

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CHOLLA ENERGY LLC AND OBM, INC.

Petitioners

U.S. Patent No. 11,016,456

**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	5
II.	QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	9
III.	LEVEL OF ORDINARY SKILL IN THE ART	19
IV.	RELEVANT LEGAL STANDARDS	20
V.	BACKGROUND TECHNOLOGY.....	21
A.	Electrical Grid Stability in View of Renewable Generator Supply	21
B.	Using Green Datacenters to Harness Curtailed Power From Renewable Power Sources	23
VI.	OVERVIEW OF THE '456 PATENT	26
VII.	CLAIM CONSTRUCTION	29
VIII.	IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE....	30
IX.	CHALLENGE 1: THE COMBINATION OF PELIO AND CHAPEL RENDERS OBVIOUS CLAIMS 1-4, 14, 17-18, 20, 25 AND 27	30
A.	Summary of Pelio	30
B.	Summary of Chapel.....	32
C.	Reasons to Combine Pelio with Chapel	34
D.	Detailed Analysis of Claims.....	42
1.	Claim 1	42
2.	Claim 2.....	61
3.	Claim 3	65
4.	Claim 4.....	67
5.	Claim 14.....	69
6.	Claim 17	71
7.	Claim 18.....	72
8.	Claim 20.....	74
9.	Claim 25	76

10.	Claim 27	85
X.	CHALLENGE 2: THE COMBINATION OF PELIO, CHAPEL, AND YANG RENDERS OBVIOUS CLAIMS 5-8, 10, 12, AND 28-30	89
A.	Summary of Yang	89
B.	Reasons to Combine Pelio and Chapel with Yang	91
C.	Detailed Analysis of Claims.....	95
1.	Claim 5	95
2.	Claim 6	106
3.	Claim 7	109
4.	Claim 8	112
5.	Claim 10	115
6.	Claim 12	118
7.	Claim 28	119
8.	Claim 29	130
9.	Claim 30	133
XI.	CHALLENGE 3: THE COMBINATION OF PELIO, CHAPEL, AND BELADY RENDERS OBVIOUS CLAIMS 19, 21-24, AND 26.....	134
A.	Summary of Belady	134
B.	Reasons to Combine Pelio and Chapel with Belady	136
C.	Detailed Analysis of Claims.....	140
1.	Claim 19	140
2.	Claim 21	143
3.	Claim 22	148
4.	Claim 23	153
5.	Claim 24	155
6.	Claim 26	157
XII.	CHALLENGE 4: THE COMBINATION OF PELIO, CHAPEL, AND RASMUSSEN RENDERS OBVIOUS CLAIM 9	157

A.	Summary of Rasmussen	157
B.	Reasons to Combine Pelio and Chapel with Rasmussen	160
C.	Detailed Analysis of Claim	162
1.	Claim 9	162
XIII.	CHALLENGE 5: THE COMBINATION OF PELIO, CHAPEL, AND CORHODZIC RENDERS OBVIOUS CLAIMS 15-16	165
A.	Summary of Corhodzic	165
B.	Reasons to Combine Pelio and Chapel with Corhodzic	166
C.	Detailed Analysis of Claims.....	169
1.	Claim 15	169
2.	Claim 16.....	170
XIV.	CONCLUSION.....	173

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. 11,016,456 (“the ’456 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 1-10, 12, and 14-30 (“the Challenged Claims”) of the ’456 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 ’456 patent, Ex.1001;
- b. the prosecution history of the ’456 patent (“’456 File History”), Ex.1002;
- c. U.S. Patent Pub. No. 2011/0316337 to Pelio et al. (“Pelio”), Ex.1005;

- d. U.S. Patent Pub. No. 2018/0041071 to Chapel et al. (“Chapel”), Ex.1006;
- e. 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.1007; and
- f. U.S. Patent No. 8,849,469 to Belady et al. (“Belady”), Ex.1008.
- g. Rasmussen, “Impact of Leading Power Factor on Data Center Generator Systems,” Schneider Electric, 2014 (“Rasmussen”), Ex.1010; and
- h. U.S. Patent No. 8,080,900 to Corhodzic et al. (“Corhodzic”), Ex.1011.

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. Martinot, E., “Grid Integration of Renewable Energy: Flexibility, Innovation, and Experience,” *Annu. Rev. Environ. Resour.*, vol. 41:223-251, 2016 (“Martinot”), Ex.1012;
- b. 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.1013;

c. 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.1014;

d. 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;

e. 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;

f. 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;

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

h. 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.1019;

i. U.S. Patent Publication No. 2010/0088261 to Montalvo (“Montalvo”), Ex.1020;

j. IEEE Xplore, <https://ieeexplore.ieee.org/document/7516102>, Ex.1021;

k. Wayback Machine capture of www.apc.com white papers index, https://web.archive.org/web/20141222215407/http://www.apc.com/prod_docs/results.cfm?DocType=White+Paper&Query_Type=10, Ex.1022;

l. Zhabelova et al., “Data Center Power Dynamics Within the Settings of Regional Power Grid,” 2015 IEEE 20th Conference on Emerging Technologies & Factory Automation (ETFA), September 8-11, 2015, Ex.1023;

m. Metadata from Rasmussen (Ex.1010), Ex.1024; and

n. Schneider Buys APC for \$6.1 Billion, wayback machine archive of October 30, 2006 report, <https://web.archive.org/web/20090606055449/https://www.datacenterknowledge.com/archives/2006/10/30/schneider-buys-apc-for-61-billion>; last accessed April 2, 2024, Ex.1025.

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, which is now Intel 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. Srinivasamurthy, 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 '456 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 ’456 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 '456 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 '456 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 '456 patent, as of its earliest alleged effective filing date of January 11, 2018, would have been someone knowledgeable and familiar with green datacenter efforts and renewable power co-location 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 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 '456 patent.

IV. RELEVANT LEGAL STANDARDS

23. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '456 patent, I am relying on certain basic legal principles that Cisco's counsel has explained to me.

24. I understand that prior art to the '456 patent includes patents and printed publications in the relevant art that predate the priority date of the '456 patent. For purposes of this Declaration, I am applying January 11, 2018, as the priority date of the '456 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.1020, [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.1012, 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.1012, 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.1012, p. 15. These proactive measures involve managing the amount of power arriving at the grid from various generation sources to maintain a balanced system. Ex.1012, 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.1012, p. 15.

29. During curtailment, certain power sources may be temporarily disconnected from the grid or operated at reduced capacities. Ex.1020, [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.1012, 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.1012, p. 15. Such curtailment in response to market conditions is also referred to as economic curtailment. *See, e.g.*, Ex.1007, 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.1014, p. 1. Co-locating at the source reduces or eliminates the costly transmission infrastructure and associated losses. Ex.1014, 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.1014, p. 6.

33. Containers were a well-known solution for providing flexible co-location ability to power sources. *See, e.g.*, Ex.1005, Abstract, [0024], [0048]-[0050], [0066] (flexible, modular datacenter consisting of one or more datacenter containers); Ex.1007, 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.1013, p. 7 (“using containers and collocating at generation sites and avoiding power distribution costs”); Ex.1014, p. 6 (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.1014, 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.1013, pp. 1, 2. 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.1017, 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.1007, 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.1006, [0020], [0037], [0049], [0050]. A controller was used to “control[] delivery of power to” the computers. Ex.1006, [0019], [0027]-[0028].

VI. OVERVIEW OF THE '456 PATENT

37. The '456 patent focuses on “a method and system for dynamic power delivery to a flexible datacenter us[ing] unutilized behind-the-meter power sources without transmission and distribution costs.” Ex.1001, 4:17-20. The '456 patent

explained that the flexible datacenter is “configured to modulate power delivery to one or more computing systems based on the availability of unutilized behind-the-meter power or an operational directive.” Ex.1001, 4:21-24.

38. The “flexible datacenter” includes “a mobile container 205, a behind-the-meter power input system 210, a power distribution system 215, a climate control system ..., a datacenter control system 220, and a plurality of computing systems 100.” Ex.1001, 6:15-22; *see also* FIG. 2:



39. The flexible datacenter “may be a storage trailer disposed on wheels,” a “storage container,” “inflatable container, a floating container, or any other type of container suitable for housing a mobile datacenter.” Ex.1001, 6:30-38.

40. The datacenter control system is “configured to dynamically modulate power delivery to one or more computing systems” within the flexible datacenter “based on unutilized behind-the-meter power availability or an operational directive” from a control system. Ex.1001, 6:22-29. Modulating includes an action to “dynamically reduce or withdraw power delivery to one or more computing systems,” limit the number of powered computing systems to a smaller subset, or reduce the operating frequency or power mode of one or more of the computing systems. Ex.1001, 10:37-47. Modulating also includes ramping up the computing systems to a desired power consumption level. *See, e.g.*, Ex.1001, 9:19-42; 12:17-27.

41. The '456 patent relies upon one or more of the datacenter control system, a remote master control system, and/or a local generator control system in managing the dynamic use of behind-the-meter power for datacenter operations. The datacenter control system is used to “modulate power delivery by either ramping-up flexible datacenter 200 to fully operational status, ramping-down flexible datacenter 200 to offline status ..., reducing power consumption by withdrawing power delivery from, or reducing power to, one or more computing systems 100 or groups

240 of computing systems 100, or modulating a power factor correction factor for the local station by controllably adjusting which phases of three-phase nominal AC voltage are used” by the computing systems. Ex.1001, 7:7-22. The datacenter control system “may independently, or cooperatively with one or more of local station control system 410, remote master control system 420, and grid operator 440” perform the power modulation. Ex.1001, 8:41-49.

42. As I will explain below, these concepts were well known as of the priority date of the ’456 patent.

VII. CLAIM CONSTRUCTION

43. It is my understanding that in order to properly evaluate the ’456 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 ’456 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

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

45. 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.

46. 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 PELIO AND CHAPEL RENDERS OBVIOUS CLAIMS 1-4, 14, 17-18, 20, 25 AND 27

A. Summary of Pelio

47. U.S. Patent Pub. No. 2011/0316337 to Pelio et al. (Ex.1005, “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 ’456 patent.

48. Like the ’456 patent, 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.1005, 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.1005, [0002]. Power generation sources include, for example, “solar photo-voltaics, wind, wave, and tidal.” Ex.1005, [0033]. Figure 1A provides an illustrative example with the datacenter 160 “collocated with and powered by a wind generator farm 162” (Ex.1005, [0044]):



Ex.1005, Fig. 1A (annotated).

49. The modular datacenter 160 includes one or more “data center containers” that come in “[v]arious sizes ... e.g., ISO-Standard 20, 40, and 53-foot

length types.” Ex.1005, [0050]. For example, a “standard-size shipping container” datacenter can be “home to 1,000 or more servers each.” Ex.1005, [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.1005, [0051]. Each datacenter container distributes the power thus input via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1005, [0079].

50. 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.1005, [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.1005, [0106].


B. Summary of Chapel

51. U.S. Patent Pub. No. 2010/0145542 to Chapel et al. (Ex.1006, “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

“Premises Power Usage Monitoring System.” Chapel was not before the Examiner during prosecution of the ’456 patent.

52. 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.1006, [0002]. Intelligent local devices are used to “control power delivered to an electrical device via a circuit,” including electrical devices in a datacenter. Ex.1006, [0047].

53. 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.1006, [0056]. Such customer premises include a “data center environment,” which includes “data devices” such as servers and the like. Ex.1006, [0080]. The servers are plugged into the receptacles, which are in communication with a local controller. Ex.1006, [0082]. The local controller is further in communication with a remote controller. Ex.1006, [0083]; FIG. 6:



Ex.1006, Fig. 6.

54. 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 “an appropriate message ... to the local controller,” which in turn conveys the message to the appropriate receptacle. Ex.1006, [0090].

C. Reasons to Combine Pelio with Chapel

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

56. First, both Pelio and Chapel are analogous art to the '456 patent because they are in the same field of endeavor and reasonably pertinent to the problems being solved by the '456 patent. The Challenged Claims of the '456 patent focus on a “flexible datacenter” with a behind-the-meter power input system, a datacenter control system, and multiple computing systems, with the power delivery

to the computing systems being modulated based on unutilized behind-the-meter power availability or an operational directive. Ex.1001, 19:5-16 (claim 1). The operational directive can be based on a variety of conditions, including those relating to behind-the-meter power availability. Ex.1001, 19:17-20:46 (dependent claims). Thus, the '456 patent is in the field of power delivery to computers in a datacenter.

57. Pelio is in the same field of endeavor, namely power delivery to computers in a datacenter. Pelio's approach is to co-locate datacenter containers at a power generating station in order to harness the station's output behind-the-meter, before any step-up transformations for transmission and distribution on a grid. Ex.1005, Abstract, [0002]. Further, Pelio's power generating station examples include various renewable types that have variable output (due to the different cycles of wind and sunshine, for example). Ex.1005, [0037]. Finally, 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].

58. Pelio is also pertinent to problems identified by the '456 patent. According to the '456 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, 18:25-31. The '456 patent had to also address

the problem of controlling power at the datacenter based upon how much power was available at the time from the energy source. *See* Ex.1001, 4:24-27, 18:45-50. The '456 patent also addressed the problem of cost imposed by transmission and distribution of power over a grid, by co-locating the datacenter with an energy source. *See* Ex.1001, 18:51-57. 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.1005, Abstract, [0002], [0019]. Pelio also teaches using renewable power only when it is available for compute purposes at the datacenter. Ex.1005, [0106]. Thus, a POSITA would have found Pelio's teachings pertinent to the problems identified in the '456 patent.

59. Chapel is also in the same field of endeavor of power delivery to computers in a datacenter. 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.1006, [0002], [0047]. Chapel teaches controlling the power delivery for multiple reasons, including, for example, when power supply is insufficient (either actual or projected). Ex.1006, [0048].

60. Chapel is also pertinent. 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.1006, [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.1006, [0090]. Thus, a POSITA would have found Chapel’s teachings pertinent to the problems identified in the ’456 patent.

61. Second, Pelio describes ways in which datacenter containers obtain behind-the-meter power from renewable generating sources, avoiding the costs of transmission and distribution as well as the losses associated therewith. Ex.1005, 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.1005, [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.1005, [0079].

62. 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.1006, [0027]; [0082]. Power distribution units were a well-known aspect of datacenter design, including datacenter containers as taught in Pelio. Indeed, a POSITA would have understood that Pelio's behind-the-meter renewable power sources were well-known for powering datacenters well before the earliest effective filing date of the '456 patent. *See, e.g.*, Ex.1005, Abstract; Ex.1007, pp. 1, 4; Ex.1013, p. 7; Ex.1014, p. 6; Ex.1015, pp. 3, 4.

63. Third, Pelio recognizes the variability of renewable power sources, cycling “between full operation” and as far as “zero” output due to lack of the renewable resource (e.g., wind, solar, waves, etc.). Ex.1005, [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.1005, [0106]. Pelio's teachings are more focused on placing datacenter containers behind-the-meter to harness the power output before transmission costs and losses arise, as well as providing equipment configuration flexibility, than on the specifics of how to handle the variability of the power source itself. Ex.1005, [0002]; [0070]. Thus, a POSITA would have been further motivated to turn to Chapel for Chapel's power management teachings.

64. 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.1006, [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.1006, [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.1006, [0090].

