CCN Forwarding Engine Based on Bloom Filters

M. Tortelli, L. A. Grieco, G. Boggia,

Dipartimento di Elettrotecnica ed Elettronica (DEE)

Politecnico di Bari - Italy

CFI 2012

The 7th International Conference on Future Internet Technologies

11 – 12 Sept. 2012

Imperial Palace Hotel, Seoul, Korea

Outline

- From “WHERE” to “WHAT”
- Content Centric Networking (CCN)
- Bloom Filter Basics
- Update Times Evaluation
- The proposed Forwarding Engine
- Conclusion
- Future works

From “WHERE” to “WHAT”

… *ARPAnet* (1969) ⇨ **TCP/IP** … *INTERNET* (1971) ⇨ …2012

No Architectural Change!

✓Host Centric

✓Conversational Model

However…the **SCENARIO** is changed!!

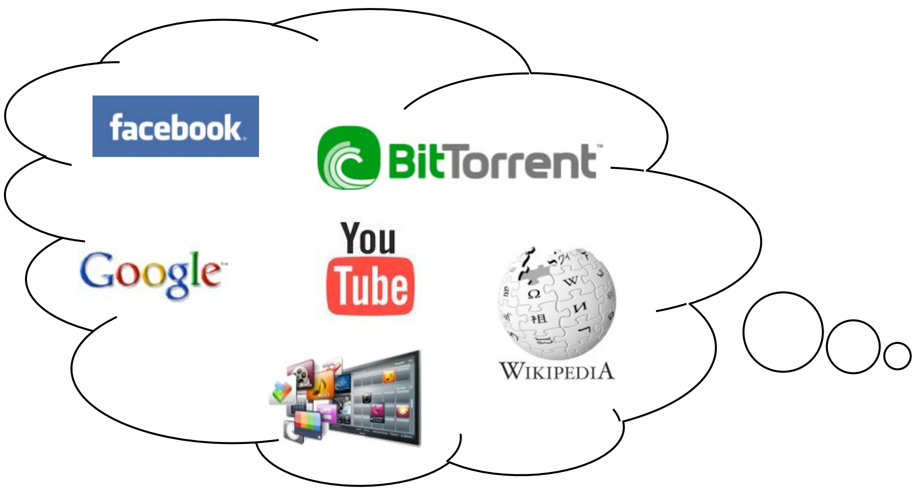

It is all about **CONTENTS!**

From “WHERE” to “WHAT”

Future Internet Goals:

- Simplify the usability
- Increase the efficiency

- Secure the privacy
- Enhance the media experience of the users

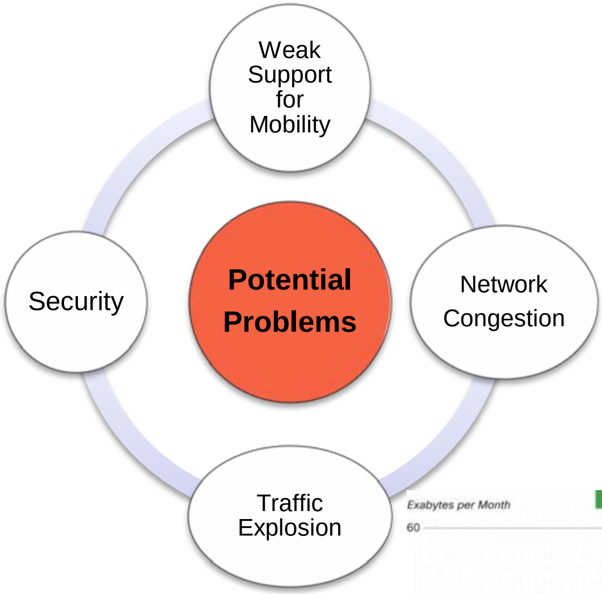

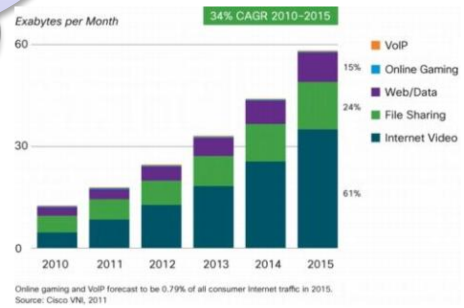

CONTENTS as a **“PRIMITIVE”**

telematics lab

Content Centric Networking (CCN)

