

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SAMSUNG ELECTRONICS CO. LTD.
Petitioner,

v.

HEADWATER RESEARCH LLC,
Patent Owner.

IPR2024-01407
Patent 9,179,359

**DECLARATION OF ERIK DE LA IGLESIA IN SUPPORT OF PATENT
OWNER'S RESPONSE TO THE PETITION**

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

Executed on December 26, 2024, at Mountain View, California.



Erik de la Iglesia

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	BACKGROUND AND QUALIFICATIONS	2
III.	LEGAL PRINCIPLES	6
A.	Claim Construction	6
B.	Burden of Proof	6
C.	Anticipation	6
D.	Obviousness	7
IV.	PERSON OF ORDINARY SKILL IN THE ART	8
V.	DR. BUTLER’S CHALLENGE TO CLAIM 26 IS UNSUPPORTED.	9

I, Erik de la Iglesia, hereby declare as follows:

I. Introduction

1. I am over the age of eighteen (18) years and otherwise competent to make this declaration.

2. I have been retained as an expert witness on behalf of Patent Owner for the above-captioned *inter partes* review (“IPR”). I understand that the petition for *inter partes* review involves U.S. Patent No. 9,179,359 (“the ’359 patent”).

3. I make this declaration based on my personal knowledge, educational background and training, consideration of the materials I discuss herein, and my expert opinions.

4. My work on this matter is being billed at my normal hourly consulting rate of \$750, with reimbursement for actual expenses. My compensation is not related to the outcome of any proceeding involving the ’359 patent. I have no personal interest in the outcome of the case.

5. In preparing this Declaration, I have reviewed and considered the ’359 patent, the ’359 patent’s prosecution history, the Petition and prior art references submitted with the Petition, the declaration submitted by Dr. Butler in this proceeding, and each document cited in my declaration.

II. Background and Qualifications

6. My qualifications for forming the opinions given in this declaration are summarized here and are addressed more fully in my curriculum vitae, which is submitted as Exhibit 2011.

7. I hold a Bachelor of Science degree in Electrical Engineering from the University of Florida, and a Master of Science degree in Electrical Engineering from Stanford University. While at Stanford, I was an NSF Graduate Research Fellow. My graduate focus included VLSI CMOS architecture, high-speed circuit design, computer networks and protocols, and artificial intelligence. I have worked in several computer and electrical engineering related fields for over 25 years resulting in 68 issued U.S. patents that have been cited over 2850 times in USPTO patent applications. Since 2003, I have served in a number of chief architect and founder roles of companies that have been acquired or gone public. My experiences include the design or implementation of networks and network-based computation and protocol systems similar to the technology of the '359 patent, including systems for content targeting based on user behavior across devices.

8. Between December 1997 and November 1999, I was a circuit design engineer at Intel Corporation within the Mobile and Handheld Products Group. My responsibilities included working on several aspects of mobile processors and chipsets including clock and power distribution, dynamic clock frequency

selection, and the circuit design, device simulation, and validation of such features. Until March of 2000, I worked on similar functionality within the Microprocessor Products Group for the Itanium 3 and 4 products.

9. My work in power distribution included thermal modeling and predictive control of silicon devices and their packaging such as multi-tiered cooling systems for embedded, mobile, and desktop environments. Such systems included multiple tiers of thermal sensors, active and passive thermal control systems, and the feedback and analytic firmware and software control. My models were part of the development of the Thermal Design Power (TDP) metrics used by OEM and system designers to plan thermal control solutions for consumer and industrial products.

10. Between March 2000 and September 2003, I worked as a logic designer and architect for WebStacks and, through acquisition, Extreme Networks. WebStacks built a fully hardware-based TCP and HTTP processing stack for network processing functions including load balancing, proxy processing, and content rewriting. I was responsible for the architecture and design of the HTTP processing engine including header extraction and processing.

