A Resource-Constrained Asymmetric Redundancy Elimination Algorithm

Yu-Sian Li

June, 2013

Outline

- Introduction
- RCARE Algorithm
- Dynamic Adaptation Algorithm
- Conclusion and Future Work
- Achievements

What We Consider ?

- **Bandwidth asymmetric channels**

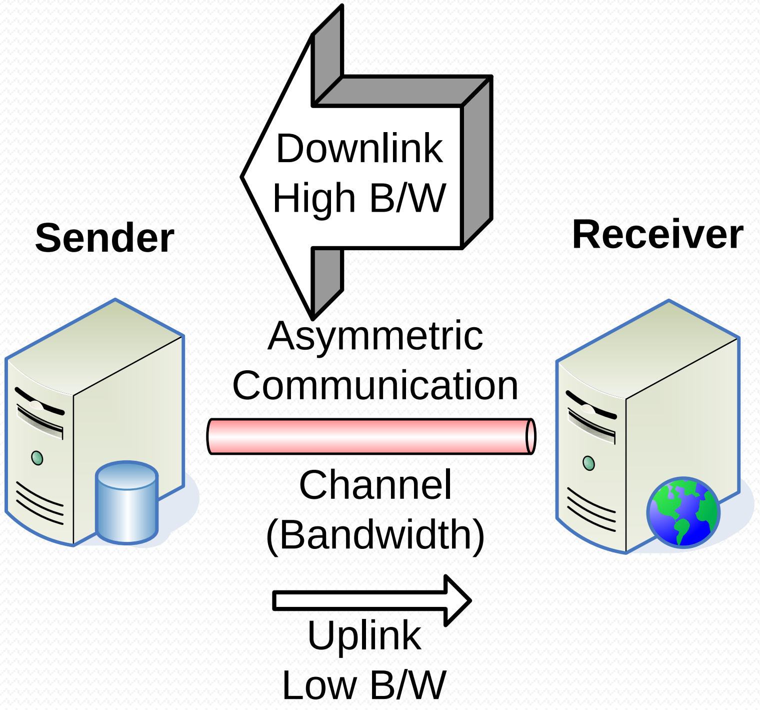

What We Consider ?

- **Capability asymmetric channels**

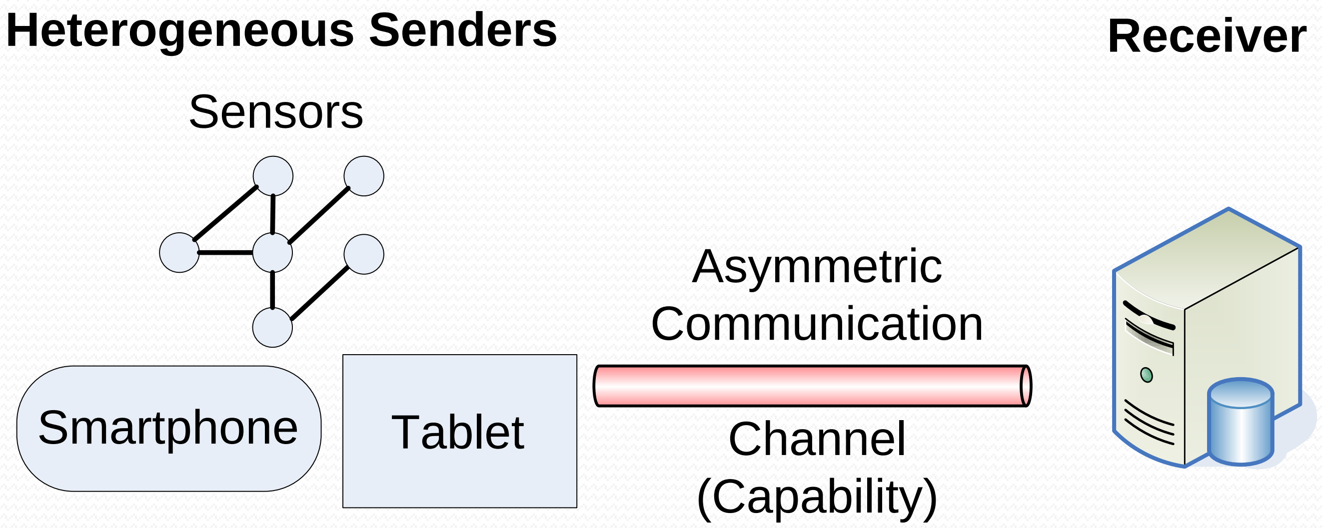

- Slower CPU
- Smaller memory
- Battery powered

- Faster CPU
- Larger memory
- Power line powered

Problem Statement

- We study the problem of increasing uplink goodput gain* in asymmetric communications
- Capitalizing the otherwise wasted downlink bandwidth and/or receiver capability.

* the relative goodput improvement compared to the standard TCP transfer

Existing Asymmetric Communication Algorithms

- ListQuery[1]
  - Dictionary: a list of seen messages, in non-increasing order by their frequencies
  - Guessed message: sublist of the $t^{1/k}$ most popular messages
  - Hints: the identification of the matched message in sublist
- Weakness
  - Cannot identify the partial-matches
  - High downlink cost

[1] T. Gagie, "Dynamic Asymmetric Communication," in ***Proc. of Information Processing Letters***, 2008

Previous Works

- Trang et al.
  - Implement ListQuery and DBES on NS-2[2]
  - Simulations on synthetic traces[2]
  - Implement EndRE[3] on NS-2

[2] C. Trang, X. Huang, and C. Hsu, "Pushing Uplink Goodput of An Asymmetric Access Network Beyond its Uplink Bandwidth," in ***Proc. of ICC'12****, Ottawa, Canada, June 2012*

[3] B. Aggarwal et al., "EndRE: An endsystem redundancy elimination service for enterprises," in ***Proc. of NSDI'10***, San Jose, CA, June 2010

Contributions

- Propose an asymmetric redundancy elimination algorithm RCARE
- The first redundancy elimination algorithm tailored for resource-constrained scenarios
- Study the correlation between unlink goodput gain and data stream features
- Design an adaptation algorithm for allocating the cache size

RCARE

- Sits on top of TCP
- Increases the uplink goodput from a sender to one or more receivers
- Flexible matching mechanism to identify partial-matches

