Policy-Agile Encrypted Networks Via Secure Function Computation

Rajesh Krishnan
Senior Member, IEEE
Front Royal, VA, USA
krash@ieee.org

Ravi Sundaram
College of Computer Science
Northeastern University
Boston, MA, USA
koods@ccs.neu.edu

***Abstract*—Recent developments in cryptography in the areas of secure multi-party function computation and homomorphic encryption enable new policy-agile encrypted networking capabilities. By applying algorithms that can securely perform encrypted operations on encrypted data without decryption of either, context- and content-aware decisions can be performed (or policy rules applied) within a new advanced "black core" network while preserving confidentiality of the mission context and content. We discuss three application areas—with applicability to defense, law enforcement, intelligence community, and commercial networks, especially when resource sharing and collaboration across multiple organizations must occur over out-sourced third party infrastructures—that are enabled.**

Keywords-content-oriented internetworking; secure multiparty computation; black core; policy-agility; content/context-awareness

I. INTRODUCTION

Consider the following scenario: the router within an encrypted network must decide which of two encrypted IP-based video streams should be prioritized for sending over a bandwidth-constrained satellite link. The two video streams have the same quality-of-service markings, yet one carries mission-critical information, a distinction that requires context/content awareness that is typically not available.

One approach would be to add more fine-grained priorities. This assumes quasi-static scenarios and tight coordination of the ingress points and the routers within the encrypted networks. Moreover, this approach addresses only the symptom, and not the core issue: *How can we provide policy agility within an encrypted "black core" network?* On the one hand, encrypted networks allow management and operation of infrastructure to be delegated while preserving "need to know" considerations. On the other hand, dynamic policy-agile operation responsive to commander's intent requires context and content awareness, unavailable in the "black core" network. Resolving this tension between two conflicting requirements in current operational networks requires manual intervention that is slow and expensive.

What if we could perform encrypted operations on encrypted data without decrypting either? For example, encrypted data could carry additional encrypted metadata, and network management can distribute encrypted policy rules, and network nodes can apply the encrypted policy rule on encrypted metadata without decryption. This will indeed enable a new advanced policy-agile encrypted networking capability. In this paper, we explore this possibility and its potential applications.

Secure multi-party computation and homomorphic encryption are areas of cryptography that have traditionally been ignored by the field of communications and networking. These areas however are rapidly seeing numerous applications in private information retrieval, private searching, secure voting, and other areas.

In cryptography, **secure multi-party computation** is a problem that was suggested by Andrew C. Yao in [2]. In that publication, Yao introduces and proposed a solution to the millionaire problem: Alice and Bob are two millionaires who want to find out who is richer without revealing the precise amount of their wealth. This problem and result gave way to a generalization called multi-party computation (MPC) protocols. In an MPC, a given number of participants (p1, p2 ... pN) each with private data (respectively d1, d2 ... dN). The participants want to compute the value of a public function F on N variables at the point (d1, d2 ... dN). An MPC protocol is dubbed secure if no participant can learn more from the description of the public function and the result of the global calculation than what he/she can learn from his/her own entry — under particular conditions depending on the model used. Variations such as those proposed by Feige et al. [4] and [8] allow a third party to perform the computation. Secure MPC provides solutions to various real-life problems such as distributed voting, private bidding and auctions, sharing of signature or decryption functions, private information retrieval, etc. In our proposed work, we plan to apply secure multi-party communication to allow an intermediate node to make content-oriented caching, forwarding, and retrieval decisions.

Homomorphic encryption is a form of encryption where one can perform a specific algebraic operation on the plaintext by performing a (possibly different) algebraic operation on the ciphertext. For example, it is possible for a third party computer to compute encrypted circuits on encrypted inputs without gaining any information about the inputs or the circuits being computed. Rivest, Adleman, and Dertouzos [1] posed the problem of whether any arbitrary computation (consisting of both addition and multiplication operations in a field) can be performed using homomorphic encryption. The existence of a fully homomorphic cryptosystem would have great practical implications in the outsourcing of private computations, for instance, in the context of cloud computing. A solution proved

more elusive; for more than 30 years, it was unclear whether fully homomorphic encryption was even possible. In 2009, the first fully homomorphic cryptosystem was constructed by Craig Gentry [10] which removes theoretical barriers to fully homomorphic encryption. In our work, we plan to apply these concepts to allow intermediate nodes to aggregate encrypted interest predicates for scalable and secure content networking; we do not however require arbitrary computations and will identify a simpler metadata language for which efficient homomorphic schemes (such as the Paillier cryptosystem [27]) are known.