PIVOTAL POINTS

1. Every node is **Content-Aware**
2. Routing and Forwarding strategies are based on **Hierarchical Content Names** and are applied only to *Interests*
3. **Distributed Caching Operations** permit Content Retrieval to be Location Independent

TARGET = Internet-scale deployment of CCN

… PROS

- ✓ Servers load reduction
- ✓ Improvement of the *QoE* perceived by users
- ✓ Security issues independent of contents locations
- ✓ Compliant with the existing infrastructure

Content Centric Networking (CCN)

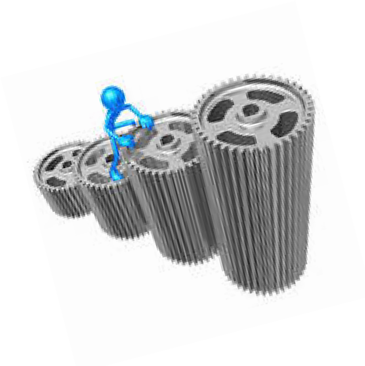

There are **INSIGHTS** along the way to go!

NAMING CONVENTIONS

SCALABILITY ISSUES

SIGNALLING OVERHEAD

Content Centric Networking (CCN)

NAMING

EVERY CONTENT OBJECT MUST BE UNIQUELY IDENTIFIABLE!

HIERARCHICAL **youtube.com/video/Olympics/London_2012**	**FLAT LABEL** **18D73B01_598A6DFF9117CEDA**
Human Readable	Self-certifying
Simple Aggregation	Easily Globally Unique
Longest –Match Lookup	More Flexible
Lack of Global Uniqueness	No Aggregation
	Need for External Binding

Content Centric Networking (CCN)

There are **INSIGHTS** along the way to go!

NAMING CONVENTIONS

SCALABILITY ISSUES

SIGNALLING OVERHEAD

SCALABILITY

BGP TODAY	**≈ $4 \cdot 10^5$ routes in a BGP Routing Information Base**
CONTENT-BASED ROUTING PROTOCOL	**Google reports a lower bound of 10^{12} unique URLs → we need at least $O(10^{12})$ routes to account for all the unique content objects on the Web!!**

Lookup times will degrade forwarding performances!

Content Centric Networking (CCN)

There are **INSIGHTS** along the way to go!

