Nos. 20-2011, -2094

IN THE
United States Court of Appeals for the Federal Circuit

APPLE INC.,
Plaintiff-Appellant,

v.

WI-LAN INC.,
Defendant-Cross-Appellant.

On Appeal from the United States District Court
for the Southern District of California
No. 3:14-cv-02235-DMS-BLM
Hon. Dana M. Sabraw

OPENING BRIEF AND NONCONFIDENTIAL ADDENDUM OF APPELLANT APPLE INC.

Sean C. Cunningham
Stanley J. Panikowski
Erin P. Gibson
DLA PIPER LLP (US)
401 B Street, Suite 1700
San Diego, CA 92101
(619) 699-2700

Mark S. Davies
Thomas Fu
James Anglin Flynn
Katherine M. Kopp
ORRICK, HERRINGTON & SUTCLIFFE LLP
1152 15th Street, NW
Washington, DC 20005
(202) 339-8400

Max Carter-Oberstone
ORRICK, HERRINGTON & SUTCLIFFE LLP
405 Howard Street
San Francisco, CA 94105

Counsel for Appellant

CLAIM LANGUAGE AT ISSUE

U.S. Patent No. 8,457,145, Claim 9 (Representative)

9. A subscriber unit for a wireless communication system, wherein the wireless communication system includes a plurality of subscriber units in communication with an associated base unit, comprising:

a plurality of queues, each queue for grouping data based on the QoS; and

a media access (MAC) module configured to

set an initial value for a timer associated with a queue, and

periodically, on expiration of the value of the timer, transmit a bandwidth request indicating an amount of bandwidth required for transmitting the data from the queue.

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

CERTIFICATE OF INTEREST

Case Number 2020-2011; 2020-2094

Short Case Caption Apple Inc. v. Wi-LAN Inc.

Filing Party/Entity Apple Inc.

Instructions: Complete each section of the form. In answering items 2 and 3, be specific as to which represented entities the answers apply; lack of specificity may result in non-compliance. **Please enter only one item per box; attach additional pages as needed and check the relevant box**. Counsel must immediately file an amended Certificate of Interest if information changes. Fed. Cir. R. 47.4(b).

I certify the following information and any attached sheets are accurate and complete to the best of my knowledge.

Date: 11/20/2020

Signature: /s/ Mark S. Davies

Name: Mark S. Davies

1. Represented Entities. Fed. Cir. R. 47.4(a)(1).	**2. Real Party in Interest.** Fed. Cir. R. 47.4(a)(2).	**3. Parent Corporations and Stockholders.** Fed. Cir. R. 47.4(a)(3).
Provide the full names of all entities represented by undersigned counsel in this case.	Provide the full names of all real parties in interest for the entities. Do not list the real parties if they are the same as the entities. ☑ None/Not Applicable	Provide the full names of all parent corporations for the entities and all publicly held companies that own 10% or more stock in the entities. ☑ None/Not Applicable
Apple Inc.		

☐ Additional pages attached

4. Legal Representatives. List all law firms, partners, and associates that (a) appeared for the entities in the originating court or agency or (b) are expected to appear in this court for the entities. Do not include those who have already entered an appearance in this court. Fed. Cir. R. 47.4(a)(4).

☐ None/Not Applicable ☑ Additional pages attached

Please see attachment		

5. Related Cases. Provide the case titles and numbers of any case known to be pending in this court or any other court or agency that will directly affect or be directly affected by this court's decision in the pending appeal. Do not include the originating case number(s) for this case. Fed. Cir. R. 47.4(a)(5). See also Fed. Cir. R. 47.5(b).

☑ None/Not Applicable ☐ Additional pages attached

6. Organizational Victims and Bankruptcy Cases. Provide any information required under Fed. R. App. P. 26.1(b) (organizational victims in criminal cases) and 26.1(c) (bankruptcy case debtors and trustees). Fed. Cir. R. 47.4(a)(6).

☑ None/Not Applicable ☐ Additional pages attached

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

Case No. 2020-2011; 2020-2094
Apple Inc. v. Wi-LAN Inc.

ADDITIONAL PAGE TO CERTIFICATE OF INTEREST (FORM 9)

4. Legal Representatives. List all law firms, partners, and associates that (a) appeared for the entities in the originating court or agency or (b) are expected to appear in this court for the entities. Do not include those who have already entered an appearance in this court. Fed. Cir. R. 47.4(a)(4).

Mark D. Fowler, John Allcock, Tiffany C. Miller, Robert Buergi, Robert C. Williams, Andrew Stein, Jacob Anderson, Amy H. Walters, Peter P. Maggiore (all currently or formerly with DLA Piper LLP (US))

Milbank, Tweed, Hadley & McCloy LLP, Christopher J. Gaspar, Mark C. Scarsi, Dat T. Nguyen, Jennifer L. Miremadi, Miguel J. Ruiz, Michael K. Sheen

Eisner APC, Ashlee N. Lin

TABLE OF CONTENTS

Page

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

Rules

Other Authorities

STATEMENT OF RELATED CASES

On August 1, 2016, a panel of this Court (Lourie, Bryson, Chen, JJ.) decided *Wi-LAN USA, Inc. v. Apple Inc.*, No. 2015-1256. That appeal involved the same parties and involved patents related to one of the two patents at issue here. The decision is reported at 830 F.3d 1374.

Apple and its counsel are unaware of any cases pending in this or any other court that will directly affect or be directly affected by this Court's decision in this case.

INTRODUCTION

This is a case about a patent owner overreaching. At each stage of the litigation—claim construction, infringement, and damages—Wi-LAN Inc. has sought to claim more, to capture more, and to cash-in more than it had earned.

This Court has already rejected Wi-LAN's prior attempt to benefit from expansive claim constructions related to similar wireless communication technology. Wi-LAN's prior constructions, this Court found, were "incompatible" with the specifications' "consistent descriptions" of the inventions. *Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1376-77 (Fed. Cir. 2016). The same result is appropriate here. All of the asserted claims are directed to one component of a *three*-component system: a wireless "base station", a "subscriber unit", and "user" devices. But Wi-LAN's construction reads the "subscriber unit" term out of the claims. Once again, Wi-LAN's expansive constructions are incompatible with the intrinsic evidence.

Not only is Wi-LAN stretching its patents, it's also proposing an expansive (and incorrect) understanding of the meaning of "direct" infringement. Under Wi-LAN's theory of direct infringement, *Apple* created the voice and data connections on its iPhones. Appx10251. Apple does no such thing. Rather, the user and the base station alter the phone to create those connections. Wi-LAN's expansive theory of direct infringement is inconsistent with settled direct infringement case law.

Wi-LAN's final reach comes when money is on the table. Wi-LAN has never made a serious attempt to value the economic worth of its claimed inventions. Wi-LAN's damages model applied a series of generic and arbitrary "adjustments" to the royalty rates of thousand-patent Wi-LAN portfolio licenses based on three obscure, niche market phone providers. But Wi-LAN's generic and arbitrary "adjustments" do not even try to apportion the value of the claimed inventions. And the licenses that form the basis of its analysis are radically dissimilar from the hypothetical licenses into which Apple and Wi-LAN would have entered.

This Court should reverse the judgment below or, at a minimum, remand for a new trial on damages.

JURISDICTIONAL STATEMENT

The district court had jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a) and entered final judgment resolving all claims on June 16, 2020, Appx64-66. Apple filed its notice of appeal on July 15, 2020, Appx800-801; *see also* Fed. R. App. P. 4(a)(1)(A). This Court has jurisdiction pursuant to 28 U.S.C. § 1295(a)(1).

STATEMENT OF THE ISSUES

1. Are claims to a "subscriber unit" limited to a three-component system where the patents exclusively describe the invention as a three-component system and exclusively describe the benefits of the invention as arising from adding a third component?

2. Can a product directly infringe a wireless communications patent when the product does not meet the claim limitations absent further alteration by the device user and a network's base station?

3. Must a damages award be vacated where a damages expert simply relies on a generic and arbitrary assumption that any subset of a thousand-patent wireless portfolio (no matter how big or small the subset) has value equal to 75% of the whole?

4. Must a damages award be vacated where a damages expert relies on a "comparable" license analysis that ignores licenses negotiated by companies of similar size and market position to the defendant (all of which take the form of a lump-sum payment) and relies instead exclusively on per-unit licenses negotiated by small, niche-market players?

STATEMENT OF THE CASE

Apple's iPhones Connect Users Across The World

Although now easy to take for granted, Apple's iPhone is one of the most revolutionary inventions of the twenty-first century. As *Time Magazine* puts it, the iPhone "changed the world." Lisa Eadicicco, *This Is Why the iPhone Upended the Tech Industry*, Time (June 29, 2017), https://tinyurl.com/y3hgtr4o. Today, there are more than a billion iPhone users worldwide, with Apple selling an additional 200 million iPhones each year. Neil Cybart, *A Billion iPhone Users*, Above Avalon (Oct. 26, 2020), https://tinyurl.com/y3yas6qe.

This case involves two chips in every Apple iPhone: the applications processor (indicated below in blue with an Apple logo) and the baseband processor (in purple, labeled Qualcomm):

Appx818. The applications processor is the "brains of the iPhone." Appx10917. It runs the phone's applications, like voice, text, iMessage, and FaceTime. The baseband processor wirelessly sends and receives data from the base station. Appx10942.

When an Apple customer receives a new iPhone (or cellular-capable iPad), the phone comes pre-loaded with a range of innovative features, including powerful computing chips, Retina displays, high-definition cameras, Apple Pay, and a health-monitoring application. *E.g.*, Appx11068-11069. But one key thing is missing from the iPhones as sold: they have not yet established a connection to a wireless network.

Instead, when the user decides to connect the phone to his or her network of choice, the phone makes an initial connection with and "receives a set of parameter values from the base station." Appx10258. The base station provides the phone with the instructions and templates necessary for the phone to establish its connections to the network and transmit data. *See* Appx10258; Appx10948; Appx11018-11019; Appx11039. Moreover, the networks themselves determine what kinds of connections the phone will be able to make. For example, Sprint historically did not offer the service known as Voice over LTE (VoLTE)—a technology for making phone calls over the internet—and therefore did not provide the parameters necessary for a new phone to establish a VoLTE connection. *See* Appx10413-10415; Appx10952-10953. Once established based on the network-provided parameters, these network connections allow data to flow to and from the base station and then to and from the applications processor running inside the iPhones.

Wi-LAN Asserts Wireless Communication Patents Against Apple

Wi-LAN, a former Canadian device-manufacturing company, now centers its business around patent licensing and litigation efforts. *See, e.g.*, Appx10454-10457; Appx10487-10492. In particular, Wi-LAN licenses and asserts a portfolio of patents related to wireless communications. The portfolio includes roughly a thousand patents, which purportedly cover a range of technologies including WiFi, Bluetooth, CDMA, LTE, and HSPA (a technology that improves the performance of the 3G network). *See* Appx15203.

In 2014, Wi-LAN started experiencing financial trouble. According to its CEO, it had "suffered from recent significant setbacks," including "multiple litigation losses" and an "inability to close [a] sufficient number of significant de[als]." Appx15191-15192.

Also in 2014 and while already engaged in a patent-infringement suit with Apple, Wi-LAN served Apple a notice of infringement as to, among others, the two patents at issue in this appeal, U.S. Patent No. 8,457,145 ("'145 patent") and U.S. Patent No. 8,537,757 ("'757 patent"). Wi-LAN had purchased the patents from a small wireless company, Ensemble Communications. Appx857-859. Both the '145 and '757 patents purportedly teach improvements to wireless communications systems using similar terms. In particular, both patents claim a component of a wireless communications system called a "subscriber unit" or "subscriber station." Appx323-324 (32:32, 34:13, 34:26); Appx342 (17:46).[1]

The '145 patent. The '145 patent is directed to improvements in allocating bandwidth in a wireless communications system. Such a system typically provides for wireless transmission of data between a base station (such as a cell tower) and user devices (such as flip phones). *See* Appx308 (1:28-30); Appx311 (7:41-53). In the prior art, a major constraint on wireless systems was scarce radio bandwidth between the base station and the user devices. *See, e.g.*, Appx308 (2:40-60). In some settings, meta-communication (i.e., instructions) between the base station and user devices about how to allocate that scarce bandwidth was taking up so much bandwidth that it was encroaching

---

[1] The specifications of the asserted patents use another term, "customer premises equipment," to refer to the claimed subscriber units. The relationship between these terms is discussed below. *See infra* pp. 29-32.

on what was available for actual, substantive data transmission. *See* Appx308 (2:54-60).

The ’145 patent proposes adding a third component—the subscriber unit—to the wireless communications system in order to perform the task of allocating bandwidth. As the specification describes, the system consists of three components: (1) a base station, which links the system to the wider network, like the Internet; (2) a “subscriber unit,” which communicates with the base station; and (3) a user device, which transmits data to and from the base station via the subscriber unit. *See* Appx308 (1:32-54, 2:15-19). The claimed inventive subscriber unit performs the bandwidth-allocation function itself, using a protocol intended to reduce the strain on the system’s bandwidth. *See, e.g.*, Appx310-311 (6:6-39, 7:39-53).

Figure 1 of the ’145 patent, reproduced below, provides an exemplary communications system, depicting all three components. *See* Appx308 (2:4-19). Each set of buildings has a base station [106] that communicates with the network. The subscriber unit (interchangeably called the subscriber station) [110] is “positioned at fixed customer sites,” here on top of each customer’s building [112]. Appx308 (2:13-14).

The subscriber unit requests bandwidth from the base station based on the demands of the users inside each customer building [112], their devices, and their services. *See* Appx308 (2:20-25). “The users of the system,” who are not depicted within their buildings in Figure 1, “may include both residential and business customers.” Appx308 (2:15-16).

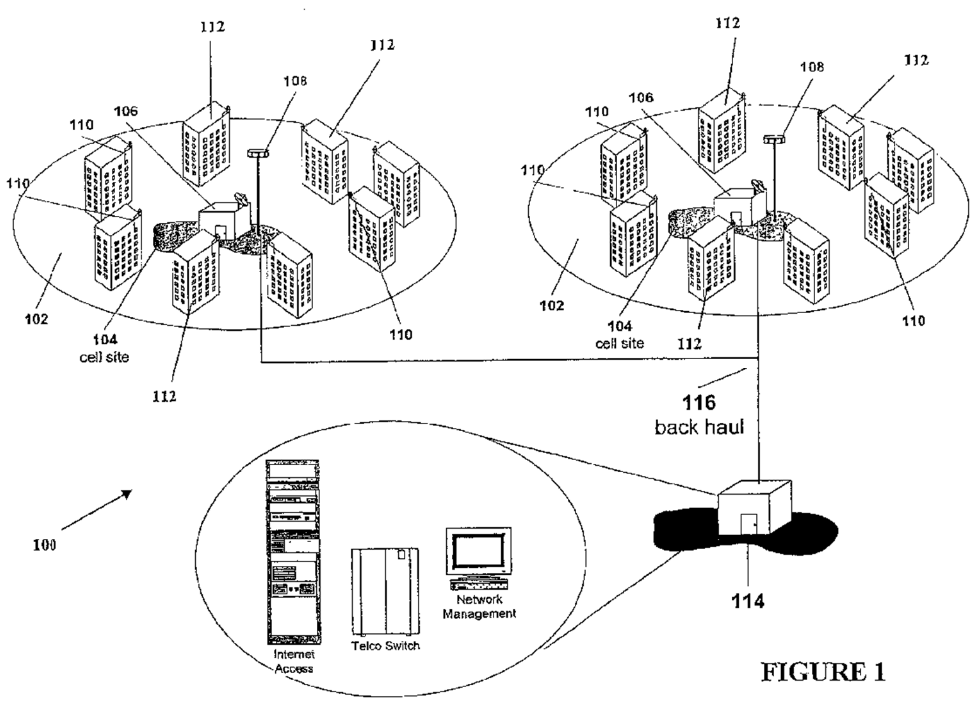

FIGURE 1

Appx291.

The ’145 patent claims three efficiency advantages from using its bandwidth-allocation technology in this configuration. According to the specification, these advantages are:

1. "One advantage of having the [subscriber unit] determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task." Appx311 (7:46-48); *accord* Appx314 (14:8-10). In other words, the base station's computer power is no longer tied up with this task.

2. "In addition, the communication overhead that is required by having the base station instruct the [subscriber unit] how to distribute its allocated bandwidth is eliminated." Appx311 (7:48-51); *accord* Appx314 (14:10-13). Instead of tying up wireless bandwidth with allocation instructions, that bandwidth is free to be used for substantive data transmission.

3. "The [subscriber unit] is free to use the uplink bandwidth that was allocated to it in a manner that is different than that originally requested or granted by the base station." Appx311 (7:41-43). As a result, the subscriber unit can change the allocation dynamically, in real-time response to the users' demands. "[T]he [subscriber unit] is in a much better position to respond to the[se] varying uplink bandwidth allocation needs…." Appx314 (14:14-16).

Three claims of the ’145 patent are at issue here: claims 9, 26, and 27. Claim 9, which is representative, is directed to a “subscriber unit” that sets a timer for each of its data queues and, “periodically,” “transmit[s] a bandwidth request indicating an amount of bandwidth required for transmitting the data from the queue.” Appx323 (32:32-44).

The ’757 patent. The ’757 patent deals with the same kind of wireless communications system, including a base station, subscriber unit, and user devices. *See* Appx334 (1:65-2:20). In such a system, “buildings or other physical obstructions (such as trees or hills, for example), … may cause channel interference” between the base station and the subscriber unit. Appx335 (3:40-42). “[E]nvironmental conditions,” including the distance “between some [subscriber units] … and the base station” may also affect the wireless communications between the two. Appx336 (5:17-19). The same is true for “a rain or snow storm.” Appx336 (5:38). When faced with such interference, typically a “less robust” transmission mode must be used, meaning less data can be transmitted at any moment. Appx336 (5:19). As a result, the subscriber unit must manage its competing needs for data-transmission volume and signal quality. *See* Appx335 (3:26-62).

Claim 1 of the '757 patent is asserted here. It claims a "subscriber station" that "monitor[s]" the quality of data transmissions from the base station and can respond by "determin[ing] a preferred … mode for the downlink data" among the various available transmission modes. Appx342 (17:46-18:5).

The District Court And This Court Reject Wi-LAN's Overbroad Claim Constructions Of Similar Wireless Communication Patents

One of the "litigation losses" Wi-LAN encountered was a suit against Apple for infringement of U.S. Patent Nos. 8,311,040 and 8,315,640. *See Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1376-77 (Fed. Cir. 2016). There, the district court's constructions resulted in summary judgment of non-infringement in Apple's favor. *Id.* at 1376.

Wi-LAN appealed, urging broad constructions of the "connection" terms in its claims. For example, one of the asserted patents claimed a "node for a communications system" that would "allocate bandwidth for the specified connection." *Id.* at 1377. Wi-LAN urged a broad construction encompassing a node that allocates bandwidth to one or many "specified connections." *Id.* at 1381. Emphasizing that the specified-connection term was used in the singular in the claims, this Court nonetheless rejected Wi-LAN's construction. *Id.* at 1381-82. Its reasoning was that "[t]he specification's consistent descriptions of *multiple* specified connections … suggest that the patent's claims do not encompass an embodiment contrary to these descriptions," i.e., with only one specified connection. *Id.* at 1382 (emphasis added). "Because [it] credit[ed] the specification's consistent descriptions," this Court rejected Wi-LAN's attempts to stretch the claims broader than the specification would allow. *Id.* at 1384.

Wi-LAN Tries Again

Less than two weeks after claim construction briefing concluded in the "connections" litigation, Wi-LAN sent Apple the notice of infringement referenced above, asserting the two patents at issue in this appeal. Apple promptly sought declaratory judgment of noninfringement and invalidity; Wi-LAN answered and asserted counterclaims of infringement against a wide range of iPhones and iPads. *See* Appx278-287; Appx344-379.

As relevant to this appeal, the parties' claim-construction dispute turned on the meaning of the "subscriber unit" term. Apple explained that the subscriber unit must be a distinct intermediate device, situated between the base station and the user devices. Appx446-449. Apple relied on the specifications' consistent descriptions of the claimed subscriber unit. Apple observed that the specifications use the subscriber-unit term interchangeably with another term—customer premises equipment ("CPE")—which, the parties agreed, is a device distinct from the user's device. *See* Appx393-394; Appx517; Appx522. And the specifications consistently describe the invention and its purported benefits in terms of CPE, i.e., fixed or portable (but not mobile) equipment situated between the base station and user devices. *See* Appx308 (1:38-39, 2:3, 2:13-14); Appx314 (14:8-18); Appx334 (1:28-29); Appx337 (7:46-53). Wi-LAN, by contrast, maintained that the subscriber unit need not be a distinct device but instead could be a "subcomponent" of, i.e., contained within, the user's device. Appx395-399. Ignoring the specifications' consistent, interchangeable usage of the terms "CPE" and "subscriber unit," the district court adopted Wi-LAN's construction. Appx6.

A Jury Awards Wi-LAN $145 Million, And The District Court Orders A New Damages Trial

Ultimately, four claims went to the jury: claims 9, 26, and 27 of the '145 patent and claim 1 of the '757 patent. *See* Appx629.[2] The district court held a seven-day jury trial. Appx47. The jury returned a verdict for Wi-LAN of $145 million. Appx629-630.

Apple moved for judgment as a matter of law, renewing its claim-construction arguments and arguing that Wi-LAN had failed to prove infringement even under the district court's constructions. In particular, the district court's construction of the subscriber-unit term required the unit to allocate bandwidth across multiple "user connections." Appx6-7. Wi-LAN's infringement theory was that Apple's devices did so by allocating bandwidth across two connections: data and VoLTE. But, as Apple explained in its motion, neither a data connection nor a VoLTE connection exists in its devices when they are

---

[2] Prior to trial, the district court granted Apple's motion for summary judgment of no infringement by Apple devices containing an Intel chipset based on a license to the patents-in-suit that Intel took after it had been sued by Wi-LAN. Appx23-25. Then, on the eve of trial, Wi-LAN voluntarily dismissed four patents from the lawsuit. "Despite Wi-LAN's decision to proceed on only those claims, its damages numbers remained the same." Appx46 n.2.

sold to customers. Appx646. Instead, both the user and the base station must make significant alterations to the device to create those connections. The district court denied Apple's motion without substantive discussion. Appx30-31.

The court then granted Apple's motion for a new trial on damages. David Kennedy, Wi-LAN's damages expert, had offered a novel "direct valuation" methodology which no court had blessed. Appx34. Wi-LAN claimed that this "direct valuation" methodology satisfied its obligation to apportion damages—i.e., to ensure that its damages compensated it for "only … the infringing features" of the accused products. Appx32-34. But the district court found otherwise. Appx34. Notably, Wi-LAN's damages model equated the claimed inventions with VoLTE—a technology Wi-LAN undisputedly did not invent. *See* Appx15076 ("Q. Did you invent VoLTE? A. No."); *see also* Appx48. This approach "overstated the footprint of the invention," resulting in a grossly inflated and unapportioned damages calculation that lacked "factual basis." Appx35-38.

The court offered Wi-LAN a choice between a second damages trial or a remittitur of $10 million. Appx38. Wi-LAN rejected the offer of remittitur, preferring a second trial. *See* Appx38.

A Jury Awards Wi-LAN $85 Million, And The District Court Declines To Order Another Damages Trial

At the damages retrial, Kennedy opined that Wi-LAN was entitled to $85 million in damages. Appx15294. In an attempt to support this conclusion, Kennedy used, as a starting point, portfolio licenses that Wi-LAN had entered into with three small, niche players in the cellphone manufacturing business ("Small Niche Co." licenses): Unnecto, Doro, and Vertu.

- Unnecto "target[ed] frugal shoppers with budget smartphones." Appx15323. It went out of business in 2014 and sold not a single licensed phone. Appx15587.
- Doro is based in Sweden and focuses primarily on the European market. Appx15327. It targets senior citizens with low-tech, affordable phones, such as its $50 flip phone. Appx15325-15326. It paid approximately $482,000 in royalties to Wi-LAN for its U.S. sales. Appx15587.

- Vertu sells ultra-luxury phones for tens of thousands of dollars, including models plated in gold and outfitted in stingray leather. Appx15327-15328; Appx15585. It has been "in and out of financial trouble," and undergone multiple "reorganizations." Appx15166. It paid $13,000 in royalties to Wi-LAN. Appx15585.

The Small Niche Co. portfolio agreements each included roughly a thousand of Wi-LAN's patents, including patents that purportedly practice important technologies not at issue here, such as CDMA, WiFi, Bluetooth, and HSPA. Appx15203; Appx15331-15333. Each license ostensibly called for a running royalty of $0.50 per phone. Appx15163-15167.

Kennedy applied a series of "adjustments" to that $0.50-per-phone figure, supposedly to account for the substantial differences between these portfolio licenses and the "hypothetical license" into which Wi-LAN and Apple would have entered for just the two patents-in-suit. Appx15229-15231; Appx15239-15245. He first assumed that "75 percent of the value of the patents" contained in the Small Niche Co. licenses could be "assign[ed] … to the two patents in this case."

Appx15335. He then increased that figure through additional "adjustments." All told, Kennedy concluded that Apple would have paid $0.45 per phone—90% of the Small Niche Co. royalty rate for a thousand-patent portfolio—for a license to just the '145 and '757 patents. Appx15216. To "make the math easier" for the jury, Kennedy clarified that multiplying the $0.45 royalty by the number of infringing iPhones resulted in a total damages figure of $85 million. Appx15294.

Apple filed a *Daubert* motion, arguing that Kennedy's latest damages model shared the same critical flaw as his "direct valuation" model: it inflated the damages calculation by failing to apportion the value of the claimed inventions. Appx704. This time, Kennedy had based his damages calculation on the Small Niche Co. licenses, without accounting for the important differences between those licenses and the hypothetical two-patent license between Wi-LAN and Apple. The district court denied that motion orally, explaining that in its view, these concerns "go to weight" not admissibility. Appx754. The court did not directly respond to Apple's argument that Kennedy's model is legally defective because it "is untethered to the facts of this case, [and] to the value of [the '145 and '757] patents." Appx753.

While the jury was considering its second damages verdict, the court remarked that the trial proceedings, regardless of the outcome, were little more than a "speed bump on the way to the Federal Circuit." Appx15731. The following day, the jury found for Wi-LAN in the amount of $85 million—precisely what Kennedy had advocated. Appx756.

Apple filed post-trial motions, renewing the arguments it raised in its *Daubert* motion, and also arguing that the Small Niche Co. licenses were radically dissimilar to the hypothetical negotiation and so could not sustain the jury's verdict. Appx763; Appx765.

The court did not address the substance of Apple's arguments. Instead, it summarily rejected Apple's motion, noting only that "Apple raise[d] many of the same arguments on damages it ha[d] raised throughout th[e] case," and that the Court would not "rehash" its views on them. Appx62.

SUMMARY OF ARGUMENT

1. The asserted patents' specifications consistently describe a three-component system, in which the claimed inventive subscriber unit is an intermediary between the base station and the user devices. Yet

under Wi-LAN's constructions, the patents also cover a two-component system in which the subscriber unit can be a "subcomponent" of the user's device. Wi-LAN's constructions ignore the context of the patents and invention, are incompatible with the patents' consistent descriptions of the invention and its benefits, and are unsupported by the prosecution history. *Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1381-84 (Fed. Cir. 2016); *On Demand Mach. Corp. v. Ingram Indus., Inc.*, 442 F.3d 1331, 1339-40 (Fed. Cir. 2006).

2. Wi-LAN failed to offer evidence sufficient to establish direct infringement by Apple. Both sides' experts agreed that the accused products, as sold, require additional configuration before they can make VoLTE and data connections. Because those connections are required by the district court's claim constructions, there is no infringement without them. Accordingly, Apple's mere sale of these products is not an act of direct infringement. *ParkerVision, Inc. v. Qualcomm Inc.*, 903 F.3d 1354, 1361-62 (Fed. Cir. 2018); *Finjan, Inc. v. Secure Computing Corp.*, 626 F.3d 1197, 1203-05 (Fed. Cir. 2010).

3. Wi-LAN's damages model failed to meet this Court's standards. Wi-LAN did not attempt to apportion the value actually created by the

particular patents-in-issue, as 35 U.S.C. § 284 and a century-and-a-half of precedent require, *e.g.*, *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1312-18 (Fed. Cir. 2011); *VirnetX, Inc. v. Cisco Sys., Inc.*, 767 F.3d 1308, 1331-34 (Fed. Cir. 2014). Rather, Wi-LAN's expert simply assigned those two patents 75% of the value of Wi-LAN's entire thousand-patent wireless portfolio—a figure that was inflated to 90% through additional "adjustments." Indeed, he assumed that *any* subset of that portfolio—no matter how big or small, or what technologies it included—would be worth that same amount. The law demands more than that kind of generic and abstract heuristic.

Moreover, the prior licenses that form the basis of the analysis are radically dissimilar to the hypothetical license between Wi-LAN and Apple. The evidence showed that a large company like Apple would have obtained a lump-sum royalty from Wi-LAN. Yet Wi-LAN's analysis focuses exclusively on running royalty rate agreements with small cellphone providers. Because these running royalty rate agreements bear no resemblance to the hypothetical lump-sum license, they too do not support the damages verdict. *LaserDynamics, Inc. v. Quanta Comput., Inc.*, 694 F.3d 51, 80-81 (Fed. Cir. 2012); *Wordtech*

Sys., Inc v. Integrated Networks Sols., Inc. 609 F.3d 1308, 1320-21 (Fed. Cir. 2010).

ARGUMENT

I. Wi-LAN's Overbroad Claim Construction Expands The Claims Beyond What Was Invented

Over the course of two lawsuits and a half-dozen patents, the parties to this appeal have litigated variations on a single, central question. Do the claimed subscriber units operate only in three-component systems—the systems consistently and exclusively described by and embodied in the specifications? Or do the patents also encompass subscriber units in two-component systems, the systems set out in the prior art? [3]

The relevant claims and specifications can be read in one common-sense, consistent manner to describe marginal alleged advancements in

[3] The claim construction issues in this appeal were decided entirely on the intrinsic record, namely, the patents and the prosecution history. In such cases, the district court's ruling was "solely to a determination of law, and the Court of Appeals will review that construction de novo." *Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1381 (Fed. Cir. 2016) (quoting *Teva Pharm. USA, Inc. v. Sandoz, Inc.*, 574 U.S. 318, 331 (2015)).

wireless technology by introducing a third component into wireless communication systems.

But Wi-LAN advances the two-component position. Its proposed construction of the subscriber-unit term is any "module that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections." Appx393. The district court agreed, rejecting Apple's construction that requires the subscriber unit to be a standalone component between the base station and the user devices. With its construction, Wi-LAN persuaded the district court to construe its claims in a way that is "incompatible" with the specifications' "consistent descriptions" of the inventions. *Wi-LAN*, 830 F.3d at 1384, 1389. This Court should reject Wi-LAN's construction, just as it rejected Wi-LAN's similarly impermissible constructions in the prior case. *See id.*

A. The patents describe a three-component wireless communications system that includes a "base-station" unit, a "subscriber unit," and a "user" device.

The parties propose competing constructions of the subscriber-unit term in each of the claims. *See, e.g.*, Appx323 (32:32) (Claim 9: "A subscriber unit for a wireless communication system…."). The dispute turns on whether the subscriber unit can be a subcomponent of the user's device.

Apple contends that a subscriber unit is "fixed or portable customer premises equipment that wirelessly receives UL bandwidth from a base station, and allocates the bandwidth across connected user devices." Appx446. In other words, the subscriber unit must be a standalone device ("fixed or portable customer premises equipment") between the base station and the user devices.

Wi-LAN proposes that a subscriber unit is a "module that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections." Appx393. Wi-LAN's construction allows it to contend that the baseband processor—an internal subcomponent of the iPhone—is the infringing subscriber unit.

The patents dictate the three-component construction. The asserted patents' specifications and figures describe a wireless communications system with three key components: the base station, the subscriber unit, and user devices. First, the "base stations," which provide the system's connection "with the fixed network infrastructure," like the phone lines and the Internet. Appx308 (1:45-47). Second, the "subscriber unit," described as equipment "positioned at fixed customer sites." Appx308 (2:13-14). Third, "[t]he users of the system," who "may include both residential and business customers," and who use "bandwidth requesting device[s]" to transmit data. Appx308-309 (2:15-16, 3:11). The customers and their devices are "served by the [subscriber unit]," from which they "request[]" "broadband services [of] different bandwidth and latency requirements." Appx308 (2:24-26). The subscriber unit, in turn, "request[s] bandwidth allocations" from the base station. Appx308 (2:22-23).

In familiar cellular telephone systems, cellphones communicate directly with a base station. Bandwidth allocation must be performed by the base station across its connected phones. *See* Appx311 (7:46-48). But as explained above (pp. 10-12), that setup yields three disadvantages: (1) the burden imposed on base stations to perform bandwidth allocation for all the user devices it served, (2) the bandwidth usage required to communicate those allocations to the user devices, and (3) the inability to adapt those base-station-derived allocations to real-time, local changes in the user devices' bandwidth demands. *See* Appx311 (7:46-53); Appx314 (14:8-18).

By positioning an intermediate device (the subscriber unit) between the base station and the user devices, the asserted patents introduce a claimed efficiency. The subscriber unit is able to allocate bandwidth locally and dynamically, purportedly solving for those prior disadvantages. *See id.*

In a real-world example, suppose the Nationals are concerned that post-pandemic crowds will strain the cell towers in Southeast D.C. on opening day. They hire a local communications company to install "subscriber units"—essentially computers with antennas—around the ballpark. By installing these subscriber units, the Nats can reduce the strain on the cell towers and ensure that bandwidth will be allocated efficiently during a period when an unusually high number of users will be connecting. The subscriber units need not be fixed. For example, a portable version could be temporarily installed during parades or natural disasters to handle surges in network demand.

Taken in context, these patents' specifications describe a three-component system, with purported inventive improvements credited to the subscriber unit, an intermediate device within that system.

B. The patents equate "subscriber unit" with "customer premises equipment" and thereby confirm the patents describe a three-component system.

Within the context of this three-component system, the patents consistently and exclusively describe the subscriber unit as standalone equipment. They do so by equating the claimed subscriber unit with "customer premises equipment," which the parties agree is a standalone device. *See* Appx393-394; Appx517; Appx522.

Originally, all of Ensemble's patent claims used the term CPE rather than the subscriber-unit term. *See, e.g.*, Claim 1, U.S. Patent Application Serial No. 12/414,363 (Mar. 30, 2009). During prosecution of the '145 and '757 patents, and without explanation, the patentee replaced the CPE term with the subscriber-unit term. *See, e.g.*, Preliminary Amendment, U.S. Patent Application Serial No. 12/414,363, at 2 (May 18, 2010) (claiming "[a] method for obtaining uplink (UL) bandwidth at a subscriber station that has a wireless communication link to a base station").

Although the claim terms changed, the specifications still consistently refer to CPE, using that term interchangeably with the subscriber-unit term. In both patents, the subscriber-unit term is entirely absent from the summary of the invention and the detailed description, which exclusively refer to the invention as "CPE." Appx310-323 (6:4-31:63); Appx334-342 (1:54-17:43). By contrast, the claims use "subscriber unit" and "subscriber station" to refer to the same element. Appx323-324 (31:65-34:34); Appx342-343 (17:45-20:44).

Wi-LAN concedes that the patentee used several different terms for subscriber unit interchangeably across its many patents. Appx397 n.9 ("[T]he original independent claims of the '723 patent, claims 1 and 6, used 'wireless subscriber unit' (or 'the subscriber unit'), 'subscriber radio unit,' and 'subscriber radio station' interchangeably, further showing there is no material difference between these terms."). It provides no justification for its insistence that CPE was not *also* used interchangeably with those terms.

In adopting Wi-LAN's construction, the district court reasoned that the patentee's decision to claim subscriber units and subscriber stations in the patents asserted here—as opposed to the CPE term used in their specifications, applications, and prior patents—"counsels against construing the 'subscriber' terms as CPEs." Appx6.

But this Court has consistently rejected the blanket rule that different terms must be construed differently. In particular, "that implication is overcome where, as here, the evidence indicates that the patentee used the two terms interchangeably." *Baran v. Med. Device Techs., Inc.*, 616 F.3d 1309, 1316 (Fed. Cir. 2010). And the district court's practice of defining different terms differently "carries less weight when comparing a term in the claim to a term in the specification, especially where, as here, the specification only describes one embodiment." *SEB S.A. v. Montgomery Ward & Co.*, 594 F.3d 1360, 1369 (Fed. Cir. 2010), *aff'd sub nom. Global-Tech Appliances, Inc. v. SEB S.A.*, 563 U.S. 754 (2011).

In such cases, it is common for different sections of a patent to use different terms to refer to the same thing. *Bancorp Servs., L.L.C. v. Hartford Life Ins.*, 359 F.3d 1367, 1373 (Fed. Cir. 2004); *see also Tehrani v. Hamilton Med., Inc.*, 331 F.3d 1355, 1361 (Fed. Cir. 2003); *Tate Access Floors, Inc. v. Maxcess Techs., Inc.*, 222 F.3d 958, 968 (Fed. Cir. 2000). Where that happens, this Court has explained, there is a "correspondence between the reference … in the specified claims and the reference … in the portion of the specification relating to those claims." *Bancorp*, 359 F.3d at 1373. Such a correspondence "provides substantial support" for construing the terms to mean the same thing, especially where there is "no satisfactory alternative explanation for the apparent correspondence." *Id.*

That is precisely the case here, where the claims use one term (subscriber unit) and "other portions of the patent refer to the same" element with another term (CPE) instead. *Tehrani*, 331 F.3d at 1361. In such a case, "[t]he interchangeable use of the two terms is akin to a definition equating the two." *Edwards Lifesciences LLC v. Cook Inc.*, 582 F.3d 1322, 1329 (Fed. Cir. 2009); *see also Tate*, 222 F.3d at 968. In light of the specifications' consistent description of the subscriber unit as CPE, this Court should find these terms interchangeable and construe them accordingly.

C. The patents consistently characterize the purported invention as customer premises equipment, further confirming the patents describe a three-component system.

In addition to using "subscriber unit" and "customer premises equipment" interchangeably, the patents' characterizations of the purported invention further confirm that the subscriber unit is CPE and, therefore, that the patents describe three-component wireless systems.

The specifications describe the subscriber unit as "fixed and portable" equipment situated at a particular customer site; by contrast, they never suggest the subscriber unit could be contained within a user device. Appx308 (1:38-39) ("subscriber units (fixed and portable)"); *accord* Appx334 (1:28-29); *see also* Appx308 (2:3) ("a plurality of fixed subscriber stations"); Appx308 (2:13-14) ("positioned at fixed customer sites"). In other words, the subscriber unit might be temporarily installed and then relocated and reinstalled, as in the ballpark and parade examples above, meaning it is portable. *See, e.g.*, *Portable*, *Oxford English Dictionary* (3d ed. 2007) [hereinafter *OED*] ("[c]apable of being moved from place to place," "capable of being dismantled and re-erected elsewhere").

But the specifications are devoid of any reference to the subscriber unit being mobile, i.e., used while in motion, as it would be if it could be a subcomponent of a user's mobile device. *See, e.g.*, *Mobile*, *id.* ("capable of or characterized by movement; movable; wandering"). For example, the '145 specification explains that in a normal setup, there is a "fixed distance[]" between the base station and the CPE. Appx323 (31:41-42). In a system where the base station is replaced by a satellite, however, the distance between them would no longer be static, because the base station is now mobile relative to the fixed CPE. Appx323 (31:38-43). In neither case is the *subscriber unit* mobile (i.e., able to work in motion), as opposed to portable (i.e., movable to a new location). That description is consistent with the ordinary meaning of "premises equipment," which indicates a device located at a particular customer site. *See, e.g.*, *Premise*, *OED* ("A house or building together with its grounds, outhouses, etc., *esp.* a building or part of a building that houses a business."). And that description is inconsistent with a subscriber unit being a subcomponent of a user's mobile device.

The specifications also consistently describe the purported benefits of the invention in terms of CPE. They specifically describe the benefits of delegating tasks from the base station to CPE situated between the base station and the user devices. *See* Appx311 (7:46-48) (emphasis added) ("One advantage of having *the CPE* determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task."); *accord* Appx314 (14:8-18).

In touting the purported inventive advantages of the CPE, the specifications indicate that the CPE *is the invention*. "[W]hen the preferred embodiment"—here, the only embodiment, *e.g.*, Appx310 (6:6-23)—"is described in the specification as the invention itself, the claims are not necessarily entitled to a scope broader than that embodiment." *Edwards*, 582 F.3d at 1330.

The specifications, in fact, steer decidedly *away* from Wi-LAN's construction. Here, the patents point to the efficiencies of delegating bandwidth allocation to the CPE. If the inventors had discovered new efficiencies by further delegating that allocation task to the user's device—eliminating the need for an intermediate device at all—a person of ordinary skill in the art would expect that to be the claimed advantage over the prior art. To the contrary, the specifications *disparage* the prior-art systems in which bandwidth allocation must be performed by either the base station or the user device, and they promote the advantages of assigning that step to an intermediate device. *See* Appx311 (7:46-53); Appx314 (14:8-18). Because the specifications "disparage and, therefore, disclaim" these prior-art two-component systems, they make clear that the present invention is situated within a three-component system, as captured in Apple's proposed construction. *Openwave Sys., Inc. v. Apple Inc.*, 808 F.3d 509, 514 (Fed. Cir. 2015).

D. Wi-LAN's construction is incompatible with the patents' consistent descriptions.

Wi-LAN's construction ignores the context and consistent descriptions of the terms in the patents, just as Wi-LAN unsuccessfully urged this Court to do in the prior appeal. *See Wi-LAN*, 830 F.3d at 1382-84. "However, when the scope of the invention is clearly stated in the specification, and is described as the advantage and distinction of the invention, it is not necessary to disavow explicitly a different scope." *On Demand*, 442 F.3d at 1340. Rather, "the written description of the preferred embodiments 'can provide guidance as to the meaning of the claims, thereby dictating the manner in which the claims are to be construed, even if the guidance is not provided in explicit definitional format.'" *Bell Atl. Network Servs., Inc. v. Covad Commc'ns Grp., Inc.*, 262 F.3d 1258, 1268 (Fed. Cir. 2001) (quoting *SciMed Life Sys., Inc. v. Advanced Cardiovascular Sys., Inc.*, 242 F.3d 1337, 1344 (Fed. Cir. 2001)). Most importantly, "the claims cannot be of broader scope than the invention that is set forth in the specification." *On Demand*, 442 F.3d at 1340.

On Demand tracks this case closely. There, the claims merely described a customer, which the patentee argued could be any customer—retail, wholesale, or otherwise. *Id.* at 1339-40. But the specification at issue consistently described a *retail* customer by explaining the customer's interactions with the store clerk and the point of sale. *Id.* at 1340. By "repeatedly reinforc[ing] its usage of the term 'customer' as the retail [customer]," the specification imposed constraints on that claim term. *Id.* This Court reversed the patentee's attempt to "eliminate these constraints in order to embrace the remote large-scale production of books for publishers and retailers," i.e., wholesale customers. *Id.*

The same is true of the subscriber-unit term here. The specifications consistently describe the subscriber unit as CPE. *See, e.g.*, Appx308 (2:4); Appx334 (1:67-2:1). Indeed, the summaries of the inventions and detailed descriptions make no mention of subscriber units at all, in favor of the patentee's preferred CPE term. *See* Appx310-323 (6:6-31:63); Appx334-342 (1:54-17:43). And the specifications tout their inventive advantage as delegating a resource-intensive task from the base station to "the CPE," Appx311 (7:46), which the specifications define in plain English as "Customer Premises Equipment," Appx308 (2:4).

As in *On Demand*, "the scope of the invention" (here, CPE) "is clearly stated in the specification, and is described as the advantage and distinction of the invention." 442 F.3d at 1340. All signs point toward a construction that uses the ordinary meaning of customer premises equipment, which is distinct from a user device. The specifications "repeatedly reinforce[]" this understanding, and Wi-LAN's contrary construction is incompatible with the specifications' consistent descriptions.

None of the intrinsic evidence relied on by the district court supports a departure from this consistent description of CPE. Looking to the prosecution history, the district court cited the Examiner's treatment of the '020 patent, which is related to the '145 patent and was still asserted in this case at the time of claim construction. *See* Appx7 n.3. In that prosecution, the Examiner "read the invention claimed therein on prior art that included a 'cellular telephone network,'" the

Widegren patent. Appx7 n.3. In the district court's view, that debunked Apple's assertion that the subscriber unit (i.e., the CPE) may be fixed or portable, but not mobile. Appx7 n.3.

The district court misread Widegren. For the proposition that the "mobile radio station" in Widegren was in fact a cellphone, Wi-LAN and the district court cited a generic description of the overall communications network in that patent, not a description of the invention itself. *See* Appx424 (5:19-21) ("For example, while the present invention is described in the context of an example cellular telephone network…."). But Widegren describes its "mobile radio station" much the same way Wi-LAN's patents describe CPE: as a literal installation, manned by a "mobile station human operator," which serves user devices that include "computer servers, telephones, videophones, etc." Appx426 (9:38, 9:43-49). The Examiner's citation to Widegren, then, fails to support Wi-LAN's construction, and it is fully consistent with Apple's view that the subscriber unit is CPE.[4]

---

[4] Even if Widegren could be read to claim a mobile phone, the Examiner's action—performed under the PTO's "broadest reasonable interpretation" standard for claim construction—provides little support for Wi-LAN's construction under the "ordinary meaning" standard

For much the same reasons as Widegren, the ’145 and ’757 patents’ prefatory descriptions of the system in which their invention operates cannot salvage Wi-LAN’s construction. Both specifications state that “[e]xemplary communication systems include mobile cellular telephone systems.” Appx308 (1:40-41); Appx334 (1:29-31). Again, this is a general description of the overall system, and no one disputes that the invention can be used in the mobile-telephone setting. The ballpark example above offers one such instance. But the overall system should not be conflated with its component devices. That the system in general can be used for cellular telephone communications tells us nothing about whether CPE, as an individual component of that system, must be a standalone device.

Nor does the ’757 specification provide support for Wi-LAN’s construction when it explains that hardware components like the subscriber station “can include, for example, … [an] Application-Specific Integrated Circuit (ASIC),” i.e., a specialized computer chip. Appx335

---

applicable here. Because “the district court adopted [the Examiner’s] reasoning wholesale without accounting for the differences between the broadest reasonable interpretation standard and *Phillips*, the court erred.” *Convolve, Inc. v. Compaq Comput. Corp.*, 812 F.3d 1313, 1325 (Fed. Cir. 2016).

(3:9-11). No one disputes that the subscriber unit can include a specialized chip (or a computer with such a chip). But this tells us nothing about *where* the subscriber unit is and whether it may be a subcomponent of the user's device. If anything, the specifications weigh in the other direction by describing "software … execut[ing] on processors both in the base stations and the CPE." Appx312 (9:20-21). This passage suggests that the CPE is a distinct device that *contains* a processor *within* it—and that it cannot be a processor itself within a user's device.

* * *

Just as in the prior litigation, Wi-LAN sought a construction that is incompatible with its own patents' specifications. The specifications consistently and exclusively describe a three-component system in which purported efficiencies are derived from inserting an intermediate device—a subscriber unit/CPE—between the base station and the user devices. Wi-LAN's construction is divorced from that context, incompatible with the patents' consistent descriptions of the invention and its benefits, and unsupported by the prosecution history. It should be rejected.

Wi-LAN did not dispute that there is no infringement under Apple's proposed construction. *See* Appx643-644; Appx676-677. Indeed, because the iPhone's baseband processor communicates directly with base stations, rather than through an intermediate subscriber unit, there is no infringement. *See* Appx10251-10254; Appx10732-10733. Apple is therefore entitled to judgment as a matter of law. *E.g.*, *Comcast IP Holdings I LLC v. Sprint Commc'ns Co.*, 850 F.3d 1302, 1309 (Fed. Cir. 2017).

II. Wi-LAN's Theory Of Infringement Lacks Substantial Evidence

Wi-LAN's evidence did not come close to establishing that Apple directly infringes Wi-LAN's patents. To demonstrate infringement under the district court's constructions, Wi-LAN attempted to show that within each user device there is a subscriber unit, i.e., a "module that receives UL bandwidth from a base station, and allocates the bandwidth across its user connections." Appx5-6. But Wi-LAN is asserting its patents against Apple, not against iPhone users. Wi-LAN, therefore, needed to show that the iPhones *as sold by Apple* include multiple "user connections." But all agree the iPhones as sold by Apple do not yet have all the parameters required to establish such connections. The district court provided no reasoning in support of its order sustaining the verdict. Appx30-31.[5]

Under Wi-LAN's theory of infringement, the accused user connections in the iPhone are the data connection and the VoLTE connection. Appx10251; Appx10255. Those two connections (shown below as green and blue tubes) run between the applications processor (shown below in gray) and the baseband processor (orange). Appx11012-11014. Under the district court's claim construction, Wi-LAN contends that the iPhone's baseband processor is the claimed subscriber unit, which allocates bandwidth between these two connections. *See, e.g.*, Appx10251-10253; Appx11260.

[5] Ninth Circuit law provides the standard of review for the district court's denial of judgment as a matter of law. *Godo Kaisha IP Bridge 1 v. TCL Commc'n Tech. Holdings Ltd*, 967 F.3d 1380, 1382 (Fed. Cir. 2020). Such review is "de novo," determining "whether the evidence, construed in the light most favorable to the nonmoving party, permits only one reasonable conclusion, and that conclusion is contrary to that of the jury." *Estate of Diaz v. City of Anaheim*, 840 F.3d 592, 604 (9th Cir. 2016) (internal quotation marks omitted). Infringement is a question of fact, "reviewed for substantial evidence when tried to a jury." *Godo Kaisha*, 967 F.3d at 1383 (internal quotation marks omitted).

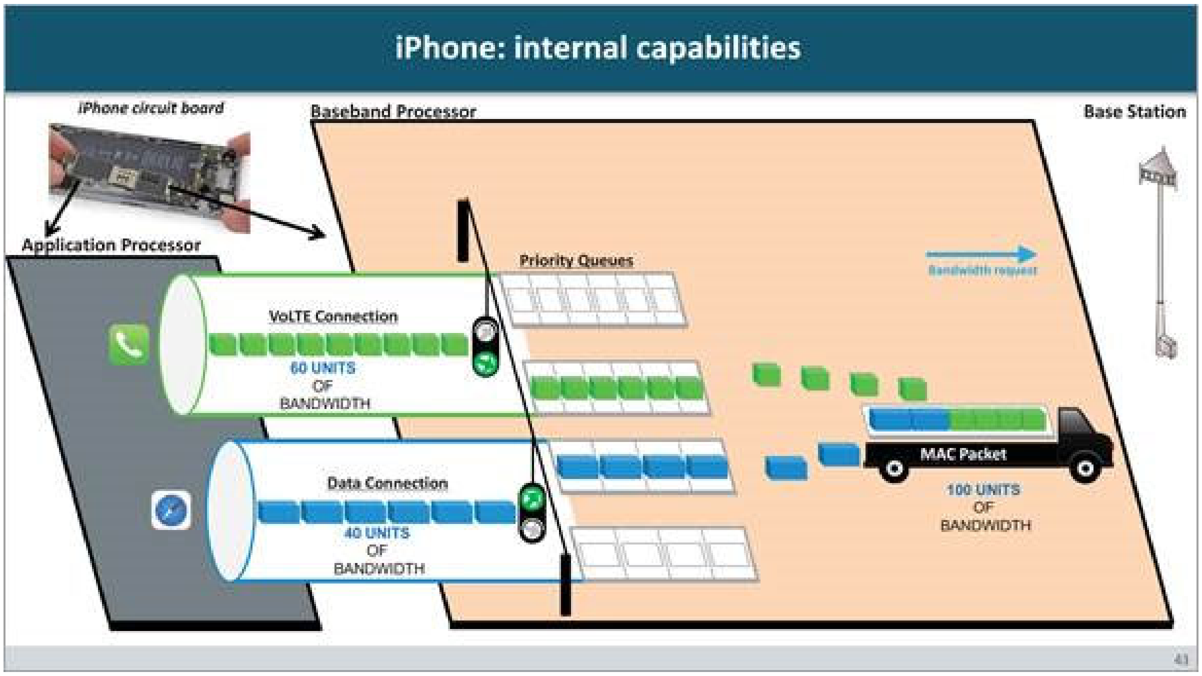

Appx628. As Wi-LAN's expert explained, the above is a "simplified diagram that shows that you have the two different connections, one for the VoLTE voice and [the] other for the browser data." Appx10253. He described data transfer across these connections as "a [Mack] truck that takes the packets and sends them along." Appx10253.

The undisputed evidence makes clear that these data and voice connections do not exist in Apple's devices when they are sold.

That is most obvious with respect to iPhones sold for use on the Sprint network. At the time of the trial, Sprint did not even offer VoLTE service to its iPhone customers. *See* Appx10413-10414; Appx10952-10953. So those phones never established a VoLTE connection at all. Although Wi-LAN's expert opined that Sprint iPhones were theoretically still capable of infringing, he conceded that in order to establish a VoLTE connection, a user would need to "unlock" it and then connect it to some "other network." Appx10415. But that would have required a significant amount of work. As Apple's engineer confirmed, even Sprint could not enable VoLTE on its own; instead it would require Apple "to make a lot of configuration changes" to the phones' software, as well as perform validation and testing of the reconfigured software on the retooled Sprint network. Appx10889; *see also* Appx10890 ("And the software changes you are referring to, do you mean software would need to be added to an iPhone [to activate VoLTE]? Yes.").

And it's not just Sprint phones that lacked the requisite user connections. Even in iPhones on other networks, those connections are not established until the user turns the phone on and connects to the user's network of choice. At that point, according to Wi-LAN's expert, the phone "receives a set of parameter values from the base station." Appx10258. Among other things, those parameters direct "the numbers of the queues that the base station would expect the voice and the data to follow." Appx10258. That is, the network provides parameters that define the number and priority of the queues in which the phone's data packets will line up for transmission. *See* Appx10948; Appx11018-11019. As Apple's expert explained, without this information from the network, a phone "cannot build" the "software structures" necessary to establish a connection. Appx11039.

In short, no voice packets and no data can flow—and there are therefore no "connections"—until the phone receives initial instructions from the network's base station. Apple's engineer confirmed this. *See, e.g.*, Appx10877-10888 (describing configuration parameters transmitted by the network upon a phone's initial set-up); Appx10882 (Q. "Are logical channels on iPhones when Apple sells them?" A. "They don't exist."); Appx10888 (Q. "Do queues exist on the iPhone when it is sold?" A. "No."); *see also* Appx10946-10952; Appx11018-11019. These parameters are network-specific and downloaded by the user after the phone is purchased; they are not preloaded or provided by Apple. *See* Appx10877-10881.

In light of this evidence, Wi-LAN did not genuinely dispute the absence of these connections in the iPhones as sold. Instead, Wi-LAN contended that the iPhones were *reasonably capable* of infringing as sold. *See* Appx11268. This Court's "reasonably capable" doctrine directs that, "when the asserted claims recite capability, our case law supports finding infringement by a 'reasonably capable' accused device on a case-by-case basis." *Ericsson, Inc. v. D-Link Sys., Inc.*, 773 F.3d 1201, 1217 (Fed. Cir. 2014). This doctrine applies *only* to functional claim limitations, i.e., those that are "drawn to capability." *Id.* at 1216; *accord Ball Aerosol & Specialty Container, Inc. v. Ltd. Brands, Inc.*, 555 F.3d 984, 994 (Fed. Cir. 2009).

But the term at issue here—"user connections"—is structural rather than functional. Reliance on the "reasonably capable" doctrine is therefore "misplaced, since that line of cases is relevant only to claim language that specifies that the claim is drawn to capability." *Ball*, 555 F.3d at 994. By contrast, when the limitation at issue is a structure, what matters is whether the structure is present. *E.g.*, *Cross Med. Prods., Inc. v. Medtronic Sofamor Danek, Inc.*, 424 F.3d 1293, 1311 (Fed. Cir. 2005) ("Here, the claim does not require that the interface be merely 'capable' of contacting bone; the claim has a structural limitation…."); *ViaTech Techs. Inc. v. Microsoft Corp.*, 733 F. App'x 542, 551 (Fed. Cir. 2018) (non-precedential) ("The claim language requires the presence of a database … and it is undisputed that there is no database….").