Yao's pioneering work on secure multiparty computation [2] is also a good introduction to the area. The related area of zero knowledge proofs and verifiable secret sharing is presented in an accessible way in [12]. The homomorphic encryption problem was proposed in Rivest et al. [1] and recently settled by Gentry [10]. A good survey of homomorphic encryption for the non-specialist can be found in [13]. A language and toolkit to support experimentation with secure multi-party computations is described in [7]. Sander, Young and Yung present a non-interactive protocol for secure crypto-computing for circuits in the NC1 complexity class (polynomial width and logarithmic depth in size of input) [5]. Rafail and Ostrovksy consider round-optimal multi-party computation in [9]. Protocols with low communication overhead that enable outsourcing the computation to a third-party can be found in [4] [8]. Rafail and Ostrovsky present an approach for private searching on streams [11].

In this paper, we look at three application areas with increasing level of complexity:

- Secure Content-Oriented Internetworking, which enables a confidential publish-subscribe capability
- Secure Datalog, which enables: (i) secure declrative networking, in particular, for confidential processing of encrypted policy, routing, and firewall rules within encrypted networks, and (ii) secure databases to support outsourcing of applications (e.g., Customer Relationship Management) while preserving confidentiality
- Crypto-Computing Arbitrary Functions in the Cloud

Secure content-oriented internetworks are a natural generalization of the cache-and-forward architecture inherent in delay-tolerant networks. Using our approach, users can specify their interests or publish content and expect infrastructure to securely match the supply and demand without loss of confidentiality. This application area addresses how nodes within a data network can make content-oriented forwarding, caching, and retrieval decisions based on encrypted metadata and encrypted interests (publish/subscribe advertisements) without decrypting them. Existing work [3] suggests that known algorithms can be applied to efficiently implement secure computation of a range of operations that can lead to a useful and scalable secure content-oriented network.

As further extensions of our work, we consider what additional new capabilities we can endow upon encrypted (black core) networks by applying secure function computation. Can we make encrypted networks more responsive and agile to mission and policy requirements without compromising confidentiality? Can we do a new style of transport-layer performance enhancing proxies? Can we do content or context-based prioritization and traffic shaping? How can we do encrypted firewall rules that can operate on encrypted packets without learning the packet headers or the firewall rules? An open question in this area whether there are practical and efficient algorithms to implement Datalog as a secure computation. We provide an approach that partially answers this question by applying known results in the area.

The third area of crypto-computing in the clouds is an active research area. Although Gentry's result has removed theoretical barriers to full homormorphic encryption, we expect it will be some time before efficient algorithms are devised. Furthermore, there will be a trade-off between time complexity and expressiveness. We expect networking systems will favor fast wire-speed operation over expressiveness, and given the rich possibilities of the other two areas, we do not explore this general case further in this paper.

II. SECURE CONTENT-ORIENTED INTERNETWORKING

A. Content-Oriented Internetworking

In a content-oriented internetwork, users describe what they want, not where it is stored, and the network moves information when and where needed. Users should be able to specify the data they need without having to know or search for where it is (or will be) located. For military applications, we can view content-oriented networking as a natural evolution of cache-and-forward networking approaches such as Disruption Tolerant Networking [15][16]. The infrastructure must efficiently support users and sub-networks that may be frequently disconnected.

Content-based networking is a paradigm for a new communication infrastructure proposed earlier by Carzaniga and Wolf [6] and others. The idea here is that the flow of information from senders to receivers is driven by the content of the messages and the interests of the receivers rather than by explicit addressing mechanisms or heavyweight publish-subscribe schemes. Carzaniga and Wolf propose a specific instantiation of the content-networking paradigm and an efficient algorithm for realizing the same. The basic idea of [6] was to suggest that the predicates be specified in Disjunctive Normal Form over a space of attribute-value pairs constituting the metadata. They provide an efficient algorithm for converting a collection of such predicates into a fast forwarding table for moving incoming datagrams to the appropriate outgoing ports.