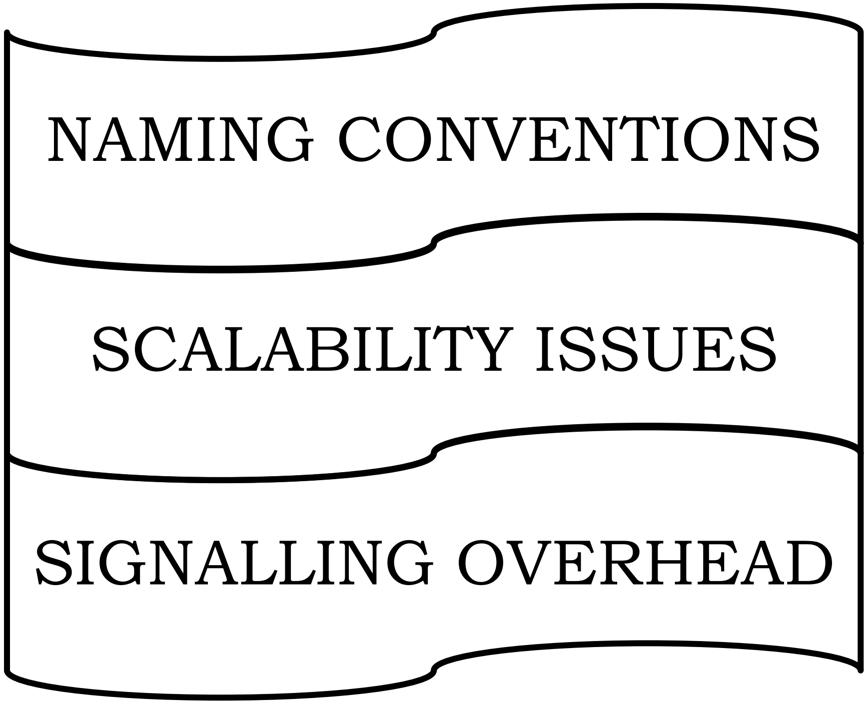

SIGNALING OVERHEAD

Universal Caching Capability

Every network entity can cache every content it forwards in order to satisfy subsequent requests.
Reduction of contents retrieval times thanks to nearest copies with respect to original servers.
Load reduction for original content providers.
BUT…
Discovering near contents outside the ordinary paths towards the original servers requires the control plane to have a capillary knowledge of the contents in the neighborhood.
Contents in caches can change very frequently, thus leading to a considerable signaling traffic if we plan to keep track of their status.
A simple "Interest Flooding" strategy or an opportunistic discovery along the known paths towards the original servers could be inefficient.

Bloom Filter Basics

BLOOM FILTERS are data structures often used to efficiently perform membership queries on large data sets.

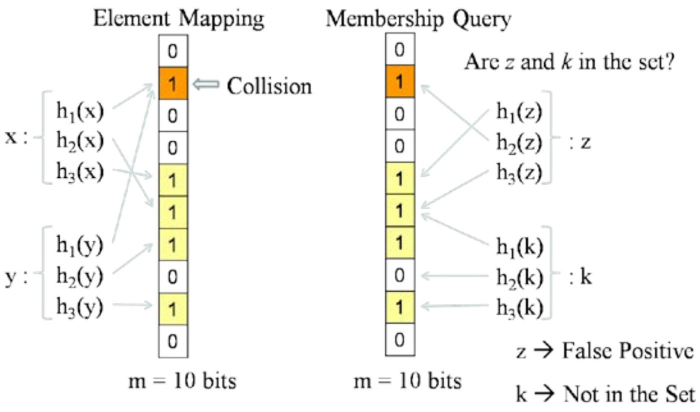

- **m**: filter length [bit].
- **n**: # of items in the considered universe.
- **k**: # of hash functions employed to map the items.

$P_{fp} = [1-(1-1/m)^{kn}]^k \approx (1-e^{-kn/m})^k$	False Positive Probability
$k_{opt} = \frac{m}{n}\ln 2$	Optimum number of Hash Functions
$P_{min} = (1/2)^k = 0.6185^{\frac{m}{n}}$	Minimum False Positive Probability
$m = -\frac{n\ln(P_{fp})}{(\ln 2)^2}$	Filter Length Under Optimum Hypotheses

Bloom Filter Basics

Bloom Filters could easily find their role in a **CCN Forwarding Engine!**

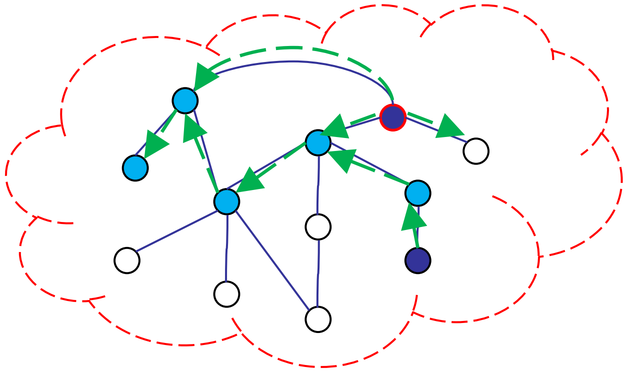

Nodes can exchange BFs representing contents in their caches, thus every one is aware of the contents in the neighborhood (multiple BFs with the same size can be merged with a simple OR operation).

Less Flooded Interests!

Compressed Routing Tables!

Fast Forwarding!

Bloom Filter Basics

FILTER LENGTH

100Gbps Ethernet ⇒ Forward an Ethernet frame in few ns!

We need to allocate BFs in fast on-chip memories like **TCAM** or **SRAM**.

