

Decentralized Scale-free Network Construction and Load Balancing in Massive Multiuser Virtual Environments

Markus Esch and Eric Tobias

University of Luxembourg

Faculté des Sciences, de la Technologie et de la Communication

Luxembourg, Luxembourg

Email: markus.esch@uni.lu, eric.tobias.001@student.uni.lu

Abstract—Immersive and interactive Massive Multiuser Virtual Environments currently attract a lot of attention. Going beyond the hype, one discovers interesting concepts and opportunities for a supposed future Web which is commonly called the *3D Web*. The global provision of such a vision is much more than just another evolutionary step in content visualization, rather severe technical challenges need to be addressed by the future Internet infrastructure and protocols. Existing commercial approaches feature only limited scalability due to mostly relying on centralized solutions. With global-scale scenarios like a 3D Web in mind, the application of distributed and decentralized techniques becomes indispensable. In [12] we have introduced the concept of a scale-free and self-organized overlay network for the interconnection of machines that build the backbone of a global-scale online environment. This paper presents a detailed analysis and evaluation of this overlay network based on simulation results. The paper shows the feasibility of the concept for self-organized and decentralized load balancing in highly dynamic environments like virtual online worlds. Moreover the scale-freedom of the network's link structure is shown, which makes it usable for the intended global-scale scenario.

Index Terms—MMVE, DVE, P2P Overlay, Self-organization, Complex Networks

I. INTRODUCTION

Massive Multiuser Virtual Environments (MMVEs) like *Second Life* or *World Of Warcraft* currently attract a lot of attention. Especially the idea of a *3D Web* as global-scale combination of MMVEs and today's Web is very popular as it provides a variety of interesting opportunities and is often seen as the future of the Internet. Envisioning a fusion of today's Web content and avatar-based interaction, a user moves an avatar through a 3D online world in order to meet friends, undertake a sightseeing tour, shop and so forth. Provided on a global scale and in combination with the expected advances in human interface technologies, such a 3D Web enables also virtual mass events and immersive interaction. While the opportunities of such a scenario are alluring, its realization poses severe technical challenges to the underlying Internet infrastructure. The central question is how scalability and interactivity in a single instance of a MMVE can be reached in a way that guarantees high responsiveness to all users.

Due to the difficulty and actuality of the problem, a lot

of research is currently conducted in this field. Although decentralized approaches clearly offer many advantages for such a scenario, current commercial precursors of a 3D Web for the most part rely on centralized client/server solutions. In order to solve the scalability issue, separate world instances are provided, each supporting at most a few thousand users. Despite centralized architectures facilitating manageability and controllability for their providers, it is still commonly recognized though, that a global-scale 3D online world can only be provided in a decentralized, Peer-to-Peer (P2P) fashion. Moreover a federated infrastructure is mandatory in order to retain scalability, fault-tolerance and freedom of the Internet. In [8] we have presented a two-tier P2P architecture incorporating a highly structured overlay of reliable server machines (so-called Public Servers) and a loosely structured overlay of user clients as basic infrastructure for an open and federated 3D Web. In this architecture, the federation of reliable server machines builds the backbone of the 3D Web and is responsible for reliable data hosting and indexing as well as client management. These servers need to collaborate closely in order to provide a single instance of a continuous and persistent virtual environment at a massive scale.

For the interconnection of the Public Servers, we have presented an overlay network, allowing for the particular properties of MMVEs, like high avatar dynamics and non uniform load distribution in [12]. This overlay utilizes the expected scale-free distribution of Public Server capacities to form a network, that emerges into a scale-free state with a power-law degree distribution in a self-organized fashion. By this means, the overlay benefits from the particular advantages of scale-free networks like short average path length and robustness. Moreover the non uniform and fluctuating load distribution in online environments is balanced in a self-organized fashion by the Public Servers. This paper presents a detailed evaluation and analysis of this network overlay, based on simulation results. After briefly introducing both the general HyperVerse infrastructure as well as the scale-free and self-organized overlay in Section II, Section III presents the analysis of all aspects of the backbone overlay. Related work is presented in Section IV and the paper concludes with

a discussion of our main contributions and open issues in Section V.

II. MMVE INFRASTRUCTURE

This section at first describes the HyperVerse architecture, a general infrastructure concept for virtual online worlds, which has been presented in [7]. Based on this concept in [12] a scale-free and self-organized network overlay has been proposed as backbone infrastructure for MMVEs. The two integral parts of the overlay are a self-organized load balancing scheme as well as an algorithm for the decentralized construction of a scale-free link structure interconnecting the network nodes. Section II-A briefly describes the basic two-tier infrastructure. Afterwards Section II-B depicts the concept of the self-organized load balancing while Section II-C presents the decentralized construction of a scale-free link structure.

