MMOX
Background and Approaches

Massively Multiparty online X (Games and applications)

David W. Levine
IBM
Thomas J. Watson Research Center
dwl@us.ibm.com

Influenced by extensive discussion on the IETF MMOX pre-bof mailing list
March 24, 2009

Outline

- Background
  - A little history
  - Defining the space
  - Range of existing approaches
  - Building blocks of immersive services
- Some of the existing code bases
  - Linden Lab / OpenSimulator.org ecosystem
  - Croquet / QWAQ ecosystem
  - Other prominent examples
- Technical Approach
  - Specific model for early work
  - Balancing short term and long term work
- Related Work
- Out of Scope
- Expectations of the proposers

Background

- Several related online services
  - Games
  - Virtual Worlds
  - Shared multi user applications
- Starting to grow ecosystems of related technologies around several popular approaches.
- Multiple implementations leads to wanting interoperability standards
- Begin to create common building blocks based on actual needs and code

A little history

- This is not entirely new terrain – VRML for example tackled similar issues – Collada provides a spec for static 3d content interchange – Other related work
- The current generation of successful virtual worlds and immersive games has led to several recent attempts to tackle standards
- After doing some preliminary work in 2008, taking this work to IETF seemed like an appropriate way to determine if it is ready for a broader arena
- Proposed WG and BOF to IETF, initial mailing list attracted a number of proposals, and lively discussion
- This work is not trying to boil the ocean, but solve current problems and then be well positioned to tackle broader issues.

Defining shared spaces

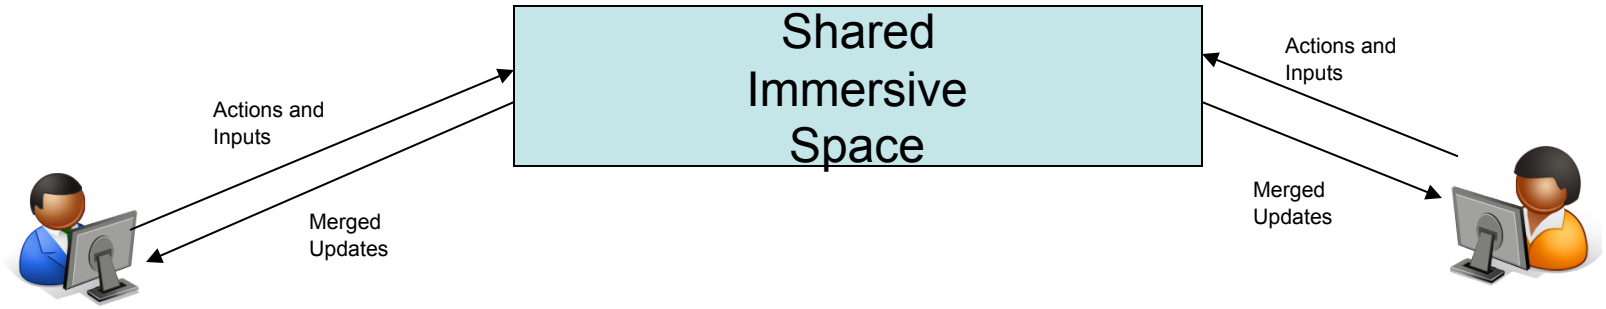

A shared space including:

- A shared setting
- A set of users
- A representation of the users projected into the space
- A shared experience of the events in the space projected back to the users currently in the space

Often, but not always:

- Objects which can be dynamically created and modified to form part of the space's setting
- Persistence of objects and setting
- Mechanisms to associate external streams of media with portions of the immersive space
- Mechanisms to run scripted behaviors in the space
- Simulated physics and behavior of objects and projected users in the spaace
- Mechanisms for users to move from one virtual space to another (by hand off between regions implicitly, or by explicit user request (teleport)
- Libraries of material the users can add to the space
- Mechanisms for users which are not in a common virtual space to interact with each other in some form (be aware of each other's locations and online status, message each other, etc)

Core Task

- Meld together the inputs from
  - A number of users and their actions
  - Media streams
  - A simulation of the underlying space
  - Scripted / Programmed artifacts in the space
- Distribute the melded results back to all the users so that a consistent shared experience results
  - Virtual worlds and current games do this with a goal of rendering consistent, visually attractive results
  - Generally this is presented by a rich client using OpenGL, DirectX, and related GPU driven rendering
- The shared experience is the core deliverable of the technology

