Presentation Abstracts

Lessons Learned Over the Last 10 Years

Lixia Zhang (UCLA)

NDNComm 2020 organized a panel to summarize the 10-year journey of the Named Data Networking (NDN) project, which was funded under the NSF Future Internet Architecture (FIA) program in 2010. This talk aims to provide a succinct summary of that 2020 panel on NDN's research results achieved, lessons learned, and articulation of next 10 years. The primary lesson we learned is that semantically named, secured data offers the right building block in constructing future networked systems, and making NDN usable to the broader community becomes an urgent next step. NDNComm 2021 will report the community's achievement over the last year and identify high-priority tasks for coming years.

NDN Codebase Overview

Beichuan Zhang (University of Arizona)

This is an overview of recent development of NDN codebase, including updated packet format, new features of NFD and other new forwarders, NDNCert, sync protocols, libraries, and others.

Developing practical E2E systems and applications using NDN

Marcin Spoczynski (Intel Labs)
Srikathyayani Srikanteswara (Intel Labs)

Tight integration of Data, Networking, and Computational Power is one of the many benefits offered by the architecture of NDN for Edge and Cloud Deployments as opposed to IP-based Networked systems. Further, multicasting, caching, or built-in security can provide many benefits to the operators and developers. We will explore how to build and deploy systems based on NDN along with the onboarding of practical applications. We will also uncover some of the gaps in the NDN stack for developing industrial IoT or cloud gaming applications and provide a recipe for future best practices for deployments of NDN networks and applications.

NDN for Data Intensive Science Experiments (N-DISE): Overview and Recent Developments

Edmund Yeh (Northeastern University)
Harvey Newman (California Institute of Technology)
Lixia Zhang (UCLA)
Jason Cong (UCLA)
Susmit Shannigrahi (Tennessee Tech)

The NSF-funded NDN for Data Intensive Science Experiments (N-DISE) project aims to accelerate the pace of breakthroughs and innovations in data-intensive science fields such as the Large Hadron Collider (LHC) high energy physics program and the BioGenome and human genome projects. N-DISE is designing and developing high-throughput caching and forwarding methods, containerization techniques, congestion control mechanisms, integrated with Field Programmable Gate Arrays (FPGA) acceleration subsystems, to produce a system capable of delivering LHC and genomic data over a wide area network at throughputs up to 100 Gbits per second, while dramatically decreasing download times. This talk will provide an overview of N-DISE and an update on recent progress, including preparations for an upcoming demo at SC21.

NDNc - A lightweight integration of ndn-cxx with NDN-DPDK to achieve high throughput performance in scientific applications

Catalin Iordache (California Institute of Technology)
Harvey Newman (California Institute of Technology)

This talk will give an overview of a newly developed NDN library, called NDNc, that adds ndn-cxx support to the NDN-DPDK forwarder. We want NDNc to be a small and efficient NDN library for C++ applications that aim for high throughput performance, something that NDN community is lacking at this moment. We will first present what XRootD framework is and how it's being used for transferring data at the Large Hadron Collider (LHC) high energy physics program, one of the world's largest big data applications, and how developers can extend this framework to introduce their own security, caching or even filesystems, all transparent to the end-users. We will then go through the plans of N-DISE to deploy Named Data Networking at LHC by implementing an NDN XRootD Open Storage System (OSS) plugin, and the results achieve during the SANDIE project through various implementations using ndn-cxx and NFD, NDNgo, or pure DPDK consumer and producer applications all in search for boosting throughput performance, which in the end serve as the precursor of NDNc. Finally, the current state of NDNc will be presented, as well as two consumer-producer applications build on it and future plans.

