

Contrace: Traceroute Facility for Content-Centric Network

<https://tools.ietf.org/html/draft-asaeda-icnrg-contrace-01>

Hitoshi Asaeda (NICT)

Xun Shao (NICT)

Thierry Turletti (Inria)

Outline

- Motivation
- Design concept
- Primary experiments
- Contrace overview and behavior
- (Packet formats and type values)
- Demo
- Conclusion

Motivation

- Protocol analysis
 - Effectiveness, robustness, and cost of designed networks
 - ▣ Hop count/RTT for content retrieval, multipath, in-network caching algorithm
 - Characteristics of content
 - ▣ Content popularity, cache hit ratio
- Operations, Administration, and Maintenance (OAM)
 - Operation
 - ▣ Cache lifetime or expiration time
 - Monitoring
 - ▣ CS capacity and usage at router, num. of interests per content
 - Trouble shooting
 - ▣ Availability of caching routers and publishers
 - (Access control and policy configuration for information disclosure)
- Active measurement tool, which could be “ICN/CCN-friendly traceroute”, is highly valuable


Our Design Concept

- Traceroute
 - Traditional tool for investigating the unicast path from source to destination
- Mtrace (Multicast traceroute)
 - Common tool for investigating the IP multicast forwarding path from destination to multicast source (or RP)
 - Mboned WG has been working for the standardization of Mtrace version 2
 - <https://tools.ietf.org/html/draft-ietf-mboned-mtrace-v2-16>
- Contrace (Primary version)
 - Tool for tracing Content-Centric Networks
 - Designed and implemented with an independent daemon that interacts with CCNx-0.8's ccnd
 - It discovers the states of in-network cache and evaluates CCN performance
 - “Contrace: A Tool for Measuring and Tracing Content-Centric Networks”, *IEEE ComMag*, Mar. 2015.

