Information-Centric Networking

Seeing the Forest for the Trees

Structure of the Talk

- Background
- Commonalities of Designs
- Differences in Design
- Directions for research
- Conclusion

Background

- Surge of interest in ICN
- Literature difficult to absorb
- Aim is to give a broader view on ICN design
- Key point of the paper: the need for architectural research

Commonalities of Design

- Basic Primitives
- Universal Caching
- Content-oriented Security Model

Fundamental Differences in the Design

- Naming

  Hierarchical human-readable *vs* self-certifying

- Interdomain Name-Based routing

  On top of BGP *vs* separate named based routing

- Narrow waist

  Retain IP *vs* replace IP

Directions for research: Topics that deserve small attention

- Required ICN Mechanisms
- HTTP can be turned into an ICN design

Directions for research: ICN design (1/2)

- Privacy:
  - Content requested available to ICN nodes
- Interdomain Policies

Directions for research: ICN design (2/2)

- Scalable routing:
  - Should handle 10^{12} objects

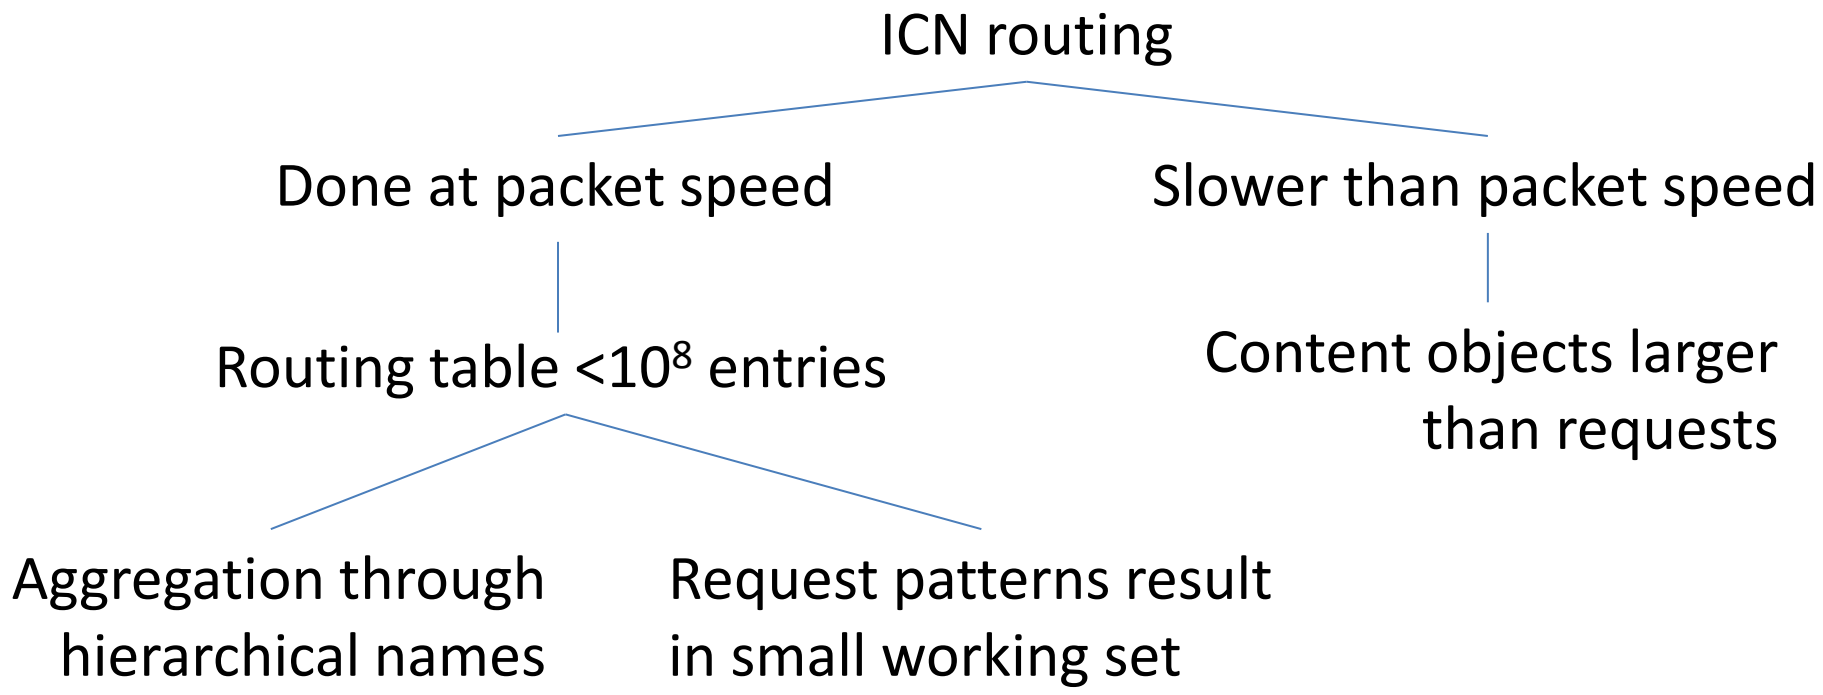

Directions for research: The performance gain of ICNs

- Empirical evidence
  - The fail of cooperative caching
  - The long tail of Content-Sharing Networks
  - Facebook image caching

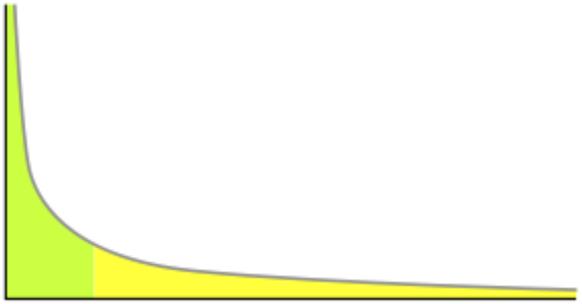

Zipf Distribution

- Efficiency of caches in the long tail

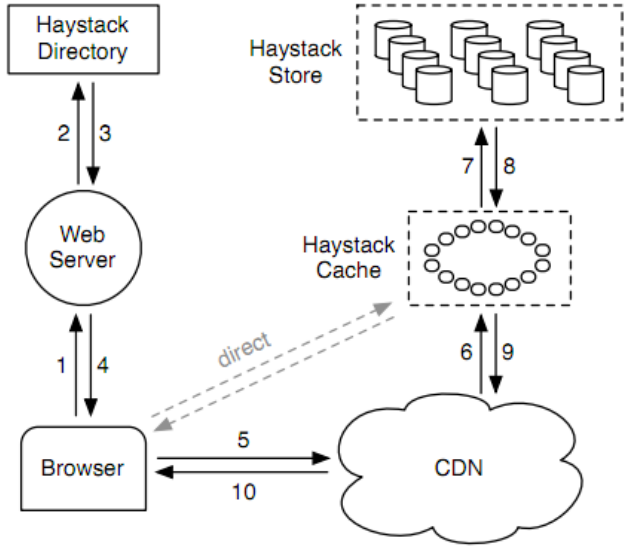

Conclusion

- Two fundamental questions:
  - What benefits do ICN designs offer?
  - Are ICN designs the best way to achieve those benefits?
- About the paper:
  - Asks the right questions
  - Does not always answer them objectively