

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

LENOVO (UNITED STATES) INC. and
MOTOROLA MOBILITY LLC,

Petitioners,

v.

HEADWATER RESEARCH LLC,

Patent Owner.

Case No. IPR2024-01180
U.S. Patent No. 9,198,076

**DECLARATION OF HARRY BIMS, Ph.D.
REGARDING U.S. PATENT NO. 9,198,076**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	SUMMARY OF OPINIONS	1
III.	BACKGROUND AND QUALIFICATIONS	2
	A. Background.....	2
	B. Compensation	7
IV.	MATERIALS CONSIDERED	8
V.	LEGAL STANDARDS.....	8
	A. Level of Ordinary Skill.....	9
	B. Claim Construction.....	10
	C. Anticipation under 35 U.S.C. § 102	11
	D. Obviousness under 35 U.S.C. § 103	11
	E. Secondary Considerations of Non-Obviousness	14
VI.	BACKGROUND AND TECHNOLOGY OF THE '076 PATENT.....	16
	A. Overview of the '076 Patent	16
	B. Claims of the '076 patent.....	18
VII.	ANALYSIS.....	18
	A. Ground 1: The Combination of <i>Rao</i> in View of <i>Araujo</i> Renders Claims 1-3 and 5-16 Obvious	19
	1. Overview of <i>Rao</i>	19
	2. Overview of <i>Araujo</i>	22
	3. Rationale for Modifying <i>Rao</i> in View of <i>Araujo</i>	28
	4. Independent Claim 1.....	32

- a. 1[pre]: “A wireless end-user device, comprising:” 32
- b. 1[a]: “a wireless wide area network (WWAN) modem to communicate data for Internet service activities between the device and at least one WWAN, when configured for and connected to the at least one WWAN; and” 33
- c. 1[b]: “one or more processors” 35
- d. 1[c]: “configured to, for a time when data communication for Internet service activities is provided by the WWAN modem and the at least one WWAN, classify whether a first end-user application associated with an Internet access request, and capable of both interacting with a user in a user interface foreground of the device, and at least some Internet service activities when not interacting with a user in the device user interface foreground, is interacting with the user in the device user interface foreground, and,” 35
- e. 1[d(1)]: “allow or disallow the Internet access request based on the classification as to whether the first end-user application is interacting with the user and based on a power control state,” 39
- f. 1[d(2)]: “such that, when the power control state is a first power control state and the first end-user application is classified as not interacting with the user, the Internet access request is disallowed, and” 44
- g. 1[d(3)]: “when the power control state is at least one other power control state, the Internet access request is allowed.” 45
- 5. Claim 2: “The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is interacting with the user in the device user interface foreground

	when the user of the device is directly interacting with that application or perceiving any benefit from that application.”	46
6.	Claim 5: “The wireless end-user device of claim 1, further comprising, when the Internet access request is disallowed, queuing the Internet access request until a power control state change occurs.”	47
7.	Claim 7: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to allow the Internet access request when the power control state is the first power control state and the first end-user application is classified as interacting with the user.”	49
8.	Claim 8: “The wireless end-user device of claim 1, further comprising a user interface, wherein the user interface is to inform the user of the device when there are options to set, control, override, or modify service usage controls that affect the processor configuration that determines whether to allow or disallow Internet access requests in the first power control state.”	50
9.	Claim 9	52
a.	9[a]: “The wireless end-user device of claim 1, further comprising: a wireless local area network (WLAN) modem to communicate data for Internet service activities between the device and at least one WLAN, when configured for and connected to the at least one WLAN,”	52
b.	9[b]: “wherein the one or more processors are further configured to allow the Internet access request when the WLAN modem provides data communication for the Internet access request and the power control state is the first power control state.”	53

10. Claim 10: “The wireless end-user device of claim 1,
wherein the power control state is a power state of the
device.” 55
11. Claim 11: “The wireless end-user device of claim 1,
wherein the first power control state is a power save state
of the device.” 57
12. Claim 12: “The wireless end-user device of claim 1,
wherein the power control state is a power state of the
WWAN modem.” 58
13. Claim 13: “The wireless end-user device of claim 1,
wherein the one or more processors are further
configured to, based on a device usage state, dynamically
change a determination of whether to allow or disallow
the Internet access request in the first power control state
when the first end-user application is classified as not
interacting with the user.” 59
14. Claim 14: “The wireless end-user device of claim 1,
wherein the one or more processors are further
configured to, based on power control state changes for
the WWAN modem, dynamically change a determination
of whether to allow or disallow the Internet access request
in the first power control state when the first end-user
application is classified as not interacting with the user.” 62
15. Claim 15: “The wireless end-user device of claim 1,
wherein, when the Internet access request is disallowed,
the one or more processors are configured to prevent the
Internet access request from causing a change to a power
state of the modem.” 63
16. Claim 16: “The wireless end-user device of claim 1,
wherein the one or more processors are configured to,
when the Internet access request is disallowed, instruct
the first end-user application to transition to a different
state.” 64

B.	Ground 2: Claim 3 Is Obvious Over <i>Rao</i> in View of <i>Araujo</i> and <i>Singh</i>	67
1.	Overview of <i>Singh</i>	67
2.	Rationale for Modifying <i>Rao</i> in View of <i>Araujo</i> and <i>Singh</i>	68
3.	Claim 3: “The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is interacting with the user in the device user interface foreground based on a state of user interface priority for the application.”	70
C.	Ground 3: Claim 4 Is Obvious Over <i>Rao</i> in View of <i>Araujo</i> and <i>Freund</i>	71
1.	Overview of <i>Freund</i>	71
2.	Rationale for Modifying <i>Rao</i> in View of <i>Araujo</i> and <i>Freund</i>	72
3.	Claim 4: “The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is not interacting with the user in the device user interface foreground when the application is providing or utilizing a background data service.”	75
D.	Ground 4: Claim 6 Is Obvious Over <i>Rao</i> in View of <i>Araujo</i> and <i>Aleksic</i>	76
1.	Overview of <i>Aleksic</i>	76
2.	Rationale for Modifying <i>Rao</i> in View of <i>Araujo</i> and <i>Aleksic</i>	77
3.	Claim 6: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to, for at least a second application or service, allow Internet access requests in the first power control	

	state without regard to whether the second application or service is classified as interacting with the user.”	79
E.	Ground 5: Claims 16 Is Obvious Over <i>Rao</i> in View of <i>Araujo</i> and <i>Montemurro</i>	80
1.	Overview of <i>Montemurro</i>	80
2.	Rationale for Modifying <i>Rao</i> in View of <i>Araujo</i> and <i>Montemurro</i>	82
3.	Claim 16: “The wireless end-user device of claim 1, wherein the one or more processors are configured to, when the Internet access request is disallowed, instruct the first end-user application to transition to a different state.”	84
VIII.	DECLARATION	86

I, Harry Bims, Ph.D., declare as follows:

I. INTRODUCTION

1. I have been asked to submit this declaration on behalf of Lenovo (United States) Inc. and Motorola Mobility LLC (“Petitioners”) in connection with a petition for *inter partes* review of U.S. Patent No. 9,198,076 (“the ’076 patent”). I have been told that the petition is being submitted to the Patent Trial and Appeal Board of the United States Patent and Trademark Office by the Petitioners.

2. Specifically, I have been retained as a technical expert by the Petitioners to study and provide my opinions on the technology claimed in, and the patentability or unpatentability of, claims 1-16 of the ’076 patent (“the Challenged Claims”). For purposes of this declaration, I was not asked to provide any opinions that are not expressed herein.

II. SUMMARY OF OPINIONS

3. This declaration is directed to the Challenged Claims of the ’076 patent, and sets forth the opinions I have formed, the conclusions I have reached, and the bases for each.

4. In forming my opinions, counsel for Petitioners told me to assume that the effective filing date of the ’076 patent is January 28, 2009.

5. Based on my experience, knowledge of the art as of the effective filing date of the ’076 patent, analysis of prior art references, and the understanding

of the claim terms in view of the specification that a person of ordinary skill in the art would have as of the effective filing date, it is my opinion that the Challenged Claims of the '076 patent are unpatentable as obvious over the prior art references discussed below.

III. BACKGROUND AND QUALIFICATIONS

A. Background

6. My educational background, career history, publications, and other relevant qualifications provided here are only a summary. My full curriculum vitae is attached to this declaration as Exhibit 1003.

7. Based on my education, research, and work experience, I am knowledgeable about the subject matter of the '076 patent and the related prior art.

8. I have worked extensively in the field of digital communications. I have studied telecommunications and systems engineering since approximately 1981. Further, I have over 30 years of industry experience in computer network design, including the design of hardware and software for computer communications in a wireless context. During this period, I have designed and implemented various products involving cellular communication technology.

9. I received a B.S. in Computer and Systems Engineering from Rensselaer Polytechnic Institute in 1985. In 1988, I received a M.S. in Electrical Engineering from Stanford University. In 1993, I received a Ph.D. in Electrical

Engineering, also from Stanford University. As a graduate student at Stanford University, I studied the principles of digital communications theory, including data modulation and demodulation, error checking and correction algorithms, and the architecture and design of semiconductor circuits used for digital communications. My Ph.D. thesis at Stanford addressed the application of trellis coding and precoding to a digital modulation system, and was titled “Trellis Coding for Multi-Level, Partial Response Continuous Phase Modulation with Precoding.”

10. After receiving my Ph.D. in 1993, I worked for Glenayre Technologies - Wireless Access Group, where I focused on hardware and software architecture and design, including inventing, designing, and building a patented computer system for real-time testing of two-way pagers and co-developing a wireless application protocol that included a CRC error checking algorithm. From 1999 to 2001, I was the Director of Engineering for Symmetry Communications, where I was responsible for managing the software development for 2.5G network equipment (the precursor to 4G network equipment) that handles communications to and from personal cellular telephones. The equipment products we developed (a “Serving GPRS Support Node” (or “SGSN”) and “Gateway GPRS Support Node” or (“GGSN”)) routed and managed cellular phone communications by relaying packets of voice, data, and video between the base station tower and the Internet,

regulating the network bandwidth allocated to different classes of cellular phone applications, and transferring control of cellular phone communications from one base station to the next as cellular phones move throughout a geographic area. I also held management responsibility for the Firmware, Hardware, Performance, and Systems Engineering Groups.

11. In January of 1999, I launched a technology consulting company, Protocomm Systems Inc., which focuses on the development of advanced wireless communications protocols and related software implementations for wireless product companies.

12. From 1999 to 2001, I was also the Director of Software Architecture for Symmetry Communications System. In this position, I was responsible for the software architecture for core SGSN (Serving GPRS Support Node) and GGSN (Gateway GPRS Support Node) products for the GPRS (Generic Packet Radio Services) market. I also held management responsibility for the Firmware, Hardware, Performance, and Systems Engineering Groups.

13. In 2001, I developed a business plan for building network infrastructure for 802.11 enterprise networks, and then later that year founded AirFlow Networks, Inc., where I invented and received over eleven patents on its core technology, which was based on the 802.11 wireless local area network specification.

14. I have been a member of the Institute of Electrical and Electronics Engineers (“IEEE”) for more than 28 years. The IEEE is an international organization of more than 400,000 engineers, and is the world’s largest technical professional society, with locations in more than 160 countries. I am currently a Senior Member of the IEEE, and serve in a variety of leadership roles, including as a reviewer and voter on draft amendments to technical standards such as 802.11 (Wi-Fi), and 802.3 (Ethernet) prior to publication by their 802 Committee. The IEEE 802 Committee is the standards development organization responsible for publishing protocol standards for Ethernet, Wi-Fi, and a variety of other wired and wireless protocols that are in commercial use globally.

15. I am currently an expert consultant for Protocomm Systems, LLC and Bims Laboratories, LLC, both of which I founded. I perform technical research in wireless technology standards, such as UMTS, LTE/4G, 5G, Wi-Fi, Bluetooth, and other network communication protocols. I am named as an inventor of twenty-three telecommunications related patents. One of those patents is U.S. Patent No. 8,935,580 (“’580 patent”) which issued January 13, 2015 and is currently assigned to Apple Inc. The application for the ’580 patent was filed with the USPTO on June 18, 2013, with priority to an earlier application filing date of July 2, 2008, which became U.S. Patent 8,468,426. The ’580 patent, titled “Multimedia-Aware Quality-of-Service and Error Correction Provisioning,” discloses a method of

encoding a stream of data packets in which the packets are routed to separate “QoS processing modules,” each possessing their own QoS profile, where the packets are processed according to multiple factors, such as priority, traffic rate, latency, tolerated jitter, transmitter scheduling, media type, etc. Subsequent to this processing, the packets are combined and transmitted.

