Management of Content-Centric Networking

Thibault Cholez

Management of Content-Centric Networking

Thibault CHOLEZ

RESCOM 2013

17/05/2013

Plan

Plan

CCN an old idea ?

A MOBICOM 2000 paper [IGE00] [AY05] "Directed Diffusion" describes the "Data-Centric" paradigm

- "Directed Diffusion is an important milestone in the data-centric routing research of sensor networks. The idea aims at diffusing data through sensor nodes by using a naming scheme for the data"
- "The main reason behind using such a scheme is to get rid of unnecessary operations of network layer routing"
- "In order to create a query, an interest is defined", "The interest is broadcast"
- "Each node receiving the data can do caching for later use"
- "Hence, by utilizing interest and gradients, paths are established between sink and sources. Several paths can be established"
- "all communication is neighbor-to-neighbor with no need for a node addressing mechanism"

Why focusing on CCN ?

Why is CCN so popular among ICN solutions ?

- Simple architecture based on simple ideas
- ACM CoNEXT 2009 paper [JST+09] : Good educational introduction, most architectural aspects covered
- Many research questions clearly highlighted (routing, key management, etc.)
- ... even if not all (scalability regarding number of contents or updates frequencies, enforcement of unique names at the Internet scale, mobility while providing content, etc.)
- CCNx implementation / community
- Lucky factor : right time / right research community ?

Network management

What is network management ?

- Wikipedia attempt : "Network management refers to the activities, methods, procedures, and tools that pertain to the operation, administration, maintenance, and provisioning of networked systems"
- NMRG functional areas : Fault management, Performance management, Security management, Configuration management, Accounting management, Service management, Event management
- Configure and control a set of resources that ensure the network is running well
- Means : monitoring (centralized, distributed, autonomous) coupled with control/optimization/economic/machine learning/stochastic theories

Network management

What is relevant for CCN ?

- Traffic management
- Cache management
- Content management
- Nodes management
- Security management

Traffic Management

Key challenges

- Traffic control for QoS [ORS12] [FRRS12]
  - Fair queuing, congestion avoidance (Interest discard, etc.)
- Traffic control for business/political purposes (easier filtering / censorship ?)
- Traffic differentiation : several hacks of the CCN architecture to handle specific traffic types
  - Private data [ACG+13] [a]
  - Ephemeral data [CPW11] [b]

---

a. "Cache Privacy in Named-Data Networking" Acs & al.

b. "Content-based publish/subscribe networking and information-centric networking" Carzaniga & al.

Cache Management

Key challenges

- What types of contents should be cached ? Where ? With which priority ? [FRRS12]
  - Video on demand
  - Long lived multimedia contents (file sharing or user generated)
  - Web
- What types of contents should not be cached ?
  - Conversational (two parties)
  - Ephemeral events (notifications from IoT world, online games)
  - Private communications (email, VOIP, etc.)
- How to use cache efficiently : size, location, caching policy (LRU, LFU, popularity, etc.)

Content Management

Key challenges

- Accountability of content's access
- Monitoring of content (diffusion, replication)
- Revocation of deprecated contents
- Access control (restriction per country, per user, etc.)

Nodes Management

Key challenges

- Monitoring of CCN nodes : collect information and status of CCN nodes, detection of anomalies
  - Strategy layer needs information
  - Define relevant information to be monitored + right granularity (per prefix, per face) ?
  - Define efficient architecture (CCN ready) for collection of information
- Design and implementation of new network tools (ping, traceroute, etc.)
- Design of new metrics (CCN/ICN flows, etc.)
- Remote configuration (no push mechanism)

Security Management

Key challenges

- Identification of new threats [WSV12] [a]
- DoS by resource exhaustion of stateful routers
  - On PIT : Interest flooding attack
  - On FIB : Announcing conflicting domains, non-existing content, non-aggregable names
  - On CS : Privacy issues (cache probing) [LLR+12] [b] [ACG+13], cache pollution
  - Cryptography attacks (long-lived content, many encryptions with the same key)

---

a. "Threats to Stability and Security in Information-Centric Networking" Wahlisch & al.

b. "Privacy risks in named data networking : what is the cost of performance ?" Lauinger & al.

