

The Second Named Data Networking Community Meeting (NDNcomm 2015)

Alexander Afanasyev, Yingdi Yu, Lixia Zhang, Jeff Burke
UCLA
aa,yingdi,lixia@cs.ucla.edu
jburke@remap.ucla.edu

kc claffy, Joshua Polterock
CAIDA/UCSD
kc,josh@caida.org

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This report is a brief summary of the second NDN Community Meeting held at UCLA in Los Angeles, California on September 28–29, 2015. The meeting provided a platform for the attendees from 63 institutions across 13 countries to exchange their recent NDN research and development results, to debate existing and proposed functionality in NDN forwarding, routing, and security, and to provide feedback to the NDN architecture design evolution.

Categories and Subject Descriptors

C.2.5 [Local and Wide-Area Networks]: Internet; C.2.1 [Network Architecture and Design]: Packet-switching networks

Keywords

Information-Centric Networking, Named Data Networking, Architecture

1. INTRODUCTION

On September 28–29, 2015, the Named Data Networking (NDN) project team held the second NDN Community Meeting [13] at the University of California, Los Angeles. The meeting was attended by 99 participants from 63 institutions across 13 countries, and streamed live using both TCP/IP- and NDN-based video streaming. Presentations and discussions covered the current state and future directions of the NDN architecture, software platform and supporting libraries, software licensing strategies, applications, and testbed. Presentations, demos, and posters showcased recent research results inside and outside of the NDN project team. The NDNcomm meeting series is part of the project team's effort to build a community around NDN research. This report summarizes discussions of the meeting.

The NDN project strives to use current and future applications to drive the development and deployment of the architecture and its supporting modules, to test prototype implementations, and to encourage an iterative cycle of design, real-world experimentation, and evaluation. In this spirit, the NDNcomm program committee encouraged contributions that considered specific case studies, application requirements, and real-world scenarios.

2. FUTURE OF STORYTELLING

Jeff Burke (UCLA REMAP) began the meeting, which was hosted at the UCLA School of Theater, Film and Television, by

describing the relationship between available technology and the fundamental human activity of storytelling, which now makes up the majority of Internet traffic. He noted that experience of stories happens through local interactions, therefore today's solutions that solely rely on cloud services are unlikely to meet the challenges of live, real-time experience that integrate local elements and context. Jeff gave several interesting examples of how NDN can expand the power of storytellers across a variety of disciplines, by lowering the barrier for experimentation, from IoT to expeditionary networks to big science and public communication of huge amounts of data. More broadly, he noted the relationship in design between problem formulation and the available solutions and suggested that attendees consider how NDN encourages not just incremental advances but reformulation of important problems, using this quote:

It is widely accepted that creative design is not a matter of first fixing the problem and then searching for a satisfactory solution concept; instead it seems more to be a matter of developing and refining together both the formulation of the problem and ideas for its solution... [15]

Van Jacobson (Google/UCLA) explained how computer scientists and network researchers “tell stories” with APIs and packet formats, and how the existing IP architecture limits the expressivity of network uses today. Application-level protocols, e.g., the web, alleviate the problem, but the client/server nature of the underlying networking architecture still imposes limitations in the face of today's communication patterns. Van described how NDN offers an entirely new higher-level networking vocabulary, idioms, and improved tools for telling stories. He illustrated this shift using three dimensions of network architecture: *transport*, *routing*, and *security*. Instead of thinking of transport as a conversation between two endpoints, NDN models transport as reconciliation of data sets [26]. Instead of conceptualizing routing as building and maintaining a set of paths through a graph, NDN allows reconceptualization of routing as maintaining information about local neighbors. With this information, interest packets can be forwarded toward points in an underlying metric space [17] via neighbors that are “closer” to locations of data. The NDN team is experimenting with this idea, which offers a new approach to routing in intermittently connected and highly dynamic (e.g., IoT) environments, but may also help address global Internet routing scalability problems. Most importantly, the NDN architecture recognizes that the various approaches to retrofitting layers of security onto fundamentally insecure IP networks are massively failing, and offers building blocks of trust (signed packets) and design patterns (schemas) [25].