RCARE Can be Deployed on

- **Two hosts**

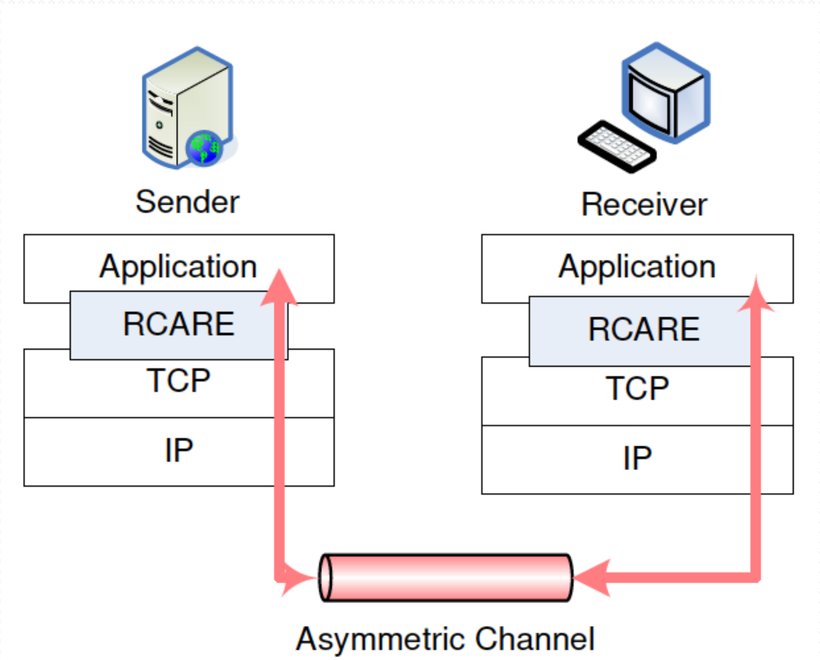

- **Network proxy**

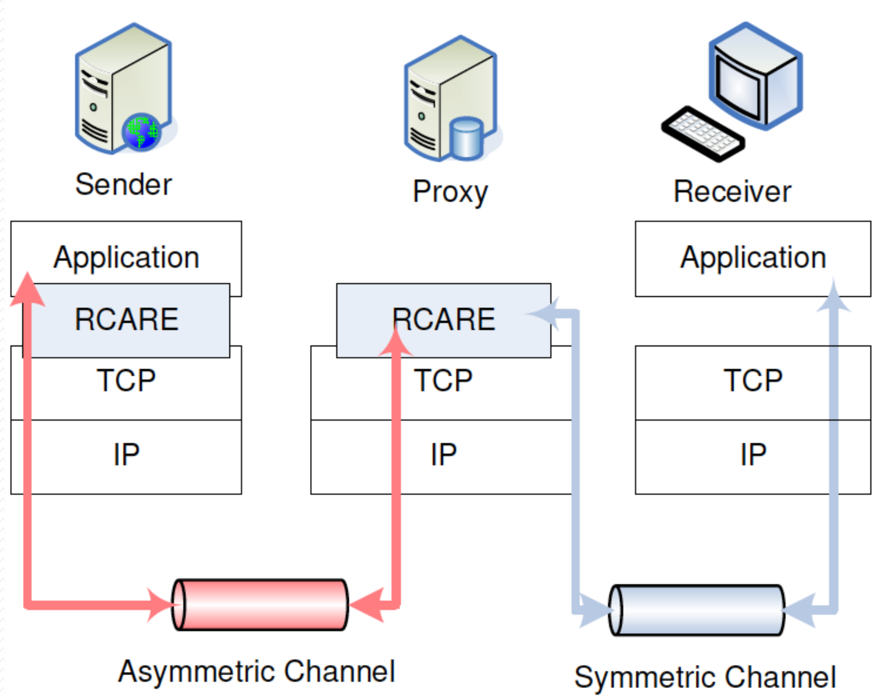

General Idea

- Asymmetric channel with low uplink bandwidth
- Weak capability on sender

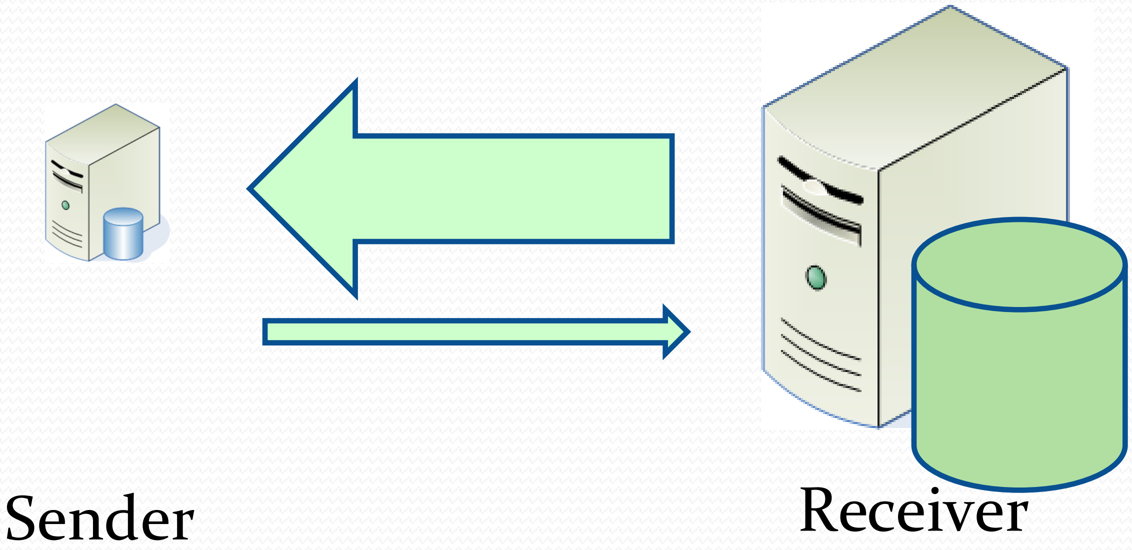

During the Transmitting

- Receiver records and analyzes incoming packets

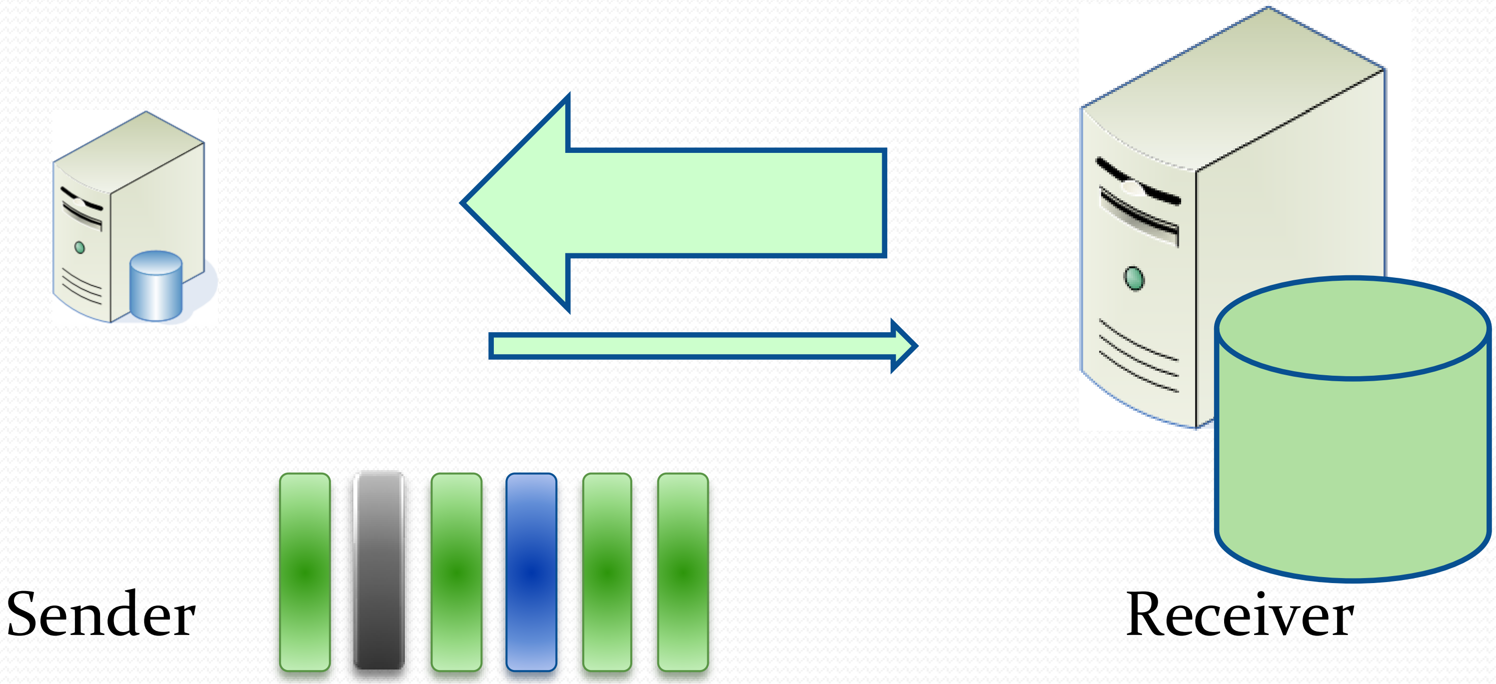

Update Sender Cache

- Receiver sends partial cache based on senders' capability

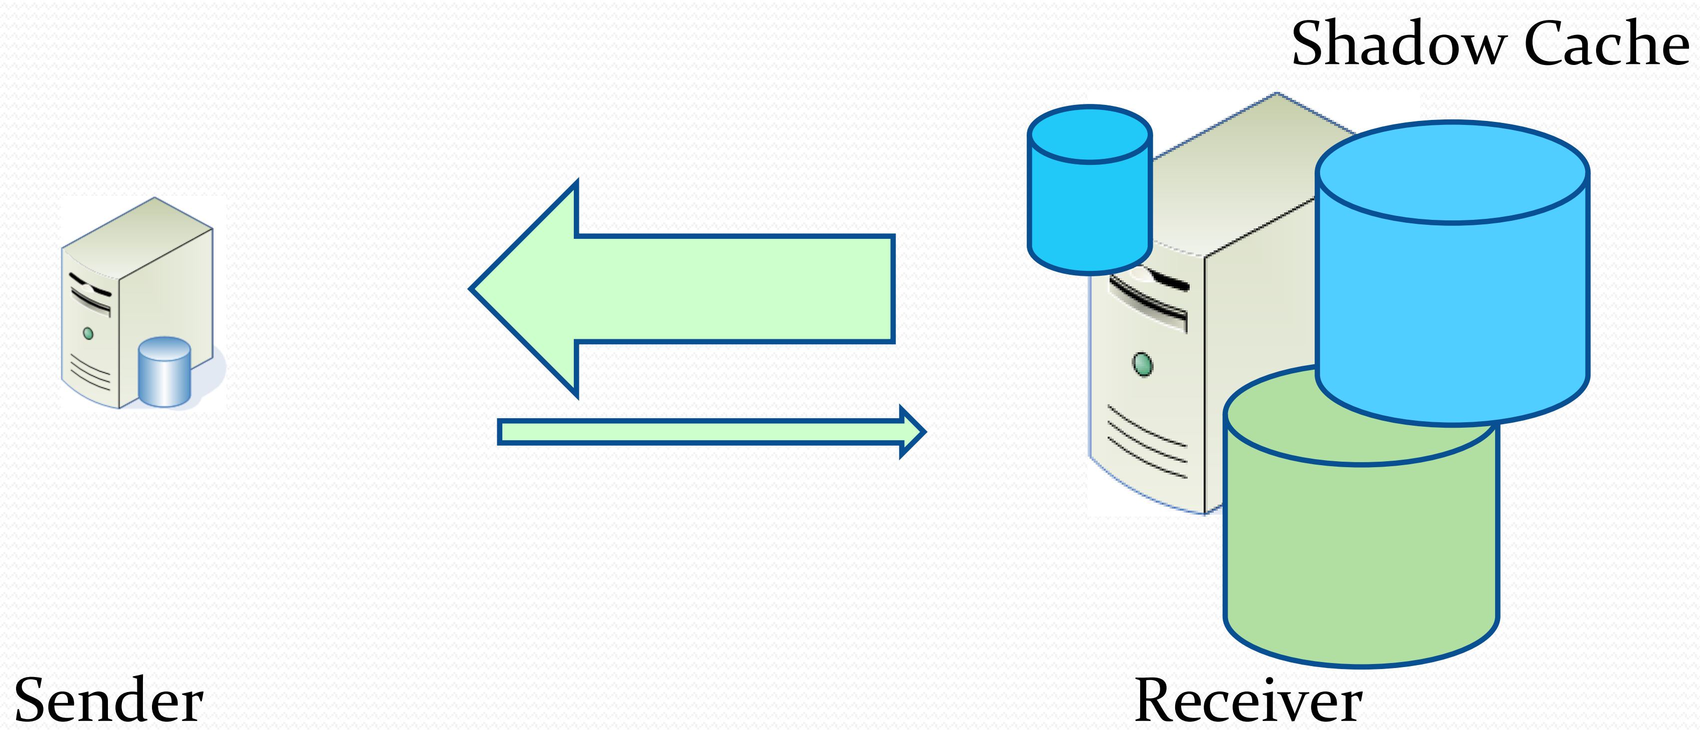

Compress Data

- Search cache, replace data stream with code word

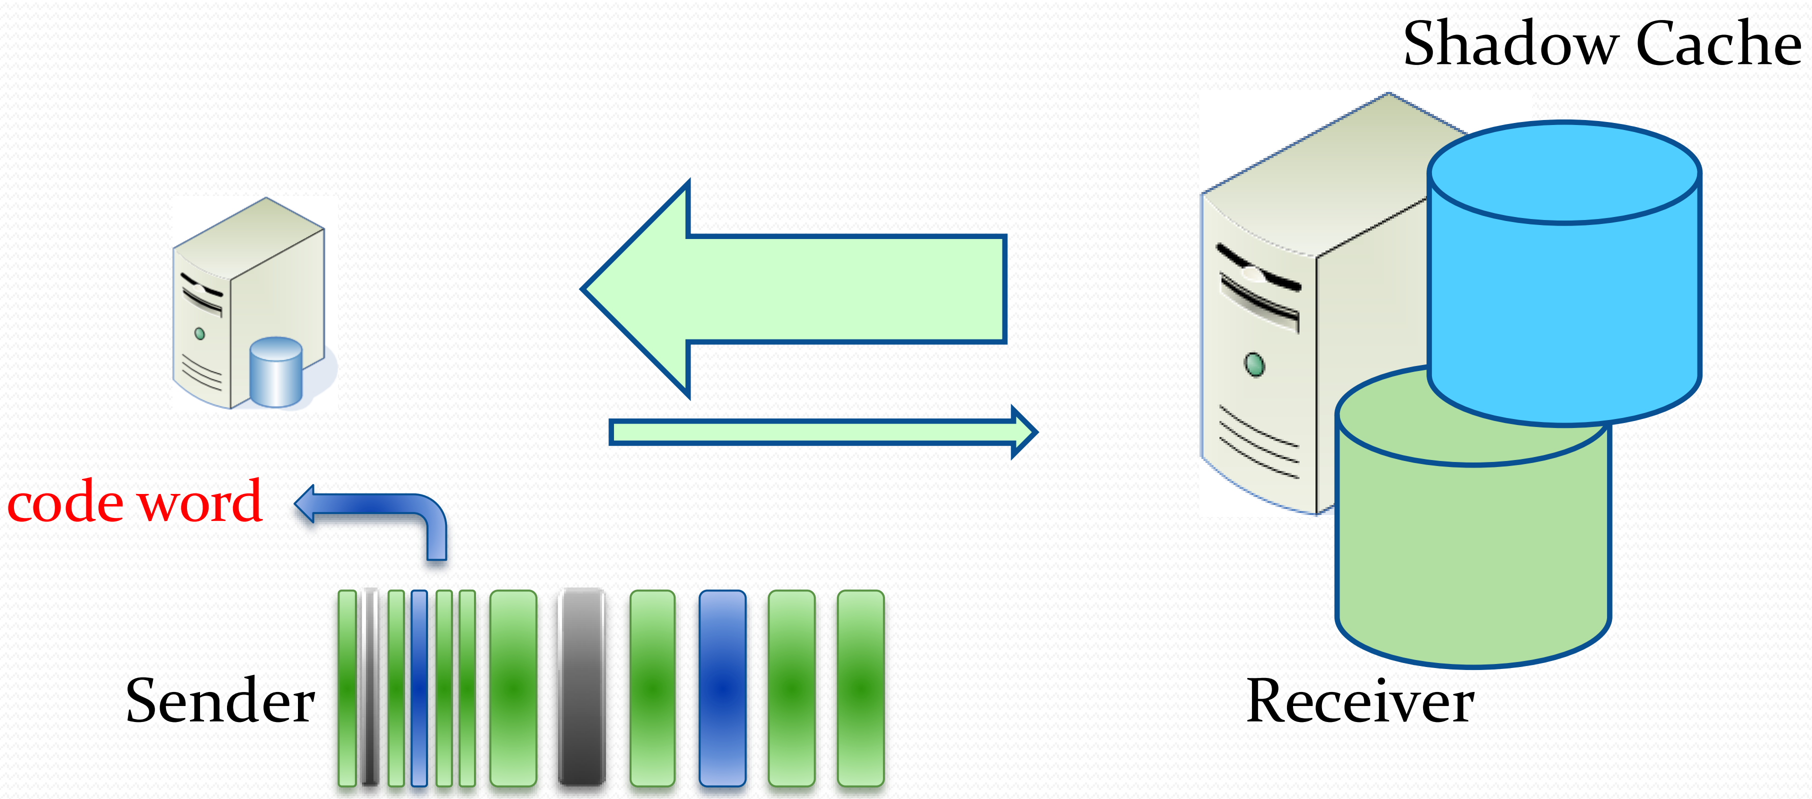

Partial Match Algorithm

- SAMPLEBYTE[3]
