A^3: Application-Aware Acceleration for Wireless Data Networks

Zhenyun Zhuang and Tae-Young Chang

GNAN Research Group, Georgia Tech, Atlanta, GA

Raghupathy Sivakumar and Aravind Velayutham

Asankya Networks, Inc., Atlanta, GA

Introduction

- Wireless Environments
  - High loss rate
  - Large delay
  - Low bandwidth
- Works focused on developing better transport protocols for wireless environments
  - TCP-ELN, WTCP, STP, etc.
  - Novel design, deals with unique characteristics
  - Improves throughput significantly
- Evaluation of transport protocols
  - FTP, or bulk of data

Analysis of Enterprise Traffic

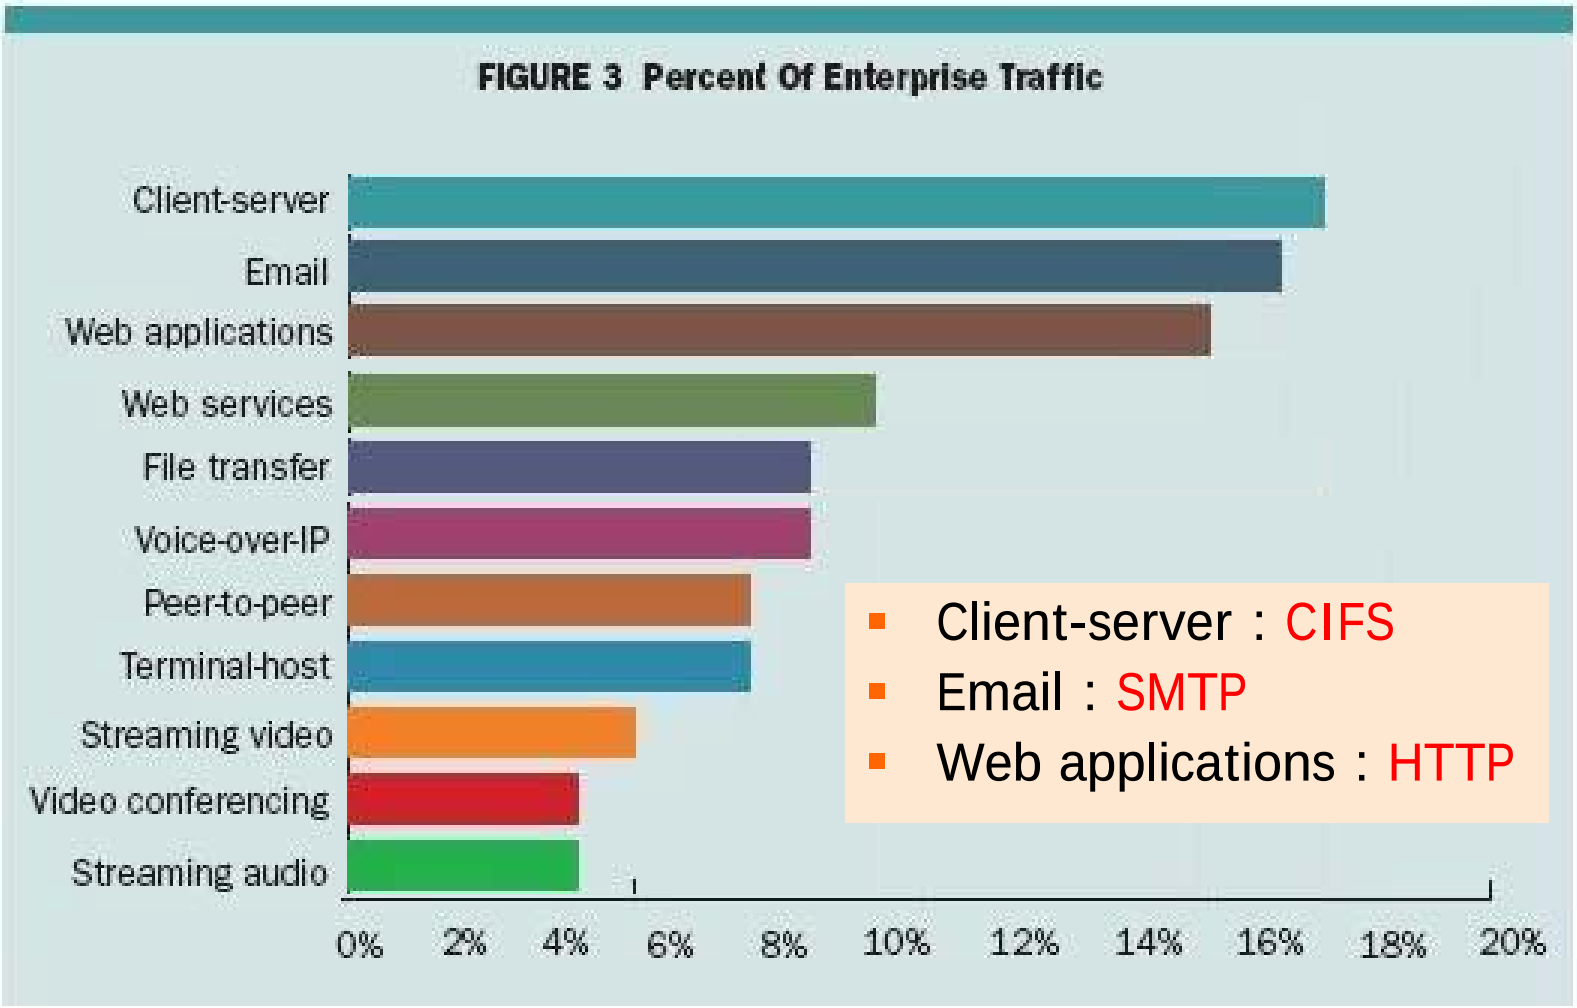

- Client-server : CIFS
- Email : SMTP
- Web applications : HTTP

Figure from Business Communications Review (April 2006)

