

UNITED STATES PATENT AND TRADEMARK OFFICE

---

BEFORE THE PATENT TRIAL AND APPEAL BOARD

---

CROWDSTRIKE, INC.,

Petitioner

v.

SKYSONG INNOVATIONS, LLC,

Patent Owner

---

*Inter Partes* Review Case No. IPR2025-01170

U.S. Patent No. 11,892,897

**DECLARATION OF NICHOLAS BAMBOS, Ph.D.**

## TABLE OF CONTENTS

|                                                                                                                                                                                                                                          |           |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| <b><i>I. INTRODUCTION</i></b> .....                                                                                                                                                                                                      | <b>8</b>  |
| <b>A. Educational Background and Professional Experience</b> .....                                                                                                                                                                       | <b>8</b>  |
| <b><i>II. Methodology; Materials Considered</i></b> .....                                                                                                                                                                                | <b>11</b> |
| <b><i>III. Legal Framework</i></b> .....                                                                                                                                                                                                 | <b>14</b> |
| <b><i>IV. Opinion</i></b> .....                                                                                                                                                                                                          | <b>21</b> |
| <b>A. The '897 Patent</b> .....                                                                                                                                                                                                          | <b>21</b> |
| <b>B. Claim Construction</b> .....                                                                                                                                                                                                       | <b>24</b> |
| <b>C. Level of a Person of Ordinary Skill in the Art</b> .....                                                                                                                                                                           | <b>24</b> |
| <b>D. Background of the Technology</b> .....                                                                                                                                                                                             | <b>26</b> |
| 1. Software Instructions for Machine Learning Models are Executed by<br>Computer Processors .....                                                                                                                                        | 26        |
| 2. Social Measures Analysis Between Users on Hacker Forums .....                                                                                                                                                                         | 28        |
| 3. Dark Web Forums Include Hacker Communications .....                                                                                                                                                                                   | 30        |
| 4. Predicting the Likelihood Software Vulnerabilities Will be Exploited<br>Before Disclosure .....                                                                                                                                       | 31        |
| <b><i>V. Opinion – Ground 1</i></b> .....                                                                                                                                                                                                | <b>32</b> |
| <b>A. Obvious Over Sabottke</b> .....                                                                                                                                                                                                    | <b>32</b> |
| 1. Overview of Sabottke .....                                                                                                                                                                                                            | 32        |
| 2. Claim 1(Pre): A method for assessing a likelihood of exploitation of<br>software vulnerabilities, comprising: .....                                                                                                                   | 34        |
| 3. Claim 1(a): utilizing a processor in operable communication with at least<br>one memory for storing instructions that are executed by the processor to<br>perform operations, including: .....                                        | 36        |
| 4. Claim 1(b): accessing a plurality of datasets associated with a<br>predetermined set of data sources, the plurality of datasets including training<br>data comprising hacker communications; .....                                    | 38        |
| 5. Claim 1(c): accessing features from the plurality of datasets that include<br>measures computed from social connections of users posting hacking-related<br>content; .....                                                            | 40        |
| 6. Claim 1(d): applying learning algorithms to the training data to generate<br>classification models that are configured to predict class labels defining a<br>likelihood of exploitation of respective software vulnerabilities; ..... | 42        |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 7. Claim 1(e): accessing one or more features associated with a software vulnerability; and; .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 44 |
| 8. Claim 1(f): computing, by applying the one or more features to the classification model, a class label defining one or more values defining a likelihood of exploitation associated with the software vulnerability.....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 45 |
| 9. Claim 1(f)(1): wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data. ....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 47 |
| 10. Claim 5 - The method of claim 1, wherein certain features correspond to a known vulnerability obtained from the plurality of datasets.....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 52 |
| 11. Claim 9 - The method of claim 1, further comprising computing mutual information from the plurality of datasets informative as to what information a given feature provides about another feature. ....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 52 |
| 12. Claim 10 - The method of claim 1, further comprising: detecting, from the plurality of datasets, vulnerabilities that appear before an associated exploitation date. ....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 53 |
| 13. Claim 14 - A computing device, configured via machine learning to apply a learned function to data associated with a software vulnerability to estimate a likelihood of exploitation of the software vulnerability, the learned function associated with a prediction model derived from at least one machine learning algorithm and a plurality of datasets associated with software vulnerabilities, the plurality of datasets including information associated with discussion of vulnerabilities by users posting hacking-related content, wherein the likelihood of exploitation predicts an actual exploitation of the software vulnerability before disclosure based on hacker communications from the plurality of datasets. .... | 55 |
| 14. Claim 15 - The computing device of claim 14, wherein the prediction model is at least one classification model that outputs from features of the software vulnerability a score indicative of the likelihood of exploitation....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 56 |
| 15. Claim 18(Pre): A method of prioritizing vulnerabilities using cyber threat intelligence, comprising: .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 57 |
| 16. Claim 18(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations including: .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 57 |
| 17. Claim 18(b): accessing a plurality of datasets associated with a predetermined set of data sources, at least a portion of the plurality of datasets defining training data including hacker communications; .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 57 |
| 18. Claim 18(c): applying learning algorithms to the training data to generate a predictive model configured to predict a likelihood of exploitation of respective software vulnerabilities; .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 58 |

|                                                                                                                                                                                                                                                                                                                                                                                                       |           |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------|
| 19. Claim 18(d): accessing one or more features associated with a software vulnerability; and.....                                                                                                                                                                                                                                                                                                    | 58        |
| 20. Claim 18(e): “computing, by applying the one or more features to the predictive model, one or more values defining a likelihood of exploitation associated with the software vulnerability, wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data.” ..... | 58        |
| <b>VI. Opinion – Ground 2 .....</b>                                                                                                                                                                                                                                                                                                                                                                   | <b>59</b> |
| <b>A. Obvious to Combine <i>Smyth</i> and <i>Phillips</i> .....</b>                                                                                                                                                                                                                                                                                                                                   | <b>59</b> |
| 1. Overview of <i>Smyth</i> .....                                                                                                                                                                                                                                                                                                                                                                     | 59        |
| 2. Overview of <i>Phillips</i> .....                                                                                                                                                                                                                                                                                                                                                                  | 62        |
| 3. Claim 1(Pre): A method for assessing a likelihood of exploitation of software vulnerabilities, comprising: .....                                                                                                                                                                                                                                                                                   | 64        |
| 4. Claim 1(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations, including: .....                                                                                                                                                                                                           | 64        |
| 5. Claim 1(b): accessing a plurality of datasets associated with a predetermined set of data sources, the plurality of datasets including training data comprising hacker communications; .....                                                                                                                                                                                                       | 66        |
| 6. Claim 1(c): accessing features from the plurality of datasets that include measures computed from social connections of users posting hacking-related content; .....                                                                                                                                                                                                                               | 72        |
| 7. Claim 1(d): applying learning algorithms to the training data to generate classification models that are configured to predict class labels defining a likelihood of exploitation of respective software vulnerabilities; .....                                                                                                                                                                    | 78        |
| 8. Claim 1(e): accessing one or more features associated with a software vulnerability; and.....                                                                                                                                                                                                                                                                                                      | 82        |
| 9. Claim 1(f): computing, by applying the one or more features to the classification model, a class label defining one or more values defining a likelihood of exploitation associated with the software vulnerability,.....                                                                                                                                                                          | 84        |
| 10. Claim 1(f)(1): wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data. ....                                                                                                                                                                                | 87        |
| 11. Claim 2: The method of claim 1, further comprising generating a plurality of estimation outputs based on the one or more values to derive an overall quantitative score.....                                                                                                                                                                                                                      | 91        |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 12. Claim 3: The method of claim 1, wherein the plurality of datasets include vulnerability data for vulnerabilities that are publicly disclosed and obtaining exploits data for exploits that were used in real world attacks.....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 92  |
| 13. Claim 5: The method of claim 1, wherein certain features correspond to a known vulnerability obtained from the plurality of datasets.....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 95  |
| 14. Claim 10: The method of claim 1, further comprising: detecting, from the plurality of datasets, vulnerabilities that appear before an associated exploitation date. ....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 96  |
| 15. Claim 13: The method of claim 1, further comprising vectorizing text features derived from textual content of the plurality of datasets using a predetermined natural language process (NPL). ....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 97  |
| 16. Claim 14: A computing device, configured via machine learning to apply a learned function to data associated with a software vulnerability to estimate a likelihood of exploitation of the software vulnerability, the learned function associated with a prediction model derived from at least one machine learning algorithm and a plurality of datasets associated with software vulnerabilities, the plurality of datasets including information associated with discussion of vulnerabilities by users posting hacking-related content, wherein the likelihood of exploitation predicts an actual exploitation of the software vulnerability before disclosure based on hacker communications from the plurality of datasets. .... | 98  |
| 17. Claim 15 - The computing device of claim 14, wherein the prediction model is at least one classification model that outputs from features of the software vulnerability a score indicative of the likelihood of exploitation...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 100 |
| 18. Claim 17: The computing device of claim 14, wherein the computing device is further configured to apply the learned function to data associated with each of a plurality of new software vulnerabilities and compute a respective likelihood of exploitation for each of the plurality of new software vulnerabilities. ....                                                                                                                                                                                                                                                                                                                                                                                                             | 102 |
| 19. Claim 18(Pre): A method of prioritizing vulnerabilities using cyber threat intelligence, comprising: .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 103 |
| 20. Claim 18(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations including: .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 103 |
| 21. Claim 18(b): accessing a plurality of datasets associated with a predetermined set of data sources, at least a portion of the plurality of datasets defining training data including hacker communications; .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 103 |
| 22. Claim 18(c): applying learning algorithms to the training data to generate a predictive model configured to predict a likelihood of exploitation of respective software vulnerabilities; .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 104 |

|                                                                                                                                                                                                                                                                                                                                                                                                    |            |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| 23. Claim 18(d): accessing one or more features associated with a software vulnerability; and.....                                                                                                                                                                                                                                                                                                 | 104        |
| 24. Claim 18(e): computing, by applying the one or more features to the predictive model, one or more values defining a likelihood of exploitation associated with the software vulnerability, wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data. .... | 104        |
| <b>VII. Opinion – Ground 3 .....</b>                                                                                                                                                                                                                                                                                                                                                               | <b>105</b> |
| <b>A. Obvious to Combine <i>Smyth, Phillips, and Nunes</i>.....</b>                                                                                                                                                                                                                                                                                                                                | <b>105</b> |
| 1. Overview of Nunes.....                                                                                                                                                                                                                                                                                                                                                                          | 105        |
| 2. Claim 4: The method of claim 3, further comprising: aligning the exploits data with the vulnerability data; and cleaning the exploits data of noise and predetermined portions of the exploits data that is irrelevant to associated software vulnerabilities. ....                                                                                                                             | 106        |
| 3. Claim 16: The computing device of claim 14, wherein the predictive model includes a Random Forest (RF) method including multiple decision tree predictors applied in combination to classify the software vulnerability which is used to estimate the likelihood of exploitation. ....                                                                                                          | 110        |
| <b>VIII. Opinion – Ground 4 .....</b>                                                                                                                                                                                                                                                                                                                                                              | <b>112</b> |
| <b>A. Obvious to Combine <i>Smyth, Phillips, and Kuperman</i>.....</b>                                                                                                                                                                                                                                                                                                                             | <b>112</b> |
| 1. Overview of Kuperman .....                                                                                                                                                                                                                                                                                                                                                                      | 112        |
| 2. Claim 6: The method of claim 1, further comprising testing the classification models by applying additional training data and one or more algorithms and evaluation metrics to optimize the classification models until the classification models compute the likelihood of exploitation according to a predefined error rate. ....                                                             | 113        |
| <b>IX. Opinion – Ground 5 .....</b>                                                                                                                                                                                                                                                                                                                                                                | <b>117</b> |
| <b>A. Obvious to Combine <i>Smyth, Phillips, and Marin</i>.....</b>                                                                                                                                                                                                                                                                                                                                | <b>117</b> |
| 1. Overview of Marin .....                                                                                                                                                                                                                                                                                                                                                                         | 117        |
| 2. Claim 7: The method of claim 1, further comprising vectorizing text features derived from the plurality of datasets using term frequency-inverse document frequency to create a vocabulary of associated words. ....                                                                                                                                                                            | 118        |
| 3. Claim 13: The method of claim 1, further comprising vectorizing text features derived from textual content of the plurality of datasets using a predetermined natural language process (NPL). ....                                                                                                                                                                                              | 120        |
| <b>X. Opinion – Ground 6 .....</b>                                                                                                                                                                                                                                                                                                                                                                 | <b>121</b> |

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |            |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <b>A. Obvious to Combine <i>Smyth, Phillips, and Bozorgi</i> .....</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>121</b> |
| 1. Overview of Bozorgi .....                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | 121        |
| 2. Claim 8: The method of claim 1, further comprising: sorting vulnerabilities associated with the plurality of datasets according to time; training the classification model using the training data, the training data defining a first subset of the plurality of datasets associated with a predetermined period of time; and testing the classification model using a second subset of the plurality of datasets associated with the predetermined period of time. ....                                                                                                           | 122        |
| 3. Claim 12: The method of claim 1, further comprising: accessing one or more features indicative of temporal connections between at least two of: a time associated with discussion of a vulnerability by users posting hacking-related content at a web forum prior to disclosure to a public vulnerability database; a time associated with disclosure of the vulnerability to the public vulnerability database; and a time associated with exploitation of the vulnerability as obtained through exploits data associated with real-world exploitation of the vulnerability. .... | 125        |
| <b>XI. CONCLUSION.....</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>128</b> |

I, Nicholas Bambos, Ph.D., declare as follows:

## **I. INTRODUCTION**

1. I am over the age of 21 and am competent to make this declaration. I make this declaration based on facts and matters within my own knowledge and on information provided to me by others, and, if called as a witness, I could and would competently testify to the matters set forth herein.

### **A. Educational Background and Professional Experience**

2. My qualifications are stated more fully in my curriculum vitae, which is attached as Appendix A. Below is a summary of my education, work experience, and other qualifications.

3. I am a R. Weiland Professor in the School of Engineering at Stanford University, having a joint appointment in the Department of Electrical Engineering and the Department of Management Science & Engineering. I have also served as the Fortinet Founders Department Chairman of the Management Science & Engineering Department from January 2016 to August 2020.

4. Before joining Stanford as an Associate Professor in 1996, I was (since 1989) an Assistant Professor and later a tenured Associate Professor in the Electrical Engineering department of the University of California at Los Angeles (UCLA).

5. I received my Ph.D. from the University of California at Berkeley (1989) in Electrical Engineering and Computer Sciences (EECS). Also, from U.C.

Berkeley, I received a M.S. in EECS (1987) and a M.A. in Mathematics (1989). I graduated in Electrical Engineering from the National Technical University of Athens-Greece (1984) with first class honors. While at U.C. Berkeley, I had been a U.C. Regents Fellow, a David Gale Fellow, an Earl Anthony Fellow, and an EECS Departmental Scholar.

6. At Stanford University, I head the Computer Systems & Networks Performance Engineering Lab, and my research group is working on high-performance designs of computer systems and networks. I have researched various high-performance design aspects of computing & networking technologies and have published numerous papers on such topics.

7. Since 2016, I have been a R. Weiland (chaired) Professor in the Stanford School of Engineering. From 1999 to 2005, I was the Director of the Stanford Networking Research Center project. I have held the Cisco Systems Faculty Development Chair (1999-2003) in computer networking at Stanford University and have won the IBM Faculty Development Award (2002) for research in performance engineering of computer systems and networks. I have also been the recipient of the National Young Investigator Award of the National Science Foundation (1992).

8. I have served as editor of various research journals, as technical reviewer for numerous computing & networking research journals, on various technical panels for the National Science Foundation, etc.

9. For over 40 years, I have done research in and taught computing and networking technology concepts and design principles (at Stanford since 1996, at UCLA during 1989-96, at UC Berkeley during 1984-89). I have supervised over 40 Ph.D. student doctoral thesis, who have later been in leadership positions in academia and in industry.

10. I have taught advanced graduate classes in computing and networking, including “Performance Engineering of Computer Systems & Networks” and “Security & Risk in Computer Systems.”

11. I have given several keynote and/or plenary research talks internationally, including the keynote talk “A Risk Management View to Information Security” in 2010 at the IEEE Conference on Decision and Game Theory for Security in Berlin, Germany. For several years starting in 2010 I was on the founding steering board of the annual IEEE Conference on Decision and Game Theory for Security.

12. Based on my technical expertise in system architecture and high-performance engineering of computer systems and networks, I have published over 300 peer-reviewed research papers (including in network security), have given

numerous technical talks world-wide, and have supervised numerous doctoral theses and research projects at Stanford and UCLA.

13. I am a named inventor on eight patents. I have also served on technical advisory boards or co-founder of technology start-up companies. I have served as a technical expert witness in numerous patent litigation cases.

14. My full qualifications and experience are set forth in my Curriculum Vitae, which is attached as Appendix A to this declaration.

15. My opinions are based on my years of education, research, and experience, as well as my study of relevant materials. In forming my opinions, I have also considered the materials identified in this declaration and in the Petition.

16. I am being compensated for my time at my ordinary hourly rate of \$525. My compensation is not dependent on the outcome of these proceedings or the content of my opinions. To the best of my knowledge, I have no financial interest in either party or in the outcome of this proceeding.

## **II. METHODOLOGY; MATERIALS CONSIDERED**

17. I have relied upon my education, knowledge and experience with the technology of computer systems and networks, as well as the other materials as discussed in this declaration in forming my opinions.

18. For this work, I have been asked to review U.S. Patent No. 11,892,897 (“’897 Patent”) (Ex. 1001) including the specification and claims, and the ’897

Patent's prosecution history ("’897 File History") (Ex. 1002) and to assess the validity of the ’897 Patent.

19. In developing my opinions relating to the ’897 Patent, I have considered at least the materials cited or discussed herein, including those listed below.

|                     |                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Exhibit 1001</b> | U.S. Patent No. 11,892,897 (" <i>’897 Patent</i> ")                                                                                                                                                                                   |
| <b>Exhibit 1002</b> | File History of the ’897 Patent (" <i>’897 File History</i> ")                                                                                                                                                                        |
| <b>Exhibit 1004</b> | Carl Sabottke et. al., <u>Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits</u> , Proceedings of the 24 <sup>th</sup> USENIX Security Symposium (2015) (" <i>Sabottke</i> ") |
| <b>Exhibit 1005</b> | U.S. Patent No. 10,264,009 to Smyth et al. (" <i>Smyth</i> ")                                                                                                                                                                         |
| <b>Exhibit 1006</b> | Elizabeth Phillips et al., <u>Extracting Social Structure from DarkWeb Forums</u> , SOTICS 2015: The Fifth International Conference on Social Media Technologies, Communication, and Informatics, (2015) (" <i>Phillips</i> ")        |
| <b>Exhibit 1007</b> | Eric Nunes et al., <u>Darknet and Deepnet Mining for Proactive Cybersecurity Threat Intelligence</u> , Proceedings of the 2016 IEEE Conference on Intelligence and Security Informatics (" <i>Nunes</i> ")                            |
| <b>Exhibit 1008</b> | U.S. Patent No. 10,652,254 to Kuperman et al. (" <i>Kuperman</i> ")                                                                                                                                                                   |
| <b>Exhibit 1009</b> | Ericsson Marin et al., <u>Product Offerings in Malicious Hacker Markets</u> , Proceedings of the 2016 IEEE Conference on Intelligence and Security Informatics (" <i>Marin</i> ")                                                     |
| <b>Exhibit 1010</b> | Mehran Bozorgi et al., <u>Beyond Heuristics: Learning to Classify Vulnerabilities and Predict Exploits</u> , Proceedings of the Sixteenth ACM Conference on Knowledge Discovery and Data Mining (KDD 2010) (" <i>Bozorgi</i> ")       |
| <b>Exhibit 1011</b> | <i>Intentionally left blank</i>                                                                                                                                                                                                       |
| <b>Exhibit 1012</b> | U.S. Patent No. 9,294,498 to Yampolskiy et al. (" <i>Yampolskiy</i> ")                                                                                                                                                                |
| <b>Exhibit 1013</b> | Victor Benjamin et al., <u>Securing Cyberspace: Identifying Key Actors in Hacker Communities</u> , Proceedings of the 2012 IEEE International Conference on Intelligence and Security Informatics (2012) (" <i>Benjamin</i> ")        |
| <b>Exhibit 1015</b> | Mitch Macdonald et al., <u>Identifying Digital Threats in a Hacker Web Forum</u> , IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, (2015) (" <i>Macdonald</i> ")                                |

|                     |                                                                                                                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Exhibit 1016</b> | <u>CTC Sentinel</u> , Volume 6, Issue 7, July 2013. Combating Terrorism Center at West Point. (“ <i>CTC Sentinel</i> ”)                                                                                                                          |
| <b>Exhibit 1017</b> | Vivian Fiona Guetler, <u>Exploring Cyberterrorism, Topic Models and Social Networks off Jihadists Dark Web Forums: A Computational Social Science Approach</u> , Graduate Theses, Dissertations, and Problem Reports (2022) (“ <i>Guetler</i> ”) |
| <b>Exhibit 1018</b> | Screenshot of VirusTotal, found at <a href="https://www.virustotal.com/getstarted/vulnerability-management">https://www.virustotal.com/getstarted/vulnerability-management</a> , captured June 12, 2025 (“ <i>VirusTotal.com Screenshot</i> ”)   |
| <b>Exhibit 1019</b> | Elizabeth Phillips et al., <u>Applying Social Network Analysis to Security</u> , International Conference on Cyber Security for Sustainable Society (2015)                                                                                       |
| <b>Exhibit 1020</b> | <i>Intentionally left blank</i>                                                                                                                                                                                                                  |
| <b>Exhibit 1021</b> | Barreras et al., <u>Improving Transparency of the Colombian Peace Treaty with NLP</u> (“ <i>Barreras</i> ”)                                                                                                                                      |
| <b>Exhibit 1022</b> | Tianjun Fu et al., <u>A Focused Crawler for Dark Web Forums</u> , Journal of the American Society for Information Science and Technology (2010) (“ <i>Fu</i> ”)                                                                                  |
| <b>Exhibit 1023</b> | Victor Benjamin et al., <u>Exploring Threats and Vulnerabilities in Hacker Web: Forums, IRC and Carding Shops</u> , IEEE Intelligence and Security Informatics (2015)                                                                            |
| <b>Exhibit 1024</b> | U.S. Patent No. 11,031,135 to Gates et. al (“ <i>Gates</i> ”)                                                                                                                                                                                    |
| <b>Exhibit 1025</b> | Mohammed Almukaynizi et al., <u>Predicting Cyber Threats through Hacker Social Networks in Darkweb and Deepweb Forums</u> , CSSSA’s Annual Conference on Computational Social Science (2017), (“ <i>Almukaynizi</i> ”)                           |
| <b>Exhibit 1026</b> | Thomas J. Holt et al., <u>Examining the Social Networks of Malware Writers and Hackers</u> , International Journal of Cyber Criminology (2012) (“ <i>Holt</i> ”)                                                                                 |
| <b>Exhibit 1027</b> | Hoa Khanh Dam et al., <u>Automatic feature learning for predicting vulnerable software components</u> , IEEE Transactions of Software Engineering, (2018) (“ <i>Dam</i> ”)                                                                       |
| <b>Exhibit 1028</b> | Michel Edkrantz, <u>Predicting Exploit Likelihood for Cyber Vulnerabilities with Machine Learning</u> , Chalmers University of Technology, (2015) (“ <i>Edkrantz</i> ”)                                                                          |

20. I have considered these materials from the viewpoint of a POSITA as of the priority date of the '897 Patent. For the purposes of this declaration, I have been asked to assume that the priority date of the '897 Patent is November 3, 2017. I note that my opinions provided in this declaration are made from the perspective of a POSITA as of this priority date of the '897 Patent, unless expressly stated otherwise. To the extent that I use any verb tense in this declaration that is present tense (e.g., “a POSITA would understand” instead of “a POSITA would have understood”), such verb tense should be understood to be my opinion as of the '897 Patent’s priority date (again, unless expressly stated otherwise). I merely use the present verb tense for ease of reading.

### **III. LEGAL FRAMEWORK**

21. I am a technical expert and do not offer any legal opinions. However, counsel has informed me as to certain legal principles regarding patentability and related matters under United States patent law, which I have applied in performing my analysis and arriving at my technical opinions in this matter.

22. I have been informed that the Patent Trial and Appeal Board (“PTAB”) now applies the claim construction standard applied by Article III courts (i.e., the *Phillips* standard) regardless of whether a patent has expired. I have been informed that under the *Phillips* standard, claim terms are to be given the meaning they would have to a person having ordinary skill in the art at the time of the invention, taking

into consideration the patent, its file history, and, secondarily, any applicable extrinsic evidence (e.g., dictionary definitions). I understand that, for purposes of this proceeding, Petitioner applies the ordinary and customary meaning of the claim terms as understood by a POSITA. I have applied the plain meaning of the claim language as instructed by *Phillips*.

23. I have also been informed that a person cannot obtain a patent on an invention if the differences between the 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 person having ordinary skill in the art. I have been informed that a conclusion of obviousness may be founded upon more than a single item of prior art. I have been further informed that obviousness is determined by evaluating the following factors: (1) the scope and content of the prior art, (2) the differences between the prior art and the claim at issue, (3) the level of ordinary skill in the pertinent art, and (4) secondary considerations of non-obviousness. In addition, the obviousness inquiry should not be done in hindsight. Instead, the obviousness inquiry should be done through the eyes of a person of ordinary skill in the relevant art at the time the patent was filed.

24. In considering whether certain prior art renders a particular patent claim obvious, counsel has informed me that I can consider the scope and content of the prior art, including the fact that one of skill in the art would regularly look to the

disclosures in patents, trade publications, journal articles, industry standards, product literature and documentation, texts describing competitive technologies, requests for comment published by standard setting organizations, and materials from industry conferences, as examples. I have been informed that for a prior art reference to be proper for use in an obviousness analysis, the reference must be “analogous art” to the claimed invention. I have been informed that a reference is analogous art to the claimed invention if: (1) the reference is from the same field of endeavor as the claimed invention (even if it addresses a different problem); or (2) the reference is reasonably pertinent to the problem faced by the inventor (even if it is not in the same field of endeavor as the claimed invention). In order for a reference to be “reasonably pertinent” to the problem, it must logically have commended itself to an inventor’s attention in considering his problem. In determining whether a reference is reasonably pertinent, one should consider the problem faced by the inventor, as reflected either explicitly or implicitly, in the specification. I believe that all of the references that my opinions in this IPR are based upon are well within the range of references a person of ordinary skill in the art would consult to address the type of problems described in the Challenged Claims.

25. I have been informed that, in order to establish that a claimed invention was obvious based on a combination of prior art elements, a clear articulation of the reason(s) why a claimed invention would have been obvious must be provided.

Specifically, I am informed that a combination of multiple items of prior art renders a patent claim obvious when there was an apparent reason for one of ordinary skill in the art, at the time of the invention, to combine the prior art, which can include, but is not limited to, any of the following rationales: (A) combining prior art methods according to known methods to yield predictable results; (B) substituting one known element for another to obtain predictable results; (C) using a known technique to improve a similar device in the same way; (D) applying a known technique to a known device ready for improvement to yield predictable results; (E) trying a finite number of identified, predictable potential solutions, with a reasonable expectation of success; (F) identifying that known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other market forces if the variations are predictable to one of ordinary skill in the art; or (G) identifying an explicit teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine the prior art references to arrive at the claimed invention.

26. I am informed that the existence of an explicit teaching, suggestion, or motivation to combine known elements of the prior art is a sufficient, but not a necessary, condition to a finding of obviousness. This so-called “teaching suggestion-motivation” test is not the exclusive test and is not to be applied rigidly in an obviousness analysis. In determining whether the subject matter of a patent

claim is obvious, neither the particular motivation nor the avowed purpose of the patentee controls. Instead, the important consideration is the objective reach of the claim. In other words, if the claim extends to what is obvious, then the claim is invalid. I am further informed that the obviousness analysis often necessitates consideration of the interrelated teachings of multiple patents, the effects of demands known to the technological community or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art. All of these issues may be considered to determine whether there was an apparent reason to combine the known elements in the fashion claimed by the patent.

27. I also am informed that in conducting an obviousness analysis, a precise teaching directed to the specific subject matter of the Challenged Claim need not be sought out because it is appropriate to take account of the inferences and creative steps that a person of ordinary skill in the art would employ. The prior art considered can be directed to any need or problem known in the field of endeavor at the time of invention and can provide a reason for combining the elements of the prior art in the manner claimed. In other words, the prior art need not be directed towards solving the same specific problem as the problem addressed by the patent. Further, the individual prior art references themselves need not all be directed towards solving the same problem. I am informed that common sense is important and should be

considered. Common sense teaches that familiar items may have obvious uses beyond their primary purposes.

28. I also am informed that the fact that a particular combination of prior art elements was “obvious to try” may indicate that the combination was obvious even if no one attempted the combination. If the combination was obvious to try (regardless of whether it was actually tried) or leads to anticipated success, then it is likely the result of ordinary skill and common sense rather than innovation. I am further informed that in many fields it may be that there is little discussion of obvious techniques or combinations, and it often may be the case that market demand, rather than scientific literature or knowledge, will drive the design of an invention. I am informed that an invention that is a combination of prior art must do more than yield predictable results to be non-obvious.

29. I am informed that for a patent claim to be obvious, the claim must be obvious to a person of ordinary skill in the art at the time of the invention. I am informed that the factors to consider in determining the level of ordinary skill in the art include (1) the educational level and experience of people working in the field at the time the invention was made, (2) the types of problems faced in the art and the solutions found to those problems, and (3) the sophistication of the technology in the field.

30. I am informed that it is improper to combine references where the references teach away from their combination. I am informed that a reference may be said to teach away when a person of ordinary skill in the relevant art, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the patent applicant. In general, a reference will teach away if it suggests that the line of development flowing from the reference's disclosure is unlikely to be productive of the result sought by the patentee. I am informed that a reference teaches away, for example, if (1) the combination would produce a seemingly inoperative device, or (2) the references leave the impression that the product would not have the property sought by the patentee. I also am informed, however, that a reference does not teach away if it merely expresses a general preference for an alternative invention but does not criticize, discredit, or otherwise discourage investigation into the invention claimed.

31. I am informed that even if a prima facie case of obviousness is established, the final determination of obviousness must also consider "secondary considerations" if presented. In most instances, the patentee raises these secondary considerations of non-obviousness. In that context, the patentee argues an invention would not have been obvious in view of these considerations, which include: (a) commercial success of a product due to the merits of the claimed invention; (b) a

long-felt, but unsatisfied need for the invention; (c) failure of others to find the solution provided by the claimed invention; (d) deliberate copying of the invention by others; (e) unexpected results achieved by the invention; (f) praise of the invention by others skilled in the art; (g) lack of independent simultaneous invention within a comparatively short space of time; (h) teaching away from the invention in the prior art.

32. I am further informed that secondary considerations evidence is only relevant if the offering party establishes a connection, or nexus, between the evidence and the claimed invention. The nexus cannot be based on prior art features. The establishment of a nexus is a question of fact. While I understand that the Patent Owner here has not offered any secondary considerations at this time, I will supplement my opinions in the event that the Patent Owner raises secondary considerations during the course of this proceeding.


#### **IV. OPINION**

##### **A. The '897 Patent**

33. The '897 Patent generally describes systems and methods “for predicting which software vulnerabilities will be exploited by malicious hackers” to prioritize patching. *'897 Patent* (Ex. 1001), Abstract, 1:24-28. Methods for prioritizing patching vulnerabilities were admittedly known, including using “social media, darkweb markets and certain white-hat websites” to predict which

vulnerabilities will be exploited in the wild. *Id.*, 1:67-2:9. These methods had limitations, thus the '897 Patent aimed to develop machine learning models “that observe data taken from online discussion posts” “to predict a likelihood of exploitation of a particular vulnerability.” *'897 File History* (Ex. 1002), at 403. In other words, the invention determines “the likelihood that a given vulnerability mentioned in a data source [e.g., forums] will be exploited[.]” *Id.*, 545.

34. The invention “leverages machine learning techniques on features derived from the social network of users participating in darkweb/deepweb (DW) forums” to assess the likelihood that software vulnerabilities will be exploited. *'897 Patent*, 4:19-27, claim 1. The patent calls this an “exploit prediction model” reproduced below in Figure 14:



*Id.*, Fig. 14, 5:66-67. The exploit prediction model collects data from multiple data sources, for example, “crawl[ing] websites on DW, both marketplaces and forums, to collect data relating to malicious hacking.” *Id.*, 7:21-26, 8:25-40. Certain features are extracted from the collected data such as NVD features, social network features, DW post language, post content, or post date. *Id.*, 11:15-38, 21:6-46, Fig. 14. These stated features broadly include data related to malicious behavior. For example, National Vulnerability Database (NVD) provides “information on the vulnerability and what it allows hackers to do when they exploit [i]t,” while dark web (DW) description “provides rich context on what the discussion is about[.]” *Id.*, 10:15-37. Similarly, the NVD provides a Common Vulnerability Scoring System (CVSS) base score which “indicat[es] the severity of each of the disclosed vulnerability[ies].” *Id.*, 10:38-57. For the DW post, the system retains the language and date of the post as a feature for analysis. *Id.*, 10:58-67, 21:65-67, Fig. 14. Additional features, such as Common Weakness Enumeration (CWE) and social network features may also be used by the system. *Id.*, 11:15-23. CWE includes enumerated common software security weaknesses (categorical), that are broken down by “flaws that can be made during the software development and can exist in the software architecture, design, or code.” *Id.*, 11:24-28. The social network features include, for example, “in-degree” and “out-degree” which indicate how active a user is, or measuring “the centrality of users in the social graph.” *Id.*, 11:29-38, 21:6-46.

35. A vulnerability analysis is then performed on these features using machine learning. *Id.*, 14:14-22:41. Finally, the system computes a class label indicating whether a given vulnerability is likely to be exploited in the wild or not. *Id.*

**B. Claim Construction**

36. I have been informed by counsel that the first step in an unpatentability analysis involves construing the claims, as necessary, to determine their scope. Second, the construed claim language is then compared to the disclosures of the prior art. I am informed that claims are generally given their ordinary and customary meaning as understood by one of ordinary skill in the art at the time of the invention in light of the patent specification and prosecution history, i.e., the intrinsic record. For purposes of my analysis, I have applied the plain and ordinary meaning as defined in this paragraph.

**C. Level of a Person of Ordinary Skill in the Art**

37. I understand that the '897 Patent, the record of proceedings at the Patent Office (which I understand is called the "File History" or "Prosecution History"), and the teachings of the prior art are evaluated from the perspective of a person of ordinary skill in the art ("POSITA") at the time the claimed invention was made. I understand that the factors considered in determining the ordinary level of skill in the art include: (i) the levels of education and experience of persons working in the

field; (ii) the types of problems encountered in the field; and (iii) the sophistication of the technology. I may also consider, if available, the education level of the inventor, prior art solutions to the problems encountered in the art, and the rapidity with which innovations are made in the relevant art.

38. I understand that a person of ordinary skill in the art is not a specific real individual, but rather a hypothetical individual having the qualities reflected by the factors above. This hypothetical person has knowledge of all prior art in the relevant field as if it were arranged on a workshop wall and takes from each reference what it would teach to a person having the skills of a POSITA.

39. Based on my review and analysis of the '897 Patent, the prior art cited herein, and the ordinary skill factors described in this section, in my opinion a POSITA in the field of the '897 Patent at the time of the earliest possible priority date (Nov. 3, 2017) would have had 1) a bachelor's degree in electrical or computer engineering, computer science, or the equivalent, 2) and (a) either a master's degree in electrical or computer engineering, computer science, or the equivalent, (b) or two years of work/research experience in computer systems and networks. Additional work and/or research experience may substitute for formal education and vice versa. When I refer to the understanding of a POSITA, I am referring to the understanding of such a person as of Nov. 3, 2017.

40. As of Nov. 3, 2017, I had more than ordinary skill in the art and meet at least these minimum qualifications to be a person having ordinary skill in the art as of the time of the claimed invention of the '897 Patent. I am, however, familiar with the skills and knowledge possessed by those I would have considered to be of ordinary skill in the art as of that date. *See above*, Background and Qualifications.

