

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CHOLLA ENERGY LLC AND OBM, INC.

Petitioners

U.S. Patent No. 11,283,261

**DECLARATION OF TAJANA S. ROSING, PH.D.,
UNDER 37 C.F.R. § 1.68 IN SUPPORT OF PETITION
FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

I.	INTRODUCTION	4
II.	QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	8
III.	LEVEL OF ORDINARY SKILL IN THE ART	18
IV.	RELEVANT LEGAL STANDARDS	19
V.	BACKGROUND TECHNOLOGY.....	20
A.	Electrical Grid Stability in View of Renewable Generator Supply	20
B.	Using Green Datacenters to Harness Curtailed Power From Renewable Power Sources	22
C.	Load Distribution Among Datacenters	26
VI.	OVERVIEW OF THE '261 PATENT	28
VII.	CLAIM CONSTRUCTION	31
VIII.	IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE....	31
IX.	CHALLENGE 1: THE COMBINATION OF KIANI AND PELIO RENDERS OBVIOUS CLAIMS 1-2 AND 9-12.....	32
A.	Summary of Kiani	32
B.	Summary of Pelio	36
C.	Reasons to Combine Kiani with Pelio	38
D.	Detailed Analysis of Claims.....	47
1.	Claim 1	47
2.	Claim 2	86
3.	Claim 9	88
4.	Claim 10	89
5.	Claim 11	90
6.	Claim 12	92
X.	CHALLENGE 2: THE COMBINATION OF KIANI, PELIO, AND CHAPEL RENDERS OBVIOUS CLAIMS 4-8	93
A.	Summary of Chapel.....	93

B.	Reasons to Combine Kiani and Pelio with Chapel	94
C.	Detailed Analysis of Claims.....	101
1.	Claim 4.....	101
2.	Claim 5.....	109
3.	Claim 6.....	111
4.	Claim 7.....	116
5.	Claim 8.....	118
XI.	CHALLENGE 3: THE COMBINATION OF KIANI, PELIO, AND BELADY RENDERS OBVIOUS CLAIM 3	119
A.	Summary of Belady.....	119
B.	Reasons to Combine Kiani and Pelio with Belady	122
C.	Detailed Analysis of Claim	127
1.	Claim 3.....	127
XII.	CHALLENGE 4: THE COMBINATION OF KIANI, PELIO, CHAPEL, AND FORESTIERO RENDERS OBVIOUS CLAIMS 13-16.....	130
A.	Summary of Forestiero.....	130
B.	Reasons to Combine Kiani, Pelio and Chapel with Forestiero	133
C.	Detailed Analysis of Claim	140
1.	Claim 13.....	140
2.	Claim 14.....	152
3.	Claim 15.....	155
4.	Claim 16.....	155
XIII.	CONCLUSION.....	159

I, Tajana S. Rosing, Ph.D., do hereby declare as follows:

I. INTRODUCTION

1. I am making this declaration at the request of Cholla Energy LLC and OBM, Inc. (“Petitioners”) in the matter of the *Inter Partes* Review of U.S. Patent No. 11,283,261 (“the ’261 patent”) to McNamara et al.

2. I am being compensated for my work in this matter at my standard hourly rate. I am also being reimbursed for reasonable and customary expenses associated with my work and testimony in this proceeding. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

3. I have been asked to provide my opinions regarding whether the subject matter of claims 1-16 (“the Challenged Claims”) of the ’261 patent would have been obvious to a person having ordinary skill in the art (“POSITA”) at the time of the alleged invention, in light of the prior art. It is my opinion that the Challenged Claims would have been obvious to a POSITA.

4. In the preparation of this declaration, I have studied:

- a. the ’261 patent, Ex.1001;
- b. the prosecution history of the ’261 patent (“’261 File History”), Ex.1002;

c. Kiani et al., “On the Fundamental Energy Trade-offs of Geographical Load Balancing,” *Green Communications and Computing*, IEEE Communications Magazine, May 2017 (“Kiani”), Ex.1005;

d. U.S. Patent Pub. No. 2011/0316337 to Pelio et al. (“Pelio”), Ex.1006;

e. U.S. Patent Pub. No. 2010/0145542 to Chapel et al. (“Chapel”), Ex.1007;

f. U.S. Patent No. 8,849,469 to Belady et al. (“Belady”), Ex.1008;
and

g. Forestiero et al., “Hierarchical Approach for Efficient Workload Management in Geo-Distributed Data Centers,” *IEEE Transactions on Green Communications and Networking*, Vol.1, No. 1, March 2017 (“Forestiero”), Ex.1009.

5. In forming the opinions expressed below, I have considered: the documents listed above; the relevant legal standards, including the standard for obviousness; and my own knowledge and experience based upon my work in the field of network communications and security as described below, as well as portions of the following additional materials:

a. U.S. Patent Publication No. 2010/0088261 to Montalvo (“Montalvo”), Ex.1010;

- b. Martinot, E., “Grid Integration of Renewable Energy: Flexibility, Innovation, and Experience,” *Annu. Rev. Environ. Resour.*, vol. 41:223-251, 2016 (“Martinot”), Ex.1011;
- c. Yang et al., “ZCCloud: Exploring Wasted Green Power for High-Performance Computing,” 2016 IEEE International Parallel and Distributed Processing Symposium, 23-27 May 2016 (“Yang”), Ex.1012;
- d. Bird et al., “Distributed (green) data centers: A new concept for energy, computing, and telecommunications,” *Energy for Sustainable Development*, Vol. 19, 2014, pp. 83-91 (“Bird”), Ex.1013;
- e. Yang et al., “Large Scale and Extreme Scale Computing with Stranded Green Power: Opportunities and Costs,” *IEEE Transactions on parallel and Distributed Systems*, December 12, 2017 IEEE Xplore Date of Publication (“Yang 2017”), Ex.1014;
- f. Chien et al., “The Zero-Carbon Cloud: High Value, Dispatchable Demand for Renewable Power Generators,” *The Electricity Journal*, Vol. 28, Issue 8, 2015, pp. 110-118 (“Chien”), Ex.1015;
- g. Brocanelli et al., “Reducing the expenses of geo-distributed datacenters with portable containerized modules,” *Performance Evaluation*, Vol. 79, 2014, pp. 104-119 (“Brocanelli”), Ex.1016;

h. Li et al., “SolarCore: Solar Energy Driven Multi-Core Architecture Power Management,” 2011 IEEE 17th International Symposium on High Performance Computer Architecture, February 12-16, 2011 (“Li 2011”), Ex.1017;

i. U.S. Patent Publication No. 2016/0109916 to Li et al. (“Li 2016”), Ex.1018;

j. Zhang et al., “GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy,” 12th International Middleware Conference (MIDDLEWARE), Dec. 2011, Lisbon Portugal (pp. 143-164) (“Zhang”), Ex.1019;

k. Blouin, “Exploring Dual-Power Feeds in Data Centers,” <https://www.belden.com/blogs/data-center/exploring-dual-power-feeds-in-data-centers>, May 17, 2018; accessed February 8, 2024 (“Blouin”), Ex.1020;

l. IEEE Xplore,
<https://ieeexplore.ieee.org/abstract/document/7926937>, Ex.1021;

m. IEEE Xplore,
<https://ieeexplore.ieee.org/abstract/document/7557052>, Ex.1022; and

n. U.S. Patent No. 9,003,216 to Sankar et al. (“Sankar”), Ex.1023.

6. Unless otherwise noted, all **emphasis** in any quoted material has been added. Claim terms are *italicized*.

II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE

7. My complete qualifications and professional experience are described in my *Curriculum Vitae*, a copy of which can be found in Exhibit 1004. The following is a brief summary of my relevant qualifications and professional experience. Among other things, my CV includes a list of my issued patents and publications for at least the last 30 years.

8. I am a distinguished professor, IEEE and ACM Fellow, and the Fratomico Endowed Chair in the Computer Science and Engineering Department at the University of California, San Diego (UCSD), and adjunct distinguished professor in the Electrical and Computer Engineering Department at UCSD. My research focuses on energy-efficient computing. My place of business is located at 9500 Gilman Drive, La Jolla, California. I am over the age of eighteen, and am a citizen of the United States.

9. I am currently a Director of System Energy Efficiency Lab where I am leading a diverse research team on projects relating to system energy-efficiency. I have over thirty years of academic and industry experience in applying, designing, studying, teaching, and writing about energy-efficient computing, including datacenters and their interaction with smart grid and renewable energy resources. My experience has spanned both hardware and software, along with measurements and simulations related to datacenters, smart grid, renewable energy resources,

power, thermal and reliability management (for which I was nominated both IEEE and ACM Fellow).

10. I received an Electrical Engineering Ph.D. degree in 2001 from Stanford University; my thesis was titled “Energy Efficient System Design and Utilization.” Before that, I received an Electrical Engineering B.S. degree in 1992 from Northern Arizona University, an Electrical and Computer Engineering M.S. degree in 1993 from the University of Arizona, and an Engineering Management M.S. degree in 2000 from Stanford University.

11. While at Stanford University, I worked in the same office and on the same machines where Yahoo was started, and just down the hall from the office where Google was started. After completing my M.S. in Electrical and Computer Engineering, I worked at Altera Corporation as a senior design engineer for four years. During and after completing my Ph.D. degree, I worked with Stanford University and Hewlett-Packard Labs leading a team of researchers to develop products for the wireless portable devices market. At HP labs, my team’s efforts also included work related to content delivery networks, streaming video over wireless, videoconferencing using mobile platforms, and energy efficiency of chips, mobile as well as server systems. A partial list of my publications during this period is below.

- T. Simunic, M. Smith, “Dynamic Power Management at HP,” Invited Paper in Special Issue of Design and Test Journal, 2001.

- G. Manjunath, V. Krishnan, T. Simunic, J. Tourrilhes, A. McReynolds, D. Das, V. Srinivasmurthy, A. Srinivasan, “Smart Edge Server: going beyond a wireless access point,” WMASH 04.
- B. Delaney, T. Simunic, N. Jayant, “Energy Aware Distributed Speech Recognition for Wireless Mobile Devices,” Special Issue on Embedded Systems for Multimedia, IEEE Design & Test, 2004.
- Acquaviva, T. Simunic, V. Deolalikar, S. Roy, “Remote Power Control of Wireless Network Interfaces,” Special Issue of Journal of Embedded Computing, No. 3, 2004.
- T. Simunic, W. Quadeer, G. De Micheli, “Managing heterogeneous wireless environments via Hotspot servers,” MMCN 05.
- T. Simunic, L. Benini, P. Glynn, G. De Micheli: “Event-Driven Power Management,” IEEE Transactions on CAD, pp.840-857, July 2001.
- T. Simunic, L. Benini, G. De Micheli: “Energy-Efficient Design of Battery-Powered Embedded Systems,” Special Issue of IEEE Transactions on VLSI, 2001.

12. Through work with Hewlett-Packard Labs, I have been granted the following United States patents:

- U.S. Patent No. 7,272,730, titled “Application-driven method and apparatus for limiting power consumption in a processor-controlled hardware platform,” which issued in 2007 from an application with a priority date in 2003;
- U.S. Patent No. 7,246,181, titled “Device and method for identifying a communication interface that performs an operating parameter closer to a desired performance level than another communication interface performs the operating parameter,” which issued in 2007 from an application with a priority date in 2004; and
- U.S. Patent No. 7,190,980, titled “Method and system for power control in wireless portable devices using wireless channel

characteristics,” which issued in 2007 from an application with a priority date in 2004.

13. I have published over 400 publications and received a number of best paper awards and nominations. I have also been an invited speaker at numerous academic and industry conferences. I have published several papers and given several presentations related to green datacenters/energy-efficient datacenters and the Smart Grid. A partial list of my journal publications prior to the '261 patent is below:

- Christine Chan, Alper Sinan Akyurek, Baris Aksanli, and Tajana S. Rosing. “Optimal Performance-Aware Cooling on Enterprise Servers,” IEEE TCAD’18.
- B. Aksanli, T. S. Rosing, “User Behavior and Flexibility based Energy Management in Residential Neighborhoods,” Special Issue of IEEE Transactions on Emerging Topics in Computing, 2017.
- A. Sinan Akyurek and Tajana Simunic Rosing, “Optimal Distributed Nonlinear Battery Control,” IEEE Journal of Emerging and Selected Topics in Power Electronics, 2017.
- P. Mercati, F. Paterna, A. Bartolini, L. Benini, T. Simunic Rosing, “WARM: Workload-Aware Reliability Management in Linux/Android,” IEEE TCAD 2016.
- R. Ayoub, R. Nath, T. Simunic Rosing, “CoMETC: Coordinated Management of Energy, Thermal, & Cooling in Servers,” TODAES’13.
- B. Aksanli, J. Venkatesh, T. Simunic Rosing, “Datacenter Modeling and Simulation with Focus on Energy Efficiency and Green Energy Integration,” IEEE Computer Special Issue on Modeling and Simulation of Smart and Green Computing Systems, 2012.

- Mohamed M. Sabry, Ayse K. Coskun, David Atienza, Tajana Simunic Rosing, Thomas Brunschwiler, “Energy-Efficient Multi-Objective Thermal Control for Liquid-Cooled 3D Stacked Architectures,” IEEE TCAD, 2011.
- R. Ayoub, K. Indukuri, T. Simunic Rosing, “Temperature Aware Dynamic Workload Scheduling in Multisocket CPU Servers,” IEEE TCAD, 2011.
- G. Dhiman, G. Marchetti, T. Simunic Rosing, “vGreen: A System for Energy Efficient Management of Virtualized Environments,” Special Issue of ACM TODAES, 2010. Top most downloaded paper ‘10-’11.
- J. Recas, C. B, T. Simunic Rosing, D. Atienza, “HOLLOWS: A Power-aware Task Scheduler for Energy Harvesting Sensor Nodes,” Journal of Intelligent Material Systems and Structures, 2010.
- A. Coskun, T. Simunic Rosing, “Utilizing Predictors for Efficient Thermal Management in Multiprocessor SoCs,” IEEE TCAD, 2009.
- G. Dhiman, T. Simunic Rosing, “Using online learning for system level power management,” IEEE TCAD, 2009.
- T. Simunic Rosing, K. Mihic, G. De Micheli, “Power and reliability management of SOCs,” IEEE Transactions on VLSI, 2007.
- G. Park, T. Simunic Rosing, M. Todd, C. Farrar, W. Hodgkiss, “Energy Harvesting for Structural Health Monitoring in Sensor Networks,” ASCE Journal, 2007.

14. A partial list of my conference publications prior to the '261 patent is

below:

- C. Chan, M. Ostertag, A. S. Akyurek, T. S. Rosing, “Context-Aware System Design,” Invited paper to SPIE 2017.

- B. Aksanli, J. Venkatesh, C. Chan, A. S. Akyurek, T. S. Rosing, “Context-Aware and User-Centric Residential Energy Management,” PerIoT, 2017.
- Akanksha Maurya, Alper Sinan Akyurek, Baris Aksanli and Tajana Rosing, “Time-Series Clustering for Data Analysis in Smart Grid,” SmartGridComm, 2016.
- B. Aksanli, A. S. Akyurek, T. S. Rosing, “User Behavior Modeling for Estimating Residential Energy Consumption,” Invited paper at SGSC’15.
- J. Venkatesh, S. Chen, P. Tinnakornsrisuphap, T. S. Rosing, “Lifetime-dependent Battery Usage Optimization for Grid-Connected Residential Systems,” MSCPES, 2015.
- A. S. Akyurek and B. Aksanli and T. S. Rosing, S2Sim: Smart Grid Swarm Simulator, IGSC 2015.
- Baris Aksanli, Alper Sinan Akyurek, Tajana Simunic Rosing, “Minimizing the Effects of Data Centers on Microgrid Instability,” IGSC’15.
- H. Rodrigues, R. Strong, A. Akyurek, T. S. Rosing. “Dynamic Optical Switching for Latency Sensitive Applications,” ACM/IEEE Symposium on Architectures for Networking and Communications Systems, 2015.
- Christine Chan, Alper Sinan Akyurek, Kalyan Vaidyanathan, Kenny Gross, Tajana Rosing, “Optimization of Energy, Cooling and IO Performance for Data-intensive Applications on Enterprise Servers,” TECHCON, 2015.
- Baris Aksanli and Tajana Rosing, “Providing Regulation Services and Managing Data Center Peak Power Budgets,” Design, Automation and Test in Europe (DATE), 2014.
- A. Sadasivarao, H. Rodrigues, S. Syed, C. Liou, S. Balakrishnan, A. Lake, E. Poyoul, C. Guok, I. Monga, T. Rosing., “Enabling Multi-Layer Provisioning and Optimization for Core Transport Networks

with Unified Packet-Optical Control Plan,” 11th USENIX Symposium on Networked Systems Design and Implementation, NSDI’14.

- H. Rodrigues, A. Akyurek, T. Rosing, “OCSEMU: SDN Enabled Fast Hybrid Optical Circuit Switch Emulator Platform to Study Application Performance in the Emerging Optical Data Center,” OIDA Software Defined Photonic and Data Center Networks Workshop, 2014.
- H. Rodrigues, R. Strong, T. Rosing, “Scheduling Optical Tunnels to Distributed Applications,” USENIX Annual Technical Conference, ATC’14.
- B. O. Akyurek and A. S. Akyurek and J. Kleissl and T. S. Rosing, “TESLA: Taylor Expanded Solar Analog Forecasting,” IEEE SmartGridComm 2014.
- Baris Aksanli, Alper Sinan Akyurek, Madhur Behl, Meghan Clark, Alexandre Donze, Prabal Dutta, Patrick Lazik, Mehdi Maasoumy, Rahul Mangharam, Truong X. Nghiem, Vasu Raman, Anthony Rowe, Alberto Sangiovanni-Vincentelli, Sanjit A. Seshia, Tajana Simunic Rosing, and Jagannathan Venkatesh, “Distributed Control of a Swarm of Buildings Connected to a Smart Grid,” 1st ACM International Conference on Embedded Systems For Energy-Efficient Buildings (BuildSys), 2014.
- F. Seracini, X. Zhang, T. S. Rosing, I. Krueger, “A Proactive Customer-Aware Resource Allocation Approach for Data Centers,” ISPA’14.
- A. S. Akyurek, B. Torre, T. S. Rosing, “ECO-DAC Energy Control Over Divide and Control,” IEEE SmartGridComm 2013.
- B. Aksanli, T.S. Rosing, “Optimal Battery Configuration in a Residential Home with Time-of-Use Pricing,” IEEE SmartGridComm 2013.
- C. Chan, B. Pan, K. Gross, K. Vaidyanathan, T. Rosing, “Correcting vibration-induced performance degradation in enterprise servers,” SIGMETRICS Performance Evaluation Review, 2013. Best paper award.

- Baris Aksanli, Eddie Pettis, Tajana Rosing, “Architecting Efficient Peak Power Shaving Using Batteries in Data Centers,” International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), 2013.
- G. Porter, R. Strong, N. Farrington, A. Forencich, P. Sun, T. Rosing, Y. Fainman, G. Papen, A. Vahdat, “Integrating Microsecond Circuit Switching into the Data Center,” SIGCOMM’13.
- P. Mercati, A. Bartolini, F. Paterna, T. Simunic Rosing, L. Benini, “Workload and User Experience-Aware Dynamic Reliability Management in Multicore Processors,” DAC 2013.
- Baris Aksanli, Eddie Pettis, Tajana Rosing, “Distributed Battery Control for Peak Power Shaving in Data Centers,” International Green Computing Conference (IGCC), 2013.
- Baris Aksanli, Jagannathan Venkatesh, Tajana Rosing, and Inder Monga, “A Comprehensive Approach to Reduce the Energy Cost of Network of Datacenters,” ISCC, 2013. Best paper award.
- L. Zhang, G. Dhiman, and T. S. Rosing. “vGreenNet: Managing Server and Networking Resources of Co-located Heterogeneous VMs,” IEEE International Parallel and Distributed Processing Symposium (IPDPS), High Performance Grid and Cloud Computing, 2013.
- Filippo Seracini; Xiang Zhang; Ingolf Krueger; Tajana Rosing; Massimiliano Menarini, “Green Web Services: Improving Energy Efficiency in Data Centers via Workload Predictions,” ICSEWS’13 GREENS.
- G. Dhiman, V. Kontorinis, R. Ayoub, L. Zhang, C. Sadler+, D. Tullsen, T. Simunic Rosing, “Themis: Energy Efficient Management of Workloads in Virtualized Data Centers,” EuroPar-VHPC’12.
- V. Kontorinis, E. Zhang, B. Aksanli, J. Samson, H. Homayoun, E. Pettis, D. Tullsen, T. Simunic Rosing, “Managing Distributed UPS Energy for Effective Power Capping in Data Centers,” ISCA 2012.

