Can We Make a Cake and Eat it Too? A Discussion of ICN Security and Privacy

Edith Ngai
Uppsala University, SE
edith.ngai@it.uu.se

Börje Ohlman
Ericsson Research, SE
Borje.Ohlman@ericsson.com

Gene Tsudik
University of California Irvine
gts@ics.uci.edu

Ersin Uzun
Xerox PARC
ersin.uzun@parc.com

Matthias Wählisch
Freie University Berlin, DE
m.waehlisch@fu-berlin.de

Christopher A. Wood
University of California Irvine
woodc1@uci.edu

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 It has NOT been peer reviewed.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

ABSTRACT

In recent years, Information-centric Networking (ICN) has received much attention from both academic and industry participants. ICN offers data-centric inter-networking that is radically different from today's host-based IP networks. Security and privacy features on today's Internet were originally not present and have been incrementally retrofitted over the last 35 years. As such, these issues have become increasingly important as ICN technology gradually matures towards real-world deployment. Thus, while ICN-based architectures (e.g., NDN, CCNx, etc.) are still evolving, it is both timely and important to explore ICN security and privacy issues as well as devise and assess possible mitigation techniques.

This report documents the highlights and outcomes of the Dagstuhl Seminar 16251 on "Information-centric Networking and Security." The goal of which was to bring together researchers to discuss and address security and privacy issues particular to ICN-based architectures. Upon finishing the three-day workshop, the outlook of ICN is still unclear. Many unsolved and ill-addressed problems remain, such as namespace and identity management, object security and forward secrecy, and privacy. Regardless of the fate of ICN, one thing is certain: much more research and practical experience with these systems is needed to make progress towards solving these arduous problems.

CCS Concepts

•Security and privacy → Systems security; Network security;
•Networks → Network architectures;

Keywords

Information-Centric Networking; Security and Privacy

1. INTRODUCTION

Dagstuhl Seminar 16251 on "Information-centric Networking and Security" was a three-day workshop held on June 19-21, 2016 in Dagstuhl, Germany. The goal was to bring together researchers with different areas of expertise relevant to Information-Centric Networking (ICN) to discuss security and privacy issues particular to ICN-based architectures. These problems have become increasingly important as ICN technology gradually matures and gets close to real-world deployment.

Brief History of Dagstuhl Seminars on ICN.

This seminar is the fourth retreat of the ICN community at Schloss Dagstuhl, which usually meets every two years at this place. In 2010, the term ICN was coined to assemble different information-centric networking approaches, such as NetInf (4WARD and SAIL projects) and PSIRP in Europe and Content-Centric Networking (CCN) in the US [1]. In 2012, reality checks were discussed [2]. In 2014, special focus was given to scalability and deployment issues [3]. And finally, at long last, the most recent seminar discussed primarily security aspects of these architectures.

Current Research Challenges.

Many ICN-based architectures have the luxury of starting with a clean slate. As a consequence, threat models in ICN are often distinct from traditional IP-based networks [4,5]. Differentiating factors between the two include new application design patterns, trust models and management, as well as a strong emphasis on object-based, instead of channel-based, security. As ICN develops, it is both timely and important to explore ICN security and privacy issues in order to devise and assess possible mitigation techniques. This was the general purpose of the Dagstuhl seminar. To that end, the seminar focused on the following issues:

- What are the relevant threat models that ICN must be concerned? How are they different from those in the IP-based networks?
- To what extent is trust management a solved problem in ICN? Have we adequately identified the core elements of a trust model, e.g., with Named-Data Networking (NDN) trust schemas?
- How practical and realistic is object-based security when framed in the context of accepted, best-practice privacy measures used in IP-based networks?
- Are there any new types of cryptographic schemes or primitives that ICN architectures could use to (a) enable more efficient or secure packet processing or (b) build a more secure architecture?

The seminar answered (entirely or partially) some of these questions and fueled discussions for others. In the remainder of this document we will review the major themes of the seminar, discuss many open problems that exist in ICN architectures, and summarize their outlook going forward.

2. DRIVING THEMES

There were several driving themes throughout the course of the seminar, including: ICN-related threat models, namespace and identity management, privacy, the so-called "locator and identifier split," ICN and IoT, and future design directions. In this section we summarize the discussions along each theme in the context of general ICN-based architectures. References to specific architectures, such as NDN or CCN, are included where necessary.

