

Lawrence Berkeley National Laboratory

LBL Publications

Title

FABRIC: A National-Scale Programmable Experimental Network Infrastructure

Permalink

<https://escholarship.org/uc/item/8pj0n2v2>

Journal

IEEE Internet Computing, 23(6)

ISSN

1089-7801

Authors

Baldin, Ilya
Nikolich, Anita
Griffioen, James
[et al.](#)

Publication Date

2019-11-01

DOI

10.1109/mic.2019.2958545

Peer reviewed

FABRIC: A National-Scale Programmable Experimental Network Infrastructure

Ilya Baldin, RENC/UNC Chapel Hill, Anita Nikolich, IIT, James Griffioen, UKY,
Indermohan (Inder) S. Monga, ESnet, LBNL, Kuang-Ching Wang, Clemson University,
Tom Lehman, Virnao, Paul Ruth, RENC/UNC Chapel Hill

Abstract—FABRIC is a unique national research infrastructure to enable cutting-edge and exploratory research at-scale in networking, cybersecurity, distributed computing and storage systems, machine learning, and science applications. It is an everywhere programmable nationwide instrument comprised of novel extensible network elements equipped with large amounts of compute and storage, interconnected by high speed, dedicated optical links. It will connect a number of specialized testbeds for cloud research (NSF Cloud testbeds CloudLab and Chameleon), for research beyond 5G technologies (Platforms for Advanced Wireless research or PAWR), as well as production high-performance computing facilities and science instruments to create a rich fabric for a wide variety of experimental activities.

Index Terms—Testbeds, Next Generation Internet.



1 INTRODUCTION

FEW research investments have had as revolutionary an impact on the development of science, technology and society as the Internet. It has transformed our lives and led to a range of groundbreaking discoveries [1]. Its ability to evolve and support new uses derives from its programmable end systems that can exchange data and collectively form novel distributed networked applications. In essence, its programmability has been the key driver for the explosive growth in innovative network-based applications that now underpin our digital society.

However, this model is starting to show its age. Designers of new applications and services are forced to work around current architectural limitations, shoehorning their ideas into what’s feasible today.

For example, a fundamental system function like data caching in the network - a function critical for digital media streaming and data downloads - required substantial investments by large companies [2] to be effectively integrated into the Internet; to support detection of security events and provide policy enforcement, network providers had to place intrusive “middleboxes” in their networks, breaking Internet architectural principles and complicating the process of troubleshooting connection problems. Such limitations have also had a stifling effect on networking and systems research, reducing the diversity of ideas. A new platform that removes limitations on a researcher’s ability to study and prototype distributed architectures and applications is desperately needed by the systems, security, network and computational science communities, as well as by those exploring new architectures for scientific disciplines that now rely on massive-scale IoT devices, sensors, mobile and wireless services, and cloud services. Today’s Internet is

ill-suited for this due to architectural constraints limiting programmability to only the edge of the network.

To address this need, we are constructing an “everywhere programmable” nationwide testbed infrastructure called **FABRIC**, that provides unprecedented programmability along with large amounts of compute and storage uniquely built into the core of the network. FABRIC explicitly fosters innovation by allowing the research community to explore new inflection points in distributed systems architectures, enabled by the ability to process and store information inside the network. Multiple architectural solutions, distinct from today’s Internet, exist – for example trading off performance, functionality, usability, security and privacy. However, the impact of these trade-offs on the architecture and its economics are not well understood today, nor can they be practically explored with existing infrastructure and testbeds. FABRIC helps answer these research questions by giving researchers a tool that can measure outcomes and observe the implications of design trade-offs at scale.

When completed FABRIC will become a widely distributed instrument both for computer science research and for the many science domains that want to explore faster and more capable distributed computational and data infrastructures. It will be a highly-programmable, tera-bits-per-second core network interconnecting existing and future resources such as GENI [3], NSF Cloud [4], [5], 5G+ PAWR facilities [6], [7], campus networks and clusters, national HPC and data facilities, scientific instruments and public clouds, running parallel to production networks and dedicated to experimentation. It will be measurable with the ability to collect fine-grained packet-level telemetries. FABRIC’s capabilities will be achieved by deploying specially designed switching nodes into the footprint of the Department of Energy’s Energy Sciences Network (ESnet) and regional networks across the US, each with substantial amounts of compute and storage capacity far exceeding that of today’s commercial network equipment. Complemented

- FABRIC is supported in part by a Mid-Scale RI-1 NSF award under Grant No. 1935966.

by programmable edge nodes on campuses and national facilities, along with connections into the commercial cloud, it will enable experiments with application-controlled data storage, computation, and forwarding anywhere in the network – controlling an Internet-scale network connecting real facilities and users – thereby revolutionizing research in distributed systems, applications, protocols and services.