Several uncommon internet tasks

- Sharing rapidly changing content in real-time with large numbers of endpoints
- On the boundary between “media streaming” and “structured document interchange” with aspects of both
- Rich, two way flows of data with long lasting connections (COMET, AJAX style)
- Collaboration of multiple services across trust boundaries to deliver full range of function
- These challenges are shared by most of the approaches currently in use – They may share state, control and content at different levels of architecture, but they share many of the unusual uses of the internet

A spectrum of approaches

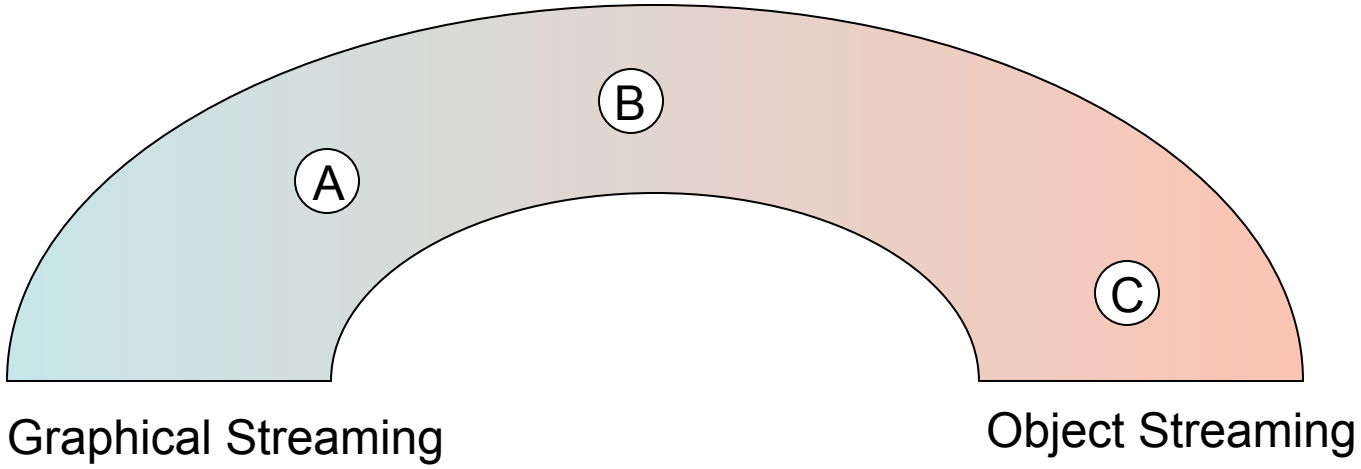

- Pure graphics delivery as one extreme
- Object level sharing (Croquet) at the other
- Some popular points
    - A Structured Graphics (Second Life like)
    - B Object/Action level synced clients/Simulations (OLIVE and Related Approaches)
    - C Object level synchronized peers (Croquet/QWAQ/Wonderland)

This is only one way to capture the differences between approaches, your metaphor may vary