2.1 IP Parity and Threat Models

ICN attempts to diverge from IP with respect to the central abstraction of hosts and point-to-point communication between them. The ICN emphasis on named data and object security instead channel security is one clear differentiating factor contributing to this divergence. To quantify the degree by which secret is improved (or worsened), threat models are needed. In general, they must capture particular design challenges in ICN, such as infrastructure protection, user-friendly key distribution and trust management, and content protection and access control. Given the wide gap between IP and ICN, there is a tremendous need for standard and well-scrutinized threat models to use in the design phase.

To give an example of a threat that is unique to ICN, consider the problem of consumer anonymity. By design, the subject and content of a ICN packets are not the facets to be considered in the context of privacy. The origin and destination details (e.g., geographical location or position within a network topology), as well as identity information (e.g., consumer identifying information), can sometimes harm the network users. As shown by [6], ICN caching and interest-collapse mechanisms make ICN itself inherently vulnerable to the possibility for an adversary to locate consumers. Therefore, the threat model must consider these vulnerabilities and adversaries capable of exploiting them successfully.

2.2 Namespace and Identity Management

In ICN, there is an intimate relationship between trust, identity, and namespace management. Furthermore, resource naming, which in the current Internet is primarily an application-layer concern, now directly affects the network layer as well. In an ideal ICN architecture, applications should be able to express their trust preferences (using policies) and let some "middleware" enforce them throughout the network. This raises two important questions: (1) what is the minimum set of policies that can be factored out of all trust models, and (2) what is the middleware that does this enforcement? The trust schema concept pioneered by the NDN architecture [7] is a prime example of a set of rules that can be used to express most trust models. Among other things, they specify what keys are allowed to sign what data. Since both keys and data are named resources in NDN and other ICN architectures, this means that a schema allows for arbitrary hierarchical trust models. It remains to be seen if other non-hierarchical trust models will be so effortlessly realized in ICN-based architectures.

With respect to (2), it is clear that the network should only be responsible for validating at most one signature per packet and doing nothing else to enforce trust models. (Further details about this limitation are provided in [8].) However, despite this limitation, network layer "trust enforcement" should not prohibit or prevent other application-layer trust models. This means that the network functionality must be simpler than that which is supported by the middleware. Functionalities such as certificate chain resolution or key retrieval should not be implemented in the core network. This behavior must be handled by other network nodes (e.g., consumers and producers) or other middleware entities.

Even with an agreement about how and to what extent the network aids in trust enforcement, we are left with the following major question: how are names registered and managed in ICN? Namespace ownership is intrinsically tied to an identity. Thus, namespace advertisements under different namespaces or in different networks must be authenticated with respect to the claimed owner's identity. In this context, an identity is a public and private key pair. The community struggled with issues about namespace scale and the practicality of a global namespace. Questions such as, "how do NATs work in a global namespace?," drove the discussion. No consensus or common understanding about how namespaces and identities should be managed was reached. This is still an area of active research.

2.3 Privacy

Privacy is and has always been an elusive property in ICN-based architectures. In many designs, there is a significant amount of information leaked by packets, including content payloads, signatures, and even the names. One major focus for this seminar was on privacy with respect to names. In this context, we defined name privacy to be the property that a so-called "network name," i.e., the name encoded in a packet, has no correlation or connection with the corresponding content object. Specifically, name privacy means that a network name reveals nothing about the data inside the content object. Ideally, names should reveal no more than what is currently revealed by an IP address and port.

Adding name privacy to ICN-based architectures is no easy task. As a thought exercise, consider how this would work if it were done cleanly, i.e., without some upper-layer service. To restrict the design space, one might make the following assumptions:

- Content may be requested by an identifier (ID) such as its cryptographic hash digest. Moreover, revealing the content ID does not compromise privacy.
- Consumers know the public key of the producer with which they want to communicate.
- Network names have an implicit routable prefix and application-specific suffix. By default, consumers do not know the locator and identifier split in a name.
- Requests may specify the ID of (1) a signature verification key or (2) the expected content.