3. NDN PROJECT STATUS

3.1 Architecture Achievements

Beichuan Zhang (U. Arizona) gave an update on the core architectural achievements in NDN forwarding, routing, and evaluation environments. The NDN protocol adopted a new flexible TLV-based packet format [21]. Over the past year the team completed the transition from the legacy code base (CCNx) to the NDN Forwarding Daemon (NFD) [10] and a set of evolving NDN libraries supporting multiple programming languages. The past year saw one major release and five minor releases of NFD, as well as multiple releases of reference libraries and prototype applications. All NDN software is open source and available at <http://www.named-data.net/>, with over 25 contributors to NFD and over 70 contributors across all projects from both inside and outside of the NDN team. The NDN platform now also runs on Android natively, as well as on a set of embedded platforms, including Raspberry Pi, DD-WRT and OpenWrt.

NDN researchers have several alternatives for evaluating NDN research and development: from simulation (ndnSIM [18]), to light-weight (Mini-NDN [3]) and full-featured (Open Network Lab [23]) emulation, to the global NDN testbed. The recently released version 2.1 of ndnSIM, the NS3-based NDN simulation package, uses the NFD code base to implement NDN forwarding functionality, which allows large-scale evaluation while providing an easy transition from simulation to real system development. The initial release of Mini-NDN, a light-weight emulation tool based on MiniNet, supports virtualized nodes that run NFD and NLSR, allowing real NDN applications to be evaluated in a smaller-scale emulated environment with high fidelity results.

Beichuan described new core architectural components, including the following:

- **LINK Object** [14], a new type of data packet that securely “redirects” from one name to a set of other names, providing a way to address routing scalability as well as mobile publishing. In mobile publishing, when a producer moves, its data name does not change, only its reachability prefix changes (e.g., globally routable prefix towards the producer). The LINK object helps routers steer interests toward the closest instance of the desired data.
- **Producer NACK** [19], a new type of data packet to notify consumers that the requested data does not exist. The producer NACK optionally includes information regarding when the data may be produced in the future.
- **Network NACK** [24], a network-level signaling mechanism to provide feedback to a downstream node that an interest cannot be further forwarded. Using this feedback, the downstream node may explore other options to forward the interest or, if all options are exhausted, notify further downstream nodes of the problem. Possible reasons to issue network NACK include detection of duplicated interests (e.g., an interest loop is detected), link failures, congestion, etc.
- **NDN Link Protocol (NDNLPv2)** [9], a link adaptation protocol between direct NDN neighbors, supporting fragmentation and reassembly, loss and link failure detection and recovery, and network NACKs.

Beichuan also summarized progress in exploration of forwarding strategies. Existing strategies have gone through multiple iterations, e.g., the best-route strategy is now at version 4, allowing application-initiated retransmissions of interests and featuring an exponential backoff of the suppression interval to prevent abuse. A

new strategy was specifically designed for access routers, and another new strategy was developed for adaptive shortest RTT-based forwarding used in conjunction with hyperbolic routing. Beichuan described the status of routing experimentation with the traditional link state routing protocol (NLSR), hyperbolic routing, and the SNAMP approach to scale global routing [14].

In the area of security, Beichuan reviewed the added support for ECDSA signatures in NDN protocols [21], the new Public key Information Base (PIB) service [12] to simplify key management and enable automatic publishing of keys, upcoming changes in NDN certificate format [4], improved signing APIs in NDN libraries for better usability, and automated testbed certificate issuance [5].

To support edge device configuration and minimize deployment complexity, NFD includes an improved implementation of automatic NDN hub discovery, routable prefix discovery, as well as a new built-in automatic prefix propagation feature of NFD to create and maintain routes on downstream (from the data producer) NDN routers on behalf of applications, i.e., when an application registers a prefix locally, the local NFD automatically propagates this registration to the downstream router.

