Named Data Networking

1. V Jacobson, D.K. Smetters, J.D. Thornton, M.F. Plass, N.H. Briggs, R.L. Braynard: Networking Named Content, CoNEXT, 2009.

2. L. Zhang, D. Estrin, J. Burke, V. Jacobson, J.D. Thornton, D.K. Smetters, B. Zhang, G. Tsudik, K.C. Claffy, D. Krioukov, D. Massey, C. Papadopoulos, T. Abdelzaher, L. Wang, P. Crowley, E. Yeh: Named Data Networking (NDN) Project, PARC Technical Report NDN-0001, 2010.
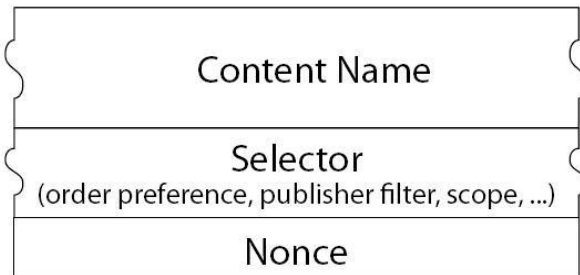
Named Hosts vs Named Data

- Originally, networks dealt with the problem of scarce resources for service provision

- However, information consumers are now less concerned with **where** the data is than **what** it is

- "Liberating" content from its originator provides a way to tackle many of the problems with current network architecture: multicast, mobility, multipath

- However, such an approach needs to be sufficiently compatible with existing technologies to be a viable alternative – i.e. it needs to be incrementally deployable
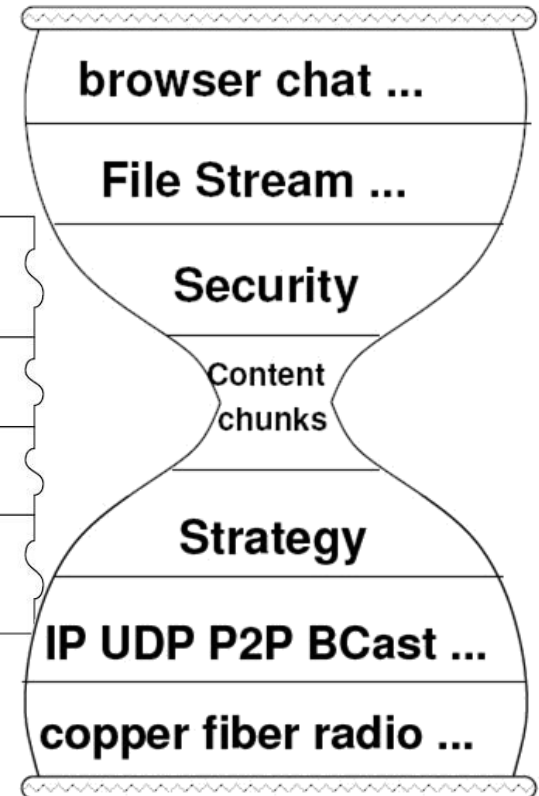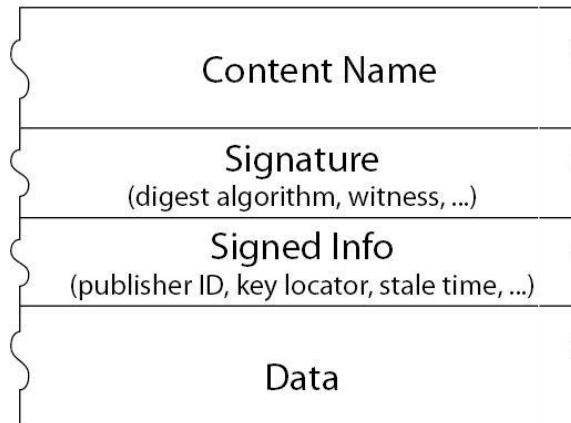
Interest and Data packets

