
Design issues for Peer-to-Peer Massively Multiplayer Online Games

Lu Fan*

School of Computing,
Edinburgh Napier University,
Edinburgh EH10 5DT, UK
E-mail: L.Fan@napier.ac.uk

*Corresponding author

Phil Trinder and Hamish Taylor

School of Mathematical and Computer Sciences,
Heriot-Watt University,
Edinburgh EH14 4AS, UK
E-mail: P.W.Trinder@hw.ac.uk
E-mail: H.Taylor@hw.ac.uk

Abstract: Massively Multiplayer Online Games (MMOGs) are increasing in both popularity and scale, and while classical Client/Server (C/S) architectures convey some benefits, they suffer from significant technical and commercial drawbacks. This realisation has sparked intensive research interest in adapting MMOGs to Peer-to-Peer (P2P) architectures. This paper articulates a comprehensive set of six design issues to be addressed by P2P MMOGs, namely Interest Management (IM), game event dissemination, Non-Player Character (NPC) host allocation, game state persistency, cheating mitigation and incentive mechanisms. Design alternatives for each issue are systematically compared, and their interrelationships discussed. We further evaluate how well representative P2P MMOG architectures fulfil the design criteria.

Keywords: P2P; peer-to-peer; MMOGs; massively multiplayer online games; IM; interest management; event dissemination; task distribution; distributed storage; anticheating; collaboration incentives.

Reference to this paper should be made as follows: Fan, L., Trinder, P. and Taylor, H. (2010) 'Design issues for Peer-to-Peer Massively Multiplayer Online Games', *Int. J. Advanced Media and Communication*, Vol. 4, No. 2, pp.108–125.

Biographical notes: Lu Fan received the BSc (2003) Degree in Computer Networking from Nanjing University of Science and Technology, Nanjing, China, and the MSc (2005) Degree with distinction in IT Software Systems and the PhD (2009) Degree in Computer Science from Heriot-Watt University, Edinburgh, UK. He is currently a Research Associate in Computer Science at the School of Computing, Edinburgh Napier University, Edinburgh, UK. His research interests include the design of distributed systems, peer to peer architectures, GRID and cloud computing technologies, massively multiplayer online games, and pervasive e-healthcare systems.

Phil Trinder is a (Full) Professor of Computer Science at Heriot-Watt University in Edinburgh. His primary research interests are the design, implementation and evaluation of high-level parallel and distributed programming languages. He has held eight major research grants to support this activity, and has over 80 refereed journal and conference publications. He received his BSc Hons Degree from Rhodes University in South Africa in 1983 and a DPhil from Oxford University in 1989 for research into parallel functional databases.

Hamish Taylor is a Lecturer in Computer Science at Heriot-Watt University in Edinburgh. Prior to that he worked as a Research Assistant at QMW in London formalising automated reasoning and before that at Heriot-Watt on implementing parallel logic programming languages. He holds an MA and MLitt from Cambridge University in Philosophy and an MSc and PhD in Computer Science from Heriot-Watt University. He lectures on distributed systems, network applications and internet technologies. He has held research grants on modelling parallel databases, visualising web data from shared caches and intelligent information discovery. His current research interests are in the design of distributed systems especially peer to peer architectures.

1 Introduction

Massively Multiplayer Online Games (MMOGs) enable thousands of players to interact simultaneously in a persistent game world over a network. C/S architectures have been predominantly employed for traditional MMOGs, because they are relatively easy to secure and implement (Mulligan and Patrovsky, 2003). However, with widespread use of the internet, MMOGs are becoming increasingly popular, and the participation has increased dramatically, reaching hundreds of thousands of players (Woodcock, 2005). At such scales, a C/S architecture exhibits various technical and commercial drawbacks, specifically in the area of reliability and cost, e.g., costs for server hardware, network bandwidth, housing, cooling, UPS systems and dedicated maintenance staff. These factors have engendered strong research interest in engineering P2P MMOGs (Lu et al., 2004; Hampel et al., 2006; Douglas et al., 2005).

The key challenge in adapting a conventional MMOG to a P2P architecture is to fulfil the functionalities of centralised game servers in a distributed fashion. This entails addressing six essential issues, namely IM (Section 2.1), Game Event Dissemination (Section 2.2), NPC Host Allocation (Section 2.3), Game State Persistency (Section 2.4), Cheating Mitigation (Section 2.5) and Incentive Mechanisms (Section 2.6).

Previous work has investigated an individual design issue, e.g., several IM schemes are compared in Boulanger et al. (2006) and Morgan et al. (2005). Other related work focuses on general requirements of a P2P MMOG, e.g., scalability, availability, interactivity (Schiele et al., 2007; Simon et al., 2007). In contrast, the most significant contributions this paper makes are to articulate a comprehensive set of six design issues to be addressed by P2P MMOGs, to survey various approaches to addressing them (Section 2), and to evaluate representative infrastructures for their integration into a single system (Section 3).

2 P2P MMOG design issues

2.1 Interest Management

The primary requirement for a P2P MMOG is to maintain a consistent, shared sense of virtual space among large numbers of players without a server's support. IM is a classical research topic that was initially addressed by Macedonia et al. in the mid-1990s (Macedonia et al., 1994). The concept of IM originates from two observations: a single player does not need to know about what is happening in the game world as long as it does not affect the player; a player's avatar only has limited movement speed and sensing capability. So, a player's view of the game world can be limited to a comparatively static Area of Interest (AOI), and the player only needs to be aware of game events that occur within its AOI. Generally speaking, existing IM schemes can be classified into three types – a spatial model, a region-based publish/subscribe model and a hybrid communication model.

2.1.1 The spatial model

The spatial model uses the properties of space as the basis for mediating interaction (Benford and Fahlen, 1993). This model is also referred to as the 'aura-nimbus' model (Boulanger et al., 2006; Morgan et al., 2005) because of its key abstractions: aura and nimbus. The 'aura' means the area that bounds the presence of an object in space, whereas the 'nimbus' means the mutual awareness levels between two objects. In other words, object *A* is only able to interact with object *B* when their auras intersect with each other, but *A* is aware of *B* when it is in *B*'s nimbus. So, every object should establish communications with other objects that fall within its nimbus, to prepare for potential interactions.