Security Management

Key challenges

- Security of contents (revocation, pollution [FMP10] [a], etc.)
- Security of the naming space (malicious names, route poisoning ∼BGP)
- Key management scheme : How to retrieve public keys quickly, efficiently, securely ?
- Collaboration for attack detection
- Enforcement of security policies within a network

---

a. "Information ranking in content-centric networks" Fotiou & al.

What is available in CCNx commands ?

Content management

- ccngetfile : retrieve a file published as CCNx content and save it to a local file
- ccnputfile : publish a file as CCNx content (local file filename or url to content with the ccnxname), manage segmentation, key signature, etc.
- ccnrm : mark as stale any locally cached content objects matching the given prefix (no further Interest response)
- ccnls : list name components available at the next level for a given CCNx name prefix
- ccnlsrepo : explore content stored under a given prefix (one or more repositories)

CCNx commands

Tools

- ccn_ccnbtoxml : convert ccn binary encoded data into XML form
- ccn_xmltoccnb : convert XML into ccn binary encoded data (ccnb)
- ccndsmoketest : testing of communications, send and receive data on sockets

Monitoring of ccnd

Monitoring commands

- ccndstatus : display the status a running ccnd (nb of active faces, stat of each face, etc)
- ccnpeek : generates an Interest, get one content item matching the name prefix and write it to stdoud (eq to IP ping)
- ccnponk : read data from stdin, send it as a single ContentObject in response to an interest

Web Interface

- http ://localhost :9695/ , similar to ccndstatus
- Limitation : sliding time window (avg of last minute) provides inaccurate results

Monitoring web interface

mailly ccnd[14550] local port 9695 api 6000 start 1338724361.760754 now 1338726149.399268

Content items: 23 accessioned, 23 stored, 11 stale, 0 sparse, 112 duplicate, 143 sent
Interests: 31 names, 2 pending, 2 propagating, 14 noted
Interest totals: 586 accepted, 447 dropped, 588 sent, 112 stuffed

Faces

- **face:** 0 **flags:** 0xc **pending:** 0
- **face:** 1 **flags:** 0x400c **pending:** 0
- **face:** 2 **flags:** 0x5012 **pending:** 0 **local:** 0.0.0.0:9695
- **face:** 3 **flags:** 0x5010 **pending:** 0 **local:** 0.0.0.0:9695
- **face:** 4 **flags:** 0x4042 **pending:** 0 **local:** [::]:9695
- **face:** 5 **flags:** 0x4040 **pending:** 0 **local:** [::]:9695
- **face:** 7 **flags:** 0x81412 **pending:** 0 **remote:** 127.0.1.1:9695 **via:** 2
- **face:** 12 **flags:** 0x1014 **pending:** 2 **activity:** 13 **remote:** 127.0.0.1:38200
- **face:** 13 **flags:** 0x1014 **pending:** 0 **activity:** 7 **remote:** 127.0.0.1:38202
- **face:** 14 **flags:** 0x21012 **pending:** 0 **activity:** 7 **remote:** 127.0.0.1:9695 **via:** 2

Face Activity Rates

	Bytes/sec In/Out	recv data/intr sent	sent data/intr recv
face: 0	259 / 31	0 / 0	0 / 0
face: 7	0 / 176	0 / 0	0 / 0
face: 12	128 / 0	0 / 0	0 / 0
face: 13	0 / 0	0 / 0	0 / 0
face: 14	440 / 263	0 / 0	0 / 0

