

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.
Petitioner,

v.

MONARCH NETWORKING SOLUTIONS, LLC,
Patent Owner.

Case IPR2024-00485
Patent 8,451,844

**DECLARATION OF WILLIAM HENRY MANGIONE-SMITH IN
SUPPORT OF PATENT OWNER'S RESPONSE TO PETITION FOR
INTER PARTES REVIEW OF U.S. PATENT NO. 8,451,844**

Mail Stop "PATENT BOARD"

Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I. Introduction	1
II. Summary of Opinions	1
III. Qualifications	2
IV. Compensation.....	6
V. Materials Considered	7
VI. Legal Understandings.....	7
A. Claim Construction	7
B. Obviousness	8
C. Person of Skill in the Art.....	11
VII. Overview of the '844 Patent	12
A. IPv4 and IPv6 Interconnection.....	12
B. The '844 Patent	15
VIII. Level of Ordinary Skill in the Art.....	17
IX. Claim Construction	18
A. “domain”	18
X. Petitioner’s References.....	20
A. Li (EX1037)	20
B. Li-2 (EX1038).....	23
C. Wu (EX1039)	26
XI. Issues Common to all Grounds	29
A. Li does not disclose the claimed IPv4 and IPv6 domains.....	29
B. The domains of Figure 7 are not separate domains	32
C. Figures 8, 10, and 11 show inter-domain embodiments	35
D. The autonomous systems of Figure 7 are under common control.....	37
XII. Ground 1	39
A. Dr. Bhattacharjee’s opinion does not support the conclusion that Li-2 was publicly available	39
B. A POSITA would not have been motivated to combine Li and Li-2 as proposed	40
1. Li and Li-2 address different problems	40

2.	Li-2 is not an extension of Li.....	42
3.	Dr. Bhattacharjee relies on motivation to combine references that are not prior art.....	43
4.	The prior art identifies NAT tables as a well-known solution for IPv4 address exhaustion	45
XIII.	Ground 2.....	46
A.	Dr. Bhattacharjee recognizes that Li fails to teach port number limitations 46	
B.	The evidence does not show simultaneous invention	47
XIV.	GROUND 3.....	50
A.	A POSITA would not have been motivated to combine Li and Wu	50
1.	Li and Wu Are Directed To Different Communication Environments..	50
2.	Li is stateless, but Wu is stateful	52
3.	Wu does not teach an a destination address constructed with a destination port number	54
XV.	Conclusion.....	58

I, William Mangione-Smith, declare as follows:

I. INTRODUCTION

1. My name is William Henry Mangione-Smith. I am over eighteen years of age, of sound mind, and qualified to make the statements set forth in this Declaration. Except for the legal principles counsel has supplied to me, all the facts and statements contained herein are within my personal knowledge and are in all respects true and correct.

2. I have been asked by Monarch Networking Solutions, LLC (“Monarch” or “Patent Owner”) to submit this declaration in response to the Petition filed by Juniper Networks, Inc. (“Juniper” or “Petitioner”) and the declaration submitted by Dr. Samrat Bhattacharjee (“Bhattacharjee Declaration”) regarding U.S. Patent No. 8,451,844 (“the ’844 Patent”).

II. SUMMARY OF OPINIONS

3. The Petition and Dr. Bhattacharjee’s declaration do not demonstrate that claims 1-10 of the ’844 Patent are invalid.

4. Dr. Bhattacharjee presents three grounds. All three Grounds rely on Li. However, Li does not disclose the claimed IPv4 and IPv6 domains because the embodiment Dr. Bhattacharjee relies on describes only one domain.

5. Ground 1 is a combination of Li and Li-2. However, Li-2 is not prior art, and even if it was, a POSITA would not combine Li and Li-2 as proposed by Dr. Bhattacharjee.

6. Ground 2 relies on Li alone. There, Dr. Bhattacharjee argues that evidence of “simultaneous invention” shows that the claims are obvious. However, Li does not describe several elements of the claimed invention, and the other evidence Dr. Bhattacharjee points to do not show simultaneous invention.

7. Ground 3 is the combination of Li and Wu. But, a POSITA would not combine Li with Wu for several reasons, including that they deal with different networking environments and that the combination of Li and Wu would eliminate the purported benefits of Li. Wu also does not disclose the destination port and constructed address claimed by the '844 patent.

III. QUALIFICATIONS

8. I have summarized my educational background, career history, publications, and other relevant qualifications below. details of my education and employment history, recent professional service, patents, publications, and other testimony are set forth in my current curriculum vitae, attached to this Declaration as Exhibit EX2013.

9. My technical background and experience include computer system design, including low level circuitry, computer architecture, computer networking,

graphics, application software, client-server application, Web technology, and system software (e.g., operating systems and compilers). I am a member of the Institute of Electrical and Electronics Engineers and the Association for Computing Machinery, which are the two most significant professional organizations in my profession. I have been employed as a design engineer, research engineer, professor and technical expert. Over my professional career, I have been an active inventor with 112 issued U.S. patents, 200 published and pending U.S. patent applications and many unpublished U.S. patent applications.

10. From 1984 until 1991, I attended the University of Michigan in Ann Arbor, Michigan, where I was awarded the degrees of Bachelor of Science and Engineering, Master of Science and Engineering, and Doctor of Philosophy in Electrical Engineering and Computer Science. My doctoral research focused on high performance computing systems including computer architecture, applications and operating system software, and compiler technology. One of my responsibilities during my graduate studies included teaching senior undergraduate students who were about to enter the profession.

11. After graduating from the University of Michigan, I was employed by Motorola in Schaumburg, Illinois. While at Motorola, I was part of a team designing and manufacturing the first commercial battery-powered product capable of delivering Internet email over a wireless (i.e., radio frequency) link and

one of the first personal digital assistants. I also served as the lead architect on the second-generation of this device with control over the entire system design including the memory subsystem architecture, embedded processor, ASIC, power system, and analog circuitry. Part of my responsibilities at Motorola involved the specification, design, and testing of system control Application-Specific Integrated Circuits (“ASICs”). I conducted the initial research and advanced design that resulted in the Motorola M*Core embedded microprocessor. M*Core was designed to provide the high performance of desktop microprocessors with the low power of contemporaneous embedded processors. The M*Core received widespread use in many communications products including various cellular handsets, advanced pagers, and embedded infrastructure. While at Motorola I was the sole inventor on one U.S. patent.

12. From 1995 until 2005, I was employed by the University of California at Los Angeles (“UCLA”) as a professor of Electrical Engineering. I was the director of the laboratory for Compiler and Architecture Research in Embedded Systems (“CARES”) and served as the field chair for Embedded Computing Systems. The CARES research team focused on research, engineering and design challenges in the context of battery-powered and multi-media mobile computing devices. One of the key developments of my lab was the Mediabench software tool, which is widely used to design and evaluate multi-media embedded devices. Key

elements of Mediabench include software that is essential for modern digital wireless communications. My primary responsibility, in addition to classroom teaching, involved directing the research and training of graduate students. I was a tenured member of the faculty and had responsibilities for teaching as well as scholarly research. While at UCLA I was a named inventor on three U.S. patent applications, one of which issued as a patent. My colleagues at UCLA were some of the leading scientists and engineers in the world with a long list of innovations from computer network security devices to the nicotine patch. The graduate student researchers in my laboratory came from a diverse set of backgrounds, all with undergraduate degrees in computer engineering, electrical engineering or computer science, many with multiple years of experience working as professional engineers in areas such as software development, computer system design and ASIC circuit design.

13. From 2005 until 2009, I was employed at Intellectual Ventures in Bellevue, Washington. My responsibilities at Intellectual Ventures include business development, technology assessment, market forecasting, university outreach, collaborative inventing, intellectual property licensing support, and intellectual property asset pricing. My colleagues and co-inventors at Intellectual Ventures included the former lead intellectual property strategist at Intel, Intel's former lead IP council, Microsoft's former chief software

architect, the founder of Microsoft research, the designer of the Mach operating system, the architect of the U.S. Defense Department's Strategic Defense Initiative, the founder of Thinking Machines (a seminal parallel processing computer system), and Bill Gates. I had responsibility for hiring and managing over 15 staff members including multiple Ph.Ds. with degrees in electrical engineering and decades of experience in product design and engineering.

14. I have been involved in design and engineering of power management and control systems for over 30 years.

15. I have also published extensively in peer-reviewed forums on topics related to the design of low-power and power-efficient systems, including hardware, software, and communications protocols.