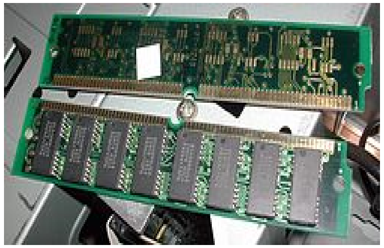

Max dimension for SRAM available nowadays: **~ 25 MB**

Scenario	Content Catalog (n)	BF Length (m) with Pfp = 0.1%
Web	10^{12}	~ 1.8 TB
UGC	10^{8}	~ 180 MB
File Sharing	10^{5}	~ 180 kB
VoD	10^{4}	~ 18 kB

$$m = -\frac{n \ln(P_{fp})}{(\ln 2)^2}$$

UGC = *User Generated Contents*

VoD = *Video on Demand*

Bloom Filter Basics

FILTER CONSISTENCY

The *false positive probability* of a BF in a CCN node is influenced by:

- The *load factor* of the filter (the proportion of 1s with respect to 0s);
- The *staleness* of the collected information.

Every node should periodically send its updated BF

Potential uncontrolled occupation of the links bandwidth!

How frequently these updated BFs should be sent?

Update Times Evaluation

We conducted simulations based on a model for general network of caches with a Least Recently Used (LRU) replacement policy *.

MODEL PARAMETERS	
$r_{i,v} = \lambda_{i,v} + \sum_{v':i \in R(v',v)} m_{i,v'}$	Incoming Request Rate for the i-*th* content at node *v*.
$m_{i,v} = r_{i,v} \cdot (1 - q_{i,v})$	Miss Rate for the i-*th* content at node *v*.
$q_{\mathrm{i,v}}$	Probability for the i-*th* content to be present in the cache of node *v* at a random point in time.
$\overrightarrow{q_v} = contents(\overrightarrow{p_v}, \lvert v \rvert)$	Function that models a single LRU cache.
$p_{i,v} = \frac{r_{i,v}}{\sum_{j=1}^{N} r_{j,v}}$	Relative portion of requests for the i-*th* content at node *v*.

$\lambda_{i,v}$ = Poisson stream of exogenous requests

* E.J Rosensweig, L. Kurose, and D. Towsley. *Approximate Models for General Cache Networks*. In *Proc*. of IEEE INFOCOM, San Diego, CA , USA, Mar. 2010

Update Times Evaluation

SIMULATION SETTINGS	
Topology	Random Graph with 50 nodes (v), 150 links and an average path length of 3.22 hops.
Content Catalog	10^4
Popularity Distribution	Zipf-like (α = 0.65)
Cache Size [contents]	500
Cache-to-Catalog Ratio	0.05
Requests Generation Rate	1 every second for each node
Request Interarrival Times	Exponential
Update Thresholds [% of cache size]	10, 20, 30

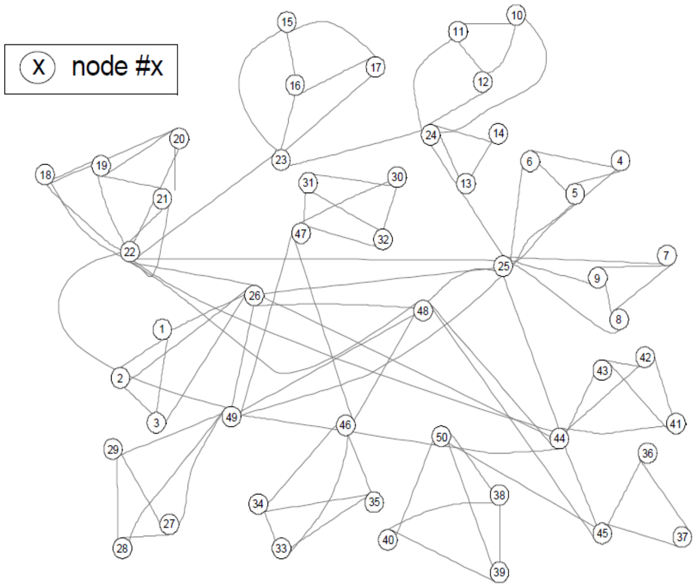

Update Times Evaluation

SIMULATION RESULTS			
Threshold	Min [s]	Avg [s]	Max [s]
10%	496	1004	1152
20%	630	1404	1512
30%	745	1692	1800

Obtained from the node with the highest degree (i.e. 12 interfaces)

We have short update times even in the presence of relaxed conditions:

- Requests are forwarded using only the shortest paths towards the original content servers (no flooding);
- The content catalog (10^4) is considerably smaller (10^{12} contents in a Web scenario);
- The Cache-to-Content Catalog ratio is bigger with respect to real scenarios.

