



ARL-TR-9500 • JULY 2022



A Unified Framework for Predictive, Adaptive, Robust, and Distributed Networking for Multidomain Operations: Second-Year Report

by Ashutosh Sabharwal, Edward Knightly, Santiago Segarra, Ananthram Swami, Eric Graves, Jake Perazzone, Chirag Rao, Gunjan Verma, Paul Yu, Nishant Mehrotra, Yirong Cheng, Nate Raymondi, Siyu Liao, C. Nicolas Barati, Chia-Yi Yeh, Furqan Ahmad, Vinicius Da Silva Goncalves, Shreya Gupta, Zhecun Li, Tarence Rice, Zhambyl Shaikhanov, Arindam Chowdhury, Abhishek Kumar, Boning Li, Zhongyuan Zhao

NOTIC

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



A Unified Framework for Predictive, Adaptive, Robust, and Distributed Networking for Multidomain Operations: Second-Year Report

Ashutosh Sabharwal, Edward Knightly, Santiago Segarra, Nishant Mehrotra, Yirong Cheng, Nate Raymondi, C. Nicolas Barati, Siyu Liao, Chia-Yi Yeh, Furqan Ahmad, Vinicius Da Silva Goncalves, Shreya Gupta, Zhecun Li, Tarence Rice, Zhambyl Shaikhanov, Arindam Chowdhury, Abhishek Kumar, Boning Li, and Zhongyuan Zhao
Electrical and Computer Engineering, Rice University

Eric Graves, Jake Perazzone, Chirag Rao, Gunjan Verma, Paul Yu, and Ananthram Swami
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
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 information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 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) July 2022		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To) August 2020–September 2021	
4. TITLE AND SUBTITLE A Unified Framework for Predictive, Adaptive, Robust, and Distributed Networking for Multidomain Operations: Second-Year Report				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Ashutosh Sabharwal, Edward Knightly, Santiago Segarra, Ananthram Swami, Eric Graves, Jake Perazzone, Chirag Rao, Gunjan Verma, Paul Yu, Nishant Mehrotra, Yirong Cheng, Nate Raymondi, C. Nicolas Barati, Siyu Liao, Chia-Yi Yeh, Furqan Ahmad, Vinicius Da Silva Goncalves, Shreya Gupta, Zhecun Li, Tarence Rice, Zhambyl Shaikhanov, Arindam Chowdhury, Abhishek Kumar, and Zhongyuan Zhao				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Rice University 6100 Main St, Houston, TX 77005 DEVCOM Army Research Laboratory ATTN: FCDD-RLC-N 2800 Powder Mill Road, Adelphi, MD 20783				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-9500	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				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 ORCID ID(s): Eric Graves, 0000-0002-4453-9134; Chirag Rao, 0000-0002-3786-6436; Paul Yu, 0000-0003-1577-3914					
14. ABSTRACT This report details the research efforts undertaken in the networks track during the second year of the cooperative agreement between Rice University and the US Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory (ARL). The driving vision of this project is to develop autonomous networks in order to support distributed multi-domain operations while being robust and resilient against near-peer adversaries. These cutting-edge networks have to be energy- and spectrally-efficient, self-adapting to diverse operating conditions, and secure-by-design against the new class of learning-based threats. A highlight of Year 2 was significant collaboration between DEVCOM ARL and Rice researchers, both in terms of overall engagement and publications.					
15. SUBJECT TERMS wireless networks, artificial intelligence, graph neural networks, security, autonomy, robustness, Network and Computational Sciences					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 88	19a. NAME OF RESPONSIBLE PERSON Ananthram Swami
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) (301) 394-0492

Contents

List of Figures	v
List of Tables	ix
1. Introduction	1
2. Major Achievements in Year 2	2
3. Year 2 Events	3
4. Supporting Infrastructure	3
5. Secure Networking from Millimeter Wave to THz	5
5.1 Enabling Mobile THz networks with Leaky-Wave Antennas	5
5.2 Secure THz Angularly Dispersive Links	10
5.3 Metasurface-in-the-Middle Attacks	13
6. AI-powered Adaptive Network Management	16
6.1 ML-aided Power Allocation for Tactical MIMO	16
6.2 Graph-based Algorithm Unfolding for Energy-Aware Power Allocation in Wireless Networks	24
6.3 Power Allocation for Wireless Federated Learning using Graph Neural Networks	31
6.4 Distributed Scheduling using Graph Neural Networks	34
6.4.1 Delay-Oriented Distributed Scheduler	36
6.4.2 Distributed Link Sparsification for Massive Access	39
6.5 Learning-based Access Control	43
7. Secure Networks against AI-based Threats	55
7.1 Estimating Network Size	55
7.2 Estimating Network Flow	60
7.3 Estimating Scene Geometry	62

8. Conclusions	64
9. References	65
Appendix. Agenda of Virtual Workshop (April 2021)	72
List of Symbols, Abbreviations, and Acronyms	74
Distribution List	77

List of Figures

Fig. 1	RENEW base station installed at multiple locations in Rice University	4
Fig. 2	Drone and controller: (left) drone equipped with sensors, battery, and antennas ready for take-off; (right) controller setup where we see the massive MIMO base station on top of a building, the software-defined radio (SDR)-Pi configuration that will be connected to the drone, and a computer as a controller.....	4
Fig. 3	LWA schematic with the angular dispersion relationship. As shown in the plot on the left, higher frequencies radiate maximally out of (and into) the antenna at lower angles. This is governed by the dispersion relationship, which is determined by satisfying boundary conditions at the slot. Here, b is the separation between the parallel plates forming the LWA, and c is the speed of light.....	6
Fig. 4	LWA-LWA misaligned wideband angularly dispersive link. If the emission and reception angles are not equal, that is, the antennas are not parallel, the signal coupling into the receiver's aperture will experience a frequency-dependent coupling loss—which translates into a frequency- and bandwidth-dependent rate loss.....	6
Fig. 5	Fall in data rate with angular offset for a receiver located at three different emission angles relative to the transmitter: 25° , 45° , and 60° , with (a) describing the fall in normalized data rate with angular offset—with a wider curve representing greater robustness to misalignment and (b) showing the absolute data rates. The larger the emission angle, the greater the resilience of the client to angular mismatch, but the smaller the achievable data rate.....	7
Fig. 6	Experimental setup with THz-TDS system and custom $4 \times 4 \text{ cm}^2$ LWAs for experimental characterization of the performance of angularly dispersive THz links with misalignment	8
Fig. 7	Comparison of the experimental and model-prediction for the end-end coupling loss—quantified through normalized received power—in the wideband angularly dispersive link. The emission angle here is 60° , with the horizontal axis being the reception angle.....	8
Fig. 8	An angularly dispersive link resulting in an insecure region in the spatial domain where the transmission is not secure against an eavesdropper Eve	10
Fig. 9	Model-driven results on insecure regions when the total transmission bandwidth scales from 0.5 to 19.5 GHz for Bob at 30°	11
Fig. 10	Experimental results on insecure regions when the total transmission bandwidth scales from 3.13 to 40.7 GHz for Bob at 40°	12
Fig. 11	Overview of the MSITM attack.....	13

Fig. 12	(a) Prototype of the MSITM. Eve employs a low-cost and rapid fabrication method that involves (b) printing the designed pattern and (c) passing printed paper patterns along with metallic foil through a laminator; (d) hot-stamped design; (e) microscopic view of the fabricated metasurface 14
Fig. 13	Eve with MSITM alters the legitimate directional transmission to receive more than 30-dB signal power relative to the baseline 15
Fig. 14	MSITM attack from perspective of Bob 16
Fig. 15	Schematic diagram of the proposed UWMMSE algorithm for MIMO wireless networks. A generic intermediate layer k is detailed..... 19
Fig. 16	Operational structure of Φ . A shared 1×1 filter kernel ω generates scalar representations for all M^2 elements of H 20
Fig. 17	Performance comparison of truncated-WMMSE, WMMSE, and proposed UWMMSE on various channel distributions and antenna configurations. Box plots of $M=20$ user sum-rate achieved by Tr-WMMSE (4 iterations), WMMSE (100 iterations), and UWMMSE (4 layers) (a) on a Rayleigh channel with $T=5$ transmit antennas and $R=3$ receiver antennas, transmitting $d=2$ symbols per channel use; (b) on a Rician channel with $T=3$ transmitters and $R=3$ receivers, transmitting $d=2$ symbols; and (c) on a Geometric channel with $T=3$ transmitters and $R=5$ receivers, transmitting $d=2$ symbols. 22
Fig. 18	Robustness analysis for UWMMSE. (a) Box plots of the achieved utilities for 10000 randomly drawn Rayleigh channel state matrices for $M = 20$ users. UWMMSE trained on Rayleigh compared with UWMMSE trained on Rician along with WMMSE. (b) Box plots of the achieved utilities for 10000 randomly drawn Rician channel state matrices for $M = 20$ users. UWMMSE trained on Rician compared with UWMMSE trained on Rayleigh along with WMMSE. (c) Mean sum-rate achieved on a geometric channel as network size varies. 23
Fig. 19	An illustration of the system model and the graph representation of CSI. (a) Multicell multiuser interference-limited wireless network considered. (b) The channel gain and multiuser interference coefficients of this network can be viewed as a graph, taking the CSI matrix H as its adjacency matrix..... 25
Fig. 20	Overall schematic view of our proposed USCA architecture for energy-efficient power allocation. (a) The embedding block. This block encodes channel information through a GCN, whose output is used as the input to the first USCA block. (b) The USCA block. This block contains two GCNs that mimic the classical operations in SCA of solving Eq. P2 and finding the step size to update the power allocation. (c) The GCN. A schematic view of a generic GCN with rectified linear unit (ReLU) nonlinearities..... 27
Fig. 21	Performance comparison of the proposed GCN-based USCA with competing approaches. (a) Performance on the test set compared with all competitor methods. Every point in the curves is given by the

	average WSEE over all test channel samples under the corresponding power constraint value. (b) Histogram of the achieved WSEE values for test channels compared with model-based competitor methods. (c) Performance at progressive training milestones compared with learning-based competitor methods.	28
Fig. 22	Illustration of the proposed PDGNet for FL power allocation	33
Fig. 23	PERs and performance of the FL model achieved by different power allocation policies. The first group of comparisons is based on the weighted-sum PER of selected workers against (a) interference strength, (b) maximum power constraint value, and (c) network size; the last one (d) displays the MNIST classification error of the FL global model at the end of each FL iteration, including the ideal FL (assuming no transmission loss) performance. The proposed PDGNet outperforms the baseline power allocation strategies in all the metrics considered.	34
Fig. 24	(left) An example of a wireless multihop network and (right) its corresponding conflict graph	35
Fig. 25	Steady-state transitions with (left) an optimal MWIS scheduler in green and (right) a greedy scheduler in blue.....	36
Fig. 26	Downstream architecture of GCN-based distributed scheduler for delay-oriented scheduling	37
Fig. 27	Four-step lookahead scheduling from time slot t	37
Fig. 28	Approximation ratios of the queue lengths of the GCN-based distributed scheduler relative to the local greedy scheduler under various conflict graph configurations. Smaller values refer to better delay. Smaller ratios indicate better performance relative to baseline.	39
Fig. 29	Wireless multihop networks with neighborhoods of different sizes: (left) dense network with large neighborhoods and (right) sparse network with small neighborhoods. The overhead, collision, and starvation of scheduling increases with the neighborhood size, while the relays and thus overhearing increase in sparser networks.	40
Fig. 30	Architecture of GCN-based distributed sparse scheduler. (a) Select cut-off utility based on the empirical cumulative density function of utility, (b) parameters Z generated by a featureless GCN based on conflict graph, and (c) sparse graph and topology-aware utility created by the parameterized functions, then the distributed scheduler generates schedules.....	40
Fig. 31	The performance of sparse schedulers, GCN and statistical method (Stat), with respect to a reference dense scheduler, given an identical input network state. (a) Approximation ratio of total utility, (b) Retention ratio of nodes, (c) Retention ratio of average degree of conflict graph, and (d) Retention ratio of P2P message complexity of scheduling contention. For (c) and (d), the smaller is better.	42

Fig. 32	Retention ratios for (a) average degree and (b) P2P message complexity of the sparse schedulers w.r.t. the vanilla local greedy scheduler vs. the average degree of the original conflict graph. The plots correspond to 500 scheduling instances with ER graphs, cut-off quantile of 0.95, and 300 time slots. Smaller retention ratios are better.	43
Fig. 33	Schematic view of our proposed deep RL agent. Node 0 relies on a rainbow DQN to select the action for the next time interval based on historical local observations.....	48
Fig. 34	Performance evaluation of the proposed RL method (RDQN) for BEB backoff mechanism. Box plots for fairness utility over 500 episodes of 50 time steps for the five methods under consideration. A Markov process with $p \in \{0.75, 1\}$ is considered for the evolution of the backoff parameters of other nodes	50
Fig. 35	Performance evaluation of the proposed RL method (RDQN) for MILD (left) and EIED (right) backoff mechanisms	50
Fig. 36	Performance evaluation with varying backoff mechanisms	51
Fig. 37	Fairness utility obtained in the first 20 steps of an episode when other nodes follow the simple policy that punishes repetition	52
Fig. 38	Performance evaluation with other nodes following a local policy. The mean utility for the three approaches are $\{0.9829, 0.9668, 0.9797\}$ and $\{0.7435, 0.7435, 0.7197\}$ for u_{fair} and u_{thr} , respectively.	53
Fig. 39	Comparison of utilities obtained by different models when the network size is varying. (left) Other nodes follow a Markov process with $p = 1$. (right) Utilities obtained by different models when evaluated on $N = 5$	54
Fig. 40	The eavesdropper (black node) passively monitoring an adjacent wireless network	56
Fig. 41	Wireless diarization pipeline. In this work, we skipped the signal detection part since our software framework dictates scheduled, non-overlapping, transmissions.....	57
Fig. 42	Locations of transmit nodes and the receiver node. On the left, the position of transmitters and receiver in Rice Stadium; on the right, the transmit nodes in the 2-D array. Transmitters used in this experiment are shown by the white circles.	58
Fig. 43	Architecture of NESTNet, the proposed Siamese neural network	58
Fig. 44	The effect of NESTNet on data. (left) 2-D projection of raw data. (right) 2-D projection of the embeddings of raw data generated by NESTNet.....	58
Fig. 45	Estimated network size vs. true network size. The bars' height is the average estimated network size from different runs. The error bars represent the 95% confidence interval of the estimate.....	59
Fig. 46	An illustration of the observed network and the eavesdropper setup .	60

Fig. 47	Algorithm's MSE under different observation durations in an IEEE 802.11 network with five independent flows (log-log plot on the right)	61
Fig. 48	Example network topology for estimating scene geometry. Ongoing data transmissions between allied nodes is decoded and repurposed for imaging the tanks.	63
Fig. 49	Simulation environment (left) and simulated image reconstruction of a tank (right)	63

List of Tables

Table 1	Average WSEE values of power allocations on test sets with different numbers of users associated with different numbers of base station (BS). The best performance for every scenario is bolded and the true optimum (when computationally feasible) is underlined.....	29
Table 2	Average WSEE values of power allocations on test sets with different channel distributions as a result of different PL propagation models.	30

1. Introduction

The ultimate goal of this project is to design a new generation of autonomous wireless networks to support distributed multidomain operations (MDO) while being robust and resilient against near-peer adversaries. The current report summarizes the advances attained during Year 2 of this project and lays down clear research directions for the upcoming years.

Autonomous networks are critical to enable the Army to conduct MDO maneuvers faster than the adversary. Moreover, these will also facilitate the convergence and exploitation of cross-domain synergies throughout the battlespace (land, air, sea, space, cyberspace, spectrum, information environment) and across various warfighting functions (intelligence, mission command, fires, sustainment, protection, movement and maneuver).

To develop the new generation of networks, the project has been organized into three interrelated research thrusts.

- **Secure Networking from Millimeter Wave to THz:** The innovation of materials enables and will continue to enable access to wide swaths of spectrum, ranging from low GHz to THz. However, the new capabilities bring new challenges, as increasingly communications are directional in the higher bands. In this thrust, we are developing novel methods to leverage special properties of emerging devices that can transmit over GHz to THz frequency bands, supported by multiple in-lab and deployed testbeds at Rice University.
- **Artificial Intelligence (AI)-powered Adaptive Network Management:** Recent advances in communication capabilities provide a rich array of tuning knobs for network management and optimization. However, network operations are scenario-dependent and are tightly tied to channel and traffic conditions and device capabilities. In this thrust, we harness AI to learn about the operating environment and steer the network's operating point to the best available, even in harsh conditions.
- **Secure Networks against AI-based Threats:** Learning-based attacks have opened a new threat vector, where the passive adversaries can learn from listening to ongoing transmissions and potentially extract crucial information even from encrypted networks. We are developing innovations that make it harder for adversaries to extract information, with innovations ranging from new antennas to novel higher-layer methods.

2. Major Achievements in Year 2

A detailed description of the advances attained in the three thrusts is given in Sections 5, 6, and 7. As a summarized preview, here we provide a list of the major achievements for each thrust during Year 2.

Secure Networking from Millimeter Wave to THz:

- Realized a method to secure THz links that exhibit angular dispersion, demonstrating that wide bandwidth and security can be simultaneously achieved.
- Experimentally demonstrated threats from a “metasurface-in-the-middle” attack in which a strong adversary attempts to intercept highly directional communication without being discovered.
- Demonstrated a foundation for robust mobile highly directional links spanning 100 GHz to 1 THz based on a multifaceted leaky-wave antenna (LWA) architecture.

AI-powered Adaptive Network Management:

- Combined the expressive power of graph neural networks and the flexibility of algorithm unfolding to devise learnable power allocation mechanisms for multiple-input multiple-output (MIMO) systems and energy-constrained scenarios.
- Proposed a provably near-optimal power allocation strategy for federated learning that incorporates specific constraints of the learning system.
- Developed a trainable and scalable link scheduling method based on graph machine learning that augments existing approaches and attains state-of-the-art performance.
- Devised a learning scheme for access control that can switch between and tune existing protocols to maximize a pre-specified utility.

Secure Networks against AI-based Threats:

- Developed a framework for unsupervised “wireless diarization,” which can identify the number of wireless nodes without prior training.
- Developed a framework to estimate link rates from a single eavesdropper in the presence of erroneous measurements.
- Developed a method to passively use the wireless energy in the air to estimate the scene geography.

3. Year 2 Events

Virtual Workshop, 04/19/2021–04/21/2021. Agenda in Appendix A, online.

Visit by Dr Baker, 09/30/2021. Rice University, in person. The visit included visits to multiple Rice networks labs, including demos.

Biweekly meetings for three parallel thrusts, online.

Deep dive student-US Army Combat Capabilities Development Command Army Research Laboratory meetings, online.