As sold, the iPhones lack the claimed user connections until the user has turned on the phone and acquired the network-provided templates and parameters necessary to establish those connections. With respect to Sprint iPhones in particular, no VoLTE connection was ever implemented, much less existing at the time of sale. Therefore, the iPhones do not infringe.

Because Wi-LAN did not—indeed, could not—close this evidentiary gap, the evidence compels the conclusion that there is no direct infringement here. The judgment below should be reversed. In the alternative, the judgment should be vacated and the case remanded for a recalculation of damages to exclude iPhones sold for use on the Sprint network.[6]

---

[6] The jury found infringement as to 189.4 million accused devices. Appx756. Sprint phones represented 27.1 million of that total, roughly 14 percent.

III. Wi-LAN's Damages Model Does Not Identify The Value Of The Claimed Inventions

The jury awarded Wi-LAN $85 million—exactly the amount to which its damages expert, Kennedy, testified that Wi-LAN was entitled. But Kennedy's opinion rested on a generic and arbitrary heuristic instead of an actual valuation of the patents-in-issue, and improperly relied on supposedly "comparable" licenses that looked nothing like the one Wi-LAN would have reached with Apple. The jury should not have been permitted to considered it.[7]

Under 35 U.S.C. § 284, damages are to be "adequate to compensate for the infringement," i.e., "for the use made of the invention by the infringer." The "invention," for "purposes of assessing damages under § 284," is "only the patented technology." *Ericsson*, 773

---

[7] A district court's rulings concerning methodologies for calculating damages in a patent case are governed by Federal Circuit law and reviewed for abuse of discretion. *Lucent Techs., Inc. v. Gateway, Inc.*, 580 F.3d 1301, 1310 (Fed. Cir. 2009). The "denial of a motion for judgment as a matter of law," and the denial of a motion for a new trial, are "procedural issue[s] not unique to patent law, reviewed under the law of the regional circuit." *Id.* at 1309. In the Ninth Circuit, denial of a motion for a new trial is reviewed for abuse of discretion. *Smith v. City & County of Honolulu*, 887 F.3d 944, 949 (9th Cir. 2018). Denial of a JMOL motion is reviewed de novo. *Hangarter v. Provident Life & Accident Ins.*, 373 F.3d 998, 1005 (9th Cir. 2004).

F.3d at 1226. Thus, the patentholder may recover only "the incremental value that the patented invention adds to the end product." *Id.* That principle—called "apportionment"—is the governing rule "in every case." *Garretson v. Clark*, 111 U.S. 120, 121 (1884).

Kennedy's damages model violated that basic principle. Kennedy opined that patents-in-suit would garner a royalty rate equal to 90% of that for the Small Niche Co. licenses—licenses that included a thousand patents relating to a range of important wireless technologies like CDMA, Bluetooth, WiFi, and HSPA; and which, unlike the lump-sum licenses that Wi-LAN negotiated with large providers, featured a per-phone royalty structure, typical for small providers. For two independent reasons, Kennedy's model fails to ensure that Wi-LAN recovers only the value of the patents-in-issue.

First, Kennedy's model is premised on a "portfolio subset discount"—an assumption that *any* subset of Wi-LAN's thousand-patent wireless portfolio is worth 75% the value of the whole. But that does nothing to disaggregate the value of the asserted patents from the thousand other portfolio patents, pertaining to other valuable (and well-known) industry standards.

Second, Kennedy's model assumes that the license Apple would have taken from Wi-LAN is comparable to the Small Niche Co. licenses, all of which involved a per-device royalty structure. But the evidence showed that a license between Wi-LAN and Apple would have involved a lump-sum payment (just like the licenses between Wi-LAN and other major companies), a fundamentally different licensing structure.

"Because [Kennedy's] proposed royalty rate lacked sufficient ties to the facts of the case," the district court "erred by not excluding [his] opinion." *Exmark Mfg. Co. v. Briggs & Stratton Power Prods. Grp., LLC*, 879 F.3d 1332, 1351 (Fed. Cir. 2018). And because, without Kennedy's opinion, there is nothing to support the jury's damages verdict, the district court "abused its discretion by denying [Apple] a new trial on damages." *Id.* Accordingly, this Court should "vacate the damages award and remand for a new trial on damages." *Id.*[8]

[8] Below, Apple also demonstrated that Kennedy's damages testimony was inadmissible because it failed to begin its analysis with the smallest salable patent-practicing unit. Appx770-772. To the extent this Court's precedent has rejected a rule "which would require all damages models to begin with the smallest salable patent-practicing unit," *Commonwealth Sci. & Indus. Rsch. Organisation v. Cisco Sys., Inc.*, 809 F.3d 1295, 1303 (Fed. Cir. 2015) (*CSIRO*), Apple contends that precedent is wrong. Apple notes this argument here to preserve it for any further review.

A. Wi-LAN's "portfolio subset discount" leads to a damages award that does not reflect the value of the claimed inventions.

The apportionment requirement reflects the commonsense idea that when a patent relates to only one feature of a complex product, any infringement damages owed to the patent-holder must reflect the value of only the patented feature. That is because the patent system represents "a carefully crafted bargain for encouraging the creation and disclosure of new, useful, and nonobvious advances in technology and design," *Bonito Boats, Inc. v. Thunder Craft Boats, Inc.*, 489 U.S. 141, 150-51 (1989), and that "carefully crafted bargain" is upset when patentees receive windfall damages out of proportion to the contribution of their patented invention. Thus, for nearly 150 years, courts have been clear that apportionment is "[t]he true rule" of patent damages. *Dobson v. Hartford Carpet Co.*, 114 U.S. 439, 445 (1885); *see also Mentor Graphics Corp. v. EVE-USA, Inc.*, 851 F.3d 1275, 1283 n.3 (Fed. Cir. 2017) ("[T]he basic principle of apportionment" continues to "appl[y] in all of patent damages.").

An uncontroversial corollary of that rule is that where a plaintiff holds multiple patents, it may recover damages only for those that the accused product actually infringes. After all, "without infringement" there can be "no damages." *Wiener v. NEC Elecs., Inc.*, 102 F.3d 534, 542 (Fed. Cir. 1996). Accordingly, where a damages model is based on prior licenses for a portfolio of patents—some of which the accused product infringes, and some of which it does not—the same "apportionment principles" apply. *Ericsson*, 773 F.3d at 1228. The model must separate out the portion of the prior licensing fees attributable to the patents infringed by the accused products from the portion attributable to the other patents in the portfolio. *Id.*; *see also VirnetX*, 767 F.3d at 1326 ("No matter what the form of the royalty, a patentee must take care to seek only those damages attributable to the infringing features."). To do otherwise would be to unfairly enrich the plaintiff (at the defendant's expense) for patents that were not infringed. Expert testimony that fails to apply these apportionment principles is inadmissible and cannot support a damages award. *E.g.*, *CSIRO*, 809 F.3d at 1301.[9]

[9] *See also, e.g.*, *Sherwin-Williams Co. v. PPG Indus., Inc.*, 2020 WL 5077547, at *5 (W.D. Pa. Aug. 27, 2020) (excluding testimony of expert who "did not perform a quantitative analysis to" arrive at conclusion that the patented features "constituted 50% of the [product's] market

In many apportionment cases, the damages dispute is about the methodology by which one side attempts to implement that principle and discern the value of the infringed patents. When must the damages analysis focus on just the smallest saleable unit that practices the patented invention, and when may damages be based on the entire product in which the invention is incorporated? *E.g., Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.*, 904 F.3d 965, 977-78 (Fed. Cir. 2018), *cert. denied*, 139 S. Ct. 1265 (2019); *Lucent*, 580 F.3d at 1336. If a patent-holder starts its analysis by focusing on more than just the patented features, can it compensate for that by taking a reduction in its royalty rate? *Exmark*, 879 F.3d at 1348. Must a damages expert directly assess the value of the patented technology himself, or may the expert rely on others who have attempted to do so? *Elbit Sys. Land & C4I Ltd. v. Hughes Network Sys., LLC*, 927 F.3d 1292, 1301 (Fed. Cir. 2019); *CSIRO*, 809 F.3d at 1301.

---

value"); *Guardant Health, Inc. v. Found. Med., Inc.*, 2020 WL 2461551, at *18 (D. Del. Apr. 7, 2020) (emphasis omitted) (excluding 50% apportionment figure because expert failed to "provide any factual foundation to support th[at] specific 50% figure"); *NetFuel, Inc. v. Cisco Sys. Inc.*, 2020 WL 1274985, at *7 (N.D. Cal. Mar. 10, 2020) (excluding expert testimony with a "lack of economic analysis" to support its apportionment).

This is not a case that implicates any fine-grained disputes about patent valuation. Rather, it asks only the more basic question whether a damages model must attempt to assess the value of the infringed patent (as opposed to the value generated by other technologies) at all. And this Court's cases uniformly agree on the answer, regardless of how they come out on specific methodological questions: Yes. *See, e.g.*, *Power Integrations*, 904 F.3d at 977. Here, however, Kennedy made clear that he did not even attempt to discern what value in Wi-LAN's wireless portfolio—comprising more than a thousand patents—was created by the two patents-in-issue and what was created by the hundreds of other non-infringed patents. Those non-infringed patents relate to important technologies that have nothing to do with the inventions at issue in this case—like Bluetooth, WiFi, and CDMA. Indeed, Kennedy admitted that the Small Niche Co. licenses included other "valuable" patents (*e.g.*, Appx15332-15333 (Q. "My question is simply, the Bluetooth patents that Doro and Vertu received rights to…are valuable." A. "Yes.")). Yet he made no effort to disaggregate the value of the patents-in-suit. He instead relied on a generic and arbitrary assumption and then applied various "adjustments," only two of which had even a marginal effect on his damages calculation. None of these adjustments even tried to distinguish the value created by the patents-in-issue from those created by other patents in Wi-LAN's portfolio.

1. Kennedy's 25% portfolio subset discount model is an abstract theory untethered from the facts of this case.

The Small Niche Co. licenses that underly Kennedy's model cover roughly a thousand wireless patents that practice a range of technologies. Wi-LAN's damages here, however, were ostensibly based on only the two patents-at-issue. Rather than trying to figure out what value was created by those two patents, and what was instead created by the (many) others, Kennedy applied a key assumption: Any subset of the wireless portfolio—no matter how big, how small, or which patents were in that subset—would be licensed at 75% of the value of the whole portfolio. The two patents-in-suit? 75%. Two different patents, say, on Bluetooth? 75%. Two WiFi patents? 75%. This generic and arbitrary rule of thumb—applied without consideration of the actual value of the '145 and '757 patents—is precisely the type of abstract theory that this

Court has consistently rejected as too attenuated from the facts of a case to be admissible.

To carry its "burden of proving damages," Wi-LAN had to "tie the expert testimony on damages to the facts of the case." *Uniloc*, 632 F.3d at 1315. Although experts may invoke "general theor[ies]," such theories are admissible only if accompanied by evidence demonstrating that they are relevant to the particular facts presented. *Id.* at 1316. Even elegant theorems discovered by Nobel laureates, *VirnetX*, 767 F.3d at 1325-26, or doctrines supported by the weight of academic research, *Uniloc*, 632 F.3d at 1312-18, cannot support a damages verdict unless it is proven that they would apply to the "facts and circumstances of the particular case at issue," *id.* at 1318. Without this factual link, "a damages model cannot meet the substantive statutory requirement of apportionment." *CSIRO*, 809 F.3d at 1302.

In *Uniloc*, for example, this Court confronted the use of the "25-percent rule": a theory that hypothesized that the "licensee [would] pay a royalty rate equivalent to 25 percent" of the value of an infringing product. 632 F.3d at 1312. The rule was espoused by academics, who concluded that its "veracity…has been confirmed by a careful examination of years of licensing and profit data, across companies and industries." *Id.* at 1313 (quotations omitted). District courts across the country had endorsed its use. *Id.* at 1314-15. This Court nevertheless held the 25 percent rule categorically inadmissible under *Daubert*. It reasoned that the rule is "an abstract and largely theoretical construct" that applies without regard to the particular facts of the case—such as the nature of the technology at issue, the industries involved, and the relative bargaining power of the parties. *Id.* at 1317. Reliance on such an "arbitrary, general rule, unrelated to the facts of th[e] case," required vacating the damages award and ordering a new trial on damages. *Id.* at 1318.

This Court reached a similar result in *VirnetX*, 767 F.3d at 1331-34, which involved an expert's invocation of the Nash Bargaining Solution—a theorem discovered by Nobel laureate John Nash (*id.* at 1325)—to justify awarding the patentee half of the incremental profits derived from the invention. The Court reasoned that because the Nash theorem presumes the existence of "a certain set of premises," an expert seeking to rely on it must establish that the facts of the case are consistent with the premises that underlie the theorem. *Id.* at 1332.

Reaffirming *Uniloc*, the Court concluded that the expert's failure to do so required vacating the damages award. *Id.* at 1334.

Kennedy's "subset discount" fares no better. It shares neither the academic plaudits nor the judicial endorsements of the theories ultimately struck down in *Uniloc* and *VirnetX*. But like those theories, it has absolutely no connection to the facts of the case. It is little more than an unscientific guess about how parties to a licensing negotiation *might* behave, which is no substitute for the fact-specific apportionment analysis that the Patent Act requires. *Ericsson*, 773 F.3d at 1228.

Kennedy testified that parties to a licensing negotiation generally focus their discussions on a small handful of patents. Appx15231-15234. Once they reach an agreement on these key patents, the licensee will then typically demand that, for a small markup, the patentholder license other relevant families of patents from its portfolio. Appx15231-15234 This practice, according to Kennedy, gives the licensee peace of mind that it will avoid future infringement disputes with the patentee. *See* Appx15231-15234. Based on the existence of this general practice, Kennedy applied a 25% "subset discount" to the royalty rate of the Small Niche Co. licenses. In other words, Kennedy assumed that "75 percent of the value" of the thousand-plus patents in the Small Niche Co. licenses was created by just "the two patents in this case." Appx15335. The remaining 25% would constitute the "markup" that Wi-LAN might have charged Apple to license the rest of its thousand-patent portfolio.

This assumption was made without any factual basis. As Kennedy freely admitted, his subset discount has no connection to the actual value of the '145 and '757 patents. Kennedy would have applied the same "25 percent discount" to "any subset of [patents in] Wi-LAN's portfolio," regardless of their actual value. Appx15336; *see also, e.g.*, Appx15337-15338 (same discount would apply if Apple had infringed Wi-LAN's Bluetooth patents). By Kennedy's own account, his subset discount is precisely the type of fact-agnostic heuristic that this Court has roundly condemned as "too crude a generalization about a vastly more complicated world." *VirnetX*, 767 F.3d at 1332; *see also CSIRO*, 809 F.3d at 1302 ("abstract recitations of royalty stacking theory" that are not "anchored to a quantitative market valuation" are "insufficiently reliable").

Indeed, there were no facts supporting the rationale underlying Kennedy's subset discount. There was, for example, no evidence that Wi-LAN took this portfolio markup approach during the negotiation of the Small Niche Co. licenses. And even if it had, there is no indication that the patents-in-suit were the focus of the negotiation such that they would account for *three-quarters* of the value of the portfolio. Just the opposite. Two of the three Small Niche Co. agreements (Vertu and Unnecto) make no mention of the '757 patent. Appx15586; Appx923-937 (Vertu); Appx884-922 (Unnecto). And the third (Doro) didn't mention the '145 patent. Appx15586; Appx860-883 (Doro). It also makes no mention of the VoLTE technology that Kennedy presumes made up the bulk of the portfolio value, while explicitly calling out that it includes patents covering "Bluetooth, Wi-Fi, CDMA, HSPA, and LTE." Appx15586. And neither the Doro nor the Unnecto license lists either patent-in-suit among the "Asserted Patents" that presumably drove the negotiation. Appx885 (Unnecto); Appx861 (Doro). There is accordingly little reason to believe that the '145 and '757 patents accounted for the lion's share of the Small Niche Co. licenses' value.

Kennedy tried to attribute his 25% figure to the testimony of James Skippen, Wi-LAN's Vice-Chairman and former CEO. Appx15335-15336. Skippen noted that, at the end of licensing negotiations, Wi-LAN "might" "sometimes" agree to license the related portions of its portfolio for a markup of between 5% to 35%. Appx15158. But Kennedy could not point to a single instance of Wi-LAN engaging in this practice. Appx15337. Nor did he testify that the patents-in-suit were the focus of the Small Niche Co. license negotiations.

To be sure, Skippen did remark—after the initiation of this lawsuit—that the patents-in-suit are the "crown jewels" of Wi-LAN's wireless portfolio. Appx15339; Appx15342. But neither Wi-LAN's actions, nor any of its prior licenses, indicate that the patents-in-suit enjoyed that sterling status. In fact, in May of 2014, the patents-in-suit were only two among fifteen wireless patents that Wi-LAN sought to "bring to [Apple's] attention" in hopes of negotiating a license. Appx807-808; Appx15341-15342. And the patents-in-suit are only two among the six Wi-LAN sued Apple on (the remainder dropping away at Wi-LAN's or the court's instigation). Moreover, to the extent that Kennedy relied on Skippen's "crown jewels" statement at all, it merely informed his distinct and unrelated adjustment recognizing that the patents-in-suit are valid and infringed. Appx15342 (patents-in-suit are "crown jewels" because they have been "found valid and infringed"); *see infra* 66-67. Skippen's remark therefore provides no support for Kennedy portfolio subset discount.

And in any event, Skippen's post-complaint qualitative judgment provides no basis for Kennedy's quantitative conclusion that the patents-in-suit are worth exactly three times the combined value of the roughly one thousand other patents in Wi-LAN's wireless portfolio. *See Exmark*, 879 F.3d at 1349-51 (damages opinion not tied to facts of the case where expert emphasized technological importance and commercial success of invention but failed to connect those factors to the specific proposed royalty numbers she advanced). Neither an "apportionment … plucked out of thin air based on vague qualitative notions of the relative importance of the [claimed] technology," nor one with a "complete lack of economic analysis to quantitively support" it, can sustain a damages verdict. *LaserDynamics*, 694 F.3d at 69.

Finally, Kennedy's 25% subset discount rationale is especially untenable in light of an important qualification that Skippen made.

Skippen emphasized that Wi-LAN would engage in this portfolio-markup practice only where the remaining portfolio being added to the license "doesn't really add much value." Appx15158; Appx15157 (likening practice to "throwing in the chaff with the wheat"); Appx15154 (practice applies where there are "only…a few patents that are actually really valuable"). It therefore could not apply to the Small Niche Co. licenses, which by Wi-LAN's account, cover an array of highly valuable inventions. Skippen's speculative testimony therefore cannot justify Kennedy's subset discount.

2. Kennedy's remaining "adjustments" are untethered from the facts of this case and do not satisfy the statutory requirement of apportionment.

In addition to his portfolio subset discount, Kennedy made four other "adjustments" to the $0.50 per-phone royalty rate of the Small Niche Co. licenses, which Wi-LAN says were sufficient to satisfy the apportionment requirement "in the context of a comparable license analysis." *See* Appx784. These adjustments relate to (1) the assumption of validity and infringement, (2) domestic versus global licensing rates, (3) the licensee's sales volume, and (4) the absence of an early adopter discount. Only the first, however, resulted in any change to the royalty rate at all; the other three merely involved Kennedy explaining why he would *not adjust* that rate to reflect various differences between the Small Niche Co. license and the hypothetical one between Apple and Wi-LAN.

But more important, none of these adjustments gets Kennedy any closer to approximating "the incremental value that the patented invention adds to the end product." *Ericsson*, 773 F.3d at 1226. Contrary to Wi-LAN's contentions below, the "apportionment requirement that a royalty should reflect the value of patented technology," applies with full force to comparable-licenses analyses, *Elbit Sys.*, 927 F.3d at 1301, just as it does "in every [patent] case," *Garretson*, 111 U.S. at 121. And here, far from isolating the value of the claimed inventions, Kennedy's adjustments amount to generic considerations that apply to any damages model.

Assumption of Validity & Infringement. Kennedy's damages model—like every model using the "hypothetical negotiation" approach—"assumes that the asserted patent claims are valid and infringed." *Lucent*, 580 F.3d at 1325. This assumption, Kennedy opined, warrants a 20% increase to the royalty rates of the Small Niche

Co. licenses—which were negotiated without the benefit of a judicial determination on validity or infringement. Appx15229-15231; Appx15245-15246. But neither this conclusion, nor the unremarkable observation underlying it, says anything about the specific value of the two patents asserted here. Nor does it say anything else about the particular facts of this case. Whatever the merits of Kennedy's 20% markup, it is not apportionment.

Domestic vs. Global Licensing Rates. Kennedy offered his views on the differences between U.S. and worldwide royalty rates. Appx15239-15240. He ultimately decided, given the varying geographic scopes of the Small Niche Co. licenses, that no adjustment was warranted. *Id.* This testimony speaks to geography, not apportionment. Again, it does nothing to "separate the value of the allegedly infringing features [of an iPhone] from the value of all other features." *CSIRO*, 809 F.3d at 1301.

Sales Volumes. Kennedy noted that companies with large sales volumes—like Apple—can typically negotiate lower royalty rates. Appx15243-15244. He nevertheless declined to apply any downward adjustment to the Small Niche Co. rate, reasoning that any discount would ultimately "cancel … out" with the absence of an early adopter discount (*infra* 68). Appx15244. These musings on the potential relationship between sales volume and royalty rates—general considerations that would apply similarly in any case—have no connection to the technology at issue here.

Early Adopter Discount. Kennedy opined that an "early adopter" discount is sometimes offered to companies that forgo litigation and instead agree to take a license early on. Appx15240-15242; Appx15245-15247. He concluded that Apple, unlike the Small Niche Co. licensees, would not qualify for such a discount. Appx15240-15242; Appx15245-15247. This observation—that a party to a patent lawsuit is not an "early adopter"—will, by definition, be true in every case where patent infringement is alleged. Pointing to this commonsense fact surely does not satisfy Wi-LAN's burden to "seek only those damages attributable to the infringing features" of the iPhone. *VirnetX*, 767 F.3d at 1326.

3. A single data point lodged in a table of Apple's expert damages report cannot support the damages verdict.

Confronted with Kennedy's failure to apportion, Wi-LAN's principal response before the district court was *not* to dispute it. Instead, Wi-LAN argued that a single figure, buried in a table of Apple's damages expert's report, was an independent basis to support the damages award. But this lone, raw data point, which Wi-LAN attempts to wield out of context, suffers from the same problem as Kennedy's overall analysis: It is not apportioned. Moreover, this figure (that was ultimately excluded as an outlier), was only one point among a constellation of data; it carried no independent significance in Apple's damages model. As Apple's expert emphasized, while any one data point may be "instructive," no single piece of data could be "dispositive." Appx15606.

Lance Gunderson—Apple's damages expert—focused on a half-dozen lump-sum royalty licenses that Wi-LAN had previously entered into with large cellphone manufacturers. He applied three different methodologies to these licenses. One methodology used a specific approach to convert these lump-sum amounts into a per-unit rate. This per-unit rate was then multiplied by the number of allegedly infringing iPhones to arrive at what a hypothetical lump-sum license between Wi-LAN and Apple might have been. Gunderson's approach resulted in a damages range of $5 million to $10 million. Appx15606.

Wi-LAN seized upon the per-unit rate of $0.495 that Gunderson calculated for the HTC license, which was higher than Kennedy's proposed rate of $0.45. But this rate did not satisfy the legal requirement of apportionment. It therefore cannot, on its own, support the damages award.

First, it was never Gunderson's testimony that any one number generated by any of his models could serve as an independent basis to support a damages award. To the contrary, he relied on dozens of data points generated from three separate damages models to arrive at his ultimate opinion. Gunderson emphasized that no single data point could be "dispositive," Appx15606—particularly not an outlier, Appx15592-15593; Appx15596-15597. Instead, Gunderson arrived at his ultimate damages calculation by assessing all of the data, in addition to other relevant "facts and circumstances" "particular" to this case. Appx15606-15607. Wi-LAN's suggestion that a single outlier data point could serve as the sole basis to support a damages award is therefore fundamentally inconsistent with the entire basis for Gunderson's testimony.

Second, the HTC figure is not apportioned to take into consideration the portfolio nature of the license—precisely the problem with Kennedy's approach. Appx15638. An unapportioned damages figure—regardless of which expert introduced it—cannot serve as the sole basis to support a damages award.

Third, Gunderson threw out the HTC figure as an outlier, and therefore did not rely on it in calculating damages. Appx15592-15593 (HTC was an "outlier[]…so I threw [it] out"); Appx15596-15597. By the terms of his own model then, it would be inappropriate to use it as an independent basis to extrapolate what the damages award should be in this case.

Finally, there is simply no indication that the jury relied upon the HTC figure. To the contrary, its $85 million verdict reflects the precise figure that Kennedy urged the jury to adopt. *See Pena v. Meeker*, 2014 WL 4684800, at *8 (N.D. Cal. Sept. 18, 2014) ("When a jury awards the exact amount of damages proposed by one party's expert, it is reasonable to infer that the damages award" was based on that expert's testimony.).

B. The Small Niche Company licenses are not comparable to the hypothetical license between Wi-LAN and Apple.

Parties may base their damages models on prior licenses only "[w]here the licenses employed are sufficiently comparable" to the hypothetical license that the parties would have entered into. *CSIRO*, 809 F.3d at 1303; *see also Georgia-Pacific Corp. v. U.S. Plywood Corp.*, 318 F. Supp. 1116, 1120 (S.D.N.Y. 1970). "[A]lleging a loose or vague comparability between different … licenses does not suffice." *LaserDynamics*, 694 F.3d at 79. Rather, "there must be a basis in fact to associate the royalty rates used in prior licenses to the particular hypothetical negotiation at issue in the case." *Uniloc*, 632 F.3d at 1317. Specifically, any use of past licenses must be supported with specific facts "that account[] for the … economic differences" between the prior licenses and the product of a hypothetical negotiation, *ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 873 (Fed. Cir. 2010). "Subsumed within" that requirement is the obligation to account for "whether the licensor and licensee would have agreed to a lump-sum payment or instead to a running royalty based on ongoing sales or usage," *Lucent*, 580 F.3d at 1326.

Here, the only licenses comparable to a hypothetical one between Wi-LAN and Apple are the licenses between Wi-LAN and other large cellphone manufacturers—all of which were for a lump-sum amount. In particular, the unrebutted evidence—including unequivocal statements made by Skippen at the time that the hypothetical negotiation would have taken place—showed that a large company like Apple would have entered into a lump-sum license with Wi-LAN for no more than $10 million.

In April of 2014, only a few months before the hypothetical negotiation would have taken place, Skippen—then the CEO of Wi-LAN—delivered a five-year strategic plan to the Board of Directors. The plan's "stated purpose" was to "to chart the general business direction for Wi-LAN in the 2014 to 2018 time frame." Appx15190. The company had "suffered from recent significant setbacks," including "multiple litigation losses" and an "inability to close [a] sufficient number of significant de[als]." Appx15191-15192.

Skippen identified a key contributor to Wi-LAN's sluggish performance: Companies were demanding to pay "lump sum" royalties, and it was "challenging to convince companies to pay license amounts that are much above single digit millions *at most*." Appx15172; Appx15193-15194 (emphasis added); *see also* Appx616 (Wi-LAN willing to "take less" in licensing negotiation "if it's (sic) been a bad year" financially).

This assessment echoed Gunderson's opinion that Apple would have likely obtained a lump-sum royalty agreement. *See* Appx15566; *see also* Appx15557 ("bigger players, especially in the telecom space almost always do lump sums"); *cf.* Appx15368 ("almost all of [Apple's] licenses are lump-sum agreements").

In a case that was otherwise characterized by vigorous disagreement, here was something that the parties could apparently agree on: Any license between Apple and Wi-LAN would have been a lump-sum royalty, and that royalty would have likely been less than $10 million.

But Kennedy ignored that evidence, and instead engaged in precisely the type of license cherry-picking, designed to "inflate the reasonable royalty analysis with conveniently selected licenses," that this Court has rejected. *Uniloc*, 632 F.3d at 1316. Kennedy declined to examine the many lump-sum licensing agreements between Wi-LAN and large companies—such as Kyocera, LG, and Motorola (*see* Appx15591; *see also* Appx15247 (Kennedy admitting that "Wi-LAN gets a lot of lump sums and maybe even prefers them at times"))—and instead focused exclusively on running royalty licenses with three very small, niche market cellphone providers: Unnecto, Vertu, and Doro.[10] Unlike Apple, these small companies possessed neither the bargaining power nor the balance sheet to obtain a lump-sum royalty, and were instead forced to take running royalty licenses on less favorable terms. As Kennedy admitted: "you have to account for differences. And there's

---

[10] Kennedy did examine a single lump-sum license that was ultimately excluded by the district court. Appx42; Appx57-58. According to Kennedy, this license was *not* comparable to the hypothetical license and was being used solely to present the jury with a specific contractual obligation. Appx42; Appx57-58. The district court excluded the license, relying on Kennedy's representation that it is not comparable, and also because Kennedy failed to apportion this license, and finally, because the contractual obligation was irrelevant and prejudicial. Appx58. In so holding, the court rejected Wi-LAN's contrary argument that the license was admissible simply because it is a "real-world" license. Appx57 ("Presumably, all of the licenses the parties rely on in the hypothetical negotiation analysis are 'real-world' licenses, but that does not automatically render them admissible.").

differences, obviously, between these small companies and a company like Apple." Appx15228; *see also* Appx15328 (admitting that "the licenses that [the Small Niche Co.'s] sign[ed] are also not like the hypothetical license that Apple would have signed"). But he failed to follow his own advice. The Small Niche Co. licenses are not sufficiently comparable to support the damages award.

This Court has vacated damages awards under similar circumstances. In *Wordtech*, for example, the Court rejected the patentholder's reliance on running royalty licenses where it was likely that the parties would have entered into a lump-sum agreement. 609 F.3d at 1320-21. This conclusion was not altered just because, as here, the prior licenses include "some…of the patents-in-suit." *Id.* at 1319. Similarly, in *Lucent* this Court held that certain running royalty licenses were not comparable because "fundamental differences exist between lump-sum agreements and running-royalty agreements," and because, as in this case, the running royalty licenses "differ substantially from the hypothetical negotiation scenario." 580 F.3d at 1330. And *LaserDynamics* clarified that an expert may not cherry-pick favorable licenses, while simply ignoring the most comparable licenses in the record. 694 F.3d at 80. That is precisely what Kennedy did here. He zeroed in on the Small Niche Co. licenses, "to the exclusion of the many licenses" that were objectively more comparable. *Id.*

Because the Small Niche Co. licenses are not sufficiently comparable, Kennedy's damages model based on them was inadmissible and could not support the damages award.

CONCLUSION

For the foregoing reasons, the Court should reverse the judgment of infringement and vacate the damages award, or, at a minimum, remand for a new trial on damages.

Respectfully submitted,

/s/ Mark S. Davies

Mark S. Davies
Thomas Fu
James Anglin Flynn
Katherine M. Kopp
ORRICK, HERRINGTON & SUTCLIFFE LLP
1152 15th Street, NW
Washington, DC 20005
(202) 339-8400

Max Carter-Oberstone
ORRICK, HERRINGTON & SUTCLIFFE LLP
405 Howard Street
San Francisco, CA 94105

Sean C. Cunningham
Stanley J. Panikowski
Erin P. Gibson
DLA PIPER LLP (US)
401 B Street, Suite 1700
San Diego, CA 92101
(619) 699-2700

Counsel for Appellant

November 20, 2020

ADDENDUM

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC.,

Plaintiff,

vs.

WI-LAN, INC.,

Defendant.

AND ALL RELATED COUNTERCLAIMS.

CASE NO. 14cv2235 DMS (BLM)

ORDER CONSTRUING CLAIMS

This matter came before the Court for a claim construction hearing on October 30, 2017. John Allcock and Sean Cunningham appeared and argued on behalf of Apple and Allison Goddard, Kevin Schubert, Robert Cote and Seth Hasenour appeared and argued on behalf of Wi-LAN. After the matter was submitted, Apple filed a Notice of Supplemental Evidence Regarding Claim Construction, to which Wi-LAN filed a response. After a thorough review of the parties' claim construction briefs and all other material submitted in connection with the hearing, the Court issues the following order construing the disputed terms of the patents at issue in this case.

I.
BACKGROUND

This case is related to another case, involving the same parties, which was previously adjudicated by this Court, *Wi-LAN v. Apple*, Case Number 13cv0798. That case involved two Wi-LAN Patents, United States Patents Numbers 8,311,040 ("the '040 Patent") and 8,315,640 ("the '640 Patent"). The Court construed the claims of the '040 Patent and the '640 Patent and then granted summary judgment of noninfringement to Apple. After that ruling, the parties stipulated to entry of final judgment so that Wi-LAN could appeal. On appeal, Wi-LAN challenged this Court's claim construction ruling, specifically the Court's constructions of the term "specified connection" in the '040 Patent and the term "UL connections" in the '640 Patent. The Federal Circuit affirmed the Court's constructions and the grant of summary judgment of noninfringement to Apple.

After the Court's claim construction ruling but before Apple filed its motion for summary judgment in the prior case, Apple filed the present case against Wi-LAN in the United States District Court for the Northern District of California alleging declaratory judgment claims for noninfringement and invalidity of five other Wi-LAN Patents, United States Patents Numbers 8,462,723 ("the '723 Patent"), 8,615,020 ("the '020 Patent"), 8,457,145 ("the '145 Patent"), 8,462,761 ("the '761 Patent") and 8,537,757 ("the '757 Patent"). Apple later filed an amended complaint adding the '040 Patent to the case. Shortly before this Court issued its summary judgment ruling in the prior case, the Northern District of California transferred this case to this Court. After Wi-LAN filed its appeal, Apple moved to stay this case pending that appeal, which the Court granted. After the appeal was decided, the stay was lifted and this case was put back on the Court's calendar.

Pursuant to Patent Local Rule 4.2.a, the parties have identified eight terms or groups of terms for construction in this case:

(1) "wireless subscriber unit"/ "subscriber unit"/ "subscriber station," which terms appear in the '145 Patent, '723 Patent, '020 Patent, '761 Patent and '757 Patent;

(2) "connections"/ "uplink connections"/ "a plurality of connections served by the subscriber unit/connections established at a [or the] subscriber unit [or

subscriber station]," which terms appear in the '145 Patent, '723 Patent, '020 Patent, '761 Patent and '757 Patent;

(3) "queue," which term appears in the '145 Patent, the '723 Patent, the '761 Patent and the '020 Patent;

(4) "packing sub-header," which term appears in the '040 Patent;

(5) "frame map"/ "sub-frame map," which terms appears in the '723 Patent, '020 Patent and the '757 Patent;

(6) "poll-me bit"/ "poll-me message," which terms appear in the '020 Patent;

(7) "fairness algorithm," which appears in the '145 Patent; and

(8) whether the preamble in Claim 26 of the '145 Patent is limiting.[1]

II.

DISCUSSION

The first four terms and groups of terms were at issue, or are similar to terms that were at issue, in the prior case. For that reason, Wi-LAN argues relitigation of these terms is barred by the doctrine of issue preclusion. The Court addresses that argument first, then turns to the construction of the claim terms and groups of terms.

A. Issue Preclusion

The term "issue preclusion" encompasses the doctrine once known as "collateral estoppel." *Taylor v. Sturgell*, 553 U.S. 880, 892 n.5 (2008). "Issue preclusion ... bars successive litigation of an issue of fact or law actually litigated and resolved in a valid court determination essential to the prior judgment ... ." *Id.* at 893 (internal quotation marks and citations omitted).

Issue preclusion, of course, is not unique to patent cases. *Aspex Eyewear, Inc. v. Zenni Optical Inc.*, 713 F.3d 1377, 1380 (Fed. Cir. 2013). Accordingly, the Federal Circuit is "guided by the precedent of the regional circuit. However, for any aspects

/ / /

[1] The parties initially requested that the Court also construe the term "QoS," but they have since agreed on the construction of that term.

that may have special or unique application to patent cases, Federal Circuit precedent is applicable." *Id.*

In the Ninth Circuit, issue preclusion applies when:

> (1) the issue necessarily decided at the previous proceeding is identical to the one which is sought to be relitigated; (2) the first proceeding ended with a final judgment on the merits; and (3) the party against whom issue preclusion is asserted was a party or in privity with a party at the first proceeding.

Paulo v. Holder, 669 F.3d 911, 917 (9th Cir. 2011) (internal quotation marks, citation and brackets omitted). The party asserting issue preclusion bears the burden of showing these elements are met. *Offshore Sportswear v. Vuarnet Int'l, B.V.*, 114 F.3d 848, 850 (9th Cir. 1997).

Wi-LAN has not met that burden here. First, for the first two terms, Wi-LAN has not shown the terms at issue here are identical to the terms at issue in the prior case. In the prior case, the Court construed the terms "wireless subscriber radio unit," "wireless communication radio unit" and "UL connections." The terms at issue here are similar, "subscriber unit," "wireless subscriber unit," "subscriber station," "connections," "uplink connections," "a plurality of connections served by the subscriber unit" and "connections established at a subscriber unit," but they are not identical to the terms construed in the prior case. Thus, issue preclusion does not apply to the first two groups of terms.

The term "queue" is identical to a term that was at issue in the prior case, but Wi-LAN has not shown the parties actually litigated that term. Rather, the parties stipulated to the construction of that term in the prior case. Thus, Wi-LAN has not shown this term is subject to issue preclusion.

The final term, "packing sub-header," was at issue in the prior case and was actually litigated. However, Wi-LAN has not shown that term was "necessarily decided" in the prior case. Indeed, the term played no part in the Court's summary judgment ruling, judgment thereon and subsequent appeal. Accordingly, "packing sub-header" is not subject to issue preclusion either.

For these reasons, the Court declines to apply issue preclusion to the above terms.

B. Claim Construction

Claim construction is an issue of law, *Markman v. Westview Instruments, Inc.*, 517 U.S. 370, 372 (1996), and it begins "with the words of the claim." *Nystrom v. TREX Co., Inc.*, 424 F.3d 1136, 1142 (Fed. Cir. 2005) (citing *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996)). Generally, those words are "given their ordinary and customary meaning." *Id.* (citing *Vitronics*, 90 F.3d at 1582). This "'is the meaning that the term would have to a person of ordinary skill in the art in question at the time of the invention.'" *Id.* (quoting *Phillips v. AWH Corp.*, 415 F.3d 1303, 1313 (Fed. Cir. 2005)). "The person of ordinary skill in the art views the claim term in the light of the entire intrinsic record." *Id.* Accordingly, the Court must read the claims "'in view of the specification, of which they are a part.'" *Id.* (quoting *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995)). In addition, "'the prosecution history can often inform the meaning of the claim language by demonstrating how the inventor understood the invention and whether the inventor limited the invention in the course of prosecution, making the claim scope narrower than it would otherwise be.'" *Id.* (quoting *Phillips*, 415 F.3d at 1318).

1. "Subscriber" Terms

The first group of terms at issue here are the "subscriber" terms, "wireless subscriber unit," "subscriber unit" and "subscriber station," which terms appear in the '145 Patent, '723 Patent, '020 Patent, '761 Patent and '757 Patent. In each of the Patents, the "subscriber" terms are described as part of a method or system of allocating, requesting or obtaining bandwidth from a base station. The parties agree these terms should be construed consistently across the Patents. Apple proposes they be construed as "fixed or portable customer premises equipment that wirelessly receives UL bandwidth from a base station, and allocates the bandwidth across connected user devices." Wi-LAN proposes that the terms be construed as a "module that receives UL

Case: 20-2011 Document: 27 Page: 98 Filed: 11/20/2020

bandwidth from a base station, and allocates the bandwidth across its user connections." As is evident from the parties' proposals, they agree the subscriber units receive UL bandwidth from a base station and allocate that bandwidth elsewhere. The disputes are two-fold: (1) Whether the subscriber unit/station should be construed as "fixed or portable customer premises equipment" or as a "module," and (2) whether the bandwidth is allocated to "connected user devices" or "user connections."

a. CPE or Module

In the prior case, there was a similar issue with respect to the terms "wireless subscriber radio unit" and "wireless communication radio unit," with Apple arguing that these units were equivalent to customer premises equipment or CPEs and Wi-LAN arguing to the contrary. The Court agreed with Wi-LAN and refused to limit these units to CPEs. Wi-LAN argues the Court should adopt that approach here, and for the following reasons, the Court agrees.

First, Wi-LAN did not use the term "CPE" in the patent claims at issue here. Rather, the claims recite "subscriber units" or "subscriber stations." As Apple pointed out in its responsive brief, Wi-LAN used the term "CPE" in the claims of two of its prior patents. (*See* Decl. of Peter Maggiore in Supp. of Apple's Responsive Br., Ex. 1 at 21; Ex. 2 at 51.[2]) Its decision not to use that term here counsels against construing the "subscriber" terms as CPEs.

Second, although the specifications are rife with the term "CPE," the patent claims "will not be confined to that example 'unless the patentee has demonstrated a clear intention to limit the claim scope using words or expression of manifest exclusion or restriction.'" *Aria Diagnostics, Inc. v. Sequenom, Inc.*, 726 F.3d 1296, 1302 (Fed. Cir. 2013) (quoting *Liebel Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 906 (Fed. Cir. 2004)). There is no evidence of that intent here. To the contrary, and as stated above,

/ / /

/ / /

[2] The page numbers cited reflect the numbers assigned by the parties.

the claims recite "subscriber units" and "subscriber stations," as does the specification. (*See* '723 Patent at 1:28-60.)[3]

Third, the prosecution history of the '723 Patent reveals the Examiner treated the "subscriber" terms at issue here as interchangeable with the "subscriber" terms at issue in the prior case. (*See* Index of Exs. in Supp. of Wi-LAN's Opening Br., Ex. G at 51.) There, the Court declined to limit the subscriber terms to CPEs, and for the reasons set out there and above, the Court declines to do so here.

b. User Devices or User Connections

Turning to the second issue, Apple also fails to cite any evidence to support its proposed construction of the terms to require the allocation of bandwidth across "connected user devices." The claims themselves do not use this language, but instead use the term "connections," ('761 Patent at 38:35-37; '020 Patent at 38:40-44; '723 Patent at 23:11-13; '145 Patent at 32:57-59; '757 Patent at 4:33-41), which is consistent with Wi-LAN's proposal.

For these reasons, the Court adopts Wi-LAN's proposed construction for the "subscriber" terms.

2. "Connections" Terms

The next group of terms is the "connections" terms, which are "connections," "uplink connections," "a plurality of connections served by the subscriber unit" and "connections established at a [or the] subscriber unit [or subscriber station]," which terms appear in the '145 Patent, '723 Patent, '020 Patent, '761 Patent and '757 Patent. As with the "subscriber" terms, the parties agree the "connections" terms should be

/ / /

[3] Inherent in Apple's understanding of a "CPE" is that it is fixed or portable, but not mobile. The specification does not limit "subscriber units" or "subscriber stations" in that way, however. (*See id.* at 1:28-38) ("Exemplary communication systems include mobile cellular telephone systems, personal communication systems (PCS), and cordless telephones.") Furthermore, the Examiner of the '020 Patent read the invention claimed therein on prior art that included a "cellular telephone network." (*See* Index of Exs. in Supp. of Wi-LAN's Opening Br., Ex. I at 75, Ex. J at 92.) Thus, this evidence also refutes Apple's proposal to construe the "subscriber terms" as CPEs.

construed consistently across the Patents. Apple proposes the terms be construed as "wired or wireless connections between the subscriber station and its connected user devices," while Wi-LAN proposes the terms be construed as "connections between the subscriber unit and its users."

Although the common term here is "connections," the real dispute centers on the meaning of "users." Consistent with its argument on the "subscriber" terms, Apple argues "users" must be separate user devices, while Wi-LAN urges a broader meaning. This dispute is similar to one raised in the prior case, namely whether the "connections" identified in the Patents can be included in one device or must be in different devices. In the prior case, the Court found the "connections" did not have to be in different devices, *i.e.*, that there could be multiple "connections" in one device. Apple attempts to move the Court away from its prior construction, but the intrinsic evidence it cites is not persuasive.[4]

Based on the specification, the Court agrees with Wi-LAN's proposed construction of the "connection" terms. Although the specification describes "users" to "include both residential and business customers," ('723 Patent at 2:10-12), it also recites "user applications." (*Id.* at 6:63-67.) The specification also goes on to describe different types of "connections," *e.g.*, VBR connections and DAMA connections, (*id.* at 20:6-8), which suggests the "users" are not confined to "user devices." For these reasons, the Court adopts Wi-LAN's proposed construction of the "connections" terms.

<u>3.</u> <u>"Queue"</u>

The next term is "queue," which appears in the '145 Patent, the '723 Patent, the '761 Patent and the '020 Patent. Apple proposes the Court construe the term as

/ / /

---

[4] Apple's reliance on the Federal Circuit's opinion in the prior case as support for its proposed construction of the "connections" terms, which reliance was first disclosed at the *Markman* hearing, is also unpersuasive. The issue here, whether the "connections" are between the subscriber station and "users" or "user devices," was not before the Federal Circuit. Thus, any reference to "user devices" in that opinion is neither helpful nor determinative of the issue here.

"buffer(s), each associated with a unique QoS, containing data to be transmitted."[5] Wi-LAN proposes the term be construed as "structure(s) containing data to be transmitted."

The term "queue" was at issue in the prior case, and the parties agreed it should be construed as "structure(s) containing data to be transmitted." Apple argues for a different construction in this case, but the intrinsic evidence it relies on does not support its proposed construction. On the contrary, the intrinsic evidence supports Wi-LAN's proposed construction. For instance, Claim 1 of the '723 Patent describes a "queue," and then modifies it with the phrase "based on the quality of service (QoS) of the data." (*Id.* at 23:11-13.) Apple's proposed construction would make that modifier redundant. Thus, the Court adopts Wi-LAN's proposed construction of "queue."

<u>4. "Packing Sub-Header"</u>

The next term is "packing sub-header," which appears in the '040 Patent. Apple proposes the Court construe this term as "a header located in a PDU payload." Wi-LAN proposes the term be construed as it was in the prior case as "a header located in a PDU."

Although issue preclusion does not bar relitigation of the construction of this term, the Court is not persuaded that its prior construction is incorrect. The Court thoroughly considered this issue in the prior case, and stands on that construction for the reasons set out there and because Apple presents no new evidence to warrant a departure from the prior construction. Thus, the Court adopts Wi-LAN's proposed construction of this term.

<u>5. "Frame Map/Sub-Frame Map"</u>

The next two terms are "frame map" and "sub-frame map," which appear in the '723 Patent, the '020 Patent and the '757 Patent. Apple proposes these terms be construed as "a single data structure that allocates physical slots for bandwidth grants

---

[5] On the Friday before the *Markman* hearing, the parties submitted an Amended Joint Claim Construction Statement, Chart and Worksheet. In that pleading, the parties informed the Court they had agreed on the construction of "QoS," and Apple amended its proposed constructions of three terms: "queue," frame/sub-frame map and poll-me bit/message. The proposals cited here are the amended proposals.

to individual subscriber stations, and that contain all UL or DL bandwidth grants for an entire [frame/sub-frame]." Wi-LAN proposes the terms be construed as "control information indicating the bandwidth allocated to a subscriber unit(s) in the uplink or downlink communications link within a frame."

The specification supports both sides' proposals, namely that the frame/sub-frame maps may both allocate bandwidth (Apple) and indicate bandwidth allocations (Wi-LAN). (*Compare* '723 Patent at 4:12-23, 9:27-29, 10:4-6, 12:9-13, 21-24, 15:7-9, 13-15, 30-41, 16:56-59, 17:14-15, 63-18:2, 19:23-29, 22:37-38 (stating frame/sub-frame maps allocate bandwidth) *with* 13:58-60, 14:25-27 (stating frame/sub-frame maps "communicate" and "indicate" bandwidth allocations.) However, neither side's proposal allows for both of those possibilities. Because each side places a limitation on the terms that is not supported by the specification, the Court declines to adopt either proposal. Rather, the Court finds each proposal includes some concepts that should be included in the proper construction of the terms. Thus, the Court construes frame/sub-frame maps as "data structures that may allocate bandwidth to subscriber station(s) and indicate the bandwidth allocated to subscriber unit(s) within a particular frame/sub-frame."

6. "Poll-Me Bit/Poll-Me Message"

The next two terms are "poll-me bit" and "poll-me message," which terms are found in the '020 Patent. Apple proposes these terms be construed as "[a bit/bits] in a MAC packet utilized to indicate to the base station that the subscriber station requires a change in UL bandwidth allocation." Wi-LAN proposes the terms be construed as "a bit sent by a currently active subscriber unit, that currently has bandwidth allocations, indicating a request to be provided an allocation of UL bandwidth in which to transmit a bandwidth request."

Of these two proposals, Wi-LAN's is more consistent with the specification. The specification confirms that "poll-me bits" and "poll-me messages" must come from "currently active" users, and "currently active" users are users that already have bandwidth allocations. (*See* '020 Patent at 11:15-26, 11:45-29, 13:7-14, 19:45-48, 37:41-44.) The specification also makes clear that the purpose of the "poll-me bits" and "poll-me messages" is to request bandwidth from the base station.

Apple's proposed construction includes a limitation that is not supported by the specification, namely that the bit be utilized in a MAC packet. Apple's proposed construction also states the purpose of the "poll-me bits" and "poll-me messages" is to "indicate to the base station that the subscriber station requires a change in bandwidth allocation," which is not exactly accurate. The purpose of these bits and messages is to request bandwidth from the base station, not to indicate a change in bandwidth allocation. Wi-LAN's terminology is more consistent with the specification, and thus the Court adopts Wi-LAN's proposed construction of these terms.

<u>7.</u> <u>"Fairness Algorithm"</u>

The last true term at issue is "fairness algorithm," which appears in the '145 Patent. Apple proposes this term be construed as "a QoS algorithm to prioritize the transmission of user data of one connection over that of at least one other connection." Wi-LAN proposes the term be construed as a "QoS algorithm to ensure fair handling of the queued data."

Between these two proposals, Wi-LAN's is more consistent with the specification. As stated therein, the purpose of the "fairness algorithm" is "to ensure fair handling of the data queued at" a QoS when "there is insufficient bandwidth to transmit all queued data during the current TDD frame." ('145 Patent at 22:65-23:3, 57-67.) Apple's proposed construction reads the "fairness" requirement out of the algorithm and replaces it with prioritization, but simple prioritization does not ensure fairness. Indeed, the process of prioritization, untethered to the concept of fairness, could result in one connection always receiving bandwidth while others go without. This is precisely the situation sought to be avoided by the use of "fairness algorithms," and thus, prioritization is not an accurate descriptor.

Furthermore, the specification provides an example of a fairness algorithm that does not involve any prioritization. (*See id.* at 23:3-12.) In that algorithm, Continuous Grant, "[a]ll data in the[ ] queues must be sent every TDD frame." (*Id.* at 10-11.)

The claims and the specification both read in terms of "fairness," not prioritization. Because's Wi-LAN's proposal is the only one that reflects that concept, the Court adopts Wi-LAN's proposed construction of this term.

8. Preamble

The final issue for the Court to decide is whether the preamble of Claim 26 of the '145 Patent is limiting. Apple asserts it is while Wi-LAN states it is not.

"'Whether to treat a preamble as a limitation is a determination resolved only on review of the entire[ ] . . . patent to gain an understanding of what the inventors actually invented and intended to encompass by the claim.'" *Poly-Am., L.P. v. GSE Lining Tech., Inc.*, 383 F.3d 1303, 1309 (Fed. Cir. 2004) (quoting *Corning Glass Works v. Sumitomo Elec. U.S.A., Inc.*, 868 F.2d 1251, 1257 (Fed. Cir. 1989)). There is no litmus test for determining when a preamble limits the invention. *Id.*; *Catalina Mktg. Int'l, Inc. v. Coolsavings.com, Inc.*, 289 F.3d 801, 808 (Fed. Cir. 2002). There are, however, some guidelines.

"'Generally, the preamble does not limit the claims.'" *Georgetown Rail Equip. Co. v. Holland L.P.*, 867 F.3d 1229, 1236 (Fed. Cir. 2017) (quoting *Allen Eng'g Corp. v. Bartell Indus., Inc.*, 299 F.3d 1336, 1346 (Fed. Cir. 2002)).

> However, a preamble may be limiting if: "it recites essential structure or steps"; claims "depend[] on a particular disputed preamble phrase for antecedent basis"; the preamble "is essential to understand limitations or terms in the claim body"; the preamble "recit[es] additional structure or steps underscored as important by the specification"; or there was "clear reliance on the preamble during prosecution to distinguish the claimed invention from the prior art."

Id. (quoting *Catalina Mktg.*, 289 F.3d at 808). In contrast, the preamble does not limit the claims when the "'patentee defines a structurally complete invention in the claim body and uses the preamble only to state a purpose or intended use for the invention.'" *Poly-Am.*, 383 F.3d at 1310 (quoting *Rowe v. Dror*, 112 F.3d 473, 478 (Fed. Cir.

1997)). "'[P]reamble language merely extolling benefits or features of the claimed invention does not limit the claim scope without clear reliance on those benefits or features as patentably significant.'" *Georgetown Rail*, 867 F.3d at 1236 (quoting *Catalina Mktg.*, 289 F.3d at 809).

Here, Claim 26 of the '145 Patent recites:

A subscriber unit for a wireless communication system, comprising:

> a plurality of queues for buffering user traffic according to a traffic parameter, each queue having an associated logical state;
>
> a media access control (MAC) element capable of
>
> > transmitting an uplink (UL) bandwidth request based on the logical state of the queues during a bandwidth request opportunity, and
> >
> > allocating between the queues a bandwidth allocation received in response to the UL bandwidth request, based on the current state of the queues.

Apple asserts the preamble is limiting because the body of the claim does not recite a structurally complete invention. Wi-LAN relies on the general rule that preambles are not limiting, and argues the preamble "merely names" the limitations set out in the body of the claim and sets out the intended use of the invention.

The Court agrees with Apple that the preamble is limiting for several reasons. First, the body of the claim does not recite a complete structure without the preamble. It is clear from the Patent as a whole that a subscriber unit is a critical part of the invention, and without that limitation, the body of the claim has no context. Second, the preamble does not "merely name" the limitations set out in the body of the claim. Rather, as stated above, it provides context for the limitations, or an "essential structure" for those elements. And finally, the preamble does not recite an intended use for the invention. It is a part of the invention, not just a use therefor. Thus, for these reasons, the Court finds the preamble of Claim 26 of the '145 Patent is limiting.

/ / /

/ / /

III.

CONCLUSION

For the reasons stated above, the disputed terms are interpreted as set forth in this Order.

IT IS SO ORDERED.

DATED: November 13, 2017

HON. DANA M. SABRAW
United States District Judge

CONFIDENTIAL MATERIAL OMITTED

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC., Plaintiff, vs. WI-LAN, INC., Defendant. AND ALL RELATED COUNTERCLAIMS.	CASE NO. 14cv2235 DMS (BLM) **ORDER:** **(1) DENYING APPLE'S MOTION TO STRIKE WI-LAN'S NEW INFRINGEMENT THEORIES;** **(2) GRANTING IN PART WI-LAN'S MOTION FOR PARTIAL SUMMARY JUDGMENT;** **(3) GRANTING IN PART AND DENYING IN PART APPLE'S MOTION FOR SUMMARY JUDGMENT;** **(4) GRANTING IN PART AND DENYING IN PART APPLE'S MOTION TO EXCLUDE CERTAIN OPINIONS OF VIJAY MADISETTI, DAVID KENNEDY AND JEFFREY PRINCE; AND** **(5) DENYING IN PART WI-LAN'S OMNIBUS MOTION TO EXCLUDE TESTIMONY OF APPLE EXPERTS** [redacted]

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

CONFIDENTIAL MATERIAL OMITTED

[redacted]

IT IS SO ORDERED.

