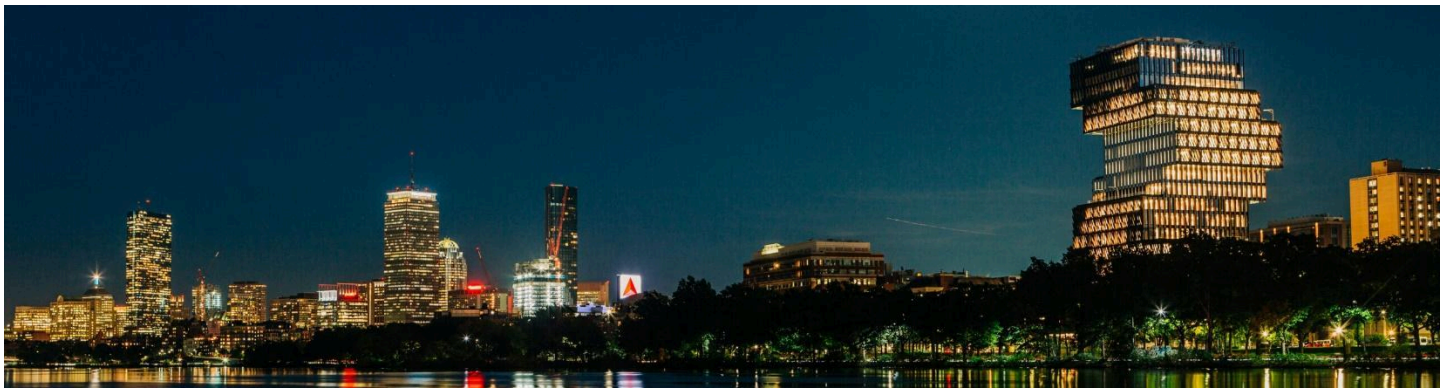


merif 2023

Midscale Experimental Research Infrastructure Forum - Boston University Center for Computing & Data Sciences, 665 Commonwealth Ave, Boston, MA 02215

AGENDA

Monday May 22, 2023

8:00 am	Breakfast - Room 1750 (Center for Computing and Data Sciences)
9:00 am	Welcoming Plenary - Room 1750
	Welcome from NSF - JD Kundu, Deputy Asst. Director, CISE Welcome from Boston University – Abraham Matta Recognition of Organizing Committee – Glenn Ricart Evolving Role of Midscale Infrastructure in Research, Education, and Public Service – Deep Medhi, NSF Program Director, CNS, CISE Roadmap to the Workshop – Glenn Ricart
9:45 am	News and Updates from Established MERIs
10:00 am	COSMOS – Dipankar Raychaudhuri
10:30 am	Morning Break
1100 am	POWDER – Kobus van der Merwe
11:30 am	FABRIC – Ilya Baldin
noon	Lunch - Room 1750

Merif 2023 Agenda Monday May 22, 2023 continued

1:00 pm	News and Updates from Newer MERIs
1:00 pm	SAGE - Nicola Ferrier 2
1:30 pm	AERPAW - Ismail Guvenc
2:00 pm	News and Sustainability Updates from Established MERIs
2:00 pm	CloudLab - Rob Ricci
2:20 pm	Chameleon - Kate Keahey
2:40 pm	Colosseum - Tommaso Melodia
3:00 pm	Afternoon Break
3:30 pm	News and Updates from Startup MERIs
3:30 pm	EduceLab - Corey Baker and Hugo Reyes-Centeno
4:00 pm	ARA - Hongwei Zhang
4:30 pm	National Internet Observatory- David Lazer
5:00 pm	Program Break
6:00 pm	Reception and Demos - Metcalf Trustee Center (1 Silber Way), 9th floor
8:30 pm	End of Day One

Tuesday, May 23, 2023

7:45 am	Breakfast - Room 1750 (Center for Computing and Data Sciences)
8:45 am	Tutorials One – Choose one of the following (3 tracks, 90 minutes each)
	SAGE - Sean Shahkarami (CDS 950)
	AERPAW - Rudra Dutta, Magreth Mushi (CDS 1646)
	Colosseum - Pedram Johari (CDS 1101)
10:15 am	Morning Break

Merif 2023 Agenda Tuesday May 23, 2023 continued

10:30 am | **Tutorials Two – Choose one of the following (3 tracks, 90 minutes each)**

POWDER - David Johnson (CDS 950)
COSMOS - Abhishek Adhikari, Jennifer Shane (CDS 1646)
Reproducibility (CDS 1101) -
Colosseum- Leonardo Bonati
AERPAW- Vuk Marojevic
CloudLab - Rob Ricci
Chameleon - Kate Keahey, Adam Cooper

Noon | **Lunch** 3

1:15 pm | **Tutorials Three – Choose one of the following (3 tracks, 90 minutes each)**

