

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

APPLE INC.,
Petitioner

v.

SONRAI MEMORY LIMITED
Patent Owner

Patent No. 6,874,014

DECLARATION OF DR. MARTIN WALKER

I.	INTRODUCTION	1
II.	BACKGROUND AND QUALIFICATIONS.....	1
III.	MATERIALS CONSIDERED AND SUMMARY OF OPINIONS.....	3
IV.	PERSON OF ORDINARY SKILL IN THE ART	4
V.	TECHNICAL BACKGROUND	5
A.	Role of operating systems	6
B.	Configuring and operating systems with multiple processors	8
C.	Multiprocessor system running multiple operating systems concurrently	10
D.	Sharing multiprocessor resources.....	13
E.	Adoption of single-chip multiprocessors	14
VI.	THE '014 PATENT	20
VII.	OVERVIEW OF THE PRIOR ART	23
A.	Overview of Gulick (EX1005).....	23
B.	Overview of AAPA (EX1001).....	28
C.	Motivation to Modify Gulick Based on AAPA	29
D.	Motivation to Modify AAPA Based on Gulick	31
VIII.	CLAIM CONSTRUCTION	32
IX.	THE PRIOR ART DISCLOSES ALL OF THE LIMITATIONS OF CLAIMS 1, 3, 5, 9, 11–13, AND 15 OF THE '014 PATENT.....	35
X.	CONCLUSION.....	107

I, Martin Walker, declare as follows:

I. INTRODUCTION

1. I have been retained by Apple Inc. (“Petitioner”) as an independent expert consultant in this proceeding before the United States Patent and Trademark Office (“PTO”) regarding U.S. Patent No. 6,874,014 (“the ’014 patent”) (EX1001). I have been asked to consider whether certain references disclose or suggest the limitations recited in claims 1, 3, 5, 7, 11-13, and 15 (“the challenged claims”) of the ’014 patent. My opinions are set forth below.

2. I am being compensated at my rate of \$650 per hour for the time I spend on this matter. My compensation is in no way contingent on the nature of my findings, the presentation of my findings in testimony, or the outcome of this or any other proceeding. I have no other interest in this proceeding.

II. BACKGROUND AND QUALIFICATIONS

3. In this section, I have summarized my educational background, career history, publications, and other relevant qualifications. My Curriculum Vitae, which states my qualifications more fully, is filed as a separate Exhibit (EX1003).

4. I received my Bachelor of Science degree in Electrical Engineering from the Massachusetts Institute of Technology in 1973. I received my Master of Science and Doctor of Philosophy degrees in Electrical Engineering in 1976 and 1979 from Stanford University.

5. Accordingly, I have over 40 years of experience with writing and analyzing computer software as well as computer system architecture.

6. Beginning in 1981, I developed computer software to aid in the development of integrated circuits. From 1983 to 1989, I was the founder and Chief Technical Officer at Analog Design Tools, Inc. Analog Design Tools' primary product was software used to design integrated circuits. I was responsible for all technical aspects of product definition and development, including the software design.

7. From 1990 to 1994, I was a founder and Executive Vice President of Symmetry Design Systems, which specialized in software product design and consulting for the electronic design marketplace.

8. In 1995, I founded a company called Frequency Technology (later called Sequence Design) that developed software for the design of advanced system-on-a-chip integrated circuits. As Chief Executive Officer, director, and Chief Scientist at Sequence, I was involved in overseeing the development of the company's products and technologies.

9. From 2000 to present, I have worked in the fields of market research and consulting in the fields of semiconductor (IC) design, including electronic design automation and field programmable gate arrays (FPGA) programming. Since that time, I have served as an expert in nearly 100 cases.

10. I am a named inventor on three patents in the fields of integrated circuit design: U.S. Patent Nos. 6,643,831, 6,381,730, and 5,901,063.

III. MATERIALS CONSIDERED AND SUMMARY OF OPINIONS

11. The opinions contained in this Declaration are based on the documents I reviewed and my professional judgment, as well as my education, experience, and/or knowledge regarding technologies relating to, among other things, networking technologies.

12. In forming my opinions expressed in this Declaration, I reviewed the following materials (and any other materials that I reference in this Declaration):

Exhibit	Description
1001	U.S. Patent No. 6,874,014
1003	Curriculum Vitae of Dr. Martin G. Walker
1004	Prosecution History of U.S. Patent No. 6,874,014
1005	U.S. Patent No. 6,314,501 to Gulick et al. ("Gulick")
1007	Joch, "HOW-TO Chip Multiprocessing," Computerworld, November 27, 2000
1008	Olukotun et al., "The Case for a Single-Chip Multiprocessor," ACM SIGPLAN Notices, Vol. 31, Issue 9, Sept. 1996
1009	U.S. Patent No. 6,550,020 to Floyd et al. ("Floyd")
1010	Barroso et al., "Piranha: A Scalable Architecture Based on Single-Chip Multiprocessing", Proc. 27th ACM Int'l Symp. Computer Architecture, 2000
1011	U.S. Patent No. 6,772,241 to George et al.
1012	U.S. Patent No. 6,526,462 to Elabd
1013	Hammond et al., "A Single-Chip Multiprocessor," Computer, vol. 30, no. 9, pp. 79-85, Sept. 1997
1014	U.S. Patent No. 6,178,503 to Madden et al.
1015	U.S. Patent No. 6,085,247 to Parsons, Jr. et al.
1016	U.S. Patent No. 6,247,109 to Kleinsorge et al.
1017	U.S. Patent No. 6,397,242 to Devine et al.

13. In support of my opinions, I have taken into account how a person of ordinary skill in the art (as defined below in Section IV) would have understood the claims and the specification of the '014 patent at the time of the alleged invention. My opinions take into account how a person of ordinary skill in the art would have understood the '014 patent, the prior art to the patent, and the state of the art at the time of the alleged invention.

14. As I discuss in detail below, it is my opinion that certain references disclose or suggest all the limitations recited in claims 1, 3, 5, 7, 11-13, and 15 of the '014 patent.

IV. PERSON OF ORDINARY SKILL IN THE ART

15. I have been asked to consider the time of the alleged invention for the '014 patent to be the early 2001 time frame, including and up to May 29, 2001, which is the filing date of an application associated with the '014 patent. (EX1001, Cover). I may refer to this time frame in this declaration as the time of the alleged invention, the relevant time, or the time of the '014 patent.

16. Based on my knowledge and experience, I understand what a person of ordinary skill in the art would have known at the time of the alleged invention of the '014 patent. In my opinion, based on my review of the '014 patent, such a person would have had an undergraduate degree in electrical engineering, computer

engineering, computer science or a related field along with at least two years of work experience in the field of computer processor architecture. (*See, e.g.*, EX1001, 1:11–2:9 (discussing the background of the '014 patent), 2:12–44 (discussing the summary of the disclosure).) Additional education could substitute for experience in the field, and vice versa.

17. My analysis of the '014 patent and my opinions in this declaration are from the perspective of one of ordinary skill in the art, as I have defined it above, at the time of the alleged invention for that patent.

V. TECHNICAL BACKGROUND

18. In this section, I discuss the state of the art relevant to the subject matter of the '014 patent as would have been known, understood, and appreciated by a person of ordinary skill in the art prior to and at the time of the alleged invention for the '014 patent (e.g., early 2001 time frame, including and up to May 29, 2001, as I mentioned above in Section IV). I rely on and incorporate the understandings below regarding such known features and technologies to support my opinions regarding the knowledge of a person of ordinary skill in the art at the relevant timeframe and the motivations relating to how the prior art discloses and/or suggests the limitations of claims 1, 3, 5, 7, 11–13, and 15 of the '014 patent. In my opinion, the understandings regarding the technologies, processes, and features as I discussed in this section (Section V) would have been in the mind of a person of

ordinary skill in the art when contemplating the disclosures and/or suggestions in the prior art that I address below in Sections VII and IX and when considering combining features from that prior art in the manner that I explain below in that section.

A. Role of operating systems

19. It was generally known at the time of the alleged invention that an operating system is a computer program that provides an environment in which other computer programs can run:

Most people are generally familiar with computer operating systems. An operating system is the code or computer program that provides an environment in which other computer programs can run. The operating system manages the system resources, such as processing resources, memory allocation, and I/O devices, to enable other programs to easily take advantage of these resources. Some operating systems are multitasking, meaning they support concurrent operation of more than one task. A multitasking operating system divides the resources among various processes, giving each process access to memory, processing resources, and at least one execution thread.

(EX1015 (Parsons), 1:14–26).

20. Also, it was generally known that certain kinds of operating systems provided an environment in which other *operating systems* could run. One example was a VMM, or Virtual Machine Monitor:

One solution that was the subject of intense research in the late 1960's and 1970's came to be known as the “virtual machine monitor” (VMM). See, for example, R. P. Goldberg, “Survey of virtual machine research,” IEEE Computer, Vol. 7, No. 6, 1974. During that time, moreover, IBM Corp. adopted a virtual machine monitor for use in its VM/370 system.

A virtual machine monitor is a thin piece of software that runs directly on top of the hardware and virtualizes all the resources of the machine. Since the exported interface is the same as the hardware interface of the machine, the operating system cannot determine the presence of the VMM. Consequently, when the hardware interface is compatible with the underlying hardware, the same operating system can run either on top of the virtual machine monitor or on top of the raw hardware.

(EX1017 (Devine), 1:37–52).

21. VMMs and similar “virtualization” OSes provided a means to run

multiple standard OSes concurrently on a single CPU or a multiple-CPU processor:

FIGS. 7 and 8, as well as the description of the invention above, have assumed the presence of a single processor in the hardware 710. This assumption has been made solely to simplify the discussion. The invention may also be used where the hardware 710 includes more than one processor.

(*Id.*, 25:2-6)

B. Configuring and operating systems with multiple processors

22. It was known at the time of the alleged invention to configure and operate a system, such as a microcomputer, with multiple processors to support additional and enhanced computing features. (EX1009 (Floyd), 1:16–27 (excerpted below); EX1011 (George), 1:19–25 (excerpted below); EX1016 (Kleinsorge), 1:34–65). Such a system was often referred to as a multiprocessor system, or just a multiprocessor.

... By contrast, microcomputers designed for more demanding tasks (e.g., network servers) commonly include multiple CPUs. While increasing the number of available processors augments computing power, coordinating the actions of the cooperating CPUs requires a concomitant increase in the complexity of the overall system architecture. ...

(EX1009 (Floyd), 1:20–27 (excerpted with emphasis added))¹.

Over the last few years, there have been many advances in computer system technology. These advances have lead to the development of computer systems having multiple processors to support additional and/or enhanced computing features. FIG. 1 illustrates, for example, one type of computer system that uses a co-processor in conjunction with a host CPU to perform complex mathematical operations.

(EX1011 (George), 1:19–25).

23. Because of the multiple processors in the multiprocessor system, these systems can execute applications more effectively because applications can be distributed among many processors instead of waiting to be executed by a single processor. *See, e.g.*, EX1009, 1:20–33 (excerpted above), *see also* EX1016, 1:41–46. For example:

Traditionally, computing speed has been addressed by using a “shared nothing” computing architecture where data, business logic, and graphic user interfaces are distinct tiers and have specific computing resources dedicated to each tier. Initially, a

¹ Unless otherwise noted, all emphases are added.

single central processing unit was used and the power and speed of such a computing system was increased by increasing the clock rate of the single central processing unit. More recently, computing systems have been developed which use several processors working as a team instead one massive processor working alone. **In this manner, a complex application can be distributed among many processors instead of waiting to be executed by a single processor.**

EX1016, 1:41–46.

C. Multiprocessor system running multiple operating systems concurrently

24. At the time of the alleged invention, it was well known that a multiprocessor system could operate with either a single operating system (OS) or multiple OSs. Running a single OS on a multiprocessor was common:

“There are a number of ways to use a multiprocessor. Today, the most common use is to execute multiple processes in parallel to increase throughput in a multiprogramming environment under the control of a multiprocessor aware operating system. We note that there are a number of commercially available operating systems that have this capability (e.g. Silicon

Graphics IRIX, Sun Solaris, Microsoft Windows NT).”