16. Additional details of my education and employment history, recent professional service, patents, publications, and other testimony are set forth in my current curriculum vitae, attached to this Declaration as **Appendix A**.

IV. COMPENSATION

17. I am being compensated for my time at my current standard rate of \$800 per hour for the services that I provide in connection with this IPR. That compensation is not contingent upon my performance, the outcome of this IPR, or any issues involved in or related to this IPR.

V. MATERIALS CONSIDERED

18. Among the materials I reviewed in forming my opinions are the '844 Patent, its prosecution history, and the patent application to which the '844 Patent claims priority, as well as the patents, references, and other materials cited in this expert declaration.

19. I also reviewed Juniper's Petition for *Inter Partes* Review of the '844 Patent and all its exhibits (EX1001 to EX1056), especially exhibit EX1003, the Declaration of Dr. Samrat Bhattacharjee submitted in support of that Petition.

VI. LEGAL UNDERSTANDINGS

20. I am not an attorney. For the purposes of this declaration, I have been informed about certain aspects of the law that are relevant to my analysis and opinions. In forming my opinions, I have not relied upon Monarch's attorneys for any purpose other than explaining the following legal principles.

A. Claim Construction

21. I understand that determining whether a patent claim is valid requires a two-step analysis. First, the claim must be construed. Second, the properly construed claim must be compared to the prior art.

22. With respect to the first step, I understand that claims are construed from the perspective of a person of ordinary skill in the art ("POSITA") at the time of the invention. I understand that in the context of an *inter partes* review proceeding, absent a disclaimer or express definition, claim terms are given their

plain and ordinary meaning to a POSITA in view of the specification and file history and the knowledge of a POSITA. Unless otherwise noted, I have applied what I consider to be the plain and ordinary meaning of all terms as understood by a POSITA at the time of the invention, which I understand to be approximately June 30, 2008 (the filing date of French Patent Application No. 08 54398, to which the '844 claims priority).

B. Obviousness

23. I understand that a patent claim may be found unpatentable in an IPR action as obvious only if the Petitioner establishes by a preponderance of the evidence that, as of the priority date, the subject matter of the claim, considered as a whole, would have been obvious to a person having ordinary skill in the field of the technology (the “art”) to which the claimed subject matter belongs.

24. I understand that a claim would have been obvious under 35 U.S.C. § 103 if one or more prior art references disclose, expressly or inherently, every claim limitation in such a way to render the claim, as a whole, obvious to a person of ordinary skill in the art at the time the purported invention was made. The relevant standard for obviousness is as follows:

A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art

to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

35 U.S.C. § 103(a).

25. In determining whether a patented invention would have been obvious, the following factors should be considered: (a) the scope and content of the prior art; (b) the differences between the prior art and the claims at issue; (c) the level of ordinary skill in the art; and (d) whatever “secondary considerations” may be present.

26. I understand that certain “secondary considerations” may be relevant in determining whether or not an invention would have been obvious, and that these secondary considerations may include commercial success of a product using the invention, if that commercial success is due to the invention; long-felt but unsolved need for the invention; evidence of copying of the claimed invention by others; industry acceptance; initial skepticism of the invention; failure of others to solve the problem; praise of the invention by others; the taking of licenses under the patents by others; and simultaneous invention by others (which may be an indication that the claimed invention was the result of ordinary skill in the art).

27. I understand that a patent composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art. While multiple prior art references or elements may, in some circumstances, be combined to render a patent claim obvious, I understand that I

should consider whether there is an apparent reason to combine the prior art references or elements in the way the patent claims. To determine whether such an apparent reason exists to combine the prior art references or elements in the way a patent claims, it will often be necessary to look to the interrelated teaching of multiple patents, to the effects of demands known to the design community or present in the marketplace, and to the background knowledge possessed by a person having ordinary skill in the art.

28. I also understand that when the prior art teaches away from combining prior art references or certain known elements, discovery of a successful means of combining them is more likely to be non-obvious. A prior art reference may be said to teach away from a patent when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the patent or would be led in a direction divergent from the path that was taken by the patent.

29. I also understand that it is not permissible to use hindsight in assessing whether a claimed invention is obvious. Rather, I understand that, to assess obviousness, one must place oneself in the shoes of a person having ordinary skill in the relevant field of technology at the time the inventions were made who is trying to address the issues or solve the problems faced by the inventor and ignore the knowledge you currently now have of the inventions.

30. I also understand that there are numerous ways in which to articulate the legal standard for obviousness, including: (a) combining prior art elements according to known methods to yield predictable results; (b) simple substitution of one known element for another to obtain predictable results; (c) use of a known technique to improve similar devices (methods, or products) in the same way; (d) applying a known technique to a known device, method, or product ready for improvement to yield predictable results; (e) choosing from a finite number of identified, predictable solutions with a reasonable expectation of success; (f) known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; and (g) some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention.

C. Person of Skill in the Art

31. I understand that the meaning of claim terms and a patent's validity are to be determined from the vantage point of one of ordinary skill in the art at the time of the invention, a "POSITA." In determining who would be one of such ordinary skill, I understand it is appropriate to consider criteria such as: (a) the type of problems encountered in the art; (b) prior art solutions to those problems; (c) the

rapidity with which innovations are made; (d) the sophistication of the technology; and (e) the education level of active workers in the field.

VII. OVERVIEW OF THE '844 PATENT

A. IPv4 and IPv6 Interconnection

32. Many telecommunication networks (including the Internet) rely on an addressing scheme defined by the Internet Protocol (“IP”) to identify the source and destination addresses for devices connected to the network. EX1001 at 1:17-21. Devices known as routers can use these IP addresses to route data packets between the source and destination devices. *See id.* The conventional version of the IP address system (known as IPv4) provided a limited number of possible addresses, which limited the total number of devices that could connect to the network. *Id.* at 1:61-2:6. Due to the limited number of available addresses, home or business networks would frequently operate as a “private network,” in which devices within the network received private addresses with access to the public Internet limited to exchanges occurring through a home gateway or application layer gateway. *Id.* at 1:43-56. The private terminals needing access to the public network share a limited number of public addresses assigned to the home or business network through techniques such as network address translation (“NAT”) process. *Id.* at 1:51-60.

33. To further extend the number of IPv4 public addresses reachable, “a solution known as Double NAT or Operator NAT has been proposed and

implemented. It entails activating a NAT function within the telecommunications network of the operator so that the home gateways use a private address (instead of a public address) in their outgoing NAT tables.” *Id.* at 2:11-22. Both home gateway NAT and Operator NAT introduce significant drawbacks, including (1) requiring the interconnection equipment to track and maintain the state of connections in order to properly route packets; (2) introducing additional modifications of data packets for each level of NAT introduced (consuming additional time and processing resources); and (3) degradations in the levels and types of services available due to the limited availability of port forwarding and the resulting need to keep IP sessions active. *Id.* at 2:23-55; 3:40-52; 13:29-46.

34. A newer IP address protocol (referred to as “IPv6”) provides a significantly larger number of public IP addresses. *Id.* at 1:61-2:1. However, IPv6 is not designed to be backward compatible with IPv4; thus, integrating IPv6 networks with existing IPv4 networks requires additional configuration beyond that dictated by the protocol standards. *See, e.g.,* RFC 1933 at 1 (detailing “Transition Mechanisms for IPv6 Hosts and Routers” that are “designed to allow IPv6 nodes to maintain complete compatibility with IPv4”). Many organizations and operators have not made thee to transition from IPv4 to IPv6 due to “financial, strategic, and technical reasons.” EX1001 at 2:7-10. “It is therefore necessary to provide for a changeover to IPv6 in the medium term. Such a changeover will necessarily lead to

a period of transition during which IPv6 domains will have to interconnect with IPv4 domains.” *Id.* at 2:46-55.

35. Because there is no single standardized strategy for transitioning from IPv4 to IPv6, there is also no single standardized solution for providing interconnection from one device or network type to another. For example, one organization may need to update its infrastructure backbone to IPv6 without breaking connectivity for existing IPv4 terminals, while another organization may need to provide a terminal with the ability to traverse an intermediate IPv4 network to exchange packets with an IPv6 terminal in a third network, while yet another organization may need to provide connectivity to mobile hosts across access points connected to an IPv6 network. This mixed technology environment is further complicated if a source or destination terminal only has a private IPv4 address because there are multiple different NAT types, each of which has different requirements regarding source and destination addresses and ports in the IP header. *See* EX1001 at 1:57-60. Because of these and other factors, it is well-known that different interconnection solutions are frequently incompatible and individual pieces of an interconnection solution cannot just be mixed and matched. Instead, the solution must be specifically tailored to the problems of a given transition strategy.

