Software Systems Research – Portfolio Review

Dr. Nelson R. Manohar Alers

nelsonmanohar@yahoo.com
nelsonmanohar.sphosting.com/research

Outline of the Talk

- Background [10%]
- Computer-Supported Collaboration [25%]
- Dynamically Customized Web Touring [25%]
- Multimedia Computing Networking [35%]
- Wrap-Up [5%]

Education

- **B.S. in Computer Engineering**
  - **The University of Puerto Rico at Mayaguez**
- **M.S. in Computer Engineering**
  - **The University of Wisconsin at Madison**
  - Computer Architecture and Organization (Advisor Dr. Yu Hen Hu)
- **M.S.E. in Industrial Engineering**
  - **The University of Michigan at Ann Arbor**
  - Statistical Quality Control/Information Systems (Advisor Dr. Dan Teichroew)
  - From January 1991 to December 1992*
- **Ph.D. in Computer Science and Engineering**
  - **The University of Michigan at Ann Arbor**
  - Software Systems Research (Advisor: Dr. Atul Prakash)
  - From January 1991* to May 1997

Work Experience

- **Senior Technical Associate**
  - **@ AT&T Bell Laboratories, Naperville, Illinois**
  - From June 1986 to August 1986
- **Member of the Technical Staff I**
  - **@ AT&T Bell Laboratories, Naperville (Indian Hill), Illinois**
  - From June 1987 to March 1991*
- **Member of the Technical Staff I**
  - **@ AT&T Bell Labs field work at La Telefonica's Spain AIN**
  - From September 1990* to December 1990*
- **Member of the Technical Staff**
  - **@ Bell Communications Research, Piscataway, New Jersey**
  - From June 1992 to August 1992
- **Research Staff Member**
  - **@ IBM Thomas J. Watson Research Center**
  - From May 1997 to December 2001*

Research Traversal

- **Computer Architecture and Organization**
  - **University of Wisconsin - Madison (87-88)**
- **Advanced Intelligent Networks**
  - **AT&T Bell Labs (S'86, S'87, 88-91)**
- **Statistical Quality Control & Systems Engineering**
  - **University of Michigan IOE Department (91-92)**
- **Distributed Computing and Distributed Systems**
  - **Bellcore (S'92);**
  - **University of Michigan EECS Department (92-93)**
- **Collaborative Systems**
  - **University of Michigan EECS Department (93-94)**
- **Collaborative Multimedia Systems**
  - **University of Michigan EECS Department (94-97)**
- **Multimedia Computing Networking**
  - **IBM T. J. Watson Research Center (97-01)**

5

Selected Publications

- **"A Framework for Programmable Multimedia Overlay Networks",**
  - N. R. Manohar, A. Mehra, M. H. Willebeek-LeMair and M. Naghshineh, in IBM Journal of Research and Development, Special Issue on Digital Video, 43(4), July/August 1999.
- **"Applying Statistical Process Control to the Adaptive Rate Control Problem",**
  - Manohar, Nelson R.; Willebeek-Lemair, Marc H.; Prakash, Atul, in Proceedings of Multimedia Computing and Networking Conference, pp. 45-60, San Jose, CA, January 1998.
- **"A Flexible Architecture for Heterogeneous Replayable Workspaces",**
  - Nelson R. Manohar and Atul Prakash, in Proceedings of the Third IEEE Int'l Conference on Multimedia Computing and Systems, pp. 274-278, Hiroshima, Japan. June 1996.
- **"Dealing with Synchronization and Timing Variability in the Playback of Interactive Session Recordings",**
  - Nelson R. Manohar and Atul Prakash, in Proceedings of the Third ACM Int'l Multimedia Conference, pp. 45-56. San Francisco, CA, November 1995.
- **"The Session Capture and Replay Paradigm for Asynchronous Collaboration",**
  - Nelson R. Manohar and Atul Prakash, in Proceedings of the Fourth ECSCW Conference, pp. 149-164. Stockholm, Sweden, September 1995.

6

Other Publications

- **"Streaming and Synchronization of Re-executable Content",**
  - N. Manohar and A. Prakash, *unpublished*, 1998.
- **"Design Issues on the Support of Tools and Media on Replayable Workspaces",**
  - N. Manohar and A. Prakash, CSE-TR-304-96, Dept. of EECS, Univ. of Michigan, September 1996.
- **"Design Considerations in Building a Distributed Collaboratory",**
  - A. Prakash, F. Jahanian, R. Hall, N. Manohar, A. Mathur, C. Rasmussen, H. Shim, T. Weymouth, G. Wu, D. Atkins, R. Clauer, and G. Olson, School of Information, Univ. of Michigan, Feb. 1995.
- **"Statistical Quality Control and Software Productivity."**
  - N. Manohar, Quals Report, (research work under Dr. Daniel Teichroew), Dept. of IOE, Univ. of Michigan, May 1992.
- **"The DCIS6 Finite State Machine Tables",**
  - Nelson R. Manohar-Alers, AT&T Bell Laboratories, Int'l 5ESS Features Development Department, Internal Memorandum, August 1989.
- **"The Computer Architecture of VLSI Digital Signal Processors",**
  - MSEE Thesis/Report, Nelson R. Manohar-Alers, Department of ECE, Graduate Engineering Library, University of Wisconsin-Madison, August 1988.

7

Intellectual Property (IP) Activity

- **IP Training:**
  - Trained with **IBM Master Inventors** Mr. Leon Lumelsky and Dr. Philip S. Yu
- **IP Performance**:
  - **Principal inventor** (and principal inventor-in-training) on seven patents.
  - **Eight USPTO patent filings**, **seven** successfully granted.
  - **Two IBM Invention Plateaus** achieved.

8

Selected Patents

- **6,572,662: dynamic customized web tours…**
  - related to: tour data mining, tour authoring, like-minded touring of multiple websites, token-based control of traversal projections over web-tours, touring clients,…
- **6,529,950: policy-based QoS negotiation…**
  - related to: brokering framework for distributed resource management, etc…
- **6,516,350: self-regulated resource management…**
  - related to: autonomous (self-regulated) distributed resource management integrating traditional demand-shaping and capacity-shaping mechanisms …
- **6,466,980: capacity shaping of distributed resources on an internet environment…**
  - related to: replication management, QoS, and capacity-shaping of a network's resources (e.g., capacity-follows-demand management distributed resource management policy), etc…

9

Complementary Patents

- **6,463,454: integrated load distribution and resource management …**
  - related to: replication and capacity policies for distributed resource management, etc…
- **6,460,082: service-oriented resource signatures…**
  - related to: low overhead resource management and measurements policy for distributed servers, resources, capacity, objects, etc…
- **6,377,996: seamless live streaming handoffs …**
  - related to: handoff of live multimedia streaming across servers, "virtual sockets", migration transparency, etc…

10

Software Systems Research Map

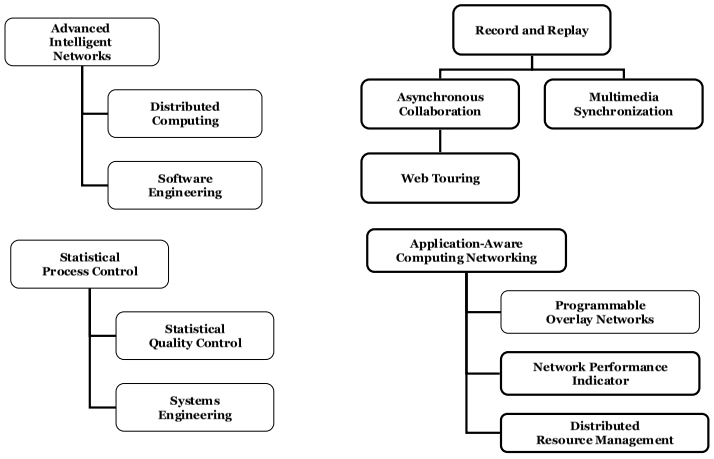

11

Outline of the Talk

- **Background [10%]**
- **Computer-Supported Collaboration (Groupware) [25%]**
  - Record and Replay (by-Re-execution) Paradigm
  - Multimedia Synchronization
- **Dynamically Customized Web Touring [25%]**
- **Multimedia Computing Networking [35%]**
- **Wrap-Up [5%]**