- There are two types of NDN packets:
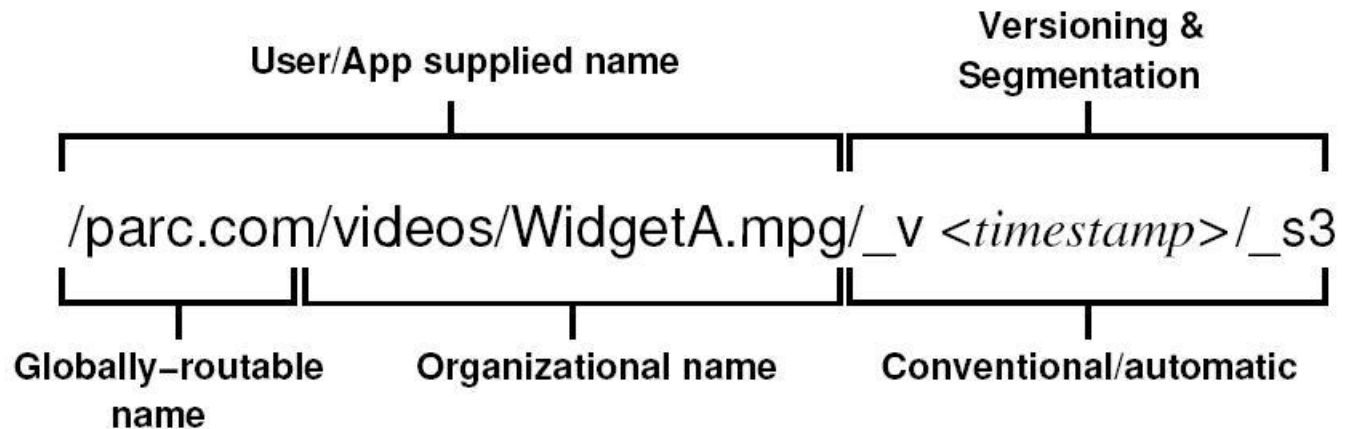 - Interest packets
  - Data packets

Interest packet

Content Name

Selector
(order preference, publisher filter, scope, ...)

Nonce

Data packet

Content Name

Signature
(digest algorithm, witness, ...)

Signed Info
(publisher ID, key locator, stale time, ...)

Data

browser chat ...

File Stream ...