Motivation: Network Setup

- Traffic Generator
  - Ixia IxChariot (IxProfile, Console, Endpoint)
- Applications
  - FTP, CIFS, SMTP and HTTP
- Wireless Networks
  - Wireless LAN (WLAN)
  - Wireless WAN (WWAN)
  - Satellite Networks (SAT)
- Transport Protocols
  - NewReno, TCP-ELN, WTCP, STP
- Parameters
  - Varying loss rate (RTT, BDP)

Motivation Results: FTP

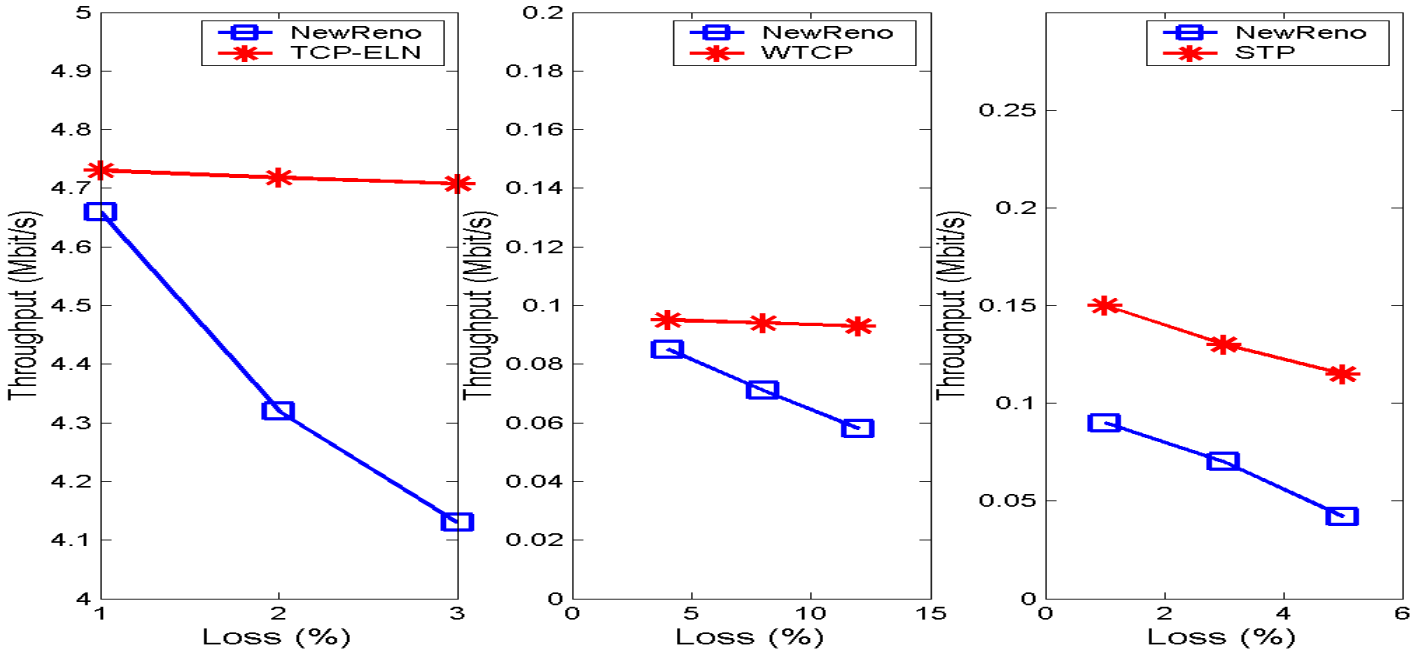

WLAN **WWAN** **SAT**

- Significant performance improvement...
  - Up to 120% in satellite networks

Motivation Results: Other Applications

- Less performance improvement in other applications...
  - 5%!