The advantage of the spatial model is that it allows fine-grained IM in which only necessary messages are transmitted among relevant peers (Boulanger et al., 2006). However, a significant drawback is that it requires all objects to exchange positional update information to identify when AOI collisions occur. For example, in Matsumoto et al. (2005), all players must share the 'frame of reference' to know their location with respect to each other. Furthermore, the frequency of these updates must be sufficient to ensure that AOI collisions can be determined in a timely fashion, which may lead to considerable communication overhead (Morgan et al., 2005). To mitigate this limitation, a Voronoi diagram can be employed to help a player find its neighbouring players in a game world (Buyukkaya and Abdallah, 2008; Hu and Liao, 2004). Each peer is required to construct and maintain a Voronoi diagram by itself, based on the spatial coordinates of neighbours. A peer only needs to keep network connections with its current neighbours, and each peer serves as the 'watchman' for one another in discovering approaching neighbours. This approach is able to reduce the communication overhead induced by a pure spatial model, but it is not ideal in that:

- A Voronoi diagram is vulnerable to the 'circular line-up' problem, which is the worst case when a peer has $n - 1$ neighbours in a diagram of n sites
- The communication overhead is not minimal, because a peer still needs to receive and process messages outside its AOI

- While the communication overhead is reduced, the computation overhead might be increased for users to construct and maintain their Voronoi diagrams.

2.1.2 Region-based publish/subscribe model

This model proposes to support a P2P MMOG using coarse-grained IM by partitioning a game world into static regions. Conceptually, it is analogous to a publish/subscribe model, in which the recipient of a message is limited to only interested participants that reside within the same, or neighbouring region as the sender (Morgan et al., 2005). The main responsibility of a region-based IM mechanism is to determine the regions that intersect a player's AOI, and to form the area-of-subscription for relevant events from the union of the intersected regions.

A region-based model offers several advantages. First, it is simpler and cheaper to compute a player's area-of-subscription than to compute AOI collisions. Second, a region's publish channel maps nicely onto a multicast group, hence gaming events can be exchanged efficiently. Third, because regions are predefined, players may carry out local IM without knowing the positions of other players. However, this model also suffers from two drawbacks:

- It is hard to determine the appropriate size of a region. A region must be of sufficient size to ensure that objects are able to disseminate messages in one region before entering another (Morgan et al., 2005). On the other hand, the granularity cannot be too large, otherwise a player's machine might be overloaded by excessive irrelevant messages.
- Region-based IM does not always work well when objects are unevenly distributed.

2.1.3 Hybrid communication model

As indicated by its name, this model is a mixture of the spatial and the region-based models. MOPAR (Yu and Vuong, 2005) and Meta-Model (Rhalibi and Merabti, 2006) are representatives of recent hybrid IM schemes. Similar to a region-based model, a game world is partitioned into multiple regions, and a super-peer (Yang and Garcia-Molina, 2003) is selected to be responsible for each region. When a player is about to join a region, it finds out the current super-peer working in that region, and updates the super-peer about its moving state. In this way, the super-peer can have a global view of the region, anticipate every peer's position in the near future, and analyse AOI collision events using an algorithm such as Storey et al. (2004). Players, whose AOIs are about to intersect, will be notified by the super-peer to establish direct P2P connections with each other, to be prepared for potential interactions.

The hybrid communication model takes advantage of both the spatial and the region-based models. On the one hand, it facilitates fine-grained IM in each region and reduces more communication overhead for ordinary players than in a pure region-based model. On the other hand, it is relatively simpler to implement and more efficient than a pure spatial model. However, it also has the following drawbacks:

- it may impose high computation and communication workloads on a super-peer for a crowded region, where dynamic zoning techniques may be needed as a remedy to distribute the total workload to a set of super-peers, each taking charge of a sub-region
- a super-peer is potentially a single failure point in a region, so fault-tolerant mechanisms are also needed to provide suitable super-peer back-ups to achieve adequate robustness for the system.

2.1.4 *IM discussion*

Of three IM approaches outlined earlier, the spatial model is the most fine-grained, and as we shall see in Section 2.2, it is a prerequisite for disseminating game events using unicast. However, owing to the lack of centralised control, the communication overhead for a player to establish a global view of the game world may be high. This can be mitigated with a hybrid model that combines the advantages of both the spatial and the region-based models. As hybrid models use super-peer networks, they must address super-peer selection, load-balancing and fault-tolerance issues.

2.2 *Game event dissemination*

While IM focuses on finding out what information is relevant to each player, game event dissemination is concerned with how relevant information is actually delivered to the players. In fact, the choice of a game event dissemination approach is largely determined by the underlying IM mechanisms used in a P2P MMOG.

2.2.1 *Unicast vs. multicast*

The spatial model supports fine-grained IM, which explicitly tells a small set of objects that a player may interact with them shortly. Therefore, the player may establish direct P2P connections just with these objects, and gaming events can be exchanged through unicast communications, e.g., in Matsumoto et al. (2005), Buyukkaya and Abdallah (2008), Hu and Liao (2004) and Hu et al. (2008). Similarly, because super-peers in a hybrid model also provide fine-grained IM services, unicast applies to Yu and Vuong (2005) and Rhalibi and Merabti (2006) as well.

However, a coarse-grained IM scheme only tells a player of some regions to which events he should subscribe. In these regions, there might be a large number of players, so unicast becomes inefficient. Instead, each region is represented by a multicast group, which offers a single medium for any region participant to publish gaming events in, and enables the events to be received by all the other region participants. In this circumstance, multicast technology becomes crucial for game event dissemination.

Traditionally, IP multicast (Deering and Cheriton, 1990) was proposed as a efficient group communication mechanism (Fiedler et al., 2002). However, owing to a number of technological, practical and business obstacles (Diot et al., 2000), IP multicast is not widely available on the internet. As an alternative, Application-Level Multicast (ALM) has been proposed to support similar functionalities, but as an application service instead of a network service (El-Sayed et al., 2003). Structured P2P overlays provide good communication infrastructures for building ALMs. For example,

Bayeux (Zhuang et al., 2001) uses Tapestry (Zhao et al., 2001), CAN Multicast (Ratnasamy et al., 2001b) uses CAN (Ratnasamy et al., 2001a), and both Borg (Zhang and Hu, 2003) and Scribe (Castro et al., 2002) use Pastry (Rowstron and Druschel, 2001a). Currently, some related work directly disseminates gaming events with a general-purpose ALM system, e.g., Lu et al. (2004), Hampel et al. (2006), Dickey et al. (2004) and Iimura et al. (2004), which use Scribe. However, others like Yamamoto et al. (2005), Rooney et al. (2005), Chen and Kalogeraki (2005) and Léty et al. (2004) have proposed their own group member management and multicast tree construction algorithms, which aim at providing better robustness, scalability and load-balancing capabilities.