65. A POSITA would have had a reasonable expectation of success in combining the teachings of Chapel with Pelio. Chapel's teachings are compatible with Pelio's teachings. Pelio teaches providing power to servers within a datacenter container (or multiple such containers) from a variable power source behind-the-meter. Ex.1005, 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.1006, [0002]; [0047]; [0080]-[0083]. Chapel's teachings are compatible with Pelio's teachings, since 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 Pelio's variable power supply from the behind-the-meter power source. *See id.*

66. Combining Chapel’s teachings with 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 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.1007, pp. 1, 4; *see also* Ex.1013, p. 7 (“using containers and collocating at generation sites and avoiding power distribution costs”); Ex.1014, p. 6 (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”).

67. 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.1007, 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 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.1006, [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 Pelio's recognition, and enable Pelio's datacenter containers to follow the power supply as it varies.

68. The combination of Pelio's and Chapel's teachings would predictably result in Pelio's datacenter containers including local controllers to control power delivered to the server devices, using the teachings adapted from Chapel. *See* Ex.1005, Abstract; [0002]; [0033]; [0037]; [0050]; [0106] (locating containers behind-the-meter with potentially varying power supply); Ex.1006, [0002]; [0047]; [0080]-[0083] (monitoring and controlling power delivery). Doing so would advantageously enable Pelio'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.).

69. Indeed, the combination of Pelio and Chapel is nothing more than the combination of known elements (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 Pelio's datacenters), to yield predictable results (Pelio'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).

70. The combination permits but does not require physical incorporation of elements from Chapel into Pelio. Nevertheless, the physical incorporation of Chapel's power monitoring and power delivery control teachings, including its controller teachings, into Pelio's datacenter teachings would teach the claim limitations.


D. Detailed Analysis of Claims

1. Claim 1

a. [1.0] *A flexible datacenter comprising:*

71. To the extent that the preamble is limiting, Pelio discloses it in several ways.


72. First, Pelio discloses "data center containers" which are individually examples of a "*flexible datacenter*." Ex.1005, 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.1005, [0049]; *see also* FIG. 2:



Ex.1005, Fig. 2 (annotated).

73. These datacenter containers come in “[v]arious sizes ... e.g., ISO-Standard 20, 40, and 53-foot length types.” Ex.1005, [0050]. Examples include datacenter containers that are “factory-built offsite, e.g., by IBM, HP, SGI/Rackable, etc.” Ex.1005, [0066]. “Such containers provide tremendous equipment configuration flexibility in the overall design.” Ex.1005, [0070]. For example, a “standard-size shipping container” datacenter can be “home to 1,000 or more servers each.” Ex.1005, [0071].

74. “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.1005, [0095]. Figure 4 illustrates these flexible datacenters as well:




75. Figure 4 illustrates “[a] number of data center containers 410-412 on a floor 414.” Ex.1005, [0097].

76. Such flexible datacenter containers, which “can be moved between trucks,” including standard-sized shipping containers, are the very type of flexible datacenters described by the ’456 patent. *See* Ex.1005, [0071], [0089]. For example, the ’456 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, 6:30-38.

77. **Second**, Pelio’s “modular datacenter” itself, which “comprises a building shell with a modular common power and cooling infrastructure to support data center containers” is also an example of a “flexible datacenter.” Ex.1005, Abstract. Pelio explicitly teaches that Figure 2 illustrates “a flexible, just-in-time, modular data center” that houses multiple containers (the other example of a

“flexible datacenter” as explained above). Ex.1005, [0024], [0048]. The modular datacenter’s interior is “partitioned into floor spaces ... designated for individual ownership and support of critical loads measurable in watts.” Ex.1005, [0048]; *see also* FIG. 2:



Ex.1005, Fig. 2 (annotated).

78. The various containers, of various sizes, are “floated-in through large roll-up back doors” or alternatively moved via a “standard crane system.” Ex.1005, [0050]. Pelio’s modular datacenter is thus able to “incrementally buy additional critical load ... container-by-container.” Ex.1005, [0055]. The result allows tenants (datacenter users) to “buy-in and/or operate a tenancy of one or more data center containers 320 of theirs to be placed in modular data center 314.” Ex.1005, [0066].

79. As a result, “[e]ach new data center container 320 is installed independent of any others” and “can be easily replaced anytime in the future.” Ex.1005, [0067], [0068]. “Such containers provide tremendous equipment

configuration flexibility in the overall design” with the containers within housing servers in cabinets/racks. Ex.1005, [0071]. The modular datacenter itself is, therefore, also a “*flexible datacenter*.”

80. Thus, Pelio discloses a “*flexible datacenter*.”

- b. [1.1] *a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;*

81. Pelio discloses a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid.


82. First, Pelio discloses that the power received by the modular datacenters/containers within the modular datacenters from one or more power generation stations are “*behind-the-meter*.” 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.1005, [0002].

83. 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.1005, [0019]. Thus, Pelio’s collocated 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 involving power losses due to heating.” Ex.1005, [0039].

84. Further, the power input to the datacenter containers themselves within the modular datacenter also are “*behind-the-meter*” – the containers receive power from the power input to the modular datacenter which is, itself, “*behind-the-meter*.” For example, Pelio discloses that “[a] number of modular power supply systems 240-[2]43 are disposed behind the common building shell and provide for the operating power predetermined for the particular modular containerized data centers.” Ex.1005, [0051]. The power supply systems can be “utility sub-panels and transformers,” or UPS systems or generators. Ex.1005, [0051]. Alternatively, **“[e]ach data center container may ... obtain its power directly from the utility** via a switchgear, transformer and conduit system.” Ex.1005, [0051].

85. Pelio discloses a “*power input system*” that receives power “*behind-the-meter*.” The “modular power supply systems 240-[2]43” are examples of a “*power input system*,” as is the utility sub-panel 245:



Ex.1005, Fig. 2 (annotated).

86. Second, Pelio provides several examples of “*behind-the-meter*” “power generation stations.” “Generating stations now include coal-fired, oil-fired, natural gas, hydro-electric, nuclear, wind farm, and solar energy types.” Ex.1005, [0005].

87. In Figure 1A’s example, “a modular data center 160 is shown collocated with and powered by a wind generator farm 162,” prior to the power being stepped up by “step-up transformer 164.” Ex.1005, [0044]; Figure 1A:






Fig. 1A
 Flexible datacenter
 Ex.1005, Fig. 1A (annotated).
 Power generation station

88. In Figure 1B’s example, “a modular data center 170 is shown collocated with and powered by solar electric generating station 150,” prior to the power being stepped up by “step-up transformer 152.” Ex.1005, [0045]; Figure 1B:



89. In Figure 1C’s example, “a modular data center 180 is shown collocated with and powered by hydro electric generating station 106,” prior to the power being stepped up by “step-up transformer 112.” Ex.1005, [0046]; Figure 1C:



90. “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.1005, [0047].

91. For each example, in view of the fact that Pelio’s modular datacenter receives the power before step-up (hence, “*behind-the-meter*”), 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.

92. Thus, Pelio discloses “*a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid.*”

c. [1.2] *a power distribution system;*


93. Pelio renders obvious, alone and in combination with Chapel, a power distribution system.


94. First, Pelio renders obvious each datacenter container including a “*power distribution system.*” Pelio explains the obvious fact that “electrical power [is] consumed by the servers and data storage units” within a datacenter container, and of course that each datacenter container can be up to “1,000 or more servers each.” Ex.1005, [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 (as explained at Ex.1005, [0051]), would be distributed to each server in a container via a “*power distribution system.*” This is fundamentally well known. *See, e.g.,* Ex.1006, [0027]; [0082].

95. For example, Pelio teaches that the power distribution to servers in a datacenter container would be via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1005, [0079]. Such power feeds are well-known to be “independent electrical feeds coming into a data center from the utility company.” Ex.1019, p. 1.

Each feed includes “[e]lectrical and distribution switchboards,” a “power distribution unit (PDU),” and “Rack-level PDUs.” Ex.1019, p. 1; *see also* Ex.1006, [0094] (datacenter using a “power distribution unit ... associated with multiple power sources 660 such as an A source and a B source”).

96. Second, Pelio discloses a “*power distribution system*” for the modular datacenter (shell) itself with the “raceway” disclosures. For example, Pelio explains that “[a]n overhead raceway 440 provides ... electrical power connections to uninterruptible power supply (UPS) unit 442 and utility power.” Ex.1005, [0099]. These “raceways” are illustrated in both figures 2 and 4:





Ex.1005, Fig. 4 (annotated).

97. Third, Chapel confirms that a datacenter, like the datacenter containers in Pelio, have a “*power distribution system*” to distribute the “*behind-the-meter*” power from sources such as those identified in Pelio.

98. For example, Chapel discloses “a power distribution unit of a data center.” Ex.1006, [0027]; *see also* [0082]. The power distribution unit outputs power received at the datacenter from one or more power sources to the servers within a datacenter. Ex.1006, [0082]. In some examples, the “power distribution unit” of a datacenter “may be associated with multiple power sources ... such as an A source and a B source.” Ex.1006, [0082]; *see also* FIG. 6:




FIG. 6
Ex.1006, Fig. 6 (annotated).

99. Thus, Pelio renders obvious, alone and in combination with Chapel, “a power distribution system.”

d. [1.3] a plurality of computing systems; and

100. Pelio discloses a plurality of computing systems.

101. Pelio discloses that the datacenter containers include inside them “network servers” that are mounted in racks and organized into rows. Ex.1005, [0071]. For example, “standard-size shipping containers” include those that are “home to 1,000 or more servers each.” Ex.1005, [0071]. This includes datacenter containers that “can be ordered from their manufacturers fully provisioned with the servers already installed.” Ex.1005, [0048]. These servers, and generally the datacenters, are “computing systems.” Ex.1005, [0008].

102. Thus, Pelio discloses “a plurality of computing systems.”

- e. **[1.4] *a datacenter control system configured to modulate power delivery to the plurality of computing systems based on an operational directive.***

103. Pelio combined with Chapel renders obvious a datacenter control system configured to modulate power delivery to the plurality of computing systems based on an operational directive.

104. First, Pelio teaches the desirability of placing datacenters near a power generating 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.1005, [0002]. As a result of doing so, “a higher percentage of the power produced by the generation source actually reaches critical load” of the datacenter. Ex.1005, [0019].

105. Pelio discloses several different types of power generating sources, including “coal-fired, oil-fired, natural gas, hydro-electric, nuclear, wind farm, and solar energy types.” Ex.1005, [0005]. Pelio 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.1005, [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. 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.1005, [0037]; *see also* [0043]; [0105]-[0106].

106. Since Pelio’s main focus is on aspects relating to the design of a modular datacenter that is collocated with a primary electric generating station (Ex.1005, Abstract), Pelio leaves it up to a POSITA as to all the various ways that power may be controlled at a datacenter. 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 Pelio (e.g., solar, wind, etc.).

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

108. Chapel relates to “electrical power distribution and management,” including how to “intelligently monitor at least a portion of the circuit and control

delivery over the circuit.” Ex.1006, [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.1006, [0019], [0047]. An example system enables “remote monitoring and/or control of multiple receptacles” which provide power to servers. Ex.1006, [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.1006, [0056]; *see also* FIG. 3:




FIG.3

Ex.1006, Fig. 3 (annotated).

109. The “local controller” is an example of a “*datacenter control system*.”

This is because it is a control system at the datacenter. Ex.1006, [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.1006, [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.1006, [0025]-[0026].

110. Chapel’s Figure 6 provides a specific example of “control in a data center environment,” with datacenter devices (e.g., servers) that are plugged into controllable receptacles. Ex.1006, [0080]-[0081]. These are connected to “a local controller 650” that is itself in communication with a remote controller 670. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083]. In an example, “an operator of a remote controller unit may choose to power down or power up a 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.1006, [0090]; *see also* FIG. 6:

Ex.1006, [0017], [0025]-[0026], [0076]. For example, a POSITA would have understood that Pelio's flexible datacenters, when collocated with a power generation station 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.

113. Thus, the combination renders obvious "*a datacenter control system configured to modulate power delivery to the plurality of computing systems based on an operational directive.*"

2. Claim 2

a. [2.0] *The flexible datacenter of claim 1,*

114. See analysis at [1.0]-[1.4].


b. [2.1] *wherein the datacenter control system is configured to receive the operational directive from a remote master control system.*

115. Pelio combined with Chapel renders obvious wherein the datacenter control system is configured to receive the operational directive from a remote master control system.

116. As explained with respect to [1.4], Chapel 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.1006, [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.”

117. Figure 3 describes “remote monitoring and/or control of multiple receptacles” which provide power to servers. Ex.1006, [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.1006, [0056]; *see also* FIG. 3:



Ex.1006, Fig. 3 (annotated).

118. The local controller communicates “with a remote controller 308 via a wide area network 306 such as the internet.” Ex.1006, [0056]. Via communication of messages from the remote controller to the local controller, Chapel teaches that the “data center equipment can be controlled from a remote location.” Ex.1006, [0027].

119. Figure 6 describes “control in a data center environment” via remote control. Ex.1006, [0079]-[0080]. A “local controller 650” is in communication with a remote controller 670. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083]; *see also* FIG. 6:

3. Claim 3

a. [3.0] *The flexible datacenter of claim 2,*

123. See analysis at [1.0]-[1.4] and [2.0]-[2.1].

b. [3.1] *wherein the operational directive depends on BTM power availability at the flexible datacenter.*

124. Pelio combined with Chapel renders obvious wherein the operational directive depends on BTM power availability at the flexible datacenter.

125. Chapel teaches a remote master control system sending operational directives in the form of the remote controller sending messages to a datacenter's local controller for controlling one or more datacenter equipment power usage. See [2.1] analysis; Ex.1006, [0027], [0056] (Figure 3), [0079]-[0083], [0090] (Figure 6). Chapel teaches sending operational directives for several reasons, including to address actual or potential "over-capacity conditions." Ex.1006, [0026]. For example, "the over-capacity condition relates to a condition potentially requiring reduction of power provided to standard ... commercial customers, e.g., conditions that have conventionally resulted in blackouts or brownouts." Ex.1006, [0026]. The conditions result from various circumstances, including "increasing industrial and residential power needs," "aging power infrastructure," and "practical limitations on new power generation." Ex.1006, [0004].

126. It would have been obvious to a POSITA, based on the combined teachings of Pelio and Chapel, that "new power generation" includes renewable

power sources such as those taught in Pelio. These include, for example, “hydro-electric ... wind farm, and solar energy types.” Ex.1005, [0005]. Moreover, Chapel’s teachings are applicable to all of the energy generator types beyond renewable, since over-capacity conditions may arise due to higher demand regardless of the generator source. *See* Ex.1005, [0005] and [0033] (e.g., coal-fired, oil-fired, natural gas, hydro-electric, nuclear, wind farm, solar energy, geothermal, wave, tidal).

127. In combination, a POSITA would have combined Chapel’s teachings to have a remote controller cause a local controller to modulate power delivery to datacenter equipment (including servers) at a datacenter with Pelio’s BTM-located generating source, according to actual or potential over-capacity conditions including those that would arise from periodic BTM-located power sources such as solar, wind, tidal, etc. as taught by Pelio (where cycles result in periods of full operation and less or zero periods of operation depending on energy availability). *See* Ex.1005, [0019], [0033], [0037]; Ex.1006, [0028], [0088]. A POSITA would have done so in order to take advantage of Chapel’s ability to control electricity delivery for more efficient distribution. *See* Ex.1006, [0007], [0047]. Further, a POSITA would have applied Chapel’s teachings to modulate power delivery at Pelio’s flexible datacenter in situations where over-capacity actually exists or is predicted to exist with respect to the BTM generator source, such as where the over-capacity condition is caused by higher demand at the datacenter.