4. Supporting Infrastructure

Developing our new generation of autonomous networks necessitates a synergy of theoretical, computational, and empirical developments where the proposed algorithms are tested on real networks and, in turn, newly developed devices motivate novel analytical problems. For example, a novel power allocation algorithm should be tested on a real wireless network. On the other hand, novel antenna technologies can reveal unexplored security-related challenges. In this section, we present infrastructure that supports the more applied directions within the project.

The Networks team is leveraging complementary large-scale research projects. The most notable among them is the RENEW project (<https://renew-wireless.org>), which is developing the world’s first fully programmable and observable wireless radio network. With RENEW, the wireless research and development community will be able to test diverse ideas and concepts, ranging from low-level hardware all the way to novel applications. The project is truly pioneering the availability of wireless infrastructure for research by providing the first: i) Programmable wide-band radios, including 5G bands of 2.5 and 3.5 GHz; ii) large-scale MIMO, including massive MIMO; iii) novel PHY and network stacks, including 5G-like and WiFi-like protocol stacks; iv) system that is observable at all layers of the protocol stack; and v) use of distributed computational resources throughout the infrastructure. The RENEW project is funded by the NSF PAWR program, and the RENEW base stations (see Fig. 1) are deployed both at Rice University and the University of Utah.

The networks team is also leveraging the ASTRO project (<http://astro.rice.edu>). This project enables a network of small and inexpensive drones to autonomously fly in order to realize mission objectives. In ASTRO, the key mission is environmental sensing in order to rapidly monitor and track plumes from extreme events such as a chemical plant fire. Here, we leverage ASTRO’s capabilities to

provide highly mobile communication relays to Soldiers. This will ensure that Soldiers always remain covered with sufficient communication resources. Figure 2 shows a close-up photo of a drone and its associated controller.



Fig. 1 RENEW base station installed at multiple locations in Rice University



Fig. 2 Drone and controller: (left) drone equipped with sensors, battery, and antennas ready for take-off; (right) controller setup where we see the massive MIMO base station on top of a building, the software-defined radio (SDR)-Pi configuration that will be connected to the drone, and a computer as a controller

Finally, for the computational needs of the project—especially the training of AI agents—two state-of-the-art computing servers are being used. Each server contains two Intel Xeon Platinum 8260 Processors with 24-Core 2.4 GHz and 36 MB Cache, twelve 16-GB PC4-21300 2666 MHz DDR4, over 5 TB of SSD storage, over 15 TB of hard disk storage, and eight PNY NVIDIA GeForce RTX 2080 Ti Blower Edition 11GB GDDR6. Having eight GPUs in each server facilitates the parallel training of multiple AI agents and enables the simultaneous development of many of the projects described in this report.

5. Secure Networking from Millimeter Wave to THz

The large swathes of bandwidth available in the THz spectrum from 0.1 to 1 THz have the potential to realize the next order of magnitude in data rate and user densities in future wireless networks. Such networks have unprecedented capabilities for security, yet also offer new threats from adversaries. This project thrust spans from performance to security for networks spanning from 100 GHz to 1 THz.

5.1 Enabling Mobile THz networks with Leaky-Wave Antennas

LWAs provide a key mechanism for THz-scale networking that is today not available with other methods: dynamic beam steering across the entire range of 0.1 to 1 THz. In other words, LWAs are today's sole technology for mobile THz beam steering. LWAs operate on the principle of angular dispersion, where higher frequencies radiate with maximum power towards smaller angles, as governed by a simple monotonic relationship, as shown in Fig. 3.

This coupling between frequency and angle has enabled dynamic beam steering through tuning the carrier frequency¹ and efficient path discovery.² Moreover, they have experimentally been shown to support links with wide bandwidth, in the order of several 10s of GHz, and data rates^{3,4} up to 50 Gb/s. To attain these benefits at both the transmitter and receiver in a THz link, it is imperative that both nodes be equipped with LWAs, as shown in Fig. 4. However, in such a link, because the frequency-angle coupling impacts both transmission and reception, a given frequency is maximally coupled both out of the transmitter's LWA and into the receiver's LWA only if the emission and reception angles are equal, that is, if the antennas are parallel.

In this project thrust, we are the first to evaluate the performance of LWA-LWA angularly dispersive THz links with misalignment. To do so, we first introduce a simple numerical model for computing the signal-to-noise ratio (SNR) and achievable data rate in a wideband LWA-LWA link, for the purpose of analyzing the link's robustness to angular mismatch. We achieve this by fusing closed-form approximations for the single-antenna far field radiation pattern of the transmitter and receiver, which treat the LWA aperture as a diffracting slot.⁵ We design the model to capture how the fall in data rate arising from coupling loss due to misalignment depends on not only the emission and reception angles but is also both frequency and bandwidth dependent through nonlinear mechanisms. Thus, the model characterizes misalignment in LWA links in a way that is fundamentally different from that in conventional directional links.




Fig. 3 LWA schematic with the angular dispersion relationship. As shown in the plot on the left, higher frequencies radiate maximally out of (and into) the antenna at lower angles. This is governed by the dispersion relationship, which is determined by satisfying boundary conditions at the slot. Here, b is the separation between the parallel plates forming the LWA, and c is the speed of light.




Fig. 4 LWA-LWA misaligned wideband angularly dispersive link. If the emission and reception angles are not equal, that is, the antennas are not parallel, the signal coupling into the receiver's aperture will experience a frequency-dependent coupling loss—which translates into a frequency- and bandwidth-dependent rate loss.

Using our model, we characterize the performance of the wideband LWA-LWA link with misalignment, in terms of robustness to misalignment, that is, how slowly does the data rate fall with mismatch, and in terms of the absolute data rate, as shown in Fig. 5. Here, we demonstrate that a receiver located at a large emission angle relative to the transmitter will be more robust to misalignment, as compared to when it is located at a small emission angle, as shown by the support of the curves increasing with emission angle in Fig. 5 (a). That is, the SNR achievable over the link changes much more slowly with higher emission angles, leading to increased robustness in the large emission angle regime. On the other hand, higher data rates are feasible at small emission angles, owing to a wider bandwidth accessible in this regime, even with significant mismatch, as compared to large emission angles, as

can be seen in Fig. 5 (b). The key reason behind this tradeoff between robustness and high data rates is the inverse and nonlinear nature of the LWA's fundamental coupling between angle and both frequency and bandwidth.




Fig. 5 Fall in data rate with angular offset for a receiver located at three different emission angles relative to the transmitter: 25° , 45° , and 60° , with (a) describing the fall in normalized data rate with angular offset—with a wider curve representing greater robustness to misalignment and (b) showing the absolute data rates. The larger the emission angle, the greater the resilience of the client to angular mismatch, but the smaller the achievable data rate.

We complement our numerical analysis by performing the first experimental characterization of a misaligned wideband angularly dispersive link, using a custom testbed consisting of LWAs, with a THz time-domain spectroscopy (TDS) system, as shown in Fig. 6. This system enables THz wideband measurements by generating a wideband THz pulse and recording time-averaged time-domain signals at the receiver.

We compare the experimental and model-based link in terms of the end-to-end coupling loss, that is, the attenuation imposed on a given frequency depending on the emission and reception angles. This is quantified through the received power, normalized to a maximum value of unity for each emission angle. We consider an emission angle of 60° —corresponding to relatively high robustness as discussed above. For this emission angle, we rotate the receiver clockwise and counterclockwise from being parallel to the transmitter so that the reception angle varies within a sector. For each angular pair, we log the magnitude spectrum of the received signal, comparing the experimental result with the best-fit model prediction, with the results shown in the form of heatmap comparisons in Fig. 7.




Fig. 6 Experimental setup with THz-TDS system and custom 4×4 cm² LWAs for experimental characterization of the performance of angularly dispersive THz links with misalignment




Fig. 7 Comparison of the experimental and model-prediction for the end-end coupling loss—quantified through normalized received power—in the wideband angularly dispersive link. The emission angle here is 60° , with the horizontal axis being the reception angle.

Observe the comparison in end-to-end coupling loss between the model and experiment in Fig. 7, with the model prediction on the left and the experimental result on the right. For either heatmap, the horizontal axis is reception angle—representing the relative orientation of the receiver—and the vertical axis represents frequency, where a heatmap value of 1 represents no coupling loss on the corresponding frequency. In Fig. 7 (left), the end-to-end coupling loss is minimized for the model-predicted peak frequency and as the reception angle increases, higher frequencies undergo greater attenuation—according to the dispersion relationship—and the spectrum shifts towards lower frequencies. When comparing Fig. 7 (right) to Fig. 7 (left), however, we note several differences—the first of which is that the experimental spectrum is shifted towards lower angles as compared to the model prediction. This indicates that the antennas being parallel surprisingly does not maximize data rate in the experimental link, with the reason being that by steering towards lower reception angles, the receiver can couple in a wider signal bandwidth, at the cost of SNR due to misalignment, but with the net result being a higher data rate.

Secondly, note that the experimental spectrum covers a much wider range of reception angles as compared to the model prediction. This is the result of several idealizations made in the model, such as the plates being infinitesimally thin, and perfect coupling at the beginning of the slot. The indication, however, is clear—the experimental link is far more robust to mismatch as compared to the model prediction.

Although not mentioned here, we have evaluated the performance of increasing the rate and coverage of the wideband angularly dispersive link in the presence of misalignment through both deterministic means—such as using multiple angularly separated LWA faces per sector as a means to provide additional options in emission and reception angle, and optimizing the physical geometry of the device for robustness—and through opportunistic means—such as non-line-of-sight paths, which may offer a different and possibly better angular match as compared to the line of sight path.

In this work, both the numerical and experimental characterization of the link’s rate and coverage performance in the presence of misalignment does not consider transmissions with modulated data streams. Prior work³ has shown that incorporating angularly dispersive devices into an end-to-end communication link can yield performance insights that are otherwise not apparent. Therefore, in the future, we intend to investigate the robustness to misalignment of the wideband angularly dispersive link when the transmissions are composed of modulated data streams.

5.2 Secure THz Angularly Dispersive Links

While angular dispersion provides new opportunities for THz communications as described above, it also introduces new security threats as increased bandwidth necessarily yields a larger signal footprint in the spatial domain and potentially benefits an eavesdropper. Since higher directivity, or a narrower signal footprint, has been shown to be more resilient against eavesdropping,⁶ an increasingly larger signal footprint of an angularly dispersive link creates security concerns as bandwidth (and data rate) increases: will THz links be fast (wideband) or secure (small footprint), but not both?

We are the first to study secure transmission strategies on angularly dispersive links to address the challenge of securing wideband transmissions with angular dispersion (Fig. 8). In particular, we propose to frequency channelize the wideband transmission and perform coding across frequency channels to secure the angularly dispersive link. The idea is to exploit the fact that for angularly dispersive links, Eve (the eavesdropper) only intercepts a subset of frequency channels well, when she is at a different angular location from Bob (the receiver).⁷ Frequency channelization and cross-channel coding together force Eve to obtain high enough signal strength across the entire transmission band to decode the message Alice (the transmitter) sends, and thus limit Eve's chance of interception.




Fig. 8 An angularly dispersive link resulting in an insecure region in the spatial domain where the transmission is not secure against an eavesdropper Eve

To demonstrate our idea, we establish angularly dispersive THz links using a parallel-plate LWA and specify a cross-channel coding strategy termed secure coding for angularly dispersive links (SCADL), which is adapted from Liu et al.⁸ and is based on information theory. The idea is to make SCADL achieve the information theoretic limit so that employing SCADL results in the best security performance among all possible cross-channel coding strategies. As a baseline, we specify independently coded baseline (ICB), which requires Alice to code independently per frequency channel. To make SCADL and ICB comparable, we let ICB have a similar construction as SCADL but without cross-channel coding. To quantify the secrecy of the transmission, we define insecure region as the set of locations where Eve can obtain a significant amount of information about the transmitted message. A larger insecure region is thus less secure. We employ a mix of theoretical analysis, model-driven evaluations, and over-the-air experiments to demonstrate the security of angularly dispersive links with the proposed coding strategy.

First, we show that Alice can use encoding of her data across different sub-bands to dramatically reduce the security disadvantage due to a widening signal footprint for angularly dispersive links. As shown in Fig. 9, when the proposed cross-coding strategy, SCADL, is employed, the insecure region remains almost identical under a bandwidth increment of almost 20 GHz, while the insecure region expands nonuniformly across angles when the baseline strategy ICB is employed. The fact that the insecure region can remain almost identical when employing SCADL is striking, as it shows that achieving a higher secrecy rate with little sacrifice in secrecy outage is possible for angularly dispersive links. Yet, we also point out that employing cross-channel coding is key to secure the angularly dispersive links, as the independently coded baseline strategy ICB results in a significant secrecy degradation.




Fig. 9 Model-driven results on insecure regions when the total transmission bandwidth scales from 0.5 to 19.5 GHz for Bob at 30°

We further perform over-the-air THz transmissions and demonstrate that the proposed cross-channel secure coding SCADL successfully manages the widening spatial footprint as observed in the model-driven approach, despite the beam asymmetry and irregularities in the measured LWA radiation pattern. As shown in Fig. 10, when SCADL is employed, the insecure region remains almost identical with increasing bandwidth. In comparison, the baseline ICB results in a significant and uneven insecure region expansion as the bandwidth increases. In particular, the insecure area expands more in the smaller angles than in the larger angles, as a result of the measured beam radiating asymmetrically stronger towards the smaller angle. In comparison, SCADL exploits the fact that a strong side-lobe does not happen at the same angle for all frequency channels, and thus preserves a relatively consistent insecure region. From the experimental results, we find that employing SCADL for angularly dispersive links can effectively reduce the disadvantage from the widening signal footprint, even under practical beam irregularities.




Fig. 10 Experimental results on insecure regions when the total transmission bandwidth scales from 3.13 to 40.7 GHz for Bob at 40°

Our results so far are based on a line-of-sight (LoS) scenario. In the future, we plan to investigate angularly dispersive link security also when non-line-of-sight (NLoS) paths exist, and thus demonstrate link secrecy in more complicated environments. In addition to realizing secure THz networks, we are also interested in generalizing our study to multiuser scenarios. In that scenario, since different users have a different SNR profile across the frequency channels, solutions to allocate or share frequency channels are required.

5.3 Metasurface-in-the-Middle Attacks

Metasurfaces are artificially engineered structures that exhibit customizable electromagnetic properties, even beyond what is available in nature.⁹ They have been used to enhance wireless communication performance in numerous ways, for example, relaying signals via transparent metasurfaces embedded in windows,¹⁰ mitigating antenna polarization mismatch with wall-integrated metasurfaces,¹¹ and extending signal coverage through metamorphic surfaces on curtains and blinds.¹² With the advancement towards 6G networks, metasurfaces are envisioned to become an even more ubiquitous part of the environment,¹³ providing highly controllable steering capability of high data rate (TB/sec), high-directional, and high-frequency (0.1 to 1 THz) wireless links.

In this project, we, for the first time, consider that the adversary employs a metasurface and demonstrate a new acute vulnerability to a diffractive metasurface-in-the-middle (MSITM) attack. In this attack, the adversary Eve places a metasurface in the path of a directive transmission between Alice and Bob as shown in Fig. 11 and targets to redirect a portion of the signal towards herself, without being detected. Specifically, she alters the radiation pattern between Alice and Bob to simultaneously i) establish a diffracted link directed towards Eve so that Eve can be located away from Alice and Bob and ii) maintain Alice and Bob's legitimate communication link so that Eve can avoid detection. To realize such an attack, we show how Eve designs a metasurface that induces abrupt phase changes at the interface of the metasurface to controllably diffract directional links and establish furtive eavesdropping links.




Fig. 11 Overview of the MSITM attack

We show that the adversary could carry out the MSITM attack in a broad set of attack scenarios and across different contexts. For example, it could be a wireless local area network (WLAN) scenario where Alice is an access point (AP) and Bob is a laptop or mobile client. Hiding the surface between THz transmissive objects, for example, paper posters or plastic souvenirs, the attacker could be eavesdropping on sensitive data such as emails, financial information, and other personal information. Similarly, the attacker could be eavesdropping on a point-to-point (P2P) rooftop backhaul link, either a person standing in between buildings and positioning the metasurface accordingly to intercept transmission or an adversary drone carrying the metasurface and assisting the attack.

Moreover, we explore the design space of the attacker and prototype the malicious metasurface. In particular, we show that Eve has a wide range of design possibilities and investigate the key principles in realizing the MSITM attack. Specifically, she first engineers subwavelength scale metallic resonators referred to as meta-atoms.¹⁴ They enable Eve to control the amplitude and phase of the transmission according to geometrical configurations and orientations of the meta-atoms. Eve then strategically arranges a group of meta-atoms to form a supercell and induce artificial and position-dependent phase discontinuity at the surface interface. By doing so, Eve can controllably scatter Alice's and Bob's transmission and generate her targeted diffraction radiation pattern to herself and Bob. We also demonstrate how the adversary exploits recent inexpensive and rapid metasurface fabrication methods such as the hot-stamping technique¹⁵ and produces malicious metasurface under 5 min at the cost of several cents as shown in Fig. 12.




Fig. 12 (a) Prototype of the MSITM. Eve employs a low-cost and rapid fabrication method that involves (b) printing the designed pattern and (c) passing printed paper patterns along with metallic foil through a laminator; (d) hot-stamped design; (e) microscopic view of the fabricated metasurface

We also perform extensive over-the-air experiments with a THz system to investigate the attack. We demonstrate that Eve is capable of diffracting sufficient signal power towards herself (more than 30-dB power relative to not having the

MSITM shown in Fig. 13) to diminish secrecy capacity by as much as 80%. In addition, we show that she can acquire significant power across a wide range of communication bands and be robust to mispositioning. She can be away from her ideal location by even 10° and still maintain normalized secrecy capacity below 0.5. Exploiting the spatial-temporal propagation of the diffraction radiation pattern, Eve can generate large insecure zones and can even take advantage of propagation irregularities to eavesdrop opportunistically.




Fig. 13 Eve with MSITM alters the legitimate directional transmission to receive more than 30-dB signal power relative to the baseline

Furthermore, our results reveal that MSITM leaves a minimal energy footprint at Bob, with only a few dB power loss as shown in Fig. 14. Because such a loss is not on the order associated with blockages, but rather with typical wireless channel variations due to, for example, a slight distance change between transmitter and receiver or a small-scale antenna orientation change, it is unlikely to trigger Alice and Bob’s beam adaptation to steer away from the threat.¹⁶ Moreover, our findings reveal that Eve is robust to metasurface tilting, even to large-scale angular changes such as 40° . She is also largely unaffected by moderate-scale surface rotation. However, large-scale rotations influence the attacker’s eavesdropping location, yet the effect observed in practice is milder than the model prediction, yielding robustness to rotational angular changes. Thus, we consider the MSITM as a strong adversary that is difficult to detect and circumvent.

