Tools for Collaboration

Colin Perkins

http://csperkins.org/teaching/2004-2005/gc5/

UNIVERSITY
of
GLASGOW

Updated Module Outline

Week 7	Lecture 13	Data Access, Integration and Management	John Watt
	Lecture 14	Data Provenance and Curation	Seamus Ross
	Tutorial 6	Discussion of Data Management/Provenance	Richard Sinnott
Week 8	Lecture 15	Data Transfer	Colin Perkins
	Lecture 16	Peer-to-Peer Communication	Colin Perkins
	Tutorial 7	Discussion of Networking Papers	Colin Perkins
Week 9	Lecture 17	Tools for Collaboration	Colin Perkins
	Lecture 18	Sample Applications	Richard Sinnott
	Tutorial 9	Discussion on the Future of Grid Computing	Richard Sinnott
Week 10	Lecture 19	The Future of Grid Computing	Richard Sinnott
	Lecture 20	Review of Major Concepts	All
	Tutorial 9	Q & A	All

- Teaching quality questionnaires available: please return at end of lecture or to the teaching office

Lecture Outline

- Requirements for building collaborative work environments
  - Session initiation, negotiation and control
  - Media transport
  - Deployment considerations
- Example: The AccessGrid
  - Comparison of AccessGrid and “typical” conferencing systems
- Other real-time applications
  - Visualization
  - Data capture and dissemination

Tools for Networked Collaborative Work

- Support collaborative work by scientists, industry and others
  - Group communication scenarios
  - Widely distributed participants
- Integrate with Grid computing resources
- Provide visualization facilities
  - Shared state and data repositories, computational resources, applications
- Provide a high-quality audio-visual experience
  - High fidelity audio, possibly with spatial positioning
  - High resolution video, wide colour gamut, high frame rate
- Provide a sense of community and presence
  - Venues as a rendezvous point; virtual meeting place
- Secure and trusted infrastructure

Building the Supporting Infrastructure

How are collaborative environments built?

