Named Data Networking
Community Meeting 2020
Sept. 10-11, 2020, virtual event hosted by NIST

Session 1: NDN Trial Deployments

Publishing Genomics Datasets into NDN Testbed and Integrating with Cloud Workflows

F. Alex Feltus Clemson University
Susmit Shannigrahi Tennessee Tech University
Stephen Ficklin Washington State University
Rini Pauly Clemson University
Cole Mcknight Clemson University
David Reddick Tennessee Tech University
Tyler Biggs Washington State University
Cameron Ogle Clemson University

Background

The genomics community has made astronomical progress in recent decades. Thanks to medium-cost DNA sequencing machines, the genomics community is rapidly approaching the computational petascale at universities and research centers. For example, in 12 years the SRA repository at the National Center for Biotechnology Information (NCBI) has accumulated >42 petabytes of high-throughput DNA sequence data. There are many other similar genomic data repositories around the world including Japan's DDBJ, EU/UK's Ensembl, and NASA's GeneLab. All datasets are complemented with valuable metadata representing evolutionary relationships, biological sample sources, measurement techniques, and biological conditions.

While there is metadata overlap between repositories, a common data aggregation site for sharing data is needed — an NDN framework. Furthermore, it is common for users of compute resources in close proximity to use the same genomic datasets. By using the NDN framework with its built-in caching mechanism, commonly used datasets can exist closer to where they are being analyzed. The result is a framework for unprecedented data aggregation between heterogeneous metadata environment and accessibility that can be applied to many disciplines across computational science.

Named Genomics Data

As part of the NSF SciDAS project (#1659300), we have developed Pynome, a Python command line interface tool that provides the user with a way to download desired genome assembly files from the Ensembl database (`https://github.com/SystemsGenetics/pynome`). We have preprocessed all multicellular genomes and given them names based upon standardized evolutionary relationship metadata as well as repository metadata. Our data naming structure is semantically meaningful, hierarchical, and will make sense to the bioinformaticists across genomics communities. Here is an example data naming hierarchy for an NCBI SRA gene expression dataset: `/BIOLOGY/SRA/9605/9606/NaN/RNA-Seq/ILLUMINA/TRANSCRIPTOMIC/PAIRED/Kidney/PRJNA359795/SRP095950/SRX2458154/SRR5139394/SRR5139394_1>SRR5139394_1.fastq.gz`

These genomes are being published into the NDN scientific testbed that was created at Colorado State and since been extended to sites across the country. We are also publishing gene expression datasets from the 42+ PB SRA repository that includes NCBI-SRA-Animal_Example (0.16 TB; 36X2 SRA Samples; Human kidney) and NCBI-SRA-Plant_Example (1.14 TB; 475X2 SRA Samples; Rice leaves). These datasets were published using NDN-Python-repo and available to genomics researchers for download.

Workflow-NDN Testbed Interaction

We have containerized all NDN software and tested those containers in the Google Cloud Platform (GCP), Pacific Research Platform (PRP), and TACC Rodeo Kubernetes (K8s) clusters. We have pulled data from the NDN testbed into genomics workflows using the Data Transfer Pod (DTP; `https://github.com/SciDAS/dtp`) that can pull data from multiple sources and land the data onto a K8s namespace PVC. Finally, we have adapted the GEMmaker `https://github.com/SystemsGenetics/GEMmaker`) containerized workflow to run Nextflow gene expression mapping containers on each K8s system.

Advantages

Publishing datasets over NDN allows for popular datasets to be automatically cached in the network. The datasets can be cached in the cloud platform running workflow containers, significantly speeding up data retrieval. Additionally, the researchers can deploy jobs to cloud locations where data is already retrieved and available, reducing time for completing a workflow.

Linking DTP with NDN is an improvement over the current methods where the containers must be pre-configured with the URLs of the datasets. Once deployed, the containers stick to the pre-configured sources of data. With NDN, the data source can change based on network conditions, data availability, and the distance between computation and data.

Since NDN provides a more flexible way to cache and distribute data to Genomics workflows, we have also created an HTTP-NDN interface. When requested data is not available in the NDN testbed, we automatically download and publish the datasets into the NDN testbed using the naming scheme we describe above. The benefits of retrieving data over NDN are location independence, in-network caching, and flexible forwarding of requests.

Data-Centric Ecosystems for Large-scale Data-Intensive Science

Edmund Yeh Northeastern University

Data-centric networking is becoming increasingly important for large-scale data-intensive science. This talk will discuss recent results and future prospects for building data-centric distribution, caching, access, and analysis systems for major science programs.

We first discuss the NSF-funded SANDIE project, a collaboration among Northeastern, Caltech and Colorado State Universities, which aims to develop and deploy a Named Data Networking (NDN)-based platform for the Large Hadron Collider (LHC) high energy physics program, one of the world's largest big data applications. This project recently finished a major demo at SC19, which showed live that NDN can efficiently index and deliver LHC high energy physics data over a transcontinental layer-2 testbed (Boston-Denver-Los Angeles) at over 6.7 Gbps (using a single thread). The demo also showed how optimized caching algorithms can decrease download times by a factor of 10.

