direction of a more expansive view of what satisfies control and direction in order to bring a third-party's actions within the purview of the alleged infringer. *See Akamai Techs., Inc. v. Limelight*

Networks, Inc., 797 F.3d 1020 (Fed. Cir. 2015) (*en banc*). Here, the evidence demonstrates that Samsung manufactured chips for Apple, using Apple's specifications, to be tested according to Apple's requirements, and to be used only in Apple's end products. This is certainly sufficient to find that Samsung acted under Apple's control and direction.

Second, Apple sought a directed verdict on damages, arguing WARF's expert Lawton had used a 50/50 split of profits attributable to the licensed feature, akin to a prohibited 25% rule of thumb. (Def.'s Mot. (dkt. #637) 7.) This does a substantial disservice to the nuanced nature of Ms. Lawton's testimony. First, WARF presented credible testimony from its managing director, Carl Gulbrandsen, as the person responsible for negotiating a license with Apple, that a 50/50 split would have been reasonable from his perspective, and why. Second, unlike cases in which the Federal Circuit has rejected expert testimony, this was not Lawton's starting point, nor even central point in arriving at a reasonable royalty opinion in this case. It was just one of many pieces of evidence that Lawton relied upon in coming up with her own royalty rate. The jury was free to discount Gulbrandsen's personal view -- and may well have in light of the ultimate damages award -- but the court sees no error in allowing his testimony to be presented to and weighed by the jury, nor does the court credit Apple's argument that WARF's damages case was otherwise legally insufficient.

ORDER

IT IS ORDERED that:

1) Defendant Apple's motion for judgment as a matter of law on liability (dkt. #555) is DENIED;

2) Plaintiff Wisconsin Alumni Research Foundation's motion for judgment as a matter of law on liability (dkt. #586) is DENIED as moot;

3) Plaintiff's motion for judgment as a matter of law on damages (dkt. #634) is DENIED; and

4) Defendant's motion for judgment as a matter of law on damages (dkt. #637) is DENIED.

Entered this 26th day of October, 2015.

BY THE COURT:

/s/

WILLIAM M. CONLEY
District Judge

APPENDIX K – District Court Rule 50(b) Opinion and Order (June 6, 2017)

IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF WISCONSIN

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff, OPINION AND ORDER

v. 14-cv-062-wmc

APPLE, INC.,

Defendant.

In this opinion and order, the court addresses a slew of post-trial motions. The jury returned a verdict in favor of plaintiff Wisconsin Alumni Research Foundation ("WARF") and awarded damages in the amount of $234 million. Invoking Federal Rules of Civil Procedure 50(b) and 59(3), defendant Apple, Inc., challenges virtually every aspect of the jury's verdicts, and myriad decisions made both before and during the trial by the court. (Dkt. #677.) For the reasons that follow, the court will deny that motion in its entirety. WARF also moves under Rule 59(e) to alter the court's grant of judgment in Apple's favor on plaintiff's willful infringement claim. Applying the new standard articulated by the United States Supreme Court in *Halo Electronics, Inc. v. Pulse Electronics, Inc.*, 136 S. Ct. 1923 (2016), the court again concludes that WARF has failed to

meet its burden of demonstrating willful infringement. Accordingly, that motion will also be denied.

The remaining motions are all WARF's: for equitable relief (dkt. #683); for an accounting, supplemental damages through the date of judgment, and pre- and post-judgment interest (dkt. #685); and for taxation of costs (dkt. ##689, 725). For the reasons that follow, the court will award an ongoing royalty rate of $2.74 per unit from the date of judgment, October 25, 2015. The court will also award supplemental damages at the per unit royalty rate awarded by the jury from June 27, 2015, to October 25, 2015.[1] The court will also award pre-judgment interest at the prime rate, compounded quarterly, and will award post-judgment interest at the statutory rate, compounded annually. The calculations for supplemental damages and pre-judgment interest will await further submissions by the parties. Finally, the court will award WARF costs in the total amount of $841,587.66.

BACKGROUND

In this patent lawsuit, WARF alleged that Apple infringed U.S. Patent No. 5,781,752 (the "'752 patent"). In response, Apple asserted various counterclaims, which challenge the validity of the patent. On the parties' cross motions for summary judgment, the court granted partial judgment to

---

[1] As described below, the court will also consider awarding supplemental damages for the A9 and A9x chips, which Apple now concedes infringe (while maintaining its objections to the jury's verdict). This issue will, however, require additional briefing.

WARF on: (1) Apple's counterclaims and defenses for anticipation under 35 U.S.C. § 102 with respect to U.S. Patent No. 5,619,662 ("Steely or the "Steely patent"); and (2) Apple's counterclaim and defense for indefiniteness under 35 U.S.C. § 112 ¶ 2 with respect to claims 5 and 6 of the '752 patent. (8/6/15 Op. & Order (dkt. #193).)

The case then proceeded to a jury trial. The jury returned a verdict in favor of WARF, finding that Apple infringed all six of the asserted claims and rejecting Apple's invalidity defense as to each of those six claims. (10/13/15 Liability Special Verdict (dkt. #603).) In the second phase of the trial, the jury answered two more questions in favor of WARF, finding Apple vicariously liable for Samsung's manufacture of Apple products, and awarded WARF $234,277,669.00 in damages. (10/19/15 Damages Special Verdict (dkt. #642).)

During the course of trial, the court also granted WARF judgment as a matter of law on one of Apple's noninfringement defenses based on the claim of a "prediction threshold detector preventing data speculation for instructions having a prediction within a predetermined range," finding that Apple had failed to put forth a factual basis for that defense to support a reasonable jury finding noninfringement on that basis. (10/16/15 Op. & Order (dkt. #639).) Finally, the court granted judgment in favor of Apple on WARF's willful infringement claim. (10/15/15 Op. & Order (dkt. #623).)

OPINION

I. Apple's Renewed Motion for Judgment as a Matter of Law and/or New Trial (dkt. #677)

Under Federal Rule of Civil Procedure 50, judgment as a matter of law may be granted where there is no "legally sufficient evidentiary basis" to find for the party on that issue. Fed. R. Civ. P. 50(a). In considering a Rule 50(a) motion, the court is to "construe the facts strictly in favor of the party that prevailed at trial." including drawing "[a]ll reasonable inferences in that party's favor and disregarding all evidence favorable to the moving party that the jury is not required to believe." *May v. Chrysler Group, LLC*, 692 F.3d 734, 742 (7th Cir. 2012) (internal citations and quotation marks omitted), *withdrawn in part on reh'g,* Nos. 11-3000, 11-3109, 2013 WL 1955682 (7th Cir. May 14, 2013). In particular, the court does not make credibility determinations or weigh the evidence, although the court must assure that "more than 'a mere scintilla of evidence' supports the verdict." *Id*. (quoting *Hossack v. Floor Covering Assocs. of Joliet, Inc.*, 492 F.3d 853, 859 (7th Cir. 2007)). Essentially, the court's "job is to decide whether a highly charitable assessment of the evidence supports the jury's verdict or if, instead, the jury was irrational to reach its conclusion." *May*, 692 F.3d at 742.

A further limitation applies as well: "Because the Rule 50(b) motion is only a renewal of the preverdict motion, it can be granted only on grounds advanced in the preverdict motion." *Wallace v. McGlothan*, 606 F.3d 410, 418 (7th Cir. 2010); *see also Thompson v. Mem'l Hosp. of Carbondale*, 625 F.3d 394, 407 (7th Cir. 2010) (refusing to consider the defendant's argument that plaintiff failed to demonstrate that he

suffered an adverse employment action, in part, because the defendant did not raise argument in Rule 50(a) motion); *see also* Fed. R. Civ. P. 50 cmt. 1991 Amendments ("A post-trial motion for judgment can be granted only on grounds advanced in the pre-verdict motion.").

Defendant also moves for a new trial under Federal Rule of Civil Procedure 59. "A new trial may be granted only if the jury's verdict is against the manifest weight of the evidence." *King v. Harrington*, 447 F.3d 531, 534 (7th Cir. 2006) (citing *ABM Marking, Inc. v. Zanasi Fratelli, S.R.L.*, 353 F.3d 541, 545 (7th Cir. 2003)). To meet this standard, defendant must demonstrate that *no* rational jury could have rendered a verdict against Apple. *King*, 447 F.3d at 534 (citing *Woodward v. Corr. Med. Servs. of Ill., Inc.*, 368 F.3d 917, 926 (7th Cir. 2004)). In making this evaluation, the court must view the evidence in a light most favorable to plaintiff, leaving issues of credibility and weight of evidence to the jury. *King*, 447 F.3d at 534. "The court must sustain the verdict where a 'reasonable basis' exists in the record to support the outcome." *Id.* (quoting *Kapelanski v. Johnson*, 390 F.3d 525, 530 (7th Cir. 2004)).

In this case, the evidence easily supports the jury's findings of infringement on all three disputed elements of the asserted claims of the patent-in-suit: (1) "detecting a mis-speculation"; (2) "the particular data consuming instruction"; and (3) "flag value." (Def.'s Mot. (dkt. #555) 4-9.) While Apple's Rule 50(b) motion extends beyond the arguments raised in its 50(a) motion, WARF does not oppose it on that basis. As such, the court will address that motion, briefly,

while ultimately rejecting all of the arguments raised by Apple.

A. Infringement

In its pending Rule 50(b) motion, Apple contends the accused products lack at least three elements required by the '752 patent, and, therefore, no reasonable jury could find that Apple literally infringed any of the asserted claims of the '752 patent. *First,* Apple contends that no reasonable jury could have found that Apple's accused products satisfy the "detecting a mis-speculation" and "mis-speculation indication" elements. Apple argues that the accused products detect only data dependence, not mis-speculations, but WARF submitted evidence -- largely through its expert Professor Conte -- showing Apple's data speculation circuit, the Load-Store Unit, is capable of both detecting data dependence and detecting mis-speculation.

Specifically, Conte testified that the Load Queue in the Load-Store Unit detects a mis-speculation between a Load and a Store by comparing the program order of the older Store and younger Load, confirming that the instructions have an address overlap, and ensuring that the younger Load has in fact executed before the older Store. (*See* Pl.'s Opp'n (dkt. #711) 21 (citing Conte testimony).) When a mis-speculation is detected, the Load Store Unit then produces a Store-Hit-Younger-Load Redirect, which is a mis-speculation indication. (*See id.*) Conte further testified that, and provided an illustration for the jury to better understand how, the timing of Apple's processor necessarily satisfies the "detecting a mis-speculation" and "mis-speculation indication"

elements. (*Id*. at 31-32.) In other words, this step is "baked into" Apple's processor. (*Id*. at 31.)

In its reply, Apple argues principally that "a processor cannot literally detect or indicate a mis-speculation absent an *explicit* check." (Def.'s Reply (dkt. #728) 15 (emphasis added).) Whatever Apple means by "explicit," the construction of "mis-speculation" agreed on by the parties contains no such requirement. (*See* 8/6/15 Op. & Order (dkt. #193) 10.) Instead, true to the language of the claim, all that is required is that (1) mis-speculations occur and (2) Apple's processor is capable of detecting and indicating such occurrences. Drawing all reasonable inferences in favor of WARF, therefore, the jury's finding of infringement of this element was certainly *not* irrational. *May*, 692 F.3d at 742.

Second, Apple challenges the jury's finding that the "particular data consuming instruction" element is satisfied. Specifically, Apple argues that the evidence demonstrates that "each entry of Apple's LSD Predictor includes a Load Tag and counter," and that the Load Tags are generated using a hashtag function and are not associated with a "particular" load instruction. (Def.'s Opening Br. (dkt. #678) 17-18.) The court rejects Apple's challenge for the same reason it denied Apple's Rule 50(a) motion, "a reasonable jury could conclude that a prediction was associated with a particular load []instruction even if that same prediction may be associated with other load instructions." (10/26/15 Op. & Order (dkt. #655) 4-5.) In its reply brief, Apple appears to step back from any defense based on aliasing and the frequency of aliasing. Even if this shift does not amount to waiver, the court agrees with WARF that the Apple's

processors are capable of operating for periods of time during which at least some of the load tags will not alias. As such, the jury reasonably rejected any non-infringement defense based on that theory. (*See* Pl.'s Opp'n (dkt. #711) 44 (citing *Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1333 (Fed. Cir. 2013) ("[A]n accused device that sometimes, but not always embodies a claim[] nonetheless infringes." (internal citation and quotation marks omitted)).) For these reasons, the court concludes that the jury reasonably found this element satisfied.

Related to its Rule 50(b) challenge, Apple also seeks a new trial based on the court's failure to instruct the jury on the meaning of "the particular." During trial, WARF moved to exclude evidence and argument on Apple's aliasing non-infringement theory, on the basis that this theory was one of claims construction, and Apple waived any construction of the term "the particular" by failing to raise it timely. In response, it was *Apple* who argued that the term "the particular" should be given its plain and ordinary meaning of "of, relating to, or being a single person or thing," no doubt at least in part to avoid a waiver for failing to seek a timely construction. (Def.'s Br. (dkt. #552) 3.) Regardless, the court agreed with Apple's interpretation, concluding that claim 1 "disclosed a prediction associated with a single load instruction," but given that this interpretation was consistent with the plain meaning of the claim terms "the" and "the particular" declined to insert a specific, untimely construction in the closing instructions. (10/8/15 Op. & Order (dkt. #559); 10/9/15 Text Order (dkt. #575).) Moreover, in the closing instructions, the jury was told that "[a]ll other claim terms should be given their

plain and ordinary meaning as viewed from the perspective of a person of ordinary skill in the art or field of the invention." (Closing Liability Instructions (dkt. #646) 5.) Finally, Apple has failed to explain adequately how it was prejudiced by the denial of its request, or why a new trial is required under Rule 59(a).

Third, Apple challenges the jury's finding that Apple's accused products satisfy the "flag value indicating whether the certain respective date producing [store] instruction has been executed" in claims 5 and 6. Apple argues that the Armed Bit in the LSD Predictor only indicates "whether the store instruction is in the Reservation Station, not whether it 'has been executed'" as would be required to have a "flag value" under those claims. (Def.'s Opening Br. (dkt. #678) 22.) As WARF explains, largely through Dr. Conte's review of RTL code, however, the change from 0 to 1 indicates that the Store instruction is in the reservation and not yet executed, and then the change from 1 to 0, further indicates that the stores are "data resolved," issued from the Reservation Station, and thus have been executed. (Pl.'s Opp'n (dkt. #711) 48-49.) In response, Apple simply argues that WARF's explanation does not account for the Armed Bit value being "0" both before and after execution. (Def.'s Reply (dkt. #728) 25-26.) Apple's argument, however, fails to consider the passage of time and Conte's testimony that the value changed from 0 to 1 and then the de-assertion from 1 to 0. Since the jury reasonably credited Conte's testimony, the court sees no basis for upsetting that finding.