A. Two-tier architecture

The HyperVerse infrastructure relies on a two-tier P2P architecture incorporating two network overlays: one loosely structured and the other structured. A detailed description of this infrastructure has been presented in [8]. Figure 1 depicts the basic architecture, where the structured backbone overlay of server machines (so-called *Public Servers*) is responsible for the reliable hosting of the virtual world, while the loosely structured overlay of user clients is primarily utilized for data distribution. The Public Servers are machines that resemble today's Web Servers, and they are not required to be under control of any centralized authority. Similarly, for their provision we rely on the incentive of being able to publish information in the 3D Web. For this reason, someone who wants to publish information in the 3D Web has to either run or have access to a Public Server that hosts this data. The federation of these Public Servers then forms the actual backbone of the infrastructure. The Public Servers are interconnected by the self-organized and scale-free overlay network that will be presented in the subsequent sections. The main tasks of the backbone overlay are indexing of the 3D objects in the virtual world, tracking of avatar positions as well as interconnecting avatars in virtual proximity to allow interaction. Due to the ever increasing client bandwidth it is mandatory to utilize the client resources as well. For this reason the clients are interconnected by a loosely structured P2P overlay that is used for torrent-based data distribution.

B. Self-organized Load Balancing

The task of the backbone service constituted by the Public Servers is indexing of world objects as well as interconnecting clients in virtual proximity in order to allow interaction and mutual rendering. For this purpose the world surface is subdivided into small cells, each cell managed by one Public Server. For subdividing the plane we are using a Voronoi diagram. To define the Voronoi diagram, a virtual position on the world surface is assigned to each Public Server. In the resulting Voronoi decomposition each Public Server manages the Voronoi cell surrounding its virtual position. In order to




Fig. 1. Two-Tier HyperVerse Infrastructure Incorporating Reliable Server Machines and User Clients

handle load imbalances occurring in virtual worlds, the load of the Voronoi cell a Public Server is hosting needs to correspond to its capacities. This means hotspots with many clients and 3D objects need to be managed by machines with high capacities, while less crowded areas have to be handled by less powerful machines. In order to be able to adapt the overlay to dynamic load shifts in the online world, the virtual positions are not fixed. Rather, the servers modify their position according to the load induced by the users of the online world. The virtual positions are managed in a self-organized fashion, based on a set of rules applied by each server. For the application of these rules a *mass* is assigned to each object or avatar, its mass reflecting the load it induces on the server backbone.

A Public Server then is able to calculate the absolute mass of the objects in its cell as well as the center of mass. Depending on its bandwidth and computational power, a maximum payload p (indicating the mass it can handle) is assigned to each Public Server. Based on information about the masses and payloads of its immediate neighbors each Public Server is able to apply a set of rules adjusting the own position. The aggregated behavior of all Public Servers then yields the decentralized distribution of the world surface according to the Public Server's capacities in a self-organized fashion. For detailed information about these rules we refer to [12].

C. Scale-free Link Structure

In addition to the distribution of the world surface, it is important how the Public Servers are linked together. In first place all nodes have links to the nodes managing the bordering Voronoi cells. Additionally, links to other nodes in the network are established in order to get a scale-free degree distribution and to allow fast routing to all positions in the virtual world. The aim is to construct a scale-free link structure in order to benefit from the particular advantages of this network type like short average path length, resilience against random node failure and small diameter. Scale-free networks are defined as networks with a power-law node degree distribution where the power-law exponent is in the range of two and three. These kind of networks have many advantageous properties like fault

tolerance and short path length, moreover for many naturally occurring networks (e.g. World Wide Web [1], metabolic network [19], protein interaction network [18]) a scale-free node degree distribution has been shown. To get a network, that emerges into a scale-free state links are established by a preferential attachment scheme, based on bandwidth and computational power.

For the decentralized construction of the scale-free overlay we proposed a preferential attachment scheme similar to the Barábasi/Albert model [4]. In order to adopt the preferential attachment scheme to our scenario the way preferential attachment evolves in networks like the Web needs to be considered. In the WWW one can assert that the number of links to popular Web sites increases faster due to their high profile. At the same time, the providers take care for allocating sufficient server capacities. That means, the scale-free link structure and the existence of hubs with sufficient capacities emerges in a self-organizing manner from the different popularity of Web sites. Transferring this observation to our scenario of a virtual online environment, it means, that Public Servers hosting very popular world objects automatically exhibit higher capacities, since this is automatically ensured by the object providers. Since the scale-free capacity distribution thus emerges automatically, we just have to construct the link structure accordingly. For this reason, we establish links to Public Servers with high capacities, with a higher probability. Hence a new node is connected to an existing node i with probability Π depending on the payload p_i of i and the maximum payload in the network p_{\max} :

$$\Pi(p_i) = \frac{p_i}{p_{\max}} \quad (1)$$

Based on this probability distribution a joining node establishes M ($M > 0$; M is a fixed parameter of the algorithm) links to existing nodes by choosing a random position in the virtual world and connecting to the node I hosting this position. Thereupon the new node selects M neighbors from the two-hop neighbors of I according to the probability distribution. The maximum payload in the network can efficiently be aggregated using a epidemic aggregation scheme [17]. For more detailed information about the algorithm we again refer to [12].

III. EVALUATION

In order to evaluate the scheme described in the previous section, the algorithm has been simulated using the network overlay simulator TopGen [31]. The results will be presented in the following sections. At first Section III-A evaluates the self-organized load balancing scheme. Subsequently Section III-B presents the evaluation of the algorithm for the decentralized construction of a scale-free link structure.