EX1008, 3.

25. However, it was also well known that a multiprocessor system could run multiple OSes concurrently. *See e.g.*, EX1011, 1:25–28 (discussing a multiprocessor system with a single OS), 2:41–47 (discussing a multiprocessor system with multiple operating systems), 2:55–61; EX1016, 1:46–65 (excerpted below). For example:

Such systems typically consist of several central processing units (CPUs) which are controlled by a single operating system.

In a variant of a multiple processor system called “symmetric multiprocessing” or SMP, the applications are distributed equally across all processors. ... In another variant called “asymmetric multiprocessing” or AMP, one processor acts as a “master” and all of the other processors act as “slaves.”

Therefore, all operations, including the operating system, must pass through the master before being passed onto the slave processors. These multiprocessing architectures have the advantage that performance can be increased by adding additional processors, but suffer from the disadvantage that the software running on such systems must be carefully written to

take advantage of the multiple processors and it is difficult to scale the software as the number of processors increases. Current commercial workloads do not scale well beyond 8-24 CPUs as a single SMP system, the exact number depending upon platform, operating system and application mix.

EX1016 (Kleinsorge), 1:46–65.

26. As an example of an AMP system, George (EX1011) discloses:

In one embodiment, the computer system may include a first processor coupled with a controller through a first bus and a second processor coupled with the controller through a second bus. **The second processor (e.g., a co-processor) operates with an operating system independent from the operating system of the first processor (e.g., a host processor).**

...

FIG. 2 illustrates one embodiment of a computer system having multiple processors running independent operating systems. Computer system 200 includes a processor 210, memory control hub (MCH) 220, an input/output control hub (ICH) 230, platform devices 240, co-processor 250, and buses 261-264. Buses 261-264 carry data and addresses to the various

components in computer system 200.

EX1011, 2:41–47, 2:55–61.

D. Sharing multiprocessor resources

27. In multiprocessor systems that execute multiple OSs, it was further known at the time of the alleged invention that such systems could be configured to share resources, such as one or more of its processors, between the different operating systems running simultaneously in the system. For example:

In accordance with the principles of the present invention, multiple instances of operating systems execute cooperatively in a single multiprocessor computer wherein all processors and resources are electrically connected together. The single physical machine with multiple physical processors and resources is adaptively subdivided by software into multiple partitions, each with the ability to run a distinct copy, or instance, of an operating system. Each of the partitions has access to its own physical resources plus resources designated as shared. In accordance with one embodiment, the partitioning of resources is performed by assigning resources within a configuration.

More particularly, software logically, and adaptively, partitions

CPUs, memory, and I/O ports by assigning them together. An instance of an operating system may then be loaded on a partition. At different times, different operating system instances may be loaded on a given partition. This partitioning, which a system manager directs, is a software function; no hardware boundaries are required. Each individual instance has the resources it needs to execute independently. **Resources, such as CPUs and memory, can be dynamically assigned to different partitions and used by instances of operating systems running within the machine by modifying the configuration.** The partitions themselves can also be changed without rebooting the system by modifying the configuration tree. The resulting adaptively-partitioned, multi-processing (APMP) system exhibits both scalability and high performance.

EX1016, 4:43–5:3.

E. Adoption of single-chip multiprocessors

28. It was known at the time of the alleged invention that multiple processors in a multiprocessor system, previously implemented on separate chips, could be advantageously combined on a single chip with sufficiently advanced chip manufacturing technology:

Due to recent advances in the manufacture of integrated circuits, however, multi-core microprocessor chips may soon be replacing single-core CPUs as the processor of choice for high-performance microcomputers. Multi-core microprocessors feature a single integrated circuit that includes two or more main processing cores, each of which may be utilized as if it was a separate central processing unit. Furthermore, in state-of-the-art multi-core processors, each of the main cores provides computing power that equals or exceeds that of a conventional high-performance single core processor.

EX1009, 1:37–48.

29. Indeed, the concept of a single chip multiprocessor has been around since at least the 1990s:

Advances in IC processing allow for more microprocessor design options. The increasing gate density and cost of wires in advanced integrated circuit technologies require that we look for new ways to use their capabilities effectively. **This paper shows that in advanced technologies it is possible to implement a single-chip multiprocessor in the same area as a wide issue superscalar processor.** We find that for applications with little

parallelism the performance of the two microarchitectures is comparable. **For applications with large amounts of parallelism at both the fine and coarse grained levels, the multiprocessor microarchitecture outperforms the superscalar architecture by a significant margin.** Single-chip multiprocessor architectures have the advantage in that they offer localized implementation of a high-clock rate processor for inherently sequential applications and low latency interprocessor communication for parallel applications.

EX1008, 1.

30. In addition, *Hammond* explains the same conclusion that for a given size of chip, a single-chip multiprocessor will outperform a wide-issue superscalar uniprocessor. EX1013, 1, 4

31. It was also expected at the time of the alleged invention that such single chip multiprocessors (also known as multi-core microprocessor chips) would replace single-processor chips in computer systems. Indeed, *Floyd* discloses this change as I noted above. EX1009, 1:37–48.

32. Barroso also relatedly discloses that:

... The abundance of explicit thread-level parallelism in commercial workloads, along with advances in semiconductor

integration density, **identify chip multiprocessing (CMP) as potentially the most promising approach for designing processors targeted at commercial servers.**

EX1010, 1.

The use of chip multiprocessing is inevitable in future microprocessor designs. Advances in semiconductor technology are enabling designs with several hundred million transistors in the near future. ... At the same time, **using the extra transistors to integrate multiple processors onto the same chip is quite promising, especially given the abundance of explicit thread-level parallelism in important commercial workloads.** At least a couple of next-generation processor designs subscribe to this philosophy by integrating two superscalar cores on a single die. **The key questions for designers of future processors will not be whether to use chip multiprocessing, but the appropriate trade-off between the number of cores and the power of each core, and how to best partition the memory hierarchy among the multiple cores.**

Id., 11.

33. Indeed, it is was well known that in the year 2000, Sun Microsystems and IBM had already incorporated chip multiprocessors into its servers:

Early next year, top-echelon power users will find out if two heads are better than one. That's when Sun Microsystems Inc., IBM, Compaq Computer Corp., Hewlett-Packard Co. and others will **start to roll out high-end servers that take advantage of chip multiprocessing (CMP)**, a step forward from current systems that load up boxes with multiple discrete chip modules.

EX1007, 1.

The age for commercially viable CMP systems has arrived, thanks to continuing refinements in chip manufacturing techniques that let engineers pack circuits more densely. The extra die space opens up room for multiple chips - two in the initial systems but perhaps up to eight in later generations of CMP modules.

Id., 2.

34. It was also well known at the time of the alleged invention of the '014 patent that there were many recognized benefits in using a single chip multiprocessor over a single processor chip. For example, each processor on a single

chip multiprocessor could be used as if it was a separate processor. EX1009, 1:37–48. Furthermore, each of the processors in the single chip multiprocessor provides a computing power that equals or exceeds that of a conventional high-performance single core processor. *Id.*; *see also* EX1010, 1, 11. Also due to the number of processors on the single chip multiprocessor, a system can provide better execution of workloads. *See, e.g.*, EX1008, Abstract, 1; EX1010, 1, 11; EX1007, 1.

35. Likewise, it was well known that by putting multiple processors on a single chip, engineers could take advantage of shorter distances and faster bus speeds when shuttling data between the processors. EX1007, 1; *see also* EX1009, 1:48–60 (“Among the benefits provided by multi-core processors is an increase in physical processing core density within multiprocessor data processing systems, in that more processing cores may operate within a given amount of space, which leads to improvements in overall multiprocessor system performance”).

36. For example:

Our detailed simulation results show that while each Piranha processor core is substantially slower than an aggressive next-generation processor, **the integration of eight cores onto a single chip allows Piranha to outperform next-generation processors by up to 2.9 times (on a per chip basis) on important workloads such as OLTP** [on-line transaction

processing]. This performance advantage can approach a factor of five by using full-custom instead of ASIC logic.

EX1010, 1.

37. In another example:


Early tests are showing that two processors in a single module outperform multiple discrete processors by 50% or more. By putting two CPUs on a single piece of silicon, engineers can take advantage of shorter distances and faster bus speeds when shuttling data between the two CPU cores. The net performance result for IBM's version - called the Power4 processor - is the ability to process 100GB of data per second, or the equivalent of 20 full-length DVDs, says Joel Tendler, director of technology strategy at IBM's server group in Austin, Texas.

EX1007, 1.

VI. THE '014 PATENT

38. The '014 patent relates to “a multiprocessing chip utilizing multiple operating systems.” EX1001, 1:6-10. The '014 patent admits that multiprocessing systems with a multiprocessing chip were well-known, and it depicts a known prior art system in Figure 2 below. EX1001, 1:34-59; 1:41-42 (explaining that it was well-known in the art to use “multiple processors on a single die” to conserve space in a

system).



EX1001, FIG. 2 (annotated).

39. The prior art multiprocessing system (200) includes a chip multiprocessor 295 that “includes multiple processors 210-240 on a single die 290” (highlighted in orange). EX1001, 1:34-37. The processors 210-240 use a single operating system 250 (highlighted in yellow) stored in memory 260 (highlighted in blue). EX1001, 1:37-40. The processors “may communicate with the memory 260 via a bus 270.” EX1001, 1:40-41. The ’014 patent states that “schemes that have placed multiple processors on a single chip typically utilize a single operating system for tying all the processors together.” EX1001, 1:44-47; 1:60-63 (the ’014 patent acknowledges that it was known in the art that multiprocessor systems could “support[] multiple processors executing multiple operating systems”).

40. Unlike the prior art system with multiple processors and a single

operating system, the '014 patent discloses a “multiprocessing system” where “multiple processors are connected to a memory storing *multiple* operating systems,” as illustrated in Figure 3 below. EX1001, 2:12-16 (emphasis added)².




FIG. 3

EX1001, FIG. 3 (annotated), 2:54-55.

41. Like the prior-art system in FIG. 2, the computer system 300 in FIG. 3 includes a chip multiprocessor 350 having multiple processors 305-320 mounted on a single die (highlighted orange). EX1001, 3:4-7. The purported advance of the multiprocessing system is that it further includes multiple operating systems 325-340 (highlighted yellow) stored in a memory 365 (highlighted blue), rather than a single operating system 250 as shown in Fig. 2 above. EX1001, 3:10-11. The '014 patent explains that “[d]uring operation, each processor 305-320 may access a respective operating system 325-340 by communicating with the memory 365, for

² Unless otherwise noted, all emphases are added.

example, via a bus 370.” EX1001, 3:11-14. The ’014 patent also explains that “[e]ach processor may be ... capable of simultaneously executing multiple operating systems, for example, by context switching, which may include rapidly switching between multiple operating systems.” EX1001, 3:47-51

42. Claims 1, 3, 5, 7, 11–13, and 15 of the ’014 patent recite limitations relating to the features discussed above. However, as I discuss below in Sections VII and IX, it is my opinion that all of these features were well known in the art at the relevant time.

VII. OVERVIEW OF THE PRIOR ART³

A. Overview of Gulick (EX1005)

43. Gulick discloses a multi-processor computer system that includes “one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system and allows the different partitions to communicate with one another through shared memory.” EX1005, 8:47-53. Gulick discloses a “multi-processor system that includes processor modules 110,

³ These summaries of the prior art references in Section VII are applicable to my analysis below in Section IX. Accordingly, I incorporate these summaries, which reflect how a person of ordinary skill in the art would have understood the references, into my analysis below where appropriate (*see, e.g.*, Section IX below).

112, and 114” that “share access to main (or global) memory 160 . . . through a system interconnection mechanism, such as system interconnection 130.” EX1005, 10:53-67. Such arrangement allows processor modules 110-114 to “communicate with each other through main memory 160.” EX1005, 10:64-67.

