


# Update on our experimental Wire-speed CCN / NDN packet format.

Massimo GALLO

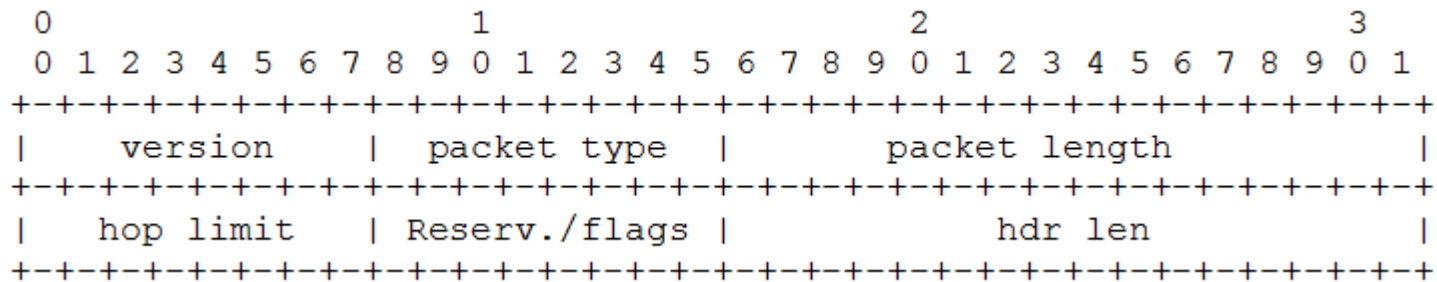
13 January 2015



- We had 4 different TLV types discriminated through first 2 bits:
  - 00 with one Byte for the Type and one Byte for the Length ( 1B-1B )
  - 01 with one Byte for the Type and one Byte for the Length ( 1B-2B )
  - 10 with one Byte for the Type and one Byte for the Length ( 2B-1B )
  - 11 with one Byte for the Type and one Byte for the Length ( 2B-2B )
- We limit TLV types to the latter one 2B-2B for simplicity.:
  - Avoid aliasing,
  - simpler parsing.
  - If needed we can add further TLV types later.



- Fixed header (Convergence Layer)
- Hop-by-Hop fragmentation/reassembly is required but not yet defined :
  - May use additional fixed header fields (through hdr length and / or reserved )
    1. Variable length (optional Hop-by-Hop TLVs) ?
    2. fixed length (additional fields)?
  - Reassembly is mandatory for both data and interests
- 