Under these assumptions, now consider how a consumer might fetch content. There are fundamentally two ways to issue a request: (1) with and (2) without a content ID. In case (1), a request name needs to only contain a routable prefix that will move the request to some cache or producer which can return the corresponding content. These locators can be completely separate from the desired content and, therefore, this approach can satisfy our name privacy goal. However, without implicit knowledge about the locator for some desired content, an upper-layer service is necessary to obtain said information.

In case (2), the application-specific suffix of a name must not reveal anything about the data. To achieve this, it must be encrypted. Name encryption introduces a number of other questions, such as how to obtain the routable prefix, what key to use for encryption, and how to "protect" the result. Assume that the routable prefix is known and that the producer public key is used for name suffix encryption. If the resultant content payload is not encrypted then one may be able to infer the name from its contents. Therefore, the content response itself must also be encrypted. This requires requests to carry a consumer-generated key that is protected in a

CCA-secure envelope. Otherwise, eavesdroppers could replay requests with the same encrypted name but their own key to obtain a decrypted response.

In all cases, it seems that to achieve name privacy then one needs some upper-layer service. Whether its role is to provide the routable prefix for a name, encrypt the response, or to separate a content ID and locator via some other means is an orthogonal issue to be resolved. Also, one critical observation is that name privacy seems to, in most cases, invalidate the utility of shared caches, which puts it at odds with the primary feature of many ICN-based architectures. Thus, it seems as though name privacy is a property that must be abandoned in pure name-based ICN designs.

2.4 Locator and Identifier Split

The notion of locators and how to fetch data with non-topological names (or even topological names that are cached off-path) was another major theme of the seminar. Routing should, possibly, only concern itself with topological names or addresses. Finding data (objects) with non-topological names should not be done in the data plane. It should be done via a service.

In CCN, this service could resolve a named root manifest to then resolve locator names by hash. In NDN, it resolves the link routing hints to allow off-path interest forwarding. In TagNet [9], there is a distinction between Locator names and Descriptor names. Locator names have a strong binding between their name and a point of attachment. Descriptor (hash) names, on the other hand, are free-form and could be present anywhere. One resolves a tag query (of either type) to a topological locator and then does data transfer on that locator.

This lack of uniformity led to one major question: should locators and identifiers be split, as in TagNet, or should they be combined? For example, in CCN, if there is a clear locator field and then a clear identifier tuple (Name, [KeyID Restriction], [Hash Restriction]), one would get full matching expressiveness with the functionality of nameless object locators. A similar approach could be done in NDN, though with a different tuple. There was no consensus on this idea, though it is worth exploring.

There was also some discussion on the benefit of ICN if one still needs to do an external name to address lookup. Specifically, is it worth it if one still needs a DNS-like function? One partial answer is that in the non-global routing space (i.e., data center, maybe IoT, some internal applications, etc.), one could inject all names into the internal routing protocol and realize the full benefit of application-specific names.

2.5 ICN and IoT

The Internet of Things (IoT) is connecting billions of smart devices (e.g., sensors) and is growing very fast. In the future, users may communicate directly to the sensors, or indirectly through the cloud or gateways. Considering the communication patterns and future trends, there may be benefits in applying ICN to the IoT. For instance, the ICN routers connecting to sensors can cache the sensor data to improve the performance of data dissemination. In addition, the users may obtain sensor data directly from the sensors or from a nearby ICN router, without going through the cloud. Although ICN for IoT may provide in-network caching and flexibility in data dissemination, it raises several security concerns:

- How can the sensors be securely configured at the time of initialization?
- How can software updates be performed securely?
- How can we handle mobility in ICN for IoT? For example, each sensor may need a unique publisher identity, which may change with its location. How does mobility affect naming of data and scalability of routing?

There are also several advantages and concerns of caching data at the sensors. Firstly, sensors are resource-limited devices, which may not have sufficient memory for caching the data. However, memory resources in the sensors may increase and the price may decrease in the future. Secondly, it is advantageous to retrieve data directly from the sensors in certain use cases. For example, it is more efficient to control home lighting without going through the cloud. Thirdly, when applying cryptography on the sensors, the encryption time could be long and cause additional delay in data retrieval. Lastly, the sensors have to be always on to listen to the interest packets in ICN, which may consume a large amount of energy. Mitigating this problem might require us to more carefully use scheduling or adaptive duty cycles.