The proposed Two-Level Forwarding Engine

OUR PURPOSE = Minimizing both signaling overhead and flooded Interests

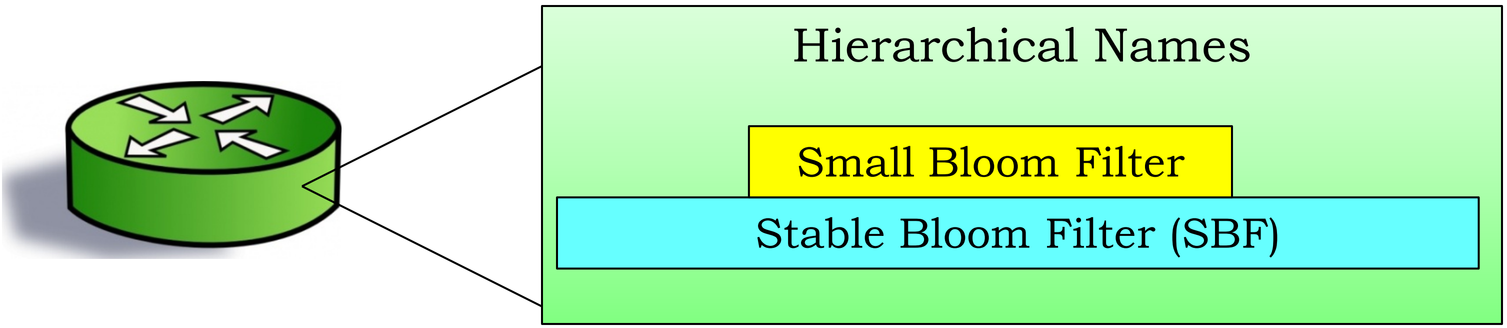

STABLE BLOOM FILTER *

- ✓ Same structure of a Counting Bloom Filter (*d* bit for every cell);
- ✓ Capacity to reinforce newest information as well as randomly delete stale ones (decrementing *P* cells).