Experiments with Our Primary Implementation

P: 100 contents

C: 2500 interests



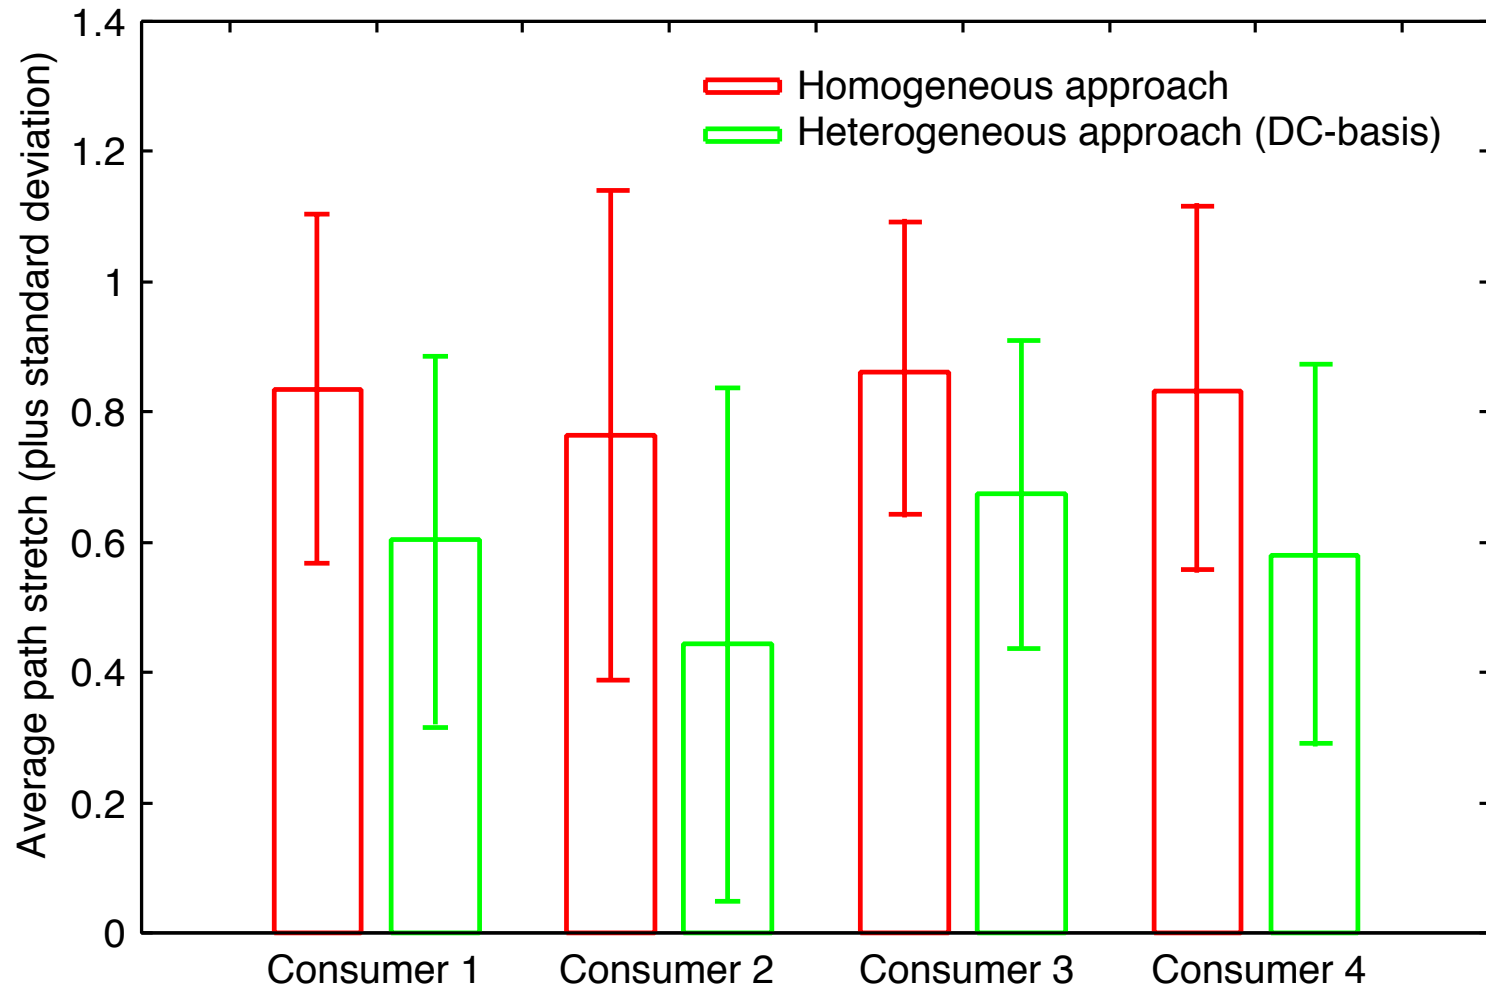
■ Setup

- GEANT network topology, which consists of 27 routers, using Rocketfuel on top of the Mini-CCNx emulator
- Contrace interacting with CCNx-0.8
- One publisher has 100 contents
- Four consumers are set up, and each consumer sends 2500 interests in total
- Content is indexed based on the Zipf popularity setting with the value of the Zipf exponent $\alpha = 1.25$


■ Measurements

- **Path stretch** and **cache efficiency** for heterogeneous approach based on router's degree centrality