Forwarding

- ccnx:/%C1.M.S.localhost/%C1.M.SRV/ccnd **face:** 0 **flags:** 0x3 **expires:** 2147481862
- ccnx:/ccnx/ping **face:** 0 **flags:** 0x3 **expires:** 2147481862
- ccnx:/ccntuto2 **face:** 7 **flags:** 0x3 **expires:** 2147482037
- ccnx:/%C1.M.S.neighborhood **face:** 0 **flags:** 0x3 **expires:** 2147481862
- ccnx:/%C1.M.S.localhost **face:** 0 **flags:** 0x23 **expires:** 2147481862
- ccnx:/ccnx/%1B%D20%5C%AD%86%99Z%1F%BE%94%09%06%FAy%12%F6%19%E4%8E%B6%F6oI%8B%17%A4%E5%A3.%05%DB **face:** 0 **flags:** 0x17 **expires:** 2147481862
- ccnx:/ccntuto **face:** 7 **flags:** 0x3 **expires:** 2147481922
- ccnx:/ccntuto2/test_chat_room/Users/tibs/Keys/%C1.M.K%00%9D%BA%9Cv%AC%DC%BE%DA%CE%80%21HAYG%1A%D1izN%3A_%2F8s%7F%FC%D1%E9%13cR/%FD%04%EAXb 12 **flags:** 0x3 **expires:** 2147483637

Experimentation for management activities

Need of better tools

- More monitored parameters
- Better accuracy
- New metrics

Need of a large scale testbed to support experiments

- Based on CCNx enabled nodes
- Solution 1 : federated testbed between academic partners (like the young Internet)
- Solution 2 : PlanetLab nodes (cf NEPI talk)

IRTF Information-Centric Networking Research Group

Main topics

- Naming schemes for ICN and scalable name resolution for flat names
- Scalable routing schemes
- Congestion control, QoS approaches, and caching strategies
- Metrics that make it possible to evaluate ICN implementations in a consistent manner
- Security and privacy, including scoping of information objects and access control to them
- Application/application-protocol design and APIs
- Business, legal and regulatory frameworks
- Deployment and interoperability (with BGP, OSPF)

IRTF Information-Centric Networking Research Group

- Very active group, mailing list : icnrg@irtf.org, web : irtf.org/icnrg
- Other related IRTF working groups : RRG (Routing Research Group), NMRG (Network Management Working Group)

Plan

Motivation

How to enforce security policies in CCN ?

- Goal : prevent users from downloading malicious/forbidden contents
- Authentication of content possible (lower layers : simple verification) but real security tools missing
- Inheritance of IP firewalls limited : no filter on IP addresses or ports
- New security features enabled by the CCN paradigm

Contribution

- Content firewall : considering content name and signature
- Use case analysis : Identification of security needs for CCN
- Design of a semantic CCN firewall : grammar definition, preprocessing for semantic enhancement
- Implementation in CCNx and performance evaluation

IP firewall main use cases

- IP_UC1 : Filtering based on the protocol (Example : http, smtp, etc.)
- IP_UC2 : Filtering based on status of the connection (new, established, etc.)
- IP_UC3 : Filtering based on a list of known blacklisted IP addresses
- IP_UC4 : Filtering unusual inbound traffic pattern (from a denial of service attack attempt)

Some use cases do not make sense in CCN, others must be adapted.

CCN-specific use cases

- CCN_UC1 : Filtering on content provider (Example : known untrustworthy or banned)
- CCN_UC2 : Filtering on bad signature
- CCN_UC3 : Filtering on content name and semantic (Example : excluding contents named with a given keyword)
- CCN_UC4 : Composition (content provider & content name)
- CCN_UC5 : Filtering on content direction (Example : avoid leakage of certain documents)
- CCN_UC6 : Filtering on heavy traffic (Preservation of QoS)
- CCN_UC7 : Filtering of stored data (Example : caching only storing specific content)