41. My opinions provided in this declaration would not change in view of minor modifications to this level of ordinary skill.

**D. Background of the Technology**

42. Aspects of the '897 Patent are discussed in a 2017 thesis paper, “Predicting Cyber Threats through Hacker Social Networks in Darkweb and Deepweb Forums” by Almukaynizi et al. (“*Almukaynizi*”). *Almukaynizi* (Ex. 1025) states that the approach presented “combines social network analysis with machine learning techniques to predict future cyber threats through darkweb/deepweb discussions with hacking-related content.” *Almukaynizi*, 1. However, as I have summarized below, the key features of this approach and of the '897 Patent were well known in the art before November 3, 2017.

*1. Software Instructions for Machine Learning Models are Executed by Computer Processors*

43. It was well known in the art before the priority date of the '897 Patent that machine learning models contained software instructions that were executed by computer processors. *Kuperman* (Ex. 1008) teaches a model that “reduces the

likelihood of successful exploit 113 of web application 120 vulnerability 121 [.]” *Kuperman*, 13:47-52, Fig. 4 (depicting processor 402 in operable communication with storage unit 416 containing instructions 424). The model is executed by “stored instructions 424 (e.g., software) embodying any one or more of the methodologies or functions [.]” *Id.*, 24:28-32. Further, *Kuperman* teaches “instructions 424 (e.g., software) may also reside, completely or at least partially, within the main memory 404 or within the processor 402... during execution thereof by the computer system 400[.]” *Id.*, 24:32-36. Therefore, *Kuperman*’s model is executed by a processor.

44. U.S. Patent No. 9,294,498 to Yampolskiy et al. (“*Yampolskiy*”) was filed on May 1, 2015. *Yampolskiy* (Ex. 1012). *Yampolskiy* teaches a system that comprises a “security signal collection module 210, a contextualization and attribution module attribution module 220, and a benchmarking module 230.” *Id.*, 7:4-6. Specifically, the contextualization and attribution module comprises a “machine learning module 228 to identify and update which factors most significantly affect an entity’s cybersecurity.” *Id.*, 7:64-67. The contextualized data coming from this module is used to calculate a “cybersecurity risk score.” *Id.*, 7:67 – 8:3. *Yampolskiy*’s system includes a “non-transitory computer-readable medium” comprising “instructions which cause the processor to perform the step of calculating a security score for at least one of the one or more of the types of data [.]” *Id.*, 2:8-16. Thus, the method of *Yampolskiy*’s system, including the machine

classifier module, is performed by software instructions that are executed by a processor.

## 2. *Social Measures Analysis Between Users on Hacker Forums*

45. Extracting and analyzing information about certain users on hacker forums is a common method of gaining further insight into a hacker community. For instance, *Benjamin* (Ex. 1013) discusses identifying the “key actors” of hacking communities by analyzing hacker communications. *Benjamin*, 4. Particularly, certain features related to a user are analyzed when attempting to identify these key actors. *See e.g., id.*, 7 (disclosing “Discussion Quality Features” as average message length and number of replies per thread, and “Community Involvement Features” as number of threads involved, tenure, sum of attachments, and total messages). These features give insight into the “hacker reputation” which can ultimately help identify “key actors” in hacker communities. *See e.g., id.*, 6 (“Two of our selected features aim to capture hacker discussion quality while the other four features are intended to model the extent of a hacker’s involvement in their community”). The identification of key actors through the analysis of user social measures and features is known to increase the ability to analyze hacker communities. *See e.g., id.*, 9 (“Further, this research contributes to advancing our understanding of hacker communities”).

46. *Macdonald* utilizes “sentiment analysis” to analyze textual communication data to better understand social interactions of users and hacker communication content. Sentiment analysis is a tool that can be used to understand the general context of discussions. *See e.g., Macdonald*, 5 (“Due to its ability to assign positive or negative scores to comments or reviews, companies can understand the general sentiment of online discussions...”). Specifically, it was known that “sentiment analysis is also a useful tool for examining social networks[.]” *Id.*, 5. Because of this, sentiment analysis is used to examine and analyze hacker communications and social interactions between users. *See e.g., id.*, (“Within a security context, sentiment analysis has been used to examine the online activities of criminal subculture, including extremist and hacker groups”).

47. User information has been analyzed and studied to understand the roles individuals play in hacker communities. For example, *Holt* was published in the International Journal of Cyber Criminology in 2012. *Holt* (Ex. 1026). *Holt* discusses the exploration of “the social networks of a group of Russian hackers using publicly accessible data to understand the nature of social relationships and the ways that they affect information sharing and action.” *Holt*, 1. *Holt* discusses how user information was extracted from forums, and a “risk index” was developed for each user. *See e.g., Holt*, 5 (describing user contact information being used to determine “their involvement in the hacker community” and assigning a risk index based on that

information). The “risk index” was used to rank users in hacker communities to understand the role each of the users played in the hacker communities. *See e.g., id.*, (“The results were then developed to create a risk index for each user based on their involvement in the creation, distribution, and use of malicious software and hacking techniques”). Thus, utilizing social and user information is known to aid in the understanding and analysis of hacker community discussions.

### 3. *Dark Web Forums Include Hacker Communications*

48. Dark web forums were known to include hacker forums before the priority date of the '897 Patent. It is well known that hackers commonly communicate on the dark web in hacking-dedicated discussion forums. *See, e.g., Benjamin*, 4 (“Hackers often congregate within underground communities, most commonly in the form of Internet Relay Chat (IRC) networks or online forums”). The dark web is able to provide security and anonymity that is often sought by forum users posting hacking-related discussions. *See e.g., Guetler* (Ex. 1017), 30 (“The Dark Web forums also provide secure communication and discussion platforms...”); *see also, Macdonald* (Ex.1015), 927 (“Hacker communities are dark networks, so members explicitly protect their anonymity, which serves to counter detection from law enforcement [29]...For this reason, discussion web forums are a popular method of communication amongst hacker communities [6]”).

49. *Macdonald* acknowledges, “Cyber-attacks targeting the computer and control systems of critical infrastructures are increasing in both in scope and prevalence.” *Macdonald* (Ex. 1015), 1. Hacker forums on the dark web are known to contain information regarding potential threats and attacks on computer systems. *See e.g., id.*, 2 (describing the discussions of threats to critical infrastructures occurring on hacker forums). The information provided by hacker communications on dark web forums is often monitored, extracted, and analyzed for the purpose of cybersecurity defense. *See e.g., id.*, 1 (“One strategy to combat this problem is to monitor and analyze the communal networks of the cyber criminals who commit these attacks. More specifically, their communication activity can be analyzed to determine potential threats”).

4. *Predicting the Likelihood Software Vulnerabilities Will be Exploited Before Disclosure*

50. Before the priority date of the '897 Patent, it was well known in the art for machine classifiers to be trained to predict if software vulnerabilities would be exploited before disclosure, e.g., before disclosure to a public database. For instance, *Bozorgi* was included in Proceedings of the 16<sup>th</sup> ACM SIGKDD International Conference on Knowledge Discovery and Data Mining in 2010. *Bozorgi*, (Ex. 1010). *Bozorgi* acknowledges that it is important to be able to determine which software vulnerabilities should be prioritized for patching. *See e.g., id.*, 1 (“a key question for systems administrators is which vulnerabilities to prioritize”). Thus, machine

classifiers are trained with previous exploit and vulnerability information to determine “how soon a vulnerability is likely to be exploited.” *Id.* Specifically, *Bozorgi* discusses training a support vector machine (SVM) classifier “to predict whether or not a vulnerability will be exploited within some time  $t$ , where  $t$  is the difference between the exploit and disclosure dates of the vulnerability.” *Id.*, 7. Therefore, it was known that machine classifiers could be used to predict the likelihood of a vulnerability being exploited.


## V. OPINION – GROUND 1

### A. Obvious Over *Sabottke*

#### 1. Overview of *Sabottke*

51. *Sabottke* teaches a “Twitter-based exploit detector” used to “assess[] which vulnerabilities are likely to be exploited” to prioritize a response. *Sabottke* (Ex. 1004), 2. “[H]ackers...discuss vulnerabilities on social media sites like Twitter” and thus these sites “constitute rich sources of information” for exploit prediction. *Id.* For “early exploit detection using information available on Twitter[,]” *Sabottke* utilizes “supervised machine learning and ground truth about exploits[.]” *Id.* *Sabottke* collects a corpus of tweets and queries the NVD and OSVDB databases for context about the vulnerabilities discussed on Twitter, “extract[s] features for training and testing a support vector machine (SVM) classifier[,]” and uses the SVM for exploit prediction. *Id.*, 3, 5, 9.

52. Specifically, *Sabottke* provides “early detection of such exploits by using information that is available publicly [e.g., from hacker Twitter streams], but is not included in existing vulnerability databases[.]” *Id.*, 3. In this way, *Sabottke* can predict vulnerabilities likely to be exploited prior to their disclosure to public vulnerability databases. *Id.*, 9 (“Information posted before disclosure”); *see also, id.*, 2 (“We characterize the exploit-related discourse on Twitter, the information posted before vulnerability disclosures, and the users who post this information”). Figure 1 below provides an overview of *Sabottke*’s system architecture:



*Id.*, Fig. 1.

53. Because *Sabottke* and the ’897 Patent are both directed to assessing cybersecurity risks for vulnerabilities, *Sabottke* is in the same field of endeavor as the ’897 Patent. *Compare Sabottke*, Abstract (teaching a system for “assessing which


vulnerabilities are likely to be exploited”), 2-3, 5 with *’897 Patent*, 1:24-28 (describing “assessing the likelihood of exploitation of software vulnerabilities”), 4:1-27, 4:41-44.

54. The *’897 Patent* describes common problems in the art related to using online discussions to predict whether certain vulnerabilities will be exploited. *’897 Patent*, 4:19-27; *see also id.*, 8:41-45 (describing DW marketplaces and forums to gather information about vulnerabilities and exploits), 13:47-50 (discussing trade of real-world exploits on DW hacking forums). *Sabottke* similarly addresses the problem of using online discussion to predict whether certain vulnerabilities will be exploited. *Sabottke*, 2-3 (discussing how on Twitter there is “a community of hackers, security vendors and system administrators discuss security vulnerabilities”), *see also*, 5 (disclosing the Twitter-based exploit detector extracting “vulnerability-related information from the Twitter stream”), 6-8, 16-17. As such, *Sabottke* is reasonably pertinent to the problems addressed by the *’897 Patent* and is analogous art.

2. *Claim 1(Pre): A method for assessing a likelihood of exploitation of software vulnerabilities, comprising:*

55. To the extent the preamble is limiting, *Sabottke* teaches a “Twitter-based exploit detector” that assesses the likelihood software vulnerabilities will be exploited. *Sabottke*, 2. *Sabottke* notes the growing number of “software vulnerabilities” has created a “need for prioritizing the response to new disclosures

by assessing which vulnerabilities are likely to be exploited” and proposes a “Twitter-based exploit detector” to fill that need. *Id.*; see also *id.*, 13 (“Twitter data provides valuable insights into the likelihood of exploitation”), 15 (“likelihood of real-world exploitation, produced by our system”). *Sabottke’s* exploit prediction system architecture is substantially similar to the ’897 Patent’s exploit prediction model:



*Sabottke*, Fig. 1; ’897 Patent, Fig. 1B.

56. *Sabottke’s* exploit prediction system architecture in Figure 1 above is substantially similar to the ’897 Patent’s exploit prediction model illustrated in Figure 1B. Both use several datasets for input data; extract features from this data; and utilize machine learning techniques along with a ground truth dataset (e.g., Symantec attack signatures) to classify that data and determine a likelihood a vulnerability mentioned in that data will be exploited in the wild.

57. In more detail, *Sabottke's* system architecture teaches the use of identified vulnerabilities that have been “exploited in the real world” as a ground truth data set. *Sabottke*, 5, Fig.1. Similarly, the '897 Patent describes that the ground truth data “comprise[s] a set of attack signatures of exploits detected in the wild[.]” '897 Patent, 7:2-4, Fig. 1B. *Sabottke* teaches assessing vulnerability features including, “CVSS Information and Database Information.” *Sabottke*, 6, Fig. 1. The '897 Patent assesses similar vulnerability features including CVSS and database information. *See e.g.*, '897 Patent, 6:60 – 7:6 (describing features that are included such as “features [] from NVD namely, CVSS Score, CVSS vector and CWE”); *see also id.*, Fig. 1B. In both *Sabottke* and the '897 Patent, the ground truth and features are used in classification techniques to predict the likelihood a vulnerability will be exploited. *See e.g.*, *Sabottke*, Fig. 1 (teaching the output of the classification as the “classification results”); *see also, e.g.*, '897 Patent, Fig. 1B (showing the output of the machine learning classifier as the “exploit prediction”).

3. *Claim 1(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations, including:*

58. *Sabottke's* “Twitter-based exploit detector” is a process implemented in software and executed by hardware that utilizes “Twitter’s Streaming API[.]” Python scripts, machine learning, and other software features for its exploit detection methods. *Sabottke*, 5 (“We present the design of a Twitter-based exploit detector,

using supervised machine learning techniques” and describing Twitter API and “Python scripts” for data collection). In my opinion, a POSITA would have understood that *Sabottke* suggests the hardware to execute the software instructions associated with its system. The software that is used by *Sabottke* would require a computer processor to execute that software and therefore perform the functions of the machine learning techniques, Twitter API, and Python scripts. *Sabottke*, 5. Namely, in my opinion a POSITA would have understood *Sabottke*’s “Twitter-based exploit detector” would have been executable on a general computing system available in 2017, including distributed computing systems, comprising a processor executing software instructions stored in memory for implementing *Sabottke*’s disclosures.


59. As I explained above in my Background of the Technology, it was well known for a processor to execute instructions to implement software machine learning techniques. For example, *Kuperman* provides an exemplary classification system executed by a computer processor. *See e.g., Kuperman* (Ex. 1008), 13:47-52, 24:28-39 (describing that the machine-readable medium contains stored instructions for implementing the machine learning model “embodying any one or more of the methodologies or functions” and the “instructions 424 (e.g., software) may also reside....within the main memory 404 or within the processor 402”), Fig. 4.

Similarly, *Yampolskiy* describes executing machine learning techniques on a processor. *See above*, ¶¶43-44.

60. At a minimum, in my opinion a POSITA would have found it obvious to implement *Sabottke*'s "Twitter-based exploit detector" on a known computing system by executing software instructions in memory via a CPU. In general, it is known in the art that software instructions can be stored in the memory of a computer and executed by the processor of that computer. *See e.g., Yampolskiy*, 2:7-29 (disclosing the system comprising a memory and processor where the processor is configured execute the steps of calculating the overall cybersecurity risk score); *see also, Kuperman*, 24:28-39 (describing the software instructions to be executed within the processor); *See above*, Section IV.D.1. Therefore, to implement *Sabottke*'s software instructions, those instructions must be executed by a CPU.

4. *Claim 1(b): accessing a plurality of datasets associated with a predetermined set of data sources, the plurality of datasets including training data comprising hacker communications;*

61. *Sabottke* teaches accessing a plurality of datasets from predetermined data sources, such as Twitter feeds and the NVD and OSVDB databases. These plurality of datasets are shown below in Figure 1:




*Sabottke*, Fig. 1 (annotated). *Sabottke*’s “detector extracts vulnerability-related information from the Twitter stream, and augments it with additional sources of data about vulnerabilities and exploits.” *Id.*, 5. *Sabottke* “collect[s] tweets mentioning vulnerabilities” using “Twitter’s Streaming API” and Python scripts. *Id.* *Sabottke* augments with additional sources to “obtain context about the vulnerabilities discussed on Twitter,” including querying “the National Vulnerability Database (NVD)” for CVSS scores, products affected, and additional references about these vulnerabilities, and crawling the “Open Sourced Vulnerability Database (OSVDB)” for disclosure dates and categories of vulnerabilities. *Id.*

62. From the corpus of collected tweets, *Sabottke* “extract[s] features for training and testing a support vector machine (SVM) classifier.” *Id.*, 2-3. *Sabottke* expressly teaches collecting tweets on Twitter that include hacker communications

relevant to security vulnerabilities. *Id.*, 2 (“The security vendors, system administrators, and hackers, who discuss vulnerabilities on social media sites like Twitter, constitute rich sources of information” for *Sabottke’s* system), 3 (“Specifically, we analyze the Twitter stream... On Twitter, a community of hackers, security vendors and system administrators discuss security vulnerabilities”).

5. *Claim 1(c): accessing features from the plurality of datasets that include measures computed from social connections of users posting hacking-related content;*

63. *Sabottke* teaches accessing features from the plurality of datasets during feature selection. *Sabottke*, 2-3 (“we extract features for training and testing a support vector machine (SVM) classifier). *Sabottke’s* feature selection process is depicted below:



*Id.*, Fig. 1. *Sabottke* accesses four categories of features: “Twitter Text, Twitter Statistics, CVSS Information and Database Information.” *Id.*, 6. The specific features are listed below Figure 2:

Figure 2: Mutual information between real world and public proof of concept exploits for CVSS, Twitter user statistics, NVD, and OSVDB features. **CVSS Information** - 0: CVSS Score, 1: Access Complexity, 2: Access Vector, 3: Authentication, 4: Availability Impact, 5: Confidentiality Impact, 6: Integrity Impact. **Twitter Traffic** - 7: Number of tweets, 8/9: # users with minimum T followers/friends, 10/11: # retweets/replies, 12: # tweets favorited, 13/14/15: Avg # hashtags/URLs/user mentions per tweet, 16: # verified accounts, 17: Avg age of accounts, 18: Avg # of tweets per account. **Database Information** - 19: # references in NVD, 20: # sources in NVD, 21/22: BUGTRAQ/SECUNIA in NVD sources, 23: allow in NVD summary, 24: NVD last modified date - NVD published date, 25: NVD last modified date - OSVDB disclosed date, 26: Number of tokens in OSVDB title, 27: Current date - NVD last modified date, 28: OSVDB in NVD sources, 29: code in NVD summary, 30: # OSVDB entries, 31: OSVDB Category, 32: Regex Category, 33: First vendor in NVD, 34: # vendors in NVD, 35: # affected products in NVD.

*Id.*, 8. *Sabottke*’s “Twitter Traffic” features include measures computed from social connections of users posting hacking-related content. Namely, at least the “# users with minimum T followers/friends,” “# retweets/replies,” and “# tweets favorited” features represent measures computed from the poster’s social connections. In my opinion, a POSITA would have understood that *Sabottke*’s “Twitter Traffic” features represent measures computed from a user’s / poster’s social connections at least because these features are representative of the social interactions between the poster and other users. *Benjamin* (Ex.1013) confirms this. *Benjamin* discloses studying “[d]iscussion quality features” which included “average message length” and “number of replies per thread[.]” *Benjamin*, 7-8. These discussion quality features have previously been studied for the effect they have on “social behavior and interactions.” *Id.*, 7. Thus, in my opinion, a POSITA would understand that features such as the “# retweets/replies” of *Sabottke* would include measures computed from

social connections of users posting hacking-related content. *See above*, Section IV.D.2.

6. *Claim 1(d): applying learning algorithms to the training data to generate classification models that are configured to predict class labels defining a likelihood of exploitation of respective software vulnerabilities;*

64. *Sabottke* teaches applying supervised learning algorithms to train its exploit detector such that the exploit detector predicts class labels defining a likelihood of exploitation of respective software vulnerabilities. *Sabottke*, 2, 4, 7. Specifically, *Sabottke*'s "techniques utilize supervised machine learning" to train linear support vector machine (SVM) classifiers. *Id.*; *see also id.*, 5 ("We present the design of a Twitter-based exploit detector, using supervised machine learning techniques"), 7 ("We train linear support vector machine (SVM) classifiers [] in a feature space" "using multiple rounds of stratified random sampling"). *Sabottke*'s supervised training techniques for its SVM classifiers are "learning algorithms" that generate a trained SVM (i.e., *classification model*). The SVM classifier is trained, "in a feature space with 67 dimensions that results from our feature selection step[.]" *Id.*, 7. So, the SVM is being trained with features that correspond to the vulnerabilities identified. Thus, the SVM will be trained based on known features of known vulnerabilities. Indeed, "supervised" training techniques are an example of well-known learning algorithms to generate a classification model. *See, e.g., Bozorgi*

at 4 (Section 4.2 describing supervised training an SVM using labeled datasets to generate a trained classification model).

65. *Sabottke*’s trained SVM classifiers are configured to predict class labels defining a likelihood of exploitation of respective software vulnerabilities. *Sabottke* “aim[s] to train a classifier that produces binary predictions: each vulnerability is classified as either exploited or not exploited.” *Id.*, 4. This teaching is consistent with the disclosures of the ’897 Patent. The ’897 Patent provides exemplary classification models, including “Support Vector Machines,” that predict a binary “class label,” which is a label that conveys a value of “exploited or not exploited.” ’897 Patent, 14:30-33, 17:64-65 (“[...] tells about another variable (here the class label  $y \in \{\text{exploited}, \text{not exploited}\}$ ).”), 21:57-59, 22:66-23:32. Analogously, *Sabottke*’s trained SVMs are a classification model configured to predict class labels as claimed. *Sabottke*, 6 (“[w]e compute...the class labels  $Y \in \{\text{exploited}, \text{not exploited}\}$ ” based on input features). *Sabottke*’s class label defines a likelihood of exploitation of respective software vulnerabilities. For example, *Sabottke* computes class labels defining one or more values defining “which vulnerabilities are likely to be exploited.” *Id.*, 2 (“assessing the risk that the vulnerabilities will be exploited”). *Sabottke*’s classifier is a statistical engine and cannot predict with *certainty*, but only with some *uncertainty*, hence, can provide only a *likelihood*. *Sabottke*, Abstract (“assessing which vulnerabilities are likely to be exploited”).

7. *Claim 1(e): accessing one or more features associated with a software vulnerability; and;*

66. *Sabottke* teaches accessing features associated with software vulnerabilities, including CVSS information and vulnerability database information from NVD and OSVDB:

To obtain context about the vulnerabilities discussed on Twitter, we **query the National Vulnerability Database (NVD) [7] for the CVSS scores, the products affected and additional references about these vulnerabilities.** Additionally, we **crawl the Open Sourced Vulnerability Database (OSVDB) [9] for a few additional attributes, including the disclosure dates and categories of the vulnerabilities in our study.**

*Sabottke*, 5. Regarding CVSS features, *Sabottke* accesses “CVSS Information - 0: CVSS Score, 1: Access Complexity, 2: Access Vector, 3: Authentication, 4: Availability Impact, 5: Confidentiality Impact, 6: Integrity Impact.” *Id.*, 8. For vulnerability database features, *Sabottke* accesses “Database Information - 19: # references in NVD, 20: # sources in NVD, 21/22: BUGTRAQ/SECUNIA in NVD sources, 23: allow in NVD summary, 24: NVD last modified date - NVD published date, 25: NVD last modified date - OSVDB disclosed date, 26: Number of tokens in OSVDB title, 27: Current date - NVD last modified date, 28: OSVDB in NVD sources, 29: code in NVD summary, 30: # OSVDB entries, 31: OSVDB Category, 32: Regex Category, 33: First vendor in NVD, 34: # vendors in NVD, 35: # affected products in NVD.” *Id.*; *see also id.*, Figure 1 (illustrating that the vulnerability features are also considered during the classification phase).

8. *Claim 1(f): computing, by applying the one or more features to the classification model, a class label defining one or more values defining a likelihood of exploitation associated with the software vulnerability*

67. The '897 Patent, referring to vulnerability features and the class label, states “[...] how much a variable (here a feature  $x_i$ ) tells about another variable (here the class label  $y \in \{\text{exploited, not exploited}\}$ ).” ’897 Patent, 17:63-65. *Sabottke* teaches or suggests computing, by applying the one or more features to the SVM classification model, a class label (e.g., exploitability) defining one or more values (e.g., “exploited” or “not exploited”) defining a likelihood of exploitation associated with the software vulnerability. Specifically, *Sabottke*’s trained SVM classifiers are configured to compute class labels for exploitation, defining values for vulnerabilities as “either exploited or not exploited.” *Supra*, Claim 1(d) (nothing *Sabottke* at 4 “aim[s] to train a classifier that produces binary predictions: each vulnerability is classified as either exploited or not exploited”); *see also Sabottke*, 6 (“[w]e compute...the class labels  $Y \dots \in \{\text{exploited, not exploited}\}$ ” based on input features). *Sabottke*’s class labels are computed by applying the extracted features described above in Claims 1(c) and 1(e) and listed in figure 2 of *Sabottke* to *Sabottke*’s trained SVM classification model. *Sabottke*, Fig. 1; *supra*, Claims 1(c) and 1(e) (describing CVSS Information, Twitter Traffic, and Database Information features).

68. *Sabottke*'s class labels define one or more values (e.g., exploited or not exploited) that define a likelihood of exploitation associated with the software vulnerability. *See e.g., Sabottke*, 10. Indeed, *Sabottke*'s classifier is a statistical engine and cannot predict with *certainty*, but only with some *uncertainty*, hence, can provide only a *likelihood*. *Id.*, Abstract (“...assessing which vulnerabilities are likely to be exploited”). For example, it was well-known that binary class labels indicating exploitation define two values, “exploited or not exploited.” *See, e.g., Bozorgi*, 3 (describing labeling vulnerabilities as positive “indicating that it has been exploited” and negative “indicating that it is not exploited”). Indeed, this is represented as a “decision rule”:

The decision rule mapping feature vectors ... to labels  $y \in \{-1, +1\}$

is given by:  $y = \text{sign}(w \cdot x + b)$

*Id.*, 4.

69. It was also well known that SVMs output a prediction indicating which of two classes a data point belongs to. *See, e.g., Bozorgi*, 8. *Bozorgi* explains that an SVM computes “the signed distance to the maximum margin hyperplane separating positive and negative examples” such that the “sign predicts whether it will be exploited, and for positively labeled examples, the magnitude indicates its severity.” *Id.*; *see also id.*, Fig. 4(b) (showing histogram of outputs compared a threshold 0, where above 0 will result in “exploited” values). In my opinion, a POSITA would

have understood that a greater magnitude of the positively labeled prediction (i.e., the “distance ( $\mathbf{w} \cdot \mathbf{x} + b$ )”) would indicate a greater likelihood of exploitation as it is further from the maximum margin of the hyperplane and therefore predicts with lower uncertainty. *Id.*, 8.

70. For *Sabottke*’s exploit prediction, a POSITA would have understood *Sabottke*’s SVM likewise indicates whether a particular vulnerability is likely to be exploited, or not likely to be exploited by, for example, computing the sign of the datapoint and its magnitude. *Sabottke*, 13 (“Twitter data provides valuable insights into **the likelihood of exploitation**”), 15 (“**likelihood of real-world exploitation**, produced by our system”). Based on this, a vulnerability receives an exploitability class label that defines a binary value – “exploited, not exploited.” *See e.g., Sabottke*, 4 (“We aim to train a classifier that produces binary predictions: each vulnerability is classified as either exploited or not exploited”). In this way, *Sabottke* computes class labels defining one or more values defining “which vulnerabilities are likely to be exploited.” *Id.*, 2 (“assessing the risk that the vulnerabilities will be exploited”), 15 (“likelihood of real-world exploitation, produced by our system”).

9. *Claim 1(f)(1): wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data.*

71. Like the '897 Patent, *Sabottke* teaches predicting actual exploitation of software vulnerabilities before those vulnerabilities are disclosed to a known vulnerability database.


72. During prosecution, Applicant provided context for analyzing vulnerabilities “before disclosure.” In a July 13, 2023, request for continued examination after a final rejection, Applicant differentiated U.S. Patent Pub. No. 2017/0214710 to Seiver et al. (“Seiver”), stating:

*Seiver does not discuss predicting a likelihood of exploitation of a vulnerability discussed online prior to disclosure to a public database or prior to actual exploitation of the vulnerability; i.e., Seiver does not disclose the features of “wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data”*

'897 File History, 480 (emphasis in original). Per Applicant, “Seiver is not directed to actively ‘scouting’ the web for vulnerabilities that have a high likelihood of exploitation and instead appears to be more directed towards defensive strategies for protecting a network.” *Id.*; *see also id.*, 480-81 (differentiating another reference based on this same limitation for not “consider[ing] text information associated with discussion of vulnerabilities by users posting hacking-related content at a web forum prior to disclosure to a public vulnerability database”). Like the claimed invention, *Sabottke* teaches actively scouting hacker communications in Twitter feeds to predict the likelihood a vulnerability will be exploited before disclosure to NVD or

another vulnerability database. For example, *Sabottke* teaches that “[f]or collecting tweets mentioning vulnerabilities, the system monitors occurrences of the ‘CVE’ keyword using Twitter’s Streaming API[.]” *Sabottke*, 5. By monitoring in this way, *Sabottke*’s system can “extract[] vulnerability-related information from the Twitter stream[.]” *Id.*

73. *Sabottke* expressly teaches methods to “characterize the landscape of threats related to **information leaks about vulnerabilities before their public disclosure**” to a public vulnerability database. *Sabottke*, 3. “For the vulnerabilities in our data set, Figure 3 compares the dates of the earliest tweets mentioning the corresponding CVE numbers with the public disclosure dates for these vulnerabilities. Using the disclosure date recorded in OSVDB, we identify 47 **vulnerabilities that were mentioned in the twitter stream before their public disclosure.**” *Id.*, 9 (“4.2 Information posted before disclosure”). Indeed, *Sabottke* teaches detecting the popular “Heartbleed” vulnerability before disclosure to the ExploitDB database. *Sabottke*, 13 (“Our Twitter-based classifier provides an ‘exploited’ output 3 hours after the first tweet appears related to this CVE on April 7, 2014” and before “Heartbleed appeared in ExploitDB on the day after disclosure (April 8, 2014)”). Figure 8c illustrates this process:



*Id.*, Fig. 8c. The dotted black line represents Twitter exploit detection and the dotted green line represents disclosure to the ExploitDB database. *Id.*, 13-14. The detection of the vulnerability in the Twitter stream occurs before the disclosure of the vulnerability to ExploitDB and well before the creation of a Symantec Signature. *Id.*, Fig. 8c. Thus, *Sabottke* teaches that the software vulnerability is predicted before disclosure.

74. In the “non-goals” section of *Sabottke*, *Sabottke* explains its invention does not “consider the detection of zero-day attacks...which exploit vulnerabilities before their public disclosure.” *Sabottke*, 3. In other words, *Sabottke* does not assess vulnerabilities that have not been made public, like vulnerabilities mentioned in “commercial or private feeds.” *Id.* *Sabottke* does, however, look at vulnerabilities that are mentioned publicly (e.g., on twitter), but that have not been reported to

public vulnerability databases like NVD or ExploitDB. *Sabottke*, Fig. 8c, 13-14. *Sabottke*'s teachings are consistent with the intrinsic record. *Sabottke* monitors the Twitter stream for "tweets mentioning vulnerabilities" to "characterize the exploit-related discourse on Twitter, **the information posted before vulnerability disclosures**, and the users who post this information." *Sabottke*, 2.

75. Similarly, the '897 Patent crawls hacker discussions on online forums and is interested in vulnerabilities mentioned on those forums that have not been disclosed to public vulnerability databases to predict whether that vulnerability will actually be exploited in the real world. Specifically, the '897 Patent describes collecting vulnerability information before it has been publicly disclosed through implementing the exploit prediction model "to crawl websites on DW, both marketplaces and forums, to collect data relating to malicious hacking." '897 Patent, 8:26-29. From these data sources, the '897 Patent can compute the likelihood of exploitation of a vulnerability. '897 Patent, 7:8-12 (describing "computing likelihood of exploitation" step). By analyzing Twitter discourse surrounding vulnerabilities, *Sabottke* is able to make predictions regarding the likelihood of a vulnerability being exploited before those vulnerabilities are reported to a public database. See e.g., *Sabottke*, 3 ("We investigate the opportunities for early detection of such exploits by using information that is available publicly, but is not included in existing vulnerability databases such as the National Vulnerability Database


(NVD) [] or the Open Sourced Vulnerability Database (OSVDB)”). *Sabottke* likewise predicts a likelihood of exploitation of vulnerabilities before disclosure to public vulnerability databases and therefore teaches this limitation.

10. *Claim 5 - The method of claim 1, wherein certain features correspond to a known vulnerability obtained from the plurality of datasets.*

76. *Sabottke* teaches accessing certain features such as CVSS information and vulnerability database information from NVD and OSVDB corresponding to a known software vulnerability. *Supra*, Claim 1(e). These features correspond to a “known vulnerability” such that they provide “context about the vulnerabilities discussed on Twitter” from vulnerability data previously disclosed in the public databases. *Sabottke*, 5. Databases containing reported vulnerabilities are crawled to obtain attributes such as “disclosure dates” and “categories of the vulnerabilities.” *Id.* After these attributes are obtained, multiple features are extracted from the data. *See e.g., Id.*, (“From the raw data collected, we extract multiple features...”). Therefore, certain features of known vulnerabilities, such as “disclosure dates” and “categories of the vulnerabilities[,]” are obtained from *Sabottke*’s dataset.

11. *Claim 9 - The method of claim 1, further comprising computing mutual information from the plurality of datasets informative as to what information a given feature provides about another feature.*

77. *Sabottke* teaches computing “[m]utual information” from the plurality of datasets:




*Sabottke*, Fig. 2; *see also id.*, 8 (“Figure 2: Mutual information between real world and public proof of concept exploits for CVSS, Twitter user statistics, NVD, and OSVDB features”). Mutual information (MI) is informative as to what information a given feature provides about another feature because it “can single out useful features suggesting that the vulnerability is exploited as well as features suggesting it is not.” *Sabottke*, 6; *see also id.*, 6 (Table I). Based on *Sabottke*’s express teachings, in my opinion a POSITA would have understood that the computation of mutual information is used to evaluate how useful a feature is with respect to another feature in predicting the exploitation label Y, such that useful features can be “single[d] out” that assist determine whether a vulnerability is exploited or not. *Sabottke*, 6.

12. *Claim 10 - The method of claim 1, further comprising: detecting, from the plurality of datasets, vulnerabilities that appear before an associated exploitation date.*

78. *Sabottke* teaches detecting vulnerabilities before an associated exploitation date by analyzing Twitter feeds for vulnerabilities mentioned that have yet to be exploited. *Sabottke*, 12-14. *Sabottke*’s classifier provides “early detection”

for real-world exploits. *Id.* Figure 8c below “illustrates early detection for the case of Heartbleed (CVE-2014-1060)”:



*Id.*, Fig. 8c. “The dashed black line represents the earliest time when our classifier is able to detect this vulnerability as a real-world-exploit...3 hours after the first tweet appears related to this CVE” and well before a Symantec attack signature was generated (red line). *Id.*, 13 (describing detecting an exploit “within 10 minutes of its first appearance on Twitter”). Further, “[t]he green line indicates the first appearance of the vulnerability in ExploitDB[.]” *Id.* ExploitDB collects exploits and is “the largest archive of public exploits.” *Id.*, 11. So, *Sabottke* is able to detect vulnerabilities on Twitter before the “first appearance of the vulnerability in ExploitDB” and before it is exploited in the real-world. *Id.*, 13. *Sabottke* thus detects vulnerabilities mentioned on Twitter before an associated exploitation date.

13. *Claim 14 - A computing device, configured via machine learning to apply a learned function to data associated with a software vulnerability to estimate a likelihood of exploitation of the software vulnerability, the learned function associated with a prediction model derived from at least one machine learning algorithm and a plurality of datasets associated with software vulnerabilities, the plurality of datasets including information associated with discussion of vulnerabilities by users posting hacking-related content, wherein the likelihood of exploitation predicts an actual exploitation of the software vulnerability before disclosure based on hacker communications from the plurality of datasets.*

79. *Sabottke* teaches or suggests a computing device for implementing the software functions described in claims 1 and 14. *Supra*, Claim 1 (describing software functions implemented by *Sabottke* and application by known hardware). The '897 Patent explains its claimed invention may be performed by an application hosted on or otherwise implemented by a “computing device 104” that “may include a server, controller, a personal computer, a terminal, a workstation...or other such computing device.” '897 Patent, 6:11-19. Likewise, *Sabottke*'s described invention would have been implemented by a processor on a similar computing device. It is known in the art that software instructions are executed by processors on computing devices. *See e.g., Kuperman*, 24:28-39 (disclosing that “instructions 424 (e.g., software) may also reside....within the main memory 404 or within the processor 402”); *see above*, Section IV.D.1. Therefore, in my opinion a POSITA would have recognized *Sabottke*'s software system is executed by hardware on a computing device.