Exploring rate-based congestion control in NDN

Sichen Song (University of California, Los Angeles)
Lixia Zhang (University of California, Los Angeles)

IP's point-to-point network delivery has resulted in a common approach by all the existing congestion control solutions: setting the congestion control window size based on the estimated pipe size of the end-to-end path between a sender and a receiver. NDN brings new challenges to congestion control because its new features of dynamic multipath forwarding and in-network caching destroy the notion of an end-to-end path and related pipe size.
In this talk, we report our work-in-progress for developing a rate-based congestion control solution to be deployed at end consumers. The basic idea of rate-based congestion control is to compare a consumer's interest sending rate with its data arrival rate in order to detect congestion and adjust the interest sending rate accordingly. We will share with the audience several observations and lessons we have learned by evaluating and revising our initial design through simulation experimentation. These observations and lessons include

1. There is a distinction between the delay in observing the effect of one's rate adjustment (which takes a round trip time) and the time period to measure the effect (the measurement period)
2. There is a tradeoff in selecting the control algorithm between its graceful sending rate increase (to minimize queue growth) and the desire to detect congestion quickly

3. When multiple consumers fetch the same data around the same time, we observed an interesting interplay between the interests from different consumers, some fetching data from caches, some from the producer directly, and some in between. These interplays create delay variations that impact congestion detection and we investigated means to overcome this issue.

NDN-DPDK File Server for Data-Intensive Science Applications

Junxiao Shi (NIST)
Davide Pesavento (NIST)
Lotfi Benmohamed (NIST)

NDN-DPDK is more than a high-speed NDN forwarder.
In support of N-DISE project, we recently developed an NDN file server based on io_uring API.
This presentation will introduce the architecture of this file server and present preliminary benchmarks.
I'll also include other updates in NDN-DPDK over the past year.

FPGA-Based Acceleration of the NDN Forwarder

Michael Lo (UCLA)
Jason Cong (UCLA)

Named Data Networking (NDN) presents a new approach of addressing content on the Internet. The current Internet architecture addresses content using a fixed length 32-bit (IPv4) or 128-bit (IPv6) IP addresses. NDN addresses can be variable in length and are not limited to numbered addresses, but more akin to a file directory. However, its variability in length presents a computational challenge for address decoding and packet forwarding. A commonly used solution is to hash the name into a fixed 64-bit value. This makes dispatching and processing downstream easier but makes this step a potential bottleneck because the string still has to be traversed entirely.

This motivates us to design an FPGA accelerator for NDN packet forwarding. Our current (initial) FPGA design uses the Alveo U250 card and focuses on accelerating the name hashing computation and the dispatching of the packet to the correct forwarding thread. The design has 8 processing elements to compute the hashes and 1 lookup table for thread dispatching. It is able to achieve up to 4x speedup compared to a Xeon Gold 6244 CPU also using 8 threads for hashing and 1 lookup table in a batched offline setting. We plan in the future to create a design that can achieve further speedup over a CPU by reducing the per packet communication overhead between CPU and FPGA to match a realistic online scenario where packets arrive randomly since it may be that batching cannot be relied on entirely.

DARPA Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE) and Mission-Integrated Network Control (MINC)

Mary Schurgot (DARPA)

The Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE) program is developing software for sharing information across multiple security levels by creating overlay networks using named data networking (NDN) technology based on individually secured packets. SHARE is securing tactical mobile handheld devices to support distributed information sharing without the need for reachback to large-scale fixed infrastructure and is integrated with the Tactical Assault Kit (TAK) application running on Android.

The Mission-Integrated Network Control (MINC) program objective is to ensure that critical data finds a path to the right user at the right time in highly contested, highly dynamic communication environments using secure control of any available communication or networking resources (communications, compute, or storage capabilities). MINC seeks to leverage recent networking advances including software-defined networking to provide flexibility via software programmability and information-centric networking to securely discover and retrieve network data.
Distribution Statement "A" (Approved for Public Release, Distribution Unlimited)

Towards Unification of Name and Address Based Communication

Mohammed Elbadry (Stonybrook University)

NFD is rich with algorithms and designs that benefit edge networks. However, industry professionals hesitate to move to a name-based stack after decades of address-based communication reliance and technical debt. In this presentation, we propose a lightweight version of NFD with NDN service APIs samples that provide address-based wireless developers the exposure to NDN while using an address-based stack to see the benefit of NDN to move to NFD. Further, we'll also cover the benefit of a unified, address, and name-based paradigm stack by going over the Medium Access Control benefits for both paradigms.