36. Accordingly, numerous solutions have been proposed to allow interconnection between IPv4 and IPv6 networks. *See* EX1001 at 2:11-22. The

various IP interconnection techniques generally fall into two categories: “stateful” and “stateless.” *See id.* at 3:40-52, 13:47-51. In a stateful solution, border or edge devices must maintain the “state” of each connection or session within their networks. *Id.* at 3:40-52. When state must be maintained within a provider network or network core there occurs a significant overhead burden on routing devices as they attempt to maintain the state for numerous connections. *Id.* at 3:40-52, 13:43-46.

37. Stateless solutions avoid many of these drawbacks. *Id.* at 13:47-56. Because there is no need to maintain the “state” of the connections within the provider network or core network, stateless solutions provide better scaling, better performance, flexible network design and better cost effectiveness.

B. The '844 Patent

38. The inventors of the '844 Patent recognized the drawbacks of stateful solutions described above. Thus, the invention of the '844 is a stateless solution for providing IPv4-to-IPv6 interconnectivity. EX1001 at 3:40-52. The invention provides methods and devices for routing data packets to and from terminals in separate domains through and IPv6 domain without the need to rely on NAT-based state tracking.

39. Figure 1 of the '844 Patent provides an example of a network implementing the claimed invention.




FIG. 1

40. As shown in Figure 1, Domain 3 is an IPv4 domain and Domain 1 is an IPv6 Domain. As a result, packets originating from terminals in Domain 3, which will be IPv4 packets, cannot be directly routed in Domain 1 or to a terminal that lacks an IPv4 address. The '844 invention addresses this problem.

41. The invention teaches that an IPv6 packet can be created from an IPv4 packet by constructing an IPv6 destination address using a predetermined operator prefix, the destination IPv4 address, and the destination port number. '844 Packet at

2:59-3:18. To create this packet, a packet from an IPv4 domain will first be routed to an interconnection device (such as Access Node 40), which will select the IPv6 prefix corresponding to the interconnection device which has previously advertised that it is responsible for routing packets to the destination IPv4 address. '844 Patent at 10:14-22. The interconnection device will then construct an IPv6 packet from the IPv4 packet and the packet will be routed as normal through the IPv6 domain to a device such as Access Router 30, which can then forward the packet to the appropriate home gateway and terminal based on the IPv4 destination address and port. The home gateway (such as Home Gateway 20) can then route the packet to its destination terminal.

42. Throughout this declaration, I refer to the claims with the same annotations as Dr. Bhattacharjee. EX1003 at 12-14.

VIII. LEVEL OF ORDINARY SKILL IN THE ART

43. I understand Dr. Bhattacharjee has stated that a POSITA “would have had a Bachelor of Science degree in electrical engineering, computer engineering, computer science, or a related subject, and three years of work experience in computer networking, including experience with the IPv4 and IPv6 protocols, or a related field. Lack of professional experience can be remedied by additional education, and vice versa.” EX1003 ¶ 55. For the purposes of this declaration, I do

not dispute Dr. Bhattacharjee's description of a POSITA. I meet Dr. Bhattacharjee's criteria, and I have applied that description in my analysis.

IX. CLAIM CONSTRUCTION

44. I understand that a district court construed the terms of the '844 patent claims. EX1036. Dr. Bhattacharjee lists these constructions in paragraph 56 of his declaration. EX1003 ¶ 56. Dr. Bhattacharjee claims to use those instructions in his analysis. *Id.* I have done the same.

45. Separately, Dr. Bhattacharjee proposes a construction of the term "domain." *Id.* ¶ 90. Dr. Bhattacharjee argues that "domain" should mean "[a] telecommunications network administered by an operator" based on discussion in the '844 patent *Id.* (citing EX1001 at 1:30-31). I also understand that Patent Owner proposes a different construction of "domain," is "a network or subnetwork and its address space that are under common control." For the reasons I explain below, Patent Owner's construction of domain is correct, and I disagree with Dr. Bhattacharjee's proposed construction.

A. "domain"

46. I have review both the intrinsic evidence and the extrinsic evidence related to the term "domain." Based on this evidence, and my own knowledge and experience, a POSITA would understand that a "domain" is "a network or subnetwork and its address space that are under common control."

47. One key aspect of a “domain” is the concept of common control. The ’844 Patent’s specification recognizes this concept and states that “[a] telecommunications network *administered by an operator* is also referred to as a domain.” EX1001 at 1:30-31 (emphasis added). In other words, a domain is only understood in the context of the party the administers it: networks administered by different operators are different domains. The specification highlights the difference between a private domain (a network that is controlled or administered by a homeowner or business entity) and a domain operated by a service provider. *Id.* at 1:32-50; 4:41-43 (“The neighboring IPv4 domain may be a public network managed by an operator or a private network, either a home network or a business network.”).

48. The concept of common administration is important, because a network operator controlling a domain can take precautions, such as restricting where in the network address translation occurs, and ensuring private addresses are not re-used within that domain. In networks that are operated and controlled by different entities, these precautions are not possible: entities (such as a service provider) may implement their own layer of address translation, and the same private IPv4 addresses can be reused in different domains. *Id.* at 2:11-22 (discussing Operator NAT).

49. Extrinsic evidence supports this definition. For example, Newton’s Telecom Dictionary defines “domain” as “[t]he part of a computer network in which


the data processing resources are under common control.” EX2003 at 3 (emphases added). The IBM Dictionary of Computing defines “domain” as “[t]hat part of a computer network in which the data processing resources are under common control.” EX2004 at 4. Both definitions highlight the need for common control within a domain.

X. PETITIONER’S REFERENCES

50. Dr. Bhattacharjee offers opinions regarding three references: (1) Li; (2) Li-2; and (3) Wu. *See* EX1003 at 2-6. I discuss each of them below.

A. Li (EX1037)

51. Li appears to be an IEEE paper titled “Forwarding IPv4 Traffics in Pure IPv6 Backbone with Stateless Address Mapping.” Li discusses a mapping scheme that allows IPv4 hosts to communicate with other IPv4 hosts across an intermediate IPv6-only network. EX1037 at 1. Specifically, Li describes a technique for “forwarding IPv4 traffics in IPv6-only backbone....” *Id.* An example of Li’s IPv6 backbone is illustrated in Figure 7, reproduced below:




Id. at 7. In this figure, AS 23910 is a service provider’s IPv6-only backbone, and AS 4538 and AS 24348 each represent customer networks that are either dual-stack or IPv4 only. *Id.* at 6. AS 23910 IPv4, in the bottom-right corner of the figure, represents other IPv4 networks “out of the scope that the provider serves.” *Id.* at 7. This would include, for instance, the IPv4 Internet. *See id.* at 3 (Figure 1 showing IPv4 Internet in this same network architecture).

52. Li discusses two ways to map an IPv4 address to an IPv6 address for transport through an IPv6-only network. The first type of mapping is reserved for IPv4 addresses controlled by the administrator of the IPv6 backbone. These addresses are mapped to IPv6 addresses by combining: (1) IPv6 prefix; (2) IPv4

network id and optionally IPv4 host id; (3) padding; and (4) and IPv4 address. *Id.* at

3. This mapping is shown in Li's Figure 2, shown below:




Id.

53. Li only describes this mapping IPv4 packets traversing the IPv6 backbone to IPv4 addresses that are held by the autonomous system (AS) that controls the IPv6 backbone. *See id.* ("It is important to emphasize that only the IPv4 addresses held by the AS can be mapped, or the mapping would run out the most space of an IPv6/32 block.").

54. Li describes a different mapping scheme for IPv4 communications traversing the IPv6 backbone that are directed to IPv4 addresses not held by the domain provider (e.g., addresses belonging to other domain providers). In this second mapping scheme, IPv4 addresses that are not administered by the IPv6 backbone provider are mapped to IPv6 addresses by appending together the

following components: (1) IPv6 prefix; (2) “0xFFFFFFFF”; (3) padding; and (4) and IPv4 address. *Id.* at 3. This mapping is illustrated below:



Annotated Version of Li’s Figure 2 (adding special mapped prefix and removing bitstring descriptive annotations). Petitioner’s allegations and Dr. Bhattacharjee’s analysis only relate to the first mapping detailed above for intra-domain communications.