- CIFS

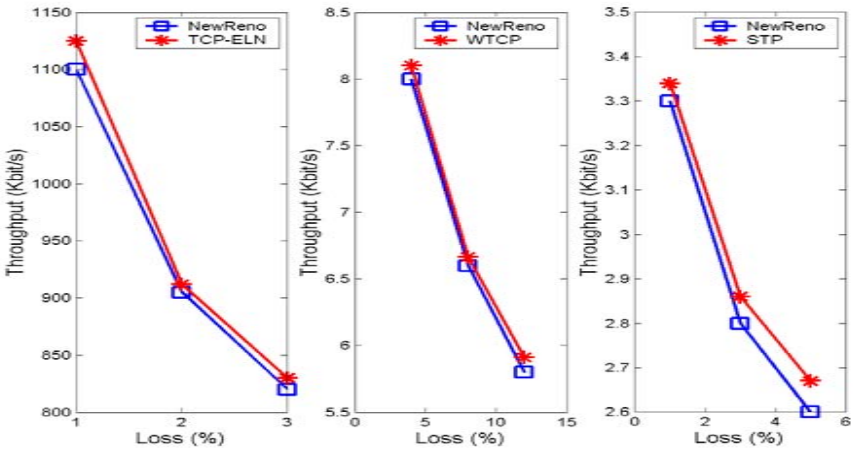

- SMTP

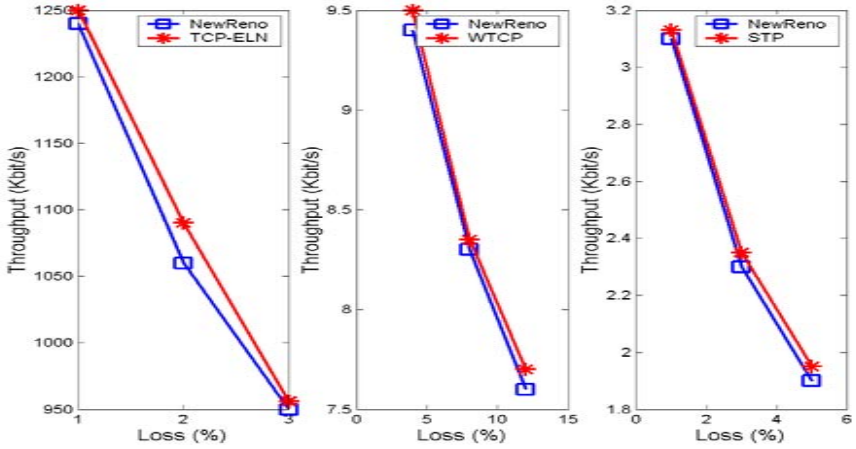

- HTTP

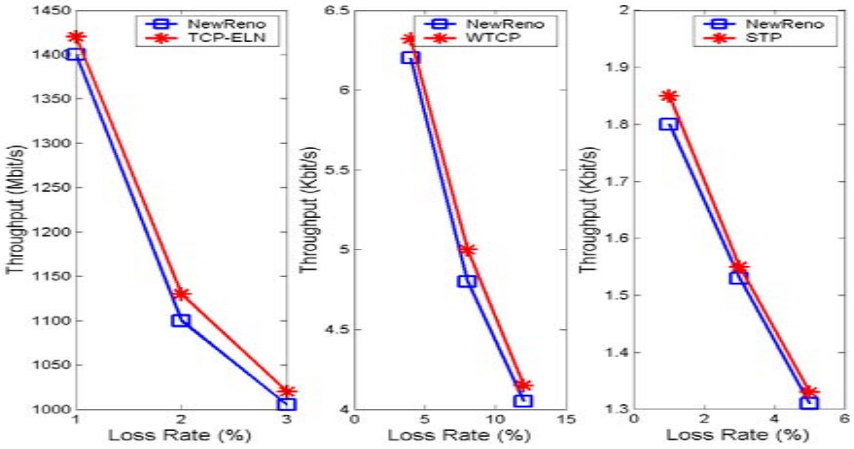

Application Traffic Patterns

FTP

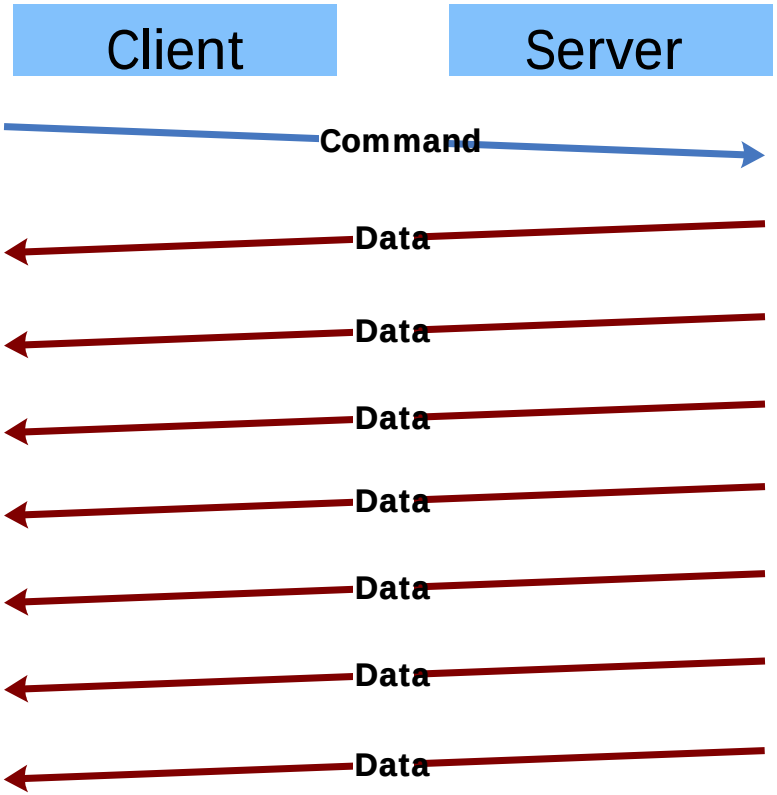

CIFS

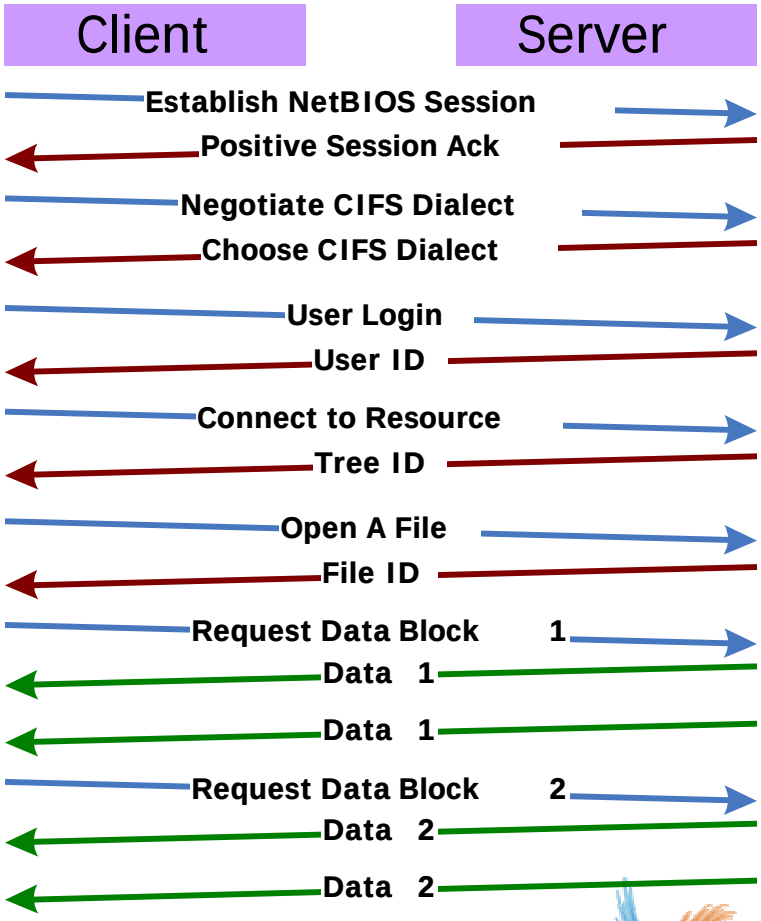

Application Traffic Patterns (cont.)

