
Deadline-Driven Auctions for NPC host allocation in P2P MMOGs

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: We present the design, implementation and evaluation of Deadline-Driven Auctions (DDAs), a novel task-mapping infrastructure for heterogeneous distributed environments. DDA is primarily designed for hosting Non-Player Characters (NPCs) in P2P Massively Multiplayer Online Games (MMOGs). Experimental and analytical results demonstrate that DDA provides four significant advantages. It is self-organising: the infrastructure is automatically managed. It efficiently allocates computing resources for large numbers (1000s) of real-time NPC tasks. It supports gaming interactivity by minimising communication latency between NPC hosts. Finally, it supports flexible matchmaking policies, and a friendly incentive policy establishes a cooperative economic model to motivate participants to contribute resources.

Keywords: NPC host allocation; P2P; MMOGs; massively multiplayer online games; real time; task mapping; heterogeneous environments; communication latency; matchmaking; incentive; simulation.

Reference to this paper should be made as follows: Fan, L., Trinder, P. and Taylor, H. (2010) 'Deadline-Driven Auctions for NPC host allocation in P2P MMOGs', *Int. J. Advanced Media and Communication*, Vol. 4, No. 2, pp.140–153.

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

A Virtual Environment (VE) is a computer-simulated environment for its users to inhabit and interact via avatars. Many VEs are MMOGs. Conventionally, VEs and MMOGs have been implemented using Client/Server (C/S) architectures, because they are relatively easy to implement and secure (Mulligan and Patrovsky, 2003). However, they also exhibit various drawbacks in reliability and cost (Fan et al., 2007), hence Peer-to-Peer (P2P) MMOGs are becoming increasingly attractive as an alternative (Hu et al., 2008).

To adapt MMOGs to P2P architectures, many key issues have to be addressed. An MMOG is different from a typical VE in that it features considerable numbers of NPCs. These NPCs need to be hosted by peers to be run (Bharambe, 2006; Yonekura et al., 2004; Hu et al., 2008). An NPC is an AI-controlled virtual actor, which drives storylines in a game, or combats Player Characters (PCs) as a monster. MMOGs have to supply their game worlds with large numbers of NPCs as required by game scenarios. Traditionally, NPCs are hosted by game servers, consuming significant processing power and network bandwidth. Therefore, one of the prerequisites for realising a P2P MMOG is to provide a mechanism that hosts such NPCs using resources available on game participant machines.

Conceptually, an NPC object is an indivisible, computational and interactive task, because:

- 1 It can only be efficiently hosted by a single computer.
- 2 It consumes processing power, as an NPC is associated with an AI programme that determines the NPC's actions. For example, a monster NPC may determine its move by examining the position of nearby players.
- 3 It needs to interact with players. For example, a monster NPC may engage in combat with a set of players, and exchange real-time gaming events with them.

A mechanism that schedules NPC tasks on game participant machines is referred to as NPC host allocation. The responsibilities of such a mechanism include the discovery of potential resource providers, the selection of suitable NPC hosts, and the migration of related AI programme and state information for an NPC. Among them, the most important one is the selection of a suitable host. First, the selected host must have adequate computing resource, such as CPU cycles, memory and network bandwidth. Second, it should also offer a low communication latency to guarantee the gaming experiences of other players that interact with the NPC. This determination must be made in due course, because the fast pace of an MMOG requires an NPC to appear at a specific position and start working in the order of seconds.

The contribution of this paper is to present the design (Section 3), analysis (Section 4), extensions (Section 5), implementation using Pastry (Rowstron and Druschel, 2001) (Section 6), and evaluation (Section 7) of DDA. DDA supports real-time NPC host allocation in P2P MMOGs using a heterogeneous task-mapping mechanism. DDA's infrastructure is a self-organised super-peer network on top of a structured P2P overlay. This infrastructure is provided by our previous research on the Mediator framework (Fan et al., 2007).