Future plans include: design, development, and evaluation of forwarding strategies to support IoT, ad hoc mobile, delay-tolerant, and otherwise constrained network environments; experimentation with hop-by-hop signaling and congestion control; packet format refinements; and performance optimizations.

3.2 NDN Application Advances

Alex Afanasyev (UCLA) introduced the *trust schema* [25], the team’s new approach to facilitating automated use of NDN security functions in applications. NDN enforces data-centric authenticity by binding a name and its content with a cryptographic signature, but leaves open the specific trust model to apply to a given packet, i.e., how keys are authorized and distributed. By naming data and defining naming conventions, NDN enables clear expressions of trust models as a set of relationships between data names and signing key names. The trust schema allows expression of virtually any trust model as a concise set of rules, the trust model (schema) can be packaged and further secured using higher-level schemas, which allows building complex yet secure chains of trust relationships. Alex illustrated how a trust schema can capture data/key name relationships by constructing a schema to power a web blog ecosystem made of blog site, administrators, authors, and articles. The trust schema approach now powers several NDN software packages: NDN forwarding (NFD), NLSR, Repo-ng, and ChronoChat.

Yingdi Yi (UCLA) summarized lessons learned on data-centric confidentiality over the last year, based on experience with the NDN-NP application environments [20]. These environments revealed the requirements of enforcing data access control without relying on intermediate entities (such as cloud servers or third-party data storage), and the need for differential confidentiality. While directly encrypting data seems like a natural solution, several challenges arise for specific application scenarios. For example, data created by a distributed set of producers may need to be accessed by another set of consumers in real time, creating challenges for access management and encryption key distribution. Yingdi presented *Name-based Access Control* (NAC), an access control mechanism based on data encryption, as a current solution for the NDN-NP application environments. Yingdi illustrated how to use the expressiveness of hierarchical NDN namespaces to define the privilege of keys (i.e., data sets that a key can encrypt/decrypt) at a fine granularity, and how to distribute data access keys via normal NDN interest/data exchanges. He noted remaining issues including access revocation, forward secrecy, and read auditing.

Jeff Burke (UCLA REMAP) described how the team is using the NDN-NP network environments as drivers to explore the value of the NDN architecture for next-generation applications. The two NDN-NP environments that Jeff presented include *Open mHealth*, an ecosystem proposed by others for personal health- and wellness, whose data-centric model for interoperability is well-aligned with NDN, and *Enterprise Building Automation and Management Systems (EBAMS)*, a research driver that considers the use of NDN for both industrial control networks and low-capability devices as envisioned for the Internet of Things. For each environment, he described pilot applications being developed, including NDNFit, a personal fitness data management system using NDN, and Mini-EBAMS, an environment for developing NDN-based building data protocols that uses Mini-NDN, integrates real data from UCLA building management systems, and bridges to the NDN testbed. These environments and pilot applications enable focus on challenging areas of architectural research that can also significantly benefit strategic application scenarios, including work in data-centric security, synchronization-based transport, and mobility support.

Jeff discussed why schematized trust and name-based access control are a natural fit for these environments, and how to make these techniques more user-friendly. He noted four open challenges: lifelong data maintenance, data discovery, publisher mobility, and coordination of data access among multiple processing entities.

Peter Gusev (UCLA) announced the project's new videoconferencing application (NdnCon [8]) and its supporting library (NDN-RTC [6]), the culmination of two years of work that now supports low-latency, consumer-driven HD-quality videoconferencing using NDN, part of the multimedia applications cluster of the NDN-NP project. NdnCon development revealed new requirements for NFD, related to interest retransmission and forwarding strategies. Future plans include expanded support for OS platforms, adaptive rate control, and scalable video encoding.

3.3 NDN Testbed

Patrick Crowley (Washington U. in St. Louis) gave an update on the status of the NDN testbed, which currently consists of 26 NDN gateways around the globe: U.S., China, South Korea, Norway, Spain, France, Italy, and Switzerland, with new nodes joining almost every month. Patrick invited participation in the testbed¹ and emphasized that the NDN team will help set up, configure, maintain, and debug operational issues, as well as assist researchers with running experiments.

