DOI:10.1145/2915915

Considering the social impact of a proposed future Internet architecture.

BY KATIE SHILTON, JEFFREY A. BURKE,
KC CLAFFY, AND LIXIA ZHANG

Anticipating Policy and Social Implications of Named Data Networking

THE INTERNET HAS become a critical platform for economic, political, cultural, and social activity. The technology behind the Internet continues to evolve, with ramifications for not only the technologies that govern network and application functions, but also for social, economic, and legal concerns. Internet protocols impact not only the basic performance and reliability of Internet services, but also impact debates about fairness issues in content delivery, free speech, trust and cybersecurity, privacy and intellectual property, and control over content.

This article discusses a proposed future Internet architecture that changes how data is delivered over the Internet. Named Data Networking (NDN) is

a prominent example within the broader research field of information-centric networking (ICN). We cannot fully predict how changing protocols will change policy outcomes: social impacts of technology are caused by an interdependent mix of technological decisions, user decisions, and social and policy contexts.^{4,24} But if we take seriously the notion that running code shapes rights, behavior, and governance,^{16,22} then analyzing how NDN would alter that code—the technical infrastructure we rely on every day—is an important challenge.

This article addresses this challenge by beginning a conversation about the social impacts of NDN, with a particular focus on content producers and consumers. We describe the building blocks of NDN; its request-response data exchange is inspired by the Web, but functions at a more fundamental level in the protocol stack. NDN uses data names for routing and forwarding, provides per-packet data signatures, and leverages in-network storage.^a We provide a scenario to illustrate the interactions of these building blocks and describe how the proposed changes could expand options for free speech,

a Many of these techniques are implemented in the application layer of today's Internet. NDN enables them at the network layer, which encourages applications to comport with them.

» key insights

- **NDN is a proposed future Internet architecture that changes the technical protocols that support applications, with implications for social, economic, and policy dimensions of today's Internet ecosystem.**
- **These implications affect a range of stakeholders, including content producers, consumers, regulators, and network operators.**
- **For consumers, NDN can expand options for free speech, security, privacy, and anonymity, while raising new challenges for data retention and forgetting. For governments and content industries, NDN raises new challenges and possibilities for control of content, and for ensuring neutrality across public networks.**

Figure 1. NDN (right) replaces the “thin waist” of the Internet; in its design, the common protocol is the exchange of named, signed data packets instead of IP packets (left).

security, privacy, and anonymity, while raising new challenges regarding data retention and forgetting. We will address impacts for governments and content industries caused by changing the way networked data is identified, handled, and routed as well as examine how these changes raise new challenges and possibilities for ensuring neutrality across public networks. Taken together, this anticipatory analysis suggests research questions and areas of technical focus for ongoing NDN research, and helps us better understand the potential consequences of information-centric networking.

Fundamental Architectural Components of NDN

A team led by Principal Investigators (PIs) from UCLA, and involving Co-PIs, staff, and students from U.S. institutions and international collaborators,^b is designing and evaluating the NDN architecture, which could serve as a new foundational layer of the Internet (see Figure 1). Today, the Internet Protocol (IP) relies on host addresses to route packets across the network. In contrast, NDN delivers based on data names directly, without using host ad-

resses of either source or destination. Rather than forwarding packets based on the *where* of IP, NDN focuses on the *what*: the named data itself. NDN relies on four key architectural components to achieve secure, efficient data delivery: names, request/response data exchange, data signatures, and in-network storage, described in detail in Zhang.³¹

Names: The crux of NDN. In NDN, applications name data at packet granularity. For example, `/edu/ucla/cs/CS217/video1/v2/s3` could refer to segment 3 of version 2 of “video1” published by the teacher of course CS217 in the UCLA-provided namespace.^c The NDN design assumes that application developers will develop standard naming conventions, such as content versioning and segmenting, to aid interoperability and code reuse. NDN also supports hierarchical name structures to facilitate trust management and scalable routing, similar to how hierarchical IP address allocation has enabled global scaling of Internet routing. Globally unique names will require coordinated management and governance,^d but the architecture also

^c Our examples show hierarchical, human-readable NDN names, though the architecture supports arbitrary byte sequences.

^d Just as IP address governance is not a part of the IP architecture, global namespace governance is not an explicit part of the NDN architecture.

supports local names intended for local use (for example, to refer to “the light switch in this room.”) So while all communication in NDN relies on data names, name mechanics will vary based on application context.

Request/response data exchange for multicast delivery. NDN dictates a closed-loop communication model based on packet-by-packet request and response (Figure 2). The model resembles Web semantics but at a per-packet granularity. A consumer sends an Interest packet specifying the name of data she wishes to receive. NDN routers may be able to use cached data to answer that Interest. All data that has previously passed through an NDN router can be cached in its Content Store. (IP routers also have packet buffers due to statistical multiplexing, however a buffered packet is removed from the buffer once it is forwarded to the intended destination). If an Interest cannot be answered with data from an NDN node’s Content Store, the node’s Forwarding Interest Base (FIB) defines where to send the Interest. Nodes use *longest prefix matching* to match data names requested in Interests to data names in the Content Store, and then forward Interests toward nodes that have registered data name prefixes, analogous to IP forwarding.

