QoS
Demands of Large Scale Scientific Visualization Applications

Jason Leigh

Electronic Visualization Laboratory
University of Illinois at Chicago

Corridor One : Research in Distance Collaboration and Visualization of Large Scale Scientific Data

- ANL, LBNL, LANL, Utah, Princeton, UIC, U Chicago
- Remote and Distributed Rendering
- Protocols for Remote Visualization
- Progressing Refinement
- Deep Images and Image Based Rendering
- Compression for Visualization Streams
- Remote Immersive Visualization
- Data Organization for Fast Remote Navigation
- High-end Collaborative Visualization Environments
- Collaborative Dataset Exploration and Analysis
- User Interfaces and Computational Steering
- Distributed Network Attached Framebuffers

Visualization Domains & Techniques

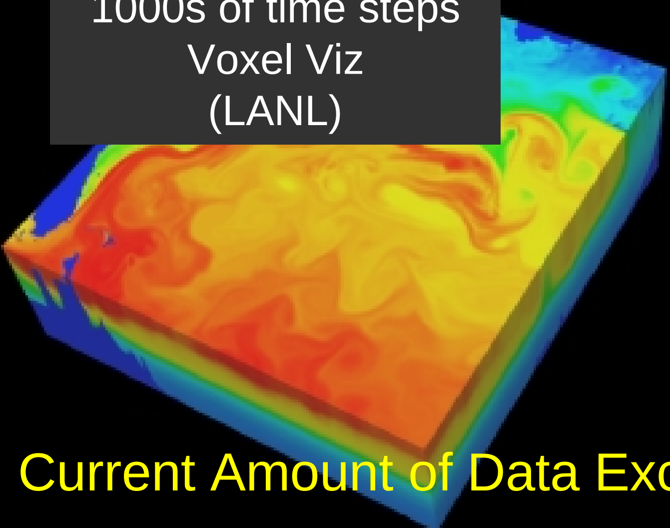

Climate Simulation
3000x4000x100 cells,
1000s of time steps
Voxel Viz
(LANL)

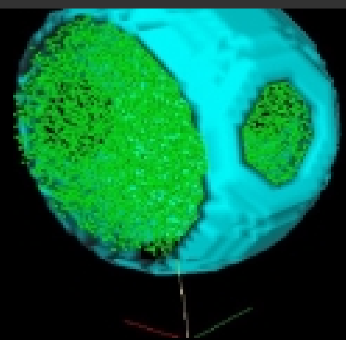

Neutron Transport in a
Uranium Sphere
Iso Viz
(LANL)

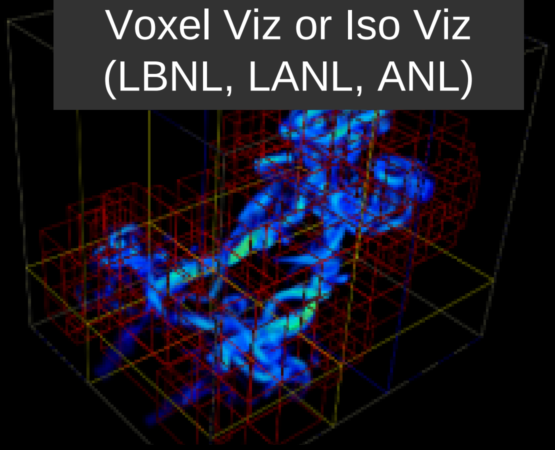

AMR from Combustion
Modeling
Voxel Viz or Iso Viz
(LBNL, LANL, ANL)

Current Amount of Data Exceeds our Ability to Visualize it in “Realtime” ~15fps

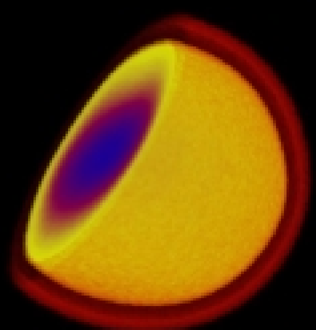

Accelerator Simulations
300M particles/time step
Particle Viz or Voxel Viz
(LANL)

