

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ACTIV FINANCIAL SYSTEMS, INC.,

Petitioner,

v.

EXEGY, INC.

Patent Owner.

Case No. IPR2020-01010
U.S. Patent No. 10,062,115

DECLARATION OF DAVID E. TAYLOR

I, David E. Taylor, declare as follows:

I. Overview

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

2. I am a named inventor on U.S. Patent No. 10,062,115 (“the ‘115 Patent”) and submit this declaration in support of the Patent Owner Preliminary Response being filed by Exegy, Inc. in the *inter partes* review of the ‘115 Patent. I am not being compensated for my time in connection with this matter.

II. My Background and Qualifications

3. I earned Bachelor of Science Degrees in Electrical Engineering and Computer Engineering in 1998, Master of Science Degrees in Electrical Engineering and Computer Engineering in 2002, and a Doctor of Science Degree in Computer Engineering in 2004 from Washington University in Saint Louis.

4. I am currently Chief Technology Officer of Exegy, Inc. in St. Louis, MO. Exegy is a global provider of technology and managed services for real-time financial market data, predictive trading signals, and historical financial market data to the financial services industry. Our solutions are considered mission-critical infrastructure by the largest financial firms in the world, including global banks, hedge funds, asset managers, market makers, brokers, proprietary traders, exchanges, dark pools, and Alternative Trading System (ATS) operators.

5. Prior to joining Exegy, I worked at the Applied Research Laboratory (ARL) at Washington University from 1999 to 2005 as a research assistant and later as a visiting assistant professor. Among its many achievements, ARL pioneered the development of high-capacity switches and Internet backbone routers. A derivative startup company, Growth Networks, commercialized the multi-stage switching technology developed at ARL and was acquired by Cisco Systems in 2002. My work at ARL included the development of early prototypes of network processors that were embedded in high-speed switches, the use of firmware and Field-Programmable Gate Array (FPGA) devices to accelerate routing and network processing tasks, and the development of packet classification algorithms that were amenable to implementation in network processors and FPGAs. I also collaborated on research of network processor architectures with the IBM Zurich Research Lab in Zurich, Switzerland in 2002.

6. In 2005, I transferred to Exegy Inc. (formerly Data Search Systems Inc.) to join the team of professors and former graduate students from Washington University in St. Louis who were commercializing hardware-accelerated computing technology. From 2005 through 2008, I served as a system architect and manager of the hardware engineering team. In this role, I led the innovation, design, and development of the integrated circuit designs for hardware-accelerated computing products, including the first Exegy products for processing financial

market data. I collaborated with senior technologists with complementary experience in real-time software and financial market data, many of whom had decades of experience from roles at Bridge Information Systems and Reuters (now Refinitiv). This collaborative work between experts in computer networks, hardware design, real-time software design, and financial market data systems led to the inventions in the '115 Patent and related patents. From 2008 through 2009, I served as Vice President of Engineering, leading all aspects of the hardware-software co-development process for all Exegy products. From 2009 through 2014, I served as Vice President of Product Management, leading product strategy and working closely with senior executives at global financial institutions to integrate Exegy products with their electronic trading infrastructure and define the roadmap for future products. In 2014, I began serving as Chief Technology Officer with general responsibilities for product strategy, intellectual property strategy and management, marketing, and sales engineering.

7. I have authored or co-authored 24 papers in the most prestigious peer-reviewed journals and conferences in the fields of computer architecture, computer communications networks, high performance switches and routers, software-hardware co-design and hardware-accelerated computing, and network packet classification algorithms, which are referenced below:

ClassBench: A Packet Classification Benchmark, David E. Taylor, Jonathan S. Turner, IEEE/ACM Transactions on Networking, Volume 15, Issue 3, 499-511, June 2007.

Longest Prefix Matching using Bloom Filters, Sarang Dharmapurikar, Praveen Krishnamurthy, David E. Taylor, ACM/IEEE Transactions on Networking, Volume 14, Issue 2, 397-409, April 2006.