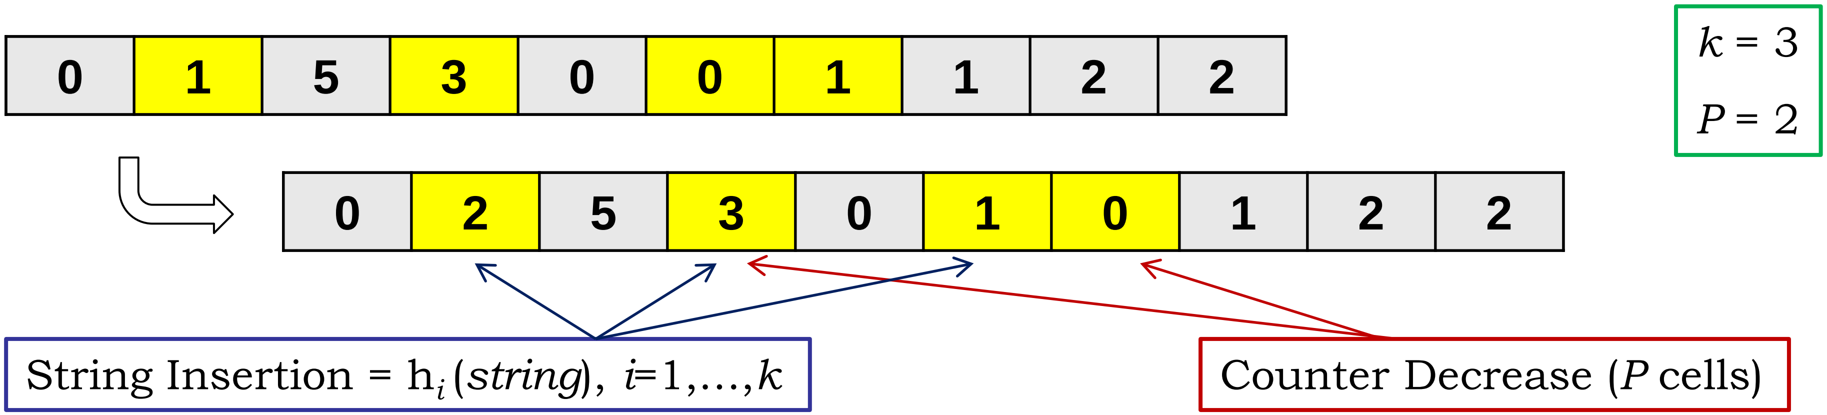

* F. Deng, D. Rafiei. *Approximately detecting duplicates for streaming data using stable bloom filters*. In *Proc*. of ACM SIGMOD, Chicago, IL, USA, 2006

The proposed Two-Level Forwarding Engine

OUR PURPOSE = Minimizing both signaling overhead and flooded Interests

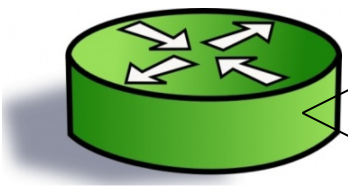

Hierarchical Names

Small Bloom Filter

Stable Bloom Filter (SBF)

STABLE BLOOM FILTER *

Its Usage…

youtube.com/video/Olympics/London_2012 → Node X, Interface 1

Dimensioning Example

n = 10^6 (ISP-scale scenario)

$P_{fp} = 0.1\%$

d = 3 bit (for every cell)

$$m = -\frac{n \ln(P_{fp})}{(\ln 2)^2}$$

⇒ $m \approx 5.4$ **MB**

NOT EXCHANGED

between nodes!

* F. Deng, D. Rafiei. *Approximately detecting duplicates for streaming data using stable bloom filters*. In *Proc*. of ACM SIGMOD, Chicago, IL, USA, 2006

The proposed Two-Level Forwarding Engine

OUR PURPOSE = Minimizing both signaling overhead and flooded Interests

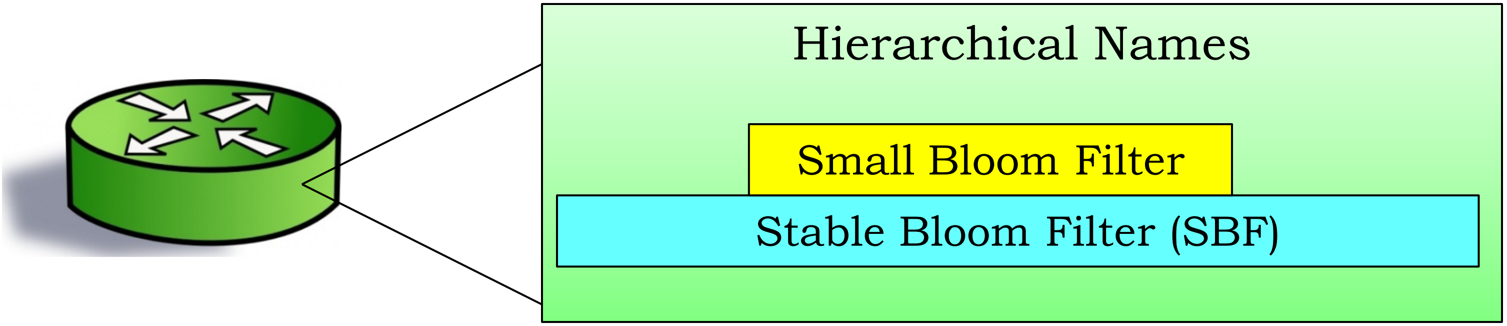

SMALL BLOOM FILTER

Hierarchical Names make easy:

- ✓ Names Aggregation;
- ✓ Regulation of name fields.