In the following sections we explore the community needs that led to the creation of FABRIC and how the FABRIC design answers those needs.

2 COMMUNITY RESEARCH NEEDS

We identified multiple problems and communities of researchers, working to solve those problems that would benefit from a deeply programmable, fast network dedicated solely to experimental activities.

2.1 Next Internet Architecture

The first Grand Challenge that FABRIC addresses is the development of the Next Internet Architecture. This covers an enormous swath of research that needs a platform on which to test new ideas and validate technical assumptions.

There is a large body of research targeted at “network programming” - developing new programming abstractions [8] and languages [9], [10] that allow experimenters to program specific behaviors into the network and achieve verifiable behavior, better predictability, resilience and security. The recent introduction of the P4 [11] language for programmable dataplanes opens new opportunities for testing this research in hardware and at scale.

In recent years, the Internet has witnessed a “tug of war” between providers and users, revolving around the desire of the former to secure their infrastructure and control the traffic in their network and the latter wanting to achieve best possible throughput for their applications. Providers continue to deploy ever-more intrusive “middleboxes” (e.g. Intrusion Detection and Prevention Systems [12], [13], [14], [15]) that block packets based on information carried in each packet, while users turn to sophisticated traffic tunneling solutions to avoid them. All the while these middleboxes violate the end-to-end principle of the IP architecture and make it difficult to diagnose traffic throughput problems for scientific applications.

The science DMZ architecture [16], for example, was designed as a solution to this problem w.r.t science data flows in the confines of today’s Internet architecture, however new solutions are needed to better define and abstract this “tussle” [17] between users and providers. Most importantly, researchers need a large “playground” on which to compare various solutions in which these tussles can be allowed to play out.

The convergence between wireless mobile and wide-area fixed communications infrastructures presents another complex area of architecture research. The wide-spread introduction of Network Function Virtualization (NFV) in 5G infrastructures, e.g. ONAP [18] provides an opportunity to study architectures bridging the two types of infrastructures allowing e.g. future 5G providers to move most of today’s “wireless edge” functions, into the core of the network at

varying geographic scales, offering novel services and reducing management overhead over the future 5G network. Testing these ideas requires geographic scale, allowing for separation of different functions in distance and across multiple domains of control.

Another example of architectures that require a lot of in-network state are ICN (Information Centric Network) architectures. Although the vast majority of Internet traffic today involves information retrieval, the existing Internet architecture is designed for point-to-point communication between hosts in the network. The ICN architectures were designed with a completely new set of protocols oriented at the task of identifying and retrieving needed information, and are incompatible with the TCP/IP stack that dominates today’s Internet [19]. This dominance also prevents wide-scale deployments of experimental ICN architectures in their pure form, instead forcing compromise solutions based on hybrid overlays.

2.2 Advanced Protocols Research and Novel Distributed Applications

Many of today’s expected scientific breakthrough discoveries are predicated on the ability of scientific applications to efficiently handle and manage massive data sets, i.e. one of NSF’s Ten Big Ideas - Harnessing Data Revolution. Even if the current Internet protocols could be made to transport these massive data sets from point A to point B quickly and efficiently, it is clear that future applications require more from the network than just transit. In order to efficiently and effectively handle today’s massive data sets, the network itself needs to offer advanced services that process the data in some way e.g., filtering, sampling, transforming, caching, storing, distributing, protecting, verifying, querying, etc.

