97-0425
Performance of Bursty World Wide Web (WWW) Sources over ABR

Bobby Vandalore, Shiv Kalyanaraman, Raj Jain, Rohit Goyal, Sonia Fahmy, and Xiangrong Cai
The Ohio State University

Seong-Cheol Kim, Samsung Electronics Co.

Overview

- Goal: Bursty TCP Performance over ABR
- Concern: ABR, UILI, Slow start policies may adversely affect bursty traffic
- Workload: SpecWeb'96 traffic (Modified)

Raj Jain

SPECWeb96

❑ Majority of traffic on the Internet is WWW

❑ Developed by Standard Performance Evaluation Corporation (SPEC)

❑ A consortium similar to ATM Forum for performance benchmarking

❑ SPECMark, SPEC CPU95, SPECInt95, SPEC SFS

❑ SPECWeb96 is for benchmarking WWW servers

❑ Ref: http://www.specbench.org/ost/web96/webpaper.html

SPECWeb96

Class 0	Class 1	Class 2	Class 3
p = 0.35	p = 0.5	p = 0.14	p = 0.01
0.1 kB	1 kB	10 kB	100 kb
0.2 kB	2 kB	20 kB	200 kB
...	...	...	...
0.9 kB	9 kB	90 kB	900 kB

❑ Each column has a given probability of access.

❑ Each row of the column then has a given probability of access (Poisson distribution centered around the midpoint of the class).

Problem with SPECWeb96

❑ Average File Size = 120.2 kB

❑ 10-sec think time
   $\Rightarrow$ Average load = 60 kbps per user

❑ Low bandwidth per user
   $\Rightarrow$ Need a very large number of users to fill 155 Mbps

❑ Solution: Use 150 or so users on 45 Mbps link.

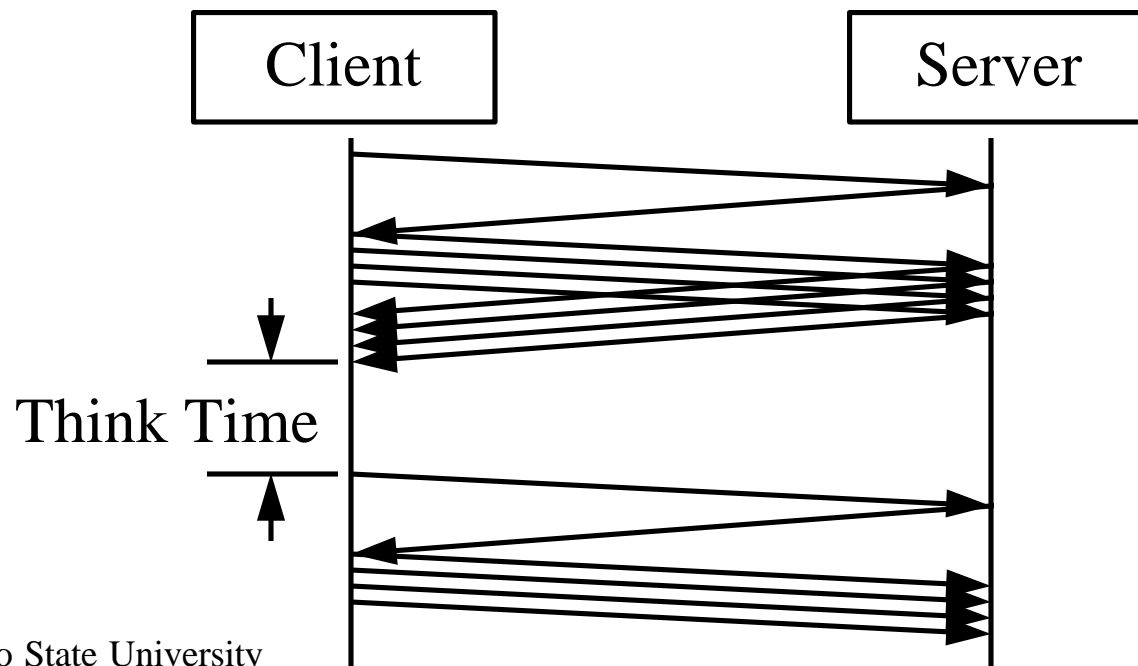❑ Also modified the workload to make 10 times more demand per user.