2.2.2 Problems with ALM

A significant problem with an ALM system is the potential latency issue, as a game event often can be relayed by multiple forwarders before it arrives at the final receiver. This process incurs unnecessary end-to-end delay, especially when the size of a multicast group is large.

To cope with this problem, it has been suggested that a multicast tree can be constructed according to the proximity of peers in the game world instead of the proximity of peers on the network (Xiang-bin et al., 2008; GauthierDickey et al., 2005; Schmiege et al., 2008). In this way, players in the vicinity are employed as the most immediate forwarders, hence multicast messages are sent to close by peers faster, while peers that are further away receive them slower. This approach attempts to provide better game interactivity by exploiting the tolerance of distant players in a game world for weak synchronisations. In addition, an expedite event dissemination mechanism has been proposed in Ahmed and Shirmohammadi (2008), which is able to reduce the overall time that a multicast tree takes to disseminate an event by utilising better the time slot between a forwarder's completion of a relaying task and the generation of the next gaming event.

2.2.3 Event dissemination discussion

Disseminating game events with ALM will typically induce longer communication latency than with unicast. However, efficiency is also a consideration for unicast as a player may not have enough bandwidth to send every game event to large numbers of recipients. Hence, a P2P MMOG requires either a fine-grained IM mechanism that enables a player to unicast game events only to necessary recipients, or a specially designed ALM mechanism that is able to exploit distant players' tolerance of weak synchronisation.

2.3 NPC host allocation

Besides Player-controlled Characters (PCs), there are also considerable numbers of AI-controlled NPCs in an MMOG, which either drive continuing storylines, or combat with PCs as monsters. MMOGs have to supply their game worlds with large numbers of such NPCs as required by game scenarios. Traditionally, NPCs are hosted by a game server, consuming significant processing power and network bandwidth. Therefore, one of the prerequisites for realising a P2P MMOG is to host such NPCs using computing resources that are available on common game participant machines.

2.3.1 *Region-based approach*

Region-based approaches (Lu et al., 2004; Iimura et al., 2004) partition a game world into multiple regions, and assign each region a super-peer, which works as an authoritative server and hosts all the NPC objects within the region. For example, in Lu et al. (2004), a live peer whose peerId is numerically closest to the regionId is selected as the ‘coordinator’ for that region, and in Iimura et al. (2004) a ‘zone owner’ is selected as the first peer that joins the zone.

These approaches have several significant drawbacks:

- because only one super-peer is selected to take charge of a region, excessive computation and communication workloads might be incurred on the super-peer
- their super-peer selection criteria are overly simple, as they do not take into consideration peers’ actual resource availabilities
- these approaches cannot guarantee to fulfil the QoS requirement for game interactivity.

2.3.2 *Virtual-distance-based approach*

The key idea of virtual-distance-based approaches (Bharambe, 2006; Yonekura et al., 2004; Hu et al., 2008) is to allocate an NPC to the player, whose avatar is closest to the NPC in a game world. Because a player that is closest to an NPC is most likely to interact with it, if the player is hosting the NPC by itself, there is no need for the player to communicate with a remote third party. It has been suggested that this approach is optimal for minimising interactive latency and communication overhead (Bharambe, 2006).

Colyseus (Bharambe, 2006) has demonstrated the feasibility of virtual-distance-based object hosting in Quake II, a well-known multiplayer first-person shooter game. The game object manager of Colyseus allocates mutable objects, e.g., NPCs, doors and weapon items, to the closest players. Similarly, A to Z (Yonekura et al., 2004) allocates each player avatar a ‘priority field’, which is analogous to the Mahalanobis distance in the domain of quadratic discriminant analysis (Anderson, 1984) to decide which player can access a shared object in the shortest time. Furthermore, the Voronoi diagram (Hu et al., 2008) discussed in Section 2.1.1 seems inherently suitable for virtual-distance-based NPC host allocation, because it is natural for each player to host the NPC objects within its own Voronoi cell.

Compared with region-based approaches, virtual-distance-based approaches are better at utilising the computing resources of more participant machines. However, they also have the following disadvantages:

- Though it is more likely that a player closest to an NPC will interact with the NPC, it does not mean that other nearby players will not. Contrarily, it is quite usual for a group of players to interact with the same NPC in an MMOG. In this case, all non-host players need to communicate with the host, and it is not guaranteed that the latency for each player is equally small.

- The computation of accurate NPC host allocation can be expensive, and because a large proportion of the players in an MMOG are constantly moving, switches of host may be frequent. Therefore, the overall computation and communication overhead may be still high.
- Cheating may become easier for unscrupulous players that might abuse their hosting of NPC objects to their own advantage. Even worse, because no third party is required in a local interaction, it is rather hard to detect such a breach.

2.3.3 Heterogeneous task sharing

A heterogeneous task-sharing mechanism (Fan et al., 2007) distributes NPCs to game participants according to the amount of computing resource and game interactivity that they can provide. Its system model involves three parties: a work source, a set of resource providers and multiple matchmaker super-peers. The work source is the virtual game world that constantly generates NPC tasks. The resource providers are game participants, which have spare computing resources available on their machines. Finally, the matchmakers bridge between resource requirements and resource availability using distributed task allocation.

Compared with virtual-distance-based approaches, a task-sharing mechanism is better at load balancing, because it takes into consideration each game participant's actual resource availability. As a result, it ensures that NPC tasks are always allocated to capable hosts, which maximises the utility of a P2P system. Furthermore, once an NPC is allocated to a game participant, the hosting relationship remains stable, unless the NPC is destroyed, or the host needs to leave the system. Therefore, NPC task migration among hosts is less frequent than in virtual-distance-based approaches.

However, in some circumstances, it would be possible to have special NPCs like shop owners that were only present to one player at a time. Such NPCs would be best hosted by a player's own machine, incurring minimal communication latency and overhead, which is unlikely to be the case in task sharing. So, in practice it could be appropriate to combine different NPC host allocation mechanisms together, so as to attain better flexibility and efficiency.

2.3.4 NPC host allocation discussion

Currently, virtual-distance-based approaches are more widely used for NPC allocation because they minimise the communication latency for NPC hosts. However, heterogeneous task-sharing mechanisms also have their merits, as they optimise overall communication latency when an NPC interacts with multiple players. Potentially, P2P MMOGs could utilise both strategies flexibly in different game scenarios.