A. Load Balancing

The self-organized plane partitioning scheme described above is used to distribute the load induced by user clients and world objects uniformly among the Public Servers. For

the usability of the load-balancing scheme the following points are important:

- The entire load induced by clients and objects needs to be distributed in a way that preferably no Public Server is overload and the global load can be handled.
- The payload of a Public Server needs to correspond to the load density in its Voronoi cell. This means, hotspots need to be hosted by high payload servers while regions with lower density have to be hosted by servers with lower payload. This is important, because if nodes with low payload host regions with a high mass density, the cells become very small which leads to a high overhead since clients often cross cell borders.
- Due to the high dynamic and unforeseeable behavior of the clients it is hard to reach, that nodes are never overloaded once a certain overall network load is reached. For this reason it is important that overloaded nodes are quickly unloaded by the neighbors in order to terminate the overload situation.

To analyze the proposed load balancing scheme with respect to these three criteria, the algorithm has been simulated. Before presenting the simulation results, we will briefly describe the simulation settings. Initially a set of 100 Public Servers has randomly been placed on a plane with a size of 800·600 units. To each Public Server a certain payload has been assigned according to a scale-free distribution with power-law exponent 2.1. Additionally a number of world objects and user clients have been added. Each of these entities induces a load corresponding to their movement speed. Because faster moving objects induce a higher load on the backbone overlay. The number of world entities depends on the requested utilization of the overall available Public Server payload. In the course of the simulation, the world entities move according to a random waypoint model. In order to simulate the hotspots occurring in virtual online worlds, certain spots are selected with a higher probability as waypoints. The Public Servers act according to the algorithm described in Section II-B in order to balance their load and to handle dynamic load shifts by adjusting their virtual position. Figure 2 shows a screenshot of the visualization of the simulation. Here one can assert a cumulation of Public Servers at the hotspots. Using these settings the simulations and measurements, discussed in the following, have been performed.

1) *Load Balancing*: At first the performance of the algorithm under different network loads has been studied. For this purpose, the algorithm was simulated with overall world entity masses of 40, 60 and 80 percent of the available payload. In the course of these simulations, the ratio of overloaded nodes as well as the ratio of the overload mass, which is the overall amount of masses exceeding the overloaded node's payload, has been measured. In order to classify these results, the same data has also been captured for a scheme that distributes the world surface uniformly and statically among the Public Servers. The results of this simulations are shown in Figure 3 and 4. Considering the ratio of overloaded nodes one can




Fig. 2. Virtual World Surface Distributed Among the Public Servers by the Proposed Self-organized Scheme. The Black Dots Represent World Entities. Each of the Voronoi Cells is Hosted by One Public Server. One can Assess a Cumulation of Public Servers at Hotspots.

assess that in any case the self-organized scheme performs better than an uniform distribution. Aside from a small organization period at the beginning the overload ratio is more or less constant during the whole simulation time. It stands out, that the overload ratio in the beginning of the simulations is considerably higher. This can be explained by the fact that the Public Servers are initially placed randomly. For this reason it takes some time till the servers have organized and the load is balanced. With 40 percent payload utilization almost no overloaded nodes occur. The rare overloaded nodes still occurring can be explained by the high dynamic of the user clients. Even with 60 percent payload utilization the amount of overloaded nodes is quite low between one and three percent. Only with a load of 80 percent, higher overload ratios between 10 and 15 percent occur. Due to the high dynamic induced by constant and unpredictable movement of the user clients overload situations are hard to avoid. But, as the analysis of the average overload periods will show, the overlay reacts quite fast on such overload situations and unloads the overloaded nodes. Moreover the overload reduction compared to an uniform and static distribution of the word surface is considerable. Our scheme could be improved by utilizing a prediction of the future client movements. This would allow a proactive handling of overload situations which could reduce the overload periods or even avoid them. Figure 4 shows the ratio of overload masses. Here similar results can be found, while the difference between the two compared schemes is yet bigger. Moreover the overload ratio using our self-organized scheme is very low, with under 3 percent even with 80 percent payload utilization.

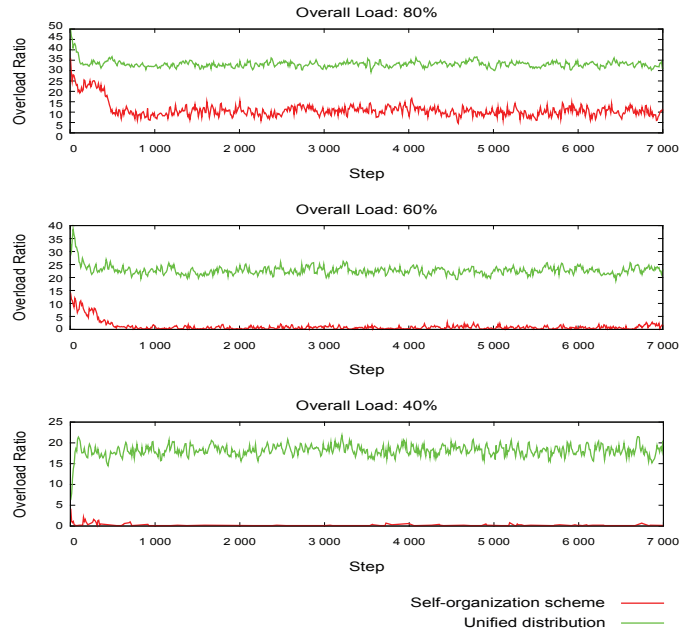


Fig. 3. Ratio of Overloaded Nodes for Networks With 40%, 60% and 80% Payload Utilization.