128. The combined teachings of Pelio and Chapel result in the remote controller sending the message (“*operational directive*”) to the data controller at the datacenter to modulate power delivery based upon how much power is available relative to demand at the datacenter (“*depends on power availability at the flexible datacenter*”).

129. Thus, Pelio combined with Chapel renders obvious “*wherein the operational directive depends on BTM power availability at the flexible datacenter.*”

4. Claim 4

a. [4.0] *The flexible datacenter of claim 3,*

130. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [4.1] *wherein the remote master control system is positioned remotely from the flexible datacenter.*


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

132. As demonstrated at [1.0], Pelio teaches a flexible datacenter. Ex.1005, Abstract, [0024], [0048], [0049].


133. As further demonstrated at [2.1], Pelio combined with Chapel renders obvious a remote controller that provides messages of instruction to the datacenter for controlling power delivery. Ex.1006, [0027], [0076].

134. As demonstrated at [2.1], Chapel teaches locating the remote controller (a “*remote master control system*”) remotely from the datacenter. Thus, Chapel’s

teaching combined with Pelio's flexible datacenter teachings would result in Pelio's flexible datacenter being located remotely from a remote controller for the system. Indeed, Chapel teaches locating the remote controller away from the datacenter premises to "remotely control" the power-consuming devices. *See* Ex.1006, [0056], [0083]; *see also* FIGs. 3, 6:



Ex.1006, Fig. 3 (annotated).



135. Thus, Pelio combined with Chapel renders obvious “*wherein the remote master control system is positioned remotely from the flexible datacenter.*”

5. Claim 14

a. [14.0] *The flexible datacenter of claim 1,*

136. See analysis at [1.0]-[1.4].

b. [14.1] *wherein the datacenter control system is further configured to modulate power delivery to the plurality computing systems based on BTM power availability at the BTM power input system.*

137. Pelio combined with Chapel renders obvious wherein the datacenter control system is further configured to modulate power delivery to the plurality computing systems based on BTM power availability at the BTM power input system.

138. First, as explained at [1.1], Pelio discloses BTM power generation with power that is input to a BTM power input system at the datacenter. *See, e.g.*, Ex.1005, [0002], [0019].

139. Second, as further explained at [1.4], the combination of Chapel and Pelio renders obvious a “*datacenter control system*” based on Chapel’s teachings of a “local controller” that can be at the customer premises. *See* Ex.1006, [0056].

140. Third, Chapel further teaches monitoring power availability (and, in combination with Pelio, renders obvious monitoring BTM power availability in particular, as demonstrated at [28.1] *infra*). *See* Ex.1006, [0004], [0047], [0056], [0083]. Moreover, a POSITA would have understood that Chapel’s load control involves not only controlling delivery of power by turning off devices, operating the devices in reduced power mode, or interrupting power distribution (Ex.1006, [0025]-[0026]), but also by turning the devices back on, ending reduced power mode, removing power interruption, etc. *See, e.g.*, Ex.1006, [0068] (turning devices on). And, since Chapel teaches the remote controller performing the power delivery control in situations where supply is expected to be less than demand (the load), it would have been obvious to a POSITA for the remote controller to also cause power supply to the load to increase when the supply becomes sufficient again (i.e., no over-capacity condition threat anymore). *See also* [28.3].

141. Thus, Pelio combined with Chapel renders obvious “*wherein the datacenter control system is further configured to modulate power delivery to the plurality computing systems based on BTM power availability at the BTM power input system.*”

6. Claim 17

a. [17.0] *The flexible datacenter of claim 1,*

142. See analysis at [1.0]-[1.4].

b. [17.1] *wherein the datacenter control system controllably enables or disables power delivery to respective computing systems of the plurality of computing systems.*

143. Pelio combined with Chapel renders obvious wherein the datacenter control system controllably enables or disables power delivery to respective computing systems of the plurality of computing systems.

144. As demonstrated with respect to [1.4], Chapel teaches that the local controller at the customer premises is used to “control[] delivery of power to the electrical device.” Ex.1006, [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.1006, [0025]-[0026].

145. Chapel’s Figure 6 provides a specific example of “control in a data center environment,” with datacenter devices (e.g., servers) that are plugged into controllable receptacles. Ex.1006, [0080]-[0081]. These are connected to “a local

controller 650” that is itself in communication with a remote controller 670. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083]. In an example, “an operator of a remote controller unit may choose to power down or power up ... a 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.1006, [0090].

146. In combination, this would include the servers from Pelio’s teachings being powered down or up by controlling power delivery to one or more servers, according to Chapel’s teachings.

147. Thus, Pelio combined with Chapel renders obvious “*wherein the datacenter control system controllably enables or disables power delivery to respective computing systems of the plurality of computing systems.*”

7. Claim 18

a. [18.0] *The flexible datacenter of claim 1,*

148. See analysis at [1.0]-[1.4].

b. [18.1] *[w]herein the datacenter control system is positioned remotely from the plurality of computing systems.*

149. Pelio combined with Chapel renders obvious [w]herein the datacenter control system is positioned remotely from the plurality of computing systems.

150. As demonstrated with respect to [1.4], Chapel teaches that receptacles – which connect with servers (the “*plurality of computing systems*”) – “communicate with a local controller 304” that can be at the “customer premises.” Ex.1006, [0056]. Implementing such controller systems separately from the servers performing datacenter operations was a well-known option to POSITAs. Indeed, it was known for the controller to be “implemented by processing equipment (e.g., one or more computer servers, data stores, etc.) which is distinct from the processing equipment provided by the data centers themselves.” Ex.1008, 8:15-20. This was a well-known alternative to implementing the control systems by the same processing equipment provided by the datacenters. *See, e.g.,* Ex.1008, 8:23-26. Thus, Chapel’s local controller implemented distinctly from the servers in a datacenter renders obvious the “*datacenter control system*” being “*positioned remotely from the plurality of computing systems.*”

151. Chapel’s teachings, in view of the scope of the “*datacenter control system,*” renders obvious claim 18 in yet another way. With respect to Figure 3’s example in Chapel, 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.1006, [0056]. Chapel further describes with respect to Figure 6 datacenter devices (e.g., servers) that are plugged into controllable receptacles. Ex.1006, [0080]-[0081]. These are

connected to “a local controller 650” that is itself in communication with a **remote controller 670**. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083]. In an example, “an operator of a remote controller unit may choose to power down or power up a 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.1006, [0090].

152. The “remote controller” devices taught in Chapel are another example of a “*datacenter control system*” according to the claim, and thus renders [18.1] obvious. This is because, as claim 24 itself acknowledges, there is no distinction between the two – “*the datacenter control system is a remote master control system.*” Thus, a POSITA would have understood the “*datacenter control system*” from claim 1 to be broad enough to encompass local and remote scenarios, which Chapel teaches.

153. Thus, Pelio combined with Chapel renders obvious “[w]herein the *datacenter control system is positioned remotely from the plurality of computing systems.*”

8. Claim 20

a. [20.0] *The flexible datacenter of claim 1,*

154. See analysis at [1.0]-[1.4].

b. [20.1] wherein the datacenter control system is
collocated with the plurality of computing systems.

155. Pelio combined with Chapel renders obvious wherein the datacenter control system is collocated with the plurality of computing systems.

156. As demonstrated with respect to [1.4], Chapel teaches that receptacles – which connect with servers (the “*plurality of computing systems*”) – “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.1006, [0056]. Being “at the customer premises” demonstrates the local controller being “*collocated with the plurality of computing systems*.” See also Ex.1006, FIGs. 3, 6:





FIG.3

Ex.1006, Fig. 3 (annotated).



157. Thus, Pelio combined with Chapel renders obvious “*wherein the datacenter control system is collocated with the plurality of computing systems.*”

9. Claim 25

a. [25.0] A system comprising:


158. To the extent that the preamble is limiting, Pelio discloses a system.

159. As explained at [1.0], Pelio discloses a “modular datacenter,” which “comprises a building shell with a modular common power and cooling infrastructure to support data center containers.” Ex.1005, Abstract. Pelio explicitly teaches that Figure 2 illustrates “a flexible, just-in-time, modular data center” that houses multiple containers. Ex.1005, [0024], [0048]. The modular datacenter’s interior is “partitioned into floor spaces ... designated for individual ownership and support of critical loads measurable in watts.” Ex.1005, [0048].


160. The modular datacenter itself, as it comprises first and second flexible datacenters as well as control elements, discloses a “system” as claimed.

b. [25.1] a first flexible datacenter that comprises:

161. Pelio discloses a first flexible datacenter. See [1.0]; Ex.1005, FIGs. 2, 4:



Ex.1005, Fig. 2 (annotated).



Ex.1005, Fig. 4 (annotated).

162. Therefore, Pelio discloses “a first flexible datacenter.”

- c. **[25.2] *a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;***

163. Pelio discloses that the first flexible datacenter comprises a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid. *See* [1.1].

- d. **[25.3] *a power distribution system; and***

164. Pelio, alone and in combination with Chapel, renders obvious the first flexible datacenter comprising a power distribution system. *See* [1.2].

- e. **[25.4] *a plurality of computing systems configured to receive power from the BTM power input system via the power distribution system;***

165. Pelio, alone and in combination with Chapel, renders obvious a plurality of computing systems configured to receive power from the BTM power input system via the power distribution system.

166. First, Pelio discloses a plurality of computing systems. *See* [1.3].

167. Second, Pelio discloses that the plurality of computing systems receive power from the BTM power input system via the power distribution system. For example, as explained with respect to [1.2] Pelio discloses that the servers in the datacenter receive power distributed via one or both of “the ‘A-side’ and ‘B-side’ power feeds.” Ex.1005, [0079]. These feeds are well-known to include a power

distribution unit and necessary hardware to provide electricity to the servers. *See* Ex.1019, p. 1; Ex.1006, [0082].

168. Since Pelio teaches that the power is provided via a BTM power input system, a POSITA would have understood that the A-side and B-side power distribution system within a given flexible datacenter container provide the BTM power to the servers. Thus, Pelio renders obvious the servers receiving the BTM power received from the BTM power input system, via the A-side and B-side power distribution.

169. In addition, Pelio also teaches that the modular datacenter (shell) itself includes a “raceway” that provides “electrical power connections to ... utility power” for the individual datacenter containers.Ex.1005, [0099]. This further renders obvious the servers receiving power from the BTM power input system (e.g., utility power panel etc.) via the raceway.


170. Third, Chapel confirms that a datacenter, like Pelio’s flexible datacenter, uses a “power distribution system” to distribute the BTM power from the BTM power input. *See* Ex.1006, [0027], [0082]; and [1.2]. In combination, the plurality of computing devices in Pelio’s datacenter receive the power distributed via Pelio’s/Chapel’s power distribution unit as received from the BTM power input system taught in Pelio.

171. Thus, Pelio, alone and in combination with Chapel, renders obvious “a plurality of computing systems configured to receive power from the BTM power input system via the power distribution system.”


f. [25.5] a second flexible datacenter; and

172. Pelio discloses a second flexible datacenter.

173. As explained with respect to [1.0], Pelio discloses multiple “data center containers,” each of which is an example of a different “flexible datacenter.” See Ex.1005, Abstract, [0049]. These are illustrated in Figures 2 and 4 of Pelio as well:



Ex.1005, Fig. 2 (annotated).



Ex.1005, Fig. 4 (annotated).

174. Figure 4 illustrates “[a] number of data center containers 410-412 on a floor 414.” Ex.1005, [0097]. Any second container is, therefore, an example of a “*second flexible datacenter.*”

175. Thus, Pelio discloses “*a second flexible datacenter.*”

g. [25.6] *a routing control system configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter.*

176. Pelio combined with Chapel renders obvious a routing control system configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter.

177. The '456 patent specification does not describe a “*routing control system*” in particular. However, according to dependent claims 26 and 27, the “*routing control system*” “*is configured to transfer a workload from the plurality of computing system to a second plurality of computing systems at the second flexible datacenter*” and “*is positioned remotely from the first flexible datacenter and the*

second flexible datacenter.” This description is consistent with the ’456 patent’s figure 5, which provides a “remote master control system 420” that is positioned remotely from the control systems for the flexible datacenters. *See, e.g.,* Ex.1001, 10:58-11:29; FIG. 5:




FIG. 5

Ex.1001, Fig. 5.

178. Pelio combined with Chapel renders obvious a routing control system, like the remote master control system setup in the ’456 patent’s Figure 5, configured to modulate power delivery to computing devices in flexible datacenters. For example, Chapel teaches using a “remote controller” and a “local controller” relative to the datacenter to control modulation of power delivery in each datacenter. *See* [2.1]. Figure 3 describes “remote monitoring and/or control of multiple receptacles” which provide power to servers. Ex.1006, [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.1006, [0056]; *see also* FIG. 3:





FIG.3

Ex.1006, Fig. 3 (annotated).

179. The local controller communicates “with a remote controller 308 via a wide area network 306 such as the internet.” Ex.1006, [0056]. Chapel teaches that the “data center equipment can be controlled from a remote location” using the messages between remote controller and local controller at the datacenter. Ex.1006, [0027].

180. Figure 6 describes “control in a data center environment” via remote control. Ex.1006, [0079]-[0080]. A “local controller 650” is in communication with

a remote controller 670. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083].



181. In an example, “an operator of a remote controller unit may choose to power down or power up ... a 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.1006, [0090].

182. In combination with Pelio’s teachings, a POSITA would implement Chapel’s local and remote controllers, and messaging control information between them, to control power consumption at each of Pelio’s flexible datacenters according to Chapel’s teachings. Indeed, Chapel itself teaches a remote controller in communication with multiple local controllers at different datacenter locations. A

“grid controller,” an example of a remote controller, “may instruct local controllers ... to operate in conservation mode. Local controllers may then execute a conservation mode of operation.” Ex.1006, [0025]; *see also* [0077] (describing grid controllers as “remote grid controller”).

183. Accordingly, a remote controller (an example of a “*routing control system*”) sends messages to each of the flexible datacenters to control the modulation of power delivery at each datacenter via one or more local controllers at each datacenter, per the combined teachings of Pelio and Chapel. *See* Ex.1005, Abstract (flexible datacenters); Ex.1006, [0018] (modulating power delivery by interrupting power distribution).

184. Thus, Pelio combined with Chapel renders obvious “*a routing control system configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter.*”

10. Claim 27

a. [27.0] *The system of claim 25,*

185. *See* analysis at [25.0]-[25.6].

b. [27.1] *wherein the routing control system is positioned remotely from the first flexible datacenter and the second flexible datacenter, and*

186. Pelio combined with Chapel renders obvious wherein the routing control system is positioned remotely from the first flexible datacenter and the second flexible datacenter. *See* [4.1] and [25.6] (demonstrating that the “routing

control system” relates to the “remote master control system”). For example, the ‘456 patent describes the remote master control system as remote from the flexible datacenters in Figure 5:






FIG. 5

Ex.1001, Fig. 5.