- R. Ayoub, R. Nath, T. Simunic Rosing, “JETC: Joint Energy Thermal and Cooling Management for Memory and CPU Subsystems in Servers,” HPCA 2012.
- Baris Aksanli, Tajana S. Rosing , Inder Monga, “Benefits of Green Energy and Proportionality in High Speed Wide Area Networks Connecting Data Centers,” DATE 2012.
- Baris Aksanli, Jagannathan Venkatesh, Liuyi Zhang, Tajana Rosing, “Utilizing Green Energy Prediction to Schedule Mixed Batch And Service Jobs in Data Centers,” HotPower 2011.
- R. Ayoub, U. Ogras, E. Gorbatoov, Y. Jin, T. Kam, P. Diefenbaugh, T. Rosing, “OS-level Power Minimization Under Tight Performance Constraints in General Purpose Systems,” ISLPED 2011.
- Denis Dondi, Piero Zappi, Tajana Šimunić Rosing, “A Scheduling Algorithm for Consistent Monitoring Results with Solar Powered High-Performance Wireless Embedded Systems,” ISLPED 2011.

15. A partial list of my book chapters prior to the '261 patent is below:

- Baris Aksanli, Jagannathan Venkatesh, Inder Monga, and Tajana Rosing, “Reliable Energy Prediction for Improved Utilization and Efficiency in Datacenters and Backbone Networks,” Computational Sustainability Springer Book Chapter, 2015; and
- G. Dhiman, R. Ayoub, T. Simunic Rosing,” Energy and Thermally Aware Scheduling in Datacenters,” in Energy-Efficient Distributed Computing, Edited by Albert Zomaya & Young Choon Lee , Wiley-Interscience 2010.

16. Since coming to UCSD, I have taught courses on a number of subjects including circuit design, embedded systems, the SmartGrid, Internet of Things, emerging computing and software for acceleration. For example, I designed the graduate course on the SmartGrid that attracted students, researchers, and faculty

from multiple departments. The course covered the following topics: Basics of Power Systems, Power System Analysis & Simulation, Integration of Renewable Energy, Energy Storage, Loads & Demand Side Management, Home Area Networks (HAN), Automated Demand Response, Electric Vehicles, Smart Grid Communication, Measurement & Control, Cyber security, Data Privacy, Distributed Sensing & Control, Wide Area Measurements (e.g., Synchrophasors), System Reliability & Ancillary Services, Microgrids, Smart Grids & Internet of Things.

17. I worked on NSF project that developed an instrument called GreenLight from 2008-2012. GreenLight was a datacenter container along with all the infrastructure needed to measure, monitor, and optimize the energy consumption of large-scale applications from many different areas. The work enables interdisciplinary researchers to understand how to make energy efficient decisions for IT computation and storage.

18. I was also a co-PI on ARPA-E project related to SmartGrid, and DARPA/Semiconductor Research Corporation StarNet project TerraSwarm, where I led efforts on SmartGrid and SmartCities. Through these efforts, I worked on energy efficiency at the grid level, and at interaction between the grid and datacenters, and also within datacenters. I collaborated directly with companies such as Microsoft and Google, and also received grants for projects related to energy efficiency of datacenters and their relationship to the grid and energy storage.

19. Based on my experience and education, I believe that I am qualified to opine as to the knowledge and level of skill of a POSITA at the time of the '261 Patent (which I further describe below), in addition to the general state of the art during that time.

III. LEVEL OF ORDINARY SKILL IN THE ART

20. I understand there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including (1) the levels of education and experience of persons working in the field at the time of the invention; (2) the sophistication of the technology; (3) the types of problems encountered in the field; and (4) the prior art solutions to those problems.

21. A POSITA in the field of the '261 patent, as of its earliest alleged effective filing date of October 30, 2018, would have been someone knowledgeable and familiar with green datacenter efforts, renewable power co-location techniques, and geographically distributed datacenter workload distribution techniques available in the late 2010's. Such a POSITA would have a bachelor's degree in computer science, computer engineering, electrical engineering, mechanical engineering, or equivalent training, and approximately two years of experience relating to the field of green datacenters, or more particularly, the use of renewable energy in powering datacenters, workload distribution amongst datacenters, and power management

techniques for datacenters. Additional work experience can substitute for specific educational background, and vice versa.

22. For purposes of this Declaration, in general, and unless otherwise noted, my statements and opinions, such as those regarding my own experience and what a POSITA would have understood or known generally (and specifically related to the references I consulted herein), reflect the knowledge that existed in the relevant field as of the priority date of the '261 patent.

IV. RELEVANT LEGAL STANDARDS

23. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '261 patent, I am relying on certain basic legal principles that Cholla's counsel has explained to me.

24. I understand that prior art to the '261 patent includes patents and printed publications in the relevant art that predate the priority date of the '261 patent. For purposes of this Declaration, I am applying October 30, 2018, as the priority date of the '261 patent.

25. I have been informed by counsel that a claimed invention is unpatentable under 35 U.S.C. § 103 if the differences between the claimed invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a POSITA. I have also been informed by counsel that the obviousness analysis considers factual inquiries, including the level

of ordinary skill in the art, the scope and content of the prior art, and the differences between the prior art and the claimed subject matter.

26. I have been further informed by counsel that there are several recognized rationales for combining references or modifying a reference to show obviousness. These rationales include: (a) combining prior art elements according to known methods to yield predictable results; (b) simple substitution of one known element for another to obtain predictable results; (c) use of a known technique to improve a similar device (method, or product) in the same way; (d) applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; (e) choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and (f) some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art or to combine prior art teachings to arrive at the claimed invention.

V. BACKGROUND TECHNOLOGY

A. Electrical Grid Stability in View of Renewable Generator Supply

27. The effective functioning of the United States electricity grid hinges on the precise management of energy supply and demand. Despite meticulous planning and advanced forecasting, maintaining a perfect balance between energy supply and demand is an ongoing challenge. The demand for electricity fluctuates throughout the day, with peaks during periods of intense weather, economic activity, or seasonal

changes. Ex.1010, [0069]. Further, the supply of electricity is becoming increasingly variable at a grid level as more renewable power sources are integrated into the grid. Ex.1011, pp. 2-3.

28. Renewable power sources – such as wind and solar – can have large swings in renewable power output over short periods of time and longer periods of time, which must be taken into account for stability and economic purposes. Ex.1011, p. 6. When the power supply exceeds the demand during periods of low consumption or abundant renewable energy generation, Independent System Operators (“ISOs”)/Regional Transmission Organizations (“RTOs”) initiate curtailment actions. Ex.1011, p. 15. These proactive measures involve managing the amount of power arriving at the grid from various generation sources and/or the amount of power being taken from the grid to maintain a balanced system. Ex.1011, p. 15. Curtailment typically occurs if there is an overgeneration of electricity from a renewable (e.g., high winds, peak sun, etc.), a lack of transmission and/or distribution line capacity to handle renewable output, or other grid conditions. Ex.1011, p. 15.

29. During curtailment, certain power sources may be temporarily disconnected from the grid or operated at reduced capacities. Ex.1010, [0045], [0055]–[0056]. For instance, wind turbines may be feathered to reduce their rotational speed and power output, or solar power plants may temporarily shut down.

Id. Curtailment can be achieved either involuntarily or voluntarily. Involuntary curtailment is a manual curtailment that is “ordered by the grid operator in response to system-reliability conditions and accomplished through exceptional dispatch.” Ex.1011, p. 15. Such curtailment for system-reliability conditions is also referred to as reliability curtailment. Voluntary curtailment “occurs through the wholesale market,” with renewable power sources “voluntarily reduc[ing] their output in response to market conditions” such as negative wholesale prices. Ex.1011, p. 15. Such curtailment in response to market conditions is also referred to as economic curtailment. *See, e.g.*, Ex.1012, p. 3.

30. By carefully managing curtailment, ISOs/RTOs work to optimize power transmission and avoid grid congestion, preventing excess electricity from going to waste, though curtailment does result in wasted electricity in the form of discarded supply. In addition to managing imbalances between supply and demand, curtailment actions are crucial for grid stability when faced with issues such as grid overloading, voltage fluctuations, or system frequency deviations. By regulating the power flow strategically, ISOs/RTOs ensure that electricity is efficiently transmitted across long distances without compromising the integrity of the grid.

B. Using Green Datacenters to Harness Curtailed Power From Renewable Power Sources

31. Unfortunately, using curtailment to regulate power flow – even strategically – results in wasted electricity from renewable power sources. That is,

power supply from a renewable power source that is not put on the grid due to curtailment (reliability-wise or economic-wise) is wasted if it is not put to any use. Many have recognized this problem, and have proposed co-locating different types of loads to harness the power that would otherwise go to waste.

32. One type of load that has been co-located with renewable power sources is the datacenter with its compute load. For example, Bird et al. describes using “renewable, constrained, or stranded energy sources combined with a network of distributed, small, cheap, and highly efficient mini data centers that are co-located at the generation source.” Ex.1013, p. 1. Co-locating at the source reduces or eliminates the costly transmission infrastructure and associated losses. Ex.1013, p. 1. Bird’s co-location was achieved by installing “rack servers into robust, modular, transportable cabinets or containers that can be placed next to the distributed electricity generation site.” Ex.1013, p. 5.

33. Containers were a well-known solution for providing flexible co-location ability to power sources. *See, e.g.*, Ex.1006, Abstract, [0024], [0048]-[0050], [0066] (flexible, modular datacenter consisting of one or more datacenter containers); Ex.1012, pp. 1, 4 (locating datacenters “on generation sites and placed in containers,” “co-located with renewable generators,” for example “on a wind farm, where computing resources are directly connected to the power backbone of wind turbines”); Ex.1014, p. 7 (“using containers and colocating at generation sites

and avoiding power distribution costs”); Ex.1013, p. 5 (modular containers); Ex.1015, pp. 3, 4 (“us[ing] containers and direct siting at renewable generation sites” to harness curtailed power); Ex.1016, p. 2 (using “portable containerized modules”).

34. Some datacenter co-location approaches rely upon switching the power source for the datacenter between the co-located renewable source and a “stable” grid power source. For example, Li (Ex.1017) described directly coupling a server load to a solar array power source. Ex.1017, pp. 1-2 (“direct-coupled PV system design ... the solar panel load is DC load and no power inverter will be used”). In Li’s arrangement, “when the solar power supply drops below a certain threshold, a secondary power supply (e.g., [grid] utilities) will be switched in (via power switch) and used as a power supply until sufficient solar power is available.” Ex.1017, p. 2. Further, “the load starts to operate using solar power when the amount of renewable energy exceeds a power-transfer threshold.” Ex.1017, p. 9.

35. Other datacenter co-location approaches rely solely upon the renewable power source for operation. For example, Chien (Ex.1015) described an approach for “us[ing] containers and direct siting at renewable generation sites” to harness curtailed power. Ex.1015, pp. 3, 4. Chien relied upon “a *pure* renewable-based computing services ... [which] accepts volatility in power as a fundamental operating constraint, and by doing so does not require support from conventional, carbon-generating sources.” Ex.1015, p. 3. Yang 2017 further expanded upon this

description, explaining that the power for the high-performance computing is “*exclusively* with stranded power” with no reliance “on any stable grid power.” Ex.1014, pp. 1, 2 (emphasis in original). Despite the volatility of the source (curtailed power), Yang 2017 demonstrated that “stranded-power based resources can be cost-effective” with benefits in duty cycle and scale for throughput. Ex.1014, pp. 2, 6.

36. To handle the volatility in power, the co-located datacenters would ramp-up or ramp-down power consumption of the servers, based on how much power was available at the time. For example, server compute resources in a co-located datacenter container may “only work when stranded power is available, and otherwise are fully powered off, producing intermittent resources.” Ex.1012, p. 4. Such “power-aware computing” further involved “controlling the amount of power consumed by the computing based on the amount of power available.” Ex.1018, [0024]. Indeed, as Chapel makes clear, it was already known to control power within a datacenter when actual or potential “over capacity conditions” exist, such as arise when the compute load exceeds the amount of power available. Ex.1007, [0020], [0037], [0049], [0050]. A controller was used to “control[] delivery of power to” the computers. Ex.1007, [0019], [0027]-[0028].

C. Load Distribution Among Datacenters


37. Concurrent to these approaches for managing the load of datacenters based on power availability, many have also implemented different strategies to manage the workload of geographically distributed datacenters to achieve better efficiencies in power consumption and/or other desired metrics.

38. It was already well-known prior to the earliest effective filing date of the '261 patent “that different geographical locations may have different and time-varying electricity prices,” such that requests “can be dynamically distributed to different data center sites.” Ex.1016, p. 1. Moreover, at least some of the different datacenter sites “may have on-site renewable energy generators (e.g., solar or wind farms) that can provide low-cost green energy.” Ex.1016, p. 1. Indeed, “different data centers in different geographical locations may have different availabilities of renewable energy depending on the local weather conditions.” Ex.1019, p. 2.

39. Thus, these prior, known approaches would “dynamically dispatch” workload requests “based on green energy availability,” if not also to further reduce operating expenses. Ex.1016, p. 1. For example, systems would “dynamically dispatch[] incoming service requests among different data centers, based on the time-varying electricity prices and availabilities of renewable energy in their geographical locations, to maximize the use of renewable energy.” Ex.1019, p. 3. At least some of these approaches further included the deployment of “portable

containerized modules ... for geo-distributed data centers” in order to better “exploit the diversity and variability in electricity price, green energy availability, and ambient temperature across different geographical locations.” Ex.1016, p. 2.

40. Prior approaches relied upon queues to managing incoming workload requests for distribution to the different datacenters. For example, one approach relied upon workload requests being “submitted to [a] shared queue” to “share a single filesystem, workload, and scheduler” so that workload was dispatched to datacenter sites with green energy when available. *See, e.g.,* Ex.1014, p. 5. This included, for example, queues at each datacenter location to handle scheduling of workload to green or brown energy, as demonstrated in Kiani. *See* Ex.1005, pp. 2-3; Figure 2:



Ex.1005, Fig. 2.

VI. OVERVIEW OF THE '261 PATENT

41. The '261 patent focuses on “a system for controlling the use of ‘behind-the-meter’ power.” Ex.1001, 1:15-16. Different computation operations assigned by enterprises “may vary in their requirements,” with some requiring “low-latency processing,” others being “extremely time sensitive,” or high reliability. Ex.1001, 5:66-6:9. Other computation operations are not time sensitive and “can be batch processed” or “distributed across multiple computational systems.” Ex.1001, 6:9-14. The '261 patent distributes computational operations “among a critical datacenter and behind-the-meter flexible datacenters.” Ex.1001, 23:14-17; *see, e.g.*, Figure 10:





FIG. 10

Ex.1001, Fig. 10 (annotated).

42. A flexible datacenter includes a mobile container, e.g. “a storage trailer disposed on wheels.” Ex.1001, 11:4-6. Other examples include a “storage container ... configured for placement on the ground” or stackable, an inflatable container, a floating container, or not even a mobile container at all (such as a datacenter “situated within a building”). Ex.1001, 11:6-17. The flexible datacenter includes a behind-the-meter power input system, a power distribution system, climate control, a datacenter control system, and multiple servers. Ex.1001, 10:56-11:3; Figure 2:



43. The '261 patent relies upon multiple datacenters, “including critical datacenters and flexible datacenters,” using a queue system to manage distribution of computational operations amongst the datacenters. Ex.1001, 6:62-66. A “queue system is a data model that can organize computational operations for subsequent access and distribution to available datacenters.” Ex.1001, 6:63-65. Any of a “remote master control system,” “local station control system,” or “datacenter control system ... may support operations of [the] queue system.” Ex.1001, 7:1-4. The queue system may be “a centralized queue system maintained centrally by a control system,” or “a decentralized queue system” that is “distributed across multiple control systems.” Ex.1001, 7:21-24, 7:40-42, 26:51-27:1. Whether centralized or distributed, datacenters can either receive distribution of computational operations from the queue system or “access and obtain computational operations” from the queue system. Ex.1001, 7:30-32 (control system distributes), 7:36-39 (datacenters access the operations), 7:55-64 (distribute or access for decentralized embodiments).

44. The organization and distribution of the computational operations may involve various criteria, including datacenter availability, datacenter type, and operational cost (including, for example, availability of cheaper behind-the-meter power). Ex.1001, 7:30-36 (criteria for distribution decision); 23:53-65 (behind-the-meter power availability). The computational operations are then assigned to the

determined datacenter for implementation. Ex.1001, 25:15-19 (datacenter control system determines when to obtain operation), 25:38-42 (remote master control system assigns operations to datacenter).

45. As I will explain below, these concepts were well known as of the priority date of the '261 patent.

VII. CLAIM CONSTRUCTION

46. It is my understanding that in order to properly evaluate the '261 patent, the terms of the claims must first be interpreted. It is my understanding that for the purposes of this *inter partes* review, the claims are to be construed under the so-called *Phillips* standard, under which claim terms are given their ordinary and customary meaning as would have been understood by a POSITA in light of the specification and prosecution history, unless the inventor has set forth a special meaning for a term. I have also been informed that claim terms only need to be construed to the extent necessary to resolve the obviousness inquiry. I have reviewed the entirety of the '261 patent, as well as its prosecution history. It is my opinion that for purposes of applying the prior art presented herein to evaluate the patentability of the Challenged Claims, no claim terms require express construction.

VIII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

47. The discussion in this Declaration provides a detailed analysis of how the asserted prior art references teach each limitation of the Challenged Claims.

48. As part of my analysis, I have considered, and discuss in detail, the scope and content of the prior art and any differences between the alleged invention and the prior art.

49. It is my opinion that the alleged invention recited in the Challenged Claims would have been obvious in view of the teachings of the asserted prior art and the knowledge of a POSITA.


**IX. CHALLENGE 1: THE COMBINATION OF KIANI AND PELIO
RENDERS OBVIOUS CLAIMS 1-2 AND 9-12**

A. Summary of Kiani

50. Kiani is an IEEE publication titled “On the Fundamental Energy Trade-offs of Geographical Load Balancing.” Kiani was published as part of the May 2017 issue of the IEEE Communications Magazine, volume 55, no. 5. Ex.1005, p. 1. The issue published May 12, 2017. Ex.1021 (IEEE Xplore listing). Kiani was added to the IEEE Xplore digital, online library on that same date, May 12, 2017. In my experience, IEEE magazine publications, including the IEEE Communications Magazine, are distributed without restriction and are publicly available as of the date of publication. Further, I have extensive experience using the IEEE Xplore online library, and in my experience publications added to the Xplore library are downloadable (upon purchase, or with the appropriate subscription) by any interested party as of the date the publication is added to IEEE Xplore.

51. I have reviewed Kiani, and it bears the typical indicia of an IEEE publication, including the title of the magazine at the bottom right, the copyright date at the bottom (2017), and document identifiers and cost for purchase at the bottom left. Ex.1005, p. 1. POSITAs were well-versed in accessing documents of interest from IEEE Xplore (if not received at a subscriber to the magazine itself), and therefore it is my opinion that a POSITA would have been able to locate Kiani with reasonable diligence. Kiani was not before the Examiner during prosecution of the '261 patent.

52. Kiani describes how to “optimize the utilization of green energy and the cost of electricity” in geographically dispersed datacenter systems. Ex.1005, p. 1. In addition to improving the energy efficiency of the servers operating in datacenters, the datacenters themselves integrate “renewable energy resources such as solar panels and wind farms into data centers, thereby promoting sustainability and green energy.” Ex.1005, pp. 1-2. One way in which datacenters integrate green energy is “called behind the meter,” by installing datacenters next to renewable power generators. Ex.1005, pp. 2-3; *see also* Figure 1:




Ex.1005, Fig. 1.