16. The services I provide include consulting in telecommunication standards setting, protocol simulation and implementation, technology assessments, engineering lab testing, and product analysis. Through Bims Laboratories, I have participated in technical discussions for several years in the IEEE 802 Committee to develop new technology that allows LTE and 5G networks to coordinate their network activities with Wi-Fi networks.

17. I have also served in several leadership capacities with the IEEE 802.16 Working Group (2009-present) and the IEEE 802.15 Working Group where I currently serve as technical editor for a wide-area cellular standard that is a competitor to LTE and 5G, and I am currently a voting member in the IEEE 802.11 Working Group. Within the IEEE 802.15 Working Group, I am the technical editor for IEEE 802.16t, an amendment to the WiMAX standard that supports long-range wireless networks for the railroad industry. Furthermore, since 2003, I have served as a wireless technology expert in patent infringement cases. I am actively involved in liaison conversations regarding the interoperation of IEEE 802.11-

based wireless networks and 3GPP networks, including 3G, LTE, and 5G networks.

18. I am an author of 9 technical publications and a named inventor of 23 telecommunication-related patents, including U.S. Patent No. 8,995,996, which describes how wireless devices that interfere with each other can coordinate their communications in an optimal manner. The entire list of those patents and publications can be found in my full CV.

19. I have been active in professional engineering associations and research bodies. From 1985 to 1991, I was the AT&T Bell Labs Cooperative Research Fellow. In 2000, I served as a Chair of IEEE Engineering Management Society, Silicon Valley Chapter. From 2007 to 2010, I served as a vice chair and secretary of IEEE 802.16h License-Exempt Task Group and have been a senior member of IEEE since 2012.

B. Compensation

20. I am being compensated for services provided in this matter at \$850 per hour. My compensation is not conditioned on the conclusions I reach as a result of my analysis or on the outcome of this matter, and in no way affects the substance of my statements in this declaration.

21. I do not have any financial interest in any Petitioner, or any of its subsidiaries. I also do not have any financial interest in Patent Owner, Headwater

Research LLC. Further, I do not have any financial interest in the '076 patent, and I have not had any contact with the named inventors of the '076 patent: Gregory G. Raleigh, James Lavine, and Alireza Raissinia.

IV. MATERIALS CONSIDERED

22. In forming my opinions, I read and considered the '076 patent and its prosecution history, the exhibits listed in the Exhibit List filed with the petition for *inter partes* review of the '076 patent, as well as any other materials referenced herein.

V. LEGAL STANDARDS

23. In forming my opinions and considering the subject matter of the '076 patent and its claims in light of the prior art, I am relying on certain legal principles that counsel in this case has explained to me. I understand that the claims define the invention. I also understand that an unpatentability analysis is a two-step process. First, the claims of the patent are construed to determine their meaning and scope. Second, once the claims have been construed, the content of the prior art is compared to the construed claims. I have been told the following legal principles apply to analysis of patentability based on 35 U.S.C. §§ 102 and 103. I also have been told that, in an *inter partes* review proceeding, a patent claim may be deemed unpatentable if it is shown by a preponderance of the evidence that the claim was anticipated by a prior art patent or publication under

§ 102 and/or rendered obvious by one or more prior art patents or publications under § 103.

A. Level of Ordinary Skill

24. I understand a person of ordinary skill in the art is determined by looking at (1) type of problems encountered in the art; (2) prior art solutions to those problems; (3) rapidity with which innovations are made; (4) sophistication of the technology; and (5) educational level of active workers in the field.

25. Therefore, based on my experience and accounting for these factors, I believe a person of ordinary skill in the art would have had, as of January 28, 2009 (the date I have been told to assume is the effective filing date of the '076 patent), (1) held at least a bachelor's degree in computer science, electrical engineering, or a related field, and (2) three to five years of experience with networking, power consumption of networked computing devices, and/or wireless digital communications systems. More education, e.g., a doctorate in an academic area emphasizing electrical engineering, computer engineering, computer science, or a related technical field, would compensate for fewer years of work experience. I meet this definition of one of ordinary skill in the art now and also at the earliest possible priority date of the '076 patent.

26. I am also familiar with the knowledge of the person of ordinary skill in the art as of the priority date of the '076 patent. I am able to opine on how the

person of ordinary skill in the art would have understood the disclosure and claims of the '076 patent, the disclosures of the prior art, the motivation to combine the prior art, and what combinations would have been obvious to one of ordinary skill in the art. In the following, I also refer to a person of ordinary skill in the art using the short-form acronym "POSITA."

B. Claim Construction

27. I understand that the United States Patent and Trademark Office interprets claim terms in an *inter partes* review proceeding under the same claim construction standard that is used in a United States federal court. I understand that under this standard, the meaning of claim terms is considered from the viewpoint of one of ordinary skill in the art at the time of the alleged invention.

28. I have been informed that claim terms are generally given their ordinary and customary meaning as understood by one of ordinary skill in the art in light of the specification and the prosecution history pertaining to the patent. I understand, however, that claim terms are generally not limited by the embodiments described in the specification.

29. I also understand that in addition to the claims, specification, and prosecution history, other evidence may be considered to ascertain the meaning of claim terms, including textbooks, encyclopedias, articles, and dictionaries. I have

been informed that this other evidence is often less significant and less reliable than the claims, specification, and prosecution history.

C. Anticipation under 35 U.S.C. § 102

30. I understand that a patent claim is invalid as anticipated if a single piece of prior art teaches every element of the claims, viewed from the perspective of a person of ordinary skill in the art.

31. I have also been told that an anticipatory reference does not have to recite word for word what is in the anticipated claim. Rather, anticipation can occur when a claimed limitation is “inherent” in the relevant reference. I have been told that if the prior art necessarily functions in accordance with, or includes, the claim limitations, it can anticipate even though the reference does not recite a limitation word for word.

D. Obviousness under 35 U.S.C. § 103

32. I understand that a patent claim is invalid as obvious if the claimed invention would have been obvious to a person of ordinary skill in the art at the time the claimed invention was made. This means that even if all of the elements of the claim cannot be found in a single prior art reference that would anticipate the claim, a person of ordinary skill in the field who knew about all the prior art would have come up with the claimed invention. I understand that in an obviousness determination, the person of ordinary skill in the art is presumed to have

knowledge of all material prior art. I understand that whether a claim is obvious is based upon the determination of several factual issues.

33. In considering obviousness, I understand that one must determine the scope and content of the prior art. I understand that, in order to be considered as prior art to a patent being considered, a prior art reference must be reasonably related to the claimed invention of that patent. A reference is reasonably related if it is in the same field as the claimed invention or is from another field to which a person of ordinary skill in the art would look to solve a known problem.

34. I understand that one must determine what differences, if any, existed between the claimed invention and the prior art.

35. I understand that a patent claim composed of several elements is not proved obvious merely by demonstrating that each of its elements was independently known in the prior art. In evaluating whether such a claim would have been obvious, one may consider whether a reason has been identified that would have prompted a person of ordinary skill in the art to combine the elements or concepts from the prior art in the same way as in the claimed invention. There is no single way to define the line between true inventiveness on the one hand (which is patentable) and the application of common sense and ordinary skill to solve a problem on the other hand (which is not patentable). For example, market forces or

other design incentives may be what produced a change, rather than true inventiveness.

36. I understand that one may consider whether (1) the change was merely the predictable result of using prior art elements according to their known functions, or whether it was the result of true inventiveness; (2) there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; (3) the innovation applies a known technique that had been used to improve a similar device or method in a similar way; (4) the claimed invention would have been obvious to try, meaning that the claimed innovation was one of a relatively small number of possible approaches to the problem with a reasonable expectation of success by those skilled in the art; (5) the invention merely substituted one known element for another known element in order to obtain predictable results; (6) the invention merely applies a known technique to a known device, method, or product to yield predictable results, or (7) known work in the field may have prompted variations of use of the same inventions in the same or different fields due to market forces or design incentives that would have been predictable to a person of ordinary skill in the art.

37. I understand that one must be careful not to determine obviousness using the benefit of hindsight; many true inventions might seem obvious after the fact. I understand that the determination should be based on the position of a

person of ordinary skill in the field at the time the claimed invention was made and should not consider what is known today or what is learned from the teaching of the patent.

E. Secondary Considerations of Non-Obviousness

38. I understand that certain secondary considerations may be considered in evaluating obviousness in order to prevent hindsight bias. These secondary considerations include commercial success of products that practice the patent, long-felt need for the patented technology, failure by others to solve the problem addressed by the patent, initial skepticism by others in the industry, industry recognition and praise of the patented products, and efforts by others to copy the patented technology.

39. I have been informed that these secondary considerations of non-obviousness must be balanced against the strength of the prima facie case of obviousness. Where the invention represents no more than the predictable use of prior art elements according to their established functions, secondary considerations are inadequate to establish non-obviousness.

40. I have also been informed that the patentee must establish a nexus between any secondary consideration factors and the claimed invention. Where a secondary consideration factor results from something other than what is both

claimed and novel in the patent, there is no nexus to the merits of the claimed invention.

41. I have been informed that, for example, the nexus requirement as it applies to the evidence of commercial success requires that the patentee must show that the sales of a product that allegedly practices the claimed invention must be a direct result of the unique characteristics of the claimed invention, as opposed to other economic and commercial factors that are unrelated to the patented invention. I understand that this means if the commercial success is due to an element that existed in the prior art or that is not claimed by the patent, then no nexus exists.

42. I have been further informed that, for example, the nexus requirement as it applies to the long-felt need element must likewise be supported by evidence. Where the differences between the prior art and the claimed invention are minimal, it cannot be said that any long-felt need was met by the patented invention or that the patent solved any technological problems that were unaddressed by others.

43. I have been informed that evidence of industry praise must also show a nexus with the claimed invention, and that self-serving statements by the patentee regarding the merits of the claimed invention are generally not given weight.

VI. BACKGROUND AND TECHNOLOGY OF THE '076 PATENT

A. Overview of the '076 Patent

44. According to the '076 patent, wireless network resources are stretched thin. “The increasing popularity of various smart phone devices, net book devices, [and] tablet computing devices” is leading to a “network capacity crunch.” '076 patent, 3:28-37. Network resources, e.g., base stations, are inundated with requests for network access (i.e., requests for Internet access), leading to “overloading” of base stations. '076 patent, 6:19-67. When this happens, base stations become “overwhelmed” by the requests and may begin “fall[ing] behind” in servicing the requests. '076 patent, 6:56-67.

45. The '076 patent explains that a device’s access requests may vary in several ways. First, they may vary in their priority. Typically, each request is associated with an application. '076 patent, 4:6-21. Because the services provided by these applications vary, the priority of requests may vary. For example, a service that is “interactive” in the foreground of a user’s device may be higher priority because it requires “reasonably low response time” for a request, while a service occurring in the background may be “lowest priority.” '076 patent, 16:35-44. For some services, it may make sense to hold off sending requests until a “scheduled transmission time.” '076 patent, 86:44-55.

46. Second, requests may also vary in their likelihood to lead to repeated requests. As one example, a device may permit an application to access the network and then later interrupt the access due to the device's "power save algorithm." '076 patent, at 7:57-62. This algorithm causes the device to "terminate[] the connection to the base station" each time the device enters a "power save state" and to reinitiate the connection when exiting the "power save state." '076 patent, 7:50-57. As a result, "each time the device enters [the] power save state and then exits [the] power save state[,] network resources are consumed." '076 patent,, 7:42-57. If the device employs "an aggressive power save algorithm that enters [the] power save state after a short idle period," repeated requests for network access occur, causing the device to "consume a proportionally large amount of resources." '076 patent, 7:50-62; *see also* '076 patent, 93:31-42.

47. To avoid "overloading" the base stations, the '076 patent explains, the device intercepts requests by applications to launch or access the network. '076 patent, 6:19-67, 11:15-37, 65:12-28. The device then handles requests differently based on conventional classifications, such as priority or likelihood to lead to repeated requests, as set forth in a "control policy." '076 patent, 11:46-49, 12:38-46. For example, one or more processors in the device classify whether an application associated with an Internet access request is interacting with a user. '076 patent, Abstract. Then, depending on a power control state (of the device or of

a wireless modem), the Internet access request is either allowed or disallowed.

'076 patent, Abstract, claim 1.

B. Claims of the '076 patent

48. I understand that the IPR Petition challenges claims 1-16 of the '076 patent. I further understand that among those claims, claim 1 is an independent claim.