Moreover, data used by today’s grand challenge problems comes from a large number of geographically distributed realtime sources: large instruments, like SKA (Square Kilometer Array) and LSST (Large Synoptic Survey Telescope), IoT (Internet of Things) devices and weather sensors, geographically distributed data repositories that must be aggregated to create a new combined data set e.g., genomics databases or social media databases, geographically distributed mobile/wireless endpoints that need to communicate with one another or the cloud like self-driving vehicle-to-cloud and vehicle-to-vehicle communication. All of this data will need to be transported, processed and stored by a wide variety of new protocols and applications. Designing protocols to solve these problems and testing them at large-enough scale requires a new continental-sized testbed that interconnects many of existing and future data sources and sinks.

Over the last few decades we have witnessed a shift between centralized and distributed data processing architectures a number of times - from mainframes to personal computers to centralized clouds to edge clouds and “fog computing”. Each time the application design followed the architecture, dictating specific ways in which applications must be implemented. It is critical that the experimental environment where these solutions are “baked-off” against one another does not mandate a specific approach to designing applications, instead offering an opportunity to design

new ones that are not constrained by predefined, and often unchangeable, locations where processing/computation is supposed to occur. By allowing applications to program the network across all layers of the protocol stack, the existing “narrow waist” of the Internet is removed, opening up the network so that applications can run their code everywhere along the path from the data’s origin to one or more of its destinations – including locations in the core that historically have been inaccessible to applications.

2.3 Intelligent Network Measurements, Control and Resilience

Autonomous network control and management, sometimes referred to as “self-driving networks” [20] is an emerging area requiring near real-time, high-fidelity telemetry linked to powerful in-network data analysis capabilities. In contrast with the opaque Internet of today, researchers require programmable measurements with an unprecedented level of monitoring capabilities. The high-volumes of data prevent it being transported to the “edge” for analysis and decision making, requiring data to be processed in-situ using significant in-network computational (e.g. CPUs, GPUs, etc.) capabilities, fully programmable dataplanes and significant storage resources. This will lead to new insights into performance measurement methods and data processing algorithms that will make future communications systems easy to manage and automate.

Network resilience is another facet of the vast field of research opportunities. For example IoT systems have largely been designed around the public cloud, streaming data to cloud services that present amazing data and analytics during good weather. However, during disruptive events such as hurricanes, these models are unable to operate without cloud access, and the model breaks down precisely at the time when IoT sensors are the most valuable. Power and network outages lead to failures in the network and partitioning of the parts that are working. What is needed is more intelligence placed throughout the network infrastructure to process and combine the data in place and provide local access, reducing the need to transport data back-and-forth – features not available in today’s networks.

2.4 Cybersecurity

Security and privacy were not explicitly built into the Internet. As a result, workarounds abound, and patchwork solutions at various layers of the Internet are cobbled together to make individual pieces more secure. Security and privacy researchers require the ability to model system vulnerabilities and attack scenarios, perform malware analysis at scale and reproduce effects of system interdependencies in a realistic environment with component fidelity not found in systems testbeds. This is critical for developing both safe and secure systems.

What is clear is that a highly-programmable network would present opportunities to respond to security events more rapidly and closer to their sources, instead of waiting for the problem to spread to its customers, but also, importantly, such a network offers a larger attack surface for hackers. Thus any new architecture based on in-network programmability paradigms must have built-in security

safe-guards that maximize the positives and mitigate or eliminate the negatives of these types of architectures, presenting rich opportunities for new network and application security research.

An experimental network environment of sufficient breadth and scale will also enable researchers to conduct social and behavioral experimentation, and combine human and technical modeling by reaching out to a number of campuses with potential opt-in users. For that it needs to allow production networks to peer with experimental topologies to e.g. steer or mirror production traffic into them for analysis. Of course, such “peering” between production and experimental networks must be done under controlled conditions with proper security procedures in place to prevent abuse.

2.5 Supporting Domain-oriented Cyber-Infrastructure Research

There are also significant opportunities to advance the state of the art when it comes to designing Cyber-Infrastructure (CI) solutions for various domain sciences.

Consider computationally-intensive research that utilizes high performance computing (HPC) facilities to carry out scientific computations, including researchers from the natural sciences, engineering, pharmacy, agriculture, business, and a growing number of other disciplines. Given the increasing demand for HPC resources, it is becoming increasingly difficult for national HPC facilities to keep up with demand. Consequently, a growing number of academic research institutions are exploring solutions to this problem by aggregating, and sharing access to the resources across institutions [21], [22]. Still others are looking to seamlessly integrate with the cloud to provide the elasticity they need [23], [24].