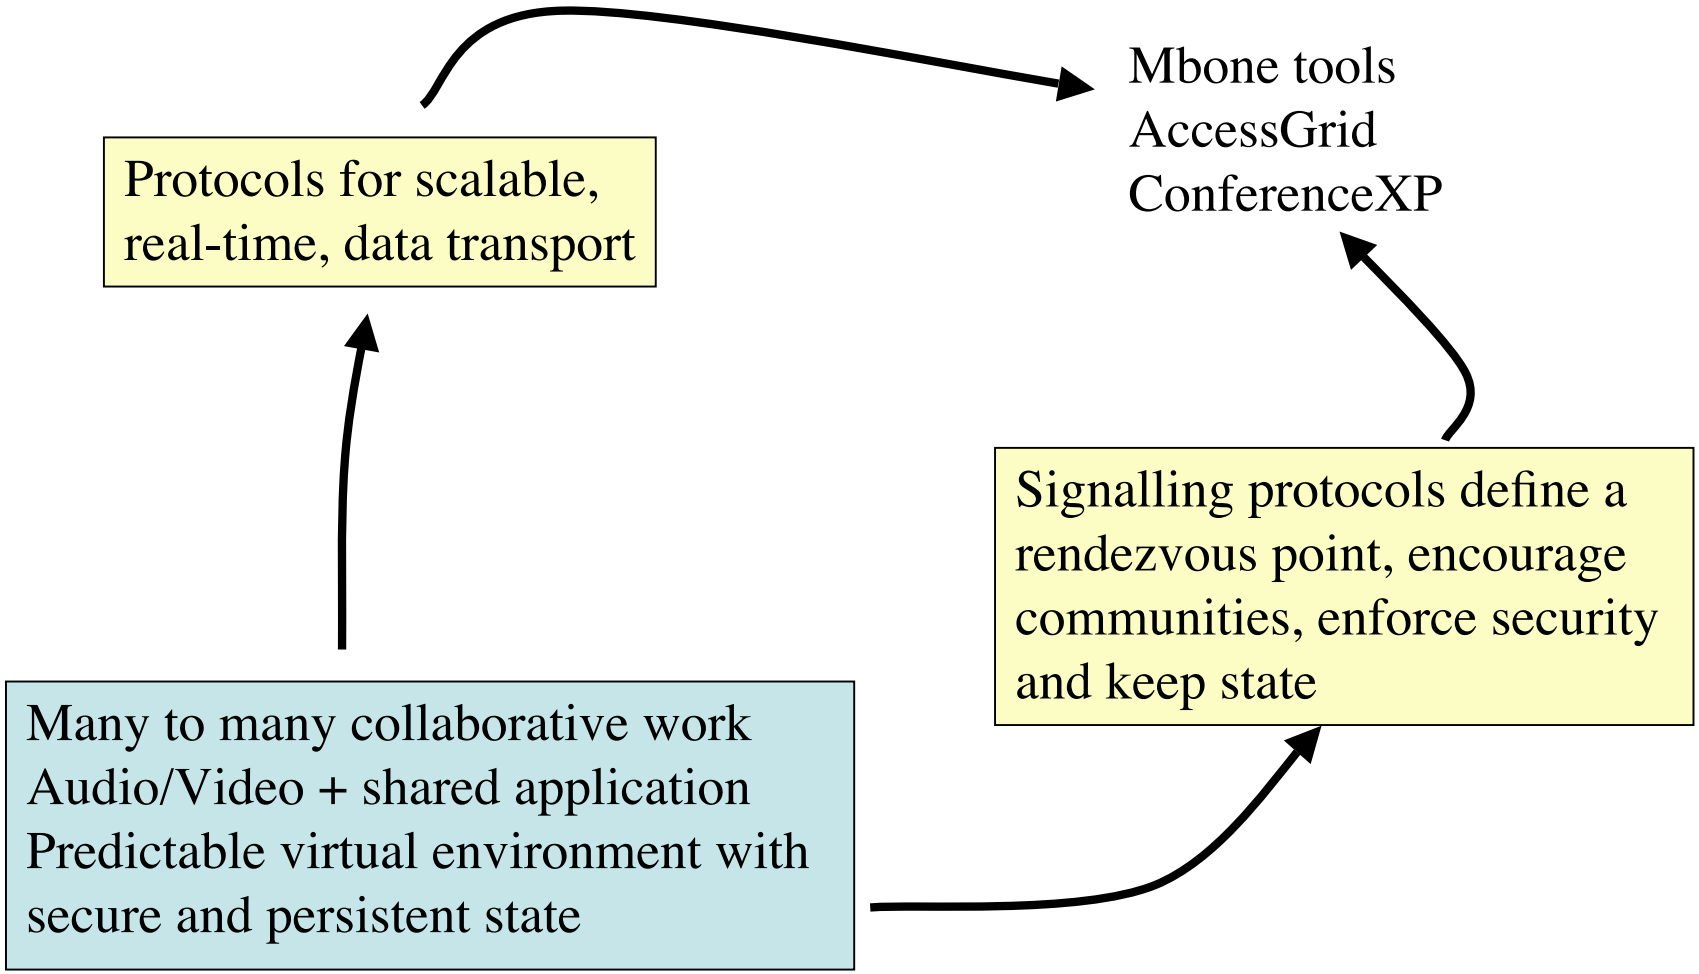

Building the Supporting Infrastructure

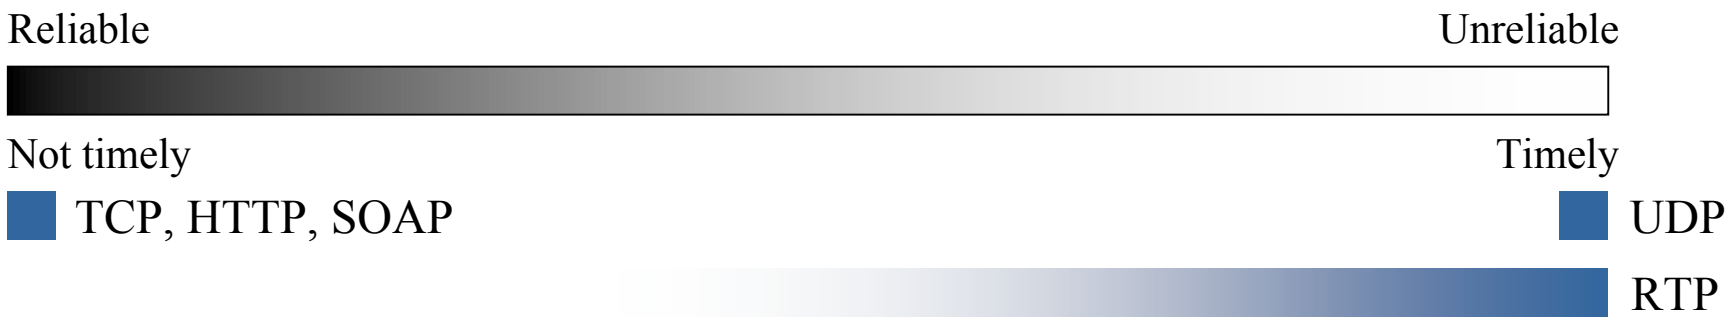

- Protocols built on unreliable packet networks must make a fundamental trade-off:
  - Timely
  - Reliable
- Real-time systems choose their transport carefully:
  - TCP for signalling, bulk transfer
  - UDP for real-time data
- Application level protocols can blur the boundary using FEC, etc.