In addition to raising challenges under Rule 50(b), Apple also argues that a new trial on infringement is

necessary because the court erred during the liability phase of trial, granting WARF judgment as a matter of law on one of Apple's non-infringement theories -- namely, Apple's defense with respect to the "prediction threshold detector" limitation. (Def.'s Opening Br. (dkt. #678) 29.) Certainly, the court granted judgment as a matter of law on this noninfringement theory in response to WARF's oral motion. The court subsequently issued an opinion and order more fully explaining its reasons for doing so. (10/16/15 Op. & Order (dkt. #639).) Since Apple's Rule 59 motion raises no new bases for reviewing that decision, the court simply rejects Apple's arguments for the reasons already stated on the record during the trial and in its subsequent written order.

Finally, in a one-paragraph throw-away challenge, Apple purports to seek a new trial on the basis that the jury's infringement verdict was against the manifest weight of the evidence. The court rejects this motion for the same reasons the court rejected Apple's challenges under Rule 50(b).

B. Invalidity

Apple also seeks judgment as a matter of law as to its defenses and counterclaims of invalidity with respect to two other arguments: (1) claims 1, 2, 3, 5, 6 and 9 of the '752 patent were obvious in view of U.S. Patent No. 5,666,506 ("Hesson") and U.S. Patent No. 5,619,662 ("Steely"); and (2) claims 1, 2, 3 and 9 of the '752 patent are anticipated by, or at least obvious in view of, the Chen prior art references.

"A party seeking to invalidate a patent based on obviousness must demonstrate 'by clear and convincing evidence that a skilled artisan would have

been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.'" *Procter & Gamble Co. v. Teva Pharm. USA, Inc.*, 566 F.3d 989, 994 (Fed. Cir. 2009) (quoting *Pfizer, Inc. v. Apotex, Inc.,* 480 F.3d 1348, 1361 (Fed. Cir. 2007)). Again, in the context of a post-verdict motion, the court is directed to "presume that the jury resolved the underlying factual disputes in favor of the verdict winner and [to] leave those presumed findings undisturbed if they are supported by substantial evidence." *Jurgens v. McKasy,* 927 F.2d 1552, 1557 (Fed. Cir. 1991). The court then "examine[s] the legal conclusion *de novo* to see whether it is correct in light of the presumed jury fact findings." *Id.* In conducting this analysis, courts are instructed to "consider all of the *Graham* factors prior to reaching a conclusion with respect to obviousness." *Kinetic Concepts, Inc. v. Smith & Nephew, Inc.,* 688 F.3d 1342, 1360 (Fed. Cir. 2012) (citing *Graham v. John Deere Co.,* 383 U.S. 1, 17-18 (1966) describing four underlying factors for deciding obviousness: (1) the scope and content of the prior art; (2) the differences between the claims and the prior art; (3) the level of ordinary skill in the art; and (4) objective indicia of nonobviousness).

As for the first defense, based on the Hesson and Steely patents, the court considered this theory in its opinion and order granting Apple judgment on WARF's willful infringement claim. (10/15/15 Op. & Order (dkt. #623) 3-4.) While the court found Apple's defense reasonable, the court also determined that the jury acted reasonably in rejecting the defense, likely "because of the investment in computer

software simulations, and time and effort required to confirm that the '752 patented invention would prove valuable in practice sometime in the future when processing speeds had increased by a factor of 10 or more." (*Id.* at 4.) As WARF's expert, Dr. Mudge, and others explained, WARF put forth sufficient evidence from which the jury reasonably could have found that one skilled in the art would not have combined Steely's memory reference tags with Hesson's solution to the problem of mis-speculations of store instructions. At the very least, Apple failed to prove this defense by clear and convincing evidence. Once again, therefore, the court sees no basis for upsetting the jury's factual findings as to Apple's obviousness defense by combining the Steely and Hesson patents.

Apple also seeks a favorable judgment on its anticipation / obviousness defense based on the Chen references. Through the testimony of the co-author of three of the four Chen references -- Professor Scott Mahlke -- WARF put forth sufficient evidence from which the a reasonable jury could have found, and indeed did find that (1) the reordering contemplated in Chen is done in a software compiler *not* by a processor; and (2) the focus of Chen is on in order, rather than out of order, executions. (Pl.'s Opp'n (dkt. #711) 134-40.) Based on this, the jury reasonably concluded that Chen did not disclose claims 1, 2, 3 and 9 of the '752 patent, or, again, the jury at least had a sound basis for finding that Apple had not met its burden of demonstrating an anticipation defense by clear and convincing evidence.

As for its related obviousness defense, Apple argued that it would have been an obvious to apply Chen to software. Putting aside WARF's challenges to

Apple's anticipation defense, experts and others skilled in the art testified persuasively to the deep divide between hardware-based and software-based approaches. (*Id*. at 143-33.) A jury reasonably could have relied on this testimony to conclude that one skilled in the art would not have combined these fundamentally different approaches. As such, the court also rejects Apple's motion to invalidate certain claims of the '752 patent based on obvious in view of the Chen prior art references.

C. Vicarious Liability

Next, Apple seeks judgment as a matter of law or a new trial as to the jury's finding that: (1) the wafers infringe when they leave the United States; and (2) Samsung's manufacturing is attributable to Apple. These findings were material to the jury's determination of damages, specifically the appropriate royalty base. Post-trial, Apple raises three core challenges with respect to these jury findings.

First, Apple argues that no reasonable jury could have found that the wafers were capable of infringing before they left the United States for further processing overseas. Specifically, Apples argues -- as it did unsuccessfully to the jury -- that the wafers "are not capable of performing the claimed functionality until 'bumping' occurs to apply power in order for the circuitry to be able to function." (Def.'s Opening Br. (dkt. #678) 51; *see also* Def.'s Reply (dkt. #728) (arguing that the wafers also must be "fused and singulated" before they can satisfy the claim limitations").)

This argument ignores the evidence presented largely through the testimony of Dr. Conte, that "the circuitry in Accused Processors are defined by their RTL code, that this circuitry contains each of the elements specific in the asserted claims, and *all* of the circuitry" in the wafers is in place before being shipped overseas. (Pl.'s Opp'n (dkt. #711) 149.) Moreover, Dr. Conte testified that the wafers can be powered on and tested on a test fixture after the wafer is manufactured in Austin, *but before being shipped overseas*. (*Id*. at 154.) This evidence provided a sufficient basis for the jury to find that the wafers were capable of infringement and indeed did infringe, before leaving the United States.

Second, Apple argues that no reasonable jury could find that Samsung's manufacturing was attributable to Apple. In so arguing, Apple contends that plaintiff must prove that Apple and Samsung have a principal-agent relationship. (Def.'s Opening Br. (dkt. #678) 55.) The court considered and previously rejected this argument as well, relying principally on the Federal Circuit's *en banc* decision in *Akamai Technologies, Inc. v. Limelight Networks, Inc.*, 797 F.3d 1020 (Fed. Cir. 2015). In that case, the court explained that "Section 271(a) is not limited solely to principal-agent relationships, contractual arrangements, and joint enterprise[.]" *Id*. at 1023. Instead, the court held that vicarious liability for purposes of establishing direct infringement "can also be found when an alleged infringer conditions participation in an activity or receipt of a benefit upon performance of a step or steps of a patented method and establishes the manner or timing of that performance." *Id*.

While recognizing that there may be important distinctions between apparatus and method claims, the court rejects Apple's attempt to distinguish *Akamai* on the basis that its holding is limited to method claims. *See Centillion Data Sys., LLC v. Qwest Commc'ns Int'l, Inc.*, 631 F.3d 1279, 1282 (Fed. Cir. 2011) (similarly focusing on "control or direct" actions of another in determining whether there was vicarious liability for non method claim). As the court explained in its opinion and order on Apple's Rule 50(a) motion, "the Federal Circuit is certainly moving in the direction of a more expansive view of what satisfies control and direction in order to bring a third party's actions within the purview of the alleged infringer." (10/26/15 Op. & Order (dkt. #655) 7.)[2]

Finally, Apple argues that if the "control or direction" standard is correct, WARF still failed to demonstrate that it controlled and directed Samsung, because Samsung independently determined to manufacture the wafers in the United States as opposed to doing so in Korea. The court agrees with WARF that this argument conflates two separate questions -- whether Samsung's actions are attributable to Apple and whether infringement occurred within the United States. (Pl.'s Opp'n (dkt. #711) 187.) There is *no* merit to Apple's argument that it needed to direct Samsung to manufacture the wafers in the United States. As for Apple's other

---

[2] For this same reason, the court rejects Apple's arguments that a new trial is warranted because of the court's instruction to the jury that Apple was vicariously liable "if Samsung performed an act of infringement under Apple's control or direction" and because the evidence goes against the manifest weight of the evidence. (Def.'s Opening Br. (dkt. #678) 65-66.)

arguments that it did not control or direct Samsung's actions, as the court indicated in its opinion and order on Apple's Rule 50(a) motion, there was more than sufficient evidence in the form of two related manufacturing contracts for a reasonable jury to conclude that Apple did control or direct Samsung's actions.

Third, Apple seeks a new trial on the basis that the court erred in allowing WARF to try its vicarious liability claim during the damages phase of trial. Apple contends that it was prejudiced by the decision. The argument is silly and warrants little attention. As Apple itself acknowledges, the court has wide discretion under Rule 42(b) to determine how best to try a case. Here, the first phase of the trial covered a lot of ground -- both infringement and invalidity. To have added WARF's vicarious liability claim would have unnecessarily complicated that first phase of trial since this issue was only material to the question of the scope of damages. If the jury had found no infringement or had found the patent invalid, there would have been no need for the evidence and argument on vicarious liability to be presented to the jury.[3]

Nor is the court convinced by Apple's argument that the jury was confused by the vicarious liability questions given statements in the instructions that

---

[3] Admittedly, all of Apple's processors were manufactured in part overseas, then the calculus would have been different. This is because the issue of vicarious liability would have been *central* to plaintiff's infringement claims and, therefore, those two questions likely would have been posed in the first phase of trial, and likely as the first two questions on the verdict form, but those were *not* the facts of this case.

the jury had already found infringement. In particular, the instructions explained that the vicarious liability questions concerned processors manufactured in part in the United States and in part overseas. (Closing Damages Instructions (dkt. #649) 1 ("In determining whether to include these [wafers] in the damages award, you must *first* consider whether the products infringe *at the time they leave the United States*, before any additional manufacturing or processing occurs outside of the United States.").) Accordingly, there appeared no risk of confusion. Indeed, the jury certainly was capable of understanding why it was being asked these questions in the second phase of the trial and how those answers impacted the damages award.

D. Damages

Apple finally raises a variety of challenges to the jury's damages award, which roughly fall into three buckets: (1) the admission of certain evidence and expert testimony; (2) errors in the instructions; and (3) the damages award was not supported by the manifest weight of the evidence. The court will address each challenge in turn.

i. Evidentiary Rulings

Apple challenges the introduction of evidence regarding patents that Apple asserted in litigation were infringed by Samsung and the royalties that Apple sought from Samsung. Specifically, Apple argues that those patents did not involve comparable technology, are reflective of the competitive relationship between Apple and Samsung, and were adopted in an unrelated litigation. The court already considered these arguments in ruling on motions *in*

limine and finds no basis for reconsidering its decision to admit such evidence now. (9/29/15 Op. & Order (dkt. #468) 25-26.) For the most part, the evidence was used by WARF to rebut Apple's damages position. Moreover, Apple was not unduly prejudiced by the evidence as it was free to present testimony and argue -- and it did both -- that: these patents are not comparable; the negotiation between Apple and Samsung involved a competitive dynamic not at issue in this litigation; and the royalties were sought in the context of litigation. All of this simply goes to the weight the jury may assign to the patents and their royalties, not the admissibility.[4] *See Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1326 (Fed. Cir. 2014) ("[W]hether these licenses are sufficiently comparable such that Motorola's calculation is a reasonable royalty goes to the weight of the evidence, not its admissibility.").

Apple also challenges the admissibility of WARF's expert Catharine Lawton's testimony that the parties would have agreed to split the profits 50/50. Here, too, the court touched on this challenge at trial and sees no basis for revisiting its decision to allow her testimony. (9/29/15 Op. & Order (dkt. #468) 36-37.) Lawton testified that she isolated the patents attributable to the patented invention, and based on "conversations with [WARF's long-serving managing director] Dr. Gulbrandsen and WARF's history of licensing and negotiation and the nature of this

---

[4] Because the court rejects each of these challenges, the court need not address WARF's argument that even without this evidence, there was sufficient support from unchallenged evidence to sustain the jury's award. (Pl.'s Opp'n (dkt. #711) 212-15.)

technology, that WARF would have sought at this hypothetical negotiation 50 to 70 percent of the incremental additional profit that Apple realized." (10/14/15 Trial Tr. (dkt. #671) 171-72.) The court found that Lawton sufficiently tied her 50% profit split to the facts of the case, and therefore her testimony is distinguishable from impermissible "rule of thumb" expert opinions. *See Carnegie Mellon Univ. v. Marvell Tech. Grp., Ltd.,* No. 09-290, 2012 WL 3679564, at *6 (W.D. Pa. Aug. 24, 2012) (distinguishing the holding in *Uniloc* from expert testimony "tied to the facts of the case"). Here, too, the jury was free to place little or no weight on this testimony, but Apple's challenge does not go to its admissibility.

Also with respect to Lawton, Apple challenges the introduction of her testimony regarding an estimated market price of the A7 chip. The court addressed this challenge in its motions *in limine* order, and again sees no reason to revisit its decision allowing her testimony and introduction of evidence. (9/29/15 Op. & Order (dkt. #468) 35-36.) Whether Lawton's cost estimate was inflated was proper fodder for cross-examination -- and Apple did explore this at trial -- but Apple's arguments fall short of demonstrating that her analysis was so unreliable that the court erred in not excluding it.

