Experiences with Wide-Area ATM Networking

Joseph B. Evans, Gary J. Minden, David W. Petr, Douglas Niehaus

Presenter: Victor S. Frost

Dan F. Servey Distinguished Professor
Electrical Engineering and Computer Science
Executive Director for Research
Information and Telecommunications Technology Center
University of Kansas
2291 Irving Hill Dr.
Lawrence, Kansas 66045
Phone: (785) 864-4833 FAX:(785) 864-7789
e-mail: frost@eecs.ukans.edu
http://www.ittc.ukans.edu/

University of Kansas

Information and Telecommunication Technology Center

Overview

- Networking experiences on the MAGIC testbed
  - Introduction to MAGIC
  - Early experiences with TCP/IP over ATM WANs
- Networking experiences on the AAI testbed
  - Overview of AAI
  - Measurement of ATM WAN performance
  - Simulation tools of WAN performance

Multidimensional Applications and Gigabit Internetwork Consortium (MAGIC)

- An architecture and implementation of a nationwide internet of high-speed IP/ATM testbeds
- A scalable, dynamically constructed, network-based, distributed storage system
- Distributed processing to enable on-demand data visualization
- Controlled access to datasets and to computing resources
- An interactive application for 3-D fusion and visualization of geo-referenced data
- Techniques for adapting application to network conditions and host capabilities

MAGIC-II Participants

- DARPA-funded
  - University of Kansas (Prime contractor)
  - Corporation for National Research Initiatives
  - Earth Resources Observation Systems Data Center
  - Lawrence Berkeley National Laboratory
  - Minnesota Supercomputer Center
  - SRI International
- Organizations contributing resources
  - Sprint
  - Splitrock Telecom