B. Li-2 (EX1038)

55. Li-2 is a cancelled Internet Draft submission titled “Bi-directional Address Borrowing (IVI)” detailing an IPv4-IPv6 transition scheme (referred to as IVI) in which “special IPv6 address blocks are defined, which can communicate with the global IPv6 networks directly and with the global IPv4 networks via stateless (or almost stateless) gateways.” EX1038 at 1, 2. Li-2 focuses on the requirement that “IPv6 hosts need to communicate with the global IPv4 networks” *Id.* at 4.

56. According to Li-2, IVI’s implementation by an ISP requires: “(1) The ISP borrows a subset of ISP’s global routable IPv4 address, maps it into IPv6 and physically uses it for ISP’s pure IPv6 hosts; [and] (2) The ISP borrows a subset of

ISP's global routable IPv6 address, maps it into IPv4 and virtually uses it for global IPv4 hosts.” *Id.* at 6. The first subset of addresses (a set of global routable IPv4 addresses referred to as “IVI4(i)”) are assigned to a set of pure IPv6 hosts within the ISP's network. *Id.* at 5. Thus, each such IPv6 host (“IVI Host”) has a global routable IPv4 address assigned to it for use in communicating with IPv4 hosts. The second subset of addresses (a set of global routable IPv6 addresses referred to as “IVIG46(i)”) are virtually assigned to every possible global IPv4 address. Thus, within the ISP's network, each possible global IPv4 address has a well-known corresponding IPv6 address.

57. The IVIG46(i) addresses are created using the following address mapping:



Id. at 6. Each IVIG46(i) address includes a 32-bit prefix followed by eight 1-bits (or “FF” in hexadecimal) to denote that it is an IVI address. *Id.* After the prefix and IVI identifier is the IPv4 address to which this IVI address corresponds followed by 56 0-bits for the remainder of the address bitstring. *Id.*


58. As detailed above, a borrowed IPv4 address (i.e., within the set IVI4(i)) is assigned to each IVI Host. Each IVI4(i) address has a corresponding address within IVIG46(i), and this IPv6 address (built using the mapping rule above) is physically assigned to an IVI Host for IPv6 communications. The set of IVIG46(i) addresses that correspond to the IVI4(i) addresses are referred to as IVI6(i) addresses. *Id.* at 5; *see id.* at 6 (“IVI4(i) is a mirror image of IVI6(i).”). “IVI6(i) is the special IPv6 address block which can communicate with both address families.” *Id.* at 6.

59. By default, there is a one-to-one mapping between each address in IVI4(i) and IVI6(i). *See id.* at 6 (“Because this mapping is 1-to-1 defined by the IVI mapping rule, it is stateless and it has feature of end-to-end address transparency.”). However, recognizing that “public-IPv4 address is a scarce resource,” Li-2 identifies port multiplexing “[t]o [further] increase the utilization ratio of the public IPv4 address.” *Id.* at 14. Li-2 explains that “[t]he mapping scheme is to use the least significant bits in the IVI6(i) to define the multiple mapping and combine the transport-layer port number to uniquely perform the mapping from IVI4(i) to IVI6(i).” *Id.* Li does not, however, propose or teach concatenating a port number. Instead, it teaches the use of zero bits for stateless mapping, one bit for two-state mapping, and four bits for 16-state mapping. *See id.* Beyond Li-2’s indication that

these multiplexing bits must appear in the least significant bits of the IPv6 address, Li-2 is silent regarding the implementation or operation of port multiplexing.

C. Wu (EX1039)


60. Wu is a Chinese Patent Application Publication titled “An automatic tunnel method capable of traversing network address translation.” EX1039 at 2. Wu’s solution is focused on “solv[ing] the communication problem between . . . two IPv6 nodes separated by the IPv4 network” in a manner that permits “NAT in the body of the tunnel.” *Id.* at 10. Wu refers to this solution as “Silkroad,” and the solution involves Silkroad clients that can communicate via a Silkroad tunnel with assistance from a Silkroad server. *See id.* at 10-11. Wu’s architecture is illustrated below:



EX1039 at 17.


61. Within Wu's Figure 4, the two endpoints of the Silkroad tunnel are the "Client end" devices, which are "IPv6/IPv4 dual-stack node[s] in the NAT field" EX1039 at 11. The "Server" corresponds to a Silkroad server outside of the NAT field. "[I]t assigns IPv6 address and the IPv6/IPv4 dual-stack node with the routing function for the client end" and "can have the route leading to the IPv6 backbone network, also the route leading to the isolated IPv6 network." *Id.* As shown in the figure, the client ends are behind a NAT function (i.e., in the NAT field) that is connected to the IPv4 Internet; thus, any tunnel constructed with that client as an endpoint will also have "NAT in the body of [the] tunnel." *Id.* at 10. The server is connected to the IPv4 Internet as well as IPv6 networks (illustrated as black clouds).

62. Wu's client ends communicate using the packet structure illustrated in Figure 2:



EX1039 at 17. This packet structure illustrates an IPv6-tunnel over IPv4; the IPv6 packets for transmission are included as IPv4 packet payload. The embedded IPv6 packets include at least one Silkroad address in the IPv6 header. The Silkroad address represents the IPv6 address assigned by the Silkroad server to the Silkroad

client end. See EX1039 at 11. The format of a Silkroad address is illustrated in Figure 1:



EX1039 at 17. The address includes an IPv6 prefix, a pseudo external port of the NAT function, a NAT type designation, a Silkroad address designation, and a pseudo external address of the NAT function. See *id.* at 11-12. The IPv6 prefix is “[t]he aggregate global unicast address prefix or the ‘site-local’ address prefix formally assigned by IANA” *Id.* at 11. The “Pseudo external port” is defined as “[t]he externa port (NAT-mapped port) and 0xFFFF and/or the subsequent 16-bit numerical value.” *Id.* at 11. The NAT type zone bit is 2-bits “used to represent the type of NAT where the client end is located” *Id.* The Silkroad address zone bit is defined as “14 bits . . . set as 0, representing that this address is a Silkroad address.” *Id.* The “Pseudo external address” is defined as “the address mapped by NAT and 0xFFFFFFFF and/or the subsequent 32-bit numerical value” *Id.* at 11-12. The Silkroad address, i.e. “an aggregate global unicast IPv6 address”, is used by the client ends to define the head node and end node of an IPv4 tunnel for transporting IPv6 packets. See *id.* at 10-11.

63. The goal of enabling communication between two nodes that are already modified to support IPv6 (as Wu purports to teach) is very different than the goal of providing IPv4 and IPv6 interconnection to hosts in a transparent manner such that hosts do not need to be modified. The terminal behind the NAT in Wu must have already been transitioned to a dual-stack node (supporting both IPv4 and IPv6) and must determine for itself whether to encapsulate and address a packet to the “SilkRoad” server or transmit it directly to the destination. EX1039 at 11-12. The terminal in Wu must also periodically send packets to the server to maintain its mapping in the NAT translation table. *Id.* at 11. This approach of offloading destination determination and connection management to the terminal (rather than handling it transparently at the router or other infrastructure) is an important design distinction that would tend to teach away from applying Wu’s teachings to an implementation both focused on operation that is transparent from the perspective of the terminal device and stateless.

XI. ISSUES COMMON TO ALL GROUNDS


A. Li does not disclose the claimed IPv4 and IPv6 domains

64. The independent claims of the ’844 Patent require “receiving a data packet from an *IPv4 domain* in an *IPv6 domain*” (’844 claims 1 and 5 (emphasis added)) or “sending a data packet from an *IPv6 domain* to a destination address of an *IPv4 domain*” (’844 claims 3 and 10 (emphasis added)). The sending and

receiving of packets between the domain types is possible because the '844 Patent enables the transparent (to the hosts) interconnection of an IPv4 domain and an IPv6 domain.