We next introduce the newly awarded NSF project N-DISE, led by Northeastern, Caltech, UCLA, and Tennessee Tech, which will produce a highly efficient and field-tested NDN-based petascale data distribution, caching, access, and analysis system serving major science programs including LHC high energy physics and the BioGenome and human genome projects. Building on the results of SANDIE, N-DISE will develop high-throughput caching and forwarding methods, containerization techniques, hierarchical memory management subsystems, congestion control mechanisms, integrated with FPGA acceleration subsystems, to produce a system capable of delivering LHC and genomic data over wide area networks at throughputs approaching 100 Gbps, with significantly decreased download times.

mGuard Project Overview

Lan Wang University of Memphis
Santosh Kumar University of Memphis
Lixia Zhang UCLA

This talk will give an overview of the new NSF funded mGuard project. mGuard aims to address two major data access challenges encountered by the NIH Center of Excellence for Mobile Sensor Data-to-Knowledge (MD2K) in its pursuit to share mobile health (mHealth) data among researchers who investigate a wide range of health and wellness issues. First, because wearable sensor data may expose privacy-sensitive information about a user, they should be accessed only by authorized users; currently this access control is largely handled manually, incurring high overhead and subject to human errors. Second, to enable real-time intervention for certain medical conditions, researchers need to retrieve and process the sensor data in real-time, which is not supported at this time. mGuard tackles the above challenges by utilizing the results from the NSF-supported Named Data Networking (NDN) initiative, in particular the solutions that automate the cryptographic key management for data access control (name-based access control, or NAC) and the solutions that enable real-time synchronization among distributed datasets (NDN Sync). First, mGuard utilizes and extends NDN NAC to automate fine-grained access control of confidential data to authorized researchers. Second, it utilizes NDN Sync to provide real-time data production notification; based on this, it enables applications to publish and subscribe to data in real time by directly using MD2K data names. These new capabilities will be deployed in the MD2K cyberinfrastructure.

FABRIC - Capabilities and Use-cases

Ilya Baldin RENCI/UNC Chapel Hill

FABRIC is a 4-year NSF-funded Mid-Scale construction project that by 2023 will deploy a world-wide experimental network infrastructure intended to support experimentation with stateful network protocols, applications and architectures. Using a combination of dedicated 100G and Terabit links the infrastructure will link FABRIC nodes together with experimental testbeds (NSF Clouds, PAWR facilities), HPC centers, scientific instruments and public clouds to create a rich environment for experimentation with a variety of applications and protocols that want to take advantage of in-network data processing, fusion and storage. This talk will briefly describe the status of FABRIC project and its expected capabilities in supporting this advanced experimentation

Session 2: ICN for Wireless Edge Networking

Introduction to the ICN-WEN Program and Learnings from NDN

Srikathyayani Srikanteswara Intel

The joint Intel/NSF ICN-WEN Center recently completed its 3 year tenure successfully. This short talk will give a brief overview of the setup of the center. The talk will also briefly touch upon some of our own experiences with NDN during this journey, and ways to increase its value proposition for industry.

SPLICE: Secure Predictive Low-Latency Information Centric Edge for Next Generation Wireless Networks

Srinivas Shakkottai Texas A&M University

The fifth generation of wireless communication promises to provide Gigabits per second data transfer rates

and communication delays of less than a millisecond. Significant challenges must be overcome in designing a system architecture such that data intensive and/or latency sensitive applications can obtain the information that they need for peak performance. Information-Centric Networking (ICN) has the potential to enable wireless efficiencies that are critical to support the strict guarantees desired by new applications such as virtual and augmented reality (VR/AR). The goal of this project is to design, develop and demonstrate SPLICE, a Secure Predictive Low-Latency Information-Centric Edge wireless network that will be able to provide information guarantees to such emerging applications. This talk will discuss the major results of the project and the way forward. The project is a collaborative effort across Texas A&M University, the Ohio State University, Purdue University, Washington University at St. Louis, and the University of Illinois at Urbana-Champaign.

Update on ICN-Enabled Secure Edge Networking with Augmented Reality (ICE-AR)

Jeff Burke UCLA REMAP

This short talk will provide an overview and update on the ICN-Enabled Secure Edge Networking with Augmented Reality (ICE-AR) project, which is supported by the NSF/Intel ICN-WEN program. ICE-AR is a collaborative effort of UCLA, Florida International University, and New Mexico State University. Finishing its third year, the project explores how Named Data Networking can support next generation wireless edge networking applications by integrating cross-layer wireless optimizations, data-centric security, acceleration-as-a-service, and new, ICN-inspired application concepts. This talk will provide a brief project overview and highlights from the last year, along with next steps for the work.

Session 3: Routing and Forwarding

On the Prefix Granularity Problem in NDN Adaptive Forwarding

Teng Liang Peng Cheng Laboratory
Junxiao Shi National Institute of Standards and Technology
Beichuan Zhang The University of Arizona

One unique architectural benefit of Named Data Networking (NDN) is adaptive forwarding, i.e., the forwarding plane is able to observe data retrieval performance of past Interests and use it to adjust forwarding decisions of future Interests. To be effective, adaptive forwarding assumes what we call Interest Routing Locality, that Interests sharing the same prefix are likely to take the same or similar forwarding path within a short time window, thus past observation can be an indicator of future performance. Since Interests can have multiple common prefixes with different lengths, the real challenge is what prefix length should be used in adaptive forwarding. The longer the common prefix is, the better Interest Path Locality, but the fewer future Interests it covers and the larger the forwarding table size. Existing implementations use static prefix length, which is known to have problems in dealing with partial network failures. In this presentation, I will introduce our work on dynamically aggregate and de-aggregate name prefixes in the forwarding table, so to use the prefixes that are the most appropriate under the current network situation. To reduce the overhead, we design mechanisms to minimize the use of longest prefix match in the processing of Data packets. Simulations demonstrate that the proposed techniques can make better forwarding decisions under partial network failures with significantly reduced overhead.

