

Enabling Technologies for Teaching an HPC Class as a Distributed Course

*Petr Holub, Ravi Paruchuri, Stacey Simmons, Tomáš Rebok,
Andrei Hutanu, Daniel Eiland and Miloš Liška*

Louisiana State University
CESNET, z. s. p. o.
Masaryk University

ViDe2007
Atlanta, 2007-03-27



Talk Overview

Introduction

Enabling Technologies

Implementation & Experiences



Distributed HPC Class – Intro

- HPC class taught by prof. Thomas Sterling at LSU
- Distributed teaching using latest high-end technologies
- Technology partners: CESNET/Masaryk Univesity and LSU
- Course participants:
 - LSU (LA)
 - Masaryk Univesity (CZ)
 - University of Arkansas (AK)
 - Lousiana Tech University (LA)
 - MCNC (NC)
 - ... others interested




Remote Lecturing

- Motivation
 - broaden course portfolio
 - increase teaching efficiency
- Without feedback
 - Lecturer is separated from students
 - Unidirectional transmission, no live interaction
- With feedback
 - enables interactions with students
 - often spans a number of physical rooms
 - lecturer is present in one of the rooms
 - media streams are transmitted to other rooms
 - media streams from all rooms are transmitted back
 - *the goal*: minimize differences for students in remote locations



Spanning Large Distances

- Time difference—e.g., mid-west USA, Czech Rep, Korea
 - USA: 19:00–20:30 (day D)
 - CZ: 1:00–2:30 (day D+1)
 - Korea: 7:00–8:30 (day D+1)
- Latency
 - light propagation in fiber comes in



Psychological Aspects

- Sound
 - focus on quality and bug-free end-to-end chain
- Video
 - high-resolution image can improve immersion feeling
- Latency for live interaction
 - both training and technologies improves on this
- Identification of who's asking (hand-rising application)
- Presence of audience



Talk Overview

Introduction


Enabling Technologies

Implementation & Experiences



Interactive Video

- Uncompressed HD video over IP
 - HD-SDI SMPTE 292M, 1080i, 4:2:2, 10 b per color plane



- 1.5 Gbps per stream, using 8,500 B Jumbo frames
- UltraGrid-based solution
 - Linux based, Centaurus HD-SDI card
 - software display (no HD-SDI monitors)
 - 10 GE NICs
 - currently merging codes to have a single implementation again
- distributed using software UDP packet reflectors

Interactive Audio

- Uncompressed audio over IP
 - currently as separate stream
 - 16 – 48 kHz sampling, 16 b quantization
 - optional FEC by redundancy
 - RAT tool
 - theoretically can sync with UltraGrid using RTCP
 - distributed using software UDP packet reflectors
 - **requires solid audio hardware and installation!!!**



Interactive Media Recording

- Need to store full-quality media for “late” participants
- Uncompressed media recording
 - 1,5 Gbps \implies 190 MBps (1.5 hr is about 1 TB)
 - storage of both video and audio network streams in a multiplexed file
 - fast disk arrays can “just” do that (RAID 0)
 - we have built our own recording/playback application
 - file indexing for fast seek capability
 - experiments with PVFS (MPI interface seems to work better than FS interface)



Other Streaming & Recording Technologies

- WebEx (<http://www.webex.com/>)
- QuickTime & VLC live streaming
- QuickTime archives
 - hand-edited
 - PiP with slides



Talk Overview

Introduction

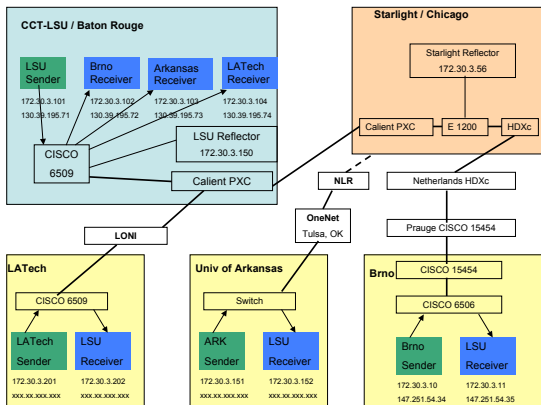
Enabling Technologies

Implementation & Experiences



Implementation

- Network setup
 - circuits meet at StarLight, uses Enlightened/NLR



- MCNC: Enlightened/NLR



Implementation

- Data distribution
 - packet reflectors running at StarLight
 - new reflectors being prepared at LSU
- Video distribution
 - 1:N distribution from LSU
 - N:1 transmission to LSU
- Audio distribution
 - N:N distribution
- Backup solution
 - AccessGrid
- Text-based chat back-channel
 - AIM groupchat (problems with Jabber gateways)