65. In all three grounds, Petitioner and Dr. Bhattacharjee point only to Li's Figure 7 to argue that Li discloses many of the requirements of the independent claims. Pet. at 28-29 (claim 1); *id.* at 39 (claim 2 referring back to claim 1 analysis); *id.* at 39-40 (claim 3); *id.* at 44 (claim 4 referring back to claim 3 analysis); *id.* at 44 (claim 5 referring back to claim 1 analysis); *id.* at 47-48 (claim 6); *id.* at 61 (claim 10 referring back to claim 3 analysis); EX1003 ¶¶87-88 (claim 1); *id.* ¶¶ 114 (claim 2 referring back to claim 1 analysis); *id.* ¶¶ 115-117 (claim 3); *id.* ¶¶ 125 (claim 4 referring back to claim 3 analysis); *id.* ¶¶ 126 (claim 5 referring back to claim 1 analysis); *id.* ¶¶ 131 (claim 6); *id.* ¶¶ 156 (claim 10 referring back to claim 3 analysis).

66. Dr. Bhattacharjee identifies IPv4 domains in red and IPv6 domains in orange in Li's Figure 7, shown below:




EX1003 ¶ 87. For claims 1 and 5, Dr. Bhattacharjee identifies the top left AS 4538 Access Network as the IPv4 domain from which an IPv4 packet originates. For claims 3 and 10, Dr. Bhattacharjee identifies the top right AS 24348 Access Network as the IPv4 domain to which an IPv4 packet is destined. See *id.* (illustrating the packet flow relied on by Dr. Bhattacharjee).

67. Under the proper construction of “domain” (see above Section IX.A), Dr. Bhattacharjee’s analysis does not show that Li discloses separate IPv4 and IPv6 domains, as required by the claim language. Instead, Li teaches that the networks identified by Dr. Bhattacharjee are all part of a single provider domain. Second, these domains all share common control and a common address space. Because of this common control, even under Dr. Bhattacharjee’s construction, the domains share a common administrator, i.e. the network provider, and thus form a single provider

domain. The '844 Patent is specifically directed to problems that arise from communications between a plurality of domains, and simply do not occur within a single domain.

B. The domains of Figure 7 are not separate domains


68. Section III of Li “illustrates the basic design, which is suitable for transition *in an intra-domain environment*.” EX1037 at 1 (emphasis added). Figure 1 presents an overview of the design, illustrating an autonomous domain comprised of two IPv4 subnets and an IPv6 backbone. *See id.* at 3. The domain is connected to the IPv4 Internet via a relay for routing packets to “IPv4 networks outside the autonomous domain”. *Id.* at 3. Annotated Figure 1, presented below, highlights Li’s domain.



Id. (highlighting Li's domain). Li labels this figure as an *intra-domain* configuration.

69. Dr. Bhattacharjee argues that Figure 7 discloses an interdomain environment, noting that Li includes figure 7 in its “Extensions to Inter-Domain environment” section. EX1003 ¶ 91. Li explains that the inter-domain environment includes “two models.” The first, includes stub autonomous systems connected to an IPv6 backbone. This is the embodiment described in Figure 1 and relied on by Dr. Bhattacharjee. The second is a “transit autonomous systems [that] are interconnected with each other.” EX1037 at 6. This model describes the inter-domain connections described in Figure 8-10 in that section. Li does not teach that Figure 7 depicts an inter-domain configuration. So, this embodiment (sending a packet from AS 4538 to AS 24348) is still an *intra-domain* setting.

70. Figure 7 the Dr. Bhattacharjee points to depicts the same network structure as Figure 1, including the same IP address distribution.




Id. at 7. Despite being in Li’s Extensions to Inter-Domain environment” section, Figure 7 (and specifically the embodiment Dr. Bhattacharjee points to) is still an intra-domain transmission. In that section, Li describes “two model: 1) stub autonomous systems are connected to an IPv6 backbone, as customer networks of the latter, and only have accesses to IPv4 hosts within the scope controlled by the provider; 2) transit autonomous systems are inter-connected with each other.” EX1037 at 6. This first model, depicted by Figure 7, is still an intra-domain embodiment. The primary difference between Figure 1 and Figure 7 is that “stub autonomous systems are connected to an IPv6 backbone, *as customer networks of the latter*, and *only have accesses to IPv4 hosts within the scope controlled by the*

provider” Id. at 6 (emphasis added). The provider controls the networks in Figure 7, just as in Figure 1.

C. Figures 8, 10, and 11 show inter-domain embodiments


71. Figures 8, 10, and 11 highlight that Figures 1 and 7 are intra-domain embodiments. Figures 8, 10, and 11 illustrate inter-domain deployments. Figure 8 depicts an inter-domain environment in which “two providers are inter-connected by native IPv6 link” EX1037 at 8.




Id. (highlighting added). In order to support the inter-domain configuration, Li explains that it is “trying to extend the address mapping scheme and the functionality of the relay router” *Id.* The relay routers can be seen connecting the two

provider domains (marked with R and highlighted orange), and relay routers are also used to connect the two provider domains to the IPv4 Internet. *See id.* at 9 (“[A] **relay router** is just a regular BGP speaker for both IPv6 and IPv4, **standing between two provider domains.**” (emphasis added)). Li teaches that its relay routers connect domains, and Dr. Bhattacharjee has not relied on communications through Li’s relay routers in its contentions.

72. Figure 10 illustrates the *partial deployment* of an *inter-domain environment*.



Id. at 9 (highlighting added). And Figure 11 illustrates “a simplified version for our *inter-domain solution* that is shown in Fig. 8.” *Id.* (emphasis added).



Id. Dr. Bhattacharjee does not identify any of the inter-domain figures as the basis of its invalidity grounds for the independent claims. In fact, Dr. Bhattacharjee repeatedly acknowledges that the communication scenario it is alleging teaches the claims is limited to the communications between the two IPv4 access networks contained in the single-provider domain illustrated in Figure 7. *See* EX1003 ¶ 62 (depicting the communication path at issue with a large black arrow); *id.* ¶ 87 (same). But a POSITA would understand that Figure 7 describes an intra-domain environment as illustrated in the later figures. Dr. Bhattacharjee has relied on an intra-domain configuration to show inter-domain communications.

D. The autonomous systems of Figure 7 are under common control

73. In addition to Li's explicit teachings that the Petitioner's identified IPv6 domain and IPv4 domain are part of the same provider domain, Li explains that these

domains are under common control of a single Internet service provider. See EX1037 at 2 (“Through out this paper, it is assumed that an Internet service provider who would like to apply this approach holds at leas[t] one IPv6 prefix not smaller than /32. A part of a /32 prefix is utilized to do the address mapping.”).

74. For instance, when describing the address mapping Dr. Bhattacharjee relies on (see EX1003 ¶¶ 96-97, 218), Li explains: “This rule of addressing is deployed without difficulty because the address assignment is fully controlled by the autonomous system itself, which is aware of all the IPv4 prefixes under its administration.” EX1037 at 3 (emphasis added). This address mapping is introduced for Li’s Figure 1, but it also applies to Li’s Figure 7, which has the same network configuration and the same IPv4 prefix assignments as Figure 1. *Compare* EX1037 at 3 (Figure 1), *with id.* at 7 (Figure 7); *id.* at 7 (“The addressing scheme is similar to that in the basic design.”); *id.* (“[T]he mapped prefixes belong to the aggregation of the provider”); *id.* at 4 (“[E]ach edge should know a complete list of IPv4 network identifiers, which really belong to the autonomous domain.” (emphasis added)).

75. Also, when Li details Figure 7, it explains that the “stub autonomous systems are connected to an IPv6 backbone, as customer networks of the latter, and only have accesses to IPv4 hosts within the scope controlled by the provider” *Id.* at 6. Li explains that the provider controls the IPv4 addresses subject to mapping

at the edge routers of Figure 7. *See id.* at 7 (“In the previous subsections, we employ a special prefix to avoid mapping the IPv4 route table explosion into IPv6. Only those IPv4 network identifiers that belong to the provider are mapped correspondingly while the others are merged into the special prefix.” (emphasis added)). The addresses outside the scope of the provider’s control use the special prefix. *See id.* (“It is also defined that a special prefix appended with ‘0xFFFFFFFF’ is correspondent to network identifiers elsewhere out of the scope that the provider serves.”); *id.* (“Meanwhile, the relay router advertises the special prefix with ‘0xFFFFFFFF’ for all the IPv4 networks outside the scope of the provider’s administration.”)