44. Gulick discloses that the processor modules 110-114 “may be configured as a separate partition within the computer system, such that multiple partitions may exist within the computer system.” EX1005, 11:1-3. Each partition operates “under the control of a separate operating system.” EX1005, 11:4-5. For example, as shown below and disclosed in Gulick, “each processor module 110, 112, and 114 . . . can be defined as a separate partition controlled via a separate operating system 170, 172, and 174.” EX1005, 11:5-10, FIG. 1.




Figure 1

EX1005, FIG. 1.

45. Gulick disclose that “[e]ach operating system 170, 172 and 174 views main memory separately as though each is the only entity accessing main memory 160.” EX1005, 11:8-10

46. FIG. 2 illustrates components of an embodiment of the computing system shown in FIG. 1 in more detail. For example, Gulick’s computer system 200 includes main memory 160 and a plurality of processing modules 240 connected to main memory via third level cache modules 230 and crossbar interconnects 290. EX1005, 11:24-26.





Figure 2

EX1005, FIG. 2.

47. As shown, computer system 200 includes pairs of processors – processor 240A and processor 240B. Processor 240A and processor 240B “share a common bus (e.g., 280A) with a single [third level cache application specific integrated circuits] TCT (e.g., 270A) within a given [third level cache] TLC (e.g., 230A).” EX1005, 11:45-54. Gulick discloses that “each TLC 230 is connected to four processors” and that “[e]ach TLC 230 and its respective four processors define a sub-Pod.” EX1005, 11:57-60. Two sub-Pods “are connected via crossbar interconnect (e.g., crossbar interconnect 290A or 290B) to form a Pod.” EX1005, 11:60-62.

48. Computer system 200 includes “multiple partitions . . . each of which may comprise one or more Pods or sub-pods, each operates under the control of a separate operating system.” EX1005, 13:21-24. As shown in the memory configuration illustrated in FIG. 5 (reproduced below), the memory configuration of computer system 200 includes “three operating systems” each of which “has its own address space 502.” EX1005, 13:30-36



EX1005, FIG. 5 (annotated).

49. As shown, Gulick discloses three memory windows (annotated red, green, blue, respectively) that “are defined within the address space 504 of main memory 160” and a shared memory (annotated yellow), “which is accessible by all

three operating systems.” EX1005, 13:37-43; FIG. 5. Gulick discloses that “the present invention contemplates an environment wherein at least two of the operating systems are different and one or more operating systems does not control or manage the second operating system.” EX1005, 13:26-29. In this manner, Gulick discloses a multi-processor computer system and method that allows for “multiple operating systems, including different operating systems, to operate in different partitions on the computer system.” EX1005, 2:37-56.

B. Overview of AAPA (EX1001)

50. The '014 patent admits that multiprocessing systems with multiple processors on a single die were well-known, and it depicts a known prior art system in Figure 2 below. EX1001, 1:34-59; 1:41-42 (explaining that it was well-known in the art to use “multiple processors on a single die” to conserve space in a system).




FIG. 2
Prior Art

EX1001, FIG. 2 (annotated).

51. The prior art multiprocessing system (200) includes a chip multiprocessor 295 that “includes multiple processors 210-240 on a single die 290” (highlighted orange). EX1001, 1:34-37. The processors 210-240 use a single operating system 250 (highlighted yellow) stored in memory 260 (highlighted blue). EX1001, 1:37-40. The processors “may communicate with the memory 260 via a bus 270.” EX1001, 1:40-41. The ’014 patent states that “schemes that have placed multiple processors on a single chip typically utilize a single operating system for tying all the processors together.” EX1001, 1:44-47. The ’014 patent also admits that it was known in the art that multiprocessor systems could “support[] multiple processors executing multiple operating systems”. EX1001, 1:60-63.

C. Motivation to Modify Gulick Based on AAPA

52. As previously noted, Gulick discloses a multi-processor computer system that includes “one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system and allows the different partitions to communicate with one another through shared memory.” EX1005, 8:47-53. Gulick thus teaches a computing system that includes multiple processors, where the multiple processors are able to concurrently execute a plurality of operating systems. Thus, it is my opinion that a person of ordinary skill in the art seeking to implement a multi-processor system pursuant to the

teachings of Gulick would consider references disclosing the arrangement or structure of such multi-processor systems, which includes AAPA.

53. It is my opinion that the AAPA discloses precisely what is suggested by Gulick, multiple processors mounted on a single die. AAPA explains that this arrangement is beneficial because providing multiple processors on a single die conserves space compared to multiprocessing schemes that utilizes multiple processors, with each processor on a separate die. EX1001, 1:28-32 (“One conventional multiprocessing scheme . . . includes a computer system 100 having multiple processors 10-40, each on a separate die (i.e., chip) 50-80 and connected to a single operating system 90 stored in a memory 95.”); 1:36-42 (“The chip multiprocessor includes multiple processors 210-240 on a single die . . . System 200 conserves space by providing multiple processors on a single die.”). It is my opinion that a person of ordinary skill in the art would immediately understand that the teachings of the AAPA regarding the fabrication process of multiple processors on a single die could be used in combination with the teachings of Gulick regarding a multiprocessor system that simultaneously executes multiple operating systems.

54. As further discussed below, the advantages of placing multiple processors on the same die were well known in the art as of the filing date of the '014 Patent. *See infra* Section IX.A.1.b. Thus, I believe that these known advantages would further motivate a POSITA to implement Gulick's processors on

a single die.

D. Motivation to Modify AAPA Based on Gulick

55. As previously noted, the AAPA discloses that multiprocessing systems with multiple processors on a single die were well-known. EX1001, 1:34-59. The AAPA discloses a multiprocessing system 200 that includes a chip multiprocessor 295 that “includes multiple processors 210-240 on a single die 290.” EX1001, 1:34-37. The processors 210-240 use a single operating system 250 stored in memory 260. EX1001, 1:37-40. The AAPA thus discloses a computing system that includes multiple processors, where the multiple processors are able to execute an operating system. The AAPA acknowledges, however, that it was well known in the art that multiprocessor system, such as computer system 200, suffer from well known scalability problems, despite its multiple processors 210-240 being placed on a single die. EX1001, 1:43-44. The AAPA provides that it was well known in the art that the limitation of these multiprocessing schemes is due to the scalability issues related to operating systems executed by a large number of processors. EX1001, 1:47-50 (“A well known limitation of this scheme and other multiprocessing schemes utilizing a single operating system is that an operating system does not scale well to large number of processors.”). Thus, it is my opinion that a person of ordinary skill in the art seeking to implement a multi-processor system pursuant to the teachings of the AAPA would consider references disclosing

scalability improvements to a multi-processor system executing an operating system, which includes Gulick.

56. It is my opinion that Gulick discloses precisely what is suggested by the AAPA, a multiprocessor system that addresses the scalability problems that were well known in the art by providing a multiprocessing system that executes multiple operating systems. Gulick explains that “[a]s modern systems grow, and the demands of the user increases, the necessity of employing a plurality of operating system increases.” EX1005, 2:32-34. Gulick then provides a solution to this limitation by providing a multiprocessing system that allows multiple operating systems to operating in different partitions on the computer system. EX1005, 2:37-43. It is my opinion that a person of ordinary skill in the art would immediately understand that the teachings of Gulick regarding use of a plurality of operating systems in a multiprocessor system to improve the scalability issues surrounding multiprocessor systems could be used in combination with the teachings of the AAPA regarding a multiprocessing system.

VIII. CLAIM CONSTRUCTION

57. I have been asked to give the terms in the challenged claims their plain and ordinary meanings, as would have been understood by a person of ordinary skill in the art, at the time of the alleged invention, taking into consideration the language of the claims, the specification, and the prosecution history (EX1004) of the '014

patent. I have applied these understandings of the claim terms in my analysis and in forming my opinions where appropriate in this Declaration.

A. “simultaneously executing two or more operating systems”

58. As I understand, in the currently pending litigation, Petitioner and PO have agreed that “simultaneously executing two or more operating systems” should be construed as “simultaneously executing two or more *independent* operating systems.”

59. Accordingly, I have been asked to analyze the challenged claims under both a plain and ordinary meaning construction and the agreed upon construction.

B. “multiple operating systems residing in a memory”

60. As I understand, in the currently pending litigation, Petitioner and PO have agreed that “multiple operations systems residing in a memory” should be construed as “multiple operating system residing in the same memory.”

61. Accordingly, I have been asked to analyze the challenged claims under both a plain and ordinary meaning construction and the agreed upon construction.

C. “multiple processors are connected to said memory via a bus”

62. As I understand, in the currently pending litigation, Petitioner and PO have agreed that “multiple processors are connected to said memory via a bus” in

dependent claim 3 should be construed as “multiple processors are connected to said memory via the same bus.”

63. Accordingly, I have been asked to analyze the challenged claims under both a plain and ordinary meaning construction and the agreed upon construction.

D. “processor means for executing a plurality of operating system means” (claim 12)

64. As I understand, in the currently pending litigation, Petitioner and PO have an agreed upon 35 U.S.C. § 112 ¶6 construction for “processor means.” The agreed-upon function is “executing a plurality of operating system means.” The agreed-upon structure is “chip multiprocessor 350 having multiple processors 305-320 mounted on a single die 360 as shown in Figure 3 and associated descriptions in the specification; or chip multiprocessor 450 having multiple processors 405-420 mounted on a single die 460 shown in Figure 4 and associated descriptions in the specification; and equivalents thereof.”

65. Further, to the extent the Board fails to adopt the agreed upon construction, I have also been asked to analyze the challenged claims by construing the function of “processor means” as “simultaneously execute two or more operating system means of said plurality of operating system means.”

66. Additionally, to the extent that “processor means” is not a means plus function limitation, I have also been asked to analyze “processor means” in accordance to its plain meaning.

67. Accordingly, for the purposes of my analysis, I have been asked to apply the agreed upon construction for “processor means.”

E. “operating system means” (claim 12)

68. For purposes of my analysis, I have been asked to apply plain and ordinary meaning to this limitation because it does not invoke means plus function.

F. “memory means for storing said plurality of operating system means” in claim 12

69. For purposes of my analysis, I have been asked to assume that the identified function of the “memory means” is “storing said plurality of operating system means,” where the corresponding structure is “SRAM and/or DRAM on the same chip as one or more processors; SRAM and/or DRAM on separate chips connected to one or more processors; magnetic media, such as tape or disk; optical media, such as CD-ROM; and the like.” EX1001, 3:15-20.

70. I have applied this understanding in my analysis

IX. THE PRIOR ART DISCLOSES ALL OF THE LIMITATIONS OF CLAIMS 1, 3, 5, 9, 11–13, AND 15 OF THE '014 PATENT

A. Ground 1: Claims 1, 3, 5, 7, 11-13, 15 are obvious over the combination of Gulick and AAPA

71. In my opinion, the combination of Gulick and AAPA renders claims 1, 3, 5, 7, 11-13, and 15 of the '014 patent obvious.

1. Claim 1

a. A multiprocessing system comprising

72. To the extent the preambles are limiting, it is my opinion that the Gulick-AAPA combination (hereinafter “Gulick-AAPA”) discloses the limitations therein.

73. Gulick discloses “a multi-processor system that includes processor modules 110, 112, and 114.” EX1005, 10:53-54.





Figure 1

EX1005, FIG. 1.

74. Gulick discloses a more detailed description of an embodiment of the multi-processor system of FIG. 1. For example, as shown below, Gulick discloses

“a computing system 200 . . . [that] includes a main memory . . . and a plurality of processing modules 240 connected to the main memory.” EX1005, 11:22-26, FIG.

2.



EX1005, FIG. 2.

75. Gulick states that “the processing modules and the main memory [as shown in FIG. 2] are arranged in a symmetrical *multiprocessing architecture*. EX1005, 11:28-32.

76. AAPA Fig. 2 likewise shows a conventional chip multiprocessor 295 in a computer system 200. EX1001, 1:34-59; 1:41-42.