VII. ANALYSIS

49. Based on my personal experience at the relevant time and my review of the prior art, it is my opinion that: (1) claims 1, 2, 5, and 7-16 would have been obvious to a POSITA as of the effective filing date of the '076 patent based on a combination of *Rao* and *Araujo*; (2) claim 3 would have been obvious to a POSITA as of the effective filing date of the '076 patent based on a combination of *Rao* in view of *Araujo* and *Singh*; (3) claim 4 would have been obvious to a POSITA as of the effective filing date of the '076 patent based on a combination of *Rao* in view of *Araujo* and *Freund*; (4) claim 6 would have been obvious to a POSITA as of the effective filing date of the '076 patent based on a combination of *Rao* in view of *Araujo* and *Aleksic*; and (5) claim 16 would have also been obvious to a POSITA as of the effective filing date of the '076 patent based on a combination of *Rao* in view of *Araujo* and *Montemurro*. In my opinion, the

Challenged Claims recite combinations of well-known prior art elements that perform their known functions to produce predictable results.

A. Ground 1: The Combination of *Rao* in View of *Araujo* Renders Claims 1-3 and 5-16 Obvious

1. Overview of *Rao*

50. Like the '076 patent, *Rao* “generally relates to optimizing peer-to-peer network communications, and in particular, to techniques for application-aware prioritizations of network communications on a client.” *Rao*, [0002]. In my opinion, because *Rao* is directed to the same field of endeavor as the '076 patent (application management in devices that can run multiple applications) and same problem as the '076 patent (optimizing use of network communications on an application-by-application basis), *Rao* is analogous art to the '076 patent.

51. As *Rao* explains, in prior art computing devices (also referred to as clients), “[c]ommunications over a network . . . are typically processed in the order generated by the activity of a user and applications of the client[.]” *Rao*, [0003]. According to *Rao*, though, this is not “always desirable” since “a network packet generated or received for an application running in the background may be processed ahead of a network packet generated or received for the application running in the foreground.” *Rao*, [0003]. As one example, “a network packet generated or received for an application running in the background may be

processed ahead of a network packet generated or received for the application running in the foreground.” *Rao*, [0003]. In another example, a voice-over-IP (VoIP) telephone call may experience increased latency and reduced quality because the packets of another application were processed ahead of the packets of the VoIP telephone call. *Rao*, [0003].

52. To address these inefficient prior art architectures, *Rao* describes a prioritization scheme that “provide[s] application-aware, client-specific prioritization of packet traffic.” *Rao*, [0003]. *Rao*’s system includes one or more computing devices 102 (also referred to as clients 105) connected to a network 104 via one or more network connections 341. *Rao*, [0086]. The prioritization scheme augments client 105 with a remote access client 120 that interacts with applications 338 and network stack 310. *Rao*, [0098], [0100]. Remote access client 120 “manages and controls network communication prioritizations on the client.” *Rao*, [0080]. Figure 5A shows a system having remote access client 120 for routing network packets from a client 105 to a network:




Fig. 5A

Rao, Fig. 5A (annotated). The client 105 runs an agent 326 and a filter 322, which in one embodiment, may together form the remote access client 120 “for routing packets via a network, or for otherwise providing remote access connectivity.” Rao, [0099].

53. Rao discloses implementing the filter 322 as a “network driver, such as a network driver operating at any layer, or portion thereof, of a network stack” at the client. Rao, [0101]. Rao’s network stack includes, for example, “one or more network drivers” associated with “any network interface cards or network access components of the computing device.” Rao, [0100].


54. *Rao*'s prioritization can be based on policies, which "apply a priority to network packets of applications 338a-338n in accordance with [specified] prioritization rules." *Rao*, [0189]. *Rao*'s policies are specified by an application's name, type, protocol, or any other suitable means. *Rao*, [0182]. These policies offer several approaches for prioritization, including whether an application is in the foreground or background. *Rao*, [0182]

2. Overview of *Araujo*

55. *Araujo* is directed to application-aware optimization of power management on a mobile device. Just as *Rao* teaches application-aware prioritization for a mobile device, *Araujo* teaches application-aware scheduling to optimize transmission timing and reduce conflicts, *Araujo* teaches a complementary application-aware power management scheme for a mobile device. Like the '076 patent, *Araujo* implements policies to control activities of "wireless telephones" on an application-by-application basis. (*Araujo*, [0014].) In my opinion, because *Araujo* is directed to the same field of endeavor as the '076 patent (application management in devices that can run multiple applications), *Araujo* is analogous art to the '076 patent.

56. Just as *Rao* uses policies to prioritize network access by certain applications over others, *Araujo* discloses using "power-management polic[ies]" to prioritize power consumption by certain applications ("programs" in *Araujo*) over

others. *Araujo*, Abstract. As *Araujo* explains, “[s]ome programs merit the use of power to service any request[,]” while “[o]ther programs do not.” *Araujo*, [0016]. Yet typical power management techniques “generally do not take [this] . . . into account when evaluating requests” from programs “to power a device” at a “machine,” such as a “wireless telephone[.]” *Araujo*, [0014], [0016]. One example of a “device” in *Araujo* is “machine 102’s wireless network card” that allows the machine “to send data over a wireless network.” *Araujo*, [0020].



Araujo, Fig. 2.

57. *Araujo*'s power management technique is illustrated in Figure 2.

Araujo, [0023]-[0035]. At 202, a "request" is received from an application, such as "a request to send data over a wireless network" using a "network card." *Araujo*, [0020], [0024]. The request is received by a "[r]equest deflector component," "interposed between" the application and a driver for the network card, that "intercepts requests and responds to the requests based on power management considerations before the request reach[es] a ... driver" for the network card. *Araujo*, [0020], [0024], [0050]. "At 204, it is determined whether servicing of the request involves a change in the [network card's] power state, such as a change from off to on, a change from one intermediate power state to another, or a change between an extreme power state (on or off) and an intermediate power state." *Araujo*, [0025]. If so, "it may be determined (at 208) whether a policy (e.g. a power management policy) allows the [network card's] power state to be changed to service the request." *Araujo*, [0025]. To make this determination, *Araujo* discloses considering a "status" of the application making the request and a "current power state of the device." *Araujo*, [0027], [0028].

58. The status "reflects its worthiness to consume power." *Araujo*, [0027]. "For example," some applications "have a 'VIP' status that allows them to consume power in situations where other [applications] would not be permitted to consume power." *Araujo*, [0027]. This status could be a "binary VIP/non-VIP

status,” or “could be indicated on a scale.” *Araujo*, [0027]. “Whether [an application] is permitted to consume power could be based on a finding as to whether the [application’s] status justifies the consumption of power under the circumstances that are present.” *Araujo*, [0027].

59. The “current power state of a device” is also considered in determining whether the network card’s power state can be changed to service the application’s request. *Araujo*, [0028]. “[A] device [e.g., a network card] may use a considerable amount of power to turn from off to on,” but “changing a device from a low power state to a high power state may use less power than turning the device from off to on.” *Araujo*, [0028]. “Thus, the current power state of the device may be a factor in deciding whether to change [the] power state, and the way in which the current power state is considered may be written into a policy.” *Araujo*, [0028].

60. Figure 2 shows a “power use policy” that “specifies how the various items of information,” including the “status” of the application and the “current power state of a device” (e.g., the network card) “are used to determine whether [the network card’s] power state will be changed.” *Araujo*, [0027], [0028], [0034]. “If the policy allows the [network card’s] power state to be changed, then the request is allowed.” *Araujo*, [0035]. But “[i]f the policy does not allow the [network card’s] power state to be changed . . . , then the request may be processed

in the manner shown in FIG. 4 (at 212).” *Araujo*, [0035]. For example, as Figure 4 shows, the application’s request may be “block[ed].” *Araujo*, [0040].




FIG. 4

Araujo, Fig. 4.

61. Thus, *Araujo* discloses a power management technique in which requests from applications at a machine to perform an action using a network card (one of *Araujo*’s “devices”) may be allowed or blocked based on an application-specific power management policy. By implementing such policies, a device better “use[s] the energy in a way that strikes a balance between providing functionality and maintaining longevity of the charge” at the wireless telephone. *Araujo*, [0001].

3. Rationale for Modifying *Rao* in View of *Araujo*

62. In my opinion, it would have been obvious to modify *Rao* to incorporate *Araujo*'s power management policy to provide a combination that better "use[s] the energy in a way that strikes a balance between providing functionality and maintaining longevity of the charge." *Araujo*, [0001]. *Rao* discloses prioritizing network packets from applications based on application-by-application "prioritization rules specified or indicated by ... policies." *Rao*, [0189]. *Rao* envisions a "multitude of ways to define client-side priorities" in its policies, including "whether an application is running in the foreground or the background." *Rao*, [0182]. But *Rao* does not expressly disclose using a device's power state in its policies to prioritize application network packets.

63. *Rao* recognized that its policy-based prioritization techniques "may improve or increase the performance, operational characteristics, and user experience on the client based on the applications running on the client." *Rao*, [0080]. In my opinion, a POSITA would have understood that "performance, operational characteristics, and user experience" each implicate prioritization to maintain the longevity of the battery in the device, and availability of network resources. *Araujo* would have motivated a POSITA to augment *Rao*'s policies to additionally consider a device's power state in prioritizing network packets from applications. In particular, in my opinion, a POSITA would have been motivated to

augment *Rao*'s policies to specify that network packets from certain applications should be allowed or blocked based not only on whether the application is currently in the foreground or the background but also on a device's power state.

64. As explained above, *Araujo* discloses a power management technique in which a request from an application to perform an action using a network card (one of *Araujo*'s "devices") may be allowed or blocked based on an application's power management policy. This policy specifies whether to allow or block the request based both on an application status and the device's power state. *Rao*'s policies, when augmented to consider the device's power state, would specify whether an application request is to be allowed or blocked based both on whether the application is currently in the device's foreground or background (as taught by *Rao*, and as similar to *Araujo*'s VIP/non-VIP status) and the device's power state. These policies are based on characteristics that may change over time, requiring dynamic association of the policies as the application's status or device's power state changes.

65. For example, *Araujo*'s request deflector functionality would be implemented in connection with or as part of the filter 322 of *Rao*'s remote access client, which—like *Araujo*'s request deflector—applies application-specific policies to determine how to handle applications' network access attempts. *Araujo*, [0024], [0050]; *Rao*, [0191]-[0193]. In this example, *Rao/Araujo*'s remote access

client would employ application-specific policies that specify whether to allow or block a request from an application based on a check of (i) whether the application is currently in the device's foreground or the background (taught by *Rao*) and (ii) the device's power state (taught by *Araujo*).

66. A POSITA would have been motivated to augment *Rao*'s policies to employ *Araujo*'s power management technique at least because *Araujo*'s technique "strike[s] the right balance between power conservation and performance"—a goal that was known in the prior art. (*Araujo*, [0015].) Battery-powered devices "have a finite amount of energy stored between charges." *Araujo*, [0001]. These devices therefore "attempt to use the energy in a way that strikes a balance between providing functionality and maintaining longevity of the charge" and "employ power management schemes to govern the use of energy." *Araujo*, [0001]. The prior art, however, recognized "a need ... for an improved power management operation in portable devices." *Aleksic*, [0004]. This was especially true for devices running multiple applications and having access to multiple network connections. *D'Amore*, [0031]; *Araujo*, [0002]-[0006]. Thus, prior art like *Araujo* recognized that one method of improving a device's power management was an application-aware prioritization of network communications that optimized the power consumption of individual applications. *Araujo*, [0002]-[0005]; *D'Amore*, [0031]; *Rao*, [0080].

67. In my opinion, a POSITA would have understood that improved, application-specific, power management schemes reduce competition among applications for other resources, including processing power at the device and network resources. *D'Amore*, [0003], [0004], [0006]; *Oestvall*, [0021]; *Aleksic*, [0022]-[0025]. By selectively allowing network access, power management schemes reduce competition among applications for network resources. *D'Amore*, [0006]. The result is that certain applications may access the network at the expense of others. *D'Amore*, [0006]. This is particularly helpful, for example, when accessing a low bandwidth network or running a high priority application requiring more processing power. Using a scheme to manage an individual application's power consumption optimizes a device's resource consumption, improving an application's operation and user experience.