- 1. Find entrance byte using marker list
- 2. Compute the representative window for fingerprint
- 3. Store the fingerprint as key which point to the raw data (if we have not seen it before)

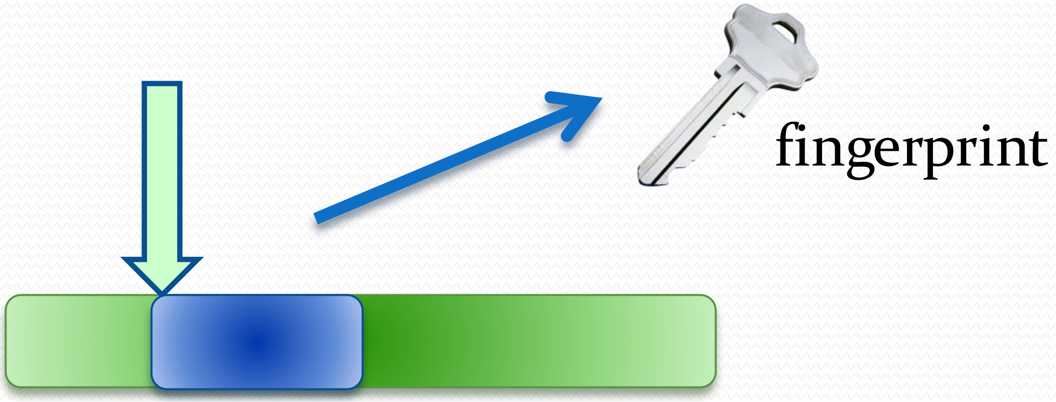

[3] B. Aggarwal et al., "EndRE: An endsystem redundancy elimination service for enterprises," in ***Proc. of NSDI'10***, San Jose, CA, June 2010

Partial Match Algorithm (cont.)

- Once the fingerprint exists in the cache
- 1. Try to extend matching stream
- 2. Replace it with <stream ID, offset, length>
- Very Efficient: encoding and decoding time is at most 0.5 and 5 msec on per packet

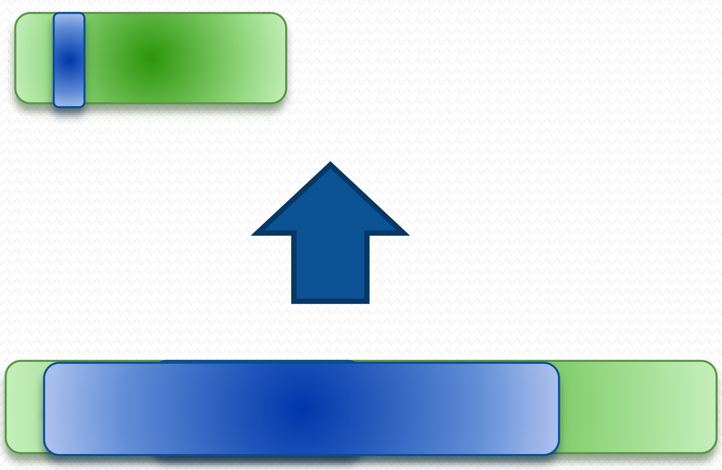