2.4 Game state persistency

An MMOG is also referred to as a persistent world (James and Walton, 2004), because its game world is always available to the users and game plots evolve even while some of the players are not playing their characters. In this case, an MMOG must store all players' profiles and inventories between login sessions. When a player comes back to the game, the player can retrieve its previous state information and continue to play.

2.4.1 Distributed storage infrastructures

Several distributed storage infrastructures have been proposed in the literature, which may facilitate game state persistency in P2P MMOGs. For example, the OceanStore project (Kubiatowicz et al., 2000) provides a global persistent data store utility designed to scale to billions of users. It supports consistent, highly available and durable storage atop an infrastructure composed of untrusted hosts.

Large-scale persistent storage services have also been built on structured P2P overlay networks, e.g., PAST (Rowstron and Druschel, 2001b) that uses Pastry (Rowstron and Druschel, 2001a). Compared with OceanStore, PAST is preferred by more P2P MMOG systems (Hampel et al., 2006; Assiotis and Tzanov, 2006; Bharambe, 2006; Iimura et al., 2004), which have already employed Scribe (Castro et al., 2002) for game event dissemination, as described in Section 2.2.1. It is fairly easy to combine support for Scribe and PAST at the same time.

2.4.2 Further considerations

Though the distributed storage infrastructures discussed earlier support many of the required facilities, their suitability to be applied to a P2P MMOG directly is still in doubt. One of the considerations is the efficiency for reading and writing data through such infrastructures. An MMOG may generate event updates frequently, but it takes a long time to modify the data using an overlay network. In this case, there will be a severe delay when a player wants to retrieve data that is still being modified (Iimura et al., 2004). One solution is to make super-peers provide a caching mechanism, so that real-time event updates take effect immediately, whereas a distributed storage infrastructure is only used as a slow medium for back-up purposes (Iimura et al., 2004).

Availability is another significant consideration. It has been argued that data should always be available for retrieval, since players would not be satisfied if their characters were unavailable because the persons that were storing their characters were not playing the game (Assiotis and Tzanov, 2006). This related work addresses the problem of separating storage needs into two categories: ephemeral and permanent data. The former can be stored using a distributed storage infrastructure, indexed by its geographical area in the virtual world. The latter should be stored locally, as it exists only when the player is in the game. Players can also back up their permanent data periodically using the distributed storage infrastructure in case their local data becomes corrupted.

2.4.3 State persistency discussion

Game state persistency is a major challenge for P2P MMOGs, as existing P2P storage infrastructures are designed to support file sharing, and seldom fulfil the performance and security requirements of an MMOG. Compared with the previous three design issues that have been heavily researched, the persistency area is still immature with many problems waiting to be investigated.

2.5 Cheating mitigation

The C/S architecture is easier to secure, because a server is able to validate every action request sent by a client before carrying it out. However, without the existence of such an authority, prevention of cheating becomes a challenging problem in P2P MMOGs.

In the literature, some propose to prevent cheating from happening by reinforcing game event ordering and state exposure protocols, whereas others just aim at detecting and remedying inconsistent simulation results after suspicious game sessions. In this section, the former are termed proactive approaches, and the latter reactive approaches.

2.5.1 Proactive approaches

Generally speaking, there are two main ways for an unscrupulous player to cheat during a game session – unfair knowledge acquisition and unwanted event suppression. A player may gain extra advantages unfairly by peeking at other players' current status, which is especially useful in strategy games. Advanced information exposure protocols such as Chambers et al. (2005) can be applied to reinforce the fair playout of a game in these cases. A more significant security weakness in a P2P MMOG is 'suppress-correct' cheating, which allows a player to gain an advantage by purposefully dropping update messages. Lockstep (Baughman et al., 2007) is the first event-ordering protocol to address fixed-delay and timestamp cheats. Lockstep orders events by rounds and increments a round only after every player has committed its move for that round. A drawback of Lockstep is that the total ordering of events suffers from the largest delay between any two players. NEO (GauthierDickey et al., 2004) improves this design by bounding the length of each round with a maximum latency, where voting is used to form a consensus on whether a given player has sent an update within a round. However, the trade-off is that a player that is slow to most nearby players will not be able to play in that area of the virtual world.

NEO requires event updates to be signed and encrypted before being sent to other players, so that a player cannot modify its own action after it has learned of others' actions. To achieve better performance, SEA (Gorawski and Stachurski, 2006) replaces NEO's encryption with a cryptographic hash function as the commitment method. Later on, a more efficient signature scheme EASES (Chan et al., 2008) further improves SEA by computing a message's digest before signing the message. Similarly, many other infrastructures have also been devised, e.g., FPS (Chen and Maheswaran, 2004) and Hack-Proof (Fung, 2006), which actively minimise the opportunity for cheating by requiring players always to react upon the same game state information for every frame.

2.5.2 Reactive approaches

Instead of applying sophisticated information exposure and event-ordering protocols, reactive approaches just aim at detecting unfair game playouts afterwards and rolling them back. For example, Log Auditing (LA) (Kabus et al., 2005) partitions a game world into multiple regions, and a super-peer called a 'Region Controller' (RC) is selected in each region. In LA, each player sends their commands as a signed sequence of packets, and the RC responds with signed game state updates. In this process, significant game events are logged by both the player and the RC. When given the same initial game state and player commands, the correct output can be reproduced by rerunning the log on a trusted machine, so that cheating can be detected. Likewise, Izaiku et al. (2006) introduce multiple 'monitor nodes' in each region, which calculate the latest game state from the previous game state, respectively, according to the game events that happened during the current timeslot. Hash values of game states are compared periodically, so as to identify potential cheats.

Both approaches mentioned earlier can be classified as referee-based mechanisms, in which the selection of non-colluding referees from untrusted peers is critical. Two secure referee selection algorithms have been proposed in Webb et al. (2008), which emphasises not only the fairness issue, but also the communication latency among referees and players.

DaCAP (Liu and Lo, 2008) and FreeMMG (Cecin et al., 2004) rely on mutual monitoring among all players concerned, rather than a limited number of referees. They organise players into ‘legal groups’ according to their locality in the game world. All members of a legal group have to compute and record all actions and status of the other members in the same group. Once cheating behaviour is detected, the cheating player is reported to a ‘check server’ with related evidence. To make collaboration to falsify player data hard, DaCAP randomly chooses a number of players from other areas of the game world to join the group, while FreeMMG allocates each group a server simulated player that can always be trusted.