MAGIC-II Core Network

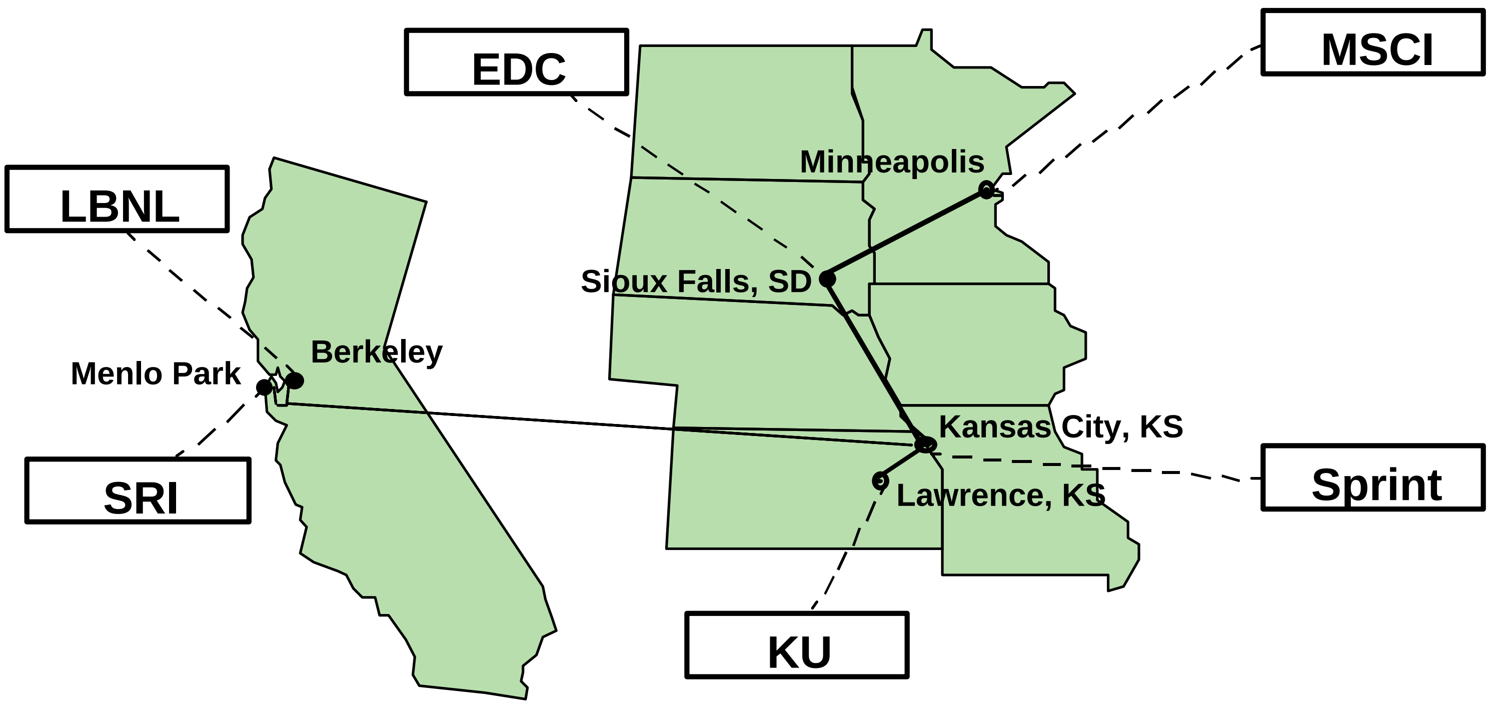

OC-48 backbone in the Midwest, OC-3 in California, DS3 connectivity between the Midwest and California

Seven sites with OC-3 or OC-12 access

Each site has an ATM LAN and multiple workstations for distributed storage and processing