80. *Sabottke* applies supervised training algorithm to train its SVM machine learning prediction models with vulnerability data obtained from Twitter hacker discussions and vulnerability databases to estimate a likelihood of exploitation of the software vulnerability. *See generally*, Claim 1. The output from this prediction model predicts it is likely to be or not an actual exploitation of the software vulnerability before disclosure based on hacker communications from the Twitter feed. *Supra*, Claim 1(f)-1(f)(1).

14. *Claim 15 - The computing device of claim 14, wherein the prediction model is at least one classification model that outputs from features of the software vulnerability a score indicative of the likelihood of exploitation.*

81. *Sabottke* teaches its Twitter exploit detector prediction model utilizes an SVM classification model that outputs a likelihood of exploitation. *Supra*, Claim 1(f) (discussing computing likelihood of exploitation). *Sabottke* teaches or suggest outputting a score indicative of the likelihood of exploitation. For example, *Sabottke* teaches outputting “predictions” that are aggregated and compared to “SVM prediction thresholds.” *Sabottke*, 12-13. These predictions are “scores” in that they are values compared to a data point’s location to a hyperplane. Indeed, in my opinion a POSITA would have understood *Sabottke*’s SVM output predictions, which are predictions based on data points separated by a hyperplane, are “scores” as claimed. E.g., positive values which indicate being on one side of the hyperplane are “exploited” and the magnitude indicates the severity and in other words, the score.

See, e.g., *Bozorgi*, 8 (discussing that the distance to the SVM hyperplane “serves as a natural score for the exploitability of a vulnerability”).

15. *Claim 18(Pre): A method of prioritizing vulnerabilities using cyber threat intelligence, comprising:*

82. To the extent the preamble is limiting, *Sabottke* teaches methods for prioritizing vulnerabilities using cyber threat intelligence. *Sabottke*, 2 (describing “detection techniques” for “response prioritization”), 3 (“Specifically, we investigate techniques for minimizing false-positive detections—vulnerabilities that are not actually exploited—which is critical for prioritizing response actions”), 7 (“we envision utilizing our classifier as a tool for prioritizing the response to vulnerability disclosures”), 15 (“Twitter-based classifiers can be employed to guide the prioritization of response actions after vulnerability disclosures”).

16. *Claim 18(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations including:*

83. *Sabottke* teaches or suggests a processor in operable communication with memory that stores instructions for performing the claimed operations. *Supra*, Claim 1(a) (further noting a POSITA would have found it obvious to implement *Sabottke*'s "Twitter-based exploit detector" on a known computing system by executing software instructions in memory via a processor).

17. *Claim 18(b): accessing a plurality of datasets associated with a predetermined set of data sources, at least a portion of the*

*plurality of datasets defining training data including hacker communications;*

84. *Sabottke* teaches accessing a plurality of datasets from predetermined data sources (*Sabottke*, Fig. 1), including hacker communications from tweets. *Supra*, Claim 1(b).

18. *Claim 18(c): applying learning algorithms to the training data to generate a predictive model configured to predict a likelihood of exploitation of respective software vulnerabilities;*

85. *Sabottke* teaches applying supervised learning algorithms to training data to generate an SVM predictive model configured to predict the likelihood of exploitation of software vulnerabilities. *Supra*, Claim 1(d).

19. *Claim 18(d): accessing one or more features associated with a software vulnerability; and*

86. *Sabottke* teaches accessing one or more features associated with a software vulnerability, including “CVSS Information” and “Database Information” features listed in Figure 2. *Supra*, Claim 1(e).

20. *Claim 18(e): “computing, by applying the one or more features to the predictive model, one or more values defining a likelihood of exploitation associated with the software vulnerability, wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data.”*

87. *Sabottke* teaches or suggests computing a value defining a likelihood of exploitation associated with the software vulnerability using its trained SVM predictive model. *Supra*, Claim 1(f). This value predicts an actual exploitation before


disclosure to a public database based on hacker communications from the training data. *Supra*, Claim 1(f)(1).

## **VI. OPINION – GROUND 2**

### **A. Obvious to Combine *Smyth* and *Phillips***

#### *1. Overview of Smyth*

88. *Smyth* teaches a “predictive engine for analyzing existing vulnerability information to determine the likelihood of a vulnerability being exploited by malicious actors[.]” *Smyth* (Ex. 1005), 1:7-11. An embodiment using this predictive engine to determine the likelihood a current vulnerability may be exploited is depicted in Figure 1 below:



*Id.*, Fig. 1. Data processor 120 collects information from a “plurality of data sources 110[.]” *Id.*, 5:40-52. Data sources 110 can include data from the U.S. National Vulnerability Database (NVD), the common vulnerabilities and exposures (CVE) system, and other public or proprietary data sources such as “online discussions” and “discussion forums.” *Id.*, 3:35-6:16, Fig. 2 (listing example data sources). The system then trains predictive models with this data, the models are combined into a “prediction ensemble 155 for use by the prediction engine 160[.]” and a “prediction 170” is generated based on the inputs that indicates “the likelihood the current

vulnerability” will be exploited. *Id.*, 4:4-34. Namely, the “prediction engine 160 generates a prediction based on processed current vulnerability information 125[,]” such as online discussion data, “and a prediction ensemble 155.” *Id.*, 10:61-64, 11:3-8 (“current vulnerability information 125 may comprise...more current information such as online discussions 240”), 14:33-15:14.

89. Because *Smyth* and the ’897 Patent are both directed to assessing cybersecurity risks for vulnerabilities, *Smyth* is in the same field of endeavor as the ’897 Patent. Compare *Smyth*, Abstract (teaching a system for “determin[ing] the likelihood of a vulnerability being exploited by malicious actors”), 1:7-11, 2:25-50, 4:4-6 with ’897 Patent, 1:24-28 (describing “assessing the likelihood of exploitation of software vulnerabilities”), 4:1-27, 4:41-44. Accordingly, in my opinion a POSITA would have recognized that *Smyth* is analogous art to the ’897 Patent.

90. The ’897 Patent describes common problems in the art for assessing cybersecurity risks for vulnerabilities, including how to “leverage[] machine learning techniques on features derived from” online discussions to predict whether certain vulnerabilities will be exploited. ’897 Patent, 4:19-27; see also *id.*, 8:41-45 (describing marketplaces and forums to gather information about vulnerabilities and exploits), 13:47-50 (discussing trade of real-world exploits on hacking forums). *Smyth* similarly addresses the problem of leveraging machine learning techniques to analyze features from discussions on online forums to “determine the likelihood a

new vulnerability may be exploited.” *Smyth*, 2:6-21; *see also id.*, 4:44-47, 5:10-15. Furthermore, *Smyth*’s method of determining the likelihood a vulnerability will be exploited is used to “prioritize allocation of the limited available personnel and resources for securing a network or a system,” just like the ’897 Patent. *Compare, id.*, 2:16-21; *with* ’897 Patent, 3:27-30 (describing use of exploit prediction for “patch prioritization”). As such, *Smyth* is reasonably pertinent to problems addressed by the ’897 Patent and is analogous art.

## 2. *Overview of Phillips*

91. *Phillips* teaches collecting “metadata” associated with “Dark Web Forums” to assess and identify “the social structure of the networks.” *Phillips* (Ex. 1006), 1. *Phillips* applies various “Social Network Analysis (SNA) techniques in order to identify a range of potentially ‘important’ members of the forums.” *Id.* In these techniques, forum posts are extracted and metadata is pulled from each forum to build a network; a social network is created representing each forum; SNA metrics are run on these social networks; and unsupervised machine learning is applied to identify important forum members. *Id.*, 2. An outline of this approach is shown below:



*Id.* Example SNA metrics ran on this metadata are described in Table I. *Id.*, 3. These include metrics described in the '897 Patent's specification. *Id.*, (describing, e.g., "Degree Centrality[,] "Betweenness Centrality Score[,] "Pagerank Score"), '897 Patent, 21:6-46 ("[i]n-degree" and "[o]ut-degree centrality[,] "[b]etweenness centrality[,] "[p]agerank centrality").

92. The '897 Patent describes certain problems associated with analyzing social connections of users posting hacker-related content in forums to identify security risks. '897 Patent, 4:19-27; *see also id.*, 11:15-38 (describing use of a "set of features [that] contains measures computed from the social connections of users posting hacking-related content"), 21:6-46 (describing the use of "Network Structure Measures," "Centrality Measures[,] and "Importance Measures" from the social network). *Phillips* similarly addresses analyzing social connections of users posting hacker-related content in forums to identify security risks. *Phillips*, 1-2 (teaching "Social Network Analysis (SNA) techniques in order to identify a range of potentially 'important' members of the forums") 3, Table 1 (disclosing applying

“SNA metrics” to metadata from dark web forums). As such, *Phillips* is reasonably pertinent to a problem addressed by the ’897 Patent and is analogous art.


3. *Claim 1(Pre): A method for assessing a likelihood of exploitation of software vulnerabilities, comprising:*

93. To the extent the preamble is limiting, *Smyth* teaches methods for using a “predictive engine for analyzing existing vulnerability information to determine the **likelihood of a vulnerability being exploited** by malicious actors against a particular computer or network of computers.” *Smyth*, 1:7-11, 2:25-50, 3:63-67, 4:4-34, 11:22-25, claims 1 and 10. *Smyth* teaches that “malicious actors may use **vulnerabilities in existing software**, hardware, or combinations of software and hardware” to compromise the integrity of computers and networked computer systems. *Id.*, 1:19-24. *Smyth*’s methods pair data input to the predictive engine with particular vulnerabilities to determine a likelihood of exploitation. *Id.*, 5:40-6:10. This includes assessing the likelihood of exploitation of “[v]**ulnerable software and versions**[.]” See, e.g., *id.*, Figure 3.

4. *Claim 1(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations, including:*

94. *Smyth* teaches a “processor unit” that performs the steps of its described methods. *Smyth*, 4:4-34, 6:17-9:17, 12:30-60, Figs. 1, 5, 6, 7. Processor 704, boxed in red, is depicted below in operable communication with main memory 708 and


secondary memory 710, boxed in orange, through communications infrastructure 706:



*Id.*, Fig. 7 (annotated). The processor contains “modules or engines” that “utilize software, such as [] program code and/or programs stored in the main memory 708 or secondary memory 710” “to perform the functions discussed” in *Smyth. Id.*, 12:30-60; *see also id.*, 11:37-42 (“the method 100 of FIG. 1 may be implemented in the computer system 700 using...instructions for performing the tasks described for the method 100”).

5. *Claim 1(b): accessing a plurality of datasets associated with a predetermined set of data sources, the plurality of datasets including training data comprising hacker communications;*

95. *Smyth's* processor accesses a plurality of datasets from multiple predetermined data sources. *Smyth*, Abstract, 4:4-24. Namely, “data processor 120 which executes on processor device 704” processes information obtained from data sources 110 “so that a combined data set 123 of historical vulnerability information is obtained” along with “current vulnerability information” 125. *Id.*, 14:33-43, 14:61-65. “The information from the data sources 110 may be accessed by the data processor 120 through application programming interface (API) calls, web page scraping, periodic database imports, or other known techniques.” *Id.*, 5:49-52. This process is shown in Figure 4 below:



*Id.*, Fig. 4 (annotated); *see also id.*, Figs. 1-2.

96. *Smyth* accesses multiple datasets from the predetermined data sources identified in Figure 4 to construct a “combined data set” used to train predictive models. *Id.*, 4:6-14 (describing training prediction models 531-534 using the combined data set). This data includes “online discussions 240,” boxed in red. *Supra*, Fig. 4. “[O]nline discussions 240” include “those on social media and discussion forums.” *Id.*, 4:44-47, 6:30-36. While *Smyth* discloses accessing training data from online “discussion forums[,]” *Smyth* does not expressly discuss accessing “hacker communications” in those forums. *Phillips* teaches this.

97. *Phillips* teaches “crawling a variety of radical websites” to collect forum data from forums associated with the “Dark Web Forum Portal[.]” *Phillips*, 3. Table II below provides an overview of exemplary dark web forums:

| TABLE II AN OVERVIEW OF THE ENGLISH LANGUAGE FORUMS USED AS PART OF OUR EXPERIMENTS |            |            |            |             |            |                                       |
|-------------------------------------------------------------------------------------|------------|------------|------------|-------------|------------|---------------------------------------|
| Forum                                                                               | #Messages: | # Threads: | # Members: | Start Date: | End Date:  | Forum URL:                            |
| Ansar AlJihad Network (Ansar1)                                                      | 29492      | 11244      | 382        | 12/08/08    | 01/20/2010 | http://www.ansar1.info/               |
| Gawaher (Gawaher)                                                                   | 372499     | 53235      | 9269       | 10/24/2004  | 06/07/12   | http://www.gawaher.com                |
| Islamic Awakening (IslamicAwakening)                                                | 201287     | 32879      | 3964       | 04/28/2004  | 05/22/2012 | http://forums.islamicawakening.com    |
| Islamic Network (IslamicNetwork)                                                    | 91874      | 13995      | 2082       | 06/09/04    | 11/10/10   | http://talk.islamicnetwork.com        |
| Islamic Web-Community (Myiwc)                                                       | 25016      | 6310       | 756        | 11/05/00    | 02/19/2010 | http://www.myiwc.com/forums/index.php |
| Turn To Islam (TurnToIslam)                                                         | 335338     | 41654      | 10858      | 06/02/06    | 05/20/2012 | http://www.turnintonislam.com/forum/  |
| Ummah                                                                               | 1491957    | 91527      | 21013      | 04/01/02    | 05/18/2012 | http://www.ummah.com/forum/           |

*Id.*, 4. In my opinion, a POSITA would have understood dark web forums include “hacker communications” for several reasons. Starting with the disclosures of the ’897 Patent, the specification explains that several data sources are mined and used by the predictive models to predict a likelihood a vulnerability will be exploited. ’897 Patent, 6:50-59. This includes “data mined from DW (darkweb and deepnet)

markets and forums focusing on malicious hacking.” *Id.*; *see also id.*, 8:26-29 (“In the present disclosure, the exploit prediction model may be implemented to crawl websites on DW, both marketplaces and forums, to collect data relating to malicious hacking.”), 8:44-57, 11:29-38, 13:47-50, 20:46-22:41.

98. In fact, during U.S. prosecution of the application that resulted in the ’897 Patent, the examiner found that accessing hacker communications from darkweb forums was “well known in the prior art before the effective filing date of applicant’s claimed invention[.]” *’897 File History*, 458. Namely, the Examiner explained the following was known prior to the filing date of the ’897 Patent:

“i. information associated with discussion of vulnerabilities by users positing hacking-related content at a web forum prior to disclosure of a public vulnerability database (paragraphs 48-54; **AAPA** [(applicant admitted prior art)] **discuss vulnerability postings in darkweb forums** and ZDI communities from at least 2015. The ZDI communities involve submitters trying to get rewards as the first to submit information about a vulnerability while **the darkweb forums may be composed of users communicating and exchanging malicious, possibly criminal information**, so in each case, the many of the content may be private and submitted before they were disclosed to a public vulnerability database)”

“ii. Vulnerability databases being public (paragraph 46-47; NVD)”

“iii. Information associated with exploitation of vulnerabilities as obtained through exploits data associated with real world exploitation of the vulnerability (paragraphs 48-54; In particular Exploitation database used by Offensive Security Training).”

*Id.*, 458-459.

The intrinsic record, therefore, supports that darkweb forums include hacker communications and this concept was widely recognized, including by applicant admitted prior art, at the time of filing the '897 Patent.

99. I agree it was common knowledge to one of skill in the art that certain darkweb forums would contain hacker communications. For instance, it was known that “[h]acker communities are dark networks[.]” *Macdonald* (Ex. 1015), 2. “Hackers often congregate within underground communities, most commonly in the form of...online forums.” *Benjamin* (Ex. 1013), 4. Darkweb forums therefore include hacker communications. With respect to the specific exemplary forums discussed in *Phillips*, while referencing mostly jihadist sites, these forums were known to include hacker communications from the dark web. *CTC Sentinel* (Ex. 1016), 1 (titled “[t]he risks posed by Jihadist Hackers” and describing “Jihadist Hacker” groups communicating in online forums). *CTC Sentinel* terms these hackers as “hactivists” “who perform[] cyber attacks for a movement or cause” for example on “online forums.” *Id.*; see also *id.*, 10 (describing publishing hacker communications on “Islamist online forums”). Moreover, other works have expressly recognized that *Phillips* was focused on Dark Web forums of terrorist and hacker groups. *Guetler* (Ex. 1017), 77 (describing a few works that “have focused on Dark Web forums of terrorist and hacker groups” including *Phillips*). As such, in my opinion a POSITA would have understood *Phillips* teaches or suggests crawling

hacker communications from the dark web. It would have thus been obvious, and a POSITA would have been motivated, to crawl dark web forums as taught by *Phillips* for hacker communications for use in *Smyth*'s predictive model for several reasons.

100. In my opinion, a POSITA would have been motivated to crawl dark web forums in *Smyth* pursuant to *Phillips* to provide a rich repository of current hacker communications for *Smyth*'s predictive engine to analyze. A POSITA would have recognized at the time that these current hacker communications would include new or otherwise not yet used malicious code for keeping *Smyth*'s predictive engine up-to-date. For example, it was commonly known that “[m]alicious software, hacking tutorials, and other resources intended to help conduct cybercrime can be commonly found within hacker communications” on forums. *Benjamin*, 4. Knowing this, a POSITA would have been motivated to use *Phillips*' method of crawling dark web forums to update and maintain *Smyth*'s predictive engine. This would facilitate *Smyth*'s goal to obtain existing threat information such that *Smyth*'s system can predict the likelihood of exploitation of “newly discovered” vulnerabilities. *Smyth*, 2:6-21, 14:53-57; see also *id.*, 6:33-34 (“more current information [] may be found in online discussions 240”).

101. In my opinion, a POSITA would have had a reasonable expectation of success modifying *Smyth* to crawl dark web forums for hacker communications as taught by *Phillips* for several reasons. *Smyth* already teaches analyzing “online

discussions 240” / “discussion forums” to determine the likelihood of a vulnerability being exploited. *Smyth*, 4:35-47, Fig. 2. *Smyth* would thus utilize the same or similar hardware to crawl dark web forums containing hacker communications to implement its methods to predict the likelihood of exploitation of vulnerabilities. For example, *Smyth* discloses its data processor 120 processes data sources 110, including online discussions 240, to place the data “into a format suitable for consumption by predictive models 531, 532, 533, 534.” *Id.*, 4:4-55; *see also id.*, 6:18-36 (teaching data processor 120 “generates processed current vulnerability information 125 from the information from the data sources”). Furthermore, *Smyth* specifies that “[t]he information from the **data sources 110 [including online discussions 240] may be accessed by the data processor 120 through** application programming interface (API) calls, **web page scraping**, periodic database imports, or other known techniques.” *Id.*, 5:40-63. As such, a POSITA would have understood the hardware of *Smyth*’s system would readily crawl dark web forums for hacker communications using the same processor that accesses online discussions.

102. *Smyth* also explains that “[o]ther data sources 260 may also be considered” and that “[e]ach of the data sources 210, 220, 230, 240, 250, 260 shown in Fig. 2 are merely examples and do not limit the number of type of data sources which may be considered by the method 100 described.” *Smyth*, 4:47-52. *Smyth* therefore contemplates examining other sources of data and it would have been well

within the skill of a POSITA to modify *Smyth* to examine hacker communications in dark web forums. In my opinion, a POSITA would have understood that *Smyth*'s disclosure of "online discussions 240" from which processor 120 may obtain information from via "web page scraping" would have been readily modified to scrape the online discussions in dark web forums. *Id.*, 5:40-63. Crawling and analyzing dark web forums was a well-known strategy for early detection of new emerging cyber threats. *See above*, Background of Technology. For example, *Macdonald* described analyzing the "communication activity" of hacker on forums to help identify threats. *Macdonald*, 1. Furthermore, as stated above, during U.S. prosecution of the application that resulted in the '897 Patent, the examiner found that accessing hacker communications from darkweb forums was "well known in the prior art before the effective filing date of applicant's claimed invention." '897 *File History*, 458. As such, a POSITA would have had success implementing that known strategy in *Smyth*'s system already equipped to analyze discussion forums.

6. *Claim 1(c): accessing features from the plurality of datasets that include measures computed from social connections of users posting hacking-related content;*

103. *Smyth* teaches accessing features from its plurality of datasets using several feature extraction techniques. For example, *Smyth* teaches using "unstructured text analysis 430" "on certain portions of the vulnerability information so that features may be derived." *Smyth*, 7:22-24. "[O]nline discussions 240 may

include forum posts formed substantially of unstructured text.” *Id.*, 7:28-30. This “unstructured text is analyzed [] to facilitate consideration by the predictive models[.]” *Id.*, 7:33-35; *see also id.*, 6:18-42 (“information may be derived from the data sources 110 for use in the combined data set 123 and the processed current vulnerability information 125”).

104. While *Smyth* teaches feature extraction techniques, *Smyth* does not expressly discuss accessing features that include “measures computed from social connections of users posting hacking-related content[.]” *Phillips*, however, teaches using “Social Network Analysis (SNA) techniques” to access social measures in dark web forums that overlap with measures described in the ’897 Patent’s specification.

105. Namely, *Phillips* uses the SNA metrics in Table I below to access measures computed from user’s posting hacker-related content in dark web forums:

| TABLE I. A TABLE OUTLINING THE METRICS USED TO EVALUATE OUR SOCIAL NETWORK |                                                                                                                                                          |
|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Attribute Name                                                             | Description                                                                                                                                              |
| Sent Messages (SM)                                                         | The number of emails sent by an employee.                                                                                                                |
| Received Messages (RM)                                                     | The number of emails received by an employee.                                                                                                            |
| Degree Centrality (DCS)                                                    | The number of distinct employees within the network that an employee has sent emails to.                                                                 |
| Betweenness Centrality Score (BCS)                                         | The betweenness centrality measure for an employee.[20]                                                                                                  |
| Pagerank Score (PRS)                                                       | The PageRank score an employee. [21]                                                                                                                     |
| Markov Ranking (MR)                                                        | The markov ranking of an employee. [22]                                                                                                                  |
| HITS Authority Score (HAS)                                                 | The authority score for an employee (if several users with high hub weights send an email t the user then they will have a higher authority score). [23] |
| HITS Hub Score (HHS)                                                       | The hub score for an employee (if the user sends emails to users with high authority scores then they will have a higher hub score). [23]                |
| Clique Score (CS)                                                          | The number of cliques (maximal subgraphs) an employee is in using the Bron and Kerbosch algorithm.[24]                                                   |
| Weighted Clique Score (WCS)                                                | The weighted clique score for each user, weighted by the number of users within each clique.                                                             |
| Average Distance Score (ADS)                                               | The average distance between the user and all other users in the graph.                                                                                  |
| Clustering Coefficient (CC)                                                | The extent to which vertices in a graph tend to cluster together. [25]                                                                                   |

*Phillips*, 3 (annotated). The highlighted features in Table I above substantively overlap with the social network measures discussed in the '897 Patent. *Compare* '897 Patent, 21:6-46 (“[i]n degree[,]” “[o]ut-degree[,]” “[i]n-degree centrality[,]” “[o]ut-degree centrality[,]” “[b]etweenness centrality[,]” “[p]agerank centrality”).

106. The SNA metrics outlined in Table I are the “top-ranked SNA metrics from previous work” by the authors of *Phillips*. *Phillips*, 2. *Phillips* utilizes those top-ranked metrics on “the forums collected as a part of the Dark Web Forum Portal dataset.” *Id.*, 3; *see also id.*, Abstract. (“For this experiment, we conducted our investigation on five forums collected in previous work as part of the Dark Web Forum portal and built upon the tool support created in our previous research in order to visualise and analyse the network” by utilizing “SNA metrics” on the “communication metadata” from these dark web forums). Namely, *Phillips* “first calculate[s] a range of SNA metrics to better understand the group members” in dark web forums and “then appl[ies] unsupervised learning in order to create clusters that would help classify the Dark Web Forum users[.]” *Id.*, *see also id.*, 1 (“we set out to trial various Social Network Analysis (SNA) techniques in order to identify a range of potentially ‘important’ members of the [dark web] forums”).

107. In this process, *Phillips* teaches collecting dark web forum posts, extracting metadata from these posts, building a social network, and then analyzing that social network with SNA techniques and machine learning. *Phillips*, 1-2 (*see*

“II. Methodology”). An outline of the “approach to extract social structure from online communications” in dark web forums is illustrated below in Figure 1:



*Id.*, 2 (Fig. 1).

108. More specifically, *Phillips* teaches evaluating “forums collected as part of the Dark Web Forum Portal dataset.” *Id.*, 3. From these dark web datasets, *Phillips* teaches creating “social networks.” *Id.*; *see also id.*, 2. Then, “to measure the graph properties of each node that reflect importance”, *Phillips* applies several SNA metrics to the social network graph. *Id.*, 3. These metrics are listed in Table I and depicted in Figure 3. *Id.*, 3. *Phillips* then applies machine learning clustering algorithms to classify forum members. *Id.*; *see also id.*, 4 (Table III. Outlining clustering algorithms and SNA metrics outputs on dark web datasets). Using the “Ansarl” dark web forum as an exemplary dataset, *Phillips*’ results demonstrated that “user 17” was the most influential actor across that forum. *Id.*, 4 (“The table shows that all four networks identified user 17 as the most influential user and agreed on the top three users (but had different orderings of users 0 and 14)”). Therefore,

*Phillip's* teaches accessing features comprising measures computed from social connections of users posting hacking-related content to classify users in dark web forums. And as described above, in my opinion a POSITA would have understood that certain dark web users post hacking-related content. *Supra*, Claim 1(b). *Phillips* therefore analyzes hacking-related content in dark web posts.

109. In my opinion, a POSITA would have been motivated to modify *Smyth* to implement *Phillips'* techniques for identifying important members on forums to improve *Smyth's* methods for exploit prediction. Correlating exploit and vulnerability information from discussion forums with social measures from the user posting in the forum would improve *Smyth's* threat predictions. A POSITA would have readily recognized that such a system would allow *Smyth* to prioritize threats associated with a user known to be more active or otherwise employ the malicious software. For example, *Smyth's* system generates threat levels for current vulnerabilities. *Smyth*, 4:29-34, 7:7-8, 11:26-33, 15:8-14. Incorporating social measures related to user postings in forums would improve the exploitation prediction quality and the reliability of threat level assessment. A POSITA would have understood, and indeed *Phillips* teaches, that this type of social network analysis would “reflect[] the importance of the individual within the network.” *Phillips*, 2. This would have provided additional features for analysis in *Smyth's* system and would have increased the performance of *Smyth's* threat predictions by

providing a correlation between the “importance” of the user posting malicious information and the likelihood of using the malicious information.


110. In my opinion, a POSITA would have had a reasonable expectation of success modifying *Smyth* to utilize *Phillips*’ SNA metrics to access features computed from social connections of users posting hacking-related content. *Smyth* already teaches feature extraction techniques that extract features from discussion forum data to feed its predictive models. *Smyth*, 7:22-35. And *Smyth* expressly states that its embodiments “may employ other techniques not specifically described here to take the information from the data sources 110 and provide it in a manner suitable for the prediction engine 160.” *Id.*, 6:39-42; see also *id.*, 8:7-10 (“The above descriptions of the techniques used to process information from data sources 110 are merely examples and should not be considered to be an exhaustive list of the contemplated techniques”). Thus, it would have been well within the skill of a POSITA to incorporate other well-known techniques, like *Phillips*’ SNA metrics, in *Smyth*’s system to process the dark web forum data for *Smyth*’s predictive models. For example, in the combined *Smyth-Phillips* system, a POSITA would have understood that similar to *Smyth*, metadata from the forum data would be extracted and prepared for processing. See generally, *Smyth*, 5:40-8:10. The extracted data would then be used “to create a social network representation of each forum,” which would be used to “calculate[] the top-ranked SNA metrics[.]” *Phillips*, 2. Once the

SNA metrics are calculated, the system would perform machine learning or network analysis on the social networks to determine the importance of the users posting malicious content.

111. Indeed, these extraction techniques and social network analyses were well known at the time of the '897 Patent. The same author of *Phillips* previously disclosed the method of extracting SNA metrics from social networks to establish a hierarchical structure of the users. Elizabeth Phillips et al., Applying Social Network Analysis to Security, (Ex. 1019), Abstract; *see also, Benjamin*, Abstract, 6 (teaching “text parsers” for “extract[ing] messages and user information from both of the communities” and using the corresponding features to capture “hacker discussion quality” and “hacker[] involvement in their community”). Furthermore, prior publications from some of the inventors of the '897 Patent establish that such extraction techniques from social networks for the purposes of identifying emerging cyber threats were well known prior to the '897 Patent's earliest effective filing date. *Nunes*, (Ex. 1007), Abstract. Being such a well-known and well described technique, in my opinion a POSITA would have had success in modifying *Smyth* to access features from the datasets that include measures computed from social connections of users posting hacking-related content.

7. *Claim 1(d): applying learning algorithms to the training data to generate classification models that are configured to predict class labels defining a likelihood of exploitation of respective software vulnerabilities;*

112. *Smyth* teaches applying various learning algorithms (e.g., 531-534) to the training data to generate prediction ensemble 155 configured to predict class labels defining a likelihood of exploitation of software vulnerabilities. *Smyth*'s training and model generation process is illustrated in Figure 5 below:



*Smyth*, Fig. 5. As shown, *Smyth* trains learning algorithms 531-534 with data from its combined data set 123. *Id.*, 9:20-26, 10:25-27. Combined data set 123 includes data from at least a portion of the data sources, including data captured from online discussions 240. *Id.*, 10:5-12 (“At least a portion of the combined data set 123 is

relied upon for retraining”), Fig. 2 (capturing data from online discussions 240), Fig. 4 (combined data set 123 generated in part from online discussions data 240). In *Smyth’s* modified system, this includes training data captured from hacker communications as taught by *Phillips. Supra*, Claim 1(b). In my opinion, a POSITA would have understood that once learning algorithms 531-534 are trained with training data from combined data set 123, algorithms 531-534 form a trained classification model. It was known learning algorithms learn from training data and features from that data to identify patterns and relationships from new input data for classification. *See, e.g., Bozorgi*, 4 (Section 4.2 describing training an SVM learning algorithm for classification). Thus, *Smyth-Phillips* teaches applying learning algorithms to training data comprising hacker communications from combined data set 123 to generate classification models.

113. Indeed, after training, learning algorithms 531-534 are combined into prediction ensemble 155. *Supra, Smith*, Fig. 5. Prediction ensemble 155 comprises trained algorithms 531-534 and therefore comprises “classification models” as claimed. The ’897 Patent provides exemplary classification models that predict “class labels[.]” ’897 *Patent*, 14:30-33, 17:64-65, 21:57-59. Some example classification models are Decision Trees, Naïve Bayes classifiers, Logistic Regressions, Support Vector Machines, and Bayesian Networks. *Id.*, 22:66-23:32. *Smyth* discloses multiple classification models including “a tree-based model 531, a

statistical model 532, a binary regression model 533, and other models relying on other algorithms 534[.]” *Smyth*, 9:42-45. *Smyth* teaches “the data processor 120 creates a combined data set 123 of historical vulnerability information obtained from data sources 110 for the prediction engine retrainer 140 so that the predictive models 531, 532, 533, 534 may be trained against the historical information.” *Id.*, 4:10-14.


114. *Smyth*’s models are combined into ensemble 155 which “generates prediction 170” that “provides a representation of the likelihood of the current vulnerability will [be] exploited by malicious actors against a particular computer network.” *Smyth*, 11:18-25. *Smyth* therefore teaches “classification models” (e.g., prediction ensemble 155) that predict “class labels” (e.g., exploitation) “defining a likelihood of exploitation of respective software vulnerabilities” (e.g., prediction 170). *See infra*, ¶118 (describing class labels and likelihood of exploitation in more detail).

115. *Smyth*’s ensemble outputs prediction 170. *Smyth*, 15:7-14. Prediction 170 computes the “likelihood of exploitation” consistent with the ’897 Patent’s disclosures. The ’897 Patent discusses certain exemplary classifiers that output a similar prediction indicating the likelihood of exploitation of a software vulnerability. *See, e.g.*, ’897 Patent, 22:66-23:32 (describing exemplary supervised learning approaches, including SVM, Naïve Bayes, Bayesian Network, Decision Tree, and Logistic Regression). Similarly, “prediction 170” “provides a

representation of the likelihood of the current vulnerability will be exploited by malicious actors against a particular computer network.” *Smyth*, 11:18-25. Thus, *Smyth*’s prediction 170 defines the likelihood of exploitation of respective software vulnerabilities. Indeed, *Smyth*’s predictor is a statistical engine and cannot predict with *certainty*, but only with some *uncertainty*, hence, can provide only a *likelihood*.

8. *Claim 1(e): accessing one or more features associated with a software vulnerability; and*

116. *Smyth* teaches accessing features associated with a software vulnerability, including features from NVD, the MITRE CVE List, VirusTotal Samples and Reports, Online Discussions, Proprietary Data Sources, and Other Data Sources. *Smyth*, 4:36-55. This process is shown below in Figure 4 below with datasets boxed in red and feature extraction techniques boxed in orange:




*Smyth*, Fig. 4 (annotated). Feature extraction techniques 420-445 extract features from datasets 210-260 that are associated with a software vulnerability. Namely, *Smyth* teaches that “data processor 120 creates processed current vulnerability information 125 by obtaining information about a current vulnerability from the data sources[.]” *Id.*, 4:21-24, 6:18-42 (“certain information may be derived from the data sources 110” including “current vulnerability information 125”). This includes NVD features such as “vulnerability descriptions” or “CVSS” base score. *Id.*, 6:56-62, 7:18-21, Fig. 3. These features overlap with vulnerability features described in the ’897 Patent. *See* ’897 Patent, Fig. 1B (“NVD” features), 6:60-67 (“features

include...vulnerability description”), 6:67-7:1 (“Additional features are included from NVD namely, CVSS score[.]”).

9. *Claim 1(f): computing, by applying the one or more features to the classification model, a class label defining one or more values defining a likelihood of exploitation associated with the software vulnerability,*

117. *Smyth* teaches or suggests a prediction engine 160 that applies one or more of the above current vulnerability features to prediction ensemble 155 to compute a class label (e.g., exploitation) defining one or more values (e.g., exploited; not exploited); defining a likelihood of exploitation (e.g., prediction 170). *Supra*, Claim 1(d). Indeed, *Smyth*’s predictor is a statistical engine and cannot predict with *certainty*, but only with some *uncertainty*, hence, can provide only a *likelihood*. This process is illustrated in Figure 6 below:



*Id.*, Fig. 6. Processed current vulnerability information 125 includes the features derived from the datasets described above. *Supra*, Claim 1(e). This includes “information from the data sources 100 that include more current information such as online discussion 240[.]” *Smyth*, 11:3-7. *Smyth* teaches “[t]he prediction engine 160 takes [] processed current vulnerability information 125 and uses the prediction

ensemble 155...to generate a prediction 170 of the likelihood the current vulnerability may be employed[.]” *Id.*, 4:24-29, 10:61-64, 11:18-25.

