

REPORT DOCUMENTATION PAGE				Form Approved OMB NO. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 11-03-2022		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 28-Sep-2018 - 30-Dec-2021	
4. TITLE AND SUBTITLE Final Report: Network Sciences: Coded Computing: A Transformative Framework for Tactical Wireless Edge Computing				5a. CONTRACT NUMBER W911NF-18-1-0400	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 611102	
				5d. PROJECT NUMBER 611104	
6. AUTHORS				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Southern California Contracts & Grants 3720 S. Flower St. Los Angeles, CA 90089 -0701				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S) ARO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 73448-NS.12	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			Amir Avestimehr
					19b. TELEPHONE NUMBER +12-137-4073

RPPR Final Report

as of 18-Mar-2022

Agency Code: 21XD

Proposal Number: 73448NS

Agreement Number: W911NF-18-1-0400

INVESTIGATOR(S):

Name: Amir Salman Avestimehr

Email: avestime@usc.edu

Phone Number: +12137407326

Principal: Y

Organization: **University of Southern California**

Address: Contracts & Grants, Los Angeles, CA 900890701

Country: USA

DUNS Number: 072933393

EIN: 951642394

Report Date: 30-Mar-2022

Date Received: 11-Mar-2022

Final Report for Period Beginning 28-Sep-2018 and Ending 30-Dec-2021

Title: Network Sciences: Coded Computing: A Transformative Framework for Tactical Wireless Edge Computing

Begin Performance Period: 28-Sep-2018

End Performance Period: 30-Dec-2021

Report Term: 0-Other

Submitted By: Amir Avestimehr

Email: avestime@usc.edu

Phone: (+12) 137-407326

Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees:

STEM Participants:

Major Goals: Our goal in this project is to develop a foundation for the design of coded computing in order to tackle the key bottlenecks that arise in tactical wireless edge computing, namely resource (in particular bandwidth) constraints and high dynamics due to mobility and disruptions in service. In particular, our proposed research plan consists of the following main thrusts. Thrust 1: Minimum Bandwidth Coded Computing (MBCC). In this thrust we focus on the bandwidth bottleneck of large-scale tactical edge computing, and develop coded computing strategies that effectively inject computation redundancy to minimize the communication load. Thrust 2: Maximum Robustness Coded Computing (MRCC). In this thrust we consider the resiliency of tactical edge computing to network dynamics, and generalize our framework to account for failing, straggling, and overloaded nodes in the network. Thrust 3: Practical Demonstrations. In parallel to the first two research thrusts, we will also design experiments that can primarily be used to test out the theory and develop optimal and near optimal resource allocation algorithms for efficient distributed edge computing.

Accomplishments: (1) Minimum Bandwidth Coded Computing

(a) Coded Private ML: A Scalable Approach for Privacy-Preserving Collaborative Machine Learning. We have considered a collaborative learning scenario in which multiple data-owners wish to jointly train a logistic regression model, while keeping their individual datasets private from the other parties. We have proposed CodedPrivateML, a fully-decentralized training framework that achieves scalability and privacy-protection simultaneously. Our protocol provides strong statistical privacy guarantees against colluding parties (adversaries) with unbounded computational power, while achieving up to 16X speedup in the training time against the benchmark protocols. This work was published in NeurIPS 2019.

(b) TACC: Topology-Aware Coded Computing for Distributed Graph Processing. This article proposes a coded distributed graph processing framework to alleviate the communication bottleneck in large-scale distributed graph processing. In particular, we propose a topology-aware coded computing (TACC) algorithm that has two novel salient features: (i) a topology-aware graph allocation strategy, and (ii) a coded aggregation scheme that combines the intermediate computations for graph processing while constructing coded messages. This work was published in IEEE Transactions on Signal and Information Processing over Networks.

(c) Turbo-Aggregate: Breaking the Quadratic Aggregation Barrier in Secure Federated Learning. Federated learning is gaining significant interests as it enables model training over a large volume of data that is distributedly stored over many users, while protecting the privacy of the individual users. However, a major bottleneck in scaling federated learning to a large number of users is the overhead of secure model aggregation across many users. In

RPPR Final Report as of 18-Mar-2022

fact, the overhead of state-of-the-art protocols for secure model aggregation grows quadratically with the number of users. We propose a new scheme, named Turbo-Aggregate, that in a network with N users achieves a linear secure aggregation overhead, while tolerating up to a user dropout rate of 50%. This work is published in IEEE Journal on Selected Areas in Information Theory.