MAGIC-II Nationwide Test Environment

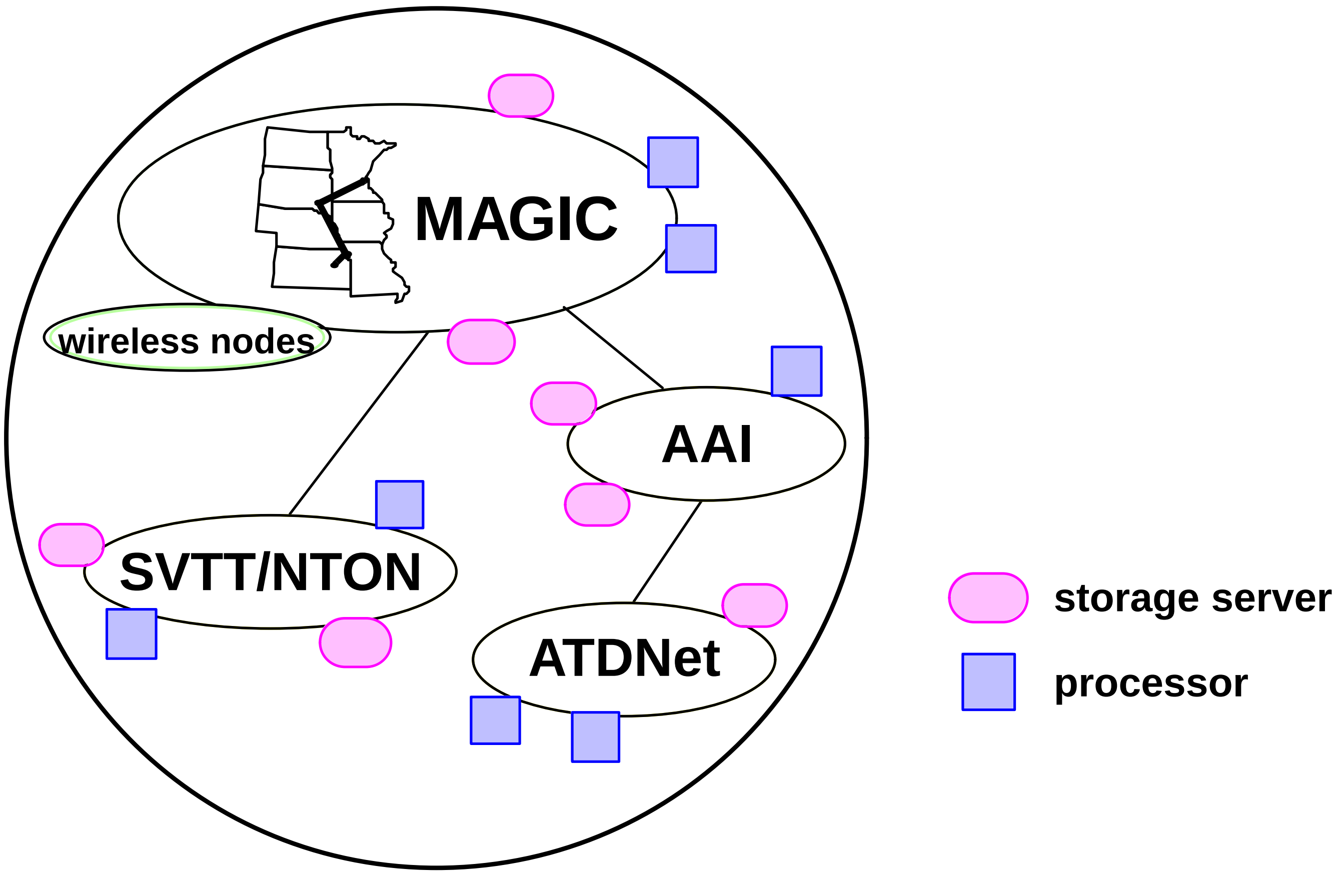

KU-ITTC MAGIC-II Research Agenda

- Create a diverse large scale network incorporating ATM wireless, distributed computing and storage technologies within the MAGIC-II internetwork, resulting in a network system with a wide range of link bandwidths and quality as well as network element capabilities.
- Develop, implement, and demonstrate technologies to monitor and distribute network 'state' to enable applications to work at their highest efficiency while satisfying users requirements in dynamic environments.
- Develop, implement, and demonstrate technologies to provide application specific services using network 'state' information to respond to dynamic environments.

TCP/IP Over ATM WANs
Early Experiences (Early 1993)

- MAGIC testbed — tests over 1000 km WAN
- High throughput hosts and interfaces
- DEC Alphas capable of 134 Mb/s TCP throughput
- DEC OTTO interface · ATM @ SONET OC-3c rates
- ATM cell-level flow control — OTTO and AN2 switch
- ATM cell-level pacing — OTTO/AN2 scheduled transmission mode

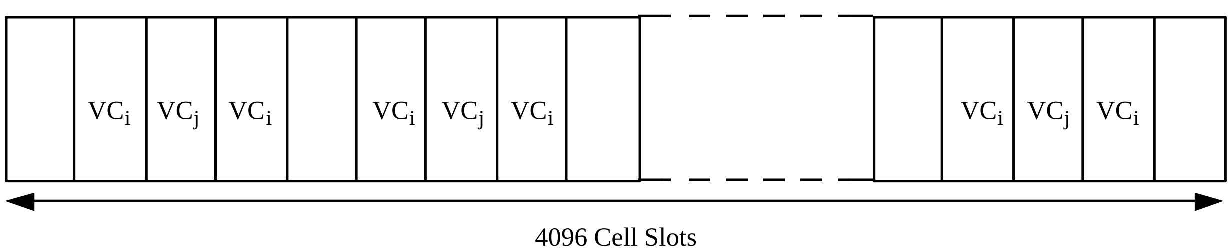

Experiment 1

- Question: WAN performance limited by TCP window size?
- Experiment: DEC Alpha 3000/400 with a DEC OTTO OC-3c interface to DEC Alpha 3000/400 over a 600 km link, 8.8 ms round-trip delay
- Results

TCP Window Size	0.5k	1k	2k	4k	8k	16k	32k	64k	128k
Throughput (Mb/s)	0.47	0.93	1.8	3.7	7.4	14.9	29.8	59.6	119

- Consistent with the theoretical limits caused by latency
- Large windows necessary for acceptable throughput

Experiment 2

- Questions: High bandwidth TCP sources will overrun ATM switch buffers at points of bandwidth mismatch? improved by pacing?
- Experiment: Alpha (OC-3c) in Lawrence, Kansas, to SPARC-10 (TAXI) in South Dakota (600 km) — a single host to another host
  - Alphas with DEC OTTO cards, SPARC-10 with FORE Systems 100 Mb/s TAXI
  - Switches --> FORE Systems ASX-100
  - 128 kB TCP windows, 64 kB write buffers
- Results

No Pacing	Pacing
0.87 Mb/s	68.20 Mb/s

Experiment 3

- Question: Will TCP rate control be more effective if TCP segment size small relative to buffers?
- Experiment: Alpha (OC-3c) in Lawrence, Kansas, to SPARC-10 (TAXI) in South Dakota (600 km), vary TCP segment size
- Results:

File Name : edc_mtu.eps
Creator : gnuplot

ACTS ATM Internetwork (AAI)