Implementation

- “Live” lower-quality streaming
 - WebEx URL is published before each class
- Supplementary materials
 - slides, problem sets
 - edited video recordings in QuickTime SD format
 - US classes share the same web
 - Czech class has its own, namely for de-synced semesters

DEPARTMENT OF COMPUTER SCIENCE & CENTER FOR COMPUTATION AND TECHNOLOGY at Louisiana State University					
High Performance Computing Concept, Method and Means					
Professor: Thomas Sterling					
CSC 7600					
Class Date at LSU	Course Number	Color	Documentation	Tools & Software	Problems Sets
	Topic	Slides	Class Videos	Lecture Notes	
Jan 10	Introduction	S000001 S000002 S000003 S000004 S000005	V000001 V000002 V000003 V000004 V000005	L000001 L000002 L000003 L000004 L000005	P000001 P000002 P000003 P000004 P000005
Jan 15	Connectivity Clusters	S000006 S000007 S000008 S000009	V000006 V000007 V000008 V000009	L000006 L000007 L000008 L000009	P000006 P000007 P000008 P000009
Jan 20	Remoteness	S000010 S000011 S000012	V000010 V000011 V000012	L000010 L000011 L000012	P000010 P000011 P000012
Jan 25	Throughput Computing	S000013 S000014 S000015	V000013 V000014 V000015	L000013 L000014 L000015	P000013 P000014 P000015
Jan 30	Parallel Computer Architecture	S000016 S000017 S000018	V000016 V000017 V000018	L000016 L000017 L000018	P000016 P000017 P000018
Feb 1	Single Node Architecture & Performance	S000019 S000020 S000021	V000019 V000020 V000021	L000019 L000020 L000021	P000019 P000020 P000021
Feb 6	Parallel Threaded Computing	S000022 S000023 S000024	V000022 V000023 V000024	L000022 L000023 L000024	P000022 P000023 P000024
Feb 9	Queue IP programming	S000025 S000026 S000027	V000025 V000026 V000027	L000025 L000026 L000027	P000025 P000026 P000027
Feb 13	Performance Measurement 1	S000028 S000029 S000030	V000028 V000029 V000030	L000028 L000029 L000030	P000028 P000029 P000030


PA177 (CSC7600) - Introduction to High Performance Computing

- Home
- Course Overview
- Course Logistics
- Why
- Prerequisites
- Course Schedule
- Support
- Resources
- Feedback

High Performance Computing

- 2017-01-10 Introduction to High Performance Computing (10:00 AM)
- 2017-01-15 Connectivity Clusters (10:00 AM)
- 2017-01-20 Remoteness (10:00 AM)
- 2017-01-25 Throughput Computing (10:00 AM)
- 2017-01-30 Parallel Computer Architecture (10:00 AM)
- 2017-02-01 Single Node Architecture & Performance (10:00 AM)
- 2017-02-06 Parallel Threaded Computing (10:00 AM)
- 2017-02-09 Queue IP programming (10:00 AM)
- 2017-02-13 Performance Measurement 1 (10:00 AM)

Implementation



Experiences

- Organizational issues
 - class time-slot scheduling (doesn't work officially for Masaryk University)
 - different semester starts at different institutions
 - recording/playback
 - requires slot for live Q/A sessions
 - problems with problem sets/test concealment
 - de-synced DST switch
- Technical issues
 - audio installation problems
 - network failures
 - manual monitoring only
 - need to deploy more reflectors as the number of participants increases



Experiences

- Audio installation problems
 - sharing components with permanent/AccessGrid installations
 - poor integrated sound cards (half-duplex, unreliable detection/operation)
 - professional sound card support in Linux



Future Challenges

- Towards more component-based integrated collaborative environment
 - having components remotely controllable (reflectors, senders, receivers, storage systems, etc.)
 - collaborative GUI for management of the system
 - continuous monitoring
 - work-flow support
- More research on user perception
 - results are supposed be presented at AccessGrid retreat in May



Conclusions

- Providing “production” service on experimental infrastructure
 - good way to push limits of the infrastructure
 - overall impression is not bad, but costs a lot of effort



Thank you for your attention!

Q?/A!

hopet@ics.muni.cz