m-ASF - An Adaptive SRTT-based Forwarding Strategy for Mobile Environments

Muktadir Chowdhury University of Memphis
Alexander Lane University of Memphis
Lan Wang University of Memphis

Adaptive SRTT based Forwarding (ASF) strategy was originally developed to aid Hyperbolic Routing which could create persistent suboptimal paths. The strategy adapts to network dynamics by utilizing NDN's stateful forwarding and actively probing the data plane. However, ASF was designed for generally stable networks with static nodes. Therefore, it does not perform well in the presence of mobile nodes, where network topology varies frequently. In this work, we propose an improved ASF strategy for mobile networks, called m-ASF. The proposed strategy uses multiple paths when the primary path is experiencing failures. However, to prevent premature and frequent changing of faces, m-ASF employs anti-oscillation mechanism. We have done a comparative analysis of m-ASF and ASF. Our results show that m-ASF is more effective than ASF in retrieving data in dynamic environment. The performance gain is attributed to m-ASF's faster exploration of new paths and finer granularity in face-ranking.

NDN-DPDK: NDN Forwarding at 100 Gbps on Commodity Hardware

Junxiao Shi NIST
Davide Pesavento NIST
Lotfi Benmohamed NIST

Since the NDN data plane requires name-based lookup of potentially large tables using variable-length hierarchical names as well as per-packet state updates, achieving high-speed NDN forwarding remains a challenge. In order to address this gap, we developed a high-performance NDN router capable of reaching forwarding rates higher than 100 Gbps while running on commodity hardware. In this paper we present our design and discuss its tradeoffs. We achieved this performance through several optimization techniques that include adopting better algorithms and efficient data structures, as well as making use of the parallelism offered by modern multi-core CPUs and multiple hardware queues with user-space drivers for kernel bypass. Our open-source forwarder is the first software implementation of NDN to exceed 100 Gbps throughput, while supporting the full protocol semantics. We also present the results of extensive benchmarking carried out to assess a number of performance dimensions and to diagnose the current bottlenecks in the packet processing pipeline for future scalability enhancements. Finally, we identify future work which includes hardware-assisted ingress traffic dispatching, dynamic load balancing across parallel forwarding threads, and novel caching solutions to accommodate on-disk content stores.

Session 4: DARPA SHARE

DARPA Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE)

Mary Schurgot DARPA

The DARPA Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE) program is developing security and networking architectures and software for sharing information across multiple security levels. SHARE addresses several technology challenges to enable tactical edge, small unit information sharing between coalition and U.S. forces. Specifically, SHARE employs named data networking (NDN) to provide resilient secure store and forward capabilities to overcome brittle end-to-end connections.

PLI-Sync: Prefetch Loss-Insensitive Sync for NDN Group Streaming

Yi Hu	Perspecta Labs
Constantin Serban	Perspecta Labs
Lan Wang	University of Memphis
Alex Afanasiev	FIU
Lixia Zhang	UCLA

In this paper we explore solutions to robust group communication in disadvantaged wireless networks, which exhibit low bandwidths, high packet loss, and frequent/permanent network partitions. More specifically, we propose a group communication protocol based on Named Data Networking (NDN). By design NDN's in-network caching and stateful forwarding plane can help improve data delivery robustness in disadvantaged networks, however, when content is generated at a non-deterministic rate, an efficient, low-overhead synchronization protocol is needed to inform group members of the new content to fetch. To address this need, we propose Prefetch Loss-Insensitive Sync (PLI-Sync) protocol, specialized for group communication in highly disadvantaged networks. PLI-Sync combines optimistic content prefetching with selective content notification and addresses the challenges of distinguishing wireless packet losses from mobility-induced disconnections, and between data availability and retrievability. Our evaluations show that leveraging the interplay between NDN's stateful data plane and a low rate sync protocol can significantly reduce communication overhead compared to solely relying on sync protocols, while maintaining low data transfer latency and robust delivery in a variety of wireless conditions and traffic load settings.

NDN in DARPA SHARE

John DeHart Washington University in St. Louis

We give an update on our team's progress in the DARPA SHARE project and highlight some of the modifications and additions to NDN we have made. Some of the unique aspects of networking with tactical radios have lead us to some application and topology specific modifications to overcome challenges in intermittent, lossy and at times segmented network environments. These modifications include a new NDN sync protocol (ICT-Sync), and new NFD strategies to help reduce the network load.

Session 5: IoT/Edge

The "Decision Maker": NDN concepts for intelligent automation

Marie-Jose Montpetit Concordia University Montreal

IoT systems are fast becoming data driven intelligent system that can accomplish real time tasks but also combine with edge and cloud computing to deliver complex decision-making. This is enabled by the availability of powerful sensors, cameras and edge devices that constantly communicate among themselves and with other systems in the "cloud-edge" continuum. The move to data networking is happening across a wide swath of potential applications from intelligent cars to self managed farming systems. Yet, not all data captured by the sensors is at the same level of importance or dedicated to the same decision system. Also, the networked micro-controllers are better addressed by name than by obscure addresses. The presentation will introduce the "Decision Maker" a layered NDN enabled architecture for distributed intelligent automation. "Decision maker" combines named entities, interest/data information filtering and discovery as well as semantic web concepts to provide a framework for next generation IoT. The presentation will focus on the main elements of the architecture as well as introducing use cases to show its usefulness.