118. *Smyth*’s prediction ensemble computes the exploitability of a vulnerability. The class label is therefore the variable “exploitability,” and that variable is defined by values such as “exploited” or “not exploited.” Indeed, it was known in the art for a prediction ensemble to output a class label defining a value defining the likelihood of the prediction. For example, the thesis by Michel Edkrantz describes an exemplary “ensemble learning algorithm.” *Edkrantz*, (Ex. 1028), 25. *Edkrantz* describes its ensemble algorithm operates by “aggregat[ing] and vot[ing] the classification of several independent algorithms” to eventually compute a “final prediction.” *Id.*, 25-26. In this case, the ensemble’s final prediction is a binary prediction “to predict whether a vulnerability will get exploited or not.” *Id.*, 14, 53; *see also, id.*, 45-47 (describing the use of datasets containing “unexploited” and “exploited” vulnerabilities and testing how accurate the ensemble computes the correct binary classification). In my opinion, a POSITA would have understood that *Smyth*’s ensemble would similarly compute an exploitation class label defining the binary values (i.e., “exploited” or “not exploited”) that predict whether a vulnerability will (likely) be exploited or not as was well known in the art. This is consistent with the ’897 Patent’s disclosures. ’897 Patent, 17:63-65 (“variable (here the class label  $y^c$  {exploited, not exploited})). In *Smyth*, *Smyth*’s prediction

“ensemble” includes various exemplary predictive models. *Smyth*, 4:14-17; *see also id.*, 9:20-45 (“prediction ensemble 155 defines a flexible structure of the predictive models 531, 532, 533, 534”). The ensemble is used to determine the exploitability of a vulnerability, e.g., its likelihood to be exploited, and is defined by prediction 170. *Smyth*, 11:22-25 (“prediction 170 provides a representation of the likelihood [] the current vulnerability will [be] exploited”). This likelihood is therefore a binary value, indicated exploited or not exploited based on prediction 170. Again, *Smyth*’s predictor is a statistical engine and cannot predict with *certainty*, but only with some *uncertainty*, hence, can provide only a *likelihood*.

10. *Claim 1(f)(1): wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data.*

119. *Smyth-Phillips* teaches this limitation. During prosecution, Applicant provided context for analyzing vulnerabilities “before disclosure.” In a July 13, 2023, request for continued examination after a final rejection, Applicant differentiated U.S. Patent Pub. No. 2017/0214710 to Seiver et al. (“Seiver”), stating:

Seiver ***does not discuss predicting a likelihood of exploitation of a vulnerability discussed online prior to disclosure to a public database or prior to actual exploitation of the vulnerability***; i.e., Seiver does not disclose the features of “wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data”.

'897 *File History*, 480-81 (emphasis in original). Per Applicant, “Seiver is not directed to actively ‘scouting’ the web for vulnerabilities that have a high likelihood of exploitation and instead appears to be more directed towards defensive strategies for protecting a network.” *Id.*; *see also id.*, 481 (differentiating another reference based on this same limitation for not “consider[ing] text information associated with discussion of vulnerabilities by users posting hacking-related content at a web forum prior to disclosure to a public vulnerability database”).

120. Unlike the references cited during prosecution, *Smyth* expressly scouts text information from “online discussions 240” as part of current vulnerability information 125 to generate a prediction on how likely that vulnerability will be exploited. *Supra*, Claim 1(b). As stated, “prediction engine 160 generates a prediction based on processed current vulnerability information 125[.]” *Smyth*, 10:61-64. “[C]urrent vulnerability information 125 contains vulnerability information about a threat that is newly discovered or currently spreading through computer networks” and includes “more current information such as online discussions 240[.]” *Id.*, 11:1-8; *see also id.*, 14:55-57 (“The current vulnerability may be newly discovered, newly relied upon by malicious actors, or the like”). In other words, *Smyth*’s prediction engine 160 proactively scouts online discussions for newly discovered vulnerabilities that have not yet been added to a public database or not yet exploited by scouting for vulnerabilities mentioned in online discussions.

121. Indeed, *Smyth* discusses associating a newly discovered vulnerability with historical vulnerabilities using a common identifier, including vendor-specific identifiers. *Smyth*, 5:64-6:10 (“certain vendors also maintain their own identifiers for vulnerabilities in their products”). Vendor-specific vulnerabilities are often private and not disclosed until the end of the process after a patch has been created. *See, e.g., Bozorgi*, 2 (“At the end of this process, vulnerabilities are documented and disclosed to the public” with information for “how [the vulnerability] can be mitigated”). *Bozorgi* explains the “precise timing of vulnerability disclosures depends considerably on how they were found and the policy of the organizations involved.” *Id.* “Some vulnerabilities are disclosed immediately upon discovery **while others may be kept private for significant periods of time to allow vendors to develop and test appropriate patches and mitigations.**” *Id.* It was therefore well known that vendor-specific vulnerabilities were not reported to public databases before a patch has been created. And *Smyth* analyzes these vendor-specific vulnerabilities. Thus, *Smyth* teaches processes for predicting the likelihood a “newly discovered” vulnerability mentioned in hacker communications in online forums will be exploited before disclosure to a public vulnerability database.

122. In my opinion, a POSITA would have understood that *Smyth*’s prediction process is consistent with the Applicant’s description of predicting the likelihood vulnerabilities will be exploited prior to disclosure to a public database or

prior to actual exploitation of the vulnerability. For example, *Smyth*'s disclosure of obtaining "newly discovered" vulnerability information, being contrasted to vulnerability information for a "currently spreading" threat, would have been understood by a POSITA to mean that the "newly discovered" vulnerability information was information "prior to actual exploitation of the vulnerability[.]" *Smyth*, 11:1-8. Furthermore, a POSITA would have understood "newly discovered" vulnerability information obtained from online discussions includes information not yet disclosed to public vulnerability repositories, like NVD.

123. Moreover, in *Smyth*'s modified system, *Smyth* scouts online discussions from dark web forums that include hacker communications and uses those hacker communications as training and testing data for *Smyth*'s predictive model. *Supra*, Claim 1(b) (proposed modification with *Phillips*). Thus, *Smyth-Phillips* teaches prediction 170 predicts an actual exploitation of the software vulnerabilities before disclosure based on hacker communications from the training data. Supporting this, *Smyth* expressly states that "[t]he prediction engine 160 takes this processed current vulnerability information 125 and uses the prediction ensemble 155 from the prediction engine retrainer 140 to generate a prediction 170 of the likelihood the current vulnerability may be employed in a customer's computer network." *Smyth*, 4:24-29. And as the modified *Smyth-Phillips* system uses hacker communications in

its training data (*see* Claim 1(b)), a POSITA would have understood the prediction 170 of *Smyth-Phillips* is based on the hacker communications from the training data.

11. *Claim 2: The method of claim 1, further comprising generating a plurality of estimation outputs based on the one or more values to derive an overall quantitative score.*

124. *Smyth* teaches its prediction ensemble generates a plurality of estimation outputs from multiple algorithms and uses those outputs to derive an overall prediction 170. Namely, *Smyth* teaches “prediction ensemble 155” uses different modeling schemes “to tailor the predictions 170 in a manner that provides a more useful prediction 170.” *Smyth*, 9:46-54. An example modeling scheme includes “stacking” the different algorithms in the ensemble to generate “predictions of multiple algorithms[.]” *Id.*, 10:57-58, 11:57-59. These plurality of predictions from the multiple algorithms in prediction ensemble 155 “generate a prediction 170 of the likelihood the current vulnerability may be employed[.]” *Id.*, 4:24-28, 11:18-21 (“the predictive models 531, 532, 533, 534” are executed “in accordance with the prediction ensemble 155 to generate a prediction 170”). Prediction 170 can be “converted into a threat level” based on defined thresholds, such as a low, medium, or high threat. *Id.*, 4:29-34, 7:6-8, 11:26-33.

125. Prediction 170 is derived from the multiple estimation outputs of the various algorithms in ensemble 155 and represents the overall likelihood the current vulnerability will be exploited. *Id.*, 11:18-25. Prediction 170 is compared to

thresholds to define a threat level and is therefore an “overall quantitative score” as claimed.

126. In my opinion, a POSITA would have understood *Smyth*’s prediction 170 to be a quantitative score. Indeed, “prediction 170 is a value that is within at least one of the thresholds previously defined by the computer administrator.” *Smyth*, 15:8-10. Therefore, it has to be quantitative and function as a quantitative score. Moreover, the prediction 170 may be used “to derive a threat level based on thresholds set by an administrator.” *Smyth*, 4:30-31. And “when the generated prediction 170 is to be converted into a threat level, the thresholds defining each threat level may be predefined by a computer administrator.” *Smyth*, 11:30-33. Therefore, a POSITA would have understood that prediction 170 can function as an overall quantitative score.

12. *Claim 3: The method of claim 1, wherein the plurality of datasets include vulnerability data for vulnerabilities that are publicly disclosed and obtaining exploits data for exploits that were used in real world attacks.*

127. *Smyth* teaches accessing a plurality of datasets, including vulnerability data for publicly disclosed vulnerabilities (e.g., from the public National Vulnerability Database, MITRE CVE List, Other Data Sources, etc.) and exploits data that includes exploits used in real world attacks (e.g., from VirusTotal Samples and Reports and Online Discussions). *Smyth*, 4:35-55, Fig. 2.

128. Regarding vulnerability data for vulnerabilities that are publicly disclosed, *Smyth* teaches “NVD 210” is “a publicly available data source” that is based on the CVE list 220, and list 220 “serves [as] a list or dictionary of CVE identifiers for publicly disclosed vulnerabilities.” *Id.*, 4:40-42, 4:56-58. NVD and the MITRE CVE list provide a public repository for publicly disclosed vulnerabilities. These publicly accessible databases are commonly known among cybersecurity experts to provide information on known cybersecurity vulnerabilities, enabling researchers, developers, and organizations to identify and address security risks. Indeed, the ’897 Patent confirms that NVD “maintains a database of publicly disclosed vulnerabilities each one identified using a unique CVE-ID.” ’897 Patent, 7:47-49.

129. Regarding exploits data for exploits used in real world attacks, *Smyth* teaches accessing data from VirusTotal Samples and Reports and online discussions (*Smyth*, 4:44-47), both of which include real world exploit data. Widely known, VirusTotal provides publicly accessible samples and analysis reports that allow users to examine files, URLs, and other artifacts for malicious behavior, supporting threat intelligence and malware research efforts. Indeed, *Smyth* teaches VirusTotal “provides a service for analyzing suspicious files and URLs, and for detecting viruses, worms, Trojans, and other kinds of malware.” *Smyth*, 5:1-9. VirusTotal is a website that allows users to submit files or URLs that are analyzed by antivirus

scanners to identify real-world security threats, including exploits. *See* Ex. 1018 (VirusTotal.com screen capture) (VirusTotal lets you “see which CVE are being exploited by attackers in the wild at this moment or during the last few weeks, what kind of malware they are distributing, and what actors are behind”); *id.*, (providing data on malware with “tag:exploit”). In my opinion, a POSITA would have understood that VirusTotal provides exploit data. VirusTotal includes publicly available analysis reports, which often include file classifications, behavior tags, and detection labels indicating the presence of known exploits. Many of these files are tagged with specific CVE identifiers and flagged by multiple antivirus engines as exploitative, confirming their alignment with real-world exploits. Furthermore, the presence of samples linked to documented attacks or threat activity shows the exploits have been observed in active use. VirusTotal therefore includes real-world exploit data. As do online discussions.

130. In my opinion, a POSITA would have understood that the online discussions disclosed by *Smyth* would have contained real-world exploit data based on its description. *See, e.g., Smyth*, 7:28-33 (online discussions can show if “a vulnerability is being actively used”), 11:1-3 (online discussions contain information “about a threat that is newly discovered or currently spreading through computer networks”). For example, information describing whether “a vulnerability is being actively used” or “currently spreading through computer networks” indicates that a

real-world exploit is being actively applied, which necessarily encompasses “exploits that were used in real world attacks.” As such, *Smyth* teaches obtaining exploits data for exploits that were used in real world attacks by its VirusTotal disclosure and its online discussions disclosure.

13. *Claim 5: The method of claim 1, wherein certain features correspond to a known vulnerability obtained from the plurality of datasets.*

131. As discussed in Claim 1(e), *Smyth* teaches accessing certain features corresponding to a known software vulnerability, including features from NVD, the MITRE CVE List, VirusTotal Samples and Reports, Online Discussions, Proprietary Data Sources, and Other Data Sources. *Supra*, Claim 1(e). Feature extraction techniques 420-445 extract features from datasets 210-260 that are associated with a software vulnerability. *Smyth*, 4:21-24, 6:18-42 (“certain information may be derived from the data sources 110” including “vulnerability information 125”). Certain features correspond to a “known vulnerability[,]” such as, for example, NVD “vulnerability descriptions” or “CVSS” base score. *Id.*, 6:56-62, 7:18-21, Fig. 3. These are features from previously disclosed or “known” vulnerabilities. As I mentioned above at Claim 3, NVD and the MITRE CVE list provide a public repository for publicly disclosed vulnerabilities. These publicly accessible databases are commonly known among cybersecurity experts to provide comprehensive information on known cybersecurity vulnerabilities, enabling researchers,

developers, and organizations to identify and address security risks. '897 Patent, 4:31-40 (NVD comprises a comprehensive list of vulnerabilities disclosed, including a CVSS score), 7:47-49 (NVD “maintains a database of publicly disclosed vulnerabilities each one identified using a unique CVE-ID”), 10:1-14 (Table 2 – NVD description). As such, *Smyth* teaches certain features that correspond to a known vulnerability obtained from the plurality of datasets.

14. *Claim 10: The method of claim 1, further comprising: detecting, from the plurality of datasets, vulnerabilities that appear before an associated exploitation date.*

132. *Smyth* detects vulnerabilities that appear before an associated “exploitation date” by analyzing online discussion data for vulnerabilities mentioned, but not yet exploited. The '897 Patent defines an “exploitation date of a vulnerability” “as the date it was first detected in the wild.” '897 Patent, 9:42-43, 11:46-48 (vulnerability mentioned online before exploitation date), 12:29-30 (same).

133. As discussed in Claim 1(f)(1), *Smyth* teaches that “current vulnerability information 125 contains vulnerability information about a threat that is newly discovered” and includes “more current information such as online discussions 240[.]” *Supra*, Claim 1(f)(1); *Smyth*, 11:1-8; *see also id.*, 14:55-57 (“The current vulnerability may be newly discovered, newly relied upon by malicious actors, or the like”). In my opinion, a POSITA would have understood that such “newly

discovered” vulnerabilities would include ones prior to a “date it was first detected in the wild” as described by the ’897 Patent because it was *newly* discovered. Indeed, it was common knowledge that hacker discussions contain information about new vulnerabilities that have not yet been exploited. *See, e.g., Yampolskiy*, 13:52-55 (“hacker discussions regarding an entity can be collected and contextualized by weighting the discussions according to the severity and immediacy of **a potential breach based on the discussions**”). Thus, *Smyth* scouts online discussions for vulnerabilities mentioned that are “newly discovered” that have not yet been exploited.


15. *Claim 13: The method of claim 1, further comprising vectorizing text features derived from textual content of the plurality of datasets using a predetermined natural language process (NLP).*

134. *Smyth* teaches processing unstructured text in online forum discussions using a predetermined natural language process (NLP) such as latent Dirichlet allocation (LDA). *Smyth*, 7:22-47 (“the unstructured text may be processed by generative statistical models such as latent Dirichlet allocation (LDA)”). LDA is a natural language processing technique that analyzes vectorized text features to “determine the relationships between topics contained in the unstructured text in a manner that may be considered by the prediction models[.]” *Smyth*, 7:40-43. Indeed, in my opinion a POSITA would have understood that “Latent Dirichlet Allocation [LDA] is an NLP technique...to reveal latent dimensions in...topics, or in general

on any collection of discrete data” and that it receives “vectoriz[ing] words and documents” as input for topic modeling. *Barreras* (Ex. 1021), 4. *Smyth* therefore suggests vectorizing text features by processing unstructured text from its plurality of datasets using LDA.

16. *Claim 14: A computing device, configured via machine learning to apply a learned function to data associated with a software vulnerability to estimate a likelihood of exploitation of the software vulnerability, the learned function associated with a prediction model derived from at least one machine learning algorithm and a plurality of datasets associated with software vulnerabilities, the plurality of datasets including information associated with discussion of vulnerabilities by users posting hacking-related content, wherein the likelihood of exploitation predicts an actual exploitation of the software vulnerability before disclosure based on hacker communications from the plurality of datasets.*

135. *Smyth* teaches a computing device, depicted below, configured to perform the functions described in claims 1 and 14:




*Smyth*, Fig. 7; *supra*, claim 1. Specifically, *Smyth* applies a prediction engine 160 (“learned function”) configured via machine learning algorithms 531-534 to vulnerability data 125 to estimate a likelihood of exploitation of the software vulnerability. *See generally*, Claim 1. *Smyth*’s prediction engine 160 is associated with a prediction ensemble 155 (“prediction model”) derived from machine learning algorithms 531-534 trained with data from the plurality of datasets, such as NVD, MITRE CVE List, VirusTotal Samples and Reports, Online Discussions, Proprietary Data Sources, and Other Data Sources (“derived from at least one machine learning algorithms and a plurality of datasets associated with software vulnerabilities”). *Id.*, Claim 1(b)-1(f). *Smyth-Phillips* further teaches accessing data from online discussions from the dark web, including discussion of vulnerabilities by users

posting hacker-related content. *Id.*, Claim 1(b) (detailing motivation to combine). *Smyth*'s prediction engine 160 generates a prediction 170 that predicts an actual exploitation of the software vulnerability before disclosure based on at least hacker communications from the plurality of datasets. *Id.*, Claim 1(f)(1).

17. *Claim 15 - The computing device of claim 14, wherein the prediction model is at least one classification model that outputs from features of the software vulnerability a score indicative of the likelihood of exploitation.*

136. *Smyth*'s prediction engine 160 includes prediction ensemble 155 (“prediction model”) which is a classification model that applies one or more current vulnerability features to compute a prediction 170 defining the likelihood the software vulnerability will be exploited. This process is depicted below in Figure 6:



*Smyth*, Fig. 6. “This prediction 170 is a value that is within at least one of the thresholds previously defined by the computer administrator.” *Id.*, 15:8-10. *Smyth* discloses that the prediction “may then be used to derive a threat level based on thresholds set by an administrator for the customer’s computer network.” *Id.*, 4:29-31. Prediction value 170 is compared to thresholds to define a threat level and is

therefore a “score indicative of the likelihood of exploitation” as claimed. *Supra*, claim 2.

18. *Claim 17: The computing device of claim 14, wherein the computing device is further configured to apply the learned function to data associated with each of a plurality of new software vulnerabilities and compute a respective likelihood of exploitation for each of the plurality of new software vulnerabilities.*

137. *Smyth*’s prediction engine 160 applies its machine learning ensemble 155 (“learned function”) to online discussion data for vulnerabilities mentioned, but not yet exploited to compute a likelihood of exploitation. *Supra*, Claim 1(f), Claim 10. I note that the claimed “learned function” is synonymous with the claimed “classification models” in claim 1 and *Smyth*’s trained prediction ensemble teaches a “learned function” for the same reasons it teaches the classification models of Claim 1. As discussed in Claim 1(f), *Smyth* teaches that “current vulnerability information 125 contains vulnerability information about a threat that is newly discovered” and includes “more current information such as online discussions 240[.]” *Supra*, Claim 1(f); *Smyth*, 11:1-8; *see also id.*, 14:55-57 (“The current vulnerability may be newly discovered, newly relied upon by malicious actors, or the like”). *Smyth* therefore applies its ensemble 155 to data from new software vulnerabilities discussed in online forums and computes the respective likelihood of exploitation of these new vulnerabilities. *Supra*, Claim 1(f)(1).

19. *Claim 18(Pre): A method of prioritizing vulnerabilities using cyber threat intelligence, comprising:*

138. *Smyth* teaches methods for using a “predictive engine for analyzing existing vulnerability information to determine the likelihood of a vulnerability being exploited by malicious actors against a particular computer or network of computers.” *Smyth*, 1:7-11, 2:25-50, 3:63-67, 4:4-34, 11:22-25 claims 1 and 10. *Smyth* explains “[t]his likelihood of exploitation information helps prioritize allocation of the limited available personnel and resources for securing a network or a system.” *Id.*, 2:19-21. In my opinion, *Smyth* thus teaches or at least suggests a method of prioritizing vulnerabilities using cyber threat intelligence by generating predictions for the likelihood a vulnerability will be exploited, where these predictions help prioritize which vulnerabilities to expend resources to patch.

20. *Claim 18(a): utilizing a processor in operable communication with at least one memory for storing instructions that are executed by the processor to perform operations including:*

139. *Smyth* discloses a processor in operable communication with a main memory and secondary memory storing instructions that are executed by the processor for performing the claimed operations. *Supra*, Claim 1(a).

21. *Claim 18(b): accessing a plurality of datasets associated with a predetermined set of data sources, at least a portion of the plurality of datasets defining training data including hacker communications;*

140. *Smyth-Phillips* teaches accessing a plurality of datasets from predetermined data sources (*see Smyth*, Fig. 4), where these datasets include online

discussions 240 comprising hacker communications as modified by *Phillips* that is used to train *Smyth*'s prediction models. *Supra*, Claim 1(b).

22. *Claim 18(c): applying learning algorithms to the training data to generate a predictive model configured to predict a likelihood of exploitation of respective software vulnerabilities;*

141. *Smyth* applies various learning algorithms (e.g., 531-534) to the training data to generate prediction ensemble 155 configured to predict a likelihood of exploitation of software vulnerabilities. *Supra*, Claim 1(d).

23. *Claim 18(d): accessing one or more features associated with a software vulnerability; and*

142. *Smyth* teaches accessing features associated with a software vulnerability, including features from NVD, the MITRE CVE List, VirusTotal Samples and Reports, Online Discussions, Proprietary Data Sources, and Other Data Sources. *Supra*, Claim 1(e).

24. *Claim 18(e): computing, by applying the one or more features to the predictive model, one or more values defining a likelihood of exploitation associated with the software vulnerability, wherein the likelihood of exploitation predicts an actual exploitation of the respective software vulnerabilities before disclosure based on the hacker communications from the training data.*

143. *Smyth* teaches a prediction engine 160 that applies one or more of the above current vulnerability features to prediction ensemble 155 to compute a prediction 170. *Supra*, Claim 1(f) (noting *Smyth* computes the exploitability of a vulnerability defined by values such as exploited or not exploited that are defined by prediction 170 that indicates the likelihood of the prediction). *Smyth-Phillips* teaches

prediction 170 predicts an actual exploitation of the software vulnerabilities before disclosure based on hacker communications from the training data. *Supra*, Claim 1(f)(1).

## VII. OPINION – GROUND 3

### A. Obvious to Combine *Smyth*, *Phillips*, and *Nunes*

#### 1. Overview of *Nunes*

144. *Nunes* teaches “an operational system for cyber threat intelligence gathering from various social platforms on the Internet particularly sites on the darknet[.]” *Nunes*, Abstract. *Nunes*’ system includes a “parser” that “extract[s] specific information from...hacker forums (discussion regarding services and threats).” *Id.*, 2. *Nunes* filters out noise from the extracted information using classifiers “to filter out products and topics relating to drugs, weapons, etc. not relevant to malicious hacking.” *Id.*, 3.

145. Because *Nunes* and the ’897 Patent are both directed to assessing cybersecurity risks for vulnerabilities, *Nunes* is in the same field of endeavor as the ’897 Patent. Compare *Nunes*, Abstract, 5 (detecting “[z]ero-day exploits [that] leverage vulnerabilities”), 1 (providing information to determine “[w]hat vulnerabilities do the latest exploits target”) with ’897 Patent, 1:24-28 (describing “assessing the likelihood of exploitation of software vulnerabilities”), 4:1-27, 8:41-45, 13:47-50. Accordingly, in my opinion a POSITA would have recognized that *Nunes* is analogous art to the ’897 Patent.

146. The '897 Patent describes common problems in the art related to using online forum discussions to predict whether certain vulnerabilities will be exploited. '897 Patent, 4:19-27; *see also id.*, 8:41-45 (describing DW marketplaces and forums to gather information about vulnerabilities and exploits), 13:47-50 (discussing trade of real-world exploits on DW hacking forums). *Nunes* similarly addresses the problem of using online forum discussions to predict whether certain vulnerabilities will be exploited. *Nunes*, Abstract (“We focus our attention to collecting information from hacker forum discussions... We have developed an operational system for obtaining information from these sites for the purposes of identifying emerging cyber threats.”), *see also*, 1 (Table 2 showing data collected from 21 different “Forums” and following discussion asking “What vulnerabilities do the latest exploits target?”). As such, *Nunes* is reasonably pertinent to a problem described by the '897 Patent and is analogous art.

2. *Claim 4: The method of claim 3, further comprising: aligning the exploits data with the vulnerability data; and cleaning the exploits data of noise and predetermined portions of the exploits data that is irrelevant to associated software vulnerabilities.*

147. *Smyth* teaches the method of claim 3. *Supra*, Ground 1, Claim 3. *Smyth* further teaches an “identifier” aligns exploits data with vulnerability data “so that interested individuals can associate the specific malicious software with the vulnerability being used by the specific malicious software.” *Smyth*, 5:4-9. Each of the data sources 210-260 “rely on a common identifier associating information with

a particular vulnerability.” *Id.*, 4:52-65. *Smyth*’s predictive engine “relies on this common identifier to associate information about the vulnerability collected from multiple disparate sources 110 in a manner that is more easily considered by the machine learning, statistical, and other predictive models used.” *Id.*, 5:32-39. “[O]ther common identifiers may be employed so long as the information from the data sources 110 can be associated with the vulnerability.” *Id.*, 5:64-6:10.

148. While *Smyth-Phillips* teaches accessing exploits data, including from online discussions in dark web forums, and aligning that exploits data with vulnerability data for analysis by *Smyth*’s predictive models, *Smyth-Phillips* does not expressly describe further processing that exploits data to remove noise and irrelevant data. I understand that the inventors of the ’897 Patent dedicated this concept to the public, however, in 2016 when they published *Nunes*. Specifically, *Nunes* teaches a “parser” that “extract[s] specific information from...hacker forums (discussion regarding services and threats).” *Nunes*, 2. To detect relevant topics only and therefore reduce noise and filter out irrelevant data, classifiers are integrated into *Nunes*’ parser. *Id.*, 3 (“classifiers are integrated into the parser to filter out products and topics relating to drugs, weapons, etc. not relevant to malicious hacking”); *see also id.*, 2 (describing specific contributions of the paper, including implementing “learning models to separate relevant information from noise in the data collected from these online platforms”).

149. In my opinion, a POSITA would have been motivated to implement *Nunes*' parser with an integrated classifier in *Smyth*'s system for at least the benefits expressed in *Nunes*. *Nunes* acknowledges that "only a fraction of [online discussion] topics with posts on these forums contain information that is relevant to malicious hacking" and thus uses its parser equipped with a classifier "to identify relevant topics." *Nunes*, 4. Because this data is "noisy" and "not suitable to use directly as input to a learning model[,]" *Nunes* cleans this data of noise and irrelevant data. *Id.*, 3. Focusing *Smyth*'s learning models only on relevant hacker data from forums and removing predetermined irrelevant data (e.g., discussions on drugs, weapons, etc.) would improve classification and reduce the computational load of *Smyth*'s leaning models. A POSITA would have understood this benefit and would have been motivated to improve *Smyth*'s system with this modification. For example, a POSITA would have recognized that limiting the data from data sources 110 to be input into data processor 120 for indexing and subsequent processing via predictive models 531-534 would have significantly decreased the computational load on data processor 120. A POSITA would have recognized this decrease in computational load would allow *Smyth*'s system to more efficiently determine the likelihood of a vulnerability being exploited in two ways. First, by decreasing computational load, *Smyth*'s data processor 120 would be able to analyze more data in a similar amount of time, allowing for an enhanced training process by increasing data exposure to

the predictive models 531-534. Second, an initial filter on the data sources would narrow the pool to more relevant data being directed to predictive models 531-534. Initially narrowing the data would improve training, testing, and implementation as the models would analyze data relevant to vulnerabilities and exploits and would therefore learn to better predict future exploitations.

150. In my opinion, a POSITA would have appreciated a reasonable expectation of success modifying *Smyth* to implement *Nunes*' parser with an integrated classifier. *Smyth* describes several known techniques to collect and process data for its predictive models. *Smyth*, 5:49-52 (collecting data by, e.g., "web page scraping" or "other known techniques"). *Nunes*' parser was a well-known technique to collect online forum data, including data from the dark web and would have integrated seamlessly into *Smyth*'s system to collect forum data 240 for *Smyth*'s predictive models. In this way, the *Smyth-Phillips* system is compatible with systems that scrape data from the dark web, such as *Nunes*, and would have utilized *Nunes*' parser to filter out the noise. And systems for reducing noise when scraping forum data were well known in the art. *See, e.g., Fu* (Ex. 1022), 2 (discussing "Dark Web forum crawling" that use "URL ordering features and techniques [to] ensure that only the desired pages are collected, and in the most efficient manner); Victor Benjamin et al., Exploring Threats and Vulnerabilities in Hacker Web: Forums, IRC, and Carding Shops, (Ex. 1023), Abstract (describing "machine learning

methodology with information retrieval techniques...to distill potential threats from the entirety of collected hacker contents”). Further, *Smyth* describes example processing performed on this collected data. *Smyth*, 6:22-55 (describing indexing and transforming data to “facilitate[] processing by the predictive models”). And *Smyth* notes embodiments may “employ other techniques not specifically described here...and provide [the data] in a manner suitable for the prediction engine 160.” *Id.* In embodiments, *Smyth* can vary processing based on the data source. *Id.*, 9:6-17. Thus, it would have been well within the skill of a POSITA to implement *Nunes*’ classifier integrated in its parser to process the collected data and filter out noise and predetermined irrelevant data. Modifying *Smyth-Phillips* in this way would have been accomplished without undue experimentation.

3. *Claim 16: The computing device of claim 14, wherein the predictive model includes a Random Forest (RF) method including multiple decision tree predictors applied in combination to classify the software vulnerability which is used to estimate the likelihood of exploitation.*

151. *Smyth-Phillips* teaches the computing device of claim 14. *Supra*, Ground 1, Claim 14. *Smyth* further teaches its prediction engine 160 (“predictive model”) comprises an ensemble of different machine learning algorithms, including “tree-based model” 531. *Smyth*, 9:20-45, Fig. 5. “The predictive engine 160 takes [] processed current vulnerability information 125 and uses the prediction ensemble 155...to generate a prediction 170 of the likelihood the current vulnerability may be

employed[.]” *Id.*, 4:24-29, 10:61-64, 11:18-25. While *Smyth* discloses using tree-based models to classify the software vulnerability, *Smyth* does not expressly state it uses a Random Forest tree-based model. *Nunes* discloses “well-known classification techniques” including “random forest (RF)[.]” *Nunes*, 2.

152. It would have been obvious to apply *Nunes*’ well-known RF classification technique as the specific tree-based model in *Smyth* because this is a simple substitution of one known element for another to obtain predictable results. In my opinion, a POSITA would have understood that applying an RF model in *Smyth*’s predictive engine would improve classification. Namely, RF models aggregate predictions with high accuracy, are robust to noise, and are non-parametric such that these models may be applied to a variety of datasets, making them suitable for *Smyth*’s invention. *Dam* (Ex. 1027), (choosing “Random Forests (RF)” for vulnerability prediction, noting RF “is one of the most effective classifiers for vulnerability prediction”). A POSITA would have appreciated a reasonable expectation of success implementing an RF model as one of the predictive algorithms in *Smyth* because this was a well-known classification technique and sub-species of the tree-based models already used in *Smyth*. See also, e.g., *Kuperman* (Ex. 1008), 13:9-14 (discussing training a model using “random forest trees”). As such, in my opinion a POSITA would have appreciated success in modifying *Smyth* to use an RF model, as taught by *Nunes*.

## VIII. OPINION – GROUND 4

### A. Obvious to Combine *Smyth*, *Phillips*, and *Kuperman*

#### 1. Overview of *Kuperman*

153. *Kuperman* teaches systems for detecting malicious activity and protecting web applications. *Kuperman*, Abstract, 1:18-21 3:30-42. *Kuperman* includes a model that “reduces the likelihood of successful exploit 113 of web application 120 vulnerability 121 due to false negatives and reduces the likelihood of false positives 115.” *Id.*, 13:47-52. To this end, *Kuperman* teaches an optimization process for this model that evaluates a desired performance level based on performance metrics including “False Positive %” and “False Negative %” such that a test and training error rate determines whether the model exceeds a predefined “desired performance level” and can be marked “live” for use. *Kuperman*, 19:61-20:22, Fig. 6.

154. Because *Kuperman* and the '897 Patent are both directed to a system for assessing cybersecurity risks for vulnerabilities, *Kuperman* and the '897 Patent are in the same field of endeavor. *Compare Kuperman*, Abstract, 13:47-52 (“[t]he model 205 itself, which is based on known malicious request attributes and known non-malicious requests attributes, reduces the likelihood of successful exploit 113 of web application 120 vulnerability 121 due to false negatives and reduces the likelihood of false positives 115”), Fig. 1, with '897 Patent, 1:24-28 (describing “assessing the likelihood of exploitation of software vulnerabilities”), 4:1-27, 8:41-

45, 13:47-50. Accordingly, a POSITA would have recognized that *Kuperman* is analogous art to the '897 Patent.

155. The '897 Patent describes common problems in the art for assessing cybersecurity risks for vulnerabilities, including attempting to develop machine learning models that “predict[] exploits in the wild with [a] True Positive Rate (TPR) of 90% while maintaining the False Positive Rate (FPR) less than 15%[.]” '897 Patent, 3:54-58; *see also, id.*, 14:42-15:10, 20:8-11. *Kuperman* similarly addresses the problem of “determin[ing] whether desired classification performance has reached at least a minimum desired performance level for the model, along with the evaluation accuracy metrics such as...Accuracy %, False Positive % and False Negative %[.]” *Kuperman*, 19:61-20:22; *see also, id.*, Fig. 6. As such, *Kuperman* is reasonably pertinent to a problem described by the '897 Patent and is analogous art.

2. *Claim 6: The method of claim 1, further comprising testing the classification models by applying additional training data and one or more algorithms and evaluation metrics to optimize the classification models until the classification models compute the likelihood of exploitation according to a predefined error rate.*

156. *Smyth* teaches a process for testing its classification models. *Smyth*, 10:5-41. Specifically, combined data set 123 is split into training 512 and testing sets 514. *Id.*, 10:13-15. If training set 512 is imbalanced, *Smyth* applies additional training data to rebalance the set “using techniques such as resampling, over sampling or under sampling, synthetic minority over-sampling, and random minority

over-sampling.” *Id.*, 10:15-27. Rebalanced training data 512 is then used to train the predictive models. *Id.* Testing sets 514 are then used to test the retrained models. *Id.*, 10:28-41. To optimize these models, *Smyth* applies a “modeling scheme 130” that employs various algorithms that overlap with the optimization algorithms disclosed in the ’897 Patent’s specification. *Compare Smyth*, 10:55-59 (describing “majority voting” and “boosting”) with ’897 Patent, Fig. 13 (205 depicting testing and improving classification models using algorithms like “Boosting” and “Majority voting”).

157. *Smyth*’s modeling scheme 130 may also include “adjustments or biases for consideration when retraining” to “tailor the predictions 170 in a manner that provides a more useful prediction 170.” *Smyth*, 9:50-54. Like the claimed invention, *Smyth* adjusts or biases data based on evaluation metrics to optimize the model “so that, for example, a higher true positive rate may be achieved while reducing the rate of false positives.” *Compare Smyth*, 9:60-64 with ’897 Patent, 15:11-30 (Table 4 listing evaluation metrics such as “TPR” true positive rate and “FPR” false positive rate).