187. Similarly, Chapel teaches locating the remote controller (a “*routing control system*”) remotely from any of the flexible datacenters (e.g., per Pelio’s teachings of multiple flexible datacenters). Thus, Chapel’s teaching combined with Pelio’s flexible datacenter teachings would result in Pelio’s flexible datacenters being located remotely from a remote controller for each system. Indeed, Chapel teaches locating the remote controller away from the datacenter premises to “remotely control” the power-consuming devices. See Ex.1006, [0056], [0083]; see also FIGs. 3, 6:



Ex.1006, Fig. 3 (annotated).



Ex.1006, Fig. 6 (annotated).

- c. [27.2] wherein the routing control system is configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter based on

***BTM power availability at the first flexible datacenter
and the second flexible datacenter.***

188. Pelio combined with Chapel renders obvious wherein the routing control system is configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter based on BTM power availability at the first flexible datacenter and the second flexible datacenter.

189. First, Pelio combined with Chapel renders obvious modulating power delivery to a flexible datacenter. *See* [14.1]. Chapel teaches monitoring power availability. *See* Ex.1006, [0004], [0047], [0056], [0083]. In particular, Chapel teaches that 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.1006, [0056].

190. In another example, datacenter devices (e.g., servers) are plugged into controllable receptacles. Ex.1006, [0080]-[0081]. These are connected to “a local controller 650” that is itself in communication with a remote controller 670. Ex.1006, [0082]. “[R]emote and/or policy based operation of the” datacenter devices is achieved via the local controller 650 and remote controller 670. Ex.1006, [0083]. In an example, “an operator of a remote controller unit may choose to power down or power up a 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.1006, [0090].

191. Thus, as demonstrated further at [1.4], Chapel teaches a remote controller sending messages for a local controller to implement to control power delivery to the computers within a datacenter, such as Pelio’s flexible datacenters.

192. Second, it would have been obvious to a POSITA to repeat the Pelio-Chapel teachings with each flexible datacenter taught by Pelio. This is nothing more than the mere duplication of parts (modulating power delivery per Chapel’s teachings for each flexible datacenter taught by Pelio).

193. Thus, Pelio combined with Chapel renders obvious “*wherein the routing control system is configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter based on BTM power availability at the first flexible datacenter and the second flexible datacenter.*”

X. CHALLENGE 2: THE COMBINATION OF PELIO, CHAPEL, AND YANG RENDERS OBVIOUS CLAIMS 5-8, 10, 12, AND 28-30

A. Summary of Yang

194. Yang is an IEEE publication titled “ZCCloud: Exploring Wasted Green Power for High-Performance Computing.” Yang was submitted as part of the 2016 IEEE International Parallel and Distributed Processing Symposium. Ex.1007, p. 1. The symposium was held May 23-27, 2016. *Id.*; *see also* Ex.1021. Yang was added to the IEEE Xplore digital, online library on July 21, 2016. In my experience,

publications for IEEE symposia are distributed without restriction to attendees at the symposium, and are publicly available as of the date of the symposium/conference. 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.

195. I have reviewed Yang, and it bears the typical indicia of an IEEE publication, including the title of the Symposium at the top, the copyright date at the bottom corner (2016), and document identifiers and cost for purchase. Ex.1007, p. 1. POSITAs were well-versed in accessing documents of interest from IEEE Xplore (if not received at the conference itself), and therefore it is my opinion that a POSITA would have been able to locate Yang with reasonable diligence. Yang was not before the Examiner during prosecution of the '456 patent.

196. Yang describes approaches for co-locating datacenters with renewable generating sources. Ex.1007, p. 1. Recognizing that renewable power generation can often be curtailed due to economic or reliability concerns, Yang proposes powering the computing at container datacenters “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3. By using containers, building costs are eliminated; and by co-locating at the generating source, transmission is eliminated. Ex.1007, pp. 1, 4. As a result, the computing resources “are activated

only when stranded power is available, and will be shutdown whenever there is insufficient power.” Ex.1007, p. 1.

B. Reasons to Combine Pelio and Chapel with Yang

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

198. First, Yang is focused on similar, analogous issues to those in Pelio, including power delivery to computers in a datacenter. Yang explains that its “computing resources are deployed on generation sites and placed in containers” with power received without any need for transmission. Ex.1007, pp. 1, 4. Thus, just like Pelio and the ’456 patent, Yang co-locates its container-based computing resources at power generating stations to harness the power output behind-the-meter. Further, the computing resources in Yang are only powered when there is excess, “stranded power” (power that the grid does not want or cannot handle at the time). Ex.1007, pp. 1, 2-3. Thus, Yang is in the same field of endeavor as the ’456 patent.

199. Further, Yang is also pertinent to the problems identified by the ’456 patent, such as what to do with “unutilized and typically wasted energy generated from renewable energy sources,” how to control power at the datacenter based on how much power was available at the time from the energy source, and dealing with transmission and distribution costs over a grid. *See* Ex.1001, 4:24-27; 18:25-31; 18:45-57. Yang’s teachings provide multiple solutions to these problems. Yang

teaches, for example, co-locating the datacenter with the generating station in order to take advantage expressly of “unutilized and typically wasted energy generated from renewable energy sources” – powering the computing at container datacenters “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3; Ex.1001, 18:25-31. Finally, Yang also teaches controlling the power at the datacenter based upon how much power is available at the time from the energy source, teaching that the computing resources “are activated only when stranded power is available, and will be shutdown whenever there is insufficient power.” Ex.1007, p. 1. Thus, a POSITA would have also found Yang’s teachings to be pertinent to the problems identified by the ’456 patent.

200. Second, a POSITA would have been motivated to combine Pelio and Chapel with Yang to take advantage of Yang’s teachings of harnessing “stranded power to enable scaling to higher levels of capability and performance” in datacenter computing. Ex.1007, p. 4. In Yang, “stranded power” refers to “power that is generated, but cannot be used” due to economic or reliability constraints on the grid, and is thus “wasted.” Ex.1007, p. 3. Yang’s teachings to use only such “stranded power” with co-located container-based computing results in “incremental scalability of infrastructure, low power cost, and negligible carbon footprint.” Ex.1007, p. 9. Thus, a POSITA would have been motivated to combine Yang’s

teachings with Pelio's and Chapel's container-based control of power delivery based on the availability of such stranded power to obtain these benefits.

201. Third, a POSITA would have had a reasonable expectation of success in combining the teachings of Yang with Pelio's system teachings and the combined teachings of Pelio and Chapel. Like Pelio, Yang provides server computing resources in a container-based platform for mobility and scalability purposes. *See* Ex.1005, [0055]-[0058] (ability to scale as-needed, with more efficient cooling management); Ex.1007, p. 4 (ability to scale as-needed, with more efficient cooling management).

202. Further, both Yang and Pelio teach locating the container-based datacenters behind-the-meter to avoid transmission and distribution issues like cost and loss. For example, Pelio teaches providing power to servers within a datacenter container (or multiple such containers) from a variable power source behind-the-meter. *See* Ex.1005, Abstract; [0002]; [0033]; [0037]; [0050]; [0106]. In like manner, Yang teaches co-locating the container-based datacenter at the generating source, eliminating transmission issues. Ex.1007, pp. 1, 4. Since both Pelio and Yang teach using container-based computing platforms, and siting those containers behind-the-meter, a POSITA would have recognized the compatibility of Pelio's and Yang's systems, and thus had a reasonable expectation of success in implementing Yang's teachings for using "stranded power" for powering the datacenter.

203. Fourth, combining Yang’s teachings with Pelio’s and Chapel’s combined teachings would have also been predictable and within a POSITA’s skillset. Again, it was already known to locate datacenters behind-the-meter with renewable power sources. Pelio demonstrates this, as does Yang. Ex.1005, Abstract, [0002]; Ex.1007, pp. 1, 4; *see also* Ex.1013, p. 7 (“using containers and collocating at generation sites and avoiding power distribution costs”); Ex.1014, p. 6 (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”).

204. Moreover, 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, Yang demonstrates that a POSITA understood to have the servers “only work when stranded power is available, and otherwise [have them] fully powered off.” Ex.1007, 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 Yang’s power control teachings at a datacenter (e.g., only working when stranded power available) would be a natural extension of Pelio’s understanding of the intermittent availability of power from Pelio’s behind-the-meter renewable power sources.

205. Implementing Yang’s teachings for using “stranded power” for powering the datacenter would provide the benefits of reducing power consumption when the power supply drops, such as happens cyclically per Pelio’s recognition, as well as “incremental scalability of infrastructure, low power cost, and negligible carbon footprint.” Ex.1007, p. 9. Further, Yang’s teachings for using stranded power “exploits un-economic grid power” and “creates dispatchable loads that consume excess generation ... benefit[ing] grid stability, cost, and management.” Ex.1007, p. 9. A POSITA would have therefore been motivated to turn from the combination of Pelio and Chapel to Yang’s teachings.

206. The combination permits but does not require physical incorporation of elements from Yang into Pelio. Nevertheless, the physical incorporation of Yang’s teachings for using “stranded power” for powering the datacenter would teach the claim limitations.

C. Detailed Analysis of Claims

1. Claim 5

a. [5.0] *The flexible datacenter of claim 3,*

207. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [5.1] *wherein BTM power availability comprises: one or more of excess local power generation at a local station level, 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 BTM power generation to power the flexible datacenter.

208. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises one or more of excess local power generation at a local station level, 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 BTM power generation to power the flexible datacenter.

“excess local power generation at a local station level”

209. First, as explained at [1.1], Pelio discloses BTM power generation. *See, e.g.,* Ex.1005, [0002], [0019].


210. Second, it would have been obvious that in different situations, BTM collocated power sources taught in Pelio would generate excess power. 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.

211. 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.1005, [0037]; *see also* Ex.1008, FIG. 3 (periods where available capacity 302, 306 exceeds actual computational needs 304, 308); 10:14-21. 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.1005, [0044]; *see also* [0045] (solar), [0046] (hydro). This describes extra generated power, which renders obvious “*excess local power generation at a local station level.*”

212. Because [5.1] provides the list of BTM power availability as “*one or more of ... or,*” any one option from the list renders the entire claim obvious. Thus, [5.1] is obvious in view of Pelio alone.

213. Third, “*excess local power generation*” is rendered obvious further in combination with Yang. For example, Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

214. Yang teaches that “excess power cannot be economically distributed through the grid” at times from intermittent renewable generating resources. Ex.1007, p. 1. “Renewables such as wind and solar have time-varying productivity,” with solar following a “diurnal cycle” and wind with variability over weeks or months. Ex.1007, p. 2. Due to this volatility, at times that results in “excess generation” from a local power source, which Yang teaches harnessing via co-located datacenters. Ex.1007, p. 3. Such “excess” generation (e.g., daytime as taught by Pelio as well, or high-wind situations), based on Yang’s collocated teachings and Pelio’s collocated teachings, demonstrates an example of “*excess local power generation at a local station level.*”

“local power generation subject to economic curtailment”

215. Further, Yang teaches that at times of excess generation, “power markets will price some generator’s outputs with a low or negative price” resulting in generated power that is “not economically useful.” Ex.1007, p. 3. “Such negative pricing signals are called ‘economic curtailment’ or dispatch down.” Ex.1007, p. 3. And POSITAs would have understood that, in response to the price signal, “generators will reduce production.” *See, e.g.*, Ex.1013, p. 3.

216. Yang’s description of “uneconomic” power was well understood as “*economic curtailment*” well before the earliest alleged priority date of the ’456 patent. For example, Martinot (“Grid Integration of Renewable Energy: Flexibility, Innovation, and Experience,” *Annu. Rev. Environ. Resour.*, vol. 41:223-251, 2016) discusses known curtailment practices and demonstrates the understanding of POSITAs:

Curtailment may occur (a) if there is too much generation at a given time (an overgeneration condition) and other generators cannot reduce their output fast enough or have reached their minimum output limits, (b) if there is not enough transmission or distribution line capacity to handle the renewable output, or (c) because of other grid conditions. Curtailment of wind and solar power plants causes a downward adjustment in their output, in spite of their potential to generate more power given prevailing

wind or solar conditions. Thus, curtailment results in “spilled” or wasted generation that could have come from the renewable generator but is instead lost. Curtailment is often viewed negatively, and the “repercussions of discarded energy” can include lost revenues, lost incentives, and increased business risk.

Ex.1012, p. 15. Martinot continues in demonstrating the understanding of POSITAs regarding curtailment:

There are different types of curtailment, sometimes causing confusion. The CPUC(23) defines manual curtailment as that ordered by the grid operator in response to system-reliability conditions and accomplished through exceptional dispatch (direct instruction to reduce/cease output); this is also called involuntary curtailment and generally regarded as undesirable. **Economic curtailment or voluntary curtailment, however, occurs through the wholesale market, when renewable generators voluntarily reduce their output in response to market conditions or protocols, for example a negative wholesale price (see Section 10).** Definitions of curtailment vary by jurisdiction (62).

Ex.1012, p. 15.

217. Thus, Yang’s description of generators voluntarily reducing production in response to a price signal teaches “*economic curtailment*.”

“local power generation subject to reliability[] curtailment”

218. The combination of Pelio, Chapel, and Yang teaches controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the availability of power generation subject to curtailment. *See* Ex.1005, [0032]-[0033]; Ex.1007, pp. 1, 3. This includes reliability curtailment, times at which electricity markets are not able to achieve “merit order” for purchasing energy from the lowest-cost generators via pricing alone. Ex.1007, pp. 2-3. “This is because as generation and load vary at each point in the grid, there can be transmission congestion.” Ex.1007, p. 3. Thus, both “oversupply” and “transmission congestion” prevent merit-order purchasing. Ex.1007, p. 3.

219. A POSITA would have understood that the mechanism taken to curtail power to address congestion issues was reliability curtailment, “to ensure stability or when there is overgeneration, ramp constraints, or transmission congestion.” Ex.1013, p. 3. That is, POSITA would have understood Yang’s description of power that is prevented from entering the grid by being curtailed by mechanisms other than mere pricing signals as “*reliability curtailment*.” This was a well-known concept well before the earliest alleged priority date of the ’456 patent. For example, Martinot demonstrates the understanding of POSITAs regarding curtailment:

There are different types of curtailment, sometimes causing confusion. The CPUC(23) defines **manual curtailment as that ordered by the grid operator in**

response to system-reliability conditions and accomplished through exceptional dispatch (direct instruction to reduce/cease output); this is also called involuntary curtailment and generally regarded as undesirable. Economic curtailment or voluntary curtailment, however, occurs through the wholesale market, when renewable generators voluntarily reduce their output in response to market conditions or protocols, for example a negative wholesale price (see Section 10). Definitions of curtailment vary by jurisdiction (62).

Ex.1012, p. 15.

220. Thus, a POSITA would have understood Yang’s description of generators reducing production in response to exceptional dispatch to address congestion issues to describe “*reliability curtailment.*”

“low local power generation”

221. The combination of Pelio, Chapel, and Yang renders obvious controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on low local power generation. *See* Ex.1005, [0037] (natural cycle of power), [0106] (times of no power); Ex.1007, pp. 1, 3.

222. First, as explained at [1.1], Pelio discloses BTM power generation. *See, e.g.,* Ex.1005, [0002], [0019].

223. Second, it would have been obvious that in different situations, BTM collocated power sources taught in Pelio would generate excess power or not enough


power. 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 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.1005, [0037]; *see also* Ex.1008, FIG. 3 (periods where available capacity 302, 306 falls below actual computational needs 304, 308); 10:14-21. Such dearth of generated power renders obvious the “*low local power generation.*”

224. Third, Yang further renders obvious a collocated energy generation station generating less power than the collocated datacenter (that receives the power BTM from the local station) would otherwise be able to consume.