Experimental results demonstrate that DDA provides the following advantages:

- 1 *Self-organisation*: DDA's infrastructure can be automatically assembled and managed. As the system evolves, NPC tasks will be distributed to suitable participants automatically (Section 7.1).
- 2 *Real-time resource allocation*: DDA efficiently processes large numbers of tasks within a few seconds, and is able to support a simulated P2P MMOG with the better part of 1000 players (Section 7.2).
- 3 *QoS desiderata*: DDA minimises the communication latency between an NPC host and ordinary players, so as to improve interactivity for a P2P MMOG (Section 7.2).
- 4 *Cooperative economic model*: DDA supports flexible resource selection policies, and it conveniently establishes a cooperative economic model that shares NPC tasks among competent resource providers fairly (Section 5.2).

2 Related work

Existing NPC host allocation schemes can be classified into static region-based approaches and dynamic virtual-distance-based approaches. The former (Lu et al., 2004; Limura 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 NPCs within the region. They have several drawbacks. Because only one super-peer is selected to take charge of a region, they might incur excessive computation and communication workloads on the super-peer. Their proposed super-peer selection criteria are also overly simple, as they do not take into consideration peers' actual resource availabilities. Furthermore, the approaches cannot guarantee to fulfil the QoS requirement for game interactivity.

In contrast, dynamic virtual-distance-based approaches distribute NPC objects to all game participants, where the key idea is to allocate an NPC to the machine of the player, whose avatar is closest to the NPC. 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 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. The game object manager of Colyseus allocates mutable objects, e.g., NPCs, doors and weapon items, to the closest players. Similarly, AtoZ (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 discussed in Hu et al. (2008) seems inherently suitable for virtual-distance-based NPC host allocation. A Voronoi diagram partitions a game world into multiple non-overlapping regions that contain exactly one player avatar in each region. In this case, it is natural for each player to host the NPC objects within its own Voronoi cell.

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

- 1 Some NPCs like shop owners may only need to be present to one player at a time. This might be enforced by environments that limit the NPC’s visibility and interactivity to one player at a time. Such NPCs would be best hosted by a player’s own machine. However, most NPCs are not like this. They can have a real-time effect on players on several hosts at a time. In this case, all the players need to communicate with the NPC host, and virtual-distance-based approaches cannot ensure that the latency for each player is equally low.
- 2 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.
- 3 Cheating may become easier for unscrupulous players who 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.

The DDA infrastructure adopts a task-mapping mechanism that is different from all the related work. DDA is good at load-balancing, because it takes into consideration each game participant’s actual resource availability and ensures that NPC tasks are always allocated to capable hosts. Also, 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. Furthermore, because a player is unlikely to host an NPC for itself, cheating is relatively harder in DDA. DDA also provides several additional advantages such as self-organisation, real-time resource allocation, QoS support and a cooperative economic model.

3 DDA design

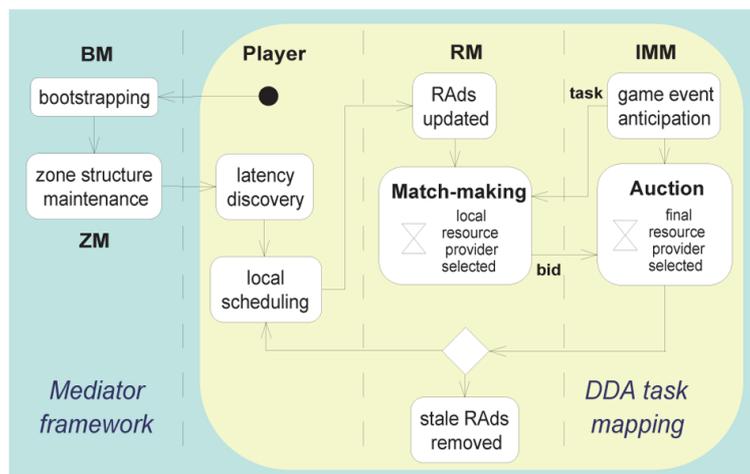
3.1 Overview

The system model for DDA involves two different parties: a work source and a set of resource providers. The work source is the virtual game world that constantly generates NPC tasks. The resource providers are game participants that have spare resource on their machines. The key idea in DDA is to bridge between resource requirements and their availability using distributed matchmaking via resource matchmakers.