77. Therefore, I submit that Gulick-AAPA discloses and/or suggests this limitation.

b. multiple processors mounted on a single die

78. It is my opinion that Gulick-AAPA discloses this limitation.

79. Gulick discloses a computer system that includes processor modules 110-114 (highlighted green) labeled as “processor(s)” in Fig. 1. EX1005, 10:53-54. Each processor module 110-114 “can . . . include a plurality of processors.” EX1005, 10:59-60.




Figure 1

EX1005, FIG. 1 (annotated).

80. Gulick discloses computer system 200 that includes a plurality of processing modules 240 comprising processors 240A/240B-240R/240S (highlighted green). EX1005, 11:24-26.




Figure 2

EX1005, FIG. 2 (annotated).


81. Computer system 200 includes “multiple partitions . . . each of which may comprise one or more Pods or sub-pods, each operates under the control of a

separate operating system.” EX1005, 13:21-24. Gulick discloses that a pod is formed from multiple sub-pods. EX1005, 11:60-62 (Two sub-Pods “are connected via crossbar interconnect (e.g., crossbar interconnect 290A or 290B) to form a Pod”). Gulick’s pods and sub-pods are “both examples of processing modules.” EX1005, 10:35-42. One or more processing modules form a partition operating under the control of its own operating system. EX1005, 9:41-46, 10:27-34.

82. Gulick also teaches that “the present invention contemplates other multiprocessing environments and configurations.” EX1005, 13:12-14. This further demonstrates to a POSITA that Gulick contemplates implementing a plurality of operating systems on a variety of conventional multiprocessors, including those disclosed in the AAPA.

83. To the extent that Gulick does not explicitly provide that the multiple processors are mounted on a single die, it is my opinion that AAPA teaches doing so.

84. The ’014 patent discloses that it is well known for a computer system to have multiple processors mounted on a single die. For example, the ’014 patent discloses a “conventional multiprocessing scheme . . . [that] includes a computer system having a chip multiprocessor 295.” EX1001, 1:34-36.



EX1001, FIG. 2 (annotated).

85. The '014 patent provides that conventional multiprocessing schemes include a chip multiprocessor 295 that “includes multiple processors 210-240 on a single die 290” (highlighted orange). EX1001, 1:34-37.

86. It is my opinion that a POSITA would have found it obvious in light of AAPA to configure the system in Gulick so that its multiple processors are “mounted on a single die,” as claimed, and would have had a reasonable expectation of success in so doing. Such an implementation would have been a straightforward combination of well-known technologies using known methods and would have had predictable results.

87. I believe that a POSITA would have recognized that Gulick and AAPA disclose features in a similar technological field. For example, as discussed above, both Gulick and AAPA relate to multi-processor computing systems. EX1005,

10:53-54 (“FIG. 1 illustrates a multi-processor system that includes processor modules 110, 112, and 114”); EX1005, 11:28-32 (“the processing modules and the main memory [as shown in FIG. 2] are arranged in a symmetrical *multiprocessing architecture*”); EX1001, 1:34-36 (“A second conventional multiprocessing scheme (shown in FIG. 2) includes a computer system 200 having a chip multiprocessor 295.”).

88. Further, I believe that a POSITA would have been motivated to configure the system in Gulick so that its multiple processors are “mounted on a single die” in order to obtain benefits associated with using multiple processors on a single integrated circuit die that are applicable to any multiprocessing architecture, including Gulick’s, that the AAPA admits were known to a POSITA. For example, utilizing a single die is beneficial because it allows system 200 to conserve space by providing multiple processors on a single die. EX1001, 1:41-42. Additionally, in the AAPA “mounting multiple processors on a single die reduces the cabling problem inherent in connecting multiple processors on separate dies in separate houses,” “reduces the latency for communication among the processors”, “improves the efficiency of message passing”, “reduces chip-to-chip communication costs”, “leads to further power efficiency”, and increases “scalability for multiprocessing.” EX1001, 2:26-38.

89. Therefore, it is my opinion that the patentee's admissions in the AAPA both support a motivation to combine particular disclosures and demonstrates the knowledge of the ordinarily-skilled artisan at the time of the invention.

90. Moreover, it is my opinion that, consistent with the AAPA, the benefits of having multiple processors on a single die in a multi-processor system were well known in the art before the '014 patent thereby motivating a POSITA to apply the AAPA's teachings regarding a single die to Gulick. For example:

- combining processors on a single piece of silicon “outperforms multiple discrete processors by 50% or more” (EX1007, 1) and allows engineers to “take advantage of shorter distances and faster bus speeds when shuttling data between” the processors (*Id.*);
- “[s]ingle-chip multiprocessor architectures have the advantage in that they offer localized implementation of a high-clock rate processor for inherently sequential applications and low latency interprocessor communication for parallel applications.” EX1008, Abstract;
- “[m]ulti-core microprocessors feature a single integrated circuit that includes two or more main processing cores, each of which may be utilized as if it was a separate central processing unit” (EX1009, 1:37-47) and “each of the main cores provides computing power that equals or exceeds that of a conventional high-performance single core processor” (*Id.*);

- “[a]mong the benefits provided by multi-core processors is an increase in physical processing core density within multiprocessor data processing systems, in that more processing cores may operate within a given amount of space, which leads to improvements in overall multiprocessor system performance” (EX1009, 1:48-60); and
- “[t]he key questions for designers of future processors will not be whether to use chip multiprocessing, but the appropriate trade-off between the number of cores and the power of each core” (EX1010, 292).⁴

91. Indeed, the '014 patent acknowledges using multiple processors mounted on a single die in a computer system as conventional and beneficial. For instance, the '014 patent describes a “conventional multiprocessing scheme (shown in FIG. 2) includes a computer system 200 having a chip multiprocessor 295” where “[t]he chip multiprocessor includes multiple processors 210-240 on a single die 290,” which saves space in the computer system. EX1001, 1:34-44, FIG. 2 (illustrating this conventional scheme).

⁴ To show that claim 1 is obvious, I rely here on the teachings of Gulick and AAPA only. Exhibits 1007-1010 demonstrate what I believe a POSITA would have known at the time of the alleged invention of the '014 patent regarding the benefits of having multiple processors on a single die. *Int'l Bus. Machines Corp. v. Intellectual Ventures II, LLC*, IPR2015-00089, Paper No. 44 at 15 (Apr. 25, 2016).

92. Thus, it is my opinion that placing Gulick's multiple processors (e.g., one or more pods) on a single die would have been a common-sense approach. For example, it was obvious to implement one or both pods comprising processors 240 in Fig 2 on a single die. EX1005, 11:33-54. As discussed in more detail hereinbelow, Gulick discloses multiple operating systems operating on the processors of exemplary system of Fig. 2, with each operating system in a partition of one or more sub-pods. EX1005, 13:21-43. I believe that a POSITA was motivated to do so by the many known advantages of single-die multiprocessors, including those expressly taught by the AAPA: "mounting multiple processors on a single die reduces the cabling problem inherent in connecting multiple processors on separate dies in separate houses," "reduces the latency for communication among the processors", "improves the efficiency of message passing", "reduces chip-to-chip communication costs", "leads to further power efficiency", and increases "scalability for multiprocessing." EX1001, 2:26-38.

93. It is my opinion that a POSITA would have had a reasonable expectation of success in implementing Gulick's multiprocessing system on a single die as taught by AAPA because doing so was within the ordinary skill of a POSITA and did not require undue experimentation. Notably, the '014 patent itself does not provide detail of how an AAPA single-die multiprocessor in FIG. 2 is configured to comprise multiple operating systems disclosed in embodiments of FIGs. 3 and 4.

This is indicative of the fact that a POSITA understood how to implement a plurality of operating systems on a multiprocessor, including a single-die multiprocessor.

94. Moreover, Gulick provides details for implementing a plurality of operating systems that is applicable to any multiprocessor, including a modification of Gulick in view of AAPA to arrange processors 240 on a single die. For example, Gulick's teaching of using dedicated memory windows for each operating system partition and a shared memory window for inter-OS communication functions (as illustrated in Gulick Fig. 5) enables implementation of multiple operating systems on a multiprocessor, including a single-die multiprocessor.

95. Therefore, based on AAPA and the knowledge in the art, I believe that it would have been obvious for the multiple processors in Gulick-AAPA to be mounted on a single die.

96. Thus, it is my opinion that Gulick-AAPA discloses and/or suggests this limitation.


c. multiple operating systems residing in a memory connected to said multiple processors

97. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

98. Gulick discloses that its multi-processor system includes main memory 160. EX1005, 10:53-67. Main memory 160 "is a directory-based memory system

and is capable of supporting various memory consistency models.” EX1005, 11:33-37. Main memory 160 “includes a plurality of memory storage units (MSUs) 220, such as memory storage units 220A, 220B, 220C, and 220D.” EX1005, 11:37-39.

99. FIG. 5 illustrates “an exemplary memory configuration that can be generated on the computer system of FIG. 2.” EX1005, 13:30-31. For example, as shown below, the memory configuration of computer system 200 includes “three operating systems” each of which “has its own address space 502.” EX1005, 13:30-36.



EX1005, FIG. 5 (annotated).

100. FIG. 5 illustrates three exclusive memory windows (annotated red, green, blue, respectively) that “are defined within the address space 504 of main memory 160” and a shared memory (annotated yellow), “which is accessible by all

three operating systems.” EX1005, 13:37-43; FIG. 5. Gulick discloses that a memory window “is defined within the main memory to which multiple partitions have shared access.” EX1005, abstract. Thus, it is my opinion that Gulick discloses that a memory window is a portion of the main memory.

101. Gulick discloses that each operating system resides in a respective memory window in main memory. Gulick discloses that “[f]or each different partition configured within the computer system, the main memory may further have defined therein an exclusive memory window to which only that partition has access and *in which* the operating system of that partition executes.” EX1005, 2:66-3:5.

102. It is my opinion that a POSITA understands that if an operating system executes within a memory window assigned to the partition, then that operating system is stored in the memory window because the processor executes the operating system from the memory address where it is stored. In this regard, Gulick also discloses that the “operating system is the first program code loaded into the main memory of a computer system after the computer system is turned on” and that the “central core of the operating system resides in the memory space *at all times*.” EX1005, 10:7-11. It is my opinion that a POSITA would understand this teaching to refer to each operating system residing in its memory window during

operation. Thus, I believe that Gulick discloses “multiple operating systems residing in a memory.”

103. Further, I believe that Gulick discloses that the memory is “connected to said multiple processors.”

104. Gulick discloses that the processor modules 110-114 “may be configured as a separate partition within the computer system, such that multiple partitions may exist within the computer system.” EX1005, 11:1-3. Each partition operates “under the control of a separate operating system.” EX1005, 11:4-5. For example, as shown below, “each processor module 110, 112, and 114 . . . can be defined as a separate partition controlled via a separate operating system 170, 172, and 174.” EX1005, 11:5-10, FIG. 1.




Figure 1

EX1005, FIG. 1.

105. Further, Gulick's computer system 200 includes main memory 160 (annotated red) and a plurality of processing modules 240 (annotated green) connected to main memory. EX1005, 11:24-28 ("Computer system 200 includes a main memory, illustrated here as main memory 160, and a plurality of processing modules 240 connected to the main memory via respective third level cache modules 230 and crossbar interconnects 290.").




Figure 2

EX1005, FIG. 2 (annotated).

106. Thus, I believe that Gulick teaches that the memory is “connected to said multiple processors.”


107. Furthermore, it is my opinion that a POSITA would have understood that placing the multiple processors on a single chip (*see supra* Section IX.A.1.b) would not have detracted from or changed the connection between the processors and memory in Gulick. Indeed, a POSITA would have recognized that in Gulick-AAPA, the processing modules 240, which comprise multiple processors mounted on a single die, would be connected to the memory (such as the main memory 160) where the multiple operating systems reside.

108. Thus, I believe that Gulick-AAPA discloses and/or suggests this limitation.

109. Further, to the extent that the Board adopts the agreed-upon District Court construction (“multiple operating systems residing in the same memory”) rather than plain meaning, it is my opinion that Gulick-AAPA still discloses and/or suggests this limitation.