76. Based on the teachings of Li, the IPv4 and IPv6 networks of Figures 1 and 7 are commonly controlled, and thus belong to the same domain. All three of Dr. Bhattacharjee’s grounds fail to identify the IPv4 domain and IPv6 domain required of every independent claim. Instead, Dr. Bhattacharjee has relied on networks within the same domain.

XII. GROUND 1

A. Dr. Bhattacharjee’s opinion does not support the conclusion that Li-2 was publicly available

77. I understand that a document was publicly available when it was “disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art exercising reasonable diligence, can

locate it.” *SRI Int’l, Inc. v. Internet Sec. Sys., Inc.*, 511 F.3d 1186, 1194 (Fed. Cir. 2008).” For the reasons below, Dr. Bhattacharjee brief analysis of IETF practices does not establish that Li-2 was publicly available.


78. Dr. Bhattacharjee contends that Li-2 was publicly accessible based on his understanding that IETF Internet Drafts “are publicly published on the IETF website.” EX1003 ¶ 67; *see also id.* ¶¶ 68-69. However, based on Dr. Bhattacharjee’s description of IETF practices, it is not clear that Li-2 was publicly available. Li-2 appears to be a cancelled and unpublished Internet Draft. *See* 1038 at 1; EX1044 at 13 (listing the document as cancelled in the file name, “draft-xli-ivi-00-7431-cancelled.txt”). Dr. Bhattacharjee does not discuss how such unpublished documents are made available, if at all. Moreover, Dr. Bhattacharjee contends that Li-2 was submitted to the IETF no earlier than June 21, 2008, only 9 days before the priority date of the ’844 Patent. He has not provided any explanation for whether a POSITA would be able to locate an unpublished Internet Draft in such a short timeframe.

B. A POSITA would not have been motivated to combine Li and Li-2 as proposed

1. Li and Li-2 address different problems

79. Li addresses the problem of allowing IPv4 hosts within an IPv4 network to communicate with other IPv4 hosts across an IPv6-only network or backbone.

This is shown in Dr. Bhattacharjee's annotated Figure 7 from Li, illustrating the communication path between two IPv4 networks traversing an IPv6-only network.



EX1003 ¶ 87.

80. Li-2, on the other hand, addresses the problem of enabling hosts within an IPv6-only network to communicate with IPv4 hosts. *See* EX1038 at 2 (“Based on [IVI’s address borrowing] scheme, special IPv6 address blocks are defined, which can communicate with the global IPv6 networks directly and **with the global IPv4 networks** via stateless (or almost stateless) gateways.” (emphasis added)). So, Li and Li-2 are focused on solving different problems regarding communication within the context of IP networks. While Li targets IPv4 host to IPv4 host communications via an IPv6 network, Li-2 targets IPv6 host communications to IPv4 hosts. These are

different problems, and Dr. Bhattacharjee has not explained how their solutions (or even components of their solutions) are interchangeable.

2. *Li-2 is not an extension of Li*

81. Dr. Bhattacharjee's motivations to combine largely rest on its incorrect contention that Li-2 is an extension of Li. Regarding the first motivation, Dr. Bhattacharjee states: "A POSITA would have had a reasonable expectation of success in combining Li and Li-2, especially because a POSITA did succeed in combining these references, since Li-2 is an extension of Li" EX1003 ¶ 79. Regarding the second motivation, Dr. Bhattacharjee states: "A POSITA not only could have applied the known improvement technique from Li-2 to Li, but a POSITA did apply this technique, and the results would therefore not only have been predictable to a POSITA but are explicitly described in Li-2." *Id.* ¶ 81. Regarding the third motivation, Dr. Bhattacharjee states: "And the results would have been predictable, as Li-2 describes the results of the substitution." *Id.* ¶ 82. And finally, regarding the fourth motivation, Dr. Bhattacharjee states: "Li-2 is an extension of the ideas described in the original Li reference" *Id.* ¶ 83.

82. But Li-2 is not an extension of Li. As detailed above, Li and Li-2 address different problems. Furthermore, Li-2 never references the Li paper, and it never suggests that Li-2 is an extension of Li. The reason is simple, Li-2 is not an extension of Li. Rather, Li-2 addresses a problem that was not mentioned by Li, i.e.

enabling IPv6 hosts to communicate with IPv4 hosts. Li, on the other hand, does not focus on IPv6 hosts at all, rather, it focuses on IPv4 hosts communicating across an IPv6 transport network. Dr. Bhattacharjee's incorrect characterization of Li-2 as an extension of Li is an attempt to provide a motivation to combine where none exists.

3. *Dr. Bhattacharjee relies on motivation to combine references that are not prior art*

83. Dr. Bhattacharjee also relies on what he calls “express teachings [in the prior art] that would have led a POSITA to combine Li and Li-2,” referring to the his “Technology Background” section for support. *See* EX1003 ¶ 79 (citing § III.C.1). Petitioner relies on that section to assert that “[m]ultiple prior art references disclose converting IPv4 addresses into IPv6 addresses” Pet. at 7 (citing EX1003 ¶¶38-39, discussing EXS1019-1020). However, Dr. Bhattacharjee's analysis does not support Petitioner's assertion. Dr. Bhattacharjee discusses two references in the those paragraphs: Samprathi (EX1019) and Rooney (EX1020). But Petitioner mischaracterizes Rooney as prior art. Rooney says on its face that “[t]his is an unpublished work protected under the copyright laws.” EX1020 at 35.

84. Dr. Bhattacharjee then states that “[b]y the priority date of the '844 patent, POSITAs understood that computers typically utilize only a small portion of their available port numbers.” EX1003 ¶ 41. In that paragraph, Dr. Bhattacharjee discusses EX1023 which is purportedly a paper presented at IEEE GLOBECOM

2008,¹ which occurred between November 30 and December 4, 2008, i.e. well after the priority date of the '844 Patent, i.e. June 30, 2008. See EX2009 (Conference Proceedings Title Page). And the paper states that it is “for publication” in the conference proceedings: “This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts *for publication* in the IEEE ‘GLOBECOM’ 2008 proceedings.” EX1023 at 5.

85. Finally, Petitioner states “the prior art teaches constructing an IPv6 address by concatenating an IPv6 prefix, an IPv4 address, and a port number.” Pet. at 8 (citing EX1003 ¶¶45-48, discussing EXS1026-1032, 1038). In those paragraphs, Dr. Bhattacharjee discusses Li-2, which is not prior art, as detailed above. Further, he discusses EX1026, which Petitioner has not shown was publicly available as of the priority date. And finally, Dr. Bhattacharjee relies on documents that he admits post-date the '844 Patent's priority date to further support his conclusions about what was purportedly known in the art. See EX1003 ¶46 (discussing EX1027 dated July 6, 2008; EX1028 dated July 29, 2008; and EX1029 dated September 2008); id. ¶170 (admitting EX1029 was published in September 2008). Since it has not been shown that these

¹ IEEE GLOBECOM refers to the IEEE Global Telecommunications Conference. See <https://globecom2008.ieee-globecom.org/>.

86. These improper references are prior art, Dr. Bhattacharjee’s conclusions here are not reliable.

4. *The prior art identifies NAT tables as a well-known solution for IPv4 address exhaustion*

87. Dr. Bhattacharjee also relies on the notion of IPv4 address exhaustion within its first two motivations to combine. *See* EX1003 ¶ 79 (explaining the “need to slow IPv4 address exhaustion by sharing a single public IPv4 address among multiple hosts.”); *id.* ¶ 80 (explaining that Li “does not solve the issue of IPv4 address exhaustion”). However, the ’844 Patent explains that NAT tables provided a well-known solution for addressing IPv4 address exhaustion. *See* EX1001 at 2:11-14 (“To limit the number of IPv4 public addresses necessary to provide an IP connectivity service to an installed base of clients, a solution known as Double NAT or Operator NAT has been proposed and implemented.”); *see also id.* at 1:51-60 (discussing different types of NAT).

88. The prior art relied on by Dr. Bhattacharjee confirms that this was an accepted approach to addressing IPv4 exhaustion. For example, he identifies Yeom (EX1021) for its teaching “there are a lot of ports in a node and only a small portion of them are used at any time.” EX1021 at 4 (cited by EX1003 ¶ 41). Yeom’s technique is to use port-address mapping tables. *See id.* at 5 (“Table 1: An example of port-address mapping table”); *id.* (Figures 4 and 5 depict a “PAT Mapping Table”). Dr. Bhattacharjee also relies on Rietz (EX1022) for the unremarkable

teaching of NAT table capabilities. *See* EX1003 ¶ 41 (citing EX1022 at 9); *see also* EX1022 at 2 (“In the current Internet, Network Address Translation (NAT) gateways that provide port address translation are quite popular. These gateways allow many hosts to be multiplexed on one single IP address and still maintain full outbound connectivity.”). The same is true for King (EX1024). *See* EX1024 at Fig. 4; *id.* at 4:16-17 (“Fig. 4 is a table showing source address and port address translation at a CPE [customer premises equipment]”).