ARA - Sarath Babu (CDS 950)
FABRIC - Paul Ruth (CDS 1646)
Education (CDS 1101) -
POWDER - Dustin Maas
AERPAW - Ismail Guvenc
CloudLab - Rob Ricci
FABRIC - Mayuri Upadhyaya
Chameleon - Chandra Shekhar Pandey

2:45 pm **Afternoon Break**

3:00 pm **Tutorials Four – Choose one of the following (2 tracks, 45 minutes each)**

Experiments Spanning Multiple MERIs - Paul Ruth (CDS 1101)
Experiments Spanning Edge + Cloud - Mike Sherman, Adam Cooper (CDS 1646)

3:45 pm **Researchers and Students Read-Out**

Researchers and Students react to the workshop so far and tell their stories

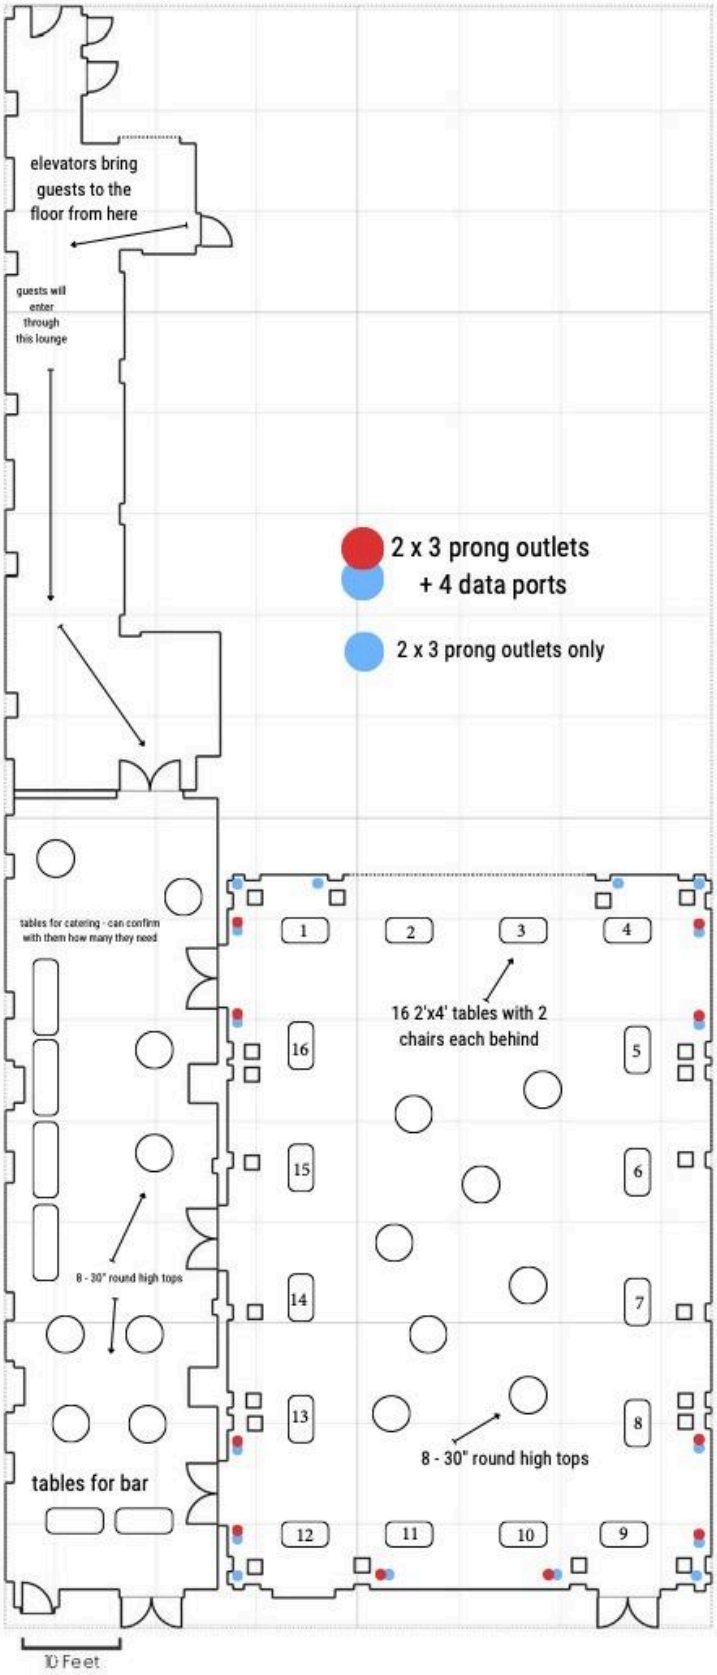
Merif 2023 Agenda Tuesday May 23, 2023 continued

