AFRL-AFOSR-VA-TR-2019-0221

Mathematical Foundations of Secure Computing Clouds

Robert Nowak
UNIVERSITY OF WISCONSIN SYSTEM

07/31/2019
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
07/29/2019

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
March 2013 - September 2018

4. TITLE AND SUBTITLE
Mathematical Foundations of Secure Computing Clouds

5a. CONTRACT NUMBER
FA9550-13-1-0138

5b. GRANT NUMBER
AFOSR 051112

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Robert Nowak

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The Board of Regents for the University of Wisconsin System
21 N. Park Street
Suite 6401

8. PERFORMING ORGANIZATION REPORT NUMBER
MSN155309

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
USAF, AFRL DUNS 143574726AF
OFFICE OF SCIENTIFIC RESEARCH
875 N. RANDOLPH ST. ROOM 3112
ARLINGTON VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
None

14. ABSTRACT
This project develops the mathematical foundations of cloud infrastructures and provides a sound theoretical foundation for offensive and defensive cloud cyberwar tactics. We assert that (1) new mathematics are required to handle the big data challenges present in cloud security, and that (2) close interaction between research agendas in system security, algorithms, and big data is needed to make this mathematics truly applicable.

15. SUBJECT TERMS
cloud computing; cyber-security; active monitoring; machine learning

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
unclassified	unclassified	unclassifie d	UU	2	Robert Nowak

19b. TELEPHONE NUMBER *(Include area code)*
6086953610

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATE COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33315-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report. e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

The Cloud is becoming the battleground of 21st century cyberwarfare. Both industry and the military are moving the majority of their computing infrastructure to cloud-based platforms, because of their low operating overhead, versatility, and high-throughput computing potential. Clouds will provide strategic advantages via real-time tactical analysis for warfighters on the front lines, data mining for intelligence analytics, cheaper high-throughput computing for simulations, and a computational substrate for cyberwar operations by enabling dynamic and flexible access to mission-critical computational capacity in the form of warehouse-scale computers. With more and more computation being moved into computational clouds, strategic advantage in the 21st century is increasingly determined by the ability to operate cloud platforms securely securely — and to disrupt those of opponents. More and more, cyberwarriors will be fighting their battles on and over cloud computing systems.

It is difficult to secure any large-scale computing infrastructure, but the challenges are particularly acute for clouds. Clouds require external connectivity (e.g., to forward operating bases), making them vulnerable to attackers, who may gain the credentials needed to access some portion of the platform. Their frequently changing workloads make it hard to define normal behavior, and conversely to flag abnormal behavior. Finally, their sheer scale (a typical cloud today consists of half a million individual servers) makes gathering, analyzing, and assessing their security status a big data problem.

Three critical obstacles were tackled in this project in order to provide a more secure and robust cloud:

1. From the attacker's perspective, we studied how one reverse engineer a large, opaque infrastructure from a sparse collection of measurements. From a defensive standpoint, we derived bounds on how actions reveal information about the infrastructure. We studied how to change the external view of the infrastructure over time, turning it into a moving target that is more difficult to track.

2. We developed capabilities for monitoring the cloud efficiently to provide insight into performance and alert us to anomalous activities and

behaviors. We characterized which features are most salient for detecting attacks and anomalies. From an offensive side, we mathematically determined probing strategies to maximize information gain and minimizes possible detection.

3. We developed new algorithms for sensing and machine learning that are well adapted to cloud platforms. We developed new statistical signal processing and machine learning algorithms for just-in-time performance from a cloud infrastructure.

The details of these advances are covered in the publications listed in this report.

AFOSR Deliverables Submission Survey

Response ID:11598 Data

1.

Report Type

Final Report

Primary Contact Email

Contact email if there is a problem with the report.

rdnowak@wisc.edu

Primary Contact Phone Number

Contact phone number if there is a problem with the report

6086953610

Organization / Institution name

University of Wisconsin-Madison

Grant/Contract Title

The full title of the funded effort.

MATHEMATICAL FOUNDATIONS OF SECURE COMPUTING CLOUDS

Grant/Contract Number

AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-13-1-0138

Principal Investigator Name

The full name of the principal investigator on the grant or contract.

Robert Nowak

Program Officer

The AFOSR Program Officer currently assigned to the award

Tristan NGUYEN

Reporting Period Start Date

03/15/2013

Reporting Period End Date

09/14/2018

Abstract

This research effort aims at developing the mathematical foundations of cloud infrastructures and providing a sound theoretical foundation for offensive and defensive cloud cyberwar tactics. We assert that (1) new mathematics are required to handle the big data challenges present in cloud security, and that (2) close interaction between research agendas in system security, algorithms, and big data is needed to make this mathematics truly applicable.

Distribution Statement

This is block 12 on the SF298 form.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your SF298 form. A blank SF298 can be found here. Please do not password protect or secure the PDF The maximum file size for an SF298 is 50MB.

SF_298.pdf

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

final_report.pdf

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

Zachary Charles, Dimitris Papailiopoulos, Jordan Ellenberg. "Approximate Gradient Coding via Sparse Random Graphs." Submitted, preprint available at https://arxiv.org/abs/1711.06771.

Jordan Ellenberg, Lalit Jain. "Convergence rates for ordinal embedding." Submitted, preprint available at https://arxiv.org/abs/1904.12994.

Quentin Berthet, Jordan Ellenberg. "Detection of Planted Solutions for Flat Satisfiability Problems," Proceedings of Machine Learning Research, PMLR 89:1303-1312, 2019.

Rao, Nikhil, Parikshit Shah, Stephen Wright, and Robert Nowak. "A greedy forward-backward algorithm for atomic norm constrained minimization." In 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, pp. 5885-5889. IEEE, 2013.