MICN, a New Perspective on Network Coding with Named Data Networking

Hirah Malik	INRIA
Cedric Adjih	INRIA
Claudio Weidmann	ENSEA
Michel Kieffer	CentraleSupelec

Overview

This submission is for a presentation in the NDN community meeting. Network coding (NC) can provide dramatic performance gains for NDN; we will present our new protocol designed to integrate it with NDN. Its goal is to maximize throughput. There are three parts: 1) a key mathematical problem that we identified that must be addressed for maximal performance, 2) our solution to it, and 3) finally, and most importantly its development as a full protocol. Emphasis will be on NDN protocol aspects, description of the design space, possible modifications, some approaches that work, others that work less well, and the ones we adopted.

Why Network Coding?

With network coding, one can add packets (XOR), or more generally make linear combinations. A consumer would send an Interest that is propagated in the network and could receive coded answers from several producers. For instance, upon receiving 3 coded Data $Q_1 = P_1 \oplus P_2$, $Q_2 = P_1 \oplus P_3$, $Q_3 = P_1 \oplus P_2 \oplus P_3$, it can recover the original content P_1, P_2, P_3 with elementary algebra (e.g. $P_3 = Q_1 \oplus Q_3$, etc.).

Thus a consumer could recover a lot of (non-redundant) content even without coordination of the providers. Without coding, redundancy is expected. As advocated by Montpetit et al. [2], NC ideally complements NDN features: decoupling the consumers from the content providers, and natural support for multiple paths. Thus, several protocols combining NC and ICN have been presented in the past, including NetcodNDN by Saltarin et al. [3], and our proposal, MICN [1].

What is the Problem with Multiple Interests?

NC is perfect to avoid (or limit) redundancy when only one Interest is sent. To increase throughput, a consumer can send several Interests in parallel: the challenge is to guarantee that the same number of non-redundant coded Data packets are retrieved. Indeed, without constraints on replies, the same Data packet could satisfy all Interests.

Assume that the Interests are numbered, and the corresponding coded Data for a specific indexed Interest is constrained to belong to a fixed set. Guaranteeing non-redundancy can be formalized as a mathematical problem (that we also denote "Sets Ensuring Linear Independent Traversals" (SELIT problem)).

What is our Novel Approach (MILIC)?

In [1], we introduced a special construction constraining the coded Data packets that satisfy a given Interest called MILIC (Multiple Interest for Linearly Independent Content). It solves the preciously explained problem by allowing multiple parallel Interests and ensures retrieval of linearly independent Data packet with each Interest. It also has other properties.

How did we Construct a new Protocol (MICN)?

MILIC is only a part of the NC solution for NDN. We designed a full protocol that integrates it, called MILIC-ICN (MICN) [1]. It has been implemented in a high-level simulator in Python to capture NDN's essential semantics and to experiment with NC-enabled variants. MICN achieves near-maximal throughput (max-flow) on tested examples.

NDN's general operation had to be modified, in the high-level semantics and the protocol functioning (PIT, FIB, caching). This includes classical NC adaptations, such as managing NC per subset of packets

("generations"), maintaining coded caches, performing decoding and recoding in the routers, and establishing conventions for names.

For high-level semantics, MICN revisits subtle questions such as merging similar Interests from different users: our current design chooses not to do so. An optional interest cancellation mechanism is also introduced.

Then support for MILIC-specific semantics requires additions: adding an index in Interest names, generating coded Data packets satisfying the MILIC constraints, pipelining multiple Interests, and proper support in case of packet losses.

Finally, getting protocol details right requires efforts: MICN is no exception. We found that a critical part was introducing careful queue management for satisfied pending PIT entries. We will also discuss the remaining issues and on-going work, such as the integration of sophisticated interest forwarding strategies.

References:

[1] H Malik, C Adjih, C Weidmann, and M Kieffer. MICN: a Network Coding Protocol for ICN with Multiple Distinct Interests per Generation.arXiv preprint arXiv:2007.01128, 2020.

[2] M.-J. Montpetit, C. Westphal, and D. Trossen. Network coding meets information-centric networking: An architectural case for information dispersion through native network coding.MobiHoc, (2):31–36, 2012.

[3] J. Saltarin, E. Bourtsoulatze, N. Thomos, and T. Braun. Adaptive Video Streaming With Network Coding Enabled Named Data Networking. IEEE transactions on multimedia, 19(10):2182–2196, 2017.

A Named Data Networking Architecture Design to Internet of Underwater Things

Qi Zhao University of California, Los Angeles
Zheng Peng The City College of New York
Xiaoyan Hong The University of Alabama

The Internet of Underwater Things (IoUT) advances our ability in exploring oceans, lakes and rivers through multiple communication technologies that connect stationary and mobile nodes underwater, at the surface and in the sky. However, characteristics such as low data rate, long propagation delay, energy-constraint, device mobility and sparsity, etc. of underwater communications remain as major challenges to the potential benefits that IoUT can bring to data availability and data sciences. This paper mainly focuses on further exploring how the Named Data Networking (NDN), a future Internet architecture, addresses the challenges of IoUT and can be adapted to potentially provide a simplified, efficient, and secure implementation of IUWT. The IoUT network and application semantics are aligned with the data-centric communication model of NDN, which enables operators to deploy and configure networks more easily, and developers to focus more on "things" and data underwater. The paper starts from introducing new challenges in IoUT and then illustrating in detail with simple examples to show how to employ NDN architecture to IoUT and how to enable additional functionalities required by IoUT. We also elaborate on the detail of our way of thinking and the remaining challenges while applying NDN to IoUT scenarios for future research directions.

Improving Existing Software Applications with a Practical and Secure NDN Publish/Subscribe Transport

Randy King Operant Networks
Jeff Thompson Operant Networks
Kathleen Nichols Pollere Inc.