12

Interactive Asynchronous Sharing of Computer-Supported Workspaces

(via the Record and Replay of Re-executable Content)*

By Nelson R. Manohar
Dissertation Supervisor: Prof. Atul Prakash
Department of Electrical Eng. and Computer Science
University of Michigan at Ann Arbor

Research sponsored* by
NSF Grant ECS-94-22701 (MedCollab),
NSF Grant IRI-92-16848 (UARC),
University of Michigan's Rackham Merit Fellowship

13

Computer-Supported Collaborative Work

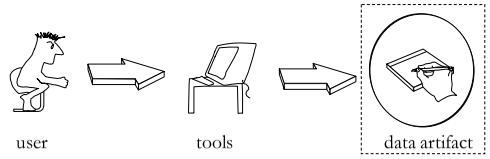

- **Computer-Supported Collaborative Work (CSCW)**
  - synchronous collaboration (i.e., working simultaneously)
  - asynchronous collaboration (i.e., working at different times)
- **Asynchronous Collaborative Work**
  - different times
  - between one or more users
  - iterative refinement of a collaboration artifact

14

Asynchronous Collaboration Problem

- **The Asynchronous Collaboration Problem:**
  - the process of reaching a conclusion may contain as much information as the conclusion itself.
  - ways to capture this "intra-task" content and make it accessible to collaborators are desirable.
- **Motivating Uses**
  - University of Michigan
    - **UARC Collaboratory (NSF Grant ~6.0M)**
    - space science workspaces
  - University of Michigan
    - **Medical Collab (NSF Grant ~1.6M)**
    - MRI/radiology workspaces
  - Applications
    - courseware, training, debugging, demos, etc.

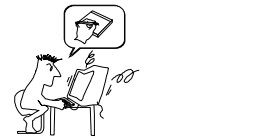

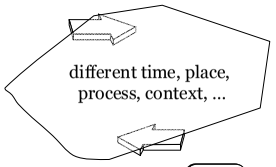

15

Capturing Intra-Task Content: Asynchronous Sharing of Workspaces

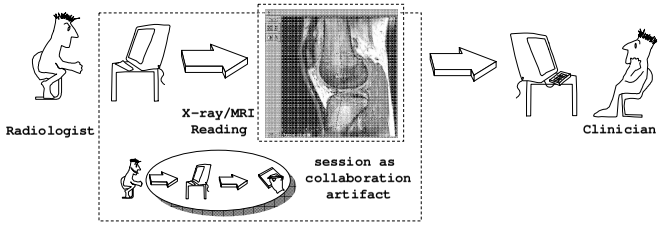

- **Goal: capture, reuse, interact, and edit an application session as just another object or artifact**
- **What components of the workspace should we capture?**
  - we want to **hear** the session ---> **voice-annotation** component
  - we want to **see** the session ---> visual component BUT
  - we want to **interact** with the session ---> **computational** component

16

Record/Replay of Application Inputs

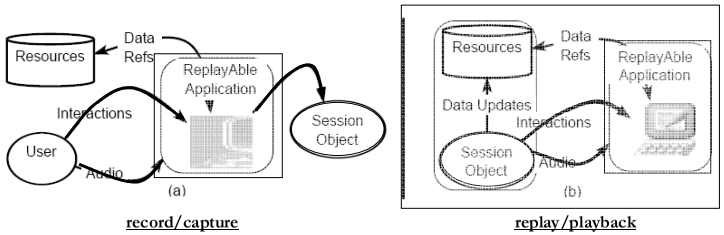

record/capture replay/playback

- **inputs modeled as temporal media streams**
  - inputs: audio-annotation and user interactions
  - during record, inputs to application(s) are captured
  - during replay, inputs are re-executed by ***replayable*** application(s)
- **captured into an interactive session object**
  - new collaboration artifact: persistent, active object
  - benefits: small size and lossless fidelity (through re-execution)

17

Session Capture and Replay: Research Problems

- **Later-time reuse of application workspace**
  - what features are needed to collaborate?
  - how to make applications ***replayable***?
  - how to make ***replayable*** sessions portable?
- **Media management and performance**
  - how is playback (i.e., replay by re-execution) affected by replay on conditions such as different load conditions?

18

Session Capture and Replay: Representation of Session Objects

- **storage of a session's streams**
  - representation & portability of streams and resources
  - time-stamp/capture of events/state

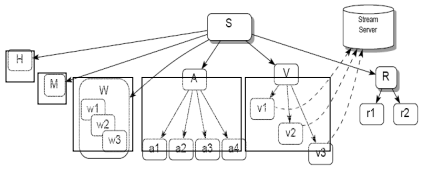

- **replay of a session's streams**
  - synchronization of asynchronous events
  - stream prefetching
  - scheduling under load conditions

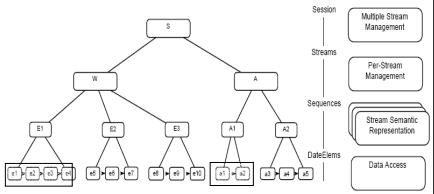

19

Replayable Application Object

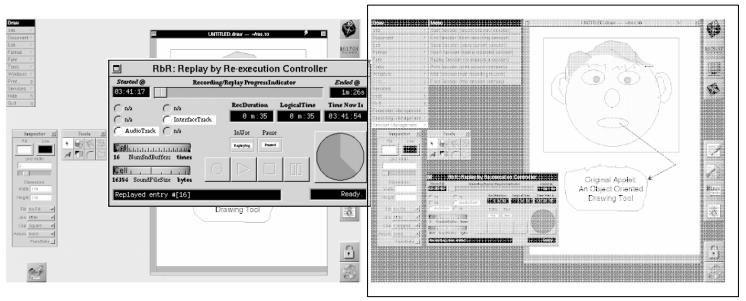

- **capture and replay of interactive session with application**
  - integrated replay of audio (**DSP**) and window (Display Postscript (**DPS**)) streams
  - API for other data streams
- **delivers replay-awareness to object-oriented (NeXTStep, ObjC) applications**
  - subclassing from ***replayable*** object class
  - record, replay, store, open, pause, continue, [browse, fast forward, fast replay, etc.]
  - transparent access to infrastructure services (next slide)

20

Replayable App's Infrastructure

- **Record Extensions**
  - stream sampling points
  - LTS services
  - data management services
  - persistency services
- **Replay Extensions**
  - stream intake points
  - prefetching services
  - scheduling services
  - synchronization services
  - measurements services
- **Playback Variability**
  - playback load/platform
  - re-execution (DPS, DSP)
  - record/replay overheads
  - timing services
  - operating system overheads

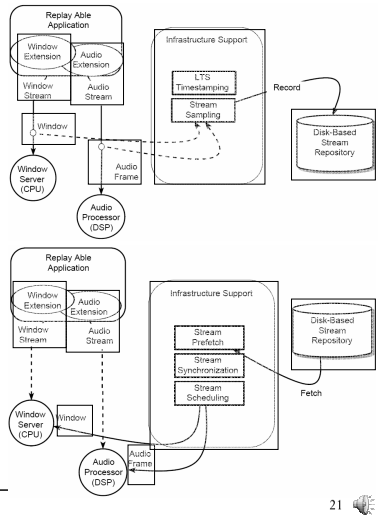

Multimedia Synchronization

"Dealing with Synchronization and Timing Variability in the Playback of Interactive Session Recordings"

by Nelson R. Manohar and Atul Prakash, in Proceedings of the Third ACM Int'l Multimedia Conference, pp. 45-56. San Francisco, CA, November 1995.

Multimedia Synchronization

- **Early experience with replay of session objects:**
  - intra-stream continuity is critical
    - playback <u>continuity</u> of audio is critical
      - that is, no gaps on the playback of continuous media
    - playback of re-executable event streams must be <u>smooth</u>
      - that is, no abrupt or sudden updates to the application state
  - while the following constraints became clear:
    - re-execution likely to occur on <u>different workstations</u>
    - re-execution likely to occur on <u>different load conditions</u>
    - re-execution likely to occur on <u>variable load conditions</u>
  - how is replay (i.e., replay by <u>re-execution</u>) affected by playback conditions?
    - what can we do to deal with that?