68. Indeed, *Rao* recognized that its policy-based prioritization techniques “may improve or increase the performance, operational characteristics, and user experience on the client based on the applications running on the client.” *Rao*, [0080]. A POSITA would have understood that the “performance, operational characteristics, and user experience” described in *Rao* would implicate, for instance, battery power, device longevity, and availability of network resources, but *Rao* does not expressly disclose its prioritization techniques serve to improve all of these aspects of the device. A POSITA would have thus been motivated to

modify *Rao* to include well-known, application-aware, power monitoring capabilities like those described in *Araujo* that reduce power consumption, decrease energy expenses, optimize network resource allocation, and increase device longevity.

69. A POSITA would have had a reasonable expectation of success in modifying *Rao* to incorporate *Araujo*'s power management functionality. Such a modification would have been well within the abilities of a POSITA. For example, *Araujo*'s policy "can be implemented as software that is stored in one or more of the data remembrance component(s) 704 and that executes on one or more of the processor(s) 702." *Araujo*, [0055]. This scheme could thus be implemented using the processor and memory units in *Rao*'s architecture. *See Rao*, [0119], [0120]. The resulting modification would have involved combining prior art elements according to known methods to yield predictable results and use of a known technique to improve similar devices in the same way.

4. Independent Claim 1

a. 1[pre]: "A wireless end-user device, comprising:"


70. In my opinion, the combination of *Rao* and *Araujo* renders obvious the preamble. *Rao* discloses "multiple computing devices 102a-102c" (also referred to as "clients 105a-105c") that connect to a network using one or more network connections. *Rao*, [0086]. "[C]lient 105 can be any type and/or form of

computing device 102 that can run one or more applications 338,” such as a “web browser, web-based client, client-server application, [or] a thin-client computing client.” *Rao*, [0088]. For example, *Rao*’s device may be a “mobile client,” such as a “notebook, personal digital assistant (PDA), a smart phone” or “telecommunication device.” *Rao*, [0198]; *see also Rao*, [0130]. *Rao*’s client “may include a network interface 118,” such as a “network interface card,” to “interface to a Local Area Network (LAN), Wide Area Network (WAN) or the Internet through a variety of connections including ... wireless connection.” *Rao*, [0125]. *Rao*’s device is an “end-user device” because *Rao* discloses a device “user.” *Rao*, [0183], [0199].

- b. 1[a]: “a wireless wide area network (WWAN) modem to communicate data for Internet service activities between the device and at least one WWAN, when configured for and connected to the at least one WWAN; and”**

71. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 1[a]. *Rao*’s device includes network interface 118, which enables connectivity to a “Local Area Network (LAN), Wide Area Network (WAN) or the Internet through a variety of connections,” such as “standard telephone lines, LAN or WAN links,” “broadband connections,” “wireless connections,” or “some combination of any or all of the above.” *Rao*, [0095], [0125]. Network interface 118 includes “a built-in network adapter, network interface card,

PCMCIA network card, card bus network adapter, wireless network adapter, USB network adapter, modem, or any other device suitable for interfacing the computing device 102 to any type of network capable of communication.” *Rao*, [0125]. This is shown in Figure 1D below:



(*Rao*, Fig. 1D (annotated).) *Rao*’s network interface, including the modem, enables connectivity and data communication for Internet service activities between the

device and a wireless wide area network (WWAN), such that *Rao* discloses a “[WWAN] modem,” as claimed.

c. 1[b]: “one or more processors”

72. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 1[b]. *Rao*’s device includes “central processing unit 102.” *Rao*, [0119], Fig. 1D (incorrectly labeled as “121” and not “102”). According to *Rao*, “[t]he central processing unit 102 is any logic circuitry that responds to and processes instructions fetched from the main memory unit 104.” *Rao*, [0119]. *Rao* discloses that, in some embodiments, “the central processing unit is provided by a microprocessor unit.” *Rao*, [0119] (listing examples). A POSITA would have understood that this central processing unit undertakes the operations claimed in 1[c] and 1[d], as discussed in more detail below.

d. 1[c]: “configured to, for a time when data communication for Internet service activities is provided by the WWAN modem and the at least one WWAN, classify whether a first end-user application associated with an Internet access request, and capable of both interacting with a user in a user interface foreground of the device, and at least some Internet service activities when not interacting with a user in the device user interface foreground, is interacting with the user in the device user interface foreground, and,”

73. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 1[c]. *Rao*’s devices each include “one or more applications 338a-338n.” *Rao*,

[0087]. These “applications 338a-338n provide one or more real-time data communications, such as VoIP” or “may provide email, collaboration, online meeting, and/or desktop sharing related services or functionality.” *Rao*, [0179]; *see also Rao*, [0088] (describing breadth of applications). In my opinion, a POSITA would have understood that these applications are associated with Internet access requests because they must request network access to access the Internet and carry out their respective services or functionality. Each of these applications are an “end-user application,” as claimed.

74. *Rao*’s central processing unit, running remote access client 120, can classify whether a given application is running in the foreground or the background of the mobile device. According to *Rao*, “remote access client 120 may determine whether the application 338a-338n associated with the network packet is running in the foreground or the background of the client 105.” *Rao*, [0188]; *see also Rao*, [0189] (disclosing how “agent 326 may use characteristics of the application 338a-338n, such as running in the foreground or background”).)

75. In my opinion, a POSITA would have understood that *Rao*’s remote access client runs on central processing unit 102. *Rao*, [0118], [0119]. A central processing unit is a hardware-based processor that is embedded in *Rao*’s devices and executes software on *Rao*’s devices, including applications 338a-338n, any operating system, and the remote access client. Because this functionality all

resides on the same central processing unit, a POSITA would have understood that the functionality of *Rao*'s remote access client is controlled via central processing unit 102.

76. A POSITA would have understood an application running in the foreground in *Rao* to be "interacting with a user in a user interface foreground of the device" because *Rao*'s remote access client classifies an application as "running in the foreground" based on a user's current interaction with that application. According to *Rao*, an "application is running in the foreground" where it is "in active use by the user." *Rao*, [0003]. In the context of *Rao*'s device, an application running in the foreground would present an interface foreground to facilitate the interaction, as claimed. For example, a POSITA would have understood that an application such as email would be capable of interacting with a user in a device's user interface foreground of the device (e.g., through an application like Microsoft Outlook) to create and send new messages. *Rao*, [0090]; *Wylde*, 5:7-19.

77. In my opinion, A POSITA would also have understood *Rao*'s applications to be "capable of both interacting with a user in a user interface foreground of the device, and at least some Internet service activities when not interacting with a user in the device user interface foreground" because it would have been well understood as of the '076 patent's priority date that an application

(including *Rao*'s applications) would be capable of at least some Internet activities, even when running in the background. For example, it was well known that an e-mail application was capable of using Internet access not only when in the foreground (e.g., when a user is sending emails) but also when in the background (e.g., to receive emails over the Internet even when the user is not interacting with the mobile device). *Rao*, [0090]; *Wyld*, 5:7-19; *Kelz*, [0040]. Similarly, an application such as a news aggregator may access the Internet in the background while the user attends to other unrelated tasks in the foreground. *Freund*, 10:24-27. This is consistent with other references, including *Oestvall*, which explain that, “without an explicit close command, applications will continue to run even when not actually in active use; they will therefore continue to use some system resources, even when residing in the background.” *Oestvall*, [0006].

78. *Rao*'s central process unit is therefore configured to classify whether an end-user application—that is capable of both interacting with a user in a user interface foreground and at least some Internet service activities when not interacting with a user in the device user interface foreground—is interacting with the user in the device user interface foreground. And the central processing unit is configured to perform this classification when data communication for Internet service activities is provided by the WWAN modem and the at least one WWAN, as discussed in Section VII.A.4.b (claim 1[a]).

- e. **1[d(1)]: “allow or disallow the Internet access request based on the classification as to whether the first end-user application is interacting with the user and based on a power control state,”**

79. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 1[d(1)]. As explained in Section VII.A.3 (Motivation), a POSITA would have been motivated and found it obvious to augment the application-specific policies taught by *Rao* to consider not only whether the application is currently in the foreground of the device (“is interacting with the user”), but also a power state of the device (“a power control state”). Thus, in an example implementation, *Rao/Araujo*’s remote access client would allow or disallow an application’s request for network access based both on (i) whether the application is currently in the device’s foreground (taught by *Rao*) and (ii) the device’s power state (taught by *Araujo*). A POSITA would have understood that the operations of *Rao/Araujo*’s remote access client are undertaken by *Rao/Araujo*’s central processing unit (“one or more processors”).

80. Indeed, *Araujo* expressly recognizes the value of considering an application’s status in connection with a power state when determining whether to allow or disallow network access. For instance, *Araujo* discloses that a “[n]etwork card” (e.g., modem) is an “an example of a device that may draw power while it is being used, and thus may have its power state managed.” *Araujo*, [0056]. This may

include, for example, either a WWAN modem or a WLAN modem. *See* Section VII.A.9 (claim 9); Section VII.A.14 (claim 14). Given that a network card “may use more power when transmitting,” a power management policy may assess whether a given application (or program) may consume power by accessing the Internet. *Araujo*, Abstract, [0021]. I note that the terms “application” (as used in *Rao*) and “program” (as used in *Araujo*) can be used interchangeably and I have done so here.

81. According to *Araujo*, “[w]hether a program is permitted to consume power,” e.g., by powering on the network card to access the network, “could be based on a finding as to whether the program’s status justifies the consumption of power under the circumstances that are present.” *Araujo*, [0027]. *Araujo* discloses that “some programs may have a ‘VIP’ status that allows them to consume power in situations where other programs would not be permitted to consume power.” *Araujo*, [0027]; *see also Araujo*, [0028], [0029], [0033] (explaining a VIP program can power on the network card). A POSITA would have understood that an application (“program” in *Araujo*) having a VIP status—a status that, like *Rao*’s foreground status, indicates “high priority”—would be permitted to consume power (e.g., power on the network card to access the network), regardless of the device’s or network card’s power state. An example of such an application is an emergency voice or alert application. *Aleksic*, [0041] (“For example, the

emergency call level 910 may be the highest priority level in that no other applications are allowed to run once this battery capacity level has been reached to allow the device at least to be used for emergency calls for safety reasons.”).

82. In contrast, *Araujo* discloses that some programs may have a non-VIP status that limits their ability to consume power in situations where VIP programs would be permitted to consume power. *Araujo*, [0027] (“Programs could be assigned binary VIP/non-VIP status, or their status could be indicated on a scale (e.g., a scale of 1-10, indicating a given program’s relative worthiness to receive a power allocation).”); *see also Araujo*, [0033] (explaining a non-VIP program cannot power on the network card). A POSITA would have understood that an application having a non-VIP status—a designation that, like *Rao*’s background status, indicates “low priority”—may not be permitted to consume power (e.g., power on the network card to access the network) when the device is in a particular power state, such as when the device or network card is off or in a low-power state.

83. Indeed, it was well known as of the priority date to a POSITA to use “priority based functional processing capabilit[ies] ... to shut down if desired, lower priority applications to allow enough battery capacity to operate the primary or higher priority applications.” *Aleksic*, [0033]; *see also D’Amore*, [0006] (“Furthermore, in one implementation, control of a resource can be taken away from an application having lower priority in order to affect control of such resource

for a higher priority application.”). As the prior art explained, “[a]s the power capacity of the device battery decreases, lower priority functions are either shut down or the quality of service of a feature is reduced in an attempt to conserve battery power.” *Aleksic*, [0039].

84. In my opinion, A POSITA would also have understood from *Araujo* that a non-VIP program would still be permitted to consume power (e.g., use the network card) when the network card is already powered on. *Araujo*, [0033] (explaining the non-VIP program may use the network card if it has already been powered on).

85. From *Araujo*, a POSITA would have been motivated and found it obvious to augment *Rao*’s policies to consider an application’s foreground/background status together with a device’s power state, just as *Araujo* describes considering an application’s VIP/non-VIP status together with a device’s power state. Further, *Araujo* describes allowing a VIP application to use the network card to access the network regardless of the device’s power state, *Rao/Araujo*’s augmented policies would allow a foreground application to access the network regardless of the device’s power state. Just as *Araujo* describes disallowing a non-VIP application from using the network card to access the network when the network card is off, *Rao/Araujo*’s augmented policies would disallow a background application from accessing the network when *Rao/Araujo*’s

network interface (or some other component needed for network access) is powered off. And just as *Araujo* describes allowing a non-VIP application to use the network card to access the network when the network card is on, *Rao/Araujo*'s augmented policies would allow a background application to access the network when *Rao/Araujo*'s network interface (or some other component needed for network access) is on. It was known in the art that permitting network access by applications in the foreground—and thus, “currently in active use by the user,” *Rao*, [0003])—was desirable as doing so would reduce perceived latency. *Oestvall*, [0021] (“When the untrusted application is brought to the foreground again [on a battery operated device], it is allowed to run again.”); *Wyld*, 1:66-2:2 (“When the computer is viewed as a time-sharing system, a task executing in foreground is given a higher priority, to use the resources in the computer, than a background task.”); *Senior*, 8:33-34, 14:64-65.