11. Between August 2003 and August 2007, I worked as Founder, Chief Architect and Director of Engineering for Reconnex, later acquired by McAfee. Reconnex built a gigabit line-rate network security analyzer capable of

classification, heuristic analysis, and policy enforcement on arbitrary protocols and content types. The Reconnex product captured and analyzed content transmitted over the internet using HTTP and other protocols, including storage protocols. Design of the Reconnex system was based on a modified and customized version of the Linux kernel which included specifically designed driver source code and interfacing between the application and kernel domains. Reconnex technology was employed in the audit and compliance verticals based on its ability to quickly analyze and classify the content of transmissions and interact with protocols and communications crossing corporate and departmental boundaries.

12. Between August 2007 and September 2008, I founded Strangways, a company focused on addressing messaging security for webmail and internet messaging protocols. Products and technology developed by Strangways were sold to SendMail Inc., Iron Port (Cisco), and other companies. The Strangways product utilized a Squid caching server and ICAP content rewrite server to translate HTTP to SMTP.

13. Between September 2008 and April 2014, I worked as Chief Architect for Gridiron Systems and, through acquisition, Senior Director of Technology for Violin Memory. I held the Senior Director position with Violin Memory both before and after the company's IPO. My responsibilities included architecture and

design of flash media controllers, adaptive caching design, and machine learning for application storage behavior analytics.

14. The Gridiron product was a storage accelerator used in many analytic environments such as big data architectures with constant ingest and processing of financial, consumer and IoT data. In my role as Chief Architect, I often worked with customers to implement and accelerate their analytics environments and spoke on related topics at trade shows and industry user groups. During product development, I worked with the data science team at Shopzilla to improve database and storage performance for real-time shopping recommendation and user tracking for advertising.

15. My work at Webstacks and Reconnex relate directly to packet processing including concepts such as priority and rules based processing relevant to the current proceeding. Approximately two dozen of my issued patents relate directly to these concepts.

16. Additional information regarding my employment history and prior expert testimony is included in Exhibit 2011. As reflected in Exhibit 2011, I have consulted on several matters involving intellectual property disputes including instances of alleged patent infringement requiring analysis of source code, reviewing and analyzing patents, writing expert reports, and testifying in court. I have testified at deposition and at trial in a number of these cases.

III. Legal Principles

A. Claim Construction

17. I understand that the first step in performing a validity analysis of the patent claims is to interpret the meaning and scope of the claims by construing the terms and phrases found in those claims. I understand that the appropriate construction of a claim term is its ordinary and accustomed meaning as understood by one of ordinary skill in the art at the time of the invention in the context of the entire patent and the prosecution history.

18. I understand that standard for claim construction in an *inter partes* review is the same standard as is applied in district court proceedings.

19. I understand that a determination of the meaning and scope of the claims is a matter of law. I have been informed that to determine the meaning of the claims, one should consider the intrinsic evidence, which includes the patent's claims, written description, and prosecution history.

B. Burden of Proof

20. I understand that in an *inter partes* review, the petitioner has the burden of proving unpatentability by a preponderance of the evidence.

C. Anticipation

21. I have been instructed by counsel and understand that a reference is anticipated if a single prior art reference discloses each and every claim element,

either explicitly or inherently, as arranged in the same way as in the claim. I understand that where even one claim element is not disclosed in a reference, a contention of anticipation fails.

22. I further understand that when a reference fails to explicitly disclose a claim element, that reference inherently discloses that element only if the reference must *necessarily* (or “inherently”) include the undisclosed claim element.

D. Obviousness

23. I have been instructed by counsel and understand that a combination of prior art references may render a claim obvious if, at the time of the invention, a person of ordinary skill in the art would have selected and combined those prior-art elements in the normal course of research and development to yield the claimed invention.

24. I understand that in an obviousness analysis, one should consider the *Graham* factors, including: the scope and content of the prior art; the differences between the claimed inventions and the prior art; and the level of ordinary skill in the art. I further understand the obviousness analysis is to be performed on a claim-by-claim basis. I understand that a person of ordinary skill in the art is a person of ordinary creativity, not an automaton.

25. I have been instructed by counsel and understand that obviousness requires more than a mere showing that the prior art includes separate references

covering each separate limitation in a claim under examination. I understand obviousness requires the additional showing that a person of ordinary skill at the time of the invention would have been motivated to combine those references in a manner that would include all limitations of the challenged claim, and, in making that combination, a person of ordinary skill in the art would have had a reasonable expectation of success.