Building the Supporting Infrastructure

- Building a collaborative work environment requires several components:
  - Signalling Protocols
    - Session advertisement
    - Session initiation, negotiation and control
    - Session description
  - Media transport protocols
- Along with a Grid computing environment/toolkit to enable shared data and application services

Signalling Protocols

The first stage of any multimedia session is signalling

- User location
- Session initiation, call setup and teardown
- Media negotiation
- Conference control

Many signalling protocols exist

- Teleconferencing: H.323
- Telephony: *SIP*
- Streaming: RTSP, *SAP/IMG*
- AccessGrid: Web services

Signalling Protocols: SAP, IMG

- "Session Announcement Protocol" RFC 2974
  - A multicast announce-listen protocol for wide area announcement of multimedia sessions
  - Announcers periodically multicast SDP descriptions to a well known group
  - Inter-announcement interval is 10+ minutes
  - Listeners slowly build up a cache of sessions
  - Suitable for announcing long-lived public sessions
    - E.g. radio/TV station, event coverage, institutional venues
  - Limited deployment on public Internet, due to lack of IP multicast
- Enhanced versions under development for 3G and cable TV
  - Private networks that support IP multicast
  - Current work on "Internet Media Guides" for IP-based television
- Others use web pages to distribute announcement
  - E.g. AccessGrid
  - Scaling problems…

Signalling Protocols: SIP

- Other sessions and tightly coupled or need control, negotiation
- Internet standard for signalling, session initiation and control
  - RFC 3261 + extensions
  - Flexible integration with other Internet services
    - Email, Web, streaming media, recording, agents, etc.
    - Protocol operation is based on HTTP
- Main Functions:
  - Invite participants to sessions
    - Find the user's current location, match with their capabilities and preferences in order to deliver invitation
    - Carry opaque session descriptions (SDP, SDPng)
  - Modification of sessions
  - Termination of sessions

SIP Protocol Components

- User Agent Client (UAC)
  - End systems (e.g. phones)
  - Send SIP requests
- User Agent Server (UAS)
  - Listens for call requests
  - Prompts user or executes program to determine response
- User Agent = UAC + UAS

- Redirect Server
  - Redirects users to try other server
- Proxy Server
  - Servers can "fork" requests to form a search tree
- Registrar
  - Receives registrations regarding current user locations

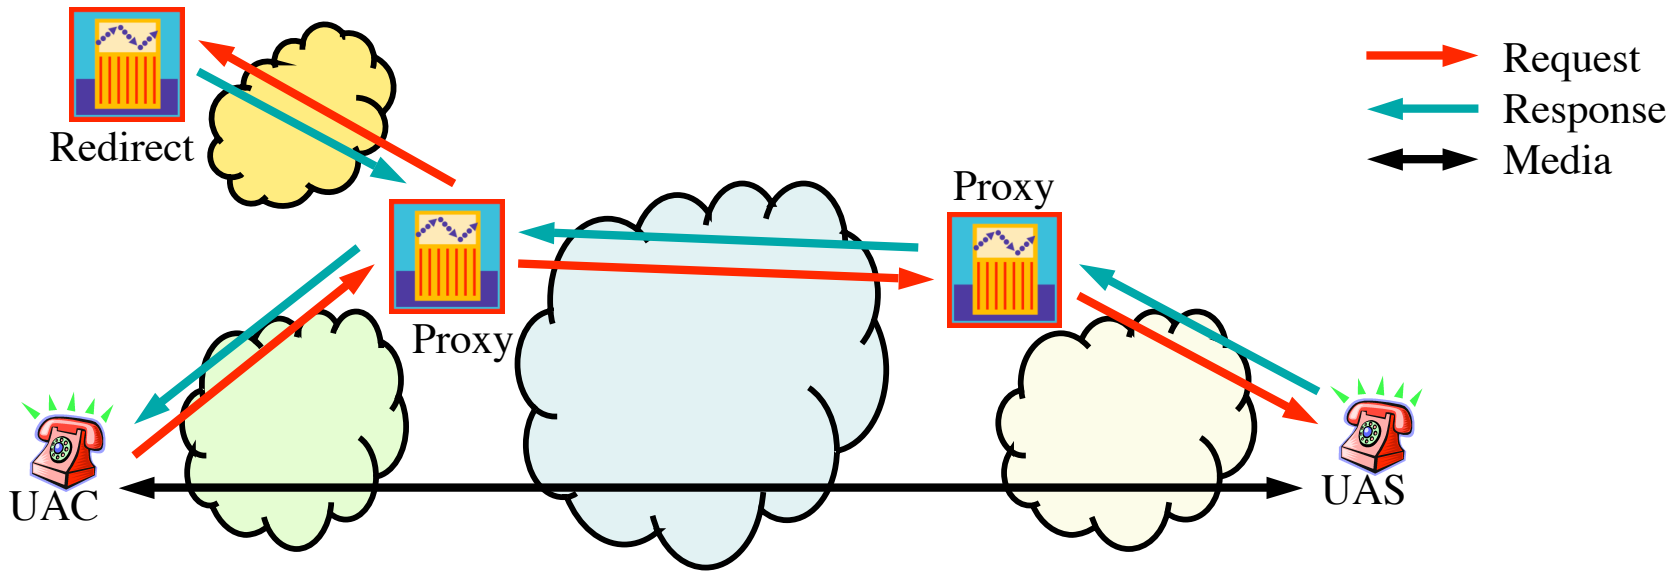

SIP Operation

- Signalling infrastructure tracks user location
  - Based on proxy servers in the network, one per site
  - Redirect calls as users move
  - Handle voice-mail, redirect to webpage, email, IM, etc.
  - Negotiate appropriate media formats
    - Needs a session description format
  - Traverse firewalls, perform accounting, etc.
- Media follows separate path to signalling
  - Media flows end-to-end, signalling goes through proxies
  - Separate media protocols
- Somewhat parallels traditional telephony signalling
  - More open and flexible
  - Based around Internet-style protocols