DATED: June 29, 2018

HON. DANA M. SABRAW
United States District Judge

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC., Plaintiff, vs. WI-LAN, INC., Defendant.	CASE NO. 14cv2235 DMS (BLM)
AND ALL RELATED COUNTERCLAIMS.	**ORDER (1) DENYING APPLE INC.'S RENEWED MOTION FOR JUDGMENT AS A MATTER OF LAW, (2) GRANTING APPLE INC.'S MOTION FOR A NEW TRIAL AND/OR REMITTITUR AND (3) DENYING WI-LAN'S MOTION FOR SUPPLEMENTAL DAMAGES, ONGOING ROYALTY, AND PREJUDGMENT AND POST JUDGMENT INTEREST**

This case comes before the Court on Apple Inc.'s renewed motion for judgment as a matter of law and/or motion for a new trial and Wi-LAN's motion for supplemental damages, ongoing royalty, and prejudgment and post judgment interest. On November 30, 2018, the Court heard argument on the damages portion of Apple's motion. Ashley Moore appeared and argued for Wi-LAN, and Sean Cunningham appeared and argued for Apple. After reviewing the parties' briefs, the record, the relevant legal authority, and after hearing argument from counsel, the Court issues the following rulings:

/ / /

I.

RENEWED MOTION FOR JUDGMENT AS A MATTER OF LAW

Apple moves for judgment as a matter of law on the issue of infringement. "A Rule 50(b) motion for judgment as a matter of law is not a freestanding motion. Rather, it is a renewed Rule 50(a) motion." *E.E.O.C. v. Go Daddy Software, Inc.*, 581 F.3d 951, 961 (9th Cir. 2009). Federal Rule of Civil Procedure 50(a)(1) provides:

> 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 the controlling law, can be maintained or defeated only with a favorable finding on that issue.

Fed. R. Civ. P. 50(a)(1). In the Ninth Circuit, "'[j]udgment as a matter of law is appropriate when the evidence presented at trial permits only one reasonable conclusion.'" *Torres v. City of Los Angeles*, 548 F.3d 1197, 1205 (9th Cir. 2008) (quoting *Santos v. Gates*, 287 F.3d 846, 851 (9th Cir. 2002)). "In other words, '[a] motion for a judgment as a matter of law is properly granted only if no reasonable juror could find in the non-moving party's favor.'" *Id.* (quoting *El-Hakem v. BJY Inc.*, 415 F.3d 1068, 1072 (9th Cir. 2005)). When considering a motion for judgment as a matter of law, the court must view the evidence "'in the light most favorable to the nonmoving party, and all reasonable inferences must be drawn in favor of that party.'" *Id.* at 1205-06 (quoting *LaLonde v. County of Riverside*, 204 F.3d 947, 959 (9th Cir. 2000)).

Here, Apple raises a number of arguments in support of its motion for judgment as a matter of law on the issue of infringement. Several of these legal arguments were raised and rejected prior to trial, *e.g.*, the claim construction arguments. Apple has failed to show that the Court's previous rulings were in error, and thus those arguments do not warrant judgment as a matter of law in Apple's favor. On the evidentiary arguments, Apple has failed to show that no reasonable juror could have found for Wi-LAN, and thus those arguments also do not warrant judgment as a matter of law in

Apple's favor. Thus, the Court denies Apple's motion for judgment as a matter of law on the issue of infringement.[1]

II.

MOTION FOR NEW TRIAL OR REMITTITUR

Apple's motion for a new trial on damages is based on Federal Rule of Civil Procedure 59, which provides: "The court may, on motion, grant a new trial on all or some of the issues-and to any party-as follows: (A) after a jury trial, for any reason for which a new trial has heretofore been granted in an action at law in federal court[.]" Fed. R. Civ. P. 59(a)(1)(A). "A trial court should grant a motion for a new trial if (1) the jury instructions were erroneous or inadequate, (2) the court made incorrect and prejudicial admissibility rulings, or (3) the verdict is contrary to the great weight of the evidence." *Chiron Corp. v. Genentech, Inc.*, 363 F.3d 1247, 1258 (Fed. Cir. 2004) (citations omitted).

As an alternative to a new trial on damages, Apple requests that the Court enter a conditional order of remittitur to a $10 million damages award. "'The Court has discretion to grant a remittitur, reducing the damages to the maximum authorized under the evidence, and then offer Plaintiffs the choice of accepting a remittitur (a reduction) of the award in lieu of a new trial on the issue of the damages only.'" *Coach, Inc. v. Celco Customs Services Co.*, No. CV 11-10787 MMM (FMOx), 2014 WL 12573411, at *23 n.128 (C.D. Cal. June 5, 2014) (quoting *Dixon v. City of Coeur d'Alene*, No. 2:10-cv-00078-LMB, 2012 WL 2923149, at *8 (D. Idaho July 18, 2012)).

/ / /

---

[1] Apple also moves for judgment as a matter of law of no damages on the ground Wi-LAN "failed to meet its burden of proving damages." (Mem. of P. & A. in Supp. of Apple's Mot. at 11.) At oral argument, Apple presented the Court with another option, namely entering judgment as a matter of law in the amount of $24 million in damages. That was the first mention of this option, and thus the Court declines to consider it here. Even if the Court considered it, however, Apple has failed to show there is sufficient evidence in the record for the Court to enter judgment as a matter of law in that amount. Furthermore, Apple's arguments on damages are directed more toward Wi-LAN's methodology, not a lack of evidence to support a damages award. Therefore, the Court addresses the issue of damages below under Apple's alternative motion for a new trial or remittitur.

CONFIDENTIAL MATERIAL OMITTED

In this case, the primary point of contention on the damages issue is apportionment.[2] Both sides agree that apportionment was required, but they disagree on the method for doing so. Apple apportioned by using the smallest salable patent practicing unit ("SSPPU"), which Apple argued was the baseband processor, while Wi-LAN used a "direct valuation" approach. Apple contends Wi-LAN's approach was riddled with legal and factual errors, and thus Apple is entitled to a new trial on damages or a remittitur to $10 million.

The general rule of apportionment is that "[a] patentee is only entitled to a reasonable royalty attributable to the infringing features." *Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.*, 904 F.3d 965, 977 (Fed. Cir. 2018). As stated above, there is no dispute that apportionment was required in this case. Thus, Wi-LAN was required, as part of its reasonable royalty analysis, to "apportion[ ] between the infringing and non-infringing features of the product." *Id.* (citations omitted).

Here, the accused product was the iPhone, and thus Wi-LAN had the burden to apportion the infringing features of the iPhone from the noninfringing features. Generally, this kind of apportionment is accomplished by ensuring the royalty base is not "larger than the smallest salable unit embodying the patented invention." *Id.*[3] If the SSPPU "itself contains several non-infringing features[,]" the patentee must apportion further by "estimat[ing] what portion of that smallest salable unit is attributed

/ / /

/ / /

/ / /

/ / /

---

[2] To be sure, Apple raises other arguments, namely, that evidence of [redacted] skewed the damages horizon, and that Wi-LAN improperly included millions of non-infringing iPhones in the royalty base. However, in light of the discussion below, the Court declines to address these other arguments.

[3] The Court notes the parties dispute what constitutes the SSPPU in this case. Apple argues it is the baseband processor while Wi-LAN asserts it is the iPhone. The Court need not resolve this issue here, however.

to the patented technology." *Id.* (citing *VirnetX, Inc. v. Cisco Systems, Inc.*, 767 F.3d 1308, 1327 (Fed. Cir. 2014).[4]

Wi-LAN argues, however, that apportionment may be accomplished by other means, and that courts should allow "flexibility in arriving at apportionment." (Wi-LAN's Opp'n to Mot. at 15) (citations omitted). There is authority to support both of these arguments, *see Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1315 (Fed. Cir. 2014), *overruled on other grounds by Williamson v. Citrix Online, LLC*, 792 F.3d 1339 (Fed. Cir. 2015), (stating party may "estimate the value of the benefit provided by the infringed features by comparing the accused product to non-infringing alternatives."); *Commonwealth Scientific and Indus. Research Org. v. Cisco Systems, Inc. ("CSIRO")*, 809 F.3d 1295 (Fed. Cir. 2015) (stating "adaptability" may be necessary in the apportionment analysis), but neither of these cases resolves the issues raised here. *CSIRO*, for instance, was a unique case wherein the parties engaged in actual license negotiations to the patent in suit. 809 F.3d at 1303. In determining a reasonable royalty, the district court used those negotiations "as a lower bound on a reasonable royalty," and the Federal Circuit affirmed that approach. *Id.* at 1304. This case does not present facts similar to those found in *CSIRO*, or facts that would necessarily call for flexibility or "adaptability" in apportionment. Nevertheless, the Court cannot say, as a matter of law, that Wi-LAN's failure to use the SSPPU in its reasonable royalty analysis requires either a new trial or remittitur on damages. Rather, whether Apple is entitled to that relief depends on whether the damages theory Wi-LAN did present to the jury was the product of a reliable methodology, and if so, whether that methodology was reliably applied to the facts of this case.

/ / /

---

[4] In exceptional cases, the entire market value of the product may be used, but "'only where the patented feature creates the basis for customer demand or substantially creates the value of the component parts.'" *Versata Software, Inc. v. SAP America, Inc.*, 717 F.3d 1255, 1268 (Fed. Cir. 2013) (quoting *SynQor, Inc. v. Artesyn Tech., Inc.*, 709 F.3d 1365, 1383 (Fed. Cir. 2013)). There is no dispute that requirement is not met in this case, and that the entire market value rule, therefore, does not apply.

Wi-LAN described its methodology in this case as apportionment through "direct valuation." Notably, Wi-LAN fails to cite any other case in which this methodology has been used to apportion the value of a patented invention as part of a reasonable royalty analysis. Nevertheless, Wi-LAN engaged three experts who each analyzed different factors as part of this methodology.

First, Wi-LAN had Dr. Madisetti study "the incremental benefits of the patented technologies and quantif[y] those technical benefits for each patent group, by comparing the Accused Products with the next-best noninfringing alternatives[.]" (Mem. of P. & A. in Opp'n to Apple's Mot. to Exclude Certain Opinions of Vijay Madisetti, David Kennedy, and Jeffrey Prince, ECF No. 352 at 9.) After doing so, Dr. Madisetti opined that voice over LTE ("VOLTE") capability increased a phone's Mean Opinion Score ("MOS") by 2.3 points.[5] In reaching that opinion, Dr. Madisetti relied on a report by Signals Ahead, which tested Samsung phones, that compared VOLTE technology with non-VOLTE technology Skype. Dr. Madisetti also conducted his own tests of VOLTE and non-VOLTE technology using iPhones. Dr. Madisetti also opined that claim 9 of the '145 Patent provided benefits in the form of 16% faster uploads, and that claim 1 of the '757 Patent provided benefits in the form of 6% faster downloads.

Professor Prince then took Dr. Madisetti's "benefits" opinions and assigned each of those purported benefits a monetary value through the use of his "willingness to pay" survey. Through that survey, Professor Prince determined the value of VOLTE technology was in the range of $69-$121, the value of increased upload speed was between $1.90 and $3.65, and the value of increased download speed was between $2.44 and $4.02.

Mr. Kennedy then took Dr. Madisetti's "benefits" opinions and Professor's Prince's valuations of those benefits to arrive at a reasonable royalty figure of $145

[5] During opening statement, Wi-LAN's counsel described VOLTE as the process of "sending voice calls over the current 4G or LTE networks which transmits the call, not through the telephone company but over the internet." (Trial Tr. at 14, July 23, 2018, ECF No. 452.)

million. That figure was based on a royalty rate of 85 cents per unit (iPhone), multiplied by the number of iPhones sold during the period of infringement (170.7 million). Mr. Kennedy explained his royalty rate by reference to Professor Prince's valuation numbers. He specifically relied on the low end valuation for the upload and download speeds ($1.90 and $2.44, respectively), but apportioned only 1% of the upper end valuation of VOLTE ($121) to Wi-LAN for a total valuation of that technology of $1.22. Mr. Kennedy used these valuations to argue the reasonableness of his 85 cents per unit royalty rate. Mr. Kennedy also used the [redacted] to show the reasonableness of his 85 cent royalty rate. He also relied on Wi-LAN's license agreement with Samsung, which provided for a lump sum payment of [redacted] to Wi-LAN, as further support for the reasonableness of his ultimate damages figure. When asked why he did not use the SSPPU as his royalty base, particularly the baseband processor, Mr. Kennedy stated that was not required because Dr. Madisetti and Professor Prince valued the patented technology. In essence, he testified that he apportioned the patented features of the iPhone through Dr. Madisetti's and Professor Prince's "direct valuation" of those features.

The problem with this approach, however, specifically as it relates to claim 26 of the '145 Patent, was that Dr. Madisetti's starting point was VOLTE, not the patented technology. This, despite the testimony of Mr. Stanwood, one of the inventors of the '145 Patent, who stated he did not invent VOLTE. (Trial Tr. at 187:14-16, July 24, 2018, ECF No. 514.)

Apple argues Wi-LAN's use of VOLTE as a starting point overstated the footprint of the invention, and that the expert testimony incorporating that argument was therefore inadmissible. *See ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 869 (Fed. Cir. 2010) (stating "trial court must carefully tie proof of damages to the claimed invention's footprint in the market place.") Wi-LAN disputes that it drew a connection between the patented technology and VOLTE and therefore overstated the footprint of the invention. (Wi-LAN's Opp'n to Apple's Mot. at 18.) However, its arguments and the evidence presented at trial refute that contention.

First, Wi-LAN used VOLTE to prove infringement of claim 26 of the '145 Patent. (Trial Tr. at 614, July 26, 2018, ECF No. 506 ("Q: ... A VOLTE to VOLTE call, that is the technology that's enabled by the '145 Patent? A: Yes.").) Although that may have been appropriate, taking that theory and simply importing it into the damages case was not.

Second, Wi-LAN does not dispute that to determine the benefits of the invention claimed in the '145 Patent, Dr. Madisetti relied primarily on a Signals Ahead test of voice call quality *using VOLTE* compared to voice call quality *using Skype*. It is undisputed this test did not mention the '145 Patent or equate its benefits with the voice call quality of VOLTE, but Dr. Madisetti drew that connection anyway. Relying on that unproven connection, Dr. Madisetti then extrapolated from the Signals Ahead test that the benefit of the invention claimed in the '145 Patent was a 2.3 unit increase to the MOS score. (*See* Trial Tr. at 266-67, July 24, 2018, ECF No. 514 ("Q: So using the patented invention, the mean opinion score, the MOS, is almost twice as good as with Skype, the alternative? A: It's actually 2.3 MOS units, so the quality is quite significant. It could be even more than twice."); *id.* at 267 (Professor Madisetti stating he did his own study of "the benefits of VOLTE over Skype" using iPhones).

Third, Wi-LAN's counsel also drew this connection in questions to their witnesses. (*See id.* at 271-72 ("Q: Apple's use of these three inventions in the patent claims in the accused iPhones, do they improve the iPhone as a whole for voice and cellular data? A: Yes. Q: Better voice and higher speeds? A: Yes."); Trial Tr. at 381, July 25, 2018, ECF No. 493 ("Q: At a high level, what are [the technical benefits of using the inventions in the accused product]? A: You get great quality from the VOLTE, that is 2.3 MOS better."); Trial Tr. at 614, July 26, 2018, ECF No. 506 ("Q: ... A VOLTE to VOLTE call, that is the technology that's enabled by the '145 Patent?

/ / /

A: Yes."). *See also* Rep. Tr. at 6, Nov. 30, 2018, ECF No. 547 (Wi-LAN's counsel stating "the 2.3 MOS score is still specific to the patented technology.")

These opinions and evidence were without factual basis. Indeed, they contradicted the testimony of Mr. Stanwood that he did not invent VOLTE, and the testimony of Mr. Kennedy that the '145 patented technology was "related to" VOLTE, not equivalent to it. (Trial Tr. at 686:10-12, July 26, 2018, ECF No. 506.) (*See also id.* at 686:13-24 (acknowledging "there are other pieces of value, lots of little pieces of value that go into VOLTE"); Trial Tr. at 14, July 23, 2018, ECF No. 452 ("VOLTE has a lot of components to it. It uses LTE networks, it uses the internet. It has a lot of different components that a lot of different people and companies invented. When you see how complex it is you will see that all had to come together.") Wi-LAN's assertion that Dr. Madisetti's "benefits" opinion was limited to the patented technology is also contradicted by Mr. Kennedy's subsequent apportionment of only 1% of VOLTE call quality to Wi-LAN. Mr. Kennedy stated he allocated 1% of this value to Wi-LAN and 99% of the value to Apple using a profit sharing analysis, "even though that's – all 121 of that is created by the Wi-LAN technology[.]" (Trial Tr. at 623:10-25, July 26, 2018, ECF No. 506.) Had Dr. Madisetti's opinion been limited to the "benefits" of the patented invention, there would have been no need for Mr. Kennedy to further apportion any value of VOLTE to Wi-LAN. That "benefit" should have been accounted for by Dr. Madisetti.[6]

Absent a sufficient factual basis, Dr. Madisetti's opinion about the "benefits" of claim 26 of the '145 Patent should not have been presented to the jury. *See LaserDynamics, Inc. v. Quanta Computer, Inc.*, 694 F.3d 51, 81 (Fed. Cir. 2012) (stating "new trial is required because the jury's verdict was based on expert opinion that finds no support in the facts in the record.") And since Dr. Madisetti's opinions were the basis for Professor Prince's opinions and Mr. Kennedy's opinions, those

---

[6] Mr. Kennedy did not apply a similar "profit sharing" deduction to the other two valuations, which further raises concerns about the reliability, perhaps, arbitrariness, of his opinions.

opinions also should have been excluded. *See Gbarabe v. Chevron Corp.*, No. 14-CV-00173-SI, 2017 WL 956628, at *17 (N.D. Cal. Mar. 13, 2017) (stating it would be "inappropriate" for expert to rely on another expert's "flawed and speculative report.")[7] Accordingly, the Court grants Apple's request for a conditional remittitur of $10 million. In the event Wi-LAN does not accept this remittitur, the Court grants Apple's motion for a new trial on damages.[8]

III.

CONCLUSION AND ORDER

For the reasons set forth above, the Court denies Apple's motion for judgment as a matter of law on the issue of infringement or damages, and denies Wi-LAN's motion for supplemental damages, ongoing royalty and prejudgment and post judgment interest. The Court grants Apple's motion for a conditional remittitur to $10 million, and orders the parties to appear for a settlement conference before Judge Major on **January 14, 2019**, at **10:30 a.m.** In the event Wi-LAN accepts the remittitur, a Notice of Acceptance of Remittitur must be filed by **January 18, 2019**. In the event Wi-LAN does not accept the remittitur, the Court grants Apple's motion for a new trial on damages, and will set a telephonic status conference with counsel to discuss dates.

IT IS SO ORDERED.

DATED: January 3, 2019

HON. DANA M. SABRAW
United States District Judge

---

[7] The Federal Rules of Evidence permit an expert to rely upon the opinions developed by another expert for the purpose of litigation if the expert independently verifies the underlying expert's work. *Fosmire v. Progressive Max Ins. Co.*, 277 F.R.D. 625, 630 (W.D. Wash. 2011). However, there is no evidence Professor Prince or Mr. Kennedy independently verified Dr. Madisetti's "benefits" opinions in this case.

[8] In light of this ruling, the Court denies Wi-LAN's motion for supplemental damages, ongoing royalty and prejudgment and post judgment interest.

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC.,

Plaintiff,

vs.

WI-LAN, INC.,

Defendant.

AND ALL RELATED COUNTERCLAIMS.

CASE NO. 14cv2235 DMS (BLM)

ORDER DENYING WI-LAN'S MOTION FOR RECONSIDERATION

On January 4, 2019, this Court issued an Order granting Apple's motion for a new trial on damages. (ECF No. 554.) Wi-LAN now moves for reconsideration of that decision. Apple filed an opposition to the motion, and Wi-LAN filed a reply.

"Reconsideration is appropriate if the district court (1) is presented with newly discovered evidence, (2) committed clear error or the initial decision was manifestly unjust, or (3) if there is an intervening change in controlling law." *School Dist. No. 1J, Multnomah County, Oregon v. ACandS, Inc.*, 5 F.3d 1255, 1263 (9th Cir. 1993). Here, Wi-LAN relies on the second prong, and argues the Court's decision to grant a new trial on damages was clearly erroneous and manifestly unjust. Specifically, Wi-LAN asserts there was substantial evidence to support the jury's damages verdict, and thus the Court's decision to grant a new trial on damages was both clearly erroneous and manifestly unjust. Wi-LAN also argues Dr. Madisetti properly measured the technical benefit of the invention recited in claim 26 of the '145 Patent, and the Court's decision to the contrary was clearly erroneous.

The burden to show a decision is clearly erroneous is a high one, "which by design is difficult to meet." *United States v. Perkins*, 850 F.3d 1109, 1125 (9th Cir. 2017). "'To be clearly erroneous, a finding must be more than possibly or even probably wrong; the error must be pellucid to any objective observer.'" *United States v. Christie*, 825 F.3d 1048, 1058 (9th Cir. 2016) (quoting *United States v. Quaintance*, 608 F.3d 717, 721 (10th Cir. 2010)) (internal quotation marks omitted). Stated otherwise, "'to be clearly erroneous, a decision must ... strike us as wrong with the force of a five-week old, unrefrigerated dead fish.'" *Ocean Garden, Inc. v. Marktrade Co., Inc.*, 953 F.2d 500, 502 (9th Cir. 1991) (quoting *Parts and Elec. Motors, Inc. v. Sterling Elec., Inc.*, 866 F.2d 228, 233 (7th Cir. 1988)). "This stringent standard 'rests on good sense and the desire to protect both court and parties against the burdens of repeated reargument by indefatigable diehards.'" *Alaimalo v. United States*, 645 F.3d 1042, 1060 (9th Cir. 2011) (quoting 18B Charles A. Wright *et al.*, *Federal Practice and Procedure* § 4478 (2d ed. 2002)).

Here, Wi-LAN argues the Court committed clear error in finding that Dr. Madisetti's opinion about the benefits of claim 26 of the '145 Patent lacked a sufficient factual basis. However, this argument is essentially a rehash of arguments Wi-LAN has raised in previous motions, namely, that Dr. Madisetti did not equate the value of claim 26 of the '145 Patent with VOLTE. Wi-LAN attempts to put a finer point on that argument here by drawing a distinction between VOLTE, in general terms, and one aspect of VOLTE, namely "improved voice quality during loading[,]" (Mot. at 10), but at its core this argument is simply a different shade of the same argument Wi-LAN has been making consistently in this case. As such, it does not warrant reconsideration of the Court's prior finding. *See Brown v. Kinross Gold, U.S.A.*, 378 F.Supp.2d 1280, 1288 (D. Nev. 2005) ("A motion for reconsideration is not an avenue to re-litigate the same issues and arguments upon which the court already has ruled.") (citing *Brogdon v. Nat'l Healthcare Corp.*, 103 F.Supp.2d 1322, 1338 (N.D. Ga. 2000)).

Furthermore, the Court does not agree with the premise of Wi-LAN's argument, namely, that Dr. Madisetti confined his opinions to that aspect of VOLTE associated with improved voice quality "during loading." As Apple points out, Dr. Madisetti testified repeatedly that the benefit of claim 26 of the '145 Patent, generally, was "much higher quality calls." (Trial Tr. at 259:20-23, July 24, 2018, ECF No. 504; *see also id.* at 261:22-24 ("And then you compare the quality of voice with VOLTE and with Skype then that gives you an idea of how much benefits Apple would have using these iPhone.")) Those opinions were not confined to improved voice quality "during loading." Thus, this argument does not show the Court's prior finding about Dr. Madisetti's opinion is clearly erroneous.

Wi-LAN's only other argument is that there was substantial evidence to support the jury's damages verdict, and thus the Court's decision to grant a new trial on damages was clearly erroneous and manifestly unjust. As an initial matter, the Court notes that Wi-LAN did not raise this "substantial evidence" argument in its opposition to Apple's motion for a new trial on damages. Indeed, the lead case in Wi-LAN's motion for reconsideration, *Landes Construction Co. v. Royal Bank of Canada*, 833 F.2d 1365 (9th Cir. 1987), is nowhere cited in Wi-LAN's opposition to Apple's motion for a new trial on damages. Wi-LAN's failure to make this specific argument in its prior brief is reason enough to deny the motion for reconsideration. *See Garber v. Embry-Riddle Aeronautical Univ.*, 259 F.Supp.2d 979, 982 (D. Ariz. 2003) ("[N]ew arguments and new legal theories that could have been made at the time of the original motion may not be offered in a motion for reconsideration.")

Even considering the merits of the argument, it does not warrant reconsideration of the Court's previous order. Contrary to Wi-LAN's argument, the jury in this case was not presented with two alternative theories of damages. Rather, the jury was presented with one theory: A reasonable royalty. The evidence Wi-LAN relies on to

CONFIDENTIAL MATERIAL OMITTED

support its assertion that there was another theory of damages, *i.e.*, [redacted], the rate sheets and the infrastructure analysis, were all part of that theory. Indeed, all of that evidence was used primarily as a "check" against Mr. Kennedy's opinion that [redacted] was a reasonable royalty in this case. It did not form the basis for an alternative theory of damages.[1]

Furthermore, the Court is persuaded that admission of this evidence, combined with Dr. Madisetti's improper opinion on the benefits of claim 26 of the '145 Patent, "skew[ed] the damages horizon for the jury." *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1320 (Fed. Cir. 2011). This is especially so with respect to [redacted] and the rate sheets. For instance, although Mr. Kennedy testified that [redacted] was "probative" to the hypothetical negotiation analysis, he did not say [redacted] was comparable to the hypothetical license the parties would have agreed to in this case, (Trial Tr. at 693-94, July 26, 2018), which was a prerequisite to its admissibility. *See LaserDynamics, Inc. v. Quanta Computer, Inc.*, 694 F.3d 51, 78-81 (Fed. Cir. 2012) (granting new trial because damages testimony relied on licenses that were not comparable and therefore not relevant). Indeed, Mr. Kennedy testified it was *not* the similarities between [redacted] that made [redacted] relevant, but [redacted] that was "most helpful." (*Id.* at 693-94.) Admission of the rate sheets was similarly prejudicial. *See Whitserve, LLC v. Computer Packages, Inc.*, 694 F.3d 10, 29-30 (Fed. Cir. 2012) (acknowledging that although "proposed licenses may have some value for determining a reasonable royalty in certain situations[,]" the evidentiary value of proposed licenses is limited by "the fact that patentees could artificially inflate the royalty rate by making outrageous offers.")

---

[1] The absence of an alternative theory of damages takes this case outside the holding of *Landes*. In that case, unlike here, the plaintiff actually presented two "alternative calculations of damages" to the jury, one based on lost profits and another based on "the difference between the purchase price and fair market value[.]" 833 F.2d at 1372-73. In light of those two theories, the court concluded "that proper respect for the role of the jury and the discretion of the trial judge favors construing a general verdict in behalf of the prevailing party." *Id.* at 1373. That presumption does not apply here, where the jury was presented with only one theory of damages.

Under these circumstances, the Court cannot say its decision to grant Apple's motion for a new trial on damages was either clearly erroneous or manifestly unjust. *See United States v. 99.66 Acres of Land*, 970 F.2d 651, 658 (9th Cir. 1992) (stating new trial is warranted "on the basis of an incorrect evidentiary ruling if the ruling substantially prejudiced a party.") Accordingly, Wi-LAN's motion for reconsideration is denied.

IT IS SO ORDERED.

DATED: March 26, 2019

HON. DANA M. SABRAW
United States District Judge

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC., Plaintiff, v. WI-LAN, INC., Defendant.	Case No.: 14cv2235 DMS (BLM) **ORDER GRANTING IN PART AND DENYING IN PART APPLE'S MOTION TO EXCLUDE THE TESTIMONY OF WI-LAN'S EXPERTS**
AND ALL RELATED COUNTERCLAIMS.	

This case comes before the Court on Apple’s motion to exclude the testimony of Wi-LAN’s experts.[1] The motion came on for hearing on July 31, 2019. Sean Cunningham appeared and argued for Apple, and Scott Cole and Warren Lipschitz appeared and argued for Wi-LAN. After the hearing, Wi-LAN submitted an unsolicited supplemental brief, to which Apple filed a response. Wi-LAN also submitted a Notice of Outstanding Dispute to alert the Court that a pending discovery motion before the Magistrate Judge “may bear” on

[1] Also pending before the Court is Wi-LAN’s omnibus motion to exclude expert testimony. The Court will address that motion in a separate order.

this ruling. After thoroughly reviewing the parties' briefs, the record in this case and the relevant legal authority, the Court grants in part and denies in part Apple's motion.

I.

BACKGROUND

This is the parties' second round of *Daubert* motions in this case. Before the trial, Apple moved to exclude the testimony of Wi-LAN's experts, Vijay Madisetti, Ph.D., Jeffrey Prince and David Kennedy pursuant to Federal Rule of Evidence 702 and *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579, 589 (1993). (*See* ECF No. 572.) In that first motion, Apple argued Dr. Madisetti's opinions about the alleged benefits of the asserted claims should have been excluded for a number of reasons. First, Apple argued Dr. Madisetti's opinions about the call quality benefits of the '761 Patent and the '145 Patent should have been excluded because they were based on insufficient facts and data[,]" (*id.* at 10), they were "an unjustified extrapolation of the limited data" he did cite, (*id.* at 11), and his "testing" of two iPhones was "unreliable because his test cannot be challenged in any objective sense." (*Id.* at 12.) Second, Apple asserted Dr. Madisetti's opinions about the upload and download speed benefits of the other asserted claims should have been excluded because "they hinge[d] solely on the alleged benefits of LTE," which this Court found not to infringe. Apple claimed that to the extent Dr. Prince and Mr. Kennedy relied on Dr. Madisetti's benefits opinions, their opinions should also be excluded. Apple also raised a separate challenge to Mr. Kennedy's opinions on the ground he failed to "'apportion the defendant's profits and the patentee's damages between the patented feature and the unpatented features[.]'" *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1318 (Fed. Cir. 2011) (quoting *Garretson v. Clark*, 111 U.S. 120, 121 (1884)). In response to that motion, Wi-LAN argued Apple's concerns with Dr. Madisetti's benefits opinions went to weight, not admissibility. (*See* ECF No. 352.) Specifically, Wi-LAN characterized Apple's arguments as raising an issue about the propriety of Dr. Madisetti using Samsung phones as the "benchmark for quantifying the technical benefits delivered by" the asserted patents. (*Id.* at 13-14.) Wi-LAN also disputed Apple's assertion that Mr.

CONFIDENTIAL MATERIAL OMITTED

Kennedy failed to apportion. It argued Mr. Kennedy apportioned through Dr. Madisetti's benefits opinions and Dr. Prince's corresponding willingness-to-pay survey and "at every step" thereafter. (*Id.* at 17-19.) Wi-LAN specifically defended Mr. Kennedy's use of the entire iPhone in his reasonable royalty analysis, and his reliance on a license agreement between Wi-LAN and [redacted] and a Wi-LAN Rate Sheet. (*Id.* at 17-20.) After reviewing the parties' briefs on these issues and hearing oral argument from counsel, the Court agreed with Wi-LAN that Apple's arguments went to weight rather than admissibility, and denied Apple's motion.

The parties then proceeded to trial where Wi-LAN presented its damages theory to the jury.[2] Dr. Madisetti presented his benefits opinions, Dr. Prince presented his survey evidence and Mr. Kennedy presented his overall opinion that Wi-LAN was entitled to damages in the range of $145 million. (Trial Tr. at 612:6-7, July 26, 2018, ECF No. 506.) Mr. Kennedy's opinion was based, in part, on Dr. Madisetti's benefits opinions, namely that the asserted claims resulted in increased call quality and faster upload and download speeds, as well as Dr. Prince's survey evidence about consumers' willingness to pay for those features. (*Id.* at 621-22.) Mr. Kennedy took that survey evidence, specifically, the amount consumers were willing to pay for the alleged benefits of the asserted claims, and applied a "profit sharing percentage" to those numbers "to determine what portion of that should be considered a royalty for Wi-LAN …." (*Id.* at 623.) As to the increase in call quality, which was valued at $122 by the survey participants, Mr. Kennedy took that number "and allocated just 1 percent of that to Wi-LAN" to arrive at a royalty rate of $1.22 per iPhone. (*Id.*) In the words of Wi-LAN's counsel, that $1.22 was "the value of the voice express lane patent that we have." (*Id.* at 624.) Mr. Kennedy used the same approach to value the benefits of the other asserted claims to arrive at valuations of $1.90 per iPhone

---

[2] By the time of trial, the only claims at issue were claims 9, 26 and 27 of the '145 Patent and claim 1 of the '757 Patent. Despite Wi-LAN's decision to proceed on only those claims, its damages numbers remained the same.

CONFIDENTIAL MATERIAL OMITTED

for the timer feature of the '145 Patent and $2.44 per iPhone for the modulation feature of the '757 Patent. (*Id.*) Mr. Kennedy then used those valuations to support the reasonableness of his 85 cent royalty rate. (*Id.* at 625.)

That was not the end of Mr. Kennedy's analysis, however. He also conducted an infrastructure analysis in which he calculated Apple's costs to improve the cellular phone infrastructure as an alternative to Apple entering into a license with Wi-LAN for the asserted patents. (*Id.* at 625-26.) Notably, Mr. Kennedy admitted that changes to infrastructure would not improve voice quality. (*Id.* at 627-28.) However, he stated Apple could obtain the benefits of the other claims through infrastructure improvements. Mr. Kennedy first calculated a total cost for those improvements, and then divided that total cost by the number of units to arrive at a "per handset cost for infrastructure now and the per handset cost for infrastructure if Apple said we'll just add infrastructure." (*Id.* at 627.) Those costs were $5.49 for the timer feature, and $2.06 for the modulation feature. (*Id.* at 627-28.) Mr. Kennedy then used those figures to, again, support the reasonableness of his 85 cent royalty rate. (*Id.* at 629.)

Mr. Kennedy then turned to Wi-LAN's Rate Sheets, which reflected licensing rates of [redacted] and [redacted]. (*Id.* at 631.) He then went through some of Wi-LAN's licenses, specifically its licenses with [redacted], Doro, Unecto and Vertu. (*Id.* at 635.) Mr. Kennedy spent the most time on Wi-LAN's Agreement with [redacted] in particular, [redacted]. (*Id.* at 650-51.) Mr. Kennedy relied on all of this evidence to support his royalty rate of 85 cents per iPhone and his total damages amount of $145 million.

After seven days of evidence and argument, the case was submitted to the jury. Less than ninety minutes after receiving the case, the jury returned a verdict for Wi-LAN in the amount of $145.1 million.

In post-trial briefing, Wi-LAN moved for supplemental damages, prejudgment and post judgment interest and an ongoing royalty, and Apple moved for judgment as a matter of law and/or for a new trial. After hearing oral argument on these motions, the Court denied Apple's motion for judgment as a matter of law, but granted Apple's motion for a conditional remittitur to $10 million, or in the alternative, for a new trial on damages. In analyzing the damages issue, the Court started with the undisputed proposition that apportionment was required, and that in cases like this involving multi-component products, apportionment is generally accomplished:

> by ensuring the royalty base is not "larger than the smallest salable unit embodying the patented invention." [*Power Integrations, Inc. v. Fairchild Semiconductor Int'l, Inc.*, 904 F.3d 965, 977 (Fed. Cir. 2018), cert. denied, ___ U.S. ___, 139 S.Ct. 1265 (2019))]. If the SSPPU "itself contains several non-infringing features[,]" the patentee must apportion further by "estimat[ing] what portion of that smallest salable unit is attributed to the patented technology." *Id.* (citing *VirnetX, Inc. v. Cisco Systems, Inc.*, 767 F.3d 1308, 1327 (Fed. Cir. 2014).

(ECF No. 548 at 4-5.) The Court recognized Wi-LAN did not use this methodology, but declined to find that its failure to do so warranted exclusion of Wi-LAN's damages theory. Instead, the Court considered whether Wi-LAN's methodology, which Wi-LAN described as "direct valuation," was reliable and reliably applied to the facts of this case. The Court found it was not, for two reasons. First, the Court found Dr. Madisetti's benefits opinion on claim 26 of the '145 Patent was unreliable because he used VoLTE as a starting point for apportionment, rather than the patented technology, despite the lack of any evidence equating the two and, indeed, evidence directly disputing that connection. The Court also found Mr. Kennedy's subsequent apportionment of Dr. Madisetti's benefits opinions was arbitrary, which further undermined the reliability of Wi-LAN's methodology. Because Dr. Prince and Mr. Kennedy relied on Dr. Madisetti's opinions in coming to their conclusions, the Court held their opinions were also inadmissible, and Apple was therefore entitled to a remittitur of $10 million.

Wi-LAN declined to accept that remittitur, and soon thereafter filed a motion for reconsideration of the Court's ruling. In that motion, Wi-LAN argued there was substantial evidence in the record to support the jury's damages award other than Dr. Madisetti's

CONFIDENTIAL MATERIAL OMITTED

opinion about the benefits of claim 26 of the '145 Patent. Wi-LAN also argued the Court's decision about Dr. Madisetti's benefits opinion was clearly erroneous. The Court rejected both of those arguments. In doing so, the Court disagreed with Wi-LAN's assertion that the jury was presented with two alternative theories of damages. Rather, the Court found the jury was presented with one damages theory based on a reasonable royalty. On that theory, the Court considered evidence other than Dr. Madisetti's benefits opinion, specifically, the [redacted], Wi-LAN's Rate Sheets and the infrastructure analysis, and found that evidence "'skew[ed] the damages horizon for the jury.' *Uniloc USA, Inc. v. Microsoft Corp.*, 632 F.3d 1292, 1320 (Fed. Cir. 2011)." (ECF No. 619 at 4.) Specifically, the Court found the [redacted] should not have been admitted because Wi-LAN failed to show it was comparable to the hypothetical license the parties would have agreed to in this case, and that Wi-LAN's Rate Sheets should not have been admitted because they were more prejudicial than probative.

In light of that ruling, the Court set the case for retrial on the issue of damages. After receiving briefs from the parties as to the scope of the retrial, the Court allowed the parties' experts to prepare supplemental reports and for the parties to take limited discovery on those reports. That Order specifically stated, "Any supplemental reports shall be confined to damages theories set out in the expert's initial reports, i.e., the parties may not raise any new damages theories in these supplemental expert reports." (ECF No. 624.)

The parties' experts prepared supplemental reports, and, save for some outstanding discovery motions, discovery is now complete. For its part, Wi-LAN's experts used the same methodology they used for the first trial. Dr. Madisetti has again attempted to isolate the benefits of the asserted claims of the patents, but this time he has equated the benefits of the voice quality invention with improved voice quality during loading.[3] Dr. Prince has recalculated consumers' willingness to pay for those modified benefits using the survey

[3] Dr. Madisetti did not modify his analysis or opinions on the upload and download speed features.

data from the first trial, and Mr. Kennedy has again applied a profit-sharing percentage to those total numbers to arrive at a value for the infringing features. This time, Mr. Kennedy did not apply a 99/1 profit-sharing percentage, but rather applied four factors to arrive at a different profit-sharing percentage, and a different value for the voice quality feature. With these modifications, Mr. Kennedy again opines that Wi-LAN is entitled to damages in the range of $145 million.

II.

DISCUSSION

"In *Daubert*, the Supreme Court set out the requirements for admissibility of expert testimony." *Summit 6, LLC v. Samsung Electronics Co., Ltd.*, 802 F.3d 1283, 1295 (Fed. Cir. 2015). "The Supreme Court stated that the trial judge plays a 'gatekeeping role,' which 'entails a preliminary assessment of whether the reasoning or methodology underlying the testimony is scientifically valid and of whether that reasoning or methodology properly can be applied to the facts in issue.'" *Id.* (quoting *Daubert*, 509 U.S. at 592-93). "The Court emphasized that the focus 'must be solely on principles and methodology, not on the conclusions that they generate.'" *Id.* (quoting *Daubert*, 509 U.S. at 595).

> This admissibility assessment, while a flexible one, may consider the following factors: (1) whether the methodology is scientific knowledge that will assist the trier of fact; (2) whether the methodology has been tested; (3) whether the methodology has been published in peer-reviewed journals; (4) whether there is a known, potential rate of error; and (5) whether the methodology is generally accepted.

Id. (quoting *Daubert*, 509 U.S. at 591-95). Under *Daubert* and Federal Rule of Evidence 702, "a district court may exclude evidence that is based upon unreliable principles or methods, legally insufficient facts and data, or where the reasoning or methodology is not sufficiently tied to the facts of the case." *Id.* In determining whether an approach is admissible, the critical question is "whether the methodology employed is reliable." *Summit 6*, 802 F.3d at 1296 (citing *Daubert*, 509 U.S. at 589-95). "[T]he question of whether the expert is credible or the opinion is correct is generally a question for the fact finder, not the court." *Id.* at 1296 (citing *Apple Inc. v. Motorola, Inc.*, 757 F.3d 1286, 1314 (Fed. Cir. 2014)). "Indeed, '[v]igorous cross-examination, presentation of contrary evidence, and careful instruction on the burden of proof are the traditional and appropriate means of attacking shaky but admissible evidence.'" *Id.* (quoting *Daubert*, 509 U.S. at 596).

Here, the *Daubert* issue relates to damages, and more particularly, the reasonable royalty method of calculating damages for patent infringement. The Federal Circuit "has recognized that estimating a reasonable royalty is not an exact science[,]" and "there may be more than one reliable method for estimating a reasonable royalty." *Id.* (citing *Apple*, 757 F.3d at 1315). For instance, "[a] party may use the royalty rate from sufficiently comparable licenses, value the infringed features based upon comparable features in the marketplace, or value the infringed features by comparing the accused product to non-infringing alternatives." *Id.* (citing *Apple*, 757 F.3d at 1315). The most common approach is the hypothetical negotiation approach, which "attempts to ascertain the royalty upon which the parties would have agreed had they successfully negotiated an agreement just before infringement began." *Lucent Techs., Inc. v. Gateway, Inc.*, 580 F.3d 1301, 1324 (Fed. Cir. 2009). All of these approaches "'necessarily involve[ ] an element of approximation and uncertainty.'" *Id.* at 1325 (quoting *Unisplay, S.A. v. American Electronic Sign Co., Inc.*, 69 F.3d 512, 517 (Fed. Cir. 1995)). "[G]iven the great financial incentive parties have to exploit the inherent imprecision in patent valuation, courts must be proactive to ensure that the testimony presented—using whatever methodology—is sufficiently reliable to support a damages award." *Commonwealth Scientific and Indus. Research Org. v. Cisco Systems, Inc. ("CSIRO")*, 809 F.3d 1295, 1301 (Fed. Cir. 2015). Furthermore, courts "must have considerable leeway in deciding how to determine whether the expert's testimony is sufficiently reliable." *Power Integrations*, 711 F.3d at 1373 (citing *Kumho Tire Co., Ltd. v. Carmichael*, 526 U.S. 137, 152 (1999)).

The touchstone for reliability, and hence admissibility, of a reasonable royalty analysis is whether it "'carefully tie[s] proof of damages to the claimed invention's footprint in the market place.'" *Uniloc*, 632 F.3d at 1317 (quoting *ResQNet.com, Inc. v. Lansa, Inc.*, 594 F.3d 860, 869 (Fed. Cir. 2010)). "As a substantive matter, it is the 'value of what was taken' that measures a 'reasonable royalty' under 35 U.S.C. § 284." *Ericsson, Inc. v. D-Link Systems, Inc.*, 773 F.3d 1201, 1226 (Fed. Cir. 2014). And in cases such as this, "where multi-component products are involved, the governing rule is that the ultimate combination of royalty base and royalty rate must reflect the value attributable to the infringing features of the product, and no more." *Id.* (citing *VirnetX*, 767 F.3d at 1326). "In each case, district courts must assess the extent to which the proffered testimony, evidence, and arguments would skew unfairly the jury's ability to apportion the damages to account only for the value attributable to the infringing features." *Id.* at 1228.

As set out above, Wi-LAN uses a three-pronged approach to arrive at a reasonable royalty rate in this case. That approach involves the use of Dr. Madisetti's benefits analysis, Mr. Kennedy's infrastructure analysis and Mr. Kennedy's analysis of Wi-LAN's licensing history. Apple moves to exclude all three of these approaches. The Court has discussed these approaches in its previous orders, but as the parties' arguments and the record have evolved, so has the Court's understanding of Wi-LAN's methodology and its application to the facts of this case. With the benefit of this expanded record, the Court addresses each of these approaches below.

A. The Benefits Methodology

The benefits methodology is Wi-LAN's lead approach in its damages case. Wi-LAN is relying on this approach for all of the claims at issue: claims 26 and 27 of the '145 Patent (the voice quality claims), claim 9 of the '145 Patent (the timer/uploading claim) and claim 1 of the '757 Patent (the modulation/downloading claim). Apple's arguments on this methodology have focused primarily on the voice quality claims, and thus, the Court's orders have also focused on that aspect of this methodology. However, Apple has argued consistently that Wi-LAN's benefits analysis of the upload and download speed claims should also be excluded. Accordingly, the Court will address both aspects of this analysis.

1. Upload and Download Speeds

Turning first to the upload and download speed claims, Apple argues the analysis of the benefits of these claims is hopelessly flawed because it relies on products that do not practice the claimed inventions. More specifically, Apple argues the products that were tested use the LTE standard rather than the VoLTE standard, which is the standard Wi-LAN relied on to prove infringement. In all of the briefs on this issue, Wi-LAN does not dispute this assertion. Indeed, in its opposition to Apple's first *Daubert* motion, it argued that Apple should have raised this issue by way of summary judgment, not through a *Daubert* motion. (*See* ECF No. 352 at 9.) Its subsequent opposition briefs relied on that initial response and the Court's silence on this particular aspect of Wi-LAN's benefits methodology. In essence, Wi-LAN offered no substantive response to Apple's argument.

After reviewing all of the parties' briefs, however, the Court agrees with Apple that this analysis was unreliable. In order to be relevant to the valuation of the asserted claims, the tested products need to practice the claimed inventions. The products Dr. Madisetti relied on did not do so. Rather than practicing the VoLTE standard, which was the standard accused of infringement, the tested products practiced the LTE standard, and iPhones practicing that standard were found not to infringe Wi-LAN's patents. Because this methodology was not reliably applied to the facts of this case, it must be excluded.

2. Voice Quality

Turning to the voice quality claims, Wi-LAN has described the methodology to determine the benefits of these claims as one of "direct valuation," and has explained that the methodology involves "building" the value of the invention "from the ground up." (ECF No. 690 at 38.) This is a unique approach to determining the value of an invention, particularly in a case like this involving a multi-component product. Indeed, neither the Court nor the parties have identified another case where this kind of "bottom-up" approach to valuation has been used with these kinds of facts. Rather, it appears the more traditional approach to determining valuation does just the opposite: It starts at the top with the accused product and works down from there to determine the value of the patented invention. This presumptive approach is the reason for the apportionment requirement and the principles surrounding that requirement, namely the SSPPU and the entire market value rule.

Wi-LAN's decision to turn this approach on its head has lead to confusion and a lack of clarity as to exactly how its approach achieves the goal of valuing the patented invention. Indeed, it begs the question why Wi-LAN would create a new and novel approach to calculating damages when there is already an established and accepted practice for doing so. The facts here are not unusual, and it is curious why a new and novel approach is necessary.

Nevertheless, Wi-LAN acknowledges the general rule of apportionment and the need to apportion in this case, and asserts it has done so. Apple disagrees, and complains that Wi-LAN is improperly relying on the entire market value rule, which Wi-LAN disputes. Indeed, Wi-LAN's response to Apple's apportionment arguments appears to be either that its unique approach does not require apportionment in the traditional sense, and/or that apportionment is inherent in its bottom-up approach. Regardless of which argument Wi-LAN is relying on, its position appears to be that neither the SSPPU nor the entire market value rule apply here given this methodology.

Notably, application of these principles to the facts of this case would not support a finding that Wi-LAN's methodology is reliable. Apple has argued repeatedly that Wi-LAN has not used the SSPPU in this case, and although Wi-LAN disputes that argument, it fails to explain why the baseband processor was the "subscriber unit" for its infringement case, but the iPhone is the measuring stick for its damages case. Also, Wi-LAN does not dispute that the infringing features of the accused product do not drive demand, which is a prerequisite to application of the entire market value rule. Thus, applying these principles does not advance Wi-LAN's position that this methodology is reliable. But Wi-LAN fares no better in avoiding them. Both of these principles, the SSPPU and the entire market value rule, are meant to aid the Court "in determining when an expert's apportionment model is reliable." *CSIRO*, 809 F.3d at 1302. Without them, the Court is left without two firmly established markers of reliability.

Lacking those markers, the Court is left to look elsewhere for indicia that Wi-LAN's benefits methodology is reliable, and was reliably applied to the facts of this case. As the proponent of this methodology, Wi-LAN "must explain the expert's methodology and demonstrate in some objectively verifiable way that the expert has both chosen a reliable scientific method and followed it faithfully." *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 43 F.3d 1311, 1319 n.11 (9th Cir. 1995). It has not met that burden here for several reasons. First, Wi-LAN can point to no other case in which this methodology has been used. Second, the starting point for this methodology, Dr. Madisetti's opinion about the benefits of the voice quality inventions, was based on a Signals Ahead report that tested Samsung phones, not the accused products. Third, Dr. Madisetti's own tests of the voice quality features of the accused products were not conducted pursuant to any scientific method.

Wi-LAN claims that Apple conducted its own tests of the voice quality features of the accused products, or products substantially similar thereto, and the results of those tests support the Signals Ahead report and Dr. Madisetti's test results,[4] but as the Court has stated in its previous orders, the problem here is Wi-LAN's correlation of the benefits of its inventions with improved voice quality, either with VoLTE in general or under loading conditions in particular, without sufficient evidence to support that correlation. As reflected in the record, numerous companies and numerous patents contributed to the VoLTE standard and the corresponding improvements in the technology. Wi-LAN continues to assert that its patents, and the particular claims being asserted here, caused these improvements and benefits, but that is not a valid assumption based on the evidence. Wi-LAN's revised methodology does not address the Court's concerns with these

---

[4] These tests are the subject of the pending motion before the Magistrate Judge.

underlying, unwarranted assumptions, specifically, that the improved voice quality with VoLTE is attributable to the asserted claims. Instead, Wi-LAN makes a different assumption, namely that its inventions caused the improvements and benefits in voice quality *during loading*, but that assumption, too, is plagued by the same problems and concerns.

Rather than addressing these concerns, Wi-LAN simply takes these problematic assumptions and wraps them in Dr. Prince's survey evidence and Mr. Kennedy's resulting valuation of the inventions. However, "[b]eginning from a fundamentally flawed premise and adjusting it based on legitimate considerations specific to the facts of the case nevertheless results in a fundamentally flawed conclusion." *Uniloc*, 632 F.3d at 1317. Based on all of the factors discussed above, the Court finds Wi-LAN's benefits methodology is not reliable and was not reliably applied to the facts of this case. Accordingly, the Court grants Apple's motion to exclude this methodology.

B. The Infrastructure Analysis

Wi-LAN's second methodology involves Mr. Kennedy's infrastructure analysis. Apple argues this analysis should be excluded as unreliable because it is not apportioned to the value of the inventions and is speculative and unscientific. The Court agrees.

As with the benefits methodology discussed above, neither the Court nor the parties have identified any other case in which the value of an invention has been measured by the cost to improve the infrastructure in the relevant field. Although the determination of a reasonable royalty "encompasses fantasy and flexibility," *Fromson v. Western Litho Plate and Supply Co.*, 853 F.2d 1568, 1575 (Fed. Cir. 1988), Mr. Kennedy's infrastructure analysis is entirely speculative. Indeed, he himself admitted that this alternative was "not economically feasible[.]" (ECF No. 506 at 628) ("So obviously, you know, Apple would not be arguing for this as an alternative, but Wi-LAN would be pointing it out to Apple, hey, here's one thing you could do, but you really can't. It's just not economically feasible to do.") Because Mr. Kennedy's infrastructure analysis is based entirely on "subjective

CONFIDENTIAL MATERIAL OMITTED

belief or unsupported speculation[,]" *Daubert*, 509 U.S. at 590, Apple's motion to exclude this methodology is granted.

C. The Licensing History Analysis

Wi-LAN's third and final methodology for calculating damages involves an analysis of Wi-LAN's licensing history. Here, Apple does not challenge the reliability of the methodology itself, but rather challenges Wi-LAN's application of the methodology to the facts of this case. Specifically, Apple takes issue with Wi-LAN's use of its Rate Sheets, the [redacted] and its licenses with Doro, Unecto and Vertu.

On the Rate Sheets, the Court has already held their admission was prejudicial. (*See* ECF No. 619 at 4.) The Court reached the same conclusion on the [redacted]. (*Id.*)

The Court also found Wi-LAN did not show the [redacted] was comparable to the hypothetical license the parties would have negotiated in this case, and that it was therefore inadmissible. (*Id.*) Wi-LAN ignores this finding and persists in arguing that comparability is a jury issue, not a basis for exclusion. However, the Court disagrees with that argument. Although the jury may decide what weight to give to the licenses that are admitted, whether the licenses are sufficiently comparable, and hence, admissible, is a question for the Court. *See Laser Dynamics*, 694 F.3d at 78-81 (granting new trial because damages testimony relies on licenses that were not comparable and therefore not relevant).

Wi-LAN also argues the [redacted] is a "real-world" license, and therefore admissible. However, the Court disagrees with that argument, as well. Presumably, all of the licenses the parties rely on in the hypothetical negotiation analysis are "real-world" licenses, but that does not automatically render them admissible. They must still be comparable to the hypothetical license being negotiated by the parties, and the [redacted]

Case: 20-2011 Document: 17 Page: 150 Filed: 11/20/2020

CONFIDENTIAL MATERIAL OMITTED

[redacted].

Even if the [redacted] was comparable, the Court would still find it inadmissible because the way it is being used, [redacted], "skew[s] unfairly the jury's ability to apportion the damages to the infringing features." *Ericsson*, 773 F.3d at 1228. As Mr. Kennedy admitted during his testimony at the first trial, the [redacted].

The only other licenses Wi-LAN relies on in its licensing history analysis are those with Doro, Unecto and Vertu. Apple argues Mr. Kennedy failed to apportion these licenses, just as he did with the [redacted], therefore Mr. Kennedy's opinions on these licenses must also be excluded. Wi-LAN again disagrees with Apple on the law, asserting there is no requirement that experts "unpack" prior licenses.

The Federal Circuit recognizes that prior licenses "are almost never perfectly analogous to the infringement action." *Id.* at 1227 (citing *VirnetX*, 767 F.3d at 1330).

However, it has also stated, "[t]estimony relying on licenses must account for such distinguishing facts when invoking them to value the patented invention." *Id.* Here, Mr. Kennedy pointed out to the jury that Doro, Unecto and Vertu were much smaller companies than Apple, but that was the only distinguishing feature he called out. Notably, he failed to account for the fact that these licenses were to Wi-LAN's entire patent portfolio, not just the patents, or four claims, at issue in this case. Clearly, this is a distinction that needs to be accounted for. However, the Court cannot say Mr. Kennedy's failure to do so renders his opinions on these licenses inadmissible. As stated above, the focus of the *Daubert* inquiry is on the methodology, not the expert's conclusions or opinions. The issues surrounding the Doro, Unecto and Vertu licenses are more properly the subject of cross-examination rather than a basis for excluding Mr. Kennedy's opinions. Accordingly, the Court denies Apple's motion to exclude this aspect of Mr. Kennedy's testimony.

III.

CONCLUSION

For the reasons set out above, the Court grants in part and denies in part Apple's motion. Specifically, the Court denies the motion as to Mr. Kennedy's opinions on the Doro, Unecto and Vertu licenses, and grants the remainder of the motion.

IT IS SO ORDERED.

Dated: September 30, 2019

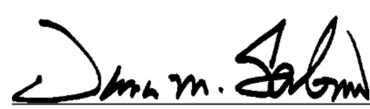

Hon. Dana M. Sabraw
United States District Judge

UNITED STATES DISTRICT COURT

SOUTHERN DISTRICT OF CALIFORNIA

APPLE INC.,

Plaintiff,

v.