The Multimedia Requirements

Continuous Media (Audio-DSP) Stream		Re-executable Event (Window-DPS) Stream	
synchronous		**asynchronous**	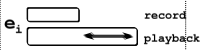
stateless	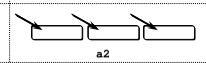	**stateful**	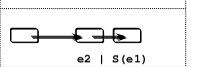
periodical/ continuous		**aperiodical/ discrete**	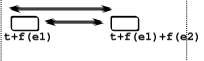

Heterogeneous Media Integration Problem:

- research heterogeneous media integration mechanisms
  - (e.g., synchronization, scheduling, storage, prefetching)
- for integrating fine-grained asynchronous (re-executable) events and continuous media

Heterogeneous Media Integration

- **Approach**
  - **model** both streams as asynchronous media where:
  - **re-execution time**: $t(e_i) = E(e_i) + f_0(e_i)$
  - **inter-event delay** $\Delta_{i,i+1}$
    - **zero** for continuous media
    - **variable** for asynchronous media
- **Validity**
  - even **continuous media** on a dedicated processor is:
    - **biased** (due to inherent bias on operating systems' timing services)
    - **asymmetrical** (between reported playback and record event duration)

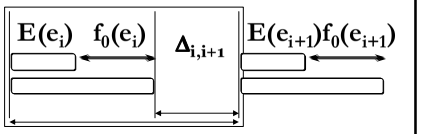

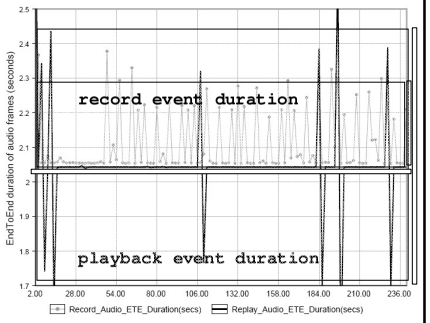

25

Synchronization Mechanism

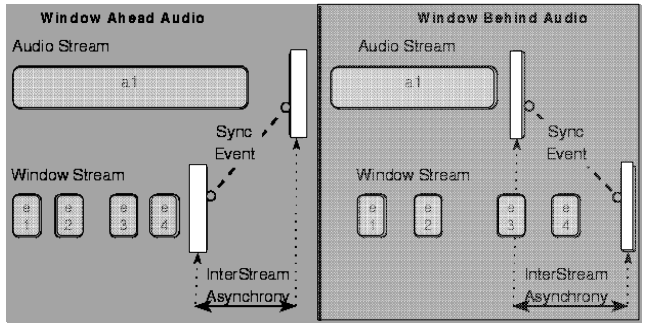

- **synchronization mechanism**
  - synchronization events S_i (W → A)
  - master/slave synchronization model (window slaved to audio)
  - slave-initiated synchronization operations

26

Intuitive Look at Adaptive Mechanism: Time Compression and Expansion

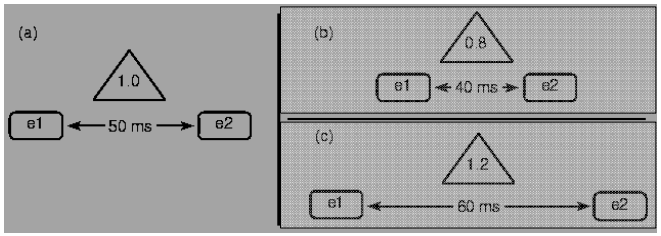

- **inter-event delay time Δ**
  - idle schedule time between two consecutive events (e_i, e_{i+1}) in a stream
- **compensation factor ω**
  - compensation to the inter-event delay time
  - based on statistical process control (SPC), for detection of inter-stream asynchrony trends on the replay of a stream

27

ω - Compensated Scheduling Intervals

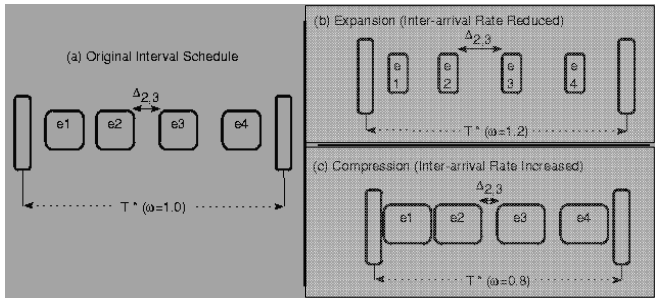

- **periodical ω-compensation of asynchrony trends**
  - loose supervisory controls, formulated independently for each re-executable stream
  - updated only once per interval, constant for all events in the same scheduling interval

28

SQC-based Adaptive Scheduling

(compensation of long term asynchrony trends)

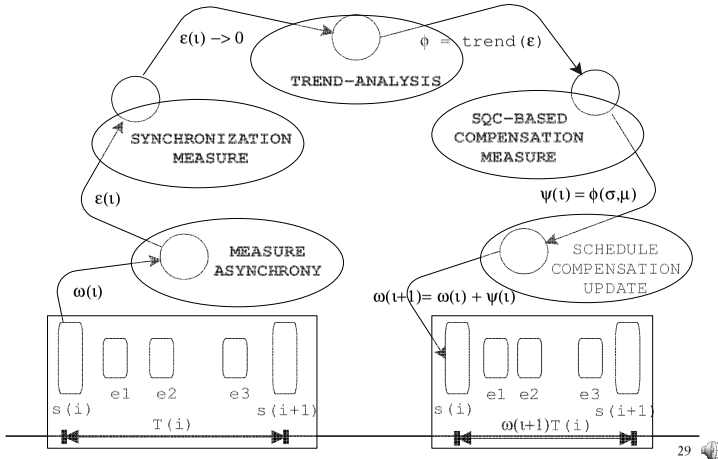

29

Performance of SQC-Adaptive Scheduling (P3): fixed control limits, ω-scheduling intervals

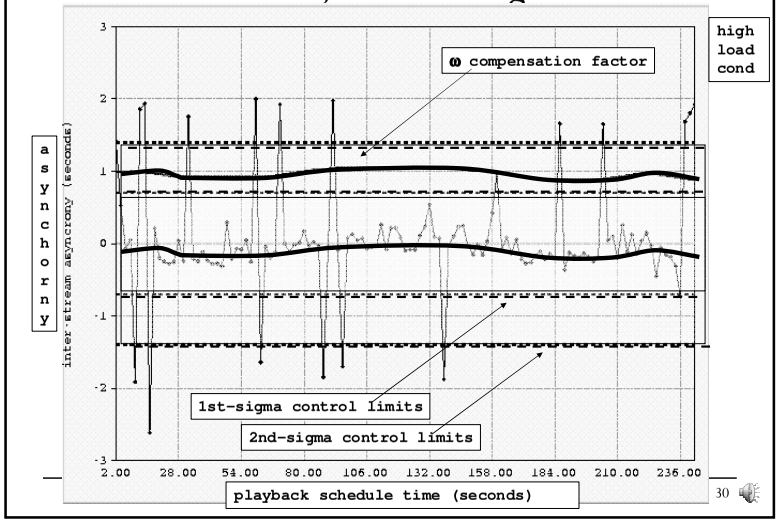

30

Comparative Protocol Performance

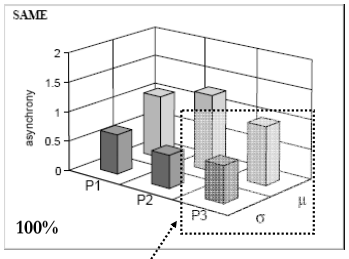

- **synch. protocols**
  - P1: timing corrections
  - P2: P1+1-way syncs
  - **P3: adaptive P2**
  - P4: P1+2-way syncs

- **desirable characteristics**
  - low mean: μ(async)
  - low std-dev: s(async)
  - across all load conditions
    - 25%, 100%, 400%

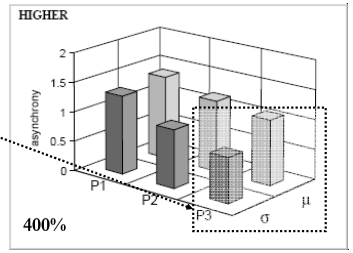

31

Observations on Performance of Protocol P3

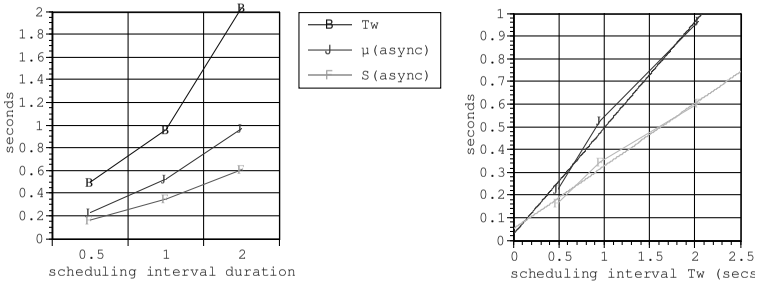

Findings:

- **μ(async) about $^{T}/_{2}$ (i.e, one-half scheduling interval)**
  - $\mu = 0.47\,T_{\omega} + 0.03$ $(R^2 = 0.99)$
- **σ(async) about $^{T}/_{3}$ (i.e, one-third scheduling interval)**
  - $\sigma = 0.28\,T_{\omega} + 0.05$ $(R^2 = 0.99)$

Research Contributions

- **delayed-sharing of sessions/workspaces**
  - a complementary paradigm for asynchronous collaboration for intra-task content capture through re-executable record and replay of an application workspace
- **protocols for scheduling and synchronizing**
  - adaptive mechanism
    - time compression and expansion (inter-event delay as degree of freedom)
    - handling of asynchrony trends
    - playback on significantly different load conditions
  - heterogeneous streams
    - fine-grained asynchronous re-executable events wrt. continuous media
    - generalized to n-ary relationships (multiple media, protocols, applications)
- **introduction/application of statistical process control**
  - sound indicators for "long term" process performance
  - application-awareness (process requirements)

33

Related Work

- **collaborative systems**
  - artifact-based collaboration (workflow) systems:
    - Prep, ObjectLens, g-Ibis, etc.
  - shared windows, screen camcorders:
    - Xtv, Xmx, Xtrap, x-teleporting, QTC, ScreenCam, etc.
- **distributed event simulations**
  - tight causal event ordering, no synchronization constraints
- **multimedia authoring systems (synchronization)**
  - type 1: tight synchronization, strong continuity, but not of asynchronous events
  - type 2: coarse-grain asynchronous events with tight synchronization

34

Outline of the Talk

- **Background [10%]**
- **Computer-Supported Collaboration (Groupware) [25%]**
  - Record and Replay Paradigm
  - Multimedia Synchronization
- **Dynamically Customized Web Touring [25%]**
- **Multimedia Computing Networking [35%]**
- **Wrap-Up [5%]**

35

Dynamic Customized Web Tours

- **Research work done at**
  - **IBM Thomas J. Watson Research Center** by Dr. Nelson R. Manohar, Dr. Philip S. Yu, and Dr. Marc H. Willebeek-Lemair
- **U.S. Patent 6,572,662**
  - related to: *like-minded touring* of multiple websites, *token-based shaping* of web-tours, touring servers, tour authoring, etc…

36

What Is A Web Tour?

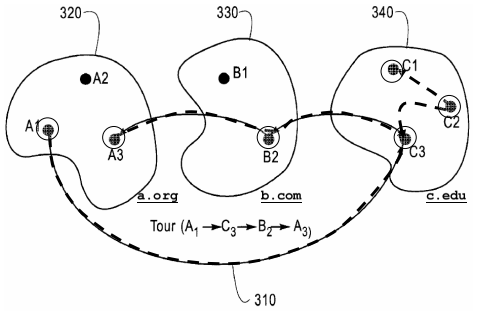

- **Merriam Webster**
  - 2 a : a journey for business, pleasure, or education often involving a series of stops and ending at the starting point; also : something resembling such a tour ² a tour of the history of philosophy³
  - **travel around**: expedition/sightsee; **trip**: exploration/outing

37

Web Touring – Basic Idea

- **From: web touring**
  - time-controlled hopping from web object to web object
  - pre-authored multimedia presentation on the web
- **To: dynamically customized web touring**
  - dynamic touring content control based on some parameter → changes over:
    - tour representation
    - touring control and presentation
    - touring experience/intelligence

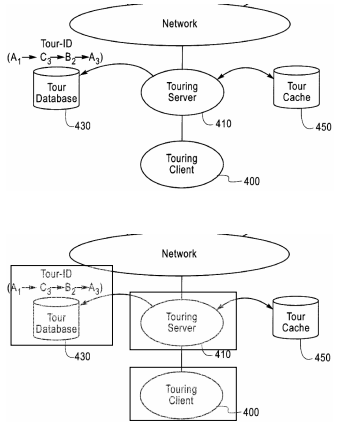

38

Motivating Research Exploration

- **Web Touring is form of an Asynchronous Collaborative Session**
  - **Session Authoring and Replay**
    - **platform independence**
    - **handling of resource references**
  - **Touring Session Server vs. Standalone Touring App**
    - **proxy approach, integration point**
    - **sharing of sessions across touring clients**
  - **Access to Session Intra-task Content**
    - **visualization of (touring) sessions – touring maps**
    - **browsing of (touring) sessions – database access**
    - **querying of (touring) sessions – token projections**
  - **Acquisition of Collaborative Session Intelligence**
    - **tour metrics applied toward authoring refinement of tour**
    - **data-mining like-minded touring experiences/decisions**
- **Motivating Uses (Where is this useful?)**
  - **e-business (catalogs shopping), distance learning (courseware), archive exploration (museums)**

39

Touring Architecture

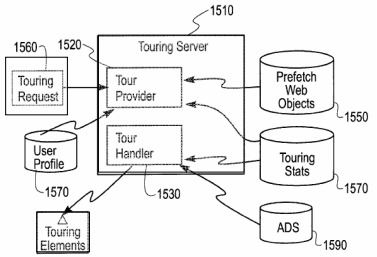

- **Touring Server**
  - provider (prefetching, versioning, database, etc.)
  - handler (scheduling, dispatching, updating, etc.)
- **Touring Client**
  - subscriber (interactions, touring requests, token management, etc.)
  - handler (tour element handling, presentation, etc.)

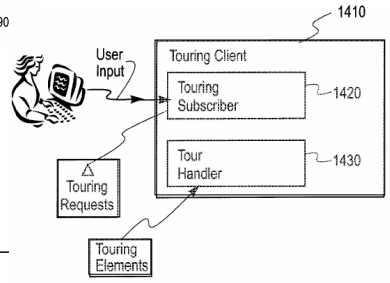

What is a tour? - Revisited

- *journey for business, pleasure, or education often involving a series of stops; … exploration; sightsee;*
- **ordered traversal of touring elements**
  - touring element represents a **tour-stop**
  - each touring element is composed of
    - **touring operator**
    - **one or more operands** (temporal URLs)
    - **presentation parameters**
- **each temporal URL**
  - temporal URLs represent **sightseeing** at the tour-stop
  - each temporal URL is composed of
    - **(data): URL** to a web resource,
    - **(control): domain-aware (visitation) tokens**

41

Tour Specification

- **Specification of Traversal Sequences and Control**
  - sequential, parallel, multiple site, dynamic, etc.
  - optional tour paths, tour fork points, tour meet points
- **Specification of Presentation Parameters**
  - presentation duration, scheduling, synchronization, etc.
  - tour visitation tokens
  - caching, persistency, versioning, plug-ins, etc.
- **Specification of Data Mining Points**
  - tracking user decisions (anchor points)

42

Specification of Traversal Sequences

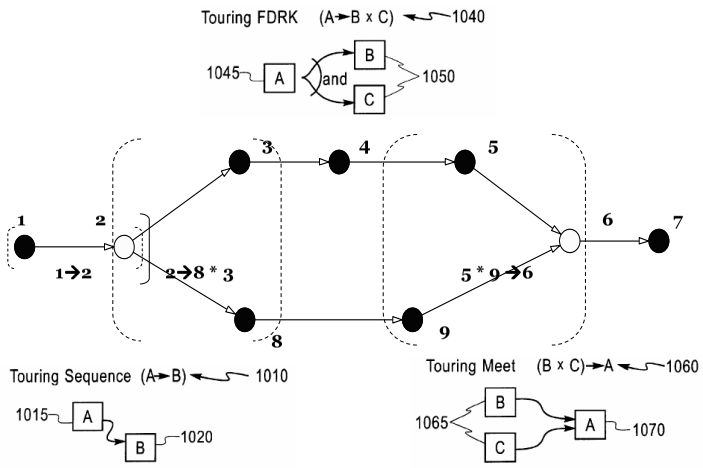

43

Anchor Points and Tour Options

- **Tracking explicit user decision making and preferences (intra-task knowledge)**
- **Tour options**
  - (A → B + B C)
  - comparative touring
- **Tour anchors**
  - (A → B + C)
  - decision making points
  - provision for side-touring interactivity during the tour

