



**ENSC 835-3: HIGH-PERFORMANCE
NETWORKS**

**DYNAMIC RIGHT-SIZING
A TCP FLOW-CONTROL INVESTIGATION IN
OPNET**

by

CAMILLE JAGGERNAUTH



ROADMAP

- Introduction
- Flow Control in TCP
- OPNET Investigation of Dynamic Right Sizing
- Results
- References



INTRODUCTION

- Networking researchers typically manually optimizing TCP buffer sizes
- To keep the network pipe full
- Achieve acceptable performance over the wide-area network.



INTRODUCTION

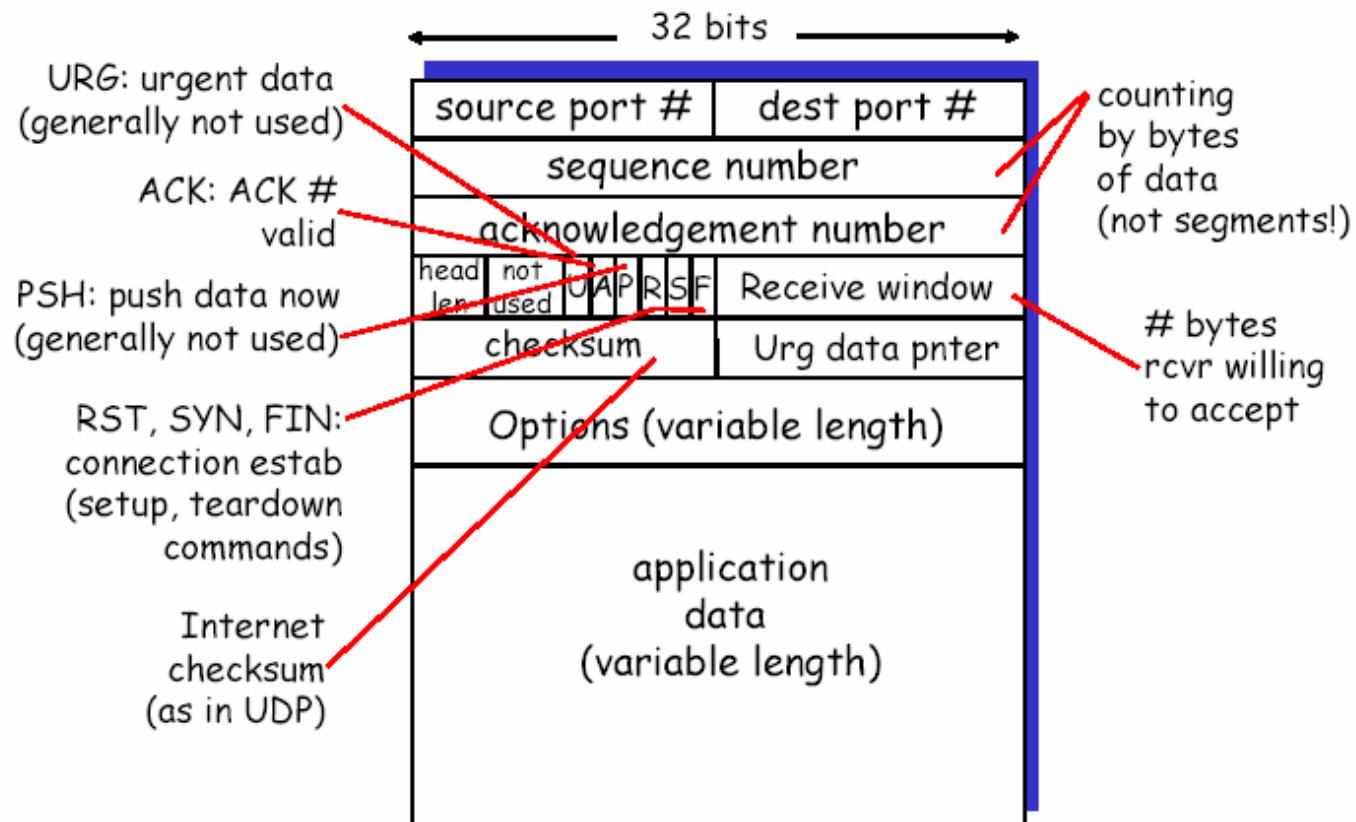
- ❑ Cumbersome
- ❑ Sub-par performance for connections with larger delay-bandwidth products
- ❑ Misappropriation of scarce resources to connections with smaller delay-bandwidth products