SMTP

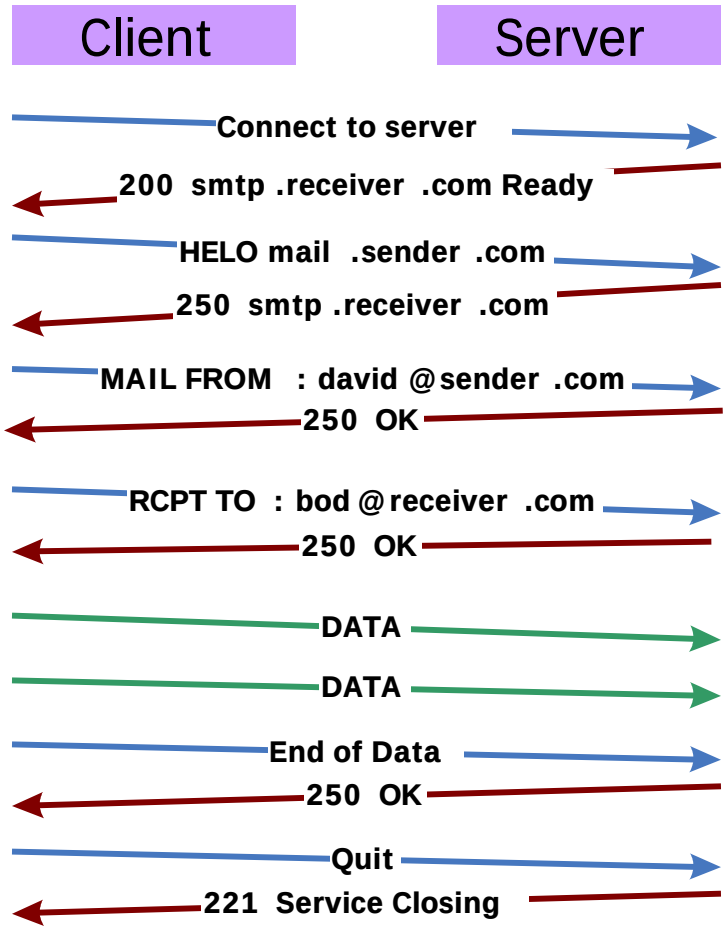

HTTP

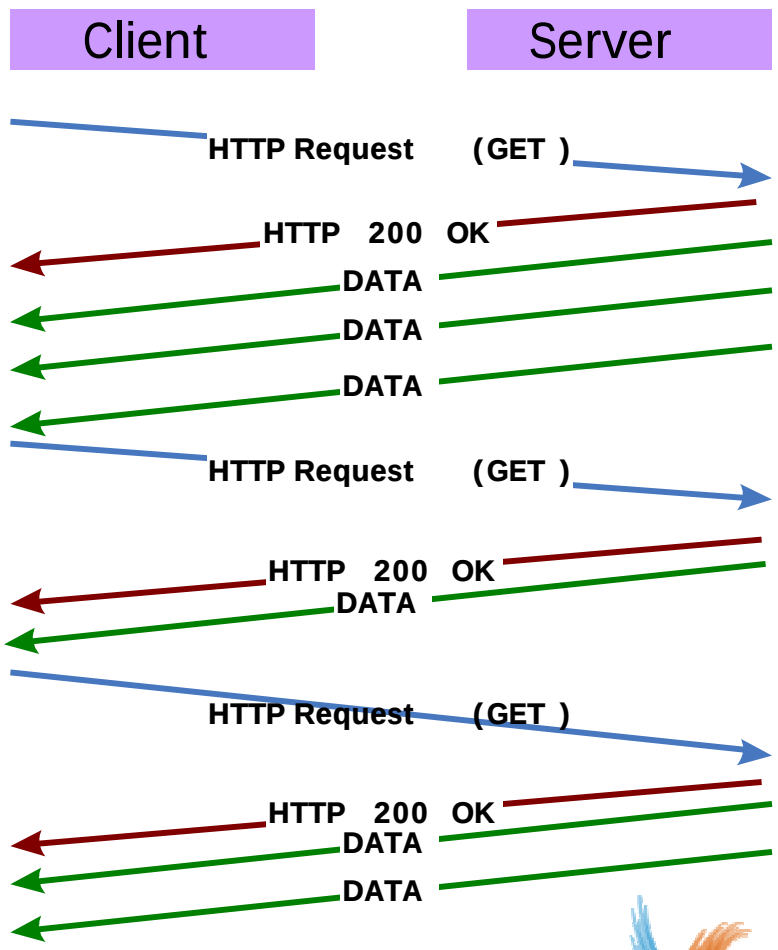

Typical Application Behaviors

- Thin Session Control Messages
  - Sent before DATA; Small
  - Retransmission timer expires to recover loss
- Batched Data Fetch
  - Data transfer is performed in batches
  - Bandwidth Delay Product cannot be fully utilized
- Flow Control Bottlenecked Operations
  - When applications are slow, receive buffer fills up
  - Flow control can kick in and be the bottleneck
- Non-prioritization of Data
  - Data are given equal importance
- Non-use of Data-reduction Techniques
  - Application-specific and user-specific information