Finally, with respect to the first category of challenges to the damages award, Apple contends that the court erred in allowing Dr. Knittel to present his regression analysis. Apple also challenges the admissibility of Dr. Knittel's testimony as part of its motions *in limine.* Here, too, the court rejected Apple's challenge, finding that it went to the weight

the jury may place on his opinion, not its admissibility. (9/29/15 Op. & Order (dkt. #468) 32-25.) The court sees no basis to reconsider that decision either.[5]

ii. Jury Instructions

Apple takes issue with two aspects of the damages instructions. First, Apple contends that the court erred in including all fifteen of the *Georgia-Pacific* factors in the instructions. Relying on *Ericsson Inc. v. D-Link Systems, Inc.*, 773 F.3d 1201 (Fed. Cir. 2014), Apple argues that the court failed to consider the facts of this case in determining which factors were relevant for the jury's consideration. While certain of the factors were more central to the parties' respective damages cases, WARF's expert did provide a slide on all fifteen factors. (Demonstrative Ex. 93 (dkt. #650-9) 74.) Moreover, Apple fails to explain how it was prejudiced from an instruction that permitted the jury to *consider* certain factors rather than require consideration. (Intro. Damages Instr. (dkt. #649) 2 ("The following is not every possible factor, but it will give you an idea of the kinds of things to consider in setting a reasonable royalty.").) More specifically and unlike *Ericsson*, Apple fails to explain

---

[5] As part of this challenge, Apple also claims that the court limited its cross-examination of Dr. Knittel, thereby undermining its attempts to challenge this methodology. The court has reviewed the portions of the trial transcript Apple cited in support of this argument, which actually reflect an attempt to manage the trial and not undue interference with Apple's cross-examination. Moreover, Apple's contention that the court required Apple to submit a proffer on further crossexamination of Knittel is belied by the record -- as WARF points out in its opposition brief. (Def.'s Opp'n (dkt. #711) 258-60.) Regardless, Apple dropped this challenge in its reply brief.

what factors would have been contrary to the licensing requirements. (Pl.'s Opp'n (dkt. #711) 261-62.)

Apple also challenges the court's refusal to provide a special instruction on non-infringing alternatives and switching costs. As an initial point, the court *did* instruct the jury to consider "the availability of other non-infringing alternatives" in determining a reasonable royalty. The court simply rejected a more detailed instruction offered by Apple, which the court deemed unnecessary and more suitable for argument. In no way did the court restrict Apple's efforts to produce evidence or argument on this factor. As for switching costs, Apple acknowledged in its reply (Def.'s Reply (dkt. #728) 151) that it failed to present any evidence on this subject which rendered the proposed instruction irrelevant.

iii. Manifest Weight of Evidence

Apple further contends that the verdict went against the manifest weight of the evidence. In support of this challenge, Apple repeats the same arguments made to the jury as to the lack of comparability of the 2009 WARF-Intel license and Apple licenses, as well as the importance of WARF's pre-litigation valuation of the patent. The jury could have accepted those arguments, but obviously opted otherwise. As for Apple's challenge to specific evidence, a party "must do more than identify favorable evidence that, if isolated from . . . opposing evidence, would support [its] conclusion." *Plyler v. Whirlpool Corp.*, 751 F.3d 509, 513 (7th Cir. 2014).

As WARF described in its opposition, the record as a whole provides sufficient support for the jury's award. Specifically, the jury reasonably could have relied on the 2009 Intel agreement for $110 million, the significant performance and energy-saving benefits Apple achieved through its use of the patented invention and the economic importance of the invention to the exponentially faster processing speed now necessary for Apple's iPhones. (Pl.'s Opp'n (dkt. #711) 212-15.) Given this, Apple has not established that the jury's award went against the manifest weight of the evidence so as to warrant a new trial.

II. WARF's Motion to Alter or Amend Judgment as to Willful Infringement (dkt. #681)

After the court granted Apple's motion for judgment as a matter of law on WARF's willful infringement claim, the United States Supreme Court articulated a different standard for proving such a claim. *Halo Electronics, Inc. v. Pulse Electronics, Inc.*, 136 S. Ct. 1923 (2016). Shortly after the Court granted *certiorari* in *Halo*, WARF filed the present motion, anticipating that the standard defined in *In re Seagate Technology, LLC*, 497 F.3d 1360 (Fed. Cir. 2007 (*en banc*), would be vacated and that the Supreme Court would adopt a similar totality of the circumstances test described in *Octane Fitness, LLC v. ICON Health & Fitness, Inc.,* 134 S. Ct. 1749 (2014), with respect to claims for attorneys' fees and costs under 35 U.S.C. § 285. WARF's motion proved prescient, but it does not change the result here.

In *Halo*, the Supreme Court rejected the "unduly rigid" objective and subjective prongs set forth in *Seagate*, instead allowing courts to award enhanced damages based on "subjective willfulness of a patent infringer, intentional or knowing, . . . without regard to whether his infringement was objectively reckless." 136 S. Ct. at 1933. Moreover, the Court rejected *Seagate*'s requirement that willful infringement be shown by a heightened clear and convincing evidence standard, instead adopting a preponderance of the evidence standard. *Id*. at 1934. In adopting a more flexible standard, however, the Court still cautioned that the award of enhanced damages should be limited to "egregious cases of misconduct beyond typical infringement." *Id*. at 1935.

In its original opinion and order granting judgment to Apple on WARF's willful infringement claim, this court applied the then-controlling two-part *Seagate* test, concluding that certain of Apple's invalidity defenses were not objectively unreasonable. Certainly, *Halo* calls into questions whether Apple, and in turn this court, may rely on an objective showing of plausible defenses absent a further showing that Apple actually believed that the '752 patent was invalid at the time it commenced infringing the patent. Indeed, the Supreme Court was quite critical of this aspect of the *Seagate* test:

> The existence of such a defense insulates the infringer from enhanced damages, even if he did not act on the basis of the defense or was even aware of it. Under that standard, someone who plunders a patent—infringing it without any reason to suppose his conduct is arguably defensible—can nevertheless

> escape any comeuppance under § 284 solely on the strength of his attorney's ingenuity.

Halo, 136 S. Ct. at 1933.

While *Halo* certainly grants more discretion in determining whether enhanced damages are appropriate under § 284, a threshold element remains proof that the defendant necessarily knew of the patent. *Id.* Prior to the court issuing its decision on WARF's willful infringement claim, the parties submitted briefs on their respective positions on this claim. (Dkt. ##587, 606.) Each side devoted a few pages to the "subjective prong," which primarily concerned Apple's knowledge of the patent. The court had no reason to previously consider these arguments, given that it rested its decision on the objective prong but will do so now.

There appears to be no dispute that Apple was aware of the '752 patent before the filing of this lawsuit. Apple engineer Stephen Meier testified at his deposition that he was given the '752 patent by outside patent prosecution counsel for Apple in November 2013. The timing of Meier's knowledge of the patent is interesting: it occurred two months *after* Apple began selling iPhones containing the accused A7 chip, and two months *before* WARF filed the present lawsuit. However, WARF argues that Apple was aware of the patent as early as 2010 based on: (1) it being briefly mentioned in an article that at least some Apple engineers reviewed; and (2) it being disclosed in one of Apple's own patent applications. As for the first basis, the patent was briefly mentioned in an academic article, without any significant description. As for the second piece of evidence,

WARF merely offers a single reference to the '752 patent among more than twenty other patent references cited in a patent issued for a memory-hazard detection and avoidance instructions for vector processing. (Def.'s Ex. 1176 (U.S. Patent No. 8,019,976B2).)

Without some linkage between the inventors of this patent or others working on that technology and the inventors of Apple's LSD Predictor, much less actual proof of Apple's copying of the '752 patented technology, this is not enough to impute knowledge to Apple. *See Potter Voice Techs., LLC v. Apple Inc.*, 24 F. Supp. 3d 882, 886 (N.D. Cal. 2014) ("In the context of willful infringement, it is safe to say that the employees required to have knowledge of the asserted patent must have some connection to the decision willfully to infringe."). Even viewed together, this limited evidence does not provide a sufficient basis for the court to find knowledge of the '752 patent pre-dating November 2013.

The timing of Apple's knowledge is material because a finding of willfulness "will depend on an infringer's prelitigation conduct." *In re Seagate Tech., LLC*, 497 F.3d at 1374, *abrogated on other grounds by Halo Elecs., Inc. v. Pulse Elecs., Inc.*, 136 S. Ct. 1923 (2016).[6] As the Federal Circuit explained:

---

[6] While the Supreme Court rejected the standard for determining willful infringement under § 284 in *Halo Elecs.*, the opinion did not upset the Federal Circuit's holding in *Seagate* that the focus of such a claim should be on prelitigation conduct. *See Dorman Prod., Inc. v. Paccar, Inc.*, No. CV 13-6383, 2016 WL 4440322, at *9 (E.D. Pa. Aug. 23, 2016), as amended (Oct. 17, 2016).

> It is certainly true that patent infringement is an ongoing offense that can continue after litigation has commenced. However, when a complaint is filed, a patentee must have a good faith basis for alleging willful infringement. Fed. R. Civ. Pro. 8, 11(b). So a willfulness claim asserted in the original complaint must necessarily be grounded exclusively in the accused infringer's pre-filing conduct. By contrast, when an accused infringer's post-filing conduct is reckless, a patentee can move for a preliminary injunction, which generally provides an adequate remedy for combating post-filing willful infringement. *See* 35 U.S.C. § 283; *Amazon.com, Inc. v. Barnesandnoble.com, Inc.*, 239 F.3d 1343, 1350 (Fed. Cir. 2001). A patentee who does not attempt to stop an accused infringer's activities in this manner should not be allowed to accrue enhanced damages based solely on the infringer's post-filing conduct. Similarly, if a patentee attempts to secure injunctive relief but fails, it is likely the infringement did not rise to the level of recklessness.

Id.; *see also Dorman Prod., Inc. v. Paccar, Inc.*, No. CV 13-6383, 2016 WL 4440322, at *9 (E.D. Pa. Aug. 23, 2016), *as amended* (Oct. 17, 2016) (rejecting willful infringement claim based primarily on post-litigation conduct, explaining Dorman's conduct during the brief pre-filing period is insufficient to allow PACCAR's claims for willful infringement 'in

the main' to be based on pre-filing conduct" (citing *Seagate*, 497 F.3d at 1374)).

Similarly, WARF's willful infringement claim is based solely on the two months period before the filing of this lawsuit. Moreover, all the evidence shows that at the time Apple learned of the '752 patent, it had already designed, manufactured and begun to sell phones containing the infringing processor. In other words, there is no evidence of copying or other egregious misconduct that would warrant a finding of willful infringement. Indeed, once the case was filed -- assuming the court can consider post-litigation conduct -- Apple developed and pursued an invalidity defense, which the court found to be objectively reasonable, albeit ultimately unsuccessful. Viewing the record as a whole, under *Halo*, therefore, the court again concludes that WARF has failed to demonstrate willful infringement by a preponderance of the evidence. Accordingly, WARF's motion to alter or amend the court's order granting judgment to Apple on that claim is denied.

III. WARF's Motion for Equitable Relief (dkt. #683)

A. Apple's Motion to Strike

In support of its motion for a equitable relief, WARF filed two reply declarations of its damages experts, Lawton and Knittel, providing *Georgia-Pacific* analysis on WARF's requested ongoing royalty rate. WARF's managing director also provided a declaration, describing WARF's desire to maintain exclusivity over its patents and submitting documents in support of that contention. Apple moved

to strike these declarations, and the portions of WARF's reply brief relying on those declarations, on the basis that they should have been submitted with the initial motion pursuant to Federal Rule of Civil Procedure 62(c)(2). (Dkt. #744.) In response, WARF argues that: (1) Apple's concession in its opposition brief that the A9 and A9x chips infringe the '752 patent justified the reply declarations; and (2) the experts were simply responding to Apple's own expert's *Georgia-Pacific* analysis.

Without going through each argument, the court generally agrees with Apple that as the party with the burden of proof, WARF should have provided Lawton and Knittel's analysis as part of its opening submission, not in reply. In particular, the court fails to see Apple's concession of infringement as any justification for the submission of otherwise untimely expert opinions in support of WARF's motion. As such, the court will grant Apple's motion to strike Lawton's and Knittel's reply declarations.

As for Gulbrandsen's declaration and Exhibits 6, 7 and 8 to Proctor's declaration (dkt. ##734-6, 734-7, 734-8), the court again agrees with Apple that WARF's focus on exclusive licenses -- as distinct from its interest in excluding infringers from practicing its patents -- is a new argument raised for the first time in reply. Accordingly, the court will also grant the motion to strike this declaration and supporting exhibits.

B. Permanent Injunction

To be entitled to a permanent injunction, a patentee must show: (1) it has suffered an irreparable injury; (2) remedies available at law are inadequate

to compensate for that injury; (3) considering the balance of hardships between the plaintiff and defendant, a remedy in equity is warranted; and (4) the public interest would not be disserved by a permanent injunction. *See Douglas Dynamics, LLC v. Buyers Prod. Co.*, 717 F.3d 1336, 1344 (Fed. Cir. 2013) (citing *eBay Inc. v. MercExchange, L.L.C.*, 547 U.S. 388, 391 (2006)). Despite these requirements, the Supreme Court has rejected a categorical rule barring non-practicing entities from seeking a permanent injunction: "University researchers or self-made inventors might reasonably prefer to license their patents, rather than undertake efforts to secure the financing necessary to bring their works to market themselves. Such patent holders may be able to satisfy the traditional four factor test, and we see no basis for categorically denying them the opportunity to do so." *eBay Inc.*, 547 U.S. at 393. Even so, the awards of injunctive relief in cases since *eBay* appear limited to cases where "a party that does not practice the asserted patent" still "sells a competing product." *Trebro Manufacturing, Inc. v. Firefly Equipment, LLC*, 748 F.3d 1159, 1171 (Fed. Cir. 2014) (citing cases).

Here, WARF does not manufacture or sell any competing product. WARF nevertheless offers various theories to support a finding of irreparable injury: (1) "[t]he meaningful threat of an injunction is necessary to support a viable voluntary licensing program for WARF for this industry"; (2) "having to initiate serious litigation against Apple to address its ongoing infringement causes significant and not quantifiable reputational harm to WARF"; and (3) litigation costs off-set WARF's financial

contributions to the University of Wisconsin. All of these theories rest on an assumption that a threat of a permanent injunction would motivate patent infringers to negotiate a license upfront thus limiting the need for future litigation to enforce its patent rights. While this theory may have some merit, it only works if the patent infringement is knowing -- absent knowledge of the patent the alleged patent infringer would have no basis for seeking a license. Otherwise, WARF is effectively proposing a new presumption in favor of the entry of a permanent injunction that would be applicable in all cases. Moreover, the award of an ongoing royalty and the threat of treble damages for willful infringement still provide significant motivation for potential patent infringers to negotiate licenses upfront.