89. Dr. Bhattacharjee does not provide a motivation to deviate from this known solution and form IPv6 addresses in the manner claimed by the '844 Patent. Instead, he use hindsight bias to leap from a desire to solve IPv4 address exhaustion to implementing the claimed solution of the '844 Patent. This attempt to rely on references that post-date the '844 Patent invention to supply a missing limitation show that the '844 Patent solution was not well-known. None of the references identified by Dr. Bhattacharjee address the problem addressed by the '844 Patent inventions, namely providing stateless solutions for routing data packets across domains using different versions of IP networks. EX1001 at Abstract, 1:15-21.

XIII. GROUND 2

A. Dr. Bhattacharjee recognizes that Li fails to teach port number limitations

90. Every independent claim requires either “constructing an IPv6 destination address by concatenating . . . [the/said] destination port number” or

“extracting . . . the destination port number from the IPv6 destination address.” Dr. Bhattacharjee admits that Li does not teach forming the claimed IPv6 address for failure to concatenate a destination port number. *See, e.g.*, EX1003 ¶ 179 (“Li teaches limitation 1[a] except for concatenating a destination port number.”). And Petitioner admits that Li fails to teach extracting the destination port number from such an IPv6 address. *Id.* at 65 (“Li teaches limitation 3[a] except for extracting the destination port number from the IPv6 address.”); *see also* Pet. at 66-67 (similar for other claims). To remedy these deficiencies, Dr. Bhattacharjee attempts to rely on evidence of simultaneous invention. *See* EX1003 ¶¶ 164-170.

B. The evidence does not show simultaneous invention

91. Dr. Bhattacharjee’s simultaneous invention argument fails to show that any of the cited evidence constitutes the claimed invention. Dr. Bhattacharjee relies on five exhibits to show simultaneous invention: (1) Li-2; (2) IVI Memo; (3) IVI Presentation; (4) Bush Paper; and (5) Bush Presentation. *See* EX1003 ¶¶ 164-170. Specifically, Dr. Bhattacharjee states that “evidence of simultaneous invention indicates that a POSITA would have been motivated to modify Li to form an IPv6 address that includes the destination port number.” *Id.* ¶ 164. However, Dr. Bhattacharjee narrows the invention to “form[ing] an IPv6 address that includes the destination port.” *Id.* This oversimplification of the ’844 Patent claims ignores numerous limitations that must be addressed to show simultaneous invention.

92. With one exception, Dr. Bhattacharjee relies on Li-2 only for the purpose of satisfying the “destination port number” of the claims. *See* EX1003 ¶ 95 (limitation 1[a] for “destination port number”); *id.* ¶ 115 (limitation 3[pre] for “destination port number”); *id.* ¶ 118 (limitation 3[a] for “destination port number”); *id.* ¶ 121 (limitation 3[b] for “destination port number”). The exception concerns claim 9, where Dr. Bhattacharjee relies on Li-2 to teach advertising IPv6 addresses purportedly containing port numbers. *See id.* ¶¶ 154-155. The remaining references to Li-2 in Dr. Bhattacharjee’s limitation-by-limitation analysis merely refer to the previous claim limitations regarding specifically the “destination port number” limitation. *See id.* ¶ 128 (limitation 5[a]); *id.* ¶ 136 (limitation 6[c]); *id.* ¶ 138 (limitation 6[d]); *id.* ¶ 153 (claim 9).

93. Dr. Bhattacharjee does not allege that Li-2 teaches any of the following limitations from independent claims 1 and 3:

- Claim 1: “generating an IPv6 data packet from the IPv6 constructed destination address and the received IPv4 packet,” EX1001 at 14:36-37;
- Claim 1: “routing the generated IPv6 data packet in the IPv6 domain using the IPv6 constructed destination address, said constructed address belonging to a range of IPv6 addresses routable to an interconnection equipment of the IPv6 domain with the IPv4 destination address,” EX1001 at 14:38-42;

- Claim 3: “extracting the IPv4 destination address . . . from the IPv6 destination address,” EX1001 at 14:55-56;
- Claim 3: “generating an IPv4 data packet from the IPv6 data packet, said packet including said IPv4 destination address . . . ,” EX1001 at 14:57-59²;
- Claim 3: “routing the IPv4 data packet to its IPv4 destination,” EX1001 at 14:60.

94. And aside from claim 9 (mentioned above), Dr. Bhattacharjee does not rely on Li-2 to show the limitations of any other challenged claims except by reference back to his port number allegations in claims 1 and 3. *See generally* EX1003 ¶¶ 86-155.

95. Dr. Bhattacharjee’s reliance on Li-2 to show simultaneous invention should be ignored because he has not meaningfully addressed any claim limitation beyond “destination port number.”

96. Similarly, for the remaining four exhibits purportedly showing simultaneous invention, Dr. Bhattacharjee only addresses the notion of forming an IPv6 address that includes a port number. *See* EX1003 ¶¶ 164-170. For the IVI Memo (EX1027), Dr. Bhattacharjee only states that it “discloses using ‘port multiplexing’ to construct an IPv6 address with a destination port number.” *Id.* ¶ 166. For the IVI presentation (EX1028), Dr. Bhattacharjee only references port multiplexing to show

² The ellipsis omissions relate specifically to “destination port number” claim language.

combining an address with port number. *id.* ¶ 167. Regarding the Bush paper (EX1029), Dr. Bhattacharjee states that Bush teaches “forming an IPv6 address by concatenating an IPv6 prefix, IPv4 address, and port number.” *Id.* ¶ 170. Finally, for the Bush presentation (EX1041), Dr. Bhattacharjee again asserts that Bush teaches forming an IPv6 address by combining a well-known prefix, IPv4 address, and port number. *Id.* As with Li-2, Dr. Bhattacharjee only focuses on the format of an IPv6 address without regard for any of the other required claim limitations.

XIV. GROUND 3

A. A POSITA would not have been motivated to combine Li and Wu

97. Dr. Bhattacharjee states that a POSITA would combine Li and Wu to render obvious the “destination port number” elements of the claims of the ’844 patent. I disagree for several reasons. First, Li and Wu address different IP routing problems. Second, Li is stateless, but Wu is stateful. Any combination of the two would eliminate the benefits of Li’s stateless system. Finally, even if a POSITA would combine Li and Wu, Wu does not teach the claimed destination port number.

1. *Li and Wu Are Directed To Different Communication Environments*

98. As I explained above, Li addresses the problem of allowing IPv4 hosts in IPv4 networks to communicate with each other across an IPv6 network. On the other hand, the goal of Wu is to enable IPv6/IPv4 dual-stack nodes located behind a network address translation (“NAT”) router to communicate with other IPv6 nodes

by using a “Silkroad” tunnel across an IPv4 network. EX1039 at 2 (abstract), 10-11.

Wu focuses primarily on traversing NAT in the IPv4 network. EX1039 at 9-11.

99. The goal of enabling communication between two nodes that are already modified to support IPv6 (as Wu purports to teach) is very different than the goal of providing IPv4 and IPv6 interconnection to hosts in a transparent manner such that hosts do not need to be modified. In fact, these two problems are almost directly opposite in nature. The client terminal in Wu must have already been transitioned to a dual-stack node (supporting both IPv4 and IPv6) and must determine for itself whether to encapsulate and address a packet to the “SilkRoad” server or transmit it directly to the destination. *See* EX1039 at 11-12. The terminal in Wu must also periodically send packets to the server to maintain its mapping in the NAT translation table. *See id.* at 11. This approach of offloading destination determination and connection management to the terminal (rather than handling it transparently at the router or other infrastructure) is an important design distinction that teaches away from applying Wu’s teachings to an implementation more focused on operation that is transparent from the perspective of the terminal device. Furthermore, the mapping maintained by Wu is by definition stateful, otherwise it would not need to be maintained.