Because a P2P system lacks centralised configuration, a self-organising mechanism is needed to set up DDA's working infrastructure. Previous work on the Mediator framework (Fan et al., 2007) supplies such an infrastructure. It automatically organises application participants into a hierarchical super-peer network using a structured P2P overlay (Rowstron and Druschel, 2001). Four super-peer roles have been defined in the framework, Boot Mediator (BM), Zone Mediator (ZM), Resource Mediator (RM) and Interest-Management Mediator (IMM).

The collaboration among the Mediator roles is illustrated by Figure 1. First, the entire game world is partitioned into multiple game zones, and a BM is selected in each zone for bootstrapping purposes. Second, the BM promotes a resource-rich peer from the existing zone members to be the ZM for that zone with responsibility for zone structure maintenance, e.g., the selection and monitoring of the IMM and multiple RMs. Third, an IMM is essentially a gaming event anticipator. By running a hybrid interest-management scheme, e.g., MOPAR (Yu and Vuong, 2005), the IMM respawns NPC objects and updates other zone members with forthcoming gaming events. Finally, RMs serve as matchmakers that facilitate real-time resource discovery. For load-balancing and performance reasons, multiple RMs are selected in each game zone and work in parallel. Once an NPC is about to be respawned, the IMM notifies the RMs in the game zone, and each RM replies with a locally optimised resource provider as a 'bid'. When the IMM has received their bids, or the deadline for the spawning task is sufficiently near, the IMM closes the 'auction' and determines which resource provider hosts the NPC.

Figure 1 The mediator framework and DDA (see online version for colours)



3.2 Local scheduling

Local scheduling refers to the activities carried out by a peer to manage its local computing resources, and to contribute spare resource to the system. Typical local scheduling activities include finding out communication latencies with other zone members, and updating resource availabilities at an RM.

DDA resorts to incentive mechanisms to convince participants to contribute their resources. DDA uses DCRC (Gupta et al., 2003) to quantify peers' contributions, so free-riders can be identified and discouraged. DDA charges peers in term of credits, according to the time that they play the game. Ideally, a peer that actively contributes usable resources to the system should be able to earn enough credits to pay for its playing time. However, if resource-rich peers eagerly volunteer for tasks, less-competitive peers will be starved of opportunities for earning credits. As a result, some virtuous peers may become poorer and poorer, and finally be regarded as free-riders. To cope with this problem, DDA adopts a friendly matchmaking policy (Section 5.2) that fairly shares credit earning opportunities among all competent resource providers.

Furthermore, only to quantify available resources is not sufficient. To fulfil the QoS requirement for game interactivity, a peer also needs to qualify its resources by finding out its communication latencies with other peers. For zone structure maintenance purposes, the ZM for each game zone publishes heartbeat messages periodically. Such messages carry a list of existing zone members, so that a peer can check whether new zone members have shown up. If so, the peer sends out a set of `Ping` messages to them, and they reply with `Pong` messages, so that communication latencies among them can be measured.

DDA adopts a matchmaking approach similar to Condor (Litzkow et al., 1988), and represents both the resource offers and the NPC tasks using ClassAds (Raman et al., 2000). A peer periodically reports its local resource quantity and quality to an RM using a Resource Ad (RAd), and an IMM notifies related RMs about an NPC task using a Job Ad (JAd). Detailed structures for a RAd and a JAd can be found in Fan et al. (2007).