There is also a flip-side to WARF's theory, as Justice Kennedy explained in his concurring opinion in *eBay*:

> An industry has developed in which firms use patents not as a basis for producing and selling goods but, instead, primarily for obtaining licensing fees. For these firms, an injunction, and the potentially serious sanctions arising from its violation, can be employed as a bargaining tool to charge exorbitant fees to companies that seek to buy licenses to practice the patent. When the patented invention is but a small component of the product the companies seek to produce and the threat of an injunction is employed simply for undue leverage in negotiations, legal damages may well be sufficient to

> compensate for the infringement and an injunction may not serve the public interest.

eBay Inc., 547 U.S. at 396–97 (Kennedy, J., concurring). As Apple points out, courts routinely refuse to award injunctive relief where the patentee's "motivation in seeking an injunction is less about preventing irreparable harm and more about extracting . . . leverage in negotiating with [the defendant]." *Hynix Semiconductor Inc. v. Rambus Inc.*, 609 F. Supp. 2d 951, 983 n.29 (N.D. Cal. 2009). (See also Def.'s Opp'n (dkt. #709) 26-28 (discussing other cases where a court's finding of irreparable harm due to damage to reputation was at least based on the patentee competing in some form with the patent infringer).)

The court is also unconvinced that WARF's pursuit of patent litigation meaningfully harms its reputation as an inventor. In seeking an injunction, WARF demonstrates that it is serious about enforcing its patent rights; the court's decision denying an injunction does not change that fact. Similarly, WARF filed this lawsuit and another lawsuit against Apple, not to mention the earlier lawsuit on the same patent against Intel -- all of which shows that WARF will pursue litigation to enforce its intellectual rights.

Of course, as an affiliate of an institution of higher learning that generally promotes the open exchange of scientific and other knowledge, there may be a downside to WARF's litigation strategy, including its expressed concern about being lumped in with so-called patent trolls as supported by newspaper articles identifying WARF as just that. (*See* Pl.'s Opening Br. (dkt. #683) 16.) However, any

reputational damage is caused by the filing of litigation itself, and would, if anything, presumably be worsened by the entry of an injunction in this case. However as Apple points out, the press's characterization of WARF as a patent troll pre-dates this litigation (Def.'s Opp'n (dkt. #709 23-24).

Given the court's skepticism that the threat of a permanent injunction would limit (as opposed to foster) litigation, the court remains unconvinced that any reputational injury caused by pursuing litigation serves as a credible basis for finding irreparable harm, not to mention the obvious distinction between a research institution tied to a public university protecting the patented work of its professors and entities from an entity that simply buys up patent rights for purposes of extracting licenses through the threat of litigation. Regardless, adding the threat of permanent injunctive relief going forward would only enhance such a reputation.

At the end of the day, the court concludes that WARF has not met its burden of demonstrating irreparable harm, and even if it had, the balance of equities and the public interest both weigh in favor of denying entry of a permanent injunction. If the court were to enter an injunction, Apple would have to disable the LDS Predictor, which would likely prevent Apple from using non-infringing features, including features covered by Apple's own patents. Most importantly, until an alternative, non-infringing alternative can be incorporated into the iPhone, removing the LSD Predictor may well deprive the public of all of the technology contained in that product, not just the infringing technology.

C. Ongoing Royalty

Likely recognizing that its request for a permanent injunction was a long-shot, WARF also seeks, as an alternative, an award of an ongoing royalty based on per unit sales. In its response, Apple also concedes that such an award is warranted. (Def.'s Opp'n (dkt. #709) 39 ("Apple does not dispute that WARF is entitled to an ongoing royalty for any infringement occurring after the entry of final judgment (and supplemental damages before that time).").) So the only question is what that royalty should be.

WARF seeks an ongoing royalty of three times the implied jury's per unit rate but provided little justification for this figure in its opening brief. WARF does offer additional support in reply -- some of which the court struck above -- mainly focused on the changed circumstances post-verdict, which obviously alters the parties' relative bargaining positions. WARF also argues that Apple's now willful infringement supports the requested per unit royalty rate.

In contrast, Apple urges the court to delay ruling on any ongoing royalty until after resolution of any appeal of the jury's findings of infringement and rejection of Apple's invalidity challenges. Alternatively, Apple proposes that the court provide an opportunity for the parties to negotiate an ongoing royalty rate. Barring either of those proposals, Apple argues that the court should simply adopt the same implied per unit rate awarded by the jury as an ongoing royalty.

As an initial matter, the court can see little efficiency or justice in forgoing a decision as to the award of an ongoing royalty. Absent a ruling, the Federal Circuit would obviously be prevented from taking up the entire case in one appeal. Such a piecemeal review on appeal seldom makes sense, especially now that this remaining issue is fully briefed. *See Nystrom v. TREX Co.*, 339 F.3d 1347, 1350 (Fed. Cir. 2003) (describing policy behind 28 U.S.C. § 1295). As for Apple's request to allow the parties to negotiate the ongoing royalty rate first, the court recognizes that the Federal Circuit has encouraged this approach, but again sees little purpose in further postponing the inevitable. *See Paice LLC v. Toyota Motor Corp.,* 504 F.3d 1293, 1313 (Fed. Cir. 2007). As WARF describes in its reply brief, the parties have already endeavored to negotiate an award and obviously failed to reach a resolution. (Pl.'s Reply (dkt. #736) 35.) Regardless, given the parties' behavior and failure to reach a settlement to date, the court finds the likelihood of a negotiated, ongoing royalty unlikely. *See Apple, Inc. v. Samsung Elecs. Co.*, No. 12-cv-00630-LHK, 2014 WL 668122, at *13 (N.D. Cal. Nov. 25, 2014) ("[T]he parties' behavior indicates that any order to negotiate ongoing royalties is likely to be futile and only delay the entry of final judgment.").

With those preliminaries aside, the court takes up the appropriate amount of an ongoing royalty. Apple's proposal that the court simply award the rate awarded by the jury for the past infringing sales is a non-starter. The Federal Circuit "easily dispose[d]" of that very argument in *Amado v. Microsoft Corp.*, 517 F.3d 1353 (Fed. Cir. 2008):

> On the other side of the dispute, Microsoft argues that the district court was entitled to award Amado no more than $0.04 per infringing unit, the amount the jury found to be a reasonable royalty. We easily dispose of this argument as well. The jury's award of $0.04 per unit was based on Microsoft's infringing conduct that took place *prior to the verdict*. There is a fundamental difference, however, between a reasonable royalty for pre-verdict infringement and damages for post-verdict infringement. *Cf. Paice LLC v. Toyota Motor Corp*., 504 F.3d 1293, 1317 (Fed. Cir. 2007) ("[P]re-suit and post-judgment acts of infringement are distinct, and may warrant different royalty rates given the change in the parties' legal relationship and other factors.") (Rader, J., concurring). Prior to judgment, liability for infringement, as well as the validity of the patent, is uncertain, and damages are determined in the context of that uncertainty. Once a judgment of validity and infringement has been entered, however, the calculus is markedly different because different economic factors are involved.

Id. at 1361-62 (emphasis added). Since then, the Federal Circuit has reiterated its holding that in the ongoing royalty context, courts should "take into account the change in the parties' bargaining positions, and the resulting change in economic circumstances, resulting from the determination of liability." *ActiveVideo Networks, Inc. v. Verizon*

Communications, Inc., 694 F.3d 1312, 1343 (Fed. Cir. 2012) (quoting *Amado*, 517 F.3d at 1362).

Predictably, WARF also stakes out an extreme position, primarily relying on a willful infringement framework in seeking an ongoing royalty rate -- thus, explaining its request for a tripling of the jury's award. The court finds this is also an ill-fit. While equitable considerations certainly come into play, the court rejects WARF's attempt to describe Apple's infringement post-jury verdict as willful, which would justify a pragmatic trebling of damages under § 284. Rather, Apple reasonably believed that the court would not enter a permanent injunction, after entry of judgment on October 26, 2015, and instead would award an ongoing royalty based on a hypothetical negotiation -- an expectation shared by this court. Given this context, Apple knew it was going to have to pay for its continued use of the infringing technology, unlike a willfully infringing party who hopes to conceal its knowing infringement. This then leaves the court with the task of arriving at the parties' bargaining positions and the outcome of the hypothetical negotiation here.

At the outset, the court credits the jury's consideration of the *Georgia-Pacific* factors in setting an implied per unit rate for past infringement. From that amount, the court is instructed to consider changes in the parties' bargaining positions. *See Amado v. Microsoft Corp.,* 517 F.3d at 1361-62. Certainly, the jury's finding of infringement bolsters WARF's bargaining position in the hypothetical negotiation which is to occur on the date of the jury's verdict. In light of that change, WARF is in a better position to demand a greater percentage of Apple's

profits that are attributable to the LSD predictor than awarded by the jury. Relying on the expert testimony of Julie L. Davis, WARF sought a royalty rate of $2.74 per unit, which the jury discounted, presumably because the jury found WARF's bargaining position during the hypothetical negotiation was not as strong as it maintained. In light of WARF's improved bargaining position after the jury's finding of infringement and validity, the court finds that the $2.74 rate is fair and reasonable. Accordingly, the court will award that amount as an ongoing royalty for all sales of iPhones containing the LSD predictor from October 26, 2015, to the end date of the patent.

IV. WARF's Motion for Accounting, Supplemental Damages through the Date of Judgment, Prejudgment Interest and Post-judgment Interest (dkt. #685)

A. Accounting

WARF seeks an accounting to determine the unit sales for purposes of establishing supplemental and ongoing royalty payments. Apple contends that the motion is unnecessary because it will voluntarily produce the financial data as available. Regardless, Apple is obligated to produce financial data showing the number of sales for the court to calculate a supplemental and ongoing damages award, and that the financial data should include sales of phones containing the A9 and A9x chips. Should WARF believe in good faith that Apple has been dilatory or inaccurate in disclosing this data, it may certainly pursue post-judgment and supplement discovery. Absent proof of either, however, the court is disinclined to order a formal accounting. The parties

are instructed to act in good faith with respect to any discovery *and* efforts to arrive at the appropriate figures, including a "meet and confer" before bringing any disputes before this court.

B. Supplemental Damages

WARF also seeks an award of supplemental damages, based on the jury's per unit award, from June 27, 2015 -- the end date of the parties' stipulation on accused units sold -- through the date of judgment, October 26, 2015. Here, too, Apple does not oppose the request, but contends that any supplemental damages rate should also cover sales through the date of the resolution of all post-trial motions. As referenced above, the court agrees with WARF that the date of judgment, October 26, 2015, is the appropriate date from which to calculate an ongoing royalty. *See Carnegie Mellon Univ. v. Marvell Tech. Grp., Ltd.*, 807 F.3d 1283 (Fed. Cir. 2015). While the court will enter an amended judgment, the amended judgment does not alter the jury's finding of liability, nor the changed circumstances from that decision, which in turn alters the ongoing royalty rate. *See Kaiser Aluminum & Chem. Corp. v. Bonjorno*, 494 U.S. 827, 835 (1990) ("By linking all post-judgment activity to the entry of judgment, the courts have been provided a uniform time from which to determine post-judgment issues." (internal quotation marks omitted)). As such, the court will award supplemental damages at the per unit royalty rate awarded by the jury from June 27 to October 25, 2015.

Preserving its right to appeal the jury's findings on liability and damages, Apple also seeks to include

its products containing the A9 and A9x chips in the supplemental damages award, since it concedes that the jury's finding of infringement covers those chips as well.[7] In response, WARF complains about Apple's last minute switch in position, especially in light of its earlier representation that these later chips may undergo design changes as well as its unwillingness to engage in discovery of the A9 and A9x chips. While the court is sympathetic to WARF's positon, it fails to provide a credible reason why the jury's royalty rate for pre-judgment infringement as well as the court's awarded ongoing royalty rate for post-judgment infringement, should not apply to the A9 and A9x chips.

While WARF speculates that the jury *may* have awarded higher damages if it had known about the A9 and A9x chips, the court is hard-pressed to understand how the continued use of the LSD predictor would appreciably have changed either parties' bargaining position at the time of the hypothetical negotiation in 2013, particularly since the negotiated royalty rate was for ongoing use. All of this is to say, that the court is inclined to include the sale of A9 and A9x (and possibly A10) chips, both in calculating a supplemental damages award and in setting an ongoing royalty rate. Both sides as directed to brief their positions for consolidating the 15-cv-621 case with this action and awarding damages for

---

[7] In the light of the above discussion, the supplemental damages award for those chips would also cover all sales up to the date of judgment, October 26, 2015, with all other sales presumably falling under the ongoing royalty rate.

infringement of the A9, A9x and A10 chips as part of this case.

C. Pre-judgment Interest

Next, WARF seeks an award of pre-judgment interest of 5.65% compounded quarterly, from the date of infringement through the date of judgment. Title 35 U.S.C. § 284 governs the award of prejudgment interest in patent infringement claims. "In the typical case an award of prejudgment interest is necessary to ensure that the patent owner is placed in as good a position as he would have been in had the infringer entered into a reasonable royalty agreement." *Gen. Motors Corp. v. Derex Corp.*, 461 U.S. 648, 655 (1983). For this reason, "prejudgment interest should be awarded under § 284 absent some justification for withholding such an award." *Id.* at 657; *see also Energy Transp. Grp., Inc. v. William Demant Holding A/S*, 697 F.3d 1342, 1358 (Fed. Cir. 2012) ("The award of pre-judgment interest is the rule, not the exception.") (quotation and citation omitted). Consistent with this case law, Apple concedes that a pre-judgment interest award is warranted, but argues that the court should award interest at the near record T-bill rate, currently 0.31%, and that the interest should be compounded annually, not quarterly.

As for the appropriate rate, the court rejects both parties' positions. Instead, it will follow the practice approved by the Federal Circuit and Seventh Circuit, which is also consistent with its own practice, by awarding prejudgment interest based on the prime rate. *See Uniroyal, Inc. v. Rudkin-Wiley Corp.*, 939 F.2d 1540, 1545 (Fed. Cir. 1991) (explaining that a

district court "is afforded wide latitude in the selection of interest rates" and "may award interest at or above the prime rate"); *First Nat. Bank of Chi. v. Standard Bank & Trust*, 172 F.3d 472, 480 (7th Cir. 1999) (ordinarily, to "award something other than the prime rate is an abuse of discretion"); *Partington v. Broyhill Furniture Indus., Inc.*, 999 F.2d 269, 274 (7th Cir. 1993) (in federal cases, "district judges should use the prime rate" for prejudgment interest); *see also Chesemore v. Alliance Holdings*, No. 09-cv-413, 2014 WL 4415919, at *8 (W.D. Wis. Sept. 5, 2014); *Nat'l Pasteurized Eggs, Inc. v. Michael Foods, Inc.,* No. 10-cv-646-wmc, slip op. at *36 (W.D. Wis. Mar. 29, 2013) (dkt. #550). As for compounding, the court credits WARF's evidence and argument that payment of running royalties on a quarterly basis is consistent with both Apple's and WARF's respective practices. (Pl.'s Br. (dkt. #685) 15-16.) Accordingly, the court will award prejudgment interest at the prime rate, compounded quarterly. Because pre-judgment interest should also apply to the award of supplemental damages, the court will await entering a pre-judgment interest award until the supplemental damages have been calculated.