NEAR Platform: Supporting Augmented Reality Over NDN

Jinghao Zhao (UCLA)
Yunqi Guo (UCLA)
Lixia Zhang (UCLA)
Songwu Lu (UCLA)

In this presentation we report our design and development of the NEAR (NDN wirEless Augmented Reality) platform, an open-source AR platform with an information-centric design. The platform integrates information-centric wireless design with acceleration and security support. In the NEAR design, mobile devices produce real time video streams, edge servers fetch these real-time camera views and publish AR contents, which are fetched by interested consumers. It supports heterogeneous AR tasks, including object detection, face detection, OpenPose, AR video overlay, 3D model rendering, multi-view video stitching, etc.

We designed the namespace for NEAR to enable the group transmission for AR contents. NEAR advantages the Wi-Fi transmission with link-layer multicast. NEAR leverages FPGA-based accelerators for YOLO and Pose Recognition to showcase the usage of acceleration-as-a-service in edge computing. NEAR also explores the NDN security design to provide encryption and integrity protection for AR named data. NEAR platform is open-source to the community and contains extensible modules for wireless, security, machine learning tasks, acceleration, and namespace design.

We also plan to demo the integration prototype of the NEAR platform. With the information-centric wireless integration based on link-layer multicast, the system scales to multiple users with marginal throughput overheads. We integrated heterogeneous AR tasks on the edge server with pythonNDN. Our prototype demonstrated the real-time multi-user AR experience with NDN. The platform also includes security support with encryption and integrity protection and acceleration support. We hope it could serve the community as the testbed to compare performance for new designs in the future.

Methods for NDN based data transfer in multi-path networking environments

Xin Tian (Intelligent Fusion Technology, Inc.)
Khanh Pham (Air Force Research Lab)
Genshe Chen (Intelligent Fusion Technology, Inc.)

In a multi-path networking environment, there are multiple network paths between a NDN data consumer and a NDN data producer. For the transfer of a large data file with the same using NDN, it is desirable to use throughputs that are available from the multiple network paths. There are a few issues that make the problem more complex. First, each end-to-end network path has its own network delay, and the network delays of the end-to-end paths may be very different. Second, available throughputs from different end-to-end network paths may or may not affect one another. Third available throughputs from the network paths are not known in prior and may not be easy to estimate during the data transfer process. One solution to the problem is to add a flow ID field in NDN interest packet at the data consumer side. NDN forwarders use the same forwarding method for NDN interest packets with the NDN file name and the same flow ID. As a result, NDN interest packets with the same flow ID will use the same path from the data consumer to the data producer, and NDN interest packets with different flow IDs may follow different network paths. The addition of flow ID allows the conversion of a multi-path networking environment into a multi-single-path networking environment. Each single network path between the data consumer and the data producer has a relatively stable delay and throughput. And independent congestion controls at the data consumer can be used to effectively receive data using the paths (corresponding to the flow IDs) through the network.

NDN for Next Generation of Factory Automation

Charif Mahmoudi (Siemens)

Factory Automation requires advanced performance guaranties to meet the timing requirements. This talk describes the complexity of the next generation of factory automation systems and the emergence of novel opportunities for NDN as industrial communication paradigm. Dr. Charif Mahmoudi will describe a holistic view on the industrial 5G research and discuss some of the challenges limiting the adoption of NDN.

Information Management in the Emerging Edge

Jeff White (Dell)

Edge computing is a changing and growing architectural deployment in Information and Computing Technology. Edge is a distributed computation system including scalar/vector compute, storage, and networking. A central driver of the emergence of Edge is the need to process data and acquire knowledge in proximity to the production of data. The ability to provide observability, classification, services and movement autonomically will be critical. NDN inherent capabilities can play a significant role in enabling the Edge Information management use case.

NDNts video streaming using QUIC and WebTransport

Junxiao Shi (NIST)

Inspired by iViSA, I built a browser-based video streaming application based on NDNts: Named Data Networking libraries for the modern web.
The app makes use of WebTransport, an emerging web specification that allows my application to connect to NDN routers over UDP-based QUIC protocol.

Measurements from real world viewers indicate that NDN-over-QUIC generally performs better than NDN-over-WebSockets.