Last but not least, a novel behavioural monitoring mechanism has been proposed in Laurens et al. (2007). This approach differs from any other methods in that it does not rely on knowledge about specific vulnerabilities and their method of exploitation to protect the system. Instead, it relies on the real-time monitoring of players’ movements and behaviours in the game world for indications of cheating play. This concept is based on the hypothesis that players engaged in cheating will exhibit characteristic behaviour, which is clearly distinguishable from normal play.

2.5.3 Cheating mitigation discussion

It is widely accepted that P2P MMOGs are more difficult to secure than conventional C/S architectures. Moreover, security issues present themselves at all stages in the design and implementation of P2P MMOGs. Hence, it is reassuring to see the intensive research into proactive and reactive cheating mitigation mechanisms starting to bear fruit.

2.6 Incentive mechanisms

P2P applications are by nature voluntary resource-sharing systems, in which there is often a tension between individual concerns and collective welfare. As the benefits of these systems are rooted in cooperation, they are inherently vulnerable to non-cooperative behaviour, and it is necessary for such systems to be designed so that participants are induced to cooperate. The mechanisms that are embedded in the system for this purpose are called Incentive Mechanisms (Zghaibeh and Anagnostakis, 2007). A P2P MMOG requires an incentive mechanism to convince its participants to contribute their resources. For example, network bandwidth is needed in game event dissemination, storage capacity is needed in game state persistency, and CPU cycles are needed in IM, NPC hosting and cheating mitigation, as discussed in previous sections.

2.6.1 Accounting mechanisms

An accounting mechanism maintains the viability of a P2P MMOG by quantifying the amount of computing resource a player has contributed to the system. On the one hand, it keeps a record of a player’s historical contribution, and on the other hand it entitles the player to consume roughly equivalent resources from other players. In this way,

selfish players can be identified and discouraged, and a sufficient level of reciprocity can be ensured to make use of a P2P MMOG beneficial.

DCRC (Gupta et al., 2003) is a fully distributed accounting system that applies to general P2P applications. The key idea in DCRC is a Debit/Credit platform using a virtual currency. By tracking a user's activities in a P2P system, DCRC bills the user according to the amount of resources that the user has consumed (i.e., Debit), and rewards the user according to the time and quality of a service that the user has offered to others (i.e., Credit). A user that stays in credit for long can be further encouraged in many different ways, e.g., service quality differentiation and application-specific privileges.

2.6.2 Reputation mechanisms

Merely quantifying a peer's contribution to a P2P system is sometimes inadequate in discouraging certain disadvantageous behaviours. For example, a player may have worked as a super-peer in a region for a long time and have contributed a lot of resources to the application. However, the player may also disconnect from the system abruptly when it decides to leave, and thus put the system into an inconsistent state, which takes much time and inconvenience to recover from. In this case, a reputation mechanism becomes valuable for qualifying a peer's dependability, honesty and overall manner towards P2P collaborations.

Many distributed reputation management systems can be used in a P2P MMOG. EigenTrust (Kamvar et al., 2003) and REPS (Huang et al., 2008) are representatives of mutual-rating-based approaches. In these systems, after each interaction peers produce either positive or negative feedbacks for each other, and keep the feedbacks in their own storage. Reputation query algorithms are provided for a peer to aggregate such local trust values from its direct friends, friends of friends, or arbitrary numbers of unacquainted peers, so that the peer can estimate approximately the trustworthiness of any other stranger peer.

In contrast, approaches like proactive (Swamynathan et al., 2008) and local (Liu et al., 2007) reputation do not depend on ratings from third parties. Instead, they provide various means for a peer to evaluate the trustworthiness of a target peer directly, hence they are inherently immune to bad-mouthing or collusion attacks. These approaches focus on addressing the challenge of anonymous reputation requests inside application traffic, because once an unscrupulous peer determines that the purpose of a request is to measure its reliability, it will be tempted to process these requests to boost its reputation.

2.6.3 Incentive discussion

The success of a P2P MMOG relies on an effective incentive mechanism that facilitates the collection of resources, and a reputation mechanism that minimises antisocial behaviours. Such incentive mechanism is a key design issue that is often unjustly ignored in the literature. As we shall see in Section 3, some P2P MMOG infrastructures may need to improve their incentive mechanisms to make them more practical.

3 Comparison of P2P MMOG architectures

In this section, six representative P2P MMOG architectures are selected and compared. Table 1 summarises their features and illustrates how they address the essential issues discussed previously. However, cheating mitigation is not included in this table, because it is a relatively separate issue, and all the architectures are potentially compatible with existing cheating mitigation techniques.

Table 1 Comparison of representative P2P MMOG architectures

<i>P2P MMOG architectures</i>	<i>Interest Management</i>	<i>Event dissemination</i>	<i>NPC host allocation</i>	<i>Game state persistency</i>	<i>Incentive mechanism</i>	<i>Overall evaluation</i>
P2P support '04	Region-based	ALM	Region-based	None	None	Simple
Distributed '04	Region-based	Unicast	Distance-based	Distributed	None	Moderate
OPeN '05	AOI-based	Unicast	None	Centralised	None	Moderate
P2P arch '06	Region-based	ALM	None	PAST	None	Simple
VAST '07	Voronoi	Unicast	Distance-based	Centralised	REPS	Complete
Mediator '07	Hybrid	Unicast	Task sharing	Distributed	DCRC	Complete

P2P support '04 (Lu et al., 2004)

This early architecture partitions a virtual game world into multiple regions, and IM is carried out by region-based publish/subscribe. Each region is associated with a publish channel, to which all region participants subscribe, and gaming events are delivered using Scribe. A super-peer called a 'region coordinator' is selected in each region, which hosts all the NPC objects in that region. A prototype application, 'SimMud', is implemented.

While representing a number of good design decisions, this architecture is evaluated as *simple* for the following reasons. The IM scheme is coarse-grained and its event dissemination relies on a general-purpose ALM middleware, which may induce high communication latency. Also, its NPC host allocation mechanism is intuitive, and neither specific game state persistency nor incentive mechanisms are provided.

Distributed '04 (Dickey et al., 2004)

This architecture adopts a region-based IM scheme, where the size of each region is quite small. In this case, the number of players in a region is limited, so players in the same region can communicate with each other using unicast. However, it is more likely that a gaming event that takes place in one region may also affect the players in neighbouring regions. Therefore, a super-peer is selected in each region to propagate local gaming events to neighbouring super-peers when necessary. Furthermore, the architecture suggests each player storing its own permanent data, while public ephemeral data is stored by a Distributed Hash Table (Stoica et al., 2001).