110. As stated, and shown in FIG. 5 above, “three exclusive memory windows 540A, 540B and 540C, one for each operating system (i.e., partition), and one shared memory window 537, which is accessible by all three operating systems 540A, 540B and 540C (i.e., partitions), are defined within the address space 504 of the main memory 160.” EX1005, 13:37-43; FIG. 5. As explained above, it is my

opinion that a POSITA understands this disclosure to teach that each operating system resides in the same main memory 160 as illustrated in the annotated FIG. 5 below.



EX1005, FIG. 5 (annotated).

111. To emphasize this further, Gulick also discloses that the “central core of the operating system resides in the memory space at all times.” EX1005, 10:7-11. Accordingly, I believe that Gulick teaches that the multiple operating systems reside in the “same memory.”

112. Furthermore, there is no teaching in Gulick that would preclude the multiple operating systems from residing in the same memory.

113. Therefore, it is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

- d. **wherein each of said multiple processors executes an operating system of said multiple operating systems, and two or more of said multiple processors are capable of simultaneously executing two or more operating systems of said multiple operating systems**

114. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

115. Gulick-AAPA includes multiple processors mounted on a single die and multiple operating systems connected to these multiple processors. *See supra* Sections IX.A.1.b-c. Each of these processors in Gulick-AAPA executes a respective operating system from the multiple operating systems, whereby two or more of the processors are capable of simultaneously executing two or more of the operating systems.

116. Gulick discloses that one or more processors is allocated to execute each operating system. For example, Gulick discloses “one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system.” EX1005, 8:47-53. Processor modules 110-114 “may be configured as a separate partition within the computer system, such that multiple partitions may exist within the computer system.” EX1005, 11:1-3. Each partition operates “under the control of a separate operating system.” EX1005, 11:4-5. For example, Gulick disclose that “each processor module 110, 112, and 114 . . . can be defined as a separate partition controlled via a separate operating system 170, 172, and 174.” EX1005, 11:5-10, FIG. 1.

117. Gulick further discloses that the “two or more of said multiple processors are capable of simultaneously executing two or more operating systems of said multiple operating systems,” as claimed. For example, Gulick discloses that the foregoing arrangement of components “allows a plurality of operating system to *concurrently execute* in different partitions within the computer system.” EX1005, 8:47-53.

118. Further, to the extent that the Board adopts the agreed-upon District Court construction of “simultaneously executing two or more operating systems” as

“simultaneously executing two or more independent operating systems,” I believe that Gulick-AAPA still teaches this limitation.

119. Gulick clearly states that “each partition operates under the control of its own operating system” and that “[t]he operating system executing on different ones of the partitions may be different operating systems, or different instances of the same operating system.” EX1005, 9:42-46; *see also* 13:27-29 (“the present invention contemplates an environment wherein at least two of the operating systems are different and one operating system does not control or manage the second operating system.”).

120. Thus, I believe that Gulick teaches this limitation under the agreed-upon construction.

121. Therefore, it is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

2. Claim 3

a. The system of claim 1, wherein said multiple processors are connected to said memory via a bus to execute said multiple operating systems

122. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

123. As discussed for claim 1, Gulick discloses that the multiple processors are connected to the memory to execute the operating systems. *See supra* Sections

IX.A.1.b-d. Gulick teaches that “[m]ain memory 160 includes a plurality of memory storage units (MSUs) 220, such as memory storage units 220A, 220B, 220C, and 220D.” EX1005, 11:37-39. Each MSU 220A-22D “includes sixteen semi-independent banks that share four double-wide data busses and eight unidirectional address busses.” EX1005, 11:41-44. Each pair of processors (e.g., 240A and 240B) (annotated green) “share a common bus (e.g., 280A [annotated blue]) with a single TCT (e.g., 270A) within a given TLC (e.g., 230A).” EX1005, 11:48-51.

124. As shown in FIG. 2 below, such common bus allows processors 240A and 240B, for example, to be connected to memory.




Figure 2

EX1005, FIG. 2 (annotated).

125. Bus 280A connects processors 240A, 240B to third level cache module 230A. EX1005, 11:55-57 (“Each third level cache module 230A through 230D is connected to a respective plurality of processors (MPs) 240A through 240S.”). Each TLC 230A, 230B, 230C, 230D is connected memory 160 via crossbar interconnects 290A, 290B. EX1005, 11:24-28 (“Computer system 200 includes a main memory, illustrated here as main memory 160, and a plurality of processing modules 240 connected to the main memory via respective third level cache modules 230 and crossbar interconnects 290.”); *see also* 11:66-12:1 (“Cross bar

interconnects 290 interface processors 240, through third level caches 230, with memory storage units 220.”).

126. Moreover, the main memory 160 comprises MSU units 220A-D where each MSU “includes sixteen semi-independent banks that share four double-wide data busses and eight unidirectional address busses.” EX1005, 11:41-44. A POSITA would understand this disclosure to teach that the MSU banks are accessed by processors 240 via a memory bus comprising one or more data and address busses, as is conventional in the art. Accordingly, I believe that a POSITA would understand that the full main memory 160 is accessed by the processors through a memory bus. This arrangement in Gulick is consistent with the AAPA memory bus 270, which a POSITA would understand to comprise address and data busses.

127. Further, to the extent that the Board adopts the agreed-upon District Court construction (“multiple processors are connected to said memory via the same bus”), I believe that Gulick-AAPA likewise discloses and/or suggests this limitation.

128. Gulick discloses that processing modules 240 are “connected to the main memory via . . . crossbar interconnects 290.” EX1005, 11:24-28. For example, as shown in FIG. 2 below, multiple sub-pods (e.g., MP 240A-D = first sub-pod; MP 240E-H = second sub-pod) are connected to memory 160 via crossbar interconnect 290A.




Figure 2

EX1005, FIG. 2 (annotated).

129. FIG. 2 of Gulick includes first communication lines (annotated purple). First communication lines are bi-directional communications lines that relay address and data communications in two directions between the sub-pods and main memory 160. Although FIG. 2 of Gulick illustrates two lines (one for each MSU 220A-B), I believe that a POSITA would find it obvious to use the same bus to access shared memory. More specifically, it is my opinion that a POSITA would understand that multiple sub-pods associated with crossbar interconnect 290A would use the same bus to access MSU pair 222, with the bus connecting processors 240 to main memory 160 via crossbar interconnect 290A.

130. Gulick also discloses several other communication lines originating from crossbar interconnect 290A. For example, a second communication line (annotated green) connect crossbar interconnect 290A to MSU pair 224. Similar to the first communication lines discussed above, second communication line is bidirectional. I believe a POSITA would understand that the first communication lines and the second communication line represent shared memory access that allows processors and sub-pods associated with crossbar interconnect 290 to access MSU pair 224 in addition to MSU pair 222. Such functionality allows for interleaved memory access, where each processor 240 distributes memory requests “over the global address space of main memory 160” in order to minimize contention. *See* EX1005, 22:24-30; *see also* 22:20-24; 23:44-24:56.

131. It is my opinion that a POSITA would find it obvious to implement processor access to memory such as MSU pairs 222 and 224 with the same memory bus. It is my opinion that this conventional implementation was well-known and is illustrated in the AAPA as bus 270. A shared bus, such as AAPA bus 270, was obvious in Gulick because it is an efficient manner of connecting multiple processors to a shared memory, such as Gulick’s main memory 160, that avoids the duplication of resources if separate busses are used to access memory. For example, a POSITA would understand that the crossbar interconnects 290 in Gulick were efficient to implement a shared memory access bus to provide access to the entire

memory space of main memory 160. On the other hand, an implementation with two busses would unnecessarily duplicate address, data and control line resources. I believe a POSITA would therefore be motivated to implement the same memory access bus in Gulick between processors 240 and main memory 160 – a conventional solution illustrated by AAPA bus 270.

132. Accordingly, it is my opinion that Gulick-AAPA discloses and/or suggests this limitation under the agreed-upon construction.

3. Claim 5

- a. The system of claim 1, wherein each of said processors executes a distinct operating system of said multiple operating systems**

133. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation.

134. Gulick discloses “one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system and allows the different partitions to communicate with one another through shared memory.” EX1005, 8:47-53. Gulick also discloses that “each partition operates under the control of its own operating system” and that “[t]he operating system executing on different ones of the partitions may be *different operating systems*, or *different instances* of the same operating system.” EX1005, 9:42-46; *see also*

EX1005, 13:27-29 (“the present invention contemplates an environment wherein ***at least two of the operating systems are different*** and one operating system does not control or manage the second operating system.”).

135. Therefore, I believe that Gulick-AAPA discloses and/or suggests this limitation.

4. Claim 7

a. A multiprocessing system comprising:

136. To the extent that the preamble of claim 7 is limiting, it is my opinion that Gulick-AAPA discloses the limitations of the preamble. *See supra* Section IX.A.1.a; *infra* Section IX.A.4.b-d.

b. a plurality of processor groups mounted on a single die

137. Claim limitation 7.b varies from claim limitation 1.b in that it recites “groups” while claim limitation 1.b does not. *See supra* Section IX.A.1.b. Gulick-AAPA discloses claim limitation 7.b for at least the same reasons discussed above from claim limitation 1.b.

138. Further, each processor in the pairs of processors (e.g., processors 240A/240B-240R/240S) are mounted on a single die. *See supra* Section IX.A.1.b. These pairs of processors are necessarily groups in and of themselves. Even further, Gulick discloses groups of processors as “one or more Pods or sub-pods” with each pod or sub-pod operating “under the control of a separate operating system.”

EX1005, 13:21-24; *see also* EX1005, 11:45-60 (“each TLC 230 is connected to four processors” and “[e]ach TLC 230 and its respective four processors define a sub-Pod.”); EX1005, 11:60-62 (Two sub-Pods “are connected via crossbar interconnect (e.g., crossbar interconnect 290A or 290B) to form a Pod”).

139. Thus, I believe that Gulick-AAPA discloses and/or suggests this limitation.

c. multiple operating systems residing in a memory connected to said groups

140. Claim limitation 7.c varies slightly from claim limitation 1.c in that it recites “groups” while claim limitation 1.c does not. *See supra* Section IX.A.1.c. It is my opinion that Gulick-AAPA discloses claim limitation 7.c for at least the same reasons discussed above from claim limitation 1.c. *See supra* Section IX.A.1.c.

d. wherein each of said groups . . . two or more operating systems of said multiple operating systems

141. Claim limitation 7.d varies slightly from claim limitation 1.d in that it recites “groups” while claim limitation 1.d does not. *See supra* Section IX.A.1.d. It is my opinion that Gulick-AAPA discloses limitation 7.d for at least the same reasons discussed above from claim limitation 1.d. *See supra* Sections IX.A.1.d.

5. Claim 11

- a. The system of claim 7, wherein each of said groups executes a distinct operating system of said multiple operating systems.**

142. Claim 11 varies slightly from claim 5 (“wherein each of said processors executes a distinct operating system of said multiple operating systems”) in that it recites “groups” while claim 5 does not. *See supra* Section IX.A.3. It is my opinion that Gulick-AAPA discloses claim 11 for at least the same reasons discussed above from claim 5. *See supra* Section IX.A.3.

6. Claim 12

- a. A multiprocessing apparatus comprising:**

143. To the extent that the preamble of claim 12 is limiting, it is my opinion that Gulick-AAPA discloses the limitations of the preamble for at least the reasons set forth above for claim limitation 1.a. *See supra* Section IX.A.1.a; *see infra* Section IX.A.6.b-d.

- b. processor means for executing a plurality of operating system means, wherein said processor means includes a plurality of processors mounted on a single die**

144. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation for at least the same reasons discussed above from claim limitations 1.b-1.d, and for the additional reasons provided below. *See supra* Sections IX.A.1.b-IX.A.1.d.

(1) operating system means

145. It is my opinion that Gulick-AAPA discloses and/or suggests an “operating system means,” as claimed. *See supra* Section VIII.E (construction of “operating system means” as a conventional operating system, such as WINDOWS NT, UNIX, and the like, or equivalents thereof).