System Parameters

- B_s : Max cache size on sender
- B_r : Max cache size on receiver
- f : Updating frequency (number of packets)
- m : Marker list size [1, 256]
- β : Selection policy [0.0, 1.0] (β MRU and 1-β MFU)

Trace-Driven Simulation

- Trace file
  - Tcpdump with payload
  - 5 trace files (T1 ~ T5)
  - Enterprise, home, and university servers
  - Size from 59 MB to 1.2 GB

- Trace-driven simulator
  - RCARE
  - ListQuery
  - EndRE
  - GZip

Results from Different Algorithms

- Fix sender B_s=B_r=256MB and f = 1 for fair comparison
- RCARE outperforms ListQuery
- Similar to EndRE

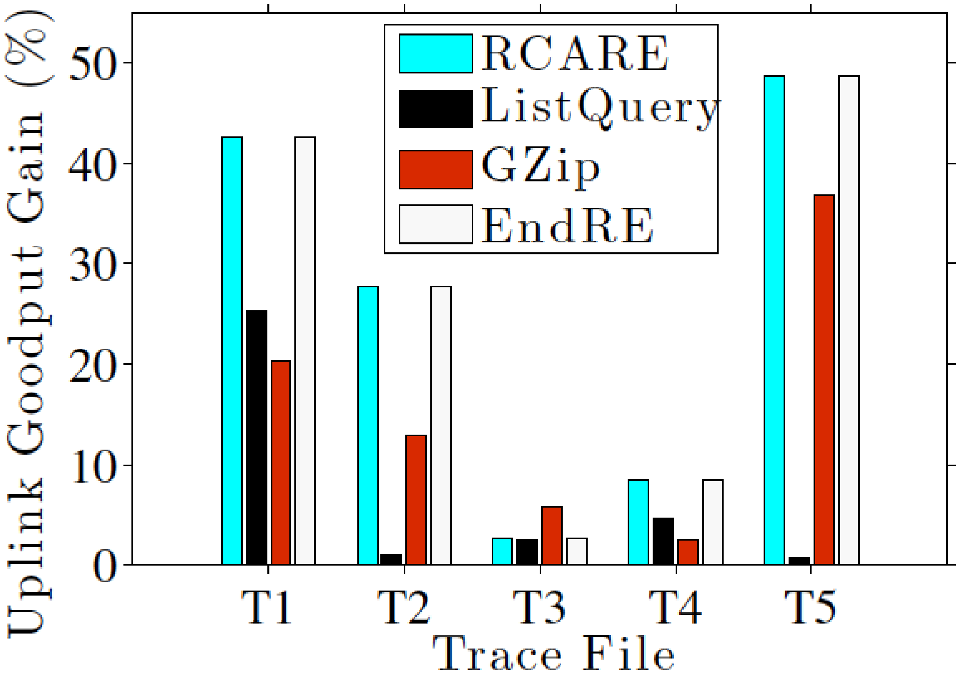

Simulation Results (cont.)