WI-LAN, INC.,

Defendant.

AND ALL RELATED COUNTERCLAIMS.

Case No.: 14cv2235 DMS (BLM)

ORDER(1) DENYING APPLE INC.'S RENEWED MOTION FOR JUDGMENT AS A MATTER OF LAW AND/OR MOTION FOR NEW TRIAL AND (2) GRANTING WI-LAN'S MOTION FOR PRE-JUDGMENT AND POST-JUDGMENT INTEREST

This case comes before the Court on Apple Inc.'s renewed motion for judgment as a matter of law and/or motion for a new trial, and Wi-LAN's motion for pre-judgment and post-judgment interest. Wi-LAN opposes Apple's motion. Apple does not oppose Wi-LAN's request for post-judgment interest, but does oppose Wi-LAN's request for pre-judgment interest. After reviewing the parties' briefs, the record, and the relevant legal authority, the Court denies Apple's motion and grants Wi-LAN's motion.

I.

APPLE'S MOTION

Apple moves for judgment as a matter of law and/or a new trial following the retrial on damages. As stated in this Court's previous orders, "[a] Rule 50(b) motion for judgment as a matter of law is not a freestanding motion. Rather, it is a renewed Rule 50(a) motion." *E.E.O.C. v. Go Daddy Software, Inc.*, 581 F.3d 951, 961 (9th Cir. 2009). Federal Rule of Civil Procedure 50(a)(1) provides:

> 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 the controlling law, can be maintained or defeated only with a favorable finding on that issue.

Fed. R. Civ. P. 50(a)(1). In the Ninth Circuit, "'[j]udgment as a matter of law is appropriate when the evidence presented at trial permits only one reasonable conclusion.'" *Torres v. City of Los Angeles*, 548 F.3d 1197, 1205 (9th Cir. 2008) (quoting *Santos v. Gates*, 287 F.3d 846, 851 (9th Cir. 2002)). "In other words, '[a] motion for a judgment as a matter of law is properly granted only if no reasonable juror could find in the non-moving party's favor.'" *Id.* (quoting *El-Hakem v. BJY Inc.*, 415 F.3d 1068, 1072 (9th Cir. 2005)). When considering a motion for judgment as a matter of law, the court must view the evidence "'in the light most favorable to the nonmoving party, and all reasonable inferences must be drawn in favor of that party.'" *Id.* at 1205-06 (quoting *LaLonde v. County of Riverside*, 204 F.3d 947, 959 (9th Cir. 2000)).

Federal Rule of Civil Procedure 59 provides: "The court may, on motion, grant a new trial on all or some of the issues-and to any party-as follows: (A) after a jury trial, for any reason for which a new trial has heretofore been granted in an action at law in federal court[.]" Fed. R. Civ. P. 59(a)(1)(A). "A trial court should grant a motion for a new trial if (1) the jury instructions were erroneous or inadequate, (2) the court made incorrect and prejudicial admissibility rulings, or (3) the verdict is contrary to the great weight of the evidence." *Chiron Corp. v. Genentech, Inc.*, 363 F.3d 1247, 1258 (Fed. Cir. 2004) (citations omitted).

Here, Apple raises many of the same arguments on damages it has raised throughout this case, namely that Wi-LAN failed to apportion its damages, its evidence was speculative and arbitrary, and its damages case was not based on the smallest saleable unit. The Court has addressed these issues numerous times in this case, and declines to rehash the parties' arguments or its own analysis here. Suffice it to say, Apple has not met the standard for either judgment as a matter of law or a new trial and/or remittitur on damages. The Court cannot say a reasonable jury did not have a sufficient evidentiary basis for the verdict in this case, or that there are grounds for a new trial. Accordingly, Apple's motion is denied.

II.

WI-LAN'S MOTION

Wi-LAN moves for pre-judgment interest at the California statutory rate of seven percent and post-judgment interest pursuant to 28 U.S.C. § 1961(a). Apple does not oppose Wi-LAN's request concerning post-judgment interest, but does oppose Wi-LAN's request for pre-judgment interest at the California statutory rate. Apple argues pre-judgment interest should be calculated at the one-year U.S. Treasury Bill rate, which Wi-LAN asserts is an average of one percent. Apple also contends Wi-LAN should not receive pre-judgment interest between the time of the first jury verdict and the judgment because Wi-LAN caused that delay.

In a recent decision, the Federal Circuit stated district courts have "'wide latitude in the selection of interest rates,'" and that permissible interest rates include "statutory rates set by states, U.S. Treasury bill rate, and the prime rate." *Schwendimann v. Arkwright Advanced Coating, Inc.*, 959 F.3d 1065, 1076 (Fed. Cir. 2020). In other patent cases, this Court has used the California statutory rate because it has better fulfilled the purpose of pre-judgment interest, which is "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." *General Motors Corp. v. DevexCorp.*, 461 U.S. 648, 655 (1983). *See KFx Medical Corp. v. Arthrex, Inc.*, Case No. 11cv1698 DMS (BLM), ECF No. 348; *Carl Zeiss*

Vision Int'l GmbH v. Signet Armorlite, Inc., Case No. 07cv0894 DMS (DHB), ECF No. 1558. The Court finds no reason to deviate from that conclusion in this case. Accordingly, the Court grants Wi-LAN's request for pre-judgment interest at the California statutory rate. The Court also rejects Apple's argument that pre-judgment interest should not be awarded for the time between the first jury verdict and the judgment. Contrary to Apple's assertion, any delay between the first jury verdict and the judgment cannot be attributed solely to Wi-LAN. All counsel, parties and the Court had scheduling issues that impacted the pace of the litigation.

III.

CONCLUSION AND ORDER

For the reasons set forth above, the Court denies Apple's motion for judgment as a matter of law and/or for a new trial, and grants Wi-LAN's motion for pre-judgment and post-judgment interest. The Court will enter a separate judgment consistent with this Order, the other orders of this Court and the jury verdicts.

IT IS SO ORDERED.

Dated: June 15, 2020

Hon. Dana M. Sabraw
United States District Judge

UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF CALIFORNIA
SAN DIEGO

WI-LAN INC., *Plaintiff*, vs. APPLE INC., *Defendant*.	No.: 3:14-cv-1507-DMS (BLM); (Lead Case No. 3:14-cv-2235-DMS-BLM) DEMAND FOR JURY TRIAL **FINAL JUDGMENT** Department: 13A Judge: Hon. Dana M. Sabraw Magistrate Judge: Hon. Barbara L. Major

Pursuant to Rule 58 of the Federal Rules of Civil Procedure, the Court enters judgment as follows:

1. Claims 9, 26, and 27 of U.S. Patent No. 8,457,145 ("the '145 Patent") and Claim 1 of U.S. Patent No. 8,537,757 ("the '757 patent) are not invalid.

2. Apple has directly infringed claims 9, 26, and 27 of the '145 Patent and claim 1 of the '757 Patent for the products accused in this litigation, namely the iPhone 7, iPhone 7 Plus, iPhone 6s, iPhone 6s Plus, iPhone SE, iPhone 6, and iPhone 6 Plus ("Infringing Products").

3. Damages in the amount of $85.23 million (at a rate of $0.45 per unit for 189.4 million units) are hereby awarded to Wi-LAN for Apple's infringement of claims 9, 26, and 27 of the '145 Patent and claim 1 of the '757 Patent occurring through May 21, 2019 for the Infringing Products.

4. For Apple's counterclaims and defenses as to the '145 and '757 Patents, judgment is entered in favor of Wi-LAN.

5. All other claims, counterclaims, and defenses as to other previously asserted patents are dismissed with prejudice.

6. Wi-LAN is entitled to pre-judgment interest at the California statutory rate of seven percent per annum from September 25, 2014 through the entry of this judgment on June 16, 2020.

7. Pre-judgment interest in the amount of $23.747 million is hereby awarded to Wi-LAN.

8. Wi-LAN is entitled to post judgment interest at the rate set forth in 28 U.S.C. § 1961 as of this date judgment is entered until the date this judgment is satisfied.

9. Wi-LAN shall file its bill of costs[1] the later of (i) 60 days after entry of a mandate from the Federal Circuit, (ii) 60 days after termination of the appeal,

[1] To the extent that Wi-LAN determines that it will file a motion for attorney's fees it shall be due at the same time as the bill of costs.

or (iii) 60 days after the time for filing an appeal has expired.

10. As to the iPhone 8, 8 Plus, X, Xs, Xs Max, and Xr sold by Apple through May 21, 2019, the parties are working on a separate stipulation and agreement to provide royalties consistent with this judgment and are reserving their respective rights.

ALL RELIEF NOT EXPRESSLY GRANTED HEREIN, OR RESERVED UNTIL FURTHER PROCEEDINGS, IS DENIED. THIS IS A FINAL JUDGMENT.

IT IS SO ORDERED.

Dated: June 16, 2020

HON. DANA M. SABRAW
United States District Court Judge

US008457145B2

(12) **United States Patent**
Zimmerman et al.

(10) **Patent No.: US 8,457,145 B2**
(45) **Date of Patent: Jun. 4, 2013**

(54) **METHOD AND APPARATUS FOR BANDWIDTH REQUEST/GRANT PROTOCOLS IN A WIRELESS COMMUNICATION SYSTEM**

(75) Inventors: **Ofer Zimmerman**, Modiin (IL); **Kenneth L. Stanwood**, Cardiff by the Sea, CA (US); **Brian Spinar**, Poway, CA (US); **Yair Bourlas**, San Diego, CA (US); **Amir Serok**, Tel Aviv (IL)

(73) Assignee: **Wi-Lan, Inc.**, Ottawa (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/487,032**

(22) Filed: **Jun. 1, 2012**
(Under 37 CFR 1.47)

(65) **Prior Publication Data**

US 2012/0243480 A1 Sep. 27, 2012

Related U.S. Application Data

(63) Continuation of application No. 12/415,698, filed on Mar. 31, 2009, now abandoned, which is a continuation of application No. 10/978,903, filed on Nov. 1, 2004, now Pat. No. 7,529,193, which is a continuation of application No. 09/783,671, filed on Feb. 14, 2001, now abandoned, which is a continuation-in-part of application No. 09/316,518, filed on May 21, 1999, now Pat. No. 6,925,068, and a continuation-in-part of application No. 09/613,434, filed on Jul. 11, 2000, now Pat. No. 6,785,252, which is a continuation-in-part of application No. 09/316,518, filed on May 21, 1999, now Pat. No. 6,925,068.

(51) **Int. Cl.**
H04W 28/20 (2009.01)

(52) **U.S. Cl.**
USPC **370/412**; 370/468

(58) **Field of Classification Search**
USPC 370/252, 412, 468
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,949,404 A 4/1976 Fletcher et al.
4,464,767 A 8/1984 Bremer

(Continued)

FOREIGN PATENT DOCUMENTS

EP 0353759 A2 2/1990
EP 0507384 A2 10/1992

(Continued)

OTHER PUBLICATIONS

United States District Court, Southern District of Florida. *Wi-LAN USA, Inc.* and *Wi-LAN, Inc.*, Plaintiffs, v. *HTC Corporation*, and *HTC America, Inc.*, Defendants. "Complaint." Entered Dec. 6, 2012. Case No. 1:12-cv-24319-KMM. 6 pages.

(Continued)

Primary Examiner — Melvin Marcelo
(74) *Attorney, Agent, or Firm* — Procopio, Cory, Hargreaves & Savitch LLP

(57) **ABSTRACT**

A method and apparatus for allocating bandwidth in a broadband wireless communication system is disclosed. One embodiment uses a self-correcting bandwidth request/grant protocol. The self-correcting bandwidth request/grant protocol utilizes a combination of incremental and aggregate bandwidth requests. CPEs primarily transmit incremental bandwidth requests to their associated base stations, followed by periodic transmissions of aggregate bandwidth requests. The use of periodic aggregate bandwidth requests (that express the current state of their respective connection queues) allows the bandwidth allocation method and apparatus to be "self-correcting". Another embodiment utilizes an abridged bandwidth request/grant protocol to allocate bandwidth. The abridged bandwidth request/grant protocol system utilizes padding packets to request a reduction in bandwidth allocation to a CPE. A base station modem alerts a base station CPU when the BS modem receives a padding packet from a CPE. After alerting the BS CPU the method can reduce the associated CPE's bandwidth allocation.

29 Claims, 17 Drawing Sheets

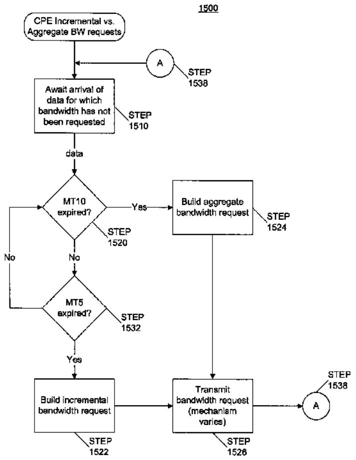

U.S. PATENT DOCUMENTS

4,495,619	A	1/1985	Acampora
4,924,461	A	5/1990	Amemiya et al.
5,130,983	A	7/1992	Heffner, III
5,297,144	A	3/1994	Gilbert et al.
5,349,342	A	9/1994	Nilles et al.
5,349,580	A	9/1994	Hester et al.
5,404,374	A	4/1995	Mullins et al.
5,412,651	A	5/1995	Gorshe
5,420,851	A	5/1995	Seshadri et al.
5,444,696	A	8/1995	Petranovich
5,444,698	A	8/1995	Kito
5,479,447	A	12/1995	Chow et al.
5,506,848	A	4/1996	Drakopoulos et al.
5,511,082	A	4/1996	How et al.
5,517,503	A	5/1996	Hess
5,560,021	A	9/1996	Vook et al.
5,592,470	A	1/1997	Rudrapatna et al.
5,594,720	A	1/1997	Papadopoulos et al.
5,594,738	A	1/1997	Crisler et al.
5,596,576	A	1/1997	Milito et al.
5,596,577	A	1/1997	Perreault et al.
5,615,212	A	3/1997	Ruszczyk et al.
5,634,006	A	5/1997	Baugher et al.
5,638,371	A	6/1997	Raychaudhuri et al.
5,638,374	A	6/1997	Heath
5,640,395	A	6/1997	Hamalainen et al.
5,644,576	A	7/1997	Bauchot et al.
5,673,031	A	9/1997	Meier
5,675,573	A	10/1997	Karol et al.
5,677,909	A	10/1997	Heide
5,710,762	A	1/1998	Petranovich
5,729,531	A	3/1998	Raith et al.
5,732,078	A	3/1998	Arango et al.
5,751,708	A	5/1998	Eng et al.
5,768,254	A	6/1998	Papadopoulos et al.
5,787,080	A	7/1998	Hulyalkar et al.
5,818,820	A	10/1998	Anderson et al.
5,818,828	A	10/1998	Packer et al.
5,818,845	A	10/1998	Moura et al.
5,828,695	A	10/1998	Webb
5,859,619	A	1/1999	Wu et al.
5,886,989	A	3/1999	Evans et al.
5,886,995	A	3/1999	Arsenault et al.
5,889,963	A	3/1999	Gopal et al.
5,890,055	A	3/1999	Chu et al.
5,896,561	A	4/1999	Schrader et al.
5,898,925	A	4/1999	Honkasalo et al.
5,917,822	A	6/1999	Lyles et al.
5,926,476	A	7/1999	Ghaibeh
5,946,322	A	8/1999	Moura et al.
5,953,344	A	9/1999	Dail et al.
5,956,330	A	9/1999	Kerns
5,956,338	A	9/1999	Ghaibeh
5,956,642	A	9/1999	Larsson et al.
5,991,287	A	11/1999	Diepstraten et al.
6,005,853	A	12/1999	Wang et al.
6,005,857	A	12/1999	Honkasalo et al.
6,006,069	A	12/1999	Langston
6,016,311	A	1/2000	Gilbert et al.
6,016,313	A	1/2000	Foster, Jr. et al.
6,023,458	A	2/2000	Tweedy et al.
6,031,827	A	2/2000	Rikkinen et al.
6,031,845	A	2/2000	Walding et al.
6,038,223	A	3/2000	Hansson et al.
6,038,455	A	3/2000	Gardner et al.
6,047,189	A	4/2000	Yun et al.
6,049,549	A	4/2000	Ganz et al.
6,049,551	A	4/2000	Hinderks et al.
6,052,594	A	4/2000	Chuang et al.
6,061,330	A	5/2000	Johansson
6,075,787	A	6/2000	Bobeck et al.
6,094,421	A	7/2000	Scott
6,097,733	A	8/2000	Basu et al.
6,104,700	A	8/2000	Haddock et al.
6,112,080	A	8/2000	Anderson et al.
6,114,968	A	9/2000	Ramakrishnan et al.
6,115,390	A	9/2000	Chuah
6,137,787	A	10/2000	Chawla et al.
6,141,534	A	10/2000	Snell et al.
6,151,312	A	11/2000	Evans et al.
6,198,728	B1	3/2001	Hulyalkar et al.
6,212,196	B1	4/2001	Momona
6,212,200	B1	4/2001	Iizuka
6,216,006	B1	4/2001	Scholefield et al.
6,222,832	B1	4/2001	Proctor
6,249,526	B1	6/2001	Loukianov
6,262,980	B1	7/2001	Leung et al.
6,275,497	B1	8/2001	Varma et al.
6,282,187	B1	8/2001	Evans et al.
6,295,285	B1	9/2001	Whitehead
6,298,049	B1	10/2001	Vanhoof et al.
6,314,110	B1	11/2001	Chin et al.
6,324,184	B1	11/2001	Hou et al.
6,337,849	B1	1/2002	Smith et al.
6,359,863	B1	3/2002	Varma et al.
6,366,761	B1	4/2002	Montpetit
6,370,153	B1	4/2002	Eng
6,373,830	B1	4/2002	Ozluturk
6,374,112	B1	4/2002	Widegren et al.
6,381,228	B1	4/2002	Prieto, Jr. et al.
6,388,999	B1	5/2002	Gorsuch et al.
6,400,684	B1	6/2002	Benmohamed et al.
6,408,005	B1	6/2002	Fan et al.
6,424,656	B1	7/2002	Hoebeke
6,434,164	B1	8/2002	Matsunaga et al.
6,438,141	B1	8/2002	Hanko et al.
6,452,933	B1	9/2002	Duffield et al.
6,469,991	B1	10/2002	Chuah
6,470,016	B1	10/2002	Kalkunte et al.
6,501,745	B1	12/2002	Turina et al.
6,546,017	B1	4/2003	Khaunte
6,567,386	B1	5/2003	Sugaya et al.
6,567,416	B1	5/2003	Chuah
6,571,101	B1	5/2003	Schulz et al.
6,594,246	B1	7/2003	Jorgensen
6,615,039	B1	9/2003	Eldering
6,621,812	B1	9/2003	Chapman et al.
6,628,668	B1	9/2003	Hutzli et al.
6,636,485	B1	10/2003	Fijolek et al.
6,650,624	B1	11/2003	Quigley et al.
6,658,010	B1	12/2003	Enns et al.
6,665,264	B1	12/2003	Davison et al.
6,678,282	B2	1/2004	Sharper et al.
6,683,866	B1	1/2004	Stanwood et al.
6,693,891	B1	2/2004	Sugita et al.
6,693,892	B1	2/2004	Rinne et al.
6,704,932	B1	3/2004	Matsunaga et al.
6,731,600	B1	5/2004	Patel et al.
6,741,614	B1	5/2004	Porter et al.
6,754,196	B1	6/2004	Daane et al.
6,771,648	B1	8/2004	Kim et al.
6,771,962	B2	8/2004	Saifullah et al.
6,785,252	B1	8/2004	Zimmerman et al.
6,791,998	B1	9/2004	Yoshihara et al.
6,795,409	B1	9/2004	Youssefmir et al.
6,829,482	B2	12/2004	Rune et al.
6,842,437	B1	1/2005	Heath
6,862,622	B2	3/2005	Jorgensen
6,895,248	B1	5/2005	Akyol et al.
6,914,890	B1	7/2005	Tobita et al.
6,925,068	B1	8/2005	Stanwood et al.
6,937,566	B1	8/2005	Forslöw et al.
6,940,874	B2	9/2005	Ruszczyk et al.
6,944,148	B1	9/2005	Gehring et al.
6,956,834	B2	10/2005	Stanwood et al.
6,970,422	B1	11/2005	Ho et al.
7,006,530	B2	2/2006	Spinar et al.
7,047,009	B2	5/2006	Laroia et al.
7,050,407	B1	5/2006	Frazer et al.
7,103,065	B1	9/2006	Quigley et al.
7,116,682	B1	10/2006	Waclawsky et al.
7,177,323	B2	2/2007	Yavatkar et al.
7,333,495	B2	2/2008	Sala et al.
7,356,339	B2	4/2008	Nam
7,408,901	B1	8/2008	Narayanabhatla
7,418,007	B1	8/2008	Liu et al.
7,486,639	B2	2/2009	Stanwood et al.

7,512,154	B2	3/2009	Quigley et al.
7,548,534	B2	6/2009	Zimmerman et al.
7,562,130	B2	7/2009	Dillon et al.
7,636,571	B2	12/2009	Lee et al.
7,693,093	B2	4/2010	Riedel et al.
7,751,437	B2	7/2010	Spinar et al.
7,783,311	B2	8/2010	Takiishi et al.
7,809,373	B2	10/2010	Park et al.
7,817,666	B2	10/2010	Spinar et al.
7,962,139	B2	6/2011	Kangude et al.
8,027,298	B2	9/2011	Stanwood et al.
8,085,891	B2	12/2011	Owen
8,131,295	B2	3/2012	Wang et al.
8,189,514	B2	5/2012	Stanwood et al.
8,243,663	B2	8/2012	Spinar et al.
8,249,014	B2	8/2012	Stanwood et al.
8,249,051	B2	8/2012	Spinar et al.
2001/0038620	A1	11/2001	Stanwood et al.
2002/0080816	A1	6/2002	Spinar et al.
2004/0213197	A1	10/2004	Zimmerman et al.
2005/0047368	A1	3/2005	Kotzin et al.
2005/0089064	A1	4/2005	Zimmerman et al.
2005/0111409	A1	5/2005	Spear et al.
2006/0002336	A1	1/2006	Stanwood et al.
2008/0049678	A1	2/2008	Chindapol et al.
2008/0232342	A1	9/2008	Spinar et al.
2008/0232391	A1	9/2008	Spinar et al.
2008/0253394	A1	10/2008	Spinar et al.
2008/0268844	A1	10/2008	Ma et al.
2009/0168802	A1	7/2009	Spinar et al.
2009/0175235	A1	7/2009	Spinar et al.
2009/0207795	A1	8/2009	Spinar et al.
2010/0150093	A1	6/2010	Stanwood et al.
2010/0150094	A1	6/2010	Stanwood et al.
2010/0157928	A1	6/2010	Spinar et al.
2011/0249585	A1	10/2011	Stanwood et al.
2011/0249586	A1	10/2011	Stanwood et al.
2011/0249645	A1	10/2011	Spinar et al.
2011/0292904	A1	12/2011	Stanwood et al.
2012/0033634	A1	2/2012	Spinar et al.

FOREIGN PATENT DOCUMENTS

EP	0622924	A2	4/1994
EP	0720405	A2	7/1996
EP	0845916	A2	6/1998
EP	0891060	A2	1/1999
EP	1045559	A1	10/2000
GB	2320162		12/1996
JP	11122289	A	4/1999
WO	92/22162	A1	12/1992
WO	98/10568	A1	3/1998
WO	9837706	A2	8/1998
WO	9907170	A2	2/1999
WO	99/38343	A2	7/1999
WO	99/39532	A1	8/1999
WO	00/01188	A1	1/2000
WO	00/72626	A1	11/2000
WO	01/33772	A1	5/2001

OTHER PUBLICATIONS

United States District Court, Southern District of Florida. *Wi-LAN USA, Inc.* and *Wi-LAN, Inc.*, Plaintiffs, v. *Apple Inc.*, Defendant. "Complaint." Entered Dec. 6, 2012. Case No. 1:12-cv-24318-KMM. 6 pages.

United States District Court, Southern District of Florida. *Wi-LAN USA, Inc.* and *Wi-LAN, Inc.*, Plaintiffs, v. *Sierra Wireless America, Inc.*, Defendant. "Complaint." Entered Dec. 6, 2012. Case No. 1:12-cv-24320-KMM. 6 pages.

United States District Court, Southern District of Florida. *Wi-LAN USA, Inc.* and *Wi-LAN, Inc.*, Plaintiffs, v. *Telefonaktiebolaget LM Ericsson*, and *Ericsson Inc.*, Defendants. "Complaint." Entered Oct. 1, 2012. Case No. 1:12-cv-23569-DMM. 7 pages.

United States District Court, Southern District of Florida. *Wi-LAN USA, Inc.* and *Wi-LAN, Inc.*, Plaintiffs, v. *Alcatel-Lucent USA Inc.*, Defendant. "Complaint." Entered Oct. 1, 2012. Case No. 1:12-cv-23568-CMA. 7 pages.

U.S. District Court, Northern District of California, "Amended Complaint for Declaratory Judgment Demand for Jury Trial" filed Sep. 30, 2008 in Case No. 5:08-cv-4555.

Dutton et al. "Asynchronous Transfer Mode (ATM) Technical Overview." 2.sup.nd Edition, Prentice Hall Oct. 1995, Chap. 3, pp. 21-25.

Shannon. "A Mathematical Theory of Communication." Bell Systems Technical Journal, pp. 379-423 (Part I), 623-656 (Part II), Jul. 1948.

Conti et al. "E-DCP, an extension of the distributed-control polling MAC protocol (DCP) for integrated services." published 1994, pp. 711-719.

Elfeitori et al. "A MAC protocol for supporting real-time VBR traffic over IEEE 802.14 based HFC access networks." Published May 9-12, 1999, pp. 197-201.

Narlikar et al. "Designing Multihop Wireless Backhand Networks with Delay Guarantees." Bell Labs, 2005, 1-4244-0222-0/06/$20.00 2006 IEEE, 12 pages.

IEEE 802.14 Proposed Standard.

J. Pons and J. Dunlop. "Bit Error Rate Characterisation and Modelling for GSM", IEEE 1998, pp. 3722-3727.

Andersen et al. "Prediction of Future Fading Based on Past Measurements." Vehicular Technology Conference, VTC 1999, vol. 1, pp. 151-155.

J.M. Torrance, L. Hanzo. "Upper Bound Performance of Adaptive Modulation in a Slow Rayleigh Fading Channel." IEEE Electronics Letters. vol. 32, Apr. 1996, p. 718.

Kim et al. "PRMA/DA: A New Media Access Control Protocol for Wireless ATM." IEEE, Jun. 23, 1996, pp. 240-244.

L.H. Charles Lee. "Convolutional Coding, Fundamentals and Applications." Artech House, Inc., 1997, p. 11-51.

Lin et al. "Error Control Coding, Fundamentals and Applications." Prentice-Hall Computer applications in Electrical Engineering Series, 1993, pp. 315-349.

Lin. "On IEEE 802.14 Medium Access Control Protocol, IEEE Communications Surveys." http://www.comsoc.org/pubs/surveys, Fourth Quarter 1998, vol. 1 No. 1, pp. 1-10.

Mollenauer et al. "An Efficient Media Access Control Protocol for Broadband Wireless Access Systems." Submission to IEEE 802.16 Broadband Wireless Access Working Group, Oct. 28, 1999, pp. 1-18.

Mollenauer et al. "MAC Proposal for IEEE 802.16.1." Submission to IEEE 802.16 Broadband Wireless Access Working Group, Dec. 24, 1999, pp. 1-84.

Bender et al. "CDMA/HDR: A Bandwidth Efficient High Speed Wireless Data Service for Nomadic Users." Communications Magazine, IEEE, vol. 38, No. 7, Jul. 2000, pp. 70-77.

Jain. "On the Impact of Channel and channel Quality Estimation on Adaptive Modulation." Dec. 2002, pp. 1-92.

Papadopoulos et al. "Reduction of Mixed Co-channel Interference in Microcellular STDD Systems." Vehicular Technology Conference, 1995 IEEE 45.sup.th, vol. 2, pp. 759-763.

Qiu et al. "Dynamic Reservation Multiple Access (DRMA): A New Multiple Access Scheme for Personal Communication System (PCS)." Wireless Networks 2 (1996), pp. 117-128, J.C. Baltzer AG, Science Publishers.

Redl et al. "Introduction to GSM," Artech House, Inc., 1995; pp. 84, 85 and 95.

Sampei et al. "Adaptive Modulation/TDMA Scheme for Personal Multi-Media Communication Systems: Proceedings of the Global Telecommunications Conference." U.S., New York, IEEE, Nov. 28, 1994, pp. 989-993.

Superseded Data-Over-Cable Service Interface Specifications, Radio Frequency Interface Specification, SP-RFIv1.1-I01-990311 (DOCSIS1.1), dated Mar. 11, 1999, http://www.cablelabs.com/specifications/SP-RFIv.1.1-I01-990311.pdf, pp. i310, 320 pages.

Ue et al. "Symbol Rate and Modulation Level Controlled Adaptive Modulation/TDMA/TDD for Personal Communication Systems." Proceedings of the Vehicular Technology Conference, U.S., New York, IEEE, vol. Conf. 45, Jul. 25, 1995, pp. 306-310.

Ulm Eet al. "Data-Over-Cable Interface Specification, Radio Frequency Interface Specification." Hewlett Packard Interim Specification, Doc. Control No. SP-RFII01-970321, Published Mar. 21, 1997, by MCNS Holdings, L.P., Section 6 (Media Access Control Specification), pp. 43-85.

Wolf et al. "On the Weight Distribution of Linear Block Codes Formed From Convolutional Codes." IEEE, IEEE Transactions on Communications, vol. 44:9, Sep. 1996, pp. 1049-1051.

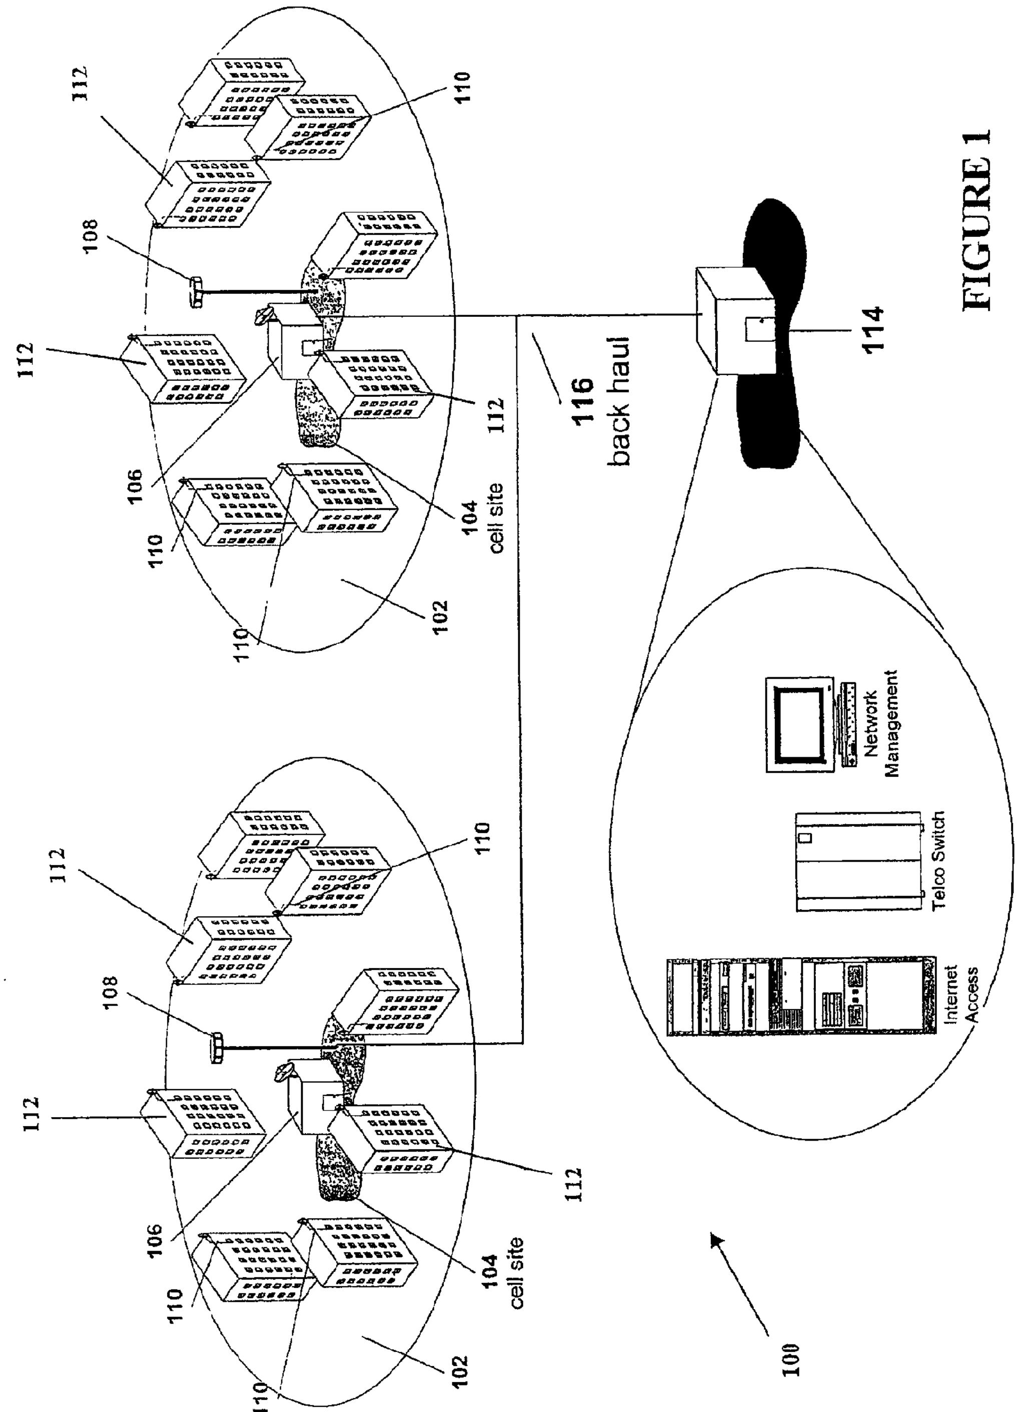
112
110
108
112
106
112
110
104
cell site
102
110
116
back haul
114
Network Management
Telco Switch
Internet Access
100

FIGURE 1

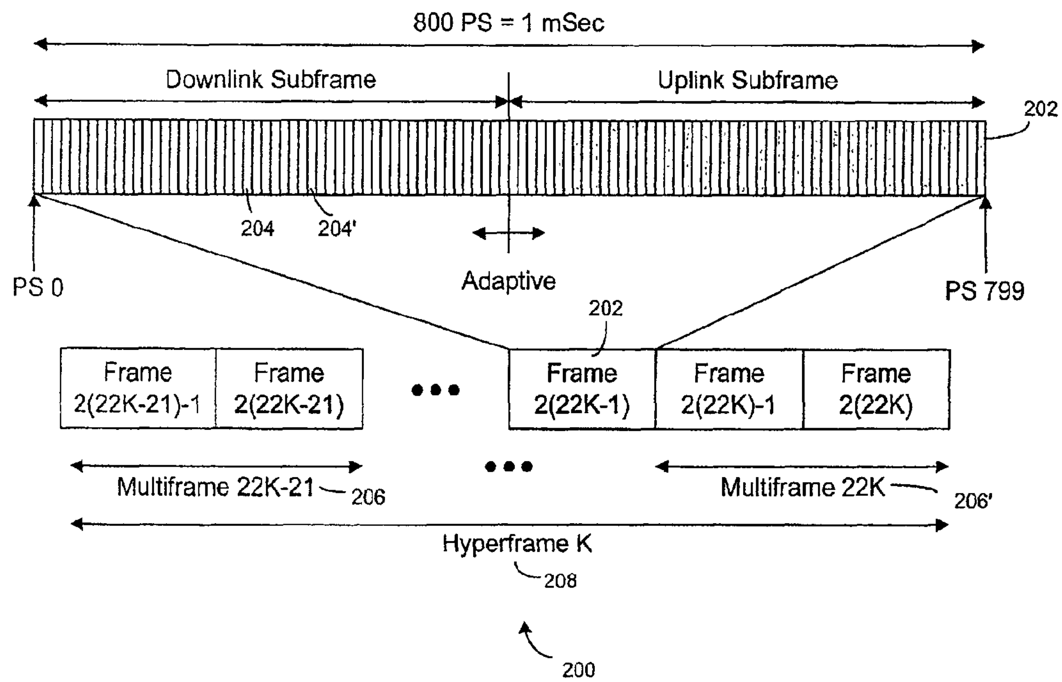
800 PS = 1 mSec
Downlink Subframe
Uplink Subframe
202
204
204'
PS 0
Adaptive
PS 799
202
Frame 2(22K-21)-1
Frame 2(22K-21)
Frame 2(22K-1)
Frame 2(22K)-1
Frame 2(22K)
Multiframe 22K-21
206
Multiframe 22K
206'
Hyperframe K
208
200

FIGURE 2

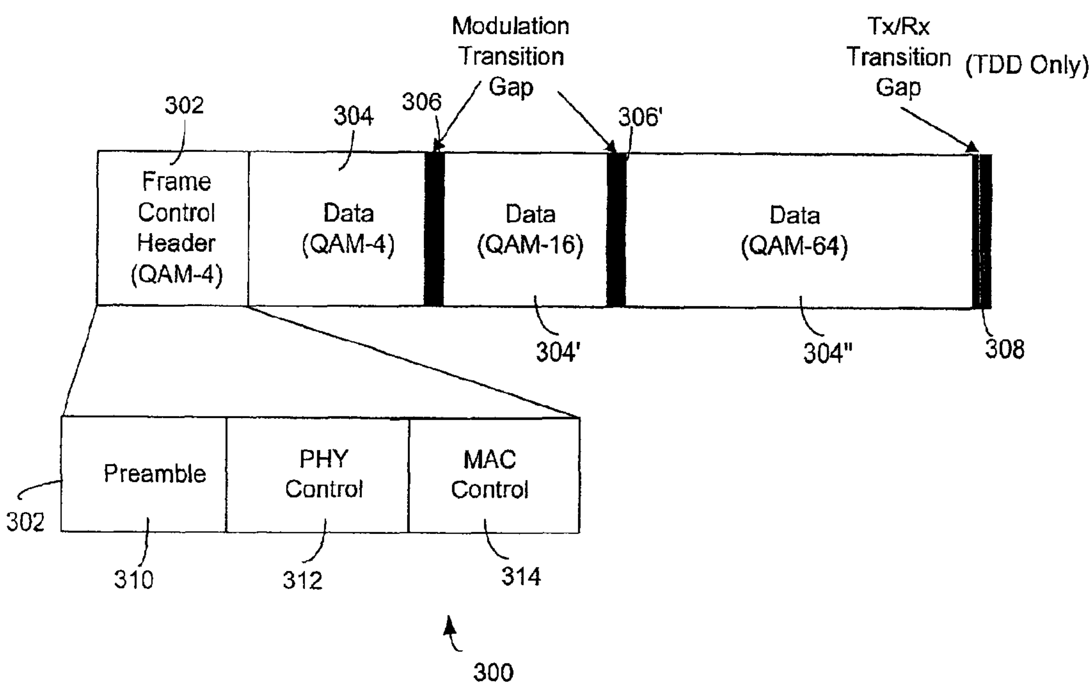
Modulation
Transition
Gap
Tx/Rx
Transition
Gap
(TDD Only)
302
304
306
306'
Frame
Control
Header
(QAM-4)
Data
(QAM-4)
Data
(QAM-16)
Data
(QAM-64)
304'
304"
308
Preamble
PHY
Control
MAC
Control
302
310
312
314
300

FIGURE 3

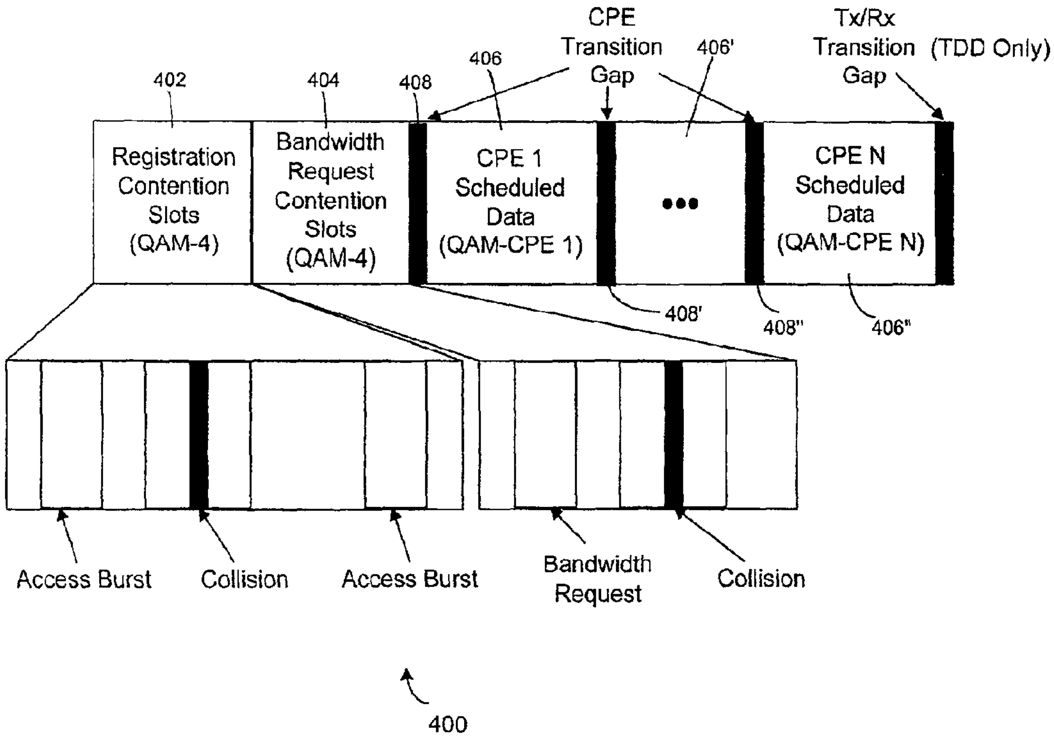
CPE
Transition
Gap
Tx/Rx
Transition (TDD Only)
Gap
402
404
408
406
406'
Registration
Contention
Slots
(QAM-4)
Bandwidth
Request
Contention
Slots
(QAM-4)
CPE 1
Scheduled
Data
(QAM-CPE 1)
CPE N
Scheduled
Data
(QAM-CPE N)
408'
408"
406"
Access Burst
Collision
Access Burst
Bandwidth
Request
Collision
400

FIGURE 4

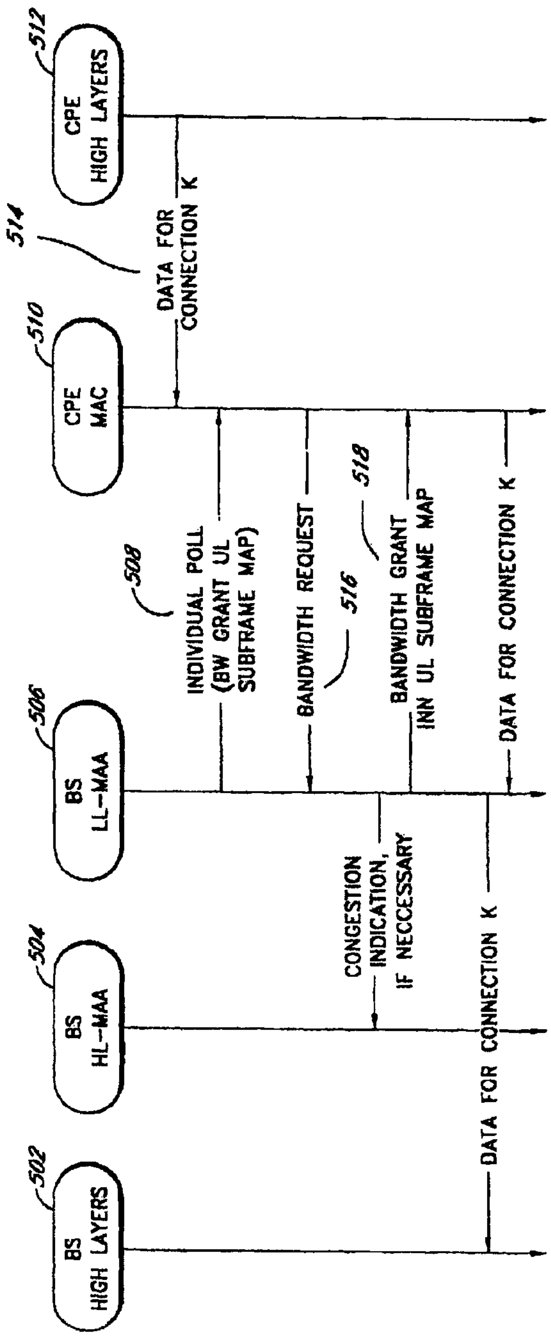
CPE HIGH LAYERS
512
CPE MAC
510
BS LL-MAA
506
BS HL-MAA
504
BS HIGH LAYERS
502
514
DATA FOR CONNECTION K
508
INDIVIDUAL POLL (BW GRANT UL SUBFRAME MAP)
516
BANDWIDTH REQUEST
518
BANDWIDTH GRANT INN UL SUBFRAME MAP
DATA FOR CONNECTION K
CONGESTION INDICATION, IF NECCESSARY
DATA FOR CONNECTION K

FIG. 5

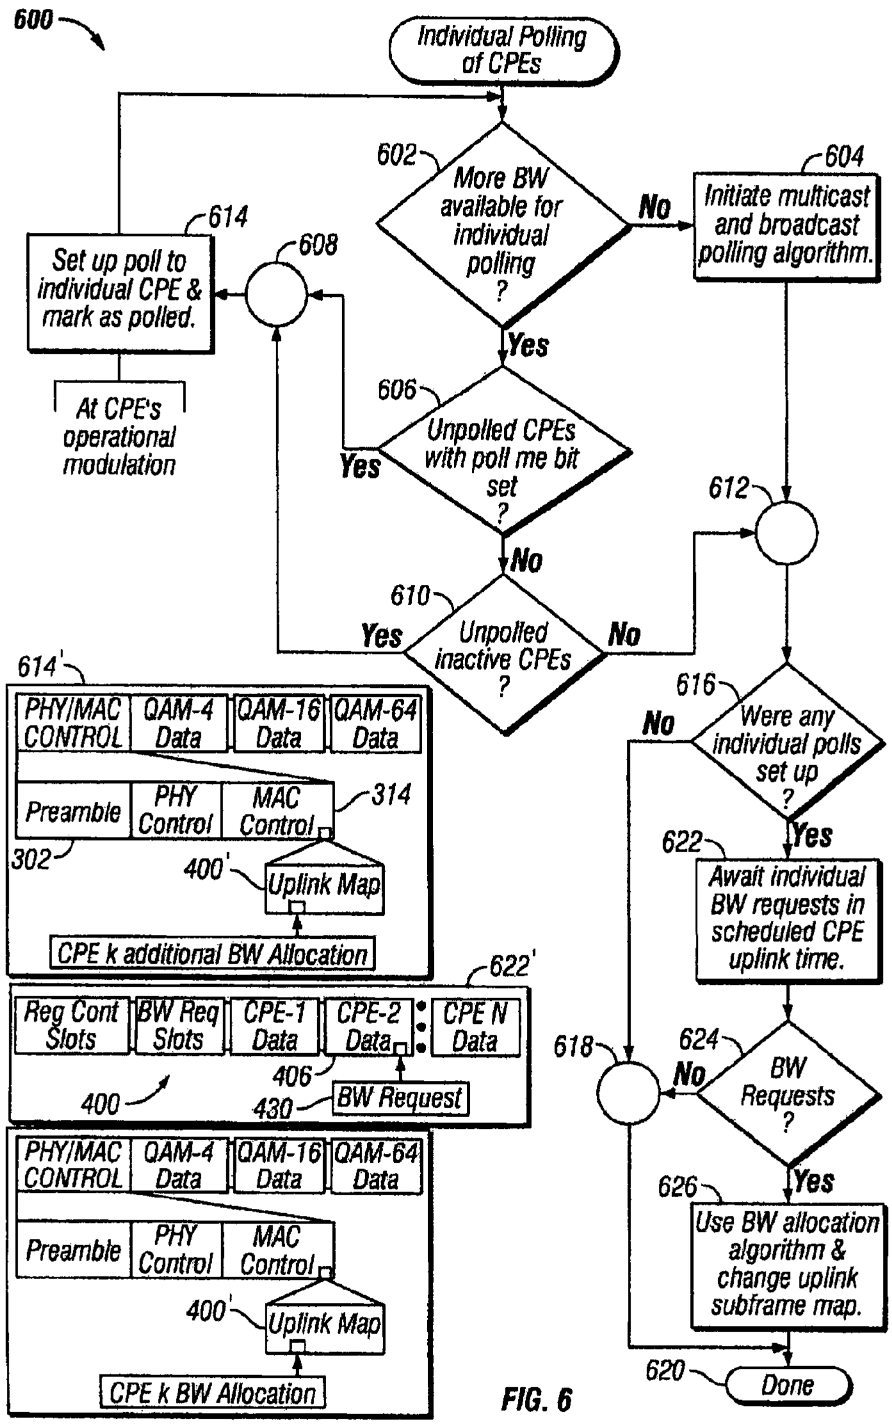
600
Individual Polling of CPEs
602
More BW available for individual polling ?
No
604
Initiate multicast and broadcast polling algorithm.
Yes
614
Set up poll to individual CPE & mark as polled.
608
At CPE's operational modulation
606
Unpolled CPEs with poll me bit set ?
Yes
No
612
610
Unpolled inactive CPEs ?
Yes
No
614'
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
314
302
400'
Uplink Map
CPE k additional BW Allocation
616
Were any individual polls set up ?
No
Yes
622
Await individual BW requests in scheduled CPE uplink time.
622'
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
406
400
430
BW Request
618
624
No
BW Requests ?
Yes
626
Use BW allocation algorithm & change uplink subframe map.
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
400'
Uplink Map
CPE k BW Allocation
620
Done

FIG. 6

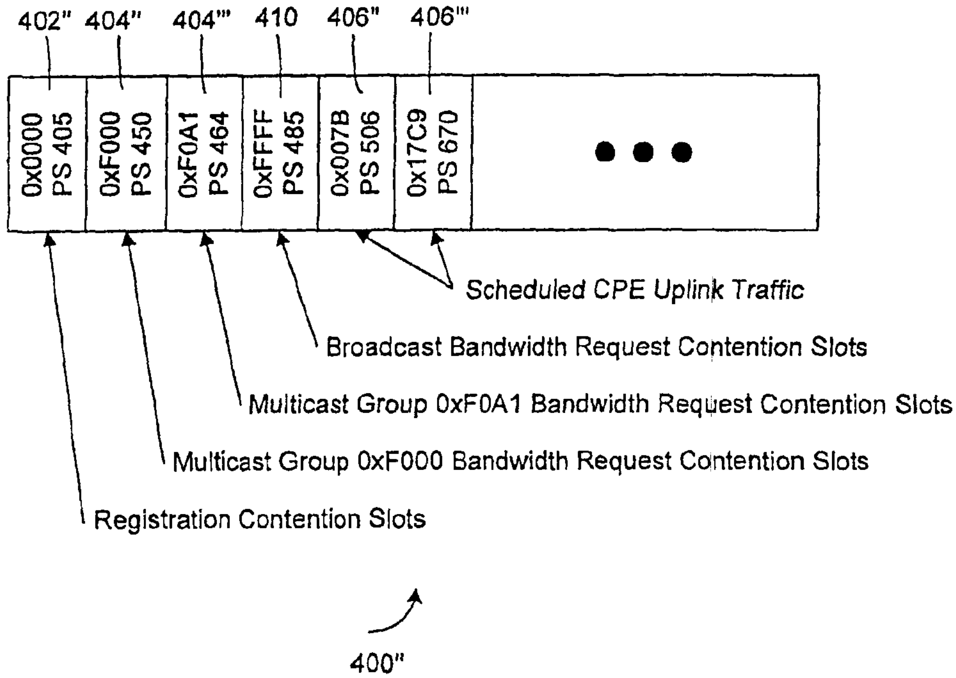
402"
404"
404'"
410
406"
406'"
0x0000
PS 405
0xF000
PS 450
0xF0A1
PS 464
0xFFFF
PS 485
0x007B
PS 506
0x17C9
PS 670
Scheduled CPE Uplink Traffic
Broadcast Bandwidth Request Contention Slots
Multicast Group 0xF0A1 Bandwidth Request Contention Slots
Multicast Group 0xF000 Bandwidth Request Contention Slots
Registration Contention Slots
400"

FIGURE 7

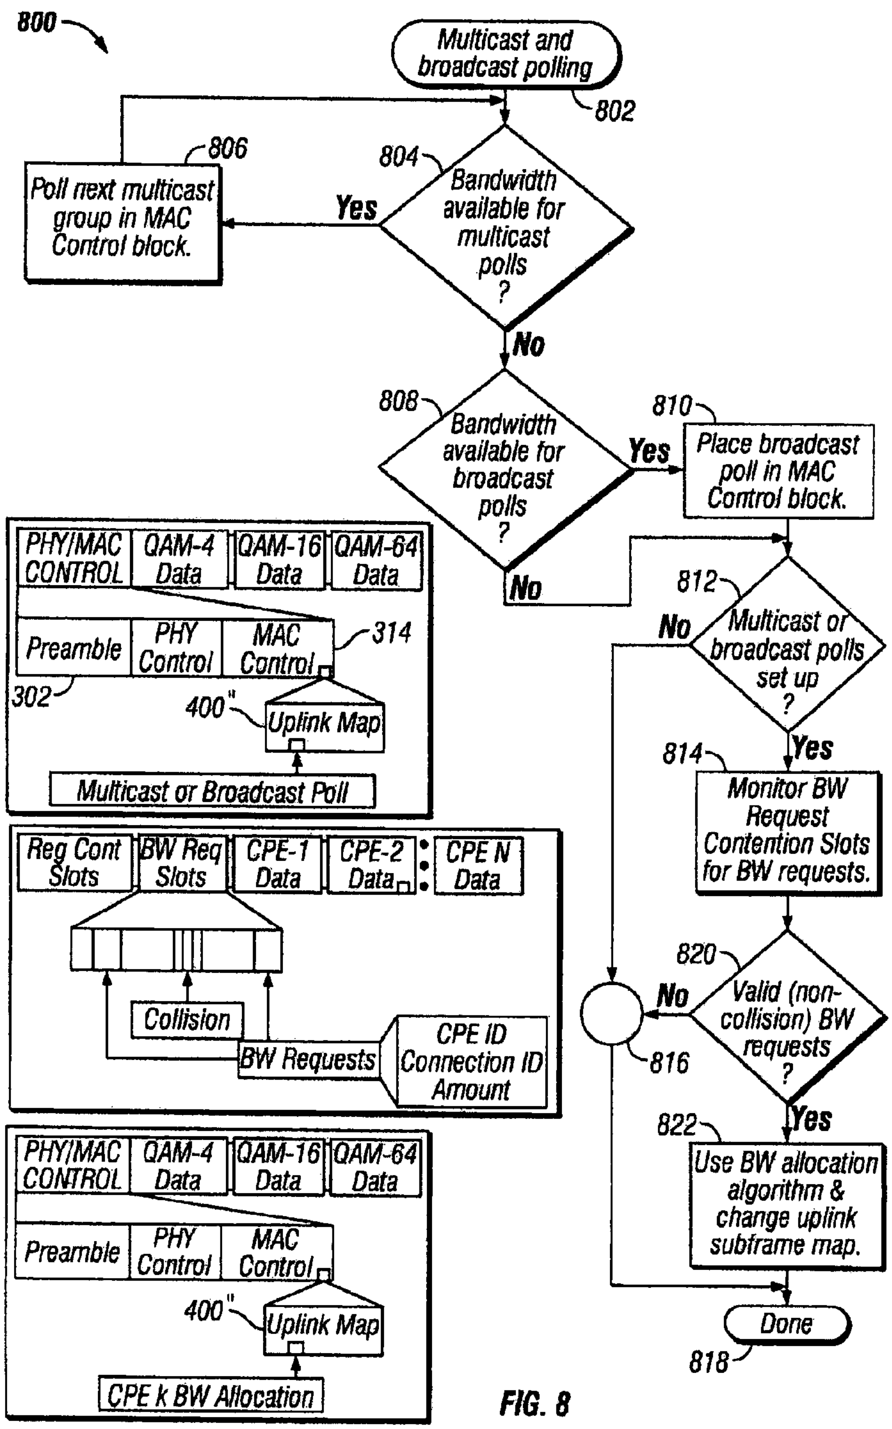
800
Multicast and broadcast polling
802
804
Bandwidth available for multicast polls ?
Yes
806
Poll next multicast group in MAC Control block.
No
808
Bandwidth available for broadcast polls ?
Yes
810
Place broadcast poll in MAC Control block.
No
812
Multicast or broadcast polls set up ?
No
Yes
814
Monitor BW Request Contention Slots for BW requests.
820
Valid (non-collision) BW requests ?
No
816
Yes
822
Use BW allocation algorithm & change uplink subframe map.
Done
818
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
314
302
400"
Uplink Map
Multicast or Broadcast Poll
Reg Cont Slots
BW Req Slots
CPE-1 Data
CPE-2 Data
CPE N Data
Collision
BW Requests
CPE ID
Connection ID
Amount
PHY/MAC CONTROL
QAM-4 Data
QAM-16 Data
QAM-64 Data
Preamble
PHY Control
MAC Control
400"
Uplink Map
CPE k BW Allocation