53. When there is “extra green energy generation at a data center,” it “can be injected into the power grid, and the data center can receive compensation for the injected power.” Ex.1005, p. 4. Alternatively, when there is insufficient green energy at a datacenter, data centers can procure “brown energy” from the grid. Ex.1005, p. 3. Or, green energy can be stored at times of excess energy availability, and used later. Ex.1005, p. 3.

54. To support a geographically dispersed datacenter system, Kiani explains that “front-end (FE) servers are co-located with users” nearby to field requests from the users. Ex.1005, p. 2. “Each FE server receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers,” such that each FE server “functions as a workload distribution center

that manages the workload by distributing user requests to the appropriate data centers.” Ex.1005, p. 2; *see also* Figure 2:



Ex.1005, Fig. 2.

55. When an FE server distributes a workload request to a datacenter, “the allocated requests to a data center are first placed in a queue before they can be processed by any available server.” Ex.1005, p. 5. There is a “queue at each data center.” Ex.1005, p. 5. This can be seen by the illustration in Figure 2 of Kiani, above.

56. When an FE server receives a workload request, the FE server uses several different criteria including “based on the availability of green energy” at the different available datacenters. Ex.1005, p. 3. “The selection of the appropriate data centers” by FE servers is “based on different parameters like server or content availability, the network distance between FE [server] and data center, the efficiency

of the data centers, the cost of the electricity, and availability of renewable energy.”
Ex.1005, p. 2.


57. Upon allocation to a datacenter, the quality of service (QoS) at each datacenter is “enforced by maintaining an upper bound” on the datacenter’s queue length. Ex.1005, p. 4. At each datacenter, the workload is organized by using a framework that “utilizes green energy as much as possible.” Ex.1005, p. 4. When brown energy is allowed for use as well (e.g., where green energy was insufficient for the workload and brown energy is allowed), the framework seeks to “minimize the cost of electricity” from the brown energy provider. Ex.1005, p. 4.

B. Summary of Pelio

58. U.S. Patent Pub. No. 2011/0316337 to Pelio et al. (Ex.1006, “Pelio”) was filed June 29, 2010, and published December 29, 2011. Pelio is titled “Power Generation Data Center.” Pelio was not before the Examiner during prosecution of the ’261 patent.

59. Pelio describes collocating datacenter containers “with a primary electric generating station” so that the datacenter containers receive their “operating power from connections upstream of any step-up transformer for transmission lines or grids.” Ex.1006, Abstract. The datacenters are located “at a power generation source ... to be powered from the primary generator at a point before the voltages are stepped up through transformers to the transmission and distribution utility

grids.” Ex.1006, [0002]. Power generation sources include, for example, “solar photo-voltaics, wind, wave, and tidal.” Ex.1006, [0033]. Figure 1A provides an illustrative example with the datacenter 160 “collocated with and powered by a wind generator farm 162” (Ex.1006, [0044]):



Ex.1006, Fig. 1A (annotated).

60. The modular datacenter 160 includes one or more datacenter containers that come in “[v]arious sizes ... e.g., ISO-Standard 20, 40, and 53-foot length types.” Ex.1006, [0050]. For example, a “standard-size shipping container[.]” datacenter can

be “home to 1,000 or more servers each.” Ex.1006, [0071]. Within the modular datacenter, each container receives power from a “modular power supply system[]” to operate with, including each container obtaining “its power directly from the utility” located behind-the-meter. Ex.1006, [0051]. Each datacenter container distributes the power thus input via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1006, [0079].

61. As is understood for several types of renewable power sources, “[s]ome energy generating stations naturally cycle between full operation and zero, for example the daily cycles of a solar energy generating station that depends on sunshine,” as well as others that cycle with other renewable sources such as wind or demands from the grid. Ex.1006, [0037]. Pelio explains that during these times of low to zero electrical output from the collocated generating station, the datacenter can either receive power via a backfeed from the grid, or drop any pending jobs in favor of another datacenter when the servers no longer have power. Ex.1006, [0106].

C. Reasons to Combine Kiani with Pelio

62. It would have been obvious to a POSITA to combine the teachings of Pelio with the teachings of Kiani for several reasons.

63. First, both Kiani and Pelio are in the same field of endeavor and are pertinent to the problems being solved by the '261 patent. The '261 patent indicates that the field of the invention “relates to a system for controlling the use of ‘behind-

the-meter’ power.” Ex.1001, 1:15-16. And the Challenged Claims fall within this field, focusing on flexible datacenters that receive power behind-the-meter before stepping up for transmission on a grid, a queue system with subqueues to handle the distribution of computational operations, and a control system to determine how to assign a computational operation based on the amount of BTM power available. *See* Ex.1001, 36:5-36, 37:19-38:17.

64. Kiani is likewise in the field of systems for controlling the use of BTM power. Kiani explains that “[g]eographical load balancing can optimize the utilization of green energy and the cost of electricity by taking advantage of greenness and price diversities at geographically dispersed data centers.” Ex.1005, p. 1. And Kiani describes one approach for greening a datacenter (harnessing the green energy availability) is “behind the meter,” where renewable power generators are installed at the data center location before the energy is placed onto the grid. Ex.1005, pp. 2-3. Kiani further includes queues to manage the distribution of computational operations amongst the datacenters, based on how much green energy is available at each datacenter. Ex.1005, pp. 2, 3. Thus, Kiani is in the same field of endeavor.

65. Pelio is also in the same field of systems for controlling the use of BTM power. Pelio “relates to data centers located at a power generation source and connected to be powered from the primary generator at a point before the voltages

are stepped up through transformers to the transmission and distribution utility grids.” Ex.1006, [0002]. Pelio’s generator examples include various renewable types that have variable output (due to the different cycles of wind and sunshine, for example). Ex.1005, [0037]. Pelio recognized that when the behind-the-meter output was insufficient, one alternative was for the co-located datacenter to process pending compute jobs only when power is available. Ex.1005, [0106]. Thus, Pelio is also in the same field of endeavor.

66. Kiani and Pelio are also pertinent to problems identified by the ’261 patent. According to the ’261 patent, its approach “provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.” Ex.1001, 35:11-15.

67. Kiani’s teachings are pertinent to these problems because they provide a solution that can “optimize the utilization of green energy and the cost of electricity by taking advantage of greenness and price diversities at geographically dispersed data centers.” Ex.1005, p. 1. Kiani further teaches harnessing otherwise wasted green energy behind the meter with opportunistic routing of computing operations to green datacenters when energy is available. Ex.1005, pp. 1-3. Kiani’s geographic load balancing approach takes into account not only price diversity or locational diversity, but also by green and brown workload and energy costs that also takes into account


the network delay associated with green computing. Ex.1005, p. 6. Thus, a POSITA would have found Kiani's teachings pertinent to the problems identified in the '261 patent.

68. Pelio's teachings are also pertinent to the problems identified by the '261 patent. Pelio's teachings provide multiple solutions to these problems, by providing for datacenter containers that are placed behind-the-meter with an energy source, in order to avoid transmission and distribution costs as well as the power loss from transmission. Ex.1006, Abstract, [0002], [0019]. Pelio also teaches using renewable power only when it is available for compute purposes at the datacenter. Ex.1006, [0106]. Thus, a POSITA would have found Pelio's teachings pertinent to the problems identified in the '261 patent.

69. Second, a POSITA considering Kiani's teachings for geographic load balancing would have been motivated to turn to Pelio for further details about how to achieve behind-the-meter operation. Kiani teaches greening a datacenter by harnessing renewable power behind the meter where available, and injecting excess power onto the grid when all of the behind the meter power is not necessary for the datacenter operations. Ex.1005, pp. 2-3, 4. Kiani's focus is on optimizing the utilization of green energy at general datacenters that are geographically distributed, and not necessarily on the details of how to implement those datacenters as behind the meter. *See, e.g.,* Ex.1005, p. 1 (Abstract). A POSITA would have been motivated

to turn to Pelio for behind the meter datacenter implementation details, including different flexible structural implementations and power delivery teachings. *See, e.g.*, Ex.1006, [0095] (flexible datacenter), [0099] (power distribution), [0079] (power feeds).


70. Kiani discloses geographically dispersed data centers, including “Internet-scale” powerful data centers. Ex.1005, p. 2. These are “generally dispersed in different geographical regions,” which achieves “redundancy, scalability, and high availability.” Ex.1005, p. 2; *see also* Figure 1:



Ex.1005, Fig. 1 (annotated).

71. Kiani specifically teaches integrating green power sources with datacenters by placing the green power sources “behind the meter” with the datacenters. Ex.1005, pp. 2-3. Pelio, in turn, teaches implementation details for behind the meter datacenter placement. For example, Pelio discloses a “flexible, just-

in-time, modular data center” that includes a building shell for housing multiple containers, each of which can be modularly added/removed, supported, cooled, and powered. *See, e.g.*, Ex.1006, [0048]-[0052]; *see also* Figure 2:



Ex.1006, Fig. 2.

72. Figure 2 illustrates two exemplary containers that each respectively serve as datacenters for one or more tenants. Ex.1006, [0049]. For example, each container can be “factory-built offsite, e.g., by IBM, HP, SGI/Rackable, etc.” Ex.1006, [0066]. “Such containers provide tremendous equipment configuration flexibility in the overall design.” Ex.1006, [0070]. For example, a “standard-size shipping container[]” data center can be “home to 1,000 or more servers each.” Ex.1006, [0071]. With multiple such containers within the building shell, each datacenter building easily provides the type of power and scale that Kiani envisioned at geographically dispersed locations, with the added benefit of “equipment configuration flexibility” provided by Pelio’s teachings.

73. Further, Pelio provides power provisioning implementation details that a POSITA would have been motivated to turn to from Kiani's general datacenter teachings. For example, while Kiani teaches providing power behind the meter to datacenters to "green" the datacenter operations, Kiani does not provide details for implementing the energy obtained by collocated, behind the meter datacenter and renewable power generation sources. *See* Ex.1005, pp. 2-3. Pelio provides power to the datacenter by a "number of modular power supply systems 240-[2]43" that are "behind the common building shell" to provide operating power for particular containerized datacenters. Ex.1006, [0051]. Pelio also teaches how to distribute the power to the different containerized datacenters within the building shell via raceways, and via A and B power feeds within each containerized datacenter. Ex.1006, [0099], [0079].

74. Third, a POSITA would have had a reasonable expectation of success in combining the teachings of Pelio with the teachings of Kiani. As an initial matter, Pelio's teachings are compatible with Kiani's teachings. Both Pelio and Kiani teach placing datacenters behind the meter to take advantage of green energy before the energy is placed on the grid for transmission and distribution. *See, e.g.,* Ex.1005, pp. 2-3, 4 (behind the meter approach places "renewable power generators at the data center location" by the datacenter operator or a third party, with excess power injected into the grid); Ex.1006, [0002], [0047] (locating datacenters at power

generation source at a point before voltages are stepped up through transformers to transmission and distribution grids).

75. Moreover, both Kiani and Pelio contemplate using a network architecture to distribute tasks/jobs between different datacenter locations. Kiani describes distributing end-user requests to different geographically located datacenters via networking (Ex.1005, p. 2), and Pelio teaches another datacenter assuming the jobs of a datacenter that has insufficient power to complete the pending jobs (Ex.1006, [0106]). Pelio's datacenter implementation teachings provide, in short, datacenter implementation details that support Kiani's stated purposes, including the placement of datacenters behind the meter and supporting workload via multiple datacenters. Yet further, both Kiani and Pelio rely upon networking infrastructure to manage workload among multiple datacenters. These factors further lead a POSITA to have a reasonable expectation of success in combining Pelio's datacenter implementation teachings with Kiani's geographically distributed datacenter load balancing teachings.

76. Fourth, combining Pelio's teachings with Kiani's teachings would have been predictable and within a POSITA's skillset. It was already well-known to place datacenters, both flexible (as taught by Pelio) and otherwise, behind-the-meter to take advantage of unutilized renewable energy (e.g., due to curtailment etc.). For example, others also taught locating datacenters "on generation sites and placed in

containers,” particularly co-located with “renewable generators.” Ex.1012, pp. 1, 4; see also Ex.1014, p. 7 (“using containers and colocating at generation sites and avoiding power distribution costs”); Ex.1013, p. 5 (modular containers at distributed site locations); Ex.1015, pp. 3, 4 (“us[ing] containers and direct siting at renewable generation sites” to harness curtailed power); Ex.1016, p. 1 (using “portable containerized modules ... for geo-distributed data centers”). Thus, a POSITA would have expected Pelio’s flexible datacenter teachings to predictably work with Kiani’s geographically distributed datacenter system teachings, just as has occurred in other known implementations such as taught by Bird and Brocanelli (to name just a few examples). Ex.1013, p. 5; Ex.1016, p. 1.

77. Specifically, a POSITA would have expected the combination of Kiani’s and Pelio’s teachings to predictably result in Kiani’s geographic load distribution to be to flexible datacenters with appropriate power distribution, as taught by Pelio. *See* Ex.1005, pp. 2-5 (managing distribution of workload to different geographically dispersed datacenters that take advantage of green energy, including behind the meter); Ex.1006, [0002], [0047] (locating flexible datacenters behind-the-meter), [0051] (modular power supplies for flexible datacenters), [0079] (A and B side power), [0099] (power distribution within datacenter building shell). Implementing Pelio’s flexible datacenter teachings would provide the benefit of avoiding power losses from stepping up/down and losses from transmission and

transformer inefficiencies via behind-the-meter implementation. Ex.1005, pp. 2-3; Ex.1006, [0019], [0039]. Further, implementing Pelio's teachings would provide Kiani's system with an enhanced ability to "evolve with developing technology" with implementation much cheaper at a container level. Ex.1006, [0060].

78. The combination of Pelio with Kiani would have therefore been nothing more than the combination of known elements (Kiani's geographically distributed, behind the meter datacenters and corresponding workload distribution, and Pelio's behind the meter, flexible datacenter implementation details), according to known methods (using Pelio's flexible datacenter teachings, and known power distribution for the datacenters of Kiani), to yield predictable results (Kiani's datacenters having added flexibility while still being located behind the meter, with power distributed to the individual servers as taught by Pelio).

79. The combination permits but does not require physical incorporation of elements from Pelio into Kiani. Nevertheless, the physical incorporation of Pelio's teachings, including behind the meter datacenter implementation teachings, into Kiani's teachings would teach the claim limitations.


D. Detailed Analysis of Claims

1. Claim 1

a. [1.0] *A system comprising:*

80. To the extent that the preamble is limiting, Kiani discloses it.

81. Kiani describes “geographical load balancing” as a tool to “optimize the utilization of green energy and the cost of electricity.” Ex.1005, p. 1. To utilize green energy, Kiani teaches “green data centers” which “integrate[] renewable energy sources” into datacenter operation. Ex.1005, pp.1-2. To accomplish geographical load balancing, Kiani describes deploying datacenters in various locations, typically “far away from a large majority of users,” with “front-end (FE) servers [] co-located with users.” Ex.1005, p. 2. “Each FE server receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers.” Ex.1005, p. 2; *see also* Figure 1:



Ex.1005, Fig. 1 (annotated).

82. An “FE server can manage the distribution of its incoming workload to different data centers based on the availability of green energy.” Ex.1005, p. 3. Datacenters can additionally use “brown energy” to power additional servers to

maintain up time. Generally, “each data center utilizes green energy as much as possible, and purchases brown energy only when the green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4.

83. The geographically distributed datacenters and FE servers together disclose a “*system*.”

- b.** **[1.1] *a flexible datacenter that is positioned to receive electrical power directly from a power generation system without the received power flowing through a grid comprising:***


84. Kiani, alone and in combination with Pelio, renders obvious a flexible datacenter that is positioned to receive electrical power directly from a power generation system without the received power flowing through a grid.

85. First, Kiani teaches using “green data centers” to service client computational requests. Ex.1005, pp. 1-2. Green datacenters receive their operating power from renewable energy sources. Ex.1005, pp. 1-2. A well-known way to “green a data center” was to “install renewable power generators at the data center location,” either owned by the datacenter operator or “a third party” that sells the generated power to the datacenter. Ex.1005, pp. 2-3. This co-location is expressly identified as “behind the meter.” Ex.1005, p. 2.

86. Kiani provides a flexible approach where, ideally, “each data center utilizes green energy as much as possible, and purchases brown energy only when the green energy generation is not adequate to serve all incoming workloads.”


Ex.1005, p. 4. A POSITA would have recognized this as teaching that there are times where the green energy is sufficient for all incoming workloads. Indeed, it was well known for flexible, “green” (i.e., powered by renewable energy) datacenters to operate only using green energy. *See, e.g.*, Ex.1012, p. 4 (operate only when stranded power available); Ex.1014, pp. 3-4 (accepting the volatility of stranded power). Nonetheless, Kiani’s setup that can obtain additional “brown energy” from the grid when the green energy is insufficient is just like the setup examples described in the ’261 patent specification. For example, the ’261 patent explains that its flexible datacenter “may additionally be connected to other power sources, such as other behind-the-meter power sources, the power grid, and/or an energy storage system.” Ex.1001, 24:17-20.

87. As illustrated in Figure 2, Kiani demonstrates that such co-location results in the datacenters receiving the generated electrical power from the generation system before that energy flows through a grid:



Ex.1005, Fig. 2 (annotated).

88. A POSITA would have understood that the overlapping circles icon represents a step-up and/or step-down transformer to get energy to the proper voltage. *See, e.g.*, Ex.1006, [0044] (transformer 164 stepping up voltage when output to grid, stepping down voltage when pulling from the grid); FIG. 1A exemplary:




Ex.1006, Fig. 1A (annotated).

89. As Kiani further teaches, the green energy generated by the renewable power generator can be “injected into the power grid” when there is more energy generated than the local datacenter can consume. Ex.1005, p. 4.

90. Second, a POSITA would have recognized that Kiani’s datacenters at each location are within a building structure. *See* Ex.1005, p. 2 (datacenters are dispersed in different geographic regions). This is well-known. Moreover, each

datacenter being within a building discloses the “*flexible datacenter*,” as that term is broadly understood per the ’261 patent specification. While a flexible datacenter includes a mobile container, e.g. “a storage trailer disposed on wheels” (Ex.1001, 11:4-6), the ’261 patent broadly states that the concept of a flexible datacenter also encompasses a datacenter “situated within a building.” Ex.1001, 11:6-17. Thus, Kiani’s disclosure alone demonstrates the “*flexible datacenter*.”


91. Third, it would have been obvious to a POSITA for Kiani’s different datacenters to be implemented as containers to provide increased flexibility in operations, as taught by Pelio. Pelio discloses “data center containers” which are individually examples of a “*flexible datacenter*.” Ex.1006, Abstract. Pelio’s Figure 2 (reproduced below) illustrates multiple examples of a flexible datacenter with each of the “containers 224A and 225B ... stacked together ... and a third data center container 225.” Ex.1006, [0049]; *see also* FIG. 2:



Ex.1006, Fig. 2 (annotated).

92. These data center containers come in “[v]arious sizes ... e.g., ISO-Standard 20, 40, and 53-foot length types.” Ex.1006, [0050]. Examples include data center containers that are “factory-built offsite, e.g., by IBM, HP, SGI/Rackable, etc.” Ex.1006, [0066]. “Such containers provide tremendous equipment configuration flexibility in the overall design.” Ex.1006, [0070]. For example, a “standard-size shipping container” data center can be “home to 1,000 or more servers each.” Ex.1006, [0071].

93. “In general, the use of containers in **flexible, just-in-time, modular data centers** of the present invention allows tremendous flexibility in the configuration and design of the other equipment needed to support such containers.” Ex.1006, [0095]. Figure 4 illustrates “[a] number of data center containers 410-412 on a floor 414.” Ex.1006, [0097] These are flexible datacenters as well:



Ex.1006, Fig. 4 (annotated).

94. Such flexible data center containers, which “can be moved between trucks,” including standard-sized shipping containers, are the very type of portable

flexible datacenters described by the '261 patent. *See* Ex.1006, [0071], [0089]. The '261 patent states that “mobile container 205 may be a storage trailer disposed on wheels and configured for rapid deployment,” including “a storage container ... configured for placement on the ground and potentially stacked in a vertical manner.” Ex.1001, 11:4-9.