- Goodput gain with various receiver cache size.
- With a quarter cache size, we can achieve similar goodput gain.

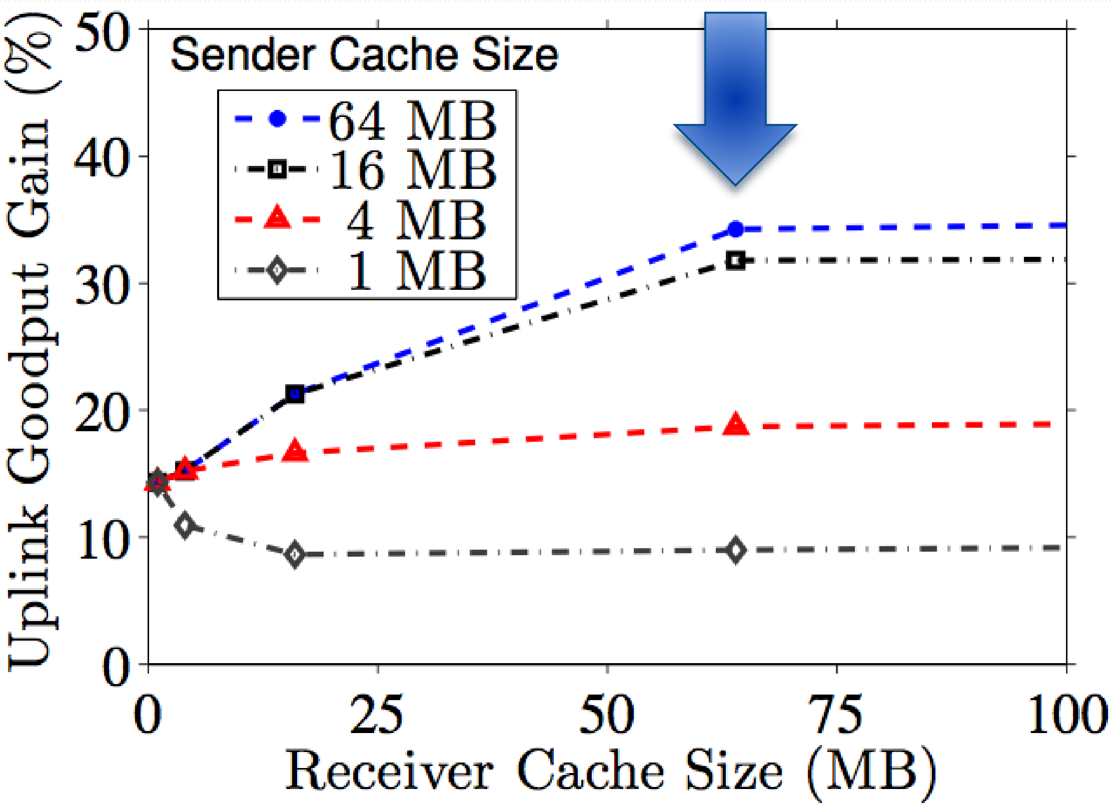

Dynamic Adaptation

- Resource-constrained hosts
- How to dynamically adapt cache sizes on individual data streams?

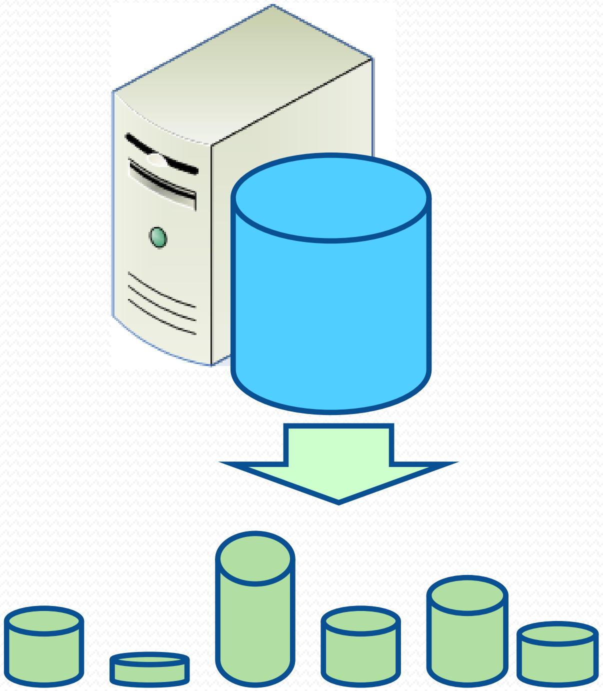

Real-Life Traffic Collection

- Dorm gateway in NCTU
- From 12:00 to 22:00 on February 20th, 2012
- 1,632 GB packet data in total
- 3,358 distinct IPs on the local network
- 3,598,829 distinct IPs from the Internet

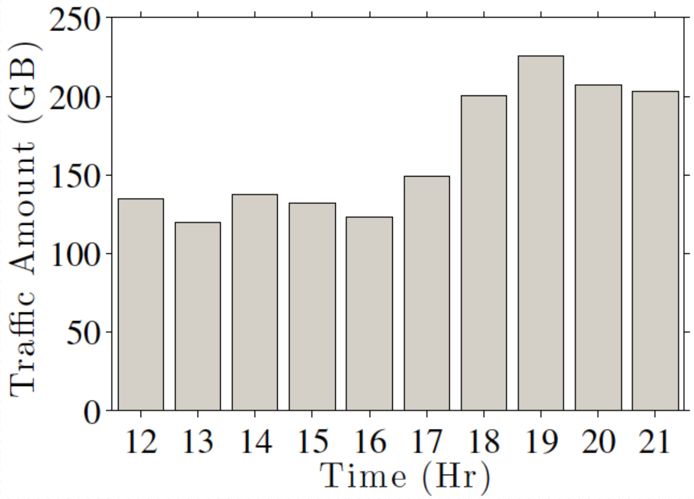

Data Stream Features

- Port number
  - Different protocols may inherently carry different amount of redundancy
  - Consider the source port
- ASCII ratio
  - ASCII data generally contains more redundancy compared to binary data
- Entropy
  - Compute the entropy of 32-byte long fixed-length data blocks
- Mean packet length
- Standard deviation of packet length