FIG. 8

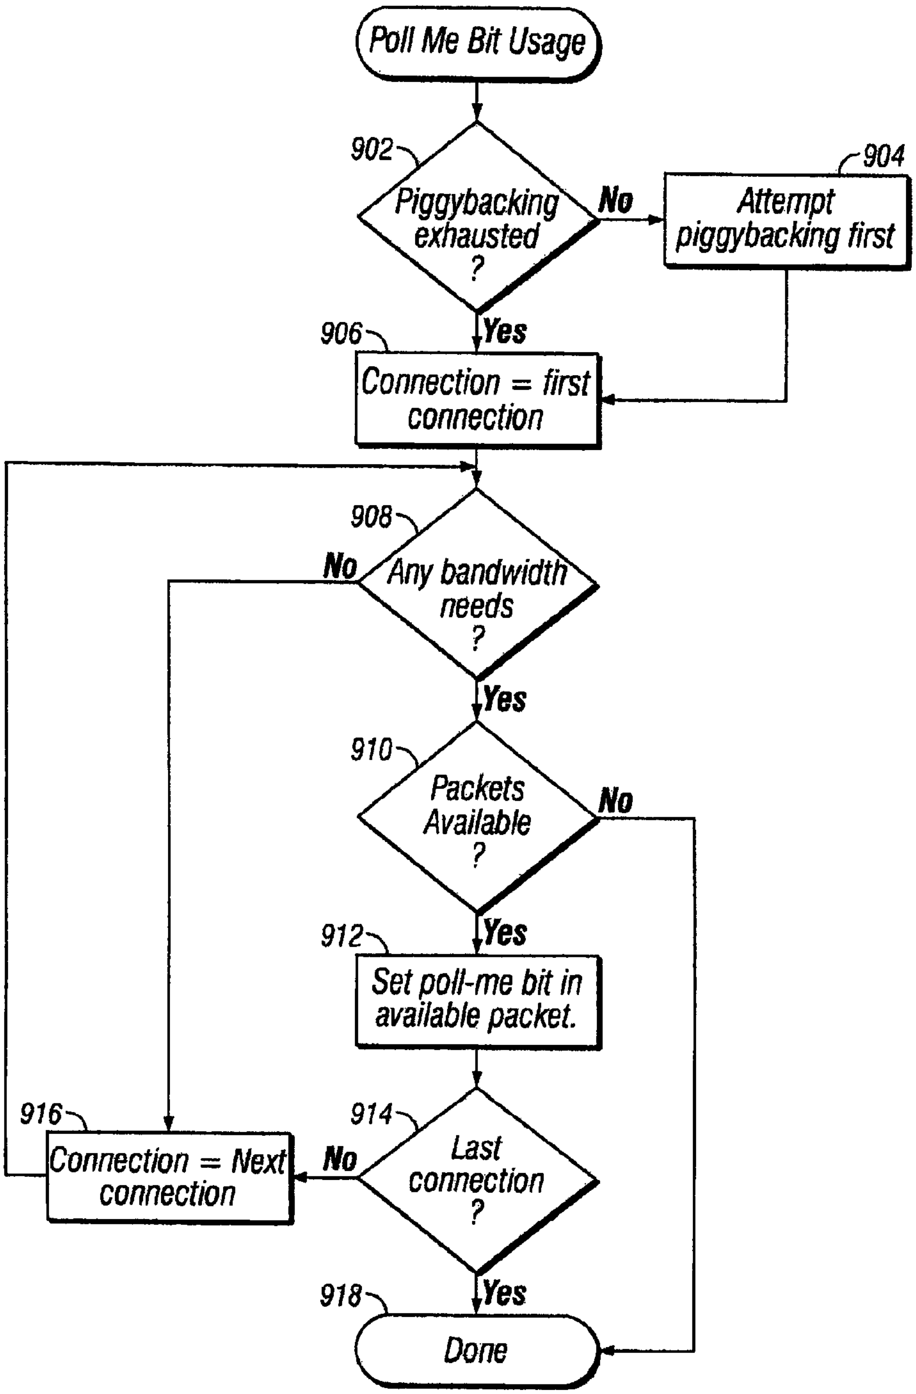
Poll Me Bit Usage
902
Piggybacking exhausted ?
No
904
Attempt piggybacking first
Yes
906
Connection = first connection
908
Any bandwidth needs ?
No
Yes
910
Packets Available ?
No
Yes
912
Set poll-me bit in available packet.
914
Last connection ?
No
916
Connection = Next connection
Yes
918
Done

FIG. 9

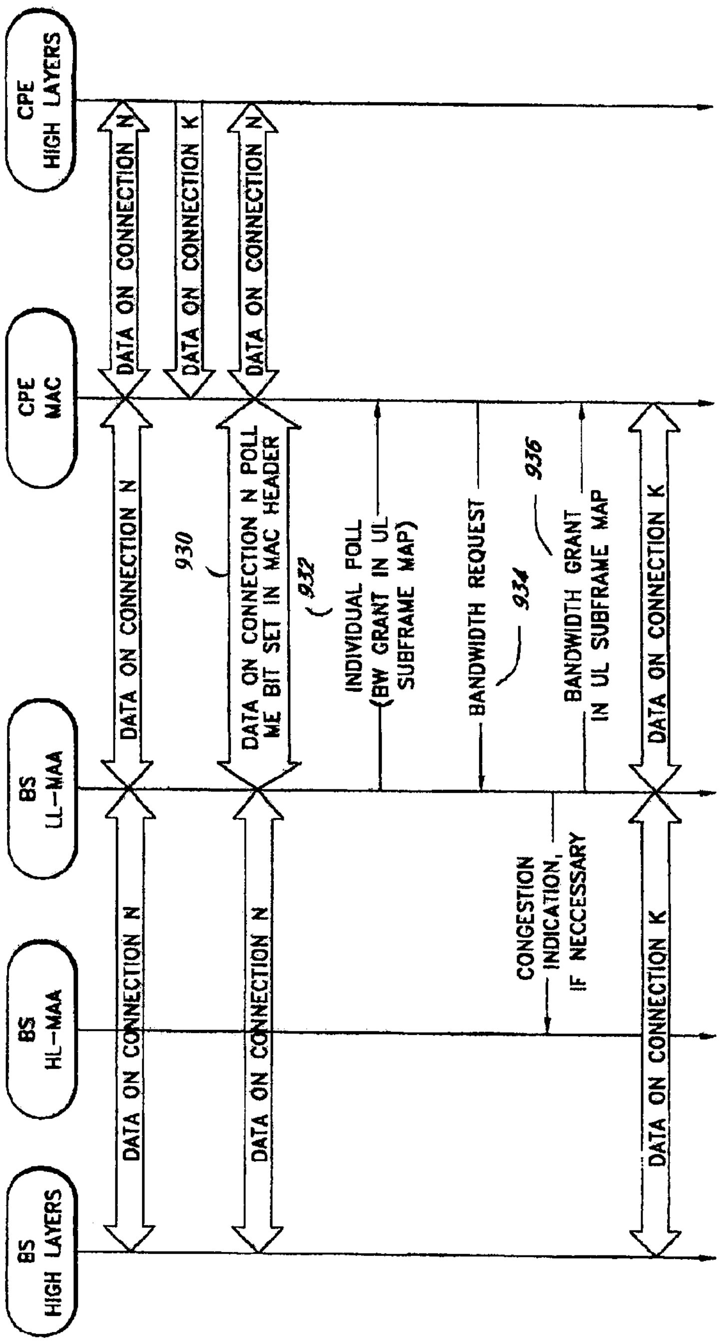
CPE HIGH LAYERS
CPE MAC
BS LL-MAA
BS HL-MAA
BS HIGH LAYERS
DATA ON CONNECTION N
DATA ON CONNECTION K
DATA ON CONNECTION N
DATA ON CONNECTION N
930
DATA ON CONNECTION N POLL ME BIT SET IN MAC HEADER
932
INDIVIDUAL POLL (BW GRANT IN UL SUBFRAME MAP)
BANDWIDTH REQUEST
934
936
BANDWIDTH GRANT IN UL SUBFRAME MAP
DATA ON CONNECTION K
DATA ON CONNECTION N
DATA ON CONNECTION N
CONGESTION INDICATION, IF NECCESSARY
DATA ON CONNECTION K

FIG. 10

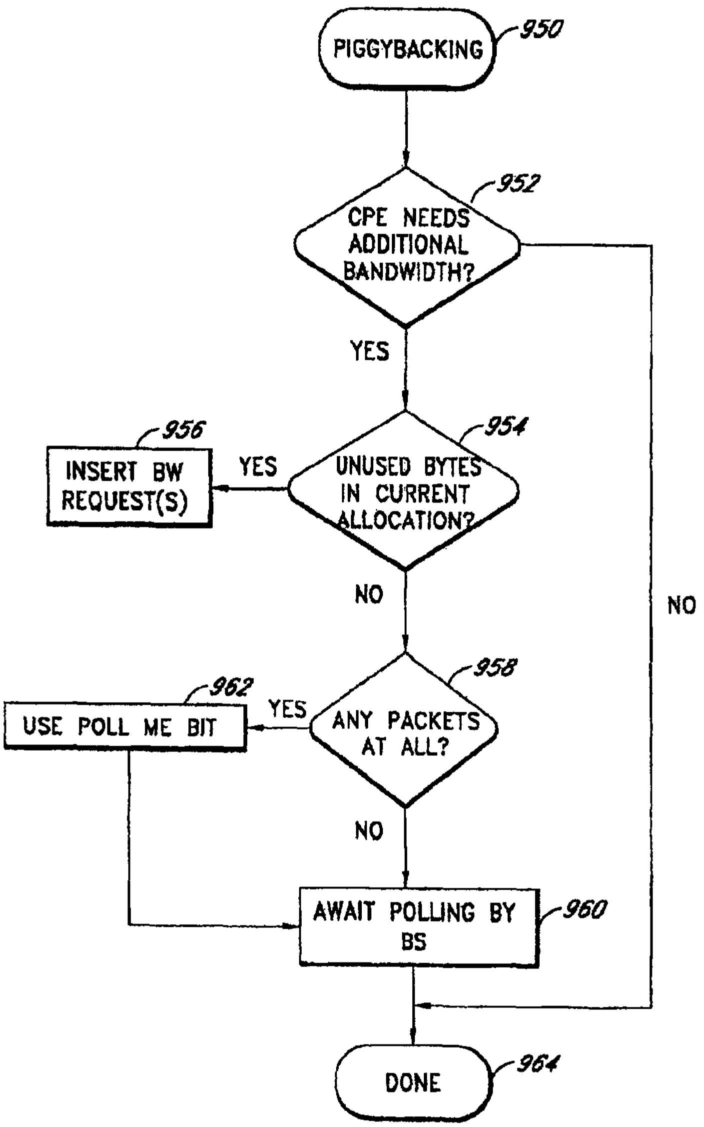
PIGGYBACKING
950
CPE NEEDS ADDITIONAL BANDWIDTH?
952
YES
NO
UNUSED BYTES IN CURRENT ALLOCATION?
954
YES
INSERT BW REQUEST(S)
956
NO
ANY PACKETS AT ALL?
958
YES
USE POLL ME BIT
962
NO
AWAIT POLLING BY BS
960
DONE
964

FIG. 11

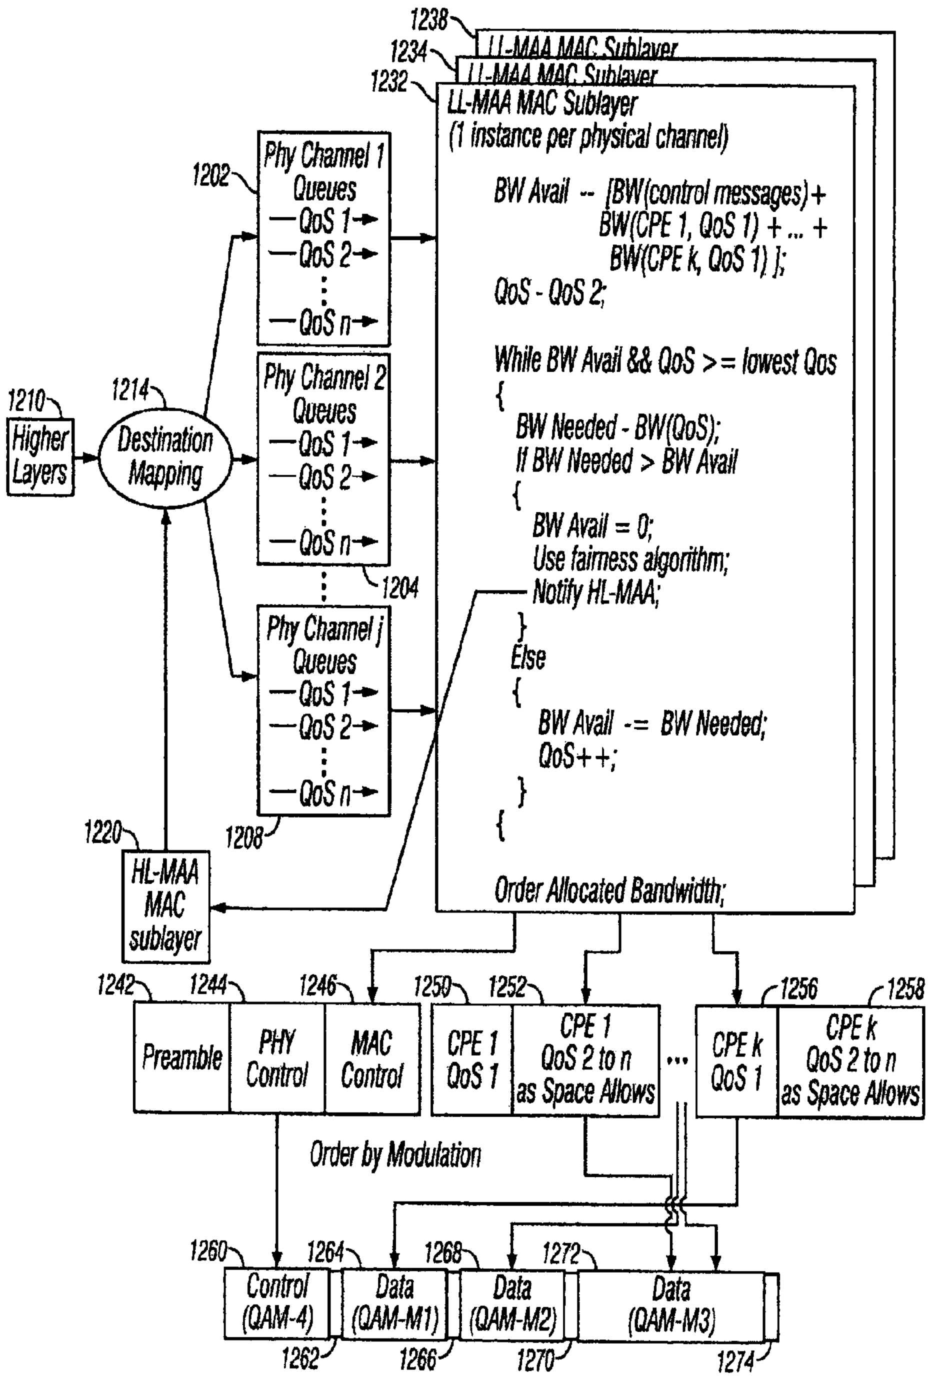
1238
LL-MAA MAC Sublayer
1234
LL-MAA MAC Sublayer
1232
LL-MAA MAC Sublayer
(1 instance per physical channel)
BW Avail -- [BW(control messages)+
BW(CPE 1, QoS 1) + ... +
BW(CPE k, QoS 1)];
QoS - QoS 2;
While BW Avail && QoS >= lowest Qos
{
BW Needed - BW(QoS);
If BW Needed > BW Avail
{
BW Avail = 0;
Use fairness algorithm;
Notify HL-MAA;
}
Else
{
BW Avail -= BW Needed;
QoS++;
}
{
Order Allocated Bandwidth;
1202
Phy Channel 1
Queues
QoS 1
QoS 2
QoS n
Phy Channel 2
Queues
QoS 1
QoS 2
QoS n
1204
Phy Channel j
Queues
QoS 1
QoS 2
QoS n
1208
1210
Higher Layers
1214
Destination Mapping
1220
HL-MAA MAC sublayer
1242
1244
1246
1250
1252
1256
1258
Preamble
PHY Control
MAC Control
CPE 1 QoS 1
CPE 1 QoS 2 to n as Space Allows
...
CPE k QoS 1
CPE k QoS 2 to n as Space Allows
Order by Modulation
1260
1264
1268
1272
Control (QAM-4)
Data (QAM-M1)
Data (QAM-M2)
Data (QAM-M3)
1262
1266
1270
1274

FIG. 12

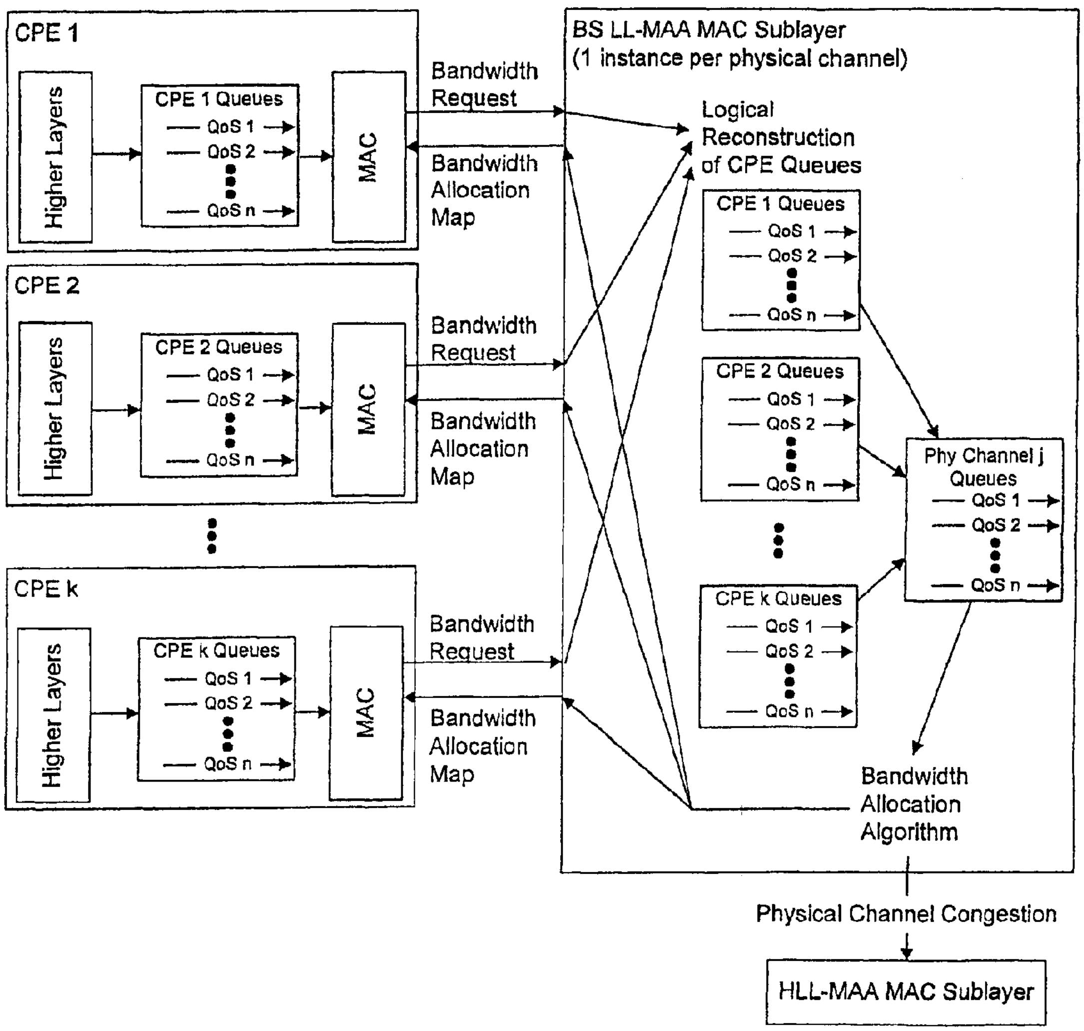
CPE 1
Higher Layers
CPE 1 Queues
QoS 1
QoS 2
QoS n
MAC
CPE 2
Higher Layers
CPE 2 Queues
QoS 1
QoS 2
QoS n
MAC
CPE k
Higher Layers
CPE k Queues
QoS 1
QoS 2
QoS n
MAC
Bandwidth Request
Bandwidth Allocation Map
Bandwidth Request
Bandwidth Allocation Map
Bandwidth Request
Bandwidth Allocation Map
BS LL-MAA MAC Sublayer
(1 instance per physical channel)
Logical Reconstruction of CPE Queues
CPE 1 Queues
QoS 1
QoS 2
QoS n
CPE 2 Queues
QoS 1
QoS 2
QoS n
CPE k Queues
QoS 1
QoS 2
QoS n
Phy Channel j Queues
QoS 1
QoS 2
QoS n
Bandwidth Allocation Algorithm
Physical Channel Congestion
HLL-MAA MAC Sublayer

FIGURE 13

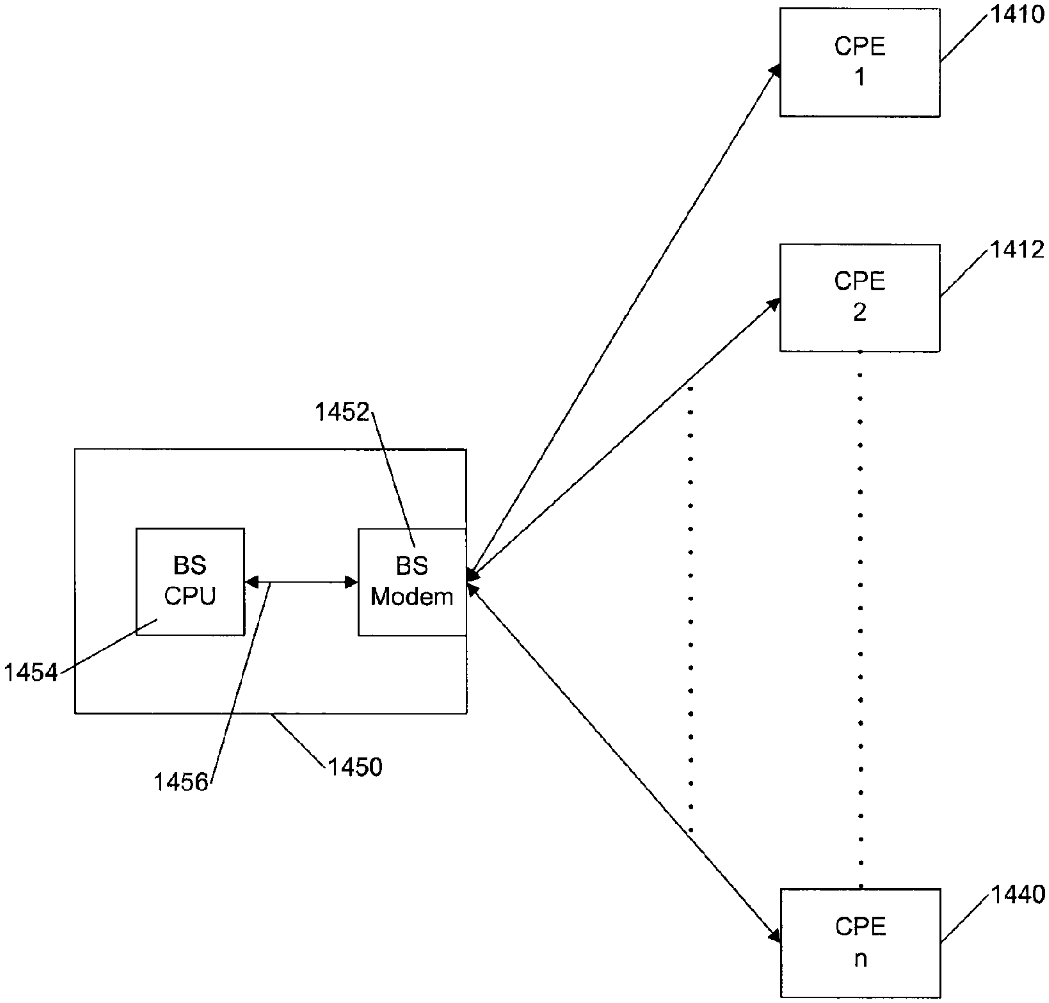
CPE
1
1410
CPE
2
1412
1452
BS
CPU
BS
Modem
1454
1456
1450
CPE
n
1440

FIGURE 14

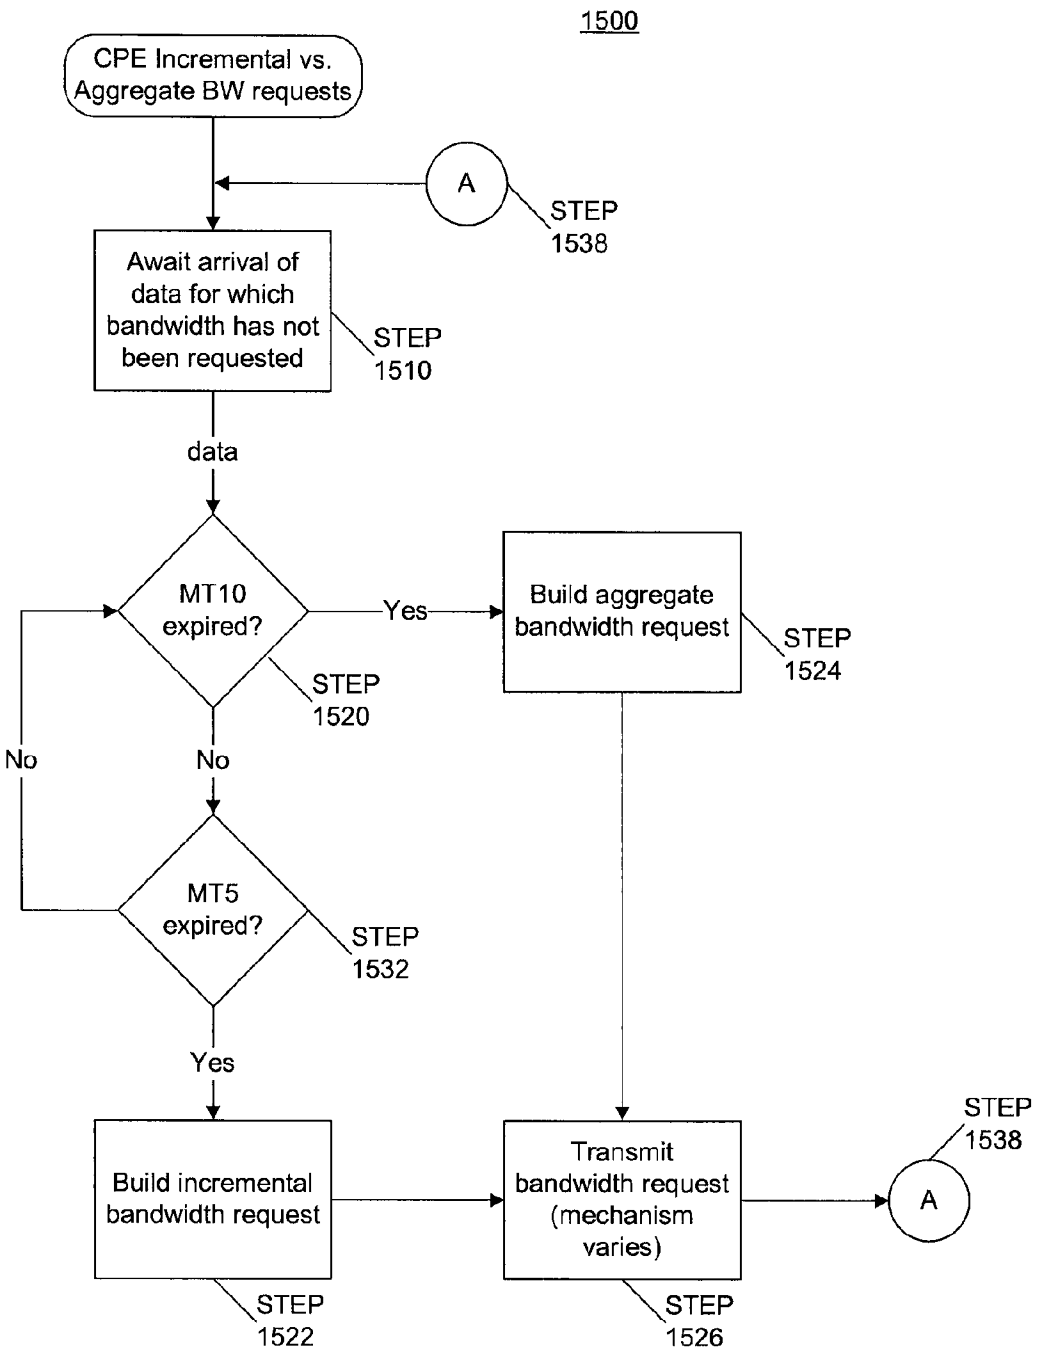
1500
CPE Incremental vs. Aggregate BW requests
A
STEP 1538
Await arrival of data for which bandwidth has not been requested
STEP 1510
data
MT10 expired?
Yes
Build aggregate bandwidth request
STEP 1524
STEP 1520
No
No
MT5 expired?
STEP 1532
Yes
Build incremental bandwidth request
Transmit bandwidth request (mechanism varies)
STEP 1538
A
STEP 1522
STEP 1526

FIGURE 15a

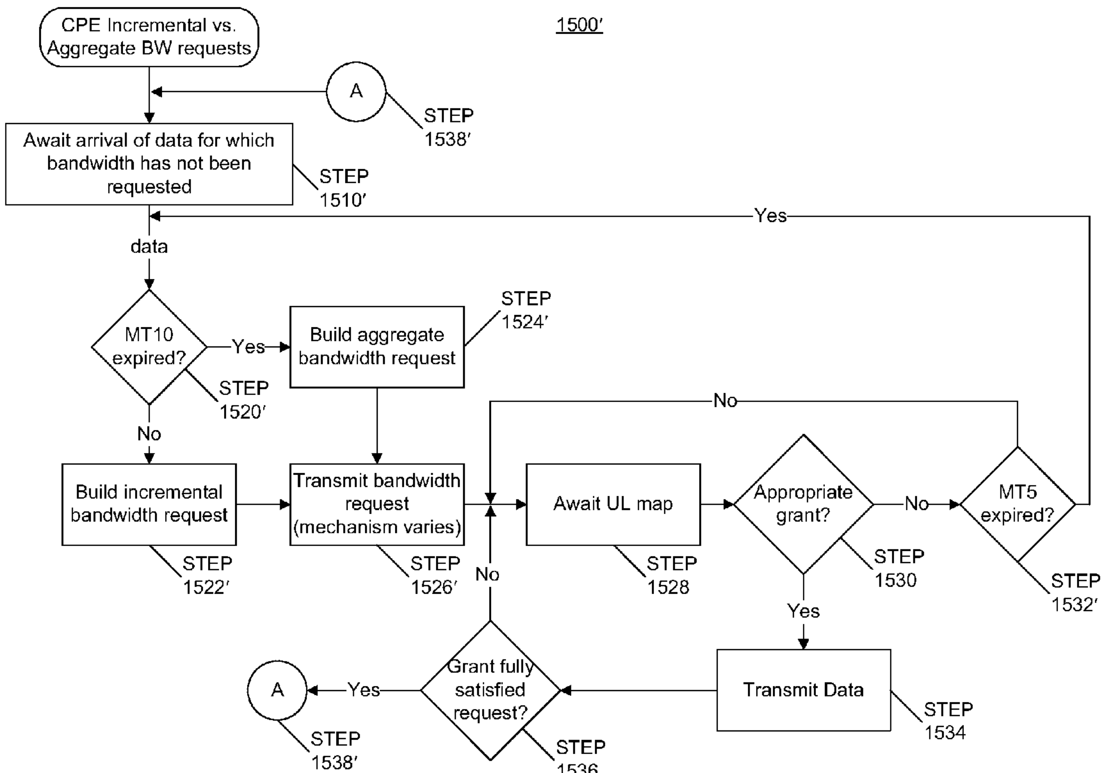

1500′
CPE Incremental vs. Aggregate BW requests
A
STEP 1538′
Await arrival of data for which bandwidth has not been requested
STEP 1510′
Yes
data
MT10 expired?
Yes
Build aggregate bandwidth request
STEP 1524′
STEP 1520′
No
No
Build incremental bandwidth request
Transmit bandwidth request (mechanism varies)
Await UL map
Appropriate grant?
No
MT5 expired?
STEP 1522′
STEP 1526′
No
STEP 1528
STEP 1530
STEP 1532′
Yes
A
Yes
Grant fully satisfied request?
Transmit Data
STEP 1534
STEP 1538′
STEP 1536

FIGURE 15b

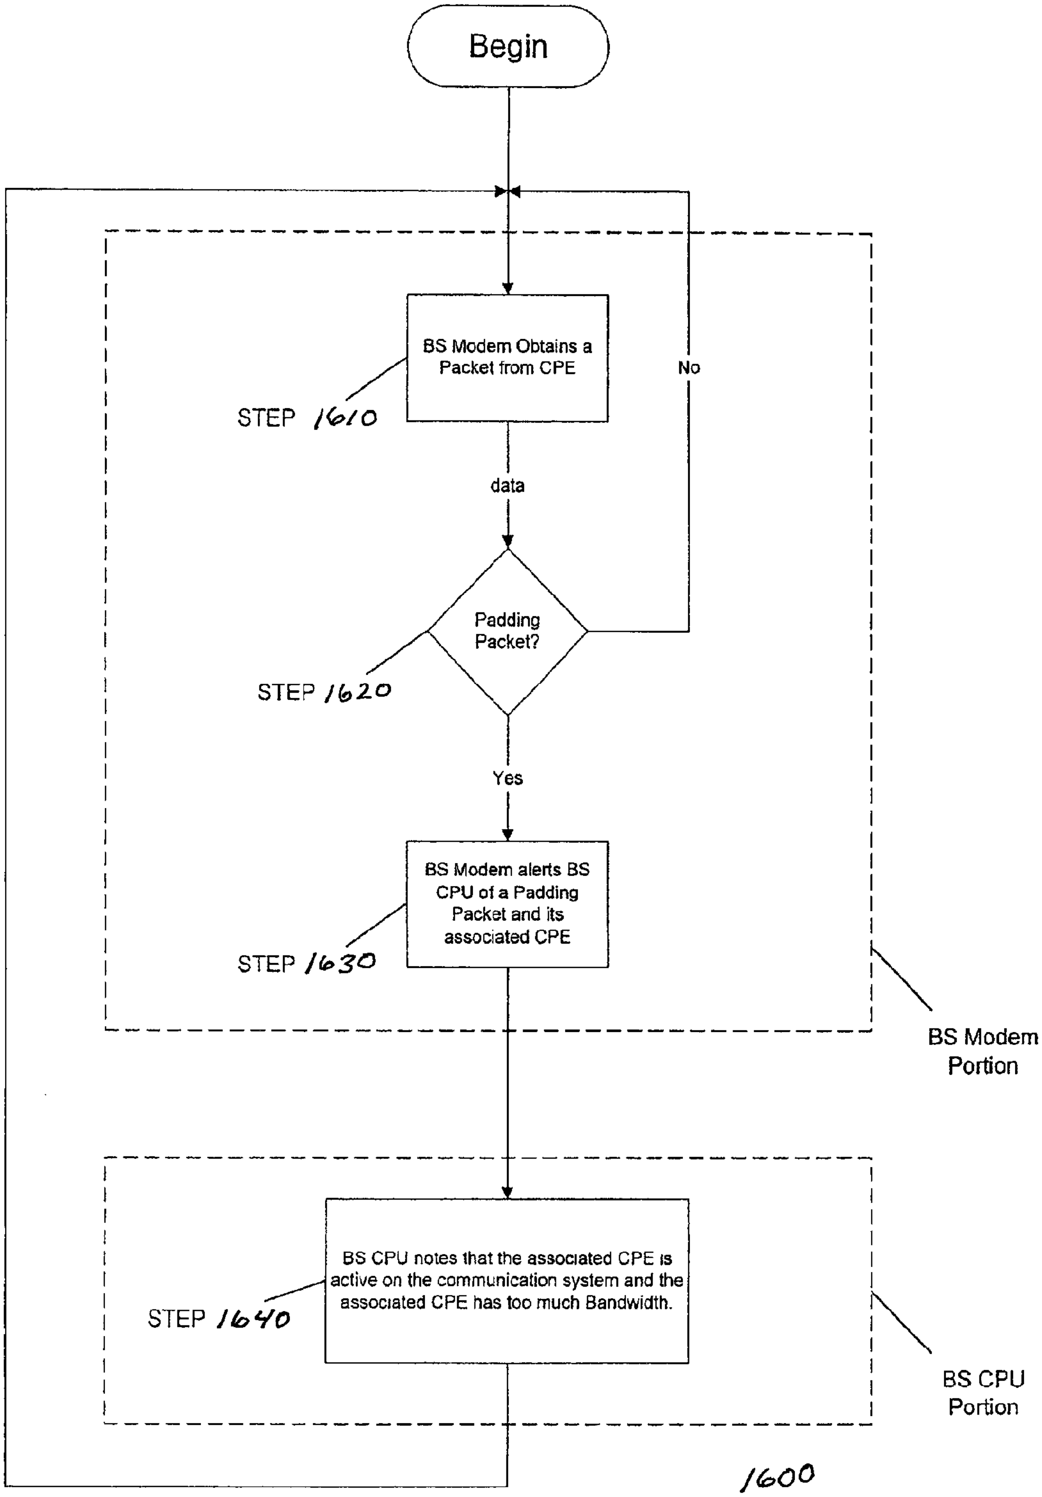
Begin
BS Modem Obtains a Packet from CPE
STEP 1610
data
Padding Packet?
STEP 1620
No
Yes
BS Modem alerts BS CPU of a Padding Packet and its associated CPE
STEP 1630
BS Modem Portion
BS CPU notes that the associated CPE is active on the communication system and the associated CPE has too much Bandwidth.
STEP 1640
BS CPU Portion
1600

FIGURE 16

METHOD AND APPARATUS FOR BANDWIDTH REQUEST/GRANT PROTOCOLS IN A WIRELESS COMMUNICATION SYSTEM

RELATED APPLICATIONS

This application is a continuation of application Ser. No. 12/415,698, filed Mar. 31, 2009, now abandoned, which is a continuation of application Ser. No. 10/978,903, filed Nov. 1, 2004, now U.S. Pat. No. 7,529,193, which is a continuation of application Ser. No. 09/783,671, filed Feb. 14, 2001, now abandoned, which is a continuation-in-part of application Ser. No. 09/316,518, filed May 21, 1999, now U.S. Pat. No. 6,925,068, all of which are hereby incorporated by reference. Application Ser. No. 09/783,671 is also a continuation-in-part of application Ser. No. 09/613,434, filed Jul. 11, 2000, now U.S. Pat. No. 6,785,252, which is a continuation-in-part of application Ser. No. 09/316,518, filed May 21, 1999, now U.S. Pat. No. 6,925,068, all of which are hereby incorporated by reference. This application is also related to and hereby incorporates by reference U. S. Pat. No. 6,016,311, issued Jan. 18, 2000.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to wireless communication systems, and to a method and apparatus of bandwidth request/grant protocols in a broadband wireless communication system.

2. Description of Related Art

As described in the commonly assigned U.S. Pat. No. 6,016,311, issued Jan. 18, 2000, titled “An Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System,” which is hereby incorporated by reference, a wireless communication system facilitates two-way communication between a plurality of subscriber radio stations or subscriber units (fixed and portable) and a fixed network infrastructure. Exemplary communication systems include mobile cellular telephone systems, personal communication systems (PCS), and cordless telephones. The key objective of these wireless communication systems is to provide communication channels on demand between the plurality of subscriber units and their respective base stations in order to connect a subscriber unit user with the fixed network infrastructure (usually a wire-line system). In the wireless systems having multiple access schemes a time “frame” is used as the basic information transmission unit. Each frame is sub-divided into a plurality of time slots. Some time slots are used for control purposes and some for information transfer. Subscriber units typically communicate with the base station using a “duplexing” scheme thus allowing the exchange of information in both directions of connection.

Transmissions from the base station to the subscriber unit are commonly referred to as “downlink” transmissions. Transmissions from the subscriber unit to the base station are commonly referred to as “uplink” transmissions. Depending upon the design criteria of a given system, the prior art wireless communication systems have typically used either time division duplexing (TDD) or frequency division duplexing (FDD) methods to facilitate the exchange of information between the base station and the subscriber units. Both the TDD and FDD duplexing schemes are well known in the art.

Recently, wideband or “broadband” wireless communications networks have been proposed for providing delivery of enhanced broadband services such as voice, data and video services. The broadband wireless communication system facilitates two-way communication between a plurality of base stations and a plurality of fixed subscriber stations or Customer Premises Equipment (CPE). One exemplary broadband wireless communication system is described in the incorporated patent, U.S. Pat. No. 6,016,311, and is shown in the block diagram of FIG. **1**. As shown in FIG. **1**, the exemplary broadband wireless communication system **100** includes a plurality of cells **102**. Each cell **102** contains an associated cell site **104** that primarily includes a base station **106** and an active antenna array **108**. Each cell **102** provides wireless connectivity between the cell's base station **106** and a plurality of customer premises equipment (CPE) **110** positioned at fixed customer sites **112** throughout the coverage area of the cell **102**. The users of the system **100** may include both residential and business customers. Consequently, the users of the system have different and varying usage and bandwidth requirement needs. Each cell may service several hundred or more residential and business CPEs.

The broadband wireless communication system **100** of FIG. **1** provides true “bandwidth-on-demand” to the plurality of CPEs **110**. CPEs **110** request bandwidth allocations from their respective base stations **106** based upon the type and quality of services requested by the customers served by the CPEs. Different broadband services have different bandwidth and latency requirements. The type and quality of services available to the customers are variable and selectable. The amount of bandwidth dedicated to a given service is determined by the information rate and the quality of service required by that service (and also taking into account bandwidth availability and other system parameters). For example, T1-type continuous data services typically require a great deal of bandwidth having well-controlled delivery latency. Until terminated, these services require constant bandwidth allocation on each frame. In contrast, certain types of data services such as Internet protocol data services (TCP/IP) are bursty, often idle (which at any one instant requires zero bandwidth), and are relatively insensitive to delay variations when active.

Due to the wide variety of CPE service requirements, and due to the large number of CPEs serviced by any one base station, the bandwidth allocation process in a broadband wireless communication system such as that shown in FIG. **1** can become burdensome and complex. This is especially true with regard to the allocation of uplink bandwidth. Base stations do not have a priori information regarding the bandwidth or quality of services that a selected CPE will require at any given time. Consequently, requests for changes to the uplink bandwidth allocation are necessarily frequent and varying. Due to this volatility in the uplink bandwidth requirements, the many CPEs serviced by a selected base station will need to frequently initiate bandwidth allocation requests. If uncontrolled, the bandwidth allocation requests will detrimentally affect system performance. If left unchecked, the bandwidth required to accommodate CPE bandwidth allocation requests will become disproportionately high in comparison with the bandwidth allocated for the transmission of substantive traffic data. Thus, the communication system bandwidth available to provide broadband services will be disadvantageously reduced.

Some prior art systems have attempted to solve bandwidth allocation requirements in a system having a shared system resource by maintaining logical queues associated with the various data sources requiring access to the shared system resource. Such a prior art system is taught by Karol et al., in U.S. Pat. No. 5,675,573, that issued on Oct. 7, 1997. More specifically, Karol et al. teach a bandwidth allocation system that allows packets or cells within traffic flows from different sources that are contending for access to a shared processing fabric to get access to that fabric in an order that is determined primarily on individual guaranteed bandwidth requirements associated with each traffic flow. In addition, the system taught by Karol et al. allow the different sources to gain access to the shared processing fabric in an order determined secondarily on overall system criteria, such as a time of arrival, or due date of packets or cells within the traffic flows. Packets or cells of data from each data source (such as a bandwidth requesting device) are queued in separate logical buffers while they await access to the processing fabric.

The bandwidth allocation techniques described in the commonly assigned and incorporated U.S. patent application Ser. No. 09/316,518, filed May 21, 2000, utilizes mechanisms referred to as "bandwidth request/grant protocols" to provide on-demand bandwidth needs of individual CPE connections. Typically, bandwidth request/grant protocols operate in accordance with the following description. A CPE typically transmits a bandwidth request to an associated base station. The request identifies the aggregate (i.e., the total) bandwidth needs of the connection. The base station receives the bandwidth request and determines whether sufficient bandwidth is available to grant the bandwidth request. If sufficient bandwidth is available, the requested bandwidth is granted to the connection, else the base station waits for sufficient bandwidth to become available before granting the requested bandwidth. As described in the parent patent application, bandwidth request/grant protocols improve bandwidth allocation efficiencies in wireless communication systems under ideal conditions.

However, as is well known, bandwidth requests (and associated grants) can be lost (i.e., never received by the associated base station) or delayed due to noise and interference effects inherent to all wireless communication systems. When bandwidth requests are lost or delayed during transmission between a CPE and a base station, bandwidth allocation efficiencies can be adversely affected. Lost or delayed bandwidth requests contribute to the reduction of bandwidth allocation efficiency in wireless communication systems by causing the base stations to inaccurately allocate bandwidth to their associated and respective CPEs.

For example, consider the situation where a selected CPE transmits a bandwidth request to its associated base station wherein the request identifies the aggregate bandwidth requirements of the selected CPE. Assume that the bandwidth request is lost in transmission due to interference on the air link between the base station and the selected CPE. In this example, the associated base station never receives the aggregate bandwidth requirements of the selected CPE, and the base station therefore never grants the CPE's bandwidth request. After waiting a suitable period of time, the CPE will determine that is has not received a bandwidth grant from the base station. Disadvantageously, the CPE will be unable to determine if the bandwidth request was lost during transmission or if the base station simply did not have sufficient bandwidth to grant the request (given the quality of service ("QoS") of the associated connection).

The CPE may then transmit a second bandwidth request for the same connection. Under certain conditions, a "race condition" may occur that could cause the bandwidth allocation technique to waste the allocation of bandwidth. If the timing of the bandwidth requests (and subsequent grants) is such that the selected CPE issues the second bandwidth request for the same connection concurrently with the base station's grant of the first request, the second request and the grant to the first request may be concurrently transmitted over the link. That is,

if the base station transmits a grant to the first request before receiving the second request from the CPE, the base station may respond to the second request and consequently grant a duplicate bandwidth request for the same connection. This disadvantageously results in an efficient allocation of bandwidth.

One alternative bandwidth request/grant protocol that prevents the occurrence of the above-described "race condition" is the so-called "guaranteed delivery protocol." As is well known, guaranteed delivery protocols make use of acknowledgment messages that are transmitted in response to bandwidth requests. In accordance with the guaranteed delivery protocol approach, a CPE transmits to its associated base station a bandwidth request that identifies the aggregate bandwidth needs of a selected connection. The base station receives the bandwidth request and transmits an acknowledgment to the CPE thereby communicating receipt of the bandwidth request. If an acknowledgment is not received by the CPE, the CPE retransmits the bandwidth request. Advantageously, guaranteed delivery protocols vastly reduce the possibility of the base station erroneously allocating duplicate bandwidth to the CPE (as described above), and thus, improves bandwidth allocation efficiencies. However, guaranteed delivery protocols disadvantageously require additional bandwidth necessary for transmitting acknowledgement messages between the base stations and the CPEs. Furthermore, response time associated with the allocation of bandwidth is reduced because the CPEs must wait to receive acknowledgements from their associated base stations.

Some bandwidth request/grant protocols known as "incremental bandwidth request/grant protocols" attempt to solve the above-described problems relating to data transmission efficiency by utilizing incremental bandwidth requests instead of aggregate bandwidth requests. Incremental bandwidth requests identify the additional bandwidth needs of a CPE connection. For example, in accordance with incremental bandwidth request methods, a base station may allocate 1000 units of bandwidth to an associated CPE connection. At a later time, the CPE connection may require 1,500 units of aggregate bandwidth (i.e., it may require an additional 500 units of bandwidth). In accordance with the incremental bandwidth request/grant protocol, the CPE will transmit an incremental bandwidth request to its associated base station indicating that it requires an additional 500 units of bandwidth. Upon receiving the incremental bandwidth request, the base station calculates the CPE connection's current aggregate bandwidth needs as 1500 units (1000 previously granted units+500 requested units).

Advantageously, systems using the incremental bandwidth request/grant protocols respond faster and require less bandwidth than do those using the guaranteed delivery protocols because acknowledgment messages are not required by the incremental bandwidth request/grant protocols. Disadvantageously, when an incremental bandwidth request is lost, the base station loses synchronization with the CPE connection, and thereby loses track of the aggregate bandwidth needs of the CPE. Synchronization is lost because the base stations typically calculate aggregate bandwidth needs by adding each incremental bandwidth request to the previous aggregate bandwidth needs estimate. Thus, the base station and the CPE connection will remain out of synchronization until the CPE connection is reset.

Some bandwidth request/grant protocol systems have attempted to solve bandwidth allocation requirements in a system having a shared system resource by utilizing "zero bandwidth request" (ZBR) messages. One such exemplary bandwidth allocation system is known as a zero bandwidth request message protocol system and is now described. ZBR message protocol systems utilize "padding packets" and the well-known TDMA multiplexing scheme. In the well-known TDMA multiplexing scheme, a BS designates a portion of its uplink sub-frame (i.e., bandwidth) to an associated CPE. The associated CPE transmits data to the BS on the uplink. When a CPE does not have enough uplink data to utilize its entire portion of bandwidth (i.e., it has too much bandwidth allocation), it transmits padding packets to "pad" or fill its unused portion of bandwidth. The CPE then transmits a ZBR message to its associated base station (BS) to request a reduction in the CPE's bandwidth allocation. The CPE's associated BS then reduces the CPE's bandwidth allocation accordingly.

Disadvantageously in ZBR message protocol systems, utilization of ZBR messages decreases a communication systems overall speed. Base stations and CPEs require increased processing time to process and transmit ZBR messages, respectively. A BS requires increased processing time to process ZBR messages. This disadvantage is magnified in typical communication systems because a BS typically receives ZBR messages from hundreds of associated CPEs. Thus, each BS in the communication system requires relatively large amounts of time to process these ZBR messages.

Another disadvantage of zero bandwidth request message protocol systems is that CPEs can become "confused" when deciding whether to transmit zero bandwidth requests to their associated base stations. For example, a CPE has the following status: a CG connection rate of one cell per second and a DAMA connection with no data available. When the CPE's associated BS allocates one cell within a one-second time interval the cell may not yet be available within the CPE's ATM controller queues. In accordance with the ZBR protocol system, the CPE should transmit a ZBR message because of the "no data available" status of the DAMA connection. However, the CPE does not know whether or not a CG cell is going to be sent and thus it does not know whether or not to transmit a ZBR message to its associated BS. Thus, the CPE becomes confused and can erroneously transmit or refrain from transmitting a ZBR message.

A need exists for a bandwidth request/grant protocol method and apparatus that efficiently processes and responds to bandwidth allocation requests. The bandwidth allocation method and apparatus should accommodate an arbitrarily large number of CPEs generating frequent and varying bandwidth allocation requests on the uplink of a wireless communication system. For example, in the system shown in FIG. **1**, as many as one hundred CPEs may be allowed to be simultaneously active, coordinating their transmissions on the uplink. Furthermore, the system can accommodate approximately one thousand CPEs on the physical channel. Such a bandwidth allocation method and apparatus should be efficient in terms of the amount of bandwidth consumed by the bandwidth request and grant messages that are exchanged between the plurality of base stations and the plurality of CPEs. That is, the plurality of bandwidth requests generated by the CPE should consume a minimum percentage of available uplink bandwidth. In addition, the bandwidth allocation method and apparatus should respond to the bandwidth allocation requests in a timely and accurate manner. The method and apparatus should be responsive to the needs of a particular communication link. The bandwidth needs may vary due to several factors, including the type of service provided over the link and the user type. Bandwidth should be allocated to high priority services in a sufficiently short time frame to maintain the quality of service specified by the CPE. The bandwidth request/grant protocol method and apparatus should correct itself when a bandwidth request is lost due to the noise or interference effects present on an air link.

SUMMARY OF THE INVENTION

The present invention is a method and apparatus for requesting and allocating bandwidth in a broadband wireless communication system. In one embodiment, the method and apparatus is a self-correcting bandwidth request/grant protocol that efficiently allocates bandwidth in the wireless communication system. The self-correcting bandwidth request/grant protocol utilizes a combination of incremental and aggregate bandwidth requests. CPEs primarily transmit incremental bandwidth requests to associated and respective base stations. The CPEs also periodically transmit aggregate bandwidth requests to their associated base stations. By primarily requiring incremental bandwidth requests, the possibility that a base station will erroneously issue duplicate bandwidth allocations to the same CPE for the same connection is vastly reduced. Race conditions that may have occurred when only aggregate bandwidth requests are transmitted are eliminated by requiring the CPEs to request bandwidth in an incremental manner.

In addition, by requiring the CPEs to periodically transmit aggregate bandwidth requests (that express the current state of their respective connection queues), the bandwidth allocation method and apparatus is "self-correcting". The bandwidth allocation method and apparatus is self-correcting because any lost incremental bandwidth requests are detected by the base stations when the aggregate bandwidth requests are received from their respective CPEs. Upon receipt and processing of the aggregate bandwidth requests, the base stations reset their records to correctly reflect the current bandwidth requirements of their associated CPEs. Periodic use of aggregate bandwidth requests provides a self-correcting bandwidth allocation protocol yet without the bandwidth overhead (e.g., bandwidth required to transmit acknowledgements) typically associated with the prior art self-correcting protocols (such as the guaranteed delivery protocols).

In wireless communication systems, DAMA services are allocated bandwidth on a demand-assignment basis. When a CPE requires additional bandwidth on a DAMA connection, it transmits a bandwidth request message to an associated base station. The CPE transmits an incremental bandwidth request to its associated base station. Periodically (e.g., every fifth bandwidth request), the CPE transmits an aggregate bandwidth request to its associated base station. The aggregate bandwidth request is used by the base station to "reset" (or update) its records to reflect the CPE's current bandwidth requirements. In this manner, the present bandwidth allocation protocol is said to be "self-correcting." In additional embodiments, four consecutive incremental bandwidth requests are transmitted, followed by one aggregate bandwidth request transmission. In alternative embodiments, alternative patterns of consecutive incremental and aggregate bandwidth requests can be used.