Each node also uses a Pending Interest Table (PIT) to record the interface, or *face* in NDN parlance, from which it received the Interest. Unlike the FIB and Content Store, the PIT is a fundamentally new entity without an analogy in IP. PIT entries track Interest packets that have been forwarded, to enable data to be returned along the path taken by the Interests. Each PIT entry records the requested data name, the incoming face(s) of the Interest(s), and the outgoing face(s) to which the Interest has been forwarded. Interest propagation creates a hop-by-hop trail of “breadcrumbs” back to the consumer for each path the Interest takes. When the Interest packet reaches a node with matching data, the node responds with a data packet, which is forwarded back along the trail, consuming (such as, deleting) the PIT breadcrumbs along the way.

The request/response model of NDN enables inherent multicast data delivery, as requests for the same data packet from multiple consumers are

^b For a full list of participants and collaborators, see the Named Data Networking website: <http://named-data.net/>.

collapsed into a single PIT entry when they flow through the same router. For example, if a router receives Interests with the same name from five of its faces, the router only forwards the first Interest for that name while recording the incoming faces for the other four Interests in its PIT. When the corresponding data packet comes back, the router forwards that matching data back out to all five faces.

PIT state enables control of traffic load by limiting the number of pending Interests to achieve flow balance. (Only one Interest and one data need to traverse any link for all requestors to be satisfied.) The PIT state can also be used to mitigate DDoS attacks by setting an upper bound on the number of PIT entries allowed.

An NDN network is *loop-free* because each node keeps an entry for each outstanding Interest in its PIT, detecting and discarding duplicates. Each node forwards an Interest to multiple upstream nodes simultaneously and uses the feedback loop created by the request/response structure to evaluate packet delivery performance across its faces—for example, different networks peering with a router or different wireless links on a mobile handset.

Data signatures for provenance and security. Another fundamental aspect of NDN is its use of cryptographic signatures within data packets. NDN requires each data packet to be signed by a key that binds the content to its name. A key locator field encodes the name of the packet's signing key. NDN does not dictate how the consuming application evaluates whether to trust the key. This *data-centric* approach secures the data packet independently of how it is communicated, in contrast with channel-based models such as TLS/SSL on the current Internet.

An active area of research focuses on defining a set of well-understood trust models from which application developers can choose. Within a given trust model, signatures enable determination of data packet provenance, and serve as the basic building block of security in NDN,²¹ including encryption-based access control.³⁰ A valid signature by a trusted key is a strong indication that the data is what it purports to be, regardless of from

where the data was retrieved. The NDN research team is experimenting with a variety of hierarchical, web-of-trust, and evidentiary trust models that use features of NDN for efficient key dissemination and evaluation of trust relationships.²⁹

Diverse and pervasive storage at the network layer. Because NDN applications do not care from where requested data is retrieved, any NDN node can answer an Interest if it has corresponding data. This feature enables an NDN network to take advantage of diverse and pervasive forms of storage to yield performance and scalability enhancements, and also provides support for disruption-tolerant networking (DTN). NDN networks can republish data from the local storage of any nearby device, use router memory as data caches, and deploy persistent repositories that

work with any NDN content. Through these means, NDN provides features similar to today's content distribution networks^e (CDNs), but at the network layer, and thus available consistently for *all* data, without contractual agreements between content producers and CDN providers. This is an active area of exploration; for example, NDN researchers are developing new primitives to interact with repos and support efficient synchronization among named data collections.¹⁸

These four abstractions combine and interact to form an NDN network. Naming data necessitates the request/

e CDN services replicate data across a geographically distributed network connected to the IP Internet, moving content close to high concentrations of users to provide faster data access over a broader area (often globally) than a traditional Web hosting model.

Figure 2. Request/response data exchange.

response data exchange. Stored named data can serve future requests, unlike destination-specific IP packets. And because data can be served from anywhere, it must be signed to protect its provenance and integrity.

An NDN Scenario: The Internet of Things

A use case that illustrates the possibilities of NDN is the Internet of Things (IoT). The IoT concept envisions every device, and many objects, as network-enabled, context-aware (to varying extents), and often integrated with Web and mobile applications. We introduce this case, which we will draw on throughout the article, to orient readers to the ways in which NDN's technical changes shape a wide variety of social issues in a realistic application environment.