The architecture is evaluated as *moderate*, because its IM, event dissemination and state persistency mechanisms are well designed. However, it assumes that there are always adequate peers donating computing resources, and thus when an NPC becomes active, a random capable peer is selected to host that NPC. Such NPC host allocation and incentive mechanisms are overly simple.

OPeN '05 (Douglas et al., 2005)

The OPeN architecture proposes a distributed spatial data index service, which is built on top of a structured P2P overlay network. With this service, players can register their current locations in a game world, and query about other entities in their AOIs. Nearby entities establish direct UDP packet flows with each other to exchange gaming events. Persistent game data is stored and managed by a centralised database server. A simple P2P MMOG is implemented for demonstration purposes.

The architecture is evaluated *moderate*, because its IM and event dissemination mechanisms seem to be adequate. However, the infrastructure still depends on game servers for state persistency, and NPC host allocation and incentive mechanisms are not supported.

P2P arch '06 (Hampel et al., 2006)

This architecture is purely Pastry based, as it uses Scribe for game event dissemination, and PAST for game state persistency. The architecture employs coarse-grained region-based IM, but does not provide details about NPC host allocation and incentive mechanisms.

The architecture is evaluated as *simple*, because it directly employs ALM and distributed storage middlewares built on top of Pastry. As these middleware components are designed for general P2P applications, their performances may not be adequate for a P2P MMOG. Neither the arguments for the suitability of the middleware nor the demonstrations of its effectiveness are provided.

VAST '07 (Hu et al., 2006)

A unique Voronoi-assisted IM mechanism is employed in the VAST project. The Voronoi diagram is also inherently suitable for virtual-distance-based NPC host allocation. Furthermore, the architecture provides a native incentive mechanism as well. A prototype application 'ASCEND' is implemented.

Though currently the architecture still requires a game server for peer bootstrapping, load-balancing, fault-tolerance and state persistency purposes, it is evaluated as *complete*, because its application of the Voronoi technology is quite remarkable, which offers a consistent way of fine-grained IM, efficient game event dissemination and convenient NPC host allocation.

Mediator '07 (Fan et al., 2007)

The Mediator framework is the authors' ongoing research. The framework employs a hybrid IM scheme like MOPAR (Yu and Vuong, 2005) and disseminates gaming events through unicast communication. A novel task-mapping mechanism, Deadline-Driven Auctions (DDA), is devised to support the sharing of real-time NPC tasks. DDA is inherently compatible with reactive cheating mitigation approaches, e.g., LA (Kabus et al., 2005), and also supports a DCRC-like (Gupta et al., 2003) incentive mechanism that motivates application participants to contribute their resources to the system. Furthermore, a membership-aware multicast mechanism (MAMBO) (Fan et al., 2008) is developed for maintaining game zone infrastructures efficiently. It is also convenient to support game state persistency using PAST (Rowstron and Druschel, 2001b), as both PAST and MAMBO use the same overlay network. Key components of this framework as well as a test-bed application are implemented.

The architecture is evaluated as *complete*, because it addresses all of the design issues surveyed by this paper. In particular, its NPC host allocation mechanism is good at minimising communication latency among NPC hosts and ordinary players.

4 Conclusion and future work

To adapt MMOGs from conventional C/S architectures to P2P architectures is a challenging and active research area. This paper articulates a comprehensive set of six key issues for the design of P2P MMOGs. Design alternatives for each issue are systematically compared, the relationships between design decisions are discussed, and areas for further research are identified (Section 2). Finally, representative P2P MMOG infrastructures are evaluated against the design issues (Section 3).

The authors are currently refining the Mediator framework (Section 3) and intend to evaluate it for a prototype MMOG in the near future.

References

- Ahmed, D.T. and Shirmohammadi, S. (2008) 'An expedite state dissemination mechanism for MMOGs', *Proc. I-Span*, IEEE, Sydney, Australia, pp.199–203.
- Anderson, T.W. (1984) *An Introduction to Multivariate Analysis*, 2nd ed., John Wiley & Sons, Hoboken, NJ, USA.
- Assiotis, M. and Tzanov, V. (2006) 'A distributed architecture for MMORPG', *Proc. NetGames*, ACM, Singapore, Article No. 4, pp.1–7.
- Baughman, N.E., Liberatore, M. and Levine, B.N. (2007) 'Cheat-proof payout for centralized and peer-to-peer gaming', *IEEE/ACM Transactions on Networking*, Vol. 15, No. 1, pp.1–13.
- Benford, S. and Fahlen, L.E. (1993) 'A spatial model of interaction in large virtual environments', *Proc. ECSCW*, IEEE, Milan, Italy, pp.107–123.
- Bharambe, A. (2006) 'Colyseus: a distributed architecture for online multiplayer games', *Proc. NSDI*, USENIX, San Jose, CA, USA, pp.3–6.
- Boulanger, J-S., Kienzle, J. and Verbrugge, C. (2006) 'Comparing interest management algorithms for massively multiplayer games', *Proc. NetGames*, ACM, Singapore, Article No. 6, pp.1–12.
- Buyukkaya, E. and Abdallah, M. (2008) 'Data management in Voronoi-based P2P gaming', *Proc. CCNC*, IEEE, Las Vegas, USA, pp.1050–1053.
- Castro, M., Druschel, P., Kermarrec, A. and Rowstron, A. (2002) 'SCRIBE: a large-scale and decentralized application-level multicast infrastructure', *IEEE Journal on Selected Areas in Communications*, Vol. 20, No. 8, pp.1489–1499.
- Cecin, F.R., Real, R., de Oliveira Jannone, R., Geyer, C.F.R., Martins, M.G. and Barbosa, J.L.V. (2004) 'FreeMMG: a scalable and cheat-resistant distribution model for internet games', *Proc. DS-RT*, IEEE, Budapest, Hungary, pp.83–90.
- Chambers, C., Chang Feng, W., Chi Feng, W. and Saha, D. (2005) 'Mitigating information exposure to cheaters in real-time strategy games', *Proc. NOSSDAV*, ACM, Washington, USA, pp.7–12.
- Chan, M-C., Hu, S-Y. and Jiang, J-R. (2008) 'An efficient and secure event signature (EASES) protocol for peer-to-peer massively multiplayer online games', *Computer Networks*, Vol. 52, No. 9, pp.1838–1845.
- Chen, B.D. and Maheswaran, M. (2004) 'A fair synchronization protocol with cheat proofing for decentralized online multiplayer games', *Proc. NCA*, IEEE, Cambridge, MA, USA, pp.372–375.