Survey & Taxonomy of Packet Classification Techniques, David E. Taylor, ACM Computing Surveys, Volume 37, Issue 3, 238-275, September 2005.

Robust Header Compression (ROHC) in Next-Generation Network Processors, David E. Taylor, Andreas Herkersdorf, Andreas Doering, Gero Dittmann, IEEE/ACM Transactions on Networking, Volume 13, Issue 4, 755-768, August 2005.

Models, Algorithms, and Architectures for Scalable Packet Classification, David E. Taylor, doctoral dissertation, Department of Computer Science and Engineering, Washington University in Saint Louis, WUCSE-2004-40, July 2004.

Scalable IP Lookup for Internet Routers, David E. Taylor, John W. Lockwood, Todd Sproull, Jonathan S. Turner, David B. Parlour, IEEE Journal on Selected Areas in Communications, Volume 21, Number 4, May 2003.

Dynamic Hardware Plugins (DHP): Exploiting Reconfigurable Hardware for High-Performance Programmable Routers, David E. Taylor, Jonathan S. Turner, John W. Lockwood, Computer Networks, February 2002, Volume 38, Issue 3, pp. 295-310, Elsevier Science.

A Capstone Computer Engineering Design Course, William D. Richard, David E. Taylor, David M. Zar, IEEE Transactions on Education, November 1999, Volume 42, Number 4, pp. 288-294.

A Scalable Architecture for High-Throughput Regular Expression Pattern Matching, Benjamin C. Brodie, Ron K. Cytron, David E. Taylor, IEEE International Symposium on Computer Architecture (ISCA) 2006, 6/06.

Diversifying the Internet, Jonathan S. Turner, David E. Taylor, IEEE GlobeCom 2005, St. Louis, MO, 11/05.

Scalable Packet Classification using Distributed Crossproducting of Field Labels, David E. Taylor, Jonathan S. Turner, IEEE INFOCOM 2005, Miami, FL, 3/05.

ClassBench: A Packet Classification Benchmark, David E. Taylor, Jonathan S. Turner, IEEE INFOCOM 2005, Miami, FL, 3/05.

System-on-Chip Packet Processor for an Experimental Network Services Platform, David Taylor, Alex Chandra, Yuhua Chen, Sarang Dharmapurikar, John Lockwood, Wenjing Tang, Jonathan Turner, IEEE Globecom 2003, San Francisco, CA, 12/03.

Packet Classification using Extended TCAMs, Ed Spitznagel, David Taylor, Jonathan Turner, IEEE International Conference on Network Protocols (ICNP), November 4-7, 2003, Atlanta, GA, USA.

Longest Prefix Matching using Bloom Filters, Sarang Dharmapurikar, Praveen Krishnamurthy, David E. Taylor, ACM SIGCOMM'03, August 25-29, 2003, Karlsruhe, Germany.

Scalable IP Lookup for Programmable Routers, David E. Taylor, John W. Lockwood, Todd S. Sproull, Jonathan S. Turner, David B. Parlour, IEEE INFOCOM 2002: 21st Annual Joint Conference of the IEEE Computer and Communications Societies, New York, NY, 6/02.

Dynamic Hardware Plugins in an FPGA with Partial Run-time Reconfiguration, Edson L. Horta, John W. Lockwood, David E. Taylor, David Parlour, Design Automation Conference (DAC), New Orleans, LA, 6/02.

Design of a High Performance Dynamically Extensible Router, Sumi Choi, John Dehart, Ralph Keller, Fred Kuhns, John Lockwood, Prashanth Pappu, Jyoti Parwatikar, W. David Richard, Ed Spitznagel, David Taylor, Jonathan Turner, and Ken Wong, Proceedings of the DARPA Active Networks Conference and Exposition, 5/02.