D. Post-judgment Interest

Finally, WARF seeks an award of post-judgment interest in the amount of 0.23% compounded annually, as provided under 28 U.S.C. § 1961(a). Apple does not oppose this request. As such, that motion will be granted as unopposed.

V. WARF's Motion for Taxation of Costs (dkt. #689); Amended Bill of Costs (dkt. #725)[8]

This brings us to the last pending motion. WARF seeks an award of costs allowed to the prevailing part under Federal Rule of Civil Procedure 54(d)(a). Specifically, WARF seeks reimbursement of the following costs:

(1) fees of the clerk and *pro hac vice fees*, totaling $750.00;

(2) fees for the service of summons and subpoena, totaling $339.72;

(3) fees for printed and electronically recorded transcripts, totaling $154,814.68;

(4) witness travel fees, totaling $26,943.96;

(5) fees for exemplification and the costs of printing and photocopying, totaling $878,709.35;

(6) rental fees of photocopiers, totaling $5,475.24; and

(7) Apple infringing device purchases of $5,497.74.

(Pl.'s Mot. (dkt. #690).)

In response, Apple takes issue with certain categories of costs, arguing that the fees requested are excessive or are inappropriate and should be eliminated. *First*, with respect to WARF's request for $878,709.35 in exemplification and copying, Apple contends that the fees should be reduced by over

---

[8] The Amended Bill of Costs moots the original submission (dkt. #688).

$500,000. Specifically, Apple challenges WARF's inclusion of office supplies and freight costs, as non-taxable costs under Seventh Circuit law. (Def.'s Opp'n (dkt. #705) 8.) While WARF concedes in its reply that the $4,500 cost for transporting documents (e.g., the freight cost) should be deducted, it contends that its request for custom tabs and binders prepared by third-party copy vendors are taxable, and distinguishable from general office supplies not allowed by the Seventh Circuit. (Pl.'s Reply (dkt. #726) 4-5 (citing cases).) The court agrees with WARF that the $18,023.13 actually incurred in "office supplies" charges used to compile documents by a third-party vendor for use in this case are taxable. Accordingly, the court rejects Apple's objection as to this category. Accordingly, the amended bill of costs will reflect only WARF's deduction of the $4,500 cost for transporting documents.

In that same category, Apple objects to photocopying costs associated with five of WARF's depositions, and specifically, objects to the number of pages of exhibits that were printed by WARF, pointing out that the number of pages of exhibits actually used during the depositions was substantially less than that printed, and WARF opted for tens of thousands of expensive color copies, rather than black and white. As for this request, Apple seeks a reduction in taxable costs of $28,993.53. The court agrees with WARF that requiring a justification on a document-by-document basis is "preposterous." (Pl.'s Reply (dkt. #726) 5 (citing *NOW, Inc. v. Scheidler*, 750 F.3d 696, 698 (7th Cir. 2014)).) While perhaps WARF could have been more measured in its printing of exhibits, the court will not fault it for thorough

preparation to the extent the costs are amply documented. Moreover, color copies, especially in the context of a deposition concerning technical issues, appear reasonable. Accordingly, the court rejects this challenge to WARF's cost request.[9]

Next, Apple challenges the invoices submitted by local counsel, Godfrey & Kahn, and by WARF's graphics vendor, arguing that they do not provide sufficient detail to know whether the costs are appropriately taxable, and requests a reduction of $3,660.45 for the former and $202,478.31 for the latter. The court credits Attorney Gregor's affidavit, and sees no basis for requiring additional detail as to the copying costs incurred by Godfrey & Kahn. As for the graphics vendor, here, too, the invoice is sufficiently detailed to award costs, with the exception of certain categories identified by WARF in its reply, totaling $3,274.64. WARF's amended bill of costs also reflects this reduction. Both objections are, however, otherwise overruled.

Under the same exemplification and copying category, WARF seeks $115,475.55 in data storage costs for its e-discovery database. Since the parties' briefing on WARF's bill of costs, this court adopted the majority opinion, which "interpret[es] narrowly the meaning of 'making copies' in § 1920(4) in the context of electronic discovery." *Split Pivot, Inc. v. Trek Bicycle Corp.*, 154 F. Supp. 3d 769, 780 (W.D. Wis. 2015). Under that approach, the court will award costs only for the copying of electronic data, including

---

[9] Even so, the court notes that WARF's amended bill of costs appropriately reduced the amount for photocopying after finding some duplication. (Pl.'s Reply (dkt. #726) 6.)

copying metadata and hard drives. *Id*. As such, the court rejects WARF's request for costs of storing its e-discovery database. Apple's objection, therefore, is sustained, and the court will deduct $115,475.55 from WARF's request for fees for exemplification and copying.

Relatedly, Apple also challenges WARF's request for $147,757.00 for ESI and electronic discovery work performed by Irell & Manella's litigation and database support department personnel. The timesheets describe tasks ranging from converting documents to PDF format, OCR'ing and uploading data to the database, or preparing documents for production, and creating "review bins" based on attorney instructions, among other tasks. (Def.'s Opp'n (dkt. #705) 16.) The court agrees with Apple that costs for creating review bins and database management are *not* taxable for the same reasons that the court does not tax data storage costs. The court further agrees with Apple that a reduction by 50% is appropriate in light of the difficulty, if not impossibility, of determining which of the costs are taxable and which are not. Accordingly, this objection also is sustained, and the court will reduce the fees requests by an additional $73,878.50.

Finally, with respect to exemplification costs, Apple also challenges WARF's request for reimbursement of $30,616 in data purchases made by two of WARF's experts. For reasons set forth above, the court again agrees with Apple that these costs do not constitute copying under § 1920(4). Accordingly, the court sustains this objection, and will deduct an additional $30,616.00. Based on all of the decisions above, this means the court will award fees for

exemplification and copying in the total amount of $658,739.30.

Second, with respect to WARF's request for $154,814.68 in fees for printed or electronically recorded transcripts, Apple challenges various transcript fees as unreasonable or not necessary. WARF seeks $4,220.65 for deposition videos of its ten experts, which Apple inexplicably contends was unnecessary because WARF intended to call its experts live at trial. While these videos were not used at trial, the court finds that the expense was reasonable for trial preparation purposes. Therefore, that objection is overruled.

Apple also challenges WARF's expenses for rough transcripts, expedited transcripts, Realtime and other miscellaneous deposition expenses, and seeks a reduction of $51,580.35. Here the court agrees with WARF that the complexity of patent litigation justifies these expenses. Indeed, both sides appear to have benefitted from the use of all of these technologies. Accordingly, this objection is overruled as well.

Third, Apple objects to certain witness travel fees as being unreasonably high, and seeks a reduction of $6,034.78. Specifically, Apple contends that WARF scheduled depositions in Los Angeles for the convenience of counsel, thereby requiring WARF's witnesses to travel for depositions. Apple maintains that it was willing to travel to the witness's city for those depositions. The court agrees with WARF that in light of the complexity of this case, this challenge is silly and the objection is overruled.

Fourth, and finally, Apple objects to WARF's request for $5,497.74 for reimbursement of device purchases, listed as an "other" cost category. WARF does not maintain that these devices were used as demonstrative evidence or introduced into evidence, unlike the costs for devices allowed in the *Apple, Inc. v. Samsung Electronics* case. Moreover, WARF fails to cite to controlling case law allowing this category of costs as taxable. As such, the court will sustain the objection and deduct that amount from the bill of costs.

In sum, the court will award costs in the total amount of $841,587.66.

ORDER

IT IS ORDERED that:

1) Defendant Apple Inc.'s renewed motion for judgment as a matter of law and/or new trial (dkt. #677) is DENIED.

2) Plaintiff Wisconsin Alumni Research Foundation's motion to alter or amend judgment as to willful infringement (dkt. #681) is DENIED.

3) Plaintiff's motion for equitable relief (dkt. #683) is GRANTED IN PART AND DENIED IN PART. Plaintiff's motion for a permanent injunction is denied, but its motion for an award of an ongoing royalty is granted. The going royalty rate is set at $2.74 per unit.

4) Defendant's motion to strike declarations and portions of reply brief (dkt. #744) is GRANTED.

5) Plaintiff's motion for accounting, supplemental damages through the date of judgment,

prejudgment interest and post-judgment interest (dkt. #685) is GRANTED IN PART AND RESERVED IN PART.

a. Plaintiff's unopposed motion for an accounting is denied without prejudice to plaintiff renewing the motion if discovery efforts fail.

b. Plaintiff's motion for supplemental damages through the date of judgment is granted. Once the accounting is complete, the court will award plaintiff supplemental damages on infringing units sold from June 27, 2015 through the date of judgment, October 26, 2015. The court reserves on whether the supplemental damages award should cover Apple products containing the A9, A9x and A10 chips.

c. Plaintiff's motion for prejudgment interest is granted. The court awards plaintiff prejudgment interest at the prime rate, compounded quarterly. The court reserves on the amount of prejudgment interest, awaiting calculation of supplemental damages award.

d. Plaintiff's motion for post-judgment interest is granted. The court awards plaintiff post-judgment interest at the statutory rate, compounded annually, pursuant to 28 U.S.C. § 1961(a).

6) Plaintiff's motion for taxation of costs (dkt. #689) and amended bill of costs (dkt. #725) is GRANTED IN PART AND DENIED IN PART. The court awards plaintiff costs in the total amount of $841,587.66.

7) On or before June 20, 2017, the parties should submit a joint statement, if possible, updating the court on the status of discovery requests material to the supplemental damages award.

8) On or before June 20, 2017, plaintiff should submit a brief responding to the court's proposal to award supplemental damages at the jury-awarded per unit royalty rate for all sales of infringing Apple products containing A9 and A9x chips prior to the entry of judgment and awarding an ongoing royalty for all sales of Apple products containing the A9, A9x and A10 chips post-judgment. Defendant may have until July 7, 2017, to respond.

Entered this 6th day of June, 2017.

BY THE COURT:

/s/

WILLIAM M. CONLEY
District Judge

APPENDIX L – Judgment (July 27, 2017)

IN THE UNITED STATES DISTRICT COURT FOR THE WESTERN DISTRICT OF WISCONSIN

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff,

v.

APPLE, INC.,

Defendant.

SECOND AMENDED JUDGMENT IN A CIVIL CASE

Case No. 14-CV-062-WMC

This action came for consideration before the court and a jury with District Judge William M. Conley presiding. Partial judgment was granted by the court. Other issues were tried to a jury, which rendered its verdict, and the original judgment was entered based on that verdict. The court has now resolved all post-judgment issues and enters this amended judgment.

IT IS ORDERED AND ADJUDGED that judgment is entered in favor of plaintiff Wisconsin Alumni Research Foundation against defendant Apple, Inc., in the amount of $506,084,992.66 on plaintiff's claim of literal infringement of claims 1,

2, 3, 5, 6 and 9 of U.S. Patent No. 5,781,752 (the "'752 patent"), consisting of the following:

(1) damages awarded by the jury;

(2) supplemental damages at the rate of $1.61 per infringing unit for accused processors sold through the initial entry of judgment on October 26, 2015;

(3) ongoing royalties at the rate of $2. 74 per infringing unit for accused processors from October 27, 2015, through expiration of the '752 patent on December 26, 2016;

(4) prejudgment interest calculated at the prime rate compounded quarterly through the initial October 26, 2015, judgment;

(5) costs as taxed on June 6, 2017; and

(6) post-judgment interest at the statutory rate of 0.232% compounded annually through June 30, 2017.

IT IS FURTHER ORDERED AND ADJUDGED that judgment is entered in favor of plaintiff Wisconsin Alumni Research Foundation against defendant Apple, Inc., on defendant's

claims of invalidity of claims 1, 2, 3, 5, 6 and 9 of the '752 patent.

IT IS FURTHER ORDERED AND ADJUDGED that judgment is entered in favor of defendant Apple, Inc., against plaintiff Wisconsin Alumni Research Foundation on plaintiff's willful infringement claim.

IT IS FURTHER ORDERED AND ADJUDGED, in addition to the amount set forth in the first paragraph above, that judgment shall bear post-judgment interest at the statutory rate of 0.232% compounded annually from July 1, 2017, until finally satisfied in full.

Approved as to form this 26th day of July, 2017.

/s/

_____

William M. Conley, District Judge

/s/ 7/27/17

_____ _____

Peter Oppeneer, Clerk of Court Date

APPENDIX M – Excerpt from Corrected Non-Confidential Brief of Appellee Wisconsin Alumni Research Foundation (January 16, 2018)

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

WISCONSIN ALUMNI RESEARCH FOUNDATION,

Plaintiff-Appellee,

-v-

APPLE, INC.,

Defendant-Appellant.

17-2265, -2380

Appeals from the United States District Court for the Western District of Wisconsin in case no. 3:14-cv-00062-wmc, Judge William M. Conley

---

CORRECTED NON-CONFIDENTIAL BRIEF OF APPELLEE WISCONSIN ALUMNI RESEARCH FOUNDATION

MORGAN CHU
GARY N. FRISCHLING
JOSEPH M. LIPNER
JASON SHEASBY
ALAN J. HEINRICH
CHRISTOPHER T. ABERNETHY
AMY E. PROCTOR

IRELL & MANELLA LLP
1800 Avenue of the Stars, Suite 900
Los Angeles, California 90067
(310) 277-1010

Attorneys for Appellee
Wisconsin Alumni Research Foundation

January 16, 2018

Excerpt from Pages 13-26

* * *

ARGUMENT

I. SUBSTANTIAL EVIDENCE SUPPORTS THE INFRINGEMENT VERDICT FOR THE "PARTICULAR" LIMITATIONS

A. Substantial Evidence Supports the Jury's Finding that Apple's Processors Associate a Prediction with a "Particular" Load

The claims recite "a prediction associated with the particular [load] instruction." Appx265. Apple contends ***no reasonable jury*** could find this satisfied. But substantial evidence supports the verdict.

Professor Conte testified Apple's processors practice this limitation, evidenced by Apple's RTL code and technical documentation. Appx1441-1442; Appx1483. He explained: When a load and store mis-speculate, a 3-bit "prediction" counter is added to an LSDP entry, along with a Load Tag and Store Tag for the instructions. Appx1477-1483. A "hash function" generates each Tag by combining an instruction's memory address and register number, creating a 12-bit identifier. Appx1484-1485. And because the load that mis-speculated "***always*** generates the same 12-bit number" for its Load Tag, the "prediction" will ***always*** remain "associated with [the] particular [load] instruction" that mis-speculated (and thereby created the prediction). Appx2518; Appx2248-2249. Professor Conte also highlighted Apple's technical documentation, which states the Tags are sufficient to "uniquely identify a given store and load." Appx1489-1490; Appx10125-10138(PX45); Appx10131(PX45.7); Appx19284. Substantial evidence thus supports the verdict for "a prediction associated with the particular [load] instruction."