- Chen, F. and Kalogeraki, V. (2005) 'Adaptive real-time update dissemination in distributed virtual simulation environments', *Proc. ISORC*, IEEE, Seattle, WA, USA, pp.233–236.
- Deering, S.E. and Cheriton, D.R. (1990) 'Multicast routing in datagram internet works and extended LANs', *ACM Transactions on Computer Systems*, Vol. 8, No. 2, pp.85–110.
- Dickey, C.G., Zappala, D. and Lo, V. (2004) 'A fully distributed architecture for massively multiplayer online games', *Proc. NetGames*, ACM, Portland, Oregon, USA, p.171.
- Diot, C., Levine, B.N., Lyles, B., Kassem, H. and Balensiefen, D. (2000) 'Deployment issues for the IP multicast service and architecture', *IEEE Network*, Vol. 14, pp.78–88.
- Douglas, S., Tanin, E. and Harwood, A. (2005) 'Enabling massively multi-player online gaming applications on a P2P architecture', *Proc. ICIA*, IEEE, Hong Kong and Macau, China, pp.7–12.
- El-Sayed, A., Roca, V. and Mathy, L. (2003) 'A survey of proposals for an alternative group communication service', *IEEE Network*, Vol. 17, No. 1, pp.46–51.
- Fan, L., Taylor, H. and Trinder, P. (2007) 'Mediator: a design framework for P2P MMOGs', *Proc. NetGames*, ACM, Melbourne Australia, pp.43–48.
- Fan, L., Trinder, P. and Taylor, H. (2008) *MAMBO: Membership-Aware Multicast with Business Optimisation*, Short Papers for DEBS'08, Rome, Italy, pp.1–2, Available online at www.debs.org
- Fiedler, S., Wallner, M. and Weber, M. (2002) 'A communication architecture for massive multiplayer games', *Proc. NetGames*, ACM, Braunschweig, Germany, pp.14–22.
- Fung, Y.S. (2006) 'Hack-proof synchronization protocol for multi-player online games', *Proc. NetGames*, ACM, Singapore, Article No. 47, pp.1–6.
- GauthierDickey, C., Lo, V. and Zappala, D. (2005) 'Using N-trees for scalable event ordering in peer-to-peer games', *Proc. NOSSDAV*, Washington, USA, ACM, pp.87–92.
- GauthierDickey, C., Zappala, D., Lo, V. and Marr, J. (2004) 'Low latency and cheat-proof event ordering for peer-to-peer games', *Proc. NOSSDAV*, ACM, Cork, Ireland, pp.134–139.
- Gorawski, M. and Stachurski, K. (2006) 'A secure event agreement (SEA) protocol for peer-to-peer games', *Proc. ARES*, IEEE, Vienna, Austria, pp.34–41.
- Gupta, M., Judge, P. and Ammar, M. (2003) 'A reputation system for peer-to-peer networks', *Proc. NOSSDAV*, ACM, Monterey, California, USA, pp.144–152.
- Hampel, T., Bopp, T. and Hinn, R. (2006) 'A peer-to-peer architecture for massive multiplayer online games', *Proc. NetGames*, ACM, Singapore, Article No. 48, pp.1–4.
- Hu, S-Y. and Liao, G-M. (2004) 'Scalable peer-to-peer networked virtual environment', *Proc. NetGames*, ACM, Portland, Oregon, USA, pp.129–133.
- Hu, S-Y., Chang, S-C. and Jiang, J-R. (2008) 'Voronoi state management for peer-to-peer massively multiplayer online games', *Proc. CCNC*, IEEE, Las Vegas, Nevada, USA, pp.1134–1138.
- Hu, S-Y., Chen, J-F. and Chen, T-H. (2006) 'VON: a scalable peer-to-peer network for virtual environments', *IEEE Network*, Vol. 20, No. 4, pp.22–31.
- Huang, G-Y., Hu, S-Y. and Jiang, J-R. (2008) 'Scalable reputation management for P2P MMOGs', *Proc. MMVE*, IEEE, Reno, Nevada, USA, pp.1–5.
- Iimura, T., Hazeyama, H. and Kadobayashi, Y. (2004) 'Zoned federation of game servers – a peer-to-peer approach', *Proc. NetGames*, ACM, Portland, Oregon, USA, pp.116–120.
- Izaiku, T., Yamamoto, S., Murata, Y., Shibata, N., Yasumoto, K. and Ito, M. (2006) 'Cheat detection for MMORPG on P2P environments', *Proc. NetGames*, ACM, Singapore, Article No. 46, pp.1–4.
- James, D. and Walton, G. (2004) *2004 Persistent Worlds Whitepaper*, Technical Report, IGDA Online Games SIG, www.igda.org
- Kabus, P., Terpstra, W.W., Cilia, M. and Buchmann, A.P. (2005) 'Addressing cheating in distributed MMOGs', *Proc. NetGames*, ACM, Hawthorne, NY, USA, pp.1–6.