(ACTS = Advanced Communications Technology Satellite)

- Objectives
  - Evaluate use of ATM WAN for joint use of parallel and vector processors
  - Evaluate use of national-scale, high-speed terrestrial/satellite ATM network
  - Evaluate ATM WAN for congestion, signaling, and multicast technologies

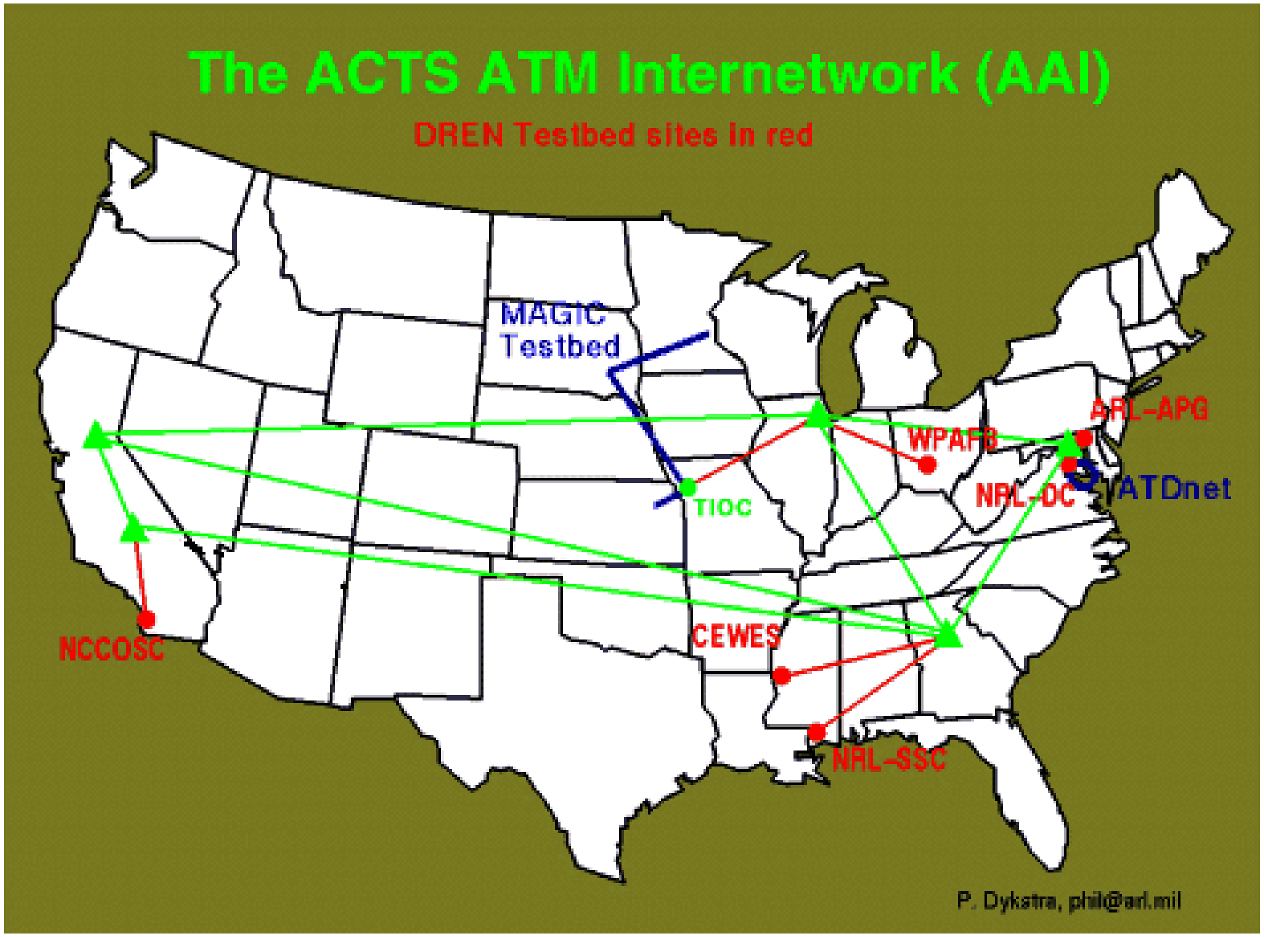
The ACTS ATM Internetwork (AAI)
DREN Testbed sites in red
MAGIC
Testbed
TIOC
WPAFB
ARL-APG
NRL-DC
ATDnet
NCCOSC
CEWES
NRL-SSC
P. Dykstra, phil@arl.mil