 - D. Rossi and G. Rossini, "On Sizing CCN Content Stores by Exploiting Topological Information," *Proc. IEEE INFOCOM NOMEN Workshop*, March 2012.

Path Stretch



Cache Efficiency




X : Content index based on popularity
Y : Number of accesses at publisher

(New) Contrace Overview

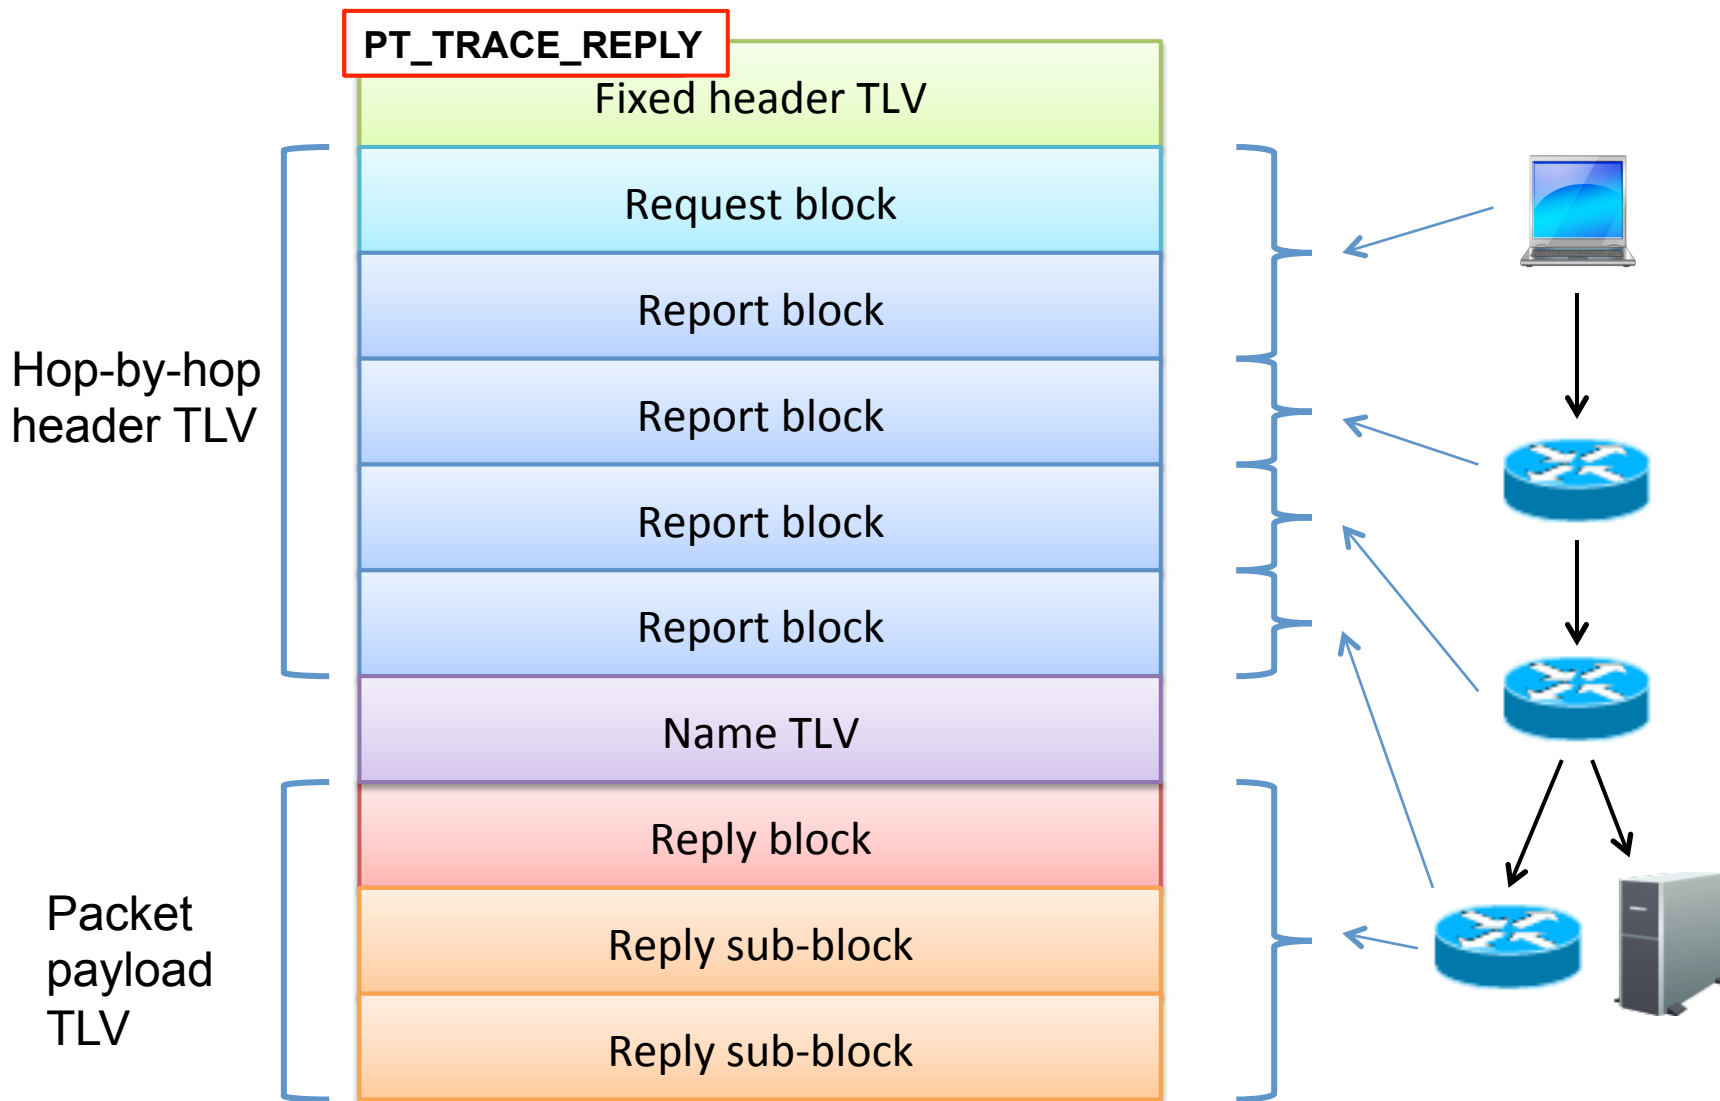
- Components
 - Contrace user program (e.g., contraceptive command running on a consumer)
 - Contrace forwarding function implemented on a content forwarder (i.e., router, publisher)
- Messages
 - Request and Reply
 - Compatible with CCNx-1.0 TLV format
- Contrace provides various information such as;
 - Hop count/RTT for content retrieval
 - Cache lifetime or expiration time at router

