FloRA - Flock-Based Resource Allocation for Decentralized Distributed Virtual Environments

Jean Botev
University of Luxembourg
L-1359 Luxembourg
jean.botev@uni.lu

Steffen Rothkugel
University of Luxembourg
L-1359 Luxembourg
steffen.rothkugel@uni.lu

ABSTRACT

The growth of Massively Multiuser Virtual Environments (MMVEs), increasingly interactive social networking platforms and in particular their likely convergence render today's centralized hosting approaches impracticable. To handle potentially single-instance virtual environments of such massive scale, decentralized systems are necessary that also involve the resources of clients.

The expedient design of techniques for enabling this kind of next-generation Decentralized Distributed Virtual Environments (DDVEs) is a growing field of research. We aim at the provision of an infrastructure enabling such DDVEs in the HyperVerse project, focusing on collaboration and self-organization as means to achieve maximum scalability.

In this paper we present FloRA, a flock-based resource allocation scheme that helps alleviate the load imposed by regions with a higher user density as they often occur in DDVEs. Exploiting the heterogeneity of clients, only local information is utilized to tackle exigencies.

Evaluations show that both for the discovery of these regions and their alleviation the local views converge well to a global one, with favorable effects on the overlay topology.

Categories and Subject Descriptors

C.2.4 [**Distributed Systems**]: Distributed Applications; C.2.1 [**Network Architecture and Design**]: Distributed networks, Peer-to-peer computing

General Terms

Algorithms, Design, Performance

1. INTRODUCTION

Over the last decade a growing popularity of Massively Multiuser Virtual Environments (MMVEs) can be observed, be it social-centric online communities like for instance *Second Life* or game-centric environments, e.g. *World of Warcraft*, to name the largest. Particularly the idea of a three-dimensional Internet experience attracts much attention, and with the recent growth of social networking communities embracing interactive aspects, their convergence into some even larger, hybrid form becomes an easily understandable assumption and attractive idea. It is widely accepted that ultimately the user numbers of such a global-scale virtual environment render today's mostly centralized hosting approaches impractical and that the amount of users and associated data can only be handled on such a massive scale by decentralized - or at least hybrid systems - with a higher degree of commitment from the clients in terms of resources.

In the HyperVerse research project[1], we investigate fundamental principles suitable for the realization of such extremely large-scale and highly interactive Decentralized Distributed Virtual Environments (DDVEs) with the goal of creating a self-organizing and sustainable middleware service as a basis for future virtual environments.

While there is consensus on the fact that Peer-to-Peer (P2P) technologies are particularly promising to enable DDVEs and a range of approaches presented later in this paper exist, our approach differs in many aspects from them.

One key feature of the HyperVerse middleware is the underlying two-tier P2P architecture consisting of a highly structured backbone network of reliable, server-like machines and a loosely structured overlay on top. Detailed for instance in [2], the architecture is depicted schematically in Figure 1 with T_1 forming a loosely-coupled, geometric client overlay especially for real-time data distribution, and T_0 constituting a federated and highly structured backbone network underneath.

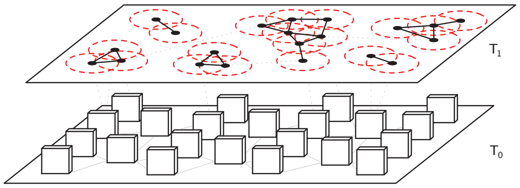

Figure 1: Two-tier HyperVerse architecture.

The dissemination of data in the HyperVerse is based on a Torrent-like technology [5], taking into account virtual geography and thus exploiting access locality. In order to provide the scalability needed to handle virtual environments of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
DISIO 2011, March 21, Barcelona, Spain
Copyright © 2011 ICST 978-1-936968-00-8
DOI 10.4108/icst.simutools.2011.245544

---

[1] http://hyperverse.syssoft.uni-trier.de

such large scale and dynamics, inter-client exchange of data is effected whenever possible. Clients are interconnected based on their proximity in virtual space via a geographic indexing service, effectuating specific qualities of the network, e.g. the power-law degree distribution, and allowing for self-adaptation of the statistical structure. By maintaining these direct links between nearby peers, costly routing mechanisms can be avoided as - due to the nature of virtual environments - most communication takes place in direct interaction or mutual visibility between peers and therefore is covered for most part by nearest neighbor connections.
An important objective is to develop adaptive, self-organized techniques already on the geometric overlay level to unburden and support the backbone which is responsible for neighbor discovery and the dynamic, fallback provision of data in the event of clients entering empty regions.