(d) "CodedReduce: A fast and robust framework for gradient aggregation in distributed learning". We focus on the commonly used synchronous Gradient Descent paradigm for large-scale distributed learning, for which there has been a growing interest to develop efficient and robust gradient aggregation strategies that overcome two key system bottlenecks: communication bandwidth and stragglers' delays. We proposed a joint communication topology design and data set allocation strategy, named CodedReduce (CR). We showed that CodedReduce achieves speedups of up to 27.2X and 7.0X, respectively over the benchmarks. This work is now published in IEEE/ACM Transactions on Networking.

(2) Maximum Robustness Coded Computing

(a) Lagrange Coded Computing (LCC): We have developed Lagrange Coded Computing (LCC) that significantly expands the set of computations for which coded computing can be applied to. This work was presented/published in AISTATS 2019.

(b) Timely-Throughput Optimal Coded Computing: In tactical edge computing platforms we face a critical bottleneck of unpredictability, unreliability, and high variability of computing resources. Meanwhile, there is significantly increasing demand for timely and event-driven services with deadline constraints. We developed a dynamic computation strategy called Lagrange Estimate-and Allocate (LEA) strategy, which achieves the optimal timely computation throughput. This work was published in ACM Mobihic 2019, and was a finalist for best paper award.

(c) Speed-Adaptive Coded Computing: We have developed a dynamic workload distribution strategy for coded computation called Slack Squeeze Coded Computation (S2C2). This work was published in Proc. of 2019 Supercomputing Conference, and is a finalist of both best paper and best student paper award.

(d) "Analog Lagrange Coded Computing." We built on our strong work from last year on Lagrange Coded Computing, and extended it to the analog domain. This work was published in IEEE Journal on Selected Areas in Information Theory.

(e) "ApproxIFER: A Model-Agnostic Approach to Resilient and Robust Prediction Serving Systems". Due to the surge of cloud-assisted AI services, the problem of designing resilient prediction serving systems that can effectively cope with stragglers/failures and minimize response delays has attracted much interest. The common approach for tackling this problem is replication which assigns the same prediction task to multiple workers. This approach, however, is very inefficient and incurs significant resource overheads. Hence, a learning-based approach known as parity model (ParM) has been recently proposed which learns models that can generate parities for a group of predictions in order to reconstruct the predictions of the slow/failed workers. While this learning-based approach is more resource-efficient than replication, it is tailored to the specific model hosted by the cloud and is particularly suitable for a small number of queries (typically less than four) and tolerating very few (mostly one) number of stragglers. Moreover, ParM does not handle Byzantine adversarial workers. We propose a different approach, named Approximate Coded Inference (ApproxIFER), that does not require training of any parity models, hence it is agnostic to the model hosted by the cloud and can be readily applied to different data domains and model architectures. Compared with earlier works, ApproxIFER can handle a general number of stragglers and scales significantly better with the number of queries. Furthermore, ApproxIFER is robust against Byzantine workers. Our extensive experiments on a large number of datasets and model architectures also show significant accuracy improvement by up to 58% over the parity model approaches. This work is accepted for presentation in AAAI'2022 (a premier conference in machine learning with acceptance rate of 15%).

Training Opportunities: There were 4 phd students and 2 postdoc working on this project. This provided an excellent opportunity for training junior phd student, as well as the postdoc for mentoring students. PI Avestimehr is also a recipient of the 2020 USC Mentoring Award.

RPPR Final Report as of 18-Mar-2022

Results Dissemination: (1) Publications at top conference venues in machine learning and information theory:

- NeurIPS conference
- AAAI conference
- AISTATS Conference
- ICML Conference
- MobiHoc Conference
- ISIT Conference

(2) Publications at top journals in information theory, communications, networking, and signal processing

- IEEE Journal on Selected Areas in Information Theory
- IEEE/ACM Transactions on Networking. Published
- IEEE Transactions on Signal Processing over Networks

(3) publication of a monograph on "Coded Computing" at Now Foundations and Trends in Information Theory.

(4) many keynote talks and presentations by PI Avestimehr and his students.

Honors and Awards: (1) PI Avestimehr has received the 2019 IEEE James L. Massey Research & Teaching Award from IEEE Information Theory Society. This award is the highest honor given by the IEEE Information Theory Society to researchers under 40 years of age, and is given annually to a single recipient.

(2) PI Avestimehr was elected as an IEEE Fellow for contributions to the analysis of communication and computation over wireless networks.