Security

Content chunks

Strategy

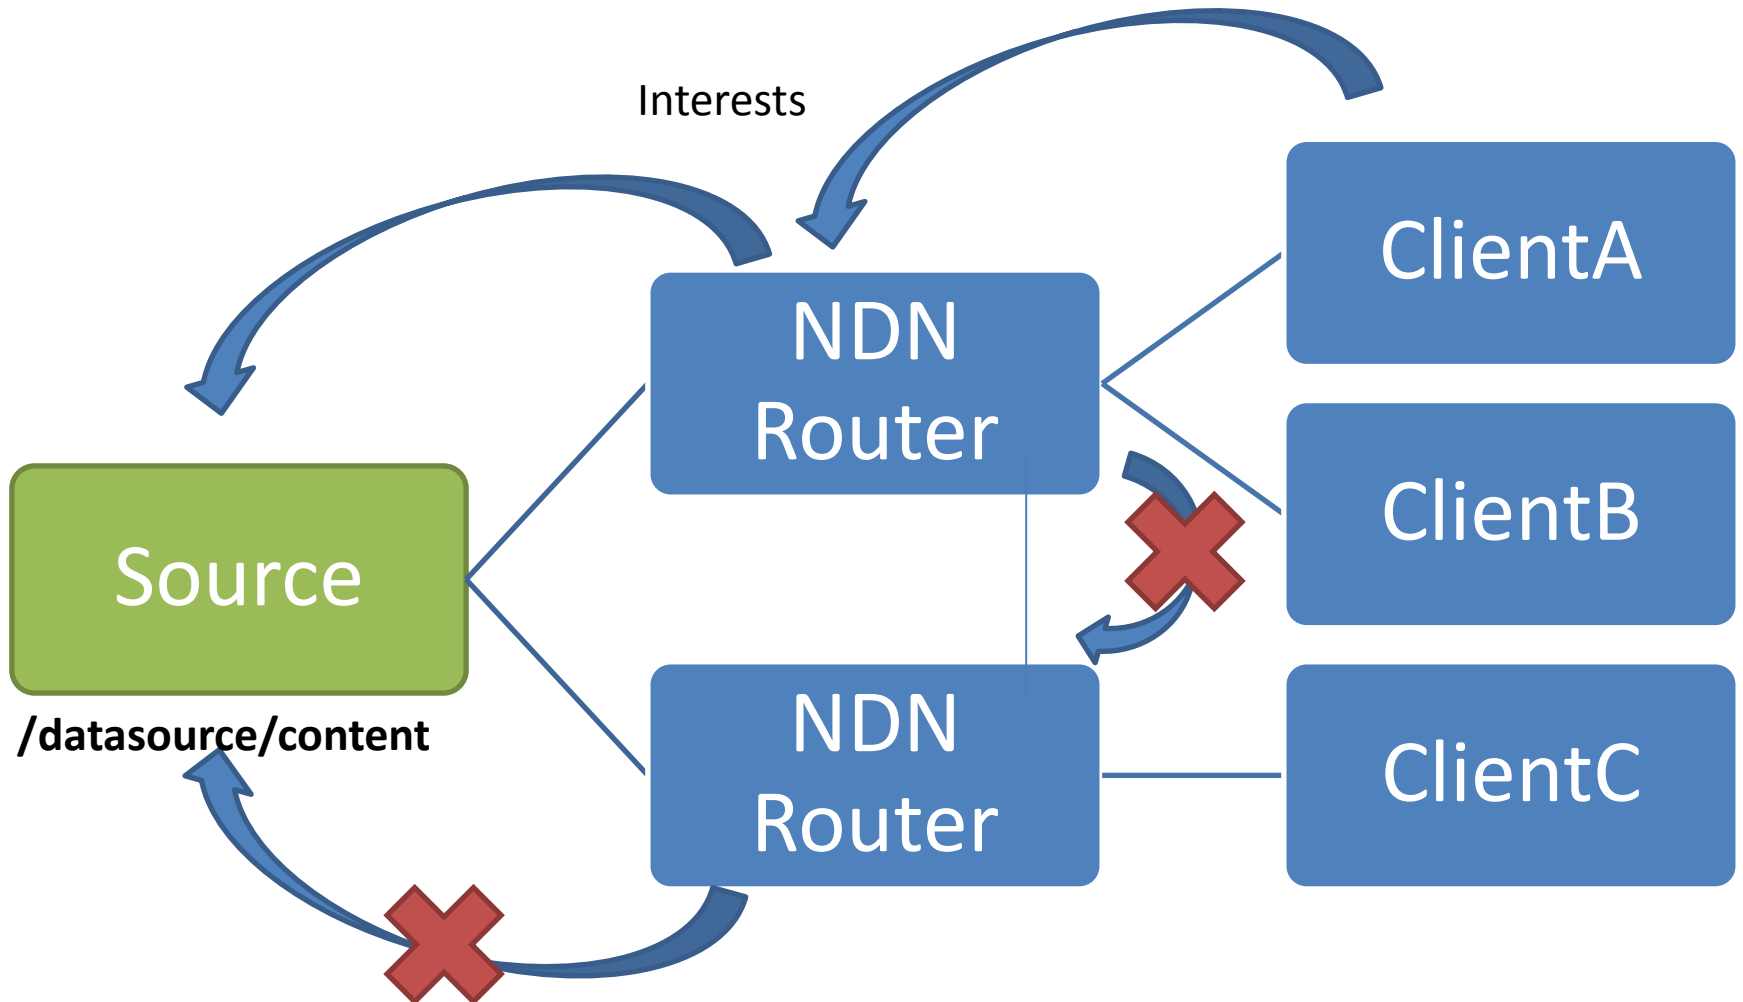