Building Blocks

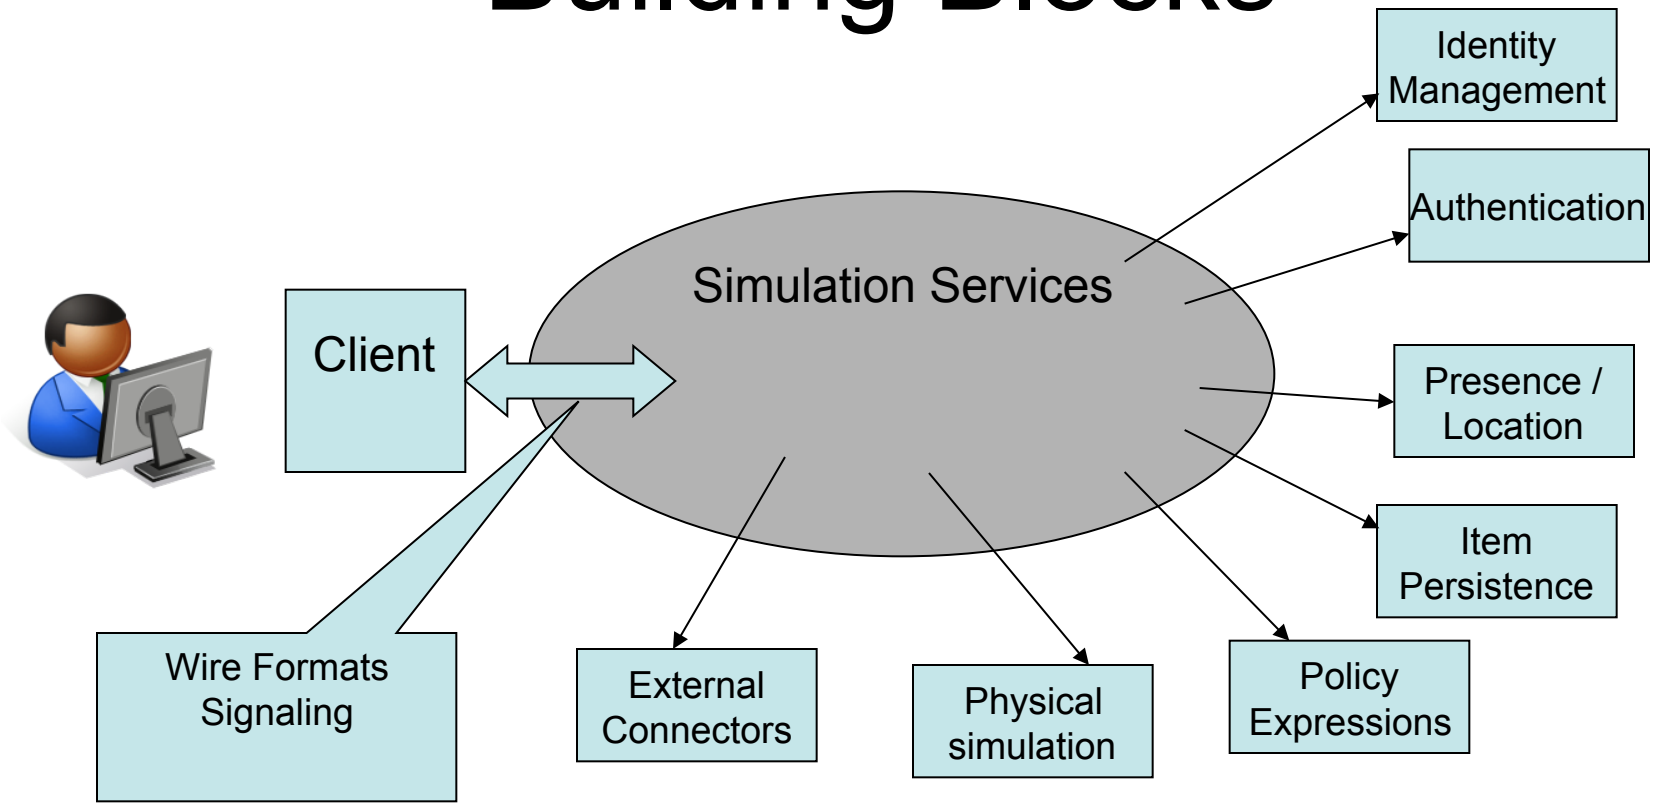

This is a logical, not structural view of the parts which comprise these spaces

Several ways to combine the parts

- Clients generally run some local effects
- Clients can include copies of the virtual space, ranging from a set of graphical elements to a complete synchronized model of the space.
- The degree to which there is a definitive copy of the virtual space is variable
- The forms of signaling actions, and how the actions are merged into the shared experience is variable, ranging from centralized, to peer to peer structures
- Different approaches yield different tradeoffs
- Bridging different affordances in the systems is fundamentally hard.

Content models vary

- Content can be modeled in different types of 3d building blocks, and at different degrees of composition
- Some systems work primarily on relatively monolithic models, others on sets of smaller elements combined to produce similar visual effects
- Systems tend to be tightly coupled, between content and modeling both at the service and client level
- The degree to which content changes, varies dramatically across different current systems

Deep Coupling is an issue

- Most of the current deployed systems are complete, closed environments with deep coupling between the components, explicit and implicit
- Most of the systems take advantage of these deep couplings to significantly increase the power of their systems and lower the bandwidth and computational costs of providing shared experiences
- It will take time, and effort to factor out the various couplings, and some will most likely limit interoperability for some time.