Control and Configuration Software for a Reconfigurable Networking Hardware Platform, Todd S. Sproull, John W. Lockwood, David E. Taylor, FCCM'02: 2002 IEEE Symposium on Field-Programmable Custom Computing Machines, Napa, CA, 4/02.

Dynamic Hardware Plugins (DHP): Exploiting Reconfigurable Hardware for High-Performance Programmable Routers, David E. Taylor, Jonathan S. Turner, John W. Lockwood, IEEE OPENARCH 2001, Anchorage, AK, 4/01. Selected for publication in Computer Networks.

Reprogrammable Network Packet Processing on the Field Programmable Port Extender (FPX), John W. Lockwood, Naji Naufel, David E. Taylor, Jon S. Turner, FPGA 2001: Ninth ACM International Symposium on Field-Programmable Gate Arrays, Monterey, CA, 2/01.

Development of an FPGA-Based South Bridge using Spectrum and ModelSim, William D. Richard, David E. Taylor, Mentor Graphics User's Group International Conference, Portland, OR, 10/00.

Field Programmable Port Extender (FPX) for Distributed Routing and Queueing, John W. Lockwood, Jon S. Turner, David E. Taylor, FPGA 2000: Eighth ACM International Symposium on Field-Programmable Gate Arrays, Monterey, CA, 2/00.

Developing an FPGA Workflow to Ease Novice and Experienced Designers, David E. Taylor, David M. Zar, Mentor Graphics User's Group International Conference, Portland, OR, 10/98. Awarded Best Paper for University/Research Track.

8. As a recognized expert, I have served as a reviewer of scholarly articles for the following refereed journals and conferences: IEEE/ACM Transactions on Networking, IEEE Journal on Selected Areas in Communications, IEEE Micro, IEEE Communications Letters, Computer Networks (Elsevier), ACM Sigcomm, IEEE Infocom, IEEE Globecom, and IEEE International Conference on Communications.

9. Further, I am a named inventor on 68 issued patents and many active patent applications relating to high speed processing of financial market data.

10. Through my professional work, I am an expert in the field of high speed processing of financial market data, which is the field of the '115 Patent. I have been invited to speak and sit on expert panels by the United States Securities and Exchange Commission (SEC), Financial Industry Forum (FIF), Securities Technology Assessment Center (STAC), Options Insider Radio Network, and The Trading Show.

III. Claim Interpretation

11. In my opinion, a price engine “configured to (1) access a plurality of price point records based on the streaming limit order event data, and (2) compute updated price point data based on the accessed price point records and the streaming limit order event data” should be defined to mean price point records in memory, where each price point record contains at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, and (c) an order count indicating the total number of orders for the financial instrument at that price.

12. My opinion is supported by the specification of the '115 Patent, which discloses that “at minimum, the set of price point records preferably contain the price, volume (sum of order sizes at that price, which can be referred to as the price volume), and order count (total number of orders at that price).” '115 Patent at column 18, lines 10-16. Additional support can be found in Figures 21(a)-(b),

which illustrate price point records 2100, 2150 as containing various data fields including price, volume, and order count data fields. For example, price point record 2100 of Figure 21(a) has: (1) “a price field 2110” associated with a financial instrument, (2) “a volume field 2112, identifying the volume of shares across all limit orders for the financial instrument (see the symbol field 2104) on the regional exchange (see the GEID field 2106) at the price identified in the price field 2110,” and (3) “a count field 2114, which comprises a count of how many limit orders make up the volume 2112.” ‘115 Patent at column 19, lines 15-30.

13. Similarly, the price point record 2150 of Figure 21(b) has: (1) “a price field 2158,” (2) “a volume field 2160,” and (3) “a count field 2162.” ‘115 Patent at column 19, lines 52-67. The ‘115 Patent also teaches that, although shown in Figures 21(a)-(b) as illustrating a specific number of defined data fields, the “price point records 2100 and 2150 can be configured to have more or fewer and/or different data fields.” ‘115 Patent at column 20, lines 22-24.