Contrace Basic Behavior

- **Request message** is initiated by Contrace user and forwarded toward caching router or publisher based on the FIB in a hop-by-hop manner
- Request message includes **Request block** and **Report block(s)**
- **Reply message** is initiated by caching router or publisher and forwarded toward Contrace user based on the PIT entry
- Reply message includes **Reply block** and **Reply sub-block(s)**



Contrace Message Flow



Unique Behavior

- Multipath support
 - Unlike the regular Interest-Data communication, PIT entry created by Contrace Request is removed only when the [Contrace Reply Timeout] value is expired
 - Unlike the regular Interest-Data communication, Request message is sent to all neighbor routers simultaneously (even if some strategy schedules the delay)
- Contrace Requests SHOULD NOT result in PIT aggregation in routers during the Request message transmission.
- Contrace Replies MUST NOT be cached in routers upon the Reply message transmission.

What Contrace Can Provide (Re: Motivation)

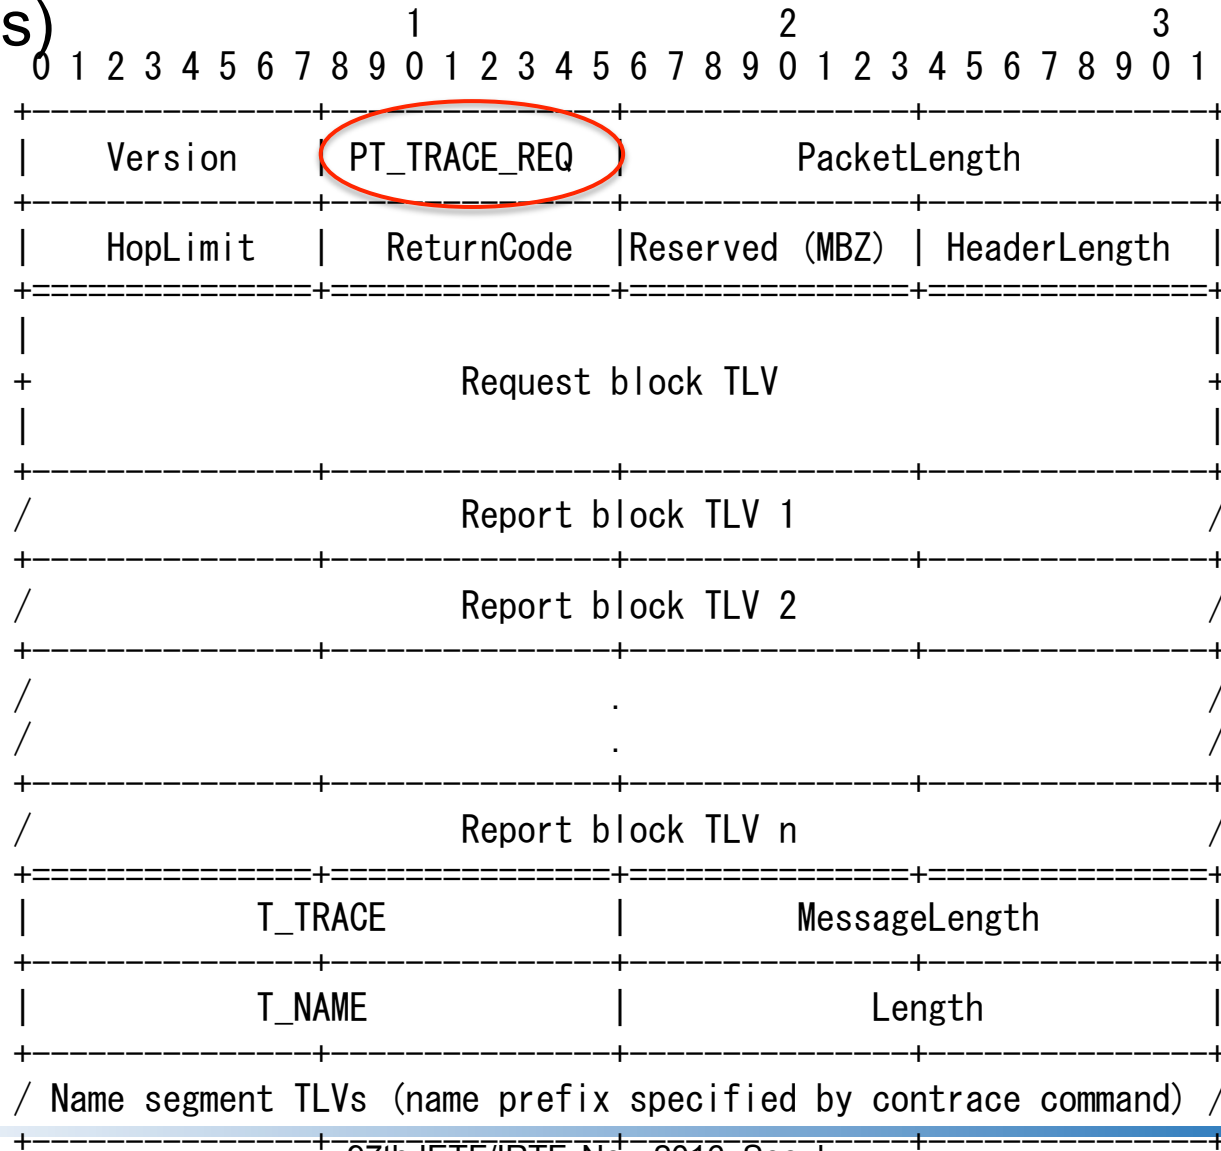
- Protocol analysis
 - Effectiveness, robustness, and cost of designed networks
 - Hop count/RTT for content retrieval, multipath, in-network caching algorithm
 - Characteristics of content
 - Content popularity, cache hit ratio
- Operations, Administration, and Maintenance (OAM)
 - Operation
 - Cache lifetime or expiration time
 - Monitoring
 - CS capacity and usage at router, num. of interests per content
 - Trouble shooting
 - Availability of caching routers and publishers
 - (Access control and policy configuration for information disclosure)

Gateway Discovery

- “Gateway” is a router supporting different scheme names (e.g., ccnx:/ and ndn:/) and having both FIBs
 - Discussing about interoperability or integration of CCN and NDN is not the aim of this document, but ...
 - It is good if we have a chance to discover the gateway, so ...
- Contrace could also support discovering the closest gateway

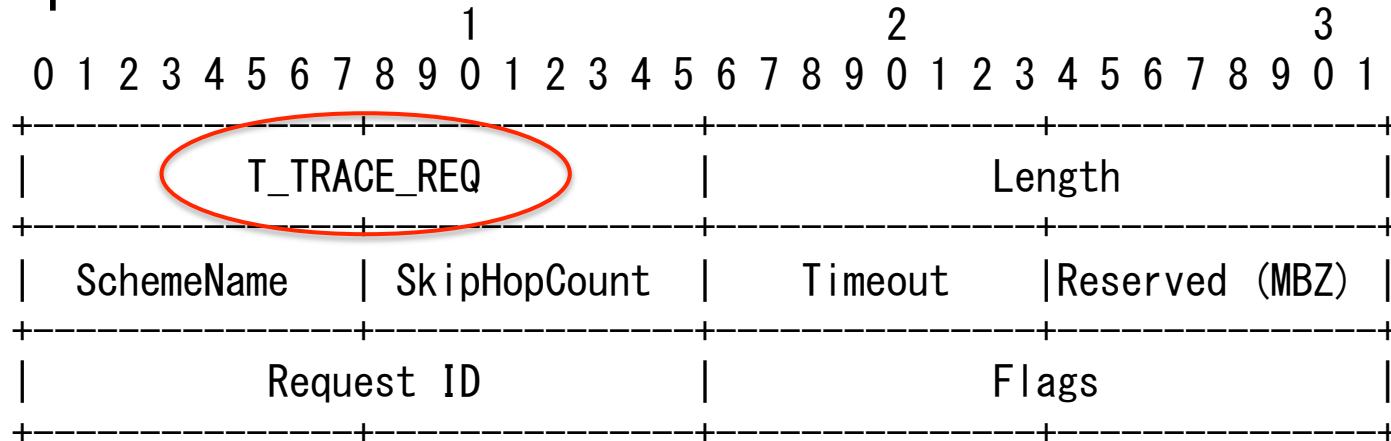
Request Message

- Request message, including Request block, and Report block(s)