Comparison : IP vs CCN use cases

IP use cases	CCN use cases	Filtering on
IP_UC1	CCN_UC3	Protocol / Content name
IP_UC2	--	Status of the connection
IP_UC3	CCN_UC1	Listed IP / Content provider
IP_UC4	CCN_UC6	Unusual / Heavy traffic
--	CCN_UC2	Bad signature
--	CCN_UC4	Composition of filters
--	CCN_UC5	Content direction
--	CCN_UC7	Stored data

Syntax definition

- Syntax based on iptables for ease of use and readability
- 3 different types of rules

```
rule = r_interest | r_data | r_face

r_interest = "interest" SP direction SP
             match_interest SP "pit" SP action

r_data = "data" SP direction SP match_data
         SP ["cs"|"pit"] SP action

r_face = "face" SP number
```

r_interest

Main rule

```
r_interest = "interest" SP direction SP
             match_interest SP "pit" SP action
```

Syntactic elements

```
direction =  "*"|"int"|"ext"
action = "forward"|"drop"
match_interest = content_name
```

Example

```
interest * \@game|play|fun\@ 15 pit drop
```

r_data

Main rule

```
r_data = "data" SP direction SP match_data
         SP ["cs"|"pit"] SP action
```

Syntactic elements

```
direction =  "*"|"int"|"ext"
action = "forward"|"drop"
match_data = content_name SP provider
content_name = "*"|reg_exp
provider  = sign_check SP provider_sign
sign_check = "0" | "1"
provider_sign = "*"|first_sign *next_signs
```

Example

```
data * \@game|fun\@ 0 0 123456789A;FFFF0000AA pit drop
```

Pre-processing with Disco

- Sequences of more than 3 characters are extracted
- Segmented as real human-readable words
- For each word, x semantically similar words are found...
- ... and included into an extended regular expression