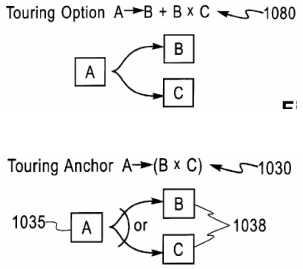

44

Tour Representation

- **URL Resource Reference**
- **Touring Tokens**
  - Application/Domain Aware
  - Server/Client/User-Controllable
- **Temporal URLs**
  - URL + Tokens
- **Touring Operators**
  - Sequence, fork, meet, etc.
- **Touring Elements**
  - A touring operator and its operands (temporal URLs)
- **Tour**
  - Collection of touring elements
- **Touring Client**
  - Tour evaluated at stateful conditions (e.g., token bag)

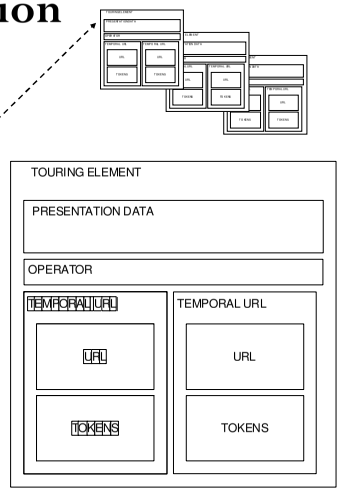

Tour Representation Example

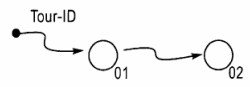

- **Sequence**
  - (01 → 02)
- **Tokens**
  - (Shirt, Shirt)
- **Duration:**
  - (7 minutes)

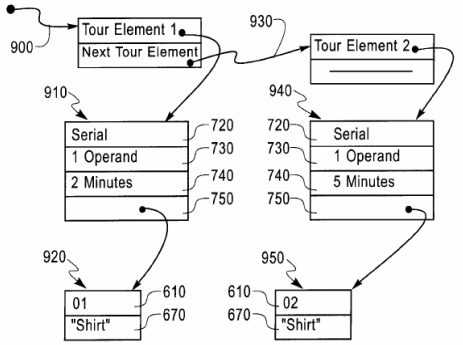

Tokens: Touring Control Layer

- **Touring State**
  - captured through user-visible touring token bag
  - *Sort-of* "traversal memories/souvenirs"
  - Also, as rating points
- **User-Control of Touring**
  - enable/disable tokens
  - add/drop tokens

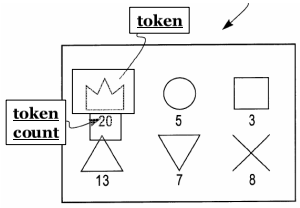

- **Example:**
  - **TOKEN DOMAIN**: Ancient Empires
  - **TOKEN BAG**: [**Greeks**: Parthenon, Theater; **Phoenicians**: Trading, Shipping; **Sumerians**: Urban Planning, Masonry; **Egyptians**: Sculpture, Unification; **Indians**: Aryan Invasion; **Romans**: Law; **Semitic**: Monoteism]

Token-based Touring Projections

14 minute presentation
2 paths
8 nodes

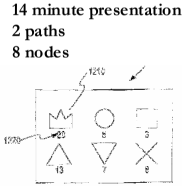

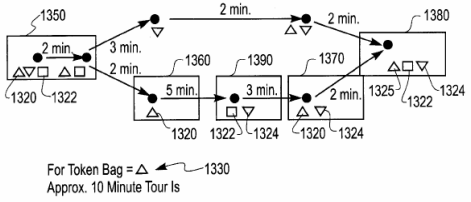

11 minute presentation
1 path
5 nodes

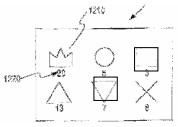

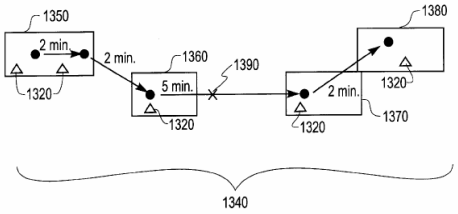

Data Mining Touring Experience

- **Like-minded Touring Statistics Over**
  - touring elements
  - anchor points
  - touring tokens
- **Like-minded Touring Suggestions**
  - path/decision-based like-minded
  - token-based like-minded

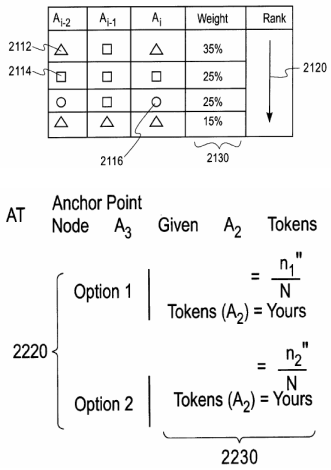

Touring Presentation

- **replay control**
  - forward, back, stop, etc…
- **presentation control**
  - java, frames,
- **content control**
  - token bag
- **user feedback**
  - statistics, state
- **data**
  - touring elements

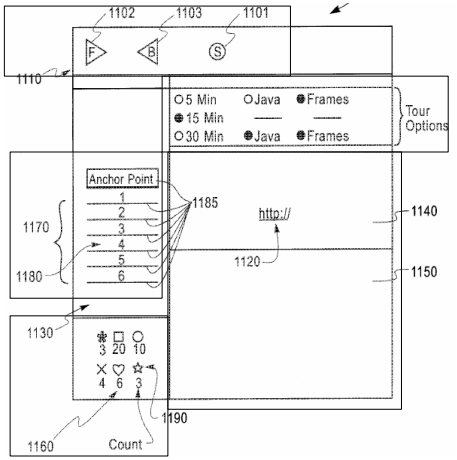

Research Contributions

- **Dynamically Customizable Web Touring**
  - token bags – application/domain abstraction (sort-of touring "souvenirs")
    - accumulated over touring experience
    - as well as through direct user control over token-based state
  - touring projections – tuning control transformations of touring content
    - controlled through user management of touring bags
  - collaborative access to touring intelligence
    - comparative (simultaneous multiple tour element) exploration
    - data mining of touring experience (anchor points)
    - like-minded exploration suggestions over touring content

Related Work

- **Web Touring**
  - multiple window/site comparative touring
    - e.g., synchronization of objects/sites
  - dynamic touring visualization projections
- **Multimedia Presentation and Authoring**
  - self-contained (closed collection) static presentations
    - handling of external resource references
    - handling of ephemeral presentation elements
    - handling of versioning and caching
- **Intelligent Training Systems**
  - self-contained dynamic presentations over closed collection
    - i.e., even though some with adaptive user interface and like-minded content exploration
  - no provision for side-touring or user-control of presentation state
    - e.g., impact of side-touring or token management over touring state

Outline of the Talk

- **Background [10%]**
- **Computer-Supported Collaboration (Groupware) [25%]**
- **Dynamically Customized Web Touring [25%]**
- **Multimedia Computing Networking [35%]**
  - **Network Performance Envelope Estimation**
  - Distributed Resource Management
- **Wrap-Up [5%]**

53

Multimedia Computing Networking –

"Network Performance Envelope Estimation"

Work at IBM Thomas J. Watson Research Center

"Applying Statistical Process Control to the Adaptive Rate Control Problem", by Manohar, Nelson R.; Willebeek-Lemair, Marc H.; Prakash, Atul, in Proceedings of Multimedia Computing and Networking Conference, pp. 45-60, San Jose, CA, January 1998.