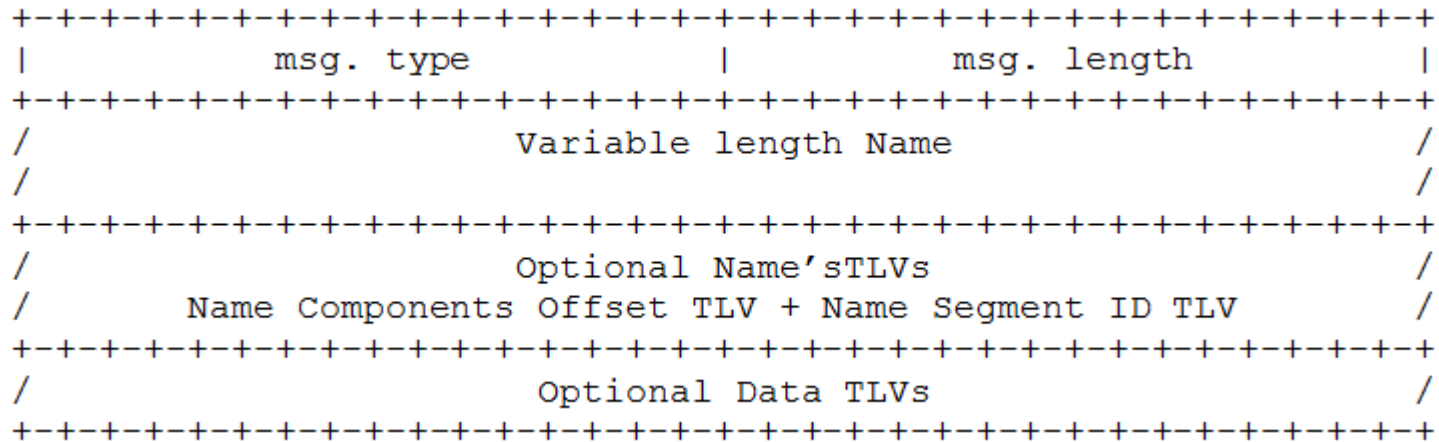


```

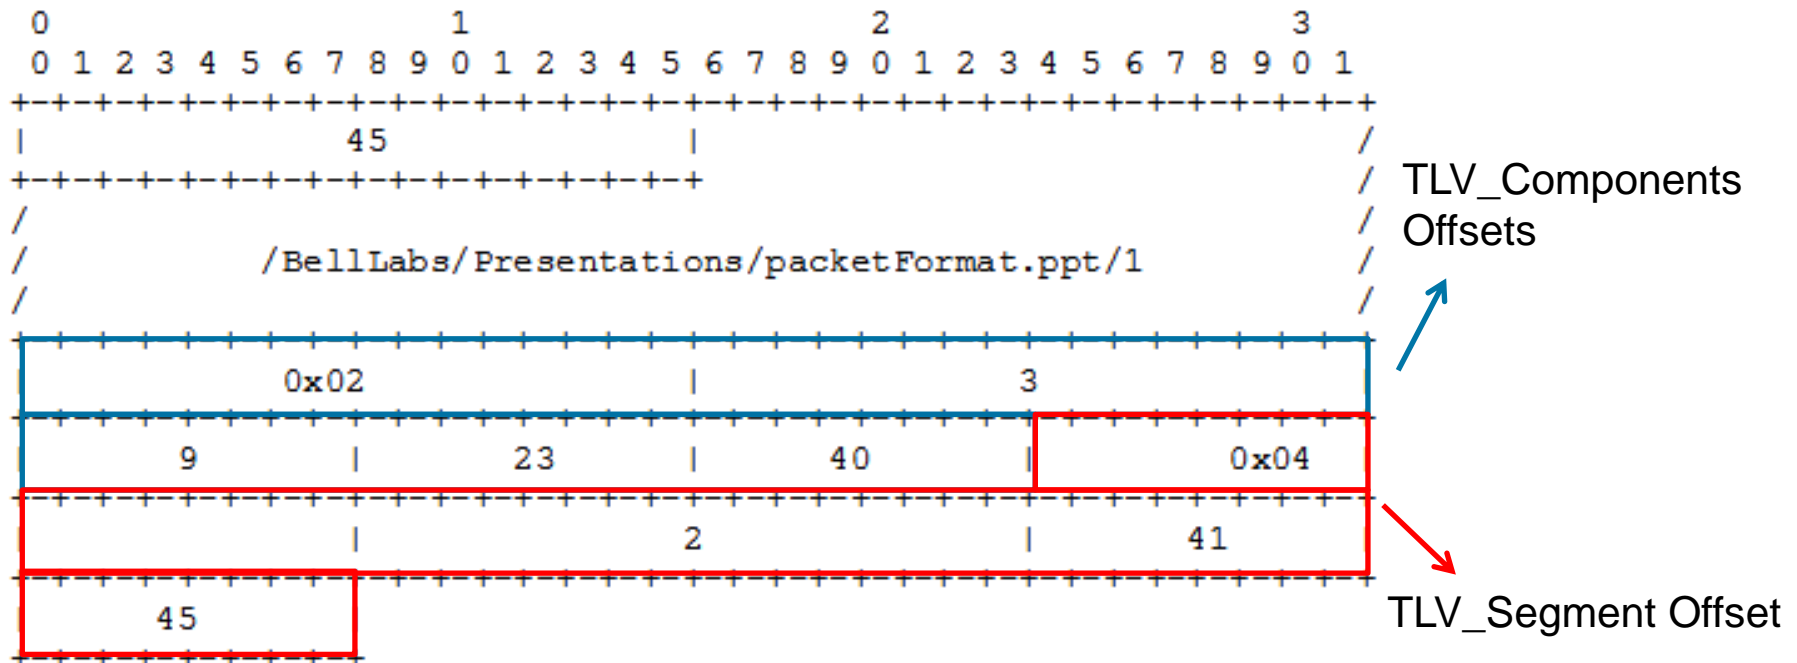
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           msg. type           |           msg. length           |
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/                               Variable length Name                               /
/                                                                                               /
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/                               Optional Name's TLVs                               /
/      Name Components Offset TLV + Name Segment ID TLV      /
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/                               NONCE TLV                               /
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
/                               Optional Interest TLVs                               /
+++++-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

- A TLV for the name is not required.