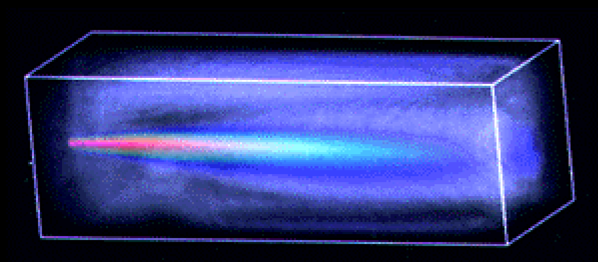

Smoke Viz of Gas Combustion
Voxel Viz
(Utah CSAFE)

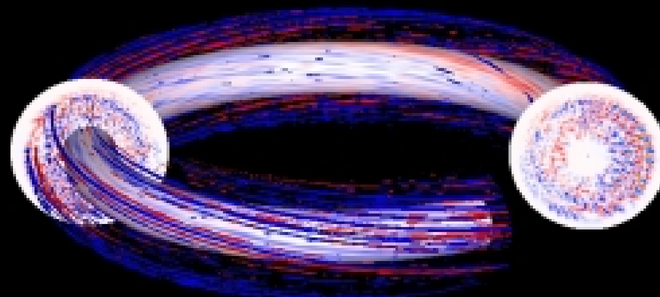

Plasma Flow Viz
Voxel Viz or Vector Viz
(Princeton)

Visualization Displays

Minnesota Power Wall (5x2)

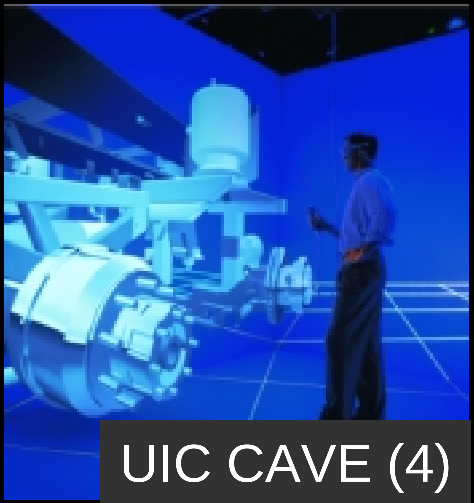

UIC CAVE (4)

ANL Active Mural (5x3)

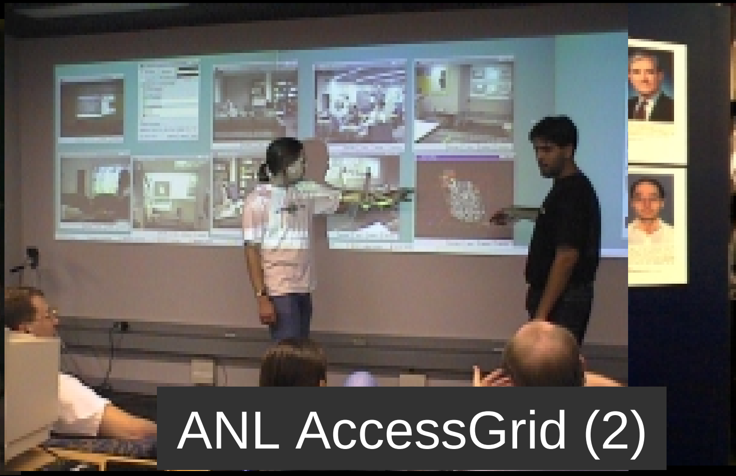

ANL AccessGrid (2)

UIC ImmersaDesk (1)

Graphics Capabilities and Bandwidth Requirements to Sustain Streaming Distance Visualization *for 1 user*

		Desktop or ImmersaDesk (1)	CAVE (4)	Active Mural (5x3)	Powerwall (5x2)
Geometry	3M lit polygons/s per pipe	864M	864M	864M	864M
Volume	10M cell structured mesh @ 10fps	1.2G	1.2G	1.2G	1.2G
Volume	1M cell unstructured mesh @ 10fps	125M	125M	125M	125M
Image	1Kx1K frame buffer @ 15fps with 10:1 compression	50M	200M	750M	500M
Stereo Image	1Kx1K frame buffer @ 15fps with 10:1 compression	100M	400M	1.5G	1G

UIC's work in Corridor One

- Collaboration in Teleimmersive Environments
- CAVERNsoft – a toolkit for building Teleimmersive Applications
- Two examples:
  - Teleimmersive Data Explorer (TIDE) – Nikita Sawant, Chris Scharver
  - Collaborative Image Based Rendering Viewer (CIBR View) – Jason Leigh, Steve Lau [LBL]
- Plan to test these applications over EMERGE DiffServ testbed.