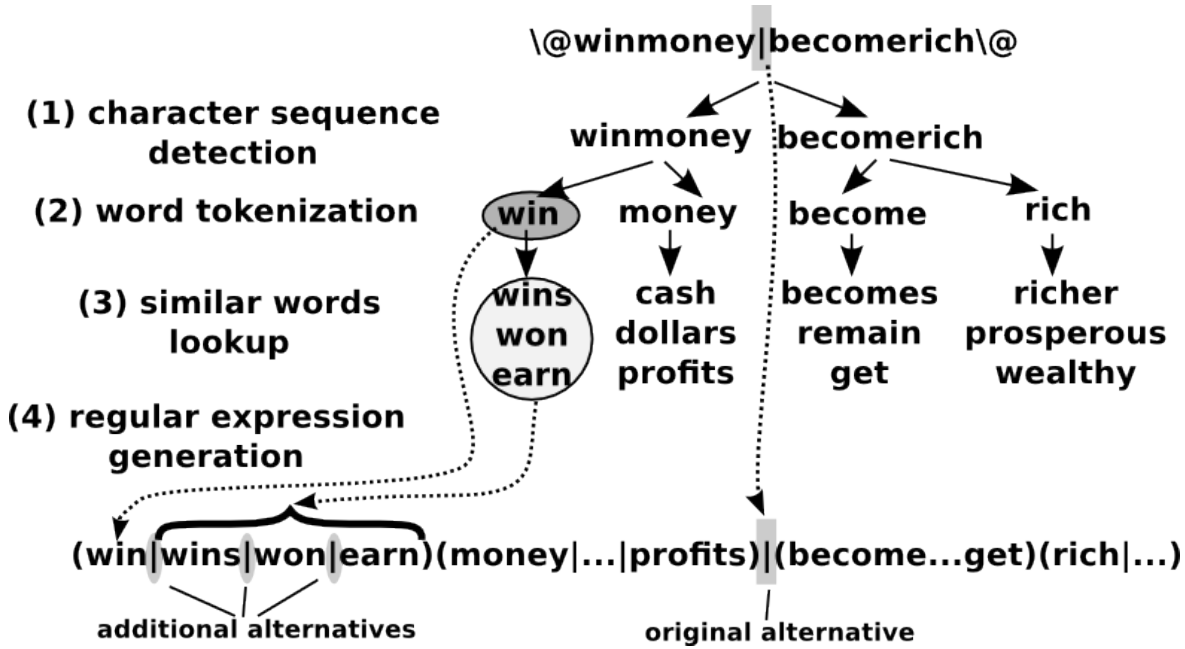

Implementation

Integration within the CCN stack

- Firewall directly processes content chunks : captures packets arriving on a face, applies rules on it, eventually calls standard CCN process

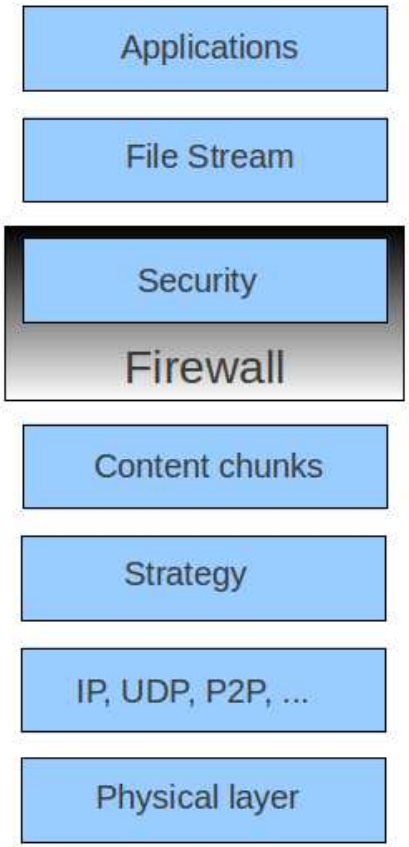

Evaluation (1/3)

CCN firewall evaluation setup

- 6 nodes
- Intermediate routers do not cache
- Consumer request single binary file
- Measurement of transmission time

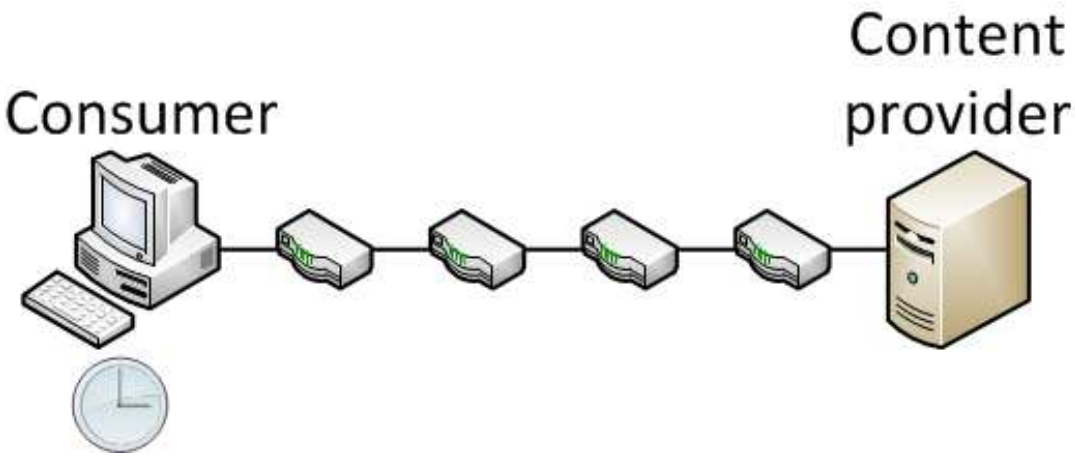

Evaluation (2/3)

Impact of the number of rules on the transfer time

- Increasing step 100 MB
- Requested files size : 500 MB and 1 GB
- Shows small to no impact on transfer time