Currently, cheating mitigation is not a primary focus of DDA's design, but obviously it is possible for unscrupulous peers to cheat on local scheduling. For example, to earn more credits, a peer may lie about the spare computing resource it can provide. As a result, the peer will be allocated tasks that are beyond its capability, which consequently damages other peers' gaming experiences. One avenue of future work is to enhance DDA's security by discouraging disadvantageous peer behaviour using a reputation system.

3.3 Zone-level scheduling

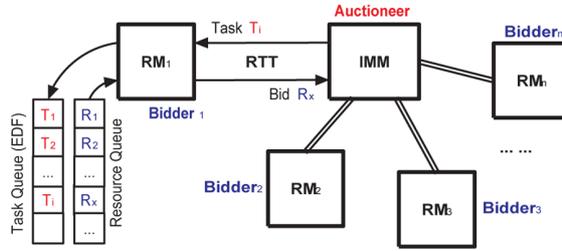
Figure 2 depicts DDA's zone-level scheduling, where an IMM works as an auctioneer, and multiple RMs serve as bidders. Each RM maintains two queues: a Resource Queue (RQ) for storing RAds from resource providers, and a Task Queue (TQ) for buffering JAds given by the IMM. A TQ suits being made a priority queue, where tasks are ordered on an Earliest Deadline First (EDF). The processing cycle of the RM involves selecting the task with the shortest deadline from the TQ, scanning the RQ to find the most adequate resource provider, and returning it to the IMM as a bid. From the IMM's perspective, every task T_i must be completed within its deadline D_i . This requirement is captured in equation (1), in which $C(T_i)$ means the completion time for task i .

$$\forall_{i \in 1 \dots n} C(T_i) < D_i. \quad (1)$$

According to Figure 2, $C(T_i)$ comprises two periods of time: the Round-Trip Time (RTT) between the IMM and an RM, and the time that an RM takes to discover a locally optimised resource provider. To simplify the analysis, it is assumed that by carefully selecting super-peers from candidates, each RM has an equally short RTT with the IMM. Furthermore, the RQs maintained by different RMs are equally long, each containing l RAds. If it takes t milliseconds for an RM to match one RAd to the JAd, it will take $l * t$ milliseconds in total for the RM to complete matching of all the RAds to decide which one is the best. Some preordering of RAds might reduce the search time a little, but would incur the overhead of computing a preordering for every new RAd received. So, for task T_i :

$$\forall_{i \in 1 \dots n} C(T_{i+1}) = RTT + i * l * t. \quad (2)$$

Figure 2 DDA zone-level scheduling (see online version for colours)



4 Design analysis

The respawning algorithm used in an MMOG is important to estimate the reasonable scale of i in equation (2). Currently, well-known commercial MMOGs (WoW, 2001; WoL, 2004) employ a timer-based approach that respawns NPCs to keep the number of NPCs and PCs to a stable ratio. Suppose that the number of players in a game zone is P , the NPC to PC ratio is R , and on average an NPC's lifetime is TTL seconds, $P * R / TTL$ NPCs should be respawned every second. Accordingly, if a timer is set to respawn NPCs every Int th second (i.e., the respawning interval), each time $P * R * Int / TTL$ NPCs are respawned. In fact, the respawning interval determines the deadline for a set of NPC tasks. If the last task can be processed before its deadline, previous tasks can meet their deadlines as well. In this case, equation (1) is changed to:

$$C(T_{last}) = RTT + \frac{P * R * Int * l * t}{TTL} < Int. \quad (3)$$

In the Mediator framework, a player may also trigger other events that require resource providers to be located. Assuming that each player triggers r such events in every second, there would be $r * P$ events in total. It is more likely that the player-triggered events have shorter deadlines than regular respawning events, and correspondingly, tasks for the former have higher priorities in an RM's TQ. The worst case is that in every respawning interval, the RM needs to complete all player-triggered tasks before handling the last respawning task, so equation (3) is changed to:

$$RTT + \left(\frac{P * R}{TTL} + r * P \right) * Int * l * t < Int. \quad (4)$$

Table 1 lists all the key parameters of the DDA model, together with nominal values typical of commercial MMOGs (WoW, 2001; WoL, 2004). Applying the nominal values to equation (4) creates equation (5), and let us draw the following inferences:

$$(600 - P) * Int > 300. \quad (5)$$

- based on these assumptions, DDA can support up to 600 peers in each game zone, which corresponds to a zone size of around 500 peers in many typical MMOGs
- when implementing a P2P MMOG with a zone size of 500 peers, the minimal respawning interval is 3 s
- on average, a peer may obtain a task in every ten respawning intervals, so the credits rewarded for running a task should cover the payment charged for corresponding playing time.

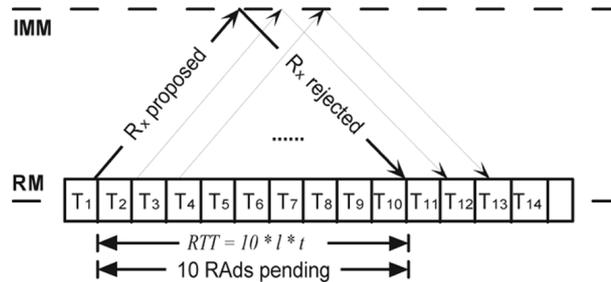
Table 1 DDA variables and nominal values

<i>Variable</i>	<i>Meaning</i>	<i>Nominal value</i>
<i>P</i>	Zone population	
<i>R</i>	NPC : PC ratio	5 : 1
<i>TTL</i>	NPC life time expectation	300 s
<i>R</i>	Event triggering rate	1/60 per second
<i>Int</i>	Respawning interval	
<i>RTT</i>	Round Trip Time	0.5 s
<i>L</i>	RM resource queue length	50 RAds
<i>T</i>	Matchmaking time	1 ms per RAd

5 DDA extensions

5.1 Multilevel Feedback Delay Queue

Though theoretically DDA supports up to 600 peers in each game zone, the actual population might be smaller, especially when dynamic zoning is supported. Suppose that there are only 200 peers in a zone, and a respawning interval of 6 s is used. Hence, at the beginning of every interval, the IMM sends out 20 tasks to the RMs, specifying the deadline as 6 s. For T_1 , RM Alice proposes RAd R_x to the IMM as a bid, then carries on processing T_2 . Unfortunately, at the IMM end, R_y proposed by RM Bob is better than R_x , so R_x is rejected and returned to Alice. In spite of the time for IMM's decision-making, it will take at least RTT time before Alice finds out that her bid has been rejected. Here, a potential problem is that in expression 2, $l * t$ is too small compared with RTT . Figure 3 demonstrates this 'resource tie-up' problem, in which 10 RAds are tied up throughout every respawning interval. To address this problem, an RM may need to slow down its matchmaking, but still guarantees that all tasks meet their deadlines. This can be achieved with a Multilevel Feedback Delay Queue (MFDQ).

Figure 3 The resource tie-up problem

An MFDQ is similar to a multilevel feedback queue for Unix process scheduling, where the main difference is that in the MFDQ, each run queue is executed only when its delay has expired. If the inter-run queue delay is d , then an MFDQ comprises $n = \text{ceil}(\text{Int} = d)$ run queues in total. Each run queue represents a different urgency level. Queue `Instant` buffers the most urgent tasks. Other run queues are arranged in a circle, with a `Head` pointer indicating the queue to be executed next. An MFDQ mainly supports two operations:

1) *offer(Task t)* — the enqueue operation:

```
offer(Task t){
    t.enqueue_deadline = t.original_deadline-RTT-l*t;
    t.urgency_level = floor(t.enqueue_deadline/d);
    if(t.urgency_level == 0)
        Instant.add(t);
    else if(t.urgency_level >= (n-1))
        Tail.add(t);
    else
        getQueue((t.urgency_level+Head.index)%n).add(t);
}
```

The algorithm recalculates a task's enqueue deadline by subtracting RTT and $l * t$ from its original deadline. Then, it evaluates the task's urgency level. If the task is not able to bear one inter-queue delay, the task is added to queue `Instant`. Otherwise, the task is added to a run queue.

