

## Extending Locales: Awareness Management in MASSIVE-3

James Purbrick, Chris Greenhalgh  
*School of Computer Science and Information Technology,  
University of Nottingham, UK.  
Email: {jcp, cmg}@cs.nott.ac.uk*

### Abstract

*Locales (as first embodied in the SPLINE system) are a very good approach to constructing and organising large virtual environments, especially for distributed and multi-user systems. In the MASSIVE-3 system we have adopted the locale approach, and extended it in a number of significant areas: we have integrated the notion of awareness (drawing on our previous work with awareness management in virtual environments); we have added support for abstractions, i.e. simplified or alternative representations of locales or sets of locales; we have integrated data management based on functional and organisation factors; and we have introduced a flexible and general framework for defining dynamic locale selection policies. This paper describes our approach, illustrates our current implementation with a sample application (a virtual sports arena), and gives a qualitative evaluation of our approach to locales.*

### 1. Introduction

This paper is concerned with awareness management in MASSIVE-3, which is a distributed multi-user virtual environment system (i.e. a Collaborative Virtual Environment or CVE). Awareness management – also called interest management – is concerned with identifying which objects and information in a system are of relevance to a particular observer. This can be used to determine: which virtual world state should be replicated on the observing machine; which network communication (e.g. updates, audio or video streams) should be received by the observing machine, and with what data fidelity and network quality of service; which virtual objects should be rendered, and at what level of detail; which virtual sounds should be played out to the user, and at what volume; and other medium- and application- dependent adaptations to the available information

All virtual environment systems have some form of awareness management, although in many cases it is implicit - a side effect of other design decisions - rather than explicit and reasoned. For example, a system which provides undivided virtual worlds will use virtual worlds

as the basis for awareness management: the virtual world that the user is in will determine what data is replicated, what network communication is received and what the user sees and hears.

MASSIVE-3 is the third generation of CVE system which we have developed, with a continual emphasis on powerful and flexible awareness management. As described in this paper, MASSIVE-3 builds upon and integrates several elements of the state of the art in awareness management, including locales, explicit awareness, abstractions, and functional and organisational support.

The next section reviews related work, and highlights concepts and techniques that we are employing here. Section 3 describes how we have combined these into an integrated model of awareness management used in MASSIVE-3, while section 4 explains key details of our current implementation. Section 5 illustrates the system's (and the model's) capabilities through a demonstration application. Section 6 then reflects upon our approach. Finally, section 7 identifies areas for future work and gives our conclusions.

### 2. Background and Related Work

#### 2.1. MASSIVE-1 and MASSIVE-2

Our own previous systems have been based on the so-called Spatial Model of Interaction, an explicit model of awareness in space-based systems such as CVEs. MASSIVE-1 [6] is based on an explicit measure of desired awareness, where 0.0 indicates no awareness (i.e. not replicated, not visible and/or not audible) and 1.0 indicates full awareness (i.e. fully replicated, highest available fidelity, best level of detail, etc.). We have brought an explicit notion of awareness to MASSIVE-3, to allow the system to reason about awareness in a reflective fashion.

MASSIVE-2 [7] went beyond MASSIVE-1 to include contextual effects (e.g. rooms) and abstractions (i.e. alternative and simpler representations of virtual world content). This was realised over an IP multicast

infrastructure. Contextual effects, and abstractions are also supported in MASSIVE-3, although the details differ.

Both of these previous systems used a single global coordinate system per world, and had looser notions of content grouping than are provided by MASSIVE-3.

## 2.2. SPLINE

Locales, introduced in the SPLINE system [1], provide a powerful and expressive means of structuring and composing a virtual world. The work presented here is based on and extends the notion of locales as found in SPLINE. The key concepts of locales are reviewed below.

- In SPLINE a locale is the fundamental unit of world composition. Each locale typically corresponds to a distinct region of the virtual world, such as a room, corridor, open space or vehicle, and can contain virtual objects as well as users' embodiments (their representations with the virtual world).
- A key feature of the locale approach is that there is no single global coordinate for the whole virtual world. Instead, each locale defines its own independent coordinate system.
- A complete virtual world is composed by linking together a number of locales. Each link from one locale to another includes a 3D transformation which defines the relationship between the locales' coordinate systems. Note that these links and transformations can be asymmetric.
- These transformations can be used to construct non-Euclidean worlds as demonstrated in Diamond Park [1]. This includes buildings that are larger on the inside than on the outside, and transportation obelisks which allow users to move rapidly between points in the virtual world that are otherwise distant.
- The system's built-in awareness management (replication) policy assumes that the system should replicate, monitor and display the user's current locale and its immediate topological neighbours only, i.e. those which are linked directly to the current locale.