Van Jacobson, in an influential talk in 2007 at Google, Inc. on "A New Way to Look at Networking," presented his vision for a radical content-centric networking paradigm and articulated the imperative to replace the current conversation-centric paradigm [18]. Van's talk touches upon some security issues for such a network, for example, secure and authentic binding of content to the tags.

We use content-oriented internetworking as a general term covering content-centric networking proposed by Van

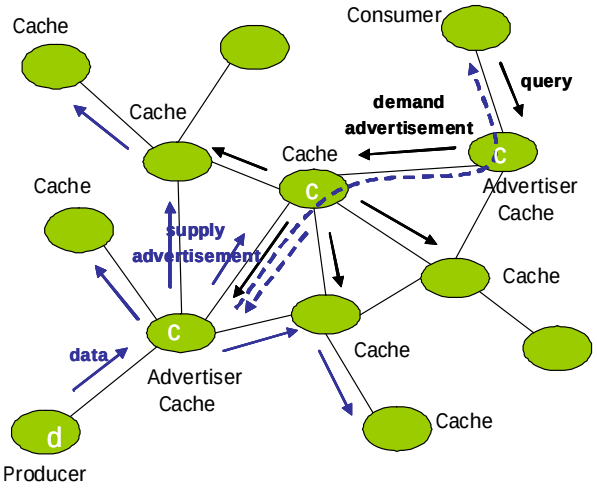

Figure 1: Content-Oriented Internetwork

Jacobson, content-based networking proposed by Carzaniga and Wolf, and other approaches such as CCSDS AMS [14].

A general content-oriented network is shown in Figure 1. Content producers produce content and content advertisers generate metadata (e.g., attribute-value pairs describing content) and disseminate supply advertisements, and optionally the content to selected nodes. Consumers broadcast their demand (interests) and intermediate nodes must match the demand with the supply and push content. Consumers may also generate queries for stored content. Content is forwarded along ports over which matching interests were heard; nodes en route may optionally cache the content to satisfy future queries. Thus both a push and a pull model are supported.

The demand and supply advertisements respectively describe who provides what and who requires what. The advertisements — who, what, and additional constraints on delivery — are described using extensible ontologies. The advertisement data bases are synchronized using Epidemic or other disruption-tolerant protocols. Rules specify policies for caching, forwarding, purging, and other content management operations.

B. Secure Content-Oriented Internetworking Problem

A major problem left open by several earlier content-oriented internetworking works is that of confidentiality. It is particularly relevant in the context of defense and sensitive information networks.

Communities of users must use internetworking infrastructure operated by untrusted parties in order to provide and consume information. For example, in multi-national coalition mission scenarios, data may need to be cached, carried, and forwarded by nodes belonging to partner nations. Therefore, it becomes necessary that content confidentiality is protected for the lifetime of the information as it is cached and forwarded by third party infrastructure, and not just during transit. A new model called temporal security (illustrated in Figure 2) was originally proposed by Preston Marshall to capture this requirement [26]. An implication of this architecture is that end host systems need operating system and hardware support to ensure applications keep clear-text in ephemeral memory, with safeguards that prevent persistent storage of clear-text.

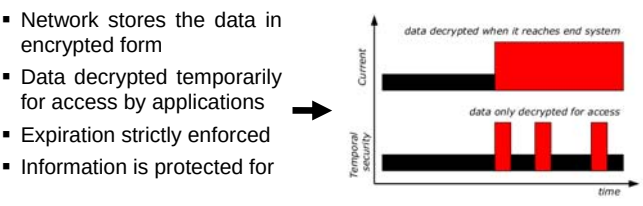

Figure 2: A new security model focused on content-oriented networking [26]

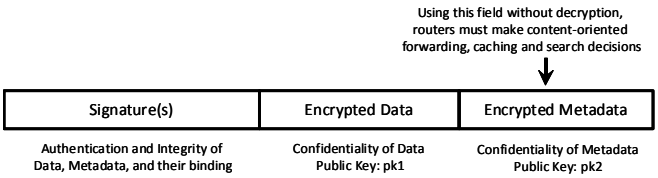

Figure 3: Abstract protocol data units for secure content-oriented networking

Developing the secure content networking concept further, in this paper we require that intermediate nodes must not only blindly disseminate encrypted content, but actually be able to make decisions about how to store or move the data in a content-oriented manner.