146. Gulick’s system includes operating systems. *See supra* Section IX.A.1.c; EX1005, 11:1-10 (“According to the present invention, one or more processor modules may be configured as a separate partition within the computer system . . . each partition operating under the control of a separate operating system. For example, each processor module 110, 112 and 114 of FIG. 1 can be defined as a separate partition controlled via a separate operating system 170, 172 and 174. Each operating system 170, 172 and 174 views main memory separately as though each is the only entity accessing main memory 160.”). Gulick teaches that “the term ‘operating system’ refers to the program code that controls and coordinates the use of the hardware among the various application programs for various users.” EX1005, 10:4-7. It is my opinion that the operating system can be inclusive of “off-the-shelf” operating systems “such as a version of Unix or Windows NT.” EX1005, 11:33-44.

147. Thus, I believe that Gulick-AAPA teaches “operating system means.”

(2) processor means

148. It is my opinion that Gulick-AAPA discloses and/or suggests a “processor means,” as claimed. *See* Section VIII.D (construction of “processor means”). For example, Gulick-AAPA discloses a computer system that includes processor modules 110-114. *See supra* Section IX.A.1.b; EX1005, 10:53-54. Gulick-AAPA also discloses a computer system that includes processor pairs 240A/240B- 240R/240S. *See supra* Section IX.A.1.b; EX1005, FIG. 2. As discussed above, a POSITA would have understood that processor modules 110-114 and/or processor pairs 240A/240B-240R/240S can be mounted on a single die. *See supra* Section IX.A.1.b.

149. Thus, for this reason alone, I believe that Gulick-AAPA discloses “processor means.”

150. Further, to the extent the Board adopts the agreed-upon District Court construction, it is my opinion that Gulick-AAPA likewise discloses this limitation.

151. As I understand, agreed-upon function for “processor means” is “executing a plurality of operating system means.” *See supra* Section VIII.D. Gulick-AAPA system discloses that processors 110-114 and/or processors 240A/240B-240R/240S execute a plurality of operating systems. *See supra* Section IX.A.1.c-IX.A.1.d; *see also* EX1005, 8:47-53 (“one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality

of operating systems to concurrently execute in different partitions within the computer system and allows the different partitions to communicate with one another through shared memory.”).

152. Thus, I believe that Gulick-AAPA discloses the agreed-upon function for “processor means.”

153. Further, to the extent the Board does not adopt the agreed upon function, it is my opinion that Gulick-AAPA also teaches the function of “simultaneously execute two or more operating system means of said plurality of operating system means” under the proposed alternative construction for the reasons stated above. As indicated above, Gulick teaches that “one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system.” EX1005, 8:47-53. Thus, Gulick-AAPA discloses this function as well. The structure corresponding to this alternative function is the same as the agreed-upon construction and is met for the reasons stated below for that construction.

154. It is my opinion that Gulick-AAPA also discloses the agreed-upon structure for “processor means,” which is the same as the agreed upon construction.

155. As I understand, the agreed-upon structure for “processor means” is “chip multiprocessor 350 having multiple processors 305-320 mounted on a single

die 360 as shown in Figure 3 and associated descriptions in the specification; or chip multiprocessor 450 having multiple processors 405-420 mounted on a single die 460 shown in Figure 4 and associated descriptions in the specification; and equivalents thereof.” *See supra* Section VIII.D.

156. The '014 Patent discloses chip multiprocessor 350 (left, green). Similarly, Gulick discloses a chip multiprocessor (right, green).




FIG. 3

The '014 Patent





Figure 1


Gulick

157. The '014 Patent discloses that chip multiprocessor 250 includes four processors 305-320 (left, red). Similarly, Gulick discloses that computer system 200 includes pairs of processors – processors 240A/240B-240R/240S. EX1005, FIG. 2. Gulick discloses that computer system 200 includes “multiple partitions . . . each of which may comprise one or more Pods or sub-pods, each operates under the control of a separate operating system.” EX1005, 13:21-24. Gulick explicitly

discloses that each TLC 230 and its four respective processors define a sub-pod.
EX1005, 11:45-60.



The '014 Patent



Gulick

158. Further, as indicated in Section IX.A.1.b, I believe it would be obvious to modify Gulick's teachings with AAPA to have processors 240A/240B-240R/240S mounted on a single die.

159. In addition, the structure disclosed by Gulick-AAPA corresponds to the structure disclosed in the '014 patent. The '014 Patent discloses chip multiprocessor 450 (left, green). Similarly, Gulick discloses a chip multiprocessor (right, green).




FIG. 4
The '014 Patent




Figure 1
Gulick

160. The '014 Patent discloses that chip multiprocessor 450 has multiple processors 405-420 (left, red). Similarly, Gulick discloses that computer system 200 includes pairs of processors – processors 240A/240B-240R/240S. EX1005, FIG. 2. Gulick discloses that computer system 200 includes “multiple partitions . . . each of which may comprise one or more Pods or sub-pods, each operates under the control of a separate operating system.” EX1005, 13:21-24. Gulick explicitly discloses that each TLC 230 and its four respective processors define a sub-pod. EX1005, 11:45-60.




FIG. 4
The '014 Patent




Figure 2

Gulick

161. Further, as indicated in Section IX.A.1.b, I believe it would be obvious to modify Gulick's teachings with AAPA to have processors 240A/240B-240R/240S mounted on a single die.

162. Additionally, to the extent that "processor means" is not a means plus function limitation, it is my opinion that "processor means" should be construed according to its plain meaning and that Gulick-AAPA discloses "processor means" under its plain meaning. *See supra* Section IX.A.6.b(2). The plain meaning is disclosed by Gulick-AAPA for the reasons already discussed.

163. Thus, it is my opinion that Gulick-AAPA system discloses the corresponding structure.

- c. and wherein said processor means is operable to simultaneously execute two or more operating system means of said plurality of operating system means**

164. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation for at least the same reasons discussed above from claim limitation 1.d. *See supra* Section IX.A.1.d.

- d. memory means for storing said plurality of operating system means**

165. It is my opinion that Gulick-AAPA discloses and/or suggests this limitation for at least the same reasons discussed above from claim limitations 1.c-1.d and for the additional reasons explained below. *See supra* Section IX.A.1.c-IX.A.1.d.

166. It is my opinion that Gulick-AAPA discloses and/or suggests a “memory means for storing said plurality of operating system means” as claimed.

167. It is my opinion that Gulick-AAPA teaches the identified function of “memory means.” As I understand, the identified function of “memory means” is “storing said plurality of operating system means.” *See supra* Sections IX.A.1.c-IX.A.1.d.

168. Gulick discloses that main memory 160 “includes a plurality of memory storage units (MSUs) 220, such as memory storage units 220A, 220B, 220C, and 220D.” EX1005, 11:37-39. In FIG. 5 of Gulick, the memory

configuration of computer system 200 includes “three operating systems” each of which “has its own address space 502.” EX1005, 13:30-36.

169. Thus, it is my opinion that Gulick-AAPA discloses the identified function of “memory means.”

170. It is my opinion that Gulick-AAPA also discloses the identified structure of “memory means.” As I understand, the identified structure of “memory means” is “SRAM and/or DRAM on the same chip as one or more processors; SRAM and/or DRAM on separate chips connected to one or more processors; magnetic media, such as tape or disk; optical media, such as CD-ROM; and the like.” *See supra* Section VIII.F.

171. Gulick defines the term “memory storage unit” as “a memory space capable of storing information.” EX1005, 10:21-22. Gulick teaches that “[e]ach memory storage unit includes a plurality of memory storage units, sometimes referred to as banks of DRAM (Dynamic Random Access Memory).” EX1005, 10:22-25.

172. Thus, I believe that Gulick-AAPA discloses the identified structure of “memory means.”

173. Therefore, it is my opinion that Gulick-AAPA discloses this limitation.

7. Claim 13

- a. The apparatus of claim 12, wherein said processor means is connected to said memory means via a bus to execute said multiple operating systems.**

174. Claim 13 varies slightly from claim 3 in that it recites a “processor means” and “memory means” while claim 3 does not. *See supra* Section IX.A.3. As discussed above, it is my opinion that Gulick-AAPA discloses a “processor means” (*see supra* Section IX.A.6.b) and “memory means” (*see supra* Section IX.A.6.d).

175. Accordingly, I believe that Gulick-AAPA discloses claim 13 for at least the same reasons discussed above from claim 3. *See supra* Section IX.A.2.

8. Claim 15

- a. The apparatus of claim 12, wherein each of said processors of said processor means is operable to execute a distinct one of said plurality of operating system means.**

176. Claim 15 varies slightly from claim 5 in that it recites a “processor means” and “operating system means” while claim 5 does not. *See supra* Section IX.A.3. As discussed above, it is my opinion that Gulick-AAPA system discloses a “processor means” (*see supra* Section IX.A.6.b(2)) and “operating system means” (*see supra* Section IX.A.6.b(1)).

177. Accordingly, I believe that Gulick-AAPA discloses claim 15 for at least the same reasons discussed above from claim 5. *See supra* Section IX.A.3.

B. Ground 2: Claims 1, 3, 5, 7, 11-13, 15 are obvious over the combination of AAPA and Gulick

178. In my opinion, the combination of AAPA and Gulick renders claims 1, 3, 5, 7, 11-13, and 15 of the '014 patent obvious.

1. Claim 1

a. A multiprocessing system comprising

179. To the extent the preambles are limiting, it is my opinion that the AAPA-Gulick combination (hereinafter “AAPA-Gulick”) discloses the limitations therein.

180. AAPA discloses that multiprocessing systems with a multiprocessing chip were well-known, and it depicts a known prior art system single-die multiprocessor in Figure 2 below. EX1001, 1:34-59.




FIG. 2
Prior Art

EX1001, FIG. 2 (annotated).

181. AAPA discloses a prior art multiprocessing system 200 that includes a chip multiprocessor 295 that “includes multiple processors 210-240 on a single die 290” (highlighted orange). EX1001, 1:34-37.

182. Therefore, I believe that AAPA-Gulick discloses and/or suggests this limitation.

b. multiple processors mounted on a single die

183. It is my opinion that AAPA-Gulick discloses this limitation.

184. AAPA discloses that its chip multiprocessor 295 includes “multiple processors 210-240 on a single die 290.” EX1001, 1:36-37, 1:41-42.




FIG. 2
Prior Art

EX1001, FIG. 2 (annotated).

185. Therefore, I believe that AAPA-Gulick discloses and/or suggests this limitation.

c. multiple operating systems residing in a memory connected to said multiple processors

186. It is my opinion that AAPA-Gulick discloses and/or suggests this limitation.

187. AAPA discloses that “processors 210-240 are connected to a single operating system 250 stored in a memory 260.” EX1001, 1:38-40.




FIG. 2
Prior Art

EX1001, FIG. 2 (annotated).

188. Accordingly, I believe that AAPA discloses that multiple processors 210-240 on a single die 290 communicate with an operating system 250 residing in memory 260.

189. To the extent that AAPA does not disclose that there are multiple operating systems residing in memory, it is my opinion that Gulick teaches doing so.

190. Gulick discloses, as shown below, “a multi-processor system that includes processor modules 110, 112, and 114.” EX1005, 10:53-54.




Figure 1

EX1005, FIG. 1.

191. Gulick discloses a more detailed description of an embodiment of the multi-processor system of FIG. 1 in FIG. 2. For example, as shown below, Gulick discloses “a computing system 200 . . . [that] includes a main memory . . . and a plurality of processing modules 240 connected to the main memory.” EX1005, 11:22-26, FIG. 2.





Figure 2

EX1005, FIG. 2.

192. Gulick discloses that its multi-processor system includes main memory 160. EX1005, 10:53-67. Main memory 160 “is a directory-based memory system and is capable of supporting various memory consistency models.” EX1005, 11:33-37. Main memory 160 “includes a plurality of memory storage units (MSUs) 220, such as memory storage units 220A, 220B, 220C, and 220D.” EX1005, 11:37-39.