Malloy, Matthew L., and Robert D. Nowak. "Near-Optimal Adaptive Compressed Sensing." IEEE Transactions on Information Theory 7, no. 60 (2014): 4001-4012.

Jamieson, Kevin, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. "lil'ucb: An optimal exploration algorithm for multi-armed bandits." In Conference on Learning Theory, pp. 423-439. 2014.

Vats, Divyanshu, and Robert D. Nowak. "A junction tree framework for undirected graphical model selection." The Journal of Machine Learning Research 15, no. 1 (2014): 147-191.

Dasarathy, Gautam, Parikshit Shah, Badri Narayan Bhaskar, and Robert D. Nowak. "Sketching sparse matrices, covariances, and graphs via tensor products." IEEE Transactions on Information Theory 61, no. 3 (2015): 1373-1388.

Jamieson, Kevin G., Lalit Jain, Chris Fernandez, Nicholas J. Glattard, and Rob Nowak. "Next: A system for real-world development, evaluation, and application of active learning." In Advances in neural information processing systems, pp. 2656-2664. 2015.

Figueiredo, Mario, and Robert Nowak. "Ordered weighted l1 regularized regression with strongly correlated covariates: Theoretical aspects." In Artificial Intelligence and Statistics, pp. 930-938. 2016.

Pimentel-Alarcón, Daniel L., Nigel Boston, and Robert D. Nowak. "A characterization of deterministic sampling patterns for low-rank matrix completion." IEEE Journal of Selected Topics in Signal Processing 10, no. 4 (2016): 623-636.

Jain, Lalit, Kevin G. Jamieson, and Rob Nowak. "Finite sample prediction and recovery bounds for ordinal embedding." In Advances In Neural Information Processing Systems, pp. 2711-2719. 2016.

Liu, J. and Wright, S. J., "An accelerated randomized Kaczmarz algorithm," Mathematics of Computation 85, pp. 153-178, 2016

Balzano, L., and Wright, S. J., "Local convergence of an algorithm for subspace identification from partial data," Foundations of Computational Mathematics, pp. 1-36, 2014.

Liu, J., Wright, S. J., Re, C., Sridhar, S., and Bittorf, V., "An asynchronous parallel stochastic coordinate descent method," Journal of Machine Learning Research 16, pp. 285-322, 2015.

Liu, J. and Wright, S. J., "Asynchronous stochastic coordinate descent: Parallelism and convergence properties", SIAM Journal on Optimization 25, pp. 351-376, 2015.

Wright, S. J., "Coordinate descent algorithms," Mathematical Programming, Series B, 151 pp. 3-34, 2015.

Rao, N., Shah, P. and Wright, S. J., "Forward-backward greedy algorithms for atomic-norm minimization," IEEE Transactions on Signal Processing 63, pp. 5798-5811, 2015.
Kennedy, R., Balzano, L., Wright, S. J., and Taylor, C. J., "Online algorithms for factorization-based structure from motion," Computer Vision and Image Understanding 150, pp. 139-152, 2016.

Lim, C.~H. and Wright, S. J., "Beyond the Birkhoff polytope: Convex relaxations for vector permutation problems." In Advances In Neural Information Processing Systems, 2014.

Sridhar, S., Wright, S. J., Re, C., Liu, J, Bittorf, V., and Zhang, C., "An approximate, efficient LP solver for LP rounding." In Advances In Neural Information Processing Systems, 2013.

WhoWas: A Platform for Measuring Web Deployments on IaaS Clouds
Liang Wang, Antonio Nappa, Juan Caballero, Thomas Ristenpart, and Aditya Akella
Internet Measurement Conference - IMC 2014

A Placement Vulnerability Study in Multi-tenant Public Clouds
Venkatanathan Varadarajan, Yinqian Zhang, Thomas Ristenpart, and Michael Swift
USENIX Security 2015

Stealing Machine Learning Models via Prediction APIs
Florian Tramer, Fan Zhang, Ari Juels, Michael Reiter, and Thomas Ristenpart
USENIX Security 2016

Side-Channel Attacks on Shared Search Indexes
Liang Wang, Paul Grubbs, Jiahui Lu, Vincent Bindschaedler, David Cash, and Thomas Ristenpart
IEEE Symposium on Security and Privacy - Oakland 2017

Beyond worst-case analysis for joins with minesweeper Hung Q. Ngo, Dung T. Nguyen, Christopher Ré, Atri Rudra:. PODS 2014: 234-245

Incremental Knowledge Base Construction Using DeepDive. Jaeho Shin, Sen Wu, Feiran Wang, Christopher De Sa, Ce Zhang, Christopher Ré:
PVLDB 8(11): 1310-1321 (2015)

Ensuring Rapid Mixing and Low Bias for Asynchronous Gibbs Sampling. Christopher De Sa, Christopher Ré, Kunle Olukotun:

ICML 2016: 1567-1576. Best paper.

Gaussian Quadrature for Kernel Features. Tri Dao, Christopher De Sa, Christopher Ré: NIPS 2017: 6107-6117

Learning Compressed Transforms with Low Displacement Rank Anna T. Thomas, Albert Gu, Tri Dao, Atri Rudra, Christopher Ré:. NeurIPS 2018: 9066-9078

New discoveries, inventions, or patent disclosures:

Do you have any discoveries, inventions, or patent disclosures to report for this period?

No

Please describe and include any notable dates

Do you plan to pursue a claim for personal or organizational intellectual property?

Changes in research objectives (if any):

Nothing to report.

Change in AFOSR Program Officer, if any:

Tristan Nguyen

Extensions granted or milestones slipped, if any:

Nothing to report

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, $K)

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Jul 29, 2019 13:35:11 Success: Email Sent to: rdnowak@wisc.edu