- Due to Hop-By-Hop fragmentation and reassembly through the fixed header, msg. Type and msg. Length are needed
- Nonce is mandatory and inserted after the name
  - Used to prevent loops coupled with name (only associated to a PIT entry, any specific data structure)



- Fixed Header is the same of Interest
- NONCE is not needed
- Signature and security TLVs (Signature is mandatory)

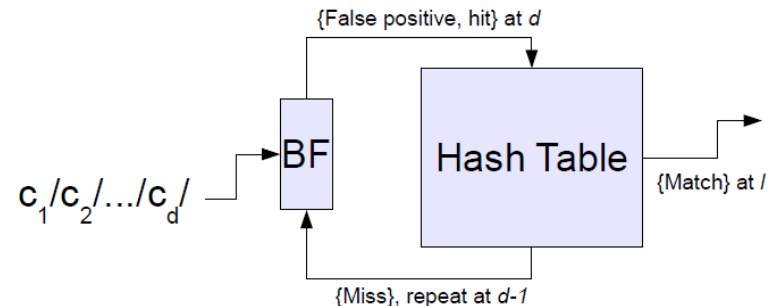


- No TLV inside a name container. Field (component) separator needed
- Offset TLV contains the list of offsets of the different components 2B (or 1B depending on the compact/extended Component offset TLV used) :
  - Fast parsing form Exact Match (PIT, CS) and LPM (FIB)
  - Name segment id is encoded with a numeric type (i.e. 4byte unsigned int)
  - Other *special* components (e.g., version) may be added similarly to segm id)

- Intel DPDK version of our content router (credits goes to L. Saino - UCL)
- General purpose server, Intel Xeon processor:
  - L1 Cache 32 KB (each core)
  - L2 Cache 256 KB (each core)
  - L3 Cache 8MB
  - Intel 82599EB 10Gbps card
- Test on **FIB lookup** and **Hash computation+Name decoding** cost
  - 60B names composed of 5 or 10 components with different Name encodings
  - FIB with 1 million prefixes

FIB Lookup, two stages approach [1]:

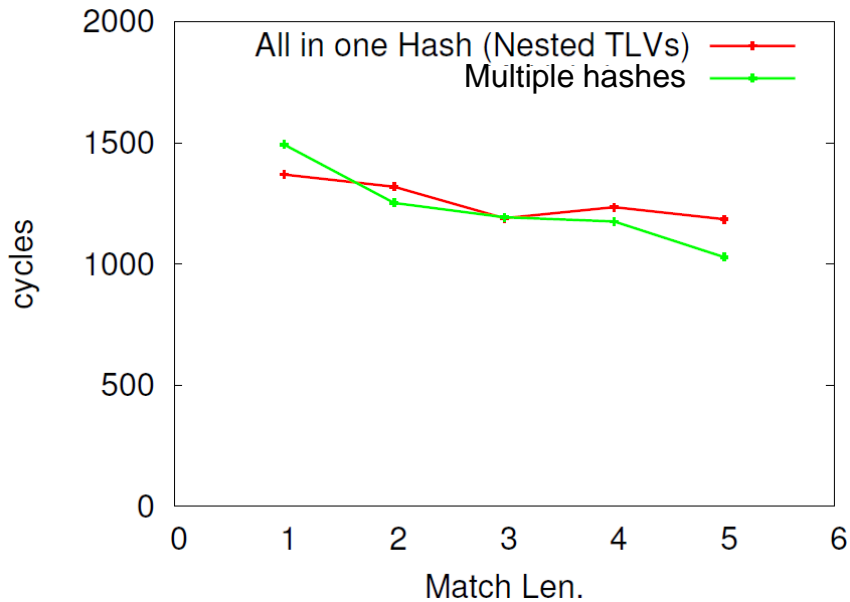
1. Check the Bloom filter  
(Check length  $n$ ,  $n-1$ , ...)