Data Stream Features - Analysis

- R^2 value of single-variable regression

Feature	Linear	Quadratic
Entropy H	0.74	0.85
ASCII ratio θ	0.08	0.67
Standard diviation of packet length	0.04	0.16
Mean packet length	0.28	0.29

- R^2 value of two-variable regression

Feature	Linear	Quadratic
Entropy H and ASCII ratio θ	0.39	0.59
Entropy H and Mean packet length	0.45	0.61
ASCII ratio θ and Mean packet length	0.39	0.48

- We use quadratic regression of entropy to build the prediction model

Prediction Model

- Use regression of entropy to predict goodput gain

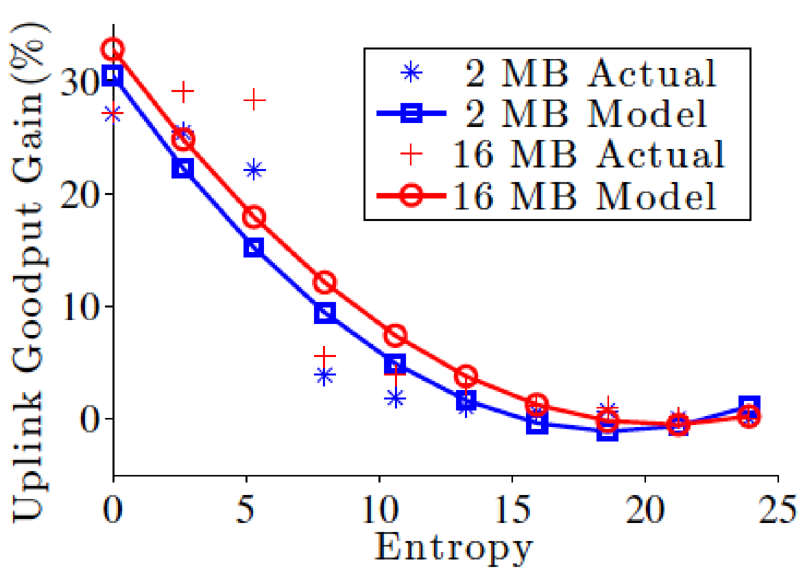

The prediction model closely follows the actual goodput gain.

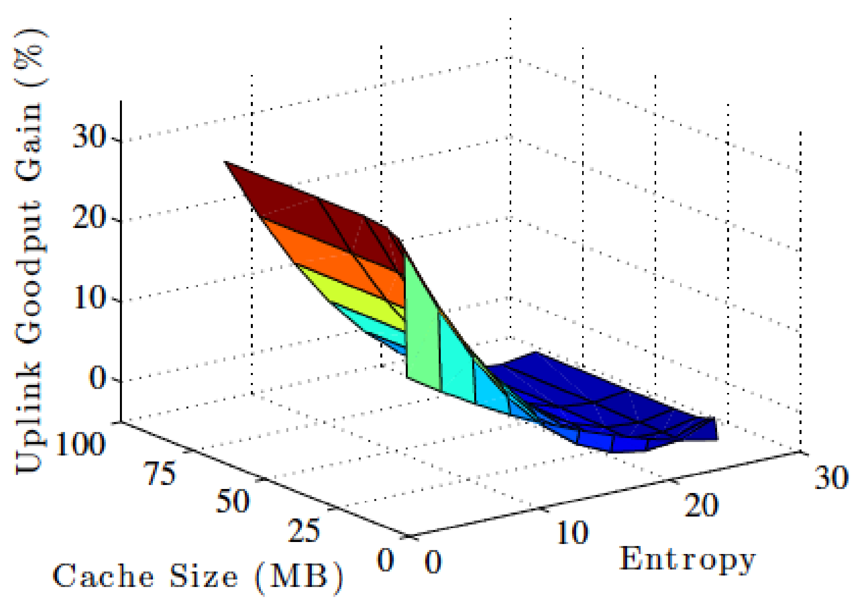

The interpolated surface of our proposed prediction model.

Adaptation Formulation

- The prediction model can be written as a piecewise linear function
- Formulate the optimization problem
  - Objective: Max goodput gain
  - Decision variable: stream cache size

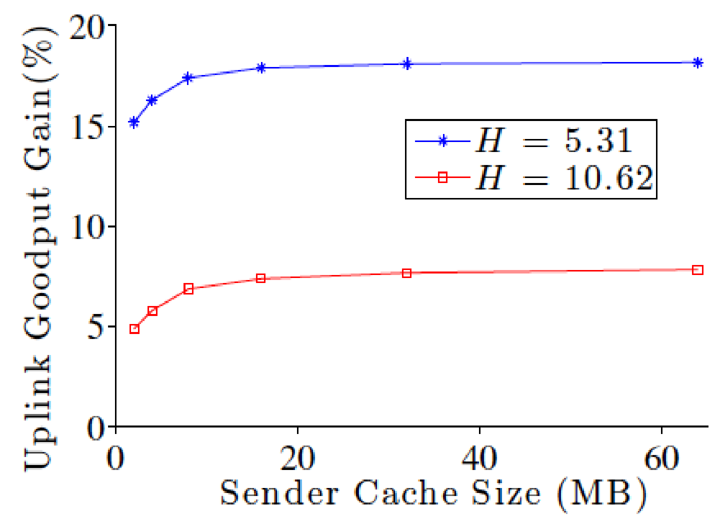

Adaptation Algorithm

- OPT
  - Solve the optimal problem with CPLEX
- EFF
  - Invests the remaining cache size on the data stream that is estimated to achieve the highest goodput gain
- AVG
  - Equally divides the total cache size to each data stream
  - Baseline

Evaluation on Adaptation

- Goodput gain (normalized to OPT)
  - OPT and EFF outperform AVG
  - EFF achieves very similar goodput than that of OPT