In an NDN IoT, names provide a richer and more versatile approach to addressing potentially billions of devices across the world, and the architecture's use of cryptographic signatures for each packet provide a valuable security building block not present in IP. NDN enables the Internet-connected "things," and the data they create and consume, to be addressed by one or more application-specific names at the network layer, often without requiring further middleware or gateways.^{3,8,28} For example, a manufacturer-assigned name, such as `/local/appliance/kitchen/toaster/Black&Decker/<serial_number>`, might be used to address a kitchen appliance from another device in the same smart home. That appliance would be configured in this namespace at the factory and respond to Interests in its prefix `/local/appliance` using a power line or wireless interface. In a simple scenario, other devices in a home (for example, a user's phone) could issue Interests on a regular basis. Interests for `/local/appliance` would be used to discover the device when first plugged in; then, its more specific name could be used for direct communication. In this case, NDN enables applications to use the network layer directly to discover nearby devices in these well-known namespaces (for example, `/local/appliance`), without needing the devices to be connected to the global Internet. At the same time, they share the same

network layer protocol as all other NDN Internet applications, providing opportunities for straightforward integration with local or global Web applications, using data signatures and encryption-based access control for security. This example in the IoT domain illustrates that semantic classification can facilitate discovery of new devices on a network—from a new lightbulb to a digital television—using names.

Policy and Social Implications Of NDN's Components

By fundamentally altering the concepts used to design networked applications and the components available to build them, a transition from IP to NDN could impact policy issues including free speech, security, privacy, content regulation, and network neutrality. Some changes are difficult to predict because Internet infrastructure purposefully provides adaptable mechanisms and interpretive flexibility.¹² But even during the design stage, we can articulate a few important ways NDN would likely change the nature of Internet interactions. Here, we explore how the NDN architecture could improve free speech; improve trust and security; both improve and challenge privacy; complicate content regulation by governments and industry; and introduce open questions for network neutrality.

Improvements to free speech. As the IoT example illustrates, NDN facilitates the development of environments where local devices can transmit content without reliance upon global infrastructure providers. Data packets can be stored and republished by anyone using any device, expanding the options for data dissemination and enhancing and expanding opportunities for communication and free speech.

Consider a regime with authoritarian tendencies that allows Internet access but constrains what is published. NDN makes it easier than IP to share data via alternative communications paths and opportunistic connectivity (toasters and phones as well as laptops and routers), without global infrastructure or complex intermediate services providing indirection or anonymization. Users moving in cars or planes or people with ad hoc wireless on their mobile devices can exchange

data via NDN by leveraging storage on their devices and intermittent connectivity to pass content around, without leaving traces of where the data originated. Any NDN node with access to multiple networks, for example, wireless and wired connections, can bridge those networks by forwarding and/or satisfying Interests, increasing the number of paths data can take to a consumer.³¹ Moreover, namespaces can be locally scoped or encrypted, which can render NDN's data exchange mechanisms and decentralized communication capabilities even more tolerant of disrupted connectivity than IP.

Today, blocking a small number of well-known websites is an effective censorship scheme.⁶ Enabling decentralized communication at the lowest layers of the network can allow users to route around censorship, creating positive impacts for free speech. For example, NDN would enable a group of phones at a protest to use data *muling*—a combination of data storage and direct device-to-device communication in which the phones carry data (and keys) from place to place rather than relying on infrastructure that might be subject to global surveillance. Individually signed packets of a sensitive video, or the keys to verify that video, can be reassembled by any device based on common naming conventions, and verified as being from the same publisher using data signatures. Such peer-to-peer muling can occur in IP networks, but is more complicated at the network and application layers. In addition, NDN content producers could encapsulate or encrypt data names to hide traffic and thwart attempts to block content based on its name.

NDN's emphasis on data signatures could complicate a social mechanism often relied upon to protect free speech: anonymous content production. Therefore, NDN's improvements for free speech must be weighed against its challenges to anonymous speech.

Improvements for trust and security. NDN requires all data be signed so that applications can verify the publisher of received content. In the IoT scenario, each networked device in a home would sign content, enabling applications such as lighting control or energy monitoring services to ver-

ify that data they receive, including commands, originates from a trusted source. Because per-packet signatures are part of the architecture and therefore not dependent on an application or domain, NDN will increase recognition of, and reliance on, data provenance to improve data security and thus consumer trust in content. In the IP Internet, provenance must be established on a per-application basis, and is currently established intermittently and inconsistently. NDN's signature mechanisms can help verify provenance even for orphaned data (data with no online application). Content signatures can also reduce risks such as spoofed data and phishing. Including such provenance explicitly in packets mitigates concerns about data tampering en route.

To take advantage of NDN's security features (in particular, the per-packet cryptographic signatures), application developers will require new trust models that can be used by classes of applications, as well as frameworks for establishing, exchanging, and revoking keys within data-centric networks. These challenges, as discussed earlier, are the most significant for NDN architecture development. Fortunately, we believe there will be increasing incentives to develop such trust models and key distribution mechanisms over time, as they are necessary not only for NDN, but for better security in all networked communications.

Finally, NDN's request/response data exchange provides benefits for network security by mitigating common problems in today's IP Internet, such as distributed denial of service (DDoS) attacks. Since each Interest retrieves at most one data packet, a router can use the PIT (as described previously) to control the number of pending Interests to achieve flow balance, mitigating volumetric DDoS attacks. Techniques for NDN DDoS mitigation have been explored extensively in other work.^{1,14}

Improvements and challenges for privacy. NDN's four fundamental architectural departures from IP have implications that can both challenge and benefit user privacy. The request/response data exchange increases *anonymity of information seekers*, while content signatures and names compli-

Because NDN applications do not care from where requested data is retrieved, any NDN node can answer an Interest if it has corresponding data.

cate *anonymity for content producers*. The architectural emphasis on in-network storage presents new challenges for *limiting data retention*.