NDN Play

Varun Patil (UCLA, Computer Science)
Tianyuan Yu (UCLA, Computer Science)
Lixia Zhang (UCLA, Computer Science)

NDN Play simulates a simple NDN network entirely in the browser using NDNts. It serves as a teaching tool for introducing NDN to newcomers, and as a visualizer for NDN networks for researchers. NDN Play features a topology editor, a TLV visualizer and running, bundling and replaying arbitrary experiments in the browser. It also supports running as a frontend for MiniNDN, providing easy real time visualization of NDN experiments. The demo will showcase the web application and its capabilities.

Multiverse: Designing a Network Management System with NDN

Amar Abane (NIST)
Davide Pesavento (NIST)
Mheni Merzouki (NIST)
Junxiao Shi (NIST)
Lotfi Benmohamed (NIST)
Abdella Battou (NIST)

Multiverse is a research platform intended for managing NDN, Quantum, and IP networks. Currently, each of these technologies has a different level of maturity, complexity, and management expectations. We present an overview of Multiverse network management use cases, requirements, and design principles. We then highlight the benefits exchanged between Multiverse and NDN. On the one hand, Multiverse provides centralized management for NDN-DPDK forwarders. On the other hand, Multiverse leverages NDN concepts to define its own management APIs with a pub/sub abstraction. This allows the specification of various network management schemes which can be implemented in either NDN-based or IP-based pub/sub frameworks. The pub/sub abstraction and tooling are illustrated through an example of Quantum network management.

NDN Forwarder Manager Demo

Xinyu Ma (UCLA)

Up to now, configuring NDN for use on end hosts has generally been difficult due to the absence of graphical configuration interfaces. To improve the usability of NDN, we have developed the NDN Forwarder Manager (NDN-FM). NDN-FM provides a graphical interface for users to manage a local instance of the NDN Forwarding Daemon (NFD), allowing them to monitor the status of the forwarder; create, update, and delete faces and routes; manage certificates; and run basic NDN debugging tools. NDN-FM can also be used to manage other NDN packet forwarders that support the NFD Management Protocol. In this demo, we want to show the functionalities of NDN-FM and possible use cases.

mGuard: A Secure Real-time Data Distribution System with Fine-Grained Access Control for mHealth Research

Lan Wang (University of Memphis)
Saurab Dulal (University of Memphis)
Chaudhry Nasir Ali (University of Memphis)
Siqi Liu (UCLA)
Adam Robert Thieme (University of Memphis)
Santosh Kumar (University of Memphis)
Lixia Zhang (UCLA)

In this talk, I will give an overview of the mGuard project funded by NSF. I will also talk about the issues we encountered and future plan.

The mGuard project aims to address two major data access challenges encountered by the NIH Center of Excellence for Mobile Sensor Data-to-Knowledge (MD2K, https://md2k.org/) in its pursuit to share mobile health (mHealth) data among researchers who investigate a wide range of health and wellness issues. First, since wearable sensor data may expose privacy-sensitive information about a user, it should only be accessed by authorized users; currently, this access control is largely handled manually, incurring high overhead while being subject to human error. Second, in order to enable real-time intervention for certain medical conditions, researchers need the ability to retrieve and process the sensor data in real-time, which is not currently supported. mGuard tackles these 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).

In the first year of this project, we have developed the overall mGuard system design, designed an NDN naming scheme for MD2K's mHealth data, developed an access control policy specification and implemented policy parser, extended the NAC-ABE library to support key policy attribute-based encryption (KP-ABE), and designed a pub-sub API using PSync and NAC-ABE for publishing and subscribing to mHealth data. In the second year, we will finish the system implementation and testing, and deploy mGuard on the NDN testbed.

Hydra - A secure, distributed, and federated storage for large science data

Susmit Shannigrahi (Tennessee Tech University)

This work describes Hydra – a framework that supports big-science communities with secure and resilient access to big datasets. Different from distributed databases, Hydra is a loose federation of repositories potentially owned by multiple administrative entities between organizations, and will have security built into its file operations and data access natively. Hydra addresses the following two problems at once: providing data publishers an automatically replicated, location independent publication platform for new data, and providing data consumers a uniform name-based interface to access all data, including both newly published data and existing data stored in legacy storage systems. Hydra will (a) provide long term data storage, (b) automatically replicate data across multiple geographically distributed repositories and maintain a desired degree of data replication in face of individual repository failures, (c) provide secure and scalable file access, and (d) provide a unified interface to access data stored in both Hydra and existing storage systems.