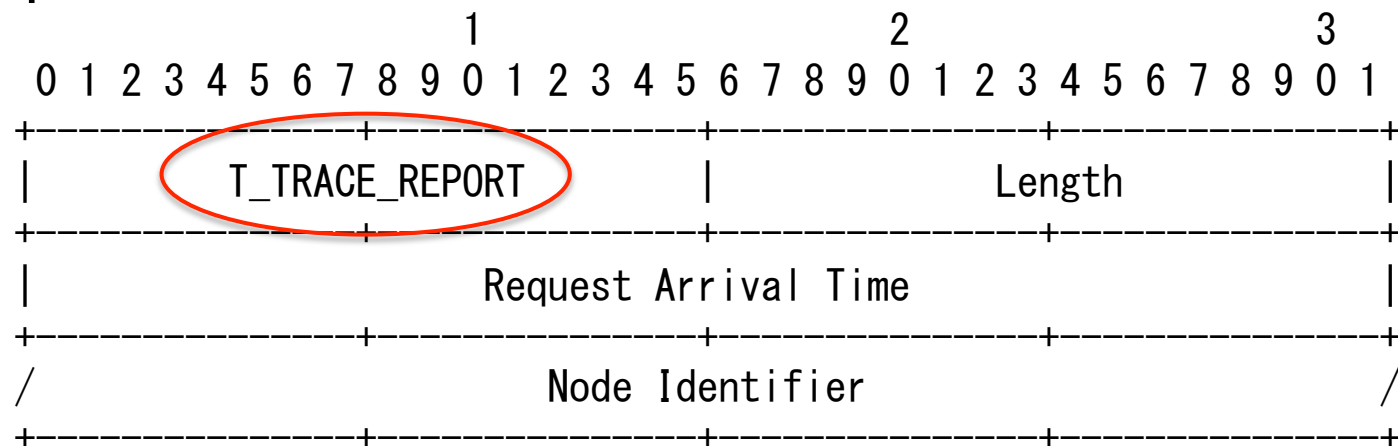


Request Block and Report Block

- Request block

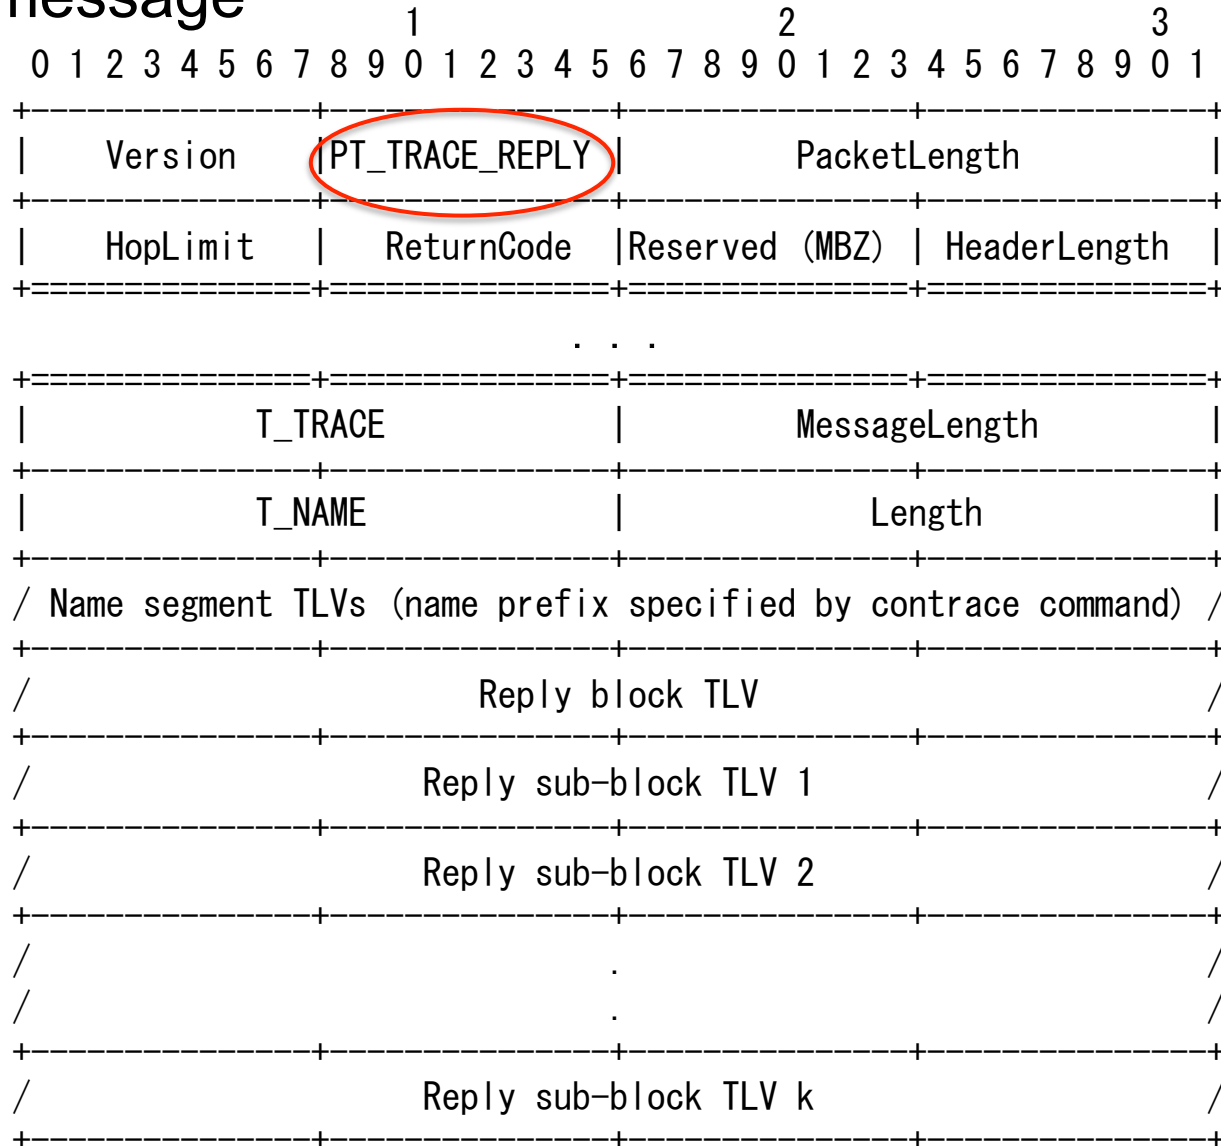


- Report block