193. FIG. 5 (reproduced below) illustrates “an exemplary memory configuration that can be generated on the computer system of FIG. 2.” EX1005, 13:30-31. As shown, the memory configuration of computer system 200 includes

“three operating systems” each of which “has its own address space 502” in the main memory. EX1005, 13:30-36.



EX1005, FIG. 5 (annotated).

194. FIG. 5 illustrates three exclusive memory windows (annotated red, green, blue, respectively) that “are defined within the address space 504 of main memory 160” and a shared memory (annotated yellow), “which is accessible by all three operating systems.” EX1005, 13:37-43; FIG. 5. Gulick discloses that a memory window “is defined within the main memory to which multiple partitions have shared access.” EX1005, abstract. Thus, Gulick discloses that a memory window is a portion of the main memory.

195. Gulick discloses that each operating system resides in a respective memory window in main memory. Gulick discloses that “[f]or each different partition configured within the computer system, the main memory may further have defined therein an exclusive memory window to which only that partition has access and in which the operating system of that partition executes.” EX1005, 2:66-3:5.

196. It is my opinion that a POSITA understands that if an operating system executes within a memory window assigned to the partition, then that operating system is stored in the memory window because the processor executes the operating system from the memory address where it is stored. In this regard, Gulick also discloses that the “operating system is the first program code loaded into the main memory of a computer system after the computer system is turned on” and that the “central core of the operating system resides in the memory space at all times.” EX1005, 10:7-11. It is my opinion that a POSITA would understand this teaching to refer to each operating system residing in its memory window during operation. Thus, I believe that Gulick discloses “multiple operating systems residing in a memory.”

197. It is my opinion that a POSITA would have been motivated in light of Gulick to configure the system in the AAPA so that multiple processors reside in its memory as claimed.

198. Specifically, Gulick discloses that conventional computer systems typically employ a single operating system. EX1005, 2:29-31. Gulick further discloses that there are limitations with conventional computer systems employing a single operating system. As one example, Gulick discloses that “[a]s modern computer systems grow, and the demands of the user increases, the necessity of employing a plurality of operating systems increases.” EX1005, 2:32-34. Gulick then provides that “[w]hat is needed is a computer system and method for allowing multiple operating systems, including different operating systems, to operate in different partitions on the computer systems.” EX1005, 2:37-40.

199. I believe that the solution in Gulick addresses the very problems described in the '014 Patent. For example, AAPA describes that as modern systems grow, the efficiency of operating systems goes down dramatically. EX1001, 1:50-53 (“as the number of processors managed by a single operating increases, the efficiency of the operating system goes down dramatically.”).

200. In other words, the solution in Gulick provides a known technique (e.g., increasing the number of operating systems employed by a multiprocessor computing device) to a known device (e.g., AAPA’s multiprocessor computer system 200) ready for improvement (as conceded by the AAPA) to yield predictable results (e.g., increasing the efficiency of the operating system of computer system

200 by increasing the number of operating systems employed by computer system 200). *See* MPEP 2143(I)(D).


201. For this reason alone, I believe that a POSITA would have been motivated in light of Gulick to configure the memory as described by AAPA so that its memory includes multiple operating systems residing therein as claimed.

202. Further, as indicated above, in order to facilitate use of multiple operating systems, Gulick partitions its main memory 160. EX1005, 2:46-50 (“The present invention is directed to a computer system and methods for allowing multiple operating systems to operating in different partitions within a single computer architecture and for allowing the different partitions to communicate with one another through shared memory”); 8:47-53 (“The present invention is directed to a multi-processor computer system, having . . . a main memory having one or more storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system”); 8:53-56 (“The main memory is divided into a plurality of memory storage units”); 8:55-56 (“The main memory is allocated among the different partitions.”); 8:67-9:2.

203. Gulick broadly defines the term partition as “one or more processing module(s) that are under the control of a single operating system. The term partition is used herein to refer, in whole or in part, to the processing module(s) of the partition, the operating system executing on the partition, any exclusive memory

window assigned to the partition, other clients or application programs executing on the partition, or any combination thereof.” EX1005, 10:28-34.

204. Gulick shows an exemplary partitioned memory in FIG. 5, which illustrates “an exemplary memory configuration that can be generated on the computer system of FIG. 2.” EX1005, 13:30-31. For example, as shown in the memory configuration illustrated in FIG. 5 (reproduced below), the memory configuration of computer system 200 includes “three operating systems” each of which “has its own address space 502.” EX1005, 13:30-36.



EX1005, FIG. 5 (annotated).

205. It is my opinion that a POSITA would be motivated to partition memory 250 of the AAPA into the multiple memory windows disclosed by Gulick in order

to allow computer system 200 to employ multiple operating systems taught by Gulick. For example, as described above and exemplified in an annotated FIG. 2 of the AAPA below, partitioning memory 260 of computer system 200 would involve applying the known partitioning techniques described in Gulick in order for computer system 200 to employ multiple operating systems in its memory 260.




FIG. 2
Prior Art

Annotated
based on
Gulick

FIG. 2 (annotated).

206. I believe that such an implementation would have been a straightforward combination of well-known technologies using known methods and would have had predictable results.

207. Further, to the extent that the Board adopts the narrower agreed-upon District Court construction (“multiple operating systems residing in the same memory”) rather than plain meaning, it is my opinion that AAPA-Gulick still discloses and/or suggests this limitation. As illustrated above, the multiple operating systems are stored in the same memory 260. *See supra* Section IX.A.1.d.

208. Therefore, I believe that AAPA-Gulick discloses and/or suggests this limitation.

- d. wherein each of said multiple processors executes an operating system of said multiple operating systems, and two or more of said multiple processors are capable of simultaneously executing two or more operating systems of said multiple operating systems**

209. It is my opinion that AAPA-Gulick discloses and/or suggests this limitation.

210. AAPA-Gulick includes multiple processors mounted on a single die and multiple operating systems residing in memory that are connected to these multiple processors. *See supra* Sections IX.B.1.b-c. I believe that each of these processors executes an operating system from the multiple operating systems, whereby two or more of the processors are capable of simultaneously executing two or more of the operating systems.

211. AAPA discloses that processors 210-240 are connected to memory 260. EX1001, 1:37-40. It is my opinion that a POSITA would understand that, when AAPA's memory 260 is partitioned as taught by Gulick to include multiple operating systems, each of the processors 210-240 would execute at least one of the operating systems.

212. To the extent that AAPA does not explicitly disclose that processors 210-240 are capable of simultaneously executing two or more of the operating systems, I believe Gulick teaches doing so.

213. Gulick discloses that one or more processors is allocated to execute an operating system. Processor modules 110-114 "may be configured as a separate partition within the computer system, such that multiple partitions may exist within the computer system." EX1005, 11:1-3. Each partition operates "under the control of a separate operating system." EX1005, 11:4-5. Gulick discloses that "each processor module 110, 112, and 114 . . . can be defined as a separate partition controlled via a separate operating system 170, 172, and 174." EX1005, 11:5-10, FIG. 1.

214. Gulick further discloses that the "two or more of said multiple processors are capable of simultaneously executing two or more operating systems of said multiple operating systems," as claimed. For example, Gulick disclose that the foregoing arrangement of components "allows a plurality of operating system

to *concurrently execute* in different partitions within the computer system.” EX1005, 8:47-53. Thus, Gulick teaches multiple processors capable of concurrently or simultaneously executing two or more operating systems.

215. It is my opinion that a POSITA would be motivated by the teachings of Gulick to modify the operations of the AAPA such that the operating systems 210-240 of AAPA could simultaneously execute two or more operating systems residing in memory 260. Such combination involves applying a known technique (e.g., executing two or more operating systems concurrently as described in Gulick) to a known device (e.g., AAPA’s computer system 200) to yield predictable results (e.g., executing two or more operating systems within AAPA’s computer system 200).

216. Thus, I believe that AAPA-Gulick discloses this limitation.

217. Further, to the extent that the Board adopts the agreed-upon District Court construction of “simultaneously executing two or more operating systems” as “simultaneously executing two or more independent operating systems,” I believe that AAPA-Gulick still teaches this limitation.

218. Gulick clearly states that “each partition operates under the control of its own operating system” and that “[t]he operating system executing on different ones of the partitions may be different operating systems, or different instances of the same operating system.” EX1005, 9:42-46; *see also* EX1005, 13:27-29 (“the present invention contemplates an environment wherein at least two of the operating

systems are different and one operating system does not control or manage the second operating system.”). I believe that a POSITA would recognize that Gulick’s disclosure of different operating systems that do not control or manage other operating systems is an example of independent operating systems.

219. Thus, it is my opinion that AAPA-Gulick teaches this limitation under the agreed-upon construction.

2. Claim 3

a. The system of claim 1, wherein said multiple processors are connected to said memory via a bus to execute said multiple operating systems

220. It is my opinion that AAPA-Gulick discloses and/or suggests this limitation.

221. AAPA-Gulick discloses that multiple processors are connected to memory to execute the operating systems. *See supra* Section IX.B.1. AAPA further discloses that the computer system 200 establishes connections between processors 210-240 and memory 260 via a bus 270. EX1001, 1:40-41 (“The processors 210-240 may communicate with the memory 260 via a bus 270”).

222. Accordingly, I believe a POSITA would have understood that the system in AAPA-Gulick would have included a bus that connects the multiple processors with the main memory to execute instructions from the various operating systems and process data.

223. Thus, it is my opinion that AAPA-Gulick discloses and/or suggests this limitation.

224. Further, to the extent that the Board adopts the agreed-upon District Court construction (“multiple processors are connected to said memory via the same bus”), I believe that AAPA-Gulick likewise discloses and/or suggests this limitation.

225. AAPA discloses that “processors 210-240 may communicate with the memory 260 via a bus 270.” EX1001, 1:40-41; FIG. 2. Only a single bus 270 (highlighted purple) is illustrated in FIG. 2 of the '014 Patent. EX1001, FIG. 2.




FIG. 2
Prior Art

EX1001, FIG. 2.

226. Accordingly, under the agreed-upon construction, I believe that AAPA-Gulick still discloses and/or suggests this limitation.

3. Claim 5

- a. The system of claim 1, wherein each of said processors executes a distinct operating system of said multiple operating systems**

227. It is my opinion that AAPA-Gulick discloses and/or suggests this limitation.

228. AAPA-Gulick discloses that multiple processors are connected to memory to execute multiple operating systems. *See supra* Section IX.B.1. AAPA-Gulick further discloses that the computer system 200, or, more specifically, processors 210-240, are capable of “simultaneously executing two or more independent operating systems.” *See supra* Section IX.B.1.d.

229. Gulick further discloses that the operating systems residing in memory may be different operating systems. For example, Gulick discloses “[t]he operating systems executing on different ones of the partitions may be different operating systems, or different instances of the same operating system.” EX1005, 9:42-46. Gulick further discloses that “the present invention contemplates an environment wherein at least two of the operating systems are different and one operating system does not control or manage the second operating system.” EX1005, 13:27-29.

230. It is my opinion that a POSITA would be motivated by the teachings of Gulick to modify the operations of the AAPA such that the processors 210-240 of AAPA could execute a distinct operating system of the operating systems residing in memory 260. As discussed for claim 1 above, I believe such a combination involves a simple substitution of one known element for another to obtain predictable results. *See supra* Section IX.A.1.c.

4. Claim 7

a. A multiprocessing system comprising:

231. To the extent that the preamble of claim 7 is limiting, it is my opinion that AAPA-Gulick discloses the limitations of the preamble for at least the reasons set forth above for claim limitation 1.a. *See supra* Section IX.B.1.a; *infra* Section IX.B.4.b-d.

b. a plurality of processor groups mounted on a single die

232. Claim limitation 7.b varies from claim limitation 1.b in that it recites “groups” while claim limitation 1.b does not. *See supra* Section IX.B.1.b. AAPA-Gulick discloses claim limitation 7.b for at least the same reasons discussed above for claim limitation 1.b.