4:30 pm	Challenges and Opportunities for MERIs
	Group Discussion of the Challenges Facing MERIs and the Opportunities that Lie Ahead
5:00 pm	Program Break
6:00 pm	Reception and Demos - Metcalf Trustee Center (1 Silber Way), 9th floor
8:00 pm	End of Day Two

Wednesday, May 24, 2023

8:00 am	Breakfast - Room 1750 (Center for Computing and Data Sciences)
9:00 am	Plenary Demos - best demos of previous two days
9:45 am	Workshop Feedback, Recommendations, and Report Writing
	MERI successes, stumbling blocks, key policies and services, factors leading to Researcher success, paths to inclusivity and diversity, publicity avenues, recognition For researchers, workshop improvements, etc.
11:00 am	Workshop Closing Session
Noon	End of workshop

Room Layout for Demo/Poster Session


elevators bring guests to the floor from here
guests will enter through this lounge
2 x 3 prong outlets + 4 data ports
2 x 3 prong outlets only
tables for catering - can confirm with them how many they need
16 2'x4' tables with 2 chairs each behind
8 - 30" round high tops
tables for bar
8 - 30" round high tops
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
10 Feet

Demos/Posters on Monday, May 22, 2023

Vote for your favourite here: https://tinyurl.com/merifdemo1

1. **(Tables 1 and 2) "Open RAN and 5G & Beyond Experiments on the POWDER platform," David Johnson, Dustin Maas, and Kobus Van der Merwe (University of Utah)**

   Abstract: This demonstration will show the POWDER experimental workflow and a number of example configurations that enable Open RAN and 5G & beyond research on the platform. First, we will demonstrate how to create Open RAN experiments on POWDER, using the O-RAN Software Community's near-real time RIC, and our NexRAN xApp. We will demonstrate both closed-loop RAN slicing and spectrum sharing use cases on real radios in different POWDER radio environments. We will further demonstrate how to use POWDER's TOTA (test orchestration, test automation) framework to automate these experiments across a parameter space. Second, we will demonstrate how to create POWDER experiments that deploy a complete 5G network using OpenAirInterface5G. We will show how users can experiment with 5G and OAI in a variety of POWDER environments, from passing IQ samples over sockets, to a bench-top setup with SDRs and wired RF connections, to indoor and outdoor over-the-air scenarios with SDRs and COTS UEs.

2. **(Table 3) "AERPAW Front-End Control Framework Automation - Resource Control and Orchestration for a Computing-Supported Physical Research Platform," Magreth Mushi and Rudra Dutta (North Carolina State University)**

   Abstract: As part of Phase-2 goals, AERPAW team has automated the platform's resource allocation, scheduling, and access control eliminating most of the operators' involvement. In this demo, we will focus on front-end automation as it relates to experimenters and their initial stage of executing virtual experiments. Specifically, the most necessary capabilities that we will focus on are: 1) Scheduling the usage of various resources, to support or for access by or on behalf of specific users of the platform, 2) Enabling and disabling access dynamically to various resources by use of appropriate credentials, 3) Dynamically configuring each resource to enable appropriate access and usage at any time. In this demo, we will demonstrate these three main functionalities that allow an experimenter to start and coordinate their interaction with the AERPAW testbed.

3. **(Table 4) "SLICES-RI Plain Orchestrating System (pos) - Reproducible Experiment Workflows by Design," Sebastian Gallenmueller, Serge Fdida, and Georg Carle**

   Abstract: Our poster describes a well-structured experiment workflow to design, execute, and document experiments. By creating experiments according to this structure, experiments become reproducible without additional effort for the experimenter, a property we call reproducibility by design. We further show our ideas to port the workflow to European testbeds as part of the SLICES-RI initiative and US-based testbeds.

4. **(Tables 5 and 6) "Self-Driving in Miniature Autonomous RC Cars using CHI@Edge," Rick Anderson (Rutgers University), Michael Sherman (University of Chicago)**

   Abstract: Machine learning and self-driving cars are both hot topics, but with a high barrier to entry for many students. We leverage Chameleon's CHI@Edge to demonstrate an edge-to-cloud workflow for device provisioning, data collection, training, and autonomous driving around a self-contained course. This framework acts as a building block for "hands-on" lab courses, even if students and hardware aren't local to each other, or to the instructor.

5. **(Table 7) "Using FABRIC to Evaluate TCP Congestion Control Algorithms for High-Speed WANs," Jose Gomez, Elie Kfoury, Jorge Crichigno (University of South Carolina)**

   Abstract: This poster presentation will discuss the results of a study on TCP congestion control algorithms over FABRIC, including the latest algorithms: the Bottleneck Bandwidth and Round-trip Time (BBR) version 1 (BBRv1) and BBR version 2 (BBRv2). BBRv1 and BBRv2 are rate-based congestion control algorithms that do not adhere to the Additive Increase Multiplicative Decrease (AIMD) control law, which has been used by TCP since the 1990s. Previous evaluations have been conducted on emulated environments (e.g., Mininet) and have not studied the performance of BBR on high-throughput (100 Gbps) high-latency WAN environments. The TCP congestion control algorithm is one of the main tuning elements to attain high throughput when transferring big science data (one of the main motivations to deploy Science DMZs). Thus, the results of this study may provide guidelines for network operators and research communities, such as the NSF Campus Cyberinfrastructure (CC*) and FABRIC communities.