TIDE

CIBR View

Common Characteristics of both Teleimmersive Applications

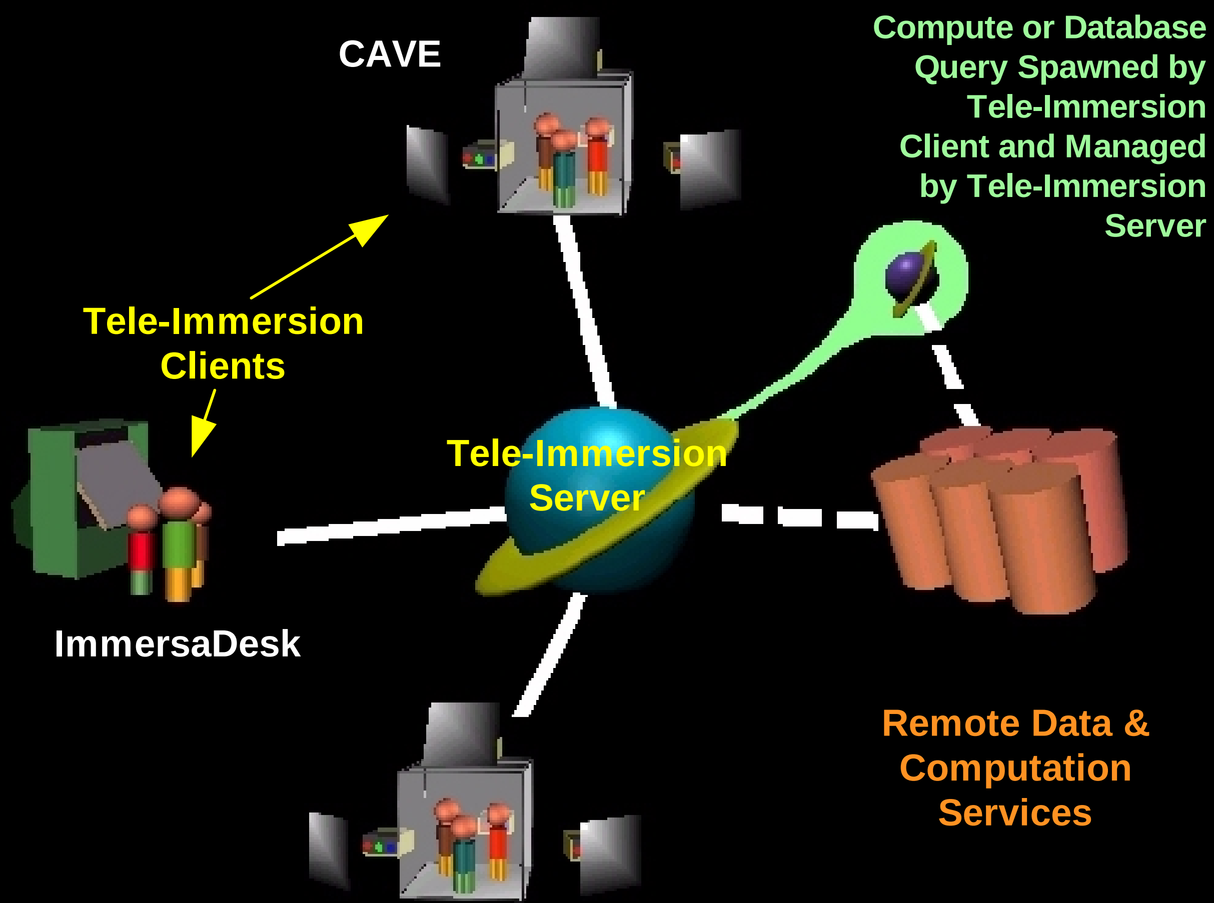

Characterization of TIDE & CIBRview streams

	Estimated bandwidth (bits/s)	DiffServ Types	Burstiness	Latency sensitive	Jitter sensitive	Error sensitive
UDP avatar	6K x n (15fps)	Interactive Real-time	**Constant**	**Y**	**Y**	**N**
UDP audio stream	64K x n		**Brief**	**Y**	**Y**	**N**
UDP video stream	10M (2-way only)		**Constant**	**Y**	**Y**	**YN**
Viz stream Playback (UDP or TCP)	See prev slide	Non-interactive Real-time	**Constant**	**N**	**N**	**YN**
TCP control data	7K x n	Reliable	**Brief**	**YN**	**YN**	**Y**
TCP bulk data	depends	Best Effort or Deadline Delivery	**Sustained burst**	**N**	**N**	**Y**

Compensating for Lack of QoS

- Forward Error Correction to expedite TCP control data
- Parallel Socket Striping to expedite TCP bulk data transfer

Forward error correction scheme for low-latency delivery of error sensitive data

- Ray Fang, Dan Schonfeld, Rashid Ansari
- Transmit redundant data over high bandwidth networks that can be used for error correcting UDP streams to achieve lower latency than TCP.
- Transmit redundant data to improve quality of streamed video by correcting for lost packets.

FEC Experiments

- EVL to SARA- Amsterdam (40Mb/s 200ms RT latency)
- Broader Ques:
  - Can FEC provide a benefit? How much?
  - Tradeoff between redundancy and benefit?
- Specific Ques:
  - TCP vs UDP vs FEC/UDP
  - How much jitter does FEC introduce?
  - High thru put UDP vs FEC/UDP to observe loss & recovery
  - UDP vs FEC with background traffic
  - FEC over QoS: WFQ or WRED congestion management- hypothesis: WRED is bad for FEC

Latency of transmitting 100 packets under UDP, TCP, FEC/UDP with 3:1 redundancy.