Support for anonymous information-seeking. NDN's request/response data exchange improves support for anonymous information-seeking: there is no source address in an Interest. Though Interest packets create a trail in the PIT as they travel toward a data packet, each router's table indicates only the next hop and these PIT entries are erased as soon as a data packet satisfies the outstanding Interest(s). Although routers could log such trails of breadcrumbs, users are not likely to have their Interests traced back to them unless an actor (an authoritarian regime, for example) can access and correlate state across all routers in the (possibly many) paths that data packets have taken. The IoT scenario illustrates how difficult enacting this level of control would be: those paths would likely include privately-owned devices in homes and buildings, in addition to routers owned by Internet Service Providers (ISP)s. So while ISPs might log Interests and forward them to governments, decreased reliance on ISPs as the sole source of connectivity would circumvent such logging. Providing anonymous data retrieval could substantially benefit privacy, allowing individuals to consume controversial content without fear of embarrassment or harm.^{13,26}

Challenges for anonymous content production. Compared to *consumer anonymity*, *content producer anonymity* in an NDN network is difficult to achieve. Data producers can be identified in more than one way—for example, by the key used to sign the data, the namespace in which the data or key are published, or by the content itself.

While NDN data must be signed, it may be signed with ephemeral keys or keys unlinked to real-world identities. Encryption of both names and data can be used to provide confidentiality. But NDN's pervasive use of signatures may make it easier for infrastructure providers and content consumers to *require* signatures that use verified, real-world identities. For example, online forum moderators struggling with trolls and sock puppet accounts—or trying to discriminate against certain

users—might not accept comments sent in packets without verified, real-world signatures. Namespace ownership records may also reveal publisher identities, similar to today's WHOIS database. Thus, another important area of NDN research is trust schemes that provide alternatives to real-world identity for content authentication.

NDN researchers have explored special routing approaches to preserve content source anonymity.¹⁷ Content producers might desire anonymity to participate in free speech, evade censorship, and experiment with multiple online identities.²⁵ Unfortunately, anonymity is also used to evade prosecution for criminal behavior or support mob behavior and hate crimes.¹⁰ Though designing a network architecture to prevent all criminal behavior is an impossible (and, we believe, undesirable) goal, it is worthwhile to consider the benefits and costs of measures to increase content producer anonymity as the project goes forward.

Improvements for content access control. As mentioned earlier, the NDN architecture encourages applications to secure data by encrypting it rather than relying on channel-based security over which the data flows, as is currently done through secure sockets layer/transport layer security, (SSL/TLS), virtual private networks (VPN)s, and similar schemes on IP networks. In the IoT example, there is no need to set up secure connection between two communicating devices, because any potentially sensitive data is encrypted by the application. Securing the data directly should reduce the impact of now-common perimeter and channel security compromises, while still leveraging NDN caching for group communication.

Once published, encrypted data can be replicated and hosted in many (potentially hostile) locations, although only those with access to the right keys can decrypt the information. In this way, NDN makes explicit what is already implicit in schemes like SSL/TLS: encrypted data in transit can be sniffed and stored by others. NDN makes it easier to request a chunk of someone's encrypted data (for example, by sending Interests for common namespaces like /local/appliance), and that encrypted data might be

cached anywhere. Encrypted data may be widely available for extended periods of time, increasing the long-term potential for unauthorized decryption. Content access control will thus require careful design and integration of modern encryption mechanisms and techniques, such as forwarding secrecy and long-term encrypted storage. Further, NDN's integrated use of cryptography also will require navigating open challenges such as the computational burden of encryption in resource-constrained environments (like the IoT) and the challenges of key distribution and revocation.⁹

Challenges for the right to be forgotten. As personal data proliferates on the Internet, there is increasing concern that such data cannot be erased or forgotten. The specter of total accountability for our past actions is considered unpleasant at best and potentially limiting to social interaction and democracy at worst.^{7,23} International privacy scholars as well as policymakers in Europe have been paying increased attention to data retention and disposal, or the "right to be forgotten."^{7,23,27} More recently, California adopted Senate Bill 568, which requires websites to enable minors to easily remove their own posts from websites.

IP routers purge data from buffers as soon as it leaves the routers. That is, they default toward "forgetting" at the infrastructure level, with substantial data retention occurring at the application layer, to support targeted advertising and other purposes. In contrast, NDN routers default toward remembering at the infrastructure level, via content stores and repos. In IP, parties can request that publishers remove data from hosting sites at the edges of the network. Although copies may proliferate elsewhere on user machines, any new request to the hosting site will go unsatisfied. Returning to the IoT example, in NDN, cached copies of data from baby monitors or mobile devices may proliferate on routers, repositories, as well as application-specific stores, and thus remain accessible to Interests. Architectural support for "forgetting" in an NDN world will require mitigation measures, such as time-to-live information in packets, protocols that respect those limits, and further research into self-destructing data.