225. Yang teaches that renewable power generation, like wind and solar, has volatility. Ex.1007, p. 2. Due to this volatility, there can be “time-varying productivity,” resulting in power that “is only intermittently available.” Ex.1007, pp. 1, 3. Indeed, Yang recognizes that there are times where the datacenter has to be “shutdown whenever there is insufficient power.” Ex.1007, p. 1.

226. Such downtimes (e.g., nighttime as taught by Pelio as well, or low-wind situations), based on Yang’s collocated teachings and Pelio’s collocated teachings, demonstrates an example of “*low local power generation.*”

227. Fourth, Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

228. Thus, Pelio combined with Chapel and Yang renders obvious “*low local power generation.*”

“transient local power generation situations”

229. The combination of Pelio, Chapel, and Yang teaches controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the availability of power generation. *See* Ex.1005, [0032]-[0033]; Ex.1007, pp. 1, 3. 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.1005, [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.1005, [0037]. The cycling between full operation and zero demonstrates a transient power generating situation.

230. Further, Yang powers the computing at container datacenters “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3. That is, the computing resources “are activated **only when stranded power is available, and will be shutdown whenever there is insufficient power.**” Ex.1007, p. 1. This also describes a “*transient power generating situation,*” where the power is either available or not (and, thus, transient).

231. Thus, a POSITA would have understood Pelio's description, and additionally Yang's description, of generators having varying levels of power generation to describe a "*transient power generating situation.*"

2. Claim 6

a. [6.0] *The flexible datacenter of claim 3, wherein BTM power availability comprises:*

232. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [6.1] *excess local power generation at a local station level.*

233. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises excess local power generation at a local station level.

234. First, as explained at [1.1], Pelio discloses BTM power generation. See, e.g., Ex.1005, [0002], [0019].

235. Second, it would have been obvious that in different situations, BTM collocated power sources taught in Pelio would generate excess power. 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 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.1005, [0037]; see also Ex.1008, FIG. 3 (periods where available capacity 302, 306 exceeds actual


computational needs 304, 308); 10:14-21. Such extra generated power renders obvious the “*excess local power generation at a local station level.*”

236. Third, Yang further renders obvious a collocated energy generation station generating excess power that the collocated datacenter (that receives the power BTM from the local station) can consume.

237. Yang teaches that “excess power cannot be economically distributed through the grid” at times from intermittent renewable generating resources. Ex.1007, p. 1. “Renewables such as wind and solar have time-varying productivity,” with solar following a “diurnal cycle” and wind with variability over weeks or months. Ex.1007, p. 2. Due to this volatility, at times that results in “excess generation” from a local power source, which Yang teaches harnessing via co-located datacenters. Ex.1007, p. 3. Such “excess” generation (e.g., daytime as taught by Pelio as well, or high-wind situations), based on Yang’s collocated teachings and Pelio’s collocated teachings, demonstrates an example of “*excess local power generation.*”

238. Fourth, Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

239. Yang therefore confirms Pelio’s BTM teaching, and renders obvious that the “*excess local power generation*” (as taught by Yang) is indeed local (collocated datacenter/renewable generator). In combination, therefore, a POSITA would understand the combination as teaching a controller (the remote controller of Chapel) monitoring BTM power availability (in its example, from a solar-powered generator, or wind generator, per Pelio and Yang), and determining a course of action based thereon based on how much power is available for the datacenter to process tasks (such as taught by Chapel combined with Yang).

240. Thus, the combination of Pelio, Chapel, and Yang renders obvious “*wherein BTM power availability comprises excess local power generation at a local station level.*”

3. Claim 7

a. [7.0] *The flexible datacenter of claim 3,*

241. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [7.1] *wherein BTM power availability comprises: local power generation subject to economic curtailment.*

242. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises local power generation subject to economic curtailment.

243. First, Pelio recognizes that it is important “to keep the grid in balance,” which would result in times where there is generated electricity that exceeds demand on the grid. Ex.1005, [0032]-[0033].

244. Second, power generated but subject to economic curtailment would have been an obvious type of power availability for the combined teachings of Pelio and Chapel to monitor.

245. For example, Yang teaches that at times of excess generation, “power markets will price some generator’s outputs with a low or negative price” resulting in generated power that is “not economically useful.” Ex.1007, p. 3. “[S]uch negative pricing signals are called ‘economic curtailment’ or dispatch down.”

Ex.1007, p. 3. And POSITAs would have understood that, in response to the price signal, “generators will reduce production.” *See, e.g.*, Ex.1013, p. 3.

246. Yang’s description of “uneconomic” power was well-understood as “*economic curtailment*” well before the earliest alleged priority date of the ’456 patent. For example, Martinot demonstrates the understanding of POSITAs:

Curtailment may occur (a) if there is too much generation at a given time (an overgeneration condition) and other generators cannot reduce their output fast enough or have reached their minimum output limits, (b) if there is not enough transmission or distribution line capacity to handle the renewable output, or (c) because of other grid conditions. Curtailment of wind and solar power plants causes a downward adjustment in their output, in spite of their potential to generate more power given prevailing wind or solar conditions. Thus, curtailment results in “spilled” or wasted generation that could have come from the renewable generator but is instead lost. Curtailment is often viewed negatively, and the “repercussions of discarded energy” can include lost revenues, lost incentives, and increased business risk.

Ex.1012, p. 15. Martinot continues:

There are different types of curtailment, sometimes causing confusion. The CPUC(23) defines manual curtailment as that ordered by the grid operator in response to system-reliability conditions and accomplished through

exceptional dispatch (direct instruction to reduce/cease output); this is also called involuntary curtailment and generally regarded as undesirable. **Economic curtailment or voluntary curtailment, however, occurs through the wholesale market, when renewable generators voluntarily reduce their output in response to market conditions or protocols, for example a negative wholesale price (see Section 10).** Definitions of curtailment vary by jurisdiction (62).


Ex.1012, p. 15.

247. Thus, Yang’s description of generators voluntarily reducing production in response to a price signal describes “*economic curtailment*.”

248. Third, the combined teachings of Pelio, Chapel, and Yang further renders obvious controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the availability of power generation subject to curtailment including economic curtailment.

249. For example, Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

250. Thus, Pelio combined with Chapel and Yang renders obvious “*wherein BTM power availability comprises local power generation subject to economic curtailment.*”

4. Claim 8

a. [8.0] *The method according to claim 3,*

251. *See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].*

b. [8.1] *wherein BTM power availability comprises: local power generation subject to reliability curtailment.*

252. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises local power generation subject to reliability curtailment.

253. As demonstrated with respect to [5.1], the combination of Pelio, Chapel, and Yang teaches controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the availability of power generation subject to curtailment. *See* Ex.1005, [0032]-[0033]; Ex.1007, pp. 1, 3. This includes reliability curtailment, times at which electricity markets are not able to achieve “merit order” for purchasing energy from the lowest-cost generators via pricing alone. Ex.1007, pp. 2-3. “This is because as generation and load vary at each point in the grid, there can be transmission congestion.” Ex.1007, p. 3. Thus, both “oversupply” and “transmission congestion” prevent merit-order purchasing. Ex.1007, p. 3. A POSITA would have understood that the mechanism taken to curtail power to address congestion issues included reliability curtailment, “to ensure stability or when there is overgeneration, ramp constraints, or transmission congestion.” Ex.1013, p. 3.

254. That is, a POSITA would have understood Yang’s description of power that is prevented from entering the grid by being curtailed by mechanisms other than mere pricing signals as “*reliability curtailment*.” This was a well-known concept well before the earliest alleged priority date of the ’456 patent. For example, Martinot demonstrates the understanding of POSITAs:


There are different types of curtailment, sometimes causing confusion. The CPUC(23) defines **manual curtailment as that ordered by the grid operator in**

response to system-reliability conditions and accomplished through exceptional dispatch (direct instruction to reduce/cease output); this is also called involuntary curtailment and generally regarded as undesirable. Economic curtailment or voluntary curtailment, however, occurs through the wholesale market, when renewable generators voluntarily reduce their output in response to market conditions or protocols, for example a negative wholesale price (see Section 10). Definitions of curtailment vary by jurisdiction (62).

Ex.1012, p. 15.

255. Thus, a POSITA would have understood Yang's description of generators reducing production in response to exceptional dispatch to address congestion issues to describe "*reliability curtailment*."

256. For example, Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

257. Thus, Pelio combined with Chapel and Yang renders obvious “*wherein BTM power availability comprises local power generation subject to reliability curtailment.*”

5. Claim 10

a. [10.0] *The flexible datacenter of claim 3,*

258. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [10.1] *wherein BTM power availability comprises: low local power generation.*

259. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises: low local power generation.

260. First, as explained at [1.1], Pelio discloses BTM power generation. See, e.g., Ex.1005, [0002], [0019].

261. Second, it would have been obvious that in different situations, BTM collocated power sources taught in Pelio would generate excess power or not enough


power. 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 generates less power than the datacenter may need, 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.1005, [0037]; *see also* Ex.1008, FIG. 3 (periods where available capacity 302, 306 falls below actual computational needs 304, 308); 10:14-21. Such dearth of generated power renders obvious the “*low local power generation*.”

262. Third, Yang further renders obvious a collocated energy generation station generating less power than the collocated datacenter (that receives the power BTM from the local station) would otherwise be able to consume.

263. Yang teaches that renewable power generation, like wind and solar, has volatility. Ex.1007, p. 2. Due to this volatility, there can be “time-varying productivity,” resulting in power that “is only intermittently available.” Ex.1007, pp. 1, 3. Indeed, Yang recognizes that there are times where the datacenter has to be “shutdown whenever there is insufficient power.” Ex.1007, p. 1.

264. Such downtimes (e.g., nighttime as taught by Pelio as well, or low-wind situations), based on Yang’s collocated teachings and Pelio’s collocated teachings, demonstrates an example of “*low local power generation.*”

265. Fourth, Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

266. Thus, Pelio combined with Chapel and Yang renders obvious “*wherein BTM power availability comprises: low local power generation.*”

6. **Claim 12**

a. **[12.0] *The flexible datacenter of claim 3,***

267. See analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. **[12.1] *wherein BTM power availability comprises: transient local power generation situations.***

268. Pelio combined with Chapel and Yang renders obvious wherein BTM power availability comprises: transient local power generation situations.

269. 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.1005, [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.1005, [0037]. The cycling between full operation and zero demonstrates a transient power generating situation.

270. Further, Yang powers the computing at container datacenters “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, pp. 1, 2-3. That is, the computing resources “are activated **only when stranded power is available, and will be shutdown whenever there is insufficient power.**” Ex.1007, p. 1. This also describes a “*transient power generating situation,*”

where the power is either available or not (and, thus, transient). A POSITA would have understood Pelio's description, and additionally Yang's description, of generators having varying levels of power generation to describe a "*transient power generating situation*."

271. Thus, Pelio combined with Chapel and Yang renders obvious "*wherein BTM power availability comprises: transient local power generation situations*."

7. Claim 28

- a.** [28.0] *A method of dynamic power delivery to a flexible datacenter using behind-the-meter (BTM) power comprising:*

272. To the extent the preamble is limiting, Pelio, alone and in combination with Chapel, renders obvious a method of dynamic power delivery to a flexible datacenter using behind-the-meter (BTM) power.

273. First, Pelio discloses a flexible datacenter that uses BTM power. *See* [1.0] (flexible datacenter), [1.1] (BTM power). Pelio, in particular, teaches that collocating the datacenters with the generation source results in "a higher percentage of the power produced by the generation source actually reach[ing] critical load." Ex.1005, [0019]. As a result, "computer processing is done at the source of the power." Ex.1005, [0019].

274. Second, Pelio teaches that power from renewable sources of energy "naturally cycle between full operation and zero, for example the daily cycles of a

solar energy generating station that depend on sunshine.” Ex.1005, [0037]. “Others cycle with the wind, or demands from the grid.” Ex.1005, [0037]. It would have been obvious to a POSITA that the cycling of power output from the generating station for renewables like solar, wind, etc., results in a dynamic delivery of power to the collocated datacenters taught by Pelio, thus rendering this limitation obvious.

275. Third, Pelio combined with Chapel further renders this limitation obvious. Chapel teaches that “[d]uring operation of the system, a controller such as a local controller may identify (512) a condition governed by a policy ... such as a change in operating mode dictated by the grid power supplier” and received from a remote controller. Ex.1006, [0076], [0088], [0089] (e.g., remote grid controller); [0058] and [0094]-[0095]. Responsive to the instruction from a remote controller, the local controller at the datacenter dynamically modifies power delivery by controlling delivery of power to the datacenter loads (one or more servers). Ex.1006, [0027]-[0028], [0054]-[0056].

276. Thus, Pelio, alone and in combination with Chapel, renders obvious “*a method of dynamic power delivery to a flexible datacenter using behind-the-meter (BTM) power.*”


b. [28.1] *monitoring, by a control system, BTM power availability at the flexible datacenter,*

277. Pelio combined with Chapel renders obvious monitoring, by a control system, BTM power availability at the flexible datacenter.

278. As demonstrated at [1.1], Pelio teaches a datacenter receiving power via a BTM power source. Ex.1005, [0002], [0019], [0039]. Pelio expressly recognizes that some BTM power sources, “such as solar photo-voltaics, wind, wave, and tidal, by their nature can only add power to the grid independent of the demand.” Ex.1005, [0033]. It would have been obvious to a POSITA to monitor the varying output of these power sources in order to “throttle up and down” other power sources (Ex.1005, [0033]), or the load itself (Ex.1005, [0106] (dropping load when generating output drops), Ex.1006, [0089] (powering up or down devices depending upon over-capacity condition)). That is, in order to know whether there is a need to throttle up or down another power source, or the level of load itself, a POSITA would recognize that the current status would need to be monitored in Pelio.

279. 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.1006, [0004]. The power distribution is from power generation sources. Ex.1006, [0004], [0082] (power sources). Chapel teaches reporting information about loads to “allow[] for remotely monitoring and/or controlling electrical devices” in datacenters. Ex.1006, [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.1006, [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.1006, [0056]. Chapel’s Figure 6 includes another example, with local controller
650 sending information about power availability and load to remote controller 670.
Ex.1006, [0083].



Ex.1006, Fig. 3 (annotated).

280. 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 Pelio and recognized by Chapel). See Ex.1005, [0033], [0106]; Ex.1006, [0004], [0089]. With the power availability being monitored as taught by Chapel, Pelio’s system would thus monitor availability of power at the datacenter via the one or more BTM power sources.

281. Thus, Pelio combined with Chapel renders obvious “*monitoring, by a control system, BTM power availability at the flexible datacenter.*”

- c.** [28.2] *wherein the flexible datacenter includes a BTM power input system configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;*

282. Pelio discloses wherein the flexible datacenter includes a BTM power input system configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid. *See* [1.1].

- d.** [28.3] *determining that a ramp-up condition is met, wherein the ramp-up condition depends on BTM power availability at the flexible datacenter; and*

283. Pelio combined with Chapel and Yang renders obvious determining that a ramp-up condition is met, wherein the ramp-up condition depends on BTM power availability at the flexible datacenter.