Based on these observations, there are several major outstanding questions. Firstly, what are some critical use cases for ICN in the IoT and how can we experiment with them to better understand the application communication patterns and the related security requirements. Secondly, beyond network stack simplicity, what are other ways in which the IoT might benefit from ICN? Thirdly, what is the best way to configure and bootstrap the sensors securely? And finally, what is the cost of providing security for IoT data with ICN? Answers to some of these questions are already topics of active research [10].

2.6 Access Control

Another challenge unique to ICN is how to develop scalable object-based access control mechanisms. A variety of encryption techniques have been used in the past to protect access to confidential network data [11]. Many design approaches, particularly in CCN and NDN, exist well above the network layer [12–18]. In contrast, publish-subscribe ICNs such as ENCODERS [19] integrate access control into the network layer. It uses multi-authority attribute-based encryption [20] to allow content access to be scoped to selected nodes in the system. Since the system is completely decentralized, peers serve as brokers that match content from the publishers with interests expressed by the subscribers. In order to perform such a match, an intermediate node must be authorized to see both the relevant content tags and the subscriber interests. In this case, access control policies applied to the metadata (content tags and subscriber interests) effectively create reachability constraints that are independent from the one defined by the routing protocols. Consequently, when performed at the network layer, this security-routing interaction must be treated carefully during policy definition.

2.7 Design Exploration

There are many infrastructure security and privacy problems in IP-based networks that we could try to remedy in ICN-based architectures. Can ICNs use the NSEC3 strategy of *authenticated denial* [21] to limit incoming requests for non-existent content as a way of deterring sophisticated distributed DoS (DDoS) attacks? Can we levereage recent advances in deep packet inspection over encryption data [22] to let forwarders blindly route packets without seeing the content or sensitive information, e.g., application names? Can we leverage information-theoretic PIR techniques [23] to support truly private content queries? Can we prevent correlation of static content across multiple consumers by adopting randomizable encryption schemes [24]? And if so, how can we do so while maintaining the integrity of content? There are many features that could improve beyond what is done in IP-based networks. Moreover, there are cryptographic algorithms, schemes, and protocols that could allow us to realize these features. However, it is unclear whether or not these more esoteric cryptographic schemes can and should be applied in the network layer of future networks.

3. OUTLOOK AND OPEN PROBLEMS

After nearly a decade of research, there still exists an abundance of security and privacy problems that the ICN community has yet to adequately address. This leaves much room for future work on a variety of topics, which are elaborated upon below.

3.1 Object-Based Security with Forward Secrecy

Summary. We need to resolve the channel- versus object-based security debate. Should we continue to focus on securing the data instead of securing the pipe? The former promotes a "take what you want" model of group access control which is dependent on long-lived keys. One factor why channel-based security is preferred is because it allows for forward secrecy. Forward secrecy is the property that exposure of a principal's long-term secret keys does not compromise the secrecy of their previously generated ephemeral (session) keys. This is a useful property to have in the presence of eavesdropping attackers intercepting and logging traffic. It minimizes data and key compromise windows and therefore reduces the overall attack surface. However, it requires protocols and techniques for deriving ephemeral keys and then updating them regularly. Without forward secrecy, packet confidentiality is reduced to the efficacy of key management. That is, if private keys are leaked, then intercepted packets can be decrypted.

Given the IP-based push for TLS everywhere that leads to applications built on top of (D)TLS rather than plain TCP (or UDP), object-based security is a significant departure from what is accepted as best-practice today. Can ICN-based architectures make a compelling enough case to motivate applications to revert to this less-secure form of data delivery?

Outlook. Given that ICN focuses on object security, the need for transport protocols that provide forward secrecy should be implemented in higher layers. Attendees found that while most ICN-based architectures do not preclude forward secrecy, forward secrecy should not be a requirement in the network layer. However, a worthy research question is whether or not we can design an object-based security scheme that also provides forward secrecy. Reviewing existing schemes such as the forward-secure public-key encryption scheme of Canetti et al. [25] might prove useful here.

3.2 Namespaces, Identities, and Routing

Summary. As previously discussed, in ICNs, namespace ownership, identities, and the routing fabric are all intimately coupled. Currently, we do not know of a way to break away from a centralized model for namespace management and arbitration. Routing updates therefore have a dependency on this management oracle. Without it, any producer application would be able to advertise any namespace it wants. Schemes such as [26] can be used to remove malicious producers from advertising under incorrect namespaces, but they *do not* resolve conflicts over namespace ownership.