233. Further, to the extent that “group” is construed as containing more than one processor (which Petitioner does not concede), it is my opinion that AAPA-Gulick still discloses claim limitation 7.B.

234. AAPA discloses that multiple processors are managed by a single operating system. EX1001, 1:50-54 (“as the number of processors managed by a single operating system increases, the efficiency of the operating system goes down dramatically.”). Further, AAPA shows, in FIG. 2 reproduced below, a group of processors 210-240 on a single die 290. EX1001, FIG. 2.




FIG. 2
Prior Art

EX1001, FIG. 2.

235. As used in the '014 patent, the term “processor group may include one or more processors mounted on the die.” EX1001, 2:24-25. Therefore, a POSITA would understand that AAPA processors 210-240 in FIG. 2 comprise processor groups.

236. Accordingly, I believe that AAPA discloses “processor groups mounted on a single die.”

237. To the extent that AAPA does not explicitly disclose *a plurality* of processor groups mounted on a single die, I believe that Gulick teaches doing so.

238. For example, AAPA-Gulick discloses each processor in the pairs of processors (e.g., processor 240A/240B-240R/240S) are mounted on a single die. *See supra* Section IX.B.1.b. These pairs of processors are groups. For example, Gulick discloses groups of processors as “one or more Pods or sub-pods” with each pod or sub-pod operating “under the control of a separate operating system.” EX1005, 13:21-24; *see also* EX1005, 11:45-60 (“each TLC 230 is connected to four processors” and “[e]ach TLC 230 and its respective four processors define a sub-Pod.”); EX1005, 11:60-62 (Two sub-Pods “are connected via crossbar interconnect (e.g., crossbar interconnect 290A or 290B) to form a Pod”).

239. It is my opinion that a POSITA would be motivated by the teachings of Gulick to modify the multiprocessing system of the AAPA such that the processors 210-240 of AAPA could be combined into groups (e.g., pods or sub-pods) as taught by Gulick to improve, for example, overall performance of the AAPA’s multiprocessor system by devoting multiple processors to each operating system instance. *See e.g.*, EX1004, 13:21-24; 11:45-62.

240. Further I believe that such modification of AAPA based on the teachings of Gulick involves a simple substitution of one known element for another to obtain predictable results.

241. Thus, I believe that AAPA-Gulick discloses this limitation.

c. multiple operating systems residing in a memory connected to said groups

242. Claim limitation 7.c varies slightly from claim limitation 1.c in that it recites “groups” while claim limitation 1.c does not. *See supra* Section IX.B.1.c. It is my opinion that AAPA-Gulick discloses claim limitation 7.c for at least the same reasons discussed above from claim limitation 1.c. *See supra* Section IX.B.1.c.

d. wherein each of said groups . . . two or more operating systems of said multiple operating systems

243. Claim limitation 7.d varies slightly from claim limitation 1.d in that it recites “groups” while claim limitation 1.d does not. *See supra* Section IX.B.1.d. It is my opinion that AAPA-Gulick system discloses limitation 7.d for at least the same reasons discussed above from claim limitation 1.d. *See supra* Section IX.B.1.d.

5. Claim 11

- a. The system of claim 7, wherein each of said groups executes a distinct operating system of said multiple operating systems.**

244. Claim 11 varies slightly from claim 5 in that it recites “groups” while claim 5 does not. *See supra* Section IX.B.3. It is my opinion that AAPA-Gulick discloses claim 11 for at least the same reasons discussed above from claim 5. *See supra* Section IX.B.3.

6. Claim 12

- a. A multiprocessing apparatus comprising:**

245. To the extent that the preamble of claim 12 is limiting, it is my opinion that AAPA-Gulick discloses the limitations of the preamble for at least the reasons set forth above for claim limitation 1.a. *See supra* Section IX.B.1.a; *see infra* Section IX.B.6.b-d.

- b. processor means for executing a plurality of operating system means, wherein said processor means includes a plurality of processors mounted on a single die**

246. It is my opinion that AAPA-Gulick discloses and/or suggests this limitation for at least the same reasons discussed above for claim limitations 1.b-1.d, and for the additional reasons provided below. *See supra* Sections IX.B.1.b-IX.B.1.d.

(1) Operating system means

247. It is my opinion that AAPA-Gulick discloses and/or suggests “operating system means,” as claimed. *See supra* Section VIII.E (construction of “operating system means” as a conventional operating system, such as WINDOWS NT, UNIX, and the like, or equivalents thereof).

248. AAPA discloses computing system 200 that includes an “operating system 250 stored in memory 260.” EX1001, 1:39-40. AAPA describes computing system 200 and its components as a “second conventional multiprocessing scheme.” EX1001, 1:34. I believe that a POSITA would understand that, at the time of invention, “operating system 250” of the AAPA would be inclusive of “a conventional operating system, such as WINDOWS NT, UNIX, and the like, or equivalents thereof.”

249. Thus, it is my opinion that AAPA-Gulick teaches “operating system means.”

(2) Processor means

250. It is my opinion that AAPA-Gulick discloses and/or suggests “a processor means” as claimed. *See supra* Section VIII.D (construction of “processor means”).

251. I believe AAPA-Gulick discloses the agreed-upon function. As I understand, the agreed-upon function for “processor means” is “executing a

plurality of operating system means.” *See supra* Section VIII.D. AAPA discloses that “processors 210-240 are connected to a single operating system 240 stored in memory 260.” EX1001, 1:38-40. As a POSITA understands, processors 210-240 execute operating system 240 stored in memory 260.

252. Although AAPA does not explicitly state that processors 210-240 execute a *plurality* of operating system means, I believe that a POSITA would find it obvious to modify the AAPA to do so based on the teachings of Gulick, as described above. *See supra* Section IX.B.1.c.



253. Further, to the extent the Board does not adopt the agreed upon construction, it is my opinion that AAPA-Gulick also teaches the function of “simultaneously execute two or more operating system means of said plurality of operating system means” under Petitioner’s proposed alternative construction for the reasons stated above. For example, as stated above, I believe a POSITA would find it obvious to combine the multiple operating systems of Gulick into the multi-processor system described in AAPA. *See supra* Section IX.B.1.c. Further, Gulick discloses that the ability to simultaneously execute two or more of its operating systems. EX1005, 8:47-53 (“one or more processor modules and a main memory having one or more memory storage units . . . that allows a plurality of operating systems to concurrently execute in different partitions within the computer system and allows the different partitions to communicate with one another through shared

memory.”). Thus, I believe that AAPA-Gulick discloses this function as well. It is also my position that the structure corresponding to this alternative function is the same as the agreed upon construction and is met for the reasons stated below for that construction.



254. Thus, it is my opinion that AAPA-Gulick discloses the agreed-upon function for “processor means.”

255. It is my opinion that AAPA-Gulick also discloses the corresponding structure for “processor means.” As indicated above, it is my understanding that the agreed-upon structure for “processor means” is “chip multiprocessor 350 having multiple processors 305-320 mounted on a single die 360 as shown in Figure 3 and associated descriptions in the specification; or chip multiprocessor 450 having multiple processors 405-420 mounted on a single die 460 shown in Figure 4 and associated descriptions in the specification; and equivalents thereof.” *See supra* Section VIII.D.

256. As shown below, the '014 Patent discloses computing system 300 having chip multiprocessor 350 (left, green). Similarly, AAPA discloses computing system 200 having chip multiprocessor 295 (right, green).




257. The '014 Patent discloses that chip multiprocessor 350 includes four processors 305-320 (left, red). Similarly, AAPA discloses that computer system 200 includes four processor 210-240 (right, red).




258. Thus, I believe that AAPA-Gulick system discloses the agreed upon structure of chip multiprocessor 350.

259. Further, as shown below, the '014 Patent discloses chip multiprocessor 450 (left, green). Similarly, AAPA discloses a chip multiprocessor (right, green).



260. The '014 Patent discloses that chip multiprocessor 450 includes four processors 405-420 (left, red). Similarly, AAPA discloses that computer system 200 includes four processor 210-240 (right, red).



261. Further, as indicated in Section IX.B.4.b, I believe that it would be obvious to modify AAPA's teachings with Gulick to include groups of processors, similar to chip multiprocessor 450. *See supra* Section IX.B.4.b.

262. Thus, believe that AAPA-Gulick discloses the corresponding structure of chip multiprocessor 450.

263. Therefore, to the extent "processor means" is construed under the agreed-upon construction, it is my opinion that AAPA-Gulick teaches this limitation.

264. Additionally, to the extent that "processor means" is not a means plus function limitation, "processor means" should be construed according to its plain meaning. It is my opinion that the plain meaning is disclosed by AAPA-Gulick for the reasons already discussed.

- c. and wherein said processor means is operable to simultaneously execute two or more operating system means of said plurality of operating system means**

265. It is my opinion that AAPA-Gulick discloses and/or suggests limitation 12.c for at least the same reasons discussed above from claim limitation 1.d. *See supra* Section IX.B.1.d.

- d. memory means for storing said plurality of operating system means**

266. It is my opinion that AAPA-Gulick discloses and/or suggests limitation 12.d for at least the same reasons discussed above for claim limitations 1.c-1.d and for the additional reasons explained below. *See supra* Sections IX.B.1.c-IX.B.1.d.

267. I believe that AAPA-Gulick teaches the identified function of “memory means.” It is my understanding that the identified function of the “memory means” is “storing said plurality of operating system means.” *See supra* Section VIII.F. The combined AAPA-Gulick system discloses that the memory means is configured for “storing said plurality of operating systems.” *See supra* Sections IX.B.1.c-IX.B.1.d.

268. Thus, it is my opinion that AAPA-Gulick discloses the identified function of “memory means.”

269. It is my opinion that AAPA-Gulick also discloses the identified structure of “memory means.” It is my understanding that the identified structure of “memory means” is “SRAM and/or DRAM on the same chip as one or more processors; SRAM and/or DRAM on separate chips connected to one or more processors; magnetic media, such as tape or disk; optical media, such as CD-ROM; and the like.” *See supra* Section VIII.F. AAPA discloses computing system 200 that includes a “memory 260.” EX1001, 1:39-40. AAPA describes computing system 200 and its components as a “second conventional multiprocessing scheme.” EX1001, 1:34. I believe that a POSITA would understand that, at the time of

invention, “memory 260” of the AAPA would be inclusive of “SRAM and/or DRAM on the same chip as one or more processors; SRAM and/or DRAM on separate chips connected to one or more processors; magnetic media, such as tape or disk; optical media, such as CD-ROM; and the like.”

270. Thus, I believe that AAPA-Gulick discloses the corresponding structure of “memory means.”

271. Therefore, it is my opinion that AAPA-Gulick discloses this limitation.

7. Claim 13

- a. The apparatus of claim 12, wherein said processor means is connected to said memory means via a bus to execute said multiple operating systems.**

272. Claim 13 varies slightly from claim 3 in that it recites a “processor means” and “memory means” while claim 3 does not. *See supra* Section IX.B.3. It is my opinion that AAPA-Gulick discloses a “processor means” (*see supra* Section IX.B.6.b) and “memory means” (*see supra* Section IX.B.6.d).

273. Accordingly, it is my opinion that AAPA-Gulick discloses claim 13 for at least the same reasons discussed above from claim 3. *See supra* Section IX.B.2.

8. Claim 15

- a. The apparatus of claim 12, wherein each of said processors of said processor means is operable to execute a distinct one of said plurality of operating system means.**


274. Claim 15 varies slightly from claim 5 in that it recites a “processor means” and “operating system means” while claim 5 does not. *See supra* Section IX.B.3. As discussed above, it is my opinion that Gulick-AAPA system discloses a “processor means” (*see supra* Section IX.B.6.b(2)) and “operating system means” (*see supra* Section IX.B.6.b(1)).

275. Accordingly, it is my opinion that AAPA-Gulick discloses claim 15 for at least the same reasons discussed above from claim 5. *See supra* Section IX.B.3.

X. CONCLUSION

I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: 26 May 2023

By: 
Martin Walker, PhD.