2) *poll()* — the dequeue operation:

```
poll(){
    int to_move = ceil(total_number_of_tasks/n);
    int done = 0;
    for(each Task t in Head){
        Instant.add(Head.remove());
        done++;
    }
    Head = getQueue((Head.index+1)%n);
    int current = Head.index;
    while(done < to_move){
        if(! getQueue(current).isEmpty()){
            Instant.add(getQueue(current).remove());
            done++;
        }
        else
            current = (current+1)%n;
    }
}
```

The algorithm first moves all tasks in queue `Head` to queue `Instant`. Because each run queue may have a different length, it moves the same number of tasks each time. If less tasks are moved from `Head` to `Instant`, more tasks in lower-level queues are moved as well. If a task's deadline can be satisfied in a common EDF queue, the MFDQ can satisfy its deadline as well.

5.2 Matchmaking policies

A matchmaking policy is the set of criteria that an RM uses for selecting locally optimised resource providers (by default, the term 'optimised' refers to minimised communication latency). In the Mediator framework, an IMM anticipates forthcoming gaming events in a game zone, as discussed in Section 3.1. So, an IMM is able to provide in a JAd a list of target players, which are able to interact with a newly respawned NPC. At matchmaking time, an RM takes out every available RAd from its RQ, matching it to the JobAd using two conditions:

- Does the resource provider have adequate computing resource for running the task?
- Is the mean communication latency between the resource provider and the target players the shortest?

Because this policy strictly selects the resource provider that provides minimum latency, it is called a 'strict incentive policy'. However, a problem is that less-competitive peers are more likely to be starved of opportunities for earning credits, and finally to be regarded as free-riders, as discussed in Section 3.2. To stop this problem from happening, a factor τ can be introduced to relax the selection criteria as follows:

- Does the resource provider have adequate computing resource for running the task?
- Is the mean communication latency within the range of *shortest* * τ ?
- Is the resource provider low on credit?

This policy favours the poorest peer providing a latency that is not greater than *shortest* * τ , and making it a 'friendly incentive policy'. Experimental results demonstrate that with $\tau=1.2$, DDA shares tasks among the peers more fairly than using the strict incentive policy of $\tau=1$ (Section 7.3). Actually, DDA allows flexible matchmaking policies to be adopted. For example, when a reputation mechanism is brought into effect in future work, it will allow policies that favour more dependable resource providers.

6 Implementation

Currently, a proof-of-concept prototype for DDA has been implemented using FreePastry 2.1, an open-source implementation of Pastry (Rowstron and Druschel, 2001). The RM's matchmaking mechanism has been implemented using ClassAds 2.2 (Litzkow et al., 1988). Furthermore, to evaluate the prototype, a test-bed application has been developed, which simulates a P2P MMOG using a two-dimensional game world and hundreds of virtual player avatars moving according to a random waypoint algorithm.

The test-bed employs the Direct discrete event simulator integrated in FreePastry. The simulator takes in a network topology model generated by GT-ITM, and simulates

internet-scale communication latencies among peers in a P2P network. Each peer is assigned some virtual computing resource, e.g., CPU cycles and network bandwidth, and local scheduling activities are simulated by virtual resource managers.