Section 3 will describe how we have built on and extended this framework.

## 2.3. NPSNET

NPSNET-IV [8] implements a spatial awareness management scheme which tiles a virtual world into fixed-size hexagonal cells, and uses these as the units of communication. In addition, it describes the use of functional and temporal class as further factors which can inform the awareness management process. In the work presented here we have included the notion of functional class, as described in section 3.3, but employ locales to

structure the spatial aspects of awareness, giving greater flexibility and dynamic control.

## 2.4. PARADISE

The PARADISE system [10] introduces "projection aggregations" to structure awareness management and communication. These combine spatial subdivision (based on octrees) with organisational scope (based on the organisational structure of the entities in question, in that case military units). A significant aspect of their approach is that everything is (conceptually) an aggregate, i.e. a composite representation which can (in principle) be expanded to give more underlying detail in terms of sub-components. In our terminology this is a form of hierarchical abstraction, which is also supported in the approach presented here. However our approach situates this in a locale-based framework (for more flexible spatial organisation), and integrates an explicit notion of awareness.

## 2.5. Other Systems

The above are the five main systems and approaches which relate most directly to this work, and from which we have synthesised a new hybrid approach. As already noted, some form of awareness management is implicitly present in all virtual environment systems, even if is only a notion of "virtual world". Several other systems have some region-like notion (e.g. [3], [4]), which we believe can be subsumed by a mapping on to locales. A less straight-forward comparison can be made with visibility-based awareness management as exemplified by the RING system [5]; here, the individual rooms are analogous to locales, but the actual awareness management process reasons about line-of-sight visibility. We argue that this could be encapsulated within a custom locale selection policy that could be integrated seamlessly with our model (see sections 3.5 and 4.3).

## 3. Model

We begin by briefly defining the main components of our extended locale-based awareness management model. We then describe its main features in more detail.

### 3.1. Main Components

As in SPLINE, a MASSIVE-3 world is composed spatially from multiple *locales*, each of which defines its own independent co-ordinate system. Locales are linked together using *boundaries*, which establish a connection and (unidirectional) geometric relationship between two locales. A boundary is defined by the following properties:

**Table 1. Properties of a boundary.**

Target	Locale this link references
Transform	Transform used to place objects in target co-ordinate system
Polygon	Polygon which user can move through to cross boundary
Type	Whether the target is adjacent or a parent or child locale (see section 3.4)
Awareness	Effect of boundary on awareness

Each locale can be further divided into several *aspects*. These are the fundamental unit of awareness management in MASSIVE-3. All distributed data in MASSIVE-3 must be contained in a particular aspect of a particular locale. An aspect is defined by the following properties:

**Table 2. properties of a locale aspect.**

Functional class(es)	Type or class of content, e.g. EMBODIMENT, TERRAIN, BASE
Organisational scope(s)	E.g. PLATOON-A, BATTALION-3, TEAM-A, ...
Fidelity	“Quality” of this aspect, e.g. resolution, accuracy, ...
Cost	Cost of the aspect, e.g. some function of state size, bandwidth, rendering, ...
Contents	Replicated data, including objects, attributes, links, annotations

The following sub-sections give further explanation of boundaries, functional and organisational scopes, fidelity, and locale selection.

### 3.2. Boundaries and Awareness

The awareness properties of a boundary allow it to represent, for example, closed vs. open doorways, windows that impair awareness, or video screens that enhance awareness. This awareness information could supplement awareness derived from the spatial model in much the same way as third party objects do in MASSIVE-2 [7]. Alternatively, these awareness modifiers allow basic reasoning about awareness to be performed on locales as a whole. Note that awareness is often medium-dependent (e.g. different for audio vs. visual).

### 3.3. Functional and Organisational Scopes

[8] proposes the use of functional class for partitioning awareness and communication, an idea which we have integrated here. Typically, different functional classes of data (e.g. different media) will only be of interest to a subset of observers, and this can be exploited if those functional classes are separated within the awareness

management framework. A single aspect can contain one or more functional classes. At one extreme a single aspect can contain an entire locale; at the other extreme there can be a separate aspect for every functional class.