Session Description: SDP

- Simple text-based format to describe addresses, ports, codecs, etc.
- Limited syntax, poor expressive power
- Limited support for grouping and labelling media lines
- Limited support for negotiation of actual/proposed capabilities
  - Offer/answer model + extensions

```
v=0
o=jdoe 2890844526 2890842807 IN IP4 10.47.16.5
s=SDP Seminar
i=A Seminar on the session description protocol
u=http://www.example.com/seminars/sdp.pdf
e=j.doe@example.com (Jane Doe)
c=IN IP4 224.2.17.12/127
t=2873397496 2873404696
a=recvonly
m=audio 49170 RTP/AVP 0
m=video 51372 RTP/AVP 99
a=rtpmap:99 h263-1998/90000
```

Session Description: SDPng

- SDP widely deployed, community suffering from its limitations
  - Numerous extensions to kludge around these
  - Difficult to maintain, understand
  - Required for backwards compatibility

- Initial proposal for “SDPng”, an XML-based replacement for SDP
  - Expressive syntax; allows grouping, labelling, can be validated
- Makes explicit difference between:
  - Describing the actual session parameters (original SDP design)
  - Describing and negotiating the capabilities of a system
- Much more expressive; more complex

SDPng Example

```
<?xml version="1.0" encoding="UTF-8"?>
<sdpng xmlns="http://www.iana.org/sdpng"
       xmlns:audio="http://www.iana.org/sdpng/audio"
       xmlns:rtp="http://www.iana.org/sdpng/rtp"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xsi:schemaLocation="http://www.iana.org/sdpng sdpng-base.xsd
                           http://www.iana.org/sdpng/audio sdpng-audio-pkg.xsd
                           http://www.iana.org/sdpng/rtp sdpng-rtp-pkg.xsd"
       owner="A@example.com" id="98765432" version="1">
  <cap>
    <audio:codec name="avp:pcmu">
      <audio:encoding>PCMU</audio:encoding>
      <audio:channels>1</audio:channels>
      <audio:sampling>8000</audio:sampling>
    </audio:codec>
    <audio:codec name="avp:gsm">
      <audio:encoding>GSM</audio:encoding>
      <audio:channels>1</audio:channels>
      <audio:sampling>8000</audio:sampling>
    </audio:codec>
  </cap>
  ...
```

Potential capabilities

SDPng Example

```
...
  <def>
    <rtp:udp name="rtp-cfg1" ref="rtpudpip4">
      <rtp:ip-addr>192.168.47.11</rtp:ip-addr>
      <rtp:rtp-port>51400</rtp:rtp-port>
    </rtp:udp>
  </def>
```

Definitions used below

```
  <cfg>
    <component>
      <alt>
        <audio:codec ref="avp:pcmu"/>
        <rtp:udp ref="rtp-cfg1"><rtp:pt>0</rtp:pt></rtp:udp>
      </alt>
      <alt>
        <audio:codec ref="avp:gsm"/>
        <rtp:udp ref="rtp-cfg1"><rtp:pt>3</rtp:pt></rtp:udp>
      </alt>
    </component>
  </cfg>
</sdpng>
```

Codecs currently in use

Media Transport Protocols

- Once the session has been setup, media flows
- Desire convergence on a single media transport protocol for:
  - Voice over IP
  - Teleconferencing
  - Streaming media

- Not simple: IP not well suited to real-time data transfer

Media Transport Protocols

- The challenge:
  - build a mechanism for robust, real-time media delivery above an unreliable and unpredictable transport layer
  - without changing the transport layer