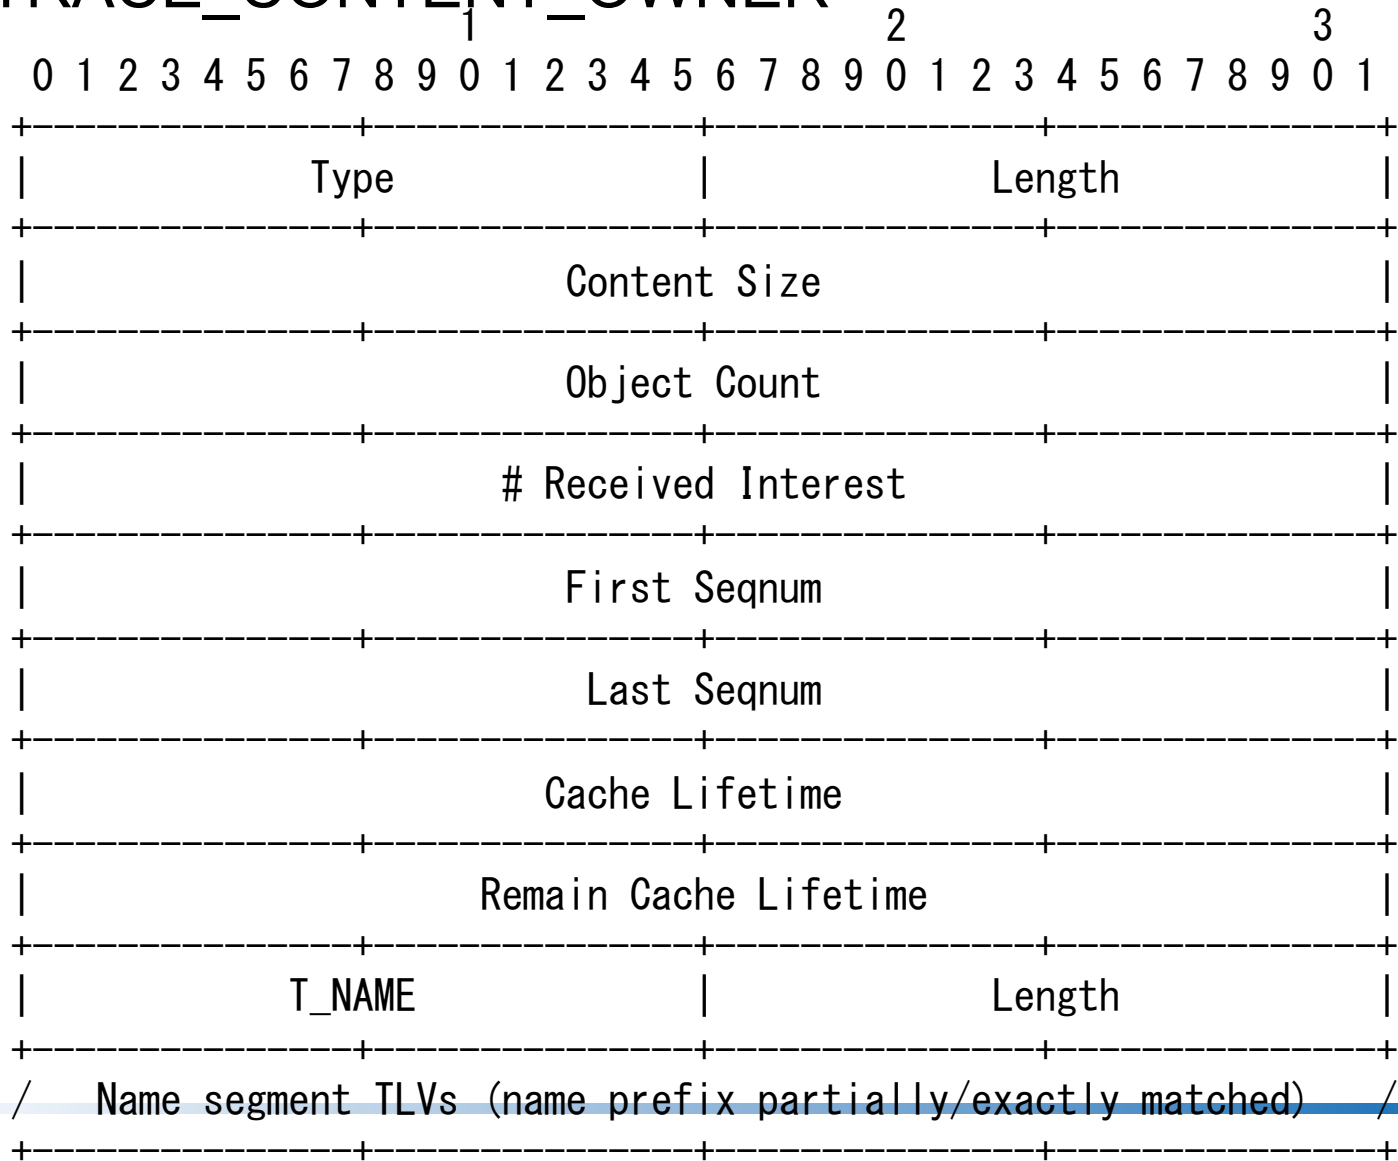
Reply Message

■ Reply message



Reply Sub-Block

- Reply sub-block (for T_TRACE_CONTENT and T_TRACE_CONTENT_OWNER)