Researchers are looking for innovative approaches for efficient sharing of a wide variety of computational resources across institutions and commercial providers by programming the network to cache data it is likely to need the most, establishing optimized paths and using cross-layer protocols based on the type of data being moved, leveraging measurement information about the network to identify the best location to schedule tasks, or providing feedback based on history of past workloads that leads multiple institutions to make joint investments to build-out “hot spots” in the network.

Alternatively, consider research based on big data. Researchers of all types now have access to massive geographically dispersed data sets being developed as part of national, public, or privately hosted repositories. Examples range from satellite imagery used to map the Arctic [25], to LIDAR datasets used to build 3D maps to study geography [26], to genomics repositories used to study cancer and other diseases [27], [28], [29], to weather data used to study climate change [30], to traffic data used to study traffic patterns [31], [32], to linguistic databases used to study dialects across time and space [33], to social media feeds used to study human behaviors and interactions [34], [35], to business databases used, for example, to record company reports, SEC filings, and other government mandated reporting [36]. Researchers working with these massive data sets are all

facing a set of common challenges associated with efficiently analyzing the data originating from multiple sources and are coming up with their own solutions that work at the edge of the network [37]. A permanent solution to these problems requires an ability to deploy data processing capabilities in the network to operate on data while it is being collected, efficiently cache it, fuse and compress multiple datastreams on the fly, filter or deduplicate the data will open up new possibilities for scientific application design, improve efficiency of scientific applications and reduce time to discovery for many disciplines.

Similarly there has been an explosion in the number of environmental sensors and IoT devices used by researchers – not to mention video and audio IoT devices – to collect streaming time-series data that can give complete and highly precise information about the environment being studied. Being able to create optimized data streams between these instruments/devices and researchers will be critical to handle the massive amounts of information they are generating on a daily basis.

Real-time or near-real-time communications are becoming also critical whether it is to support emergency response, tele-medicine or streaming data for AI workloads. All these applications require efficient data transport and in-network processing with some timing guarantees to support a variety of feedback control loops with robust security and integrity guarantees. Exploring the various design points in a realistic environment can only be done with a mix of production and experimental resources and only at scales that approach the latencies and traffic volumes of today’s Internet.

3 FABRIC DESIGN

The national testbed we are constructing – FABRIC – will answer many of these and other needs from various parts of the research community.

FABRIC will consist of several components: (a) a Terabit network supercore and coast-to-coast 100Gbps core, with embedded FABRIC nodes, utilizing the fiber footprint and co-location spaces of ESnet, (b) a set of programmable edge nodes hosted by a number of campuses at US Universities and connected via regional network providers and (c) connections to a number of external existing and future cyber-resources contributed by the community. FABRIC links between nodes will use isolated optical waves or other dedicated capacity to provide predictable performance to experimenters. FABRIC will link these facilities providing a connected infrastructure parallel to available production networks, but with an explicit focus on supporting experimental activities in support of systems, network and other domain science research.

By Metcalfe’s law, the effect of a network is proportional to the square of the number of connected entities. For this reason the experimenter network topologies will include FABRIC nodes as well as any of the connected infrastructures, like NSF-funded PAWR and NSFCloud platforms, resources at scientific computing centers (TACC, SDSC, PSC, NCSA), instruments like LSST, public clouds like Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP) and others. The architecture is extensible: the set of resources described above is a starting point and is

expected to grow over time to encompass new scientific facilities, instruments and elements of cyber-infrastructure, as well as Bring Your Own Equipment (BYOE) extensions, allowing integration of experimenter-provided equipment. The planned topology with some of the anticipated connected resources is shown in Figure 1.

Key enabling technologies for FABRIC nodes are: (a) Hardware supporting the P4 programming language [38], OpenFlow [8], and eBPF [39], as well as other network programming abstractions which allow experimenters to program the data path behavior, even creating completely new packet formats incompatible with today’s standards, introduce new measurement capabilities into the network and process network traffic in novel ways at “wire speed”; (b) GPUs with tensor core support which can turn the network into a Machine Learning/AI-enabled system capable of rapidly introspecting on its own behavior for security, optimization or in reaction to user data coming from multiple sources for fusion and processing directly in the network; (c) 100’s of CPU cores, terabytes of memory and 10’s of terabytes of storage that can be ‘put in the path of the packets’ to provide support for in-network data processing, buffering or hosting of applications inside the network.