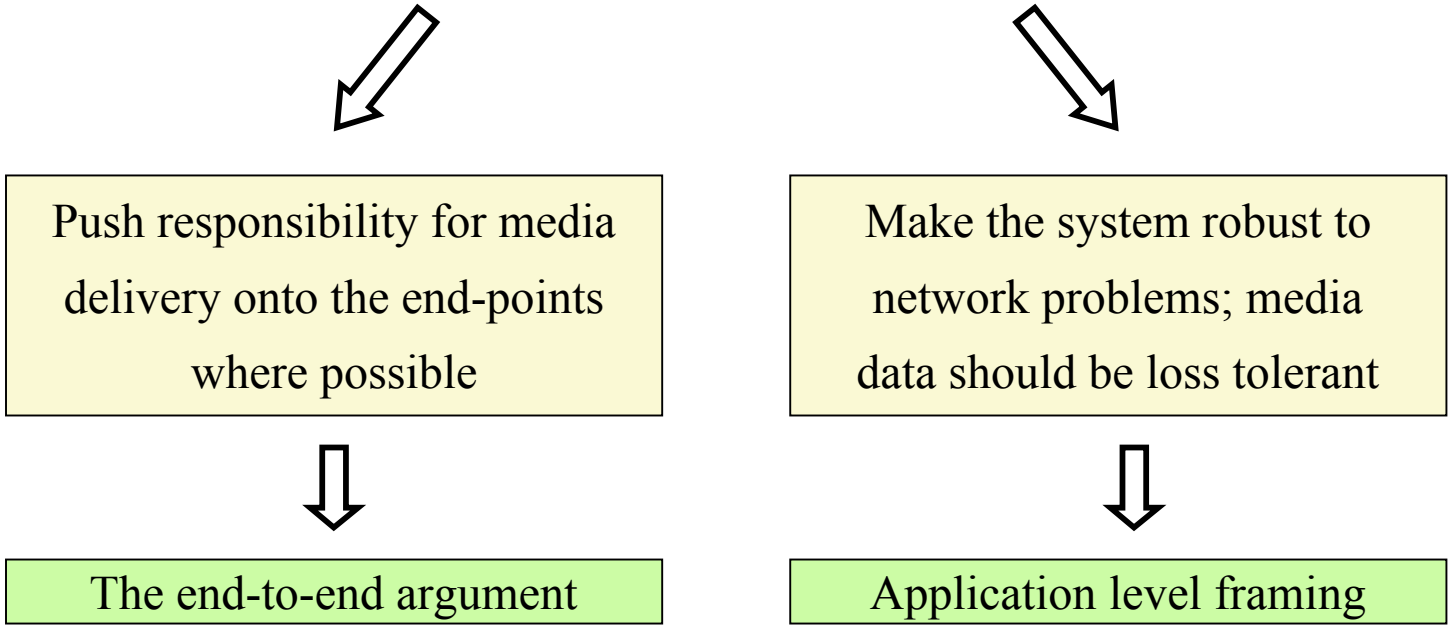

Media Transport Protocols

- End-to-end Argument
  - Responsibility remains with the end points, ensure delivery or note failure
  - Allows provenance tracking
  - Allows choice of reliability/timeliness trade-off
- Application Level Framing
  - Only the application has sufficient knowledge of its data to make an informed decision about how that data should be transported
  - The transport protocol should accept data in application meaningful chunks ("ADUs")
  - The transport protocol should expose details of delivery, allowing the applications to react intelligently if there are problems

Both desirable for real-time applications…

RTP: Real-time Transport Protocol

- There is a standard real-time protocol built around these ideas
- Real-time Transport Protocol (RTP)
  - The standard for real-time transport over IP networks
    - Video conferencing; VoIP/telephony; Streaming audio and video
  - Published as an IETF "draft standard" RFC
    - RFC 3550, 3551
    - Widespread use for video conferencing (H.323 and SIP-based)
    - 3GPP mobile phones
    - Some use in streaming

RTP: Real-time Transport Protocol

- Delivers a single data stream to one, or more, receivers
  - Sequencing and timing recovery
  - Framing and error detection
  - Lip-synchronisation
  - Participant and data format identification
  - Reception quality statistics
  - Extensions for security, authentication, etc.
- Payload formats and profiles
  - Adaptation layers for particular data types, application classes

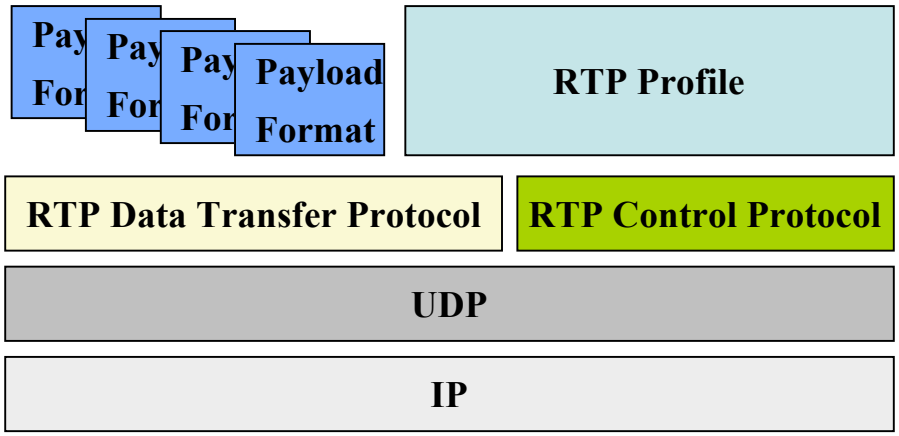

Deployment Considerations

- Multimedia traffic has separate signalling and media flows
  - Initial signalling flow to negotiate the type and format of media, determine end-point addresses, etc.
  - Media flows start as a result of signalling
- Doesn't work through NAT boxes
  - The IP address an end-point thinks it has, and negotiates in the signalling, doesn't match that seen by it's peer
- Several proposals to work around the issue:
  - STUN: detection and characterisation of NAT
    - Has to run on the same port as the media, multiplexing problem
    - Only works through a limited subset of NAT types
  - TURN: relay protocol when both parties behind NAT ⇒ triangle routing
  - ICE: Methodology for systematically trying alternatives to detect the NAT
- Only real solution to remove NAT!