Outlook. The ICN community still does not have a clear answer on how to handle namespace and identity management. While trust management in ICN can be distributed and function without a global PKI, it seems difficult to break away from this model for namespace management and arbitration. This has strong implications on how names are propagated in the routing fabric. Can any producer application advertise any name, anywhere in the network? If not, how can name prefix advertisements be constrained or limited?

3.3 Network Management

Summary. ICN names are user-generated content in FIBs. In effect, FIBs serve as a (globally) replicated name set wherein any name owner can write into the set. The complexity of this state is influenced by the fact that prefix owners can always de-aggregate and create arbitrary names, even if prefixes are restrictively assigned. However, this raises questions of resource exhaustion attacks on FIBs and general complexity attacks (e.g., hash collisions). Newer attacks try to leak information from the FIB contents to target the forwarding plane. Can ICNs be managed to avoid these types of scalability or security problems, or do they necessitate an ecosystem in which any producer can inject any type of information into the network state?

Outlook. This class of problem is tied to how namespaces, identities, and routing are to be handled in ICN-based architectures. Until we have a shared understanding of how namespace ownership, advertisement, and propagation will be controlled, we cannot expect to manage the network state to prevent state exhaustion or mitigate FIB scalability problems.

3.4 Privacy

Summary. Privacy of consumers, producers, and content all remain significant challenges in most ICN-based architectures. As of yet, we have not adequately addressed these privacy problems. In particular, since data names reveal large amount of information to the passive eavesdropper, privacy demands that names and payloads have no correlation. However, achieving this seems infeasible without the presence of an upper-layer service akin to the one that would resolve non-topological identifiers to topological names. Moreover, the trend in the IETF and other standard bodies is to put privacy as a *primary* goal going forward. Consequently, to meet future privacy expectations, many architectures may need to make certain techniques such as onion-based routing [27], name encryption [28], or secure sessions [29] a de facto part of the architecture. Otherwise, it is difficult to foresee the incentives to use these architectures in the real world.

Outlook. Privacy seems difficult to achieve without major architectural changes to ICN-based systems. As of now, it seems as though ICN-based architectures trade privacy for efficiency, which contradicts the perspective of the IETF and beyond. More research is needed to determine if both privacy and efficiency can be achieved in future without such drastic tradeoffs.

3.5 Locator and Identifier Split

Summary. There is still a high uncertainty about whether ICN should split the content locator and identifier. Names in architectures such as NDN and CCN serve as both locator and identifier of the data, though there are extensions that permit explicit locators (e.g., through the use of NDN LINK objects). This distinction is necessary under the common understanding that routing could be more efficient with topological names. Finding data through non-topological names should not be implemented in the data plane as part of the global routing space. However, if we revert to a distinction between topological locators and identifiers, then the unique features of ICN become much more limited.

Outlook. By focusing on named data instead of hosts, many ICN-based architectures blur the line between content locators and identifiers. This has implications on how routing and discovery occurs in the network. Existing architectures differ in packet format and protocol semantics in how these two mechanisms are performed. More research is needed before we, as a community, declare one approach superior to another.

3.6 Common Crypto

Summary. ICN-based architectures have the unique privilege of being able to start from a clean slate without an inheriting any legacy cruft. This often leads to a strong desire to explore the use of young cryptographic primitives and protocols, such as those built on pairings [30]. The mistake designers often make is that certain architectural features or system characteristics become dependent on these cryptographic schemes. Thus, in the off-chance that they should be found insecure, then the architecture or system is no longer valid. The use of such cryptographic techniques is dubious at best. The security and cryptography communities need more time to assess emerging primitives and protocols before they are adopted in any major way. One avenue is to funnel designs through the CFRG [31], which is often responsible for bridging the gap between academia and industry. Recent schemes under consideration by this research group include password-authenticated key exchange protocols and post-quantum-secure hash-based signature schemes. The latter of which is particularly relevant with respect to content authenticators.