AAI Network Topology

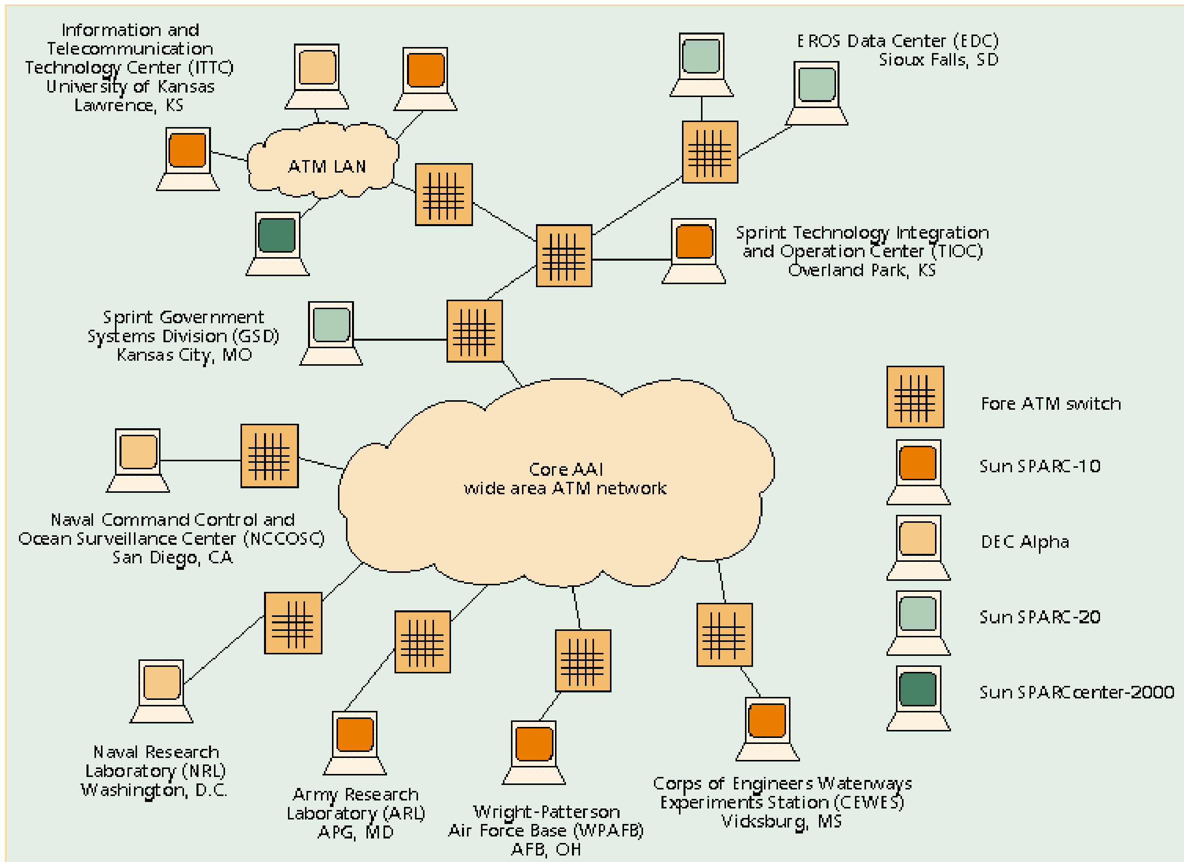

KU-ITTC AAI Research Agenda

- Determine performance characterization of ATM WANs
  - Measurement
  - Simulation
- Characterize ATM WAN traffic profiles
- Evaluate performance of ATM WAN congestion controls

WAN Measurement Tools

- NetSpec: A first step toward network benchmarking
  - Multiple host network loading
  - Automated execution
  - Reproducible experiments
  - Multiple traffic types
    - Full speed (as fast as the source can transmit to the network)
    - Constant bit rate, CBR (transmission of a periodic pattern of bursts
    - Random (transmission of a random pattern of bursts)
      - WWW
      - FTP
      - MGEG Video
      - Teleconferencing video
      - Voice
      - Telnet