Challenges for law enforcement and content regulation. The Internet's vital role in cross-border commerce means it contends with diverse national and international policies regulating publication and use of content. Content produced by illegal activities may be restricted (for example, bans on the sale of Nazi memorabilia in France have led to restrictions on content listed in online marketplaces); other forms of content may have use restrictions designed to guarantee a profit to content creators. Enforcing publication and use regulations on content across the global Internet is a challenging task in today's IP Internet. Corporate interests often use the *where* of IP source addresses to enforce market-based restrictions on content access via IP geolocation heuristics. Law enforcement uses a range of tactics—ranging from IP address tracing to deep packet inspection—to track and prosecute both producers and consumers of illegal or pirated content. A transition to NDN will change the tools needed for tracing individuals and monitoring and restricting communications, making current forms of content regulation more challenging, but also potentially more equitable.

Complications for law enforcement. NDN's emphasis on semantic names and data signatures may make certain types of law enforcement easier. For example, keys used to sign data provide strong provenance. In the IoT scenario, the publisher of critical content might be traced by matching the key to identifiable (perhaps registered) devices. And if clear-text data names reflect actual content (for example, data prefixed with /local/PIR was known to be generated by passive infrared security sensors), network-level packet-sniffing and therefore, network regulation could become less computationally intensive. On the other hand, encryption of both NDN names and packet content could mitigate the risk of packet-sniffing. A social shift toward widespread data encryption would raise new challenges for law enforcement. Police and regulatory regimes have long been wary of widespread use of encryption, while developers have resisted providing back doors for law enforcement to inspect or wiretap communications. Encryption would limit the capabili-

Social and policy impacts of NDN for content producers, consumers, regulators, and network operators.

	Free speech	Privacy	Control of content	Network neutrality
Named, signed data available from any node willing to provide it	<i>Improvement for content producers and consumers: can route data around censorship attempts</i>	<i>Improvement for content consumers: surveillance of content more difficult to achieve</i>	<i>Challenge for content producers and regulators: complicates geographic content restrictions</i>	<i>Improvement for content producers & consumers: diversifies interests in tussle over Internet resources</i> <i>Challenge for network operators: diversifies competition</i>
Strong provenance built on data signatures and straightforward key distribution	<i>Improvement for content producers and consumers: increases trust in provenance of speech</i> <i>Improvement for network operators: Increases information available for network strategies</i>	<i>Challenge for content producers: may identify content producers</i>	<i>Improvement for content producers and regulators: may help identify infringing content</i>	<i>Challenge for content producers and consumers: may enable discrimination based on data type or origin.</i> <i>Improvement for network operators: Increases information available for network strategies</i>
Data persistence via uniformly accessed, pervasive storage	<i>Improvement for content producers and consumers: data persists even when subject to takedowns</i>	<i>Challenge for content producers and consumers: may increase likelihood of decryption by unauthorized parties</i>	<i>Challenge for content producers and regulators: complicates content control</i>	<i>Improvement for content producers and consumers: diversifies interests in tussle over Internet resources</i> <i>Challenge for network operators: Incentives for hosting caching unclear¹</i>
Request/response model of data exchange	<i>Improvement for content consumers: ensures anonymity for content seekers; can route requests around censorship attempts</i>	<i>Improvement for content consumers: ensures anonymity for content seekers</i>	<i>Improvement for content regulators: may suppress requests by name</i>	<i>Improvement for network operators: can control traffic load by controlling the number of pending interests</i>

ties of deep packet inspection, used for everything from security concerns to managing traffic flow.⁵

NDN will also change how governments assert regional jurisdiction on the Internet. Today, IP addresses are often used to target law enforcement action.¹⁵ Countermeasures to such targeting in IP include content encryption, encapsulation, and use of third-party resources such as botnets. NDN further disassociates communication from location, as demonstrated by the IoT scenario, which allows communication between devices without any reference to physical geography. This disassociation complicates the identification and geolocation of suspicious activity based on network data. By complicating the use of network data for identification and geolocation, NDN may encourage law enforcement methods that are more effective, such as following financial trails rather than Internet traffic.

Digital rights management. Law enforcement personnel are not the only stakeholders that rely on IP address geolocation capabilities. Sports franchises use them to restrict subscribers

in local markets from watching games online. Gambling operations restrict participation from countries in which such operations are illegal. Search results are tailored to locations. However, one level of indirection, such as Virtual Private Networks (VPNs), can often circumvent IP-address-based control. In NDN, stakeholders might need to rely on application-layer identity and location information to enforce such content restrictions. Although interests can come from anywhere, stakeholders could build systems of encryption and key distribution based on location-verified subscribers.