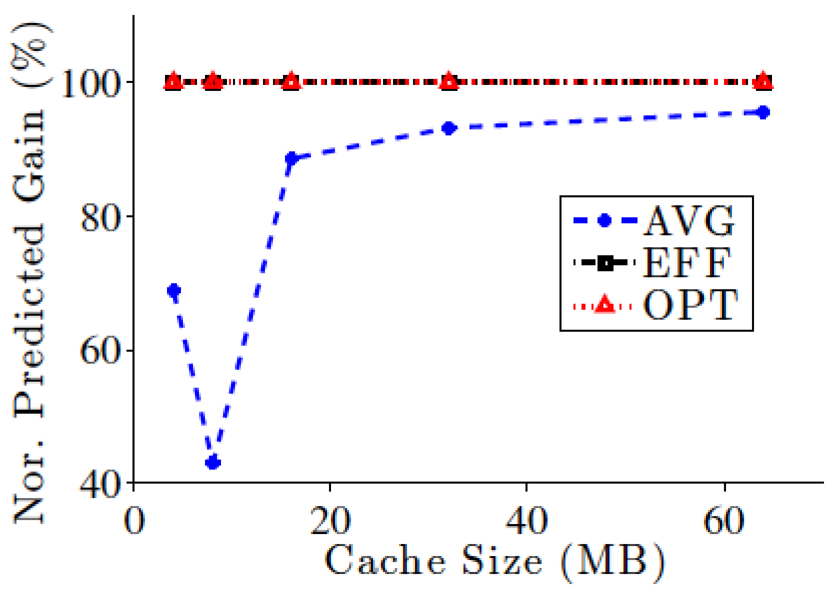

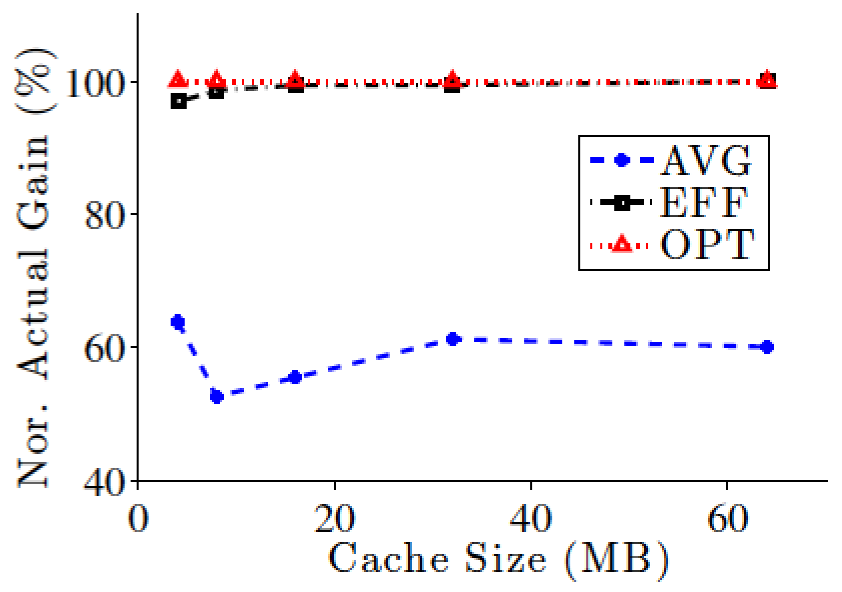

Evaluation on Adaptation (cont.)

- Overhead
  - EFF algorithm runs as fast as AVG
  - OPT may consume more than 20 MB memory