6. **(Table 8) "Federating University Compute Clusters with EdgeNet," Berat Senel, Rick McGeer, Ufuk Bombar, Matt Hemmings, Ciro Scognamiglio, Olivier Fourmaux, Timur Friedman**

   Abstract: EdgeNet is a set of extensions to Kubernetes that adapt this free, industry-standard tool so that software can be deployed beyond a single large datacenter cloud to run in edge clouds. We show how EdgeNet can be used to interconnect instances of one ubiquitous type of edge cloud: the university compute cluster. Rather than allow high performance computing resources to stand idle, why not make them available to researchers at other institutions during down times? In exchange, when researchers at your institution need greater capacity, their workloads can scale to resources that are on offer elsewhere. This system can serve as the basis for frictionless collaboration between researchers at different institutions that have federated clusters, with simple sharing of containers that will operate anywhere and the ability to bring tools to the data rather than the other way around when the data is voluminous. This is preliminary work with a functioning prototype and we welcome your observations!

7. **(Table 9) “Open RAN Experiments with Hardware-in-the-loop on the Colosseum Wireless Network Emulator,” Leonardo Bonati (Northeastern University)**

   Abstract: In this demonstration, we will showcase how the Colosseum wireless network emulator facilitates repeatable experiments in the Open Radio Access Network (RAN) ecosystem. As a publicly accessible white-box platform with 256 software-defined radios and a massive channel emulator, Colosseum enables the design, development, and testing of solutions at scale in various deployments and channel conditions through virtualized and softwarized waveforms and protocol stacks. We will demonstrate the platform's ability to automatically initiate "batch jobs" that instantiate a cellular network with User Equipments (UEs) connecting to a RAN base station. The base station can allocate different network slices based on traffic demands and quality of service requested by the users and utilizes the Multi-generator (MGEN) toolkit to serve traffic to the UEs.

8. **(Table 10) “NSDF@CHESS Democratizing Data Access at the Cornell Light Source,” Giorgio Scorzelli, Valerio Pascucci (University of Utah)**

   Abstract: The National Science Data Fabric (NSDF) is a National Science Foundation project focused on democratizing access to scientific data across a diversity of disciplines, institutions, and people https://nationalsciencedatafabric.org. NSDF introduced a novel trans-disciplinary approach for integrated data delivery and access to shared storage, networking, computing, and educational resources. CHESS is a state-of-the-art seven-beamlines synchrotron radiation facility at Cornell University open to all scientists by peer-reviewed proposals. It serves 250 on-campus users and 1000 visitors yearly and currently stores about two petabytes of data. The NSDF@Chess collaboration addresses the problem of assisting CHESS in publishing its on-site storage capacity to external users, including those with extremely limited resources. In this demo, we show how NSDF is providing CHESS with stream services for just-in-time access and analytics during remote experimenters' data acquisitions: as soon as a beamline produces data, remote users can access it without delay. Data flows are also monitored to detect pain points and circumvent bottlenecks. Moreover, we demonstrate how the NSDF Entry Point facilitates CHESS’ connection and use of other distributed resources such as Chameleon Cloud; CloudLab; the PRISM Center at the University of Michigan; the Center For High-Performance Computing at the University of Utah; the Open Science Grid; the Open Storage Network. The live demonstration will provide the MERIF 2023 Workshop attendees with a hands-on experience of accessing and exploring in real-time very large imaging datasets without the need for any bulk data transfer.

9. **(Table 11) "Large-Scale Human Genome Sequence Analysis on FABRIC and CloudLab," Manas Das, Khawar Shehzad, and Praveen Rao (University of Missouri)**

   Abstract: Genomics is regarded as a Big Data science. Human genome sequencing has become economically feasible for conducting large-scale genomic studies and for use in clinical practice. Variant calling is a fundamental task that is performed to identify variants in an individual's genome compared to a reference human genome. This task can enable better understanding of an individual's risk to diseases and eventually lead to new innovations in precision medicine and drug discovery. However, variant calling on large number of human genome sequences requires significant computing and storage resources. In this demo, we will show how to leverage FABRIC and CloudLab testbeds for large-scale human genome sequence analysis. In particular, we will demonstrate the use of cluster computing (e.g., bare metal servers, virtual machines) and hardware accelerators (e.g., GPUs) to accelerate standard variant calling pipelines.