In one embodiment, the method and apparatus utilizes an abridged bandwidth request/grant protocol to allocate bandwidth. This embodiment utilizes padding packets to request a reduction in bandwidth allocation to a CPE. In one embodiment, a base station modem alerts a base station CPU when the BS modem receives padding packets from a CPE. After alerting the BS CPU the present inventive method can reduce the associated CPE's bandwidth allocation.

The embodiments herein disclosed reduce the amount of bandwidth that must be allocated for bandwidth request and bandwidth allocation purposes. The opportunities for allowing a CPE to request bandwidth are very tightly controlled and utilize a combination of a number of bandwidth request and allocation techniques to control the bandwidth request process. There are a number of means by which a CPE can transmit a bandwidth request message to an associated base station.

One such means uses a "polling" technique whereby a base station polls one or more CPEs and allocates bandwidth specifically for the purpose of allowing the CPEs to respond with a bandwidth request. The polling of the CPEs by the base station may be in response to a CPE setting a "poll-me bit" or, alternatively, it may be periodic. In accordance with the present invention, periodic polls may be made to individual CPEs, to groups of CPEs, or to every CPE on a physical channel. When individually polling a CPE, the base station polls an individual CPE by allocating uplink bandwidth in an uplink sub-frame map to allow the CPE to respond with a bandwidth request. Similarly, in group polling, the base station polls several CPEs by allocating uplink bandwidth in the uplink sub-frame map to allow the CPEs to respond with a bandwidth request. The CPEs must contend for the allocated bandwidth if collisions occur. Bandwidth allocations are not in the form of an explicit message that is communicated by the base station to the CPEs, but rather the bandwidth allocations are transmitted implicitly by allocating bandwidth in the uplink sub-frame map.

Another means used in reducing bandwidth consumed by the bandwidth request messages is the technique of "piggybacking" bandwidth requests on bandwidth already allocated to a CPE. In accordance with this technique, currently active CPEs request bandwidth using previously unused portions of uplink bandwidth that is already allocated to the CPE. Alternatively, the bandwidth requests can be piggybacked on uplink bandwidth already allocated and currently being used by a data service. In accordance with this alternative, the CPE "steals" bandwidth already allocated for a data connection by inserting bandwidth requests in time slots previously used for data.

The CPE is responsible for distributing the allocated uplink bandwidth in a manner that accommodates the services provided by the CPE. The CPE is free to use the uplink bandwidth that was allocated to it in a manner that is different than that originally requested or granted by the base station. The CPE advantageously determines which services to give bandwidth to and which services must wait for subsequent bandwidth requests. One advantage of having the CPE determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task. In addition, the communication overhead that is required by having the base station instruct the CPE how to distribute its allocated bandwidth is eliminated. By using a combination of bandwidth allocation techniques, the system advantageously makes use of the efficiency benefits associated with each technique.

The base station media access control ("MAC") allocates available bandwidth on a physical channel on the uplink and the downlink. Within the uplink and downlink sub-frames, the base station MAC allocates the available bandwidth between the various services depending upon the priorities and rules imposed by their quality of service ("QoS"). The base station MAC maintains a set of queues for each physical channel that it serves. Within each physical channel queue set, the base station maintains a queue for each QoS. The queues hold data that is ready to be transmitted to the CPEs present on the physical channel. The base station higher MAC control layers are free to implement any convenient fairness or traffic shaping algorithms regarding the sharing of access between connections at the same QoS, without impacting the base station lower MAC control layers. In determining the amount of bandwidth to allocate at a particular QoS for a particular CPE, the base station takes into account the QoS, modulation, and the fairness criteria used to keep an individual CPE from using up all available bandwidth. In one embodiment, the base station attempts to balance the uplink/downlink bandwidth allocations using an adaptive time-division duplexing technique (ATDD).

The uplink bandwidth allocation method is very similar to the downlink bandwidth allocation except that, rather than being maintained by the base station, the data queues are distributed across and maintained by each individual CPE. Rather than check the queue status directly, the base station preferably receives requests for bandwidth from the CPEs using the techniques described above.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1** shows a broadband wireless communication system adapted for use with the present invention.

FIG. **2** shows a TDD frame and multi-frame structure that can be used by the communication system of FIG. **1** in practicing the present invention.

FIG. **3** shows an example of a downlink sub-frame that can be used by the base stations to transmit information to the plurality of CPEs in the wireless communication of FIG. **1**.

FIG. **4** shows an exemplary uplink sub-frame that is adapted for use with the present bandwidth allocation invention.

FIG. **5** is a flow diagram showing the information exchange sequence used in practicing the individual polling technique of the present invention.

FIG. **6** is a flow diagram showing the individual polling technique of the present invention.

FIG. **7** shows an exemplary uplink sub-frame map that is used to facilitate the present multicast/broadcast bandwidth allocation technique.

FIG. **8** is a flow diagram showing the multicast and broadcast polling technique of the present invention.

FIG. **9** is a flow diagram showing use of a "poll-me" to stimulate polling of a CPE in accordance with the present invention.

FIG. **10** shows the message sequence that is used by the present invention in requesting polls using the "poll-me" bit.

FIG. **11** is a flow diagram showing the bandwidth request piggybacking process of the present invention.

FIG. **12** shows the downlink bandwidth allocation method used by the present invention.

FIG. **13** shows the uplink bandwidth allocation method used by the present invention.

FIG. **14** shows a simplified block diagram of an exemplary communication system adapted for use with embodiments disclosed herein.

FIG. **15***a* is a flow diagram showing one embodiment of the self-correcting bandwidth request/grant protocol method.

FIG. **15***b* is a flow diagram showing another embodiment of the self-correcting bandwidth request/grant protocol method.

FIG. **16** shows a flowchart for one embodiment of the abridged bandwidth request/grant protocol method.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION OF THE INVENTION

Throughout this description, the preferred embodiment and examples shown should be considered as exemplars, rather than as limitations on the present invention.

The embodiments of the present invention are methods and apparatuses for bandwidth request/grant protocols that allocate bandwidth in a broadband wireless communication system. One very important performance criterion of a broadband wireless communication system, and any communication system for that matter having a physical communication medium shared by a plurality of users, is how efficiently the system uses the physical medium. Because wireless communication systems are shared-medium communication networks, access and transmission by subscribers to the network must be controlled. In wireless communication systems a Media Access Control ("MAC") protocol typically controls user accesses to the physical medium. The MAC determines when subscribers are allowed to transmit on the physical medium. In addition, if contentions are permitted, the MAC controls the contention process and resolves any collisions that occur.

In the system shown in FIG. **1**, the MAC executed by software present in the base stations **106** (in some embodiments, the software may execute on processors both in the base stations and the CPE) control the transmission time for all of the CPEs **110**. The base stations **106** receive requests for transmission rights and grant these requests within the time available taking into account the priorities, service types, quality of service and other factors associated with the CPEs **110**. As described above in the background of the invention, the services provided by the CPEs **110** TDM information such as voice trunks from a PBX. At the other end of the service spectrum, the CPEs may uplink bursty yet delay-tolerant computer data for communication with the well-known World Wide Web or Internet.

The base station MAC maps and allocates bandwidth for both the uplink and downlink communication links. These maps are developed and maintained by the base station and are referred to as the Uplink Sub-frame Maps and Downlink Sub-frame Maps. The MAC must allocate sufficient bandwidth to accommodate the bandwidth requirements imposed by high priority constant bit rate (CBR) services such as T1, E1 and similar constant bit rate services. In addition, the MAC must allocate the remaining system bandwidth across the lower priority services such as Internet Protocol (IP) data services. The MAC distributes bandwidth among these lower priority services using various QoS dependent techniques such as fair-weighted queuing and round-robin queuing.

The downlink of the communication system shown in FIG. **1** operates on a point-to-multi-point basis (i.e., from the base station **106** to the plurality of CPEs **110**). As described in U.S. Pat. No. 6,016,311, the central base station **106** includes a sectored active antenna array **108** which is capable of simultaneously transmitting to several sectors. In one embodiment of the system **100**, the active antenna array **108** transmits to six independent sectors simultaneously. Within a given frequency channel and antenna sector, all stations receive the same transmission. The base station is the only transmitter operating in the downlink direction, hence it transmits without having to coordinate with other base stations, except for the overall time-division duplexing that divides time into upstream (uplink) and downstream (downlink) transmission periods. The base station broadcasts to all of the CPEs in a sector (and frequency). The CPEs monitor the addresses in the received messages and retain only those addressed to them.

The CPEs **110** share the uplink on a demand basis that is controlled by the base station MAC. Depending upon the class of service utilized by a CPE, the base station may issue a selected CPE continuing rights to transmit on the uplink, or the right to transmit may be granted by a base station after receipt of a request from the CPE. In addition to individually addressed messages, messages may also be sent by the base station to multicast groups (control messages and video distribution are examples of multicast applications) as well as broadcast to all CPEs.

Within each sector, in accordance with the present invention, CPEs must adhere to a transmission protocol that minimizes contention between CPEs and enables the service to be tailored to the delay and bandwidth requirements of each user application. As described below in more detail, this transmission protocol is accomplished through the use of a polling mechanism, with contention procedures used as a backup mechanism should unusual conditions render the polling of all CPEs unfeasible in light of given delay and response-time constraints. Contention mechanisms can also be used to avoid individually polling CPEs that are inactive for long time periods. The polling techniques provided by the present inventive method and apparatus simplifies the access process and guarantees that service applications receive bandwidth allocation on a deterministic basis if required. In general, data service applications are relatively delay-tolerant. In contrast, real-time service applications such as voice and video services require that bandwidth allocations be made in a timely manner and in adherence to very tightly-controlled schedules.

Frame Maps—Uplink and Downlink Sub-frame Mappings

In one embodiment, the base stations **106** maintain sub-frame maps of the bandwidth allocated to the uplink and downlink communication links. As described in the related and incorporated U.S. Pat. No. 6,016,311, the uplink and downlink are multiplexed in a time-division duplex (or "TDD") manner. In one embodiment, a frame is defined as comprising N consecutive time periods or time slots (where N remains constant). In accordance with this "frame-based" approach, the communication system dynamically configures the first N_1 time slots (where N is greater than or equal to N_1) for downlink transmissions only. The remaining N_2 time slots are dynamically configured for uplink transmissions only (where N_2 equals $N-N_1$). Under this TDD frame-based scheme, the downlink sub-frame is preferably transmitted first and is prefixed with information that is necessary for frame synchronization.

As is well known in the communications art, in another communication system, the uplink and downlink can be multiplexed using the well known frequency-division duplex (or "FDD") approach. FDD techniques are well known in the communication arts and thus are not described in more detail herein. However, one exemplary TDD system is described in detail herein. The invention present invention can be used in both the described TDD communication system, or in an FDD communication system.

FIG. **2** shows a TDD frame and multi-frame structure **200** that can be used by a communication system (such as that shown in FIG. **1**) in practicing the embodiments. As shown in FIG. **2**, the TDD frame is subdivided into a plurality of physical slots (PS) **204**. In the embodiment shown in FIG. **2**, the frame is one millisecond in duration and includes 800 physical slots. Alternatively, the present invention can be used with frames having longer or shorter duration and with more or fewer PSs. The available bandwidth is allocated by a base station in units of a certain pre-defined number of PSs. Some form of digital encoding, such as the well-known Reed-Solomon encoding method, is performed on the digital information over a pre-defined number of bit units referred to as information elements (PI). The modulation may vary within the frame and determines the number of PS (and therefore the amount of time) required to transmit a selected PI.

As described in the incorporated U.S. Pat. No. 6,016,311, in one embodiment of the broadband wireless communication system shown in FIG. **1**, the TDD framing is adaptive. That is, the number of PSs allocated to the downlink versus the uplink varies over time. The present bandwidth allocation method and apparatus can be used in both adaptive and fixed TDD systems using a frame and multi-frame structure similar to that shown in FIG. **2**. As described above, the embodiments disclosed herein can also be used in FDD communication systems. As shown in FIG. **2**, to aid periodic functions, multiple frames **202** are grouped into multi-frames **206**, and multiple multi-frames **206** are grouped into hyper-frames **208**. In one embodiment, each multi-frame **206** comprises two frames **202**, and each hyper-frame comprises twenty-two multi-frames **206**. Other frame, multi-frame and hyper-frame structures can be used with the present invention. For example, in another embodiment of the present invention, each multi-frame **206** comprises sixteen frames **202**, and each hyper-frame comprises thirty-two multi-frames **206**. Exemplary downlink and uplink sub-frames are shown respectively in FIGS. **3** and **4**.

Downlink Sub-frame Map

FIG. **3** shows one example of a downlink sub-frame **300** that can be used by the base stations **106** to transmit information to the plurality of CPEs **110**. The base station preferably maintains a downlink sub-frame map that reflects the downlink bandwidth allocation. The downlink sub-frame **300** preferably comprises a frame control header **302**, a plurality of downlink data PSs **304** grouped by modulation type (e.g., PS **304** data modulated using a QAM-4 modulation scheme, PS **304'** data modulated using QAM-16, etc.) and possibly separated by associated modulation transition gaps (MTGs) **306** used to separate differently modulated data, and a transmit/receive transition gap **308**. In any selected downlink sub-frame any one or more of the differently modulated data blocks may be absent. In one embodiment, modulation transition gaps (MTGs) **306** are 0 PS in duration. As shown in FIG. **3**, the frame control header **302** contains a preamble **310** used by the physical protocol layer (or PHY) for synchronization and equalization purposes. The frame control header **302** also includes control sections for both the PHY (**312**) and the MAC (**314**).

The downlink data PSs are used for transmitting data and control messages to the CPEs **110**. This data is preferably encoded (using a Reed-Solomon encoding scheme for example) and transmitted at the current operating modulation used by the selected CPE. Data is preferably transmitted in a pre-defined modulation sequence: such as QAM-4, followed by QAM-16, followed by QAM-64. The modulation transition gaps **306** contain preambles and are used to separate the modulations. The PHY Control portion **312** of the frame control header **302** preferably contains a broadcast message indicating the identity of the PS **304** at which the modulation scheme changes. Finally, as shown in FIG. **3**, the Tx/Rx transition gap **308** separates the downlink sub-frame from the uplink sub-frame which is described in more detail below.

Uplink Sub-frame Map

FIG. **4** shows one example of an uplink sub-frame **400** that is adapted for use with the present bandwidth allocation invention. In accordance with the present bandwidth allocation method and apparatus, the CPEs **110** (FIG. **1**) use the uplink sub-frame **400** to transmit information (including bandwidth requests) to their associated base stations **106**. As shown in FIG. **4**, there are three main classes of MAC control messages that are transmitted by the CPEs **110** during the uplink frame: (1) those that are transmitted in contention slots reserved for CPE registration (Registration Contention Slots **402**); (2) those that are transmitted in contention slots reserved for responses to multicast and broadcast polls for bandwidth allocation (Bandwidth Request Contention Slots **404**); and those that are transmitted in bandwidth specifically allocated to individual CPEs (CPE Scheduled Data Slots **406**).

The bandwidth allocated for contention slots (i.e., the contention slots **402** and **404**) is grouped together and is transmitted using a pre-determined modulation scheme. For example, in the embodiment shown in FIG. **4** the contention slots **402** and **404** are transmitted using a QAM-4 modulation. The remaining bandwidth is grouped by CPE. During its scheduled bandwidth, a CPE **110** transmits with a fixed modulation that is determined by the effects of environmental factors on transmission between that CPE **110** and its associated base station **106**. The downlink sub-frame **400** includes a plurality of CPE transition gaps (CTGs) **408** that serve a similar function to the modulation transition gaps (MTGs) **306** described above with reference to FIG. **3**. That is, the CTGs **408** separate the transmissions from the various CPEs **110** during the uplink sub-frame. In one embodiment, the CTGs **408** are 2 physical slots in duration. A transmitting CPE preferably transmits a 1 PS preamble during the second PS of the CTG **408** thereby allowing the base station to synchronize to the new CPE **110**. Multiple CPEs **110** may transmit in the registration contention period simultaneously resulting in collisions. When a collision occurs the base station may not respond.

By using the bandwidth allocation method and apparatus disclosed herein, scheduled uplink traffic data is bandwidth allocated to specific CPEs **110** for the transmission of control messages and services data. The CPE scheduled data is ordered within the uplink sub-frame **400** based upon the modulation scheme used by the CPEs **110**. In accordance with the present invention and in the manner described in detail below, bandwidth is requested by a CPE **110** and is subsequently granted by an associated base station **106**. All of the bandwidth allocated to a selected CPE within a given TDD frame (or alternatively an adaptive TDD frame, as the case may be) is grouped into a contiguous CPE scheduled data block **406**. The physical slots allocated for the CTGs **408** are included in the bandwidth allocation to a selected CPE **110** in the base station uplink sub-frame map.

In addition to the bandwidth that is allocated for the transmission of the various types of broadband services (i.e., the bandwidth allocated for the CPE scheduled data slots **406**), and the bandwidth allocated for CPE registration contention slots, bandwidth must also be allocated by the base station MAC for control messages such as requests for additional bandwidth allocations. As described in more detail below, in accordance with the present invention, CPEs **110** request changes to their bandwidth allocations by making bandwidth requests of their associated base stations **106**. The present inventive method and apparatus reduces the amount of bandwidth that must be set aside for these bandwidth allocation requests. In accordance with the present invention, the opportunities for requesting bandwidth are very tightly controlled. The present invention advantageously utilizes a combination of a number of techniques to tightly control the bandwidth request process. There are a number of means by which a CPE can transmit a bandwidth request message to its associated base station.

For example, one such means uses a "polling" technique whereby a base station polls one or more CPEs and allocates bandwidth specifically for the purpose of allowing the CPE(s) to transmit bandwidth requests. In accordance with this method, the polling of CPEs by the base station may be in response to a CPE setting a "poll-me bit" in an upstream direction or it may be periodic. In accordance with the present invention, periodic polls may be made to individual CPEs (referred to as "reservation-based" polling), to groups of CPEs ("multicast" polling), or to every CPE on a physical channel ("broadcast" polling). In reservation-based polling, the base station polls an individual CPE and then allocates uplink bandwidth to allow the CPE to respond with a bandwidth request. Similarly, in multicast and broadcast polling, the base station polls several CPEs and then allocates uplink bandwidth to allow the CPEs to respond with a bandwidth request. However, the CPEs must contend for the allocated bandwidth if collisions occur. Advantageously, neither the bandwidth polls nor the bandwidth allocations are in the form of explicit messages that are communicated by the base station to the CPEs. Rather, the bandwidth polls comprise unsolicited grants of bandwidth sufficient for transmitting bandwidth requests. Bandwidth allocations are implicit via bandwidth allocations occurring in the uplink sub-frame map. The polling techniques are described in more detail below with reference to FIGS. **4-10**.

As shown in FIG. **4**, a portion of the uplink bandwidth may periodically be allocated for these bandwidth allocation or CPE connection requests. The uplink sub-frame **400** includes a plurality of bandwidth request contention slots **404**. A CPE **110** must first be registered and achieve uplink synchronization with a base station before it is allowed to request bandwidth allocation. Therefore there is no need to allow for transmit time uncertainties in the length of the bandwidth request contention period. Consequently the bandwidth request contention period may be as small as a single PI, which, in one embodiment, at QAM-4 requires 6 PS. As with the registration requests, if a collision occurs, the base station may not respond to the CPE. If, however, the base station successfully receives a bandwidth request message from a CPE, it responds by allocating the CPE additional scheduled data **406** bandwidth in the uplink sub-frame **400**. The various polling techniques used by the present invention help to minimize the need to use the contention slots **404**. These techniques are described in more detail below.

Another means used by the present invention in reducing the bandwidth consumed by the bandwidth request messages is the technique of "piggybacking" bandwidth requests on bandwidth already allocated to a CPE. In accordance with this technique, currently active CPEs request bandwidth using previously unused portions of uplink bandwidth that is already allocated to the CPE. The necessity of polling CPEs is thereby eliminated. In an alternative embodiment of the present invention, bandwidth requests are piggybacked on uplink bandwidth allocated and actively being used by a data service. In accordance with this alternative, the CPE "steals" bandwidth already allocated for a data connection by inserting bandwidth requests in time slots previously used for data. The details of these piggybacking techniques are described in more detail below with reference to FIG. **11**.

Once a CPE is allocated bandwidth by the base station, the CPE, not the base station, is responsible for using the uplink bandwidth in a manner that can accommodate the services provided by the CPE. The CPE is free to use the uplink bandwidth that was allocated to it in a manner that is different than originally requested or granted by the base station. For example, the service requirements presented to a selected CPE can change after the selected CPE requests bandwidth from its associated base station. The CPE advantageously determines which services to give bandwidth to and which services must wait for subsequent bandwidth requests. To this end, the CPE maintains a priority list of services. Those services having higher priority (e.g., those services having high quality of service demands) will be allocated bandwidth before those services having lower priority (e.g., IP-type data services). If the CPE does not have sufficient bandwidth to meet its service requirements, the CPE will request additional bandwidth allocations by either setting its poll-me bit or by piggybacking a bandwidth allocation request.

One advantage of having the CPE determine how to distribute its allocated bandwidth is that it relieves the base station from performing this task. In addition, the communication overhead that is required by having the base station instruct the CPE how to distribute its allocated bandwidth is thereby eliminated, thus increasing usable system bandwidth. In addition, the CPE is in a much better position to respond to the varying uplink bandwidth allocation needs of high quality of service data services. Therefore, the CPE can better accommodate the needs of these types of service requirements than can the base station.

The various techniques used to enhance the efficiency of the bandwidth allocation request process are described in more detail below in the sub-sections that follow. Although these techniques are described in separate sub-sections, the present inventive method and apparatus advantageously uses all of the techniques in combination to reduce the bandwidth consumed by the bandwidth allocation requests.

Thus, the present invention advantageously makes use of the efficiency benefits associated with each bandwidth allocation technique. For example, although an individual polling technique is beneficial with regard to the ability to provide fast response times to bandwidth allocation requests, it is relatively inefficient with regard to the amount of bandwidth consumed by the bandwidth allocation process. In contrast, the group polling method is relatively efficient with regard to the bandwidth consumed by the bandwidth allocation process, but it is less efficient with regard to the ability to respond to bandwidth allocation requests. Use of a "poll-me" bit is relatively efficient when considered from both the bandwidth consumption and response time perspectives. In addition, the piggybacking technique further enhances bandwidth consumption efficiency by using previously unused portions of the bandwidth to send the bandwidth allocation requests. In contrast to the prior art approaches, the present invention advantageously uses all of these bandwidth allocation techniques in combination to maximize efficiency.

Polling

In one embodiment of the broadband wireless system **100** of FIG. **1** designed for use with the present invention, a CPE **110** is assigned a dedicated connection identifier (ID) when the CPE **110** first registers with the system **100**. The ID is used when the base station **106** exchanges control messages with the plurality of CPEs **110**. As described above, variations in bandwidth requirements (i.e., increases or decreases to bandwidth requirements) are necessary for all services transported by the system **100** with the exception of uncompressible constant bit rate, or continuous grant (CG) services. The bandwidth requirements of uncompressible CG services do not change between connection establishment and termination. The requirements of compressible CG services, such as channelized-T1 services, may increase or decrease depending on traffic.

In contrast, many of the data services facilitated by the system **100** of FIG. **1** are bursty and delay-tolerant. Because bandwidth is provided to these services on a demand assignment basis as needed these services are commonly referred to as Demand-Assigned Multiple Access or "DAMA" services. When a CPE **110** needs to request bandwidth for a DAMA service it transmits a bandwidth request message to the base station **106**. The bandwidth request messages communicate the immediate bandwidth requirements for the DAMA service. The bandwidth requirements can and typically do vary over time. The quality of service or "QoS" for the DAMA connection is established when the CPE connection is initially established with the base station. Therefore, the base station has the ability to access or "look-up" the QoS for any DAMA service that it is currently accommodating.

As described above, in accordance with the present invention, the CPEs **110** have a number of different techniques available to them for communicating bandwidth request messages to their associated base stations. One such technique is by transmitting a bandwidth request message in response to being polled by a base station. In accordance with the polling technique taught by the present invention, the base station allocates bandwidth to selected CPEs specifically for the purpose of making bandwidth requests. The bandwidth allocations may be to individual CPEs or to groups of CPEs. As described in more detail below in the sub-section that describes the group polling technique, allocations to groups of CPEs define bandwidth request contention slots that are used in resolving bandwidth request collisions. Advantageously, the bandwidth allocations are not made in the form of explicit messages, but rather they are made in the form of bandwidth allocation increases in the transmitted map describing the uplink sub-frame **400** (FIG. **4**). Polling is performed on a per-CPE basis, bandwidth is requested on a per-connection-ID basis, and bandwidth is allocated on a per-CPE basis. These concepts are described in more detail below.

Reservation-based Polling Technique (Individual Polling)

In accordance with the present inventive method and apparatus, when a CPE is polled individually, no explicit message is transmitted to poll the selected CPE. Rather, the CPE is allocated bandwidth in the uplink sub-frame map that is sufficient to allow the CPE to respond with the bandwidth request. Specifically, the base station allocates bandwidth in the CPE scheduled data block **406** (FIG. **4**) for the selected CPE that is sufficient to allow the selected CPE to respond with a bandwidth request message. If the selected CPE does not require more bandwidth, it returns a request for zero bytes. A zero byte request (rather than no request) is used in the individual polling process because explicit bandwidth for a reply is allocated.

In accordance with the present invention, only inactive CPEs and active CPEs that explicitly request to be polled are eligible for individual polling. Active CPEs that do not set their respective "poll-me" bits in the MAC packet header will not be polled individually. These restrictions are imposed upon the bandwidth request process by the present invention and they advantageously save bandwidth compared with polling all of the CPEs individually. In one embodiment of the present invention, active CPEs respond to polling using the modulation scheme currently in use. However, inactive CPEs may respond using a QAM-4 or similarly robust modulation scheme to ensure that their transmission is sufficiently robust to be detected by the base station even under adverse environmental conditions.

The present invention advantageously ensures timely responses to requests for more bandwidth for a constant bit rate service such as a channelized T1 service in which channels may be added or dropped dynamically. To ensure that the base station responds quickly to requests for more bandwidth for a constant bit rate service, the uplink bandwidth allocated to a constant bit rate service that is not currently operating at a maximum rate is made sufficiently large to accommodate the service's current rate and a bandwidth request.

The information exchange sequence for individual polling is shown in the flow diagram of FIG. **5**. As shown in FIG. **5**, the base station preferably has several layers of control mechanisms or protocol stacks **502**, **504** and **506** that control, among other things, the bandwidth request and allocation process. The base station MAC is sub-divided into two sub-domains: (1) the HL-MAA MAC domain **504** and the LL-MAA Mac domain **506**. The LL-MAA MAC domain spans exactly a physical channel. Each physical channel requires an instance of the LL-MAA MAC domain. The HL-MAA MAC domain spans multiple physical channels, typically all in the same sector. A MAC domain comprises an HL-MAA MAC domain and the LL-MAA MAC domains associated with the physical channels within the HL-MAA MAC domain.

As shown in FIG. **5**, the base station individually polls (as indicated by control arrow **508**) a CPE by allocating bandwidth sufficient for the CPE to respond with a bandwidth request message. This bandwidth is allocated in the uplink sub-frame **400**. If the CPE MAC **510** determines that there is data to be sent for a selected connection k (typically determined by being instructed by a higher CPE control layer **512** via a control path **514**), then the CPE MAC control mechanism issues a bandwidth request **516** to the base station MAC **506**. If there is insufficient bandwidth available to the CPE **110** as determined by the base station's LL-MAA **506**, the bandwidth request will not be granted. Else, the bandwidth request will be granted and this will be implicitly communicated to the CPE MAC **510** by the base station allocating additional bandwidth to the CPE in the uplink sub-frame **400**. This is shown in FIG. **5** via the control path **518**. The CPE will then begin transmitting data to the base station over the uplink using the bandwidth that has been allocated to it.

FIG. **6** is a flow diagram showing the individual polling technique **600** provided by the present invention. As shown in FIG. **6**, the method starts at decision STEP **602** to determine whether bandwidth is available for the purpose of individually polling the CPEs. If no more bandwidth is available for individually polling the CPEs **110** then the method proceeds to STEP **604** and initiates a multicast or broadcast polling method. This multicast and broadcast polling method is described in detail in the sub-section below. However, if sufficient bandwidth is available for the purpose of individually polling CPEs, the method proceeds to a decision STEP **606** whereat a determination is made whether there are any un-polled active CPEs that have a "poll-me" bit set. If so, the method proceeds to a control point **608**. If not, the method proceeds to a decision STEP **610** whereat it determines whether there are any un-polled inactive CPEs present. If so, the method proceeds to the control point **608**. If not, the method proceeds to a control point **612**.

The present inventive method proceeds from the control point **608** to STEP **614** to individually poll the selected CPE. Thus, the method ensures that only un-polled active CPEs requesting more bandwidth (by setting their respective "poll-me" bits) and inactive CPEs are individually polled. This reduces bandwidth as compared with a polling method that would individually poll all CPEs.

As shown in FIG. **6**, at STEP **614** the base station initiates the polling of the selected CPE and marks the CPE as polled. This is shown diagrammatically in FIG. **6** in the caption box **614'**. The caption box **614'** of FIG. **6** shows the downlink sub-frame map **300** described above in FIG. **3**. The MAC control portion **314** of the MAC frame control header **302** preferably includes an uplink sub-frame map **400'**. The uplink sub-frame map **400'** is communicated to the CPE MAC when the base station transmits this information to the CPE via the downlink. As shown in FIG. **6**, and responsive to the polling

STEP **614**, the base station MAC allocates additional bandwidth to the selected CPE (in FIG. **6** this CPE is referred to as CPE "k") in the uplink. This increased bandwidth allocation is communicated to the CPE k via the uplink sub-frame map **400'**. Thus, no additional bandwidth is needed to respond to the need to poll the selected CPE.

As shown in FIG. **6**, the method then returns to the decision STEP **602** to determine whether there is more bandwidth available for individually polling the CPEs. When it is determined (at the decision STEPS **606** and **610**, respectively) that there are no active CPEs having a poll-me bit set and that there are no un-polled inactive CPEs present, the method proceeds to a decision STEP **616**. At the decision STEP **616**, the method determines whether any individual polls were performed. If not, the method proceeds to a control point **618** and the method subsequently terminates at the termination step **620**. However, if individual polls were performed, the method proceeds to a STEP **622** to await the individual bandwidth requests from the CPE that was polled (e.g., CPE "k"). As shown in the caption **622'** of FIG. **6**, this bandwidth request **430** is generated by the polled CPE (e.g., CPE "k") during the CPE scheduled data block **406** scheduled for the selected CPE in the uplink sub-frame **400**. In one embodiment, all data includes a header that indicates the type of data being transmitted. For example, in this embodiment, control messages have associated CPE-unique connection identifiers that are assigned to them when the CPE registers. The structure of the control messages allows a base station to determine that a control message is a bandwidth request.

As shown in FIG. **6**, the method proceeds from STEP **622** to a decision STEP **624** to determine whether any bandwidth requests were received. If not, the method terminates. However, if so, the method proceeds to a STEP **626** whereat a bandwidth allocation method is initiated. As described in more detail below the base station uses a bandwidth allocation method to allocate bandwidth to the requesting CPE. In one embodiment, a self-correcting bandwidth request/grant protocol is also used. Details of the self-correcting bandwidth request/grant protocol are described in more detail below with reference to FIGS. **15***a* and **15***b*. In another embodiment, the bandwidth allocation method uses "padding packets" to improve bandwidth allocation efficiencies in wireless communication systems. This embodiment is described in more detail below with reference to FIG. **16**. The bandwidth allocation is indicated to the CPE by making appropriate changes to the uplink sub-frame map **400'**. The method then terminates at STEP **620**.

Contention-based Polling Technique (Multicast and Broadcast Polling)

As described above with reference to STEP **604** of the individual polling method of FIG. **6**, if there is not sufficient bandwidth available for the purpose of individually polling the CPEs, the present invention may be used to poll the CPEs in multicast groups and a broadcast poll may be issued by the base station. Also, if there are more inactive CPEs than there is bandwidth available to individually poll them, some CPEs may be polled in multicast groups and a broadcast poll may be issued.

In accordance with one embodiment, the addressing of CPEs is preferably performed as follows: each CPE is assigned a unique permanent address (e.g., in one embodiment the CPE has a 48-bit address) that is used in the registration process; and each CPE is also given a basic connection ID (e.g., in one embodiment the CPE is given a 16-bit basic connection ID and a 16-bit control connection ID during the registration process). Each service that is provisioned for a selected CPE is also assigned a connection ID. Connection IDs are generated by the base station MAC (specifically, by the base station HL-MAA) and are unique across an HL-MAA MAC domain. The basic connection ID that is assigned when the CPE is registered with a base station is used by the base station MAC and the CPE MAC to exchange MAC control messages between the CPE and the base station. The control connection ID (also assigned during registration) is used by the base station and the CPE to exchange control and configuration information between the base station and the CPE higher levels of control.

In accordance with one embodiment, certain connection IDs are reserved for multicast groups and broadcast messages. Of all of the addresses available a portion of them are preferably reserved for multicast use. For example, in one embodiment of the present invention, if the four most-significant bits of the connection ID are set to logical ones (hex "Fxxxx") the address is interpreted as being set aside for multicast use. In this embodiment, a total of 4K distinct multicast addresses are available. One example of such a multicast use is for the distribution of a video service. In one preferred embodiment, the connection ID used to indicate a broadcast to all stations is (0xFFFF) (i.e., all 16 bits are set to a logical one).

Similar to the individual polling technique described above with reference to FIGS. **5** and **6**, the multicast polling message is not explicitly transmitted by the base station to the CPE. Rather, the multicast poll message is implicitly transmitted to the CPE when the base station allocates bandwidth in the uplink sub-frame map. However, rather than associating allocated bandwidth with a CPE's basic connection ID as done when performing an individual poll, the base station associates the allocated bandwidth to a multicast or broadcast connection ID. This multicast/broadcast bandwidth allocation is shown in the multicast/broadcast uplink sub-frame map **400"** shown in FIG. **7**. It is instructive to compare the uplink subframe **400** (FIG. **4**) used by the base station when individual polling the CPEs with the uplink sub-frame map **400"** of FIG. **7**. FIG. **7** shows the uplink sub-frame map which is transmitted in the MAC control portion of the downlink.

As shown in FIG. **7**, the multicast/broadcast uplink subframe map **400"** used includes registration contention slots **402"** that map the registration contention slots **402** of FIG. **4**. However, rather than associating allocated bandwidth with a selected CPE's basic connection ID, the allocated bandwidth is associated with a reserved registration ID. As shown in FIG. **7**, the uplink sub-frame map **400"** preferably includes a plurality of multicast group bandwidth request contention slots **404"**, **404'"**, etc. The uplink sub-frame map **400"** also includes broadcast bandwidth request contention slots **410**. Finally, similar to the uplink sub-frame of FIG. **4**, the uplink subframe map used by the present invention to initiate multicast or broadcast polls includes a plurality of CPE scheduled data blocks **406"**, **406'"**, etc., that are used to transport uplink traffic data.

In accordance with the present inventive method and apparatus, when a poll is directed to a multicast or broadcast connection ID, CPEs belonging to the polled group request bandwidth using the bandwidth request contention slots (either the multicast contention slots for the group specified or the broadcast bandwidth request contention slots **410**) allocated in the uplink sub-frame map **400"**. In order to reduce the likelihood of collisions only CPE's needing bandwidth are allowed to reply to multicast or broadcast polls. Zero-length bandwidth requests are not allowed in the bandwidth request contention slots. In one embodiment, CPEs transmit the bandwidth requests in the bandwidth request contention slots (e.g., contention slots **404**) using QAM-4 modulation. In this embodiment, the contention slots are sized to hold a 1-PS preamble and a bandwidth request message. Due to physical resolution characteristics, the message requires 1 PI (or 6 PS) using QAM-4 modulation. In this embodiment, multiple bandwidth request messages from the same CPE fit in a single bandwidth request contention slot without increasing the bandwidth utilization or the likelihood of collisions occurring. This allows the same CPE to make multiple bandwidth requests in the same slot.

If an error occurs when performing either a multicast or broadcast poll (such as the detection of an invalid connection ID) the base station transmits an explicit error message to the CPE. If the base station does not respond with either an error message or a bandwidth allocation within a predefined time period, the CPE will assume that a collision occurred. In this case the CPE uses a selected pre-defined contention resolution process. For example, in one preferred embodiment, the CPE uses the well known "slotted ALOHA" contention resolution process to back off and try at another contention opportunity.

Contention Resolution Process

Contention is necessary when there is insufficient time to poll all of the CPEs individually within a suitable interval. The base station is able to define contention periods both for multicast groups and also for all CPEs generally (i.e., broadcast). After CPE scheduled data, control messages, and polling are allowed for, the base station allocates all unused time in the upstream part of the TDD frame to contention, either for bandwidth requests or for registration purposes. Typically the bandwidth request interval will be many PIs long (e.g., 1 PI=6 PS using QAM-4 modulation). The CPEs must transmit their requests at a random time (on burst boundaries) within this interval to reduce the likelihood of collisions occurring.

In accordance with the present invention, a CPE needing to transmit in a request interval preferably randomly selects a PI within the interval, and makes a request in the associated starting PS. This randomization minimizes the probability of collisions. A collision is presumed if there is no response from the base station to the request within a pre-defined time period. If the base station does not respond within the predefined time period the collision resolution process of the present invention is initiated.

One embodiment uses the following resolution process: Assuming that the initial backoff parameter is i and that the final backoff parameter is f,

On the first collision, the CPE waits a random interval between zero and 2^i contention opportunities and then tries again.

If another collision occurs, then the interval is doubled and the CPE tries again, repeating until the interval 2^f is reached.

If the CPE is still unsuccessful, an error is reported to the system controller and the contention process is aborted. Other contention resolution mechanisms can be used to practice the present invention. For example, the well-known Ternary tree mechanism could be used to resolve contentions.

FIG. **8** is a flowchart showing the multicast and broadcast polling method **800** of the present invention. As shown in FIG. **8**, the group polling method **800** proceeds from an initial step **802** to a decision STEP **804** whereat the method determines whether there is sufficient bandwidth available for multicast polls. If sufficient bandwidth is available for multicast polls, the method proceeds to a STEP **806** to poll the next multicast group in the MAC control portion **314** of the MAC frame control header **302**. However, if there is insufficient bandwidth available to perform a multicast poll, the method proceeds to a decision STEP **808** whereat the method determines whether there is sufficient available bandwidth for performing a broadcast poll. If so, the method proceeds to a STEP **810**. If not, the method proceeds to a decision STEP **812**.

As shown in FIG. **8**, at the STEP **810** a broadcast poll is initiated by placing the broadcast poll in the MAC control portion **314** of the MAC frame control header **302**. Similar to the individual polling technique, the multicast poll message is implicitly transmitted to the CPE by allocating bandwidth in the uplink sub-frame map **400"**. The allocated bandwidth is associated with a multicast or broadcast connection ID.

At the decision STEP **812** the method determines whether a broadcast or multicast poll was initiated. If so, the method proceeds to a STEP **814** whereat the method monitors the appropriate bandwidth request contention slots (e.g., as defined by the bandwidth contention slot descriptions **404"**, **404'"**, and the broadcast bandwidth request contention slot descriptions **410** of FIG. **7**). If no broadcast or multicast poll was initiated, the method proceeds to control point **816** and subsequently terminates at a termination STEP **818**.

The method proceeds from the monitoring STEP **814** to a decision STEP **820** to determine whether valid (i.e., non-colliding) bandwidth requests were detected. If no valid bandwidth requests were detected at STEP **820**, the method proceeds to the control point **816** and terminates at termination STEP **818**. However, if the method detects valid bandwidth requests, the method proceeds from STEP **820** to STEP **822**. At STEP **822** the method uses a convenient bandwidth allocation algorithm to allocate bandwidth to the CPE that requested bandwidth. The preferred bandwidth allocation algorithm is described below in more detail with reference to FIGS. **12-13**. The bandwidth is allocated in the uplink subframe map **400"** as shown in FIG. **8**.

Poll-Me Bit

As described above with reference to FIGS. **3-8**, and in accordance with the present invention, a currently active CPE sets a "poll-me" bit or a "priority poll-me" in a MAC packet in order to indicate to the base station that it requires a change in bandwidth allocation. For example, in one embodiment of the present invention, a selected CPE requests a poll by setting a poll-me ("PM") bit in the MAC header. Similarly, in accordance with the present invention, a selected CPE sets a priority poll-me ("PPM") bit in the MAC header in order to indicate that a priority poll is desired.

In order to reduce the bandwidth requirements associated with individually polling every active CPE, the active CPEs are individually polled if and only if one of the poll-me bits is set by the CPE. When the base station detects a request for polling (when the CPE sets its poll-me bit), the individual polling technique shown in FIG. **9** is activated in order to satisfy the request. The procedure by which a CPE stimulates a base station to poll the CPE is shown in FIG. **9**. In an alternative embodiment, multiple packets having "poll-me" bits set indicate that the CPE needs to make bandwidth allocation requests for multiple connections.

FIG. **9** is a flow chart that shows how the poll-me bit is used to stimulate polling. As shown in FIG. **9**, the method first determines at a decision STEP **902** whether the piggybacking technique described in more detail below has been exhausted. If not, the method proceeds to STEP **904** and attempts to perform "piggybacking" first. The method then proceeds to a STEP **906** whereat the connection is set equal to a first connection. In this manner, the poll-me bits are scanned for each connection within the CPE. The method shown in FIG. **9** then proceeds to a decision STEP **908** to determine whether any bandwidth needs exist. If not, the method proceeds to a STEP **916** and scans for the next connection. If a bandwidth need exists, the method proceeds to a decision STEP **910**. At STEP

910 the method determines whether any more packets are available for accommodating the poll-me bit. If not, the method terminates at the STEP 910. However, if packets are available, the method proceeds to a STEP 912 and sets a poll-me bit in an available packet.

FIG. 10 shows the message sequence that is used by the present invention in requesting polls using the "poll-me" bit described above. As shown in FIG. 10 at data connection 930, the CPE initiates a polling sequence by setting its associated poll-me bit in the MAC header. The base station MAC responds via data message 932 by individually polling the selected CPE. This response is made by allocating bandwidth to the selected CPE in the uplink sub-frame map as shown in FIG. 10. The selected CPE subsequently responds with a bandwidth request as shown in communication path 934. In response to the CPE's bandwidth request, the base station grants bandwidth and allocates bandwidth to the CPE in the uplink sub-frame map as shown in communication path 936. The selected CPE then transmits its data to the base station via an associated connection link.

"Piggybacking" Technique

As described above with reference to the present inventive method and apparatus, in order to further reduce overhead bandwidth necessary for the bandwidth allocation process, currently active CPEs may "piggyback" a bandwidth request (or any other control message) on their current transmissions. The CPEs accomplish this piggybacking of bandwidth by using unused bandwidth in TC/PHY packets of existing bandwidth allocations. The procedure for using excess bandwidth in this manner is shown in FIG. 11.

As shown in FIG. 11, the method initiates the piggybacking process at STEP 950. The method proceeds to a decision STEP 952 to determine whether the CPE requires additional bandwidth. If so, the method proceeds to a decision STEP 954, if not, the method proceeds to a termination STEP 964 whereat the method terminates. At the decision STEP 954 the method determines whether any unused bytes exist in the current allocation. If so, the method proceeds to insert bandwidth requests into the unused bytes at STEP 956. If not, the method proceeds to a decision STEP 958. At the decision STEP 958, the method determines whether any packets at all are allocated to the CPE. If there are no packets found at the decision STEP 958, the method proceeds to STEP 960. However, if packets are allocated, the method proceeds to a STEP 962 whereat the CPE sets its poll-me bit. The method then proceeds to the STEP 960 whereat the CPE awaits polling by the associated base station. The method then terminates at the STEP 964.

Bandwidth Allocation

As described above, the base station MAC is responsible for allocating the available bandwidth of a physical channel on the uplink and the downlink. Within the uplink and downlink sub-frames, the base station MAC scheduler allocates the available bandwidth between the various services depending upon the priorities and rules imposed by their quality of service (QoS). Additionally, the higher control sub-layers of the base station MAC allocate across more than one physical channel.

Downlink Bandwidth Allocation—One Embodiment

The downlink bandwidth is allocated as shown in FIG. 12. The base station MAC maintains a set of queues for each physical channel that it serves. Within each physical channel queue set, the base station maintains a queue for each QoS. The queues hold data that is ready to be transmitted to the CPEs present on the physical channel. The higher layers of the base station protocol stack are responsible for the order in which data is place in the individual queues. The base station higher control layers are free to implement any convenient fairness or traffic shaping algorithms regarding the sharing of access between connections at the same QoS, without impacting the base station lower MAC control layers. Once data is present in the queues it is the responsibility of the base station lower levels of control (e.g., the BS LL-MAA of FIGS. 5 and 10) to allocate bandwidth based on the QoS.

In one embodiment of the present invention, in determining the amount of bandwidth to allocate at a particular QoS for a particular CPE, the base station takes into account the QoS, modulation, and the fairness criteria used to keep an individual CPE from using up all available bandwidth. Bandwidth is preferably allocated in QoS order. If there is a queue that cannot be transmitted entirely within a particular TDD frame, a QoS specific fairness algorithm, such as fair-weighted queuing, is used within that queue. Each connection is given a portion of the remaining available bandwidth based upon its relative weight. The derivation of weights is QoS-dependant. For example, ATM traffic may be weighted based upon contractual bandwidth limits or guarantees, while IP connections may all receive identical weights. Once the bandwidth is allocated the data is transmitted in a manner whereby the data is sorted by modulation type.

Uplink Bandwidth Allocation—One Embodiment

The uplink bandwidth allocation method is very similar to the downlink bandwidth allocation method described above with reference to FIG. 12. However, rather than being maintained by the base station, the data queues are distributed across and maintained by each individual CPE. Rather than check the queue status directly, the base station preferably receives requests for bandwidth from the CPEs using the techniques described above with reference to FIGS. 3-11. Using these bandwidth requests, the base station reconstructs a logical picture of the state of the CPE data queues. Based on this logical view of the set of queues, the base station allocates uplink bandwidth in the same way as it allocates downlink bandwidth. This uplink bandwidth allocation technique is shown in FIG. 13.

As described above, the bandwidth allocated to any selected CPE is transmitted to the selected CPE in the form of bandwidth being allocated in the uplink sub-frame map. Starting at a point in the TDD, the uplink sub-frame map allocates a certain amount of bandwidth to the selected CPE. The selected CPE then allocates this bandwidth across its connections. This allows the CPE to use the bandwidth in a different manner than requested if it receives higher priority data while awaiting the bandwidth allocation. As described above, the bandwidth allocations are in a constant state of change owing to the dynamic nature of bandwidth requirements. Consequently, a selected CPE may receive unsolicited modifications to the bandwidth granted on a frame-by-frame basis. If the selected CPE is allocated less bandwidth for a frame than is necessary to transmit all waiting data, the CPE must use the QoSs and fairness algorithms to service its queues. The CPE may "steal" bandwidth from lower QoS connections to piggyback request for more bandwidth using the piggybacking technique described above. TDM connections not already at maximum bandwidth are allocated enough extra bandwidth in the uplink to piggyback a request for additional bandwidth.

QoS Specific Fairness Algorithms

Data for transmission on the uplink and the downlink is preferably queued by quality of service (QoS) designations. The data is transmitted in order of a QoS queue priority as described above. As the queued data is transmitted, there may be a QoS queue for which there is insufficient bandwidth to transmit all queued data during the current TDD frame. When this situation occurs, a QoS specific fairness algorithm is initiated to ensure fair handling of the data queued at that QoS. There are 3 basic fairness algorithms that can be implemented: (1) Continuous Grant; (2) Fair-weighted queuing; and (3) Round Robin.

The MAC preferably does not police connections for bandwidth usage. Policing should be performed by higher control layers. The MAC assumes that all pending data has met contractual restrictions and can be transmitted. Continuous Grant queues have the simplest fairness algorithm. All data in these queues must be sent every TDD frame. Insufficient bandwidth indicates an error in provisioning.

Fair Weighted Queuing

Fair weighted queuing requires that all connections at a given QoS have a weight assigned to them to determine the percentage of the available bandwidth they are eligible to receive. This weight value is preferably derived from one of three data rate parameters, depending upon the contractual parameters of the provisioned connection. These three parameters are: (1) Data Pending; (2) Guaranteed Rate; and (3) Average Rate.

Real-time VBR connections are established as DAMA connections with fair-weighted queuing based upon data pending. For a QoS queue of this type in a TDD frame having insufficient bandwidth to transmit all of the data in the queue, a weight for each connection in the queue is determined. In one embodiment, this weight is the amount of data pending for the connection expressed as a percentage of the total data pending in the queue. Because the amount of data pending is dynamic, the weights for these types of queues must be determined every TDD frame where there is insufficient bandwidth to send all data in the affected queue.

For DAMA connections contracted at a guaranteed rate the weights are calculated based on the guaranteed rate. In this case, the weight preferably is expressed as a percentage of the total guaranteed rate of all connections with data pending in the queue. Because the guaranteed rate is provisioned the weights need not be determined each TDD frame where they are used. Rather, the weights for a queue are only determined when there is a provisioning change (i.e., a new connection, a change in connection parameters, or a connection termination) for one of the connections in the queue.

For DAMA connections contracted at an average rate the weights are preferably calculated based on the average rate. The weight is the average rate expressed as a percentage of the total average rate of all connections with data pending in the queue. Because the average rate is provisioned the weights need not be determined each TDD frame where they are used. Rather, the weights for a queue are only recalculated when there is a provisioning change for one of the connections in the queue.

In all of the cases described above, the granularity of the bandwidth allocations may be too coarse to provide a perfect percentage-based weighted allocation across the connections in the queue. This may result in some queues not receiving any bandwidth in a particular TDD frame. To ensure that the occurrence of this condition is fairly distributed across the connections in the queue, the connection that did not receive bandwidth is given priority the next time the insufficient bandwidth condition exists for the queue. For queues with weights based upon guaranteed or average rates some connections may not have sufficient data pending to use all of the bandwidth that they are entitled to based upon their calculated weight. In these cases, the connection's unused bandwidth is fairly distributed across the connections having excess data pending.

Some QoSs require that data be aged. For queues at these QoSs there is an associated queue of one step higher priority. If data is not transmitted by the provisioned aging parameter, the data is moved to the higher QoS queue and given priority over newer data in the original queue regardless of the relative weights of the connections.

Round Robin

The Round Robin fairness algorithm is used for best effort connections where all connections have equal weight. When insufficient bandwidth exists to transmit all data in the queue in a particular TDD frame connections are allocated bandwidth in a round-robin fashion with each connection receiving a block of bandwidth up to a queue-specific maximum. Connections that did not receive bandwidth are given priority the next time the insufficient bandwidth condition exists.

Bandwidth Allocation Algorithm

For each TDD frame, the base station allocates the downlink portion of the TDD frame and it performs an estimate of the uplink traffic to allocate uplink bandwidth to the CPEs. The CPEs individually allocate their allotted bandwidth across their pending data connections.

Base Station Downlink

As shown in FIG. **2**, in one preferred embodiment of the present invention, based on the ATDD split (i.e., the percentage of bandwidth allocated to the uplink and downlink) the base station has some number of the 800 PS in the TDD frame available for downlink transmissions. The downlink bandwidth allocation algorithm preferably proceeds as follows.

First, the base station allocates PSs to the PI for PHY Control and enough PSs for at least 1 PI for the MAC Control. The base station preferably performs uplink bandwidth allocation before downlink bandwidth allocation in order to determine the number of PIs to allocate for the MAC Control. In one preferred embodiment, the PHY Control and MAC Control are always sent using QAM-4 modulation.

For connections with downlink continuous grant data pending, the base station determines the number of PIs required to transmit the data. This number is then converted to PSs as a function of the modulation used for the CPE associated with each connection. For each remaining QoS or until available bandwidth is entirely allocated, the base station determines if there is enough bandwidth to satisfy the entire need of the QoS queue. If so, the base station allocates the required bandwidth. Otherwise, if there is not enough bandwidth to satisfy the queue, the base station implements the queue-specific fairness algorithm described above.

Base Station Uplink

In one preferred embodiment, based upon the ATDD split described above with reference to FIG. **2**, the base station has a pre-determined number of PSs in the TDD frame available for uplink transmissions. The base station must maintain an estimate of the data and control messages pending at each QoS for the CPEs that it serves. The base station estimates the data traffic based upon the bandwidth requests received from the CPEs and based upon an observation of actual data traffic. The base station estimates the uplink control message traffic based upon the protocols currently engaged (i.e., connection establishment, "poll-me" bit usage, etc.) and based upon the base station's polling policy (i.e., individual, multicast, and broadcast). The uplink bandwidth allocation algorithm proceeds as follows.

For connections with uplink continuous grant data pending, the base station preferably determines the number of PIs required to transmit the data. This number is then converted to a number of PSs as determined by the modulation used for the CPE associated with each connection. Continuous grant connections having a current bandwidth that is less than the maximum bandwidth are always allocated uplink bandwidth that is the smaller of: 1) their maximum bandwidth or 2) their current bandwidth plus the bandwidth necessary to send a CG bandwidth change message.

For each remaining QoS, or until available bandwidth is entirely allocated, the base station determines if there is bandwidth sufficient to satisfy the entire need of the QoS queue and it then allocates the required bandwidth. Otherwise, if there is not bandwidth sufficient to satisfy the queue, the base station implements the queue-specific fairness algorithm described above.

CPE Uplink

As described above, for each TDD frame, the CPEs are allocated a portion of the uplink sub-frame in which to transmit their respective data. Because the bandwidth requirements of the CPE may have changed since the base station received the bandwidth request information that it used to allocate the uplink bandwidth, the CPEs themselves are responsible for allocating their allotted bandwidth based upon their current bandwidth requirements. That is, the CPEs are not constrained to distribute allocated bandwidth to their data connections in the same manner that the CPE used in requesting the bandwidth from the base station. The CPE's uplink bandwidth allocation algorithm preferably proceeds as follows.

For connections having uplink continuous grant data pending, the CPE determines the number of PIs that are required to transmit the data. This number is then converted to a PS number based upon the modulation scheme used by the CPE. For each remaining QoS, or until available bandwidth is entirely allocated, the CPE determines if there is bandwidth sufficient to satisfy the entire need of the QoS queue. If so, the CPE allocates the required bandwidth. Otherwise, if there is not bandwidth sufficient to satisfy the queue, the CPE implements the queue-specific fairness algorithm described above.

Exemplary Communication System

An exemplary communication system adapted for use with the bandwidth allocation protocols is now described. FIG. **14** shows a simplified block diagram of an exemplary communication system adapted for use with the embodiments described herein. As shown in FIG. **14**, the exemplary communication system **1400** comprises a base station **1450** and a plurality of CPEs **1410-1440**. Although only three CPEs are shown in FIG. **14** one skilled in the communication art shall recognize that the base station **1450** is typically associated with hundreds of CPEs. The base station (BS) **1450** comprises a base station modem **1452**, a base station CPU **1454** and a CPU/Modem interface **1456**.

As shown in FIG. **14**, the BS **1450** communicates with the plurality of CPEs **1410-1440** via the BS modem **1452**. The BS modem **1452** establishes an uplink and downlink with any of the plurality of CPEs **1410-1440** in a manner well known to one skilled in the communication art. As described above, the BS **1450** transmits information such as bandwidth allocation messages and CG and ATM data to the plurality of CPEs **1410-1440** on the downlink. On the uplink the BS receives information such as bandwidth request messages and CG and ATM data from the plurality of CPEs **1410-1440**.

In one embodiment, the exemplary communication system **1400** utilizes "padding packets" and the well-known TDMA multiplexing scheme. In the well-known TDMA multiplexing scheme, the BS **1450** designates a portion of its uplink sub-frame (i.e., bandwidth) to an associated CPE. In accordance with the exemplary communication system **1400**, the associated CPE transmits data to the BS **1450** on the uplink. When a CPE does not have enough uplink data to utilize its entire portion of bandwidth, it transmits padding packets to "pad" or fill the unused portion of bandwidth.