Linden Lab emerging Ecosystem

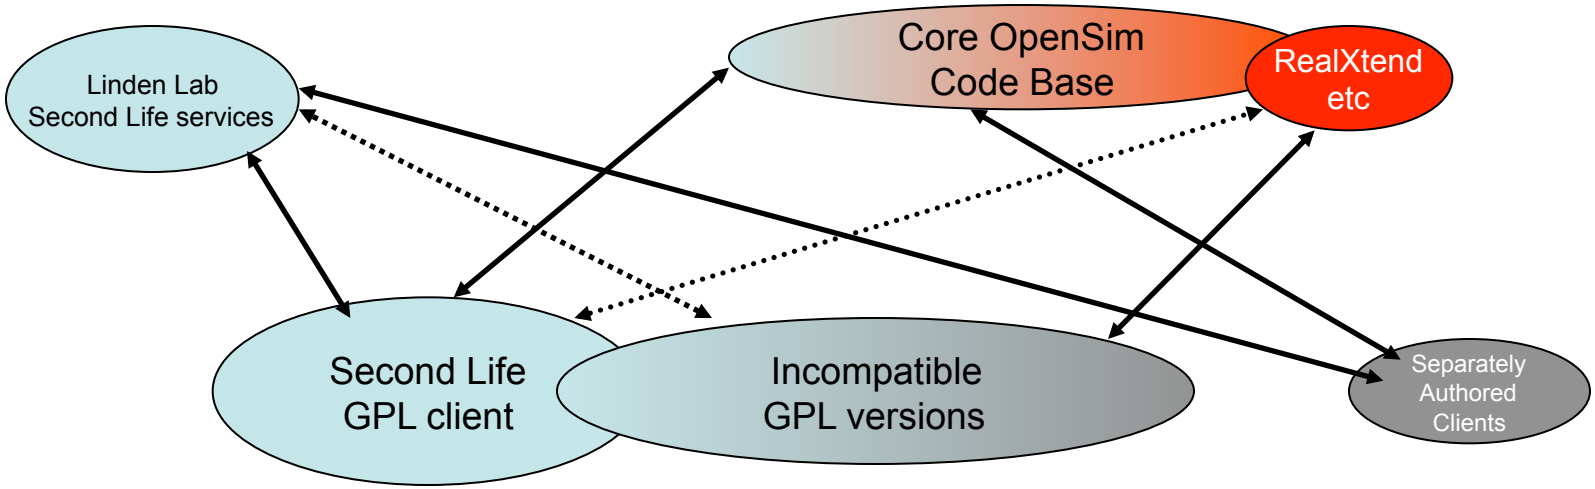

- Anchored by Linden Lab's GPL licensed client
  - Growing collection of diverging clients
  - RealXtend merges Linden Client with OGRE rendering
  - Increasingly distinct code bases
- Second Life Service, OpenSimulator.Org
  - Two completely separate code bases
  - OpenSimulator code base has multiple diverging forks, some of which requires specialized clients
  - Includes IBM's Sametime 3d beta offerin
- Several lightweight, Entirely separate code base clients
- Classic interoperability issues
- Several groups exploring extending interop
  - Linden Lab sponsored Open Grid Protocol project
  - OpenSimulator.org Hypergrid project
  - Intel's Cable Beach
  - MXP work in the OpenSimulator community.

Second Life model

- Virtual space simulated with rich host side physical simulation, serving up stream of graphic model updates to clients
- Rich set of visual effects delegated to client – Almost no physical modeling on client
- Virtual Space servers connect to a broad set of back end servers for asset, inventory, messaging services
- Fine grained, additive geometry 3d model, with fully dynamic user modifiable content
- Content and messaging on extensible XML serialization of a dynamic typed model

Croquet / QWAQ / Squeak

- Croquet is built on top of Squeak distributed Smalltalk environment
- Distributed Objects as building blocks of virtual shared spaces
- Peer to Peer simulation approach
- Extended to form QWAQ

Wonderland

- Sun sponsored Java based open source project
- Distributes Java objects to clients
- Blend of served spaces and distributed objects
- Rich media model