(3) Best Paper Award finalist at IEEE/ACM Mobihoc 2019 (top 6 papers – the conference itself has acceptance rate of 23%)

(4) Both Best Paper Award and Best Student Paper Award finalists at IEEE/ACM High Performance Computing, Networking, Storage, and Analysis (SC19) for paper (top 4 papers)

(5) Baidu Best Paper Award at NeurIPS -SpicyFL 2020 Workshop on Scalability, Privacy, and Security in Federated Learning.

(6) PI Avestimehr was an Invited Speaker at the 2021 National Academy of Engineering (NAE) EU-US Frontiers of Engineering Symposium on the topic of privacy-preserving edge learning.

(7) PI Avestimehr and two of PhD students working on the project received a 2021 Qualcomm Innovation Award for on-device federated learning.

(8) PI Avestimehr was a recipient of 2021 "USC Mentoring Award".

(9) PI Avestimehr is named as inaugural director of the USC-Amazon Center on Secure and Trusted Machine Learning (<https://trustedai.usc.edu>)

Protocol Activity Status:

Technology Transfer: Filed a patent on Lagrange Coded Computing.

PARTICIPANTS:

Participant Type: Graduate Student (research assistant)

Participant: Qian Yu

Person Months Worked: 8.00

Funding Support:

Project Contribution:

National Academy Member: N

Participant Type: Graduate Student (research assistant)

Participant: Chien-Shen Yang

Person Months Worked: 4.00

Funding Support:

Project Contribution:

National Academy Member: N

RPPR Final Report
as of 18-Mar-2022

Participant Type: Graduate Student (research assistant)

Participant: Saurav Prakash

Person Months Worked: 4.00

Funding Support:

Project Contribution:

National Academy Member: N

Participant Type: Graduate Student (research assistant)

Participant: Jinhyun So

Person Months Worked: 4.00

Funding Support:

Project Contribution:

National Academy Member: N

Participant Type: Postdoctoral (scholar, fellow or other postdoctoral position)

Participant: Basak Guler

Person Months Worked: 4.00

Funding Support:

Project Contribution:

National Academy Member: N

Participant Type: Postdoctoral (scholar, fellow or other postdoctoral position)

Participant: Ramy Ali

Person Months Worked: 4.00

Funding Support:

Project Contribution:

National Academy Member: N

ARTICLES:

RPPR Final Report as of 18-Mar-2022

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE International Symposium on Information Theory

Publication Identifier Type: Publication Identifier:

Volume: Issue: First Page #:

Date Submitted: 10/7/20 12:00AM Date Published:

Publication Location:

Article Title: Coded computing for distributed graph analytics

Authors: S. Prakash, R. Pedarsani, and S. Avestimehr

Keywords: Coded computing, distributed graph analytics

Abstract: There is a growing interest in development of in-network dispersed computing paradigms that leverage the computing capabilities of heterogeneous resources dispersed across the network for processing massive amount of data is collected at the edge of the network. We consider the problem of task scheduling for such networks, in a dynamic setting in which arriving computation jobs are modeled as chains, with nodes representing tasks, and edges representing precedence constraints among tasks. In our proposed model, motivated by significant communication costs in dispersed computing environments, the communication times are taken into account. More specifically, we consider a network where servers are capable of serving all task types, and sending the results of processed tasks from one server to another server results in some communication delay that makes the design of optimal scheduling policy significantly more challenging than classical queueing networks. As the main contributions of

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Transactions on Signal and Information Processing over Networks

Publication Identifier Type: DOI Publication Identifier: 10.1109/TSIPN.2020.2998223

Volume: 6 Issue: First Page #: 508

Date Submitted: 10/7/20 12:00AM Date Published:

Publication Location:

Article Title: TACC: Topology-Aware Coded Computing for Distributed Graph Processing

Authors: Basak Guler, A. Salman Avestimehr, Antonio Ortega

Keywords: Coded Computing; Graph Analytics

Abstract: There is a growing interest in development of in-network dispersed computing paradigms that leverage the computing capabilities of heterogeneous resources dispersed across the network for processing massive amount of data is collected at the edge of the network. We consider the problem of task scheduling for such networks, in a dynamic setting in which arriving computation jobs are modeled as chains, with nodes representing tasks, and edges representing precedence constraints among tasks. In our proposed model, motivated by significant communication costs in dispersed computing environments, the communication times are taken into account. More specifically, we consider a network where servers are capable of serving all task types, and sending the results of processed tasks from one server to another server results in some communication delay that makes the design of optimal scheduling policy significantly more challenging than classical queueing networks. As the main contributions of

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

RPPR Final Report as of 18-Mar-2022

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Transactions on Information Theory