So far, we have explored the attack and investigated its impact on Bob (energy footprint) as it has implications for Alice and Bob’s re-steering. In future work, we plan to study the attack detection mechanisms. In particular, we will consider that Bob will try to infer the attack at different stages of the link establishment processes. For instance, Bob will try to measure ongoing time-varying channel and infer if a change is due to the addition of MSITM or nodal or environmental mobility.




Fig. 14 MSITM attack from perspective of Bob

6. AI-powered Adaptive Network Management

In this second thrust, we report on five projects that derive AI solutions to fundamental problems in the management of wireless networks at the PHY-MAC layer, namely, power allocation (for MIMO systems, energy-constrained scenarios, and federated learning systems), distributed scheduling, and access control. A common denominator in these projects is the formulation of model-based deep learning techniques—in particular, graph neural networks (GNNs)—that augment classical solutions to these problems with specially designed learning modules and achieve state-of-the-art performance.

6.1 ML-aided Power Allocation for Tactical MIMO

The emerging area of massive MIMO is a potential solution to increase rate and survivability, with increasing DOD interest in adapting commercial 5G solutions to tactical network challenges. Concomitant with this are challenges related to computational complexity of implementing (near-)optimal solutions, as the number of devices and/or number of antennas increase. Domain-aware machine learning algorithms have recently emerged as promising solution approaches. The idea here is to combine relevant domain information available in the form of a structured algorithmic solution for a specific problem with neural architectures that learn directly from data. Some of these techniques have been successfully applied in designing model-aware deep receivers to carry out symbol detection and recovery in single-input single-output (SISO) as well as MIMO systems. Another crucial problem that we have investigated within this framework is that of power allocation in wireless networks for interference management.

Optimal power allocation in wireless networks, especially under fast-changing channel conditions, is a challenging task that requires solving a non-convex

constrained sum-rate optimization problem. A classical solution approach is to reformulate it as a weighted minimum mean square error¹⁷ (WMMSE) optimization problem involving three variables. The WMMSE objective is tri-convex, that is, convex in each of its three variables when the others are fixed and, therefore, more tractable than the original problem. However, its performance is limited by the cumbersome iterative steps that tend to be computationally expensive (quadratic in the number of users) and time consuming. This is especially detrimental for fast-changing channels.

Our earlier work, unfolded WMMSE (UWMMSE)-SISO,^{18,19} addressed this drawback for a single-antenna ad-hoc interference network by truncating WMMSE iterations and embedding a GNN-based learnable component within the structure to compensate for the truncation. However, this method in its original form cannot handle MIMO systems. In this project, we design an unfolding scheme for WMMSE to make it amenable to MIMO systems.

Specifically, we make two key contributions. First, we propose the UWMMSE-MIMO architecture for power allocation in multi-antenna interference scenarios.²⁰ Then, we empirically evaluate the performance of the proposed method in comparison to WMMSE, demonstrate its generalization to networks of unseen sizes, and illustrate its robustness with respect to varying channel distributions.

To formalize the problem, we consider a single-hop ad-hoc interference network with M distinct transmitter-receiver pairs. Each transmitter is equipped with T antennas while the number of receiver antennas is R . Let $\mathbf{V}_i \in \mathbb{R}^{T \times d}$ denote the beamformer that transmitter i uses to transmit a signal $\mathbf{s}_i \in \mathbb{R}^d$ to its corresponding unique receiver $r(i)$ for all $i \in \{1, \dots, M\}$. Assuming a linear channel model, the received signal $\mathbf{y}_i \in \mathbb{R}^R$, at $r(i)$ is given by

$$\mathbf{y}_i = \mathbf{H}_{ii} \mathbf{V}_i \mathbf{s}_i + \sum_{j=1, j \neq i}^M \mathbf{H}_{ij} \mathbf{V}_j \mathbf{s}_j + \mathbf{n}_i \quad (1)$$

where $\mathbf{H}_{ij} \in \mathbb{R}^{R \times T}$ represents the channel from transmitter j to receiver $r(i)$, while $\mathbf{n}_i \in \mathbb{R}^R$ denotes additive white Gaussian noise sampled from a standard normal distribution $N(0, \sigma^2 \mathbf{I})$. Finally, the transmitted signal is estimated at $r(i)$ using a receiver beamformer $\mathbf{U}_i \in \mathbb{R}^{R \times d}$, to obtain $\hat{\mathbf{s}}_i = \mathbf{U}_i^T \mathbf{y}_i$ for all $i \in \{1, \dots, M\}$.

The performance of a wireless network, such as the one described above, is measured in terms of capacity of the network, defined as a weighted combination of the instantaneous data rate c_i achievable at each receiver $r(i)$, given by Shannon's theorem as a function of the signal-to-interference-plus-noise-ratio. More precisely, if we define the channel state information (CSI) tensor $\mathcal{H} \in \mathbb{R}^{M \times M \times R \times T}$ such that $[\mathcal{H}]_{ij::} = \mathbf{H}_{ij}$ and the transmitter beamformer tensor $\mathcal{V} \in \mathbb{R}^{M \times T \times d}$ such that $[\mathcal{V}]_{i::} = \mathbf{V}_i$, we have that for every user i its rate is given by

$$c_i(\mathbf{V}, \mathcal{H}) := \log_2 \det \left(\mathbf{I} + \mathbf{H}_{ii} \mathbf{V}_i \mathbf{V}_i^T \mathbf{H}_{ii}^T (\sigma^2 \mathbf{I} + \sum_{j \neq i} \mathbf{H}_{ij} \mathbf{V}_j \mathbf{V}_j^T \mathbf{H}_{ij}^T)^{-1} \right). \quad (2)$$

Hence, our objective is to determine the transmit beamformer $\mathbf{V}$ that maximizes the network capacity under a given power constraint

$$\begin{aligned} & \max_{\mathbf{V}} \sum_{i=1}^M \alpha_i c_i(\mathbf{V}, \mathcal{H}) \\ & \text{s.t. } \text{Tr}(\mathbf{V}_i \mathbf{V}_i^T) \leq P_{\max}, \forall i. \end{aligned} \quad (3)$$

Henceforth, for simplicity we focus on the case where every user is given the same weight $\alpha_i = 1$ in the objective. However, the optimization problem in Eq. 3 is nonconvex and has been shown to be nondeterministic polynomial-time hardness (NP-hard). The standard approach to solve this problem is to reformulate it as a constrained WMMSE optimization problem. Specifically, introducing the tensors $\mathbf{W} \in \mathbb{R}^{M \times d \times d}$ and $\mathbf{U} \in \mathbb{R}^{M \times R \times d}$ the reformulated problem is:

$$\begin{aligned} & \min_{\mathbf{U}, \mathbf{W}, \mathbf{V}} \sum_{i=1}^M (\text{Tr}(\mathbf{W}_i \mathbf{E}_i) - \log \det(\mathbf{W}_i)) \\ & \text{s.t. } \text{Tr}(\mathbf{V}_i \mathbf{V}_i^T) \leq P_{\max}, \forall i, \end{aligned} \quad (4)$$

where $\mathbf{W}_i = [\mathbf{W}]_{i::} \geq 0$ is a weight matrix at receiver $r(i)$ while the $d \times d$ matrix $\mathbf{E}_i$ is the error covariance between transmitted and received signals, which depends on $\mathbf{V}$ and $\mathbf{U}$.

The optimization problem (Eq. 4) is equivalent to Eq. 3 as it can be shown that the variable $\mathbf{V}^*$ in the global optimal solution $\{\mathbf{W}^*, \mathbf{U}^*, \mathbf{V}^*\}$ of the former is the same as the optimal transmit beamformer $\mathbf{V}^*$ in the latter. Moreover, the problem in Eq. 4 is *tri-convex*, that is, the objective is convex in each of the three variables when the other two are fixed. This makes Eq. 4 amenable to a block coordinate descent (BCD) solution, which we leverage for our unfolded architecture. In spite of this tri-convexity, WMMSE performance is limited by its cumbersome iterative steps, which are computationally complex (per-iteration complexity scales $\mathcal{O}(M^2)$ where M is the number of users). It can also be time consuming depending on the size and complexity of the wireless network, making it particularly ineffective for fast-changing channel conditions.

Algorithm unfolding refers to constructing a learnable connectionist architecture using the update structure of an iterative algorithm. The process essentially involves truncating the iterations of a classical algorithm to form a cascade of a limited number of layers, each of which copies the update operations of the algorithm in addition to carrying a neural module that approximates one or more parameters of the algorithm. In this way, an unfolded algorithm synergistically combines the domain information available in the form of a structured solution method for a given problem with neural-network-based modules that can be learned from data.

Recalling that our goal is to solve Eq. 3, we seek a parametric function Λ such that $\mathbf{v} = \Lambda(\mathcal{H}; \theta)$ approximates the solution to the problem for a given CSI $\mathcal{H}$. In this work, we propose a neural network architecture for Λ inspired by the unfolding of WMMSE into a set of hybrid layers, each of which carries an embedded GNN learning module to accelerate convergence. More specifically, by setting the matrices $\mathbf{V}_i^{(0)} = \sqrt{P_{\max}} \mathbf{1}_{T \times d}$ for every user i , we have that for layers $k = 1, \dots, K$

$$\mathbf{a}^{(k)} = \Psi(\mathcal{H}; \theta_a), \quad \mathbf{b}^{(k)} = \Psi(\mathcal{H}; \theta_b) \quad (5)$$

$$\mathbf{U}_i^{(k)} = \sum_{j \neq i} \left(\sigma^2 \mathbf{I} + \sum_{j \neq i} \mathbf{H}_{ij} \mathbf{V}_j^{(k-1)} \mathbf{V}_j^{(k-1)T} \mathbf{H}_{ij}^T \right)^{-1} \mathbf{H}_{ii} \mathbf{V}_i^{(k-1)} \quad (6)$$

$$\mathbf{W}_i^{(k)} = [\mathbf{a}^{(k)}]_i \left(\mathbf{I} - \mathbf{U}_i^{(k)T} \mathbf{H}_{ii} \mathbf{V}_i^{(k-1)} \right)^{-1} + [\mathbf{b}^{(k)}]_i \quad (7)$$

$$\mathbf{V}_i^{(k)} = \beta \left(\sum_{j \neq i} \left(\sigma^2 \mathbf{I} + \sum_{j \neq i} \mathbf{H}_{ji}^T \mathbf{U}_j^{(k)} \mathbf{W}_j^{(k)} \mathbf{U}_j^{(k)T} \mathbf{H}_{ji} \right)^{-1} \mathbf{H}_{ii} \mathbf{U}_i^{(k)} \mathbf{W}_i^{(k)} \right) \quad (8)$$

where updates (Eqs. 6–8) are repeated for every user i , Ψ in Eq. 5 is a GNN architecture with a set of trainable weights θ , $\mathcal{H}$ is a transformation to the CSI to be described later, and $\beta(\cdot)$ is a nonlinear activation. The output of our layered architecture is then given by $\mathbf{v} = \Lambda(\mathcal{H}; \theta)$ such that $[\mathbf{v}]_{i::} = \mathbf{V}_i^{(K)}$ for all i . A schematic view of the variable dependence of UWMMSE is given in Fig. 15.




Fig. 15 Schematic diagram of the proposed UWMMSE algorithm for MIMO wireless networks. A generic intermediate layer k is detailed.

The first notable point about the UWMMSE equations in Eqs. 5–8 is that if we ignore Eq. 5 and instead set $\mathbf{a}^{(k)} = \mathbf{1}$ and $\mathbf{b}^{(k)} = \mathbf{0}$ for every layer k , then Eqs. 6–8 boil down to the classical BCD closed-form updates of WMMSE. However, by providing the additional flexibility to UWMMSE of learning $\mathbf{a}^{(k)} \in \mathbb{R}^M$ and $\mathbf{b}^{(k)} \in \mathbb{R}^M$ in Eq. 5—which are implemented as weights of an affine

transform in Eq. 7—we enable faster convergence and better performance compared with the classical WMMSE.

The choice of the specific GNN Ψ in Eq. 5 and the size of the trainable weights θ are arbitrary and can be made depending on the nature of the problem. Independent of the choice of architecture, we treat the channel matrix as a weighted adjacency matrix of a directed graph, which is used to aggregate information from the neighboring nodes. In the multi-antenna setting that we consider here, the channel between a transmitter i and a receiver $r(j)$ is described by an $R \times T$ matrix that depends on the number of transmitter and receiver antennas. However, since the GNN Ψ requires the CSI between i and $r(j)$ to be represented by a scalar, we propose the use of a single-layered 1×1 depthwise convolution operation with shared filter weights to reduce $\mathcal{H}$ to an amenable structure. Essentially, we define an additional fully connected neural layer $\Phi(\mathcal{H}; \omega): \mathbb{R}^{M \times M \times R \times T} \rightarrow \mathbb{R}^{M \times M \times 1}$ where $\omega \in \mathbb{R}^{RT}$ and

$$\underline{\mathcal{H}} = \Phi(\mathcal{H}; \omega) \quad (9)$$

forms the input to Ψ . Figure 16 illustrates the operational structure of Φ where we first reshape the last two dimensions in H into a single dimension of size RT and then apply a convolution. Indeed, this operation can be interpreted as a learnable weighted combination of the RT antenna coefficients for each channel matrix $\mathbf{H}_{ij} = [\mathcal{H}]_{ij}$ to generate a scalar representation $\mathbf{H}_{ij}$ for the channel. It is important to note that the combining set of weights ω are identical for all channel elements and learned during training. This is appropriate as all channel representations must have identical functional mapping from their respective antenna coefficients in a way that is analogous to shared 1×1 convolutions of image pixels. Additionally, having a shared filter kernel allows for an $\mathcal{O}(M^2)$ reduction in the number of learnable weights.




Fig. 16 Operational structure of Φ . A shared 1×1 filter kernel ω generates scalar representations for all M^2 elements of H .

The proposed architecture has very few trainable weights, making it easy to train and likely to generalize as illustrated in the numerical examples (see Fig. 18). The number of weights θ in each of the two Ψ s is $6h + 2$, where h is the number of hidden layers. Further, the linear layer Φ has $RT + 1$ parameters, where R and T represent the number of receive and transmit antennas. The total number of trainable weights is thus $12h + RT + 6$, and is independent of the number of users M .

We now explain the nonlinear function β in Eq. 8. To incorporate the power constraint in the optimization problem, the proposed UWMMSE uses a saturation nonlinearity β in its unfolded layers, ensuring that individual layer outputs obey the power constraint. For the multi-antenna setup that we consider here, this involves constraining the transmit beamformer V such that all its elements $\mathbf{V}_i = [\mathbf{v}]_i$ satisfy $\text{Tr}(\mathbf{V}_i \mathbf{V}_i^T) \leq P_{\max}$. To attain this, the activation β applied to an arbitrary matrix $\mathbf{A}$ would generate identity if the above inequality is satisfied. Otherwise, a box nonlinearity is imposed that involves normalizing $\mathbf{A}$ with its Frobenius norm.

We train UWMMSE in an unsupervised manner by simply observing channel samples and applying stochastic gradient ascent to maximize the objective in Eq. 3 with respect to the trainable parameters in our architecture.

Our proposed architecture is composed of four unfolded WMMSE layers with two 2-layered GCNs in each unfolded layer modeling the function Ψ . The hidden layer dimension of all GCNs is set to 5. A maximum of 15000 training iterations are performed with early stopping. Batch size is fixed at 64 and the learning rate is set to 10^{-3} . We run about 160 test iterations with the same batch size.

We use randomly generated channel realizations to evaluate the model performance. We compare the performance attained by our method with that of established baselines in the low-noise regime where the achievable capacity depends strongly on the interference between users making the power allocation process more challenging. We choose the following baselines for comparison:

- 1) WMMSE forms the baseline for our experiments. We set a maximum of 100 iterations per sample.
- 2) Truncated WMMSE (Tr-WMMSE) provides a performance lower bound to UWMMSE. We fix the number of iterations to 4 to match UWMMSE layers.

The comparisons are shown in Fig. 17. Since the channel realizations are sampled randomly, we observe significant variation in the utility values for individual samples under optimal power allocation.




Fig. 17 Performance comparison of truncated-WMMSE, WMMSE, and proposed UWMMSE on various channel distributions and antenna configurations. Box plots of $M=20$ user sum-rate achieved by Tr-WMMSE (4 iterations), WMMSE (100 iterations), and UWMMSE (4 layers) (a) on a Rayleigh channel with $T=5$ transmit antennas and $R=3$ receiver antennas, transmitting $d=2$ symbols per channel use; (b) on a Rician channel with $T=3$ transmitters and $R=3$ receivers, transmitting $d=2$ symbols; and (c) on a Geometric channel with $T=3$ transmitters and $R=5$ receivers, transmitting $d=2$ symbols.

The comparison shows that the performance of our method exceeds that of WMMSE, on average, for all three channel models under various antenna configurations. Moreover, there is a significant gap in achieved mean sum-rate between UWMMSE and Tr-WMMSE with the same number of iterations as unfolded layers in all cases, proving the effectiveness of a learning-based module in a hybrid method of this form. The superiority of our method can be further illustrated using the inference time comparisons. While WMMSE requires ~ 600 msec on average to converge, per sample, our method has an average inference time of ~ 30 msec, effectively providing a $20 \times$ speedup. Time taken by Tr-WMMSE is ~ 24 msec; however, the achieved performance is poor. Our method, therefore, exceeds the performance of WMMSE with a full set of iterations under a time complexity that is comparable to its truncated version.

We also investigated the robustness of our proposed method to variations in channel distributions and wireless network size. These are realistic scenarios since the channel distribution depends largely on geographical and climatic conditions as well as the community structure (rural, urban, suburban) in a given area. Further, an ad-hoc wireless network may undergo modifications in size as new users join the network and old users leave. It is imperative that the power allocation method be able to maintain a steady performance despite these variations. To this end, we compare the performances of two versions of our model, one of which is trained on Rayleigh and the other is trained on Rician channel realizations. These comparisons are performed independently on a Rayleigh and a Rician test set. Figure 18 (a) and (b) show that both models are effective in achieving comparable performance for out-of-distribution data points. This observation indicates that our model can endure changes in channel conditions without significant performance degradation. In Fig. 18 (c), we consider a slightly more challenging scenario wherein, under a

geometric channel setting (a geometric channel model has a composite structure with path loss and Rayleigh fading components), the network size varies within a fixed range. To evaluate the interpolation behavior of our model, we sample only odd-number network sizes in training while the test set comprises all integer values in the specified range. In this case, the Rayleigh and Rician versions of UWMMSE fail to maintain a steady performance across the range of sizes, doing slightly better for a few smaller sizes but degrading fast as the network grows. However, UWMMSE trained on geometric channels is able to maintain a constant gain over WMMSE at all sizes. This superiority can be attributed to the GNN learning module, which is regularized by the inherent graph structure of geometric channels. Since real-world channel models are more aligned toward the geometric setup composed of path loss and fading, our model has a natural advantage in dealing with more challenging scenarios that will be considered in future versions of this work.