86. Thus, *Rao*'s policies, augmented according to *Araujo*, would allow or disallow an application's request for network access (“Internet access request”) (1) based on whether an application is in the foreground (“interacting with the user”) and (2) device's power state (“power control state”). In an example implementation, e.g., as described in Section VII.A.3 (Motivation), these augmented policies would be applied by *Rao/Araujo*'s remote access client (implemented to perform the functions of *Araujo*'s request deflector), which sits

between the applications and a driver for the network interface card (or other network access components) in *Rao*'s client. *Rao*, [0101], Fig. 5. In this implementation, when a request from an application was received, *Rao*'s remote access client would "intercept[] requests" for network access "and respond[] to the requests ... before the request reach[es] [the] device driver." *Araujo*, [0024], [0050]. Applying the augmented policies, *Rao/Araujo*'s remote access client would allow or disallow (e.g., "block," *Araujo*, [0040]) the request (1) based on whether the application is in the foreground and (2) based on a power control state.

- f. **1[d(2)]: "such that, when the power control state is a first power control state and the first end-user application is classified as not interacting with the user, the Internet access request is disallowed, and"**

87. As explained above, just as *Araujo* describes disallowing a non-VIP application from using the network card to access the network when the network card is off, *Rao/Araujo*'s augmented policies would disallow a background application from accessing the network when *Rao*'s network interface is powered off. Thus, when *Rao/Araujo*'s network interface is powered off ("first power control state") and the application requesting network access is in the background ("not interacting with the user"), *Rao/Araujo*'s remote access client (implemented to include *Araujo*'s request deflector functionality) would disallow (e.g., "block," *Araujo*, [0040]) the request for network access. Section VII.A.3 (Motivation);

Section VII.A.4.e (claim 1[d(1)]); *see also Araujo*, [0038], Fig. 4 (illustrating all options if a power state change is not allowed). A POSITA would have understood that the operations of *Rao*'s remote access client are undertaken by *Rao/Araujo*'s central processing unit ("one or more processors").

- g. **1[d(3)]: “when the power control state is at least one other power control state, the Internet access request is allowed.”**

88. *Rao/Araujo* renders obvious claim 1[d(3)]. As explained above, just as *Araujo* describes allowing a non-VIP application to use the network card to access the network when the network card is on ("at least one other power control state"), *Rao/Araujo*'s augmented policies would allow a background application to access the network when *Rao/Araujo*'s network interface is on. Thus, when *Rao/Araujo*'s network interface is powered on, *Rao/Araujo*'s remote access client (implemented to include *Araujo*'s request deflector functionality) would allow the network access request, regardless of whether the application is in the foreground or the background. Section VII.A.3; Section VII.A.4.e (claim 1[d(1)]); *Araujo*, [0037] ("the program's use of the device does not affect power management, so the already-powered device may be permitted to service the program's request (at 308)"); *see also Araujo*, Fig. 3. A POSITA would have understood that the operations of *Rao/Araujo*'s remote access client are undertaken by *Rao/Araujo*'s central processing unit ("one or more processors").

5. **Claim 2: “The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is interacting with the user in the device user interface foreground when the user of the device is directly interacting with that application or perceiving any benefit from that application.”**

89. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 2. As detailed in Section VII.A.4.d (claim 1[c]), *Rao/Araujo*’s central processing unit classifies whether a given application is running in a device’s foreground. *Rao* discloses that an “application is running in the foreground” when it is “in active use by the user.” *Rao*, [0003]. In my opinion, a POSITA would have understood that a user is directly (and currently) interacting with an application, as claimed, when the application is in the foreground and in active use, as disclosed by *Rao*.

90. *Rao/Araujo*’s central processing unit determines whether a user is currently interacting with a first application (e.g., by engaging in “active use” of the application, as taught by *Rao*) and/or is perceiving any benefit from that application (e.g., by viewing the application as it is running in the foreground and/or perceiving that the application is ready for input by the user, as also taught by *Rao*). *Rao*, [0003].

6. Claim 5: “The wireless end-user device of claim 1, further comprising, when the Internet access request is disallowed, queuing the Internet access request until a power control state change occurs.”

91. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 5. As discussed in Section VII.A.3 (Motivation) and Section VII.A.4.e (claim 1[d(1)]), in an example implementation, *Rao/Araujo*’s remote access client (implemented to include *Araujo*’s request deflector functionality) would disallow certain requests for network access according to *Rao/Araujo*’s augmented policies.

92. *Araujo* discloses multiple possible actions when a request for network access is disallowed (e.g., initially blocked). *Araujo*, [0038], Fig. 4. In my opinion, a POSITA would have understood that the actions (as illustrated in Figure 4) are not exclusive, and additional actions may be taken when a request is disallowed. “[B]lock[ing]” is one possible action, as discussed above. *Araujo*, [0040].

93. According to *Araujo*, a further “such action is to queue a program’s request to be processed at some later time (block 404).” *Araujo*, [0039], Fig. 4. For example, if an application “requests to use a device that is turned off (e.g., a disk, a network card, etc.), and the power management policy does not allow the device to be turned on at the time of the request, the request may be queued so that it can be serviced at some later [time].” *Araujo*, [0039], claim 8. Then, “if at some later time the power management policy allows the device to be turned on, then the

program's request may be dequeued and processed.” *Araujo*, [0039]. In my opinion, a POSITA would have understood that a change in the power state (for example, when the network card has been powered on (e.g., due to a request from a VIP application as taught by *Araujo*)) is an example of the sort of event that may lead to a request being dequeued and processed.

94. Similarly, *Araujo* discloses that another potential response to a disallowed request “is to impose a delay on carrying out the request (block 414).” *Araujo*, [0043], Fig. 4. According to *Araujo*, there are some scenarios in which the device's power management policy “might call for the network card to be turned on at specific intervals.” *Araujo*, [0043]. *Araujo* discloses that “[i]f a request to transmit data over a network is made between these intervals, the request may be delayed until the next time that the network card is scheduled to be powered on, at which time the request could be allowed.” *Araujo*, [0043]. A POSITA would have understood that a “delay” is a time duration experienced by a packet that is stored in a queue. Thus, *Araujo* teaches that the queued packet is assigned a time duration that it must wait in the queue. The result is that, when the Internet access request is denied, *Rao/Araujo* also discloses delaying the request (which is implemented by causing the internet access request that is subject to the claimed “queuing” to wait for a time duration) via *Rao/Araujo*'s request deflector functionality until the device turns back on (such that “a power control state change occurs”).

7. **Claim 7: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to allow the Internet access request when the power control state is the first power control state and the first end-user application is classified as interacting with the user.”**

95. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 7. As explained above in Section VII.A.3 (Motivation), *Rao*’s policies, when augmented to consider the device’s power state, would specify whether an application request is to be allowed or blocked based both on whether the application is currently in the device’s foreground (taught by *Rao*, and similar to *Araujo*’s VIP status) and the device’s power state. Just as *Araujo* describes allowing a VIP application to use the network card to access the network even when the network card is off, *Rao/Araujo*’s augmented policies would allow a foreground application (akin to *Araujo*’s VIP application) to access the network even when *Rao/Araujo*’s network interface is off (“the first power control state”). Thus, when the application requesting network access is in the foreground (“interacting with the user”) and *Rao/Araujo*’s network interface is off (“the first power control state”), *Rao/Araujo*’s remote access client would allow the request for network access.

- 8. Claim 8: “The wireless end-user device of claim 1, further comprising a user interface, wherein the user interface is to inform the user of the device when there are options to set, control, override, or modify service usage controls that affect the processor configuration that determines whether to allow or disallow Internet access requests in the first power control state.”**

96. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 8. User interfaces were well known as of as of January 28, 2009, and commonly used for resource management for coordinating among multiple applications. For example, *Rao* discloses “agent 326 may provide a configuration mechanism such as a user interface, graphical or otherwise, design[ed] and constructed for configuring or specifying the policies 520.” *Rao*, [0183]. Likewise, other references disclosed user interfaces that allowed a user to customize resource management for each application, including in the context of power management. *Aleksic*, [0036] (“A graphic user interface 412 is shown which allows a user to select on a per application basis for example, whether the controller should control features to allow the maximum application runtime that the user desires or a maximum quality of service that the application provides or any suitable level therebetween.”), [0048], [0049]. *Araujo* also disclosed a user interface. *Araujo*, [0041].

97. A POSITA augmenting *Rao*’s policies to include *Araujo*’s power management policy would have likewise incorporated *Araujo*’s user interface

functionality into *Rao*'s user interface to enable full use of *Araujo*'s power management policy. Given *Rao*'s architecture already includes user interface functionality, such a modification would be well within the abilities of a POSITA and a POSITA would have had a reasonable expectation of success.

98. A POSITA would have recognized that a user interface for specifying *Rao/Araujo*'s augmented policies would include "options to set, control, override, or modify service usage controls that affect the processor configuration that determines whether to allow or disallow Internet access requests in the first power control state" because any changes to the augmented policies would affect the processor configuration that determines whether to allow or disallow Internet access requests in the first power control state.

99. Moreover, a POSITA would have been motivated and found it obvious to implement the user interface to "inform the user of the device when there are options" related to *Araujo*'s policy, because *Araujo* discloses a user interface that prompts the user for input on its policies. According to *Araujo*, when an application cannot access the Internet (due to the device's power state), central processing unit may "[p]rompt[] a user (or other person) for input." *Araujo*, [0041]. "[I]f a program requests to send data over a network card that has been turned off, a dialog box may be generated on a display that informs the person in front of the display that the card is off." *Araujo*, [0041]. A "dialog box" is a user interface, as

claimed, because it allows interaction with a user. *Araujo* further discloses that this dialog box “may also solicit input from the person (e.g., ‘Would you like to proceed anyway?’), or may simply inform the person that the card is off and suggest that the user try again later.” *Araujo*, [0041]. A POSITA would have understood that soliciting user input, including asking whether the user would like to override the prevailing policy, would entail informing the user of an option to override the service usage control, as claimed.

9. Claim 9

- a. 9[a]: “The wireless end-user device of claim 1, further comprising: a wireless local area network (WLAN) modem to communicate data for Internet service activities between the device and at least one WLAN, when configured for and connected to the at least one WLAN,”**

100. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 9[a]. As discussed in Section VII.A.4.b (claim 1[a]), *Rao/Araujo*’s device includes a wireless wide area network (WWAN) modem to communicate data for Internet service activities. The device “can connect to one or more networks 104.” *Rao*, [0095]; *see also Rao*, [0125]. Thus, in addition to a WWAN modem (identified for claim 1), a POSITA would have understood that *Rao/Araujo*’s device would include other modems, including a WLAN modem (i.e., a second modem), to connect to other networks.

101. *Rao*'s device, like other devices known at the time, "can connect to one or more networks 104." *Rao*, [0095]; *see also Rao*, [0125]. *Rao* discloses that its network interface 118 also enables connectivity to a "Local Area Network (LAN) . . . through a variety of connections," such as "wireless connections." *Rao*, [0125]; *see also Rao*, [0095]. Network interface 118 includes a "modem . . . suitable for interfacing the computing device 102 to any type of network capable of communication," including WLAN. *Rao*, [0125].

102. Given *Rao* discloses that its device can connect to multiple networks and a network interface, including a second modem, that enables connectivity to a WLAN, *Rao/Araujo* also discloses a "[WLAN] modem," as claimed.

- b. 9[b]: "wherein the one or more processors are further configured to allow the Internet access request when the WLAN modem provides data communication for the Internet access request and the power control state is the first power control state."**

103. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 9[b]. As discussed in Section VII.A.4 (claim 1), in *Rao/Araujo*, the central processing unit controls the remote access client to handle the Internet access requests according to *Rao*'s policies, augmented to include *Araujo*'s power management policy. (From *Araujo*, a POSITA would have been motivated to implement *Rao/Araujo*'s augmented policies to specify that, when the WWAN network interface card is off ("first power control state") and an application's

request is disallowed for the WWAN network interface, the request may be referred to *Rao*'s WLAN network interface card for transmission.