2) *Hotspot Accuracy*: In order to verify the second requirement it has been measured how accurate the high payload servers are placed. For this purpose the upper ten percent of servers in regard to payload have been considered. Figure 5 depicts the ratio of these nodes hosting the cells with the fifteen percent highest cell masses. This measurement has also been done for overall network loads of 40, 60 and 80 percent. Aside from the organization period in the beginning,




Fig. 4. Overload Mass Ratio for Networks With 40%, 60% and 80% Payload Utilization.

the hotspot accuracy is quite good and becomes better with increasing network load. While the accuracy is between 70 and 80 percent for 40 percent network utilization it rises to 90–100 percent for 60 and 80 percent network load. Moreover the accuracy is relatively constant during the whole simulation time. This shows that the algorithm reaches a good assignment of nodes with a high payload to the hotspot regions, which is an important criterion for the usability of the scheme.




Fig. 5. Hotspot Accuracy in Percent for Networks With 40%, 60% and 80% Payload Utilization.

3) *Overload Duration*: Above it has been shown that overloaded nodes cannot totally be avoided, especially in the presence of high network loads. This can be explained by the high client dynamic. In order to show, that the proposed algorithm is able to react quickly if a node becomes overloaded, the average overload period has been measured. The results are shown in Figure 6 which depicts the average duration of a overload situation for the different network loads and for the

self-organized scheme as well as for the uniform distribution. For all network loads, the average overload period is low using the self-organized scheme and is actually decreasing after the organization period in the beginning. For the uniform distribution overload periods last considerably longer and do not decrease.

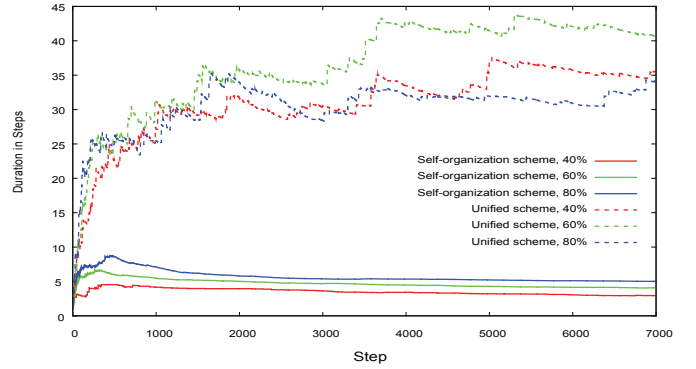


Fig. 6. Average Overload Period in Simulation Steps for Networks With 40%, 60% and 80% Payload Utilization.

This section presented the analysis of the self-organized load balancing scheme which is used to distribute the world surface corresponding to the given payload distribution of the Public Servers. The evaluation has shown that the algorithm can fulfill the three main requirements stated at the beginning of this section quite good. This illustrates the usability of the approach for load distribution in a dynamic MMVE scenario.

B. Link Structure

The purpose of the algorithm described in Section II-C is the decentralized construction of a scale-free link structure. This scheme has also been analyzed by TopGen simulations. The particular aim of this evaluation was to show that the algorithm indeed constructs a scale-free overlay network with a power-law node degree distribution. Additionally, important properties of the resulting graph like the clustering coefficient, the diameter, the average shortest path length and the correlation have been measured. Two different situations have been simulated and analyzed: (a) the network construction and (b) the behavior of the network under churn. The results will be presented in the subsequent sections.

1) *Construction*: For the evaluation of the network construction scheme, networks with 10000 nodes have been generated and analyzed. Since the power-law exponent of the payload distribution affects the resulting network, simulations with the three different exponents 2.1, 2.4, and 2.7 have been performed. These exponents have been chosen, since analyses have detected these exponents in important computer networks and since it is expected that a future 3D Web exhibits similar statistical properties. For the WWW an indegree power-law exponent of 2.1 [24] has been shown, while the outdegree exponent is 2.4 [9]. In addition it has been measured that the degree (in- and outdegree) of the Internet router network is 2.7 [14]. Another important parameter of the algorithm is the

number of links (M), assigned to each node when joining the network. To analyze the effect of this parameter simulations with M -values of 10, 20 and 30 have been executed.

Figure 7 shows the node degree distribution for networks generated with the given power-law exponents and M -values on a double logarithmic scale. For each exponent, networks with 5000 to 10000 nodes are depicted. The monotonic decrease of the curves shows the power-law degree distribution. Note, that the noise at the tail of the distribution is common in empirical data. In order to show the scale-freedom of the networks a maximum likelihood power-law fit, as described in [11] has been made. The results for networks with 1000 to 10000 nodes are depicted in Figure 8. One notices, that the power-law exponent is between 2 and 3 for $M = 10$ at each network size, proving the scale-free property of the network. For $M = 20$ and $M = 30$ the power-law exponent increases and actually temporarily exceeds the upper bound of 3. This shows, that M is a critical parameter of the algorithm. The increased power-law exponent can be explained by the fact that more links get assigned to each node, which leads to more nodes with a high degree. This shows, that the construction of a scale-free network can efficiently be done by drawing just a small number of edges per node.