Fig. 18 Robustness analysis for UWMMSE. (a) Box plots of the achieved utilities for 10000 randomly drawn Rayleigh channel state matrices for $M = 20$ users. UWMMSE trained on Rayleigh compared with UWMMSE trained on Rician along with WMMSE. (b) Box plots of the achieved utilities for 10000 randomly drawn Rician channel state matrices for $M = 20$ users. UWMMSE trained on Rician compared with UWMMSE trained on Rayleigh along with WMMSE. (c) Mean sum-rate achieved on a geometric channel as network size varies.

To conclude, we proposed a MIMO version of UWMMSE, a hybrid learnable power allocation algorithm for wireless networks. The main contribution of this method is the use of GNN-based learning modules that leverage the connectivity structure of the transceivers to learn a mapping for faster parameter updates of an unfolded iterative algorithm tuned to approximate a global solution of the NP-hard power allocation problem. It is this hybrid architecture that captures on the one hand the dynamic structure of the iterative solution while on the other hand approximates the cumbersome computational steps with learnable weights and drives the model towards faster convergence. We note that our GNN modules have very few learnable parameters, and thus are easy to train. In the future, we will go beyond synthetic data and validate empirically the effectiveness of our proposed method on real-world data sets.

6.2 Graph-based Algorithm Unfolding for Energy-Aware Power Allocation in Wireless Networks

In this project, we develop a novel graph-based trainable framework to maximize the weighted sum energy efficiency (WSEE) for power allocation in wireless communication networks. To address the non-convex nature of the problem, the proposed method consists of modular structures inspired by a classical iterative suboptimal approach and enhanced with learnable components. More precisely, we propose a deep unfolding of the successive concave approximation (SCA) method. In our unfolded SCA (USCA) framework, the originally preset parameters are now learnable via graph convolutional neural networks (GCNs) that directly exploit multiuser CSI as the underlying graph adjacency matrix. The USCA framework is trained through a stochastic gradient descent (SGD) approach using a progressive training strategy. The unsupervised loss is carefully devised to preserve the monotonic property of the objective under maximum power constraints.

Energy efficiency wireless power allocation is motivated by the fact that it can prolong network lifetime, help reduce greenhouse gas emissions and electromagnetic pollution,²¹ and is also vital to reduce electricity budgets and prolong the battery life of mobile devices.²² However, due to the nonconvex nature of energy efficiency (EE) metrics, EE-oriented problems in interference networks (Fig. 19) are typically NP-hard.²³ Thus, traditional convex optimization is practically incapable of directly handling the problem at hand. With the aid of monotonic programming, works such as Matthiesen et al.²⁴ reduced the computational time of globally optimal branch-and-bound procedures, but it scales poorly with the size of the network.




Fig. 19 An illustration of the system model and the graph representation of CSI. (a) Multicell multiuser interference-limited wireless network considered. (b) The channel gain and multiuser interference coefficients of this network can be viewed as a graph, taking the CSI matrix $\mathbf{H}$ as its adjacency matrix.

Hence, several suboptimal methods based on optimization theory have been developed to obtain approximate power policies while promoting scalability, one of which, namely the SCA algorithm, appears the most interesting due to its popular adoption and guaranteed convergence.²⁵ SCA has proven state-of-the-art performance in the maximization of some EE metrics such as the system global EE²⁶ and the minimum of individual EE.²⁷ However, for metrics involving the weighted sum or product of node-wise EE there is a notable gap in performance between SCA solutions and global optima even with only seven users.²⁴ In addition, the time-varying channel characteristics pose a harsher requirement on computational efficiency as the allocation needs to be faster than the channel coherence time. In this context, the nested iterative nature of SCA can preclude its practical implementation.

In overcoming these identified challenges, we seek to enhance SCA with data-driven machine learning elements. Deep neural networks (DNNs) have shown promising success in wireless power allocation and other communication problems.^{28–30} Most of these approaches are model-free, meaning that they parameterize some function of interest with established multipurpose architectures such as multilayer perceptrons (MLPs),²⁸ convolutional neural networks,²⁹ and GNNs.³⁰ Hence, these approaches rely on the approximation capability of DNNs to mimic the optimal solutions of classical optimization techniques based on solved examples. On the downside, these deep models do not tend to generalize well to unseen scenarios such as shifts in channel distributions, varying network sizes, or

user mobility. Moreover, obtaining a large set of labeled data, or optimal solutions in our case, as are needed to train DNNs in a supervised manner, may be infeasible for computationally intractable problems.

Our goal is to achieve higher performance and faster computation than classical methods by combining the strengths of expert knowledge and statistical learning. Motivated by this, we adapt the paradigm of algorithm unfolding³¹ to accomplish this synergistic combination. Before diving into any technical details, we give a formal statement and the SCA solution of the maximum WSEE power allocation problem as follows, with μ_i and $P_{c,i}$ being system constants at user i .

$$\begin{aligned} \max_{\mathbf{p}} \quad & \sum_{i=1}^L w_i \text{EE}_i, \text{ where } \text{EE}_i = B \frac{\log \left(1 + \frac{H_{i,i} p_i}{1 + \sum_{j \neq i} H_{i,j} p_j} \right)}{\mu_i p_i + P_{c,i}}, \\ \text{s.t.} \quad & 0 \leq p_i \leq P_{m,i}, \quad i = 1, \dots, L. \end{aligned} \quad (\text{P1})$$

Here, EE—the energy efficiency in bits/Joule—is a function of channel bandwidth B , the channel gain from transmitters to both intended and unintended receivers H_{ij} , and transmit powers p_j . By the classical SCA algorithm, a first-order optimal solution to Eq. P1 can be found by iteratively solving the following approximating Eq. P2 instead, in which $\{\tilde{\text{EE}}_i(\mathbf{p}; \mathbf{p}^{(t)})\}_{t=0}^T$ is a sequence of (pseudo-)concave surrogate functions satisfying a series of assumptions,²⁵ whose expression is omitted here, but can be found in prior work.²⁴

$$\begin{aligned} \max_{\mathbf{p}} \quad & \sum_{i=1}^L w_i \tilde{\text{EE}}_i(\mathbf{p}; \mathbf{p}^{(t)}), \\ \text{s.t.} \quad & 0 \leq p_i \leq P_{m,i}, \quad i = 1, \dots, L. \end{aligned} \quad (\text{P2})$$

We sequentially solve Eq. P2 following the update policy on $\mathbf{p}(t)$ given by the line search

$$\mathbf{p}^{(t)} = \mathbf{p}^{(t-1)} + \gamma^{(t)} (\mathbb{B} \mathbf{p}^{(t)} - \mathbf{p}^{(t-1)}), \quad (10)$$

where $\mathbb{B} \mathbf{p}^{(t)}$ is a global maximizer of Eq. P2 at the t -th iteration. Using the Armijo rule, a scalar step size $\gamma^{(t)}$ can be picked for a substantial update in the ascent direction, specific to the current iteration. Leveraging the convergence of the subproblems in Eq. P2, it has been shown²⁵ that any limit point of $\{\mathbf{p}^{(t)}\}$ is guaranteed to be a stationary point of Eq. P1. However, SCA requires nested iterations to solve Eq. P2 and update Eq. 10 for convergence and is thus especially inefficient in large networks. Additionally, it has been demonstrated²⁴ that SCA depends on sophisticated initializations to land in good local optima, which makes its computation even more inefficient. Intuitively, if we can find steeper ascent directions (i.e., a $\mathbb{B} \mathbf{p}^{(t)}$ in Eq. 10 that need not be the exact global maximizer of the

concave approximation but can be closer to the true non-approximated objective) and better step sizes (i.e., self-adaptive step sizes that can be different in every dimension) at each iteration, we can achieve higher performance with fewer iterations of these enhanced SCA blocks. With this motivation, we illustrate in Fig. 20 our proposed unfolded architecture USCA,³² in which key functions of SCA are augmented by learnable neural networks.




Fig. 20 Overall schematic view of our proposed USCA architecture for energy-efficient power allocation. (a) The embedding block. This block encodes channel information through a GCN, whose output is used as the input to the first USCA block. (b) The USCA block. This block contains two GCNs that mimic the classical operations in SCA of solving Eq. P2 and finding the step size to update the power allocation. (c) The GCN. A schematic view of a generic GCN with rectified linear unit (ReLU) nonlinearities.

To sum up, our proposed solution, namely the unfolded SCA (USCA) framework, combines classical (model-based) and learning-based (model-free) methodologies in a deep unfolding manner. USCA is learning-based in the sense that it leverages SGD to learn GNN parameters from simulated channel data. GNNs are able to exploit topological patterns of interference between users to locally process instantaneous CSI. At the same time, our method is model based since it takes advantage of the classical SCA algorithm by informing a stacked neural network architecture that imitates the iterative structure of SCA. In training this framework, we adopt a progressive learning rate decaying strategy³³ to stabilize the learning process.

Through numerical experiments, we demonstrate the superior performance of the proposed USCA method and its generalizability over different network sizes and channel distributions. In Fig. 21, we compare USCA with the following competitors and benchmarks on the basis of the wideband spatial (WBS) path-loss (PL) model, including:

- 1) SCA: Classical SCA with a sufficiently large enough upper limit of iterations.

- 2) Tr-SCA: The truncated SCA, with a smaller number of iterations matching the depth of USCA.
- 3) MLP-USCA: An MLP-based USCA framework, simply obtained by replacing GCNs Ψ with MLPs.
- 4) GCN: Plain GCN without the unfolding architecture.
- 5) Max-Pow: Maximum power policy.
- 6) Optimal design (OPT): Globally optimal solution given by the branch-and-bound procedure using monotonic optimization.




Fig. 21 Performance comparison of the proposed GCN-based USCA with competing approaches. (a) Performance on the test set compared with all competitor methods. Every point in the curves is given by the average WSEE over all test channel samples under the corresponding power constraint value. (b) Histogram of the achieved WSEE values for test channels compared with model-based competitor methods. (c) Performance at progressive training milestones compared with learning-based competitor methods.

So far, the results are focused on a test set with $M = 4$ receivers (base stations) and $L = 8$ transmitters (users). To demonstrate the scalability of the proposed method, we also showcase the models' performance with different network sizes. In particular, the scenario of varying L is of special relevance since, due to mobility, users might join and leave the system at any time. We also consider the case of varying M to assess whether a model trained on a smaller communication infrastructure can scale to a larger system. The models being tested are directly taken from the previous comparison (thus, trained on $M = 4$ and $L = 8$), and their performance is shown in Table 1.

Table 1 Average WSEE values of power allocations on test sets with different numbers of users associated with different numbers of base station (BS). The best performance for every scenario is bolded and the true optimum (when computationally feasible) is underlined.

#BS (M)	4			9	16	
#Users (L)	6	8	12	18	48	100
GCN-USCA	5.876	6.766	8.843	15.518	31.198	32.359
MLP-USCA	-	5.445	-	-	-	-
GCN	5.750	6.541	8.543	15.046	30.468	33.298
Tr-SCA	5.663	6.131	7.358	13.401	25.813	29.591
SCA	5.701	6.148	7.407	13.418	25.823	29.595
Opt	<u>6.087</u>	-	-	-	-	-

The first thing to notice is that, excluding the method computing the true optimum, GCN-USCA performs the best for every configuration except for the case $L = 100$, where it ranks second behind GCN. Importantly, GCN-USCA beats the classical baseline SCA for every considered configuration.

For the case $M = 4$ and $L = 6$, the system is small enough for us to compute the true optimal power allocation. Thus, we can see that GCN-USCA reduces the suboptimality gap of classical SCA by 45% while being computationally more efficient. Lastly, note that MLP-USCA can only be tested in the same setting on which it was trained since its input and output dimensions are fixed. This is to show the versatility of graph-based learning principles in wireless communication networks, which further underscores the importance of considering GCNs for the Ψ functions in the unfolded architecture. GCNs not only lead to better performance within the configuration seen during training but also enable the generalization to systems of varying size, a necessary requirement in practical deployments.

Finally, we depart from the WBS model for our propagation model and study how generalizable the allocation methods are to changes in the test channel distribution. The test data sets are generated following the Hata-COST231 propagation model.³⁴ Four types of path loss effects are considered—urban or suburban environment, with or without the log-normal shadowing—with acronyms of Urb-noSF (Hata-COST231 urban model without shadowing), Sub-SF (Hata-COST231 suburban model with shadowing), Sub-noSF, and Urb-SF. We assume 1.9 GHz for the carrier frequency and 30 m for the BS height; if applicable, the standard deviation of shadowing is 8 dB. Again, for each candidate method, the model being tested here is the exact model from Fig. 21 without any further processing. Numerical results are reported in Table 2.³⁵

Table 2 Average WSEE values of power allocations on test sets with different channel distributions as a result of different PL propagation models

PL propa- gation model	Wideband spatial	Hata Urban		Hata Suburban	
		w/o SF	SF	w/o SF	SF
GCN-USCA	6.766	7.427	8.501	10.137	10.997
MLP-USCA	5.445	5.536	6.419	6.770	7.461
GCN	6.541	7.198	8.260	9.770	10.699
Tr-SCA	6.131	6.532	7.791	9.132	10.426
SCA	6.148	6.661	8.056	9.768	11.298
GCN-USCA ⁺	6.771	7.548	8.676	10.737	11.699

Again, without any retraining, end-to-end GCN-USCA achieves the best performance in almost every setting, only ranking second in sub-SF behind the classical SCA baseline. If the same initialization as SCA is applied to enforce the monotonicity of the resulting WSEE values achieved by GCN-USCA (and we call it GCN-USCA⁺), it would take a dominant lead over all candidate methods in all cases considered. Besides, a few observations are in order. First, the fact that GCN-USCA consistently outperforms MLP-USCA and GCN demonstrates good generalizability properties of our method compared with other learning-based approaches. Second, notice that the gap between Tr-SCA and SCA increases for the suburban channels, possibly indicating a more challenging allocation scenario. This is consistent with our (pre-trained) GCN-USCA underperforming in such a setting. Overall, GCN-USCA has shown reasonably good generalizability to channel models unseen during training. If the mismatch between training and testing data proves to be a major problem, the learning-based method is easily improvable by transfer learning given adequate (or even imperfect) data from the target distribution.

In conclusion, we developed USCA, a novel graph-based learning approach to allocate transmit powers with optimized EE in wireless interference networks. Using algorithm unfolding directed by classical SCA, we proposed the block-wise modular architecture of USCA that intrinsically incorporates problem-specific modules augmented by trainable components. To parameterize these components, GCNs are leveraged to make explicit use of the underlying topology of communication networks. Through extensive analyses and experiments, we have demonstrated that USCA yields solutions that are: i) Accurate, since it consistently outperforms the original algorithm and well-established benchmarks; ii) Efficient, because the feedforward neural architecture, once trained, costs less time and computation than model-based methods; and iii) Distributed, owing to the decentralized implementation proposed for its modular GCN-based architecture.

A natural direction to explore next would be to obtain energy-efficient power allocation in case of higher degrees of heterogeneity, such as nonuniform system constants, user-specific power or other constraints, and different user weights. From a methodological perspective, this work also paves the way for many intriguing avenues for future research in wireless resource allocation not limited to power. First, we can dig deeper into understanding the transformed loss landscape of the sum-of-ratios fractional objective by, for example, visualizing the morphology and dynamics of loss function during training or fine-tuning. Second, in lieu of the mentioned post-training fine-tuning, we are interested in ways to inherently promote the model’s generalizability—for example, alternative designs of learning components and distribution diversity in training data. Finally, a direction of practical importance is to consider scenarios where channel information is missing, noisy, or even adversarial, which may raise broader problems to address network security.

6.3 Power Allocation for Wireless Federated Learning using Graph Neural Networks

Federated learning (FL) is a collaborative machine learning methodology in which a model is trained based on multiple local data sets without explicitly exchanging data samples.³⁶ This makes it particularly relevant in settings where privacy concerns or policy constraints are important as in a coalition military network. To converge, a typical FL process may involve multiple iterations where the server broadcasts a global model to selected workers who update this model using their on-device data and transmit the updated local models back to the server who aggregates them into a new global model, so on and so forth. Thus, the quality of communications from the workers to the server is critical to good learning performance. In this project, we aim to maximize the uplink transmission success rate under an overall data rate constraint in each iteration, which should ultimately improve FL in interference-limited wireless networks with fading channels.

To accommodate the real-time requirement in dynamic environments, current state-of-the-art works in power allocation are mainly suboptimal optimization methods or their parameterized trainable versions. For example, existing work³⁷ analyzes the impact of the wireless channel on FL convergence rate, based on which a joint optimization scheme of power, worker, and resource management is proposed for static channels. Besides, machine learning has been successful in optimizing power allocation with various objectives (e.g., to maximize system utility in data rate, energy efficiency, or proportional fairness). For example, to maximize sum-rate in interference networks, a suboptimal power allocation algorithm was enhanced with augmented parameters learned by GCNs.^{18,20}

The empirical success of GCNs inspires us to leverage their learning capabilities to develop a power allocation policy that is efficient and tailored to the specific requirements of FL. Nonetheless, the existing works cannot be directly applied to our case due to two reasons. First, most cases of power allocation with general purpose are optimized under only trivial constraints, such as a box constraint on the transmit power,^{32,37} whereas complicated constraints (e.g., delay or energy) may arise in response to the specific demands of FL. Second, a common limitation is the assumption of orthogonal channels without multiuser interference, which simplifies the problem of optimal power control but reduces the optimality of the derived solutions.^{38,39} Instead, we formulate our problem as determining the power allocation policy $p^*: \mathbb{R}^{L \times L} \rightarrow \mathbb{R}^L$ that solves the following optimization problem

$$p^* = \operatorname{argmax}_p g(\mathbb{E}_{\mathbf{H} \sim \mathcal{H}}[\mathbf{1} - \text{PER}(p(\mathbf{H}), \mathbf{H})]), \quad (\text{P3})$$

$$\text{s.t. } r_{0,i} \leq \mathbb{E}_{\mathbf{H} \sim \mathcal{H}}[R_i(p(\mathbf{H}), \mathbf{H}) \mid R_i(p(\mathbf{H}), \mathbf{H}) > 0], \forall i, \quad (\text{P3-a})$$

$$p(\mathbf{H}) \in [0, P_{\max}], \quad \forall \mathbf{H}, \quad (\text{P3-b})$$

where $\mathbf{H}$ is a CSI instance drawn from some channel distribution $\mathcal{H}$, and $r_{0,i}$ is a prespecified minimum level of desired expected data rate for user i . The packet error rate (PER) and achievable data rate R of user i have the following definitions,

$$\text{PER}_i(\mathbf{p}, \mathbf{H}) = 1 - e^{-\frac{m}{\text{SINR}_i}} \text{ and } R_i(\mathbf{p}, \mathbf{H}) = B \log(1 + \text{SINR}_i),$$

$$\text{SINR}_i = \frac{H_{i,i}p_i}{1 + \sum_{j \neq i} H_{i,j}p_j}, \quad \forall i = 1, \dots, L.$$

where

The concave and increasing function g is defined by $g(\mathbf{q}) = \sum_i w_i g(q_i)$. Hence the power allocation seeks to maximize the weighted sum of successful transmission probabilities, which is a proxy for the quality of the global FL model. Optimally solving Eq. P3 is challenging. Apart from having a non-convex objective, there are at least two additional sources of complexity: i) We are optimizing over the space of power allocation functions, which is inherently infinite-dimensional, and ii) We are not given a specific channel or sets of channels but rather we have to consider expected values over a whole distribution of channels.