Cross-layer experimentation will be supported by the integration of Dense Wavelength-Division Multiplexing (DWDM) functionality in supercore nodes allowing users to observe the optical layer.

Unlike other experimental facilities, FABRIC infrastructure will not be completely isolated from the Internet. While this isolation, as implemented e.g. in GENI [3] or the National Cyber Range (NCR) [40], led to improvements in reproducibility of networking and security experiments, it severely limited what experimenters could do on the infrastructure, isolating them from the rest of the world. To overcome this problem, FABRIC will offer programmable peering with production networks in multiple locations, allowing experimenter topologies to be joined with production networks in a natural way, vastly expanding the realm of possibilities for the types of resources and users that can utilize the infrastructure. This feature could support a deployment of persistent FABRIC services that interact with real users or enable researchers to mirror or redirect traffic into FABRIC infrastructure for analysis, based on proper arrangements with campus IT personnel.

Public clouds, like AWS, Azure and GCP have become an integral part of modern computational science and the future cyberinfrastructure ecosystem cannot be imagined without them. Current research on them, however, is limited to off-the-shelf programmability offerings or black-box testing. In FABRIC we plan to provide users with the ability to peer their experimental topologies with public clouds using Internet2 Cloud Connect capability at speeds up to 100Gbps, enabling them to combine the best features of both experimental and production infrastructures for the needs of their applications or experiments. We will enable research into hybrid infrastructures by placing FABRIC virtual nodes inside cloud provider infrastructure and linking them to experimenter-owned cloud-based virtual infrastructure, application stacks or “serverless” applications.

FABRIC will enable researchers to measure all layers and aspects of hardware across the network footprint, including



Fig. 1. FABRIC core topology.

performance parameters, utilization and electrical power consumption. Some of the unique envisioned measurement capabilities not available elsewhere include nanosecond-precision time measurements of packet arrival times, all packet and error counts available in the hardware, custom measurement capabilities afforded by integrating P4-programmed dataplanes, and physical layer measurements from the optical hardware. Taken together the packet measurements will allow experimenters to precisely pinpoint the position of packets in analyzed flows within the network, a capability we call PacketGPS.

Data artifacts collected by experimenters in facilities distributed storage will be made available to the wider community through FABRIC experimenter portal with the purpose of supporting various kinds of meta-research on e.g. experiment reproducibility or enabling new types of network security research, by giving researchers a trove of data on user actions captured from FABRIC peerings with production networks for learning/analytics applications.

FABRIC is radically different from past experimental testbeds like PlanetLab [41], Emulab [42], GENI [3] and DETER [43]. The critical differentiators are the significant storage and compute resources and unprecedented level of programmability offered not just at the edge, where applications commonly reside today, but in the core of the network, its nationwide scale, extremely high and predictable wide-area performance and the ability to collect measurements with microsecond accuracy. Past projects created fairly isolated testbeds with limited connectivity to other resources. In contrast, FABRIC will become a national-scale programmable interconnect for a variety of experimental resources, providing different research communities and industry users a range of resource options, allowing them to deploy experimental software and test out production-ready architectures and services in a more realistic environment before offering them to end users. This will allow FABRIC to serve a much broader community of researchers working on cyber-infrastructure-related problems in a wide variety of scientific disciplines.