7 Experimental results

7.1 Implementation adequacy

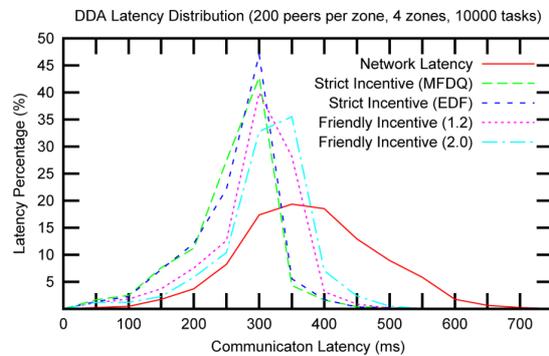
On a workstation with 2.4 GHz Intel Core2 Duo CPU and 2 GB memory, the test-bed simulates game sessions with 800 players. The maximum population supported in the experiment is mainly limited by the scalability of the Direct simulator. Experimental results show that during the bootstrapping time, various super-peer roles are selected successfully. Furthermore, under an average churn rate of 5% per minute, game zone structures are maintained correctly, and the leaving of both super and common peers are handled properly.

7.2 Real-time resource allocation

Figure 4 depicts the distribution of communication latencies among resource providers located by DDA, and corresponding resource consumers. First, the experimental results demonstrate that for all the tasks that have been measured, none of them missed their deadlines. DDA is capable of processing large number of tasks, whose dead lines vary between 3 s and 8 s. So, DDA seems able to satisfy the real-time requirement on resource discovery and allocation needed by a P2P MMOG.

Second, the ‘Network Latency’ curve demonstrates the actual communication latency distribution for the network topology used by the Direct simulator. By default, the distribution is approximately a Gaussian distribution $N(350, 100)$. If resource providers are randomly selected, the latency distribution for DDA should be similar to the default curve. However, as demonstrated by Figure 4, when RMs schedule tasks using an EDF queue, around 90% of the latencies are below mean latency. Furthermore, an MFDQ increases the probability for DDA to discover optimal resource providers. As a result, around 95% of the latencies are below mean latency. So, by carefully selecting resource providers, DDA is able to satisfy the QoS requirement on interactive NPC tasks needed by a P2P MMOG.

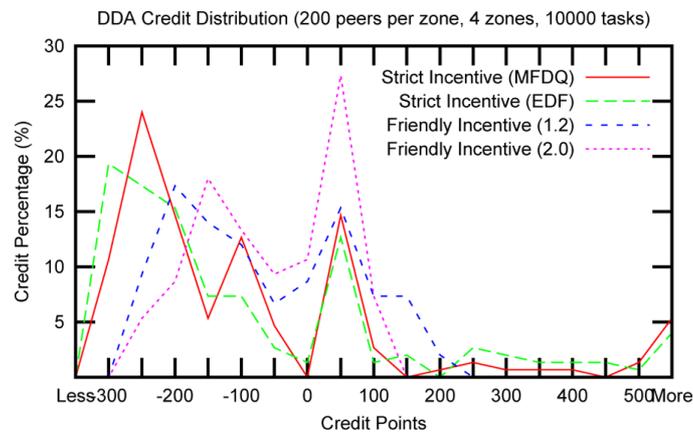
Figure 4 DDA latency distribution (see online version for colours)



7.3 Cooperative economic model

Figure 5 depicts the distribution of credits earned by the 800 players after 10,000 tasks. With the strict incentive policy, the gap between the richest and poorest players is wide. Specifically, around 5% of the players have earned more than 500 credits, but around 20% other players are in debt to the system for -300 credits. In contrast, with a friendly incentive policy ($\tau=1.2$), the gap between the rich and poor can be significantly narrowed (some peers are still in debt, owing to their inability to provide usable resources).

Figure 5 DDA credit point distribution (see online version for colours)



The impact of different incentive policies on resource quality is displayed in Figure 4. Using the strict incentive policy, only 5% of resource providers are selected when their latencies are higher than the mean. However, this ratio increases to about 30% using the friendly incentive policy ($\tau=1.2$). Similarly, when τ is more relaxed, e.g., $\tau=2.0$, the gap can be further narrowed in Figure 5, but the resource quality becomes worse. A P2P MMOG can use τ to customise its own cooperative economic model to strike a balance between fairness in task sharing and QoS for maintaining acceptable gaming experience.