To remedy these shortcomings, we propose an algorithm called *PDGNet* (Primal-Dual Graph convolutional power Network) that directly learns a power policy via a GCN aided primal-dual learning architecture. PDGNet essentially comprises two parts. It first parameterizes a subset of all possible power policies p using a GCN such that $p(\mathbf{H}) = \Psi(\mathbf{H}; \Theta)$, and instead optimizes over the trainable weights Θ of this neural architecture. The choice of GCN is driven by the following three appealing characteristics, i) permutation equivariance, meaning that the allocated power is not affected by the indexing of workers, ii) flexibility to varying number

of workers, and iii) distributed deployment, because GCNs are essentially a bank of local graph filters.

We restate Eq. P3 in a way that is amenable to a primal-dual solution

$$\max_{\Theta, \tilde{\mathbf{q}}, \mathbf{r}} g(\tilde{\mathbf{q}}), \quad (\text{P4})$$

$$\text{s.t. } \tilde{\mathbf{q}} \leq \mathbb{E}_{\mathbf{H} \sim \mathcal{H}}[1 - \text{PER}(\Psi(\mathbf{H}; \Theta); \mathbf{H})], \quad (\text{P4-a})$$

$$r_i \leq \mathbb{E}_{\mathbf{H} \sim \mathcal{H}}[R_i(\Psi(\mathbf{H}; \Theta), \mathbf{H}) \mid R_i(\Psi(\mathbf{H}; \Theta), \mathbf{H}) > 0], \quad (\text{P4-b})$$

$$r_i \in [r_{0,i}, +\infty), \quad \forall i, \quad (\text{P4-c})$$

$$\Psi(\mathbf{H}; \Theta) \in [0, P_{\max}], \quad \forall \mathbf{H}, \quad (\text{P4-d})$$

whose Lagrangian is given by

$$\mathcal{L}_{\Psi}(\Theta, \tilde{\mathbf{q}}, \mathbf{r}, \lambda_q, \lambda_r) = g(\tilde{\mathbf{q}}) + \lambda_q^{\top} (\mathbb{E}_{\mathbf{H}}[f_q] - \tilde{\mathbf{q}}) + \lambda_r^{\top} (\mathbb{E}_{\mathbf{H}}[f_r \mid f_r > 0] - \mathbf{r}).$$

Following standard primal-dual solutions of optimization problems,⁴⁰ this Lagrangian motivates an iterative primal-dual learning approach. It can be shown that the duality gap of Eq. P4 is dependent on the expressive power of the parameterization³⁰ of p . Figure 22 gives a high-level illustration of the PDGNet architecture. To the best of our knowledge, this is the first GCN-based power allocation policy for FL over (interference-limited) wireless networks.




Fig. 22 Illustration of the proposed PDGNet for FL power allocation

We provide extensive evaluation of the proposed method in two aspects (Fig. 23). We compare PDGNet to several benchmarks including: i) a random power policy (Rand), ii) a maximum power policy (optimal in the orthogonal channel case, Orth), and iii) a topology-agnostic primal-dual learning based on feedforward MLP (PDF). First, we show the resulting PER values as a proxy of the communication performance with various wireless configurations, such as interference strength, maximum power, and size. In all cases considered, PDGNet always demonstrates

a lower PER than others. Additionally, the final FL performance is evaluated via a classification task using the MNIST data set. The ideal FL convergence curve (with error-free transmission) is also plotted to show that the proposed PDGNet can result in near optimal FL performance, significantly outperforming the competing power allocation methods.




Fig. 23 PERs and performance of the FL model achieved by different power allocation policies. The first group of comparisons is based on the weighted-sum PER of selected workers against (a) interference strength, (b) maximum power constraint value, and (c) network size; the last one (d) displays the MNIST classification error of the FL global model at the end of each FL iteration, including the ideal FL (assuming no transmission loss) performance. The proposed PDGNet outperforms the baseline power allocation strategies in all the metrics considered.

In our next steps, we will enforce additional constraints such as bounds on the expected energy consumption and incorporate real-time prediction accuracy of the global model as a criterion for worker selection or power allocation at each iteration. By doing so, we can force the resource allocation to directly optimize the desired FL performance, thus facilitating heterogeneous workers and more sophisticated tasks. We are also interested in exploring more challenging application scenarios such as large-scale communication networks, imperfect channel information, and heterogeneous data at the workers.

6.4 Distributed Scheduling using Graph Neural Networks

During Year 2, we continued our study of distributed link scheduling in wireless multihop networks in a time-slotted system with orthogonal multiple access. On top of the throughput-oriented scheduling studied in our previous work,⁴¹ we focused on improving the delay performance and signaling overhead of distributed link scheduling. In many scenarios, the delay performance of wireless ad-hoc networks is more important than throughput, whereas the maximum weighted independent set (MWIS)-formulated link scheduler has poor latency performance. Moreover, in wireless ad-hoc networks with high density, such as Internet-of-Things, the MWIS-formulated link scheduler would have unreasonably high signaling overhead, since the signaling overhead is proportional to the average degree of the conflict graph.

We developed two schedulers improving the delay performance and signaling overhead, respectively. These two schedulers share similar formulation and downstream architecture with the throughput-optimal scheduler,⁴¹ described as follows:

We first construct a conflict graph of links in the wireless network, as exemplified in Fig. 24. A vertex of the conflict graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, as shown in Fig. 24 (right), represents a link in the connectivity graph of the wireless multihop network shown in Fig. 24 (left), and an edge in the conflict graph represents a conflict relationship between the corresponding two links in the physical graph, that is, two active links interfering each other when they are too close, or two links sharing the same user.




Fig. 24 (left) An example of a wireless multihop network and (right) its corresponding conflict graph

The MWIS problem is NP-hard and needs to be approximated by heuristic solvers in practical scheduling. The baseline distributed heuristic is the local greedy solver,⁴² which builds a solution in an iterative fashion: in each iteration, a node is selected if its utility is greater than all of its neighbors, and the neighbors of the selected nodes are then excluded from the solution:

$$v \in \hat{\mathbf{v}}_{\text{Gr}}, \text{ if } u(v) > \max_{v_i \in \mathcal{N}(v)} u(v_i).$$

This process is repeated on the residual graph $\mathcal{G}' = \mathcal{G} \setminus (\hat{\mathbf{v}} \cup \mathcal{N}(\hat{\mathbf{v}}))$ until all nodes are either in or excluded from the solution, or in other words, when the residual graph is empty. The local greedy solver has good scalability due to its efficiency.

6.4.1 Delay-Oriented Distributed Scheduler

The throughput-optimal scheduler adopts a queue-based per-link utility, for example, the product of queue length and link rate.⁴² Moreover, the memoryless nature of queue length-based utility makes it difficult to optimize performance metrics that depend on serial decisions, such as delay. As an example, consider a wireless network with six links with a conflict graph of star topology shown in Fig. 25, an initial state of all empty queues, a constant arrival rate of 1 and a link rate of 2 (both in packets per time slot) for each link, and set the per-link utility as the queue length so that links with longer queues are preferentially scheduled.




Fig. 25 Steady-state transitions with (left) an optimal MWIS scheduler in green and (right) a greedy scheduler in blue

The optimal MWIS scheduler will schedule the set that achieves the maximum sum of queue lengths, leading to the alternation of two network states, as in Fig. 25 (left). In contrast, the greedy scheduler always starts with the link with the largest queue, leading to the alternation of the two network states in Fig. 25 (right). As a result, the average queue length under the optimal MWIS scheduling is 2.17, which is greater than that of the greedy scheduler, 1.5. In fact, for unsaturated network traffic, the optimal MWIS solver has poorer delay performance than a simple greedy heuristic.

Existing approaches to reduce the delay of scheduling include setting delay as a constraint of optimization,^{43,44} replacing the queue length in the conventional utility functions with metrics of delay (e.g., virtual queues of congestion,⁴⁵ sojourn time,⁴⁶ age-of-information⁴⁷), and machine learning-based solutions.^{48–50} Conventional analytical approaches generally do not fully consider network topology, while the machine learning-based solutions are either limited to centralized scheduling or have poor scalability. To address these problems, we develop a GCN-based distributed scheduler. The downstream architecture of our GNN-based scheduler contains a graph convolutional neural network followed by a local greedy solver,⁵¹ as shown in Fig. 26.




Fig. 26 Downstream architecture of GCN-based distributed scheduler for delay-oriented scheduling

The downstream architecture in Fig. 26 is similar to that of the throughput-optimal scheduler in our previous work.⁴¹ The difference is that in our delay-oriented scheduler, the GCN is no longer featureless, and the node embeddings are directly fed to the local greedy solver. The input feature $\mathbf{S}$ could be the queue length-based per-link utility, along with other features, such as link hierarchy. Both the GCN and local greedy solver can generalize to networks of various sizes and topologies, and be deployed in a distributed manner.

To train the GCN in Fig. 26, we develop a reinforcement learning approach with a reward signal based on relative performance in K -step lookahead scheduling. Specifically, we first create an interactive virtual environment, where the conflict graph, packet arrivals, and link rates are drawn from stationary distributions of interests. Then, at each time slot, we execute two different scheduling policies for K -steps under identical environmental (arrivals and link rates) conditions, as shown in Fig. 27. The baseline policy is the local greedy scheduler with per-link utility given by $\mathbf{u}(t) = \mathbf{q}(t) \odot \mathbf{r}(t)$, where $\mathbf{q}(t)$ and $\mathbf{r}(t)$ represent the queue length and link rate at t .




Fig. 27 Four-step lookahead scheduling from time slot t

The reward signal is the ratio of average queue length over K steps between the baseline policy and GCN policy under training. Then, we have the following return vector at t ,

$$\boldsymbol{\rho}(t) = \varphi \left(\frac{\sum_{k=1}^K \|\tilde{\mathbf{q}}(t+k)\|_1}{\sum_{k=1}^K \|\mathbf{q}(t+k)\|_1} \right) \mathbf{v}(t) + \mathbf{u}_{\text{GCN}}(t) \odot [\mathbf{1} - \mathbf{v}(t)] ,$$

where $\mathbf{v}(t)$ is the indicator vector of the schedule $\hat{\mathbf{v}}_{\text{GCN}}(t)$, and ϕ is an activation function that can be either linear or a Heaviside step function. Intuitively, whenever the GCN policy beats the baseline in a K -step lookahead scheduling, the schedule $\hat{\mathbf{v}}_{\text{GCN}}(t)$ is encouraged for the network state at t . Consequently, a root-mean-square loss between the node embedding vector and the return vector is minimized to train our GCN.

Through numerical experiments, we demonstrate that the GCN-based distributed scheduler can reduce the delay performance, when the conflict graph is more centralized, and the central links are more likely to be congested by the memoryless local greedy baseline. Figure 28 shows the boxplot of the approximation ratios of the queue lengths of the GCN-based scheduler relative to the local greedy scheduler on synthetic random graphs of different graph models, where the number under the graph model labeled on x-axis is the peak-to-average ratio of the node degree in the conflict graph (larger peak-to-average ratio refers to a more centralized conflict graph). The Erdős-Rényi (ER) conflict graphs seek to represent networks of uniformly distributed users with identical transmit power (i.e., unit-disk interference zone). The star and Barábasi-Albert (BA) models arise from networks with several connected components and links of heterogeneous transmit power, for example, macro-cells surrounded by micro-cells and device-to-device links. The power-law trees represent wireless backhaul networks where some links have interfering neighbors.

To summarize, our GCN-based distributed scheduler is scalable and improves the delay performance in wireless networks by combining the efficiency of a local greedy scheduler and the ability of GCNs to encode the network state and topological information. The proposed scheduler exhibits a marked improvement in performance over classical baselines, especially in wireless networks with several connected components and heterogeneous transmit power, and shows good generalizability over graph types, graph sizes, and traffic loads.




Fig. 28 Approximation ratios of the queue lengths of the GCN-based distributed scheduler relative to the local greedy scheduler under various conflict graph configurations. Smaller values refer to better delay. Smaller ratios indicate better performance relative to baseline.

6.4.2 Distributed Link Sparsification for Massive Access

Existing distributed schedulers, such as the baseline local greedy solver and carrier-sense multiple access (CSMA), face the challenge of high signaling overhead when it comes to massive access. Massive access refers to the multiple access for wireless networks with numerous devices and high network density, especially due to the proliferation of wireless devices and emerging machine-to-machine traffic, such as in the Internet of Things. Being designed for throughput or utility maximization, existing distributed schedulers introduce contention overheads or collision rates proportional to the average size of the neighborhood in the network, as illustrated in Fig. 29. For example, in the local greedy solver, a link needs to exchange messages with every link in its neighborhood. This can lead to prohibitively high scheduling overhead in massive access, as well as increased congestion, instability, energy consumption, radio footprint, and security vulnerability in network operations.




Fig. 29 Wireless multihop networks with neighborhoods of different sizes: (left) dense network with large neighborhoods and (right) sparse network with small neighborhoods. The overhead, collision, and starvation of scheduling increases with the neighborhood size, while the relays and thus overhearing increase in sparser networks.

Existing approaches to the aforementioned problems, such as topology control,^{52–54} sleep scheduling,^{55,56} and cross-layer optimization,⁵⁷ are not applicable to massive access in general wireless multihop networks. Therefore, we develop a GCN-based distributed link sparsification scheme to improve the massive access and coexistence of the wireless network. To reduce the scheduling overhead, our scheme assesses a link using both the utility value and topological information, and withholds it from a scheduling contention that it could not win. Our optimization addresses two conflicting objectives: to retain as much network capacity as possible, while removing as many vertices in the original conflict graph as possible.

We seek to (approximately) solve the problem by restraining ourselves to a parametric family of functions given by

$$h_v(u(v); \mathbf{Z}) = z_0(v)u(v)H\left(z_0(v)u(v) - z_1(v)u^{(\eta)}\right),$$

where $\mathbf{Z} = [\mathbf{z}_0, \mathbf{z}_1] \in \mathbb{R}^{|\mathcal{V}| \times 2}$ is a matrix collecting the local parameters for all vertices in the conflict graph, $H(\cdot)$ is the Heaviside (step) function, and $u^{(\eta)}$ is the η -quantile utility under the network state distribution, as illustrated in Fig. 30 (a).




Fig. 30 Architecture of GCN-based distributed sparse scheduler. (a) Select cut-off utility based on the empirical cumulative density function of utility, (b) parameters $\mathbf{Z}$ generated by a featureless GCN based on conflict graph, and (c) sparse graph and topology-aware utility created by the parameterized functions, then the distributed scheduler generates schedules.

The parameters $\mathbf{Z}$ are generated by a featureless GCN as node embedding in a distributed manner, as shown in Fig. 30 (b). Since the GCN is featureless, $\mathbf{Z}$ only needs to be updated when the topology of the network changes, which happens several orders of magnitude slower than the change of per-link utility. The family of parameterized function, $\{h_v(u(v); \mathbf{Z})\}$, creates a sparse conflict graph $\mathcal{G}^s$ and topology-aware utility vector $\mathbf{W}^s$, based on which a schedule is generated by the distributed scheduler $\hat{\mathbf{v}}^s = c(\mathcal{G}^s, \mathbf{W}^s)$, as illustrated in Fig. 30 (c). The training of the GCN is carried out in two stages. Intuitively, the first stage trains the GCN to generate threshold with favorable selectivity, the second stage adjusts the node embeddings so that they can outperform the statistical baseline under different cut-off thresholds.

Two numerical experiments are conducted to evaluate the performance of the GCN-based scheme and the statistical baseline. In the first experiment, we draw a random network state (conflict graph and utility vector) from the distribution of interest, and the network state is fed to three schedulers, GCN-based sparsification, statistical baseline, and the reference dense scheduler (vanilla local greedy scheduler⁴²). The relative performance of the first two sparse schedulers with respect to the reference dense scheduler is illustrated in Fig. 31. With the same input, the GCN-based method can achieve higher network capacity (total utility in Fig. 31 [a]), lower average node degree (Fig. 31 [c]), and lower P2P message complexity (Fig. 31 [d]), than the statistical baseline.




Fig. 31 The performance of sparse schedulers, GCN and statistical method (Stat), with respect to a reference dense scheduler, given an identical input network state. (a) Approximation ratio of total utility, (b) Retention ratio of nodes, (c) Retention ratio of average degree of conflict graph, and (d) Retention ratio of P2P message complexity of scheduling contention. For (c) and (d), the smaller is better.

In a dynamic scheduling simulation experiment with a very large cut-off quantile of 0.95, the GCN-based sparse scheduler can achieve a lower collision rate that is only one-fifth to one-third of the statistical baseline, in random access networks, as illustrated by the retention ratio of average degree of the sparse conflict graph in Fig. 32 (a). For time-slotted networks, the GCN-based sparse scheduler has lower P2P message complexity as the network density increases, while the P2P message complexity of the statistical baseline has the opposite trend, as shown in Fig. 32 (b). This shows that the developed GCN-based sparse scheduler can work better than the statistical baseline in denser networks.




Fig. 32 Retention ratios for (a) average degree and (b) P2P message complexity of the sparse schedulers w.r.t. the vanilla local greedy scheduler vs. the average degree of the original conflict graph. The plots correspond to 500 scheduling instances with ER graphs, cut-off quantile of 0.95, and 300 time slots. Smaller retention ratios are better.

In summary, our GCN-based distributed link sparsification scheme reduces the scheduling overhead in wireless networks with orthogonal access. The GCN can encode the network topology into a parameterized local decision process to keep a link from contending if its chance of winning is low. The proposed scheme significantly reduces the scheduling overhead for time-slotted and random-access networks while retaining most of the network capacity. Thus, our approach can improve access, energy efficiency, and radio footprint of wireless multihop networks.

6.5 Learning-based Access Control

Users in a wireless network share a common channel, and efficient utilization of the available channel bandwidth is necessary for reliable communications. Medium access control (MAC) layer protocols are designed to provide the rules of channel access in a WLAN. In particular, the IEEE 802.11 protocol has emerged as the most popular MAC layer protocol for WLANs, mainly because of its simplicity. The protocol specifies a distributed coordination function (DCF)⁵⁸ technique that combines CSMA/collision avoidance (CA) with a binary exponential backoff (BEB) mechanism to minimize packet collisions among users. However, the DCF has been shown to achieve suboptimal channel bandwidth utilization in complex settings such as high load, dynamic network sizes, and degraded transmission conditions. To improve channel utilization, a class of works focuses on modifying the exponential backoff mechanism. Two popular approaches include multiplicative increase linear decrease (MILD)⁵⁹ and exponential increase exponential decrease (EIED).⁶⁰ These modified backoff mechanisms allow users to deviate from standard BEB in order to improve channel utilization in dense networks. Like BEB, the success of these approaches relies on the coordination between all users in terms of following the same backoff mechanisms. With the

growing number of diverse devices, a fully coordinated system where every device follows the same backoff mechanism might not be tenable in practice.