Operant Networks is a startup founded in 2015 to initially address the problem of unreliable and insecure communications to renewable energy sites through NDN. We have since broadened our scope to include a wide range of energy industry applications, all with a strong focus on cybersecurity. Working in partnership with Pollere, UCLA, US Air Force, US Department of Energy, and key industry partners we have developed

an interesting NDN-based transport to replace an existing TCP/IP publication/subscribe framework in a cybersecurity intrusion detection system.

Historically, industrial networks are secured through perimeter firewalls to exclude malicious traffic. It is well known that these efforts are fraught with difficulty and unlikely to be completely effective at excluding determined attacks. In response, intrusion detection systems ('IDS') have become a growing market segment, aided by two realizations:

- Nearly all network traffic is encrypted and difficult to inspect in a firewall
- Advanced hackers tend to lurk in networks for weeks, learning the details of the site, before implementing their final attack; allowing time to detect their presence

Probably the most advanced IDS systems are based on distributed internal sensors: network packet sniffers deployed throughout a campus to monitor internal traffic as well as that leaving the site. A leading example of this type of IDS is Zeek (zeek.org), developed over decades as an open source project. It was developed from early on as a distributed system; a cluster of multiple packet sniffer instances on multiple physical servers collaborate to aggregate and analyze large amounts of distributed information.

Recently, Zeek re-architected their IDS framework to utilize a publish/subscribe model over discrete TCP/IP links secured by SSH. This communications paradigm provided Operant an opportunity to replace it with an NDN-based multicast pub/sub transport. In addition to the known security benefits of NDN, this transport demonstrates the first commercial use of fine grained and deployable trust schemas within a pub/sub framework. Our belief is that the intrusion monitoring solution must be much more secure than the underlying customer network or it provides little added security value.

Additionally, the NDN transport utilizes VLAN UDP multicast to remove many of the practical headaches of installing a distributed system within a TCP/IP network; reducing the need to maintain certificates or static IP addresses. A large IDS installation might utilize hundreds of sensor appliances, this could become an overwhelming maintenance burden in the TCP/IP world.

Most generally, we have learned that deploying NDN in an existing communications model (publish/subscribe) is a much easier path to commercialization than switching applications or their users from a TCP-based model to an Interest/Data exchange paradigm. We implemented the NDN Pub/Sub with an MQTT-like API (the most popular IoT Pub/Sub) and look forward to discovering other existing applications that can benefit from our lightweight secure transport and its complex trust models and broker-less pub-sub framework.

Session 6: Security/MAC

Rolling out NDN for DDoS Mitigation

Zhiyi Zhang	UCLA
Sichen Song	UCLA
Angelos Stavrou	George Mason University
Eric Osterweil	George Mason University
Lixia Zhang	UCLA

Distributed Denial of Service (DDoS) attacks have plagued the Internet for decades, but the basic defense approaches have not fundamentally changed. Rather, the size and rate of growth in attacks have actually outpaced carriers' and DDoS mitigation services' growth, calling for new solutions that can be, partially or fully, deployed imminently and exhibit effectiveness.

In this talk, we examine the basic functions in Named Data Networking (NDN), a newly proposed Internet architecture, that can address the principle weaknesses in today's IP networks.

We demonstrate by a new DDoS mitigation solution over NDN, Fine-grained Interest Traffic Throttling FITT, that NDN's architectural changes, even when incrementally deployed, can make DDoS attacks fundamentally more difficult to launch and less effective. FITT leverages the NDN design to enable the network to

detect DDoS from victim's feedback, throttles DDoS traffic by reverse its exact paths through the network, and enforces control over the misbehaving entities at their sources.

Our extensive simulation results show that FITT can throttle attack traffic with one-way time delay from the victim to the NDN gateway; upon activation, FITT effectively stop attack traffic from impacting benign flows, resulting in over 99% of packets reaching victims being legitimate ones. We further demonstrate that service providers may implement NDN/FITT on existing CDN nodes as an incrementally deployable solution to effectuate the application- level remediation at the sources, which remains unattainable in today's DDoS mitigation approaches.

Secure Sharing of Spatio-temporal Data through Name-based Access Control

Laqin Fan University of Memphis
Lan Wang University of Memphis

Named Data Networking (NDN) is proposed as a future Internet architecture, which provides name-based data publishing and fetching primitive. Compared to TCP/IP, NDN removes the need to manage IP address and has semantically meaningful names, NDN has stateful and name-based forwarding which simplifies the network stack, NDN's data-centric security and in-network caching are more efficient than cloud-based data delivery and encrypted channel. Name-based Access Control (NAC) is a content-based access control in NDN, which requires access control by encrypting content at the time of production directly without relying on a third-party (i.e., Cloud) to host the contents, and utilizes NDN's hierarchical naming convention to express access control policy and accomplish automating key distribution. As more sensitive data is generated and stored, there will be a need to share the data securely, such as mobile health (mHealth) data, data from smart-home systems. These data are generated over time and/or location, and could be collected as ongoing data streams. The data owners may want to share entire data or a time-location-specific subset of the data based on their preferences. Therefore, access control to those data must address the spatio-temporal attributes, and support secure real-time data sharing as well. An effective and secure access control solution is required to ensure that only authorized users can access to certain data. Inspired by Named-based Access Control model with data-centric security, we take into account the data attributes to make access decisions. By specifying access control policies with time and/or location attributes, we could limit data access to a given time bound and/or location area. In this work, we make three contributions: (1) we design an NDN's semantically structured naming convention to express fine-grained access control policy on spatio-temporal data, (2) for real-time data sharing, we deploy PSync to have a publish-subscribe module, (3) we have implemented a practical spatial-temporal data access control prototype based on NAC library in NDN codebase. Moreover, we evaluate the performance using Mini-NDN for different content key granularity.