We define a specialised functional class called “BASE” which must exist for every locale. The base aspect contains all boundaries to other locales and all references to the other aspects of the locale. This allows the system to learn about locales and world topology just by replicating base aspects, which will normally be relatively compact and static. When the topology has been mapped out and the available aspects are known the system can choose which non-base aspects to join.

As well as standard (built-in) functional classes, additional classes may be defined by an application. For example, a “mark-up” class could be used to contain information that annotates objects in other aspects, allowing a user performing online editing to view features that would normally be hidden.

[10] adds the notion of organisational scopes, orthogonal to functional class. These will normally be defined by the application to reflect organisational associations between objects. For example, a game may define an organisational scope for each team, while a military simulator may define an organisational structure to mirror the military hierarchy. The system could also make use of organisational scope to split overcrowded aspects into sub-parts.

Dividing the contents of a virtual environment by space (into locales), and by function and organisation (into aspects) allows a fine grained grouping of objects within the environment. This is considered further in section 6.2.

### 3.4. Fidelity and Abstractions

There are cases in which an observer is interested in a particular locale or aspect but does not have the resources (of whatever form) to cope with it. Our model allows a locale to contain aspects with the same functional and organisational scopes, but different fidelities. An observer can then choose the aspect with the fidelity that suits their requirements and available resources. We call these lower fidelity representations “abstractions”, as in [7] (c.f. “aggregates” in [10]).

Abstractions may display data at a lower resolution, for example using geometries with low polygon counts or objects with low update rates. Alternatively, abstractions may contain composite representations of multiple objects. A currently implemented embodiment abstraction (a re-implementation from [2]) scales the height of a single humanoid object to represent the number of users in the locale (this can be seen in section 5 and figure 4).

For the purposes of rendering, aspects with different functional or organisational scopes are complimentary

(providing additional non-overlapping information) while aspects that differ only in fidelity are alternatives (providing variants of the same underlying information). So two aspects differing only in fidelity will not be rendered at the same time. However, aspects that differ only in fidelity *may* be replicated at the same time, to allow rendering to switch quickly between the alternatives without waiting for the new aspect to be loaded across the network.

In addition to abstractions of single locales we also allow abstractions of multiple locales, or abstractions of abstractions (as in [10] and [7]). An adjacent-type boundary between two locales indicates that those locales are linked in some (unspecified) way. However a parent locale (defined by parent/child boundary relationships) is regarded as an abstraction of its child locales. In other words we have introduced a mechanism for nesting topologies. The outside view of the arena in section 5 is an example of this.

### 3.5. Awareness Management: Selection Policies

The main “work” of awareness management is to choose which information should be replicated, which communication received, which information rendered, etc. In the approach presented here awareness management is performed primarily at the level of locales and aspects. E.g. the awareness management system must choose which locales (and which aspects of those locales) should be replicated and which of these should be rendered. Note that these choices are at least partially independent: one collection of representations may be replicated over the network, another (subset) rendered visually and a third (subset) rendered as audio.

SPLINE’s build-in awareness management policy replicated and rendered the user’s current locale and its immediate neighbours. We have abstracted away the exact policy to give a policy-independent framework into which application-defined policies can be dynamically inserted, as required. In this framework all policies have the same basic structure, outlined below.

1. Begin new selection round: reset budgets and internal state.
2. Sort known locales and aspects according to a policy-specific ordering criterion (i.e. assigning priorities)
3. Iterate over each aspect in the sorted list: determine whether to select it, based on aspect/locale characteristics, choices made so far, and remaining budget(s).

For example, SPLINE’s adjacent neighbour strategy can be seen as a special case of a more general strategy, which replicates all locales within N steps of the current locale. This policy has no specific budget, sorts the known locales according to topological distance from the

current locale, and accepts all locales (aspects) whose distance is 0 (current) or 1 (neighbour). Section 4.3 gives examples of the standard selection policies currently available in MASSIVE-3.

## 4. Implementation Details

In this section we give a brief outline of the MASSIVE-3 system, before focussing on the realisation of locales, aspects and boundaries, and the standard locale selection policies.

### 4.1. Overview

MASSIVE-3 is a development of the HIVE project Kernel (HIVEK), which combines work on consistency in CVEs from the University of Reading [9] with world structuring and awareness management as described here. HIVEK provides, first and foremost, a shared data service, designed specifically for CVEs.