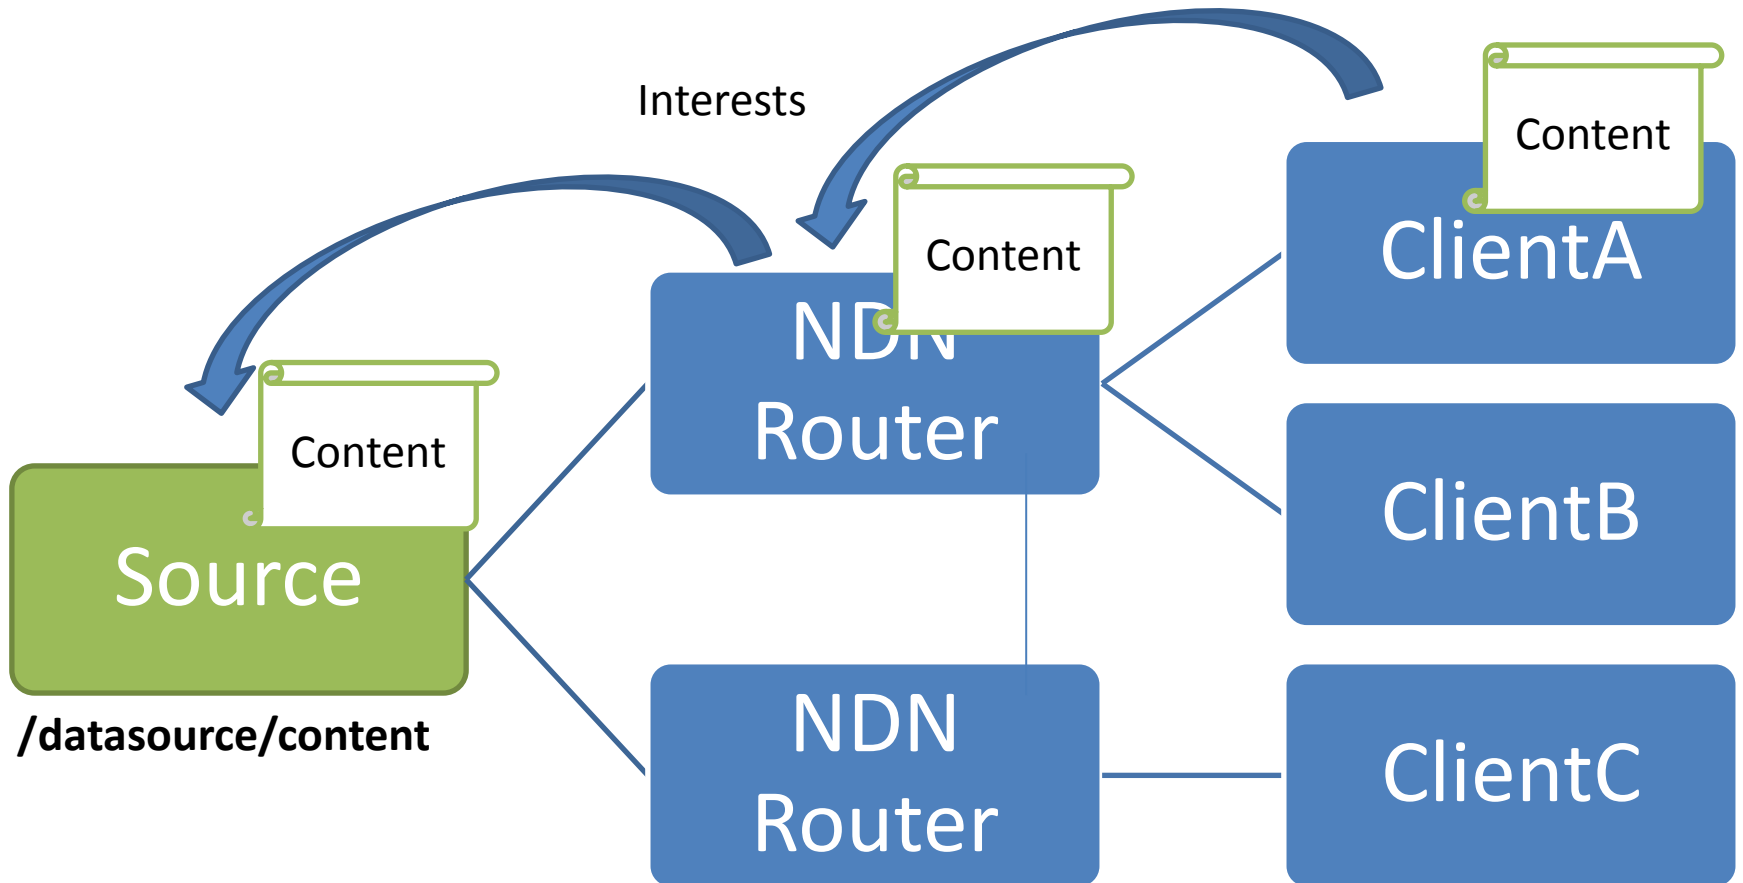IP UDP P2P BCast ...