284. First, as demonstrated at [1.0], Pelio teaches a flexible datacenter. As further demonstrated at [1.1], Pelio also teaches providing available BTM power to the flexible datacenters for datacenter operation. *See, e.g.,* Ex.1005, [0002] (BTM), [0005] (BTM power availability dependent upon conditions such as water, wind, sunshine), [0037] (solar energy cycles).

285. Second, Chapel further teaches monitoring power availability (and, in combination with Pelio, renders obvious monitoring BTM power availability in particular, as demonstrated at [28.1]). *See* Ex.1006, [0004], [0047], [0056], [0083]. Moreover, a POSITA would have understood that Chapel's load control involves not only controlling delivery of power by turning off devices, operating the devices in reduced power mode, or interrupting power distribution (Ex.1006, [0025]-[0026]), but also turning the devices back on, ending reduced power mode, removing power interruption, etc. And, since Chapel teaches the remote controller performing the power delivery control in situations where supply is expected to be less than demand (the load), it would have been obvious to a POSITA for the remote controller to also cause load to increase when the supply becomes sufficient again (i.e., no over-capacity condition threat anymore).

286. Third, this is confirmed by the teachings of Yang. Yang teaches that it was known to follow the supply of available power from a renewable source in powering servers in a datacenter, including determining that a condition has been met to increase the load.

287. Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches locating a datacenter at the same location as the renewable generation, prior to the energy reaching the grid – just as Pelio does. Yang teaches ways in which to

harness what Yang coins as stranded power – power that is curtailed or otherwise available from renewable power sources but not needed at that time on the grid. Ex.1007, p. 3. Yang does so by 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.1007, pp. 1, 4; *see also* Fig. 3:



Ex.1007, Fig. 3 (annotated).

288. Yang’s main focus is on powering the servers at the co-located datacenters “exclusively by wasted power generated from intermittent renewable sources.” Ex.1007, p. 1. Such “wasted power” is only wasted because the grid to which the renewable sources are otherwise connected does not use it at times. Ex.1007, pp. 1-3. Further, Yang teaches that the co-location of the datacenters “eliminates transmission” costs typically arising from normal grid use. Thus, a POSITA would have understood that Yang’s teachings – particularly in combination

with Pelio's BTM teachings – provide an example of a BTM power source that has a varying “*BTM power availability at the flexible datacenter.*”

289. Yang explains, in particular, that “ZCCloud [the datacenter] resources are activated only when stranded power is available, and will be shutdown whenever there is insufficient power.” Ex.1007, p. 1; *see also* p. 4 (“ZCCloud computing resources only work when stranded power is available”). Yang's teaching of using renewable power once it is available demonstrates “*determining that a ramp-up condition is met,*” namely that sufficient power is available from the renewable system is met in order to run the co-located datacenter.

290. In combination, a POSITA would understand that the remote controller taught by Chapel, while monitoring power availability from the BTM power source at a flexible datacenter (such as taught by Pelio), would determine that a condition is met to increase the load imposed by the datacenter servers when sufficient power exists from the power source to operate, as taught by Yang. Moreover, based on the combined teachings of Chapel and Yang, a POSITA would have further understood that Chapel's power modulation teachings would apply based on the variability of the renewable source's power output to also operate on a per-computer or per-rack basis (as Chapel teaches) to follow the amount of power output during operation. For example, Li (“SolarCore: Solar Energy Driven Multi-Core Architecture Power Management”) teaches that it was known to adjust the load imposed by servers in a

datacenter to follow the supply of available power from a renewable source, where “the load starts to operate using solar power when the amount of renewable energy exceeds a power-transfer threshold” and follows the supply availability from the PV system thereafter. Ex.1017, p. 9.


291. Thus, Pelio combined with Chapel and Yang renders obvious “*determining that a ramp-up condition is met, wherein the ramp-up condition depends on BTM power availability at the flexible datacenter.*”

- e. **[28.4]** *based on determining that the ramp-up condition is met, providing an operational directive to a datacenter control system at the flexible datacenter,*


292. Pelio combined with Chapel and Yang renders obvious based on determining that the ramp-up condition is met, providing an operational directive to a datacenter control system at the flexible datacenter.

293. As demonstrated at [28.3], the combination of Pelio, Chapel, and Yang teaches determining that a ramp-up condition is met by a remote controller. Further, Chapel generally teaches modulating power delivery by the local controller (a “*datacenter control system*”) by policies including those communicated from a remote controller. “For example, the local controller may receive a message from the electrical utility provide specifying transition to an efficiency mode or a brownout mode.” Ex.1006, [0076]; *see also* [0058], [0066], [0090]. Such “messages,” and the policies communicated from remote controller(s), render

obvious “*providing an operational directive.*” The messages are sent to the local controller, an example of a “*datacenter control system.*” See analysis at [1.4], Figures 3, 6:



Ex.1006, Fig. 3 (annotated).



294. Further, Yang teaches starting to operate using the renewable energy once stranded power becomes available from the co-located renewable source. Ex.1007, pp. 1, 4. It would have been obvious to a POSITA for Yang's starting to operate to occur from an operational directive, such as the message from the remote controller taught by Chapel.

295. Thus, the combination of Pelio, Chapel, and Yang renders obvious “*based on determining that the ramp-up condition is met, providing an operational directive to a datacenter control system at the flexible datacenter.*”

- f. **[28.5] wherein the operational directive instructs the datacenter control system to enable BTM power delivery to one or more computing systems at the flexible datacenter.**

296. Pelio combined with Chapel and Yang renders obvious wherein the operational directive instructs the datacenter control system to enable BTM power delivery to one or more computing systems at the flexible datacenter.

297. As explained at [28.4] and [1.4], Pelio combined with Chapel renders obvious an “*operational directive*” that instructs the “*datacenter control system*” (e.g., a “local controller” taught by Chapel at Pelio’s datacenter) to modulate power delivery to one or more computing systems at the datacenter. *See, e.g.,* Ex.1006, [0025]-[0026], [0076]; *see also* [0058], [0066], [0090].

298. Further, Yang teaches that “ZCCloud [the datacenter] resources are activated only when stranded power is available, and will be shutdown whenever

there is insufficient power.” Ex.1007, p. 1; *see also* p. 4 (“ZCCloud computing resources only work when stranded power is available”). In combination, therefore, a POSITA would understand Chapel’s teaching of power modulation (e.g., “*power delivery*”) to enable delivery of Pelio’s BTM power to one or more servers in the datacenter, in response to the amount of renewable energy available as stranded power being greater than a threshold as taught by Yang (i.e., when stranded power is available versus when it is not). In particular, a POSITA would apply Chapel’s teaching of modulating power delivery (interrupting power distribution) to remove the interruption once sufficient stranded power was again available (such as taught by Yang, which would meet Chapel’s needs of avoiding an over-capacity condition).

299. Thus, Pelio combined with Chapel and Yang renders obvious “*wherein the operational directive instructs the datacenter control system to enable BTM power delivery to one or more computing systems at the flexible datacenter.*”

8. Claim 29

a. [29.0] *The method of claim 28, further comprising:*

300. *See analysis at [28.0]-[28.5].*

b. [29.1] *determining that a ramp-down condition is met, wherein the ramp-down condition depends on BTM power availability at the flexible datacenter; and*

301. Pelio combined with Chapel, alone and further combined with Yang, renders obvious determining that a ramp-down condition is met, wherein the ramp-down condition depends on BTM power availability at the flexible datacenter.

302. As explained at [1.4], Pelio combined with Chapel teaches modulating power delivery to reduce the load in response to the variable availability of power, such as from a renewable source. *See, e.g.*, Ex.1005, [0106] (throttling load by dropping it when generating output drops, such as from a BTM source); Ex.1006, [0025]-[0026], [0076]; *see also* [0058], [0066], [0090]. Thus, both Pelio and Chapel teach that the load throttling is dependent upon power availability. It would have been obvious to a POSITA that these sections from Pelio (e.g., where load is dropped) and Chapel (similar) demonstrate a determination that a ramp-down condition has been met that was dependent upon the lack of power availability at the datacenter. That is, it would have been obvious to a POSITA that Pelio's and Chapel's teachings of dropping load when there is insufficient power availability would have included a determination that the power availability is insufficient prior to, or along with, the implementation of the described action.

303. Further, Yang teaches that "ZCCloud [datacenter] resources are activated only when stranded power is available, and will be **shutdown whenever there is insufficient power.**" Ex.1007, p. 1; *see also* p. 4 ("ZCCloud computing resources only work when stranded power is available, and otherwise are fully powered off"). Thus, the combined teachings demonstrate determining a "*ramp-down condition is met*" (e.g., dropping generator output per Pelio, over-voltage condition per Chapel, power supply from stranded renewable power unavailable per

Yang) that depends on Pelio's BTM power availability, such as of a renewable power source as taught by both Chapel and Yang.

304. Thus, Pelio combined with Chapel and Yang renders obvious “*determining that a ramp-down condition is met, wherein the ramp-down condition depends on BTM power availability at the flexible datacenter.*”

c. **[29.2] *based on determining that the ramp-down condition is met, providing a second operational directive to the datacenter control system at the flexible datacenter,***

305. Pelio combined with Chapel and Yang renders obvious based on determining that the ramp-down condition is met, providing a second operational directive to the datacenter control system at the flexible datacenter.

306. As demonstrated at [29.1], the combination of Pelio, Chapel, and Yang teaches determining that a ramp-down condition is met by a remote controller. As further demonstrated at [28.3], Chapel teaches modulating power delivery by the local controller (a “*datacenter control system*”) by policies including those communicated from a remote controller. “For example, the local controller may receive a message from the electrical utility provide specifying transition to an efficiency mode or a brownout mode.” Ex.1006, [0076]; *see also* [0058], [0066], [0090]. Such “messages,” and the policies communicated from remote controller(s), render obvious “*providing a second operational directive.*” The messages are sent to the local controller, a “*datacenter control system.*” *See* analysis at [1.4].

307. Further, Yang teaches stopping use of the renewable power source if its availability as stranded power (i.e., curtailed by the grid) falls below an amount sufficient for operation. Ex.1007, pp. 1, 4; *see also* [29.1]. It would have been obvious to a POSITA for Yang’s “shutdown” to occur from another operational directive, such as a message from the remote controller taught by Chapel, which provides several examples of instructions to shut down individual or multiple servers. *See, e.g.*, Ex.1006, [0058], [0066], [0076], [0090].

308. Thus, Pelio combined with Chapel and Yang renders obvious “*based on determining that the ramp-down condition is met, providing a second operational directive to the datacenter control system at the flexible datacenter.*”

- d.** [29.3] *wherein the second operational directive instructs the datacenter control system to disable BTM power delivery to one or more computing systems at the flexible datacenter.*

309. Pelio combined with Chapel renders obvious wherein the second operational directive instructs the datacenter control system to disable BTM power delivery to one or more computing systems at the flexible datacenter. *See* analysis at [1.4] (demonstrating modulating power delivery by disabling power delivery per Chapel’s teachings, combined with Pelio’s BTM teachings).

9. Claim 30

- a.** [30.0] *The method of claim 28,*

310. *See* analysis at [28.0]-[28.5].

- b. **[30.1] *wherein the control system is positioned remotely from the flexible datacenter.***


311. Pelio combined with Chapel renders obvious wherein the control system is positioned remotely from the flexible datacenter. *See* [4.0].

XI. CHALLENGE 3: THE COMBINATION OF PELIO, CHAPEL, AND BELADY RENDERS OBVIOUS CLAIMS 19, 21-24, AND 26

A. Summary of Belady

312. 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 ’456 patent.

313. 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).

314. 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, 2:15-22.

315. The management system temporally moves computational tasks by “adjust[ing] a time at which a data component, associated with a computational tasks, 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.

316. 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 Pelio and Chapel with Belady

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

318. First, Belady, like Pelio, focuses on analogous issues to those in the ’456 patent, including power delivery to computers in a datacenter. 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 Pelio and the ’456 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 '456 patent.

319. Further, Belady is pertinent to the problems identified by the '456 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, 4:24-27; 18:25-31; 18:45-57. 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 '456 patent.

320. Second, a POSITA would have been motivated to turn from Pelio to Belady in order to take advantage of Belady’s teachings of how to move computational tasks due to power availability issues, and how to detect and determine that it is appropriate to move. Pelio already recognized that when there is

insufficient renewable power for a datacenter container, it is desirable to move pending compute jobs to another datacenter. Ex.1005, [0106]. Pelio explains that “a cloud network architecture in the Internet could be relied upon to seamlessly assume the data center jobs” from the container when power is insufficient or gone. Ex.1005, [0106]; *see also* Ex.1014, 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 Pelio’s suggestion.

321. For example, Belady teaches using a management system to manage the execution of tasks – like 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 on how to implement Pelio’s movement of datacenter jobs.

322. Third, POSITA would have had a reasonable expectation of success in applying Belady’s movement of computational tasks teachings to the combined teachings of Pelio and Chapel. Like 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, [0002]. Moreover, both Pelio and Belady recognize the utility of moving computing jobs/tasks between datacenters 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, [0106]; Ex.1008, 2:15-22. Since both 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 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.

323. In combination, Pelio's datacenter containers would implement Belady's real-time monitoring teachings, and computational task movement teachings, by a control system such as a local controller and/or remote controller such as those taught by Chapel, with Belady's management system teachings for implementing the task movement. Doing so would have been predictable and within a POSITA's skillset. Indeed, Pelio already recognized the desirability of moving datacenter jobs due to 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, Abstract, [0002]; Ex.1007, pp. 1, 4; Ex.1013, p. 7; Ex.1014, p. 6; Ex.1015, pp. 3, 4; Ex.1016, p. 2.

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

C. Detailed Analysis of Claims

1. Claim 19

a. [19.0] *The flexible datacenter of claim 1,*

325. See analysis at [1.0]-[1.4] in Challenge 1.

b. [19.1] *wherein the operational directive is a workload directive based on a current BTM power availability and a projected BTM power availability at the flexible datacenter.*

326. Pelio combined with Chapel and Belady renders obvious wherein the operational directive is a workload directive based on a current BTM power availability and a projected BTM power availability at the flexible datacenter.


327. First, as demonstrated with respect to [1.4], Pelio combined with Chapel renders obvious an “*operational directive*” with Chapel's teachings of messages from a remote/grid controller. “For example, the local controller may receive a message from the electrical utility provide specifying transition to an efficiency mode or a brownout mode.” Ex.1006, [0076]; *see also* [0058], [0066], [0090]. The local controller at the customer premises is used to “control[] delivery of power to the electrical device.” Ex.1006, [0018].

328. Second, Chapel further renders obvious that the operational directive is a “*workload directive*” that is based on power availability. Chapel teaches that the modulation of the power delivery by the local controller is done by policies including those communicated from a remote controller. 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.1006, [0025]-[0026]. Thus, since the directive turns on or off, or puts into reduced power mode, computers in a datacenter, a POSITA would have understood the directive to be one related to the workload of the computers.

329. Third, as explained with respect to [14.1], Pelio combined with Chapel renders obvious modulating power delivery to the plurality computing systems based on BTM power availability at the BTM power input system. It would have been obvious to a POSITA to further base the directive on projected BTM power availability as well, since this would allow the Pelio-Chapel combined teachings to operate with future availability taken into account. This is confirmed by the teachings of Belady.

330. Belady teaches using “real-time information provided by the data centers” about the availabilities 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). Belady uses this information in

determining how to direct operations at one or more datacenters to proceed. *See* Ex.1008, 10:60-62 (management system uses the information to manage execution of computational tasks). Belady teaches this as implemented by the management side, for example by a management functionality 602. *See* Figure 6:



Ex.1008, Fig. 6 (annotated).