10. **(Table 12) "Dynamic Spectrum Access with IEEE 1900.5.2 Spectrum Consumption Model," Prasad Netalkar (Rutgers University)**

    Abstract: The future of wireless services and applications such as Augmented Reality, Internet-of-Things, and SmartCities will increasingly rely on Dynamic Spectrum Access (DSA) methods to manage spectrum resources in a rapid and efficient manner. In order to achieve this, we propose a novel spectrum management architecture, algorithm design, and experimentation framework that utilizes IEEE 1900.5.2 Spectrum Consumption Models (SCMs) to allow RF devices to indicate their spectrum usage and interference protection needs, as well as assess compatibility with other devices. The algorithm leverages Collaborative Interaction Language (CIL), developed by DARPA in the control plane, to coordinate and exchange spectrum usage information between multiple wireless domains. The simulation results and experimental validation demonstrates the advantages of adopting SCMs together with our proposed architecture enabling precise spectrum assignments in dynamic and dense communication environments.

11. **(Table 13) "Teaching on Testbeds: Ready-to-use Lab Materials for Teaching Computer Networks," Mayuri Upadhyaya, Fraida Fund (New York University)**

    Abstract: This demo will highlight a collection of ready-to-use open educational resources using FABRIC, CloudLab, and Chameleon for teaching computer networks, network security, and related topics. The demo will show some of the materials we have prepared, with a special emphasis on how testbed-specific features are used to improve the experience for students and instructors. We will also discuss planned future enhancements.

12. **(Table 14) "Using MFlib to Obtain Precise and Comparable Timestamps in FABRIC," Charles Carpenter, Yongwook Song, Hussamuddin Nasir, Pinyi Shi, Mami Hayashida, Jack Hancock, Zongming Fei, and James Griffioen (University of Kentucky)**

    Abstract: A key feature of the FABRIC network is the ability to accurately timestamp packets and user-defined events. Because clocks all across the FABRIC infrastructure are driven by a GPS-based time signal, timestamps are highly accurate and can be compared between racks in different locations. This demo will highlight how to use the FABRIC MFlib to take measurements enabled by FABRIC that are not possible in conventional networks. Examples include precisely measuring the one-way latency of a link, and timestamping user-defined events occurring at different nodes so they can be compared.

13. **(Table 15) "Artifact Evaluation With CloudLab," Rob Ricci (University of Utah)**

    Abstract: Artifact evaluation committees (AECs) are becoming common at top systems, networking, and security conferences. Papers can submit their code and data ("artifacts") to be evaluated for various "badges" showing the level to which their work can be reproduced. For many types of artifacts, using a testbed provides a good way to make sure that both author and evaluator have access to the same computing environment, easing the process and helping to achieve repeatable results. This demo shows how paper authors and AEC members can use CloudLab to make the AEC process simple and repeatable.

14. **(Table 16) "Sending and Receiving Encrypted Messages between Network Attached FPGAs in OCT," Suranga Handagala, Miriam Leeser (Northeastern University), and Michael Zink (University of Massachusetts, Amherst)**

    Abstract: In this demo, we will show how to use the Alveo U280 FPGAs in the Open Cloud Testbed (OCT). These FPGAs have network connections so the FPGAs can communicate directly through a network switch. The demo will consist of two FPGA nodes made available to the research community via CloudLab, where one will serve as a sender and the other as a receiver. The sender optionally encrypts a message before sending it to the receiver FPGA which receives and decrypts it using the AES128 standard.

Demos/Posters on Tuesday, May 23, 2023

Vote for your favourite here: https://tinyurl.com/2p89zfw2

1. **(Table 1) "srsRAN Multi-Node LTE SISO Experiment (SE1) in AERPAW's Development Environment," Keith Powell, Ozgur Ozdemir, Vuk Marojevic, and Mihail Sichitiu (North Carolina State University)**

   Abstract: This experiment sets up one node with the EPC and eNB and one or more nodes as UEs. Ping and iperf3 can be used between any UE and the EPC to determine latency measurements and channel capacity. The experiment outputs timestamped results each second including channel quality indicator, modulation and coding scheme, block error rate, and SNR. The experiment can be configured to use different frequencies and channel bandwidths. We will demonstrate how an experimenter can log in to our AERPAW portal and run this experiment remotely in AERPAW's emulation environment. Further description of this experiment can be found in AERPAW's user manual at https://sites.google.com/ncsu.edu/aerpaw-wiki/aerpaw-user-manual/4-sample-experiments-repository/4-1-radio-software/4-1-1-srsran-experiments/se1-multi-node-lte-siso

2. **(Table 2) "ARA: Wireless Living Lab for Smart and Connected Rural Communities", Sarath Babu, Hongwei Zhang, Daji Qiao et al. (Iowa State University)**

   Abstract: As a part of the NSF PAWR program, ARA is an at-scale platform for advanced wireless research, being deployed across the Iowa State University (ISU) campus, City of Ames (where ISU resides), and surrounding research and producer farms as well as rural communities in central Iowa, spanning a rural area with diameter over 60km. It serves as a wireless living lab for smart and connected rural communities, enabling the research and development of rural-focused wireless technologies that provide affordable, high-capacity connectivity to rural communities and industries such as agriculture. In the demo, we will demonstrate the experiment workflow (e.g., testbed access, resource reservation, experiment control, and data acquisition) and example experiments in ARA.