95. In combination, it would have been obvious for a POSITA to implement Pelio’s flexible datacenter teachings with Kiani’s system teachings in order to provide more flexibility in datacenter provisioning. This would provide, as Pelio recognizes, “tremendous flexibility in the configuration and design of the other equipment needed to support such containers.” Ex.1006, [0095].

96. Fourth, a POSITA would have understood from the teachings in Kiani that the green datacenters are, by being located at the same location as the energy source, “*positioned to receive electrical power directly from a power generation system.*” *See* Ex.1005, pp. 2-4 (extra green energy generation can be injected into the power grid, demonstrating that the datacenter receives it first (“extra green energy”)). Alone and further in combination with Pelio, this results in behind the meter green datacenters that are also flexible by, e.g., being housed in portable containers.

97. Moreover, a POSITA would have understood that “behind the meter” power as taught by Kiani would be power that did not “*flow[] through a grid*” before

reaching the datacenter. Ex.1005, pp. 2-3. A POSITA would have understood that the claim itself explains what it means for the datacenter to receive electrical power “*directly*” from the power generation system from what is then recited in the claim - receiving the power “*without the received power flowing through a grid.*”

98. Pelio confirms that a POSITA would have understood Kiani’s behind the meter teaching as disclosing the claimed matter. Pelio discloses that the flexible datacenter receives electrical power directly from a power generation system. Pelio teaches the flexible datacenter receiving the electrical power “*directly*” from the generator because Pelio teaches receiving the power “*without the received power flowing through a grid.*”

99. For example, Pelio discloses that the power received by the containers within the modular data centers from one or more power generation stations prior to stepping up the power to a grid – i.e., behind-the-meter, as Kiani teaches. Pelio explains that “[t]he present invention relates to **data centers located at a power generation source** and connected to be powered from the primary generator **at a point before the voltages are stepped up through transformers to the transmission and distribution utility grids.**” Ex.1006, [0002]. “The important common thread between FIGS. 1A-1C is that **the modular data centers 160, 170, and 180 take their operating power directly from the electrical generating units before any corresponding step-up transformer.**” Ex.1006, [0047]. In other words,

the power received by the datacenter has not passed through a grid before it reaches the datacenter – the power expressly avoids the grid.


100. As a result, power is provided directly to the collocated datacenter “such that the power **does not have to be stepped-up to transmission level** and then stepped-down to distribution and end-users.” Ex.1006, [0019]. Thus, Pelio’s collocated datacenters, examples of Kiani’s collocated green datacenters, are able to avoid power losses from “[e]ach stepping up and stepping down of voltages with transformers, and hauling electrical power long distances over transmission lines involve[ing] power losses due to heating.” Ex.1006, [0039].

101. Thus, Kiani alone and combined with Pelio renders obvious “*a flexible datacenter that is positioned to receive electrical power directly from a power generation system without the received power flowing through a grid.*”

- c. **[1.2] *a behind-the-meter (“BTM”) power input system configured to receive power from the power generation system prior to the power undergoing step-up transformation for transmission to the grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;***

102. Kiani, alone and combined with Pelio, renders obvious a behind-the-meter (“BTM”) power input system configured to receive power from the power generation system prior to the power undergoing step-up transformation for transmission to the grid and provide the received power to the flexible datacenter instead of sending the received power to the grid.

103. First, as demonstrated at [1.1], Kiani teaches locating datacenters behind the meter with a power generating station. Ex.1005, pp. 2-3. It would have been obvious to a POSITA that a datacenter located behind the meter, as taught in Kiani, would have a corresponding “BTM” power input system, as the power input system necessarily receives the power BTM when so implemented. As Kiani illustrates in Figure 2, for example, exemplary datacenters receive the power from the renewable source before step-up by a transformer, instead of the received power (received at the datacenter) being sent to the grid:




Ex.1005, Fig. 2 (annotated).

104. Thus, Kiani alone renders obvious this limitation.

105. Second, that a collocated datacenter would have a BTM power input system is further demonstrated in combination with Pelio. Pelio discloses the datacenter receiving power from a co-located power generation station “before the

voltages are stepped up through transformers to the transmission and distribution utility grids.” Ex.1006, [0002]. This means that the datacenter receives the “operating power directly from the electrical generating units before any corresponding step-up transformer.” Ex.1006, [0047]. This discloses behind-the-meter power input, as it is exactly what the ’261 patent defines as behind-the-meter: “‘Behind-the-meter’ power includes power that is received from a power generation system (for instance, but not limited to, a wind or solar power generation system) prior to the power undergoing step-up transformation to High Voltage class AC power for transmission to the grid.” Ex.1001, 4:62-67.

106. Third, Pelio expressly discloses a “*power input system*” that receives power “*behind-the-meter*.” In an embodiment, the “modular power supply systems 240-[2]43” are “utility sub-panels and transformers,” and/or generators or UPS systems, and are therefore examples of a “*power input system*,” as is the utility sub-panel 245 (all of which receive the power prior to step-up transformation, and thus respectively qualify as a behind-the-meter power input system):



Ex.1006, Fig. 2 (annotated).

107. In view of the fact that Pelio’s modular datacenter receives the power before step-up (hence, “*behind-the-meter ... prior to the power undergoing step-up transformation for transmission to the grid*”), Pelio’s power input system at the modular datacenter is likewise “*behind-the-meter.*” So, too, is the power input system at each datacenter container “*behind-the-meter*” in Pelio. Indeed, because the power is provided to the modular datacenter in Pelio, the power so-provided is, by definition, provided to the “*flexible datacenter instead of sending the received power to the grid.*” In combination, Kiani’s system teachings with Pelio’s BTM power input teachings result in Kiani’s flexible datacenters receive the BTM power via Pelio’s BTM power input teachings.

108. Thus, Kiani alone, and further in combination with Pelio, renders obvious “*a behind-the-meter (“BTM”) power input system configured to receive power from the power generation system prior to the power undergoing step-up*

transformation for transmission to the grid and provide the received power to the flexible datacenter instead of sending the received power to the grid.”

d. [1.3] a power distribution system; and



109. Kiani, alone and combined with Pelio, renders obvious a power distribution system.

110. First, Kiani explains that servers in datacenters impose both a base load and a proportional load. Ex.1005, p. 1. The base load is the power consumption from servers even when they are idle. Ex.1005, p. 1. The proportional load “is the extra power consumption that is proportional to the CPU utilization of the servers and accordingly to the load.” Ex.1005, p. 1. Servers at each datacenter rely on “energy utilized to power them,” where servers are “powered by the available green energy.” Ex.1005, p. 4. It would have been obvious to a POSITA that the servers, by being powered by the green power source in Kiani, would have received that power via a “*power distribution system.*” This is a simple and fundamentally well-known concept – any way in which the servers receive the power to operate would be recognized as a “*power distribution system.*” See, e.g., Ex.1007, [0027]; [0082]. Thus, Kiani alone renders obvious this limitation.

111. Second, the combination with Pelio further renders this limitation obvious. Pelio discloses a power distribution system for power to reach flexible datacenters within Pelio’s modular datacenter building shell. Pelio explains that

“[a]n overhead raceway 440 provides ... electrical power connections to uninterruptible power supply (UPS) unit 442 and utility power.” Ex.1006, [0099].

These raceways are illustrated in both figures 2 and 4, distributing the received BTM power to the different flexible datacenter containers within the modular datacenter:



112. Pelio further renders obvious a power distribution system in each datacenter container, such as those that would be found in Kiani’s datacenters. Pelio explains the obvious fact that “electrical power [is] consumed by the servers and

data storage units” within a data center container, and of course that each data center container can be “home to 1,000 or more servers each.” Ex.1006, [0071], [0075]. It would have been obvious to a POSITA that the electrical power, which is input via the “modular power supply systems 240-[2]43” and the utility sub-panel 245, would be distributed to each server in a container via a “*power distribution system*.” This is fundamentally well known. *See, e.g.*, Ex.1007, [0027]; [0082].


113. For example, Pelio teaches that the power distribution to servers in a data center container would be via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1006, [0079]. Such power feeds are well-known to be “independent electrical feeds coming into a data center from the utility company.” Ex.1020, p. 1. Each feed includes “[e]lectrical and distribution switchboards,” a “power distribution unit (PDU),” and “Rack-level PDUs. Ex.1020, p. 1; *see also* Ex.1007, [0082] (data center using a “power distribution unit ... associated with multiple power sources 660 such as an A source and a B source”).

114. Thus, Kiani alone, and further combined with Pelio, renders obvious “*a power distribution system*.”

- e. **[1.4] *a plurality of computing systems powered by the BTM power input system via the power distribution system;***

115. Kiani renders obvious, alone and combined with Pelio, a plurality of computing systems powered by the BTM power input system via the power distribution system. This is further rendered obvious in combination with Pelio.

116. First, Kiani teaches that datacenters support “[o]nline services ... on hundreds of thousands of servers spread across server farms provisioned for the peak load.” Ex.1005, p. 1. A “server farm” is an exemplary datacenter. Ex.1023, 3:18-20 (“one or more systems and/or techniques for regulating the amount of power on a power grid via a datacenter(s) (e.g., a server farm) are provided.”). More specifically, with respect to Figure 1, Kiani teaches that computational requests are sent to “the back-end servers at geo-dispersed data centers.” Ex.1005, p. 2. Each datacenter has “back-end servers.” Ex.1005, p. 2. This results in “servers at each data center” that are powered by green energy to perform desired computational requests. Ex.1005, p. 4 (also explaining that “different data centers have different servers with different service capabilities”); Figure 1:



Ex.1005, Fig. 1 (annotated).

117. As further explained at [1.2] and [1.3], the servers at each datacenter are powered by the behind-the-meter power, with it being obvious that each server would receive the power via a power distribution system and BTM power input. Thus, Kiani alone renders obvious *“a plurality of computing systems powered by the BTM power input system via the power distribution system.”*

118. **Second**, this claim limitation is further rendered obvious in combination with Pelio. Pelio discloses that the datacenter containers include inside them “network servers” that are mounted in racks and organized into rows. Ex.1006, [0071]. For example, “standard-size shipping containers” include those that are “home to 1,000 or more servers each.” Ex.1006, [0071]. This includes data center containers that “can be ordered from their manufacturers fully provisioned with the

servers already installed.” Ex.1006, [0048]. These servers, and generally the data centers, are “computer systems.” Ex.1006, [0008].

119. Pelio teaches that the network servers are powered by the power input system via the power distribution system. As noted at [1.2], the “modular power supply systems 240-[2]43” are examples of a “*power input system*,” as is the utility sub-panel 245 (all of which receive the power prior to step-up transformation, and thus qualify as a behind-the-meter power input system). In particular, the “modular power supply systems 240-[2]43” (e.g., utility sub-panels and transformers, and/or generators and/or UPS systems) and utility sub-panel 245 “provide for the operating power predetermined for the particular modular containerized data centers then or soon to be resident.” Ex.1006, [0051]. Pelio teaches that the containers “may alternatively obtain [their] power directly from the utility via a switchgear, transformer and conduit system.” Ex.1006, [0051].


120. The power is provided from the utility/power input system to the containers via “electrical cable raceways.” Ex.1006, [0049]. The raceways “provide[] for ... electrical power connections” between the containers and the utility power via the power input system. Ex.1006, [0099]. Further, the A and/or B side power feeds provide that power from the raceways to the servers themselves. Ex.1006, [0079].

121. Thus, Kiani alone, further combined with Pelio, renders obvious “*a plurality of computing systems powered by the BTM power input system via the power distribution system.*”

f. [1.5] *a queue system configured to organize a plurality of computational operations, wherein the queue system comprises multiple queue subsystems, and*

122. Kiani discloses, and renders obvious, a queue system configured to organize a plurality of computational operations, wherein the queue system comprises multiple queue subsystems.

123. First, Kiani discloses using a queue system to organize the assignment of computational operations across multiple datacenters. Kiani describes “geographical load balancing” as a tool to “optimize the utilization of green energy and the cost of electricity.” Ex.1005, p. 1. To utilize green energy, Kiani teaches “green data centers” which “integrate[] renewable energy sources” into datacenter operation. Ex.1005, pp.1-2. To accomplish geographical load balancing, Kiani describes deploying datacenters in various locations, typically “far away from a large majority of users,” with “front-end (FE) servers [] co-located with users.” Ex.1005, p. 2. “Each FE server receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers.” Ex.1005, p. 2; *see also* Figure 1:




Ex.1005, Fig. 1.

124. Further, “each FE server functions as a workload distribution center that manages the workload by distributing user requests to the appropriate data centers.” Ex.1005, p. 2. The FE servers select appropriate data centers “based on different parameters like server or content availability, the network distance between FE and data center, the efficiency of the data centers, the cost of the electricity, and availability of renewable energy.” Ex.1005, p. 2. It would have been obvious to a POSITA that the FE server would have a queue system to handle receipt of incoming requests and distribution of outbound assignments to respective datacenters, based on Kiani’s teachings. This, alone, renders obvious the claimed “*queue system*.” Yet further, there can be multiple FE servers each with its own queue for distributing workload to datacenters, as can be seen in Figure 1. Ex.1005.

125. Second, the “*queue system*” is further disclosed by Kiani’s teachings of queues at the different datacenters. Kiani teaches that the FE server distributes allocated requests to the assigned datacenter, where the requests “are first **placed in a queue** before they can be processed by any available server” at the datacenter. Ex.1005, p. 5. There is a “queue at each data center” which receives workload requests from an FE server and then distributes the workload. Ex.1005, p. 5.


126. In a particular example, Kiani teaches that an “FE server can manage the distribution of its incoming workload to different data centers based on the availability of green energy.” Ex.1005, p. 3. Datacenters can additionally use “brown energy” to power additional servers to maintain up time. Generally, “each data center utilizes green energy as much as possible, and purchases brown energy only when the green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4. Kiani illustrates the queues associated with each datacenter in Figure 2:



Ex.1005, Fig. 2 (annotated).

127. As can be seen in Figure 2, Kiani illustrates N datacenters, each with a plurality of servers. All of the queues together, one at each datacenter, constitutes another example of the claimed “*queue system*.” The combined teachings of Kiani and Pelio would result in each flexible datacenter having a queue to receive incoming workload requests from an FE server, as taught by Kiani.

128. Third, Kiani further discloses the queue system comprising multiple queue subsystems. Kiani teaches that there is “a queue at each data center” to receive workload allocation from one or more FE servers. Ex.1005, pp. 3, 5. Specifically, from an FE server allocation, the “allocated requests to a data center are first placed in a queue before they can be processed by any available server” at the datacenter. Ex.1005, p. 5; Figure 2:



Ex.1005, Fig. 2 (annotated).

129. This is consistent with the claimed scope of the '261 patent. According to claim 10, “a first queue subsystem of the queue system is located at the first flexible datacenter,” and claim 11 “a second queue subsystem of the queue system is located at the second datacenter.” Thus, according to the claims of the '261 patent itself, aspects of the queue system are distributed amongst the datacenters, just as Kiani teaches.

130. Even if one were to ignore the fact that Kiani’s FE server would have its own queue system, a POSITA would have appreciated that Kiani’s teaching of the FE server operating in cooperation with the subqueues at each datacenter together demonstrates a “queue system” that performs the claimed operations.

131. Thus, Kiani discloses, and renders obvious, “*a queue system configured to organize a plurality of computational operations, wherein the queue system comprises multiple queue subsystems.*”

g. [1.6] *wherein each queue subsystem organizes a set of computational operations according to rules associated with the queue subsystem; and*

132. Kiani discloses wherein each queue subsystem organizes a set of computational operations according to rules associated with the queue subsystem.

133. First, Kiani discloses organizing computational operations at each datacenter in each datacenter’s queue according to rules. As noted at [1.5], each datacenter has a queue subsystem that receives workload allocations from the FE server. See Ex.1005, pp. 3, 5. Kiani further teaches that each datacenter has a “QoS at each data center [that] is enforced **by maintaining an upper bound**” on the queue length of that datacenter. Ex.1005, p. 4. This is a rule associated with the queue subsystem at that datacenter. Further, because it impacts the queue itself, the upper bound is a rule that causes the queue to organize the allocated workload (“*a set of computational operations*”) at the datacenter (in the queue) for execution.

134. Further, at each datacenter, Kiani teaches that “for green workload allocation” the rule is to “maximize the utilization of green energy,” while limiting “the allocated green workload by the available green resources.” Ex.1005, p. 4. Further, “[i]f the green energy is not adequate to satisfy the QoS requirements,

another optimization framework is then designed to allocate the brown workload.” Ex.1005, p. 4. For the brown workload, in turn, the objective is “to minimize the cost of electricity” while enforcing the “QoS requirements at each data center.” Ex.1005, p. 4. Each rule causes the queue at the datacenter to organize the allocated workload, for example for implementation either by green energy or brown energy, subject to the QoS requirements. Thus, Kiani discloses a plurality of rules associated with the queue subsystem at the datacenter to achieve “utiliz[ing] green energy as much as possible, and purchas[ing] brown energy only when the green energy generation is not adequate to serve all incoming workloads” (Ex.1005, p. 4):

- An “upper bound” on the queue length at the datacenter;
- A QoS for the datacenter that determines the “upper bound”;
- Maximize green energy utilization for workload allocation at the datacenter;
- Apply brown energy if green energy unavailable/insufficient;
- Minimize the cost of brown energy for workload at the datacenter; and
- Maintaining the QoS requirement for the datacenter for brown energy operations.

135. Second, Kiani further teaches that each FE server “can manage the distribution of its incoming workload to different data centers based on the

availability of green energy.” Ex.1005, p. 3. Kiani further explains that “[t]he selection of the appropriate data centers” by FE servers is “based on different parameters like server or content availability, the network distance between FE [server] and data center, the efficiency of the data centers, the cost of the electricity, and availability of renewable energy.” Ex.1005, p. 2. Thus, the very allocation of workload to different datacenters is done according to “rules” (according to the “different parameters”), which is another act of “*organiz[ing] a set of computational operations*”). Moreover, the organizing is done “*according to rules associated with the queue subsystem*” since each datacenter can have different limitations “due to different service efficiency parameters.” Ex.1005, p. 5.

136. Thus, by virtue of the FE server sending a workload allocation to a particular queue subsystem at a datacenter due to rules dictating that particular allocation as the best option, the organizing is done “*according to*” rules that are “*associated with*” that queue subsystem (since they cause the workload to be allocated to that queue subsystem’s datacenter). These additional rules, either alone or in addition to the rules for execution at the datacenter itself, therefore disclose this claim limitation.


137. Accordingly, Kiani discloses “*wherein each queue subsystem organizes a set of computational operations according to rules associated with the queue subsystem.*”

h. [1.7] *a routing control system configured to:*

138. Kiani discloses a routing control system.

139. The '261 patent specification does not describe a “*routing control system*” in particular. However, according to dependent claim 2, the “*routing control system*” “*is a remote master control system.*” The remote master control system can be either “*positioned remotely from the flexible datacenter*” (claim 2) or “*positioned at the flexible datacenter*” (claim 3).

140. Kiani discloses a routing control system, like the remote master control system setup in the '261 patent's Figure 5, configured to perform operations as recited in claim 1. For example, Kiani teaches that the FE server manages “the distribution of its incoming workload to different data centers based on the availability of green energy.” Ex.1005, p. 3. The FE server distributes allocated requests to the assigned datacenter. Ex.1005, p. 5. Thus, the FE server is an example of a “*routing control system.*”



Ex.1005, Fig. 2 (annotated).


144. The FE server “receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers,” such that each FE server “functions as a workload distribution center that manages the workload by distributing user requests to the appropriate data centers.” Ex.1005, p. 2. The “incoming workload” to the FE server represents a “*computational operation to be performed,*” because the workload constitutes operation requests from users. Ex.1005, pp. 2-3. By knowing what workload the FE server has to distribute, the FE server has “*determine[d], using the queue system,*” the computational operation to be performed (and, thus, distributed to a datacenter for performance).

145. Thus, Kiani discloses “[a routing control system configured to:] *determine, using the queue system, a computational operation to be performed.*”

j. [1.9] *[a routing control system configured to:] determine an amount of BTM power that is available to the plurality of computing systems;*

146. Kiani discloses, and in combination with Pelio renders obvious, [a routing control system configured to:] determine an amount of BTM power that is available to the plurality of computing systems.