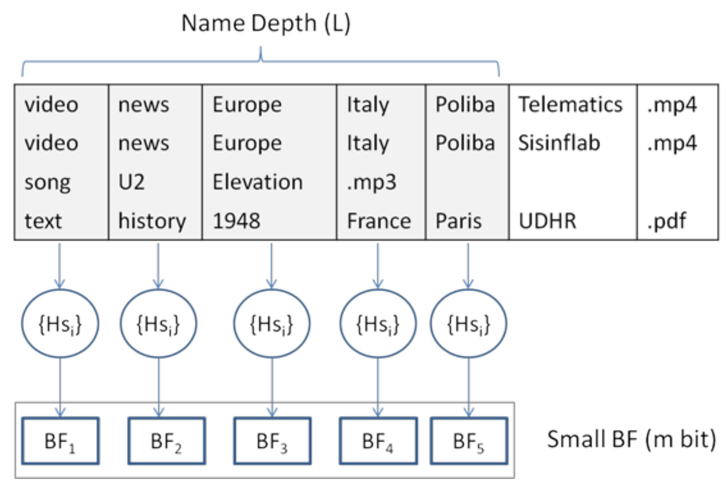

Separate BFs for each name field until a tunable name depth (L)…

The proposed Two-Level Forwarding Engine

OUR PURPOSE = Minimizing both signaling overhead and flooded Interests

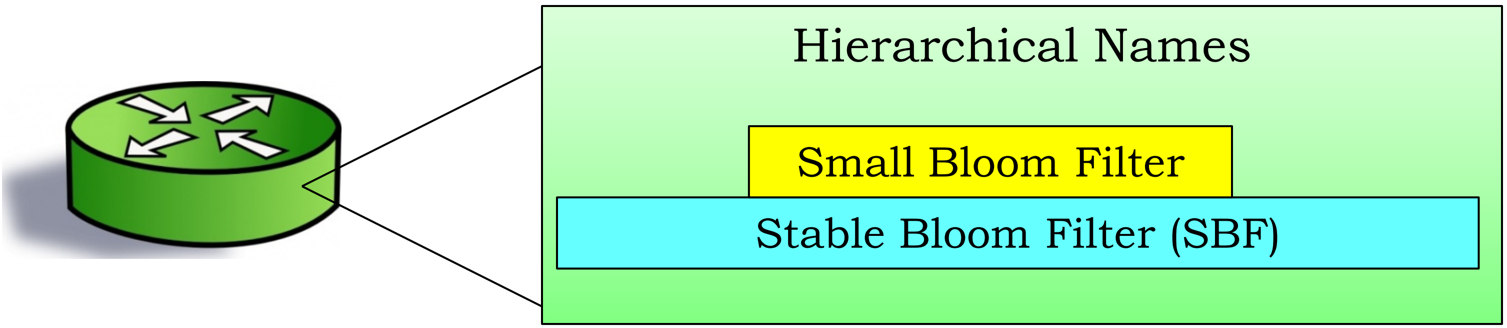

SMALL BLOOM FILTER

Dimensioning Issues…

The estimation of the catalog population for each field is not so trivial!

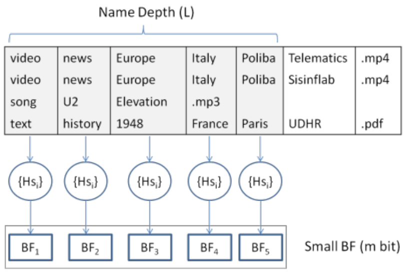

HINT

From the BNL* (Billiruglu & Neufeld List, 2007)

~ **2500 word families** cover the 90% of the English Language!

* N. Hancioglu et. al. *Through the looking glass and into the land of lexico-grammar*. English for Specific Purposes, Vol. 27, No. 4. (2008), pp. 459-479

The proposed Two-Level Forwarding Engine

OUR PURPOSE = Minimizing both signaling overhead and flooded Interests

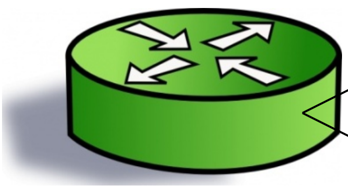

Hierarchical Names

Small Bloom Filter

Stable Bloom Filter (SBF)

SMALL BLOOM FILTER

Dimensioning Issues…

The estimation of the catalog population for each field is not so trivial!