Modified SPECWeb96

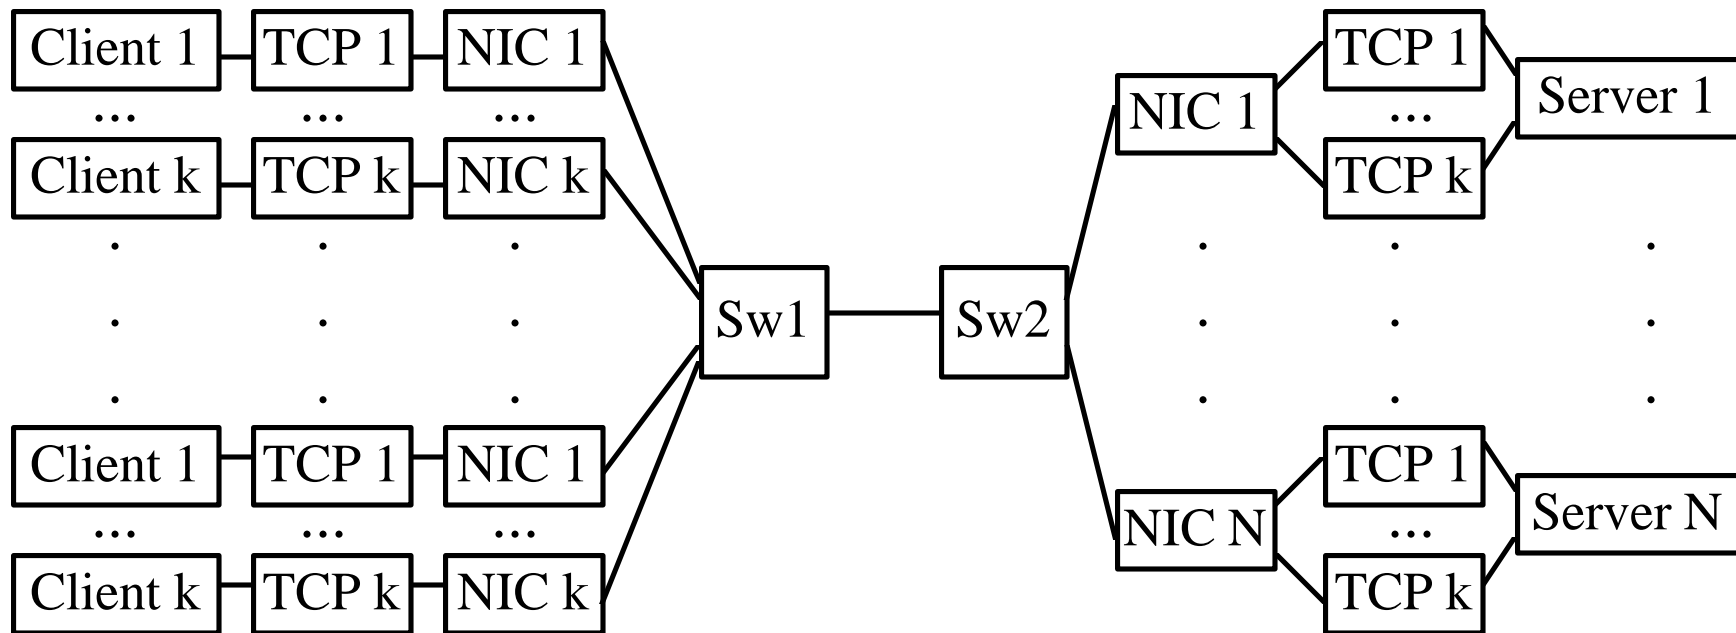| Class 0	Class 1	Class 2	Class 3	Class 4
p = 0.2	p = 0.28	p = 0.40	p = 0.112	p = 0.008
0.1 kB	1 kB	10 kB	100 kb	1 MB
0.2 kB	2 kB	20 kB	200 kB	2 MB
...	...	...	...	...
0.9 kB	9 kB	90 kB	900 kB	9 MB

❑ Each Web page consists of one index page and 4 images.

❑ First Column: Index page (p = 1/5)

❑ Other Columns: p = 0.8 Probability of class n-1 in the original SPECweb96

Raj Jain

Modified SPECWeb96 (Cont)

❑ Average File size = 120.3 kB

❑ Bandwidth per user = 0.48 Mbps

❑ HTTP 1.1 ⇒ All components of a web page are fetched in one TCP connection.

Client   Server

Think Time

K-N Client-Server Configuration

- ❑ K clients per server, K=15
- ❑ N servers, N=1, 2, 5, 10
- ❑ Total K*N VCs
- ❑ Note: Typo in contribution 10 ms = 3667 cells (not 3680)

Parameters

❑ TCP Parameters:

    ❑ Window = 16×64 kB (on WANs with 30 ms RTT)

    ❑ MSS = 512 Bytes $\Rightarrow$ 12 cells per segment
    $\Rightarrow$ Maximum TCP payload = 78% of link rate
    = 35.1 Mbps on T3

❑ ERICA Parameters:

    ❑ Averaging interval = Min{500 cells, 5 ms}

    ❑ Queue thresholds:
    T0 = 500 μs, a=1.15, b=1.05, QDLF = 0.5

❑ Workload Parameters:

    ❑ Inter-batch time = 10 seconds
    $\Rightarrow$ 72 Mbps on the T3 link with 150 clients

Simulation Results

# of Servers	Max Switch Q (Cells)	TCP Throughput	Efficiency
1	3677 = 1.0 F	6.1 Mbps	17.4%
2	6154 = 1.7 F	14.2 Mbps	40.3%
5	14057 = 3.8 F	34.1 Mbps	97.1%
10	17269 = 4.7 F	32.7 Mbps	93/2%

F = Feedback delay

- Queues increase linearly first, then stabilize
- Queue lengths are a small multiple of the feedback delay.
- Efficiency increases linearly first, then stabilizes
- Efficiency is high.
- ABR is stable even under bursty TCP traffic.