Apple relies on competing testimony. Dr. August testified that, because Load Tags are generated by a "hash function," there is a potential for "***aliasing***" events, where "[t]wo different loads" happen to "produce the same load tag," causing both loads to update the same prediction. Appx2160-2165; Appx2179-2180; Appx19877; Appx2155-2160. Because Load Tags can alias, Dr. August contended, Apple's processors never associate a "prediction" with a "particular" load. *Id.*; Appx2237-2238 (agreeing "the basis" of his non-infringement theory is "hashing can result in aliasing"). This did not entitle Apple to

JMOL, because "faced with competing expert testimony . . . the jury was free to disbelieve [Apple's] expert and credit [WARF's] expert." *Intellectual Ventures I LLC v. Motorola Mobility LLC*, 870 F.3d 1320, 1327 (Fed. Cir. 2017).

Importantly, even if Apple had a cognizable non-infringement defense for aliasing events (it does not), infringement would still be established by the times aliasing does ***not*** occur. "***[A]n accused device that 'sometimes, but not always, embodies a claim[] nonetheless infringes***.'" *Broadcom Corp. v. Emulex Corp.*, 732 F.3d 1325, 1333 (Fed. Cir. 2013). It was undisputed aliasing is "***very rare***" in Apple's processors. Appx1500-1501. Apple engineer Mr. Williams estimated aliasing occurs "less than 0.1 percent" of the time. Appx26467-26468(169:25-171:14); Appx2238-2244; App19286-19287. Dr. August agreed the frequency of aliasing is "very low." Appx2238. Likewise, Professor Conte estimated loads are "uniquely identified" by their Load Tags "99.98 percent of the time." Appx2517; Appx10131(PX45.7). Thus, as the District Court found, substantial evidence shows "Apple's processors are capable of operating for periods of time during which at least some of the load tags will not alias," and thus "the jury reasonably found this element satisfied." Appx205.

Further, even during rare aliasing events, the claims remain satisfied because they include the transition "***comprising***," which is "***open-ended and does not exclude additional, unrecited elements***." *CollegeNet, Inc. v. ApplyYourself, Inc.*, 418 F.3d 1225, 1235 (Fed. Cir. 2005); *see Vulcan Engineering Co., Inc. v. Fata Aluminum, Inc.*, 278

F.3d 1366, 1375 (Fed. Cir. 2002) ("It is irrelevant whether an element has capabilities in addition to that stated in the claim."). As Professor Conte testified—and Dr. August agreed—the load that mis-speculated ***always*** generates the same Load Tag, remaining associated with the same "prediction." Appx2518; Appx2248-2249. If another load aliases by generating the same Load Tag, the "prediction" merely becomes associated with ***two*** loads, including "the particular [load]" that mis-speculated. Appx1482-1484; Appx2515-2516. Thus, as the District Court found, "a reasonable jury could conclude that a prediction was associated with a particular load []instruction ***even if that same prediction may be associated with other load instructions***." Appx205. Apple says this "treats the 'particular' requirement as though it does not exist." Br. 23. Not so. The claims require a "prediction" associated with "the particular [load]" that mis-speculated, but ***do not exclude*** association with additional loads.

Finally, Apple argues "[t]he Load Tags are hashed—and therefore associated with multiple load instructions—regardless of whether 'aliasing' occurs." Br. 24. Apple contends "there are only 4,096 Load Tags available," while some programs have more loads. *Id.* This argument lacks merit.

First, as a factual matter, Apple has not contended ***every*** program has multiple loads associated with ***every*** Load Tag. Br. 24 (citing Appx2168-2169; Appx2296-2298; Appx2056-2061; Appx1605-1606). Programs can have fewer than 4,096

loads, and programs with many more loads can still have some Load Tags associated with only one load.

Second, Apple is requesting a ***claim construction*** never sought below—that the claims require a guaranteed unique Load Tag for each load. As discussed in Part I.B.2 below, this would exclude the preferred embodiment. But more directly, the claims do not require "Load Tags" at all, much less guaranteed unique Load Tags. They recite "a ***prediction*** associated with the particular [load] ***instruction***." Appx265. As Professor Conte explained, until a first aliasing event where two loads update the same prediction, that prediction will reflect the history of ***only one*** load instruction—*i.e.*, "the particular [load] instruction" that mis-speculated (and created the prediction). Appx1488-1489; Appx1502-1503; Appx2516-2518. Even Apple's Dr. August confirmed such scenarios exist (Appx2244-2250), and he agreed "a ***prediction*** will be associated with more than one load ***instruction***" only when "***aliasing*** occurs." Appx2239-2240. Apple's argument regarding 4,096 Load Tags is thus a red herring, as Dr. August conceded the frequency of aliasing remains "very low." Appx2238.[3]

Accordingly, substantial evidence supports the infringement verdict for "a prediction associated with the particular [load] instruction."

---

[3] As Professor Conte explained, aliasing remains rare in part because the LSDP holds only 192 predictions (far fewer than the 4,096 available Load Tags), which change frequently. Appx1500-1501; Appx1605-1606; Appx1436-1437; Appx1478-1479.

B. The District Court's "Plain Meaning" Construction for "Particular" Does Not Entitle Apple to a New Trial

1. The District Court Did Not Abuse its Discretion in Finding "Apple Has Waived Any Request" to Construe "Particular"

Apple contends the District Court "construed" the term "particular," then "refused to provide its construction to the jury" in violation of *O2 Micro Int'l Ltd. v. Beyond Innovation Tech. Co.*, 521 F.3d 1351 (Fed. Cir. 2008). Br. 20-21, 26. Apple misstates the record. The court ***declined*** to construe "particular," finding it needs no construction. Appx145. The court further ruled "***Apple has waived any request***" for a construction, as the court had "sided with Apple in its request that the term simply be given its plain and ordinary meaning." Appx147.

Apple's waiver defeats its *O2 Micro* argument. *Broadcom Corp. v. Qualcomm Inc.*, 543 F.3d 683, 694 (Fed. Cir. 2008) ("Unlike in *O2 Micro*, where the parties disputed the proper construction of a term at a pre-trial *Markman* hearing, Qualcomm here has failed to [timely] offer its proposed construction of 'networks'" and thus "has waived its right to request a construction of 'networks' and . . . implicitly conceded that the meaning of 'networks' is clear and not in need of construction."). This Court applies its own precedent for claim construction waiver issues, and reviews district court waiver findings for ***abuse of discretion***. *Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1384-85 (Fed. Cir. 2016). Apple claims error, but does not attempt to explain how the court abused its discretion. Br. 25-30.

As Apple concedes, "***Apple maintained prior to trial that the 'particular' limitations required no construction and should have their plain meaning.***" Br. 25. This was a strategic choice Apple made, even while advancing a non-infringement theory it now contends merited a specific construction. As the District Court observed, "the dispute ha[d] been known to the parties for some time," as "Apple's expert Dr. August articulated its reading of claim 1a in his March 2015 report [Appx26665-26675] and reiterated the theory in his August 2015 deposition [Appx26152]." Appx142. Specifically, Dr. August interpreted "a prediction associated with the particular [load] instruction" as requiring "unique identifier[s]," each guaranteed to be associated with "one and only one" load at all times, leaving zero possibility of aliasing events. Appx26152. WARF's expert Professor Conte opposed this theory in both his report and deposition. Appx15782-15788; Appx17817-17819. Because this issue was "central to Apple's theory of noninfringement," the court explained, "Apple should have raised the need for a construction once the dispute presented itself." Appx142. Yet Apple remained silent, apparently finding the option of arguing non-infringement under "plain meaning" preferable to risking an unfavorable claim construction. Br. 25.

At trial, WARF moved *in limine* to exclude Apple's flawed theory, arguing Apple planned to invoke an unsupported claim construction it had never requested. Appx18646-18724. In opposition, Apple reiterated its position that "particular" needs no

construction. Appx18725-18746. For example, Apple argued:

> Neither party asked the Court to construe the term "particular" during claim construction, at summary judgment, or in pre-trial motions. Apple has always maintained that the phrase—which uses only an ordinary word—***does not require any construction; the plain and ordinary meaning should apply***.

Appx18728.

> Apple believes that the word "particular" ***does not require any construction***, because the '752 patent uses the word in its ordinary sense.

Appx18730.

> Apple did not propose a claim construction for the "particular" limitation because it believes that ***the plain and ordinary meaning of this understandable word should apply***.

Appx18734.

The District Court ruled in Apple's favor, agreeing "the plain meaning of the claim terms 'the' and 'the particular'" needed no explanation and thus "there is no need for instructing the jury on the meaning of this term." Appx145. The court also denied WARF's motion to exclude Apple's non-infringement theory. Appx146. Apple thus presented its non-infringement argument to the jury, resting on the fourth day of trial. Appx2017; Appx2299.

Then, on the fifth trial day, Apple attempted an opportunistic gambit. Appx18762. Ignoring the court's "plain meaning" construction—which Apple had repeatedly requested—Apple parsed the court's order and plucked a partial statement from its reasoning. The court had stated the patent "discloses a prediction associated with a single load instruction, albeit one that is 'dynamic.'" Appx145. Apple took this sentence, altered it by striking "~~albeit one that is 'dynamic'~~," then transformed it into a requested jury instruction. Appx18762.

The District Court denied Apple's request, finding waiver:

> Consistent with Apple's position, the court denied WARF's request to exclude expert testimony and argument on this issue. ***The court also sided with Apple in its request that the term simply be given its plain and ordinary meaning.*** As such, ***Apple has waived any request*** to now insert a construction of the term into the closing jury instructions.

Appx147.

The Federal Circuit will "generally support a district court's case-management authority to set a schedule for claim construction that requires parties to take positions on various dates and holds the parties to these positions." *Wi-LAN USA*, 830 F.3d at 1385. This Court has repeatedly affirmed waiver findings on facts similar to those here. *Id.* (collecting cases).

Nuance Communications, Inc. v. Abbyy USA Software House, Inc., 813 F.3d 1368 (Fed. Cir. 2016), is particularly instructive. At *Markman*, Nuance argued that "identifying" should receive its "commonly understood meaning." *Id.* at 1372. The "court agreed with Nuance and thus construed the term 'identifying' to mean 'identifying,'" *i.e.*, no construction. *Id.* But it "became apparent during summary judgment . . . that the parties disagreed as to the plain and ordinary meaning of 'identifying.'" *Id.* "Nuance asked the district court to allow briefing on the meaning of 'identify' so that the dispute could be resolved before trial." *Id.* But the court refused, saying "it was 'too late to do construction' and that it was 'unnecessary.'" *Id.* Nuance contended this violated *O2 Micro*, but this Court disagreed, explaining: "the ***district court found in Nuance's favor*** by adopting the plain and ordinary meaning of the term 'identifying,'" and "[t]he fact that shortly before trial Nuance became ***dissatisfied with its own proposed construction and sought a new one*** does not give rise to an *O2 Micro* violation." *Id.* at 1373.

The court here similarly "sided with Apple in its request that the term simply be given its plain and ordinary meaning," then rejected Apple's mid-trial request for a construction, finding "Apple has waived any [such] request." Appx147. Indeed, the support for finding waiver is even stronger here than in *Nuance*. While Nuance changed its position "before trial" (*Nuance*, 813 F.3d at 1372), Apple maintained its "plain meaning" request until the fifth trial day (Appx18762), a day after Apple had rested (Appx2017; Appx2299), avoiding cross-examination of Apple's

witnesses on the application of its untimely proposed construction.

Apple cites no contrary authority, while citing a case that refutes its position. *See Function Media, L.L.C. v. Google Inc.*, 708 F.3d 1310, 1322 (Fed. Cir. 2013) ("FM may not object to the court's decision to instruct the jury to apply the claim construction that FM itself proposed."). Apple's remaining cases are inapposite. None involved a party advocating "plain meaning," the court agreeing, then the party requesting a different construction mid-trial. And none reviewed a district court's waiver ruling for abuse of discretion; they addressed waiver first on appeal. *GPNE Corp. v. Apple Inc.*, 830 F.3d 1365, 1372 (Fed. Cir. 2016) (first addressing "[o]n appeal," and finding no waiver because "at *Markman*, the parties disputed the construction"); *Creative Internet Advertising Corp. v. Yahoo!, Inc.*, 476 Fed. App'x 724, 728-29 (Fed. Cir. 2011) (non-precedential) (first addressing waiver on appeal); *Orix Credit All., Inc. v. Taylor Mach. Works, Inc.*, 125 F.3d 468, 478 (7th Cir. 1997) (not involving claim construction).

Apple contends it "could not have waived the request to provide the court's construction to the jury, because the construction did not exist earlier." Br. 27. But the court repeatedly stated its construction of "particular" was "***plain meaning***." Appx146; Appx147. Apple acknowledged as much when seeking a new construction on the eve of charging the jury. Appx18762 ("Apple recognizes that the Court has likely ruled on this issue already.").

Apple also tries to distinguish waiver of a ***claim construction*** from waiver of a ***jury instruction*** on

claim construction. Br. 27-28. This is a distinction without a difference. *Akamai Techs., Inc. v. Limelight Networks, Inc.*, 805 F.3d 1368, 1376 (Fed. Cir. 2015) (waiver in context of jury instructions); *Nuance*, 813 F.3d at 1373 (waiver in context of summary judgment). The court found waiver based on Apple's "request that the term simply be given its plain and ordinary meaning." Appx147. The court did not err, much less abuse its discretion.

* * *

APPENDIX N – Plaintiff-Appellee Wisconsin Alumni Research Foundation's Combined Petition for Panel Rehearing and Rehearing En Banc (November 28, 2018)

Nos. 17-2265, -2380

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

WISCONSIN ALUMNI RESEARCH FOUNDATION,
Plaintiff - Appellee,
v.
APPLE INC.,
Defendant - Appellant.