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TIT.2019.2963864

Volume: 66

Issue: 3

First Page #: 1920

Date Submitted: 10/7/20 12:00AM

Date Published: 3/1/20 8:00AM

Publication Location:

Article Title: Straggler Mitigation in Distributed Matrix Multiplication; Fundamental Limits and Optimal Coding

Authors: Qian Yu, Mohammad Ali Maddah-Ali, A. Salman Avestimehr

Keywords: Coded Computing, Distributed Matrix Multiplication

Abstract: There is a growing interest in development of in-network dispersed computing paradigms that leverage the computing capabilities of heterogeneous resources dispersed across the network for processing massive amount of data is collected at the edge of the network. We consider the problem of task scheduling for such networks, in a dynamic setting in which arriving computation jobs are modeled as chains, with nodes representing tasks, and edges representing precedence constraints among tasks. In our proposed model, motivated by significant communication costs in dispersed computing environments, the communication times are taken into account. More specifically, we consider a network where servers are capable of serving all task types, and sending the results of processed tasks from one server to another server results in some communication delay that makes the design of optimal scheduling policy significantly more challenging than classical queueing networks. As the main contributions of

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE/ACM Transactions on Networking

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TNET.2019.2919553

Volume: 27

Issue: 4

First Page #: 1330

Date Submitted: 10/7/20 12:00AM

Date Published: 8/1/19 2:00PM

Publication Location:

Article Title: Communication-Aware Scheduling of Serial Tasks for Dispersed Computing

Authors: Chien-Sheng Yang, Ramtin Pedarsani, A. Salman Avestimehr

Keywords: Coded Computing; edge computing; dispersed computing

Abstract: There is a growing interest in development of in-network dispersed computing paradigms that leverage the computing capabilities of heterogeneous resources dispersed across the network for processing massive amount of data is collected at the edge of the network. We consider the problem of task scheduling for such networks, in a dynamic setting in which arriving computation jobs are modeled as chains, with nodes representing tasks, and edges representing precedence constraints among tasks. In our proposed model, motivated by significant communication costs in dispersed computing environments, the communication times are taken into account. More specifically, we consider a network where servers are capable of serving all task types, and sending the results of processed tasks from one server to another server results in some communication delay that makes the design of optimal scheduling policy significantly more challenging than classical queueing networks. As the main contributions of

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

RPPR Final Report as of 18-Mar-2022

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE/ACM Transactions on Networking

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TNET.2021.3109097

Volume: 30

Issue: 1

First Page #: 148

Date Submitted: 3/11/22 12:00AM

Date Published: 2/1/22 8:00AM

Publication Location:

Article Title: CodedReduce: A Fast and Robust Framework for Gradient Aggregation in Distributed Learning

Authors: Amirhossein Reisizadeh, Saurav Prakash, Ramtin Pedarsani, Amir Salman Avestimehr

Keywords: Coded Computing, Gradient Aggregation

Abstract: A distributed computing scenario is considered, where the computational power of a set of worker nodes is used to perform a certain computation task over a dataset that is dispersed among the workers. Lagrange coded computing (LCC), proposed by Yu et al., leverages the well-known Lagrange polynomial to perform polynomial evaluation of the dataset in such a scenario in an efficient parallel fashion while keeping the privacy of data amidst possible collusion of workers. This solution relies on quantizing the data into a finite field, so that Shamir's secret sharing, as one of its main building blocks, can be employed. Such a solution, however, is not properly scalable with the size of dataset, mainly due to computation overflows. To address such a critical issue, we propose a novel extension of LCC to the analog domain, referred to as analog LCC (ALCC). All the operations in the proposed ALCC protocol are done over the infinite fields of $\mathbb{R}$ / $\mathbb{C}$ but for practical implementat

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Journal on Selected Areas in Information Theory

Publication Identifier Type: DOI

Publication Identifier: 10.1109/JSAIT.2021.3054610

Volume: 2

Issue: 1

First Page #: 479

Date Submitted: 3/11/22 12:00AM

Date Published: 3/1/21 8:00AM

Publication Location:

Article Title: Turbo-Aggregate: Breaking the Quadratic Aggregation Barrier in Secure Federated Learning