Bootstrapping Remote NDN Entity Leveraging CA-based Authentications

Tianyuan Yu (UCLA)
Philipp Moll (UCLA)
Zhiyi Zhang (UCLA)
Alex Afanasyev (Florida International University)
Lixia Zhang (UCLA)

Bootstrapping enables an NDN entity sending Interest and receiving Data packets securely.
The new challenge in NDN application deployment is to securely bootstrap a remote NDN entity, which requires mutual authentication and name assignment.
In this talk, we present a design for remote NDN entity bootstrapping that leverages existing authentication systems and minimizes manual operations.
Our design addresses the two needs with (1) new NDN entity authenticating the trust anchor controller via software distribution with today's web security support, and trust zone controller authenticating a new entity by verifying its CA rooted certificate; (2) new entity self-obtaining a name from its existing authentications such as CA rooted certificates, and applying identity bundle from trust zone controller that provides initial trust relations and security credentials.
We apply this design to a federated data repository project to bootstrap NDN application instances on remote file servers from different campuses.

Industrial Applications and NDN

Kathleen Nichols (Pollere, Inc.)

Operational Technology (OT) networks support applications ranging from interconnecting a single home's Internet of Things (IoT) devices to distributed control of massive regional power distribution systems. Until recently, OT networks were a small niche with application development and deployment heavily constrained by proprietary hardware and communication systems that required every application to manage all the low-level details of interconnection
and signaling.

Today, most OT systems are built from low-cost commodity parts using open, application-focused 'communication' standards like Zigbee, Zwave, BT-mesh, etc. Without exception, these standards solved early OT's app development barrier, not by adopting IT's endpoint communication model, but instead by using an internal layer 2 mesh to insulate applications from communication details and only exposing a simple pub/sub-like API mediated by application-oriented logical identifiers instead.

This application-centric communication approach has accelerated the OT market growth by minimizing system NRE and deployment cost while ensuring interoperability and flexibility. Unfortunately, solving the communication problems also vastly increased the attack surface, highlighting OT's almost total lack of a security architecture (early systems relied on a proprietary system's inherent “security through obscurity" and limited connectivity). Since the middle layer of modern OT stacks is IP6-based, a first thought was to apply endpoint security architectures evolved for the Internet and IT like (D)TLS (https://buildwithmatter.com). Unfortunately, this still resulted in a number of spectacular failures. (https://blog.checkpoint.com/2020/02/05/the-dark-side-of-smart-lighting-check-point-researchshows-how-business-and-home-networks-can-be-hacked-from-a- lightbulb,
https://www.washingtonpost.com/business/2021/05/10/colonial-pipeline-gas-oil-markets).

More recent efforts note there are fundamental differences in the communications in OT networks that can be exploited to create more secure networks. (RFC8520, NIST SP 1800-15B, “Trust Schemas and ICN: Key to Secure Home IoT” ACM ICN 2021) Unlike general Internet applications, all the elements of an OT network have a role (a particular job to do) which makes it possible to characterize "who can say what to which" (as well as who can’t). ICN, particularly NDN, provides primitives that can be employed to enforce this characterization while supporting and often simplifying the robust multicast communications known to work well for OT.

This talk aims to point out that OT has problems NDN-based approaches can solve while conventional networking either can’t or would have inferior performance. A further aim is to suggest the community consider this as the way that NDN can contribute to networking in the short term.

Considerations for Higher Level Transports over Sync

Varun Patil (UCLA, Computer Science)
Lixia Zhang (UCLA, Computer Science)

Application developers looking to build data-centric applications over NDN may not be familiar with the usage of NDN Sync for namespace synchronization. We build a publish-subscribe API over the State Vector Sync protocol as a higher-level transport, providing developers with a familiar pattern for multiparty communications. We discuss the design of the Pub/Sub API and its capabilities such as resiliency and performance. We also discuss further considerations for such a higher-level transport, such as the choice of the underlying Sync protocol, relationship between application data units and network packets, segmentation, signing patterns and data aggregation and its effects on protocol performance.