A Full Data-centric Network Stack Integrating V-MAC and NFD

Mohammed Elbadry Stony Brook University
Fan Ye Stony Brook University
Peter Milder Stony Brook University
YuanYuan Yang Stony Brook University

Current NFD performance on wireless networks is severely obstructed by the lack of a true data-centric MAC layer. Building on top of V-MAC, our novel data-centric MAC prototype that offers high rate (up to 65Mbps on pi and 900 mbps on another platform), low loss (1-3%) multicast support (dozens of receivers), we present V-MAC NFD Link Protocol (V-NLP), a Link Protocol integrating V-MAC and NFD for a fully data-centric network stack. The new stack eliminates the network grouping concept, and enables pub/sub abstraction at the radio level. V-NLP allows researchers and developers to run NFD over V-MAC with multiple frame types (Interest, Data, Announcement, Implicit Interest), customized data rates, and configurable multicast support. This fully data-centric stack will enable NFD to achieve high performance

for applications in vehicles, drones and IoT environments. The source code for V-NLP, V-MAC and firmware for supported chipsets will be released to the community.

Session 7: Discovery/Configuration

Plug-n-Play NDN

Eric Newberry	UCLA
Tianyuan Yu	UCLA
Zhiyi Zhang	UCLA
John Dellaverson	UCLA
Lixia Zhang	UCLA

A major barrier to the adoption of NDN is the high manual effort required to set up and configure a functioning NDN network consisting of multiple hosts. To date, such configuration has required manual configuration on every host to establish initial content reachability and manual reconfiguration of links between forwarders whenever there is a topology change. These configuration tasks are particularly difficult when one is working with embedded or IoT devices. Meanwhile, with IP networks, one can connect multiple end devices together in a local network and have them "just work". Therefore, to help encourage the adoption of NDN beyond the research community, we have developed "Plug-n-Play NDN", where an NDN environment can easily established in a LAN environment. We demonstrate its effectiveness using a small example scenario.

NDNSD: Service Discovery in NDN

Saurab Dulal	The University of Memphis
Lan Wang	The University of Memphis

Service discovery (SD) is a fundamental requirement of contemporary networking systems. modern web, IoT application, building management, smart device, LAN, WAN, mobile application, etc depends on SD in a way or another. The proliferation of the applications towards the edge, rapid expansion and advancement of IoT and sensor networks, cloud computing, etc have pushed the service discovery to the next level of importance. It is also changing the dynamics of how the services and resources were discovered and used in the past. Furthermore, the pervasive nature of wireless networks and mobile devices has underscored its importance and is expected to be even more in the future. This presentation introduces NDNSD, a general-purpose, fully distributed, and scalable service publishing and discovery mechanism for NDN using the NDN synchronization protocol.

Data synchronization protocol plays a crucial role in NDN. The original NDN architecture envisioned combining sync protocol with application accessible libraries to provide transport functionalities to the applications. The application accessible libraries should hide the core network functionalities and primitives such as interest and data from the applications, and the sync should help the transfer of data from one application to another. The Internet protocol stack is a well-known example of such a model – best known as the hourglass model. Remaining in the realm of the internet hourglass model, releasing the importance of SD in NDN, and taking inspiration from some of the previous works, we design and developed NDNSD. It is an application accessible reusable library that uses synchronization protocol for service announcements and discovery. We view service discovery problems as data synchronization problems. Some applications actively looking to discover services while others trying to advertise the services. There are several benefits of using sync for SD i) no external dependencies or demanding change in network layer – sync comes with NDN natively ii) inherently supports multi-way communications, iii) Flexibility to implement application semantic and so on. Furthermore, two or more parties can agree on a common sync group. This applies to local as well as global applications or devices. A mobile application can agree on sync group "/letschat", whereas printer services can agree on "/printers". Similarly, IoT and edge applications can have their sync

group. Unlike the limitation in TCP/IP, these names (sync group) can be directly used in the network layer which gives huge flexibility to applications implementing their semantics and reflective names that can identify their services. Through this presentation, we will also share some of the results and experience gained through our real-world experiments and testing of the NDNSD at The University of Memphis.

NDN Repo for Genomics Datasets - Progress and Future Directions

Zhaoning Kong	UCLA/Purdue
Zixuan Zhong	UCLA
Alex Feltus	Clemson University
Susmit Shannigrahi	Tennessee Tech University
Lixia Zhang	UCLA

The genomics community publishes and retrieves a large amount of datasets. Currently, such large datasets are collected to, and distributed from, centralized repositories such as the National Center for Biotechnology Information (NCBI) which hosts more than 42 petabytes of DNA sequence data; similar repositories exist in other countries, such as Japan and UK, as well. Due to the advancement in sequencing technology and simulation capacities, institutional laboratories are also generating and hosting a large amount of geographically distributed datasets. Making NDN the substrate for genomic workflows is expected to make big data movement and sharing simpler and more efficient. Given that genomics dataset names already follow hierarchical and semantically meaningful naming schemes that are agreed upon by the community, these names can be easily converted to NDN names, which can then be used to publish, discover, and retrieve these datasets.