In this paper we present *FloRA*, a flock-based resource allocation scheme for discovering and alleviating critical regions with a high user density in DDVEs to cushion effects otherwise leading to costly readjustment operations at backbone level. By harnessing the heterogeneity of peers, more powerful ones particularly in terms of bandwidth and reliability provide so-called *virtual peers* to facilitate data distribution. Operating exclusively on local information, an extended flocking mechanism enables these virtual peers to monitor the environment and adaptively react to exigencies. Aggregate knowledge about current critical regions forms within the flock by means of a gossip protocol, ensuring a high responsiveness in dynamic environments.

The remainder of this paper is structured as follows: Section 2 gives a brief overview of user mobility and concentration in virtual environments, outlining the problem and substantiating the necessity for action. Section 3 then focuses on the FloRA scheme, introducing relevant concepts and detailing algorithmic aspects. Subsequently, an evaluation of the scheme will be provided in Section 4, followed by a discussion of related work in the context of P2P overlay networks, structural detection, load balancing and super-peer approaches for DDVEs in Section 5. Finally, we summarize our contributions and outline perspectives in Section 6.

2. USER MOBILITY AND CONCENTRATION

The movement of avatars, i.e. instances representing users and their behavior in virtual environments, can differ from real-life motion patterns due to the possibility of fast travel, teleportation or in general the fluctuation of users known as churn. Still, a large number of users is likely to remain stable in terms of locality as many virtual environments focus on user-to-user interaction or longer activities (also involving user-to-object interaction) that take place in one specific area. This collaborative aspect has the effect that motion will focus upon a specific set of locations, often with a self-energizing character, drawing in further users. These regions can form in an extremely dynamic fashion but on the other hand also become very stable and static. Generally, they are called *hot spots*, with their dynamic, burst-like variation being commonly known as *flash crowds*. A hot spot can informally be defined as a region within a virtual environment where a large number of users gather for a variety of possible reasons. For instance, hot spots can be landmarks of constant interest (e.g. a capital of a virtual region, sights etc.), or temporary interest (unique or spontaneous events like concerts, meeting places etc.). Similar phenomena can also be observed in social networks.

Efficiently and timely alleviating traffic bursts in hot spot regions is one of the prime concerns in the design of technologies for DDVEs.

3. MODEL

Prior to detailing the actual FloRA protocol itself, we will clarify some concepts, assumptions and terminology important to its mode of operation.

3.1 Peer Heterogeneity

Asymmetries regarding peer resources are inherent to large-scale P2P networks. Individual peers exhibit differences in terms of bandwidth, reliability, processing power et cetera, and exploiting this heterogeneity to the benefit of the overall system without increasing its susceptibility to node failure is an area of research attracting lots of interest. The goal is not to soften down the decentralization and create bottlenecks, when certain peers take over additional responsibilities in the network in order to harness existing heterogeneities and tap the full potential in terms of resources.

3.1.1 Multiplexers

We call the more powerful peers in our geometric overlay *multiplexers*. They are selected by the backbone when joining the network using a metric considering mainly bandwidth, but also taking into account previous session lengths. In certain P2P systems, it has been shown that previous session lengths can be used as an estimator for future session lengths [26]. This way, the existing heterogeneity can be harnessed to additionally improve overall system scalability and reliability. Because of the aforementioned two-tier structure of the HyperVerse, multiplexers can be identified without expensive collaborative calculation of a selection-parameter [7]. Apart from dealing with the visible representation of their user's avatar, the function of multiplexers is to split off and host one or more additional peers, referred to as virtual peers.

3.1.2 Virtual Peers

Constituting invisible representations of multiplexer resources but equaling the other peers in their responsibilities concerning data distribution, instead of being controlled by a user, virtual peers are driven by flock-based motion with an underlying epidemic scheme as described in the next section. The other significant difference to ordinary peers lies in the fact that virtual peers are inducing preferential connections, i.e. if a virtual peer is in proximity, a peer always rather connects to the virtual peer than to another peer. The rationale behind this preferential connection lies in the fact that virtual peers are provided by those peers with highest bandwidth and reliability. This way, they facilitate the distribution of object content and the dissemination of movement updates in a range around their virtual position.