Appeals from the United States District Court for the Western District of Wisconsin in case no. 3:14-cv-00062-wmc, Judge William M. Conley

---

PLAINTIFF-APPELLEE WISCONSIN ALUMNI RESEARCH FOUNDATION'S COMBINED PETITION FOR PANEL REHEARING AND REHEARING EN BANC AND SUPPLEMENTAL APPENDIX

MORGAN CHU
GARY N. FRISCHLING
JASON SHEASBY
ALAN J. HEINRICH
CHRISTOPHER T. ABERNETHY
AMY E. PROCTOR

IRELL & MANELLA LLP
1800 Avenue of the Stars, Suite 900
Los Angeles, California 90067 (310) 277-1010

Attorneys for Appellee Wisconsin Alumni Research Foundation

November 28, 2018

CERTIFICATE OF INTEREST

Counsel for Appellee Wisconsin Alumni Research Foundation ("WARF") certifies the following:

1. The full name of every party represented by me is: Wisconsin Alumni Research Foundation

2. The parties named in the caption are the real parties in interest represented by me.

3. All parent corporations and any publicly held companies that own 10 percent or more of the stock of the party represented by me are: None.

4. The names of all law firms and the partners or associates that appeared for the party now represented by me in the trial court or are expected to appear in this Court are: Irell & Manella LLP: Morgan Chu, Alan Heinrich, Gary Frischling, Jason Sheasby, Joseph Lipner (former), Amy Proctor, Christopher Abernethy, Anthony Rowles, H. Annita Zhong, Jackson Trugman (former), Brian Eggleston (former); Godfrey & Kahn S.C.: Jennifer Gregor, Bryan Cahill.

5. The title and number of any case known to me to be pending in this or any other court of agency that will directly affect or be directly affected by this court's decision in the pending appeal are:

W*isconsin Alumni Research Foundation v. Apple Inc.*,

Case Nos. 2017-2265, -2380 (Fed. Cir.);

Wisconsin Alumni Research Foundation v. Apple Inc.,

Case No. 3:14-cv-00062-WMC (W.D. Wis.); and

Wisconsin Alumni Research Foundation v. Apple Inc.,

Case No. 3:15-cv-00621-WMC (W.D. Wis.).

Dated: November 28, 2018 Respectfully submitted,

IRELL & MANELLA LLP

By: */s/ Morgan Chu*
Morgan Chu

Attorneys for Appellee Wisconsin Alumni Research Foundation

TABLE OF CONTENTS

TABLE OF AUTHORITIES

RULE 35(B)(2) STATEMENT

Based on my professional judgment, I believe:

The panel decision is contrary to the following decisions of the Supreme Court and this Court: *Teva Pharmaceuticals USA v. Sandoz*, 135 S.Ct. 831 (2015); *Apple v. Samsung Electronics,* 839 F.3d 1034 (Fed. Cir. 2016) (en banc); *Wi-LAN USA v. Apple*, 830 F.3d 1374 (Fed. Cir. 2016); *Nuance Commc'ns v. ABBYY USA Software House*, 813 F.3d 1368 (Fed. Cir. 2016); *O2 Micro Int'l v. Beyond Innovation Tech.*, 521 F.3d 1351 (Fed. Cir. 2008); *Markman v. Westview Instruments*, 52 F.3d 967 (Fed. Cir. 1995).

This appeal requires answers to precedent-setting questions of exceptional importance:

1. May this Court override a district court's case-management authority by adopting a claim construction the court ruled the appellant had "waived," without explanation or finding abuse of discretion, disregarding well-settled principles of appellate review?

2. Must this Court adhere to Teva, and give deference to a jury's fact finding regarding a POSITA's understanding, when the jury was properly instructed to apply claim language's "plain and ordinary meaning as viewed [by a POSITA]"?

3. On substantial evidence review, may this Court make findings contrary to admitted facts and substitute its own technical analysis for expert testimony supporting the verdict?

4. May this Court apply a new construction to grant JMOL on a record developed by nearly two

years of litigation and a trial held without that construction, instead of remanding for proceedings consistent with the new construction?

/s/Morgan Chu
ATTORNEY OF RECORD FOR WARF

INTRODUCTION

A jury found Apple's microprocessors infringed WARF's 5,781,752 patent after the district court ruled Apple had waived any construction of a key limitation. But on "substantial evidence" review, the panel construed that limitation, applied its construction, and granted JMOL of non-infringement, exceeding its "appellate function." *Apple*, 839 F.3d at 1039. WARF requests panel rehearing and rehearing *en banc*.

Microprocessors can run faster by executing instructions "out-of-order." But this is not always possible, as some instructions require data from prior instructions. For example, "store" instructions write data to memory and "load" instructions read data from memory. "Data dependence" is when a "load" must read data previously written by a "store." At execution time, the processor may not know if a store and load are dependent. But it may still execute them out-of-order, "speculating" they are probably not dependent. If correct, performance improves. But if incorrect, a "mis-speculation" occurs, causing an error and harming performance. WARF's '752 invention claims a "Predictor" circuit in a processor that predicts when it is likely safe to execute a load before a store, enabling significantly faster execution.

The claims recite "detecting a mis-speculation" by "a [load] instruction," then creating "a prediction associated with the particular [load] instruction." Appx265. Given Apple's waiver, "associated with the particular" was not construed below. Yet the panel held, with little analysis: "In our view, the plain meaning of 'particular,' as understood by a [POSITA]

after reading the '752 patent, requires the prediction to be associated with a *single* load instruction." Op. 11. The panel made four significant errors in reaching and applying this construction.

First, it disregarded the district court's waiver ruling. Apple maintained "the 'particular' limitations required no construction." Apple Br. 25. But after both parties rested their cases-in-chief, Apple sought construction in a jury instruction. Appx18762. Judge Conley refused, holding "Apple has waived any… construction." Appx147. Waiver rulings are reviewed for "abuse of discretion." *Wi-LAN*, 830 F.3d at 1385. Yet, without mentioning waiver or abuse of discretion, the panel adopted Apple's waived construction. This undercuts a trial court's ability to conduct fair proceedings and encourages gamesmanship, rewarding parties for avoiding construction below then pursuing new "plain meaning" constructions on appeal. ***The en banc Court should clarify that overriding a "district court's case- management authority" requires finding "abuse of discretion."*** *Id.*[1]

Second, the panel disregarded the Supreme Court's *Teva* decision. Claim construction can include both (1) a "factual finding" regarding "meaning to a [POSITA]," and (2) a "legal analysis" of whether a POSITA "would ascribe that same meaning to that term *in the context of the specific patent claim*." *Teva*, 135 S.Ct. at 841. Because Apple argued the limitation "does not require construction" (Appx18728), Judge Conley properly instructed the jury to apply "plain

---

[1] All bold emphasis added.

and ordinary meaning as viewed [by a POSITA]." Appx163. The jury heard expert testimony regarding a POSITA's view, and its decision is subject to the "ordinary rule governing appellate review of factual matters." *Teva*, 135 S.Ct. at 838. But rather than review the jury's findings under the "ordinary rule" (substantial evidence), the panel substituted Apple's proposed "plain meaning" *de novo*. *Teva* does not permit "turning what [was] fundamentally a factual question for the jury regarding whether the accused systems and features infringe the patent claims into a legal one for the court." *NobelBiz, Inc. v. Global Connect, L.L.C.*, 876 F.3d 1326, 1327 (Fed. Cir. 2017) (O'Malley, J., dissenting). ***"It is time [for] much-needed guidance en banc" on the review of "plain-and-ordinary-meaning construction[s]."*** *Id.*

Third, the verdict was well-supported even under the panel's construction requiring "a prediction associated with a *single* load instruction" (Op. 11). WARF's expert testified that, per his independent calculations, loads are ***"uniquely identified for 99.98 percent of the time."*** Appx2516-2517; *see* Appx1620. Apple's expert offered corroborating testimony: "the chances in which a prediction will be associated with ***more than one*** load instruction is ***.0 -- .0.1 percent."*** Appx2239- 2240; *see* Appx2243. ***The panel ignored all this testimony.***

The panel also misunderstood key evidence regarding "aliasing," which is when one prediction is updated by multiple loads. Apple's Reply admits its processors "***only rarely … alias***" (Apple Reply 4), and Apple's expert testified the "***frequency of***

aliasing is fairly low" (Appx2238). ***Yet the panel found contrary to the admitted facts***, stating WARF's expert "jumped to the conclusion" that "aliasing is very rare" and that this was "unreasonable to infer." Op. 14-15.

These errors flow from Apple's sleight-of-hand argument that "the fact that there are only 4,096 Load Tags available even though Apple's iOS software contains millions of load instructions confirms that each Load Tag *always* represents multiple load instructions during a processor's operation." Apple Br. 24; Op. 16. A load tag is an abbreviation (or "hash") of a load instruction's address. Apple misleadingly focused on load tags, while the claims recite a "prediction" associated with a "load instruction." There are at most ***192 predictions*** maintained in Apple's Predictor at any time; stale predictions are constantly discarded to free space. *See* Part III.A. It is irrelevant that "Apple's operating system alone contains millions of load instructions" (Op. 16), because ***only a small subset of load instructions execute during a prediction's lifespan***. A deleted prediction is not associated with anything.

The jury heard both parties' arguments and evidence and rejected Apple's position, reasonably concluding that the "millions" of load instructions executed after a given prediction's deletion are not "associated with" it, because instructions cannot affect, or be affected by, a prediction that no longer exists. ***The panel exceeded its appellate function by assuming the role of factfinder***, substituting its theoretical analyses for expert testimony that in

actual operation loads are "uniquely identified" more than 99.9% of the time.

Fourth, at minimum, remand is required for proceedings consistent with the panel's new construction. WARF formulated its trial strategy in reliance on Apple having sought no construction. This included choosing not to try doctrine of equivalents and narrowing the evidence presented. The Court should make clear *en banc* that, where the proposed construction was not requested until the very end of trial, remand is necessary to maintain basic fairness. Apple should not benefit from hiding its construction request until it was too late for a response. ***En banc guidance is needed to establish a clear rule that avoids panel-dependent outcomes***.

ARGUMENT

I. THIS COURT CANNOT IGNORE AN EXPRESS WAIVER RULING

A district court's ability to enforce claim construction procedures, including timely disclosure of positions, is essential to fair proceedings. This Court accordingly reviews exercise of a trial court's "case-management authority" for "abuse of discretion." *Wi-LAN*, 830 F.3d at 1385. The panel abrogated its appellate role, reinstating Apple's waived claim construction (Op. 11) without discussing waiver, much less finding abuse of discretion. Case-management decisions should not be discarded on a whim.

"Apple maintained prior to trial that the 'particular' limitations required no construction" (Apple Br. 25) and that its expert would apply "plain

and ordinary meaning." Appx18728. Then, on the fifth trial day—after both parties rested their cases-in-chief—Apple reversed course and sought this jury instruction: "'a prediction associated with the particular [load] instruction' means 'a prediction associated with a single load instruction.'" Appx18762. Judge Conley refused, explaining: "The court [] sided with Apple in its request that the term simply be given its plain and ordinary meaning. As such, ***Apple has waived any request to now insert a construction*** of the term into the closing jury instructions." Appx147.

This Court routinely affirms waiver rulings. *Wi-LAN*, 830 F.3d at 1385 (collecting cases). *Nuance*, 813 F.3d at 1373, is instructive. At *Markman*, the "court found in Nuance's favor by adopting [] plain and ordinary meaning." *Id.* The court did not abuse its discretion by then declining to construe when Nuance sought a construction "shortly before trial." *Id.* Here, Apple's request came much later, hours before closing arguments. *See also Broadcom v. Qualcomm*, 543 F.3d 683, 694 (Fed. Cir. 2008) ("waived" construction eliminates "*O2 Micro*" concerns).

The panel should have declined to entertain Apple's waived construction. If litigants can argue construction on appeal after waiver, the "district court's case-management authority to set a schedule for claim construction" that "holds the parties to the[ir] positions" will be toothless. *Wi-LAN*, 830 F.3d at 1385. The Court should clarify that overriding a waiver ruling requires finding "abuse of discretion."

II. *TEVA* REQUIRED DEFERENCE TO THE JURY'S FACT FINDINGS ON "PLAIN AND ORDINARY MEANING" TO A POSITA

The panel contravened the Supreme Court's *Teva* decision by ignoring that "plain meaning" to a POSITA is a fact finding. "It is time [for] much-needed guidance *en banc*" on the review of "plain-and-ordinary-meaning construction[s]." *NobelBiz*, 876 F.3d at 1327 (O'Malley, J., dissenting).

Under *Teva*, claim construction may include (1) a "factual finding" regarding "meaning to a [POSITA]," and (2) a "legal analysis" of whether a POSITA "would ascribe that same meaning to that term *in the context of the specific patent claim*." 135 S.Ct. at 841. By arguing the limitation "does not require construction" (Appx18728) and uses language "in its ordinary sense" (Appx18730), Apple conceded a POSITA would not ascribe a unique meaning in this patent. Judge Conley, finding waiver, properly instructed the jury to apply "plain and ordinary meaning as viewed [by a POSITA]." Appx163.

The jury heard expert testimony on how a POSITA would understand "detecting a mis-speculation" by "a [load] instruction," then creating "a prediction associated with the particular [load] instruction." Appx265; *Teva*, 135 S.Ct. at 841 ("Experts may be examined to explain terms of art, and the state of the art."). Dr. Conte testified a POSITA would understand "a prediction associated with the particular [load] instruction" means a prediction associated with "the load that produced the mis-speculation." Appx1481-1482. Unlike the panel's

construction, this did not exclude occasional association with additional loads. Appx1491-1492.

As the jury was properly instructed to apply "plain and ordinary meaning," review was for "substantial evidence." *Teva*, 135 S.Ct. at 838 (applying "ordinary rule governing appellate review of factual matters"); *Asetek Danmark v. CMI USA*, 852 F.3d 1352, 1359-60 (Fed. Cir. 2017) (where parties "provide the jury only with the claim language … the only question is one of substantial evidence"). This required "presum[ing] the jury resolved all underlying factual disputes in favor of the verdict," *Apple*, 839 F.3d at 1040, including crediting Dr. Conte's testimony.

Judge Conley applied the correct standard in denying JMOL, finding the "jury could conclude that a prediction was associated with a particular load instruction ***even if that same prediction may be associated with other load instructions***." Appx205. The panel cryptically "note[d]" this, but said nothing. Op. 11 n.6. Judge Conley properly deferred to the jury. Appx1491 (allowing Dr. Conte's testimony because "it is ***the jury*** who will decide … plain meaning"). The panel did not.

Instead, the panel determined plain meaning "[i]n our view" (Op. 11), undermining the role of trial judge and jury. If this Court may always construe "plain meaning" *de novo*, litigants may strategically seek no construction, then pursue "plain meaning" constructions on appeal. Gamesmanship will abound as litigants try to make purported requirements express, or shift positions, after it is too late to

respond. *See* Part IV. Adhering to *Teva* provides the solution.