Other notables

- Forterra's OLIVE
  - Lockstep coupled simulation
  - Identical simulation model in all clients
  - Most control flows as low level events which are then simulated in all clients
- Blizzard, and Game derived approaches
  - Content almost entirely client side
  - Bulk of simulation locally managed, with servers managing state coupling and enforcing "no-cheating" validation of Shared State
- Metaplace
- Sirikata
- Steady stream of new entrants at various function points in the environment

Related Standards

- Collada
  - Primary 3d content interchange format
  - Mostly static import/export focus
  - Gradually changing
- OpenID, OAuth
- X.509
- Bi-Directional HTTP / Reverse - HTTP

MMOX Technical approach

- Define virtual spaces as web resources with URI/URL style addreses
  - Well defined set of public RESTful web services
  - Leverage existing web standards for security, naming, transport
  - Work to standardize approaches where existing standards are insufficient
- Provide a set of components focused on the user's stable identity, content and presence/location information
- Provide a rich set of affordances for permitting deployers to offer different content, policies, and deployment approaches
- Start with relatively compatible clouds of service and grow out to a broader set of services over time
- When in doubt, focus on the web/resource style approach to describing and solving problems
- Hand off between clients (parallel to handoff from clients to content plugins) may well be part of the realistic interop world for quite some time
- Explicitly build an extensible, dynamic protocol markup
- This is not a "superset client" approach. Handoffs between clients, or between extension sets in clients is likely
- Use web style content negotiation, content streaming to deliver various types of content across a consistent set of services

Clusters of services

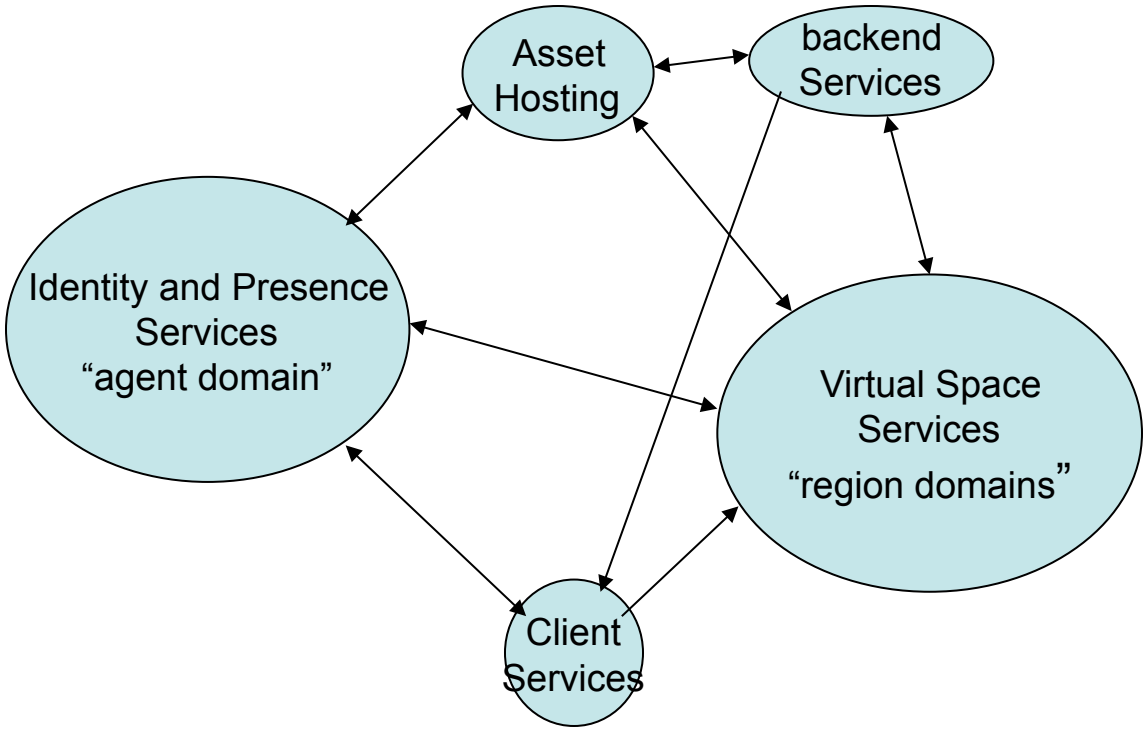

This is one clustering of the services, currently described by the Open Grid Protocol. It may represent one profile for using these services, with others being defined over time