sV2Pc: On Scaling LTE-based Vehicle-to-Pedestrian Communication using NDN

Proyash Podder (Florida International University)
Sanjeev Kaushik Ramani (Florida International University)
Somak Datta Gupta (Ford Motor Company)
Azin Neishaboori (Ford Motor Company)
Alex Afanasyev (Florida International University)

Recent advances in vehicular communication technologies, including DSRC, LTE, enable cars to observe the surroundings, communicate (between each other, roadside units, and pedestrians), and analyze various potential hazards more effectively. However, for now, such communication is largely limited to broadcasting alerts. We believe it is mostly because of the mismatch between the current standards of the networking architecture (IP) and vehicular application semantics. Named Data Networking (NDN) is a proposed networking architecture that is designed to network the world of computing devices by names and secured data items instead of point-to-point packet delivery.

This paper aims to demonstrate the feasibility of using NDN to build a scalable mechanism for vehicle-to-pedestrian (V2P) communication in the promising base station independent communication using the recently developed CV2X standard.
Using a simplified vehicle intersection awareness application, representing a class of V2P communications, we identified advantages of data-centric communication of NDN, as well as challenges that need to be addressed to scale the communication, including the needs for distributed mechanisms for similar Interest and Data suppression and geographical bounds for Interest forwarding. To confirm

the findings and effectiveness of the proposed sV2Pc solution, we performed a simulation-based study, evaluating bandwidth use and comparing it with a currently proposed approach used for pedestrian safety.

BLEnD: Improving NDN Performance Over Wireless Links Using Interest Bundling

Md Ashiqur Rahman <marahman@cs.arizona.edu> (University of Arizona)
Teng Liang <liangt@pcl.ac.cn> (Peng Cheng Laboratory)
Beichuan Zhang <bzhang@cs.arizona.edu> (University of Arizona)

Named Data Networking (NDN) employs small-sized Interest packets to retrieve large-sized Data packets. Given the half-duplex nature of wireless links, Interest packets frequently contend for the channel with Data packets, leading to throughput degradation over wireless links. In this work, we present a novel idea called BLEnD, an Interest-bundling technique that encodes multiple Interests into one at the sender and decodes at the receiver. The major design challenges are to reduce the number of Interest transmissions without impacting the one-Interest one-Data principle embedded everywhere in NDN architecture and implementation, and support flow/congestion control mechanisms that usually use Interest packets as signals. BLEnD achieves these by bundling/unbundling Interests at the link adaptation layer, keeping all NDN components unaware and unaffected. Over a one-hop WiFi link, BLEnD improves application throughput by 30%. It may also be used over multiple hops and be improved in a number of ways.

Adaptive Duplicate Suppression for Multicasting in a Multi-Access NDN Network

Saurab Dulal (University of Memphis)
Lan Wang (University of Memphis)

Multicasting [1] is an efficient choice for multi-party (one-to-many, many-to-many) communication that is desired by modern applications such as video conferencing, gaming, vehicular networks, disaster management, IoT and sensor networks. It provides a great way to disseminate information simultaneously to users interested in the same data. Multicast can also eliminate the requirement of fixed infrastructure (e.g., access points) in a wireless ad-hoc network.

Name Data Networking (NDN) provides native support for data multicast. In NDN, every piece of content is named and signed (optionally secured) during the creation of the packet. This enables decoupling of the data packets from the producer. Any intermediate node can cache and serve the data because the receiver can always verify the data provenance through the signature. However, the current NDN forwarder lacks a duplicate suppression mechanism for multicast in a multi-access network, which can cause network congestion and significantly degrade the overall packet delivery performance. In our earlier work [1], we experienced enormous traffic overhead in many of our experiments due to the absence of multicast duplicate suppression in NDN. When all the nodes exchanged multicast messages with each other at the same time, we were unable to scale the network beyond 13 nodes because an enormous number of duplicate packets caused extreme network congestion.

Through this presentation, we would like to share our work “Adaptive Duplicate Suppression for Multicasting in a Multi-Access NDN Network”. It includes our design, implementation, and some preliminary evaluation results. Our early evaluation has shown a substantial reduction in duplicate traffic in NDN multicast communication.

[1] Dulal, Saurab. NDNSD: Service Publishing and Discovery in NDN. Master’s Thesis. University of Memphis, 2020.