A^3: Application-Aware Acceleration

- Application aware
  - Recognize applications
- Application transparent
  - No modifications to applications
- A set of design elements
  - Transaction Prediction (TP)
  - Redundant and Aggressive Retransmissions (RAR)
  - Prioritized Fetching (PF)
  - Infinite Buffering (IB)
  - Application-aware Encoding (AE)

A³: Transaction Prediction (TP)

- Transaction Prediction
  - Deterministically predict future requests
  - Issue them ahead of time
  - Designed for protocols that divide data into blocks
    - Examples, CIFS, HTTP

CIFS: Throughput

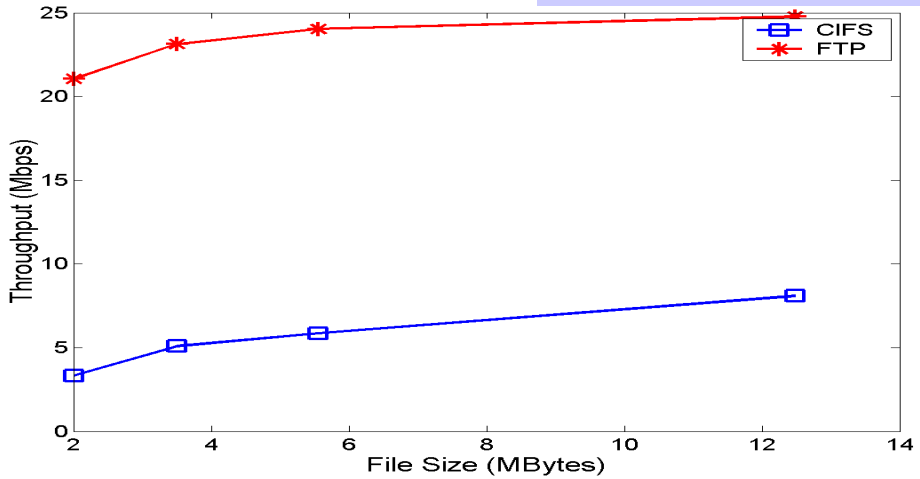

CIFS: Number of Requests

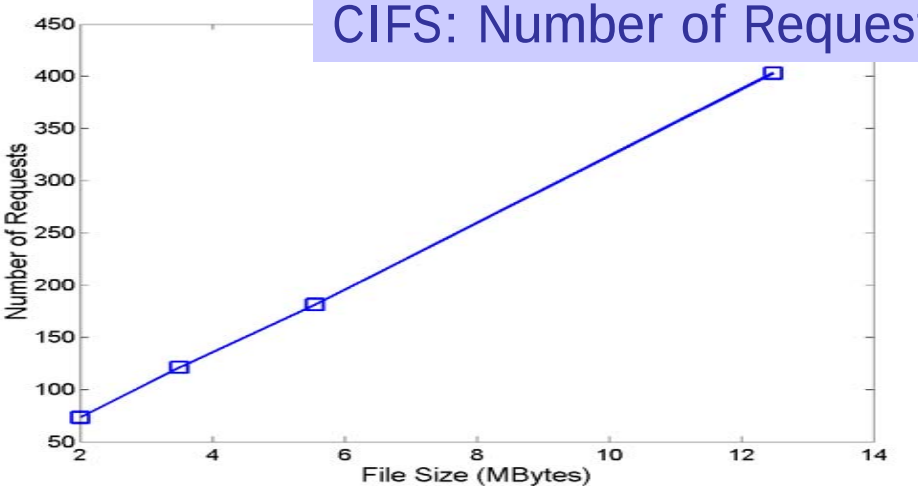

Redundant and Aggressive Retransmission (RAR)

- Helps protect thin session control messages from losses
- Packet-level redundancy
- Aggressive retransmission
- Not applying to DATA
  - Loss recovery is masked by subsequent packets
  - High overhead

SMTP Throughput

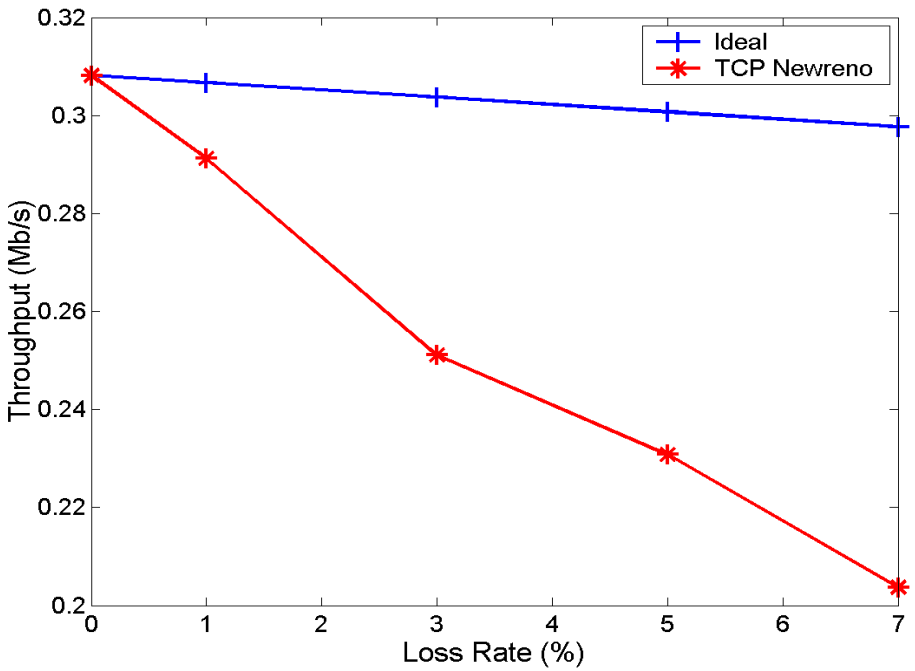

Prioritized Fetching (PF)