NetSpec Example

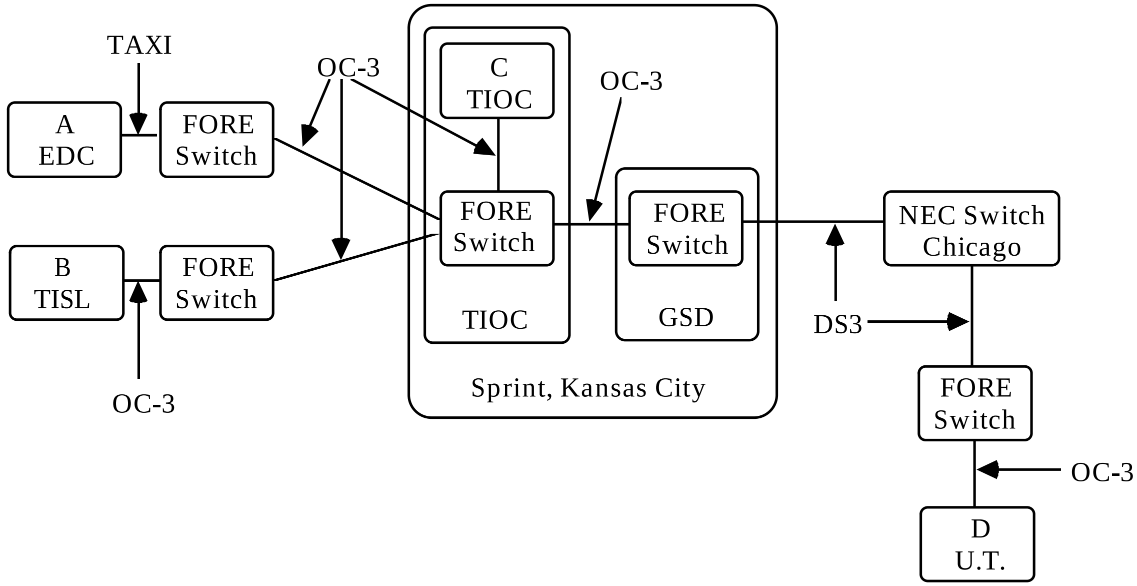

	Tx (Mbps)	Rx(Mbps)
A-C	29.319	29.287
B-D	29.366	29.204

Aggregate Network Throughput Performance

- Throughput metrics
  - Maximum losses throughput
  - Peak throughput
  - Full load throughput