- Kamvar, S.D., Schlosser, M.T. and Garcia-Molina, H. (2003) 'The eigentrust algorithm for reputation management in P2P networks', *Proc. WWW*, ACM, Budapest, Hungary, pp.640–651.
- Kubiatowicz, J., Bindel, D., Chen, Y., Czerwinski, S., Eaton, P., Geels, D., Gummadi, R., Rhea, S., Weatherspoon, H., Weimer, W., Wells, C. and Zhao, B. (2000) 'OceanStore: an architecture for global-scale persistent storage', *Proc. ASPLOS*, ACM, Cambridge, MA, USA, pp.190–201.
- Laurens, P., Paige, R.F., Brooke, P.J. and Chivers, H. (2007) 'A novel approach to the detection of cheating in multiplayer online games', *Proc. ICECCS*, IEEE, Auckland, New Zealand, pp.97–106.
- Léty, E., Turletti, T. and Baccelli, F. (2004) 'SCORE: a scalable communication protocol for large-scale virtual environments', *IEEE/ACM Transactions on Networking*, Vol. 12, No. 2, pp.247–260.
- Liu, H-I. and Lo, Y-T. (2008) 'DaCAP – a distributed anti-cheating peer to peer architecture for massive multiplayer on-line role playing game', *Proc. CCGRID*, IEEE, Lyon, France, pp.584–589.
- Liu, S., Li, J. and Wang, X. (2007) 'Local reputation for P2P MMOG design', *Proc. PDCAT*, IEEE, Adelaide, Australia, pp.523–528.
- Lu, H., Knutsson, B., Xu, W. and Hopkins, B. (2004) 'Peer-to-peer support for massively multiplayer games', *Proceeding of INFOCOM*, IEEE, Hong Kong, China, pp.7–11.
- Macedonia Michael, R., Zyda Michael, J., Pratt David, R., Barham Paul, T. and Steven, Z. (1994) 'NPSNET: a network software architecture for large scale virtual environments', *Presence*, Vol. 3, No. 4, pp.265–287.
- Matsumoto, N., Kawahara, Y., Morikawa, H. and Aoyama, T. (2005) 'A scalable and low delay communication scheme for networked virtual environments', *Proc. GLOBECOM*, IEEE, St. Louis, Missouri, USA, pp.529–535.
- Morgan, G., Lu, F. and Storey, K. (2005) 'Interest management middleware for networked games', *Proc. i3D*, ACM, Washington DC, USA, pp.57–64.
- Mulligan, J. and Patrovsky, B. (2003) *Developing Online Games – An Insiders Guide*, New Riders Publishing, Indianapolis, IN, USA, ISBN: 1592730000.
- Ratnasamy, S., Francis, P., Handly, M., Karp, R. and Schenker, S. (2001a) 'A scalable content addressable network', *Proc. SIGCOMM*, ACM, San Diego, USA, pp.161–172.
- Ratnasamy, S., Handley, M., Karp, R. and Shenker, S. (2001b) 'Application level multicast using content addressable networks', *Proc. NGC*, London, UK, pp.14–29.
- Rhalibi, A.E. and Merabti, M. (2006) 'Interest management and scalability issues in P2P MMOG', *Proc. CCNC*, IEEE, Las Vegas, NV, USA, pp.1188–1192.
- Rooney, S., Bauer, D. and Deydier, R. (2005) 'A federated peer-to-peer network game architecture', *Communications*, Vol. 42, No. 5, pp.114–122.
- Rowstron, A. and Druschel, P. (2001a) 'Pastry: scalable, decentralized object location and routing for large scale peer-to-peer systems', *Proc. Middleware*, ACM, Heidelberg, Germany, pp.329–350.
- Rowstron, A. and Druschel, P. (2001b) 'Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility', *Proc. SOSP*, Chateau Lake Louise, Banff, Canada, pp.188–201.
- Schiele, G., Suselbeck, R., Wacker, A., Hahner, J., Becker, C. and Weis, T. (2007) 'Requirements of peer-to-peer-based massively multiplayer online gaming', *Proc. CC-GRID*, IEEE, Rio de Janeiro, Brazil, pp.773–782.
- Schmieg, A., Stieler, M., Jeckel, S., Kabus, P., Kemme, B. and Buchmann, A. (2008) 'Sense – maintaining a dynamic localized peer-to-peer structure for position based multicast in games', *Proc. P2P*, IEEE, Aachen, Germany, pp.247–256.

- Simon, R., Klaus, W., Marc, F., Heiko, N., Leo, P. and Georg, C. (2007) 'Peer-to-peer-based infrastructure support for massively multiplayer online games', *Proc. CCNC*, Las Vegas, NV, USA, IEEE, pp.763–767.
- Stoica, I., Morris, R., Karger, D. and Kaashoek, F. (2001) 'Chord: a scalable peer-to-peer lookup service for internet applications', *Proc. NetGames*, ACM, San Deigo, CA, USA, pp.149–160.
- Storey, K., Lu, F. and Morgan, G. (2004) 'Determining collisions between moving spheres for distributed virtual environments', *Proc. CGI*, IEEE, Crete, Greece, pp.140–147.
- Swamynathan, G., Zhao, B.Y. and Almeroth, K.C. (2008) 'Exploring the feasibility of proactive reputations', *Concurrency and Computation: Practice and Experience*, Vol. 20, No. 2, pp.155–166.
- Webb, S.D., Soh, S. and Trahan, J. (2008) 'Secure referee selection for fair and responsive peer-to-peer gaming', *Proc. PADS*, IEEE, Rome, Italy, pp.63–71.
- Woodcock, B.S. (2005) *An Analysis of MMOG Subscription Growth*, Technical Report, www.mmogchart.com
- Xiang-bin, S., Yue, W., Qiang, L., Ling, D. and Fang, L. (2008) 'An interest management mechanism based on N-tree', *Proc. SNPD*, IEEE, Phuket, Thailand, pp.917–922.
- Yamamoto, S., Murata, Y., Yasumoto, K. and Ito, M. (2005) 'A distributed event delivery method with load balancing for MMORPG', *Proc. NetGames*, ACM, Hawthorne, NY, USA, pp.1–8.
- Yang, B.B. and Garcia-Molina, H. (2003) 'Designing a super-peer network', *Proc. ICDE*, IEEE, Bangalore, India, pp.49–60.
- Yonekura, T., Kawano, Y. and Hanawa, D. (2004) 'Peer-to-peer networked field-type virtual environment by using atoz', *Proc. CW*, IEEE, Tokyo, Japan, pp.241–248.
- Yu, A.P. and Vuong, S.T. (2005) 'MOPAR: a mobile peer-to-peer overlay architecture for interest management of massively multiplayer online games', *Proc. NOSSDAV*, ACM, Skamania, Washington, USA, pp.99–104.
- Zghaibeh, M. and Anagnostakis, K.G. (2007) 'On the impact of P2P incentive mechanisms on user behavior', *Proc. NetEcon+IBC*, ACM, San Diego, California, USA, pp.1–23.
- Zhang, R. and Hu, Y.C. (2003) 'Borg: a hybrid protocol for scalable application-level multicast in peer-to-peer networks', *Proc. NOSSDAV*, ACM, Monterey, California, USA, pp.172–179.
- Zhao, B., Kubiawicz, J. and Joseph, A. (2001) *Tapestry: An Infrastructure for Fault-tolerant Wide Area Location and Routing*, Technical Report, UC Berkeley.
- Zhuang, S.Q., Zhao, B.Y., Joseph, A.D., Katz, R.H. and Kubiawicz, J.D. (2001) 'Bayeux: an architecture for scalable and fault-tolerant wide-area data dissemination', *Proc. NOSSDAV*, ACM, New York, USA, pp.11–20.