The recently developed NDN python-repo by UCLA provides a promising solution for serving genomic datasets. In an effort to evaluate NDN for serving genomics data, we published actual datasets from NCBI's Sequence Read Archive (SRA) database into the NDN science testbed (originally created by Colorado State University and later extended to Caltech and Northeastern). The purpose of this experiment is to investigate location transparent data retrieval and insertion into genomics workflows. We created containerized genomics workflows on the Pacific research platform (PRP). We configured these containers to utilize NDN tools to retrieve SRA sequences from the NDN science testbed; in this process, we have successfully published and retrieved up to 1.4TB of science data using python-repo.

The current python-repo stores a single copy of datasets only, which is subject to a single point of failure. To address this issue, we have sketched out an initial design to enhance python-repo with automated replication function among a given set of repo instances. This design utilizes anycast to attract user data insertion requests to the repo instance R that is closest to the requester, then R uses one-hop DHT to determine the other repo instances to replicate the received data file. Following the same anycast routing, data fetching requests are attracted to a repo instance R2 close to the requester, and R2 again uses one-hop DHT to find out the replica of the requested data file. We will also report our initial design of inter-repo communication protocol which handles the dynamics of the repo set, such as repo failures and recoveries, as well as new repo insertions and removals.

Posters and Demos

NDNViber: Vibration-Assisted Automated Bootstrapping of IoT Devices

Sanjeev Kaushik Ramani	Florida International University
Proyash Podder	Florida International University
Alex Afanasyev	FIU

The rapid proliferation of sensors and their use in the modern Internet of Things (IoT) environment has

led to a highly connected environment. For these inexpensive and connected devices to function efficiently, they need to communicate with each other as per the application they support. Communication with the correct entity and joining the correct network to share information are necessary operations. The action of pairing such devices securely so that they can trust the information exchange between them is termed as onboarding / trust bootstrapping. Bootstrapping is usually a highly cumbersome process, especially in resource-constrained and interface- less devices, which may not be accessible even physically after installation. In this paper, we propose NDNViber which compliments the existing bootstrapping techniques used in NDN based IoT networks. NDNViber provides NDN based networks with a dynamic, usable and secure out-of-band communication scheme using modulated vibrations to bootstrap multiple devices simultaneously and works in devices without any user interfaces. We implement a prototype that involves a commodity smartphone as the controller that can bootstrap many small IoT devices that possess accelerometer sensors. With NDNViber, we also analyze the bootstrapping of IoT devices that are inaccessible due to their physical orientation and deployment locations.

NDNts: Named Data Networking libraries for the Modern Web

Junxiao Shi yoursunny.com

NDNts enables Named Data Networking in modern web applications. It works both at the frontend in modern browsers and at the backend in Node.js server runtime, and can be used in both TypeScript and JavaScript projects.

This short presentation / demo showcases some of the best features in NDNts and how they make developer's life easier. This includes:

- automatic Interest retransmissions
- automated certificate issuance via NDNCERT
- declarative trust schema

The demo will use a combination of code samples and screenshots. Participants will be able to interact with some web applications on the Internet.

Towards a Distance Vector Routing Protocol for Named Data Networking

Italo Valcy Da Silva Brito Federal University of Bahia (UFBA)
Leobino Sampaio Federal University of Bahia (UFBA)
Lixia Zhang UCLA

Ad hoc mobile scenarios desire a lightweight routing protocol that can help propagate rapidly changing data reachability information in a highly dynamic environment. The currently deployed Named-Data Networking (NDN) routing protocol, NLSR, is based on link-state algorithms, which require synchronization of the link-state database which can be difficult to achieve in the above-intended scenario. We are developing a distance-vector routing protocol that enables each node to selectively propagate a data reachability vector containing the named-data prefixes current reachable to their neighbors. Such reachability information can be propagated transitively, allowing all reachable nodes at the time to estimate their reachability to desired data in a distributed and asynchronous manner. This poster presents our work in progress and prototype of the NDN Distance Vector Routing (NDVR) protocol. The initial design of NDVR consists of the simplest possible way to propagate name prefix reachability information based on distance vector algorithms. The protocol lets dynamically identified neighbor nodes use NDN's Interest and Data packets to propagate routing updates and runs in two main phases: (i) dynamic neighbor discovery (hello) and (ii) distance-vector information exchange (dvinfo). NDVR prototype is being developed on the ndnSIM simulation environment. We will provide a preliminary evaluation of NDVR's efficiency and effectiveness as compared to the cases of not using a routing protocol such as DDSN (Distributed Dataset Synchronization over disruptive Networks).

CertCoalesce: Efficient Certificate Pool for NDN-Based Systems

Sanjeev Kaushik Ramani Florida International University
Alex Afanasyev FIU

Named Data Networking (NDN) relies on public key sign- ing to ensure integrity and authenticity for all data packets fetched in the network. One of the considerations for reliability of such signing is limiting the scope (what the key can sign) and time (how long the key can sign) of the public keys and their certificates, usually referred to as "least privilege principle." Traditionally, the public key certificates are issued for relative long periods of times measured in months or years; which requires considerations for certificate revocation, e.g, when the private key is lost or compromised. However, if the validity periods can be reduced to days or hours, the complex (and sometimes semi-broken) revocation mechanisms can be completely eliminated. This poster proposes such a mechanism—CertCoalesce certificates—to efficiently manage virtually unlimited pools of short-term certificates with limited networking, storage, and computational overheads. Specifically, a single certificate request with a "primary" key can be used to bootstrap the process of creating an unlimited number of short-term certificates for derivative private/public keys. Moreover, such certificates can be issued asynchronously—periodically pre-provisioned or upon request with an Interest—terminating issuance of future certificates when necessary. Moreover, CertCoalesce design owing to the underlying elliptic curve cryptography ensures that a compromised key from the pool of keys will not reveal information about other keys/certificates in the pool.