copper fiber radio ...

3

Naming

- Data can be named using URI-like hierarchical variable length identifiers.

- At the lowest levels of the hierarchy, content is separated into versions and segments

- Content can be selected by specifying traversal rules in the interest packet

User/App supplied name

Versioning & Segmentation

/parc.com/videos/WidgetA.mpg/_v *<timestamp>*/_s3

Globally–routable name

Organizational name

Conventional/automatic

Forwarding Model

Interests

Source

/datasource/content

NDN Router

NDN Router

ClientA

ClientB

ClientC

Forwarding Model

Interests

Content

ClientA

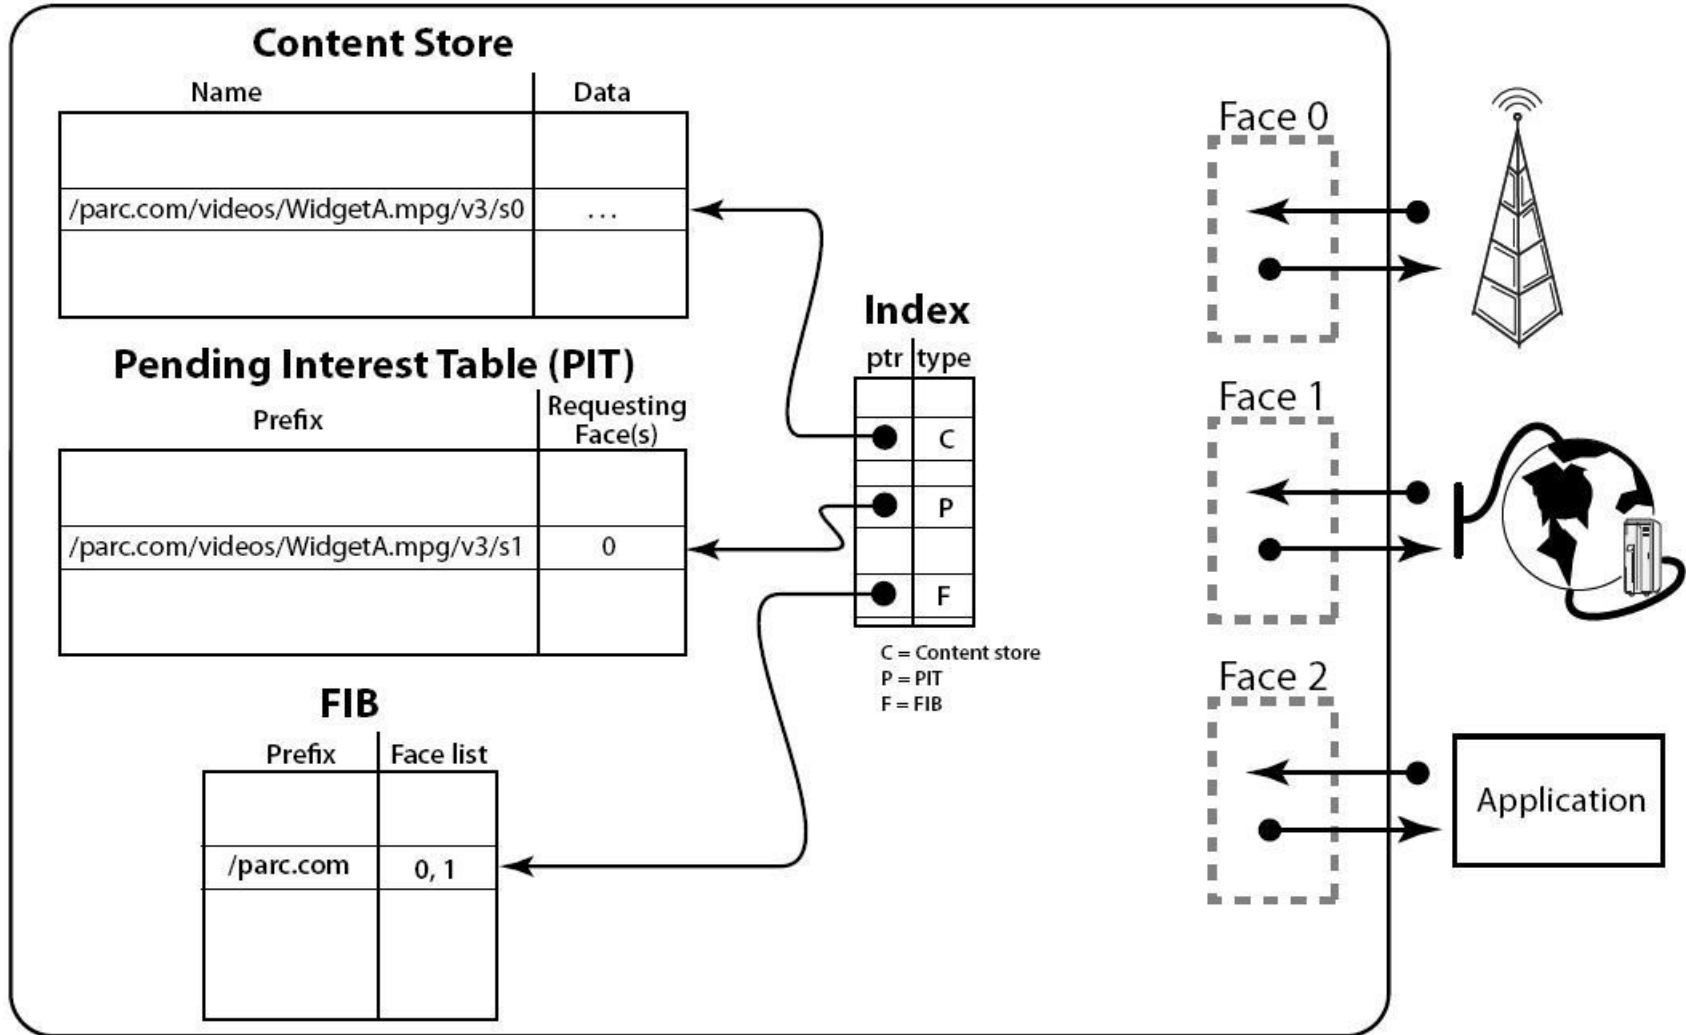Content

NDN Router

Content

Source

ClientB

/datasource/content

NDN Router

ClientC

6

NDN Forwarding Engine

Transport Strategy

- Each interest retrieves a single data packet. Applications retransmit unresolved interests.

- Multiple interests can be sent simultaneously to provide flow control functionality

- Each node in the network has its own strategy layer, which allows end-to-end flow control

- Content segment numbers also act like TCP sequence numbers. Out-of-sequence segments are buffered in the content cache.

8

Routing

- NDN needs a mechanism to allow link state advertisements that include the naming prefixes

- Broadcast-based intra-domain routing protocols like OSPF include a type-label value field in their protocol that can be used

- The FIB is updated with all matching interfaces for advertised name prefixes

- For inter-domain routing, current versions of BGP also allow propagation of name prefix

Security and Privacy

- Rather than authenticating endpoints, content segments are signed and optionally encrypted

- Certificates linking namespaces to content can be propagated like any other content, allowing key exchange

- Related content forms trust relationships and can be used to build evidence as to the authenticity of related content.