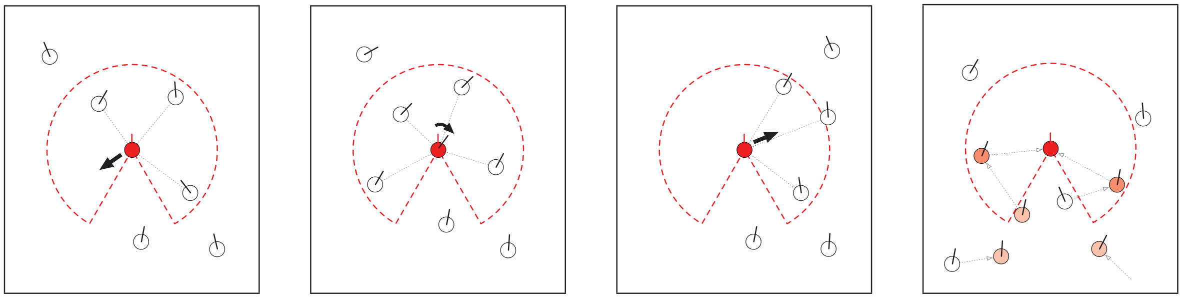

(a) Separation (b) Alignment (c) Cohesion (d) Hierarchization

Figure 2: Extended Flocking Rule Set

3.2 Flocking

The FloRA protocol draws inspiration from nature and relies on a flocking model for virtual peers. Essentially, it shapes a two-step process with the hot spot region discovery on the one hand and a subsequent alleviation on the other.

Flocking has attracted continuous interest and is studied in various scientific disciplines. From an algorithmic point of view as regards FloRA, it is particularly relevant that flocking is entirely decentralized and that it can be modeled based on three simple proximity-based rules as identified by Reynolds in [24]:

Separation - Assures that the virtual peers keep some distance between each other and thus prevents too tight formations or collisions.

Alignment - An important element that averages the overall course of neighboring virtual peers and thus generates and adds an aligned mean direction.

Cohesion - Allows for group formation of virtual peers with their neighbors by introducing a heading towards the average position of the virtual peers in proximity.

Recent empirical research results of studies involving pigeons suggest that real flocking behavior displays well-defined hierarchical structures which - from an evolutionary perspective, e.g. when foraging - are likely to be more efficient [19]. Therefore, we extend the above set by another rule inspired by the state-dependent social foraging model as described in [23]:

Hierarchization - Through the manifestation of leader-follower relationships in pair-wise interactions, these natural hierarchies between virtual peers lead to the emergence of alternating leadership, additionally influencing the flock heading.

Figure 2 illustrates the different rules, which together create an emergent and complex, yet efficient target-oriented flocking behavior. Centered in all cases is the virtual peer affected by the exemplary application of the particular rule. It is visualized together with its interaction radius, i.e. the area within which other virtual peers are considered neighbors. The short lines originating from each virtual peer depict the current orientation, while the missing piece at the back of the interaction area symbolizes that it can be anisotropic and therefore does not necessarily have to be fully circular. Like in nature, it could rather reflect some specific physical properties and perceptive capabilities. Unlike the other three illustrations that focus on positional updates, Figure 2d does not explicitly show the corrective influence on motion. It focuses on the establishment of a leader through implicit interactions, whose current direction then influences the heading of its neighbors. The developing hierarchical structure is reflected in the directed arrows along with the different nuances in the fills.

Having detailed the rule set inducing the flocking motion, the two principal algorithmic steps constituting FloRA will now be described in more detail.

3.2.1 Foraging / Discovery

While moving as a flock, the virtual peers gather information about the environment. In nature, this would be motivated through foraging, and analogous to that notion in FloRA densely populated regions, i.e. hot spots, constitute food sources of different size. Clearly, the larger the food sources, the more attractive they are, and their efficient discovery is crucial.

In between the virtual peers forming the flock, a gossip-based epidemic protocol is used to identify hot spots in proximity. Such protocols have emerged as an efficient communication paradigm especially for large-scale and dynamic distributed systems, maintaining both simplicity and scalability at constant communication cost [16]. Generally, they perform particularly well in network topologies with a good expansion [4] as exhibited also by the overlay we are operating on. The discovery is based on an aggregate of virtual peers' local perception of regional density, and information exchange with a random neighbor within the flock is effected in periodic intervals.