14. Nothing in the ‘115 specification, however, teaches or suggests removing the “price field,” the “volume field,” or the “count field” from the price point records 2100, 2150.

15. The definition of the subject limitation is also informed by the fact that the ‘115 specification teaches that price point records are in memory. Data structures are “used to store price point records.” ‘115 Patent at column 19, lines 4-

7. For example, the price point records may be stored on a single shared memory or across multiple memories. ‘115 Patent at column 20, lines 46-66, Figs. 17-18. The price point records are then retrieved from memory based on streaming limit order events. ‘115 Patent at column 20, lines 61-66.

16. The above evidence demonstrates that “price point records” means price point records in memory, where each price point record contains at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, and (c) an order count indicating the total number of orders for the financial instrument at that price.

IV. Parsons Does Not Render the Challenged Claims Obvious

17. As I discuss below, there are significant differences between the teachings of Parsons and the claimed subject matter of the ‘115 Patent. In general, Parsons requires longer compute times and redundant computations as compared to the approach described and claimed in the ‘115 Patent because Parsons requires price aggregations be generated from limit order records each and every time a price aggregate view is requested. As a result of these differences, the ‘115 Patent claims an approach that significantly reduces the work required to provide price aggregate data for a price-aggregated view.

A. Ground 1A of the Petition Fails Because There are no Price Point Records in Parsons

18. A number of teachings and disclosures in Parsons are key to understand the differences between Parsons and the challenged claims in this proceeding. First, Parsons discloses an order book cache (OBC) for maintaining a cache of financial instrument records. Parsons at ¶96. Second, each financial instrument record “consists of an array of sub-records that define individual price or order entries for that financial instrument.” Parsons at ¶96. Third, Parsons’ “OBC includes the ability to generate” a price-aggregated view to show orders at the same price point “aggregated together to create entries that are indicative of the total number of orders available at each price point.” Parsons at ¶¶71, 96.

19. Claim 43 requires “at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, **and** (c) an order count indicating the total number of orders for the financial instrument at that price.” Parsons’ sub-records defining price entries do not disclose this limitation. Instead, Parsons only discloses sub-records defining price entries that are used for a sort order. Parsons at ¶96. As a result, Parsons fails to disclose sub-elements (b) and (c) above.

20. In addition, Parsons does not disclose or suggest the aggregated data shown in its generated price-aggregated view is stored in memory. To that end, Parsons’ price-aggregated view does show “the total number of orders available at

each price point” (Parsons at ¶¶71, 96), but does not disclose or suggest the “total number of orders available at each price point” is a stored entry in its sub-record.

21. Instead, Parsons teaches generating a view showing “the total number of orders available at each price point” by retrieving records for a financial instrument, grouping order entries in that record by price point, and then aggregating various fields that share the same price point to compute the data for the price-aggregated view. Parsons at ¶¶71, 96. For each price aggregate view, each of these steps are repeated. Parsons at ¶¶71, 96. Thus, the aggregated data shown in Parsons’ price-aggregated views is repeatedly generated, not stored in memory.

22. Thus, Parsons fails to disclose price point records in memory, where each price point record contains at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, and (c) an order count indicating the total number of orders for the financial instrument at that price.

B. Ground 1A of the Petition Also Fails Because a Person Having Ordinary Skill in the Art Would Not Have Been Motivated to Access and Update Sub-Records in Parson’s OBC Using the LVC Architecture

23. A person having ordinary skill in the art would not have been motivated to access and update sub-records in Parsons’ OBC using the last value cache (“LVC”) architecture because the OBC and LVC are performing

fundamentally different tasks. To that end, Parsons' OBC performs sorting functions. In contrast, Parsons' LVC does perform sorting functions. As described in paragraphs 71 and 96 of Parsons, Parsons' OBC maintains a sorted listing of limit order records for each financial instrument. Each limit order record corresponds to one limit order resting on the market for that financial instrument, and the position of a limit order record in each sorted listing, for a given financial instrument, changes when limit order records are added or removed from that sorted listing. Parsons at ¶96. In view of this arrangement, there is a one-to-many relationship between each financial instrument and order records.