Patrick highlighted an important aspect of the NDN testbed: its routing protocol. Unlike today's global routing protocols that rely on an honor system for successful operations, the NDN testbed uses the NLSR [11], which now fully utilizes the built-in security of NDN. NLSR is the first continuously operating service to rely on an NDN trust schema to define and enforce named-based trust relationships between testbed entities. When a site joins the NDN testbed, the site operator gets a key pair, signed by the NDN testbed root authority. Testbed site operators create and certify keys for running routing software. All routing messages exchanged on NDN testbed are fully signed, with trust schema automating strict authentication of the trust chains.

4. SCIENCE AND NDN

The first morning of the meeting ended with a session of four talks by **Inder Monga** (ESnet), **Susmit Shannigrahi** (Colorado State U.), **Phil DeMar** (FermiLab), and **Huhnkuk Lim** (KISTI)

¹<http://named-data.net/ndn-testbed/policies-connecting-nodes-ndn-testbed/>

on scientific "big data" uses of NDN. All presenters recognized the need for new architectural and protocol support for managing, analyzing, and sharing extremely large volumes of scientific data. Each talk illustrated the same needs: ways to abstract storage and network capability from user-data interaction; ability for users to specify and retrieve needed portions of data; high end-to-end throughput; and a secure, scalable framework to integrate data management and network transport. ESnet staff worked with Christos Papadopolous (Colorado State U.) earlier this year to demonstrate the use of NDN to address some of the above needs. His group is now interested in expanding a federation of NDN devices and exploring open questions, e.g., caching strategies, content discovery, load balancing, and how to build knowledge of underlay into the NDN overlay.

5. INTELLECTUAL PROPERTY APPROACH

Eben Moglen (Columbia Law School/Software Freedom Law Center) gave an inspiring keynote on intellectual property considerations for future Internet architecture projects. Eben explained why GPL version 3 is the most appropriate license for core NDN software elements, in particular the reference implementation of NDN forwarder (NFD), while LGPLv3 and other more permissive licenses are appropriate for supporting tools and libraries.

Large-scale software ecosystems are increasingly occupied by stakeholders who want a core within which everyone can be assured of fair collaboration, and avoidance of commercial misappropriation of elements of that core, more than they want to promote product differentiation or competitive edge in that core. Commercial stakeholders also obviously find value in having a periphery in which they can differentiate themselves, and where they can avoid agreements about forms of code that may limit the ability to compete. This cooperate-in-core, compete-at-edges model characterizes the Linux ecosystem, and the Open Invention Network (OIN) patent pool has evolved as a vehicle to support it. Eben explained why, in the case of the NDN ecosystem, which has not yet accumulated a large enough pool of patents for an OIN model to gain traction, GPL version 3 is the most effective tool to accomplish the goals it shares with the OIN ecosystem.

6. LIGHTNING TALKS

The community meeting agenda included two sessions of lightning talks, many of which advertised posters and demos for the evening reception. **Minsheng Zhang** (U. of Memphis) described how to synchronize a partial dataset. **Vince Lehman** (U. of Memphis) presented the evaluation of hyperbolic routing as a viable solution to routing scalability challenges. Since an NDN forwarding table (FIB) could grow to an unmanageable size, hyperbolic routing offers an appealing solution because it can diminish the need for a FIB. Vince's group is exploring how to use the smart forwarding plane of NDN to mitigate some inherent drawbacks of hyperbolic routing, e.g., detecting suboptimal paths, getting around local minima, and enabling reacting to network dynamics.

Anders Plymoth (TelHoc) presented an interesting commercial use of NDN to create an application service platform supported with the NDN architecture. He provides a simple API to NDN-based secure data storage and data services, using geographically distributed servers connected via TLS. Each data is assigned a unique ID (which can use URN, JSON, or SQL naming), which can be used to retrieve the data securely.