The Internet Multimedia Protocol Stack

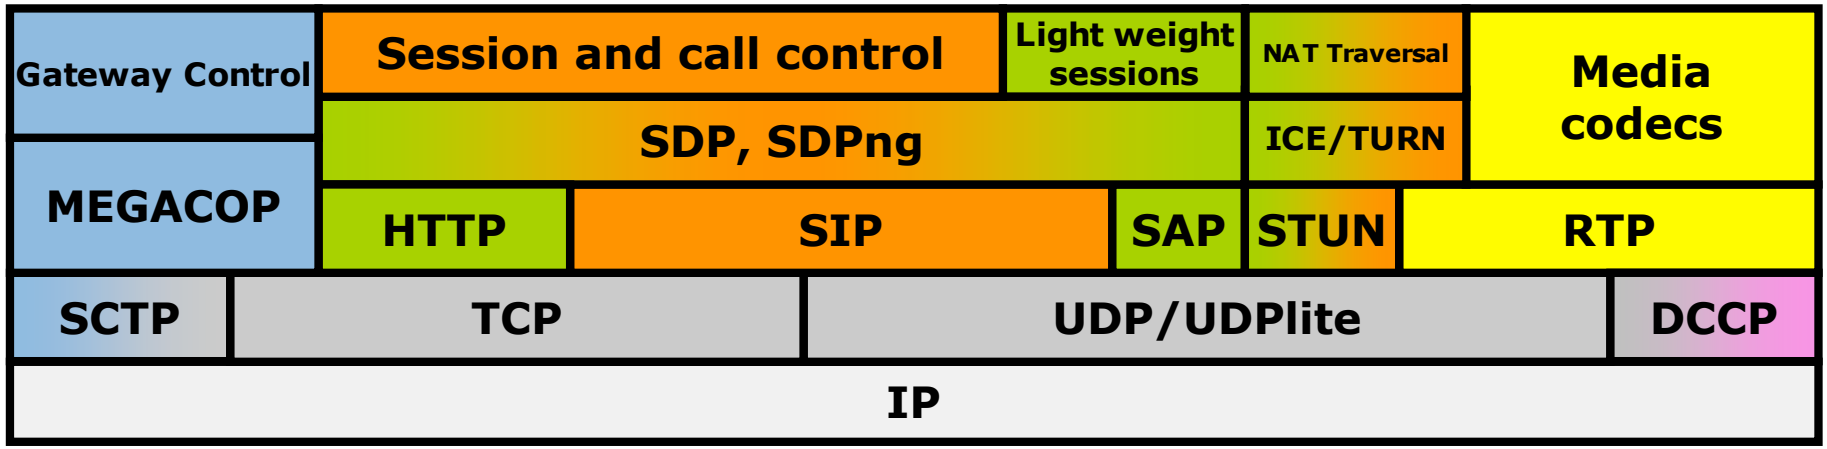

Example: The AccessGrid

- An attempt to integrate multicast RTP audio/video conferencing with a web-services framework
  - Standard transport and media codecs
  - Non-standard web-services based session initiation and control
    - Should eventually incorporation of other Grid services
- http://www.accessgrid.org/