24. By contrast, Parsons' LVC performs the function of looking up a field to update a record representing the current state of a financial instrument. Parsons at ¶¶62, 95. Thus, Parsons' LVC maintains a one-to-one relationship between each financial instrument and an order record. Nothing in Parsons discloses or suggests that its LVC maintains a sorted listing of records, or a one-to-many relationship between each financial instrument and order records. A person having ordinary skill in the art would not have had any reason to use the LVC architecture in Parsons' OBC (that requires sorting and maintaining one-to-many relationship between each financial instrument and order records) because Parsons' LVC does not perform sorting or maintaining one-to-many relationship between each financial instrument and order records.

C. Ground 1A of the Petition Also Fails Because a Person Having Ordinary Skill in the Art Would Not Have Expected Petitioner's Modified OBC to Operate as Intended

25. As explained above, Parsons' OBC and LVC perform fundamentally different tasks. As such, a person having ordinary skill in the art would have no reasonable expectation of success in using Parsons' LVC architecture in Parsons OBC that is required to perform specific fundamental tasks, none of which are performed by Parsons' LVC architecture.

D. Ground 1B of the Petition Fails Because Parsons Does Not Disclose Price Point Records

26. Parsons discloses an LVC that maintains a cache of financial instrument records. Parsons at ¶95. Among the types of fields included in the records are "last trade price, last trade size, last trade time, best bid price, best bid size, best bid time, best ask price, best ask size, best ask time, total trade volume, daily change, tick direction, price direction, high trade price, high price time, low trade price, low price time, and close price," and derived fields such as "total trade volume at bid, total trade volume at ask, traded value, traded value at bid, traded value at ask, and volume-weighted average price (VWAP)." Parsons at ¶126.

27. The fields associated with Parsons' LVC are not equivalent to "price point records," where each record contains the required limitations of claim 43: "at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, and (c) an

order count indicating the total number of orders for the financial instrument at that price.”

28. The “price point records” of claim 43 require an order count indicating a total number of orders for a particular financial instrument at a particular price. None of the fields of Parsons relates to this requirement. Instead, “last trade price, best bid price, best ask price, high trade price, low trade price, and close price” are fields relating to a single action, such as an executed trade, or an updated quote containing the best bid or ask price, at a particular price. Parsons simply does not teach or disclose an aggregated number of orders for a particular financial instrument at each unique price.

29. In addition, “total trade volume at bid” and “total trade volume at ask” are fields quantifying the aggregate size of executed trades occurring at a price equal to the best bid price and the best ask price at the time of the trade execution. The “total trade volume at bid” and “total trade volume at ask” fields do not relate to an order count indicating a total number of resting limit orders for a particular financial instrument at a particular price. Thus, the fields associated with Parsons’ LVC do not include fields relating to an order count indicating a total number of orders for a particular financial instrument at a particular price as required by Claim 43’s “price point records.”

30. In summary, it is my opinion that Parsons' fields for "total trade volume at bid," "total trade volume at ask," "last trade price, best bid price, best ask price, high trade price, low trade price, and close price" are not equivalent to the claimed "price point records," each containing "at least (a) a price associated with a financial instrument, (b) a price volume indicating a sum of order sizes for the financial instrument at that price, and (c) an order count indicating the total number of orders for the financial instrument at that price"

E. Ground 1B of the Petition Also Fails Because Parsons Does Not Disclose the Claimed "Price Engine"

31. The '115 Patent makes clear the claimed "price engine" is not equivalent to an LVC such as Parsons' LVC. This is clear because the '115 Patent discloses an LVC that is separate from a price engine. Figures 29 and 30 of the '115 Patent illustrate pipelines including "an Order Normalization and Price Aggregation (ONPA) module" having an order engine and a price engine, and a value cache update (VCU) module 2900 for updating "a Last Value Cache (LVC) data structure." '115 Patent at column 17, lines 52-55, column 18, lines 6-12, column 20, line 61 to column 21, line 3, column 25, lines 6-20, and column 28, lines 19-34. The VCU module 2900 updates its LVC data structure based on an output of the ONPA module (which includes the price engine). '115 Patent at column 28, lines 22-34 and Figures 29-30.