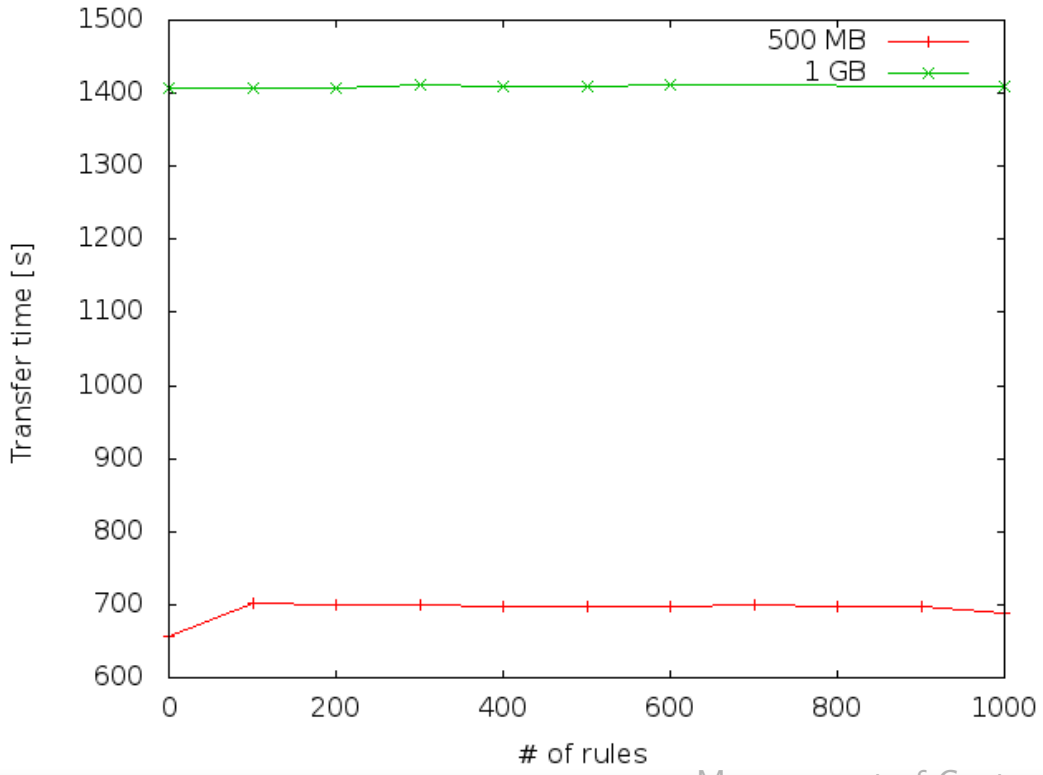

Evaluation (3/3)

Impact of a 1000-rules firewall on the transfer time

- Repeated experiment (500 MB file transfer) to obtain significant results
- Applied Chi-square and KS-test on obtain result
- Overhead of the firewall is insignificant