3. **(Table 3) "Performance Benchmarking of NDN-DPDK File Server on FABRIC," Junxiao Shi, Davide Pesavento, Lotfi Benmohamed (National Institute of Standards and Technology)**

   Abstract: NDN-DPDK is a set of high-performance Named Data Networking (NDN) applications including network forwarder, file server, and traffic generator. The NDN-DPDK forwarder is the first NDN router capable of reaching forwarding rates higher than 100 Gbps while running on commodity hardware and supporting the full protocol semantics. NDN-DPDK has been used in data-intensive science data distribution use cases, such as the distribution of Large Hadron Collider (LHC) datasets. We found that two of the midscale experimental research infrastructure projects, FABRIC and Cloudlab, to have sufficient hardware and network performance for running NDN-DPDK benchmarks. In this demo, we will show a file transfer scenario running on a three-node topology in the FABRIC testbed, in which files are downloaded from two NDN-DPDK file servers through an NDN-DPDK

forwarder. The scenario is controlled by a webapp, which gathers statistics about download completion time and transfer throughput.

4. **(Table 4) "Turning the Block in NYC and Still Getting 5G Coverage? mmWave Around-the-Corner Measurements for Dense Urban Deployment," Abhishek Adhikari, Shivan Mukherjee, Aahan Mehta, Manav Kohli, Dmitry Chizhik, Jinfeng Du, Rodolfo Feick, Reinaldo Valenzuela, Gil Zussman (Columbia University)**

   Abstract: Following prior outdoors measurements to quantify and understand propagation of 28 GHz mmWave signal in urban environments, a street-based measurement campaign in the deployment area of the PAWR COSMOS testbed in New York City was conducted in around-corner scenarios. Nearly 700 links and 10 million measurements providing results for 15 diverse urban environments are presented, consistent with the variability of city streets, at 3 different receiver elevations. Results can inform urban city planning on projects, such as Link5G, by providing essential and unique data important for cost-effective mmWave deployment.

5. **(Table 5 and 6) "Sage Edge AI," Sean Shahkarami, Rajesh Sankaran, Nicola Ferrier, Pete Beckman, Neal Conrad, Yongho Kim, Sergey Shemyakin, Joseph Swantek, Seongha Park, Bhupendra Raut, Ismael Perez, and Omar Zorob (Argonne National Laboratory)**

   Abstract: Sage is a new kind of NSF Mid-scale Research Infrastructure (MSRI) that supports "AI at the Edge". It provides cyberinfrastructure for running cutting edge machine learning and AI algorithms near sensors, instruments and large data sources with a focus on scientific applications. Sage is deploying cyberinfrastructure in environmental test-beds in California, Montana, Colorado, Oklahoma, and Kansas, in the National Ecological Observatory Network, and in urban environments in Illinois and Texas. This demo will provide a brief overview of Sage's infrastructure and will provide hands-on experience with Sage nodes, developing applications and accessing data. The goal is to illustrate how Sage can help solve your own problems or open new avenues for research.

6. **(Table 7) "Scalable Cybersecurity Experimentation with the Merge Testbed Platform," Brian Kocoloski, Chris Tran, Geoff Lawler, Joe Barnes, Michael Elkins, Yuri Pradkin, Calvin Ardi, Mahek Savani, Alefiya Hussain, Steve Schwab, Terry Benzel (University of Southern California)**

   Abstract: Cybersecurity researchers require scalable reconfigurable testbeds to emulate high fidelity experimental environments. The Merge Testbed Platform, developed at USC Information Sciences Institute allows researchers to define and deploy network topologies with thousands of nodes in a matter of minutes. In this demonstration, we highlight the scalability and fidelity of the Merge platform, and showcase an intuitive dashboard frontend designed to let interested potential users explore the system.

7. **(Table 8) "Data Collection and Validation of Auto ESN over Colosseum", Jiangqi Hu, Sabarish Krishna Moorthy (State University of New York, Buffalo)**

   Abstract: With growing research interest in data-driven decision-making, researchers have used different types of machine learning algorithms as a possible solution to solve sophisticated network control problems. Additionally, to verify the custom-built machine learning algorithms, researchers use community shared platforms for high-fidelity network emulation compared to the use of simulators. In this demonstration, we will showcase the (i) collection of data from Colosseum network emulator considering for different network topologies; (ii) use of collected data to train the Echo State Network (ESN) model with automated hyperparameter tuning (AutoESN) and (iii) prediction accuracy of the trained AutoESN model by testing it using data collected from Colosseum. Finally, we will present some future research directions and applications of AutoESN based network control scenarios.