In this project, we consider an uncertain environment scenario where users in a wireless network are free to choose one of the three backoff mechanisms (BEB, MILD, and EIED). Moreover, users are allowed to choose the backoff parameters locally. In such a setting, we view the network from the perspective of an *intelligent user* that joins such a heterogeneous network. The intelligent user does not know the backoff mechanism and backoff parameters employed by other users. In such an uncertain setting, the intelligent user's objective is to choose its backoff mechanism and parameters only based on local observations to maximize a specified global utility. We consider two utilities in this work, namely, throughput and fairness.

We consider a wireless local area network with N nodes transmitting packets to a central AP. Inspired by the IEEE 802.11 protocol, nodes are modelled to follow contention-based mechanisms for channel access. The BEB mechanism enables efficient channel sharing in a distributed manner by specifying the number of time slots that each node must wait before a transmission attempt.⁵⁸ Under the BEB mechanism, a node, having a packet to transmit, draws the number of slots to wait from a uniform distribution in $(0, c - 1)$ where c is chosen based on the history of transmission attempts. The initial value of c for a node is chosen equal to the minimum contention window ($CW_{\min}$) ω , and after every collision, the value of c is doubled. In general, c is chosen as $2^j \omega$ after j successive collisions and saturates at a maximum permissible prespecified value. Once the node successfully transmits a packet, c is reset to ω . The key idea behind BEB is to use collision as an indicator for network congestion and to provide a randomized rule that adapts to the congestion.

We also consider two extensions of BEB, that is, MILD⁵⁹ and EIED.⁶⁰ These mechanisms modify the vanilla BEB by introducing new update rules for c for both cases, that is, collision as well as successful transmission. In particular, MILD increases the value of c to $\gamma_I c$ after every collision and decreases the value to $c - \gamma_D$ after successful transmissions, with γ_I and γ_D set to 1.5 and 1, respectively.⁵⁹ The use of linear decrease after successful transmission improves the performance over BEB when the network is heavily loaded. Unlike BEB, where the contention window resetting mechanism causes large variation in the contention window size, the slow change in contention window in MILD is useful when collisions are high. Moreover, an intermediate approach (EIED) between contention window reset (BEB) and linear decrease (MILD) has been proposed that uses exponential decrease of contention window after successful transmission. In particular, EIED updates c to $r_I c$ and c/r_D in case of collision and successful transmission,

respectively, for some $r_I, r_D \in \mathbb{R}^+$. Furthermore, parameters r_I and r_D are selected depending on traffic load and network size.⁶⁰

We represent N nodes indexed by $0, 1, \dots, N-1$. All nodes are assumed to be in the radio range of all other nodes, and there are no hidden nodes. The network is observed in discrete intervals $k = 1, 2, \dots$, each of duration T seconds. At each time step k , a node can select one of the three backoff mechanisms and can choose values for corresponding backoff parameters, that is, ω for BEB, γ_I and γ_D for MILD, and r_I and r_D for EIED. We represent these choices for node i at time step k as action a_i^k . The action can take either scalar or vector values depending on the problem setting. Within each time interval k , we assume that the backoff mechanism and corresponding parameters for all nodes remain constant.

We consider the problem of optimal action selection from the point of view of an intelligent node (here node 0). Node 0 seeks to optimize its action a_0^k without explicit knowledge of the actions selected by other nodes and, even more challenging, in a setting where these values might change from one interval k to the next. Actions of other nodes are modelled using a discrete-time stochastic process that can take values in a finite state space Ω . In selecting the right action, node 0 must rely exclusively on local information.

For each interval k , we define τ_i^k as the fraction of time that the i^{th} node transmits packets without collision. Furthermore, we use τ_{-i}^k to represent the total fraction of time in interval k that all nodes except the i^{th} node transmit without collision, that is, $\tau_{-i}^k = \sum_{j=0, j \neq i}^{N-1} \tau_j^k$. In our formulation, apart from knowing its own a_0^k , node 0 observes τ_0^k and τ_{-0}^k . Based on this, we define the observation vector $\mathbf{o}^k = [\tau_0^k, \tau_{-0}^k]^T$. Notice that $\mathbf{o}^k$ is a random vector that depends on $\{\mathbf{a}_i^k\}_{i=0}^{N-1}$ through the (stochastic) backoff mechanism.

Lastly, to discuss *optimal* backoff mechanism and parameter design, we need to focus on the notion of utility. In our case, we associate interval k with two instantaneous network utility functions $u(\mathbf{o}^k, N)$ defined as

$$u_{fair}(\mathbf{o}^k, N) = 1 - \left| \frac{\tau_0^k}{\tau_0^k + \tau_{-0}^k} - \frac{1}{N} \right|,$$

and

$$u_{thr}(\mathbf{o}^k, N) = \tau_0^k + \tau_{-0}^k.$$

The utility function $u_{fair}(\mathbf{o}^k, N)$ corresponds to the instantaneous fairness observed from the perspective of node 0. For interval k , ratios $\frac{1}{N}$ and $\frac{\tau_0^k}{\tau_0^k + \tau_{-0}^k}$ represent desired and observed share of channel occupancy, respectively, for node 0. Fairness is maximized when the observed share of channel time is close to the desired share, and hence $u_{fair}(\mathbf{o}^k, N)$ achieves the maximum value of 1 when node

0 achieves $\frac{1}{N}$ share of the total channel time. Apart from the fairness utility, we also consider the normalized throughput $u_{thr}(\mathbf{o}^k, N)$ as a utility function. Both utility functions take values in $[0,1]$.

With this notation in place, we formally state our optimal action selection problem:

Given N and the history $\{\mathbf{o}^k, \mathbf{a}_0^k\}_{k=1}^{K-1}$, determine the action $\mathbf{a}_0^K$ to maximize $\mathbb{E}[\sum_{k=K}^{\infty} \alpha^{k-K} u(\mathbf{o}^k, N)]$.

At every discrete-time instant K , we assume that node 0 has knowledge of the network size N and it has access to past action-observation pairs $\{\mathbf{a}_0^k, \mathbf{o}^k\}_{k=1}^{K-1}$. With this information available, at every time-step, our intelligent node seeks to maximize the discounted sum of utilities. In general, it seeks to maximize not only the immediate utility at time K but rather an infinite-horizon discounted sum, where $\alpha \in (0,1)$ establishes the relative importance of future utilities. Since $\mathbf{o}^k$ is a random vector, future utilities are random variables, as is their discounted sum. Hence, the objective of node 0 is to maximize the expected value of this discounted sum.

We propose to frame our problem from a reinforcement learning (RL) perspective and solve it via deep Q-learning. In general, an RL problem is formulated as a Markov decision process (MDP).⁶¹ MDPs are used to model discrete-time, stochastic control problems where an agent learns an optimal control policy by interacting with an environment. For our problem, node 0 is modelled as the RL agent and the wireless network acts as the environment. In particular, an environment under an MDP is defined as a tuple (S, A, P, r, α) where S and A are, respectively, the state and action spaces. Moreover, P models the state transition dynamics such that $P(s'|s, a)$ gives the probability that the system switches to state s' when action a is taken in state s . At every transition, $r(s, a)$, a stochastic reward function, represents the reward obtained if action a is taken at state s . Finally, α is the discount factor that takes value in $[0,1]$. Furthermore, an agent is modelled as a stochastic policy π , which is a probability distribution $\pi(\cdot | s)$ over action space at a given state s . For instance, at state s , $\pi(a|s)$ represents the probability of taking action a under policy π .

For our problem, the reward function is chosen to be either $u_{fair}(\mathbf{o}^k, N)$ or $u_{thr}(\mathbf{o}^k, N)$. The next state under the MDP framework only depends on the current state and current action. In our setting, node 0 does not have information about dynamics of other nodes (environment), which could depend on the history of more than one time interval. For this reason, at every time interval k , state is defined as the collection of observations $\mathbf{o}^k$ and action $\mathbf{a}_0^k$ over M previous time intervals. Moreover, the action space for our RL agent is either the set of backoff parameters

or backoff mechanisms depending on the setting. With these specifications, solving our problem is equivalent to finding an optimal policy under the MDP model.

To find the optimal policy, we rely on the deep Q-learning algorithm which is a model-free approach for solving RL problems. Deep Q-learning learns a mapping between every state-action pair and the corresponding Q value using a neural network called deep Q-network (DQN). The Q value for the state-action pair (s, a) is defined as the maximum discounted reward obtained by the agent when it takes action a in state s . Once trained, the DQN can be used to find the action that achieves the highest discounted reward at every state. It should be noted that, in our implementation, instead of using a simple multilayer perceptron to learn the optimal Q-function, we rely on recent modifications that have shown empirical improvements in terms of training stability and amount of data needed for training. Specifically, the rainbow agent⁶² incorporates six improvements (double DQN, prioritized replay, dueling networks, multistep learning, distributional RL, and noisy nets) that led to superior performance over vanilla DQN in gaming environments. We have also witnessed those benefits in the implementation of deep Q-learning for our problem.

At a generic time interval K , the communication of packets to the AP is governed by the choices of actions of every user. In the example illustrated in Fig. 33, users select different actions locally at time K . Different colors of users represent different actions selected by them. These actions selected by other users are not known to node 0. Instead, node 0 observes the operation of the wireless network $\mathbf{o}^K$ and, of course, knows its own $\mathbf{a}_0^K$. The concatenation of these observations for the last M time intervals are used as the input to a trained rainbow DQN (RDQN). The output of this neural network consists of a vector containing the Q-values corresponding to every possible action that node 0 can select (shown as bar plot for different actions). Node 0 will select as a_0^{K+1} the action that attains the largest Q-value at the output of the DQN (depicted by the red bar), while a_i^{K+1} for $i \neq 0$ evolves following a local policy. This same process is repeated and node 0 relies again on the neural network to select a_0^{K+2} , and so on.




Fig. 33 Schematic view of our proposed deep RL agent. Node 0 relies on a rainbow DQN to select the action for the next time interval based on historical local observations.

Next, we demonstrate the effectiveness of the proposed approach through a range of experiments. First, we investigate the ability of the proposed DQN to select optimal backoff parameters for node 0 when the backoff parameters of other nodes are modelled to follow a discrete-time stochastic process. Further, we explore a different setting where other nodes select their backoff mechanisms instead of backoff parameters using the same stochastic process. Next, we study the selection of backoff parameters by the DQN when the choice of backoff parameters for other nodes depends on the actions of node 0. Finally, we demonstrate the generalization capabilities of the DQN model with respect to varying network sizes.

We simulate a WLAN with N transmitting nodes along with an AP using the ns-3 simulator. Nodes in the network are programmed to select and follow one of the three backoff mechanisms. For all the experiments, we use a constant packet size of 1400 bytes and a constant data rate of 1 Mbps. The observation interval length T is fixed at 20 s.

As previously discussed, node 0 acts as an intelligent agent and uses a rainbow DQN for selecting its optimal choices. The DQN constitutes two coupled feed-forward neural networks, each having four layers with 32 nodes per layer and ReLU nonlinearities. Furthermore, the DQN model is implemented using PyTorch⁶³ library in python, and the model weights are trained using mini-batch gradient descent with a batch size of 32. Other DQN parameters such as experience replay buffer size, number of training episodes, episode length, discount factor α , and learning rate are selected separately for every experiment using grid-search based hyperparameter tuning.

In evaluating the performance of our proposed RL approach,⁶⁴ we compare it with the following four baselines:

- i) Optimal design (OPT): This benchmark assumes full knowledge of the underlying stochastic process and computes explicitly the expected value in our problem. Although unrealistic in practice, it sets the upper bound of what is achievable by any method.

ii) Random forest (RF)⁶⁵: The classifier is trained using the local observations $\{\mathbf{o}^k, \mathbf{a}_0^k\}_{k=K-M+1}^K$ as input and $\mathbf{a}_0^{*K+1}$ as target labels. The training labels $\mathbf{a}_0^{*K+1}$ are selected as the optimal actions that maximize the fairness utility at the next time step.

iii) Optimal constant policy (CP): Just like for OPT, we assume complete knowledge of the system. However, we restrain the selection of the action to be constant, that is, it cannot change across time steps. Improvements upon CP illustrate the value of an adaptive design.

iv) Standard policy (SP): Node 0 follows a static policy and selects the default backoff parameters for every backoff mechanism.

We begin by considering a scenario where all nodes employ the same backoff mechanism and the backoff parameters of other nodes follow a Markov process with two states. We choose specific values for backoff parameters for the two states in order to capture two extreme behaviors of nodes. For all mechanisms, parameter values in state 1 are chosen such that other nodes wait for a short time before every transmission, and state 2 corresponds to a scenario with a large wait time. At each discrete time step, the system switches states with probability p , and we analyze the cases where $p \in \{0.75, 1\}$. Moreover, to evaluate the RL agent's performance on multiple backoff mechanisms, we examine three separate cases corresponding to BEB, MILD, and EIED. For this scenario, we consider u_{fair} as the utility function.

For MILD and EIED, the action set for node 0 is dependent on the choices of two parameters. In these cases, we define an action as joint value of the two parameters, that is, $a_0^k = [\gamma_I^k, \gamma_D^k]^T$ for MILD and $a_0^k = [r_I^k, r_D^k]^T$ for EIED. The architecture of the DQN is kept the same for all three backoff mechanisms, with only differences present in the last layer to account for different sizes of action sets. We select N to be 10, that is, we have one intelligent node that uses DQN to select the optimal parameters and 9 other nodes. The discount factor is set as $\alpha = 0.9$, and the DQN is trained for 5000 episodes. At the start of each episode, node 0 chooses an action uniformly from the set of actions, and the trajectory evolves for a specified number of time steps. We fix the number of time steps in each episode to be equal to 50. Since the underlying process is Markovian, we consider a memoryless implementation with $M = 1$.

We can observe in Fig. 34 that, when $p = 1$, our proposed RDQN method and RF achieve optimal performance even in the absence of knowledge of the underlying Markov process for all the backoff mechanisms. The advantage of RDQN over the supervised RF cannot be appreciated in this deterministic process. The lack of

adaptability of CP and SP makes them lag behind. More importantly, for $p = 0.75$, RDQN is still able to learn the optimal policy while the performance of RF is degraded. As a consequence of the stochasticity of the process, similar observations $\{o^k, a_0^k\}$ can lead to drastically different optimal values for a_0^{*k+1} , thus confusing the supervised RF classifier.




Fig. 34 Performance evaluation of the proposed RL method (RDQN) for BEB backoff mechanism. Box plots for fairness utility over 500 episodes of 50 time steps for the five methods under consideration. A Markov process with $p \in \{0.75, 1\}$ is considered for the evolution of the backoff parameters of other nodes

We indeed observe the superior performance of RDQN over RF, CP, and SP for the cases when the backoff mechanism is chosen as EIED and MILD, as in Fig. 35.




Fig. 35 Performance evaluation of the proposed RL method (RDQN) for MILD (left) and EIED (right) backoff mechanisms

We next consider a setting where other nodes select their backoff mechanisms using the Markov process mentioned in the previous experiment. In this case, the two states represent the selection of BEB and MILD backoff mechanisms. At each time step, node 0 chooses its action from $\{BEB, MILD, EIED\}$. Throughout this experiment, we fix the parameter values for different backoff mechanisms as standard parameters. For example, if a node chooses to follow BEB, then its $CW_{\min}$

is fixed as 32. Similar to the previous experiment, we consider two cases where $p \in \{0.75, 1\}$. We select fairness as the utility function and network size $N = 10$. The DQN architecture and the training procedure are kept the same as in the previous experiment. Figure 36 portrays the average fairness utility over 500 episodes obtained by the five methods under consideration.




Fig. 36 Performance evaluation with varying backoff mechanisms

RDQN achieves utility closest to the optimal utility for both values of p . Moreover, RDQN clearly outperforms the supervised method RF when $p = 0.75$. CP, in this case, corresponds to the MILD mechanism, and SP corresponds to BEB. One can observe that a fixed choice of backoff mechanism does not yield utility as high as utility achieved by dynamic choices of RDQN.

In this next experiment, we investigate a setting where other nodes select their actions using a policy dependent on the action of node 0. We fix the backoff mechanism as BEB for all the nodes and study the optimal selection of the minimum contention window (MCW) for node 0. For this experiment, we fix N as 3, that is, we have one intelligent node along with two other nodes in the network. We start with a simple policy for the other nodes that punishes the repetition of actions selected by node 0. Under this policy, the other nodes are modelled to select their backoff parameters in two states. More specifically, they start in state 1 at the beginning of each episode, and if node 0 selects the same action for five consecutive time intervals, other nodes switch to state 2 and stay in the same state for all time.

At the beginning of each episode, node 0 selects its $\omega_0 \in \{32, 512\}$. For the other nodes, the parameters $[\omega_1, \omega_2]$ are fixed as $[128, 128]$ and $[32, 512]$ for state 1 and state 2, respectively. The different values are chosen in state 2 such that state 2 leads to low fairness utility irrespective of choices of node 0, thus, providing a mechanism of delayed punishment to our intelligent node. At every time interval, node 0 selects its action, that is, ω_0 in $\{32, 64, 128, 256, 512\}$. Under this setting,

the optimal policy for node 0 is not only to select the action that provides the highest utility at the current time but also to ensure that other nodes stay in state 1.

We use the same DQN architecture as in the previous experiments with inputs $\{\mathbf{o}^k, \omega_0^k\}_{k=K-M+1}^K$. As the other nodes change their states based on the choices of actions for node 0 in five previous time intervals, we select M as 5.

Figure 37 illustrates the utility obtained by RDQN and CP with $\omega_0 = 128$, for the first 20 steps in an episode. It can be observed that the RDQN policy gets high utility for four consecutive time intervals followed by a low utility in the fifth time interval. This apparently suboptimal choice in the immediate term is actually optimal in the long run since it avoids the system to fall back onto state 2. On the other hand, a constant ω_0 performs well for the initial five time intervals and acts poorly when the system switches to state 2 for all following time intervals. After training, the DQN learns a policy that selects ω_0 as 128 for four consecutive time intervals followed by 256 in the fifth time interval. The average fairness utility attained under this policy is 0.9645, which matches that of the optimal policy.