Additionally some other network properties like the clustering coefficient, the average shortest path length as well as the diameter have been measured. Figures 9, 10 and 11 show these results for networks generated with different payload distribution exponents and network sizes between 1000 and 10000 nodes. The clustering coefficient and the average shortest path length of the scale-free network are compared to random graphs with the same number of nodes and edges. This comparison shows the *small-world* property of the generated networks, because according to [35] a graph is a small-world graph if (a) the graph has a small characteristic path length similar to a random graph and (b) the clustering coefficient is considerably higher than for a random graph with the same number of vertices and edges. Additionally Figure 11 shows the small diameter of the network for each configuration.

For the characterization of complex networks, the degree-degree correlation is an important property that is well studied. This property indicates the average degree of the neighbors of a given node. With $P(k'|k)$ being the probability that a node with degree k is connected to a node with degree k' the average degree of the nearest neighbors (nn) of a vertex with degree k is defined as [28]:

$$\bar{k}_{nn}(k) = \sum_{k'} k' P(k'|k). \quad (2)$$

According to [27] a network is correlated if the function $\bar{k}_{nn}(k)$ is increasing or decreasing and uncorrelated if the average neighbor degree is independent of k . As it has been shown in [27] many real world networks like the Internet exhibit correlations. In addition [21] investigates growing networks and asserts the inevitable presents of correlations in this kind of networks. For this reason, it is not surprising,




Fig. 7. Node Degree Distribution on a Double Logarithmic Scale for Networks With 5000 and 10000 Nodes.




Fig. 8. Fitted Power-Law Exponent

that the network generated by our approach exhibits degree-degree correlations as well. This can be seen by the monotonic

decreasing curve of the function $\bar{k}_{nn}(k)$ depicted in Figure 12 which is typical for a *disassortative mixing* according to [27]. If a uncorrelated overlay network is desired, we propose to




Fig. 9. Clustering Coefficient for the Scale-Free Graphs and for Random Graphs With the Same Number of Nodes and Links

utilize the rewiring scheme described in [25] to eliminate the correlation by a so-called *link randomization* scheme. This algorithm induces just a constant overhead by periodically swapping links between random nodes. This way, the degree of each node remains unmodified and the scale-free node degree distribution is preserved while eliminating the correlation.

The analysis of the network construction process has shown, that the algorithm is able to construct an overlay network with a scale-free node degree distribution. It has also been shown that the network is a small-world network having a small characteristic path length and a high clustering coefficient. This shows the usability of the overlay network in our scenario, since it allows fast routing and node lookup. In addition, correlations, being typical for growing networks, have been asserted. In the next section the behavior of the overlay under churn will be analyzed.

2) *Churn*: As shown above, the proposed algorithm generates scale-free small-world networks. In order to show, that these properties remain if nodes join and leave the network we analyzed the behavior of the overlay under churn. For this purpose, we constructed overlay networks with 5000 nodes, simulated 500 steps and measured the same data as above. In each step a random number of nodes joined and left the network. These simulations have been conducted with $M = 10$ (Because it has been shown above that this is a promising

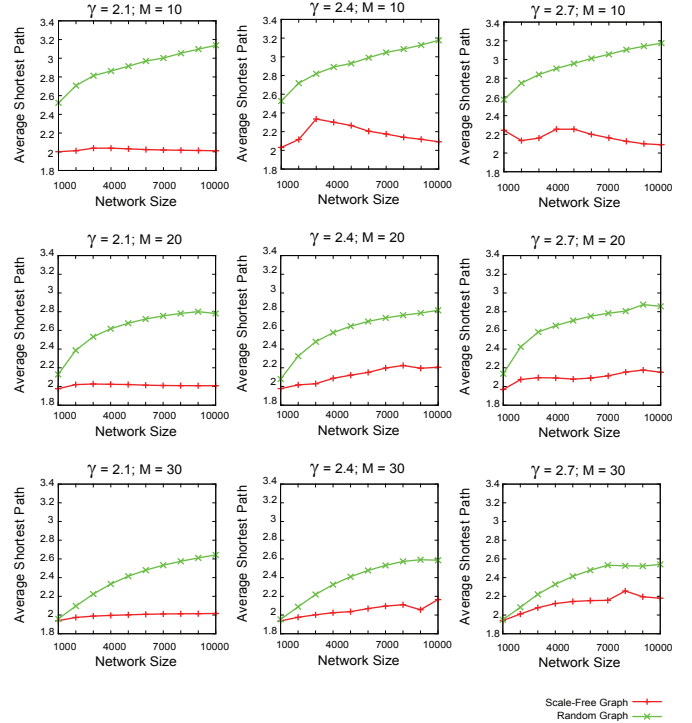


Fig. 10. Average Shortest Path Length for the Scale-Free Graphs and for Random Graphs With the Same Number of Nodes and Links