8. **(Table 9) "Dealing with Changes: Resilient Routing via Graph Neural Networks and Multi-Agent Deep Reinforcement Learning," Sai Shreyas Bhavanasi, Lorenzo Pappone, Flavio Esposito (Saint Louis University)**

   Abstract: The networking community is increasingly exploring machine learning for tasks like routing, traffic prediction, and resource management. Reinforcement Learning (RL) has gained traction in network management, including packet routing. However, retraining a RL model is often required when network topology changes, hindering real-world deployment. In this paper, we propose two novel approaches to reinforcement learning-based routing that minimize flow set collisions and enable routing in dynamic conditions without retraining. We compare these approaches to other routing protocols, using various Quality-of-Service metrics, and share our findings.

9. **(Table 10) "Science Traffic as a Service," Ayomikun Gbadamosi, Jack Brassil (Princeton University)**

   Abstract: We are developing a prototype of a decentralized computing and networking system to create, collect and distribute a diverse collection of real and synthetic science traffic flows to experimental networking testbeds. We will first demonstrate the Science Traffic as a Service (STAAS) prototype on the FABRIC midscale research infrastructure. Our key project insight is that many science flows are already in transit at any moment on campuses. Using new campus cyberinfrastructure including passive optical Test Access Points, Network Packet Brokers, and data-plane programmable ethernet switches, STAAS will safely tap and forward copies of these flows onto the experimental testbed, while preserving both the timing integrity of the flows and the data privacy of their payloads. Large scale, high bandwidth experiments will be achieved by enlisting participation of many or all STAAS edge nodes on multiple campuses. By introducing a service-based model STAAS can reduce the burden of traffic generation on experimenters, improve experiment reproducibility, and help the operators of scientific instruments improve efficiency and broaden participation.

10. **(Table 11) "National Internet Observatory," David Lazer, Christo Wilson, and Dave Choffnes**

    Abstract: The Observatory aims to help researchers understand how people behave online and how online platforms structure what people see. This is accomplished by creating a large panel of individuals whose behaviors and interactions with platforms will be recorded. Data is collected from personal computers via a browser extension and from smartphones via apps; examples of data products include search queries, YouTube recommendations, mobile app usage data, and survey responses. We are building infrastructure to make this data and associated analysis tools available to researchers so they can incorporate them in their empirical studies. We expect to open up our data analysis infrastructure to outside researchers in the coming year.

11. **(Table 12) "GNU Radio Channel Sounder Experiment (GE2) in AERPAW's Development Environment," Ozgur Ozdemir, Anil Gurses, and Mihail Sichitiu (North Carolina State University)**

    Abstract: This experiment is a GNU Radio based channel sounder between two AERPAW Nodes one node acting as a transmitter and one node acting as a receiver. The experiment generates time-stamped relative received signal strength and signal quality. A pseudo-random sequence of bits is transmitted from a transmitter at a fixed node using GNURadio with a certain rate and center frequency. The transmitted signal is received and processed in GNURadio. The processing involves correlating the IQ samples with the sequence of transmitted bits and measuring the signal strength after the correlation. The signal strength as well as a measure of the signal quality are timestamped and logged in a file. Running the receiver at a portable node following a trajectory, the channel as a function of the location of the portable node can be measured. We will demonstrate how an experimenter can log in to our AERPAW portal and run this experiment remotely in AERPAW's emulation environment. Further description of this experiment can be found in AERPAW's user manual at https://sites.google.com/ncsu.edu/aerpaw-wiki/aerpaw-user-manual/4-sample-experiments-repository/4-1-radio-software/4-1-3-gnu-radio-experiments/ge2-channel-sounder

12. **(Table 13) "Open Educational Resources for Teaching Machine Learning Systems Deployment," Chandra Shekhar Pandey, Fraida Fund (New York University)**

    Abstract: Cross-curricular connections can be a powerful way to motivate students learning. In this demo, we will show a set of experiential learning materials for NSFCloud platforms that integrate cloud computing and machine learning concepts, and that are suitable for use in graduate-level coursework on either topic.

13. **(Table 14) "Open-Access Full-Duplex Wireless in the COSMOS Testbed", Manav Kohli, Tingjun Chen, Ivan Seskar, Harish Krishnaswamy, Gil Zussman (Columbia University)**

    Abstract: Full-duplex (FD) wireless enables simultaneous transmission and reception on the same frequency channel, which can substantially increase the throughput in wireless networks. To support experimentation with FD wireless at the higher layers of the network, we integrated four open-access FD radios in Sandbox 2 of the COSMOS testbed. These FD radios are available for use by the community, and this demo shows the baseline experiments which are currently available to researchers.