54

Introduction and Motivation

- **Goal**
  - integration meaningful feedback (into the network) about the requirements and performance of adaptive multimedia applications
- **We would like to explore**
  - mechanisms for characterizing long-term network variability (e.g., delay, bandwidth, buffer size)
  - mechanisms for inducing application requirements into adaptive multimedia networking
- **We would these mechanisms**
  - to be robust and easy to implement

55

Media Adaptation Envelope

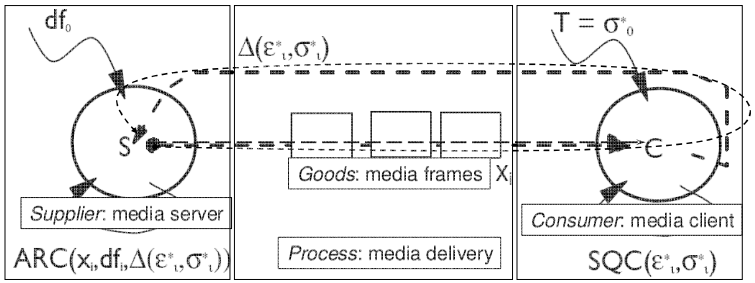

- **Delivery Process**
  - network delivery and/or media processing processes
- **Adaptive Rate Control Problem (ARC)**
  - ARC (media, degree of freedom, feedback)
- **Statistical Quality Control (SQC)**
  - SQC (process indicator, process variability)

56

SPC-ARC Multimedia Applications

- **Media Integration**
  - generation and enforcing of process (long-term) performance envelope
- **Adaptive Coding**
  - application-oriented
  - media-independent
  - reactive to long-term (persistent) trends

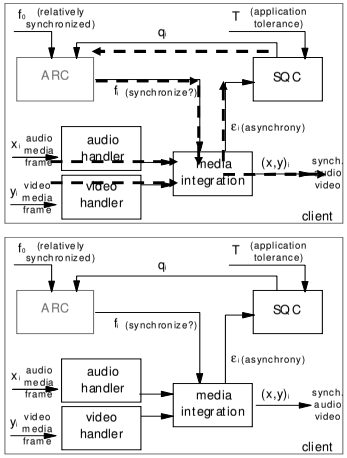

57

Building Reliable Network Performance Indicator

- **Applications of Statistical Process Control (SPC) to Adaptive Rate Control (ARC)**
  - Complementary (Envelope) Approach
  - Adaptive Media Streaming & Synchronization
  - Statistical Quality Control Adaptation Principles
- **Application of SPC to Multimedia Computing Networking**
  - SPC Is By Definition Application-Awareness
  - Exposes Long-Term vs. Short-Term Variability
  - Exposes Performance Envelope
  - Application: Indicator Stability Monitor
- **Conclusions**

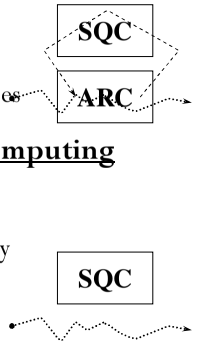

58

Network Performance Envelope: Long-Term vs. Short-Term Variability

- **From:**
  - fast indicators
  - low setup cost
  - short-term smoothing
  - stateless
  - no outlier analysis
- **To:**
  - forecast-strength
  - setup cost
  - process performance
  - process state
  - outlier analysis

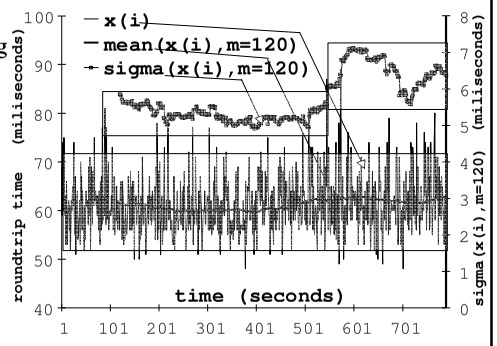

59

SPC-based Performance Envelope

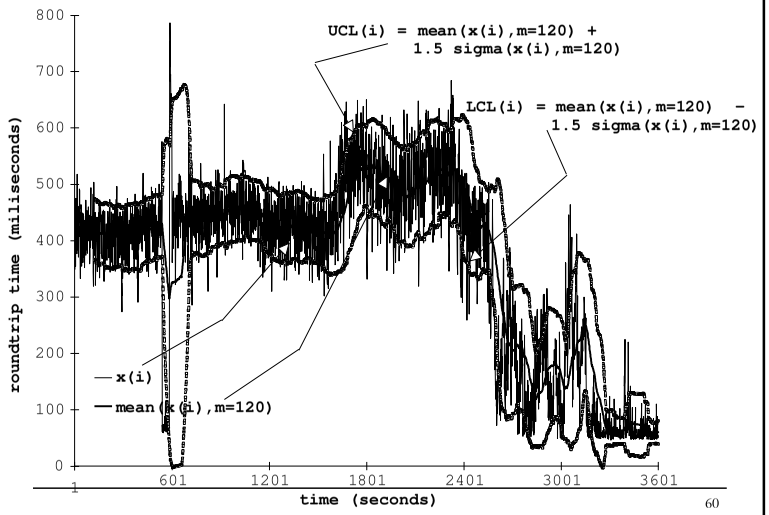

60

SPC-based Envelope Formulation

- Stationarity Test
- Robust Indicators/Forecast/State
- Smoothed Process Indicators (scale, sampling, type)
  - UWMA(x_i,m) smoother
  - $\mu(i,m') = \mu(x_i .. x_{m'-i})$
  - $\mu(i,m) = \mu(x_i .. x_{m-i})$
  - $\sigma(i,m) = \sigma(x_i .. x_{m-i})$
- Hypothesis Testing (comparison-of-means)
  - H0 : $\mu(i,m') = \mu(i,m)$
  - Z0 = $\mu(i,m') - \mu(i-{}^{m}/_{2}, m)$
- Confidence Interval (forecast estimation)
  - if $|Z0| < k * \sigma(i,m)$
  - then $mon_i = mon_{i-1}$
  - else $mon_i = \mu(i,m)$
  - $mon^{*}_{i+1} = mon_i$

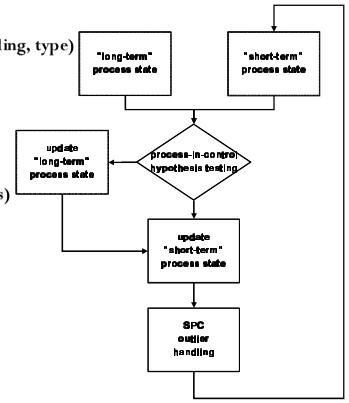

61

Long-Term Stability Monitor

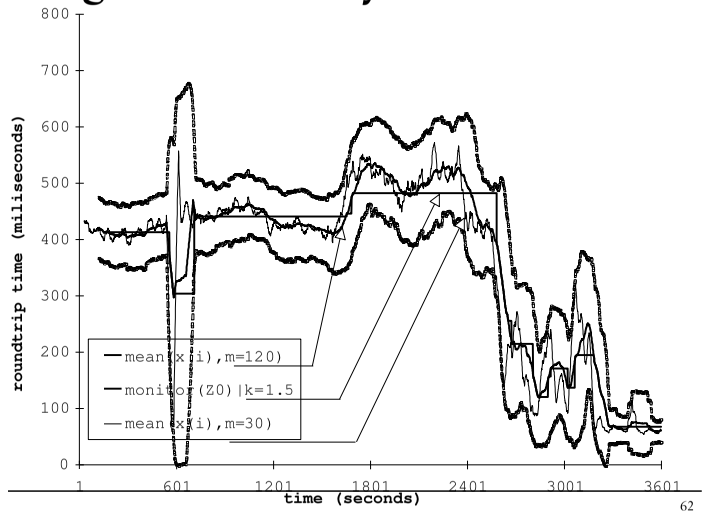

62

Comparative Performance

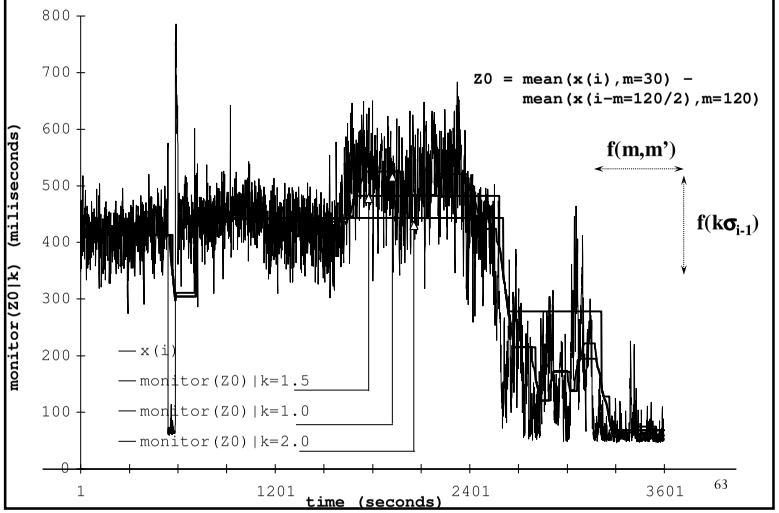

63

Research Contributions

- Introduced application and relevance of SPC to multimedia computing networking
  - online generation of process performance envelope
  - process-performance guidance over ARC problem
- Developed network state performance indicator
  - online SPC-based monitor of network state (e.g., delay)
  - reliable statistical indicator:
    - adaptable setup cost (m vs. m')
    - process performance/process state aware (μ, σ, k, mon)
    - forecast-strength (confidence interval, process performance/state)
    - adaptable reaction time (m vs. m'; k)
    - adaptable outlier recognition and handling (μ, σ, k)
  - detection and forecast of "long-term" stationary conditions