Transfer from local to remote host memory as fast as possible

Throughput versus Aggregate Load

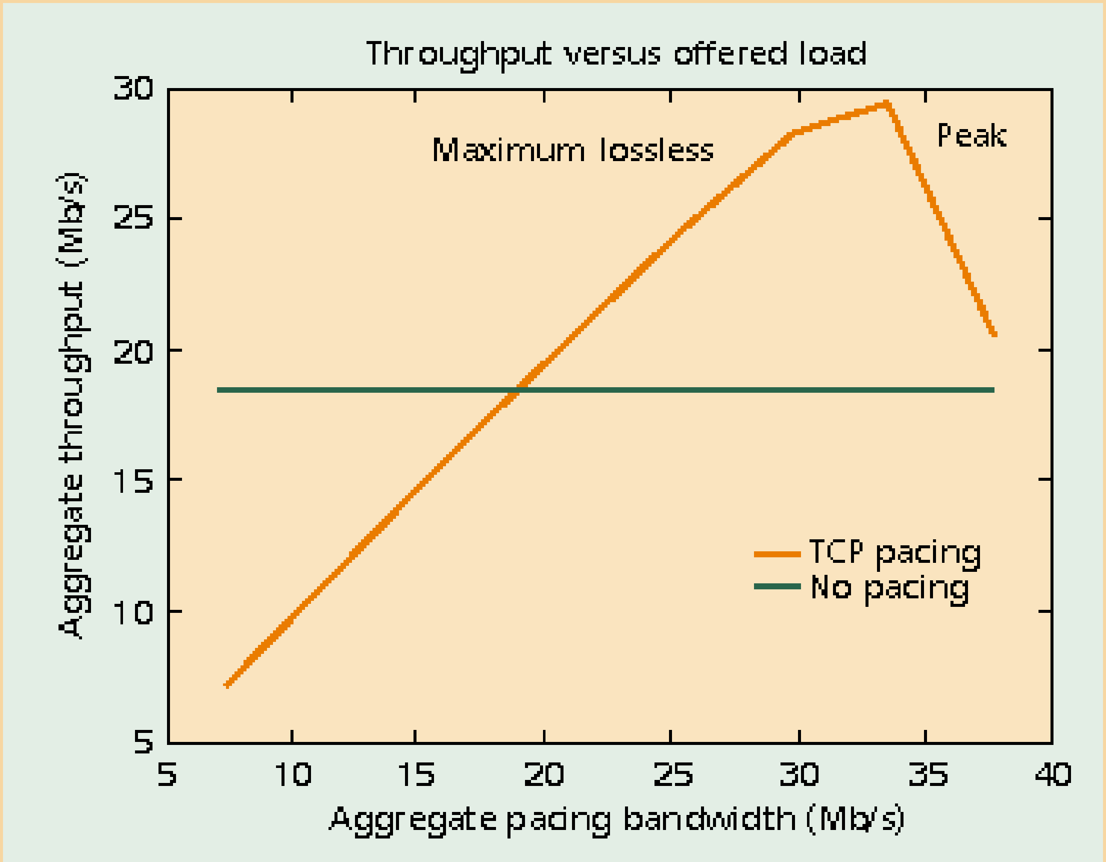

Performance of FTP over ATM WAN's

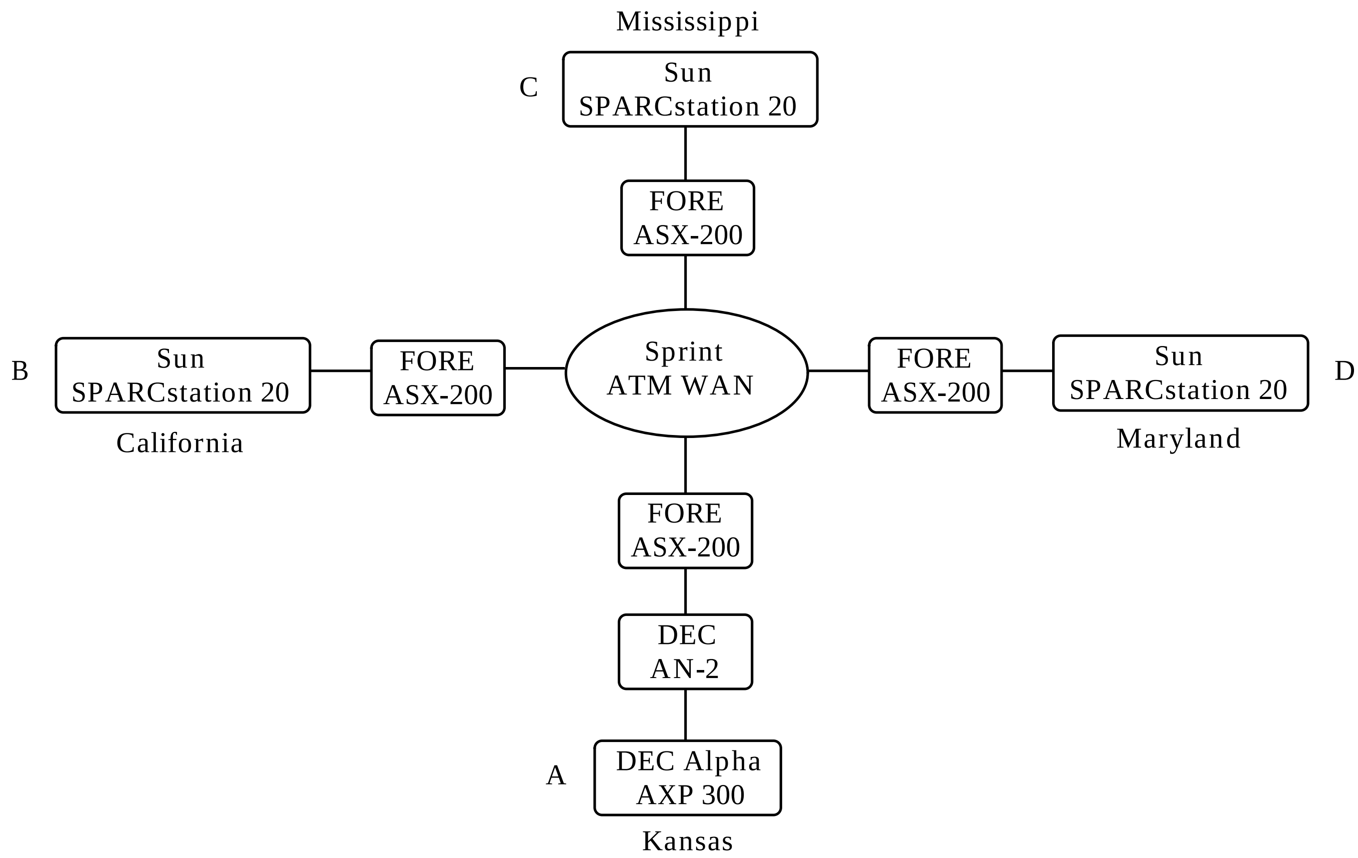

Throughput Performance with Standard FTP

Source/Destination	File Size (million bytes)	Throughput (Mbps)
B-C	14	2.56
B-A	17	3.60
	55	3.63
D-B	75	2.88
D-A	75	4.61
D-C	75	3.39
A-D	333	2.56
A-B	333	1.76
A-C	333	3.28

Throughput Performance with Modified FTP

Source/Destination	MTU Size (Bytes)	File Size (million bytes)	Throughput (Mbps)
B-D	4470	17	15.2
B-D	4470	55	12.8
D-B	9188	17	22.4

FTP Throughput Results with Large Windows

Simulation of ATM WAN's

- Goals
  - Determine the level of model fidelity required to accurately predict ATM WAN performance
  - Determine the feasibility of measurement based validation of ATM WAN simulation models
  - Identify factors influencing ATM WAN performance

Simulation Parameters

System Parameter	Value
TCP MTU size	9180 bytes
TCP processing and OS overhead time	
- DEC 3000 AXP	200-300 µs
- SGI	550 µs
- SPARC 10	550 µs
- SPARC 5	700µs
TCP user send buffer size	64 kBytes
Slow-timer period	0.5 s
Fast-timer period	0.2 s
Minimum RTO	1.0 s
AAL5 SAR processing time	0.2 µs
AAL5 cell payload size	48 Bytes
Switch processing time	4 µs
Switch output buffer size per VC	256 cells
OC-3c link speed	155 Mb/s
TAXI link speed	100 Mb/s
DS-3 link speed	45 Mb/s

Network Configuration