Digital Rights Management (DRM) typically involves controlling distribution of content, and controlling whether consumers can redistribute that content. NDN supports the first kind of DRM well, but makes republishing easier than it is with IP. As in the IP Internet, copyright holders can distribute verified, encrypted media, and consumers can access the content with the proper key. However, widespread encryption challenges the benefits of in-network caching, reducing economic

incentives to provide such caching.² Reliance on encryption for copyright enforcement also hinders legitimate reuses of content, such as fair use in educational contexts, critique, and parody. Content producers might enable fair use by giving copies of keys to libraries, or providing portions of the content in the clear for scholarship, critique, parody, or other protected fair uses. But once consumers have received and decrypted verified content, they may distribute unauthorized versions in the clear, a task made easier by NDN.

NDN's in-network storage and caching means that many segments of both licensed (presumably encrypted) and unlicensed (presumably decrypted) media could reside on routers and repos. A world where countless copies proliferate across the Internet challenges assumptions embedded in copyright law, as well as the current mechanisms of copyright enforcement, such as the Digital Millennium Copyright Act (DMCA) takedown notice.¹⁹ On the IP Internet, videos are commonly hosted by major providers such as YouTube or Hulu, which

respond to DMCA takedown notices. However, even these major providers struggle with the scale: as of August 2015, Google was receiving over 12 million URLs requested to be removed from search per week.²⁰ Who would be responsible for taking down an infringing video distributed on thousands of routers by thousands of different organizations across the world? The political economy of repos—who owns them, and in what jurisdictions—will impact the future efficacy of such takedown notices. NDN's in-network storage may increase pressure on lawmakers to redefine intellectual and political understandings of copyright already challenged by pervasive digital duplication.

Network neutrality: An uncertain outcome. The network neutrality debate focuses on what actors pay for Internet resources such as bandwidth and storage, and whether those actors providing resources (for example, ISPs) may throttle or privilege traffic to increase revenue. Consideration of NDN's impact on network neutrality motivates a deeper discussion of NDN node operation. The algorithms and parameters for configuring network-forwarding policy for given data prefixes and links in NDN are typically referred to as *strategies*.^f Strategies are an evolving part of NDN research. In every NDN device, strategies control the operation of three tables—the FIB, the Content Store, and the PIT. The strategies for these tables affect performance by enabling node owners to express traffic shaping policies in terms of namespaces and faces to other nodes.

Interest forwarding strategies. Routing protocols and/or manual setup of static routes are used to configure forwarding strategies in the FIB. The resulting configuration expresses the policies of router administrators, who may choose to discriminate based on data types (indicated within data names, for example, /local/toaster) or namespace of publication (ucla/cs/local/toaster). Such traffic discrimination may occur in IP

but at higher layers, for example, HTTP or via the Domain Name System (DNS) names. NDN routers will be capable of such choices at the network layer.

Content store strategies. All data that passes through an NDN router can be cached in the content store, and persists according to a router's configured caching policy. NDN spreads caching and its costs across the Internet infrastructure, which democratizes content storage functions and introduces new stakeholders into the tussle over Internet resources. Researchers are considering economic incentives for deploying caches and markets for cache participation.² In-network storage will also impact the political economy of content dissemination. Given enough in-network caching, content producers on an NDN Internet can use a cheap server and low-bandwidth connection to make their viral videos reachable by millions of interested viewers, with the network providing scalability to handle content requests. Thus, NDN could reduce dependence on third-party services to scale content distribution. Users could share content on their own terms, rather than being subject to a third-party provider or hosting service's terms.

Pending interest table (PIT) strategies. Because the PIT records which Interest packets have been forwarded, and then waits for data packets to return, policies that modify how long to retain Interests in the PIT could impact data retrieval performance. (While a field in each Interest specifies a lifetime, it is up to each forwarder to obey that field.) Whether consumers or namespace providers are able to influence the quality of service through longer Interest storage in the PIT, or more aggressive re-issuing of Interests across multiple outgoing faces, are strategy configuration questions that could impact a node's neutrality.

Neutrality implications of NDN node strategies. The actors controlling NDN traffic routing decisions are likely to be more diverse than on an IP Internet. Nonetheless, an NDN-based Internet's ISPs will continue to have incentives (if not obligations) to author strategy modules to manage the tables in their routers and prioritize data with certain types or names. Data names may reveal types of content, such as IoT, video, scientific data, or emergency response data. Globally routable name prefixes

expressing data origin, such as /edu/ucla or /com/nytimes, may indicate institutional power or status. Signing keys may also reveal data origin. Future standards might use names or signing keys to prioritize particular interest/data exchanges, such as emergency response traffic. Executing such prioritization across multiple providers would bear the same policy complexity and risk as with attempts to do multi-provider QoS on today's IP networks.¹¹

In an NDN network, routing provides only one of the input factors for forwarding decisions; locally configured forwarding strategies make the final decision on which Interest is forwarded along which path, or forwarded at all. In this way, NDN's inherent support for mobility and disruption-tolerant networking could mitigate the threat of harmful traffic discrimination. Even if prioritized networking evolves using semantically-meaningful names or pay-for-retention policies on routers, NDN's ability to forward requests around providers that do not respond efficiently will give consumers more options for data transmission. Small ISPs can use multipath forwarding to choose forwarding paths based on performance measurements. NDN enables providers to route around ISPs that throttle traffic based on certain names, which will provide a disincentive for such throttling.