Lei Liu (Fujitsu) overviewed issues and solutions related to producer mobility in NDN. He described a proposed solution to im-

prove bandwidth efficiency during the hand-over by using a centralized mobility coordinator.

Diego Perino (Alcatel-Lucent) motivated implementation of the high-speed software NDN router on commodity hardware, and described a prototype implementation (Augustus).

Maliheh Shirvanian (U. of Alabama at Birmingham/Verisign) focused on the implementation of access control mechanisms for NDN. Access control through content encryption can be challenging to implement with respect to key distribution and credential revocation. In some environments, active participation of routers can simplify access control, at the cost of some NDN features. For example, for interests marked using agreed upon naming conventions, routers can actively check permissions with the access control provider, or routers can require consumers to solve puzzle challenges in order to access data.

Andrew Brown (Intel) described an application of NDN to enable the integrated, secure, end-to-end telemetry and analytics solutions for IoT. NDN addresses multiple requirements of these solutions, including flexibility in the underlying transports, ability to process data locally, fine-grained application management of the control plane, and seamless handling of intermittent connectivity. Andrew also highlighted Intel's contributions to the NDN platform: bug fixes and minor feature implementation in the client libraries (jNDN and NDN-JS); leading development of several open-source NDN utility libraries for Java [2, 1]; and minimalist NDN forwarder implementation for embedding in Java applications.

G.Q. Wang (Huawei) introduced several possible network environments in which ICN could be useful. He argued the need for flexible data packet format to support network environments with different MTUs, especially extremely short MTUs, and proposed a packet format to accommodate this requirement.

7. BREAKOUTS

To enable focused conversations, we held parallel breakout sessions on both days of the community meeting.

- **Revisiting name-content binding: In-network name space operations.** Moderator: Christian Tschudin (U. Basel)
The group assessed the relevance of name-related in-network operations, where names are rewritten, or remapped, and can be bound to dynamic content.
- **NDN Development Environments and Support - Common Client Libraries.** Moderators: Alex Afanasyev (UCLA) and Jeff Thompson (UCLA REMAP)
The moderators gave updates on recent development progress, implementation motivations, and design tradeoffs. The breakout group also discussed features desired by the community, opportunities for participation in development, and feedback on the development environment(s).
- **Security Support: Schematized Trust; Name Based Encryption.** Moderator: Yingdi Yi (UCLA)
This group compared security functionality in IP and the current stage of NDN, and cataloged desired security properties for NDN. The discussion focused on privacy issues including: identifying privacy requirements and threat models in different application scenarios; how to enable forward secrecy in NDN; and how to mitigate privacy exposure caused by data names. There was a follow-up debate on whether it would be beneficial for NDN to offer privacy parity with IP (i.e., point-to-point encrypted channels) in addition to NDN's content-based user privacy and confidentiality support, and how to make the security and privacy design tradeoffs *explicit*.

- **Defense Applications of NDN.** Moderator: Tamer Refaai (The MITRE Corporation)
This group focused on determining whether and to what extent ongoing NDN research is applicable to contested (tactical) environments, and what adaptations of the current NDN implementation may be needed to support such environments.
- **Evaluation of NDN Architecture.** Moderators: Nicholas Proferes (U. Maryland) and kc claffy (UCSD/CAIDA)
This group discussed approaches for assessing progress toward the broad technical as well as social goals of NDN and ICN more generally.
- **NDN startups.** Moderator: Patrick Crowley (WUSTL)
This breakout session focused on establishing a shared understanding of why NDN startups are inevitable, and on brainstorming startup directions and opportunities.
- **NDN research funding opportunities and team-building.** Moderator: kc claffy (UC San Diego/CAIDA)
The session provided participants with pointers to funding agency and industry sources of funding and hints for building teams and proposals.
- **ICN-over-UDP.** Moderator: Christian Tschudin (U. of Basel)
The group discussed approaches on how to locate and create UDP tunnels to the first hop ICN router. The discussion favored the mDNS auto-discovery approach, as DNS and DHCP services are usually not available or cannot be reliably configured for a given environment.