Contrace User Command

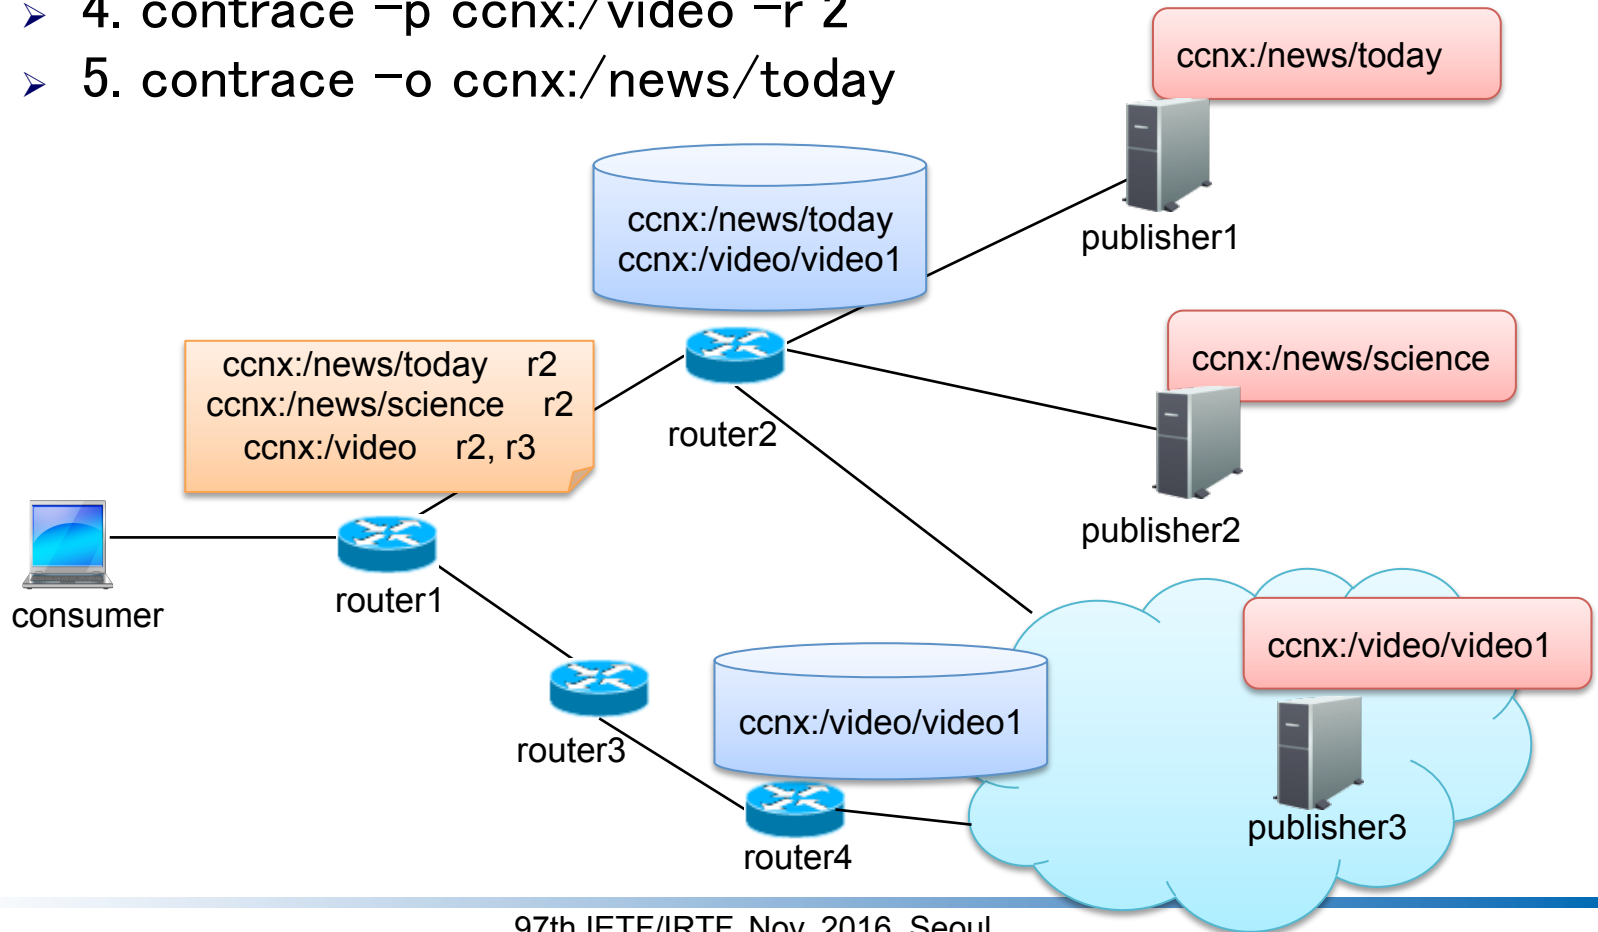
Usage: `contrace [-p] [-g] [-n] [-o] [-r hop_count] [-s hop_count] [-w wait_time] name_prefix`