INTRODUCTION


- ❑ Originally implemented in ns
- ❑ Dynamic Right sizing lets the receiver estimate the sender's congestion window size
- ❑ Uses this estimate to dynamically change the size of the receiver's window advertisements.
- ❑ As a result, the sender will be congestion-window-limited rather than flow-control-window-limited

TCP SEGMENT STRUCTURE




FLOW CONTROL IN TCP

- Receive side of TCP connection has a receive buffer
- App process may be slow reading from buffer
- Won't overflow receiver's buffer by transmitting too much, too fast




FLOW CONTROL IN TCP

- **RcvWindow = RcvBuffer - [LastByteRcvd - LastByteRead]**
- Rcvr advertises spare room by including value of **RcvWindow** in segments
- Sender limits unACKed data to **RcvWindow** guarantees receive buffer doesn't overflow



OPNET SETUP



OPNET SETUP

The screenshot displays the OPNET software interface. The main window is titled '(Application Config) Attributes' and shows a tree view of application configurations. The 'Type' is set to 'utility'. The tree view includes the following items:

- name: Application Config
- model: Application Config
- Application Definitions: (...)
- rows: 1
 - row 0
 - Name: File Transfer (custom)
 - Description: (...)
 - Custom: Off
 - Database: Off
 - Email: Off
 - Ftp: (...)
 - Http: Off
 - Print: Off
 - Remote Login: Off
 - Video Conferencing: Off
 - Voice: Off
 - Voice Encoder Schemes: All Schemes

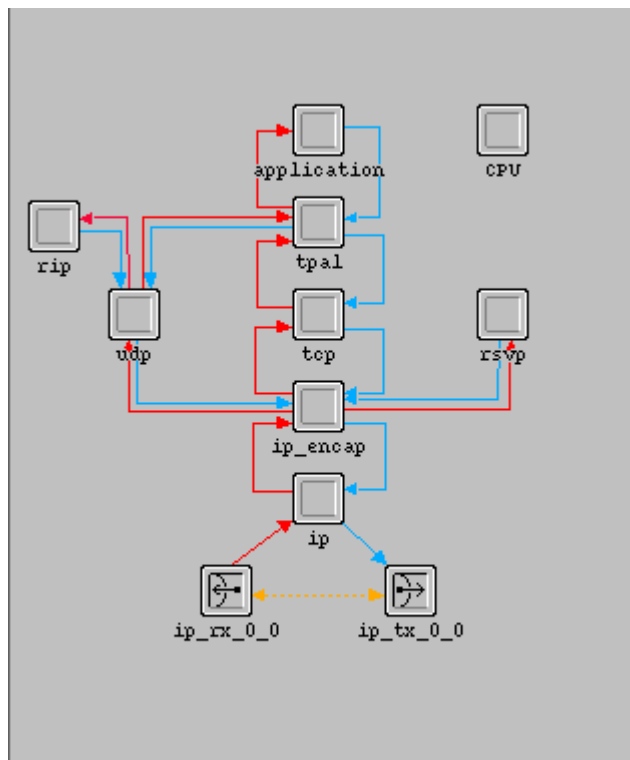
At the bottom of the dialog, there are checkboxes for 'Apply changes to selected objects' and 'Advanced', and buttons for 'End Next', 'OK', and 'Cancel'.

Overlaid on the right is a '(ftp) Table' dialog box with the following data:

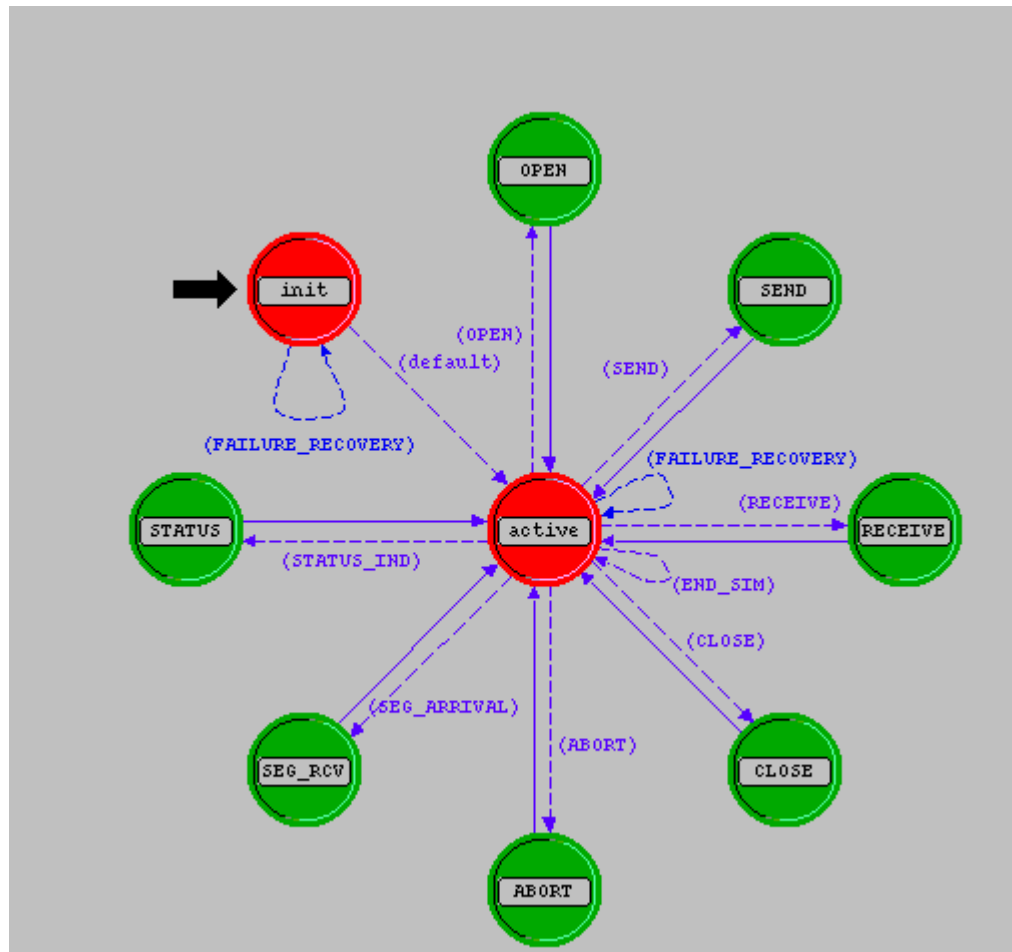
Attribute	Value
Command Mix (Get/Total)	100%
Inter-Request Time (seconds)	constant (3600)
File Size (bytes)	constant (1600000)
Symbolic Server Name	FTP Server
Type of Service	Best Effort (0)
RSVP Parameters	None
Back-End Custom Application	Not Used

Buttons at the bottom of the '(ftp) Table' dialog include 'Details', 'Promote', 'OK', and 'Cancel'.

TCP NODE MODEL



TCP PROCESS MODEL





TCP MODEL PROCESSES

- ❑ `tcp_manager_v3.c`
- ❑ Represents the root process of the “tcp” module.
- ❑ It manages a set of TCP connections by invoking the appropriate api processes

IMPLEMENTATION STRATEGY

- Receiver can determine the sender size by below equation
- d is bytes of data received
- round trip times between N_{min} and N_{max}

$$\frac{d}{n_{max}} \leq w \leq \frac{d}{n_{min}}$$



IMPLEMENTATION STRATEGY

- Observing the Congestion Window variable change
- Altering the receiver buffer window accordingly



STATISTICS OBSERVED

- Delay
- Receiver Window Size
- Congestion Window Size
- Flight Size (sent but unacked data in sender buffer)
- Traffic Received



RESULTS

- The first 25 transfers used the default window sizes of 64 kilobytes for both the sender and receiver.
- The second 25 transfers, shown in dotted lines, used dynamically sized windows