8. INDUSTRY PANEL

Eve Schooler (Intel) moderated an insightful industry panel with the following experts: **Ralph Droms** (Cisco), **Luca Muscariello** (Orange Labs/IRT-SystemX), **Borje Ohlman** (Ericsson Research), **Ignacio Solis** (PARC), **Guo Qiang Wang** (Huawei Technologies), and **Greg Rutz** (CableLabs). She opened the panel by asking the following questions:

- How long have you/your company been focused on some flavor of ICN?
- Why is ICN interesting to your company?
- What one aspect of NDN you would like to change? And
- When is it appropriate to push for standardization?

Common themes were that NDN might be an auspicious approach to dealing with relentless traffic growth, especially over emerging mobile (5G) and IoT networks. Industry representatives also wanted to explore ICN's potential for providing innovative services, lowering the cost of operations, or managing and securing networks. There was a diversity of opinion on how soon ICN protocols would be ready for standardization. Some believed that interoperability was more important at this point than standardization, others thought that NDN was still too deep into fundamental research to even prioritize interoperability. Some that the community was missing a real content provider, who could talk about how new revenue streams could be made possible with NDN.

There was an interest in Anders Pymoth's (Telhoc) use of NDN to make a real business, in this case an application service platform selling data services. His key differentiator is that all data is stored and transported in encrypted form. His service does not mention NDN being part of the system, yet NDN enables his service to flexibly and easily manage and serve user content in a secure way, without having to deal with CDNs or cloud infrastructure. He has also virtualized the service such that when he needs to add more capac-

ity he can easily spin up a new VM. His service had some specific requirements, which NDN met perfectly.

9. PARTICIPANT FEEDBACK

During the second day of the meeting, we ran an interactive session to give everyone an opportunity to articulate the most interesting thing they learned from the previous day, and provide feedback and suggestions to the NDN team.

Several participants found it pleasantly surprising to see projects using the NDN software base (especially NDN-JS) they had never seen before, which motivated library developers to continue to invest time in software development. Although there was interest in more tangible use cases and demos that could ignite broad industry interest, there was strong appreciation for the scientific community's interest in use of NDN, and questions about why agencies such as the U.S. Department of Energy (funding three of the four projects presented in the "Science and NDN" session, Section 4) were not investing more heavily in NDN to address their pressing mission-critical scientific data problems.

Some industry participants expressed fascination with the scientific uses of NDN, but also wanted to see more work on how to handle namespace design to support security in "small data" (IoT/sensor networks) contexts. There was also acknowledgment of recent defense-related interest in NDN, e.g., for use in tactical environments.

One participant encouraged more active discussions about NDN network NACK support at ICNRG meetings, and hoped for more coordination on a consistent set of primitives.

There was recognition of many gaps remaining between application needs and the existing results, in particular related to usable security and supporting expected privacy needs of users. One person observed that the meeting agenda seemed to emphasize presentation of the accomplishments of massive software development efforts over the past year, rather than on more detailed reviews of specific research results. (Part of the motivation for that was that four NDN research papers [25, 19, 16, 22] were presented later in the week at the ICN conference, and the NDNcomm organizers avoided repeating those since there was substantial overlap in the NDNcomm and ICN audiences.)

Lixia Zhang, the leading PI of the NDN research team, gave final remarks of the session. She noted that the NDN community is rapidly growing, as indicated by the participation in the community meeting (e.g. 63 organizations vs. 39 at NDNcomm 2014), amplifying the need for better tutorial materials to explain to newcomers the basic NDN concepts and operations. Many people still consider the appeal of NDN to be relatively narrow, e.g., as an effective solution to content distribution problems, rather than a new network architecture. But NDN can enable secure and efficient communication and application development for a large range of applications, from very big to very small, and makes new applications possible and practical.

There is high demand for better documentation of the NDN team's technical design decisions, to share with, and collect feedback from, the larger networking community. The team started a NDN tech memo series for this purpose, but the number of design issues that need documentation has grown faster than the resources available to document them. There was community consensus that the NDN team must prioritize this challenge.