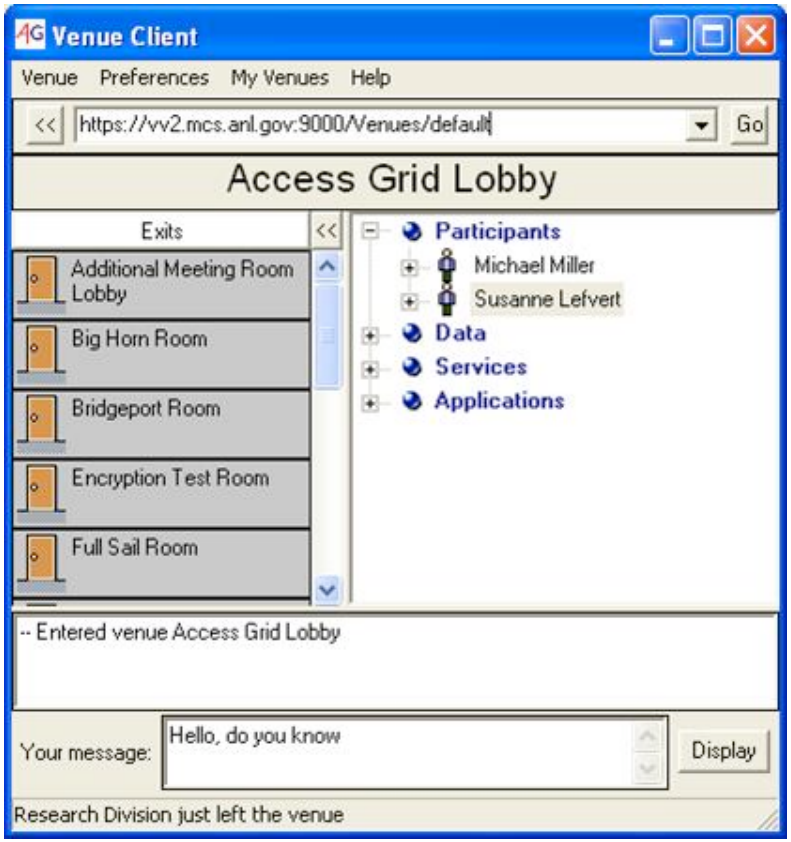

- Widely used by the e-science community
- Other communities use SIP or H.323 based solutions

Other Real-Time Applications

- Grid community becoming interesting in real-time applications:
  - Visualization
  - Data Capture and Dissemination

- Remote real-time visualization obviously fits within the same framework as video conferencing
- What about real-time data dissemination?

Data Capture and Dissemination

- Particle physics and astronomy
- Monitoring earthquakes
- Environmental sensor networks
- …

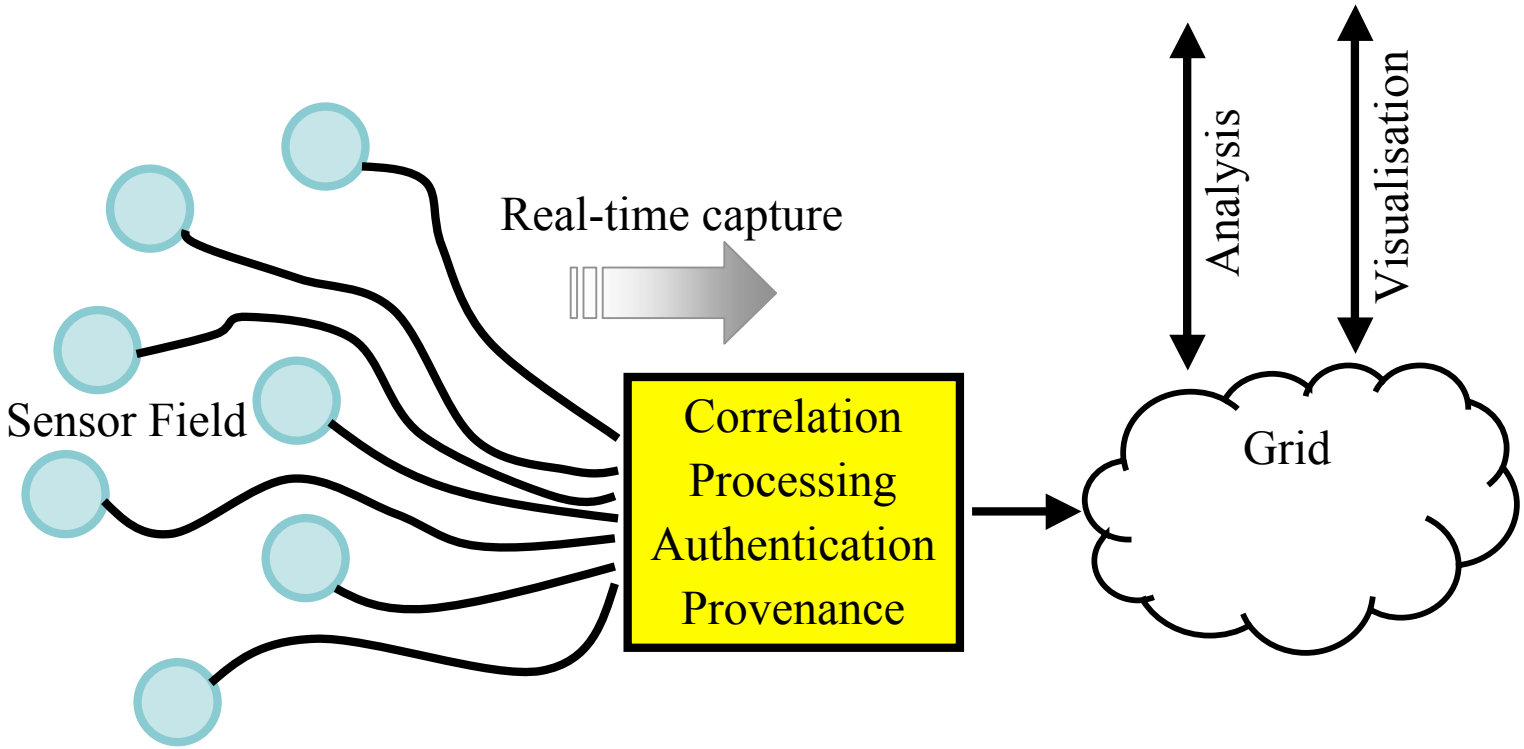

Data Capture and Dissemination

- Initial data capture steps use ad-hoc methods
  - Unreliable and not robust to changes
  - Difficult to ensure consistent processing and data provenance
- To enable these classes of application, we need:
  - Transport protocols and middleware for time sensitive data
    - Real-time flow from sensor to capture device
    - Real-time flow from analysis tool to visualisation device
  - Service description and negotiation
    - Including real-time services/resources; e.g. capture devices
  - Service invocation and control
    - Critical time dependencies