Fig. 37 Fairness utility obtained in the first 20 steps of an episode when other nodes follow the simple policy that punishes repetition

We next move to a more realistic setting where the other nodes select their actions independently using a local policy (LP). Under this setting, at each time step k , node 0 uses the DQN model to select its MCW ω_0^k in $\{32, 64, 128, 256, 512\}$ and the other nodes select their MCWs by following a preset policy based on local observations. More precisely, at every time step k , the i^{th} node updates its MCW from ω_i^k to ω_i^{k+1} following

$$\omega_i^{k+1} = \begin{cases} 2\omega_i^k & \text{if } \tau_i^k > \frac{\tau_i^k + \tau_{-i}^k}{N} + \epsilon, \\ \frac{\omega_i^k}{2} & \text{if } \tau_i^k < \frac{\tau_i^k + \tau_{-i}^k}{N} - \epsilon, \\ \omega_i^k & \text{otherwise.} \end{cases}$$

Here $\frac{\tau_i^k + \tau_{-i}^k}{N}$ is the fair share of channel time for node i in interval k . According to the local policy, a node doubles its MCW in the next interval if it achieves more than the fair share of channel time in the interval k . Similarly, the node reduces its MCW to half of the previous value when the node achieves less than the fair share of the channel time. Moreover, we limit the possible values of ω_i^k in $\{32, 64, 128, 256, 512\}$, and we choose the boundary values 512 and 32 if the updated ω_i^{k+1} overshoots 512 or goes below 32, respectively. A positive parameter ϵ is introduced to make the system stable and avoid frequent changes of MCWs. The local policy discussed here is inspired by an existing contention window adjustment algorithm.⁶⁶

It should be noted that under the local policy, the choice of MCW for other nodes is dependent indirectly on the choice of MCW for node 0. The dependence comes from the fact that the choice of MCW for node 0 affects its occupancy of channel time τ_0 , which in turn affects the fair share threshold $\frac{\tau_i^k + \tau_{-i}^k}{N}$ for all other nodes. We train two separate DQNs for fairness and throughput utility with only current observations, that is, $M = 1$. Moreover, the parameter ϵ is fixed as 0.05 for all nodes. We compare the performance of the proposed RDQN approach with approaches where node 0 uses CP and LP for selecting ω_0 .

Figure 38 illustrates the fairness and throughput utility obtained by the three approaches over 500 episodes. RDQN achieves the highest fairness utility among the three approaches with a mean utility value of 0.9829. Furthermore, RDQN outperforms LP and achieves the same utility as CP when we consider throughput as the utility. The throughput is maximized when all nodes select a constant action of 32, and learns to choose 32 at all time intervals.




Fig. 38 Performance evaluation with other nodes following a local policy. The mean utility for the three approaches are $\{0.9829, 0.9668, 0.9797\}$ and $\{0.7435, 0.7435, 0.7197\}$ for u_{fair} and u_{thr} , respectively.

Finally, we evaluate the generalization capabilities of the proposed approach when the network size changes. Specifically, we consider the Markov process analyzed earlier for the BEB mechanism with $p \in \{0.75, 1\}$ and vary the number of nodes $N \in \{4, 5, 6\}$. We train three separate models of RDQNs, namely RDQN₄, RDQN₅, and RDQN₆ for the three network sizes and a fourth model RDQN₄₊₆, which is jointly trained on $N = 4$ and $N = 6$. To jointly train the RDQN on two network sizes, we randomly select one $N \in \{4, 6\}$ at the beginning of each episode and keep the N fixed for that episode. Moreover, we append the value of N to the input of the RDQN, and thus the new input is $\{\mathbf{o}^k, \omega_0^k, N\}_{k=K-M+1}^K$. Similar to the first experiment, we select $M = 1$.

Figure 39 (left) illustrates the fairness utility over 500 episodes for $p = 1$. Moreover, we evaluate the trained models as well as OPT, CP, and SP, separately, on $N = 4$ and $N = 6$. It can be observed that the jointly trained model RDQN₄₊₆ performs closest to optimal for both network sizes and both values of p . However, the model trained on only one network size does not perform optimally on other network sizes. We compare the performance of models tested on $N = 5$ for both values of p in Fig. 39 (right). As expected, the RDQN model trained on 5 nodes performs optimally when tested on a 5-node network. One could observe that the model jointly trained on $N \in \{4, 6\}$ for $p = 1$ also achieves utility closest to OPT when tested on $N = 5$. For $p = 0.75$, the jointly trained model achieves lower than the optimal utility. In this case the environment becomes more stochastic. For $p = 1$, the jointly trained RDQN is able to transfer knowledge to network sizes not observed during training.




Fig. 39 Comparison of utilities obtained by different models when the network size is varying. (left) Other nodes follow a Markov process with $p = 1$. (right) Utilities obtained by different models when evaluated on $N = 5$.

To summarize, we formulated the problem of optimal selection of backoff mechanism and backoff parameters as an RL problem. Within this framework, we leveraged a state-of-the-art DQN architecture to select the optimal value of the mechanism and parameters in a sequential manner exclusively from local

observations. Through ns-3 simulations, we showed that the proposed RL agent achieves close-to-optimal behavior and outperforms other learning-based and rule-based methods. Ongoing work includes: i) the combination of RL with graph neural networks to accommodate the case where every node is intelligent and can be trained based on its local observations, and ii) the incorporation of more realistic topologies that go beyond star connections to a common AP.

7. Secure Networks against AI-based Threats

In this thrust, we are developing methods to extract information about the network from completely encrypted networks. The driving idea is that since the nodes are transmitting energy, they are inadvertently leaking information. In Year 2, we developed three ideas to demonstrate this concept. First, we demonstrated that the network size can be estimated to a high accuracy for small networks without any prior training. Second, we demonstrate how this information can then in turn be used to estimate network flow rates using the same eavesdropper. Third, we demonstrated that the energy from the objects can be used to image the scene passively, demonstrating that the networks expose information about the scene.

7.1 Estimating Network Size

We consider the problem of estimating the size of an encrypted wireless network – defined as the number of active nodes in the network. We envision a scenario where an adversarial entity passively eavesdrops on the encrypted communication among nodes of a legitimate network, shown in Fig. 40, to estimate the network size. Should the adversary be successful, it can use the network size estimate as the first step in analyzing network operations, for example, to access probabilities and network protocol parameters.⁶⁷ Furthermore, it can build on this information to estimate network flows, thereby attaining a measure of the importance of each node in the network. The adversary equipped with this information can then devise and implement targeted attacks to disrupt and even completely bring down the transmission links of the legitimate network. In this project, we show that it is possible to estimate the network size down to an error margin of 1.38 nodes in networks of sizes of five to six nodes.




Fig. 40 The eavesdropper (black node) passively monitoring an adjacent wireless network

The main challenge is posed by the complete lack of identifiers under the total encryption assumption. In most communication protocols, MAC headers contain such identifiers to tag a transmitter and its transmissions. Without direct access to any such conventional identifier, the adversary must use the communication itself, the signals on the channel that reach the adversary, to tell transmissions and transmitters apart. Hence the adversary must separate the transmitters based solely on patterns extracted from their transmissions. These patterns may be caused by idiosyncratic errors due to process variability of the RF components, known as RF fingerprints, the channel variability across different node locations, or a combination of both. A different version of this problem has been proposed and studied in the literature in the context of authentication and intrusion detection. There, the transmission characteristics of the wireless nodes in the legitimate network are extracted, either as handcrafted features⁶⁸ or via deep neural networks,⁶⁹ before deployment, and new transmissions are compared against known transmissions. In this context, the network is not encrypted. If a node sends data whose characteristics do not match any of those in the list of known nodes, that node is considered an intruder. This is the description of a supervised learning problem where labeled data are available for training and comparison during deployment. Our problem, however, is one of unsupervised learning. This poses a second challenge. To estimate the size of the network, we must separate nodes we have never seen before. To the best of our knowledge, we are the first to attempt to solve this unsupervised problem.

To solve this problem and address the above challenges, we devised a wireless diarization pipeline inspired by a similar methodology in the speech processing community⁷⁰—there, the goal is to distinguish speakers and assign correctly detected speech segments to speakers. The wireless diarization pipeline, shown in Fig. 41, is composed of i) a neural network to generate embeddings of the input signals, ii) a distance metric to measure similarity between produced embeddings, and iii) a clustering method. The latter takes the distances between embeddings as input and outputs cluster memberships. The number of these clusters will then be our estimate for the number of transmitters in the data set.




Fig. 41 Wireless diarization pipeline. In this work, we skipped the signal detection part since our software framework dictates scheduled, non-overlapping, transmissions.

For embedding generation, we proposed and implemented NESTNet, a Siamese neural network. These neural networks are designed for discerning similarities and differences among inputs and the embeddings they produce minimize the distance among similar (i.e., from the same class) input, while maximizing the distance between dissimilar ones. In Figs. 43 and 44, we show the architecture of NESTNet along with a low-dimension projection of the data before and after embedding generation, illustrating, thus, its effect on the input. Next, we use the cosine distance metric to measure similarities between embeddings. Finally, we cluster through an agglomerative hierarchical clustering (AHC) algorithm. The number of resulting clusters is then our estimate of the network size. Both NESTNet and AHC are trained over a training and a validation set before being tested on data from unseen nodes. Our data set consists of packets of 10 orthogonal frequency-division multiplexing (OFDM) symbols, each containing 52 binary phase-shift keying (BPSK) symbols from 26 5G-grade IRIS SDR transmitters. Our data set has 450 such packets from every one of the 26 transmitters. The transmitters are in a 2-D array on top of Rice stadium while a single IRIS node at a 200 m distance is our receiver—the position of the transmitters is shown in Fig. 42. The transmitted symbol packets are identical across all nodes and thus do not carry any distinguishing information. Even in an encrypted network, there are still unencrypted pilots for network discovery/association and channel estimation. This configuration is a particularly challenging one: our methodology exploits subtle

hardware and channel variations. However, our SDRs' hardware is very precise, at a 5G standard level, and they are in tight proximity allowing for little channel variability among transmitters.




Fig. 42 Locations of transmit nodes and the receiver node. On the left, the position of transmitters and receiver in Rice Stadium; on the right, the transmit nodes in the 2-D array. Transmitters used in this experiment are shown by the white circles.




Fig. 43 Architecture of NESTNet, the proposed Siamese neural network




Fig. 44 The effect of NESTNet on data. (left) 2-D projection of raw data. (right) 2-D projection of the embeddings of raw data generated by NESTNet.

Despite these challenges, inherent to the problem but also stemming from the configuration of our resources, our methodology manages to estimate the size of unseen networks within an error margin of 1.38 nodes and a median error of one. Furthermore, a comparison between a pipeline without NESTNet and ours reveals the value of our proposed methodology. The estimation error for the case with no embedding generation goes to nearly three nodes. Despite this success, our method has its limitations. As the size of the test network increases, the already narrow space due to tight proximity and highly accurate hardware becomes even more congested, making size estimation even harder. This leads to higher errors in network size estimation as shown in Fig. 45.




Fig. 45 Estimated network size vs. true network size. The bars' height is the average estimated network size from different runs. The error bars represent the 95% confidence interval of the estimate.

Our next step is to focus on diarization accuracy. Rather than merely trying to estimate the size of the network, we will estimate the number of data packets assigned to each cluster/transmitter with high accuracy. We will increase the generalizability of our approach by collecting data from more transmitters, thus allowing for measuring the diarization accuracy of larger networks. Our goal is to have data from at least 100 transmitters. Instead of using only BPSK symbols, we intend to include higher modulation orders such as QPSK, 16-QAM, and 64-QAM. We will use oversampling at the receiver to capture nonlinearities of each transmitter leaking to the adjacent bands. Furthermore, we will emulate various degrees of signal overlap among different nodes and evaluate the effect of interference on the diarization accuracy. We will also replace cosine distance-based similarity measurements with Probabilistic Linear Discriminant Analysis (PLDA). We have noticed improved network size estimation on our previous data set and expect PLDA to help achieving high diarization accuracy over a more extensive data set. Finally, we will test and compare other clustering strategies, either density-based (e.g., HDBSCAN) or kernel-based.

7.2 Estimating Network Flow

We consider the problem of adversaries launching denial-of-service (DoS) attacks towards the legitimate network utilizing the network traffic flow information, such that adversaries could select to attack the most critical/busy paths in the legitimate network for most significantly disrupting network communication. In this project we aim to investigate the possibility of estimating the network flow distributions in a random-access network with collision avoidance mechanism (e.g., CSMA/CA network) with passive eavesdroppers. The goal is to design a passive estimation algorithm for the eavesdropper to estimate the flow packet rates (number of packets transmitted per unit time) for each flow in the network, without using any side information other than the eavesdropped RF signals over-the-air.

We formulate our network and eavesdropper model as follows: A single eavesdropper is passively monitoring the network as in Fig. 46, and we assume that the eavesdropper has high sensitivity such that it could hear the transmission of all the nodes in the network. It turns out that computing the flow packet rates is difficult for the eavesdropper as it faces two main challenges: i) The eavesdropper suffers from packet collisions at its receiver, as multiple non-interfering flows transmitting simultaneously would collide at the eavesdropper's receiver (the spatial-reuse phenomenon). So, the eavesdropper needs to estimate the flow packet rates from partial non-collided packet observations, yet the statistical relationship between packet rates with and without spatial reuse is complex for each flow. ii) The eavesdropper has no knowledge of the network interference pattern that determines which pairs of flows could transmit simultaneously. Lacking the information of the network interference pattern makes designing the estimator difficult as it is necessary for understanding the statistical relationship between the partial packet observation and the true flow packet rate.




Fig. 46 An illustration of the observed network and the eavesdropper setup

To solve these challenges, we have proposed a passive flow packet rate estimation algorithm utilizing some prior analysis over the stationary throughput distribution of the CSMA/CA network.^{71–73} Specifically, even though formulating a statistical model of the flow throughput distribution in an arbitrary CSMA/CA network is still an open problem, when the CSMA/CA network has only single-hop flows, a simple product-form expression of the flow throughput distribution can be found. Existing results^{70–72} imply the time-reversibility property of the stochastic process of network traffic status evolution. We use the time-reversibility of our network traffic model to design a simple closed-form estimator for the eavesdropper to estimate flow packet rates with only partial observation under spatial-reuse noise. We have also designed a co-estimation pipeline for estimating the network contention graph (which embeds the network interference patterns) in parallel. Our algorithm works robustly in single-hop-flow CSMA/CA networks with stationary flow traffic, regardless of the specific packet arrival distributions and transmission time distributions of each flow. We have further provided systematic convergence analysis over our estimator’s performance to prove that it is asymptotically optimal, and its asymptotic convergence rate under the metric of mean-squared error (MSE) is $O(1/T)$ with respect to the observation duration T .

We have also provided numerical simulations of our algorithm performance in the ns-3 platform to test our algorithm’s optimality and its empirical convergence rate. We test our algorithm under a 5-flow CSMA/CA network with IEEE802.11 MAC protocol. The five flows in the networks are mutually independent such that they could all transmit simultaneously (the worst case computationally as the spatial-reuse collisions at eavesdropper’s receiver are most severe). The result of MSE over the observation duration T is shown in Fig. 47. Based on the log-log plot, we have found that the best-fit line slope for each MSE curve is approximately 1, which coincides with the asymptotic convergence rate being $O(1/T)$.




Fig. 47 Algorithm’s MSE under different observation durations in an IEEE 802.11 network with five independent flows (log-log plot on the right)

We have designed a centralized algorithm in which a single eavesdropper takes charge of monitoring the whole network. The scalability of the centralized algorithm is limited since the algorithm would have high computational complexity when the network size becomes large, and the performance would become worse due to heavy spatial-reuse collisions. In the future, we will focus on extending the current work to a distributed setting where multiple eavesdroppers monitor the network collaboratively, such that each eavesdropper only takes charge of monitoring a manageable size of observed network nodes to enhance the scalability of the algorithm.

7.3 Estimating Scene Geometry

We consider the problem of estimating a scene’s spatial information—information about the scene’s physical geometry such as vehicles, buildings, or other obstructions—by taking advantage of ongoing wireless communications already present within the scene. Such spatial information is quite valuable, providing context of the battlefield to deployed allies in situations where drone imagery is not effective (such as imaging line-of-sight occluded targets or imaging in adverse weather conditions) while also benefiting the performance of the underlying communications network by informing decisions (such as scheduling and channel estimation).

An illustration of this project is shown in Fig. 48, wherein repeated uplink and downlink data transmissions between pairs of allied nodes (here depicted by a single overhead drone and multiple ground-based nodes) are taking place. There exists a line-of-sight path between each transmit/receiver pair, as well as many other paths (i.e., multipath) where the transmitted signal is reflected off objects in the scene (tanks). Under the assumption that the line-of-sight data transmission is decodable at the receiver, this signal can then theoretically be used as the reference signal and the multipath reflections containing the scene information can be harnessed using specialized image processing algorithms.^{74,75} An example of such a system setup, as well as the simulated imaging output, is shown in Fig. 49, where the blue lines represent locations of ground-based transmitter locations (belonging to a foot-soldier, for example), the red lines are corresponding receiver locations for a flyover drone collecting the data transmissions, and the black target points form the shape of a tank to be estimated.




Fig. 48 Example network topology for estimating scene geometry. Ongoing data transmissions between allied nodes is decoded and repurposed for imaging the tanks.




Fig. 49 Simulation environment (left) and simulated image reconstruction of a tank (right)

In the past year, significant steps have been taken towards experimentally demonstrating the operation of such a system using Rice’s RENEW infrastructure. We have addressed the ever-present issue of transmitter/receiver phase synchronization (a problem which is significantly complicated in the presence of multipath) using a graph-based method that jointly denoises the signal and compensates for phase offsets by taking advantage of smoothness in spatial and measurement domains (i.e., that the scene and network and varying slowly with respect to data transmissions) as

$$\min_{X, \phi} \left\| X \odot e^{-j\phi} \mathbf{1} - Y \right\|_F^2 + \lambda \left\| X - \underline{A}X \right\|_F^2 + \beta \sum_{i=1}^N |\nabla x_i(t)|,$$

where the first term is a signal fidelity term that isolates the data matrix without the phase offset X , from the phase shift ϕ , the second term is a graph total variation term that promotes smoothness in the spatial domain, and the third term promotes sparsity in the time-domain gradients to denoise the signal.

The current thrust of this project is to collect a large and diverse experimental data set to evaluate our scene estimation technique. Rice is uniquely positioned among research institutions to collect data sets within a large network setting—hundreds of meters in scale—using the RENEW platform. Evaluating our technique on real data will provide useful insight into the validity of assumptions made in our model in real-world settings, such as location knowledge requirements and the effects of multiple-scattering events. Beyond this, once the scene geometry estimator’s performance is sufficient, we can finally incorporate the scene information back into the network engine to improve the overall performance.

8. Conclusions

Year 2 saw an increased joint engagement and activity between DEVCOM Army Research Laboratory and Rice researchers. The joint research discussions have revealed new directions to explore, and we plan to expand the productive collaborative relationships.