Options	Descriptions
<code>[-p] <i>name_prefix</i></code>	Mandatory option. Name prefix of the content (e.g., <code>ccnx:/news/today/</code>) the Contrace user wants to trace. “ <code>ccnx:/</code> ” can be allowed for permitted users. The <code>-p</code> option allows a partial match for the name prefix; otherwise, an exact match is required
<code>-n</code>	No cache information is requested. Only the routing path information to the specified content/cache and RTT between Contrace user and content forwarder informed.
<code>-o</code>	This option can be specified if a Contrace user needs to trace the path to the content publisher.
<code>-r <i>hop_count</i></code>	Number of traced routers. If the Contrace user specifies this option, only the specified number of hops from the Contrace user trace the Request.
<code>-s <i>hop_count</i></code>	Number of skipped routers. If the Contrace user specifies this option, the number of hops from the Contrace user simply forward the Contrace Request messages without adding its own Report block and without replying the Request, and the next upstream router starts the trace.
<code>-w <i>wait_time</i></code>	Contrace Reply Timeout value (in seconds) that the <code>contrace</code> user waits for Response. After the timeout, the <code>contrace</code> user silently discards the Response. Default: 4 s.

Demo


■ Commands

- 1. `contrace ccnx:/news/science`
- 2. `contrace -n ccnx:/video/video1`
- 3. `contrace -p ccnx:/video`
- 4. `contrace -p ccnx:/video -r 2`
- 5. `contrace -o ccnx:/news/today`




1. contrace ccnx:/news/science

```
cefuser@ceflab:~$  
cefuser@ceflab:~$ contrace ccnx:/news/science  
contrace to ccnx:/news/science with HopLimit=32, SkipHopCount=0, Flag=0x0000 and Request ID=142  
response from publisher2: no error, time=1.699000 ms  
route information:  
1 consumer          0.119 ms  
2 router1           0.442 ms  
3 router2           0.261 ms  
4 publisher2        0.376 ms  
cache information:  prefix      size    cobs    interests  start-end  lifetime  expire  
1 p ccnx:/news/science/  8658 KB   8456      0    0-8455    634 secs  2951 secs  
cefuser@ceflab:~$
```




2. contrace -n ccnx:/video/video1

```
cefuser@ceflab:~$  
cefuser@ceflab:~$ contrace -n ccnx:/video/video1  
contrace to ccnx:/video/video1 with HopLimit=32, SkipHopCount=0, Flag=0x0002 and Request ID=11717  
  
response from router2: no error, time=1.516000 ms  
  
route information:  
1 consumer          0.715 ms  
2 router1           0.228 ms  
3 router2           0.275 ms  
  
response from router4: no error, time=2.136000 ms  
  
route information:  
1 consumer          0.715 ms  
2 router1           0.228 ms  
3 router3          -2.128 ms  
4 router4           3.231 ms  
  
cefuser@ceflab:~$
```




3. contrace -p ccnx:/video

```
cefuser@ceflab:~$  
cefuser@ceflab:~$ contrace -p ccnx:/video  
contrace to ccnx:/video with HopLimit=32, SkipHopCount=0, Flag=0x0001 and Request ID=5417  
  
response from router2: no error, time=1.648000 ms  
  
route information:  
1 consumer          0.735 ms  
2 router1           0.226 ms  
3 router2           0.713 ms  
  
cache information:  prefix      size    cobs    interests  start-end  lifetime  expire  
1 c ccnx:/video/video1/      4908 KB    4793      1    0-4792    2920 secs  652 secs  
  
response from router4: no error, time=2.037000 ms  
  
route information:  
1 consumer          0.735 ms  
2 router1           0.226 ms  
3 router3           0.994 ms  
4 router4           2.871 ms  
  
cache information:  prefix      size    cobs    interests  start-end  lifetime  expire  
1 c ccnx:/video/video1/      4908 KB    4793      0    0-4792    2932 secs  652 secs  
  
cefuser@ceflab:~$
```




4. contrace -p ccnx:/video -r 2

```
cefuser@ceflab:~$  
cefuser@ceflab:~$ contrace -p ccnx:/video -r 2  
contrace to ccnx:/video with HopLimit=2, SkipHopCount=0, Flag=0x0001 and Request ID=29312  
response from router3: no info, time=1.566000 ms  
route information:  
1 consumer          0.799 ms  
2 router1           0.300 ms  
3 router3           0.568 ms  
response from router2: no error, time=2.131000 ms  
route information:  
1 consumer          0.799 ms  
2 router1           0.300 ms  
3 router2           0.342 ms  
cache information:  prefix      size    cobs    interests  start-end    lifetime    expire  
1 c ccnx:/video/video1/      4908 KB    4793      1    0-4792    1314 secs    2257 secs  
cefuser@ceflab:~$  
cefuser@ceflab:~$
```



5. contrace -o ccnx:/news/today

```
cefuser@ceflab:~$  
cefuser@ceflab:~$ contrace -o ccnx:/news/today  
contrace to ccnx:/news/today with HopLimit=32, SkipHopCount=0, Flag=0x0004 and Request ID=29932  
response from publisher1: no error, time=1.888000 ms  
  
route information:  
1 consumer          0.142 ms  
2 router1           0.357 ms  
3 router2           0.340 ms  
4 publisher1        1.198 ms  
  
cache information:  prefix      size      cobs      interests  start-end  lifetime  expire  
1 p ccnx:/news/today/ 1565 KB    1529      5    0-1528    204 secs  3319 secs  
cefuser@ceflab:~$
```



Conclusion

- Contrace is a powerful tool for tracing Content-Centric Networks
- Contrace uses two messages, Contrace Request and Reply messages, which are compatible with CCNx-1.0 TLV format
- Contrace provides various information, e.g.,
 - Hop count/RTT for content retrieval
 - Cache lifetime or expiration time
 - Num. of received interests
- Implementation is on-going (and partially done)
- Next step: Revise the I-D
 - Improve the quality and fix editorial errors
 - Address [TBD]s