147. First, as demonstrated at [1.7], Kiani discloses a “*routing control system*” with the “FE server” teachings. *See* Ex.1005, pp. 3, 5.



Ex.1005, Fig. 2 (annotated).

148. Second, Kiani discloses that the FE server determines how much “green energy,” which in one implementation is “behind the meter” (thus, “*an amount of BTM power*”) with a datacenter, is available (“*available to the plurality of computing systems*”). For example, a way to “green a data center” is “behind the meter,” and the “FE server can manage the distribution of its incoming workload to

different data centers based on the availability of green energy” generated by the BTM energy source (e.g., wind, solar) taught in Kiani. Ex.1005, pp. 2-3 (“wind turbine,” “solar generation”).

149. Since Kiani teaches that each datacenter has a plurality of servers (“*computing systems*”), of which multiple are powered by “green energy” located BTM, Kiani therefore teaches that the FE server’s determination of the amount of green energy available to a datacenter is really a determination of BTM power available to the plurality of computing systems at the datacenter. Indeed, Kiani teaches that green workload is defined as “the workload served by a green server” in a datacenter, with “some servers [] turned on and powered by the available green energy (green servers).” Ex.1005, p. 4.

150. Kiani illustrates this concept with an example based on simulation data about the trends of wind power at three different datacenters. *See* Ex.1005, p. 4. In one slot of time, the FE server determines that “before hour 8 and from hours 15 to 24, the available wind power is adequate to serve all the service requests” at the three exemplary datacenters, such that “the brown workload is almost not allocated to the data centers.” Ex.1005, p. 4. “In other words, in these hours, the available wind power is the key decision factor to allocate workloads among the data centers” by the FE server. Ex.1005, p. 4.

151. Third, this limitation is further rendered obvious in combination with Pelio's teaching of receiving power BTM by a collocated datacenter. *See* analysis at [1.1] and [1.2] *supra*. In the combination, Kiani's FE server would take into account the BTM power available according to both Kiani's and Pelio's BTM teachings for powering a datacenter.

152. Thus, Kiani alone, and further in combination with Pelio, renders obvious “[a routing control system configured to:] determine an amount of BTM power that is available to the plurality of computing systems.”

k. [1.10] *[a routing control system configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available; and*


153. Kiani combined with Pelio renders obvious [a routing control system configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available.

154. First, the FE server distributes workload to datacenters “based on the availability of green energy” at the different datacenters. Ex.1005, p. 3. And, as noted at [1.1] and [1.2], Kiani discloses providing green energy via “behind the meter” with collocated renewable power generators and datacenters. Ex.1005, pp. 2-3.

155. The “availability of green energy” determination, together with Kiani's “behind the meter” teaching for that green energy, demonstrates the claimed “*determine whether to route the computational operation*” to a datacenter “*using the*

amount of BTM power that is available.” Indeed, in an example based on availability of wind power as the “green energy,” Kiani explains that “the available wind power is **the key decision factor** to allocate workloads among the data centers.” Ex. 1005, p. 4. As will be explained at [1.11] below, Kiani further teaches the act of allocating workload to a datacenter, which demonstrates actually routing a computational operation to that datacenter. *See* Ex.1005, pp. 2, 4.

156. Second, as explained at [1.1], the combined teachings of Kiani and Pelio render obvious a flexible datacenter placed BTM per Kiani’s and Pelio’s teachings. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2:



Ex.1006, Fig. 2 (annotated).

157. Accordingly, Pelio’s flexible datacenter implementation detail would be used in combination with Kiani’s BTM datacenter teachings, as well as Kiani’s teaching of distributing (routing) workload (computational operations) to the BTM


datacenter based on availability of BTM green energy (amount of BTM power available).

158. Thus, Kiani combined with Pelio renders obvious “[a routing control system configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available.”

- 1. [1.11] [a routing control system configured to:] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.**

159. Kiani combined with Pelio renders obvious [a routing control system configured to:] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

160. First, Kiani teaches actually allocating the workload to a determined flexible datacenter (per the flexible datacenter teachings of Pelio). “[E]ach FE server functions as a workload distribution center that manages the workload **by distributing user requests to the appropriate data centers.**” Ex.1005, p. 2.




Ex.1005, Fig. 2 (annotated).

161. The allocation of workload in Kiani discloses the “*assign[ing] the computational operation*” per the claim, since the workload is an example of computational operations (user requests for operations, per Kiani at p. 2) and Kiani’s distribution of workload is an act of assigning. Specifically, each “FE server receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers.” Ex.1005, p. 2. The act of distributing workload demonstrates an assignment of the computational operation for performance at the receiving datacenter.


162. Moreover, Kiani generally teaches assigning workload to datacenters based on availability of green energy with an example based on wind power. Based on the trend of wind power, after hour 14 “more of the green workload **is assigned**

to data center 1 where the highest wind power is available.” Ex.1005, p. 4. This can be seen in the example of wind power trends in Figure 4:



Ex.1005, Fig. 4 (annotated).

163. The allocation of workload based on this sample set of wind power availability from Figure 4 is illustrated in Kiani’s Figure 3. Following the example from Figure 4 at hour 14, Kiani explains that green workload is assigned to datacenter 1 based on the higher availability of wind power at that location at that hour:



Ex.1005, Fig. 3 (annotated).

164. Thus, Kiani’s distribution of requests (workload), also referred to as its allocation of workload, was expressly an act of “assign[ing]” the workload to a given datacenter.

165. Second, as explained at [1.10], Kiani teaches determining to route workload (computational operations) to a datacenter, which in combination with Pelio’s teachings would be a flexible datacenter. *See* Ex.1005, pp. 2-4; Ex.1006, Abstract, [0049], [0071]. Thus, Kiani’s allocation of workload, based on determining the amount of BTM green energy available to a datacenter, demonstrates the actual allocation being based on that determination.

166. Third, as explained at [1.1], the combined teachings of Kiani and Pelio render obvious a flexible datacenter placed BTM per Kiani’s and Pelio’s teachings.

See Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2. Accordingly, Pelio's flexible datacenter implementation detail would be used in combination with Kiani's BTM datacenter teachings, resulting in the one or more of the datacenters taught in Kiani being a "*flexible datacenter*."

167. Thus, Kiani combined with Pelio renders obvious "[a routing control system configured to:] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter."

2. Claim 2

a. [2.0] *The system of claim 1,*


168. See analysis at [1.0]-[1.11].

b. [2.1] *wherein the routing control system is a remote master control system positioned remotely from the flexible datacenter.*

169. Kiani, alone and combined with Pelio, renders obvious wherein the routing control system is a remote master control system positioned remotely from the flexible datacenter.

170. First, Kiani teaches that the FE server (an example of a "*routing control system*") is located remotely from the datacenters. Often, "powerful data centers are generally deployed far away from a large majority of users." Ex.1005, p. 2. In response, "front-end (FE) servers are co-located with users" in order to receive user requests and "distribute[] the requests to back-end servers at geo-dispersed data

centers.” Ex.1005, p. 2. Kiani further illustrates this in Figure 1, with the FE servers remote from the datacenters they route workload requests to:



Ex.1005, Fig. 1 (annotated).

171. Any one of the FE servers demonstrates the control system. Since Kiani teaches routing control by the FE server, remotely from the datacenters, the FE server teaches a “*remote master control system*.”

172. **Second**, as explained at [1.1], the teachings of Kiani alone and combined with Pelio render obvious a flexible datacenter placed BTM. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2. Accordingly, Kiani’s datacenters are flexible by being within a building, just as the example in the ’261 patent, as well as further where Pelio’s flexible datacenter implementation detail is used in combination with Kiani’s BTM datacenter teachings, resulting in the one or

more of the datacenters taught in Kiani being a “*flexible datacenter*,” which is remote from any of the FE servers.

173. Thus, Kiani, alone and combined with Pelio, renders obvious “*wherein the routing control system is a remote master control system positioned remotely from the flexible datacenter.*”

3. Claim 9

a. [9.0] *The system of claim 1,*

174. See analysis at [1.0]-[1.11].

b. [9.1] *wherein the routing control system controls distribution of the plurality of computational operations from the queue system.*

175. Kiani discloses wherein the routing control system controls distribution of the plurality of computational operations from the queue system.

176. As demonstrated at [1.7], Kiani discloses a routing control system with its description of the FE server. Specifically, “each FE server receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers,” such that “each FE server functions as a workload distribution center that **manages the workload by distributing user requests to the appropriate data centers.**” Ex.1005, p. 2. Managing the workload distribution demonstrates that the FE server controls distribution of computational operations from the FE server – a part of the “*queue system.*” See analysis at [1.5].

177. Thus, Kiani discloses “*wherein the routing control system controls distribution of the plurality of computational operations from the queue system.*”

4. Claim 10


a. [10.0] *The system of claim 1,*

178. See analysis at [1.0]-[1.11].

b. [10.1] *wherein a first queue subsystem of the queue system is located at the flexible datacenter.*

179. Kiani combined with Pelio renders obvious wherein a first queue subsystem of the queue system is located at the flexible datacenter.

180. **First**, as demonstrated at [1.5], Kiani discloses a queue at each of the datacenters. When requests are allocated by the FE server, the requests (workload) are “placed in a queue” at the assigned datacenter, with a “queue at each data center” for workload requests. Ex.1005, p. 5. Thus, with multiple datacenters, Kiani discloses a first queue at a first datacenter. *See also* Ex.1005, Figure 2:



Ex.1005, Fig. 2 (annotated).

181. Second, as explained at [1.1], the combined teachings of Kiani and Pelio render obvious a flexible datacenter placed BTM per Kiani’s and Pelio’s teachings. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract, [0049], [0071]; Figure 2. Accordingly, Pelio’s flexible datacenter implementation detail would be used in combination with Kiani’s BTM datacenter queue teachings, resulting in the one or more of the datacenters taught in Kiani being a “*flexible datacenter*” with a queue (a “*first queue subsystem*”).

182. Thus, Kiani combined with Pelio renders obvious “*wherein a first queue subsystem of the queue system is located at the flexible datacenter.*”

5. Claim 11


a. [11.0] *The system of claim 10, further comprising:*

183. *See* analysis at [1.0]-[1.11] and [10.0]-[10.1].

- b. [11.1] a second datacenter, wherein a second queue subsystem of the queue system is located at the second datacenter.**

184. Kiani discloses a second datacenter, wherein a second queue subsystem of the queue system is located at the second datacenter.

185. Kiani discloses multiple datacenters, with each datacenter having its own queue. Specifically, there is a “queue at each data center” for workload requests. Ex.1005, p. 5. Thus, with multiple datacenters, Kiani disclose a second queue at a second datacenter. *See also* Ex.1005, Figure 2:



186. Notably, claim 11 does not recite that the second datacenter must be a “flexible” datacenter. Thus, Kiani alone discloses this limitation with its different datacenter disclosures. However, even if the second datacenter had to be a flexible datacenter as well, this would have been further obvious in combination with Pelio’s

teachings of a flexible datacenter being applied to one or more additional datacenters in Kiani. Indeed, Pelio itself teaches that it was known to have multiple flexible datacenters. *See* Ex.1006, [0049]; *see also* FIG. 2.

187. Thus, Kiani discloses “*a second datacenter, wherein a second queue subsystem of the queue system is located at the second datacenter.*”

6. Claim 12

a. [12.0] *The system of claim 11,*

188. *See* analysis at [1.0]-[1.1], [10.0]-[10.1], and [11.0]-[11.1].

b. [12.1] *wherein the second datacenter is a second flexible datacenter.*

189. Kiani combined with Pelio renders obvious wherein the second datacenter is a second flexible datacenter.

190. As explained at [11.1], Kiani teaches multiple datacenters, with each datacenter having its own queue. Ex.1005, p. 5. As further explained at [11.1], it would have been obvious for Kiani’s additional datacenters to likewise be flexible datacenters. This would have been nothing than the mere duplication of Pelio’s flexible datacenter teaching for additional ones of the datacenters taught by Kiani. Moreover, Pelio itself teaches that it was known to have multiple flexible datacenters. *See* Ex.1006, [0049]. Accordingly, in combination one or more additional datacenters per Kiani would be flexible datacenters, based on Pelio’s combined teachings.

191. Thus, Kiani combined with Pelio renders obvious “*wherein the second datacenter is a second flexible datacenter.*”

X. CHALLENGE 2: THE COMBINATION OF KIANI, PELIO, AND CHAPEL RENDERS OBVIOUS CLAIMS 4-8


A. Summary of Chapel

192. U.S. Patent Pub. No. 2010/0145542 to Chapel et al. (Ex.1007, “Chapel”) was filed September 29, 2009, and published June 10, 2010. Chapel claims priority to PCT/US2008/57150 filed March 14, 2008. Chapel is titled “Smart Electrical Outlets and Associated Networks.” Chapel was not before the Examiner during prosecution of the ’261 patent.

193. Chapel “relates generally to electrical power distribution and management and, in particular, to [a] ... device associated with a local (e.g., single or multiple residential or business premises) circuit, to intelligently monitor at least a portion of the circuit and control delivery of electricity over the circuit.” Ex.1007, [0002]. Intelligent local devices are used to “control power delivered to an electrical device via a circuit,” including electrical devices in a datacenter. Ex.1007, [0047].

194. A system with a number of smart receptacles (into which computers are connected for power, for example) “communicate with a local controller” located at the customer premises, which in turn communicates “with a remote controller” to remotely control power to the computers. Ex.1007, [0056]. Such customer premises include a “data center environment,” which includes “data devices” such as servers

and the like. Ex.1007, [0080]. The servers are plugged into the receptacles, which are in communication with a local controller. Ex.1007, [0082]. The local controller is further in communication with a remote controller. Ex.1007, [0083]; FIG. 6:



Ex.1007, Fig. 6.

195. For example, when an operator of the remote controller desires to “power down or power up a device or set of devices,” the operator has the remote controller send “[a]n appropriate message ... to the local controller,” which in turn conveys the message to the appropriate receptacle. Ex.1007, [0090].

B. Reasons to Combine Kiani and Pelio with Chapel

196. A POSITA would have been motivated to combine Kiani and Pelio with Chapel for several reasons.

197. First, Chapel describes approaches for “electrical power distribution and management” of electrical devices connected to one or more circuits, including

controlling delivery of power to those devices within a customer premises like a datacenter. Ex.1007, [0002], [0047]. Chapel teaches controlling the power delivery for multiple reasons, including, for example, when power supply is insufficient (either actual or projected). Ex.1007, [0048]. Chapel is, therefore, pertinent.

198. For example, Chapel teaches solutions to the need of modulating power delivery in response to power supply changes. In particular, Chapel teaches ways in which to “intelligently monitor at least a portion of the circuit and control delivery of electricity over the circuit,” be that a single circuit or an entire premises such as a datacenter. Ex.1007, [0002], [0019], [0056], [0079]. Chapel’s solution includes control for both powering up and powering down multiple devices at the premises based on predicted or actual power supply status. Ex.1007, [0090]. Thus, a POSITA would have found Chapel’s teachings pertinent to the problems identified in the ’261 patent.

199. Second, Kiani and Pelio both describe datacenters obtaining behind-the-meter power from renewable generating sources, avoiding the costs of transmission and distribution as well as the losses associated therewith. *See* Ex.1005, pp. 2-3; Ex.1006, Abstract; [0002], [0019]. Pelio further teaches various aspects of how a modular datacenter that houses one or more datacenter containers conveys the behind-the-meter power to the containers, including for example raceways to convey power from power supply system(s). Ex.1006, [0051], [0099]. Pelio also teaches that

electrical power is distributed to servers within a datacenter container via “power feeds,” but leaves the details up to POSITAs. Ex.1006, [0079].

200. A POSITA would have therefore been motivated to turn to Chapel, which provides additional details of how to implement various elements within a datacenter including power distribution systems. For example, Chapel teaches that a datacenter includes “a power distribution unit” associated with one or multiple power sources. Ex.1007, [0027]; [0082]. Power distribution units were a well-known aspect of datacenter design, including datacenters as taught in Kiani and Pelio. Indeed, a POSITA would have understood that Kiani’s and Pelio’s behind-the-meter renewable power sources were well-known for powering datacenters well before the earliest effective filing date of the ’261 patent. *See, e.g.*, Ex.1006, Abstract; Ex.1012, pp. 1, 4; Ex.1014, p. 7; Ex.1013, p. 6; Ex.1015, pp. 3, 4.

201. Third, Kiani and Pelio both recognize the variability of renewable power sources. Kiani describes this as the “availability of renewable energy,” including “the weather dependency of green energy” such that “green energy is not a constantly available resource.” Ex.1005, pp. 2, 3. Thus, Kiani explains that 24/7 datacenters “have to be connected to on-grid brown energy” to procure when green energy is insufficient. Ex.1005, p. 3. Pelio describes this variability as the cycling “between full operation” and as far as “zero” output due to lack of the renewable resource (e.g., wind, solar, waves, etc.). Ex.1006, [0037]. As part of this recognition,

Pelio explains that there are times where a datacenter located behind-the-meter would receive insufficient electricity from the generating source to adequately power the datacenter, in which case the servers would stop operating until sufficient power is available again. Ex.1006, [0106].

202. Kiani's and Pelio's teachings contemplate placing datacenter containers behind-the-meter to harness the power output before transmission costs and losses arise, as well as providing equipment configuration flexibility, leaving up to a POSITA the specifics of how to handle the variability of the power source itself. *See* Ex.1005, pp. 2-3; Ex.1006, [0002]; [0070]. Thus, a POSITA would have been further motivated to turn to Chapel for Chapel's power management teachings.

203. For example, Chapel teaches ways in which to manage electrical power distribution within a datacenter, including monitoring and controlling delivery of electricity to servers within the datacenter based on the power supply. Ex.1007, [0002]; [0047]; [0080]-[0083]. Chapel teaches accomplishing this power management via a combination of the receptacles into which servers plug into for power, one or more local controllers, and a remote controller to implement power management commands and strategies. Ex.1007, [0080]-[0083]; [0090]. This includes a command in the form of a message from the remote controller, which the local controller then conveys to the relevant device(s) within the datacenter in order to control delivery of power to the servers. Ex.1007, [0090].

204. A POSITA would have had a reasonable expectation of success in combining the teachings of Chapel with Pelio and Kiani. Chapel's teachings are compatible with Kiani's and Pelio's teachings. For example, Kiani teaches providing power to datacenters so that the servers within can process user requests (workload). Ex.1005, p. 3. Pelio teaches providing power to servers within a datacenter container (or multiple such containers) from a variable power source behind-the-meter. Ex.1006, Abstract; [0002]; [0033]; [0037]; [0050]; [0106]. Chapel, in turn, teaches devices and methods by which to modulate the power delivery to one or more devices within a datacenter. Ex.1007, [0002]; [0047]; [0080]-[0083]. Chapel's teachings are compatible with Kiani's and Pelio's teachings, since Kiani's and Pelio's variable power supply would be monitored according to Chapel's monitoring teachings and proper control implemented per Chapel's power management teachings, in response to the variable power supply from the behind-the-meter power source. *See id.*

205. Combining Chapel's teachings with Kiani's and Pelio's teachings would have been predictable and within a POSITA's skillset. It was already known to locate datacenters behind-the-meter with renewable power sources. For example, beyond Kiani's and Pelio's teachings on this matter, others also taught locating datacenters "on generation sites and placed in containers," particularly co-located with "renewable generators." Ex.1012, pp. 1, 4; see also Ex.1014, p. 7 ("using

containers and colocating at generation sites and avoiding power distribution costs”); Ex.1013, p. 5 (modular containers at distributed site locations); Ex.1015, pp. 3, 4 (“us[ing] containers and direct siting at renewable generation sites” to harness curtailed power); Ex.1016, p. 1 (using “portable containerized modules ... for geo-distributed data centers”).

206. In particular, it was already well-known to control the load within the datacenter in response to the varying power supply of a renewable power source. For example, a POSITA understood that one way in which to implement server compute resources in a co-located datacenter container is to have the servers “only work when stranded power is available, and otherwise [have them] fully powered off.” Ex.1012, p. 4. In doing so, “power-aware computing” was able to control “the amount of power consumed by the computing based on the amount of power available.” Ex.1018, [0024]. A POSITA would have recognized that Chapel’s power monitoring and power control teachings at a datacenter would be responsive to the amount of power available from Kiani’s/Pelio’s behind-the-meter renewable power sources. Chapel teaches, for example, powering down when supply is lower than demand in order to avoid an over-capacity condition, or powering up when supply is sufficient for more load. *See* Ex.1007, [0020], [0037], [0049], [0050]. Implementing Chapel’s power control teachings would provide the benefit of reducing power consumption when the power supply drops, such as happens cyclically per Kiani’s and Pelio’s

recognition, and enable behind the meter datacenters to follow the power supply as it varies.