10. MOVING FORWARD

Participants made several suggestions for next year:

- Community office hours for assistance on installation of NDN software, in particular to assist more users of NdnCon, NDN-native video-conference software.
- More time between sessions for talking with each other at next community meeting.
- Feedback from broader user communities of software base.
- Hyperbolic routing explanation (or make it an online NDN seminar).
- Interest group meetings in between community meetings.
- European hackathon.
- Tutorial material organization.
- Finish design memos to explain NDN architecture design decisions

11. RESULTS OF 1ST NDN HACKATHON

The weekend before the community meeting, the NDN team held the first NDN Hackathon [7], with the goal of fostering collaboration and advance the state of NDN in multiple directions. During the community meeting, **Steve Benedetto** (Colorado State) summarized the results of the hackathon, which was attended by 25 participants who worked on the following seven (of 19 originally proposed) projects:

- Access control with NFN, led by Christian Tschudin (U. Basel)
- Adopt Mini-NDN for testing any project based on NFD, led by Vince Lehman (U. Memphis)
- NFD on Windows, led by Junxiao Shi (U. Arizona)
- New cache decision policies for NFD, led by Klaus Schneider (U. Arizona)
- Port NDN-RTC to Ubuntu Linux, led by Peter Gusev (UCLA REMAP)
- Demonstrate NDN over Bluetooth LE on the Arduino, led by Jeff Thompson (UCLA REMAP)
- Modify NDNS code to execute in the latest ndnSIM, led by Spyros Mastorakis (UCLA)

On Sunday night, each group presented the final results of their project to NDN Consortium members, who voted to select the winning project: "Demonstrate NDN over Bluetooth LE on the Arduino". The winners received a Raspberry Pi2 pre-loaded with the NDN software and toolkits. The resulting code is available online under <https://github.com/ndncomm>. The projects are expected to continue to be developed beyond the hackathon.

12. POSTERS AND DEMONSTRATIONS

The evening session of the first day showcased ten posters and demos (some of which also had earlier lightning talks):

- Patrick Crowley (Washington U., St. Louis), Experiments with the Emulated NDN Testbed in ONL
- Susmit Shannigrahi (Colorado State), NDN for scientific data
- Lei Liu (Fujitsu Laboratories of America), Bandwidth-Efficient Solutions for Seamless Producer Mobility Support in Named Data Networking
- Ryota Ohnishi (Panasonic), Adaptive Rate Control integration on NDN-RTC
- Minsheng Zhang (U. Memphis), PartialSync: Synchronizing a Partial Namespace in NDN

- Andrew Brown (Intel), NDN and the Internet of Things: Analytics Everywhere
- Spyridon Mastorakis (UCLA), ndnSIM 2.0: A new version of the NDN simulator
- Divyashri Bhat (U. Massachusetts Amherst), Load Balancing Approach for Adaptive Bit-Rate Streaming in Information Centric Networks
- Zhehao Wang (UCLA REMAP), Hierarchical storage in Building Management System
- Haitao Zhang (UCLA), NDNFit architecture and progress
- Jiachen Wang (UCLA REMAP), NDN-RTC testing and updates
- Vince Lehman (U. Memphis), Hyperbolic vs. Link-State Routing in ICN
- Anders Plymouth (TelHoc), TelHoc Borrego - An Application Service Platform supported with NDN networking
- Dustin O'Hara (UCLA REMAP), UX Design for NDN ID Manager & Data Sharing App
- Teng Liang (U. Arizona), NDN-Home: NDN based Smart Home System

13. ACKNOWLEDGMENTS

The workshop was partially supported by the National Science Foundation CNS-1345286, CNS-1345318, and CNS-1457074. We thank the NDNcomm Program Committee members for their effort of putting together an excellent program. We thank all participants for their insights and feedback at the workshop.