Data Capture and Dissemination

- Many similarities to network multimedia systems; distinct from traditional Grid protocols
  - RTP profiles and payload formats
    - E.g. RTP for visualization
    - E.g. use of RTP with eVLBI
  - SDP and SDPng to describe data capture/transmission sessions
  - SIP to negotiate data formats
- Use Grid virtual organization model for security, authentication, later data analysis and processing
- May need new toolkit elements…

Example: e-VLBI radio astronomy

- Data read from multiple radio telescopes combined to give a single, more accurate, output
- Many terabytes of data
- Currently boxes of magnetic tapes shipped around the world using FedEx
  - This doesn't improve response time for the science!
  - Can miss important events
- The astronomy community would like to capture, correlate, process data in real-time using the Grid
  - Exploring use of RTP payload formats
  - Gigabit rate real-time transfer

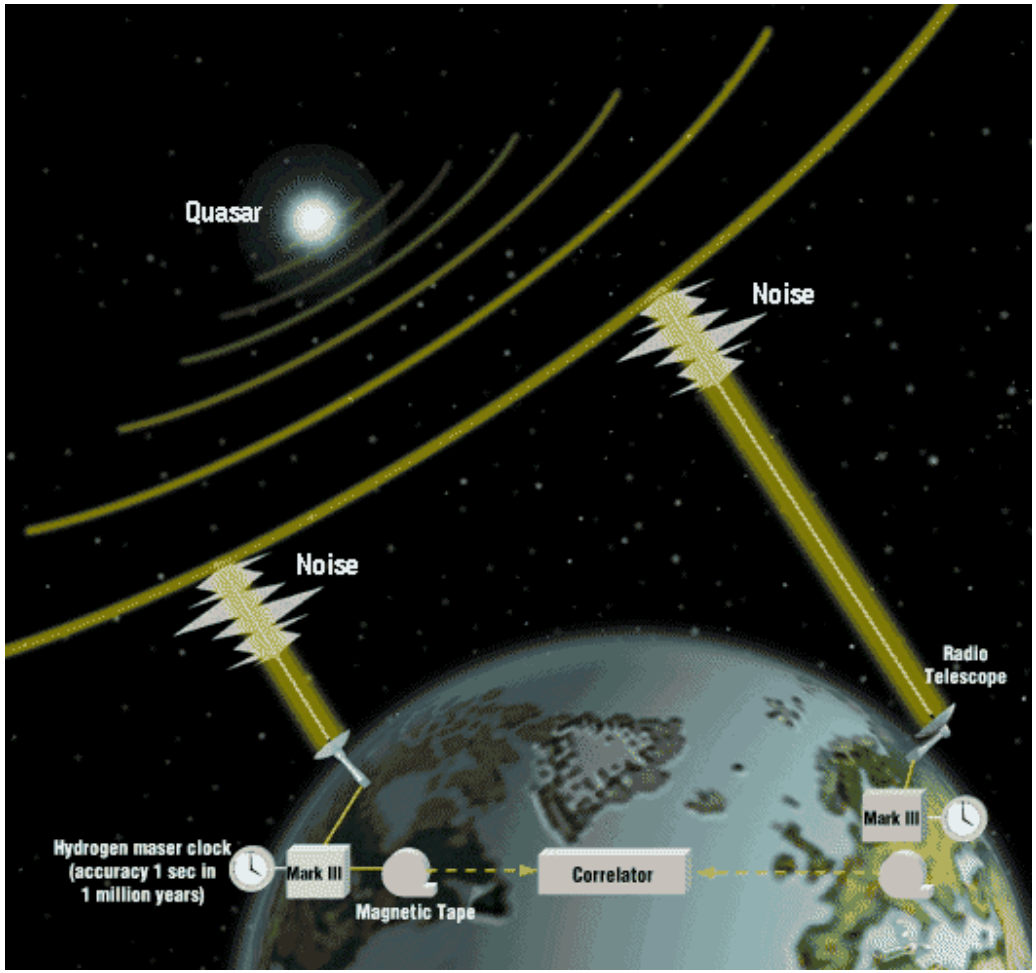

Lecture Summary

- You should know…
  - What components make up collaborative work applications
  - Outline of signalling protocols
    - SIP
    - SAP
    - SDP/SDPng
  - Outline of media transport protocols
    - RTP
  - Examples of real-time/collaborative work applications
    - AccessGrid
    - eVLBI