64

Related Work

- **Network Indicator Problem**
  - network delay fractal
    - random variable regardless of timescale
  - long-term indicators
    - leverages statistical process control for robust indicator
- **Adaptive Rate Control**
  - buffer smoothing
    - not application aware, short-term horizon
  - end-to-end application awareness
  - statistical process control
    - typically long-term horizon, industrial processes
    - typically run-to-run feedback control

65

Outline of the Talk

- **Background [10%]**
- **Computer-Supported Collaboration (Groupware) [25%]**
- **Dynamically Customized Web Touring [25%]**
- **Multimedia Computing Networking [35%]**
  - Network Performance Envelope Estimation
  - **Distributed Resource Management**
- **Wrap-Up [5%]**

66

Distributed Resource Management

Work at IBM T. J. Watson Research Center

By N. R. Manohar, L. Lumelsky, and S. Wood
U.S. Patents 6,516,350; 6,463,454; 6,529,950; 6,377,996; 6,460,082; 6,466,980

67

Motivation and Goals

- **Evaluation of Multimedia Computing Networking at:**
  - internet2 networking level (i.e., very high bandwidth)
- **What could benefit from this?**
  - valuable asset large object model (e.g., movies)
- **What could we do now?**
  - distributed resource management
  - utility computing (e.g., service plug-ins)

68

Distributed Resource Management

- **From: Ad-Hoc Resource Management**
  - ad-hoc client-services
    - streaming, metering
  - ad-hoc server-services
    - caching, load-balancing
- **To: Resource Management Plane**
  - distributed management plane
  - plug-in compliant servers
  - standard plane server-services
    - caching, replication, load-balancing, dynamic hosting, etc.
  - brokering of clients to servers

FIG. 3

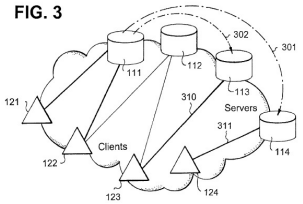

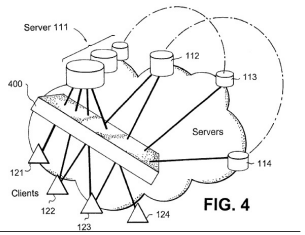

FIG. 4

69

Distributed Resource Management

- **Resource Management Policies**
  - traditional (load-balancing) demand shaping
    - (demand regulated to capacities) policy
  - capacity-shaping
    - (capacities regulated to demand) policy
- **Resource Management Mechanisms**
  - self-regulated capacity shaping
    - object is associated (allocated) object replicas
    - object replica specifies (serving) resource commitment
    - placed object replica commits binding resource allocation
    - expiration time associated with object replica
    - server placement of object replicas varies over time
    - total number of object replicas varies over time
  - thus, allocated object serving capacity shaped over time

70

Basic Idea: Resource Brokering Plane

- **Brokering Infrastructure Services**
  - Data: Object Demand
  - Resources: Allocated Object Capacities
  - Functions: Server-Services

Servers	**Objects**
geographically distributed	objects
multiple level hosting:	object replicas
global resource servers (e.g., dedicated server farms)	allocated object replicas
plug-in resource servers (e.g., spare hosting capacity)	number of allocated replicas
	placement of allocated replicas

71

Service Management Layer

- **brokering of plug-in server-services**
  - application-level services for hosted objects
- **drives capacity shaping**
  - brokers requests to objects
- **interfaces with systems management layer**

FIG. 6

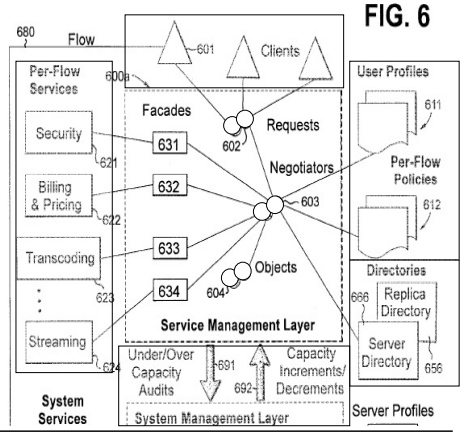

72

Systems Management Layer

- **systems management**
  - of dynamic object hosting at brokered servers
- **interfaces with service management layer**
- **implements capacity shaping**
  - brokers objects to servers

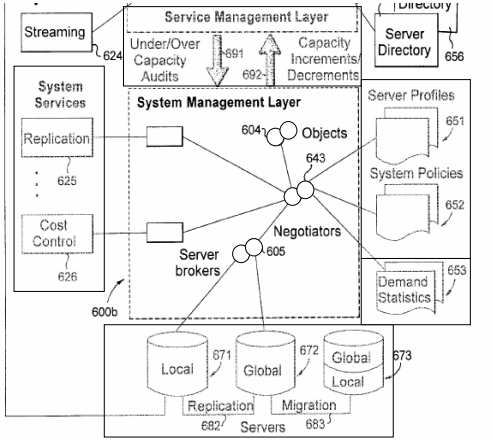

Self-Regulated Capacity Shaping

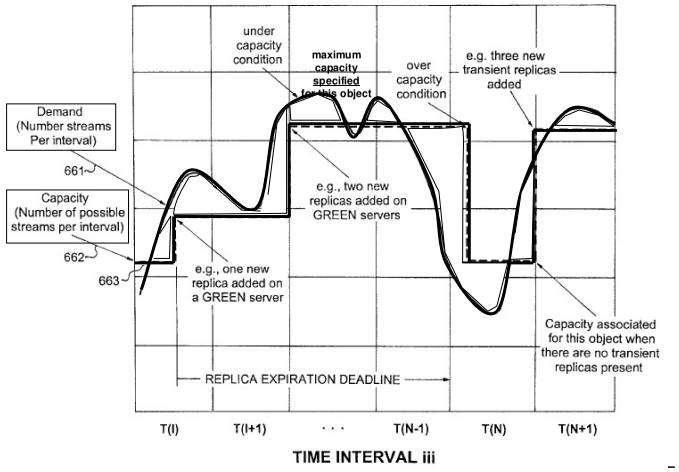

Multi-Level Dynamic Hosting: Global/Local Servers

- **O2 Demand**
  - G1
- **O2 Servers**
  - GS(G2)
  - LS(G1)
- **O2 State**
  - red @ G1
  - red @ G4
- **O2 Shaping**
  - green @ G3

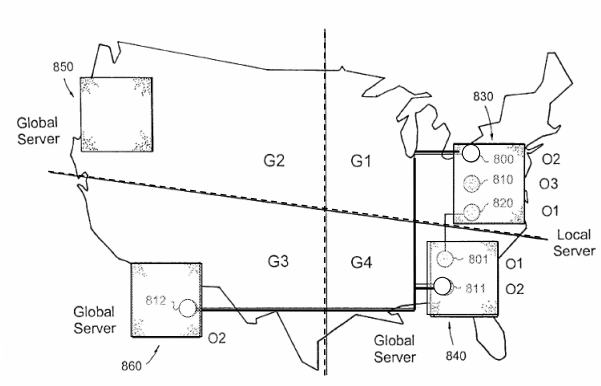

Replica Lifecycle Management

- **scarcity check**
  - demand/capacity analysis (1)
- **replica creation**
- **replica placement**
- **replication**
- **oversupply check**
  - demand/capacity analysis (2)
- **replica deletion**