26. I also understand that an obviousness analysis must be conducted with awareness of the distortion caused by hindsight bias and with caution of arguments reliant upon *ex post* reasoning. For instance, I understand that when considering obviousness, I should put myself in the position of a person of ordinary skill in the field at the time of the invention, rather than considering new information that is known today, but was not known before the priority date of the challenged patent.

IV. Person of Ordinary Skill in the Art

27. Dr. Butler proposes that a POSITA in the field of the '359 patent “would have had (1) at least a bachelor’s degree in computer science, computer engineering, electrical engineering, or a related field, and (2) at least two years of industry experience in services and application implementation in communication networks. Additional graduate education could substitute for professional experience, and vice versa.” Ex. 1003, ¶15. For purposes of this my analysis, I do not disagree with this proposed level of skill in the art. I meet at least these

requirements of a POSITA, and have met those requirements since before the priority date of the '359 patent.

V. Dr. Butler's challenge to claim 26 is unsupported.

28. Claim 26 depends from claim 1 – those claims are reproduced below:

1. A wireless end-user device, comprising:

a wireless wide area network (WWAN) modem to receive and transmit Internet data between the device and at least one WWAN when configured for and connected to the at least one WWAN;

a wireless local area network (WLAN) modem to receive and transmit Internet data between the device and at least one WLAN when configured for and connected to the at least one WLAN; and

one or more processors configured to apply a stored WWAN differential traffic control policy to Internet data service provided using the WWAN modem and the at least one WWAN, and indicate to a particular application, via an application program interface (API), one or more network access conditions based on the applied WWAN differential traffic control policy, wherein the one or more network access conditions includes a network access condition that indicates the unavailability to the particular application of an Internet data service that is currently available via the WWAN modem to a different application, but that is indicated as unavailable to the particular application.

26. The wireless end-user device of claim 1, wherein the one or more processors are further configured to determine that the different application is an active application based on a user interaction with the different application.

29. As shown above, Claim 26 requires that the system be configured to “determine that the different application” (i.e., the application for which the Internet data service is currently available to) “is an active application based on a user interaction with the different application.”

30. A POSITA would understand this limitation to be straightforward, and as requiring that the “determin[ation] that the different application is an active application” would be “based on a user interaction with the different application”—in other words, the status of user interaction with the different application would be an *input* to the “active application” determination, such that the “active application” determination would be *based on* the user interaction with the application.

31. This understanding is the straightforward meaning of this language, and is also supported by the '359 patent. Specifically, the '359 patent teaches that the system can entail “classification rules” for activities / applications “can include, for example, examining if one or more of the following has taken place within a specified period of time: user has interacted with the device, *user has interacted*

with the service usage activity, [etc.]” Ex. 1001 at 101:27-36 (emphasis added). In other words, user interaction with the service activity (which a POSITA would understand includes an application that uses network service) is used as an input to the classification rules, such that the classification of an application is “based on a user interaction with the application” as recited in claim 26.

32. In my opinion, Dr. Butler does not show that his Shell-Cole combination discloses or renders obvious that the “determinat[ion] that the different application is an active application” would be “based on a user interaction with the application” as recited by claim 26.

33. Dr. Butler alleges that a “POSITA would have... understood and found obvious that an application that provides a voice-related connection request ([i.e., the alleged] *different application*) involves a determination that a user is likely to interact with the application (e.g., voice call) (*a user interaction with the different application*), and thus, prioritization by [Shell’s] connection manager modules involves *determin[ing] that the different application is an active application.*” Ex. 1003, ¶140 (emphases in original).

34. Dr. Butler’s assertion in this regard does not show the obviousness of claim 26 for three reasons, which are discussed below in turn.