8 Conclusion and future work

This paper presents DDAs, a novel task-mapping infrastructure for heterogeneous environments. The strength of DDA is to support self-organised real-time NPC host allocation in P2P MMOGs, as well as to meet the QoS requirements for game interactivity.

A prototype for DDA and a test-bed application have been implemented (Section 6). Experimental results show that, compared with previous work (Section 2), DDA has four main advantages. First, DDA is self-organising. Its infrastructure can be automatically assembled and maintained as peers join and leave the system (Section 7.1). Second, DDA is able to process large numbers of real-time NPC tasks and to support a simulated

P2P MMOG with 800 players effectively (Section 7.2). Third, DDA satisfies the QoS requirement for game interactivity by keeping the communication latency between NPC hosts and ordinary players low, e.g., 95% of the latencies are below mean latency (Section 7.2). Fourth, by applying a friendly incentive policy, DDA can establish a cooperative economic model that shares tasks among application participants fairly (Section 7.3).

In future work, DDA's security will be enhanced by a distributed reputation system, and a new resource matchmaking policy will take into consideration a resource provider's trustworthiness and dependability. Furthermore, though DDA is primarily designed to support NPC host allocation, it may also apply to general P2P applications with real time computational or interactive tasks. In future work, DDA's potential usage for other application types will be explored.

Acknowledgement

The authors gratefully acknowledge the contributions of Jeff Hoyer for tirelessly providing vital software maintenance for the FreePastry, the Direct simulator, and for many helpful suggestions that expedited this research.

References

- Anderson, T.W. (1984) *An Introduction to Multivariate Analysis*, 2nd ed., John Wiley & Sons, Hoboken, NJ, USA.
- Bharambe, A. (2006) 'Colyseus: a distributed architecture for online multiplayer games', *Proceedings of NSDI*, USENIX, San Jose, CA, USA, pp.3–6.
- Fan, L., Taylor, H. and Trinder, P. (2007) 'Mediator: a design framework for P2P MMOGs', *Proceedings of NetGames*, ACM, Melbourne, Australia, pp.43–48.
- Gupta, M., Judge, P. and Ammar, M. (2003) 'A reputation system for peer-to-peer networks', *Proceedings of NOSSDAV*, ACM, Monterey, California, USA, pp.144–152.
- Hu, S-Y., Chang, S-C. and Jiang, J-R. (2008) 'Voronoi state management for peer-to-peer massively multiplayer online games', *Proceedings of CCNC*, IEEE, Las Vegas, NV, USA, pp.1134–1138.
- Limura, T., Hazeyama, H. and Kadobayashi, Y. (2004) 'Zoned federation of game servers – a peer-to-peer approach', *Proceedings of NetGames*, ACM, Portland, Oregon, USA, pp.116–120.
- Litzkow, M., Livny, M. and Mutka, M. (1988) 'Condor – a hunter of idle workstations', *Proceedings of ICDCS*, San Jose, California, USA, pp.104–111.
- 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.
- Mulligan, J. and Patrovsky, B. (2003) *Developing Online Games – An Insiders Guide*, New Riders Publishing, Indianapolis, IN, USA, ISBN: 1592730000.
- Raman, R., Livny, M. and Solomon, M. (2000) 'Resource management through multilateral matchmaking', *Proceedings of HPDC*, IEEE, Pittsburgh, Pennsylvania, USA, pp.290, 291.
- Rowstron, A. and Druschel, P. (2001) 'Pastry: scalable, decentralized object location and routing for large scale peer-to-peer systems', *Proceedings of Middle Ware*, ACM, Heidelberg, Germany, pp.329–350.

- World of Legend (2004) Shanda Entertainment, wool.sdo.com
- World of Warcraft (2001) Blizzard Entertainment, worldofwarcraft.com
- Yonekura, T., Kawano, Y. and Hanawa, D. (2004) 'Peer-to-peer networked field-type virtual environment by using AtoZ', *Proceedings of 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', *Proceedings of NOSSDAV*, ACM, Skamania, Washington, USA, pp.99–104.