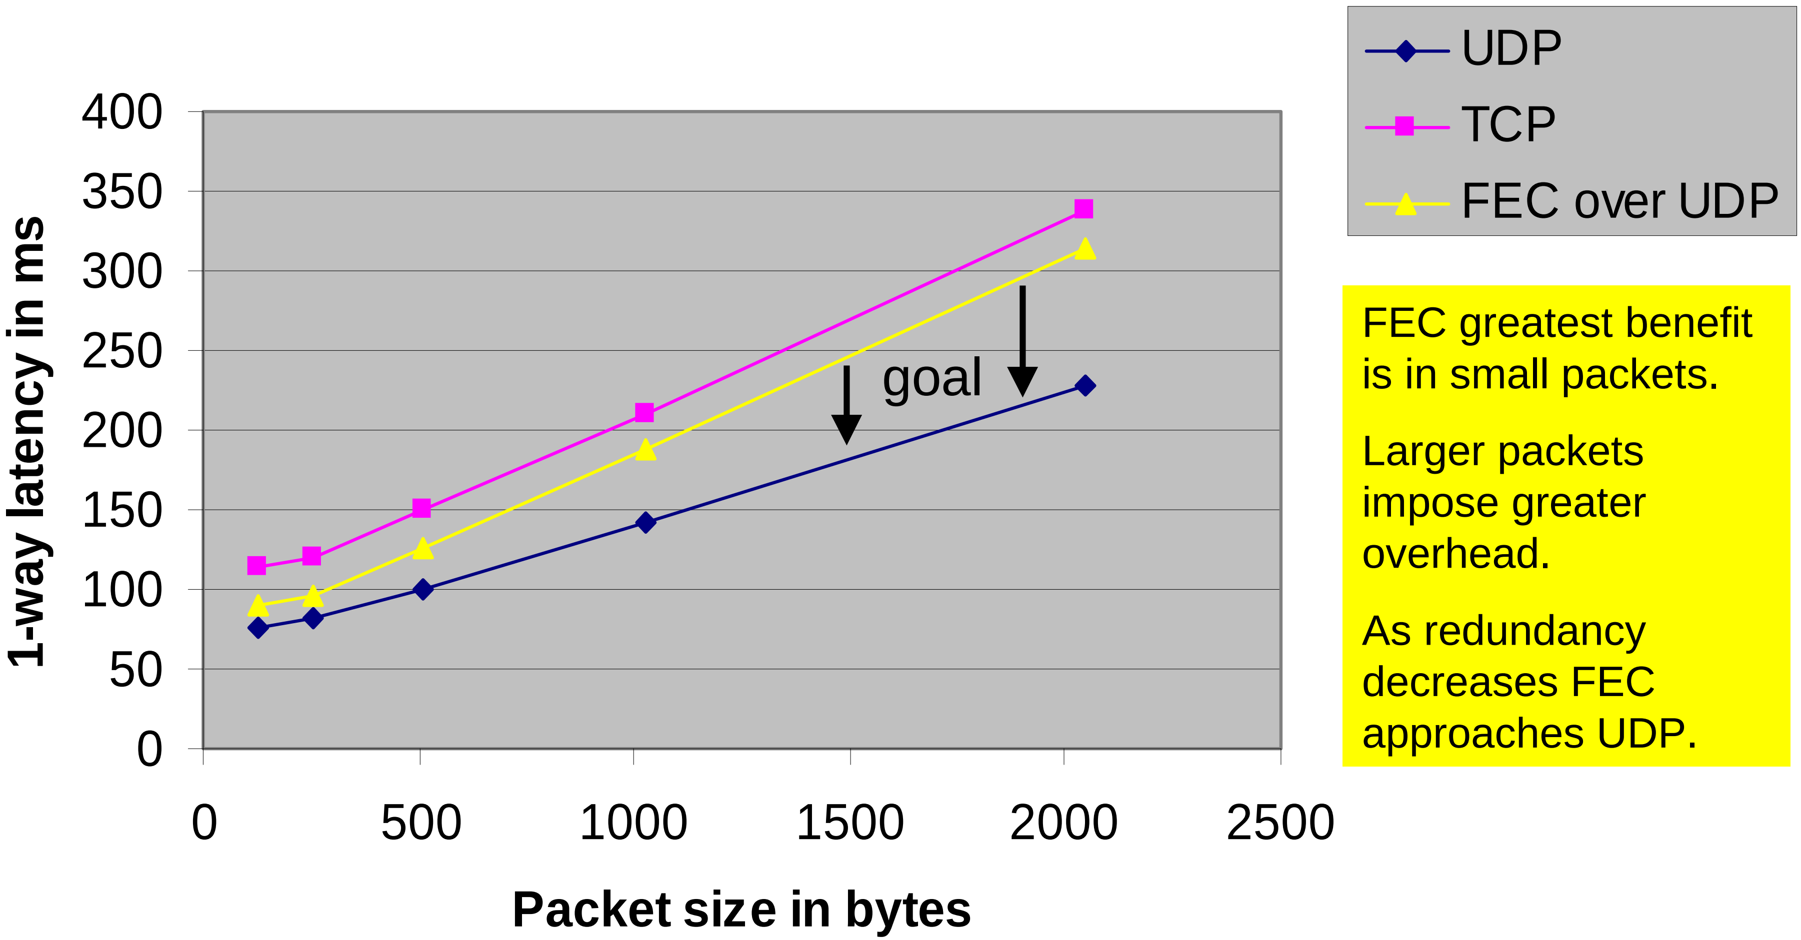

FEC greatest benefit is in small packets.

Larger packets impose greater overhead.

As redundancy decreases FEC approaches UDP.