2. Check hash at identified length  $d$




[1] D. PERINO, M. VARVELLO, L. LINGUAGLOSSA, R. LAUFER, R. BOISLAIGUE,  
"Caesar: A Content Router for High-Speed Forwarding on Content Names", ACM/IEEE ANCS 2014.

- 60B names with components 5:
- 2 name encoding under test:
  - Name +offsets (1 Byte component separator needed)  
Total length: 65B, Components length: 13B 7,6% overhead
  - Nested TLVs (T+L = 2B+2B)  
Total length: 80B, Components length: 12B (+4B of T+L). 25% overhead

### Lookup cost




### Hash + parsing cost



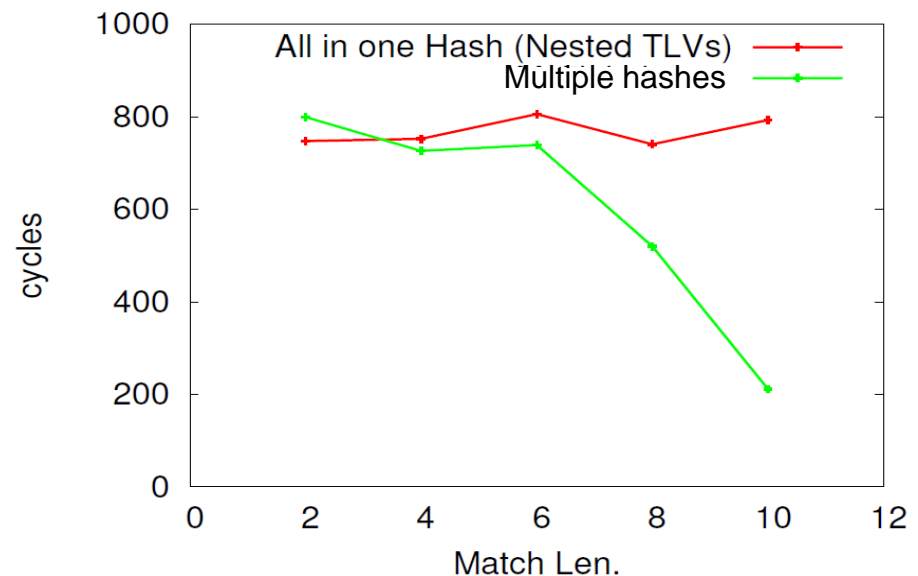


- 60B names with components 10:
- 2 name encoding under test:
  - Name +offsets (1 Byte component separator needed)  
Total length: 70B, Components length: 7B. 12.25% overhead
  - Nested TLVs (T+L = 2B+2B)  
Total length: 100B, Components length: 6B (+4B of T+L) 40% overhead

### Lookup cost



### Hash + parsing cost



## Fixed header + extension

- Topological information (e.g., Hop limit)
- Hop-by-hop fragmentation and reassembly

## Naming

- Only few type of *special* components (segment, what else?)

## Under discussion

- TLV format: for the moment 2B+2B. Need more?
- Payload for Interest packet ? Not for the moment; security issues.
- ICN architecture interoperability
- Network management/processing commands