331. The management functionality 602 “can be implemented by the management system 120 of FIG. 1.” Ex.1008, 12:20-30. And per Figure 1, management system 120 can be “distinct from the individual data centers,” including “by a separate system.” Ex.1008, 8:15-45.

332. In combination, a POSITA would have implemented Belady’s teachings of using both real-time and predicted power availability at the datacenters

when determining the message that is sent to the local controller at the datacenter for implementation, as taught by the combination of Pelio and Chapel.

333. Thus, the combination of Pelio, Chapel, and Belady renders obvious “*wherein the operational directive is a workload directive based on a current BTM power availability and a projected BTM power availability at the flexible datacenter.*”

2. Claim 21

a. [21.0] *The flexible datacenter of claim 1,*

334. See analysis at [1.0]-[1.4] in Challenge 1.

b. [21.1] *wherein the datacenter control system is further configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems, and*


335. Pelio combined with Chapel and Belady renders obvious wherein the datacenter control system is further configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems.

336. First, Pelio generally teaches a plurality of flexible datacenters, each with its own plurality of computing systems. *See* [1.3]; Ex.1005, [0048], [0049] (multiple containers), [0071].

337. Second, it would have been obvious to transfer workload between flexible datacenters, whether on the same premises or at different locations, and thus transfer between different computing systems. For example, Pelio explains that

when there is insufficient renewable power for a datacenter container, “a cloud network architecture in the Internet could be relied upon to seamlessly assume the data center jobs” from the container when power is insufficient or gone. Ex.1005, [0106]. Belady also demonstrates this and expands upon Pelio’s teaching by providing implementation detail teachings.

338. Belady teaches that “a renewable power source can be generally co-located with the datacenter to which it provides power,” to thereby “eliminate some of the infrastructure associated with conventional power distribution architecture.” Ex.1008, 6:53-7:17. Notably, while in some examples Belady’s power sources are “off-grid,” in other examples such collocated power sources “can be connected to the grid.” Ex.1008, 6:53-67; *see also* FIG. 1:



Ex.1008, Fig. 1 (annotated).

339. Belady teaches that, for a collocated datacenter to a renewable power source such as wind or solar, “computational capacity [of the collocated datacenter] has an undulating shape that reflects the availability of power throughout a 24 hour cycle” from the renewable power source. Ex.1008, 9:19-21. For example, for a solar power source, “[i]n the daylight hours, sunlight is available, so computational capacity 202 is available,” while at nighttime sunlight and, thus, power, is not available. Ex.1008, 9:21-24.

340. Belady further explains that the “actual computational needs 204” does not always match the capacity (and, thus, power output available from the renewable power source). Ex.1008, 9:24-29. To address this, Belady teaches (1) “temporal

movement of data components” or (2) “location-based (i.e., spatial) movement of data components,” where a “data component refers to any part of a computational task.” Ex.1008, 9:14-16, 6:37-38. Figure 3 relates to a situation where a controller determines to shift computational tasks between datacenters at different locations (X and Y in the example) to take advantage of different periods of excess power generation at each location. Ex.1008, 10:14-34; Figure 3:




FIG. 3

Ex.1008, Fig. 2 (annotated).

341. Specifically, Belady teaches “sending data components” from location X to location Y at the time of the lulls at location X, which corresponds to a time of excess generation at location Y, as shown by “movement 310” in Figure 3 above. Ex.1008, 10:25-28. Similarly, Belady teaches taking advantage of lulls in capacity

at location Y “by sending data components to location X” as shown by movement 312. Ex.1008, 10:29-32.

342. In combination, a POSITA would have implemented Belady’s data movement teachings (between different datacenters) with Chapel’s datacenter control teachings.


343. Thus, Pelio combined with Chapel and Belady renders obvious “*wherein the datacenter control system is further configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems.*”

c. [21.2] *wherein the second plurality of computing systems is positioned at a second flexible datacenter.*

344. Pelio combined with Chapel and Belady renders obvious wherein the second plurality of computing systems is positioned at a second flexible datacenter.

345. First, Pelio generally teaches a plurality of flexible datacenters, each with its own plurality of computing systems. *See* [1.3]; Ex.1005, [0048], [0049] (multiple containers), [0071]. Thus, among the plurality of datacenters in Pelio would be one that demonstrates a “first flexible datacenter” as well as another one that demonstrates a “*second flexible datacenter.*” Moreover, each datacenter container in Pelio would be understood to have its own sets of servers. *See, e.g.,* Ex.1005, [0048], [0071].

346. Second, as demonstrated at [21.1], Belady renders obvious moving data components between two different datacenters. *See* Ex.1008, 10:25-28.



Ex.1008, Fig. 1 (annotated).

347. In combination, it would have been obvious for each datacenter taught by Belady to be flexible datacenters, as taught by Pelio. This would result in the first datacenter in Figure 1 of Belady being a “*first flexible datacenter*” and either of the other datacenters being a “*second flexible datacenter*.” See Ex.1008, FIG. 1 above.

348. Thus, Pelio combined with Chapel and Belady renders obvious “*wherein the second plurality of computing systems is positioned at a second flexible datacenter.*”

3. Claim 22

a. [22.0] *The flexible datacenter of claim 21,*

349. See analysis at [1.0]-[1.4] in Challenge 1 and [21.0]-[21.2].

- b. [22.1] *wherein the flexible datacenter is coupled to a first power generation station such that the plurality of computing systems operates using BTM power from the first power generation station; and*

350. Pelio combined with Chapel and Belady renders obvious wherein the flexible datacenter is coupled to a first power generation station such that the plurality of computing systems operates using BTM power from the first power generation station.

351. In general, Pelio teaches multiple datacenters connected with at least one BTM power source. *See* Ex.1005, [0044]-[0047]. Further, Belady teaches that the first datacenter (a first flexible datacenter in the Pelio-Belady combined teachings) “receives power from ... the wind power source 112.” Ex.1008, 6:43-67; 7:31-40.




FIG. 1

Ex.1008, Fig. 1 (annotated).

352. As Belady expressly teaches, “a renewable power source can be generally co-located with the data center to which it provides power,” to thereby “eliminate some of the infrastructure associated with conventional power distribution architecture.” Ex.1008, 6:53-7:17. Notably, while in some examples Belady’s power sources are “off-grid,” in other examples such collocated power sources “can be connected to the grid.” Ex.1008, 6:53-67. This includes datacenter A receiving power from renewable power source 112 (which includes operation as a BTM power source) as an example. Ex.1008, 7:31-40. This is consistent with Pelio’s behind-the-meter datacenter location teachings, and thus a POSITA would have found the Pelio-Belady combination to render this limitation obvious, where

Ex.1003


for example different of Pelio's flexible datacenters are powered by different behind-the-meter power sources, as taught by Belady.

353. Thus, Pelio combined with Chapel and Belady renders obvious *“wherein the flexible datacenter is coupled to a first power generation station such that the plurality of computing systems operates using BTM power from the first power generation station.”*

- c. **[22.2]** *wherein the second flexible datacenter is coupled to a second power generation station such that the second plurality of computing systems operates using BTM power from the second power generation station.*

354. Pelio combined with Chapel and Belady renders obvious wherein the second flexible datacenter is coupled to a second power generation station such that the second plurality of computing systems operates using BTM power from the second power generation station.

355. Belady teaches that the second datacenter (a second flexible datacenter in the Pelio-Belady combined teachings) receives power from the “solar power source 114.” Ex.1008, 6:43-67; 7:31-40; *see also* Figure 1:



Ex.1008, Fig. 1 (annotated).

356. Again, since Belady teaches that “a renewable power source can be generally co-located with the data center to which it provides power,” and that collocated power sources can be connected to the grid as well, Belady teaches the datacenter B receiving power from renewable power source 114 for the computing systems to operate using BTM power from the source 114. Ex.1008, 6:53-7:17, 7:31-40.

357. Thus, Pelio combined with Chapel and Belady renders obvious “wherein the second flexible datacenter is coupled to a second power generation station such that the second plurality of computing systems operates using BTM power from the second power generation station.”

4. Claim 23

a. [23.0] *The flexible datacenter of claim 22,*

358. See analysis at [1.0]-[1.4] in Challenge 1, [21.0]-[21.2], and [22.0]-[22.2].

b. [23.1] *wherein the datacenter control system is configured to transfer the workload from the plurality of computing systems to the second plurality of computing systems in response to detecting a change in BTM power availability at the plurality of computing systems.*


359. Pelio combined with Chapel and Belady renders obvious wherein the datacenter control system is configured to transfer the workload from the plurality of computing systems to the second plurality of computing systems in response to detecting a change in BTM power availability at the plurality of computing systems.

360. First, the combined teachings of Pelio and Chapel render obvious the “*datacenter control system*” that operates based on BTM power availability. *See* [1.4].

361. Second, the combined teachings of Pelio, Chapel, and Belady render obvious transferring workload between pluralities of computing systems at different flexible datacenters. *See* [21.1].

362. Third, Belady specifically teaches performing the transfer between datacenters in response to detecting a change in BTM power availability at a datacenter (and, hence, at a plurality of computing systems). For example, Belady

teaches “sending data components” from location X to location Y at the time of the lulls at location X, which corresponds to a time of excess generation at location Y, as shown by “movement 310” in Figure 3 above. Ex.1008, 10:25-28. Similarly, Belady teaches taking advantage of lulls in capacity at location Y “by sending data components to location X” as shown by movement 312. Ex.1008, 10:29-32. Each lull, particularly where the lull falls below the actual computational need at a given datacenter, illustrates a time of “change in BTM power availability” at that datacenter. *See* Ex.1008, FIG. 3:



363. Thus, Pelio combined with Chapel and Belady renders obvious “wherein the datacenter control system is configured to transfer the workload from the plurality of computing systems to the second plurality of computing systems in

response to detecting a change in BTM power availability at the plurality of computing systems.”

5. Claim 24

a. [24.0] *The flexible datacenter of claim 23,*


364. See analysis at [1.0]-[1.4] in Challenge 1, [21.0]-[21.2], [22.0]-[22.2], and [23.0]-[23.1].

b. [24.1] *wherein the datacenter control system is a remote master control system that is positioned remotely from the flexible datacenter and the second flexible datacenter.*

365. Pelio combined with Chapel alone, and further combined with Belady, renders obvious wherein the datacenter control system is a remote master control system that is positioned remotely from the flexible datacenter and the second flexible datacenter.

366. First, Pelio combined with Chapel renders obvious a remote master control system that is positioned remotely from the datacenters (flexible datacenters, per Pelio’s teachings). *See* [1.4] (remote controller). It would have been obvious to a POSITA for the local controller and remote controller to be implemented together. This would have been nothing more than the routine implementation of a controller at a remote location, a well-known alternative for implementing controllers. Thus, Pelio combined with Chapel alone renders this limitation obvious.

367. **Second**, [24.1] is further rendered obvious in combination with Belady. Belady teaches that the management functionality 602 “can be implemented by the management system 120 of FIG. 1.” Ex.1008, 12:20-30. And per Figure 1, management system 120 can be “distinct from the individual data centers,” including “by a separate system.” Ex.1008, 8:15-45; Figure 1:



Ex.1008, Fig. 1 (annotated).

368. Thus, Pelio combined with Chapel, alone and further combined with Belady, renders obvious “*wherein the datacenter control system is a remote master control system that is positioned remotely from the flexible datacenter and the second flexible datacenter.*”

6. Claim 26

a. [26.0] *The system of claim 25,*

369. See analysis at [25.0]-[25.6] in Challenge 1.

b. [26.1] *wherein the routing control system is configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems at the second flexible datacenter in response to detecting a change in BTM power availability at the first flexible datacenter.*

370. Pelio combined with Chapel and Belady renders obvious wherein the routing control system is configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems at the second flexible datacenter in response to detecting a change in BTM power availability at the first flexible datacenter. *See* [23.1] and [21.1].

XII. CHALLENGE 4: THE COMBINATION OF PELIO, CHAPEL, AND RASMUSSEN RENDERS OBVIOUS CLAIM 9

A. Summary of Rasmussen

371. Rasmussen is a white paper published by Schneider Electric titled “Impact of Leading Power Factor on Data Center Generator Systems.” Rasmussen has a copyright marking dated 2014. Ex.1010, p. 1. Further, the metadata of Rasmussen indicates that the white paper was created November 12, 2014. Ex.1024.

372. Rasmussen was a white paper published by Schneider Electric. Back at the time which Rasmussen was published, Schneider Electric used www.apc.com as a support website, to which many white papers about datacenter needs to address

challenges in planning, designing, operating, and maintaining datacenter facilities. White papers were a well-known source of information for POSITAs to reference regarding datacenter operations, including power control, power factor control, and other related issues such as those taught by Rasmussen. Indeed, APC was known as a leading supplier of equipment in the datacenter industry at the time that Schneider Electric acquired it in 2007. *See, e.g.*, Ex.1025, p. 1.

373. Further, given APC's reputation in the datacenter industry well before the earliest priority date of the '456 patent, it is my opinion that a POSITA would have been easily able to locate Rasmussen with little effort. APC's white papers were all available and easily locatable at APC's support website. *See* Ex.1022 (providing a listing of white papers, with each identified by an indexed number WP-[number of paper] v[version number]). Rasmussen is white paper 200, locatable on page 3 of the PDF of Ex.1022, identified as version 0. Ex.1022, p. 3. Further, the description of the Rasmussen white paper from the index listing in Ex.1022 matches the Executive Summary of Rasmussen. *See* Ex.1010, p. 1. As a result, it is my opinion that Rasmussen was publicly accessible well before the earliest alleged priority date of the '431 patent, and that a POSITA would have been able to locate Rasmussen with reasonable diligence from the APC white paper website. This is confirmed by the fact that others accessed Rasmussen well before the Critical Date.

See Ex.1023, pp. 2, 6 (citation [6] to the Rasmussen white paper). Rasmussen was not before the Examiner during prosecution of the '456 patent.

374. Rasmussen is related to the problem of “IT loads” exhibiting “a characteristic called leading power factor’,” where the generator source “can tolerate a small amount of leading power factor” but the IT loads imposed by a datacenter on the generator source can affect the power factor over time. Ex.1010, p. 2. Datacenters can cause an associated generator source to trip off line even when “under load conditions that appear to be well within the generator ratings” where “IT load becomes directly supplied by the generator” such as occurs where there is a power path to the load that does not have UPS protection. Ex.1010, p. 2.

375. “[C]apacitive loads” have a “current waveform [that] is out-of-phase because it is actually earlier in time than the voltage waveform,” leading to capacitive loads “said to have ‘leading’ power factor.” Ex.1010, p. 4. While “IT equipment” (e.g., servers) have power supplies that “are intended to eliminate harmonic currents [lagging power factor], [they] were not intended to fix out-of-phase currents.” Ex.1010, p. 4. “It is these out-of-phase reactive currents, that are present in power factor corrected power supplies, that impact generators.” Ex.1010, p. 4. To mitigate the impact such capacitive loads can have on the associated generator, Rasmussen teaches to remove IT load (e.g., one or more servers) from the generator so that “sources of out-of-phase current are removed.” Ex.1010, p. 14.