Yet in our assertion there seems to be an apparent conflict between the desire for confidentiality of content and the desire to make decisions based on the content. Fortunately, this conflict can be resolved by drawing upon areas of cryptography —secure multi-party computation and homomorphic encryption, in particular—that have traditionally been ignored by the field of communications and networking. These areas however are rapidly seeing numerous applications in private information retrieval, private searching, secure voting, and other areas.

We summarize the secure content-oriented internetworking problem as follows: ***Can nodes within the network make content-oriented forwarding, caching, and retrieval decisions based on encrypted metadata and encrypted interests (publish/subscribe advertisements) without decrypting them?***

C. Proposed Approach

In our approach, we expect that content and metadata are protected separately. Content is decrypted by the end-application as needed and is always stored encrypted. Metadata is never decrypted, but intermediate routers operated upon metadata (using homomorphic encryption) to make content-oriented caching and forwarding decisions. In addition cryptographic signatures are included to provide guarantee the integrity and authenticity of the data, the metadata, and the binding between the two. An abstract protocol data unit (PDU) for content that illustrates the separation of security concerns is shown in Figure 3.

Central to our approach is the problem of Public Evaluation of Private Predicates on Encrypted Data (PEPPED). Consider a content network in which senders are broadcasting data on different topics. Receivers have interests and wish to get the data corresponding to their interests. A receiver specifies its interests in terms of a predicate. A straightforward and non-secure implementation is for each receiver to communicate its predicate to the upstream router who then applies the predicate as a filter to decide the appropriate datagrams to forward on a given port. Now, for obvious privacy reasons the data itself needs to be encrypted so that other entities cannot tell what content is being forwarded to a receiver. In the Carzaniga-Wolf model [6] there is metadata in the form of attribute-value pairs to which the predicate is then applied. Now, straightaway this leads to a privacy problem because the metadata is in the clear and so is the predicate. This will reveal all kinds of things about the receiver that they may not wish to have revealed. A natural requirement is that the metadata be encrypted as well and that the receiver communicates an obfuscated version of their predicate such that the router is able to evaluate the predicate on the encrypted metadata but is unable to tell much either about the predicate or the metadata (and associated data). This will allow the router to determine which ports to forward a datagram on.

A trivial solution to the privacy problem would be to forward all the encrypted datagrams to all receivers and have each receiver decrypt the data of interest to them. But this violates the raison d' etre of content-based networking which is to optimize the movement of information by forwarding only relevant datagrams. Another advantage of forwarding relevant information is that receivers explicitly do not receive information about interests not specified by their predicates whereas the trivial solution of forwarding all data to everybody all the time does not provide this security guarantee. To summarize, the basic problem is to make a cryptographic scheme that enables senders to encrypt data and receivers to create obfuscated predicates that routers can then evaluate on the encrypted data to enable them to make their forwarding decisions. One can also consider additional natural requirements such as the need of a router to propagate various disjunctions of the predicates from its downstream ports to its upstream ports while still maintaining anonymity of its downstream ports which can be generalized as a problem of secure delegation.