104. As explained in Section VII.A.6 (claim 5), *Araujo* discloses multiple possible actions when a request for network access is disallowed. *Araujo*, [0038], Fig. 4. One possible action is referring the request to another device. *Araujo*, [0042]. According to *Araujo*, “[t]here may be a table, or some other data structure, that lists various devices that could service a request if a particular device is not powered on,” and “[s]uch a table could be consulted.” *Araujo*, [0042]. For example, *Araujo* explains, some mobile devices (like *Rao*'s device) may have “plural network cards,” and if a request to use a first network card (for example, a WWAN modem) is denied because the card is turned off, then the request may be sent to a second network card (e.g., a WLAN modem). *Araujo*, [0042].

105. In evaluating an Internet access request, a POSITA would have understood that communicating through a WLAN modem may generally consume less power than communicating through a WWAN modem and would be preferred for Internet access in power-conservation modes as part of a device's power management policy. *Aerrabottu*, 1:62-66 (“In one embodiment, the device switches from operation on a cellular communication protocol while in a call to operation on a wireless local area network protocol in the event that the battery power level drops to the predetermined threshold.”); *Cole*, [0011] (“Another issue

related to using different connections is the power consumption associated with each connection type. While a particular connection may have a desirably high bandwidth, the connection modality may require significant processing power from the mobile device, thereby greatly shortening battery life.”), [0088] (explaining that a “WWAN modem . . . consumes significant battery power”).

106. The power consumption is due to the fact that a distance of transmission may be less for a WLAN modem (e.g., transmitting to a nearby access point with a home) than for a WWAN modem (e.g., transmitting to an access point on a cell tower). As the distance for the transmission increases, the amount of power necessary for the transmission also increases due to the associated path loss. Thus, in *Rao/Araujo*, a background application’s Internet access request via WWAN may be denied if the WWAN network card is powered off, but the power management policy may allow that same request via WLAN.

10. Claim 10: “The wireless end-user device of claim 1, wherein the power control state is a power state of the device.”

107. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 10. As explained in Section VII.A.4 (claim 1), a POSITA would have been motivated and found it obvious from *Araujo* to augment *Rao/Araujo*’s policies to allow or disallow an application’s request for network access (“Internet access request”) (1) based on whether an application is in the foreground (“interacting

with a user”) and (2) based on a power control state, such as the power state of a network card, which is a power state of the device.

108. To the extent “power state of the device” refers to the power state of the entire mobile device (rather than a single component), a POSITA would have understood from *Araujo* that the power management policies may implicate the entire device and/or a component of the device that employs power. As *Araujo* explains, “[e]lectronic devices (e.g., computers, wireless telephones, audio/video equipment, etc.) normally employ some type of power management framework.” *Araujo*, [0014]; *see also Araujo*, [0003] (“[p]ower management for a machine”). For example, *Araujo* describes a “lower-power-consumption mode” that applies to various components of the device (“a first power control state”). *Araujo*, [0014]. A POSITA would have understood that a device would likewise include a non-lower-power consumption mode (“at least one other power control state”).

109. A POSITA, when reviewing *Araujo*, would also read *Araujo*’s disclosure in the context of the prior art that formed the background knowledge of an ordinary artisan. This would include, for example, *Cole*, which teaches that a “power state”—the same phrase used in *Araujo*—applies to an entire mobile device (rather than just a component, such as a modem). Like *Araujo*, *Cole* teaches “chang[ing] the configuration of the communication functions in direct response to the power state of the mobile device 110.” *Cole*, [0090]. A POSITA would have

been motivated and found it obvious to use the power state of the entire mobile device, rather than the power state of the network card, because this would account for power consumption outside of the network card. A POSITA would have understood that this would advantageously enable *Rao/Scahill* to allow or disallow an application's request for network access based on the power state of the entire mobile device, which may not be the same power state of the network card and may yield different a different decision regarding whether to allow or disallow an application's request for network access. *See* Section VII.A.3 (Motivation).

110. As a result, a POSITA would have understood that in *Rao*'s policies, augmented to include *Araujo*'s power management policy, a power state (a "power control state") includes a power state of the device.

11. Claim 11: "The wireless end-user device of claim 1, wherein the first power control state is a power save state of the device."

111. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 11. As discussed in Section VII.A.10 (claim 10), *Rao/Araujo*'s policies allow or disallow an application's request for network access based, in part, on a power control state, such as the power state of a network card, which is a power state of the device. As discussed in Section VII.A.4.f (claim 1[d(2)]), the "first power control state" is when *Rao/Araujo*'s network interface is powered off. A POSITA would have understood that being powered off is a power save state of

Rao/Araujo's device at least because powering off the network card, even if it is temporary, saves power in the wireless end-user device by enabling the device to consume less power.

112. In addition, *Araujo* discloses a “lower-power-consumption mode” in which a device (i.e., the entire mobile device) turns down “brightness” and “processor speed when its remaining battery power drops below a certain threshold.” *Araujo*, [0014]. A POSITA would have understood that this power state is a “power save state” because the amount of power consumed is less than during standard operations. *Cole*, [0099] (describing a “reduced power state” or “low power mode”). In my opinion, A POSITA would have understood that, in the context of *Rao/Araujo*'s augmented policies, a power save state (“the first power control state”) is a power state of the device, as claimed.

12. Claim 12: “The wireless end-user device of claim 1, wherein the power control state is a power state of the WWAN modem.”

113. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 12. As discussed in Section VII.A.4.b (claim 1[a]), *Rao/Araujo*'s device includes a WWAN modem, as claimed. *Araujo* discloses that “power may be controlled for devices,” such as “disks, network adapters, processors, monitors, etc.” *Araujo*, [0018], [0019]. In my opinion, a POSITA would have understood that a WWAN modem is an example of a network adapter. A POSITA would have

further understood that a WWAN modem may be “on or off” or “an intermediate power usage between fully-on and fully-off.” *Araujo*, [0019]. Each of these is the power state of the WWAN modem. *Araujo* details how its power management policy can be driven by the “power state” (the “power control state”) of the network adapter (that is, the WWAN modem). *Araujo*, [0004], [0015], [0018]-[0022]. Thus, in the context of *Rao/Araujo*’s architecture, the power control state used to make the policy decisions is the WWAN modem’s power state.

13. Claim 13: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to, based on a device usage state, dynamically change a determination of whether to allow or disallow the Internet access request in the first power control state when the first end-user application is classified as not interacting with the user.”

114. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 13. Applying *Rao/Araujo*’s augmented policies, *Rao/Araujo*’s remote access client would allow or disallow an application’s request (1) based on whether the application is in the foreground and (2) based on a power control state. *Rao* discloses that, in addition to making a determination based on whether the application is in the foreground, *Rao*’s remote access client “may determine any other characteristics or statistics of the application 338a-338n [for the purpose of assessing priority] such as ... total execution time[] and/or frequency of use.” *Rao*,

[0188]-[0189]. A POSITA would have understood these characteristics reflect a “device usage state” because they reflect how the device is being used.

115. Indeed, a POSITA would have understood that assessing priority based on other characteristics or statistics of the application 338a-338n such as total execution time and/or frequency of use would advantageously enable *Rao/Araujo*’s remote access client to be more responsive to how a user engages with an application on a given device (for examples, see the discussion at the end of this section). This would, in turn, improve the user’s experience when using the device.

116. A POSITA would have had a reasonable expectation of success in modifying *Rao/Araujo*’s policies to adjust how priority is assessed. *Rao* discloses that a POSITA “will recognize and appreciate the multitude of ways to define client-side application priorities.” *Rao*, [0182]. *Rao* explains that policies: (1) “may be provided by or downloaded via the gateway 340,” (2) “may comprise any type and format of syntax and/or language for specifying a policy,” and (3) “may be provided via any type and/or form of medium.” *Rao*, [0183]. These policies also “may be configurable by a user by any suitable means and/or mechanism.” *Rao*, [0183]. The result is that the modification would have been well within the abilities of an ordinary artisan to implement.

117. From *Rao/Araujo*, a POSITA would have understood that, based on these “device usage state” characteristics, *Rao*’s remote access client may dynamically change its determination of whether to allow or disallow an application’s request (1) based on whether the application is in the foreground and (2) based on a power control state. For instance, when the application requesting network access is in the background (“not interacting with the user”) and the network interface is powered off (“first power control state”), *Rao/Araujo*’s remote access client may determine that a request from the application for network access should be disallowed in light of how the device is being used (e.g., a total execution time or frequency of use does not exceed a predetermined threshold). But *Rao/Araujo*’s remote access client may dynamically change this determination (e.g., to allow the application’s request) if how the device is being used changes (e.g., the total execution time or frequency of use now exceeds a predetermined threshold). Section VII.A.3 (Motivation); Section VII.A.4.e (claim 1[d(1)]).) In my opinion, a POSITA would have understood that the operations of *Rao/Araujo*’s remote access client are undertaken by *Rao/Araujo*’s CPU (“one or more processors”)

- 14. Claim 14: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to, based on power control state changes for the WWAN modem, dynamically change a determination of whether to allow or disallow the Internet access request in the first power control state when the first end-user application is classified as not interacting with the user.”**

118. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 14. As discussed in Section VII.A.4 (claim 1), *Rao/Araujo*’s central processing unit controls the remote access client to handle Internet access requests according to *Rao*’s policies, augmented to include *Araujo*’s power management policy. From *Araujo*, a POSITA would have been motivated and found it obvious to implement *Rao/Araujo*’s augmented policies to specify that, where a first network card such as a WLAN modem is turned off, an Internet access request from a background application (“not interacting with the user”) may be allowed or disallowed based on whether a second, WWAN modem is turned on.

119. As explained in Section VII.A.9 (claim 9), *Araujo* discloses multiple possible actions when a request for network access is disallowed. *Araujo*, [0038], Fig. 4. One possible action is referring the request to another device. *Araujo*, [0042], Fig. 4. For example, *Araujo* explains, some mobile devices (like *Rao*’s device) may have “plural network cards,” and if a request to use a first network card (e.g., a WLAN modem) is denied because the card is turned off, then the request may be sent to a second modem (e.g., a WWAN modem). *Araujo*, [0042].

If the second, WWAN modem is on, the request would be allowed, because the application may access the Internet.

120. Thus, from *Rao/Araujo*, a POSITA would have understood that, based on power control state changes for the WWAN modem, *Rao*'s remote access client may dynamically change its determination of whether to allow or disallow an application's request (1) based on whether the application is in the foreground and (2) based on a power control state. For instance, when the application requesting network access is in the background ("not interacting with the user") and the WLAN modem is powered off ("first power control state"), *Rao/Araujo*'s remote access client may determine that a request from the application for network access should be disallowed if the second, WWAN modem is also powered off. But *Rao/Araujo*'s remote access client may dynamically change this determination (e.g., to allow the application's request via the WWAN modem) if the WWAN modem is powered on.

15. Claim 15: "The wireless end-user device of claim 1, wherein, when the Internet access request is disallowed, the one or more processors are configured to prevent the Internet access request from causing a change to a power state of the modem."

121. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 15. As discussed in Section VII.A.4.f (claim 1[d(2)]), when *Rao/Araujo*'s network interface is powered off and the application requesting network access is

in the background, *Rao/Araujo*'s remote access client would disallow (e.g., "block," *Araujo*, [0040]) the request for network access. *Araujo* explains that such disallowing involves preventing the request for network access from causing a change to the network card.

122. According to *Araujo*, "if a particular request does not justify powering a device under the power management policy, the request deflector may prevent the request from being passed to the device driver." *Araujo*, [0005]; *see also Araujo*, Fig. 2 (208, 212), Fig. 4. The result is that "the change of a device's power state [(the modem's power state)] may be physically prevented or deferred." *Araujo*, [0050]. Indeed, as *Araujo* details to explain its power management policy, "state transitions" between power states entail "power usage," and preventing such a state change preserves power. *Araujo*, [0029]. Accordingly, in *Rao/Araujo*, when a request is disallowed, *Rao/Araujo*'s remote access client (controlled by the CPU) would prevent the request from causing a change to a power state of the network interface.

16. Claim 16: "The wireless end-user device of claim 1, wherein the one or more processors are configured to, when the Internet access request is disallowed, instruct the first end-user application to transition to a different state."

123. In my opinion, the combination of *Rao* and *Araujo* renders obvious claim 16. *Araujo* discloses that "if the power state change is not allowable under

the policy,” the device can “block the request” and “notif[y] [the program] that the request will not be processed.” *Araujo*, [0038], [0040]. For example, “a program could be notified that a request to use a device is not being put through to the device because the request would be contrary to power-usage policy.” *Araujo*, [0017]. *Araujo* discloses that “[t]he program may also be notified that the reason for the request is that the device is off. For example, the program may receive an error code or result code indicating that the request calls for a device that is not turned on.” *Araujo*, [0040].