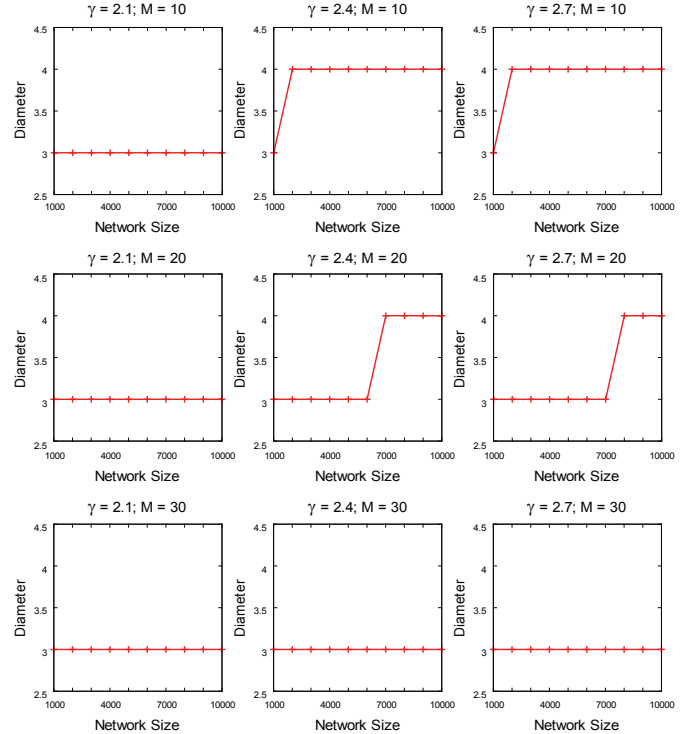


Fig. 11. Network Diameter

M -value) and the same power-law exponents of the payload distribution as above.

The development of the fitted power-law exponent, depicted




Fig. 12. Average Neighbor Degree on a Double Logarithmic Scale in a Network With 10000 Nodes and $M = 10$, Showing the Disassortative Mixing

in Figure 13 shows that the scale-freeness of the overlay remains under churn. This can also be noticed for the small-world property as the development of the clustering coefficient (Figure 14) and the average shortest path length (Figure 15) shows. Moreover the diameter of the network stays small as depicted in Figure 16. Figure 17 depicts the average neighbor degree for a network with 5000 nodes and $M = 10$ after 500 churn steps and shows that the degree-degree correlation remains under churn.

If nodes leave a network it is possible that the network falls apart in isolated partitions and the network becomes unconnected. This cannot happen in our overlay since each node is connected to all its Voronoi neighbors and hence the connectedness of the network is always guaranteed.

The churn simulations have shown that the advantageous properties of the overlay remain under churn which is important for the highly dynamic MMVE setting.

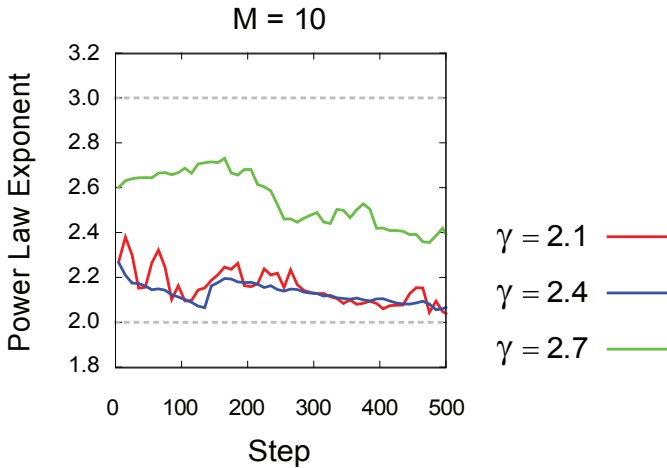


Fig. 13. Fitted Power-Law Exponent Development Under Churn




Fig. 14. Clustering Coefficient Development Under Churn for the Scale-Free Graphs and for Random Graphs With the Same Number of Nodes and Links




Fig. 15. Average Shortest Path Length Development Under Churn for the Scale-Free Graphs and for Random Graphs With the Same Number of Nodes and Links




Fig. 16. Diameter Development Under Churn




Fig. 17. Average Neighbor Degree on a Double Logarithmic Scale in a Network With 5000 Nodes and $M = 10$ After 500 Churn Simulation Steps, Showing the Disassortative Mixing

IV. RELATED WORK

This section presents work and approaches related to the topic of this paper. At first several other P2P-based approaches

for virtual online worlds will be discussed before presenting related approaches for the construction of scale-free networks.

The impracticality of centralized approaches for the provision of a global-scale 3D Web scenario is well recognized. Hence several P2P-based approaches for Massive Multiuser Virtual Environments exist. Since most of them rely on a pure P2P scheme without a backbone infrastructure, the avatar tracking in these systems is implicitly handled in a pure P2P fashion. This contrasts to our HyperVerse approach, where we propose the utilization of a federated backbone service in order to guarantee the reliable and persistent provision of the online environment.