The two most fundamental classes in MASSIVE-3 are *Agent* and *Environment*. An Agent encapsulates the notion of activity or computation, and each Agent will correspond to exactly one process or thread within the system. An Environment encapsulates the concept of shared data and is the system realisation of a locale aspect. An Environment implements a tree of typed data items, which are normally a (partial) scene graph with additional annotations for semantics and behaviours. A complete running system will comprise a number of Agents, typically distributed across a number of machines, interacting and communicating via one or more shared Environments. Environments are fully replicated, (via an initial non-blocking streamed transfer). Consistency between these replicas is maintained by a logical multicast-based distribution of updates to all Environment replicas.

### 4.2. Locales, Boundaries and Aspects

As noted above, each aspect of a locale is realised as one Environment object. One of the standard data types supported by an Environment is a direct realisation of a boundary, having the properties described in section 3.1. Each locale will have at least one aspect (which includes the base functional class). This will be realised as one Environment. All of the boundaries to other locales will be represented by boundary data items within this Environment. Where separate aspects exist these are referenced via further (specialised) boundaries from the base aspect.

An application joins a virtual world by locating a starting locale (implicitly, its base aspect) and replicating it. Once this is replicated the application discovers the boundaries which identify the locale’s other aspects and

which link that locale to its immediate neighbours. Using this boundary information the application constructs a *locale map* which is a complete representation of known locales and aspects. This locale map is passed to the current replication policy (an arbitrary locale selection policy), and the aspects selected by that policy are then replicated. As new locale base aspects are replicated (and as locales are changed, e.g. to add or remove boundaries) the replication policy is re-invoked for the updated locale map. Thus the system incrementally discovers and manages locale replication.

The renderer has access to the same current locale map, but applies its own selection policy to it, to select an arbitrary subset of the replicated aspects for rendering. The rendering policy will make the choice between available levels of fidelity (i.e. abstractions).

### 4.3. Locale Selection Policies

There are a number of standard policies available in MASSIVE-3, and others can be added by the application. The built-in policies differ in their spatial selection, but operate identically on functional and organisational scopes, selecting functional and organisational classes based on an explicit set to accept. Note that policies can be used in combination (e.g. "policy A OR policy B") to implement compound selection criteria.

Examples of the standard policies are described below.

- N Step Selection.** This is the generalisation of nearest neighbour selection described in section 3.5. Increasing the size of N allows users to see and interact with more of the world at one time. Allowing different values of N creates a potentially adaptive strategy: users on powerful machines or with fast network links can choose a higher value to fully utilise their resources, while users struggling with poor frame rates or a congested network can decrease N to concentrate resources on their immediate vicinity. Note that some parts of Diamond Park [1] *require* limited values of N in order to work as expected; in these situations additional awareness information on boundaries could be used to ensure correct rendering.




Figure 1. N-step selection examples.

- N Nearest Selection.** While N step replication can provide a degree of flexibility and performance control it is very coarse grained: increasing N from 1 to 2 may result in many additional locales and aspects being included in the replicated or rendered world model. An improvement to this method is to implement a strategy that includes only the N locales closest to the user in the local world model (potentially independent of the topological distance) limiting it to an exact number of locales.
- N Most Aware, and Cost/Benefit Selection.** The previous policies only make use of locale topology and distance. The N Most Aware policy replicates the current locale and the N locales with the highest awareness from the current locale. An enhancement of this strategy is based on a cost/benefit model, where an aspect's cost, published in its base aspect, is considered alongside its awareness: awareness over cost gives the value of the representation (which is the sorting criterion used by the policy). The policy selects as many high value aspects as possible given a finite budget (which typically represents some part of available system resources).