For the internal prioritization of hot spot regions as well as automatic motion adjustment of the virtual peers, we define the weight of a hot spot by the amount of contained peers:

Let h_i be a hot spot region within the virtual environment containing n peers. With $w_1, ..., w_n$ denoting each peer's individual weight, we define the cumulative weight W_i of the hot spot as

$$W_i = \sum_{i=1}^{n} w_i \quad . \tag{1}$$

With $l_1, ..., l_n$ representing the location of the individual virtual peers, a hot spot h_i's center C_i is then defined as

$$C_i = \frac{\sum_{j=1}^{n} l_j \cdot w_j}{W_i} \quad . \tag{2}$$

Each virtual peer maintains a fixed-size local list H_i of identified potential hot spot regions. L_i being a virtual peer vp_i's maximum lookahead distance, the following algorithm describes the aggregation in case that there is an exchange between two neighboring virtual peers vp_i and vp_j:

```
if $\exists\ (C_k, W_k) \in H_i$ with $dist(C_k, p_j) < L_j$ and $\max(W_k)$ then
    send $(C_k, W_k)$
else
    send $(l_i, w_i)$
end if
if $H_j$ full $\wedge\ \exists\ (C_l, W_l) \in H_j$ with $W_l < W_{proposal}$ then
    replace $(C_l, W_l)$ with $(C_{proposal}, W_{proposal})$
else
    add $(C_{proposal}, W_{proposal})$
end if
if $\exists\ (C_k, W_k) \in H_j$ with $dist(C_k, p_i) < L_i$ and $\max(W_k)$ then
    send $(C_k, W_k)$
else
    send $(l_j, w_j)$
end if
```

This constitutes a modified and extended version of the epidemic aggregation protocol presented in [14]. Range constraints are introduced to the maximum aggregation to reflect the desired locality in DDVEs. Only information within a virtual peer's lookahead range is considered, and hot spots transcending that range will be removed from the local list.

Accurate knowledge about hot spots is particularly important both for assuming flock leadership and the second part of the FloRA protocol described in the following section.

3.2.2 *Perching / Alleviation*

After their discovery, the actual resource allocation measures have to be taken in order to alleviate hot spots. Our protocol models this alleviation in analogy to the flocking behavior of birds as a perching step. In nature, perching can be triggered by fatigue and as a result of foraging, and we model it in a similar fashion.

Simulating fatigue can be done by assigning a random value within the range between 0 and a certain perching base which - unless modified by the following rules - will cause a virtual peer to perch from time to time. The foraging model however always leads to perching a minimum amount of time modified by factors influenced by the current situation of a virtual peer. In line with the previously introduced notation, the minimum perching duration D_{perch} for a virtual peer vp_i is modified according to the following two conditions:

```
if $\exists\ h_i \in H_i$ with $vp_i$ in $h_i$ then
    $D_{perch}$ $\times$= factor($W_i$)
end if
```

```
if $T_{min} \le deg_i$ then
    $D_{perch}$ += norm($deg_i$)
end if
```

It is sufficient that one of the two is fulfilled to trigger the perching of the affected virtual peer. The duration when actually sojourning within a hot spot known to a virtual peer through aggregate information thus will be influenced by its weight as well as the actual degree of incoming connections. Otherwise, if the amount of incoming connections exceeds the minimum threshold T_{min}, the virtual peer only speculatively perches based on more immediate, spontaneous assessment of the situation.

Together with the preferential connection to virtual peers mentioned in Section 3.1.2, resources are allocated where needed and - as a positive side effect - the overall network structure shifts in a favorable fashion. We will now give a concrete example of the different elements' interplay.

3.2.3 *Cumulative Behavior*

The overall behavior ensuing from the adaptive combination of foraging and perching is illustrated in Figure 3, which comprises a chronological series of FloRA visualization snapshots. The full video of this showcase simulation is available online[2] for download and viewing.

User mobility does not follow a uniform distribution that would lead to random motion patterns. Instead, as motivated in Section 2, peers are more likely to move towards high user densities in their proximity. In order to define the hot spot distribution, we devised a graphical probability map where the hot spot location probability is encoded over the area by grayscale values ranging from black (hot spot) to white (other). The map utilized here induces a hot spot distribution similar to the dots on a regular dice depicting the number five. Hot spot regions thus form over time in the center of the area as well as towards its corners.