A well-known approach in this field is the *Solipsis* project [20], providing a virtual online world in a pure P2P fashion. *Solipsis* performs node discovery by mutual neighbor notification. Here the current neighbors of each node act as *watchmen* which interconnect the node to new nodes moving into its *Area of Interest*. This scheme does only work if a node resides inside of the convex hull formed by its neighbors, hence neighbor discovery can not be guaranteed, especially in a highly dynamic scenario with high churn rates. *FLoD* [16] provides a framework for pure P2P-based 3D scene streaming that is built upon *VON* [15], a Voronoi-based P2P overlay network distributing the virtual world among peers and relying on mutual notification for neighbor discovery. Since leaving nodes may result in extensive reorganizations of the Voronoi overlay, this approach is also vulnerable to high churn rates of the clients. An approach similar to *VON* is *QuON* [3], but instead of using a Voronoi decomposition of the virtual world surface, *QuON* uses local point quadtrees for this purpose. In [30], an architecture similar to *FLoD* is described that uses super peers, so-called *connectivity peers*, to interconnect the peers. In [33] a highly-structured user client overlay is proposed in order to form a P2P virtual environment. Here a *quadtree* in combination with a *Chord* [32] overlay is applied. A hybrid approach combining a structured Pastry [29] overlay with unstructured P2P message exchange among immediate neighbors has been made by *MOPAR* [36]. This approach divides the world surface into hexagonal cells of fixed size, each managed by a *master node* inside the cell. The *Badumna* [23] middleware incorporates three different protocols, utilized in different situations depending on the network load in order to realize a scalable MMVE network infrastructure that can dynamically be adapted to the current load.

Currently a lot of research work is done in the field of scale-free networks. This is on the one hand due to the advantageous properties of this kind of networks in terms of reliability, resilience and short path length. On the other hand for many large networks it has been shown that they exhibit a scale-free degree distribution. Be it constructed networks like the World Wide Web [1] and the Internet [13], social networks like the graph of scientific collaborations [26] or cellular networks like the metabolic network [19]. The question, which fundamental mechanism leads to the emergence of scale-freeness in large networks has first been answered by the *Barabasi/Albert model* [5]. This model mimics the two basic mechanisms responsible

for the emergence of scale-freeness: *growth* and *preferential attachment*. The basic idea is to constantly add new nodes to the network and connect these nodes to existing nodes in the network with a probability equivalent to the existing nodes' degree. This means a new node is connected to a node with high degree with a higher probability, resembling the principle *the rich get richer*. As a generalization of the Barabasi/Albert model *Bianconi* and *Barabasi* have presented a preferential attachment model based on fitness values assigned to each node in [6]. Another prominent fitness based model that draws edges depending on the benefit of the edge for the involved vertices has been presented by *Caldarelli et al.* [10]. Interestingly this model produces scale-free networks without growth and preferential attachment. A model relying on edge growth instead of node growth has been presented in [22] by *Krapivski, Rodgers* and *Redner*. In [34], *Vázquez et al.* a model has been proposed that creates scale-free network by copying nodes and rewiring its edges with a certain probability. In opposite to the approaches mentioned above, this model works with local rules only. The other models require global knowledge of the network structure and rely on the random selection of network nodes for the creation of edges, which is non trivial with just local knowledge [2]. Our approach is different from those mentioned above since no global instance for the random selection of nodes is required. Rather, for the decentralized random node selection, the existence of the virtual geography in our setting is utilized. In general our approach can be classified as fitness-based model oriented on the Public Server's payloads.

V. CONCLUSION

This paper presented a detailed analysis and evaluation of the self-organized and scale-free backbone overlay network for the HyperVerse infrastructure that has been proposed in [12]. The overlay allows for non-uniform object and avatar distribution as well as for the high avatar dynamics expected in an MMVE scenario. The world surface is distributed among the Public Servers by a Voronoi diagram, each server managing one Voronoi cell. Additionally a scale-free link structure between the network nodes is constructed in a decentralized fashion.

It has been shown that the self-organized load balancing scheme is able to handle the dynamic of the user clients and to react to the formation of hotspots by distributing the world surface accordingly. The scheme guarantees, that spots with a high object mass density are managed by the Public Servers with high payload. If a node becomes overloaded though, it is quickly unloaded by its neighbors and the overload situation is terminated.

The simulation-based analysis of the link structure generation algorithm has shown the scale-free node degree distribution of the resulting network. Moreover a low diameter as well as a short average path length have been observed. The high clustering coefficient measured shows the small-world property of the network. Hence the network is resilient, fault tolerant and allows efficient routing, which are important

criteria for the usability in a global-scale scenario. As churn simulations have shown, these properties remain as nodes join and leave the network. Additionally, the network correlation has been analyzed with the for a growing networks expected result that the overlay exhibits disassortative mixing which however can be eliminated with constant overhead using a link randomization scheme.