158. While *Smyth* discloses using additional training data, algorithms, and evaluation metrics to optimize the classification models to verify their performance, *Smyth* does not expressly describe optimizing the classification models “according to a predefined error rate.” *Kuperman* teaches the concept of model optimization,

for example, stating that “[o]nce an initial model is created, the model generator 209 may perform one or more testing iterations on a model to verify and/or **optimize** the model.” *Kuperman*, 19:21-23. Further, *Kuperman* teaches that “[d]uring training, the model generator 209 may generate a graph from results of training iterations to representing [sic] the cost function of the algorithm over the number of iterations trained.” *Id.*, 19:61-64. Additionally, *Kuperman* provides an example of such a graph in its Figure 6, stating that “[t]urning to FIG. 6, an example learning curve visulation graph 600 is shown that aids administrators in **visualizing the effectiveness** of a model at predicting malicious/non-malicious requests for a given number of requests in the training set.” *Id.*, 19:66-20:3. Moreover, *Kuperman* explains that “[t]he learning curve visualization 600 allows a system administrator to visually determine whether desired classification performance has **reached at least a minimum desired performance level** for the model.” *Id.*, 20:6-9. Note that *Kuperman*’s Figure 6 shows an increasing “Training Error” curve crossing a dashed “**Desired Performance Level**” horizontal straight line. *Id.*, Fig. 6. In fact, *Kuperman* explains that “[t]he **training error curve** represents the **error rate** as measured by evaluating the examples used to train.” *Id.*, 20:16-17. Finally, *Kuperman* mentions that “[o]nce **a model exceeds a desired performance level**, the model generator 209 and/or administrator may mark the model **as live** in the model database 311.” *Id.*, 20:20-22. For example, when the rising Training Error curve exceeds the Desired Performance

Level in Figure 6 of *Kuperman*, the model is marked as live and where the Desired Performance Level is a predefined training error rate. In other words, *Kuperman* teaches an optimization process for a malware classification model, where a minimum desired performance level is evaluated based on metrics such that a training error rate determines whether the model exceeds a predefined “desired performance level” and can be marked “live” for use. *Kuperman*, 19:61-20:22.

159. In my opinion, a POSITA would have been motivated to modify *Smyth* to optimize its models according to a predefined error rate, as taught by *Kuperman*, such that the models would be live after their performance exceeds a predefined desired performance level, for example, using the training error rate of *Kuperman*. *Smyth* already discloses testing its models “against the testing set 514 so that their performance can be verified.” *Smyth*, 10:44-47. Verifying performance according to a predefined error rate provides a simple mechanism to interpret results and provides a good indicator for model performance. A POSITA would have understood that *Smyth* describes techniques to acquire data evaluating the performance of the system. *Smyth*, 9:60-64. A POSITA would have understood that establishing an acceptable and/or desirable rate for these metrics would have improved *Smyth*’s system by providing a target error rate for the system to reach. Such an evaluation would have provided an objective standard by which to gauge the success of the system on the test set. Moreover, a POSITA would have understood that attaining an acceptable

and/or desirable performance level before marking the model live would be good engineering and operational practice.

160. In my opinion, a POSITA would have appreciated a reasonable expectation of success making this modification because *Smyth* already describes optimizing and verifying model performance and a predefined error rate was a widely used metric for model optimization and verification that would have been suitable for *Smyth*'s methods. Such a modification would have simply used the known technique of error rate threshold to improve *Smyth*'s system in the same way to provide an objective standard for its test set. *Kuperman*'s teachings would easily be incorporated into *Smyth*'s system without undue experimentation.

## **IX. OPINION – GROUND 5**

### **A. Obvious to Combine *Smyth*, *Phillips*, and *Marin***

#### *1. Overview of Marin*

161. *Marin* teaches analyzing websites on the darkweb for malicious hacking products and using machine learning to categorize and study those products. *Marin* (Ex. 1009), 1. *Marin*'s system includes clustering and grouping categories of online hacking tools, including "exploit kits[.]" *Marin*, 3.

162. The '897 Patent describes common problems in the art for assessing cybersecurity risks for vulnerabilities, including attempting to implement the exploit prediction model to "crawl websites on DW, both marketplaces and forums, to collect data relating to malicious hacking." '897 Patent, 8:25-29; see also, *id.*, 6:50-

59, 7:22-26, 8:30-9:17. *Marin* similarly addresses this problem by “scrap[ing]” “darkweb and deepweb” “sites and collection[ing] information about [malicious hacking] products...in order to understand their various categories and how they become specialized with respect to vendor and marketplace.” *Marin*, Abstract; *see also, id.*, 3 (determining prevalence of certain hacking techniques, such as “keyloggers”). As such, *Marin* is reasonably pertinent to a problem described by the ’897 Patent and is analogous art.

2. *Claim 7: The method of claim 1, further comprising vectorizing text features derived from the plurality of datasets using term frequency-inverse document frequency to create a vocabulary of associated words.*

163. *Smyth* teaches vectorizing text features derived from the plurality of datasets, including forum discussions, using “unstructured text analysis 430.” *Smyth*, 7:22-47 (“unstructured text analysis 430 is performed on certain portions of the vulnerability information so that features may be derived” such as “online discussions 240”), Fig. 4. *Smyth* further teaches in embodiments, “the unstructured text may be processed by generative statistical models[.]” *Id.*, 7:35-47 (listing example statistical models such as latent Dirichlet allocation (LDA), which uses vectorized text, and sentiment analysis). While *Smyth* suggests using certain text vectorization techniques (e.g., LDA, which uses vectorized text), *Smyth* does not expressly discuss using “term frequency-inverse document frequency” (TF-IDF) to vectorize text features.

164. TF-IDF was a known text vectorization technique taught by *Marin*, a paper co-authored by a co-inventor of the '897 Patent. *Marin*, 1. Indeed, TF-IDF was well-known in the cybersecurity space for text vectorization. *See, e.g., Gates* (Ex. 1024), 1:15-19, 16:7-17:10 (describing a system for “determining tiered responses to [cybersecurity] threats” which uses “TF.IDF process that can generate or modify term vectors formed in the data conditioning modules 210 and 220, with the TF.IDF being a metric that assigns numerical values to unstructured text”). *Marin* teaches analyzing websites on the darkweb for malicious hacking products and using machine learning to categorize and study those products. *Marin*, 1. *Marin* teaches vectorizing text features for its machine learning models, such as title, description, posting date, product name, using “standard term frequency – inverse document frequency (TF-IDF)[.]” *Id.*, 2; *see also id.*, 3 (Table II showing “performance of each TF-IDF vectorization”). In my opinion, it would have been obvious to vectorize text information from *Smyth-Phillips*’ online forums using TF-IDF in conjunction with LDA to create a vocabulary of associated words and uncover latent topics. A POSITA would have recognized that *Smyth*’s system would benefit from using TF-IDF in conjunction with TF-IDF. TF-IDF is a simple and computationally inexpensive vectorization technique that accurately determines the relevance of words within text. LDA provides benefits of uncovering latent topics and semantic structures. Accordingly, a POSITA would have been motivated to

implement TF-IDF and LDA together to enhance classification by correlating frequent keywords (TF-IDF) with latent topics and themes (LDA). Using LDA and TF-IDF together would have allowed the system to capture both high-level semantic structure (via LDA) and term-specific importance (via TF-IDF). A POSITA would have recognized that the combination would enrich the classification by leveraging the strengths of both models to improve interpretability and performance.

165. A POSITA would have had a reasonable expectation of success making this simple modification because *Smyth* already uses Latent Dirichlet Allocation (LDA), which relies on text vectorization, like that provided by TF-IDF and TF-IDF was widely known and commonly used technique for the same with understood benefits. A POSITA would have understood that *Smyth* teaches one example of a model for processing unstructured text, “such as latent Dirichlet allocation[.]” *Smyth*, 7:35-40. A POSITA would have understood that the addition of a text vectorization model to a system already performing natural language processing on vectorized text would have been simply applying the known technique of TF-IDF to *Smyth*’s system, which, in conjunction with LDA, would have improved the system for the reasons described above.

3. *Claim 13: The method of claim 1, further comprising vectorizing text features derived from textual content of the plurality of datasets using a predetermined natural language process (NPL).*

166. The '897 Patent describes exemplary NLP preprocessing techniques, including both “term frequency – inverse document frequency” (TF-IDF) and “Latent Dirichlet Allocation” (LDA). '897 Patent, Fig. 13. As stated above, TF-IDF was a known text vectorization method used to vectorize text features, as taught by *Marin*, *See*, Claim 7. Moreover, *Smyth* discloses using LDA. *Id.* For the same reasons as above, a POSITA would have been motivated and appreciated a reasonable likelihood of success in modifying *Smyth* with *Marin* to vectorize text features derived from textual content of the plurality of datasets using a predetermined natural language process (i.e., TF-IDF/LDA).

## **X. OPINION – GROUND 6**

### **A. Obvious to Combine *Smyth*, *Phillips*, and *Bozorgi***

#### *1. Overview of Bozorgi*

167. *Bozorgi* teaches a vulnerability prioritization scheme that predicts “whether and how soon a vulnerability is likely to be exploited.” *Bozorgi*, 1-2. *Bozorgi* further teaches extracting various vulnerability features for classifying and prioritizing the vulnerabilities. *Id.*, 4.

168. Because *Bozorgi* and the '897 Patent are both directed to systems for assessing cybersecurity risks for vulnerabilities, *Bozorgi* and the '897 Patent are in the same field of endeavor. *Compare Bozorgi*, Abstract (“we show how to train classifiers that predict whether and how soon a vulnerability is likely to be exploited”), 2, 5-6, with '897 Patent 1:24-28 (describing “assessing the likelihood

of exploitation of software vulnerabilities”), 4:1-27, 8:41-45, 13:47-50. Accordingly, in my opinion, a POSITA would have recognized that *Bozorgi* is analogous art to the ’897 Patent.

169. The ’897 Patent describes common problems in the art for assessing cybersecurity risks for vulnerabilities, including prioritizing vulnerabilities for purposes of predicting exploitation. ’897 *Patent*, 7:13-17 (describing the generation of estimation outputs, and presenting them by “ranking vulnerabilities”), Fig. 13. *Bozorgi* similarly addresses “a key question” of “which vulnerabilities to prioritize.” *Bozorgi*, Abstract; *see also, id.*, 6-8 (“the models provide real-valued scores that practitioners can use to prioritize vulnerabilities”). As such, *Bozorgi* is reasonably pertinent to a problem described by the ’897 Patent and is analogous art.

2. *Claim 8: The method of claim 1, further comprising: sorting vulnerabilities associated with the plurality of datasets according to time; training the classification model using the training data, the training data defining a first subset of the plurality of datasets associated with a predetermined period of time; and testing the classification model using a second subset of the plurality of datasets associated with the predetermined period of time.*

170. *Smyth* teaches training its classification model with training data from the plurality of datasets in combined data set 123. *Supra*, Claim 1(b), 1(d). *Smyth* splits combined data set 123 into training and testing sets. *Supra*, Claim 6. “One such split of the combined data set 123 may separate information for which impacts are known from information for which impacts are not quantified.” *Smyth*, 10:9-12. For

example, “testing set 514 may be formed of information for which the impacts of certain vulnerabilities are known, so an actual threat level is derived from such impacts.” *Id.*, 10:29-32. *Smyth* therefore teaches sorting and dividing vulnerability data for training and testing based on whether the impact is known or not. *Smyth*, however, does not sort and divide vulnerability data for training and testing based on time.

171. *Bozorgi* teaches a vulnerability prioritization scheme that predicts “whether and how soon a vulnerability is likely to be exploited.” *Bozorgi*, 1. Like *Smyth*, *Bozorgi* trains a classifier “on known vulnerabilities to make predictions about new ones.” *Id.*, 5. To accomplish this, *Bozorgi* trains and initializes its classifier with all vulnerabilities “reported between time 0 and time  $t$  that have known outcomes” and then further tests this classifier with “vulnerabilities that appear in the next time interval” between  $t+1$  and  $t+T$  based on new vulnerabilities arriving between  $t+1$  and  $t+T$ , and then “update[s] the classifier with known vulnerability outcomes.” *Id.*, 6. *Bozorgi*, therefore, trains and initializes its classifier with a first subset of data (i.e., of vulnerabilities arriving between time 0 and  $t$ ) “associated with a predetermined period of time” (i.e., between time 0 and  $t+T$ ) and further tests its classifier with a second subset of data (i.e. of vulnerabilities arriving between time  $t+1$  and  $t+T$ ) within that “associated predetermined period of time.”

172. In my opinion, a POSITA would have understood that the time between 0 and  $t$  plus the “next time interval” (i.e., time  $t+1$  to  $t+T$ ) would constitute a “predetermined time period of time” (i.e., time 0 to  $t+T$ ) as claimed. A POSITA would have thus understood that *Bozorgi* sorts its vulnerability data based on time to train and test data from different time intervals, that is, from 0 to  $t$  and from  $t+1$  to  $t+T$  correspondingly, where  $T$  can be, for example, a week or a month.

173. In my opinion, a POSITA would have been motivated to modify *Smyth* pursuant to *Bozorgi* to sort (chronologically), train, and test *Smyth*’s vulnerability data according to time because this would be a natural splitting of the data, based on time ordering (e.g., past and future) would better reflect real-world applications and would preserve temporal dependencies.

174. Further, a POSITA would have appreciated a reasonable expectation of success making this minor modification. *Smyth* teaches its data sources 110 can additionally include the “date a vulnerability was released” and has a similar process for splitting training and test data. *Smyth*, 7:9-21, 10:5-12. Thus, it would have required simple programming changes to sort vulnerability data by dates/times and split the training and test data according to the same. Since *Smyth* is 1) already equipped to acquire the information necessary to divide the data in a time-based manner, and 2) already divides the data based on certain factors (e.g., whether it was

known or not), a POSITA would have appreciated success in modifying *Smyth* to divide the data in a time-based manner, as taught by *Bozorgi*.

175. It would also have been obvious to modify *Smyth* pursuant to *Bozorgi* to sort vulnerabilities by time and to train and test with different subsets of datasets over the same time interval to evaluate the classifier's performance. As discussed above, *Smyth* teaches dividing a combined dataset into separate training and test sets. *Smyth*, Fig. 5. That combined dataset is made up of data from a predetermined period of time. *Smyth*, 4:10-14 (describing combined data set of "historical vulnerability information obtained from data sources 110"). *Smyth* trains with a portion of data from that predetermined period of time and tests with a different portion of data from that predetermined period. *Smyth*, 10:12-27. It would have been obvious to train with a specific portion of data from those datasets, such as NVD data, and test with a different portion, such as online discussions 240, to determine how reliably *Smyth*'s classifier detects vulnerabilities mentioned in the test dataset. This modification would provide beneficial metrics to assess classifier performance. And a POSITA would have appreciated success making this simple modification for the same reasons detailed above.

3. *Claim 12: The method of claim 1, further comprising: accessing one or more features indicative of temporal connections between at least two of: a time associated with discussion of a vulnerability by users posting hacking-related content at a web forum prior to disclosure to a public vulnerability database; a time associated with disclosure of the vulnerability to the public*

*vulnerability database; and a time associated with exploitation of the vulnerability as obtained through exploits data associated with real-world exploitation of the vulnerability.*

176. *Smyth-Phillips* teaches deriving features from various data sources 100, including features from web forums containing hacking content. *Supra*, Claim 1(c). For example, *Smyth* teaches using “unstructured text analysis 430” on portions of the vulnerability information such as online discussions so that features may be derived for *Smyth*’s predictive models. *Id.*; *Smyth*, 7:22-35. *Phillips* also derives features from metadata of hacker forums. *Supra*, Claim 1(c); *Phillips*, 1-2 (“Methodology”). This includes, or would have been obvious to include, “timestamps” for forum posts. *Phillips*, 1. For example, *Phillips* cites to “[e]xisting work” that establishes a “hierarchy from the metadata of email communications alone[,]” which includes “the post’s author and timestamp.” *Id.* *Phillips* further emphasizes that it “set[s] out to investigate whether similar techniques could be applied within Dark Web Forums[.]” *Id.* In my opinion, a POSITA would have understood that, by using techniques similar to the existing art, *Phillips*’ system would similarly collect “the [Dark Web Forum’s] author and timestamp.” Additionally, or alternatively, a POSITA would have found it obvious to include the timestamp that *Phillips* references as it sought out to mimic similar communication-based systems that collected timestamps of posts for distinguishing the data. *Smyth-Phillips* therefore teaches or renders obvious accessing features indicative of “a time

associated with discussion of a vulnerability by users posting hacking-related content at a web forum prior to disclosure to a public vulnerability database[.]”

177. While *Smyth* teaches deriving features from known vulnerability data, including CVE entries, and notes CVE entries include “the date the entry was created[.]” *Smyth* does not expressly state it derives CVE date information for its classification process. *Bozorgi* teaches extracting CVE time stamps “as features for classification[.]” *Bozorgi*, 3. Table 2 provides a list of features extracted for each vulnerability *Bozorgi* analyzes, including “CVE Dates[.]” *Id.*, 4 (describing “integer-valued features” that “encode useful information, such as the date when a vulnerability was first disclosed”). Since CVE is a “well-known, online source of vulnerability data,” and *Bozorgi* expressly teaches extracting the “date when a vulnerability was first disclosed” such as “CVE Dates[.]” *Bozorgi* teaches accessing features indicative of a “time associated with disclosure of the vulnerability to the public vulnerability database[.]” A POSITA would have been motivated and would have appreciated a reasonable expectation of success modifying *Smyth-Phillips* pursuant to *Bozorgi* to access temporal features for the same benefits discussed above in claim 8 regarding improving classification with simple programming changes. *Supra*, Claim 8.

## **XI. CONCLUSION**

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

Dated: 

By: 2025-07-25

# CURRICULUM VITAE

## Nicholas Bambos

R. Weiland Professor in the School of Engineering  
Prof. of Electrical Engineering  
Prof. of Management Science & Engineering

Stanford University

April 2025

Nick Bambos is a R. Weiland Professor in the School of Engineering at Stanford University, having a joint appointment in the Department of Electrical Engineering and the Department of Management Science & Engineering. He has also been the Fortinet Founders Department Chairman of the Management Science & Engineering Department (2016-20). He heads the Computer Systems & Networks Performance Engineering Lab at Stanford. This is comprised of PhD students and industry visitors engaged in research projects, which require solid understanding of the physical principles of computing/networking and hardware/software architectures (e.g., digital design, file systems), but largely focus on system architectures and high-performance engineering of computer systems and networks. His research contributions span the areas of wireless & wireline networking, the Internet, cloud computing and data centers, high-speed networking and packet switching, multimedia streaming, information service engineering, computer systems security, etc. His methodological interests and contributions span the areas of stochastic modeling and control, online task scheduling, routing and distributed processing, network control, etc.

He has graduated over 40 Ph.D. students (including two postdocs), who have moved on to key positions in academia (Stanford, CalTech, Michigan, GaTech, NYU, USC, Columbia, etc.), the information technology industry (Cisco, IBM, Broadcom, Qualcomm, Nokia, MITRE, ST Micro, Intel, Samsung, Google, Facebook, Twitter, etc.), technology startups and the finance tech industry. From 1999 to 2005 he served as Director of the Stanford Networking Research Center project, a partnership between Stanford and information technology industries, involving tens of corporate members, faculty and doctoral students.

He received his Ph.D. in Electrical Engineering and Computer Sciences (EECS) from the University of California at Berkeley (1989), as well as the M.S. in EECS (1987) and the M.A. in Mathematics (1989) from the same University. He graduated in Electrical Engineering from the National Technical University of Athens-Greece (1984) with first class honors. Before joining Stanford as an Associate Professor in 1996, he served as Assistant (1990-95) and tenured Associate Professor (1995-96) in the Electrical Engineering Department of the University of California at Los Angeles (UCLA).

Nick Bambos has received several best research paper awards, has held the Cisco Systems Faculty Development Chair (1999-2003) in computer networking at Stanford, and has won the IBM Faculty Award (2002) for high-impact research in performance engineering of computer systems and networks, as well as the Griffin Award (1997). He has been the David Morgenthau Faculty Scholar (1996-99) at Stanford, and has received the National Young Investigator Award (1992) from the National Science Foundation (NSF) for research in computer networks and distributed computing architectures, as well as the NSF Research Initiation Award (1990) for studies in performance modeling of computer systems. He has also been a U.C. Regents Fellow, a David Gale Fellow, and an Earl Anthony Fellow at U.C. Berkeley.

He has served as editor of various research journals, on international technical and scientific committees, and on review panels for networking and computing technologies. He has also served on the technical boards of start-up companies, as a consultant for high-technology development and management, and as an expert witness in high-profile patent litigation cases involving networking, computing and information technologies.

## EDUCATION

- **Ph.D. in Electrical Engineering and Computer Sciences**, University of California at Berkeley, 1989.
- **M.A. in Mathematics**, University of California at Berkeley, 1989.
- **M.S. in Electrical Engineering and Computer Sciences**, University of California at Berkeley, 1987.
- **Diploma in Electrical Engineering - First Class Honors**, National Technical University of Athens-Greece, 1984.

## PROFESSIONAL EXPERIENCE

- **R. Weiland Professor** in the School of Engineering, Stanford University, 2016 – present.
- **Fortinet Founders Department Chairman**, Department of Management Science & Engineering, Stanford University, 2016 – 2020.
- **Professor**, Department of Management Science & Engineering, and Department of Electrical Engineering, Stanford University, 2003 – present.
- **Head**, Computer Systems & Networks Performance Engineering Research Lab (NetLab), Stanford University, 1998 – present.
- **Director**, Stanford Networking Research Center (SNRC) Project Initiative, Stanford University, 1999 – 2005.
- **Associate Professor**, Department of Management Science & Engineering, and Department of Electrical Engineering, Stanford University, 1999 – 2003.
- **Associate Professor**, Department of Engineering-Economic Systems & Operations Research, and (courtesy appointment at) Department of Electrical Engineering, Stanford University, 1996 – 1999.
- **Associate Professor** (tenured), Department of Electrical Engineering, University of California at Los Angeles, 1995 – 1996.
- **Assistant Professor**, Department of Electrical Engineering, University of California at Los Angeles, 1989 – 1995.

## HONORS AND AWARDS

- **Best Paper Award in Green Communications and Networking** for the paper “Power-Managed Data Centers for Sustainable Computing” (E. Zeger, N. Bambos, M. Pilanci). Proceedings of the 2024 IEEE International Conference in Communications (ICC 2024), pp. 2833-2839, Denver, CO, June 2024.
- **2022 Best Paper Award for Telecommunications and Network Analytics** from INFORMS on for the paper “Robust Power Management via Learning and Game Design” (Z. Zhou, P. Mertikopoulos, A. Moustakas, N. Bambos, P. Glynn), INFORMS Annual Meeting, IN, October 2022.
- **Outstanding Paper Award at Healthcom** for the paper “Physiological Waveform Imputation of Missing Data using Convolutional Autoencoders” (with D. Miller, A. Ward, D. Scheinker, Andrew Shin). *Proceedings of the 2018 IEEE 20<sup>th</sup> International Conference on e-Health Networking, Applications and Services (Healthcom’18), Ostrava, Czech Republic, September 2018.*
- **Best Paper Award in Green Communication and Computing** by the TAOS (Transmission, Access & Optical Systems) Technical Committee for the paper: *Dynamic Resource Management in Virtualized Data Centers with Bursty Traffic* (with M. Valdez-Vivas, J. Apostolopoulos) presented in the IEEE International Communications Conference (ICC’14) in Sydney, Australia, June 2014.
- **“Best of Globecom” Paper Award** (top 2%) for the paper: *Dynamic Resource Management in Cloud Computing with Frictions and Congestion “Weather”* (with M. Valdez-Vivas, J. Apostolopoulos) presented in IEEE Global Communications Conference (Globecom’14) in Austin, TX, December 2014.
- **2013-14 Eugene L. Grant Teaching Award**, Stanford University.
- **Best Journal Paper Award for 2011** from the IEEE Multimedia Communications Technical Committee for the paper: *Channel, Deadline, and Distortion Aware Scheduling of Video Streams over Wireless Links* (with A. Dua, C. Chan, J. Apostolopoulos), IEEE Trans. on Wireless Communications, 9(3):1001-1011, 2010.
- **HP Faculty Award**, Hewlett-Packard Laboratories, Palo Alto, 2006
- **Cisco Systems Faculty Scholar**, School of Engineering, Stanford University, 1999 - 2003.
- **IBM Faculty Award**, IBM Corporation Research Labs, 2002.
- **Dana Adams Griffin Award**, School of Engineering, Stanford University, 1997.
- **David Morgenthaller Faculty Scholar**, School of Engineering, Stanford University, 1996 - 1999.
- **National Young Investigator Award**, National Science Foundation, 1992 - 1997.
- **Research Initiation Award**, National Science Foundation, 1990 - 1992.
- **EECS Departmental Scholar**, University of California at Berkeley, 1988 - 1989.
- **U.C. Regents Fellow**, University of California at Berkeley, 1986 - 1988.
- **Earl C. Anthony Fellow**, University of California at Berkeley, 1985 - 1986.
- **David and Sylvia Gale Fellow**, University of California at Berkeley, 1984 - 1985.

- **First Class Honors** for top graduating students in 1984 (among 150 students), Electrical Engineering Department, National Technical University of Athens-Greece, 1984.
- **Academic Excellence Awards** (annual), Hellenic State Scholarships Foundation, Greece, 1979 - 1984.
- **Academic Prizes** (annual), Hellenic Engineering Society, Greece, 1979 - 1984.
- **First Class Honors** for top students admitted to the University in 1979 (among tens of thousands of students participating in a national-level entrance exam), National Technical University of Athens-Greece, 1979.

#### **DISTINGUISHED INVITED LECTURES AND TALKS** (selected; additional talks on pg. 19)

- *Keynote Talk – Geometry of Resource Scheduling in Service Systems*. The 15<sup>th</sup> International Conference on Service Systems and Service Management (ICSSSM), Hangzhou, China, July 21, 2018.
- *Plenary Talk -- Transmitter Power Control in Wireless Networking*. The 2018 International Conference on Computing, Networking and Communications (ICNC), Maui, Hawaii, March 6, 2018.
- *Keynote Address – Power Control and Bandwidth Management in Wireless Networking*. IEEE Wireless On-Demand Network Systems and Services Conference (WONS), Jackson Hole, WY, February 2017.
- *Plenary Address – Digital Living 2030*, Shandong Chinese Academy of Sciences, Jinan, China, January 2017.
- *Keynote Address – A Risk Management View to Information Security*, IEEE Conference on Decision and Game Theory for Security, Berlin, Germany, Nov. 2010.
- *Keynote Address – Wireless Computing: From Infrastructure to Services*, IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Athens, September 2007.
- *Invited Distinguished Lecture - Three Control Problems in Packet Switching*, Workshop on Mathematics of Networks, Ecole Normale Supérieure, Paris, France, June 2007.
- *Invited Distinguished Talk - Throughput and Power Issues in Packet Switches*, Center for Ultra-Broadband Information Networks, University of Melbourne, Australia, August 2007.
- *Keynote Talk - Opportunities and Challenges in Wireless Computing and Multimedia Networking*, Network Ventures Conference, Redwood Shores, CA, March 2005.
- *IBM Distinguished Talk - Resource Management and Load Balancing in Flexible Computing/Networking Service Platforms*, IBM Research Labs, Almaden, CA, June 2005
- *Distinguished Speaker - Power Control in Wireless Networks*, Wireless Industry Forum, Electrical Engineering Department, University of Texas at Austin, April 2004.
- *Plenary Talk - Power Control in Wireless Networks* and member of Panel of Experts on Network Technology Trends, IEEE Workshop on Local and Metropolitan Area Networks (LANMAN), pp. 69-74, CA, April 2004
- *Invited Lecture - Geometry of Packet Scheduling in Communication Switches*, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Nov. 2004.

- *Plenary Talk – Next Generation Wireless System Architectures* and member of Panel of Experts on Wireless Technologies in the IEEE International Conference in Broadband Networks (BroadNets), San Jose, CA, October 2004.
- *Plenary Talk – Next Generation Wireless Data Networks: Control Plane Principles and Architectures.* Strategic Architecture Council of Sun Microsystems, Nov. 2003
- *International Expert* invited to serve on the International Evaluation Committee for network engineering research at INRIA (the French/European National Institute of Informatics and Automation of 900 scientists working in six campuses), 2003-2004.

## UNIVERSITY AND PROFESSIONAL SERVICE

- *Academic Senate* of Stanford University (Member), 2005 – 2010
- *Committee on Research*, Stanford University (Member), 2009 – 2012.
- *Member of the Steering Board*, IEEE Conference on Decision and Game Theory for Security (GameSec), 2010-2017.
- *Member of the Scientific Board*, LINC – Laboratory on Information, Networking and Communication Sciences, Paris, France (founded by INRIA, Institut Telecom, France Telecom and Alcatel-Lucent), 2010-2023.
- *Member the Editorial Boards* of the Wireless Networks Journal (2001-07), Computer Networks Journal (2000-02), Queueing Systems Theory and Applications (1999-2007).
- *Member of the Technical Program Committees* of the various networking/computing conferences (IEEE Infocom 2001-03, IEEE Globecom 2001-03, IEEE Mobicom 2002, IEEE VTC 2003, IEEE MobiWac 2005, SpasWin 2006, BroadNets 2006-07, Wireless Broadband World Forum 2006-07, GameComm 2008)
- *Technical Reviewer* for several research journals in computing and networking, as well as in queueing systems and applied probability (IEEE/ACM Transactions on Networking, IEEE Tran. on Vehicular Technology, IEEE Tran. on Automatic Control, Advances in Applied Probability, Journal of Applied Probability, Annals of Applied Probability, Operations Research, Journal of the ACM, Math of OR, IEEE Trans. on Communications)
- *Member of various Technical Panels* of the National Science Foundation (NSF) and other research and technology development agencies (NSF Information Technology Research Panel, 2002; INRIA Project Evaluation Panel, 2003; NSF CAREER Panel, 2005; NSF Foundations Panel, 2006; NSF NeTS CAREER Panel, 2006; NSF FIND Panel, 2007; NSF CDI Panel, 2008).

## RESEARCH PUBLICATIONS, REPORTS, TALKS, ETC.

---

### BOOK CHAPTERS:

1. **Personalized Monitors for Real-Time Detection of Physiological States** (L. Chow, N. Bambos, A. Gilman, A. Chander). *Wearable Technologies: Concepts, Methodologies, Tools and Applications*, chapter 42, pp. 931-951, IGI Global, 2018. [Included in this 3-volume book, but previously also published in: *The International Journal of E-Health and Medical Communications*, 5(4):1-19, 2014.]
2. **Power Control in Wireless Networking: Basic Principles and Core Algorithms** (with S. Gitzenis). *Signal Processing for Mobile Communications Handbook* (M. Ibnkahla, Editor), CRC Press, 2004.
3. **Transmitter Power Control in Wireless Computing** (with S. Gitzenis). *Handbook of Mobile Computing* (I. Mahgoub, Editor), CRC Press, 2004.
4. **Adaptive Batch Scheduling for Packet Switching with Delays** (with K. Ross). *High-Performance Packet Switching Architectures* (Editors: I. Elhanany, M. Hamdi), Springer-Verlag, pp. 66-80, 2006.
5. **Geometry of Packet Scheduling: Maximal Throughput Cone Scheduling Algorithms** (with K. Ross). *High-Performance Packet Switching Architectures* (Editors: I. Elhanany, M. Hamdi), Springer-Verlag, pp. 82-100, 2006.
6. **Autonomic Admission Control for Networked Information Servers** (with G. Paleologo). *Telecommunications Network Design and Management, Operations Research and Computer Science Interfaces*, Kluwer, vol. 23, pp. 227-244, 2003

---

### JOURNAL PAPERS:

7. **Gameseeker: Online Learning for Admission Control of Networked Open Multiagent Systems** (I. Bistritz, N. Bambos). *IEEE Transactions on Automatic Control*, 69(11):7894-7709, Nov. 2024.
8. **Power is Knowledge: Distributed and Throughput Optimal Power Control in Wireless Networks** (I. Bistritz, N. Bambos). *IEEE/ACM Transactions in Networking* 12(6):4722-4734, Dec. 2024.
9. **Association of Common Pediatric Surgeries with New Onset Chronic Pain in Patients 0-21 Years of Age in the United States.** (I. Thapa, E. De Souza, An. Ward, N. Bambos, T.A. Anderson). *The Journal of Pain*, 24(2):320-331, 2023.
10. **WAVES – The Lucile Packard Children’s Hospital Pediatric Physiological Waveforms Dataset** (D. Miller, G. Dhillon, N. Bambos, A. Shin, D. Scheinker). *Nature Scientific Data*, 10(124):1-7, March 2023.
11. **Personalized and Energy-Efficient Health Monitoring: A Reinforcement Learning Approach** (B. Eden, I. Bistritz, N. Bambos, I. Ben-Gal, E. Khmelnitsky). *IEEE Control Systems Letters*, 7:955-960, 2023 (and also presented at the 2023 American Control Conference, San Diego, CA, Dec. 2023).
12. **Smart Greedy Distributed Energy Allocation: a Random Games Approach** (I. Bistritz, A. Ward, Z. Zhou, N. Bambos). *IEEE Transactions on Automatic Control* 67(5):2208-2220, May 2022 (published

online May 2021).

13. **No Weight-Regret Learning in Adversarial Bandits with Delays** (I. Bistritz, Z. Zhou, X. Chen, N. Bambos). *Journal of Machine Learning Research (JMLR)*, 23(139):1-43, March 2022.
14. **Distributed Stochastic Optimization with Large Delays** (Z. Zhou, P. Mertikopoulos, N. Bambos, P. Glynn and Y. Ye). *Mathematics of Operations Research*, 47(3):2082-2111, Aug. 2022 (published online Dec. 2021).
15. **Wireless Body-Area Network Control Policies for Energy-Efficient Health Monitoring** (Y.B. David, T. Geller, I. Bistritz, I. Ben-Gal, N. Bambos, E. Khmelnitsky). *Sensors Journal*, 21/4245, pp. 1-21, June 2021.
16. **Prediction of Prolonged Opioid Use After Surgery in Adolescents: Insights From Machine Learning** (A. Ward, T. Jani, E. De Souza, D. Scheinker, N. Bambos, A. Anderson). *Journal of Anesthesia & Analgesia*, 133(2):304-313, 2021.
17. **Consensus-Based Stochastic Control for Model-Free Battery Cell Balancing** (I. Bistritz, N. Bambos). *IEEE Transactions on Control of Network Systems*, 8(3):1139-1150, Sept. 2021.
18. **Robust Power Management via Learning and Game Design** (Z. Zhou, P. Mertikopoulos, A. Moustakas, N. Bambos, P. Glynn). *Operations Research*, 69(1):331-345, 2021.
19. **One for All and All for One: Distributed Learning of Fair Allocations with Multi-Player Bandits** (I. Bistritz, T.Z. Baharav, A. Leshem, N. Bambos). *IEEE Transactions on Selected Areas in Information Theory*, 2(2):584-598, June 2021.
20. **Incidence of and Factors Associated with Prolonged and Persistent Postoperative Opioid Use in Children 0 to 18 Years of Age** (A. Ward, E. De Souza, D. Miller, E. Wang, E. Sun, N. Bambos, T. Anderson). *Journal of Anesthesia & Analgesia* 131(4):1237-1248, Oct. 2020.
21. **On the Convergence of Mirror Descent Beyond Stochastic Convex Programming** (with Z. Zhou, P. Mertikopoulos, S. Boyd, P. Glynn). *SIAM Journal of Optimization*, 30(1):687-716, 2020.
22. **Clustering Users by Their Mobility Behavioral Patterns** (with I. Ben-Gal, S. Weinstock, G. Singer). *ACM Transactions on Knowledge Discovery from Data*, 13(4)45:1-28, 2019.
23. **Deterministic and Stochastic Wireless Network Games: Equilibrium, Dynamics and Price of Anarchy** (with Z. Zhou, P. Glynn). *Operations Research*, 66(6):1498-1516, 2018.
24. **Infinite Time Horizon Maximum Causal Entropy Inverse Reinforcement Learning** (with Z. Zhou, M. Bloem). *IEEE Transactions on Automatic Control*, 63(9): 2787-2802, September 2018.
25. **Stochastic Models of Ground Delay Program Implementation for Prediction, Simulation and Insight** (with M. Bloem). *Journal of Aerospace Operations*, 5(1-2):85-117 (2017), April 2018.
26. **Video Streaming Schemes for Industrial IoT** (with H. Kanzaki, K. Schubert). *Journal of Reliable Intelligent Environments*, 3(4):233-241, Dec. 2017.
27. **Improving Predictions of Pediatric Surgical Durations with Supervised Learning** (with N. Master, Z. Zhou, D. Miller, D. Scheinker, P. Glynn). *Int. Journal of Data Science and Analytics*, 4:35-52, May 2017.
28. **Service Rate Control of Tandem Queues with Power Constraints** (with L. Xia, D. Miller, Z. Zhou). *IEEE Transactions on Automatic Control*, March 2017.