File Name : 11tiocedcarl.epsi
Title : /tmp/xfig-fig000874
Creator : fig2dev
CreationDate : Thu May 16 14:09:32 1996
Pages : 1

Simulation Model

File Name : 12tiocarledc.ps
Creator : BONeS Designer 3.0.1 / Block Diagram Editor Printer 1.0
CreationDate : Friday, 3/15/96 02:00:38 pm CST
Pages : 1

University of Kansas

Information and Telecommunication Technology Center

Comparison of Experimental and Simulation Performance Predictions

Connection	Experimental Results	Simulation Results
Baseline results: Point-to-point connections		
TIOC to ARL	4.2 Mb/s	7.18 Mb/s
TIOC to EDC	64.2 Mb/s	65.98 Mb/s
Simultaneous traffic streams: Single source, two destinations		
TIOC to ARL	4.45 Mb/s	4.60 Mb/s
TIOC to EDC	64.36 Mb/s	61.37 Mb/s
Simultaneous traffic streams: Two sources, single destination		
ARL to TIOC	2.15 Mb/s	4.87 Mb/s
EDC to TIOC	52.42 Mb/s	65.01 Mb/s
Simultaneous full duplex traffic streams		
TIOC to ARL	4.34 Mb/s	5.16 Mb/s
ARL to TIOC	4.3 Mb/s	5.16 Mb/s
TIOC to EDC	22.18 Mb/s	41.80 Mb/s
EDC to TIOC	31.18 Mb/s	41.30 Mb/s

Lessons Learned

- ATM wide-area networking is a reality
- High throughput is achievable with TCP/IP over ATM WANS
- Complex traffic control is feasible at high speeds
- There is a growing need for network-wide benchmarking tools, e.g., NetSpec
- Simulation of large and complex ATM networks is computationally intensive
- Computer simulation can be used to predict the performance of some aspects of ATM WANs