Conclusion

NDN brings some of the semantics of the current Internet's application layers to the network layer, providing technical benefits to application developers, network operators, and end users. NDN's architectural design decisions are also likely to have implications for social and policy issues in both layers, including some of today's most pressing challenges: free speech, privacy, control of content, and network neutrality. Whether these impacts are benefits or drawbacks depends upon stakeholder perspective. The accompanying table maps NDN's technical implications to social and policy impacts for a range of stakeholders: content producers, consumers, regulators, and network operators.

By diversifying the nodes that can provide data, NDN will likely improve conditions for free and anonymous speech and information seeking for consumers and producers. NDN's

^f Here, we use strategies more broadly than the NDN architects have so far, using this term to cover any policy choice that can be made in an NDN node that does not violate the "thin waist" of the architecture as currently understood.

strong provenance built on data signatures will identify content producers. Strong provenance will enhance trust in, and security of, content, while simultaneously complicating anonymous information production. Strong provenance may also help content producers identify infringing content, and signatures provide a mechanism to help producers secure content with encryption-based access control. But pervasive storage and request/response data exchange will challenge producers interested in content control and geographic access restrictions. Finally, network neutrality is a complicated outcome to predict. Future decisions in naming and routing may hinder network neutrality, as the use of names for routing could facilitate new forms of traffic discrimination. At the same time, NDN will promote increased competition among network operators by enabling applications to efficiently route around infrastructure that constrains their traffic.

This article has sought to address policy and social implications of the network that are significant departures from today's IP Internet. As such, it has not addressed Internet policy topics that remain closely tied to existing challenges in IP, which are areas for future work. For example, NDN faces challenges in globally routable naming rights management similar to those of IP. We have also not addressed application-level policy issues such as the relationship between advertising data collection and privacy or application-level regulations such as accessibility requirements or required geolocation services such as E911.

However, identifying open questions relevant to the network layers illustrates an advantage of anticipatory policy studies. Analyzing potential social and policy impacts of the NDN architecture can help prioritize research questions within the NDN project and broader content-centric networking initiatives. The practical impact of NDN will depend on future directions in several open research areas: (1) balancing meaningful names to simplify application development with opaque names to protect privacy; (2) standardizing mechanisms for cryptographic key assignment, distribution and revocation; (3) developing usable design patterns for managing trust in a broad range of applications; (4) providing usable, secure imple-

mentations of more complex multi-participant encryption schemes; and (5) creating fair congestion management to enable network neutrality.

Most of NDN's potential policy impacts are speculative, in part because we are exploring them while the architecture design is still evolving. Yet imagining the social changes NDN might encourage during the design process provides opportunities for pro-social computing research. We hope this work will spark continuing discussion of the current and future Internet's impact on society. Thinking creatively about changes can help us better understand the relationship between infrastructure and our world.

Acknowledgments. The authors thank Van Jacobson, David D. Clark, Paul Ohm, Charles Duan, Steven Bellovin, and anonymous reviewers for feedback and ideas that shaped this work. This research was supported by the National Science Foundation under grant numbers CNS-1040868, CNS-1421876, and CNS-1345318.