29. **Power Control for Packet Streaming with Head-of-Line Deadlines** (with N. Master). Journal of Performance Evaluation, 106:1-18, 2016.
30. **Power Optimization in Random Wireless Networks** (with A. Moustakas, P. Mertikopoulos). IEEE Transactions on Information Theory, 62(9):5030-5058, September 2016.
31. **Adaptive Prefetching in Wireless Computing** (with N. Master, A. Dua, D. Tsamis, J.P. Singh). IEEE Transactions on Wireless Communications, 15(5):3296-3310, 2016.
32. **Myopic Policies for Non-Preemptive Scheduling of Jobs with Decaying Value** (with N. Master). Journal of Probability in Engineering and Informational Sciences, November 2016.
33. **Maximum Weight Matching with Hysteresis in Overloaded Queues with Setups** (with C. Chan, M. Armony). Queueing Systems, 82:315-351, 2016,
34. **Cone Schedules for Processing Systems in Fluctuating Environments** (with K. Ross, G. Michailidis). IEEE Transactions of Automatic Control, 60(10):2710-2715, 2015.
35. **Personalized Monitors for Real-Time Detection of Physiological States** (with L. Chow, A. Gilman, A. Chander). International Journal of e-Health and Medical Communications, 5(4):1-19, October 2014.
36. **Air Traffic Control Area Configuration Advisories from Near-Optimal Distinct Paths** (with M. Bloem). Journal of Aerospace Information Systems, 11(11):764-784, November 2014.
37. **Ground Delay Program Analytics with Behavioral Cloning and Inverse Reinforcement Learning** (with M. Bloem). AIAA Journal of Aerospace Information Systems, 2015.
38. **Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions** (with C. Chan, V.F. Farias, G. Escobar). Operations Research (INFORMS), 60(6):1323-1341, Jan. 2013.
39. **BEST-AP: Non-Intrusive Estimation of Available Bandwidth and its Application for Dynamic Access Point Selection** (with P. Dely, L. Chow, N. Bayer, H. Einsiedler, C. Peylo). Computer Communications (Elsevier), 39:78-91, Jan 2014.
40. **A Software-Defined Networking Approach for Handover Management with Real-Time Video in WLANs** (with P. Dely, A. Kassler, L. Chow, N. Bayer, H. Einsiedler, C. Peylo, D. Mellado, M. Sanchez). Journal of Modern Transportation (Springer), 21:58-65, March 2013.
41. **Network Assisted Mobile Computing with Optimal Uplink Query Processing** (with C. Chan, J. Singh). IEEE Transactions on Mobile Computing, 60(6):1323-1341, Jan. 2013.
42. **Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions** (with C. Chan, V.F. Farias, G. Escobar). Operations Research (INFORMS), 60(6):1323-1341, Jan. 2013.
43. **Power and Delay Aware Management of Packet Switches** (with L. Mastroleon, D. O'Neill, B. Yolken). IEEE Transactions on Computers, , 61(12):1789-1799, Dec. 2012.
44. **Optimal State Surveillance under Observation Budget Constraints** (with P. Bommannavar). Int. Journal on Advances in Intelligent Systems, 4(4):57-67, Dec. 2011.
45. **A Scalable Delay-Power Control Algorithm for Bandwidth Sharing in Wireless Networks** (with F. Baccelli and N. Gast). IEEE/ACM Transactions on Networking, 19(5):1458-1471, October 2011.

46. **An Integrated Approach to Security Risk Management for IT-Intensive Organizations** (with Jeff Mounzer and Tansu Alpcan). Journal of Information Assurance and Security, 6(2):115-123, 2011.
47. **Channel, Deadline, and Distortion Aware Scheduling of Video Streams over Wireless Links** (with A. Dua, C. Chan, J. Apostolopoulos). IEEE Trans. on Wireless Communications, 9(3):1001-1011, March 2010.
48. **A Characterization of Max-Min SIR-Balanced Power Allocation with Applications** (with S. Stanczak, M. Kaliszan). ACM Wireless Networks, 16: 2335-2347, Nov. 2010.
49. **Game Based Capacity Allocation for Utility Computing Environments** (with B. Yolken). Telecommunications Systems Journal, pp. 1-17, May 2010.
50. **Scheduling Algorithms for Broadcasting Media with Multiple Distortion Measures**, (with C. W. Chan, S. Wee, J. Apostolopoulos), IEEE Trans. on Wireless Communications, 8(8) 4188-4199, 2009.
51. **Adaptive Data-Aware, Utility-Based Scheduling in Resource-Constrained Systems** (with D. Vengerov, L. Mastroleon, D. Murphy). Journal of Parallel and Distributed Computing, 70:871-879, 2009.
52. **Projective Cone Scheduling (PCS) Algorithms for Maximal Throughput Packet Switches** (with K. Ross). IEEE/ACM Transactions on Networking, 17(3):976-989, June 2009.
53. **Joint Task Migration and Power Management in Wireless Computing** (with S. Gitzenis). IEEE Transactions of Mobile Computing. 8(9): 1189-1204, Sept. 2009.
54. **Content-Aware Payout and Packet Scheduling for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos) IEEE Trans. on Multimedia, 10(5):885-895, Aug. 2008.
55. **Joint Transmitter Power Control and Mobile Cache Management in Wireless Computing** (with S. Gitzenis). IEEE Transactions of Mobile Computing, 7(4):498-512, April 2008.
56. **Transmission Power and Duration-Aware Payout Control for Packetized Media Streaming over Wireless Links** (with Y. Li). Wireless Communications and Mobile Computing, 8:309-326, Jan. 2008
57. **Downlink Wireless Packet Scheduling with Deadlines** (with A. Dua). IEEE Trans. on Mobile Computing, 6(12), pp. 1410-1425, 2007.
58. **TCP Performance Dynamics and Link-Layer Adaptation Based Optimization Methods for Wireless Networks** (with J. P. Singh, Y. Li, A. Bahai, B. Xu, G. Zimmerman). IEEE Trans. Wireless Communications, 6(5), pp. 1864-1878, 2007.
59. **Transmission Power and Duration-Aware Payout Control for Packetized Media Streaming over Wireless Links** (with Y. Li). Wireless Communications and Mobile Computing, 2008.
60. **Joint Power-Payout Control for Media Streaming over Wireless Links** (with Y. Li, A. Markopoulou, J. Apostolopoulos), IEEE Transactions in Multimedia, 2006.
61. **Optimal Processor Allocation to Differentiated Job Flows** (with G. Michailidis and K. Wasserman), Performance Evaluation, v. 63(1), pp. 1-14, Jan. 2006.
62. **Queueing Networks of Random Link Topology: Stationary Dynamics of Maximum Throughput Schedules** (with G. Michailidis), Queueing Systems Theory and Applications, v. 50(1), pp. 5-52, May 2005.

63. **A Fuzzy Reinforcement Learning Approach to Power Control in Wireless Transmitters** (with D. Vengerov and H. Berenji). IEEE Trans. SMC (Part B), v. 35-4, pp. 768-778, Aug. 2005.
64. **Queueing and Scheduling in Random Environments** (with G. Michailidis). Advances in Applied Probability, vol. 36 (1), 293-317, 2004.
65. **Scheduling Optical Traffic Bursts in Time-Domain Wavelength Interleaved Optical Networks** (with K. Ross, K. Kumaran, I. Saniee, I. Widjaja). IEEE Journal of Selected Areas in Communications, vol. 21(9), pp. 1441-1451, 2003.
66. **Queueing Dynamics and Maximal Throughput Scheduling in Switched Processing Systems** (with M. Armony). Queueing Systems Theory and Applications, vol. 44(3), pp. 209-252, 2003.
67. **On Stability of Queueing Networks with Job Deadlines** (with A. Ward). Journal of Applied Probability, vol. 40(2), pp. 293-304, 2003.
68. **On Parallel Queueing with Random Server Connectivities and Routing Constraints** (with G. Michailidis). Probability in the Engineering and Informational Sciences, vol. 16, pp. 185-203, 2002.
69. **Optimal Control of Parallel Queues with Batch Service** (with C. Xia, G. Michailidis, P. Glynn). Probability in the Engineering and Informational Sciences, vol. 16, pp. 289-307, 2002.
70. **Power Controlled Multiple Access (PCMA) Schemes in High Performance Wireless Packet Networks** (with S. Kandukuri). IEEE Personal Communications, vol. 9(3), pp.58-64, 2002.
71. **Dynamic Online Task Scheduling on Parallel Processors** (with C. Xia, G. Michailidis). Performance Evaluation, vol. 46, pp. 219-233, 2001.
72. **Globally Constrained Power Control Across Multiple Channels in Wireless Packet Networks** (with S. Kandukuri). ACM Mobile Networks, vol. 6(5), pp. 427-434, 2001. (invited paper)
73. **Channel Access Algorithms with Active Link Protection for Wireless Communication Networks with Power Control** (with S. Chen, G. Pottie). IEEE/ACM Transactions on Networking, vol. 8(5), pp. 583-597, 2000.
74. **On Synchronization of Poisson Processes and Queueing Networks with Service and Synchronization Nodes** (with B. Prabhakar, T. Mountford). Advances in Applied Probability, vol. 32(3), pp. 824-843, 2000.
75. **Power-Induced Time Division on Asynchronous Channels** (with J. Rulnick). ACM Wireless Networks, vol. 5(2), pp. 71-80, 1999. (invited paper)
76. **Towards Power-Sensitive Network Architectures in Wireless Communications: Concepts, Issues and Design Aspects**. IEEE Personal Communications, vol. 5(3), pp. 50-59, 1998. (invited paper)
77. **Mobile Power Management for Wireless Communication Networks** (with J. Rulnick). ACM Wireless Networks, vol. 3(1), pp. 3-14, 1997. (invited paper)
78. **On a Singular Feature of Critical G/M/1 Queues** (with B. Prabhakar). Systems and Control Letters, vol. 28(5), pp. 239-245, 1996.
79. **Wireless, Mobile, Multimedia Networks** (with A. Alwan, R. Bagrodia, N. Bambos, J. Cong, M. Gerla, L. Kleinrock, J. Villasenor). IEEE Personal Communications, vol. 3(2), pp. 34-51, 1996.

80. **Convergence of Departures in Tandem Networks of  $^*/GI/Infinity$  Queues** (with B. Prabhakar, T. Mountford). Probability in the Engineering and Informational Sciences, vol. 10(4), pp. 487-500, 1996.
81. **The Supercomputer Supernet Testbed: A WDM-Based Supercomputer Interconnect** (with L. Kleinrock, M. Gerla, J. Cong, E. Gafni, L. Bergman, J. Bannister, S.P. Monacos, T. Bujewski, P-C. Hu, B. Kannan, B. Kwan, E. Leonardi, J. Peck, P. Palnati, S. Walton). IEEE/OSA Journal of Lightwave Technology, vol. 14(6), pp. 1388-1399, 1996.
82. **Optimal Server Allocation to Parallel Queues with Finite Capacity Buffers** (with K. Wasserman). Probability in the Engineering and Informational Sciences, vol. 10(2), pp. 279-285, 1996.
83. **Asymptotic Optimality of Statistical Multiplexing in Pipelined Processing** (with K. Wasserman). Queueing Systems Theory and Applications, vol. 21, pp. 97-123, 1995.
84. **OPTIMIC: A Scalable, Distributed, All-Optical Terabit Network** (with L. Kleinrock, J. Cong, E. Gafni, M. Gerla, L. Bergman, J. Bannister). Journal of High Speed Networks, vol. 4(4), pp. 407-424, 1995.
85. **On Infinite Queueing Tandems** (with B. Prabhakar). Systems and Control Letters, vol. 23, pp. 305-314, 1994.
86. **Optimality Aspects of Greedy Schemes in Parallel Processing of Random Graph-Structured Jobs** (with S. Chen). Probability in the Engineering and Informational Sciences, vol. 8, pp. 229-243, 1994.
87. **On Stationary Tandem Queueing Networks with Job Feedback** (with K. Wasserman). Queueing Systems Theory and Applications, vol. 15, pp. 137-164, 1994.
88. **Scheduling and Stability Aspects of a General Class of Parallel Processing Systems** (with J. Walrand). Advances in Applied Probability, vol. 25, pp. 176-202, 1993.
89. **On Closed Ring Queueing Networks**. Journal of Applied Probability, vol. 29(4), pp. 979-995, 1992.
90. **On Flows in Stochastic Marked Graphs** (with F. Baccelli & J. Walrand). Probability in the Engineering and Informational Sciences, vol. 5, pp. 145-157, 1991.
91. **On Stability and Performance of Parallel Processing Systems** (with J. Walrand). Journal of the Association for Computing Machinery, vol. 38(2), pp. 429-452, 1991.
92. **An Invariant Distribution for the G/G/1 Queueing Operator** (with J. Walrand). Advances in Applied Probability, vol. 22, pp. 254-256, 1990.
93. **On the Asymptotic Execution Time of Multi-Tasked Processes on Tandem Processors** (with J. Walrand). Systems and Control Letters, vol. 13, pp. 391-396, 1989.
94. **On State-Dependent Queues and Acyclic Queueing Networks** (with J. Walrand). Advances in Applied Probability, vol. 21, pp. 681-701, 1989.
95. **On Queues with Periodic Inputs** (with J. Walrand). Journal of Applied Probability, vol. 26, pp. 381-389, 1989.

---

## PEER-REVIEWED CONFERENCE PAPERS:

96. **Optimal Control of Remote Patient Monitoring with Multidimensional Health States** (S. Chandak, I. Thapa, N. Bambos, D. Scheinker). Proceedings of the 2025 IEEE International Conference in Communications (ICC 2025). Montreal, Canada, June 2025. [accepted]
97. **Accelerated Regularized Learning in Finite N-Person Games** (K. Lotidis, A. Giannou, P. Mertikopoulos, N. Bambos). Proceedings of the 38<sup>th</sup> Conference of Neural Information Processing (NeurIPS 2024), pp. 1-32, Vancouver, Canada, Dec. 2024.
98. **Power-Managed Data Centers for Sustainable Computing** (E. Zeger, N. Bambos, M. Pilanci). Proceedings of the 2024 IEEE International Conference in Communications (ICC 2024), pp. 2833-2839, Denver, CO, June 2024. **Best Paper Award** in Green Communications and Networking.
99. **Tiered Service Architecture for Remote Patient Monitoring** (S. Chandak, I. Thapa, N. Bambos, D. Scheinker). Proceedings of the 2024 IEEE International Conference on e-Health Networking, Application & Services (HealthCom 2024), pp. 1-7, Nara, Japan, Nov. 2024.
100. **Federated Learning with Social Data Sharing** (M. Jazi, I. Bistriz, N. Bambos, I. BenGal). Proceedings of the 7<sup>th</sup> International Conference on Artificial Intelligence in Information and Communication, Fukuoka, Japan, Feb. 2025.
101. **Payoff-based Learning with Multiplicative Weights in Quantum Games** (K. Lotidis, P. Mertikopoulos, N. Bambos, J. Blanchet). Proceedings of the 37<sup>th</sup> Conference on Neural Information Processing Systems (NeurIPS 2023), pp. 1-33, New Orleans, LA, Dec. 2023.
102. **The Stability of Matrix Multiplicative Weights Dynamics in Quantum Games** (K. Lotidis, P. Mertikopoulos, N. Bambos). Proceedings of the 2023 62<sup>nd</sup> IEEE Conference on Decision and Control (CDC 2023), pp. 1225-1232, Marina Bay Sands, Singapore, Dec. 2023.
103. **Tug of Peace: Distributed Learning for Quality-of-Service Guarantees** (S. Chandak, I. Bistriz, N. Bambos). Proceedings of the 2023 62<sup>nd</sup> IEEE Conference on Decision and Control (CDC 2023), pp. 2346-2351, Marina Bay Sands, Singapore, Dec. 2023.
104. **Admission Control for Games with a Dynamic Set of Players** (S. Chandak, I. Bistriz, N. Bambos). Proceedings of the 2023 62<sup>nd</sup> IEEE Conference on Decision and Control (CDC 2023), pp. 1219-1224, Marina Bay Sands, Singapore, Dec. 2023.
105. **Personalized and Energy-Efficient Health Monitoring: A Reinforcement Learning Approach** (B. Eden, I. Bistriz, N. Bambos, I. Ben-Gal, E. Khmelnitsky), Presented at the 2023 American Control Conference (ACC) (and published in the IEEE Control Systems Letters, 7:955-960, 2023), San Diego, CA, Dec. 2023.
106. **Equilibrium Bandits: Learning Optimal Equilibria of Unknown Dynamics** (S. Chandak, I. Bistriz, N. Bambos). Proceedings of the 22<sup>nd</sup> International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), pp. 1-18; London, UK, May 2023.
107. **Wasserstein Distributionally Robust Linear-Quadratic Estimation under Martingale Constraints** (K. Lotidis, N. Bambos, J. Blanchet, J. Li). Proceedings of the 21<sup>st</sup> International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 1-15; Valencia, Spain, April 2023.

108. **Power-Optimized Processor Slowdown Control** (A.J. Mann, N. Bambos). *Proceedings of the IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2022.
109. **Learning in Games with Quantized Payoff Observations** (K. Lotidis, P. Mertikopoulos, N. Bambos). *Proceedings of the IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2022.
110. **Framed Projective Cone Scheduling: Latency vs. Context-Switching Tradeoff in Data Centers** (E. Zeger, A.J. Mann, N. Bambos). *Proceedings of the IEEE Global Communications Conference*, Rio de Janeiro, Brazil, Dec. 2022.
111. **Queue Up your Regrets: Achieving the Dynamic Capacity Region of Multi-player Bandits** (I. Bistriz, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
112. **Active Testing for an Emerging Epidemic** (A.J. Mann, I. Bistriz, N. Bambos). *Proceedings of the IEEE International Conference on e-Health Networking, Applications & Services (HealthCom)*, pp. 7-12, Genoa, Italy, 2022.
113. **Online Learning for Load Balancing of Unknown Monotone Resource Allocation Games** (I. Bistriz, N. Bambos). *Proceedings of the 38<sup>th</sup> International Conference on Machine Learning (ICML)*, PMLR 139:968-979, July 2021.
114. **Controlling Epidemics via Testing** (K. Lotidis, A. Moustakas, N. Bambos). *Proceedings of the 2021 60<sup>th</sup> Conference on Decision and Control (CDC)*, pp. 2092-2097, Austin, TX, Dec. 2021.
115. **An OpenAI-OpenDSS Framework for Reinforcement Learning on Distribution-Level Microgrids** (K. Moy, C. Tae, Y Wang, G. Henri, N. Bambos, R. Rajagopal). *Proceedings of the 2021 IEEE Power & Energy Society General Meeting (PESGM)*, pp. 1-5, Washington, DC, July 2021.
116. **Power Controlled Slowdown in Data Centers** (A. Mann, N. Bambos). *Proceedings of IEEE International Conference on Communications (ICC)*, pp. 1214-1219, Seoul, South Korea, May 2022.
117. **Cooperative Multi-Player Bandit Optimization** (I. Bistriz, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 22 pages, 2020.
118. **Distributed Distillation for On-Device Learning** (I. Bistriz, A. Mann, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 26 pages, 2020.
119. **My Fair Bandit: Distributed Learning of Max-Min Fairness with Multi-Player Bandits** (I. Bistriz, T.Z. Baharav, A. Leshem, N. Bambos). *Proceeding of the International Conference on Machine Learning (ICML)*, PMLR 119: 930-940, 10 pages, 2020.
120. **Distributed Scheduling of Charging for On-Demand Electric Vehicle Fleets** (I. Bistriz, M. Klein, N. Bambos, O. Maimon, R. Rajagopal). *Proceedings of the IFAC Workshop on Discrete Event Systems (WODES)*, 6 pages, 2020. [**Best student paper finalist award**; 1 best paper award and 3 finalist awards out of 75 papers presented at the conference.]
121. **Operationally-Informed Hospital-Wide Discharge Prediction Using Machine Learning** (by A. Ward, A. Mann, J. Vallon, G. Escobar, N. Bambos, A. Schuler). *Proceedings of the 22<sup>nd</sup> IEEE Int. Conf. on e-Health Applications & Services (HealthCom)*, 6 pages, Shenzhen, China, Dec. 2020. [**Best Paper Runner-Up Award**; 1 best paper award and 2 runner-up awards out of 100 papers presented at the conference.]

122. **Rates of, and Risk Factors for, Substance Use Disorder in Children and Young Adults Before and After Surgery** (I. Thapa, A. Ward, B. De Souza, N. Bambos, T. Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020; extended abstract published in the Journal of Anesthesia & Analgesia Supplement, 130(5)S:689-690.*
123. **Rates of, and Risk Factors for, Chronic Pain in Children and Young Adults Before and After Surgery** (I. Thapa, A. Ward, B. De Souza, N. Bambos, T. Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020; extended abstract published in the Journal of Anesthesia & Analgesia Supplement, 130(5)S:691.*
124. **Development and Validation of a Machine Learning Algorithm for Prediction of Prolonged Opioid Use After Surgery in Adolescents and Young Adults** (Trisha Jani, A. Ward, B. De Souza, N. Bambos, Thomas A Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020; extended abstract published in the Journal of Anesthesia & Analgesia Supplement, 130(5)S:712-713.*
125. **Distributed Scheduling of Charging for On-Demand Electric Vehicle Fleets** (with I. Bistritz, M. Klein, O. Maimon, R. Rajagopal). *Proceedings of the 2020 Workshop on Discrete-Event Systems (WODES), Rio de Janeiro, Brazil, Nov. 2020.*
126. **Online EXP3 Learning in Adversarial Bandits** (with I. Bistritz, Z. Zhou, X. Chen, J. Blanchet). *Proceedings of the Conference on Advances in Neural Information Processing Systems (NeurIPS), Vancouver, Canada, December 2019.*
127. **Controlling Contact Network Topology to Prevent Measles Outbreaks** (with I. Bistritz, D. Kahana, I. Ben-Gal, D. Yamin). *Proceedings of the 2019 IEEE Global Communications Conference (Globecom), pp. 1-6, Waikoloa, Hawaii, December 2019.*
128. **Asymptotically Optimal Distributed Gateway Load-Balancing for the Internet of Things** (with I. Bistritz). *Proceedings of the 2019 IEEE International Conference on the Network of the Future (NoF), pp. 98-101, Rome, Italy, October 2019.*
129. **Noninvasive Identification of Hypotension Using Convolutional-Deconvolutional Networks** (with D. Miller, A. Ward, A. Shin, D. Scheinker). *Proceedings of 2019 IEEE International Conference on E-Health Networking, Applications and Services (HealthCom), pp. 1-6, Bogota, Colombia, October 2019.*
130. **Learning to Emulate an Expert Projective Cone Scheduler** (with A. Ward, N. Master). *Proceedings of the 2019 IEEE American Control Conference (ACC), pp. 292-297, Philadelphia, PA, July 2019.*
131. **Incidence of Persistent Opioid Use in Children after Surgery** (with T. Anderson, A. Ward, B. De Souza, E. Wang, D. Miller). *Proceedings of the International Anesthesia Research Society 2019 Annual Meeting and International Science Symposium, May 17-20, Montreal, Quebec, Canada; published in the Journal: Anesthesia & Analgesia Supplement 128(5)S:354:764-765.*
132. **Smart Greedy Distributed Allocation in Microgrids** (with I. Bistritz, A. Ward, Z. Zhou). *Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1-6, Shanghai, China, May 2019.*
133. **Anesthesiologist Surgery Assignments Using Policy Learning** (with A. Ward, Z. Zhou, Ellen Wang, D. Scheinker). *Proceedings of the 2019 IEEE International Conference on Communications (ICC), 6 pages, Shanghai, China, May 2019.*

134. **Learning Health State Transition Probabilities via Wireless Body Area Networks** (with Tal Geller, Yair Bar David, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1-6, Shanghai, China, May 2019.
135. **The Power of Consensus: Optimal Distributed Multichannel Wireless Transmitter Power Control** (with I. Bistriz). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1- 6, Shanghai, China, May 2019.
136. **Bidding-Based Dynamic Power Pricing Scheme in Smart Grids** (with A. Ward, Z. Zhou). Proceedings of the 2019 IEEE International Conference on Computing, Networking and Communications (ICNC), pp. 729-734, Honolulu, Hawaii, February 2019.
137. **Learning in Games with Lossy Feedback** (with Z. Zhou, P. Mertikopoulos, S. Athey, P. Glynn and Y. Ye). Proceedings of 32<sup>nd</sup> Neural Information Processing Systems (NIPS) Conference, 11 pages, Montreal, Canada, December 2018.
138. **Distributed Asynchronous Stochastic Optimization under Unbounded Delays: How Slow Can You Go?** (with Z. Zhou, P. Mertikopoulos, P. Glynn and Y. Ye, Li-Jia Li, Li Fei-Fei). Proceedings of the 35<sup>th</sup> International Conference on Machine Learning (ICML), 10 pages, Stockholm, Sweden, July 2018.
139. **Power Control with Random Delays: Robust Feedback Averaging** (with A. Ward, Z. Zhou, P. Mertikopoulos). Proceedings of the 57<sup>th</sup> Conference on Decision and Control (CDC), pp. 7040-7045, Miami Beach, FL, December 2018.
140. **Robustness of Join-the-Shortest-Queue Scheduling to Communication Delay** (with S. Mehdian, Z. Zhou). Proceedings of the 2018 American Control Conference (ACC), pp. 3708-3713, Milwaukee, WI, June 2018.
141. **Optimal Health Monitoring via Wireless Body Area Networks** (with Yair Bar David, Tal Geller, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 57<sup>th</sup> Conference on Decision and Control (CDC), pp. 7040-7045, Miami Beach, FL, December 2018.
142. **Sensing-Constrained Power Control in Digital Health** (with D. Miller, Z. Zhou, N. Bambos, I. Ben-Gal). Proceedings of the 2018 American Control Conference (ACC), pp. 3708-3713, Milwaukee, WI, June 2018.
143. **Optimal Sensing for Patient Health Monitoring** (with D. Miller, Z. Zhou, I. Ben-Gal). Proceedings of the 2018 International Communications Conference (ICC), 7 pages, Kansas City, MO, May 2018.
144. **Physiological Waveform Imputation of Missing Data using Convolutional Autoencoders** (with D. Miller, A. Ward, D. Scheinker, Andrew Shin). Proceedings of the 2018 IEEE 20<sup>th</sup> International Conference on e-Health Networking, Applications and Services (Healthcom), pages 6, Ostrava, Czech Republic, September 2018.
145. **Automatic Sleep Arousal Identification from Physiological Waveforms Using Deep Learning** (with D. Miller, A. Ward, N. Bambos). Proceedings of 2018 Computing in Cardiology (CinC) Conference, pages 5, Maastricht, Netherlands, September 2018.
146. **Bidding-Based Dynamic Power Pricing Scheme in Smart Grids** (with A. Ward, Z. Zhou). Proceedings of the 2018 IEEE International Conference on Computing, Networking and

Communications (ICNC), pp. 729-734, Honolulu, Hawaii, February 2019.

147. **Smart Greedy Distributed Allocation in Microgrids** (with I. Bistriz, A. Ward, Z. Zhou). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
148. **Anesthesiologist Surgery Assignments Using Policy Learning** (with A. Ward, Z. Zhou, Ellen Wang, D. Scheinker). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
149. **Learning Health State Transition Probabilities via Wireless Body Area Networks** (with Tal Geller, Yair Bar David, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
150. **The Power of Consensus: Optimal Distributed Multichannel Wireless Transmitter Power Control** (with I. Bistriz). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
151. **Countering Feedback Delays in Multi-Agent Learning** (Z. Zhou, P. Mertikopoulos, N. Bambos, P. Glynn, C. Tomlin), Proceedings of Advances for Neural Information Processing Systems (NIPS), pp.1-11, Long Beach, CA, Dec. 2017.
152. **Stochastic Mirror Descent in Variationally Coherent Optimization Problems** (with Z. Zhou, P. Mertikopoulos, S. Boyd, P. Glynn), Proceedings of Advances for Neural Information Processing Systems (NIPS), pp. 1-11, Long Beach, CA, Dec. 2017.
153. **Classical Communication Network Design Via Quantum Heuristics** (with A. Ward, N. Master). Proceeding of the Conference on Decision and Control (CDC), pp. 1-6, Long Beach, CA, Dec. 2017.
154. **Mirror Descent Learning in Continuous Games** (with Z. Zhou, P. Mertikopoulos, A. Moustakas, P. Glynn). Proceeding of the Conference on Decision and Control (CDC), pp. 1-6, Long Beach, CA, Dec. 2017.
155. **Stable Power Control in Wireless Networks via Dual Averaging** (with Z. Zhou, P. Mertikopoulos, A. Moustakas, S. Mehdian, Peter Glynn). Proceedings of IEEE Global Communications Conference (GLOBECOM), pp. 1-6, Singapore, Dec. 2017.
156. **Longest-Queue-First Scheduling with Intermittent Sampling** (with S. Mehdian, Z. Zhou). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
157. **Least Action Routing: Identifying the Optimal Path in a Wireless Relay Network** (with A. Moustakas, P. Mertikopoulos, Z. Zhou). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
158. **Market-Based Dynamic Service Mode Switching in Virtualized Wireless Networks** (with M. Dimakopoulou, M. Valdez-Vivas, J. Apostolopoulos). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
159. **Video Streaming Services for Industrial IoT** (with H. Kanzaki, K. Schubert). Proceedings of International Conference on Computer Communications and Networks (ICCCN), pp. 1-7, Vancouver, CA, Aug. 2017.

160. **Dynamic Control of Data Center Network and Computation Resources** (with D. Miller, Li Xia, Z. Zhou, N. Bambos). Proceedings of the International Conference on Computing, Networking and Communications (ICNC), pp. 1-6, San Jose, CA 2017.
161. **Asynchronous Best Response Dynamics for Resource Allocation Games in Utility Computing** (with K. Schubert, N. Master, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.
162. **An Infinite Dimensional Model for a Single-Server Priority Queue** (with N. Master, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.
163. **Join-the-Shortest Queue Scheduling with Delay: A Case Study** (with S. Mehdian, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.
164. **Cost-Sensitive Security Risk Management for Large-Scale Computing Infrastructures** (with N. Master). Proceedings of the IEEE International Conference on Computing, Networking and Communications (ICNC), 6 pages, San Jose, CA, January 2017.
165. **An Infinite Dimensional Model for a Many Server Priority Queue** (with N. Master, Z. Zhou). Proceeding of the Conference on Information Sciences and Systems (CISS), 6 pages, Johns Hopkins University, Baltimore, MD, March 2017.
166. **Wireless Channel Selection with Non-Invasive Power Probing** (with K. Schubert). Proceedings of IEEE Wireless Telecommunications Symposium, 6 pages, Chicago, IL, April 2017.
167. **Infinite Server Queueing Networks with Deadline Bases Routing** (with N. Master), IEEE Conference on Decision and Control (CDC), pp. 5360-5366, Las Vegas, NV, December 2016.
168. **Repeated Games for Power Control in Wireless Communications: Equilibrium and Regret** (with Z. Zhou, P. Glynn). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 3603-3610, Las Vegas, NV, December 2016.
169. **A Game-Theoretical Formulation of Influence Networks** (w Z. Zhou, B. Yolken, R.A. Miura-Ko). Proceeding of IEEE American Control Conference, pp. 3802-3807, Boston, MA, July 2016.
170. **Dynamics of Linear Influence Networks under Stochastic Environments** (with Z. Zhou, P. Glynn). Proceeding (Springer) of Conference on Decision and Game Theory for Security (GameSec), 6 pages, New York, NY, November 2016.
171. **A Stochastic Stability Characterization of the Foschini-Miljanic Algorithm in Random Wireless Networks** (with Z. Zhou, D. Miller, P. Glynn). Proceedings of IEEE Global Communications Conference (GLOBECOM), 6 pages, Washington, DC, December 2016.
172. **Wireless Power-Controlled TCP with Holdover** (with K. Schubert). Proceedings of IEEE International Conference on Communications (ICC), 6 pages, Kuala Lumpur, Malaysia, May 2016.
173. **Data Aggregation for Low Power Wireless Devices** (with K. Schubert). Proceedings of IEEE MILCOM, 6 pages, November 2016.
174. **Detecting Inaccurate Predictions of Pediatric Surgical Durations** (with Z. Zhou, D. Miller, N. Master, P. Glynn). IEEE International Conference on Data Science and Advanced Analytics (DSAA), pp. 452-457, Montreal, Canada, October 2016.

175. **Reliable and Efficient Performance Monitoring in Linux** (with M. Dimakopoulou, S. Eranian, N. Koziris). Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis, 11 pages, Salt Lake City, Utah, November 2016.
176. **Scalable Data Center Power Management via a Global Stress Signal** (with D. Miller, N. Master, Z. Zhou, N. Bambos). IEEE Global Communications Conference (GLOBECOM), 7 pages, 2015.
177. **Target-Rate Driven Resource Sharing in Queueing Systems** (with Z. Zhou). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 4490-4945, Osaka, Japan, December 2015.
178. **Wireless Communication Games in Fixed and Random Environments** (with Z. Zhou). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 1637-1642, Osaka, Japan, December 2015.
179. **Dynamic Management of Network Risk from Epidemic Phenomena** (A. Sinha, J. Ducci). IEEE Conference on Decision and Control (CDC), 6 pages, Osaka, Japan, December 2015.
180. **Low Latency Policy Iteration via Parallel Processing and Randomization** (with N. Master). IEEE Conference on Decision and Control (CDC), pp. 1084-1091, 2015.
181. **Channel Responsive Wireless TCP** (with K. Schubert, H. Endo). IEEE International Communications Conference (ICC'14), 6 pages, London, UK, 2015.
182. **A General Model for Resource Allocation in Utility Computing** (with Z. Zhou). American Control Conference, pp. 1746-1751, Chicago, IL, 2015.
183. **Service Rate Control for Jobs with Decaying Value** (with N. Master). American Control Conference, 6 pages, Chicago, IL, 2015.
184. **Ground Delay Program Analytics with Behavioral Cloning and Inverse Reinforcement Learning** (with M. Bloem). Proceedings of AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, 6 pages, Atlanta, GA, June 2014.
185. **Infinite Time Horizon Maximum Causal Entropy Inverse Reinforcement Learning** (with M. Bloem). IEEE Conference on Decision and Control (CDC'14), pp. 4911-4916, Los Angeles, CA, December 2014.
186. **Power-Controlled Multiple Access with Queue-Dependent Backoff Threshold** (with J. Mounzer, K. Schubert, A. Goldsmith). IEEE Global Communications Conference (Globecom'14), pp. 4738-4744, Austin, TX, December 2014.
187. **Randomized Iterations for Low Latency Fixed Point Computation** (with N. Master). IEEE Conference on Decision and Control (CDC'14), pp. 5208-5215, Los Angeles, CA, December 2014.
188. **Distributed Smart Grid Architecture for Delay and Price Sensitive Power Management** (with N. Master, J. Mounzer), accepted to the IEEE International Communications Conference (ICC'14), pp. 3670-3675, Sydney, Australia, June 2014.
189. **Dynamic Resource Management in Virtualized Data Centers with Bursty Traffic** (with M. Valdez-Vivas, J. Apostolopoulos), accepted to the IEEE International Communications Conference (ICC'14), Sydney, Australia, June 2014.
190. **Power Control for Wireless Streaming with HOL Packet Deadlines** (with N. Master), accepted to the IEEE International Communications Conference (ICC'14), pp. 2263-2269, Sydney, Australia, June 2014.