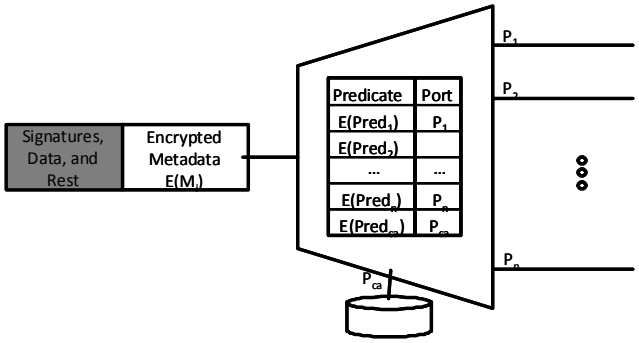

Figure 4: Forwarding and caching in a secure content-oriented network router

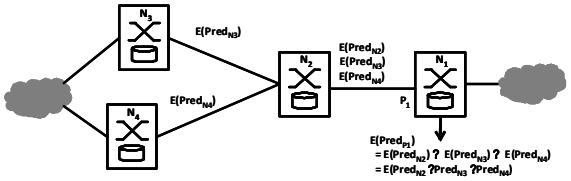

Figure 5: Aggregation of interest predicates using homomorphic encryption

There is a extant body of literature in this active research area we can draw upon to implement PEPPED in this context. One approach is provided by Raiciu and Rosenblum [3], who address confidentiality issues for content-based publish-subscribe systems in well-connected networks. The main contribution of [3] is to show that confidentiality can be preserved for a certain limited form of Content Based Publish Subscribe (CBPS). They assume a model in which the publisher and subscriber share symmetric keys and the content broker is honest but curious. They repurpose a scheme from [20] for equality match and schemes from [19] and [21] for keyword match ([19] is a more constrained scheme that requires documents to have the same number of keywords). Their schemes maintain the confidentiality of both the subscriber and the publisher. They show how the basic schemes for equality and keyword match can be used as primitives for performing numeric matches as well as arbitrary predicates. Unfortunately, their schemes do not scale and require exponential time in the worst case for arbitrary predicates. Nevertheless, they made an important contribution by showing how existing schemes enable confidentiality for two simple and useful predicate forms: equality and keyword match.

A secure content router receiving a content PDU must decide whether to forward the content, and if so, on which ports. For this purpose, the forwarder maintains a table of encrypted interest predicates associated with the ports as shown in Figure 4. Using PEPPED described earlier, the forwarder determines if the encrypted metadata on the PDU matches the encrypted predicate for a port; if there is a match, the PDU is forwarded along the corresponding port. The construction of this table of predicates is described later.

The content router must also determine whether to cache the content. By mapping the cache to a designated port, this decision reduces to a similar computation as the forwarding decision. The cache predicate may be updated by a process different from the routing process that updates the forwarding table entries.

Following the work by Carzaniga and Wolf, we present an approach to compute the forwarding table, but securely in our case. The forwarder computes a disjunction of interest predicates received from neighbors on a port are assigns the aggregate predicate to that port. This computation is done by applying homomorphic encryption as illustrated in Figure 5. Other sophisticated routing approaches will be needed in mobile and disruption-prone networks.

Another related issue is key management. For the purposes of this paper, we assume a public key infrastructure and secure group key management approach such as GSAKMP [17] is in

place. Producers and consumers must be able to negotiate group keys for encrypting data which does not involve the routers. Public keys of the producers/advertisers that generate metadata and public keys of the consumers that generate interests, however, must be available to the routers. We envision multiple communities of interest with separate group keys, and built-in expiration of encrypted content which must be refreshed periodically.

In summary, a key significance of the proposed approach is that it enables third party infrastructure to serve as cache and forward relay nodes that make content-oriented decisions on caching and forwarding (so that resources are appropriately utilized) without loss of confidentiality. In other words, using the proposed approach, an intermediate node can match content to user subscriptions without learning anything about the content.

III. Feasibility Of Secure Datalog

Datalog is a general purpose language that is used to query databases (see [25]). It can also be used in CBPS schemes – the interests of the subscriber are expressed in Datalog which is then matched against the "database" generated by the publisher. In the typical case the subscriber's interests are consider to be longer-lived whereas the publisher is presumed to generate output at higher frequency. Thus the natural measure of complexity is the data complexity of Datalog where the data (put out by the publisher) is considered the input while the query (put out by the subscriber) is considered fixed. (This is in contrast to the combined complexity measure where both query and data are considered part of the input.) Datalog is known to be P-complete under the data complexity measure.

We now turn to a formalization of CBPS based on the work of [4]. The publisher generates an encrypted version x of his output while the subscriber generates an encrypted version y of his predicate. The content broker then computes f(x,y) and based on the output decides whether to forward x to the subscriber or not. Here f is assumed to be known to the content broker. The question is what is the class of predicates y that can be handled effectively, i.e., in polynomial-time. In [4] the authors show that f can be computed securely for functions in the class NC1 using a randomized blinding scheme based on Barrington's theorem [22]. Now, we need that $f(x,y) = Cx(y)$ where Cx is the predicate corresponding to the encrypted encoding x. So the question now becomes: what is the class of circuits for which f in NC1 is a universal circuit? The concept of universal circuit was defined in [23] (see also [24]) and it is shown in [5] that the class of circuits for which an NC1 circuit is universal is strictly contained in NC1. Observe that since Datalog is P-complete, then under the well accepted notion that P <> NC we cannot hope to handle the subscriber's interests expressed in general Datalog. Nevertheless, we are able to show that we can handle a very rich class of predicates. Developing on ideas from [23] we are able to show that there exist NC1 universal circuits for the class of circuits that are log-depth and bounded symmetric function aggregates of NC0. This class of circuits not only includes equality and keyword match but also intersection of sets of keywords, and even more complex functions such as the parity of set intersection or the tribes function.