14. REFERENCES

- [1] jndn-management.
<https://github.com/01org/jndn-management>.
- [2] jndn-utils.
<http://github.com/01org/jndn-utils>.
- [3] Mini-NDN.
<https://github.com/named-data/mini-ndn>.
- [4] NDN certificate format version 2.0.
<http://named-data.net/doc/ndn-cxx/0.4.0/tutorials/certificate-format.html>.
- [5] NDN Certifications.
<http://ndncert.named-data.net>.
- [6] NDN-RTC. <https://github.com/remap/ndnrctc>.
- [7] NDNcomm 2015 Hackathon participant guide.
<https://ndncomm.github.io/>.
- [8] NdnCon.
<https://github.com/remap/ndncon/wiki>.
- [9] NDNLpv2. <http://redmine.named-data.net/projects/nfd/wiki/NDNLpv2>.
- [10] NFD - Named Data Networking forwarding daemon.
<http://named-data.net/doc/NFD/current/>.
- [11] NLSR - Named Data link state routing protocol.
<http://named-data.net/doc/NLSR/current/>.
- [12] Public key info base (PIB) service.
http://redmine.named-data.net/projects/ndn-cxx/wiki/PublicKey_Info_Base.
- [13] NDN community meeting (NDNcomm 2015): Architecture, applications, and collaboration. <https://www.caida.org/workshops/ndn/1509/>, 2015. Accessed on October 12, 2015.
- [14] Alexander Afanasyev, Cheng Yi, Lan Wang, Beichuan Zhang, and Lixia Zhang. SNAMP: Secure namespace mapping to scale NDN forwarding. In *Proceedings of 18th IEEE Global Internet Symposium (GI 2015)*, April 2015.
- [15] Cross, N. and Dorst, K. Co-evolution of problem and solution space in creative design. 1999.
- [16] Peter Gusev and Jeff Burke. NDN-RTC: Real-time videoconferencing over Named Data Networking. In *Proceedings of the 2nd International Conference on Information-Centric Networking*, pages 117–126, 2015.
- [17] Vince Lehman. Hyperbolic vs. link-state routing in ICN. Poster at NDNcomm 2015, 2015.
- [18] Spyridon Mastorakis, Alexander Afanasyev, Ilya Moiseenko, and Lixia Zhang. Technical report.
- [19] Ilya Moiseenko, Lijing Wang, and Lixia Zhang. Consumer/producer communication with application level framing in Named Data Networking. In *Proceedings of the 2nd International Conference on Information-Centric Networking*, pages 99–108, 2015.
- [20] NDN Team. FIA-NP: Collaborative research: Named Data Networking Next Phase (NDN-NP) proposal. Technical Report NDN-0026, NDN Project, August 2014.
- [21] NDN Team. NDN packet format specification.
<http://named-data.net/doc/ndn-tlv/>, 2015.
- [22] Tian Song, Haowei Yuan, Patrick Crowley, and Beichuan Zhang. Scalable name-based packet forwarding: From millions to billions. In *Proceedings of the 2nd International Conference on Information-Centric Networking*, pages 19–28, 2015.
- [23] Charlie Wiseman, Jonathan Turner, Michela Becchi, Patrick Crowley, John DeHart, Mart Haitjema, Shakir James, Fred Kuhns, Jing Lu, Jyoti Parwatarikar, et al. A remotely accessible network processor-based router for network experimentation. In *Proceedings of the 4th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, pages 20–29, 2008.
- [24] Cheng Yi, Alexander Afanasyev, Ilya Moiseenko, Lan Wang, Beichuan Zhang, and Lixia Zhang. A case for stateful forwarding plane. *Computer Communications*, 36(7):779–791, 2013.
- [25] Yingdi Yu, Alexander Afanasyev, David Clark, kc claffy, Van Jacobson, and Lixia Zhang. Schematizing and automating trust in Named Data Networking. In *2nd ACM Conference on Information-Centric Networking*, September 2015.
- [26] Zhenkai Zhu and Alexander Afanasyev. Let's ChronoSync: Decentralized dataset state synchronization in Named Data Networking. In *Proceedings of the 21st IEEE International Conference on Network Protocols (ICNP 2013)*, 2013.