100. Because of these differences, a POSITA would not have been motivated to combine the references. First, Li is concerned with communications *between* IPv4

hosts; however, Wu is concerned with communications *between* IPv6 hosts. Second, Li is concerned with communications that traverse an IPv6-only network; in contrast, Wu, is concerned with communications that traverse an IPv4-only network that also includes a NAT field. Dr. Bhattacharjee's contention that a POSITA would have been motivated to combine these references treats IPv4 networks and IPv6 networks as if they are mix-and-match components. However, the very problem with IPv4-IPv6 transition technologies is that the two types of networks are inherently incompatible. To suggest that an IPv4-focused tunneling solution can simply be transplanted to an IPv6 transit domain solution, as Petitioner does, ignores this inherent incompatibility.

2. Li is stateless, but Wu is stateful

101. Stateless solutions represented a key advantage over prior art prior IP transition techniques. In a stateful solution, border or edge devices must maintain the “state” of each connection or session within their networks. As the '844 Patent explains, “A solution based on NAT-PT that is not stateless is not to be recommended” and “would degrade the service provided”:

To summarize, the invention makes possible a first phase of migration of IP connectivity services towards IPv6 using packet routing mechanisms that takes the following factors into account:

Service operators suffer a problem of running out of IPv4 addresses.

The final migration to IPv6 will take a plurality of years (10 years minimum), mainly because of the large number of participants to “convince” to go over to IPv6, the large number of autonomous

systems (AS), and the diversity of interconnection mechanisms. Moreover, it should be pointed out that the address problem is a priori of concern only to operators. Clients have no reason to modify the architecture of their local networks, in particular its FTP, HTTP, etc. servers, to IPv6. Operators must expect a long haul to assist clients with the changeover to IPv6.

Simply migrating a given domain to IPv6 does not solve the problem of global connectivity (contacting any remote machine present on the Internet). Interconnection with the IPv4 world must therefore be provided.

A solution based on NAT-PT that is not stateless is not to be recommended. This would degrade the service provided as value-added services linked to port forwarding configured by clients would not work.

'844 Patent 13:23-51 (emphasis added). The ability to implement stateless translation and encapsulation represents a significant advance over the prior approaches to routing IPv4 data packets over an intermediate IPv6 domain.

102. Wu describes what it calls a “Silkroad Tunnel.” Tunnels are typically stateful. Figure 1 of Wu shows “the IPv6 address adopted when the client end communicates.” EX1039 at 3 (Claim 3. The Brief Description of the Drawings states that Figure 1 is the form of “SilkRoad address **assigned to the client end** by the server.” *Id.* at 11. And Wu must maintain mapping of the several ports and addresses. *Id.* at 11-12. Thus, Wu is a stateful system.

103. On the other hand, Li is a stateless solution. Li explains that an important aspect of Li’s disclosure is the stateless nature of the system. EX1037 at 3 (“With the address-mapping rule commonly known within the autonomous

domain, IPv4 source and destination addresses in a packet can be translated into their IPv6 version without any explicitly stored states . . .”). This approach allows packets to be transmitted without the overhead required to initialize a connection, maintain and update information about the connection (e.g. Wu’s mapping tables), and modify incoming and outgoing packets.

104. A POSITA would not attempt to combine a stateless system with a stateful one. Any resulting system (if one could be created) would result in a stateful system that would remove the benefits of the stateless system described above. As a result, a POSITA looking to modify Li would never look to Wu.

3. *Wu does not teach an a destination address constructed with a destination port number*

105. Wu also does not teach the claimed destination port number of the ’844 Patent. As Dr. Bhattacharjee acknowledges, Wu’s “NAT [] translates a source port number into an external port number.” EX1003 ¶ 205 (emphasis added). Dr. Bhattacharjee does not acknowledge this difference between the claims and Wu, nor does he explain how or why Wu could be modified to teach a destination port.

106. Figure 1 of Wu shows the IPv6 address adopted when the client end communicates. EX1039 at 3 (Claim 3), 17. The Brief Description of the Drawings in Wu states that Figure 1 is the form of “SilkRoad address assigned to the client end by the server.” *Id.* at 11. The specification states that “The IPv6 address adopted by the client ends during communication is a kind of address with prescribed form” that

includes a 64-bit IPv6 address prefix and a 64-bit interface ID. *Id.* Wu indicates that the address shown in Figure 1 is the address of the client, and thus it will be used as the source address in communications from the client. This is confirmed by the discussion of the IPv4 header, which indicates that “The source IPv4 address is the IPv4 address of the client end.” *Id.* at 12. This SilkRoad address is generated and assigned to the client as part of a registration process that occurs before it conducts transmissions on the network. *Id.* Thus, the IPv6 address assigned to the client occurs in a preliminary operation. It is not an IPv6 address that is constructed in response to receiving an IPv4 data packet.

107. The IPv6 address of Figure 1 includes a pseudo external port field that appears to be a NAT-mapped port value (the private IPv4 address of the client has been mapped to a public “external” address). This field reflects the value that the “client end sends [in] an initialized data packet to the server”. The UDP port value of the client is mapped to a different port when the packet passes the NAT, and this value is used to generate the IPv6 address as part of the initialization configuration of the client. EX1039 at 11-12. The NAT translation table maintains the mapping of the client port and the external port. *Id.* at 3, 11. The pseudo port field in Figure 1 is at best the source port of the client, not a destination port.

108. The IPv6 address in Figure 1 also includes a pseudo external address. This field appears to reflect the value of the client’s private IPv4 address, mapped to

a public IPv4 address when the packet passes the NAT. EX1039 at 12. The NAT table maintains a state table of the mapping between the private IPv4 address and the public (external) IPv4 address. *Id.*

109. I also note that a POSITA would not understand “pseudo external port” or “pseudo external address” to be terms of art, and notably Dr. Bhattacharjee made no attempt to explain what a “pseudo external port” or “pseudo external address” is. I conducted a Google search and found no references to either term, matching my own understanding of the art. Wu’s characterization of a pseudo external port is “[t]he external port (NAT-mapped port) and 0xFFFF and/or the subsequent 16-bit numerical value.” EX1039 at 11. However, a POSITA would never implement an “external port (NAT-mapped port) and 0xFFFF and/or the subsequent 16-bit numerical value.” A POSITA would understand that 0xFFFF in binary would be a 16-bit number consisting entirely of ones. If a POSITA treated the definition from Wu as a suggested bit manipulation, an “and” bitwise operation between 0xFFFF and any 16-bit number simply results in the same value as the original number. An “or” bitwise operation between 0xFFFF and any 16-bit number always results in the value 0xFFFF. Neither operation would make sense to a POSITA, doing extra work to always get back either the same value or 0xFFFF is simply illogical. Based on the definition in Wu, a POSITA would not necessarily understand a “pseudo external port” to even be a port number, much less understand from Wu how a destination

port number should be used when constructing an IPv6 destination address. The same logic applies to the “pseudo external address,” which a POSITA would also not necessarily understand to be an IPv4 destination address as claimed. EX1039 at 12.

110. As these discussions indicate, at best the pseudo external address and pseudo external ports represent source information of the client that is configured as part of an initialization of the client device. These external source and port values do not describe destination address and port information. EX1039 at 3, 11-12.

111. The server in Wu also does not perform IPv6 destination address construction as specified in the '844 claims. The '844 Patent claims refer to constructing an IPv6 destination address by concatenating certain parameters from a received IPv4 data packet, including the IPv4 destination address and port information included in the received IPv4 data packet. This concatenation operation that forms part of the address constructing process requires an operation in which each of these pieces of information are joined together to form a single unit in a contiguous pattern. The server in Wu does not perform any sort of concatenating operation, or otherwise append any of the discrete values required into an IPv6 address. Rather, it simply strips the IPv4 header from the packet and forwards the un-encapsulated IPv6 data packets. EX1039 at 15. The server itself does not construct an address and removing an encapsulating header or copying an IPv6

address is different than constructing an address through a concatenation operation.

As a result, the operation performed in Wu of retrieving the encapsulated IPv6 address also does not meet the requirement of the constructing limitation because no concatenation occurs at the server.


XV. CONCLUSION

112. For at least the reasons set forth herein, it is my opinion that Petitioner has failed to establish a reasonable likelihood that any of the claims of the '844 Patent are unpatentable over the cited prior art.

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

Executed on October 30, 2024, in Kirkland WA

Respectfully Submitted,



William Henry Mangione-Smith