Outlook. There are no compelling reasons to apply esoteric (and often immature) cryptographic techniques in ICN, at least at the network layer. Computationally bounded and traditional cryptographic primitives, such as elliptic curve digital signatures, hash functions, etc., could be the extent of per-packet cryptographic processing done by the routers. Anything more would become fodder for Denial-of-Service attacks that could render the entire infrastructure ineffective. However, architecture designs should not restrict themselves to specific algorithms. There should be flexibility in accommodating multiple (and evolving) cryptographic primitives. This could be useful if, for example, post-quantum digital signature schemes become necessary for the longevity of content authenticators.

3.7 Evolving Network Services

Summary. The Internet has a history of adapting the existing law system to new business paradigms. One such paradigm is in-network processing, which, in recent years, has expanded to aid and impact routing, forwarding, packet replication, packet splitting and merging, quality of control, caching, and others. The relationship between these services and existing laws has been a continual tussle. When and how does caching affect copyright laws? When do other services violate the Secrecy of Correspondence (SoC) statute? Moreover, Deep Packet Inspection on SSL/TLS connections, while technically feasible, may violate the SoC statute and various other privacy rules. Thus, there has been a recent push for all-or-nothing secrecy, which unfortunately stifles network business opportunities.

Outlook. Privacy should be controllable in that it allows secrecy preferences to be expressed by senders in packet headers. This is one area where ICN can innovate to allow in-network processing to continue without violating existing laws. This type of flexibility is not always possible in IP-based protocols such as TCP and TLS, both of which have fixed packet frames and do not easily permit any sort of privacy expression. One notable exception to this claim is the presence of TCP option fields. These may be (mis)used to allow for privacy preferences to be conveyed. For example, tcpcrypt [32] uses these option fields to indicate that a TCP connection will be encrypted. So, while it may be feasible to express privacy preferences in some packet headers, many existing protocols were certainly not designed with this in mind.

4. CONCLUSION

This paper described in detail the discussions and outcomes of the Dagstuhl 16251 seminar on ICN security and privacy. Despite significant research over the past half decade, there are still many open problems with solutions that are difficult to be completely realized with the existing architectures. Are we too invested in the current architectures to make significant design changes to solve these problems? Is there something to be gained by sacrificing properties such as privacy in favor of features such as object security? If so, is this the right tradeoff to make today? Only future research and development will tell.

Acknowledgements

We thank Schloss Dagstuhl for providing a stimulating setting for this seminar. Much progress was made over the course of the seminar and since its completion. This is mainly because of the ease of face-to-face collaboration and interaction at Dagstuhl.

5. REFERENCES

[1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A Survey of Information-Centric Networking (Draft)," in *Information-Centric Networking*, ser. Dagstuhl Seminar Proceedings, B. Ahlgren, H. Karl, D. Kutscher, B. Ohlman, S. Oueslati, and I. Solis, Eds., no. 10492. Dagstuhl, Germany: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2011. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2011/2941

[2] A. Ghodsi, B. Ohlman, J. Ott, I. Solis, and M. Wählisch, "Information-centric networking – Ready for the real worldl (Dagstuhl Seminar 12361)," *Dagstuhl Reports*, vol. 2, no. 9, pp. 1–14, 2013. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2013/3787

[3] D. Kutscher, T. Kwon, and I. Solis, "Information-Centric Networking 3 (Dagstuhl Seminar 14291)," *Dagstuhl Reports*, vol. 4, no. 7, pp. 52–61, 2014. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2014/4785

[4] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure," *Computer Networks*, vol. 57, no. 16, pp. 3192–3206, Nov. 2013.

[5] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in Named Data Networking," in *Proc. of ICCCN*. IEEE, 2013, pp. 1–7.

[6] A. Compagno, M. Conti, P. Gasti, L. V. Mancini, and G. Tsudik, "Violating consumer anonymity: Geo-locating nodes in named data networking," in *International Conference on Applied Cryptography and Network Security*. Springer, 2015, pp. 243–262.

[7] Y. Yu, A. Afanasyev, D. Clark, V. Jacobson, L. Zhang *et al.*, "Schematizing trust in named data networking," in *Proceedings of the 2nd International Conference on Information-Centric Networking*. ACM, 2015, pp. 177–186.

[8] C. Ghali, G. Tsudik, and E. Uzun, "Network-layer trust in named-data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 12–19, 2014.

[9] M. Papalini, "Tagnet: A scalable tag-based information-centric network," Ph.D. dissertation, Università della Svizzera Italiana, 2015.