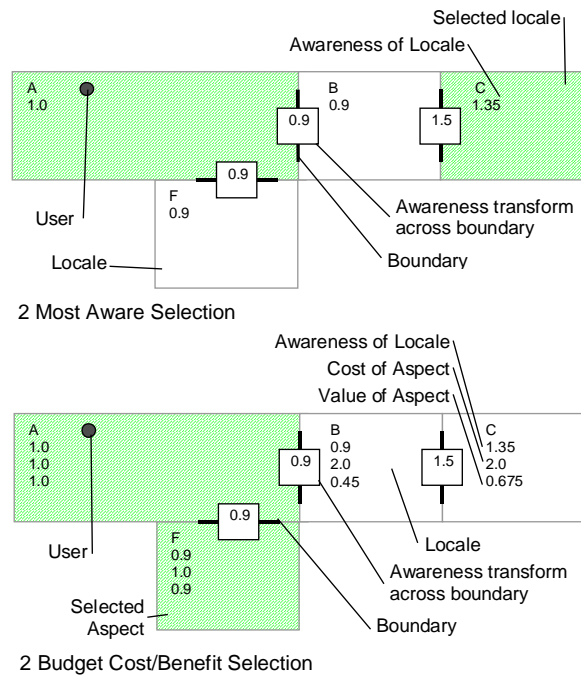


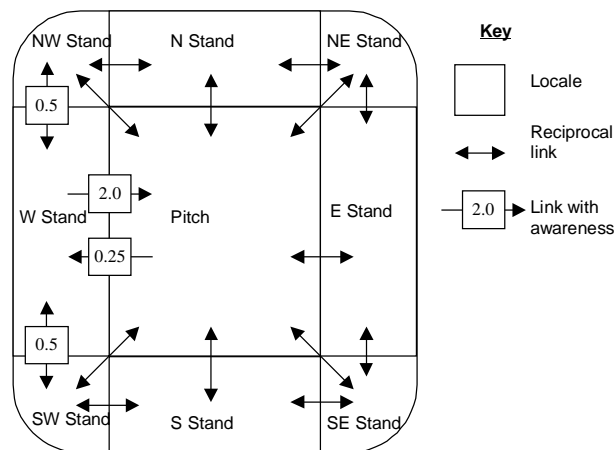
Figure 2. N Most Aware and Cost/Benefit selection examples.

## 5. Demonstration

To demonstrate our approach we have constructed a world representing a virtual sports arena. This example is particularly appropriate as it is a densely populated, but

very open environment, and so cannot be easily coped with by systems that use visibility (RING) or arbitrary spatial subdivision (NPSNET). In addition there are objects of the same functional type (embodiments) that should not be treated equally (the players are – in some senses – more important than the spectators).

The world is split into 9 locales, each with 6 aspects: a base aspect containing boundaries and references to other aspects; two embodiment aspects containing users' representations within the world for each team; a terrain aspect containing the stadium geometry; and two embodiment abstractions (again one for each team). The locales and boundaries that make up the world topology are shown below.



**Figure 3. Locale structure of the example arena.**

Each stand locale is linked to the stand next to it and to the pitch. All links are reciprocal allowing users to move between them freely. However the awareness across these links is not symmetrical: stand-to-pitch links have a higher awareness than pitch-to-stand links (as shown on the west stand in the diagram). This linkage means that spectators are very aware of the players on the pitch and have peripheral awareness of neighbouring crowds, while the players have reduced awareness of the spectators.

The arena also has an overall locale (not shown) that represents the arena to the rest of the virtual world. All external boundaries to the arena link to this overall locale, while all links between this overall locale and the arena's component locales are parent/child links. Locale selection policies can make use of this hierarchical information, for example a user outside the arena might render only this overall representation and not move down parent-child links.

To demonstrate the flexibility of our approach we will look at the consequences of using the different selection policies defined in MASSIVE-3 to select locales for replication and rendering within the arena.


If a spectator used a 1MostAware policy they would see the pitch, a 2MostAware policy would select the pitch

and the locale they occupy, while higher values of N would start to reveal the neighbouring stands.

A spectator using a 1Step selection policy (as used in SPLINE) would see the pitch and neighbouring stands, moving to 2Step replication would include the entire stadium. Using the NNearest policy, the pitch, and then the stands, would be progressively revealed as N is increased.

The problem with these simple approaches is that if a user is unable to cope with all of the embodiments in the stands then they must also discard large parts of the stadium. To address this problem we require MASSIVE-3's facilities to select different aspects of locales in different ways. In the stadium example a user might use a 2Step policy to select terrain aspects of the world, allowing them to view the entire stadium structure, while using a 1Step policy to select embodiment aspects and so see the players on the pitch and spectators in neighbouring locales. If this was still too much to cope with, a 1MostAware policy used for embodiments would limit the awareness of other users to the players on the pitch. If organisational scoping was used organisations could be set up for the players, fans for each side and potentially even smaller groups of friends, allowing more refined choices to be made.