- Divide data into categories of different priorities
- Fetch them with different speeds
- Helps protocols that treat data with equal importance
  - Example, HTTP

Web page data vs Screen

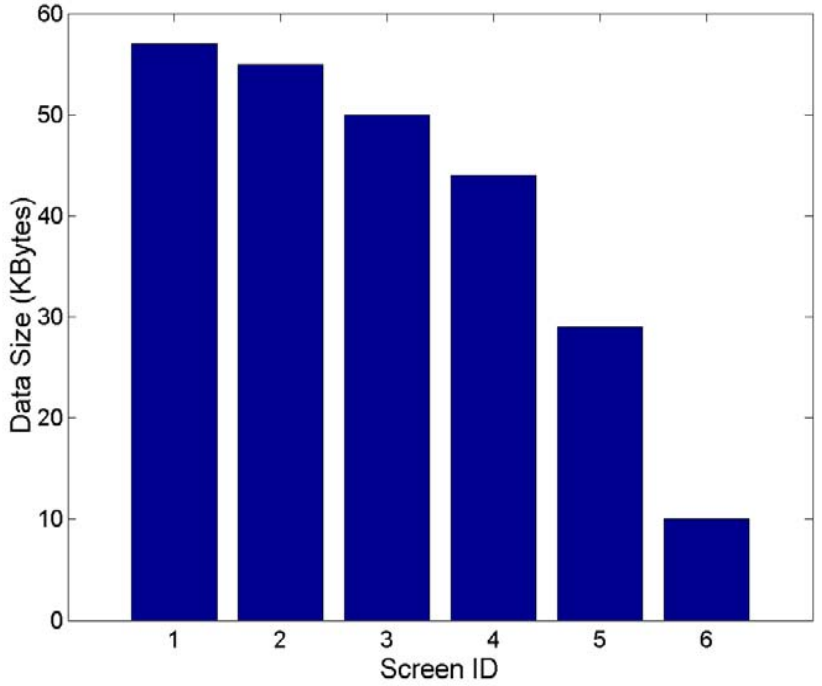

Infinite Buffering (IB)

- Prevents flow control from throttling the transmissions
- Uses local storage to store data at the receiver
- Flow control never kicks in

CIFS: Throughput vs Application Rate

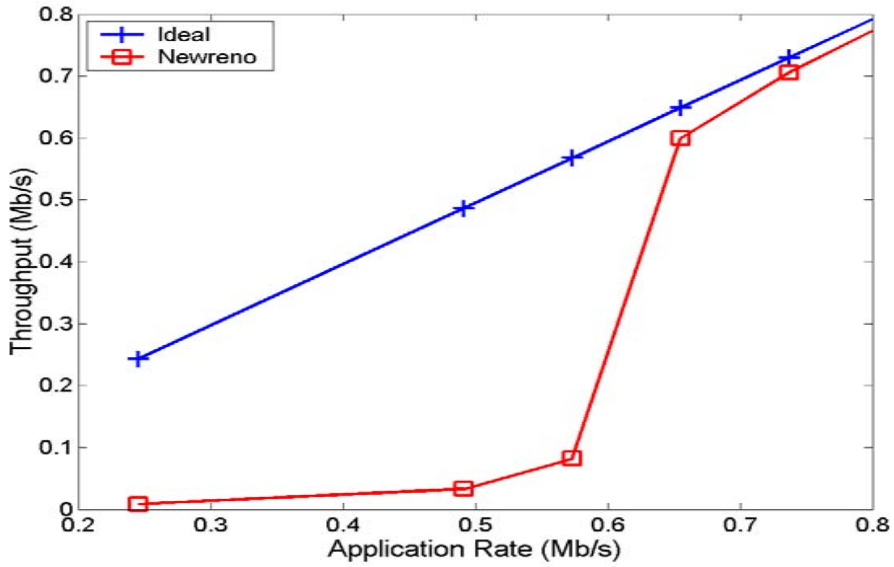

Application-aware Encoding (AE)

- Uses application and user specific information
- Better compress data
- Example, SMTP

10 Persons (100 emails)

ID	Unique Word	Total Word	Char. per Word	Bits per Email of Binary Coding	Bits per Email of Simple Coding
1	1362	6383	6.22	3176	664.6
2	3554	19284	7.12	10984	2274.6
3	2645	12653	7.08	7167	1438.5
4	4536	25481	6.15	12537	3095.2
5	966	4728	11.46	4335	468.8
6	1205	6413	5.48	2811	656.4
7	798	3346	4.40	1178	322.6
8	1527	6836	5.72	3128	723.0
9	1758	9171	4.91	3602	988.6
10	1402	8320	7.3	4859	869.7