124. In my opinion, given *Araujo*’s disclosure and what would otherwise been known to a POSITA, a POSITA would have understood or found it obvious that, when a request is blocked, *Rao/Araujo*’s remote access client (controlled by *Rao/Araujo*’s central processing unit) instructs an application to transition to a different state.

125. Indeed, *Araujo* details an example where a notification would be beneficial. According to *Araujo*, “if an application with a low status is requesting use [of] a network card that is presently turned off, the system might decide . . . to wait until a high-status daemon makes a request to use the network card.” *Araujo*, [0033]. According to *Araujo*, “[i]f such a request is expected (e.g., based on past patterns of usage) within, say, the next two seconds, then the low-status application’s request may be . . . delayed in view of the anticipated high-status

daemon's request." *Araujo*, [0033]. From *Araujo*, a POSITA would have understood that, when an application request is disallowed, *Rao/Araujo*'s remote access client would instruct the application to move into an idle state and such an instruction would prevent the low-status application from making repeated requests or otherwise unnecessarily consuming device resources. This idle state would only be temporary because the idled application would have its request allowed once the high-status daemon requests Internet access.

126. As a result, a POSITA would have been motivated by *Araujo* to implement *Rao/Araujo*'s central processing unit to control *Rao/Araujo*'s remote access client in a manner that instructs an application to transition to an idle state where that application's Internet access request has been disallowed. This modification would have been well within the abilities of a POSITA, and would merely require enabling *Rao/Araujo*'s remote access client to transmit a code (e.g., error code) that is received by the requesting application (as taught by *Araujo*) indicating that the application's Internet access request calls for a device that is not turned on and instructing the application to enter a standby mode. As a result, a POSITA would have had a reasonable expectation of success in implementing this modification.

B. Ground 2: Claim 3 Is Obvious Over *Rao* in View of *Araujo* and *Singh*

1. Overview of *Singh*

127. *Singh* is directed to “displaying windows on a graphical user interface based on relative priorities.” *Singh*, Abstract. Like the ’076 patent, *Singh* recognizes that “[c]omputers and other electronic devices are typically operable to execute multiple applications concurrently.” *Singh*, 1:21-22. In my opinion, because *Singh* is directed to the same field of endeavor as the ’076 patent (application management in devices that can run multiple applications), *Singh* is analogous art to the ’076 patent. *Singh* is directed to addressing issues raised by a device’s ability to execute multiple applications concurrently and is analogous art to the ’076 patent.

128. According to *Singh*, an issue with the “simultaneous display of multiple windows” on the GUI is that “[t]he display area of the graphical user interface available to [the] windows is limited.” *Singh*, 1:21-28, 1:37-40. Thus, “it is desirable to position, size, and/or configure the windows on the GUI in order to maximize the benefit of the display area of the GUI for the user.” *Singh*, 1:40-42. *Singh* thus teaches using the “relative priorities” of windows for arranging the windows on the GUI. *Singh*, Abstract. “A higher priority window may cover a lower priority portion of a lower priority window to provide for a simultaneous

viewing of the higher priority window and a higher priority portion of the lower priority window.” *Singh*, Abstract.

2. Rationale for Modifying *Rao* in View of *Araujo* and *Singh*

129. In my opinion, it would have been obvious to a POSITA to modify *Rao*’s remote access client to include *Singh*’s windows management function.

130. *Rao* describes a scenario where an “application is running in the foreground and currently in active use by the user, [and] a network packet generated or received for an application running in the background may be processed ahead of a network packet generated or received for the application running in the foreground.” *Rao*, [0003]. According to *Rao*, this scenario is not desirable because “network packets for [background] applications may be processed ahead of the real-time network packets of [a foreground] VoIP telephone call thereby increasing the latency and reducing the quality of the voice communication.” *Rao*, [0003]. To prevent network activity from background applications from interfering with network activity from foreground applications, *Rao* teaches policies that “define prioritization based on whether an application is running in the foreground or the background of the client” and “the network packets are communicated according to the determined priority.” *Rao*, [0182], [0184].

131. A POSITA would have recognized that, in cases where multiple windows are displayed simultaneously (e.g., with the highest-priority window having “focus” and covering portions of lower-priority windows), it is challenging for *Rao*’s device to prioritize based on determining which application is running in the foreground. In such cases, multiple windows are in the foreground but partially displayed and one window has the focus (and thus should have the highest priority). For example, in the window with focus, a user is engaged in a video call that requires real-time communication. Another foreground window displays the user’s email, but the user interacts with email only occasionally (e.g., during pauses in the video call). While both windows are in the foreground, network traffic related to the video call should receive priority when the video call has focus. Thus, in my opinion, a POSITA implementing *Rao/Araujo* would have been motivated to look to other references to improve the prioritization of Internet access requests when multiple applications are displayed in a user interface of the device.

132. In my opinion, a POSITA would have recognized “the multitude of ways to define client-side application priorities” and further recognized that *Singh*’s window priorities represent such an example. *Rao*, [0182]. *Rao/Araujo*’s policies would thus be further augmented to permit the remote access client to classify whether an application is interacting with the user in the device user

interface foreground using *Singh*'s window prioritization. *Rao*, [0182], [0183]; *Singh*, 4:52-55, Abstract.

133. A POSITA would have had a reasonable expectation of success in doing so, which would have required only routine programming knowledge well within the skill of a POSITA before the '076 patent's priority date. Indeed, the change would have amounted to nothing more than the use of a known technique to improve similar mobile devices in a similar way and combining prior art elements according to known methods to yield the predictable results described above. *Rao/Araujo/Singh*'s central processing unit would be configured to perform the operations discussed herein.

3. Claim 3: "The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is interacting with the user in the device user interface foreground based on a state of user interface priority for the application."

134. In my opinion, the combination of *Rao*, *Araujo*, and *Singh* renders obvious claim 3. As discussed in Section VII.A.4.d (claim 1[c]), *Rao/Araujo*'s central processing unit classifies "whether an application is running in the foreground or the background." *Rao*, [0182]. *Rao/Araujo*'s central processing unit would further use *Singh*'s window management function to improve its classification of foreground applications where multiple windows are displayed simultaneously. Section VII.B.2 (Motivation).

135. *Singh* teaches a “window management function [that] determine[s] a relationship of windows with one another.” *Singh*, 9:19-21. According to *Singh*, “when a user interacts with a window, the window becomes the highest priority window.” *Singh*, 9:33-38. A POSITA would have known that users interact with foreground windows, and in *Singh*, users only interact with the highest priority window (because any interaction causes the window to be assigned the highest priority). Therefore, a POSITA would have understood that *Singh*’s window management function—integrated into *Rao/Araujo*’s remote access client—uses a state of user interface priority (e.g., highest priority) to classify that the application is interacting with the user in the user interface foreground.

C. Ground 3: Claim 4 Is Obvious Over *Rao* in View of *Araujo* and *Freund*

1. Overview of *Freund*

136. *Freund* relates to using “access rules” for efficient use of computing resources and controlling network access. *Freund*, 3:60-67. In my opinion, because *Freund* is directed to the same field of endeavor as the ’076 patent (application management in devices that can run multiple applications) and same problem as the ’076 patent (optimizing use of network communications on an application-by-application basis), *Freund* is analogous art to the ’076 patent.

137. *Freund* recognized that, as May 6, 1997, access rules and “monitoring based on an individual client application (or versions thereof) [were] not supported [in the prior art].” *Freund*, 2:63-3:1. As a result, *Freund* teaches application-specific access rules to determine if an application or process “should have access to the Internet and what kind of access (i.e., protocols, Internet addresses, time limitations, and the like) is permissible.” *Freund*, 4:45-63. According to *Freund*, “[a]ccess rules can include criteria such as total time a user can be connected to the Internet,” “time a user can interactively use the Internet,” and “a list of applications or application versions that a user can or cannot use in order to access the Internet.” *Freund*, Abstract, 13:2-13.

138. To apply the rules, *Freund* gives “[s]pecific consideration[] for detecting interactive or ‘active’ use by the user versus background or ‘passive’ use by certain applications.” *Freund*, 8:63-67. *Freund* teaches the use of “a preferred user interface or ‘wizard’ dialogs for configuring rules.” *Freund*, 24:16-17. This interface displays “rules governing operation of the Internet access monitor” and allows a user to create, edit, or delete rules. *Freund*, 24:17-24.

2. Rationale for Modifying *Rao* in View of *Araujo* and *Freund*

139. In my opinion, it would have been obvious to a POSITA to modify *Rao/Araujo*’s remote access client to include *Freund*’s method for classifying

whether an application is running in the background, to improve the operation of *Rao*'s remote access client.

140. *Rao* contrasts “an application [that] is running in the foreground and currently in active use by the user” with “an application running in the background.” *Rao*, [0003]. *Rao*, however, does not expressly describe how to determine whether an application is running in the background. As a result, a POSITA implementing *Rao*'s teachings would have looked to similar references, including *Freund*, to understand how to classify an application as running in the background.

141. *Freund* teaches network access policies that give consideration to whether an application is operating in a device's foreground or background. *E.g.*, *Freund*, 10:5-43. In particular, *Freund* teaches that “‘background’ use” should be determined, at least in part, by evaluating “when an application accesses the Internet without user interaction.” *Freund*, 3:12-16.

142. *Freund* discloses a number of benefits to its classification method. According to *Freund*, its method “for detecting interactive or ‘active’ use by the user versus background or ‘passive’ use by certain applications ... prevents the system from mistakenly blocking access of a user based on excessive use, when in fact a background application is instead responsible for the activity.” *Freund*, 8:63-9:3. Indeed, according to *Freund*, “[c]onventional technologies cannot distinguish

between this ‘active’ versus ‘passive’ access and therefore cannot present an accurate picture of the actual time the user has spent on on-line activities.” *Freund*, 10:30-33.

143. The resulting combination would further *Rao*’s goal of prioritizing the transmission of network packets from foreground applications over background applications to reduce the detrimental latency impacts that result from a failure to prioritize. *E.g.*, *Rao*, [0003]; *see also* Section VII.A.1 (Overview of *Rao*). *Freund*’s method allows for a more accurate distinction between background and foreground applications, including certain “bandwidth-intensive [background] applications” (*Freund*, 2:12-14) that should be de-prioritized relative to certain applications operating in the foreground (e.g., a VoIP application to provide a telephone call). As a result, a POSITA would have been motivated to integrate *Freund*’s method for classifying whether an application is running in the background into *Rao/Araujo*’s remote access client.

144. A POSITA would have also had a reasonable expectation of success in modifying *Rao/Araujo*’s architecture in view of *Freund*’s classification method. *Rao/Araujo*’s remote access client already evaluates whether a device is in active use (e.g., to classify whether an application is running in the foreground), and *Freund*’s classification method adds additional functionality for classifying whether an application is running in the background. Such a modification would

have been well within the abilities of an ordinary artisan to implement and would not have conflicted with the classification of applications performed by *Rao/Araujo*'s remote access client.

- 3. Claim 4: “The wireless end-user device of claim 1, wherein the one or more processors are configured to classify that the first end-user application is not interacting with the user in the device user interface foreground when the application is providing or utilizing a background data service.”**

145. In my opinion, the combination of *Rao*, *Araujo*, and *Freund* renders obvious claim 4. As discussed in Section VII.A.4.d (claim 1[c]), *Rao/Araujo*'s central processing unit, via the remote access client, classifies whether a given application is running in a device's foreground or the background. *Rao/Araujo*'s remote access client would implement *Freund*'s method for determining when an application is running in the background. Section VII.C.2 (Motivation).

146. *Freund* defines “‘background’ use” as “when an application accesses the Internet without user interaction.” *Freund*, 3:12-16. *Freund* further explains that background use “occurs when an application or process executing in the background (i.e., does not have ‘focus’) accesses the Internet, such as a mail client which intermittently polls an Internet-based mail server.” *Freund*, 10:10-24; *see also Freund*, 10:40-43 (“A given application itself can be examined for determining whether it is ‘active’ by determining whether the application receives ‘focus’ and/or receives user input (e.g., mouse clicks or key strokes).”). Since

Freund defines background use as occurring when an application or process executing in the background accesses the Internet (when an application uses background data services), then by *Freund*'s definition, an application would be classified as not interacting with a user when the application uses background data services.

147. In my opinion, a POSITA would have thus understood that when *Rao/Araujo*'s remote access client (controlled by its central processing unit) classifies a first end-user application as running in the background, as disclosed by *Freund*, the classification would be performed based on determining that the application is providing or utilizing a background data service (e.g., by generating network packets in the background).