- The nature of the routing model makes attacks such as DDoS less effective

Evaluation

- The evaluation focuses on a small set of familiar applications, and shows how these operate under NDN

- Focus is more on proving the concept, rather than on proving performance

- Small scale trial of HTTP/S, VoIP and multicast

Evaluation (cont.)

Total Download Time

Throughput
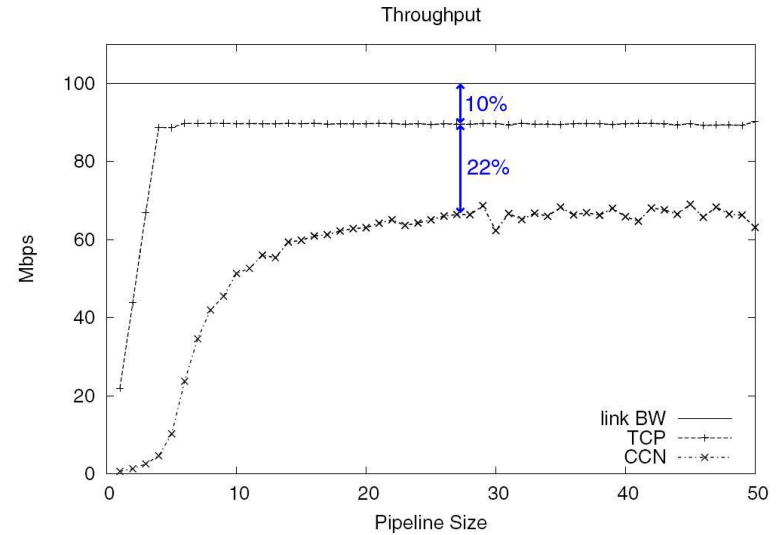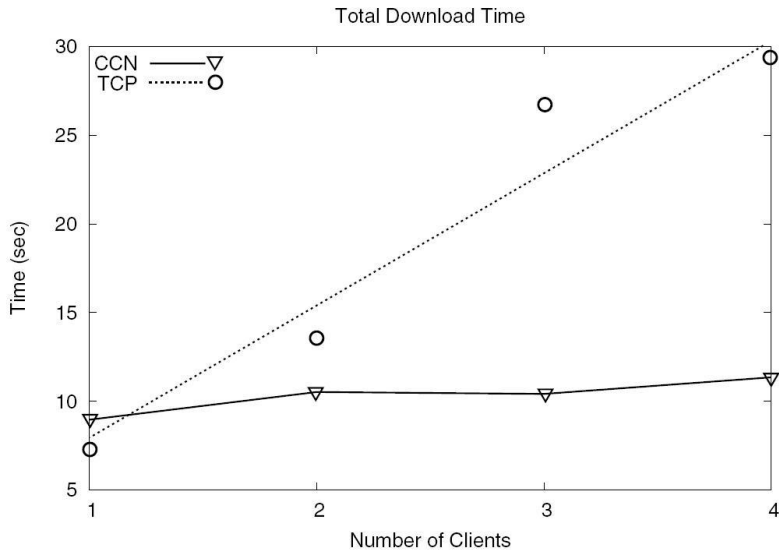
	HTTPS	CCN/UDP	CCN/ETH
App Payload	16944	16944	16944
Packets Out	16	5	16
Bytes Out	1548	629	1791
Avg. Size Out	96.75	125.8	111.94
Packets In	22	14	14
Bytes In	21232	18253	20910
Overhead In	25.31 %	7.73 %	23.41 %
Avg. Size In	965.09	1303.79	1493.57

Pros

- Conceptually elegant in the way it addresses some of the key issues affecting networks
- Can also be used with real-time streams such as VoIP and
- Can integrate with existing network hardware, utilising existing protocols
- Evaluation shows that the scheme does provide performance benefits to content distribution

Cons

- The scalability of flooding interests is questionable
- The unbounded address space could be difficult to implement on current hardware
  - [2] proposes a way to tailor names to a global/site model to allow use with TCAM hardware
- Arbitrary hierarchy could lead to untenable routing table growth
- The current evaluation with less than ten clients does not provide very strong evidence of scalability
- Evaluation also does not compare against other information-centric or traditional approaches

Questions / Comments?