207. The combination of Kiani's and Chapel's teachings would predictably result in Kiani-Pelio's datacenters (e.g., containers) including local controllers to control power delivered to the server devices, using the teachings adapted from Chapel. *See* Ex.1005, pp. 2-3 (behind the meter with varying power supply); Ex.1006, Abstract; [0002]; [0033]; [0037]; [0050]; [0106] (locating containers behind-the-meter with potentially varying power supply); Ex.1007, [0002]; [0047]; [0080]-[0083] (monitoring and controlling power delivery). Doing so would advantageously enable Kiani's datacenters to better follow the power supply and avoid grid electricity costs when the renewable supply dwindles (such as from low wind, darkness, drawing from the grid, etc.).

208. Indeed, the combination of Kiani, Pelio, and Chapel is nothing more than the combination of known elements (Kiani's and Pelio's behind-the-meter datacenter teachings and Chapel's controller teachings within a datacenter), according to known methods (using Chapel's power monitoring and power delivery control teachings with Kiani's datacenters), to yield predictable results (Kiani's datacenters' power consumption following the amount of supply available from the behind-the-meter renewable generating source, according to Chapel's power delivery control teachings).

209. The combination permits but does not require physical incorporation of elements from Chapel into Kiani and Pelio. Nevertheless, the physical incorporation of Chapel's teachings for power monitoring and power delivery control would teach the claim limitations.

C. Detailed Analysis of Claims

1. Claim 4

- a.** [4.0] *The system of claim 1, wherein the flexible datacenter further comprises:*

210. See analysis at [1.0]-[1.11].

- b.** [4.1] *a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions.*


211. Kiani combined with Pelio and Chapel renders obvious a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions.

212. First, as demonstrated at [1.3], the combination of Kiani and Pelio renders obvious a power delivery system that delivers power to the plurality of computing systems.

213. Second, both Kiani and Pelio teach the desirability of placing datacenters near a power generating source "behind the meter." Ex.1005, pp. 2-3;

see also Ex.1006, [0002]. As a result of doing so, “a higher percentage of the power produced by the generation source actually reaches critical load” of the data center. Ex.1006, [0019]. Kiani and Pelio both describe several different types of power generating sources, including “wind turbine or solar panel” in Kiani, as well as “coal-fired, oil-fired, natural gas, hydro-electric, nuclear, wind farm, and solar energy types” taught by Pelio. Ex.1005, p. 2; Ex.1006, [0005].

214. Kiani and Pelio both recognize that renewable sources can have varying power output, due for example to “the weather dependency of green energy.” Ex.1005, p. 3. Kiani illustrates this variability with a wind power example in Figure 4. Figure 4 describes “the trends of wind power” at three different exemplary datacenters over the course of 24 hours, with each horizontal line showing the power output from wind turbines located BTM with the different datacenters:



Ex.1005, Fig. 4.

215. Pelio, similarly, recognizes that “[s]ome energy generating stations naturally cycle between full operation and zero, for example the daily cycles of a solar energy generating station that depends on sunshine,” as well as others that cycle with other renewable sources such as wind or demands from the grid. Ex.1006, [0037]. Thus, Pelio recognizes that there are times when power is less available from renewables (if at all) than other times, and requires some action to account for that shortfall. Both Kiani and Pelio’s solutions include to “draw inexpensive power back through” its connections to transmission and distribution grids, such as at night, or to pass jobs to other datacenters/power down. Ex.1006, [0037]; *see also* [0043]; [0105]-[0106]. Kiani describes this as using “brown energy.” Ex.1005, p. 4. The

amount of green energy available, the wind speed, sunlight conditions, cost of brown electricity, are all different examples of “*power system conditions*.”

216. Kiani and Pelio both leave it up to a POSITA as to all the various ways that power may be controlled at a data center due to the varying availability of power due to the varying conditions present at the datacenters. Thus, a POSITA would have been motivated to turn to Chapel, which provides details about how to control power within a datacenter to address times of actual or potential power shortfalls, such as would occur with renewable sources of power with variable output like some of those taught in Kiano and/or Pelio (e.g., solar, wind, etc.).

217. Third, Chapel teaches how to control power within a datacenter when actual or potential “over capacity condition[s]” exist. Ex.1007, [0048]. The available supply of power from a renewable generating station falling below the power demand of a data center, as mentioned in both Kiani and Pelio, is an example where Chapel’s control teachings are applicable. *See* Ex.1005, p. 3; Ex.1006, [0035]. Chapel teaches controlling power in a datacenter by modulating power delivery to one or more of the servers in the datacenter.

218. Chapel relates to “electrical power distribution and management,” including how to “intelligently monitor at least a portion of the circuit and control delivery of electricity over the circuit.” Ex.1007, [0002]. Chapel’s teachings are applicable in a variety of environments, including “in the context of a data center,”

with Chapel's focus being on "data center equipment control and power distribution systems." Ex.1007, [0019], [0047]. An example system enables "remote monitoring and/or control of multiple receptacles" which provide power to servers. Ex.1007, [0056]. Receptacles – which connect with servers – "communicate with a local controller 304" that can be at the "customer premises." Ex.1007, [0056]. The "local controller 304" can, in turn, communicate with a remote controller 308" that is operated "by a private or public party" such as an owner of the customer premises, a controller of the power grid, etc. Ex.1007, [0056]; *see also* FIG. 3:





FIG.3

Ex.1007, Fig. 3 (annotated).

219. The "local controller" is an example of a "*datacenter control system*." This is because it is a control system at the datacenter. Ex.1007, [0056] (can be located at the customer premises, which is a datacenter). Further, the local controller

at the customer premises is used to “control[] delivery of power to the electrical device.” Ex.1007, [0018]. Controlling delivery of power includes, per Chapel, turning off devices or operating the devices in reduced power mode – for example by interrupting power distribution. Ex.1007, [0025]-[0026].

220. Chapel’s Figure 6 provides a specific example of “control in a data center environment,” with data center devices (e.g., servers) that are plugged into controllable receptacles. Ex.1007, [0080]-[0081]. These are connected to “a local controller 650” that is itself in communication with a remote controller 670. Ex.1007, [0082]. “[R]emote and/or policy based operation of the” data center devices is achieved via the local controller 650 and remote controller 670. Ex.1007, [0083]. In an example, “an operator of a remote controller unit may choose to power down or power up a device or set of devices,” in which case “[a]n appropriate message is transmitted from the remote controller to the local controller,” which then implements the action. Ex.1007, [0090]; *see also* FIG. 6:



221. Fourth, Chapel teaches that the modulation of the power delivery by the local controller (a “*database control system*”) is done based on monitored conditions of the power system.

222. Kiani and Pelio both teach a datacenter receiving power via a BTM power source, with corresponding variability in renewable power sources. Ex.1005, pp. 2-3; Ex.1006, [0002], [0019], [0033], [0039]. It would have been obvious to a POSITA to monitor the varying output of these power sources in order to know how much green energy versus brown energy is available for workload allocation per datacenter (Ex.1005, p. 4), or to “throttle up and down” either other power sources (Ex.1006, [0033]), or the load itself (Ex.1006, [0106] (dropping load when generating output drops), Ex.1007, [0089] (powering up or down devices depending upon over-capacity condition)). That is, in order to know how much workload to

allocate as green workload/for green energy processing (per Kiani), and/or whether there is a need to throttle up or down another power source (per Kiani and Pelio), or the level of load itself, a POSITA would recognize that the current status would need to be monitored in both Kiani and Pelio.

223. Indeed, Chapel teaches monitoring power availability conditions. **“Power distribution may be monitored and controlled in relation to addressing actual or potential over capacity conditions.”** Ex.1007, [0004]. The power distribution is from power generation sources. Ex.1007, [0004], [0082] (power sources). Chapel teaches reporting information about loads to “allow[] for **remotely monitoring** and/or controlling electrical devices” in datacenters. Ex.1007, [0047]. In particular, Figure 3’s system 300 enables “remote monitoring and/or control of multiple receptacles” to deal with potential over-capacity situations. Ex.1007, [0056]. Chapel achieves this in Figure 3 via local controller 304 that communicates monitoring information to remote controller 308 (an example of a “*control system*”). Ex.1007, [0056]. Chapel’s Figure 6 includes another example, with local controller 650 sending information about power availability and load to remote controller 670. Ex.1007, [0083].

224. In combination, it would have been obvious to a POSITA for a remote controller (as taught by Chapel) to receive power information from one or more local controllers (and other data sources) about how much power is available from the one

or more BTM power sources (e.g., as taught by Kiani and Pelio, and recognized by Chapel). See Ex.1005, pp. 2-4; Ex.1006, [0033], [0106]; Ex.1007, [0004], [0089]. Other examples of monitored conditions include wind speed, solar cover, cost of brown electricity, and so on, as taught by Kiani. With the power conditions being monitored as taught by Chapel, Kiani's flexible datacenter (as further modified by Pelio's teachings) would thus monitor availability of power at the datacenter via the one or more BTM power sources.

225. The local controller at the datacenter (a "*datacenter control system*") per the combined teachings of Kiani, Pelio, and Chapel would modulate the power delivery to the servers, as taught by Chapel, based on monitored conditions (e.g., monitored at the remote controller of Chapel, an example of the FE server of Kiani with instructions sent to the local controller to implement).

226. Thus, Kiani combined with Pelio and Chapel renders obvious "*a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on one or more monitored power system conditions.*"

2. Claim 5

a. [5.0] *The system of claim 4, wherein the one or more monitored power system conditions comprises:*

227. See analysis at [1.0]-[1.11], and [4.0]-[4.1].

b. [5.1] *behind-the-meter (“BTM”) power availability at the plurality of computing systems.*

228. Kiani combined with Pelio and Chapel renders obvious wherein the one or more monitored power system conditions comprises: behind-the-meter (“BTM”) power availability at the plurality of computing systems.

229. First, as explained at [4.1], the combination of Kiani, Pelio, and Chapel teaches “*one or more monitored power system conditions.*”

230. Second, Kiani particularly teaches that the availability of BTM power at the datacenter – and, thus, at the “back-end servers” at the datacenter – is one of the monitored power system conditions. *See* Ex.1005, p. 3 (“based on the availability of green energy”); p. 2 (“back-end servers at geo-dispersed data centers”).

231. As further explained at [4.1], in order to know how much workload to allocate as green workload/for green energy processing (per Kiani), and/or whether there is a need to throttle up or down another power source (per Kiani and Pelio), or the level of load itself, a POSITA would recognize that the current status of green energy availability BTM at the datacenter (and, thus, to computing systems in the datacenter) would need to be monitored in both Kiani and Pelio. The combination with Chapel renders obvious such monitoring of Kiani’s BTM green energy availability. Ex.1007, [0004], [0047], [0082].

232. Thus, Kiani combined with Pelio and Chapel renders obvious “*wherein the one or more monitored power system conditions comprises: behind-the-meter (“BTM”) power availability at the plurality of computing systems.*”

3. Claim 6

- a.** **[6.0]** *The system of claim 4, wherein the one or more monitored power system conditions comprises:*

233. See analysis at [1.0]-[1.11], and [4.0]-[4.1].

- b.** **[6.1]** *one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation.*

234. Kiani combined with Chapel alone, and further combined with Pelio, renders obvious wherein the one or more monitored power system conditions comprises one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power

generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation.

235. Because [6.1] provides the list of monitored power system conditions as “*one or more of ... or,*” any one option from the list renders the entire claim obvious.

Low local power

236. First, it would have been obvious that in different situations, BTM collocated power sources taught in Kiani (and Pelio) would generate excess power or not enough power. Kiani teaches that “green energy is not a constantly available resource to power” datacenters. Ex.1005, p. 3. This is because green energy can have a “weather dependency,” such as wind speed for wind turbines, and cloud cover for solar generation. Ex.1005, p. 3. Thus, a POSITA would have recognized that times would arise when the collocated generation station generates less green energy than the datacenter needs for all of its workload. In Kiani, this corresponds to times where “green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4. This renders obvious the “*one or more monitored power system conditions comprises: ... low local power generation.*” Thus, Kiani alone renders obvious claim 6.

237. Second, the combination with Pelio further renders obvious low local power generation monitored conditions. As already noted, it would have been

obvious that in different situations, BTM collocated power sources would generate excess power or not enough power (including those taught in Pelio). For example, Pelio recognizes the well-known fact that some types of renewable power generating stations, such as solar and wind, “naturally cycle between full operation and zero.” Ex.1006, [0037]. Thus, a POSITA would have recognized that times would arise where the collocated generation station generates less power than the datacenter needs, such as nighttime when there typically is no sunshine during “the daily cycles of a solar energy generating station,” or low wind for wind generation stations. Ex.1006, [0037]; *see also* Ex.1008, FIG. 3 (periods where available capacity 302, 306 falls below actual computational needs 304, 308); 10:14-21. Pelio thus further renders obvious the “*low local power generation*.”

Excess local power generation

238. First, as a POSITA would have recognized, there would be times where there would be an excess of local power generation. Indeed, Kiani teaches the desirability of using “green energy as much as possible,” and only reverting to brown energy “when the green energy generation is not adequate to serve all incoming workloads.” Ex.1005, p. 4. More to the point, in addition to handling incoming workloads, Kiani teaches that “extra green energy generation at a data center can be injected into the power grid.” Ex.1005, p. 4. Thus, Kiani teaches to a POSITA times where there is enough green energy generated to service incoming workload that

there is excess green energy injected into the grid. This further renders obvious the “one or more monitored power system conditions comprises: ... excess local power generation at a local station level,” namely at the collocated renewable power source.

239. Second, the combination with Pelio further renders this limitation obvious. For example, Pelio recognizes the well-known fact that some types of renewable power generating stations, such as solar and wind, “naturally cycle between full operation and zero.” Ex.1005, [0037]. Thus, a POSITA would have recognized that times would arise where the collocated generation station would have varying levels of power availability for various reasons. For example, it would have been obvious from Pelio’s teachings of a collocated datacenter to a renewable power source that there are times when the power source generates more power than even the datacenter may need, such as peak sunshine during “the daily cycles of a solar energy generating station” or high wind for wind generation stations. Ex.1006, [0037]; *see also* Ex.1008, FIG. 3 (periods where available capacity 302, 306 exceeds actual computational needs 304, 308); 10:14-21.

240. As an example, when collocated with a wind farm, Pelio teaches that there are times where there is more than sufficient power for the datacenter that Pelio teaches further supplying “power to the distribution and/or transmission grids when the wind is good.” Ex.1006, [0044]; *see also* [0045] (solar), [0046] (hydro). This

describes extra generated power, which renders obvious “*excess local power generation at a local station level.*”

Transient local power

241. First, Kiani further teaches the monitored power system condition being transient local power generation situations. Kiani teaches that the local power generation by a renewable source is “*transient*” by describing, as already explained, that “green energy is not a constantly available resource to power” datacenters. Ex.1005, p. 3. This is because green energy can have a “weather dependency,” such as wind speed for wind turbines, and cloud cover for solar generation. Ex.1005, p. 3. This “weather dependency” describes a “*transient power generation situation*” since the power varies in its availability (and, thus, is transient).

242. Second, the combination of Kiani with Pelio further renders obvious this limitation. Pelio teaches that the availability of renewable power is “*transient*,” by describing how “a local data center’s electric generating plant went down due to lack of fuel, sun, wind, gas.” Ex.1006, [0106]. In short, “[s]ome energy generating stations naturally cycle between full operation and zero, for example the daily cycles of a solar energy generating station that depend on sunshine.” Ex.1006, [0037]. The cycling between full operation and zero demonstrates a transient power generating situation.

243. Any one of the above renders obvious the entire claim. Thus, Kiani alone, and further in combination with Pelio, renders obvious claim 6.

4. Claim 7

- a.** [7.0] *The system of claim 1, wherein the flexible datacenter further comprises:*

244. See analysis at [1.0]-[1.11].

- b.** [7.1] *a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive.*

245. Kiani combined with Pelio and Chapel renders obvious wherein the flexible datacenter further comprises: a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive.

246. First, as demonstrated with respect to [4.1], the combination of Kiani, Pelio, and Chapel renders obvious a “*datacenter control system ... configured to modulate power delivery to the plurality of computing systems*” based on monitored power system conditions.

247. Second, the combination of Kiani, Pelio, and Chapel further renders obvious the datacenter control system modulating power delivery “*based on an operational directive.*”

248. Chapel teaches that the modulation of the power delivery by the local controller (a “*datacenter control system*”) is done by policies including those

communicated from a remote controller. “For example, the local controller may receive a message from the electrical utility provider specifying transition to an efficiency mode or a brownout mode.” Ex.1007, [0076]; *see also* [0058], [0066], [0090]. Such “messages,” and the policies communicated from remote controller(s), render obvious the claimed modulation of the power delivery “*based on an operational directive*.”

249. In combination, therefore, Kiani, Pelio, and Chapel teach to a POSITA a “local controller” as in Chapel determining to control power delivery by interrupting power distribution from the power distribution system to one or more servers within Kiani-Pelio’s flexible datacenters according to a message received from a remote controller. Ex.1007, [0017], [0025]-[0026], [0076]. For example, a POSITA would have understood that Kiani-Pelio’s flexible datacenters, when collocated with a power generation station (BTM, per both Kiani and Pelio) that has a reduction in power output (e.g., solar or wind conditions changing), would have times when an overload condition could arise unless power is received from the grid (instead of from the BTM source) or demand (load) is reduced. As a result, a POSITA would have implemented Chapel’s teachings for reducing delivery of power to one or more servers within a datacenter to reduce demand (the load from the operating servers) and avoid an overload condition.

250. Thus, the combination of Kiani, Pelio, and Chapel renders obvious “*wherein the flexible datacenter further comprises: a datacenter control system, wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive.*”

5. Claim 8

a. [8.0] *The system of claim 7,*

251. See analysis at [1.0]-[1.11] and [7.0]-[7.1].

b. [8.1] *wherein the operational directive originates from the routing control system and depends on behind-the-meter (“BTM”) power availability at the plurality of computing systems.*

252. Kiani combined with Pelio and Chapel renders obvious wherein the operational directive originates from the routing control system and depends on behind-the-meter (“BTM”) power availability at the plurality of computing systems.

253. First, as demonstrated with respect to [7.1], the Kiani-Pelio-Chapel combination renders obvious modulating power delivery based on an “*operational directive*” with Chapel’s teaching of a remote controller sending a message to the local controller to modulate the power distribution to servers within a datacenter. See Ex.1007, [0076]; *see also* [0058], [0066], [0090]. As further explained at [4.1], Chapel’s “remote controller” is an example of the “FE server” taught in Kiani. See Ex.1005, p. 2; Ex.1007, [0056], [0083].

254. Second, as explained at [5.1], the combined Kiani-Pelio-Chapel teachings further render obvious monitoring BTM power availability at the servers within Kiani's datacenters (which are flexible datacenters, in combination with Pelio's teachings). Ex.1005, p. 2; Ex.1006, Abstract, [0049], [0071]; Figure 2; Ex.1007, [0004], [0047], [0056], [0082]-[0083].

255. In combination, therefore, Kiani-Pelio-Chapel renders obvious Kiani's FE server monitoring BTM power availability as a remote controller according to Chapel's teachings, and providing a message to locally control a flexible datacenter (per Chapel's and Pelio's teachings), depending on the amount of BTM power available to the servers within the flexible datacenter.


256. Thus, the combination of Kiani, Pelio, and Chapel renders obvious *"wherein the operational directive originates from the routing control system and depends on behind-the-meter ("BTM") power availability at the plurality of computing systems."*

XI. CHALLENGE 3: THE COMBINATION OF KIANI, PELIO, AND BELADY RENDERS OBVIOUS CLAIM 3

A. Summary of Belady

257. U.S. Patent No. 8,849,469 to Belady *et al.* (Ex.1008, "Belady") was filed October 28, 2010, and issued September 30, 2014. Belady is titled "Data Center System That Accommodates Episodic Computation." Belady was not before the Examiner during prosecution of the '261 patent.