RESULTS

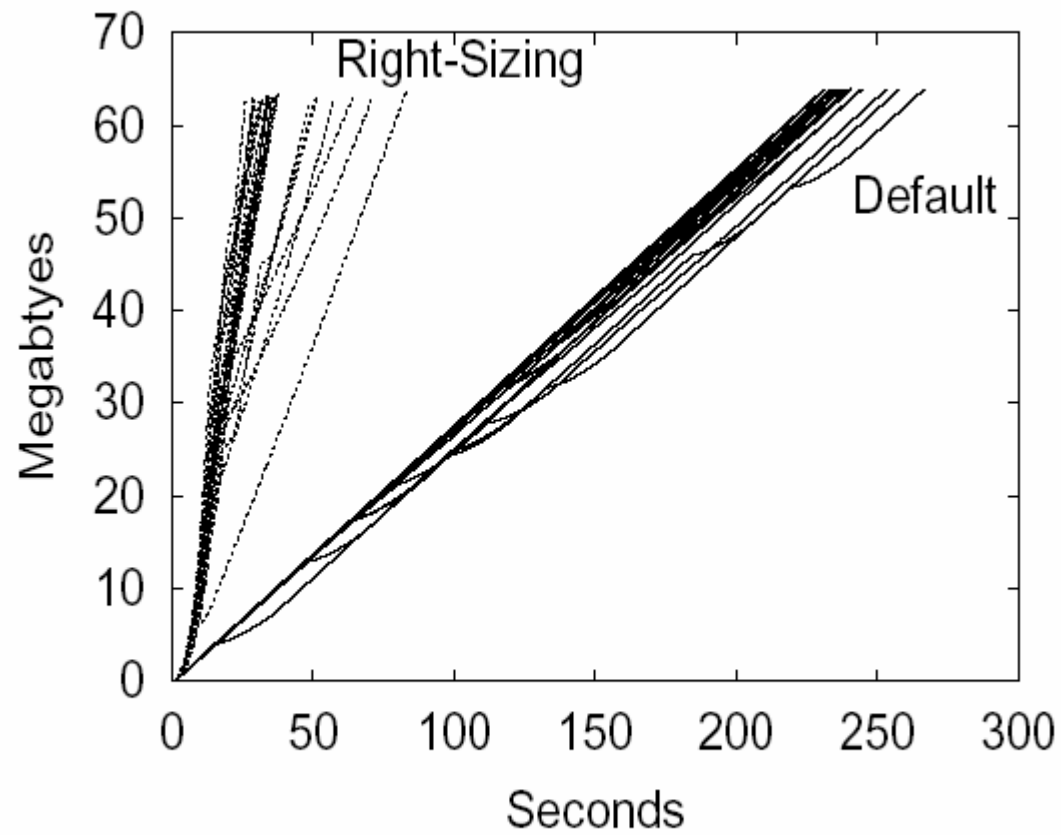


Figure 1: Progress of data transfers

RESULTS

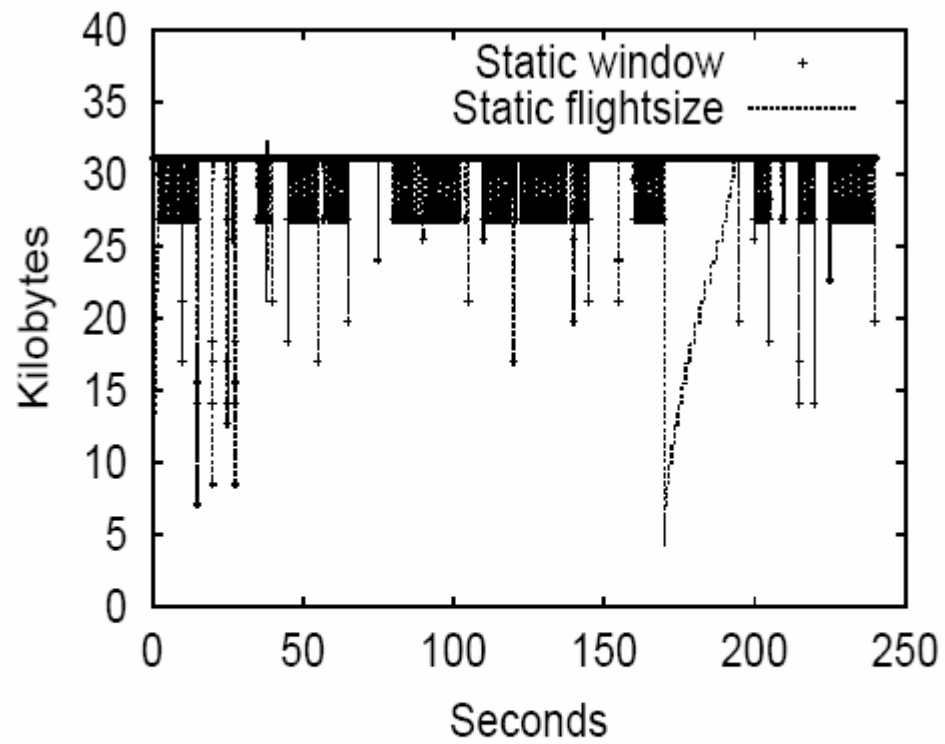


Figure 2: Default window size
Flight & window sizes

RESULTS

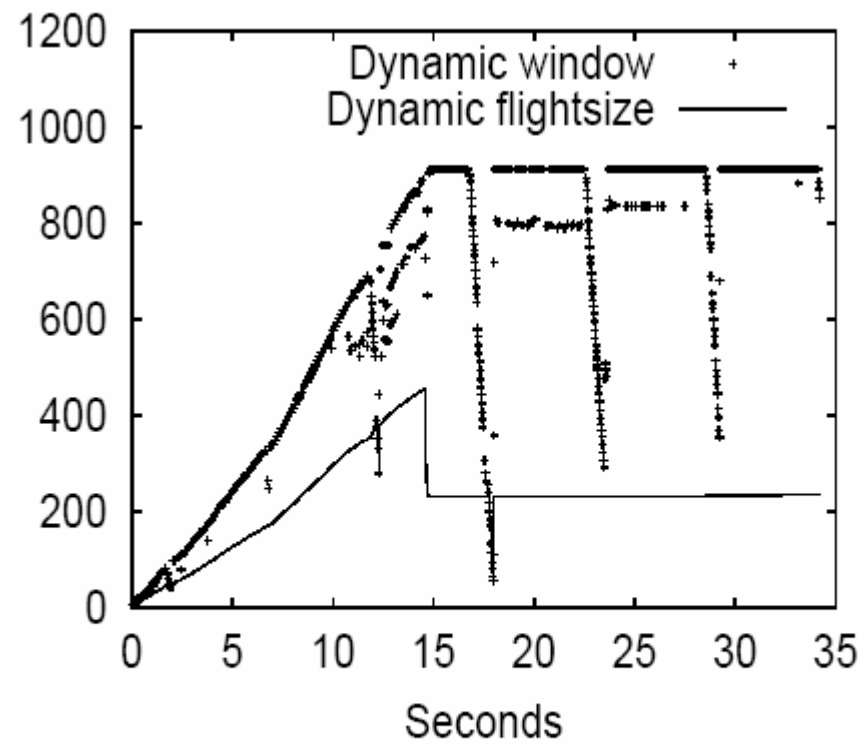


Figure 3: Dynamic right-sizing
Flight & window sizes



REFERENCES

- Mike Fisk and Wu-chun Feng, Dynamic Right-Sizing: TCP Flow-Control Adaptation, in Proceedings of the 14th Annual ACM/IEEE SC2001 Conference, November 2001.
<http://public.lanl.gov/radiant/pubs/drs/sc2001-poster.pdf>
- RFC1323 – TCP Window Scale Option
- Kurose and Ross – Chapter 3 presentation slides pg 72, 73 Computer Networking: A Top Down Approach Featuring the Internet, 3rd edition. Jim Kurose, Keith Ross Addison-Wesley, July 2004.
- Van Jacobson, Congestion Avoidance and Control, in Proceedings, SIGCOMM 1988 Workshop. ACM SIGCOMM, Aug. 1988, pp. 314-329, ACM Press, Stanford, CA.
- Jeff Semke, Jamshid Mahdavi, and Matt Mathis., Automatic TCP buffer tuning, Computer Communications Review, vol. 28, no. 4, pp. 315-323, Oct. 1998.
- Jian Lui and Jim Ferguson, "Automatic TCP socket buffer tuning," in Supercomputing 2000 Research Gems, Nov. 2000, Awarded Best Research Gem of the Conference
Brian L. Tierney