These functional and organisational scoping mechanisms allow the user to see the stadium and embodiments of interest to them, however they are left without a sense of the number of people present in the stadium, which looks largely empty. The abstract embodiment aspects used in this demonstration address this problem: the abstractions represent the number of people in a locale, in this case using a single humanoid, which changes size with the occupancy of the standard embodiment aspect. This requires much less bandwidth and geometry than the many individual embodiments in the full representation, while still conveying the size of the gathering.



**Figure 4. A view from the south stand.**

The screenshot above shows a view of the demonstration world from the south stand. A 2Step selection policy is used to replicate terrain and embodiment abstraction aspects, allowing the user to see the entire stadium and low fidelity abstractions of the fans (which are the large humanoids in the distance). A 1Step selection policy is used to replicate full fidelity versions of embodiment aspects that contain fans from the same team as the user (shown in the foreground). Players on the pitch are replicated using a 1MostAwarePolicy for embodiment aspects. In this example all replicated aspects are being selected for rendering at the highest available fidelity. The system also allows the rendering policy to select a dynamic subset of this.

## 6. Evaluation

In this section we give a qualitative evaluation of the approach presented and the specific features which we have integrated. We consider in turn: locales, aspects, abstractions, awareness and general selection policies.

### 6.1. Locales

The locale approach is currently unique in allowing the construction of virtual worlds with non-Euclidean structures (e.g. virtual mirrors, toroids, worm-holes, directed topographies). It also supports high local precision even in very large worlds. For these reasons we have selected a locale approach over the more common single Cartesian co-ordinate system.

Using locales also gives a very well defined basis for awareness management and communication: every data item is in exactly one locale at a time, and each locale can be readily designed to reflect the intrinsic and semantic structure of the virtual world (e.g. rooms, buildings, vehicles, regions). We argue that this explicit grouping of items (in locales) facilitates awareness management, communication (including group communication), and resource management, as well as authoring and world composition.

However, locales on their own provide only a spatial form of awareness management, and no notion of abstractions or alternative representations.

### 6.2. Aspects

In our initial implementation we had only locales, and not aspects. This resulted in two main problems:

- Replicating a locale was “all or nothing”. In particular, an observer that was interested in only part of a locale (e.g. its boundaries, its static geometry) would have to replicate the whole locale content.

- There was no well-defined method for integrating abstractions, i.e. no principled means of introducing alternative representations of locales. Even where abstractions could be introduced they could only be applied to the locale as a whole, whereas we might want to abstract only part of the locale (e.g. the embodiments in it).

Introducing aspects as an explicit concept has addressed both of these issues: the base abstract of a locale is a light-weight method of finding locale topology and general information, and aspects with different fidelities support abstractions of arbitrary subsets of a locales content.

### 6.3. Abstractions

Similarly, a simple abstraction (such as that used in section 5) can give a user a limited awareness of the presence and activity of others within a locale. The abstraction’s complexity and bandwidth requirements are independent of the number of users in that locale (unlike the full-fidelity version, which has  $O(N)$  requirements).

The possibility of nested topologies (i.e. hierarchical abstractions) is extremely powerful for creating structured virtual worlds which span extremely large ranges of scale (e.g. planet, country, city, district, building, room); without such a mechanism it would not be possible to gain an overview of multi-locale structures.

### 6.4. Awareness

Making awareness explicit is an effective means of encoding more of the semantics of the virtual environment. For example, it can represent the difference between an opaque doorway, a translucent window and an open space. This is particularly important to inform the rendering or presentation process when different locale selection policies are being used. For example, SPLINE examples such as the Desert House or the Obelisks [1] would “break” with some replication policies, by rendering locales which have no mutually consistent single co-ordinate system. Awareness annotations on boundaries can be used to preserve these presentational constraints, to allow an open choice of replication policies while maintaining a consistent appearance.

Explicit awareness can also be used to facilitate user interaction and resource management, as demonstrated in MASSIVE-1 and MASSIVE-2 [6,7].

### 6.5. General Selection Policies

By defining a general framework for locale and aspect selection we have gained the following benefits compared to any single strategy:

- We have de-coupled the various uses of awareness management – replication, communication, rendering, etc. – and allowed each to be handled in a specific and appropriate manner.
- Different observers can tailor their own policies to reflect the particular resources they have, e.g. how powerful their machine is, how good their network connection is, etc. This provides a key mechanism to support deployment on heterogeneous hardware platforms and infrastructures.
- This approach allows the open evolution of selection policies, so that we can integrate current and emerging resource management and quality of service based strategies into our system.