35. **First**, Dr. Butler’s statement relates to “a determination that a user is *likely to interact* with the application” which is obviously not the same thing as a

processor basing any determination that an application is “active” “based on a user interaction with [the] application” as claimed. The claim does not recite “determinin[ing] that the different application is an active application based on [whether the user might hypothetically interact with] the different application”; instead, it requires “determinin[ing] that the different application is an active application based on *a user interaction with* the different application.”

36. The fact that a POSITA might hypothetically be able to infer that if a voice application is issuing a voice-related connection request, a user would be somewhat more likely to interact with that application, does not mean that there has been a determination “based on a user interaction” with that application. Instead, under Dr. Butler’s own articulation of his theory, any determination would at *most* be based on the *probability* of interaction with an application (e.g., a user might pick up the phone in response to an incoming voice call), not based on an *actual* user interaction with the application.

37. **Second**, even *if* Dr. Butler’s theory related to *actual* user interaction (as opposed to just *potential* user interaction as Dr. Butler alleges), Dr. Butler provides no reason why a POSITA would believe that Shell discloses or renders obvious “that an application that provides a voice-related connection request... involves a determination that a user is likely to interact with [or has interacted with] the application.” Ex. 1003, ¶140. While a POSITA may be able to guess that

if an application issues a “voice-related connection request” that a user is at least somewhat likely to interact with the application, he provides no reason why a *processor* would make such an evaluation.

38. More realistically, a POSITA would have programmed the system to simply prioritize voice requests higher than data requests, *regardless of user interaction*, as Shell teaches. See Ex. 1005, ¶57 (“[V]oice requests may be assigned a higher priority than data requests.”) (cited by Ex. 1003, ¶140). Dr. Butler provides no reason why a data request would receive a higher priority than a voice request based on user interaction. For example, to the extent Dr. Butler’s theory is based on the theory that voice connection requests are inherently more likely to involve user interaction than data requests, the prioritization scheme can be implemented without the processor actually evaluating any probability of “user interaction” with the voice application. Instead, by simply always prioritizing a voice application’s connection request over a data connection request, the processor would be implementing an effective prioritization scheme, and it would be unnecessary and inefficient for the processor to separately (or as part of the prioritization scheme) evaluate the specific likelihood of user interaction.

39. Similarly, to the extent Dr. Butler’s theory is about *actual* user interaction (and not just likelihood of user interaction), Dr. Butler also does not provide any reason why a POSITA would believe Shell to disclose or render

obvious the processors evaluating actual user interaction in the context of the Shell-Cole combination. For instance, a processor might hypothetically evaluate whether a user placed an outgoing call via the application (which would have involved prior user interaction with the application) as opposed to the application receiving an incoming call (which may not have involved prior user interaction with the application). But in the context of the Shell-Cole combination, Dr. Butler does not explain why a POSITA would have thought it obvious to care at all about the status of user interaction in prioritizing these two (or any other) scenarios. Shell teaches that voice connections receive priority over other types of connections, and it would not have been obvious in the context of Shell-Cole to change the priority of voice connection requests based on a user interaction with the requesting voice application. Instead, a POSITA would understand Shell to disclose that voice connection requests should be prioritized over data requests, *regardless* of the status of user interaction, and Dr. Butler provides no evidence or reasoning to the contrary.

40. **Third**, Dr. Butler does not provide any explanation as to how his assertion in paragraph 140 of his declaration relates to specific requirements of the claim language, which requires that the “one or more processors are further configured to determine that the different application is an active application *based on a user interaction* with the different application.” Ex. 1001, claim 26 (emphasis

added). Dr. Butler cites to a portion of Shell that discloses that “a request from an application that is currently active may be assigned a higher priority than a request from an application running in the background.” Ex. 1003, ¶140 (quoting Ex. 1005, ¶57). But the claim does not require simply determining whether an application is active, it requires determining whether the application is active ***based on a user interaction*** with that application. Dr. Butler’s analysis does not explain why the “active application” determination would in any way rely on or otherwise be “based on” “a user interaction with [the] application” as claim 26 requires.