III. ON SUBSTANTIAL EVIDENCE REVIEW, THE PANEL IMPROPERLY IGNORED EXPERT TESTIMONY AND CONTRADICTED ADMITTED FACTS

Even under the panel's construction requiring "a prediction associated with a *single* load instruction," there was substantial evidence that Apple's Predictor infringes "99.98 percent of the time" according to independent calculations by WARF's expert (Appx2516-2157; Appx1620) and "99.9 percent of the time" according to Apple's witnesses (Appx1501-1502; Appx2238-2240). *Cf. Broadcom v. Emulex*, 732 F.3d 1325, 1333 (Fed. Cir. 2013) (infringing "sometimes" suffices).

A. Apple's Predictor Circuit Can Only Track 192 Predictions At Any Time

The '752 circuit predicts whether pairs of store and load instructions are likely dependent. Appx264(11:26-31). Predictions (109) reside in a prediction table (44):

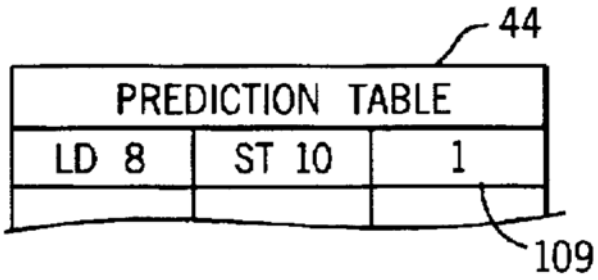

Appx255(Fig.5). A table entry is created when a store-load pair mis-speculates. Appx265(13:38-45). It

then prevents speculation by subsequent occurrences of that store-load pair if the prediction is above a threshold (Appx264(11:26-40)), while also updating the prediction. Appx264-265(12:10-17,12:50-13:3). The table is "stored in a small, high speed memory." Appx260(3:57-60). Once full, stale predictions are "deleted" and "replace[d]" as space is needed. Appx265(13:3-17).

Apple's processors include a "Load-store dependency predictor" circuit ("Predictor") for "predicting which loads are likely dependent on older stores." Appx10125. It maintains a "***192 entry***" prediction table for "store, load pairs." Appx2055. Like the '752 patent (Appx265(13:3-17)), because Apple's table is small, "stale entries" are deleted by a "replacement policy" and "'age-out' scheme" as needed. Appx10131; Appx10136. Thus, at any time, there are at most 192 predictions in Apple's Predictor, and they are constantly being replaced. *Id*. Deleted predictions are lost and are not associated with anything.

In the load portion of each table entry, Apple's Predictor stores a "load tag," which is a 12-bit number generated by a "hashing function." Op. 7. There are 4,096 possible 12-bit hash values. *Id*. A given load "always generates [the] same 12-bit number." Appx2518. But occasionally more than one load instruction can "update the same prediction." Op. 7. This is referred to as "aliasing." *Id*.

Aliasing harms performance, so Apple designed its hash to "minimiz[e] the frequency of aliasing." Appx2243. Apple admits that two loads "only rarely cause the same prediction to be updated ('alias')."

Apple Reply 4. Because the table has only "192 entries, you have less entries, less chance of aliasing." Appx1500-1502.

B. The Panel Disregarded Substantial Evidence That Apple's Predictor Uniquely Identifies Loads More than 99% of the Time

In granting JMOL of non-infringement under its construction, the panel held: "given that only 4,096 load tags are possible, and that ***Apple's operating system alone contains millions of load instructions***, the only reasonable inference to draw is that ***load tags*** will always represent multiple ***load instructions***." Op. 16. The panel erred in at least four ways.

First, the panel lost sight of the claim language, which as construed requires "a prediction associated with a *single* load instruction." Op. 11. The claimed association is between "prediction" and "load instruction," not between "***load tags***" and "load instructions." Yet the panel granted JMOL because it found "the only reasonable inference to draw is that ***load tags*** will always represent multiple load instructions." Op. 16; *see* Op. 15 (finding "it is not reasonable to infer that ***load tags***, in practice, uniquely identify load instructions.").

Second, the panel's focus on "millions of load instructions" (Op. 16) was misplaced, because ***only a small subset of instructions execute during a prediction's lifespan***. With only a ***192 entry*** table, predictions are short-lived, as they are constantly deleted and replaced. *See* Part III.A. The panel was concerned multiple load instructions may "update the

same prediction" and that a prediction "may impact" multiple instructions. Op. 7. But only instructions executed during a prediction's lifespan can possibly update, or be impacted by, that prediction.

Appx10136; Appx1500-1502. A deleted prediction cannot interact with instructions executed after it no longer exists. *Id*. The jury was not compelled to find a prediction is "associated with" a load instruction executed after the prediction's deletion.

That is why the low frequency of "aliasing" matters. "Aliasing" is "when two load instructions actually update the same prediction in operation because they share the same load tag." Op. 15. As a corollary, if a prediction did not alias, it means the prediction was updated by—and the load tag belonged to—***only one*** load instruction executed during that prediction's lifespan. Appx1500-1502; Appx1620; Appx2516- 2517; Appx2238-2240; Appx2243. The jury could find a prediction absent aliasing is "associated with a *single* load instruction," because that prediction is only updated by, and only ever impacts, one load instruction. And as WARF's expert Dr. Conte explained, because "aliasing is very rare," predictions are virtually always associated with a single load instruction. Appx1501-1502.

Third, the panel erred by finding "WARF's expert jumped to the conclusion that 'aliasing is very rare'" and that this was "unreasonable to infer." Op. 14-15. This flatly contradicted Apple's admissions and both parties' witnesses. Apple admitted in its briefing: "two instructions within a group ***will only rarely*** cause the same prediction to be updated ('***alias***')." Apple Reply 4; *see also* SAppx0002 (Apple JMOL brief:

"WARF pointed to evidence suggesting 'aliasing' is rare"). Apple's expert concurred: "***frequency of aliasing is fairly low***." Appx2238. Apple's engineer, Mr. Williams, also concurred, testifying the "average amount of aliasing" is "***in a similar range***" to the "***less than .1 percent***" performance impact caused by aliasing. Appx26467-26468(169:25-170:11). Dr. Conte concluded the same, testifying "aliasing is very rare" (Appx1501-1502) based on his independent calculation:

> Q. What is the chance of there being an alias in the A7 processor?
>
> A. So Mr. Williams said that it was in the range of 0.1 percent. ***I actually calculated 0.02 when I calculated it out by hand***.

Appx1620.

Rather than acknowledging the jury could reasonably credit the testimony that aliasing is in fact "rare," the panel focused on emails from Apple's engineer Stephan Meier, from which it concluded that the "0.1% statistic" represented only aliasing's "performance impact," not frequency. Op. 14. Of course, the jury was free to ignore these emails; Dr. Conte's independent "by hand" calculations provide substantial evidence and stand on their own. But Mr. Meier also testified ***performance impact and frequency are closely correlated*** (SAppx0004-0005); performance impact is low because "frequency of aliasing is fairly low." SAppx0005(65:8-15,65:21-24). Dr. August confirmed Apple "minimize[d] the performance impact" by "minimizing the frequency of aliasing." Appx2243. In short, every witness who addressed the issue confirmed aliasing is rare.

The panel also suggested it wouldn't matter if aliasing were infrequent, stating that one cannot conclude "that load instructions rarely hash to the same load tag, merely because the frequency of load instructions actually updating the same prediction during operation is low." Op. 15. But this had it backwards, again because the claimed association is between "prediction" and "load instruction," not between "load tag" and "load instruction." It is irrelevant how often multiple load instructions hash to the same tag; what matters is that most ***predictions*** in Apple's circuit are associated with a single ***load instruction***.

Finally, the panel ignored testimony from both parties' experts that predictions are almost always associated with a single load. WARF's Dr. Conte offered his calculation that, given how Apple's 192-entry Predictor operates, load instructions are "***uniquely identified for 99.98 percent of the time***." Appx2516- 2517; see Appx1620. Apple's Dr. August offered corroborating testimony:

> Q. And if the jury accepts what Mr. Williams said in his deposition was truthful, that would mean that the chances in which this aliasing occurs, ***the chances in which a prediction will be associated with more than one load instruction is .0 -- .0.1 percent?***
>
> A. ***Correct***.

Appx2239-2240; see Appx2243; Appx26467-26468(169:25-170:11) (Williams: aliasing in "range" of "less than .1 percent"). ***The panel ignored all this testimony***.

Substantial evidence review does not permit this Court to re-weigh evidence or choose which experts to credit—that was for the jury. *Knapp v. Eagle Property Mgmt.*, 54 F.3d 1272, 1283 (7th Cir. 1995) (must "view all evidence and reasonable inferences in the light most favorable to the prevailing party"). The jury heard substantial evidence speaking to the panel's construction: more than 99% of the time, predictions are associated with a single load. The verdict should be affirmed.

IV. REMAND IS THE APPROPRIATE REMEDY IF NOT AFFIRMED

Without remanding, the panel granted JMOL under its construction, finding "insufficient evidence" in a trial record developed without that construction. Op. 12-13. WARF relied on the fact that Apple sought "no construction" (Apple Br. 25) until mere hours before closing arguments. WARF thus "tried the case and presented appellate arguments based on this understanding." *Apple*, 834 F.3d at 1034 n.4.

With the panel's construction, the trial record would have been different. Based on Apple's positions and the court's rulings, WARF did not assert the doctrine of equivalents. Under the panel's construction, ***WARF would have tried DoE***.

Similarly, because there are 4,096 load tags, the panel suggested WARF could have prevailed on literal infringement had it offered evidence "programs can have fewer than 4,096 load instructions." Op. 16 n.9. While not featured at trial—as no construction necessitated it—this is undisputed. WARF Br. 17; Apple Reply 2 n.1. ***WARF would present such evidence on remand***.

In addition to being unfair, not remanding violated the Seventh Amendment. Claim construction preserves the "right to a jury trial on the application of the properly construed claim." *Markman*, 52 F.3d at 984. Remand was necessary because the first construction was on appeal. *Odetics v. Storage Tech.*, 185 F.3d 1259, 1272 (Fed. Cir. 1999) (Court "cannot, consistent with the Seventh Amendment, evaluate a jury's verdict based on…a legal standard not given to the jury."); *Malta v. Schulmerich Carillons*, 952 F.2d 1320, 1345 (Fed. Cir. 1991) ("This court, by retroactively imposing new requirements for proving infringement, then re-finding the facts…has denied this litigant's [Seventh Amendment] right.") (Newman, J., dissenting).

In a case such as this, "even if the court erred in [not construing], the remedy [sh]ould be, at most, a new trial." *Eon Corp. IP Holdings v. Silver Spring Networks*, 815 F.3d 1314, 1329 (Fed. Cir. 2016) (Bryson, J., dissenting). Yet panel-dependent outcomes abound. Many remand if "the jury was not presented with … the [panel's] construction." *Virnetx v. Cisco Sys.*, 767 F.3d 1308, 1319 (Fed. Cir. 2014). Some "reverse the district court's determination with respect to JMOL without remand." *SimpleAir v. Sony Ericsson Mobile Commc'ns*, 820 F.3d 419, 425 (Fed. Cir. 2016). Others consider vague factors like "the degree of difference between" constructions. *Finistar v. DirecTV Group*, 523 F.3d 1323, 1333 (Fed. Cir. 2008). ***En banc guidance is needed to establish a consistent rule***.

This is not a close case. Apple sought "no construction" and was found to have "waived any

request" for a construction. See Part I. The proper roles of an appellate court, trial court and jury—and basic fairness—at minimum require a remand for proceedings consistent with the panel's new construction.

Dated: November 28, 2018 Respectfully submitted,

IRELL & MANELLA LLP
Morgan Chu
Gary N. Frischling
Jason Sheasby
Alan Heinrich
Christopher T. Abernethy
Amy E. Proctor

By: */s/ Morgan Chu*
Morgan Chu

Attorneys for Appellee Wisconsin Alumni Research Foundation

APPENDIX O – Rule 50 of the Federal Rules of Civil Procedure

RULE 50. JUDGMENT AS A MATTER OF LAW IN A JURY TRIAL; RELATED MOTION FOR A NEW TRIAL; CONDITIONAL RULING

(a) JUDGMENT AS A MATTER OF LAW

(1) *In General.* If a party has been fully heard on an issue during a jury trial and the court finds that a reasonable jury would not have a legally sufficient evidentiary basis to find for the party on that issue, the court may:

(A) resolve the issue against the party; and

(B) grant a motion for judgment as a matter of law against the party on a claim or defense that, under controlling law, can be maintained or defeated only with a favorable finding on that issue.

(2) *Motion.* A motion for judgment as a matter of law may be made at any time before the case is submitted to the jury. The motion must specify the judgment sought and the law and facts that entitle the movant to the judgment.

(b) RENEWING THE MOTION AFTER TRIAL; ALTERNATIVE MOTION FOR A NEW TRIAL. If the court does not grant a motion for judgment as a matter of law made under Rule 50(a), the court is considered to have submitted the action to the jury subject to the court's later deciding the legal

questions raised by the motion. No later than 28 days after the entry of judgment—or if the motion addresses a jury issue not decided by a verdict, no later than 28 days after the jury was discharged—the movant may file a renewed motion for judgment as a matter of law and may include an alternative or joint request for a new trial under Rule 59. In ruling on the renewed motion, the court may:

(1) allow judgment on the verdict, if the jury returned a verdict;

(2) order a new trial; or

(3) direct the entry of judgment as a matter of law.

(c) GRANTING THE RENEWED MOTION; CONDITIONAL RULING ON A MOTION FOR A NEW TRIAL.

(1) *In General*. If the court grants a renewed motion for judgment as a matter of law, it must also conditionally rule on any motion for a new trial by determining whether a new trial should be granted if the judgment is later vacated or reversed. The court must state the grounds for conditionally granting or denying the motion for a new trial.

(2) *Effect of a Conditional Ruling*. Conditionally granting the motion for a new trial does not affect the judgment's finality; if the judgment is reversed, the new trial must proceed unless the appellate court orders otherwise. If the motion for a new trial is conditionally denied, the appellee may assert

error in that denial; if the judgment is reversed, the case must proceed as the appellate court orders.

(d) TIME FOR A LOSING PARTY'S NEW-TRIAL MOTION. Any motion for a new trial under Rule 59 by a party against whom judgment as a matter of law is rendered must be filed no later than 28 days after the entry of the judgment.

(e) DENYING THE MOTION FOR JUDGMENT AS A MATTER OF LAW; REVERSAL ON APPEAL. If the court denies the motion for judgment as a matter of law, the prevailing party may, as appellee, assert grounds entitling it to a new trial should the appellate court conclude that the trial court erred in denying the motion. If the appellate court reverses the judgment, it may order a new trial, direct the trial court to determine whether a new trial should be granted, or direct the entry of judgment.