258. Belady relates to “[a] data center system ... which includes plural data centers powered by plural power sources.” Ex.1008, 2:6-7. “The plural power sources can include any combination of renewable power sources (such as wind power sources, solar power sources, etc.)” that can be off-grid or on-grid. Ex.1008, 2:7-14. Different datacenters are placed at different geographic locations and connected with different ones of the plural power sources. Ex.1008, 10:14-34. The datacenters can be co-located with the renewable power sources to avoid the losses and transmission limits associated with transformation equipment. Ex.1008, 6:53-7:17; FIG. 1:



Ex.1008, Fig. 1 (annotated).

259. Instead of moving power, Belady teaches moving computational tasks either temporally or spatially, i.e. between different datacenters based on power availability or other considerations. Ex.1008, 7:49-60 (temporally), 7:61-8:8 (spatially). A “management system ... manag[es] execution of a plurality of computational tasks by the plurality of data centers ... by moving data components to the sites where power is available and cost-effective.” Ex.1008, 2:14-22.

260. The management system temporally moves computational tasks by “adjust[ing] a time at which a data component, associated with a computational task, is to be performed within a data center.” Ex.1008, 2:23-26. The management system spatially moves “a data component from a first data center to a second data center,” more generally described as moving from one location to another for computation. Ex.1008, 2:36-41. For example, when there is a “lull[] in the available capacity” of a renewable power source at one location, the management system transfers data components for processing to another datacenter at a location that has more available capacity. Ex.1008, 10:20-32.

261. The management system uses “real-time information provided by the data centers” about the availability of power from “various power sources at each data center” as well as additionally “predict[ed] power source availability.” Ex.1008, 12:41-47 (real-time), 12:48-53 (predicted). This information is used to determine how to direct operations at one or more datacenters to proceed. Ex.1008, 10:60-62

(management system uses the information to manage execution of computational tasks).

B. Reasons to Combine Kiani and Pelio with Belady

262. A POSITA would have been motivated to combine Belady with Kiani and Pelio for several reasons.

263. First, Belady, like Kiani and Pelio, focuses on analogous issues to those in the '261 patent, including systems for controlling the use of BTM power. Belady teaches powering different datacenters at different locations with respective power sources, including renewable power sources that are behind-the-meter connected with a given datacenter. Ex.1008, 2:6-14, 6:53-7:17. Thus, just like Kiani, Pelio and the '261 patent, Belady teaches co-locating computing resources at power generating stations to harness the power output behind-the-meter. Thus, Belady is in the same field of endeavor as the '261 patent.

264. Further, Belady is pertinent to the problems identified by the '261 patent, such as “the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources,” and dealing with transmission and distribution costs over a grid. *See* Ex.1001, 35:11-15. Belady provides multiple solutions to these problems. For example, Belady also teaches co-locating the datacenters to the renewable power sources, in order to reduce issues like infrastructure costs and transmission losses.

Ex.1008, 7:1-9. Further, Belady also teaches moving computational tasks either temporally (earlier or later), spatially (to different location), or both temporally and spatially in order to take advantage of “where power is available and cost-effective.” Ex.1008, 2:15-22. Thus, a POSITA would have also found Belady’s teachings to be pertinent to the problems identified by the ’261 patent.

265. Second, a POSITA would have been motivated to turn from Kiani and Pelio to Belady for Belady’s teaching of implementing a “management system 120” in multiple ways, including in “distributed fashion ... by the processing equipment provided by the individual data centers.” Ex.1008, 8:23-26. The placement of a management system, of which Kiani’s FE server is an example (because both manage the distribution of tasks/workload amongst datacenters based on various criteria and needs), was a mere design choice. There were a limited number of options available to a POSITA that were obvious to try, which Belady lists. According to Belady, the options include (1) “distinct from the individual data centers,” (2) implemented in by the data centers, or (3) split between remote and local locations. *See* Ex.1008, 8:15-31. Kiani demonstrated the implementation of option (1) for the FE servers, and Belady demonstrated options (2) and (3) as well. A POSITA would have been motivated to implement, for example, option (2) with the management system at a datacenter as well, in order to benefit from already-existing hardware, instead of requiring additional hardware at another location.

Moreover, a POSITA would have had a reasonable expectation of success in implementing Belady's teaching of the management system at a datacenter with Kiani's FE server functions, since there were only a limited number of options, all of which had been successfully implemented before as demonstrated by Belady. Ex.1008, 8:15-31; *see also* Ex.1009. In combination, it would have been obvious for Kiani's FE server operations to occur at a datacenter, per Belady's teachings, by operating for example on top of the datacenter's functionality.

266. Third, a POSITA would have also been motivated to turn from Kiani and Pelio to Belady in order to take advantage of Belady's teachings of how to further move computational tasks due to power availability issues, and how to detect and determine that it is appropriate to move. For example, Kiani noted that in addition to geographical load balancing for the distribution of computational operations, "load shifting algorithms" that shift the load after already having been distributed is also an option. Ex.1005, p. 2. As another example, Pelio already recognized the desirability of moving pending compute jobs to another datacenter when there is insufficient renewable power at the currently assigned datacenter. Ex.1006, [0106]; *see also* Ex.1013, p. 1 (demonstrating it was known to shift computational load between datacenters due to non-availability of intermittent power). A POSITA would be motivated to turn to Belady because Belady's teachings provide implementation details for Kiani's and Pelio's suggestions.

267. For example, Belady teaches using a management system to manage the execution of tasks – like Kiani’s “requests” and Pelio’s “data center jobs” amongst multiple datacenters in view of variable power availability. Ex.1008, 2:15-22. This includes, for example, moving computation tasks from one datacenter powered by a solar generating source in one location that is experiencing a lull (e.g., due to nighttime) to another datacenter powered by another solar generating source that is geographically in daytime. Ex.1008, 10:14-34. A POSITA would have been motivated to turn to Belady’s teachings for details of an example of a “load shifting algorithm” per Kiani, and Pelio’s example of datacenter job movement.

268. Fourth, POSITA would have had a reasonable expectation of success in applying Belady’s movement of computational tasks teachings to the combined teachings of Kiani and Pelio. Like both Kiani and Pelio, Belady teaches placing the datacenter co-located with the renewable generating source, whether the generating source is off-grid or on-grid. Ex.1008, 6:53-7:17; *see also* Ex.1005, pp. 2-3 (behind the meter), Ex.1006, [0002] (same). Moreover, each of Kiani, Pelio and Belady recognize the utility of moving computing jobs/tasks between datacenters for cost advantage (e.g., Kiani) and/or in response to power availability issues at the current datacenter (such as due to a variable power output from the renewable generating source). Ex.1005, p. 2; Ex.1006, [0106]; Ex.1008, 2:14-22. Since all of Kiani, Pelio and Belady teach moving computational jobs/tasks between datacenters, and siting

those datacenters behind-the-meter, a POSITA would have recognized the compatibility of Kiani's, Pelio's and Belady's systems, and thus had a reasonable expectation of success in implementing Belady's implementation details teachings for moving tasks between datacenters.

269. In combination, Kiani's datacenters (e.g., implemented as Pelio's flexible datacenters) would implement Belady's real-time monitoring teachings, and computational task movement teachings, with Belady's management system teachings for implementing task movement. Doing so would have been predictable and within a POSITA's skillset. Indeed, Kiani and Pelio already recognized the desirability of moving datacenter jobs (e.g., due to cost and/or power availability issues), and Belady teaches specific ways in which to do so. And it was already well-known, in particular, to use behind-the-meter renewable power sources for the datacenters' tasks. Ex.1005, pp. 2-3; Ex.1006, Abstract, [0002]; Ex.1012, pp. 1, 4; Ex.1014, p. 7; Ex.1013, p. 6; Ex.1015, pp. 3, 4.

270. The combination permits but does not require physical incorporation of elements from Belady into Kiani. Nevertheless, the physical incorporation of Belady's operations into Kiani's system for implementation would teach the claim limitations.

C. Detailed Analysis of Claim

1. Claim 3

a. [3.0] *The system of claim 1,*

271. See analysis at [1.0]-[1.11] in Challenge 1.

b. [3.1] *wherein the routing control system is a remote master control system positioned at the flexible datacenter.*

272. Kiani combined with Pelio and Belady renders obvious wherein the routing control system is a remote master control system positioned at the flexible datacenter.

273. First, as demonstrated with respect to [2.1] in Ground 1, Kiani teaches that the “*routing control system is a remote master control system.*” See Ex.1005, p. 2.

274. Second, as further demonstrated with respect to [1.1], the combination of Kiani and Pelio renders obvious the “*flexible datacenter.*”

275. Third, while Kiani’s examples illustrate the FE server as being positioned separately from the datacenters (flexible datacenters, in combination with Pelio), it would have been obvious to a POSITA that the placement of the control system is a mere design choice, with options including placing the control system remotely from the datacenter, at the datacenter, or splitting the control functions between remote and local systems. Belady demonstrates this.

276. Belady teaches a “management system” that “manag[es] execution of a plurality of computational tasks by the plurality of data centers ... by moving data components to the sites where power is available and cost-effective.” Ex.1008, 2:14-22. In an example, Belady teaches that management functionality 602 “can be implemented by the management system 120 of FIG. 1.” Ex.1008, 12:20-30. And, in particular, Belady teaches that the management system 120 – an example of a control system like Kiani’s FE server – “the functionality provided by the management system 120 can be implemented, in distributed fashion, by the processing equipment provided by the individual data centers.” Ex.1008, 8:23-26; *see also* 8:15-20, 8:26-31 (demonstrating the other two design choice options of placing the control system remote from the datacenter, or splitting the control functions between remote and local systems).

277. In combination, it would have been obvious to a POSITA to implement Kiani’s FE server in a “distributed fashion” by the processing equipment at each datacenter. This was one of several options that were obvious to try. In particular, while Kiani’s original concept of datacenters was to have “powerful data centers,” this was by way of example. For example, Kiani explained that “Internet-scale powerful data centers are few because of the scale and cost of deployment and operation,” and are “generally dispersed in different geographic regions.” Ex.1005, p. 2. To achieve this diversity in location (for scalability, availability, and

price/location diversity), Kiani explained that it was typical for “power datacenters” to be “generally deployed far away from a large majority of users.” Ex.1005, p. 2. But Kiani does not preclude the possibility of datacenters – whether “Internet-scale” or not – being placed closer to users such that implementing an FE server at the datacenters themselves (per Belady’s teachings) would result in the FE server still being close to users while being at the datacenter.

278. Further, when Kiani’s teachings are combined with Pelio’s teaching of flexible, highly-portable datacenters-on-wheels, the result would be much more mobile datacenters that can be placed in a variety of locations. *See* Ex.1005, p. 2 (placing datacenters to achieve redundancy, scalability, high availability, price diversity, location diversity, etc.); Ex.1006, [0055]-[0060] (modular, containerized datacenters provide flexibility in upgrading and location placement, as well as more efficient cooling). With the smaller, more flexible datacenters according to Pelio’s teachings, a POSITA would have therefore recognized that Pelio’s teachings would result in the benefit of more flexibility in location for the datacenters, including at locations coincident with the FE servers taught in Kiani. Further, the function of the FE server would remain the same as Kiani teaches, since even with co-location of the FE server functionality there may still be differences in “server or content availability ... efficiency of the data centers, the cost of the electricity, and availability of renewable energy” with respect to different datacenters, in which case

the FE server still must perform its workload distribution function based on those parameters and the workload strategies of the FE server. Ex.1005, p. 2.

279. Thus, Kiani combined with Pelio and Belady renders obvious “*wherein the routing control system is a remote master control system positioned at the flexible datacenter.*”

XII. CHALLENGE 4: THE COMBINATION OF KIANI, PELIO, CHAPEL, AND FORESTIERO RENDERS OBVIOUS CLAIMS 13-16

A. Summary of Forestiero

280. Forestiero is an IEEE publication titled “Hierarchical Approach for Efficient Workload Management in Geo-Distributed Data Centers.” Forestiero was not before the Examiner during prosecution of the ’261 patent. Forestiero was first published via the IEEE Xplore digital, online library on August 31, 2016. *See* Ex.1009, p. 1; *see also* Ex.1022 (Xplore listing). I have extensive experience using the IEEE Xplore online library, and in my experience publications added to the Xplore library are downloadable (upon purchase, or with the appropriate subscription) by any interested party as of the date the publication is added to IEEE Xplore. In this case, that was August 31, 2016, which in my opinion is the date on which Forestiero was publicly disseminated via IEEE Xplore as an indexed, searchable public document (in IEEE Xplore’s system) available for purchase and download by any interested party.

281. I have reviewed Forestiero, and it bears the typical indicia of an IEEE publication, including the title of the Journal at the top, the copyright date at the bottom (2016), and document identifiers and cost for purchase. Ex.1009, p. 1. POSITAs were well-versed in accessing documents of interest from IEEE Xplore, and therefore it is my opinion that a POSITA would have been able to locate Forestiero with reasonable diligence as of August 31, 2016.

282. Yet further, Forestiero was also included in Volume 1, Issue 1 of the IEEE Transactions on Green Communications and Networking, which issued in March 2017. Ex.1009, p. 1; *see also* Ex.1022 (Xplore listing). I have extensive experience publishing in various IEEE Journals, and in my experience papers included in IEEE Journals are disseminated to recipients well before the end of the month of the issue's distribution (here, March 2017). At the latest, the Journal is available in print by the end of the identified month, meaning that Forestiero (as part of the IEEE Journal) was publicly disseminated in print form, and available for review by recipients and/or in libraries with reasonable diligence, as of March 31, 2017 (the end of the month of Journal dissemination).

283. Forestiero describes a “hierarchical approach for workload management in geographically distributed [data centers, or DCs].” Ex.1009, p. 1 (Abstract). Forestier's framework includes “two layers: at the *lower layer*, each site adopts its own strategy to distribute and consolidate the workload internally,” and

an “*upper layer*” with “a set of algorithms – shared by all the sites – ... used to evaluate the behavior of single sites and distribute the workload among them.” Ex.1009, pp. 1-2 (emphasis in original). Figure 1 illustrates the upper and lower layers:




Fig. 1. EcoMultiCloud scenario: upper and lower layer of two interconnected data centers.

Ex.1009, Fig. 1 (p. 4).

284. Forestiero teaches using a decentralized approach for load management (instead of a centralized approach, e.g. a centralized server) to avoid the risks of “poor scalability,” “poor ability to adapt to changing conditions,” and “lack of autonomy” at the different datacenters. Ex.1009, p. 3. Forestiero therefore teaches “decentraliz[ing] part of the intelligence and distribute the decision points, while still exploiting the centralized architecture and functionalities offered by virtualization infrastructures in single data centers.” Ex.1009, p. 3.

285. Each datacenter includes a “data center manager,” or DCM, that operates with the upper and lower layers. The DCM “runs the algorithms of the upper layer,” with a local manager (LM) operating the lower layer – where typically the managers “may be deployed on the same host.” Ex.1009, p. 4. At the upper layer, the DCM “acquire[s] detailed knowledge about the current state of the local [datacenter],” from which the DCM extracts “high level information” about the local datacenter’s state, “transmit[s]/receive[s] such high level information to/from all the other DCMs” at other, interconnected datacenters, and “execute[s] the algorithms of the upper layer to combine the collected information and take decisions about the distribution of the workload among the [datacenters].” Ex.1009, p. 4.

286. As a result, “all the DCMs [at the different datacenters] execute the upper layer algorithms” to choose the target datacenter for a workload task (described as a virtual machine in Forestiero, which is Forestiero’s way of describing workload allocation). Ex.1009, p. 4. Forestiero teaches that “the amount of information to be distributed” between datacenters (the DCMs at each) “is tiny, i.e., of the order of a few packets,” such that the information exchange is not burdensome between the datacenter sites. Ex.1009, p. 4.

B. Reasons to Combine Kiani, Pelio and Chapel with Forestiero

287. A POSITA would have been motivated to combine Kiani, Pelio, and Chapel with Forestiero for several reasons.

288. First, Forestiero describes approaches for handling workload in geo-distributed datacenter systems, including a datacenter manager at each datacenter that communicates with the other datacenters' DCMs to collectively choose a target datacenter to assign a workload to, and a local manager at each datacenter to manage workload at a datacenter site once workload is assigned there. Ex.1009, p. 4. Forestiero's approaches are designed for multi-datacenter scenarios that can balance multiple business objectives and constraints and adapt workload distribution to varying conditions. Ex.1009, p. 4. Forestiero's teachings enable different datacenters to better manage their own workload to further optimize performance and cost, which are issues pertinent to the '261 patent. For example, the '261 patent identifies the different costs of different locations and types of power as a problem in the management of workload. *See, e.g.*, Ex.1001, 6:14-17 (different datacenters having different costs of operations), 6:40-42 (aiming to reduce cost of using a datacenter). A POSITA would have therefore found Forestiero pertinent to the problems identified in the '261 patent.

289. Second, a POSITA would have been motivated to turn to Forestiero's geo-distributed workload management teachings when considering Kiani's geographical load balancing teachings in order to take advantage of Forestiero's benefits of decentralized workload management. Implementing the decision intelligence in a decentralized manner via a datacenter manager (DCM) at each

datacenter would provide the benefit of removing the limitations of centralized approaches, including the risks of poor scalability, poor ability to adapt to changing conditions, and lack of autonomy of single datacenters. *See* Ex.1009, p. 3. For example, Kiani's geographical load balancing teachings would benefit by the FE servers taught by Kiani having their decision-making burden reduced to simply fielding new requests, with the decision logic implemented by the network of datacenter managers at the different datacenter locations. *See* Ex.1005, p. 2 (FE servers receiving user requests); Ex.1009, p. 4 (all the DCMs collectively choose the target datacenter for the user request (described as a VM in Forestiero)).

290. It would have been obvious to a POSITA that workload management is typically handled in only a small number of ways, including either via a centralized approach (e.g., Kiani's FE servers making workload distribution determinations), a decentralized approach (e.g., Forestiero's approach), or split between the two. Such approaches were already well-recognized, as Forestiero explains. *See* Ex.1009, p. 3 (comparing centralized and decentralized approaches); *see also* Ex.1008, 8:23-26 (the management system functionality can be implemented "in distributed fashion"), 8:15-20, 8:26-31 (other options of implementation being remote from datacenters or splitting between the two). Thus, it would have been obvious to a POSITA to try implementing the decision-making for load distribution in a decentralized manner, such as taught by Forestiero. Kiani already contemplated the benefit of distributing

the decision-making among multiple FE servers, including “different workload distribution strategies [that] can be adopted at each FE [server] by considering different objectives” unique to each site. Ex.1005, p. 2.

291. Turning to Forestiero’s distributed decision-making teachings would have beneficially provided Kiani’s system with a hierarchical architecture (e.g., upper and lower layers to take into account system-wide and local issues when determining how to allocate and/or re-allocate workload). Ex.1009, pp. 3, 4. The combination would have further provided the benefit of “exploiting the centralized architecture and functionalities” of Kiani’s system while decentralizing the decision-making for autonomy and increased ability to adapt to changing conditions. Ex.1009, p. 3.

292. Third, a POSITA would have had a reasonable expectation of success combining Forestiero’s teachings with Kiani’s system teachings. A POSITA would have recognized that Forestiero’s distribution of the eventual workload distribution decision to a controller at each respective datacenter (e.g., in a collective decision-making process with all other datacenter controllers) as a natural extension of Kiani’s recognition of the benefits of distributing decision-making. *See* Ex.1009, p. 4 (distributed decision-making at each datacenter in a collective fashion); *see also* p. 3 (decentralized decision-making improves ability for locations to adapt to changing conditions locally). Thus, centralized and decentralized decision-making

approaches each already had demonstrated a reasonable expectation of success in their implementation.