41. For example, in the context of Shell-Cole, a voice application may be deemed “active” if it is requesting a voice connection (e.g., to receive an incoming call) even if there has been no user interaction with that application. Dr. Butler provides no evidence or analysis that a POSITA would find it obvious to make the “determin[ation] that the different application is an active application” to be “based on a user interaction” with that application. Indeed, Dr. Butler’s assertion that “a voice-related connection request” inherently involves the user being more likely to interact with the application (Ex. 1003, ¶140) shows that the system need ***not*** evaluate a user interaction with the application to determine the appropriate priority. If the processor simply prioritizes voice-related connection requests (even if it does not consider a user interaction with the application, or even the likelihood of potential user interaction with the application), the operation of the processor

would be consistent with Shell’s teachings that voice requests should have higher priority than data requests (*see* Ex. 1005, ¶57), and would also be consistent with Dr. Butler’s implied assertion that applications that are likely to be interacted with by users should be prioritized (*see* Ex. 1003, ¶140).

42. Again, this prioritization of voice requests may ultimately end up prioritizing applications that are likely to be interacted with (or even applications that have been interacted with) by a user, but this is just a side effect of “blindly” prioritizing voice connection requests. Prioritizing voice requests does not require (and does not render obvious) that a processor would make an “active application” determination that was actually “based on a user interaction” with the application issuing the request.

43. By way of analogy, one might assign school children a priority based on their age. As children who are older are, on average, more likely to be taller, this would have the effect of giving higher priority on average to taller children. But the prioritization ***based on age*** would not be a prioritization “based on height” even if it ultimately had the side effect of assigning higher priority to taller children, because the assignment system never took into account (i.e., was not “based on”) the height of any child. Likewise, assigning higher priority to voice-related connection requests may have the side effect of prioritizing applications that are more likely to be interacted with by users, but the processors assigning

high priority to voice-related connection requests are not making their priority determination (and are not making an “active application” determination) “based on a user interaction” with the application.

44. Finally, I note that Dr. Butler is not precise as to whether he believes Shell-Cole **discloses** the additional limitation of claim 26, or whether it would be obvious. *See* Ex. 1003, ¶140 (alleging: “A POSITA would have therefore understood and found obvious that...”). I note that Shell-Cole does not disclose (either expressly or inherently) the limitation of claim 26, because the cited portions of Cole that Dr. Butler relies on all relate only to prioritizing “active” applications or prioritizing voice applications, but not to making a determination that an application is an “active application” that is “based on a user interaction” with the application, as claim 26 requires. And the claim 26 requirement is certainly not inherent in the context of Shell-Cole, as the one or more processors could simply be configured to prioritize voice communications (or prioritize active applications) **without** evaluating user interaction with an application.

45. I also note that Dr. Butler’s reference to Shell at paragraphs 8 and 13-14 does not support the obviousness of claim 26. As an initial matter, paragraph 8 of Shell is in the “Background and Relevant Art” section of Shell’s disclosure (ex. 1005 at ¶4), and is generally discussing “problems associated with connection methods” (Ex. 1005 ¶9). Thus, whatever purported relevance paragraph 8 of Shell

may have to the claim language (which Dr. Butler does not clearly specify), the teaching of Shell's paragraph 8 is not something a POSITA reading Shell would have viewed as *desirable*, because paragraph 8 is discussing problems that existed in the prior art.

46. Paragraphs 13-14 of Shell also do not provide any support for the obviousness of claim 26. While Dr. Butler does not explain *why* these paragraphs of Shell support his obviousness opinion, paragraph 13 generally restates that “a voice call may be given higher[] priority over data calls that share the same modem,” similar to the teaching from paragraph 57 of Shell I have addressed above. *See* Ex. 1005, ¶15. And paragraph 14 does not appear to address any teachings disclosing prioritization or classification techniques at all. *See* Ex. 1005, ¶14.

47. Regarding obviousness, Dr. Butler does not provide any motivation to modify Shell-Cole to perform the claimed determination based on a user interaction with the application. As I have explained above, a POSITA reviewing Shell would have found it obvious to simply prioritize voice communications *without* looking to whether a user has interacted with the requesting voice application, and I do not believe a POSITA would have found it obvious to modify Shell (or Shell-Cole) to arrive at the requirements of claim 26.