## 7. Conclusions and Future Work

We have presented a new approach to awareness management in CVEs, which integrates elements from several state of the art systems, including locales, functional and organisation scoping, awareness, and abstractions. This approach has been implemented in the MASSIVE-3 system, a demonstration of which has been presented. We commend locales as a general approach to virtual environment composition, but suggest that other researchers strongly consider provision for other issues of awareness management, including those that we have addressed in this paper.

We are continuing to develop and enhance our implementation of the model, in particular: refining the support for hierarchical abstractions, creating general and reusable abstractions, and developing more sophisticated selection policies (e.g. for pro-active resource management, and for visibility-based culling).

Our approach provides a sound basis for system scalability based on world subdivision and abstraction. However, further issues need to be addressed at the locale or aspect level, including: group adaptive controls, e.g. group flow control and congestion control (for communication), and adaptive level of detail. Related work is also required on the automated management of world structure, e.g. automatic generation of aspects and abstractions, and locale subdivision.

## Acknowledgements

We gratefully acknowledge the funding support of the projects “large-scale real time multi-user virtual reality” (UK EPSRC research grants GR/L 01527 and GR/L 02074) and “the nature and utility of persistence in virtual environments” (UK EPSRC research grant GR/M09223).

## References

- [1] Barrus, J.W., Waters, R.C., and Anderson, D.B. (1996), “Locales: Supporting Large Multiuser Virtual Environments”, in *IEEE Computer Graphics and Applications*, 16 (6), Nov., 50-57.
- [2] Benford, S. D., Greenhalgh, C. M., Lloyd, D. L., *Crowded Collaborative Virtual Environments*, in *Proc. 1997 ACM Conference on Human Factors in Computer Systems (CHI97)*, Atlanta, Georgia, March 22-27, 1997, ACM Press.
- [3] Broll, W., “SmallTool – a toolkit for realizing shared virtual environments on the internet”, *Distributed Systems Engineering*, Volume 5, Number 3, September 1998, pp. 118-128, ISSN 0967-1846.
- [4] Tapas K. Das, Gurminder Singh, Alex Mitchell, P Senthil Kumar, Kevin McGee, “NetEffect: A Network Architecture for Large-Scale Multi-user Virtual World”, in *Proc. IEEE Sixth Workshops of Enabling Technologies: Infrastructure for Collaborative Enterprises*, June 18-20, 1997, MIT, Cambridge, MA.
- [5] Funkhouser, T.A., “RING: A Client-Server System for Multi-User Virtual Environments.” *Computer Graphics*, 1995 SIGGRAPH Symposium on Interactive 3D Graphics, Monterey, CA., April 1995, pp. 85-92.
- [6] Chris Greenhalgh and Steve Benford, MASSIVE: A Virtual Reality System for Tele-conferencing, *ACM Transactions on Computer Human Interfaces (TOCHI)*, Volume 2, Number 3, pp. 239-261, ISSN 1073-0516, ACM Press, September 1995.
- [7] Greenhalgh, C. M. and Benford, S. D., *Supporting Rich And Dynamic Communication In Large Scale Collaborative Virtual Environments*, *Presence: Teleoperators and Virtual Environments*, Vol. 8, No. 1, February 1999, pp. 14-35, MIT Press.
- [8] M.R. Macedonia, M.J. Zyda, D.R. Pratt, D.P. Brutzman, and P.T. Barham, “Exploiting Reality with Multicast Groups: A Network Architecture for Large-scale Virtual Environments”, in *Proc. 1995 IEEE Virtual Reality Annual International Symposium (VRAIS95)*, 11-15 March, 1995, RTP, North Carolina.
- [9] Roberts, D.J., and Sharkey, P.M. (1997), “Maximising Concurrency and Scalability in a Consistent, Causal, Distributed Virtual Reality System, Whilst Minimising the Effect of Network Delays”, *Proceedings of the 6th International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET-ICE)*, June 18-20, 1997, Cambridge, Massachusetts, USA, IEEE Computer Society, ISBN 0-8186-7967-0, pp 161-166.
- [10] S.K. Singhal and D.R. Cheriton, “Using Projection Aggregations to Support Scalability in Distributed Simulation”, in *Proc. 1996 International Conference on Distributed Computing Systems (ICDCS'96)*, 1996.