Word Frequencies of a Person

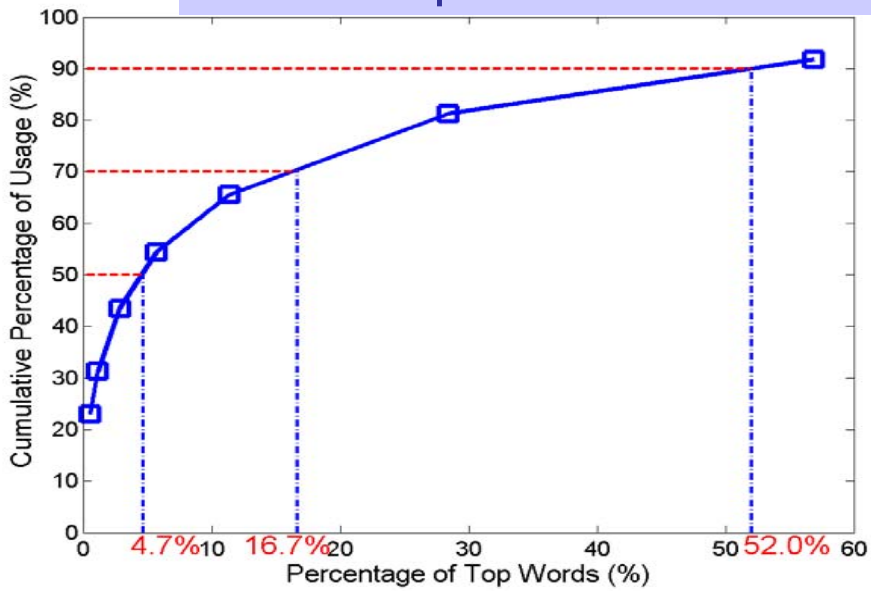

A³ Deployment Model

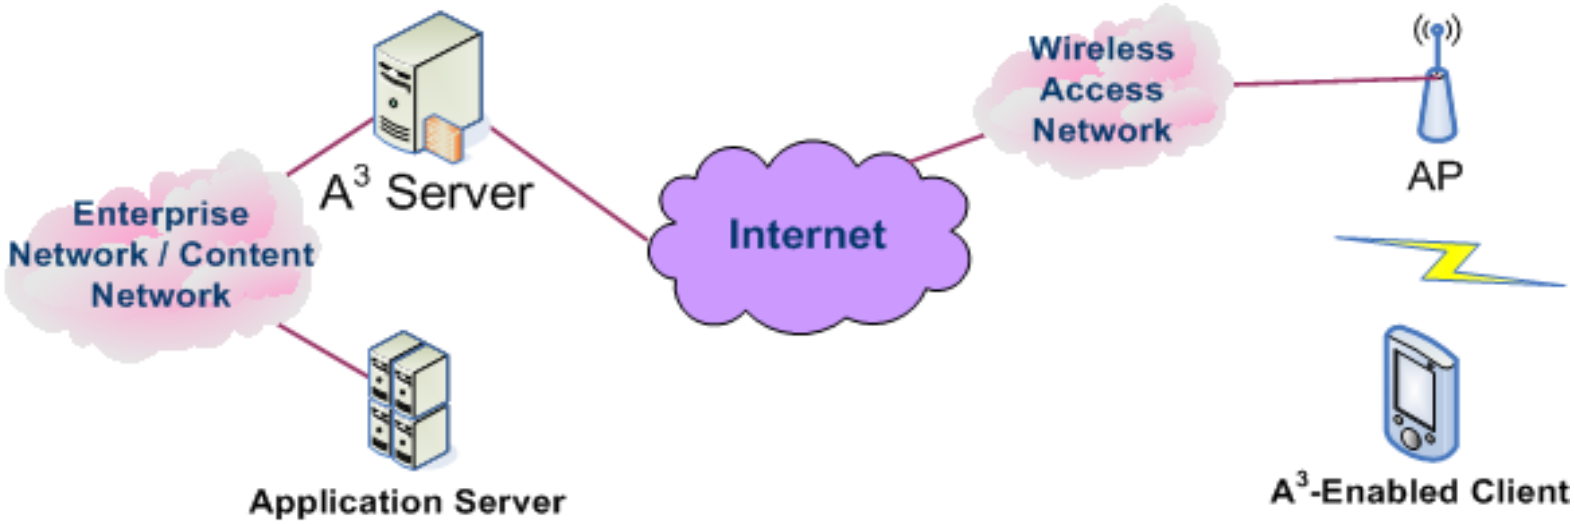

- Client side is a software module
- Server side can be software modules installed on application servers, or packet processing appliances
- Point solution is also possible
- Implementation with Netfilter for Linux Systems