References

- Afanasyev, A., Mahadevan, P., Moiseenko, I., Uzun, E. and Zhang, L. Interest flooding attack and countermeasures in Named Data Networking. *IFIP Networking Conference*, 2013, 1–9.
- Agyapong, P.K. and Sirbu, M. Economic incentives in information-centric networking: Implications for protocol design and public policy. *IEEE Commun. Mag.* 50, 12 (2012), 18–26; <http://doi.org/10.1109/MCOM.2012.6384447>
- Bannis, A. and Burke, J.A. *Creating a secure, integrated home network of things with Named Data Networking*. UCLA, 2015, Los Angeles, CA.
- Baxter, G. and Sommerville, I. 2011. Socio-technical systems: From design methods to systems engineering. *Interacting with Computers* 23, 1 (2011), 4–17; <http://doi.org/10.1016/j.intcom.2010.07.003>
- Bendrath, R. and Mueller, M. The end of the net as we know it? Deep packet inspection and Internet governance. *New Media & Society* 13, 7 (2011), 1142–1160.
- Best, M.L. and Wade, K.W. Democratic and anti-democratic regulators of the Internet: A framework. *The Information Society* 23, 5 (2007), 405–411; <http://doi.org/10.1080/01972240701575684>
- Blanchette, J.-F. and Johnson, D.G. Data retention and the panoptic society: the social benefits of forgetfulness. *The Information Society* 18 (2002), 33–45.
- Burke, J.A., Gasti, P., Nathan, N. and Tsudik, G. Securing instrumented environments over content-centric networking: the case of lighting control. In *Proceedings of IEEE INFOCOMM 2013 NOMEN Workshop*. Retrieved Aug. 7, 2015; <http://named-data.net/publications/nomen13/>
- Chaabane, A., De Cristofaro, E., Kaafar, M.A. and Uzun, E. 2013. Privacy in content-oriented networking: Threats and countermeasures. *SIGCOMM Comput. Commun. Rev.* 43, 3 (2013), 25–33; <http://doi.org/10.1145/2500098.2500102>
- Citron, D.K. Civil rights in our information age. In *The offensive internet: privacy, speech, and reputation*. Harvard University Press, Cambridge, MA and London, 2010, 31–49.
- Clark, D.D., Bauer, S., Lehr, W. et al. Measurement and analysis of Internet interconnection and congestion. In *Proceedings of the 42nd Research Conference on Communication, Information, and Internet Policy*. Social Science Research Network, 2014. Retrieved Aug. 17, 2015; <http://papers.ssrn.com/proxy-um.researchport.umd.edu/abstract=2417573>
- Clark, D.D., Wroclawski, J., Sollins, K.R. and Braden, R. Tussle in cyberspace: Defining tomorrow's Internet. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ACM, 347–356; <http://doi.org/10.1145/633025.633059>
- Cohen, J.E. A right to read anonymously: A closer look at "copyright management" in cyberspace. *Connecticut Law Review* 28 (1996), 981–1039.
- Compagno, A., Conti, M., Gasti, P. and Tsudik, G. Poseidon: Mitigating interest flooding DDoS attacks in Named Data Networking. In *Proceedings of the 2013 IEEE 38th Conference on Local Computer Networks*, 630–638; <http://doi.org/10.1109/LCN.2013.6761300>
- Cooke, L. 2007. Controlling the net: European approaches to content and access regulation. *Journal of Information Science* 33, 3 (2007), 360–376; <http://doi.org/10.1177/0165551506072163>
- DeNardis, L. Hidden levers of internet control. *Information, Communication & Society* 15, 5 (2012), 720–738; <http://doi.org/10.1080/1369118X.2012.659199>
- DiBenedetto, S., Gasti, P., Tsudik, G. and Uzun, E. ANDaNA: Anonymous Named Data Networking application. In *Proceedings of the 19th Annual Network & Distributed System Security Symposium*, Internet Society, Retrieved June 26, 2012; <http://arxiv.org/abs/1112.2205>
- Fu, W., Abraham, H.B. and Crowley, P. Synchronizing namespaces with invertible bloom filters. In *Proceedings of the 11th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*. IEEE Computer Society, 123–134. Retrieved Aug. 17, 2015; <http://dl.acm.org/citation.cfm?id=2772722.2772740>
- Gillespie, T. *Wired Shut: Copyright and the Shape of Digital Culture*. The MIT Press, 2009.
- Google transparency report. 2005. Retrieved Aug. 17, 2015; <http://www.google.com/transparencyreport/removals/copyright/>
- Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M., Briggs, N. and Braynard, R. Networking named content. *Commun. ACM* 55, 1 (Jan. 2012), 117–124; <http://doi.org/10.1145/2063176.2063204>
- Lessig, L. *Code: Version 2.0*. Basic Books, New York, NY, 2006.
- Mayer-Schoenberger, V. *Useful Void: The Art of Forgetting in the Age of Ubiquitous Computing*. Harvard University, Cambridge, MA, 2007.
- Pfaffenberger, B. Technological dramas. *Science, Technology & Human Values* 17, 3 (1992), 282–312.
- Phillips, D.J. From privacy to visibility: Context, identity, and power in ubiquitous computing environments. *Social Text* 23, 2 (2005), 95–108.
- Richards, N.M. The perils of social reading. *Georgetown Law Journal* 101, 3 (2013). Retrieved Apr. 15, 2012; http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2031307
- Rosen, J. The right to be forgotten. *Stanford Law Review Online* 64 (2012), 88.
- Shang, W., Ding, Q., Marianantoni, A., Burke, J. and Zhang, L. Securing building management systems using named data networking. *IEEE Network* 28, 3 (2014), 50–56; <http://doi.org/10.1109/MNET.2014.6843232>
- Yu, Y., Afanasyev, A., Clark, D., Claffy, K., Jacobson, V. and Zhang, L. Schematizing and automating trust in Named Data Networking. In *Proceedings of the 2nd ACM Conference on Information-Centric Networking*. ACM, 2015.
- Yu, Y., Afanasyev, A. and Zhang, L. *Name-based access control*. University of California Los Angeles, CA, 2015.
- Zhang, L., Estrin, D., Burke, J.A. et al. *Named Data Networking Tech Report 001*. University of California Los Angeles, CA, 2010. Retrieved Sept. 8, 2014; <http://named-data.net/publications/techreports/tr001ndn-proj/>

Katie Shilton (kshilton@umd.edu) is an assistant professor of information studies at the University of Maryland, College Park.

Jeffrey A. Burke (jburke@remap.ucla.edu) is the assistant dean, technology and innovation, in the School of Theater, Film and Television at the University of California, Los Angeles.

kc claffy (kc@caida.org) is the founder and director of the Center for Applied Internet Data Analysis (CAIDA) at the University of California, San Diego.

Lixia Zhang (lixia@cs.ucla.edu) is the Jonathan B. Poste; Professor of Computer Science at the University of California, Los Angeles.

Copyright held by authors.