The BS modem **1452** and the BS CPU **1454** transmit uplink and downlink data via the CPU/Modem interface **1456**. Examples of uplink and downlink data include control messages, bandwidth request messages, bandwidth allocation messages, CG data, DAMA data, padding packets and flag packets. Embodiments of the bandwidth request protocols will now be described.

Self-Correcting Bandwidth Request/Grant Protocol

The self-correcting bandwidth request/grant protocol provides a mechanism for efficiently allocating bandwidth. The self-correcting bandwidth request/grant protocol can be used in both TDD and FDD communication systems. The self-correcting bandwidth request/grant protocol utilizes both incremental and aggregate bandwidth requests. CPEs primarily transmit incremental bandwidth requests followed by periodically transmitting aggregate bandwidth requests. The use of incremental bandwidth requests reduces risks that a base station will erroneously issue duplicate bandwidth allocations to the same CPE for the same connection. Race conditions that may have occurred when using only aggregate bandwidth requests are eliminated by requiring the CPEs to request bandwidth in an incremental manner. However, use of periodic aggregate bandwidth requests (that express the current state of their respective connection queues) allows the bandwidth allocation method and apparatus to be "self-correcting".

The method is self-correcting because errors that may have been introduced due to lost incremental bandwidth requests are corrected by the base stations when the aggregate bandwidth requests are received and processed. Upon receipt of the aggregate bandwidth requests, the base stations reset their records to correctly reflect the current bandwidth requirements of their associated CPEs. Periodic use of aggregate bandwidth requests provides a self-correcting bandwidth allocation protocol yet without the bandwidth overhead (e.g., bandwidth required to transmit acknowledgements) typically associated with the prior art self-correcting protocols (such as the guaranteed delivery protocols).

In wireless communication systems, DAMA services are allocated bandwidth on a demand-assignment basis. When a CPE requires additional bandwidth on a DAMA connection, it transmits a bandwidth request message to its associated base station. Depending upon the mode used by the communication system, the CPE transmits the bandwidth request directed toward either the specific connection or to the QoS class to which the connection belongs. In an exemplary embodiment, the QoS for the connection is established at the time that the connection is established and is subsequently referenced by the base station or the CPE depending upon the mode of operation being used.

In accordance with one embodiment of a communication system adapted for use with the self-correcting protocol, CPEs preserve bandwidth with the assistance of various bandwidth request timers (e.g., bandwidth request timers "MT**5**" and "MT**10**", described in more detail below with reference to FIGS. **15***a* and **15***b*). The bandwidth request timers are used in these communication systems to prevent the CPEs from requesting bandwidth too frequently. Because the bandwidth request timer values can be varied dynamically based upon certain system characteristics, the bandwidth request timers are considered to be "dynamic." For example, the bandwidth request timer values can be dynamically varied to reflect the current flow of data received via the CPE ports. The bandwidth request timers can also dynamically vary their values based upon the connection priority in the system. In systems using dynamic bandwidth request timers, the CPEs typically use the bandwidth request timers in accordance with the following description.

At the beginning of a data burst, the CPE resets the dynamic bandwidth request timers. During the data burst, the CPE decreases or shortens the bandwidth request timer values (thereby reducing the time period between the transmission of consecutive bandwidth requests) based upon the current data flow and the connection priority of the CPE. The CPEs prevent bandwidth requests from being transmitted too frequently by controlling the bandwidth request timer values. More specifically, the CPEs ensure that the bandwidth request timer values are never decreased below a pre-determined threshold value, thereby ensuring that the bandwidth requests are not transmitted too frequently. One of ordinary skill in the communications art shall recognize that "static" timers can also be used to practice the present invention. Unlike the dynamic timers described above, static timers are not dynamically varied based upon certain system characteristics. Therefore, those skilled in the art shall recognize that the utilization of dynamic bandwidth request timers is not meant to be a limitation as static timers can also be used.

In accordance with one embodiment, the CPE has an option of transmitting a message that contains either the total immediate requirements of the DAMA connection/QoS (i.e., the aggregate bandwidth requirements) or the incremental immediate requirements of the DAMA connection/QoS (i.e., the incremental bandwidth requirements). The CPE primarily transmits an incremental bandwidth request, but periodically (e.g., every fifth request) transmits an aggregate bandwidth request. The incremental bandwidth requests specify the increased or decreased bandwidth requirements for the CPE (i.e., the amount that the bandwidth requirements have changed since the most recent bandwidth request was made). The aggregate bandwidth requirements specify the total immediate bandwidth requirements of the DAMA connection/QoS. The aggregate bandwidth request is used by the base stations to reset their records to reflect the current total bandwidth requirements of an associated CPE. In one embodiment, four consecutive incremental bandwidth requests are transmitted. These incremental bandwidth requests are followed by an aggregate bandwidth request. The method of the self-correcting bandwidth request/grant protocol is now described in more detail with reference to FIGS. **15***a* and **15***b*.

FIG. **15***a* is a flow diagram showing one embodiment of the self-correcting bandwidth request/grant protocol **1500**. As shown in FIG. **15***a*, the method begins at a STEP **1510** whereat the CPE waits for the arrival of data for which bandwidth has not yet been requested. When data arrives, the method proceeds to a decision STEP **1520** to determine whether a timer (referred to in FIG. **15***a* as the "MT**10** timer") has expired. The MT**10** timer is preferably a dynamic bandwidth request timer as described above. The MT**10** timer is associated with the connection/QoS. The value of the MT**10** timer is used to determine the number of frames that the method waits between the transmission of aggregate bandwidth requests (in other words, the MT**10** timer determines the duration of the timer period between the periodic transmission of aggregate bandwidth requests). The MT**10** timer can have different values for each QoS used by the system. The MT**10** timer can be dynamically updated to different values depending upon current data flow and connection priority. Any bandwidth requests that are sent before the expiration of the MT**10** timer are, by definition, incremental bandwidth requests. In one embodiment, the number of frames between transmission of aggregate bandwidth requests is dependent upon the QoS. In one exemplary embodiment, the MT**10** timer duration equals 30 frames (for "real-time VBR" connections, or "rt-VBR"). In another exemplary embodiment, the MT**10** timer duration equals 100 frames (for "non-real-time VBR" connections, or "nrt-VBR").

Referring again to FIG. **15***a*, if the MT**10** timer is determined to have expired at STEP **1520**, the method proceeds to a STEP **1524** whereat the method builds an aggregate bandwidth request. If the MT**10** timer has not expired at STEP **1520**, the method proceeds to a decision STEP **1532** whereat the method determines whether a second timer, referred to in FIG. **15***a* as an "MT**5** timer", has expired. Similar to the MT**10** (aggregate bandwidth request timer) described above, the MT**5** timer also is preferably a dynamic bandwidth request timer. The MT**5** timer is also associated with the connection/QoS. The MT**5** timer is used to determine the number of frames that the method will wait for a response to a bandwidth request. As described above, bandwidth requests can be transmitted using either a "piggybacking" technique, or via individual polls. The MT**5** timer measures the duration (in frames) that the method waits for a response to one of these bandwidth requests. The MT**5** timer may use different values associated with different quality of services (QoS) used by the system (in other words, the duration of the MT**5** timer is QoS-dependent). The timer values of the MT**5** timer can be dynamically varied based upon current data flow and connection priority. In one embodiment, the method waits for duration of 10 frames (for rt-VBR). Alternatively, the method waits for duration of 20 frames (for nrt-VBR).

In one embodiment, the period counted by the MT**5** timer is less than the period counted by the MT**10** timer for the same QoS. If the MT**5** timer is determined not to have expired at the decision STEP **1532**, the method returns to the STEP **1520** to determine whether the MT**10** timer has expired. However, if the MT**5** timer is determined to have expired at the decision STEP **1532**, the method proceeds to a STEP **1522** whereat the method builds an incremental bandwidth request. The format used in building incremental and aggregate bandwidth requests is described in more detail below with reference to Table 1.

As shown in FIG. **15***a*, the method proceeds to a STEP **1526** whereat the bandwidth request built at the STEP **1522** (incremental bandwidth request) or at the STEP **1524** (aggregate bandwidth request) is transmitted to a selected base station. As indicated in the notation of the STEP **1526**, the mechanism used to transmit the bandwidth request can vary. Those skilled in the communications system design art shall recognize that various methods of transmitting bandwidth request messages can be used without departing from the scope or spirit of the present invention. The method proceeds to a control point (i.e., indicated in FIG. **15***a* by a STEP **1538**) where the method returns to the STEP **1510** to await the arrival of data for which bandwidth has not been requested.

In one embodiment, the units of requested bandwidth that are indicated in the bandwidth request/grant messages are the "natural" units of the system. In a variable length packet system, the "natural" units of the system are measured in bytes and bandwidth is therefore requested using bytes as the unit of bandwidth measurement. In a fixed length packet system, the "natural" units of the system are measured in packets and bandwidth is therefore requested using packets as the unit of bandwidth measurement.

FIG. **15***b* is a flow diagram showing an alternative embodiment of the self-correcting bandwidth request/grant protocol **1500'** of the present invention. As shown in FIG. **15***b*, the method begins at a STEP **1510'** whereat the CPE waits for the arrival of data for which bandwidth has not yet been requested. When data arrives, the method proceeds to a decision STEP **1520'** to determine whether a timer (referred to in FIG. **15***b* as the "MT**10** timer") has expired. The MT**10** timer in FIG. **15***b* is substantially similar to the MT**10** timer described above regarding FIG. **15***a* and thus is not described herein in more detail.

Referring again to FIG. **15***b*, if the MT**10** timer is determined to have expired at the STEP **1520'**, the method proceeds to a STEP **1524'** whereat the method builds an aggregate bandwidth request. If the MT**10** timer is determined not to have expired at the STEP **1520'**, the method proceeds to a STEP **1522'** whereat the method builds an incremental bandwidth request. The format used in building incremental and aggregate bandwidth requests is described in more detail below with reference to Table 1.

The method then proceeds to a STEP **1526'** whereat the bandwidth request built at the STEP **1522'** (incremental bandwidth request) or at the STEP **1524'** (aggregate bandwidth request) is transmitted to a selected base station. The mechanism used to transmit the bandwidth request can vary. Those skilled in the communications system design art shall recognize that various methods of transmitting bandwidth request messages can be used without departing from the scope or spirit of the present invention.

The method proceeds to a STEP **1528** whereat the method waits to receive an uplink sub-frame map indicating that a bandwidth grant has been provided. The method then proceeds to a decision STEP **1530** to determine whether the bandwidth grant is appropriate for the purpose of either transmitting data or waiting for an appropriate grant. If the grant was appropriate, the method proceeds to a STEP **1534**. Else, the method proceeds to a decision STEP **1532'**. At the decision STEP **1532'**, the method determines whether a second timer, referred to in FIG. **15***b* as an "MT**5** timer", has expired. The MT**5** timer in FIG. **15***b* is substantially similar to the MT**5** timer described above with reference to FIG. **15***a* and thus is not described in more detail herein.

In one embodiment, the method waits for duration of 10 frames (for rt-VBR). Alternatively, the method waits for duration of 20 frames (for nrt-VBR). In one embodiment, the period counted by the MT**5** timer is less than the period counted by the MT**10** timer for the same QoS. If the MT**5** timer is determined to have expired at the decision STEP **1532'**, the method returns to the decision STEP **1520'** to determine whether the MT**10** timer has expired. Else, the method returns to the STEP **1528** to await receipt of an uplink Map. If the bandwidth grant was determined to be appropriate at the decision STEP **1530**, the method proceeds to the STEP **1534**.

As shown in FIG. **15***b*, at the STEP **1534**, the CPE transmits data across the connection and proceeds to a decision STEP **1536** to determine whether the bandwidth grant fully satisfied the bandwidth request. If the grant fully satisfied the bandwidth request, the method proceeds to a control point (i.e., indicated in FIG. **15***b* by a STEP **1538'**) where the method returns to the STEP **1510'**. Else, the method returns to the STEP **1528** where the method waits to receive an uplink sub-frame map containing a bandwidth grant from the base station. As described above with reference to FIG. **15***a*, in one embodiment, the units of bandwidth that are indicated in the bandwidth request/grant messages are the "natural" units of the system. In a variable length packet system, the "natural" units of the system are measured in bytes and bandwidth is therefore requested using bytes as the unit of bandwidth measurement. In a fixed length packet system, the "natural" units of the system are measured in packets and bandwidth is therefore requested using packets as the unit of bandwidth measurement.

In one embodiment, bandwidth requests from the CPE to the base station are preferably made using the following

5-byte format (summarized below in Table 1). In one embodiment, the bandwidth request comprises four fields: a "Message Type ID" field, a "Connection ID" or "QoS" field, a "Request Type" field and an "Amount Requested" field. The Message Type ID field is one byte in length and contains a hexadecimal value of "0xB0". The Connection ID (or QoS) field is two bytes in length and contains the connection ID or QoS of the connection that is requested. The contents of the Connection ID or QoS field depends upon the BW Request Scope field of a MAC Control message. The Request Type field is one bit in length and contains information indicating whether a bandwidth request is an incremental or an aggregate bandwidth request. The Request Type field is set to a logical one to indicate that the bandwidth request is an incremental request. The Request Type field is set to a logical zero to indicate that the bandwidth request is an aggregate bandwidth request. In one embodiment, the Amount Requested field is fifteen bits in length and contains either (1) the amount of bandwidth (measured in the natural units of the system) that is requested per frame for CG connections or (2) the amount of incremental/aggregate bandwidth (measured in the natural units of the system) that is requested for DAMA connections. As described above with reference to FIG. **15***b*, for DAMA connections, the bandwidth units are measured in bytes in variable-length packet communication systems and in packets in fixed-length packet communication systems.

TABLE 1

Exemplary Bandwidth Request Format

Field	Size	Comments
Message Type ID	1 byte	Value = 0xB0
Connection ID or QoS	2 bytes	Depends upon BW Request Scope field of MAC Control message.
Request Type	1 bit	0 = total needs for connection/QoS 1 = incremental needs for connection/QoS
Amount Requested	15 bits	In units per frame for CG connections and total or incremental bytes pending for DAMA connections. Units = bytes for variable length packet system. Units = packets for fixed length packet system.
Total	5 bytes	

Abridged Bandwidth Request/Grant Protocol System

In embodiments that use an abridged bandwidth request/grant protocol system, the systems utilize padding packets to request a reduction in bandwidth allocation to a CPE. This embodiment is "abridged" because it does not require zero bandwidth request messages. The abridged bandwidth request/grant protocol system is now described in detail.

FIG. **16** shows a flowchart for the abridged bandwidth request/grant protocol system. The abridged bandwidth request/grant protocol method **1600** comprises a BS modem portion and a BS CPU portion. The BS modem portion comprises STEPs **1610**-**1630** and the BS CPU portion comprises STEP **1640**. In accordance with one embodiment, steps within the BS modem portion (e.g., STEP **1610**) are performed by the BS modem **1452**. Similarly, steps within the BS CPU portion are performed by the BS CPU **1454**. One skilled in the communication art shall recognize that the steps of the present inventive method may be performed in other configurations without departing from the scope or spirit of the present invention. For example, the present inventive method may perform STEP **1610** in the BS modem and STEPs **1620**-**1640** in the BS CPU.

Referring to FIG. **16**, the method begins at a STEP **1610** whereat the BS modem obtains a packet from a CPE. The method then proceeds to a decision STEP **1620** whereat the method determines whether the packet obtained in the STEP **1610** is a padding packet. If so, the method proceeds to a STEP **1630** whereat the BS modem alerts the BS CPU of the padding packet and its associated CPE, else the method returns to the STEP **1610**. In one embodiment of STEP **1630**, the alert sent from the BS modem to the BS CPU comprises a flag packet comprising data identifying an associated CPE that transmitted the padding packet. One skilled in the communication art shall recognize that other means for alerting the BS CPU may be utilized without departing from the scope or spirit of the present invention. For example, a shared memory means may be utilized.

In accordance with the STEP **1630** of the present method, a BS modem identifies a CPE that has transmitted a padding packet from the transmission portion of the uplink sub-frame. One skilled in the communication art shall recognize other methods of identifying CPE's without departing from the scope or spirit of the present invention. The BS modem can identify CPEs at every uplink sub-frame or uplink scheduling interval. After the STEP **1630** the method proceeds to a STEP **1640**.

As shown in FIG. **16**, the BS CPU portion comprises the STEP **1640** whereat the BS CPU notes that the associated CPU is active on the communication system and has too much bandwidth. In accordance with one embodiment, the BS CPU can reduce the associated CPE's bandwidth allocation. BS CPU resets all requested bandwidth for the associated CPE. One skilled in the communication art shall recognize that other methods of reducing CPE bandwidth allocation can be utilized with this embodiment without departing from its scope or spirit. This embodiment does not require zero bandwidth request messages.

A number of embodiments have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, the method and apparatus can be used in any type of communication system. Its use is not limited to a wireless communication system. One such example is use of the invention in a satellite communication system. In such a communication system, satellites replace the base stations described above. In addition, the CPEs are no longer at fixed distances from the satellites. Therefore, it will be more difficult to schedule DAMA services for the CPEs. Alternatively, the present invention can be used in a wired communication system. The only difference between the wired system and the wireless system described above is that the channel characteristics vary between the two. However, the bandwidth allocations do not change as between the two types of systems.

Although embodiments of the self-correcting bandwidth request/grant protocol described above uses a four-one incremental bandwidth request-aggregate bandwidth request pattern (i.e., four incremental bandwidth requests are transmitted, followed by one aggregate bandwidth request), alternative incremental/aggregate bandwidth request patterns can be used without departing from the scope or spirit of the present invention. For example, in one alternative embodiment, three incremental bandwidth requests are transmitted followed by one aggregate bandwidth request. Another alternative uses a 3-to-2 incremental-to-aggregate bandwidth request transmission pattern. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiment, but only by the scope of the appended claims.

The invention claimed is:

1. A method for requesting bandwidth in a wireless communication system, wherein the wireless communication system includes a plurality of subscriber units in communication with an associated base unit, the method comprising:

placing data received from various sources into a queue based on the quality of service (QoS) of the data;

setting an initial value of a timer; and

periodically, on expiration of the value in the timer, transmitting a bandwidth request indicating an amount of bandwidth required for transmitting the data from the queue.

2. A method as claimed in claim **1**, further comprising:

receiving a bandwidth grant in response to the bandwidth request; and

distributing the bandwidth grant between the queues at the subscriber unit using a fairness algorithm.

3. A method as claimed in claim **1**, wherein the value of the timer is dynamically changed.

4. A method as claimed in claim **1**, wherein the value of the timer is changed to reflect the current flow of data received on a plurality of connections established at the subscriber unit.

5. A method as claimed in claim **4**, wherein the subscriber unit transmits a padding packet to notify the base unit that the subscriber unit does not require any additional bandwidth allocation.

6. A method as claimed in claim **1**, wherein the value of the timer is changed based on the priority of the data in the associated queue.

7. A method as claimed in claim **1**, wherein the bandwidth request is inserted in an unused portion of the bandwidth allocated to the subscriber unit for data traffic.

8. A method as claimed in claim **1**, wherein the timer is reset upon receipt of the bandwidth grant.

9. A subscriber unit for a wireless communication system, wherein the wireless communication system includes a plurality of subscriber units in communication with an associated base unit, comprising:

a plurality of queues, each queue for grouping data based on the QoS; and

a media access (MAC) module configured to

set an initial value for a timer associated with a queue, and

periodically, on expiration of the value of the timer, transmit a bandwidth request indicating an amount of bandwidth required for transmitting the data from the queue.

10. A subscriber unit as claimed in claim **9**, wherein the value of the timer is dynamically changed.

11. A subscriber unit as claimed in claim **9**, wherein the value of the timer is changed to reflect the current flow of data received on a plurality of connections established at the subscriber unit.

12. A subscriber unit as claimed in claim **9**, wherein the value of the timer is changed based on the priority of the data in the associated queue.

13. A subscriber unit as claimed in claim **9**, wherein the bandwidth request is inserted in an unused portion of the bandwidth allocated to the subscriber unit for data traffic.

14. A method of allocating uplink (UL) bandwidth to connections established at a subscriber unit, the method comprising:

a) setting a value for a first timer;

b) determining a bandwidth amount requirement based on user traffic and transmitting a bandwidth request with the bandwidth amount requirement when the first timer expires;

c) receiving a bandwidth grant in response to the bandwidth request and distributing the bandwidth between the connections; and

d) transmitting the user traffic on receipt of the bandwidth grant.

15. A method as claimed in claim **14**, further comprising:

receiving a bandwidth grant in response to the bandwidth request; and

distributing the bandwidth grant between the queues at the subscriber unit using a fairness algorithm.

16. A method as claimed in claim **15**, wherein the bandwidth amount requirement comprises a total immediate bandwidth requirement for enabling transmission of the traffic received on the connections.

17. A method as claimed in claim **14**, further comprising:

e) determining an incremental bandwidth amount requirement, and

f) transmitting a bandwidth request with the incremental amount requirement when a second timer expires.

18. A method as claimed in claim **17**, wherein the incremental bandwidth requirement provides the change in bandwidth requirements since a last bandwidth request.

19. A method as claimed in claim **17**, wherein, when the bandwidth requirement for the connection is insufficient for accommodating additional traffic received on the connection, the incremental bandwidth requirement identifies a number of additional bandwidth units.

20. A method as claimed in claim **17**, wherein, when the bandwidth requirement for the connection exceeds the bandwidth necessary to accommodate the traffic received on the connection, the incremental bandwidth requirement includes padding packets to request a reduction in the bandwidth allocation.

21. A method as claimed in claim **17** further comprising:

setting the second timer to measure the time elapsed since the transmission of the last incremental bandwidth request; and

transmitting a new incremental bandwidth request upon expiration of the second timer if the UL bandwidth grant has not been received.

22. A method as claimed in claim **21**, wherein the first and the second timers are associated with a connection established at the subscriber unit.

23. A method as claimed in claim **22**, wherein the first and the second timers have an initial value selected based on a traffic parameter of the connection.

24. A method as claimed in claim **23**, wherein the traffic parameter is the quality of service of the traffic carried by the connection.

25. A method as claimed in claim **23**, wherein an initial value of the first and the second timers is dynamically selected based on current data flow and connection priority.

26. A subscriber unit for a wireless communication system, comprising:

a plurality of queues for buffering user traffic according to a traffic parameter, each queue having an associated logical state;

a media access control (MAC) element capable of

transmitting an uplink (UL) bandwidth request based on the logical state of the queues during a bandwidth request opportunity, and

allocating between the queues a bandwidth allocation received in response to the UL bandwidth request, based on the current state of the queues.

27. A subscriber unit as claimed in claim **26**, further comprising a first timer for determining the transmission time for an aggregate bandwidth request.

28. A subscriber unit as claimed in claim **26**, further comprising a second timer for determining the transmission time for an incremental bandwidth request.

29. A subscriber unit as claimed in claim **26**, wherein the MAC element sets the initial value of the first timer based on the logical state of the queue.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE

CERTIFICATE OF CORRECTION

PATENT NO.	: 8,457,145 B2	Page 1 of 1
APPLICATION NO.	: 13/487032	
DATED	: June 4, 2013	
INVENTOR(S)	: Zimmerman et al.	

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Specification

At column 33, Line 6, please delete “queues” and insert therefor, --connections--.

Signed and Sealed this
Twelfth Day of November, 2013

Teresa Stanek Rea
Deputy Director of the United States Patent and Trademark Office

US008537757B2

(12) **United States Patent**
Arviv et al.

(10) **Patent No.: US 8,537,757 B2**
(45) **Date of Patent: Sep. 17, 2013**

(54) **ADAPTIVE CALL ADMISSION CONTROL FOR USE IN A WIRELESS COMMUNICATION SYSTEM**

(75) Inventors: **Eli Arviv**, Modiin (IL); **Brian Spinar**, Poway, CA (US); **Kenneth L. Stanwood**, Cardiff by the Sea, CA (US); **David Gazelle**, La Jolla, CA (US); **Ofer Zimmerman**, Modiin (IL); **Penny Efraim**, Kfar Sava (IL); **Yair Bourlas**, San Diego, CA (US); **Sheldon L. Gilbert**, San Diego, CA (US)

(73) Assignee: **Wi-Lan, Inc.**, Ottawa (CA)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/487,013**

(22) Filed: **Jun. 1, 2012**
(Under 37 CFR 1.47)

(65) **Prior Publication Data**

US 2013/0064097 A1 Mar. 14, 2013

Related U.S. Application Data

(60) Continuation of application No. 12/414,363, filed on Mar. 30, 2009, now Pat. No. 8,213,359, which is a continuation of application No. 11/693,546, filed on Mar. 29, 2007, now Pat. No. 7,529,204, which is a continuation of application No. 11/350,464, filed on Feb. 8, 2006, now Pat. No. 7,289,467, which is a division of application No. 10/032,044, filed on Dec. 21, 2001, now Pat. No. 7,023,798.

(60) Provisional application No. 60/258,428, filed on Dec. 27, 2000.

(51) **Int. Cl.**
H04W 4/00 (2009.01)

(52) **U.S. Cl.**
USPC **370/328**; 370/465

(58) **Field of Classification Search**
None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,949,404 A 4/1976 Fletcher et al.
4,495,619 A 1/1985 Acampora
(Continued)

FOREIGN PATENT DOCUMENTS

EP 0507384 B1 10/1992
EP 0845916 B1 6/1998
(Continued)

OTHER PUBLICATIONS

"Asynchronous Transfer Mode (ATM) Technical Overview", 2nd Edition, Prentice Hall, Oct. 1995, Chapter 3, pp. 21-25.

(Continued)

Primary Examiner — Kevin C Harper
(74) *Attorney, Agent, or Firm* — Procopio, Cory, Hargreaves & Savitch LLP

(57) **ABSTRACT**

The invention relates to communication systems and to systems and methods for implementing adaptive call admission control (CAC) in such systems. Adaptive call admission control can determine what CPE to base station calls (connections) are allowed at any given time. CAC, coupled with precedence, can further determine what connections are suspended if less bandwidth is available than is currently committed. Multiple techniques are disclosed to select connections for suspension. These techniques include suspending enough connections through the affected CPE until there is enough bandwidth to meet the remaining commitment, randomly (or in a round robin fashion) choosing connection to suspend from the entire set of connection, and using precedence priority levels.

24 Claims, 6 Drawing Sheets

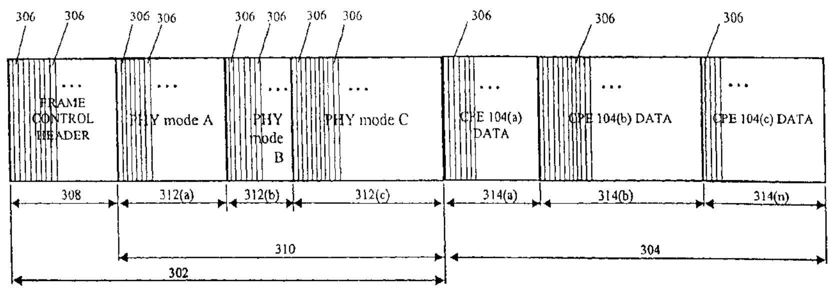

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,130,983	A	7/1992	Heffner, III	
5,297,144	A	3/1994	Gilbert et al.	
5,420,851	A	5/1995	Seshadri et al.	
5,444,698	A	8/1995	Kito	
5,471,645	A	11/1995	Felix	
5,511,082	A	4/1996	How et al.	
5,615,212	A	3/1997	Ruszczyk et al.	
5,638,371	A	6/1997	Raychaudhuri et al.	
5,638,374	A	6/1997	Heath	
5,675,573	A	10/1997	Karol et al.	
5,751,708	A	5/1998	Eng et al.	
5,768,254	A	6/1998	Papadopoulos et al.	
5,828,695	A	10/1998	Webb	
5,857,147	A	1/1999	Gardner et al.	
5,859,619	A	1/1999	Wu et al.	
5,890,055	A	3/1999	Chu et al.	
6,006,069	A	12/1999	Langston	
6,016,311	A	1/2000	Gilbert et al.	
6,016,313	A	1/2000	Foster, Jr. et al.	
6,038,455	A	3/2000	Gardner et al.	
6,094,421	A	7/2000	Scott	
6,112,080	A	8/2000	Anderson et al.	
6,141,317	A	10/2000	Marchok et al.	
6,519,462	B1	2/2003	Lu et al.	
6,636,500	B2	10/2003	Krishnamoorthy et al.	
7,023,798	B2	4/2006	Bourlas et al.	
7,050,407	B1	5/2006	Frazer et al.	
7,289,467	B2	10/2007	Bourlas et al.	
7,379,441	B2 *	5/2008	Stanwood et al.	370/337
2001/0028678	A1 *	10/2001	Kato et al.	375/222
2002/0061006	A1	5/2002	Soliman	
2002/0086707	A1 *	7/2002	Struhsaker et al.	455/561

FOREIGN PATENT DOCUMENTS

EP	0891060	A3	1/1999
WO	9222162	A1	12/1992
WO	9909779	A1	2/1999
WO	9953700	A1	10/1999
WO	0049824	A1	8/2000

OTHER PUBLICATIONS

C.E Shannon"A Mathematical Theory of Communication", Bell System Technical Journal, pp. 379-423 (Part 1), 623-656 (Part II), Jul. 1948.

H.C. Papadopoulos et al., Reduction of Mixed Co-channel Interference in Microcellular STDD Systems, Vehicular Technology Conference, 1995 IEEE 45th, vol. 2, pp. 759-763.

L.H. Charles Lee, "Convolutional Coding, Fundamentals and Applications", Artech House, Inc., 1997, p. 11-51.

Lin et al., "Error Control Coding, fundamentals and Applications", Prentice-Hall Computer Applications in Electrical Engineering Series., 1993, pp. 315-349.

Redl et al., "An Introduction to GSM", Artech House, Inc., 1995; pp. 84, 85 and 95.

Sampei et al., "Adaptive Modulation/TDMA Scheme for Personal Multi-Media Communication Systems", Proceedings of the Global Telecommunications Conference, U.S., New York, IEEE, Nov. 28, 1994, pp. 989-993.

Ue et al., "Symbol Rate and Modulation Level Controlled Adaptive Modulation/TDMA/TDD for Personal Communication Systems", Proceedings of the Vehicular Technology Conference, U.S. New York, IEEE, vol. Conf. 45, Jul. 25, 1995, pp. 306-310.

Ulm et al., "Data-Over-Cable Interface Specifications, Radio Frequency Interface Specification", Hewlett Packard Interim Specification, Doc. Control No. SP-RFII01-970321, published Mar. 21, 1997 by MCNS Holdings, L.P., Section 6, pp. 43-85.

Wolf et al., "On the Weight Distribution of Linear Block Codes Formed From Convolutional Codes", IEEE, IEEE Transactions on Communications, vol. 44:9, Sep. 1996.

U.S. District Court, Northern District of California, "Amended Complaint for Declaratory Judgment Demand for Jury Trial" filed Sep. 30, 2008 in Case No. 5:08-cv-4555.

J.M. Torrence, L. Hanzo, "Upper Bound Performance of Adaptive Modulation in a Slow Rayleigh Fading Channel," IEEE Electronics Letters, vol. 32, p. 718, Apr. 1996.

J. Pons and J. Dunlop, "Bit Error Rate Characterisation and Modelling for GSM", IEEE 1998, pp. 3722-3727.

P. Bender et al., "CDMA/HDR: A Bandwidth Efficient High Speed Wireless Data Service for Nomadic Users", Communications Magazine, IEEE, vol. 38, No. 7, Jul. 2000, pp. 70-77.

P. Jain, "On the Impact of Channel and Channel Quality Estimation on Adaptive Modulation", Dec. 2002.

J.B. Andersen et al., "Prediction of Future Fading Based on Past Measurements", Vehicular Technology Conference, VTC 99, vol. 1, pp. 151-155.

G. Narlikar et al., "Designing Multihop Wireless Backhand Networks with Delay Guarantees", Bell Labs, 2005.

* cited by examiner

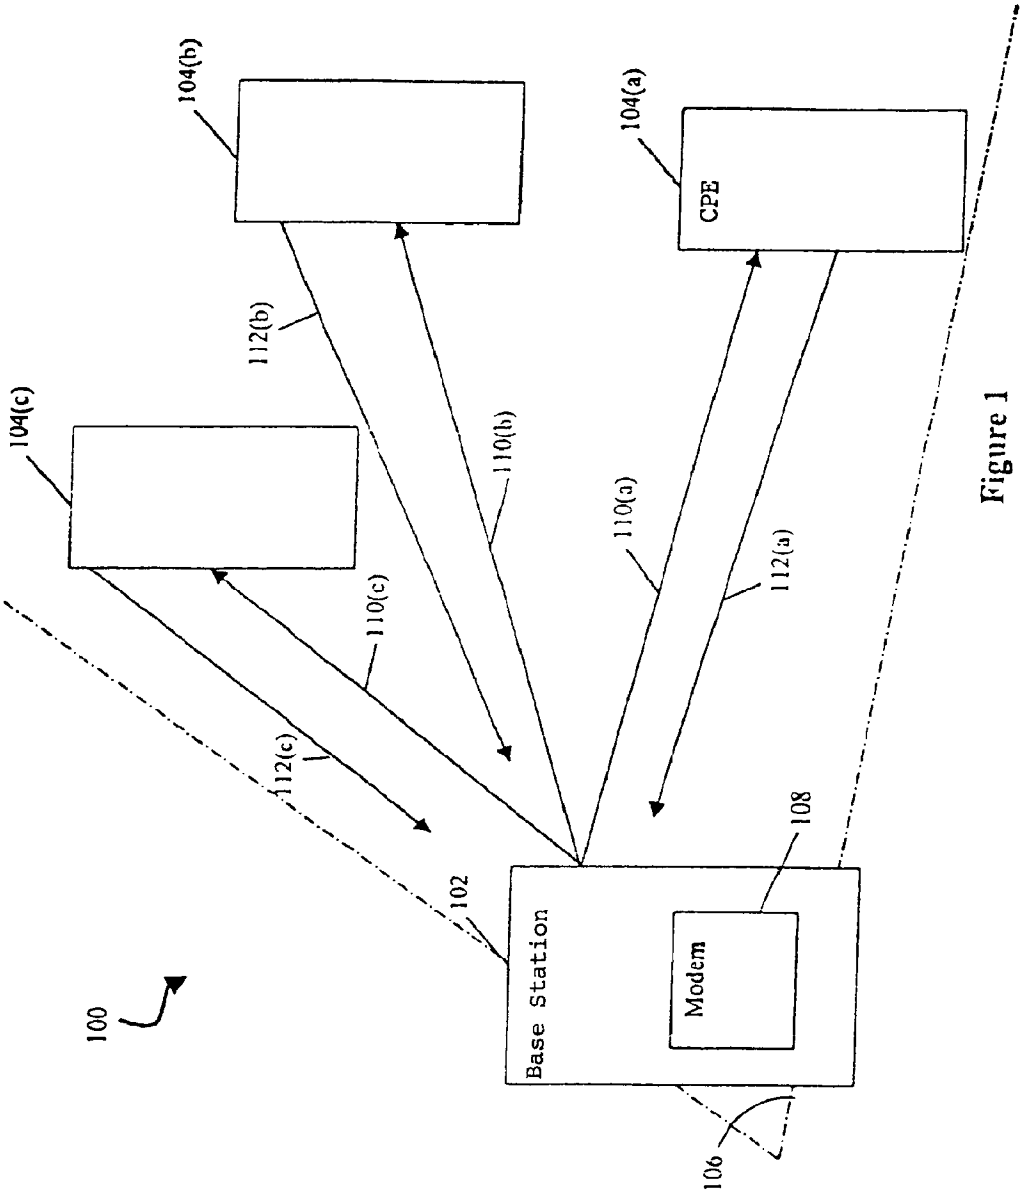
100
102
104(a)
CPE
104(b)
104(c)
106
108
110(a)
110(b)
110(c)
112(a)
112(b)
112(c)
Base Station
Modem

Figure 1

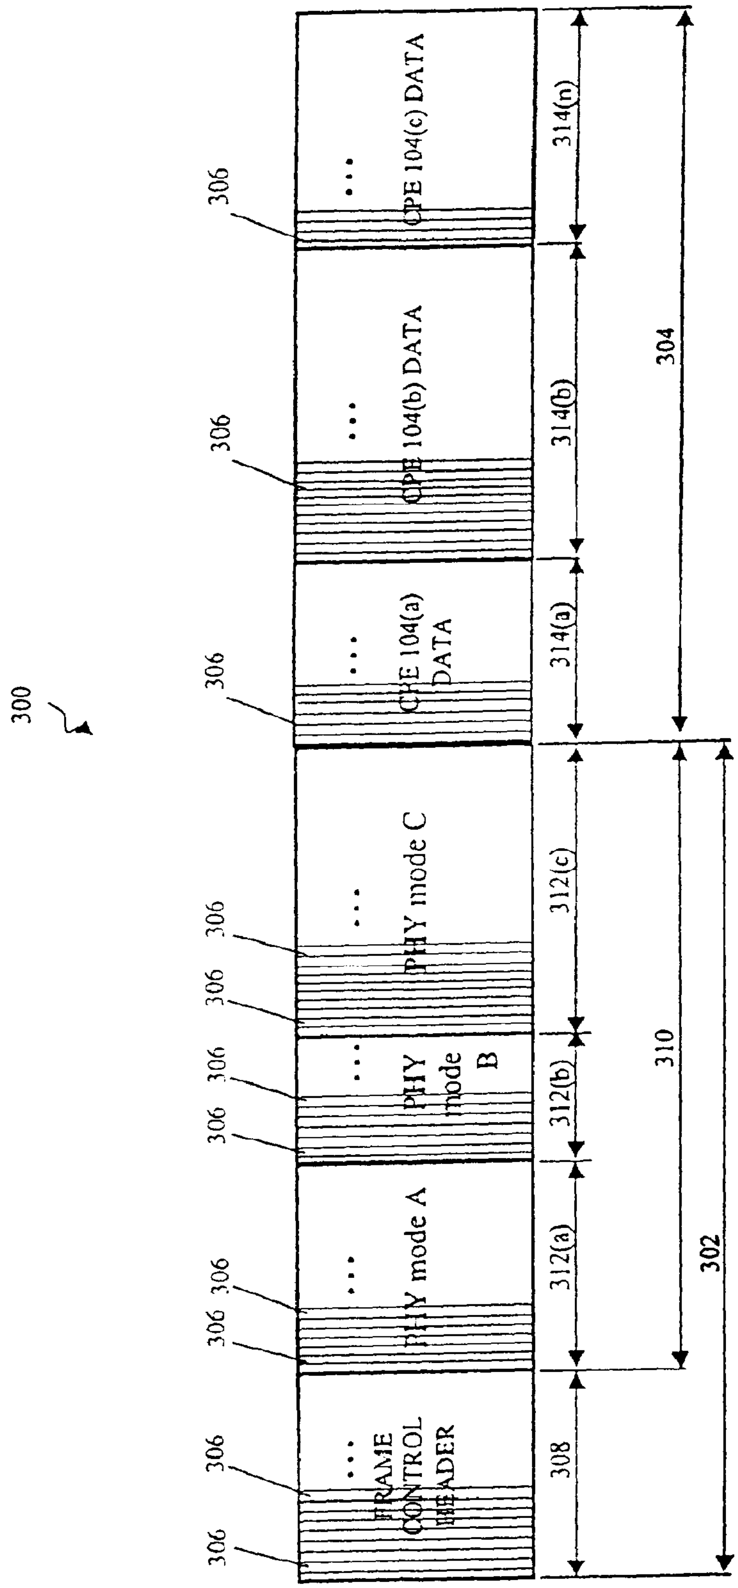
300
306
CPE 104(c) DATA
CPE 104(b) DATA
CPE 104(a) DATA
PHY mode C
PHY mode B
PHY mode A
FRAME CONTROL HEADER
314(n)
314(b)
314(a)
312(c)
312(b)
312(a)
308
304
310
302

FIGURE 2

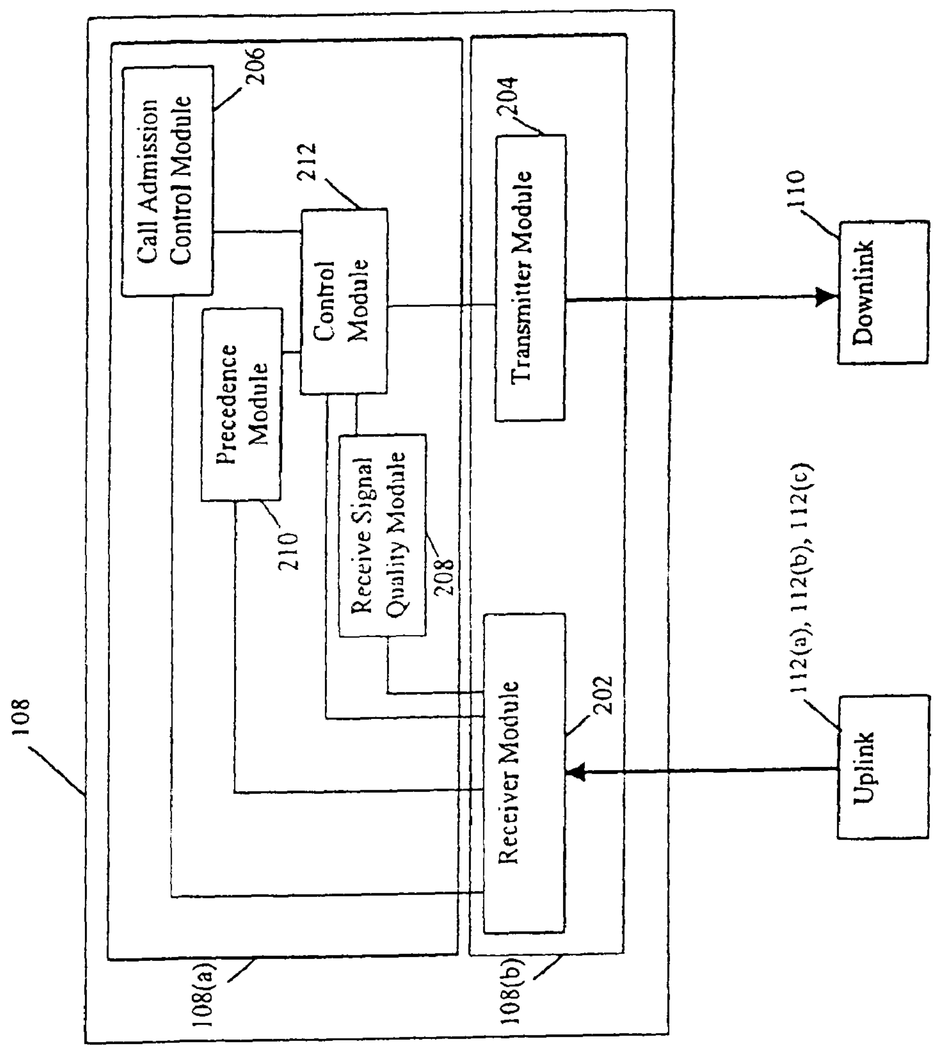
108
108(a)
108(b)
Call Admission Control Module
206
212
Control Module
Precedence Module
210
Receive Signal Quality Module
208
204
Transmitter Module
Receiver Module
202
110
Downlink
Uplink
112(a), 112(b), 112(c)

Figure 3

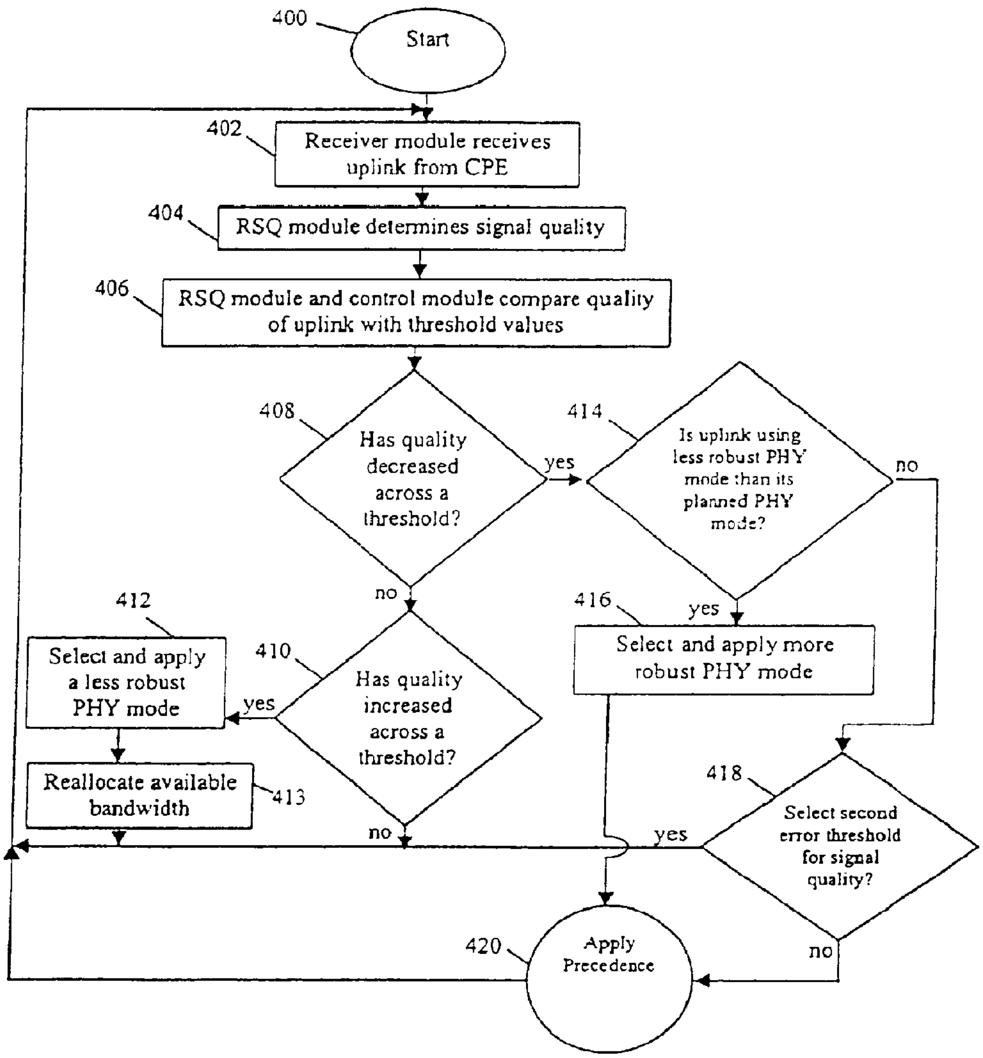
400
Start
402
Receiver module receives uplink from CPE
404
RSQ module determines signal quality
406
RSQ module and control module compare quality of uplink with threshold values
408
Has quality decreased across a threshold?
yes
414
Is uplink using less robust PHY mode than its planned PHY mode?
no
no
416
yes
Select and apply more robust PHY mode
412
Select and apply a less robust PHY mode
410
yes
Has quality increased across a threshold?
Reallocate available bandwidth
413
no
418
Select second error threshold for signal quality?
yes
no
420
Apply Precedence

FIGURE 4

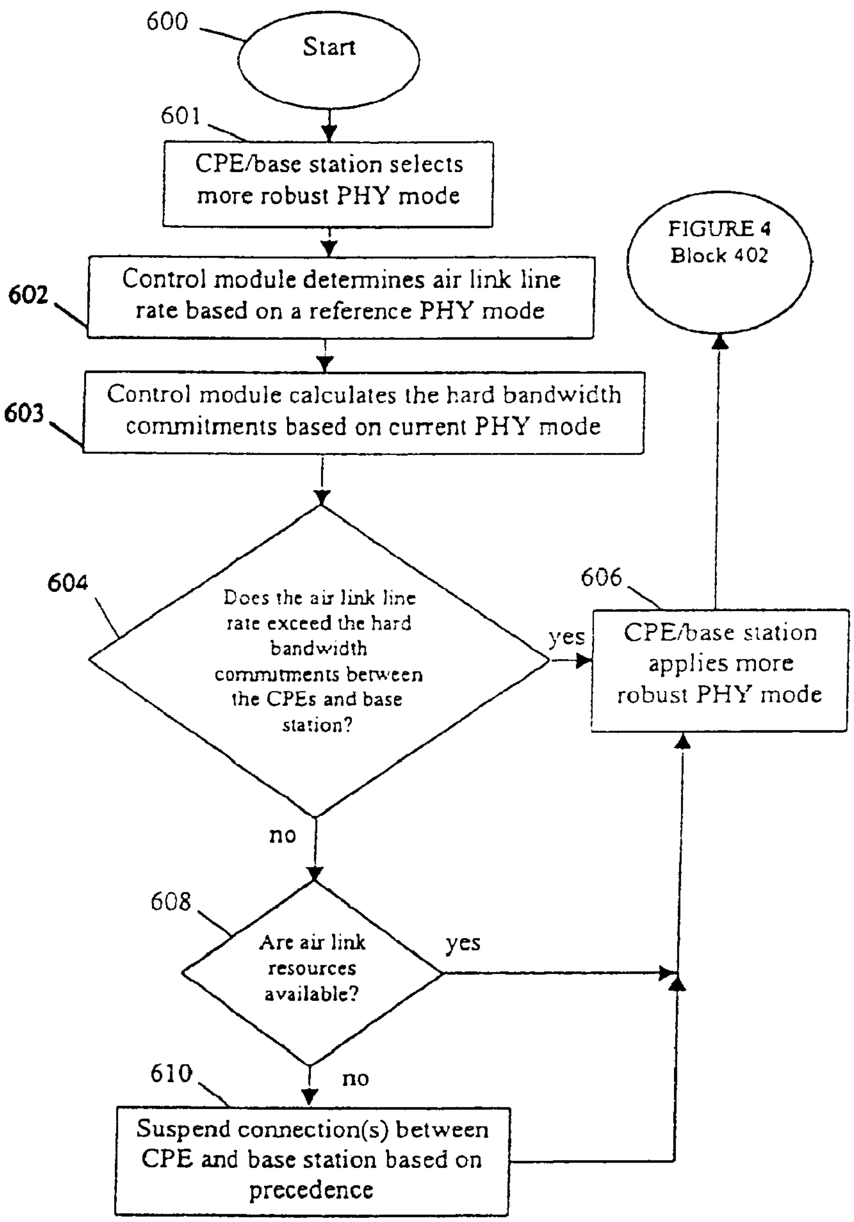
600
Start
601
CPE/base station selects more robust PHY mode
602
Control module determines air link line rate based on a reference PHY mode
603
Control module calculates the hard bandwidth commitments based on current PHY mode
604
Does the air link line rate exceed the hard bandwidth commitments between the CPEs and base station?
yes
no
606
CPE/base station applies more robust PHY mode
FIGURE 4 Block 402
608
Are air link resources available?
yes
no
610
Suspend connection(s) between CPE and base station based on precedence

FIGURE 5

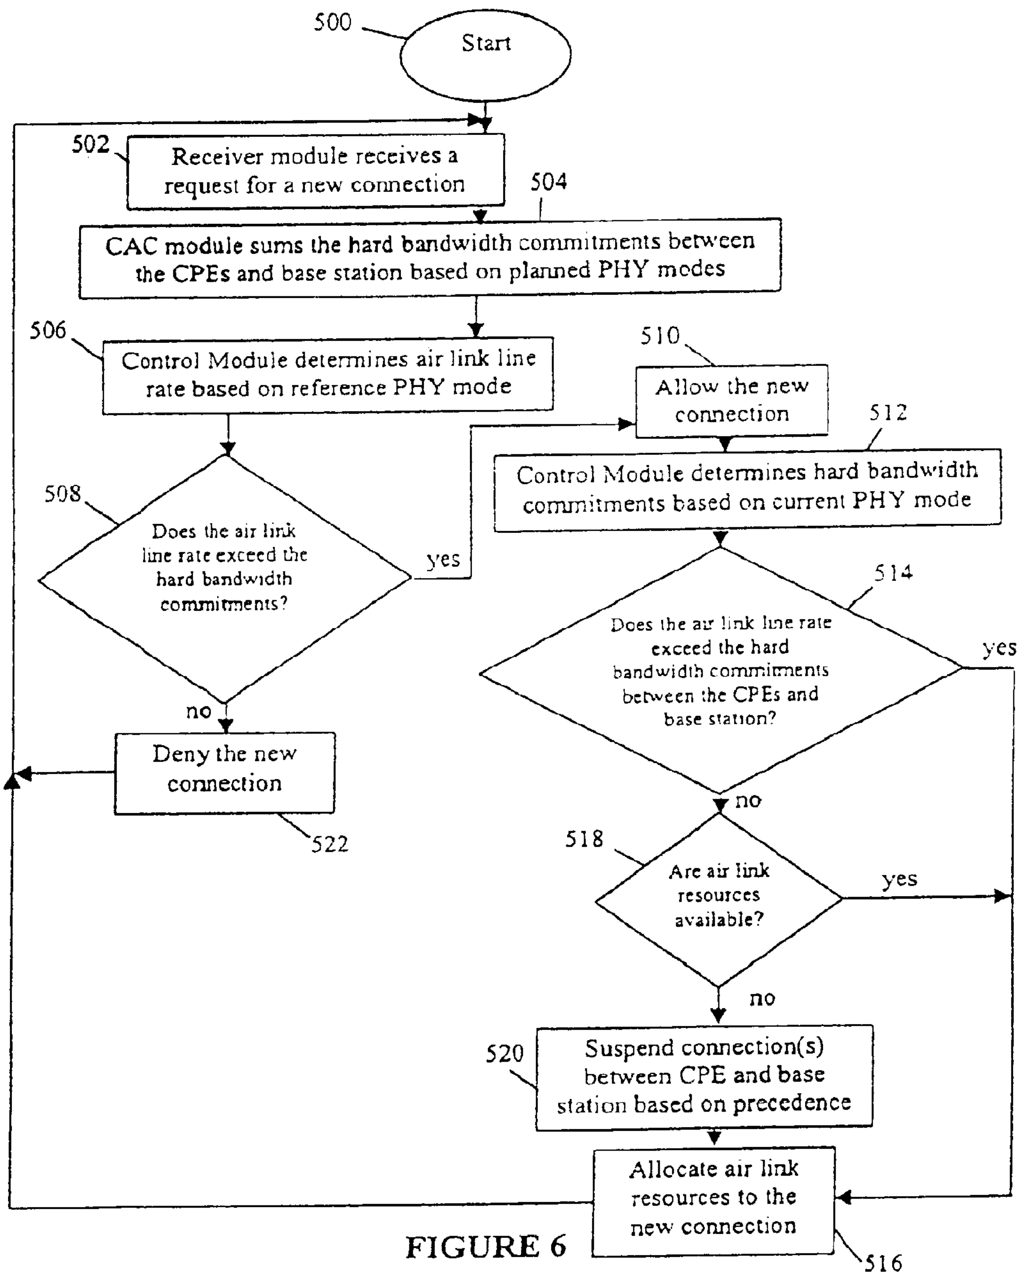

500
Start
502
Receiver module receives a request for a new connection
504
CAC module sums the hard bandwidth commitments between the CPEs and base station based on planned PHY modes
506
Control Module determines air link line rate based on reference PHY mode
508
Does the air link line rate exceed the hard bandwidth commitments?
yes
no
510
Allow the new connection
512
Control Module determines hard bandwidth commitments based on current PHY mode
514
Does the air link line rate exceed the hard bandwidth commitments between the CPEs and base station?
yes
no
Deny the new connection
522
518
Are air link resources available?
yes
no
520
Suspend connection(s) between CPE and base station based on precedence
Allocate air link resources to the new connection
516

FIGURE 6

ADAPTIVE CALL ADMISSION CONTROL FOR USE IN A WIRELESS COMMUNICATION SYSTEM

RELATED APPLICATIONS

This application is a continuation of U.S. application Ser. No. 12/414,363, filed Mar. 30, 2009, which is a continuation of U.S. application Ser. No. 11/693,546, filed Mar. 29, 2007, now U.S. Pat. No. 7,529,204, which is a continuation of U.S. application Ser. No. 11/350,464, filed Feb. 8, 2006, now U.S. Pat. No. 7,289,467, which is a divisional of U.S. application Ser. No. 10/032,044, filed Dec. 21, 2001, now U.S. Pat. No. 7,023,798, which claims priority to U.S. provisional patent application Ser. No. 60/258,428, filed Dec. 27, 2000, all entitled ADAPTIVE CALL ADMISSION CONTROL FOR USE IN A COMMUNICATION SYSTEM, all of which are incorporated herewith in their entirety by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to communication systems and to a system and method for implementing adaptive call admission control in such systems.