REFERENCES

- [1] R. Albert, H. Jeong, and A.-L. Barabási. Internet: Diameter of the world-wide web. *Nature*, 401(6749):130–131, September 1999.
- [2] Y. Azar, A. Z. Broder, A. R. Karlin, N. Linial, and S. Phillips. Biased random walks. In *STOC '92: Proceedings of the twenty-fourth annual ACM Symposium on Theory of Computing*, pages 1–9, New York, NY, USA, 1992. ACM.
- [3] H. Backhaus and S. Krause. QuON: a quad-tree-based overlay protocol for distributed virtual worlds. *International Journal of Advanced Media and Communication*, 4(2):126 – 139, Mar. 2010.
- [4] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509, 1999.
- [5] A.-L. Barabási, R. Albert, and H. Jeong. Mean-field theory for scale-free random networks, July 1999.
- [6] G. Bianconi and A.-L. Barabási. Competition and multiscaling in evolving networks. *EPL (Europhysics Letters)*, 54(4):436, 2001.
- [7] J. Botev, M. Esch, A. Höhfeld, H. Schloss, and I. Scholtes. The HyperVerse - concepts for a federated and torrent-based "3D Web". The 1st International Workshop on Massively Multiuser Virtual Environments (MMVE), 2008.
- [8] J. Botev, A. Höhfeld, H. Schloss, I. Scholtes, P. Sturm, and M. Esch. The HyperVerse - concepts for a federated and torrent-based "3D Web". *International Journal of Advanced Media and Communication*, 2(4):331–350, 2008.
- [9] A. Z. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. L. Wiener. Graph structure in the web. *Computer Networks*, 33(1-6):309–320, 2000.
- [10] G. Caldarelli, A. Capocci, P. De Los Rios, and M. A. Muñoz. Scale-free networks from varying vertex intrinsic fitness. *Phys. Rev. Lett.*, 89(25):258702, Dec 2002.
- [11] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data, 2009.
- [12] M. Esch and I. Scholtes. A scale-free and self-organized P2P overlay for massive multiuser virtual environments. In *Proceedings of the 5th International Conference on Collaborative Computing Network, Applications and Worksharing (CollaborateCom)*, Washington D.C., USA, November 2009.
- [13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 251–262, New York, NY, USA, 1999. ACM.
- [14] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery, 2000.
- [15] S.-Y. Hu, J.-F. Chen, and T.-H. Chen. VON: a scalable peer-to-peer network for virtual environments. *IEEE Network Magazine*, 20(4):22–31, 2006.
- [16] 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 *The 27th Conference on Computer Communications (IEEE INFOCOM '08)*, 2008.
- [17] M. Jelasity and A. Montresor. Epidemic-style proactive aggregation in large overlay networks. In *ICDCS '04: Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04)*, pages 102–109, Washington, DC, USA, 2004. IEEE Computer Society.
- [18] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai. Lethality and centrality in protein networks. *Nature*, 411(6833):41–42, May 2001.
- [19] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–654, October 2000.
- [20] J. Keller and G. Simon. Solipsis: A massively multi-participant virtual world. In *PDPTA*, pages 262–268, 2003.
- [21] P. L. Krapivsky and S. Redner. Organization of growing random networks. *Phys. Rev. E*, 63(6):066123, May 2001.
- [22] P. L. Krapivsky, G. J. Rodgers, and S. Redner. Degree distributions of growing networks. *Phys. Rev. Lett.*, 86(23):5401–4, 2001.
- [23] S. Kulkarni. Badumna network suite: A decentralized network engine for massively multiplayer online applications. In H. Schulzrinne, K. Aberer, and A. Datta, editors, *Peer-to-Peer Computing*, pages 178–183. IEEE, 2009.
- [24] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tompkins, and E. Upfal. The web as a graph. In *PODS '00: Proceedings of the nineteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 1–10, New York, NY, USA, 2000. ACM.
- [25] S. Maslov and K. Sneppen. Specificity and stability in topology of protein networks. *Science*, 296(5569):910–913, May 2002.
- [26] M. E. J. Newman. The structure of scientific collaboration networks. *PROC.NATL.ACAD.SCI.USA*, 98:404, 2001.
- [27] M. E. J. Newman. Assortative mixing in networks. *Phys. Rev. Lett.*, 89(20):208701, Oct 2002.
- [28] R. Pastor-Satorras, A. Vázquez, and A. Vespignani. Dynamical and correlation properties of the internet. *Phys. Rev. Lett.*, 87(25):258701, Nov 2001.
- [29] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware*, volume 2218 of *Lecture Notes in Computer Science*, pages 329–350. Springer, 2001.
- [30] J. Royan, P. Gioia, R. Cavagna, and C. Bouville. Network-based visualization of 3d landscapes and city models. *IEEE Comput. Graph. Appl.*, 27(6):70–79, 2007.
- [31] I. Scholtes, J. Botev, M. Esch, A. Höhfeld, H. Schloss, and B. Zech. TopGen - internet router-level topology generation based on technology constraints. In *Simutools '08: Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*, pages 1–10, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
- [32] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 149–160, New York, NY, USA, 2001. ACM Press.
- [33] E. Tanin, A. Harwood, and H. Samet. Using a distributed quadtree index in peer-to-peer networks. *The VLDB Journal*, 16(2):165–178, 2007.
- [34] A. Vázquez, A. Flammini, A. Maritan, and A. Vespignani. Modeling of protein interaction networks. *ArXiv Condensed Matter e-prints*, Aug. 2001.
- [35] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.
- [36] A. P. Yu and S. T. Vuong. Mopar: a mobile peer-to-peer overlay architecture for interest management of massively multiplayer online games. In *NOSSDAV '05: Proceedings of the international workshop on Network and operating systems support for digital audio and video*, pages 99–104, New York, NY, USA, 2005. ACM.