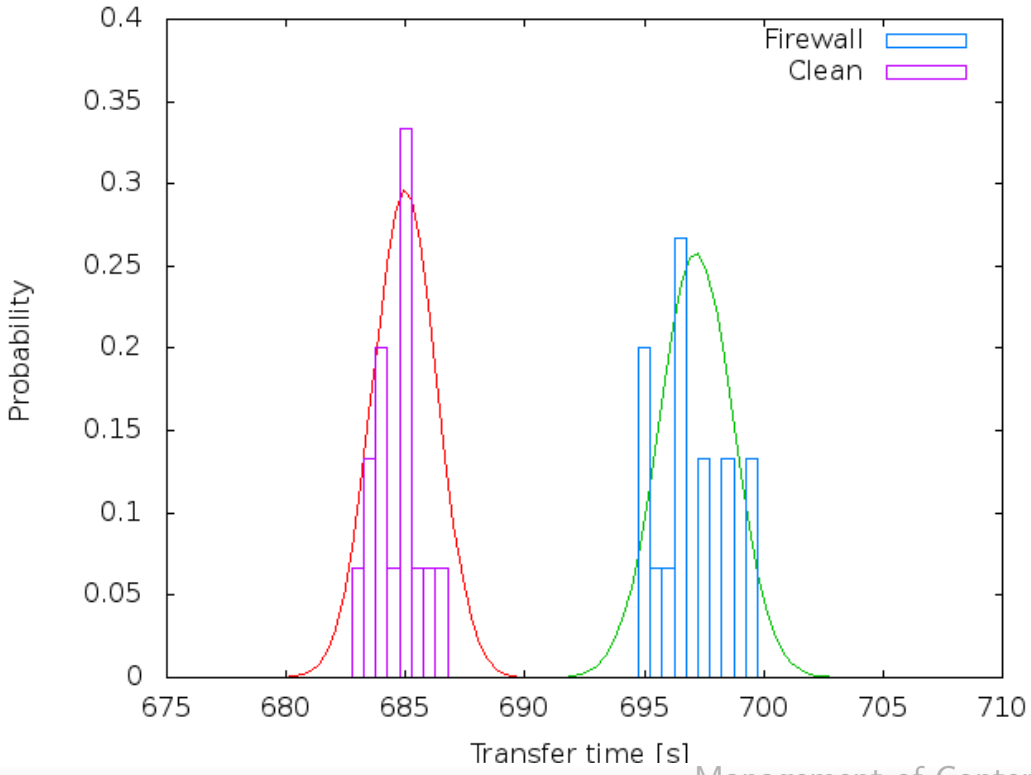

Questions ?

G. Acs, M. Conti, P. Gasti, C. Ghali, and G. Tsudik.
Cache privacy in named-data networking.
In *the 33rd International Conference on Distributed Computing Systems (ICDCS)*, 2013.

Kemal Akkaya and Mohamed Younis.
A survey on routing protocols for wireless sensor networks.
Ad Hoc Networks, 3 :325–349, 2005.

Antonio Carzaniga, Michele Papalini, and Alexander L. Wolf.
Content-based publish/subscribe networking and information-centric networking.
In *Proceedings of the ACM SIGCOMM workshop on Information-centric networking*, ICN '11, pages 56–61, New York, NY, USA, 2011. ACM.

N. Fotiou, G. F. Marias, and G. C. Polyzos.
Information ranking in content-centric networks.

pages 1–7, June 2010.

Christine Fricker, Philippe Robert, James Roberts, and Nada Sbihi.
Impact of traffic mix on caching performance in a content-centric network.
In *INFOCOM Workshops*, pages 310–315. IEEE, 2012.

Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin.
Directed diffusion : a scalable and robust communication paradigm for sensor networks.
In *Proceedings of the 6th annual international conference on Mobile computing and networking*, MobiCom '00, pages 56–67, New York, NY, USA, 2000. ACM.

Van Jacobson, Diana K. Smetters, James D. Thornton, Michael F. Plass, Nicholas H. Briggs, and Rebecca L. Braynard.
Networking named content.

In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, CoNEXT '09, pages 1–12, New York, NY, USA, 2009. ACM.

Tobias Lauinger, Nikolaos Laoutaris, Pablo Rodriguez, Thorsten Strufe, Ernst Biersack, and Engin Kirda.
Privacy risks in named data networking : what is the cost of performance ?
Computer Communication Review, 42(5) :54–57, 2012.

Sara Oueslati, James Roberts, and Nada Sbihi.
Flow-aware traffic control for a content-centric network.
In Albert G. Greenberg and Kazem Sohraby, editors, *INFOCOM*, pages 2417–2425. IEEE, 2012.

Matthias Wählisch, Thomas C. Schmidt, and Markus Vahlenkamp.
Backscatter from the data plane — threats to stability and security in information-centric networking.
CoRR, abs/1205.4778, 2012.