The larger, hatched circular areas illustrate the currently aggregated local knowledge, i.e. where the virtual peers at that specific discrete time step of the simulation suspect hot spot regions. Like in the graphical depiction of the rules in Figure 2, a virtual peer's heading is symbolized by a line coming out of its center. For better visibility, in the snapshots virtual peers are marked additionally to the red color and thin line by a conic-shaped border stressing their current orientation. The current leader is identified by its dark green color. If there is none, the leader that previously emerged

---

[2] http://mocca.uni.lu/FloRA

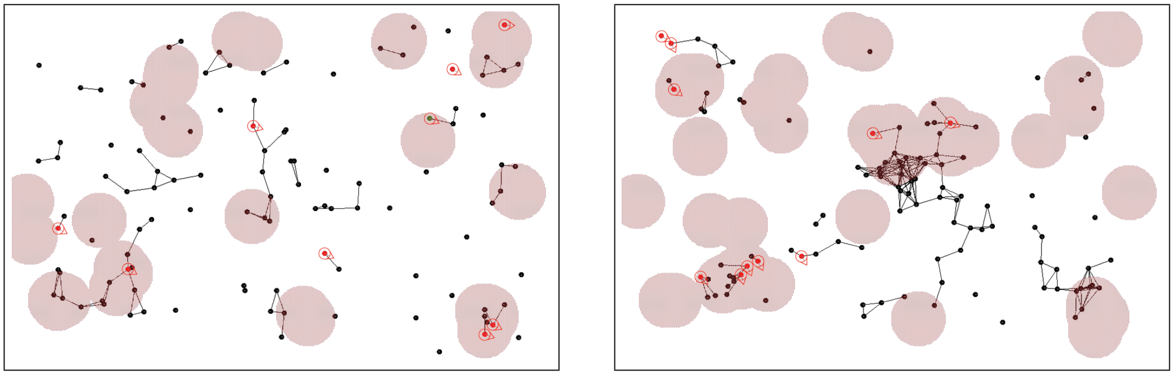

(a) Simulation Step 50

(b) Simulation Step 250

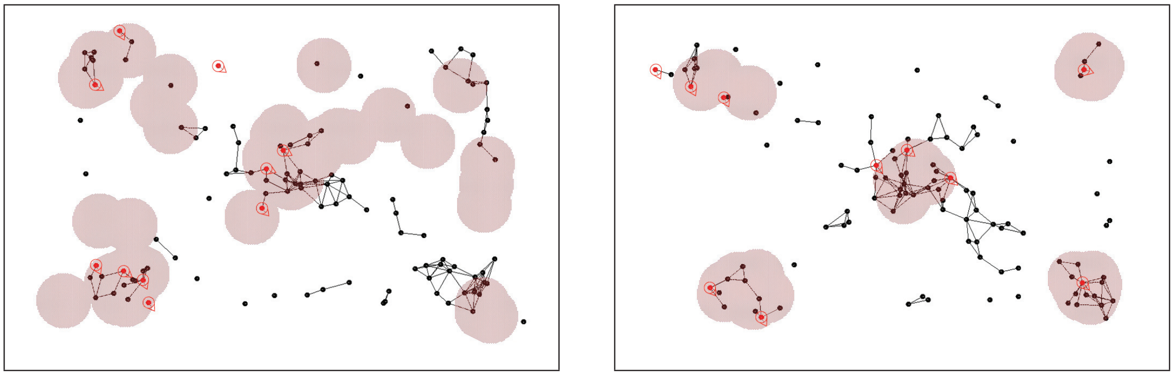

(c) Simulation Step 500

(d) Simulation Step 2500

Figure 3: FloRA visualization snapshots (virtual peers are marked additionally to guide the eye)

just perched with another one taking over that role in a subsequent step.

Figure 3a shows the initial situation where all peers began moving in the virtual environment after initially having been distributed randomly in the plane. If at all, only little information has been exchanged so far between the virtual peers which is why local knowledge about hot spots is rather based on the individual virtual peers' more immediate surroundings. For this reason, they appear to move freely at different velocities emerging from the basic rule set described earlier. In Figure 3b, the perceived hot spots grow in size and coverage by virtual peers improves as they speculatively perch based on incoming connections while discovering more and building up their knowledge. Hot spot locations in the process become more and more pronounced as depicted in the following Figure 3c, and the tendency towards consolidation about the assumed hot spot locations is clearly visible. As a consequence, virtual peers - if not moving - perch in the 'heaviest' hot spots which at this point are still partly indistinct but identifiable large patches. The other peers (including multiplexers) adopted a motion towards locations of interest according to the underlying statistical model defined by the probability map. Figure 3d shows a later view of the particular situation where hot spots have been discovered almost perfectly and covered by virtual peers.