Crypto-computing Datalog opens up a variety of possibilities. It can support a secure database in which both data and operations are confidential. This allows outsourcing of applications (such as Customer Relationship Management) by DoD vendors without loss of confidentiality and risk of opening up to Open Source Intelligence collection.

Secure Datalog can also enable secure declarative networking capability, in particular, for confidential processing of encrypted policy, routing, and firewall rules within encrypted networks. Secure Datalog holds the key to making encrypted networks more responsive and agile to mission and policy requirements without compromising confidentiality. Possible applications to explore include a new style of transport-layer performance enhancing proxies to allow dynamic splicing, content or context-based prioritization and traffic shaping, and encrypted firewall rules that can operate on encrypted packets without learning the packet headers or the firewall rules.

For now, however, which interesting and useful Datalog subsets can be crypto-computed remains open.

IV. Summary And Future Work

In this paper we propose exploring the use of secure function computation and homomorphic encryption to support novel policy-agile encrypted networking capabilities. We propose a way to resolve the conflicting needs between security and context-aware policy-agility ineherent in current encrypted networking architectures. Two key technologies with wide-ranging application possibilities are secure content-oriented networking and secure Datalog subsets.

We propose radical innovations that represent departures from current approaches: (i) we move from the dominant end-to-end conversation paradigm to a content-oriented paradigm that is a natural generalization of delay-tolerant networking, and (ii) we move to a security model that focuses on protecting the content rather than the channel. Current approaches do not allow intermediate nodes to securely make caching, routing, and retrieval decisions in a content-oriented fashion — current approaches require the content/metadata and the interest predicates to be in clear-text in order to make content-oriented decisions. Our approach removes this restriction and allows third party infrastructure to perform these decisions without compromising confidentiality of the content or the interest predicates.

Even though cryptocomputing Datalog is not possible (because Datalog is P-complete whereas even Gentry's scheme is NC1), nevertheless, we can compute a very rich subclass of NC1, and future work on whether these correspond to useful restricted versions of Datalog will be interesting.

Our survey of the literature indicates a body of literature that suggest both of these key technologies are feasible. Further work in algorithm and systems development and vulnerability analyses are necessary before the envisioned capabilities can be achieved.

Acknowledgment

Rajesh Krishnan thanks Dr. Preston Marshall for leading him to some of the key insights on rethinking security, especially for content-oriented networking.

REFERENCES

[1] R. L. Rivest, L. Adleman, and M. L. Dertouzos, "On data banks and privacy homomorphisms," In Foundations of Secure Computation, 1978.

[2] A.C. Yao, "How to Generate and Exchange Secrets," Proc. of Twenty-seventh IEEE Symposium on Foundations of Computer Science, Toronto, Canada, October 1986, 162–167.

[3] Costin Raiciu and David S. Rosenblum, "Enabling Confidentiality in Content-Based Publish/Subscribe Infrastructures," Second International Conference on Security and Privacy in Communication Networks (SecureComm 2006), Baltimore, MD, USA, Aug 28–Sep 1, 2006

[4] Uri Feige, Joe Kilian, and Moni Naor, "A Minimal Model for Secure Computation," Proceedings of the Twenty-Sixth Annual ACM Symposium on theory of Computing (STOC '94), Montreal, Quebec, Canada, May 23–25, 1994

[5] Tomas Sander, Adam Young, and Moti Yung, "Non-Interactive CryptoComputing For NC1," Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS '99), Washington, DC, USA, Oct 17–18, 1999