191. **Channel Exploration for Wireless Media Streaming with Handoff and Rebuffering Control** (with L. Chow, J.P. Singh), accepted to the IEEE International Communications Conference (ICC'14), pp. 2281-2288, Sydney, Australia, June 2014.
192. **Channel Aware Rebuffering for Wireless Media Streaming with Handoff Control** (with L. Chow, B. Collins, P. Dely, A. J. Kassler, C. Peylo, H. J. Einsiedler, N. Bayer). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), pp. 4434-4439, Atlanta, Georgia, December 2013.
193. **Congestion versus Accuracy Tradeoffs in IP Traffic Classification** (with M. Valdez-Vivas). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), Atlanta, Georgia, December 2013.
194. **Playout Buffer Responsive Wireless Streaming for Multiple Clients** (with P. Bommannavar, J. Apostolopoulos). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), pp. 2276-2281, Atlanta, Georgia, December 2013.
195. **QoS-Aware Admission Control in Heterogeneous Data Centers** (with C. Delimitrou, C. Kozyrakis). Proceedings of USENIX International Conference on Autonomic Computing, San Jose, CA, June 2013.
196. **Delay-Sensitive Power Management for Packet Switches** (with M. Valdez-Vivas, D. O'Neill). Proceedings of IEEE International Communications Conference (ICC'13), Budapest, Hungary, June 2013.
197. **Latency and Energy in Quality-Driven Applications for Networked Wireless Devices** (with M. Valdez-Vivas). Proceedings of IEEE International Communications Conference (ICC'13), Budapest, Hungary, June 2013.
198. **Resource Allocation and Scheduling for Energy Efficient Tracking** (with P. Bommannavar, J. Apostolopoulos). Proceedings of IEEE International Communications Conference (ICC'13), pp. 1532-1537, Budapest, Hungary, June 2013.
199. **Deadline Aware Packet Scheduling in Switches for Multimedia Streaming Applications** (with P. Bommannavar, J. Apostolopoulos). Proceedings of IEEE International Communications Conference (ICC'13), pp. 3797-3802, Budapest, Hungary, June 2013.
200. **Power Optimization in Random Wireless Networks (with A.L. Moustakas)**. Proceedings of IEEE International Symposium on Information Theory (ISIT'13), Istanbul, Turkey, July 2013.
201. **Real-Time Physiological Stream Processing for Health Monitoring Services** (with L. Chow). Proceedings of International Conference on e-Health Networking, Applications and Services (HealthCom'13), pp. 611-616, Lisbon, Portugal, Oct. 2013.
202. **An Approach for Finding Multiple Areas of Specialization Configuration Advisories** (with M. Bloem). Proceedings of AIAA Aviation Technology, Integration, and Operations (ATIO), 6 pages, Los Angeles, CA, August 2013.
203. **Playout-Buffer Aware Hand-Off Control for Wireless Video Streaming** (with L. Chow, B. Collins, C. Peylo, H. Einsiedler, N. Bayer, P. Dely, A. Kassler), IEEE Global Communications Conference (GlobeComm), pp. 5237-5242, Los Angeles, CA, 6 pages, Dec. 2012.
204. **Resource Constrained Failure Management in Networked Computing Systems** (with P. Bommannavar). IEEE Global Communications Conference (GlobeComm), pp. 1884-1889, Los

Angeles, CA, 6 pages, Dec. 2012.

205. **Approximating the Likelihood of Historical Airline Actions to Evaluate Airline Delay Cost Functions** (with M. Bloem, H. Huang). IEEE Conference on Decision and Control (CDC), Maui, HI, pp. 508-513, Dec. 2012.
206. **Power Optimization on a Network: The Effect of Randomness** (with A. Moustakas), IEEE Int. Symposium of Information Theory, pp. 606-610, June 2012.
207. **Nonlinear Cooperative Dynamics in Distributed Power Control for Wireless Networks** (with J. Mounzer). IEEE Int. Communications Conference (ICC), pp. 5042-5047, June 2012.
208. **Evaluation of Algorithms for a Miles-in-Trail Decision Support Tool** (with M. Bloem, D. Hattaway, and N. Bambos). Int. Conference on Research in Air Transportation, Berkeley, CA, 8 pages, May 2012.
209. **Power Budgeted Packet Scheduling for Wireless Multimedia** (with P. Bommannavar, J. Apostolopoulos). International Communications Conference (ICC), pp. 2108-2113, June 2012.
210. **Power Control in Random Networks: The Effect of Disorder in User Positions** (with A. Moustakas). IEEE Int. Workshop for Spatial Stochastic Models for Wireless Networks (SPASWIN), pp. 380-385, May 2012.
211. **Intrusion Detection in IT Networks with Limited Observations** (with P. Bommannavar) IEEE International Conference on Computing, Networking and Communications (ICNC), Maui, Hawaii, pp. 6, Nov. 2012.
212. **Coordinated Tactical Air Traffic and Airspace Management** (with M. Bloem) IEEE Conference on Decision and Control (CDC), pp. 5287-5292, Orlando, FL, Dec. 2011.
213. **Implementation of Fine-Grained Power/Bitrate Control for 802.11 Devices** (with Min-Wook Jeon, A. Kabani). Int. Tech. Conference on Circuits/Systems, Computers and Communications, Gyeongju, Korea, June 2011.
214. **Security Risk Management in Computing Systems with Constraints on Service Disruption** (with P. Bommannavar), IEEE International Conference on Computer Communication Networks (ICCCN), pp. 6, Maui, Hawaii, June 2011.
215. **Security Risk Management via Dynamic Games with Learning** (with Praveen Bommannavar and Tansu Alpcan. IEEE Int. Communications Conference (ICC), pp. 1-6, Kyoto, Japan, June 2011.
216. **Dynamic Control and Mitigation of Interdependent IT Security Risks** (with J. Mounzer, T. Alpcan) IEEE Int. Communications Conference (ICC), pp. 1-6, Cape Town, South Africa, May 2010 (Best Paper Award).
217. **Optimal State Surveillance under Budget Constraints** (with P. Bommannavar) Int. Conf. on Emerging Network Intelligence (Emerging'10), pp. 1-6, Florence, Italy, Oct. 2010 (Best Paper Award).
218. **Hybrid Power Control Algorithms for Streaming and Data Traffic in Wireless Networks** (with J. Mounzer) IEEE Int. Communications Conference (ICC), pp. 1-6, Cape Town, South Africa, May 2010.

219. **Integrated Security Risk Management for IT-Intensive Organizations** (with J. Mounzer, T. Alpcan) Sixth Int. Conf. on Information Assurance and Security (IAS), pp. 329-334, Atlanta, GA, Aug. 2010.
220. **Locking Dynamics and Mitigation Schemes in Distributed Power Control for Wireless Networks** (with J. Mounzer) IEEE Global Communications Conference (GlobeComm), pp. 1-6, Dec. 2010 .
221. **Fairness of Heterogeneous Queues through Cone Scheduling** (with C. Chan, M. Armony). Manufacturing and Service Operations Management Society Conference, pp. 1-3, Technion, Haifa, Israel, June 2010.
222. **Patch Scheduling for Risk Exposure Mitigation under Service Disruption Constraints** (with P. Bommannavar) 4<sup>th</sup> Int. Conf. on Emerging Security Information Systems and Technologies (SecureWare), pp. 176-181, Venice, July 2010.
223. **Admission Control for Autonomous Wireless Links with Power Constraints** (with M. Kaliszan, S. Stanczak). IEEE Int. Conf. in Acoustics Speech and Signal Processing, pp. 3138-3141, Dallas, TX, March 2010.
224. **Backlog aware scheduling for ingress memories in high-radix, single-stage switches** (with D. Tsamis, B. Yolken, W. Olesinski, H. Eberle, N. Gura). IEEE Global Com. Conference (Globecom), Honolulu, HI, pp. 1-6, Dec. 2009
225. **Throughput Loss in Task Scheduling due to Server State Uncertainty** (with C. Chan). ACM Int. Conf. in Performance Evaluation Methodologies and Tools (ValueTools), pp. 1-10, Pisa, Italy, 2009.
226. **Game Theoretic Rate Control for Mobile Devices** (with D. Tsamis, T. Alpcan). IEEE Int. Conf. on Game Theory for Networks (Gamenets), Instabul, Turkey, 2009.
227. **Dynamic Resource Modeling for Heterogeneous Wireless Networks** (with D. Tsamis, T. Alpcan, J. Singh). IEEE Int. Conf. on Communications (ICC), Dresden, Germany, pp. 1-6, 2009.
228. **Backlog aware scheduling for ingress memories in high-radix, single-stage switches** (with D. Tsamis, B. Yolken, W. Olesinski, H. Eberle, N. Gura). IEEE Global Com. Conference (Globecom), Honolulu, HI, pp. 1-6, 2009.
229. **Cost and Target-Based Scheduling for Switch Power Control** (with B. Yolken, D. Tsamis). IEEE International Conference on Communications (ICC), Dresden, Germany, 2009.
230. **Modeling Dependencies in Security Risk Management** (with T. Alpcan) ACM Int. Conf. on Risk and Security of Internet and Systems (CRiSIS), Toulouse, France, 6 pgs, 2009.
231. **A Characterization of Max-Min Power Allocation with Applications** (with Slawomir Stanczak, Michal Kaliszan, Marcin Wicznanowski). IEEE Int. Symposium on Information Theory, Seoul, Korea, 6 pgs, June 2009.
232. **Maximizing Throughput of Hospital Intensive Care Units with Patient Readmissions** (with C. W. Chan, V. F. Farias, G. J. Escobar). INFORMS Manufacturing and Service Operations Management (MSOM) Conference, Boston, MA, 5 pgs, June 2009.
233. **Target-Based Power Control for Queueing Systems with Applications to Packet Switches** (with B. Yolken, D. Tsamis). IEEE Global Telecom. Conference (Globecom), New Orleans, pp. 1-6, Dec. 2008.

234. **Security Investment Games in Interdependent Organizations** (with A. Miura, A. Yolken, J. Mitchell), Proc. 46<sup>th</sup> Allerton Conference on Communication, Control and Computing, pp. 1-10, Sept. 2008.
235. **Game Based Admission Control for Wireless Systems** (with B. Yolken) Internationa Wireless Internet Conference (WICON), pp. 1-7, Nov. 2008.
236. **Admission Control for Power-Controlled Wireless Networks under General Interference Functions** (with S. Stanczak, M. Kaliszan), 42nd Asilomar Conference on Signals, Systems and Computers, *Monterey, USA, Oct. 2008*.
237. **Fast Power-Up Active Link Protection in Autonomous Distributed Transmitter Power Control** (with S. Gitzenis), Proc. IEEE International Wireless Communication and Computing Conference (IWCMC), pp. 1-6, Chania, Greece, Aug. 2008.
238. **Game Based Capacity Allocation for Utility Computing Environments** (with B. Yolken), Proc. GameComm, pg. 1-10, Athens, Greece, Oct. 2008.
239. **Network-Assisted Wireless Computing** (with C. Chan and J. P. Singh), Proc. IEEE PIMRC, 2008.
240. **Target-Based Power Control for Queueing Systems with Applications to Packet Switches** (with B. Yolken and D. Tsamis), IEEE Globecom, New Orleans, 2008.
241. **Security Decision-Making amongst Interdependent Organizations** (with A. Miura, B. Yolken, J. Mitchell), Proc. IEEE Computer Security Foundations Symposium (CSF), Pittsburgh, PA, 2008
242. **Wireless Broadcasting to Diverse Users** (with C.W. Chan, S. Wee, J. Apostolopoulos). Proc. IEEE International Conference on Communications, Beijing, China, 2008.
243. **Receiver Based Optimization of Video Delivery over Wireless Links** (with C.W. Chan, S. Wee, J. Apostolopoulos), Proc. IEEE International Conference on Multimedia, 2008.
244. **Backlog Aware Scheduling for Large Buffered Crossbar Switches** (A. Dua, B. Yolken, W. Olesinski, H. Eberle, N. Gura). Proc. IEEE International Communications Conference, Beijing, China, May 2008.
245. **Target-Driven and Incentive-Aligned Power Control for Wireless Networks** (with B. Yolken), Proc. IEEE Globecom, Washington, DC, Nov. 2007.
246. **Optimal Scheduling of Media Packets with Multiple Distortion Measures** (with C.W. Chan, S. Wee, J. Apostolopoulos), Proc. IEEE International Conference on Multimedia, Beijing, China, June 2007.
247. **Power Aware Management of Packet Switches** (with L. Mastroleon, D. O'Neill, B. Yolken), Proc. IEEE Hot Interconnects, Stanford, CA, Aug. 2007.
248. **Dynamic Control Approach to Risk Mitigation in Computing Infrastructures** (with A. Miura-Ko). Proc. IEEE International Symposium on Information Assurance and Security, Manchester, UK, Aug. 2007.
249. **Backlog Aware Low Complexity Schedulers for Input Queued Packet Switches** (with A. Dua, W. Olesinski, H. Eberle, N. Gura). Proc. IEEE Hot Interconncts, pp. 39-46, Stanford, CA, 2007.
250. **Performance Tradeoffs in Mobile Computing: To Fetch or Not to Fetch?** (with A. Dua, J. Singh) Proc. ACM MobiWac, Crete, Greece, Oct. 2007.

251. **Buffer Management in Wireless Media Streaming** (with A. Dua). Proc. IEEE Globecom, Washington, DC, Dec. 2007
252. **Distributed Backlog-Driven Power Control in Wireless Networking** (with A. Dua). Proc. IEEE LanMan, Princeton, NJ, June 2007.
253. **Power Management of Packet Switches via Differentiated Delay Targets** (with B. Yolken). Proc. IEEE International Communications Conference, Beijing, China, May 2007.
254. **Power-Managed Packet Switching** (with A. Dua, B. Yolken). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007.
255. **Wireless Packet Scheduling with Soft Deadlines** (with A. Dua). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007
256. **SecureRank: An Efficient Algorithm for Prioritizing Vulnerabilities in Computing Systems** (with R. A. Miura-Ko). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007.
257. **Job Scheduling for Maximal Throughput in Autonomic Computing Systems** (with K. Ross), Proc. of International Workshop of Self-Organizing Systems (IWSOS), Passau, Germany, September 2006 (and in Lecture Notes in Computer Science 4124 Springer 2006).
258. **Capacity Maximizing Packet Scheduling Algorithms for Interconnection Networks with Finite Buffers** (with K. Cross). IEEE Global Communications Conference (GlobeCom), San Francisco, CA, 2006.
259. **Receiver-Based Optimization for Video Delivery over Wireless Links** (with C. Chan, J. Apostolopoulos, Y. Li). Proceedings of IEEE International Conference on Multimedia (ICME'06), Toronto, Ontario, July 2006.
260. **Optimal Throughput and Routing for Wireless Link Arrays** (with F. Baccelli, C. Chan). Proceedings of Infocom 2006, 12 pages, Barcelona, Spain, April 2006.
261. **Joint Power Allocation and Scheduling for Deadline Constrained Wireless Traffic** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, San Francisco, Dec. 2006.
262. **Low-Jitter Scheduling Algorithms for Deadline-Aware Packet Switches** (with A. Dua) Low-Jitter Scheduling Algorithms for Deadline-Aware Packet Switches. Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, Dec. 2006.
263. **Patching Rate Management for Controlled Service Disruption in Data Centers** (with L. Mastroleon, R. A. Miura-Ko). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, San Francisco, Dec. 2006.
264. **Packet Transmission and Content-Dependent Payout for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of IEEE Multimedia Signal Processing (MMSP'05), 10 pages, Shanghai, China, Oct. 2005.
265. **Power Management in Battery-Constrained Handheld Devices**, (with A. Markopoulou, Y. Li, C. Chan) Proceedings of IEEE International Conference on Broadband Networks (BroadNets'05), pp. 441-450, Boston, MA, Oct. 2005.

266. **Dynamic Risk Mitigation for ‘Self-Defending’ Network Security**, SecureComm, Athens, May 2005.
267. **On the Fairness-Delay Tradeoff in Wireless Packet Scheduling** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom’05), v. 5, pp. 2544-2548, St. Louis, MO, Nov. 2005.
268. **Power Control and QoS Tradeoffs for Real-Time Wireless Traffic** (with A. Dua). Proceedings of IEEE Wireless Communications and Networking Conference (WCNC’06), vol. 2, pp. 1148-1153, Las Vegas, NV, April 2006.
269. **Downlink Scheduling of Heterogeneous Traffic** (with A. Dua). Proceedings of IEEE International Conference on Communications (ICC’06), Instabul, Turkey, June 2006.
270. **Scheduling with Soft-Deadlines for Input-Queued Switches** (with A. Dua). Proc. of the 44<sup>th</sup> Allerton Conference on Communication, Computation and Control, Allerton, IL, September 2006.
271. **Cross-Layer Multi-hop Wireless Routing for Inter-vehicle Communication** (with J.P. Singh, B. Srinivasan, D. Clawin) Proceedings of IEEE International Conference on Network Testbeds and Research Infrastructures (TridentCom’06), Barcelona, Spain, March 2006.
272. **Autonomic Power Management Schemes for Internet Servers and Data Centers** (with L. Mastroleon, C. Kozyrakis, D. Economou) Proceedings of IEEE Global Communications Conference (GlobeCom’05), v. 2, pp. 943-947, St. Louis, MO, Nov. 2005.
273. **Power-Managed Block Level File Decryption in Wireless Network Computing** (with S.Gitzenis), Proceedings of IEEE International Symposium on Modeling and Optimization in Mobile, Ad-Hoc, and Wireless Networks (WiOpt’06), Boston, MA, April 2006.
274. **On the Singular Behavior of a Queuing System with Random Connectivity** (with G. Michailidis). Proceedings of the IEEE Conference on Decision and Control (CDC’05) , pp. 5348-5353, Seville, Spain, Dec. 2005
275. **A Control Formulation of the Network Security Problem via a Risk Management Approach**. Proceedings of the 2005 IEEE International Conference on Security and Privacy for Emerging Areas in Communication Networks, 3 pages, Athens, Greece, Sept. 2005.
276. **Packet Transmission and Content-Dependent Payout for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of IEEE Multimedia Signal Processing (MMSP’05), 10 pages, Shanghai, China, Oct. 2005.
277. **Power Management in Battery-Constrained Handheld Devices**, (with A. Markopoulou, Y. Li, C. Chan) Proceedings of IEEE International Conference on Broadband Networks (BroadNets’05), pp. 441-450, Boston, MA, Oct. 2005.
278. **Joint Power/Payout Control Schemes for Media Streaming over Wireless Links** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of 14<sup>th</sup> IEEE Packet Video Workshop, 12 pages, Irvine, CA, Dec. 2004.
279. **Optimizing Timeout-Based Sleep Algorithms** (with A. Markopoulou). IEEE/ACM International Symposium on Information Processing in Sensor Networks (IPSN’05), UCLA, Los Angeles, CA, April 2005.

280. **Deadline Constrained Packet Scheduling in Wireless Networks** (with A. Dua). Proceedings of IEEE Vehicular Technology Conference VTC-F'05, v. 1, pp. 196-200, Dallas, TX, Sept. 2005.
281. **On the Fairness-Delay Tradeoff in Wireless Packet Scheduling** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'05), v. 5, pp. 2544-2548, St. Louis, MO, Nov. 2005.
282. **Channel-State-Awareness Based Transmission Power Adaptation for Efficient TCP Dynamics in Wireless Networks** (with J.P. Singh, Y. Li) Proceedings of IEEE International Conference on Communications (ICC'05), v. 5, pp. 3553-3559, Seoul, Korea, March 2005.
283. **Empirical Observations on Wireless LAN Performance in Vehicular Traffic Scenarios and Link Connectivity Based Enhancements for Multihop Routing** (with J.P. Singh, B. Shrinivasan, D. Clawin, Y. Yan). Proceedings of IEEE Wireless Communications and Networking Conference (WCNC'05), v. 3, pp. 1676-1682, New Orleans, LA, March 2005
284. **Transmission Power and Payout Control for Packetized Audio Streaming over Wireless Links**, (with Y. Li) Proceedings of IEEE International Conference on Wireless Networks, Communications and Mobile Computing, Workshop on Mobility Management and Wireless Access (MobiWac'05), v. 1, pp. 55-62, Maui, Hawaii, June 2005.
285. **Autonomic Power Management Schemes for Internet Servers and Data Centers** (with L. Mastroleon, C. Kozyrakis, D. Economou) Proceedings of IEEE Global Communications Conference (GlobeCom'05), v. 2, pp. 943-947, St. Louis, MO, Nov. 2005.
286. **Media and Data Traffic Coexistence in Power-Controlled Wireless Networks** (with S. Gitzenis). Proceedings of ACM Wireless Multimedia Networking and Performance Modeling (WMuNeP'05), pp. 78-85, Montreal, Quebec, CA, Oct. 2005.
287. **Dynamic Quality of Service Control in Packet Switch Scheduling** (with K. Ross). Proceedings of IEEE International Conference on Communications (ICC'05), v. 1, pp. 396-401, Seoul, Korea, May 2005.
288. **Packet Scheduling Across Networks of Switches** (with K. Ross) Proceedings of IEEE International Conference on Networking (ICN'05), Reunion Island, April 2005.
289. **On the Singular Behavior of a Queuing System with Random Connectivity** (with G. Michailidis). Proceedings of the IEEE Conference on Decision and Control (CDC'05) , pp. 5348-5353, Seville, Spain, Dec. 2005
290. **Efficient Data Prefetching for Power-Controlled Wireless Packet Networks** (with S. Gitzenis). Proceedings of ACM/IEEE Mobile and Ubiquitous Networking Conference (MobiQuitous), pp.64-73, Boston, MA, August 22-25, 2004.
291. **Power-Controlled Packet Relays in Wireless Data Networks** (with S. Gitzenis). Proceedings of IEEE Global Communications Conference (GlobeCom), vol. 1, pp. 464-469, San Francisco, CA, December 1-5, 2003.
292. **Integrated Power Control for Circuit-Switched & Packet-Switched Traffic in Wireless Networks** (with S. Gitzenis). Proceedings of the IEEE Vehicular Technology Conference (VTC), vol. 4, pp. 2604-2609, Orlando, FL, Fall 2003.
293. **Power Control for TCP Adaptation to High-mobility Broadband Systems** (with J.P. Singh, K. Radermacher, V. Scharf-Katz). Proceedings of IEEE Vehicular Technology Conference (VTC), Los

Angeles, CA, 9 pages, Fall 2004.

- 294. **Power-Controlled Wireless Links for Media Streaming Applications** (with Y. Li). Proceedings of the IEEE Wireless Telecommunications Symposium 2004, pp. 102-111, Pomona, CA, May 14-15, 2004
- 295. **Power-Controlled Media Streaming in Interference-Limited Wireless Networks** (with Y. Li). Proceedings of the 2004 IEEE International Conference in Broadband Networks, pp. 560-568, San Jose, CA, October 25-29, 2004.
- 296. **Power-Controlled Wireless Links for Video Streaming Applications** (with Y. Li). Proceedings of the IEEE Vehicular Technology Conference (VTC), Los Angeles, CA, 9 pages, Sept. 26-29 2004.
- 297. **Local Search Scheduling Algorithms for Maximal Throughput in Packet Switches** (with K. Ross). Proceedings of the IEEE INFOCOM 2004, Hong Kong, China, 12 pages, March 7-11, 2004.
- 298. **Optimizing Quality of Service in Prioritized Packet Switch Scheduling** (with K. Ross). Proceedings of the IEEE International Conference on Communications (ICC), Paris, France, 5 pages, June 20-24, 2004.
- 299. **A Service Risk Management Approach to Capacity Protection in Optical Networks** (N. Bambos, S. Gitzenis, A. Miura, O. Gerstel, L. Paraschis). Proceedings of IEEE Workshop in Local and Metropolitan Area Networks (LANMAN'04), Mill Valley, CA, pp. 69-74, April 2004.
- 300. **Mobile-to-Base Task Migration in Wireless Computing** (S. Gitzenis, N. Bambos). Proceedings of IEEE Pervasive Communications Conference *PerCom* 2004, pp. 187-196, Orlando, FL, March 14-17, 2004.
- 301. **Distributed Power and Admission Control for Wireless Networks** (with T. Holliday, A. J. Goldsmith, P. Glynn), ISIT 2004.
- 302. **Dynamic Scheduling of Optical Data Bursts in Time-Domain Wavelength Interleaved Networks** (Kevin Ross, K. Kumaran, I. Saniee, I. Widjaja). Proceedings of the 11<sup>th</sup> IEEE Symposium on High-Performance Interconnects, Stanford, CA, pp. 108-113, 2003.
- 303. **Local/Remote Task Execution in Network Computing** (with Savvas Gitzenis). Proceedings of the International Conference on Internet Computing, IC '03, Las Vegas, NV, pp. 842-848, June 2003.
- 304. **Integrated Power Control for Circuit-Switched and Packet-Switched Traffic in Wireless Networks** (with S. Gitzenis). Proceedings of the IEEE Fall Vehicular Technology Conference VTC'03, Orlando, FL, vol. 4, pp. 2604-2609, Oct. 2003.
- 305. **Proposal and Demonstration of Link-Connectivity-Assessment-Based Enhancements to Routing in Mobile Ad-Hoc Networks** (with J. Singh, D. Clawin, B. Shrinivasan). Proceedings of the IEEE Fall Vehicular Technology Conference VTC'03, Orlando, FL, Oct. 2003.
- 306. **Power Management of Packet Switch Architectures with Speed Modes** (with D. O'Neill). Proceedings of the 41<sup>st</sup> Annual Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, Monticello, IL, Oct. 2003. (invited paper)
- 307. **Distributed Power Control for Time Varying Wireless Networks: Optimality and Convergence** (with T. Holliday, A. J. Goldsmith, P. Glynn), Proceedings of the 41<sup>st</sup> Annual Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, Monticello, IL,

Oct. 2003. (invited paper)

308. **Capacity Protection/Restoration in Tree Optical Networks for Enterprise Content Distribution** (with S. Gitzenis, O. Gertsel, L. Paraschis). Proceedings of the Workshop “Protection and Restoration: from SONET/SDH to Next-Generation Networks” of IEEE Global Communications Conference (GlobeCom '03), San Francisco, CA, 2003.
309. **Power-Controlled Packet Relays in Wireless Data Networks** (with Savvas Gitzenis). Proceedings of the IEEE Global Communications International Conference (GlobeCom '03), San Francisco, CA, vol. 1, pp. 464-469, Dec. 2003.
310. **Learning Policies for Power Control in Wireless Networks.** (with D. Vengerov, H. Berenji). Proceedings of the IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 2390-2394, Oct. 2002.
311. **Power Efficient MAC Scheme Using Channel Probing in Multirate Wireless Ad-Hoc Networks** (with JW. Kim). Proceedings of IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 2382-2386, Oct. 2002.
312. **Geometry of Scheduling Algorithms for Packet Switches** (with K. Ross). Proceedings of the Allerton Conference on Communication, Control and Computing, University of University at Urbana-Champaign, Monticello, Illinois, Oct. 2002 (invited paper)
313. **Wireless LAN Performance Under Varied Stress Conditions in Vehicular Traffic Scenarios** (with J. Singh, B. Shrinivasan, D. Clawin). Proceedings of the IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 743-747, Oct. 2002.
314. **A Practical Adaptive Algorithm for Power Control in Wireless Networks** (with D. Vengerov, H. Berenji). Proceedings of the 2002 Symposium on Autonomous Intelligent Networks and Systems, University of California at Los Angeles, Los Angeles, CA, May 2002.
315. **Power-Controlled Data Prefetching/Caching in Wireless Data Networks** (with S. Gitzenis). Proceedings of the 20th IEEE Conference on Computer Communications (INFOCOM 2002), New York, NY, vol. 3, pp. 1405-1414, June 2002.
316. **Power Control for Multirate Wireless Networks with Groupwise Serial Multiuser Detection** (with J-W. Kim). Proceedings of the 2001 IEEE Global Communications Conference (GLOBECOM 2001), San Antonio, Texas, vol. 5, pp. 3201-3205, November 2001.
317. **Multi-Modal Dynamic Multiple Access (MDMA) in Wireless Packet Networks** (with S. Kandukuri). Proceedings of the 20th IEEE Conference on Computer Communications (INFOCOM 2001), Anchorage, Alaska, vol. 1, pp. 199-208, April 2001.
318. **Power Controlled Multiple Access (PCMA) in Wireless Communication Networks** (with S. Kandukuri). Proceedings of the 19th IEEE Conference on Computer Communications (INFOCOM 2000), Tel Aviv, Israel, vol. 2, pp. 386-395, March 2000.
319. **Multi-Channel Power Control for Data Traffic in Wireless Networks** (with S. Kandukuri). Proceedings of the 1999 IEEE International Workshop on Mobile Multimedia Communications (MoMuC '99), pp. 83-92, Nov. 1999.
320. **Queueing Networks with Interacting Service Resources** (with M. Armony). Proceedings of the 37<sup>th</sup> Allerton Conference on Communication, Control and Computing. Monticello, Illinois. Sept. 1999.

(invited paper)

321. **Routing in Networks with Random Topologies** (with K. Scott). Proceedings of the 1997 International Communications Conference (ICC '97), vol. 2, pp. 862-866, 1997.
322. **Power Control and Time Division Multiple Access: The TDMA vs. CDMA Question** (with J. Rulnick). Proceedings of the 1997 IEEE Conference on Computer Communications (INFOCOM '97), Kobe, Japan, vol. 2, pp. 635-642, April 1997.
323. **Modeling Communication Links in Percolating Environments** (with N. Alexopoulos). Proceedings of the 1997 IEEE International Conference on Advances in Communications and Control (COMCON '97), pp. 249-261, 1997.
324. **Formation and Maintenance of Self-Organizing Wireless Networks** (K. Scott). Proceedings of the 31<sup>st</sup> Asilomar Conference on Signals, Systems and Computers, Monterey, CA, vol. 1, pp. 31-35, Nov. 1997.
325. **Algorithmic Speed-up of All Pairs Shortest paths Computation with Reconfigurable Optical Interconnection** (with S.M.P. Yip). Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '96), San Jose, CA, vol. 1, pp. 355-366, August 1996.
326. **Performance Evaluation of Power-Managed Mobile Communication Devices** (with J. Rulnick). Proceedings of the 1996 International Conference on Communications (ICC '96), Dallas, Texas, vol. 3, pp. 1477-1481, June 1996.
327. **Scalable Communication Schemes for Massively Parallel Processing Using Reconfigurable Optical Interconnects** (with S.M.P. Yip). Proceedings of the 1996 International Conference on Parallel & Distributed Systems (ICPADS '96), Tokyo, Japan, pp. 500-507, June 1996.
328. **Best Effort Bandwidth Reservation in High Speed LANs Using Wormhole Routing** (with B. Kwan, P-C. Hu, L. Kleinrock, Hong Xu, Joseph Touch). Proceeding of 5th IEEE International Conference on Computer Communications and Networks (ICCCN '96), Rockville, Maryland, 1996.
329. **Mobile Power Management for Maximum Battery Life in Wireless Communication Networks** (with J. Rulnick). Proceedings of the 1996 IEEE Conference on Computer Communications (INFOCOM '96), San Francisco, CA, vol. 2, pp. 443-450, March 1996.
330. **The Supercomputer Supernet (SSN): A High-Speed Electro-Optical Campus and Metropolitan Network** (with L. Kleinrock, J. Bannister, L. Bergman, J. Cong, E. Gafni, M. Gerla). Proceedings of the 1996 IEEE/ SPIE Conference on Optical Interconnects in Broadband Switching Architectures, San Jose, pp. 22-33, Feb. 1996.
331. **Optimal Server Allocation to Parallel Queues with Randomly Modulated Service Rates** (with G. Michailidis). Proceedings of the Conference on Information Sciences and Systems, Princeton University, Princeton, New Jersey, pp. 692-698, March 1996.
332. **Routing and Channel Assignment for Low Power Transmission in Personal Communication Systems** (with K. Scott). 1996 IEEE International Conference on Universal Personal Communications, vol. 2, pp. 498-502, 1996.
333. **On the Stationary Dynamics of Parallel Queues with Random Server Connectivities** (with G. Michailidis). Proceedings of the 34<sup>th</sup> IEEE Conference on Decision and Control (CDC '95), New

Orleans, LA, vol. 4, pp. 3638-3643, Dec. 1995.

334. **The Self-Organizing Wireless Adaptive Network (SWAN) Protocol for Communication Among Mobile Users** (with K. Scott). Proceedings of the 1995 IEEE Global Telecommunications Conference (GLOBECOM '95), Singapore, vol. 1, pp. 355-359, Nov. 1995.
335. **Channel Probing for Distributed Access Control in Wireless Communication Networks** (with S. Chen, D. Mitra). Proceedings of the 1995 IEEE Global Telecommunications Conference (GLOBECOM '95), Singapore, v. 1, pp.322-326, November 1995.
336. **Radio Link Admission Algorithms for Wireless Networks with Power Control and Active Link Quality Protection** (with S. Chen, G. Pottie). Proceedings of the 1995 IEEE Conference on Computer Communications (INFOCOM '95), Boston, MA, v. 1, pp. 97-104, June 1995.
337. **Entropy Methods for High Speed Communications** (with B. Prabhakar). Proceedings of the 1995 Conference on Information Sciences and Systems, The Johns Hopkins University, Baltimore, Maryland, , pp. 448-453, March 1995.
338. **Adaptive (T1/T2)-Multiplexing Transmission Schemes for Voice/Data Networks** (with A. Nguyen, M. Sherif) Proceedings of the 1995 IEEE International Symposium on Computers and Communications (ISCC '95), pp. 430-435, June 1995.
339. **A Distributed, Mobile, Wireless Infrastructure for Multimedia Applications** (with S. Chen, M. Gerla, J. Tsai). Proceedings of the 1995 WINLAB Conference, Rutgers University, East Brunswick, NJ, March 1995.
340. **Admission Control Schemes for Wireless Communication Networks with Adjustable Transmitter Powers** (with S. Chen, G. Pottie). Proceedings of the 1994 IEEE Conference on Computer Communications (INFOCOM '94), Toronto, Canada, vol. 1, pp. 21-28, May 1994.
341. **On Distributed Power Control for Radio Networks** (with S. Chen, G. Pottie). Proceedings of the 1994 International Communications Conference (ICC '94), New Orleans, LA, vol. 3, pp. 1281-1285, March 1994.
342. **On Power Control with Active Link Quality Protection in Wireless Communication Networks** (with S. Chen and G. Pottie). Proceedings of the 1994 Conference on Information Sciences and Systems, Princeton, NJ, pp. 426-467, March 1994.
343. **The Asymptotics of Traffic Processes in Large Queueing Networks** (with T. Mountford and B. Prabhakar). Proceedings of the 1994 Allerton Conference on Communication, Control and Computing, University of Illinois at Urbana-Champaign, Urbana, Illinois, pp. 563-572, Oct. 1994.
344. **Queueing Dynamics and Throughput of Ring Networks with Spatial Reuse** (with A. Nguyen). Proceedings of the 1993 Allerton Conference on Communication, Control and Computing, University of Illinois at Urbana-Champaign, Urbana, Illinois, pp. 576-587, Oct. 1993 (invited paper).
345. **On Parallel Processing of Random Graph Structured Jobs** (with S. Chen). Proceedings of the 1993 NSF Design and Manufacturing Systems Conference, University of North Carolina at Charlotte, Charlotte, NC, pp. 3-11, Jan. 1993.
346. **Power Control Based Admission Policies in Cellular Radio Networks** (with G. Pottie). Proceedings of the 1992 IEEE Global Telecommunications Conference (GLOBECOM '92), Orlando, Florida, vol. 2, pp. 863-867, Dec. 1992.