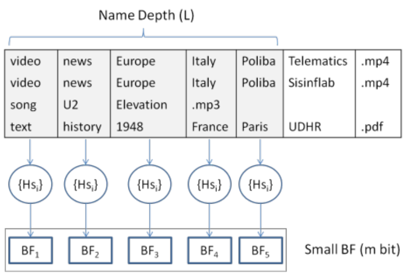

Example

n = 2500 (unique names)

$P_{fp} = 0.05\%$

$L = 10$

$$m = -\frac{n \ln(P_{fp})}{(\ln 2)^2}$$

$m \approx 50$ kB

EASLY EXCHANGABLE

between nodes!

LESS OVERHEAD!

The proposed Two-Level Forwarding Engine

NAMES AGGREGATION **MEANS** → INACCURACY

BUT…

- ✓ A greedy forwarding policy could be implemented even with a partial name match;
- ✓ We can compensate for this inaccuracy thanks to the complete information hold in the larger SBFs.

FORWARDING SCHEME

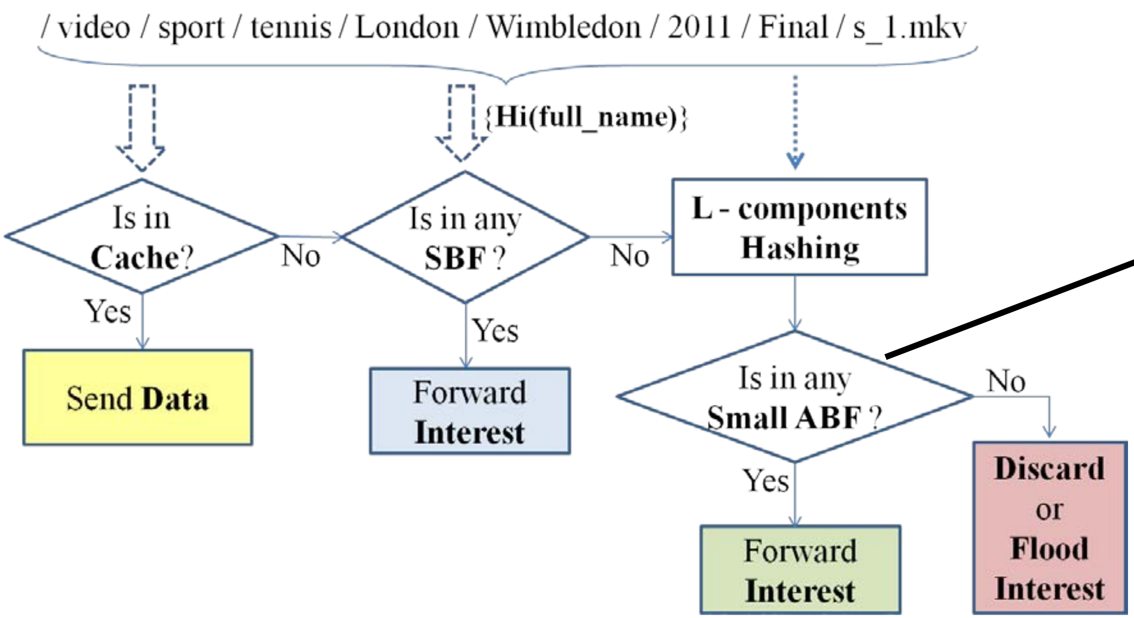

Membership Query

Keyword-based

OR

Consecutive

The road ahead …

Develop a customizable CCN simulator to highlight the pros and cons of the proposed forwarding strategy

Stress the attention on the scalability issues that could arise in the presence of a content catalog much greater than 10^6 contents, as well as on the comparison between hierarchical and flat label names

Evaluate the effects of the aggregation and propagation of BFs within different sized neighborhoods (in terms of number of hops from the nodes advertising their contents)

Refine and test the proposed Two-Level Forwarding Engine using a large scale test-bed (such as PlanetLab)

Thanks a lot for your kind attention !!!

Michele Tortelli

- **Telephone number:** +39 080 5963301
- **Fax:** +39 080 5963410
- **Web:** http://telematics.poliba.it/tortelli
- **E-mail:** m.tortelli@poliba.it
- **Personal e-mail:** tortellimichele@gmail.com
- **Skype contact:** michele.tortelli

Politecnico di Bari

via Orabona, 4 – 70125 Bari (Italy)