4. EVALUATION

For an initial evaluation of FloRA, we made use of the network topology generation and simulation environment TopGen [25] which allows for deterministic and event-based simulation of user movement. The mobility model underlying the simulations on hand bears resemblance to the modeling of power law graphs described in [1], with the probability of visiting a certain location proportional to the number of peers in its surrounding area. The simulated region is sized 1000 x 650 pixels and constitutes an excerpt of a DDVE populated with 100 peers and an additional 10 virtual peers hosted by multiplexers, which is a conservative estimate of the multiplexer ratio as described in Section 3.1.

Because of being proportionately few in numbers, the virtual peers need slightly more time to aggregate knowledge on the 'heaviest' hot spots compared to when all peers would be involved in the process. Still defining a steep learning curve, as depicted in Figure 4, after around 730 steps (i.e. only about 6 exchanged messages per virtual peer) they achieve an average discovery rate of more than 80%. This is cushioned though by the speculative perching, which ensures an overall good accuracy of virtual peer positions under high dynamics but also in the built-up phase. After that initial bootstrapping due to the initially random distribution, virtual peer position reaches an average accuracy above 90%. The fluc-

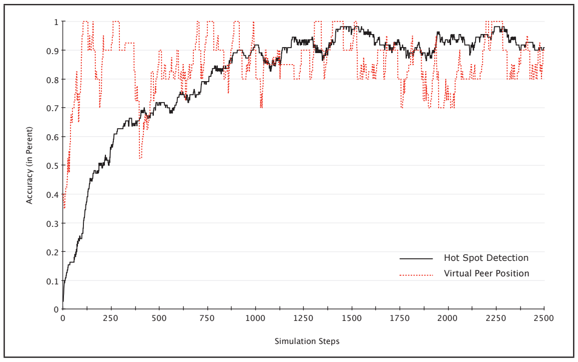

Figure 4: Accuracy Development

tuation here is generated by the dynamics itself and the perching mechanism with intentional moving on of perching virtual peers after a specific interval as described previously.

Figure 5 depicts the development of the connectivity, showcasing the influence of the preferential connection to virtual peers on the overlay. With increasing positional accuracy of the virtual peers, i.e. the coverage of relevant hot spots, they level the overall node degree out to an average value of circa 2.7, leading to an overall sparser topology further unburdening the backbone.

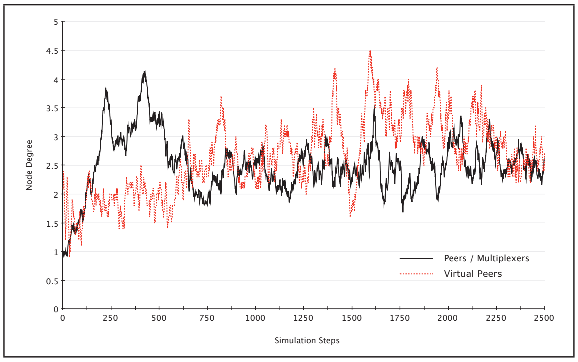

Figure 5: Node Degree Development

The fact that hot spot information is aggregated only within the small group of virtual peers leads to a message reduction in comparison to approaches involving each peer in the aggregation. Still, the protocol leads to a fast-converging global discovery at constant communication cost. The maximum period relevant hot spots remain uncovered by virtual peers furthermore is kept low due to the high dynamics of the flocking scheme.

5. RELATED WORK

The issue of identifying and alleviating hot spots and, as their more dynamic variant, flash crowds occurring in distributed systems is a widely studied subject. In client-server settings, traditionally a detection would be followed - if at all - by measures usually entailing substantial loss of quality of service. For instance in [3] performance degradations are being monitored to detect hot spots, with the consequence being that no more users are admitted to the identified regions, resulting in a seriously impaired user experience. Later client-server based approaches like [22] and [21] have already identified P2P technologies as complementary measure, utilizing P2P networks if necessary to unburden the initial centralized network while not substantially derogating the experience.