347. **Optimal Message Flow on Ring Networks with Spatial Reuse** (with A. Nguyen). Proceedings of the 31<sup>st</sup> IEEE Conference on Decision and Control (CDC '92), Tuscon, AZ, vol. 2, pp. 2360-2363, Dec. 1992.
348. **On Power Control in High Capacity Cellular Radio Networks** (with G. Pottie). Proceedings of the 1992 WINLAB Conference, Rutgers University, East Brunswick, New Jersey, pp. 239-247, April 1992.
349. **On Parallel Processing in Tandem Queueing Networks with Feedback**. Proceedings of the 1992 NSF Design and Manufacturing Systems Conference, Georgia Institute of Technology, Atlanta, GA, pp. 495-498, Jan. 1992.
350. **On Tandem Queueing Networks with Feedback** (with K. Wasserman). Proceedings of the 1991 Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, pp. 310-312, Oct. 1991.
351. **Queueing Theoretic Analysis of a General Class of Concurrent Processing Systems**. Proceedings of the 1991 NSF Design and Manufacturing Systems Conference, The University of Texas at Austin, Texas, pp. 1-5, January 1991.
352. **Maximal Throughput for Stability of a Class of Parallel Processing Systems** (with J. Walrand). Proceedings of the 29<sup>th</sup> IEEE Conference on Decision and Control (CDC '90), Honolulu, Hawaii, vol. 1, pp. 161-166, Dec. 1990.
353. **Flow Analysis of Stochastic Marked Graphs** (with F. Baccelli, J. Walrand). Proceedings of the 28<sup>th</sup> IEEE Conference on Decision and Control (CDC '89), Tampa, Florida, vol. 2, pp. 1528-1531, Dec. 1989.
354. **On Bifurcated Shortest Path Routing in Communication Networks** (with G. Stassinopoulos). Proceedings of the 1984 IEEE European Conference on Control and Communications, Patras, pp. 604-609, July 1984.

---

## REPORTS:

355. **ARQ: A Multiclass Admission Control Protocol for Heterogeneous Datacenters** (with C. Delimitrou, C. Kozyrakis). Stanford Technical Multiscale Architecture & Systems Lab. Stanford Technical Report, TR-2013-002, Jan. 2013.
356. **Adaptive Data-Aware Utility-Based Scheduling in Resource-Constrained Systems** (with D. Vengerov, L. Mastroleon, D. Murhpy). Sun Microsystems Labs Technical Report SMLI TR-2007-164, April 2007.
357. **Architecting Novel Information Services and Markets**, Detecon Management Report, 2006.
358. **Wireless Packet Scheduling with Deadlines**. Stanford Technical Report, 2005

359. **Projective Processing Schedules in Queueing Structures: Applications to Packet Scheduling in Communication Networks Switches.** Stanford Technical Report. 2002.
360. **Network Security and Infrastructure Assurance – Architectural and Performance Engineering Issues.** Stanford Technical Report, 2004
361. **Noninvasive Channel Probing for Distributed Admission Control and Channel Allocation in Wireless Networks** (with J-W. Kim, D. Mitra, S. Chen). Stanford Technical Report, 2002
362. **Power-Controlled Data Prefetching and Caching in Wireless Mobile Networking** (with S. Gitzenis). Stanford Technical Report, 2002.
363. **Queueing Networks with Random Link Topologies; Stationary Dynamics of Maximal Throughput Schedules** (with G. Michailidis). Stanford Technical Report, 2001.
364. **On Stationary Multi-Processor Queues with Threshold-Based Server Activation** (with J. Rulnick). Technical Report: NetLab 1997-09-01, Stanford University, 1997.
365. **Convergence of Departures in Tandem Networks of  $/GI/Infinity$  Queues** (with T. Mountford, B. Prabhakar). Technical Report: HPL-BRIMS-96-15, BRIMS, HP Labs, UK, 1996.
366. **Processor Migration in Distributed Queueing: An Analysis of the two Processor Case** (with K. Scott). Technical Report UCLA-ENG-95-138, School of Engineering and Applied Science, UCLA, 1995.
367. **Multicopy Parallel Processing Networks: Throughput Analysis** (with K. Scott). Technical Report UCLA-ENG-95-137, School of Engineering and Applied Science, UCLA, 1995.
368. **Critical Phenomena in Massively Parallel Processing Networks: Throughput Analysis - Initial Results** (with A. Zarkesh). Technical Report UCLA-ENG-95-141, School of Engineering and Applied Science, UCLA, 1995.
369. **A Fully Distributed Algorithm for the Creation and Maintenance of Wireless Networks of Mobile Users** (with K. Scott). Technical Report UCLA-ENG-95-112, Department of Electrical Engineering, University of California at Los Angeles, 1995
370. **Probing the Collective Behavior of Massively Parallel Processing Systems by the Renormalization Group Method** (with A. Zarkesh). Technical Report UCLA-ENG-95-140, School of Engineering and Applied Science, UCLA, 1995.
371. **Performance of a Switch State Dependent Timeout Scheme in a Wormhole Switching LAN** (with B. Kwan). UCLA Technical Report UCLA- ENG-95-121, School of Engineering and Applied Science, UCLA, 1995.
372. **Branching Processes in Rollback Synchronization of Parallel Processing Systems: Structural Results** (with J. Rulnick). Technical Report UCLA-ENG-95-136, School of Engineering and Applied Science, UCLA, 1995.
373. **Optimal Choice of Queue Length Thresholds for Processor Activation in the M/M/K Queue** (with J.M. Rulnick). Technical Report UCLA-ENG-95-134, School of Engineering and Applied Science, UCLA, 1995.

374. **Partial Worm Retraction Routing Scheme for Wormhole Switching LANs** (with B. Kwan). UCLA Technical Report UCLA ENG-94-108, November 1994.
375. **Admission Control and Routing Protocols in Stationless Multi-Hop Radio Networks** (with K. Scott). Technical Report UCLA-ENG-94-106, School of Engineering and Applied Science, U.C. Los Angeles, 1994.
376. **Optimistic (Partial Retraction) Routing Schemes for Wormhole Switching LANs** (with B. Kwan). Technical Report UCLA-ENG-94-108, School of Engineering and Applied Science, U.C. Los Angeles, 1994.
377. **Radio Link Admission Algorithms for Wireless Networks with Power Control and Active Link Quality Protection** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-94-25. School of Engineering and Applied Science, U.C. Los Angeles, 1994.
378. **Distributed Dynamic Power Control Schemes in High Capacity Wireless Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-57, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
379. **A Localized Power and Admission Control Scheme Adapted to Varying Channel Conditions for High Capacity Wireless Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-56, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
380. **Admission Control Schemes for Guaranteed Link Quality in High Capacity Cellular Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-55, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
381. **On Generalized Multiserver Queues** (with J. Walrand). Technical Report UCLA-ENG-9-24-90, School of Engineering and Applied Science, U.C. Los Angeles, 1990.
382. **Recent Results on the Asymptotic Behavior of Concurrent Processing Systems Under Stationary and Ergodic Inputs** (with J. Walrand). Technical Report UCB/ERL M88/78, Electronics Research Laboratory, U.C. Berkeley, 1988.
383. **An Alternative Approach to the Asymptotic Makespan of Serially Executed Multitasked Jobs** (with J. Walrand). Technical Report UCB/ERL M88/79, Electronics Research Laboratory, U.C. Berkeley, 1988.

---

#### TALKS:

1. **Invited Seminar – A Risk Management Approach to Cyber-Security**, Operations Research Program, School of Engineering, University of Texas at Austin, March 2017.
2. **Predicting Pediatric Surgical Case Durations** (published abstract with N. Master, D. Scheinker). INFORMS International Meeting, Hawaii, 2016.
3. **Myopic Scheduling of Jobs with Decaying Value with Applications**, (abstract with N. Master, C.W. Chan) INFORMS International Meeting, Hawaii, 2016. (content essentially the same as of talk below.)

4. **Games on Influence Networks: Equilibria, Free Riding and Dynamics** (with Z. Zhou, P. Glynn). INFORMS Annual Meeting, Nashville, Tennessee, 2016.
5. **Myopic Scheduling of Jobs with Decaying Value with Applications in Patient Scheduling** (published abstract with N. Master, C.W. Chan). INFORMS Annual Meeting, Philadelphia, PA, 2015.
6. **Distinguished Invited Seminar: Power and Throughput Issues in Next-Generation Packet Switches**, National Technical University of Athens, Greece, Feb. 2009.
7. **Distinguished Invited Seminar: Wireless Multimedia Computing**, DoCoMo Research Labs, Palo Alto, Feb. 2009.
8. **Plenary Talk: Architecting Novel Information Services and Markets**, 4<sup>th</sup> International IMS Workshop, Fraunhofer Institute for Open Communication Systems – FOKUS, Berlin, Germany, Nov. 2008.
9. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, NEC Research Labs, Tokyo, July 2008.
10. **Power Management in Networking and Computing**, Hitachi Labs, Tokyo, July 2008.
11. **Power Management of Packet Switches**, Sun Labs, Menlo Park, April 2008.
12. **Power and Throughput Issues in Next-Generation Packet Switches**, UC Irvine, February 2008.
13. Invited Distinguished Lectures: 1) **Robust Maximal Throughput Schedules in Processing Systems**, 2) **Transmitter Power Control in Wireless Multimedia Networking**, Advanced Network Sciences Institute, University of California, San Diego, August 2008.
14. **Autonomous Admission (and Power) Control in Wireless Networking**, Technical University of Berlin, December 2007.
15. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Samsung Labs, Santa Clara, December 2008.
16. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Stanford OR Seminar, November 2007.
17. **Throughput and Power Issues in Packet Switching**, University of Sydney, Australia, August 2007.
18. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Bogazici University, Istanbul, Turkey, August 2007.
19. **Outsourcing in Computing and Data Storage Services**, Deutsche Telekom Laboratories, Berlin, August 2007.
20. **Dynamic Control of Packet Switches**, Sun Labs, Menlo Park, February 2007.
21. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Deutsche Telekom Laboratories, Berlin, January 2007.
22. **Network Control Issues in Video Transmission and Distribution**, Deutsche Telekom Labs, August 2006.

23. **Power Control in Wireless Networks**, IBM Research Labs, Hawthorne, July 2006.
24. **Scheduling and Power Management Issues in Data Centers and High-Performance Computing**, Sun Labs, Menlo Park, May 2006.
25. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Athens Information Technology Institute, May 2006.
26. **Power Control in Wireless Networks**, ENS, Paris.
27. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Technical University of Athens, December 2005.
28. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Athens University of Economics and Business, November 2005.
29. **Power Control in Wireless Networks**, Sprint Labs, Burlingame, December 2004.
30. **‘Geometry’ of Packet Scheduling in Communication Switches**, MIT LIDS Seminar, November 2004.
31. **Power(Rate)/Playout-Controlled Video over Wireless**, France Telecom Research Labs, Paris, June 2004.
32. **Next-Generation Wireless Data/Media Networks**, Athens Information Technology Institute, February 2004.
33. **Projective Packet Scheduling in Switching Systems**, France Telecom Research Labs, Paris, February 2004.
34. **Next-Generation Wireless Data/Media Networks**, France Telecom Research Labs, February 2004.
35. **Geometry of Packet Scheduling Algorithms for Input-Queued Communication Switches**. INFORMS Conference, San Jose, CA, Nov. 2002.
36. **Adaptive Admission Control for Multiclass Web Servers** (with G. Paleologo). INFORMS Telecommunications Conference, Boca Raton, FL, March 2002.
37. **Power-Controlled Data (Pre)Fetching and Caching in Wireless Networking** (with S. Gitzenis). INFORMS Telecommunications Conference, FL, March 2002.
38. **On Stability of Queueing Networks with Job Deadlines** (with A. Ward). INFORMS Applied Probability Conference, New York City, NY, July 2001.
39. **Throughput Maximizing Properties of Statistical Multiplexing in Stationary Ergodic Networks with Feedback** (with C. Westphal). INFORMS Applied Probability Conference, New York City, NY, July 2001.
40. **Predictive Caching Techniques in Wireless Networks** (with A. Ward). INFORMS Telecommunications Conference, Boca Raton, FL, March 2000.
41. **Dynamic Cone Policies for Skilled Based Routing in Call Centers** (with M. Armony). INFORMS Telecommunications Conference, Boca Raton, FL, March 2000.

42. **Dynamic Cone Policies for Queueing Networks with Interfering Resources** (with M. Armony). INFORMS Conference, Philadelphia, PA, November 1999.
43. **Stability and Performance Issues of a General Heterogeneous Parallel Processing System** (with G. Michailidis, K. Wasserman, S. Akin). INFORMS Applied Probability Conference, Ulm, Germany, July 1999.
44. **Bandwidth Allocation in Wireless Networks** (with G. Michailidis). INFORMS Conference, Seattle, WA, October 1998.
45. **Stability and Large Buffer Asymptotics of a General Parallel Processing System** (with G. Michailidis, K. Wasserman, S. Akin). Conference of the Canadian Operational Research Society, Windsor, Ontario, June 1999.
46. **Dynamic Cone Policies for Queueing Networks with Interfering Resources** (with M. Armony). INFORMS Conference, Philadelphia, PA, November 1999.
47. **Power Control and QoS in Wireless Communication Networks** (with M. Armony). INFORMS Conference, Cincinnati, OH, May 1999.
48. **Scheduling Algorithms for High Performance Switches** (with S. Yip). INFORMS Conference, Seattle, WA, Oct. 1998.
49. **Bandwidth Allocation in Wireless Networks** (with G. Michailidis). INFORMS Conferences., Seattle, WA, October 1998.
50. **Queueing in Interfering Resource Environments** (with M. Armony). INFORMS Conference, Tel Aviv, Israel, June 1998.
51. **Allocation of Interdependent Resources in Queueing Systems with Applications to Wireless Networks** (with M. Armony). INFORMS Telecommunications Conference, Boca Raton, FL, March 1998.
52. **On Performance and Limitations of Ultra Large Scale Digital Interconnects with Reconfigurable Topologies** (with S. Yip). IEEE/LEOS Conference on Interconnection of High Speed Digital Systems, Santa Fe, NM, May 1997.
53. **Queueing in Random Environments** (with G. Michailidis). INFORMS Applied Probability Conference, Boston, MA, June 1997.
54. **Network Flows in Random Environments** (with G. Michailidis). INFORMS Applied Probability Conference, Atlanta, GA, June 1995.
55. **On Queueing Networks with Random Link Topologies** (with G. Michailidis). INFORMS Telecommunications Conference, Boca Raton, FL, March 1995.
56. **Power Control Issues in Wireless Networks** (with S. Chen, G. Pottie). INFORMS Telecommunications Conference, Boca Raton, FL, March 1995.
57. **The Entropy and Delay of Traffic Processes in ATM Networks** (with B. Prabhakar). 1995 IEEE Information Theory Workshop on Multiple Access and Queueing, St. Louis, Missouri, April 1995.

58. **Scalable Supercomputer Network** (with L. Kleinrock, N. Bambos, L. Bergman, J. Cong, E. Gafni, M. Gerla). 1994 IEEE Gigabit Networking Workshop, Toronto, Canada, June 1994.
59. **The Supercomputer Supernet Hardware Design** (with L. Kleinrock, N. Bambos, L. Bergman, J. Cong, E. Gafni, M. Gerla). DARPA Networking Systems Meeting, Santa Fe, NM, September 1994.
60. **Dynamics of Parallel Processing: A Non-Causal Queue** (with J. Rulnick). Annual meeting of the Society of Industrial and Applied Mathematics (SIAM), San Diego, CA, July 1994.
61. **On the Synchronization of Poisson Processes and Queueing Networks with Synchronization Nodes** (with B. Prabhakar and T. Mountford). Fifth Workshop on Stochastic Analysis and Related Fields, Silivri, Turkey, July 1994.
62. **On Flow Control in Ring Structured Local Area Networks with Spatial Reuse** (with A. Nguyen). INFORMS Telecommunications Conference, Boca Raton, Florida, March 1992.
63. **Stable Policies for Processing Systems with Precedence Constraints**. ORSA/TIMS Joint National Meeting, Anaheim, California, November 1991.
64. **On Stability of Parallel Processing Systems** (J. Walrand). SIAM Conference on Applied Probability in Science and Engineering, New Orleans, LA, March 1990.
65. **State-Dependent Queues and Queueing Networks** (with J. Walrand). International Conference on Stochastic Programming, Ann Arbor, MI, August 1989.
66. **Recent Developments in Power Control and their Impact on Next-Generation Dense Wireless Network** (invited distinguished talk). Wireless/Sensor Network Track of the 2002 IEEE Wireless Communications and Networking Conference (WCNC 2002), Orlando, Florida, March 2002.
67. **Power Control and Management Issues in Next-Generation Wireless Networks** (invited keynote talk). 2001 IEEE Symposium on Ad Hoc Wireless Networks (SAWN 2001) of the 2001 IEEE Global Communications Conference (GLOBECOM 2001), San Antonio, Texas, November 2001.
68. **Queueing Models and Problems in Wireless Network Control** (invited plenary talk). Stochastic Networks Conference, Madison, WI, June 2000.
69. **Power Controlled Multiple Access (PCMA) in Wireless Networking** (invited talk). IEEE Computer Communications Workshop, Estes Park, Colorado, November 1999.
70. **Queueing Networks with Interdependent Resources** (invited plenary talk). Oberwolfach Workshop on Applied Probability,, Mathematisches Forschungsinstitut, Oberwolfach, Germany, December 1998.
71. **Power Control Issues in Wireless Communication Networks** (invited plenary talk). Workshop on System and Control Issues in Communication Networks, National Science Foundation and Army Research Office, Airlie Center, Airlie, Virginia, August 1996.
72. **On Queueing Networks with Random Link Topologies: Bandwidth Allocation in Wireless Communication Networks** (invited plenary talk). BRIMS Workshop on Mathematical Trends in Communication Networks, Hewlett-Packard Labs, Bristol, UK, June 1996.
73. **Power Control with Active Link Protection in Wireless Networking** (invited plenary talk). EPFL Workshop on Computer Networking, EPFL, Switzerland, Summer 1996.

- 74. **Power Control Algorithms in High Performance Wireless Networks** (invited talk). Networking Seminar, IBM Watson Research Center, NY, July 2001.
- 75. **Queueing Networks in Random Environments** (invited talk). Networking Seminar, Dept. of EECS, U.C. Berkeley, March 1998.
- 76. **On Queueing Networks with Randomly Modulated Service Rates** (invited talk). Research Colloquium, Center for Operations Research, MIT, September 1997.
- 77. **High-Performance Wireless Networking**. Networking Seminar, Dept. of EECS, U.C. Berkeley, Nov. 1996

--- CONFIDENTIAL ---

Please note that most engagements identified below are governed by confidentiality agreements.

April 15, 2025

**A) EXPERT WITNESS ENGAGEMENTS OF NICHOLAS BAMBOS**

I, Nicholas Bambos, have had the following temporary expert witness engagements in the litigation cases listed below:

1. (2002-2003) Technical consultant and expert witness (*deposed*) for *Altima Communications, Inc.* in the case “Level One Communications, Inc. v. Altima Communications, Inc.” Case No. CIV. S-99-2488 GEB (GGH), US District Court, Eastern District of California
2. (2004-2005) Technical consultant and expert witness (*deposed*) for *Microsoft Corporation* in the litigation case “Lextron Systems, Inc. vs. Microsoft Corporation” Case No. CV 04-00588-VRW, US District Court, Northern District of California, San Francisco Division
3. (2004-2005) Technical consultant and expert witness (*not deposed*) for *Quova, Inc.* in the case “Quova, Inc. v. Digital Envoy, Inc.” Case No. C04 02674JL, US District Court, Northern District of California, San Jose Division.
4. (2004-2005) Technical consultant and expert witness (*deposed*) for *American Honda Motor Co., Inc.* in the litigation case “Donner, Inc. (Plaintiff and Counter-Defendant) v. American Honda Motor Co. Inc. and McDavid Plano-Acra, L.P. and The Beaumont Company d/b/a Classic Acura, (Defendants and Counter-Plaintiffs) v. Irah H. Donner (Counter-Defendant)” Civil Action File No. 5:03-CV-253 (DF), US District Court, Eastern District of Texas, Texarkana Division.
5. (2005-06) Technical consultant and expert witness (*deposed*) for *Pulse Engineering Inc. and Full Rise Electronics Co.* in the litigation case “Regal Electronics, Inc. v. Pulse Engineering, Inc., Fahrner-Miller Associates, Inc., Full Rise Electronics Co. Ltd., Max Lion Electronics, Inc., Bell Fuse, Inc., Stewart Connector Systems, Inc.” Civil Case No: 03-01296 JW, US District Court, Northern District of California, San Jose Division.
6. (2005-2008) Technical consultant and expert witness (*not deposed*) for *Realtek Semiconductor Corporation* in the litigation case “3COM Corporation v. D-Link Systems, Inc. and Realtek Semiconductor Corporation” Case No. CV-03-2177 VRW BZ, US District Court, Northern District of California, San Francisco Division.
7. (2006-2007) Technical consultant and expert witness (*deposed*) for *Tekelec, Inc.* in the litigation case “Bouygues Telecom, S.A. vs. Tekelec, Inc.” Case No. 4:05 CV 78-FL3, US District Court, Eastern District of North Carolina.
8. (2006-2008) Technical consultant and expert witness (*deposed*) for *Nissan North America, Inc.* in the litigation case “Mobile Micromedia Solutions LLC vs. Nissan North America, Inc.” Civil Action File No. 5:05-CV-230, US District Court, Eastern District of Texas, Texarkana Division.

9. (2006-2013) Technical consultant and expert witness (*not* deposed) for *J2 Global Communications, Inc.* and *Catch Curve, Inc.* in the litigation cases (all related) “J2 Global Communications Inc. v. CallWave, Inc.” Civil Action No. CVO4-706D8 DP (AJWx); and “J2 Global Communications Inc. v. Venali, Inc.” Civil Action No. CVO4-1172DDP (AJWx); and “Catch Curve, Inc. (Plaintiff/Counterclaim-Defendant) v. Venali, Inc. (Defendant/Counterclaimant)” Civil Action No. 05-4820 DDP (AJWx), US District Court, Central District of California; and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Protus IP Solutions, Inc.” Civil Action No. CV05-5610 DDP (AJWx); and “COA Network, Inc. v. J2 Global Communications, Inc.”; and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Captaris, Inc. and Open Text Corp.” Civil Action No. CV09-4150 DDP (AJWx); and J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Packetel, Inc.” Civil Action No. CV09-3240 DDP (AJWx); and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Easylink Services International Corporation” Civil Action No. CV 09-4189 DDP (AJWx); and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Venali, Inc.” CV04-1172 DDP (AJWx); and the related case in Canada “J2 Global Communications, Inc. and Catch Curve Inc. v. Protus IP Solutions Inc.” Canadian Federal Court Action No. T-139-06.
10. (2007) Technical consultant and expert witness (*not* deposed) for *Netgear* in the litigation case “Serconet, Ltd. (Plaintiff) v. Netgear, Inc. (Defendant/Third-Party Plaintiff) v. Cameo Communications, Inc., Delta Networks, Inc., Hon Hai Precision Industry Co., Ltd., Sercomm Corp., and Z-Com, Inc. (Third Party Defendants)” Civil Action No. C-06-04646 PJ, US District Court, Northern District of California, San Francisco Division.
11. (2008) Technical consultant and expert witness (*not* deposed) for *Quantum Corporation* in the litigation case “Quantum Corporation v. Riverbed Technology, Inc.” Case No. C-08-0927-WHA, US District Court, Northern District of California.
12. (2008) Technical consultant and expert witness (*deposed*) for Vero Systems (Teoco) in the litigation case “Telarix, Inc. vs. Vero Systems, Inc.” Eastern District of Virginia Case No. 1-08-cv-587 (CMH)
13. (2009-2010) Technical consultant and expert witness (*not* deposed) for ActiveIdentity Corporation in the litigation case “ActiveIdentity vs. Intercede Group PLC and Intercede Ltd.” Northern District of California, Case No. C-08-4577 VRW.
14. (2010-2011) Technical consultant and expert witness (*not* deposed) for Perfect 10 Satellite Distributing Company in the litigation case “John Mezzalingua Associates, Inc. vs. Pace Electronics, Inc., Perfect 10 Antenna Company and DS Engineering, LLC” Case Number: 0.10-cv-00064-MJD-JJG.
15. (2010) Technical consultant and expert witness (*not* deposed) for Belkin International, Inc., Cisco-Linksys L.L.C., D-Link Systems, Inc., Netgear, Inc., and SMC Networks, Inc. in the litigation case “OptimumPath, L.L.C. v. Belkin Int’l, Inc., et al” Case No. 4:09-CV-1398-CW (N.D. Cal)
16. (2011) Technical consultant and expert witness (*not* deposed) for Infoblox, Inc. in the litigation case “Infoblox v. BlueCat” Case Number CV 10-01962 JVS.
17. (2011-2012) Technical consultant and expert witness (*deposed*) for Twitter, Inc. in the litigation case “Cooper Notification, Inc., v. Twitter, Inc., Everbridge, Inc., Rave Wireless, Inc., Federal Signal Corp” Civil Action No. 09-865-LPS.
18. (2011-2014) Technical consultant and expert witness (*deposed*) for Barnes & Noble, Inc. in the litigation case “Barnes & Noble, Inc. v. LSI Corp.” No. 3:11-CV-2709-EMC, U.S.D.C. N.D. Cal.

19. (2011-2013) Technical consultant and expert witness (*not* deposed) for Google, Inc. in the litigation case “MasterObjects, Inc. v. Coogole, Inc.” Case No. CV 11-01054 PJH, Northern District of California, Oakland Division.
20. (2012-2013) Technical consultant and expert witness (*not* deposed) for Money Suite Company in the litigation case “The Money Suite Company v. Insurance Answer Center, LLC, et al,” USDC Central District of California – Case No. SACV 11-1847-AG (JPRx)
21. (2012-2013) Technical consultant and expert witness (*deposed*) for Brilliant Telecommunications / Juniper Networks in the litigation case “Alcatel Lucent Technologies USA Inc. v. Brilliant Telecommunications, Inc., et al,” Case No.: 1-11-CV-206910, Superior Court of California, County of Santa Clara.
22. (2012-2013) Technical consultant and expert witness (*deposed & testified at ITC trial*) for HTC Corporation / HTC America Inc. (Respondents) and for Google, Inc. (Intervener) in “The Matter of Certain Electronic Devices, Including Mobile Phones, Tablet Computers, and Components Thereof, ITC Inv. No. 337-TA-847.” (Complainants: Nokia Corp., Nokia Inc., Intellisync Corp.) Worked on three patents, two of which were dropped by Complainants and the Judge found no infringement on the third.
23. (2013) Technical consultant and expert witness (*not* deposed) for HTC Corporation / HTC America, Inc. in the litigation case “Smartphone Technologies vs. AT&T Mobility LLC, HTC Corporation, HTC America, Inc.” Case No. 6:11-CV-561 (E.D. Tex.).
24. (2013) Technical consultant and expert witness (*deposed*) for Aeritas, LLC in the litigation cases: 1) Aeritas, LLC v. Alaska Air Group, Inc., Case No. 11-967-SLR (D. Del.); 2) Aeritas, LLC v. Continental Airlines, Inc., Case No. 11-968-SLR (D. Del.); 3) Aeritas, LLC v. Delta Airlines, Inc., Case No. 11-969-SLR (D. Del.); 4) Aeritas, LLC v. United Airlines, Inc., Case No. 11-970-SLR (D. Del.); 5) Aeritas, LLC v. US Airways, Inc., Case No. 11-1267-SLR (D. Del.); and 6) Aeritas, LLC v. Virgin America, Inc., Case No. 12-1492-SLR (D. Del.).
25. (2013-2014) Technical consultant and expert witness (*deposed*) for IpLearn, LLC in the litigation case “IpLearn, LLC v. Blackboard Inc.” Case No. 1:11-CV-00876-RGA (D. Del.).
26. (2013-2015) Technical consultant and expert witness (*not* deposed) for Infoblox, Inc. in the litigation case “Versata Software, Inc., et al v. Infoblox, Inc.” Case No. 1:13-cv-00678-LPS (D. Del.)
27. (2014-2015) Technical consultant and expert witness (*not* deposed) for Kyocera Corp. and Kyocera Communications Inc. in the litigation case “Adaptix, Inc. v. Kyocera Communications Inc.” Case No(s) 5:14-cv-02894-PSG (and -02895-PSG), Northern District of California, and related IPR case(s).
28. (2015-2016) Technical consultant and expert witness (*deposed*) for IBM Corporation in the litigation case “Twin Peaks Software Inc. v. IBM Corporation” Case No: 14-CV-03933-JST.
29. (2016-2019) Technical consultant and expert witness (*deposed*) for Samsung in the litigation case “Huawei Technologies, Co. Ltd. Et al v. Samsung Electronics Co. Ltd. Et al.” Case No: 3:16-CV-02787 N.D. Cal.
30. (2016-2017) Technical consultant and expert witness (*not* deposed) for Bunjie, Inc. in the litigation case(s) “Bunjie, Inc. v. Acceleration Bay, LLC.” Case No: IPR 2017-01600 (U.S. Patent No. 6,910,069) and Case No: IPR 2017-01601 (U.S. Patent No. 6,829,634).

31. (2016-2018) Technical consultant and expert witness (*not deposed*) for BlackBerry in the litigation case “BlackBerry Limited v. Nokia Corporation, Nokia Solutions and Networks Oy, and Nokia of America Corporation,” Case No. 1:17-CV-00155 (RGA), United States District Court for the District of Delaware. (Also provided consultation services for related IPRs, but did not participate as a disclosed expert witness).
32. (2017) Technical consultant and expert witness (*not deposed*) for Amazon.com and Amazon Web Services, Inc. in IPR proceedings related to the litigation case “Broadcom Corporation, et al. v. Amazon.com, Inc. et al.” Case No: 8:16-cv-01774-JVS-JCG (C.D. Cal.)
33. (2017-2019) Technical consultant and expert witness (*not deposed*) for Wiko SAS in the patent litigation case “Koninklijke Philips N.V. vs. Wiko SAS,” Case No. 200.219.487/01, Court of Appeal of The Hague in the Netherlands; and technical consultant and expert witness (*not deposed*) for AsusTek Computer Inc., Asus Europe B.V. and Asus Holland B.V. in the related patent litigation case “Koninklijke Philips N.V. vs. AsusTek Computer Inc., Asus Europe B.V. and Asus Holland B.V.,” Case No. 200.221.250/01, Court of Appeal of The Hague in the Netherlands.
34. (2018-2019) Technical consultant and expert witness (*not deposed*) for Juniper Networks, Inc. in the patent litigation case “Parity Networks, LLC v. Juniper Networks, Inc.,” Case No. 6:17-CV-00495-RWS-KNM, United States District Court for the Eastern District of Texas.
35. (2019) Technical consultant and expert witness (*not deposed*) for Metaswitch Networks Ltd. and Metaswitch Networks Corporation in the patent litigation case “Sonus Networks, Inc. (d/b/a Ribbon Communications Operating Company) v. Metaswitch Networks Ltd. and Metaswitch Networks Corporation (counterclaim case),” Civil Action Nos. 2:18-CV-57-RWS, Eastern District of Texas.
36. (2020 -2021) Technical consultant and expert witness (*not deposed*) for Dell Technologies Inc., Dell Inc., and EMC Corp. in the patent litigation case “WSOU Investments, LLC v. Dell Technologies Inc. et al., No 6:20-cv-477 (W.D. Tex.).
37. (2022 - 2023) Technical consultant and expert witness (*not deposed*) for Ericsson Inc., AT&T Corp., AT&T Communications LLC., AT&T Mobility, AT&T Mobility II LLC, AT&T Services Inc. in the patent litigation case “K. Mizra LLC v. AT&T Corp. et al., 2:21-cv-00241 (E.D. Tex.) (Lead Case).
38. (2022 - 2023) Technical consultant and expert witness (*not deposed*) for Dell in the patent litigation case Corrigent Corp. v. Dell Technologies Inc. and Dell Inc., Case No. 22-cv-496 (D. Del.); and also for Cisco in the related patent litigation case Corrigent Corp. v. Cisco Systems, Inc. Case No. 22-cv-396 (W.D. Tex.)
39. (2025 - ) Technical consultant and expert witness (*not deposed*) for Amazon Web Services, Inc. in the patent litigation case Croga Innovations, Ltd. v. Amazon Web Services, Inc., Case No 1:24-cv-00398-DII (W.D.Tx), and also in the corresponding IPR case.

## **B) OTHER CONSULTING, ETC. ENGAGEMENTS OF NICHOLAS BAMBOS**

Additionally, I have had the following startup-founding, technical advising, consulting, and other engagements (with compensation, unless otherwise specified)) in the past.

1. (1994-1996) Research visitor to the AT&T Bell Labs Mathematics Center for two weeks in the summer of 1994 and two weeks in the summer of 1996.
2. (1996-present) Research proposal review panelist at the National Science Foundation.
3. (1998) Research consultant for Sun Microsystems for two weeks in 1998.
4. (1999-2001) Co-founder of “AON Networks” (optical networking technology startup).
5. (2001-2003) Member of the Technical Advisory Board of ChainCast (peer-to-peer audio streaming technology startup; now StreamAudio).
6. (2001-2003) Member of the Technical Advisory Board and consultant for Greenfield Networks (packet switching technology startup; eventually acquired by Cisco).
7. (2001-2004) Advisor at ITU Ventures, Los Angeles, CA (technology investment firm).
8. (2004-2004) Member (uncompensated) of the Advisory Board of Detecon USA (consulting branch of Deutsche Telekom).
9. (2008-2009) Co-founder of Dhaani (initially Verde) Systems (technology startup focusing on power management of desktop computers).
10. (2008) Presentation for Cyber Ark International, Inc. on “science and technology innovation in the Silicon Valley” to Chinese delegation in Oct. 2008.
11. (2010 - 2023) Member (uncompensated) of the Scientific Committee, Laboratory for Information, Networking and Communication Sciences (LINCS), Paris, France.
12. (2012) Consulting for Quinn-Emanuel on patent analysis in Oct. & Nov. of 2012.
13. (2013) Consulting for Kilpatrick-Townsend-Stockton on prior art for defendant in Sept. & Oct. 2013, regarding multimedia messaging services (MMS) technology.
14. (2015) Consulting for Quinn-Emanuel on patent analysis in Jan. to May of 2015.
15. (2015) Consulting for Kirkland-Ellis on patent analysis in April 2015, regarding network resource management technology.
16. (2015) Consulting for Winston-Strawn on prior art for defendant in the period May-September 2015.
17. (2016) Consulting for Quinn-Emanuel on patent litigation in January 2016 (engaged, but performed no work and received no compensation).
18. (2018 - 2020) Member (uncompensated) of the Scientific Advisory Board, German-Turkish Advanced Research Center (GT-ARC) for Information and Communication Technologies, Berlin, Germany.
19. (2019 - 2021) Consulting for Desmarais LLP on a patent litigation matter beginning September 2019, regarding wireless network communications.

20. (2021 - 2022) Consulting for an international law firm on a litigation matter beginning February 2021, regarding networking technology.
21. (2023) Consulting for Haynes and Boone LLP on a patent litigation matter, regarding wire-line interconnections technology.
22. (2025 January to March) Consulting for Fenwick & West on a patent litigation matter, regarding computer systems and networks. (engaged, but performed no work and received no compensation)

## ISSUED PATENTS

Nicholas Bambos

October 22, 2018

1. **US 6,711,280 B2** 2004/2001 – Method and Apparatus for Intelligent Ranging via Image Subtraction.
2. **US 7,317,730 B1** 2008/2001 – Queueing Architecture and Load Balancing for Parallel Packet Processing in Communication Networks.
3. **US 7,634,287 B1** 2009/2001 – Power Controlled Multiple Access (PCMA) in Wireless Communication Networks.
4. **US 8,320,269 B2** 2012/2009 – Scalable Delay-Power Control Algorithm for Bandwidth Sharing in Wireless Networks.
5. **US 8,423,070 B2** 2013/2004 – Method and System to Model TCP Throughput, Assess Power Control Measures, and Compensate for Fading and Path Loss, for Highly Mobile Broadband Systems.
6. **US 8,488,500 B2** 2013/2008, **US 9,454,209 B2** 2016/2008, **EP 2,272,287 B1** 2017/2008 (European counterpart with English, French and German versions), **JP 5,997,794 B2** 2016/2008 (Japanese counterpart) – Power Management of Networked Devices – (and pending application **US 15/909,619**, filed on 2018-03-01, which is a continuation of pending application **US 15/276,573**, filed on 2016-09-26, which is a continuation of application **US 13/942,586**, filed on 2013-07-15 (now patent **US 9,457,209**), which is a continuation of application **US 12/114,721**, filed on 2008-05-02 (now patent **US 8,488,500**); and pending application **IN 6987/CHENP/2010** in India).
7. **US 9,614,958 B2** 2017/2007, **EP 2,132,921 B1** 2018/2007 (European counterpart) – Predictive Computer Network Services Provisioning for Mobile Users.