293. Kiani's main reason for using the FE servers is to place them near to users for speed and efficiency. Ex.1005, p. 2. That each FE server can further implement different workload distribution strategies is an added benefit. *See* Ex.1005, p. 2. In combination with Forestiero's distributed decision-making teachings, it would have been obvious for Kiani's FE servers to receive user workload requests so that requests are received close to the users, and the decision-making of final distribution to occur in a distributed yet cooperative manner amongst the different, geographically-dispersed datacenters. Such a combination would include the FE servers contributing "high level information" as part of the information exchange that Forestiero teaches amongst the datacenter managers at each datacenter, either with the FE servers operating as effectively peers to the datacenter managers, or in similar capacity as the local managers taught in Forestiero.

294. For example, Forestier teaches that each datacenter manager communicates with its local manager to obtain "detailed knowledge about the current state of the local DC," which is then used as part of the information exchange and decision-making process performed collectively by the datacenter managers. *See* Ex.1009, p. 4. It would be a trivial detail to a POSITA to implement Kiani's FE

servers as information sources in a manner similar to Forestiero's local managers, as Forestiero's teachings are more about exchanging and using the information as it is about the particular sources of information. Indeed, Forestiero contemplates the local manager not necessarily being deployed in the same host as the datacenter manager, which a POSITA would have appreciated to mean other locations (and, thus, communication between the two) are possible as well – such as Kiani's FE server locations. *See* Ex.1009, p. 4 (DCM and LM “may be deployed on the same host”).

295. As another example, Forestiero teaches that each datacenter manager shares “high level information” with all of the other datacenter managers in the network. Ex.1009, p. 4. It would have been obvious to a POSITA for Kiani's FE servers to be a peer source of appropriate “high level information” shared with all of the datacenter managers, without the FE servers necessarily having to be part of the decision-making process. *See* Ex.1005, p. 2 (FE servers implementing decision-making based on different types of high-level information). With the FE servers sharing their information (such as green energy utilization, electricity cost minimization, profit maximization, etc. as taught by Kiani, p. 2), controllers at datacenters would be able to make even more efficient decisions according to Forestiero's distributed decision-making teachings.

296. Based on these facts, a POSITA would have found the combination of Kiani's (including Pelio's and Chapel's further teachings) and Forestiero's teachings

predictable and within a POSITA's skillset. Specifically, based on the above facts, a POSITA would have expected the combination of Kiani's system teachings and Forestiero's distributed decision-making teachings to result in Kiani's FE servers sharing their information with the other datacenter controllers to contribute to Forestiero's collective decision-making. This would provide Kiani's system with the benefits identified by Forestiero, including the benefits of a hierarchical architecture, multiple-objective and multiple-constraint balancing among locations, and increased adaptability between locations. *See* Ex.1009, pp. 3, 4.

297. The combination of Forestiero and Kiani would have therefore been nothing more than the combination of known elements (Kiani's distributed FE server decision-making and Forestiero's further-decentralized information sharing and decision-making) according to known methods (using Forestiero's distributed decision-making at the different datacenters with Kiani's different datacenters and already-existing distributed decision-making at different FE servers), to yield predictable results (Kiani's FE servers fielding user requests with distributed decision-making at the datacenters taking into account information shared from all the datacenters and FE servers, as taught by Forestiero).

298. The combination permits but does not require physical incorporation of elements from Forestiero into Kiani. Nevertheless, the physical incorporation of Forestiero's teachings would teach the claim limitations.

C. Detailed Analysis of Claim

1. Claim 13

a. [13.0] *A flexible datacenter comprising:*

299. See analysis at [1.1] in Challenge 1.

b. [13.1] *a behind-the-meter (“BTM”) power input system configured to receive power from a power generation system prior to the power undergoing step-up transformation for transmission to a grid and provide the received power to the flexible datacenter instead of sending the received power to the grid;*

300. See analysis at [1.2] in Challenge 1.

c. [13.2] *a power distribution system;*

301. See analysis at [1.3] in Challenge 1.

d. [13.3] *a datacenter control system;*

302. See analysis at [4.1] and [7.1] in Challenge 2.

e. [13.4] *a plurality of computing systems, wherein the plurality of computing systems are powered by the BTM power input system via the power distribution system; and*

303. See analysis at [1.4] in Challenge 1.

f. [13.5] *a queue subsystem configured to organize a set of computational operations for performance by the plurality of computing systems, wherein the queue subsystem is part of a distributed queue system;*

304. Kiani discloses, and renders obvious, a queue subsystem configured to organize a set of computational operations for performance by the plurality of

computing systems, wherein the queue subsystem is part of a distributed queue system.

305. First, as explained at [1.5], Kiani discloses a queue subsystem at each datacenter that receives workload assignments from an FE server. Ex.1005, p. 5 (requests are “first placed in a queue” at a datacenter before being processed by any server at the datacenter, with a “queue at each data center”).

306. Second, as further explained at [1.5], the queues at each datacenter, together with or without the FE servers queues (that would have been obvious to a POSITA in view of Kiani’s teachings), demonstrates that the “*queue subsystem*” at a given datacenter “*is part of a distributed queue system*” (Kiani’s system that includes a queue at each datacenter).

307. Third, as explained at [1.6], Kiani explains that the queue at each datacenter organizes computational operations at that datacenter for execution (“*performance*”) by servers at that datacenter. *See* Ex.1005, pp. 3-5 (managing workload according to availability of green energy versus brown energy at a datacenter, maintaining an upper bound on a datacenter queue, maintaining QoS requirements).

308. Thus, Kiani discloses, and renders obvious, “*a queue subsystem configured to organize a set of computational operations for performance by the*

plurality of computing systems, wherein the queue subsystem is part of a distributed queue system.”

- g. [13.6] wherein the datacenter control system is configured to: determine, using the queue system, a computational operation to be performed;**

309. Kiani combined with Chapel and Forestiero renders obvious wherein the datacenter control system is configured to: determine, using the queue system, a computational operation to be performed.

310. First, as demonstrated with respect to [1.8], Kiani discloses a control system that determines, using the queue system, a computational operation to be performed. The FE server “receives requests from its nearby users and distributes the requests to the back-end servers at geo-dispersed data centers,” such that each FE server “functions as a workload distribution center that manages the workload by distributing user requests to the appropriate data centers.” Ex.1005, p. 2. The “incoming workload” to the FE server represents a “*computational operation to be performed,*” because the workload constitutes operation requests from users. Ex.1005, pp. 2-3. By knowing what workload the FE server has to distribute, the FE server has “*determine[d], using the queue system,*” the computational operation to be performed (and, thus, distributed to a datacenter for performance).

311. Second, as noted at [13.3] and explained in detail with respect to [4.1] and [7.1], Chapel teaches a datacenter control system with Chapel’s local controllers,

which are used to control power within a datacenter, including power distribution and management. Ex.1007, [0002]; *see also, e.g.*, Figure 2:




FIG.3

Ex.1007, Fig. 3 (annotated).

312. Given the presence of a local controller at each datacenter, according to the combined teachings of Kiani and Chapel, it would have been obvious to a POSITA to implement additional functionality at the datacenter's local controller to further achieve power management, including by distributing the control of workflow allocation to the datacenters' local controllers.

313. Indeed, it would have been obvious to a POSITA that there are several different ways in which to implement the determination of a computational operation to be performed. Kiani, for example, describes a centralized approach where an FE server both fields the requests from users (the workload) and assigns the workload

to appropriate datacenters. It would have been obvious to POSITAs to instead implement a decentralized management approach for geo-distributed datacenters, as demonstrated by Forestiero, via the local controllers taught by Chapel.

314. Forestiero describes the allocation of workload in a decentralized manner by allocating a “virtual machine” associated with a given job request. Thus, when Forestiero describes allocating virtual machines, it is describing the allocation of workload – according to Forestiero, “workload assignment.” Ex.1009, pp. 1, 2.

315. In particular, Forestiero teaches a two-layer framework for handling workload in geo-distributed datacenter systems. The lower layer deals with a “local manager (LM) [that] performs the functionalities of the lower layer,” including managing the workload of that datacenter where the LM is located once the workload is assigned there by the upper layer. Ex.1009, p. 4. The upper layer includes a datacenter manager (DCM) at each datacenter that “runs the algorithms of the upper layer.” Ex.1009, p. 4; *see also* Figure 1:




Fig. 1. EcoMultiCloud scenario: upper and lower layer of two interconnected data centers.

Ex.1009, Fig. 1 (annotated).

316. A given DCM – such as the DCM at datacenter 1 in Figure 1 – “integrates the information coming from the lower layer and uses it to implement the functionalities of the upper layer,” including “communicate with the local LM” to acquire the details about the datacenter’s current state, “extract relevant high level information about the state of the [datacenter],” “transmit/receive such high level information to/from all the other DCMs [at other datacenters],” and “execute the algorithms of the upper layer to combine the collected information and take decisions about the distribution of the workload among the [datacenters].” Ex.1009, p. 4. Forestiero explains that the “framework is designed so that all the DCMs execute the upper layer algorithms and, for example, choose the target [datacenter] for a VM originated locally.” Ex.1009, p. 4. Each DCM makes its determination

Ex.1003

based on the information received locally and from other datacenter sites. The framework is illustrated in Figure 2 with each DCM executing the upper layer algorithm for choosing the target datacenter for a workload assignment:




Fig. 2. EcoMultiCloud scenario: the DCMs of four data centers exchange high level information about the state of local data centers. Such information is used, for example, to decide which site should accommodate a new VM.

Ex.1009, Fig. 2 (annotated).

317. In combination, it would have been obvious for the geographically distributed datacenters taught in Kiani to each include DCM functionality within the local controller (as taught by Chapel) according to the decentralized DCM teachings of Forestiero. *See* Ex.1009, p. 3. Kiani's FE servers would still receive incoming workload requests from nearby users, thus taking advantage of Kiani's teaching of FE servers receiving requests from nearby users, while the decision intelligence would be decentralized per Forestiero's teachings to the local controllers taught by

Chapel that are used for power distribution and management. Indeed, as Forestiero teaches, the DCM “may be deployed on the same host” as other local entities at the datacenter, such as a “local virtualization infrastructure” or, as a POSITA would recognize, Chapel’s local controller functions for modulating power delivery. *See* Ex.1009, p. 4. As a result, the combination teaches the “*datacenter control system*” (the DCM implemented at a local controller, such as taught by the combination of Chapel and Forestiero) determining a computational operation to be performed (the workload/virtual machine associated with a job per the Kiani and Forestiero combined teachings).

318. Implementing the decision intelligence in a decentralized manner via a DCM at each datacenter would provide the benefit of removing the limitations of centralized approaches, including the risks of poor scalability, poor ability to adapt to changing conditions, and lack of autonomy of single datacenters (as explained by Forestiero). *See* Ex.1009, p. 3. And, while the particular embodiment of Forestiero’s Figure 2 provides a whole mesh architectural example, Forestiero explains that other network architectures are also possible to accommodate a larger number of datacenters. *See* Ex.1009, pp. 4-5.

319. Thus, Kiani combined with Chapel and Forestiero renders obvious “*wherein the datacenter control system is configured to: determine, using the queue system, a computational operation to be performed.*”

- h. [13.7] *[wherein the datacenter control system is configured to:] determine an amount of BTM power that is available to the plurality of computing systems;***

320. Kiani combined with Pelio, Chapel, and Forestiero renders obvious [wherein the datacenter control system is configured to:] determine an amount of BTM power that is available to the plurality of computing systems.

321. First, as demonstrated with respect to [1.9], Kiani discloses, and combined with Pelio renders obvious, a control system configured to determine an amount of BTM power that is available to the plurality of computing systems. In Kiani, this is accomplished via the FE server, which a POSITA would have viewed as an example of a “*routing control system*.” Further, Kiani teaches that the FE server determines how much “green energy” is available to a datacenter, where “green energy” includes energy produced with the datacenter “behind the meter.” See Ex.1005, pp. 2-3.

322. Second, it would have been obvious to a POSITA for the control to be implemented in a decentralized manner, including by a datacenter control system as taught by Chapel’s local controller, further modified by Forestiero’s decentralized DCM teachings, with Kiani’s FE server fielding user requests and the distribution of those requested jobs occurring according to the determination by the DCMs (examples of datacenter control systems, implemented by Chapel’s local controllers for example) taught by Forestiero. See, e.g., Ex.1009, Figure 2:




Fig. 2. EcoMultiCloud scenario: the DCMs of four data centers exchange high level information about the state of local data centers. Such information is used, for example, to decide which site should accommodate a new VM.

Ex.1009, Fig. 2 (annotated).

323. In combination, the information collected by Forestiero’s DCMs (as implemented via Chapel’s local controller teachings) about the conditions at the datacenter would include the availability of power – “green energy” availability per Kiani’s teachings, which is BTM power as taught by both Kiani and Pelio.

324. Thus, Kiani combined with Pelio, Chapel, and Forestiero renders obvious “[wherein the datacenter control system is configured to:] determine an amount of BTM power that is available to the plurality of computing systems.”

- i. **[13.8] [wherein the datacenter control system is configured to:] determine whether to route the**

***computational operation to the flexible datacenter using
the amount of BTM power that is available; and***

325. Kiani combined with Pelio, Chapel, and Forestiero renders obvious [wherein the datacenter control system is configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available.

326. First, as explained at [1.10], the combination of Kiani and Pelio renders obvious determining to route the computational operation to the flexible datacenter using the amount of BTM power available.

327. Second, it would have been further obvious for the distribution functionality to be implemented at the datacenters in a decentralized manner, instead of via a centralized controller, as further demonstrated and taught by Forestiero, as explained above with respect to [13.6].

328. Thus, the combination of Kiani with Pelio, Chapel, and Forestiero renders obvious “[wherein the datacenter control system is configured to:] determine whether to route the computational operation to the flexible datacenter using the amount of BTM power that is available.”

- j. **[13.9] [wherein the datacenter control system is configured to:] based on a determination to route the computational operation to the flexible datacenter,**

assign the computational operation for performance at the flexible datacenter.

329. Kiani combined with Pelio, Chapel, and Forestiero renders obvious [wherein the datacenter control system is configured to:] based on a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.

330. First, as explained at [1.11], the combination of Kiani and Pelio renders obvious assigning the computational operation for performance at the flexible datacenter based on a determination to route the computational operation there.

331. Second, it would have been further obvious for the distribution functionality to be implemented at the datacenters in a decentralized manner, instead of via a centralized controller, as further demonstrated and taught by Forestiero, as explained above with respect to [13.6]. For example, Forestiero teaches that each DCM (each implemented by a local controller according to Chapel at each of Kiani's datacenters) "execute the algorithms of the upper layer to combine the collected information and take decisions about the distribution of the workload among the DCs," including the assignment of the workload amongst the datacenters. Ex.1009, p. 4.

332. Thus, the combination of Kiani with Pelio, Chapel, and Forestiero renders obvious "[wherein the datacenter control system is configured to:] based on

a determination to route the computational operation to the flexible datacenter, assign the computational operation for performance at the flexible datacenter.”

2. Claim 14


a. [14.0] *The flexible datacenter of claim 13,*

333. See analysis at [13.0]-[13.9].

b. [14.1] *wherein the distributed queue system includes at least two queue subsystems having different rules for organizing computational operations.*

334. Kiani discloses wherein the distributed queue system includes at least two queue subsystems having different rules for organizing computational operations.

335. First, as demonstrated at [1.5], Kiani discloses at least two queue subsystems with Kiani’s description of a queue at each datacenter:



Ex.1005, Fig. 2 (annotated).

Ex.1003

336. Second, Kiani discloses organizing computational operations at each datacenter in each datacenter's queue according to rules – as explained with respect to [1.6]. *See* Ex.1005, pp. 3, 5. Kiani teaches a “**QoS at each data center** [that] is enforced by maintaining an upper bound” on the queue length of that datacenter. Ex.1005, p. 4. This is a rule associated with the queue subsystem at that datacenter, with each datacenter having its own rule for its respective queue subsystem. Further, because it impacts the queue itself, the upper bound is a rule that causes the queue to organize the allocated workload (“*computational operations*”) at each datacenter (in the respective queue) for execution.

337. In an example, Kiani teaches that the “upper bound” for the QoS at each datacenter “is **varied at different data centers** due to different service efficiency parameters.” Ex.1005, p. 5. The upper bound, again, is a limit on the queue length for a given datacenter, which influences the QoS at each datacenter. Ex.1005, p. 4. But the datacenters “have different servers with different service capabilities” and different network delays, resulting in different datacenters handling workload differently from others. Accordingly, Kiani discloses the different datacenters having their own respective queues (“*at least two queue subsystems*”), each having its own “upper bound” (an example of “*different rules for organizing computational operations*” for each datacenter) imposed on queue length based on the unique parameters of each datacenter.

338. And, as further noted at [1.6], Kiani’s variety of rules are determined to facilitate accomplishing the targeted QoS for the workload at each datacenter, and includes for each datacenter:

- An “upper bound” on the queue length at the datacenter;
- A QoS for the datacenter that determines the “upper bound”;
- Maximize green energy utilization for workload allocation at the datacenter;
- Apply brown energy if green energy unavailable/insufficient;
- Minimize the cost of brown energy for workload at the datacenter; and
- Maintaining the QoS requirement for the datacenter for brown energy operations.

339. *See* Ex.1005, p. 4. The queue length of the subqueue at each datacenter, and the other rules for that subqueue, thus influence the organizing of workload (“*computational operations*”) at that datacenter.

340. Thus, Kiani discloses “*wherein the distributed queue system includes at least two queue subsystems having different rules for organizing computational operations.*”

3. Claim 15

a. [15.0] *The flexible datacenter of claim 13,*

341. See analysis at [13.0]-[13.9].

b. [15.1] *wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive received from a remote master control system.*

342. Kiani combined with Pelio and Chapel renders obvious wherein the datacenter control system is configured to modulate power delivery to the plurality of computing systems based on an operational directive received from a remote master control system. *See analysis at [7.1].*

4. Claim 16

a. [16.0] *The flexible datacenter of claim 15,*

343. See analysis at [13.0]-[13.9] and [15.0]-[15.1].

b. [16.1] *wherein the remote master control system is positioned remotely from the flexible datacenter.*


344. Kiani combined with Pelio and Chapel renders obvious wherein the remote master control system is positioned remotely from the flexible datacenter.

345. For example, Kiani teaches a remote master control system (FE server) positioned remotely from the datacenters, which in combination with Pelio are flexible datacenters. *See analysis at [2.1].*

346. As yet another example, Chapel further teaches a local controller (a “*datacenter control system*”) receiving one or more messages from a remote

controller that identifies a power delivery modulation action to take, such as to “transition to an efficiency mode or a brownout mode” (hence teaching an “operational directive”). See Ex.1007, [0076]; *see also* [0058], [0066], [0078]. Figures 3 and 6 of Chapel illustrate the local controller receiving the one or more messages from a remote controller, an example of a “remote master control system.”


347. Figure 3 describes “remote monitoring and/or control of multiple receptacles” which provide power to servers. Ex.1007, [0056]. Receptacles – which connect with servers – “communicate with a local controller 304” that can be at the “customer premises.” Ex.1006, [0056]. The “local controller 304 can, in turn, communicate **with a remote controller 308**” that is operated “by a private or public party” such as an owner of the customer premises, a controller of the power grid, etc. Ex.1007, [0056]; *see also* FIG. 3:



Ex.1007, Fig. 3 (annotated).

348. The local controller communicates “with a remote controller 308 via a wide area network 306 such as the internet.” Ex.1007, [0056]. Via communication of messages from the remote controller to the local controller, Chapel teaches that the “data center equipment can be conveniently controlled from a remote location.” Ex.1007, [0027].

349. Figure 6 describes “control in a data center environment” via remote control. Ex.1007, [0079]-[0080]. A “local controller 650” is in communication with a remote controller 670. Ex.1007, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1007, [0083]; *see also* FIG. 6:



350. In an example, “an operator of a remote controller unit may choose to power down or power up a device or set of devices,” in which case “[a]n appropriate message is transmitted from the remote controller to the local controller,” which then implements the action. Ex.1007, [0090].

351. In combination with Pelio’s and Kiani’s teachings, a POSITA would implement Chapel’s local and remote controllers, and messaging control information between them, to control power consumption at the datacenter according to Chapel’s teachings.

352. Thus, Kiani combined with Pelio, alone and further in combination with Chapel, renders obvious “*wherein the remote master control system is positioned remotely from the flexible datacenter.*”

XIII. CONCLUSION

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

Dated: 8/28/24


Tajana S. Rosing, Ph.D.