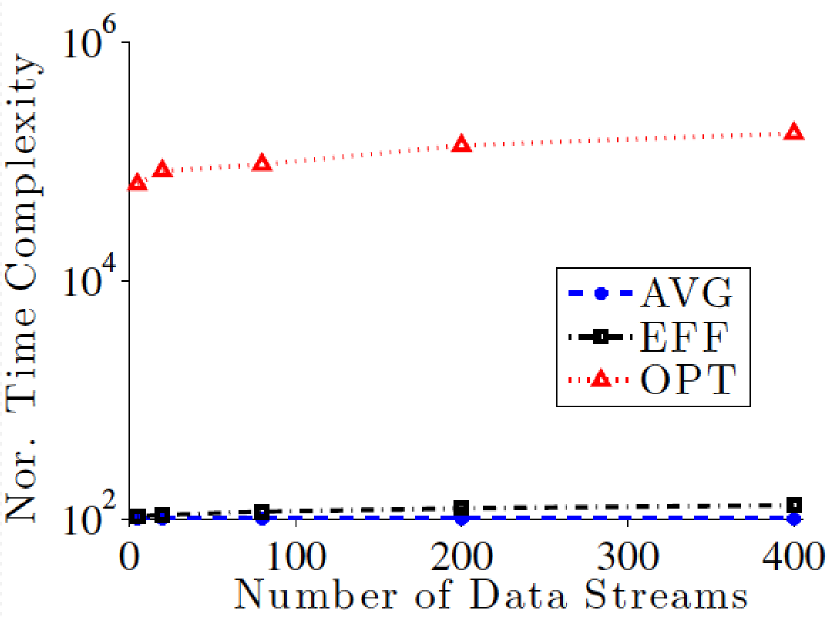

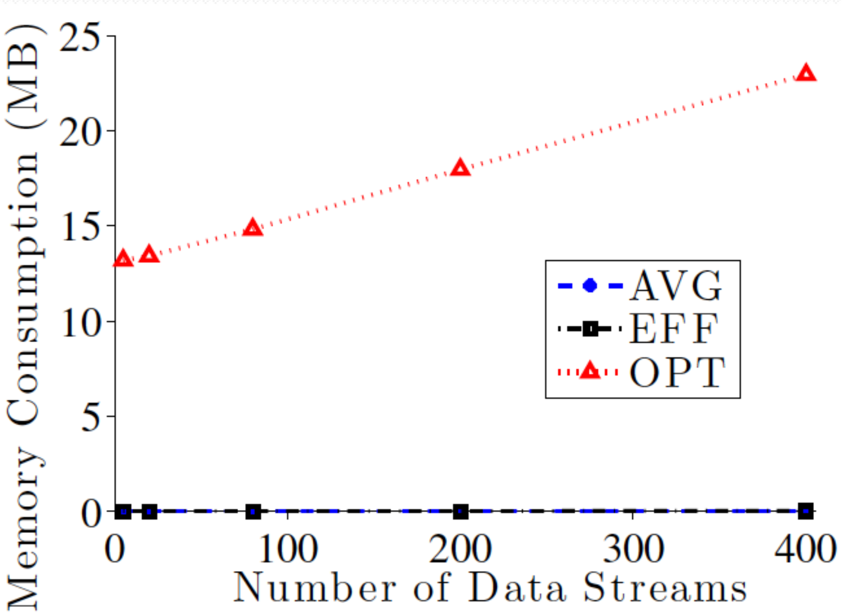

Performance Gain

- Deployment configurations

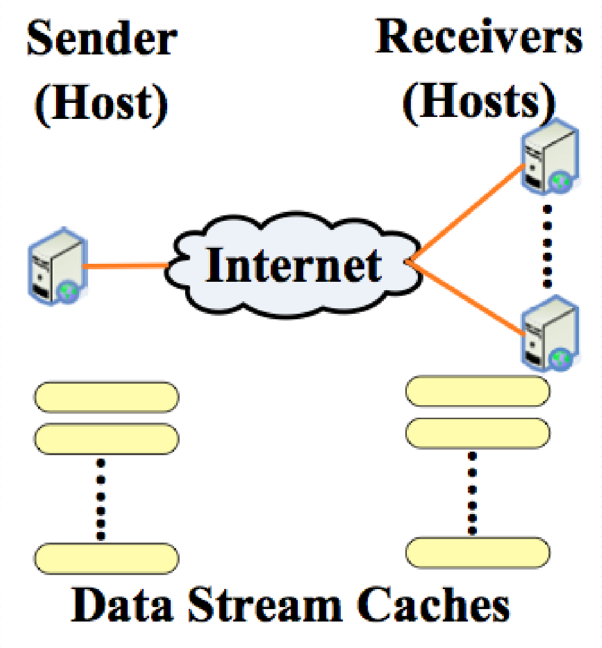

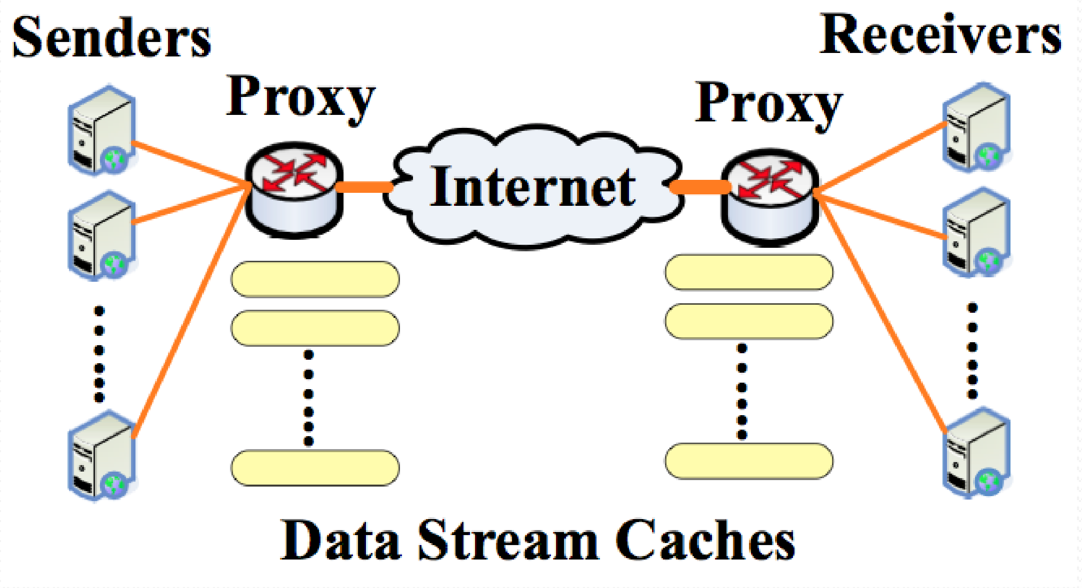

Host based **Proxy based**

Performance Gain – Host

- Consider the number of data streams ≥ 2 on each host
- There are 876 GB trace data and 1280 hosts in total
- EFF with 4 MB cache size achieves almost the same goodput gain of the AVG with 32 MB cache size
- EFF always outperforms AVG

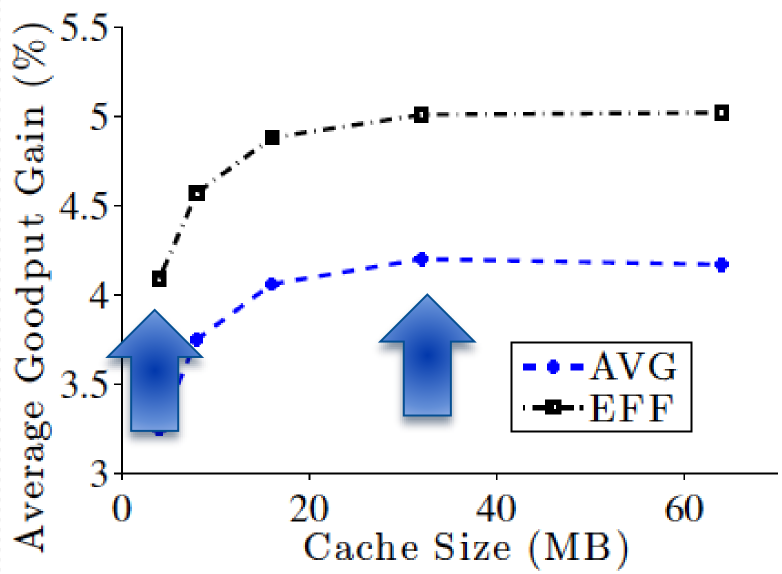

Performance Gain – Host (cont.)

- Zoom in to the first 10% hosts achieving the highest goodput gain
- EFF algorithm achieves over 40% goodput gain on average

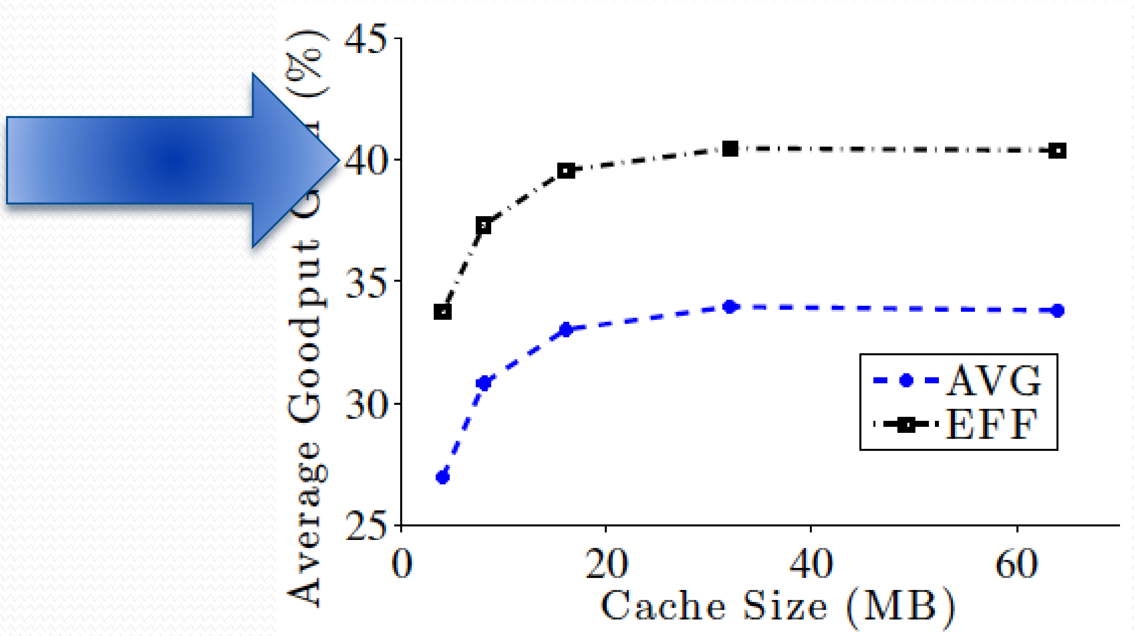

Performance Gain - Proxy

- Update frequency ***f*** = 10,000
- Sender cache size **Bs** ∈ {0 .25 , 0 .5 , 1 , 4 , 16} GB
- EFF is at least 10 times and at most 22 times higher than that of AVG

Algorithm	$B = 0.25\ GB$	$0.5\ GB$	$1\ GB$	$4\ GB$	$16\ GB$
AVG	0.13%	0.12%	0.12%	0.12%	0.13%
EFF	0.60%	1.12%	1.78%	2.66%	2.87%

Conclusions

- We proposed a new asymmetric communication algorithm, RCARE
- RCARE outperforms existing Asymmetric Communication Algorithms (ListQuery) by up to 50 times goodput and reduces 384 times downlink traffic amount
- RCARE is flexible on cache size adaptation
- Our adaptation algorithm improves up to 87% uplink goodput gain compared to a baseline

Future Work

- Inter-sender redundancy elimination
  - Share caches among all clients
- User behavioral patterns
  - Employ multiple versions of cache and use them on different days/hours
- Implement RCARE in a real network stack and conduct experiments

Achievements

- Y. Li, C. Trang, X. Huang, C. Hsu, and P. Lin, "CacheQuery: A practical asymmetric communication algorithm," in **Proc. of IEEE Global Communications Conference (GLOBECOM'12)**, Anaheim, CA, December, 2012, **TAOS Best Paper Award**.
- Y. Li, C. Trang, S. Wang, X. Huang, C. Hsu, and P. Lin, "A Resource-Constrained Asymmetric Redundancy Elimination Algorithm," **IEEE/ACM Transactions on Networking**, in preparation.

- Y. Li, C. Chen, T. Lin, C. Hsu, Y. Wang, and X. Liu, "An End-to-End Testbed for Scalable Video Streaming to Mobile Devices over HTTP," in ***Proc. of IEEE International Conference on Multimedia and Expo (ICME'13)***, San Jose, CA, July, 2013.

Q & A