P2P architectures specifically designed for the provision of DDVEs are for instance VON [12], Solipsis [10], FLoD [13] or Mediator [8]. While particularly the more recent approaches consider the heterogeneity of peers to allocate different roles to them within the network, they basically constitute a single P2P network which then is structurally rearranged to mitigate effects arising from hot spot regions. Common techniques are based on interest management, for instance the reduction of the size of the individual Area of Interest (AoI) as performed in the Voronoi-based clustering in [12], or the adjustment of AoI shapes in FLoD [13]. Also in [10], the mitigation involves immediate and more costly readjustments of the supporting network.
The idea to exploit heterogeneity in P2P networks in a similar way to cushion effects arising from hot spots can be mainly found in super-peer load balancing approaches. These super-peers as investigated e.g. in [15], [27] or [11] resemble the multiplexers in our scheme and take additional tasks according to their capability and then reshape the networks accordingly. In [27] the idea of so-called *virtual servers* is introduced, which are altruistically maintained data hosting blocks kept within peers. Each peer maintains at least one of these, and an algorithm constructs another overlay that efficiently distributes the virtual servers in a way that super-peers shoulder a large amount of the data to balance the load. Mainly in connection with spatial computing, some approaches introduce a distance measure for their protocols such as in [15]. Super-peer election is handled in a decentralized fashion, and has also been considered in [18] for unstructured networks. The proposed H_2O protocol is an interesting option in case that there is no reliable backbone facilitating super-peer election, but while mentioning a possible utilization of super-peers to allocate resources where needed, no concrete method is presented. Another super-peer approach operating on an unstructured network level is described in SOSPNet [11]. Load balancing is tackled here by automatically discovering semantic relationships between files to which pointers are maintained by the super-peers which form the second layer in a dual-layer network.
Most super-peer mechanisms dealing with phenomena similar to hot spots have been considered in conjunction with file sharing scenarios, and thus do not consider specific properties of virtual environments influencing their formation like the inherent spatio-temporal aspects and characteristics of user mobility.

Spatiality however is intrinsic to flocking, and some research in that field has been devoted for instance on harnessing this property in order to detect structures [9] or to generate clusters [6] in data sets projected onto a two-dimensional plane. However, in both cases data are rather static and the approaches are targeted mostly at data and/or pattern mining in post-processing. A general overview of flocking techniques in distributed dynamic systems is for instance provided in [20].

6. CONCLUSION AND PERSPECTIVES

In this paper we have presented FloRA, an efficient flock-based resource allocation scheme for P2P-based DDVEs. Exploiting the heterogeneity of clients, peers that are more reliable and stronger particularly in terms of bandwidth contribute excess resources for the distribution of data by splitting off virtual peers. These follow an entirely self-organized, bipartite flocking scheme for the discovery and alleviation of hot spot regions as they commonly occur in such environments. Aggregating necessary information in a gossiping manner by exclusively utilizing locally available data, the individual views converge well to a global one. Showing high resilience towards peer dynamics, FloRA effectuates favorable structural overlay adjustments.

Shortcomings of the presented approach will be discussed in the following. Firstly, we utilized an artificial mobility model for the conducted simulations not involving churn explicitly. It is therefore planned to perform additional simulations and comparative measurements on the basis of real-world avatar traces from Second Life gathered at the University of Singapore [17]. The results of such a large-scale evaluation could deliver interesting insights for optimizations of our scheme. Secondly, further spatio-temporal analysis also on potential anomalies similar to locality-based failures in greedy routing should be conducted. Albeit borderline situations, they might occur under certain circumstances and thus have to be investigated and ruled out as far as possible.

Initially designed for the application in DDVEs, we believe FloRA opens up interesting perspectives for highly dynamic spatio-temporal systems and their monitoring in general.

7. ACKNOWLEDGMENTS

The work presented in this paper was supported by the Fonds National de la Recherche (Luxembourg).

8. REFERENCES

[1] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

[2] J. Botev, M. Esch, A. Höhfeld, H. Schloss, I. Scholtes, and P. Sturm. The HyperVerse - concepts for a federated and torrent-based ”3D web”. *International Journal of Advanced Media and Communication*, 2(4):331–350, 2008.

[3] X. Chen and J. S. Heidemann. Flash crowd mitigation via adaptive admission control based on application-level observations. *ACM Transactions on Internet Technology (TOIT)*, 5(3):532–569, 2005.

[4] F. R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, Providence, RI, 1997.

[5] B. Cohen. Incentives build robustness in BitTorrent. In *Proceedings of the 1st Workshop on the Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, 2003.

[6] X. Cui, J. Gao, and T. E. Potok. A flocking based algorithm for document clustering analysis. *Journal of Systems Architecture*, 52(8-9):505–515, 2006.

[7] M. Esch, J. Botev, H. Schloss, and I. Scholtes. GP3 - a distributed grid-based spatial index infrastructure for massive multiuser virtual environments. In *Proceedings of the 14th IEEE International Conference on Parallel and Distributed Systems (ICPADS 2008)*, pages 811–816, Melbourne, Australia, 2008.