Update Stream pattern

- Managing update streams between components is a core pattern repeated in most implementations
- The content of the stream varies
  - Actions
  - Results of actions
  - Triggers of operation in other components (clients or peer simulations)
  - Graphic model updates
- There are often implicit mechanisms to express interest
- Providing a web style approach to this pattern permits us to leverage URIs, Multiple representations of content, content negotiation, and other common web patterns
- Standardizing this pattern is a long term effort, with large potential benefits

Managing access to virtual spaces

- Much of the Open Grid Protocol work is about two tasks
  1. Managing a user's identity, location and "presence" information
  2. Managing hand off between and access to virtual spaces
- When we define virtual spaces as web addressable resources, this approach makes sense, and allows the use of existing web approaches for many issues, rather then re-creating them.
- The separation of the problem into identity management and virtual space management helps limit the complexity and scope of some of the security and policy issues we need to solve

Managing access to content

- Current OGP work provides less focus on access to content
- Straightforward web service approaches can be adopted
- The hard issues are security and policy driven
- Current specifications in the proposed charter provide the security underpinnings to apply policy based solutions to these problems
- We do not expect to solve unsolvable problems

Working Group Approach

- Two tracks
  - Focus on existing ecosystem with clear needs
  - Refine emerging proposals as fodder for future concrete work
- Linden Lab Open Grid Protocol offers a staring point to solve real problems
  - Multiple code bases
  - Existing interoperability problems to solve
- The next layer out requires additional work before it is ready for spec writing and code
  - Foster this work
  - Provide a clear path to incorporate mature ideas via charter revision if appropriate

Concrete work products

- Open Grid Protocol Derived Specs
  - Protocol encodings
  - Core services
  - Suite of specs necessary to enable interop within the emerging Ecosystem, and grow out from these
- Core problems, Use Cases and technology document
  - Grow into set of use cases for validating work
  - Capture immature but important work
  - Basis for creating new spec drafts and feeding them into possible re-chartering as they mature
- Expect a specify, bake, refine approach, with possible charter updates if new work matures

Dynamic, extensible protocols

- Highly distributed long running deployed services
- Focus on markups which permit us to define extensible, nested low level protocol elements
- Balance profiles of elements, which combine to describe specific forms of interoperability with required elements permitting content negotiation and certainty of understanding base requests
- Similar to the HTTP transport, with registered content types approach

Focus on pragmatic solutions

- Event Queue – Not elegant, but solves the core problem
  - The specifications model a two way pipe
  - The Event queue provides a specific mapping of this onto current, firewall friendly technology
  - Future specifications may leverage emerging solutions
- Homogenous edge touching regions
  - Expose properties such as variable geometry regions,
  - Permit clients to only deal with these issues in limited circumstances
- Accept heavy weight handoffs when they provide useful results
- Start from learned lessons in running real deployed solutions and grow out from this base

Security Considerations

- The working group proposes to address X.509 mechanisms for managing component to component authentication
- The working group is investigating the use of OpenID and Oauth
- The working group is define the affordances necessary to permit a variety of access, authentication, rights management and permissions polices to be used in conjunction with the base specifications.

Out of Scope

- The MMOX working group does not intend to mandate specific security, DRM, Authentication or Authorization policies, beyond ensuring the affordances are in place to permit policies to be in place, and determine the policies which are in effect, or absence of same
- The MMOX working group does not intend to address social, political or legal structures associated with the creation and management of deployed immersive spaces, beyond the level of use cases derived from actual deployed environments
- There are a great many interesting ideas which have been proposed by various interested parties which are not currently embodied in code, or related to existing running systems. We look forward to future work when such ideas are more concretely defined.

Realistic Expectations

- We will provide real, concrete value for significant portions of the virtual worlds landscape
- The creation of additional factorings of the services, and additional extension points based on these factorings will increase the scope of services which can interoperate based on this work.
- Bridging the broader gaps between some of the existing approaches to immersive spaces will require substantial work, and falls outside the scope of the working groups concrete deliverables
- The Working Group plans to foster emerging work and will contemplate charter revisions if new work enters sufficiently mature states, as evidenced by actual code and emerging draft documents.
- The outcome of the IETF 74 BOF should provide sufficient progress towards rough consensus to yield revised charter proposals and a focused discussion so we can charter at or before IETF 75