D. Ground 4: Claim 6 Is Obvious Over *Rao* in View of *Araujo* and *Aleksic*

1. Overview of *Aleksic*

148. Like the '076 patent, *Aleksic* teaches monitoring and prioritizing usage of device resources on an application-by-application basis. (E.g., *Aleksic*, [0022]-[0025], [0033].) In my opinion, because *Aleksic* is directed to the same field of endeavor as the '076 patent (application management in devices that can run multiple applications), *Aleksic* is analogous art to the '076 patent. *Aleksic* is analogous art to the '076 patent.

149. *Aleksic* recognizes that when a portable device has limited device resources, “users are typically most concerned with ... the ability to use the portable device for at least an emergency call” (i.e., a voice application) or other critical applications. *Aleksic*, [0002], [0050]. To address this concern, *Aleksic* teaches a controller that allows a user to “prioritize which functional processing capability features of the portable device [e.g., which applications] take priority over others.” *Aleksic*, [0022]. *Aleksic* teaches “a graphic user interface 1100 that provides selectable priority information shown as ‘critical’ or ‘noncritical’ for various applications (e.g., functional processing capability features) of the device.” *Aleksic*, [0050]. Thus, according to *Aleksic*, if the available device resources are not sufficient “in comparison to an emergency call threshold, the device will enter the emergency call mode and effectively shut down all applications to reserve remaining power for voice communications (e.g., emergency calls)” or other critical applications, leaving those applications to run in either the background or foreground, as necessary. *Aleksic*, [0043].

2. Rationale for Modifying *Rao* in View of *Araujo* and *Aleksic*

150. In my opinion, it would have been obvious to a POSITA to modify *Rao/Araujo*’s remote access client to include *Aleksic*’s prioritization of critical or emergency applications and further augment *Rao/Araujo*’s policies.

151. The prior art recognized “a need ... for an improved power management operation in portable devices.” *Aleksic*, [0004]. As described in Section VII.A.3, a POSITA would have modified *Rao* in view of *Araujo* to further improve *Rao*’s application-aware power management. *See Rao*, [0179]-[0195]. In implementing this modification, a POSITA would have recognized opportunities to provide further prioritization for applications designated by a device user as “critical” or “emergency” so that these applications would be available, even when device resources are limited. *Aleksic*, [0043], [0048]. This would have motivated a POSITA to seek solutions beyond *Rao* and *Araujo* to achieve such additional efficiencies and further augment *Rao*’s policies.

152. *Aleksic*’s teachings allow a user to “prioritize which functional processing capability features of the portable device [e.g., which applications] take priority over others.” *Aleksic*, [0022], [0050]. This includes emergency or other critical applications, wherein a user would require prioritizing Internet access at all times, including for both incoming and outgoing network communications. *Aleksic*’s prioritization would have enabled *Rao/Araujo*’s prioritization architecture to further account—at a policy level—for “critical” and “emergency” applications. *Aleksic*, [0043], [0048].

153. A POSITA would have had a reasonable expectation of success in modifying *Rao/Araujo*’s policies in view of *Aleksic*. *Rao* discloses that a POSITA

“will recognize and appreciate the multitude of ways to define client-side application priorities.” *Rao*, [0182]. *Rao* explains that policies: (1) “may be provided by or downloaded via the gateway 340,” (2) “may comprise any type and format of syntax and/or language for specifying a policy,” and (3) “may be provided via any type and/or form of medium.” *Rao*, [0183]. These policies also “may be configurable by a user by any suitable means and/or mechanism.” *Rao*, [0183]. The result is that the modification would have been well within the abilities of an ordinary artisan to implement.

3. Claim 6: “The wireless end-user device of claim 1, wherein the one or more processors are further configured to, for at least a second application or service, allow Internet access requests in the first power control state without regard to whether the second application or service is classified as interacting with the user.”

154. In my opinion, the combination of *Rao*, *Araujo*, and *Aleksic* renders obvious claim 6. As discussed in Section VII.D.2 (Motivation), *Rao/Araujo/Aleksic*’s remote access client’s augmented policies allow for prioritizing certain critical applications regardless of the power control state of *Rao*’s network interface and whether the application is running in the background (“not interacting with the user”). Thus, when *Rao/Araujo/Aleksic*’s network interface is powered off and a critical or emergency application requests network access, *Rao/Araujo/Aleksic*’s remote access client would allow an Internet access

request regardless of whether the critical or emergency application is in the background (“not interacting with the user”). Section VII.A.3 (Motivation); Section VII.A.4.e (claim 1[d(1)]).) A POSITA would have understood that the operations of *Rao/Araujo/Aleksic*’s remote access client are undertaken by its central processing unit (“one or more processors”).

E. Ground 5: Claims 16 Is Obvious Over *Rao* in View of *Araujo* and *Montemurro*

1. Overview of *Montemurro*

155. Like the ’076 patent, *Montemurro* is directed to application-aware optimization of network transmissions from a mobile device Just as *Rao* and *Araujo* teach application-aware prioritization of network access, *Montemurro* discloses application-aware “policy-based routing of communications among two or more modes of wireless communication.” *Montemurro*, [0001]. In my opinion, because *Montemurro* is directed to the same field of endeavor as the ’076 patent (application management in devices that can run multiple applications) and the same problem as the ’076 patent (optimizing use of network communications on an application-by-application basis), *Montemurro* is analogous art to the ’076 patent.

156. Devices with such connectivity capabilities are “dual or multi-mode devices,” with “radio access technologies that provide access to” multiple network types, including WLAN and WWAN. *Montemurro*, [0003]. But, according to

Montemurro, “[t]here are costs associated with application access from these different network[s],” and thus, “[i]t is therefore desirable to have a mechanism that seeks to optimize communications” over the different network types.

Montemurro, [0004].

157. To optimize communications, *Montemurro* discloses “a rules engine 206A (sometimes called an inference engine) as a policy-based mechanism using rules to configure TCP/IP network stack 206B and connection table 206C to coordinate communications for applications 204 using communications interfaces (also referenced as network interfaces) 208.” *Montemurro*, [0027]. *Montemurro* discloses that this rules engine “configures the communication operations with a set of rules/policies that could include various factors such as radio access technology (e.g. for high/low bandwidth properties), cost, presence, time of day, location (e.g. geo-based policies, network roaming), destination IP address, application type, and Quality of Service (QoS) requirements, among others.” *Montemurro*, [0027].

158. Rules engine 206A “configures (i.e. modifies, periodically in response to changes of state and the evaluation of its rules) connection table 207C and routing table 210 to optimize the flow of communications over multiple communications modes (e.g., interfaces 208 and respective networks 104 and 106).” *Montemurro*, [0034]. On a state change, e.g., “time of day,”

“connecting/disconnecting of device 102 with a specific network 104 and 106,” rules engine 206A executes and modifies “routing table 210 to ensure that data goes out to the most appropriate network (via respective network interface 208).” *Montemurro*, [0035]. Rules engine 206A further “determine[s] which interface would be best to service a particular application.” *Montemurro*, [0029].

2. Rationale for Modifying *Rao* in View of *Araujo* and *Montemurro*

159. As described in Section VII.A.3 (Motivation), a POSITA would have modified *Rao* in view of *Araujo* to further improve *Rao*’s architecture to include application-aware power management. *See Rao*, [0179]-[0195]. In implementing this modification, a POSITA would have recognized opportunities to provide further optimization in scenarios where *Rao/Araujo*’s device has access to multiple networks. *Rao*, [0095]; *see also Montemurro*, [0003], [0004]. This would have motivated a POSITA to seek solutions beyond *Rao* and *Araujo* to achieve such additional efficiencies.

160. According to *Montemurro*, devices with multiple radio access technologies can access more than one network, including at least “wireless local area networks (WLAN)” and “wireless wide area networks (WWAN) (e.g. cellular technologies[]).” *Montemurro*, [0003]. But “[t]here are costs associated with application access from these different networks.” *Montemurro*, [0004]. And

“[t]here are also quality considerations such as the speed at which content can be delivered, or in the case of streaming media, the quality of service at which the content is delivered.” *Montemurro*, [0004]. “It is therefore desirable to have a mechanism that seeks to optimize communications for multi-mode capable devices, that is, that seeks to improve communications for multi-mode capable devices.” *Montemurro*, [0004].

161. Given *Rao*’s disclosure of multi-mode devices that “connect to one or more networks” (*Rao*, [0095]), a POSITA would have been motivated to improve *Rao/Araujo*’s architecture to include *Montemurro*’s policy-based data routing for multimode operations. In my opinion, a POSITA would have understood that incorporating these teachings into *Rao/Araujo* would have resulted in the ability to “offer different usage models depending on the mode of wireless operation selected.” *Montemurro*, [0003]. *Montemurro*’s policies would have enabled *Rao*’s prioritization architecture to further account—at a policy level—for “radio access technology (e.g. for high/low bandwidth properties), cost, presence, time of day, location (e.g. geo-based policies, network roaming), destination IP address, application type, and Quality of Service (QoS) requirements.” *Montemurro*, [0029].

162. A POSITA would have had a reasonable expectation of success in modifying *Rao/Araujo*’s policies in view of *Montemurro*’s policies. *Rao* discloses

that a POSITA “will recognize and appreciate the multitude of ways to define client-side application priorities.” *Rao*, [0182]. *Rao* explains that policies: (1) “may be provided by or downloaded via the gateway 340,” (2) “may comprise any type and format of syntax and/or language for specifying a policy,” and (3) “may be provided via any type and/or form of medium.” *Rao*, [0183]. These policies also “may be configurable by a user by any suitable means and/or mechanism.” *Rao*, [0183]. The result is that the modification would have been well within the abilities of an ordinary artisan to implement.

3. Claim 16: “The wireless end-user device of claim 1, wherein the one or more processors are configured to, when the Internet access request is disallowed, instruct the first end-user application to transition to a different state.”

163. In my opinion, the combination of *Rao*, *Araujo*, and *Montemurro* renders obvious claim 16. *Montemurro* discloses that “applications can be developed and operated to provide different usage models that may vary depending on the mode of operation that is available at runtime.” *Montemurro*, [0024]; *see also Montemurro*, [0003]. The rules engine “configures the communication operations with a set of rules/policies that could include various factors such as radio access technology (e.g. for high/low bandwidth properties), cost, presence, time of day, location (e.g. geo-based policies, network roaming), destination IP address, application type, and Quality of Service (QoS) requirements, among

others.” *Montemurro*, [0029]. As discussed in Section VII.A.4.c (claim 1[b]), in *Rao/Araujo*’s remote access client, policies are administered via central processing unit 102.

164. *Montemurro* teaches that a processor can, when an Internet access request is disallowed (e.g., because the available network for the network request is not allowable under the policy), instruct the application to transition to a different state by connecting to a different network and/or altering its behavior to account for the different network. Thus, for example, “a policy rule could be ‘the MP3 download application 204D runs on the lowest cost network 104 and 106 available’.” *Montemurro*, [0029]. Thus, for example, if the MP3 download application is running on a WWAN network when a lower cost WLAN network becomes available, the MP3 download application would be directed to use the WLAN network. In doing so, *Rao/Araujo*’s processor would instruct the MP3 download application to transition to a different state of using the WLAN network.

165. According to *Montemurro*, “an application may interact with the rules engine 206A (e.g. via a query or other mechanism) to acquire access network information.” *Montemurro*, [0047]. As an example, “[s]ome access technology switches will result in IP address changes. If an application using the connection needs to be aware of the IP address, the change may not be transparent and may need to be communicated to (or otherwise discoverable by) the application.”

Montemurro, [0048]. *Montemurro* gives another example where connection “differences are significant (e.g. between GSM/GPRS and WAN 802.11 technologies),” the application may transition to a different state of behavior by “tailoring its behaviour to account for the new access technology in use,” such as changing “[b]uffer sizes, retransmission timers[,] etc.” *Montemurro*, [0048]. A POSITA would have recognized that these behavior changes are examples of transitions between different states, and these transitions can occur in response to a disallowed Internet access request (where a more suitable network becomes available).

VIII. DECLARATION

166. For the reasons set forth above, it is my opinion that claims 1-16 of the '076 patent would have been obvious in view of the prior art.


167. In signing this Declaration, I understand it will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I understand I may be subject to cross-examination in this case, and that any cross-examination will take place in the United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

168. I declare that all statements made herein of my knowledge are true and that all statements made on information and belief are believed to be true; and

further, that these statements were made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

169. I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed on this 8th day of August 2024.

By: 
Harry Bims, Ph.D.