REFERENCES

- [1] T. Hey, S. Tansley, and K. Tolle, eds., *The Fourth Paradigm: Data-Intensive Scientific Discovery*. Redmond, Washington: Microsoft Research, 2009.
- [2] E. Nygren, R. K. Sitaraman, and J. Sun, "The Akamai network: a platform for high-performance internet applications," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, 2010.
- [3] R. McGeer, M. Berman, C. Elliott, and R. Ricci, eds., *The GENI Book*. Springer International Publishing, 2016.
- [4] R. Ricci, E. Eide, and The CloudLab Team, "Introducing CloudLab: Scientific Infrastructure for Advancing Cloud Architectures and Applications," *USENIX ;login;*, vol. 39, Dec. 2014.
- [5] J. Mambretti, J. Chen, and F. Yeh, "Next Generation Clouds, the Chameleon Cloud Testbed, and Software Defined Networking (SDN)," in *2015 International Conference on Cloud Computing Research and Innovation (ICCCRI)*, pp. 73–79, Oct 2015.
- [6] "Powder: Platform for Open Wireless Data-driven Experimental Research." <https://powderwireless.net/>.
- [7] "COSMOS: Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment." <https://cosmos-lab.org/>.
- [8] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, 2008.
- [9] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker, "Frenetic: A network programming language," *ACM Sigplan Notices*, vol. 46, no. 9, pp. 279–291, 2011.
- [10] B. T. Loo, T. Condie, M. Garofalakis, D. E. Gay, J. M. Hellerstein, P. Maniatis, R. Ramakrishnan, T. Roscoe, and I. Stoica, "Declarative networking," *Communications of the ACM*, vol. 52, no. 11, pp. 87–95, 2009.
- [11] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: Programming Protocol-independent Packet Processors," *SIGCOMM Comput. Commun. Rev.*, vol. 44, pp. 87–95, July 2014.
- [12] R. Sommer, "The Bro network intrusion detection system." <http://www.icir.org/robin/rwth/bro-intro.pdf>, 2007.
- [13] Zeek, "The Zeek network security monitor." <https://www.zeek.org/>.
- [14] "Snort Network Intrusion Detection Prevention System." <https://www.snort.org/>.
- [15] "Palo-Alto Networks." <https://www.paloaltonetworks.com/>.
- [16] "Science DMZ." <https://fasterdata.es.net/science-dmz/>.
- [17] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: defining tomorrow's internet," in *ACM SIGCOMM Computer Communication Review*, vol. 32, pp. 347–356, ACM, 2002.
- [18] "ONAP: Open Network Automation Platform." <https://www.onap.org/>.

- [19] J. Burke, A. Afanasyev, T. Refaei, and L. Zhang, "NDN impact on tactical application development," in *Proc. of MILCOM*, Oct. 2018.
- [20] N. Feamster and J. Rexford, "Why (and how) networks should run themselves," *arXiv preprint arXiv:1710.11583*, 2017.
- [21] B. von Oehsen, "The eastern regional network (ERN): Simplifying multi-campus research collaboration." <https://meetings.internet2.edu/media/medialibrary/2019/03/05/20190306-van-oehsen-emerging-platforms.pdf>, 2019.
- [22] "CCC secure share: An information superhighway for our industry." <https://www.cccis.com/secure-share/>.
- [23] Jeffrey Burt, "An Adaptive Approach to Bursting HPC to the Cloud." <https://www.nextplatform.com/2018/02/26/adaptive-approach-bursting-hpc-cloud/>, 2018.
- [24] M. Basilyan, W. Gorman, K. Binder, and A. Ma-Weaver, "Easy HPC clusters on GCP with Slurm." <https://cloud.google.com/blog/products/gcp/easy-hpc-clusters-on-gcp-with-slurm>, 2018.
- [25] University of Minnesota, "Polar Geospatial Center." <https://www.pgc.umn.edu/>.
- [26] "Kentucky Geological Survey." <http://www.uky.edu/KGS/>.
- [27] National Cancer Institute, "Genomic data commons." <https://gdc.cancer.gov/>.
- [28] SevenBridges, "Actionable informatics for biomedical research." <https://www.sevenbridges.com/>.
- [29] Broad Institute, "Genomics." <https://www.broadinstitute.org/genomics>.
- [30] NCDC, "Climate information." <https://www.ncdc.noaa.gov/climate-information>.
- [31] "Waze Traffic Reporting App." <https://www.waze.com/>.
- [32] Bureau of Transportation Statistics, "State transportation statistics." <https://www.bts.gov/product/state-transportation-statistics>.
- [33] "Linguistic Atlas Project (LAP)." <http://www.lap.uga.edu/>.
- [34] "Tweetsets: Twitter datasets for research and archiving." <https://tweetsets.library.gwu.edu/>.
- [35] Facebook, "Graph API." <https://developers.facebook.com/docs/graph-api>.
- [36] U.S. Securities and Exchange Commission, "Reports and publications." <https://www.sec.gov/reports>.
- [37] "XRootD software framework for fast, low latency and scalable data access." <http://xrootd.org/>.
- [38] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, *et al.*, "P4: Programming protocol-independent packet processors," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 87–95, 2014.
- [39] J. Kicinski and N. Viljoen, "eBPF Hardware Offload to SmartNICs: cls bpf and XDP," *Proceedings of netdev*, vol. 1, 2016.
- [40] B. Ferguson, A. Tall, and D. Olsen, "National cyber range overview," in *Military Communications Conference (MILCOM), 2014 IEEE*, pp. 123–128, IEEE, 2014.
- [41] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "Planetlab: an overlay testbed for broad-coverage services," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.
- [42] M. H. R. R. L. Stoller, J. Duerig, S. Guruprasad, T. Stack, K. Webb, and J. Lepreau, "Large-scale virtualization in the emulab network testbed," in *USENIX Annual Technical Conference, Boston, MA*, 2008.
- [43] J. Mirkovic, T. Benzel, T. Faber, R. Braden, J. Wroclawski, and S. Schwab, "The DETER Project: Advancing the Science of Cyber Security Experimentation and Test," in *Proceedings of the 2010 IEEE International Conference on Technologies for Homeland Security (HST '10)*, November 2010.