B. Reasons to Combine Pelio and Chapel with Rasmussen

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

377. First, Rasmussen is in the same field of endeavor of the '456 patent, namely power delivery to computers in a datacenter. Rasmussen describes issues related to power delivery from a generator to a server IT load in a datacenter. Ex.1010, p. 2. Thus, Rasmussen is in the same field of endeavor as the '456 patent and Pelio.

378. Rasmussen is also reasonably pertinent to the problems identified by the '456 patent, such as how to control power at the datacenter and ensure stability for the generator. *See* Ex.1001, 4:24-27; 18:25-31; 18:45-57. For example, Rasmussen teaches controlling power at the datacenter by removing IT devices from the load (e.g., which includes by turning them off). Ex.1010, p. 14. By implementing Ramsussen's teachings, a POSITA would find a solution to the '456 patent-identified problems of controlling power and ensuring stability (by removing IT load from the generator).

379. Second, a POSITA would have been motivated to combine Rasmussen's teachings for correcting power factor issues to protect the associated generator source with Pelio's and Chapel's power delivery control teachings, because doing so would ensure the long-term stability of the generator operation.

For example, the out-of-phase reactive currents caused by the power supplies of servers can create issues when more load from the datacenter (by increasing the IT load of the servers) occurs much later, potentially many years after the datacenter is finished. Ex.1010, p. 9. By implementing Rasmussen's mitigation teachings (e.g., by removing IT devices from the load when necessary), a POSITA would be able to lessen the impact of increasing load at the datacenter which might otherwise cause the generator to exceed its allowed power factor limit. *See* Ex.1010, pp. 9, 14.

380. Third, a POSITA would have had a reasonable expectation of success in combining Rasmussen's teachings with the combined teachings of Pelio and Chapel. The combination of Pelio and Chapel already results in a datacenter container (as taught by Pelio) being controlled by local and/or remote controllers to turn off or on one or more servers at the container (as taught by Chapel), consistent with Rasmussen's teachings of removing IT load when necessary. In fact, all that Rasmussen's teachings add to the combination of Pelio and Chapel is another reason for Chapel's control teachings to act upon (namely, for power factor correction to protect the generator).

381. Fourth, combining Rasmussen's teachings with the combined teachings of Pelio and Chapel would have also been predictable and within a POSITA's skillset. Pelio and Chapel already recognize the benefit of monitoring the power supply and changing the power status of one or more servers in a datacenter.

See Ex.1006, [0002], [0047] (monitoring and controller power delivery). Similarly, Rasmussen teaches removing IT load to implement power factor correction. Ex.1010, p. 14. A POSITA would have recognized that removing load from a generator means that the power supply is controlled by reducing the load. It would have been a trivial matter for a POSITA to monitor the power supply, per Chapel's teachings, and control power delivery to one or more servers (again per Chapel's teachings) in order to correct for power factor (per Rasmussen's teachings). Doing so would provide the benefit of generator stability by preventing the power factor from exceeding an allowed limit for the generator, per Rasmussen. *See* Ex.1010, p. 14.

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

C. Detailed Analysis of Claim

1. Claim 9

a. [9.0] *The flexible datacenter of claim 3,*

383. *See* analysis at [1.0]-[1.4]; [2.0]-[2.1]; and [3.0]-[3.1].

b. [9.1] wherein BTM power availability comprises: local power generation c to power factor correction.

384. Pelio combined with Chapel and Rasmussen renders obvious wherein BTM power availability comprises: local power generation c to power factor correction.

385. First, Pelio recognizes that it is important “to keep the grid in balance,” which would result in times where there is generated electricity that does not meet one or more needs of the grid. *See* Ex.1005, [0032]-[0033].

386. Second, power generated but subject to power factor correction would have been an obvious type of power availability for the combined teachings of Pelio and Chapel to monitor and control. For example, Rasmussen teaches that datacenters can cause an associated power generator to trip off line even when “under load conditions that appear to be well within the generator ratings,” specifically in situations where “IT load becomes directly supplied by the generator” such as occurs where there is a power path to the load that does not have UPS protection. Ex.1010, p. 2. This is caused by the leading power factor that IT loads (e.g., servers) exhibit, which generators typically do not tolerate beyond a certain range. Ex.1010, p. 2.

387. Third, Rasmussen specifically teaches that it was known to control power factor for a local generator, such as a collocated generator as taught in combination with Pelio, by turning off one or more servers. According to Rasmussen, “capacitive loads” have a “current waveform [that] is out-of-phase

because it is actually earlier in time than the voltage waveform,” leading to capacitive loads “said to have ‘leading’ power factor.” Ex.1010, p. 4. While servers (“IT equipment” in Rasmussen, p. 2) have power supplies that “are intended to eliminate harmonic currents [lagging power factor], [they] were not intended to fix out-of-phase currents.” Ex.1010, p. 4. “It is these out-of-phase reactive currents, that are present in power factor corrected power supplies, that impact generators.” Ex.1010, p. 4. A POSITA would have appreciated that this would include collocated generators as taught by Pelio.

388. Recognizing the issue, Rasmussen teaches strategies to mitigate the problem of leading power factor on generator operation. One approach is to **remove load** from the generator, so that “sources of out-of-phase current are removed.” Ex.1010, p. 14. Rasmussen’s teaching of removing IT load from the generator to implement power factor correction is the very thing that the ’456 patent contemplated for achieving “the power factor correction required by the local station.” Ex.1001, 8:33-40; *see also* 9:3-12 (monitored power availability includes power generation subject to power factor correction).

389. In combination, Pelio, Chapel, and Rasmussen further renders obvious controlling power delivery to computing systems based on an operational directive (per Pelio-Chapel, *see* [1.4]) based on the monitoring of the BTM power for power factor correction needs, per Rasmussen’s teachings.

390. Thus, Pelio combined with Chapel and Rasmussen renders obvious “*wherein BTM power availability comprises: local power generation c to power factor correction.*”

XIII. CHALLENGE 5: THE COMBINATION OF PELIO, CHAPEL, AND CORHODZIC RENDERS OBVIOUS CLAIMS 15-16

A. Summary of Corhodzic

391. U.S. Patent No. 8,080,900 to Corhodzic *et al.* (Ex.1011, “Corhodzic”) was filed July 18, 2007, and issued December 20, 2011. Corhodzic is titled “Direct-Coupled IT Load.” Corhodzic was not before the Examiner during prosecution of the ’456 patent.

392. Corhodzic relates to “electrical power distribution in data centers.” Ex.1011, 1:5-8. Power arrives at the datacenter as “a substantially sinusoidal signal (e.g., 50 Hz, 60 Hz) that was generated, transmitted, and distributed by the electric utility.” Ex.1011, 6:28-31. “Substantially unconditioned AC power may, of course, pass through various circuits typically provided by an electric utility.” Ex.1011, 6:35-41. Racks of the datacenter “may be powered by three phase AC power that is delivered to the data center 105 from an electric utility 130.” Ex.1011, 6:46-50.

393. From the three phase AC power, the datacenter distributes “one phase voltage and a neutral line” to each rack. Ex.1011, 7:16-18. Each tray in a rack may receive the same phase voltage line (e.g., phase A voltage), while other racks receive other phases, such as a different rack receiving a phase B voltage and yet another

rack receiving phase C voltage (out of the three-phase voltage input to the datacenter). Ex.1011, 7:21-36.

B. Reasons to Combine Pelio and Chapel with Corhodzic

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

395. First, Corhodzic is in the same field of endeavor of the '456 patent, namely power delivery to computers in a datacenter. Corhodzic teaches that it relates to “electrical power distribution in data centers,” and particularly to providing different phases of AC voltage to different racks in the datacenter. Ex.1011, 1:5-8, 6:46-50, 7:21-36. Thus, Corhodzic is in the same field of endeavor as the '456 patent and Pelio.

396. Corhodzic is also reasonably pertinent. For example, Corhodzic teaches providing “an integrated uninterruptible power system in which only one AC-to-DC rectification occurs between the utility AC grid and the processing circuit (e.g., microprocessor) loads.” Ex.1011, 1:54-58. This is in contrast to prior approaches, which did not use a “central UPS ... to condition power or provide back-up power to data center 104.” Ex.1011, 6:1-3. “Instead, the power” in such prior systems “proceeds to various loads 118 in the data center 104 without any substantial conditioning,” and a “local UPS” for each load. Ex.1011, 6:3-13. Corhodzic addresses this problem by integrating the UPS with the DC load, such that only one

rectification between the power source and the processor occurs, resulting in improved power conversion efficiency. Ex.1011, 2:20-36. By implementing Corhodzic's teachings, a POSITA would understand that datacenters according to the '456 patent benefit from the improved efficiency of power conversion, since in large datacenters the percentage of loss from conversion can result in "millions of dollars of electricity annually." Ex.1011, 1:40-46.

397. Second, a POSITA would have recognized that it was well-known for datacenters to receive 3-phase power from a generating source, and for the datacenter to distribute one or more phases of the received power to different racks (and, thus, the computing systems at those racks) for the computing systems to operate. Corhodzic confirms that this was a well-known fact. A POSITA would have been motivated to turn to Corhodzic for details about receiving and distributing 3-phase power, since Corhodzic teaches receiving 3-phase power from a generator and distributing the phases to different racks. Ex.1011, 6:28-31, 6:46-50, 7:16-18. 7:21-36. It would have therefore been obvious for a POSITA to turn from Pelio and Chapel to supplement with Corhodzic's teachings.

398. Third, a POSITA would have had a reasonable expectation of success in combining Corhodzic's teachings with the combined teachings of Pelio and Chapel. The combination of Pelio and Chapel already results in a datacenter container (as taught by Pelio) receiving power from a generator power source, and

that the computing systems within the datacenter container perform their operations from the power distributed within the datacenter container upon receiving the power from the source, consistent with Corhodzic's teachings of receiving 3-phase power and distributing the phases of power to the different computing systems within the datacenter. Such combination of Corhodzic's power distribution teachings would have been predictable and within a POSITA's skillset, since Pelio and Chapel already recognized that power was distributed to the different computing systems.

399. Finally, a POSITA would have been motivated to turn to Corhodzic's phase distribution teachings, and integrating the UPS with the load (e.g., on a tray with a processor), in order to achieve Corhodzic's benefits. These include, for example, improving "electrical power conversion efficiency ... by providing a modular processing unit with an uninterruptible power system in which on a single AC-to-DC rectification occurs between the utility mains and the microprocessor." Ex.1011, 2:20-24. This would result in "substantial energy savings, which may yield further benefits such as reduced conductor ... requirements and reduced thermal ... loads." Ex.1011, 2:24-28.

400. The combination permits but does not require physical incorporation of elements from Corhodzic into Pelio. Nevertheless, the physical incorporation of Rasmussen's teachings would teach the claim limitations.

C. Detailed Analysis of Claims

1. Claim 15

a. [15.0] *The flexible datacenter of claim 1,*

401. See analysis at [1.0]-[1.4].

**b. [15.1] *wherein the BTM power input system comprises:
an input configured to receive three-phase BTM AC
voltage from the power generation station.***

402. Pelio combined with Chapel and Corhodzic renders obvious wherein the BTM power input system comprises: an input configured to receive three-phase BTM AC voltage from the power generation station.

403. First, Pelio discloses a BTM power input system. See [1.1].

404. Second, Pelio renders obvious that the BTM power input system would be configured to receive three-phase AC voltage from the collocated BTM power source. For example, Pelio teaches that “[i]n alternating current circuits, the inductance and capacitance of the phase conductors can be significant.” Ex.1005, [0016]. These “phase conductors” demonstrate multiple phases, which a POSITA would recognize three phases as a well-known option.

405. Third, Corhodzic demonstrates what a POSITA would have understood and recognized from Pelio’s teachings about multiple-phase conductors. In particular, Corhodzic teaches that “data center 105 receives unconditioned 3-phase AC power from the transmission system 195.” Ex.1011, 10:38-40; *see also* 10:4-11. In combination, therefore, Pelio and Corhodzic teach Pelio’s BTM power

input system receiving three-phase voltage (per Corhodzic), which would be from a BTM source as Pelio teaches.

406. Thus, Pelio combined with Chapel and Corhodzic renders obvious “*wherein the BTM power input system comprises: an input configured to receive three-phase BTM AC voltage from the power generation station.*”

2. Claim 16

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

407. See analysis at [1.0]-[1.4] and [15.0]-[15.1].

b. [16.1] *wherein the BTM power input system is configured to provide a single-phase AC voltage to the datacenter control system.*

408. Pelio combined with Chapel and Corhodzic renders obvious wherein the BTM power input system is configured to provide a single-phase AC voltage to the datacenter control system.

409. First, Pelio discloses a BTM power input system. See [1.1]. Further, Pelio combined with Chapel also teaches a local controller at the datacenter. See [1.4]. In combination, a POSITA would have recognized that the BTM power taught by Pelio would be provided to the local controller as well, since in embodiments of Chapel it is at the datacenter as well. See Ex.1006, [0056]. For example, the power would be provided to the local controller via one or both of Pelio’s “A-side” and “B-side” power feeds. See Ex.1005, [0079].

410. Second, Corhodzic further renders obvious providing a single-phase AC voltage from the BTM power input system. Corhodzic teaches that “one phase voltage and a neutral line are distributed to each rack 125” in a datacenter. Ex.1011, 7:16-17. Corhodzic continues:

As an example, the exemplary tray 110 (shown in magnified detail) in the rack 125A receives a phase A voltage and the neutral line. Each of the trays 110 in the rack 125A receives the same AC input voltage signal, namely the Phase A-to-neutral voltage.

Similarly, each of the trays 110 in the rack 125B receives a Phase B-to-neutral as the AC input voltage signal, and each of the trays 110 in the rack 125C receives Phase C-to-neutral as the AC input voltage signal. In other implementations, different phase voltages may be distributed among the trays 110 in one of the racks 125A-125C, and/or the AC input voltage signal to each of the trays 110 may be line-to-line voltages instead of line-to-neutral voltages. In various embodiments, any practical number of phases (e.g., 1, 2, 3, 4, 5, 6, . . . 12 or more) may be distributed to provide operating power to individual trays 110.

Ex.1011, 7:21-36.

FIG. 2 shows an exemplary power distribution architecture in a system 200, which may be implemented in a large facility with large power requirements, for

example. The system 200 includes a utility AC mains generator 205 to supply AC mains voltage from a utility such as the electric utility 130. The exemplary system 200 also includes two back-up AC generators, including a diesel fuel powered generator 210 and a co-located (e.g., gas-fired generator 215). Power from the generators 205, 210, 215 may be combined and/or selected by the AC switchgear 220, and then delivered to the tray 110 via an AC bus 225. In the event of a fault on the AC mains from the generator 205, the generators 210, 215 may provide a backup AC input voltage signal on the AC bus 225.

Ex.1011, 10:55-67.

411. In combination, therefore, Pelio and Corhodzic teach Pelio's BTM power input system providing single-phase voltage (per Corhodzic), which would be from a BTM source as Pelio teaches. Since the datacenter control system (taught by Chapel) is part of the datacenter, it would have therefore been obvious that the datacenter control system would be one of the systems that would receive the single-phase voltage as taught by Corhodzic.

412. Thus, Pelio combined with Chapel and Corhodzic renders obvious *"wherein the BTM power input system is configured to provide a single-phase AC voltage to the datacenter control system."*

XIV. CONCLUSION

414. 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: 5/5/24


Tajana S. Rosing, Ph.D.