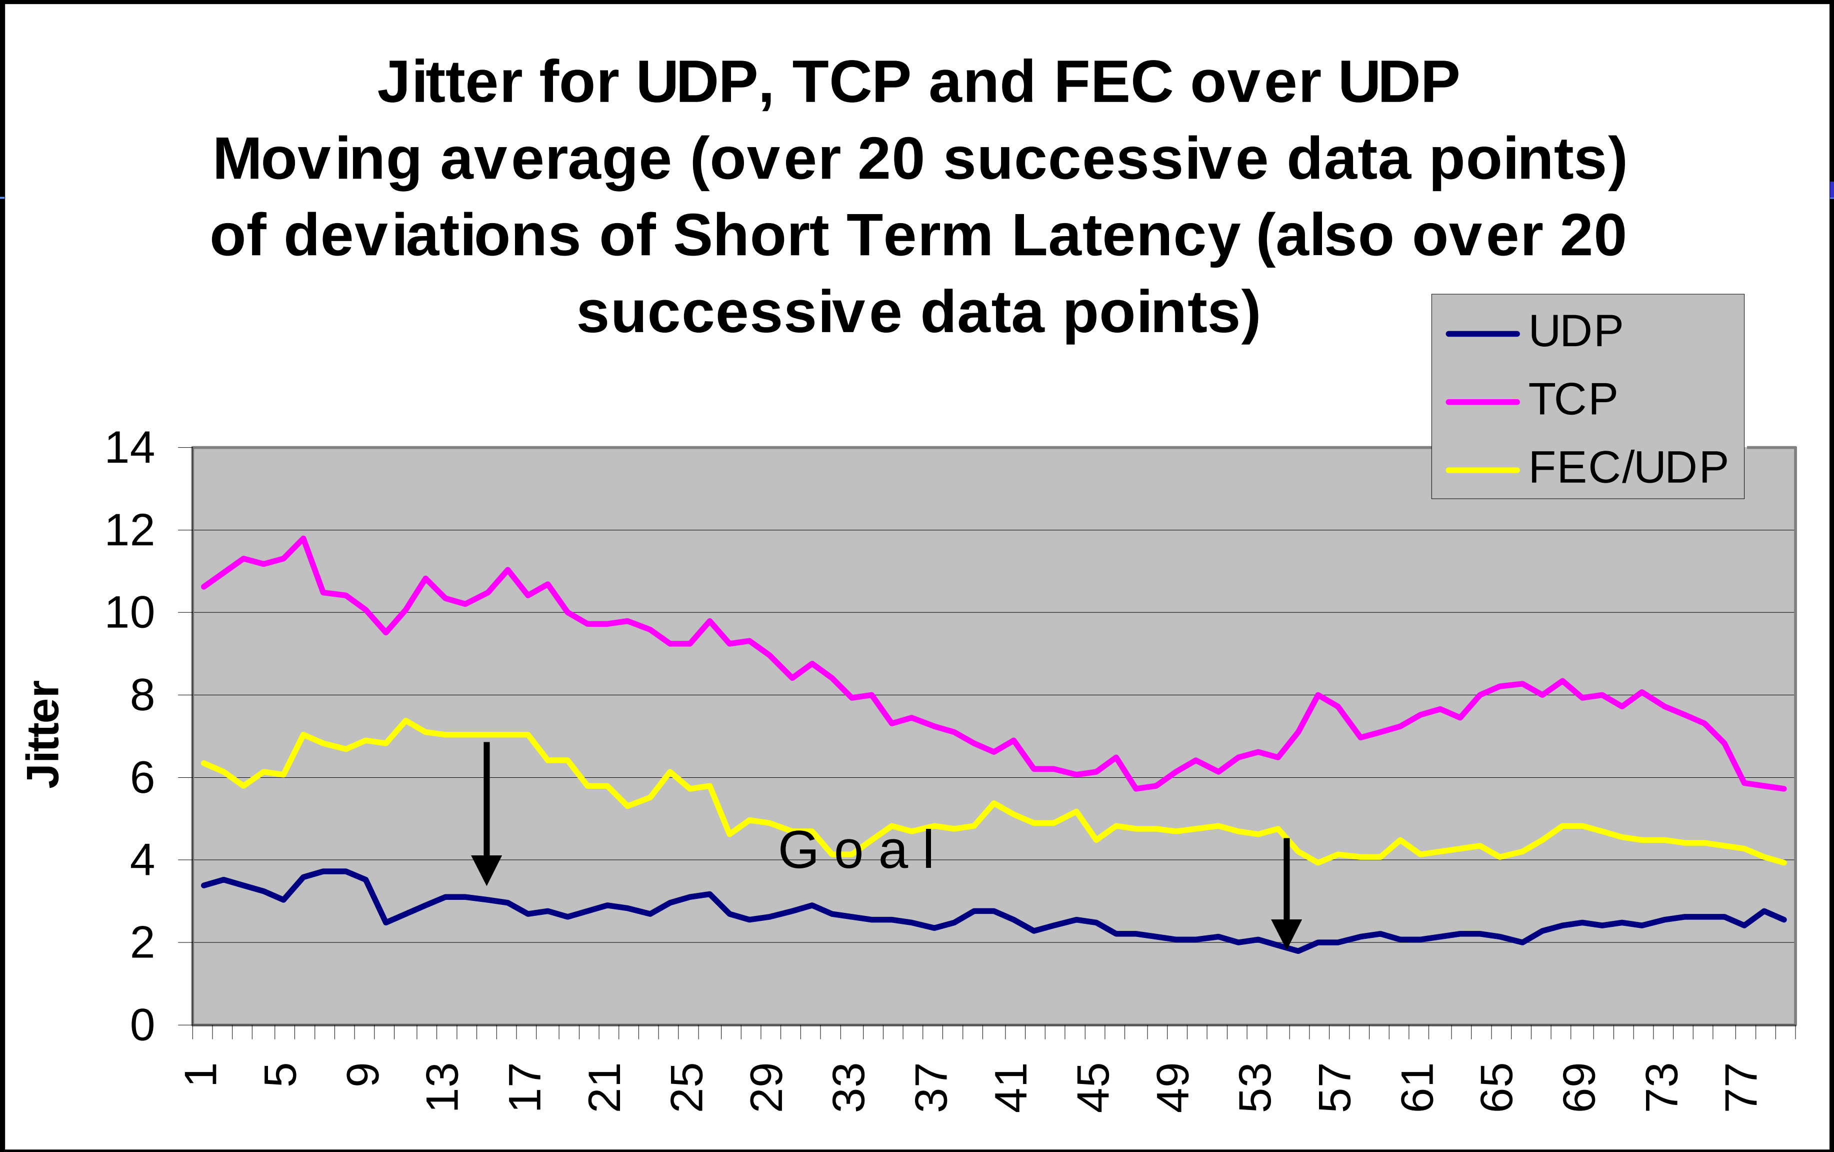
Jitter for UDP, TCP and FEC over UDP
Moving average (over 20 successive data points) of deviations of Short Term Latency (also over 20 successive data points)
UDP
TCP
FEC/UDP
Jitter
Goal
0
2
4
6
8
10
12
14
1
5
9
13
17
21
25
29
33
37
41
45
49
53
57
61
65
69
73
77