Authors: Jinhyun So, Basak Guler, A. Salman Avestimehr

Keywords: Federated learning, secure aggregation

Abstract: A distributed computing scenario is considered, where the computational power of a set of worker nodes is used to perform a certain computation task over a dataset that is dispersed among the workers. Lagrange coded computing (LCC), proposed by Yu et al., leverages the well-known Lagrange polynomial to perform polynomial evaluation of the dataset in such a scenario in an efficient parallel fashion while keeping the privacy of data amidst possible collusion of workers. This solution relies on quantizing the data into a finite field, so that Shamir's secret sharing, as one of its main building blocks, can be employed. Such a solution, however, is not properly scalable with the size of dataset, mainly due to computation overflows. To address such a critical issue, we propose a novel extension of LCC to the analog domain, referred to as analog LCC (ALCC). All the operations in the proposed ALCC protocol are done over the infinite fields of $\mathbb{R}$ / $\mathbb{C}$ but for practical implementat

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

RPPR Final Report as of 18-Mar-2022

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Journal on Selected Areas in Information Theory

Publication Identifier Type: DOI

Publication Identifier: 10.1109/JSAIT.2021.3056377

Volume: 2

Issue: 1

First Page #: 283

Date Submitted: 3/11/22 12:00AM

Date Published: 3/1/21 4:00PM

Publication Location:

Article Title: Analog Lagrange Coded Computing

Authors: Mahdi Soleymani, Hossam Mahdavi, A. Salman Avestimehr

Keywords: coded computing

Abstract: A distributed computing scenario is considered, where the computational power of a set of worker nodes is used to perform a certain computation task over a dataset that is dispersed among the workers. Lagrange coded computing (LCC), proposed by Yu et al., leverages the well-known Lagrange polynomial to perform polynomial evaluation of the dataset in such a scenario in an efficient parallel fashion while keeping the privacy of data amidst possible collusion of workers. This solution relies on quantizing the data into a finite field, so that Shamir's secret sharing, as one of its main building blocks, can be employed. Such a solution, however, is not properly scalable with the size of dataset, mainly due to computation overflows. To address such a critical issue, we propose a novel extension of LCC to the analog domain, referred to as analog LCC (ALCC). All the operations in the proposed ALCC protocol are done over the infinite fields of $\mathbb{R}$ / $\mathbb{C}$ but for practical implementat

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published

Journal: IEEE Journal on Selected Areas in Information Theory

Publication Identifier Type: DOI

Publication Identifier: 10.1109/JSAIT.2021.3053220

Volume: 2

Issue: 1

First Page #: 441

Date Submitted: 3/11/22 12:00AM

Date Published: 3/1/21 4:00PM

Publication Location:

Article Title: CodedPrivateML: A Fast and Privacy-Preserving Framework for Distributed Machine Learning

Authors: Jinhyun So, Basak Guler, A. Salman Avestimehr

Keywords: coded computing, privacy preserving machine learning

Abstract: How to train a machine learning model while keeping the data private and secure? We present CodedPrivateML, a fast and scalable approach to this critical problem. CodedPrivateML keeps both the data and the model information-theoretically private, while allowing efficient parallelization of training across distributed workers. We characterize CodedPrivateML's privacy threshold and prove its convergence for logistic (and linear) regression. Furthermore, via extensive experiments on Amazon EC2, we demonstrate that CodedPrivateML provides significant speedup over cryptographic approaches based on multi-party computing (MPC).

Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: Y

CONFERENCE PAPERS:

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published

Conference Name: the Twentieth ACM International Symposium

Date Received: 12-Aug-2019

Conference Date: 02-Jul-2019

Date Published:

Conference Location: Catania, Italy

Paper Title: Timely-Throughput Optimal Coded Computing over Cloud Networks

Authors: C. Yang, R. Pedarsani, and S. Avestimehr

Acknowledged Federal Support: Y

RPPR Final Report
as of 18-Mar-2022

Publication Type: Conference Paper or Presentation **Publication Status:** 0-Other
Conference Name: Proceedings of Machine Learning Research
Date Received: 12-Aug-2019 Conference Date: 10-Apr-2019 Date Published:
Conference Location: Japan
Paper Title: Lagrange Coded Computing: Optimal Design for Resiliency, Security and Privacy
Authors: Q. Yu, S. Li, N. Raviv, M. Mousavi Kalan, M. Soltanolkotabi, S. Avestimehr,
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 3-Accepted
Conference Name: ACM Super Computing 2019
Date Received: 12-Aug-2019 Conference Date: 01-Oct-2019 Date Published:
Conference Location: Denver, Colorado
Paper Title: Distributed Matrix Multiplication Using Speed Adaptive Coding
Authors: K Narra, Z Lin, M Kiamari, S Avestimehr, M Annavaram
Acknowledged Federal Support: **Y**

Partners

,

I certify that the information in the report is complete and accurate:
Signature: Salman Avestimehr
Signature Date: 3/11/22 4:51PM

This page is blank.