Managing NDN with the Multiverse Network Management System

Amar Abane National Institute of Standards and Technology
Abdella Battou National Institute of Standards and Technology
Omar El Mimouni National Institute of Standards and Technology
Asmaa Hailane National Institute of Standards and Technology
Mheni Merzouki National Institute of Standards and Technology
Davide Pesavento National Institute of Standards and Technology
Junxiao Shi National Institute of Standards and Technology
Lotfi Benmohamed National Institute of Standards and Technology

As an emerging networking architecture for the future Internet, Named Data Networking (NDN) needs innovative management tools to facilitate the configuration and monitoring of NDN deployments, such as for testbed experimentation or research evaluations. The Multiverse Network Management System (MNMS) project at NIST aims at providing a feature-rich solution for configuring, monitoring and managing NDN networks. Although MNMS is designed with commercial (IP) network management features in mind (e.g., FCAPS), a great attention is dedicated to natively supporting the information-centric concepts of NDN. MNMS defines two main entities. First, the management agent running on each managed node (e.g., a forwarder or a server) provides an interface to interact with it either locally or remotely. Second, the controller, which represents the central controlling entity of MNMS, provides a set of services such as topology, routing, configuration, notification, and telemetry. Considering potential NDN deployments with current forwarding implementations (e.g., NDN-DPDK forwarder), MNMS entities are designed to communicate either using the NDN protocol, or over an IP-based protocol such as HTTPS. The agents and the controller interact in order to provide two sets of services accessible via a Web-based user interface: (i) telemetry which allows the operator to trigger customized measurements and collect results, and (ii) network management through which the operator configures and manages the network, assisted by the automatic processing features of the controller.

A Novel P4 Target Architecture for Runtime-Reconfigurable NDN Data Planes

Ouassim Karrakchou University of Ottawa
Nancy Samaan University of Ottawa
Ahmed Karmouch University of Ottawa

One main limitation that is currently receiving much attention in both the IP and Named Data Networking (NDN) research communities is the limited ability to dynamically reconfigure the network data plane in a top-bottom approach. Rather than allowing network operators to define and configure their switch functionalities, network switches are still designed with pre-dictated functionalities and offer little flexibility with respect to their reconfigurations. In these, switches, data plane algorithms/protocols are mostly hardcoded on chips and can only be changed by the switch manufacturers. This, in turn, limits their capacity to implement new protocols and hinders network operators' ability to adapt to evolving applications requirements.

The programming protocol-independent packet processing language (P4) represents a promising solution to the above limitation. Using P4, the control plane can program the forwarding behavior of the data plane independently from the switch hardware. P4 is a high-level programming language that can be used to describe how packets should be parsed, processed, and forwarded in the data plane. Network programmers write P4 programs to describe different behaviors (e.g., L3/L2 switches and load balancers) which are then compiled and translated into switch-specific configurations and executables. Nonetheless, a main challenge that is still facing current P4 architectures is switch runtime re-programmability. A P4 switch runs a single P4 program that contains all its forwarding functionalities. Hence, a modification related to the forwarding behavior of a single flow necessitates the recompilation of the program and reconfiguration of the switch which results in a switch down-time. Another additional limitation is that P4 is designed for fixed-size byte-based IP protocol headers, which makes P4 parsers, deparsers and match-action pipelines incompatible with the variable-size string-based headers in NDN packets.

In this poster, we address these limitations by presenting a novel P4-based NDN data plane that takes advantage of the P4 programmability while maintaining the traditional formats of the NDN packets and forwarding plane tables of NDN switches. NDN packets are first matched against extended NDN PIT and FIB tables, similar to traditional NDN switches, to determine their forwarding ports and to possibly associate them with P4 functions preassigned to their namespaces. The packet headers are then parsed and forwarded to one of the P4 functions. These P4 functions can perform measurements, execute advanced forwarding strategies (e.g., decide next-hops based on measurements or forward packets to controllers) or selectively modify packet headers. To facilitate runtime configurability, each P4 function runs in isolation from the other programs and maintains its separate state. By controlling the granularity of the P4 program sharing among the namespaces, and in turn, the forwarding paths, we can achieve a balance between function scalability and runtime reconfigurability. To address the limitation of name processing in P4, content-name processing is achieved at line-speed by designing a specialized NDN packet parser and deparser that are shared by all the P4 programs match-action pipelines. Finally, to allow for advanced NDN packet processing, we show how several behaviors can be implemented by extending the P4 language using extern functions. Our design can be implemented on standard FPGA-based P4 hardware like the NetFPGA-SUME card.

On Using NDN to Vertically Secure Smart Power Distribution

Sanjeev Kaushik Ramani Florida International University
Alex Afanasyev FIU

Smart Grids is a modern adoption of power systems with the potential to revolutionize power distribution and management. In a general setting, a smart grid system is composed of multiple stakeholders involved in power generation, distribution and consumption. Effective, secure and trustworthy communication among these entities is crucial for the optimal operation of the system. Named Data Networking (NDN), the prominent Information Centric Networking (ICN) model provides a conducive architectural design to support

the needs of this system. In this poster, we thus advocate the use of NDN which provides the consumers with better service (QoS), experience (QoE) and vertical security when used in Smart power distribution. We also discuss the data-centric security, data immutability and opportunistic in-network caching that provide necessary accessories for enhanced grid operation.