[6] Antonio Carzaniga and Alexander L. Wolf, "Forwarding in a Content-Based Network," Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols For Computer Communications (SIGCOMM '03), Karlsruhe, Germany, Aug 25–29, 2003

[7] Assaf Ben-David, Noam Nisan, and Benny Pinkas, "FairplayMP: A System for Secure Multi-Party Computation," Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS '08), Alexandria, VA, USA, Oct 27–31, 2008

[8] Jaideep Vaidya and Chris Clifton, "Leveraging the Multi in Secure Multi-Party Computation," Proceedings of the 2003 ACM Workshop on Privacy in the Electronic Society (WPES '03), Washington, DC, USA, Oct 30, 2003

[9] Jonathan Katz and Rafail Ostrovsky, "Round-Optimal Secure Two-Party Computation," Proceedings of CRYPTO 2004: Advances in Cryptology, Santa Barbara CA, Aug 15–19, 2004

[10] Craig Gentry, "Fully Homomorphic Encryption using Ideal Lattices," Proceedings of the 41st Annual ACM Symposium on theory of Computing (STOC '09), Bethesda, MD, USA, May 31–Jun 02, 2009

[11] Rafail Ostrovsky, William E. Skeith, III, "Private Searching on Streaming Data," Journal of Cryptology , Volume 20 Issue 4, Oct 2007

[12] Jean-Jacques Quisquater, Myriam Quisquater, Muriel Quisquater, Michael Quisquater, Louis Guillou, Marie Annick Guillou, Gaid Guillou, Anna Guillou, Gwenole Guillou, Soazig Guillou, and Thomas A. Berson, "How to Explain Zero-Knowledge Protocols to your Children," Proceedings on Advances in Cryptology Santa Barbara, CA, USA, Aug 20–24, 1989

[13] Caroline Fontaine and Fabien Galand, "A Survey of Homomorphic Encryption for Nonspecialists," EURASIP Journal on Information Security, Jan 2007, Pages 1–15

[14] Scott Burleigh, "Overview of AMS (CCSDS Asynchronous Message Service)," DARPA DTN Phase 2 Kickoff, Arlington, VA, USA, Aug 9, 2006

[15] Vint Cerf, Scott Burleigh, Adrian Hooke, Leigh Torgerson, Robert Durst, Keith Scott, Kevin Fall, and Howard Weiss, "Delay-Tolerant Networking Architecture," RFC 4838, Apr 2007

[16] Keith Scott and Scott Burleigh, "Bundle Protocol Specification," RFC 5050, Nov 2007

[17] H. Harney, U. Meth, A. Colegrove, and G. Gross, "GSAKMP: Group Secure Association Key Management Protocol," June 2006.

[18] Van Jacobson, "A New Way to Look at Networking," Google Tech Talks, August 30, 2006, http://video.google.com/videoplay?docid=-6972678839686672840#

[19] Eu-Jin Goh, "Secure indexes", Cryptology ePrint Archive, Report 2003/216, 2003.

[20] Practical Techniques for Searches on Encrypted Data, with Dawn Song and David Wagner. In Proc. of IEEE Security and Privacy Symposium, May 2000

[21] Yan-Cheng Chang and Michael Mitzenmacher, "Privacy preserving keyword searches on remote encrypted data." in ACNS, 2005.

[22] David A. Mix Barrington : Bounded-Width Polynomial-Size Branching Programs Recognize Exactly Those Languages in NC^1, STOC, 1986

[23] Stephen A. Cook, H. James Hoover: A Depth-Universal Circuit. SIAM J. Comput. 14(4): 833-839 (1985)

[24] Leslie G. Valiant: Universal Circuits (Preliminary Report) STOC 1976: 196-203

[25] S. Ceri, G. Gottlob, L. Tanca, "What you always wanted to know about Datalog (and never dared to ask)," IEEE Transactions on Knowledge and Data Engineering (1989) Volume: 1, Issue: 1, Pages: 146-166

[26] Preston Marshall, "From Self-Forming Mobile Networks to Self-Forming Mobile Content Services," Keynote speech, ACM MobiCom 2008, San Francisco, CA, USA, September 16, 2008, http://www.sigmobile.org/mobicom/2008/MobiCom08-Marshall-Keynote.pdf

[27] Pascal Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," EUROCRYPT 1999, pp. 223–238.