2. Description of the Related Art

A wireless communication system facilitates two-way communication between a plurality of subscriber units (fixed and portable) and a fixed network infrastructure. Exemplary communication systems include mobile cellular telephone systems, personal communication systems ("PCS"), and cordless telephones. An objective of these wireless communication systems is to provide communication channels on demand between the subscriber units and their respective base stations in order to connect a subscriber unit end user with the fixed network infrastructure (usually a wire-line system). In the wireless systems having multiple access schemes, a time "frame" is used as the basic information transmission unit. Each frame is sub-divided into a plurality of time slots. Subscriber units typically communicate with their respective base station using a "duplexing" scheme thus allowing for the exchange of information in both directions of the connection.

Transmissions from the base station to the subscriber units are commonly referred to as "downlink" transmissions. Transmissions from the subscriber units to the base station are commonly referred to as "uplink" transmissions. Depending upon the design criteria of a given system, wireless communication systems have typically used either time division duplexing ("TDD") or frequency division duplexing ("FDD") methods to facilitate the exchange of information between the base station and the subscriber units.

SUMMARY OF THE INVENTION

The systems and methods have several features, no single one of which is solely responsible for its desirable attributes. Without limiting the scope as expressed by the claims which follow, its more prominent features will now be discussed briefly. After considering this discussion, and particularly after reading the section entitled "Detailed Description" one will understand how the features of the system and methods provide several advantages over traditional communication systems.

One aspect is a communication system that is configured to control the admission of new connections and the suspension of existing connections between a base station and customer premise equipments (CPEs), wherein the base station and the CPEs are each configured to increase or decrease the robustness of their transmission modulation technique by adapting their PHY mode. The system comprises a first CPE having a first modem configured to modulate data in a communication link using a first current PHY mode and a first planned PHY mode, a second CPE having a second modem configured to modulate data in a communication link using a second current PHY mode and a second planned PHY mode, and a base station having a third modem configured to transmit and receive data to and from the first and second CPEs. The system further comprises a call admission control (CAC) module configured to determine whether to allow a new connection between the first CPE and the base station or between the second CPE and the base station based on a comparison of a total air link line rate between the first and second CPEs and the base station, wherein the total air link line rate is based on a reference PHY mode, with a bandwidth commitment value between the base station and the first and second CPEs, wherein the bandwidth commitment is based on the first and second planned PHY modes.

Another aspect is a method for controlling the admission of connections in a wireless communication system between a base station and associated CPEs, including a requesting CPE. The method comprises receiving a request for a new connection from a requesting CPE, summing the hard bandwidth commitments between a base station and associated CPEs, including the new connection and existing connections, based on a planned PHY mode for each connection, and determining an air link line rate between the base station and the associated CPEs based on a reference PHY mode. The method further includes if the air link line rate exceeds the hard bandwidth commitments, accepting the new connection and determining a second hard bandwidth commitments for the existing connections between the base station and the associated CPEs based on a current PHY mode for each connection, else denying the new connection. The method still further includes if the air link line rate exceeds the second hard bandwidth commitments, allocating air link resources to the new connection, else determining whether additional air link resources are available, and if additional air link resources are available, allocating the air link resources to the new connection, else suspending at least one of the existing connections between the base station and the associated CPEs.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. **1** is a simplified block diagram of a wireless communication system including a base station and one or more CPEs.

FIG. **2** is an illustration of the structure of a Time Division Duplex ("TDD") frame.

FIG. **3** is a block diagram of a modem.

FIG. **4** is a flowchart illustrating the process of adaptively adjusting a PHY mode for an uplink connection between the base station and a CPE.

FIG. **5** is a flowchart illustrating the process of precedence being applied to existing connections between the CPE and the base station.

FIG. **6** is a flowchart illustrating the process of call admission control to a new connection between a CPE and the base station.

DETAILED DESCRIPTION

The following detailed description is directed to certain specific embodiments of the invention. However, the invention can be embodied in a multitude of different systems and methods. In this description, reference is made to the drawings wherein like parts are designated with like numerals throughout.

In connection with the following description many of the components of the various systems, some of which are referred to as a "module," can be implemented as software, firmware or a hardware component configured to perform one or more functions or processes. Hardware components can include, for example, a Field Programmable Gate Array (FPGA) or Application-Specific Integrated Circuit (ASIC). Such components or modules may reside on the addressable storage medium and configured to execute on one or more processors. Thus, a module may include, by way of example, components, such as software components, object-oriented software components, class components and task components, processes, functions, attributes, procedures, subroutines, segments of program code, drivers, firmware, microcode, circuitry, data, databases, data structures, tables, arrays, and variables. The functionality provided for in the components and modules may be combined into fewer components and modules or further separated into additional components and modules. Additionally, the components and modules may advantageously be implemented to execute on one or more computers.

FIG. **1** is a block diagram of an exemplary wireless communication system **100**. Alternatively, the methods and systems herein disclosed can be implemented in wired communication systems (not shown). One exemplary broadband wireless communication system is described in U.S. Pat. No. 6,016,311, by Gilbert et al., issued Jan. 18, 2000, entitled "Adaptive Time Division Duplexing Method and Apparatus for Dynamic Bandwidth Allocation within a Wireless Communication System," hereby incorporated by reference. The system **100** includes a base station **102** and at least one customer premise equipment. The system depicted in FIG. **1** shows three CPEs **104**(*a*)-(*c*). More or fewer CPEs can be used. The CPEs and the base station receive and transmit data along wireless communication links **110**(*a*)-(*c*), **112**(*a*)-(*c*).

FIG. **1** does not show buildings or other physical obstructions (such as trees or hills, for example), that may cause channel interference between data from communication links **110**, **112**. The CPEs **104** and the base station **102** communicate by transmitting their data as radio frequency signals. The term channel refers to a band or range of radio frequencies of sufficient width for communication. For example, the range of frequencies from 26.500 GHz to 26.525 GHz would provide a 25 MHz wide channel. Although the following discussion uses the example of a system that transmits information within the Local Multi-Point Distribution Services (LMDS) band at frequencies of approximately 28 GHz, the invention is not so limited. Information can be transmitted at various frequencies and ranges including, for example, 10 GHz to 66 GHz using Quadrature Amplitude Modulation (QAM) symbols. The systems and methods described herein can also be used in a Multichannel Multi-point Distribution Service (MMDS) which operates below 10 GHz. In the MMDS, Orthogonal Frequency Division Multiplexing (OFDM) symbols may be transmitted between the base station and CPEs as an alternative to single carrier QAM modulation. In such a system, the methods and systems are applied to one or more of the OFDM subchannels.

Referring again to FIG. **1**, the communication links **110**(*a*), **110**(*b*), **110**(*c*) are referred to as downlinks (i.e., from the base station **102** to the CPE's **104**) and can operate on a point (base station)-to-multi-point (CPE's) basis. Transmissions to and from the base station **102** can be directional in nature, and thus limited to a particular transmission sector **106** of the base station **102**. Within a given sector **106**, CPEs **104**(*a*), **104**(*b*), **104**(*c*) receive the same transmission along their respective downlinks **110**(*a*), **110**(*b*), **110**(*c*). To distinguish between data intended for a specific CPE, the CPEs can monitor control information in their respective downlink **110**(*a*), **110**(*b*), **110**(*c*) and typically retain only the data intended for them. In communication systems that have multiple sectors, the base station **102** can include a sectored active antenna array (not shown) which is capable of simultaneously transmitting to multiple sectors. In one embodiment of the system **100**, the active antenna array transmits to four independent sectors.

The communication links **112**(*a*), **112**(*b*), **112**(*c*) are referred to as an uplink (i.e., from the CPEs **104** to the base station **102**) and can operate on a point-to-point basis. Thus, in FIG. **1**, each CPE **104**(*a*), **104**(*b*), **104**(*c*) originates its own uplink **112**(*a*), **112**(*b*), **112**(*c*). Communication with the base station **102** is bi-directional and can be multiplexed on the basis of Time Division Duplexing (TDD). For a TDD transmission from, for example, CPE **104**(*a*), CPE **104**(*a*) would send its data along communication link **112**(*a*) to the base station **102** during a preassigned time slot in a transmission frame. The specific frame structures of the uplink and downlink will be discussed further below.

Alternatively, the system can employ Frequency Division Duplexing (FDD). In such an FDD system, duplexing of transmissions between the base station and the CPEs is performed in the frequency domain. Different sets of frequencies are allocated for uplink and downlink transmissions. The systems and methods described herein can be used in such an FDD system.

Each CPE **104** is further coupled to a plurality of end users that may include both residential and business customers. Each customer can have one or more connections between the CPE and the base station. Consequently, each end user connection can have different and varying usage and bandwidth requirements. Each CPE **104**(*a*)-(*c*) may service several hundred or more end users, but at least one end user will be assigned to transmit and receive data via at least one connection through each CPE **104**.

The data transmitted along the communication links **110**, **112** is in analog form, and thus a modem **108** is used to modulate the digital data prior to transmission. FIG. **1** illustrates the modem **108** being located at the base station **102**, however, a similar or identical modem **108** may be used at the other end of the downlinks **110**(*a*), **110**(*b*), **110**(*c*) to demodulate the received analog data. Thus, the modems **108** in the base station and each CPE are used for uplinking data from the CPEs to the base station and for downlinking data from the base station to the CPEs.

The base station and CPEs can use adaptive modulation and forward error correction (FEC) schemes to communicate. Adaptive modulation, or adaptable modulation density, includes varying the bit per symbol rate modulation scheme, or modulation robustness, of downlinks and uplinks transmitted between CPEs and the base station. Examples of such modulation schemes include quadrature amplitude modulation-4 (QAM-4), QAM-16, QAM-64, and QAM-256. If QAM-4 is used, each resulting symbol represents two bits. If QAM-64 is used, each resulting symbol represents six bits. Adaptive FEC includes varying the amount of error correction data that is transmitted in the downlink and/or uplink. Channel characteristics, for example the modulation and FEC, for the downlink and/or uplink can be varied independently. For ease of explanation, the phrase "PHY mode" is used to indicate characteristics of a communication channel or link, including for example, modulation scheme and/or an FEC.

The PHY mode(s) planned for use in the sector **106** is normally determined as a function of the geographical relationship between the base station **102** and the CPEs, the rain region, and the implementation or modem complexity of the CPEs. Examples of rain regions include rain regions A-Q. Recommendations for modeling the rain region's effect on signal propagation can be found in Rec. ITU-R PN.837.1. Thus, a planned PHY mode may be different for the CPEs depending on the capabilities and transmission quality of each CPE **104** and base station **102** pair. For ease of explanation, the phrase "planned PHY mode" is used to indicate the planned PHY mode for a CPE **104** and base station **102** pair as described above.

Better environmental conditions, e.g., less distance, between some CPEs (such as CPE **104**(*c*) for example) and the base station **102** may permit the use of a less robust PHY mode by such CPEs as compared to a PHY mode used by CPEs located farther from the base station. For example, if CPE **104**(*c*) is capable of receiving QAM-64 data coupled with achieving adequate transmission quality between CPE **104**(*c*) and the base station **102**, all data transmitted between the CPE and the base station can be modulated using QAM-64. In the same system CPEs **104**(*a*), **104**(*b*), which, for example, are only capable of receiving QAM-4 data, will only transmit and receive QAM-4 data. By using different or variable PHY modes for different CPEs associated with a single base station, the communication system **100** as a whole increases its bandwidth utilization.

The transmission quality between the base station **102** and a CPE **104** may not only vary between each CPE and base station pair as described above, but may also vary over time, or between the uplink and downlink transmissions of a single pair (i.e. asymmetrical transmissions). For example, in FIG. **1**, the transmission quality may significantly decrease during a rain or snow storm. When the link quality is decreased, there is an increased chance that transmitted data along communication links **110**(*a*), **110**(*b*), **110**(*c*), **112**(*a*), **112**(*b*), **112**(*c*) may be unrecognizable or lost to the receiving base station or CPE. To accommodate these time variations in link quality, the communication system **100** can dynamically adjust or "adapt" the PHY mode for each base station **102** and CPE **104**. In such an adaptive system, the bandwidth utilization of the communication system **100** further increases.

FIG. **2** represents a time division duplexing ("TDD") frame and multi-frame structure for use in communication system **100**. Frame **300** includes a downlink subframe **302** and an uplink subframe **304**. The downlink subframe **302** is used by the base station **102** to transmit information to the CPEs **104**(*a*)-(*c*). In any given downlink subframe **302**, all, some, or none of the transmitted information is intended for a specific CPE **104**. The base station **102** may transmit the downlink subframe **302** prior to receiving the uplink subframe **304**. The uplink subframe **304** is used by the CPEs **104**(*a*)-(*c*) to transmit information to the base station **102**.

Subframes **302**, **304** are subdivided into a plurality of physical layer slots (PS) **306**. Each PS **306** correlates with a duration of time. In FIG. **2**, each subframe **302**, **304** can be one-half millisecond in duration and include 400 PS for a total of 800 PS per frame **300**. Alternatively, subframes having longer or shorter durations and with more or fewer PSs can be used. Additionally, the size of the subframes can be asymmetrical and can be varied over time.

Each downlink subframe **302** can include a frame control header **308** and downlink data **310**. The frame control header **308** includes information for the CPEs to synchronize with the base station **102**. The frame control header **308** can include control information indicating where a PHY mode change occurs in the downlink. The frame control header **308** can also include a map of a subsequent uplink subframe **304**. This map allocates the PSs **306** in the uplink subframe **304** between the different CPEs. The frame control header **308** can further include a map of attributes of the downlink data **310**. For example, attributes may include, but are not limited to, the locations of the PSs **306** in the subframe **302** that are intended for each individual CPE.

The downlink data **310** is transmitted using a pre-defined PHY mode or a sequence of PHY modes with three PHY modes A, B, and C depicted in FIG. **2** as an example. Individual or groups of PSs **306** in the downlink subframe **302** are assigned to data intended for specific CPEs **104**. For example, the base station **102** could assign PSs in one, some, or all of the PHY modes A, B, and C for transmitting data to CPE **104**(*a*). In FIG. **2**, the data is divided into three PHY modes, where PHY mode A (**312**(*a*)) is the most robust modulation (i.e. least prone to transmission errors caused by signal interference) and while PHY mode C (**312**(*c*)) is the least robust (i.e. most prone to transmission errors caused by signal interference). In between these PHY modes is PHY mode B (**312**(*b*)). Additional PHY modes can also be used.

Still referring to FIG. **2**, the uplink subframe **304** comprises uplink data **314**(*a*)-(*n*). The uplink subframe **304** is used by the CPEs **104**(*a*)-(*c*) to transmit information to the base station **102**. The subframe **304** is subdivided into a plurality of PSs **306**. Each CPE **104**(*a*)-(*c*) transmits its information during its allocated PS **306** or range of PSs **306**. The PSs **306** allocated for each CPE can be grouped into a contiguous block of a plurality of data blocks **314**(*a*)-(*n*). The CPEs use data blocks **314**(*a*)-(*n*) to transmit the uplink subframe **304**. The range of PSs **306** allocated to each block in the plurality of data blocks **314**(*a*)-(*n*) can be selected by the base station **102**. The data transmitted in each data block **314**(*a*)-(*n*) is modulated by the transmitting CPE. For example, CPE **104**(*a*) modulates and transmits uplink data block **314**(*a*). The same or different PHY modes can be used for each data block **314**(*a*)-(*n*). The data blocks **314**(*a*)-(*n*) can also be grouped by PHY mode.

During its data block, the CPE transmits with a PHY mode that is selected based on measured channel parameters from its prior transmission(s). Similarly, the base station can select a downlink PHY mode for a communication link based on measured channel parameters from its prior transmission(s). The process for selecting a PHY mode will be explained in more detail below. The measured channel parameters can be included in the uplink subframe **304** for transmission by the CPEs to the base station or can be included in the downlink subframe **302** for transmission by the base station to the CPE. Once received, the base station or CPE can utilize the channel parameters to determine if the PHY mode of the downlink subframe **302** or the uplink subframe **304** should be changed.

Each CPE **104** can receive all downlink transmissions that are modulated using its current PHY mode or are modulated using a more robust PHY mode than its current PHY mode. The frame control header **308** is typically modulated using the most robust PHY mode to ensure that all CPEs **104**(*a*)-(*c*) may receive it. Because each CPE receives the frame control header, each CPE **104** is initially synchronized with the downlink subframe **302** at the beginning of the frame **300**. The downlink subframe can be sorted by robustness, which allows each CPE to maintain synchronization during the subsequent portion of the downlink that could include data for that CPE.

FIG. **3** is a block diagram of a modem **108** which can be used to modulate/demodulate data in the wireless communication system **100** described above. The modem **108** is used to control the number and quality of existing and new connections between the CPEs and base station. Modems **108** are used by the base station **102** and CPEs **104** to modulate and demodulate data. For ease of description, the modem **108** will now be described with reference to the base station **102**.

The modem **108** can include a control section **108**(*a*) and a modem section **108**(*b*). The modem section **108**(*b*) includes a receiver module **202** and a transmitter module **204**. The control section **108**(*a*) includes a call admission control (CAC) module **206**, a Receive Signal Quality (RSQ) module **208**, a precedence module **210**, and a control module **212**. Alternatively, the functionality provided for by the control section **108**(*a*) can be separate from the modem **108**. Further, the control section **108**(*a*) components and modules may be combined into fewer components and modules or further separated into additional components and modules within the base station **102** and/or CPE **104**.

At a base station **102**, the transmitter module **204** converts digital data to an appropriately modulated analog signal communicated as a downlink **110**, using for example, QAM modulation and FEC. The analog signal may also be up converted to a carrier frequency prior to transmission. The receiver module **202** at the base station **102** demodulates an uplink **112**(*a*), **112**(*b*), **112**(*c*) and converts it back to digital form. When configured as a CPE **104**(*a*), the transmitter module **204** converts digital data to an appropriately modulated analog signal communicated as an uplink **112**, using for example, QAM modulation and FEC. The analog signal may also be up converted to a carrier frequency prior to transmission. The receiver module **202** at the CPE **104** demodulates a downlink **110** and converts it back to digital form.

The wireless communication system **100** can provide "bandwidth-on-demand" to the CPEs. Thus, the uplink can include bandwidth requests for new and existing connections from end users. The CPEs request bandwidth allocations from their respective base station **102** based upon the type and quality of service requested by the end users served by the CPE. A CPE or base station can continue an existing connection or allow a new connection depending on, for example, a user's defined quality of service, bandwidth needs, and transmission quality. Thus, each end user potentially uses a different broadband service having different bandwidth and latency requirements. Moreover, each user can select a portion(s) of their bandwidth to have variable priority levels, or precedence.

To this end, the type and quality of service available to the end users are variable and selectable. The amount of bandwidth dedicated to a given service can be determined by the information rate and the quality of service required by that service (and also taking into account bandwidth availability and other system parameters as will be described below). For example, T1-type continuous data services typically require a great deal of bandwidth having well controlled delivery latency. Until terminated, these services require constant bandwidth allocation for each downlink subframe **302** and uplink subframe **304** in a frame **300** (see FIG. **2**). In contrast, certain types of data services such as Internet Protocol data services ("TCP/IP") are bursty, often idle (which at any one instant may require zero bandwidth), and are relatively insensitive to delay variations when active.

Referring again to FIG. **3**, the Receive Signal Quality (RSQ) module **208** interfaces with the receiver module **202** and the control module **212**. The RSQ module **208** is configured to monitor signal quality of the received uplink signal. In a communication system that adapts PHY modes, the selection of a PHY mode can be based on channel parameters monitored/measured by the RSQ module **208**. These channel parameters can include the signal to noise ratio (SNR) of the modulated data at the receiver module **202** at the base station **102**. A bit error rate (BER), at the base station **102** or CPE **104**, can also be used in selecting the PHY mode. For example, when the received signal drops below a threshold value for a SNR, a more robust PHY mode can be selected by the modem **108** for the connection. Signal quality can be measured over a period of time by the RSQ module **208**, and, in response to changes in the signal quality, the control module **212** determines if the PHY mode for the transmitting CPE should be changed. The control module **212** at the base station **102** interfaces with the transmitter module **204** to control the PHY mode for the modem **108**. Further, the control module **212**, via the transmitter module **204**, can alert the transmitting CPE to change its PHY mode. Measuring signal quality over time helps avoid cyclic changes in the PHY mode due to transient changes in the communication link's quality.

The RSQ module at the CPE can measure signal quality for a signal that is transmitted by the base station **102** and received by the CPE. The CPE can alert the base station to change the base station's transmitting PHY mode. In one embodiment, only the modem **108** at the base station **102** includes the control module **212**. In this embodiment, each CPE measures its own signal quality and transmits its value within its uplink **112** to the base station **102**. The control module **212** is then able to monitor the signal quality of the signal received by the CPEs to determine if the downlink **110** PHY modes should be changed.

The call admission control (CAC) module **206** determines what CPE to base station connections are allowed at any given time. For example, the receiver module **202** can receive a request for a new connection between the CPE and base station in the uplink **112**. The CAC module determines whether to grant that request. This determination can be based on intrinsic factors relating to the new connection as well as communication system level factors. Examples of intrinsic factors are a quality of service and a type of service requested by the end user for the new connection. The extrinsic factors are external to the new connection. The extrinsic factors can include the type and quality of service for the existing connections, whether available bandwidth is allocated to the requesting CPE, the available bandwidth in the communication link, and the portion of the frame that is allocated for the uplink and downlink. An example of a type and quality of service that can be evaluated by the CAC module **206** are hard bandwidth commitments.

The CAC module **206** can be configured to determine whether there will be enough bandwidth to support all of the connections between the CPEs **104** and the base station **102**. For example, the CAC module **206** can determine whether there will be enough bandwidth for hard bandwidth commitments between the base station and CPEs. These hard bandwidth commitments can include, for example, constant bit rate (CBR) connections, the minimum cell rate (MCR) portion of a guaranteed frame rate (GFR) connections, and some function of sustainable cell rate (SCR) for variable bit rate (VBR) and variable bit rate real-time (VBR-rt) connections. Alternatively, hard bandwidth commitments could be the bandwidth measured, rather than calculated, that is necessary to provide the quality of service (QoS) desired for the connection. For ease of explanation, the following description uses hard bandwidth commitments as an exemplary type of connection. However, the systems and methods disclosed herein are not so limited and can be applied to any type of connection. Further, the systems and methods can be applied to one or more types of connections.

The CAC module **206** determines whether there is enough bandwidth to allow the new connection. This can be determined by summing the hard bandwidth commitments for each connection on each CPE **104**(*a*), **104**(*b*), **104**(*c*) (see FIG. **1**). Thus, each CPE will have a hard bandwidth commitment for its existing connections. All of the hard bandwidth commitments from the CPEs can then be summed to get the total hard bandwidth commitments for all of the existing connections through base station **102**. The control module **212** can perform these calculations. The CAC module **206** compares the total hard bandwidth commitments to an air link line rate. The air link line rate is the amount of bandwidth available between the CPEs and base station. If the air link line rate exceeds the total hard bandwidth commitments, the new connection is allowed. If the total hard bandwidth commitments meet or exceed the air link line rate, the CAC module **206** denies the new connection.

In the communication system described above, each connection between the CPE **104** and base station **102** will have a planned PHY mode. The planned PHY mode is used by the CAC module **206** in determining whether to allow the new connection. As will be explained below, the calculation of the total hard bandwidth commitments for any given sector **106** (see FIG. **1**) presents additional difficulties for communication systems **100** which adapt PHY modes.

In communication systems **100** that adapt, or vary, their PHY modes, the available bandwidth necessary for existing connections can vary. Since each PHY mode used by the base station **102** and/or CPE **104** for its communication link **110** (*a*)-(*c*), **112**(*a*)-(*c*) is adaptive, the robustness of each communication link can vary (see FIG. **1**). As the robustness varies, the bandwidth allocated for an existing connection or new connection will also vary.

In such communication systems, connections are allowed to be modulated with PHY modes that are more or less robust than the planned PHY mode. Each end user connection can dynamically select its current PHY mode. This current PHY mode can be different than the planned PHY mode that was planned for the connection. If a connection is modulated using a more robust PHY mode than the planned PHY mode, the connection will exceed its allocated bandwidth.

In an embodiment of a communication system **100** that adapts PHY modes, the CAC module **206** allows new connections with reference to a minimum air link line rate. The minimum air link line rate is a measure of bandwidth that would be required if all of the existing connections between the CPEs and base station were modulated using a least efficient PHY mode regardless of whether the least efficient PHY mode is actually used. The least efficient PHY mode can include, for example, QAM-4 modulation with a maximum amount of FEC overhead bits. This method ensures that during adverse weather conditions each CPE will be able to select its least efficient PHY mode and transmit its data within its assigned bandwidth without losing its connection with the base station. In this embodiment, the CAC module **206** will deny a new connection if the new connection will cause the CPE to exceed its minimum air link line rate. The CAC module **206** can determine whether to allow or deny a new connection in conjunction with the control module **212**. During spells of good weather, the CPE can select a less robust PHY mode for its current PHY mode. By selecting a less robust PHY mode, additional bandwidth between the CPE and base station would be freed up. However, the communication system **100** is constrained from taking advantage of the freed up bandwidth when the decision to allow new connections is based upon the minimum air link line rate.

In another embodiment of the communication system **100** that adapts PHY modes, the CAC module **206** allows the CPE to take advantage of the freed up bandwidth. The CAC module **206** limits new connections based on a comparison of the bandwidth required for the connection if it is modulated using the CPE's planned PHY mode with the available bandwidth. The available bandwidth is determined by summing the CPE's hard bandwidth commitments that would be used by the existing connections if those connections were modulated using the planned PHY mode of the CPE. If the available bandwidth is equal to or exceeds the bandwidth required for the new connection, the CAC module **206** will allow the connection. However, if the CPE operates using a less robust PHY mode than its preferred PHY mode, there is the potential that data through the CPE will be lost.

In the presence of adaptive PHY modes and to take advantage of the CPE's planned PHY mode, the bit rate associated with each connection's PHY mode is compared. Connections at different PHY modes (modulation and FEC) effectively have different bit rates, or air link line rates, and thus are not directly compared. One method for comparing these bit rates is to normalize the PHY modes associated with each connection.

Equation 1, below, can be used to normalize the bandwidth used for connections through an individual CPE.

$$W_{CPEi} = \sum_{i=1}^{n} ER * \text{mod} \qquad \text{Equation 1}$$

Where W_{CPEi} is a normalized value or weight for the entire bandwidth used by an individual CPE. W_{CPEi} is proportional to the equivalent bandwidth of its connections and the current modulation associated with each connection. Er is the number of bits per unit time that are transmitted by the CPE for a connection. Each connection is modulated using an associated PHY mode. The term mod is the inverse of the associated PHY mode efficiency that is used to modulate the connection. The bit/symbol rate for QAM-64 is 6, for QAM-16 is 4, and for QAM-4 is 2. For example, if during a first connection between CPE **104**(*a*) and the base station **102**, 10,000 bits/s were transmitted using QAM-4, and during a second connection between CPE **104**(*a*) and the base station, 18,000 bits/s were transmitted using QAM-64, Equation 1 would be:

$W_{CPE104(a)}$=(10,000 bits/s*½ symbol/bit)+(18,000 bits/s*⅙ symbol/bit)=8,000 bits/s.

The 8,000 bits/s for CPE **104**(*a*) is then added to $W_{CPE104(b)}$ and $W_{CPE104(c)}$ to determine a total normalized bandwidth for the CPEs in sector **106**.

Normalization is used to determine the effective hard bandwidth commitment usage through the modem **108**. The CAC module **206** interfaces with the control module **212** to compare the different PHY modes for the existing connections and the new connection with the available bandwidth between the base station **102** and CPEs **104**. In this embodiment, the control module **212** is configured to normalize each CPE's air link line rate. Once the control module **212** has determined the normalized value for each CPE's committed bandwidth requirements, the CAC module **206** can sum and compare them against a common air link line rate.

Equation 2, below, can be used by the CAC module **206** to determine the total bandwidth used, i.e. W_{Link}−W, by all of the CPEs in the sector.

$$W_{Link} = \sum_{i=1}^{n} W_{CPEi} \quad \text{Equation 2}$$

Where W_{CPEi} is a normalized value or weight for the entire bandwidth used by an individual CPE in the sector.

For example, with PHY modes of QAM-4, QAM-16, and QAM-64, each using the same FEC, QAM-4 requires 3 times the air link resources, or bandwidth, of QAM-64 and QAM-16 requires 1.5 times the air link resources of QAM-64. In this example, the control module **212** can normalize to QAM-64. Thus, CPEs operating at QAM-64 would have their hard bandwidth commitments multiplied by a weight of 1, CPE's operating at QAM-16 would have their hard bandwidth commitments multiplied by a weight of 1.5, and CPE's operating at QAM-4 would have their hard bandwidth commitments multiplied by a weight of 3. The CAC module **206** then sums these hard bandwidth commitments and compares the total against a line rate of a communication link operating entirely at the selected normalized PHY mode, QAM-64 with the single FEC. Alternatively, the control module **212** normalizes to QAM-4 by applying weights of ⅓ to QAM-64, ½ to QAM-16, and 1 to QAM-4. The selection of QAM-64 and QAM-4, each with a single FEC, for use as a normalization PHY mode are only examples. Any PHY mode could be used to define the air link line rate for normalizing the connections between the CPEs and base station.

Still referring to FIG. **3**, the precedence module **210** will now be described. The precedence module **210** interfaces with the receiver module **202** and the control module **212** to apply a priority, or precedence, to one or more connections when less bandwidth is available than required to meet the hard bandwidth commitments. This can occur when the CAC module **206** is configured as described above to limit new connections based on planned PHY modes of the CPEs but some or all of the CPEs are operating at a more robust (less efficient) current PHY mode. The precedence module **210** determines which connection(s) are to be suspended. However, before connections are suspended, the base station **102** can re-allocate bandwidth, that is not intended for hard bandwidth commitments, among the CPEs to increase the available bandwidth for hard bandwidth commitments. Alternatively or in addition to, in TDD systems, the base station **102** can adjust the portion of a downlink subframe **302** and of an uplink subframe **304** in the frame **300** (see FIG. **2**) to increase the available bandwidth for a CPE that requires additional bandwidth due to a change in the current PHY mode or the addition of a connection. However, if additional bandwidth is not available, the precedence module **210** selects which connections from among the CPEs are suspended.

Bandwidth problems can arise when one or more CPEs are using more robust PHY modes than their planned PHY modes for their connections. For example, if communication system **100** was designed for 99.99% availability, a comparison would be made between a CPE's geographical proximity to the base station and the communication system's rain region. Based on this comparison, a planned PHY mode is selected for that CPE that allows it to operate at that planned PHY mode or a less robust PHY mode the entire year except for approximately 53 minutes. If a CPE exceeds a SNR or BER threshold and transmits its uplink using a more robust PHY mode than its planned PHY mode, it will require additional bandwidth for these 53 minutes. At least two things can occur during this 53 minutes depending on whether additional air link resources in the communication system **100** are available. Should additional bandwidth be needed when only a few existing connections, between the base station **102** and CPEs **104** in sector **106**, select a more robust PHY mode, the base station **102** may be able to reallocate the available bandwidth. Thus, if the communication system is sufficiently under subscribed, the CPE **104** can use the additional air link resources it requires when using a more robust PHY mode than its planned PHY mode during the 53 minutes. If many existing connections between the base station and CPEs are subject to similar adverse environmental conditions, the base station **102** may be unable to accommodate the CPEs' bandwidth requests. When the air link resources aren't available, the precedence module **210** selects which of the existing connections from the CPEs **104** (*a*)-(*c*) to suspend.

The precedence module **210** interfaces with the control module **212** to compare the bit rates for the existing connections through each CPE based on each CPE's current PHY mode. While the CAC module **206** compares the planned PHY modes of the CPEs to determine whether a new connection is allowed, the precedence module **210** compares the current PHY modes to the selected reference air link line rate to determine if a suspension should occur. The control module **212** is configured to compare the current PHY modes of the CPEs. As explained above, one method for comparing the PHY modes is normalization. Once normalized, the precedence module **210** determines if additional bandwidth between the CPEs and base station is available. If additional bandwidth is available, the precedence module **210** can determine a margin value. If additional bandwidth is not available, the precedence module **210** selects which connections are going to be suspended.

The precedence module **210** can be configured to suspend enough connections through the CPE that is requesting additional bandwidth until there is enough bandwidth to meet the remaining demand. The amount of outage during the year for the connections through the affected CPE **104** is planned based on the availability and rain region as discussed above. CPEs **104** located at greater distances from the base station **102** or having limited visibility of the base station would more likely be subject to the application of precedence. In this embodiment, CPE's are penalized by their geographic proximity to the base station **102**. For example, the same CPEs, those that are barely able to meet their availability numbers at their planned PHY modes, would be the first to have their hard bandwidth connections with the base station **102** suspended. These CPEs may receive the full brunt of the planned 53 minutes per year outage. In contrast, other CPEs (in particular, those barely unable to meet the availability number at the next less robust PHY mode) would have plenty of bandwidth because connections through the geographically challenged CPE's would be suspended before they need to drop to a more robust PHY mode and request additional bandwidth.

Alternatively, the precedence module **210** can also randomly select connections for suspension or select them in a round robin fashion. The precedence module **210** chooses connection to suspend from the entire set of connections that have hard bandwidth commitments through the CPEs in the sector **106**. The CPEs subject to potential suspension include CPEs that may still be operating at their planned PHY mode. In this embodiment, the communication system **100** as a whole, and each individual connection still meets its availability numbers since the planned outage is evenly shared. For example, if a rain fade caused the base station **102** and CPEs to lose half of their bandwidth, each connection from among all of the CPEs would, on average, see only 26 minutes outage per year rather than 53 minutes. Thus, the precedence aspect of adaptive CAC can allow you to increase system availability

(26 minutes outage vs 53 minutes outage) or capacity. For example, operating a CPE **104** at a less robust PHY mode than would typically be planned for the CPE increases the system's capacity. The communication system can rely on adaptive CAC coupled with precedence to distribute the outage among all of the CPEs. This achieves the planned 53 minutes outage, but with increased modulation efficiency for the CPE operating at the less robust PHY mode.

Further, the precedence module **210** can use levels in conjunction with the random selection method discussed above when selecting which connections to suspend. In this embodiment, each connection between the CPEs **104**(*a*)-(*c*) and base station **102** is assigned a precedence level. Alternatively, each CPE is assigned a precedence level for its connections. For example, there are five levels, one through five, with precedence level one being assigned to the most important connections and precedence level five being assigned to the least important connections. The random selection of connections for suspension is applied as discussed above with reference to the second embodiment. However, instead of applying the method of the second embodiment to all connections simultaneously, the precedence module **210** applies it based on each connection's assigned precedence level. Continuing with the example above, the random selection would be initially applied to connections assigned to precedence level five. If and when the precedence level five connections are exhausted, the precedence module **210** applies the random selection process to connections assigned to precedence level four and so on until there is adequate bandwidth available for the remaining connections that have hard bandwidth commitments. Thus, individual connections can be selected to have their uplink or downlink transmissions suspended in favor of other connections.

Further, the precedence module **210** can allow connections to continue to operate with their current PHY mode even when a first SNR or BER threshold is exceeded. Instead, a second threshold is implemented to maintain the connection at the same PHY mode. However, the error rate associated with the connection may increase.

FIG. **4** is a flowchart illustrating the process of adaptively adjusting a PHY mode for a connection between the base station **102** and a CPE. This process can be implemented by a modem **108** at a base station. Alternatively, this process is performed by a modem **108** at the CPE. A specific CPE **104** can change its uplink PHY mode independent of that CPE's downlink PHY mode. The specific CPE's PHY mode can also be independent of the uplink PHY modes used by other CPEs **104** within the same sector **106**. Because the base station **102** must synchronize with each individual CPE **104** that uplinks data, the uplink quality may be different than the downlink quality with a specific CPE **104**. The base station **102** can perform the process of adaptively adjusting the uplink PHY mode used by a specific CPE **104**. As such, a similar process may be completed for each CPE **104** within the sector **106** in order to adaptively adjust each CPEs **104** uplink modulation.

The following description describes a process for adaptively adjusting a PHY mode for an uplink from a CPE to a base station. The same process is used for adaptively adjusting a PHY mode for a downlink from the base station to the CPE.

In particular, flow begins in start block **400**. Flow moves to a block **402** where a receiver module **202** at a base station **102** receives an uplink from a CPE **104**. Flow proceeds to block **404**, where the quality of the channel parameters for the uplink **112** is determined by a receive signal quality (RSQ) module **208**. The quality may be a function of the state of the transmission medium (e.g. air, foggy air, wet air, smoky air, etc.) and the ability of both the transmitting and receiving components (e.g. CPE **104** and base station **102**) to respectively transmit and receive data. The base station **102** can determine the quality of each uplink **112**(*a*)-(*c*). Alternatively, the base station **102** periodically transmits channel parameter measurements, which are indicative of the quality of a CPE's uplink **112**, to that CPE **104**. The CPE **104** then uses these channel parameter measurements to determine the quality of its uplink. These channel parameter measurements can include a SNR and/or a BER measurement of the uplink **112**(*a*)-(*c*). For example, base station **102** can determine the quality of uplink **112**(*c*) based on a measurement by its RSQ module **208** (see FIG. **3**). A single SNR measurement or a series of several SNR measurements taken during a frame **300** (see FIG. **2**) or during multiple frames may be used to determine the uplink quality. The control module **212** can analyze multiple measurements to determine an uplink's quality.

Continuing to block **406**, the base station **102** or CPE **104** compares the calculated uplink quality with a current PHY mode threshold. The current PHY mode threshold can include an upper threshold and a lower threshold at which the PHY mode is changed. For example, if CPE **104**(*a*) is currently uplinking data to base station **102** using PHY mode B, the PHY mode will change when the uplink quality exceeds an upper threshold or goes below a lower threshold.

Next at decision block **408**, the CPE determines whether the uplink quality has decreased and crossed a PHY mode lower threshold according to the comparison made in block **406**. Continuing with the example above, if the PHY mode lower threshold associated with PHY mode B has not been crossed, flow proceeds to decision block **410** where the system determines whether the uplink quality has crossed an upper PHY mode threshold associated with PHY mode B. If the current modulation upper threshold has been exceeded, flow continues to block **412** where the PHY mode is changed to a less robust, denser modulation. For example, PHY mode C is selected for CPE **104**(*a*). The base station **102** can send a request to the CPE **104** indicating a desired uplink PHY mode change. Alternatively, the base station **102** transmits an uplink map to all CPEs **104** in the downlink subframe **302** (see FIG. **2**) indicating which CPEs have been allotted uplink PS's and the PS's associated PHY modes. The base station **102** indicates to an individual CPE **104** that the PHY mode has been changed by allotting uplink subframe **304** PSs to that CPE that use a less robust PHY mode. For example, if the uplink PHY mode for CPE **104**(*a*) is to be changed from PHY mode B to PHY mode C, the base station **102** assigns uplink subframe PS's which are to be modulated using PHY mode C. This uplink assignment serves as an indicator to the CPE that its uplink PHY mode has been change. Flow continues to a block **413** where the system can reallocate the newly available bandwidth. For example, the newly available bandwidth can be allocated for new or existing hard bandwidth commitments, new connections, or connections that had been previously suspended. Flow then returns to block **402** as described above.

Returning to decision block **410**, if the current PHY mode upper threshold has not been exceeded, flow continues to block **402** as described above.

Returning to decision block **408**, if the PHY mode lower threshold has been crossed, flow proceeds to a decision block **414** where the system determines whether the connections, between the CPE and base station that have a hard bandwidth commitment, are using a less robust PHY mode than the planned PHY mode for the connections. If the connection(s) is using a less robust PHY mode than its planned PHY mode, the process proceeds to block **416** where a more robust PHY mode is selected for the connection(s). If the base station determines whether the uplink quality has crossed a threshold, the base station **102** can send a request to the CPE **104** indicating a desired uplink PHY mode change. Alternatively, the base station **102** can transmit an uplink map to all CPEs **104** in the downlink subframe **302** indicating which CPEs have been allotted uplink PS's along with the PS's associated PHY modes. This allows the base station **102** to indicate to an individual CPE **104** that the PHY mode has been changed by allotting uplink subframe **304** PSs to that CPE that uses a more robust PHY mode. For example, if the uplink PHY mode for CPE **104**(*a*) is to be changed from PHY mode B to PHY mode A, the base station **102** assigns uplink subframe PS's which are to be modulated using PHY mode A. This uplink assignment serves as an indicator to CPE **104**(*a*) that its uplink PHY mode has been change. Flow then continues to block **420** where a precedence module **210** (see FIG. **3**) determines whether connections between the base station and the CPEs are to be suspended. Precedence will be explained with reference to FIG. **5**. Flow then continues to block **402** as described above.

Returning to decision block **414**, if the connection's current PHY mode is at least as robust as its planned PHY mode, the process continues to decision block **418** where the control module **212** can replace the lower threshold associated with the current PHY mode of the connection that has the hard bandwidth commitment with a second lower threshold. The process continues to block **402** as described above except that at block **406** the RSQ module **208** and the control module **212** use the second lower threshold to compare with the measured signal quality of the connection.

Returning to decision block **418**, if the control module does not select the second lower threshold, the process moves to a block **420**, as described above, where the precedence module **210** (see FIG. **3**) determines whether connections between the base station and the CPEs are to be suspended. Precedence will be explained with reference to FIG. **5**. Once precedence has been applied, the process returns to state **402** as described above.

FIG. **5** is a flowchart illustrating the process of applying precedence to existing connections between the CPEs **104** and the base station that have hard bandwidth commitments. This process can be implemented by a modem **108** at a base station. Alternatively, this process is performed by a modem **108** at the CPE. Flow begins in start block **600**. Flow moves to block **601** where a more robust PHY mode is selected for the existing connection. Flow proceeds to block **602** where the control module **212** determines an air link line rate based on a reference PHY mode. Flow moves to block **603** where the control module calculates the hard bandwidth commitments for the existing connections between the base station **102** and CPEs **104** based on the current PHY mode for each connection. Flow moves to a decision block **604** where the precedence module **210** determines whether the air link line rate determined at block **602** exceeds the hard bandwidth commitments between the CPEs and base station. If the air link line rate exceeds the hard bandwidth commitments, the process continues to a block **606** where the more robust PHY mode selected in block **601** is applied for the existing connection. Flow then returns to block **402** of FIG. **4** where the base station **102** receives the next uplink from a CPE **104**.

Returning to decision block **604**, if the air link line rate does not exceed the hard bandwidth commitments, flow proceeds to a decision block **608** where the precedence module **210** determines whether additional air link resources are available. These additional air link resources can include available bandwidth in the uplink subframe **302** and available bandwidth in the downlink subframe **304** (see FIG. **2**). If additional air link resources are available, flow proceeds to block **606** where the more robust PHY mode is applied for the existing connection. Flow then returns to block **402** of FIG. **4** where the base station **102** receives the next uplink from a CPE **104**.

Returning to decision block **608**, if additional air link resources are not available, flow moves to a block **610** where the precedence module **210** suspends existing connections between the base station **102** and the CPEs **104**. As described above, the precedence module **210** can, for example, suspend connections only between the base station and the affected CPE, randomly suspend connections between the base station and all of the CPEs in a sector **106**, or suspend connections between the base station and all of the CPEs in the sector in a round-robin fashion. Further, the precedence module **210** can randomly suspend connections between the base station and the CPEs that have a lower precedence priority than other connections. Alternatively, the precedence module **210** can suspend the connections that have a lower precedence priority in a round-robin fashion. The process moves to block **606** as described above where the more robust PHY mode is applied for the existing connection. The process then returns to block **402** of FIG. **4** where the base station **102** receives the next uplink from a CPE **104**.

FIG. **6** is a flowchart illustrating the process of call admission control for a new connection between a CPE and the base station. This process can be implemented at a base station. Alternatively, this process is performed at the CPE. Flow begins in start block **500**. Flow proceeds to block **502** where the base station receiver module receives a request for a new connection. The process continues to block **504** where the CAC module **206** sums the hard bandwidth commitments between the CPEs and base station based on the planned modulations of the CPEs. Next, at a block **506**, the control module **212** determines an air link line rate for the existing connections between the base station and CPEs based on the reference PHY mode. Flow moves to a decision block **508** where the CAC module **206** determines whether the air link line rate determined at block **506** exceeds the hard bandwidth commitments determined at block **504**. If the air link line rate exceeds the hard bandwidth commitments, the process continues to a block **510** where the CAC module **206** allows the new connection. However, air link resources are not initially allocated to the connection since the connection has been allowed based on the planned PHY modes of the CPEs and base station. The CPEs and base station could be operated at a more robust PHY mode than their planned PHY mode.

Flow proceeds to block **512** where the control module **212** determines the hard bandwidth commitments for the existing connections between the base station **102** and CPEs **104** based on the current PHY mode for each connection. Flow moves to a decision block **514** where the precedence module **210** determines whether the air link line rate determined at block **506** exceeds the hard bandwidth commitments between the CPEs and base station determined at block **512**. If the air link line rate exceeds the hard bandwidth commitments, the process continues to a block **516** where the base station allocates air link resources to the new connection. Flow then returns to block **502** where the base station **102** receives a request for a new connection.

Returning to decision block **514**, if the air link line rate does not exceed the hard bandwidth commitments, flow proceeds to a decision block **518** where the precedence module **210** determines whether additional air link resources are available. These additional air link resources can include available bandwidth in the uplink subframe **302** and available bandwidth in the downlink subframe **304** (see FIG. **2**). If additional air link resources are available, flow proceeds to block **516** where the base station allocates air link resources to the new connection. Flow then returns to block **502** where the base station **102** receives a request for a new connection.

Returning to decision block **518**, if additional air link resources are not available, flow moves to a block **520** where the precedence module **210** suspends existing connections between the base station **102** and the CPEs **104**. As described above, the precedence module **210** can, for example, suspend connections only between the base station and the affected CPE, randomly suspend connections between the base station and all of the CPEs in a sector **106**, or suspend connections between the base station and all of the CPEs in the sector in a round-robin fashion. Alternatively, the new connection is accepted into a suspended state since the precedence module **210** has already determined which of the other connections are to be suspended. Further, the precedence module **210** can randomly suspend connections between the base station and the CPEs that have a lower precedence priority than other connections. Alternatively, the precedence module **210** can suspend the connections that have a lower precedence priority in a round-robin fashion. The process moves to block **516** where the base station allocates air link resources to the new connection. Flow then returns to block **502** where the base station **102** awaits a request for a new connection.

Returning to decision block **508**, if the air link line rate does not exceed the hard bandwidth commitments, flow proceeds to a block **522** where the CAC module **206** denies the new connection. The process then returns to block **502** to await the next request for a new connection.

The foregoing description details certain embodiments of the invention. It will be appreciated, however, that no matter how detailed the foregoing appears in text, the invention can be practiced in many ways. As is also stated above, it should be noted that the use of particular terminology when describing certain features or aspects of the embodiments should not be taken to imply that the terminology is being re-defined herein to be restricted to including any specific characteristics of the features or aspects of the embodiment with which that terminology is associated. The scope of the embodiments should therefore be construed in accordance with the appended claims and any equivalents thereof.

What is claimed is:

1. A subscriber station for a wireless communication system comprising:

a modem section configured to receive downlink data from a base station on a downlink link and to transmit uplink data to the base station on an uplink link shared with other subscribers stations;

a receive signal quality module configured to monitor a downlink (DL) quality parameter for the downlink data providing a parameter value; and

a control section configured to:

determine a preferred downlink physical (PHY) mode for the downlink data among a plurality of PHY modes of different degrees of robustness, the preferred downlink PHY mode being defined between a first and a second threshold for the parameter value;

instruct the modem section to transmit to the base station an indication of the preferred downlink PHY mode;

identify in a DL sub-frame map received from the base station, a current downlink PHY mode selected for the subscriber station based on the preferred downlink PHY mode and the bandwidth available to the subscriber station on the downlink link; and

instruct the modem section to receive the downlink data based on the current downlink PHY mode,

wherein the downlink PHY mode specifies a modulation format and a forward error correction technique used for transmission of downlink data.

2. A subscriber station as claimed in claim **1**, wherein the current downlink PHY mode is different from the preferred downlink PHY mode.

3. A subscriber station as in claim **1**, wherein the DL quality parameter is the signal to noise ratio (SNR).

4. A subscriber station as in claim **3**, wherein a less robust preferred downlink PHY mode is selected by the subscriber station if the SNR value is above the first and the second threshold.

5. A subscriber station as in claim **1**, wherein the DL quality parameter is the bit error rate (BER).

6. A subscriber station as in claim **5**, wherein a less robust preferred downlink PHY mode is selected by the subscriber station if the BER value is under the second and the first threshold.

7. A subscriber station as claimed in claim **1**, wherein the control section further selects the downlink PHY mode based on the subscriber station capabilities.

8. A subscriber station as claimed in claim **1**, wherein the control module is further configured to:

identify in an UL sub-frame map received from the base station, a current uplink PHY mode selected for the subscriber station based on a UL quality parameter for the uplink data and the bandwidth available to the subscriber station on the uplink link; and

instruct the modem section to transmit the uplink data using the uplink PHY mode.

9. A subscriber station as claimed in claim **1**, further comprising a precedence module that interfaces with the modem module and the control module to apply a priority to one or more uplink connections established at the subscriber station when less uplink bandwidth is available for the subscriber station on the uplink link than required to serve all uplink connections.

10. A subscriber station as claimed in claim **9**, wherein the control module is configured to limit establishment of a new uplink connection based on the uplink bandwidth allocated to the subscriber station and a planned uplink PHY mode.

11. A subscriber station as claimed in claim **10**, wherein the control module is configured to suspend an uplink connection based on the bandwidth allocated to the subscriber station and the priority of all uplink connections served by the subscriber station.

12. A method of communication in a wireless communication system, comprising:

receiving at a subscriber station downlink data from a base station on a downlink link and transmitting uplink data to the base station on an uplink link shared with other subscribers stations;

monitoring a quality parameter for the downlink data and providing a parameter value;

determining a preferred downlink physical (PHY) mode for the downlink data among a plurality of PHY modes of different degrees of robustness, the preferred PHY mode being defined between a first and a second threshold for the parameter value;

transmitting to the base station an indication of the preferred PHY mode;

identifying in a DL sub-frame map received from the base station, a current downlink PHY mode selected for the subscriber station based on the preferred downlink PHY mode and the bandwidth available to the subscriber station on the downlink link; and

receiving the downlink data having the current downlink PHY mode,

wherein the downlink PHY mode specifies a modulation format and a forward error correction technique used for transmission of downlink data.

13. A method as claimed in claim **12**, wherein the quality parameter is the signal to noise ratio (SNR), and wherein a more robust preferred downlink PHY mode is provided by the subscriber station if the SNR value is under the first and the second threshold.

14. A method as claimed in claim **12**, wherein the quality parameter is the bit error rate (BER).

15. A method as claimed in claim **14**, wherein a more robust preferred downlink PHY mode is selected by the subscriber station if the BER value is above the second and the first threshold.

16. A method of communication in a wireless communication system, comprising:

receiving at a subscriber station downlink data from a base station on a downlink link and transmitting uplink data to the base station on an uplink link shared with other subscribers stations;

monitoring a parameter for the downlink data and providing a parameter value;

transmitting to the base station the parameter value for enabling the base station to establish a current downlink PHY mode for the subscriber station based on a first and a second threshold for the parameter value and the downlink resources available to the subscriber station on the down-link; and

identifying in a DL sub-frame map received from the base station, the current downlink PHY mode; and

receiving the downlink data having the current downlink PHY mode,

wherein the downlink PHY mode specifies a modulation format and a forward error correction technique used for transmission of downlink data.

17. A method as claimed in claim **16**, wherein the parameter value is one of a BER value and a SNR value.

18. A method as claimed in claim **16**, further comprising establishing a current uplink PHY mode for transmission of uplink data, wherein the uplink PHY mode specifies a modulation format and a forward error correction technique used for transmission of uplink data.

19. A method as claimed in claim **18**, wherein the current uplink and downlink PHY modes can be adapted independently based on one or more parameters for the uplink and respectively downlink data.

20. A method as claimed in claim **18**, wherein the step of establishing a current uplink PHY mode comprises:

receiving from the base station an indication of the quality of the uplink data and determining if the quality of uplink data is between a first and a second thresholds;

selecting a less robust UL PHY mode than the current UL PHY mode if the quality of uplink data increased over the first and second threshold; and

selecting a more robust UL PHY mode than the current UL PHY mode if the quality of uplink data is decreased under the first and second threshold.

21. A method as claimed in claim **18**, wherein the step of establishing a current uplink PHY mode further comprises maintaining the current UL PHY mode if the quality of the uplink data between the first and second thresholds.

22. A method as claimed in claim **20**, further comprising, when a more robust UL PHY mode is selected:

determining the amount of UL bandwidth available to the subscriber station using a reference PHY mode for the subscriber station;

calculating a hard bandwidth commitment for all UL connections currently established on the CPE, based on the current PHY mode; and

transmitting the uplink data using the more robust PHY mode if the amount of UL bandwidth available to the subscriber station exceeds the hard bandwidth commitment.

23. A method as claimed in claim **22**, further comprising, if the amount of UL bandwidth available to the subscriber station is less than the hard bandwidth commitment for the CPE,

re-select a less robust uplink PHY mode whenever additional bandwidth is allocated to the subscriber station from the base station; and

suspend a connection established on the subscriber station if additional bandwidth is not available to the subscriber station.

24. A method as claimed in claim **20**, further comprising, wherein the bandwidth resulting from using a more robust UL PHY mode is reallocated to other CPEs or connection for UL transmission.

* * * * *

CERTIFICATE OF COMPLIANCE

The brief complies with the type-volume limitation of Fed. Cir. R. 32(b)(1) because this brief contains 13,725 words, excluding the parts of the brief exempted by Fed. R. App. P. 32(f) and Fed. Cir. R. 32(b)(2).

This brief complies with the typeface requirements of Fed. R. App. P. 32(a)(5) and the type style requirements of Fed. R. App. P. 32(a)(6) because this brief has been prepared in a proportionally spaced typeface using Microsoft Word 2016 in Century Schoolbook 14-point font.

ORRICK, HERRINGTON & SUTCLIFFE LLP

/s/ Mark S. Davies
Mark S. Davies
Counsel for Appellant

CERTIFICATE OF SERVICE

I hereby certify that I electronically filed the foregoing with the Clerk of the Court for the United States Court of Appeals for the Federal Circuit by using the appellate CM/ECF system on November 20, 2020.

A copy of the Opening Brief with Confidential Addendum was served upon the following counsel of record via FedEx and electronic mail:

Jeffrey A. Lamken
Lucas M. Walker
Rayiner I. Hashem
Leonid Grinberg
MOLOLAMKEN LLP
600 New Hampshire Ave., NW,
Suite 500
Washington, DC 20037
jlamken@mololamken.com
lwalker@mololamken.com
rhashem@mololamken.com
lgrinberg@mololamken.com

Mike McKool
MCKOOL SMITH, P.C.
300 Crescent Court, Suite 1500
Dallas, TX 75201
mmckool@mckoolsmith.com

Scott Cole
MCKOOL SMITH, P.C.
300 West 6th Street, Suite 1700
Austin, TX 78701
scole@mckoolsmith.com

Brett E. Cooper
MCKOOL SMITH, P.C.
One Manhattan West
395 9th Avenue, 50th Floor
New York, NY 10001
bcooper@mckoolsmith.com

ORRICK, HERRINGTON & SUTCLIFFE LLP

/s/ Mark S. Davies
Mark S. Davies
Counsel for Appellant