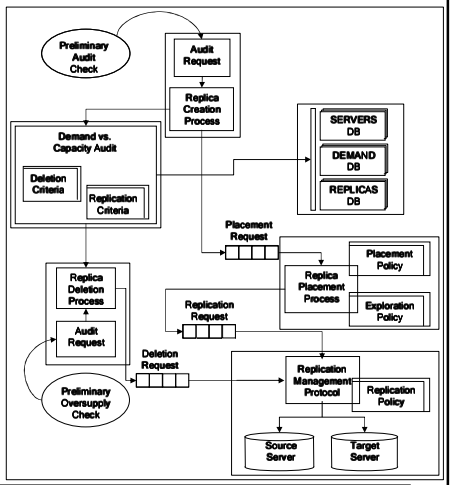

Low-Overhead State Tracking

- **Object to Demand**
  - demand volume
  - demand rating
- **Server to Capacity**
  - capacity rating
  - utilization state
  - globality
- **Object to Replica**
  - server placement
  - server capacity
  - expiration time

ObjectID	Demand rate req/s	Volume t_{i-1} req	Volume t_{i-2} req	Hot Object	Time Stamp
420	10	120	60	yes	t_i
425	5	60	55	no	t_i
428	5	30	62	no	t_i

`object 420: hot, large demand`

Server	IP Address	Capacity Rating	Utilization State	Timestamp	Globality
1211	209.09.9.127	Low	Red	t1	local
1221	128.0.0.1	High	Green	t2	global

`global server 1221: high capacity, low use`

Object_ID	Replica	Server	Transient Replica	Time-to-Live
420	421	1211	NO	
	422	1221	YES	060599-133000
440	441	1211	NO	

`transient replica for 420 at server 1221`

(100-X)% Brokered Capacity

allocation regions for server-serving capacity

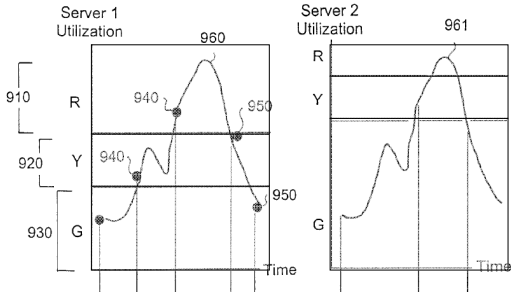

- **red: local critical spare (X%)**
  - non-allocatable server-serving capacity
- **yellow: trigger control for plane (server-tracking) updates**
  - safe region to smooth out transient state changes and messaging delays
- **green: low-overhead server-serving capacity**
  - safe server-serving capacity

Resource Brokering of Server

- **server resource bins**
  - storage bin
    - reserved replica storage units
  - service bin
    - reserved replica serving resources
  - binding
    - allocation of pre-allocated resources (service bin) to an allocated object replica (at a storage unit)
- **global/local resource pools**
  - resource bins further divided into:
    - reserved for global availability – used for global system-wide capacity shaping
    - reserved for local availability

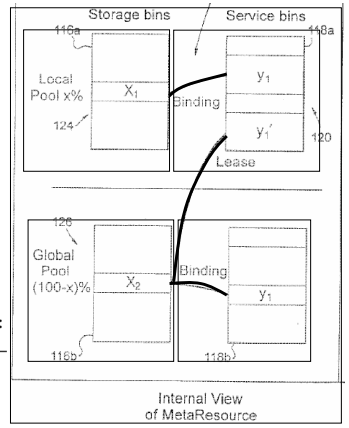

Object Service Unit Abstraction

- **Resource Envelope**
  - specifies resource reqs for object-to-server binding
- **Service Signature**
  - brokering parameters
  - adaptation parameters

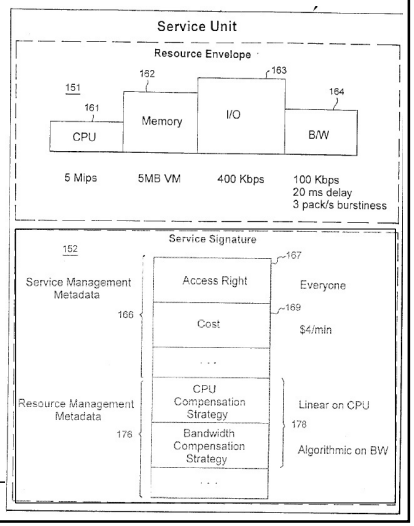

Contributions

- **Self-Regulated Resource Management**
  - capacity-shaping in addition to traditional demand-shaping
  - policy-based application-level brokering
  - complementary to existing technologies (multicast, QoS, copyrights, etc.), but levering internet2 resources
- **Resource Management Plane**
  - low-overhead distributed resource state tracking
  - (100-X)% slack: above best effort, below hard QoS guarantees
- **Global/Local Resources**
  - serving-servers (100% global or 100% local)
  - serving-server resource pools (fractionally global/local; useful for capacity-shaping)

81

Related Work

- **Virtual Hosting (server farms – e.g., IBM Global Services)**
  - demand-shaping of capacity at centralized center(s) vs.
  - capacity-shaping of demand through placement of distributed capacities
- **Distributed Programmable Planes (grid computing – e.g., Globus)**
  - resource management for entirely different problem,
  - small number of large (computations and datasets) objects with low viewership vs.
  - large number of large (value-asset serving) objects with high and wide viewership
- **Distributed Replica Management**
  - distributing edge-caching network (e.g., Akamai) – on-demand caching of relatively small objects vs. resource-managed replication of valuable asset very large object
- **Multicast Streaming/Programmable Overlay Networks**
  - multicast is complementary technology
  - plane represents application-level intelligent network management over the placement of serving-servers, that is, the self-regulated placement of the sources of the multicast trees
- **Brokered Distributed QoS Architectures**
  - distributed QoS brokering is complementary technology
  - delivering per- server-session connection management once distributed resource management plane self-regulates placement of server-sources
  - delivering end-to-end QoS feasibility information for brokering of clients to servers

82

Outline of the Talk

- **Background [10%]**
- **Computer-Supported Collaboration (Groupware) [25%]**
  - Record and Replay Paradigm
  - Multimedia Synchronization
- **Dynamically Customized Web Touring [25%]**
- **Multimedia Computing Networking [35%]**
  - Network Performance Envelope
  - Distributed Resource Management
- **Wrap-Up [5%]**

83

Software Systems Research Map

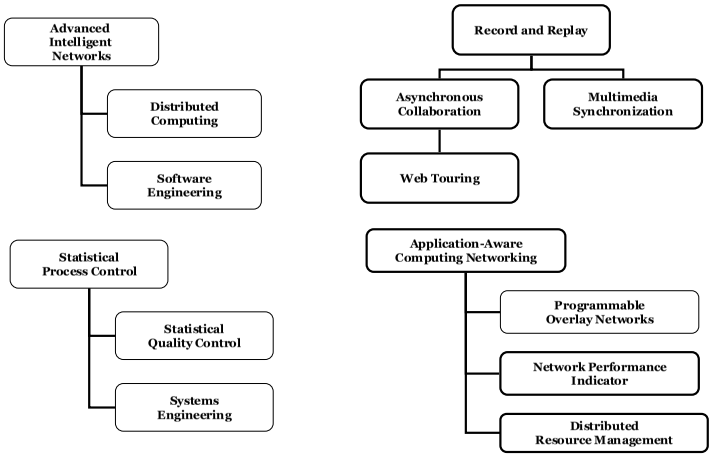

84

Selected Research Contributions

- **Record And Replay (By Re-execution) Paradigm**
  - for asynchronous collaboration, capture of intra-task content
  - manipulation of computer sessions as first class objects
- **Multimedia Scheduling And Synchronization Protocols**
  - for fine-grained re-executable events and continuous media
- **Application Of Statistical Process Control**
  - long-term process performance indicators, application awareness
  - for adaptive rate control problems (multimedia, networking)
- **Dynamically Customizable Web-touring**
  - token-based projections over touring content
  - like-minded exploration suggestions over touring content
- **Distributed Resource Management**
  - service management plane, internet2-centered service model,
  - replica management (time variant number/placement of replica capacity)

85

Proposed Courses

- **Classes**
  - Introduction to Databases
  - Introduction to Software Engineering
- **Special Topics**
  - Special Topics: Software Design
  - Special Topics: Collaborative Systems
- **Advanced Classes**
  - Software Systems Principles
- **Seminars**
  - Software Systems Research

86

Proposed Focus Areas

- **CSCW/HCI Applications**
  - groupware, collaborative intelligence, information management, etc.
- **Distributed Resource Management**
  - utility computing, pervasive computing, sensor-based computing, etc.
- **Software Systems Research**
  - systems support for the above
  - formalization, experimentation, and simulation

87