Ilya Baldin is the Director for Network Research and Infrastructure and RENCI/UNC Chapel Hill. His research interests include networked and distributed systems, signaling and control protocols and infrastructure information models. He has a PhD in Computed Science from North Carolina State University. He is also a Senior Member of IEEE. He can be contacted at ibaldin@renci.org.

Anita Nikolich Anita Nikolich is a Cybersecurity Researcher in the Computer Science Department at Illinois Tech, a Fellow at the CyberPolicy Institute at the Harris School of Public Policy at the University of Chicago, Co-Chair of the DEFCON AI Village and a member of the ARIN Advisory Council.

James Griffioen is a Professor of Computer Science and the Director of the Laboratory for Advanced Networking at the University of Kentucky. He also serves as the Director of the University of Kentucky Center for Computational Sciences. His research interests include future internet architectures, network measurement and instrumentation, network policy and economics, and cloud computing. He has a PhD in Computer Science from Purdue University, and can be contacted at griff@uky.edu.

Indermohan (Inder) S. Monga serves as the Division Director for Scientific Networking Division, Lawrence Berkeley National Lab and Executive Director of Energy Sciences Network, a high-performance network user facility optimized for large-scale science, interconnecting the National Laboratory System in the United States. Mr. Monga's research interests include developing and deploying advanced networking services for collaborative and distributed "big-data" science. He currently holds 23 patents and has 20+ years of industry and research experience in telecommunications and data networking. His undergraduate degree is in electrical/electronics engineering from Indian Institute of Technology in Kanpur, India, with graduate studies in Computer Engineering from Boston University.

Kuang-Ching Wang is Professor of Electrical and Computer Engineering and Associate Director of Research for the Watt Family Innovation Center at Clemson University. His research interests include software defined networking, Internet architecture and protocols, computing system architecture, to artificial intelligence. Wang has a PhD in Electrical Engineering from the University of Wisconsin-Madison. He's a Senior Member of the IEEE. Contact him at kwang@clemson.edu.

Tom Lehman is a Systems Architect for Virnao. His current research and development interests include multi-resource orchestration of networked and distributed cyberinfrastructure systems in service of big data driven domain science workflows. Lehman was previously the Director of Research at the University of Maryland Mid-Atlantic Crossroads and a Computer Scientist and Project Leader at the University of Southern California, Information Science Institute. He has an M.S. in Electrical Engineering from The Johns Hopkins University.

Paul Ruth is an Assistant Director of Network Research and Infrastructure at RENCI - University of North Carolina at Chapel Hill. He is a core contributor to NSF GENI and is a co-PI of the NSF Cloud Chameleon testbed leading its programable networking efforts. He earned his PhD in Computer Science from Purdue University in 2007. Contact him at pruth@renci.org.