

The Implications of Network-level Packet Caching

Ashok Anand*, Athula Balachandran†, Aditya Akella* and Srinivasan Seshan†
*UW-Madison and †CMU

1. MOTIVATION

Internet traffic is increasing at a tremendous rate. Some studies predict that the net traffic might grow 5-fold by 2013 [1]. Our recently proposed framework of deploying packet level redundancy elimination (RE) on network elements (e.g., routers) ([2], [3]) can serve as an effective way to improve network efficiency without expensive upgrades. These proposals expand the scope of point deployments (e.g., access links of stub network) of “WAN Optimizers” (Riverbed, Peribit, etc.) to network-wide settings.

In the SmartRE design [3], network elements keep track of recently forwarded packets in a cache, strip redundant content from packets and encode them on the fly by comparing against the recently forwarded packets. The downstream network elements decode or reconstruct full packets by looking up their cache in a coordinated fashion. A centralized controller allocates the responsibility of encoding, decoding and caching across network elements, while minimizing the network-wide footprint of duplicate traffic thereby improve network-wide utilization.

In this work, we demonstrate the benefits of network-level packet caching using *RE-enabled* software routers (based on Click). We take as example the scenario of an on-demand video service, and show on a GENI slice that network-wide RE can significantly reduce the network-wide utilization.

2. DEMONSTRATION

Our demonstration (Figure 1) presents the benefits of packet level RE within a GENI slice, with video servers hosted on Wisconsin, and video clients at another site, labeled UCSD/BBN; these nodes should ideally be located in campuses that have deployed OpenFlow-enabled switches (the demonstration set-up can be expanded across multiple other sliceable nodes across other campuses, if available). The slice connects *RE-enabled* software routers, and is remotely controlled by a NOX-based controller that runs on a separate PC.

In our demonstration, we show that network-level RE empowers the network to accommodate more users by reducing the network wide link utilizations. The users can watch the same video at different times. Regardless of the workload imposed, RE will dynamically absorb

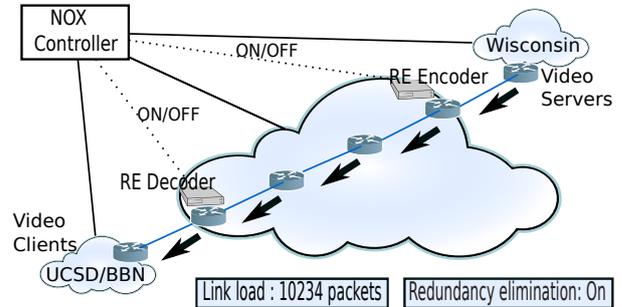


Figure 1: The setup is shown. NoX controller controls enabling/disabling of RE and setting up path between client and server. Using this, we show the benefit of RE. The frontend GUI shows the link loads.

the additional bandwidth demand by eliminating redundant content. We compare the performance against a baseline network with no RE. To highlight the benefits, we provide a simple GUI to show the load on network links.

We implement the encoding and decoding operations on software routers. The packets are stored on DRAM, and a hashtable is maintained to compare a packet against stored packets. We use Rabin fingerprinting to generate packet hashes.

We propose to leverage OpenFlow’s functionality to keep track of traffic demand, sharing of video across users, and network topology. The NoX controller collects these data from OpenFlow switches, and allocates encoding, caching and decoding operations on software routers. It also computes routes traversing RE-enabled software routers. In this demo, our clients are located at a single site, so route computation is not needed. For future demo, we will have clients located at multiple sites and we will also use redundancy aware routes to show further benefits.

3. REFERENCES

- [1] Cisco forecast. http://blogs.cisco.com/sp/comments/cisco_visual_networking_index_forecast_annual_update/.
- [2] A. Anand, A. Gupta, A. Akella, S. Seshan, and

S. Shenker. Packet caches on routers: the implications of universal redundant traffic elimination. In *SIGCOMM*, 2008.

- [3] A. Anand, V. Sekar, and A. Akella. Smartre: an architecture for coordinated network-wide redundancy elimination. In *SIGCOMM*, 2009.