14. **(Table 15) "EduceLab: Infrastructure for Next-Generation Heritage Science," H. Reyes-Centeno, C. E. Baker, J. Balk, S. Smith, B. Seales (University of Kentucky)**

    Abstract: EduceLab, headquartered within the University of Kentucky's William S. Webb Museum for Anthropology, is a data-centric instrument platform for Heritage Science, applying the tools, techniques, and rigor of scientific measurement and inquiry to cultural artifacts, biological remains, and human-impacted environments. The EduceLab ecosystem, currently under development, includes four operational clusters: (1) BENCH laboratory equipment for advanced imaging (e.g. microscopy, computed tomography, photogrammetry); (2) MOBILE field-deployable instrumentation for in-situ data acquisition (e.g. optical scanners, LIDAR, ground-penetrating radar, unmanned aviation vehicles); (3) CYBER infrastructure for capturing, structuring, and processing large-scale data sets; and (4) a FLEX protean environment for envisioning, building, and testing custom instrument configurations co-designed by researchers and diverse stakeholders.

15. **(Table 16) "Experiments Spanning FABRIC and Chameleon," Paul Ruth (RENCI)**

    Abstract: In this demo, you will see an example of deploying experiments that span Chameleon and FABRIC. Specifically, we will run a single Jupyter notebook that deploys an isolated network that connects Chameleon servers to the FABRIC FABnet network. This `stitched` network can be used to directly communicate between servers on Chameleon and FABRIC, as well as, between servers on different Chameleon sites. This serves as an example of using FABRIC as a core network to connect many of the infrastructures that are discussed at MERIF.

Tutorial pre-work and links

SAGE

- **Pre-work:** The only pre-work is registering in our portal ahead of time: go to https://sagecontinuum.org/, click Portal in the upper right corner, click Sign In in the upper right corner and complete registration.
- The main prerequisites are basic experience with and access to Python, as well as some minor Git experience. The tutorial will start out as a crash course, guided tour of our portal and then transition to more hands-on work.

AERPAW

- **Pre-work:** Before the tutorial, you can install software on their laptops, access the tutorial environment, setup accounts, and request project membership. Pre-work instructions are available at AERPAW MERIF 2023.
- **Link:** Tutorial instructions are at AERPAW MERIF 2023.

Colosseum

- **Pre-work:** We encourage the interested attendees to sign up for an account through this online form. Please note that it requires to be submitted through their PIs.
- **Link:** Given the limited time of the tutorial, we are not planning for any in-session hands-on activities, but we encourage the attendees to follow the instructions on this page before the tutorial and come with questions.
- All the above information (and more) can be found on our website at www.colosseum.net which may be helpful for attendees to look at before they attend the tutorial.

POWDER

- **Pre-work:** Please create a POWDER account and join the MERIF project before the tutorial, following the instructions at POWDER MERIF 2023.
- **Link:** Tutorial materials are linked at POWDER MERIF 2023.

COSMOS

- **Pre-work:** Create an account and set up SSH keys following the instructions at COSMOS MERIF 2023 Signup.
- **Link:** During the tutorial session, you will be assigned to a group. Instructions for each group are available at COSMOS MERIF 2023 Instructions.

ARA

- **Pre-work;** Participants should have a laptop with Mozilla Firefox installed.
- **Link**: Refer to ARA User Manual.

FABRIC

- **Pre-work;** Sign up for a FABRIC account on the portal and also sign in on this Google form.
- **Link:** Slides for Hello, FABRIC and for wide-area slices.

Reproducibility (Colosseum, AERPAW, CloudLab, Chameleon)

- **Colosseum:** Link to sign up for Colosseum account and tutorial instructions.
- **AERPAW:** Refer to the reproducibility section of AERPAW MERIF 2023.
- **CloudLab:** We will review the material at CloudLab and Repeatable Research.
- **Chameleon:** Use https://bit.ly/merif-23-chameleon to register for a Chameleon Account, and join the project for MERIF 2023.

Education (POWDER, AERPAW, CloudLab, FABRIC, Chameleon)

- **POWDER:** Education tutorial and slides are linked at POWDER MERIF 2023.
- **AERPAW:** Refer to the education section of AERPAW MERIF 2023.
- **CloudLab:** We will review the material at CloudLab for Classes.
- **FABRIC:** We will demo this experiment. If you have not previously used FABRIC, you can do Hello, FABRIC as pre-work.
- **Chameleon:** We will demo this experiment sequence, with these slides. If you have not previously used Chameleon, you can do Hello, Chameleon as pre-work.

Experiments spanning multiple MERIs

- **Pre-work;** Sign up for a FABRIC account on the portal and also sign in on this Google form.
- **Link:** Slides for FABRIC facility ports.

Experiments Spanning Edge + Cloud

- **Pre-work;**.Use https://bit.ly/merif-23-chameleon to register for a Chameleon Account, and join the project for MERIF 2023.
- **Link:** Slides, Intro Tutorial, Edge-to-Cloud Tutorial