[8] L. Fan, H. Taylor, and P. Trinder. Mediator: A design framework for P2P MMOGs. In *Proceedings of the 6th ACM SIGCOMM workshop on Network and system support for games*, pages 43–48, Melbourne, Australia, 2007.

[9] G. Folino and G. Spezzano. An adaptive flocking algorithm for spatial clustering. In *Proceedings of the 7th International Conference on Parallel Problem Solving from Nature (PPSN VII)*, pages 924–933, Granada, Spain, 2002.

[10] D. Frey, J. Royan, R. Piegay, A.-M. Kermarrec, E. Anceaume, and F. L. Fessant. Solipsis: A decentralized architecture for virtual environments. In *Proceedings of the 1st International Workshop on Massively Multiuser Virtual Environments (MMVE)*, Reno, Nevada, USA, 2008.

[11] P. Garbacki, D. H. J. Epema, and M. van Steen. Optimizing peer relationships in a super-peer network. In *Proceedings of the 27th International Conference on Distributed Computing Systems (ICDCS ’07)*, page 31, Toronto, Ontario, Canada, 2007.

[12] S.-Y. Hu, J.-F. Chen, and T.-H. Chen. Von: A scalable peer-to-peer network for virtual environments. *IEEE Network*, 20(4):22–31, 2006.

[13] S.-Y. Hu, T.-H. Huang, S.-C. Chang, W.-L. Sung, J.-R. Jiang, and B.-Y. Chen. FLoD: A framework for peer-to-peer 3D streaming. In *Proceedings of the 27th Conference on Computer Communications (IEEE InfoCom ’08)*, pages 1373–1381, Phoenix, AZ, USA, 2008.

[14] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems*, 23(3):219–252, 2005.

[15] G. P. Jesi, A. Montresor, and O. Babaoglu. Proximity-aware superpeer overlay topologies. *IEEE Transactions on Network Science and Service Management*, 4(2):74–83, 2007.

[16] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS’03)*, pages 482–491, Cambridge, MA, USA, 2003.

[17] H. Liang, R. N. Silva, W. T. Ooi, and M. Motani. Avatar mobility in user-created networked virtual worlds: measurements, analysis, and implications. *Multimedia Tools and Applications*, 45(1-3):163–190, 2009.

[18] V. Lo, D. Zhou, Y. Liu, C. Gauthierdickey, and J. Li. Scalable supernode selection in peer-to-peer overlay networks. In *Proceedings of the Second International Workshop on Hot Topics in Peer-to-Peer Systems (HOT-P2P ’05)*, pages 18–27, La Jolla, CA, USA, 2005.

[19] M. Nagy, Z. Ákos, D. Biro, and T. Vicsek. Hierarchical group dynamics in pigeon flocks. *Nature*, 464(7290):890–893, 2010.

[20] R. Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006.

[21] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai. Distributing streaming media content using cooperative networking. In *Proceedings of the 12th International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV 2002)*, pages 177–186, Miami Beach, FL, USA, 2002.

[22] C. Pan, M. Atajanov, M. B. Hossain, T. Shimokawa, and N. Yoshida. FCAN: Flash crowds alleviation network. In *Proceedings of the 2006 ACM Symposium on Applied Computing (SAC)*, pages 759–765, 2006.

[23] S. A. Rands, G. Cowlishaw, R. A. Pettifor, J. M. Rowcliffe, and R. A. Johnstone. Spontaneous emergence of leaders and followers in foraging pairs. *Nature*, 423(6938):432–434, 2003.

[24] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4):25–34, 1987.

[25] I. Scholtes, J. Botev, M. Esch, A. Höhfeld, H. Schloss, and B. Zech. TopGen - internet router-level topology generation based on technology constraints. In *Proceedings of the 1st International Conference on Simulation Tools and Techniques for Communications, Networks and Systems (SIMUTools 2008)*, pages 1–10, Marseille, France, 2008.

[26] D. Stutzbach and R. Rejaie. Understanding churn in peer-to-peer networks. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement (IMC '06)*, pages 189–202, Rio de Janeiro, Brazil, 2006.

[27] S. Surana, B. Godfrey, K. Lakshminarayanan, R. Karp, and I. Stoica. Load balancing in dynamic structured peer-to-peer systems. *Performance Evaluation*, 63(3):217–240, 2006.