[10] W. Shang, A. Bannis, T. Liang, Z. Wang, Y. Yu, A. Afanasyev, J. Thompson, J. Burke, B. Zhang, and L. Zhang, "Named data networking of things," in *2016 IEEE First International Conference on Internet-of-Things Design and Implementation (IoTDI)*. IEEE, 2016, pp. 117–128.

[11] R. Tourani, T. Mick, S. Misra, and G. Panwar, "Security, privacy, and access control in information-centric networking: A survey," *arXiv preprint arXiv:1603.03409*, 2016.

[12] D. K. Smetters, P. Golle, and J. D. Thornton, "CCNx access control specifications," PARC, Tech. Rep., Jul. 2010.

[13] S. Misra, R. Tourani, and N. E. Majd, "Secure content delivery in information-centric networks: Design, implementation, and analyses," in *ICN*, 2013.

[14] M. Ion, J. Zhang, and E. M. Schooler, "Toward content-centric privacy in ICN: Attribute-based encryption and routing," in *ICN*, 2013.

[15] C. A. Wood and E. Uzun, "Flexible end-to-end content security in CCN," in *CCNC*, 2014.

[16] J. Kurihara, C. Wood, and E. Uzuin, "An encryption-based access control framework for content-centric networking," *IFIP*, 2015.

[17] Y. Yu, A. Afanasyev, and L. Zhang, "Name-based access control," *Named Data Networking Project, Technical Report NDN-0034*, 2015.

[18] C. Ghali, M. A. Schlosberg, G. Tsudik, and C. A. Wood, "Interest-based access control for content centric networks," in *International Conference on Information-Centric Networking*. ACM, 2015.

[19] M. Raykova, H. Lakhani, H. Kazmi, and A. Gehani, "Decentralized authorization and privacy-enhanced routing for information-centric networks," in *Proceedings of the 31st Annual Computer Security Applications Conference*. ACM, 2015, pp. 31–40.

[20] M. Chase, "Multi-authority attribute based encryption," in *Theory of Cryptography Conference*. Springer, 2007, pp. 515–534.

[21] B. Laurie, G. Sisson, R. Arends, and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence," IETF, RFC 5155, March 2008.

[22] J. Sherry, C. Lan, R. A. Popa, and S. Ratnasamy, "Blindbox: Deep packet inspection over encrypted traffic," in *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4. ACM, 2015, pp. 213–226.

[23] C. Tschudin, "Private information retrieval over icn," in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, April 2016, pp. 534–539.

[24] O. Blazy, G. Fuchsbauer, D. Pointcheval, and D. Vergnaud, "Signatures on randomizable ciphertexts," in *International Workshop on Public Key Cryptography*. Springer, 2011, pp. 403–422.

[25] R. Canetti, S. Halevi, and J. Katz, "A forward-secure public-key encryption scheme," in *International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2003, pp. 255–271.

[26] S. DiBenedetto and C. Papadopoulos, "Mitigating poisoned content with forwarding strategy," in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, April 2015.

[27] E. Uzun, S. DiBenedetto, G. Tsudik, and P. Gasti, "Anonymous named data networking application," in *19th Annual Network and Distributed System Security Symposium (NDSS)*, 2012.

[28] C. Ghali, G. Tsudik, and C. A. Wood, "(The Futility of) Data Privacy in Content-Centric Networks," in *ACM CCS Workshop on Privacy in the Electronic Society (WPES)*, 2016.

[29] C. Wood, E. Uzun, and M. Mosko, "CCNx Key Exchange Protocol Version 1.0," Internet Engineering Task Force, Internet-Draft draft-wood-icnrg-ccnxkeyexchange-01, Oct. 2016, work in Progress. [Online]. Available: https://tools.ietf.org/html/draft-wood-icnrg-ccnxkeyexchange-01

[30] C. A. Wood and E. Uzun, "Flexible end-to-end content security in ccn," in *2014 IEEE 11th Consumer Communications and Networking Conference (CCNC)*. IEEE, 2014, pp. 858–865.

[31] "Crypto Forum Research Group (CFRG)," https://irtf.org/cfrg, accessed: 2016-11-21.

[32] A. Bittau, M. Hamburg, M. Handley, D. Mazieres, and D. Boneh, "The case for ubiquitous transport-level encryption." in *USENIX Security Symposium*, 2010, pp. 403–418.