Packet Loss over UDP vs FEC/UDP

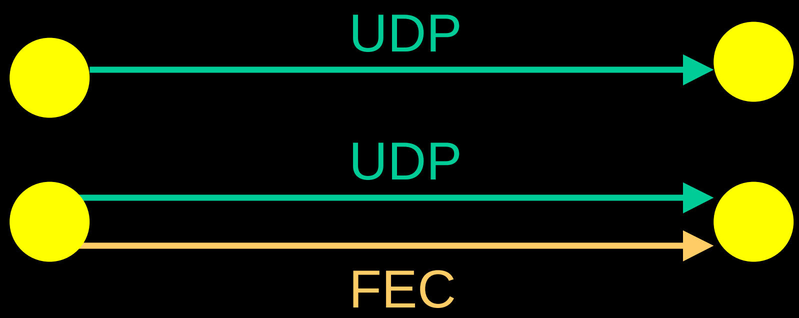

Data Rate (bits/s)	**Packet Size** (Bytes)	**Packet Loss Rate in UDP (%)**	**Packet Loss Rate in FEC over UDP (%)**
1M	128	0.4	0
1M	256	0.2	0
1M	1024	0.2	0
10M	128	30	4
10M	256	25	3
10M	1024	21	1.5

Parallel Sockets, psocket API in CAVERNsoft

Robert Grossman, Stuart Bailey, Emory Creel,
Georg Reinhart- National Center for Data Mining
Between UIC and Highway1 in Washington, D.C. over vBNS.
At the time Highway1 had a DS-3/45Mbps to vBNS

Transport Mechanism	Secs	Mb/s
Traditional Socket	96	8
Psocket 2	57	14
Psocket 3	34	24
Psocket 4	30	27
Psocket 5	26	30
Psocket 6	26	30

QoSiMoto: QoS Internet Monitoring Tool

- Kyoung Park
- Reads Netlogger data sets from file or from netlogger daemon.
- CAVE application runs on SGI, Linux and NT
- Information Visualization problem.
- How to leverage 3D.
- Averaging of data points over long traces.
- www.evl.uic.edu/cavern/qosimoto

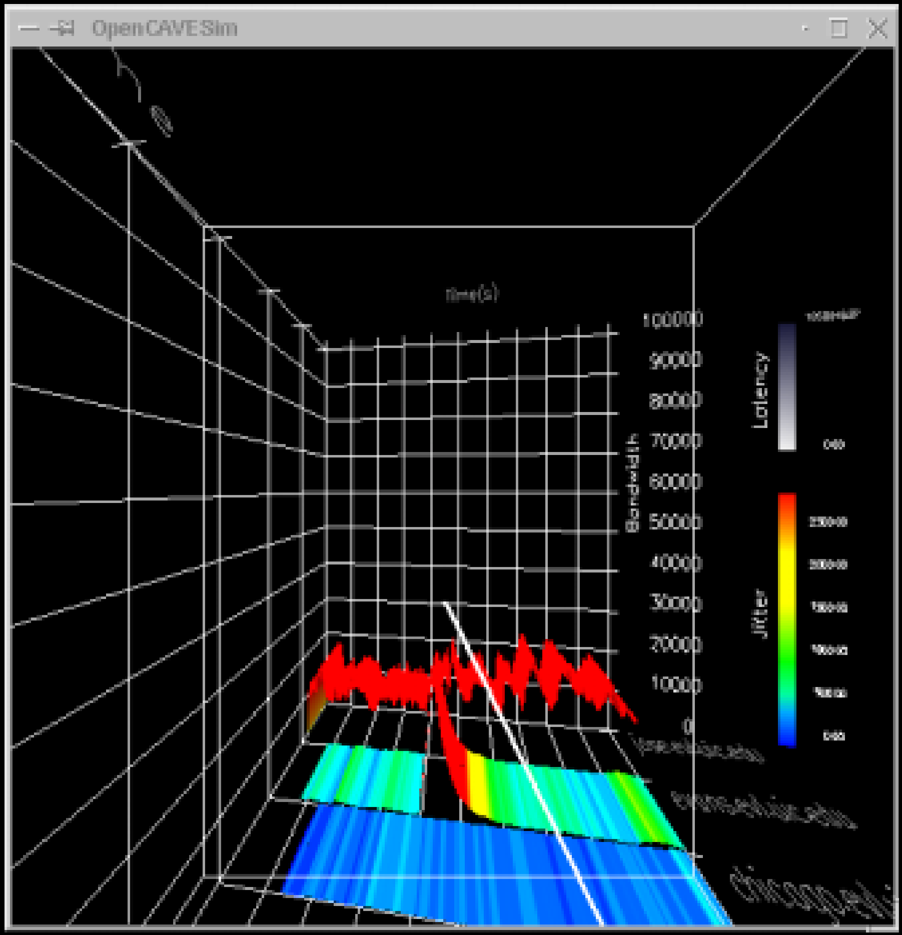

Additional Info

- Teleimmersion, CAVERNsoft, TIDE, CIBR View, QoSIMoto:
  www.evl.uic.edu/cavern
- Email:
  cavern@evl.uic.edu
- EMERGE:
  www.evl.uic.edu/cavern/EMERGE