9. References

1. Karl NJ, McKinney RW, Monnai Y, Mendis R, Mittleman DM. Frequency-division multiplexing in the terahertz range using a leaky-wave antenna. *Nat Photonics*. 2015;9(11):717.
2. Ghasempour Y, Yeh C-Y, Shrestha R, Mittleman DM, Knightly E. Single shot single antenna path discovery in THz networks. *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*; 2020.
3. Ma J, Karl NJ, Bretin S, Ducournau G, Mittleman DM. Frequency-division multiplexer and demultiplexer for terahertz wireless links. *Nat Commun*. 2017;8(1):1–8.
4. Ghasempour Y, Shrestha R, Charous A, Knightly E, Mittleman DM. Single-shot link discovery for terahertz wireless networks. *Nat Commun*. 2020;11(1):1–6.
5. Sutinjo A, Okoniewski M, Johnston RH. Radiation from fast and slow traveling waves. *IEEE Antennas Propag Mag*. 2008;50(4):175–181.
6. Ma J, Shrestha R, Adelberg J, Yeh C-YTY, Hossain Z, Knightly E, Jornet JM, Mittleman DM. Security and eavesdropping in terahertz wireless links. *Nature*. 2018;563(7729):89–93.
7. Yeh C-Y, Ghasempour Y, Amarasinghe Y, Mittleman DM, Knightly E. Security in terahertz WLANs with leaky wave antennas. *Proceedings of the 13th Conference on Security and Privacy in Wireless and Mobile Networks*; 2020.
8. Liu T, Prabhakaran V, Vishwanath S. The secrecy capacity of a class of parallel Gaussian compound wiretap channels. *IEEE International Symposium on Information Theory*; 2008 July 6–11; Toronto, Ontario, Canada. *IEEE Information Theory Society*; c2008. p. 116–120.
9. Kildishev AV, Boltasseva A, Shalaev VM. Planar photonics with metasurfaces. *Science*. 2013;339(6125).
10. Kitayama D, Hama Y, Goto K, Miyachi K, Motegi T, Kagaya O. Transparent dynamic metasurface for a visually unaffected reconfigurable intelligent surface: controlling transmission/reflection and making a window into an RF lens. *Opt Express*. 2021;29(18):29292–29307.

11. Chen L, Hu W, Jamieson K, Chen X, Fang D, Gummesson J. Pushing the physical limits of IoT devices with programmable metasurfaces. *USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*; 2021. p. 425–438.
12. Zelaya RI, Ma R, Hu W. Towards 6G and beyond: Smarten everything with metamorphic surfaces. *ACM Workshop on Hot Topics in Networks*; 2021. p. 155–162.
13. Akyildiz IF, Kak A, Nie S. 6G and beyond: the future of wireless communications systems. *IEEE Access*. 2020;8:133995–134030.
14. Zhang X, Tian Z, Yue W, Gu J, Zhang S, Han J, Zhang W. Broadband terahertz wave deflection based on c-shape complex metamaterials with phase discontinuities. *J Adv Mater*. 2013;25(33):4567–4572.
15. Guerboukha H, Amarasinghe Y, Shrestha R, Pizzuto A, Mittleman DM. High-volume rapid prototyping technique for terahertz metallic metasurfaces. *Opt Express*. 2021;29(9):13806–13814.
16. IEEE 802.11 Working Group. Enhanced throughput for operation in license-exempt bands above 45 GHz. *IEEE*; 2017. IEEE P802.11ay/D0.3.
17. Shi Q, Razaviyayn M, Luo Z-Q, He C. An iteratively weighted MMSE approach to distributed sum-utility maximization for a MIMO interfering broadcast channel. *IEEE Trans Signal Process*. 2011;59(9):4331–4340.
18. Chowdhury A, Verma G, Rao C, Swami A, Segarra S. Unfolding WMMSE using graph neural networks for efficient power allocation. *IEEE Trans Wirel Commun*. 2021;20(9):6004–6017.
19. Chowdhury A, Verma G, Rao C, Swami A, Segarra S. Efficient power allocation using graph neural networks and deep algorithm unfolding. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*; 2021. p. 4725–4729.
20. Chowdhury A, Verma G, Rao C, Swami A, Segarra S. ML-aided power allocation for Tactical MIMO. *IEEE Military Communications Conference (MILCOM)*; 2021. p. 273–278.
21. Calvillo CF, Sanchez-Miralles A, Villar J. Energy management and planning in smart cities. *Renew Sustain Energ Rev*. 2016;55:273–287.
22. Yu Q, Zhao Y, Zhang L, Yang K, Leng S. A fair resource allocation algorithm for data and energy integrated communication networks. *Mobile Inf Syst*. 2016.

23. Li Y, Sheng M, Yang C, Wang X. Energy efficiency and spectral efficiency tradeoff in interference-limited wireless networks. *IEEE Commun Lett.* 2013;17(10):1924–1927.
24. Matthiesen B, Zappone A, Besser K-L, Jorswieck EA, Debbah M. A globally optimal energy-efficient power control framework and its efficient implementation in wireless interference networks. *IEEE Trans Signal Process.* 2020;68:3887–3902.
25. Yang Y, Pesavento M. A unified successive pseudoconvex approximation framework. *IEEE Trans Signal Process.* 2017;65(13):3313–3328.
26. Long K, Li W, Jiang M, Lu J. Non-cooperative game-based power allocation for energy-efficient NOMA heterogeneous network. *IEEE Access.* 2020;8:49596–49609.
27. Li Y, Sheng M, Wang X, Zhang Y, Wen J. Max-min energy-efficient power allocation in interference-limited wireless networks. *IEEE Trans Vehicular Tech.* 2014;64(9):4321–4326.
28. Rajapaksha N, Manosha K, Rajatheva N, Latva-aho M. Deep learning-based power control for cell-free massive MIMO networks. 2021. arXiv preprint arXiv:2102.10366.
29. Xu D, Chen X, Wu C, Zhang S, Xu S, Cao S. Energy-efficient subchannel and power allocation for HetNets based on convolutional neural network. *IEEE Vehicular Tech Conf*; 2019. p. 1–5.
30. Eisen M, Zhang C, Chamon LF, Lee DD, Ribeiro A. Learning optimal resource allocations in wireless systems. *IEEE Trans Signal Process.* 2019;67(10):2775–2790.
31. Monga V, Li Y, Eldar YC. Algorithm unrolling: interpretable, efficient deep learning for signal and image processing. 2019. arXiv preprint arXiv:1912.10557.
32. Li B, Verma G, Segarra S. Graph-based algorithm unfolding for energy-aware power allocation in wireless networks. 2022. arXiv preprint arXiv:2201.11799.
33. Wu K, Guo Y, Li Z, Zhang C. Sparse coding with gated learned ISTA. *IEEE Int Conf on Learning Repr*; 2019.
34. Damosso E, Correia LM. COST Action 231: digital mobile radio towards future generation systems. Final report. European Commission, Directorate-General for the Information Society and Media; 1999.

35. Li B, Verma G, Segarra S. Graph-based Algorithm Unfolding for Energy-aware Power Allocation in Wireless Networks. arXiv preprint arXiv:2201.11799
36. McMahan B, Moore E, Ramage D, Hampson S, Aguera y Arcas B. Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics. PMLR. 2017;54:1273–1282.
37. Chen M, Yang Z, Saad W, Yin C, Poor HV, Cui S. A joint learning and communications framework for federated learning over wireless networks. IEEE Trans Wirel Commun. 2020;20(1):269–283.
38. Xia W, Quek T, Guo K, Wen W, Yang HH, Zhu H. Multi-armed bandit-based client scheduling for federated learning. IEEE Trans Wirel Commun. 2020;19(11):7108–7123.
39. Yoshida N, Nishio T, Morikura M, Yamamoto K. MAB-based client selection for federated learning with uncertain resources in mobile networks. IEEE Globecom Workshops; 2020. p. 1–6.
40. Boyd SP, Vandenberghe L. Convex optimization. Cambridge University Press; 2004.
41. Zhao Z, Verma G, Rao C, Swami A, Segarra S. Distributed scheduling using graph neural networks. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP); 2021. p. 4720–4724.
42. Joo C, Shroff NB. Local greedy approximation for scheduling in multihop wireless networks. IEEE Trans Mob Comput. 2012;11(3): 414–426.
43. Jaramillo JJ, Srikant R, Ying L. Scheduling for optimal rate allocation in ad hoc networks with heterogeneous delay constraints. IEEE J Sel Area Comm. 2011;29(5):979–987.
44. Hou I-H, Kumar PR. Utility-optimal scheduling in time-varying wireless networks with delay constraints. Proceedings of the 11th ACM International Symposium on Mobile ad hoc networking and computing; 2010 Sep. p. 31–40.
45. Xue D, Ekici E. Delay-guaranteed cross-layer scheduling in multihop wireless networks. IEEE ACM Trans Netw. 2012;21(6):1696–1707.
46. Hai L, Gao Q, Wang J, Zhuang H, Wang P. Delay-optimal back-pressure routing algorithm for multihop wireless networks. IEEE Trans Vehicular Tech. 2018;67(3):2617–2630.

47. Hsu Y-P, Modiano E, Duan L. Age of information: design and analysis of optimal scheduling algorithms. *IEEE Int Symposium on Info Theory (ISIT)*; 2017. p. 561–565.
48. Gupta M, Rao A, Visotsky E, Ghosh A, Andrews JG. Learning link schedules in self-backhauled millimeter wave cellular networks. *IEEE Trans Wirel Commun.* 2020;19(12):8024–8038.
49. Lee M, Yu G, Li GY. Graph embedding-based wireless link scheduling with few training samples. *IEEE Trans Wirel Commun.* 2021;20(4):2282–2294.
50. Gao J, Shen Y, Ito M, Shiratori N. Multi-agent Q-learning aided backpressure routing algorithm for delay reduction. 2017. arXiv preprint arXiv:1708.06926.
51. Zhongyuan Z, Verma G, Swami A, Segarra S. Delay-oriented distributed scheduling using graph neural networks. 2021. arXiv preprint arXiv:2111.07017.
52. Santi P. Topology control in wireless ad hoc and sensor networks. *ACM Comput Surv.* 2005;37:164–194.
53. Ramanathan R. Antenna beamforming and power control for ad hoc networks. *Mobile Ad Hoc Networking.* 2004;139–173.
54. Ray NK, Turuk AK. A hybrid energy efficient protocol for mobile ad hoc networks. *J Comput Netw Commun.* 2016.
55. Guha S, Basu PB, Chau C-KC, Gibbens R. Greenwave sleep scheduling: Optimizing latency and throughput in duty cycling wireless networks. *IEEE J Sel Areas Commun.* 2011;29(8):1595–1604.
56. Long J, Büyüköztürk O. Collaborative duty cycling strategies in energy harvesting sensor networks. *Computer-Aided Civ Infrastruct Eng.* 2020;35:534–548.
57. Lin L, Lin X, Shroff NB. Low-complexity and distributed energy minimization in multihop wireless networks. *IEEE/ACM Trans Netw.* 2010;18(2):501–514.
58. Bianchi G. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J Sel Areas Commun.* 2000;18(3):535–547.
59. Bharghavan V, Demers A, Shenker S, Zhang L. MACAW: a media access protocol for wireless LANs. *ACM SIGCOMM Comp Commun Rev.* 1994;24(4):212–225.

60. Song NO, Kwak BJ, Song J, Miller ME. Enhancement of IEEE 802.11 distributed coordination function with exponential increase exponential decrease backoff algorithm. *IEEE Veh Tech Conf.* 2003;4:2775–2778.
61. Sutton RS, Barto AG. Reinforcement learning: an introduction. MIT Press Cambridge; 1998.
62. Hessel M, Modayil J, Van Hasselt H, Schaul T, Ostrovski G, Dabney W, Horgan D, Piot B, Azar M, Silver D. Rainbow: combining improvements in deep reinforcement learning. 2017. arXiv preprint arXiv:1710.02298.
63. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, et al. PyTorch: an imperative style, high-performance deep learning library. *Adv in Neur Inf Process Syst.* 2019;32:8024–8035.
64. Kumar A, Verma G, Rao C, Swami A, Segarra S. Adaptive contention window design using deep Q-learning. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*; 2021. p. 4950–4954.
65. Abyaneh AHY, Hirzallah M, Krunz M. Intelligent-CW: AI-based framework for controlling contention window in WLANs. *IEEE Int Symp on Dynamic Spectrum Access Netw*; 2019. p. 1–10.
66. Bensaou B, Wang Y, Ko CC. Fair medium access in 802.11 based wireless ad-hoc networks. *IEEE Wrkshp Mobile and Ad Hoc Netw and Comp (MobiHOC)*; 2000. p. 99–106.
67. Mehrotra N, Graves E, Swami A, Sabharwal A. Minimax bounds for blind network inference. *IEEE International Symposium on Information Theory (ISIT)*; 2021. p. 1823–1828.
68. Brik V, Banerjee S, Gruteser M, Oh S. Wireless device identification with radiometric signatures. *ACM International Conference on Mobile Computing and Networking*; 2008. p. 116–127.
69. Al-Shawabka A, Restuccia F, D’Oro S, Jian T, Rendon BC, Dy J, Ioannidis S, Chowdhury K, Melodia T. Exposing the fingerprint: dissecting the impact of the wireless channel on radio fingerprinting. *IEEE Conference on Computer Communications*; 2020. p. 646–655.
70. Singh P, Vardhan V H, Ganapathy S, Kanagasundaram A. LEAP diarization system for the second DIHARD challenge. *Annual Conference of the International Speech Communication Association*; 2019. p. 983–987.

71. Liew SC, Kai CH, Leung HC, Wong. Back-of-the-envelope computation of throughput distributions in CSMA wireless networks. *IEEE Trans Mob Comput.* 2010;9(9):1319–1331.
72. Jiang L, Walrand J. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE ACM Trans Netw.* 2010;18(3):960–972.
73. Laufer R, Kleinrock L. The capacity of wireless CSMA/CA networks. *IEEE ACM Trans Netw.* 2016;24(3):1518–1532.
74. Sit YL, Zwick T. MIMO OFDM radar with communication and interference cancellation features. *IEEE Radar Conf*; 2014. p. 0265–0268.
75. Zhang JA, Rahman ML, Huang X, Guo YJ, Chen S, Heath RW. Perceptive mobile networks: cellular networks with radio vision via joint communication and radar sensing. *IEEE Veh Technol Mag.* 2021;16(2):20–30.

Appendix. Agenda of Virtual Workshop (April 2021)

ARL–Rice Agenda

April 19–21, 2021

All times are ET

April 19

3:00 – 3:10 pm ET	Prof. Ashu Sabharwal, Introductions and quick overview
3:10 – 4:00 pm ET	Prof. Santiago Segarra, “Graph-based Intelligence for Autonomous Network Management: An Overview”
4:00 – 4:35 pm ET	Arindam Chowdhury, “Graph Neural Networks and Algorithm Unfolding for Power Allocation”
4:35 – 5:10 pm ET	Dr. Zhongyuan Zhao, “Link Scheduling using Graph Neural Networks”
5:10 – 5:30 pm ET	Closing questions/discussion

April 20

2:00 – 2:30 pm ET	Prof. Edward Knightly, Overview
2:30 – 3:00 pm ET	Chia-Yi Yeh, “Secure Wireless from 100 GHz to 1 THz”
3:00 – 3:30 pm ET	Zhambyl Shaikhanov, “Adversarial Metasurfaces”
3:30 – 4:00 pm ET	Zhecun Liu, “Directional Modulation to Thwart Eavesdroppers”
4:00 – 4:30 pm ET	Closing questions/discussion

April 21

2:00 – 2:30 pm ET	Prof. Ashu Sabharwal, Overview
2:30 – 3:00 pm ET	Dr. Nico Barati, “Unsupervised RF Fingerprinting”
3:00 – 3:30 pm ET	Nishant Mehrotra, “Blind Network Inference”
3:30 – 4:00 pm ET	Yirong Cheng, “Network Traffic Eavesdropping from Imprecise Measurements”
4:00 – 4:30 pm ET	Closing questions/discussion

List of Symbols, Abbreviations, and Acronyms

2-D	two-dimensional
AHC	agglomerative hierarchical clustering
AI	artificial intelligence
AP	access point
ARL	Army Research Laboratory
BA	Barábasi-Albert
BCD	block coordinate descent
BEB	binary exponential backoff
BPSK	binary phase-shift keying
BS	base station
CA	collision avoidance
CP	constant policy
CSI	channel state information
CSMA	carrier-sense multiple access
CW	contention window
DCF	distributed coordination function
DEVCOM	US Army Combat Capabilities Development Command
DNN	deep neural network
DOD	Department of Defense
DoS	denial-of-service
DQN	deep Q-network
EE	energy efficiency
EIED	exponential increase exponential decrease
ER	Erdős-Rényi
FL	federated learning
GCN	graph convolutional network
GNN	graph neural network
ICB	independently coded baseline
LoS	line-of-sight
LP	local policy
LWA	leaky-wave antenna
MAC	medium access control

MCW	minimum contention window
MDO	multidomain operation
MDP	Markov decision process
MILD	multiplicative increase linear decrease
MIMO	multiple-input multiple-output
ML	machine learning
MLP	multilayer perceptron
MSE	mean-squared error
MSITM	metasurface-in-the-middle
MWIS	maximum weighted independent set
NP-hard	nondeterministic polynomial-time hardness
NSF	National Science Foundation
OFDM	orthogonal frequency-division multiplexing
OPT	optimal design
P2P	point-to-point
PAWR	Platforms for Advanced Wireless Research
PDF	primal-dual learning based on feedforward MLP
PDGNet	primal-dual graph convolutional power network
PER	packet error rate
PL	path loss
PLDA	Probabilistic Linear Discriminant Analysis
RDQN	rainbow deep Q-network
ReLU	rectified linear unit
RF	radio frequency
RFo	random forest
RL	reinforcement learning
SCA	successive concave approximation
SCADL	secure coding for angularly dispersive links
SDR	software-defined radio
SGD	stochastic gradient descent
SISO	single-input single-output
SNR	signal-to-noise ratio
SP	standard policy

TDS	time-domain spectroscopy
Tr	truncated
USCA	unfolded successive concave approximation
UWMMSE	unfolded weighted minimum mean squared error
WBS	wideband spatial
WLAN	wireless local area network
WMMSE	weighted minimum mean squared error
WSEE	weighted sum energy efficiency

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLD DCI
TECH LIB

17 DEVCOM ARL
(PDF) FCDD RLD
P BAKER
A KOTT
FCDD RLD FR
H MAUPIN
FCDD RLC
C BEDELL
FCDD RLC N
B RIVERA
A SWAMI
FCDD RLC NT
R DROST
E GRAVES
J PERAZZONE
C RAO
G VERMA
P YU
FCDD RLR
B HALPERN
J ALEXANDER
FCDD RLR EL
J QIU
FCDD RLR IN
C WANG
R ULMAN

2 DEVCOM CSISR
(PDF) J BOKSINER
A QURESHI