Evaluation Setup

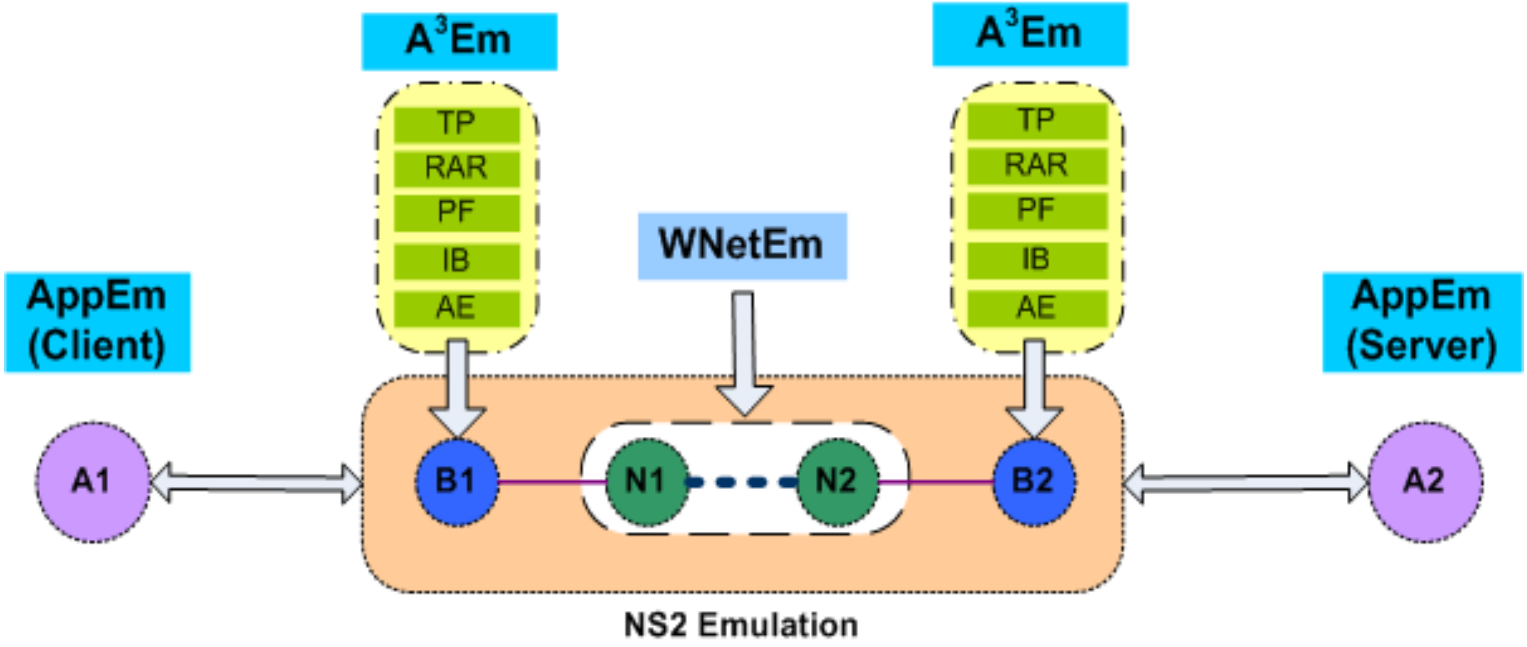

- Application Emulator (AppEm)
- A^3 Emulator (A^3Em)
- Wireless Network Emulator (WNetEm)

	WLAN	WWAN	SAT
BW (Mb/s)	5	0.1	1
RTT(ms)	5	200	1000
Loss (%)	1	8	3

Performance Evaluation

- CIFS
  - RAR, IB and TP
  - 70% improvement
- SMTP
  - RAR, IB and AE
  - 110% improvement
- HTTP
  - RAR, IB and PF
  - 30% improvement

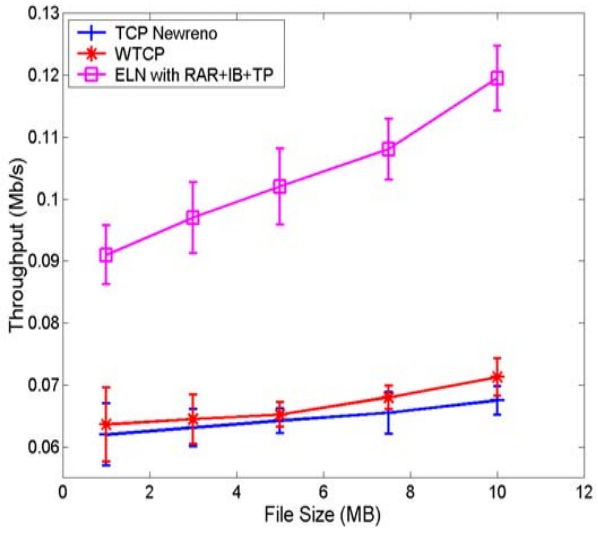

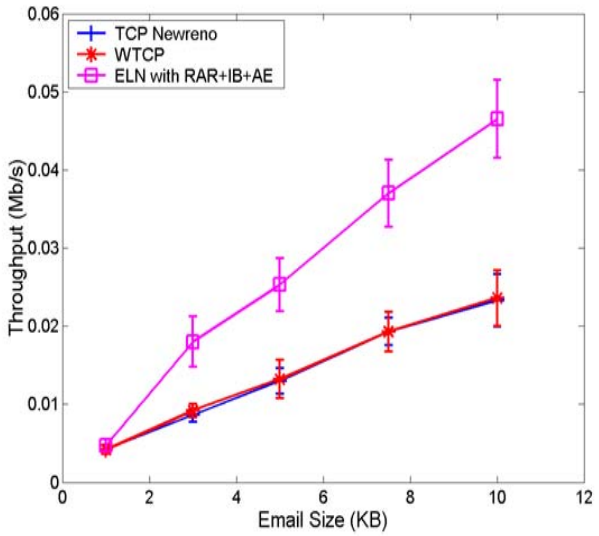

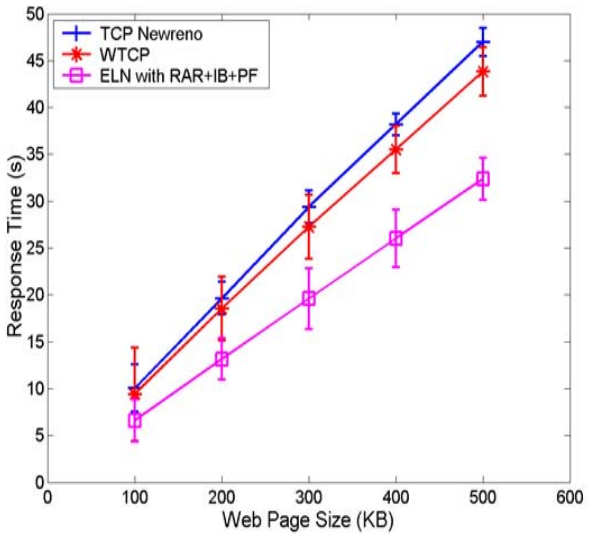

Related Works

- WAP: Wireless Application Protocol
- PIE: Pocket Internet Explorer
- Odyssey [Noble 1997]
- Coda [Satyanarayanan 1990]
- “Upload” client-tasks to server side [Czerwinski 2001]
- Out-of-order HTTP objects transmitted in UDP [Mohomed 2006]
- Commercial WAN optimizers [Riverbed, etc]

Summary

- Use emulations to test performances of several popular used applications
- Identify five application behaviors, and analyze their impacts on performances
- Propose an application-aware acceleration solution
  - TP, RAR, PF, IB, and AE
- Evaluate its effectiveness

GNAN Research Group : http://www.ece.gatech.edu/research/GNAN/

Asankya Networks, Inc. : http://www.asankya.com/