32. Additionally, the ‘115 Patent makes clear that the VCU module and LVC data structure are described in Parsons. ‘115 Patent at column 28, lines 46-49. Thus, one skilled in the art would not equate the claimed “price engine” to Parsons’ LVC.

F. Ground 1B of the Petition Fails Because a Person Having Ordinary Skill in the Art Would Not Have Been Motivated To Implement Parsons’ OBC and LVC in Parallel

33. I understand that the Petition states that Parsons’ OBC includes the claimed “order engine” and Parsons’ LVC includes the claimed “price engine.” The Petition also admits that Parsons does not disclose accessing and updating records in the OBC and LVC in parallel. The Petition is incorrect, however, in arguing that a person having ordinary skill in the art would have been motivated to implement Parsons’ OBC and LVC so that their operations are performed in parallel.

34. The modified system shown in the Petition would result in sending the same input messages to Parsons’ OBC and LVC for processing. A person having ordinary skill in the art would have no reason to implement Parsons’ OBC and LVC because Parsons’ OBC and LVC process different types of input messages. To that end, Parsons’ OBC consumes input messages relating to multiple entries showing each outstanding limit order to buy and sell (bids and offers) resting on an exchange for each financial instrument. Parsons at ¶¶71, 96. Those skilled in the

art commonly refer to the types of input messages consumed by Parsons' OBC as "Level 3" messages.

35. Parsons' LVC consumes a different type input message than the OBC. Parsons' LVC consumes input messages relating to a single entry showing a current state to buy and sell (e.g., a best bid, a best offer, a best price level) resting on an exchange for each financial instrument. Parsons at ¶¶62, 126. Those skilled in the art commonly refer to the types of input messages consumed by Parsons' LVC as "Level 1" messages.

36. If Parsons' OBC and LVC are implemented so that their operations are performed in parallel as suggested by the Petition, the OBC and LVC would both process needless and unnecessary input messages.

37. A person having ordinary skill in the art would not have been motivated to implement Parsons' OBC and LVC in a way to cause needless and unnecessary processing.

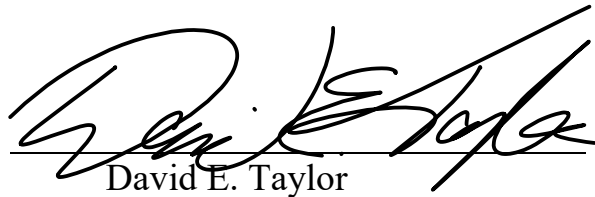
G. Ground 1B of the Petition Also Fails Because a Person Having Ordinary Skill in the Art Would Not Have Expected Petitioner's Modified OBC and LVC to Operate as Intended

38. As explained above, Parsons' LVC consumes a different type input message than the OBC. However, implementing Parsons' OBC and LVC as suggested by the Petition would result in sending the same input messages to Parsons' OBC and LVC for processing. This leads to the OBC and LVC

processing needless and unnecessary input messages. As such, because the modified OBC and the LVC in Parsons would process needless and unnecessary input messages, a person having ordinary skill in the art would not have had a reasonable expectation of success in modifying Parsons' OBC and LVC so that their operations are performed in parallel as suggested by the Petition.

The matters stated in this declaration are true and accurate to the best of my personal knowledge. Pursuant to 28 U.S.C. § 1746, I declare under penalty of perjury under the laws of the United States that the foregoing is true and correct.

Dated: March 1, 2021



David E. Taylor